

English



FUJITSU Software BS2000

openSM2 V10.0

Software Monitor

User Guide

Edition April 2015

Comments... Suggestions... Corrections...

The User Documentation Department would like to know your opinion on this manual. Your feedback helps us to optimize our documentation to suit your individual needs.

Feel free to send us your comments by e-mail to:

manuals@ts.fujitsu.com

Certified documentation according to DIN EN ISO 9001:2008

To ensure a consistently high quality standard and user-friendliness, this documentation was created to meet the regulations of a quality management system which complies with the requirements of the standard DIN EN ISO 9001:2008.

cognitas. Gesellschaft für Technik-Dokumentation mbH

www.cognitas.de

Copyright and Trademarks

Copyright © 2015 Fujitsu Technology Solutions GmbH.

All rights reserved.

Delivery subject to availability; right of technical modifications reserved.

All hardware and software names used are trademarks of their respective manufacturers.

Contents

1	Preface	15
1.1	Objectives and target groups of this manual	17
1.2	Summary of contents	17
1.3	Changes since the last edition of this manual	18
1.4	Notational conventions	19
2	The SM2 monitoring program	21
2.1	Overview	21
2.2	Operating modes	23
2.3	Users	24
2.4	Table of authorizations	25
2.5	Activating a monitoring task	26
2.6	Acquisition of monitored data	27
2.7	Monitoring cycle	29
2.8	Outputting and saving monitored data	30
2.9	Monitored variables	32
2.10	Use in a computer network	36

3	System monitoring with openSM2	41
3.1	Performance expectations of the user	41
3.2	Applications	43
3.3	Ensuring efficient DP system usage	45
3.4	Analyzing monitored data	46
4	SM2 monitoring programs	49
4.1	Overview	49
4.2	Privileged SM2 monitoring programs	52
	BCAM-CONNECTION	
	Monitored data on connection sets	52
	CHANNEL-IO	
	Monitored data on the channel load	57
	CMS	
	Monitored data on the catalog management system	57
	COSMOS	
	Monitored data on the system for bottleneck analysis	58
	DAB	
	Monitored data on DAB activities	58
	DISK-FILE	
	Monitored data on file access for selected disks	59
	DLM	
	Monitored data on lock requests	59
	FILE	
	Monitored data on file access for selected files	60
	GS	
	Monitored data on access to global storage	60
	GSVOL	
	Monitored data on volumes in global storage	60
	HSMS	
	Monitored data on file migration	60
	ISAM	
	Monitored data on ISAM pools	61
	MSCF	
	Monitored data on communication between hosts	61
	NSM	
	Monitored data on HIPLEX MSCF®	61

OPENFT	
Monitored data on openFT	61
PERIODIC-TASK	
Monitored data on tasks	62
PFA	
Monitored data on caches under User PFA	62
POSIX	
Monitored data on POSIX	62
PUBSET	
Monitoring data for SF pubsets and volume sets	63
RESPONSETIME	
Monitored data on the BCAM pool	63
SAMPLING-DEVICE	
Monitored data on I/Os, data volume and utilization of devices	68
SERVICETIME	
Monitored data on the service time	69
SESAM-SQL	
Monitored data for the SESAM/SQL database system	71
STORAGE-SYSTEM	
Monitored data for Symmetrix systems	71
SVC	
Monitored data on SVC calls	72
SYSTEM	
Global system and category-specific monitored data	72
TASK	
Task-specific monitored data	73
TCP-IP	
Monitored data on TCP/IP connections	74
TLM	
Monitored data on locks	74
UDS-SQL	
Monitored data for the UDS/SQL database system	74
UTM	
Monitored data on openUTM applications	75
VM	
CPU shares of guest systems under VM2000	76
4.3 Nonprivileged SM2 monitoring programs	77
FILE	
Monitored data on file access	77
ISAM	
Monitored data on ISAM pools	77
TASK	
Task-specific monitored data	78

5	SM2 operation	79
5.1	Starting and terminating SM2	79
5.2	Overview of SM2 operation	81
5.3	Statements for SM2 administrators	83
	ADD-BCAM-CONNECTION-SET	
	Add connection set for BCAM-CONNECTION monitoring	88
	ADD-CONNECTION-SET	
	Add connection set for RESPONSETIME monitoring	91
	ADD-COSMOS-EVENT	
	Define events for monitored-data acquisition	93
	ADD-FILE	
	Define file to be monitored	95
	ADD-ISAM-FILE	
	Define an ISAM pool to be monitored	96
	ADD-ISAM-POOL	
	Define ISAM pool to be monitored	97
	ADD-OPENFT-INSTANCE	
	Add openFT instance for monitoring	98
	CALL-ADMINISTRATION-PART	
	Switch from evaluation part to administration facility	99
	CALL-EVALUATION-PART	
	Switch from administration facility to evaluation part	99
	CHANGE-MEASUREMENT-PROGRAM	
	Stop current monitoring program run and restart it with new monitored objects	100
	CLOSE-LOG-FILE	
	Close SM2 output file	101
	END	
	Terminate SM2 run	101
	INITIATE-COSMOS	
	Prepare COSMOS monitoring program	102
	MODIFY-ADMINISTRATOR-ADMISSION	
	Admit additional SM2 administrators	102
	MODIFY-COSMOS-PARAMETERS	
	Modify parameters for COSMOS monitoring	103
	MODIFY-MEASUREMENT-PERIODS	
	Modify SM2 monitoring cycle	112
	MODIFY-RESPONSETIME-PARAMETERS	
	Modify parameters for RESPONSETIME monitoring	114
	MODIFY-USER-ADMISSION	
	Define authorizations for nonprivileged users	117

OPEN-LOG-FILE	
Open SM2 output file	120
REMOVE-BCAM-CONNECTION-SET	
Exclude connection set(s) from BCAM-CONNECTION monitoring	121
REMOVE-CONNECTION-SET	
Exclude connection set(s) from RESPONSETIME monitoring	122
REMOVE-COSMOS-EVENT	
Remove events from COSMOS monitoring program	123
REMOVE-FILE	
Remove file(s) from FILE monitoring program	124
REMOVE-ISAM-FILE	
Remove ISAM pool from ISAM monitoring program	125
REMOVE-ISAM-POOL	
Remove ISAM pool(s) from ISAM monitoring program	126
REMOVE-OPENFT-INSTANCE	
Remove openFT instance for OPENFT monitoring program	127
SELECT-HOSTS	
Define hosts for SM2 statements and screen output	128
SET-BCAM-CONNECTION-PARAMETERS	
Define BCAM-CONNECTION monitoring parameters	130
SET-CHANNEL-IO-PARAMETERS	
Define CHANNEL-IO monitoring parameters	132
SET-COSMOS-PARAMETERS	
Define COSMOS monitoring parameters	133
SET-DISK-FILE-PARAMETERS	
Define DISK-FILE monitoring parameters	141
SET-PERIODIC-TASK-PARAMETERS	
Define PERIODIC-TASK monitoring parameters	141
SET-RESPONSETIME-PARAMETERS	
Define RESPONSETIME monitoring parameters	143
SET-SAMPLING-DEVICE-PARAMETERS	
Monitoring program definition for SAMPLING-DEVICE monitoring	145
SET-SERVICETIME-PARAMETERS	
Define SERVICETIME monitoring parameters	146
SET-STORAGE-SYSTEM-PARAMETERS	
Define monitoring parameters for the monitoring program STORAGE-SYSTEM	147
SET-SYSTEM-PARAMETERS	
Define SYSTEM monitoring parameters	149
SET-TASK-PARAMETERS	
Define TASK monitoring parameters	151
SHOW-ACTIVE-PARAMETERS	
Output active monitoring parameters	154
SHOW-DEFINED-PARAMETERS	
Output defined monitoring parameters	155

SHOW-MEASUREMENT-STATUS	
Output monitoring status	156
SHOW-SELECTED-HOSTS	
Output selected hosts	156
SHOW-SM2-STATUS	
Output status of SM2 system tasks	157
SHOW-USER-MEASURED-OBJECTS	
Output monitored objects and associated users	159
START-MEASUREMENT-PROGRAM	
Start monitoring program run	160
STOP-MEASUREMENT-PROGRAM	
Terminate monitoring program run	161
5.4 Statements for nonprivileged users	162
BREAK	
Switch to system mode	170
CHANGE-ISAM-STATISTICS	
Include and exclude ISAM pools	171
END	
Terminate monitoring	173
FILE	
Monitoring files	174
HELP	
Request user help information	175
OUTPUT	
Define output mode	176
REMARK	
Insert remarks	178
REPORT	
Select reports	179
RESTART	
Start output of selected reports	183
SELECT-CHANNEL-PARAMETERS	
Specify sort criterion for CHANNEL report	183
SELECT-CMS-PUBSET	
Select pubsets/private disks	184
SELECT-DAB-CACHE	
Select DAB caches	185
SELECT-DEVICE-DISK-PARAMETERS	
Specify sort criterion for DEVICE DISK report	186
SELECT-PERIODIC-TASK-PARAMETERS	
Select sort criterion and output information for PERIODIC TASK report	187

SELECT-SYMMETRIX-PUBSET	
Select pubsets for SYMMETRIX DEVICE report	189
SELECT-SYMMETRIX-UNIT	
Select Symmetrix systems for SYMMETRIX PUBSET report	190
SELECT-UTM-APPLICATION	
Select UTM applications	191
SHOW-USER-MEASURED-OBJECTS	
Output monitored objects	192
START	
Start report output	192
START-ISAM-STATISTICS	
Monitor ISAM pools	193
STATUS	
Output monitoring status	195
STOP-ISAM-STATISTICS	
Terminate monitoring of ISAM pools	195
5.5	BS2000 commands for activating user task monitoring
	196
5.6	Writing to the SM2 output file
	202
5.7	SM2 monitoring program: examples
	206
6	SM2 screen output
	209
<hr/>	
6.1	SM2 reports
	218
ACF report	218
ACTIVITY report	220
BCAM CONNECTION report	224
BCAM MEMORY report	227
CATEGORY report	228
CATEGORY QUEUE report	230
CATEGORY WSET report	233
CHANNEL report	235
CMS report	238
CPU report	241
DAB reports	243
DAB report	244
DAB CACHE report	246
DEVICE DISK report	249
DEVICE OTHER report	252
DEVICE TAPE report	254
DEVICE TD report	256
DISK FILE report	257

DLM report	258
FILE reports	260
GLOBAL report	262
GS report	263
GSVOL report	264
ISAM reports	265
ISAM FILE report	268
MEMORY report	270
MSCF report	273
NSM report	275
OPENFT report	277
PCS report	280
PERIODIC TASK report	282
PFA CACHE report	284
POSIX report	286
PUBSET report	288
RESPONSETIME report	290
SESAM SQL report	294
SHARED PUBSET report	295
SVC report	297
SYMMETRIX reports	298
SYMMETRIX report	300
SYMMETRIX PUBSET report	302
SYMMETRIX DEVICE report	304
TCP/IP report	306
TLM report	308
UDS SQL report	309
UTM reports	310
UTM report	311
UTM APPLICATION report	313
VM report	318
VM CPU POOL report	321
VM GROUP report	323
6.2 SM2 information screens	324
ACTIVE PARAMETER	324
DEFINED PARAMETER	334
MEASUREMENT STATUS	335
SELECTED HOSTS	339
STATUS TABLE	340
USER MEASURED OBJECTS	343

7	Installation and application of SM2	345
7.1	Installation	345
7.1.1	Loading and unloading the subsystem	347
7.1.2	Installation of the manager of INSPECTOR and ANALYZER	347
7.2	System resource utilization by SM2	348
7.2.1	External storage utilization	348
7.2.2	CPU utilization	349
7.2.3	Main memory utilization	351
7.3	Accuracy of the SM2 data	352
7.3.1	Causes of inaccuracies	352
7.3.2	Notes on particular variables	354
7.3.3	Measures for reducing the error sources	355
7.4	Special applications	356
7.4.1	SM2 together with VM2000	356
7.4.2	VOLUME load with DRV	363
7.5	Sample procedures	364
8	SM2U1 utility routine	373
8.1	Preparing the SM2U1 session	375
8.2	Starting and terminating SM2U1	377
8.3	SM2U1 statements	378
	END	
	Terminate SM2U1	379
	MERGE-FILES	
	Merge files	379
	SELECT-MEASUREMENT-GROUPS	
	Select records	381
	SEPARATE-FILES	
	Separate SM2 output files	383
	SET-COMPRESSION	
	Turn SM2U1 compression function on or off	384
	SHOW-INFORMATION	
	Output information on monitoring programs and monitoring cycles	386
8.4	Notes on the output files	388
8.5	Plausibility checks	388
8.6	Notes on implementation	389
8.7	Examples for the SM2U1 utility routine	389

9	SM2R1 analysis routine	391
9.1	Starting and terminating SM2R1	397
9.2	Monitoring environment output	401
9.3	Structure of time series and statistical data	404
9.3.1	Computation methods	404
9.3.2	Bar chart layout	406
9.3.3	Representation of the statistical data	410
9.3.4	Variables reports	411
9.3.5	Additional outputs	411
9.4	General structure of the statements	412
9.5	Statements	414
	CREATE-JOBCLASS-SET	
	Define set name for job classes	414
	CREATE-JOBNAME-SET	
	Specify set name for job names	415
	CREATE-TSN-SET	
	Define set name for specific TSNs	416
	CREATE-USERID-SET	
	Define set name for user IDs	417
	END	
	Terminate statements	417
	MODIFY-REPORT-CONDITIONS	
	Modify preset range for monitored data of subreport	418
	PRINT-CONFIGURATION	
	Output system configuration	420
	PRINT-HSMS-STATISTICS	
	Print data on file migration and retrieval	422
	PRINT-QUEUE-TRANSITION	
	Print statistics on task queues	424
	PRINT-REPORTS	
	Print report groups	426
	PRINT-SUMMARY	
	Print summary analysis	449
	PRINT-TASK-STATISTICS	
	Print task statistics	461
	SET-EVALUATION-PERIOD	
	Set analysis period and analysis subinterval	473
	SET-EXCEPTION-PERIOD	
	Exclude periods	478

	SET-REPORT-FOCUS	
	Select time slots from analysis period	482
	SET-TITLE	
	Output title	484
	START-AUTOMATIC-ANALYSIS	
	Start automatic bottleneck analysis	485
9.6	SM2R1 analysis routine examples	488
9.7	SM2R1 transfer file records	499
10	Other analysis routines	507
<hr/>		
10.1	openSM2 Manager	507
10.1.1	Calling the openSM2 Manager	508
10.1.2	Working with the openSM2 Manager	509
10.1.3	Functions of the openSM2 Manager	510
10.2	ANALYZER	
	Analysis with a Windows PC	512
10.3	INSPECTOR	
	Online monitoring with a Windows PC	519
10.4	SM2-PA	
	Program analyzer	522
11	SM2 program interfaces	525
<hr/>		
11.1	C interfaces	525
11.1.1	The SM2GMS macro	526
11.1.2	The SM2GDAT macro	528
11.1.3	Evaluating the return codes	540
11.1.4	Example	544
11.1.5	Structures of the SM2GMS macro	557
11.1.6	Structures of the SM2GDAT macro	558
11.2	Program interface for querying system utilization	604

Contents

12	Variables reports	607
12.1	Table of variables	607
12.2	Table of report groups	611
12.3	Table of report names	614
12.4	Table of variables reports	623
	Glossary	661
	Related publications	675
	Index	677

1 Preface

The software product openSM2 (BS2000) permits comprehensive online monitoring with configurable alarm management, detailed bottleneck analyses and forecasts of the future performance requirements. openSM2 (BS2000) thus enables users to optimize the performance of their BS2000 systems.

openSM2 (BS2000), together with openSM2 (Open Systems), forms the product line openSM2, a powerful solution for monitoring heterogeneous IT environments.

openSM2 (Open Systems) supports the server platforms Solaris, Linux, Microsoft Windows, VMware vSphere, Xen and X200, the ETERNUS DX and Symmetrix disk storage systems and all SNMP-capable systems.

Details of the distributions and versions supported are provided in the Data Sheet on openSM2 (Open Systems). openSM2 (Open Systems) is not described in this manual.

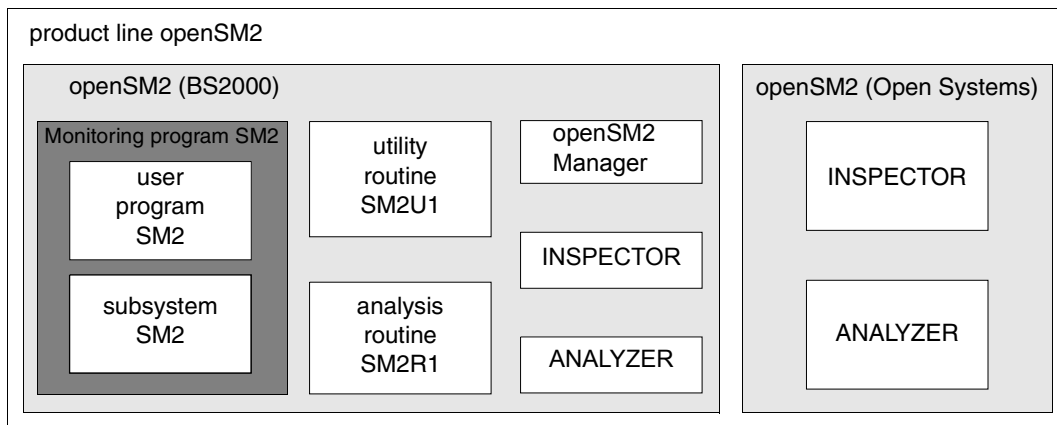


Figure 1: Product line openSM2

Overview of openSM2 (BS2000)

openSM2 (BS2000) provides the user with statistics on the performance of the BS2000 operating system and on resource utilization.

openSM2 (BS2000) consists of the following components:

- SM2 monitoring program
- SM2U1 utility routine
- SM2R1 analysis routine
- openSM2 Manager
- INSPECTOR
- ANALYZER

The SM2 monitoring program consists of a privileged subsystem and a nonprivileged user program.

- In the SM2 subsystem the monitored data is recorded, made available at a program interface for real-time monitoring (online), and on request written to an output file which can subsequently be analyzed (offline).
- The SM2 user program offers functions for controlling monitored data acquisition and for presenting the monitored data on the screen.

The SM2U1 utility routine is used for managing SM2 output files.

The SM2R1 analysis routine analyzes SM2 output files.

The openSM2 Manager is available as add-on software in the SE manager for the FUJITSU Servers BS2000 SE Series (for short: SE servers). The openSM2 Manager is the web-based user interface for central monitoring of all components of an SE server.

INSPECTOR is used to simultaneously monitoring multiple BS2000 systems online on a Windows PC.

ANALYZER is used to analyzing SM2 output files on a Windows PC.

In addition, the chargeable software product SM2-PA is available for analyzing user-specific output files.

1.1 Objectives and target groups of this manual

This manual describes the software product openSM2 (BS2000).

It is intended for systems support staff who needs information on the current system workload. It will be of particular interest to systems support staff who wish to assess the performance of their configuration and isolate possible bottleneck areas by means of long-term monitoring.

1.2 Summary of contents

This manual describes the SM2 monitoring program, the SM2U1 utility routine and the SM2R1 analysis routine, and gives an overview of the other analysis routines. These routines allow you to carry out performance analyses in BS2000.

At the end of the manual you will find various lists which are designed to make it easier for you to use the manual.

Readme file

The functional changes to the current product version and revisions to this manual are described in the product-specific Readme file.

Readme files online

Readme files are available to you online in addition to the product manuals under the various products at <http://manuals.ts.fujitsu.com>.

Readme files under BS2000

On your BS2000 system you will find Readme files for the installed products under the file name:

```
SYSRME.<product>.<version>.E
```

Please refer to your system administrator for the user ID under which the required Readme file can be found. You can also obtain the path name of the Readme file directly by entering the following command:

```
/SHOW-INSTALLATION-PATH INSTALLATION-UNIT=<product>,LOGICAL-ID=SYSRME.E
```

You can view the Readme file on screen with `/SHOW-FILE` or by opening it in an editor, or print it on a standard printer using the following command:

```
/PRINT-DOCUMENT <filename>, LINE-SPACING=*BY-EBCDIC-CONTROL
```

Additional product information

Current information, version and hardware dependencies, and instructions for installing and using a product version are contained in the associated Release Notice. These Release Notices are available at <http://manuals.ts.fujitsu.com>.

1.3 Changes since the last edition of this manual

The following major changes have been made since the last edition of this manual:

- The manual has been brought into line with BS2000 OSD/BC V10.0.
- The openSM2 Manager is available as add-on software in the SE Manager. New [section “openSM2 Manager”](#).
- New overflow screen in the UTM APPLICATION report with further data for the UTM application (including TAC class values).
- The SM2GDAT macro has been enhanced accordingly for this purpose.
- When Live Migration takes place, a new session with the new server is started in the SM2 output file. Automatic performance analysis across the session boundary leads to separate analyses for the server concerned.
- The operand value SCOPE=*USER(...) has been omitted from the SM2 administrator statement ADD-ISAM-POOL, as has the operand value SCOPE=*USER in the CHANGE-ISAM-STATISTICS statement (nonprivileged users).

1.4 Notational conventions

All statements for controlling monitoring operations (SM2 administration statements) and the statement interfaces of the SM2U1 and SM2R1 utilities are fully supported by the SDF statement syntax. The SDF syntax is described in the “Commands” manual [3].

The SM2 functions for selecting and controlling screen output can only be addressed via ISP. The ISP syntax is described on [page 169](#).

On account of the frequency with which names are used in this manual, the following abbreviations have been adopted for the sake of simplicity and clarity:

- **BS2000 servers** for the servers with /390 architecture and the servers with x86 architecture. These servers are operated with the corresponding BS2000 operating system.
- **Servers with /390 architecture** for the Server Unit /390 of the FUJITSU Server BS2000 SE Series and the Business Servers of the S Series
- **Servers with x86 architecture** for the Server Unit x86 of the FUJITSU Server BS2000 SE Series and the Business Servers of the SQ Series (x86-64 architecture)
- **SE servers** for the FUJITSU Server SE Series (Server Units /390 and x86)
- **S servers** for the Business Servers of the S Series (/390 architecture)
- **SQ servers** for the Business Servers of the SQ Series (x86-64 architecture)

In examples the strings `<date>`, `<time>` and `<version>` specify the current outputs for date, time and version of a software product when the examples are otherwise independent of the date, time and version.

The string `<ver>` specifies a version number in file names; for openSM2, `<ver>` = 190, for SM2-TOOLS, `<ver>` = 090.

The following typographical elements are used in this manual:



For notes on particularly important information

[]

References to other publications within the text are given in abbreviated form followed by numbers; the full titles are listed in the “References” section at the back of this manual.

input

Inputs and system outputs in examples are shown in typewriter font.

2 The SM2 monitoring program

2.1 Overview

SM2 provides the user with information on the status and performance of the system, making it possible to improve the performance of both the system and its applications. The mode and scope of monitored data acquisition in SM2 can be controlled using commands and statements.

SM2 consists of a privileged subsystem and a nonprivileged user program. The SM2 subsystem generates a number of system tasks which are responsible for collecting data and writing it to the SM2 output files.

The following figure shows the exchange of information between the various SM2 tasks.

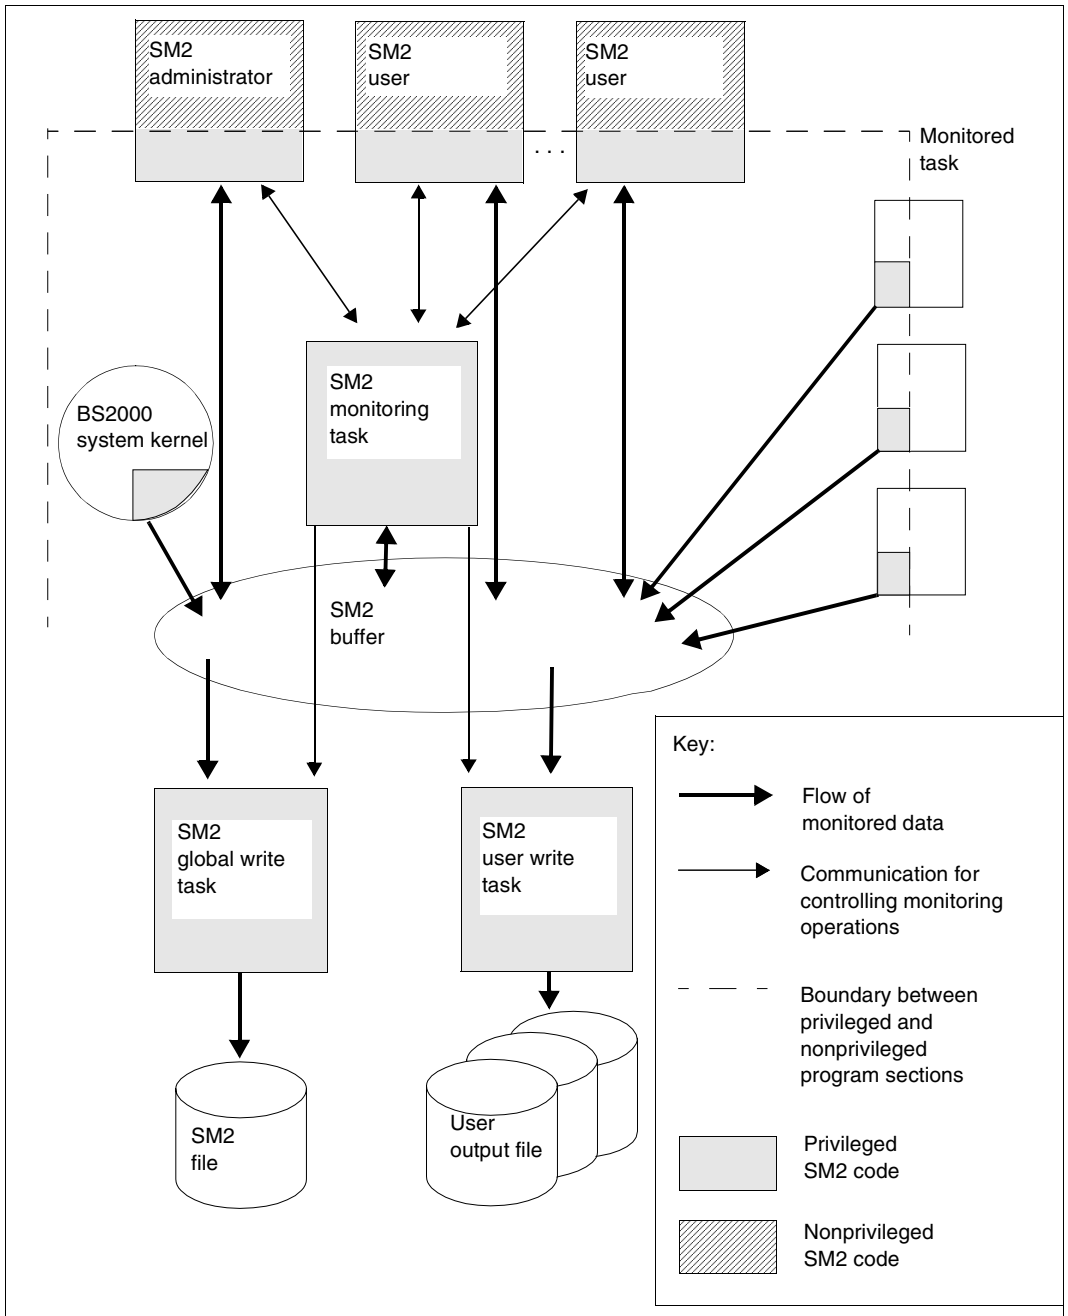


Figure 2: SM2 configuration

2.2 Operating modes

SM2 offers the following operating modes. They can be used simultaneously.

- Real-time monitoring

Here SM2 periodically provides monitored data for online monitoring. The SM2 user program enables the monitored data to be displayed on the terminal in the form of reports. This data refers to the latest monitoring cycle (e.g. 150 seconds) and can be used as snapshots for assessing the current system status. Only selected SM2 monitored data is output in this mode. INSPECTOR enables multiple BS2000 systems to be monitored simultaneously on one Windows PC.

- Background monitoring (information output to a file for subsequent analysis)

Here the collected data is output in the form records to a file, the SM2 output file. This data can be analyzed later using SM2R1 or ANALYZER.

This mode allows for selective use of the SM2 operands and is equally suitable for long-term monitoring with a low system load and for brief special-purpose monitoring with a correspondingly higher system load.

- User-specific task monitoring

Here the nonprivileged user (using the START-TASK-MEASUREMENT command) can register tasks under his/her own ID for monitoring by SM2. In addition to the task-specific monitored data, a command counter and SVC task statistics can be requested. In the case of user-specific task monitoring, all information is written to user-specific SM2 output files.

2.3 Users

Privileged users

Privileged SM2 users are users who have been assigned the system privilege SW-MONITOR-ADMINISTRATION. These fall into the following categories: “primary” administrators, “secondary” administrators and other privileged users.

Entry of an administration statement gives the first privileged SM2 user the status of SM2 primary administrator (except the statements SHOW- and SELECT-HOSTS). This user retains this status until he/she terminates his/her program or switches over to online analysis using the CALL-EVALUATION-PART statement. The SM2 primary administrator has all privileges, and there can only be one primary administrator at any one time. The primary administrator is the only person who is authorized to admit other (secondary) administrators (MODIFY-ADMINISTRATOR-ADMISSION statement) and to permit nonprivileged users to run monitoring programs (MODIFY-USER-ADMISSION).

With the exception of the MODIFY-ADMISSION statements, the secondary administrators have the same rights as the primary administrator. All other privileged users are only offered the SHOW functions and the SELECT-HOSTS statement in the administration section.

All administrators can:

- create and close the SM2 output file
- set monitoring parameters
- activate or deactivate optional monitoring runs.

In addition to these privileges, the SM2 administrators can also use all statements of a nonprivileged user.



SM2 cannot prevent the various administrators from carrying out operations which conflict with one another. It is only possible to make sensible use of the system if the different administrators agree on a common course of action.

Modification of the variables to be monitored and functions which can be added may influence the real-time monitoring of other SM2 users. Modification of the monitoring cycle applies to all SM2 users.

Nonprivileged users

All users who call the SM2 user routine from a user ID without the system privilege SW-MONITOR-ADMINISTRATION are categorized as nonprivileged users. These users are only authorized to use the statements for controlling output during real-time monitoring.

2.4 Table of authorizations

The following table indicates which users are permitted to carry out which functions and under what circumstances.

Function (group)	Privilege SWMONADM			Priv. usr	User
	Prim. adm	Sec. adm	Other		
Start/stop functions					
Start SM2 monitor	y	n	n	n	n
Terminate SM2 monitor	5	5	5	5	5
Administration functions					
MODIFY-ADMINISTRATOR-ADMISSION	y	n	n	n	n
MODIFY-USER-ADMISSION	y	n	n	n	n
MODIFY-MEASUREMENT-PERIODS	y	y	n	n	n
OPEN-/CLOSE-LOG-FILE	y	y	n	n	n
ADD/REMOVE measurement object	y	y	n	n	n
SET/MOD measurement PARAMETER def.	y	7	n	n	n
INITIATE-COSMOS	y	n	n	n	n
START-/CHANGE-/STOP-MEASUREMENT-PROGRAM	y	7	n	n	n
SHOW statements	y	y	y	n	n
CALL-EVALUATION-PART	y	y	y	n	n
SELECT-HOSTS	y	y	y	n	n
Analysis functions					
OUTPUT / START	y	y	y	y	y
SELECT-CHANNEL / -CMS / -DAB / -DEVICE / -PERIODIC-TASK / -PFA / -SYMMETRIX / -UTM	y	y	y	y	y
STATUS	y	y	y	y	y
REPORT	y	y	y	1	1
RESTART	y	y	y	1	1
FILE	6	6	6	2	n
START-/CHANGE-/STOP-ISAM-STATISTICS	6	6	6	2	n
SHOW-USER-MEASURED-OBJECTS	4	4	4	4	4
CALL-ADMINISTRATION-PART	y	y	y	n	n
BS2000 commands					
START-/STOP-TASK-MEASUREMENT	6	6	6	2	n

Table 1: Authorization table

Other SM2 features					
Monitor foreign files	3	3	n	n	n
Monitor foreign ISAM pools	3	3	n	n	n
Monitor foreign tasks	6	6	6	n	n
Monitor SVCs / PCounter	6	6	6	6	6

Table 1: Authorization table

Key

- Prim. adm: Primary administrator
- Sec. adm: Secondary administrator
- Other: Other privileged user
- Priv. user: Users permitted to execute a monitoring program as specified by MODIFY-USER-ADMISSION
- User: Users not permitted to execute a monitoring program
- y: Yes
- n: No
- 1: Yes, but some reports are reserved for privileged users
- 2: Yes if the corresponding monitoring process is permitted for the caller
- 3: Yes, but only for the corresponding privileged monitoring program
- 4: Yes, but only for objects registered by this user
- 5: System privilege SUBSYSTEM-MANAGEMENT required; SM2 privileges are irrelevant
- 6: Yes, provided that "Priv. user" is also active
- 7: Yes, except for the COSMOS monitoring program

2.5 Activating a monitoring task

When SM2 is called for the first time during a session, the SM2 monitoring task is activated. This is an internal task which collects the monitored data, edits it, and stores it in a central buffer. From this buffer, the data is transferred to the various SM2 user tasks or to the SM2 output file for output. If different offline and online periods are selected, the monitored data is written to two central buffer stores.

The precise time at which the SM2 monitoring task was initiated and the online or offline period selected is given in the SM2 MEASUREMENT STATUS under SM2 GATHERING TASK CREATED AT, ONLINE PERIOD, OFFLINE PERIOD.

2.6 Acquisition of monitored data

SM2 records a wide range of monitored data and outputs it to either a screen and/or a file. The data is collected at regular intervals, called monitoring cycles. Some of the monitoring tasks are carried out by default, while others are carried out by special monitoring programs which can be activated as and when required.

The monitored data is then processed and displayed internally by SM2 or by independent openSM2 monitoring and analysis routines.

There are three methods of collecting data:

- Based on the monitoring cycle

Most of the monitored data (e.g. the CPU load) is collected at the end of each monitoring cycle. The current count (or time stamp) is taken from the SM2 or system tables, and the difference between this value and the value at the end of the last monitoring cycle is calculated.

The same procedure is used to calculate the intermediate values for the methods described below at the end of the monitoring cycle (except for the monitoring program TASK).

A monitoring cycle can be set to between 10 seconds and 1 hour.

- Based on the sampling cycle

Because the status is constantly changing, it may not be appropriate to wait until the end of the monitoring cycle to query some monitored data (e.g. device utilization). When acquiring this type of monitored data, it is recommended that the monitoring cycle be broken down into several sampling cycles. At the end of each sampling cycle, the current monitored variable is obtained. Then, at the end of each monitoring cycle, the mean value across all sampling cycles is computed.

A sampling cycle can be set between 200 milliseconds and 10 seconds.

- Event-driven

Some monitored data (e.g. the disk operation times) is determined by monitoring events within the system. When an event occurs, such as the start of an input/output operation, an SM2 routine is activated which collects the event-specific data. This data is then used to calculate the monitored variables.

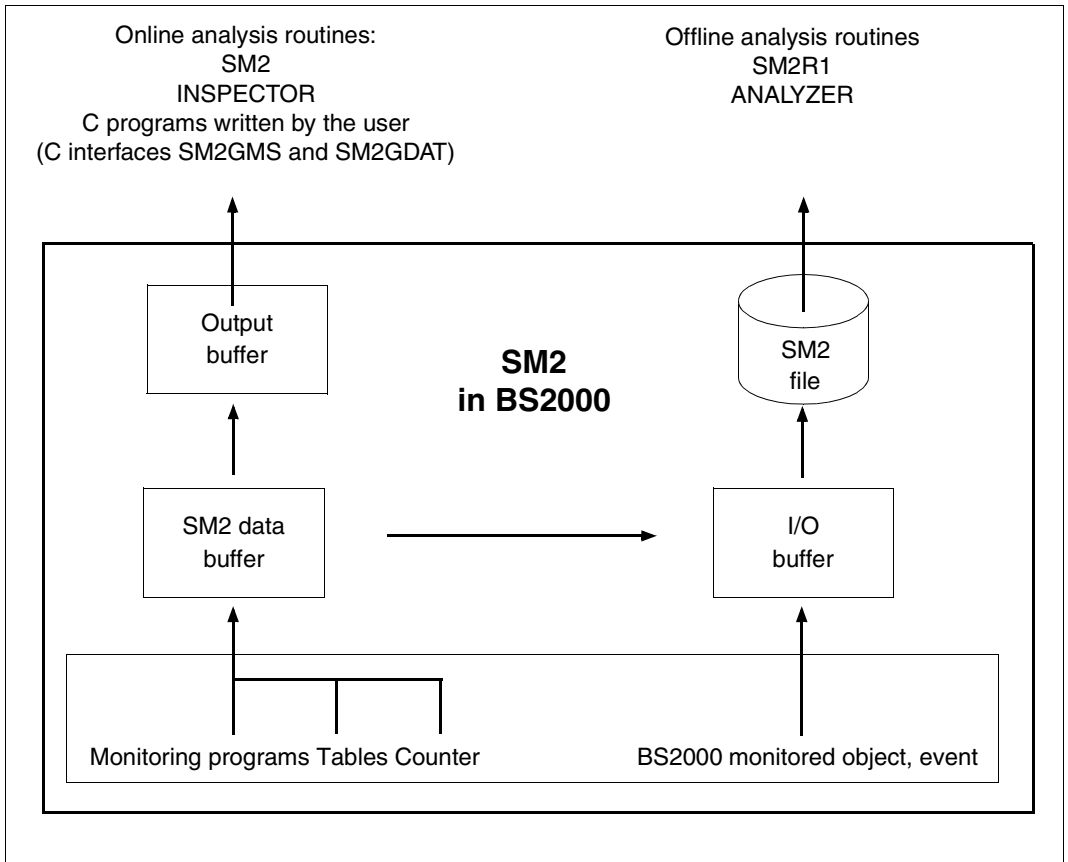


Figure 3: Collecting monitored data

2.7 Monitoring cycle

As mentioned above, SM2 collects data at regular intervals known as monitoring cycles (except for the monitoring program TASK). At the end of each monitoring cycle, the processed data is written to the data buffer and to the SM2 output file, if one is available. The length of the monitoring cycle is preset at 150 seconds when the subsystem is started. The SM2 administrator can change this setting using the OFFLINE-PERIOD operand of the MODIFY-MEASUREMENT-PERIODS statement. The value is entered under OFFLINE PERIOD in SM2 MEASUREMENT STATUS.

Online monitoring cycle

The online cycle also defines a monitoring cycle in which the monitored data for online analysis is collected and written to another buffer. The online cycle is deactivated by default. It can be modified using the ONLINE-PERIOD parameter in the MODIFY-MEASUREMENT-PERIODS statement. When the online cycle is deactivated, the ONLINE-PERIOD corresponds to the OFFLINE-PERIOD. The CYCLE column in each report shows the online value currently set. In addition, the value under the ONLINE-PERIOD column is entered in SM2 MEASUREMENT STATUS.

The online cycle should only be activated if the delay at the terminal is too great when carrying out online monitoring. This may happen if a long monitoring cycle is set for output to an SM2 output file. Please note that by activating the online cycle, you also increase the workload on the system.

Sampling cycle

For monitored-data acquisition based on random sampling, SM2 subdivides the monitoring cycle into a large number of small, regular time slots called sampling cycles. The monitoring task is activated at the specified interval for data collection. At the end of each sampling cycle, snapshots are taken for some variables. From the large number of snapshots, a mean value is computed for each monitoring cycle.

The accuracy of the computed value thus depends on the length of the sampling cycle:

The shorter the sampling cycle, the greater the number of snapshots that can be taken during the monitoring cycle. The SM2 workload on the system, however, increases simultaneously.

The sampling cycle is preset to 800 milliseconds when the subsystem is started.

The SM2 administrator can change this value in the SAMPLING-PERIOD operand of the MODIFY-MEASUREMENT-PERIODS statement. The SAMPLING PERIOD column in the SM2 MEASUREMENT STATUS indicates the current setting, while the SAMPLES column in the individual reports contains the number of samples taken during the latest monitoring cycle.

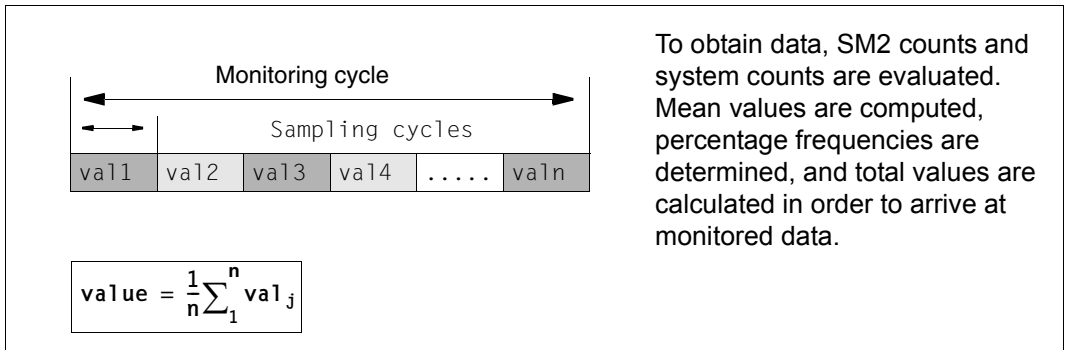


Figure 4: Relationship between the sampling cycle and the monitoring cycle

2.8 Outputting and saving monitored data

At the end of a monitoring cycle, the monitored data is summarized into groups which have similar content, extended by the SM2 monitoring task to include additional information (time stamp, host,...), and then written to a central SM2 data buffer.

If a user has activated online monitoring, the monitored data is transferred from the central data buffer to the buffer of the user task. Note that sensitive data is transferred to privileged users only. The monitored data is then processed and output in the address space of the user task.

For background monitoring, the SM2 administrator can open an SM2 output file. At the end of a monitoring cycle, the monitored data is grouped into data records and written to the SM2 output file. The information in the SM2 output file can be used for subsequent problem and trend analyses. Special analysis routines are provided for this purpose.

Methods used for compressed output of monitored data

1. Mean values

SM2 adds up the sampled values at the end of each sampling cycle and computes a mean value for each monitoring cycle (e.g. the number of tasks) as follows:

$$\text{Average} = \frac{\text{Sum of sampled values}}{\text{Number of samples}}$$

2. Percentage frequency

At the end of each sampling cycle, SM2 records the occurrence of specific events or conditions and computes the percentage frequency of events (conditions) in relation to the total number of samples (e.g. channel utilization) as follows:

$$\text{Frequency} = \frac{\text{Sum of occurrences}}{\text{Number of samples}} * 100 \%$$

3. Ascending system counters and SM2 counters

The activity during the monitoring cycle can be determined by calculating the differences between values gathered at the beginning and the end of the monitoring cycle (e.g. idle time).

4. Computation

SM2 computes some data from monitored data, e.g. the summation values for all categories.

Uninterruptible clock resetting

SM2 works internally with UTC time. During the switch from daylight saving time to winter time and vice versa, SM2 can continue to work without interruption.

2.9 Monitored variables

The table below shows the main monitored variables of SM2.

A table showing the links between the monitored variables and monitoring programs, ONLINE reports, and SM2R1 report groups can be found in [chapter “Variables reports” on page 607](#).

Monitored variable	Description
Cache	
Access statistics and hit rates for DAB caches	Number of read/write accesses to DAB cache areas per second and the percentage frequency of read/write accesses to DAB subareas without disk access.
Access statistics for ISAM pools	Number of ISAM accesses per second to pages in the ISAM pool and directly to disk, as well as the number of ISAM accesses with a delay per second. Number of accesses to the index per second and the hits as a percentage when the index is accessed. Size of the ISAM pool and reserved pages in PAM pages.
Access statistics and hit rates for hiperfiles (PFA)	Number of read/write accesses to DAB cache areas per second, and their percentage share of the total number of read/write accesses. Number of failed attempts to use the cache per second.
Accesses and hit rates for Symmetrix systems	Number of read/write accesses to Symmetrix systems, pubsets and devices of Symmetrix systems. Proportion of read/write accesses handled successfully by the cache.
CPU	
CPU utilization	Time during which the processor is in one of the following statuses: TU, TPR, SIH, IDLE or STOP.
Number of system calls	Number of SVC calls in TU/TPR per second, and the sum of all SVC calls per second.

Table 2: SM2 monitored variables

(part 1 of 4)

Monitored variable	Description
Communication	
Response times	Average time in seconds between the receipt of a message on the system and the sending of a response to the application.
Transaction times	Average time in seconds between input and the last output.
Thinking time	Time between an output and the resulting subsequent input.
Wait time	Average wait time for incoming messages before being processed by the application.
Transaction rates	Number of transactions per second.
Length of transaction message	Average length in bytes of input/output messages for selected connection sets.
Number of inputs/outputs of communication devices	Number of read/write accesses per second, and the number of bytes transferred.
Network transmission rates	Transmission rate for all TCP/IP connections; data transfer information for specific connections.
Memory for data transfer	Current and maximum sizes of the resident memory; threshold value from which BCAM issues warnings
Databases	
Utilization in the SESAM/SQL database system	Number of transactions; number of SQL plan generations, number of read/write accesses to buffers
Utilization in the UDS/SQL database system	Number of CODASYL and SQL statements; number of transactions; number of read/write accesses; number of wait states
Files	
Catalog access statistics	Number of read/write accesses to catalog entries/JV entries of files per second, as well as average access times in milliseconds.
File access statistics	Number of PAM-WAIT, PAM-CHECK, PAM-READ, PAM-WRITE operations and input/output operations per second for one file. Distribution of input/output operations to files for selected disks.
File access times	Average access time in milliseconds for each input/output.

Table 2: SM2 monitored variables

(part 2 of 4)

Monitored variable	Description
IOs	
Number of input/output operations	Number of input/output operations per second.
Number of paging input/output operations	
Channel utilization and channel transmission rates	Number of input/output operations or transferred PAM blocks per channel.
Device utilization and transmission rates	Device utilization in % without/because of paging activities.
Length of device queues	Number of tasks waiting for the execution of inputs/outputs before the device.
Duration of input/output operations	Hardware operating time between start subchannel and device end for input/output operations of a specific device, measured in milliseconds.
Access distribution to PAM blocks of disks	PAM block addressed for input/output operation.
Memory/disk storage	
Main memory utilization	Number of 4-KB pages in main memory.
Utilization of paging area	Number of pageable 4-KB pages on the devices.
Utilization of virtual address space	Number of class 1 through class 4 pages in the virtual address space.
Page fault rate	Number of page fault interrupts per second.
Access statistics for global storage	Number of read/write accesses to global storage per second and the number of bytes transferred per second.
Utilization of pubsets	Capacity and utilization of SF pubsets and volume sets.

Table 2: SM2 monitored variables

(part 3 of 4)

Monitored variable	Description
Subsystem	
PCS data	Service rate of the affected categories in %, job delays and Service units per second.
UTM response times and transaction rates	Number of dialog and asynchronous transactions per second and the average time in seconds taken for the transaction.
Number and duration of send jobs via MSCF	Number of send jobs per second, as well as average total and wait times.
POSIX data	File accesses, message accesses, semaphore activities, buffer utilization and all types of system accesses per second.
Lock requests to the DLM	Number of enqueue, convert, dequeue and information lock requests per second, as well as the number of lock allocations and releases per second.
Data on synchronization functions in HIPLEX (NSM)	
HSMS data	Data on the migration of files to the background level and the retrieval of files to the processing level.
openFT data	Load values of selected openFT instances.
Task	
Number of tasks	Number of BATCH, DIALOG, TP and system tasks.
Length of task queues	Number of tasks in a category.
Dwell times in task queues	Dwell times in task queues per category.
Frequency of task activation and deactivation	
Task-specific utilization of resources	Service units per second, percentage CPU utilization, input/output operations per second, used pages in 4-KB pages, paging read per second.
Utilization and queues for task locks	Percentage utilization of a task lock and the number of tasks in the task lock queue.
VM2000	
Hypervisor activities (servers with /390 architecture)	Percentage active and idle times of the server.
Guest system activities	Planned and measured CPU utilization as a percentage.
CPU pools	Utilization of the CPU pools
VM groups (servers with /390 architecture)	Utilization of VM groups

Table 2: SM2 monitored variables

(part 4 of 4)

2.10 Use in a computer network

SM2 provides functions for the central online monitoring and control of monitored-data acquisition in a computer network.

The monitored data is independently recorded on each host in the network by a local SM2 monitoring program, and is exchanged between the hosts via a LAN. In this way, the online reports of all hosts can be output on any computer in the network.

There are also reports which output the combined monitored data of the various hosts.

All control statements for the acquisition of monitored data (changing monitoring cycles, switching monitoring programs, etc.) can be entered on any host for all computers in the network.

The SM2 functions for monitored-data acquisition in a network are also available in a HIPLEX network.

Prerequisites for using SM2 in computer networks

1. An MSCF connection of type CCS (Closely Coupled System) must exist between all hosts in the network.
For further information on this connection type, refer to the “HIPLEX MSCF” manual [8].
2. The SM2 functions for monitored-data acquisition in a network can only be used for those computers in the network which run the same version of SM2.
3. SM2 must be started once on all computers in the network (e.g. with START-SM2) or at least loaded via the DSSM command /START-SUBSYSTEM
SUBSYSTEM-NAME=SM2.

Selecting hosts with SM2

The privileged statement SELECT-HOSTS (see [page 128](#)) selects the computers to which all subsequently entered SM2 statements are to be sent. This statement always applies to the user who issues it. Following the END statement and a restart, the default value applies again, i.e. SM2 statements are only executed on the local host.

The SHOW-SELECTED-HOSTS statement (see [page 156](#)) provides information on the selected hosts. This includes the time of the last monitored variable queried.

Users in a computer network

SM2 distinguishes between privileged users (“primary” and “secondary” administrators and other privileged users) and nonprivileged users. This also applies in a computer network.

It does not distinguish between the rights of local and remote users.

Users can have different rights on different hosts, depending on their privileges and the rights they have already been assigned. In particular, the primary administrator of a host can come from a different host. These primary administrators can be identified on the MEASUREMENT STATUS screen by the additional output of the host name.

The set of administration statements available may be restricted according to the rights of the user. As a result, the set of permitted statements may vary from one host to the next. Prohibited statements are rejected by the relevant hosts. If several hosts are selected, the statements of the first specified host are displayed in the user dialog.

Screen output

The screens containing the local data of a particular host are output in succession. These are followed by the screens of the next host, and so on. The system name in the top left-hand corner of the screen indicates which host to which the screen belongs.

Before the local screens of the individual hosts, screens showing combined data from the whole network can be output (GLOBAL, NSM and SHARED PUBSET reports):

- The GLOBAL report provides an overview of the utilization of all hosts in the network.
- The NSM report outputs data of the NSM subsystem for all hosts in the network.
- The SHARED PUBSET report shows the access to shared disk devices.

These reports only cover the hosts which have been previously selected by means of the SELECT-HOSTS statement.

Standard SM2 output file

The standard SM2 output file is created on every host on the home pubset. To enable the SM2 output files of the various hosts to be distinguished later on, the file name also includes the system name.

Restrictions for the computer network

1. The BS2000 /START-TASK-MEASUREMENT command can only be used on the local host.
2. Scroll commands
 - Scrolling backward with - or **-R** is not supported across host boundaries; i.e. you cannot scroll back from the first report of one host to the last report of the previous host.
 - You can only change to the first report of the previous host using the scroll command **-N** (node) (see [page 164](#)).
 - You can only access the reports containing combined monitored data using **-N** while scrolling backward.
 - The scroll command **--** always displays the first selected report.
3. The monitored values of the individual hosts are only output together if the time also corresponds. The system clocks on all hosts should be synchronized as far as possible and the same monitoring cycle must be set in SM2.

If no valid or new data exists, the following outputs are possible:

- SHOW-SELECTED-HOSTS statement
In the LAST BUFFER column, RSLT NOT VALID is output for the last monitored variable queried instead of the time.
- GLOBAL and NSM reports
In the column for the monitored variables, *** is output.
- SHARED PUBSET report
The message SOME DATA MISSED is output.
- Local reports
For local reports, the message NO DATA FROM xyz is output.

Dynamic IO configuration change

SM2 recognizes a dynamic IO configuration change and, if need be, automatically modifies the range of objects monitored.

The following table describes the behavior of the monitoring programs when adding a device/channel in the context of a dynamic IO configuration change:

Monitoring program	Behavior when adding a device/channel
DISK-FILE	No action
SAMPLING-DEVICE	Monitoring program is terminated internally and then restarted
SERVICETIME	Monitoring program is not restarted, even if the device belongs to the monitoring program definition
SYSTEM	Device is involved in the measurement, if it belongs to the monitoring program definition
TASK	Monitoring program is not restarted, even if the device belongs to the monitoring program definition
CHANNEL-IO	Monitoring program is terminated internally and then restarted, if monitoring program definition is *ALL
SAMPLING-CHANNEL	Monitoring program is terminated internally and then restarted

If a device/channel is removed, no monitoring program is terminated internally and then restarted. The device/channel is anymore visible in the SM2 datas, but the measurement values are always zero.

Parallel Access Volumes (PAV)

SM2 supports maximally 7 alias devices of a base device.

3 System monitoring with openSM2

3.1 Performance expectations of the user

When assessing the performance of a DP system, the user is generally interested in:

- the time it takes for the system to respond to requests,
- how many requests the system can handle within a given period of time, and
- what resources are used.

The performance level of the hardware (e.g. the number of instructions processed per second) and of the operating system (e.g. task management efficiency) are generally of secondary interest.

Parameters for online use

A distinction is made between inquiry-and-transaction mode and timesharing mode:

In **inquiry-and-transaction mode** the terminal users can communicate only with programs which are specific to a given application. Normally, a large number of users interact with a relatively small number of application programs.

In **timesharing mode** the terminal users enter their own application programs for interactively processing a given task. The application programs are generally controlled by system programs for creating, testing and updating files and programs.

In both inquiry-and-transaction mode and timesharing mode, the unit of DP operation is a transaction.

The transaction time is the delay between the arrival of the input at the host and the output of the acknowledgment. The DP system can output several responses with different response times for a single transaction. If resources are not released when a transaction is completed but are reserved for the next transaction, the transaction is known as a multi-step transaction.

The sum of single-step and multi-step transactions required for handling one application is called an operation.

The characteristic criteria for describing performance expectations are:

- Transaction rate
Sum of all successfully completed transactions per unit of time
- Response time
Time required for processing by the DP system
- Number of simultaneously active users.

The monitored values for these characteristics are determined in SM2 by the monitoring program BCAM-CONNECTION ([page 52](#)), RESPONSETIME ([page 63](#)) and UTM ([page 75](#)).

Batch processing criteria

In batch processing, the unit of DP processing is a job.

- Throughput rate
Number of successfully processed jobs per unit of time
- Dwell time
Time required for processing a job

For the individual user, a satisfactory performance (short response time, short dwell time) is obtained if the required resources are available when they are requested. This requirement can easily be met if utilization of the requested resources (in particular of the input/output devices) is kept to a minimum.

The DP system performance, on the other hand, is characterized by the transaction and throughput rates, where maximum utilization of the resources is desirable for economic reasons.

These contradictory performance requirements relating to resource utilization can be fulfilled only by carefully planning the workload and the associated resource requirements.

3.2 Applications

We must consider two different applications:

- Trend monitoring (= long-term monitoring) for obtaining data for system planning
- Bottleneck analysis for locating and eliminating performance problems

The type of monitored data acquisition (frequency, scope) must be adapted to the application.

Trend monitoring

The utilization data of the following resources is required in order to carry out long-term system planning:

- CPUs
- channels
- devices
- main memory

Additional monitoring routines need not be activated.

It is advisable to use the following settings for monitoring periods:

Sampling cycle (SAMPLING-PERIOD):	1000 milliseconds
Monitoring cycle (OFFLINE-PERIOD):	5 minutes
Analysis subinterval	1 hour

The monitoring period should cover the entire period from SYSTEM READY through to SHUTDOWN. If output of the online screen report takes too long during the session, you can shorten the online monitoring cycle.

Monitoring times are set using the MODIFY-MEASUREMENT-PERIODS statement.

It is a good idea to create a new SM2 output file every day (OPEN-LOG-FILE / CLOSE-LOG-FILE statements). The SM2U1 routine can be used to combine (and split) daily SM2 output files to create one large file, known as the master SM2 output file. The daily SM2 output files must be added to the master SM2 output file in chronological order.

Bottleneck analysis

Before monitoring is started, you must clarify any performance problems, i.e. performance expectations that are not satisfied. The following problems may exist:

- System-oriented performance problems
These arise if the system throughput rate is unsatisfactory, and are indicated by a low transaction rate and/or throughput rate. The most likely cause is the overloading of one or more resources.
- User-oriented performance problems
These occur due to long delays when handling specific load requirements.

The following monitored variables should be used to analyze bottlenecks. openSM2 also allows for more extensive analysis through addition monitoring programs.

These monitored variables and monitoring programs make it easier to locate overloaded resources:

Monitored variable	Monitoring program
Number of tasks in the system queues and at devices	Monitored by default
Number of input/output operations per device	Monitored by default
Working set per category	Monitored by default
CPU utilization and number of input/output operations per category	SYSTEM
Number of input/output operations and volume of data transferred per channel	CHANNEL-IO
Access to catalog entries	CMS
Number of transactions	RESPONSETIME, BCAM-CONNECTION, and UTM

Table 3: Monitoring programs used to locate overloaded resources

The following settings are recommended for monitoring times (MODIFY-MEASUREMENT-PERIODS statement):

Sampling cycle (SAMPLING PERIOD):	400 milliseconds
Monitoring cycle (OFFLINE-PERIOD):	60 seconds
Analysis subinterval:	1 – 5 minutes
Monitoring period:	0.5 – 5 hours

Monitoring must be carried out during peak load periods.

Due to the shorter monitoring cycle and the activated monitoring programs, bottleneck analysis produces a large volume of data compared to trend monitoring. The volume of data corresponds to the number of objects monitored. The resulting SM2 output file may be very large.

Because of the volume of data generated, it does not make sense to copy all data record types into the master SM2 output file. SM2U1 can be used to suppress certain data records when updating the master SM2 output file.

To investigate delays when handling special load requirements, you will need further information in addition to the system utilization data described above. To begin with, the monitoring program PERIODIC-TASK or TASK can be used to select a task. The DISK-FILE monitoring program can be used for overloaded disks to determine the files accessed most frequently. It is not possible to list general guidelines for the additional selection of monitoring programs. For further information, please refer to the “Performance Handbook” [5].

3.3 Ensuring efficient DP system usage

To minimize performance analysis problems, the following procedure should be adhered to:

- Define the performance expectations (see [page 41](#)).
- Check the extent to which performance expectations are satisfied once productive operation has begun. This involves carrying out a “basis” test.
- If some performance expectations are not satisfied, check whether these are system-oriented or user-oriented.
- Check for errors in the system setup.
- Carry out a bottleneck analysis (see the section “[Automatic performance analysis](#)” on [page 394](#) and the description of the START-AUTOMATIC-ANALYSIS statement on [page 485](#)). Concentrate on those bottlenecks whose elimination promises the greatest performance improvement.
- After eliminating the detected bottlenecks, repeat the bottleneck analysis. Many bottlenecks remain hidden until after some kind of intervention.
- Monitor the system at regular intervals (see [page 43](#)) to detect saturation symptoms in the main resources (due to increasing loads) and to avoid critical system conditions.

The procedure is described in detail in the “Performance Handbook” [5], which also explains how to interpret the results.

3.4 Analyzing monitored data

There are several tools on various different platforms which can be used to process and display monitored data.

SM2 (BS2000)

At the end of each monitoring cycle, it outputs monitored data in the form of reports.

See the [chapter “SM2 screen output” on page 209](#).

SM2R1 (BS2000)

The program SM2R1 analyzes the SM2 output file. SM2R1 outputs statistics in tabular form and time series in chart form to a file which is ready to print.

See the [chapter “SM2R1 analysis routine” on page 391](#).

openSM2 Manager (SE servers)

The openSM2 Manager is the web-based user interface for performance monitoring of the SE servers.

With the openSM2 Manager the monitored data of all components and systems of an SE server and of other systems outside the SE server can be presented and monitored with user-defined rules. A summary presentation of the most important utilization values of all monitored systems provides a quick overview of the total utilization of the SE server. The monitored data can be arranged in a straightforward manner in various freely configurable report views. In addition to the presentation of the current monitored data, it is also possible to display historical data.

See [section “openSM2 Manager” on page 507](#).

ANALYZER (Microsoft Windows)

ANALYZER is a development of SM2R1. It makes use of the graphical options offered by a PC when displaying monitored data. SM2 output files from various BS2000 computers can be analyzed in a single analysis run. The result data can be displayed in tabular form and can be edited graphically. Charts can be copied into the clipboard and transferred to any Windows application that can access the clipboard.

See the [section “ANALYZER Analysis with a Windows PC” on page 512](#).

INSPECTOR (Microsoft Windows)

INSPECTOR is used to simultaneously display and monitor SM2 data from several BS2000 computers online. The monitored data is supplied in the form of charts which are updated at the end of every monitoring cycle. Rules can be defined for monitoring the monitored data in which actions are defined which are to be taken when certain conditions occur.

See the [section “INSPECTOR Online monitoring with a Windows PC” on page 519](#).

SM2-PA (BS2000)

The SM2-PA product can be purchased separately and is used to analyze user-specific SM2 output files from user task monitoring operations. The result data is displayed in the form of statistics and offers information on task resource utilization and/or the performance of user programs.

See the [section “SM2-PA Program analyzer” on page 522](#) and the “SM2-PA” manual [15].

4 SM2 monitoring programs

4.1 Overview

SM2 offers a range of optional monitoring programs that record additional data. A distinction is drawn between privileged monitoring programs, which only an SM2 administrator can start and stop, and monitoring programs for user-specific monitored objects, for which the permission of the SM2 primary administrator is required (see the MODIFY-USER-ADMISSION statement, [page 117](#)).

The monitoring programs are started and stopped by means of START-MEASUREMENT-PROGRAM and STOP-MEASUREMENT-PROGRAM. For some monitoring programs, monitoring parameters must be defined with the SET-<programname>-PARAMETERS statement and/or the objects to be monitored must be selected with ADD-... and REMOVE-... before the program is started.

The data supplied by the monitoring programs is output on screen and/or to the SM2 output file. To start monitoring programs that output their data exclusively to the SM2 output file, an SM2 output file must be open. All other monitoring programs also write their data to the SM2 output file when it is open. In the case of screen output with some monitoring programs, it is possible to use a SELECT-... statement to select additional, object-specific online reports.

The following table indicates the monitoring programs:

- for which parameters must be defined or the objects to be monitored must be selected
- that output their data exclusively to the SM2 output file
- that output data in an online report on the screen
- that output data in an online report only to privileged users
- for which additional, object-specific reports can be selected.

Monitoring programs	Parameter/ monitored object definition	Output exclu- sively to file	Online output		
			Name of the online report	Only privi- leged users	Selection of additional reports
BCAM-CONNECTION	x		BCAM CONNECTION		
CHANNEL-IO (4)(6)	x		BCAM MEMORY		
CMS (5)			CHANNEL		
COSMOS	x	x (1)	CMS		
DAB			DAB		x
DISK-FILE	x		DISK FILE		
DLM			DLM		
FILE			FILE	x	
GS			GS		
GSVOL			GSVOL		
HSMS					
ISAM	x		ISAM	x	
			ISAM FILE	x	
MSCF			MSCF		
NSM			NSM		
OPENFT	x		OPENFT		
PERIODIC-TASK (4)(5)	x		PERIODIC TASK		
PFA			PFA		
POSIX			POSIX		
PUBSET			PUBSET		
RESPONSETIME	x		RESPONSETIME		
SAMPLING-DEVICE (4)(6)(7)	x (3)		DEVICE DISK		
			DEVICE TAPE		
			DEVICE TD		
			DEVICE OTHER		
SESAM-SQL			SESAM SQL		
SERVICETIME	x	x			
STORAGE-SYSTEM	x		SYMMETRIX		x
SVC			SVC		
SYSTEM	x		CATEGORY		
TASK	x	x			
TCP-IP			TCP/IP		
TLM			TLM		
UDS-SQL			UDS SQL		
UTM			UTM		x
VM			VM		
			VM CPU POOL,		
			VM GROUP		

Table 4: Overview of the SM2 monitoring programs

(part 1 of 2)

Monitoring programs	Parameter/ monitored object definition	Output exclu- sively to file	Online output		
			Name of the online report	Only privi- leged users	Selection of additional reports
USER FILE USER ISAM USER TASK		x (2)	FILE ISAM		

Table 4: Overview of the SM2 monitoring programs

(part 2 of 2)

- (1) COSMOS outputs data to a separate SM2 output file (see the description of COSMOS, [page 58](#)).
- (2) Output is to a user-specific SM2 output file.
- (3) Parameter definition is optional because there is a default.
- (4) The online output can be controlled via sort criteria.
- (5) The volume of output information can be controlled.
- (6) The monitoring program is started for all channels when the SM2 subsystem is started.
- (7) The monitoring program cannot be terminated with the STOP-MEASUREMENT-PROGRAM TYPE=*ALL statement, but only with the STOP-MEASUREMENT-PROGRAM TYPE=*SAMPLING-DEVICE statement.

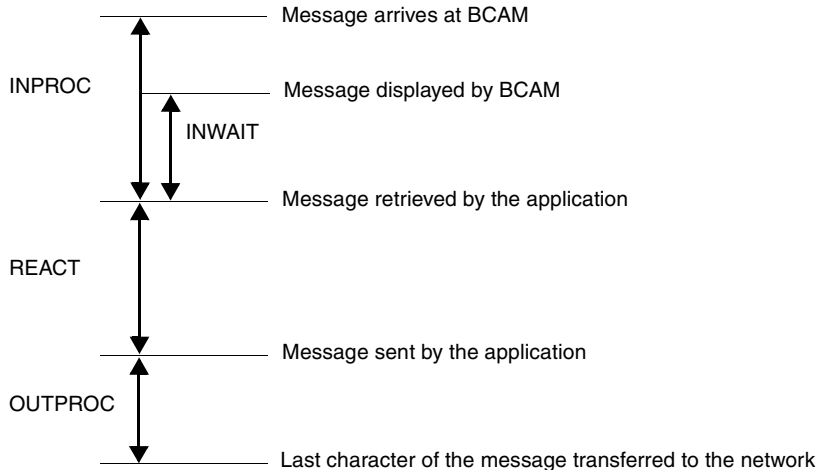
4.2 Privileged SM2 monitoring programs

BCAM-CONNECTION

Monitored data on connection sets

The BCAM-CONNECTION monitoring program supplies monitored data for connections that are combined in connection sets. It provides monitored variables for DCAM applications (with NEA names) and socket applications (with port numbers) only. Information is recorded on INWAIT, REACT, INPROC and OUTPROC times.

- The INWAIT time is defined as the time between the display of a BCAM message at an application and the retrieval of this message by the application.
- The REACT time is defined as the time between an application's send call and the receive call which immediately precedes it. If the send and receive calls are logically linked, then the REACT time can be seen as the time it takes for the application to respond. In dialog-oriented applications, each input is normally followed by the appropriate output.
- The INPROC time is defined as the time between the arrival of a message at BCAM and the retrieval of this message by the application. The INPROC time includes the INWAIT time.
- The OUTPROC time is defined as the time between a send call and the transfer of the last byte of a message to the network. Data transfer to the network may also be delayed by the receiving party.

Representation of time definitions

In addition, statistics are recorded on incoming and outgoing data (jobs, the volume of data, packets, etc.).

The SM2 administrator can use the SET-BCAM-CONNECTION-PARAMETERS and ADD-BCAM-CONNECTION-SET statements to define global and connection-specific monitoring parameters respectively.

The SET-BCAM-CONNECTION-PARAMETERS statement must be entered before the first ADD-BCAM-CONNECTION-SET statement.



Short connections are also recorded which are generated and terminated during the monitoring cycle. Furthermore, new connections are taken into account in the monitoring cycle in which they are generated and terminated connections in the monitoring cycle in which they are terminated.

BCAM keeps a connection open for up to 4 minutes even if a task has already closed the connection. From the SM2 viewpoint, too, such a connection is still open, but no further activity takes place on it.

It can also occur that a connection is rejected by the partner. No connection has been established. The number of connections of this type is determined and output.

Global parameters

Time statistics are recorded as an average value and bucket-specifically. The BUCKET parameters consist of a list of the maximum INWAIT, REACT, INPROC and OUTPROC times, arranged in order of magnitude. The ranges defined in these lists apply to all connection sets defined in the ADD-BCAM-CONNECTION-SET statement.

<integer 1 .. 99999> is an integer which defines a limit value. The number of limit values in this list implicitly determines the number of ranges. This may vary from list to list, but four limit values is the maximum.

The last value in the list of ranges specifies the overflow limit. All monitored variables that exceed this limit are included under the last range. The bucket parameters in BCAM can be specified by both SM2 and other applications. It is also possible to disable monitored data acquisition entirely.

The SET-BCAM-CONNECTION-PARAMETERS statement contains the *UNCHANGED operand value. This causes SM2 to leave the bucket parameters set in BCAM as they are, and to adopt these values for its own monitoring program definition.

If the SET-BCAM-CONNECTION-PARAMETERS statement is used to define alternative bucket parameters, then these will also be set in BCAM when you start the monitoring program. If you modify the bucket parameters in BCAM during the monitoring process, the new parameters will also apply to SM2.

If monitored data acquisition is disabled in BCAM during monitoring, SM2 will reenable it using the parameters defined in SM2.

Connection-specific parameters

A connection is defined by an NEA name or a port number, the type of connection, and the host name. A connection set is understood as a number of connections defined by specifying these objects.

The objects may be specified using

- a fully qualified name (name)
- a partially qualified name (name*)
- or an unqualified name (*ANY, *LOCAL, *REMOTE, *BOTH)

Using an unqualified name may lead to fluctuations in the number of connections in a connection set during monitoring. The BCAM-CONNECTION screen shows the number of connections of the connection set which are currently being monitored. A connection set is defined by specifying a list of up to 16 name definitions. You can specify either NEA application names or port numbers, but not both.

You must also specify the type of connection:

*LOCAL	Only connections within the host
*REMOTE	Only remote connections
*BOTH	All connections

Another distinguishing feature used to select connections is the names of the hosts between which the connection exists. It is possible to define up to 32 connection sets.

*Examples of monitoring program definitions**Example 1*

```
SET-BCAM-CONNECTION-PARAMETER
    INWAIT-BUCKETS = *UNCHANGED,
    REACT-BUCKETS  = *UNCHANGED,
    INPROC-BUCKETS = *UNCHANGED,
    OUTPROC-BUCKETS = *UNCHANGED

ADD-BCAM-CONNECTION-SET
    SET-NAME = setname1,
    CONNECTION-SELECTION = *BY-NEA-NAME (
        CONNECTION-NAME = *SPECIFIED (
            LOCAL-APPLICATION = $DIALOG,
            PARTNER-APPLICATION = *ANY)),
    CONNECTION-TYPE = *REMOTE,
    HOST-SELECTION = *SPECIFIED (
        LOCAL-HOST-NAME = *LOCAL,
        PARTNER-HOST-NAME = hostname1)
START-MEASUREMENT-PROGRAM TYPE = *BCAM-CONNECTION
```

In these statements, the monitoring parameters set in BCAM are transferred. A connection set is defined which records all connections with the name type NEA and the local application name \$DIALOG. The connections are remote and lead to the specified hosts.

Example 2

```
REMOVE-BCAM-CONNECTION-SET SET-NAME = *ALL
SET-BCAM-CONNECTION-PARAMETER
ADD-BCAM-CONNECTION-SET
    SET-NAME = setname2,
    CONNECTION-SELECTION = *BY-PORT-NUMBER(
        PORT-NUMBER = *SPECIFIED (
            LOCAL-PORT-NUMBER = 5000,
            PARTNER-PORT-NUMBER = *ANY)),
    CONNECTION-TYPE = *BOTH,
    HOST-SELECTION = *SPECIFIED (
        LOCAL-HOST-NAME = *LOCAL,
        PARTNER-HOST-NAME = hostname2)
CHANGE-MEASUREMENT-PROGRAM TYPE = *BCAM-CONNECTION
```

These statements end the current monitoring process and restart it using new parameters. The bucket parameters are set to their default values. A connection set is defined which records all connections with the local port number 5000 and which lead to the specified host.



Regardless of the defined connection sets, the BCAM-CONNECTION monitoring program also collects data on the size of the resident memory for data transfer.

CHANNEL-IO

Monitored data on the channel load

The CHANNEL-IO monitoring program provides precise information on the channel load. For each channel to be monitored, it returns the number of I/O operations for PAM block transfer, byte transfer, and without data transfer (NODATA). In addition, it indicates the volume of data transferred for PAM block and byte transfer.

The monitoring program CHANNEL-IO provides RSC IOs (servers with x86 architecture) of TD/openCRYPT-BOX devices and disks.

The monitoring program is started for all channels when the SM2 subsystem is started.

- servers with /390 architecture
For channels of the type FC (FibreChannel) the correct type is output.
- servers with x86 architecture
The FibreChannel peripherals with BS2000 are operated by emulating the devices in X2000 as bus peripherals. The bus periphery is specified as type BUS in SM2.

The real, physical FibreChannels can be different from the channels that are visible for the BS2000 operating system (and for SM2).

The SM2 channel-specific measurement data is not meaningful. If summation values are formed for all channels, valid values are still obtained.

Further information to the measurement values is in [section "CHANNEL report" on page 235](#).

CMS

Monitored data on the catalog management system

The CMS monitoring program provides information on the performance of the catalog management system. It outputs monitored variables for each catalog and for all private disks.

COSMOS

Monitored data on the system for bottleneck analysis

COSMOS is an event-driven monitoring program for recording detailed system data that is required particularly for bottleneck analyses of OLTP applications. It collects monitored data for all tasks, for task groups (which can be selected by user ID, category, or job name), or for individual tasks. Because of the high volume of very detailed monitored data involved, COSMOS should only be used for short-term monitoring. Intimate knowledge of the internal workings of the system is required to interpret the monitored data.

COSMOS is an SM2 add-on product and only available on special release. A detailed description is supplied with the product.

DAB

Monitored data on DAB activities

SM2 outputs highly detailed information on DAB activity in all DAB cache areas.

DISK-FILE

Monitored data on file access for selected disks

This monitoring program provides information on physical access to files on a selected disk, provided that access was by way of the privileged I/O output system PPAM. This includes file access using the methods EAM (Evanescent Access Method), ISAM (Indexed Sequential Access Method), SAM (Sequential Access Method) and UPAM (User Primary Access Method).

In the case of multiprocessor systems, monitoring is restricted to I/O operations originating in the local system running SM2.

If the monitoring program is used for parallel access volumes (PAV), the following applies:

- Basic devices and alias devices can be selected.
- If a basic device is selected, the data refers to the assigned alias devices.

Notes

- For each disk the DISK-FILE monitoring program writes only the 20 file names with the most inputs/outputs in the monitoring cycle to the SM2 output file. This means that in monitoring cycles with few inputs/outputs, null values may be output in report 320.
- The file name must always have the format :<catid>:<\$userid>.<filename>. A different format can be used for special disk accesses (e.g. from DAB).
- *OVERRUNS is a reserved file name. Here all inputs/outputs which could not be assigned to a specific file are counted.



Due to the extension of the I/O path, the monitoring program should only be used for short periods of time.

DLM

Monitored data on lock requests

The DLM (Distributed Lock Manager) monitoring program provides information on lock requests from TU, TPR and NSM.

FILE

Monitored data on file access for selected files

This monitoring program provides information on the number and average duration of file access, provided that access was by way of the privileged I/O output system PPAM. This includes file access using the methods EAM (Evanescent Access Method) and ISAM (Indexed Sequential Access Method).

In addition, data is recorded for the access methods SAM (Sequential Access Method) and UPAM (User Primary Access Method), all disk access, and all access to magnetic tape files with standard blocks.

The FILE monitoring program also records the access times to files. The monitored time covers the logical duration of I/O operations between start and end of a job from the point of view of the software. For asynchronous I/O operations, this time can be considerably greater than the hardware time.

In the case of multiprocessor systems, monitoring is restricted to I/O operations originating in the local system running SM2.

Overall, a maximum of 32 files can be monitored with this monitoring program.

GS

Monitored data on access to global storage

This monitoring program provides information on read and write access to global storage (GS), categorized by partitions.

GSVOL

Monitored data on volumes in global storage

This monitoring program provides information on read and write access to volumes which are emulated on global storage.

HSMS

Monitored data on file migration

This monitoring program supplies information on the migration of files to the background level and the retrieval of files to the processing level.

ISAM

Monitored data on ISAM pools

When the access method NK-ISAM is installed, the monitoring program for ISAM pools provides information on the performance of the ISAM buffer management facility. ISAM buffer management is aimed at reducing the physical I/O operations involved in accessing ISAM files. Using the information recorded in this monitoring program, the access behavior of this ISAM pool can be optimized by modifying the parameters of an ISAM pool, as required (for more details on the ISAM access method, please refer to the “DMS Macros” manual [2]).

Up to 16 ISAM pools (defined via pool name, catalog ID and scope) and in addition the pool data of up to 16 ISAM files can be monitored.

MSCF

Monitored data on communication between hosts

The MSCF (Multiple System Control Facility) monitoring program provides information on communication between the local host and other hosts.

NSM

Monitored data on HIPLEX MSCF®

The NSM (Node Synchronization Manager) monitoring program provides information on the basic functions of a HIPLEX MSCF® network.

OPENFT

Monitored data on openFT

The OPENFT monitoring program supplies data on file transfer with openFT.

openFT’s monitored data acquisition must have been enabled beforehand using the openFT command `MODIFY-FT-OPTIONS . . . , MONITORING=*ON`. See the openFT manual “Installation and Administration” [9].

The SM2 administrator can use the `ADD-/REMOVE-OPENFT-INSTANCE` statements to define up to 16 openFT instances which are to be monitored.

PERIODIC-TASK

Monitored data on tasks

This monitoring program records the most important utilization data for all tasks. The privileged statement SET-PERIODIC-TASK-PARAMETERS is used to specify the tasks whose data is to be written to the SM2 output file. This prevents the SM2 output file from becoming too large.

Nonprivileged users receive information only on those tasks that are under their user ID.



The support of RSC IOs (servers with x86 architecture) of TD/openCRYPT-BOX devices is not possible, because there is no connection between RSC IOs and tasks.

PFA

Monitored data on caches under User PFA

The PFA monitoring program records monitored data on caches under User PFA.

User PFA (User Controlled Performant File Access) allows users to assign performance attributes to files, thus converting them into hiperfiles (High Performance Files). The idea behind hiperfiles is that file access is accelerated by buffering the file in a high-speed cache in order to avoid the slow disk input/output operations. Main memory (MM) or global storage (GS) are used as cache media for HIPERFILES.

POSIX

Monitored data on POSIX

This monitoring program provides information on the POSIX subsystem.

Data corresponding to the sar options a, m, b and c is output.

PUBSET

Monitoring data for SF pubsets and volume sets

The PUBSET monitoring program records the utilization of SF pubsets and volume sets. Furthermore, the number of volumes, the saturation level and the capacity are output.

The SF pubsets and volume sets must be imported.

Volumes with an allocation lock (restricted volumes) are not taken into account.

RESPONSETIME

Monitored data on the BCAM pool

The RESPONSETIME monitoring program records information on response times, think times, transactions times, and wait times in the BCAM pool. It supplies useful monitored data for dialog-oriented DCAM applications only (with NEA names).

The **response time** is understood to mean the time between the receipt of a message by the host and the sending of a response to the partner application.

SM2 makes a distinction between two kinds of response time:

Definition 1 The time between an input (receipt of a message) and the next output (sending of a message) via the same connection.


Definition 2 The time between an input and the next output or all follow-up outputs sent before a further input.

The SM2 administrator can choose between these two definitions (using the SET-RESPONSETIME-PARAMETERS or MODIFY-RESPONSETIME-PARAMETERS statement).

The **think time** is understood to mean the time taken between output and the next (subsequent) input. The **transaction time** is understood to mean the total time elapsed between an input and the final output. The **wait time** is understood to mean the time between entry of a message in the BCAM pool and retrieval of the message from this pool.

For simplicity, **interaction** is the term used as an overall definition covering the wait, response, think, transaction times.

Example: Calculating the time values



T_I Input time stamp

T_O Output time stamp

T_R Time stamp of the current receive action in the BCAM application

$RT_1 = T_{O1} - T_I$ Response time (definition 1)

$RT_2 = T_{O2} - T_{O1}, T_{O3} - T_{O2}, T_{O4} - T_{O3}$ Response time (definition 2)

$TrT = T_{O4} - T_I$ Transaction time

$ThT = T_{I,next} - T_{O4}$ Think time

$WT = T_R - T_I$ Wait time

The outputs at times $T_{O1,2}$ cannot refer to the input at time T_I because the latter was not retrieved until time T_R .

The outputs at times $T_{O3,4}$ can (but need not) refer to the input at time T_I .

Because the monitoring is performed in the BCAM transport system, no definite statement can be made about the logical (application-specific) connection between the individual inputs and outputs.

With dialog-oriented applications, each input is generally followed by the associated output, so the correct response times can be established.

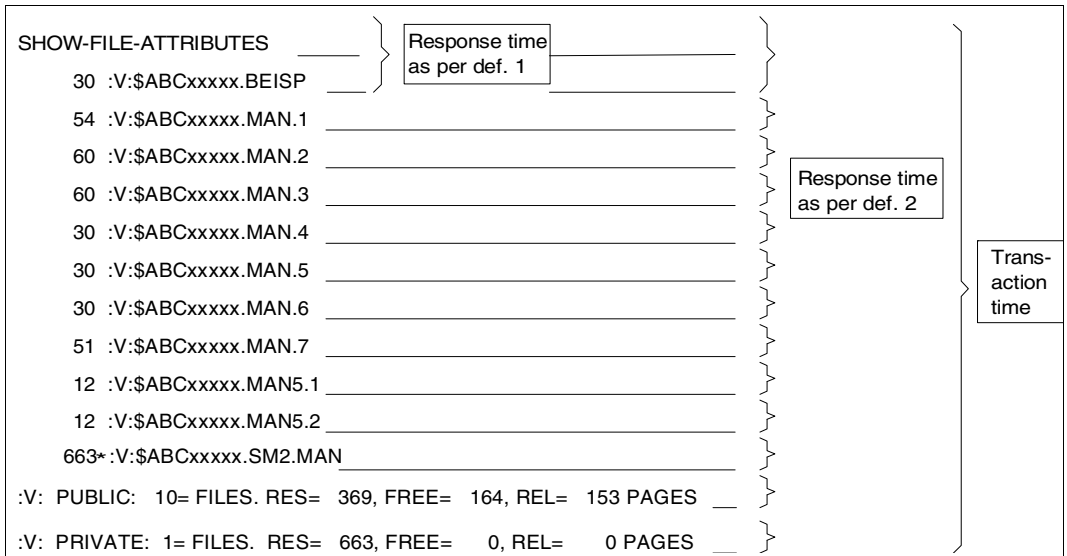
SM2 interprets each message from a partner application to a BCAM application as an input and each message from a BCAM application to a partner application as an output. No distinction is made as to whether the communications partner of the BCAM application is a terminal or an application in a partner system.

In order to restrict monitoring to individual connections, SM2 employs the BCAM application name. SM2 assumes that the transmitter or receiver (terminal) has been uniquely assigned to a specific application name.

There are some special applications for which this application name has not been defined. Specific response time measurement by SM2 is not suitable for such cases, but the corresponding values are recorded in the total.

Wait times within the communications host and data transfer times cannot be taken into consideration.

Example: Responses to a /SHOW-FILE-ATTRIBUTES command



The SM2 administrator can use the SET- or MODIFY-RESPONSETIME-PARAMETERS statement and the ADD-CONNECTION-SET statement to define global and connection-specific monitoring parameters respectively.

Global parameters

If response times are recorded for specific categories (SCOPE=*CATEGORY), the calculation of means values per category only involves values which are below the maximum corresponding bucket value. If no bucket-specific analysis is required, only one value is required to determine the overflow limit. If more than one bucket has been defined, values below this limit are ignored as being irrelevant, and the category assignment valid at the end of interaction remains applicable

Using the parameter CONNECTION-NUMBER=<integer 1..8187>, the maximum number of connections to be monitored is defined. Once this number is reached, response time monitoring continues, in which case SM2 will react to any unrecorded interaction in the following manner:

- The message NOTE: CONNECTION NUMBER TOO LOW will be issued in the SM2 output report. Each of the RESPONSETIME screens (see description of the RESPONSETIME report) will show the same message. The message is output only if no interactions were recorded during the monitoring cycle in question.
- An indication to this effect will be given in the SM2 response record at the end of this monitoring cycle.

- With SM2R1, the message SOME INTERACTIONS MISSED will be issued at the end of the analysis. This message is output only if no interactions were recorded during the analysis period in question.

The BUCKET parameters consists of a list of the maximum response, think, transaction and wait times, arranged in order of magnitude. The ranges defined in the specified lists apply to all connection sets specified in the ADD-CONNECTION-SET statement.

<integer 1 .. 99999> is an integer which defines the limit value. The number of limit values in this list implicitly determines the number of ranges. This may vary from list to list, but five limit values is the maximum. The last value in the list of ranges specifies the overflow limit. All monitored variables that exceed this limit value are included under the last value, but are not included in the calculation of mean values given in the RESPONSE TIME report.

Connection-specific parameters

The SM2 administrator can monitor the response time both globally, i.e. for all possible connections, and for specific individual connections, connection groups, or connection sets.

A connection is uniquely defined in the local the host by specifying the following objects:

- the name of an application
- the name of a partner
- the host name of the partner
- the type of connection

A connection group is a number of connections that are defined by the specification of the application and of the type of the connection. The partner name and computer name of the partner cannot be specified. For example, (\$DIALOG) defines a connection group which includes the \$DIALOG application.

A connection set is understood as a combination of connection groups. It is defined by specifying a list of up to five connection groups.

For example, ((\$DIALOG),(UTM1)) defines a connection set which combines two connection groups with the \$DIALOG and UTM1 applications.

This monitoring program supplies statistics for each connection set. In addition, the set of all connections is monitored on an ongoing basis. When selecting which connections are to be monitored, the SM2 administrator can define up to 15 connection sets. (*ALL) is always set automatically as the additional connection set. Furthermore, the SM2 administrator must specify whether the selected connections are to be monitored or excluded from monitoring using the *BY-CONNECTION or *EXCEPT-CONNECTION operand of the ADD-CONNECTION-SET statement. With *EXCEPT-CONNECTION, all connections except for those in the specified connection groups are monitored.

The SM2 administrator can also specify the type of the connection. If *LOCAL is specified, connections within the host are taken into account. This makes it possible to record the response times for individual applications, even if it would normally appear to the system that everything is running under the OMNIS carrier application.

Each defined connection set can consist of a maximum of five connection groups. The total number of connection groups which can be defined is, in turn, limited to 15 at most.

Examples of the RESPONSETIME monitoring program

Example 1

```
SET-RESPONSETIME-PARAMETER
ADD-CONNECTION-SET SET-NAME=setname1,SET-DEFINITION=*BY-CONNECTION,
  CONNECTION-SET=((utm1),(utm2),(utm3))
ADD-CONNECTION-SET SET-NAME=setname2,SET-DEFINITION=*BY-CONNECTION,
  CONNECTION-SET=(utm3),CONNECTION-TYPE=*BOTH,
ADD-CONNECTION-SET SET-NAME=setname3,SET-DEFINITION=*EXCEPT-CONNECTION,
  CONNECTION-SET=((utm1),(omnis),(apl1))
START-MEASUREMENT-PROGRAM TYPE=*RESPONSETIME
```

The above statements are used to explicitly define three connection sets: the first by specifying a list of three connection groups. These three connection groups are taken as a whole. The second connection set takes account of local and remote connections. All the other connection sets only take account of remote connections (≅ default setting). The third connection set combines all connections which belong to neither the UTM1 or OMNIS application nor to the explicit connection (apl1).

Example 2

```
REMOVE-CONNECTION-SET SET-NAME=*ALL
MODIFY-RESPONSETIME-PARAMETERS DEFINITION=*2
ADD-CONNECTION-SET SET-NAME=setname,SET-DEFINITION=*EXCEPT-CONNECTION,
  CONNECTION-SET=($OMNIS)
CHANGE-MEASUREMENT-PROGRAM TYPE=*RESPONSETIME
```

The above statements terminate the RESPONSETIME monitoring program (apart from *GLOBAL) and restart it, monitoring all connections with the exception of the \$OMNIS application. In addition, the response time as defined in definition 2 is monitored. All other monitoring parameters from the response time monitoring process just terminated are retained.

Example 3

```
SET-RESPONSETIME-PARAMETERS CONNECTION-NUMBER=100,RESPONSETIME-BUCKETS=  
  (50,100,200),THINKTIME-BUCKETS=(50,100,150,200,250)  
START-MEASUREMENT-PROGRAM TYPE=*RESPONSETIME
```

The SET-RESPONSETIME-PARAMETERS statement is used to automatically define a connection set with the name *GLOBAL; all connections in this set are monitored (CONNECTION(*ALL)). In addition, here the number of connections is limited to 100. Three limits (50,100,200) are set for the response time, and five limits (50,100,150,200,250) for the think time. The transaction time and the wait time in the BCAM pool are set to their default values.

SAMPLING-DEVICE**Monitored data on I/Os, data volume and utilization of devices**

The SAMPLING-DEVICE monitoring program measures the number of I/Os, the data volume and the utilization of the devices. Monitored data of the IORM subsystem is also recorded. For disk devices the hardware and software service times can also be output.

RSC IOs for disks (servers with x86 architecture) are gathered and gone into the IOs for read and write. Additionally the average number of parallel handled RSC IOs is a new value.

RSC IOs for TD/openCRYPT-BOX devices (servers with x86 architecture) are gathered and gone into the IOs for read and write.

The monitoring program is started when the SM2 subsystem is started without the recording of service times.

The monitoring program cannot be terminated using the STOP-MEASUREMENT-PROGRAM TYPE=*ALL statement, but only with the STOP-MEASUREMENT-PROGRAM TYPE=*SAMPLING-DEVICE statement.

If the monitoring program is used for parallel access volumes (PAV), the following applies:

- Only basic devices are output in the DEVICE DISK report.
- The monitored data refers to the basic device and the assigned alias devices (summation, averaging).

If an alias device changes from one basic device to another, then the datas of the alias device are not included in the datas of any basic device in the actual measurement interval.

The data of the alias devices is also output at the SM2GDAT data interface.

SERVICETIME

Monitored data on the service time

The SERVICETIME monitoring program allows the service times of individual devices to be measured in detail. A check is carried out to establish that the disk devices are “attached”. The service times are recorded by DCS (Dynamic Channel Subsystem).

Distinctions are drawn between the following service times (see “service times definition” in the glossary on [page 671](#)).

Service time	System	Definition
DQT Device queue time	SM2	Time from the I/O chain to the subchannel start (waiting time before the device)
FPT Function pending time	DCS	Time from the subchannel start to I/O start (waiting time for the assignment of a path; part of the hardware service time)
DDT Device disconnect time	DCS	Physical positioning time (part of the hardware service time)
DCT Device connect time	DCS	Data transfer time (part of the hardware service time)
RST Remaining service time	SM2	Time from channel end to device end (part of the hardware service time) or if DCS delivers no data: time from subchannel start to device end (hardware service time)

Table 5: Service times

In certain circumstances (see the notes below), DCS cannot provide any detailed service times. The service time is subdivided into the waiting time before the device and the hardware service time. The hardware service time is referred to as the remaining service time.

The data of the monitoring program is only output to the SM2 output file.

Note on Bus peripherals

DCS cannot deliver measurement data of devices serviced by this interface.

Notes on VM2000

- Under VM2000 the recording of detailed service times of DCS can be started by only one guest system. If a second guest system attempts to activate the function as well, a warning is issued. The monitoring program is started, but detailed service times of DCS are not delivered.
- The VM2000 command `/SHOW-VM-RESOURCES INFORMATION=*STD/*ALL` can be used to check whether SERVICETIME measurement is already active in a guest system. If it is, message VMS2035 is output.
- DCS cannot deliver detailed service times for disks with indirect I/O.

Note on servers with x86 architecture

DCS does not deliver detailed service times on servers with x86 architecture. Only the times measured by SM2 itself (device queue time, remaining service time) can be recorded.

Notes on disk storage systems

- Given an overall hit rate of 100%, the device disconnect time is zero.
- The following applies to parallel access volumes (PAV):
 - Basic and alias devices can be selected.
 - Each device defined is measured locally. In particular, no monitored data of an alias device is added to a basic device.

SESAM-SQL

Monitored data for the SESAM/SQL database system

The SESAM-SQL monitoring program supplies monitored data of the database system SESAM/SQL-Server.

See also the SESAM/SQL “Database Operation” manual [13].

Prerequisite in SESAM/SQL

To transfer statistical data from SESAM/SQL to openSM2, start SESMON in batch mode:

```
/START-SESAM-PERF-MONITOR  
//SET-MONITOR-OPTIONS . . . ,OUTPUT=*SM2
```

When `OUTPUT=*SM2` only one DBH can be specified per SESMON instance. A new SESMON instance must be started each time data of another DBH is to be output.

The interval at which SESMON transfers the data to openSM2 is automatically set to approx. 30% of the SM2 monitoring cycle. It is not possible to set this value manually.



The monitored data is supplied by SESAM/SQL asynchronously to openSM2 and applies for one or more cycles which are defined by SESAM/SQL and need not match the SM2 cycle exactly. Here differences can exist in the duration of the cycles, and time displacements can also exist between the SESAM/SQL and SM2 cycles.

The duration of one or more SESAM/SQL cycles is used to standardize the monitored data to one second. The data is therefore exact, but it only matches the SM2 cycle to a certain degree.

STORAGE-SYSTEM

Monitored data for Symmetrix systems

Monitored data is only collected for directly accessible, generated controllers of Symmetrix systems.



In BS2000 the software product SHC-OSD is required. You require the software product openSM2 (Open Systems) to monitor ETERNUS DX storage systems.

The monitoring program STORAGE-SYSTEM gathers also the meta volume information for Symmetrix systems. The information is only output at the SM2GDAT data interface and written into SM2 output file.

The SM2 online reports and the SM2R1 reports do not provide the meta volume information of Symmetrix systems. This is only supported from INSPECTOR and ANALYZER.

SVC

Monitored data on SVC calls

This monitoring program records all the SVC calls in the system.

The monitored data returned will be of interest only to performance specialists, since detailed knowledge of the internal workings of the system is required to interpret it.

SYSTEM

Global system and category-specific monitored data

The SYSTEM monitoring program provides important system-global and category-specific data. The most important aspect of these variables is that they can be obtained using event-driven methods. These monitored values make it possible to assign values which are otherwise global by default (CPU, I/O) to individual categories, i.e. the category-related share of the load can be determined. For precise determination of category dilation, the SET-SYSTEM-PARAMETERS statement is required to include all devices in the monitoring process. In addition, detailed monitored data on queues is recorded. The monitored data is written to the SM2 output file; some of the values are displayed on screen in the CATEGORY report.

The support of RSC IOs (servers with x86 architecture) of TD/openCRYPT-BOX devices is not possible, because there is no connection between RSC IOs and tasks.

If the monitoring program is used for parallel access volumes (PAV), the following applies:

- Basic and alias devices can be selected.
- The monitored data refers to the basic device and the assigned alias devices. If a basic device gets a new alias device, then the alias device is involved in the measurement.

TASK

Task-specific monitored data

The purpose of the monitoring program is to provide task-specific variables, from which the resource requirements of individual tasks can be determined.

This data can be used for the following purposes:

- to determine which tasks most frequently use specific resources
- to optimize operation by relocating loads
- as input data for simulation models or analytic calculations.

SM2 writes this data to the SM2 output file, from which subsequent analysis is possible. As with SM2 reports, data cannot be output during the monitoring process.

When defining the monitoring program, the tasks to be monitored must be specified. Both user tasks and system tasks can be monitored. SM2 can monitor both current tasks already executing or further tasks created in the course of monitoring operations.

The support of RSC IOs (servers with x86 architecture) of TD/openCRYPT-BOX devices is not possible, because there is no connection between RSC IOs and tasks. Instead of this the number of physical accesses to coded files are output.

Monitoring information can be written to the SM2 output file at two points in time:

1. At task termination time, if the task is terminated during monitoring.
2. At monitoring termination time, if the task is still executing when the monitoring program is terminated.

If the devices to be monitored are defined in the DEVICES operand of the SET-TASK-PARAMETERS statement, the number and service time of I/O operations for these devices are output for each task according to whether they are hardware- or software-related. Due to the time and effort involved in gathering this monitored data, this monitoring program should only be used for brief periods of time (e.g. one hour) or for a small number of tasks. For each task, it is possible to monitor up to 64 devices.

If the monitoring program is used for parallel access volumes (PAV), the following applies:

- Basic and alias devices can be selected.
- The monitored data refers to the basic device and the assigned alias devices (summation, averaging). If a basic device is selected, the data also applies to the assigned alias devices.
If a basic device gets a new alias device, then the alias device is involved in the measurement.

TCP-IP

Monitored data on TCP/IP connections

This monitoring program provides information on TCP/IP connections. IPv4 as well as IPv6 connections are recorded.



A large number of connections can result in high CPU utilization by the SM2 monitoring task.

TLM

Monitored data on locks

The TLM (TASK LOCK MANAGER) monitoring program records monitored data on locks managed by the Task Lock Manager. The occupation and average queue length of each lock are determined.

The monitored data returned will be of interest only to performance specialists, since detailed knowledge of the internal workings of the system is required to interpret it.

UDS-SQL

Monitored data for the UDS/SQL database system

The UDS-SQL monitoring program supplies monitored data of the UDS/SQL database system. See also the UDS/SQL “Database Operation” manual [16].

The monitoring program can only collect data which has been transferred to SM2 by a UDS monitor instance. Data transfer is initiated either by means of the `MEDIUM=S,n` statement when the UDS monitor is started or by means of the `INFORM-PROGRAM MSG='ADD MEDIUM=S,n'` command during operation. It can be terminated again using the `INFORM-PROGRAM MSG='FINISH MEDIUM=S'` command

The interval at which the UDS monitors supplies the data to SM2 is defined in seconds ($5 \leq n \leq 999$) with n . This should be considerably lower than the monitoring cycle set in SM2 so that data can be transferred several times within an SM2 monitoring cycle.



The monitored data is supplied by UDS/SQL asynchronously to openSM2 and applies for one or more cycles which are defined by UDS/SQL and need not match the SM2 cycle exactly. Here differences can exist in the duration of the cycles, and time displacements can also exist between the UDS/SQL and SM2 cycles.

The duration of one or more UDS/SQL cycles is used to standardize the monitored data to one second. The data is therefore exact, but it only matches the SM2 cycle to a certain degree.

UTM

Monitored data on openUTM applications

The UTM monitoring program records application-specific monitored data.

The following conditions have to be fulfilled:

- The UTM-SM2 subsystem is started.
- The UTM applications are delivering data.
- BS2000 accounting is started (only for DB-specific user values).

The UTM-SM2 subsystem is started automatically.

The UTM application must be ready to deliver data to SM2.

In the KDCDEF run (operand MAX SM2=ON / OFF / NO), it is possible to influence the transfer of monitored data.

- SM2=OFF (default setting)
No monitored data is delivered.
If you use the UTM administration interface with KDCAPPL and SM2=ON, the transfer of monitored data can be activated subsequently for each UTM application.
- SM2=ON
Data is delivered.
Additional administration is not required.
- SM2=NO
No data is delivered. It is not possible to start the transfer of data subsequently.

DB-specific utilization figures (UDS and SESAM only) are only supplied if the following conditions are met:

- BS2000 accounting is active.
- The UTM accounting record UTMA is switched on (using the BS2000 command MODIFY-ACCOUNTING-PARAMETERS SET-RECORD-TYPE=UTMA).
- UTM Accounting is switched on (using the KDCAPPL statement, parameter ACC=ON).
- Monitored-data acquisition is enabled in SESAM (using the statement ACC,TP=ON,CPU).

You will find more detailed information on this topic in the openUTM manuals.

VM CPU shares of guest systems under VM2000

The VM monitoring program collects the monitored data of virtual machines, CPU pools and VM groups of a VM2000 system.

Only on the monitor system are the values for all virtual machines, CPU pools and VM groups supplied.

On servers with x86 architecture the data for the VM2000 hypervisor and for the VM groups are dispensed with.

4.3 Nonprivileged SM2 monitoring programs

The following instructions must be heeded when using nonprivileged monitoring programs. You should also refer to the descriptions of the corresponding privileged monitoring programs for a better understanding of monitoring-program-specific features and variants.

FILE

Monitored data on file access

With this monitoring program, a user can only include files in the monitoring process if the SM2 primary administrator has permitted the monitoring of file access values. This permission can be granted either to all users or to users with certain IDs.

In this case, the SM2 MEASUREMENT STATUS screen will have the entry FILE in the line entitled USER MEASUREMENTS ALLOWED.

Each authorized user can include and exclude files which are set up under his own ID. It is also possible to include files which do not exist. Privileged users are not treated differently from normal users.

Any monitored files are excluded from the monitoring process once SM2 is terminated.

All together, up to 32 files can be monitored by this monitoring program; no values for these files are written to the SM2 output file.

ISAM

Monitored data on ISAM pools

With this monitoring program, a user can only include ISAM pools in the monitoring process if the SM2 primary administrator has permitted ISAM pool monitoring. This permission can be granted either to all users or to users with certain IDs. In this case, the SM2 MEASUREMENT STATUS screen will have the entry ISAM in the line entitled USER MEASUREMENTS ALLOWED.

Each authorized user can include or exclude global ISAM pools and all task-local ISAM pools which have been created or are being created under his user ID.

SM2 does not check whether an ISAM pool exists when a user tries to include it; in other words an ISAM pool can be included in the monitoring process even before it has been created. If an ISAM pool which has been included does not exist during the monitoring cycle, all its monitored data will be listed with blanks. Since the total number of ISAM pools that can be monitored by SM2 for nonprivileged users is limited to 16, care must be taken not to unnecessarily include too many non-existent ISAM pools in the monitoring operations.

No values are written to the SM2 output file for ISAM pools monitored with the aid of this monitoring program.

Indirect selection of a pool by specifying a file contained in the pool is reserved for the SM2 administrator.

TASK

Task-specific monitored data

A user can start this monitoring program only if the primary SM2 administrator has permitted user-specific task monitoring. This permission can be granted either to all users or to users with certain IDs. In this case, the SM2 MEASUREMENT STATUS screen will contain the entry TASK in the line entitled USER MEASUREMENTS ALLOWED.

With the /START-TASK-MEASUREMENT command, the user is then given the option of including tasks under his/her own ID for monitoring by SM2. This task information is stored in a user-specific SM2 output file and it can then be analyzed using the SM2-PA analysis routine (see the "SM2-PA" manual [15]).

5 SM2 operation

5.1 Starting and terminating SM2

Start SM2

The SM2 user program is called with the BS2000 command `START-SM2`.

Format

START-SM2
<code>VERSION = <u>*STD</u> / <product-version mandatory-man-corr> / <product-version mandatory-man-without-corr> / <product-version without-man-corr></code>
<code>,MONJV = <u>*NONE</u> / <filename 1..54 without-gen-vers></code>
<code>,CPU-LIMIT = <u>*JOB-REST</u> / <integer 1..32767></code>

Operands

VERSION =

Specifies the version of SM2 to be called.

VERSION = *STD

The current version is called.

VERSION = <product-version>

The specified version is called.

MONJV =

Specifies the name of the job variable to be monitored in the SM2 run. The job variable must already be cataloged.

MONJV = *NONE

No job variable is specified.

MONJV = <filename 1..54 without-gen-vers>

Specifies the name of a job variable already cataloged.

CPU-LIMIT =

Specifies the CPU time required for the SM2 run. In interactive mode, the user is notified by the system if this time is exceeded. In batch mode, the SM2 run is terminated.

CPU-LIMIT = *JOB-REST

There is no limit on the CPU time available to the program.

CPU-LIMIT = <integer 1..32767>

Specifies the CPU time required for the SM2 run.

When the SM2 program is started the following messages can be output:

- BLS0500 PROGRAM 'SM2', VERSION '<version>' OF '<date>' LOADED
This message is only output if SM2 is started with the
/START-(EXECUTABLE-)PROGRAM SM2 command. It provides information on the version of
the SM2 prephase (loader and starter of the SM2 program).
- NPSLOAD Program 'SM2', Version '<version>' of '<date>' loaded from file
' :ZZZZ:\$TSOS.SYSLNK.SM2.<ver>.SM2'
This message provides information on the SM2 program version (nonprivileged SM2
program).
- NPS0555 VERSION OF SM2 SUBSYSTEM (TPR-PART) IS <version>
This message provides information on the SM2 subsystem version (privileged SM2
program). This message is only output for privileged SM2 users.

The two last digits for the version in the NPSLOAD and NPS0555 messages indicate the current revision level of the nonprivileged/privileged SM2 program.

Terminate SM2

The SM2 user program is terminated using the END statement.

- The automatic output mode is terminated with the BREAK function. The END statement is entered after returning to program mode (using /INFORM-PROGRAM).
- In controlled mode, END can be entered instead of an output control character just like any other output statement.

SM2 acknowledges program termination with the message

SM2 TERMINATED NORMALLY.



Monitored data acquisition is never terminated by the END statement. Termination of the SM2 subsystem is only possible by explicitly issuing the command
/STOP-SUBSYSTEM SUBSYSTEM-NAME=SM2.

5.2 Overview of SM2 operation

After the program load message, SM2 prompts you to enter statements.

Each caller (privileged or nonprivileged user, see [section “Users” on page 24](#)) is initially taken to the analysis subinterval of SM2. Here the user can do the following:

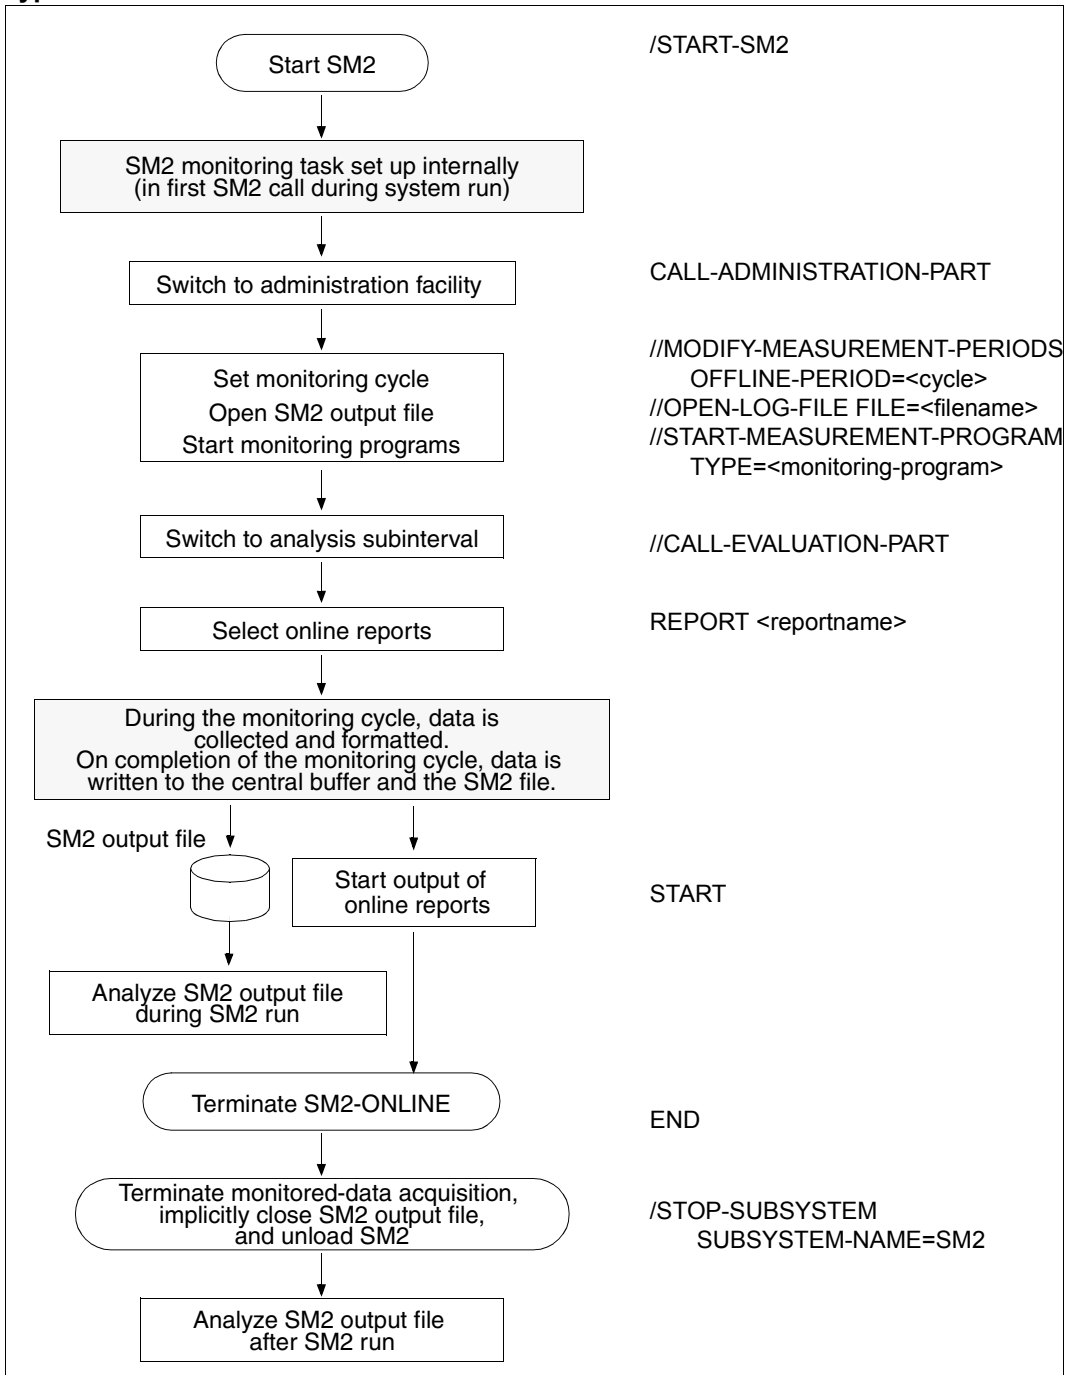
- select monitored data for output on-screen
- control screen output
- output information on the monitoring process

In addition, the privileged user (SM2 call with the system privilege SW-MONITOR-ADMINISTRATION) also has the option of using the CALL-ADMINISTRATION-PART statement to change to the administration facility. Here the SM2 administrator can perform the following administration tasks:

- define monitoring parameters
- activate/deactivate optional monitoring runs
- set up the SM2 output file
- assign privileges for nonprivileged users
- output information on the monitoring process
- specify a host for SM2 statements and reports (only applies for privileged users themselves)

In order to be able to issue the statements used to control real-time monitoring, the privileged user must switch into the analysis subinterval (using the CALL-EVALUATION-PART statement). In doing so, the privileged user defines his/her status as an SM2 administrator.

Typical SM2 run



5.3 Statements for SM2 administrators

These statements are used by SM2 administrators to control data acquisition by SM2. The statements for assigning authorizations to nonprivileged users and for admitting other SM2 administrators are only available to the SM2 primary administrator

Statements for switching between the analysis subinterval and the administration facility

Statement	Function
CALL-ADMINISTRATION-PART	Switch from the analysis subinterval to the administration facility
CALL-EVALUATION-PART	Switch from the administration facility to the analysis subinterval

Statements for defining monitoring programs

Some monitoring programs require you to specify which objects are to be monitored and with what parameters. The ADD, REMOVE, SET and MODIFY statements are provided for this purpose.

The ADD statements each define an object to be monitored by a particular monitoring program. The REMOVE statements can be used to exclude a set of monitored objects, defined with the aid of ADD statements, from the monitoring process. The SET statements define a list of objects to be monitored or the monitoring parameters. The MODIFY statements can be used to modify the default values. The objects to be monitored or the monitoring parameters defined using the statements listed above can be output using the SHOW-DEFINED-PARAMETERS statement. Monitoring of the defined objects does not begin until the monitoring programs are started (using the START-MEASUREMENT-PROGRAM statement) or restarted (using CHANGE-MEASUREMENT-PROGRAM). The monitoring program definitions become active and can be output using the SHOW-ACTIVE-PARAMETERS statement.

Statement	Function
ADD-BCAM-CONNECTION-SET	Add a connection set for BCAM-CONNECTION monitoring
ADD-CONNECTION-SET	Add a connection set (applications) for RESPONSETIME monitoring
ADD-COSMOS-EVENT	Define events for monitored-data acquisition
ADD-FILE	Define a file to be monitored
ADD-ISAM-FILE	Define an ISAM pool (data space) to be monitored
ADD-ISAM-POOL	Define an ISAM pool to be monitored

Statement	Function
ADD-OPENFT-INSTANCE	Define an openFT instance to be monitored
MODIFY-COSMOS-PARAMETERS	Modify the parameters for COSMOS monitoring
MODIFY-RESPONSETIME- PARAMETERS	Modify the parameters for RESPONSETIME monitoring
REMOVE-BCAM-CONNECTION-SET	Exclude a connection set from BCAM-CONNECTION monitoring
REMOVE-CONNECTION-SET	Exclude a connection set from RESPONSETIME monitoring
REMOVE-COSMOS-EVENT	Remove events from the COSMOS monitoring parameters
REMOVE-FILE	Remove file(s) from the FILE monitoring parameters
REMOVE-ISAM-FILE	Remove ISAM pool (data space) from the ISAM monitoring parameters
REMOVE-ISAM-POOL	Remove ISAM pool(s) from the ISAM monitoring parameters
REMOVE-OPENFT-INSTANCE	Remove an openFT instance from the OPENFT monitoring parameters
SET-BCAM-CONNECTION-PARAMETERS	Define BCAM-CONNECTION monitoring parameters
SET-CHANNEL-IO-PARAMETERS	Define CHANNEL-IO monitoring parameters
SET-COSMOS-PARAMETERS	Define COSMOS monitoring parameters
SET-DISK-FILE-PARAMETERS	Define DISK-FILE monitoring parameters
SET-PERIODIC-TASK-PARAMETERS	Define PERIODIC-TASK monitoring parameters
SET-RESPONSETIME-PARAMETERS	Define RESPONSETIME monitoring parameters
SET-SAMPLING-DEVICE-PARAMETERS	Specify recording of service times for disk devices
SET-SERVICETIME-PARAMETERS	Define SERVICETIME monitoring parameters
SET-STORAGE-SYSTEM-PARAMETERS	Define the scope of additional measured data for the STORAGE-SYSTEM monitoring program
SET-SYSTEM-PARAMETERS	Define SYSTEM monitoring parameters
SET-TASK-PARAMETERS	Define TASK monitoring parameters

Statements for preparing, starting, and stopping monitoring programs

Once you have defined the objects to be monitored and the monitoring parameters, the monitoring program run can be started.

Certain preparations are recommended for the COSMOS monitoring program to ensure that it starts quickly and without errors. After defining new monitored objects/monitoring parameters, the monitoring program run must be stopped and then started again in order to actually add the newly defined objects and parameters to the monitoring program. The INITIATE, START, CHANGE and STOP statements are used for this purpose.

Statement	Function
INITIATE-COSMOS	Prepare the COSMOS monitoring program
START-MEASUREMENT-PROGRAM	Start the monitoring program run
CHANGE-MEASUREMENT-PROGRAM	Stop the current monitoring program run and restart it with new monitored objects/monitoring parameters
STOP-MEASUREMENT-PROGRAM	Terminate the monitoring program run

The START statement is used to start the monitoring programs. The CHANGE statement is required following the redefinition of monitored objects/monitoring parameters of the monitoring program run in order to add the new objects/parameters to the monitoring program. The monitoring program run is stopped and then started with the new monitored objects/monitoring parameters. The STOP statement terminates the specified monitoring programs.

The START, CHANGE and STOP statements merely interrupt the current monitoring cycle.

Statements for opening and closing the SM2 output file and controlling monitored-data acquisition

In addition to the output of reports on the screen, monitored data can be continuously written to file. Creating the SM2 output file (OPEN-LOG-FILE) and closing it (CLOSE-LOG-FILE) is a privilege reserved for SM2 administrators. The OPEN-LOG-FILE statement can be used to modify the predefined SM2 file attributes.

The MODIFY-MEASUREMENT-PERIODS statement can be used to modify the monitoring cycles for background monitoring (collecting and writing monitored data to the SM2 output file), and to modify the screen output and the sampling cycle.

Statement	Function
OPEN-LOG-FILE	Open the SM2 output file
CLOSE-LOG-FILE	Close the SM2 output file
MODIFY-MEASUREMENT-PERIODS	Modify the SM2 monitoring cycle

Statements for assigning privileges

After the subsystem is started, monitoring of tasks, files and ISAM pools by the nonprivileged user is not permitted. But by using the MODIFY-USER-ADMISSION statement, the primary SM2 administrator can grant or withdraw permission for all BS2000 users to monitor tasks under their own ID, files and ISAM pools. The MODIFY-ADMINISTRATOR-ADMISSION statement allows the primary SM2 administrator to admit additional (secondary) administrators.

Statement	Function
MODIFY-USER-ADMISSION	Define authorizations for nonprivileged users
MODIFY-ADMINISTRATOR-ADMISSION	Admit additional SM2 administrators

Statements for outputting information

The SHOW statements provide information on the monitoring status, the objects, and the monitoring parameters for selected monitoring programs, as well as on the status of the SM2 system tasks.

Statement	Function
SHOW-MEASUREMENT-STATUS	Output the monitoring status
SHOW-DEFINED-PARAMETERS	Output defined monitoring parameters
SHOW-ACTIVE-PARAMETERS	Output active monitoring parameters
SHOW-USER-MEASURED-OBJECTS	Output monitored objects and associated users
SHOW-SM2-STATUS	Output the status of SM2 system tasks
SHOW-SELECTED-HOSTS	Output selected hosts

Statement for specifying hosts for administration statements and screen outputs

This statement is used to specify the hosts to which the SM2 administration statements apply and for which the screens are to be output.

Statement	Function
SELECT-HOSTS	Define hosts for SM2 statements and screen output

Statement for terminating the SM2 program

The END statement is used to terminate the SM2 session. However, monitoring operations which have already been initiated are not affected, i.e. they continue to run.

Statement	Function
END	Terminate the SM2 run

The following sections describe the statements for SM2 administrators in alphabetical order.



In all statements which allow you to use wildcard syntax, only the character * is supported as the last character. All other constructs possible in SDF are rejected with a error message.

The internal program name for syntax validation in SM2 statements in EDT is SM2190.

ADD-BCAM-CONNECTION-SET

Add connection set for BCAM-CONNECTION monitoring

This statement is used to define a connection set which is to be added to the BCAM-CONNECTION monitoring program definition.

The following features must be specified:

- the type of connection with local and partner (connection) names
- the names of the local and partner systems
- an indication of whether the connections are local or remote

It is possible to define up to 32 connection sets.

This statement must be preceded by a SET-BCAM-CONNECTION-PARAMETERS statement.

Format

ADD-BCAM-CONNECTION-SET

```

SET-NAME = <alphanum-name 1..16>
,CONNECTION-SELECTION = *BY-NEA-NAME(...) / *BY-PORT-NUMBER(...)
  *BY-NEA-NAME(...)
    |
    | CONNECTION-NAME = list-poss(16): *SPECIFIED(...)
    |   *SPECIFIED(...)
    |     LOCAL-APPLICATION = *ANY / <alphanum-name 1..8 with wild>
    |     ,PARTNER-APPLICATION = *ANY / <alphanum-name 1..8 with wild>
  *BY-PORT-NUMBER(...)
    |
    | PORT-NUMBER = list-poss(16): *SPECIFIED(...)
    |   *SPECIFIED(...)
    |     LOCAL-PORT-NUMBER = *ANY / <integer 1..65535>
    |     ,PARTNER-PORT-NUMBER = *ANY / <integer 1..65535>
,CONNECTION-TYPE = *REMOTE / *LOCAL / *BOTH
,HOST-SELECTION = *ANY / *SPECIFIED(...)
  *SPECIFIED(...)
    |
    | LOCAL-HOST-NAME = *LOCAL / <alphanum-name 1..8 with-wild>
    | ,PARTNER-HOST-NAME = *ANY / <alphanum-name 1..8 with-wild>

```


Operands**SET-NAME = <alphanum-name 1..16>**

Defines a name for the selected connection set.

CONNECTION-SELECTION =

Defines the type of connection set to be monitored.

CONNECTION-SELECTION = *BY-NEA-NAME(...)

Defines a selection of connections in accordance with the application name.

CONNECTION-NAME = list-poss(16): *SPECIFIED(...)**LOCAL-APPLICATION = *ANY / <alphanum-name 1..8 with wild>**

*ANY: No special local application is selected.

<alphanum-name 1..8>: Specifies the name of the selected local application. An asterisk * is also permitted as the last character in a wildcard specification.

,PARTNER-APPLICATION = *ANY / <alphanum-name 1..8 with wild>

*ANY: No special partner application is selected.

<alphanum-name 1..8>: Specifies the name of the selected partner application. An asterisk * is also permitted as the last character in a wildcard specification.

CONNECTION-SELECTION = *BY-PORT-NUMBER(...)

Defines a selection of connections in accordance with the port number.

PORT-NUMBER = list-poss(16): *SPECIFIED(...)**LOCAL-PORT-NUMBER = *ANY / <integer 1..65535>**

*ANY: No special local transport address is selected.

<integer 1..65535>: Local transport address.

PARTNER-PORT-NUMBER = *ANY / <integer 1..65535>

*ANY: No special partner transport address is selected.

<integer 1..65535>: Partner transport address.

CONNECTION-TYPE =

Specifies the type of connection to be taken into consideration for the connection set.

CONNECTION-TYPE = *REMOTE

Only remote connections are taken into consideration.

CONNECTION-TYPE = *LOCAL

Only local connections are taken into consideration.

CONNECTION-TYPE = *BOTH

Both local and remote connections are taken into consideration.

HOST-SELECTION =

Defines the names of the hosts between which the connections to be monitored should exist.

HOST-SELECTION = *ANY

No special host name is selected.

HOST-SELECTION = *SPECIFIED(...)**LOCAL-HOST-NAME = *LOCAL / <alphanum-name 1..8 with-wild>**

*LOCAL: The local host name should be used.

<alphanum-name 1..8>: Specifies the name of the local host. With some connections, the local processor name may differ from the host name. An asterisk * is also permitted as the last character in a wildcard specification.

PARTNER-HOST-NAME = *ANY / <alphanum-name 1..8 with-wild>

*ANY: No special partner host should be assumed.

<alphanum-name 1..8>: Specifies the name of the partner host. An asterisk * is also permitted as the last character in a wildcard specification.

ADD-CONNECTION-SET

Add connection set for RESPONSETIME monitoring

This statement is used to define a connection set which is to be added to the RESPONSETIME monitoring program definition. A connection set consists of one or more (as many as five) connection groups. The connection set can be either positively defined (i.e. all groups specified in this connection set are monitored) or negatively defined (i.e. all groups except those specified in the connection set are monitored).

For each connection set, a freely selectable name must be assigned; the name is the one under which this set can be addressed in the REMOVE-CONNECTION-SET statement and also addressed by SM2R1 for analysis purposes.

If a connection set is defined, it is also possible to specify whether remote, local or both types of connections are to be taken into account for the set. SM2 then passes the correct parameter value to BCAM-SM2. This value then applies to this connection set only.

This statement cannot be entered until after the SET-RESPONSETIME-PARAMETERS statement has been entered. The SET-RESPONSETIME-PARAMETERS statement automatically defines a connection set with the name *GLOBAL in which all connections are monitored (*CONNECTION (*ALL)). The value CONNECTION-TYPE=*REMOTE applies for this connection set and cannot be changed.

A total of 16 connection sets (*GLOBAL inclusions) may be defined with a total of 16 different connection groups ((*ALL) inclusion). If this number is exceeded, the statement will be rejected and a message to this effect issued. If an attempt is made to declare more than one set under one set name, the statement will be rejected and a message will be issued.

Format

ADD-CONNECTION-SET

```

SET-NAME = <alphanum-name 1..16>
,SET-DEFINITION = *BY-CONNECTION / *EXCEPT-CONNECTION
,CONNECTION-SET = list-poss(5): *CONNECTION(...)
    *CONNECTION(...)
        | APPLICATION = *ALL / <alphanum-name 1..8 with-wild>
,CONNECTION-TYPE = *REMOTE / *LOCAL / *BOTH

```

Operands

SET-NAME = <alphanum-name 1..16>

Defines a name for the selected connection set.

SET-DEFINITION =

Defines the connection set to be monitored.

SET-DEFINITION = *BY-CONNECTION

All connection groups defined in the connection set are monitored.

SET-DEFINITION = *EXCEPT-CONNECTION

All connection groups except those defined in the connection set are monitored.

CONNECTION-SET = list-poss(5): *CONNECTION(...)

Defines a selection of up to five connection groups.

APPLICATION = *ALL / <alphanum-name 1..8 with-wild>

*ALL: No special application is selected.

<alphanum-name 1..8>: Specifies the name of the selected application. An asterisk * is also permitted as the last character in a wildcard specification.

CONNECTION-TYPE =

Specifies the type of connection to be taken into consideration for the connection set.

CONNECTION-TYPE = *REMOTE

Only remote connections are taken into consideration.

CONNECTION-TYPE = *LOCAL

Only local connections are taken into consideration.

CONNECTION-TYPE = *BOTH

Both local and remote connections are taken into consideration.



No application name is defined for applications that use the SOCKETS(BS2000) or ICMX(BS2000) interface. Connections of these applications cannot be selected in SM2 using the ADD-CONNECTION-SET statement. They are contained in the *GLOBAL connection set.

The BCAM-CONNECTION supports the selection of the partner computer, the local application and the partner application.

ADD-COSMOS-EVENT

Define events for monitored-data acquisition

This statement is used to define events which are to be added to the COSMOS monitoring program definition.

Format

ADD-COSMOS-EVENT

```
EVENT-NAME = *STANDARD-EVENTS / list-poss(71): *STANDARD-EVENTS / *ACF / *BCAM / *BCPT /
  *BLS / *BOUR / *CHTM / *CMD / *CMS / *DAB / *DCAM / *DLM / *DSM / *EIA / *EIA2 / *EIA3 /
  *FITC / *GSAC / *HAL / *IDLE / *INTR / *IONQ / *ISEV / *ISPL / *KAI / *LOCK / *MSG /
  *NSM / *PAGE / *PAM / *PCCC / *PCTC / *PDEA / *PEND / *PIO / *PMIO / *PRGS / *PRGT /
  *PRTY / *RELM / *REQM / *RSCS / *RSCT / *SDV / *SLOT / *SNAP / *STD1 / *STDI /
  *STDN / *SVC / *SWSR / *TGMA / *TGMP / *TGMT / *TIC / *TINF / *TLM / *TLT / *TSKI /
  *TSVC / *UTM / *VMCH / *VMH / *VMI / *VMLK / *VMPD / *VMPR / *VMS / *VM2 /
  *WSCT / *XEIA
```

Operands

EVENT-NAME =

Specifies the name of the events to be recorded.

EVENT-NAME = *STANDARD-EVENTS

The following events are to be recorded by default: ACF, BLS, BOUR, CHTM, CMS, DAB, EIA2, EIA3, FITC, IDLE, INTR, IONQ, PAGE, PAM, PCCC, PCTC, PEND, PMIO, PRGS, PRGT, RELM, REQM, SDV, SVC, TSKI, TSVC, WSCT.

EVENT-NAME = list-poss(71): *STANDARD-EVENTS / ...

The events specified in the list are to be recorded. The mandatory events CREA, DEST, INIT, LGON, MMRC, PTSK and STAT are always recorded.

Behind some events (known as group events) there lie more than one event. The following table provides an overview:

Group event name	Events
EIA	EIA2, EIA3
PEND	PEND, UNPN
DAB	All DAB events
DLM	All DLM events
NSM	All NSM events
VM2	All VM2000 hypervisor events (servers with /390 architecture only)
VMH	VMHS, VMHE
VMI	VMIS, VMIE
VMS	VMSS, VMSE
TIC	TICS, TICE
KAI	KCOL, KRST

Table 6: Group events

ADD-FILE

Define file to be monitored

This statement is used to define a file which is to be added to the FILE monitoring program definition. This monitoring program allows you to define up to 32 files for monitoring. If this number is exceeded, the statement is rejected and a corresponding message is issued.

Format

ADD-FILE
FILE-NAME = <filename 1..54>

Operands

FILE-NAME = <filename 1..54>

Name of the file whose access values are to be monitored. The file name must be specified as a fully qualified name.

ADD-ISAM-FILE

Define an ISAM pool to be monitored

This statement is used to define the pathname of a NK-ISAM file, which is to be added to the ISAM monitoring program definition.

This is necessary when the file is located in a global ISAM pool which was created automatically (i.e. without a preceding CREATE-ISAM-POOL command). Such pools are located in a data space and have no name. Consequently they cannot be addressed with the ADD-ISAM-POOL statement.

If the file is not located in a global ISAM pool, SM2 records no data. If the pool contains further files, SM2 records only accesses to the specified file.

This monitoring program allows you to define up to 16 ISAM files for monitoring. Independently it is allowed to define up to 16 ISAM pools through their poolname for monitoring.

Format

ADD-ISAM-FILE
FILE-NAME = <filename 1..54>

Operands

FILE-NAME = <filename 1..54>

Name of the NK-ISAM file to be added to the monitoring program definition. The file name must be specified as a fully qualified name.

ADD-ISAM-POOL

Define ISAM pool to be monitored

This statement is used to define an ISAM pool which is to be added to the ISAM monitoring program definition. This statement allows you to define up to 16 ISAM pools for monitoring. If this number is exceeded, the statement is rejected and a message is issued.

Format

ADD-ISAM-POOL

```
POOL-NAME = <alphanum-name 1..8>
,SCOPE = *HOST-SYSTEM / *TASK(...)
    *TASK(...)
        | TSN = <alphanum-name 1..4>
,CAT-ID = *HOME / <catid 1..4>
```

Operands

POOL-NAME = <alphanum-name 1..8>

Specifies the name of the ISAM pool to be monitored.

SCOPE =

Defines the type of ISAM pool.

SCOPE = *HOST-SYSTEM

The pool here is a global ISAM pool.

SCOPE = *TASK(...)

The pool here is a task-local ISAM pool.

TSN = <alphanum-name 1..4>

Specifies the TSN of the task under which the task-local ISAM pool was set up.

CAT-ID =

Specifies the catalog ID of the pubset to which the ISAM pool has been assigned.

CAT-ID = *HOME

The ID here is the catalog ID of the home pubset.

CAT-ID = <cat-id 1..4>

Catalog ID of the pubset to which the ISAM pool has been assigned.

ADD-OPENFT-INSTANCE

Add openFT instance for monitoring

This statement is used to define an openFT instance which is added to the monitoring program definition of OPENFT. A total of up to 16 openFT instances which are to be monitored can be defined for this monitoring program. If this number is exceeded, the statement is rejected and a corresponding message is issued.

Format

ADD-OPENFT-INSTANCE
INSTANCE-NAME = <alphanum-name 1..8>

Operands

INSTANCE-NAME = <alphanum-name 1..8>

Specifies the name of the openFT instance.

CALL-ADMINISTRATION-PART

Switch from evaluation part to administration facility

This statement is used to switch from the evaluation part to the administration facility, which provides an SDF statement interface.

To switch back to the analysis subinterval, use the CALL-EVALUATION-PART command.

Format

CALL-ADMINISTRATION-PART

CALL-EVALUATION-PART

Switch from administration facility to evaluation part

This statement is used to switch from the administration facility to the evaluation part, i.e. after this statement is entered all statements required for online analysis can be entered. To switch back to SDF statement mode for administration purposes, use the CALL-ADMINISTRATION-PART command.

Format

CALL-EVALUATION-PART

CHANGE-MEASUREMENT-PROGRAM

Stop current monitoring program run and restart it with new monitored objects

This statement is used to stop the current monitoring program run and restart it using newly defined monitored objects/monitoring parameters. The current monitoring cycle is interrupted.

Format

CHANGE-MEASUREMENT-PROGRAM
TYPE = list-poss(14): *BCAM-CONNECTION / *CHANNEL-IO / *COSMOS / *DISK-FILE / *FILE / *ISAM / *PERIODIC-TASK / *RESPONSETIME / *SAMPLING-DEVICE / *SERVICETIME / *STORAGE-SYSTEM / *SYSTEM / *TASK

Operands

TYPE =

Specifies the name of the monitoring program to be restarted using the modified monitored objects/monitoring parameters, see [chapter “SM2 monitoring programs” on page 49](#).



Even if errors are detected when processing the CHANGE statement before the monitoring programs are stopped, the monitoring programs continue to execute. If, after stopping the program, errors are detected when restarting the monitoring programs, the monitoring programs specified in the statement will still be stopped.

CLOSE-LOG-FILE

Close SM2 output file

This statement is used to close an SM2 output file. The current monitoring cycle is interrupted.

Format

CLOSE-LOG-FILE



When the SM2 output file is closed, all monitoring programs are stopped and the monitored data is written exclusively to the SM2 output file (see [table “Overview of the SM2 monitoring programs” on page 50](#)).

END

Terminate SM2 run

This statement is used to terminate SM2. All current measurements are continued using the currently defined monitoring parameters. This is also the case if no further user is working with SM2. The only way to terminate the SM2 subsystem along with the privileged system tasks is to issue the command `/STOP-SUBSYSTEM SUBSYSTEM-NAME=SM2`.

Format

END

INITIATE-COSMOS

Prepare COSMOS monitoring program

The COSMOS monitoring program is fully initiated (prepared), but no events are recorded yet – not even mandatory events. This statement is used to ensure to a large extent that the COSMOS monitoring program subsequently starts without errors and quickly (without mount messages).

Format

INITIATE-COSMOS

MODIFY-ADMINISTRATOR-ADMISSION

Admit additional SM2 administrators

This statement is used by the SM2 primary administrator to admit additional secondary administrators.

Format

MODIFY-ADMINISTRATOR-ADMISSION
CONCURRENCY = * <u>NO</u> / *YES

Operands

CONCURRENCY =

Specifies whether additional SM2 administrators are to be admitted. After the start of the SM2 subsystem no additional SM2 administrators are permitted.

CONCURRENCY = *NO

No additional SM2 administrators are to be admitted.

CONCURRENCY = *YES

Additional administrators are admitted. These administrators are allowed to execute all SM2 statements with the exception of MODIFY-ADMINISTRATOR-ADMISSION and MODIFY-USER-ADMISSION.

MODIFY-COSMOS-PARAMETERS

Modify parameters for COSMOS monitoring

This statement is used to modify the COSMOS monitoring program definition. It is useful when the SHOW-DEFINED-PARAMETERS statement finds parameters that have to be modified.

Format

MODIFY-COSMOS-PARAMETERS

```
TITLE = *UNCHANGED / <c-string 1..80>
,BUFFER-SIZE = *UNCHANGED / <integer 1..40>
,NUMBER-OF-BUFFERS = *UNCHANGED / <integer 2..512>
,ADDITIONAL-INFO = *UNCHANGED / *NONE / list-poss(2): *CONFIGURATION / *VM2000(...)
    *VM2000(...)
        | FILE-NAME = <filename 1..54 without-gen-vers>
,OUTPUT = *UNCHANGED / *DISK(...) / *WRAP-AROUND(...) / *TAPE(...) / *STREAM-TAPE(...)
    *DISK(...)
        | FILE-NAME = list-poss(16): <filename 1..54 without-gen-vers>
    *WRAP-AROUND(...)
        | FILE-NAME = <filename 1..54 without-gen-vers>
    *TAPE(...)
        | FILE-NAME = list-poss(16): <filename 1..54 without-gen-vers>
    *STREAM-TAPE(...)
        | FILE-NAME = list-poss(16): <filename 1..54 without-gen-vers>
```

(part 1 of 2)

```
,TASK-SELECTION = *UNCHANGED / *ALL / *SPECIFIED(...)
    *SPECIFIED(...)
        ,JOB-NAME = *NOT-SPECIFIED / list-poss(8): <alphanum-name 1..8 with-wild>
        ,CATEGORY = *NOT-SPECIFIED / list-poss(8): <alphanum-name 1..7 with-wild>
        ,USER-ID = *NOT-SPECIFIED / list-poss(8): <alphanum-name 1..8 with-wild>
        ,TSN = *NOT-SPECIFIED / list-poss(16): <alphanum-name 1..4 with-wild>
        ,TYPE = *NOT-SPECIFIED / list-poss(4): *SYSTEM / *DIALOG / *BATCH / *TP
,EVENT-SELECTION = *UNCHANGED / *ALL-BY-ADD-COSMOS-EVENT / *SPECIFIED(...)
    *SPECIFIED(...)
        EIA-INTERRUPT-CLASS = *UNCHANGED / *ANY / list-poss(5): *SVC / *PROGRAM /
            *MACHINE-CHECK / *IO / *EXTERNAL
        ,EIA-SVC-NUMBER = *UNCHANGED / *ANY / list-poss(8): <integer 1..255>
        ,IO-DEVICE = *UNCHANGED / *ANY / list-poss(8): <alphanum-name 2..4>
        ,DAB-CACHE-ID = *UNCHANGED / *ANY / list-poss(8): <alphanum-name 1..32>
        ,MEMORY-CLASS = *UNCHANGED / *ANY / list-poss(4): *3 / *4 / *5 / *6
        ,SLOT-MEMORY-CLASS = *UNCHANGED / *ANY / list-poss(4): *3 / *4 / *5 / *6
        ,PEND-CODE = *UNCHANGED / *ANY / list-poss(16): <integer 1..22>
        ,LOCK-ID = *UNCHANGED / *ANY / list-poss(4): <alphanum-name 1..2>
        ,TLT-DESCRIPTOR = *UNCHANGED / *ANY / list-poss(8): <alphanum-name 1..3>
        ,TSKI-SWITCH = *UNCHANGED / *ANY / *TASK
        ,TSVC-SVC-NUMBER = *UNCHANGED / *ANY / list-poss(8): <integer 1..255>
        ,CPU-NUMBER = *UNCHANGED / *ANY / list-poss(32): <integer 0..31>
,UNLOAD = *UNCHANGED / *AT-MEASUREMENT-PROGRAM-STOP / *AT-SM2-STOP
,MEASUREMENT-TIME = *UNCHANGED / *NOT-SPECIFIED / <integer 1..60>
```

(part 2 of 2)

Operands

TITLE =

Specifies the title of the COSMOS monitoring process.

TITLE = *UNCHANGED

The currently specified title is not changed.

TITLE = <c-string 1..80>

The specified title is assigned.

BUFFER-SIZE =

Specifies the number of 4K pages per buffer.

BUFFER-SIZE = *UNCHANGED

The currently specified number of pages per buffer is not changed.

BUFFER-SIZE = <integer 1..40>

Buffers of the specified size (number of pages) are created.

NUMBER-OF-BUFFERS =

Specifies the number of buffers.

NUMBER-OF-BUFFERS = *UNCHANGED

The currently specified number of buffers is not changed.

NUMBER-OF-BUFFERS = <integer 2..512>

COSMOS creates the specified number of buffers for writing events.

ADDITIONAL-INFO =

Specifies additional data to be recorded.

ADDITIONAL-INFO = *UNCHANGED

The existing data specified for ADDITIONAL-INFO is not changed.

ADDITIONAL-INFO = *NONE

No additional data is recorded.

ADDITIONAL-INFO = *CONFIGURATION

The configuration at the start of monitoring is also recorded.

ADDITIONAL-INFO = *VM2000(...)

The VM2000 events are also recorded. For this purpose, they must be explicitly opened (ADD-COSMOS-EVENT statement).

FILE-NAME = <filename 1..54 without-gen-vers>

Specifies the name of the file to which the VM2000 events are written.

OUTPUT =

Specifies how the COSMOS output files are written.

OUTPUT = *UNCHANGED

The currently specified value for OUTPUT is not changed.

OUTPUT = *DISK(...)

The COSMOS monitoring data is written sequentially to the files on hard disk; i.e. there is a write task for each output file. If a monitoring data buffer is full, a write task is activated that writes the buffer to the corresponding file.

FILE-NAME = list-poss(16): <filename 1..54 without-gen-vers>

Specifies the file(s) to which the COSMOS monitored data is written.

OUTPUT = *WRAP-AROUND(...)

COSMOS monitored data is written to a file on disk. If the disk space allocated (PRIMARY-ALLOCATION) is insufficient, the file is overwritten from the top. The secondary allocation (SECONDARY-ALLOCATION) must be set to zero.

FILE-NAME = <filename 1..54 without-gen-vers>

Specifies the file to which COSMOS monitored data is written.

OUTPUT = *TAPE(...)

The COSMOS monitored data is written sequentially to the files on tape (works as described for OUTPUT=*DISK.)

FILE-NAME = list-poss(16): <filename 1..54 without-gen-vers>

Specifies the file(s) to which COSMOS monitored data is written.

OUTPUT = *STREAM-TAPE(...)

The COSMOS output files are written in streaming mode. This reduces the chances of events not being recorded, because the files are written to the tape continuously without any repositioning of the tape.

FILE-NAME = list-poss(16): <filename 1..54 without-gen-vers>

Specifies the file(s) to which COSMOS monitored data is written.

TASK-SELECTION =

Specifies the tasks to be monitored.

TASK-SELECTION = *UNCHANGED

The current selection of tasks to be monitored is not changed.

TASK-SELECTION = *ALL

All tasks are monitored.

TASK-SELECTION = *SPECIFIED(...)

Specifies the tasks to be monitored by means of their TSN, user ID, job name, category or type.

JOB-NAME =

The tasks are selected by means of their job name. An asterisk * is also permitted as the last character in a wildcard specification.

JOB-NAME = *NOT-SPECIFIED

No tasks are selected by means of their job name.

JOB-NAME = list-poss(8): <alphanum-name 1..8 with-wild>

Specifies the job names of the tasks to be monitored.

CATEGORY =

The tasks are selected by means of their category. An asterisk * is also permitted as the last character in a wildcard specification.

CATEGORY = *NOT-SPECIFIED

No tasks are selected by means of their category.

CATEGORY = list-poss(8): <alphanum-name 1..7 with-wild>

Specifies the categories of the tasks to be monitored.

USER-ID =

The tasks are selected by means of their user ID.

USER-ID = *NOT-SPECIFIED

No tasks are selected by means of their user ID.

USER-ID = list-poss(8): <alphanum-name 1..8 with-wild>

Specifies the user IDs of the tasks to be monitored. These must be specified without the \$ sign. An asterisk * is also permitted as the last character in a wildcard specification.

TSN =

The tasks are selected by means of their TSN.

TSN = *NOT-SPECIFIED

No tasks are selected by means of their TSN.

TSN = list-poss(16): <alphanum-name 1..4 with-wild>

Specifies the TSNs of the tasks to be monitored. An asterisk * is also permitted as the last character in a wildcard specification.

TYPE =

The tasks are selected by means of their type.

TYPE = *NOT-SPECIFIED

No tasks are selected by means of their type.

TYPE = *SYSTEM

All system tasks are to be monitored.

TYPE = *DIALOG

All interactive tasks are to be monitored.

TYPE = *BATCH

All batch tasks are to be monitored.

TYPE = *TP

All TP tasks are to be monitored.

EVENT-SELECTION =

Specifies the events to be recorded. Their recording is linked to specific conditions.

EVENT-SELECTION = *UNCHANGED

The existing values for EVENT-SELECTION are not changed.

EVENT-SELECTION = *ALL-BY-ADD-COSMOS-EVENT

All open events are recorded.

EVENT-SELECTION = *SPECIFIED(...)

Only those events that satisfy the specified conditions are recorded.

EIA-INTERRUPT-CLASS =

Controls the recording of EIA events on the basis of their interrupt class (IC).

EIA-INTERRUPT-CLASS = *UNCHANGED

The currently specified ICs are not changed.

EIA-INTERRUPT-CLASS = *ANY

The recording of EIA events is independent of their IC.

EIA-INTERRUPT-CLASS = list-poss(5): *SVC / *PROGRAM / *MACHINE-CHECK / *IO / *EXTERNAL

Only those EIA events with the specified IC are recorded.

EIA-SVC-NUMBER =

Controls the recording of EIA events on the basis of their SVC.

EIA-SVC-NUMBER = *UNCHANGED

The currently specified SVCs are not changed.

EIA-SVC-NUMBER = *ANY

The recording of EIA events is independent of their SVC.

EIA-SVC-NUMBER = list-poss(8): <integer 1..255>

Only those EIA events with the specified SVC(s) are recorded.

IO-DEVICE =

Controls the recording of SDV, CHTM, IONQ and PMIO events on the basis of their mnemonic device name.

IO-DEVICE = *UNCHANGED

The currently specified mnemonic device names are not changed.

IO-DEVICE = *ANY

The recording of SDV, CHTM, IONQ and PMIO events is independent of their mnemonic device name.

IO-DEVICE = list-poss(8): <alphanum-name 2..4>

Only those SDV, CHTM, IONQ and PMIO events with the specified mnemonic device names are recorded.

DAB-CACHE-ID =

Controls the recording of DAB events on the basis of their DAB cache IDs.

DAB-CACHE-ID = *UNCHANGED

The currently specified DAB cache IDs are not changed.

DAB-CACHE-ID = *ANY

The recording of DAB events is independent of their DAB cache IDs.

DAB-CACHE-ID = list-poss(8): <alphanum-name 1..32>

Only those DAB events with the specified DAB cache IDs are recorded.

MEMORY-CLASS =

Controls the recording of RELM and REQM events on the basis of their memory class.

MEMORY-CLASS = *UNCHANGED

The currently specified memory classes for RELM and REQM events are not changed.

MEMORY-CLASS = *ANY

The recording of RELM and REQM events is independent of their memory class.

MEMORY-CLASS = list-poss(4): *3 / *4 / *5 / *6

Only those RELM and REQM events with the specified memory classes are recorded.

SLOT-MEMORY-CLASS =

Controls the recording of SLOT events on the basis of their memory class.

SLOT-MEMORY-CLASS = *UNCHANGED

The currently specified memory classes for SLOT events are not changed.

SLOT-MEMORY-CLASS = *ANY

The recording of SLOT events is independent of their memory class.

SLOT-MEMORY-CLASS = list-poss(4): *3 / *4 / *5 / *6

Only those SLOT events with the specified memory classes are recorded.

PEND-CODE =

Controls the recording of PEND events on the basis of their pend code.

PEND-CODE = *UNCHANGED

The currently specified pend codes are not changed.

PEND-CODE = *ANY

The recording of PEND events is independent of their pend code.

PEND-CODE = list-poss(16): <integer 1..22>

Only those PEND events with the specified pend code are recorded.

LOCK-ID =

Controls the recording of LOCK events on the basis of their lock ID.

LOCK-ID = *UNCHANGED

The currently specified lock IDs are not changed.

LOCK-ID = *ANY

The recording of LOCK events is independent of their lock ID.

LOCK-ID = list-poss(4): <alphanum-name 1..2>

Only those LOCK events with the specified lock IDs are recorded.

TLT-DESCRIPTOR =

Controls the recording of TLT on the basis of their TLT descriptor (TLT = Task Location Table).

TLT-DESCRIPTOR = *UNCHANGED

The currently specified values for the TLT descriptor are not changed.

TLT-DESCRIPTOR = *ANY

The recording of TLT events is independent of their TLT descriptor.

TLT-DESCRIPTOR = list-poss(8): <alphanum-name 1..3>

Only those TLT events with the specified TLT descriptors are recorded.

TSKI-SWITCH =

Controls the recording of TSKI events on the basis of their TIC (Task in Control).

TSKI-SWITCH = *UNCHANGED

The currently specified value for TSKI-SWITCH is not changed.

TSKI-SWITCH = *ANY

The recording of TSKI events is independent of their TIC.

TSKI-SWITCH = *TASK

Only the first TSKI event of a TIC is recorded.

TSVC-SVC-NUMBER =

Controls the recording of TSVC events on the basis of their SVC numbers.

TSVC-SVC-NUMBER = *UNCHANGED

The currently specified SVC numbers are not changed.

TSVC-SVC-NUMBER = *ANY

The recording of TSVC events is independent of their SVC number.

TSVC-SVC-NUMBER = list-poss(8): <integer 1..255>

The TSVC events are recorded for the specified SVC numbers.

CPU-NUMBER =

Controls the recording of events according to the CPU number.

CPU-NUMBER = *UNCHANGED

The currently specified CPU numbers are not changed.

CPU-NUMBER = *ANY

The recording of events is independent of their CPU number.

CPU-NUMBER = list-poss(32): <integer 0..31>

Only the events for the specified CPU numbers are recorded.

UNLOAD =

Specifies the time at which the COSMOS subsystem is to be unloaded.

UNLOAD = *UNCHANGED

The currently specified value is not changed.

UNLOAD = *AT-MEASUREMENT-PROGRAM-STOP

The COSMOS subsystem is unloaded at the end of COSMOS monitoring.

UNLOAD = *AT-SM2-STOP

The COSMOS subsystem is unloaded at the end of SM2 monitoring.

MEASUREMENT-TIME =

Specifies the duration of the monitoring process.

MEASUREMENT-TIME = *UNCHANGED

The currently specified duration is not changed.

MEASUREMENT-TIME = *NOT-SPECIFIED

Monitoring is terminated by the user with the STOP-MEASUREMENT-PROGRAM
TYPE = *COSMOS statement.

MEASUREMENT-TIME = <integer 1..60>

Specifies the time in minutes after which monitoring is to be terminated automatically.

MODIFY-MEASUREMENT-PERIODS

Modify SM2 monitoring cycle

This statement is used to modify the SM2 time periods which control the cyclical acquisition of SM2 monitored data.

Format

MODIFY-MEASUREMENT-PERIODS

```

OFFLINE-PERIOD = *UNCHANGED / <integer 10..3600>
,ONLINE-PERIOD = *UNCHANGED / *SAME-AS-OFFLINE / <integer 10..3600>
,SAMPLING-PERIOD = *UNCHANGED / <integer 200..10000>

```

Operands

OFFLINE-PERIOD =

Specifies the SM2 monitoring cycle during which monitored data is collected and written to a file.

OFFLINE-PERIOD = *UNCHANGED

The monitoring cycle defined here remains unchanged. When the SM2 subsystem is started, this value is set to 150 seconds.

OFFLINE-PERIOD = <integer 10..3600>

Defines the SM2 monitoring cycle in seconds.

ONLINE-PERIOD =

Specifies the SM2 monitoring cycle during which the monitored data is collected for screen output and then displayed. The cycle of the background monitoring operations (namely gathering and writing the monitored data to the SM2 output file) is not changed.

ONLINE-PERIOD = *UNCHANGED

The online cycle remains unchanged. When the SM2 subsystem is started, this value is set to *SAME-AS-OFFLINE (i.e. no separate online cycle is defined).

ONLINE-PERIOD = *SAME-AS-OFFLINE

The online cycle is the same as for background monitoring. This setting is the optimum setting for internal management of data. If the online cycle is not required for a longer period of time, this offline setting is selected.

ONLINE-PERIOD = <integer 10..3600>

Defines an online cycle in seconds.



If the same value is selected for OFFLINE-PERIOD and ONLINE-PERIOD, this is not the equivalent of "ONLINE-PERIOD = *SAME-AS-OFFLINE". What it does mean is that internally separate buffers are maintained and provided with data. For reasons relating to the system load, such a parameter setting should be avoided.

SAMPLING-PERIOD =

Changes the sampling cycle during which samples are taken of some of the SM2 monitored variables.

SAMPLING-PERIOD = *UNCHANGED

The sampling cycle remains unchanged. When the subsystem is started, this value is set to 800 milliseconds.

SAMPLING-PERIOD = <integer 200..10000>

Defines a sampling cycle in milliseconds. Values which are entered are rounded down to multiples of 100 (milliseconds).

MODIFY-RESPONSETIME-PARAMETERS

Modify parameters for RESPONSETIME monitoring

This statement is used to modify the RESPONSETIME monitoring program definition.

Format

MODIFY-RESPONSETIME-PARAMETERS

```
SCOPE = *UNCHANGED / list-poss(2): *BUCKET / *CATEGORY
,DEFINITION = *UNCHANGED / *1 / *2
,CONNECTION-NUMBER = *UNCHANGED / <integer 1..8187>
,RESPONSETIME-BUCKETS = *UNCHANGED / *STD-LIMITS / list-poss(5): <integer 1..99999>
,THINKTIME-BUCKETS = *UNCHANGED / *STD-LIMITS / list-poss(5): <integer 1..99999>
,TRANSACTIONTIME-BUCKETS = *UNCHANGED / *STD-LIMITS / list-poss(5): <integer 1..99999>
,WAITTIME-BUCKETS = *UNCHANGED / *STD-LIMITS / list-poss(5): <integer 1..99999>
```

Operands

SCOPE =

Defines whether recording of the response time data is to be bucket-specific or category-specific.

SCOPE = *UNCHANGED

The scope defined by this time will not be changed.

SCOPE = *BUCKET

The response time data is to be recorded based on buckets.

SCOPE = *CATEGORY

The response time data is to be recorded based on categories.

DEFINITION =

Defines the type of response time to be monitored.

DEFINITION = *UNCHANGED

The type of response time defined by this time will not be changed.

DEFINITION = *1

The response time valid is the period between user input and the first relevant output.

DEFINITION = *2

In addition to the response time defined by the value 1, the periods between output following any entry are output individually as response times.

CONNECTION-NUMBER =

Specifies the maximum number of connections to be monitored.

CONNECTION-NUMBER = *UNCHANGED

The number of connections defined at this time for the RESPONSETIME monitoring program remains unchanged.

CONNECTION-NUMBER = <integer 1..8187>

The specified maximum number of connections is recorded.

RESPONSETIME-BUCKETS =

Defines the upper limits of up to five ranges (in units of 100 ms) in which the response times are to be stored by order of magnitude.

RESPONSE-BUCKETS = *UNCHANGED

The currently defined upper limits of the ranges for response times are not changed.

RESPONSE-BUCKETS = *STD-LIMITS

Sets the upper limits of the ranges for responses time to 5, 10, 20, 50 and 100.

RESPONSE-BUCKETS = list-poss(5): <integer 1..99999>

Defines the upper limits of the ranges for response times. These must be specified in ascending order.

THINKTIME-BUCKETS =

Defines the upper limits of up to five ranges (in units of 100ms) in which the think times are to be stored by order of magnitude.

THINKTIME-BUCKETS = *UNCHANGED

The currently defined upper limits of the ranges for think times are not changed.

THINKTIME-BUCKETS = *STD-LIMITS

Sets the upper limits of the ranges for think times to 50, 150, 300, 600 and 1200.

THINKTIME-BUCKETS = list-poss(5): <integer 1..99999>

Defines the upper limits of the ranges for think times. These must be specified in ascending order.

TRANSACTTIME-BUCKETS =

Defines the upper limits of up to five ranges (in units of 100ms) in which the transaction times are to be stored by order of magnitude.

TRANSACTTIME-BUCKETS = *UNCHANGED

The currently defined upper limits of the ranges for transaction times are not changed.

TRANSACTTIME-BUCKETS = *STD-LIMITS

Sets the upper limits of the ranges for transaction times to 5, 10, 20, 50 and 100.

TRANSACTTIME-BUCKETS = list-poss(5): <integer 1..99999>

Defines the upper limits of the ranges for transaction times. These must be specified in ascending order.

WAITTIME-BUCKETS =

Defines the upper limits of up to five ranges (in units of 100ms) in which the wait times in the BCAM pool are to be stored by order of magnitude.

WAITTIME-BUCKETS = *UNCHANGED

The currently defined upper limits of the ranges for wait times in the BCAM pool are not changed.

WAITTIME-BUCKETS = *STD-LIMITS

Sets the upper limits of the ranges for wait times in the BCAM pool to 1, 2, 5, 10 and 20.

WAITTIME-BUCKETS = list-poss(5): <integer 1..99999>

Defines the upper limits of the ranges for wait times in the BCAM pool. These must be specified in ascending order.

MODIFY-USER-ADMISSION

Define authorizations for nonprivileged users

This statement is used by the primary SM2 administrator to specify which users are permitted to execute user-specific monitoring programs.

Format

MODIFY-USER-ADMISSION
<pre> TASK = <u>*UNCHANGED</u> / *ALLOW(...) / *INHIBIT *ALLOW(...) USER-ID = <u>*ALL</u> / list-poss(16): <alphanum-name 1..8> ,FILE = <u>*UNCHANGED</u> / *ALLOW(...) / *INHIBIT *ALLOW(...) USER-ID = <u>*ALL</u> / list-poss(16): <alphanum-name 1..8> ,ISAM = <u>*UNCHANGED</u> / *ALLOW(...) / *INHIBIT *ALLOW(...) USER-ID = <u>*ALL</u> / list-poss(16): <alphanum-name 1..8> </pre>

Operands

TASK =

The authorization is set for the TASK user-specific monitoring program.

TASK = *UNCHANGED

The authorization specified at this time for the TASK user-specific monitoring program remains unchanged. After the SM2 subsystem is started, the TASK user-specific monitoring program is not permitted.

TASK = *ALLOW(...)

The TASK user-specific monitoring program is permitted. The user uses the BS2000 commands /START-TASK-MEASUREMENT and /STOP-TASK-MEASUREMENT to start and stop tasks for monitoring.

USER-ID =

Specifies the users for whom the TASK user-specific monitoring program is permitted.

USER-ID = *ALL

The TASK user-specific monitoring program is permitted for all BS2000 users.

USER-ID = list-poss(16): <alphanum-name 1..8>

The TASK user-specific monitoring program is permitted for the users of the specified user IDs. The user IDs must be specified without \$.

TASK = *INHIBIT

The TASK user-specific monitoring program is suppressed. All currently running measurements are terminated.

FILE =

The authorization is specified for the FILE user-specific monitoring program.

FILE = *UNCHANGED

The authorization specified at this time for the FILE user-specific monitoring program remains unchanged. After the SM2 subsystem is started, the FILE user-specific monitoring program is not permitted.

FILE = *ALLOW(...)

The FILE user-specific measurement program is permitted. A file is signed on for monitoring by means of the SM2 statement FILE.

USER-ID =

Specifies the users for which the FILE user-specific monitoring program is permitted.

USER-ID = *ALL

The FILE user-specific monitoring program is permitted for all BS2000 users.

USER-ID = list-poss(16): <alphanum-name 1..8>

The FILE user-specific monitoring program is permitted for the users of the specified user IDs. The user IDs must be specified without \$.

FILE = *INHIBIT

The FILE user-specific monitoring program is not permitted. All currently running measurements are terminated.

ISAM =

The authorization is specified for the ISAM user-specific monitoring program.

ISAM = *UNCHANGED

The authorization specified at this point for the ISAM user-specific monitoring program remains unchanged. After the SM2 subsystem is started, the ISAM user-specific monitoring program is not permitted.

ISAM = *ALLOW(...)

The ISAM user-specific monitoring program is not permitted for nonprivileged users. ISAM pools can be included in or excluded from monitoring using the SM2 statements START-/STOP-/CHANGE-ISAM-STATISTICS.

USER-ID =

Specifies the users for which the ISAM user-specific monitoring program is permitted.

USER-ID = *ALL

The ISAM user-specific monitoring program is permitted for all BS2000 users.

USER-ID = list-poss(16): <alphanum-name 1..8>

The ISAM user-specific monitoring program is permitted for the users of the specified user IDs. The user IDs must be specified without \$.

ISAM = *INHIBIT

The ISAM user-specific monitoring program is suppressed. All currently running measurements are terminated.



The restriction regarding the maximum number of objects to be monitored is not changed.

If the *INHIBIT parameter is specified to withdraw execution permission from all users, any objects monitored by a user are automatically deactivated for monitoring purposes. If the list of user IDs is updated (*ALLOW(USER-ID=...)), users who had previously started a monitoring operation can complete this operation, i.e. monitoring is not deactivated automatically.

OPEN-LOG-FILE

Open SM2 output file

This statement is used to open an SM2 output file. The current monitoring cycle is interrupted.

Format

OPEN-LOG-FILE
FILE = <u>*STD</u> / *BY-LINK-NAME ,BUFFER-OUTPUT = <u>*NORMAL</u> / *IMMEDIATE

Operands

FILE =

An SM2 output file is opened; it includes the options listed below. If any other SM2 output file has already been opened at this point in time, it will be closed immediately.

FILE = *STD

A SAM file with the name SM2.hostname.yyyy-mm-dd.sss.nn is opened, where hostname stands for the name of the system as it appears on the screens, yyyy.mm.dd for the day and the point in time the file was created, sss for the number of the BS2000 session and nn for the current number of the SM2 output file within this session (counting from 1). The file is created under the ID from which the OPEN statement was issued. For information on the file attributes, see [section "Writing to the SM2 output file" on page 202](#).

FILE = *BY-LINK-NAME

Opens a SAM or PAM file with the name and file attributes previously defined in a ADD-FILE-LINK command. The link name SMLINK is to be used.

BUFFER-OUTPUT =

Defines file output.

BUFFER-OUTPUT = *NORMAL

A buffer, when full, is always written to a file.

BUFFER-OUTPUT = *IMMEDIATE

At the end of each monitoring cycle, a record for identifying the end of the cycle is written, after which the output buffer is output immediately. Even buffers which are not completely filled with data are output. This could lead to block corruption. Consequently, a file generated in this way should not be processed using SM2U1; the sections of the block not filled with data are removed in this case. This method of file output should be employed when using ANALYZER if the opened file is to be analyzed.

REMOVE-BCAM-CONNECTION-SET**Exclude connection set(s) from BCAM-CONNECTION monitoring**

This statement is used to define the connection set(s) to be removed from the BCAM-CONNECTION monitoring program definition.

Format

REMOVE-BCAM-CONNECTION-SET
SET-NAME = <u>*ALL</u> / <alphanum-name 1..16>

Operands**SET-NAME =**

Specifies the name of the connection set to be removed from the monitoring program definition.

SET-NAME = *ALL

All connection sets are to be removed from the monitoring program definition.

SET-NAME = <alphanum-name 1..16>

Name of the connection set to be removed from the monitoring program definition.

REMOVE-CONNECTION-SET

Exclude connection set(s) from RESPONSETIME monitoring

This statement is used to define the connection set(s) to be removed from the RESPONSETIME monitoring program definition. The connection set *GLOBAL (see also ADD-CONNECTION-SET on [page 91](#)) cannot be removed.

Format

REMOVE-CONNECTION-SET
SET-NAME = <u>*ALL</u> / <alphanum-name 1..16>

Operands

SET-NAME = *ALL

All connection sets except *GLOBAL are to be removed.

SET-NAME = <alphanum-name 1..16>

Name of the connection set to be removed from the monitoring program definition.

REMOVE-COSMOS-EVENT**Remove events from COSMOS monitoring program**

This statement is used to define events which are to be removed from the COSMOS monitoring program definition.

Format**REMOVE-COSMOS-EVENT**

```
EVENT-NAME = *STANDARD-EVENTS / *ALL / list-poss(71): *STANDARD-EVENTS / *ACF / *BCAM /
*BCPT / *BLS / *BOUR / *CHTM / *CMD / *CMS / *DAB / *DCAM / *DLM / *DSM / *EIA / *EIA2 /
*EIA3 / *FITC / *GSAC / *HAL / *IDLE / *INTR / *IONQ / *ISEV / *ISPL / *KAI / *LOCK / *MSG /
*NSM / *PAGE / *PAM / *PCCC / *PCTC / *PDEA / *PEND / *PIO / *PMIO / *PRGS /
*PRGT / *PRTY / *RELM / *REQM / *RSCS / *RSCT / *SDV / *SLOT / *SNAP / *STD1 /
*STDI / *STDN / *SVC / *SWSR / *TGMA / *TGMP / *TGMT / *TIC / *TINF / *TLM / *TLT /
*TSKI / *TSVC / *UTM / *VMCH / *VMH / *VMI / *VMLK / *VMPD / *VMPR / *VMS / *VM2 /
*WSCT / *XEIA
```

Operands**EVENT-NAME =**

Specifies the names of the events to be removed from the monitoring program definition.

EVENT-NAME = *STANDARD-EVENTS

The following standard events are to be removed from the monitoring program definition: ACF, BLS, BOUR, CHTM, CMS, DAB, EIA2, EIA3, FITC, IDLE, INTR, IONQ, PAGE, PAM, PCCC, PCTC, PEND, PMIO, PRGS, PRGT, RELM, REQM, SDV, SVC, TSKI, TSVC, WSCT.

EVENT-NAME = *ALL

All events – except the mandatory events CREA, DEST, INIT, LGON, MMRC, PTSK and STAT – are removed from the monitoring program definition.

EVENT-NAME = list-poss(71): *STANDARD-EVENTS / ...

The events specified in the list are removed from the monitoring program definition. The mandatory events CREA, DEST, INIT, LGON, MMRC, PTSK and STAT are not removed.

REMOVE-FILE

Remove file(s) from FILE monitoring program

This statement is used to define the file(s) to be removed from the FILE monitoring program definition.

Format

REMOVE-FILE
FILE-NAME = <u>*ALL</u> / <filename 1..54>

Operands

FILE-NAME = *ALL

All files are to be removed from the monitoring program definition.

FILE-NAME = <filename 1..54>

The name of the file to be removed from the monitoring program definition. The file name must be specified fully qualified.

REMOVE-ISAM-FILE**Remove ISAM pool from ISAM monitoring program**

This statement defines the pathname of a NK-ISAM file which is to be removed from the ISAM monitoring program definition.

Format

REMOVE-ISAM-FILE
FILE-NAME = <u>*ALL</u> / <filename 1..54>

Operands**FILE-NAME = *ALL**

All names of NK-ISAM files are to be removed from the monitoring program definition.

FILE-NAME = <filename 1..54>

Specifies the name of the NK-ISAM file to be removed from the monitoring program definition. The file name must be specified fully qualified.

REMOVE-ISAM-POOL**Remove ISAM pool(s) from ISAM monitoring program**

This statement is used to define the ISAM pool(s) to be removed from the ISAM monitoring program definition.

Format**REMOVE-ISAM-POOL**

```

POOL-NAME = *ALL / <alphanum-name 1..8>
,SCOPE = *ANY / *HOST-SYSTEM / *TASK(...) / *USER(...)
      *TASK(...)
      |   TSN = <alphanum-name 1..4>
,CAT-ID = *ANY / *HOME / <cat-id 1..4>

```

Operands**POOL-NAME = *ALL**

All ISAM pools are to be removed from the monitoring program definition. SCOPE and CAT-ID are not taken into account in this case.

POOL-NAME = <alphanum-name 1..8>

Specifies the name of the ISAM pool to be removed from the monitoring program definition.

SCOPE =

Defines the type of the ISAM pool to be removed from the monitoring program definition.

SCOPE = *ANY

Global and task-local pools are to be removed from the monitoring program definition.

SCOPE = *HOST-SYSTEM

The pool is a global ISAM pool.

SCOPE = *TASK(...)

The pool is a task-local ISAM pool.

TSN = <alphanum-name 1..4>

Specifies the TSN of the task under which the task-local ISAM pool has been set up.

CAT-ID =

Specifies the catalog ID of the pubset of the ISAM pool to be removed from the monitoring program definition.

CAT-ID = *ANY

ISAM pools with any catalog ID are to be removed from the monitoring program definition.

CAT-ID = *HOME

The catalog ID is the catalog ID of the home pubset.

CAT-ID = <cat-id 1..4>

Catalog ID of the pubset to which the ISAM pool is assigned.

REMOVE-OPENFT-INSTANCE**Remove openFT instance for OPENFT monitoring program**

This statement is used to define openFT instances which are to be removed from the monitoring program definition of OPENFT.

Format

REMOVE-OPENFT-INSTANCE
INSTANCE-NAME = <u>*ALL</u> / <alphanum-name 1..8>

Operands**INSTANCE-NAME = *ALL**

All openFT instances are to be removed from the monitoring program definition.

INSTANCE-NAME = <alphanum-name 1..8>

Name of the openFT instance which is to be removed from the monitoring program definition.

SELECT-HOSTS

Define hosts for SM2 statements and screen output

This statement is used to define the hosts for which SM2 administration statements are to apply and for which screens are to be output. If reports are output, the selected remote hosts send their monitored data to the local host.

The statement only works locally for the user who issues it. Any caller with administration rights can define his or her own list of hosts. The END statement deletes the definition.

Format

SELECT-HOSTS

```
HOST-NAME = *ALL / *LOCAL / list-poss(16): <alphanum-name 1..8>  
,PARTNER-TYPE = *XCS / *CCS
```

Operands

HOST-NAME =

Specifies the appropriate host names.

HOST-NAME = *ALL

Selects all hosts.

HOST-NAME = *LOCAL

Selects the user's own host.

HOST-NAME = list-poss(16): <alphanum-name 1..8>

Selects the hosts specified via the host names.

PARTNER-TYPE =

Specifies the MSCF-specific partner type of the host.

PARTNER-TYPE = *XCS

Selects hosts of partner type XCS.

PARTNER-TYPE = *CCS

Selects hosts of partner type CCS.



The host name uniquely identifies a host in a processor network. This name must be specified in the SELECT-HOSTS statement and is output on the screens as the host name.

Note that when an HIPLEX MSCF connection is set up, a processor name must be specified.

The processor name can be different from the host name. The processor name only identifies a system locally for one other system, not globally for the whole network. The host name is the same as the local processor name.

SET-BCAM-CONNECTION-PARAMETERS

Define BCAM-CONNECTION monitoring parameters

This statement is used to define the global monitoring parameters for the BCAM-CONNECTION monitoring program.

Format

SET-BCAM-CONNECTION-PARAMETERS

```
INWAIT-BUCKETS = *STD-LIMITS / *UNCHANGED / list-poss(4):<integer 1..999999>
,REACT-BUCKETS = *STD-LIMITS / *UNCHANGED / list-poss(4):<integer 1..999999>
,INPROC-BUCKETS = *STD-LIMITS / *UNCHANGED / list-poss(4):<integer 1..999999>
,OUTPROC-BUCKETS = *STD-LIMITS / *UNCHANGED / list-poss(4):<integer 1..999999>
```

Operands

INWAIT-BUCKETS =

Defines the upper limits of up to four ranges (in units of 1 ms) in which the INWAIT times are to be stored by order of magnitude.

INWAIT-BUCKETS = *STD-LIMITS

Sets the upper limits of the ranges for INWAIT times to 100, 200, 500 and 1000 ms.

INWAIT-BUCKETS = *UNCHANGED

The upper limits of the ranges for INWAIT times currently defined in BCAM are not changed.

INWAIT-BUCKETS = list-poss(4): <integer 1..999999>

Defines the upper limits of the ranges for INWAIT times. These must be specified in ascending order.

REACT-BUCKETS =

Defines the upper limits of up to four ranges (in units of 1 ms) in which the REACT times are to be stored by order of magnitude.

REACT-BUCKETS = *STD-LIMITS

Sets the upper limits of the ranges for REACT times to 500, 1000, 2000 and 4000.

REACT-BUCKETS = *UNCHANGED

The upper limits of the ranges for REACT times currently defined in BCAM are not changed.

REACT-BUCKETS = list-poss(4):<integer 1..999999>

Defines the upper limits of the ranges for REACT times. These must be specified in ascending order.

INPROC-BUCKETS =

Defines the upper limits of up to four ranges (in units of 1 ms) in which the INPROC times are to be stored by order of magnitude.

INPROC-BUCKETS = *STD-LIMITS

Sets the upper limits of the ranges for INPROC times to 100, 200, 500 and 1000 ms.

INPROC-BUCKETS = *UNCHANGED

The upper limits of the ranges for INPROC times currently defined in BCAM are not changed

INPROC-BUCKETS = list-poss(4):<integer 1..999999>

Defines the upper limits of the ranges for INPROC times. These must be specified in ascending order.

OUTPROC-BUCKETS =

Defines the upper limits of up to four ranges (in units of 1 ms) in which the OUTPROC times are to be stored by order of magnitude.

OUTPROC-BUCKETS = *STD-LIMITS

Sets the upper limits of the ranges for OUTPROC times to 100, 200, 500 and 1000 ms.

OUTPROC-BUCKETS = *UNCHANGED

The upper limits of the ranges for OUTPROC times currently defined in BCAM are not changed.

OUTPROC-BUCKETS = list-poss(4):<integer 1..999999>

Defines the upper limits of the ranges for REACT times. These must be specified in ascending order.

SET-CHANNEL-IO-PARAMETERS

Define CHANNEL-IO monitoring parameters

This statement is used to define the channels for the CHANNEL-IO monitoring program.

Format

SET-CHANNEL-IO-PARAMETERS
CHANNELS = <u>*ALL</u> / list-poss(64): <x-string 1..4>

Operands

CHANNELS = *ALL

All channels are to be added to the monitoring program definition.

CHANNELS = list-poss(64): <x-string 1..4>

The channel addresses (CHANNEL-PATH-IDs) of the channels to be added to the monitoring program definition are specified. A check is carried out to establish whether these channels are defined in the configuration.



The CHANNEL-IO monitoring program is started for all channels when the SM2 subsystem is started.

SET-COSMOS-PARAMETERS

Define COSMOS monitoring parameters

This statement is used to define the monitoring parameters for the COSMOS monitoring program. The mandatory events CREA, DEST, INIT, LGON, MMRC, PTSK and STAT are defined automatically. Events specified beforehand with the ADD-COSMOS-EVENT statement remain in the monitoring program definition.

The files specified with OUTPUT and ADDITIONAL-INFO=*VM2000 must already be created.

The PRIMARY-ALLOCATION value of the SPACE operand (CREATE-FILE command) must be at least twice the specified number of buffers for writing the events (NUMBER-OF-BUFFERS) when creating the COSMOS output file(s).

The SECONDARY-ALLOCATION value of the SPACE operand (CREATE-FILE command) must be set to zero when creating the COSMOS output file for WRAP-AROUND.

Format

SET-COSMOS-PARAMETERS
<pre> TITLE = C 'COSMOS' / <c-string 1..80> ,BUFFER-SIZE = <u>7</u> / <integer 1..40> ,NUMBER-OF-BUFFERS = <u>32</u> / <integer 2..512> ,ADDITIONAL-INFO = *<u>CONFIGURATION</u> / *NONE / list-poss(2): *CONFIGURATION / *VM2000(...) *VM2000(...) FILE-NAME = <filename 1..54 without-gen-vers> ,OUTPUT = *<u>DISK</u> (...) / *WRAP-AROUND(...) / *TAPE(...) / *STREAM-TAPE(...) *DISK(...) FILE-NAME = list-poss(16): <filename 1..54 without-gen-vers> *WRAP-AROUND(...) FILE-NAME = <filename 1..54 without-gen-vers> </pre>

(part 1 of 2)

```

*TAPE(...)
  | FILE-NAME = list-poss(16): <filename 1..54 without-gen-vers>
*STREAM-TAPE(...)
  | FILE-NAME = list-poss(16): <filename 1..54 without-gen-vers>
,TASK-SELECTION = *ALL / *SPECIFIED(...)
  *SPECIFIED(...)
    | JOB-NAME = *NOT-SPECIFIED / list-poss(8): <alphanum-name 1..8 with-wild>
    | ,CATEGORY = *NOT-SPECIFIED / list-poss(8): <alphanum-name 1..7 with-wild>
    | ,USER-ID = *NOT-SPECIFIED / list-poss(8): <alphanum-name 1..8 with-wild>
    | ,TSN = *NOT-SPECIFIED / list-poss(16): <alphanum-name 1..4 with-wild>
    | ,TYPE = *NOT-SPECIFIED / list-poss(4): *SYSTEM / *BATCH / *DIALOG / *TP
,EVENT-SELECTION = *ALL-BY-ADD-COSMOS-EVENT / *SPECIFIED(...)
  *SPECIFIED(...)
    | EIA-INTERRUPT-CLASS = *ANY / list-poss(5): *SVC / *PROGRAM / *MACHINE-CHECK / *IO /
    |                                     *EXTERNAL
    | ,EIA-SVC-NUMBER = *ANY / list-poss(8): <integer 1..255>
    | ,IO-DEVICE = *ANY / list-poss(8): <alphanum-name 2..4>
    | ,DAB-CACHE-ID = *ANY / list-poss(8): <alphanum-name 1..32>
    | ,MEMORY-CLASS = *ANY / list-poss(4): *3 / *4 / *5 / *6
    | ,SLOT-MEMORY-CLASS = *ANY / list-poss(4): *3 / *4 / *5 / *6
    | ,PEND-CODE = *ANY / list-poss(16): <integer 1..22>
    | ,LOCK-ID = *ANY / list-poss(4): <alphanum-name 1..2>
    | ,TLT-DESCRIPTOR = *ANY / list-poss(8): <alphanum-name 1..3>
    | ,TSKI-SWITCH = *ANY / *TASK
    | ,TSVC-SVC-NUMBER = *ANY / list-poss(8): <integer 1..255>
    | ,CPU-NUMBER = *ANY / list-poss(32): <integer 0..31>
,UNLOAD = *AT-MEASUREMENT-PROGRAM-STOP / *AT-SM2-STOP
,MEASUREMENT-TIME = *NOT-SPECIFIED / <integer 1..60>

```

(part 2 of 2)

Operands**TITLE =**

Specifies the title of the COSMOS monitoring process.

TITLE = C'COSMOS'

Sets the title of the COSMOS monitoring process to "COSMOS".

TITLE = <c-string 1..80>

Sets the title of the COSMOS monitoring process to the value specified here.

BUFFER-SIZE =

Specifies the number of 4K pages per buffer.

BUFFER-SIZE = 7

The buffers created have a default size of seven pages.

BUFFER-SIZE = <integer 1..40>

The buffers created have the specified number of pages.

NUMBER-OF-BUFFERS =

Specifies the number of buffers.

NUMBER-OF-BUFFERS = 32

By default, 32 buffers are created for writing events.

NUMBER-OF-BUFFERS = <integer 2..512>

COSMOS creates the specified number of buffers for writing the events.

ADDITIONAL-INFO =

Specifies the additional data to be recorded.

ADDITIONAL-INFO = *CONFIGURATION

The configuration at the start of monitoring is also recorded.

ADDITIONAL-INFO = *NONE

No additional data is recorded.

ADDITIONAL-INFO = *VM2000(...)

The VM2000 events are also recorded. For this purpose, they must be explicitly opened (ADD-COSMOS-EVENT statement).

FILE-NAME = <filename 1..54 without-gen-vers>

Specifies the name of the file to which the VM2000 events are written.

OUTPUT =

Specifies how the COSMOS output files are written. The files must have already been created. File link names are no longer necessary.

OUTPUT = *DISK(...)

The COSMOS monitored data is written sequentially to the files on hard disk; i.e. there is one write task for each output file. If a monitored data buffer is full, a write task is activated that writes the buffer to the appropriate file.

FILE-NAME = list-poss(16): <filename 1..54 without-gen-vers>

Specifies the file(s) to which the COSMOS monitored data is written.

OUTPUT = *WRAP-AROUND(...)

COSMOS monitored data is written to a file on disk. If the disk space allocated (PRIMARY-ALLOCATION) is insufficient, the file is overwritten from the top. The secondary allocation (SECONDARY-ALLOCATION) must be set to zero.

FILE-NAME = <filename 1..54 without-gen-vers>

Specifies the file to which COSMOS monitored data is written.

OUTPUT = *TAPE(...)

The COSMOS monitored data is written sequentially to the files on tape (this works as described for OUTPUT=*DISK).

FILE-NAME = list-poss(16): <filename 1..54 without-gen-vers>

Specifies the file(s) to which COSMOS monitored data is written.

OUTPUT = *STREAM-TAPE(...)

The COSMOS output files are written in streaming mode. This reduces the chances of events not being recorded, because the files are written to the tape continuously without any repositioning of the tape.

FILE-NAME = list-poss(16): <filename 1..54 without-gen-vers>

Specifies the file(s) to which COSMOS monitored data is written.

TASK-SELECTION =

Specifies the tasks to be monitored.

TASK-SELECTION = *ALL

All tasks are monitored.

TASK-SELECTION = *SPECIFIED(...)

Specifies the tasks to be monitored by means of their TSN, user ID, job name, category or type.

JOB-NAME =

The tasks are selected by means of their job name.

JOB-NAME = *NOT-SPECIFIED

No tasks are selected by means of their job name.

JOB-NAME = list-poss(8): <alphanum-name 1..8 with-wild>

Specifies the job names of the tasks to be monitored. An asterisk * is also permitted as the last character in a wildcard specification.

CATEGORY =

The tasks are selected by means of their category.

CATEGORY = *NOT-SPECIFIED

No tasks are selected by means of their category.

CATEGORY = list-poss(8): <alphanum-name 1..7 with-wild>

Specifies the categories of the tasks to be monitored. An asterisk * is also permitted as the last character in a wildcard specification.

USER-ID =

The tasks are selected by means of their user ID.

USER-ID = *NOT-SPECIFIED

No tasks are selected by means of their user ID.

USER-ID = list-poss(8): <alphanum-name 1..8 with-wild>

Specifies the user IDs of the tasks to be monitored. These must be specified without the \$ sign. An asterisk * is also permitted as the last character in a wildcard specification.

TSN =

The tasks are selected by means of their TSN.

TSN = *NOT-SPECIFIED

No tasks are selected by means of their TSN.

TSN = list-poss(16): <alphanum-name 1..4 with-wild>

Specifies the TSNs of the tasks to be monitored. An asterisk * is also permitted as the last character in a wildcard specification.

TYPE =

The tasks are selected on the basis of their type.

TYPE = *NOT-SPECIFIED

No tasks are selected on the basis of their type.

TYPE = *SYSTEM

All system tasks are to be monitored.

TYPE = *BATCH

All batch tasks are to be monitored.

TYPE = *DIALOG

All interactive tasks are to be monitored.

TYPE = *TP

All TP tasks are to be monitored.

EVENT-SELECTION =

Specifies the events to be recorded. Their recording is linked to specific conditions.

EVENT-SELECTION = *ALL-BY-ADD-COSMOS-EVENT

All open events are recorded.

EVENT-SELECTION = *SPECIFIED(...)

Only those events that satisfy the specified conditions are recorded.

EIA-INTERRUPT-CLASS =

Controls the recording of EIA events on the basis of their interrupt class (IC).

EIA-INTERRUPT-CLASS = *ANY

The recording of EIA events is independent of their IC.

EIA-INTERRUPT-CLASS = list-poss(5): *SVC / *PROGRAM / *MACHINE-CHECK / *IO / *EXTERNAL

Only those EIA events with the specified IC are recorded.

EIA-SVC-NUMBER =

Controls the recording of EIA events on the basis of their SVC.

EIA-SVC-NUMBER = *ANY

The recording of EIA events is independent of their SVC.

EIA-SVC-NUMBER = list-poss(8): <integer 1..255>

Only those EIA events with the specified SVC(s) are recorded.

IO-DEVICE =

Controls the recording of SDV, CHTM, IONQ, and PMIO events on the basis of their mnemonic device name.

IO-DEVICE = *ANY

The recording of SDV, CHTM, IONQ, and PMIO events is independent of their mnemonic device name.

IO-DEVICE = list-poss(8): <alphanum-name 2..4>

Only those SDV, CHTM, IONQ, and PMIO events with the specified mnemonic device names are recorded.

DAB-CACHE-ID =

Controls the recording of DAB events on the basis of their DAB cache IDs.

DAB-CACHE-ID = *ANY

The recording of DAB events is independent of their DAB cache IDs.

DAB-CACHE-ID = list-poss(8): <alphanum-name 1..32>

Only those DAB events with the specified DAB cache IDs are recorded.

MEMORY-CLASS =

Controls the recording of RELM and REQM events on the basis of their memory class.

MEMORY-CLASS = *ANY

The recording of RELM and REQM events is independent of their memory class.

MEMORY-CLASS = list-poss(4): *3 / *4 / *5 / *6

Only those RELM and REQM events with the specified memory classes are recorded.

SLOT-MEMORY-CLASS =

Controls the recording of SLOT events on the basis of their memory class.

SLOT-MEMORY-CLASS = *ANY

The recording of SLOT events is independent of their memory class.

SLOT-MEMORY-CLASS = list-poss(4): *3 / *4 / *5 / *6

Only those SLOT events with the specified memory classes are recorded.

PEND-CODE =

Controls the recording of PEND events on the basis of their pend code.

PEND-CODE = *ANY

The recording of PEND events is independent of their pend code.

PEND-CODE = list-poss(16): <integer 1..22>

Only those PEND events with the specified pend code are recorded.

LOCK-ID =

Controls the recording of LOCK events on the basis of their lock ID.

LOCK-ID = *ANY

The recording of LOCK events is independent of their lock ID.

LOCK-ID = list-poss(4): <alphanum-name 1..2>

Only those LOCK events with the specified lock IDs are recorded.

TLT-DESCRIPTOR =

Controls the recording of TLT on the basis of their TLT descriptor (TLT = Task Location Table).

TLT-DESCRIPTOR = *ANY

The recording of TLT events is independent of their TLT descriptor.

TLT-DESCRIPTOR = list-poss(8): <alphanum-name 1..3>

Only those TLT events with the specified TLT descriptors are recorded.

TSKI-SWITCH =

Controls the recording of TSKI events on the basis of their TIC (Task in Control).

TSKI-SWITCH = *ANY

The recording of TSKI events is independent of their TIC.

TSKI-SWITCH = *TASK

Only the first TSKI event of a TIC is recorded.

TSVC-SVC-NUMBER =

Controls the recording of TSVC events on the basis of their SVC numbers.

TSVC-SVC-NUMBER = *ANY

The recording of TSVC events is independent of their SVC number.

TSVC-SVC-NUMBER = list-poss(8): <integer 1..255>

The TSVC events are recorded for the specified SVC numbers.

CPU-NUMBER =

Controls the recording of events according to the CPU number.

CPU-NUMBER = *ANY

The recording of events is independent of their CPU number.

CPU-NUMBER = list-poss(32): <integer 0..31>

Only the events for the specified CPU numbers are recorded.

UNLOAD =

Specifies the time at which the COSMOS subsystem is to be unloaded.

UNLOAD = *AT-MEASUREMENT-PROGRAM-STOP

The COSMOS subsystem is unloaded at the end of COSMOS monitoring.

UNLOAD = *AT-SM2-STOP

The COSMOS subsystem is unloaded at the end of SM2 monitoring.

MEASUREMENT-TIME =

Specifies the duration of the monitoring process.

MEASUREMENT-TIME = *NOT-SPECIFIED

Monitoring is terminated by the user with the STOP-MEASUREMENT-PROGRAM statement.

MEASUREMENT-TIME = <integer 1..60>

Specifies the time in minutes after which monitoring should be terminated automatically.

SET-DISK-FILE-PARAMETERS

Define DISK-FILE monitoring parameters

This statement is used to define the disk devices for the DISK-FILE monitoring program.

Format

SET-DISK-FILE-PARAMETERS
DEVICES = list-poss(8): <alphanum-name 2..4>

Operands

DEVICES = list-poss(8): <alphanum-name 2..4>

The mnemonic device names of the disk devices to be monitored are added to the monitoring program definition. A check is carried out to establish whether these disk devices are defined in the configuration.

SET-PERIODIC-TASK-PARAMETERS

Define PERIODIC-TASK monitoring parameters

This statement is used to define the tasks from which monitored data is to be written to the SM2 output file for the PERIODIC-TASK monitoring program.

Format

SET-PERIODIC-TASK-PARAMETERS
LOG-TASKS = * <u>NONE</u> / *ALL / *SPECIFIED(...)
*SPECIFIED(...)
USER-ID = * <u>NOT-SPECIFIED</u> / list-poss(64): <alphanum-name 1..8>
,JOB-NAME = * <u>NOT-SPECIFIED</u> / list-poss(64): <alphanum-name 1..8>
,TSN = * <u>NOT-SPECIFIED</u> / list-poss(64): <alphanum-name 1..4>

Operands

LOG-TASKS =

Specifies the tasks in the monitoring program definition whose monitored data is to be written to the SM2 output file.

LOG-TASKS = *NONE

No records are written to the SM2 output file.

LOG-TASKS = *ALL

Monitored data on all tasks is written to the SM2 output file.

LOG-TASKS = *SPECIFIED(...)

Monitored data on selected tasks is written to the SM2 output file.

USER-ID =

Tasks are selected on the basis of their user ID.

USER-ID = *NOT-SPECIFIED

Tasks are not selected on the basis of their user ID.

USER-ID = list-poss(64): <alphanum-name 1..8>

Monitored data on the tasks with the specified user IDs is written to the SM2 output file. These user IDs must be specified without the \$ sign.

JOB-NAME =

Tasks are selected on the basis of their job names.

JOB-NAME = *NOT-SPECIFIED

Tasks are not selected on the basis of their job names.

JOB-NAME = list-poss(64): <alphanum-name 1..8>

Monitored data on the tasks with the specified job names is written to the SM2 output file.

TSN =

The tasks are selected on the basis of their TSN.

TSN = *NOT-SPECIFIED

No tasks are selected on the basis of the TSN.

TSN = list-poss(64): <alphanum-name 1..4>

The monitored data of the tasks with the specified TSNs is written to the SMS output file.

SET-RESPONSETIME-PARAMETERS

Define RESPONSETIME monitoring parameters

This statement is used to define the monitoring parameters for the RESPONSETIME monitoring program.

Format

SET-RESPONSETIME-PARAMETERS
SCOPE = <u>*BUCKET</u> / list-poss(2): *BUCKET / *CATEGORY ,DEFINITION = <u>*1</u> / *2 ,CONNECTION-NUMBER = <u>*1024</u> / <integer 1..8187> ,RESPONSETIME-BUCKETS = <u>*STD-LIMITS</u> / list-poss(5): <integer 1..99999> ,THINKTIME-BUCKETS = <u>*STD-LIMITS</u> / list-poss(5): <integer 1..99999> ,TRANSACTION-BUCKETS = <u>*STD-LIMITS</u> / list-poss(5): <integer 1..99999> ,WAITTIME-BUCKETS = <u>*STD-LIMITS</u> / list-poss(5): <integer 1..99999>

Operands

SCOPE =

Defines whether the response time data is to be recorded by bucket or by category.

SCOPE = *BUCKET

The response time data is to be recorded by bucket.

SCOPE = *CATEGORY

The response time data is to be recorded by category.

DEFINITION =

Defines the type of response time to be monitored.

DEFINITION = *1

This response time is the time between user input and the next output operation.

DEFINITION = *2

In addition to the response time defined with 1, the times between follow-up output operations to an input are recorded individually as response times.

CONNECTION-NUMBER =

Defines the maximum number of connections to be monitored.

CONNECTION-NUMBER = *1024

A maximum of 1024 connections are recorded.

CONNECTION-NUMBER = <integer 1..8187>

The maximum specified number of connections is recorded.

RESPONSE-TIME-BUCKETS =

Defines the upper limits of up to five ranges (in units of 100ms) in which the response times are to be stored by order of magnitude.

RESPONSE-TIME-BUCKETS = *STD-LIMITS

Sets the upper limits of the ranges for response times to 5, 10, 20, 50 and 100.

RESPONSE-TIME-BUCKETS = list-poss(5): <integer 1..99999>

Defines the upper limits of the ranges for response times. These must be specified in ascending order.

THINK-TIME-BUCKETS =

Defines the upper limits of up to five ranges (in units of 100ms) in which the think times are to be stored by order of magnitude.

THINK-TIME-BUCKETS = *STD-LIMITS

Sets the upper limits of the ranges for think times to 50, 150, 300, 600 and 1200.

THINK-TIME-BUCKETS = list-poss(5): <integer 1..99999>

Defines the upper limits of the ranges for think times. These must be specified in ascending order.

TRANSACTION-TIME-BUCKETS =

Defines the upper limits of up to five ranges (in units of 100ms) in which the transaction times are to be stored by order of magnitude.

TRANSACTION-TIME-BUCKETS = *STD-LIMITS

Sets the upper limits of the ranges for transaction times to 5, 10, 20, 50 and 100.

TRANSACTION-TIME-BUCKETS = list-poss(5): <integer 1..99999>

Defines the upper limits of the ranges for transaction times. These must be specified in ascending order.

WAITTIME-BUCKETS =

Defines the upper limits of up to five ranges (in units of 100ms) in which the wait times in the BCAM pool are to be stored by order of magnitude.

WAITTIME-BUCKETS = *STD-LIMITS

Sets the upper limits of the ranges for wait times in the BCAM pool to 1, 2, 5, 10 and 20.

WAITTIME-BUCKETS = list-poss(5): <integer 1..99999>

Defines the upper limits of the ranges for wait times in the BCAM pool. These must be specified in ascending order.

SET-SAMPLING-DEVICE-PARAMETERS**Monitoring program definition for SAMPLING-DEVICE monitoring**

This statement is used to specify the monitoring parameters for the SAMPLING-DEVICE monitoring program.

Format

SET-SAMPLING-DEVICE-PARAMETERS
DISK-SERVICETIME = * <u>OFF</u> / *ON

Operands**DISK-SERVICETIME =**

It is specified in the monitoring program definition whether the service times for disk devices are recorded.

DISK-SERVICETIME = *OFF

It is specified in the monitoring program definition that the service times for disk devices are not recorded.

DISK-SERVICETIME = *ON

It is specified in the monitoring program definition that the service times for disk devices are recorded.



The SAMPLING-DEVICE monitoring program is started when the SM2 subsystem is started without recording of service times for disk devices.

SET-SERVICETIME-PARAMETERS

Define SERVICETIME monitoring parameters

This statement is used to define the devices for the SERVICETIME monitoring program.

Format

SET-SERVICETIME-PARAMETERS
<pre> DEVICES = *SPECIFIED(...) *SPECIFIED(...) DEVICE = *NOT-SPECIFIED / list-poss(256): <alphanum-name 2..4 with-wild> ,PUBSET = *NOT-SPECIFIED / list-poss(64): <alphanum-name 1..4> </pre>

Operands

DEVICES = *SPECIFIED(...)

The specified devices are added to the monitoring program definition.

DEVICE = *NOT-SPECIFIED

No devices are added to the monitoring program definition on the basis of their device name.

DEVICE = list-poss(256): <alphanum-name 2..4 with-wild>

The mnemonic devices names of the devices to be monitored are added to the monitoring program definition. A check is carried out to establish whether the devices are defined in the configuration. "*" is permitted as a wildcard as the last character.

PUBSET = *NOT-SPECIFIED

No devices of a pubset are added to the monitoring program definition on the basis of their pubset name.

PUBSET = list-poss(64): <alphanum-name 1..4>

All devices of a pubset are added to the monitoring program definition on the basis of their pubset name.



A maximum of 256 devices are added to the monitoring program definition. If there is a partial qualification or a selection via pubsets, the devices are selected in the following order of priority:

- fully qualified devices
- devices who belong to a pubset (in the sequence in the PDT (Physical Device Table))
- partially qualified devices (in the sequence of the partially qualified device-names and in the PDT (Physical Device Table))

SET-STORAGE-SYSTEM-PARAMETERS

Define monitoring parameters for the monitoring program STORAGE-SYSTEM

The statement specifies the monitoring parameters for the STORAGE-SYSTEM monitoring program.

Format

SET-STORAGE-SYSTEM-PARAMETERS
<pre> ADDITIONAL-DATA = <u>*NONE</u> / *ALL / *SYMMETRIX(...) *SYMMETRIX(...) TYPE = <u>*NONE</u> / *ALL / list-poss(3): *DEVICE(...) / *PHYSICAL-DISK / *DIRECTOR *DEVICE(...) DEVICES = *ALL / <u>*DEVICES-WITH-MN</u> ,BACK-END-DATA = *YES / <u>*NO</u> </pre>

Operands

ADDITIONAL-DATA =

Defines which data should be measured in addition to the global measured data of storage systems. The global measured data is the data which relates to the storage system as a whole.

ADDITIONAL-DATA = *NONE

No additional data is recorded in addition to the global data. This setting means currently means that only the data of the Symmetrix report is collected.

ADDITIONAL-DATA = *ALL

Beside the global data for storage systems, currently the data for the individual devices (logical volumes), physical disks and directors (adapters) is also collected for Symmetrix systems.

ADDITIONAL-DATA = *SYMMETRIX(...)

Defines the additional data which is to be collected for Symmetrix systems.

TYPE = *NONE

No additional data is collected for Symmetrix systems.

TYPE = *ALL

Beside the global data for Symmetrix systems data for the individual devices (logical volumes), physical disks and directors (adapter) is also collected for Symmetrix systems.

TYPE = list-poss(3):*DEVICE(...)

Defines the amount of devices (logical volumes) and the additional data that are to be collected.

DEVICES =

Defines the amount of devices (logical volumes) for for which data is to be collected.

DEVICES = *ALL

Data is to be collected for all devices (logical volumes).

DEVICES = *DEVICES-WITH-MN

Data is to be collected only for devices (logical volumes) whose (BS2000) mnemonic is known.

Knowledge of the BS2000 mnemonic depends on the connection type:

- S-CON
At least one device (logical volume) of at most 256 devices must be attached to the control unit, then all BS2000 mnemonics of this control unit are visible for all devices.
- FibreChannel
Only the BS2000 mnemonics of the attached devices (logical volume) are known.

BACK-END-DATA =

Defines whether the back-end data is also to be collected for the devices (logical volume) selected above. The back-end data relates to the IOs between the cache and the physical disks in the Symmetrix system per device (logical volume).

BACK-END-DATA = *YES

The back-end data is to be collected.

BACK-END-DATA = *NO

The back-end data is not to be collected.

TYPE = list-poss(3):*PHYSICAL-DISK

Beside the global data for Symmetrix systems the data for the physical disks in the Symmetrix system is also collected.

TYPE = list-poss(3):*DIRECTOR

Beside the global data for Symmetrix systems the data for the directors (adapters) in the Symmetrix system is also collected.

SET-SYSTEM-PARAMETERS

Define SYSTEM monitoring parameters

This statement is used to define the devices for the SYSTEM monitoring program.

Format

SET-SYSTEM-PARAMETERS
<pre> DEVICES = <u>*NONE</u> / *ALL / *SPECIFIED(...) *SPECIFIED(...) DEVICE = <u>*NOT-SPECIFIED</u> / list-poss(256): *DISK / *TAPE / <alphanum-name 2..4 with-wild> ,PUBSET = <u>*NOT-SPECIFIED</u> / list-poss(64): <alphanum-name 1..4> </pre>

Operands

DEVICES =

The specified devices are added to the monitoring program definition.

DEVICES = *NONE

No devices are added to the monitoring program definition.

DEVICES = *ALL

All devices accept the communication components are added to the monitoring program definition.

DEVICES = *SPECIFIED(...)

The specified devices and/or pubsets are added to the monitoring program definition.

DEVICE = *NOT-SPECIFIED

No devices are added to the monitoring program definition on the basis of their device name.

DEVICE = list-poss(256): *DISK

All disk devices are added to the monitoring program definition.

DEVICE = list-poss(256): *TAPE

All tape devices are added to the monitoring program definition.

DEVICE = list-poss(256): <alphanum-name 2..4 with-wild>

The mnemonic device names of the devices to be monitored are added to the monitoring program definition. A check is carried out to establish whether the devices are defined in the configuration. * is permitted as a wildcard as the last character.

PUBSET = *NOT-SPECIFIED

No pubsets are added to the monitoring program definition on the basis of their pubset name.

PUBSET = list-poss(64): <alphanum-name 1..4>

Pubsets are added to the monitoring program definition on the basis of their pubset name.



SM2R1 uses the devices defined here to determine, among other things, the time equivalent for the productive performance of the peripherals (ANALYZER (report group RST, report "IO resource service time for category"); SM2R1 (report group RST, report 61)) and the dilation factor (ANALYZER (report group Dilation, report "DILATION for categories"); SM2R1 (report group DILATION, report 57)).

If possible, all devices should be monitored to capture precise values. If these reports are required, it is recommended that you specify DEVICE=(*DISK,*TAPE).

SET-TASK-PARAMETERS

Define TASK monitoring parameters

This statement is used to define the tasks and devices for the TASK monitoring program. All tasks can be selected, but this results in a higher system load. Tasks can be explicitly defined using a list of TSNs, user IDs, job names and job classes. Tasks can also be selected by specifying a task attribute.

If tasks are defined by specifying multiple selection criteria, the selection is made by logical ORing, i.e. a task is selected if at least one of the specified criteria has been satisfied.

Format

SET-TASK-PARAMETERS
<pre> TASK-SELECTION = *ALL / *SPECIFIED(...) *SPECIFIED(...) TSN = *NOT-SPECIFIED / list-poss(64): <alphanum-name 1..4> ,USER-ID = *NOT-SPECIFIED / list-poss(32): <alphanum-name 1..8 with-wild> ,JOB-NAME = *NOT-SPECIFIED / list-poss(32): <alphanum-name 1..8 with-wild> ,JOB-CLASS = *NOT-SPECIFIED / list-poss(32): <alphanum-name 1..8 with-wild> ,TYPE = *NOT-SPECIFIED / list-poss(4): *SYSTEM / *BATCH / *DIALOG / *TP ,DEVICES = *NONE / *ALL / *SPECIFIED(...) *SPECIFIED(...) DEVICE = *NOT-SPECIFIED / list-poss(64): *DISK / <alphanum-name 2..4 with-wild> ,PUBSET = *NOT-SPECIFIED / list-poss(64): <alphanum-name 1..4> </pre>

Operands

TASK-SELECTION =

The specified tasks are added to the monitoring program definition.

TASK-SELECTION = *ALL

All tasks are added to the monitoring program definition.

TASK-SELECTION = *SPECIFIED(...)

Tasks are selected on the basis of their TSN, user ID, job name, job class or type and added to the monitoring program definition.

TSN =**TSN = *NOT-SPECIFIED**

No tasks are added to the monitoring program definition on the basis of their TSN.

TSN = list-poss(64): <alphanum-name 1..4>

Tasks are added to the monitoring program definition on the basis of their TSN. A task with the specified TSN is only monitored provided it already exists at startup of the TASK monitoring program.

USER-ID =**USER-ID = *NOT-SPECIFIED**

No tasks are added to the monitoring program definition on the basis of their user ID.

USER-ID = list-poss(32): <alphanum-name 1..8>

Tasks are added to the monitoring program definition on the basis of their user ID. * is permitted as a wildcard as the last character. All tasks of the specified user ID that already exist at startup of the TASK monitoring program are monitored, as well as all those that are created during the monitoring program run. The user IDs are to be specified without \$.

JOB-NAME =**JOB-NAME = *NOT-SPECIFIED**

No tasks are added to the monitoring program definition on the basis of their job names.

JOB-NAME = list-poss(32): <alphanum-name 1..8>

Tasks are added to the monitoring program definition on the basis of the job name. * is permitted as a wildcard as the last character. All tasks with the specified job name that have the specified job name at the time of selection (START-MEASUREMENT-PROGRAM or LOGON) are monitored.

JOB-CLASS =**JOB-CLASS = *NOT-SPECIFIED**

No tasks are added to the monitoring program definition on the basis of their job class.

JOB-CLASS = list-poss(32): <alphanum-name 1..8>

Tasks are added to the monitoring program definition on the basis of the job class. * is permitted as a wildcard as the last character. All tasks of the specified job class that belong to the job class at the time of selection (START-MEASUREMENT-PROGRAM or LOGON) are monitored.

TYPE =

Tasks are added to the monitoring program definition on the basis of their task attribute. All tasks that have the specified task attribute at the time of selection (START-MEASUREMENT-PROGRAM or LOGON) are monitored. If a task changes the task attribute during the monitoring program run, this has no effect on monitoring. In other words, it is not added to or removed from what is monitored as a result of this.

TYPE = *NOT-SPECIFIED

No tasks are added to the monitoring program definition on the basis of the task attribute.

TYPE = list-poss(4): *SYSTEM / *BATCH / *DIALOG / *TP

*SYSTEM: All system tasks are added to the monitoring program definition.

*BATCH: All batch tasks are added to the monitoring program definition.

*DIALOG: All dialog tasks are added to the monitoring program definition.

*TP: All TP tasks are added to the monitoring program definition.

DEVICES =

The specified devices are added to the monitoring program definition.

DEVICES = *NONE

No devices are added to the monitoring program definition.

DEVICES = *ALL

All devices are added to the monitoring program definition.

DEVICES = *SPECIFIED(...)

The specified devices and/or pubsets are added to the monitoring program definition.

DEVICE = *NOT-SPECIFIED

No devices are added to the monitoring program definition on the basis of their device name.

DEVICE = list-poss(64): *DISK

All disk devices are added to the monitoring program definition.

DEVICE = list-poss(64): <alphanum-name 2..4 with-wild>

The specified devices are added to the monitoring program definition. * is permitted as a wildcard as the last character.

PUBSET = *NOT-SPECIFIED

No pubsets are added to the monitoring program definition on the basis of their pubset name.

PUBSET = list-poss(64): <alphanum-name 1..4>

Pubsets are added to the monitoring program definition on the basis of their pubset name.



A maximum of 64 devices are added to the monitoring program definition. If there is a partial qualification, or a selection via pubset or *DISK or *ALL, the devices are selected in the following order of priority:

- fully qualified devices
- devices who belong to a pubset (in the sequence in the PDT (Physical Device Table) and in the pubsets)
- partially qualified devices (in the sequence of the partially qualified device-names and in the PDT (Physical Device Table)
- *DISK (disks in the sequence in the PDT)
- *ALL (alle devices in the sequence in the PDT, except remote data processing devices)

SHOW-ACTIVE-PARAMETERS

Output active monitoring parameters

This statement is used to output the current objects and monitoring parameters for each active monitoring program.

Format

SHOW-ACTIVE-PARAMETERS

```
TYPE = *ALL / list-poss(14): *BCAM-CONNECTION / *CHANNEL-IO / *COSMOS / *DISK-FILE / *FILE /
    *ISAM / *OPENFT / *PERIODIC-TASK / *RESPONSETIME / *SAMPLING-DEVICE /
    *SERVICETIME / *STORAGE-SYSTEM / *SYSTEM / *TASK
```

Operands

TYPE =

Specifies the name of the active monitoring program for which the current objects and monitoring parameters are to be output, see [chapter “SM2 monitoring programs” on page 49](#).

TYPE = *ALL

Outputs the objects and monitoring parameters for all active monitoring programs.

TYPE = *BCAM-CONNECTION



Only the bucket values currently set in BCAM are displayed.

SHOW-DEFINED-PARAMETERS

Output defined monitoring parameters

This statement is used to output the objects and monitoring parameters currently defined for each definable monitoring program.

Format

SHOW-DEFINED-PARAMETERS

```
TYPE = *ALL / list-poss(14): *BCAM-CONNECTION / *CHANNEL-IO / *COSMOS / *DISK-FILE / *FILE /  
*ISAM / *OPENFT / *PERIODIC-TASK / *RESPONSETIME / *SAMPLING-DEVICE /  
*SERVICETIME / *STORAGE-SYSTEM / *SYSTEM / *TASK
```

Operands

TYPE =

Specifies the name of the monitoring program for which the currently defined objects and monitoring parameters are to be output, see [chapter “SM2 monitoring programs” on page 49](#).

TYPE = *ALL

Outputs the objects and monitoring parameters currently defined for all definable monitoring programs.

TYPE = *BCAM-CONNECTION

Outputs the bucket values currently set in BCAM, provided the SET-BCAM-CONNECTION-PARAMETERS statement has not yet been issued. If the SET-BCAM-CONNECTION-PARAMETERS statement is entered with *UNCHANGED, the bucket values currently set in BCAM are adopted by SM2 and can be output using the SHOW-DEFINED-PARAMETERS statement.

SHOW-MEASUREMENT-STATUS

Output monitoring status

This statement is used to output the MEASUREMENT STATUS screen.

Format

SHOW-MEASUREMENT-STATUS

SHOW-SELECTED-HOSTS

Output selected hosts

This statement is used to output the hosts selected with the SELECT-HOSTS statement. The display includes the host name, the processor name (from the point of view of the host from which the statement was issued), and the time of the last monitored data queried.

If there is no valid or new data, RSLT NOT VALID is output instead of the time. Possible reasons for this are listed below:

- Monitored data was not queried since the SELECT-HOSTS statement.
- The remote host is not sending.
- The remote host shows a different system time or uses a different SM2 monitoring cycle. Data from the respective hosts (including the local host) is output separately.

Format

SHOW-SELECTED-HOSTS

SHOW-SM2-STATUS

Output status of SM2 system tasks

This statement is used to output information on the status of individual SM2 system tasks and of the subsystems used by SM2.

Format

SHOW-SM2-STATUS
<pre> INFORMATION = *STD / *TASK / *EVENT-TRACE(...) *EVENT-TRACE(...) FROM = *ACTUAL / <integer 1..32767> </pre>

Operands

INFORMATION =

Provides information on the status of the SM2 system components.

INFORMATION = *STD

Provides information on the status of the SM2 system tasks and of the subsystems used by SM2. The most recent event trace entries are also shown. The event trace entries contain important events during execution of SM2. These events include:

- start and termination of SM2 tasks
- beginning and end of the use of other subsystems
- any errors which occurred during execution of SM2 components

INFORMATION = *TASK

The status of the SM2 system tasks and of the subsystems used by SM2 is output.

INFORMATION = *EVENT-TRACE(...)

Event trace entries are output.

FROM =

Specifies the position as of which the trace entries are to be output. The position specified corresponds to the most recent entry (the one with the highest number). If trace entries which do not yet exist are selected, the most recent ones are output. The fact that the trace buffer is overwritten cyclically means that it is possible that the required entries are no longer all available. The last available entry is marked with the symbol "===>".

FROM = *ACTUAL

The most recent trace entries are output.

FROM = <integer 1..32767>

The trace entries are output as of the specified number.



For those users who are interested, the layout of the event trace entries is described under the new STATUS TABLE screen.

SHOW-USER-MEASURED-OBJECTS

Output monitored objects and associated users

This statement is used to output the objects currently being monitored by users, together with the associated users.

Format

SHOW-USER-MEASURED-OBJECTS
TYPE = <u>*ALL</u> / list-poss(3): *FILE / *ISAM / *TASK

Operands

TYPE =

Specifies the name of the user monitoring program for which the currently monitored objects and corresponding users are to be output.

TYPE = *ALL

Outputs the objects currently monitored and the associated users for all user monitoring programs.

TYPE = *FILE

Outputs the objects currently monitored and the associated users for the FILE monitoring program.

TYPE = *ISAM

Outputs the objects currently monitored and the associated users for the ISAM monitoring program.

TYPE = *TASK

Outputs the objects currently monitored and the associated users for the TASK monitoring program.

START-MEASUREMENT-PROGRAM

Start monitoring program run

This statement is used to start the monitoring programs specified under TYPE.

Format

START-MEASUREMENT-PROGRAM

```
TYPE = list-poss(32): *BCAM-CONNECTION / *CHANNEL-IO / *CMS / *COSMOS / *DAB / *DISK-FILE
      *DLM / *FILE / *GS / *GSVOL / *HSMS / *ISAM / *MSCF / *NSM / *OPENFT / *PERIODIC-TASK /
      *PFA / *POSIX / *PUBSET / *RESPONSETIME / *SAMPLING-DEVICE / *SERVICETIME /
      *SESAM-SQL / *STORAGE-SYSTEM / *SVC / *SYSTEM / *TASK / *TCP-IP / *TLM / *UDS-SQL /
      *UTM / *VM
```

Operands

TYPE =

Specifies the name of the monitoring program to be started, see [chapter "SM2 monitoring programs" on page 49](#).



If the COSMOS monitoring program is to be started, it must have been defined. If the monitoring program has not been prepared, it is prepared implicitly (INITIATE-COSMOS).

For the BCAM-CONNECTION, CHANNEL-IO, COSMOS, DISK-FILE, FILE, ISAM, OPENFT, PERIODIC-TASK, RESPONSETIME, SERVICETIME, STORAGE-SYSTEM, SYSTEM and TASK monitoring programs, the monitoring parameters/monitored objects must be defined using the appropriate statements before the program is started.

STOP-MEASUREMENT-PROGRAM

Terminate monitoring program run

This statement is used to terminate the monitoring programs specified under TYPE.

Format

STOP-MEASUREMENT-PROGRAM

```
TYPE = *ALL / list-poss(34): *BCAM-CONNECTION / *CHANNEL-IO / *CMS / *COSMOS / *DAB / *DISK /  
*DISK-FILE / *DLM / *FILE / *GS / *GSVOL / *HSMS / *ISAM / *MSCF / *NSM / *OPENFT /  
*PERIODIC-TASK / *PFA / *POSIX / *PUBSET / *RESPONSETIME / *SAMPLING-DEVICE /  
*SERVICETIME / *SESAM-SQL / *STORAGE-SYSTEM / *SVC / *SYSTEM / *TASK / *TCP-IP /  
*TLM / *UDS-SQL / *UTM / *VM / *X2000
```

Operands

TYPE =

Specifies the name of monitoring program to be terminated, see [chapter “SM2 monitoring programs” on page 49](#).

TYPE = *ALL

Terminates all active monitoring programs (except SAMPLING-DEVICE).

5.4 Statements for nonprivileged users

These statements enable the nonprivileged user to select monitored data for output and control the screen output. They can also be used by the SM2 administrator provided he/she switches to the analysis subinterval using the CALL-EVALUATION-PART statement. This automatically resets the status of the administrator; privileges and special authorizations associated with some statements remain unchanged (see [section “Users” on page 24](#)).

These SM2 functions cannot be addressed through SDF.

For a description of the ISP syntax, see [page 169](#).

Default values for statements

All functions covered by the SM2 statements are predefined by means of default values at program start time. The user must enter statements himself only if he wishes to change the defaults or activate additional (optional) output operations.

An exception to the above is the START statement, which the user must employ to initiate output of a report, as well as the END statement, which then terminates the SM2 session for the particular user.

Statements for selecting monitored data for screen output

The following statements define the monitored data which is to be output to the individual reports or which need not be output.

Statement	Function
FILE	Select files to be included in or excluded from monitoring in the next monitoring interval
SELECT-CHANNEL-PARAMETERS	Specify sort criterion for the CHANNEL report
SELECT-CMS-PUBSET	Select pubsets for the CMS report
SELECT-DAB-CACHE	Select DAB cache
SELECT-DEVICE-DISK-PARAMETERS	Specify sort criterion for the DEVICE DISK report
SELECT-PERIODIC-TASK-PARAMETERS	Select sort criterion and output information for the PERIODIC TASK report
SELECT-SYMMETRIX-PUBSET	Select pubset of a Symmetrix system for the SYMMETRIX DEVICE report
SELECT-SYMMETRIX-UNIT	Select a Symmetrix system for the SYMMETRIX PUBSET report

Statement	Function
SELECT-UTM-APPLICATION	Select UTM applications
SHOW-USER-MEASURED-OBJECTS	Output monitored objects
START-ISAM-STATISTICS	Monitor ISAM pools
CHANGE-ISAM-STATISTICS	Include and exclude ISAM pools
STOP-ISAM-STATISTICS	Terminate the monitoring of ISAM pools

Statements for controlling reports

The following statements are used to control reports.

Statement	Function
OUTPUT	Define the output mode. Reports are output either automatically at regular intervals (automatic output mode) or at the request of the user (controlled mode).
REPORT	Select reports which can be output using START or RESTART.
RESTART	Start the output of selected reports. Unlike the START statement, this statement does not allow for the retrieval of new monitored data.
START	Start report output

In **automatic output mode**, the selected reports are output regularly over a defined period of time. The overflow screens are not output in this mode.

The period of time over which reports are output is known as the **output cycle**. When the program is started, this is preset to the current length of the monitoring cycle, but can be changed by the user. The interval between reports depends on the number of reports output in the last output cycle. In the first output cycle, reports are output without a delay. The user must ensure that a practical relationship is maintained between the monitoring cycle and the output cycle. Further information can be found under [“Relationship between the monitoring cycle and the output cycle” on page 165](#).

In **controlled mode**, the individual reports are output at the request of the user, i.e. the user can control SM2 report output or enter another statement.

The following scroll commands are available for controlling the output of SM2 reports which may comprise one or more screens:

1. **++** selects the first screen of the first report of the next monitoring cycle.
2. **--** selects the first screen of the first report of the current monitoring cycle.
3. **+R** selects the first screen of the next report.
If the current screen is the last for this monitoring cycle, the first screen of the first report of the next monitoring cycle is displayed.
4. **-R** selects the first screen of the preceding report.
If the current screen is the first for this monitoring cycle, the current screen is displayed again.
5. **+** selects the next sequential screen of the current report.
If the current screen is the last of the current report, the first screen of the next report is displayed. The DUE or DÜ1 key can also be pressed instead of the + key.
6. **-** selects the preceding screen of the current report. If the current screen is the first of this report, the first screen of the preceding report is displayed.
7. **>** selects the overflow screen for the current report. Overflow reports can exist for the DEVICE DISK reports if it is not possible for all the monitored objects to fit on the same screen.
If there is no overflow screen, > has the same effect as +.
8. **<** selects the previous overflow screen.
If there is no overflow screen, < has the same effect as -.
9. **<<** selects the first screen of a report with overflow screens.
If there is no overflow screen, << has the same effect as -.
10. **+N** selects the first report of the next host.
If the current screen is the last one of this monitoring cycle, the first screen of the first report from the next monitoring cycle is displayed.
In this case, the reports with network-specific data are displayed first, provided such reports were selected.
11. **-N** selects the first report of the preceding host.
If the current report is the first one of the first host, the first report with network-specific data is displayed, provided such a report was selected.

This method is useful for checking that suitable output and control statements have been entered for a specific task or when a specific report is to be displayed for extended investigation at the terminal.

Starting report output

Report output is started with the START or RESTART statement. In the case of the START statement, new monitoring data is obtained from the central buffer of SM2 and then output. In the case of the RESTART statement, reports are output with the data of the last START statement.

SM2 enters the monitoring data in output forms called **reports**. Their format is fixed and cannot be modified by the user. The user uses statements to specify whether the selected reports are output at the terminal and/or SYSLST.

The reports contain factually related data (see the report types in [chapter "SM2 screen output" on page 209](#)).

Monitoring then proceeds cycle after cycle. The reports are output either automatically or upon user request.

Input during monitoring

The automatic output mode can be interrupted with the BREAK function (K2 key). To allow interruption, the time interval between two consecutive reports is at least 2 seconds.

The /INFORM-PROGRAM command permits the user to return from system mode to SM2 program mode. All output statements can be entered as at program start time. Monitoring is continued in the meantime. The output cycle begins when the START statement is entered. The data monitored during the last completed monitoring cycle is again output following the /INFORM-PROGRAM command.

In controlled mode, any output statement can be entered at any time instead of a scroll statement without monitoring being interrupted. After the START statement, the reports for the last completed monitoring cycle can be requested.

Relationship between the monitoring cycle and the output cycle

The first monitoring cycle begins when the SM2 monitoring task is initiated. In each cycle, all of the required data is collected, edited and entered in a central buffer at the end of the cycle. The SM2 user tasks fetch the output data from this buffer.

The output cycle of an SM2 user task begins when the START statement is entered. The data collected during the last completed monitoring cycle is then fetched and output.

This means that the first reports can be output at the earliest one monitoring cycle length after SM2 monitoring task initiation.

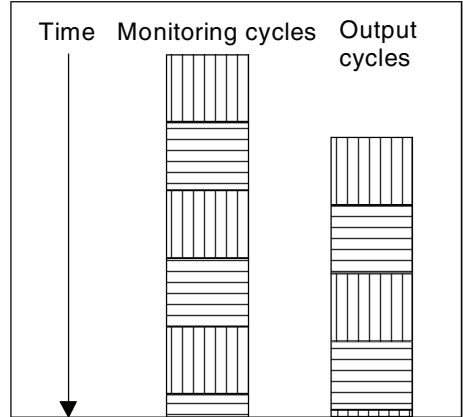
Depending on the output mode and the length of the two cycles, the following situations are possible:

A Automatic output mode

A1 Monitoring cycle and output cycle have the same length


Monitored data acquisition and report output take place at intervals of the same length

If the system load is high, the time required for output to the terminal may be considerably increased; this may result in the data of a particular monitoring cycle being skipped during output.




A2 Monitoring cycle longer than output cycle

After all reports for a monitoring cycle have been output, SM2 waits until data of the next monitoring cycle becomes available. This causes the output interval to be extended.



A3 Monitoring cycle shorter than output cycle

Monitored data is lost because the central buffer of the monitoring task is cleared faster than the data is output during the output cycle.




B Controlled mode

If the user does not succeed in having all the reports of a monitoring cycle output before the central buffer is filled again by the monitoring tasks, some monitoring cycles will be lost. When a report is requested again, the first report of the new monitoring cycle is output.

Note

The ACTIVITY, CHANNEL, and MEMORY reports are requested in this example.

The data of the monitoring cycles marked with * was overwritten in the central buffer before the user requested all the reports of the monitoring cycle shown first.



Statement for outputting information

The STATUS statement outputs the SM2 MEASUREMENT STATUS screen.

Statement	Function
STATUS	Output the monitoring status

General statements

The following statements control the SM2 run of the user (BREAK, END) and provide assistance as required (HELP, REMARK).

Statement	Function
BREAK	Change to system mode
HELP	Request user help information
REMARK	Insert remarks
END	Terminate monitoring

ISP syntax description

The statements issued to SM2 by nonprivileged users comprise the statement names and, in some statements, additional operands. The statement names and operands must be separated by at least one blank. Depending on the statement concerned, either positional or keyword operands can be used. Keyword operands can be entered in any order. If a statement only contains keyword operands, the comma shown in front of the first operand in the syntax is omitted. In the statement descriptions, certain metacharacters are used. These are described in the table below:

Format notation	Meaning	Example
UPPERCASE	Uppercase letters indicate constants which must be entered by the user in precisely this form.	NAME=
lowercase	Lowercase letters indicate variables for which the user must substitute appropriate values upon input, i.e. their contents can vary from case to case.	NAME=filename
{ }	Braces are used to indicate alternatives, i.e. one entry can or must be selected from the enclosed values.	$\left. \begin{array}{l} \text{TERMINAL} \\ \text{TER} \end{array} \right\}$
[]	Square brackets indicate that the enclosed entries are optional, i.e. may be omitted. (Parentheses must be entered.)	[,STATS= $\left. \begin{array}{l} \text{ON} \\ \text{OFF} \end{array} \right\}$]
<u>Underscoring</u>	Underscoring (underlining) indicates the default value, i.e. the value used by the system when none has been explicitly specified.	DISP= $\left. \begin{array}{l} \text{ON} \\ \underline{\text{OFF}} \\ \text{NO} \end{array} \right\}$
...	Dots indicate repetition, i.e. the preceding syntactical unit may be repeated several times in succession.	$\left[\begin{array}{l} + \\ - \end{array} \right\} \text{chn-no} \dots$

Length of the statements

A statement issued to SM2 can extend over more than one line. Continuation lines must be indicated as such by means of a hyphen (continuation character).

When statements are entered via a **terminal**, the continuation character must be the last character in the line, i.e. the character immediately preceding the ETX character. When SM2 recognizes a continuation character during interactive data input, it requests the next input line using the ENTER command.

When statements are entered from a **file**, the continuation character must be in column 72 of the input line. Characters in column 73 and beyond are ignored in all lines.

The maximum permissible length of an input line is 256 characters. The maximum permissible length of a statement is 512 characters, a string of blanks being counted as a single blank. SM2 acknowledges the entry of a line or statement that is too long with the message LINE TRUNCATED or COMMAND OVERLENGTH.

The following description of the statements for nonprivileged users is arranged in alphabetical order.

BREAK

Switch to system mode

This statement is used to interrupt SM2 processing and switch to the system's command mode.

Format

Operation	Operands
BREAK	

CHANGE-ISAM-STATISTICS Include and exclude ISAM pools

Every SM2 user can use the CHANGE-ISAM-STATISTICS statement to select one or more ISAM pools whose monitored data is to be output or no longer output in the next monitoring interval. The user can only include ISAM pools for monitoring provided the SM2 administrator has permitted the user-specific monitoring program ISAM. If this is the case, the entry ISAM appears in the USER MEASUREMENTS ALLOWED line of the SM2 MEASUREMENT STATUS screen.

If ISAM pool monitoring has been permitted, each user can activate or deactivate global ISAM pools and all task-local ISAM pools which have been or will be set up under his/her user ID.

SM2 supports the simultaneous monitoring of up to 16 ISAM pools for the set of all users.

Indirect selection of a pool by specifying a file contained in the pool is reserved for the SM2 administrator.

Format

Operation	Operand
CHANGE-ISAM-STATISTICS	<pre> ADD-POOL = { *NONE (*POOL(POOL-NAME=poolname1 [,SCOPE = { *HOST-SYSTEM *TASK(TSN=tsn) }] [,CAT-ID = { *HOME catid }])[*POOL()]...) } ,REMOVE-POOL = { *NONE (*POOL(POOL-NAME=poolname1 [,SCOPE = { *HOST-SYSTEM *TASK(TSN=tsn) }] ,CAT-ID = { *HOME catid })[*POOL()]...) *OWN } </pre>

Operands

ADD-POOL	Defines which ISAM pools are to be included in the monitoring process.
= <u>NONE</u>	No ISAM pools are included in the monitoring process.
=*POOL()	Defines which ISAM pools are to be monitored. Up to 16 ISAM pools can be defined in one list.
POOL-NAME	
=poolname1	Defines a pool name up to 8 characters in length.
SCOPE	Defines the type of ISAM pool.
= <u>HOST-SYSTEM</u>	The pool is a global ISAM pool.
=*TASK()	The pool is a task-local ISAM pool. The TSN operand defines the task under which the ISAM pool has been set up.
TSN=tsn	TSN of the task under which the task-local ISAM pool has been set up.
CAT-ID	Catalog ID of the pubset to which the ISAM pool has been assigned. This specification is part of the identification of the ISAM pool (in addition to the specifications for POOL-NAME and SCOPE).
= <u>HOME</u>	The catalog ID is the one of the home pubset.
=catid	Catalog ID of the pubset to which the ISAM pool has been assigned.
REMOVE-POOL	Defines which ISAM pools are to be excluded from the monitoring process.
= <u>NONE</u>	No ISAM pools are to be excluded from the monitoring process.
=*OWN	Only those ISAM pools included by a user in the monitoring process are excluded.
=*POOL()	See the description of ADD-POOL=*POOL()

The *POOL structure name can be omitted, but if so it must not occur anywhere else in the statement.

END
Terminate monitoring

This statement is used to terminate the SM2 run for the user.

Format

Operation	Operands
END	

FILE

Monitoring files

This statement is used to select one file which monitored data is to be output in the next monitoring cycle or is no longer to be output.

The user may only include files in the monitoring process if the FILE monitoring program is permitted by the SM2 administrator. If this is the case, the entry FILE appears in the USER MEASUREMENTS ALLOWED line of the SM2 MEASUREMENT STATUS screen.

Provided file monitoring is allowed, each user can include and exclude files which have been set up under his user ID in the monitoring process. It is also possible to include and exclude files which do not exist.

SM2 supports the simultaneous monitoring of up to 32 files for all users.

Format

Operation	Operands
FILE	NAME=filename ,STATS= <u>ON</u> /OFF

Operands

NAME

=filename Specifies the name of the file whose access values are to be monitored or to be excluded from further monitoring.
As a general rule, fully qualified file names must be specified, e.g. :A:\$USERID.FILE1 (see the "DMS Macros" manual [2]).

STATS

=ON Includes the specified file in the monitoring process if this has not already been done.

=OFF Excludes the specified file from the monitoring process.

HELP**Request user help information**

This statement is used to request an overview of all SM2 statements for nonprivileged users.

Format

Operation	Operands
HELP	

Output includes the statement name followed by a brief explanation of the statement.

OUTPUT

Define output mode

This statement is used to:

- define the device/file (terminal and/or SYSLST) to which the reports are output, and
- define the mode of output.

Reports are output either automatically at regular intervals (automatic output mode) or at the request of the user (controlled mode).

Format

Operation	Operands
OUTPUT	$\left[\left\{ \begin{array}{l} \text{TERMINAL} \\ \text{TER} \end{array} \right\} = \left\{ \begin{array}{l} \text{Y[ES]} \\ \text{N[O]} \end{array} \right\} \right]$ $\left[\left\{ \begin{array}{l} \text{HARDCOPY} \\ \text{HC} \end{array} \right\} = \left\{ \begin{array}{l} \text{Y[ES]} \\ \text{N[O]} \end{array} \right\} \right]$ $\left[\left\{ \begin{array}{l} \text{INTERVAL} \\ \text{INT} \end{array} \right\} = \left\{ \begin{array}{l} [\text{number}] \\ \text{C[ONTROLLED]} \end{array} \right\} \right]$

Operands

TERMINAL

- =YES Directs output of the reports to the terminal.
- =NO Deactivates report output to the terminal.

HARDCOPY

- =YES Directs output of the reports to the system file SYSLST. If SM2 runs in batch mode, the reports are always output to SYSLST.
- =NO Prevents output of the report to the system file SYSLST.

INTERVAL	Specifies the output mode.
=number	<p>Specifies automatic output mode (see page 165).</p> <p>“number” defines the output cycle during which the reports are to be output at regular intervals.</p> <p>This can be an integer between 10 and 3600 seconds.</p> <p>The INTERVAL=number specification is not mandatory. If it is omitted, SM2 inserts the value for the output cycle by using the value of the online monitoring cycle (default). An exception to this is the first output cycle, in which the reports are output in direct succession. With a normal workload, the system requires approximately one second to output a report. The value for the output cycle should be at least as large as the sum of the output times of the selected reports.</p>
=CONTROLLED	Each report is output only at the request of the user (controlled mode). The reports of the last completely executed cycle are available for output. Each report can be requested using a paging statement (see page 164) once SM2 displays the ENTER COMMAND message. In place of the paging statements, any other statement can also be used to request the report. The statement OUTPUT INTERVAL=number can be used to switch to automatic output mode.



Before changing the output cycle, please refer to [“Relationship between the monitoring cycle and the output cycle”](#) on [page 165](#).

Example 1

```
OUTPUT INT=C
```

SM2 output is prompted via the terminal. Following every output operation, the user is requested to enter a command. SM2 then executes the statement. The user also has the option of controlling report output.

Example 2

```
OUTPUT TER=NO,HC=YES
```

SM2 output is to the SYSLST system file. There is no output to the terminal. Output takes place in automatic mode. The analysis period is equal to the monitoring cycle.

REMARK

Insert remarks

This statement is used to insert remarks in order to document procedures.

Format

Operation	Operands
REMARK	[any text]

Operands

any text

Any text can be inserted; the text can extend beyond more than one line. Continuation lines must be marked by means of a hyphen (continuation symbol). See also [“Length of the statements” on page 170](#).

REPORT
Select reports

This statement is used to select the reports to be output.

Format

Operation	Operands
REPORT	$\left\{ \begin{array}{l} \text{STD} \\ \text{ALL} \\ \left[\left\{ \begin{array}{l} + \\ - \end{array} \right\} \text{report-type} \right] \end{array} \right\} \left[\left\{ \begin{array}{l} + \\ - \end{array} \right\} \text{report-type} \right]$

Operands

STD

Selects the following reports:

- ACTIVITY report
- MEMORY report
- CHANNEL report
- DEVICE DISK report

for output in the same order as shown above.

This operand must be specified only if the default value has been changed and is to be reactivated.

ALL

All reports are output.

If a monitoring program has not been activated or is not supplying any data at the moment, an appropriate message is output.

[{ + }
 { - } report-type]...

Updates the list of reports to be output. “report-type” designates the output reports, the identifiers of which are as follows:

Report ID	Guaranteed abbrev.	Report name
ACF		ACF
ACTIVITY	A, ACT	ACTIVITY
BCAM		BCAM CONNECTION
BCAM_MEMORY	BCAMM	BCAM MEMORY
CATEGORY	CAT	CATEGORY
CATQ		CATEGORY QUEUE
CATW		CATEGORY WSET
CHANNEL	C, CHA	CHANNEL
CMS		CMS
CPU		CPU
DAB		DAB
DEVICE_DISK	D, DDISK	DEVICE DISK
DEVICE_OTHER	DOTH	DEVICE OTHER
DEVICE_TAPE	DTP	DEVICE TAPE
DEVICE_TD	DTD	DEVICE TD
DISK_FILE	DFILE	DISK FILE ¹
DLM		DLM
FILE	F, FIL	FILE ²
GLOBAL	G	GLOBAL
GS		GS
GSVOL		GSVOL
ISAM		ISAM ²
ISAM_FILE		ISAM FILE
MEMORY	M, MEM	MEMORY
MSCF		MSCF
NSM		NSM
OPENFT	OFT	OPENFT
PCS		PCS

Table 7: Report identifiers, abbreviations and names

(part 1 of 2)

Report ID	Guaranteed abbrev.	Report name
PERIODIC	PER	PERIODIC TASK ¹
PFA		PFA
POSIX	POS	POSIX
PUBSET	PUB	PUBSET
RESPONSE	R, RES	RESPONSETIME
SESAM_SQL	SESAM	SESAM SQL
SHARED_PUBSET	SHA	SHARED PUBSET
SVC		SVC
SYMMETRIX	SYM	SYMMETRIX
TCP_IP	TCP	TCP/IP
TLM		TLM
UDS_SQL	UDS	UDS SQL
UTM		UTM
VM		VM2000
VM_CPU_POOL	VMCP	VM CPU Pool
VM_GROUP	VMG	VM GROUP

Table 7: Report identifiers, abbreviations and names

(part 2 of 2)

¹ For nonprivileged users, the report only shows them the tasks or files of their own user IDs.

² For nonprivileged users, only the user-specific reports are output with the monitoring objects of the user.

The user can specify both an individual report identifier with or without prefix and multiple report identifiers (with prefix) one after the other. The sign preceding the first report type of a line is not mandatory. If omitted, the newly specified report types of this line completely replace the ones previously set.

If the first report type of a line is specified using a prefix, the reports specified in this line are added to (+) or removed from (-) the previous ones.

All further reports of an input line must be prefixed by a sign.

Example 1

```
REPORT +DAB+CMS
```

The reports for DAB and CMS are added to the reports already activated.

Example 2

```
REPORT ALL-DAB
```

All existing reports are activated, with the exception of the DAB report.

RESTART

Start output of selected reports

This statement is used to start the output of reports with the data of the last START statement.

Thus, unlike the START statement, it does not retrieve new monitored data from the central buffer of SM2. The advantage of this for SM2 users is that new control and output statements can take effect before the next monitoring cycle.

It only makes sense to use the RESTART statement in controlled mode.

Format

Operation	Operands
RESTART	

SELECT-CHANNEL-PARAMETERS

Specify sort criterion for CHANNEL report

Every SM2 user can use the SELECT-CHANNEL-PARAMETERS statement to specify a sort criterion for the CHANNEL report.

If the statement is not entered, sorting is carried out on the basis of BUSY.

Format

Operation	Operands
SELECT-CHANNEL-PARAMETERS	SORT = <u>*BUSY</u> / *IO

Operands

SORT Specifies the sort criterion for the CHANNEL report.

=*BUSY The channels are sorted on the basis of BUSY.

=*IO The channels are sorted on the basis of the number of I/Os.



BUSY is always zero in the case of bus peripherals. In other words, if BUSY is used for sorting, these channels are always at the end of the CHANNEL report. For channels of type FC see hints in the [section "CHANNEL report" on page 235](#).

SELECT-CMS-PUBSET

Select pubsets/private disks

Every SM2 user can use the SELECT-CMS-PUBSET statement to select pubsets on the basis of their pubset ID or the set of all private disks for which a CMS report is to be output.

If the statement is not entered, all private disks are output first, followed by all other pubsets (corresponds to *ALL).

Format

Operation	Operands
SELECT-CMS-PUBSET	$\text{PUBSET-ID} = \left\{ \begin{array}{l} \text{cat-id} \\ \text{(cat-id1[,cat-id2][, *PRIVATE].)} \\ \underline{*ALL} \end{array} \right\}$

Operands

PUBSET-ID

- =cat-id Specifies the pubsets for which CMS reports are to be output on the basis of the pubset ID. A maximum of 16 pubsets can be selected.
- =*PRIVATE All private disks are output.
- =*ALL All pubsets and all private disks are output.

SELECT-DAB-CACHE

Select DAB caches

This statement is used to select one or more DAB caches to be displayed in the DAB report. For each selected DAB cache, a DAB-CACHE report is then output containing monitored data on the subareas or files.

Format

Operation	Operands
SELECT-DAB-CACHE	$[\text{ADD-CACHE-ID} = \left\{ \begin{array}{l} \text{cache-id} \\ (\text{cache-id1}[\text{,cache-id2}]...) \end{array} \right\}]$ $[\text{REMOVE-CACHE-ID} = \left\{ \begin{array}{l} \left\{ \begin{array}{l} \text{cache-id} \\ (\text{cache-id1}[\text{,cache-id2}]...) \end{array} \right\} \\ \text{*ALL} \end{array} \right\}]$

Operands

ADD-CACHE-ID

=cache-id

Specifies up to 16 names of DAB caches. The monitored data of the corresponding caches is output starting with the next monitoring cycle. "cache-id" is the name of a DAB cache and can be up to 32 characters in length.

REMOVE-CACHE-ID

=cache-id

Specifies up to 16 names of DAB caches in a list. The monitored data of the corresponding DAB caches is no longer output starting with the next monitoring cycle. "cache-id" is the name of a DAB cache and can be up to 32 characters in length.

=*ALL

Deletes all previously specified DAB caches from the name list, i.e. their monitored data is no longer output starting with the next monitoring cycle.

SELECT-DEVICE-DISK-PARAMETERS

Specify sort criterion for DEVICE DISK report

Every SM2 user can use the SELECT-DEVICE-DISK-PARAMETERS statement to specify a sort criterion for the DEVICE DISK report.

If the statement is not entered, I/O is the sort criterion.

Format

Operation	Operands
SELECT-DEVICE-DISK-PARAMETERS	SORT = *BUSY / *QLENGTH / * <u>IO</u> / *SWTIME / *RSC

Operands

SORT	Specifies the sort criterion for the DEVICE DISK report.
= *BUSY	The disk devices are sorted on the basis of BUSY.
= *QLENGTH	The disk devices are sorted on the basis of QUEUE length.
= * <u>IO</u>	The disk devices are sorted on the basis of the number of I/Os.
= *SWTIME	The disk devices are sorted on the basis of the software service time.
= *RSC	The disk devices are sorted on the basis of the average number of RSC IOs processed in parallel.

Operands

OUTPUT-INFORMATION

Selects user IDs or job names whose data is to be output on the screen.

=USER-ID(...) Specifies the tasks to be output via the user ID.

*ALL:

All tasks are output. A user ID can only be output if it exists.

userid:

The tasks with the specified user IDs (up to 16) are output. User IDs must be specified without the \$ sign.

=JOB-NAME(...) Specifies the tasks to be output via the job names.

*ALL:

All tasks are output. The job name can only be output if it exists.

jobname:

The tasks with the specified job names (up to 16) are output.

SORT

Specifies sort criteria. A PERIODIC TASK report is created for each sort criterion specified.

=SERVICE-UNITS, CPU, IO, UPG, PAGING-READ

The tasks are sorted on the basis of the specified criteria. If there are several sort criteria, the variously sorted reports are output one after the other.

If no task is found with the selected output information, the following message appears:

NO PERIODIC TASK DATA FOR SELECTED <output-information>.

If no task is found with the specified sort criterion, *NONE is output in the first line under TSN.

SELECT-SYMMETRIX-PUBSET

Select pubsets for SYMMETRIX DEVICE report

Every SM2 user can use the SELECT-SYMMETRIX-PUBSET statement to select up to 16 pubsets that are displayed in the SYMMETRIX PUBSET report. A SYMMETRIX DEVICE report is then output for each selected pubset.

Format

Operation	Operands
SELECT-SYMMETRIX-PUBSET	$\text{PUBSET-ID} = \left\{ \begin{array}{l} \left\{ \begin{array}{l} \text{cat-id} \\ (\text{cat-id1}[\text{,cat-id2}][\text{,*NOID}]..) \end{array} \right\} \\ \text{*NONE} \\ \text{*NOID} \\ \text{*ALL} \end{array} \right\}$

Operands

PUBSET-ID

- =cat-id Specifies the pubsets for which SYMMETRIX DEVICE reports are to be output by means of the pubset ID. A maximum of 16 pubsets can be selected.
- =*NOID A SYMMETRIX DEVICE report is output with all devices that are not assigned to a pubset.
- =*NONE The SYMMETRIX DEVICE report is not to be output.
- =*ALL SYMMETRIX DEVICE reports are to be output for all pubsets.

SELECT-SYMMETRIX-UNIT**Select Symmetrix systems for SYMMETRIX PUBSET report**

This statement is used to select up to 16 Symmetrix systems to be displayed in the SYMMETRIX report.

For each selected Symmetrix system, a SYMMETRIX PUBSET report is then output containing monitored data on the individual pubsets supported by the system.

Format

Operation	Operanden
SELECT-SYMMETRIX-UNIT	$\text{SERIAL-NUMBER} = \left\{ \begin{array}{l} \left\{ \begin{array}{l} \text{number} \\ (\text{number1}[\text{,number2}]..) \end{array} \right\} \\ \text{*NONE} \\ \text{*ALL} \end{array} \right\}$

Operands**SERIAL-NUMBER**

- =number Specifies the Symmetrix systems for which SYMMETRIX PUBSET reports are to be output by means of the serial number of the Symmetrix system.
The serial number can be entered in the length 1 to 12 (must be unambiguous).
Up to 16 Symmetrix system serial number can be specified.
- =*NONE The SYMMETRIX PUBSET report is not to be output.
- =*ALL SYMMETRIX PUBSET reports are to be output for all Symmetrix systems.

SELECT-UTM-APPLICATION

Select UTM applications

This statement is used to select up to 16 UTM applications to be displayed in the UTM report. For each selected UTM application, a UTM APPLICATION report is then output.

Format

Operation	Operands
SELECT-UTM-APPLICATION	<pre>[ADD-APPLICATION = { name (name1[,name2]...) }]</pre> <pre>[,REMOVE-APPLICATION = { { name (name1[,name2]...) } } *ALL]</pre>

Operands

ADD-APPLICATION

=name Specifies up to 16 UTM applications. "name" is the name of a UTM application and can be up to 8 characters in length.

REMOVE-APPLICATION

=name Specifies up to 16 UTM applications in a list. "name" is the name of a UTM application and can be up to 8 characters in length.

=*ALL All the UTM application names specified previously are deleted from the name list, i.e. the monitored data for these applications will no longer be output as of the next monitoring cycle.

SHOW-USER-MEASURED-OBJECTS**Output monitored objects**

This statement is used to show an SM2 user the objects he/she is currently monitoring.

Format

Operation	Operands
SHOW-USER-MEASURED-OBJECTS	TYPE = TASK

Operands

TYPE

=TASK

Outputs the user's tasks currently being monitored by the user-specific task monitoring program.

START**Start report output**

This statement is used to start report output.

Format

Operation	Operands
START	

START-ISAM-STATISTICS

Monitor ISAM pools

This statement is used to select one or more ISAM pools whose monitored data is to be output in the next monitoring cycle.

The user can only include ISAM pools for monitoring provided the SM2 administrator has permitted the user-specific monitoring program ISAM. If this is the case, the entry ISAM appears in the USER MEASUREMENTS ALLOWED line of the SM2 MEASUREMENT STATUS screen.

If ISAM pool monitoring has been permitted, each user can activate or deactivate global ISAM pools and all task-local ISAM pools which have been or will be set up under his/her user ID.

SM2 supports the simultaneous monitoring of as many as 16 ISAM pools for all nonprivileged users.

Indirect selection of a pool by specifying a file contained in the pool is reserved for the SM2 administrator.

Only the input of the last START-ISAM-STATISTIC statement applies here. If you wish to monitor several ISAM pools, these must be specified in a START-ISAM-STATISTIC statement.

Format

Operation	Operands
START-ISAM-STATISTICS	$ \text{POOL} = \left\{ (*\text{POOL} \left\{ \begin{array}{l} \text{POOL-NAME}=\text{poolname1} \\ \text{[SCOPE}=\left\{ \begin{array}{l} *HOST-SYSTEM \\ *TASK(TSN=\text{tsn}) \end{array} \right\} \right\} \\ \text{[CAT-ID}=\left\{ \begin{array}{l} *HOME \\ \text{catid} \end{array} \right\} \right\} \end{array} \right\} \right\} \text{[, *POOL()]} \dots \right\} $

Operands

POOL

=(*POOL(...)) Defines which ISAM pools are to be monitored. Up to 16 ISAM pools can be defined in one list.

POOL-NAME

=poolname1 Defines a pool name of up to 8 characters in length.

SCOPE

Defines the type of ISAM pool.

=*HOST-SYSTEM The pool is a global ISAM pool.

=*TASK(...) The pool is a task-local ISAM pool. The TSN operand defines the task to which the ISAM pool belongs.

TSN=

TSN of the task to which a task-local ISAM pool belongs.

CAT-ID

Catalog ID of the pubset to which the ISAM pool is assigned. The specification is part of the identification of an ISAM pool (in addition to the POOL-NAME and SCOPE specifications).

=*HOME

The ID is the catalog ID of the home pubset.

=catid

Catalog ID of the pubset to which the ISAM pool has been assigned.

The *POOL structure name can be omitted, but if so it must not occur anywhere else in the statement.

STATUS

Output monitoring status

This statement is used to output the MEASUREMENT STATUS screen.

Format

Operation	Operands
STATUS	

STOP-ISAM-STATISTICS

Terminate monitoring of ISAM pools

This statement is used to terminate monitoring of the ISAM pools previously included in the monitoring process.

Format

Operation	Operands
STOP-ISAM-STATISTICS	



If the SM2 administrator has suppressed the user-specific monitoring program ISAM, all ISAM pools included for output are automatically excluded.

5.5 BS2000 commands for activating user task monitoring

Any user can have SM2 monitor tasks under his/her own ID, provided the SM2 primary administrator has given permission for the user-specific monitoring program (by issuing the SM2 statement `MODIFY-USER-ADMISSION TASK=*ALLOW`. The total number of tasks which can be monitored simultaneously is limited to 16).

Only users who have been assigned the system privilege `SW-MONITOR-ADMINISTRATION` can start or terminate monitoring of any tasks.

Function

The BS2000 command `/START-TASK-MEASUREMENT` is used to specify the tasks to be included in the monitoring process. The SM2 monitoring program then records task-specific monitored data and enters this data in a user-specific file.

In addition to task-specific monitoring data, program counter statistics and SVC statistics on program runs can also be requested.

In the case of program counter statistics for counting commands, the program is interrupted at defined intervals (using a timer) and the address of the next command to be executed and contained in the program counter is transferred to the user-specific SM2 output file (registering the number of times the counter sampled the program).

In the case of SVC statistics, all SVC program macros called during task monitoring are recorded, and SVC numbers and macro addresses are also stored in the SM2 output file.

If the BS2000 command `/STOP-TASK-MEASUREMENT` is issued or the task is terminated, user task monitoring ends and the user-specific SM2 output file is closed. In addition, withdrawal of the primary SM2 administrator's permission for the user-specific monitoring program `TASK (MODIFY-USER-ADMISSION TASK=*INHIBIT)` as well as deactivation of SM2 (`STOP-SUBSYSTEM SUBSYSTEM-NAME=SM2`) will terminate any currently executing task measurement operations.

The product `SM2-PA` is provided to analyze the SM2 output file generated for user task monitoring (for further information, please refer to the "SM2-PA" manual [15]).



The user-specific SM2 output file can exceed the limit specified for the user by `PUBLIC-SPACE-LIMIT` in the `ADD-USER-` or `MODIFY-USER-ATTRIBUTES` command, even if `PUBLIC-SPACE-EXCESS=*NO` is specified.

Formats

START-TASK-MEASUREMENT

TSN = *OWN / <alphanum-name 1..4>
,FILE = *STD / ***BY-LINK-NAME**
,PCOUNTER-INTERVAL = *NONE / <integer 1..10000>
,SVC-STATISTICS = *OFF / ***ON**
,LOAD-INFO = *STD / ***DETAILED**

STOP-TASK-MEASUREMENT

TSN = *OWN / <alphanum-name 1..4>

Operands

TSN =

Specifies the task to be monitored or for which monitoring is to be terminated.

TSN = *OWN

The task invoking the command will be monitored or monitoring of that task will be terminated.

TSN = <alphanum-name 1..4>

The task with the TSN specified is monitored or monitoring of that task is terminated. Only users who have been assigned the system privilege SW-MONITOR-ADMINISTRATION can start or terminate monitoring of any tasks. All other users can only monitor their own tasks.

Restriction: The task SM2U, which processes the user-specific SM2 output files, cannot be monitored.

FILE =

Specifies the user-specific SM2 output file.

FILE = *STD

The user-specific SM2 output file contains the file attributes which are defined in the following table.

FILE = *BY-LINK-NAME

The user defines the file name and further file attributes using the corresponding BS2000 commands. The file is assigned via the link name. The link name for the user-specific SM2 output file is PALINK.

PCOUNTER-INTERVAL =

Defines the sampling cycle for command counter statistics (in CPU milliseconds).

PCOUNTER-INTERVAL = *NONE

Command counter statistics are not activated.

PCOUNTER-INTERVAL = <integer 1..10000>

Sampling cycle in CPU milliseconds. The only samples recorded are those with the status TU.

SVC-STATISTICS =

The SVC macros of the task to be measured are recorded and written to the user-specific SM2 output file.

SVC-STATISTICS = *OFF

SVC statistics are not activated.

SVC-STATISTICS = *ON

SVC statistics are activated. The only SVCs recorded are those with the status TU.

LOAD-INFO =

Specifies the point at which the module loading information is to be recorded.

LOAD-INFO = *STD

The module loading information is only recorded when the program is started or terminated (or when monitoring is started or terminated).

LOAD-INFO = *DETAILED

The module loading information is recorded when the program is started (or when monitoring is started) and every time the task performs a load or unload operation.

Command return codes

(SC2)	SC1	Maincode	Meaning
	0	CMD0001	No errors
	32	NPS0050	System error in SM2 modules, command rejected
	64	NPS0044	No authorization for monitoring task, command rejected
	64	NPS0045	Task already monitored by SM2
	64	NPS0046	Task monitoring not started by /START-TASK-MEASUREMENT, command rejected
	64	NPS0047	Specified task does not exist
	64	NPS0051	User cannot monitor specified task
	64	NPS0065	Cannot monitor any more tasks
	64	NPS0066	Invalid file attributes for user-specific SM2 output file
	64	NPS0067	DMS code '(&00)' in the case of a macro '(&01)' for user-specific SM2 output file



You will find a general description of the command return codes in the “Commands” manual [3].

Attributes of the user-specific SM2 output file

Meaning	Attribute	Attribute
Operand for file creation	FILE =*STD	FILE =*BY-LINK-NAME
Definition of file attributes	Predefined	Optional
File name	\$userid.SM2. TASKSTATISTIK.nnnn ¹	filename
Block length	BUFFER-LENGTH = *STD(SIZE=16)	BUFFER-LENGTH = *STD(SIZE=number) ²
Storage space allocation	SPACE = *RELATIVE (PRIMARY-ALLOCATION=48, SECONDARY- ALLOCATION=48)	SPACE = *RELATIVE (PRIMARY- ALLOCATION=number1, SECONDARY-ALLOCATION= number2) ^{2, 3}
Access method	SAM	No selection option, SAM format is mandatory.
Open mode	EXTEND	No selection option, the existing file is always extended
Format of the data blocks	BLOCK-CONTROL-INFO is defined by the corresponding system parameter	BLOCK-CONTROL-INFO = *PAMKEY/ *WITHIN-DATA-BLOCK/ *NO ⁴

Table 8: Attributes of the user-specific SM2 output file

¹ nnnn = TSN of the task to be monitored

² In the case of an SM2 output file assigned using link names, the BUFFER-LENGTH and SPACE attributes are taken from the standard file, unless specified otherwise by the user.

³ PRIMARY-ALLOCATION and SECONDARY-ALLOCATION each 48 unless defined otherwise by the user.

⁴ The BLOCK-CONTROL-INFO=*NO operand is treated like BLOCK-CONTROL-INFO=*WITHIN-DATA-BLOCK, because only SAM format is permitted for the file.



To generate module-related program counter and SVC statistics, the load information (the load address and length) of each module is required for each module which has been linked into the monitored program or is to be loaded by the monitored program. When the program or measurement is terminated by means of the /STOP-TASK-MEASUREMENT command, SM2 determines this module information and writes it to the SM2 output file.

Under certain circumstances SM2 can only supply incomplete information or none at all:

- If the program to be monitored uses overlay techniques, the module loading information should be recorded using LOAD-INFO=DETAILED. If this is not done, the samples and SVCs for all overlaid modules are assigned to the module loaded at the time monitoring is terminated.
- No information will be supplied for programs which are linked and loaded without LSD (List for Symbolic Debugging).
- No information will be supplied for code sections which were not generated with the aid of load events (e.g. if the code was put in memory dynamically requested by means of MOVE commands).

5.6 Writing to the SM2 output file

Introduction

In addition to report output at the terminal and on SYSLST, consecutive output of monitored data to the SM2 output file is provided for.

At file creation time, a system task is generated for writing to the SM2 output file. This system task terminates when the SM2 output file is closed.

Use of the SM2 output file is recommended

- when system utilization is to be monitored for an extended period of time (standard case, trend monitoring)
- for bottleneck analysis (peak load analysis) if special monitored data is required in addition to the global system monitored data. The corresponding monitoring programs must be activated for this purpose.

Only the SM2 administrator can create and close the SM2 output file.

The file attributes are preset by SM2, but the SM2 administrator can change the preset default options.

The SM2R1 program is available in BS2000 and the application ANALYZER in MS-Windows for analyzing output files. The SM2R1 data interface (SM2R1 transfer file) is available to users who wish to analyze the SM2 output file using their own programs. For information on the record structure of this file, refer to [chapter “SM2R1 analysis routine” on page 391](#)).

Attributes of the SM2 output file

The SM2 administrator can choose to create an SM2 output file with the preset file attributes or to open an SM2 output file with freely selected attributes.

SM2 output file with preset file attributes

This file is opened with the statement `OPEN-LOG-FILE FILE=*STD`. SM2 opens a SAM file in output mode and catalogs it under the name `SM2.hostname.yyyy-mm-dd.sss.nn`.

Here, `hostname` indicates the assignment to a host, `yyyy.mm.dd` is the date on which the file was created, `sss` is the number of the BS2000 session and `nn` is the consecutive number of the SM2 output file within that session (counted from 1).

SM2 output file with freely selected file attributes

Such a file is created with the statement OPEN-LOG-FILE FILE=*BY-LINK-NAME after the name and attributes have been defined with the ADD-FILE-LINK command.

It is possible to define

- a PAM file in output mode with a freely selectable name, or
- a SAM file with freely selectable open mode and file name.

The following table contains the file attributes which are preset in SM2 or can be defined with the FILE command:

Meaning	Attribute	Attribute	
File creation operand	FILE=*STD	FILE=*BY-LINK-NAME	
File attribute definition	Preset	Optional	
File name	SM2.hostname. yyyy-mm-dd.sss.nn	filename ¹	
File link name		SMLINK	
Block size ²	BUFFER-LENGTH = *STD (SIZE=16)	BUFFER-LENGTH = *STD (SIZE=16)	
Storage space allocation	SPACE=*RELATIVE (PRIMARY- ALLOCATION=576, SECONDARY- ALLOCATION=576)	selectable ²	
Access method	SAM	ACCESS-METHOD = *UPAM [,SHARED-UPDATE = YES] ³	ACCESS-METHOD = *SAM
Open mode	OUTIN	OPEN-MODE = *OUTIN	OPEN-MODE = *OUTPUT / *EXTEND
Format of the data blocks	BLOCK-CONTROL- INFO = *WITHIN- DATA-BLOCK	BLOCK-CONTROL-INFO = *PAMKEY / *WITHIN-DATA-BLOCK / *NO ⁴	

Table 9: Possible attributes of the SM2 output file

¹ It is advisable to select a file name that indicates the date and time of file creation, so as to facilitate file sequencing in the input order required by SM2U1.

² Only block lengths of 16 are used. If no block length is specified, this is set to 16 by SM2. If BUFFER-LENGTH is specified without the SPACE operand, SM2 uses the default value 576 for the primary storage allocation and for the secondary storage allocation.

- ³ ACCESS-METHOD=*UPAM,SHARED-UPDATE=*YES should be selected if the open SM2 output file is to be analyzed using ANALYZER. This means that BUFFER-OUTPUT=*IMMEDIATE must be selected in the OPEN-LOG-FILE statement.
- ⁴ SM2 treats the operand BLOCK-CONTROL-INFO= *NO in the same way as BLOCK-CONTROL-INFO= *WITHIN-DATA-BLOCK.

If the SM2 output file is created on a private volume (additional operands in the CREATE-FILE command VOLUME=...,DEVICE-TYPE=...), allocation of the private volume for other systems must be authorized using /SET-DISK-PARAMETERS ... USER-ALLOCATION=*SHARE.



The high block length may result in mismatches, which in turn cause the file to become very large, particularly if BUFFER-OUTPUT=*IMMEDIATE was selected in the OPEN-LOG-FILE statement. However, this setting allows you to minimize the number of "MISSED RECORDS". The majority of mismatches can be eliminated by executing an SM2U1 run, which means the advantages of this setting far outweigh the disadvantages.

MISSED RECORDS are records that could not be written to the SM2 output file due to an input/output bottleneck and were therefore lost. They are indicated in the STATUS TABLE screen and by SM2R1 in EVALUATION STATISTICS.

Managing the SM2 output file

PAM output file

PAM output files are formatted in blocks.

The records are not separated by block boundaries.

An SM2 output file created with SHARED-UPDATE=*YES can be analyzed by ANALYZER during the monitoring process or converted to a SAM file by SM2U1 and then analyzed using SM2R1.

SAM output file

The SAM output file format is logically similar to the PAM output file format.

An advantage of SAM output files is that they need not be converted before the SM2R1 analysis routine is called (refer to [chapter “SM2U1 utility routine” on page 373](#)).

In the OUTPUT open mode, a new SAM output file can be created or a file can be replaced by another file with the same name.

In the EXTEND open mode, an existing file can be expanded to accommodate further record groups. This obviates the need to call the SM2U1 routine for merging several SM2 output files.

Opening and closing the SM2 output file

The OPEN-LOG-FILE statement opens the SM2 output file, while the CLOSE-LOG-FILE statement closes the SM2output file; likewise the SM2 output file is implicitly closed when SM2 is stopped (STOP-SUBSYSTEM SUBSYSTEM-NAME=SM2).



After a system breakdown or if the SM2 output file was not closed properly for some other reason, the file must be processed with the SM2U1 routine.

The SM2 output file can exceed the limit specified by PUBLIC-SPACE-LIMIT in the ADD-USER- or MODIFY-USER-ATTRIBUTES command, even if PUBLIC-SPACE-EXCESS=*NO has been specified.

If the SM2 output file is already open, the OPEN-LOG-FILE statement closes it and then opens a new SM2 output file. The monitoring programs SERVICETIME and TASK, which write monitored data exclusively to the SM2 output file (see [table “Overview of the SM2 monitoring programs” on page 50](#)), are not terminated in the process.

5.7 SM2 monitoring program: examples

Example 1

This simple example illustrates how SM2 is started and stopped again later on.

The following default options are used:

- monitoring cycle: 150 seconds
- sampling cycle: 800 milliseconds

Reports A (ACTIVITY), M (MEMORY), C (CHANNEL) and D (DEVICE DISK) are output at the terminal in automatic output mode. SM2 is interrupted after some time with the BREAK function (K2 key) and terminated.

```
/START-SM2
START
.           (Output of the reports)
.
.
BREAK or K2
/INFORM-PROGRAM
END
```

Example 2

In this example, SM2 is started with data output to the SM2 output file. The file is given the default name SM2.hostname.yyyy.mm.dd.sss.nn. The monitoring cycle is preset to 150 seconds and the sampling cycle to 800 milliseconds. The SYSTEM monitoring program is activated for all devices.

The SM2 user task is then terminated and monitoring continues in the background.

SM2 is called again after some time to terminate the SYSTEM monitoring program and close the SM2 output file.

```
/START-SM2
CALL-ADMINISTRATION-PART
SET-SYSTEM-PARAMETERS DEVICES=*ALL
START-MEASUREMENT-PROGRAM TYPE=*SYSTEM
OPEN-LOG-FILE FILE=*STD
END
.
.
.
.
/START-SM2
CALL-ADMINISTRATION-PART
STOP-MEASUREMENT-PROGRAM TYPE=*SYSTEM
CLOSE-LOG-FILE
END
```

Example 3

In this example, SM2 is started with data output to the SM2 output file. The file is given the default name SM2.hostname.yyyy.mm.dd.sss.nn. The monitoring cycle is set to 20 seconds and the sampling cycle to 200 ms. The BCAM-CONNECTION monitoring program is activated. Here the defaults are set for the range limits for the INWAIT, REACT, INPROC and OUTPROC times. All connections of the \$DIALOG application are monitored. All monitored data is written to the SM2 output file. At the terminal, the A (ACTIVITY), M (MEMORY), C (CHANNEL) and D (DEVICE DISK) reports are output in automatic output mode.

After some time, SM2 is interrupted with the BREAK function (K2 key) and terminated. However, the SM2 output file remains open and the BCAM-CONNECTION monitoring program continues running.

```

/START-SM2
CALL-ADMINISTRATION-PART
SET-BCAM-CONNECTION-PARAMETERS
ADD-BCAM-CONNECTION-SET           -
    SET-NAME=DIALOG,              -
        CONNECTION-SELECTION=*BY-NEA-NAME( -
            CONNECTION-NAME=*SPECIFIED( -
                LOCAL-APPLICATION=$DIALOG, -
                PARTNER-APPLICATION=*ANY)), -
        HOST-SELECTION=*ANY
START-MEASUREMENT-PROGRAM TYPE=*BCAM-CONNECTION
OPEN-LOG-FILE FILE=*STD
MODIFY-MEASUREMENT-PERIODS OFFLINE-PERIOD=20,SAMPLING-PERIOD=200
CALL-EVALUATION-PART
START
.      (Output of the Reports)
.
.
BREAK or K2
/INFORM-PROGRAM
END

```

6 SM2 screen output

The monitoring task collects monitored data during an online monitoring cycle, edits it, and enters it in the central buffer when the online monitoring cycle is completed.

From the central buffer, each SM2 user task transfers the data into its own buffer and enters it in output forms called reports.

These monitored data reports can then either be output at regular intervals throughout the output cycle (automatic output mode), or can be requested individually by the user (controlled mode). See [“Statements for controlling reports” on page 163](#).

In addition, SM2 information screens are available, some in the administration facility only and others in the analysis subinterval and the administration facility. Instead of monitored data, these contain information on the monitoring process and on the status of SM2. The following information screens can be output at the request of the user:

- MEASUREMENT STATUS
- USER MEASURED OBJECTS
- DEFINED PARAMETER
- ACTIVE PARAMETER
- STATUS TABLE
- SELECTED HOSTS

SM2 reports

The monitored data is broken down into factually related groups and entered in various reports. The following reports may be output without the SM2 administrator having to activate an optional monitoring program.

Report	Meaning
ACF report	Overview of activation control of system task management
ACTIVITY report	Overview of the overall system load
CATEGORY QUEUE report	Overview of the queue occupancy of all categories
CATEGORY WSET report	Overview of the working sets of all categories
CHANNEL report	Overview of the current channel occupancy.
CPU report	Overview of CPU utilization for the various CPUs
DEVICE DISK report	Overview of the activity of disk devices
DEVICE OTHER report	Overview of the activity of all devices not contained in the DEVICE DISK, DEVICE TAPE or DEVICE TD report (e.g. console)
DEVICE TAPE report	Overview of the activity of tape devices
DEVICE TD report	Overview of the activity of communication devices
GLOBAL report	Overview of host-specific monitored data on important system activities
MEMORY report	Overview of main memory load and virtual address space
PCS report	Overview of PCS (Performance Control Subsystem) activity (used only if PCS is installed)
SHARED PUBSET report	Overview of the activity of disk devices that belong to shared pubsets

Table 10: Reports without optional monitoring program

The following reports can only be output if the SM2 administrator has activated the corresponding optional monitoring program beforehand or authorized the appropriate user monitoring operations (MODIFY-USER-ADMISSION):

Report	Meaning
BCAM CONNECTION report	Overview of bucket-specific time statistics and information on data sent and received
BCAM MEMORY report	Overview of the maximum size, the threshold value from which BCAM issues warnings on the console, and the current size of the resident memory for data transfer.
CATEGORY report	Overview of the consumption statistics of all categories
CHANNEL report	Overview of the current channel occupancy. The CHANNEL report is a special case: The values BUSY and NON OVERLAP are output by default. The remaining monitored data is collected by the CHANNEL-IO monitoring program which is started for all channels when the SM2 subsystem is started.
CMS report	Overview of the activities of the catalog management system (CMS)
DAB report	Overview of the activities of the disk access buffer function (for a more detailed overview, see the DAB-CACHE report)
DISK FILE report	Overview of the file activities of selected disk devices
DLM report	Overview of the activities of the Distributed Lock Manager
FILE report	Overview of the access values of files previously designated
GS report	Overview of the utilization of global storage
GSVOL report	Overview of the utilization of emulated volumes in global storage (GS)
ISAM report	Overview of ISAM buffer management performance
ISAM FILE report	Overview of the monitored data for selected files in ISAM pools
MSCF report	Overview of the monitored data of the MSCF subsystem
NSM report	Overview of host-specific monitoring values of the Distributed Lock Manager
OPENFT report	Overview of the load values for one openFT instance
PERIODIC TASK report	Overview of the most important utilization data of all tasks: privileged users receive information on all tasks, while nonprivileged users receive information relating only to those tasks under their own user IDs
PFA CACHE report	Overview of the use of caches under User PFA
POSIX report	Output of the monitoring data of the POSIX subsystem

Table 11: Reports that require an optional monitoring program and/or the permission of the SM2 administrator
(part 1 of 2)

Report	Meaning
PUBSET report	Output of the occupancy of imported SF pubsets or volume sets
RESPONSETIME report	Overview of the response time behavior throughout the system
SESAM SQL report	Overview of data of the SESAM/SQL database system
SVC report	Overview of the number of SVCs called
SYMMETRIX reports	Overview of the use of Symmetrix systems (detailed information is provided by the SYMMETRIX DEVICE and SYMMETRIX PUBSET reports)
TCP/IP report	Overview of the volume of data transmitted per IP address and port number
TLM report	Overview of the Task Lock Manager calls (for privileged users only)
UDS SQL report	Overview of data of the UDS/SQL database system
UTM report	Overview of the consumption statistics for UTM applications and detailed values for the individual applications in the UTM APPLICATION report
VM report	Overview of the values of guest systems of a VM2000 system
VM CPU POOL report	Overview of the monitored data for CPU pools
VM GROUP report	Overview of the values of VM groups
X2000 report	Overview of the activity of the X2000 part of servers with x86 architecture

Table 11: Reports that require an optional monitoring program and/or the permission of the SM2 administrator
(part 2 of 2)

SM2 information screens

The information screens are special in that they contain no monitored data, but provide information on the monitoring run and on the SM2 status. These screens are only output if explicitly requested by the user.

The first two screens listed are available in both the analysis subinterval and the administration facility. The others are only available to privileged users in the administration facility.

Screen	Information
MEASUREMENT STATUS	Overview of all defined and active monitoring programs
USER MEASURED OBJECTS	Overview of user-monitored objects
DEFINED PARAMETER	Overview of the monitoring parameters and monitored objects defined for the individual monitoring programs
ACTIVE PARAMETER	Output of the monitoring parameters and monitoring objects valid for the active monitoring programs
STATUS TABLE	Status of the SM2 system tasks and of the subsystems used by SM2
SELECTED HOSTS	Overview of the hosts selected using the SELECT-HOSTS statement

Table 12: SM2 information screens

Displaying follow-up screens to information screens

Entry	Meaning
'+' or key K3 or DUE	Page forward
'-' or key K1	Page backward (if you page backward from the first page, the first page is output again)
'-.'	Page back to the first page
'END'	Terminate output of follow-on screens

All other entries have the same effect as +.

The scroll commands can also be used to output the information screens of various hosts.



It is not possible to page in batch and procedure modes. All follow-on screens are always output, but no overflow screens.

The header line of each screen contains global information, namely:

CYCLE Duration of the (online) monitoring cycle in seconds, i.e. period after which the collected values are available for output

SAMPLES Number of samples taken within the (online) monitoring cycle

These values are output if they are relevant to the current screen.

The host name is output in the top left-hand corner.

The date and time are output in the top right-hand corner of the screen. In the case of monitored data reports, this is the point at which the last (online) monitoring cycle was terminated. For all other screens, it is the current time when the screen is output.

In the screens which concern the host network (GLOBAL, NSM, SHARED-PUBSET), the header line displays the data of the first host listed in the report or of the first selected host. If the list of selected hosts contains the local host, this host is always the first one in the list. The remaining hosts appear in the order in which they were specified by the user in the SELECT-HOSTS statement.

Representation of values in SM2 output screens

The string ***** or blanks are output in place of a value if no value is available.

The string <<<<< is output if computation is not possible. The string >digit digit digit ... is output if the value does not fit into the available space.

Terminal emulations on the PC

In the case of terminal emulations on the PC (e.g. MT9750), the character set "International" must be set in the keyboard configuration in order to output the delimiter "[]."

Output sequence

Report name	Identifiers			
GLOBAL	N	O		
NSM	N	O	M	
SHARED PUBSET	N	O		
ACTIVITY				
MEMORY				
CHANNEL		O	M ¹⁾	
DEVICE DISK		Ü	M ²⁾	
DEVICE TD		Ü	M ²⁾	
Table 13: Output sequence				(part 1 of 2)
DEVICE TAPE		Ü	M ²⁾	
DEVICE OTHER		Ü	M ²⁾	
PERIODIC TASK		O	M	
PRIVILEGED FILE		O	M	
USER FILE		O		P
RESPONSETIME			M	
ACF				
CATEGORIE QUEUE				
CATEGORIE WSET				
PCS				
DAB		O	M	
DAB CACHE	F	O	M	
CMS		O	M	
UTM		O	M	
UTM APPLICATION	F	O	M	
PRIVILEGED ISAM			M	
USER ISAM				P
VM2000		O	M	
TLM			M	
CPU		O		
SVC			M	
PFA CACHE		O	M	
POSIX			M	

Report name	Identifiers			
CATEGORY			M	
MSCF			M	
DEVICE TD		O		
DLM			M	
GSVOL		O	M	
GS		O	M	
TCP IP		O	M	
BCAM CONNECTION		O	M	
DISK FILE		O	M	
SYMMETRIX		O	M	
SYMMETRIX PUBSET	F	O	M	
SYMMETRIX DEVICE	F	O	M	
DEVICE TAPE		O		
DEVICE OTHER		O		
PUBSET		O	M	
VM CPU POOL		O	M	
ISAM FILE			M	
VM GROUP		Ü	M	
BCAM MEMORY			M	
SESAM SQL		Ü	M	
UDS SQL		Ü	M	
OPENFT		Ü	M	

Table 13: Output sequence

(part 2 of 2)

- N: Network report
 F: Follow-up screen requested by means of a special statement
 O: Report containing an overflow screen for further monitored objects
 M: Monitoring program must be activated
 P: User-specific monitoring program must not be permitted for the nonprivileged user
- 1) The CHANNEL report contains information from the CHANNEL-IO monitoring program as well as standard monitored data.
 - 2) The associated SAMPLING-DEVICE monitoring program is started when the SM2 subsystem is started.

This sequence is maintained even if the user has suppressed output of some reports. The suppressed records are simply skipped without any time being lost. The same applies to output in controlled mode.

The reports of a given monitoring cycle can be requested in the above sequence. The data of the next monitoring cycle is output only after the last record of the preceding cycle has been output.

The reports are listed in alphabetical order in the following description.

6.1 SM2 reports

ACF report

This report supplies information on the activation control function (ACF) of the PRIOR task scheduler. The data can be used to assess the internal activation decisions. Data interpretation and consequential actions presuppose intimate familiarity with the activation algorithm used in the PRIOR task scheduler and should therefore be left to system specialists.

Report output

The REPORT ACF statement is used to request the output of the ACF report.

D016ZE04 SM2 ACF								<date>	<time>
COUNTERS PER MIN									
ACTIVA	FORCDA	PREEMT	INVOCL	INVOCS	MICTSR	YSYSERV	FVI		
7125.1	0.0	0.0	41.4	7125.1	14439	38.3	9125.0		
RESOURCE UTILIZATION IN %									
RESOURCE	LOW	MEDIUM	HIGH						
CPU	87.5	5.0	7.5						
MEMORY	100.0	0.0	0.0						
PAGING	100.0	0.0	0.0						

Monitoring information

COUNTERS PER MIN

The number of calls per minute is output for each of the variables listed in the table.

ACTIVA	Total number of task activations per minute
FORCDA	Number of task deactivations per minute enforced by ACF
PREEMT	Number of preemptions per minute
INVOCL	Number of "long" ACF invocations per minute
INVOCS	Number of "short" ACF invocations per minute
MICTSR	Number of micro-time-slice runouts per minute

YSERV	Number of system services runouts per minute
FVI	Number of "Fremd vor Idle" (remote before idle) accesses per minute. "Fremd vor Idle" access: a processor initiates a task from the local Q1 of another processor if its own Q1 is empty, so as not to become idle.

If all counts in a monitoring cycle are zero, SM2 outputs the message NO ACTIVATION CONTROL FUNCTION ACTIVITY IN LAST INTERVAL.

RESOURCE UTILIZATION IN %

Utilization of the resources CPU, memory and paging is classified as low, medium and high utilization. The percentages given in the table represent the share of total calls in the monitoring cycle.

ACTIVITY report

This report provides the user with an overview of the most important system activities.

Report output

The REPORT ACTIVITY statement is used to request the output of the ACTIVITY report. POSIX data can only be output if the monitoring program is activated.

When analyzing the monitoring results, it should be remembered that the SM2 monitoring task is a system task that is kept active for the duration of the whole monitoring process. Another system task is created for data entry in the SM2 output file, which is likewise kept active for the duration of the whole monitoring process; the same applies to user task monitoring.

When systems comprising a number of logical machines are being monitored, the mean values of the monitored data for all active logical machines are output.

```

D016ZE04 SM2 ACTIVITY (VM)  CYCLE:  60 S  SAMPLES:  144      <date> <time>

```

CPU UTILIZATION % (2 LM'S)				PCS DATA			MEMORY		
TU	TPR	SIH	IDLE	SRACT	RDACT	NPP	WS-ACT	WS-WACT	
8.4	26.5	26.3	38.6	91231	1.1	820838	26972	0	

TASKS (CURRENT TASK TYPE)					TASKS IN SCHEDULER QUEUES				
BATCH	DIALOG	TP	SYS	#-TASKS	#-ACT	CPU-Q	IO-Q	PAG-Q	
105.2	167.0	58.0	169.0	499.1	148.6	1.8	0.2	0.0	

IO'S PER SEC							POSIX DATA	
TOTAL	DISK	PAGE	TD	CRYPT	TAPE	PRINT	SCALL/S	FORK/S
3724.0	294.9	0.0	3423.6	0.0	5.3	0.0	35.6	0.0

Monitoring information

(VM) This abbreviation will be output in the first line if the system monitored is a VM2000 system.

CPU UTILIZATION

(2 LM'S) This additional information indicates how many logical machines were active.

Proportion of CPU time during which the CPU was in the following processor states:

TU	User programs
TPR	Processing of program interrupts
SIH	Analysis of program interrupts
IDLE	Inactive

For systems comprising a number of logical machines, mean values for all active logical machines are output.

The sum of the TU, TPR, SIH and IDLE times is always 100%. For evaluation of the TU and TPR percentages, see [section "Accuracy of the SM2 data" on page 352](#).

In the monitoring cycle in which reconfiguration took place, these values and the number of I/O operations are marked as invalid (*).

SM2 supplies a STOP time for each logical machine. This value indicates the proportion of time during which the relevant logical machine was inoperable. The STOP time is output by:

- SM2 in the CPU report
- SM2R1 in the CPU report group in reports 1 and 2 when the parameter PROCESSOR-SPLITTING=*YES is specified
- INSPECTOR in the CPU report group, "STOP" report
- ANALYZER in the CPU report group, "Utilization normed (processor splitting)" and "Utilization real (processor splitting)" reports

PCS DATA

SRACT	Current overall system service rate
RDACT	Current delay of all jobs in the overall system

These fields will not contain any data if PCS was inactive during the last monitoring cycle.

MEMORY

NPP	Number of available pageable pages (4 KB) in main memory
WS-ACT	Working set of all active tasks (task queues 0 – 4) in 4 KB pages
WS-WACT	Working set of all ready, inactive tasks (task queue 5, in the case of PCS also task queue 6) in 4 KB pages

TASKS (CURRENT TASK TYPE)

BATCH	Average number of batch tasks during the monitoring cycle
DIALOG	Average number of interactive tasks during the monitoring cycle
TP	Average number of TP tasks during the monitoring cycle
SYS	Average number of system tasks during the monitoring cycle
#-TASKS	Average number of all tasks logged on to the system during the monitoring cycle



All batch, interactive, TP and system tasks are assigned to the TYPE to which they belong at the time of sampling.

TASKS IN SCHEDULER QUEUES

#-ACT	Average number of all active tasks (including SM2 tasks)
CPU-Q	Average number of tasks waiting to use the processor(s) and of tasks currently using the processor(s) (task queues 0 and 1, without the SM2 monitoring task)
IO-Q	Average number of tasks waiting for I/O termination (tasks in task queue 4 with I/O pend code)
PAG-Q	Average number of tasks waiting for paging (tasks in task queue 3)

IO'S PER SEC

These counts indicate the number of input/output operations (EXCP calls).

TOTAL	Sum of all I/O operations per second (including paging)
DISK	Number of DMS I/O operations to disk units per second (without paging I/O operations)
PAGE	Number of paging I/O operations per second. All paging operations (reading and writing) are counted.
TD	Number of IO operations to communication devices per second (family name TD), without the IOs of openCRYPT-BOX devices
CRYPT	Number of I/O operations to openCRYPT-BOX devices per second
TAPE	Number of I/O operations to disk devices per second (family name TAPE, BIMTAPE, UNMTAPE or MBK)
PRINT	Number of I/O operations to high-speed printers per second (family name PRINTER)

POSIX DATA

SCALL/S	Number of all system calls per second
FORK/S	Number of FORK system calls per second

These two ACTIVITY report values are the same as the corresponding values of the POSIX screen.

These fields will not contain any values if the POSIX monitoring program is not activated.

BCAM CONNECTION report

This report supplies information on the data sent and received as well as bucket-specific time statistics for defined connection sets.

Report output

The REPORT BCAM statement is used to request the output of the BCAM CONNECTION report.

The report can only be output if the SM2 administrator has

- defined the monitoring parameters using the SET-BCAM-CONNECTION-PARAMETERS statement
- defined the connection sets to be monitored using the ADD-BCAM-CONNECTION-SET statement
- started the monitoring process using the START-MEASUREMENT-PROGRAM TYPE=*BCAM-CONNECTION statement

A separate screen is output for each of the connection sets to be monitored (up to 32). If the number of connections to be monitored within a connection set is zero, no report is output for that connection set.

```

D016ZE04 SM2 BCAM CONNECTION      CYCLE:  60 S                      <date> <time>
SET: DIALOG                        # CONNECTIONS:  140  # REJECTIONS:  0

```

	TSDU	DATA	BUFFER	%	PWIN	PDATA	SCOM	ZWR
IN	0.7	0.0	0.0	0	2.4	2.0	0.0	0.0
OUT	2.3	0.8	0.0	0	0.3	5.1		

INWAIT TIME		REACT TIME		INPROC TIME		OUTPROC TIME	
BUCKET	COUNT	BUCKET	COUNT	BUCKET	COUNT	BUCKET	COUNT
5	42	500	40	5	42	5	73
10	0	1000	0	10	0	10	16
100	0	5000	1	100	0	100	30
1000	0	10000	0	1000	0	1000	15
> 1000	0	> 10000	0	> 1000	0	> 1000	0
AVG	0	AVG	49	AVG	0	AVG	33

Monitoring information

SET	Name of the connection set
# CONNECTIONS	Number of connections in the connection set. The computation includes connections that were active both at the current monitoring time and at the end of the preceding monitoring cycle, or were opened and/or closed during the monitoring cycle.
# REJECTED	Number of rejected attempts to open a connection
IN/OUT	Data sent and received
TSDU	Number of TSDUs per second TSDUs (Transport Service Data Units) correspond to application jobs on BCAM.
DATA	Average volume of data in KB per TDSU
BUFFER %	Amount of cache space occupied by unretrieved messages in KB and as a percentage of the maximum buffer size. This value describes the status at the end of the monitoring cycle, and not the mean value over the entire monitoring cycle.
PWIN	Number of packets without user data (Packet WINDOW) per second
PDATA	Number of packets with user data (Packet DATA) per second
SCOM	Number of send jobs per second in the event of a resource bottleneck (Send Call Over Maximum)
ZWR	Number of packets with Zero Window Information per second, i.e. where the partner application does not allow the sending of any further data
INWAIT TIME	Time between the display of a message and the retrieval of that message from the application
REACT TIME	Time between an application's send call and the receive call which immediately precedes it
INPROC TIME	Time between the arrival of a message at BCAM and the retrieval of that message by the application
OUTPROC TIME	Time between a send call and the transfer of the last byte of a message to the network
BUCKET	Limit values in milliseconds of the ranges within which the monitored data is arranged by order of magnitude. > indicates overflow values.

COUNT Number of messages sent/received in the individual ranges

AVG Average time per transmitted/received message in milliseconds.
Times under one millisecond for a transmitted/received message
are valued as zero when calculating the average value.



Connections which were established and/or cleared down during the monitoring cycle are also recorded.

BCAM MEMORY report

This report contains monitored data concerning the size of the resident memory for data transfer.

Report output

The REPORT BCAM_MEMORY statement is used to request the output of the BCAM MEMORY report. The report can only be output if the BCAM-CONNECTION monitoring program is activated (see BCAM CONNECTION report).

MEMORY POOL	(MB)	INPUT	OUTPUT
LIMIT		48	48
LIMIT TRAP		35	34
USED		1	0

Monitoring information

LIMIT	Maximum size of the resident memory for data transfer in MB
LIMIT TRAP	Threshold value in MB for the size of the resident memory from which BCAM issues warnings on the console
USED	Current size of the resident memory in MB



In the case of USED the monitored data only reflects the status at the end of the monitoring cycle, not the average value over the monitoring cycle.

CATEGORY report

This report provides an overview of the consumption statistics of the individual categories. For each category, the percentage CPU time and the I/O operations for paging and disk devices are output.

Report output

The REPORT CATEGORY statement is used to request the output of the CATEGORY report.

The report can only be output if the SM2 administrator has

- defined the devices to be monitored using the SET-SYSTEM-PARAMETERS statement
- started the monitoring process using the START-MEASUREMENT-PROGRAM TYPE=*SYSTEM statement

D016ZE04 SM2 CATEGORY CYCLE: 60 S					<date> <time>
CATEGORY NAME	CPU-TIME (%)	PAGING-I/O (1/S)	I/O (1/S)	(2 LM'S)	
SUM	34.8	0.0	299.1		
SYS	11.6	0.0	21.0		
DIALOG	0.0	0.0	0.0		
BATCH	10.9	0.0	144.4		
TP	2.6	0.0	14.0		
TP1	0.0	0.0	0.0		
BATCHDB	0.0	0.0	0.0		
BATCHF	4.0	0.0	7.1		
DIALOG1	1.1	0.0	40.4		
DIALOG2	4.3	0.0	72.2		

Monitoring information

This report contains the totals (SUM) for all categories in the first line following the table header. The data for the individual categories (up to 16) is output in the following lines.

The columns contain the following data:

CATEGORY NAME	Task category designation
CPU-TIME (%)	Each category's percentage of the time (TU+TPR+SIH+IDLE)
PAGING-IO (1/S)	Total number of paging I/Os per second on the monitored devices The pseudo-category SUM contains all paging I/Os (read, write). In the individual categories, only the paging I/Os of type read are output.
DISK-IO (1/S)	Total number of I/O operations performed per second on the monitored devices
(2 LM'S)	Number of active logical machines

CATEGORY QUEUE report

The PRIOR task scheduler uses three criteria for controlling main memory and CPU utilization by the tasks. These criteria are

- the task category
- the multiprogramming level
- the task priority

Detailed information on these criteria can be found in the manuals “Introduction to System Administration” [6] and “Performance Handbook” [5]. In the following, only information that is relevant for the SM2 reports is discussed.

Task category

In BS2000, 16 task categories are distinguished at present, i.e. the 4 standard categories

- SYS (system tasks)
- TP (inquiry-and-transaction tasks)
- DIALOG (interactive tasks)
- BATCH (batch tasks)

and up to 12 categories that can be defined freely by the system administrator.

Multiprogramming level

The multiprogramming level denotes the number of tasks of a given category that may use main memory concurrently, i.e. the number of active tasks of a category.

In the /MODIFY-TASK-CATEGORIES command the system administrator uses the category attributes

- MIN MPL
- MAX MPL
- WEIGHT

to specify the relative priority of each category in order to make decisions regarding activation (= allocation of authorization to use main memory).

MIN MPL is used to specify a minimum number of active tasks of a category. The system tries to reach the specified MIN MPL value first.

MAX MPL is used when there is no fixed limit, i.e. activation continues even after the maximum MPL value as long as no resource bottleneck occurs.

WEIGHT is used to control the sequence of activation.

Report output

The REPORT CATQ statement is used to request the output of the CATEGORY QUEUE report.

```

D016ZE04 SM2 CATEGORY QUEUE CYCLE: 60 S SAMPLES: 144 <date> <time>
|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|
| CATEGORY |          TASKS IN SCHEDULER QUEUES          |          MPL          |
| NAME     | CPU-Q | IO-Q | PAG-Q | ACT  | WACT | NADM | ALL  | MIN  | MAX  | WEIGHT |
|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|
| SUM      | 1.8   | 0.2   | 0.0   | 148.6 | 0.0  | 0.0  | 499.1 |      |      |        |
| SYS      | 0.6   | 0.0   | 0.0   | 64.0  | 0.0  | 0.0  | 168.9 | 30   | 64   | 512    |
| DIALOG   | 0.0   | 0.0   | 0.0   | 0.0   | 0.0  | 0.0  | 11.0  | 1    | 3    | 126    |
| BATCH    | 0.4   | 0.1   | 0.0   | 16.2  | 0.0  | 0.0  | 36.0  | 16   | 17   | 6      |
| TP       | 0.2   | 0.0   | 0.0   | 44.0  | 0.0  | 0.0  | 58.0  | 45   | 47   | 52     |
| TP1      | 0.0   | 0.0   | 0.0   | 0.0   | 0.0  | 0.0  | 0.0   | 0    | 2    | 126    |
| BATCHDB  | 0.0   | 0.0   | 0.0   | 0.0   | 0.0  | 0.0  | 0.0   | 0    | 2    | 126    |
| BATCHF   | 0.3   | 0.0   | 0.0   | 22.0  | 0.0  | 0.0  | 69.2  | 22   | 24   | 2      |
| DIALOG1  | 0.1   | 0.0   | 0.0   | 0.6   | 0.0  | 0.0  | 154.2 | 1    | 2    | 96     |
| DIALOG2  | 0.2   | 0.1   | 0.0   | 1.8   | 0.0  | 0.0  | 1.8   | 2    | 3    | 70     |

```

Monitoring information

This report contains the totals (SUM) for all categories in the first line following the table header. The data for the individual categories (up to 16) is output in the following lines.

The columns contain the following data:

CATEGORY NAME Task category designation

TASK IN SCHEDULER QUEUES

CPU-Q Average number of tasks of a category waiting to use the processor(s) and of tasks currently using the processor(s) (task queues 0 and 1, excluding the SM2 monitoring task)

IO-Q Average number of tasks of a category waiting for I/O termination (tasks in task queue 4 with I/O pend code)

PAG-Q Average number of tasks of a category waiting for paging (tasks in task queue 3)

ACT Average number of active tasks of a category

WACT Average number of inactive ready tasks of a category (tasks in task queue 5)

NADM	Average number of tasks of a category waiting for activation (tasks in task queue 6; this value is supplied only if PCS is used)
ALL	Total number of all tasks of a category
MPL	
MIN	Average value of the minimum multiprogramming level used to determine the activation of tasks of a category
MAX	Average value of the maximum multiprogramming level used to determine the activation of tasks of a category. This value is used primarily to prevent overloading by setting a load limit.
WEIGHT	Average weight used to determine the activation sequence for the task categories

In PCS mode the MIN MPL, MAX MPL and WEIGHT values are modified dynamically, i.e. they no longer contain the system administrator specifications.

CATEGORY WSET report

Report output

The REPORT CATW statement is used to request the output of the CATEGORY WSET report.

```

D016ZE04 SM2 CATEGORY WSET CYCLE: 60 S SAMPLES: 144 <date> <time>

```

CATEGORY NAME	WS-ALL		WS-ACT		WS-WACT	
	PPC	UPG	PPC	UPG	PPC	UPG
SUM	57190	586563	26972	367537	0	0
SYS	10931	264105	9012	257540	0	0
DIALOG	873	2290	0	0	0	0
BATCH	5642	50189	3487	36466	0	0
TP	13264	61549	11593	56282	0	0
TP1	0	0	0	0	0	0
BATCHDB	0	0	0	0	0	0
BATCHF	5461	25098	2414	13760	0	0
DIALOG1	20613	180039	60	196	0	0
DIALOG2	406	3293	406	3293	0	0

Monitoring information

This report contains the totals (SUM) for all categories in the first line following the table header. The data for the individual categories (up to 16) is output in the following lines.

The columns contain the following data:

CATEGORY NAME	Task category designation
WS-ALL	Working sets of all tasks
PPC	Sum of the average planned page counts of all tasks in a category
UPG	Sum of the average used page counts of all tasks in a category
WS-ACT	Working sets of active tasks
PPC	Sum of the average planned page counts of the active tasks in a category
UPG	Sum of the average used page counts of the active tasks in a category
WS-WACT	Working set of ready active tasks
PPC	Sum of the average planned page counts of the tasks in a category waiting for activation.

UPG Sum of the average used page counts of the tasks in a category waiting for activation.

For an explanation of the values PPC and UPG, see [“paging in BS2000” on page 669](#) of the glossary.

CHANNEL report

This report contains monitored data on the activity of I/O channels.

Report output

The REPORT CHANNEL statement is used to request the output of the CHANNEL report.

The values BUSY and NON OVERLAP are recorded by default for type S channels.

Other monitored data can be output only if the SM2 administrator has

- defined the channels to be monitored using the SET-CHANNEL-IO-PARAMETERS statement
- started the monitoring process using the statement START-MEASUREMENT-PROGRAM TYPE=*CHANNEL-IO.

The monitoring program is started for all channels when the SM2 subsystem is started.

Overflow screens may be requested for this report.

The monitored data is output on the basis of the sort criterion. The sort criterion is defined by means of the SELECT-CHANNEL-PARAMETERS statement (see [page 183](#)). By default, the sort criterion is BUSY.

D016ZE04 SM2 CHANNEL (I) CYCLE: 60 S SAMPLES: 144 <date> <time>									
PATH ID	TYPE	BUSY (%)	N-OVER (%)	PAM		BYTES		NODATA	
				(I/O/S)	(KB/S)	(I/O/S)	(KB/S)	(I/O/S)	
09C	TYP FC	0		0.0	0.0	3423.8	8036.5	0.0	
098	TYP FC	1		67.3	591.9	0.0	0.0	7.1	
058	TYP FC	1		66.9	606.8	0.0	0.0	7.2	
018	TYP FC	1		67.3	598.8	0.0	0.0	6.6	
0D8	TYP FC	1		65.7	591.7	0.0	0.0	6.5	
05C	TYP FC	0		0.0	0.0	4.3	7125.5	0.7	
00F	TYP S	0	0	0.0	0.0	0.0	0.0	0.4	
0DC	TYP FC	0		0.0	0.0	0.0	0.0	0.4	
049	TYP S	0	0	0.0	0.0	0.3	0.0	0.0	
0C8	TYP S	0	0	0.2	0.7	0.0	0.0	0.0	
048	TYP S	0	0	0.1	0.5	0.0	0.0	0.0	
003	TYP 2	0	0	0.0	0.0	0.1	0.0	0.0	
08C	TYP S	0	0	0.0	0.0	0.1	0.0	0.0	
04F	TYP S	0	0	0.0	0.0	0.0	0.0	0.0	
0CF	TYP S	0	0	0.0	0.0	0.0	0.0	0.0	
008	TYP S	0	0	0.0	0.0	0.0	0.0	0.0	
04B	TYP S	0	0	0.0	0.0	0.0	0.0	0.0	

Monitoring information

In addition to the name of the report, the sort criterion used is output in the header: B (busy) or I (total number of input/output of PAM block transfer, byte transfer and input/output without data transfer).

The values BUSY and N-OVER are recorded by default for channels of type S. All other values are output only if the CHANNEL-IO monitoring program has been activated for the monitored channels. In the case of non-monitored channels, blanks are entered.

PATH-ID	Channel address.
TYPE	Channel type; possible outputs: – BUS – TYP FC – TYP S
BUSY	Percentage of the monitoring cycle during which the corresponding channel was active.
N-OVER	Percentage of the monitoring cycle during which the corresponding channel was active and all processors were inactive (NON OVERLAP).
PAM	PAM block transfer (disk devices): Number of I/O operations or data volume in KB transferred per sec.
BYTES	Byte transfer (MBK): Number if I/O operations or data volume in KB transferred per second.
NODATA	Number of I/O operations without data transfer per second.

Servers with /390 architecture

For channels of hte type FC (Fibre Channel) the correct type is output.

For channels of TYP FC the following restrictions apply:

- A value for BUSY(%) can be output only if the channel is monitored with the CHANNEL-IO monitoring program. Additionally the SORT=*IO operand should be set in the SELECT-CHANNEL-PARAMETERS statement. The utilization is calculated from the number of inputs/outputs. Direct, sample-based recording of the utilization is not possible.
- There are no values for N-OVER(%).
- When using VM2000, please note that BUSY(%) does not indicate the total utilization of the channel (in contrast to other channel types) but the local utilization of the relevant guest system.

Servers with x86 architecture

The FibreChannel periphery is operated under BS2000 by emulating the devices in X2000 as bus peripherals. The bus periphery is always specified as type BUS in SM2.

The following restrictions must be borne in mind:

- The real, physical FibreChannels can be different from the channels that are visible for the BS2000 operating system (and those displayed for SM2). Consequently the measured data supplied by SM2 is not meaningful with regard to the individual channel. Valid values can be obtained by forming sum values.
- For channels of type BUS no values are output for BUSY(%) and N-OVER(%).



SM2 is not notified of device management commands (see the “Commands” manual [3]). In the monitoring cycle in which the operator enters such commands, the data for the affected channels may be incorrect

When using VM2000, please note that the channel occupancy values BUSY and N-OVER for type S channels always indicate the overall load of the channel. The percentage utilization by an individual guest system cannot be determined. The value N-OVER is therefore insignificant. The monitored data supplied by the CHANNEL-IO monitoring program always refers to the respective guest system.

CMS report

This report supplies performance data on the CMS (Catalog Management System). Data is collected separately for each catalog. A report is output for each pubset imported during the last monitoring cycle and for all private volumes. The home pubset is marked. Detailed information on CMS is contained in the manuals “HIPLEX MSCF” [8] and “Introduction to System Administration” [6].

Multiprocessor systems

In a multiprocessor system, the CMS report supplies access data for the tasks of the system which manages the catalog (LOCAL) as well as access data for tasks running on remote computers. This data is listed under REMOTE.

In an integrated shared pubset, access by so-called slave sharers is also recorded on the pubset master (SHARED identifier).

Report output

The REPORT CMS statement is used to request the output of the CMS report. The report can only be output if the SM2 administrator has started the monitoring process using the statement START-MEASUREMENT-PROGRAM TYPE=*CMS.

The SELECT-CMS-PUBSET statement (see [page 184](#)) can be used to specify the set of reports to be output. By default, all pubsets and all private disks are output.

D016ZE04 SM2 CMS		CYCLE: 60 S		SAMPLES: 144		<date> <time>		
CAT-ID: 20S6(SLAVE)		HIT-RATE: 0 %		#-BMT: 0 (PAGEABLE)				
REQ. QU	SER. QU	BLOCK QU	USER-ID Q	CE QU	SFP QU	SMP QU	PH-RD/S	PH-WRT/S
0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.0	0.0
ITEMS	READ	READ-LBN	SCAN	UPD-RENAME	WRITE/CLEAR			
RESPONSETIME (MS)								
LOCAL-FILE (/S)	12.8	0.9	10.3	0.0	0.0	9.6		
JV (/S)	0.6	0.0	0.0	0.0	0.0	0.2		
REMOTE-FILE (/S)	0.0	0.0	0.0	0.0	0.0	0.0		
JV (/S)	0.0	0.0	0.0	0.0	0.0	0.0		
SHARED-FILE (/S)	0.0	0.0	0.0	0.0	0.0	0.0		
JV (/S)	0.0	0.0	0.0	0.0	0.0	0.0		

Monitoring information

CAT-ID	Catalog identifier
POOL OF PRIVATE DISKS	This is output for private disks in place of CAT-ID. The monitoring data refers to the F1 label accesses of all private disks.
HOME	The catalog is marked with home pubset if it belongs to the home pubset.
MASTER	If the catalog belongs to an integrated shared pubset and the SM2 executes on the associated master processor, the catalog is identified as the MASTER catalog.
SLAVE	If the catalog belongs to an integrated shared pubset and the SM2 executes on the associated slave processor, the catalog is identified as the SLAVE catalog.
HIT-RATE	No read access is required if a wanted catalog entry is already in a CMS storage area. HIT-RATE is the percentage of such (non-physical) accesses of all read accesses.
#-BMT	Number of buffer management tables or buffers used by CMS (refer to the “Introduction to System Administration” [6]).
PAGEABLE	Buffers are pageable (class 4 memory)
or	
RESIDENT	Buffers are resident (class 3 memory). Refer to the “Introduction to System Administration” [6] .
REQ. QU	Average number of requests waiting for the release of a buffer management table or currently using a buffer management table. If the request queue is always long, this means that there are too few buffers compared to the number of CMS calls, i.e. that the number of buffer management tables should be increased. If the maximum number of buffer management tables is already being used, the pubset can be subdivided into smaller pubsets. The desired number of buffer management tables can be specified at system setup using the BMTNUM parameter.
SER. QU	Average number of requests waiting for a serialization lock for exclusively searching a suitable buffer management table or currently using this lock (relevant for system engineers)

BLOCK QU	Average number of requests waiting for an exclusive lock for processing a block of a partition or currently using this lock (a partition corresponds to a user ID)
USER-ID QU	Average number of requests waiting for a partition lock or currently using this lock. This lock is not normally exclusive.
CE QU	Average number of requests waiting for an exclusive lock for processing a catalog entry of a block of a partition or currently using this lock
SFP QU	Average number of requests to the speedcat for single-feature subsets
SMP QU	Average number of requests to the catalog index for system-managed subsets
PH-RD/S	Number of physical read accesses per second
PH-WRT/S	Number of physical write accesses per second

The table columns contain the following data:

ITEMS	Name of the monitored variable
READ	Number of read accesses to catalog entries without LBN specification
READ-LBN	Number of read accesses to catalog entries with LBN specification
SCAN	Number of SCAN accesses to catalog entries
UPD-RENAME	Number of write accesses to catalog entries merely involving manipulation of file attributes
WRITE/CLEAR	Number of write or erase accesses to catalog entries

The individual lines contain the following data:

RESPONSETIME (MS)	Global response times for the above data for files and job variables and for all accesses, regardless of their origin (local or remote computer).
LOCAL-FILE (/S)	File entry accesses originating from the local computer.
JV (/S)	Job variable entry accesses originating from the local computer.
REMOTE-FILE (/S)	File entry accesses originating from a remote computer.
JV (/S)	Job variable entry accesses originating from a remote computer.
SHARED-FILE (/S)	File entry accesses by a slave.
JV (/S)	Job variable entry accesses by a slave.

CPU report

The CPU report provides you with an overview of the utilization of the individual processors.

Report output

The REPORT CPU statement is used to request the output of the CPU report. Overflow screens may be requested for this report.

```

D016ZE04 SM2 CPU                                CYCLE:  60 S                                <date> <time>
| NORMED TO 100 %                                | REAL
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
LM | TU % | TPR % | SIH % | IDLE % | STOP % | TU % | TPR % | SIH % | IDLE %
-----+-----+-----+-----+-----+-----+-----+-----+-----+
AVG |  8.4 | 26.5 | 26.3 | 38.6 |  0.0 |  8.2 | 26.0 | 25.8 | 37.8
-----+-----+-----+-----+-----+-----+-----+-----+-----+
 1 |  9.9 | 25.1 | 26.1 | 38.7 |  0.0 |  9.7 | 24.6 | 25.6 | 38.0
 2 |  0.0 |  0.0 |  0.0 |  0.0 | 100.0 |  0.0 |  0.0 |  0.0 |  0.0
 3 |  6.9 | 27.9 | 26.5 | 38.4 |  0.0 |  6.8 | 27.4 | 26.0 | 37.7
-----+-----+-----+-----+-----+-----+-----+-----+-----+

```

Monitoring information

- LM** Logical machine number.
- Percentage of time during which the processor was in the following states:
- TU % user programs
 - TPR % processing of program interrupts
 - SIH % analysis of program interrupts
 - IDLE % inactive
 - STOP % Percentage of time during which the processor was not operable
- NORMED** In the NORMED TO 100 % section, the sum of the percentage values for TU, TPR, SIH and IDLE is 100 %.

- REAL The percentage values measured for TU, TPR, SIH and IDLE are output in the REAL section.
When using VM2000, there may be discrepancies between the NORMED value and the REAL value (see also [section “SM2 together with VM2000” on page 356](#)). In particular, you should remember that the times specified refer to the assigned CPUs of the VM and not to the system as a whole.
- AVG The average values across all processors are output in the AVG row.



While SM2 records the sum of the TU and TPR times precisely, this total is divided into TU and TPR times on the basis of the proportion of the samples that fall to TU and TPR. This causes statistical inaccuracies in the calculation of TU and TPR times, which can be noticeable in the case of small TU and TPR values (i.e. when a small number of samples fall to TU and TPR). TU or TPR values of 0.0 % (online reports) or 0.0 milliseconds (TIME_IO buffer of the C interface), in particular, should therefore be interpreted with caution.

The connection between the CPU values under REAL and under CPU MEAS in the VM report is explained in [section “Special applications” on page 356](#).

DAB reports

These reports provide information on DAB (Disk Access Buffer) activity (for an explanation of the DAB function, please refer to the “DAB” manual [7]).

- The DAB report outputs access statistics for each DAB cache.
- The DAB CACHE report contains detailed information on the subareas and files served by a particular DAB cache.

The report can only be output if the SM2 administrator has started the monitoring process using the statement `START-MEASUREMENT-PROGRAM TYPE=*DAB` statement.

The figure below illustrates the hierarchy of reports and shows the statements used to request them:

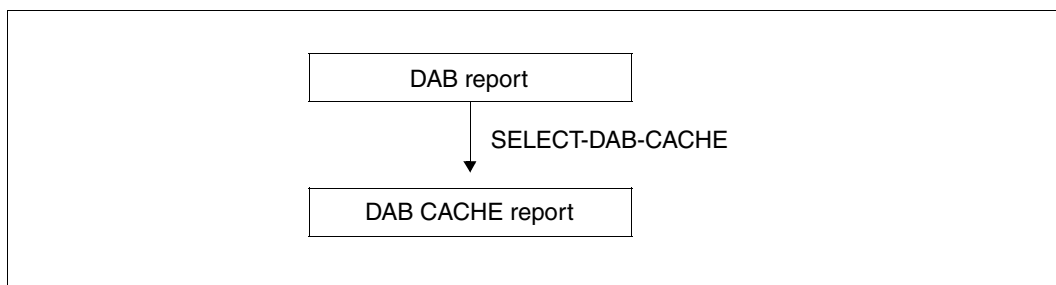


Figure 5: Hierarchy of DAB reports

DAB report

This report supplies monitored data on the various DAB caches.

Report output

The REPORT DAB statement is used to request the output of the DAB report.

Overflow screens may be requested for this report.

The monitored data output is sorted under the categories READ and WRITE.

D016ZE04 SM2 DAB		CYCLE: 60 S		<date> <time>			
CACHE-ID	SIZE	READ (1/S)	RD HIT (%)	WRITE (1/S)	WR HIT (%)	OVER (1/S)	
MMRD#SHVOL	50.0 MB	64.2	99.4	0.0	0.0	0.0	
MMRD#20SH	150 MB	30.8	100.0	2.3	0.0	0.0	
20SQ	20.0 MB	0.0	0.0	0.0	0.0	0.0	
10SZ	10.0 MB	0.0	0.0	0.0	0.0	0.0	
MMRD#SHPUB	250 MB	0.0	0.0	0.0	0.0	0.0	
10SY	10.0 MB	0.0	0.0	0.0	0.0	0.0	

Monitoring information

Each line of the report contains data on the different DAB caches created by the BS2000 system administrator using the /START-DAB-CACHING command.

The individual columns contain the following data:

CACHE-ID	Name of the DAB cache defined by the user in the CACHE-ID operand of the /START-DAB-CACHING command or assigned internally by DAB. If the DAB cache is located in GS (global storage), the name of the partition is also output in brackets.
SIZE	Size of the DAB cache in KB, MB or GB.
READ (1/S)	Number of read operations per second in all subareas or files served by the DAB cache.
RD HIT (%)	Percentage of read operations without disk access. SM2 compares the number of read operations for which disk access was not required (as the data to be read was already in the cache) with the total number of read operations in all subareas or files served by the DAB cache.

WRITE (1/S)	Number of write operations per second to all subareas or files served by the DAB cache.
WR HIT (%)	Percentage of write operations in which data was written to the cache in relation to the total number of write operations to all subareas or files served by the DAB cache.
OVER (1/S)	Number of failed attempts to use the DAB cache per second due to an overload.



While the SM2 DAB report contains the data of the last monitoring cycle, the BS2000 command /SHOW-DAB-CACHING supplies the READ, WRITE and HIT values since the DAB buffer was set up. The values of the DAB report and those of the SHOW-DAB-CACHING command cannot therefore be compared directly.

DAB CACHE report

This report supplies information on the subareas or files served by a DAB cache.

Report output

The DAB CACHE report is output only at the explicit request of the SM2 user by means of the SELECT-DAB-CACHE statement (see [page 185](#)). After this statement has been entered, the user receives not only the DAB report, but also the DAB CACHE report for the DAB caches he or she has selected.

The DAB CACHE report provides detailed information on the subareas or files served by DAB.

Overflow screens may be requested for this report.

The monitored data output is sorted under the categories READ and WRITE.

It is possible to output information on up to 16 DAB caches.

```

D016ZE04 SM2 DAB CACHE          CYCLE:  60 S          <date> <time>
ID: MMRD#SHPUB                  MOD: R          MED:MM SIZE(FIX): 250.0 MB
  VSN / FILE          | FIRST-HP, LAST-HP | S | READ | RD-HIT | WRITE | WR-HIT | OVER
  -----|-----|---|-----|-----|-----|-----|-----
 *SUMMARY            | 1-          1 | | 0.0 | 0.0 | 0.0 | 0.0 | 0.0
  -----|-----|---|-----|-----|-----|-----|-----
    
```

Monitoring information

- ID Name of the appropriate DAB cache as selected via the ADD-CACHE-ID operand of the SELECT-DAB-CACHE statement. If the DAB cache is located in GS (global storage), the name of the partition is also output in brackets.
- MOD R: The cache is used as a read cache.
- W: The cache is used as a write cache.
- W/R: The cache is used as a read/write cache.
- PFA: The cache is used as a PFA cache.

In the case of the R/W, W and PFA modes, the Force Out parameter is also specified. This indicates whether and at what fill level data written to the cache should be written back to the external data volumes.

N	Data is not written back (No)
L	Low fill level
H	High fill level
MED	Cache medium
MM:	Main memory
GS:	Global storage
SIZE(FIX)	Size of the DAB cache. This can be fixed (FIX) or variable (VAR).
SIZE(VAR)	In the case of (VAR), the size at the time of data acquisition is output.

Each line in this report contains data on the different subareas or files served by the cache with the specified CACHE-ID.

The individual columns contain the following data:

VSN / FILE	VSN of the volume containing the subarea or, if DAB supports a file, the name of the file containing the subarea. If there is more than one subarea on the same volume or in the same file, these subareas are listed in successive lines. The relevant VSN or file name can be found only in the first line of this list. If the file name is longer than 21 characters, it is truncated on the right.
FIRST-HP, LAST-HP	First and last <u>physical</u> block numbers of the subarea, if a volume is listed in column 1. First and last <u>logical</u> block numbers of the subarea, if a file name is listed in column 1.
S	Any subarea which was not served during the last monitoring cycle is marked with an asterisk (*) in this column (S = served). No data is output for this subarea.

The columns READ, RD-HIT, WRITE, WR-HIT and OVER have the same meaning as in the DAB report. However, the information they contain refers to individual subareas or files and not to the DAB cache as a whole.



In order to reduce the amount of data recorded, the following applies to the DAB-CACHE report when AutoDAB is used:

If the cache is created with `AREA=*BY-SYSTEM(ADM-PFA)` or with `CACHED-FILES=*BY-SYSTEM(USER-PFA)`, the output does not include a list of the files served. Instead, it contains a summary of the entire cache. The `(VSN/FILE=*SUMMARY)` row of the DAB CACHE report thus contains the same data as the DAB report. As a result, the information in the FIRST-HP, LAST-HP column is irrelevant.

Monitored data on the files served by the cache can be output using the DAB statement `SHOW-DAB-CACHING CACHE-ID=*ALL, INF=*SYSTEM-CACHED-FILES (CACHING=*ACTIVE)`.

DEVICE DISK report

This report contains data on the I/O operations performed with peripheral devices during one monitoring cycle. It counts the number of EXCP calls per second (see glossary, [page 666](#)).

Report output

The REPORT DEVICE_DISK statement allows the user to request the output of the DEVICE DISK report.

The monitored data under SERVICETIME is only output if the SM2 administrator:

- uses the SET-SAMPLING-DEVICE-PARAMETERS statement to define the monitoring of the service times and
- uses the START-MEASUREMENT-PROGRAM TYPE=*SAMPLING-DEVICE statement to start monitoring.

When the SM2 subsystem is started, the SAMPLING-DEVICE monitoring program is started without the service times being recorded.

Overflow screens can be requested for this report.

The output of the monitored data is sorted on the basis of the sort criterion. The sort criterion is defined by means of the SELECT-DEVICE-DISK-PARAMETERS statement (see [page 186](#)). Sorting on the basis of the number of I/O operations is specified by default.



- SM2 is not notified of device management commands (see the “Commands” manual [\[3\]](#)). In the monitoring cycle in which such commands are entered, incorrect data may be produced for the affected devices.
- GS volumes are only output in the GSVOL report.
- For parallel access volumes (PAV), the following applies:
 - Only basic devices are output.
 - The monitored data refers to the basic device and the assigned alias devices (summation, averaging).
 - If an alias device changes from one basic device to another, then the datas of the alias device are not included in the datas of any basic device in the actual measurement interval.
 - For information on PAV see the “Introduction to System Administration” [\[6\]](#).

```

D016ZE04 SM2 DEVICE DISK (I) CYCLE: 60 S SAMPLES: 150 <date> <time>
  MN | VOLUME | P | UTIL | Q | RSC | READ | WRITE | SERVICE TIME
  IO% | PG% | LGTH | (IO/S) | (KB/S) | (IO/S) | (KB/S) | HW (MS) | SW (MS)
-----|-----|---|-----|---|-----|-----|-----|-----|-----|
B3A8 | 6VS4.0 | | 13 | 0.1 | | 84.5 | 194.2 | 83.1 | 166.3 | 0.8 | 0.8
B3A9 | 6VS4.1 | | 12 | 0.1 | | 83.2 | 166.4 | 66.7 | 138.8 | 0.8 | 0.8
D0E4 | QVS1.5 | | 9 | 0.1 | | 131.7 | 2174.7 | 0.0 | 0.0 | 0.5 | 0.5
D0E3 | QVS1.4 | | 2 | 0.0 | | 70.2 | 193.8 | 0.0 | 0.0 | 0.3 | 0.3
D101 | CVS3.3 | B | 2 | 0.0 | | 9.4 | 195.6 | 19.4 | 313.7 | 1.8 | 1.8
D336 | | | 1 | 0.0 | | 12.9 | 6.8 | 1.1 | 2.3 | 1.3 | 1.3
B336 | | | 2 | 0.0 | | 10.2 | 6.0 | 1.0 | 2.0 | 1.6 | 1.6
D10B | CVS6.1 | | | 0.0 | | 11.2 | 92.9 | 0.0 | 0.1 | 0.4 | 0.4
D100 | CVS3.2 | B | | 0.0 | | 2.5 | 11.4 | 3.2 | 11.8 | 1.2 | 1.2
D0FE | CVS3.0 | | | 0.0 | | 1.6 | 6.6 | 2.1 | 7.2 | 1.1 | 1.1
D10C | CVS6.2 | | | 0.0 | | 3.5 | 19.0 | 0.0 | 0.0 | 0.5 | 0.5
D005 | 2OSH.0 | | | 0.0 | | 1.1 | 0.8 | 2.4 | 8.5 | 1.3 | 1.3
D27B | GVS5.5 | | | 0.0 | | 2.7 | 14.6 | 0.2 | 0.5 | 0.4 | 0.4
D0FF | CVS3.1 | | | 0.0 | | 1.1 | 4.6 | 1.6 | 6.0 | 1.1 | 1.1
D279 | GVS4.5 | | | 0.0 | | 1.8 | 21.6 | 0.1 | 0.2 | 0.7 | 0.7
D43D | QVS1.3 | B | | 0.0 | | 1.6 | 7.3 | 0.0 | 0.0 | 0.6 | 0.6
D435 | QVS2.1 | B | | 0.0 | | 1.4 | 12.1 | 0.0 | 0.0 | 0.4 | 0.4
  
```

Monitoring information

In addition to the name of the report, the sort criterion is also output in the header: B (busy), Q (Qlength), I (IO) or S (SWTIME).

- MN Mnemonic name of the disk device.
- VOLUME VSN of the disk device or spaces (e.g. when the disk belongs to a subset that has not been imported).
- P PAV information.
Basic devices are marked with a "B".
- UTIL
- IO% PG% Percentage of time during which the device was active during the monitoring cycle due to DMS input/output operations (IO) or paging input/output operations (PG)
- Q-LGTH Average length of the device queue, including jobs currently being serviced
- RSC Average number of parallel processed RSC IOs (only servers with x86 architecture)
- READ Number of input operations or number of KB transferred per second
- WRITE Number of output operations or number of KB transferred per second

SERVICETIME

HW (MS)

Average hardware service time in milliseconds

SW (MS)

Average software service time (including hardware service time) in milliseconds

DEVICE OTHER report

In the DEVICE OTHER report, data is output for all devices that are not contained in the DEVICE DISK, DEVICE TD and DEVICE TAPE reports (e.g. consoles).

The volume of data or distinction between the IOs on the basis of READ and WRITE can only be correct when the corresponding device drivers supply the data.

Report output

The REPORT DEVICE_OTHER statement allows the user to request the output of the DEVICE OTHER report.

The report can only be output when the SAMPLING-DEVICE monitoring program is switched on. The monitoring program is started automatically when the SM2 subsystem is started or by the SM2 administrator by means of the statement START-MEASUREMENT-PROGRAM TYPE=*SAMPLING-DEVICE.

Overflow screens can be requested for this report. The output is sorted on the basis of the number of I/O operations.

D016ZE04 SM2 DEVICE OTHERS CYCLE: 60 S SAMPLES: 144 <date> <time>							
MN	TYPE	Q LGTH	IO%	READ		WRITE	
				(IO/S)	(BYTES/S)	(IO/S)	(BYTES/S)
CR	CON3027	0.0		0.0	0.0	0.0	0.0
CU	CON3027	0.0		0.0	0.0	0.0	0.0
CB	CON3027C	0.0		0.0	0.0	0.0	0.0
CC	CON3027C	0.0		0.0	0.0	0.0	0.0
CD	CON3027C	0.0		0.0	0.0	0.0	0.0
CE	CON3027C	0.0		0.0	0.0	0.0	0.0
CV	CON3027C	0.0		0.0	0.0	0.0	0.0
CZ	CON3027C	0.0		0.0	0.0	0.0	0.0
NA	CON3027C	0.0		0.0	0.0	0.0	0.0
NB	CON3027C	0.0		0.0	0.0	0.0	0.0
NO	CON3027C	0.0		0.0	0.0	0.0	0.0
N1	CON3027C	0.0		0.0	0.0	0.0	0.0
N2	CON3027C	0.0		0.0	0.0	0.0	0.0
N9	CON3027C	0.0		0.0	0.0	0.0	0.0
N3	CON38	0.0		0.0	0.0	0.0	0.0
N8	CON38	0.0		0.0	0.0	0.0	0.0
N4	CON3803	0.0		0.0	0.0	0.0	0.0

Monitoring information

MN	Mnemonic device name
TYPE	Short name for the device type
Q LGTH	Average length of the device queue including the jobs that are currently being dealt with
IO%	Percentage of the monitoring interval in which the device was active with I/Os
READ	Number of input operations or number of bytes transferred per second
WRITE	Number of output operations or number of bytes transferred per second

DEVICE TAPE report

The DEVICE TAPE report contains monitored data on I/O operations involving magnetic tape cartridge devices during a monitoring interval. The number of EXCP calls per second are counted (see glossary, [page 666](#)).

Report output

The REPORT DEVICE_TAPE statement allows the user to request the output of the DEVICE TAPE report. The report can only be output when the SAMPLING-DEVICE monitoring program is switched on. The monitoring program is started automatically when the SM2 subsystem is started or by the SM2 administrator by means of the statement START-MEASUREMENT-PROGRAM TYPE=*SAMPLING-DEVICE.

Overflow screens can be requested for this report. The output is sorted on the basis of the number of I/O operations.

D016ZE04 SM2 DEVICE TAPE/MBK CYCLE: 60 S SAMPLES: 144 <date> <time>							
MN	VOLUME	TYPE	IO%	READ		WRITE	
				(IO/S)	(KB/S)	(IO/S)	(KB/S)
AF	FB4024	3590E	6	0.2	0.0	3.9	6939.8
AD		3590E	0	0.6	185.9	0.0	0.0
AE		3590E		0.1	0.0	0.0	0.0
AQ		3590E	0	0.1	0.0	0.0	0.0
BD		3590E		0.1	0.0	0.0	0.0
BE		3590E		0.1	0.0	0.0	0.0
BF		3590E		0.1	0.0	0.0	0.0
BQ		3590E		0.1	0.0	0.0	0.0
MD		3590E		0.0	0.0	0.0	0.0
MA		3590E		0.0	0.0	0.0	0.0
MC		3590E		0.0	0.0	0.0	0.0
MH		3590E		0.0	0.0	0.0	0.0
I0		3590E		0.0	0.0	0.0	0.0
I1		3590E		0.0	0.0	0.0	0.0
I2		3590E		0.0	0.0	0.0	0.0
I3		3590E		0.0	0.0	0.0	0.0
I4		3590E		0.0	0.0	0.0	0.0

Monitoring information

MN	Mnemonic device name
VOLUME	VSN of the volume connected during the whole of the last monitoring interval. If no volume is specified, either the volume was changed during the last monitoring interval or at the time of the query it was not possible to access the volume table for system-internal reasons.
TYPE	Short name for the device type
IO%	Percentage of the monitoring interval in which the device was active with I/Os
READ	Number of input operations or number of bytes transferred per second
WRITE	Number of output operations or number of bytes transferred per second

DEVICE TD report

The DEVICE TD report provides data on the I/O operations of communication devices. The data is recorded per connection (i.e. on the basis of the mnemonic device name).

Report output

The REPORT DEVICE_TD statement allows the user to request the output of the DEVICE TD report. The report can only be output when the SAMPLING-DEVICE monitoring program is switched on. The monitoring program is started automatically when the SM2 subsystem is started or by the SM2 administrator by means of the statement START-MEASUREMENT-PROGRAM TYPE=*SAMPLING-DEVICE.

Overflow screens can be requested for this report. The output is sorted on the basis of the number of I/O operations.

D016ZE04 SM2 DEVICE TD		CYCLE: 60 S		<date> <time>	
MN	DEVICE TYPE	READ (IO/S)	WRITE (IO/S)	DATA (KB/S)	(BYTES/IO)
U4	ZAS-LAN	0.0	1712.5	167.3	97.6
U5	ZAS-LAN	1710.4	0.0	7969.1	4659.2
RP	ZAS-LAN	0.3	0.0	0.0	53.3
CI	SKP2	0.2	0.0	0.0	0.0
CH	SKP2	0.2	0.0	0.0	0.0
Y4	ZAS-BCAM	0.0	0.0	0.0	0.0
Y3	ZAS-BCAM	0.0	0.0	0.0	0.0
64AD	ZAS-LAN (C)	0.0	0.0	0.0	0.0
RO	ZAS-LAN	0.0	0.0	0.0	46.2
XD	ZAS-BCAM	0.0	0.0	0.0	
YW	ZAS-DUMP	0.0	0.0	0.0	
XE	ZAS-DUMP	0.0	0.0	0.0	
Y0	ZAS-BCAM	0.0	0.0	0.0	
Y1	ZAS-BCAM	0.0	0.0	0.0	
YR	ZAS-BCAM	0.0	0.0	0.0	
YS	ZAS-BCAM	0.0	0.0	0.0	
YX	ZAS-BCAM	0.0	0.0	0.0	

Monitoring information

- MN Mnemonic name of the connection of the device
- DEVICE TYPE Short name for the device type.
openCRYPT-BOX devices are marked with a '(C)'
- READ (IO/S) Number of read operations in the last monitoring interval
- WRITE (IO/S) Number of write operations in the last monitoring interval
- DATA Transferred data in number of KB per second or bytes per I/O operation

DISK FILE report

This report contains information on the input/output operations on all files of selected disk devices.

Report output

The REPORT DISK_FILE statement is used to request the output of the DISK FILE report. The report can only be output if the SM2 administrator has

- defines the disk devices to be monitored using the SET-DISK-FILE-PARAMETERS statement
- started the monitoring process using the statement START-MEASUREMENT-PROGRAM TYPE=*DISK-FILE

The statistics for each disk devices are output in a separate screen.

All files are displayed to the privileged user. Nonprivileged users only get to see files with their own ID.

Overflow screens can be requested for this report. The output of the monitored data is sorted on the basis of READ and WRITE.

D016ZE04 SM2 DISK FILE			CYCLE: 60 S	<date> <time>
READ	WRITE	FILENAME	(MN: B3A3, VSN: 6VS1.1)	
0.0	0.0	*OVERRUNS		
0.4	4.3	:20S6:\$NEI.DOCLST.SM2.17.0A00		
3.0	0.0	:20S6:\$NEI.LIB.P		
1.1	1.6	:20S6:\$NEI.INTLIB.SM2.17.0A00		
0.6	0.0	:20S6:\$NEI.LIB.SM2.TESTPROCS		

Monitoring information

READ Number of read operations per second

WRITE Number of write operations per second

FILENAME Name of the monitored file.
 The first row always contains the value *OVERRUNS. This entry includes all input/output operations that could not be assigned to a file entry because the internal SM2 table was full during the monitoring cycle. It is generally set to zero or to a low value. In the FILENAME column, not all names are in the format: <catid>:<\$userid>.<filename>. In the case of special disk accesses (e.g. by DAB), this format is not applied.

MN Mnemonic device name of the monitored disk device

VSN VSN of the monitored disk device

DLM report

This report contains monitored data from the DLM (Distributed Lock Manager).

Report output

The REPORT DLM statement is used to request the output of the DLM report.

The report can only be output if the SM2 administrator has activated the monitoring process using the statement START-MEASUREMENT-PROGRAM TYPE=*DLM.

D016ZE04 SM2 DLM		CYCLE: 60 S	<date> <time>	
		TU	TPR	NSM
NUMBER ENQUEUE	(1/S)	0.0	8.6	8.5
NUMBER CONVERT	(1/S)	0.0	67.5	67.4
NUMBER DEQUEUE	(1/S)	0.0	8.6	8.5
NUMBER INFORMATION	(1/S)	0.0	1.1	0.0
NUMBER GRANT EVENTS	(1/S)	0.0	0.0	8.5
NUMBER RELEASE EVENTS	(1/S)	0.0	0.0	0.0

Monitoring information

The individual entries have the following meanings:

NUMBER ENQUEUE (1/S)

Number of enqueue lock requests per second from TU, TPR or NSM

NUMBER CONVERT (1/S)

Number of convert lock requests per second from TU, TPR or NSM

NUMBER DEQUEUE (1/S)

Number of dequeue lock requests per second from TU, TPR or NSM

NUMBER INFORMATION (1/S)

Number of information lock requests per second from TU or TPR. The "NSM" column always contains the value 0.

NUMBER GRANT EVENTS (1/S)

Number of grant events (information about lock allocations)
per second

via TU contingency,
via bourse (or user eventing),
from NSM (information about lock allocations on the local host).

NUMBER RELEASE EVENTS (1/S)

Number of release events (requests for lock releases)
per second

via TU contingency,
via bourse (or user eventing),
from NSM (requests for lock release from remote hosts).

FILE reports

The FILE reports contain monitored data on access to a file.

Report output

The REPORT FILE statement is used to request the output of the FILE reports.

The FILE report consists of the following sections:

- USER FILE report for the FILE user-specific monitoring program
- PRIVILEGED FILE report for the FILE monitoring program FILE (for privileged users only)

The USER FILE report can only be output if the SM2 administrator has

- permitted the FILE user-specific monitoring program using the statement MODIFY-USER-ADMISSION FILE=*ALLOW and
- included files in the monitoring process using the FILE statement

The PRIVILEGED FILE report can only be output if the SM2 administrator has

- defined files using the ADD-FILE statement
- started file monitoring using the statement START-MEASUREMENT-PROGRAM TYPE=*FILE

Outputting monitored data in the FILE reports

In the **USER FILE report**, the nonprivileged user receives the values for the files of the FILE user-specific monitoring program he or she has included.

In the **PRIVILEGED FILE report**, the privileged user is shown the values for all the files included in the monitoring process by an SM2 administrator for the monitoring program. These values are also written to the SM2 output file.

Overflow screens may be requested for this report.

The layout of the FILE reports for nonprivileged and privileged users is identical, except for the headers (the former are entitled SM2 USER FILE REPORT, the latter are entitled SM2 PRIVILEGED FILE REPORT).

D016ZE04 SM2 PRIVILEGED FILE					CYCLE: 60 S	<date> <time>
WAIT	CHECK	READ	WRITE	TIME	FILENAME	
1.5	0.0	0.0	1.5	2.23	:20SH:\$TSOS.TSOSCAT	
0.0	0.0	0.0	0.0	*	:2S06:\$TSOS.TSOSCAT	
0.0	0.0	0.0	0.0	*	:20S7:\$TSOS.TSOSCAT	

Monitoring information

WAIT The number of WAIT operations per second is incremented if:

- reading from the file to main memory was initiated and the task is waiting for the end of the read operation, or
- writing from main memory to the file was initiated and the task is waiting for the end of the write operation.

WAIT operations not combined with read or write operations are ignored.

CHECK The number of CHECK operations per second is incremented whenever a check is made to see whether the preceding write operation was performed properly (read-after-write check). Only those read-after-write checks are counted which were initiated together with a write operation.

READ Number of read operations per second

WRITE Number of write operations per second

TIME Average duration of an input/output operation in milliseconds;
* means that no input/output operation has taken place.



The value shown here represents the logical duration of the input/output operation from initiation to completion of the job from the point of view of the software. In the case of asynchronous input/output operations, this may be considerably greater than the hardware time.

The < sign indicates that the value cannot be fully accommodated in the field.

FILENAME Name of the monitored file

In the case of composite PAM operations such as Read and Wait for Termination, both values are incremented (READ and WAIT).

GLOBAL report

The GLOBAL report contains data on a host network. It provides the user with an overview of the most important system activities (CPU utilization, I/O rates, size of pageable memory) for each selected host.

Report output

The REPORT GLOBAL statement is used to request the output of the GLOBAL report.

The privileged user can use the SELECT-HOSTS statement to select the hosts to be displayed in the GLOBAL report. If these include the local host, this is output in the first line, followed by the remaining hosts in the order specified in the SELECT-HOSTS statement. The header contains information on the first host for which valid data is found. If valid information cannot be found on a particular host, the string "***" appears in place of the monitored data.

Only the local computer is displayed to the nonprivileged user.

Overflow screens may be requested for this report.

For further information, please refer to the description of the SELECT-HOSTS statement on [page 128](#).

D016ZE04 SM2 GLOBAL		CYCLE: 60 S					<date> <time>			
HOST	LM	CPU UTILIZATION (%)				MEMORY NPP	IO'S (1/S)			
		TU	TPR	SIH	IDLE		TOTAL	DISK	PAGE	
D016ZE04	2	4.8	23.1	4.3	67.6	819183	800.5	578.8	0.0	
D016ZE07	2	46.9	2.6	1.6	48.8	441327	474.6	397.5	0.0	
D016ZE15	3	4.0	21.4	7.9	66.5	450743	224.1	130.5	0.0	

Monitoring information

The monitored data displayed here is the same as that of the ACTIVITY report.

GS report

This report supplies partition-specific monitored data on the utilization of global storage.

Report output

The REPORT GS statement is used to request the output of the GS report.

The report can only be output if the SM2 administrator has activated the monitoring process using the statement START-MEASUREMENT-PROGRAM TYPE=*GS.

Overflow screens may be requested for this report.

The monitored data output is sorted in accordance with the total number of READ and WRITE operations.

D017ZE03 SM2 GS			CYCLE: 60 S			<date> <time>		
! PARTITION !	! UNIT !	!	! READ !	!	! WRITE !	!	!	!
!	!	!	(1/S)	!	(KB/S)	!	(1/S)	!
!	!	!	!	!	!	!	!	!
!	!	!	!	!	!	!	!	!
! DAB120SH !	! 1 !	!	29.1	!	45.4	!	1.4	!
! GSTRACE !	! DUAL !	!	0.0	!	0.0	!	0.0	!
! DAB220SH !	! 2 !	!	0.0	!	0.0	!	0.0	!

Monitoring information

PARTITION	Name of the partition in global storage
UNIT	Unit on which the partition lies. A dual partition consists of a contiguous address space on both GS units, whose start addresses are differentiated by the HW constant "DUPLICATION-BOUNDARY".
READ	Number of read operations per second and number of KB transferred per second
WRITE	Number of write operations per second and number of KB transferred per second

GSVOL report

This report provides monitored data on emulated volumes in global storage.

Report output

The REPORT GSVOL statement is used to request the output of the GSVOL reports. This report can only be output if the SM2 administrator has activated the monitoring process using the statement START-MEASUREMENT-PROGRAM TYPE=*GSVOL.

Overflow screens may be requested for this report.

The monitored data output is sorted under the categories READ and WRITE.

D017ZE03 SM2 GSVOL		CYCLE: 60 S		<date> <time>				
!	DEVICE-MN	!	UNIT	!	READ (1/S)	!	WRITE (1/S)	!
!	0381	!	1	!	822.2	!	355.2	!
!	0380	!	1	!	765.2	!	336.4	!
!	0382	!	1	!	787.3	!	290.3	!
!	0383	!	1	!	0.2	!	0.1	!

Monitoring information

DEVICE-MN Mnemonic device name of the emulated volume

UNIT GS unit on which the volume was created
 A dual partition consists of a contiguous address space on both GS units, whose start addresses are differentiated by the HW constant "DUPLICATION-BOUNDARY".

READ (1/S) Number of read operations per second

WRITE (1/S) Number of write operations per second

When comparing with the values of the GS report, note that in the GSVOL report the I/O operations refer to the emulated volumes. One I/O operation can hide several GS accesses (additional key, block length).

ISAM reports

The ISAM reports provide monitored data for ISAM pools which were specified using the pool name, catalog ID and scope. Monitoring covers all files contained in the pool.

File-specific analyses within global ISAM pools in the data space can be output in the ISAM FILE report (see [page 268](#)).

Report output

The REPORT ISAM statement is used to request the output of the ISAM report.

The ISAM report consists of the following sections:

- USER ISAM report for the ISAM user-specific monitoring program
- PRIVILEGED ISAM report for the ISAM monitoring program (for privileged users only).

The USER ISAM report can only be output when the SM2 administrator

- has used the statement MODIFY-USER-ADMISSION ISAM=*ALLOW to permit the ISAM user-specific monitoring program and
- included ISAM pools in the monitoring process using the START-/CHANGE-ISAM-STATISTICS statement

The PRIVILEGED ISAM report can only be output if the SM2 administrator has

- defined ISAM pools using the ADD-ISAM-POOL statement
- started ISAM pool monitoring using the START-MEASUREMENT-PROGRAM TYPE=*ISAM statement

Outputting monitored data in the ISAM reports

In the **USER ISAM report**, the nonprivileged user receives the values of the ISAM pool of the ISAM user-specific monitoring program that he or she has included.

In the **PRIVILEGED ISAM report**, the privileged user is shown the values for all the ISAM pools included in the monitoring process by an SM2 administrator for the ISAM monitoring program using ADD-ISAM-POOL. These values are also written to the SM2 output file.

The layout of the ISAM reports for nonprivileged users and for privileged users is identical, except for the header (report entitled SM2 USER ISAM REPORT for nonprivileged users, and SM2 PRIVILEGED ISAM REPORT for privileged users).

D016ZE04 SM2 PRIVILEGED ISAM				CYCLE: 60 S				<date> <time>		
POOLNAME	SC	CAT-ID	USER-ID/ TSN	OP	FIX HIT%	WAIT%	SLOT WAIT%	INDEX OP	HIT%	SIZE
SRPMPPOOL	TA	20SH	RP01	34.9	100.0	0.0	0.0	23.3	100.0	138
\$SYS01	HS	20SH								

Monitoring information

POOLNAME	Name of the monitored ISAM pool
SC	Identifies the relevant ISAM pool as global (HOST-SYSTEM, HS) or task-local (TASK, TA)
CAT-ID	Catalog ID of the pubset specified for the monitored ISAM pool
USER-ID/TSN	Task sequence number (TSN) of the task under which a task-local ISAM pool was created. This column remains empty for global ISAM pools.
FIX	
OP	Number of FIX operations per second (number of FIX hits + number of FIX waits + number of FIX IOs). This also contains the number of accesses to the index.
HIT%	Percentage of FIX operations where the desired buffer page is already in the ISAM pool (i.e. no read operation from disk is required) relative to the total number of all FIX operations
WAIT%	Percentage of FIX operations requiring a waiting period until one or more buffer pages are released relative to the total number of all FIX operations
SLOT	
WAIT%	Percentage of RESERVE-SLOT operations leading to a wait state for the requested task due to a slot bottleneck relative to the total number of all RESERVE-SLOT operations

INDEX

OP	Number of accesses to the index per second
HIT%	Percentage of all accesses to the index accounted for by accesses to the index where the index searched for is already in the ISAM pool
SIZE	Size of the ISAM pool in 2 KB pages (excluding some administration data)

If an ISAM pool is never addressed during a monitoring cycle, the last seven columns contain blanks.

ISAM FILE report

The ISAM FILE report provides performance data on NK-ISAM file of ISAM pools in the Data Space.

The ISAM pools specified using a pool name, catalog ID and scope can be output in the ISAM report (see [page 265](#)).

Report output

The REPORT ISAM_FILE statement is used to request the output of the ISAM FILE report.

The ISAM FILE report can only be output if the SM2 administrator has

- defined the name of a NK-ISAM file using the ADD-ISAM-FILE statement
- started ISAM pool monitoring using the START-MEASUREMENT-PROGRAM TYPE=*ISAM statement

The privileged user receives values for all measured ISAM files.

The nonprivileged user receives values only on the measured ISAM files belonging to their user ID.

D016ZE04 SM2 ISAM FILE		CYCLE: 60 S				<date> <time>	
FILENAME	OP	FIX HIT%	WAI%	SLOT WAI%	INDEX OP	HIT%	SIZE
:20S6:\$NEI.SM2.ISAMFILE.1	1.4	72.2	0.0	0.0	0.7	97.4	512

Monitoring information

FILENAME	Name of the NK-ISAM file. It is possibly shortened spent; there is output a '>' as last character.
FIX	
OP	Number of FIX operations per second (number of FIX hits + number of FIX waits + number of FIX IOs). This also contains the number of accesses to the index.
HIT%	Percentage of FIX operations where the desired buffer page is already in the ISAM pool (i.e. no read operation from disk is required) relative to the total number of all FIX operations
WAI%	Percentage of FIX operations requiring a waiting period until one or more buffer pages are released relative to the total number of all FIX operations
SLOT	
WAI%	Percentage of RESERVE-SLOT operations leading to a wait state for the requested task due to a slot bottleneck relative to the total number of all RESERVE-SLOT operations
INDEX	
OP	Number of accesses to the index per second
HIT%	Percentage of all accesses to the index accounted for by accesses to the index where the index searched for is already in the ISAM pool
SIZE	Size of the ISAM pool in 2 KB pages (excluding some administration data)

MEMORY report

This report provides an overview of main memory and virtual address space utilization.

Report output

The REPORT MEMORY statement is used to request the output of the MEMORY report.

```

D016ZE04 SM2 MEMORY          CYCLE:  60 S  SAMPLES:  144      <date> <time>
                                MAIN MEMORY FRAMES                BIG PAGES
+-----+-----+-----+-----+-----+-----+-----+-----+
| TOTAL | PAGEABLE | R-ONLY-Q | RD-WR-Q | EMPTY-Q | SYS-GLOB | TOTAL | USED |
+-----+-----+-----+-----+-----+-----+-----+-----+
| 983040 | 820838 |      285 |      642 | 233348 | 586562 |      0 |      0 |
+-----+-----+-----+-----+-----+-----+-----+-----+

                                # PAGES IN VIRTUAL MEMORY        WORK. SET (PPC)        USED PAGES
+-----+-----+-----+-----+-----+-----+-----+-----+
| CL1 | CL2 | CL3 | CL4 | CL4-S | ALL | ACT | ACT |
+-----+-----+-----+-----+-----+-----+-----+-----+
| 1031 | 2680 | 14407 | 45616 |      80 | 57190 | 26972 | 367537 |
+-----+-----+-----+-----+-----+-----+-----+-----+

                                PAGING AREA FRAMES                PAGE FAULTS                PAGE TRANSFERS
+-----+-----+-----+-----+-----+-----+-----+-----+
| TOTAL | USED | TOTAL | 1ST-READ | RECLAIMS | READ | WRITE |
+-----+-----+-----+-----+-----+-----+-----+-----+
| 3840006 | 864129 | 946.4 | 946.4 | 0.0 | 0.0 | 0.0 |
+-----+-----+-----+-----+-----+-----+-----+-----+
    
```

Monitoring information

MAIN MEMORY FRAMES

- TOTAL Total number of 4-Kb pages in main memory
- PAGEABLE Number of pageable 4-Kb pages in main memory
- R-ONLY-Q Average number of 4-Kb read-only pages in the free pool (read-only queue)
- R-WR-Q Average number of 4-Kb read/write pages in the free pool (read/write queue)
- EMPTY-Q Average number of 4-Kb pages in the free pool, which are not assigned to a virtual page
- SYS-GLOB Average number of system-globally managed 4-Kb page

BIG PAGES

TOTAL	Total number of Big Pages (on servers with x86 architecture only)
USED	Number of used Big Pages (on servers with x86 architecture only)

PAGES IN VIRTUAL MEMORY

CL1	Number of class 1 pages (4 Kb) in the virtual address space
CL2	Number of class 2 pages (4 Kb) in the virtual address space
CL3	Number of class 3 pages (4 Kb) in the virtual address space
CL4	Number of all class 4 pages (4 Kb) in the virtual address space
CL4-S	Number of class 4 pages (4 Kb) for shareable modules in the virtual address space; CL4-S is included in CL4

WORK. SET (PPC)

ALL	Sum of planned page counts (PPCs) of all tasks in the system
ACT	Sum of planned page counts of all active tasks in the system (including the SM2 tasks) (Task queues 0-4)

USED PAGES

ACT	Sum of used page counts (UPG) of all active tasks in the system
-----	---

PAGING AREA FRAMES

TOTAL	Maximum number of pageable 4-Kb pages in all paging devices
USED	Number of pageable 4-Kb pages used in all paging devices

PAGE FAULTS

TOTAL	Number of page fault interrupts per second (without "genuine" paging errors)
1ST-READ	Number of page fault interrupts per second caused by the first access to a 4-Kb page
RECLAIMS	Number of page fault interrupts per second for which the addressed 4-Kb page is still in main memory

PAGE TRANSFERS

READ Number of 4-Kb pages read per second

WRITE Number of 4-Kb pages written per second



The sum of read and written pages does **not** match the PAGE value under IO'S PER SEC of the ACTIVITY report. Memory management tries to group up to 16 4-Kb page-out pages (WRITE) for one output operation. Each READ causes one 4-Kb page to be transferred, i.e. causes one input operation (see also page fault, [page 668](#) in the glossary).

MSCF report

This report contains monitored data from the MSCF subsystem.

Report output

The REPORT MSCF statement is used to request the output of the MSCF reports. The report can only be output if the SM2 administrator has activated the monitoring process using the statement START-MEASUREMENT-PROGRAM TYPE=*MSCF.

D016ZE04 SM2 MSCF		CYCLE: 60 S		<date> <time>		
			# SERVER TASKS		TASK LIMITS	
CALLS	83.8 /S	ORIGINAL	4	MAXIMAL	50	
SHORTAGES	1.2 /S	ACTUAL	4	FLOW SET	37	
HOSTS	6	OCCUPIED	0	FLOW RESET	25	
REQUEST WITH REPLY			FLOW STATE			
AVG TIME	11 MS	AVG TIME	***** MS	# FLOWS	0.0 /S	
AVG WAIT	9 MS	ACTUAL STATE	NO FLOW			
# REQUESTS	82.8 /S					

Monitoring information

The entries have the following meanings:

CALLS Number of send jobs per second

SHORTAGES Number of send jobs rejected per second due to line overload

HOSTS Number of hosts entered in MSCF with which the local host maintains a connection (excluding the local host)

SERVER TASKS ORIGINAL Number of server tasks originally requested

ACTUAL Number of current server tasks

OCCUPIED Number of current server tasks which could cause a bottleneck

These values merely indicate the status at the end of the monitoring cycle and do not represent the mean value over the entire monitoring cycle.

TASK LIMITS

MAXIMUM	Maximum number of server tasks
FLOW SET	Limit value for the number of occupied server tasks as of which the FLOW state is set
FLOW RESET	Limit value for the number of occupied server tasks as of which the FLOW state is reset

The MSCF subsystem sets its own FLOW or NO FLOW state, depending on the number of occupied server tasks. If the FLOW state is set, the maximum number of server tasks can be exceeded.

For further information on server tasks and their limit values, see the "HIPLEX MSCF" manual [8].

REQUEST WITH REPLY

AVG TIME	Average total time of REQUEST WITH REPLY jobs from the beginning of the job to the first reply, in milliseconds
AVG WAIT	Average wait time with REQUEST WITH REPLY jobs for the first reply from the receiving host, in milliseconds
# REQUESTS	Number of REQUEST WITH REPLY jobs per second

FLOW STATE

AVG TIME	Average dwell time in the FLOW state; only the transitions from FLOW to NO FLOW are recorded
# FLOWS	Number of changes from the FLOW state to the NO FLOW state per second
ACTUAL STATE	Current state: FLOW or NO FLOW

For the average times, ***** is output if no events have occurred.

NSM report

This report provides data from the NSM (Node Synchronization Manager) subsystem for a computer network.

Report output

The REPORT NSM statement is used to request the output of the NSM report.

The report can only be output if the SM2 administrator has activated the monitoring process using the statement START-MEASUREMENT-PROGRAM TYPE=*NSM.

The privileged user can use the SELECT-HOSTS statement to select the hosts to be displayed in the NSM report. If these include the local host, this is output in the first line, followed by the remaining hosts in the order specified in the SELECT-HOSTS statement. The header contains information on the first host for which valid data is found. If valid information cannot be found on a particular host, the string “***” appears in place of the monitored data.

Only the local host is displayed to the nonprivileged user.

Overflow screens may be requested for this report.

For further information, please refer to the description of the SELECT-HOSTS statement on [page 128](#).

```

D016ZE04 SM2 NSM                CYCLE:   60 S                <date> <time>
+-----+-----+-----+-----+-----+-----+
| MAX MSG LENGTH:   5 KB | USED MSG LENGTH:  0.7 % |
+-----+-----+-----+-----+-----+-----+
| CIRC TIME:       53.3 MS |
+-----+-----+-----+-----+-----+-----+
| HOST | DURATION | #WAIT | # LOCK- | TOTAL REQ | LOCAL REQ |
|      | TIME (MS)| MSG   | SERVER  | (1/S)     | (%)       |
+-----+-----+-----+-----+-----+-----+
| D016ZE04 | 1.0 | 0.0 | 59 | 72.0 | 91.7 |
| D016ZE07 | 1.0 | 0.0 | 11 | 0.0  | ***** |
| D016ZE15 | 1.0 | 0.0 | 5  | 6.0  | 91.0 |
+-----+-----+-----+-----+-----+-----+

```

Monitoring information

MAX MSG LENGTH Length of the NSM message buffer in the token in KB

USED MSG LENGTH Percentage of the NSM message buffer used

CIRC TIME Circulation time of the token in milliseconds

This data is output by the first host.

HOST	Name of the host on which the data is collected
DURATION TIME (MS)	Dwell time of the token in the host in milliseconds
#WAIT MSG	Number of messages waiting for space in the token
#LOCKSERVER	Number of lock servers. Each lock server represents a resource. It is created at the first enqueue and released at the last dequeue.
TOTAL REQ (1/S)	Total number of requests (enqueue, dequeue, convert, cancel) per second
LOCAL REQ (%)	Percentage of all requests that are local requests. A local request is a request in which the lock server already exists and is located on the local host.

OPENFT report

This report provides load values for openFT instances.

Report output

The REPORT OPENFT statement is used to request the output of the OPENFT report.

The OPENFT report can only be output if

- openFT has been prepared for data to be sent to SM2 (see [page 61](#))
- the SM2 administrator has defined openFT instances using the ADD-OPENFT-INSTANCE statement
- the SM2 administrator has activated the monitoring process using the START-MEASUREMENT-PROGRAM TYPE=*OPENFT statement

The user is shown a screen for each openFT instance which is to be monitored (max. 16).

BALU2	SM2	OPENFT	CYCLE:	60 S	<date>	<time>
INSTANCE: STD						
THROUGHPUT			STATUS			
NET DATA SENT	(KB/S)	695.5	SYNC REQ ACTIVE	0		
NET DATA RECEIVED	(KB/S)	331.7	ASYN REQ ACTIVE	1		
NET DATA TOTAL	(KB/S)	1035.2	REQ WAIT	0		
			REQ HOLD	0		
DISK DATA READ	(KB/S)	760.9	REQ SUSPEND	0		
DISK DATA WRITTEN	(KB/S)	177.3	REQ LOCK	0		
DISK DATA TOTAL	(KB/S)	938.2				
REQ ACCEPTED	(/S)	0.0	CONFIGURATION			
REQ SUCCESSFUL	(/S)	0.0	PAR SELECT	ALL		
REQ ABORTED	(/S)	0.0	REQ SELECT	ALL		
REQ INTERRUPTED	(/S)	0.0				
REQ AUTH ERROR	(/S)	0.0				
CON FAILED	(/S)	0.0	CON LIM/USED%	16 / 6		
CON ABORTED	(/S)	0.0	REQ LIM/USED%	2000 / 0		

Monitoring information

INSTANCE	Name of the openFT instance
NET DATA SENT	Number of KB sent per second for send jobs
NET DATA RECEIVED	Number of KB sent per second for receive jobs
NET DATA TOTAL	Number of KB sent per second
DISK DATA READ	Number of KB read from files per second

DISK DATA WRITTEN	Number of KB written to files per second
DISK DATA TOTAL	Number of KB read from files or written to files per second
REQ ACCEPTED	Number of jobs received per second
REQ SUCCESSFUL	Number of successfully completed jobs per second
REQ ABORTED	Number of aborted jobs per second
REQ INTERRUPTED	Number of interrupted jobs per second
REQ AUTH ERROR	Number of jobs rejected per second because of incorrect user identification
CON FAILED	Number of unsuccessful connection attempts per second
CON ABORTED	Number of connections aborted per second
SYNC REQ ACTIVE	Current number of synchronous jobs in ACTIVE state
ASYN REQ ACTIVE	Current number of asynchronous jobs in ACTIVE state
REQ WAIT	Current number of jobs in WAIT state
REQ HOLD	Current number of jobs in HOLD state
REQ SUSPEND	Current number of jobs in SUSPEND state
REQ LOCK	Current number of jobs in LOCK state
PAR SELECT	Selection of the partners who are to be included in monitoring (see openFT command MODIFY-FT-OPTIONS). Possible outputs: – OPENFT – FTAM – FTP – ALL
REQ SELECT	Selection of the jobs which are to be included in monitoring (see openFT command MODIFY-FT-OPTIONS). Possible outputs: – LOC (jobs issued locally) – REM (jobs issued in the remote system) – SYNC (synchronous jobs) – ASYNC (asynchronous jobs) – ALL (all jobs)
CON LIM	Maximum number of connections active in parallel which can be used to execute the file transfer requests (openFT parameter CONNECTION-LIMIT)

USED%	Ratio of the occupied connections to the maximum number of connections which are active in parallel. This value can temporarily be greater than 100 if CONNECTION-LIMIT has just recently been reduced.
REQ LIM	Maximum number of jobs which can be saved in the request queue (asynchronous jobs, openFT parameter REQUEST-LIMIT)
USED%	Ratio of the jobs stored in the request queue to the maximum number of jobs which can be stored in the request queue (asynchronous jobs). This value can temporarily be greater than 100 if REQUEST-LIMIT has just recently been reduced.



The monitored values in the case of USED% always reflect the status at the end of the monitoring cycle, but not an average value over the monitoring cycle.

PCS report

This report provides an overview of the activities of the Performance Control Subsystem.

Report output

The REPORT PCS statement is used to request the output of the PCS report. The report can only be output if PCS was active during the last monitoring cycle.

D016ZE04 SM2 PCS			CYCLE: 60 S		SAMPLES: 144		<date> <time>	
CATEGORY NAME	SRACT (SU/S)			SQPLN (%)	SQACT (%)	RDACT	DUR-RO (1/S)	DUR-RO-P (1/S)
	TOTAL	IO	CPU					
SUM	91232	39521	46741		100.0	1.1	0.138	0.000
SYS	44593	28735	15619	0.0	48.8	1.0	0.000	0.000
DIALOG	0	0	0	20.0	0.0	**	0.000	0.000
BATCH	19357	2170	14853	4.4	21.2	1.0	0.000	0.000
TP	4203	101	3580	8.4	4.6	1.5	0.000	0.000
TP1	0	0	0	22.6	0.0	**	0.000	0.000
BATCHDB	0	0	0	22.6	0.0	**	0.000	0.000
BATCHF	13929	7122	5438	0.4	15.2	1.0	0.000	0.000
DIALOG1	2381	733	1578	10.0	2.6	1.0	0.138	0.000
DIALOG2	6767	659	5672	11.2	7.4	1.0	0.000	0.000

Monitoring information

This report contains the totals (SUM column) for all categories (except for the SQPLN value) in the first line below the table header. The data for the individual categories (up to 16) is output in the following lines.

CATEGORY NAME	Name of a category
SRACT (SU/S)	Service unit per second of this category, specified as TOTAL, IO and CPU
TOTAL	Total number of service units per second of this category (sum of IO, CPU and memory service rates)
IO	IO service units per second of this category
CPU	CPU service units per second of this category
SQPLN (%)	System performance percentage planned for this category. (This variable is irrelevant for the SUM category and is therefore replaced by blanks.)

SQACT(%)	Current percentage of system performance for this category
RDACT	Current delay of all jobs in this category. If no dilation was sensed, ** is output (see section on “dilation factor” on page 665 of the glossary). (A system-global value is given in the SUM line.)
DUR-RO (1/S)	Number of DURATION time slice runouts for tasks in this category. (The given service rate has been used up, the task is placed in the defined continuation category.)
DUR-RO-P (1/S)	Number of DURATION time slice runouts with preemption for tasks or jobs in this category. (Meaning as above; since the defined continuation category is not permitted because of overloading, the task is entered in task queue 6.)

PERIODIC TASK report

This report provides information on the most important consumption statistics of all tasks.

Report output

The REPORT PERIODIC statement is used to request the output of the PERIODIC TASK report.

The report can only be output if the SM2 administrator has

- defined the tasks whose monitored data is to be written to the SM2 output file using the SET-PERIODIC-TASK-PARAMETERS statement
- activated the monitoring process using the statement START-MEASUREMENT-PROGRAM TYPE=*PERIODIC-TASK

Privileged users receive monitored data on all tasks.

Nonprivileged users receive monitored data only on the tasks belonging to their user ID.

Overflow screens may be requested for this report.

The monitored data output is sorted in accordance with the sort criterion defined in the SELECT-PERIODIC-TASK-PARAMETERS statement (see [page 187](#)). By default, it is sorted in accordance with the service units.

```

D016ZE04 SM2 PERIODIC TASK (S-U) CYCLE: 60 S <date> <time>
  TU + TPR= 35.0 %  SIH= 26.4 %  IDLE= 38.6 %  IO(1/S)= 3724.1  ( 2 LM'S)
  TSN | USER-ID | JOBNAME | T |SERV-RATE| CPU% | IO(1/S)| UPG | PAG RD| CRYPT
  -----|-----|-----|---|-----|-----|-----|-----|-----|-----
BCAM | | | S | 28898.1 | 0.2 | 3424.3 | 392 | 0.0 | 0.0
4H5N | NEI | ENNO | B | 18085.9 | 10.4 | 142.9 | 206 | 0.0 | 0.0
4HZF | TSOS | PALL301N | B | 13351.9 | 3.8 | 4.1 | 2767 | 0.0 | 0.0
BCAO | | | S | 10240.1 | 7.6 | 0.0 | 6 | 0.0 | 0.0
4IB1 | ZAC | E | D | 8629.4 | 5.0 | 98.8 | 2053 | 0.0 | 0.0
DM | | | S | 3208.2 | 2.3 | 11.4 | 23 | 0.0 | 0.0
4H1Y | TSOS | SHCUSERT | T | 1337.1 | 0.9 | 6.2 | 83 | 0.0 | 0.0
4H1X | TSOS | SHCUSERT | T | 1238.0 | 0.8 | 5.7 | 83 | 0.0 | 0.0
3RFE | NEI | INSPECTR | B | 941.3 | 0.5 | 0.0 | 3791 | 0.0 | 0.0
SM2G | | | S | 792.6 | 0.5 | 0.0 | 794 | 0.0 | 0.0
MSCF | | | S | 540.0 | 0.4 | 0.0 | 18 | 0.0 | 0.0
NSMS | | | S | 400.3 | 0.3 | 0.0 | 30 | 0.0 | 0.0
OHA7 | ORAQAX09 | MMNLQX9 | T | 396.9 | 0.2 | 0.0 | 3308 | 0.0 | 0.0
4HOJ | NEI | ENNO | D | 338.3 | 0.2 | 0.4 | 1257 | 0.0 | 0.0
4ICU | TSOS | RMF#4ICS | B | 296.3 | 0.1 | 0.4 | 2404 | 0.0 | 0.0
9DK8 | TSOS | HSMSSERV | B | 250.5 | 0.2 | 2.4 | 266 | 0.0 | 0.0
    
```

Monitoring information

The first header line shows the report name and the sort criterion used (CPU, IO, UPG or S-U for service units).

The second header line displays the following global system values (analogous to REPORT ACTIVITY):

TU+TPR	Percentage of the (TU+TPR+SIH+IDLE) time accounted for by the (TU+TPR) time
SIH	Percentage of the (TU+TPR+SIH+IDLE) time accounted for by the SIH time
IDLE	Percentage of the (TU+TPR+SIH+IDLE) time accounted for by the IDLE time
IO(1/S)	Total number of I/Os per second in the last cycle
(2 LM'S)	Number of logical machines which were active

The PERIODIC TASK report contains data for the tasks selected with the SELECT-PERIODIC-TASK-PARAMETERS statement (see [page 187](#)). By default, the data of all tasks is output.

TSN	TSN of the task
USER-ID	User ID of the task
JOB-NAME	Job name of the task
T	Current task type at the end of the monitoring cycle
SERV-RATE	Elapsed service units of the task per second
CPU%	Task's percentage of the (TU+TPR+SIH+IDLE) time (without the share of the hypervisor under VM)
IO(1/S)	Number of I/Os for the task per second
UPG	Mean UPG of the task
PAG RD	Number of page reads per second for the task
CRYPT	Number of physical accesses to coded files per second



Physical writing of pages to the paging disk is assigned to the system task with the TSN "PGE" for the purposes of consumption statistics.

PFA CACHE report

By means of the PFA monitoring program, SM2 records monitored data on the use of caches under User PFA.

The PFA CACHE report provides an overview of the use of all PFA cache areas of the product DAB. The data of the PFA CACHE report is also output in the DAB report.

Main memory (MM) or global storage (GS) are used as a cache medium.

Report output

The REPORT PFA statement is used to request the output of the PFA CACHE report. The report can only be output if the SM2 administrator has activated the monitoring process using the statement START-MEASUREMENT-PROGRAM TYPE=*PFA.

Overflow screens may be requested for this report.

The monitored data output is sorted under the categories READ and WRITE.

D016ZE04 SM2 PFA CACHE				CYCLE: 60 S		<date> <time>			
CACHE ID	CACHE MED	BLK SIZE	SIZE	READ (1/S)	RD HIT (%)	WRITE (1/S)	WR HIT (%)	OVERRUNS (1/S)	
10SY	MM	2 K	10 MB	0.0	0.0	0.0	0.0	0.0	
10SZ	MM	2 K	10 MB	0.0	0.0	0.0	0.0	0.0	
20SQ	MM	2 K	20 MB	0.0	0.0	0.0	0.0	0.0	

Monitoring information

CACHE ID Cache ID. This is identical to the catalog ID of the pubset served.

CACHE MED Cache medium MM: Main memory
GS: Global storage

BLK SIZE Block size of the cache

SIZE Size of the cache

READ (1/S) Number of read operations to hiperfiles per second

RD HIT (%) Percentage of hits in the total number of read operations to hiperfiles

WRITE (1/S) Number of write operations to hiperfiles per second

WR HIT (%) Percentage of fast write operations in the total number of write operations to hiperfiles

OVERRUNS (1/S) Number of unsuccessful attempts to use the cache due to overload per second



The number of read and write accesses is output to one decimal place. Read and write rates less than 0.05 per second are rounded down to 0.0 per second and hit rates are calculated.

POSIX report

Report output

The REPORT POSIX statement is used to request the output of the POSIX report. The report can only be output if the SM2 administrator has activated the monitoring process using the statement START-MEASUREMENT-PROGRAM TYPE=*POSIX.

```

D016ZE04 SM2 POSIX          CYCLE:  60 S                               <date> <time>

----- OPTION A -----
  IGET/S      NAMEI/S
    0.0        2.0

----- OPTION M -----
  MSG/S      SEMA/S
    0.0        2.1

----- OPTION B -----
  BREAD/S    BWRITE/S    LREAD/S    LWRITE/S    RCACHE %    WCACHE %
    0.0        0.0        0.1        0.1        100.0       100.0

  PREAD/S    PWRITE/S
    0.0        0.0

----- OPTION C -----
  SCALL/S    SREAD/S    SWRITE/S    FORK/S    EXEC/S    RCHAR/S
   35.6       4.3        0.3        0.0       0.0      986.4

  WCHAR/S
   22.7
    
```

Monitoring information

The individual entries have the following meanings:

OPTION A provides information on the use of system routines for file access:

- IGET/S Number of files determined by means of the Inode entry in the UFS (Berkeley "fast file system") file system
- NAMEI/S Number of times per second a path name is searched for in the file system

OPTION M provides information on message and semaphore activities:

- MSG/S Number of accesses to messages per second
- SEMA/S Number of semaphore activities per second

OPTION B provides information on buffer utilization:

BREAD/S, BWRITE/S Number of data transmissions per second between the system buffer and the hard disk or other block-oriented devices

LREAD/S, LWRITE/S Accesses to the system buffer per second

RCACHE, WCACHE Cache-memory hit percentage, i.e.
RCACHE 1-BREAD/LREAD (in percent)
WCACHE 1-BWRITE/LWRITE (in percent)

PREAD/S, PWRITE/S Number of physical data transmissions per second (raw device)

OPTION C provides information on system calls:

SCALL/S All types of system call per second

SREAD/S, SWRITE/S, FORK/S, EXEC/S
Specific system calls

RCHAR/S, WCHAR/S Characters transmitted by means of read() and write() system calls

PUBSET report

Information on the utilization of imported SF pubsets or volume sets is output in the PUBSET report.

Report output

The REPORT PUBSET statement is used to request the output of the PUBSET report.

The report can only be output if the SM2 administrator has activated the monitoring process using the statement START-MEASUREMENT-PROGRAM TYPE=*PUBSET.

Overflow screens may be requested for this report.

The monitored data output is sorted under the category USED.

D016ZE04 SM2 PUBSET		CYCLE: 60 S			<date> <time>	
CATID	TYP	# VOL	SAT LEV	CAPACITY GB	USED %	
2PFS	SF	1	0	2.3	98	
4VS1 (1004)	VS	18	0	41.7	98	
MID1	SF	11	1	25.5	97	
1CV4 (1004)	VS	32	0	74.1	96	
CVS7 (20SC)	VS	6	1	13.9	93	
CVS3 (20SC)	VS	6	0	13.9	92	
CVS6 (20SC)	VS	6	0	13.9	91	
GVS4 (20SG)	VS	6	0	13.9	91	
2RZV	SF	16	2	37.1	90	
QVS2 (10SQ)	VS	4	0	9.3	87	
2OPP	SF	2	1	4.6	84	
2OP1	SF	2	1	4.6	84	
2OP2	SF	2	1	4.6	84	
2OWI	SF	2	0	4.6	82	
CVS5 (20SC)	VS	6	0	13.9	79	
GVS1 (20SG)	VS	6	0	13.9	79	
QVS4 (10SQ)	VS	2	0	4.6	79	

Monitoring information

The individual entries have the following meanings:

CATID Identifier of an SF pubset or volume set. In the case of volume sets the ID of the SM pubset is subsequently output.

TYP Type of monitored object:
 - SF (SF subset)
 - VS (volume set)

VOL Number of volumes

SAT LEV	Saturation level (0-6)
CAPACITY (GB)	Capacity in GB
USED (%)	Utilization in %



Only data of imported SF pubsets or volume sets is output. Furthermore, only volumes without an allocation lock (see MODIFY-PUBSET-RESTRICTIONS) are taken into account.

RESPONSETIME report

This report contains data on response times, think times, transaction times and wait times for messages in the BCAM pool.

Report output

The REPORT RESPONSE statement is used to request the output of the RESPONSETIME report.

The report can only be output if the SM2 administrator has

- defined the monitoring parameters using the SET-RESPONSETIME-PARAMETERS statement
- specified the connection sets to be monitored using the ADD-CONNECTION-SET statement
- activated the monitoring process using the statement START-MEASUREMENT-PROGRAM TYPE=*RESPONSETIME

Depending on the definition of the SCOPE operand in the SET-RESPONSETIME-PARAMETERS or MODIFY-RESPONSETIME-PARAMETERS statement, SM2 output bucket-specific or category-specific screens.

In the case of SCOPE=*BUCKET, the user receives a separate screen for each of the connection sets to be monitored (up to 16, including global statistics) with bucket-specific values. Global statistics are output first.

In the case of SCOPE=*CATEGORY, the user receives a separate screen for each of the connection sets to be monitored (up to 16, including global statistics) with category-specific values. Global statistics are output first.

Categories for which no values have been specified are not output. If values are specified for more than 13 categories, the screen is scrolled upwards by the corresponding number of lines so that the topmost lines of the report are not visible on the screen.

If both SCOPE=*CATEGORY and SCOPE=*BUCKET are specified, a screen is output for bucket-specific values first, followed by one for category-specific values.

Bucket-specific screen

```

D016ZE04 SM2 RESPONSETIME (BUCKETS) CYCLE: 60 S <date> <time>
*GLOBAL APPLICATION: *ALL
PARTNER: *ALL
CONN-TYPE: REMOTE PROCESSOR: *ALL

```

RESP. TIME (1)		THINK TIME		TRANS. TIME		WAIT TIME	
BUCKET	COUNT	BUCKET	COUNT	BUCKET	COUNT	BUCKET	COUNT
0.5	4386	5.0	4393	0.5	4357	0.1	345
1.0	8	15.0	21	1.0	11	0.2	0
2.0	8	30.0	4	2.0	7	0.5	0
5.0	11	60.0	1	5.0	16	1.0	0
10.0	3	120.0	8	10.0	11	2.0	0
> 10.0	18	> 120.0	7	> 10.0	32	> 2.0	0

```

AVG SEC | 0.1 | 0.3 | 0.1 | 0.0
NR. INTER | 4416 | 4427 | 4402 | 345
RATE 1/SEC | 76.1 | 76.3 | 75.9 | 5.9

```

```

INPUT-LEN:570.7 OUTPUT-LEN:451.7

```

Monitoring information

For each screen, the associated connections or connection groups are output above the monitoring data, as well as the connection set name ("GLOBAL" is the name given in the example), when the ADD-CONNECTION-SET statement is defined).

CONN-TYPE	Types of connections taken into account during monitoring: REMOTE: Remote connections only LOCAL: Local connections only BOTH: Local and remote connections
RESP. TIME (1)	Response time data; the type of response time monitored (as specified via the DEFINITION operand of the SET- and MODIFY-RESPONSETIME-PARAMETERS statements) is output (1): Response time as per definition 1 (2): Response time as per definition 2
THINK TIME	Think time data
TRANS. TIME	Transaction time data
WAIT TIME	Wait time data (in the BCAM pool)
AVG SEC	Average duration of an interaction in seconds (without overflow data)

NR. INTER	Total number of interactions in the monitoring cycle (without overflow data)
RATE 1/SEC	Number of monitored interactions per second (without overflow data)
BUCKET	Corresponds to the bucket limits in seconds in which monitored data is arranged in accordance with duration. The last line of the columns (>) contains the overflow data. Data displayed there is not used for computing AVG SEC, NR. INTER and RATE 1/SEC.
COUNT	Corresponds to the number of interactive operations in the individual buckets

The SM2 administrators can modify the bucket limits and the number of buckets by means of the SET- and MODIFY-RESPONSETIME-PARAMETERS statements if the default options are not suitable. If no interactions were recorded in the last monitoring cycle, the message NOTE: SOME INTERACTIONS MISSED:<number> is displayed in the bottom line of each of the RESPONSE TIME screens. This may be the case, for example, if the number of active connections is greater than that defined in the CONNECTION-NUMBER parameter of the SET- or MODIFY-RESPONSETIME-PARAMETERS statement.

INPUT-LEN	The average lengths in bytes of the input or
OUTPUT-LEN	output messages

Category-specific screen

```

D016ZE04 SM2 RESPONSETIME (CATEGORY) CYCLE: 60 S <date> <time>
*GLOBAL APPLICATION: *ALL
PARTNER: *ALL
CONN-TYPE:REMOTE PROCESSOR: *ALL

```

CATGORY NAME	RESPONSE TIME			THINK TIME			TRANSACTION TIME			WAIT TIME		
	AVG	#INT	RATE	AVG	#INT	RATE	AVG	#INT	RATE	AVG	#INT	RATE
SUM	0.1	4416	76.1	0.3	4427	76.3	0.1	4402	75.9	0.0	345	5.9
SYS	0.1	78	1.3	8.9	85	1.5	0.1	4328	74.6	0.0	150	2.6
BATCH	0.1	1456	25.1	0.0	1457	25.1	0.0	0	0.0	0.0	0	0.0
TP	0.9	33	0.6	2.6	38	0.7	1.0	31	0.5	0.0	15	0.3
BATCHF	0.0	422	7.3	0.0	423	7.3	0.0	2	0.0	0.0	180	3.1
DIALOG1	0.1	310	5.3	0.7	306	5.3	0.2	41	0.7	0.0	0	0.0
DIALOG2	0.1	2117	36.5	0.1	2118	36.5	0.0	0	0.0	0.0	0	0.0

Monitoring information

For each screen the associated connections or connection groups are output over the monitoring values; the connection set name (in the example it is *GLOBAL) is output, too. (A connection set with the name *GLOBAL, in which all connections are monitored (CONNECTION (*ALL,*ALL,*ALL)), is automatically defined using the SET-RESPONSETIME-PARAMETERS statement.)

CONN-TYPE:	Types of connections taken into account during monitoring: REMOTE: Remote connections only LOCAL: Local connections only BOTH: Local and remote connections
CATGORY NAME	Name of the category
RESPONSETIME	Monitoring data for the response time
THINK TIME	Monitoring data for the think time
TRANSACTION TIME	Monitoring data for the transaction time
WAIT TIME	Monitoring data for the wait times (in the BCAM pool)
AVG	Average duration of the transaction (in seconds) (without overflow values)
#INT	Total number of transactions during the monitoring cycle (without overflow values)
RATE	Number of interactions monitored per second (without overflow values)



If the BCAM-CONNECTION monitoring program is started at the same time, incorrect values may be output for WAIT TIME.

SESAM SQL report

This report contains data of the SESAM/SQL database system.

Report output

The REPORT SESAM_SQL statement is used to request the output of the SESAM SQL report.

The SESAM SQL report can only be output if

- SESAM/SQL has been prepared for data to be sent to SM2 (see [page 71](#))
- the SM2 administrator has activated the monitoring process using the START-MEASUREMENT-PROGRAM TYPE=*SESAM-SQL statement.

Overflow screens may be requested for this report.

The monitored data output is sorted under the category DB CONF.

AILEEN2 SM2 SESAM SQL		CYCLE: 30 S				<date> <time>			
DB CONF	UPDTA (1/S)	RETTA (1/S)	RESTA (1/S)	PLGEN (1/S)	RDSYS (1/S)	WRSYS (1/S)	RDUSR (1/S)	WRUSR (1/S)	
SESAMEW	2.4	2.4	4.8	1.4	70.3	10.0	3837.2	7.4	

Monitoring information

DB CONF	Name of the DB configuration
UPDTA (1/S)	Number of update transactions per second
RETTA (1/S)	Number of retrieval transactions per second
RESTA (1/S)	Number of reset transactions per second
PLGEN (1/S)	Number of SQL plan generations per second
RDSYS (1/S)	Number of logical read accesses to the system data buffer
WRSYS (1/S)	Number of logical write accesses to the system data buffer
RDUSR (1/S)	Number of logical read accesses to the user data buffer
WRUSR (1/S)	Number of logical write accesses to the user data buffer



The monitored data is supplied by SESAM/SQL asynchronously to SM2 and applies for one or more cycles which are defined by SESAM/SQL and need not match the SM2 cycle. Here differences can exist in the duration of the cycles, and time displacements can also exist between the SESAM/SQL and SM2 cycles. The

duration of one or more SESAM/SQL cycles is used to standardize the monitored data to one second. The data is therefore exact, but it only matches the SM2 cycle to a certain degree.

SHARED PUBSET report

The SHARED PUBSET report contains monitored data on I/O operations to disk devices of shared pubsets for a computer network.

Report output

The REPORT SHARED_PUBSET statement is used to request the output of the SHARED PUBSET report.

The monitored data under SERVICETIME is only output when the SM2 administrator

- uses the SET-SAMPLING-DEVICE-PARAMETERS statement to define the monitoring of the service times and
- uses the statement START-MEASUREMENT-PROGRAM TYPE=*SAMPLING-DEVICE to start monitoring.

When the SM2 subsystem is started, the SAMPLING-DEVICE monitoring program is started without the service times being recorded.

The privileged user can use the SELECT-HOSTS statement (see [page 128](#)) to select the hosts whose monitored data is to be displayed in the report.

The header contains information on the local host (if this is selected) or on first host for which valid data is found. If valid information cannot be found on a particular host, the string `SOME DATA MISSED` appears in the second line.

Only the monitored data of the local host is displayed to the nonprivileged user.

For parallel access volumes (PAV), the following applies:

- Only basic devices are output. The monitored data refers to the basic device and the assigned alias devices (summation and averaging).
- For basic devices the average number of assigned alias devices is output. If a device is not a basic device, spaces are output.

Overflow screens may be requested for this report.

The monitored data output is sorted in accordance with the number of I/O operations.

D016ZE04 SM2 SHARED PUBSET CYCLE: 60 S SAMPLES: 75 <date> <time>										
MN	VOLUME	Q	LGTH	RSC	READ		WRITE		SERVICETIME	
					(IO/S)	(KB/S)	(IO/S)	(KB/S)	HW	SW
D114	CVS8.2	0.6	0.0		32.2	1045.4	395.8	797.1	1.2	1.2
D101	CVS3.3	0.1	0.1		106.2	1988.3	0.0	0.0		
D179	1DQM.4	0.0	0.0		71.3	285.2	0.0	0.0		
D32F	2RZV.F	0.0	0.0		24.8	589.2	0.0	0.0	0.7	0.7
D2F9	NVS1.3	0.0	0.0		20.9	53.4	2.2	6.6	0.4	0.4
D434	QVS2.0	0.0			0.0	0.0	20.0	40.0		
D2F6	NVS1.0	0.0	0.0		2.0	7.3	13.8	55.4	1.2	1.2
D2F5	1CVN.0	0.0	0.0		0.3	2.1	12.1	47.5	1.4	1.4
D112	CVS8.0	0.0	0.0		0.7	4.0	8.2	32.9	1.3	1.3
D2FA	NVS1.4	0.0	0.0		2.3	5.2	6.3	18.8	1.1	1.1
D0FC	CVS2.2	0.0	0.0		7.8	204.1	0.0	0.0	0.7	0.7
D113	CVS8.1	0.0	0.0		0.3	17.5	3.3	13.2	1.3	1.3
D110	CVS7.2	0.0	0.0		0.0	0.0	3.5	21.1	1.5	1.5

Monitoring information

The values correspond to those of the DEVICE DISK report and form the sum of the values from the individual hosts.

SM2 cannot automatically determine the percentage utilization of the disks (as in the DEVICE DISK report). Each host recognizes only its own accesses to the disk. The disk is already considered active when the I/O chaining starts, even if the system still has to wait until another host has completed its access. This can cause a dilation of the local utilization, which depends on the level of the local or remote loads and their distribution over time.

- MN Mnemonic name of the disk device
- VOLUME VSN of the disk device or spaces (e.g. if the disk belongs to a pubset that has not been imported)
- Q LGTH Average length of the device queue including the jobs that are currently being dealt with
- RSC Average number of parallel processed RSC IOs (servers with x86 architecture)
- READ Number of input operations or number of KB transferred per second
- WRITE Number of output operations or number of KB transferred per second
- SERVICETIME
- HW (MS) Average hardware service time in milliseconds
- SW (MS) Average software service time (including hardware service time) in milliseconds

SVC report

This report provides an overview of the SVC calls.

Report output

The REPORT SVC statement is used to request the output of the SVC report. The report can only be output if the SM2 administrator has activated the monitoring process using the statement START-MEASUREMENT-PROGRAM TYPE=*SVC.

The monitored data output is sorted in accordance with the total number of calls from TU and TPR. The maximum number of entries that can be output is the number that will fit on the screen. In addition, the sum is output for all SVCs and, if necessary, a value for SVCs that do not appear on this page. Only one screen is output.

D016ZE04 SM2 SVC			CYCLE: 60 S			<date> <time>		
SVC	TU	TPR SVC	TU	TPR SVC	TU	TPR SVC	TU	TPR
SUM	1767.7	5252.0 250	0.0	2782.2 234	0.0	1259.9 39	1025.0	1.6
246	0.0	476.7 69	209.2	0.0 47	43.0	157.1 215	0.0	186.6
188	74.1	64.8 170	6.9	126.1 1	69.4	4.6 183	63.7	0.6
214	0.0	64.0 241	0.0	49.2 146	34.9	2.3 44	35.8	0.0
168	15.5	14.8 123	28.1	0.0 128	22.5	0.0 186	0.8	19.2
124	15.9	0.0 175	15.4	0.0 34	11.9	0.0 125	11.7	0.0
89	4.8	4.8 152	8.7	0.0 151	8.5	0.0 171	0.0	8.4
161	0.1	7.9 85	7.0	0.0 144	5.5	1.2 40	6.5	0.0
7	3.0	2.4 213	0.0	5.4 35	4.0	0.0 187	3.9	0.0
172	3.7	0.0 58	3.7	0.0 229	0.0	3.0 135	2.7	0.0
121	2.5	0.0 45	2.5	0.0 228	0.0	2.3 38	0.4	1.4
96	1.4	0.3 66	1.6	0.0 191	1.3	0.1 145	1.0	0.4
72	1.3	0.0 67	1.2	0.0 159	1.2	0.0 140	0.0	1.1
222	0.0	1.1 102	0.9	0.0 73	0.9	0.0 70	0.8	0.0
83	0.7	0.0 245	0.0	0.7 233	0.0	0.6 9	0.5	0.0
52	0.5	0.0 156	0.4	0.0 190	0.0	0.4 88	0.4	0.0
162	0.4	0.0 43	0.4	0.0 74	0.3	0.0 80	0.3	0.0
158	0.0	0.2 16	0.0	0.2 28	0.2	0.0 RST	0.6	0.2

Monitoring information

- SVC SVC number
- TU Number of SVC calls in TU per second
- TPR Number of SVC calls in TPR per second
- SUM
- TU Sum of all SVCs per second from TU
- TPR Sum of all SVCs per second from TPR
- RST
- TU Remaining SVCs per second from TU
- TPR Remaining SVCs per second from TPR

SYMMETRIX reports

Data concerning the utilization of Symmetrix systems is output in the SYMMETRIX reports.

- The SYMMETRIX report provides summary data on all accessible Symmetrix systems.
- The SYMMETRIX PUBSET report provides monitored data on the pubsets of selected Symmetrix systems.
- The SYMMETRIX DEVICE report provides monitored data on the devices of selected pubsets on Symmetrix systems.

The reports can only be output if the SM2 administrator

- has defined the monitoring operands with the SET-STORAGE-SYSTEM-PARAMETERS statement
- has activated the monitoring process using the START-MEASUREMENT-PROGRAM TYPE=*STORAGE-SYSTEM statement.

The STORAGE-SYSTEM monitoring program requires the chargeable product SHC-OSD. The monitored data is recorded only for the directly accessible, generated controllers of the Symmetrix systems.

Unlike other SM2 monitoring programs, the information recorded on Symmetrix systems is not host-specific, but refers globally to one Symmetrix system. This must be taken into account particularly when interpreting the data.

The figure below illustrates the hierarchy of reports and the statement used to request them:

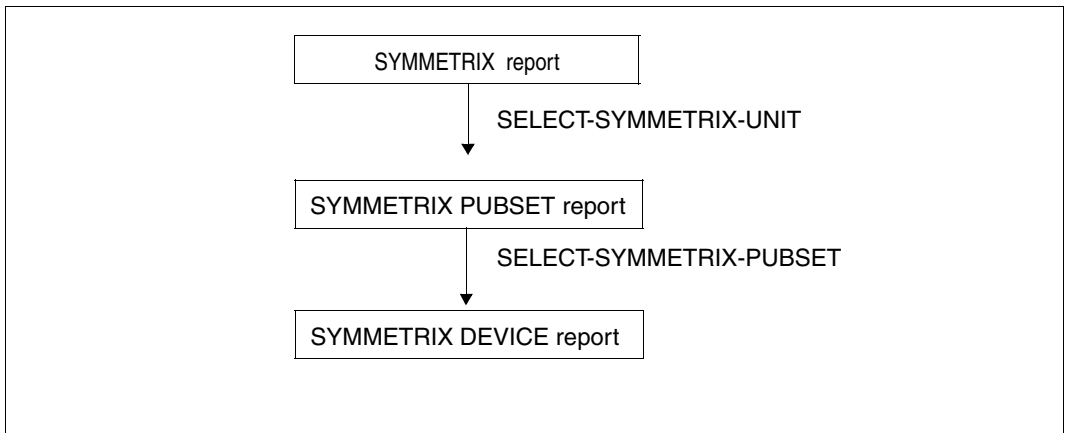


Figure 6: Hierarchy of Symmetrix reports

Every SM2 user can use the SELECT-SYMMETRIX-UNIT statement to select control units of Symmetrix systems for which a SYMMETRIX PUBSET report is then output. Only Symmetrix systems displayed in the SYMMETRIX report can be selected.

Every SM2 user can use the SELECT-SYMMETRIX-PUBSET statement to select pubsets for which a SYMMETRIX DEVICE report is then output. Only pubsets that are displayed in the SYMMETRIX PUBSET report can be selected. It is also possible to request SYMMETRIX DEVICE reports for the devices which are not assigned to a subset.

SYMMETRIX report

The SYMMETRIX report provides an overview of the I/O activity and the cache utilization of the generated control units of Symmetrix systems.

Report output

The REPORT SYMMETRIX statement is used to request the output of the SYMMETRIX report.

Overflow screens may be requested for this report.

The monitored data output is sorted under the categories READ and WRITE.

```

D016ZE04 SM2 SYMMETRIX                CYCLE: 60 S                <date> <time>
SERIAL-NO | A | SIZE |      READ      |      WRITE      | TOT
           |   | (GB) | (1/S) | HIT(%) | (KB/S) | (1/S) | HIT(%) | (KB/S) | HIT(%)
-----+---+-----+-----+-----+-----+-----+-----+-----+-----+
000287461761|D| 12.0 | 1151.1 | 99.2 | 8741.3 | 590.2 | 99.8 | 3786.3 | 99.4
000287460352|D| 16.0 | 646.5  | 99.4 | 4996.4 | 783.4 | 100.0 | 3547.6 | 99.7
-----+---+-----+-----+-----+-----+-----+-----+-----+-----+
    
```

Monitoring information

Serial-No	Serial number of the Symmetrix system Mnemonic name for each generated control unit
A	Access to the Symmetrix system – D (direct via an attached device)
SIZE	Cach size in GB
READ (1/S)	Number of read accesses per second
READ HIT (%)	Percentage of read operations completed by the cache in relation to the total number of read operations
READ (KB/S)	Number of KB transferred per second for read operations
WRITE (1/S)	Number of write accesses per second
WRITE HIT (%)	Percentage of write operations completed by the cache in relation to the total number of write operations
WRITE (KB/S)	Number of KB transferred per second for write operations

TOT HIT (%) Percentage of read and write operations completed by the cache in relation to the total number of access operations executed



The monitored data output here refers to the Symmetrix system as a whole, i.e. it includes access operations from other hosts.

SYMMETRIX PUBSET report

The SYMMETRIX PUBSET report provides monitored data on input/output activities and the cache usage of pubsets of a Symmetrix system. It is output for each Symmetrix system.

Report output

The SYMMETRIX PUBSET report is output only on the explicit request of the SM2 user with by means of the SELECT-SYMMETRIX-UNIT statement (see [page 190](#)). After issuing this statement, the user receives in addition to the SYMMETRIX report the SYMMETRIX PUBSET report for the selected Symmetrix systems.

A separate screen containing the monitored data for the pubsets is output for each Symmetrix system selected.

Overflow screens can be requested for this report.

The output of the monitored data is sorted on the basis of READ and WRITE.

```

D016ZE04 SM2 SYMMETRIX PUBSET      CYCLE:   60 S                <date> <time>
                                     SERIAL-NO: 000287460352
    PUBSET |      READ      |      WRITE      | TOT-H
    (1/S)  | HIT(%) | (KB/S) | (1/S) | HIT(%) | (KB/S) | (%)
-----|-----|-----|-----|-----|-----|-----
*NOID | 517.9 | 99.5 | 1944.3 | 561.1 | 100.0 | 1984.4 | 99.8
20DS  | 12.2 | 95.0 | 76.1  | 162.2 | 100.0 | 936.7  | 99.7
2RZV  | 57.6 | 99.6 | 1620.5 | 1.1  | 100.0 | 7.4    | 99.6
20RZ  | 33.4 | 100.0 | 552.1 | 1.0  | 100.0 | 2.9    | 100.0
40SQ  | 10.1 | 99.3 | 503.5 | 23.4 | 100.0 | 115.2  | 99.8
10SQ  | 13.1 | 99.2 | 275.2 | 5.7  | 100.0 | 212.2  | 99.4
20WI  | 1.1  | 59.1 | 11.2  | 13.1 | 100.0 | 111.2  | 96.7
0     | 0.1  | 100.0 | 1.1   | 4.3  | 100.0 | 135.3  | 100.0
10SZ  | 0.2  | 100.0 | 1.9   | 2.0  | 100.0 | 5.0    | 100.0
20SG  | 0.2  | 100.0 | 1.8   | 2.0  | 100.0 | 8.3    | 100.0
20SH  | 0.1  | 0.0   | 0.5   | 2.2  | 100.0 | 15.5   | 97.7
2ATS  | 0.2  | 100.0 | 4.3   | 1.4  | 100.0 | 3.5    | 100.0
32GB  | 0.1  | 100.0 | 1.4   | 1.0  | 100.0 | 2.4    | 100.0
20S7  | 0.1  | 100.0 | 0.6   | 1.0  | 100.0 | 2.5    | 100.0
    
```

Monitoring information

SERIAL-NO Serial number of the Symmetrix system

PUBSET ID of the pubset

*NOID is output for PUBSET if the value *NOID has been specified for SELECT-SYMMETRIX-PUBSET (for all devices that do not belong to a pubset).

READ (1/S)	Number of read accesses per second
READ HIT (%)	Percentage of read operations completed by the cache in relation to the total number of read operations
READ (KB/S)	Number of KB transferred per second for read operations
WRITE (1/S)	Number of write accesses per second
WRITE HIT (%)	Percentage of write operations completed by the cache in relation to the total number of write operations
WRITE (KB/S)	Number of KB transferred per second for write operations
TOT HIT (%)	Percentage of read and write operations completed by the cache in relation to the total number of access operations executed



The monitored data output here refers to the Symmetrix system as a whole, i.e. it includes access operations from other hosts.

SYMMETRIX DEVICE report

The SYMMETRIX DEVICE report provides monitored data on the I/O activities and the cache utilization of the devices of a Symmetrix system. It is output per pubset.

Report output

The SYMMETRIX DEVICE report is output only at the explicit request of the SM2 user with the help of the SELECT-SYMMETRIX-PUBSET statement (see [page 189](#)). This statement is used to output a SYMMETRIX report together with a SYMMETRIX DEVICE report for the selected pubsets. If the SYMMETRIX PUBSET report has been requested in addition, this is output after the SYMMETRIX report but before the SYMMETRIX DEVICE report.

A separate screen is output for each pubset selected.

Overflow screens may be requested for this report.

The monitored data output is sorted under the categories READ and WRITE.

```

D016ZE04 SM2 SYMMETRIX DEVICE      CYCLE:   60 S                <date> <time>
                                SERIAL-NO: 000287461761  PUBSET: 20S6

  SYM  |  MN  |  VSN  |  (1/S)  |  READ  |  (KB/S)  |  WRITE  |  (1/S)  |  (KB/S)  |  TOT-H  |
  DEV  |      |      |          |  HIT(%) |          |  HIT(%) |          |          |  (%)   |
-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|
3A7   | B3A7 | 6VS3.1 | 20.4   | 99.2   | 77.1     | 97.7    | 14.9    | 44.6     | 98.5   |
3A6   | B3A6 | 6VS3.0 | 8.9    | 99.6   | 25.5     | 98.4    | 21.0    | 161.1    | 98.8   |
3A4   | B3A4 | 6VS2.0 | 8.8    | 99.6   | 56.1     | 99.8    | 8.7     | 26.5     | 99.7   |
3A3   | B3A3 | 6VS1.1 | 8.7    | 98.0   | 44.3     | 99.0    | 6.8     | 61.8     | 98.4   |
3A0   | B3A0 | 6CV1.0 | 2.0    | 96.6   | 8.3      | 98.2    | 5.7     | 22.3     | 97.8   |
3A2   | B3A2 | 6VS1.0 | 3.0    | 98.9   | 22.0     | 96.3    | 1.4     | 9.9      | 98.0   |
3A8   | B3A8 | 6VS4.0 | 4.4    | 100.0  | 22.1     | 0.0     | 0.0     | 0.0      | 100.0  |
3A1   | B3A1 | 6CV1.1 | 0.6    | 78.1   | 2.6      | 96.9    | 0.6     | 1.1      | 87.5   |
3A5   | B3A5 | 6VS2.1 | 0.7    | 100.0  | 6.4      | 100.0   | 0.4     | 1.3      | 100.0  |
3A9   | B3A9 | 6VS4.1 | 0.0    | 0.0    | 0.0      | 0.0     | 0.0     | 0.0      | 0.0    |
3AA   | B3AA | 6VS5.0 | 0.0    | 0.0    | 0.0      | 0.0     | 0.0     | 0.0      | 0.0    |
3AB   | B3AB | 6VS5.1 | 0.0    | 0.0    | 0.0      | 0.0     | 0.0     | 0.0      | 0.0    |
    
```

Monitoring information

SERIAL-NO Serial number of the Symmetrix system

PUBSET ID of the pubset

*NOID is output for PUBSET if the value *NOID has been specified for SELECT-SYMMETRIX-PUBSET (for all devices that do not belong to a pubset).

SYM DEV Symmetrix system-internal number of device

MN Mnemonic name of the device (logical volume) or blanks

VSN	VSN of the device (logical volume) or blanks
READ (1/S)	Number of read accesses per second
READ HIT (%)	Percentage of read operations completed by the cache in relation to the total number of read operations
READ (KB/S)	Number of KB transferred per second for read operations
WRITE (1/S)	Number of write operations per second
WRITE HIT (%)	Percentage of write operations completed by the cache in relation to the total number of write operations
WRITE (KB/S)	Number of KB transferred per second for write operations
TOT HIT (%)	Percentage of read and write operations completed by the cache in relation to the total number of access operations executed



General notes on the monitored data

The monitored data above can differ from the values in the DEVICE DISK report because

- the data in the DEVICE DISK report comprises local data of the host and the data of the Symmetrix system comprises global data of the device (logical volume)
- the Symmetrix system does not always count every IO immediately
- there is a time lag between the gathering of the DEVICE DISK and the Symmetrix system data

TCP/IP report

This report provides the IP and port numbers as well as data on input and output for each TCP/IP connection. IPv4 as well as IPv6 connections are provided.

Report output

The REPORT TCP statement is used to request the output of the TCP/IP report. The report can only be output if the SM2 administrator has activated the monitoring process using the statement START-MEASUREMENT-PROGRAM TYPE=*TCP-IP.

Overflow screens may be requested for this report.

The monitored data output is sorted in accordance with the total number of bytes sent and received.

D016ZE04 SM2 TCP/IP		CYCLE: 60 S		<date> <time>				
	REMOTE IP		REMOTE PORT		LOCAL PORT			
1.	172.25.85.22		4444		4138			
2.	172.26.32.111		1975		20			
3.	172.25.81.9		102		4102			
4.	172.25.81.9		102		4163			
5.	172.25.81.30		102		4260			
6.	172.25.83.26		55579		28080			
7.	172.25.83.77		4432		1110			
		INPUT		OUTPUT				
	TSDU(/S)	(KB/S)	BUF (KB)	WIN (KB)	TSDU(/S)	(KB/S)	BUF (KB)	WIN (KB)
1.	296.1	6933.5	16.2	52.6	0.0	0.0	0.0	63.9
2.	4.7	94.0	3.1	26.4	0.0	0.0	0.0	63.9
3.	84.6	41.9	0.0	72.0	83.1	32.5	0.3	72.0
4.	21.3	3.5	0.0	22.8	0.0	0.0	0.0	22.0
5.	0.0	0.0	0.0	7.1	21.3	3.5	0.9	72.0
6.	0.1	0.0	0.0	8.1	0.0	1.6	0.0	255.9
7.	0.0	0.0	0.0	8.1	0.5	1.0	2.4	63.9

Monitoring information

The individual entries have the following meanings:

REMOTE IP	IP address of the remote host.
REMOTE PORT	Port number via which the application on the remote host communicates
LOCAL PORT	Port number via which the application on the local host communicates

INPUT

TSDU (/S)	Number of TSDUs received. TSDUs (Transport Service Data Units) correspond to jobs of the application on BCAM.
(KB/S)	Number of KB received
BUF (KB)	Cache space occupied by messages which have not yet been retrieved
WIN (KB)	Last window size received from the partner

OUTPUT

TSDU (/S)	Number of TSDUs sent TSDUs (Transport Service Data Units) correspond to jobs of the application on BCAM.
(KB/S)	Number of KB sent
BUF (KB)	Cache space occupied by messages which have not yet been sent or acknowledged
WIN (KB)	Last window size sent to the partner



The values for sent and received window sizes only apply for remote connections. The values BUF and WIN merely reflect the status at the end of the monitoring cycle, and do not represent mean values over the entire monitoring cycle.

TLM report

This report provides monitored data on the occupation of locks managed by the Task Lock Manager (TLM).

Report output

The REPORT TLM statement is used to request the output of the TLM report. The report can only be output if the SM2 administrator has activated the monitoring process using the statement START-MEASUREMENT-PROGRAM TYPE=*TLM.

The monitored data output is sorted in accordance with OCCUPATIONS. Only one screen is output.

```

D016ZE04 SM2 TLM          CYCLE:  60 S  SAMPLES:  144          <date> <time>

```

NAME	WAITING TASKS	OCCUPATIONS
LKSI007B	0.00	100.0
LKSI0483	0.00	100.0
LKSI007F	0.00	100.0
LKSI0076	0.00	100.0
LKSI0077	0.00	100.0
LKSI007D	0.00	100.0
LKSI0078	0.00	100.0
LKSI0071	0.00	100.0
LKSI0075	0.00	100.0
LKSI0073	0.00	100.0
LKSI006V	0.00	100.0
LKSI0070	0.00	100.0
LKSI004F	0.00	100.0
LKSI006T	0.00	100.0
LKSI000F	0.00	100.0
LKSI000E	0.00	100.0

Monitoring information

NAME	Name of the monitored task lock (non-printable characters are represented as blanks)
WAITING TASK	Mean number of tasks in the queue of the lock
OCCUPATIONS	Relative frequency with which the task lock is occupied (in percent)

UDS SQL report

This report contains data of the UDS/SQL database system.

Report output

The REPORT UDS_SQL statement is used to request the output of the UDS SQL report.

The UDS SQL report can only be output if

- UDS/SQL has been prepared for data to be sent to SM2 (see [page 74](#))
- the SM2 administrator has activated the monitoring process using the START-MEASUREMENT-PROGRAM TYPE=*UDS-SQL statement.

Overflow screens may be requested for this report.

The monitored data output is sorted under the category DB CONF.

SQHAV101 SM2 UDS SQL	CYCLE: 30 S									<date>	<time>
DB CONF	DML (1/S)	TA (1/S)	UPDTA (1/S)	LREAD (1/S)	PREAD (1/S)	PWRIT (1/S)	LCKWT (1/S)	DLOCK (1/S)			
Q1254KON	229.2	11.4	10.4	1676.4	366.2	185.2	21.6	9.1			

Monitoring information

DB CONF	Name of the DB configuration
DML (1/S)	Number of CODASYL and SQL statements per second
TA (1/S)	Number of transactions per second
UPDTA (1/S)	Number of update transactions per second
LREAD (1/S)	Number of logical read accesses to the databases per second
PREAD (1/S)	Number of physical read accesses to the databases per second
PWRIT (1/S)	Number of physical write accesses to the databases per second
LCKWT (1/S)	Number of wait states for locks per second
DLOCK (1/S)	Number of data deadlocks per second

i The monitored data is supplied by UDS/SQL asynchronously to SM2 and applies for one or more cycles which are defined by UDS/SQL and need not match the SM2 cycle. Here differences can exist in the duration of the cycles, and time displacements can also exist between the UDS/SQL and SM2 cycles. The duration of one or more UDS/SQL cycles is used to standardize the monitored data to one second. The values are therefore exact, but they only match the SM2 cycle to a certain degree.

UTM reports

These reports output application-specific data. In the following description, you should note that the terms used are the same as those used in UTM. This results in some discrepancies compared with the terms used in the RESPONSETIME reports (e.g. “transaction”). Please refer to the UTM manuals for an explanation of the terms used in UTM.

- The UTM report provides an overview of all UTM applications which supply data to SM2. The decision as to whether a UTM application supplies data to SM2 is taken when the application is generated and/or by means of a statement from the UTM administrator.

For further information, please refer to [section “UTM Monitored data on openUTM applications” on page 75](#) and the “Using openUTM Applications under BS2000” manual [11].

- The UTM APPLICATION report provides detailed information on a UTM application.

The reports can only be output if the SM2 administrator has activated the monitoring process using the statement `START-MEASUREMENT-PROGRAM TYPE=*UTM`.

The figure below illustrates the hierarchy of reports and the statement used to request them:

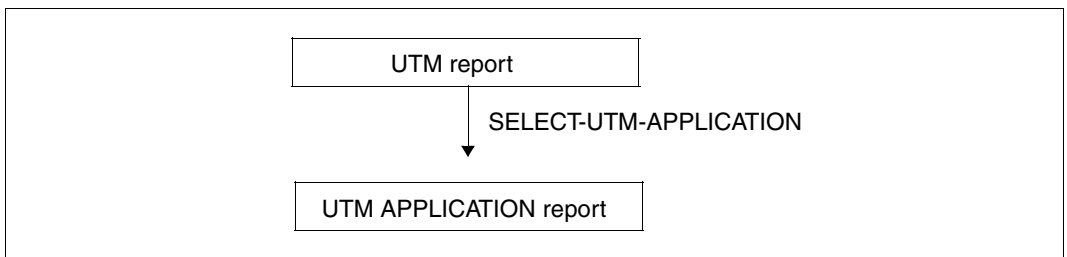


Figure 7: Hierarchy of UTM reports

UTM report

This report provides monitored data on UTM applications.

Report output

The REPORT UTM statement is used to request the output of the UTM report.

The UTM report contains only dialog-oriented monitored data. The data relating to database calls and distributed processing (DDP) is averaged out for all the steps in a dialog, even for those steps which contain no DB calls or distributed processing calls.

Overflow screens may also be requested for this report.

The monitored data output is sorted in accordance with the number of dialog steps performed (#DS).

D017ZE29 SM2 UTM		CYCLE: 300 S					<date> <time>		
APPLI-NAME	M	TIME (S/DS)	DB (%)	TAC-CL (%)	DDP (%)	# DS (1/S)	#DB-CALL (1/DS)	# TSK	# USER
SEI10	S	0.1	0.0	0.0	0.0	11.3	0.0	10	12
SDB1	S	0.0	0.0	0.0	0.0	0.0	0.0	4	3

Monitoring information

APPLI-NAME Name of the UTM application

M Mode of the UTM application
 S: UTM-S application
 F: UTM-F application

TIME Average total dwell time in seconds for a dialog step (DS) measured by UTM as the time elapsed between UTM accepting a message and sending a message

The following three values indicate average percentages of the total time shown under TIME:

DB	Wait time in UTM for execution of database calls by database systems
TAC-CL	Wait time as a result of TAC class bottlenecks
DDP	Wait time in UTM for messages from remote applications (DDP = distributed data processing)
# DS	Number of dialog steps executed per second
#DB-CALL	Number of database calls executed per dialog step
# TSK	Number of tasks currently running for the application
# USER	Number of users currently signed on to the application

UTM APPLICATION report

This report provides detailed information on a UTM application.

Report output

The UTM APPLICATION report is output only at the explicit request of the SM2 user with the help of the SELECT-UTM-APPLICATION statement (see [page 191](#)). This statement is used to output a UTM report together with a UTM APPLICATION report for the selected applications. The UTM APPLICATION report supplies detailed information on the UTM applications.

It is possible to output data on up to 16 applications.

An overflow screen with additional data for the UTM application can be requested for this report, see [page 316](#).

```

D017ZE29 SM2 UTM APPLICATION      CYCLE: 300 S      <date> >time>
APPL: SEI10      JOBN: SEI      UTM-VERS: V06.3M02      MODE: UTM-S      # DB'S: 0

```

MEAN VALUE FOR	DIALOG	ASYNCH	# DIAL STEPS (1/S)	11.3
TOTAL TIME (S/ST)	0.1	0.0	# DIAL TA'S (1/S)	11.3
TOT TIME WITH DB (S/ST)	0.0	0.0	# ASYN TA'S (1/S)	0.0
STEPS WITH DB (1/S)	0.0	0.0	# ASYN CONV (1/S)	0.0
TIME IN DB (S/ST)	0.0	0.0	# TASKS	10
# DB CALLS (1/ST)	0.0	0.0	# MAX ASYN TASKS	3
CPU TIME IN DB (MS/ST)	0.000	0.000	# CONNECTED USERS	12
# IO IN DB (1/ST)	0.0	0.0	# WAITING PRINTS	0
TOT TIME WITH DDP (S/ST)	0.0	0.0	# WAITING ATACS	0
STEPS WITH DDP (1/S)	0.0	0.0	# WAITING DPUTS	0
TIME IN DDP (S/ST)	0.0	0.0	# DIAL CURR CONV	9
CPU TIME (MS/ST)	1.53	0.000	# ASYN CURR CONV	0
# IO (1/ST)	8.7	0.0	CACHE HIT RATE (%)	97
			FREE PAGE POOL (%)	99
			BOURSE WAIT T (MS)	0

Monitoring information

APPL	Name of the UTM application
JOBN	BS2000 job name for the first task started for the UTM application
UTM-VERS	UTM version
MODE	UTM-S or UTM-F mode
# DB'S	Number of database systems the application works with

The following entries are all average values which refer to dialog steps (DIALOG) and asynchronous conversations (ASYNCH). In the explanations of the fields, "step" is used as a generic term for this. It is abbreviated in the report as "ST".

TOTAL TIME	For dialog steps: the total time for each dialog step from acceptance of the entry by UTM to the sending of a dialog message by UTM. For asynchronous conversations: the total time per second for each asynchronous conversation from the start of processing to the end of processing (not including the wait time before the start of processing). TOTAL TIME also includes waiting times that arise as a result of a TAC class bottleneck or as a result of waiting for messages from remote applications. On the other hand, a waiting time in the job queue of the UTM application (i.e. before a UTM task accepts the job for the first time) is not included.
TOT TIME WITH DB	Value analogous to TOTAL TIME, but only for steps with database calls (*)
STEPS WITH DB	Number of steps per second with database calls
TIME IN DB	Time per second per step UTM waits for the execution of database calls (*)
# DB CALLS	Number of database calls per step (*)
CPU TIME IN DB	CPU time per millisecond consumed per step in database systems (*) (***)
# IO IN DB	Number of I/Os per step in the called database systems (*) (***)
TOT TIME WITH DDP	Value analogous to TOTAL TIME, but only for steps with distributed processing calls (**)
STEPS WITH DDP	Number of steps per second with distributed processing calls
TIME IN DDP	Time per second per step UTM waits for the arrival of a message from a remote application (**)
CPU TIME	CPU time per millisecond consumed by UTM for processing the step, including user subroutines
# IO	Number of I/Os to/from UTM tasks which occur during processing of the step, including user subroutines
(*)	Only those steps are included where database calls occur
(**)	Only those requests are included where distributed processing calls occur

(***) Not all database systems which coordinate with UTM supply these values. Some database systems allow the database administrator to activate the provision of monitoring data by issuing a statement. For further details see the “Using openUTM Applications under BS2000” manual [11] and the relevant database system manuals. The monitored data is set to zero for database systems which do not supply any values. If several database systems are being used, and only some of these provide monitoring data, this should be taken into account when interpreting the average values.

The next four values show the application’s throughput (performance data).

# DIAL STEPS	Number of completed dialog steps per second
# DIAL TA'S	Number of completed dialog transactions per second
# ASYN TA'S	Number of completed asynchronous transactions per second
# ASYN CONV	Number of completed asynchronous conversations per second

The remaining values indicate the current status of the application.

# TASKS	Number of tasks running for the application
# MAX ASYN TASKS	Maximum number of tasks available for asynchronous processing
# CONNECTED USERS	Number of users signed on to the application
# WAITING PRINTS	Number of print jobs waiting for execution
# WAITING ATACS	Number of jobs for asynchronous programs waiting for processing
# WAITING DPUTS	Number of time-driven jobs waiting (background and output jobs)
# DIAL CURR CONV	Number of open dialog conversations
# ASYN CURR CONV	Number of open asynchronous conversations
CACHE HIT RATE	Hit rate in % when searching for a UTM page in the UTM cache
FREE PAGE POOL	Percentage of free pages in the UTM page pool
BOURSE WAIT T	Time in milliseconds a request waits in the job queue for a UTM application. This value is approximated by generating an internal test message.



Asterisks (***) are output in place of invalid or unavailable monitoring data.

Report output (overflow screen)

D017ZE29 SM2 UTM APPLICATION		CYCLE: 300 S		14-09-03 16:00:00		
APPL: SEI10		JOBN: SEI		TACCLASS INFO		
VALUE FOR	MEAN	CNT	NR	ELAPSED TIME	WAIT TIME	WAITING
				MEAN	CNT	MSG
CACHE WAIT	0.0	22.6				
CACHE HIT	1.0	9.2	01	111.4	13.5	0
PER WRITE	0.0	9.2	02	0.0	0.0	0
USER RES W	0.4	4.4	03	0.0	0.0	0
SYS RES W		680	04	0.0	0.0	0
INPUT MSG		11.3	05	0.0	0.0	0
OUTPUT MSG		11.3	06	0.0	0.0	0
LOGFILE WR		0.0	07	0.0	0.0	0
CURR LOAD		14	08	0.0	0.0	0
			09	0.0	0.0	0
			10	0.0	0.0	0
			11	0.0	0.0	0
			12	0.0	0.0	0
			13	0.0	0.0	0
			14	0.0	0.0	0
			15	0.0	0.0	0
			16	0.0	0.0	0

Monitoring information

CACHE WAIT

MEAN

Percentage of the wait situations when requesting cache buffers.

CNT

Number of requests for cache buffers per second.

CACHE HIT

MEAN

Percentage of the found cache buffers.

CNT

Number of cache buffers searched for per second.

PER WRITE

MEAN

In the case of periodic writes, number of pages written per periodic write.

CNT

Number of periodic writes per second.

USER RES W

MEAN

Percentage of the wait situations when accessing secondary storage areas (GSSB, TLS, ULS).

CNT

Number of transaction resource requests per second.

SYS RES W

CNT

Percentage of the wait situations for system locks in the entire number of lock requests in requests per thousand. The value is updated by UTM every 100 seconds.

INPUT MSG CNT	Number of all messages which the application has received from clients or partner applications (per second).
OUTPUT MSG CNT	Number of all messages which the application has sent to clients, printers or partner applications (per second).
LOGFILE WR CNT	Number of requests to write log records in the user log file (USLOG) (per second).
CURR LOAD CNT	Utilization of the application in percent in the last completed interval (100 seconds). The value is updated by UTM every 100 seconds.
TACCLASS INFO NR	Contains values for 16 TAC classes.
ELAPSED TIME MEAN CNT	Mean duration of the subroutine runs in milliseconds. Number of subroutine runs per second.
WAIT TIME MEAN CNT	Mean wait time of messages in milliseconds. Number of wait situations per second.
WAITING MSG CNT	Number of waiting messages (status at the end of the monitoring period).

VM report

This report provides monitored data for the individual virtual machines of a VM2000 system.

Report output

The REPORT VM statement is used to request the output of the VM report. The report can only be output if the SM2 administrator has activated the monitoring process using the statement START-MEASUREMENT-PROGRAM TYPE=*VM. The values for all VMs are only output if SM2 is running on the monitor VM. If this is not the case, only the values for the VM on which SM2 is running are output.

Overflow screens may also be requested for this report.

D016ZE04 SM2 VM		CYCLE: 60 S		<date>	<time>				
HYPERVISOR:		IDLE: 8.3%	ACTIVE: 5.1%	MEAS: 77.5%	# USABLE CPUS: 3				
IND	NAME	VM-GROUP	CPU-POOL	I	MEMORY (MB)	MAX	CPU (%) PLAN	MEAS	# CPUS
2	D016ZE04		*STDPOOL		3840	66	14	59.7	2

Monitoring information

In the description below, some values refer to the TOTAL TIME. This is the whole available CPU time during the monitoring period.

The following applies: TOTAL TIME = CYCLE * # USABLE CPUS with

CYCLE

Duration of the monitoring interval

USABLE CPUS

Number of real CPUs, available to BS2000 guest systems (also extra CPUs)

HYPERVERSOR	
IDLE	Percentage of TOTAL TIME accounted for by IDLE time. IDLE outputs the CPU share which is fundamentally still available. This share is not necessarily available to all VMs.
ACTIVE	Percentage of TOTAL TIME spent by the hypervisor in active time
# USABLE CPUS	Number of real CPUs, available to BS2000 guest systems (also extra CPUs)
MEAS %	Percentage of the active time of all BS2000 guest systems in the TOTAL TIME
IND	Index of the VM
NAME	Name of the VM
VM-GROUP	Name of the VM group the VM belongs to
CPU-POOL	Name of the CPU pool assigned to the VM
I	<p>The monitoring value MEAS contains also the hypervisor idle share of the CPUs available to the VM, if</p> <ul style="list-style-type: none"> – a VM runs in the scheduling mode “Dedicated CPUS” (DC) and – uses VM-ACTIVE-IDLE=*AT-DEDICATED-CPUS (see SHOW-VM-RESOURCES INF=*CPU). <p>This is indicated with Y.</p> <p>The VM active share then is only visible in the guest systems CPU report.</p>
MEMORY (MB)	Main memory size of the VM in megabytes

CPU (%)	
MAX	Maximum percentage of the VM on the CPU performance. This value relates to all available real CPUs (also extra CPUs). In VM2000 the value can be set with a precision of 0.01%. SM2 only outputs the unrounded, integral part here.
PLAN	Planned percentage of the VM on the CPU performance. In VM2000 the value can be set with a precision of 0.01%. SM2 only outputs the unrounded, integral part here.
MEAS	<p>Percentage of TOTAL TIME that can be accounted for by the CPU time of the VM. VMs which have used no CPU time during the monitoring cycle are not shown. Newly started VMs are not recorded until the subsequent monitoring cycle is completed.</p> <p>If a VM runs in the scheduling mode "Dedicated CPUS" (DC) and uses VM-ACTIVE-IDLE=*AT-DEDICATED-CPUS (see SHOW-VM-RESOURCES INF=*CPU), then the monitoring value MEAS contains also the hypervisor idle share of the CPUs available to the VM. The VM active share then is only visible in the guest systems CPU report.</p>
# CPUS	Number of active virtual CPUs assigned to a VM



The connection between CPU MEAS and the CPU values under REAL in the CPU report is explained in detail in [section "Special applications" on page 356](#).

The HYPERVISOR IDLE and ACTIVE values are only available on servers with /390 architecture.

VM CPU POOL report

This report provides monitored data for CPU pools.

Report output

The REPORT VM_CPU_POOL statement is used to request the output of the VM CPU POOL report. The report can only be output if the SM2 administrator has activated the monitoring process using the statement START-MEASUREMENT-PROGRAM TYPE=*VM.

The values for all CPU pools are only output if SM2 is running on the monitor VM. If this is not the case, only the values for the CPU pool which is assigned to the VM on which SM2 is running are output.

Overflow screens may also be requested for this report.

```

D017ZE19 SM2 VM CPU POOL          CYCLE:  60 S          <date> <time>

+-----+-----+-----+-----+
| CPU-POOL | #   | HPV CPU (%) | CPU (%) |
|          | CPUS| ACT   IDLE  |         |
+-----+-----+-----+-----+
| *STDPOOL |  3  |  7.2  15.3  |  77.5  |
+-----+-----+-----+-----+
    
```

Monitoring information

- CPU POOL Name of the CPU pool
- # CPUS Number of real CPUs, available to BS2000 guest systems (also extra CPUs) which are assigned to the CPU pool

The following CPU values relate to the CPU-pool-specific TOTAL TIME, i.e. the total CPU time made available by the CPU pool during a monitoring cycle.

The following applies: TOTAL TIME = CYCLE * # CPUS with

- CYCLE Duration of the monitoring interval

HPV CPU (%)

ACT	Percentage of TOTAL TIME spent by the hypervisor in active time
IDLE	Percentage of TOTAL TIME accounted for by hypervisor IDLE time
CPU (%)	Percentage of TOTAL TIME accounted for by the CPU time of all BS2000 VMs which are assigned to the CPU pool



If a CPU pool runs in the scheduling mode “dedicated CPUS” (DC) (see SHOW-VM-CPU-POOL INF=*STD) and one or more VMs of the assigned pool use the mode VM-ACTIVE-IDLE=*AT-DEDICATED-CPUS (see SHOW-VM-RESOURCES INF=*CPU) then the following applies: The idle percentage of these VMs will be output under CPU (%) and not under IDLE.

The HPV CPU (%) ACT and IDLE values are only available on servers with /390 architecture.

VM GROUP report

This report provides monitored data for VM groups (servers with /390 architecture).

Report output

The REPORT VM_GROUP statement is used to request the output of the VM GROUP report. The report can only be output if the SM2 administrator has activated the monitoring process using the statement START-MEASUREMENT-PROGRAM TYPE=*VM.

The values for all VM groups are only output if SM2 is running on the monitor VM. If this is not the case, only the values for the VM group which is assigned to the VM on which SM2 is running are output.

Overflow screens may also be requested for this report.

```

LEIBNIZ1 SM2 VM GROUP                CYCLE: 60 S                <date> <time>

+-----+-----+-----+-----+-----+
! VM-GROUP ! CPU-POOL ! MAX   CPU (%) !
!          !          !       PLAN  MEAS !
+-----+-----+-----+-----+-----+
! VMGR1    ! *STDPOOL ! 100   14   ! 9.7 !
! VMGR2    ! *STDPOOL ! 80    14   ! 0.4 !
! VMGR3    ! *STDPOOL ! 20    14   ! 0.2 !
    
```

Monitoring information

- VM-GROUP Name of the VM-Gruppe
- CPU-POOL Name of the CPU pool the VM group belongs to
- CPU (%) See also [page 322](#)
- MAX Maximum percentage of the VM group
- PLAN Planned percentage of the VM group
- MEAS Percentage of TOTAL TIME that can be accounted for by the CPU time of the VM group
- TOTAL TIME is the whole available CPU time during the monitoring period. The following applies:
 TOTAL TIME = CYCLE * # REAL CPUS with
- CYCLE Duration of the monitoring interval
- # USABLE CPUS Number of real CPUs, available to BS2000 guest systems (also extra CPUs)

6.2 SM2 information screens

ACTIVE PARAMETER

This screen displays the parameters set for the active monitoring programs. As many screens as are necessary are output. If sufficient space is available, all the parameters for one monitoring program are displayed on a single screen.

It is only possible to request output from the administration section using the SHOW-ACTIVE-PARAMETERS statement.

```

D016ZE04 SM2 ACTIVE PARAMETER                                     <date> <time>
-----
BCAM CONNECTION PARAMETER
-----
INWAIT-BUCKETS           : 5, 10, 100, 1000
REACT-BUCKETS           : 500, 1000, 5000, 10000
INPROC-BUCKETS          : 5, 10, 100, 1000
OUTPROC-BUCKETS         : 5, 10, 100, 1000

BCAM CONNECTION SET
SET-NAME                 : DIALOG                                PORT1110
CONNECTION-TYPE          : REMOTE                                REMOTE
HOST-SELECTION           :
(LLOCAL / PARTNER )     (*ANY / *ANY )                        (*ANY / *ANY )
CONNECTION-SELECTION    : *BY-NEA-NAME                          *BY-PORT-NUMBER
(LLOCAL / PARTNER )     ($DIALOG / *ANY )                      ( 1110 / *ANY )

CHANNEL-IO PARAMETER
-----
CHANNELS                 : *ALL

```

Monitoring information

BCAM-CONNECTION-PARAMETER

- | | |
|-----------------|---|
| INWAIT-BUCKETS | Definition of the upper limits of up to four ranges in which the INWAIT times are listed by order of magnitude |
| REACT-BUCKETS | Definition of the upper limits of up to four ranges in which the REACT times are listed by order of magnitude |
| INPROC-BUCKETS | Definition of the upper limits of up to four ranges in which the INPROC times are listed by order of magnitude |
| OUTPROC-BUCKETS | Definition of the upper limits of up to four ranges in which the OUTPROC times are listed by order of magnitude |

BCAM-CONNECTION-SET

SET-NAME Name of the selected connection set

CONNECTION-TYPE Type of connected to be monitored (REMOTE, LOCAL or BOTH)

HOST-SELECTION Hosts between which the connections are established

CONNECTION-SELECTION Type of connection set (*BY-PORT-NUMBER or *BY-NEA-NAME) and the selected port numbers or application names

CHANNEL-IO PARAMETER

CHANNELS Channel addresses for the channels monitored

```

D016ZE04 SM2 ACTIVE PARAMETER                                <date> <time>
COSMOS PARAMETER
-----
TITLE                : COSMOS
BUFFER-SIZE          : 7
NUMBER-OF-BUFFERS    : 32
ADDITIONAL-INFO      : *CONFIGURATION
OUTPUT               : *DISK
                    : :20S6:$NEI.COSMOS.PAM.V18.0A

TASK-SELECTION
TSN                  : *NONE
USER-ID              : *NONE
JOB-NAME             : *NONE
CATEGORY             : *NONE
TYPE                 : SYSTEM, BATCH, DIALOG, TP

EVENT-SELECTION
EIA-INTERRUPT-CLASS : *ANY

```

Monitoring information

COSMOS PARAMETER

TITLE	Title of the COSMOS monitoring process
BUFFER-SIZE	Number of 4K pages per buffer
NUMBER-OF-BUFFERS	Number of buffers
ADDITIONAL-INFO	Additional data to be included
OUTPUT	Way in which the COSMOS output files are to be written
TASK-SELECTION	Conditions on the basis of which the tasks are to be monitored
TSN	Selection by TSN
USER-ID	Selection by user ID
JOB-NAME	Selection by job name
CATEGORY	Selection by category
TYPE	Selection by task type
EVENT-SELECTION	Conditions on the basis of which the events are selected
EIA-INTERRUPT-CLASS	Selection of EIA events on the basis of the interrupt code

```

D016ZE04 SM2 ACTIVE PARAMETER                                <date> <time>
EIA-SVC-NUMBER      : *ANY
IO-DEVICE           : *ANY
DAB-CACHE-ID       : *ANY
MEMORY-CLASS        : *ANY
SLOT-MEMORY-CLASS   : *ANY
PEND-CODES          : *ANY
LOCK-ID             : *ANY
TLT-DESCRIPTOR      : *ANY
TSKI-SWITCH         : *ANY
TSVC-SVC-NUMBER     : *ANY
CPU-NUMBER          : *ANY

UNLOAD              : AT-MEASUREMENT-PROGRAM-STOP
MEASUREMENT-TIME    : *NOT-SPECIFIED

OPENED EVENT        : INIT, STAT, PTSK, CREA, DEST, MMRC, LGON, ACF, BOUR,
                    CHTM, CMS, DABA, DABC, DABE, DABF, DABI, DABS, EIA2,
                    BLS, EIA3, SVC, FITC, IDLE, INTR, IONQ, PAGE, PAM,
                    PCCC, PCTC, PEND, PMIO, PRGS, PRGT, RELM, REQM, SDV,
                    TSKI, TSVC, WSCT

```

Monitoring information

EIA-SVC-NUMBER	Selection of EIA events on the basis of the SVC number
IO-DEVICE	Selection of PMIO, SDV, CHTM and IONQ events on the basis of the mnemonic device name
DAB-CACHE-ID	Selection of DAB events on the basis of the device VSN
MEMORY-CLASS	Selection of the REQM and RELM event on the basis of the memory class
SLOT-MEMORY-CLASS	Selection of the SLOT event on the basis of the memory class
PEND-CODES	Selection of the PEND event on the basis of the pend code
LOCK-ID	Selection of the LOCK event on the basis of the lock ID
TLT-DESCRIPTOR	Selection of the TLT event on the basis of the TLT descriptor
TSKI-SWITCH	Selection of the TSKI event
TSVC-SVC-NUMBER	Selection of the TSVC events on the basis of the SVC numbers
CPU-NUMBER	Selection of the events on the basis of the CPU numbers
UNLOAD	Time at which the COSMOS subsystem is to be unloaded
MEASUREMENT-TIME	Timing of the monitoring process
OPENED EVENT	Events opened for recording

```
D016ZE04 SM2 ACTIVE PARAMETER <date> <time>
DISK PARAMETER
-----
NO PARAMETERS DEFINED
DISK-FILE PARAMETER
-----
DEVICES          : B3A2, B3A3, B3A4, B3A5, B3A6, B3A7, B3A8, B3A9
FILE PARAMETER
-----
FILENAME         : :20SH:$TSOS.TSOSCAT
                  : :2S06:$TSOS.TSOSCAT
                  : :20S7:$TSOS.TSOSCAT
```

Monitoring information

DISK-FILE-PARAMETER

DEVICES Mnemonic device names of the monitored disk devices

FILE PARAMETER

FILENAME Name of the file whose access values are monitored


```

D016ZE04 SM2 ACTIVE PARAMETER                                <date> <time>
ISAM PARAMETER
-----
POOL-NAME (SCOPE)      : :20SH:SRMPPOOL(TA=RP01      )
FILE-NAME              : :20S6:$NEI.SM2.ISAMFILE.1
PERIODIC-TASK PARAMETER
-----
TSN                    : *NONE
USER-ID                : NEI, TSOS
JOB-NAME               : *NONE

```

Monitoring information

ISAM PARAMETER

POOL-NAME Name of the monitored ISAM pool with the catalog ID of the pubset to which the ISAM pool is assigned

(SCOPE) Defined type of the ISAM pool:
*HOST or TA=tsn

FILE-NAME Name of the monitored ISAM files

PERIODIC-TASK PARAMETER

TASKS FOR LOG FILE

Tasks whose monitored data is to be written to the SM2 output file
Selected tasks are specified under the selection criteria USER-ID and JOB-NAME.

```

D016ZE04 SM2 ACTIVE PARAMETER                                <date> <time>
RESPONSETIME PARAMETER
-----
SCOPE                : (BUCKET,CATEGORY)
DEFINITION           : 1
CONNECTION-NUMBER    : 1024
RESPONSE-BUCKETS     : 5, 10, 20, 50, 100
THINKTIME-BUCKETS    : 50, 150, 300, 600, 1200
TRANSACT-BUCKETS     : 5, 10, 20, 50, 100
WAITTIME-BUCKETS     : 1, 2, 5, 10, 20

CONNECTION SET
SET-NAME             : *GLOBAL                               DIALOG
SET-DEFINITION       : BY-CONNECTION                       BY-CONNECTION
CONNECTION-TYPE      : REMOTE                               REMOTE
CONNECTION           : (*ALL ,*ALL ,*ALL )($DIALOG ,*ALL ,*ALL )
(APPL,PARTNER,PROC)

```

Monitoring information

RESPONSETIME PARAMETER

SCOPE Indicates whether the response time data has been recorded according to buckets (BUCKET) or categories (CATEGORY)

DEFINITION Definition of the type of response time monitored

CONNECTION-NUMBER
Maximum number of connections monitored

RESPONSE-BUCKETS
Definition of the upper limits of up to five ranges in which the response times are stored by order of magnitude

THINKTIME-BUCKETS
Definition of the upper limits of up to five ranges in which the think times are stored by order of magnitude

TRANSACT-BUCKETS
Definition of the upper limits of up to five ranges in which the transaction times are stored by order of magnitude

WAITTIME-BUCKETS
Definition of the upper limits of up to five ranges in which wait times in the BCAM pool are stored by order of magnitude

CONNECTION-SET

SET-NAME	Name of the selected connection set
SET-DEFINITION	Type of connection set monitored (BY or EXCEPT)
CONNECTION-TYPE	Type of connection monitored (REMOTE, LOCAL or BOTH)
CONNECTION	Selection of up to five connection groups
APPL	Name of the selected application
PARTNER	Name of the selected partner
PROC	Name of the selected partner computer

```

D016ZE04 SM2 ACTIVE PARAMETER                                <date> <time>
SAMPLING-DEVICE PARAMETER
-----
DEVICE SERVICETIME      : *ON
STORAGE-SYSTEM PARAMETER
-----
ADDITIONAL DATA
SYMMETRIX              : DEVICES WITH MN, BACK END DATA, DIRECTOR
SERVICETIME PARAMETER
-----
DEVICES                : *NONE
PUBSETS                : 20S6
SYSTEM PARAMETER
-----
DEVICES                : *DISK, *TAPE
PUBSETS                : *NONE

```

Monitoring information

SAMPLING-DEVICE PARAMETER

DEVICE SERVICETIME

Specifies the recording of the service times

STORAGE-SYSTEM PARAMETER

ADDITIONAL DATA SYMMETRIX

Angabe über den Umfang der zu messenden Daten

SERVICETIME PARAMETER

DEVICES Specifies the mnemonic device names of the monitored devices

PUBSETS IDs of the pubsets

SYSTEM PARAMETER

DEVICES Mnemonic device names of the monitored devices or *NONE, *ALL, *TAPE or *DISK

PUBSETS IDs of the pubsets

```

D016ZE04 SM2 ACTIVE PARAMETER                                <date> <time>
TASK PARAMETER
-----
TASK SELECTION          : *ALL
DEVICES                 : *NONE
PUBSETS                 : 20S6

```

Monitoring information

TASK PARAMETER

TASK SELECTION *ALL: All tasks are monitored.

Tasks whose monitored data is to be written to the SM2 output file.

Selected tasks are specified under the selection criterion TSN, USER-ID, JOB-NAME, JOB-CLASS or TYPE.

DEVICES Mnemonic device names monitored by the task monitoring program for each task or *NONE, *ALL or *DISK

PUBSETS IDs of the pubsets



If a selected monitoring program is not active, the message MEASUREMENT NOT ACTIVE is displayed in place of the parameters.

DEFINED PARAMETER

This screen displays the parameters set for the defined monitoring programs. As many screens as are necessary are output. If sufficient space is available, all the parameters for one monitoring program are displayed on a single screen.

It is only possible to request output from the administration section using the SHOW-DEFINED-PARAMETERS statement. The layout of the screen is the same as the ACTIVE PARAMETER screen, with the difference that the first line contains the header DEFINED PARAMETER.



If no parameters are defined for a selected monitoring program, the message NO PARAMETERS DEFINED is displayed in place of the parameters.

MEASUREMENT STATUS

This screen provides the user with information on the current status of the measurement at any time. Output comprises three parts. The first two parts contain general administration data and the third part provides information on the functions permitted and on monitoring programs which have been defined and which are running.

In the analysis subinterval, output of the MEASUREMENT STATUS screen is requested by means of the STATUS statement, and in the administration facility it is requested by means of the SHOW-MEASUREMENT-STATUS statement.

```

D016ZE04 SM2 MEASUREMENT STATUS                                     <date> <time>

TSN OF SM2 PRIMARY ADMINISTRATOR : *NONE
NUMBER OF PRIVILEGED SM2 USERS.. :    4
NUMBER OF NONPRIVILEGED SM2 USERS:    2

SM2 GATHERING TASK CREATED AT... : <date> <time>
OFFLINE PERIOD.....              :    240 S
ONLINE PERIOD.....               :    60 S
SAMPLING PERIOD.....             :    400 MS
SM2 LOGGING FILE.....            : :20S6:$NEI.SM2.SAM.MANUAL.V18.0A

SECONDARY ADMINISTRATOR ALLOWED. : YES
USER MEASUREMENTS ALLOWED.....   : FILE    ISAM    TASK
USER MEASURED OBJECTS.....       : FILES ( 0) ISAM-POOLS ( 0) TASKS ( 0)
PRIVILEGED MEASUREMENTS DEFINED. : BCAM    CHA-IO  COSM    DISK    D-FILE
  FILE    ISAM    PERTSK  RTIME  S-DEV    ST-SYS  SVTIME  SYSTEM  TASK
PRIVILEGED MEASUREMENTS ACTIVE.. : BCAM    CHA-IO  CMS     DAB     DISK
  D-FILE  DLM    FILE    GS     HSMS    ISAM    MSCF    NSM    PERTSK
  PFA     POSIX  PUBSET  RTIME  S-DEV    SESAM   ST-SYS  SVC     SVTIME
  SYSTEM  TASK   TCPIP   TLM    UDS     VM
    
```

General administration data

TSN OF SM2 PRIMARY ADMINISTRATOR

Task sequence number (TSN) of the SM2 primary administrator or, if none exists yet, *NONE.



In a host network, it is possible to perform administration from a remote host. In this case, the host name is output after the TSN.

NUMBER OF PRIVILEGED SM2 USERS

Number of privileged users currently working with SM2 (including the primary administrator)

NUMBER OF NONPRIVILEGED SM2 USERS

Number of nonprivileged users currently working with SM2

Either

SM2 GATHERING TASK CREATED AT

Date (using the ISO format) and time the monitoring task was created
(provided the SM2 monitoring task is active)

or

SUBSYSTEM SM2 IN DELETE !!!

Subsystem SM2 in the DELETE state.

With this status, no further users are allowed to log on. Permission for user-specific task monitoring is withdrawn. The SM2 output file and user-specific SM2 output files are closed automatically. Repeated attempts to open the file are rejected. Subsystem SM2 then terminates.

or

SM2 GATHERING TASK TERMINATED !!!

SM2 monitoring task terminated (due to error)

OFFLINE PERIOD

Length of the monitoring cycle for background monitoring (in seconds).

ONLINE PERIOD

Length of the monitoring cycle for online analysis (in seconds)

or

SAME AS OFFLINE

no separate monitoring cycle for online analysis exists.

SAMPLING PERIOD

Length of the sampling period in milliseconds.

SM2 LOGGING FILE

*NONE if no SM2 output file is defined

or

*OPEN if an SM2 output file is open
(output for nonprivileged users)

or

name of the open SM2 output file
(output for privileged users)

Status data on additional functions

SECONDARY ADMINISTRATOR ALLOWED:

YES: The SM2 primary administrator has permitted other SM2 administrators

NO: Parallel SM2 administration is not permitted

USER MEASUREMENTS ALLOWED

List of the user monitoring programs permitted. Each user sees precisely those user monitoring programs he/she is allowed to use.

The following entries are possible: FILE, ISAM, TASK

USER MEASURED OBJECTS

The total number of monitored objects is output for each permitted user monitoring program. If this number is zero, the corresponding user monitoring program is permitted (for some or all users), but nobody has yet activated an object.

The following entries are possible: FILES, ISAM-POOLS, TASKS

PRIVILEGED MEASUREMENTS DEFINED

List of monitoring programs with monitoring definition

The following entries are possible: BCAM, CHA-IO, COSM, D-FILE, FILE, ISAM, PERTSK, RTIME, S-DEV, ST-SYS, SVTIME, SYSTEM, TASK

PRIVILEGED MEASUREMENTS ACTIVE

List of active monitoring programs.

The following entries are possible: BCAM, CHA-IO, CMS, COSM, DAB, D-FILE, DLM, FILE, GS, GSVOL, HSMS, ISAM, MSCF, NSM, PERTSK, PFA, POSIX, PUBSET, RTIME, S-DEV, SESAM, ST-SYS, SVC, SVTIME, SYSTEM, TASK, TCPIP, TLM, UDS, UTM, VM

Meaning of abbreviations for the monitoring programs:

Abbreviations	Monitoring program
BCAM	BCAM-CONNECTION monitoring program
CHA-IO	CHANNEL-IO monitoring program
CMS	CMS monitoring program
COSM+	COSMOS (“+” means that COSMOS has been initialized but not started)
DAB	DAB monitoring program
D-FILE	DISK-FILE monitoring program
DLM	DLM monitoring program
FILE	FILE monitoring program (privileged users)

Table 14: Abbreviations for monitoring programs

(part 1 of 2)

Abbreviations	Monitoring program
GS	GS monitoring program
GSVOL	GSVOL monitoring program
HSMS	HSMS monitoring program
ISAM	ISAM monitoring program (privileged users)
MSCF	MSCF monitoring program
NSM	NSM monitoring program
OFT	openFT monitoring program
PERTSK	PERIODIC-TASK monitoring program
PFA	PFA monitoring program
POSIX	POSIX monitoring program
PUBSET	PUBSET monitoring program
RTIME	RESPONSETIME monitoring program
S-DEV	SAMPLING-DEVICE monitoring program
SESAM	SESAM monitoring program
ST-SYS	STORAGE-SYSTEM monitoring program
SVC	SVC monitoring program
SVTIME	SERVICETIME monitoring program
SYSTEM	SYSTEM monitoring program
TASK	TASK monitoring program
TCPIP	TCP-IP monitoring program
TLM	TLM monitoring program
UDS	UDS monitoring program
UTM	UTM monitoring program
VM	VM2000 monitoring program

Table 14: Abbreviations for monitoring programs

(part 2 of 2)

If a monitoring program is followed by an asterisk (e.g. UTM*), this means that the monitoring program is active as far as SM2 is concerned, but that the corresponding subsystem is currently not supplying any data (e.g. because it has not been started).

SELECTED HOSTS

This screen displays the hosts which were selected by means of the SELECT-HOSTS statement.

The output can only be requested in the administration facility using the SHOW-SELECTED-HOSTS statement.

D016ZE04 SM2 SELECTED HOSTS			<date> <time>
HOST-NAME	PROCESSOR-NAME	LAST BUFFER	
*LOCAL	D016ZE04	12:31:01	
D016ZE07	D016ZE07	12:31:00	
D016ZE15	D016ZE15	12:31:00	

Information on the selected hosts

HOST-NAME	Name of the host
PROCESSOR-NAME	Processor name from the point of view of the local host (for information on host and processor names, see the note in the description of the SELECT-HOSTS statement on page 128)
LAST BUFFER	Time stamp of the last monitoring data called; this time stamp is identical to the time in the top right-hand corner of the reports RSLT NOT VALID means that no (up-to-date) data is available. Possible reasons: <ul style="list-style-type: none"> – No data has been requested. – Data has been requested but not transferred because the monitoring cycle on the addressed host was not yet terminated or the transfer was unsuccessful.

STATUS TABLE

No monitoring data is output in the STATUS TABLE.

Output comprises two parts. The first part provides the user with information on which SM2 tasks exist and on the status of the subsystems used by SM2. The line containing "MISSED RECORDS" is only output if unwritten records exist.

The second part outputs as many of the most recent trace entries as can be accommodated on one screen.

The SHOW-SM2-STATUS statement is used to request the output of the STATUS TABLE. This statement is only available to privileged users in the administration facility.

TASK / SS		START	TIME	STATUS	END	TIME
GATHERER		<date>	<time>	RUNNING		
WRITE-TASK		<date>	<time>	ENDED	<date>	<time>
USER-WRITE-TASK		<date>	<time>	RUNNING		
BCAM-SM2		<date>	<time>	IN USE		
UTM-SM2		<date>	<time>	UNAVAIL		
SHC-OSD		<date>	<time>	IN USE		

LAST TRACED EVENTS							
CUR#	TIME	TSN	REASON	TYPE	MODUL	ID	CODE
8	15:12:03	SM2U	CREATE	TASK	UDM	0000	00000000
7	15:08:53	SM2G	UTM-NLD	SS	UTM	0010	00400008
6	15:08:39	SM2G	BCM-STA	SS	DSSM	0000	00000000
5	15:07:58	SM2G	STS-STA	SS	STS	00A0	00000000
4	10:21:04	SM2W	TERM	TASK	GDM	1F54	00000000
3	10:19:53	SM2W	CREATE	TASK	GDM	016E	00000000
2	10:19:22	SM2G	CREATE	TASK	GAT4	0000	00000000
1	10:19:22	038E	SM2-STA	SS	84S	0000	00000000

Monitoring information

TASK / SS	Name of the task or subsystem
START TIME	Date (in ISO4 format) and time at which the task or subsystem was created
STATUS	Status of the task or subsystem

Meanings of the entries for tasks:

RUNNING	Task running normally
ENDED	Task terminated normally
ABENDED	Task terminated abnormally

Meanings of the entries for subsystems:

IN USE	SM2 is working with the subsystem
UNAVAIL	SM2 requires data from the subsystem, but the subsystem is not running
UNUSED	SM2 is no longer working with the subsystem (of its own accord)
ABENDED	SM2 is no longer working with the subsystem because the subsystem is no longer running
STOPPED	SM2 has unloaded the subsystem
END TIME	Date (in ISO4 format) and time at which the task was terminated or at which the subsystem stopped being used
MISSED RECORDS	This information is omitted if SUM = 0.
SUM	Total number of unwritten records (in the current SM2 output file)
LAST INTERVAL	Unwritten records in the last cycle
MISSED EVENTS IN LAST COSMOS SESSION	Number of messed events during the last Cosmos session
LAST TRACED EVENTS	
CUR#	Number of the trace entry (descending)
TIME	Time of the trace entry
TSN	TSN of the task in which a trace entry was written. The SM2 system tasks have the following default TSNs: GATHERER: SM2G WRITE-TASK: SM2W USER-WRITE-TASK: SM2U

REASON	Reason for the trace entry. In many cases (with TYPE=S-ER) this is the name of the system function which reported an error.	
CREATE	Creation of task	
TERM	Termination of task	
Subsystem-STA	Start of subsystem use	
Subsystem-STO	End of subsystem use	
Subsystem-NLD	Subsystem not loaded	
TYPE		
I-ER:	Internal error	(internal error during execution of SM2)
S-ER:	System error	(error calling a BS2000 system function)
TASK:	Task event	(start/end/crash of an SM2 task)
SS:	Subsystem event	(start/end of subsystem use by SM2)
MODUL	Abbreviated name of the SM2 module writing the trace entry (without "NPS", "NPFS")	
ID	Identifies the location at which the error occurred. In many cases this is the offset within the module. If the error is unique within the module, this field contains no entry.	
CODE	Return code of the system function called (for TYPE=S-ER) or additional information (for TYPE=I-ER) or zero	



The trace entries (particularly those with TYPE=I-ER,S-ER) are only relevant for SM2 diagnosis.

USER MEASURED OBJECTS

This screen displays the parameters set for the monitoring programs defined by the user. As many screens as are necessary are output. If sufficient space is available, all the parameters for one monitoring program are displayed on a single screen.

In the analysis section, only the parameters for the user monitoring program TASK are displayed, together with the objects users have activated for themselves.

The SHOW-USER-MEASURED-OBJECTS statement is used to request the output of this screen.

```

BALU2      SM2 USER MEASURED OBJECTS                                <date> <time>
-----
FILE PARAMETER
-----
FILENAME           : :4V07:$TSOS.SM2.SAM.TEST
BY TASK            : 038E

ISAM PARAMETER
-----
POOL               : :4V07:YDBP0001(*HOST      )
BY TASK            : 038E

TASK PARAMETER
-----
MEASURED TSN       : 038E
MEASURED USER-ID  : TSOS
MEASURING TSN      : 038E
MEASURING USER-ID : TSOS
PC-INTERVAL        : 10
SVC-STATISTICS     : ON
  
```

Monitoring information

FILE PARAMETER

FILENAME Name of the file for which access values are to be monitored
 BY TASK Task sequence number of the tasks monitoring the file

ISAM PARAMETER

POOL Name of the ISAM pool being monitored
 BY TASK Task sequence number of the tasks monitoring the ISAM pool

TASK PARAMETER

MEASURED TSN	TSN of the task being monitored
MEASURED USER-ID	User ID of the task being monitored
MEASURING TSN	TSN of the task which initiated monitoring
MEASURING USER-ID	User ID of the task which initiated monitoring
PC-INTERVAL	Sampling cycle for program counter statistics in milliseconds, or 0 if the program counter statistics are not activated.
SVC-STATISTICS	Indicates whether SVC statistics are activated (ON) or deactivated (OFF).



If no objects have been defined for a selected monitoring program, the message “NO OBJECTS DEFINED” is issued in place of the parameters. If the user is not authorized for monitoring or if the monitoring is not permitted for this user, the message “MEASUREMENT NOT ALLOWED” is issued in place of the parameters.

7 Installation and application of SM2

7.1 Installation

openSM2 is supplied using the delivery system SOLIS and consists of the delivery groups SM2 and SM2-TOOLS.

By default openSM2 is installed using the installation monitor IMON.

Delivery group SM2

Delivery group SM2 contains the files for the SM2 software monitor and the SM2R1 and SM2U1 utility routines. <ver> = 190 applies for openSM2 V10.0.

Filename	Function
SM2 SM2U1 SM2R1	Prephases for loading and starting the corresponding phases from the libraries listed below
SYSLNK.SM2.<ver>	Loadable parts of SM2 (for servers with /390 architecture only)
SKMLNK.SM2.<ver>	Loadable parts of SM2 (for servers with x86 architecture only)
SYSLNK.SM2.<ver>.SM2	Loadable parts of SM2
SYSLNK.SM2.<ver>.SM2U1	Loadable parts of SM2U1
SYSLNK.SM2.<ver>.SM2R1	Loadable parts of SM2R1
SYSRMS.SM2.<ver>	Loader package for SM2
SYSNRF.SM2.<ver>	Auxiliary file for REP processing
SYSSII.SM2.<ver>	Declaration file for IMON
SYSLIB.SM2.<ver>	C header files for program interface, ISM2CALL module, sample procedures
SKULIB.SM2.<ver>	ISM2CALL module (for native servers with x86 architecture only)
SYSDAT.SM2.<ver>.MTFILE	Control file for SM2R1
SYSMSP.SM2.<ver>.D	PLI1 text file for SM2R1 (German)
SYSMSP.SM2.<ver>.E	PLI1 text file for SM2R1 (English)

Table 15: Files of delivery group SM2

(part 1 of 2)

Filename	Function
SYSSPR.SM2.<ver>.SM2R1	Procedure for the START-SM2R1 command
SYSSDF.SM2.<ver>	Syntax file with all statements and commands (SM2, SM2U1, SM2R1)
SYSMES.SM2.<ver>	Message file for SM2,SM2U1,SM2R1
SYSSSC.SM2.<ver>	DSSM declarations for SM2
SIPLIB.SM2.<ver>	Contains the “restricted macros” of SM2
SYSFGM.SM2.<ver>.D	Release notice (German)
SYSFGM.SM2.<ver>.E	Release notice (English)

Table 15: Files of delivery group SM2

(part 2 of 2)

Delivery group SM2-TOOLS

Delivery group SM2-TOOLS contains the files for the INSPECTOR and ANALYZER applications. <ver> = 090 applies for openSM2 V10.0.

Filename	Function
SYSLNK.SM2-TOOLS.<ver>	BS2000 agents
SYSSDF.SM2-TOOLS.<ver>	Syntax file with all commands
SYSFGM.SM2-TOOLS.<ver>.D	Release notice (German)
SYSFGM.SM2-TOOLS.<ver>.E	Release notice (English)
SYSSII.SM2-TOOLS.<ver>	Declaration file for IMON
SYSSPR.SM2-TOOLS.<ver>	Procedures for the SDF commands
SYSDAT.SM2-TOOLS.<ver>.IN.CONF	Configuration file for the INSPECTOR agents
SPCDAT.SM2-TOOLS.<ver>	Manager

Table 16: Files of delivery group SM2-TOOLS

Further information on configuring, starting and terminating the agents of INSPECTOR and ANALYZER is provided in the help for INSPECTOR and ANALYZER manager, section “The BS2000 agent”.

Messages of openSM2

You obtain information on the meaning of the messages issued by openSM2 with the BS2000 command /HELP-MSG-INFORMATION.

All messages can be found in the message file using the software product MSGMAKER (see the “Utility Routines” manual [1]).

7.1.1 Loading and unloading the subsystem

The SM2 subsystem is generated and started for the first time under an ID which has been assigned the privilege SW-MONITOR-ADMINISTRATION. This loads the SM2 modules. The BCAM-SM2 subsystem required for the RESPONSETIME monitoring program is only loaded when it is required. The required subsystem UTM-SM2 is not started automatically (see [page 75](#) and the “Using openUTM Applications under BS2000” manual [11]).

The SM2 subsystem can also be loaded via the DSSM command `/START-SUBSYSTEM SUBSYSTEM-NAME=SM2`. SM2 can then be started (`/START-SM2`) under any user ID.

The SM2 subsystem exists until it is unloaded via the DSSM command `/STOP-SUBSYSTEM SUBSYSTEM-NAME=SM2`. The SM2 output file is implicitly closed when `/STOP-SUBSYSTEM SUBSYSTEM-NAME=SM2` is specified. Likewise, permission for task monitoring by the user is retracted. Authorization beyond that granted by the monitoring privilege SW-MONITOR-ADMINISTRATION is required to execute the DSSM commands (SUBSYSTEM-MANAGEMENT). Unconditional unloading with `/STOP-SUBSYSTEM SUBSYSTEM-NAME=SM2,FORCED=*YES` is possible, but it should only be applied in the event of problems.

The system administrator can obtain information on the current status of the SM2 subsystem by issuing the command `/SHOW-SUBSYSTEM-STATUS SUBSYSTEM-NAME=SM2`.

7.1.2 Installation of the manager of INSPECTOR and ANALYZER

When the SM2-TOOLS release group is installed with IMON, the managers are stored as BS2000 file `SPCDAT.SM2-TOOLS.<ver>`. The installation procedure is described in the Release Notice for SM2-TOOLS.

7.2 System resource utilization by SM2

To perform monitoring, SM2 must use the resources of the system to be monitored. This section provides an overview of SM2 resources requirements and information on assessing and restricting the induced system utilization.

With regard to system utilization, only the following basic resources are considered:

- external storage
- CPU
- main memory

During monitoring, only the resources utilized by the monitoring task are analyzed (i.e. without considering any parallel SM2 user programs), this being the way SM2 is most usually operated.

7.2.1 External storage utilization

The disk storage space requirements for SM2 output files can be restricted by:

- appropriate organizational measures (e.g. transfer of old SM2 output files to tape and creation of a new SM2 output file for each monitoring task)
- modifying SM2 parameters (e.g. specification of the monitoring cycle with the OFFLINE-PERIOD operand)
- specifying the monitoring programs to be activated
- compressing the output files with the SM2U1 statement SET-COMPRESSION

7.2.2 CPU utilization

To facilitate understanding of the factors affecting CPU utilization, readers should be familiar with the monitoring methods (see [section “Acquisition of monitored data” on page 27](#)).

Method based on monitoring cycle

At the end of the monitoring cycle, the measured values of the previous monitoring cycle are calculated for all monitoring programs (with the exception of TASK). In contrast to the sample cycle, the monitoring cycle is so much longer that system utilization by this method based can be disregarded.

Method based on samples

A monitoring task routine is activated at specific intervals (see the SAMPLING-PERIOD operand of the MODIFY-MEASUREMENT-PERIODS statement on [page 112](#)) to take samples.

For each activation of the sampling routine, a basic load has to be processed, regardless of the variables (devices and tasks) to be monitored.

In addition to this basic load, further instructions are processed which depend on the number of objects to be monitored (devices, channels, tasks).

If this number remains constant in the monitoring cycle, CPU utilization is almost directly proportional to the sampling rate (i.e. halving the sampling cycle, for example, causes the induced system load to be doubled).

Lengthening the sampling cycle should be balanced by lengthening the monitoring cycle to prevent a deterioration in sampling precision.

The sample-driven method is used for device and channel utilization, the length of queues, and the monitoring programs CMS and TLM.

Method based on events

When this method is used, the monitor is the “passive” component in contrast to the other system components, which are “active”. When specific events occur in the system (e.g. starting of an I/O operation), specific monitoring programs are activated which collect the relevant data (e.g. which device, which task, etc.).

While the monitor is inactive, no system utilization is induced.

If, however, events occur which are to be monitored, utilization increases in proportion to the load (i.e. to the number of calls).

When this monitoring method is used, the system load can be reduced only by reducing the number of objects to be monitored.

The event-driven method is used by the monitoring programs CHANNEL-IO, DISK-FILE, FILE, ISAM, PERIODIC-TASK, RESPONSETIME, SAMPLING-DEVICE, SERVICETIME, SYSTEM and TASK and by all user-specific monitoring programs.

UTM monitoring program

If UTM applications are running on the system and these applications provide data for SM2, the following additional load occurs in each UTM task: an additional 500 instructions (approx.) are required at the end of each dialog step and each asynchronous conversation in order to provide this data. For a typical application, this amounts to considerably less than 1% of the entire processing volume.

The resulting additional load on the system thus depends on the throughput in the applications, but can generally be ignored.

If values are also provided from the database systems, an additional load arises in these database systems in order to capture the monitored data. This load depends on the database system itself and the version used. For this reason, no general rule can be given.

Write task and I/O buffer

SM2 creates a write task (system task with TSN=SM2W) for writing to the SM2 output file. This task exists only from OPEN to CLOSE and is activated only when an I/O buffer is full. The input/output operation is controlled by the write task, and the CPU is required for task execution (CPU time for TPR and SIH states). The CPU time required for writing to the I/O buffer is assigned not to the write task but to the initiating task (system task or SIH processor state).

The data rate for writing to the I/O buffer depends on whether the record is written

- by the monitoring task at the end of a monitoring cycle,
- by another task, or
- in the SIH processor state

In the first case, the data rate depends on the number of active monitoring programs (data records) and monitored objects, as well as the duration of the monitoring cycle (OFFLINE-PERIOD operand in the MODIFY-MEASUREMENT-PERIODS statement). If the same data records remain active during monitoring, the workload is inversely proportional to the duration of the monitoring cycle. The second and third cases only apply to the monitoring program TASK.

TASK user-specific monitoring program

The system load caused by the TASK user-specific monitoring program activated by means of the /START-TASK-MEASUREMENT command is primarily due to the PCounter statistics and the SVC statistics.

Program counter statistics

The system load caused by program counter statistics depends on the following factors:

- the sampling cycle selected
- the number of tasks monitored

7.2.3 Main memory utilization

SM2 requires storage space for its code and for tables and I/O buffers.

The subsystems for the monitoring programs RESPONSETIME and UTM are loaded subsequently if required.

Some system modules are resident, others are paging modules. The size of the tables depends on the number of objects to be monitored.

7.3 Accuracy of the SM2 data

This section indicates the major factors affecting the accuracy of the data supplied. The accuracy of some particularly important variables is also discussed.

Inaccuracies resulting from rounding problems are not considered here.

7.3.1 Causes of inaccuracies

Like all software monitors, SM2 runs under the system to be monitored and thus requires certain resources for its own operation; therefore, strictly speaking, it modifies the system to be monitored. However, this influence is small and can generally be disregarded. For a description of system utilization, see [section "System resource utilization by SM2" on page 348](#).

Inaccuracies resulting from marginal problems

Ideally, the SM2 activities during monitoring activation and deactivation, during the taking of samples and at the end of a monitoring cycle operand, should take no time at all. This being impossible, certain inaccuracies result. However, this effect is minimal and decreases in direct proportion to the number of actions that have to be performed at any given time (small number of tasks, devices etc. to be monitored).

Inaccuracies resulting from classification


Some SM2 values are gathered on a system-global basis, on a category-specific basis, and/or on a task-specific basis. For category-specific data, SM2 uses the category assignment that is valid at the time the data is collected (sample or event). However, SM2 does not recognize category switches for the TASK monitoring program. This is why comparisons between task data accumulated by category and category-specific SM2 data may lead to interpretation errors.

Monitoring method inaccuracies

Different inaccuracies can occur depending on the monitoring method used:

1. Event-driven monitoring method

This method supplies very precise data at the cost of increased system workload. Problems can occur only when the duration of events is monitored.



As shown above, the duration of an event (and, if applicable, an activity count) is assigned to the second monitoring cycle even though part of it should be assigned to the first cycle. The relative magnitude of the monitoring error decreases in inverse proportion to the length of the monitoring cycle.

2. Sample-driven monitoring method

The accuracy of this monitoring method is subject to the laws of statistics.

A requirement for the validity of the monitored data is that the samples are independent of the monitored events. SM2 uses the system timers to control sampling by having itself activated at regular intervals.

Hardware interrupts are used for this purpose. However, such an interrupt is not permitted whenever the CPU is in a non-interruptible state. This results in a sampling delay and thus in a certain dependency on system events.

Certain unavoidable system activities cause further delays between interrupt acceptance and sampling by SM2. If statistical independence of the samples is assumed, the accuracy of the monitored data depends on the number of samples.

An assessment of the accuracy can be obtained by using confidence intervals (e.g. deviation of not more than 1% in 99% of all cases).

It should be noted that a high sampling rate increases the system workload. Therefore, a long monitoring cycle is preferable to excessively frequent sampling.

7.3.2 Notes on particular variables

CPU utilization

The term “percentage deviation” used in the following always denotes the absolute deviation, in other words:

X and Y have a deviation of P if $\text{abs}(X-Y) \leq P$.

TU/TPR/SIH time, IDLE time

The operating system uses timers to determine three time components: the total for TU and TPR times, the SIH time and the IDLE time. Consequently, falsifications due to the non-interruptible SIH code cannot occur. Investigations have shown that this timer is sufficiently precise for realistic monitoring cycles. As a rule, comparative measurements made with hardware monitors under typical load conditions have shown deviations of less than 1%.

SM2 divides the total for TU and TPR times into separate values for the TU time and the TPR time on the basis of samples taken during the sampling period. The accuracy of this procedure depends on the number of samples, which means that large deviations are possible if few samples are taken.

Device utilization

All devices except disks which support RSC I/Os

The utilization is recorded by SM2 using the sample-driven monitoring method.

As regards the hardware, the device is busy from the time an I/O request is received until data transfer terminates. As regards the software, the device is busy from the time an I/O request arrives in the I/O system module until channel termination.

The values determined by the software monitor should therefore always be higher than the actual device utilization. However, this systematic monitoring error is virtually balanced out by the sampling delays caused primarily by the non-interruptible system states.

Disks which support RSC I/Os

The device utilization is calculated by SM2 from the average number of RSC I/Os per second, where: 4 RSC-I/Os = 100% utilization.

Channel utilization

Type S channels on S servers

The utilization is recorded by SM2 using the sample-driven monitoring method.

SM2 uses a special machine instruction to determine whether a channel is active or not. Thus there are no deviations between hardware and software monitors in the interpretation of this condition (see also restrictions on the ability of monitoring the channel load in [section "CHANNEL report" on page 235](#)). The sampling delays caused primarily by non-interruptible system states result in the values determined by SM2 always being smaller than the actual values.

Type FC channels on servers with /390 architecture

The utilization is calculated by SM2 from the data rate and the blocking of the I/Os.

You will find the data on which the calculation is based in the Performance Handbook [5].

7.3.3 Measures for reducing the error sources

From the above explanations, the following recommendations can be deduced; they should be observed if high importance is placed on the accuracy of the data.

1. Minimization of marginal problems

The monitoring cycles should be long but not too long. Short-term peak loads will not be detected if the monitoring cycles are too long because SM2R1 only supplies the total activity or the average value for the whole monitoring cycle.

All desired monitoring functions should be activated before (not during) the monitoring cycle.

Similarly, the desired monitoring functions should be kept active somewhat longer than the monitoring cycle. For later analysis with SM2R1 or ANALYZER the desired monitoring cycle should be specified as the analysis period.

2. Reduction of the number of objects to be monitored

This reduces the induced system workload, which results in less falsification of the monitored system.

3. Sufficiently large number of samples

As a rule, this should be implemented not by increasing the sampling rate but by lengthening the monitoring cycle. The sampling rate should be increased only if the system workload permits or if it is not possible to lengthen the monitoring cycle.

7.4 Special applications

7.4.1 SM2 together with VM2000

If monitoring is executed under VM2000 control within BS2000 systems, the following special considerations are applicable for the implementation of the SM2 program and for the interpretation of the monitored data. Generally speaking, SM2 is executable under any guest system and provides monitored data relevant to the respective local guest system.

Notes on the CPU utilization values

- The sum of TU + TPR + SIH + IDLE is relativized in every guest system to 100%. This applies for
 - SM2 in the ACTIVITY, PERIODIC TASK and CPU reports (NORMED column)
 - SM2R1 in the report groups CPU (report 1)
(but not in the report groups PERIODIC-TASK-JOBNAME (report 161), PERIODIC-TASK-TSN (report 153) and PERIODIC-TASK-USERID (report 157))
 - INSPECTOR in the report groups CPU (TotalTime, NORM, CPUTUTime[%], CPUTPRTime[%], CPUSIHTime[%] and CPUIdleTime[%] reports) and PERIODIC-TASK (JOBCPU, TSNCPU and UIDCPU reports)
 - ANALYZER in the report groups CPU (“Utilization normed” and “Utilization normed (processor splitting)” reports)
(ut not in the report groups PERIODIC-TASK-JOBNAME, PERIODIC-TASK-TSN and PERIODIC-TASK-USERID (“CPU time” report))

The result of this relativation is the manual formation of the sums of the percentage load for all guest systems, **not** the total CPU utilization.

- The non-relativized values for TU, TPR, SIH and IDLE, i.e. the real utilization values ascertained by the guest system concerned, are output in
 - SM2 in the CPU report (REAL column)
 - SM2R1 in the report groups CPU (report 2) and PERIODIC-TASK-JOBNAME (report 161), PERIODIC-TASK-TSN (report 153) and PERIODIC-TASK-USERID (report 157))
 - INSPECTOR in the report group CPU
(TotalTimeReal, REAL, CPUTUTimeReal[%], CPUTPRTimeReal[%], CPUSIHTimeReal[%] and CPUIdleTimeReal[%] reports)

- ANALYZER in the report groups CPU (“Utilization real” und “Utilization real (processor splitting)” reports) and PERIODIC-TASK-JOBNAME, PERIODIC-TASK-TSN and PERIODIC-TASK-USERID (“CPU time” report)

The relationship between the data in the CPU report (REAL column), the data in the VM report (CPU MEAS) and the data in the PERIODIC TASK report (CPU%) is explained in detail in the following section.

- Furthermore, non-relativized values for TU + TPR are output on a category-specific basis in
 - SM2 in the CATEGORY report
 - SM2R1 in the report group CATEGORY-CPU (62 report)
 - INSPECTOR in the report group CATEGORY (CPU report)
 - ANALYZER in the report group CATEGORY-CPU (“CPU utilization (TU+TPR) for category” report)
- The time equivalent of the CPU time per category is also based on non-relativized times in
 - SM2R1 in the report group RST (60 report)
 - ANALYZER in the report group RST (“CPU resource service time for category” report)
- The CPU time given in the task statistics of SM2R1 is likewise the real CPU time consumed by a task in processor states TU + TPR.
- A summary of the overall CPU utilization, distribution of the CPU capacity among the individual guest systems, and the CPU time used by the individual guest systems is represented for the SM2 which is brought to execution on the monitor VM.
- If SM2 is running on the “normal” guest system, only the values of the local guest system, the local CPU pool and the local VM group are output. The utilization of the channels type S is an exception, as the overall utilization is always output here.
- For information on the accuracy of the TU and TPR times, see [“Notes on particular variables” on page 354](#).



The corresponding values can also be output by using the VM2000 command SHOW-VM-STATUS (see the “VM2000” manual [17]).

Correlation between VM, CPU and PERIODIC TASK reports

VM report:

D016ZE04 SM2 VM		CYCLE: 60 S		<date>	<time>				
HYPERVISOR: IDLE: 15.3%		ACTIVE: 7.2%		MEAS: 77.5%	# USABLE CPUS: 3				
IND	NAME	VM-GROUP	CPU-POOL	I	MEMORY (MB)	MAX	CPU (%) PLAN	MEAS	# CPUS
2	D016ZE04		*STDPPOOL		3840	66	40	41.0	2

USABLE CPUS is the number of real CPUs, available to BS2000 guest systems (also extra CPUs). This can be less than the number of real CPUs which exist.

CPU MEAS (%) is the VM's share of the performance provided by the CPUs currently available:

$$\text{CPU MEAS (\%)} = (\text{CPU time} * 100) / (\text{CYCLE} * \# \text{ USABLE CPUS})$$

where CPU time is the total measured CPU time for the VM and CYCLE is the length of the SM2 monitoring cycle.

CPU PLAN (%) also relates to the server performance currently available.

Servers with /390 architecture

The sum of the hypervisor's active time and the shares of all VMs shows the overall utilization of the server:

$$\text{CPU total (\%)} = \text{HYPERVISOR ACTIVE (\%)} + \text{sum CPU MEAS (\%)}$$

or

$$\text{CPU total (\%)} + \text{HYPERVISOR IDLE (\%)} = 100 \%$$

Servers with x86 architecture

The sum of the active time of the domain DOM0 (which can be obtained from the VM CPU POOL report) and the percentages of all VMs presents the total utilization of the server:

$$\text{CPU total (\%)} = \text{CPU (\%)} \text{ of CPU pool * POOL0} + \text{Sum CPU MEAS (\%)}$$

CPU report:

D016ZE04 SM2 CPU					CYCLE: 60 S		<date> <time>			
NORMED TO 100 %					REAL					
LM	TU %	TPR %	SIH %	IDLE %	STOP %	TU %	TPR %	SIH %	IDLE %	
AVG	8.4	26.5	26.3	38.6	0.0	8.2	26.0	25.8	37.8	
1	9.9	25.1	26.1	38.7	0.0	9.7	24.6	25.6	38.0	
2	0.0	0.0	0.0	0.0	100.0	0.0	0.0	0.0	0.0	
3	6.9	27.9	26.5	38.4	0.0	6.8	27.4	26.0	37.7	

The real CPU utilization relates to the number of virtual CPUs of the VM which are currently active.

$$TU\% + TPR\% + SIH\% = (\text{CPU time} * 100) / (\text{CYCLE} * \# \text{CPUS})$$

This provides the following correlation with CPU MEAS (%) from the VM report:

$$\text{CPU MEAS} (\%) = (TU\% + TPR\% + SIH\%) * \# \text{CPUS} / \# \text{USABLE CPUS}$$

The following correlation exists between the normed and real CPU times in the CPU report:

$$TU \text{ normed } \% = (TU\% * 100) / (TU\% + TPR\% + SIH\% + IDLE\%)$$

and analogously for TPR normed % and SIH normed %.

PERIODIC TASK report:

D016ZE04 SM2 PERIODIC TASK (S-U) CYCLE: 60 S									
<date> <time>									
TU + TPR= 35.0 % SIH= 26.4 % IDLE= 38.6 % IO(1/S)= 3724.1 (2 LM'S)									
TSN	USER-ID	JOBNAME	T	SERV-RATE	CPU%	IO(1/S)	UPG	PAG RD	CRYPT
BCAM			S	28898.1	0.2	3424.3	392	0.0	0.0
4H5N	NEI	ENNO	B	18085.9	10.4	142.9	206	0.0	0.0
4HZF	TSOS	PALL301N	B	13351.9	3.8	4.1	2767	0.0	0.0
BCAO			S	10240.1	7.6	0.0	6	0.0	0.0
41B1	ZAC	E	D	8629.4	5.0	98.8	2053	0.0	0.0
DM			S	3208.2	2.3	11.4	23	0.0	0.0
4H1Y	TSOS	SHCUSERT	T	1337.1	0.9	6.2	83	0.0	0.0
4H1X	TSOS	SHCUSERT	T	1238.0	0.8	5.7	83	0.0	0.0
3RFE	NEI	INSPECTR	B	941.3	0.5	0.0	3791	0.0	0.0
SM2G			S	792.6	0.5	0.0	794	0.0	0.0
MSCF			S	540.0	0.4	0.0	18	0.0	0.0
NSMS			S	400.3	0.3	0.0	30	0.0	0.0
OHA7	ORAQAX09	MMNLQX9	T	396.9	0.2	0.0	3308	0.0	0.0
4HOJ	NEI	ENNO	D	338.3	0.2	0.4	1257	0.0	0.0
4ICU	TSOS	RMF#4ICS	B	296.3	0.1	0.4	2404	0.0	0.0
9DK8	TSOS	HSMSSERV	B	250.5	0.2	2.4	266	0.0	0.0

CPU% is calculated like the standardized CPU time in the CPU report.

The values from the CPU% column totalled for all tasks provide TU + TPR.

The following applies for the CPU time of a task:

$$\text{CPU\%} < 100 / \text{LM'S}$$

where LM'S = number of active logical machines.

The values for TU + TPR, SIH and IDLE are normed values, i.e. they match the "NORMED TO 100 %" values in the CPU report:

The CPU time of a task can be calculated from the PERIODIC TASK report and CPU report:

$$\text{CPU\% real} = \text{CPU\%} * (\text{TU\%} + \text{TPR\%} + \text{SIH\%} + \text{IDLE\%}) / 100$$

Non-relativized values are output in

- SM2R1 in the report groups PERIODIC-TASK-JOBNAME (161 report), PERIODIC-TASK-TSN (153 report) and PERIODIC-TASK-USERID (157 report)
- ANALYZER in the report groups PERIODIC-TASK-JOBNAME, PERIODIC-TASK-TSN and PERIODIC-TASK-USERID ("CPU-time" report)


Notes on the IDLE values

The overall CPU utilization is determined by taking the HYPERVISOR-IDLE value and extrapolating the difference to 100%.


The IDLE value for a guest system given by the SM2 monitoring program loses its original meaning, in which case it merely acts as an indicator of the percentage for a given time period, one during which the respective guest system was inactive of its own volition.

Example: 2 guest systems, VMx and VMy
 Each guest system is to be active for 50 s within the time period of 100 s.

- Ideal situation: No idling occurs.



- Realistic case: idling does occur for half of the required active period (***)



The SM2 idle value of 33% is explained by the fact that SM2 precisely monitors the active time (50 s), adds the voluntary inactive time (reduced to 25 s as the result of idling), and relativizes it to 100%.

Even if server utilization is 100%, it makes sense for SM2 to provide an IDLE value for the following reasons:

Full utilization of server load can be the result of a low-priority guest system which has highly intense CPU time (in the extreme case: CPU loop), the system always obtaining control as soon as the other guest systems voluntarily revert to IDLE mode.

The SM2 IDLE value permits the operator of a high-priority guest system to estimate the extent to which additional load can be utilized. In this case the lower-priority system is idle to the appropriate degree specified.

Influence of other guest systems on monitored data

The general rule is: dilation resulting from other guest systems is always included in all SM2 monitoring values formed by the difference between two time values.

This includes the following monitoring information:

- all I/O times for the monitoring programs TASK, SYSTEM, SAMPLING-DEVICE and FILE
- response times for the monitoring program RESPONSETIME and BCAM-CONNECTION
- catalog access times for the monitoring program CMS
- QUEUE values in the task statistics
- QUEUE values for queues Q1 to Q13 for the QUEUE-TRANSITION operand

Q0 (CPU-QUEUE) is without dilation, as the real CPU requirements from the product TU + TPR are output!

The monitoring values supplied by the SERVICETIME monitoring program (only on servers with /390 architecture), namely the values

- DEVICE DISCONNECT TIME
- DEVICE CONNECT TIME
- FUNCTION PENDING TIME

represent the real hardware times (determined by DMS) without dilation by other guest systems. Dilation of I/O times is reflected in the (REMAINING) SERVICETIME value.

All SM2 monitored data with “per second” (1/s) units refers to the elapsed time, and not to the time for a guest system. As a result, the values determined are less than those which actually occur during the time the guest system was active (for example I/Os/sec. per disk, paging rate).

Disk utilization is determined by the sampling method while the corresponding guest system is still active, in other words the values are not corrupted by other guest systems (in contrast, the number of I/O operations does not refer to the active time per second, rather to the seconds elapsed, and therefore the value appears to be lower).

In the case of the type S channel utilization values, the overall utilization per channel is output. The percentage load caused by any one guest system cannot be determined here; as a consequence the NON OVERLAP value (CPU and CHANNEL simultaneously active)

has no significance when an individual guest system is involved. In the case of channels of the type FC the utilization is output by the local guest system as it is calculated from the number of inputs/outputs on the guest system.

In the case of the CHANNEL-IO monitoring program, the values supplied (e.g. the I/O operations per channel) always refer to the respective guest system. Correlation of these values on utilization is only possible for channels of type FC. It is not possible for channels of type S/2, due to the above-mentioned restriction.

The SERVICETIME monitoring program can start the recording of the detailed service times of the DCS from only one guest system at any particular time. If an attempt is made to activate the function from a second guest system as well, a corresponding warning appears. The monitoring program is started, but no DCS-specific monitored data is provided. The VM2000 command `/SHOW-VM-RESOURCES INFORMATION=*STD/*ALL` can be used to check whether SERVICETIME monitoring is already active in a guest system. If it is, message VMS2035 is output.

No DCS-specific monitored data are supplied for disks with indirect I/O (VM2000).

7.4.2 VOLUME load with DRV

The product DRV (Dual Recording by Volume) allows dual disk drives to be used. This increases the availability factor for data stored on the disks. Each write request made by DVS is sent to both disks, and each read request is executed on the disk with the shortest access time available (shortest positioning path, lowest queue length prior to the disk drive).

The duplicate of the disk has its own mnemonic device address, yet the same VSN as the original disk. When interpreting the monitoring values, the consequences of this are the following ones:

- The VSN is duplicated in the CONFIGURATION table (with different mnemonic name values).
- A VSN can also appear twice in the DEVICE report (with different monitored data).
- In Report 12 (percentage of device utilization per VSN), SM2R1 lists the monitored data of both drives; in extreme cases utilization values of up to 200% may be the case.

7.5 Sample procedures

The following procedures serve as examples for automating frequently recurring operations when the SM2 is used in normal computer center operation. They can also be found in the SYSLIB.SM2.<ver> library.

The first procedure starts short-term monitoring with output to a separate file, defining and starting monitoring programs.

The second procedure presupposes that the preceding monitoring operation has already been in progress for some time (e.g. 30 minutes). It terminates monitoring and closes the SM2 output file.

The third procedure makes the preparations for analysis with SM2R1 and updates the master output file. In the example it is assumed that the master output file is on disk. If it is on tape, the procedure has to be modified.

Input files Old master output file and the SM2 output file that was just closed.

Output file New master output file.

The fourth procedure causes the SM2 output file just closed to be analyzed.

1) Start short-term monitoring (SM2.START.MEASUREMENT)

```

/BEGIN-PROCEDURE LOGGING=C,PARAMETERS=YES( PROCEDURE-PARAMETERS=( -
/      &CYCLE   = 300,                -
/      &SAMPLE  = 500,                -
/      &TSN     = '(BCAM,BCAT)',      -
/      &DEVICE  = '(E018,E019,E01A)'  -
/      ),ESCAPE-CHARACTER=C'&' )
/ASSIGN-SYSDTA TO-FILE=*SYSCMD
/
/REMARK =====
/REMARK *** EXAMPLE 1 ***
/REMARK *** DEFAULT FOR MONITORING CYCLE   = 300 S AND ***
/REMARK *** FOR SAMPLING CYCLE             = 500 MS ***
/REMARK =====
/
/REMARK =====
/REMARK *** DEFINE SM2 OUTPUT FILE AND LOAD SM2 ***
/REMARK =====
/
/DELETE-FILE FILE-NAME=SM2.OUTPUT
/SET-JOB-STEP
/

```

```

/CREATE-FILE FILE-NAME = SM2.OUTPUT,           -
/          SUPPORT   = PUBLIC-DISK (SPACE = RELATIVE (           -
/          PRIMARY-ALLOCATION   = 576,           -
/          SECONDARY-ALLOCATION = 576))
/
/ADD-FILE-LINK LINK-NAME      = SMLINK,           -
/          FILE-NAME      = SM2.OUTPUT,           -
/          ACCESS-METHOD = SAM,           -
/          OPEN-MODE      = OUTPUT
/
/START-SM2
REMARK =====
REMARK ****   BRANCH TO ADMINISTRATION PART           ****
REMARK =====
CALL-ADMINISTRATION-PART

REMARK =====
REMARK ****   DEFINE SYSTEM MONITORING PROGRAM           ****
REMARK ****   MONITOR ALL DISKS AND TAPES           ****
REMARK =====

SET-SYSTEM-PARAMETERS DEVICES = *SPECIFIED(DEVICE = (*TAPE,*DISK))

REMARK =====
REMARK ****   DEFINE TASK MONITORING PROGRAM           ****
REMARK ****   MONITOR TASKS WITH TSNS 1111 AND 1112 AND           ****
REMARK ****   WITH THE JOB NAME TEST           ****
REMARK ****   MONITOR DISK FOOB ONLY           ****
REMARK =====

SET-TASK-PARAMETERS TASK-SELECTION = *SPECIFIED(TSN           = &TSN,           -
                                                JOB-NAME = TEST),           -
DEVICES = *SPECIFIED(DEVICE = &DEVICE)

REMARK =====
REMARK ****   DEFINE FILE MONITORING PROGRAM           ****
REMARK =====

ADD-FILE FILE-NAME=:4V05:$TSOS.TSOSCAT

REMARK =====
REMARK ****   DEFINE BCAM-CONNECTION MONITORING PROGRAM           ***
REMARK =====
SET-BCAM-CONNECTION-PARAMETERS INWAIT-BUCKETS = *STD-LIMITS, -
                                REACT-BUCKETS  = *STD-LIMITS, -
                                INPROC-BUCKETS = *STD-LIMITS, -
                                OUTPROC-BUCKETS = *STD-LIMITS

```

```

ADD-BCAM-CONNECTION-SET                                     -
    SET-NAME=DIALOG,                                     -
    CONNECTION-SELECTION = *BY-NEA-NAME(                 -
        CONNECTION-NAME = *SPECIFIED(                   -
            LOCAL-APPLICATION = $DIALOG,                 -
            PARTNER-APPLICATION = *ANY)),                 -
    HOST-SELECTION = *ANY

REMARK =====
REMARK ****  DEFINE STORAGE-SYSTEM MONITORING PROGRAM    ****
REMARK =====

SET-STORAGE-SYSTEM-PARAMETERS                             -
    ADDITIONAL-DATA = *SYMMETRIX(                       -
        TYPE = *DEVICE(                                  -
            DEVICES      = *DEVICES-WITH-MN,            -
            BACK-END-DATA = *NO ))

REMARK =====
REMARK ****  OPEN SM2 FILE, SPECIFY MONITORING CYCLE AND  ****
REMARK ****  SAMPLING CYCLE                               ****
REMARK =====

MODIFY-MEASUREMENT-PERIODS OFFLINE-PERIOD = &CYCLE,    -
    SAMPLING-PERIOD = &SAMPLE

OPEN-LOG-FILE FILE = *BY-LINK-NAME

REMARK =====
REMARK ****  START MONITORING PROGRAMS                    ****
REMARK =====

START-MEASUREMENT-PROGRAM TYPE=( *SYSTEM, *TASK, *FILE,  -
    *BCAM-CONNECTION, *STORAGE-SYSTEM)

REMARK =====
REMARK ****  TERMINATE SM2 USER PROGRAM                  ****
REMARK ****  MONITORING CONTINUES IN THE BACKGROUND !!!  ****
REMARK =====

END
/
/END-PROCEDURE

```

2) *Terminate short-term monitoring and resume former monitoring operations
(SM2.STOP.MEASUREMENT)*

It is assumed that SM2 was started with the preceding procedure.

```

/BEGIN-PROCEDURE LOGGING=C,PARAMETERS=YES( PROCEDURE-PARAMETERS=(      -
/      &CYCLE   = 300,              -
/      &SAMPLE  = 500              -
/      ),ESCAPE-CHARACTER=C'&' )
/ASSIGN-SYSDTA TO-FILE=*SYSCMD
/
/REMARK  =====
/REMARK  ***   EXAMPLE 2                               ***
/REMARK  ***   DEFAULT FOR MONITORING CYCLE           = 300 S AND   ***
/REMARK  ***   FOR SAMPLING CYCLE                     = 500 MS     ***
/REMARK  =====
/
/REMARK  =====
/REMARK  ***   DELETE NON-CATALOGED FILE              ***
/REMARK  =====
/
/DELETE-FILE FILE-NAME=SM2.CHANGE
/SET-JOB-STEP
/
/REMARK  =====
/REMARK  ***   LOAD SM2 USER PROGRAM                  ***
/REMARK  =====
/
/START-SM2
REMARK  =====
REMARK  ****   BRANCH TO ADMINISTRATION PART AND TERMINATE   ****
REMARK  ****   ALL ACTIVE MONITORING PROGRAMS                ****
REMARK  =====
CALL-ADMINISTRATION-PART

STOP-MEASUREMENT-PROGRAM TYPE=( *SYSTEM,*TASK,*FILE, -
                                *BCAM-CONNECTION, *STORAGE-SYSTEM)

REMARK  =====
REMARK  ****   CLOSE SM2 FILE                             ****
REMARK  =====

CLOSE-LOG-FILE

REMARK  =====
REMARK  ****   THE MONITORING PROGRAM DEFINITIONS ARE       ****
REMARK  ****   STILL VALID                                  ****
REMARK  =====

```

```
REMARK =====
REMARK ****  TERMINATE SM2 USER PROGRAM          ****
REMARK ****  RENAME CLOSED FILE                  ****
REMARK =====

END
/
/MODIFY-FILE-ATTRIBUTES FILE-NAME = SM2.OUTPUT,      -
/                               NEW-NAME  = SM2.CHANGE,  -
/                               SUPPORT   = PUBLIC-DISK (SPACE=RELEASE(100))
/
/END-PROCEDURE
```


3) Prepare for SM2R1 analysis with the SM2U1 utility (SM2U1.PREPARE)

Update the master output file with all the records with the exception of those for the TASK monitoring program.

```

/BEGIN-PROCEDURE LOGGING=C,PARAMETERS=YES( PROCEDURE-PARAMETERS=(      -
/      &SM2UIN      = SM2.CHANGE,          -
/      &SM2UTASK    = SM2.SAM.TASK,        -
/      &SM2MASTER  = SM2.MASTER          -
/      ),ESCAPE-CHARACTER=C'&')
/ASSIGN-SYSDTA TO-FILE=*SYSCMD
/
/REMARK =====
/REMARK ***   ASSIGN INPUT FILE                               ***
/REMARK =====
/
/ADD-FILE-LINK LINK-NAME = SM2U1,      -
/      FILE-NAME = &SM2UIN
/
/REMARK =====
/REMARK ***   ASSIGN OUTPUT FILE (SM2 MASTER FILE)         ***
/REMARK =====
/
/CREATE-FILE FILE-NAME = &SM2MASTER
/SET-JOB-STEP
/
/ADD-FILE-LINK LINK-NAME = SM2U0,      -
/      FILE-NAME = &SM2MASTER
/
/REMARK =====
/REMARK ***   LOAD SM2U1, UPDATE SM2 MASTER FILE AND ALL   ***
/REMARK ***   RECORDS EXCEPT THOSE FOR THE TASK AND DISK ***
/REMARK ***   MONITORING PROGRAMS                           ***
/REMARK =====
/
/START-SM2U1

SELECT-MEASUREMENT-GROUPS SELECTION = *ALL(EXCEPT = *TASK-STATISTICS)
END
/
/REMARK =====
/REMARK ***   ASSIGN INPUT FILE                               ***
/REMARK =====
/
/ADD-FILE-LINK LINK-NAME = SM2U1,      -
/      FILE-NAME = &SM2UIN
/

```

```
/REMARK =====
/REMARK ***   ASSIGN OUTPUT FILE FOR TASK RECORDS           ***
/REMARK =====
/
/DELETE-FILE FILE-NAME = &SM2UTASK
/SET-JOB-STEP
/
/CREATE-FILE FILE-NAME = &SM2UTASK
/
/ADD-FILE-LINK LINK-NAME = SM2U0,      -
/          FILE-NAME = &SM2UTASK
/
/REMARK =====
/REMARK ***   LOAD SM2U1, CREATE SM2 FILE FOR TASK           ***
/REMARK ***   ANALYSIS.                                     ***
/REMARK =====
/
/START-SM2U1
      SELECT-MEASUREMENT-GROUPS SELECTION=*TASK-STATISTICS

      END
/
/END-PROCEDURE
```

4) Perform analysis with SM2R1 (SM2R1.EVALUATION)

This procedure analyzes the SM2 output file just closed.

```

/BEGIN-PROCEDURE LOGGING=C,PARAMETERS=YES( PROCEDURE-PARAMETERS=(      -
/      &SM2OUT      = SM2.CHANGE,      -
/      &SM2R1OUT    = SM2R1.OUT,      -
/      &SM2UTASK    = SM2.SAM.TASK,    -
/      &SM2R1OTASK  = SM2R1.OUT.TASK   -
/      ),ESCAPE-CHARACTER=C'&')
/
/MODIFY-TERMINAL-OPTIONS OVERFLOW-CONTROL = NO-CONTROL
/
/ASSIGN-SYSDTA TO-FILE = *SYSCMD
/
/DELETE-FILE FILE-NAME = &SM2R1OUT
/SET-JOB-STEP
/
/REMARK =====
/REMARK ***          LOAD SM2R1 AND ANALYZE SM2 FILE          ***
/REMARK =====
/
/START-SM2R1 MONITOR-FILE-NAME = &SM2OUT,      -
/          LIST-FILE-NAME      = &SM2R1OUT
PRINT-CONFIGURATION
SET-TITLE TEXT='*** STANDARD-STATISTIKEN ***'
PRINT-REPORTS REPORT-LIST = (*STD, *FILE, -
                          *STORAGE-SYSTEM-SYMMETRIX)

PRINT-QUEUE-TRANSITION
END
/
/REMARK ====='
/REMARK ***          PRINT SM2R1 OUTPUT FILE          ***'
/REMARK ====='
/
/PRINT-DOCUMENT FROM-FILE          = &SM2R1OUT,      -
/          DELETE-AFTER-PRINT     = *YES,          -
/          DOCUMENT-FORMAT        = *TEXT(          -
/          LINE-SPACING=*BY-EBCDIC-CONTROL)
/
/REMARK ====='
/REMARK ***          ANALYZE TASK          ***'
/REMARK ====='
/
/DELETE-FILE FILE-NAME = &SM2R1OTASK
/SET-JOB-STEP
/

```

```
/REMARK ====='  
/REMARK   ***      LOAD SM2R1 AND ANALYZE SM2 FILE          ***'  
/REMARK   ***      (TASK RECORDS)                          ***'  
/REMARK ====='  
/  
/START-SM2R1 MONITOR-FILE-NAME = &SM2UTASK,      -  
/          LIST-FILE-NAME      = &SM2R1OTASK  
  SET-TITLE TEXT='*** TASKSTATISTIK ***'  
  PRINT-TASK-STATISTICS INFORMATION=*HIGH  
  END  
/  
/REMARK ====='  
/REMARK   ***      PRINT SM2R1 OUTPUTFILE                  ***'  
/REMARK ====='  
/  
/PRINT-DOCUMENT FROM-FILE          = &SM2R1OTASK,  -  
/          DELETE-AFTER-PRINT     = *YES,         -  
/          DOCUMENT-FORMAT        = *TEXT(        -  
/                                  LINE-SPACING=*BY-EBCDIC-CONTROL)  
/  
/MODIFY-TERMINAL-OPTIONS OVERFLOW-CONTROL=USER-ACKNOWLEDGE  
/END-PROCEDURE
```

8 SM2U1 utility routine

The SM2U1 utility routine serves to edit and manage the SM2 output files. SM2U1 support:

- conversion of PAM output files to SAM output files
- merging of several SM2 output files
- output of information on the contents of SM2 output files
- selection of SM2 output records
- splitting of SM2 output files
- compression of output files.

Depending on the specifications made by the SM2 administrator, SM2 enters the monitoring data in a PAM file or an SAM file.

PAM files must be converted to SAM files for processing by the SM2R1 analysis routine. The SM2U1 utility routine performs this conversion operation. After conversion, the file format corresponds to that of SAM output files.

The SM2U1 utility routine can also be used if the SM2 administrator wants to merge several SM2 output files into a single file. Up to 99 PAM or SAM output files can be merged into one SAM output file during an SM2U1 run.

With SM2U1, an SM2 output file containing several SM2 measurements can be split up into individual files.

Furthermore, records can be selected which are to be transferred to the master output file for long-term statistics.

The compression function of SM2U1 allows disk space to be saved when output files are merged. It does this by compressing the output file, producing a result as if SM2 had run with a monitoring cycle of an hour.



SM2U1 also supports the merging of multiple SM2 output files created as the result of user task measurements, merging them into a single SM2 output file. The other SM2U1 functions cannot be employed for user-specific SM2 output files.

Users

SM2U1 does not require a privileged user ID. Any user who has access to the SM2 output file and to SM2U1 can use this utility routine.

Interruption-free clock resetting

SM2U1 works internally with UTC time. The program can also process files with “double” hours (i.e. when a clock has been put back).

8.1 Preparing the SM2U1 session

Prior to calling SM2U1, the program must be provided with the input and output files, together with the corresponding link names. In the case of input files, SM2U1 checks whether a user-specific file is assigned to the file link name. If it is, only the MERGE-FILES statement is provided. In all other cases (including errors on opening the first input file), SM2U1 assumes that the file is a system-global SM2 output file and makes all SM2U1 statements available.

Depending on the SM2U1 function required, three different applications are possible:

a) Merge files function

(Statements: MERGE-FILES (for user-specific and system-global files), SELECT-MEASUREMENT GROUPS, END):

SM2U1 processes as many input files as there are link names, in ascending order. The input files are specified using the ADD-FILE-LINK command. The corresponding file names are: SM2U11, SM2U12, .. , SM2U199. Up to 99 files can be merged; they can be either PAM or SAM type files. The output file is likewise assigned via a ADD-FILE-LINK command using the link name SM2U0.

```
/ADD-FILE-LINK LINK-NAME=SM2U11,FILE-NAME=SM2-output-file1
/ADD-FILE-LINK LINK-NAME=SM2U12,FILE-NAME=SM2-output-file2
.
.
.
/CREATE-FILE FILE-NAME=SM2-output-file,
              SUPPORT=...(VOLUME=...,DEVICE-TYP=...,SPACE=...)
/ADD-FILE-LINK LINK-NAME=SM2U0,FILE-NAME=SM2-output-file,
              OPEN-MODE=...
```

b) Information output function

(Statement: SHOW-INFORMATION):

SM2U1 processes as many input files as there are link names, in ascending order. The ADD-FILE-LINK command is used to specify the input files. The link names of the files are then SM2U11, SM2U12, ..., SM2U199 (i.e. up to another 99 input files).

```
/ADD-FILE-LINK LINK-NAME=SM2U11,FILE-NAME=SM2-output-file1
/ADD-FILE-LINK LINK-NAME=SM2U12,FILE-NAME=SM2-output-file2
.
.
.
```

c) Separate files function

(Statement: SEPARATE-FILES):

SM2U1 processes an input file with the link name SM2UI. Up to 99 output files can be specified using the link names SM2UO1... SM2UO99.

```

/ADD-FILE-LINK LINK-NAME=SM2UI,FILE-NAME=SM2-output-file
/CREATE-FILE FILE-NAME=SM2-output-file1,
              SUPPORT=...(VOLUME=...,DEVICE-TYP=...,SPACE=...)
/CREATE-FILE FILE-NAME=SM2-output-file2,
              SUPPORT=...(VOLUME=...,DEVICE-TYP=...,SPACE=...)
.
.
.
/ADD-FILE-LINK LINK-NAME=SM2UO1,FILE-NAME=SM2-output-file1,
              OPEN-MODE=...
/ADD-FILE-LINK LINK-NAME=SM2UO2,FILE-NAME=SM2-output-file2,
              OPEN-MODE=...
.
.
.

```



Following successful processing of the statement, file names SM2UI1 through SM2UI99 of the input file(s) or SM2UO1 through SM2UO99 of the output files are released by the program. SM2U1 outputs the following message for each output file: NUMBER OF RECORDS WRITTEN: number

8.2 Starting and terminating SM2U1

The SM2U1 program is called by means of the BS2000 command /START-SM2U1

Format

START-SM2U1

```
VERSION = *STD / <product-version 6..10> / <product-version 4..8 without-corr> /
           <product-version 3..7 without-man>
,MONJV = *NONE / <full-filename 1..54 without-gen-vers>
,CPU-LIMIT = *JOB-REST / <integer 1..32767>
```

Operands

VERSION =

Specifies the program version of SM2U1 to be called.

VERSION = *STD

The current program version is called.

VERSION = <product-version>

The specified program version is called.

MONJV =

Specifies the name of the job variable to monitor the SM2U1 run. The job variable must already be cataloged.

MONJV = *NONE

No job variable is specified.

MONJV = <full-filename 1..54 without-gen-vers>

Specifies the name of an already cataloged job variable.

CPU-LIMIT =

Specifies the CPU time that SM2U1 can use during execution. If this time is exceeded in interactive mode, the user is notified by the system; in batch mode, the SM2U1 run is terminated.

CPU-LIMIT = *JOB-REST

There is no time limit for the program.

CPU-LIMIT = <integer 1..32767>

Specifies the CPU time that SM2U1 can use during execution.

The SM2U1 statement END in turn, terminates SM2U1 once processing of the statements is completed.

8.3 SM2U1 statements

Using the statements of the SM2U1 utility, SM2 output files are readied for further processing.

Statement	Function
SHOW-INFORMATION	Request information on the monitoring programs and monitoring cycles given in the system-global SM2 output files.
MERGE-FILES	Merge user-specific and system-global monitoring files generated during various monitoring operations.
SET-COMPRESSION	Turn compression function on or off for system-global monitoring files on an off to merge multiple cycles in the input files into a single cycle in the output file.
SELECT-MEASUREMENT-GROUPS	Select records for the conversion of system-global SM2 output files.
SEPARATE-FILES	Separate SM2 output files into their components and distribute these to the output files. Separate faulty files.
END	Terminate SM2U1 or start the merging of system-global SM2 output files.

The statements are described in alphabetical order.



The internal program name for the syntax check of SM2U1 statements is SM2U1 (or SM2T1 for the processing of user-specific SM2 output files).

END

Terminate SM2U1

The END statement marks the end of statement entry and the beginning of processing. Following termination of processing, the program session is terminated.

Format

END



Provided that other functions are not explicitly requested using the SM2U1 statements described in the following sections, the END statement is the one always used to initiate merging of system-global SM2 output files. (In this case END has the same effect as MERGE-FILES TYPE=*MONITOR-FILE; END.)

The MERGE-FILES statement must be specified explicitly when processing a user-specific SM2 output file.

MERGE-FILES

Merge files

The two MERGE-FILES statements can be used to merge system-global SM2 output files or to merge files resulting from various monitoring operations as part of the TASK user-specific monitoring program (/START-TASK-MEASUREMENT). SM2U1 checks the first input file and, depending on the type, assumes PA-FILE or MONITOR-FILE by default. Each input file must be assigned a link name SM2U1...SM2UIn, and the output file must be assigned the link name SM2UO.

Format 1

MERGE-FILES
TYPE = *PA-FILE

Operands**TYPE =**

Selects the type of SM2 output file.

TYPE = *PA-FILE

Merges user-specific SM2 output files (to the TASK user-specific monitoring program).

Format 2

MERGE-FILES
TYPE = *MONITOR-FILE

Operands**TYPE =**

Selects the type of SM2 output file.

TYPE = *MONITOR-FILE

Merges system-global SM2 output files.

SELECT-MEASUREMENT-GROUPS

Select records

This statement is used to select records to be transferred or not to be transferred to the output file when system-global SM2 output files are converted. The user specifies the names of report groups or SM2R1 statistics (see the relevant statements under SM2R1). The appropriate records are then transferred to the output file. If the statement is not specified, all records are transferred.

Format

SELECT-MEASUREMENT-GROUPS

```
SELECTION = *STD / *ALL(...) / list-poss(59): *AUTOMATIC-ANALYSIS / *BCAM-CONNECTION /
*BCAM-MEMORY / *CATALOG-MANAGEMENT / *CATEGORY-CPU / *CATEGORY-IO /
*CATEGORY-QUEUE / *CATEGORY-WORKING-SET / *CHANNEL / *CONFIGURATION /
*CPU / *DAB / *DEVICE / *DILATION / *DISK / *DISK-FILE / *DLM / *FILE / *GS / *GSVOL /
*HSMS-STATISTICS / *IO / *ISAM / *ISAM-FILE / *MEMORY / *MSCF / *NSM / *OPENFT /
*PCS / *PERIODIC-TASK / *PFA / *POSIX / *PRIOR-ACF / *PUBSET /
*QUEUE-TRANSITION / *RESPONSE-TIME / *RST / *STD / *SAMPLING-DEVICE /
*SERVICETIME / *SESAM-SQL / *STORAGE-SYSTEM-SYMMETRIX / *SUMMARY-ACTIVITY /
*SUMMARY-CMS / *SUMMARY-DAB / *SUMMARY-PCS / *SUMMARY-POSIX /
*SUMMARY-UTM / *SVC / *SYSTEM / *TASK / *TASK-STATISTICS / *TCP-IP / *TLM /
*UDS-SQL / *UTM / *VM2000 / *VOLUME / *WORKING-SET
```

*ALL(...)

```
EXCEPT = *NONE / list-poss(59): *AUTOMATIC-ANALYSIS / *BCAM-CONNECTION
*BCAM-MEMORY / *CATALOG-MANAGEMENT / *CATEGORY-CPU /
*CATEGORY-IO / *CATEGORY-QUEUE / *CATEGORY-WORKING-SET /
*CHANNEL / *CONFIGURATION / *CPU / *DAB / *DEVICE / *DILATION / *DISK /
*DISK-FILE / *DLM / *FILE / *GS / *GSVOL / *HSMS-STATISTICS / *IO / *ISAM /
*ISAM-FILE / *MEMORY / *MSCF / *NSM / *OPENFT / *PCS / *PERIODIC-TASK /
*PFA / *POSIX / *PRIOR-ACF / *PUBSET / *QUEUE-TRANSITION /
*RESPONSE-TIME / *RST / *STD / *SAMPLING-DEVICE / *SERVICETIME /
*SESAM-SQL / *STORAGE-SYSTEM-SYMMETRIX / *SUMMARY-ACTIVITY /
*SUMMARY-CMS / *SUMMARY-DAB / *SUMMARY-PCS / *SUMMARY-POSIX /
*SUMMARY-UTM / *SVC / *SYSTEM / *TASK / *TASK-STATISTICS / *TCP-IP / *TLM /
*UDS-SQL / *UTM / *VM2000 / *VOLUME / *WORKING-SET
```

Operands

SELECTION =

Selects the records to be transferred to the output file.

SELECTION = *STD

All records required by SM2R1 are transferred if PRINT-REPORTS REPORT-LIST = *STD is specified in the statement.

SELECTION = *ALL(...)

All records are selected except those belonging to the report groups specified with EXCEPT.

EXCEPT =

Specifies the records which are not to be transferred to the output file.

EXCEPT = *NONE

No records are excluded from transfer.

EXCEPT = list-poss(59): *AUTOMATIC-ANALYSIS / ...

The items in this list correspond to the report groups or to the SM2R1 statistics and summaries. The records of these items are not transferred.

SELECTION = list-poss(59): *AUTOMATIC-ANALYSIS / ...

The items in this list correspond to the report groups or to the SM2R1 statistics and summaries. The records which belong to these items are selected.



The SELECT-MEASUREMENT-GROUPS statement is not cumulative, in other words, if it is entered more than once, only the last statement applies.

SEPARATE-FILES

Separate SM2 output files

This statement can be used to split up a system-global SM2 output file created in the course of several monitoring cycles into its original components and distribute them to various output files. Even errored files can be split up.

The expression or operand ***BLOCK** used below defines a set of monitoring cycles in temporal succession. A block is understood to consist of at least one monitoring period and at most all monitoring periods.

Blocks are defined with the aid of indexes assigned to the monitoring time periods according to their position in the SM2 output file.

The interrelationship between index and monitoring period can be displayed using the statement **SHOW-INFORMATION=*MONITORING-PERIODS**.

The order in which the blocks are input is decisive in assigning them to the corresponding output files. Since, as a standard procedure, output files are opened and written in **EXTEND** mode, multiple blocks of the input file can be transferred to an output file in one pass. To do so, link names **SM2UO1** to **SM2UOn** must be used for the same physical output file.

Format

SEPARATE-FILES

```
FILES = list-poss(99): *BLOCK(...)
    *BLOCK(...)
        | FIRST-INDEX = <integer 1..32000>
        | ,LAST-INDEX = *SAME / *LAST / <integer 1..32000>
```

Operands

FILES = list-poss(99): *BLOCK(...)

The specified file blocks are separated. An output assignment must be effected for each file block using the **/ADD-FILE-LINK** command.

FIRST-INDEX = <integer 1..32000>

Specifies the index of the first monitoring period in a block to be separated. The index must be greater than the **LAST-INDEX** value of the preceding value.

LAST-INDEX =

Specifies the index of the last monitoring period in the block to be separated.

LAST-INDEX = *SAME

The last index is the same as the first one, i.e. precisely one monitoring cycle is separated for this block.

LAST-INDEX = *LAST

The file is separated up to the end of the input file.

LAST-INDEX = <integer 1..32000>

The index of the last monitoring period in the block to be separated is specified explicitly; the index must be greater than or equal to the associated FIRST-INDEX value.

Example

The first and second monitoring period of a merged SM2 output file (with at least three monitoring periods) with the link name SM2U1 is to be assigned to a file with the link name SM2UO1, and the remaining monitoring periods are to be assigned to a file with the link name SM2UO2.

```
SEPARATE-FILES FILES = (*BLOCK(1,2), *BLOCK(3,*LAST))
```

SET-COMPRESSION**Turn SM2U1 compression function on or off**

The SET-COMPRESSION statement turns the SM2U1 compression function on or off. The statement for compression can be specified when the SM2 output files are merged. As a result of compression, a number of cycles in the input files are merged into a single cycle in the output file. The monitoring cycle of the output file is an hour. The output file corresponds to an SM2 output file written by SM2 in the same period of time with a monitoring cycle of an hour.

Format

SET-COMPRESSION
COMPRESSION = <u>*ON</u> / *OFF

Operands

COMPRESSION =

Specifies whether the compression function is turned on or off.

COMPRESSION = *ON

The compression function is turned on.

COMPRESSION = *OFF

The compression function is turned off.

Restrictions

- In SM2R1 and ANALYZER report 74 and the “Request delay for category” report (report group PCS) cannot be output with a compressed file.
- In the compressed file, only groups of monitored objects present for the whole hour in the uncompressed file are output. Groups of monitored objects that would not be output by SM2 in an offline cycle of an hour in the same time period are not included either.
- If there are missed records in the uncompressed file, even more records are missing in the compressed file. It is essential to avoid missed records.
- Record 66 is not compressed (monitoring data for the QUEUE-TRANSITION).
- Record 67 (monitoring data for MSCF) is not compressed.

As a result of compression, there is not enough data for an accurate bottleneck analysis, since existing maximum and minimum values are smoothed out. Thus, the compressed file is suitable for long-term statistics and trend analysis.

As a result of the compression of output files with the default monitoring cycle of 150 seconds, a compression factor of approximately 15 to 20 is achieved. The evaluation of a compressed file with SM2R1 or ANALYZER is accelerated accordingly.

SHOW-INFORMATION

Output information on monitoring programs and monitoring cycles

The user can employ this statement to obtain information on the monitoring programs and monitoring cycles contained in a system-global SM2 output file. Output covers all input files with link names SM2U11 to SM2U199. If SHOW-INFORMATION is entered as the only SM2U1 statement, then the link name of the output file does not have to be specified. In this case a simple screen output of the requested information is generated.

If the link name is supplied for output or if SHOW-INFORMATION is entered in combination with SELECT-MEASUREMENT-GROUPS or MERGE-FILES, merging of input files is also initiated once the END statement has been completely processed.

Format

SHOW-INFORMATION
INFORMATION = <u>*MONITORING-PERIODS</u> / list-poss(2): *STATISTIC-PERIODS / *MONITORING-PERIODS

Operands

INFORMATION =

Selects the type of information.

INFORMATION = *MONITORING-PERIODS

Outputs information on the individual monitoring cycles as well as current file names. This operand is particularly important for output of the directories of files where more than one measurement is involved.

The following individual information is output:

- date and time the file was opened
- date and time the file was closed
- number of SM2 monitoring cycles involved
- index of the monitoring period (important for splitting up SM2 output files).

INFORMATION = *STATISTIC-PERIODS

Information pertaining to start and stop times is output for the monitoring programs contained in the SM2 output file. This information is required in order to be able to select individual monitoring programs for analysis.

The following individual information is output:

- date and time the monitoring program started; name of the monitoring program (for the TASK and SYSTEM monitoring programs)
- date and time the monitoring program stopped; name of the monitoring program (for the TASK and SYSTEM monitoring programs)
- date and time of the CLOSE record (after abnormal termination of SM2 and after monitoring programs have been started, provided there is a CLOSE record).

8.4 Notes on the output files

The user can control the output medium, memory allocation and OPEN mode (OUTPUT or EXTEND mode). As a standard procedure, SM2U1 output files are opened using OPEN-MODE=*EXTEND.

Similarly to the SM2 output files, the following applies to the output files: restrictions apply to the block length specification.

The user must either assign BUFFER-LENGTH=*STD(SIZE=16) or, if the user does not define a block length, the value 16 is assigned by SM2U1. The primary allocation must be at least twice as large and the secondary allocation must be at least the same size as the block length. If these conditions are not met, the primary allocation is set to 32 and the secondary allocation is set to 96.

After the output file or files are closed, PAM pages are released provided the following conditions have been satisfied:

- the file is not a dummy file
- the pages were already reserved for the file prior to the SM2U1 run

In the case of output files on tape, the following consideration must be borne in mind.

- Owing to the extensive checking required, continuous writing to a tape (OPEN-MODE=*EXTEND) often involves extensive time-consuming spooling procedures.

8.5 Plausibility checks

The SM2U1 utility routine is aware of the SM2 output file format and performs the following checks for system-global SM2 output files:

1. Are the input files SM2 output files?

The presence of the OPEN record at the beginning of a file indicates that the input file is an output file.

2. Have the individual input files been closed with a CLOSE record?

If not, a system failure CLOSE record is created and a message is output to this effect.

3. Are the dates in the SM2 cycles in ascending order?

The following plausibility check is performed in conjunction with the “merge user-specific SM2 output files” function:

Are the input files user-specific SM2 output files?

The criterion for this is the corresponding START record of the user task monitoring facility.

8.6 Notes on implementation

If an SM2 output file has not been closed properly (e.g. due to a system failure), the file must be repaired prior to processing with SM2U1 by issuing the command

```
/REPAIR-DISK-FILES FILE-STATUS=*ANY(FILE-NAME=filename)
```

SM2 output files are opened using the operand OPEN-MODE=*INPUT. In the event of DMS error messages on any file, the SM2U1 session will be aborted.

8.7 Examples for the SM2U1 utility routine

The following examples can also be found in the SYSLIB.SM2.<ver> library.

Example: SM2U1.PAM.TO.SAM

Convert a system-global PAM to a SAM output file.

All of the SM2 records are to be copied.

```
/BEGIN-PROCEDURE LOGGING=CMD,PARAMETERS=*YES(PROCEDURE-PARAMETERS=
  (&FILEIN,&FILEOUT),ESCAPE-CHARACTER=C'&')
/ASSIGN-SYSDTA TO-FILE=*SYSCMD
/ADD-FILE-LINK LINK-NAME=SM2U1,FILE-NAME=&FILEIN
/CREATE-FILE FILE-NAME=&FILEOUT,SUPPORT=*PUBLIC-DISK(SPACE=
  *RELATIVE(PRIMARY-ALLOCATION=576,SECONDARY-ALLOCATION=576))
/ADD-FILE-LINK LINK-NAME=SM2UO,FILE-NAME=&FILEOUT,BUFFER-LENGTH=*STD(SIZE=16)
/START-SM2U1
//END
/END-PROCEDURE
```

Example: SM2U1.SELECT.MEASUREMENT.GROUPS

Extend the system-global SM2 output file and copy all records required for an SM2R1 analysis using PRINT-REPORTS ...,REPORT-LIST=*STD.

```
/BEGIN-PROCEDURE LOGGING=CMD,PARAMETERS=*YES(PROCEDURE-PARAMETERS=
  (&FILEIN,&FILEOUT),ESCAPE-CHARACTER=C'&')
/ASSIGN-SYSDTA TO-FILE=*SYSCMD
/ADD-FILE-LINK LINK-NAME=SM2U1,FILE-NAME=&FILEIN
/CREATE-FILE FILE-NAME=&FILEOUT,SUPPORT=*PUBLIC-DISK(SPACE=
  *RELATIVE(PRIMARY-ALLOCATION=576,SECONDARY-ALLOCATION=576))
/ADD-FILE-LINK LINK-NAME=SM2U0,FILE-NAME=&FILEOUT,BUFFER-LENGTH=*STD(SIZE=16)
/START-SM2U1
//SELECT-MEASUREMENT-GROUPS SELECTION = *STD
//END
/END-PROCEDURE
```

Example: SM2U1.SEPARATE.FILES

A file which is a merged file is to be split up into two files: the first monitoring period and the rest of the monitoring period.

```
/BEGIN-PROCEDURE LOGGING=CMD,PARAMETERS=*YES(PROCEDURE-PARAMETERS=
  (&FILEIN,&FILEFIRST,&FILEREST),ESCAPE-CHARACTER=C'&')
/ASSIGN-SYSDTA TO-FILE=*SYSCMD
/ADD-FILE-LINK LINK-NAME=SM2UI,FILE-NAME=&FILEIN
/CREATE-FILE FILE-NAME=&FILEFIRST
/ADD-FILE-LINK LINK-NAME=SM2U01,FILE-NAME=&FILEFIRST
/CREATE-FILE FILE-NAME=&FILEREST
/ADD-FILE-LINK LINK-NAME=SM2U02,FILE-NAME=&FILEREST
/START-SM2U1
//SEPARATE-FILES FILES=(*BLOCK(1,*SAME),*BLOCK(2,*LAST))
//END
/END-PROCEDURE
```

9 SM2R1 analysis routine

The SM2R1 routine analyzes the data collected by the SM2 monitoring program and stored in a system-global SM2 output file. The result of the analysis is output in the form of charts and/or tables containing statistical data or is stored in a transfer file for further processing by the corresponding programs.

Users

SM2R1 is not confined to privileged users. Any user who has access to the SM2 output file and SM2R1 can use the analysis routine.

Control file

SM2R1 is a table-based analysis routine. The analysis control data is stored in the SYSDAT.SM2.<ver>.MTFILE control file. A valid control file is supplied with SM2 and SM2R1. SM2R1 can be executed only with the associated SYSDAT.SM2.<ver>.MTFILE.

File to be analyzed

The SM2 output file can be processed by SM2R1 only in SAM format, i.e. PAM files must be converted to SAM files using SM2U1.

Text file

SM2R1 is written in PL/1 and therefore requires the PL/1 text file for error exits. This file is supplied together with SM2 and SM2R1.

Output file

The results of an SM2R1 run are written to a file which can be output on the printer by means of the PRINT-DOCUMENT command. The output file can be freely selected.

The period of time for which the SM2 output file is to be analyzed is called the analysis period in the following.

An analysis period may encompass several SM2 monitoring cycles plus interleaved periods of time during which no monitoring took place.

ISO format

SM2R1 outputs all dates in ISO format (yy-mm-dd) and expects input of the date in the same format. June 13, 2012, for example, has to be entered as follows: 12-06-13. Dates can also be entered in the form yyyy-mm-dd.

SM2R1 result lists

1. Header page
2. Monitoring environment of first analyzed session (see [page 401](#))
3. System configuration output (see [page 420](#))
4. Actual analysis results; these are broken down as follows:
 - Time series and statistics (REPORTS)
For statistics, the average, minimum, maximum and standard deviation values are supplied for the whole output period.
For time series, the average, minimum, maximum and standard deviation values are supplied for the whole output period. In addition to the statistics display, these mean values are displayed in charts as a function of the time.
 - Summaries
The essential data of a monitoring cycle is summarized.
 - Task statistics (TASK STATISTICS)
 - Queue statistics (QUEUE STATISTICS)
 - HSMS statistics (HSMS-STATISTICS)
 - Automatic performance analysis
5. If the SM2 output file contains multiple sessions then the monitoring environment for the last analyzed session is output.
6. Information on the analysis run (see [page 403](#))
7. Table of contents (see [page 395](#))

SM2R1 transfer file

The values determined by SM2R1 are transferred to a so-called SM2R1 transfer file (SM2R1 data interface). For each SM2R1 monitoring value requested the following information is written to this transfer file:

- mean value for the overall analysis period
- maximum value for the overall analysis period
- minimum value for the overall analysis period
- standard deviation for the overall analysis period
- number of subcycles in the analysis period
- mean values of the individual subcycles.

Structure of the transfer file:

- The first record of the transfer file contains information pertaining to the analysis period and the duration of the subcycle. This record is given the record identifier `TIM2`.
- The next record of the transfer file contains general system information pertaining to the SM2 monitoring environment. This record is given the record identifier `SYST`.
- Afterwards, the configuration valid for the analysis period is output. The configuration requires several records. These records are given the record identifier `CONF`.
- This is followed by “pairs of records” each of which provides a description (first part) and monitoring data (second part) of an SM2R1 monitored variable. The records which contain the description are given the identifier `DSCR`. The records which contain the monitoring information are given the identifier `DATA`. In addition, these record pairs supply the report number and the monitored variable number for the respective SM2R1 monitored variable (for a description of the records see [section “SM2R1 transfer file records” on page 499](#)).

Interruption-free clock resetting

When time-series charts are created, “double” time stamps that are the result of the local clock being set back are left out.

In task, queue and HSMS statistics, the ELAPSED TIME of the titles is determined from the UTC time.

The records of the transfer file are not affected by the interruption-free clock resetting, since the “double” time stamps are already left out when the time-series charts are created.

Automatic performance analysis

SM2R1 detects bottlenecks in the system by means of the automatic performance analysis facility and reports them to the user via appropriate messages. This considerably reduces the need for large quantities of paper, as is normally required for diagnosing a performance problem: after the resource causing the bottleneck has been identified by SM2R1 during the automatic analysis, the manual file analysis can restrict itself to the data relevant to this resource.

The automatic analysis can be started on its own or in conjunction with other SM2R1 analyses.

The specification of one or more analysis periods is also valid for the performance analysis, i.e. only the selected periods are checked for bottlenecks.

Alongside the records that are present by default, automatic analysis also analyzes the records from the monitoring programs SERVICETIME, SYSTEM, TASK and, if present, VM.

Automatic analysis is activated by means of the START-AUTOMATIC-ANALYSIS statement ([page 485](#)).



When Live Migration takes place, a new session with the new server is started in the SM2 output file. Automatic performance analysis across the session boundary leads to separate analyses for the server concerned.

Table of contents for SM2R1

SM2R1 outputs a table of contents at the end of the entire analysis. This table begins on a new page under the header TABLE OF CONTENTS and now comprises four levels. The individual levels can be identified by the depth of indentation.

Level 1

specifies the individual runs of the SM2 output file (a normal case includes only one run). The following lines are output:

```
FIRST EVALUATION
SECOND EVALUATION
```

etc.

The SM2R1 title page is output to indicate a run. The following header also belongs to level 1:

```
EVALUATION STATISTICS ...
```

It supplies information about the SM2 output file.

Level 2

refers to general data, the configuration and subanalyses.

Possible entries for this level:

```
1. DATA FOR FIRST SESSION      ...
2. DATA FOR LAST SESSION      ...
3. CONFIGURATION-TABLE        ...
4. QUEUE STATISTICS           ...
5. AUTOMATIC ANALYSIS         ...
6. TIME-SERIES EVALUATION     ...
7. STATISTICS EVALUATION      ...
8. TASK STATISTICS            ...
9. HSMS STATISTICS            ...
10. SUMMARY STATISTICS        ...
```

The first five entries for level 2 do not include any subelements of level 3.

Levels 3 and 4

provide more detailed subanalyses of the entries (headers) output in level 2.

Example of a table of contents

T A B L E O F C O N T E N T S

FIRST EVALUATION		1
DATA FOR FIRST SESSION		1
CONFIGURATION TABLE		2
SUMMARY STATISTICS		25
ACTIVITY REPORT		25
PCS REPORT		28
QUEUE STATISTICS		31
TASK STATISTICS		40
CPU-TIME AND #IOS SORTED, FIRST 20 TASKS		40
CLASS-5 + CLASS-6-PAGES-SORTED, FIRST 20 TASKS		41
CATEGORIES		42
TIME-SERIES EVALUATION		50
*CPU		50
UTILIZATION NORMED	(REPORT 1)	50
SUM SVC CALLS	(REPORT 137)	51
*IO		52
IO'S FOR DEVICE CLASSES	(REPORT 3)	52
*DISK		53
UTILIZATION	(REPORT 124)	53
IO'S	(REPORT 125)	54
PAM IO'S	(REPORT 127)	55
*SERVICETIME		56
DURATION OF IO'S FOR DEVICE	(REPORT 231)	56
*CHANNEL		57
UTILIZATION	(REPORT 10)	57
AUTOMATIC ANALYSIS		59
EVALUATION STATISTICS		60

9.1 Starting and terminating SM2R1

The SM2R1 analysis program is called by means of the BS2000 command /START-SM2R1. The PL/1 text file and the SM2R1 control file are assigned automatically via link names.

Format

START-SM2R1

```

VERSION = *STD / <product-versionmandatory-man-corr> / <product-version mandatory-man-without-corr> /
          <product-version without-man-corr>
,MONJV = *NONE / <filename 1..54 without-gen-vers>
,CPU-LIMIT = *JOB-REST / <integer 1..32767>
,MONITOR-FILE-NAME = *BY-LINK-NAME / <filename 1..54 without-gen-vers>
,LIST-FILE-NAME = *BY-LINK-NAME / <filename 1..54 without-gen-vers>
,EVALUATION-FILE-NAME = *BY-LINK-NAME / <filename 1..54 without-gen-vers>

```

Operands

VERSION =

Specifies the program version of SM2R1 to be called.

VERSION = *STD

The current program version is called.

VERSION = <product-version>

The specified program version is called.

MONJV =

Specifies the name of the job variable to monitor the SM2R1 run. The job variable must already be cataloged.

MONJV = *NONE

No job variable is specified.

MONJV = <filename 1..54 without-gen-vers>

Specifies the name of an already cataloged job variable.

CPU-LIMIT =

Specifies the CPU time SM2R1 can use during execution. If this time is exceeded in interactive mode, the user is notified by the system; in batch mode, the SM2R1 run is terminated.

CPU-LIMIT = *JOB-REST

There is no time limit for the program.

CPU-LIMIT = <integer 1..32767>

Specifies the CPU time SM2R1 can use during execution.

MONITOR-FILE-NAME = *BY-LINK-NAME

The SM2 output file to be analyzed was assigned by means of the MONDTA file link name.

MONITOR-FILE-NAME = <filename 1..54 without-gen-vers>

Specifies the name of the SM2 output file to be analyzed which is assigned by means of the MONDTA file link name.

LIST-FILE-NAME = *BY-LINK-NAME

The SM2R1 output file was assigned by means of the DIALST file link name.

LIST-FILE-NAME = <filename 1..54 without-gen-vers>

Specifies the name of the SM2R1 output file which is assigned by means of the DIALST file link name.

EVALUATION-FILE-NAME = *BY-LINK-NAME

The SM2R1 transfer file was assigned by means of the EVALDTA file link name.

EVALUATION-FILE-NAME = <filename 1..54 without-gen-vers>

Specifies the name of the SM2R1 transfer file which is assigned by means of the EVALDTA file link name.

The SM2R1 statement `END` initiates processing, terminating SM2R1 after processing of the statement is completed.

Preparing the program call

Before SM2R1 is called, all the input files have to be specified and the output file which is to incorporate the reports and statistics generated has to be defined. A transfer file for the monitoring data which can be requested as part of a PRINT-REPORTS statement must likewise be assigned prior to the program run.

The input and output files are specified either by means of the corresponding operands of the START-SM2R1 command or by means of file link names assigned with the ADD-FILE-LINK command.

If the output files were not created with CREATE-FILE, they are created implicitly by SM2R1.

The input files are

- The SYSDTA system file containing the user's input statements. SYSDTA can also be assigned to a SAM or ISAM file in addition to SYSCMD.
- The SAM output file to be analyzed. The link name to be used is MONDTA.
- The SYSDAT.SM2.<ver>.MTFILE control file defining the format of the records in the SM2 output file. The link name to be used is MTFILE (assigned automatically by the START-SM2R1 command).
- The PL/1 text file which includes the message texts of the PL/1 runtime system. The link name to be used is TEXTLINK (assigned automatically by the START-SM2R1 command).

The last two files are supplied together with SM2R1.

- The user can choose any name for the output file. Its link name is DIALST.
- The name for the SM2R1 transfer file is also freely selectable. The link name to be used is EVALDTA.
BUFFER-LENGTH=*STD(SIZE=16) is assigned to the transfer file by default. The user's own assignments are not changed, but can lead to the program being aborted since the DATA record cannot exceed 32032 bytes in length.

The output file or transfer file need be defined only if it is to be accessed at a later point in time.

Sample command sequence

```

/ASSIGN-SYSDTA TO=*SYSCMD / <statement file>
/CREATE-FILE FILE-NAME=<outputfile>,SUPPORT=...(SPACE=...)
/CREATE-FILE FILE-NAME=<transferfile>
/START-SM2R1 MONITOR-FILE-NAME=<SAM file of monitored data>,
              LIST-FILE-NAME=<outputfile>,
              EVALUATION-FILE-NAME=<transferfile>

...          (SYSDTA statement sequences)

//END

```



If /MODIFY-JOB-SWITCHES ON=1 is entered prior to SM2R1 execution, SM2R1 will issue the message *RUNOPT OR *END EXPECTED from the PL/1 runtime system. After input of the PL/1 control statements (*RUNOPT) and (*END), SM2R1 statements can be entered again.

Program execution

After the call, SM2R1 displays its version number and the version number of the MTFILE control file and expects statement input. These statements determine:

- the period of time for which the SM2 output file is to be analyzed (analysis period)
- the degree of detail in which the analysis is to be executed
- which variables are to be analyzed
- in what form the analyzed data is to be output.

Analysis itself does not begin until the user has sent off his/her statement sequence with an END statement. No more statements can be input after this.

Program, warning and error messages are always output on SYSOUT. SM2R1 also lists all entered statements on SYSOUT if it is running as a batch task or if the input medium is a file.

Program termination / printing the output file

R1 terminates automatically at the end of the analysis process. If output of monitoring data to the output file is requested, the file can be printed with the following command:

```

/PRINT-DOCUMENT FROM-FILE=outputfile,
                DOCUMENT-FORMAT=*TEXT(LINE-SPACING=*BY-EBCDIC-CONTROL)

```


9.2 Monitoring environment output

Following the header page, SM2R1 outputs monitoring environment information. This data is subdivided into static system information and monitoring information.

Static system information is information that does not change during a BS2000 session. It facilitates identification of the monitored hardware and software and provides relevant data for this purpose.

The monitoring information facilitates identification of the monitoring and analysis routines used. It also comprises the monitoring program operand values at the beginning of SM2 monitoring.

The static system information and the monitoring information appear a maximum of two times: once with the header DATA FOR FIRST SESSION, and another time with the header DATA FOR LAST SESSION. The first appearance refers to the first BS2000 session in the analysis period and the second to the last BS2000 session in the analysis period. DATA FOR LAST SESSION is output only if the data has changed during the analysis period.

The system name is output in the monitoring environment, in all reports and in the summaries. When the name is changed (e.g. by mixing files of different versions), the last name is output in each case.

The data is described in detail in the following:

Static system information

- | | |
|--------------------------|--|
| a) SYSTEM IDENTIFICATION | Data identifying the monitored BS2000 system |
| SYSTEM NAME | Name of the BS2000 system |
| BS2000 VERSION | 3-digit version number of the operating system |
| GENERATION DATE | Date in ISO format on which the BS2000 system was generated |
| b) SYSTEM INFORMATION | Information on the generated system |
| MAX USER ADDRESS SPACE | Maximum user address space in Kb |
| CLASS 1 MEMORY | Size of virtual class 1 memory in Kb |
| CLASS 2 MEMORY | Size of virtual class 2 memory in Kb |
| # LOGICAL MACHINES | Number of logical machines |
| # ACTIVE LM'S | Number of active logical machines (only output for files if not identical to # LOGICAL MACHINES) |

c) HARDWARE IDENTIFICATION	Data identifying the CPU
MACHINE TYPE	CPU type
CPU 0 ID	Serial number of the first or only CPU
d) HARDWARE INFORMATION	CPU-related data (only the first CPU ID is output)
MAIN MEMORY	Main memory size in Kb (peripheral configuration information is output separately)
e) STARTUP IDENTIFICATION	Information identifying the startup operation
VOLUME # IPL DEVICE	Data medium from which the operating system was loaded
SESSION #	Number of the monitored BS2000 session
f) STARTUP INFORMATION	Information on the startup time
DATE OF STARTUP	Startup date
TIME OF STARTUP	Startup time

Monitoring information

a) MONITOR IDENTIFICATION	Data identifying the monitoring program and the analysis routine
SM2.OML VERSION	SM2 module library version number
SM2R1 VERSION	Version number of the SM2R1 analysis routine
MTFILE VERSION	Version number of the SYSDAT.SM2.<ver>.MTFILE command file
b) MONITOR INFORMATION	Operand settings for the monitoring program at the time the SM2 output file was opened
OFFLINE PERIOD	Monitoring cycle in seconds
SAMPLING PERIOD	Sampling cycle in milliseconds
NAME OF SM2-FILE	Name of the SM2 output file

Information on the analysis run

The second-last (numbered) sheet of the SM2R1 output contains the following information:

NO. OF SM2-SESSIONS	Number of monitored sessions
NO. OF SM2-INTERVALS	Number of SM2 monitoring cycles within the analysis period
NO. OF MISSED RECORDS	Number of SM2 records excluded from monitoring
TOTAL NO. OF SM2-RECORDS	Total number of SM2 records read
DAY AND TIME OF EVALUATION	Date and time at end of analysis
CPU-TIME OF EVALUATION	CPU time used for analysis in seconds

The number of records processed for each SM2 record is then output under the header SM2 RECORDS PROCESSED.

9.3 Structure of time series and statistical data

9.3.1 Computation methods

The analysis period specifies the period for which the SM2 output file is to be analyzed. Individual sections of this time frame are called analysis subintervals. When time series are output, an analysis subinterval is represented as a bar.

Computation of the value of an analysis subinterval

The monitoring cycle records are used to compute the monitored variables for one subinterval.

SM2R1 uses only those records for each requested variable that contain the required monitoring data.

From this monitoring data SM2R1 computes the value of an analysis subinterval using the following formula:

$$\bar{x} = \frac{\sum_{i=1}^n x_i g_i}{\sum_{i=1}^n g_i}$$

where:

- \bar{x} mean value for an analysis subinterval
- x_i single monitored value
- g_i weighting
- n number of monitored values

Depending on the meaning of the individual monitored data, a weighting on the basis of the length of the monitoring cycle or a different monitoring variable is performed. For example, the monitored data for the duration of an input/output may be weighted using the number of inputs/outputs.

The computed value is entered in the **bar chart** as the value for one analysis subinterval.

SM2R1 processes the data of those monitoring cycles whose time stamp record falls within the relevant analysis subinterval.

Computation of statistical data

For each monitored variable, SM2R1 can output a table with statistical data instead of a bar chart. The data covers the whole analysis period. If it is output in addition to a bar chart, it refers to one bar chart page.

The analysis subinterval data computed beforehand is used for the computation.

Average

The average value of the analysis period is calculated in a way similar to the values for the analysis subintervals.

Minimum

Minimum of the averages of the various analysis subintervals:

$$\min(\bar{x})$$

where: \bar{x} average for an individual analysis subinterval (see above)

Maximum

Maximum of the averages of the various analysis subintervals:

$$\max(\bar{x})$$

where: \bar{x} average for an individual analysis subinterval

Standard deviation

$$S = \sqrt{\frac{\sum_{i=1}^n (x_i - \bar{x})^2}{n - 1}}$$

where: \bar{x} average value of the analysis period
 x_i value of an analysis subinterval
 n number of analysis subintervals



The values for minimum and maximum depend greatly on the length of the analysis subintervals. The longer these are, the greater the extent to which extreme short-term values are compensated for: this makes the maximum values smaller and the minimum values larger.

Minimum and maximum values from different measurements can therefore only be compared with each other when the analysis subintervals are the same length.

The size of the analysis subintervals can be set with the SM2R1 statement `SET-EVALUATION-PERIOD EVALUATION-PERIOD=*PERIOD(TIME-STEPPS=...)`.

9.3.2 Bar chart layout

Each report has a header containing a name identifying the report, and the associated report group. SM2R1 subdivides the analysis period into small time slots of equal length called analysis subintervals.

The user can control this subdivision by specifying either the number or the length of the analysis subintervals. He can also direct SM2R1 to perform the subdivision. In this case the analysis period is subdivided into as many as 100 subintervals.

Bar chart

Each analysis subinterval is represented in the bar chart by a bar over the horizontal axis, which itself is divided up into corresponding segments. The vertical axis is the monitoring data scale.

The analysis subinterval is the time represented by one bar.




Figure 8: Schematic representation of a bar chart

In the bar chart, SM2R1 outputs the values computed for each analysis subinterval. The computation method is described in the section [“Computation of the value of an analysis subinterval”](#) on page 404.

Division and annotation of the axes

The horizontal axis (abscissa) is the time axis. It is subdivided in analysis subinterval slots and annotated accordingly.

If an analysis subinterval consists of a number of monitoring cycles, the time of the first time stamp record of the analysis subinterval is output as the annotation.

The vertical axis (ordinate) is the monitoring data scale. It is subdivided and annotated appropriately by the analysis routine.

The monitoring data is entered in the bar chart over the horizontal axis in bar form. For each variable, a suitable unique symbol (letter or other character) is used (see [section "Table of variables reports" on page 623](#)). Each symbol is briefly explained in the bar chart legend.

If the ACCUMULATED variables are entered, the totals for the variables are always rounded. This means that the following inaccuracies can arise with very small values:

- If, for instance, the rounded value for the first variable is exactly the same as the rounded value for the first two variables, no symbol is output for the second variable.
- If the rounded total is greater than the rounded value for the first variable, a symbol is output.

It is nevertheless possible that the value for the second variable is greater in the first of these two cases than in the second.

Entry of variables: ACCUMULATED and NON ACCUMULATED

The specification ACCUMULATED or NON-ACCUMULATED in the header line of the bar chart indicates how the variables are entered in a chart.

ACCUMULATED means that the variables are sequenced, i.e. the bar for one variable begins at the point where the bar for the preceding one ends.

NON ACCUMULATED means that each variable is entered beginning at zero, i.e. the symbols for a small variable overlay the symbols in the lower part of a large variable. If two variables in a chart are the same size, one of them will supersede the other.

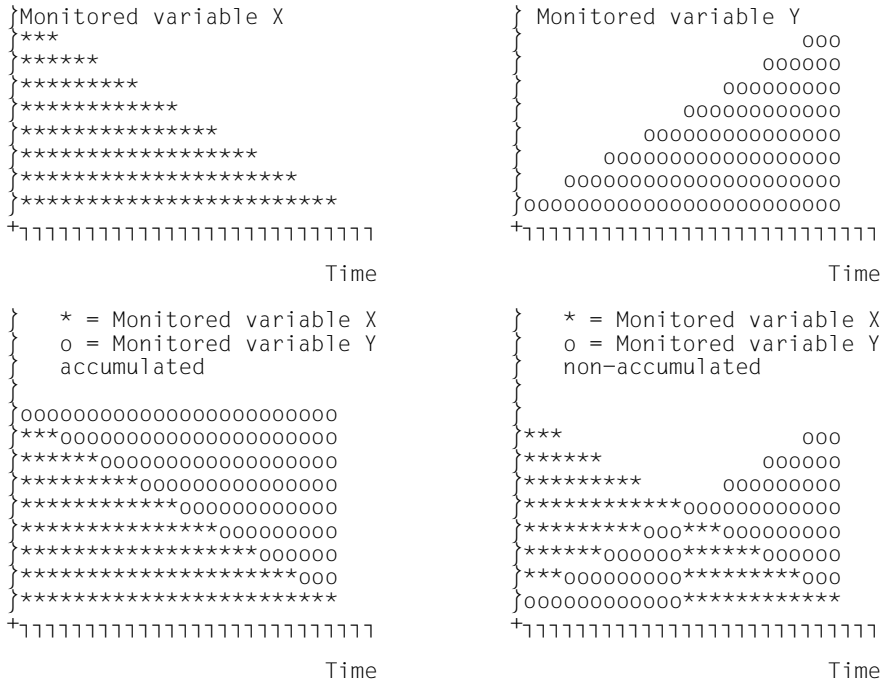


Figure 9: ACCUMULATED/NON ACCUMULATED representation of variables

Representation of averages

In charts using non-accumulated representation, the average of the monitored variables is represented by the appropriate symbol on the right-hand side.

The symbol < represents ' '. As many print columns are used on the right-hand side as are required for unequivocal representation if the symbols overlap.

Variables for which no values can be determined

It is possible that SM2 cannot determine the value of a certain variable in an interval, e.g. because

- the variable was not monitored by SM2 or its value was not entered in the SM2 output file for later analysis, or
- the variable was monitored but no activity occurred in the interval.

Variables for which no values could be determined in an analysis subinterval are not used for averaging. Blanks appear in the graphical representation instead of a bar.

Subcharts

If the user has specified the size or the number of analysis subintervals with the TIME-STEPS operand of the SET-EVALUATION-PERIOD statement so that more than 100 subintervals (time slots in the bar chart) are required for each analysis period, subcharts have to be output. This is because only 100 time slots can be accommodated on the horizontal axis in the printer listing. SM2R1 does this automatically.

Requesting bar chart output

The user can request bar chart output with the INFORMATION=*DIAGRAMS operand of the PRINT-REPORTS statement.

9.3.3 Representation of the statistical data

In these tables, SM2R1 displays the following values for the entire analysis period:

- average (AVG)
- minimum (MIN)
- maximum (MAX)
- standard deviation (SDEV)
- number of analysis subintervals (CNT).

SM2R1 computes these values from the data of all analysis subintervals of an analysis period (from one chart page if the operand INFORMATION=*DIAGRAMS was specified).

The computation methods are described in section [“Computation of statistical data” on page 405](#).

Table format

SM2R1 outputs a separate line for every variable contained in the requested variables report. The statistical data is entered in columns. The table header lines contain the page number, the variables report number, analysis period information (date/time TO date/time) and analysis subinterval information (TIME-STEPS). The analysis period information refers to the first or last time stamp record. Statistical data output can be requested with the operand INFORMATION=*STATISTICS of the PRINT-REPORTS statement. In the case of analysis time periods which are either large or small, the number of analysis subintervals may exceed the maximum value displayed, in which case the value 9999 will be entered in column CNT.

Example statistical data table

```

***** SM2R1 - EXAMPLE *****

*CPU: UTILIZATION NORMED (REPORT 1)
ALL PROCESSORS
ACCUMULATED                EVALUATION FROM 08-09-08, 15:09:56: TO 16:27:31 BY 0:00:47

ID - MEASURED TERM , UNIT IS PERCENT ----- AVG ---- MIN ---- MAX ---- SDEV -- CNT
1 = TU TIME                13.553    1.479    37.548    11.391    23
. = TPR TIME                11.856    2.110    55.603    11.078    23
3 = SIH TIME                3.778     0.740     7.513     1.919    23
  = IDLE TIME              70.813    25.163    95.343    19.748    23
- = STOP TIME                0         0         0         0         23
    
```

9.3.4 Variables reports

SM2R1 analyzes the variables contained in the SM2 output file individually. For output, SM2R1 combines logically related variables in groups.

These groups are called variables reports.

Subreports (referring to various monitored objects) for a monitoring data report are always sorted in alphabetical order, except for reports whose subreports are category-specific. These reports begin with the SM2 pseudo-category SUM followed by the standard categories SYS, DIALOG, BATCH and TP, and finally the remaining categories in alphabetical order.

9.3.5 Additional outputs

The requesting of separate reports by means of PRINT-REPORTS, together with dynamic I/O configuration change, result in the output of the following tables.

- **LEGEND-LIST OF RESPONSE-TIME REPORTS**
For the report group RESPONSE-TIME, the LEGEND-LIST OF RESPONSE-TIME REPORTS is output on a separate printed page following the last report. This list contains all the allocated names together with the associated connection groups and/or connection group sets.
- **DEVICE-LIST OF EXTENDED SYSTEM STATISTIC REPORTS**
For the report groups CATEGORY-CPU and CATEGORY-IO, the DEVICE-LIST OF EXTENDED SYSTEM STATISTIC REPORTS is output on a separate page following the last report. This latter report provides a list of mnemonic device names to which the output, category-specific values of the monitoring program SYSTEMS refer. Alongside the mnemonic device names, the time of the device list is also output.
- **LEGEND-LIST OF BCAM-CONNECTION REPORTS**
For the report group BCAM-CONNECTION, the LEGEND-LIST OF BCAM-CONNECTION REPORTS is output on a separate sheet following the last report. This list contains all the connection sets together with their definitions.
- **LIST OF ADDED/REMOVED I/O UNITS**
In the event of a dynamic I/O configuration change, the tables are output either directly after the system configuration or after the reports. The date, time, mnemonic device name, channel number and device type are output. In the case of devices, the channel number is also output (as specified in the PREFERRED-PATH operand of the ADD-IO-UNIT command).

9.4 General structure of the statements

All statements issued to SM2R1 can be divided into three categories:

- SET/MODIFY statements

Statement	Function
SET-TITLE	Output title
SET-EVALUATION-PERIOD	Define analysis period and analysis subinterval
SET-EXCEPTION-PERIOD	Fade out time periods
SET-REPORT-FOCUS	Select time slots from analysis period
MODIFY-REPORT-CONDITIONS	Modify preset range for the monitoring data of a subreport

Each of the SET/MODIFY statements sets a condition which is valid for subsequent analyses requested with PRINT statements. A SET statement remains in effect until it has been replaced by a new SET statement. The MODIFY statement applies to one program session. If a number of MODIFY statements are issued for the same variable, the settings specified in the last MODIFY statement are used for the current program session.

For each statement which modifies the analysis period a new SM2R1 header sheet is output, indicating the defined periods. Merely the basic evaluation period is indicated in the separate reports.

- CREATE statements

Statement	Function
CREATE-TSN-SET	Declare SET names for TSNs
CREATE-USERID-SET	Declare SET names for user IDs
CREATE-JOBNAME-SET	Declare SET names for job names
CREATE-JOBCLASS-SET	Declare SET names for job classes

Each of the CREATE statements defines a certain set of connections, TSNs, user IDs, job names or job classes under a freely selectable name; afterwards an analysis requested via a PRINT statement is performed for them. Each further CREATE statement defines a new set. Sets already defined are kept until the end of the statement sequence, as long as the maximum number of definitions has not been exceeded.

- PRINT/START statements

Statement	Function
PRINT-CONFIGURATION	Output system configuration
PRINT-REPORTS	Output report groups
PRINT-SUMMARY	Output summary analysis
PRINT-HSMS-STATISTICS	Output HSMS analysis
PRINT-TASK-STATISTICS	Output task analysis
PRINT-QUEUE-TRANSITION	Output statistics on task queues
START-AUTOMATIC-ANALYSIS	Start automatic bottleneck analysis

A specific analysis in accordance with the conditions defined by the SET statements is not started until one of the PRINT statements is issued.



The internal program name for the syntax checking of SM2R1 statements in EDT is SM2R1-190.

9.5 Statements

The following statements are described and listed in alphabetical order.

CREATE-JOBCLASS-SET **Define set name for job classes**

The user employs this statement to merge tasks to form a group of job classes. This group is assigned a set name. Using this set name, the user addresses this set of tasks in the PRINT-TASK-STATISTICS statement, then obtaining the analysis report. A total of up to eight different job class sets may be defined here; in other words, the CREATE-JOBCLASS-SET statement must not be specified more than 8 times.

Format

CREATE-JOBCLASS-SET
SET-NAME = <alphanum-name 1..16> ,JOBCLASS = list-poss(32): <alphanum-name 1..8>

Operands

SET-NAME = <alphanum-name 1..16>

Specifies the name for a set of tasks belonging to one group.

JOBCLASS = list-poss(32): <alphanum-name 1..8>

Specifies the set of tasks by defining job classes by groups.

CREATE-JOBNAME-SET

Specify set name for job names

The user employs this statement to merge tasks belonging to a group of job names. A set name is assigned to this group. Using the set name, the user addresses this task set with the aid of the PRINT-TASK-STATISTICS statement, obtaining an analysis. A total of up to eight different sets of job names can be defined; in other words, the CREATE-JOBNAME-SET statement must not be specified more than 8 times.

Format

CREATE-JOBNAME-SET

```
SET-NAME = <alphanum-name 1..16>  
,JOBNAME = list-poss(32): <alphanum-name 1..8>
```

Operands

SET-NAME = <alphanum-name 1..16>

Defines a name for a set of tasks belonging to a group.

JOBNAME = list-poss(32): <alphanum-name 1..8>

Specifies the set of tasks by the names of jobs belonging to a group.

CREATE-TSN-SET

Define set name for specific TSNs

The user employs this statement to merge tasks belonging to a group of task sequence numbers (TSNs). A set name is assigned to this group. Using the set name, the user specifies this task set in the PRINT-TASK-STATISTICS statement, obtaining an analysis. A total of up to eight different sets of TSNs can be defined; in other words, the CREATE-JOBNAME-SET statement must not be specified more than 8 times.

Format

CREATE-TSN-SET

```
SET-NAME = <alphanum-name 1..16>  
,TSN = list-poss(32): <alphanum-name 1..4>
```

Operands

SET-NAME = <alphanum-name 1..16>

Defines the name for a set of tasks belonging to a class.

TSN = list-poss(32): <alphanum-name 1..4>

Defines the set of tasks by task sequence number (TSN), forming a class.

CREATE-USERID-SET

Define set name for user IDs

The user employs this statement to merge tasks belonging to a group of user IDs. This group is given a set name. Using the set name, the user specifies this task set in the PRINT-TASK-STATISTICS statement, obtaining an analysis. A total of up to eight different sets of user IDs can be defined; in other words, the CREATE-USERID-SET statement must not be specified more than 8 times.

Format

CREATE-USERID-SET
SET-NAME = <alphanum-name 1..16> ,USER-ID = list-poss(32): <alphanum-name 1..8>

Operands

SET-NAME = <alphanum-name 1..16>

Defines the name for a set of tasks belonging to one class.

USER-ID = list-poss(32): <alphanum-name 1..8>

Defines the set of tasks belonging to one class by user ID. User IDs must be entered without \$.

END

Terminate statements

All of the statement sequences issued to SM2R1 must be terminated by an END statement. Not until then does SM2R1 commence analysis operations. The user is then no longer able to intervene in the program session.

If a procedure contains a syntax error, the statements which have been read correctly up to that point are processed.

Format

END

MODIFY-REPORT-CONDITIONS

Modify preset range for monitored data of subreport

This statement allows the user to modify the preset range for the monitored data of a subreport. The modifications are valid for a program session. It is also possible to modify threshold values to reduce the output volume (see section [“Reducing the number of monitored objects to be output” on page 446](#)).

Format

MODIFY-REPORT-CONDITIONS

```
REPORT-NUMBER = <keyword-number>
,ITEM-NUMBER = <integer 1..5>
,LOWER-LIMIT = *UNCHANGED / <fixed 0..2147483647>
,UPPER-LIMIT = *UNCHANGED / <fixed 0..2147483647>
,SUPPRESS-CONDITION = *UNCHANGED / *INSIDE-RANGE / *OUTSIDE-RANGE
,COMPARE-VALUE = *UNCHANGED / *MEAN-VALUE / *MINIMUM / *MAXIMUM / *STANDARD-DEVIATION
```

Operands

REPORT-NUMBER = <keyword-number>

Specifies the report number.

ITEM-NUMBER = <integer 1..5>

Specifies the item number. The item number is determined by sequentially numbering the items for each report number in the table of variables reports (see [page 623](#)). If a non-existent item number is specified, an error message is issued containing the report number and the item number.

LOWER-LIMIT =

Specifies the lower limit of the range.

LOWER-LIMIT = *UNCHANGED

The lower limit of the range is unchanged.

LOWER-LIMIT = <fixed 0..2147483647>

The lower limit of the range is reset.

UPPER-LIMIT =

Specifies the upper limit of the range.

UPPER-LIMIT = *UNCHANGED

The upper limit of the range is unchanged.

UPPER-LIMIT = <fixed 0..2147483647>

The upper limit of the range is reset.

SUPPRESS-CONDITION =

Specifies when a subreport is to be suppressed.

SUPPRESS-CONDITION = *UNCHANGED

The setting is unchanged.

SUPPRESS-CONDITION = *INSIDE-RANGE

The subreport is suppressed if the monitoring data lies within the range (including the limit values).

SUPPRESS-CONDITION = *OUTSIDE-RANGE

The subreport is suppressed if the monitoring data is outside the range (excluding the limit values).

COMPARE-VALUE =

Specifies the value to be compared with LOWER-LIMIT and UPPER-LIMIT.

COMPARE-VALUE = *UNCHANGED

The setting is unchanged.

COMPARE-VALUE = *MEAN-VALUE

The mean value of the monitoring data is used for comparison.

COMPARE-VALUE = *MINIMUM

The minimum value in the monitoring data is used for comparison.

COMPARE-VALUE = *MAXIMUM

The maximum value in the monitoring data is used for comparison.

COMPARE-VALUE = *STANDARD-DEVIATION

Standard deviation is used for comparison.



The MODIFY-REPORT-CONDITIONS statement can only be issued once for any given report and item number within an SM2R1 session. The last statement issued applies.

PRINT-CONFIGURATION

Output system configuration

The user requests output of the system configuration using the PRINT-CONFIGURATION statement. The print appears after the first output of the monitoring environment. The first configuration which SM2R1 finds in the form of SM2 records in the first session to be analyzed is printed. This is an output of the configuration statistics. Modifications relating to dynamic I/O configuration change are output in a separate table (see [section “Additional outputs” on page 411](#)).

Format

PRINT-CONFIGURATION

The configuration is output in hierarchical form. The hierarchical levels are formed by the channels (CHANNEL-PATH-ID), device controllers (CONTROLLER) and attached devices (ATTACHED DEVICES).

The following information is output in the headers for the individual levels:

1. CHANNEL-PATH-ID: nnnn channel-type
 where nnnn is the channel number in hexadecimal format and channel-type is the channel type (BUS CHANNEL, TYP FC, TYP S).
2. CONTROLLER: mn
 where mn is the name of the device controller specified at generation time.
3. In the case of devices directly attached to the channel and on the same controller level, the following header appears:
 DIRECT ATTACHED DEVICES

The header contains the same information as for the level which follows: This level further delineates the devices attached to a controller.

The header is then:
 ATTACHED DEVICES

The individual devices are then listed.



Here “attached” implies use within the configuration. See also the DETACHED entry.

The following entries appear under the header:

DEVICE MNEM

The mnemonic device name specified for generation.

VOLUME (VSN)

Volume header, if any.

PATH INFO

Provides the input/output path, i.e. channel (4 characters) and device address (2 characters).

DEVICE TYPE

External designation of the device.

INTERNAL CODE

System-internal device code.

4. The following eight entries represent possible device attributes or information. If a device had the corresponding attribute at the time of monitoring, this is indicated by * in some columns. Otherwise, the corresponding information is output.

DETACHED DEVICE

The device exists within the system configuration, but it is not available at the present time (e.g. /DET UNIT=(mm,mn,...)).

PAGING DEVICE

The device is one used for paging.

PRIVATE DEVICE

Exclusive device assignment for a task. In the case of disks and tapes, marking takes place only if a volume has been mounted.

SHARED PRIVATE

Identifies a private disk which can function as a shareable private disk in MSCF mode (SHARED PRIVATE DISK).

SYSTEM PRIVATE

Identifies a private disk which can be used by multiple tasks.

BLOCK FORMAT

Identifies a disk's block format (K2, NK2, NK4).

PAV

Identifies a basic device (B) or alias device (A); otherwise, spaces are output.

BASE MNEM

Specifies the mnemonic device name of the basic device if it is an alias device; otherwise, spaces are output.

PRINT-HSMS-STATISTICS

Print data on file migration and retrieval

This analysis provides a table indicating the migration of files to the background level and the retrieval of files to the processing level.

A temporal distribution of the Recall operations (in intervals of 2 minutes) is generated. The data is supplied both for the entire Recall operation and for the individual steps of the Recall operation.

Format

PRINT-HSMS-STATISTICS

Output data

GENERAL INFORMATION

# MIGRATION RUNS	Number of migration runs
# MIG FILES PER RUN	Number of files per migration run
# RECALL RUNS	Number of (explicit and implicit) Recall operations
# RCL FILES PER RUN	Number of files per Recall operation
MOUNTS	Number of mount operations concerning the background level media

MIGRATION OPERATIONS

# FILES	Number of migrated files
AVG SIZE	Average size of migrated files in 2 KB pages
AVG # EXTENTS	Average number of file extents

MIGRATED FILES

# FILES	Number of files migrated between the levels S0-S1, S0-S2, S1-S2
# DAYS	Average number of days per file between the last access and migration

RECALL OPERATIONS

# FILES	Number of files retrieved (explicitly and implicitly) from background levels S1 and S2
# DAYS	Average number of days per file on the background level for the Recall operation

DURATION OF RECALL REQUESTS

The Recall operations are distributed across different time slots according to their duration. In the case of recall operations lasting between 0 and 18 minutes, the time slots are each of two minutes. All operations lasting longer than 18 minutes are counted in the last time slot. Both the percentages and absolute numbers of Recall operations are output. Alongside the full Recall operation (SUM), monitoring values for 4 separate steps of the Recall operation are supplied.

Description of the steps:

SUM	Total Recall operation
HSM0120	From start of job to start of processing
ARC0000	From start of processing to ARCHIVE call
ARC0018	From ARCHIVE call to end of ARCHIVE
TAPE MOUNTED	Mount operation on the background level device

PRINT-QUEUE-TRANSITION

Print statistics on task queues

This analysis is initiated by the PRINT-QUEUE-TRANSITION statement. A chart of system-global output and category-specific output (designated E in the chart) as well as times spent in task queues is printed. In the case of duration times (D in the chart) a percentage indicating the duration of all tasks in this queue category relative to the duration of all tasks in this category of queues based on overall duration is output (i.e. the duration of all tasks of this category in all queues).

The first 30 (maximum) categories occurring in the SM2 output file are printed, including the SM2 pseudo-category SUM.

The categories are listed in the following order:

SUM, standard categories SYS, DIALOG, BATCH and TP as well as the remaining categories sorted in alphabetical order.

If more than 30 categories occur in an analysis period, the others are ignored.

The above values are output for all tasks in the system under the SUM entry.

In the case of queues used for various purposes a breakdown is provided (e.g. for Q4 and Q12).

Basically only those transitions are counted for which the queue has changed (e.g. from WS1 to WS2).

The prerequisite for this analysis is that the monitoring program SYSTEM has been activated during the SM2 monitoring run.

Duration time in the queue for the CPU (Q1) does not include the time spent with the CPU itself (Q0).

Output from the queue for the CPU includes output operations from the CPU (also to Q1), but not transitions from Q1 to Q0 (initiations).

The SM2 START-MEASUREMENT-PROGRAM and STOP-MEASUREMENT-PROGRAM statements are used to activate and deactivate analysis. If more than one pair of START/STOP-MEASUREMENT-PROGRAM statements are issued during an analysis period, multiple output reports are issued.

If the analysis period is selected to place the beginning of analysis time between the START- and STOP-MEASUREMENT-PROGRAM statements, the start time used is the time of the next time stamp record.

If no STOP-MP record has been located by the time the end of the analysis period is reached, the time of the next time stamp record is used as the end time.

Format

PRINT-QUEUE-TRANSITION

Output data

PRINT-QUEUE-TRANSITION provides statistics on the task queues for the specified analysis period.

E =	Category-specific output from task queues.
D =	Duration times per category and queue, measured as a percentage of the total time all tasks in this category spend in all queues.
SUM	Total values of all tasks in the system.
SYS	SYSTEM category.
DIALOG	DIALOG category.
BATCH	BATCH category.
TP	TP category.
xx	Further categories sorted in alphabetical order (where xx is the name of the category).

The extreme left-hand column lists the queues used for the different purposes. A breakdown, for example for Q4 and Q12, is provided. The CPU time is the same as the value for queue Q0.

PRINT-REPORTS

Print report groups

The user employs the PRINT-REPORTS statements to select the report groups to be analyzed. The INFORMATION operand is used to define the type of output.

The user employs the group names in the REPORT-LIST operand to address report groups logically related to one another. Which variable reports belong to which report group is depicted in [table "Report groups" on page 611](#).

The REPORT-NUMBER operand allows selection of individual reports in a report group using their report numbers. The default value REPORT-NUMBER=*STD means that the most important reports of a report group are analyzed. The relevant report numbers can be found in the [table "Report groups" on page 611](#).

The CONDITIONED-REPORTS operand specifies whether or not subreports whose monitoring data is within a certain range are to be output for a given report group.

For an overview of the reports and a detailed list of the individual variables, see [table "Variable reports" on page 623](#).

Format

PRINT-REPORTS

```

INFORMATION = *DIAGRAMS / *STATISTICS / *INTERFACE
,REPORT-LIST = *STD / *ALL(...) / list-poss(48): *BCAM-CONNECTION(...) / *BCAM-MEMORY(...) /
  *CATALOG-MANAGEMENT(...) / *CATEGORY-CPU(...) / CATEGORY-IO(...) /
  *CATEGORY-QUEUE(...) / *CATEGORY-WORKING-SET(...) / *CHANNEL(...) / *CPU(...) / *DAB(...) /
  *DEVICE(...) / *DILATION(...) / *DISK(...) / *DISK-FILE(...) / *DLM(...) / *FILE(...) / *GS(...) /
  *GSVOL(...) / *IO(...) / *ISAM(...) / *ISAM-FILE(...) / *MEMORY(...) / *MSCF(...) / *NSM(...) /
  *OPENFT(...) / *PCS(...) / *PERIODIC-TASK-JOBNAME(...) / *PERIODIC-TASK-TSN(...) /
  *PERIODIC-TASK-USERID(...) / *PFA(...) / *POSIX(...) / *PRIOR-ACF(...) / *PUBSET(...) /
  *RESPONSE-TIME(...) / *RST(...) / *SESAM-SQL(...) / *SERVICETIME / *STD(...) /
  *STORAGE-SYSTEM-SYMMETRIX(...) / *SVC(...) / *TASK(...) / *TCP-IP(...) / *TLM(...) /
  *UDS-SQL(...) / *UTM(...) / *VM2000(...) / *VOLUME(...) / *WORKING-SET(...)

```

(part 1 of 8)

```

*ALL(...)
  | REPORT-NUMBER = *ALL / *STD
  | ,CONDITIONED-REPORTS = *NO / *YES
*BCAM-CONNECTION(...)
  | CONNECTION-SET = *ALL / list-poss(32): <alphanum-name 1..16>
  | ,REPORT-NUMBER = *STD / *ALL / list-poss(17): *192 / *193 / *194 / *195 / *196 / *197 / *198 / *199 /
  |               *200 / *201 / *202 / *203 / *226 / *258 / *259 / *260 / *261
  | ,CONDITIONED-REPORTS = *NO / *YES
*BCAM-MEMORY(...)
  | REPORT-NUMBER = *STD / *ALL / list-poss(2): *280 / *281
  | ,CONDITIONED-REPORTS = *NO / *YES
*CATALOG-MANAGEMENT(...)
  | CATALOG-ID = *ALL / list-poss(32): <catid 1..4>
  | ,REPORT-NUMBER = *STD / *ALL /
  |               list-poss(10): *66 / *67 / *68 / *69 / *70 / *71 / *72 / *103 / *104 / *185
  | ,CONDITIONED-REPORTS = *NO / *YES
*CATEGORY-CPU(...)
  | CATEGORY = *ALL / list-poss(32): <alphanum-name 1..7>
  | ,REPORT-NUMBER = *STD / *ALL / *62
  | ,CONDITIONED-REPORTS = *NO / *YES
*CATEGORY-IO(...)
  | CATEGORY = *ALL / list-poss(32): <alphanum-name 1..7>
  | ,REPORT-NUMBER = *STD / *ALL / list-poss(3): *63 / *64 / *65
  | ,CONDITIONED-REPORTS = *NO / *YES
*CATEGORY-QUEUE(...)
  | CATEGORY = *ALL / list-poss(32): <alphanum-name 1..7>
  | ,REPORT-NUMBER = *STD / *ALL / list-poss(3): *28 / *30 / *31
  | ,CONDITIONED-REPORTS = *NO / *YES
*CATEGORY-WORKING-SET(...)
  | CATEGORY = *ALL / list-poss(32): <alphanum-name 1..7>
  | ,REPORT-NUMBER = *STD / *ALL / list-poss(2): *29 / *58
  | ,CONDITIONED-REPORTS = *NO / *YES

```

(part 2 of 8)

***CHANNEL(...)**

```

CHANNEL-PATH-ID = *ALL / list-poss(32): <x-string 1..4>
,REPORT-NUMBER = *STD / *ALL / list-poss(4): *10 / *101 / *102 / *257
,CONDITIONED-REPORTS = *YES / *NO

```

***CPU(...)**

```

PROCESSOR-SPLITTING = *NO / *YES
,REPORT-NUMBER = *STD / *ALL / list-poss(5): *1 / *2 / *6 / *137 / *204
,CONDITIONED-REPORTS = *NO / *YES

```

***DAB(...)**

```

CACHE-ID = *ALL / list-poss(32): <alphanum-name 1..32>
,REPORT-NUMBER = *STD / *ALL / list-poss(8): *79 / *80 / *81 / *82 / *189 / *190 / *205 / *206
,CONDITIONED-REPORTS = *NO / *YES

```

***DEVICE(...)**

```

DEVICE = *ALL / *SPECIFIED (...)
  *SPECIFIED (...)
    DEVICE = *NOT-SPECIFIED / list-poss(256): <alphanum-name 1..4>
    ,DEVICE-TYPE = *NOT-SPECIFIED / *TAPE / *TD / *OTHER
,REPORT-NUMBER = *STD / *ALL / list-poss(8): *11 / *35 / *36 / *100 / *230 / *282 / *283 / *319
,CONDITIONED-REPORTS = *YES / *NO

```

***DILATION(...)**

```

CATEGORY = *ALL / list-poss(32): <alphanum-name 1..7>
,REPORT-NUMBER = *STD / *ALL / *57
,CONDITIONED-REPORTS = *NO / *YES

```

(part 3 of 8)

```

*DISK(...)
  SPECIFIED = *ALL / *DEVICE(...) / *VOLUME-AND-DEVICE(...)
    *DEVICE(...)
      | DEVICE = list-poss(256): <alphanum-name 2..4>
    *VOLUME-AND-DEVICE(...)
      | DEVICE-VOLUME = list-poss(256): *SELECT(...)
        *SELECT(...)
          | DEVICE = <alphanum-name 2..4>
          | ,VOLUME = <vsn 1..6>
      ,REPORT-NUMBER = *STD / *ALL / list-poss(8): *124 / *125 / *126 / *127 / *227 / *228 / *229 / *270
      ,CONDITIONED-REPORTS = *YES / *NO

*DISK-FILE(...)
  | REPORT-NUMBER = *STD / *ALL / list-poss(1): *320
  | ,CONDITIONED-REPORTS = *NO / *YES

*DLM(...)
  | REPORT-NUMBER = *STD / *ALL / list-poss(4): *170 / *171 / *172 / *173
  | ,CONDITIONED-REPORTS = *NO / *YES

*FILE(...)
  | FILE-NAME = *ALL / list-poss(32): <filename 1..54>
  | ,REPORT-NUMBER = *STD / *ALL / list-poss(3): *13 / *14 / *191
  | ,CONDITIONED-REPORTS = *NO / *YES

*GS(...)
  | PARTITION = *ALL / list-poss(32): <alphanum-name 1..8>
  | ,REPORT-NUMBER = *STD / *ALL / list-poss(2): *177 / *178
  | ,CONDITIONED-REPORTS = *YES / *NO

*GSVOL(...)
  | DEVICE-MNEMONIC = *ALL / list-poss(32): <alphanum-name 1..4>
  | ,REPORT-NUMBER = *STD / *ALL / *174
  | ,CONDITIONED-REPORTS = *YES / *NO

```

(part 4 of 8)

```

*IO(...)
  PROCESSOR-SPLITTING = *NO / *YES
  ,REPORT-NUMBER = *STD / *ALL / list-poss(2): *3 / *4
  ,CONDITIONED-REPORTS = *NO / *YES

*ISAM(...)
  ISAM-POOL = *ALL / list-poss(32): <alphanum-name 1..8>
  ,REPORT-NUMBER = *STD / *ALL / list-poss(4): *85 / *86 / *87 / *224
  ,CONDITIONED-REPORTS = *NO / *YES

*ISAM-FILE(...)
  FILE-NAME = *ALL / list-poss(32): <alphanum-name 1..54>
  ,REPORT-NUMBER = *STD / *ALL / list-poss(4): *271 / *272 / *273 / *274
  ,CONDITIONED-REPORTS = *NO / *YES

*MEMORY(...)
  REPORT-NUMBER = *STD / *ALL / list-poss(10): *8 / *9 / *52 / *53 / *54 / *55 / *56 / *94 / *95 / *275
  ,CONDITIONED-REPORTS = *NO / *YES

*MSCF(...)
  REPORT-NUMBER = *STD / *ALL / list-poss(4): *166 / *167 / *168 / *169
  ,CONDITIONED-REPORTS = *YES / *NO

*NSM(...)
  REPORT-NUMBER = *STD / *ALL / list-poss(5): *179 / *180 / *181 / *182 / *184
  ,CONDITIONED-REPORTS = *NO / *YES

*OPENFT(...)
  REPORT-NUMBER = *STD / *ALL /
                    list-poss(10): *309 / *310 / *311 / *312 / *313 / *314 / *315 / *316 / *317 / *318
  ,CONDITIONED-REPORTS = *NO / *YES

*PCS(...)
  CATEGORY = *ALL / list-poss(32): <alphanum-name 1..7>
  ,REPORT-NUMBER = *STD / *ALL / list-poss(5): *73 / *74 / *75 / *76 / *77
  ,CONDITIONED-REPORTS = *NO / *YES

```

(part 5 of 8)

***PERIODIC-TASK-JOBNAME(...)**

```

  JOBNAME = *ALL / list-poss(32): <alphanum-name 1..8>
  ,REPORT-NUMBER = *STD / *ALL / list-poss(4): *160 / *161 / *162 / *163
  ,CONDITIONED-REPORTS = *YES / *NO

```

***PERIODIC-TASK-TSN(...)**

```

  TASK = *ALL / *SPECIFIED(...)
  *SPECIFIED(...)
    TSN = *NOT-SPECIFIED / list-poss(32): <alphanum-name 1..4>
    ,USER-ID = *NOT-SPECIFIED / list-poss(32): <alphanum-name 1..8>
  ,REPORT-NUMBER = *STD / *ALL / list-poss(4): *152 / *153 / *154 / *155
  ,CONDITIONED-REPORTS = *YES / *NO

```

***PERIODIC-TASK-USERID(...)**

```

  USERID = *ALL / list-poss(32): <alphanum-name 1..8>
  ,REPORT-NUMBER = *STD / *ALL / list-poss(4): *156 / *157 / *158 / *159
  ,CONDITIONED-REPORTS = *YES / *NO

```

***PFA(...)**

```

  CACHE-ID = *ALL / list-poss(32): <alphanum-name 1..4>
  ,REPORT-NUMBER = *STD / *ALL / list-poss(3): *134 / *135 / *136
  ,CONDITIONED-REPORTS = *NO / *YES

```

***POSIX(...)**

```

  REPORT-NUMBER = *STD / *ALL /
    list-poss(10): *141 / *142 / *143 / *144 / *146, *147 / *148 / *149 / *150 / *151
  ,CONDITIONED-REPORTS = *NO / *YES

```

***PRIOR-ACF(...)**

```

  RESOURCES = *ALL / list-poss(3): *CPU / *MEM / *PAG
  ,REPORT-NUMBER = *STD / *ALL / list-poss(3): *32 / *33 / *34
  ,CONDITIONED-REPORTS = *NO / *YES

```

***PUBSET(...)**

```

  REPORT-NUMBER = *STD / *ALL / list-poss(3): *262 / *263 / *264
  ,CONDITIONED-REPORTS = *NO / *YES

```

(part 6 of 8)

***RESPONSE-TIME(...)**

CONNECTION-SET = *ALL / list-poss(32): *GLOBAL / <alphanum-name 1..16> /
*ALL-EXCEPT(...)

*ALL-EXCEPT(...)

SET-NAME = <alphanum-name 1..16>

,REPORT-NUMBER = *STD / *ALL / list-poss(25): *19 / *20 / *21 / *22 / *23 / *24 / *25 / *26 / *27 /
 *46 / *47 / *48 / *49 / *50 / *83 / *88 / *89 / *90 / *91 / *92 / *93 / *107 / *108 /
 *109 / *110

,CONDITIONED-REPORTS = *NO / *YES

***RST(...)**

CATEGORY = *ALL / list-poss(32): <alphanum-name 1..7>

,REPORT-NUMBER = *STD / *ALL / list-poss(3): *59 / *60 / *61

,CONDITIONED-REPORTS = *NO / *YES

***SERVICETIME(...)**

DEVICE = *ALL / list-poss(256): <alphanum-name 1..4>

,REPORT-NUMBER = *STD / *ALL / *231

,CONDITIONED-REPORTS = *YES / *NO

***SESAM-SQL(...)**

REPORT-NUMBER = *STD / *ALL / list-poss(15): *294 / *295 / *296 / *297 / *298 / *299 / *300 / *301 /
 *302 / *303 / *304 / *305 / *306 / *307 / *308

,CONDITIONED-REPORTS = *NO / *YES

***STD(...)**

REPORT-NUMBER = *STD / *ALL

,CONDITIONED-REPORTS = *NO / *YES

***STORAGE-SYSTEM-SYMMETRIX(...)**

SERIAL-NUMBER = *ALL / list-poss(32): <alphanum-name 3..12>

,REPORT-NUMBER = *STD / *ALL / list-poss(25): *232, *233, *234, *235, *236,
 *237, *238, *229, *240, *241, 242, *243, *244, *245, *246, *247, *248, *249,
 *250, *251, *252, *253, *254, *255, *256

,CONDITIONED-REPORTS = *NO / *YES

(part 7 of 8)


```

*SVC(...)
  SVC-NUMBER = *ALL / list-poss(64): <integer 0..255>
  ,REPORT-NUMBER = *STD / *ALL / *123
  ,CONDITIONED-REPORTS = *YES / *NO

*TASK(...)
  REPORT-NUMBER = *STD / *ALL / *5
  ,CONDITIONED-REPORTS = *NO / *YES

*TCP-IP(...)
  REPORT-NUMBER = *STD / *ALL / list-poss(2): *186 / *187
  ,CONDITIONED-REPORTS = *YES / *NO

*TLM(...)
  LOCK-NAME = *ALL / list-poss(32): <alphanum-name 1..8>
  ,REPORT-NUMBER = *STD / *ALL / list-poss(2): *96, *97
  ,CONDITIONED-REPORTS = *YES / *NO

*UDS-SQL(...)
  REPORT-NUMBER = *STD / *ALL /
    list-poss(10): *284 / *285 / *286 / *287 / *288 / *289 / *200 / *291 / *292 / *293
  ,CONDITIONED-REPORTS = *NO / *YES

*UTM(...)
  APPLICATION = *ALL / list-poss(32): <alphanum-name 1..8>
  ,REPORT-NUMBER = *STD / *ALL / list-poss(7): *128 / *129 / *130 / *131 / *132 / *133 / *225
  ,CONDITIONED-REPORTS = *NO / *YES

*VM2000(...)
  REPORT-NUMBER = *STD / *ALL /
    list-poss(6): *98 / *99 / *164 / *267 / *268 / *269
  ,CONDITIONED-REPORTS = *NO / *YES

*VOLUME(...)
  VOLUME = *ALL / list-poss(32): <vsfn 1..6> / <alphanum-name 1..6>
  ,REPORT-NUMBER = *STD / *ALL / *NONE / *12
  ,CONDITIONED-REPORTS = *YES / *NO

*WORKING-SET(...)
  REPORT-NUMBER = *STD / *ALL / *15
  ,CONDITIONED-REPORTS = *NO / *YES

```

(part 8 of 8)



In connection with certain variables, a value “0” computed during the analysis subinterval is not interpreted. As a result these analysis subintervals are ignored in the calculation of average values. If all the subintervals have the value “0”, output of the monitoring data is suppressed. This could lead to entire subreports not being output.

Operands

INFORMATION =

Defines how the analyzed data is to be output.

INFORMATION = *DIAGRAMS

The monitoring data is output as time series in the form of charts and statistical tables.

INFORMATION = *STATISTICS

The monitoring data is output in the form of statistical tables.

INFORMATION = *INTERFACE

The requested monitoring data is output to a transfer file. The link name for the transfer file is EVALDTA.

REPORT-LIST =

Defines the report groups to be analyzed (in a list).

REPORT-LIST = *STD

The CHANNEL, CPU, DISK, IO, MEMORY, RESPONSE-TIME, TASK and WORKING-SET report groups are analyzed.

REPORT-LIST = *ALL(...)

All report groups (except DEVICE and VOLUME) are analyzed.

REPORT-NUMBER =

The report numbers can be used to select specific reports in a group (in this case all reports or the standard reports).

REPORT-NUMBER = *ALL

All reports in a report group are analyzed.

REPORT-NUMBER = *STD

The most important reports in a report group are analyzed. The corresponding report numbers are contained in the report group table, see [page 611](#).

CONDITIONED-REPORTS =

It is possible to suppress subreports whose monitoring data lies within the preset range. See also section "[Reducing the number of monitored objects to be output](#)" on [page 446](#) for details on the preset ranges.

CONDITIONED-REPORTS = *NO /*YES

The subreports whose monitoring data lies within the predefined range are either output or suppressed.



The REPORT-LIST, REPORT-NUMBER and CONDITIONED-REPORTS operands in this statement always have the same meaning. They are only shown in shortened form or not at all below.

REPORT-LIST = *BCAM-CONNECTION(...)

The report group BCAM-CONNECTION is analyzed.

CONNECTION-SET =

Defines the selected connection sets and calculates statistics and charts for them.

CONNECTION-SET = *ALL

Analyzes the response data for all the connection sets present in the SM2 output file.

CONNECTION-SET = list-post(32): <alphanum-name 1..16>

Specifies the name of the connection set that is to be analyzed. The name of the connection set is defined in SM2 (ADD-BCAM-CONNECTION-SET).

REPORT-LIST = *CATALOG-MANAGEMENT(...)

The CATALOG-MANAGEMENT report group is analyzed.

CATALOG-ID =

Specifies the catalog IDs to be analyzed.

CATALOG-ID = *ALL

Analyzes all monitored catalog IDs.

CATALOG-ID = list-poss(32): <catid>

Specifies the catalog ID to be analyzed.

REPORT-LIST = *CATEGORY-CPU(...)

The CATEGORY-CPU report group is analyzed.

CATEGORY =

Specifies the categories to be analyzed. For each category, a separate chart is output and a separate group entry is made in the table.

CATEGORY = *ALL

Analyzes all monitored categories.

CATEGORY = list-poss(32): <alphanum-name 1..7>

Specifies the name of the category to be analyzed. The following abbreviations can be used for the category name:

TP	Transaction tasks category
SYS	SYSTEM category
DIALOG	DIALOG category
BATCH	BATCH category
SUM	Total values are output for all categories.

The abbreviations defined during system generation can be used for categories other than the standard categories.

REPORT-LIST = *CATEGORY-IO(...)

The CATEGORY-IO report group is analyzed. The structure and meaning of this operand value is the same as described under *CATEGORY-CPU.

REPORT-LIST = *CATEGORY-QUEUE(...)

The CATEGORY-QUEUE report group is analyzed. The structure and meaning of this operand value is the same as described under *CATEGORY-CPU.

REPORT-LIST = *CATEGORY-WORKING-SET(...)

The CATEGORY-WORKING-SET report group is analyzed. The structure and meaning of this operand value is the same as described under *CATEGORY-CPU.

REPORT-LIST = *CHANNEL(...)

The CHANNEL report group is analyzed.

CHANNEL-PATH-ID =

Specifies channels to be analyzed. For each specified channel, a chart is output or a group is entered in the table.

CHANNEL-PATH-ID = *ALL

Analyzes all the monitored channels.

CHANNEL-PATH-ID = list-poss(32): <x-string 1..4>

Analyzes only the specified channel addresses. The values from x'0' to x'1FF' are permitted as channel addresses.

In SM2 the channel path ID is always output in hexadecimal form.

REPORT-LIST = *CPU(...)

The CPU report group is analyzed.

PROCESSOR-SPLITTING =

Defines whether the values of individual processors are to be analyzed if SM2 has monitored a multiprocessor system.

PROCESSOR-SPLITTING = *NO

The average values of the processors are indicated in the chart or in the group of tables.

PROCESSOR-SPLITTING = *YES

The individual values of the processors are indicated as a subreport in one chart each or in one group of tables.



CONDITIONED-REPORTS=*YES always applies to report 6, irrespective of the specifications made in the statement.

REPORT-LIST = *DAB(...)

The DAB report group is analyzed.

CACHE-ID =

Specifies the DAB cache areas to be analyzed.

CACHE-ID = *ALL

Analyzes all monitored DAB cache areas.

CACHE-ID = list-poss(32): <alphanum-name 1..32>

Specifies the name of the DAB cache area to be analyzed.

REPORT-LIST = *DEVICE(...)

The DEVICE report group is analyzed.

DEVICE =

Specifies devices for which an analysis is to be carried out. A chart or a group in the table is output for each device specified.

DEVICE = *ALL

Analyzes all monitored devices.

DEVICE = *SPECIFIED(...)

Analyzes only the specified devices. The mnemonic device name and/or the device type defines the individual devices.

DEVICE = *NOT-SPECIFIED

No devices are selected on the basis of their device names.

DEVICE = list-poss(256): <alphanum-name 1..4>

The mnemonic device names of the devices to be analyzed are specified.

DEVICE-TYPE = *NOT-SPECIFIED

No devices are specified on the basis of their device type.

DEVICE-TYPE = *TAPE

All devices with the FAMILY name TAPE, UNMTAPE, MBK and BIMTAPE are analyzed.

DEVICE-TYPE = *TD

All devices with the FAMILY name TD are analyzed.

DEVICE-TYPE = *OTHER

All devices that are not disks and do not belong to DEVICE-TYPE=*TD or *TAPE are analyzed.

REPORT-LIST = *DILATION(...)

The DILATION report group is analyzed. The dilation value output in this report is an internally computed value, i.e. not the PCS dilation value. The structure of this operand value and its meaning are the same as described for the operand *CATEGORY-CPU.

REPORT-LIST = *DISK(...)

The DISK report group is analyzed.

SPECIFIED =

Specifies the devices and volumes to be analyzed. A new subreport is output if the device or volume changes.

SPECIFIED = *ALL

All devices and volumes of the *DISK report group are analyzed.

SPECIFIED = *DEVICE(...)

The specified devices and all the corresponding volumes are analyzed.

DEVICE = list-poss(32): <alphanum-name 2..4>

Specifies the device.

SPECIFIED = *VOLUME-AND-DEVICE(...)

The specified devices and volumes are analyzed.

DEVICE-VOLUME = list-poss(32): *SELECT(...)

The specified devices and volumes are analyzed.

DEVICE = <alphanum-name 2..4>

Specifies the device.

VOLUME = <vsn 1..6>

Specifies the volume assigned to the device.

REPORT-LIST = *FILE(...)

The FILE report group is analyzed.

FILE-NAME =

Specifies the files for which a special analysis is to be performed. A chart page or a group in the table is output for each specified file.

FILE-NAME = *ALL

Analyzes all files which have been monitored.

FILE-NAME = list-poss(32): <filename>

Analyzes only the specified files. The name of the file specified must be a fully qualified name. Values of files not specified are not included in the list.

REPORT-LIST = *GS(...)

The GS report group is analyzed.

PARTITION =

Specifies the partitions to be analyzed.

PARTITION = *ALL

Analyzes all partitions.

PARTITION = list-poss(32): <alphanum-name 1..8>

Specifies a list of partitions to be analyzed.

REPORT-LIST = *GSVOL(...)

The GSVOL report group is analyzed.

DEVICE-MNEMONIC =

Specifies the devices to be analyzed via their mnemonic device names.

DEVICE-MNEMONIC = *ALL

Analyzes all devices which are specified via their mnemonic device names.

DEVICE-MNEMONIC = list-poss(32): <alphanum-name 1..4>

Specifies a list of devices to be analyzed.

REPORT-LIST = *IO(...)

The IO report group is analyzed.

PROCESSOR-SPLITTING =

Specifies whether the I/O values of the individual processors are to be analyzed if SM2 has monitored a multiprocessor system.

PROCESSOR-SPLITTING = *NO

The cumulative I/O values of the processors are shown in the chart or in the group of tables.

PROCESSOR-SPLITTING = *YES

The individual I/O values of the processors are shown as one subreport in each chart or in the group of tables.

REPORT-LIST = *ISAM(...)

The ISAM report group is analyzed.

ISAM-POOL =

Specifies the ISAM pools to be analyzed.

ISAM-POOL = *ALL

Analyzes all ISAM pools.

ISAM-POOL = list-poss(32): <alphanum-name 1..8>
Specifies a list of ISAM pools to be analyzed.

REPORT-LIST = *ISAM-FILE(...)
The ISAM-FILE report group is analyzed.

FILE-NAME =
Specifies the ISAM file to be analyzed.

FILE-NAME = *ALL
Analyzes all ISAM files.

FILE-NAME = list-poss(32): <alphanum-name 1..54>
Specifies a list of ISAM files to be analyzed.

REPORT-LIST = *PCS(...)
The PCS report group is analyzed. The structure of this operand value and its meaning are the same as described for the operand *CATEGORY-CPU.

REPORT-LIST = *PERIODIC-TASK-JOBNAME(...)
The PERIODIC-TASK-JOBNAME report group is analyzed.

JOBNAME =
Specifies the job names of the tasks to be analyzed.

JOBNAME = *ALL
All tasks selected via job names are analyzed.

JOBNAME = list-poss(32): <alphanum-name 1..8>
The tasks with the specified job names are analyzed.

REPORT-LIST = *PERIODIC-TASK-TSN(...)
The PERIODIC-TASK-TSN report group is analyzed.

TASK =
Specifies the tasks to be analyzed.

TASK = *ALL
Analyzes all the monitored tasks.

TASK = *SPECIFIED(...)
Analyzes only the specified tasks.

TSN =
Specifies the TSNs of the tasks to be analyzed.

TSN = *NOT-SPECIFIED

No tasks are selected on the basis of their TSN.

TSN = list-poss(32): <alphanum-name 1..4>

The tasks with the specified TSNs are analyzed.

USER-ID =

Specifies the user IDs of the tasks to be analyzed.

USER-ID = *NOT-SPECIFIED

No tasks are selected on the basis of their user ID.

USER-ID = list-poss(32): <alphanum-name 1..8>

The tasks with the specified user IDs are evaluated.

REPORT-LIST = *PERIODIC-TASK-USERID(...)

The PERIODIC-TASK-USERID report group is analyzed.

USERID =

Specifies the user IDs of the tasks to be analyzed.

USERID = *ALL

All tasks selected via user ID are analyzed.

USERID = list-poss(32): <alphanum-name 1..8>

The tasks with the specified user IDs are analyzed.

REPORT-LIST = *PFA(...)

The PFA report group is analyzed.

CACHE-ID =

Specifies the cache areas.

CACHE-ID = *ALL

All cache areas are analyzed.

CACHE-ID = list-poss(32): <alphanum-name 1..4>

The specified cache areas are analyzed.

REPORT-LIST = *PRIOR-ACF(...)

The PRIOR-ACF report group is analyzed.

RESOURCES =

Specifies the resources to be analyzed.

RESOURCES = *ALL

A chart containing the table of statistical values per group entry is output for the resources central processor, main memory and paging.

RESOURCES = *CPU

Analyzes the entries for the CPU.

RESOURCES = *MEM

Analyzes the entries for main memory.

RESOURCES = *PAG

Analyzes the entries for paging.

REPORT-LIST = *RESPONSE-TIME(...)

The RESPONSE-TIME report group is analyzed. Reports 89 to 93, which are part of this group, are broken down by categories, provided that SM2 has registered category-specific information (in record 39). If so, the category name is included in the page header, and SM2R1 outputs a separate chart and separate statistics for each monitoring definition. Unless otherwise specified, multiple definitions are processed in the order in which they occur. Only the first 32 definitions are used.



In the RESPONSE-TIME report group, the name of the connection set or connection set group is output after the monitoring data. This name is assigned by the user by means of the SM2 statement ADD-CONNECTION-SET.

CONNECTION-SET =

Specifies selected connection sets and computes statistics and draws charts for them.

CONNECTION-SET = *ALL

Analyzes the response data of all the connection sets included in the SM2 output file.

CONNECTION-SET = *GLOBAL

Analyzes the response data for the connection of all applications, partners and processors.

CONNECTION-SET = list-poss(32): <alphanum-name 1..16>

Specifies the name of the connection set to be analyzed. The name of the connection set is defined by SM2 (ADD-CONNECTION-SET).

CONNECTION-SET = *ALL-EXCEPT(...)

Selects a connection set that is not to be analyzed.

SET-NAME = <alphanum-name 1..16>

Specifies the name of the connection set that is not to be analyzed.

REPORT-LIST = *RST(...)

The RST report group is analyzed. The structure of this operand value and its meaning are the same as described for the operand *CATEGORY-CPU.

REPORT-LIST = *SERVICETIME(...)

The SERVICETIME report group is evaluated.

DEVICE =

Specifies devices for which an analysis is to be carried out. A chart or a group in the table is output for each device specified.

DEVICE = *ALL

Analyzes all the monitored devices.

DEVICE = list-poss(256): <alphanum-name 1..4>

Analyzes the devices that are specified by means of their mnemonic device name.

REPORT-LIST = *STD(...)

The CHANNEL, CPU, DISK, IO, MEMORY, RESPONSE-TIME, TASK and WORKING-SET report groups are analyzed.

REPORT-LIST = *STORAGE-SYSTEM-SYMMETRIX(...)

The STORAGE-SYSTEM-SYMMETRIX report group is analyzed.

SERIAL-NUMBER =

Specifies the symmetrix systems to be analyzed.

SERIAL-NUMBER = *ALL

Analyzes all symmetrix systems.

SERIAL-NUMBER = list-poss(32): <alphanum-name 3..12>

Analyzes the symmetrix systems that are specified by their serial number.

REPORT-LIST = *SVC(...)

The SVC report group is analyzed.

SVC-NUMBER =

Specifies the SVC numbers to be analyzed.

SVC-NUMBER = *ALL

All SVC numbers are to be analyzed.

SVC-NUMBER = list-poss(64): <integer 0..255>

The specified SVC numbers are analyzed.

REPORT-LIST = *TLM(...)

The TLM report group is analyzed.

LOCK-NAME =

Specifies the names of the locks to be analyzed.

LOCK-NAME = *ALL

All locks are analyzed.

LOCK-NAME = list-poss(32): <alphanum-name 1..8>

The locks with the specified names are analyzed.

REPORT-LIST = *UTM(...)

The UTM report group is analyzed.

APPLICATION =

Specifies the UTM applications.

APPLICATION = *ALL

All UTM applications are analyzed.

APPLICATION = list-poss(32): <alphanum-name 1..8>

The specified UTM applications are analyzed.

REPORT-LIST = *VOLUME(...)

The VOLUME report group is analyzed.

VOLUME =

Specifies the volumes for which an analysis is to be performed. A chart or a table entry is output for every volume specified.

VOLUME = *ALL

Analyzes all volumes which have been monitored.

VOLUME = list-poss(32): <vsn 1..6> / <alphanum-name 1..6>

Analyzes only the specified volumes. The individual volumes are specified by name.

Report groups

Group names allow selection of a group of monitoring data reports for analysis. Monitored objects can be specified for report groups in the same way as for monitoring data reports. A new page is used for each monitored object.

In all reports which allow variable numbers of monitored objects, ALL is preset as the default value. This takes all monitored objects in the SM2 output file into account.

You will find an overview of all group names and reports starting on [page 611](#).



If PROCESSOR-SPLITTING=*YES is specified as the monitored object in the *CPU (reports 1, 2, 137) and *IO report groups, the following applies.

The following subreports are output when a multiprocessor system is involved:

- values for all processors
- values for each individual processor

If no monitored object is specified or in the case of PROCESSOR-SPLITTING=*NO, only the accumulated data for the processors is output.

If REPORT-LIST=*ALL is specified, the *DEVICE and *VOLUME report groups are not analyzed.

Report names

The report name is contained in the heading of each report page. For a list of report names see [section “Table of report names” on page 614](#).

Reducing the number of monitored objects to be output

The user can specify whether subreports which contain “normal” monitored data, i.e. whose data lies within a defined range, are to be output or not

A range is defined for every value monitored. If even one item of monitoring data lies outside this range, the subreport is still output and contains all monitoring data.

This reduction in the number of subreports to be output applies to charts and statistics. All reports are always output in the interface file and the summaries.

The ranges are preset as follows (see MODIFY-REPORT-CONDITIONS on [page 418](#) for an explanation of the names):

REPORT-GROUP	REPORT-NUMBER	ITEM-NUMBER	LOWER-LIMIT	UPPER-LIMIT	SUPPRESS-CONDITION	COMPARE-VALUE
*VOLUME	12	1	0	15	*INSIDE	*MAXIMUM
		2	0	15	*INSIDE	*MAXIMUM
*DEVICE	11	1	0	15	*INSIDE	*MAXIMUM
		2	0	15	*INSIDE	*MAXIMUM
	35	1	0	15	*INSIDE	*MAXIMUM
	36	1	0	0,3	*INSIDE	*MEAN-VALUE
	100	1	0	100	*INSIDE	*MEAN-VALUE
		2	0	50	*INSIDE	*MEAN-VALUE
230	1	0	1	*INSIDE	*MEAN-VALUE	
	2	0	0,5	*INSIDE	*MEAN-VALUE	
*DISK	124	1	0	15	*INSIDE	*MAXIMUM
		2	0	15	*INSIDE	*MAXIMUM
	125	1	0	15	*INSIDE	*MAXIMUM
	126	1	0	0,3	*INSIDE	*MEAN-VALUE
	127	1	0	2	*INSIDE	*MEAN-VALUE
	227	1	0	1	*INSIDE	*MEAN-VALUE
		2	0	1	*INSIDE	*MEAN-VALUE
	228	1	0	10	*INSIDE	*MEAN-VALUE
		2	0	5	*INSIDE	*MEAN-VALUE
229 ¹	-	-	-	*INSIDE	*MEAN-VALUE	
*CHANNEL	10	1	0	10	*INSIDE	maximum
		2	0	10	*INSIDE	*MEAN-VALUE
	101	1	0	80	*INSIDE	*MEAN-VALUE
		2	0	250	*INSIDE	*MAXIMUM
	102	1	0	40	*INSIDE	*MAXIMUM
		2	0	40	*INSIDE	*MAXIMUM
3	0	40	*INSIDE			
*PERIODIC-TASK	152,156,160	1	0	1000	*INSIDE	*MAXIMUM
		1	0	5	*INSIDE	*MEAN-VALUE
	153,157,161	1	0	10	*INSIDE	*MEAN-VALUE
		2	0	10	*INSIDE	*MAXIMUM
-JOBNAME/	154,158,162	1	0	10	*INSIDE	*MEAN-VALUE
-TSN		2	0	10	*INSIDE	*MAXIMUM
-USERID	155,159,163	1	0	1000	*INSIDE	*MEAN-VALUE
*TLM	96	1	0	2	*INSIDE	*MAXIMUM
	97	1	0	10	*INSIDE	*MEAN-VALUE
*SVC	123	1	0	20	*INSIDE	*MEAN-VALUE
		2	0	20	*INSIDE	*MEAN-VALUE

REPORT-GROUP	REPORT-NUMBER	ITEM-NUMBER	LOWER-LIMIT	UPPER-LIMIT	SUPPRESS-CONDITION	COMPARE-VALUE
*MSCF	166	1	0	50	*INSIDE	*MEAN-VALUE
		2	0	0,1	*INSIDE	*MEAN-VALUE
	167	1,2	0	4	*INSIDE	*MAXIMUM
	168	1	0	50	*INSIDE	*MEAN-VALUE
	169 ¹	2	0	0,1	*INSIDE	*MEAN-VALUE
*GSVOL	174	1,2	0	100	*INSIDE	*MAXIMUM
*GS	177	1,2	0	100	*INSIDE	*MEAN-VALUE
	178	1,2	0	100	*INSIDE	*MAXIMUM
*TCP-IP	186	1,2	0	1	*INSIDE	*MAXIMUM
	187	1,2	0	300	*INSIDE	*MAXIMUM
*SERVICE-TIME	231	1	0	15	*INSIDE	*MEAN-VALUE
		2	0	5	*INSIDE	*MEAN-VALUE
		3	0	1	*INSIDE	*MEAN-VALUE
		4	0	10	*INSIDE	*MEAN-VALUE
		5	0	2	*INSIDE	*MEAN-VALUE

¹ Information on a reduction of the number of monitored objects to be output is not meaningful.

With these report groups, the CONDITIONED-REPORTS operand in the PRINT-REPORTS statement is preset to *YES.

The ranges for all other monitoring data are preset as follows (these settings mean that all reports are suppressed):

LOWER-LIMIT = 0

UPPER-LIMIT = $1 * 10^{63}$

SUPPRESS-CONDITION = *INSIDE-RANGE

COMPARE-VALUE = *MEAN-VALUE

In addition, the user can change the preset ranges for a session using the MODIFY-REPORT-CONDITIONS statement.

PRINT-SUMMARY

Print summary analysis

Summary analysis groups the essential data of given areas, e.g. for catalog management or PCS monitoring purposes. It provides a quick overview of system performance. The monitoring data is output in the form of tables.

For some monitoring data, only the average for the whole analysis period is provided (mean values), while in other cases the minimum and maximum values are provided as well.

Format

PRINT-SUMMARY
<pre> PARTITION = <u>*EVALUATION-PERIOD</u> / <integer 1..1000>(…) <integer 1..1000>(…) DIMENSION = <u>*MINUTES</u> / *SECONDS / *HOURS / *DAYS ,SUMMARY-REPORTS = <u>*ACTIVITY</u> / *ALL / list-poss(6): *ACTIVITY / *CMS / *DAB / *PCS / *POSIX / *UTM </pre>

Operands

PARTITION =

Defines the number of output operations. Outputs are made at regular intervals as soon as the specified period has elapsed, or more specifically, at the end of the last SM2 monitoring cycle to be started.

PARTITION = *EVALUATION-PERIOD

Only one output is made at the end of the analysis period (or at the end of the file if the file does not cover the whole analysis period).

PARTITION = <integer 1..1000>

Defines the period between the individual outputs.

DIMENSION = *MINUTES / *SECONDS / *HOURS / *DAYS

Defines the time units which determine the period between the individual outputs. If *HOURS or *DAYS is specified, this also stipulates full hours/days for the beginning of the first output cycle.

SUMMARY-REPORTS = *ACTIVITY / *ALL / list-poss(6): *ACTIVITY / *CMS / *DAB / *PCS / *POSIX / *UTM

Defines the SUMMARY reports to be output. These are the ACTIVITY, CMS, DAB, PCS, POSIX and UTM reports. If the user specifies the operand value *ALL, all the existing SUMMARY reports are output.

Example

```
//PRINT-SUMMARY PARTITION=60,SUMMARY-REPORTS=*PCS
```

Within the analysis period, this statement supplies a PCS summary report every 60 minutes.

SUMMARY ACTIVITY report

The summary report on activity contains an overview of important data monitored during the analysis period and combined in various groups. For most of these values, the minimum/maximum values and the standard deviation are output in addition to the average over the observed period.

REPORTING PERIOD

SYSTEM DATA

#TASKS

QUEUE STATISTICS

RESPONSE TIME STATISTICS

CPU STATISTICS

IO STATISTICS

MEMORY STATISTICS

PAGINGAREA STATISTICS

CHANNEL STATISTICS

DEVICE STATISTICS

All monitoring data refers to the specified analysis period.

REPORTING PERIOD

The first and last time stamp found within the analysis period are output. The BY entry corresponds to the SM2 gathering cycle. FOCUS and EXCEPT specifications valid for the analysis period must be taken from the preceding cover page.

If EXCEPT has been specified for the analysis, the statistics contain the message

```
*** EXCEPTION PERIOD EXISTS ***.
```

Next, the system data applying to the appropriate partition is output:

SYSTEM DATA

For a data description see [section "Monitoring environment output" on page 401](#).

Then the following monitoring data is output:

#TASKS

This lists the average number of tasks within the analysis period. Assignment to the respective categories occurs at task creation.

Monitored variable	Meaning
SYS	See report 5
BATCH	See report 5
DIALOG	See report 5
TP	See report 5

QUEUE STATISTICS

The number of queued tasks is output for the entire system, i.e. all categories are consolidated.

Monitored variable	Meaning
IN CPU QUEUE	See report 31
IN PAGING QUEUE	See report 31
IN IO QUEUE	See report 31
ACTIVE	See report 30
INACT READY	See report 30
NOT ADMITTED	See report 30
INACT NOT READY	See report 30
TOTAL	See report 28

RESPONSE TIME STATISTICS

The monitoring data is not output unless the RESPONSETIME monitoring program was active during the analysis period. If only one definition was active, blanks are output for the other definitions. The output comprises system-global values, i.e. the think, response and transaction times over all connections.

Monitored variable	Meaning
MEAN THINK TIME	See report 46
MEAN RESPONSE TIME (1)	See report 47
MEAN RESPONSE TIME (2)	See report 48
MEAN TRANSACTION TIME	See report 49
TRANSACTION RATE	See report 50
RESPONSE RATE (1)	See report 50
RESPONSE RATE (2)	See report 50

CPU STATISTICS

For multiprocessor servers, the average value of all processors is output.

Monitored variable	Meaning
TU TIME	See report 1
TPR TIME	See report 1
SIH TIME	See report 1
IDLE TIME	See report 1
STOP TIME	See report 1
ACTIVE LOGICAL MACHINES	See report 6

IO STATISTICS

For multiprocessor servers, the total of all processors is output.

Monitored variable	Meaning
NON PAGING DISK IO' S	See report 3
PAGING IO' S	See report 3
TAPE IO' S	See report 3
PRINTER IO' S	See report 3
OTHER	See report 3

MEMORY STATISTICS

For the page fault statistics and the quotient from working set and available main memory, only the values for AVERAGE are output.

Monitored variable	Meaning
#CLASS 3 PAGES	See report 53
#CLASS 4 PAGES	See report 53
AVAILABLE PAGES (NPP)	See report 15
WSET ACT TASKS (PPC)	See report 15
WSET INACT READY TASKS (PPC)	See report 15
USED PAGES ACT TASKS	See report 58
USED PAGES INACT TASKS	See report 58
(WSET ACT + INACT READY)/NPP	Quotient from the planned working set of active and inactive ready tasks and the available main memory
TOTAL # PAGE FAULTS	See report 54
# PAGE RECLAIMS	See report 54
# FIRST PAGE ACCESS	See report 54
# PAGE WRITES TO DISK	See report 55
# PAGE READS FROM DISK	See report 56

PAGINGAREA STATISTICS

Monitored variable	Meaning
PAGES ON PAGING DEVICE(S)	See report 9
USED PAGES ON PAGING DEVICE(S)	See report 9
PAGES ON EXPANDED STORAGE ¹	See report 9

¹ The monitored variable has no meaning for output files created in SM2 V15.0 or higher.

CHANNEL STATISTICS

The values for the ten most heavily loaded channels are output in the channel statistics. The values for CHANNEL IO RATE and CHANNEL PAM PAGES are only output if the CHANNEL-IO monitoring program was active during the monitoring period.

Monitored variable	Meaning
CHANNEL BUSY STATE	See report 10
CHANNEL IO RATE	See report 102
CHANNEL PAM PAGES	See report 101

DEVICE STATISTICS

The values for the ten most heavily loaded devices are output in the device statistics. To this end, the value for DEVICE BUSY STATE (NON PAGING) is used.

Monitored variable	Meaning
DEVICE BUSY STATE (NON PAGING)	See report 11
DEVICE BUSY STATE (PAGING)	See report 11
DEVICE IO RATE	See report 35
DEVICE PAM PAGES	See report 100

CMS summary report

All monitored catalog management data is analyzed if the user specifies the operand value *CMS.

Up to five catalog identifiers are output on one page. The catalog identifiers are sorted in alphabetical order.

The text POOL OF PRIVATE DISKS is output instead of the \$ character (see reports 66 to 72) to denote the complete range of private disks.

If available, monitoring data for private disks is always output first. The average as well as the minimum and the maximum values are supplied for all monitored variables.

Monitored variable	Meaning
LENGTH OF QUEUES SERIAL REQUEST USERID BLOCK CATALOG ENTRY SPEEDCAT REQUEST CAT INDEX REQ	See report 66 See report 66 See report 66 See report 66 See report 66 See report 185 See report 185
ACCESSES [1/S] PHYSICAL READ PSEUDO READ PHYSICAL WRITE IO-ERROR	See report 67 See report 67 See report 67 See report 67
RESP. TIME [MS] READ NO LBN READ LBN SCAN UPDATE / RENAME WRITE / CLEAR	See report 72 See report 72 See report 72 See report 72 See report 72
LOC FILE/JV [1/S] READ NO LBN READ LBN SCAN UPDATE / RENAME WRITE / CLEAR	See reports 68 and 70 See reports 68 and 70 See reports 68 and 70 See reports 68 and 70 See reports 68 and 70

continued

Monitored variable	Meaning
REM FILE/JV [1/S] READ NO LBN READ LBN SCAN UPDATE / RENAME WRITE / CLEAR	See reports 69 and 71 See reports 69 and 71 See reports 69 and 71 See reports 69 and 71 See reports 69 and 71
SH FILE/JV [1/S] READ NO LBN READ LBN SCAN UPDATE / RENAME WRITE / CLEAR	See reports 103 and 104 See reports 103 and 104 See reports 103 and 104 See reports 103 and 104 See reports 103 and 104



The data from LOCAL FILE and LOCAL JV, REMOTE FILE, as well as from REMOTE JV, SHARED PUBSET FILE and SHARED PUBSET JV is added together.

DAB summary report

If the user specifies the operand value *DAB, the summary statistics for the various DAB cache areas are output.

As with all SM2R1 summaries, the average as well as the minimum and the maximum values are supplied for all monitoring data.

The following information is also output for each DAB cache area:

MEDIUM Cache medium:
 MM = main memory
 GS = global storage

SIZE Size of the DAB cache area in MB.

Monitored variable	Meaning
DAB READ REQUESTS HITS	See report 79 See report 79
DAB WRITE REQUESTS HITS	See report 80 See report 80

PCS summary report

The operand value *PCS serves to select the summary statistics for the Performance Control Subsystem.

This report supplies more comprehensive data than the report group PCS. The categories are sorted in alphabetical order except for the standard categories SUM, SYS, DIALOG, BATCH and TP, which appear (in this order) at the beginning of the list.

Data not available for the pseudo-category SUM is replaced by blanks. The average as well as the minimum and the maximum values are supplied for all monitored variables.

As an additional option the name of the PCS option is output.

The following overview lists all the monitored data. The SUM column provides information about monitored variables that are output for the pseudo-category SUM.

Monitored variables	SUM	Meaning
SERVICE-QUOTA- ACTUAL PLANNED MAX MIN		See report 73 See report 73 See report 73 See report 73
REQUEST-DELAY- ACTUAL MAX MIN	x x	See report 74 See report 74 See report 74
THROUGHPUT QUOTA	x	Percentage specifying the throughput optimization share 0% → no optimization (full response time optimization) 100% → full local optimization (no response time opt.) For SUM, this variable corresponds to the global PCS parameter with the same name
DURATION # SERVICE UNITS		Number of used SERVICE UNITS after which an automatic category switch occurs
RUNOUTS RUNOUTS PREEMT	x x	See report 75 See report 75
SERVICE-RATE- ACTUAL CPU MEMORY IO PLANNED	x x x x x	See report 77 SERVICE UNITS used by the category, subdivided into CPU, MEM and I/O See report 77

POSIX summary report

If the user specifies the operand value *POSIX, the summary statistics for POSIX are output.

As in the case of all other SM2R1 summaries, the mean, maximum and minimum values are supplied for all monitored variable.

Monitored variables	Meaning
OPTION A [1/S] IGET NAMEI	See report 141 See report 142
OPTION B [1/S] LREAD BREAD LWRITE BWRITE PREAD PWRITE	See report 143 See report 143 See report 144 See report 144 See report 146 See report 146
OPTION C [1/S] SCALL SREAD SWRITE FORK EXEC RCHAR WCHAR	See report 147 See report 148 See report 148 See report 148 See report 148 See report 149 See report 149
OPTION M [1/S] MSG SEMA	See report 150 See report 151

UTM summary report

If the user specifies the operand value *UTM, the summary statistics for the universal transaction monitor are output.

As with all SM2R1 summaries, the average as well as the minimum and the maximum values are supplied for all monitoring data.

The following information is also output for every application:

- MODE Mode of the UTM application:
 S = UTM-S application
 F = UTM-F application
- #DB Number of database systems with which the application is coordinated.

Monitored variables	Meaning
# USERS	See report 129
# TASKS	See report 130
MAX # ASYN TASKS	See report 130
# DIAL STEPS [1/S]	See report 128
# DIAL TA	See report 128
# DIAL WITH DB	Number of dialog steps per second with database calls
# DIAL WITH DDP	Number of dialog steps per second with distributed processing calls
# ASYN CONV	Number of completed asynchronous conversations per second
# ASYN TA	See report 128
# ASYN WITH DB	Number of asynchronous conversations with database calls per second
# ASYN WITH DDP	Number of asynchronous conversations with distributed processing calls per second

continued

Monitored variables	Meaning
DIALOG [S/STEP] TOTAL TIME TOTAL T WITH DB TOTAL T WITH DDP TIME IN DB TIME IN DDP TACCL WAIT TIME	<p>In the case of dialog steps: total time in seconds per dialog step from acceptance of the input by UTM to the sending of the dialog message by</p> <p>In the case of asynchronous conversations: total time in seconds per asynchronous conversation from the beginning to the end of processing (without waiting time until the start)</p> <p>TOTAL TIME also includes waiting times that arise as a result of a TAC class bottleneck or of waiting for messages from remote applications. On the other hand, it does not include a waiting time in the job queue of the UTM application (i.e. before a UTM task accepts the job for the first time).</p> <p>Value as for TOTAL TIME, but only for dialog steps with database calls</p> <p>Value as for TOTAL TIME, but only for dialog steps with distributed processing calls</p> <p>Time UTM waits for execution of database calls per dialog step</p> <p>Time UTM waits for a message from a remote application per dialog step</p> <p>Average time UTM waits as a result of TAC class bottlenecks per dialog step</p>
ASYNCHR. [S/CONV] TOTAL TIME TOTAL T WITH DB TOTAL T WITH DDP TIME IN DB TIME IN DDP	<p>Total time in seconds for asynchronous conversation from the beginning to the end of processing (excluding wait time before the start)</p> <p>Value as for TOTAL TIME, but only for asynchronous conversations with database calls</p> <p>Value as for TOTAL TIME, but only for asynchronous conversations with distributed processing calls</p> <p>Time UTM waits for execution of database calls per asynchronous conversation</p> <p>Time UTM waits for a message from a remote application per asynchronous conversation</p>
DIALOG [1/STEP] # IO # IO IN DB # DB CALLS	<p>Number of IOs of UTM tasks during processing of the dialog step, including subroutines of the user</p> <p>Number of IOs in the called DB systems per dialog step</p> <p>Number of DB calls per dialog step</p>

continued

Monitored variables	Meaning
ASYNCHR. [1/CONV] # IO # IO IN DB # DB CALLS	Number of IOs of UTM tasks during processing of the asynchronous operation, including subroutines of the user Number of IOs in the called DB systems per asynchronous operation Number of DB calls per asynchronous operation
DIALOG [MS/STEP] CPU TIME CPU TIME IN DB	CPU time used by UTM tasks in milliseconds per dialog step, including subroutines of the user CPU time used in DB systems in milliseconds per dialog step
ASYNCHR. [MS/CONV] CPU TIME CPU TIME IN DB	CPU time used by UTM tasks in milliseconds per asynchronous operation, including subroutines of the user CPU time used in DB systems in milliseconds per asynchronous operation

PRINT-TASK-STATISTICS

Print task statistics

This analysis can cover some or all monitored tasks. The INFORMATION operand of the PRINT-TASK-STATISTICS statement defines the amount of detail desired in the output. Further classification by user IDs (USERID-SET), TSNs (TSN-SET), job name (JOBNAME-SET) and job class (JOBCLASS-SET) is possible.

A set name can be defined for a set of tasks via the CREATE-USERID-SET, CREATE-TSN-SET, CREATE-JOBNAME-SET and CREATE-JOBCLASS-SET statements. Output of these sets is initiated with the PRINT-TASK-STATISTICS statement.

Task statistics analysis

A requirement for task analysis is that the TASK monitoring program was active during SM2 monitoring. The task records in the SM2 output file begin with a START-MP record and end with a STOP-MP record for task statistics.

If several START-MP/STOP-MP pairs occur in an analysis period, several task analyses are output, i.e. each set of task statistics is output separately.

The individual tasks to be analyzed during a monitoring operation are output in the order in which they appear in the SM2 output file. If the analysis period is specified so that it begins or ends between a START-MP and a STOP-MP record, any measurement recorded by the TASK monitoring program whose START-MP record falls within the analysis period is output.

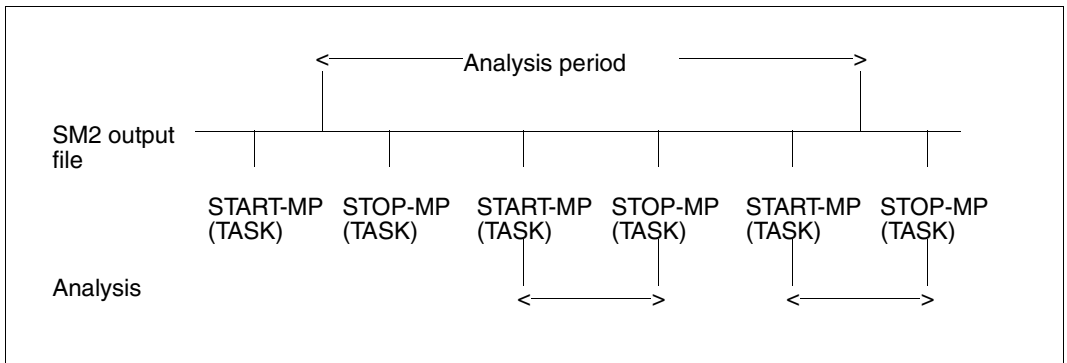


Figure 10: Monitored task data in the SM2 output file

Each monitoring operation is analyzed completely, i.e. from START- through STOP-MEASUREMENT-PROGRAM. Information on the times for START- and STOP-MEASUREMENT-PROGRAM can be obtained using the SM2U1 function SHOW-INFORMATION.

Format**PRINT-TASK-STATISTICS**

```

INFORMATION = *LOW / *HIGH
,TSN = *ALL / *BY-SET-NAME / *NONE / list-poss(64): <alphanum-name 1..4>
,USERID-SET = *NONE / list-poss(8): <alphanum-name 1..16>
,TSN-SET = *NONE / list-poss(8): <alphanum-name 1..16>
,JOBNAME-SET = *NONE / list-poss(8): <alphanum-name 1..16>
,JOBCLASS-SET = *NONE / list-poss(8): <alphanum-name 1..16>

```

Operands**INFORMATION =**

Defines the degree of detail for the task statistics.

The following data is supplied for each task to be analyzed:

- a) Consecutive task number (INDEX)

This number is supplied by the analysis routine to facilitate task identification. This index is incremented by one each time a task record is found.
- b) Identification section

This part serves to identify the task. It contains data such as user ID, task sequence number, job name, job class, etc..
- c) Monitoring data

The degree of detail is specified by the following operands:

INFORMATION = *LOW

The monitoring data is output without the data for the monitored devices.

INFORMATION = *HIGH

The monitoring data is output with the data for the monitored devices (ACCESSES and SERVICETIME).

TSN =

Defines the tasks to be analyzed.

TSN = *ALL

All tasks in the SM2 output file are analyzed according to the specified degree of detail.

TSN = *BY-SET-NAME

Specifies exactly the tasks that have been selected by means of a SET name in one of the operands USERID-SET, TSN-SET, JOBNAME-SET or JOBCLASS-SET.

TSN = *NONE

None of the tasks defined by TSN is analyzed. Only those tasks affected by class formation (see the CREATE-USERID-SET, CREATE-TSN-SET, CREATE-JOBNAME-SET, CREATE-JOBCLASS-SET statements) are analyzed if set names are specified for the operands described below.

TSN = list-poss(64): <alphanum-name 1..4>

The tasks defined by the task sequence number (TSN) are analyzed.

USERID-SET =

Selects tasks defined by the user ID for analysis.

USERID-SET = *NONE

None of the tasks defined by the user ID is analyzed.

USERID-SET = list-poss(8): <alphanum-name 1..16>

Set name defined in the CREATE-USERID-SET statement. The set name appears in the class header. Up to 8 such identifiers with 32 user IDs each may be formed. When these limits are exceeded, the surplus elements are ignored. SM2R1 lists the set names and the associated user IDs again on a separate sheet. The output values are total values for all the tasks in a class.

TSN-SET =

Selects tasks defined by means of the task sequence number (TSN) for analysis.

TSN-SET = *NONE

None of the tasks defined by the TSN is analyzed.

TSN-SET = list-poss(8): <alphanum-name 1..16>

Set name defined in the CREATE-TSN-SET statement. The set name appears in the class header. Up to 8 classes with 32 TSNs each may be formed. When these limits are exceeded, the surplus elements are ignored. SM2R1 lists the set names and the associated TSNs again on a separate sheet. The output values are total values for all the tasks in a class.

JOBNAME-SET =

Selects tasks specified by job name.

JOBNAME-SET = *NONE

No tasks specified by job name are selected.

JOBNAME-SET = list-poss(8): <alphanum-name 1..16>

Set name defined by the CREATE-JOBNAME-SET statement. The set name appears as the header for this class. Up to 8 such classes with 32 associated job names can be formed. SM2R1 lists the set names with associated job names again on a separate sheet. The output values are total values for all the tasks in a class.

JOBCLASS-SET =

Selects tasks specified by job name.

JOBCLASS-SET = *NONE

No tasks specified by job name are selected.

JOBCLASS-SET = list-poss(8): <alphanum-name 1..16>

Set name defined by the CREATE-JOBCLASS-SET statement. The set name appears as the header of this class. Up to 8 such classes with 32 associated job classes can be formed. SM2R1 lists the set names with the associated job classes on a separate sheet. The output values are total values for all tasks belonging to the class.

If the user specifies the set name, all tasks which do not belong to the above classes are combined to form a residual class. If the residual class is devoid of tasks, there is no output.

Example

```
//PRINT-TASK-STATISTICS TSN=(BCAM,PGE)
```

This statement supplies an analysis output for the tasks with the TSNs BCAM and PGE.



The statements SET-REPORT-FOCUS (time slots) and SET-EXCEPTION-PERIOD are ignored for task statistics analysis.

Task analysis supplies the following summaries by default.

To facilitate understanding, the following terms are defined:

- **TASK START TIME:** time at which task monitoring is started. For tasks which did not exist when monitoring was started, this is the LOGON time. Only for system tasks is this the task creation time.
- **TASK STOP TIME:** time at which monitoring of the task terminates. For tasks which are still running at monitoring termination time, this time is the task statistics STOP time (see the STOP-MP record). For tasks which were terminated before monitoring termination, this time is the task termination time (LOGOFF or task destruction).
- **TASK MONITORING CYCLE:** difference between task stop time and task start time.
- **VOLUNTARY ACTIVE WAIT STATE:** when a task is active, SM2 interprets the following times as voluntary wait times:
 - time spent in queue 2
 - time spent in queue 4 for boursing
 - time spent in queue 4 for inter-task communication (ITC)
 - time spent in queue 4 for PASS and VPASS
- **VOLUNTARY INACTIVE WAIT STATE:** When a task is inactive, SM2 interprets the following times as voluntary wait times:
 - time spent in queue 12 for boursing (this time is displayed separately)
 - time spent in queue 13 (PASS, VPASS, ITC)

Summaries supplied by task statistics by default

1. List sorted in descending order of the 20 tasks which used up most CPU time in the task monitoring cycle. For each task the consecutive task number (INDEX), the user ID (USERID), the task sequence number (TSN), the task category (CATEGORY), the CPU time (CPU-TIME) used in seconds, the number of I/O operations (#I/O'S), the name of the program with the highest CPU utilization (PROGRAM-NAME) and its CPU utilization (PGM-CPU) are supplied.
2. List sorted in descending order of the 20 tasks which performed most I/O operations during the task monitoring cycle (without paging I/Os). The same data is output for each task as in the preceding summary, referring this time to the number of I/O operations.
3. List sorted in descending order of the 20 tasks which occupy the maximum user address space. The tasks are sorted according to the maximum amount of class 5 and class 6 memory.

4. The monitored tasks are grouped according to category and the category data is output. Totals are supplied on a system basis (SUM) and per category. The degree of detail is the same as for LEVEL=HIGH or LEVEL=LOW without the identification part. It should be noted that the PGE task for device-specific ACCESSES and SERVICETIME values (data supplied per task) is not included in this summary.
5. Not more than the first 30 categories found in the SM2 output file (including the pseudo-category SUM) are output for sum and class; they are output in the following order: SUM, standard categories SYS, DIALOG, BATCH and TP, remaining categories (in alphabetical order). If more than 30 categories occur in a monitoring cycle, the surplus categories are ignored.
This does not affect task category output. The category to which the task was assigned at termination time is always output.

Output data per task

1. Identification part

The identification part comprises the following data:

USERID	–	User identification Blanks are output if the user ID cannot be determined (e.g. for system tasks).
TSN	–	Task sequence number.
JOBNAME	–	Job name (blanks are entered for system tasks).
JOBCLASS	–	Name of job class (blanks are entered for system tasks).
CATEGORY	–	Task category at task stop time.
START-DATE	–	Date at task start time.
START-TIME	–	Time of day at task start time.
END-DATE	–	Date at task stop time.
END-TIME	–	Time of day at task stop time.

2. Data supplied for LEVEL=LOW

ELAPSED TIME(S)

Same as task monitoring cycle.

TOTAL CPU-TIME(S)

Accumulated CPU time (TU and TPR shares of CPU time) in seconds since task initiation.

TOTAL # IO'S

Total number of input/output operations (except for paging) since task initiation (number of EXCP calls).

IO'S

Total number of input/output operations (except for paging) in the task monitoring cycle (number of EXCP calls).

SVC'S IN TU STATE

Number of SVC calls from the TU state in the task monitoring cycle.

SVC'S IN TPR STATE

Number of SVC calls from the TPR state in the task monitoring cycle.



This number does not contain the FUNCTION DETACH calls.

For the four following items refer to [“page fault” on page 668](#).

PAGE FAULTS

Number of page fault interrupts in the task monitoring cycle. “Real” page faults (access to virtually unassigned pages) are not included in this number.

PAGE READS

Number of pages read from background memory during the task monitoring cycle.

PAGE RECLAIMS

Number of page fault interrupts in the task monitoring cycle for which the addressed page was still in main memory.

1ST PAGE ACCESSES

Number of page fault interrupts for first access to a page in the task monitoring cycle.

WSET (PPC) IN PAGES

Average working set value (PPC) during the task monitoring cycle.

ESA PAGES

Maximum occupancy of DATA SPACE pages.

CLASS 5 + 6 PAGES

Maximum occupancy of user address space (total amount of class 5 and class 6 memory, including any common memory pools present).

CPU ALLOCATIONS

Number of requests to the CPU(s) during the task monitoring cycle. This number does not contain the requests for which the task retains the CPU after interrupts. (all exits from queue 0 to other queues are counted.)

CPU-TIME (S)

CPU time (TU and TPR shares of the CPU time) in the task monitoring cycle in seconds.

CPU WAITS

Number of wait states for CPU allocation (queue Q1)

DURATION (S)

Total time in seconds spent waiting for CPU allocation.

This also includes periods in which the CPU is active for other tasks in the SIH state or is occupied with management tasks.

ACTIVE WAITS

Number of active wait states in the TASK-MESSPERIODE (queues Q2 and Q4 except for time for I/Os)

DURATION (S)

Total dwell time (in seconds) during voluntary active wait states in the task monitoring cycle.

DISK IO WAITS

Number of wait states for I/O operations on disk devices in the task monitoring cycle (queues Q4 and Q12).

DURATION (S)

Total dwell time (in seconds) for I/O operations to disk devices in the task monitoring cycle.

NON DISK IO WAITS

Number of wait states for I/O operations during the wait states in the task monitoring cycle (queues Q4 and Q12).

DURATION (S)

Total dwell time (in seconds) during voluntary inactive wait states in the task monitoring cycle except for boursing.

INACTIVE WAITS

Number of voluntary inactive wait states in the task monitoring cycle except for boursing (queues Q10, Q11, Q13).

DURATION (S)

Total dwell time (in seconds) during voluntary inactive wait states in the task monitoring cycle except for boursing.

ADMISSIONS

Number of admissions in the task monitoring cycle (sum of queues Q5 and Q6).

DURATION (S)

Wait time for admission in seconds (sum of queues Q5 and Q6).

BOURSE LONG WAITS

Number of voluntary inactive wait states in the task monitoring cycle due to boursing (queue 12).

DURATION (S)

Total dwell time (in seconds) during voluntary inactive wait states in the task monitoring cycle due to boursing.

SERVICE UNITS

Service units which have been included in the task monitoring cycle.

CPU SERVICE UNITS

CPU service units which have been included in the task monitoring cycle.

IO SERVICE UNITS

I/O service units which have been included in the task monitoring cycle.

MEMORY SERVICE UNITS

Memory service units which have been included in the task monitoring cycle.

For the following four items, the name of the most frequently used program is output. The program name is specified as follows:

- the file name, in the case of linked programs (START-EXECUTABLE-PROGRAM FROM-FILE= <filename>). If the file name including catalog ID and user ID is longer than 20 characters, the catalog and user IDs are omitted from the file name and only the first 20 characters are output.
- the module name, in the case of programs started via DLL (START-EXECUTABLE-PROGRAM FROM-FILE=*MODULE(LIBRARY=*STD,ELEMENT=...))
- EAM OMF in the case of programs started from the EAM library (START-EXECUTABLE-PROGRAM FROM-FILE=*MODULE(LIBRARY=*OMF, ELEMENT=...))
- *NONE, if no program was executed within the task during the monitoring cycle.

HIGHEST CPU CONSUMER

The program with the highest CPU consumption (in seconds).

HIGHEST IO CONSUMER

The program with the most I/O operations.

HIGHEST S-U CONSUMER

The program with the highest service unit consumption.

HIGHEST WS CONSUMER

The program with the highest working set consumption.

HIGHEST PRIORITY

The highest priority assigned for the task.

LOWEST PRIORITY

The lowest priority assigned for the task.

INPUT MESSAGES/SEC

Number of input messages per second for this task. Contrary to the RESPONSETIME monitoring program, messages between applications within a host are also included.

AVG INP.-LGTH (BYTES)

Average input length in bytes.

OUTPUT MESSAGES/SEC

Number of output messages per second from this task. Contrary to the RESPONSETIME monitoring program, messages between applications within a host are also included.

AVG OUTP.-LGTH (BYTES)

Average output length in bytes.

PAMPAGES PER DISK IO

Average number of PAM pages per I/O operation on disk devices during the task monitoring cycle.

KB PER NON DISK IO

Average amount of data transported in Kb per I/O output operation on other devices during the task monitoring cycle.

3. Values supplied if LEVEL=*HIGH

In this case the same values as for LEVEL=*LOW are supplied together with those described below.

Two values, subdivided on DVS input/output) and paging (DVS or PAGING) are supplied for each monitored device specified using a mnemonic device name. The mnemonic device name and the pubset ID are output.

This data is also supplied, broken down as to PAGING and DMS values, as total values for all devices.

ACCESSES PER SEC

Number of accesses (EXCP calls) per second and per device

ACCESSES TOTAL

Number of accesses (EXCP calls) per device

SERVICETIME HW (MS)

Average hardware service time per device.

SERVICETIME SW (MS)

Average software service time per device. (including hardware service time)

In the case of paging I/Os, only input/output operations are assigned to the task; these operations relate to the paging error interrupt caused by the task, i.e. in the case of read chains, these I/O operations are assigned to the relevant task (the counter is incremented) and the hardware time is included in the calculations.

For parallel access volumes (PAV), the devices specified for the monitoring program definition are output. The data of the basic devices refers to the basic device and the assigned alias devices.

Write-only inputs/outputs are assigned to the PGE task (see [“assignment of paging activities to the initiating task” on page 661](#)).

4. Data supplied in addition for classes (user ID, TSN, categories, job name, job class).

TASKS

Number of tasks in this class.

AVERAGE MPL

Average MPL (Multi Programming Level).

This value is computed from the elapsed time for all tasks in the class, divided by the time from START-MP to STOP-MP. Inaccuracies can occur as a result of the fact that ELAPSED TIME and START/STOP-MP time are measured in different units.



The number of device accesses per second is computed from the total number of IOs per device, divided by the time from START-MEASUREMENT-PROGRAM to STOP-MEASUREMENT-PROGRAM or from task creation at time of task terminating.

The monitoring data # ADMISSIONS and DURATION (S) has only values different from zero if the PCS was in operation for part of the time or the entire time. For the interpretation of the data see the "PCS" manual [12].

For SPOOL tasks the accumulated CPU time and the number of I/O operations (TOTAL CPU-TIME(S) or TOTAL # IO'S) may be incorrect because the system counters monitored by SM2 are reset to zero by SPOOL each time a job terminates.

If a SPOOL task is sensed at STOP-MEASUREMENT-PROGRAM time and a SPOOL job is in progress at this time, the TSN and user ID of the initiator are output.

SET-EVALUATION-PERIOD

Set analysis period and analysis subinterval

The SET-EVALUATION-PERIOD statement defines the analysis period and analysis subinterval of an analysis run. The analysis period (evaluation period) is the period during which the SM2 output file is to be monitored.

The analysis subinterval defines the division within the analysis period over which the individual values for the bar chart output are to be averaged.

Within an SM2R1 run different SET-EVALUATION-PERIOD statements can be issued for different analyses.

Format

SET-EVALUATION-PERIOD

```
EVALUATION-PERIOD = *DAY(...) / *MONTH(...) / *YEAR(...) / PERIOD(...)
```

```
*DAY(...)
```

```
  | DAY = EARLIEST / *CURRENT / <date>
```

```
*MONTH(...)
```

```
  | YEAR = *THIS-YEAR / <integer 0..2059>
```

```
  | ,MONTH = LAST-MONTH / <integer 1..12>
```

```
*YEAR(...)
```

```
  | YEAR = LAST-YEAR / <integer 0..2059>
```

```
PERIOD(...)
```

```
  | FROM = *START / *DATE(...)
```

```
    *DATE(...)
```

```
      | DAY = EARLIEST / *CURRENT / <date>
```

```
      | ,TIME = 00:00:00 / <time>
```

```
  | ,TO = *STOP / *DATE(...)
```

```
    *DATE(...)
```

```
      | DAY = START-DATE / <date>
```

```
      | ,TIME = 23:59:59 / <time>
```

```
  | ,TIME-STEPS = *STD / <integer 1..1000>(...
```

```
    <integer 1..1000>(...
```

```
      | DIMENSION = MINUTES / *SECONDS / *DAYS / *HOURS / *NO-OF-STEPS
```

General time specifications

time	Time in the form hh:mm:ss. The seconds need not be specified; analysis then begins and ends at the minutes boundary.
date	The date can be specified in the form yy-mm-dd or yyyy-mm-dd.

Operands**EVALUATION-PERIOD =**

Defines the time frame covered by the analysis period and the analysis subinterval.

EVALUATION-PERIOD = *DAY(...)

The analysis period covers the specified day. SM2R1 sets the length of a subinterval to 15 minutes.

DAY =

Defines the day to be analyzed.

DAY = *EARLIEST

The analysis period covers the first day in the SM2 output file.

DAY = *CURRENT

The analysis period covers the current day from 00:00:00 to 23:59:59.

DAY = <date>

The analysis periods covers the date specified as 00:00:00 to 23:59:59.

EVALUATION-PERIOD = *MONTH(...)

The analysis period covers the current month. SM2R1 sets the length of a interval to 8 hours.

YEAR =

Defines the year during which a month is to be analyzed.

YEAR = *THIS-YEAR

A month of the current year is to be analyzed.

YEAR = <integer 0..2059>

A month in the specified year is to be analyzed.

MONTH =

Defines the month of the specified year.

MONTH = *LAST-MONTH

The analysis period is the month before the current month of the year defined.

MONTH = <integer 1..12>

The analysis period is the specified month from 00:00:00 of the first day to 23:59:59 of the last day.

EVALUATION-PERIOD = *YEAR(...)

The analysis period covers the specified year. SM2R1 sets the length of a subinterval to 7 days. In general the week and year boundaries are not identical. The period actually analyzed is therefore somewhat longer than the specified year.

YEAR =

Defines the year in which a month is to be analyzed.

YEAR = *LAST-YEAR

The analysis period is the previous year.

YEAR = <integer 0..2059>

The analysis period is the specified year.

EVALUATION-PERIOD = *PERIOD(...)

The analysis period is specified explicitly.

FROM =

Defines the start of the analysis period.

FROM = *START

The analysis period starts at the beginning of the SM2 output file.

FROM = *DATE(...)

Defines the start of the analysis period by specifying the date and time of day.

DAY =

Defines the start date of the analysis period.

DAY = *EARLIEST

The start date is the first day in the SM2 output file.

DAY = *CURRENT

The start date is the current day.

DAY = <date>

The start date is the specified day.

TIME =

Defines the time the analysis period starts.

TIME = *00:00:00

Analysis begins at 00:00:00 of the specified start date.

TIME = <time>

Analysis begins at the specified time on the specified start date.

TO =

Defines the time at which the analysis period ends.

TO = *STOP

The analysis period stops at the end of the SM2 output file.

TO = *DATE(...)

Defines the end of the analysis period by specifying the date and time of day.

DAY =

Defines the end of the analysis period.

DAY = *START-DATE

The end date of the analysis period is the same as the start date.

DAY = <date>

The end date of the analysis period is the specified day.

TIME =

Defines the time at which the analysis period ends.

TIME = *23:59:59

The analysis ends at 23:59:59 on the specified end date.

TIME = <time>

The analysis ends at the specified time on the specified end date.

TIME-STEPS =

Defines the size of the analysis subintervals.

TIME-STEPS = *STD

The entire analysis period is subdivided into 100 subintervals. This corresponds to single-page bar chart output.

TIME-STEPS = <integer 1..1000>(...)

The user can define the extent of an analysis subinterval either directly by specifying a number of time units (e.g. TIME-STEPS=5(DIM=*MINUTES)) or indirectly by specifying a specific number of analysis subintervals, based on which the entire analysis period is to be subdivided. (TIME-STEPS=50(DIM=*NO-OF-STEPS)).

DIMENSION =

Defines the time unit of the analysis subinterval.

DIMENSION = *MINUTES

Time unit in minutes.

DIMENSION = *SECONDS

Time unit in seconds.

DIMENSION = *DAYS

Time unit in days.

DIMENSION = *HOURS

Time unit in hours.

DIMENSION = *NO-OF-STEPS

Denotes the subdivision of the analysis period into the specified number of subintervals.

Example

```
//SET-EVALUATION-PERIOD EVALUATION-PERIOD=*DAY(*CURRENT)
```

This statement designates the current day as the analysis period.



If the selected subintervals are too small (TIME-STEPS.. less than the SM2 monitoring cycle), blanks will appear in the chart instead of a bar.

Up to 100 analyzed subintervals can be represented graphically on one printed page. The statistical data of the same 100 subintervals is output in table form on the same page. If the entire analysis period comprises more than 100 subintervals, report output is spread over several pages, but the statistical data refers to the associated chart and not to the whole analysis period.

SET-EXCEPTION-PERIOD

Exclude periods

The SET-EXCEPTION-PERIOD statement offers the user a convenient option for excluding either fixed periods or certain recurring periods such as weekdays or weekends. The statement thus especially supports long-term analysis with SM2R1.

The statement takes effect during execution of the PRINT-QUEUE-TRANSITION, PRINT-REPORTS, PRINT-SUMMARY and START-AUTOMATIC-ANALYSIS statements.

Format

SET-EXCEPTION-PERIOD
<pre> EXCEPTION = <u>*NONE</u> / list-poss(3): *PERIOD(...) *PERIOD(...) FROM = *DATE (...) *DATE(...) DAY = <date with-compl> ,TIME = <u>*00:00:00</u> / <time> ,TO = *DATE (...) *DATE(...) DAY = <u>*START-DATE</u> / <date with-compl> ,TIME = <u>*23:59:59</u> / <time> ,REPEAT-EXCEPTION = <u>*NONE</u> / *PERIOD(...) *PERIOD(...) FROM = *DEFINED (...) *DEFINED(...) DAY-OF-THE-WEEK = *MONDAY / *TUESDAY / *WEDNESDAY / *THURSDAY / *FRIDAY / *SATURDAY / *SUNDAY ,TIME = <u>*00:00:00</u> / <time> ,TO = *DEFINED (...) *DEFINED(...) DAY-OF-THE-WEEK = <u>*SAME-DAY</u> / *MONDAY / *TUESDAY / *WEDNESDAY / *THURSDAY / *FRIDAY / *SATURDAY / *SUNDAY ,TIME = <u>*23:59:59</u> / <time> </pre>

General time specifications

time	Time in the form hh:mm:ss. The seconds need not be specified; analysis then begins and ends at the minutes boundary.
date	The data can be specified in the form yy-mm-dd or in the form yyyy-mm-dd.

Operands**EXCEPTION =**

The specified period is excluded from analysis.

EXCEPTION = *NONE

No period is specified.

EXCEPTION = list-poss(3): *PERIOD(...)

The exclusion period comprises the period explicitly specified.

FROM = *DATE(...)

Defines the start of the exclusion period by specifying the day and time.

DAY = <date>

The specified day is the start date of the exclusion period.

TIME =

Specifies the start time of the exclusion period.

TIME = *00:00:00

Exclusion starts at 00:00:00 on the defined start date.

TIME = <time>

Exclusion starts at the specified time on the specified start date.

TO = *DATE(...)

Defines the end of the exclusion period by specifying the day and time.

DAY =

Specifies the end date of the exclusion period.

DAY = *START-DATE

The end date of the exclusion period is equal to the start date.

DAY = <date>

The specified day is the end date of the exclusion period.

TIME =

Defines the end time of the exclusion period.

TIME = 23:59:59

Exclusion ends at 23:59:59 of the defined end date.

TIME = <time>

Exclusion ends at the specified time on the define end date.

REPEAT-EXCEPTION =

The specified period is periodically excluded from analysis.

REPEAT-EXCEPTION = *NONE

No period is specified.

REPEAT-EXCEPTION = *PERIOD(...)

The exclusion period comprises the period explicitly specified.

FROM = *DEFINED(...)

Defines the start of the exclusion period by specifying a day of the week and a time.

DAY-OF-THE-WEEK = *MONDAY / ... / *SUNDAY

The start day of the exclusion period is the specified day of the week.

TIME =

Defines the start time of the exclusion period.

TIME = *00:00:00

Exclusion begins at 00:00:00 of the specified day of the week.

TIME = <time>

Exclusion begins at the specified time of the specified day of the week.

TO = *DEFINED(...)

Defines the end of the exclusion period for the specified day of the week and time.

DAY-OF-THE-WEEK =

Defines the end of the exclusion period.

DAY-OF-THE-WEEK = *SAME-DAY

The end date of the exclusion period is the same day it began.

DAY-OF-THE-WEEK = *MONDAY / ... / *SUNDAY

The end date of the exclusion period is the specified day of the week.

TIME =

Defines the end time of the exclusion period.

TIME = *23:59:59

Exclusion ends at 23:59:59 of the specified day of the week.

TIME = <time>

Exclusion ends at the specified time of the specified day of the week.

Example

```
//SET-EXCEPTION-PERIOD REPEAT-EXCEPTION=*PERIOD(FROM=*DEFINED(*FRIDAY,  
16:00:00),TO=*DEFINED(*MONDAY,08:00:00))
```

This statement excludes weekends from long-term analysis.

SET-REPORT-FOCUS

Select time slots from analysis period

The SET-REPORT-FOCUS statement enables the user to select time slots of a (larger) analysis period for analysis (PRINT-REPORTS, PRINT-SUMMARY, START-AUTOMATIC-ANALYSIS and PRINT-QUEUE-TRANSITION statements).

Only the data used in the slots is used for analysis. If the analysis period comprises several days, the selected time slots are analyzed for each of these days.

Up to 3 slots can be specified, in ascending order of time, but they must not be contiguous.




Figure 11: Time slots within the analysis period

Format

SET-REPORT-FOCUS(SET-WINDOW)

FOCUS(WINDOW) = *NONE / list-poss(3): *DEFINED(...)

*DEFINED(...)

 | FROM = <time>

 | ,TO = <time>

Operands**FOCUS(WINDOW) =**

Defines time slots.

FOCUS(WINDOW) = *NONE

No time slots are output.

FOCUS(WINDOW) = list-poss(3): *DEFINED(...)

Defines up to 3 time slots.

***DEFINED(...)**

Defines the start and end times of a slot.

FROM = <time>

Specifies the start time of the slot.

TO = <time>

Specifies the end time of the same slot.



All specified times must be within one calendar day (00:00:00 to 23:59:59).

SET-TITLE

Output title

The title specified by the user is output on every page of the analysis. The SET-TITLE statement triggers output of the title.

Format

SET-TITLE
TEXT = <c-string 1..100>

Operands

TEXT = <c-string 1..100>

Any text enclosed in apostrophes can be specified. An apostrophe in the print format must be denoted by two apostrophes in the input text. The text must not exceed 100 characters, excluding the apostrophes.

Example

```
//SET-TITLE TEXT='SM2 MANUAL EXAMPLE'
```

START-AUTOMATIC-ANALYSIS

Start automatic bottleneck analysis

This statement enables the user to start automatic analysis of bottlenecks for the selected time interval (see [“Automatic performance analysis” on page 394](#)). For a complete analysis, monitored data from the monitoring programs SERVICETIME, SYSTEM, TASK and (if available) VM must be additionally available for the relevant analysis period. SM2R1 outputs the results of the analysis to the selected output medium.

Format

START-AUTOMATIC-ANALYSIS

```

MAIN-APPLICATION = *TP / *DIALOG / *BATCH
,OUTPUT-MEDIUM = *BOTH / *SYSLST
,IMPORTANT-CATEGORIES = *STD / list-poss(16): <alphanum-name 1..7>

```

Operands

MAIN-APPLICATION = *TP / *DIALOG / *BATCH

Specifies the main application mode of the system: transaction processing (TP), dialog mode (DIALOG) or batch processing (BATCH).

OUTPUT-MEDIUM = *BOTH / *SYSLST

Specifies the output medium SYSLST or the system files SYSLST and SYSOUT (*BOTH).

IMPORTANT-CATEGORIES =

Specifies the names of the important categories.

IMPORTANT-CATEGORIES = *STD

The SM2R1 program selects the important categories itself. The criterion for the importance of a category is its weight. The third of all categories with the greatest weight are selected.

IMPORTANT-CATEGORIES = list-poss(16): <alphanum-name 1..7>

Lists the names of the important categories.

If there are no important categories, the following is output:

```
IMPORTANT CATEGORIES:
    *NONE
```

This happens when the corresponding record does not exist or the user specifies categories that do not exist in the output file.

Example

```
//START-AUTOMATIC-ANALYSIS MAIN-APPLICATION=*DIALOG
```

This statement starts automatic analysis for a system with dialog mode as the main application.

Conventions and results

For the basic principles behind automatic performance analysis, refer to the “Performance Handbook” [5].

The rules governing the detection of resource bottlenecks can be subdivided into the groups CPU, IO, Paging and VM2000.

Example

```
START OF AUTOMATIC ANALYSIS
MAIN APPLICATION: TP
IMPORTANT CATEGORIES:
    BATCHDB
    TP
    DIALOG2

!!! PARAMS UNBALANCED CATEGORY,
    - CATEGORY TP      WAITS FOR CPU LONGER THAN CATEGORY ...
    - PRIORITIES ARE NOT BALANCED.
    - RAISE PRIORITY FOR CATEGORY.

!!! CPU BOTTLENECK, TP
    - WAIT TIME FOR CPU EXCEEDS RECOMMENDED LIMIT FOR CATEGORY TP
    - TOO HEAVY CPU USAGE OF CATEGORY. SYSTEM PARAMETERS ARE BALANCED.
    - INCREASE CPU SPEED.

!!! PUBLIC DISK OVER UTILIZED, (A832,20S7.2)
    - UTILIZATION OF PUBLIC DEVICE (A832,20S7.2) EXCEEDS RECOMMENDED LIMIT.
    - TOO MANY USERS FOR SAME DISK.
    - TAKE CARE FOR UNIFORM UTILIZATION OF PUBLIC VOLUME SET.

END OF AUTOMATIC ANALYSIS
```

For each bottleneck diagnosed, three records are output to SYSLST and optionally output to SYSOUT:

Record 1 includes the description of the bottleneck, record 2 lists the possible reasons, while record 3 gives instructions on how to tune the system.

In addition, it should be noted that the SM2 values are checked against other configuration-specific performance data, having previously determined the SM2 values for the set time period. For this reason automatic analysis is not performed over the entire time period specified, otherwise the probability of any peak times being corrupted by any underload times would be too great; rather, automatic analysis is only performed for time periods which have already been recognized as being of critical importance.

If no potential bottleneck is indicated, this does not imply that no bottleneck actually exists, rather it merely indicates that SM2R1 was not able to identify any bottleneck using the resources available to it. If the system response nevertheless remains unacceptable, a systems analyst or technician must be called in to analyze system performances. Nevertheless, the SM2 system does allow the user to avoid needless routine and drudgery.

9.6 SM2R1 analysis routine examples

All of the examples provided for the SM2R1 analysis routine presuppose specification of the input and output files either in the START-SM2R1 command or assignment via file link name.

The following examples are also in the SYSLIB.SM2.170 library.

Example SM2R1.REPORTS.1

In this example, a header (SET-TITLE) is specified and system configuration output (PRINT-CONFIGURATION) is requested. Analysis begins on 08.12.09 at 12.40 and ends on the same day at 14.30. The analysis period is subdivided into subintervals with a uniform duration of 60 seconds each. Charts and statistics are output for report groups CPU, IO and STORAGE-SYSTEM-SYMMETRIX.

Response time reports are to be analyzed only for specific connection groups already provided by SM2.

```

/START-SM2R1 MONITOR-FILE-NAME = SM2.SAM.MANUAL.V17.0A
  SET-TITLE TEXT = 'SM2R1 MANUAL EXAMPLE'
  SET-EVALUATION-PERIOD EVALUATION-PERIOD = *PERIOD (          -
                                FROM = *DATE (DAY=08-12-09,   TIME=12:40:00), -
                                TO   = *DATE (DAY=*START-DATE, TIME=14:30:00), -
                                TIME-STEPS = 60 (DIMENSION=*SECONDS))
PRINT-CONFIGURATION
PRINT-REPORTS INFORMATION = *DIAGRAMS,          -
  REPORT-LIST = (*CPU,*IO,*STORAGE-SYSTEM-SYMMETRIX)
END

```


Example SM2R1.REPORTS.2

The analysis period is divided into subintervals with a uniform duration of 5 minutes each. Analysis begins on 08.12.09 and ends on the same day. SM2 data from 0:00 to 13:00 and from 14:00 to 23:59:59 is masked out, so that only statistics averaged in the time slot (over 5 minutes) are output. All variables reports to which monitored data can be assigned are output. No charts are output.

In addition, the SM2 output file is analyzed from the beginning of the file till 08.12.09 (23:59:59). Bar charts and statistics are output for the report groups CPU, IO and CHANNEL. The values are averaged over 2 minutes.

```

/START-SM2R1 MONITOR-FILE-NAME = SM2.SAM.MANUAL.V17.0
  SET-TITLE TEXT = 'SM2R1 MANUAL EXAMPLE'
  SET-EVALUATION-PERIOD EVALUATION-PERIOD = *PERIOD(   -
      FROM = *DATE (DAY=08-12-09),   -
      TO   = *DATE (DAY=08-12-09),   -
      TIME-STEPS = 5 (DIMENSION = *MINUTES))
  SET-REPORT-FOCUS FOCUS = *DEFINED (FROM=13:00:00, TO=14:00:00)
  PRINT-REPORTS INFORMATION = *STATISTICS,   -
      REPORT-LIST = *ALL

  SET-EVALUATION-PERIOD EVALUATION-PERIOD= *PERIOD(   -
      FROM = *START,   -
      TO   = *DATE (DAY=08-12-09),   -
      TIME-STEPS = 2 (DIMENSION = *MINUTES))
  PRINT-REPORTS INFORMATION = *DIAGRAMS,   -
      REPORT-LIST = (*CPU, *IO, *CHANNEL)

END

```

Example SM2R1.TASK.STATISTICS

In this example, a header (SET-TITLE) is specified. Analysis extends over the whole SM2 output file. The tasks whose TSNs are specified under TSN= are analyzed with the maximum amount of detail.

```

/START-SM2R1 .....
  SET-TITLE TEXT='SM2R1 MANUAL EXAMPLE'
  PRINT-TASK-STATISTICS INFORMATION=*HIGH,TSN=(1M0Q,1M55,1PAN)
END

```

Example SM2R1.CREATE.USERID.SET

In this example, a header (SET-TITLE) is specified. Analysis extends over the whole SM2 output file. The data of all the tasks specified via the user IDs given under USERID-SET is grouped in classes ABC and XY and output only once per class with the maximum amount of detail. No separate reports are output for the individual tasks. The remaining tasks are grouped and output in a remainder class. Task queue statistics are also output.

```
/START-SM2R1 .....
  SET-TITLE TEXT='SM2R1 MANUAL EXAMPLE'
  CREATE-USERID-SET SET-NAME=ABC,USER-ID=(ABCSP00L,ABCDIAG,ABCKLI)
  CREATE-USERID-SET SET-NAME=XY,USER-ID=(XY333,XYSWN,XY2631V8)
  PRINT-TASK-STATISTICS INFORMATION=*HIGH,TSN=*NONE,USERID-SET=(ABC,XY)
  PRINT-QUEUE-TRANSITION
  END
```

Example SM2R1.SUMMARY

In this example, a header (SET-TITLE) is specified. Analysis extends over the period of august 2008. A summary ACTIVITY report is output every day.

```
/START-SM2R1 .....
  SET-TITLE TEXT = 'SM2R1 MANUAL EXAMPLE'
  SET-EVALUATION-PERIOD EVALUATION-PERIOD = *MONTH (YEAR=08, MONTH=8)
  PRINT-SUMMARY PARTITION          = 1 (DIMENSION = *DAYS), -
                                SUMMARY-REPORTS = *ACTIVITY
  END
```

Example SM2R1.ALL

```
/BEGIN-PROCEDURE LOGGING=A
/ASSIGN-SYSDTA TO-FILE=*SYSCMD

/CREATE-FILE FILE-NAME = SM2R1.OUT.MANUAL, -
/          SUPPORT = PUBLIC-DISK (SPACE = *RELATIVE( -
/          PRIMARY-ALLOCATION = 576, -
/          SECONDARY-ALLOCATION = 576)) ----- (1)

/START-SM2R1 MONITOR-FILE-NAME = SM2.SAM.MANUAL.V17.0A, -
/          LIST-FILE-NAME = SM2R1.OUT.MANUAL, -
/          EVALUATION-FILE-NAME = SM2R1.OUT.EVALDATA ----- (2)
  SET-TITLE TEXT = '***** SM2R1 - EXAMPLE *****' ----- (3)
  SET-EVALUATION-PERIOD EVALUATION-PERIOD = *PERIOD (
                                FROM = *DATE (DAY=2008-12-09,TIME=12:36:00),
                                TO = *DATE (DAY=2008-12-09,TIME=14:38:00),
                                TIME-STEPS = 60 (DIMENSION = *SECONDS)) ----- (4)
  PRINT-CONFIGURATION ----- (5)
  PRINT-SUMMARY SUMMARY-REPORTS = (*ACTIVITY, *PCS) ----- (6)
  PRINT-QUEUE-TRANSITION ----- (7)
```

```

PRINT-TASK-STATISTICS INFORMATION = *HIGH,           -
                                TSN                 = *NONE ----- (8)

MODIFY-REPORT-CONDITIONS REPORT-NUMBER = 10,         -
                          ITEM-NUMBER  = 1,         -
                          UPPER-LIMIT  = 10,        -
                          COMPARE-VALUE = *MEAN-VALUE ----- (9)

MODIFY-REPORT-CONDITIONS REPORT-NUMBER = 10,         -
                          ITEM-NUMBER  = 2,         -
                          UPPER-LIMIT  = 10,        -
                          COMPARE-VALUE = *MEAN-VALUE

MODIFY-REPORT-CONDITIONS REPORT-NUMBER = 102,        -
                          ITEM-NUMBER  = 1,         -
                          UPPER-LIMIT  = 80,        -
                          COMPARE-VALUE = *MEAN-VALUE

MODIFY-REPORT-CONDITIONS REPORT-NUMBER = 102,        -
                          ITEM-NUMBER  = 2,         -
                          UPPER-LIMIT  = 80,        -
                          COMPARE-VALUE = *MEAN-VALUE ----- (10)

MODIFY-REPORT-CONDITIONS REPORT-NUMBER = 102,        -
                          ITEM-NUMBER  = 3,         -
                          UPPER-LIMIT  = 80,        -
                          COMPARE-VALUE = *MEAN-VALUE

PRINT-REPORTS INFORMATION = *DIAGRAMS,           -
                          REPORT-LIST = (*CPU (REPORT-NUMBER = (*1,*137)) -
                                          *IO,           -
                                          *DISK (SPECIFIED = *DEVICE( -
                                                  DEVICE=D192), -
                                                  REPORT-NUMBER = (*124,*125,*127),-
                                                  CONDITIONED-REPORTS = *NO), -
                                          *DEVICE (DEVICE = *SPECIFIED( -
                                                  DEVICE = D192)), -
                                          *CHANNEL) ----- (11)

START-AUTOMATIC-ANALYSIS MAIN-APPLICATION = *DIALOG, -
                          OUTPUT-MEDIUM  = SYSLST ----- (12)

END

/PRINT-DOCUMENT FROM-FILE      = SM2R1.OUT.MANUAL, -
/                               DOCUMENT-FORMAT = *TEXT( -
/                               LINE-SPACING=*BY-EBCDIC-CONTROL) ----- (13)
/END-PROCEDURE

```

- (1) Create the SM2 output file (SM2R1 list file).
- (2) Start SM2R1 and assign the input and output files.
- (3) Output a header for all pages of the SM2R1 analysis report.
- (4) Define the analysis period and analysis subinterval:
analysis dated <date> from 12:36:00 to 14:38:00 in subintervals of 60 seconds.
- (5) Output category-specific queue statistics.
- (6) Output a SUMMARY for the most important global monitoring data (ACTIVITY) and category-specific monitoring data regarding the performance control system (PCS).
- (7) Output category-specific queue statistics.
- (8) Output task statistics (summarized by categories) and a hit list of the TSNs with the greatest consumption.
- (9) Modify threshold values for channel statistics. Report 10 is only to be output if the average values in the analysis period are greater than 10% for “CPU OVERLAP STATE FOR CHANNEL” (ITEM-NUMBER =1) or for “BUSY STATE FOR CHANNEL” (ITEM-NUMBER=2).
- (10) Modify threshold values for channel statistics. Report 102 is only to be output if the average values in the analysis period are greater than 80 I/Os for “IO’S (BY PAM-TRANSFER) FOR CHANNEL” (ITEM-NUMBER=1), “IO’S (BY BYTE-TRANSFER) FOR CHANNEL” (ITEM-NUMBER=2) or for “IO’S (BY NODATA-TRANSFER) FOR CHANNEL” (ITEM-NUMBER=3).
- (11) Output monitoring data in chart form and as statistical values for the report groups *CPU, *IO, *DISK, *DEVICE and *CHANNEL. Report 1 (CPU load) and report 137 (TU and TPR SVCs per second) are to be output for *CPU. In the case of *DISK, only reports 124, 125 and 127 are to be output for DEVICE=D192 and the default threshold values should not be taken into account. Only device D192 is to be taken into account for *DEVICE. The threshold values as modified under points 9 and 10 are valid for *CHANNEL.
- (12) Start automatic bottleneck analysis with output of results to the SM2R1 report file (main application type is dialog mode).
- (13) Print results.



The default threshold values or modified threshold values used for suppressing reports are not usually into account when the monitored data is output to the transfer file or when PRINT-SUMMARY is specified.

***** SM2R1 - EXAMPLE *****

SM2 SUMMARY ACTIVITY REPORT 1

REPORTING PERIOD

=====

FROM : <date> , 12:36:02
 TO : 14:36:01
 BY : 0:01:00

SYSTEM DATA

=====

SYSTEM NAME : I11BXS
 MACHINE TYPE : S180-30
 SESSION # : 192
 BS2000 VERSION : 180
 HOST : D016ZE04
 MAIN MEMORY [KB]: 3932164
 CLASS 1 MEMORY [KB]: 4124
 CLASS 2 MEMORY [KB]: 10720

TASK AND RESPONSE TIME SUMMARY

=====

#TASKS

	AVG	MIN	MAX	SDEV
SYS	: 169.256	168.700	170.400	0.448
BATCH	: 108.200	102.000	117.600	4.918
DIALOG	: 167.334	153.000	175.800	5.785
TP	: 57.999	57.900	58.100	0.025

QUEUE STATISTICS [#TASKS]

	AVG	MIN	MAX	SDEV
IN CPU QUEUE	: 1.638	0.200	5.800	1.178
IN PAGING QUEUE	: 0	0	0	0
IN IO QUEUE	: 0.193	0.000	0.600	0.164
ACTIVE	: 150.185	140.000	158.200	4.817
INACT READY	: 0	0	0	0
NOT ADMITTED	: 0	0	0	0
INACT NOT READY	: 352.603	340.900	362.100	5.830
TOTAL	: 502.789	490.000	506.900	3.332

RESPONSE TIME STATISTICS

		AVG	MIN	MAX	SDEV
MEAN THINK TIME	[S] :	0.556	0.118	6.652	1.140
MEAN RESPONSE TIME (1)	[S] :	0.202	0.035	2.182	0.386
MEAN RESPONSE TIME (2)	[S] :				
MEAN TRANSACTION TIME	[S] :	0.471	0.088	2.852	0.738
TRANSACTION RATE	[1/S]:	104.349	19.264	282.774	68.113
RESPONSE RATE (1)	[1/S]:	104.485	19.438	282.884	68.081
RESPONSE RATE (2)	[1/S]:				

***** SM2R1 - EXAMPLE *****

SM2 SUMMARY ACTIVITY REPORT 2

SYSTEM STATISTICS

=====

CPU STATISTICS [NORMED %]

		AVG	MIN	MAX	SDEV
TU TIME	:	7.133	0.706	17.283	4.176
TPR TIME	:	22.172	7.103	39.664	6.190
SIH TIME	:	15.582	1.526	27.749	8.892
IDLE TIME	:	55.114	26.969	88.905	14.450
STOP TIME	:	0	0	0	0
ACTIVE LOGICAL MACHINES	:	2.000	2.000	2.000	0.000

IO STATISTICS [1/S]

		AVG	MIN	MAX	SDEV
NON PAGING DISK IO'S	:	294.006	69.638	922.578	179.042
PAGING IO'S	:	0.003	0.000	0.017	0.005
TAPE IO'S	:	3.053	0.798	5.148	1.734
PRINTER IO'S	:	0	0	0	0
OTHER	:	1954.298	124.230	3793.191	1288.260

MEMORY STATISTICS

	AVG	MIN	MAX	SDEV
#CLASS 3 PAGES :	14440	14220	14678	119
#CLASS 4 PAGES :	45645	45602	46001	79
AVAILABLE PAGES (NPP) :	821291	820540	822536	570
WSET ACT TASKS (PPC) :	26562	25556	27446	557
WSET INACT READY TASKS (PPC) :	0	0	0	0
USED PAGES ACT TASKS :	366867	358103	376332	5640
USED PAGES INACT READY TASKS :	0.132	0.000	3.000	0.545
(WSET ACT + INACT READY) / NPP :	0.032			
TOTAL # PAGE FAULTS [1/S]:	716.030	103.488	1789.060	397.869
# PAGE RECLAIMS [1/S]:	0.273	0.000	7.737	1.636
# FIRST PAGE ACCESS [1/S]:	715.754	103.488	1783.566	397.292
# PAGE WRITES TO DISK [1/S]:	0.001	0.000	0.012	0.003
# PAGE READS FROM DISK [1/S]:	0.002	0.000	0.017	0.004

PAGINGAREA STATISTICS

	AVG	MIN	MAX	SDEV
PAGES ON PAGING DEVICE(S) :	3840006	3840006	3840006	0
USED PAGES ON PAGING DEVICE(S):	1074711	846253	1879596	386794
PAGES ON ES / GS :	0	0	0	0

***** SM2R1 - EXAMPLE *****

SM2 SUMMARY ACTIVITY REPORT 3

CHANNEL STATISTICS

=====

CHANNEL BUSY STATE [%]

	AVG	MIN	MAX	SDEV
CHANNEL PATH ID: 00CE (TYP S) :	21.167	0.166	81.818	20.121
CHANNEL PATH ID: 004E (TYP S) :	20.812	0.166	87.879	20.330
CHANNEL PATH ID: 000D (TYP S) :	20.639	0.166	84.848	20.389
CHANNEL PATH ID: 0018 (TYP FC) :	1.495	0.000	4.160	0.865
CHANNEL PATH ID: 0058 (TYP FC) :	1.478	0.000	4.160	0.849
CHANNEL PATH ID: 0098 (TYP FC) :	1.478	0.000	4.160	0.857
CHANNEL PATH ID: 00D8 (TYP FC) :	1.473	0.000	4.160	0.857
CHANNEL PATH ID: 00C8 (TYP S) :	0.670	0.000	8.472	1.732
CHANNEL PATH ID: 0048 (TYP S) :	0.604	0.000	6.478	1.491
CHANNEL PATH ID: 0009 (TYP S) :	0.172	0.000	1.827	0.333

CHANNEL IO RATE [1/S]

	AVG	MIN	MAX	SDEV
CHANNEL PATH ID: 00CE (TYP S) :	0.002	0.000	0.120	0.022
CHANNEL PATH ID: 004E (TYP S) :	0.002	0.000	0.040	0.009
CHANNEL PATH ID: 000D (TYP S) :	0.002	0.000	0.160	0.029
CHANNEL PATH ID: 0018 (TYP FC) :	73.693	17.475	230.922	44.784
CHANNEL PATH ID: 0058 (TYP FC) :	73.181	17.113	230.153	44.739
CHANNEL PATH ID: 0098 (TYP FC) :	73.567	17.508	230.494	44.747
CHANNEL PATH ID: 00D8 (TYP FC) :	73.230	17.296	230.644	44.802
CHANNEL PATH ID: 00C8 (TYP S) :	0.157	0.087	0.474	0.097
CHANNEL PATH ID: 0048 (TYP S) :	0.173	0.116	0.521	0.108
CHANNEL PATH ID: 0009 (TYP S) :				

CHANNEL PAM PAGES [1/S]

	AVG	MIN	MAX	SDEV
CHANNEL PATH ID: 00CE (TYP S) :	0.001	0.000	0.080	0.014
CHANNEL PATH ID: 004E (TYP S) :	0.001	0.000	0.040	0.008
CHANNEL PATH ID: 000D (TYP S) :	0.002	0.000	0.120	0.022
CHANNEL PATH ID: 0018 (TYP FC) :	270.177	30.922	778.061	147.763
CHANNEL PATH ID: 0058 (TYP FC) :	268.779	30.390	774.729	147.475
CHANNEL PATH ID: 0098 (TYP FC) :	269.036	31.559	772.905	147.903
CHANNEL PATH ID: 00D8 (TYP FC) :	268.295	30.889	765.916	147.185
CHANNEL PATH ID: 00C8 (TYP S) :	0.272	0.125	0.762	0.123
CHANNEL PATH ID: 0048 (TYP S) :	0.333	0.216	0.698	0.104
CHANNEL PATH ID: 0009 (TYP S) :				

***** SM2R1 - EXAMPLE *****

SM2 SUMMARY ACTIVITY REPORT 4

DEVICE STATISTICS

=====

DEVICE BUSY STATE (NON PAGING) [%]

	AVG	MIN	MAX	SDEV
D186 20DS.0 :	3.246	0.000	44.925	7.891
D189 20DS.3 :	3.219	0.000	37.752	6.724
D156 WORK01 :	2.769	0.000	9.983	3.476
D1FB PUB002 :	2.673	0.000	57.807	10.059
D336 :	1.489	0.000	4.464	0.711
D32B 2RZV.B :	1.085	0.000	3.328	1.177
B336 :	1.068	0.000	2.899	0.998
B337 :	0.958	0.000	5.824	1.285
D1FC PUB003 :	0.958	0.000	16.162	3.130
B3A3 6VS1.1 :	0.919	0.000	3.968	0.940

DEVICE BUSY STATE (PAGING) [%]

-----				AVG	MIN	MAX	SDEV	---
D186	20DS.0	:		0	0	0	0	0
D189	20DS.3	:		0	0	0	0	0
D156	WORK01	:		0	0	0	0	0
D1FB	PUB002	:		0	0	0	0	0
D336		:		0	0	0	0	0
D32B	2RZV.B	:		0	0	0	0	0
B336		:		0	0	0	0	0
B337		:		0	0	0	0	0
D1FC	PUB003	:		0	0	0	0	0
B3A3	6VS1.1	:		0	0	0	0	0

DEVICE IO RATE [1/S]

-----				AVG	MIN	MAX	SDEV	---
D186	20DS.0	:		29.353	0.000	411.786	72.042	
D189	20DS.3	:		37.029	0.000	377.904	82.775	
D156	WORK01	:		30.777	0.000	122.808	37.459	
D1FB	PUB002	:		24.476	0.000	592.332	102.823	
D336		:		11.332	0.120	14.911	2.234	
D32B	2RZV.B	:		12.583	0.000	42.473	14.107	
B336		:		7.243	0.120	15.589	5.493	
B337		:		5.361	0.000	12.712	5.761	
D1FC	PUB003	:		2.651	0.000	26.563	6.878	
B3A3	6VS1.1	:		6.466	0.882	14.606	3.993	

DEVICE DATA RATE [KB/S]

-----				AVG	MIN	MAX	SDEV	---
D186	20DS.0	:		117.397	0.000	1647.127	288.171	
D189	20DS.3	:		135.432	0.000	1237.548	250.348	
D156	WORK01	:		113.828	0.000	452.796	139.827	
D1FB	PUB002	:		169.605	0.000	2258.543	436.049	
D336		:		8.151	0.160	10.845	1.588	
D32B	2RZV.B	:		347.183	0.000	1183.730	394.708	
B336		:		5.412	0.160	10.845	3.349	
B337		:		3.675	0.000	8.253	3.928	
D1FC	PUB003	:		162.673	0.000	3000.107	568.334	
B3A3	6VS1.1	:		27.664	2.246	76.189	23.557	

***** SM2R1 - EXAMPLE *****

*CPU: UTILIZATION NORMED (REPORT 1)


ALL PROCESSORS

D016ZE04

ACCUMULATED

EVALUATION FROM <date>, 12:36:02 TO 14:12:01 BY 0:01:00

ID - MEASURED TERM , UNIT IS PERCENT	----- AVG	---- MIN	---- MAX	---- SDEV	-- CNT
1 = TU TIME	6.466	0.706	17.283	3.886	27
. = TPR TIME	23.336	16.955	39.664	5.264	27
3 = SIH TIME	18.529	2.846	27.749	7.168	27
= IDLE TIME	51.669	26.969	75.907	12.538	27
- = STOP TIME	0	0	0	0	27



9.7 SM2R1 transfer file records

The detailed description of the records provides the following information on every record field:

- contents of the record field
- data format: binary, floating or alphanumeric representation
- length of the field (in bytes)
- field displacement, counting from the beginning of the data area; for fields within a monitored object group or repeat group, counting starts from the beginning of the group (marked by a “+”).

A 4-byte control field (record length field), whose first two bytes contain in binary form the length of the record, including the control field, always comes before the actual data area of each record.

TIM2 record

The TIM2 record contains the data from the TIME record (Version 10.0). The only difference is that dates and times are output in ISO4 format. The date and time are separated by a “T”.

Displacement	Length	Format	Meaning
0	4	character	TIM2 identification
4	19	character	Start of the analysis period: (= 1st time stamp in the period) in the format yyyy-mm-ddThh:mm:ss
23	19	character	End of the analysis period: (= last timestamp in the period) in the format: yyyy-mm-ddThh:mm:ss
42	4	binary	Length of a subinterval: Number of days
46	4	binary	Number of seconds
50	19	character	Start of grid subdivision in the format: yyyy-mm-ddThh:mm:ss
69	2	binary	Number of selected time slots
from 71 to 111	8 8	character character	Description of time slot (3 times): Start of the time slot in the format: hh:mm:ss End of the time slot in the format: hh:mm:ss
119	2	binary	Number of excluded time periods

Displacement	Length	Format	Meaning
from121	19	character	Description of excluded time periods (3 times): Start of the exclusion period in the format: yyyy-mm-ddThh:mm:ss
to 216	19	character	End of the exclusion period in the format: yyyy-mm-ddThh:mm:ss
235	2	binary	Description of a regularly excluded time period (once only) Start of the exclusion period: Day of week (*)
237	8	character	Time in the format hh:mm:ss
245	2	binary	End of the exclusion period: Day of week (**)
247	8	character	Time in the format hh:mm:ss
255	2	binary	Number of entries per DATA record

- (*) The day of the week can have the values 0 (= no regular exclusion period) and 1 through 7 (1 being Monday, 7 being Sunday).
- (**) The day of the week can have the values 1 through 7 (one being Monday, 7 being Sunday), and also the values 8 (Monday) through 13 (Saturday). The values 8 through 13 are required if the excluded time period selected extends over a weekend.

SYST record

A “T” is used to separate the date and time when they are represented in ISO4 format.

Displacement	Length	Format	Meaning
0	4	character	SYST identifier
4	2	binary	Point in time at which this system information was obtained Year Month Day Hour Minute Second
6	2	binary	
8	2	binary	
10	2	binary	
12	2	binary	
14	2	binary	
16	8	character	Name of BS2000 system
24	3	character	Operating system version number
27	8	character	Creation date of the generated system in the form yy/mm/dd
35	8	character	SM2 version designation
43	8	character	SM2R1 version designation
51	8	character	MTFILE version designation
59	4	binary	Size of main memory (in 4-Kb pages, minus 1) Configuration name Size of virtual class 1 memory (4 Kb) Size of virtual class 2 memory (4 Kb) Size of expanded storage/global storage (4 Kb)
63	21	character	
84	4	binary	
88	4	binary	
92	4	binary	
96	4	binary	Task address space (in 1-Kb pages)
100	2	binary	Length of a monitored object group Displacement between the first monitored object group and the beginning of the record Number of monitored object groups
102	2	binary	
104	2	binary	
106	19	character	Time this system information was obtained in the format: yyyy-mm-ddThh:mm:ss
125	10	character	Creation date for the generated system in the format: yyyy-mm-dd
135	8	character	Host name

The identification is supplied for each processor in the monitored object group.

Monitored object group:

Displacement	Length	Format	Meaning
+0	4	binary	Processor ID reserved for the processor
+4	4	–	

CONF records

Each CONF record consists of the record header and at least one repeat group. The repeat groups can be global information, channel, controller or device. The entry under the “record identifier” header indicates the repeat group in question.

If all CONF records are consecutively read to memory and the header is omitted, the individual repeat groups can be located with the aid of the displacement values and the number given in the repeat group entitled “global information”.

The “device” repeat group can only occur once in a CONF record, followed by at least one of its path repeat groups.

A set of CONF records might then have the following appearance:

Record header	global		
Record header	CHA	CHA	CHA
Record header	CTL	CTL	CTL
Record header	DEV	Path repeat group	
Record header	DEV	Path repeat group	Path repeat group
Record header	DEV	Path repeat group	
Record header	DEV	Path repeat group	

CONF record

Record header:

Displacement	Length	Format	Meaning
0	4	character	CONF' identification
4	2	binary	Length of a repeat group entry
6	2	binary	Displacement between the first repeat group and the beginning of the record
8	2	binary	Number of repeat groups in this record
10	1	binary	Record identifier X' 00' global information X' 02' channel record X' 03' controller record X' 04' device record

Record identifier X'00' – global information

Repeat group:

Displacement	Length	Format	Meaning
+0	2	binary	Length
+2	2	–	Not used
+4	8	character	Identification (C'\$DSTATUS')
+12	8	–	Not used
+20	4	binary	Displacement for the first channel information value
+24	4	binary	Number of channel information values
+28	4	binary	Displacement for the first controller information value
+32	4	binary	Number of controller information values
+36	4	binary	Displacement for the first device information value
+40	4	binary	Number of device info. values
+44	4	–	Not used

Record identifier X'02' – channel

Repeat group:

Displacement	Length	Format	Meaning
+0	2	binary	Length
+2	2	–	Not used
+4	4	character	Identification (C'CHN')
+8	4	–	Reserved
+12	1	binary	Channel type
+13	1	binary	Channel path ID
+14	4	character	IO-SIDE number or X'FFFFFFFF'

Record identifier X'03' – controller

Repeat group:

Displacement	Length	Format	Meaning
+0	2	binary	Length
+2	2	–	Not used
+4	4	character	Identification (C'CTL')
+8	4	character	Controller mnemonic
+12	1	binary	Controller type
+13	1	binary	Difference in controller type

Record identifier X'04' – device

Repeat group:

Displacement	Length	Format	Meaning
+0	2	binary	Length
+2	2	–	Not used
+4	4	character	Identification (C'DEV')
+8	4	character	Device mnemonic
+12	1	binary	TSOS device type
+13	4	–	Reserved
+17	1	binary	Device information: X'01': unit record device X'02': disk device X'04': tape device X'10': paging device X'20': public device X'40': device switchable X'80': shared private disk device
+18	4	–	Reserved
+22	6	character	VSN
+28	4	character	TSN of the device user task
+32	4	character	TSN of the device owner task
+36	2	binary	Length of the path repeat group
+38	2	binary	Displacement for the first path repeat group
+40	2	binary	Number of path repeat groups

Path repeat group:

Displacement	Length	Format	Meaning
+0	1	binary	Channel path ID
+1	1	binary	CTL#/DEV# or controller and device address
+2	1	binary	Display of path availability
+3	1	–	Reserved
+4	4	character	Controller mnemonic
+8	4	–	Reserved
+12	4	character	X'FFFFFFFF'

DSCR record

Displacement	Length	Format	Meaning
0	4	character	DSCR identification
4	2	binary	Report number
6	2	binary	Variable number (1 – 5)
8	40	character	Description of monitoring value
48	2	binary	Length of monitoring group description (l1)
50	(l1)	character	Description of monitoring group

DATA record

Displacement	Length	Format	Meaning
0	4	character	DATA identification
4	2	binary	Report number
6	2	binary	Monitoring value number (1 – 5)
8	4	float	Average value over the entire analysis period
12	4	float	Maximum value over the entire analysis period
16	4	float	Minimum value over the entire analysis period
20	4	float	Standard deviation over the entire analysis period
24	4	float	Number of subintervals in the analysis period, including monitored values
from 28	4 each	float	Values of individual subintervals

As of a displacement of 28, the DATA contains the values of the individual subintervals.

The total number of data entries is contained in the TIM2 record.

The number of data entries not equal to 0 is contained in the DATA record. If this value is 0, no values follow.

10 Other analysis routines

The web-based user interface of openSM2 (openSM2 Manager) and the programs ANALYZER and INSPECTOR contain a convenient online help system. For a description of the product SM2-PA, see the “SM2-PA” manual [15].

10.1 openSM2 Manager

The openSM2 Manager is the web-based user interface for performance monitoring of SE servers.

The openSM2 Manager is available as add-on software in the SE Manager. It runs on the Management Unit and enables central monitoring on the Server Units /390 and x86, the Application Units (x86), of storage systems and all devices with SNMP capability.

The openSM2 Manager is operated using a browser, either locally on the Management Unit or on a remote operator terminal. You will find the requirements for the browser in the description of the SE Manager in the “Operation and Administration” manual [18].

10.1.1 Calling the openSM2 Manager

- ▶ Start the SE Manager and log in.

Working with the SE server is described in the online help of the SE Manager and in the “Operation and Administration” manual [18].

- ▶ In the main window of the SE Manager select the *Performance* menu.



To obtain access to the openSM2 Manager, you must be a member of one of the following user groups:

Administrator or Operator user group of the SE Manager or a user group which is permitted to access the openSM2 Manager.

Otherwise access is denied.

- ▶ Following successful login, the main window of the openSM2 Manager is displayed.

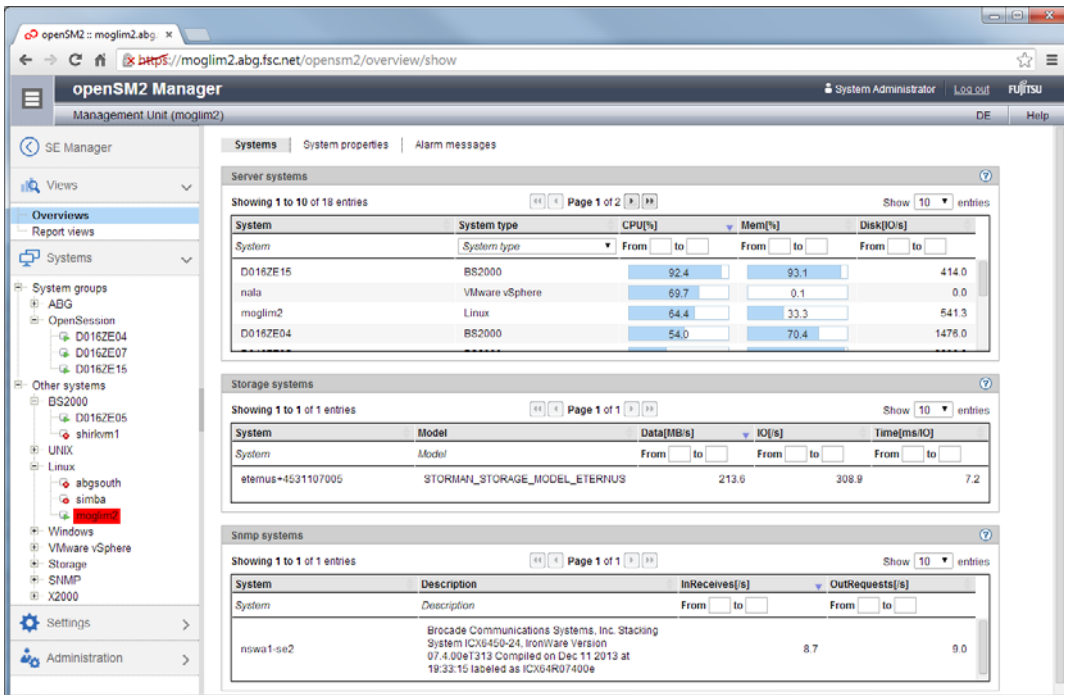


Bild 12: openSM2 Manager: Overviews

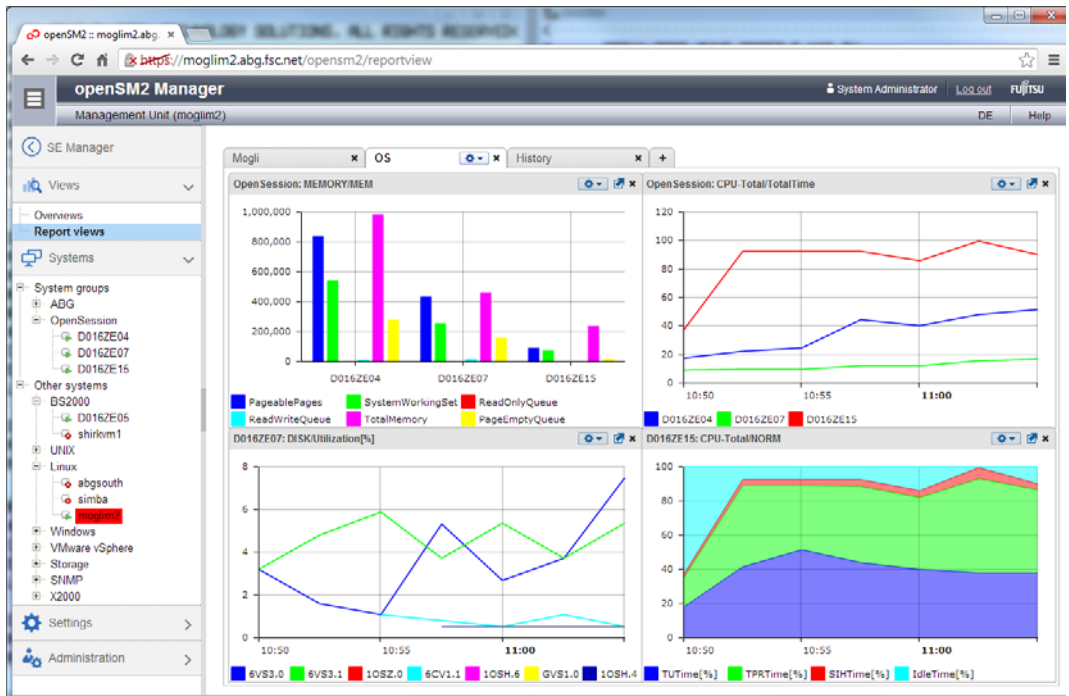


Bild 13: openSM2 Manager: Report views

10.1.2 Working with the openSM2 Manager

- At the top of the tree structure of the main window you will find the entry < **SE Manager**, which enables you to return to the SE Manager.

When you return to the SE Manager, you are taken to the last screen you selected in the SE Manager.

- The openSM2 Manager always uses the language setting and the session timeout value of the SE Manager. You can modify the language setting in the openSM2 Manager.
- When you log off from the openSM2 Manager, you also log off implicitly from the SE Manager.
- Further information and help on working with the openSM2 Manager is provided in the online help of the openSM2 Manager.

10.1.3 Functions of the openSM2 Manager

The functional scope of the openSM2 Manager is largely the same as that of the INSPECTOR program. In addition, the openSM2 Manager offers the following functions:

- Automatic monitoring of all components and systems of an SE server
- Summary presentation of the most important utilization values
- Extension of the system group concept by inhomogeneous and dynamic system groups
- Display of historical data

The openSM2 Manager is used primarily to monitor the components and systems of an SE server. However, the openSM2 Manager also enables you to monitor systems outside the SE server provided these have been configured as required.

The openSM2 Manager supports the server systems BS2000, Solaris, Linux, Microsoft Windows, VMware vSphere, Xen and X2000, the ETERNUS DX and Symmetrix storage systems, and all systems with SNMP capability.

At selectable intervals (monitoring intervals) agents collect monitored data on the current status of the monitored systems and store this in a database. The agents run on the Management Unit and collect the data of the monitored systems remotely, which means that openSM2 does not need to be installed on these systems. Only on BS2000 and Solaris systems does an agent have to be installed locally on the monitored system. This agent transfers the monitored data to the master agent on the Management Unit over a TCP/IP connection. Furthermore, if monitoring is also required for Windows systems, a Windows agent must be installed on any selected Windows system.

The monitored data is stored in two databases. One database contains the monitored data from the online monitoring of the last 24 hours. The second (optional) database contains archived monitored data for offline analyses. The scope of archiving (i.e. the systems and monitored variables whose monitored data is to be archived) can be configured.

Furthermore, you can compress the archived monitored data by combining a number of monitoring intervals.

When the SE server is appropriately configured, the components and systems in it are automatically determined by the openSM2 Manager and included in monitoring. The openSM2 administrator must, if necessary, enter authentication data for the systems to permit the agent to establish a connection to the systems. Furthermore, the openSM2 administrator can enter further systems which are to be monitored.

Multiple systems - also with different system types - can be combined to form a system group. For example a system group can be formed with all guest systems (VM2000) on a virtualized server. Dynamic system groups permit systems to be selected and grouped according to specific criteria. The members of these system groups are not permanently defined, but are determined dynamically by filter functions, which means that new systems which are added can automatically be assigned to a system group.

The monitored systems are displayed in a tree structure which makes the assignment of the systems to system platforms and system groups visible. The color of a system entry shows the status of the system.

A summary presentation of the most important utilization values of all monitored systems with filter and sort functions provides a quick overview of the total utilization of the SE server.

Snapshot reports containing the monitored data of the current monitoring intervals and time series reports showing the time period of the monitored data are available to present the monitored data. Either the monitored data of one single system or of a system group can be presented in a report. The reports can be arranged, stored and opened again in various report views which can be freely configured by the user.

User management with a role concept assigns the users different rights. In addition to presenting the monitored data for the monitored systems, all users can define user-specific settings. Administrators can in addition define or modify global configuration settings.

Each user can modify the display names of the systems and define system groups. He/She can modify predefined report groups, reports and monitored variables, and also define new report groups, reports and monitored variables.

The monitored data can be monitored on the basis of user-defined rules. A rule defines conditions and actions. When all the conditions of a rule are satisfied, the defined actions are performed. The alarm status is displayed by the color of the system entry in the system list which is defined in the rule.

An administrator can define which systems are to be monitored and specify settings for the agents. Furthermore, he/she can define the role and thus the rights of other users. He/She can also export, archive and delete monitored data.

10.2 ANALYZER

Analysis with a Windows PC

ANALYZER is the component of openSM2 for analyzing openSM2 output files. The convenient graphical user interface of the manager enables you to specify the type and scope of the analysis. The analysis request so defined is processed by the agent, and the result is then presented by the manager. The output files to be analyzed can be located on any servers on which agents are running.

ANALYZER allows you to generate informative diagrams on the basis of the monitored data with a minimum of effort. The monitored data can thus be interpreted easily, enabling performance bottlenecks to be detected quickly and future performance requirements to be forecast. The wide range of options for creating the diagrams and the ease with which the data can be transferred to other Windows applications also make ANALYZER an indispensable tool for creating monitoring reports.

Various output files – also from different servers – can be analyzed simultaneously in an analysis run. The result of the analysis can be stored in a file to be processed further, and it can also be transferred to standard programs such as Excel and Word or to the clipboard.

The user can use macros in ANALYZER to run defined analyses entirely automatically and create daily monitoring reports, for example, without the need for any manual intervention.

ANALYZER allows the monitored data of selected variables over extended periods of time to be saved in compressed form on the PC in long-term files.

Just like an output file on a server, a long-term file can be analyzed and thus provide the database for trend analyses and long-term capacity planning.

The following analyses can be carried out with ANALYZER:

- diagrams
- automatic analysis
- configuration

Diagrams

The analysis result is supplied in the form of a diagram and a table.

ANALYZER offers many options for creating analyses in accordance with the wishes and requests of the user:

- selection of the analysis period with definition of time windows and exclusion of periods within them (e.g. weekends)
- selection of monitored variables on the basis of report group, report and variable
- combination of several monitored variables by means of formulae
- selection of monitored objects (e.g. devices, channels, categories, etc.)
- creation of a hit list of the monitored objects with the highest values in the case of partially qualified monitored objects
- identifying the processes involved in the case of peak system load
- monitoring of defined service levels through analysis of frequencies (e.g. in what percentage of the monitoring intervals was the response time between 0.5 and 1.0 seconds?)
- revelation of functional relationships between monitored variables through being displayed in a correlation chart
- customized design of graphics through selection of different chart types (bars, lines, logarithmic lines, etc.) and colors

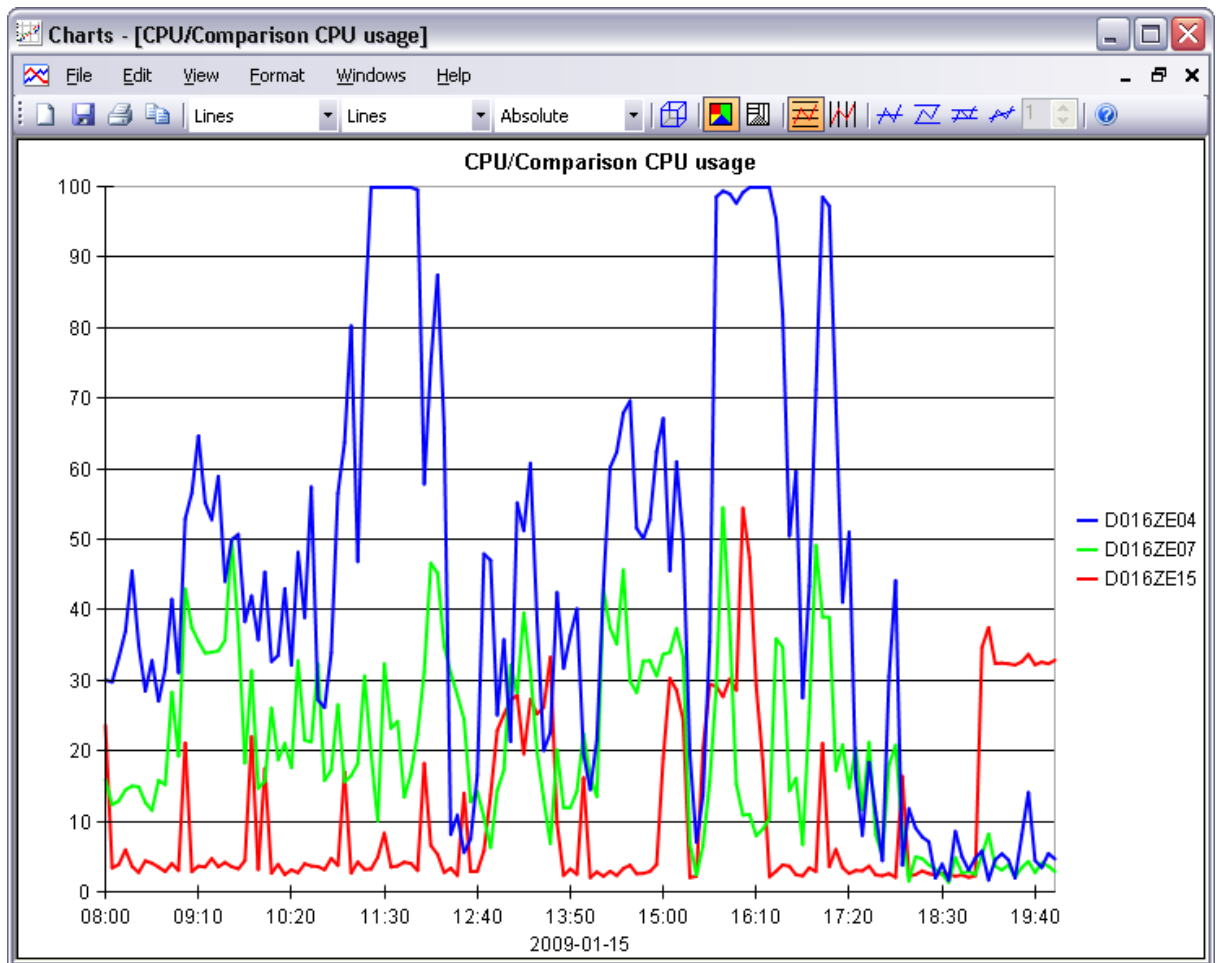


Figure 14: ANALYZER result of the analysis of several output files

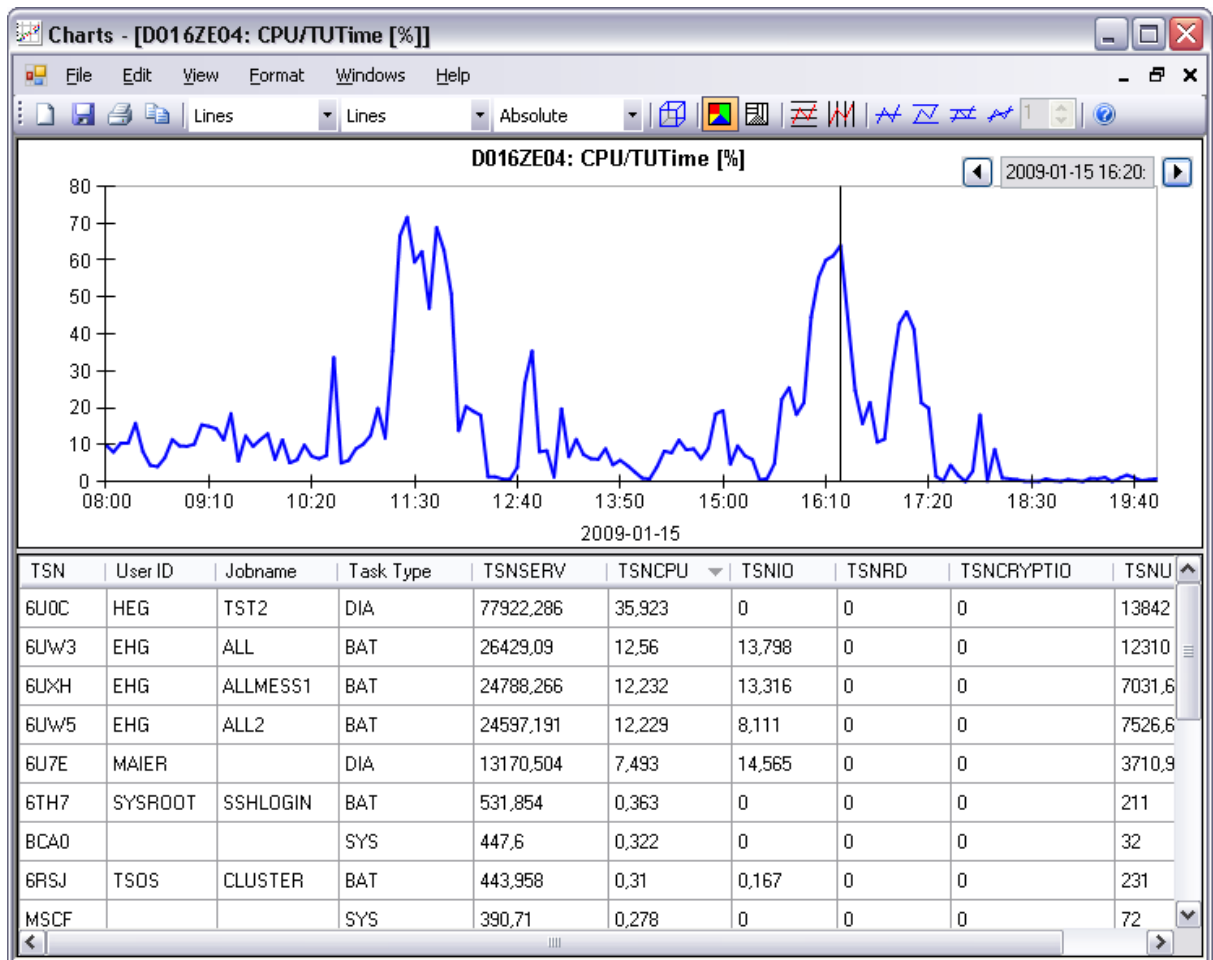


Figure 15: ANALYZER result of the analysis of an output file with a process hit list




Figure 16: ANALYZER status diagram for a VM2000 system

Automatic analysis

Automatic analysis allows program areas in BS2000 systems to be identified quickly and countermeasures to be taken to deal with them.

In bottleneck analysis, monitored data is examined on the basis of specified conditions and rules. If the conditions for a rule are met, this is a sign of a system bottleneck, and a corresponding message is created, indicating how the bottleneck can be eliminated.

The result of the automatic analysis is displayed as a structured list of messages that can be printed and saved.

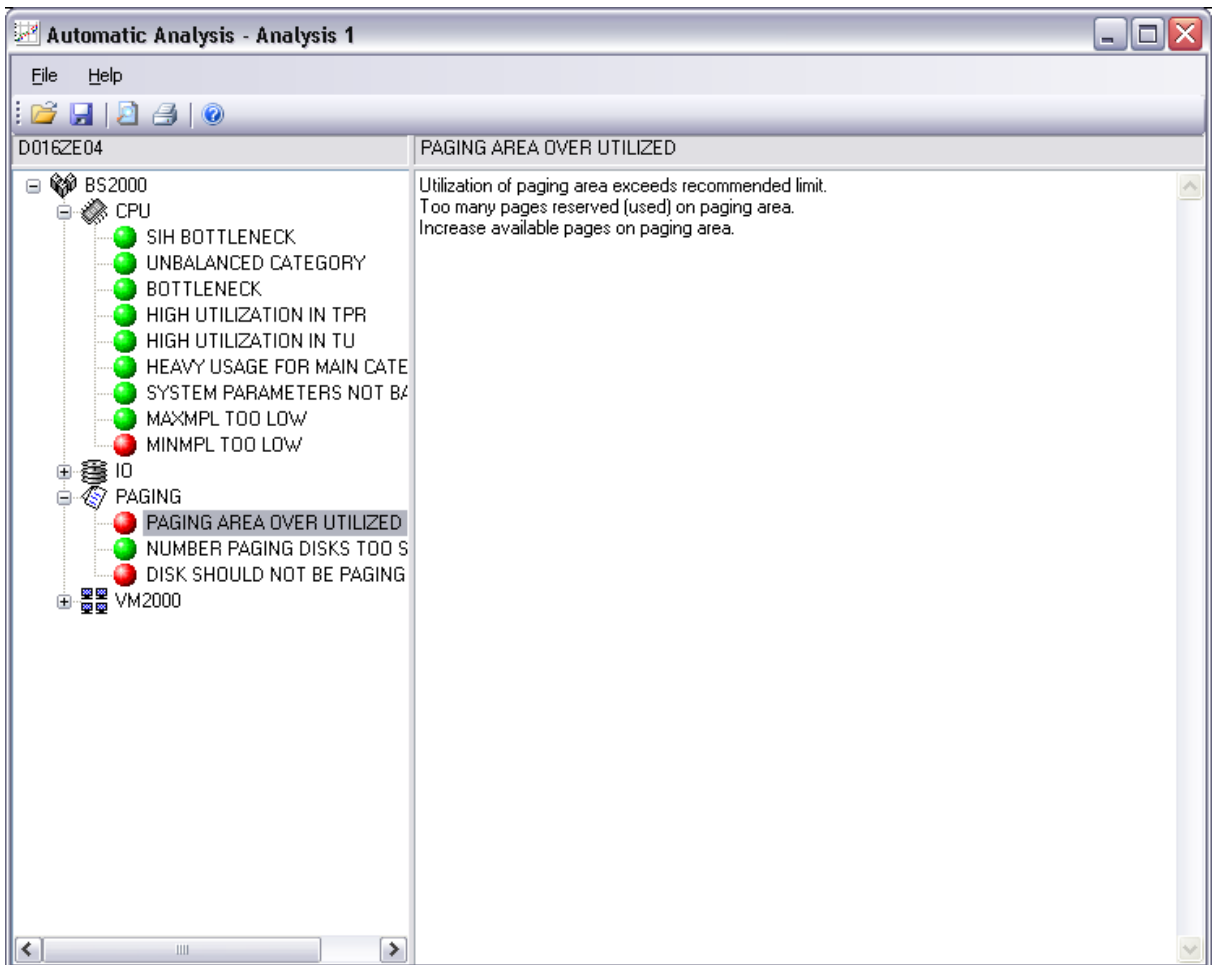


Figure 17: ANALYZER result of an automatic analysis

The structure tree in the left-hand half of the window displays the areas CPU, IO, PAGING and VM2000, which are subdivided into message areas. The message areas for which messages exist are highlighted in red.

The text area in the right-hand half of the window displays information and messages. If a red message area is clicked, the associated messages are displayed.



When Live Migration takes place, a new session with the new server is started in the SM2 output file. An automatic performance analysis across the session boundary leads to separate analyses for the server concerned.

Configuration

The user can display the system configuration of a BS2000 server in a tree structure and thus obtain an overview of the server's peripherals.

The configuration can be saved in a file and read in again at any time for further processing or printed out in clear graphical form.

It is also possible to search the configuration for specific devices. Any of the attributes of the device (such as MN, VSN, device type, etc.) can be used as a search criterion.

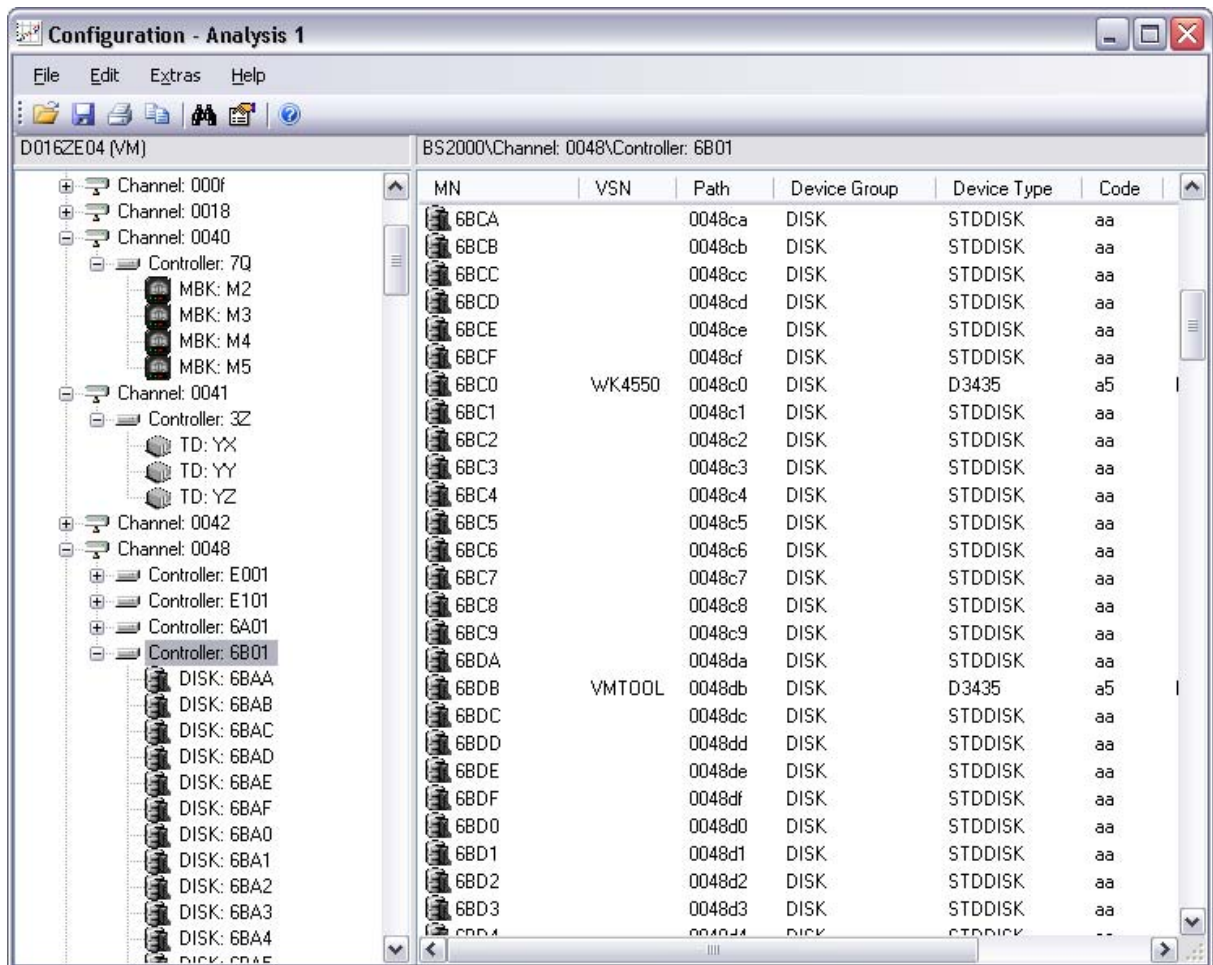


Figure 18: ANALYZER display of a configuration with marked controller 6B01

The list view on the right shows detailed information on all devices that are connected to the selected 6B01 controller.

10.3 INSPECTOR

Online monitoring with a Windows PC

INSPECTOR is the component of openSM2 which enables multiple systems to be monitored simultaneously and the monitored data to be presented graphically.

The clear graphical display and a rule-based check of the monitored data permit efficient monitoring of the systems. The user can trace the monitored data at any point in the procedure and thus assess the behavior of the system and identify problems at an early stage.

The rule-based check of the monitored data supports the extensive automation of monitoring by means of configurable alarm messages and actions which are triggered automatically.

System list and reports

The systems monitored are displayed on the manager in the system list. The system list has a tree structure in which the assignment of the systems to system platforms and system groups can be seen. The color of a system entry shows the status of the system. The following values are possible: "no data", "values in the normal range" and "alarm". The status display is continuously updated and permits exception situations to be detected and reacted to immediately.

If the user wants to follow the development more closely on the hosts, there are various types of report available offering graphical or tabular views of the monitored data.

- Snapshot reports give the current system status. The user has the choice between the graphical display of the monitored data and a display that corresponds to the screen outputs of SM2 in BS2000.
- The time series reports show the monitored data over the last one to four hours.
- System reports can be opened in order to view selected monitored data of a system.
- Global reports offer an overview of the monitored data of several systems in a chart or table.

The monitored variables and objects can be defined whose monitored data is to be included in a report. In addition, new monitored variables can be defined by combining variables in formulae.

The charts in the reports can be customized by selecting different chart types (bars, lines, logarithmic lines, etc.) and colors. You can use the mouse to print the charts, save them to graphics files or copy them to the clipboard in order to transfer them to standard programs (MS Word, MS PowerPoint, etc.). For the printout it is possible to insert a header and footer and modify the labeling of the the charts.

The large number of report types and design options for the charts offer great scope for designing a monitoring screen that suits the requirements of users.



Figure 19: INSPECTOR monitoring screen with different report types

Rule-based monitoring

The monitored data of each monitored variable can be monitored on the basis of user-defined rules.

Conditions and actions are specified in a rule. A condition is defined by a monitored variable and threshold values for the monitored data. If all the conditions for the rule are satisfied, an alarm is triggered for the rule and the actions are executed. The alarm status for the rule is terminated when at least one of the conditions is no longer satisfied.

The user can define different rules for different times of the day in order to take into account the different ways in which the system is used – e.g. interactive mode during the day and batch mode at night.

In the event of an alarm the entry in the system list takes on the color assigned to the rule's alarm level.

Furthermore, actions defined for the rule are performed.

The following actions are possible:

- A report with the monitored variable concerned is opened automatically for each of the rule's conditions.
- An audible alarm is triggered, which is either a beep or a (monitored-variable-specific) sound file, depending on what is selected.
- An Enter job is started on the BS2000 host in order, for example, to start an SM2 monitoring program.
- A batch file is started on the PC.
- The user is notified by e-mail or SMS.
- An SNMP trap is sent to a Management Station.

10.4 SM2-PA Program analyzer

The program analyzer is an analysis routine for user-specific SM2 output files generated by the SM2 monitoring routine.

If permitted to do so by the SM2 administrator, any user can monitor a task using the SM2 monitoring routine by registering it for SM2 user task monitoring. The SM2 monitoring routine records the task-specific statistical values and session-specific monitoring data and writes this information to a user-specific SM2 output file.

SM2-PA analyzes this file. The results are supplied in the form of statistics which inform the user about the task's resource consumption or the performance characteristics of user programs and thus form the basis for tuning measures.

The most important performance statistics of a user task are output during task analysis, e.g. CPU time used, memory allocation, DMS and paging I/Os, the number of SVC calls and wait times.

SVC and program counter statistics can be requested during program analysis. The program counter statuses and SVC calls are supplied for separate modules and freely definable address ranges.

These session-specific statistics help in the more detailed examination of user program characteristics and make it possible to identify program sections which are very frequently executed or use a high level of CPU time.

All the statistics can be output to a file in a form ready for printing.

Example SM2-PA program counter statistics

SM2-PA PCOUNTER STATISTICS (SUMMARY TU EVALUATION)					
PROGRAM	:	PERSM			
PCOUNT AREA	:	*STD	- *STD		
NUMBER OF SAMPLES:		1099688			
					NO. MODULES WITHOUT PCOUNTER: 6969
					SAMPLING INTERVAL (MSEC) : 1
I MODULE	I FROM	I TO	I ABS	I REL(%)	I
I NTIMGTIM	I 00F50000	I 00F549FF	I 836585	I 76.07	I XX
I PEZIX	I 01274138	I 0127721F	I 202567	I 18.42	I XXXXXXXXXXXX
I PERLAZ@	I 3E0D16E8	I 3E0D2F67	I 20132	I 1.83	I X
I PERRS41@	I 3E0F2B28	I 3E0F40AF	I 9061	I 0.82	I
I PERPUV@	I 3E377E20	I 3E37AD77	I 4261	I 0.39	I
I PERAKL@	I 3E0C0000	I 3E0C1367	I 3004	I 0.27	I
I PERUCS@	I 3E3986C0	I 3E39ADF7	I 2685	I 0.24	I
I PERLAZ	I 012ED0F0	I 012ED1E7	I 1354	I 0.12	I
I PERRS23@	I 3E0E9070	I 3E0EA94F	I 887	I 0.08	I
I PERUPER@	I 3E39ADF8	I 3E39DD8F	I 872	I 0.08	I
I PECFI	I 01279E68	I 0127D1AF	I 830	I 0.08	I
I PERRS15@	I 3E0DDE00	I 3E0DF83F	I 818	I 0.07	I
I PERPCH@	I 3E3915E0	I 3E3940C7	I 693	I 0.06	I
I PERRS12@	I 3E0DA660	I 3E0DB19F	I 687	I 0.06	I
I PERDBZ@	I 3E3722A0	I 3E373A8F	I 678	I 0.06	I
I PERRS14@	I 3E0DBFA0	I 3E0DDDF	I 645	I 0.06	I
I PERKDM@	I 3E3F5500	I 3E3F833F	I 627	I 0.06	I
I PECKF@	I 3E3AB000	I 3E3AC517	I 612	I 0.06	I
I PERRF0@	I 3E0D62C0	I 3E0D736F	I 547	I 0.05	I
I PEZDIFP	I 0128BB78	I 01290A9B	I 500	I 0.05	I
I PERDBHV@	I 3E2530D0	I 3E255DC7	I 498	I 0.05	I
I PERSNP@	I 3E380888	I 3E38255F	I 475	I 0.04	I
I PEZDI	I 01283158	I 012858FB	I 471	I 0.04	I
I PERRS19@	I 3E0DF840	I 3E0E1B87	I 441	I 0.04	I

11 SM2 program interfaces

In this chapter, the C interfaces SM2GMS and SM2GDAT and the assembly language interface PFMON are described.

The C interface permits access to all the data of the report screens. PFMON supplies a subset of this data: data on CPU utilization and the number of input/output operations.

New applications should use the C interface, since PFMON will not be supported in the long term.

11.1 C interfaces

The C interfaces allow users to call all the current SM2 monitoring data in their C programs; they provide status information and all the data of the report screens.

In addition, two C macros are made available to the user:

- SM2GMS supplies all the data of the MEASUREMENT STATUS screen.
- SM2GDAT supplies the data of the most recently completed monitoring cycle of all report screens.

11.1.1 The SM2GMS macro

The SM2GMS macro (SM2 Get Measurement Status) supplies all the data of the MEASUREMENT STATUS screen.

The macro stores the data in a C structure of the type SM2GMS_get_measurement_status_md1. This structure is defined by the caller and must be passed as a parameter when the macro is called.

```
-----
#include "FHDR.H"
#include "SM2GMS.H"
#include "SM2RC.H"

SM2GMS( struct SM2GMS_get_measurement_stat_md1 SM2_STATUS, char *host_name );
-----
```

```
struct SM2_GMS_get_measurement_stat_md1 SM2_STATUS;
```

This structure must be defined before the macro is called. The macro stores the return codes and the data of the MEASUREMENT STATUS screen in this structure.

There is a description of the structure in [section “Structures of the SM2GMS macro” on page 557](#).

char *host_name

This parameter, which supports a host network, can be used to call the SM2 data of those hosts to which a connection exists through the MSCF subsystem.

However, this is dependent on the following conditions:

- the MSCF connection type must be CCS and the MSCF partner type must be CCS or XCS
- the SM2 subsystem must be loaded on the remote host
- the SM2 version on the remote host must be higher than or the same as that of the SM2GMS version of the local machine.

The parameter specifies the address of the field which contains the name of the host. If the SM2 data of a remote host is to be called, the host name of that host must be specified in the field in 8 characters (see the “HIPLEX MSCF” manual [8]).

If the SM2 data of the local host is to be called, the 8-byte field is to be initialized with zero (binary) or a blank.

Return code

The return codes are stored in the standard header.

The standard header is a structure of the type ESMFHDR. It can be addressed by means of the “hdr” structure element of the SM2_STATUS structure.

The evaluation of the return codes is described in [section “Evaluating the return codes” on page 540](#).



- SM2GMS is implemented as a macro.
- When using the macro, the ISM2CALL module from the SYSLIB.SM2.<ver> library must be linked to the program (for native code on servers with x86 architecture: SKULIB.SM2.<ver>).
- SM2GMS.H is in the SYSLIB.SM2.<ver> library and can be printed out with the following command:

```
/PRINT-DOCUMENT  
  FROM-FILE=*LIBRARY-ELEMENT(LIBRARY=SYSLIB.SM2.<ver>,  
  ELEMENT=SM2GMS.H,TYPE=S)
```
- There is an example of the use of this macro in [section “Example” on page 544](#).

11.1.2 The SM2GDAT macro

The SM2GDAT macro (SM2 Get Data) supplies all the monitoring data for the most recently completed monitoring cycles; this data is also output in the SM2 report screens. There is an SM2 data buffer for each area (e.g. TIME-IO, DAB, MEMORY). The caller uses a macro call to request all SM2 data buffers or a freely selectable subset of all SM2 data buffers.

The macro copies the requested SM2 data buffers to an output area made available by the caller; the size and address of the output area are passed as parameters to the macro when it is called.

```
-----
#include "FHDR.H"
#include "SM2GDAT.H"
#include "SM2RC.H"

SM2GDAT( struct SM2GDAT_get_data_md1 SM2_DATA, long length_buffer,
         void *buffer_ptr, unsigned long buffer_flags, char *host_name );
-----
```

struct SM2GDAT_get_data_md1 SM2_DATA

This structure of the type SM2GDAT_get_data_md1 must be defined before the macro is called. After the call, the return code and the actually required size of the output area can be queried. There is a description of the structure in [section "Structures of the SM2GMS macro" on page 557](#).

long length_buffer

Size of the output area to which the macro is to copy the SM2 data buffer, specified in multiples of 4 Kb.

void *buffer_ptr

Address of the output area to which the macro is to copy the SM2 data buffer.

unsigned long buffer_flags

Integer variable whose binary value specifies the SM2 data buffers to be copied by the macro to the output area. Symbolic constants are available for specifying the SM2 data buffers; if several SM2 data buffers are to be copied to the output area, the corresponding symbolic constants of the SM2 data buffers must be added ([section "Example" on page 544](#)).

Caution

The same symbolic constant cannot be re-added multiple times. If this were to happen, then the required SM2 data buffers would not be supplied. In particular, no symbolic constants may be added to SM2GDAT_BUFFER_ALL.

The following applies to programs which were generated using an earlier version of the macro:

- The programs remain executable and do not have to be recompiled.
- If they are recompiled then the logical OR operator must be replaced by an addition for the logical constants as otherwise compilation errors will occur.

char *host_name

This parameter, which supports a host network, can be used to call the SM2 data of those hosts to which a connection exists through the MSCF subsystem.

However, this is dependent on the following conditions:

- the MSCF connection type must be CCS and the MSCF partner type must be CCS or XCS
- the SM2 subsystem must be loaded on the remote host
- the SM2 version must be higher than or the same as that of the SM2GDAT version of the local host.

The parameter specifies the address of the field which contains the name of the host. If the SM2 data of a remote host is to be called, the host name of that host must be specified in the field in 8 characters (see the “HIPLEX MSCF” manual [8]).

If the SM2 data of the local host is to be called, the 8-byte field is to be initialized with zero (binary) or a blank.

Names of the symbolic constants

Constant	Meaning
SM2GDAT_BUFFER_ACF	Requests the ACF data buffer
SM2GDAT_BUFFER_ALL	Requests all the SM2 data buffers
SM2GDAT_BUFFER_BASIC	Requests the BASIC data buffer
SM2GDAT_BUFFER_BCAM	Requests the BCAM-CONNECTION data buffer
SM2GDAT_BUFFER_CATEGORY	Requests the CATEGORY data buffer
SM2GDAT_BUFFER_CHANNEL_IO	Requests then CHANNEL-IO data buffer
SM2GDAT_BUFFER_CMS	Requests the CMS data buffer
SM2GDAT_BUFFER_DAB	Requests the DAB data buffer
SM2GDAT_BUFFER_DISK_FILE	Requests the DISK-FILE data buffer
SM2GDAT_BUFFER_DLM	Requests the DLM data buffer
SM2GDAT_BUFFER_FILE	Requests the FILE data buffer
SM2GDAT_BUFFER_GS	Requests the GS data buffer
SM2GDAT_BUFFER_GSVOL	Requests the GSVOL data buffer
SM2GDAT_BUFFER_HSMS	Requests the HSMS data buffer
SM2GDAT_BUFFER_ISAM	Requests the ISAM data buffer
SM2GDAT_BUFFER_ISAM_FILE	Requests the ISAM-FILE data buffer
SM2GDAT_BUFFER_MEMORY	Requests the MEMORY data buffer
SM2GDAT_BUFFER_MSCF	Requests the MSCF data buffer
SM2GDAT_BUFFER_NSM	Requests the NSM data buffer
SM2GDAT_BUFFER_OPENFT	Requests the OPENFT data buffer
SM2GDAT_BUFFER_PERTASK	Requests the PERIODIC-TASK data buffer
SM2GDAT_BUFFER_PFA	Requests the PFA data buffer
SM2GDAT_BUFFER_POSIX	Requests the POSIX data buffer
SM2GDAT_BUFFER_PUBSET	Requests the PUBSET data buffer
SM2GDAT_BUFFER_RUNTIME	Requests the RESPONSETIME data buffer
SM2GDAT_BUFFER_SCHANNEL	Requests the SCHANNEL data buffer
SM2GDAT_BUFFER_SDEVICE	Requests the SDEVICE data buffer
SM2GDAT_BUFFER_SESAM_SQL	Requests the SESAM-SQL data buffer
SM2GDAT_BUFFER_ST_SYSTEM	Requests the STORAGE-SYSTEM data buffer
SM2GDAT_BUFFER_SVC	Requests the SVC data buffer
SM2GDAT_BUFFER_SYSTEM	Requests the SYSTEM data buffer
SM2GDAT_BUFFER_TCP_IP	Requests the TCP_IP data buffer

Constant	Meaning
SM2GDAT_BUFFER_TIME_IO	Requests the TIME_IO data buffer
SM2GDAT_BUFFER_TLM	Requests the TLM data buffer
SM2GDAT_BUFFER_USERFILE	Requests the USERFILE data buffer
SM2GDAT_BUFFER_USERISAM	Requests the USERISAM data buffer
SM2GDAT_BUFFER_UDS_SQL	Requests the UDS-SQL data buffer
SM2GDAT_BUFFER_UTM	Requests the UTM data buffer
SM2GDAT_BUFFER_VM	Requests the VM data buffer
SM2GDAT_BUFFER_VM_CPU_POOL	Requests the VM cpu pool data buffer
SM2GDAT_BUFFER_VM_GROUP	Requests the VM-GROUP data buffer

Return code

The return codes are stored in the standard header.


The standard header is a structure of the type ESMFHDR. It can be addressed by means of the “hdr” structure element of the SM2_DATA structure. The evaluation of the return codes is described in [section “Evaluating the return codes” on page 540](#).

Notes

- SM2GDAT is implemented as a macro.
- When the macro is used, the ISM2CALL module from the SYSLIB.SM2.<ver> library must be linked to the program (for native code on servers with x86 architecture: SKULIB.SM2.<ver>).
- SM2GDAT.H is in the SYSLIB.SM2.<ver> library and can be printed out with the following command:

```
/PRINT-DOCUMENT
  FROM-FILE=*LIBRARY-ELEMENT(LIBRARY=SYSLIB.SM2.<ver>,
  ELEMENT=SM2GDAT.H,TYPE=S)
```

- “length_buffer”
SM2 checks whether the size specified in the “length_buffer” parameter for the output area is large enough for the requested SM2 data buffer (but not whether the caller has really made available an output area – e.g. using “malloc” – of the size specified in “length-buffer”).

- Size of the output area
The size of some data buffers (e.g. the PERIODIC-TASK data buffer) can change dynamically (the number of tasks is not constant). If the output area is not large enough, a corresponding return code is set, and the actually required size is stored in the “length_buffer” structure element of the SM2_DATA structure. In this case, the “length_buffer” parameter must be supplied with the value of the “length_buffer” structure element, enough space must be made available, and the macro call must be repeated with the new “length_buffer” value.
The caller can use this to determine the required size for the output area by calling the macro initially with a value of 0 for the “length_buffer” parameter.
- Making the output area available
The output area for the requested SM2 data buffer must be made available by the caller before calling the macro.
 Before the macro is called with a sufficiently high value in the “length_buffer” parameter, an output area must always be made available that is at least as large as the value specified in the “length_buffer” parameter. When the macro is called in order to determine how large the output area has to be, a value of 0 should always be specified for the “length_buffer” parameter to prevent a situation in which the size specified in the “length_buffer” parameter happens to be large enough but the actually available output area is not.
- Validity of the SM2 data buffers
The data buffers BASIC, TIME_IO, MEMORY, CATEGORY, ACF and SCHANNEL can always be supplied. For all other data buffers, the corresponding monitoring program must be active; whether a monitoring program is active can be checked using the SM2GMS macro. For all requested data buffers, after SM2GDAT has been called and before the data buffer has been evaluated, the “valid” bit should be checked. This is because there are situations in which SM2 supplies the data buffer but it does not contain any valid data. The “valid” bit can be checked with the “state” structure element in the BUFFER HEADER of the data buffer. The BUFFER HEADER is described in [section “Structures of the SM2GDAT macro” on page 558](#), and [section “Example” on page 544](#) contains an example of the checking of the “valid” bit.
- Monitoring data of the data buffers
The [chapter “SM2 screen output” on page 209](#) contains a detailed description of the monitored data supplied in the data buffers.
The screen outputs use the same interface. The monitored data supplied in the data buffers is “raw data” that first has to be converted to units such as percent, accesses per second, and so on. For example, the CPU idle, TU, TPR and SIH times are supplied in units of “0.1 milliseconds”; to convert them to a percentage, they must be divided by the duration of the monitoring cycle (“elapsed_time”) (unit “1/300 seconds!”).
- The [section “Example” on page 544](#) contains an example of the use of the macro.

Structure of the output area

In the output area, the caller receives copies of the requested SM2 data buffers. These copies are obtained from the central data buffer to which the SM2 monitoring task writes the monitoring data at the end of each online cycle (or offline cycle, when no online cycle is defined).

The structure of the output area is shown in the following figure.

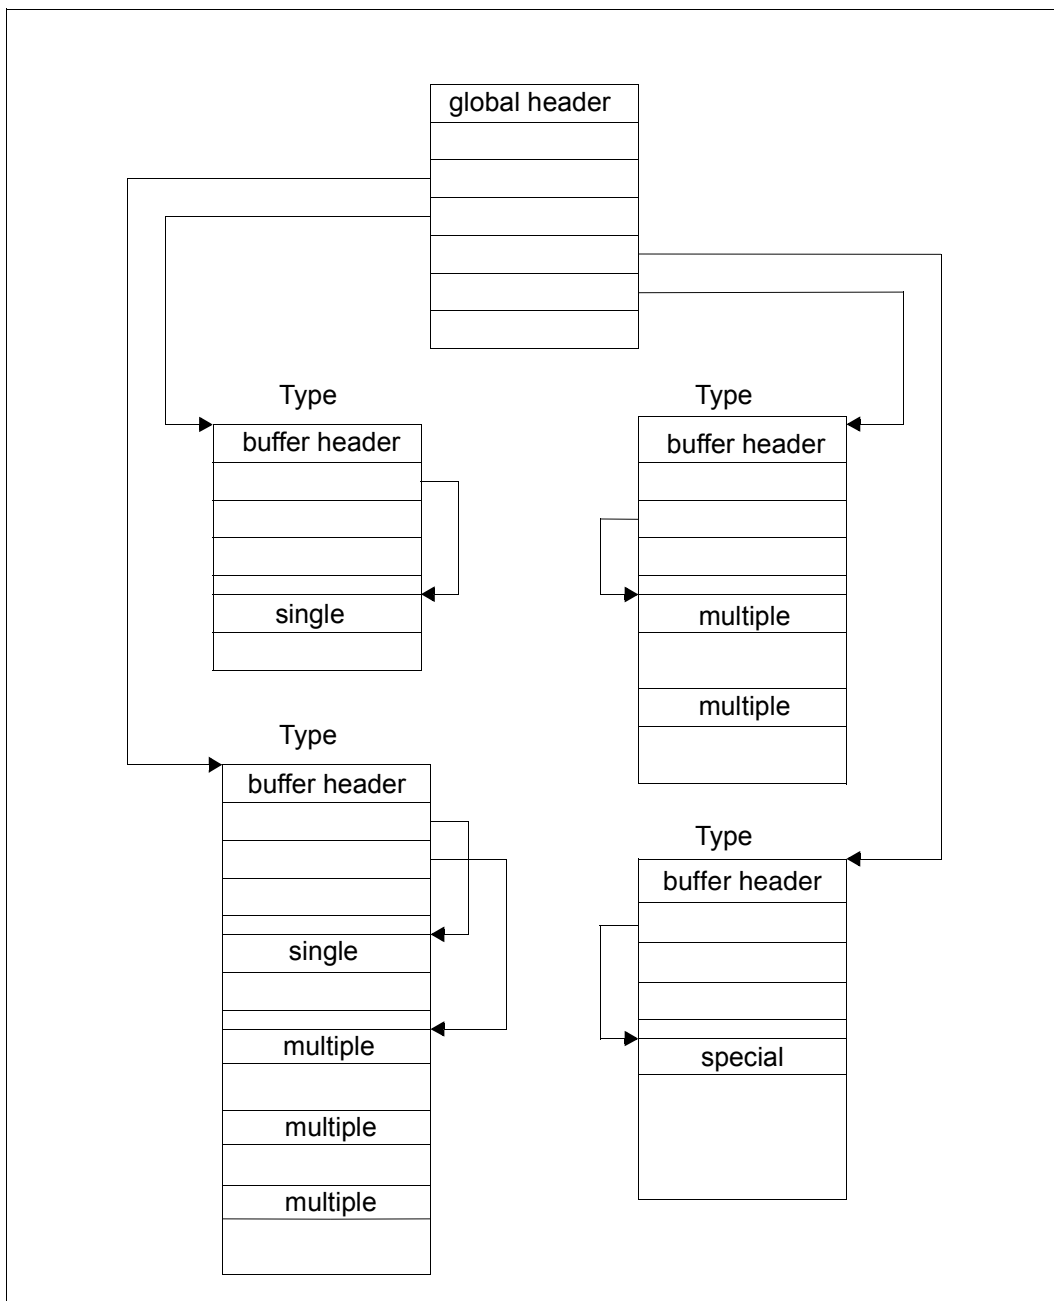


Figure 20: Structure of the output area of the SM2GDAT macro

The output area begins with a global header that contains pointers to the individual data buffers. The “SM2GDAT_global_header_md1” structure describes the global header. The pointers to unrequested data buffers are supplied with NULL.

All data buffers have a standard BUFFER HEADER and a data area. The “SM2GDAT_buffer_header_md1” structure describes the BUFFER HEADER, which contains information on the validity and position of the monitoring data in the data buffer. The structure of the data area depends on the type of the data buffer.

There are four different types of data buffer (see [figure 20](#) above), which are described briefly below; positioning on the data area with the information from the BUFFER HEADER is also explained for the individual types.

In [section “Structures of the SM2GDAT macro” on page 558](#) you will find detailed descriptions of the different data buffers.

Type 1 The data area has a fixed length (referred to as “single” in the diagram).

In the case of this type, the “fixed_part_dsp” element of the BUFFER HEADER supplies the displacement of the data area from the beginning of the BUFFER HEADER.

Data buffers: ACF, BASIC, DLM, HSMS, MEMORY, MSCF, NSM, POSIX

Type 2 The data area consists of a number of repeat groups (referred to as “multiple” in the diagram), which all have the same structure; each repeat group contains the monitoring data for a specific monitored object.

In the case of this type, the “first_group_dsp” element of the BUFFER HEADER supplies the displacement of the first repeat group from the beginning of the BUFFER HEADER. The other repeat groups come after the first repeat group. The “length_group” element supplies the length of a repeat group, and the “number_groups” element supplies the number of repeat groups in the data area. Data buffers: GSVOL, SCHANNEL, SDEVICE, SVC, UTM

Type 3 The data area consists of an area of fixed length (“single”) and a number of repeat groups (“multiple”), which all have the same structure. The area of fixed length contains data that does not refer to a specific object, whereas each repeat group contains the monitoring data for a specific object.

In the case of this type, the “fixed_part_dsp” element of the BUFFER HEADER supplies the displacement of the fixed-length data area from the beginning of the BUFFER HEADER, and the “first_group_dsp” element supplies the displacement of the first repeat group from the beginning of the BUFFER HEADER. The other repeat groups come after the first repeat group. The “length_group” element supplies the length of a repeat group.

In the case of the BCAM-CONNECTION, CMS and GS data buffers, the number of repeat groups is supplied by the “number_groups” element of the BUFFER HEADER and, in the case of the CATEGORY, CHANNEL-IO, DISK-FILE, FILE, ISAM, PERIODIC-TASK, PUBSET, SESAM-SQL, TCP-IP, TIME-IO, TLM, UDS-SQL, USERFILE, USERISAM, VM, VM-CPU-POOL and VM-GROUP data buffers, by the “used_groups” element in the fixed-length data area. In the case of the PFA data buffer, the number_bcb_groups element supplies the number of repeat groups and the first_bcb_group_dsp element the displacement from the first repeat group. The number of “used_groups” repeat groups that are actually filled may be smaller than the number of “number_groups” present.

Data buffers: BCAM-CONNECTION, CATEGORY, CHANNEL-IO, CMS, DISK-FILE, FILE, GS, ISAM, PERIODIC-TASK, PFA, PUBSET, SESAM-SQL, TCP-IP, TIME-IO, TLM, UDS-SQL, USERFILE, USERISAM, VM, VM-CPU-POOL, VM-GROUP

Type 4 The data area has a “special” structure. This structure is described in [section “Structures of the SM2GDAT macro” on page 558](#).

In the case of this type, the “fixed_part_dsp” element of the BUFFER HEADER supplies the displacement of the data area from the beginning of the BUFFER HEADER.

Data buffers: DAB, RESPONSETIME, STORAGE-SYSTEM

The following table provides an overview of all the data buffers and the structures that describe the data area of the data buffer:

Data buffer	Type	Symbolic constant for selecting the data buffer in “buffer_flags”	Structures
ACF	1	SM2GDAT_BUFFER_ACF	SM2GDAT_acf_single_mdl
BASIC	1	SM2GDAT_BUFFER_BASIC	SM2GDAT_basic_single_mdl
BCAM- CONNEC- TION	3	SM2GDAT_BUFFER_BCAM	SM2GDAT_bcam_single_mdl SM2GDAT_bcam_measurement_mdl SM2GDAT_bcam_multiple_mdl SM2GDAT_bcam_description_mdl SM2GDAT_bcam_nea_mdl SM2GDAT_bcam_port_number_mdl
CATEGORY	3	SM2GDAT_BUFFER_CATEGORY	SM2GDAT_category_single_mdl SM2GDAT_category_multiple_mdl SM2GDAT_category_pcs_mdl
CHANNEL- IO	3	SM2GDAT_BUFFER_CHANNEL_I O	SM2GDAT_chio_single_mdl SM2GDAT_chio_multiple_mdl
CMS	3	SM2GDAT_BUFFER_CMS	SM2GDAT_cms_single_mdl SM2GDAT_cms_multiple_mdl
DAB	4	SM2GDAT_BUFFER_DAB	SM2GDAT_dab_single_mdl SM2GDAT_dab_partial_area_mdl SM2GDAT_dab_cache_transfers_mdl SM2GDAT_dab_buffer_area_mdl
DISK-FILE	3	SM2GDAT_BUFFER_DISK_FILE	SM2GDAT_disk_file_single_mdl SM2GDAT_disk_file_multiple_mdl
DLM	1	SM2GDAT_BUFFER_DLM	SM2GDAT_dlm_single_mdl
FILE	3	SM2GDAT_BUFFER_FILE	SM2GDAT_file_single_mdl SM2GDAT_file_multiple_mdl
GS	3	SM2GDAT_BUFFER_GS	SM2GDAT_gs_single_mdl SM2GDAT_gs_multiple_mdl
GSVOL	2	SM2GDAT_BUFFER_GSVOL	SM2GDAT_gsvol_multiple_mdl
HSMS	1	SM2GDAT_BUFFER_HSMS	SM2GDAT_hsms_single_mdl
ISAM	3	SM2GDAT_BUFFER_ISAM	SM2GDAT_isam_single_mdl SM2GDAT_isam_multiple_mdl
ISAM-FILE	3	SM2GDAT_BUFFER_ISAM_FILE	SM2GDAT_isam_file_single_mdl SM2GDAT_isam_file_multiple_mdl
MEMORY	1	SM2GDAT_BUFFER_MEMORY	SM2GDAT_memory_single_mdl
MSCF	1	SM2GDAT_BUFFER_MSCF	SM2GDAT_mscf_single_mdl SM2GDAT_mscf_time_count-mdl

Table 17: Overview of all the data buffers and the structures

(part 1 of 3)

Data buffer	Type	Symbolic constant for selecting the data buffer in “buffer_flags”	Structures
NSM	1	SM2GDAT_BUFFER_NSM	SM2GDAT_nsm_single_md1 SM2GDAT_tokenetab_md1 SM2GDAT_tokenoptab_md1 SM2GDAT_lockconvtab_md1 SM2GDAT_enqlockmodetab_md1 SM2GDAT_lockoptypetab_md1 SM2GDAT_lockservicetab_md1
OPENFT	3	SM2GDAT_BUFFER_OPENFT	SM2GDAT_openft_multiple_md1
PERIODIC-TASK	3	SM2GDAT_BUFFER_PERTASK	SM2GDAT_pertask_single_md1 SM2GDAT_pertask_multiple_md1
PFA	3	SM2GDAT_BUFFER_PFA	SM2GDAT_pfa_single_md1 SM2GDAT_pfa_bcb_md1
POSIX	1	SM2GDAT_BUFFER_POSIX	SM2GDAT_posix_single_md1
PUBSET	3	SM2GDAT_BUFFER_PUBSET	SM2GDAT_pubset_single_md1 SM2GDAT_pubset_multiple_md1
RESPONSE-TIME	4	SM2GDAT_BUFFER_RUNTIME	SM2GDAT_rtime_single_md1 SM2GDAT_rtime_connection_md1 SM2GDAT_rtime_connectionset_md1 SM2GDAT_rtime_category_list_md1 SM2GDAT_rtime_bucket_md1 SM2GDAT_rtime_category_md1
SCHANNEL	2	SM2GDAT_BUFFER_SCHANNEL	SM2GDAT_schannel_multiple_md1
SDEVICE	2	SM2GDAT_BUFFER_SDEVICE	SM2GDAT_sdevice_multiple_md1
SESAM-SQL	3	SM2GDAT_BUFFER_SESAM_SQL	SM2GDAT_sesam_sql_single_md1 SM2GDAT_sesam_sql_multiple_md1
STORAGE-SYSTEM	4	SM2GDAT_BUFFER_ST_SYSTEM	SM2GDAT_sts_single_md1 SM2GDAT_sts_header_md1 SM2GDAT_sts_sym-group_md1 SM2GDAT_sts_sym-dev_group_md1 SM2GDAT_sts_sym-b_end_group_md1 SM2GDAT_sts_sym-disk_group_md1 SM2GDAT_sts_sym-dir_group_md1
SVC	2	SM2GDAT_BUFFER_SVC	SM2GDAT_svc_multiple_md1
SYSTEM	3	SM2GDAT_BUFFER_SYSTEM	SM2GDAT_system_single_md1 SM2GDAT_system_multiple_md1 SM2GDAT_system_queue_data_md1

Table 17: Overview of all the data buffers and the structures

(part 2 of 3)

Data buffer	Type	Symbolic constant for selecting the data buffer in “buffer_flags”	Structures
TCP-IP	3	SM2GDAT_BUFFER_TCP_IP	SM2GDAT_tcp_ip_single_md1 SM2GDAT_tcp_ip_multiple_md1 SM2GDAT_local_addr_md1 SM2GDAT_remote_addr_md1
TIME-IO	3	SM2GDAT_BUFFER_TIME_IO	SM2GDAT_time_io_single_md1 SM2GDAT_time_io_multiple_md1
TLM	3	SM2GDAT_BUFFER_TLM	SM2GDAT_tlm_single_md1 SM2GDAT_tlm_multiple_md1
UDS-SQL	3	SM2GDAT_BUFFER_UDS_SQL	SM2GDAT_uds_sql_single_md1 SM2GDAT_uds_sql_multiple_md1
USERFILE	3	SM2GDAT_BUFFER_USERFILE	SM2GDAT_userfile_single_md1 SM2GDAT_userfile_multiple_md1
USERISAM	3	SM2GDAT_BUFFER_USERISAM	SM2GDAT_userisam_single_md1 SM2GDAT_userisam_multiple_md1
UTM	2	SM2GDAT_BUFFER_UTM	SM2GDAT_utm_multiple_md1 SM2GDAT_utm_constant_md1 SM2GDAT_utm_periodic_md1 SM2GDAT_utm_event_md1 SM2GDAT_utm_avg_md1 SM2GDAT_utm_ext_v2_md1 SM2GDAT_utm_ext_v3_md1 SM2GDAT_utm_ext_tacclass_md1
VM	3	SM2GDAT_BUFFER_VM	SM2GDAT_vm_single_md1 SM2GDAT_vm_multiple_md1
VM-CPU-POOL	3	SM2GDAT_BUFFER_VM_CPU_POOL	SM2GDAT_vm_cpupool_single_md1 SM2GDAT_vm_cpupool_multiple_md1
VM-GROUP	3	SM2GDAT_BUFFER_VM_GROUP	SM2GDAT_vm_group_single_md1 SM2GDAT_vm_group_multiple_md1

Table 17: Overview of all the data buffers and the structures

(part 3 of 3)

11.1.3 Evaluating the return codes

The SM2GMS and SM2GDAT macros store the return codes in a structure of the type ESMFHDR, which can be addressed via the “hdr” structure element of the structure of the type “SM2GMS_get_measurement_stat_md1” (in the case of SM2GMS) or of the structure of the type “SM2GDAT_get_data_md1” (in the case of SM2GDAT).

The ESMFHDR structure is declared in FHDR.H and structured as follows:

```
struct ESMFHDR {
    struct FHDRifid_md1 if_id;      /* interface identifier      */
    struct FHDRretc_md1 returncode; /* return code                */
};
```

The “if_id” structure element of the type “FHDRifid_md1” is used internally by the SM2GMS and SM2GDAT macros and has no significance for the evaluation of the return codes.

The “returncode” structure element of the type “FHDRretc_md1” has the following structure:

```
/* returncode structure */
struct FHDRretc_md1 {
    union /* rc */ {
        struct {
            struct {
                unsigned char subcode2;
                unsigned char subcode1;
            } subcode;
            union /* mc */ {
                unsigned short maincode;
                struct {
                    unsigned char maincode2;
                    unsigned char maincode1;
                } main_returncode;
            } mc;
        } structured_rc;
        unsigned long rc_nbr; /* general return code: */
    } rc;
};
```

Symbolic names and constants are available for the evaluation of the return codes. These are defined in FHDR.H and SM2RC.H, and their use is explained below.

To check the return codes, the user must first query the “returncode.rc.structured_rc.mc.maincode” structure element of the ESMFHDR structure. This structure element (referred to subsequently as MAINCODE) can be addressed by means of the symbolic name FHDR_RC_MAINCODE.

Examples

```
SM2GMS.hdr.FHDR_RC_SUBCODE1
SM2GDAT.hdr.FHDR_RC_SUBCODE1
```

When the macro is executed successfully, MAINCODE has the value "FHDRsuccessful_processing".

In the event of an error, MAINCODE and the "returncode.rc.structured_rc.subcode.subcode1" structure element supply error information.

The "returncode.rc.structured_rc.subcode.subcode1" structure element (referred to subsequently as SUBCODE1) can be addressed by means of the symbolic name FHDR_RC_SUBCODE1.

Examples

```
SM2GMS.hdr.FHDR_RC_SUBCODE1
SM2GDAT.hdr.FHDR_RC_SUBCODE1
```

The following table explains the most important return codes:

SUBCODE1	MAINCODE	Meaning and action to be taken
0	0	The call was successful.
1	65535	The requested function is not supported (incorrect entry for UNIT or FUNCTION in the standard header). Unrecoverable error. Action: None
3	65535	The specified version of the interface is not supported (incorrect version entry in the standard header). Unrecoverable error. Action: None
4	65535	Parameter list is not aligned on a word boundary. Unrecoverable error. Action: None
65	65535	The SM2 subsystem does not exist. Action: The subsystem must be explicitly created.
128	65535	The SM2 subsystem is incapable of responding for a short period. Action: Wait for a short time and then repeat the call.
129	65535	The SM2 subsystem is incapable of responding for a longer period. Action: Wait for a sufficiently long time and then repeat the call.

continued ►

SUBCODE1	MAINCODE	Meaning and action to be taken
1	1	<p>Can only occur in the case of the SM2GDAT macro.</p> <p>The size specified in the “length_buffer” parameter is too small for the requested SM2 data buffers.</p> <p>In this case, the macro stores the actually required size in the “length_buffer” structure element of the structure of the type “SM2GDAT_get_data_md1”, which was passed to the macro as a parameter in the call.</p> <p>Action:</p> <p>Supply the “length_buffer” parameter with the value of the “length_buffer” structure element, make an output area of this size available, and repeat the call.</p>
1	2	<p>Can only occur in the case of the SM2GDAT macro.</p> <p>The size specified in the “length_buffer” parameter is large enough for the requested data buffers, but SM2 cannot validate the output area; i.e. it recognizes that (calculated from the beginning of the output area) it cannot write to the area required for the requested data buffers.</p> <p>This error code does not cover the situation in which the caller specifies a large enough size in “length_buffer” but does not make enough space available for the output area. SM2 cannot determine whether the caller is overwriting his or her own data areas or program code.</p> <p>Action:</p> <p>Correct the program – always make available at least as much space as is specified in the “length_buffer” parameter.</p>
64	3	<p>Can only occur in the case of the SM2GDAT macro.</p> <p>SM2 has no new data.</p> <p>Action:</p> <p>Wait (e.g. 5 seconds) and then try again. If necessary, synchronize the call cycle with the SM2 cycle.</p> <p>The SM2GMS macro supplies the currently set SM2 monitoring cycle.</p>
32	4	<p>System error.</p> <p>Action: None</p>
32	5	<p>The SM2 monitoring task is not running.</p> <p>Action: None</p>
64	6	<p>Specified host unavailable.</p>

The following symbolic constants are available in FHDR.H and SM2RC.H for the values for MAINCODE and SUBCODE1 listed in the table:

Symbolic constants for MAINCODE

```
#define FHDRsuccessful_processing    0
#define SM2RCbuffer_too_short       1
#define SM2RCbuffer_not_valid       2
#define SM2RCno_data                 3
#define SM2RCsystem_error           4
#define SM2RCgatherer_down          5
#define SM2RChost_not_available     6
#define FHDRlinkage_error           65535
```

Symbolic constants for SUBCODE1

```
#define FHDRfct_not_supported        1
#define FHDRver_not_supported       3
#define FHDRalignment_error        4
#define FHDRss_not_created         65
#define FHDRwait_short_term        128
#define FHDRwait_long_term         129
```

The example in the following section provides a practical guide to evaluating return codes.



- SM2RC.H is in the SYSLIB.SM2.<ver> library and can be printed out with the following command:

```
/PRINT-DOCUMENT
  FROM-FILE=*LIBRARY-ELEMENT(LIBRARY=SYSLIB.SM2.<ver>,
    ELEMENT=SM2RC.H,TYPE=S)
```

- FHDR.H is in the SYSLIB.BS2CP.<ver> library and can be printed out analogously.

11.1.4 Example

The following example, which is also in the SYSLIB.SM2.<ver> library. It clarifies the use of the SM2GMS and SM2GDAT macros. To begin with, the sample program outputs the host name and obtains the SM2 online monitoring cycle. The data of the SM2 data buffers BASIC, SDEVICE, TIME_IO and DAB are then output in the cycle of the SM2 online monitoring cycle. Each of these data buffers serves as an example for one of four data buffer types:

- BASIC data buffer (type 1)
The time stamp of the monitoring cycle is output and the number of samples in the cycle obtained.
- SDEVICE data buffer (type 2)
The utilization of each device whose utilization with regard to DMS I/Os exceeds a threshold value is output.
- TIME-IO data buffer (type 3)
The number of active logical machines is output, together with the CPU utilization for each active logical machine whose CPU utilization exceeds a threshold value.
- DAB data buffer (type 4)
The number of reads and the number of read hits are output for each DAB buffer.

Note in the SM2GDAT example that a value of 0 is specified for the “length_buffer” parameter in the first call in order to determine the actually required size of “length_buffer”.

```
#include <stdio.h>
#include <stdlib.h>
#include <signal.h>

#include "FHDR.H"
#include "SM2GMS.H"
#include "SM2GDAT.H"
#include "SM2RC.H"

#define SHORT_TIME 1
#define LONG_TIME 60

#define WAIT_SHORT_AND_TRY_AGAIN 0
#define WAIT_LONG_AND_TRY_AGAIN 1
#define BUFFER_TOO_SHORT 2
#define TERMINATE_PROGRAM 3

#define NOT_DEFINED 0
```



```

#define NOT_SET    0
#define SET        1

#define SIZE_4K_PAGE    4096

#define DMS_IOS_THRESHOLD_VALUE    60.0
#define CPU_UTILIZATION_THRESHOLD_VALUE    50.0

typedef struct SM2GDAT_buffer_header_md1    StrBufferHeader;
typedef struct SM2GDAT_basic_single_md1    StrBASICsingle;
typedef struct SM2GDAT_sdevice_multiple_md1 StrSDEVICEmultiple;
typedef struct SM2GDAT_time_io_single_md1  StrTIME_IOSingle;
typedef struct SM2GDAT_time_io_multiple_md1 StrTIME_IOMultiple;
typedef struct SM2GDAT_dab_single_md1      StrDABsingle;
typedef struct SM2GDAT_dab_buffer_area_md1 StrDAB_BUFFER_AREAmultiple;
typedef struct SM2GDAT_dab_partial_area_md1 StrDAB_PARTIAL_AREAmultiple;
typedef unsigned char                       Uchar;

struct SM2GMS_get_measurement_stat_md1 SM2_STATUS;
struct SM2GDAT_get_data_md1 SM2_DATA;

struct SM2GDAT_global_header_md1 *GlobalHeader_ptr;
struct StrBufferHeader *BufferHeader_ptr;

void *buffer_ptr;
long length_buffer = 0L; /* initialized with 0L for the first call */
                          /* of SM2GDAT */

int Cycle;
int NumberSamples;
char DABData;
char LocalHostName[] = "          ";

void TerminateProgram( char * );

void main( void )
{
    void GetSM2StatusInformations( void );
    void CallSM2GDAT( double );
    void ShowBASICData( void );
    void ShowSDEVICEData( void );
    void ShowTIME_IOData( void );
    void ShowDABData( void );
}

```

```

GetSM2StatusInformations( );

for( ; ; )
{
    CallSM2GDAT( SM2GDAT_BUFFER_BASIC          /* data buffer of type 1 */
                + SM2GDAT_BUFFER_SDEVICE      /* data buffer of type 2 */
                + SM2GDAT_BUFFER_TIME_IO     /* data buffer of type 3 */
                + SM2GDAT_BUFFER_DAB );      /* data buffer of type 4 */

    ShowBASICData( );
    ShowSDEVICEData( );
    ShowTIME_IOData( );
    ShowDABData( );

    sleep( Cycle );
}

exit( 0 );
}

void TerminateProgram( char *Message )
{
    printf( "\n\n%s\n\n", Message );
    printf( "Program abnormally terminated.\n" );
    exit( -1 );
}

void GetSM2StatusInformations( void )
{
    char TRY_AGAIN;
    char Message[100];
    int Errorhandling( int, int, char * );

    TRY_AGAIN = 'y';
    while( TRY_AGAIN == 'y' )
    {
        SM2GMS( SM2_STATUS, LocalHostName ); /* calling macro SM2GMS */
    }
}

```

```

if( SM2_STATUS.hdr.FHDR_RC_MAINCODE == FHDRsuccessful_processing )
    TRY_AGAIN = 'n';
else
    {
    switch( Errorhandling( ( int )SM2_STATUS.hdr.FHDR_RC_MAINCODE,
                          ( int )SM2_STATUS.hdr.FHDR_RC_SUBCODE1,
                          Message ) )
        {
        case WAIT_SHORT_AND_TRY_AGAIN:
            sleep( SHORT_TIME );
            TRY_AGAIN = 'y';
            break;

        case WAIT_LONG_AND_TRY_AGAIN:
            sleep( LONG_TIME );
            TRY_AGAIN = 'y';
            break;

        case TERMINATE_PROGRAM:
            TerminateProgram( Message );
            break;

        default:
            TerminateProgram( "Unexpected return code from SM2GMS.\n" );
            break;
        }
    }
}

printf( "Host: %8.8s\n", SM2_STATUS.status.endsystem_name );

if( SM2_STATUS.status.online_cycle == NOT_DEFINED )
    Cycle = SM2_STATUS.status.offline_cycle;
else
    Cycle = SM2_STATUS.status.online_cycle;

if( SM2_STATUS.status.active_programs.flag.dab == SET )
    DABData = 'y';      /* measurement program DAB is active */
else
    DABData = 'n';      /* measurement program DAB is not active */
}

```

```
void CallSM2GDAT( double BufferSelection )
{
    char Message[100];
    int Errorhandling( int, int, char * );
    void GetMemory( void );
    char TRY_AGAIN;

    TRY_AGAIN = 'y';
    while( TRY_AGAIN == 'y' )
    {
        /* calling macro SM2GDAT to get the data buffers */
        SM2GDAT( SM2_DATA, length_buffer, buffer_ptr, BufferSelection, LocalHostName );

        if( SM2_DATA.hdr.FHDR_RC_MAINCODE == FHDRsuccessful_processing )
            TRY_AGAIN = 'n';
        else
        {
            switch( Errorhandling( ( int )SM2_DATA.hdr.FHDR_RC_MAINCODE,
                                   ( int )SM2_DATA.hdr.FHDR_RC_SUBCODE1,
                                   Message ) )
            {
                case WAIT_SHORT_AND_TRY_AGAIN:
                    sleep( SHORT_TIME );
                    TRY_AGAIN = 'y';
                    break;

                case WAIT_LONG_AND_TRY_AGAIN:
                    sleep( LONG_TIME );
                    TRY_AGAIN = 'y';
                    break;

                case BUFFER_TOO_SHORT: /* needed buffer_length has increased */
                    /* since the last call of SM2GDAT - this return code */
                    /* is especially expected for the first call of */
                    /* SM2GDAT, because length_buffer is initialized with 0 */

                    length_buffer =
                        SM2_DATA.length_buffer; /* copy the needed size */

                    GetMemory( ); /* allocate memory for the output area */

                    TRY_AGAIN = 'y';
                    break;
            }
        }
    }
}
```

```
        case TERMINATE_PROGRAM:
            TerminateProgram( Message );
            break;

        default:
            TerminateProgram( "Unexpected return code from SM2GDAT.\n" );
            break;
    }
}

/* initialize pointer to evaluate the global header */
GlobalHeader_ptr = ( struct SM2GDAT_global_header_md1 * )buffer_ptr;
}

int Errorhandling( int MAINCODE, int SUBCODE1, char *Message )
{
    switch( MAINCODE )
    {
        case FHDRlinkage_error:
            switch( SUBCODE1 )
            {
                case FHDRfct_not_supported:
                case FHDRfct_not_available:
                case FHDRver_not_supported:
                case FHDRalignment_error:
                case FHDRss_not_created:
                    break;

                case FHDRwait_short_term:
                    return( WAIT_SHORT_AND_TRY_AGAIN );

                case FHDRwait_long_term:
                    return( WAIT_LONG_AND_TRY_AGAIN );

                default:
                    break;
            }
            break;

        case SM2RCbuffer_too_short:
            return( BUFFER_TOO_SHORT );

        case SM2RCno_data:
            return( WAIT_SHORT_AND_TRY_AGAIN );
    }
}
```

```
    case SM2RCbuffer_not_valid:
    case SM2RCsystem_error:
    case SM2RCgatherer_down:
        break;
    default:
        break;
}

sprintf( Message, "MAINCODE: %d, SUBCODE1: %d",
        MAINCODE, SUBCODE1 );

return( TERMINATE_PROGRAM );
}

void GetMemory( void )
{
    static char FirstCall = 'y';

    if( FirstCall == 'y' )
        FirstCall = 'n';
    else
        free( buffer_ptr );

    if( ( buffer_ptr =
        ( Uchar *)malloc((unsigned int)( SIZE_4K_PAGE * length_buffer )) ) == NULL )
        TerminateProgram( "malloc( ): Not enough memory." );
}

/*
 *   evaluates the BASIC data buffer - this is an example to evaluate
 *   a data buffer of type 1
 */

void ShowBASICData( void )
{
    StrBufferHeader *BufferHeader_ptr;
    StrBASICsingle *BASICsingle_ptr;
```

```

        /* initialize pointer to evaluate the buffer header of the */
        /* BASIC data buffer */
    BufferHeader_ptr = ( StrBufferHeader * )GlobalHeader_ptr->basic_ptr;

    if( BufferHeader_ptr->state.valid == NOT_SET ) /* no valid data in data buffer */
    {
        printf( "Data in BASIC-Buffer not valid.\n" );
        return;
    }

        /* initialize pointer to evaluate the fixed data area of the */
        /* BASIC data buffer */
    BASICsingle_ptr =
    ( StrBASICsingle * ) ( ( Uchar * ) BufferHeader_ptr + BufferHeader_ptr->fixed_part_dsp );

    printf( "\n\nTime stamp of cycle:  %10.10s %8.8s\n",
           BASICsingle_ptr->date, BASICsingle_ptr->time );

    NumberSamples = BASICsingle_ptr->samples;
}

/*
 *   evaluates the SDEVICE data buffer - this is an example to evaluate
 *   a data buffer of type 2
 */

void ShowsDEVICEData( void )
{
    StrBufferHeader *BufferHeader_ptr;
    StrSDEVICEmultiple *SDEVICEmultiple_ptr;
    Uchar *Multiple_ptr;
    int LengthMultipleGroup;
    int NumberMultipleGroups;
    float DmsBusy;
    int i;

        /* initialize pointer to evaluate the buffer header of the */
        /* SDEVICE data buffer */
    BufferHeader_ptr = ( StrBufferHeader * )GlobalHeader_ptr->sdevice_ptr;

```

```

if( BufferHeader_ptr->state.valid == NOT_SET ) /* no valid data in data buffer */
{
    printf( "\nData in SDEVICE-Buffer not valid.\n" );
    return;
}

LengthMultipleGroup = BufferHeader_ptr->length_group;
NumberMultipleGroups = BufferHeader_ptr->number_groups;

    /* initialize pointer with the address of the first repeat group */
Multiple_ptr = ( Uchar * ) BufferHeader_ptr + BufferHeader_ptr->first_group_dsp;

printf( "\n" );
for( i = 0; i != NumberMultipleGroups; i++, Multiple_ptr += LengthMultipleGroup )
{
    /* initialize pointer to evaluate the repeat group of the */
    /* SDEVICE data buffer */
    SDEVICEMultiple_ptr = ( StrSDEVICEMultiple * ) Multiple_ptr;
    DmsBusy = ( SDEVICEMultiple_ptr->busy_dms * 100.0 ) / NumberSamples;
    if( DmsBusy >= DMS_IOS_THRESHOLD_VALUE )
        printf( "Utilization of VSN %6.6s, MN %4.4s, Type %02x: %f\n",
                SDEVICEMultiple_ptr->vsn,
                SDEVICEMultiple_ptr->mnemonic,
                SDEVICEMultiple_ptr->type[1],
                DmsBusy );
}
}

/*
 *   evaluates the TIME_IO data buffer - this is an example to evaluate
 *   a data buffer of type 3
 */

void ShowTIME_IOData( void )
{
    StrBufferHeader *BufferHeader_ptr;
    StrTIME_IOsingle *TIME_IOsingle_ptr;
    StrTIME_IOMultiple *TIME_IOMultiple_ptr;
    Uchar *Multiple_ptr;
    int LengthMultipleGroup;
    int NumberMultipleGroups;
    float SUM, CPU_Utilization;
    int i;

```



```

    /* initialize pointer to evaluate the buffer header of the */
    /* TIME_IO data buffer */
    BufferHeader_ptr = ( StrBufferHeader * )GlobalHeader_ptr->time_io_ptr;

    if( BufferHeader_ptr->state.valid == NOT_SET ) /* no valid data in data buffer */
    {
        printf( "\nData in TIME_IO-Buffer not valid.\n" );
        return;
    }

    /* initialize pointer to evaluate the fixed data area of the */
    /* TIME_IO data buffer */
    TIME_IOsingle_ptr = ( StrTIME_IOsingle * ) ( ( Uchar * ) BufferHeader_ptr
        + BufferHeader_ptr->fixed_part_dsp );

    printf( "\nNumber active logical machines: %d\n",
        TIME_IOsingle_ptr->active_lm );

    LengthMultipleGroup = BufferHeader_ptr->length_group;
    NumberMultipleGroups = BufferHeader_ptr->number_groups;

    /* initialize pointer with the address of the first repeat group */
    Multiple_ptr = ( Uchar * ) BufferHeader_ptr + BufferHeader_ptr->first_group_dsp;

    for( i = 0; i != NumberMultipleGroups; i++, Multiple_ptr += LengthMultipleGroup )
    {
        /* initialize pointer to evaluate the repeat group of the */
        /* TIME_IO data buffer */
        TIME_IOfmultiple_ptr = ( StrTIME_IOfmultiple * ) Multiple_ptr;
        SUM = ( TIME_IOfmultiple_ptr->tu_time
            + TIME_IOfmultiple_ptr->tpr_time
            + TIME_IOfmultiple_ptr->sih_time );

        if( ( SUM + TIME_IOfmultiple_ptr->idle_time ) > 0 )
            CPU_Utilization = ( SUM * 100.0 )
                / ( SUM + TIME_IOfmultiple_ptr->idle_time );
        else
            CPU_Utilization = 0.0;
    }

```

```

    if( i == 0 ) /* the first repeat group contains the average values */
        /* of all logical machines */
        printf( "Average CPU utilization of all logical machines:  %f %%\n",
            CPU_Utilization );
    else
        if( CPU_Utilization >= CPU_UTILIZATION_THRESHOLD_VALUE )
            printf( "CPU utilization logical machine %d:  %f %%\n",
                i, CPU_Utilization );
    }
}

/*
 *   evaluates the DAB data buffer - this is an example to evaluate
 *   a data buffer of type 4
 */

void ShowDABData( void )
{
    StrBufferHeader *BufferHeader_ptr;
    StrDABsingle *DABsingle_ptr;
    StrDAB_BUFFER_AREAmultiple *DAB_BUFFER_AREAmultiple_ptr;
    Uchar *Multiple_ptr;
    int LengthMultipleGroup;
    int NumberMultipleGroups;
    long reads, read_hits;
    int i;
    void EvaluatePartialAreas( int, StrDABsingle *, long *, long *);

    if( DABData == 'n' )
    {
        printf( "\nMeasurement program DAB is not active.\n" );
        return;
    }

    /* initialize pointer to evaluate the buffer header of the */
    /* DAB data buffer */
    BufferHeader_ptr = ( StrBufferHeader * )GlobalHeader_ptr->dab_ptr;

    if( BufferHeader_ptr->state.valid == NOT_SET ) /* no valid data in data buffer */
    {
        printf( "\nData in DAB-Buffer not valid.\n" );
        return;
    }
}

```

```

        /* initialize pointer to evaluate the fixed data area of the */
        /* DAB data buffer */
DABsingle_ptr =
( StrDABsingle * ) ( ( Uchar * ) BufferHeader_ptr + BufferHeader_ptr->fixed_part_dsp );

if( DABsingle_ptr->reconfigurations != 0 ) /* DAB reconfiguration */
{
    printf( "\nDAB reconfiguration in last cycle.\n" );
    return;
}

LengthMultipleGroup = DABsingle_ptr->length_buffer_area_group;
NumberMultipleGroups = DABsingle_ptr->number_buffer_area_groups;

        /* initialize pointer with the address of the first buffer area */
        /* repeat group */
Multiple_ptr =
    ( Uchar * ) DABsingle_ptr + DABsingle_ptr->first_buffer_area_group_dsp;

printf( "\n" );
for( i = 0; i != NumberMultipleGroups; i++, Multiple_ptr += LengthMultipleGroup )
{
    reads = read_hits = 0L;
        /* initialize pointer to evaluate the buffer area repeat group */
        /* of the DAB data buffer */
DAB_BUFFER_AREAmultiple_ptr =
    ( StrDAB_BUFFER_AREAmultiple * ) Multiple_ptr;
EvaluatePartialAreas( i + 1, DABsingle_ptr, &reads, &read_hits );

    printf( "DAB buffer %s: reads: %ld, read_hits: %ld\n",
        DAB_BUFFER_AREAmultiple_ptr->id, reads, read_hits );
}
}

void EvaluatePartialAreas( int BufferIndex, StrDABsingle *DABsingle_ptr,
    long *reads, long *read_hits )
{
    StrDAB_PARTIAL_AREAmultiple *DAB_PARTIAL_AREAmultiple_ptr;
    Uchar *Multiple_ptr;
    int LengthMultipleGroup;
    int NumberMultipleGroups;
    int i;

```

```
LengthMultipleGroup = DABsingle_ptr->length_partial_area_group;
NumberMultipleGroups = DABsingle_ptr->number_partial_area_groups;

    /* initialize pointer with the address of the first partial area */
    /* repeat group */
Multiple_ptr =
    ( Uchar * ) DABsingle_ptr + DABsingle_ptr->first_partial_area_group_dsp;

for( i = 0; i != NumberMultipleGroups; i++, Multiple_ptr += LengthMultipleGroup )
{
    /* initialize pointer to evaluate the buffer area repeat group */
    /* of the DAB data buffer */
    DAB_PARTIAL_AREAmultiple_ptr =
        ( StrDAB_PARTIAL_AREAmultiple * ) Multiple_ptr;

        /* the partial area belongs to the buffer and was served in */
        /* the monitoring cycle */
    if( DAB_PARTIAL_AREAmultiple_ptr->buffer_index == BufferIndex
        && DAB_PARTIAL_AREAmultiple_ptr->state == SM2GDATserved )
    {
        *reads += DAB_PARTIAL_AREAmultiple_ptr->transfers.read;
        *read_hits += DAB_PARTIAL_AREAmultiple_ptr->transfers.read_hit;
    }
}
}
```

11.1.5 Structures of the SM2GMS macro

The declarations of the structures of the SM2GMS macro are in SM2GMS.H, and the definitions of the symbolic constants for evaluating the return codes are in SM2RC.H.

The “SM2GMS_get_measurement_stat_md1” structure, which is used to call the macro, contains the “hdr” structure element of the structure type ESMFHDR (this structure is declared in FHDR.H) and the “status” structure element of the structure type “SM2GMS_status_data_md1”.

The macro sets parameters internally in “hdr”, and it also stores its return codes there. The evaluation of the return codes is described in [section “Evaluating the return codes” on page 540](#).

The macro supplies the “host_name” element internally with the value passed by the caller in the “host_name” parameter.

In the “status” structure element, the macro stores its information on the status of SM2. This information corresponds to that in the MEASUREMENT STATUS screen (see [page 335](#)).

The “status” structure element contains four elements of the structure type “SM2GMS_measurement_program”, in which displays for the defined and active monitoring programs of SM2 are set. By querying these flags, the user can check whether the associated monitoring programs for the data buffers requested with the SM2GDAT macro are activated in SM2.

The “ss_sm2_status” should also be noted. This indicates whether the SM2 subsystem should be unloaded.

11.1.6 Structures of the SM2GDAT macro

The declarations of the structures of the SM2GDAT macro are in SM2GDAT.H, and the definitions of the symbolic constants for evaluating the return codes are in SM2RC.H.

In addition to the “SM2GDAT_get_data_md1” structure (and its substructures) used in the macro call, there are a number of structures for describing the data buffers of the output area.

All the structures of the SM2GDAT macro are described below. The data buffers are described in alphabetical order.

The “SM2GDAT_get_data_md1” structure

The “SM2GDAT_get_data_md1” structure is passed to the SM2GDAT macro as a parameter in the macro call.

The macro sets parameters internally in “hdr”, and it also stores its return codes there. The evaluation of the return codes is described in [section “Evaluating the return codes” on page 540](#).

The macro supplies the “selected1”, “selected2”, “length_buffer”, “buffer_ptr” and “host_name” elements internally with the values passed by the caller in the “buffer_flags”, “length_buffer”, “buffer_ptr” or “host_name” parameters.

After an error-free call and in the case of the return code SM2GDAT_buffer_too_short, the “buffer_length” element contains the actually used or required length of the output area.

Structure of the global header

The global header is located at the beginning of the output area and is described by the “SM2GDAT_global_header_md1” structure.

For each data buffer, it contains a pointer to the beginning of the data buffer. Only the pointers for those data buffers requested by the caller are supplied. All other pointers are filled with null.

Structure of the BUFFER HEADER

Each data buffer begins with the BUFFER HEADER, which is described by the “SM2GDAT_buffer_header_md1” structure. The BUFFER HEADER contains administration information that enables the actual monitoring data (in the data area) of the respective data buffer to be accessed. The administration information is evaluated differently, depending on the type of data buffer (see the description of the relevant data buffer).

The “valid” bit indicates whether valid data exists for the data buffer. It must always be queried first. A “valid” bit that is not set means both that the remaining administration elements in the BUFFER HEADER and the monitoring data in the data area may be invalid.

The possible causes of an unset “valid” bit are:

- The monitoring program is not activated.
- SM2 cannot create any monitoring data because different monitoring objects occurred in two successive monitoring cycles.
- SM2 cannot create any monitoring data because SM2 did not receive any basic data from the system.
- SM2 has an error.

The “size” element specifies the length of the data buffer (including that of the BUFFER HEADER).

Structure of the ACF data buffer

The ACF data buffer supplies data for activation control of the PRIOR task administration facility. The data provides information for the internal decision concerning activation.

The data area of the ACF data buffer comprises an area of fixed length.

Data buffer type: 1

Fixed-length data area

The fixed-length data area contains monitoring data of the PRIOR task scheduler (number of task activations, deactivations, preemptions, etc.). In addition, 3 counters for each of the resources CPU, MEMORY and PAGING provide information on how often the utilization of these resources was low, medium or high.

Descriptive structure: struct SM2GDAT_ACF_single_md

Displacement from the
beginning of the BUFFER “fixed_part_dsp” from the BUFFER HEADER
HEADER:



Corresponding screen report in the online component: ACF
Monitoring program: –

Structure of the BASIC data buffer

The BASIC data buffer contains basic information on SM2 monitoring.

The data area of the BASIC data buffer comprises an area of fixed length.

Data buffer type: 1

Fixed-length data area

The fixed-length data area supplies basic data for the most recently completed monitoring cycle (elapsed time, data, time, number of samples, etc.). The number of logical machines set, the number of data records not written to the output file (and thus lost), a VM2000 indicator and the host name are also made available.

Descriptive structure: struct SM2GDAT_basic_single_md1

Displacement from the
beginning of the BUFFER "fixed_part_dsp" from the BUFFER HEADER
HEADER:



Corresponding screen report in the online component: ACTIVITY
Monitoring program: –

Structure of the BCAM-CONNECTION data buffer

The BCAM-CONNECTION data buffer supplies data for defined connection sets and on the size of the resident memory for data transfer.

The data area of the BCAM-CONNECTION comprises an area of fixed length and a number of repeat groups.

Data buffer type: 3

Fixed-length data area.

The fixed length data area contains the number of subsequent repeat groups and the data on the size of the resident memory for data transfer.

Descriptive structure: SM2GDAT_bcam_single_md1

Displacement from the
beginning of the BUFFER "fixed_part_dsp" from the BUFFER HEADER
HEADER:

Repeat groups

A repeat group is written for every monitored connection set.

The repeat group contains

- the name of the connection set,
- the definition of the connection set
(the structures SM2GDAT_bcam_description_mdl, SM2GDAT_bcam_nea_mdl, SM2GDAT_bcam_port_number_mdl),
- the number of active connections, and
- the monitored data (the structure SM2GDAT_bcam_measurement_mdl) for the connection set.

The “description_data_dsp” element contains the displacement from the beginning of the repeat group to the definition of the connection set. Among other things, the definition of the connection set contains a list of application names or port numbers which can be accessed via length, number and displacement using the familiar logic. The access elements are “length_description”, “number_description” and “first_description_dsp”. The displacement is measured from the start of the definition of the connection set.

The connection set data comprises only those connections that were active both at observation time and during the immediately preceding monitoring cycle.

Descriptive structure:	SM2GDAT_bcam_multiple_mdl SM2GDAT_bcam_measurement_mdl SM2GDAT_bcam_description_mdl SM2GDAT_bcam_nea_mdl SM2GDAT_bcam_port_number_mdl
------------------------	---

Displacement of the first repeat group from the beginning of the BUFFER HEADER:	“first_group_dsp” from the BUFFER HEADER
---	--

Length of a repeat group:	“length_group” from the BUFFER HEADER
---------------------------	---------------------------------------

Number of repeat groups:	“number_groups” from the BUFFER HEADER
--------------------------	--



Corresponding screen reports in the online component:
BCAM CONNECTION, BCAM MEMORY
Monitoring program: BCAM-CONNECTION

Structure of the CATEGORY data buffer

The CATEGORY data buffer contains data for task administration and the working set and data of the PCS system. It supplies monitoring data for each category.

The data area of the CATEGORY data buffer consists of an area of fixed length and a number of repeat groups.

Data buffer type: 3

Fixed-length data area

The fixed-length data area contains monitoring data not assigned to any specific category.

The mean number of tasks is made available for the SYS, DIALOG, BATCH and TP task types. The sum of the values provides the mean number of all tasks in the system.

The “pcs_active” indicator reveals whether data of the PCS subsystem exist. In the case of PCS reconfiguration (“pcs_reconfigurations” not equal to 0), no PCS data can be supplied.

Most of the values of the CATEGORY data buffer are obtained by sampling. You should bear in mind that the accuracy of the data varies with the number of samples taken.

Descriptive structure: SM2GDAT_category_single_mdl

Displacement from the beginning of the BUFFER HEADER: “fixed_part_dsp” from the BUFFER HEADER

Repeat groups

The first repeat group contains the sum values of all categories. SUM is output as “name”. “type” and “number” are undefined. Each subsequent repeat group contains the monitoring data for a single category.

Each repeat group provides information on the number of tasks in the different queues and on the working set of different task groups.

The data of the PCS subsystem is only valid when the “pcs_data” indicator is set in the SM2GDAT_category_single_mdl structure.

Descriptive structures: SM2GDAT_category_multiple_mdl
SM2GDAT_category_pcs_mdl

Displacement of the first repeat group from the beginning of the BUFFER HEADER: “first_group_dsp” from the BUFFER HEADER

Length of the repeat group: “length_group” from the BUFFER HEADER

Number of repeat groups: "used_groups" from
SM2GDAT_category_single_mdI



Corresponding screen reports in the online component:
ACTIVITY, CATEGORY QUEUE, CATEGORY WSET, PCS
Monitoring program: –

Structure of the CHANNEL-IO data buffer

The CHANNEL-IO data buffer contains data on I/O channel data transmission.

Fixed-length data area.

The number of repeat groups that are subsequently used is located in "number_used_groups".

Descriptive structure: SM2GDAT_hsms_single_mdI

Displacement from the beginning of the BUFFER HEADER: "fixed_part_dsp" from the BUFFER HEADER

Repeat groups

A repeat group is written for each monitored channel. The repeat group provides information about the number of I/Os or the volume of the transferred data in the case of PAM block or byte transfer. The I/Os without data transfer are also output.

Descriptive structure: SM2GDAT_chio_multiple_mdI

Displacement of the first repeat group from the beginning of the BUFFER HEADER: "first_group_dsp" from the BUFFER HEADER

Length of a repeat group: "length_group" from the BUFFER HEADER

Number of repeat groups: "number_groups" from the BUFFER HEADER



Corresponding screen report in the online component: CHANNEL
Monitoring program: CHANNEL-IO

Structure of the CMS data buffer

The CMS data buffer supplies data on the performance of the catalog management system.

The data area of the CMS data buffer comprises an area of fixed length and a number of repeat groups.

Data buffer type: 3

Fixed-length data area

The fixed-length data area contains the number of CMS reconfigurations. In the case of a reconfiguration, no data can be supplied to the repeat groups.

Descriptive structure: SM2GDAT_cms_single_mdI

Displacement from the beginning

of the BUFFER HEADER: "fixed_part_dsp" from the BUFFER HEADER

Repeat groups

A repeat group is written for each imported public volume set. In addition, a repeat group is written for the entire set of all private disks. Data is supplied for queues, accesses and response times.

The data for the queues is obtained by sampling. The number of samples is contained in "samples". You should bear in mind that the accuracy of the data varies with the number of samples taken.

Descriptive structure: SM2GDAT_cms_multiple_mdI

Displacement of the first repeat group

from the beginning of the BUFFER HEADER: "first_group_dsp" from the BUFFER HEADER

Length of a repeat group: "length_group" from the BUFFER HEADER

Number of repeat groups: "number_groups" from the BUFFER HEADER



Corresponding screen report in the online component: CMS
Monitoring program: CMS

Structure of the DAB data buffer

The DAB data buffer supplies data on the activity of the disk access buffer function.

The data area of the DAB data buffer comprises an area of fixed length and two types of repeat group, which are arranged in two chains.

Data buffer type: 4

Fixed-length data area

For each of the two chains of repeat groups, the fixed-length data area contains the length of a repeat group, the number of repeat groups and the displacement from the first repeat group. The displacements are measured from the beginning of the “SM2GDAT_dab_single_md1” structure.

The “reconfigurations” element provides information on the possible changes to the configuration of the DAB buffers. In this case, the repeat groups contain only information on the new configuration of the DAB buffers but no monitoring data on the accesses.

Descriptive structure:	SM2GDAT_dab_single_md1
Displacement from the beginning of the BUFFER HEADER:	“fixed_part_dsp” from the BUFFER HEADER

The BUFFER AREA repeat group

A repeat group exists for each DAB cache area. This contains the name and size of the DAB cache area as well as various attributes. The repeat groups come one after the other in the data buffer.

Descriptive structure:	SM2GDAT_dab_buffer_area_md1
Displacement of the first repeat group from the beginning of the structure SM2GDAT_dab_single_md1:	“first_buffer_area_group_dsp” from SM2GDAT_dab_single_md1
Length of a repeat group:	“length_buffer_area_group” from SM2GDAT_dab_single_md1
Number of repeat groups:	“number_buffer_area_groups” from SM2GDAT_dab_single_md1

The PARTIAL AREA repeat group

A repeat group exists for each subarea or file of each DAB buffer. The repeat groups come one after the other in the data buffer.

The assignment of the PARTIAL AREA repeat groups to the DAB cache areas is as follows:

If the “buffer_index” element of the PARTIAL AREA repeat group has the value n, the subarea belongs to the DAB buffer of the nth repeat group in the chain of BUFFER AREA repeat groups.

A PARTIAL AREA repeat group contains administration information and monitoring data on the accesses (“SM2GDAT_cache_transfers_mdl” structure). The “state” indicator reveals whether the subarea was used in the monitoring cycle. The “assignment” indicator reveals whether a file or an area on disk is used. The file name or VSN is also output. The “first_page” and “last_page” elements specify the first and last block number, respectively. Depending on “assignment”, they are physical or logical block numbers.

Descriptive structures:	SM2GDAT_dab_partial_area_mdl SM2GDAT_cache_transfers_mdl
Displacement of the first repeat group from the beginning of the structure SM2GDAT_dab_single_mdl:	“first_partial_area_group_dsp” from SM2GDAT_dab_single_mdl
Length of a repeat group:	“length_partial_area_group” from SM2GDAT_dab_single_mdl
Number of repeat groups:	“number_partial_area_groups” from SM2GDAT_dab_single_mdl



Corresponding screen report in the online component: DAB, DAB CACHE.
Monitoring program: DAB

Structure of the DISK-FILE data buffer

For the selected disk devices, the DISK-FILE data buffer contains an overview of the files which have triggered physical I/Os.

Fixed-length data area.

The number of repeat groups that are subsequently used is located in “number_used_groups”.

Descriptive structure: SM2GDAT_disk_file_single_mdl

Displacement from the beginning of the BUFFER HEADER: “fixed_part_dsp” from the BUFFER HEADER

Repeat groups

A maximum of 300 repeat groups are created for each monitored disk device. Each repeat group contains the mnemonic of the disk device, the VSN and the file name together with the number of physical I/Os for reading and writing. If multiple disk devices are monitored then the repeat groups are always collated for each individual device on output, i.e. if there are two devices, then all the repeat groups for the first device are output first followed by all the repeat groups for the second device.

For each device, a repeat group with the filename “*OVERRUNS” is always output first. These group gathers together the I/Os for those files for which there was no free entry in the SM2 table.

Descriptive structure: SM2GDAT_disk_file_multiple_mdl

Displacement of the first repeat group from the beginning of the BUFFER HEADER: “first_group_dsp” from the BUFFER HEADER

Length of a repeat group: “length_group” from the BUFFER HEADER

Number of repeat groups: “number_groups” from SM2GDAT_disk_file_single



Corresponding screen report in the online component: DISK-FILE
Monitoring program: DISK-FILE

Structure of the DLM data buffer

The DLM data buffer provides general information on Distributed Lock Managers events.

The data area of the DLM data buffer comprises an area of fixed length.

Data buffer type: 1

Fixed-length data area

The fixed-length data area contains the number of different lock requests from TU, TPR and NSM in the system.

Descriptive structure: SM2GDAT_dlm_single_mdI

Displacement from the beginning of the BUFFER HEADER: "fixed_part_dsp" from the BUFFER HEADER



Corresponding screen report in the online component: DLM
Monitoring program: DLM

Structure of the FILE data buffer

The FILE data buffer contains data on file accesses.

The data area of the FILE data buffer comprises an area of fixed length and a number of repeat groups.

Data buffer type: 3

Fixed-length data area

The fixed-length data area contains only the number of repeat groups to follow.

Descriptive structure: SM2GDAT_file_single_mdl

Displacement from the beginning of the BUFFER HEADER: "fixed_part_dsp" from the BUFFER HEADER

Repeat groups

A repeat group is written for each file monitored. The repeat group provides information on the number of read, write, wait and check accesses as well as the average duration of I/Os.

Descriptive structure: SM2GDAT_file_multiple_mdl

Displacement of the first repeat group from the beginning of the BUFFER HEADER: "first_group_dsp" from the BUFFER HEADER

Length of a repeat group: "length_group" from the BUFFER HEADER

Number of repeat groups: "used_groups" from SM2GDAT_file_single_mdl



Corresponding screen report in the online component: FILE.
Monitoring program: FILE

Structure of the GS data buffer

The GS data buffer contains information about global storage.

The data area of the GS data buffer comprises an area of fixed length and a number of repeat groups.

Data buffer type: 3

Fixed-length data area

The fixed-length data area provides the number of GS reconfigurations. During a reconfiguration, no data can be supplied in the repeat groups.

Descriptive structure: SM2GDAT_gs_single_mdI

Displacement from the beginning of the BUFFER HEADER: "fixed_part_dsp" from the BUFFER HEADER

Repeat groups

A repeat group is written for each GS partition. The repeat group provides the number of accesses and the number of bytes read and written for each GS partition.

Descriptive structure: SM2GDAT_gs_multiple_mdI

Displacement of the first repeat group from the beginning of the BUFFER HEADER: "first_group_dsp" from the BUFFER HEADER

Length of a repeat group: "length_group" from the BUFFER HEADER

Number of repeat groups: "number_groups" from the BUFFER HEADER



Corresponding screen report in the online component: GS.

Monitoring program: GS

Structure of the GSVOL data buffer

The GSVOL data buffer contains information about emulated volumes in global storage.

The data area of the GSVOL data buffer comprises only repeat groups.

Data buffer type: 2

Repeat groups

A repeat group is written for each volume. The repeat group provides the device name, the GS unit and the number of I/Os.

Descriptive structure:	SM2GDAT_gsvol_multiple_mdI
Displacement of the first repeat group from the beginning of the BUFFER HEADER:	“first_group_dsp” from the BUFFER HEADER
Length of a repeat group:	“length_group” from the BUFFER HEADER
Number of repeat groups:	“number_groups” from the BUFFER HEADER



Corresponding screen report in the online component: GSVOL
Monitoring program: GSVOL

Structure of the HSMS data buffer

The HSMS data buffer provides information about the migration of files to the background level and their retrieval into the processing level. The data area of the HSMS data buffer comprises an area of fixed length.

Data buffer type: 1

Fixed length data area

Descriptive structure:	SM2GDAT_hsms_single_mdI
Displacement from the beginning of the BUFFER HEADER:	“fixed_part_dsp” from the BUFFER HEADER



Corresponding screen report in the online component: ---
Monitoring program: HSMS

Structure of the ISAM data buffer

The ISAM data buffer supplies data on access to ISAM pools.

The data area of the ISAM data buffer comprises an area of fixed length and a number of repeat groups.

Data buffer type: 3

Fixed-length data area

The fixed-length data area contains only the number of repeat groups to follow.

Descriptive structure: SM2GDAT_isam_single_mdl

Displacement from the beginning
of the BUFFER HEADER: "fixed_part_dsp" from the BUFFER HEADER

Repeat groups

A repeat group is written for each ISAM pool monitored. The repeat group provides information on the accesses and the pages used for buffering.

Descriptive structure: SM2GDAT_isam_multiple_mdl

Displacement of the first repeat group
from the beginning of the BUFFER
HEADER: "first_group_dsp" from the BUFFER
HEADER

Length of a repeat group: "length_group" from the BUFFER HEADER

Number of repeat groups: "used_groups" from
SM2GDAT_isam_single_mdl



Corresponding screen report in the online component: ISAM
Monitoring program: ISAM

Structure of the SAM-FILE data buffer

The ISAM-FILE data buffer supplies data on access to NK-ISAM files in a global ISAM pool in the Data Space.

The data area of the ISAM-FILE data buffer comprises an area of fixed length and a number of repeat groups.

Data buffer type: 3

Fixed-length data area

The fixed-length data area contains only the number of repeat groups to follow..

Descriptive structure: SM2GDAT_isam_file_single_mdI

Displacement from the beginning

of the BUFFER HEADER: "fixed_part_dsp" from the BUFFER HEADER

Repeat groups

A repeat grup is written for each ISAM file monitored. The repeat group provides on the accesses and the pages used for buffering.

Descriptive structure: SM2GDAT_isam_file_multiple_mdI

Displacement of the first repeat group
from the beginning of the BUFFER
HEADER:

"first_group_dsp" from the BUFFER
HEADER

Length of a repeat group:

"length_group" from the BUFFER HEADER

Number of repeat groups:

"used_groups" from
SM2GDAT_isam_file_single_mdI



Corresponding screen report in the online component: ISAM FILE
Monitoring program: ISAM

Structure of the MEMORY data buffer

The MEMORY data buffer supplies data on the utilization of main memory and the virtual address space.

The data area of the MEMORY data buffer comprises an area of fixed length.

Data buffer type: 1

Fixed-length data area

The fixed-length data area provides information on the size and utilization of main memory and on paging.

Descriptive structure: `struct SM2GDAT_memory_single_mdI`

Displacement from the beginning of the BUFFER HEADER: "fixed_part_dsp" from the BUFFER HEADER



Corresponding screen reports in the online component: ACTIVITY, MEMORY
Monitoring program: –

Structure of the MSCF data buffer

The MSCF data buffer provides general information on the MSCF subsystem.

The data area of the MSCF data buffer comprises an area of fixed length.

Data buffer type: 1

Fixed-length data area

The fixed-length data area contains the number of server tasks and limit values at which the MSCF subsystem sets the FLOW state (see the “HIPLEX MSCF” manual [8]). The number of executed and rejected send jobs is also shown.

For send jobs of type REQUEST WITH REPLY, the average total times of the job and the wait times for the first reply from the receiving host are displayed. The average time the system spent in the FLOW state is also shown, whereby only the transitions from FLOW to NO FLOW are taken into account.

The latter data is event-driven and is described by the structure SM2GDAT_mscf_time_count_mdl, where “average_time” already contains an average time. “count” specifies how often the event has occurred.

Descriptive structures: SM2GDAT_mscf_single_mdl
SM2GDAT_mscf_time_count_mdl

Displacement from the beginning of the BUFFER HEADER: “fixed_part_dsp” from the BUFFER HEADER



Corresponding screen report in the online component: MSCF
Monitoring program: MSCF

Structure of the NSM data buffer

The NSM data buffer contains information about the token and locks in the host network.

The data area of the NSM data buffer comprises an area of fixed length.

Data buffer type: 1

Fixed-length data area

The fixed-length data area contains the number of hosts in the host network and information about the token and locks.

Descriptive structures:	SM2GDAT_nsm_single_mdI
	SM2GDAT_tokenab_mdI
	SM2GDAT_tokenoptab_mdI
	SM2GDAT_lockconvtab_mdI
	SM2GDAT_enqlockmodetab_mdI
	SM2GDAT_lockoptypetab_mdI
	SM2GDAT_lockservicetab_mdI

Displacement from the beginning of the BUFFER HEADER:	„fixed_part_dsp“ from the BUFFER HEADER
---	---



Corresponding screen report in the online component: NSM
Monitoring program: NSM

Structure of the OPENFT data buffer

The OPENFT data buffer contains information about openFT instances.

The data area of the OPENFT data buffer comprises only of repeat groups.

Data buffer type: 2

Repeat groups

A repeat group exists for each monitored openFT instance.

Descriptive structure:	SM2GDAT_openft_multiple_mdI
Displacement of the first repeat group from the beginning of the BUFFER HEADER:	„first_group_dsp“ from the BUFFER HEADER
Length of a repeat group:	„length_group“ from the BUFFER HEADER
Number of repeat groups:	„number_groups“ from the BUFFER HEADER



Corresponding screen report in the online component: OPENFT
Monitoring program: OPENFT

Structure of the PERIODIC-TASK data buffer

The PERIODIC-TASK data buffer supplies utilization data for all tasks.

The data area of the PERIODIC-TASK data buffer comprises an area of fixed length and a number of repeat groups.

Data buffer type: 3

Fixed-length data area

The fixed-length data area contains the number of supplied repeat groups to follow (used_groups). In addition, the number of existing repeat groups is also supplied here (number_groups). This value is, however, not required for analyzing the data buffer and is used only for internal purposes.

Descriptive structure: SM2GDAT_pertask_single_mdl

Displacement from the beginning
of the BUFFER HEADER: "fixed_part_dsp" from the BUFFER HEADER

Repeat groups

A repeat group exists for each task. The repeat group provides information on service units, CPU time, I/Os, UPG (used page count) and PAGE READs. In addition, it is indicated whether the data of this task is written to the output file. In this case, the "task_in_file" bit is set.

The caller without the SWMONADM privilege receives only the data of tasks of his or her own ID.

Descriptive structure: SM2GDAT_pertask_multiple_mdl

Displacement of the first repeat group
from the beginning of the BUFFER
HEADER: "first_group_dsp" from the BUFFER
HEADER

Length of a repeat group: "length_group" from the BUFFER HEADER

Number of repeat groups: "used_groups" from
SM2GDAT_pertask_single_mdl



Corresponding screen report in the online component: PERIODIC TASK
Monitoring program: PERIODIC-TASK

Structure of the PFA data buffer

The PFA data puffer provides monitored data on caches under User PFA that are used by DAB.

The data area of the PFA data buffer comprises an area of fixed length and two types of repeat group.

There is a chain of BCB repeat groups (BCB stands for buffer control block); these are not arranged one after the other in the data buffer, as is the case with most repeat groups; instead, the displacement from the next BCB repeat group is specified in each BCB repeat group.

Data buffer type: 3

Fixed-length data area

The fixed-length data area contains the number of BCB repeat groups and the displacement from the first BCB repeat group. The following elements are no longer of importance because the PFA monitoring program no longer supplies monitored data for Symmetrix systems:

total_number_cubcb_groups
length_cubcb_group

Descriptive structure: SM2GDAT_pfa_single_mdI

Displacement from the beginning of
the BUFFER HEADER: "fixed_part_dsp" from the BUFFER HEADER

The BCB repeat group

For each cache there is a BCB repeat group.

The BCB repeat groups are not arranged one after the other in the data buffer. The "next_bcb_group_dsp" element contains the displacement from the next BCB repeat group. The following elements are no longer of importance because the PFA monitoring program no longer supplies monitored data for Symmetrix systems:

number_cubcb_groups
first_cubcb_group

The report_type element now only has the value SM2GDATpfa.

Each BCB repeat group contains information (e.g. ID, driver, medium, size) on the buffer and monitoring data (e.g. read and write accesses).

Descriptive structure:	SM2GDAT_pfa_bcb_md1
Displacement of the first repeat group from the beginning of the structure SM2GDAT_pfa_single_md1:	“first_bcb_group_dsp” from SM2GDAT_pfa_single_md1
Displacement of the next repeat group from the beginning of the BCB repeat group:	“next_bcb_group_dsp” from SM2GDAT_pfa_bcb_md1
Number of repeat groups:	“number_bcb_groups” from SM2GDAT_pfa_single_md1

CUBCB repeat group

The CUBCB repeat group is no longer of importance because the monitoring program no longer records monitored data for Symmetrix systems.



Corresponding screen report in the online component: PFA
Monitoring program: PFA

Structure of the POSIX data buffer

The POSIX data buffer supplies data on the use of various functions of the POSIX subsystem.

The data area of the POSIX data buffer comprises an area of fixed length.
Data buffer type: 1

Fixed-length data area

The monitored data is supplied in dual form: as a absolute count for the interval and as a count per second for the interval. At high load peaks, the absolute counts can overflow and thus supply incorrect values. It is advisable to work with the time-related data.

The supplied data corresponds to that of the sar command in UNIX.

Descriptive structure:	struct SM2GDAT_posix_single_md1
Displacement from the beginning of the BUFFER HEADER:	“fixed_part_dsp” from the BUFFER HEADER



Corresponding screen report in the online component: ACTIVITY, POSIX
Monitoring program: POSIX

Structure of the PUBSET data buffer

The PUBSET data buffer supplies data on the utilization of imported pubsets. Data for SF pubsets and volume sets is supplied.

The data area of the PUBSET data buffer comprises an area of fixed length and a number of repeat groups.

Data buffer type: 3

Fixed-length data area

The fixed-length data area contains only the number of the following repeat groups

Descriptive structure: SM2GDAT_pubset_single_mdI

Displacement from the beginning
of the BUFFER HEADER: "fixed_part_dsp" from BUFFER HEADER

Repeat groups

For each SF pubset and volume set a repeat group exists with

- information on the pubset
CATID, name of the SM pubset in the case of volume sets, saturation level
- monitored data
capacity and utilized storage area of the pubset

Descriptive structure: SM2GDAT_pubset_multiple_mdI

Displacement of the first repeat group
from the beginning of the BUFFER
HEADER: "first_group_dsp" from BUFFER HEADER

Length of a repeat group: "length_group" from BUFFER HEADER

Number of repeat groups: "used_groups" from
SM2GDAT_pubset_single_mdI



Corresponding screen report in online part: PUBSET
Monitoring program: PUBSET

Structure of the RESPONSETIME data buffer

The RESPONSETIME data buffer supplies data on response, think, transaction and message wait times. The monitored data is supplied bucket-specifically and/or category-specifically for defined connection sets.

The data area of the RESPONSETIME data buffer comprises an area of fixed length and three types of repeat group, which are arranged in three chains.

Data buffer type: 4

Fixed-length data area

The fixed-length data area contains for each of the three chains of repeat groups the length of a repeat group, the number of repeat groups and the displacement from the first repeat group.

The number of buckets and the data for each bucket are output for the different response time definitions.

The “connection” element (“SM2GDAT_rtime_connection_mdl” structure type) contains the defined connection groups of which the connection sets are composed.

The “connection_set” element (“SM2GDAT_rtime_connectionset_mdl” structure type) contains the definition of the connection sets, particularly how the connection sets are composed of the connection groups. A connection set consists of one or more (up to five) connection groups. The connection set can be defined either positively (i.e. all the groups specified in this connection set are monitored) or negatively (i.e. all groups are monitored except those specified in the connection set).

Each bit of the “connection_set.define_set.connection.selection” element corresponds to a connection group of the “connection” element: a set bit means that the connection group belongs to the connection set. The “connection_set.indication.indicate.positive” bit specifies whether the connection set is defined positively or negatively.

Descriptive structures:	SM2GDAT_rtime_single_mdl
	SM2GDAT_rtime_connection_mdl
	SM2GDAT_rtime_connectionset_mdl

Displacement from the beginning of the BUFFER HEADER:	“fixed_part_dsp” from the BUFFER HEADER
--	---

The CATEGORY LIST repeat group

CATEGORY LIST repeat groups are only written when the response time data is collected category-specifically.

A repeat group exists for each category monitored.

Each repeat group contains the name of the category monitored.

The repeat groups come one after the other in the data buffer

No repeat group is written for the SUM category.

Descriptive structure:	SM2GDAT_rtime_category_list_mdl
Displacement of the first repeat group from the beginning of the structure SM2GDAT_rtime_single_mdl:	“first_category_name_group_dsp” from SM2GDAT_rtime_single_mdl
Length of a repeat group:	“length_category_name_group” from SM2GDAT_rtime_single_mdl
Number of repeat groups:	“number_category_name_groups” from SM2GDAT_rtime_single_mdl

The BUCKET repeat group

BUCKET repeat groups are only written when the response time data is collected bucket-specifically. A repeat group exists for each connection set.

The repeat groups come one after the other in the data buffer and are assigned to the connection sets as follows:

the nth repeat group contains the data for the connection set that is in the nth position in the “connection_set” element of the “SM2GDAT_rtime_single_mdl” structure.

A repeat group provides information on how many interactions have occurred in the individual areas (buckets). The number and limits of the buckets are described in the “SM2GDAT_rtime_single_mdl” structure.

In addition, the time sums up to the largest bucket and the overflow values are output.

Descriptive structure:	SM2GDAT_rtime_bucket_mdl
Displacement of the first repeat group from the beginning of the structure SM2GDAT_rtime_single_mdl:	“first_bucket_group_dsp” from SM2GDAT_rtime_single_mdl
Length of a repeat group:	“length_bucket_group” from SM2GDAT_rtime_single_mdl
Number of repeat groups:	“number_bucket_groups” from SM2GDAT_rtime_single_mdl

The CATEGORY repeat group

CATEGORY repeat groups are written only when the response time data is collected category-specifically.

A chain of CATEGORY repeat groups is written for each category (including the “SUM” category). The chains come one after the other in the data buffer. Each chain supplies the category-specific data for all connection sets of a particular category. The chain of the SUM category comes at the end.

The chains with CATEGORY repeat groups are assigned to the categories as follows: the *n*th chain contains the data of the category in the *n*th position in the chain of the CATEGORY LIST repeat groups.

In a chain with CATEGORY repeat groups, the individual repeat groups are assigned to the connection sets as follows: the *n*th repeat group contains the data for the connection set in the *n*th position in the “connection_set” of the “SM2GDAT_rtime_single_mdI” structure.

A CATEGORY repeat group provides information on time sums and the number of interactions with and without overflow values.

Descriptive structure: SM2GDAT_rtime_category_mdI

Displacement of the first repeat group
from the beginning of the structure SM2GDAT_rtime_single_mdI: “first_category_group_dsp” from SM2GDAT_rtime_single_mdI

Length of a repeat group: “length_category_group” from SM2GDAT_rtime_single_mdI

Number of repeat groups: “number_category_groups” from SM2GDAT_rtime_single_mdI



Corresponding screen report in the online component:
RESPONSETIME (BUCKET, CATEGORY)
Monitoring program: RESPONSETIME

Structure of the SCHANNEL data buffer

The SCHANNEL data buffer contains data on the activity of input/output channels and of the processors. It supplies monitoring data for each channel.

The data area of the SCHANNEL data buffer consists only of repeat groups.
Data buffer type: 2

Repeat groups

A repeat group is written for each channel. The repeat group provides information on how often the channel was active at the same time as at least one processor and on how often the channel was active when no processor was active at the same time.

In addition, the repeat group provides information on whether the IOPT function is active. If required, the channel utilization is recorded by tasks of the various priority classes. For information on IOPT, see the IORM utility in the “Utility Routines” manual [1].

The data of the SCHANNEL data buffer is obtained by sampling. The number of samples is contained in “samples” in the “SM2GDAT_basic_single_md1” structure (BASIC data buffer). You should bear in mind that the accuracy of the data varies with the number of samples.

Descriptive structure:	SM2GDAT_schannel_multiple_md1
Displacement of the first repeat group from the beginning of the BUFFER HEADER:	“first_group_dsp” from the BUFFER HEADER
Length of a repeat group:	“length_group” from the BUFFER HEADER
Number of repeat groups:	“number_groups” from the BUFFER HEADER



Corresponding screen report in the online component: CHANNEL
Monitoring program: –

Structure of the SDEVICE data buffer

The SDEVICE data buffer supplies data about input/output operations on peripheral devices. Also the monitored data for communication devices are delivered. There are only ever read or write accesses for communication devices.

The data area of the SDEVICE data buffer consists only of repeat groups.
Data buffer type: 2

Repeat groups

A repeat group is written for each device. The repeat group provides information on how often the device was active on account of DMS or paging input/output operations and on the mean queue length.

This monitoring data is obtained by sampling. The number of samples is contained in “samples” in the “SM2GDAT_basic_single_mdl” structure (BASIC data buffer). You should note that the accuracy of the data varies with the number of samples.

In addition, the number of I/Os and the volume of data transferred are output.

In addition, the repeat group provides information on whether the IOPT function is active (see the IORM utility in the “Utility Routines” manual [1]). If required, the channel utilization is recorded by tasks of the various priority classes.

Monitored data for GS (global storage) is not delivered here (see the GS, GSVOL data buffer).

Descriptive structure:	SM2GDAT_sdevice_multiple_mdl
Displacement of the first repeat group from the beginning of the BUFFER HEADER:	“first_group_dsp” from the BUFFER HEADER
Length of a repeat group:	“length_group” from the BUFFER HEADER
Number of repeat groups:	“number_groups” from the BUFFER HEADER



When parallel access volumes (PAVs) are used, the data of the basic and alias devices is output. The data of basic devices relates to the basic device and the assigned alias devices (summation or averaging of the monitored data). When an alias device changes its basic device, the data of the alias device is not included in the data of the basic device in the monitoring cycle in which the change takes place.

Corresponding screen report in the online component:
DEVICE DISK, DEVICE OTHER, DEVICE TAPE, DEVICE TD
Monitoring program: SAMPLING-DEVICE

Structure of the SESAM-SQL data buffer

The SESAM-SQL data buffer supplies data on SESAM/SQL data base systems.

The data area of the SESAM-SQL data buffer comprises an area of fixed length and a number of repeat groups.

Data buffer type: 3

Fixed-length data area

The fixed-length data area contains the number of the following repeat groups

Descriptive structure: SM2GDAT_sesam_sql_single_mdl

Displacement from the beginning
of the BUFFER HEADER: "fixed_part_dsp" from BUFFER HEADER

Repeat groups

A repeat group is written for each monitored database system.

All monitored data with the unit "Number" applies for the monitoring cycle specified in "elapsed_time". Because of the asynchronous data collection (see [section "Privileged SM2 monitoring programs" on page 52](#)), the "elapsed_time" can differ from the variable of the same name in the BASIC data buffer.

Descriptive structure: SM2GDAT_sesam_sql_multiple_mdl

Displacement of the first repeat group
from the beginning of the BUFFER
HEADER: "first_group_dsp" from BUFFER HEADER

Length of a repeat group: "length_group" from BUFFER HEADER

Number of repeat groups: "used_groups" from
SM2GDAT_sesam_sql_single_mdl



Corresponding screen report in online part: SESAM SQL
Monitoring program: SESAM-SQL

Structure of the STORAGE-SYSTEM data buffer

The STORAGE-SYSTEM data buffer supplies measured data on Symmetrix systems.

The data area of the STORAGE-SYSTEM data buffer comprises an area of fixed length and up to six different types of repeat groups.

Data buffer type: 4

Fixed-length data area.

For the chain of the repeat group STORAGE-SYSTEM-HEADER the fixed-length data area contains the length of the repeat group, the number of repeat groups and the displacement from the first repeat group.

Descriptive structure: SM2GDAT_sts_single_mdl

Displacement from the beginning

of the BUFFER HEADER: "fixed_part_dsp" from BUFFER HEADER

Repeat group STORAGE-SYSTEM-HEADER

The repeat group STORAGE-SYSTEM-HEADER exists for each type of storage system.

As currently only Symmetrix systems are supported, there is only one repeat group.

For the chain of the repeat group SYMMETRIX the repeat group STORAGE-SYSTEM-HEADER contains the length of the repeat group, the number of repeat groups and the displacement from the first repeat group..

Descriptive structure: SM2GDAT_sts_header_mdl

Displacement of the first repeat group
from the beginning of the structure
SM2GDAT_sts_single_mdl: "first_header_group_dsp" from
SM2GDAT_sts_single_mdl

Length of a repeat group: "length_header_group" from
SM2GDAT_sts_single_mdl

Number of repeat groups: "number_header_groups" from
SM2GDAT_sts_single_mdl

Repeat group SYMMETRIX

For each Symmetrix system a repeat group exists with

- information on the Symmetrix system
Symmetrix serial number, size of the cache
- monitored data
number of read and write accesses, number of accesses which were completed by the cache, and number of blocks transferred

For the devices (logical volumes), the physical disks in the Symmetrix system and the directors additional repeat groups exist whose length, number and displacement from the first repeat group are also contained in the Symmetrix repeat group.

The repeat groups are written only if the data is selected when the monitoring program is defined.

Descriptive structure:	SM2GDAT_sts_sym_group_mdl
Displacement of the first repeat group from the beginning of the structure SM2GDAT_sts_header_mdl:	“first_sts_group_dsp” from SM2GDAT_sts_header_mdl
Length of a repeat group:	“length_sts_group” from SM2GDAT_sts_header_mdl
Number of repeat groups:	“number_sts_groups” from SM2GDAT_sts_header_mdl

Repeat group DEVICE

The repeat groups exist only if the recording of monitoring data for devices (logical volumes) is specified when the monitoring program is defined.

For each recorded device (logical volume) a DEVICE repeat group exists with

- information on the device
unique internal Symmetrix number of the device, the mnemonic device name and the VSN, and the ID of the pubset only if the device is “attached”
- monitored data
number of read and write accesses, number of accesses completed by the cache, and number of blocks transferred.

Descriptive structure: SM2GDAT_sts_sym_dev_group_md1

Displacement of the first repeat group from the beginning of the structure SM2GDAT_sts_sym_group_md1: “first_device_group_dsp” from SM2GDAT_sts_sym_group_md1

Length of a repeat group: “length_device_group” from SM2GDAT_sts_sym_group_md1

Number of repeat groups: “number_device_groups” from SM2GDAT_sts_sym_group_md1



The internal Symmetrix device name is also available in a length of 7 characters in the new “device_ext” variable. When the length of the device name is less than 4, the original “device” variable then only contains the first 4 characters of the original device name.

Repeat group BACK-END

The repeat groups exist only if the recording of monitoring data for devices (logical volumes), and here especially the recording of the back-end data, is specified when the monitoring program is defined.

A BACK-END repeat group with monitoring data (the number of read and write accesses and the number of transferred blocks) exists for each recorded device (logical volume). The back-end data provides information on the activity between the device (logical volume) concerned and the physical drives in the Symmetrix system. However, the data does not establish a correlation to individual physical drives.

No information on the device (logical volume) exists in this repeat group. This can be taken from the relevant DEVICE repeat group. Here the i-th BACK-END repeat group corresponds to the i-th DEVICE repeat group..

Descriptive structure: SM2GDAT_sts_sym_b_end_group_md1

Displacement of the first repeat group from the beginning of the structure SM2GDAT_sts_sym_group_md1:	“first_device_b_end_dsp” from SM2GDAT_sts_sym_group_md1
Length of a repeat group:	“length_device_b_end_group” from SM2GDAT_sts_sym_group_md1
Number of repeat groups:	“number_device_b_end_groups” from SM2GDAT_sts_sym_group_md1

Repeat group DISK

The repeat groups exist only if the recording of monitoring data for physical drives (physical disks) is specified when the monitoring program is defined.

For each physical drive a DISK repeat group exists with

- information on the physical drive
name of the physical drive
- monitored data
number of read and write accesses and number of transferred blocks.

Descriptive structure:	SM2GDAT_sts_sym_disk_group_md1
Displacement of the first repeat group from the beginning of the structure SM2GDAT_sts_sym_group_md1:	“first_disk_group_dsp” from SM2GDAT_sts_disk_group_md1
Length of a repeat group:	“length_disk_group” from SM2GDAT_sts_disk_group_md1
Number of repeat groups:	“number_disk_groups” from SM2GDAT_sts_disk_group_md1

Repeat group DIRECTOR

The repeat groups exist only if the recording of monitoring values for directors is specified when the monitoring program is defined.

For each director a DIRECTOR repeat group exists with

- information on the director number and type
- monitored data number of read and write accesses and number of transferred blocks (only RA director), and total number of accesses completed by the cache.

Descriptive structure:	SM2GDAT_sts_sym_director_group_mdl
Displacement of the first repeat group from the beginning of the structure SM2GDAT_sts_sym_group_mdl:	"first_director_group_dsp" from SM2GDAT_sts_director_group_mdl
Length of a repeat group:	"length_director_group" from SM2GDAT_sts_director_group_mdl
Number of repeat groups:	"number_director_groups" from SM2GDAT_sts_director_group_mdl



Corresponding screen report in online part:
 SYMMETRIX, SYMMETRIX PUBSET, SYMMETRIX DEVICE
 Monitoring program: STORAGE-SYSTEM

Structure of the SVC data buffer

The SVC data buffer provides information on the number of SVC calls.

The data area of the SVC data buffer comprises only repeat groups.

Data buffer type: 2

Repeat groups

A repeat group exists for each SVC.

The number of SVC calls is collected separately for TU and TPR calls.

Descriptive structure:	SM2GDAT_svc_multiple_mdI
Displacement of the first repeat group from the beginning of the BUFFER HEADER:	"first_group_dsp" from the BUFFER HEADER
Length of a repeat group:	"length_group" from the BUFFER HEADER
Number of repeat groups:	"number_groups" from the BUFFER HEADER



Corresponding screen report in the online component: SVC
Monitoring program: SVC

Structure of the SYSTEM data buffer

For each category, the SYSTEM data buffer supplies information on queues and DMS and paging input/output operations.

The data area of the SYSTEM data buffer comprises an area of fixed length and a number of repeat groups.

Data buffer type: 3

Fixed-length data area

The fixed-length data area contains the number of devices monitored and their names. In the case of the *ALL, *DISK and *TAPE monitoring program definitions, the real number of devices cannot be specified: the "number_used_devices" element contains a one, and the first four characters of the relevant name appear for the first device name.

The number of PAGE WRITES and the associated mean service times are also output. This data cannot be assigned to a category, which is why it is found here. It applies only to the devices included in monitoring.

The number of supplied repeat groups to follow is contained in “used_groups”.

Descriptive structure:	SM2GDAT_system_single_mdl
Displacement from the beginning of the BUFFER HEADER:	“fixed_part_dsp” from the BUFFER HEADER

Repeat groups

The first repeat group contains the sum values of all categories. SUM is output as “category_name”. The “category_number” is undefined. Each subsequent repeat group contains the monitoring data for a single category.

Each repeat group provides information on the CPU and system time. In multiprocessor servers, the sum of the times across all categories can be greater than the “elapsed_time” in the SM2GDAT_basic_single_mdl structure.

The number of DMS/paging I/Os and the associated service times refer only to the devices specified in the SM2GDAT_system_single_mdl structure in the “device_name” element. In the case of paging I/O, only the assignable I/Os are supplied (see above). In the case of the SUM category, the sum is formed on the basis of all I/Os and times.

For the queues, the number of tasks in the queue at the end of the monitoring cycle or the number of tasks that left the queue during the monitoring cycle is output.

Descriptive structures:	SM2GDAT_system_multiple_mdl SM2GDAT_system_queue_data_mdl
Displacement of the first repeat group from the beginning of the BUFFER HEADER:	“first_group_dsp” from the BUFFER HEADER
Length of a repeat group:	“length_group” from the BUFFER HEADER
Number of repeat groups:	“used_groups” from SM2GDAT_system_single_mdl



Corresponding screen report in the online component: CATEGORY
Monitoring program: SYSTEM

Structure of the TCP-IP data buffer

The TCP-IP data buffer provides monitored data on each TCP/IP connection of the local host.

The data area of the TCP-IP data buffer comprises an area of fixed length and a number of repeat groups.

Data buffer type: 3

Fixed-length data area

The fixed-length data area contains the number of repeat groups to follow.

Descriptive structure: SM2GDAT_tcp_ip_single_mdl

Displacement from the beginning of the BUFFER HEADER: "fixed_part_dsp" from the BUFFER HEADER

Repeat groups

A repeat group is written for each connection. IPv4 as well as IPv6 connections are reported. Depending on the version (ip-version variable), the variables of the local or remote IP address are only partially (IPv4) or fully (IPv6) filled. The repeat group provides the number of data packets sent and received, the cache area occupied by messages to be sent and received, and the size of the last window sent or received.

Descriptive structure: SM2GDAT_tcp_ip_multiple_mdl
SM2GDAT_local_ip_addr_mdl
SM2GDAT_remote_addr_mdl

Displacement of the first repeat group from the beginning of the BUFFER HEADER: "first_group_dsp" from the BUFFER HEADER

Length of a repeat group: "length_group" from the BUFFER HEADER

Number of repeat groups: "used_groups" from SM2GDAT_tcp_ip_single_mdl



Corresponding screen report in the online component: TCP-IP
Monitoring program: TCP-IP

Structure of the TIME-IO data buffer

The TIME-IO data buffer provides an overview of important system activities. It supplies monitoring data for each logical machine.

The data area of the TIME-IO data buffer comprises an area of fixed length and a number of repeat groups.

Data buffer type: 3

Fixed-length data area

The fixed-length data area contains monitoring data that is not specific to any particular logical machine. The number of CPU reconfigurations and active logical machines is output here.

Descriptive structure:	SM2GDAT_time_io_single_mdI
Displacement from the beginning of the BUFFER HEADER:	“fixed_part_dsp” from the BUFFER HEADER

Repeat groups

The first repeat group contains for all logical machines mean values for the TU, TPR, SIH and IDLE times and cumulative values for the other monitoring data. Each subsequent repeat group contains the monitoring data for a single logical machine. Consequently, when there is only one logical machine, there is only one repeat group. A repeat group contains the CPU time in the various function statuses, the number of SVC calls and the number of input/output operations. The number of the logical machine is provided in printable form (beginning with '00'). The first repeat group always contains AV for the number of the logical machine.

On servers with x86 architecture the following applies: Only in the first repeat group the element “other_io” contains the operations on devices of the family type TD (including openCRYPT-BOX devices).



The “samples_tu” and “samples_tpr” elements serve as indicators of the accuracy of the TU and TPR time – these two function statuses are recorded by SM2 by means of sampling. The accuracy of the two times varies, depending on the number of samples taken.

Descriptive structure:	SM2GDAT_time_io_multiple_mdI
Displacement of the first repeat group from the beginning of the BUFFER HEADER:	“first_group_dsp” from the BUFFER HEADER
Length of a repeat group:	“length_group” from the BUFFER HEADER
Number of repeat groups:	“number_groups” from the BUFFER HEADER



Corresponding screen reports in the online component: ACTIVITY, CPU
Monitoring program: –

Structure of the TLM data buffer

The TLM data buffer supplies data on the utilization of locks administered by the Task Lock Manager.

The data area of the TLM data buffer comprises an area of fixed length and a number of repeat groups.

Data buffer type: 3

Fixed-length data area

The number of supplied repeat groups to follow is contained in “used_groups”.

Descriptive structure: SM2GDAT_tlm_single_mdl

Displacement from the beginning of
the BUFFER HEADER: “fixed_part_dsp” from the BUFFER HEADER

Repeat groups

A repeat group exists for each lock. The repeat group provides information on the mean number of tasks in the queue and on how often the lock was identified as occupied.

The data is collected by sampling. The number of samples is contained in the SM2GDAT_tlm_single_mdl structure in the “samples” element. You should bear in mind that the accuracy of the data varies with the number of samples taken.

Descriptive structure: SM2GDAT_tlm_multiple_mdl

Displacement of the first repeat group
from the beginning of the BUFFER
HEADER: “first_group_dsp” from the BUFFER
HEADER

Length of a repeat group: “length_group” from the BUFFER HEADER

Number of repeat groups: “used_groups” from
SM2GDAT_tlm_single_mdl



Corresponding screen report in the online component: TLM
Monitoring program: TLM

Structure of the UDS-SQL data buffer

The UDS-SQL data buffer supplies data on UDS/SQL data base systems.

The data area of the SESAM-SQL data buffer comprises an area of fixed length and a number of repeat groups.

Data buffer type: 3

Fixed-length data area

The fixed-length data area contains the number of the following repeat groups

Descriptive structure: SM2GDAT_uds_sql_single_mdI

Displacement from the beginning
of the BUFFER HEADER: "fixed_part_dsp" from BUFFER HEADER

Repeat groups

A repeat group is written for each monitored database system.

All monitored data with the unit "Number" applies for the monitoring cycle specified in "elapsed_time". Because of the asynchronous data collection (see [section "Privileged SM2 monitoring programs" on page 52](#)), the "elapsed_time" can differ from the variable of the same name in the BASIC data buffer.

Descriptive structure: SM2GDAT_uds_sql_multiple_mdI

Displacement of the first repeat group
from the beginning of the BUFFER
HEADER: "first_group_dsp" from BUFFER HEADER

Length of a repeat group: "length_group" from BUFFER HEADER

Number of repeat groups: "used_groups" from
SM2GDAT_uds_sql_single_mdI



Corresponding screen report in online part: UDS SQL
Monitoring program: UDS-SQL

Structure of the USERFILE data buffer

The USERFILE data buffer contains data on access to user-specifically registered files.

The data area of the USERFILE data buffer comprises an area of fixed length and a number of repeat groups.

Data buffer type: 3

Fixed-length data area

The fixed-length data area contains only the number of repeat groups to follow.

Descriptive structure: SM2GDAT_userfile_single_mdI

Displacement from the beginning
of the BUFFER HEADER: "fixed_part_dsp" from the BUFFER HEADER

Repeat groups

A repeat group is written for each monitored file. The repeat group provides information on the number of read, write, wait and check accesses as well as the average duration of I/Os.

The number of repeat groups that are actually filled ("used_groups") may be smaller than the number of groups that exist ("number_groups"). Equally, the filled repeat groups can be distributed arbitrarily among the existing groups.

Descriptive structure: SM2GDAT_userfile_multiple_mdI

Displacement of the first repeat group
from the beginning of the BUFFER
HEADER: "first_group_dsp" from the BUFFER
HEADER

Length of a repeat group: "length_group" from the BUFFER HEADER

Number of repeat groups: "used_groups" from
SM2GDAT_userfile_single_mdI



Corresponding screen report in the online component: USER FILE
Monitoring program: FILE (USER MEASUREMENTS)

Structure of the USERISAM data buffer

The USERISAM data buffer supplies data on access to user-specifically registered ISAM pools.

The data area of the USERISAM data buffer comprises an area of fixed length and a number of repeat groups.

Data buffer type: 3

Fixed-length data area

The fixed-length data area contains only the number of repeat groups to follow.

Descriptive structure: SM2GDAT_userisam_single_mdI

Displacement from the beginning of the BUFFER HEADER: "fixed_part_dsp" from the BUFFER HEADER

Repeat groups

A repeat group is written for each monitored ISAM pool. The repeat group provides information on the accesses and the pages used for buffering.

The number of repeat groups that are actually filled ("used_groups") may be smaller than the number of groups that exist ("number_groups"). Equally, the filled repeat groups can be distributed arbitrarily among the existing repeat groups.

Descriptive structure: SM2GDAT_userisam_multiple_mdI

Displacement of the first repeat group from the beginning of the BUFFER HEADER: "first_group_dsp" from the BUFFER HEADER

Length of a repeat group: "length_group" from the BUFFER HEADER

Number of repeat groups: "used_groups" from SM2GDAT_userisam_single_mdI



Corresponding screen report in the online component: USER ISAM
Monitoring program: ISAM (USER MEASUREMENTS)

Structure of the UTM data buffer

The UTM data buffer supplies application-specific data for UTM applications.

The data area of the UTM data buffer comprises only repeat groups.

Data buffer type: 2

Repeat groups

A repeat group exists for each UTM application.

A repeat group is subdivided into constant, periodic and event-driven data. Each of these types of data has its own structures.

In the case of the event-driven data (SM2GDAT_utm_event_mdl structure), each monitoring value is calculated on the basis of a counter (“count”) and a sum element (“sum”). The counter specifies how often the event has occurred.

The sum element contains the sum of the data made available for the events. The quotient of “sum” and “count” is the mean monitoring value.

Descriptive structures:	SM2GDAT_utm_multiple_mdl
	SM2GDAT_utm_constant_mdl
	SM2GDAT_utm_periodic_mdl
	SM2GDAT_utm_event_mdl
	SM2GDAT_utm_avg_mdl

Displacement of the first repeat group from the beginning of the BUFFER HEADER:	“first_group_dsp” from the BUFFER HEADER
---	--

Length of a repeat group:	“length_group” from the BUFFER HEADER
---------------------------	---------------------------------------

Number of repeat groups:	“number_groups” from the BUFFER HEADER
--------------------------	--

With the extension extv3 the repeat group is extended by, among other things, TAC class values. The data of the TAC classes is contained in the repeat groups of the UTM application as TAC class repeat groups (structure SM2GDAT_utm_ext_tacclass_mdl). 16 TAC class repeat groups always exist which are addressed analogously to the UTM application repeat groups:

Descriptive structure:	SM2GDAT_utm_ext_tacclass_mdl
------------------------	------------------------------

Displacement of the first repeat group from the beginning of the MULTIPLE BUFFER (structure SM2GDAT_utm_multiple_mdl):	„ev3_first_tacclass_group_dsp“ aus SM2GDAT_utm_ext_tacclass_mdl
--	---

Length of a repeat group:	„ev3_length_tacclass_group“ aus SM2GDAT_utm_ext_tacclass_mdl
---------------------------	--

Number of repeat groups:	„ev3_number_tacclass_groups“ aus SM2GDAT_utm_ext_tacclass_mdl
--------------------------	---



Corresponding screen report in the online component: UTM, UTM APPLICATION
Monitoring program: UTM

Structure of the VM data buffer

The VM data buffer provides data for the individual virtual machines (VMs) of a VM2000 system.

The data area of the VM data buffer comprises an area of fixed length and a number of repeat groups.

Data buffer type: 3

Fixed-length data area

In addition to the number of repeat groups subsequently supplied with values, the fixed-length data area contains the number of real, available CPUs.

The values for the hypervisor are only available on servers with /390 architecture.

The values for the utilization of all domains, for the BS2000 guest systems and for the CPU percentages available to BS2000 guest systems are only available on servers with x86 architecture.

Displays which show the validity of the data exist for this data.

Descriptive structure: SM2GDAT_vm_single_mdI

Displacement from the beginning of
the BUFFER HEADER: "fixed_part_dsp" from the BUFFER HEADER

Repeat groups

A repeat group exists for each VM. The repeat group provides information on the CPU time utilized and the planned relative share of CPU time. The available main memory and the number of the active, virtual CPUs are also output.

Values for all VMs are output only on the monitor VM. Otherwise, only the values of your own VM are output.

Descriptive structure: SM2GDAT_vm_multiple_mdI

Displacement of the first repeat group
from the beginning of the BUFFER
HEADER: "first_group_dsp" from the BUFFER
HEADER

Length of a repeat group: "length_group" from the BUFFER HEADER

Number of repeat groups: "used_groups" from
SM2GDAT_vm_single_mdI



Corresponding screen report in the online component: VM.
Monitoring program: VM

Structure of the VM-CPU-POOL data buffer

The VM-CPU-POOL data buffer supplies data on the utilization of the CPU pool.

The data area of the VM-CPU-POOL data buffer comprises an area of fixed length and a number of repeat groups.

Data buffer type: 3

Fixed-length data area

The fixed-length data area contains only the number of the following repeat groups

Descriptive structure: SM2GDAT_vm_cpupool_single_mdI

Displacement from the beginning

of the BUFFER HEADER: "fixed_part_dsp" from BUFFER HEADER

Repeat groups

For each CPU pool a repeat group exists with

- information on the CPU pool
name, number of available real CPUs (also extra CPUs)
- monitored data
The values for the hypervisor are only available on servers with /390 architecture.
The values for the utilization of all domains, for the BS2000 guest systems and for the CPU shares available to BS2000 guest systems are only available on servers with x86 architecture.
Displays which show the validity of the data exist for this data.

Values for all CPU pools are output only on the monitor VM. Otherwise, only the values for the CPU pool of your VM are output.

Descriptive structure: SM2GDAT_vm_cpupool_multiple_mdI

Displacement of the first repeat group
from the beginning of the BUFFER
HEADER: "first_group_dsp" from BUFFER HEADER

Length of a repeat group: "length_group" from BUFFER HEADER

Number of repeat groups: "used_groups" from
SM2GDAT_vm_cpupool_single_mdI



Corresponding screen report in online part: VM CPU POOL
Monitoring program: VM

Structure of the VM-GROUP data buffer

The VM-GROUP data buffer supplies data on the utilization of VM groups (servers with /390 architecture only).

The data area of the VM-GROUP data buffer comprises an area of fixed length and a number of repeat groups.

Data buffer type: 3

Fixed-length data area

The data area of fixed length supplies the number of available real CPUs.

The number of supplied repeat groups to follow is contained in 'used_groups'.

Descriptive structure: SM2GDAT_vm_group_single_mdl

Displacement from the beginning

of the BUFFER HEADER: "fixed_part_dsp" from the BUFFER HEADER

Repeat groups

A repeat group exists for each VM group. The repeat group provides information on the CPU time utilized, the planned relative share and the maximum share of CPU time.

Values for all VM groups are output only on the monitor VM. Otherwise, only the values of your own VM group are output.

Descriptive structure: SM2GDAT_vm_group_multiple_mdl

Displacement of the first repeat group
from the beginning of the BUFFER
HEADER:

"first_group_dsp" from the BUFFER
HEADER

Length of a repeat group:

"length_group" from the BUFFER HEADER

Number of repeat groups:

"used_groups" from
SM2GDAT_vm_group_single_mdl



Corresponding screen report in the online component: VM GROUP
Monitoring program: VM

11.2 Program interface for querying system utilization

Application area: SM2 information macro

Macro type: S type (E form/L form/C form/D form)

Macro description

The **PFMON** macro provides information on the utilization of the overall system. Shares of CPU time (subdivided into idle, TU, TPR and SIH times) and the number of I/O operations for disks, magnetic tapes and other devices are output. The list of operands begins with the standard header. The user can generate a DSECT or data list for the output area.

Macro call format and operand description

Operation	Operands
PFMON	$MF = \left[\begin{array}{l} E, PARAM = \left[\begin{array}{l} adr \\ (r) \end{array} \right] \\ L \\ C[, PREFIX = prefix][, MACID = macid] \\ D[, PREFIX = prefix] \end{array} \right]$

MF	Describes the form of the macro call (see also the “Executive Macros” manual [4]).
=E,...	E form: only the command part of the macro is generated. addr = symbolic address of the operand list. r = register with the address value “addr”.
=L	L form: the standard header and the output area are generated. The output area is not structured.
=C,...	C form: the standard header and a data list for the output area are generated. prefix = prefix (one character) with which the symbolic names in the data list begin. Default: prefix=P. macid = 1 – 3 characters that replace the second to fourth characters in the symbolic names of the data list; Default: macid=FD.

=D,...

D form: a DSECT is generated for the standard header and output area.

prefix = prefix (one character) with which the symbolic names in the DSECT begin. Default: prefix=P.

Layout of the output area

PFDDATA	DS	OF	START OF DATA AREA
PFDETIM	DS	F	ELAPSED TIME IN 1/300 SEC
PFDITIM	DS	F	IDLE TIME IN 0.1 MSEC
PFD1TIM	DS	F	TU TIME IN 0.1 MSEC
PFD2TIM	DS	F	TPR TIME IN 0.1 MSEC
PFD3TIM	DS	F	SIH TIME IN 0.1 MSEC
PFDPAGW	DS	F	PAGING I/O
PFDDISW	DS	F	DISK I/O
PFDTAPW	DS	F	TAPE I/O
PFDMISW	DS	F	OTHER I/O
PF DSTIM	DS	CL6	SNAP TIME IN HHMMSS
PFDDLEN	EQU	*-PFDDATA	L'DATA-AREA

Explanation of the fields:


Field name	Meaning
PFDDATA	Specifies the beginning of the output area
PFDETIM	Duration of the last SM2 (ONLINE) monitoring cycle (unit = 1/300 seconds)
PFDITIM	Idle time in the last SM2 (ONLINE) monitoring cycle (unit = 100 microseconds)
PF D1TIM	TU time in the last SM2 (ONLINE) monitoring cycle (unit = 100 microseconds)
PF D2TIM	TPR time in the last SM2 (ONLINE) monitoring cycle (unit = 100 microseconds)
PF D3TIM	SIH time in the last SM2 (ONLINE) monitoring cycle (unit = 100 microseconds)
PFDPAGW	PFDPAGW
PFDDISW	Number of DMS disk I/O operations in the last SM2 (ONLINE) monitoring cycle
PFDTAPW	Number of DMS tape I/O operations in the last SM2 (ONLINE) monitoring cycle
PFDMISW	Number of I/O operations on other devices in the last SM2 (ONLINE) monitoring cycle
PF DSTIM	Time stamp at the moment at which the monitoring data is updated by SM2 (End of the relevant SM2 (ONLINE) monitoring cycle) Format: hhmmss (hh=hours, mm=minutes, ss=seconds)

Return code and error indicators

Register R1 contains the address of the parameter list.

A return code is supplied in the standard header relating to execution of the **PFMON** macro. The values specify hexadecimal constants. When MF=C/D applies, the structure of the return code is explained and equates are offered for the return codes.

Return
code:



cc	bb	aaaa	Meaning
00	00	0000	Normal execution
00	00	0004	Data supplied, but SM2 subsystem "in delete"
00	01	001C	Data area not accessible
00	01	FFFF	Function not supported
00	03	FFFF	Interface version not supported
00	04	FFFF	Alignment error
00	40	0004	SM2 not active; no monitoring data available
FF	FF	FFFF	Standard header not accessible



- In multiprocessor systems, the time monitoring data that is output for the function statuses represents mean values of the data obtained on the individual processors.
- The number data for the I/O operations represents the sums of the data obtained on the individual processors.
- Under VM2000 there may be large differences between the sum of the IDLE, TU, TPR and SIH times and the elapsed time on account of the CPU allocation to the different guest systems.

12 Variables reports

12.1 Table of variables

The following table presents the variables monitored by SM2 in association with monitoring program, ONLINE report and SM2R1 report group.

Monitored variable	Monitoring program	ONLINE report	SM2R1 report group
Cache			
Accesses and hit rates for DAB caches	DAB	DAB	DAB
Accesses to ISAM pools	ISAM	ISAM, ISAM-FILE	ISAM, ISAM-FILE
Accesses and hit rates for hiperfiles	PFA	PFA CACHE	PFA
Accesses and hit rates for Symmetrix systems	STORAGE-SYSTEM	SYMMETRIX/ -PUBSET/ -DEVICE	STORAGE-SYSTEM-SYMMETRIX
CPU			
CPU utilization	¹⁾	ACTIVITY, CPU	CPU
Number of system calls	¹⁾ , SVC	ACTIVITY, SVC	CPU, SVC
Files			
Catalog accesses	CMS	CMS	CATALOG-MANAGEMENT
File accesses	FILE	FILE	FILE
File access times	FILE	FILE	FILE
File accesses for selected disks	DISK-FILE	DISK FILE	

Table 18: Monitored variables in relation to the monitoring program, ONLINE report, and SM2R1 report group (part 1 of 4)

Monitored variable	Monitoring program	ONLINE report	SM2R1 report group
Databases			
Utilization in the SESAM/SQL database system	SESAM-SQL	SESAM SQL	SESAM-SQL
Utilization in the UDS/SQL database system	UDS-SQL	UDS SQL	UDS-SQL
IOs			
Number of I/O operations	SAMPLING-DEVICE	ACTIVITY, DEVICE DISK, DEVICE OTHER, DEVICE TAPE, DEVICE TD	DISK, IO, DEVICE
Number of paging I/O operations	¹⁾	ACTIVITY, MEMORY	IO
IO priorities for disks and channels	³⁾ , SAMPLING- DEVICE	-----	⁴⁾
Channel utilization and channel data transfer rates	CHANNEL-IO	CHANNEL	CHANNEL
Device utilization	SAMPLING-DEVICE	DEVICE DISK, DEVICE OTHER, DEVICE TAPE	DEVICE, DISK
Data transfer rates	SAMPLING-DEVICE	DEVICE DISK, DEVICE OTHER, DEVICE TAPE, DEVICE TD	DEVICE, DISK
Length of device queues	SAMPLING-DEVICE	DEVICE DISK, DEVICE OTHER, DEVICE TAPE	DEVICE, DISK
Duration of I/O operations	SAMPLING DEVICE	DEVICE DISK	DISK
Hardware service times	SERVICETIME SAMPLING DEVICE	----- DEVICE DISK	DEVICE, DISK

Table 18: Monitored variables in relation to the monitoring program, ONLINE report, and SM2R1 report group
(part 2 of 4)

Monitored variable	Monitoring program	ONLINE report	SM2R1 report group
Communication			
Response times, transaction times, think times, transaction rates, message lengths of transactions	RESPONSE-TIME	RESPONSE-TIME	RESPONSE-TIME
Number of I/Os by network devices	SAMPLING-DEVICE	DEVICE TD	DEVICE
Network data transfer rates	BCAM-CONNECTION	BCAM-CONNECTION	BCAM-CONNECTION
Data transfer rates for TCP-IP connection	TCP-IP	TCP/IP	TCP-IP
Network data transfer rates, memory for data transfer	BCAM-CONNECTION	BCAM CONNECTION, BCAM MEMORY	BCAM-CONNECTION, BCAM-MEMORY
Memory / disk storage			
Allocation of main memory	1)	ACTIVITY, MEMORY	MEMORY
Allocation of paging area	1)	ACTIVITY, MEMORY	MEMORY
Allocation of virtual address space	1)	ACTIVITY, MEMORY	MEMORY
Page fault rate	1)	MEMORY	MEMORY
Accesses to global memory	GS, GSVOL	GS, GSVOL	GS, GSVOL
Utilization of pubsets	PUBSET	PUBSET	PUBSET
Subsystems			
PCS data	1)	PCS	PCS
UTM response times and transaction rates	UTM	UTM/ APPLICATION	UTM
Number and duration of send requests via MSCF	MSCF	MSCF	MSCF
Data to POSIX	POSIX	POSIX	POSIX
Lock requests to the DLM	DLM	DLM	DLM
Data to HIPLEX synchro. functions	NSM	NSM	NSM
Data to HSMS	HSMS	-----	(HSMS-STATISTICS) ²⁾
Data to OPENFT	OPENFT	OPENFT	OPENFT

Table 18: Monitored variables in relation to the monitoring program, ONLINE report, and SM2R1 report group
(part 3 of 4)

Monitored variable	Monitoring program	ONLINE report	SM2R1 report group
Task			
Number of tasks	1)	ACTIVITY	TASK
Length of task queues	1), SYSTEM	CATEGORY QUEUE	CATEGORY- QUEUE
Delay times in task queues	SYSTEM	-----	(QUEUE TRANSITION) ²⁾
Frequency of task activation and deactivation	1)	ACF	PRIOR-ACF
Task-specific resource utilization	PERIODIC-TASK	PERIODIC TASK	PERIODIC-TASK- JOBNAME/ TSN/USERID
Allocation and queues of task locks	TLM	TLM	TLM
VM2000			
Hypervisor activities (servers with /390 architecture), guest system activities, CPU pools, VM groups (servers with /390 architecture)	VM2000	VM, VM CPU POOL, VM GROUP	VM2000

1) The monitored values are recorded permanently (default monitoring method).

2) The items in brackets refer to SM2R1 statistics, not report groups.

3) The monitored values are recorded permanently for channels (default monitoring method).

4) The data is available online via the SM2GDAT program interface or in the output file (analysis only possible using ANALZYER)

Table 18: Monitored variables in relation to the monitoring program, ONLINE report, and SM2R1 report group
(part 4 of 4)

12.2 Table of report groups

The columns of the following table contain, from left to right:

- the name of the (current) report group,
- an asterisk if this report group is included in the report group STD
- the monitored objects that can be specified
- the character Y if *ALL may be specified for the monitored objects
- the numbers of the reports if REPORT-NUMBER=*ALL is specified in the PRINT-REPORTS statement
- the numbers of the reports if REPORT-NUMBER=*STD is specified in the PRINT-REPORTS statement
- the SM2 monitoring programs corresponding to the report groups
(*: monitored data is entered continuously)

Group name	*STD	Monitored object	*ALL	Relevant report numbers for		SM2 monitoring program
				REPORT-NUMBER = *ALL	REPORT-NUMBER = *STD	
*ALL		–		All report numbers	Standard reports of all groups	
*BCAM-CONNECTION		Connection name	Y	192–203, 226, 258–261	196, 197	BCAM-CONNECTION
*BCAM-MEMORY		–	J	280, 281	280, 281	BCAM-CONNECTION
*CATALOG-MANAGEMENT		Catalog	Y	66–72, 103, 104, 185	66, 67, 72	CMS
*CATEGORY-CPU		Category name	Y	62	62	SYSSTAT
*CATEGORY-IO		Category name	Y	63–65	63, 64	
*CATEGORY-QUEUE		Category name	Y	28, 30, 31	31	*
*CATEGORY-WORKING-SET		Category name	Y	29, 58	29	
*CHANNEL	*	Channel address	Y	10, 101, 102, 257	10, 102	*(101, 102, 257:CH.-IO)
*CPU	*	PROCESSOR SPLITTING		1, 2, 6, 137, 204	1, 6	*

Table 19: Report groups

(part 1 of 3)

Group name	*STD	Monitored object	*ALL	Relevant report numbers for		SM2 monitoring program
				REPORT-NUMBER = *ALL	REPORT-NUMBER = *STD	
*DAB		Cache ID	Y	79–82, 189, 190, 205, 206	79, 80	DAB
*DEVICE		Devices	Y	11, 35, 36, 100, 230, 282, 283, 319	35	SAMPLING-DEVICE
*DILATION		Category name	Y	57	57	SYSTEM
*DISK	*	Device	Y	124–127, 227–229, 270	124, 128, 228	SAMPLING-DEVICE
*DISK-FILE		File names	Y	320	320	DISK-FILE
*DLM		–		170–173	170	DLM
*FILE		File name	Y	13, 14, 191	14	FILE
*GS		Partition	Y	177, 178	177, 178	GS
*GSVOL		Devices	Y	174	174	GSVOL
*IO	*	PROCESSOR-SPLITTING		3, 4	3	*
*ISAM		ISAM pool name	Y	85–87, 224	85	ISAM
*ISAM-FILE		NK-ISAM filename	Y	271–274	271	ISAM
*MEMORY	*	–	Y	8, 9, 52–56, 94, 95, 275	8, 9, 52, 54	*
*MSCF		–		166–169	166	MSCF
*NSM		–		179–182, 184	179, 184	NSM
*OPENFT		openFT-instances	Y	309-318	309, 313	OPENFT
*PCS		Category name	Y	73–77	73, 74, 77	*
*PERIODIC-TASK-JOBNAME		Job name	Y	160–163	160	PERIODIC-TASK
*PERIODIC-TASK-TSN		TSN	Y	152–155	152	
*PERIODIC-TASK-USERID		USERID	Y	156–159	156	
*PFA		Cache ID	Y	134–136	134, 135	PFA
*POSIX		–		141–144, 146–151	147, 148	POSIX

Table 19: Report groups

(part 2 of 3)

Group name	*STD	Monitored object	*ALL	Relevant report numbers for		SM2 monitoring program
				REPORT-NUMBER = *ALL	REPORT-NUMBER = *STD	
*PRIOR-ACF		CPU, MEM, PAG	Y	32–34	32	
*PUBSET		Pubset	Y	262, 263, 264	262	PUBSET
*RESPONSE-TIME	*	Connection name	Y	19–27, 46–50, 83, 89–93, 107-110	23, 47–50, 109	RESPONSE TIME
*RST		Category name	Y	59, 60, 61	60	SYSTEM
*SERVICETIME		Devices	Y	231	231	SERVICETIME
*SESAM-SQL		–	J	294–308	294–296	SESAM-SQL
*STD		–		see column 2	see column 2	
*STORAGE-SYSTEM-SYMMETRIX		Serial number	Y	232–256	234, 235	STORAGE-SYSTEM
*SVC		SVC number	Y	123	123	SVC
*TASK	*	–		5	5	*
*TCP-IP		–		186, 187	187	TCP-IP
*TLM		Lock name	Y	96, 97	96, 97	TLM
*UDS-SQL		–	J	284–293	284, 285	UDS-SQL
*UTM		Application	Y	128–133, 225	128, 131, 132	UTM
*VM2000		–		98, 99, 164, 267–269	98, 99	VM
*VOLUME		Volume name	Y	12	–	*
*WORKING-SET	*	–		15	15	*

Table 19: Report groups

(part 3 of 3)

12.3 Table of report names

The report name is given in the heading of each report page.

Report group	Name	Report
*BCAM-CONNECTION	BCAM INPROC TIME	260
	BCAM INWAIT TIME	258
	BCAM OUTPROC TIME	261
	BCAM REACT TIME	259
	CONNECTION	226
	DATA PER TSDU	197
	INPROC TIME DISTRIBUTION	194
	INPUT BUFFER	198
	INWAIT TIME DISTRIBUTION	192
	OUTPROC TIME DISTRIBUTION	195
	OUTPUT BUFFER	199
	PACKETS WITH DATA	201
	PACKETS WITH ZERO WINDOW INFORMATION	203
	PACKETS WITHOUT DATA	200
	REACT TIME DISTRIBUTION	193
	SEND CALL'S OVER MAXIMUM	202
TSDU	196	
*BCAM-MEMORY	INPUT MEMORY POOL	280
	OUTPUT MEMORY POOL	281
*CATALOG-MANAGEMENT	CMS IO'S	67
	CMS QUEUES	66
	CMS QUEUES II	185
	CMS RESPONSE TIME	72
	LOCAL CMS ACCESSES TO FILES	68
	LOCAL CMS ACCESSES TO JOB VARIABLES	70
	REMOTE CMS ACCESSES TO FILES	69
	REMOTE CMS ACCESSES TO JOB VARIABLES	71
	SH. PUBSET CMS ACCESSES TO JOB VARIABLES	104
	SHARED PUBSET CMS ACCESSES TO FILES	103

Table 20: Report names

(part 1 of 9)

Report group	Name	Report
*CATEGORY-CPU	CPU UTILIZATION (TU+TPR) FOR CATEGORY	62
*CATEGORY-IO	DURATION OF NON PAGING IO'S FOR CATEGORY	64
	DURATION OF PAGING IO'S FOR CATEGORY	65
	IO'S FOR CATEGORY	63
*CATEGORY-QUEUE	ACTIVE AND INACTIVE TASKS FOR CATEGORY	30
	TASKS IN ACTIVE QUEUES FOR CATEGORY	31
	TASKS IN CATEGORIES	28
*CATEGORY-WORKING-SET	USED PAGES (UPG) FOR CATEGORY	58
	WORKING SET (PPC) FOR CATEGORY	29
*CHANNEL	DATA TRANSFER	257
	DATA TRANSFER	101
	IO	102
	UTILIZATION	10
*CPU	ACTIVE LOGICAL MACHINES	6
	SUM SVC CALLS	137
	UTILIZATION NORMED	1
	UTILIZATION REAL	2
*DAB	DESTAGE IO'S	206
	OVERRUNS FOR DAB CACHE	189
	OVERRUNS FOR INTERNAL AREA	190
	PENDING BLOCKS	205
	READS FOR DAB CACHE	79
	READS FOR INTERNAL AREA	81
	WRITES FOR DAB CACHE	80
	WRITES FOR INTERNAL AREA	82
*DEVICE	DATA	230 283
	DATA PER IO	100 282
	IO	35
	QUEUE LENGTH	36
	RSC IO	319
	UTILIZATION	11

Table 20: Report names

(part 2 of 9)

Report group	Name	Report
*DILATION	DILATION FOR CATEGORIES	57
*DISK	ALIAS DEVICES	229
	DATA	228
	DATA PER IO	127
	IO	125
	QUEUE LENGTH	126
	RSC IO	270
	TIME	227
	UTILIZATION	124
*DISK-FILE	IO	320
*DLM	NSM GRANT/RELEASE EVENTS	173
	NSM LOCKREQUESTS	171
	USER GRANT/RELEASE EVENTS	172
	USER LOCKREQUESTS	170
*FILE	ACCESSTIME	191
	IO	14
	PAM BLOCKS	13
*GS	DATA	178
	IO	177
*GSVOL	IO	174
*IO	IO'S FOR DEVICE CLASSES	3
	PAGING IO'S (DISK)	4
*ISAM	FIX OPERATIONS	85
	INDEX OPERATIONS	224
	SLOT OPERATIONS	86
	USED PAGES	87
*ISAM-FILE	FIX OPERATIONS	271
	INDEX OPERATIONS	274
	SLOT OPERATIONS	272
	USED PAGES	273

Table 20: Report names

(part 3 of 9)

Report group	Name	Report
*MEMORY	AVAILABLE PAGEABLE MEMORY	8
	BIG PAGES	275
	CLASS 3/4 MEMORY	53
	PAGE FAULTS	54
	PAGE FRAMES	52
	PAGE READS FROM DISK	56
	PAGE READS FROM ES/GS	94
	PAGE WRITES TO DISK	55
	PAGE WRITES TO ES/GS	95
	PAGING AREA UTILIZATION	9
*MSCF	CALLS	166
	REQUESTS WITH REPLY	169
	SERVERTASKS	167
	WAIT TIME REQUEST WITH REPLY	168
*NSM	LOCK REQUESTS	179
	MESSAGE LENGTH	181
	NUMBER LOCKSERVER	184
	TOKEN DURATION TIME	180
	TOKEN TIME	182
*OPENFT	NET DATA TOTAL	309
	NET DATA SENT/RECEIVED	310
	DISK DATA TOTAL	311
	DISK DATA READ/WRITTEN	312
	REQUEST RATE	313
	CONNECTION FAULT RATE	314
	REQUESTS ACTIVE	315
	REQUESTS NOT ACTIVE	316
	CONNECTION USAGE	317
	REQUEST USAGE	318

Table 20: Report names

(part 4 of 9)

Report group	Name	Report
*PCS	DURATION RUNOUTS FOR CATEGORY	75
	REQUEST DELAY FOR CATEGORY	74
	SERVICE QUOTA DISTRIBUTION FOR CATEGORY	73
	SERVICE RATE DISTRIBUTION FOR CATEGORY	76
	SERVICE RATES FOR CATEGORY	77
*PERIODIC-TASK-JOBNAME	CPU-TIME	161
	IO	162
	SERVICE RATE	160
	UPG	163
*PERIODIC-TASK-TSN	CPU-TIME	153
	IO	154
	SERVICE RATE	152
	UPG	155
*PERIODIC-TASK-USERID	CPU-TIME	157
	IO	158
	SERVICE RATE	156
	UPG	159
*PFA	CACHE OVERRUNS	136
	CACHING READ	134
	FAST WRITE	135
*POSIX	OPTION A: IGET	141
	OPTION A: NAMEI	142
	OPTION B: NUMBER CACHE READS	143
	OPTION B: NUMBER CACHE WRITES	144
	OPTION B: NUMBER NON CACHE REQUESTS	146
	OPTION C: NUMBER CHARACTERS FOR SYSTEM CALLS	149
	OPTION C: NUMBER SPECIAL SYSTEM CALLS	148
	OPTION C: SCALL	147
	OPTION M: MSG	150
OPTION M: SEMA	151	

Table 20: Report names

(part 5 of 9)

Report group	Name	Report
*PRIOR-ACF	ACF LONG CALLS	33
	ACF SHORT CALLS	34
	ACF UTILIZATION	32
*PUBSET	RELATIVE SPACE ALLOCATION	264
	SATURATION LEVEL	263
	SPACE ALLOCATION	262
*RESPONSE-TIME	BCAM WAIT TIME	109
	BCAM WAIT TIME < LIMIT	107
	BCAM WAIT TIME < LIMIT DISTRIBUTION	108
	BCAM WAIT TIME FOR CATEGORY	110
	MESSAGE LENGTH OF TRANSACTIONS	83
	RATE OF TRANSACTION < LIMIT	23
	RESPONSE TIME (1)	47
	RESPONSE TIME (1) < LIMIT	19
	RESPONSE TIME (1) < LIMIT DISTRIBUTION	24
	RESPONSE TIME (1) FOR CATEGORY	90
	RESPONSE TIME (2)	48
	RESPONSE TIME (2) < LIMIT	20
	RESPONSE TIME (2) < LIMIT DISTRIBUTION	25
	RESPONSE TIME (2) FOR CATEGORY	91
	THINK TIME	46
	THINK TIME < LIMIT	22
	THINK TIME < LIMIT DISTRIBUTION	26
	THINK TIME FOR CATEGORY	89
	TRANSACTION RATE	50
	TRANSACTION RATE ≤ LIMIT	88
	TRANSACTION RATE FOR CATEGORY	93
	TRANSACTION TIME	49
TRANSACTION TIME < LIMIT	21	
TRANSACTION TIME < LIMIT DISTRIBUTION	27	
TRANSACTION TIME FOR CATEGORY	92	

Table 20: Report names

(part 6 of 9)

Report group	Name	Report
*RST	CPU RST FOR CATEGORY	60
	IO RST FOR CATEGORY	61
	RST FOR CATEGORY	59
*SERVICETIME	DURATION OF IO'S FOR DEVICE	231
*SESAM-SQL	LOCAL STATEMENTS	294
	LOGICAL ACCESSES CURSOR FILES	303
	LOGICAL ACCESSES SYSTEM DATA BUFFER	299
	LOGICAL ACCESSES USER DATA BUFFER	301
	NOT ACTIVE THREADS	307
	PHYSICAL ACCESSES CURSOR FILES	304
	PHYSICAL ACCESSES LOG FILES	305
	PHYSICAL ACCESSES SYSTEM BUFFER	300
	PHYSICAL ACCESSES USER BUFFER	302
	REMOTE STATEMENTS	295
	SERVICE	308
	SQL	297
	SQL PLAN	298
	THREADS	306
	TRANSACTIONS	296

Table 20: Report names

(part 7 of 9)

Report group	Name	Report
*STORAGE-SYSTEM-SYMMETRIX	DATA	235
	DEVICE BACK END DATA	252
	DEVICE BACK END IO	251
	DEVICE DATA	248
	DEVICE READS	245
	DEVICE READS AND WRITES	247
	DEVICE SEQUENTIAL READS	249
	DEVICE SEQUENTIAL WRITES	250
	DEVICE WRITES	246
	DIRECTOR IOS	255
	PHYSICAL DISK DATA	254
	PHYSICAL DISK IOS	253
	PUBSET BACK END DATA	244
	PUBSET BACK END IO	243
	PUBSET DATA	240
	PUBSET READS	237
	PUBSET READS AND WRITES	239
	*STORAGE-SYSTEM-SYMMETRIX (continued)	PUBSET SEQUENTIAL READS
PUBSET SEQUENTIAL WRITES		242
PUBSET WRITES		238
RA DIRECTOR DATA		256
READS		232
READS AND WRITES		234
SEQUENTIAL READS		236
WRITES		233
*SVC	CALLS	123
*TASK	TASKS	5
*TCP-IP	DATA	187
	TSDU	186
*TLM	OCCUPATION OF LOCK	97
	TASKS WAITING FOR LOCK	96

Table 20: Report names

(part 8 of 9)

Report group	Name	Report
*UDS-SQL	DML/SQL	284
	LOCKS	290
	LOGICAL ACCESSES DB	287
	PHYSICAL ACCESSES A/RLOG FILES	289
	PHYSICAL ACCESSES DB	288
	PPP	292
	REMOTE REQUESTS	286
	TASK COMMUNICATION	291
	TRANSACTION TIME	293
	TRANSACTIONS	285
*UTM	ASYNCHRON TIME	132
	BOURSE WAIT TIME	225
	DIALOG TIME	131
	REQUESTS	128
	TASKS	130
	USERS	129
	WAITING REQUESTS	133
*VM2000	CPU POOL CPUS	268
	CPU POOL UTILIZATION	267
	CPU'S	164
	VM2000 GROUP	269
	VM2000 HYPERVISOR	99
	VM2000 UTILIZATION	98
*VOLUME	UTILIZATION	12
*WORKING-SET	MAIN MEMORY UTILIZATION	15

Table 20: Report names

(part 9 of 9)

12.4 Table of variables reports

The following table lists the variables reports by

- their numbers
- the symbols used for the variables in the bar chart
- the variables contained in the report
- the unit used for data output.

Report No.	Chart symbol	Variable	Unit
1	1 . 3 ▣ —	TU time TPR time SIH time IDLE time STOP time A STOP time for inactive CPUs is output only when PROCESSOR-SPLITTING=*YES.	%
2	1 . 3 ▣ —	TU time TPR time SIH time IDLE time STOP time A STOP time for inactive CPUs is output only when PROCESSOR-SPLITTING=*YES. If VM2000 is used, the sum of the values does not come to 100%, since, unlike report 1, the monitored variables are not relativized.	%

Table 21: Variable reports

(part 1 of 38)

Report No.	Chart symbol	Variable	Unit
3	D	Disk I/O operations (without paging)	I/O operations per second
	*	I/O operations in response to paging requests (read and write operations are counted)	
	T	Tape I/O operations	
	.	Printer I/O operations	
	0	I/O operations with all other devices	
		Note On servers with x86 architecture this value contains the operations for devices with the FAMILY name TD only for 'ALL PROCESSORS' (i.e. also openCRYPT-BOX devices)	
4	0	I/O operations in response to paging requests (read and write operations are counted)	I/O operations per second
5	S	Existing system tasks	Number of tasks
	B	Existing batch tasks	
	D	Existing interactive tasks	
	T	Existing transaction tasks	
6	*	Active logical machines This report is only output if the number of logical machines changes during the monitoring cycle.	Number of machines
8	+	Paging memory size	4-KB pages
	-	Main memory size	
9	+	Pages used for paging on all devices (without corresponding pages in global storage)	4-KB pages
	-	Maximum number of pages for paging on all devices	
	*	Number of pages for paging in global storage	
10	+	Channel active and at least one CPU active	%
	-	Channel active	

Table 21: Variable reports

(part 2 of 38)

Report No.	Chart symbol	Variable	Unit
11	+ –	Device utilization exclusive of paging Device utilization due to paging. A bar chart is output for each device. The paging share is only output if the corresponding monitored data is not zero.	%
12	+ –	Volume utilization exclusive of paging Volume utilization due to paging A separate chart is output for each volume The paging share is output only if a value other than zero was found SM2 does not recognize a volume change during the same cycle	%
13	W C 0 I	PAM WAIT operations PAM CHECK READ operations PAM WRITE operations PAM READ operations to the file monitored by SM2; SM2R1 shortens the file name to 40 characters	Number of operations per second
14	+ – . *	Number of I/Os with PAM-WAIT Number of I/Os with PAM-CHECK Number of write I/Os Number of read I/Os	I/Os per second
15	+ – .	Planned working sets (PPC) of active and inactive ready tasks in the system Number of available main memory pages (NPP) Planned working sets (PPC) of active tasks in the system	4-KB pages
19	+ –	Average response time in accordance with definition 1 (see DEFINITION operand under SET-RESP-PARAMETER) Overflow time for the response time (see the RESP-BUCKETS operand under SET-RESP-PARAMETER) Only those response times are used for averaging which are shorter than the specified overflow time.	Seconds

Table 21: Variable reports

(part 3 of 38)

Report No.	Chart symbol	Variable	Unit
20	+	Average response time in accordance with definition 2 (see DEFINITION operand under SET-RESP-PARAMETER)	Seconds
	–	Overflow time for the response time (see RESP-BUCKETS operand under SET-RESP-PARAMETER) Only those response times are used for averaging which are shorter than the specified overflow time.	
21	T	Average transaction time	Seconds
	–	Overflow time for the response time (see TRANSACT-BUCKETS operand under SET-RESP-PARAMETER) Only those transaction times are used for averaging which are shorter than the specified overflow time.	
22	T	Average think time	Seconds
	–	Overflow time for the think time (see THINK-BUCKETS operand under SET-RESP-PARAMETER) Only those think times are used for averaging which are shorter than the specified overflow time.	
23	@	Transactions	Interactions per second
	*	Responses in accordance with definition 1	
	.	Responses in accordance with definition 2, (see DEFINITION operand under SET-RESP-PARAMETER) Data exceeding the specified overflow times is not recorded.	
24		All responses are based on definition 1 (see DEFINITION operand under SET-RESP-PARAMETER)	%
	A	Responses in first bucket	
	.	Responses in second bucket	
	B	Responses in third bucket	
	*	Responses in fourth bucket	
C	Responses in fifth bucket The bucket limits are given in the bar chart and the table.		

Table 21: Variable reports

(part 4 of 38)

Report No.	Chart symbol	Variable	Unit
25	A . B * C	All responses are based on definition 2 (see DEFINITION operand under SET-RESP-PARAMETER) Responses in first bucket Responses in second bucket Responses in third bucket Responses in fourth bucket Responses in fifth bucket The bucket limits are given in the bar chart and in the table.	%
26	A . B * C	Think times in first bucket Think times in second bucket Think times in third bucket Think times in fourth bucket Think times in fifth bucket The bucket limits are given in the bar chart and in the table.	%
27	A . B * C	Transaction times in first bucket Transaction times in second bucket Transaction times in third bucket Transaction times in fourth bucket Transaction times in fifth bucket The bucket limits are given in the bar chart and in the table.	%
28	+ . @	For each task category a chart and a separate line in the statistics table are output Number of tasks in this category Minimum multiprogramming level for this category (MIN MPL) Maximum multiprogramming level for this category (MAX MPL)	Number of tasks

Table 21: Variable reports

(part 5 of 38)

Report No.	Chart symbol	Variable	Unit
29	@ . 	For each task category a chart and a separate line in the statistics table are output Planned working sets (PPC) for active tasks in this category Planned working sets (PPC) for inactive ready tasks in this category Planned working sets (PPC) for inactive non-ready tasks in this category	4-KB pages
30	@ . –	A chart is output for each category and a separate line is output in the table of statistics Number of active tasks in this category (Q0 – Q4) Number of inactive ready tasks in this category (Q5) Number of inactive non-ready tasks in this category (Q7 – Q13) Number of tasks in this category rejected by PCS (Q6)	Number of tasks
31	1 . 4 @	A bar chart is output for each category and a separate line is output in the table of statistics Number of tasks of a category in task queue 1 (waiting for CPU) and in task queue 0 (serviced by CPU) Number of tasks of a category in the paging queue (Q3) Number of tasks of a category in the I/O queue (Q4 with I/O pend code) Number of remaining active tasks of a category (Q2) and of tasks in queue Q4 which are not waiting for I/O	Number of tasks
32	– ⌞ +	A separate chart is output for the CPU, main memory (MEM) and the paging devices (PAG), and a separate line in the statistics table. Low utilization of this resource Medium utilization of this resource High utilization of this resource	%

Table 21: Variable reports

(part 6 of 38)

Report No.	Chart symbol	Variable	Unit
33	–	Number of “long” ACF invocations	Number per second
	+	Number of forced task deactivations	
	.	Number of system service runouts	
	@	Number of active-task preemptions	
34	–	Number of “short” ACF invocations	Number per second
	+	Number of activations	
	.	Number of micro time slice runouts	
35		A separate chart is output for each device except communication controllers; in the table, the computed values are listed separately for each device.	Operations per second
	+	Number of I/O operations per device	
36		A separate chart is output for each device except communication controllers; in the table, the computed values are listed separately for each device.	Number of I/O requests
	+	Average length of device queue, including the jobs currently being serviced	
46	0	Average think time not exceeding the limit	Seconds
	*	Average think time Think times exceeding the overflow limit are also used in computing this value.	
47	0	Average response time (TYP=1) not exceeding the limit	Seconds
	*	Average response time (TYP=1) Response times exceeding the overflow limit are also used in computing this value	
48	0	Average response time (TYP=2) not exceeding the limit	Seconds
	*	Average response time (TYP=2) Response times exceeding the overflow limit are also used in computing this value	

Table 21: Variable reports

(part 7 of 38)

Report No.	Chart symbol	Variable	Unit
49	0 *	Average transaction time not exceeding the limit Average transaction time Transaction times exceeding the overflow limit are also used in computing this value	Seconds
50	@ . *	Number of transactions per second Number of responses per second in accordance with definition 1 Number of responses per second in accordance with definition 2 Transactions and responses exceeding the overflow limit are also used in computing this value	Transactions per second
52	+ . * – @	Average number of page frames managed task-locally (obsolete, no longer supplied) Average number of page frames managed system globally Average number of read-only page frames in the free pool Average number of written page frames in the free pool Average number of page frames in the free pool which are not assigned to any page	4-KB pages
53	+ . –	Number of class 4 memory pages for shareable modules in the virtual address space (below 16 Mb) Number of class 3 memory pages Number of class 4 memory pages	4-KB pages
54	+ . –	Total number of paging errors per second. “Real” paging errors (access to virtually unassigned pages) are not included Number of paging errors per second for which the addressed page is still in main memory Number of paging errors per second during first access to a page (see “Page fault” in Glossary)	Number of paging errors per second
55	*	Number of pages written to background memory (DISK) per second (see “Page fault” in Glossary)	Number of 4-KB pages per second

Table 21: Variable reports

(part 8 of 38)

Report No.	Chart symbol	Variable	Unit
56	*	Number of pages read from background memory (DISK) per second (see "Page fault" in Glossary)	Number of 4-KB pages per second
57	*	Dilation factor for the specified category. To ensure maximum accuracy in monitoring this variable, the statement SET-SYSTAT-PARAMETER DEVICES=*ALL should be issued when defining the SYSTEM monitoring program (see "Dilation factor" in the glossary)	
58	@ . 	UPG of active tasks in this category UPG of inactive ready tasks in this category UPG of inactive non-ready tasks in this category A separate chart is output for each category, and a separate line in the statistics table	4-KB pages
59	*	Time equivalent for the productive performance of the specified category A bar chart is output for each category (For RST, see "overall performance" in the glossary)	RST per second
60	*	Time equivalent for the productive performance of the specified category (CPU share) A bar chart is output for each category (For RST, see "overall performance" in the glossary)	RST per second
61	*	Time equivalent for the productive performance of the specified category (periphery share) A bar chart is output for each category (For RST, see "Overall performance" in glossary)	RST per second
62	*	Time share (TU + TPR) for the specified category A bar chart is output for each category In the case of a multiprocessor system the average value of the processors is output (values <= 100%).	%

Table 21: Variable reports

(part 9 of 38)

Report No.	Chart symbol	Variable	Unit
63	+ –	DMS I/O operations per category Paging I/O operations per category A bar chart is output for each category For paging I/Os, only the assignable I/O operations are supplied per category. On a global basis non-assignable paging I/O operations are output (for SUM) in addition to the sum for all categories (number of I/O operations in I/O chains with write operations only).	I/O operations per second
64	+ –	Hardware duration of DMS I/O operations for the specified category Software duration of DMS I/O operations for the specified category A bar chart is output for each category	Milliseconds
65	+ –	Hardware duration of paging I/O operations for the specified category Software duration of paging I/O operations for the specified category A bar chart is output for each category. For paging I/O operations only the assignable I/O operations are supplied for each category. On a global basis (for SUM), the non-assignable paging I/O operations are used in addition to the average for all categories: Hardware or software duration of the I/O operations in I/O chains with write operations only	Milliseconds

Table 21: Variable reports

(part 10 of 38)

Report No.	Chart symbol	Variable	Unit
66	+ . - @ 0	<p>Average number of requests waiting for a serialization lock for exclusively searching a partition (corresponds to a user ID) or currently using this lock</p> <p>Average number of requests waiting for the release of a buffer management table or currently using a buffer management table</p> <p>Average number of requests waiting for a lock for searching a partition or currently using this lock</p> <p>Average number of requests waiting for a lock for exclusively searching a block of a partition or currently using this lock</p> <p>Average number of requests waiting for a lock for exclusively searching a catalog entry of a block of a partition or currently using this lock</p>	Number of requests
67	+ - . @	<p>For normal pubsets the catalog ID is used (CATID, one to four alphanumeric characters). For the set of all private disks this is supplied as CATID \$.</p> <p>Number of physical read accesses (CMS) in half-pages</p> <p>Number of read accesses (half-pages) which were physically not required because the catalog entry was contained in a buffer management table (BMT)</p> <p>Number of physical write accesses in half-pages</p> <p>Number of I/O errors during physical accesses (read/write)</p> <p>For catalog ID (CATID, 1 to 4 alphanumeric characters) is output for normal pubsets. \$ is supplied as the catalog for the total of all private disks</p>	Accesses per second

Table 21: Variable reports

(part 11 of 38)

Report No.	Chart symbol	Variable	Unit
68	+	Number of local read accesses to catalog entries for files without LBN specification	Accesses per second
	–	Number of local read accesses to catalog entries for files with LBN specification	
	.	Number of local accesses to catalog entries for files during scanning (SHOW-FILE-ATTRIBUTES command)	
	@	Number of local write accesses to catalog entries for files (only manipulation of file attributes)	
	0	Number of local write accesses to catalog entries for files (not only manipulation of file attributes) - The catalog ID (1 – 4 alphanumeric characters) is used for normal pubsets. \$ is supplied as the catalog ID for the total of all private disks. - For this report a call is <i>local</i> if it does not originate from a remote processor.	
69	+	Number of non-local read accesses to catalog entries for files without LBN specification	Accesses per second
	–	Number of non-local read accesses to catalog entries for files with LBN specification	
	.	Number of non-local accesses to catalog entries for files during scanning (SHOW-FILE-ATTRIBUTES command)	
	@	Number of non-local write accesses to catalog entries for files (only manipulation of file attributes)	
	0	Number of non-local write accesses to catalog entries for files (not only manipulation of file attributes) - The catalog ID (1 – 4 alphanumeric characters) is used for normal pubsets. \$ is supplied as the catalog ID for the total of all private disks. - For this report a call is <i>local</i> if it does not originate from a remote processor.	

Table 21: Variable reports

(part 12 of 38)

Report No.	Chart symbol	Variable	Unit
70	+	Number of local read accesses to JV entries without LBN specification	Accesses per second
	–	Number of local read accesses to JV entries with LBN specification	
	.	Number of local accesses to JV entries during scanning (SHOW-JV-ATTRIBUTES command)	
	@	Number of local write accesses to JV entries (only manipulation of JV attributes)	
	0	Number of local write accesses to JV entries (only manipulation of JV attributes)	
		- The catalog ID (1 – 4 alphanumeric characters) is used for normal pubsets. \$ is supplied as the catalog ID for the total of all private disks. - For this report a call is <i>local</i> if it does not originate from a remote processor.	
71	+	Number of non-local read accesses to JV entries without LBN specification	Accesses per second
	–	Number of local read accesses to JV entries with LBN specification	
	.	Number of non-local accesses to JV entries during scanning (SHOW-JV-ATTRIBUTES command)	
	@	Number of non-local write accesses to JV entries (only manipulation of JV attributes)	
	0	Number of non-local write accesses to JV entries (only manipulation of JV attributes)	
		- The catalog ID (1 – 4 alphanumeric characters) is used for normal pubsets. \$ is supplied as the catalog ID for the total of all private disks. - For this report a call is <i>local</i> if it does not originate from a remote processor.	

Table 21: Variable reports

(part 13 of 38)

Report No.	Chart symbol	Variable	Unit
72	+ - . @ 0	<p>Average response time for read accesses to catalog entries without LBN specification</p> <p>Average response time for read accesses to catalog entries with LBN specification</p> <p>Average response time for accesses to catalog entries during scanning (SHOW-JV-ATTRIBUTES/SHOW-FILE-ATTRIBUTES command)</p> <p>Average response time for write accesses to catalog entries (only manipulation of file/JV attributes)</p> <p>Average response time for write accesses to catalog entries (not only manipulation of file/JV attributes)</p> <ul style="list-style-type: none"> - The catalog ID (1 – 4 alphanumeric characters) is used for normal pubsets. \$ is supplied as the catalog ID for the total of all private disks. - For this report a call is <i>local</i> if it does not originate from a remote processor - All calls from both local and remote processors are used for this report 	Milliseconds
73	@ . + -	<p>Current share of system performance obtained for this category (service share)</p> <p>Current system performance share provided for this category</p> <p>Maximum system performance share provided for this category</p> <p>Minimum system performance share provided for this category</p>	%
74	. + -	<p>Request delay</p> <p>Current delay of all requests in this category</p> <p>Maximum delay of all requests in this category</p> <p>Minimum delay of all requests in this category</p> <ul style="list-style-type: none"> - A summary report is output for the sum of all categories under the pseudo-category SUM. - For SUM the MINIMUM REQUEST DELAY data is undefined - For SUM the MAXIMUM REQUEST DELAY data corresponds to the global PCS parameter REQUEST DELAY 	No dimension

Table 21: Variable reports

(part 14 of 38)

Report No.	Chart symbol	Variable	Unit
75	+ –	Number of duration time slice runouts in this category Number of duration time slice runouts with preemption in this category - A summary report is output for the sum of all categories under the pseudo-category SUM. - The value of the PCS parameter DURATION is output as an additional SM2R1 identifier. Zero is always output for SUM.	Number per second
76	C M I	CPU service rate per category Memory service rate per category I/O service rate per category	SERVICE RATE
77	. –	Current SERVICE RATE provided for this category Currently used SERVICE RATE for this category A report is output for the sum of all categories under the pseudo-category SUM.	SERVICE RATE
79	+ –	Read accesses per DAB cache area Number of read accesses to all subareas supported by the DAB cache area for which no disk access was required Number of read accesses to all subareas supported by the DAB cache area A chart is output for each cache area	I/O operations per second
80	+ –	Write accesses for each DAB cache area Number of write accesses during which the cache area was written to Number of write accesses to all subareas supported by the DAB buffer A chart is output for each buffer	I/O operations per second

Table 21: Variable reports

(part 15 of 38)

Report No.	Chart symbol	Variable	Unit
81	+ –	Read accesses per DAB subarea Number of read accesses to the specified subarea for which no disk access was required Number of read accesses to the specified subarea A chart is output for each DAB subarea	I/O operations per second
82	+ –	Number of write accesses to the specified subarea Number of write accesses during which data was written to the buffer which supports the subarea A chart is output for each DAB subarea	I/O operations per second
83	I O	Average length of the input messages for all transactions in the specified connection set Average length of the output messages for all transactions in the specified connection set A chart is output for each connection set	Bytes
85	H I W	Number of ISAM accesses where the desired page is already in the ISAM pool, i.e. no read operation from disk is required (FIX-HIT) Number of ISAM accesses where a read operation from disk was required (FIX-IO) Number of ISAM accesses involving a waiting period for the release of one or more buffer pages (FIX-WAIT)	Number per second
86	W N	Number of RESERVE-SLOT operations leading to a wait state due to a slot bottleneck Number of RESERVE-SLOT operations where a slot is available and no wait state occurs	Number per second
87	T R F	Average number of pages in the ISAM pool Average number of reserved pages in the ISAM pool Average number of fixed pages in the ISAM pool	PAM pages
88	@ . *	Transactions Responses in accordance with definition 1 Responses in accordance with definition 2 (only transactions or responses below the specified overflow values are included in the calculation)	Number per second

Table 21: Variable reports

(part 16 of 38)

Report No.	Chart symbol	Variable	Unit
89	0 *	Average think time not exceeding the limit per category Average think time per category Think times exceeding the overflow limit are also used in computing this value	Seconds
90	0 *	Average response time (TYP=1) not exceeding the limit per category Average response time (TYP=1) per category Response times exceeding the overflow limit are also used in computing this value	Seconds
91	0 *	Average response time (TYP=2) not exceeding the limit per category Average response time (TYP=2) per category Response times exceeding the overflow limit are also used in computing this value	Seconds
92	0 *	Average transaction time not exceeding the limit per category Average transaction time per category Transaction times exceeding the overflow limit are also used in computing this value	Seconds
93	@ . *	Number of transactions per second for every category Number of responses in accordance with definition 1 per second for each category Number of responses in accordance with definition 2 per second for each category Transactions and responses exceeding the overflow limit are also used in computing this value	Transactions per second
94	@	Number of pages read from ES/GS during paging	Number of 4-KB pages per second
95	@	Number of pages written to ES/GS during paging	Number of 4-KB pages per second
96	*	Average number of tasks in the queue of the corresponding lock	Number of tasks

Table 21: Variable reports

(part 17 of 38)

Report No.	Chart symbol	Variable	Unit
97	*	Average utilization frequency of the corresponding task lock in percent	%
98	* . –	Maximum CPU utilization per VM planned CPU utilization per VM Monitored CPU utilization per VM	%
99	A I	Monitored active time of the server's hypervisor Monitored system idle time of the server (all data only for servers with /390 architecture)	%
100	+ –	Data volume for reading and writing Data volume for writing	Bytes per I/O operation
101	* –	Quantity of data transmitted per channel by the PAM block transfer Quantity of data transmitted by means of byte transfer	PAM pages per second
102	* - 0	Number of I/O operations per PAM block transfer Number of I/O operations per byte transfer Number of I/O operations without data transfer	I/O operations per second
103	+ . – @ 0	Number of local read accesses to catalog entries of files without LBN specification Number of local read accesses to catalog entries of files with LBN specification Number of local accesses to catalog entries for files when scanning (SHOW-FILE-ATTRIBUTES command) Number of local write accesses to catalog entries of files, whereby only the file attributes are manipulated Number of local write accesses to catalog entries of files, whereby not only file attributes are manipulated	Accesses per second

Table 21: Variable reports

(part 18 of 38)

Report No.	Chart symbol	Variable	Unit
104	+	Number of local read accesses to JV entries without LBN specification	Accesses per second
	.	Number of local read accesses to JV entries with LBN specification	
	–	Number of local accesses to JV entries during scanning (SHOW-JV-ATTRIBUTES command)	
	@	Number of local write accesses to JV entries, whereby only JV attributes are manipulated	
	0	Number of local write accesses to JV entries, whereby not only JV attributes are manipulated	
107	+	Average BCAM wait time	Seconds
	–	Overflow time for the BCAM wait time (see also WAIT-TIME-BUCKETS operand, SET-RESP-PARAMETER) Only BCAM pool wait times below the specified overflow value are included in the calculation of the average values.	
108	A	BCAM pool wait times in the 1st range	%
	.	BCAM pool wait times in the 2nd range	
	B	BCAM pool wait times in the 3rd range	
	*	BCAM pool wait times in the 4th range	
	C	BCAM pool wait times in the 5th range	
109	0	Average BCAM pool wait time not exceeding the limit	Seconds
	*	Average wait time in the BCAM pool BCAM pool wait times above the overflow range are also included in the calculation of this variable	
110	0	Average BCAM pool wait time not exceeding the limit per category	Seconds
	*	Average BCAM pool wait time per category BCAM pool wait times above the overflow range are also included in the calculation of this variable	
123	+	Number of SVC calls per SVC by TU	SVCs per second
	–	Number of SVC calls per SVC by TPR	
124	+	Load on disk device excluding paging activities	%
	–	Load on disk device due to paging activity	

Table 21: Variable reports

(part 19 of 38)

Report No.	Chart symbol	Variable	Unit
125	+	Number of I/O operations per disk device	Operations per second
126	+	Average length of the queue for the disk device including the jobs currently being handled	Number of I/O requests
127	+ -	Data volume for reading and writing Data volume for writing	PAM blocks per I/O operation
128	A D S	UTM applications Number of asynchronous transactions Number of dialog transactions Number of dialog steps	Number per second
129	U	Number of users signed on to the UTM application	Number of users
130	M T	Maximum number of tasks available for asynchronous conversations Number of tasks running for the application	Number of tasks
131	D V T .	Average time UTM waits for the execution of database calls per dialog step. All dialog steps are recorded, whether or not they contain database calls Average time UTM waits for a message to arrive from a remote application per dialog step. All dialog steps are recorded, whether or not they contain distributed processing calls Average time UTM waits as a result of TAC class bottlenecks per dialog step Average processing time per dialog step within the UTM task, measured from the time UTM accepts the message to the time UTM sends the message, excluding the parts covered by D, V and T	Seconds per dialog step

Table 21: Variable reports

(part 20 of 38)

Report No.	Chart symbol	Variable	Unit
132	D	Average time UTM waits for the execution of database calls per asynchronous conversation. All asynchronous conversations are recorded, whether or not they contain database calls	Seconds per asynchronous conversation
	V	Average time UTM waits for a message to arrive from a remote application per asynchronous conversation. All asynchronous conversations are recorded, whether or not they contain distributed processing calls	
	.	Average processing time per asynchronous conversation within the UTM task, measured from the start of the conversation to its end, excluding the parts covered by D and V	
133	A	Number of jobs for asynchronous programs waiting for processing	Number of jobs
	P	Number of print jobs waiting for execution	
	D	Number of time-driven jobs waiting for execution	
134	+	Number of hits during read accesses to hiperfiles	I/O operations per second
	–	Number of read accesses to hiperfiles	
135	+	Number of write accesses to hiperfiles	I/O operations per second
	–	Number of hits during write accesses to hiperfiles	
136	O	Unsuccessful attempts to use cache due to overload	I/O operations per second
137	–	Total number of SVC calls from TU	Number per second
	*	Total number of SVC calls from TPR	
141	F	Number of files determined in UFS file systems per second via the Inode entry	Number per second
142	S	Frequency per second with which a path name was searched for in the file system	Number per second
143	+	Read accesses to system buffers per second	Number per second
	–	Read data transfers per second between system buffer and hard disk or other block-oriented devices	

Table 21: Variable reports

(part 21 of 38)

Report No.	Chart symbol	Variable	Unit	
144	+	Write accesses to system buffers per second	Number per second	
		–		Write data transfers per second between system buffer and hard disk or other block-oriented devices
146	+	Read data transfers	Number per second	
		–		Write data transfers
147	S	All types of system calls per second	Number per second	
148	+	READ system calls	Number per second	
		.		WRITE system calls
		0		FORK system calls
		–		EXEC system calls
149	+	Number of characters per second transferred by READ system calls	Number per second	
		–		Number of characters per second transferred by WRITE system calls
150	O	Number of send and receive operations per second	Number per second	
151	O	Number of semaphore activities per second	Number per second	
152	S	Service rate (number of service units per second) of the task specified by TSN	Number per second	
153	C	Percentage of elapsed time accounted for by the task's TU+TPR time. In the case of a multiprocessor system the average value of the processors is output (values <= 100%).	%	
154	+	Number of all I/O operations per second	Number per second	
		–		Number of page reads per second
		.		Number of physical accesses to coded files
155	U	Average UPG of the task	4-KB pages	
156	S	Service rate (number of service units per second) of the task specified by the user ID	Number per second	
157	C	Percentage of elapsed time accounted for by the TU+TPR time of the task specified by user ID. In the case of a multiprocessor system the average value of the processors is output (values <= 100%).	%	

Table 21: Variable reports

(part 22 of 38)

Report No.	Chart symbol	Variable	Unit	
158	+	Number of all I/O operations per second	Number per second	
		–		Number of page reads per second
		.		Number of physical accesses to coded files
159	U	Average UPG of the task	4-KB pages	
160	S	Service rate (number of service units per second) of the task specified by the job name	Number per second	
161	C	Percentage of elapsed time accounted for by the TU+TPR time of the task specified by job name. In the case of a multiprocessor system the average value of the processors is output (values <= 100%).	%	
162	+	Number of all I/O operations per second	Number per second	
		–		Number of page reads per second
		.		Number of physical accesses to coded files
163	U	Average UPG of the task	4-KB pages	
164	*	Number of real CPUs, available to all BS2000 systems	Number of CPUs	
166	+	Number of send jobs	Number per second	
		–		Number of send jobs rejected due to overload
167	+	Number of server tasks requested originally	Number of tasks	
		.		Number of current server tasks
		–		Number of current server tasks which can cause a bottleneck
168	+	Average total time of REQUEST WITH REPLY jobs from the beginning of the job until the first reply	Milliseconds	
		–	Average wait time for REQUEST WITH REPLY jobs for the first reply from the receiving host	Milliseconds
169	*	Number of REQUESTS WITH REPLY	Number per second	

Table 21: Variable reports

(part 23 of 38)

Report No.	Chart symbol	Variable	Unit
170	+	Number of ENQUEUE LOCKREQUESTS from TU and TPR	Number per second
	–	Number of CONVERT LOCKREQUESTS from TU and TPR	
	.	Number of DEQUEUE LOCKREQUESTS from TU and TPR	
	*	Number of INFORMATION LOCKREQUESTS from TU and TPR	
171	+	Number of ENQUEUE LOCKREQUESTS to NSM	Number per second
	–	Number of CONVERT LOCKREQUESTS to NSM	
	.	Number of DEQUEUE LOCKREQUESTS to NSM	
172	+	Number of GRANT EVENTS to TU and TPR	Number per second
	–	Number of RELEASE EVENTS to TU and TPR	
173	+	Number of GRANT EVENTS from NSM	Number per second
	–	Number of RELEASE EVENTS from NSM	
174	+	Number of read accesses	Number per second
	–	Number of write accesses	
177	+	Number of read accesses to the global storage partition	Number per second
	–	Number of write accesses to the global storage partition	
178	+	Data rate for reading from the global storage partition	KB per second
	–	Data rate for writing to the global storage partition	
179	+	Total number of calls	Number per second
	.	Number of calls in which the lock server is on a remote host	
	–	Number of remote calls which had to wait for disk in the token	
180	+	Response time of the host to requests from other hosts	Milliseconds
181	.	Length of the NSM message buffer in tokens	KB
	A	Used length of NSM message buffer	
182	C	Time between two arrivals of the token	Milliseconds
184	+	Number of lock servers	Number

Table 21: Variable reports

(part 24 of 38)

Report No.	Chart symbol	Variable	Unit
185	+	Average number of requests to speedcat for single-feature pubsets	Number of requests
	.	Average number of requests to the catalog index for system-managed pubsets	
186	+	Number of TSDUs (TRANSPORT SERVICE DATA UNITS) sent	Number per second
	–	Number of TSDUs received	
187	+	Data rate for sent	KB per second
	–	Data rate for received	
189	*	Number of failed attempts to use a DAB cache area due to overload	I/O operations per second
190	*	Number of failed attempts to use a DAB subarea due to overload	I/O operations per second
191	*	Average access time for I/Os	Milliseconds per I/O
192	+	INWAIT times in 1st bucket	%
	.	INWAIT times in 2nd bucket	
	*	INWAIT times in 3rd bucket	
	–	INWAIT times in 4th bucket	
	@	INWAIT times in overflow bucket The bucket limits are indicated in the bar chart	
193	+	REACT times in 1st bucket	%
	.	REACT times in 2nd bucket	
	*	REACT times in 3rd bucket	
	-	REACT times in 4th bucket	
	@	REACT times in overflow bucket The bucket limits are listed in the bar chart	
194	+	INPROC times in 1st bucket	%
	.	INPROC times in 2nd bucket	
	*	INPROC times in 3rd bucket	
	–	INPROC times in 4th bucket	
	@	INPROC times in overflow bucket The bucket limits are listed in the bar chart	

Table 21: Variable reports

(part 25 of 38)

Report No.	Chart symbol	Variable	Unit
195	+ . * - @	OUTPROC times in 1st bucket OUTPROC times in 2nd bucket OUTPROC times in 3rd bucket OUTPROC times in 4th bucket OUTPROC times in overflow bucket The bucket limits are listed in the bar chart	%
196	* .	Number of TSDU's (Transport Service Data Units) for receive jobs Number of TSDU's for send jobs	Number per second
197	* .	Data volume received Data volume sent	KB per TSDU
198	* .	Maximum cache area for received data Occupied cache area for received data	KB
199	* .	Maximum cache area for data for sending Occupied cache area for data for sending	KB
200	* .	Number of received packets without user data Number of sent packets without user data	Number per second
201	* .	Number of received packets with user data Number of sent packets with user data	Number per second
202	*	Number of send jobs during resource bottleneck	Number per second
203	*	Number of packets with Zero Window Information; i.e. the partner application prohibits the transmission of data	Number per second
205	*	Number of blocks still to be relocated to disk	2-KB pages
206	*	Number of I/O's to write cache data back to disk	Number per second
216	+ . 0 -	Percentage of runtime in user mode Percentage of runtime in system mode Percentage of time in inactive state; process waiting for block-oriented I/O Percentage of time in inactive state	%

Table 21: Variable reports

(part 26 of 38)

Report No.	Chart symbol	Variable	Unit
218	+	Read accesses to system buffer	Number per second
	–	Read data transfers between system buffer and disk or other block-oriented devices	
219	+	Write accesses to system buffer	Number per second
	–	Write data transfers between system buffer and disk or other block-oriented devices	
220	+	Read data transfers	Number per second
	–	Write data transfers	
221	*	All types of system calls	Number per second
222	+	SREAD system calls	Number per second
	.	SWRITE system calls	
	0	FORK system calls	
	–	EXEC system calls	
223	+	Number of bytes transferred by READ system calls	Number per second
	–	Number of bytes transferred by WRITE system calls	
224	+	Number of index accesses	Number per second
	–	Number of index accesses for which no read operation from disk was required	
225	*	Time spent by a job in the job queue of the UTM application (this value is calculated by generating an internal test message)	Milliseconds
226	+	Total of the number of active connections which existed during the entire monitoring cycle and the number of connections which were opened and/or closed during the monitoring cycle	Number
	.	Number of active connections during the entire monitoring cycle	
	*	Number of connections which were opened and/or closed during the monitoring cycle	
	–	Number of rejected attempts to open a connection	
227	+	Average software service time (including hardware service time) for I/O operations of the specified disk	Milliseconds
	–	Average hardware service time for I/O operations of the specified disk	

Table 21: Variable reports

(part 27 of 38)

Report No.	Chart symbol	Variable	Unit
228	+	Data rate for reading and writing	KB per second
	–	Data rate for writing	
229	*	Number of alias device names assigned to a basic device name	Number
230	+	Data rate for reading and writing	KB per second
	–	Data rate for writing	
231	1	Device disconnect time for I/O operations of the specified device	Milliseconds
	2	Device connect time for I/O operations of the specified device	
	+	(Remaining) service time for I/O operations of the specified device	
	–	Function pending time for I/O operations of the specified device	
	*	Device queue time for I/O operations of the specified device If DCS (Dynamic Channel Subsystem) provides no data, only the (remaining) service time and the device queue time are output.	
232	+	Number of read accesses for the SYMMETRIX system	Number per second
	–	Number of read accesses for the SYMMETRIX system which were responded to by the cache	
233	+	Number of write accesses for the SYMMETRIX system	Number per second
	–	Number of write accesses for the SYMMETRIX system which were responded to by the cache	
234	+	Total number of read and write accesses for the SYMMETRIX system	Number per second
	–	Total number of read and write accesses for the SYMMETRIX system which were responded to by the cache	
235	+	Data rate for reading and writing for the SYMMETRIX system	KB per second
	.	Data rate for reading for the SYMMETRIX system	
	–	Data rate for writing for the SYMMETRIX system	

Table 21: Variable reports

(part 28 of 38)

Report No.	Chart symbol	Variable	Unit
236	+	Number of sequential read accesses for the SYMMETRIX system	Number per second
237	+	Number of read accesses for the pubset	Number per second
	–	Number of read accesses for the pubset which were responded to by the cache	
238	+	Number of write accesses for the pubset	Number per second
	–	Number of write accesses for the pubset which were responded to by the cache	
239	+	Total number of read and write accesses for the pubset	Number per second
	–	Total number of read and write accesses for the pubset which were responded to by the cache	
240	+	Data rate for reading and writing for the pubset	KB per second
	.	Data rate for reaing for the pubset	
	–	Data rate for writing for the pubset	
241	+	Number of sequential read accesses for the pubset	KB per second
	–	Number of sequential read accesses for the pubset which were responded to by the cache	
242	+	Number of sequential write accesses for the pubset	Number per second
243	+	Rate of read and write accesses for the pubset between the devices (logical volumes) and the physical disks on the SYMMETRIX system	Number per second
	.	Rate of read accesses for the pubset between the devices (logical volumes) and the physical disks on the SYMMETRIX system	
	–	Rate of write accesses for the pubset between the devices (logical volumes) and the physical disks on the SYMMETRIX system	

Table 21: Variable reports

(part 29 of 38)

Report No.	Chart symbol	Variable	Unit
244	+	Data rate for reading and writing for the pubset between the devices (logical volumes) and the physical disks on the SYMMETRIX system	KB per second
	.	Data rate for reading for the pubset between the devices (logical volumes) and the physical disks on the SYMMETRIX system	
	–	Data rate for writing for the pubset between the devices (logical volumes) and the physical disks on the SYMMETRIX system	
245	+	Number of read accesses for the device (logical volume)	Number per second
	–	Number of read accesses for the device (logical volume) which were responded to by the cache	
246	+	Number of write accesses for the device (logical volume)	Number per second
	–	Number of write accesses for the device (logical volume) which were responded to by the cache	
247	+	Total number of read and write accesses for the device (logical volume)	Number per second
	–	Total number of read and write accesses for the device (logical volume) which were responded to by the cache	
248	+	Data rate for reading and writing for the device (logical volume)	KB per second
	.	Data rate for reading for the device (logical volume)	
	–	Data rate for writing for the device (logical volume)	
249	+	Number of sequential read accesses for the device (logical volume)	KB per second
	–	Number of sequential read accesses for the device (logical volume) which were responded to by the cache	
250	+	Number of sequential write accesses for the device (logical volume)	Number per second

Table 21: Variable reports

(part 30 of 38)

Report No.	Chart symbol	Variable	Unit
251	+	Rate of read and write accesses for the device (logical volume) between the cache and the physical disk on the SYMMETRIX system	Number per second
	.	Rate of read accesses for the device (logical volume) between the cache and the physical disk on the SYMMETRIX system	
	–	Rate of write accesses for the device (logical volume) between the cache and the physical disk on the SYMMETRIX system	
252	+	Data rate for reading and writing for the device (logical volume) between the cache and the physical disk on the SYMMETRIX system	KB per second
	.	Data rate for reading for the device (logical volume) between the cache and the physical disk on the SYMMETRIX system	
	–	Data rate for writing for the device (logical volume) between the cache and the physical disk on the SYMMETRIX system	
253	+	Rate of read and write accesses for the physical disk on the SYMMETRIX system	Number per second
	.	Rate of read accesses for the physical disk on the SYMMETRIX system	
	–	Rate of write accesses for the physical disk on the SYMMETRIX system	
254	+	Data rate for reading and writing for the physical disk on the SYMMETRIX system	KB per second
	.	Data rate for reading for the physical disk on the SYMMETRIX system	
	–	Data rate for writing for the physical disk on the SYMMETRIX system	

Table 21: Variable reports

(part 31 of 38)

Report No.	Chart symbol	Variable	Unit
255	+	Rate of accesses for reading and writing of the director of the SYMMETRIX system	Number per second
	.	Rate of accesses for reading of the director of the SYMMETRIX system	
	–	Rate of accesses for writing of the director of the SYMMETRIX system	
	@	Rate of accesses for reading and writing of the director of the SYMMETRIX system which were responded to by the cache (0 for type DA (back end) and type RA (remote access))	
256	+	Data rate for sent and received of a RA (remote link) director to a remote SYMMETRIX system	KB per second
	.	Data rate for received of a RA (remote link) director from a remote SYMMETRIX system	
	–	Data rate for sent of a RA (remote link) director to a remote SYMMETRIX system	
257	+	Data volume transferred per channel via PAM-block transfer	KB per second
	–	Data volume transferred per channel via byte transfer	
258	*	Average INWAIT time	Milliseconds
259	*	Average REACT time	Milliseconds
260	*	Average INPROC time	Milliseconds
261	*	Average OUTPROC time	Milliseconds
262	+	Capacity of the SF pubsets or volume sets	GB
	–	Space allocation of the SF pubsets or volume sets	
263	*	Saturation level (0-6)	
264	*	Relative utilization of the SF pubset or volume set	%
267	–	Monitored active time of the hypervisor in relation to the CPU pool (servers with /390 architecture only)	%
	.	Monitored IDLE time in relation to the CPU pool (servers with /390 architecture only)	
	*	Utilization of the CPU pool by all virtual machines assigned to the CPU pool	
268	*	Number of available real CPUs of the CPU pool	Number of CPUs

Table 21: Variable reports

(part 32 of 38)

Report No.	Chart symbol	Variable	Unit
269	*	Maximum CPU utilization per VM group	%
	.	Planned CPU utilization per VM group	
	–	Monitored CPU utilization per VM group (all data for servers with /390 architecture only)	
270	*	Average number of RSC IOs processed in parallel	Number
	.	Maximum possible number of RSC IOs	
271	*	Number of ISAM accesses where the desired page is already in the ISAM pool, i.e. no read operation from disk is required (FIX-HIT)	Number per second
	.	Number of ISAM accesses where a read operation from disk was required (FIX-IO)	
	–	Number of ISAM accesses involving a waiting period for the release of one or more buffer pages (FIXWAIT)	
272	+	Number of RESERVE-SLOT operations leading to a wait state due to a slot bottleneck	Number per second
	–	Number of RESERVE-SLOT operations where a slot is available and no wait state occurs	
273	*	Average number of pages in the ISAM pool	PAM pages
	.	Average number of reserved pages in the ISAM pool	
	–	Average number of fixed pages in the ISAM pool	
274	+	Number of index accesses	Number per second
	–	Number of index accesses for which no read operation from disk was required	
275	+	Number of used Big Pages	Number
	–	Total number of Big Pages	
280	–	Maximum size of the resident memory for data transfer (input)	MB
	.	Threshold value for the size of the resident memory for data transfer (input) from which BCAM issues warnings on the console	
	*	Current size of the resident memory for data transfer (input)	

Table 21: Variable reports

(part 33 of 38)

Report No.	Chart symbol	Variable	Unit
281	– . *	Maximum size of the resident memory for data transfer (output) Threshold value for the size of the resident memory for data transfer (output) from which BCAM issues warnings on the console Current size of the resident memory for data transfer (output)	MB
282	+ –	Data set for reading and writing Data set for writing (In the output file report 282 uses a different data area from report 100. The data area used by report 100 von can only present data rates < 240 MB/s. Report 282 also enables higher data rates to be processed. For this purpose the data area used in report 282 only contains data rates with an accuracy of > 1 KB/s in the output file. 0 is output for lower values.)	Bytes per input/output
283	+ –	Data rate for reading and writing Data rate for writing (In the output file report 283 uses a different data area from report 230. The data area used by report 230 von can only present data rates < 240 MB/s. Report 283 also enables higher data rates to be processed. For this purpose the data area used in report 283 only contains data rates with an accuracy of > 1 KB/s in the output file. 0 is output for lower values.)	KB per second
284	– . * @	Number of DML calls Number of SQL calls Number of SQL transactions Number of SQ operations	Number per second
285	+ –	Number of transactions Number of update transactions	Number per second

Table 21: Variable reports

(part 34 of 38)

Report No.	Chart symbol	Variable	Unit
286	–	Number of update transactions with remote partners	Number per second
	.	Number of CODASYL statements sent to remote partners	
	*	Number of CODASYL statements received from remote partners	
287	+	Number of logical read accesses	Number per second
	–	Number of physical read accesses	
288	+	Number of logical write accesses	Number per second
	–	Number of physical write accesses	
289	–	Number of physical read accesses to the RLOG file	Number per second
	.	Number of physical write accesses to the RLOG file	
	*	Number of physical write accesses to the ALOG file	
290	–	Number of wait states for locks	Number per second
	.	Number of data deadlocks	
	*	Number of task deadlocks	
291	–	Number of ITC communication operations between user and server tasks	Number per second
	.	Number of ITC communication operations between server and user tasks	
	*	Number of all requests from server tasks to server tasks	
	@	Number of ITC communication operations between server tasks and server tasks	
292	*	Rate of successful accesses with PPP for the entire UDS/SQL session	%
293	*	Average transaction time for the entire UDS/SQL session	Second
294	+	Number of local CALL-DML statements	Number per second
	–	Number of local SQL statements	
295	+	Number of remote CALL-DML statements	Number per second
	–	Number of remote SQL statements	

Table 21: Variable reports

(part 35 of 38)

Report No.	Chart symbol	Variable	Unit
296	+	Number of transactions	Number per second
	.	Number of update transactions	
	*	Number of retrieval transactions	
	–	Number of reset transactions	
	@	Number of transaction locks	
297	–	Number of SQL update statements	Number per second
	.	Number of SQL retrieval statements	
	*	Number of SQL-DDL and SQL-SSL statements	
	@	Number of utility statements	
298	+	Number of SQL plan accesses	Number per second
	–	Number of SQL plan generations	
299	+	Number of logical read accesses to the system data buffer	Number per second
	–	Number of logical write accesses to the system data buffer	
300	+	Number of physical read accesses to the system data buffer	Number per second
	–	Number of physical write accesses to the system data buffer	
301	+	Number of logical read accesses to the user data buffer	Number per second
	–	Number of logical write accesses to the user data buffer	
302	+	Number of physical read accesses to the user data buffer	Number per second
	–	Number of physical write accesses to the user data buffer	
303	+	Number of logical read accesses to the cursor files	Number per second
	–	Number of logical write accesses to the cursor files	
304	+	Number of physical read accesses auf die cursor files	Number per second
	–	Number of physical write accesses to the cursor files	

Table 21: Variable reports

(part 36 of 38)

Report No.	Chart symbol	Variable	Unit
305	–	Number of physical write accesses to DA-LOG and CAT-LOG files	Number per second
	.	Number of physical read accesses to TA-LOG files	
	*	Number of physical write accesses to TA-LOG files	
306	–	Number of active threads	Number
	.	Number of ready threads	
	*	Number of unused threads	
307	–	Number of inactive threads because of IO	Number
	.	Number of inactive threads because of TA lock	
	*	Number of inactive threads because of job execution by the service task	
	@	Number of inactive threads because of semaphores, block lock or termination of a system job	
308	+	Number of active service tasks	Number
	–	Number of jobs not yet executed by the service task	
309	+	Data rate for sending to and receiving from the network	KB per second
310	+	Data rate for sending to the network	KB per second
	-	Data rate for receiving from the network	
311	+	Data rate for reading from and writing to files	KB per second
312	+	Data rate for reading from files	KB per second
	-	Data rate for writing to files	
313	+	Number of received jobs per second	Number per second
	.	Number of successfully completed jobs per second	
	*	Number of aborted jobs per second	
	-	Number of interrupted jobs per second	
	@	Number of jobs rejected per second because of incorrect user authentication	
314	+	Number of connection attempts aborted as unsuccessful per second	Number per second
	-	Number of connections aborted per second	
315	+	Number of synchronous jobs in ACTIVE state	Number
	-	Number of asynchronous jobs in ACTIVE state	

Table 21: Variable reports

(part 37 of 38)

Report No.	Chart symbol	Variable	Unit
316	- . * @	Number of jobs in WAIT state Number of jobs in HOLD state Number of jobs in SUSPEND state Number of jobs in LOCK state	Number
317	+ -	Maximum number of connections active in parallel which can be used to execute file transfer requests (openFT parameter CONNECTION-LIMIT) Number of occupied connections. This value can temporarily be greater than the maximum value if CONNECTION-LIMIT has been reduced.	Number
318	+ -	Maximum number of jobs which can be stored in the request queue (asynchronous jobs), openFT parameter REQUEST-LIMIT) Number of jobs stored in the request queue (asynchronous jobs). This value can temporarily be greater than the maximum value if REQUEST-LIMIT has been reduced.	Number
319	* .	Average number of RSC IOs processed in parallel Maximum possible number of RSC IOs	Number
320	+ -	Number of read accesses Number of write accesses See also notes on the DISK-FILE monitoring program on page 59 .	Inputs/outputs per second

Table 21: Variable reports

(part 38 of 38)

Glossary

The following terms have been used frequently in this manual. They are explained below:

activation

Whenever a task issues a request to the system after a period of inactivity (e.g. think time at a terminal), the system must take two decisions before processing can be continued:

1. The activation decision
2. The initiation decision

activation decision

This decision gives the task the right to use the CPU and to perform I/O operations. At activation time, the task is allocated as many main memory pages as it is likely to need.

Activation delays are possible in cases of resources overloading.

assignment of paging activities to the initiating task

The system initiates all paging I/O operations directly in the SIH state. However, SM2 does not assign all paging I/O operations to the SYSTEM category, but sometimes to the initiating task.

For SM2, the initiating task is that task which caused a page fault during page access. Two chains have to be distinguished:

1. Only one page is read.
The full firmware duration of the output operation is assigned to the initiating task and its category. The full software duration of the input operation is assigned to the category of the initiating task.
The task and category counters for the number of paging operations are incremented by 1.
2. Only page output occurs.
The full firmware duration of the output operation is assigned to the PGE task, but not to the SYSTEM category (although it is registered system-globally for SUM).
The PGE task counter for the number of paging operations is incremented by 1.

background storage

Storage area on peripheral devices which can be accessed by means of virtual addressing. Programs and data are transferred in pages from background storage to main memory before processing and returned (if required) to background storage after processing.

Big Pages

Big Pages (4 MB) are used for JIT compilations for a better performance on servers with /x86 architecture.

caching

Caching is understood to be the process of buffering data in a fast data medium (the cache) in an attempt to accelerate subsequent inputs and outputs to the same data areas.

Data to be written to or read from disk is buffered in the cache to avoid the longer input/output times involved in accessing the disk.

If the data to be accessed is in the cache at the time of access, this is referred to as a cache hit, otherwise as a cache miss.

The proportion of hits in the total number of accesses is known as the cache hit rate. The higher the hit rate, the greater the advantage of using the cache. The hit rate which can be achieved depends on a range of factors, such as the locality of the accesses, the size of the cache, the caching method selected (read cache, write cache, read/write cache) and the appropriate selection of files. Monitoring systems such as SM2 can be used to identify files and disks suitable for caching.

counting of I/O operations and their duration

For SM2, an I/O operation is

- any version of EXCP (EXCP, EXCPW, \$EXCP, \$EXCPW) or
- an I/O request from the memory management system for paging (no SVC call) addressed to the I/O control module. The I/O control module is the central system component for handling physical I/O operations.

The I/O control module normally processes the I/O request with one privileged instruction (Start Device or Start Subchannel). For each instruction a termination message is issued, upon which the I/O operation is counted.

In rare cases (e.g. disk connected to a channel operated in selector mode), two SDV instructions are issued. This is called “offline seek”. The first SDV prepares the I/O device and the second initiates the I/O operation proper. Only one I/O operation is counted in this case too.

The following particularities should be noted:

I/O requests which could not be started properly (SDV with condition code differing from 0) are ignored.

So-called SENSE SDVs, which are generally issued after an I/O request that resulted in an error and serve to fetch additional diagnostic information, are only included in report 100.

In SDV fast release mode supported in some servers, a properly initiated I/O operation can be rejected at a later time by a channel. The rejected SDVs are not counted.

In determining the duration of the I/O operation, the time between the SDV and the termination message is always counted. An SDV rejected because of an error is ignored; this results in an increase in waiting time for the device. An SDV rejected in SDV fast release mode is ignored; this also results in an increase in waiting time for the device.

In “offline seek” mode, the time between the first SDV and the termination message for the second SDV is counted.

CPU states

Program interrupts are caused by

- input/output requests
- calls to the Control System
- timers
- errors
- paging requests

The system distinguishes between the following CPU states when handling programs and interrupts:

TU	TPR	SIH	MEH
Processing state for user programs	Processing state for program interrupts	Recognition state for program interrupts	Recognition state for machine errors
Interruptible	Interruptible; in part non-interruptible	Interruptible by MER only	Non-interruptible
Nonprivileged	Privileged CPU states		

If the CPU is in none of the above states, it is in the IDLE state.

deactivation, forced deactivation

When a task is deactivated, it is no longer authorized to use the CPU. In the case of forced deactivation, the system withdraws the task's right to use the CPU. Forced deactivation can occur when resources are extremely overloaded.

dilation factor

Dilation factors can be determined for individual tasks, for specific task classes and on a system-global basis.

Dilation factor = Dwell time / Productive RST

Productive RST = productive CPU RST + productive I/O RST

For further information on productive RST, see [“time equivalent for the productive performance” on page 673](#).

A dilation factor smaller than 1 can occur when a task uses asynchronous I/O. In this case, the task uses 2 or more resources simultaneously.

The dilation factor output in SM2R1 report 57 is not comparable with the value “REQUEST DELAY” in the SM2 PCS report and in SM2R1 report 74. This quantity is designed for optimum response time control by PCS.

DMS I/O operations

All accesses to peripheral devices not performed for paging.

dwelt time

The time spent by a task in the system comprises the service times of the physical resources and all non-voluntary wait times of the task.

It is not always possible to distinguish between voluntary and non-voluntary wait times without unreasonable outlay.

Example

The BS2000 boursing mechanism can be used to determine both the voluntary wait times (e.g. wait for user input) and the non-voluntary wait times (wait for a busy resource).

For simplicity’s sake, SM2 considers all waits times for boursing, for the PASS and VPASS macros, and for responses to a console question (/INFORM-OPERATOR WAIT-RESPONSE=*YES or TYPIO macro with response) as voluntary wait times.

SM2 includes the following times in the dwelt time:

1. Time spent in the queue for the CPU(s).
This includes the CPU RST and the time spent in Q1.
2. Wait for paging page transfer (time spent in Q3).
3. Wait for execution of DMS I/O operations (time spent in Q4 and Q12 for I/O).
4. Further time spent in Q4 except for boursing, ITC, PASS/VPASS and waiting for response to console questions.
5. Wait for activation (time spent in Q5).
6. Wait for admission (time spent in queue Q6).

7. Wait in Q7 for hardware error recovery (HERS).
8. Further time spent in Q12 except for boursing and waiting for response to console message.
9. Hardware duration of I/O operations for paging page transfers.

EXCP

Macro on the physical I/O level which initiates a channel program (see [“counting of I/O operations and their duration” on page 663](#)).

global storage (GS)

Global storage is a form of semiconductor-based expansion memory which has been made non-volatile by a number of measures. The GS is accessed using synchronous machine instructions, i.e. no task switching takes place and the access speed is considerably increased compared to the disk I/Os.

hardware service time

Time for which devices are busy with I/O operations, also called hardware duration. The hardware duration is defined as the time between I/O initiation (start subchannel) and device end (interrupt).

See also [“service times definition” on page 671](#).

hiperfile concept

Hiperfiles (high-performance file) in BS2000 are a means of accelerating file processing. This concept is based on the use of particularly fast storage media when processing a file in order to avoid I/O bottlenecks and thus increase the performance of the entire system.

These storage media are used to buffer data to be written to or read from disk. This avoids the longer I/O times associated with disk accesses.

Main memory or global storage can be used as storage media.

In this manual, hiperfiles are files with the attribute PERFORMANCE=*HIGH or *VERY-HIGH.

initiation

After a task has been activated, the system must reach a decision on initiation. When it is initiated, the task can use a CPU, i.e. compute.

input/output, logical level

On the logical level, the user uses macro calls (e.g. GET, PUT) to initiate system functions which control data interchange with the peripheral devices, block and unblock data, and handle any errors that may occur.

interaction

Generic term for wait, response, think, and transaction processes (RESPONSETIME monitoring program).

main memory

Memory area which can be addressed directly by the CPU. The program instructions are read and processed in this area.

overall performance

The productive performance and the dilation, i.e. the relation between the time spent in the system (dwell time) and the time required for providing the productive performance, are important criteria for accessing the suitability of an IT system for a given load.

The following definition is used:

Overall performance = productive performance + housekeeping performance

Productive performance is the load handling capacity.

Housekeeping (or overhead) performance is the performance required for system control by the operating system.

Overall performance is the sum of these two quantities.

The performance is provided by the various IT system components, e.g. the CPU(s) and the peripheral devices.

Performance = work done per unit of time. A precise definition of "work done" is not attempted here. Only the time required for productive or housekeeping performance is considered.

The time equivalent of the productive performance or of the housekeeping performance is the (service) time for which resources are being used for productive work or for housekeeping.

It is also called Resources Service Time (RST).

The dwell time of a task is the sum of the physical resources service times plus all non-voluntary wait times of the task. The non-voluntary wait time of a task is thus a criteria for task obstruction by other tasks in the system.

page fault

BS2000 is an operating system with virtual addressing, i.e. it supports several address spaces simultaneously. The virtual address spaces and the real main memory are managed in units of 4-Kb pages. (On servers with x86 architecture, main memory is managed in 8-Kb pages.) The DEMAND PAGING method is used for mapping virtual pages onto real memory page frames:

When an attempt is made to access a page that is not in main memory, the hardware detects this condition and uses a page fault interrupt to notify the operating system. If the addressed page is on disk (paging device), the page is read in (page transfer).

If free page frames are required, the system tries to return modified main memory pages to disk.

SM2 counts the number of paging I/O operations (number of I/O requests to the central I/O control module of the system, which corresponds to the EXCP calls). This number is supplied in the ACTIVITY report and in SM2R1 reports 3 and 4.

For each I/O request, SM2 increments this number by 1 when

- a page is read in or
- one or more pages are written (the system tries to group up to 8 pages for one I/O operation).

The number of pages written is also given in the MEMORY report and in report 55, and the number of pages read is given in report 56.

SM2 supplies the following additional data:

1. Total number of page fault interrupts. “Real” page faults are not included in this number.
2. Number of page fault interrupts for which the addressed page was still in main memory (PAGE RECLAIMS).
3. Number of page fault interrupts for which a page transfer is required. This count is also incremented when 2 or more tasks try simultaneously to access the same virtual page and the page is not in main memory.
4. Number of page fault interrupts for the first access to a new page.

paging in BS2000

The number of addressable virtual pages is usually greater than the number of page frames available in main memory.

For the removal of pages which can no longer be kept in main memory, see [“page fault” on page 668](#).

The main memory management strategies are based on the assumption that the programs (or tasks) will only address a limited set of the total number of virtual pages within a specific period of time, i.e. that the programs are more or less local. This set of pages, whose size varies dynamically, is called the working set.

The PPC (planned page count) indicates how local a program or task is. This value is provided by the memory management system before task activation and – because it is a measure of the intensity of memory utilization by the task – is used as a task activation criterion.

The PPC value for the next activation phase depends on program behavior during the preceding activation phase.

While the task is active, the PPC value can be modified in accordance with program behavior.

The number of pages used by a task is called UPG (USED PAGE COUNT).

The main memory is managed system-global (real memory management). This means that, when determining the pages which should be removed because of a shortage of free real-memory page frames, all main memory pages are checked.

The LRU principle is used for determining the pages to be removed from main memory: To this end the main memory pages are sorted according to their access times.

The PPC value is supplied as a measure of how local a task is.

The check is made and working set pages are removed only if free page frames are required. This is why UPG may become greater than the working set, especially if the main memory is large.

Free pages or pages not actually assigned to a task are in the FREE POOL (read-only- or read-write-queue) or in the empty queue.

paging I/O operations

All I/O operations performed in response to paging requests.

paging memory

The whole main memory area available to all users for paging.

paging rate

Number of required paging I/O operations per second (calls to the I/O control module).

path info

A 3-byte address comprising the channel address (2 bytes) and the device address of a complete data path between main memory and the device.

PAV

Parallel Access Volume

The PAV function allows a number of I/Os to be executed simultaneously to a logical volume. This increases the data rate for volumes with high utilization.

PDT

The PDT (Physical Device Table) is the central device table of the input/output system which contains an entry for each device. In addition to the device name and type, the entries also contain details about device characteristics which must be borne in mind when the devices are operated.

PGE task

System task required for restarting after hardware errors during paging I/O.

response time

Time required by the system for processing a request. For the user, this is the time between an input operation and the corresponding system response (see measurement program RESPONSETIME on [page 63](#) and measurement program BCAM-CONNECTION on [page 52](#)).

RSC IOs

The term RSC IO (I/O request with “Remote System Call” on servers with x86 architecture) refers to the execution of an I/O request with a high-performance interface between BS2000 and X2000.

BS2000 defines the I/O request in the “open” format expected by the FC peripherals. This format enables multiple I/Os to be started in parallel with disks.

service times definition (DCS, I/O operations)

Detailed monitoring of service times is possible on XS systems with DCS, the interface between the operating system, the I/O processor and the physical devices.

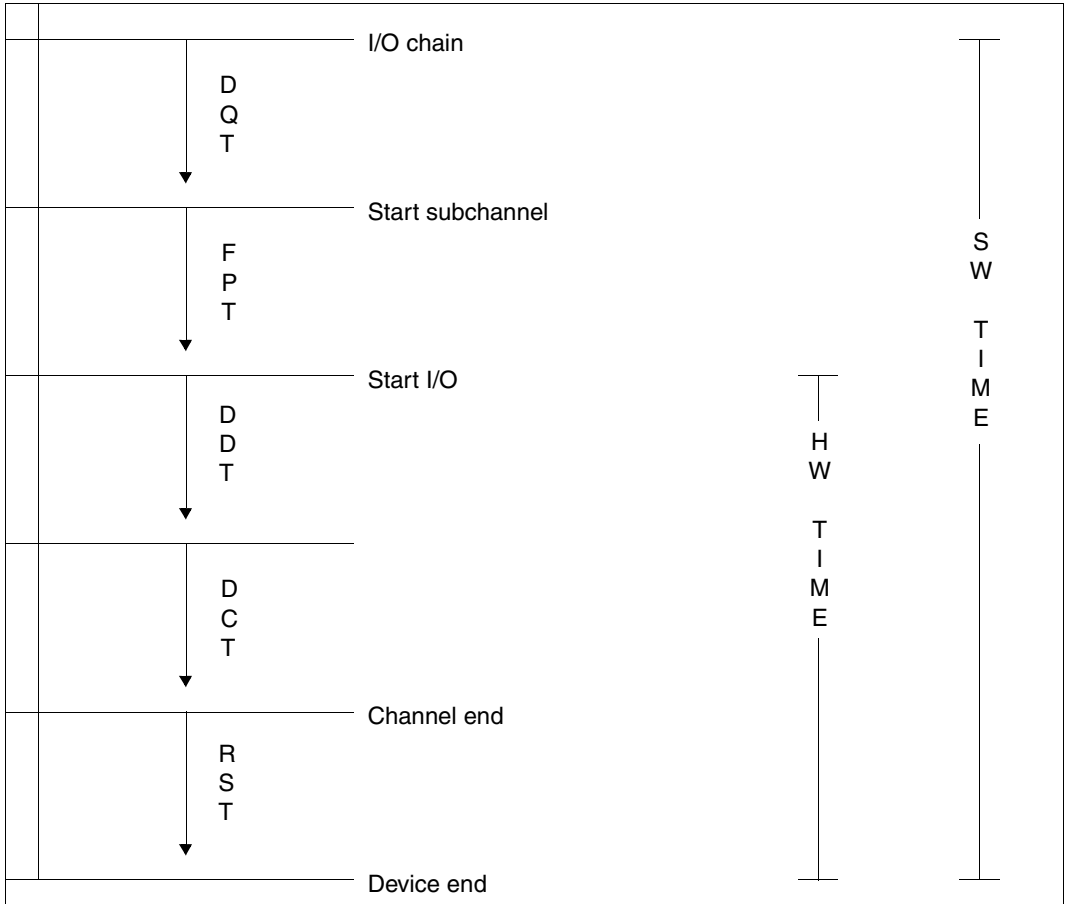


Figure 21: Service times

Service time	System	Definition
DQT Device queue time	SM2	Time from the I/O chain to the subchannel start (waiting time before the device)
FPT Function pending time	DCS	Time from the subchannel start to I/O start (waiting time for the assignment of a path; part of the hardware service time)
DDT Device disconnect time	DCS	Physical positioning time (part of the hardware service time)
DCT Device connect time	DCS	Data transfer time (part of the hardware service time)
RST Remaining service time	SM2	Time from channel end to device end (part of the hardware service time) or if DCS delivers no data: time from subchannel start to device end (hardware service time)

Table 22: Service times

software service time

The software service time (also called software duration) is obtained by adding the waiting time of an input/output request in the device queue of the system to the hardware duration. See also “service times definition” on [page 671](#).

queues

- Q0 A task is in this queue when it is using the CPU.
- Q1 Tasks in this queue are waiting to use the CPU.
- Q2 This is the queue for the SM2 write task.
- Q3 Tasks waiting for end of paging.
- Q4 Tasks waiting for peripheral I/O termination (disk, tape), for task communication events (boursing, ITC), or for VPASS where msec is specified.
- Q5 Ready tasks waiting to become active again.
- Q6 .Ready tasks waiting to be admitted (PCS).

The task queues described in the following contain tasks waiting for an event that is very remote in time before they become ready. These tasks are deactivated, i.e. they have lost the right to use main memory.

- Q7 Not used
- Q8 Not used

- Q9 Not used
- Q10 Hold queue for tasks placed in the hold state by the system or the operator in an emergency or overload situation. Tasks in the device queue. Tasks which cannot be terminated under normal conditions but have not yet been completed (error). Newly generated tasks.
- Q11 System tasks which are not called on a time basis.
- Q12 Tasks in the WHEN queue or tasks waiting for an event remote in time, e.g. time-consuming inter-task communication events (boursing, especially waiting for terminal input in interactive mode).
- Q13 Tasks waiting for VPASS or PASS end.

The task queues are implemented by chaining the TCBs (task control blocks). In the system, a separate entry refers to the first TCB of each queue. Task queues Q0 to Q4 exist once for each central processor. Queues Q5 to Q13 are system-global.

time equivalent for the housekeeping performance

The housekeeping performance of the system is a measure of the operating system overhead caused by the workload to be processed. Its definition encompasses the following factors:

1. CPU service time (SIH time) for processing paging requests.
2. Further SIH times for the operating system.
This is for simplification only, because various operations performed in the SIH state constitute productive work. Part of the productive performance for I/O execution occurs in the SIH state. System activities for processing SVC calls (SVC frame processing) are also handled in the SIH state. As far as monitoring is concerned, it would be very complicated to assign the SIH part of the productive performance to the individual tasks.
3. Hardware duration of the I/O operation for paging page transfers.

time equivalent for the productive performance

In determining the RST (resources service time), a distinction is made between the CPU and the peripheral devices as follows:

1. For the CPU the RST is the time in which instructions are processed in the TU and TPR states. This time is called productive CPU RST (the SIH share of the productive performance is not counted; see above).

2. For the peripheral equipment, the RST is the firmware service time of the devices for performing an I/O operation.

The firmware duration is defined as the time from I/O initiation (START-DEVICE or START-SUBCHANNEL instruction) to I/O termination (channel interrupt). Only I/O operations initiated by DMS are counted. This share is called productive I/O RST.

In practical operation, the firmware duration is affected by other tasks in the system. The positioning time during disk access, for example, depends on prior events, i.e. the current arm position, and on operating system strategies. In BS2000, arm movement optimization is used, i.e. the operating system tries to start I/O accesses in such a way that arm movements are minimized.

By adding the waiting time of an I/O request in the device queue of the system to the firmware duration, the software service time (also called software duration) is obtained.

transaction

Total number of system responses to a user request.

(Please refer to the normal usage in the UTM manuals for the meaning of the term “transaction” in the UTM reports).

virtual address space subdivision

The virtual address space is subdivided as follows:

1. Class 1 memory

Resident memory for system module code.

The size of this memory segment is specified at system generation time and remains constant for the session.

2. Class 2 memory

Paging memory for system module code.

The size of this memory segment depends on generation specifications and remains constant for the session.

3. Class 3 memory

Resident memory requested dynamically for tables, control blocks and overlay modules.

The size of this memory segment varies during the session.

4. Class 4 memory

Paging memory requested dynamically for tables, control blocks and overlay modules (also shared modules).

The size of this memory segment varies during the session.

working set see “[paging in BS2000](#)” on [page 669](#)

Related publications

You will find the manuals on the internet at <http://manuals.ts.fujitsu.com>. You can order printed copies of those manuals which are displayed with an order number.

- [1] **BS2000 OSD/BC
Utility Routines**
User Guide
- [2] **BS2000 OSD/BC
DMS Macros**
User Guide
- [3] **BS2000 OSD/BC
Commands**
User Guide
- [4] **BS2000 OSD/BC
Executive Macros**
User Guide
- [5] **BS2000 OSD/BC
Performance Handbook**
User Guide
- [6] **BS2000 OSD/BC
Introduction to System Administration**
User Guide
- [7] **DAB (BS2000)
Disk Access Buffer**
User Guide
- [8] **HIPLEX MSCF (BS2000)
BS2000 Processor Networks**
User Guide

- [9] **openFT for BS2000
Installation and Administration**
System Administrator Guide
- [10] **openNet Server (BS2000)
BCAM**
User Guide
- [11] **openUTM (BS2000)
Using openUTM Applications under BS2000**
User Guide
- [12] **PCS (BS2000)
Performance Control Subsystem**
User Guide
- [13] **SESAM/SQL-Server (BS2000)
Database Operation**
User Guide
- [14] **SHC-OSD / SCCA-BS2 (BS2000)
Storage Management for BS2000**
User Guide
- [15] **SM2-PA (BS2000)
SM2 Program Analyzer**
User Guide
- [16] **UDS/SQL-Server (BS2000)
Database Operation**
User Guide
- [17] **VM2000 (BS2000)
Virtual Machine System**
User Guide
- [18] **FUJITSU Server BS2000 SE Series
Operation and Administration**
User Guide

Index

A

accumulated 408
ACF report 218
activation 661
activation decision 661
ACTIVE PARAMETER screen 324
ACTIVITY report 220
ADD-CONNECTION-SET statement 91
ADD-COSMOS-EVENT statement 93
ADD-FILE statement 95
ADD-ISAM-FILE statement 96
ADD-ISAM-POOL statement 97, 98
administration facility 99
 switch to 99
admit additional SM2 administrators 102
analysis
 task statistics 461
analysis interval
 setting (SM2R1) 473
analysis period 391
 setting (SM2R1) 473
analysis routines 507
analysis subinterval 404
ANALYZER 512
attributes category 230
attributes of the SM2 output file 202
 freely selected 203
 preset 202
authorization table 26
authorizations for nonprivileged users 117
automatic analysis 516
average value of monitoring period 405
axes
 annotation 407
 division 407

B

background monitoring 23
background storage 662
bar chart
 layout 406
 structure 406
basic load 349
basic resources 348
batch processing (criteria) 42
BCAM Application Name 64
BCAM CONNECTION report 224
BCAM MEMORY report 227
BCAM-CONNECTION
 exclude connection sets 121
BCAM-CONNECTION (monitoring program) 52
Big Pages 662
bottleneck analysis 44, 516
 (autom.) starting (SM2R1) 485
boursing mechanism 665
BREAK function 165
BREAK statement 170
BUCKET parameters 53
BUSY 183

C

C interfaces 525
 example 544
 return code evaluation 540
caching 662
calculate arithmetic mean 404
CALL-ADMINISTRATION-PART statement 99
CALL-EVALUATION-PART statement 99
capacity planning 512
categories 424
category 230

- CATEGORY QUEUE report 230
- CATEGORY report 228
- CATEGORY WSET report 233
- CHANGE-MEASUREMENT-PROGRAM statement 100
- channel
 - define for monitoring 132
 - output 236
- channel load 57
- CHANNEL report 235
 - sort criterion 183
- channel utilization 355
- CHANNEL-IO (monitoring program) 57
- clock resetting
 - interruption-free 374, 394
 - uninterruptible 31
- CLOSE-LOG-FILE statement 101
- CMS (monitoring program) 57
- CMS report 238
- command return codes 199
- compression function of SM2U1 384
- computation methods 404
- computation of statistical data 405
- CONF record 502
- configuration 518
- configuration change 39
- connection group 54, 66
- connection set 52, 66, 91
 - define for monitoring 91
 - excluded, define 121, 122
- consoles 252
- continuation character 170
- control output 164
- controlled mode 164, 177
- COSMOS (monitoring program) 58
 - define monitoring program definition 133
 - modify monitoring program definition 103
 - prepare 102
- CPU pool 76
- CPU pools 321
- CPU report 241
- CPU states 664
- CPU utilization 354
- CREATE statement 412
- CREATE-JOBCLASS-SET statement (SM2R1) 414
- CREATE-JOBNAME-SET statement (SM2R1) 415
- CREATE-TSN-SET statement (SM2R1) 416
- CREATE-USERID-SET statement (SM2R1) 417
- D**
- DAB (monitoring program) 58
- DAB CACHE report 246
- DAB cache, select 185
- DAB report 243
- data
 - interface 393
 - statistical 410
- DATA record 506
- data records 30
- deactivation 664
- DEFINED PARAMETER screen 334
- degree of detail for task statistics 462
- demand paging method 668
- DEVICE DISK report 249
 - sort criterion 186
- DEVICE OTHER report 252
- DEVICE TAPE report 254
- DEVICE TD report 256
- device utilization 354
- devices
 - for extended system statistics, define 149
 - for service time statistics, define 146
- dilation factor 665
- disk devices
 - define for monitoring 141
- DISK FILE report 257
- DISK-FILE (monitoring program) 59
- DLM (monitoring program) 59
- DLM report 258
- DMS I/O operations 665
- DRV implementation 363
- DSCR record 506
- dwelt time 42, 665
- dynamic IO configuration change 39

E

- efficient DP system usage 45
- END statement 101, 173
 - (SM2R1) 417
 - (SM2U1) 379
- ES (expanded storage) 467
- ETERNUS DX 15, 510
- event-driven monitoring method 353
- events
 - define for monitored data acquisition (COSMOS) 93
 - group events 94
 - remove from monitoring program definition (COSMOS) 123
- examples
 - for SM2U1 utility routine 389
 - for the SM2R1 analysis routine 488
 - SM2 monitoring program 206
- exclude periods (SM2R1) 478
- EXCP 666, 223
- external storage utilization 348

F

- FILE
 - nonprivileged (monitoring program) 77
 - privileged (monitoring program) 60
- file
 - define for monitoring 95
 - remove from monitoring 124
 - to be removed from monitoring, define 125
- file access 59, 60
- file link name (SM2R1) 399
- file migration 60
- file names 174
- FILE reports 260
 - output monitored data 260
- FILE statement 174
- file statistics 95
 - privileged 95, 96
- file transfer 61
- files statistics 96
- forced deactivation 664
- format description (SM2) 169
- function states 221

G

- GLOBAL report 262
- global report 519
- global storage (GS) 60, 666
- group events 94
- group names (report groups) 611
- GS (monitoring program) 60
- GS report 263
- GSVOL (monitoring program) 60
- GSVOL report 264
- guest system 76

H

- hardware duration 666
- hardware service time 666
- HELP statement 175
- hiperfile concept 666
- HIPLEX MSCF® 61
- host
 - communication 61
 - define for SM2 128
- host network
 - C interfaces (char *host_name parameter) 526, 529
- HSMS (monitoring program) 60
- HYPERVISOR-IDLE value 361

I

- I/O control module 663
- I/O operations (SM2) 663
- I/O reconfiguration 411, 420
- identify resource bottlenecks 486
- IDLE
 - state 664
 - values 361
- information on monitoring programs and monitoring cycles (SM2U1) 386
- information screens 213
- INITIATE-COSMOS statement 102
- initiating task 661
- initiation 666
- INPROC time 52
- input during monitoring 165
- input files (SM2R1) 399

input/output 666
insert remarks 178
INSPECTOR 519
installation of SM2 345
interaction 63, 667
interactive processing 41
internal program name 87
INWAIT time 52
ISAM
 nonprivileged (monitoring program) 77
 privileged (monitoring program) 61
ISAM buffer management 61
ISAM FILE report 268
ISAM pools 61, 77
 define for monitoring 97, 98
 exclude from the monitoring process (nonprivileged user) 195
 include an exclude 171
 include in the monitoring process (nonprivileged user) 193
 remove 126
ISAM reports 265
 output monitored data 265
ISAM statistics 97, 98
 privileged 97, 98
ISO format 392
ISP syntax 169

J
job 42

L
language options 509
Level=High (task statistics) 471
Level=Low (task statistics) 467
logical level 666
long-term file 512

M
main memory 667
 utilization 351
MAX MPL 230
maximum 405
mean values 31
MEASUREMENT STATUS screen 156, 335
memory classes 674
MEMORY report 270
MERGE-FILES statements (SM2U1) 379
messages 346
MIN MPL 230
minimize arm movements 674
minimum 405
MISSED RECORDS 204
MODIFY 117
MODIFY-ADMINISTRATOR-ADMISSION
 statement 102
MODIFY-COSMOS-PARAMETER
 statement 103
MODIFY-MEASUREMENT-PERIODS
 statement 112
MODIFY-RESPONSETIME-PARAMETER
 statement 114
MODIFY-USER-ADMISSION statement 117
module
 information 200
 loading information 198
modules (SM2) 351
monitor files (nonprivileged user) 174
monitored data reports 209
monitored objects 446
 and associated users 159
 output 192
monitored variables 404
 entry 408
monitoring
 terminate (nonprivileged user) 173
monitoring cycle 29
 modify 112
 relationship with output cycle 165
monitoring data output 30

- monitoring environment 401
 - monitoring information 401
 - output of the system name 401
 - outputting 401
 - static system information 401
 - monitoring interval 112
 - monitoring method 27, 353
 - based on monitoring cycle 27
 - based on sampling cycle 27
 - event-driven 27, 350, 353
 - sample-driven 353
 - monitoring parameters for STORAGE-SYSTEM,
 - define 147
 - monitoring program run
 - start 160
 - terminate 161
 - monitoring programs 49, 100
 - stop and restart 100
 - monitoring screen 520
 - monitoring status, output 195
 - monitoring task, activate 26
 - monitoring, rule-based 521
 - MSCF (monitoring program) 61
 - MSCF report 273
 - MTFILE control file 391
 - multi-step transaction 41
 - multiprocessor system 238
 - multiprogramming level 230
- N**
- networks 36
 - prerequisites for using SM2 in 36
 - restrictions 38
 - screen output 37
 - selecting hosts with SM2 36
 - standard SM2 output file 37
 - users 37
 - non accumulated 408
 - notational conventions 169
 - NSM (monitoring program) 61
 - NSM report 275
- O**
- objects and monitoring parameters
 - output active 154
 - output defined 155
 - offline seek 663
 - online analysis 99
 - online help system 507
 - online monitoring 519
 - online monitoring cycle 29
 - online output 49
 - open mode 205
 - OPEN-LOG-FILE statement 120
 - OPENFT (monitoring program) 61
 - openFT instance 61, 277, 576
 - OPENFT report 277
 - openFT-Instanz 576
 - openSM2 (Open Systems) 15
 - operand
 - keyword 169
 - positional 169
 - operation 42
 - OUTPROC time 52
 - output cycle 163, 165
 - output file
 - managing 373
 - preparing 373
 - output file (SM2R1) 391
 - printing 400
 - output mode
 - automatic 163
 - define 176
 - OUTPUT statement 176
 - output title (SM2R1) 484
 - overall performance 667
 - overhead 667

P

- page fault [668](#)
 - real [668](#)
- page reclaims [668](#)
- paging [669](#)
 - activities, assign [661](#)
 - I/O operations [669](#)
 - memory [669](#)
 - rate [669](#)
- PAM output file [204](#)
- Parallel Access Volume [670](#)
- path info (ICVV) [669](#)
- PAV [670](#)
- PCS report [280](#)
- PCS summary report [456](#)
- PDT [670](#)
- percentage frequency [31](#)
- performance analysis
 - automatic (SM2R1) [394](#)
- performance expectations
 - system-oriented [44](#)
 - user-oriented [44](#)
- performance of DP system [42](#)
- PERIODIC TASK report [282](#)
 - sort criterion [187](#)
- PERIODIC-TASK (monitoring program) [62](#)
- PFA (monitoring program) [62](#)
- PFA CACHE report [284](#)
- PFA reports [284](#)
- PFMON (macro) [604](#)
 - layout of the output area [605](#)
 - return code and error indicators [606](#)
- PGE task [670](#)
- plausibility checks (SM2U1) [388](#)
- POSIX (monitoring program) [62](#)
- POSIX report [286](#)
- PPC [669](#)
- prepare SM2U1 session [375](#)
- primary administrator [24](#), [102](#)
- print report groups (SM2R1) [426](#)
- PRINT statements [413](#)
- print statistics on task queues (SM2R1) [424](#)
- print summary analysis (SM2R1) [449](#)
- print system configuration (SM2R1) [420](#)

- PRINT-CONFIGURATION statement (SM2R1) [420](#)
- PRINT-HSMS-STATISTICS statement (SM2R1) [422](#)
- PRINT-QUEUE-TRANSITION statement (SM2R1) [424](#)
- PRINT-REPORTS statement (SM2R1) [426](#)
- PRINT-SUMMARY statement (SM2R1) [449](#)
- PRINT-TASK-STATISTICS statement (SM2R1) [461](#)
- PRIOR task scheduler [230](#)
- private disk selection [184](#)
- privileged user [24](#)
- productive CPU RST [673](#)
- productive I/O RST [674](#)
- productive performance [667](#)
- program counter statistics [196](#), [351](#)
- program execution (SM2R1) [400](#)
- PUBSET (monitoring program) [63](#)
- PUBSET report [288](#)
- pubset selection [150](#), [184](#), [189](#)

Q

- querying
 - system utilization [604](#)
- queues [672](#)

R

- REACT time [52](#)
- Readme file [17](#)
- real-time monitoring [23](#)
- record
 - selecting (SM2U1) [373](#), [381](#)
 - structure [499](#)
- records (SM2R1 transfer file) [499](#)
- reducing error sources
 - measures for [355](#)
- reduction of monitored items for output
 - selection of subreports [426](#)
- reduction of the number of monitored objects for output [446](#)
 - selection of individual reports [426](#)
- REMARK statement [178](#)

- REMOVE-BCAM-CONNECTION-SET
 - statement 121
- REMOVE-CONNECTION-SET statement 122
- REMOVE-COSMOS-EVENT statement 123
- REMOVE-FILE statement 124
- REMOVE-ISAM-FILE statement 125
- REMOVE-ISAM-POOL statement 126
- REMOVE-OPENFT-INSTANCE statement 127
- report group names 611
- report groups 446, 611
- report identifiers 181
- report names 614
- REPORT statement 179
- reports 165, 209
 - select 179
- reports (SM2R1) 519
- representation of averages 409
- request queue 239
- request user help information 175
- requesting bar chart output 409
- residual class (task analysis) 464
- resource bottleneck 486
- resource service time 667
- resource utilization
 - CPU 349
 - external storage 348
- response time 42, 63, 115, 144, 670
- response time statistics 66, 114
 - bucket-specific 114, 143
 - category-specific 114, 143
 - define settings 143
 - examples 67
 - modify settings 114
- RESPONSETIME (monitoring program) 63
- RESPONSETIME report 290
- RESTART statement 183
- result lists (SM2R1) 392
- RSC IO 62, 73, 670
- RST (Resource Service Time) 667
- S**
- SAM output file 205
- sample procedures 364
- sample-driven monitoring method 353
- samples
 - number of 355
- sampling cycle 29
- SAMPLING-DEVICE (monitoring program) 68
- screen output 209
- scroll commands 164
- SDV fast release mode 663
- SE Manager 508
- secondary administrator 24, 102
- select report 179
- select timeslots (SM2R1) 482
- SELECT-CHANNEL-PARAMETER
 - statement 183
- SELECT-CMS-PUBSET statement 184
- SELECT-DAB-CACHE statement 185
- SELECT-DEVICE-DISK-PARAMETER
 - statement 186
- SELECT-HOSTS statement 128
- SELECT-MEASUREMENT-GROUPS statement (SM2U1) 381
- SELECT-PERIODIC-TASK-PARAMETER
 - statement 187
- SELECT-SYMMETRIX-PUBSET statement 189
- SELECT-UTM-APPLICATION statement 191
- SELECTED HOSTS screen 339
- SENSE-SDV 663
- SEPARATE-FILES statement (SM2U1) 383
- service time 667
 - definition (DCS) 671
 - statistics 69
- SERVICETIME (monitoring program) 69
- SESAM SQL report 294
- SESAM-SQL (monitoring program) 71, 74
- session management 509
- session timeout 509
- SET names
 - declaration for job classes (SM2R1) 414
 - declaration for job names (SM2R1) 415
 - declaration for TSNs (SM2R1) 416
 - declaration for user IDs (SM2R1) 417

- SET-BCAM-CONNECTION-PARAMETER
 - statement 130
- SET-CHANNEL-IO-PARAMETER
 - statement 132
- SET-COMPRESSION statement (SM2U1) 384
- SET-COSMOS-PARAMETER statement 133
- SET-DISK-FILE-PARAMETER statement 141
- SET-EVALUATION-PERIOD statement (SM2R1) 473
- SET-EXCEPTION-PERIOD statement (SM2R1) 478
- SET-PERIODIC-TASK-PARAMETER
 - statement 141
- SET-REPORT-FOCUS statement (SM2R1) 482
- SET-RESPONSE-TIME-PARAMETER
 - statement 143
- SET-SAMPLING-DEVICE-PARAMETER
 - statement 145
- SET-SERVICETIME-PARAMETER
 - statement 146
- SET-STORAGE-SYSTEM-PARAMETERS
 - statement 147
- SET-SYSTEM-PARAMETER statement 149
- SET-TASK-PARAMETER statement 151
- SET-TITLE statement (SM2R1) 484
- SET/MODIFY statements 412
- SF pubsets, monitoring data 63
- SHARED PVS report 295
- short-term monitoring
 - start 364
 - terminate and resume 367
- SHOW-ACTIVE-PARAMETER statement 154
- SHOW-DEFINED-PARAMETER statement 155
- SHOW-INFORMATION statement (SM2U1) 386
- SHOW-MEASUREMENT-STATUS
 - statement 156
- SHOW-SELECTED-HOSTS statement 156
- SHOW-SM2-STATUS statement 157
- SHOW-USER-MEASURED-OBJECTS
 - statement 159, 192
- SM2
 - start 79
 - terminate 80
 - use in a computer network 36
- SM2 administrator 83
- SM2 counters 31
- SM2 data
 - accuracy 352
- SM2 interfaces 525
- SM2 monitoring routine (overview) 21
- SM2 output file 23, 101, 202, 373
 - analyze 371
 - attributes 202
 - close 101, 205
 - manage 204
 - open 120, 205
 - specify 197
 - user-specific 196
 - write task 350
 - write to 202
- SM2 output files
 - separating (SM2U1) 383
- SM2 pseudo-category SUM 424
- SM2 reports 164, 209, 210
- SM2 screen output 209
 - representation of values 214
- SM2-PA 522
- SM2-PA (analysis program) 196
- SM2GDAT (macro) 528
 - overview of data buffers and structures 536
 - structure "SM2GDAT_get_data_md1" 558
 - structure of the ACF data buffer 559
 - structure of the BASIC data buffer 560
 - structure of the BCAM-CONNECTION data buffer 560
 - structure of the buffer header 558
 - structure of the CATEGORY data buffer 562
 - structure of the CHANNEL-IO data buffer 563
 - structure of the CMS data buffer 564
 - structure of the DAB data buffer 565
 - structure of the DISK-FILE data buffer 567
 - structure of the DLM data buffer 568
 - structure of the FILE data buffer 569
 - structure of the global header 558
 - structure of the GS data buffer 570
 - structure of the GSVOL data buffer 571
 - structure of the HSMS data buffer 571

- SM2GDAT (macro)
 - structure of the ISAM data buffer 572
 - structure of the ISAM-FILE data buffer 573
 - structure of the MEMORY data buffer 574
 - structure of the MSCF data buffer 575
 - structure of the NSM data buffer 576
 - structure of the OPENFT data buffer 576
 - structure of the output area 533
 - structure of the PERIODIC-TASK data buffer 577
 - structure of the PFA data buffer 578
 - structure of the POSIX data buffer 579
 - structure of the PUBSET data buffer 580
 - structure of the RESPONSETIME data buffer 581
 - structure of the SCHANNEL data buffer 584
 - structure of the SDEVICE data buffer 585
 - structure of the SESAM-SQL data buffer 586
 - structure of the STORAGE-SYSTEM data buffer 587
 - structure of the SVC data buffer 592
 - structure of the SYSTEM data buffer 592
 - structure of the TCP-IP data buffer 594
 - structure of the TIME-IO data buffer 595
 - structure of the TLM data buffer 596
 - structure of the UDS-SQL data buffer 597
 - structure of the USERFILE data buffer 598
 - structure of the USERISAM data buffer 599
 - structure of the UTM data buffer 600
 - structure of the VM data buffer 601
 - structure of the VM-CPU-POOL data buffer 602
 - structure of the VM-GROUP data buffer 603
 - structures 558
- SM2GMS (macro) 526
 - structures 557
- SM2R1 391
 - analysis routine 391
 - starting 397
 - table of contents 395
 - terminating 397
 - transfer file 393, 499
- SM2R1 analysis, prepare 369
- SM2U1 373
 - output files 388
 - starting and terminating 377
 - utility routine 373
- snapshot report 519
- software duration 672
- software service time
 - non-DCS 672
- standard categories 230
- standard deviation 405
- start report output 165, 192
 - for selected reports 183
- start SM2 79
- START statement 192
- START-AUTOMATIC-ANALYSIS statement (SM2R1) 485
- START-ISAM-STATISTICS statement 193
- START-MEASUREMENT-PROGRAM statement 160
- START-SM2 (BS2000 command) 79
- START-SM2R1 (BS2000 command) 397
- START-SM2U1 (BS2000 command) 377
- START-TASK-MEASUREMENT (BS2000 command) 196
- statement syntax (SM2R1)
 - structure 412
- statements 83
 - for SM2 administrators 83
 - length 170
 - SM2R1 414
 - SM2U1 378
 - terminating (SM2R1) 417
- static system information 401
- statistical data 405, 410
- status of SM2 system tasks, output 157
- STATUS statement 195
- STATUS TABLE 340
- STOP time 221
- STOP-ISAM-STATISTICS statement 195
- STOP-MEASUREMENT-PROGRAM statement 161
- STOP-TASK-MEASUREMENT(BS2000 command) 196
- STORAGE-SYSTEM (monitoring program) 71

- structure of
 - statistics 404
 - time series 404
 - subcharts 409
 - SVC (monitoring program) 72
 - SVC report 297
 - SVC statistics 196
 - Symmetrix 15, 510
 - SYMMETRIX DEVICE report 304
 - pubsets 189
 - SYMMETRIX PUBSET report 302
 - SYMMETRIX report 298
 - Symmetrix system, monitored data 71
 - Symmetrix system, select 190
 - SYST record 501
 - SYSTEM (monitoring program) 72
 - system configuration 518
 - system counters 31
 - system list 519
 - system mode, change to 170
 - system monitoring with openSM2 41
 - system privilege SWMONADM 24
 - system report 519
 - system resource utilization (SM2) 348
- T**
- TASK (monitoring program) 73
 - TASK LOCK MANAGER (monitoring program) 74
 - task monitoring
 - activating 196
 - TASK MONITORING CYCLE 465
 - task queues 672
 - retention period in 424
 - TASK START TIME 465
 - task statistics 461
 - outputting (SM2R1) 461
 - TASK STOP TIME 465
 - TASK user-specific (monitoring program) 78
 - tasks
 - and devices for task statistics, define 151
 - deactivated 672
 - monitored data acquisition, define 141
 - TCP-IP (monitoring program) 74
 - TCP/IP report 306
 - terminate SM2 80
 - terminate SM2 run 101
 - terminate SM2U1 379
 - terminatate
 - monitoring (nonprivileged user) 173
 - text file 391
 - think time 63, 115, 144
 - throughput rate 42
 - TIM2 record 499
 - time equivalent for the housekeeping performance 673
 - time equivalent for the productive performance 673
 - time series 513
 - time series report 519
 - timesharing mode 41
 - TLM (monitoring program) 74
 - TLM report 308
 - transaction 41, 674
 - mode 41
 - rate 42
 - time 41, 63, 115, 144
 - trend analysis 512
 - trend monitoring 43

U

UDS SQL report 309
UPG 669
user
 nonprivileged 24
 SM2 24
USER MEASURED OBJECTS screen 343
user monitoring programs 77
user task statistics 351
users
 (SM2R1) 391
 (SM2U1) 374
users and monitored objects 159
UTC time 31, 374, 394
UTM (monitoring program) 75
UTM APPLICATION report 313
UTM report 310

V

variables for which no values can be
 determined 409
variables reports 411, 623
variables, relevant 354
virtual address space
 subdivision 674
VM (monitoring program) 76
VM CPU POOL report 321
VM group 76
VM2000 GROUP report 323
VM2000 operation 356
VM2000 report 318
VM2000 system 76, 318
VOLUME load 363
volume sets, monitoring data 63
VOLUNTARY ACTIVE WAIT STATE 465
VOLUNTARY INACTIVE WAIT STATE 465

W

wait time 63
 in the BCAM pool 116, 145
 non-voluntary 665
 voluntary 665
WEIGHT 231
wildcard syntax 87
Windows PC 512
working set 669

