

Transaction Processing Facility



System Performance and Measurement Reference

Version 4 Release 1

Transaction Processing Facility



System Performance and Measurement Reference

Version 4 Release 1

Note!

Before using this information and the product it supports, be sure to read the general information under "Notices" on page xi.

Thirteenth Edition (June 2002)

This is a major revision of, and obsoletes, SH31-0170-11 and all associated technical newsletters.

This edition applies to Version 4 Release 1 Modification Level 0 of IBM Transaction Processing Facility, program number 5748-T14, and to all subsequent releases and modifications until otherwise indicated in new editions or technical newsletters. Make sure you are using the correct edition for the level of the product.

IBM welcomes your comments. Address your comments to:

IBM Corporation
TPF Systems Information Development
Mail Station P923
2455 South Road
Poughkeepsie, NY 12601-5400
USA

When you send information to IBM, you grant IBM a nonexclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

© **Copyright International Business Machines Corporation 1994, 2002. All rights reserved.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Figures	vii
Tables	ix
Notices	xi
Trademarks	xi
About This Book	xiii
Who Should Read This Book	xiii
How This Book is Organized	xiii
Conventions Used in the TPF Library	xiii
Related Information	xiv
IBM Transaction Processing Facility (TPF) 4.1 Books.	xiv
Miscellaneous IBM Books	xv
How to Send Your Comments	xv
System Measurement and Performance Introduction	1
Functions	1
TPF System Performance Overview	1
TPF Message Processing	2
Message Processing	2
TPF Communications Networks	2
Data Collection Techniques	4
Collectors	4
Sampling Mode	6
Continuous Mode	6
Data Reduction Analysis	6
Installing Data Collection/Reduction Programs	7
Using Data Collection	7
Using Data Reduction	8
General Description of Programs	8
Continuous Data Collection (CDC)	9
TPF System Performance Analysis	11
System Data Analysis	11
Environment Summary Reports	11
Pushbutton Application Summary Report	14
System Summary Report	15
System Pools Summary Report	28
TPF Logical Record Cache Summary Report	28
Coupling Facility Data Collection Reports	29
TPF Internet Mail Server Summary Report	33
TCP/IP Weighted Input Messages by Application Report	34
ECB Frame Usage Summary Report	35
ECB Heap Area Usage Summary Report	35
MPIF Configuration Report	36
Interprocessor Communication MPIF Summary Report	37
MPIF Path Activity Report	37
Frequency Distribution Reports	38
Plot Reports	39
Data Reduction Limitations Report	40
File Data Analysis	41
File Summary Reports	41

Channel Path Reports	45
Sequential File Access Report	46
Record ID Reports	46
File Accesses per Cylinder Report	47
File Comparison Plot Report	47
Storage Control Reports	48
3990 Storage Control Report: Totals	52
3990 Storage Control Report: Subtotals for Cache Subsystem nnnn	53
3990 Storage Control Report for SDA mmmm, Cache Subsystem nnnn	56
3990 Storage Control Cache Summary Report	58
3990 Storage Control Device Summary Report for Cache Subsystem	59
DP Record Anomaly Log	60
Program Data Analysis	61
Program Reports	61
Program Names and Enters Report	62
Program Details Report	62
Cumulative Program Enters Report	63
Cumulative Program On-File Enters Report	64
Package Details Report	64
Cumulative Package Enters Report	65
Cumulative Package On-File Enters Report	66
Program Legend Report	66
Weighted-Message Rate Report	67
Message Data Analysis	67
Message Summary Report	67
Action Code Summary Reports	69
Application Summary and Detail Reports	70
City Summary Report	71
Terminal Activity Report	72
Logical Unit Activity Report	74
Logical Unit Message Stream Report	75
Message Stream Report	76
Distributed Data Management (DDM) Message Stream Report	77
TCP/IP Message Summary Report	78
TPF MQSeries Data Collection Reports	79
Details of Collection/Reduction	83
Details of System Collection	83
Working Storage Usage	83
Time Stamp	83
Details of System Reduction	83
Total Working Storage Assigned	83
Milliseconds Per Weighted Message	83
Details of File Collection	83
File Intercepts	83
File Queues	83
Details of File Reduction	84
Random File Access Summary - Total Accesses Per Second	84
Random File Access Summary - Device Queues Mean	84
Random File Access Summary - Total System	84
Sequential File Access Summary - YYY Module	84
File Accesses Per Record ID Report	84
Details of Program Reduction	84
The Identification of Individual Programs	84
The System of Graphic Cues Used in the Program Reports	85
The Categorization of Programs	85

Details of Message Collection	85
Time Stamp - High-Speed Input Message	85
Time Stamp - High Speed Output Message	85
Details of Message Reduction	85
All Reports	85
Action Code Summary, Application Summary - Mean Existence Time	86
Action Code Summary - Mean Output Length	86
Action Code Summary, Application Detail - Unmatched Output	86
City Summary	86
Terminal Activity	86
Plot/Dist Option - Message	86
Weighted Message Rate	86
Message Stream - Output Messages	87
Message Summary - Max/Min	87
Use of DA Records	87
Limitations of Message Stream Reports	87
Collection Details Of Driver	87
Message Counters	87
One-Time Driver Records	88
Additional Data Collected	89
Installation Details	91
Compilation Phase (Using PL/I Optimizing Compiler)	91
Preprocessor Statements	91
Variables	93
Linkedit Phase	93
Execution Phase	94
Input	94
Files	94
Output	95
Error Messages	95
Index	97

Figures

1.	Data Collection/Reduction Overview	9
2.	Environment Summary Report	12
3.	Input Messages by Application Report	14
4.	Pushbutton Application Summary Report	14
5.	System Summary Report	16
6.	System Pools Summary Report	28
7.	TPF Logical Record Cache Summary Report	29
8.	Coupling Facility Usage Summary Report	30
9.	Coupling Facility Structure Summary Report	31
10.	Coupling Facility Locking Summary Report	32
11.	Coupling Facility Caching Summary Report	32
12.	TPF Internet Mail Server Summary Report	33
13.	TCP/IP Weighted Input Messages by Application Report	34
14.	ECB Frame Usage Summary Report	35
15.	Heap Area Usage Report	36
16.	MPIF Configuration Report	37
17.	Interprocessor Communication MPIF Report	37
18.	MPIF Path Activity Report	38
19.	Frequency Distribution Report	38
20.	Plot Report	40
21.	Data Reduction Limitations Report	41
22.	File Access Summary Report	42
23.	VFA Summary Report	43
24.	Random File Access Summary Report	44
25.	Channel Path ID Status Exception Report	45
26.	Path Activity per Channel Path ID Report	45
27.	SDA Activity Exception Report	46
28.	Sequential File Access Report	46
29.	Record ID Report	47
30.	Cylinder Analysis Report	47
31.	File Comparison Plot	48
32.	3880 Storage Control Report: Totals	49
33.	3880 Storage Control Report: Subtotals for SDID nn	50
34.	3880 Storage Control Report for SDA mmmm, SDID nn	51
35.	3990 Storage Control Report: Totals (Part 1)	52
36.	3990 Storage Control Report: Totals (Part 2)	53
37.	3990 Storage Control Report: Subtotals for Cache Subsystem nnnn (Part 1)	54
38.	3990 Storage Control Report: Subtotals for Cache Subsystem mmmm (Part 2)	55
39.	3990 Storage Control Report: Subtotals for Cache Subsystem mmmm (Part 3)	56
40.	3990 Storage Control Report for SDA nnnn, Cache Subsystem mmmm (Part 1)	57
41.	3990 Storage Control Report for SDA nnnn, Cache Subsystem mmmm (Part 2)	58
42.	3990 Storage Control Cache Summary Report	59
43.	3990 Storage Control Device Summary Report for Cache Subsystem mmmm	60
44.	DP Record Anomaly Log	61
45.	Program Names and Enters Report	62
46.	Program Details Report	63
47.	Cumulative Program Enters Report	64
48.	Cumulative Program On-File Enters Report	64
49.	Package Details Report	65
50.	Cumulative Package Enters Report	65
51.	Cumulative Package On-File Enters Report	66
52.	Program Legend Report	67
53.	Weighted-Message Rate Report	67

54.	Message Summary Report	69
55.	Action Code Summary Report	70
56.	Application Detail Report	71
57.	City Summary Report	72
58.	Terminal Activity Report	73
59.	Logical Unit Activity Report	75
60.	Message Stream Report by LU	76
61.	Message Stream Report by Agent	77
62.	Distributed Data Management (DDM) Message Stream Report	78
63.	TCP/IP Message Summary Report	79
64.	TPF MQSeries Message Channel Report	80
65.	TPF MQSeries Message Queue Report	81

Tables

1. Collectors and Their Information	5
2. ECB Fields Containing Additional Data	89
3. Preprocessor statements	91

Notices

References in this book to IBM products, programs, or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service in this book is not intended to state or imply that only IBM's product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any of IBM's intellectual property rights may be used instead of the IBM product, program, or service. Evaluation and verification of operation in conjunction with other products, except those expressly designated by IBM, is the user's responsibility.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
USA

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
Department 830A
Mail Drop P131
2455 South Road
Poughkeepsie, NY 12601-5400
USA

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

Any pointers in this book to non-IBM Web sites are provided for convenience only and do not in any way serve as an endorsement. IBM accepts no responsibility for the content or use of non-IBM Web sites specifically mentioned in this book or accessed through an IBM Web site that is mentioned in this book.

Trademarks

The following terms are trademarks of the IBM Corporation in the United States or other countries or both:

IBM
MQSeries
PR/SM
VisualAge
3090.

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

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

About This Book

This book describes the performance and measurement package for the TPF system: online data collection and offline data reduction. It provides information on analyzing the reports to improve system performance and provides diagrams and sample reports. It also provides information on installing the collection and reduction programs.

In this book, abbreviations are often used instead of spelled-out terms. Every term is spelled out at first mention followed by the all-caps abbreviation enclosed in parentheses; for example, Systems Network Architecture (SNA). Abbreviations are defined again at various intervals throughout the book. In addition, the majority of abbreviations and their definitions are listed in the master glossary in the *TPF Library Guide*.

Who Should Read This Book

This information is intended for use by system programmers who are analyzing and tuning the performance of their TPF system.

How This Book is Organized

There are five major sections in this book.

1. An introductory section describes general aspects of TPF collection and reduction.
2. The four kinds of data analysis reports (system, file, program, and message) provided by data reduction are discussed next.
3. Another section presents details of collection and reduction.
4. A brief section discusses additional data collected (entry control block (ECB) resources).
5. A fifth section describes various details required for installation of the collection and reduction programs.

Online operation of data collection is documented in *TPF Operations*. That chapter contains the options required for obtaining the reports described in this book.

Conventions Used in the TPF Library

The TPF library uses the following conventions:

Conventions	Examples of Usage
<i>italic</i>	Used for important words and phrases. For example: <i>A database</i> is a collection of data. Used to represent variable information. For example: Enter ZFRST STATUS MODULE <i>mod</i> , where <i>mod</i> is the module for which you want status.
bold	Used to represent text that you type. For example: Enter ZNALS HELP to obtain help information for the ZNALS command. Used to represent variable information in C language. For example: level

Conventions	Examples of Usage
monospaced	Used for messages and information that displays on a screen. For example: <pre>PROCESSING COMPLETED</pre> Used for C language functions. For example: <pre>maskc</pre> Used for examples. For example: <pre>maskc(MASKC_ENABLE, MASKC_IO);</pre>
<i>bold italic</i>	Used for emphasis. For example: You <i>must</i> type this command exactly as shown.
<u>bold underscore</u>	Used to indicate the default in a list of options. For example: Keyword=OPTION1 <u>DEFAULT</u>
Vertical bar	Used to separate options in a list. (Also referred to as the OR symbol.) For example: Keyword=Option1 Option2 Note: Sometimes the vertical bar is used as a <i>pipe</i> (which allows you to pass the output of one process as input to another process). The library information will clearly explain whenever the vertical bar is used for this reason.
CAPital LETters	Used to indicate valid abbreviations for keywords. For example: KEYWOrd= <i>option</i>
Scale	Used to indicate the column location of input. The scale begins at column position 1. The plus sign (+) represents increments of 5 and the numerals represent increments of 10 on the scale. The first plus sign (+) represents column position 5; numeral 1 shows column position 10; numeral 2 shows column position 20 and so on. The following example shows the required text and column position for the image clear card. <pre> ...+....1....+....2....+....3....+....4....+....5....+....6....+....7...</pre> <pre>LOADER IMAGE CLEAR</pre> Notes: <ol style="list-style-type: none"> 1. The word LOADER must begin in column 1. 2. The word IMAGE must begin in column 10. 3. The word CLEAR must begin in column 16.

Related Information

A list of related information follows. For information on how to order or access any of this information call your IBM representative.

IBM Transaction Processing Facility (TPF) 4.1 Books

- *TPF Application Requester User's Guide*, SH31-0133
- *TPF Concepts and Structures*, GH31-0139
- *TPF Operations*, SH31-0162
- *TPF System Generation*, SH31-0171
- *TPF System Macros*, SH31-0151
- *TPF System Installation Support Reference*, SH31-0149
- *TPF Transmission Control Protocol/Internet Protocol*, SH31-0120
- *TPF Program Directory*, GI11-0418.

Miscellaneous IBM Books

- *Linkage Editor Manual*, SC26-4510
- *OS PL/I Optimizing Compiler Programmer's Guide*, SC33-0006
- *PL/I Optimizing Compiler: Language Reference Manual*, SC26-4308.

How to Send Your Comments

Your feedback is important in helping to provide the most accurate and highest quality information. If you have any comments about this book or any other TPF information, use one of the methods that follow. Make sure you include the title and number of the book, the version of your product and, if applicable, the specific location of the text you are commenting on (for example, a page number or table number).

When you send comments to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate without incurring any obligation to you.

- If you prefer to send your comments electronically, do either of the following:
 - Go to <http://www.ibm.com/tpf/pubs/tpfpubs.htm>.
There you will find a link to a feedback page where you can enter and submit comments.
 - Send your comments by e-mail to tpfid@us.ibm.com
- If you prefer to send your comments by mail, address your comments to:
IBM Corporation
TPF Systems Information Development
Mail Station P923
2455 South Road
Poughkeepsie, NY 12601-5400
USA
- If you prefer to send your comments by FAX, use this number:
 - United States and Canada: 1 + 845 + 432 + 9788
 - Other countries: (international code) + 845 + 432 +9788

System Measurement and Performance Introduction

The system performance and measurement package has two components:

1. The online data collection portion
2. The offline data reduction portion (with segment names beginning with JR).

The online collection portion collects and records real-time TPF system performance information. The offline reduction portion reads, edits, and reports the recorded performance information. The data analysis is finally printed in reports. A clear understanding of these reports can lead to optimal system tuning and substantial performance increases.

This package monitors system performance by collecting predetermined indicators, at fixed time intervals, through interaction with the control program.

Performance information is collected on overall system and individual program activities. Specialized collectors retain information about the following:

- Message traffic
- Program activity
- File storage activity.

A system collector records a cross section of total system activity.

The package is designed to report on the total processor environment. The ZMEAS command can start collection on an individual processor complex or across an entire loosely coupled processor complex. The collectors can gather all activity for all subsystems in a multiple database function (MDBF) environment. The reduction programs provide for generation of systemwide summaries and detail reports for a particular subsystem.

Functions

The system performance package is designed to provide operational data on all significant activities involved in processing input messages from network resources. These resources may be terminals, logical units, applications, and others. By analyzing the reports produced, you can determine how efficiently the installation is running, discover where the bottlenecks are, and find clues about changes in system allocation (main storage, files, lines, or terminals) that could improve system performance.

This package can be used during installation and the post cutover phase to tune the system to peak efficiency. It provides a means for daily monitoring of system performance. The statistics generated allow long-term trends to be observed from the daily runs, in this way providing the basis for predicting growth in system load.

TPF System Performance Overview

An overview of the basic flow for processing a TPF message is presented before the discussion of the techniques of data collection and analysis. See *TPF Concepts and Structures* for a more detailed discussion of the Transaction Processing Facility communication network and for a day in the life of the ECB.

Do not get discouraged by the volume of data made available by these programs. Start with the summary reports and then work through the various levels of detail reports to understand the performance data in its proper perspective.

TPF Message Processing

A typical TPF system consists of a central system and the local domain network controlled by the central system. The TPF system may participate with a cross-domain network through SDLC cross-domain links if the central system has the Advanced Communication Facility (ACF). This section does not intend to describe the aspects of networks. Those details are found in *TPF Concepts and Structures*. The discussion here focuses on the major flow in the central system and the message paths in the network.

The bulk of the load presented to the TPF system comes from conversational messages that originate at resources of the network (terminals, network addressable units, and others). These messages, normally transmitted at 2400 BPS or higher speeds, are known as high-speed messages. Teletype or low-speed messages are also counted and equated to multiple high-speed messages to obtain an average or weighted message rate for the total system. All the report statistics calculated on a per-message basis use this weighted message rate as a divisor.

Message Processing

A TPF message enters the processing cycle from a network resource over a communication line after the terminal interchange is polled. If an NCP is involved, the messages or path information units are queued in the NCP and sent to the CPU on a read operation from the CPU. After the input message is assembled, the message is placed on the input list to await processing by the application programs. The CPU loop scans the various lists looking for work. Top priority is given to messages in process waiting on the cross list. With the cross list empty, the system will normally pick up existing work from the ready list. When the ready list is empty, the system will normally pick up a new message from the input list, assign an entry control block (ECB), that will be associated with the message throughout the processing cycle, and branch to start processing by application programs. When all other lists are empty, the system will pick up the deferred list. These programs will edit the message, retrieve data records from file, log data to tape, obtain additional storage when required from the pools, and release blocks. These are tasks required to develop an output message in response to the input message from the agent. When file or tape operations occur, the application program normally relinquishes control so that I/O operations can be overlapped with CPU processing.

Processing will continue on this message when the I/O operations are complete, the message (ECB) appears on the top of the ready list, and the control program selects the top item and continues processing. When processing is complete, the output message is placed on a queue for the desired network resource. At the same time, the ECB assigned to the message is returned to the system. The message is sent to the network resource when it assumes first position in the queue, and the network resource is available to receive the message.

TPF Communications Networks

Consider a network configuration of a TPF system with ACF support. The network is identified as local-domain and cross-domain networks. Because major network control characteristics of the central TPF system can be observed from the network resource addressing scheme, some aspects of the addressing schemes are described as follows:

- Actual line number, interchange address, and terminal address (LNIATA)s are used for network resources on Airlines Line Control (ALC) communication lines.
- For the SNA network, though various levels of addressing conventions exist in different SNA layers, the data collection/reduction package represents network resources using resource identifiers (RID), node names, and subarea addresses. It is assumed that analysts are knowledgeable of the relationship between node names and their physical resources in the network based on the network system generation reports. Local SNA domain resources reported by the reduction package are logical units, cluster controllers, and SDLC lines. NCP nodes are reported by their node names. For SNA cross-domain resources, the network activities are reported for cross-domain SDLC link and cross-domain logical unit names. In other words, the session activities between local SNA resources and cross-domain applications are not included in the reports of the local SNA domain.

When an NCP is shared by multiple CPUs, regardless of the ownership, data reduction reports the activities of the SDLC or ALC lines seized by the central TPF system.

- Pseudo line, interchange, and terminal addresses.
- A pseudo LNIATA address exists in a TPF system for SDLC 3270 logical units. When an SDLC 3270 logical unit starts a session with a non-SNA application in the TPF system, the previously defined pseudo LNIATA is also used in addition to the SNA address. Regardless of the physical network structure, a single pseudo line number is assigned for all messages (X'FE'). Terminals of the Network Extension Facility (NEF) network are also given pseudo line, interchange, and terminal addresses in addition to their SNA addresses.

TPF Network Extension Facility (NEF) support depends on either:

- The Network Extension Facility (NEF2 PRPQ P85025), or
- Airlines Line Control Interface (ALCI) feature of ACF/NCP.

All subsequent occurrences of NEF in the text of this publication are superseded by the previous information.

For the NEF network, terminals are addressed by Logical End Identifiers (LEID). LEIDs associated with one terminal interchange need not have the same first byte(s). For this network, the low end SNA addresses are limited to interchange devices. The relation between a pseudo LNIATA and the associated terminal depends on your definition.

This type of pseudo LNIATA is not used by data collection or reduction to report line activities. However, it is used for other network analysis reports.

In addition to the addressing schemes of the TPF system, it is helpful to review the paths between local-domain and cross-domain networks. When the central TPF system has ACF capability, the system has two communication paths between two domains. One path is supported by the Message Routing Facility and the other by an SNA cross-domain session. An SNA cross-domain session permits local SNA network resources to communicate with cross-domain applications, or cross-domain resources to communicate with local domain applications, without involving overhead to the CPU that owns the resources. This resource includes application resources.

The objective of the message routing facility is to provide the ability to:

- Forward a message received from a terminal on a TPF high-speed line to an application program running in another SCP or on another TPF processor.

Note: The term system control program (SCP) is loosely used throughout this publication to indicate a control program other than TPF. SCP will be synonymous with MVS when mentioned in the description of a function for which one or both of them can be used.

- Receive a message from an SCP or a TPF system and forward the message to a designated terminal on a TPF high-speed line.
- Transmit data from a TPF application program to an application program executing in another SCP or TPF system.
- Receive data from an SCP or a TPF system and pass the data to a designated user-written TPF application program.
- Receive data from an SCP or a TPF system and forward the data to another SCP or TPF system.

It must be reiterated that the objective of the message routing facility is to provide for the transfer of data from an origin point to a destination point in the system network domain.

A TPF system with ACF permits the message routing facility to use cross-domain SDLC links between two CPUs.

Data Collection Techniques

The basic techniques used in collecting data in the TPF system are:

1. Reading out counters imbedded in TPF code and continuously updated by TPF, such as high-speed message counts
2. Intercepting specific events such as file macros and program enters when the collector programs are active
3. Dynamic sampling of parameters that fluctuate with time, such as I/O device queues and working storage blocks in use.
4. Collecting data from certain storage controllers and certain DASD themselves. This technique is becoming increasingly more important.

In any data collection program involving real-time operations, one factor to consider is the load and interference the collection programs themselves impose on the system being measured. Minimal impact on normal processing operations is essential. Because this impact cannot be zero, every effort must be made to know the extent of the influence data collection has on the parameters to be analyzed.

Collectors

The activity in the message processing cycle is associated with 4 different collectors. These collectors retain the values of system variables involving message processing, program operation, file activity, and the overall system structure profile. These 4 collectors can be turned on singly or in combination. The data gathered by these collectors is listed in Table 1.

Table 1. Collectors and Their Information

Collector	Information Collected
Message	Arrival time (ECB created) Origin, Destination address, or both Text of message (input and output) Length of message (input and output) Type of input message (action code) Output line queue lengths Characters per line (input and output) Count of messages (input and output) Send time (ECB returned to the system) I-stream number (to which output message is routed).
System	Cross list queue length Input list queue length Ready list queue length Deferred list queue length Count of high-speed input messages Count of low-speed input messages CPU loop counters (of elapsed time in various CPU states) Working storage block usage System heap utilization Count of active ECBs Idle time accumulators Shutdown indicators Frame Usage indicators Heap Frame Usage indicators Pause time and counts SWISC creates and enters (including tpf_fork calls) System pools counts Recovery log track buffers allocated Recovery log record location Recovery log records allocated Recovery log tracks in use Recovery log track writes Commit scope buffers per commit scope WLOGC blocked conditions Commits per second Rollbacks per second Commit buffers in use in VFA.
Program	Count of program ENTRCs, ENTNCs, or ENTDCs (by program name or number) Count of program BACKCs (by program name) Program residency (main storage, file) I-stream number.
File	Count of data record accesses (by device, channel) Data record identity Count of program accesses (by device, channel) Count of tape writes (by tape ID) Count of characters written to tape Device queue counts (files and tapes) VFA synchronization data VFA RIAT option byte Caching RIAT option byte Last channel path ID used.

For most collectors, there are two kinds of information collected: *summary data* and *detailed data*. The summary data is collected according to specified intervals while the detailed data is collected by intercepts. File and message collectors have both

interval and intercept data associated with them. The system collector has no intercept data and the program collector has no interval data.

\$DCOLC macro calls imbedded in system segments centralize and assist collector implementation. (See *TPF System Macros* for more information about the \$DCOLC macro.)

Sampling Mode

In sampling mode both summary and detailed data are collected over intervals by collectors, specified in a ZMEAS command. A collector's summary data is gathered at the beginning and end of an interval, while during the body of the interval its intercept data is collected. At the conclusion of one interval another interval begins with a (possibly) new collector. In this mode each collector is run for an interval during every period in a sequential manner. For instance, a collector might run for 5 seconds every minute so that it alone is active in a 5-second interval.

The TPF system does not allow simultaneous activation of multiple collectors because the intercept techniques would significantly interfere with message processing. When details from multiple collectors are required, the collectors must be run in sampling mode.

Continuous Mode

In continuous mode all the collectors are activated sequentially for the gathering of summary data. In addition, a single collector must be specified to gather detailed data at specified intervals. For example, if the file collector was selected, file intercepts are activated at the start time and remain active until data collection is stopped. During this period the system collector and message collector are activated at specified intervals to collect the system counters and the message counts. (Because there is no interval data for the program collector, it does not get activated during these intervals.) If the program collector is selected, the system, the file, and the message collectors are invoked in sequence for summary data while program intercept data is gathered continuously.

Data Reduction Analysis

All data reduction associated with the system performance and measurement package is done under the IBM MVS system. The analysis phase should not be approached hastily. Snap judgements made from too little data usually cause more problems than leaving the system alone. Remember that the TPF system has many carefully designed cutoff levels as protection against overload conditions, and any tampering with the adjustment of these levels must be considered carefully in light of the many interacting factors involved.

An online check of the message traffic volumes the system is handling may be valuable to, for example, the ticket sales and reservations supervisors for immediate decisions, but the health of the TPF system as a programming system can only be determined after a careful diagnosis of the data reduction reports. The reports are not self-explanatory, but once you become acquainted with them, the normal state will be quickly recognized from, for example, the System Summary report. Any abnormality surfacing in this report, such as an increase in milliseconds per message, can be followed up by closer scrutiny of the message mix to determine whether or not the mix has deviated from the norm.

The primary objective of the initial analysis phase for any TPF system is to establish the normal state limits for each of the key factors affecting performance.

Once these factors are set and agreed to be realistic, a daily check on the health of the system becomes routine. However, don't let the performance analysis experts relax at this point; turn them loose on looking for possible performance improvements while they are waiting for emergency calls.

Although the data reduction reports were designed to be used by an analyst familiar with the TPF system, the analyst does not necessarily have to be a statistician. However, for the convenience of those so inclined, frequency distribution reports including means, standard deviations, variances, and so on, of each parameter, are available on an optional basis.

Remember, we are interested in the long-range health of the TPF system. The best way to assure a long life is for periodic check-ups and excellent records on the history of the patient so that seasonal changes will not confuse the diagnostician.

Installing Data Collection/Reduction Programs

Data collection programs are part of the TPF system and no action is necessary to install the collection programs. The data reduction program runs under the IBM MVS system. To install the data reduction package under the IBM MVS environment, option cards or parameters are needed to compile, link edit, and run its various phases (see *TPF Operations*). You can update the JCL and option statements to tailor the package to your installation. See "Installation Details" on page 91 for more information. The shipped package assumes that you will define each city in the proper UAT records.

Using Data Collection

Run data collection every day during the peak hour to provide a history file of key system parameters such as milliseconds per message, file accesses per message, working storage usage per message, enters per message, message rate, and message length. Any changes in system configuration, programs, or storage allocation should be carefully documented and dated so that the performance aspects of the system obtained from data collection can be correlated to these changes. This information plus the history and trend of function activity (for example, passengers boarded) can be used very effectively to predict the need for additional memory, channels, files, lines, terminals, or CPU capacity.

The daily peak hour runs should use all collectors in the sampling mode. To find the peak hour each day will involve a more extensive use of data collection initially. The system collector can be run in sampling mode during each hour of the prime shift to observe the hourly trend. The 60-minute period during which the highest mean high-speed input message rate occurs is defined as the peak hour. In most systems a double peak occurs (mid-morning and mid-afternoon). It may be necessary to track the double-peak. In such cases, alternate the running of sampling mode (all collectors) with continuous mode (system collector), for each peak, to reduce the amount of data collected.

The message format required to start and stop data collection in sampling and continuous mode can be found under system performance in *TPF Operations*.

In a system that is loosely coupled, where all processors share a common set of files, you can run data collection on each processor nonconcurrently or concurrently. If run nonconcurrently, data is collected on each processor one after the other. If run concurrently, it is unnecessary, and even undesirable, for more than

one processor to collect complexwide data from the storage controls. The ZMEAS command has a parameter for specifying which processor is to collect exclusively complexwide data.

Data collection retrieves the high watermark value at the start and end of each data collection episode. The value retrieved at the end of the episode is used as the episodic high watermark.

When data collection is run concurrently and when CFLF is supported, you should use the ZMEAS *x* command, where *x* is the CPU ID on which data is to be collected. The use of this parameter ensures that data collection has collected valid high watermark values. Data collection is activated in the basic subsystem and collects information about all subsystems in a multiple database function (MDBF) environment.

Note: CFLF storage control maintains a high watermark value for locks in use and held but resets the value after a channel program retrieves it. The command ZBUFC STATUS *dddd*, where *dddd* is a device attached to a CFLF controller, should be avoided during data collection because the command retrieves and resets the high watermark value.

Using Data Reduction

The option card formats used to specify which reduction reports are desired can be found under system performance in *TPF Operations*.

The initial analysis of any RTC tape should always start with summary reports only. The summary reports will provide all the key performance data required for history and trend analysis. When investigating a problem area, the more detailed plot and distribution reports or the specialized reports of the file and message reduction programs can be used. The plot reports showing the value of each parameter sample in chronological order can be very effective for cross-correlation of the cause and effect relationship between parameters. The file reports such as cylinder analysis and record ID analysis can be used to investigate overloading conditions on the files. The message reports such as line summary, terminal activity, logical unit activity, and message stream can be used to investigate overloading of the high-speed lines, effectiveness, and system response time.

In a loosely coupled environment, data reductions are performed for each processor. The results are similar to those obtained in a single processor environment. The primary shared resource is the file subsystem.

General Description of Programs

An overview of all the programs involved in data collection and data reduction is presented in Figure 1. Data collection passes information to data reduction through the collection tape, RTC. The factors that affect a TPF system are various kinds of inputs. Selective data is collected by various *collectors* and placed on the tape. The tape is read by the data reduction driver and various reducers provide the reports that are the product of the performance package. Data collection takes place during operation of the online TPF system, while data reduction goes on offline (to TPF system) on an IBM MVS system.




Figure 1. Data Collection/Reduction Overview

Continuous Data Collection (CDC)

Continuous data collection (CDC) is an application that collects real-time TPF 4.1 system performance information. CDC uniquely stores the data in a table in a remote relational database by using the TPF Application Requester (TPFAR) feature. An offline application is then used to interpret and display the data. The name of the table is SYSTPF.TPF_DATA. This information is available in the database on a continual basis, so you can use CDC as a monitoring tool. You can run CDC with a minimum impact on TPF 4.1 system performance. Data collected in CDC is a subset of the data collected by the system performance and measurement package.

You can use the CDCA and CDCB user exits to assist in recording additional data. See *TPF System Installation Support Reference* for more information on the CDCA and CDCB user exits.

You can use the ZCDCO command to manage CDC. See *TPF Operations* for more information on the ZCDCO command.

For more information about using an offline program to display data, go to <http://www.ibm.com/tpf/pages/tpfar.htm>.

TPF System Performance Analysis

Now that you are familiar with the philosophy and mechanics of data collection and reduction, let us focus on how to interpret the reports generated by the data reduction programs. The data collectors and the reports associated with each follow.

All the key parameters are reported on a per second or per message basis. The per second approach provides the basis for analyzing the system load characteristics; that is, the utilization of system hardware components and program segments. The per message approach provides a base for predicting the message rate the system can handle at any utilization level.

Note: When you compare numerical values in the performance reports found in this document, some values for access per second have been influenced by the skip factor. As reported in the Environment Summary Report, the skip factor that is used to generate some reports in this document decreases raw access numbers by a factor of 100. The skip factor can be specified separately for program or system collectors. The skip factor provides a system snapshot without disruptive production of huge amounts of data. Accesses per second are calculated by dividing the estimated number of events under a given skip factor by the total observation time. See the system performance and measurement information in *TPF Operations* for more information about the skip factor.

System Data Analysis

Reports follow that provide the information needed for system performance analysis. The environment and system summary reports appear first. These are followed by the reports associated with each collector.

Environment Summary Reports

The Environment Summary Report contains general information about a collection run. Besides the date, the tape volume serial number, and information identifying the complex, information is displayed that defines the system environment at the time and describes the kind of data collection being run.

```

TPF ENVIRONMENT SUMMARY REPORT                      SYSTEM WIDE
DATE COLLECTED - 27 FEB
DATE REDUCED   - 16 SEP 98
DATA COLLECTION TAPE VOLUME SERIAL NUMBER - AL3930
ENVIRONMENT - PR/SM-SH PARTITION - TPF1 CAP - N WAIT COMP - N
IBM MOD - 370XA/9672 SERIAL NUMBER - 0C034579 TPF CPUID - B
TPF RELEASE AND VERSION - TPF4.1 , PUT LEVEL - PUT07
VIRTUAL MODE - VEQV

OPERATOR INPUT REQUEST - ZMEAS I/SMPF/1023/05
MODE OF COLLECTION - SAMPLING
INTERVAL SYSTEM = 05 SEC MESSAGE = 05 SEC PROGRAM = 05 SEC FILE = 05 SEC
PERIOD = 23 SEC
FILE COLLECTION INTERCEPT SKIP FACTOR = 99
PROGRAM COLLECTION INTERCEPT SKIP FACTOR = 99
START TIME = 14:18:27
END TIME = 14:26:06
ACTIVE UTILITIES          START          END
  DATA COLLECTION        X          X

ACTIVE TRACES             START          END
  MACRO TRACE             X          X
  ENTER/BACK TRACE        X          X
  SYSTEM LOG              X          X
  IO TRACE                X          X
  VEQR MODE LOGGING       X          X

```

Figure 2. Environment Summary Report (Part 1 of 2)

OPTIONS CHOSEN		SUBOPTIONS		
COLLECTOR	OPTION			
SYSTEM				
SNA				
CONTROL	DUMP	10		
MESSAGE				
PROGRAM				
PROGRAM	PACKAGE	A	A*	
PROGRAM	PACKAGE	B	B*	
PROGRAM	PACKAGE	C	C*	
PROGRAM	PACKAGE	D	D*	
PROGRAM	PACKAGE	E	E*	
PROGRAM	PACKAGE	F	F*	
PROGRAM	PACKAGE	G	G*	
PROGRAM	PACKAGE	H	H*	
PROGRAM	PACKAGE	I	I*	
PROGRAM	PACKAGE	J	J*	
PROGRAM	PACKAGE	K	K*	
PROGRAM	PACKAGE	L	L*	
PROGRAM	PACKAGE	M	M*	
PROGRAM	PACKAGE	N	N*	
PROGRAM	PACKAGE	O	O*	
PROGRAM	PACKAGE	P	P*	
PROGRAM	PACKAGE	Q	Q*	
PROGRAM	PACKAGE	R	R*	
PROGRAM	PACKAGE	S	S*	
PROGRAM	PACKAGE	T	T*	
PROGRAM	PACKAGE	U	U*	
PROGRAM	PACKAGE	V	V*	
PROGRAM	PACKAGE	W	W*	
PROGRAM	PACKAGE	X	X*	
PROGRAM	PACKAGE	Y	Y*	
PROGRAM	PACKAGE	Z	Z*	
PROGRAM	PACKAGEREPORTS			
FILE				
FILE				
FILE	ACCESSESPERID	ALL		
FILE	CYLINDERANALYSIS	ALL		
FILE	DISTRIBUTION	ALL		
FILE	PLOT	ALL		
FILE	PATHACTIVITY			
FILE	COMPARISON			
FILE	CACHE			
FILE	CACHE	SSD		
FILE	CACHE	SALL		
FILE	CACHE	CCACHE		
FILE	CACHE	CACHESUM		
FILE	CACHE	CACHEALL		
FILE	PLOT	ALL		
FILE	DIST	ALL		
REDUCE	SS	BSS		
ALIAS	S	BSS	BSS	SUBSYSTE
ALIAS	U	HPN	BSS	SSU
CONFIGURATION SUMMARY				ONE
CPU CONFIGURATION				SYSTEM WIDE
MACH TYPE	SERIAL NR			
370XA/9672	0C034579	B		DATA COLLECTED ON THIS CPU
MDBF CONFIGURATION				
SS NAME	SS NR	SS STATE	SSU NAME	SSU NR
BSS	0	NORM		HPN

Figure 2. Environment Summary Report (Part 2 of 2)

The options listed in the *options chosen* part of the Environment Summary Report come from a typical run. This list should not be considered exhaustive. See *TPF Operations* for more information about the list of options.

The system collector is primarily concerned with the operation of the CPU proper. The items shown in the system summary (see Figure 5) may be considered in groups comprising gross input, system utilization, working storage utilization, and queues or job lists, which make up the priority of processing.

INPUT MESSAGES BY APPLICATION		SYSTEM WIDE		27 FEB	14:18:27
APPLICATION	SSU NAME	MSG / SEC	MESSAGES	PERCENT OF TOTAL	
AAAA	HPN	0.000	0	0.000	
APPA	HPN	0.000	0	0.000	
APPC	HPN	0.000	0	0.000	
AZAZ	HPN	0.000	0	0.000	
BBBB	HPN	0.000	0	0.000	
B320	HPN	0.000	0	0.000	
CBM1	HPN	0.000	0	0.000	
CBM2	HPN	0.000	0	0.000	
CBW1	HPN	0.000	0	0.000	
:					
TSIM	HPN	88.740	40732	99.948	
WEWE	HPN	0.000	0	0.000	
WRAP	HPN	0.000	0	0.000	
WWW	HPN	0.000	0	0.000	
XCXC	HPN	0.000	0	0.000	
XXXX	HPN	0.000	0	0.000	
YDYD	HPN	0.000	0	0.000	
YYYY	HPN	0.000	0	0.000	
Y0Y0	HPN	0.000	0	0.000	
ZZZZ	HPN	0.000	0	0.000	
TOTAL		88.786	40753	100.000	

Figure 3. Input Messages by Application Report

PUSHBUTTON APPLICATION SUMMARY REPORT		
SUBSYSTEM: BSS		
APPLICATION: 1		
PROGRAM ENTERED	MESSAGE ACTION TYPE	MSG / SEC
ACPF	INPUT MESSAGE	4.04
INVL	INVALID INPUT MESSAGE	0.07
WPA1	PASSENGER DATA ENTRY	12.04
WID1	ALTER TRANSACTION INFORMATION	0.94
NFA1	FLIGHT INFORMATION	18.48
PRE1	DISPLAY RECORD	2.07
NAE1	AVAILABILITY OR MISCELLANEOUS	0.00
ETA1	END TRANSACTION	8.52
FRD1	FILE RECORD	0.09
IGR1	IGNORE TRANSACTION	0.07
	SS TOTAL FOR APPLICATION 1	46.32
	HPN SUBTOTAL	46.32

Figure 4. Pushbutton Application Summary Report

Pushbutton Application Summary Report

A set of counters in the control program provides information on the number of high-speed messages started by pushbutton applications. A sample report is shown in Figure 4. The message counters used are never reset, so the counts in this report represent total messages for the applications listed over the entire life of the collection.

The system message counter table (SM) records are used as input to this report. The SM record represents a table of as many as 63 program entries that constitute one RES0 application. Each 8-byte program entry consists of a 4-byte program name and a 4-byte counter field. An SM record is written to RTC tape for each RES0 application for each subsystem user, for each usable I-stream, and for each subsystem. For example, on a system using 2 subsystems and 2 I-streams with 2

subsystem users per subsystem and with 5 RES0 applications implemented and variable MAXAPPL in segment JCD4 changed from 0 to 5, there would be $(2 \times 2 \times 2 \times 5) = 40$ SM records written to tape, both at start and at end time. TPF installations implementing RES0 applications must change segment JCD4 to set variable MAXAPPL to the number of RES0 applications implemented. MAXAPPL defaults to 0. This is done by segment JCD4 (see the SMREC DSECT in JCD4 for SM record layout). The CROSC macro with entry GLBAC is used to retrieve pointers to the global areas where the RES0 application information can be found. The RES0 application information is defined on the pilot tape and updated by the UII package. Note that although SM records are collected by subsystem user, I-stream, and RES0 application, the program groups are unique only by subsystem and by RES0 application number (1 – 5). For this reason, pushbutton applications are reported by subsystem, RES0 application number, and program name only.

Program names and descriptions for additional applications must be entered into the name and description tables DNMTBLC and DMSGDES in segment JRA3 to be reported properly. These tables have a one-to-one correspondence with each other. They are a listing of all programs used by all pushbutton applications. When additional programs are to be supported, their names and descriptions must be inserted into these tables. Because the tables are searched sequentially for each program being reported, the order of programs in the tables is not important as long as each name in DNMTBLC is in the same array position as its description in DMSGDES.

The preceding example assumes that 5 RES0 applications have been defined. Only 1 RES0 application is predefined, and you are responsible for applications 2 through 5, if they are used.

System Summary Report

This report provides a summary of the system, of I-streams, of storage, and of shutdown conditions. The items shown in the I-stream Summary (see Figure 5) may be considered in groups that comprise system utilization, activity per I-stream, and queues or job lists that comprise priority of processing.

SYSTEM SUMMARY	18 FEB	16:12:12	SYSTEM	PAGE	1
205 OBSERVATIONS					
INPUT MESSAGES PER SECOND (WORK LOAD)	MIN	MAX		MEAN	
HIGH SPEED MESSAGES (PROCESSED)	0.615	5.060		2.294	
LOW SPEED	0.000	0.000		0.000	
HIGH SPEED MESSAGES (ROUTED)	0.000	0.000		0.000	
TCP/IP WEIGHTED MESSAGES	1.800	2715.743		909.267	
CREATED ENTRIES	32.141	146.118		87.000	
SSCP INPUT MESSAGES	0.000	0.000		0.000	
UNIT RECORD LOW PRIORITY TASK	0.000	0.000		0.000	
UNIT RECORD HIGH PRIORITY TASK	0.000	0.000		0.000	
SQL REQUESTS PER SECOND	0.000	0.000		0.000	
ACTIVE SQL ECBS	0.000	0.000		0.000	
WEIGHTED MESSAGE RATE	2.415	2720.803		911.561	
RESOURCE UTILIZATION PER MESSAGE					
MILLISECONDS PER WEIGHTED MESSAGE	499.998	12975.598		2144.723	
CORE POOL BYTES PER ECB	37866.852	66018.835		49758.068	
PROCESSOR UTILIZATION					
SYSTEM WAIT STATE	20.2%	74.7%		50.8%	
TPF AVERAGE CPU UTILIZATION	25.3%	79.8%		49.2%	
PAUSE COUNT	0.000	1.000		0.102	
MAIN I-STREAM WAIT TIME (MILS)	0.000	0.810		0.021	
TIME SPENT WORKING WHILE PAUSED (MILS)	0.000	0.010		0.003	
PAUSE TIME -- I-STREAM #2 (MILS)	0.000	0.000		0.000	
PAUSE TIME -- I-STREAM #3 (MILS)	0.000	0.000		0.000	
PAUSE TIME -- I-STREAM #4 (MILS)	0.000	0.000		0.000	
PAUSE TIME -- I-STREAM #5 (MILS)	0.000	0.000		0.000	
PAUSE TIME -- I-STREAM #6 (MILS)	0.000	0.000		0.000	
PAUSE TIME -- I-STREAM #7 (MILS)	0.000	0.000		0.000	
PAUSE TIME -- I-STREAM #8 (MILS)	0.000	0.000		0.000	
PAUSE TIME -- I-STREAM #9 (MILS)	0.000	0.000		0.000	
PAUSE TIME -- I-STREAM #10 (MILS)	0.000	0.000		0.000	
IPTE INSTRUCTION RATE	6.676	24.918		14.838	
PTLB INSTRUCTION RATE	0.985	16.582		6.430	
COLLECTION INTERVAL IN SECONDS	4.161	7.036		5.202	
TPF CPU LIST LENGTH (SUMMATION OF ALL I-STREAMS)					
CROSS	0.000	200.000		4.902	
READY	0.000	236.000		3.857	
INPUT	0.000	33.000		2.402	
DEFERRED	0.000	1.000		0.878	
TOTAL SUSPEND LIST ECBS	0.000	1.000		0.009	
SUSPEND LIST LOW-PRIORITY ECBS	0.000	0.000		0.000	
SUSPEND LIST TIME-SLICED ECBS	0.000	0.000		0.000	
VCT	0.000	0.000		0.000	
LOW PRIORITY ECB CLASSIFICATIONS					
PRIORITY CLASS - BATCH	0.000	0.000		0.000	
PRIORITY CLASS - LOBATCH	0.000	0.000		0.000	
PRIORITY CLASS - IBMHI	0.000	0.000		0.000	
PRIORITY CLASS - IBMLO	0.000	0.000		0.000	

Figure 5. System Summary Report (Part 1 of 8)

SYSTEM SUMMARY		18 FEB	16:12:12	SYSTEM	PAGE	2
205 OBSERVATIONS						
INSTANTANEOUS ACTIVITY						
ACTIVE ECBS		215.000	309.000		244.112	
TIME-SLICEABLE ECBS		0.000	0.000		0.000	
TIME SLICES (PER SECOND)		0.000	0.000		0.000	
THREADS SUPPORT DATA						
MAXIMUM THREADS PER PROCESS		0.000	5.000		4.643	
HIGH WATER MARK NUMBER OF THREADS ACTIVE IN A PROCESS:						
		5				
EQUATIONS USED TO MANIPULATE DATA						
WEIGHTED MESSAGE RATE = (HS MSG.PROCESSED-HS MSG.ROUTED) + (LS MSG.RATE * 5) + (HS MSG.ROUTED * 0.3) + TCP/IP WEIGHTED MSG						
TPF CPU UTILIZATION = ((TOTAL CLOCK TIME) - (SUM OF ALL IDLE LEVELS))/ (TOTAL CLOCK TIME)						
**NOTES: 1) IF MAX OR MEAN = INFINITY, THEN THE WEIGHTED MESSAGE RATE = 0						
2) ACTIVE ECBS AND ACTIVE MESSAGE RATE ARE ANALOGOUS TERMS						
3) CPU UTILIZATION IS COMPUTED USING THE SYSTEM WAIT STATE PERCENTAGE						

Figure 5. System Summary Report (Part 2 of 8)

I-STREAM SUMMARY REPORT		18 FEB	16:12:12	SYSTEM	PAGE	3
205 OBSERVATIONS						
THE MAIN I-STREAM IS		I-STREAM 1				
NUMBER OF I-STREAMS IN USE		10				
NUMBER OF I-STREAMS CAPABLE OF BEING USED		10				
		I-STREAM 1	I-STREAM 2	I-STREAM 3	I-STREAM 4	
CPU HARDWARE ADDRESS		0	1	2	3	
PROCESSOR UTILIZATION						
IDLE LEVEL 3 - NOT ENOUGH FRMS FOR INPUT LIST		0.0%	0.1%	0.2%	0.1%	
IDLE LEVEL 11 - NO WORK FOR TPF		0.0%	0.0%	61.9%	62.3%	
CPU UTILIZATION		100.0%	99.9%	37.9%	37.6%	
SMOOTHED CPU UTILIZATION (MEAN)		13.8%	11.6%	8.0%	8.1%	
ADJUSTED CPU UTILIZATION (MEAN)		13.6%	11.1%	8.0%	8.1%	
INSTANTANEOUS ACTIVITY						
ACTIVE ECBS (MEAN)		14.244	31.678	30.976	30.971	
TPF CPU LIST LENGTH (MEAN)						
CROSS		2.001	0.452	0.338	0.343	
READY		2.150	0.600	0.126	0.149	
INPUT		1.723	0.611	0.008	0.008	
DEFERRED		0.884	0.000	0.000	0.000	
TOTAL SUSPEND LIST ECBS (MEAN)		0.005	0.000	0.000	0.000	
SUSPEND LIST LOW-PRIORITY ECBS (MEAN)		0.000	0.000	0.000	0.000	
SUSPEND LIST TIME-SLICED ECBS (MEAN)		0.000	0.000	0.000	0.000	
VCT (MEAN)		0.000	0.000	0.000	0.000	
INPUT MESSAGES PER SECOND						
CREATED ENTRIES (VIA SWISC)		0.028	0.016	0.102	0.100	
ROUTED ENTRIES (VIA SWISC)		0.067	0.002	0.390	0.374	
MILS PER ROUTED ENTRY		-	499500.000	971.794	1005.347	
EQUATIONS USED TO MANIPULATE DATA						
MILS PER ROUTED ENTRY = (CPU UTILIZATION) / (ROUTED ENTRIES)						
TPF CPU UTILIZATION = ((TOTAL CLOCK TIME) - (SUM OF ALL IDLE LEVELS))/ (TOTAL CLOCK TIME)						
**NOTES: 1) SMOOTHED CPU UTILIZATION IS COMPUTED ONCE PER SECOND USING A SMOOTHING ALGORITHM						
2) ADJUSTED CPU UTILIZATION DOES NOT INCLUDE TIME SPENT WHILE EXECUTING TASKS THAT ARE EITHER DELAYED OR DEFERRED						
3) MILS PER ROUTED ENTRY NOT MEANINGFUL FOR MAIN I-STREAM BECAUSE MAIN I-STREAM DOES WORK ON BEHALF OF APPLICATION I-STREAMS						

Figure 5. System Summary Report (Part 3 of 8)

I-STREAM SUMMARY REPORT		18 FEB	16:12:12	SYSTEM	PAGE	4
205 OBSERVATIONS						
THE MAIN I-STREAM IS		I-STREAM 1				
NUMBER OF I-STREAMS IN USE	10					
NUMBER OF I-STREAMS CAPABLE OF BEING USED	10					
		I-STREAM 5	I-STREAM 6	I-STREAM 7	I-STREAM 8	
CPU HARDWARE ADDRESS	4	5	6	7		
PROCESSOR UTILIZATION						
IDLE LEVEL 3 - NOT ENOUGH FRMS FOR INPUT LIST	0.1%	0.2%	0.2%	0.2%		
IDLE LEVEL 11 - NO WORK FOR TPF	61.4%	59.5%	60.6%	60.2%		
CPU UTILIZATION	38.5%	40.3%	39.2%	39.6%		
SMOOTHED CPU UTILIZATION (MEAN)	8.1%	8.0%	8.1%	8.2%		
ADJUSTED CPU UTILIZATION (MEAN)	8.1%	8.0%	8.1%	8.2%		
INSTANTANEOUS ACTIVITY						
ACTIVE ECBS (MEAN)	31.273	31.400	31.946	32.434		
TPF CPU LIST LENGTH (MEAN)						
CROSS	0.349	0.361	0.375	0.361		
READY	0.138	0.179	0.229	0.163		
INPUT	0.010	0.010	0.018	0.008		
DEFERRED	0.000	0.000	0.000	0.000		
TOTAL SUSPEND LIST ECBS (MEAN)	0.000	0.000	0.000	0.000		
SUSPEND LIST LOW-PRIORITY ECBS (MEAN)	0.000	0.000	0.000	0.000		
SUSPEND LIST TIME-SLICED ECBS (MEAN)	0.000	0.000	0.000	0.000		
VCT (MEAN)	0.000	0.000	0.000	0.000		
INPUT MESSAGES PER SECOND						
CREATED ENTRIES (VIA SWISC)	0.109	0.098	0.109	0.114		
ROUTED ENTRIES (VIA SWISC)	0.367	0.371	0.381	0.383		
MILS PER ROUTED ENTRY	1049.046	1086.253	1028.871	1033.942		
EQUATIONS USED TO MANIPULATE DATA						
MILS PER ROUTED ENTRY = (CPU UTILIZATION) / (ROUTED ENTRIES)						
TPF CPU UTILIZATION = ((TOTAL CLOCK TIME) - (SUM OF ALL IDLE LEVELS)) / (TOTAL CLOCK TIME)						
**NOTES: 1) SMOOTHED CPU UTILIZATION IS COMPUTED ONCE PER SECOND USING A SMOOTHING ALGORITHM						
2) ADJUSTED CPU UTILIZATION DOES NOT INCLUDE TIME SPENT WHILE EXECUTING TASKS THAT ARE EITHER DELAYED OR DEFERRED						
3) MILS PER ROUTED ENTRY NOT MEANINGFUL FOR MAIN I-STREAM BECAUSE MAIN I-STREAM DOES WORK ON BEHALF OF APPLICATION I-STREAMS						

Figure 5. System Summary Report (Part 4 of 8)

I-STREAM SUMMARY REPORT		18 FEB	16:12:12	SYSTEM	PAGE	5
205 OBSERVATIONS						
THE MAIN I-STREAM IS		I-STREAM 1				
NUMBER OF I-STREAMS IN USE		10				
NUMBER OF I-STREAMS CAPABLE OF BEING USED		10				
		I-STREAM 9	I-STREAM 10			
CPU HARDWARE ADDRESS		8	9			
PROCESSOR UTILIZATION						
IDLE LEVEL 3 - NOT ENOUGH FRMS FOR INPUT LIST		0.2%	0.1%			
IDLE LEVEL 11 - NO WORK FOR TPF		68.6%	70.5%			
CPU UTILIZATION		31.2%	29.4%			
SMOOTHED CPU UTILIZATION (MEAN)		7.7%	7.9%			
ADJUSTED CPU UTILIZATION (MEAN)		7.7%	7.9%			
INSTANTANEOUS ACTIVITY						
ACTIVE ECBS (MEAN)		4.132	5.059			
TPF CPU LIST LENGTH (MEAN)						
CROSS		0.164	0.167			
READY		0.065	0.063			
INPUT		0.005	0.004			
DEFERRED		0.000	0.000			
TOTAL SUSPEND LIST ECBS (MEAN)		0.000	0.000			
SUSPEND LIST LOW-PRIORITY ECBS (MEAN)		0.000	0.000			
SUSPEND LIST TIME-SLICED ECBS (MEAN)		0.000	0.000			
VCT (MEAN)		0.000	0.000			
INPUT MESSAGES PER SECOND						
CREATED ENTRIES (VIA SWISC)		0.097	0.129			
ROUTED ENTRIES (VIA SWISC)		0.378	0.365			
MILS PER ROUTED ENTRY		825.396	805.479			
EQUATIONS USED TO MANIPULATE DATA						
MILS PER ROUTED ENTRY = (CPU UTILIZATION) / (ROUTED ENTRIES)						
TPF CPU UTILIZATION = ((TOTAL CLOCK TIME) - (SUM OF ALL IDLE LEVELS)) / (TOTAL CLOCK TIME)						
**NOTES: 1) SMOOTHED CPU UTILIZATION IS COMPUTED ONCE PER SECOND USING A SMOOTHING ALGORITHM						
2) ADJUSTED CPU UTILIZATION DOES NOT INCLUDE TIME SPENT WHILE EXECUTING TASKS THAT ARE EITHER DELAYED OR DEFERRED						
3) MILS PER ROUTED ENTRY NOT MEANINGFUL FOR MAIN I-STREAM BECAUSE MAIN I-STREAM DOES WORK ON BEHALF OF APPLICATION I-STREAMS						

Figure 5. System Summary Report (Part 5 of 8)

BLOCK TYPE	BYTES ALLOCATED	BLOCKS ALLOCATED	MEAN BLOCKS AVAILABLE	MEAN BLOCKS IN USE	UTILIZATION		
					MIN	MAX	MEAN
LIOCB	1142784	4464	4272.317	191.6	3.71%	5.77%	4.29%
LECB	15974400	1300	1049.405	250.5	16.92%	24.30%	19.27%
LSWB	897024	1752	1364.937	387.0	13.69%	32.24%	22.09%
LCOMMON	2523136	616	589.229	26.7	3.57%	6.81%	4.34%
LFRAME	25395200	6200	3282.351	2917.6	32.54%	75.72%	47.05%
TOTAL	45932544						

MAXIMUM NUMBER OF FRAMES ON FRAME PENDING LIST: 70

SYSTEM HEAP UTILIZATION

SYSTEM HEAP SIZE: 10485760 BYTES, 2560 PAGES
 **** HIGHWATER MARK: 5906432 BYTES IN USE, 1442 FRAMES IN USE, 56.328 PERCENT IN USE
 **** THE HIGHWATER MARK OCCURRED PRIOR TO THIS DATA COLLECTION.

OBSERVATION MINIMUM: 2658304 BYTES IN USE, 649 FRAMES IN USE, 25.351 PERCENT IN USE
 MEAN: 2869596 700 27.366
 MAXIMUM: 4182016 1021 39.882

OBSERVATION MINIMUM: 11 SHARED MEMORY FRAMES IN USE, 1 SHARED MEMORY SEGMENTS IN USE
 MEAN: 11 1
 MAXIMUM: 11 1

OBSERVATION MINIMUM: 14.100 PERCENT OF ALL FRAMES IN USE
 MEAN: 24.303
 MAXIMUM: 33.934

OBSERVATION MINIMUM: 3.903 SYSTEM HEAP REQUESTS PER SECOND
 MEAN: 14.038
 MAXIMUM: 22.460

THERE WERE NOT ANY UNSUCCESSFUL SYSTEM HEAP REQUESTS.

Figure 5. System Summary Report (Part 6 of 8)

SYSTEM SUMMARY REPORT		18 FEB	16:12:12	SYSTEM	PAGE	7
205 OBSERVATIONS						
PROGRAM STORAGE UTILIZATION						
STORAGE RESERVED FOR 24-BIT CRPA	6200000 BYTES	(6.2 MEGABYTES)			
STORAGE AVAILABLE IN THE 24-BIT CRPA	5702896 BYTES	(5.7 MEGABYTES)			
24-BIT CRPA FULL, PROGRAM NOT LOADED IN 24-BIT AREA	0 OCCURRENCES					
STORAGE RESERVED FOR 31-BIT CRPA	8200000 BYTES	(8.2 MEGABYTES)			
STORAGE AVAILABLE IN THE 31-BIT CRPA	16 BYTES	(0.0 MEGABYTES)			
31-BIT CRPA FULL, PROGRAM NOT LOADED IN 31-BIT AREA	0 OCCURRENCES					
STORAGE RESERVED FOR PAT	984096 BYTES	(1.0 MEGABYTES)			
PAT SLOTS ALLOCATED	10249 SLOTS					
STORAGE RESERVED FOR EXTRA PAT	192024 BYTES	(0.2 MEGABYTES)			
EXTRA PAT SLOTS ALLOCATED FOR E-TYPE LOADER	2000 SLOTS					
COMMIT/ROLLBACK DATA						
RECOVERY LOG TRACK BUFFERS ALLOCATED	10					
RECOVERY LOG RECORDS RESIDE ON SUBSYSTEM	SSN					
RECOVERY LOG RECORDS ALLOCATED (THIS PROCESSOR)	2040					
RECOVERY LOG TRACKS ALLOCATED (THIS PROCESSOR)	170					
RECOVERY LOG TRACKS RESERVED FOR RECOVERY (THIS PROCESSOR)	50					
RECOVERY LOG TRACKS IN USE - MAXIMUM SINCE LAST IPL	8					
COMMIT SCOPE BUFFERS PER COMMIT SCOPE-USER MAX	0					
COMMIT SCOPE BUFFERS PER COMMIT SCOPE-MAX SINCE IPL	19					
	MIN		MAX			MEAN
WLOGC BLOCKED CONDITIONS, PER SECOND	0.000		0.000			0.000
RECOVERY LOG TRACK WRITES, PER SECOND	0.434		6.250			1.972
COMMITTS, PER SECOND	0.579		7.781			3.269
ROLLBACKS, PER SECOND	0.000		0.000			0.000
381-BYTE COMMIT BUFFERS IN USE IN VFA	0.000		0.000			0.000
1055-BYTE COMMIT BUFFERS IN USE IN VFA	0.000		0.000			0.000
4095-BYTE COMMIT BUFFERS IN USE IN VFA	0.000		0.000			0.000

Figure 5. System Summary Report (Part 7 of 8)

SYSTEM SUMMARY REPORT		18 FEB	16:12:12	SYSTEM	PAGE	8
205 OBSERVATIONS						
TPF SHUTDOWN CONDITIONS (TOTAL FOR ALL I-STREAMS)						
TASK DESCRIPTION	SHUTDOWN LEVEL					NUMBER OF SHUTDOWN OCCURRENCES
INPUT LIST	MORE THAN 975 ACTIVE	ECB BLOCKS				--
INPUT LIST	LESS THAN 1550 AVAILABLE	FRM BLOCKS				2
INPUT LIST	LESS THAN 92 AVAILABLE	COM BLOCKS				--
INPUT LIST	LESS THAN 325 AVAILABLE	ECB BLOCKS				--
INPUT LIST	LESS THAN 263 AVAILABLE	SWB BLOCKS				--
INPUT LIST	LESS THAN 446 AVAILABLE	IOB BLOCKS				--
DEFERRED LIST	MORE THAN 1170 ACTIVE	ECB BLOCKS				--
TIME AVAILABLE SUPERVISOR	LESS THAN 3100 AVAILABLE	FRM BLOCKS				--
TIME AVAILABLE SUPERVISOR	LESS THAN 308 AVAILABLE	COM BLOCKS				--
TIME AVAILABLE SUPERVISOR	LESS THAN 650 AVAILABLE	ECB BLOCKS				--
CREM MACRO	LESS THAN 175 AVAILABLE	SWB BLOCKS				--
CREM MACRO	MORE THAN 50 ACTIVE	RDY BLOCKS				--
CRED MACRO	LESS THAN 438 AVAILABLE	SWB BLOCKS				--
CREX MACRO	LESS THAN 438 AVAILABLE	SWB BLOCKS				--
BSC INPUT	MORE THAN 300 ACTIVE	INP BLOCKS				--
3270 LOCAL INPUT	MORE THAN 300 ACTIVE	INP BLOCKS				--
AI INPUT	MORE THAN 0 ACTIVE	INP BLOCKS				--
						NUMBER OF SLOWDOWN OCCURRENCES
SNA NCP SLOWDOWNS						--
SNA CTC INPUT SLOWDOWNS						--
SNA CTC OUTPUT SLOWDOWNS						--

Figure 5. System Summary Report (Part 8 of 8)

Maximum, minimum, and mean values are shown for each variable. For the variables that are continuous counts, reduced to a per second base, the maximum and minimum values are actually mean values for a single collection interval. The mean values are the mean of all the intervals, that is, a mean of all the interval means. The maximum and minimum values are indicative of the degree of variation that occurs in system load. Extreme values should lead one to the plot reports in

which individual interval values can be examined. Note that the maximum and minimum values for different variables do not necessarily occur during the same interval.

For mean values, a long interval has a smoothing effect on the peaks and valleys of a sample set. For example, a peak might show a much higher maximum for a one-second interval than for a 15-second interval. Appropriate run lengths, periods, and intervals must be chosen to provide detailed data for each individual requirement.

The shutdown levels for system resources are established as absolute numbers of system resources. These levels can be redefined using the ZCTKA command:

- Maximum number of ECBs allowed in system
- Minimum number of ECBs required to service the input list
- Maximum number of ECBs required to service the defer list
- Minimum number of ECBs required to service the time available supervisor
- Minimum number of SWB blocks required to service the input list
- Minimum number of IOB blocks required to service the input list
- Minimum number of Frames required to service the input list.

Note: CREMC activity is postponed whenever the ready list of the target I-stream exceeds 50 entries in length. This is a throttling mechanism and should not be considered a problem unless other TPF shutdown conditions are also reported.

The list should be kept current and readily available for the analyst working with the data reduction reports. A record of all changes to these values should be maintained in a history file so that reduction reports for a particular calendar time period can be analyzed with respect to the proper level setting.

Input Messages Per Second

Inputs are shown on a per second basis and consist of:

- High-speed processed messages (agent messages)
- Low-speed messages (teletype)
- TCP/IP messages (TCP/IP native stack only)
- System Services SSCP Messages (messages from NCP)
- High-speed routed messages (application to application)
- Created entries (CREM, CRED, CREX, and CRET macros)
- Unit Record - High Priority
- Unit Record - Low Priority.

If utilities such as schedule change or file maintenance and no unit record tasks are run during data collection, the results of the reduction will be skewed by the load supplied by these utilities (one input message could produce an inordinate amount of file activity).

The ratio of one type of message to another is constantly changing. To relate one data collection to another, or one system to another, a common denominator is needed. Therefore, all message types, except TCP/IP native stack input messages, are expressed in terms of high-speed messages. See *TPF Transmission Control Protocol/Internet Protocol* for more information about counts for TCP/IP native stack messages.

The weighted message rate is calculated using the following factors:

- Low-speed message weighting factor (WT)
- High-speed routed message weighting factor (WT2)
- TCP/IP native stack weighted messages.

System services messages are not included in the weighted message rate because they are considered network overhead and are reported solely for information.

Created entries are not included in the weighted message rate because the majority are generated by, and considered to be part of, some other, external message that is already counted.

The weighted message rate algorithm is as follows:

$$\text{WEIGHTED MESSAGE RATE} = (\text{HS MSG PROC}) - (\text{HS MSG ROUT}) + (\text{LS MSG} * \text{WT}) + (\text{HS MSG ROUT} * \text{WT2}) + (\text{TCP/IP WEIGHTED MSG})$$

Weighting factors WT and WT2 are preprocessor statements entered during SIP time which may be adjusted for each system. However, the most commonly used value for WT is 5. The most commonly used value for WT2 is 0.3; therefore, WT2 is coded as 3 in the DATACO macro.

Experience with existing systems indicates that a teletype message requires four to five times the processing required by a high-speed message; therefore, most systems use the default value. Processing required by the average routed message is roughly three-tenths (.3) of that required by the average high-speed message. Again, most systems use the default value.

Weighting factors may also be established for a given TPF system through the real-time trace facility. Message types may be sized by comparing number of ENTERs, FINDs, FILEs, MACROs, and so on. Total processing by message type is not easily available without special software or hardware monitors. The comparison method yields satisfactory approximations for weighting purposes.

Note: In the I-stream summary report, the counts for SWISC entries also include tpf_fork function calls.

Resource Utilization Per Message

Resource utilization per message involves two parameters: CPU busy time per weighted message as determined previously, and the number of bytes of working storage per active ECB. Frames allocated for system heap are not included in the number of bytes of working storage in use for this calculation. CPU busy time per weighted message is obtained by dividing TPF Processor Utilization by the Weighted Message Rate.

Processor Utilization

In the System Summary Report, PAUSE COUNT is the number of times per second that the main I-stream intentionally ran exclusively. For a multiprocessor, this is also the number of times all application I-streams were paused. Pausing has no impact for uniprocessors because there is only a main I-stream. The information is nevertheless provided as a migration aid. MAIN I-STREAM WAIT TIME (MILS) is the amount of time that the main I-stream spent waiting for other I-streams to be paused. The TIME SPENT WORKING WHILE PAUSED (MILS) is the amount of time that the main I-stream was actually working while all other I-streams were paused. The CPU utilization stated on the system summary computes utilization using idle time and collection interval time.

For the I-Stream Summary Report, data collection shows the processor utilization at each level and state and, in addition, gives TPF processor utilization. Data

Collection uses the time-of-day (TOD) clock to accumulate elapsed time in 12 idle levels for the TPF system, including the wait state. The idle levels indicate if available working storage is too low or system activity (number of active ECBs) is too high. Idle Level 11 shows the percentage of time that the system (as far as TPF is concerned) is truly idle; that is, the control program has scanned the entire CPU loop and found no work to be done. In addition, the idle times greater than zero will also be listed on the I-Stream Summary Report. The I-Stream Summary Report contains three utilization values. The first is calculated the same as for the system summary. The second (smoothed) utilization is the mean value of the utilization used by the work scheduler. The third (adjusted) utilization is calculated in the same manner as the CPU utilization on the I-Stream Summary Report, except that the time spent processing either DLAYC or DEFRC macros is subtracted from the utilization. The amount of time subtracted for each macro processed is an estimate based on the DEFRC or DLAYC and any associated processing. The intent of reporting the adjusted utilization is to show the effort the TPF system is using to process actual work.

The control parameters that trigger the various idle levels are established at system initialization time and must be adjusted for a particular system. The parameter settings are a compromise that allow maximum utilization of resources yet protect the system from catastrophic resource depletion during peak conditions. Any significant idle occupancy on levels 1-12 warrants investigation. This means that either the parameters are too restrictive or the system is approaching maximum capacity because of working storage depletion.

Finally, when there is no work for the system, the system enters the wait state.

Collection Interval in Seconds

The collection interval is the actual duration of the sample rather than the specified duration. The specified and actual duration should be very nearly the same, but there may be some conditions that will cause slight variations in the actual duration.

TPF CPU List

The various lists appearing in the System Summary represent the job queues that are associated with the CPU loop. The sequence in which they appear also indicates the priority given to each type of job. The maximum, minimum, and mean values represent an actual count of the items taken twice during each data collection interval for sampling mode, and once for continuous mode. To analyze the list data, one must be knowledgeable of the CPU loop and the types of jobs placed on each of the lists by the control program. Because systems vary widely with regard to resources and load, it is difficult to quote numbers that one would expect to see on a particular list. However, guidelines may be established, and any significant variance warrants detailed analysis.

- The absence of job queues indicates an unrestricted flow of data through the system. For a well-balanced system running at a comfortable level of utilization, the queues on the cross, ready, input, or deferred lists are expected to be very low, almost nonexistent.
- If a system is operating under a heavy load and significant queues are present, queue lengths can be expected to appear in ascending order; that is, the largest queue should represent the lowest priority job list. The cross list should have the shortest queue at all times.
- Any significant queue should correlate to other data appearing in the system summary. For example, a long deferred list should not occur unless the CPU occupancy is very high, near capacity. A long input list should not develop unless the number of active ECBs and CPU utilization indicate that the system is being heavily utilized, and has possibly reached a cutoff level.

- The number of jobs queued on the suspend list depends on the number of ECBs in the system that can be time-sliced or are low-priority. As system resource usage increases, a corresponding increase occurs in the number of low-priority ECBs that are put on the suspend list.

Low-Priority ECB Classifications

The low-priority ECB classification display shows the number of ECBs running for priority classes defined in the system. A priority class is assigned to an ECB when the LODIC macro is issued.

Instantaneous Activity

This is a snapshot of the instantaneous activity in the system. The number of active ECBs is snapped twice during each interval in sampling mode, and once per interval in continuous mode. The active ECBs include the ECBs used for data collection, and those used for unit record tasks. Data reduction subtracts the data collection ECBs, and the ECBs used for unit record tasks to determine the number of truly active ECBs.

When related to inputs, the ratio of active ECBs to weighted message rate is an indicator of message life in the system, or throughput, excluding line queue and transmission times.

Example::

Active Messages = 10

Weighted Message Rate = 20 per second

Therefore, the average message life in the system is one-half second or 500 milliseconds. Long mean message life might indicate a bottleneck, which could be because of insufficient working storage, improper program allocation, insufficient number of files, large file queues, and so on.

The maximum allowable number of active ECBs in a system is a control parameter that is established at initialization time. As with all control parameters, this number must be adjusted for each unique configuration. It must be high enough to allow full utilization of resources, yet low enough to protect the system during overloads. Correlation of active ECBs with message rates, working storage occupancy, idle occupancy, and input list helps eliminate the trial and error method of fine tuning all control parameters. ECBs that can be time-sliced are recorded here, along with the number of time slices that occur per second. ECBs are marked as being available for time slicing by the TMSLC macro.

Threads Support Data

To help ensure that there are enough threads defined to the TPF 4.1 system for any application that uses threads, data is collected from the following fields:

- The maximum number of threads for each process
The maximum number of threads active in any process at any time during an iteration of data collection
- The high-water mark number of threads active in a process
The maximum number of threads active in any process at any time during the current initial program load (IPL).

If a thread application, such as remote procedure call (RPC), is using a large number of threads compared to the number defined in keypoint A (CTKA), you can enter the ZCTKA ALTER command to change the maximum number of threads.

Equations Used to Manipulate Data

The possible complexity of the system operation and the amount and type of data collected could lead to questions of interpretation and meaning. For that reason data collection printouts include the equations used to compute:

- Weighted message rate
- Processor utilization
- Active messages
- Mils per routed Enter.

Working Storage Configuration

The total working storage and the ratio of the number of one size block to another will vary from system to system. These parameters are adjustable and established at system initialization time. The number of blocks assigned to each pool is also printed on the System Summary Report.

In the System Summary Report, block occupancy (utilization) is expressed in two forms. The minimum, maximum, and mean block occupancies are expressed as decimal fractions (blocks in use or blocks allocated). Mean available block figures are the average of the actual number of blocks allocated but unused. As with all data collection variables, block occupancy must be related to other variables in order to be meaningful.

For each block size, the mean number in use multiplied by the number of bytes per block equals the mean bytes per pool in use. When the sum of all pools is divided by the number of active messages (as shown in the System Summary Report) the amount of working storage used on a per-message basis should approach system design criteria. Any great variance should be investigated.

Example: Assume that the ECB Block Occupancy is 10% and Idle Occupancy at Level 1 is quite high. Level 1 indicates that the system is not servicing the input list because of an insufficient number of available ECB blocks, yet the ECB pool is only 10% utilized. This indicates that the CPU loop parameter stating the number of ECB blocks that must be available is set too high, or the system is experiencing very unusual peaking conditions that are not reflected in mean utilization of the working storage pool.

The maximum number of frames on the frames pending list is collected to help determine if additional 4-KB frames are needed. If the maximum number is greater than 10% of the frames in the TPF 4.1 system, a system shutdown can occur because it is low on available frames. Frames that are released by threaded ECBs are placed on the frames pending list until it is safe to reuse them; for example, after a purge of the translation look-aside buffer (PTLB) is performed on all I-streams. If a large number of frames remain on the frames pending list, additional 4-KB frames can be generated in the TPF 4.1 system. See *TPF System Generation* for more information about the CORREQ macro.

Commit/Rollback Data

Data is collected on the following activities for TPF transaction services:

- Recovery log
 - Track buffers allocated

Track buffers are main storage buffers that are used by the log manager as a holding area for recovery log data. Track buffers hold requests to write to the recovery log before the requests are written to the DASD surface. This process permits the log manager to achieve better throughput by writing 1 track at a time. Buffer allocation factors are directly related to buffer blocked conditions. Track buffers are also used to read data during log recovery.

Increase the number of buffers you have allocated (using the ZCTKA command) if you need to lower the number of blocked conditions.

Note: The number of buffers is per processor. CTKA is processor-unique.

- Location of recovery log records
The recovery log fixed file records can be located either on the basic subsystem (BSS) or on a user-defined subsystem.
- Records allocated for this processor
These records define the recovery log. They are processor unique and must reside on a single device type. Factors for allocating these records are related to the amount of outstanding commit scope data in the system. A measure of the outstanding data is the amount of recoverable data on the log, which translates into the number of tracks in use.
- Tracks allocated for this processor
This number is determined by dividing the number of records allocated for the recovery log by the number of 4K records contained on a single track of the device where the log resides. This value is useful in determining when the size of the recovery log should be increased.
- Tracks reserved for recovery for this processor
This is the number of recovery log tracks reserved for log recovery processing. This value is equal to 10% of the allocated tracks but not less than 50. This value is useful in determining when the size of the recovery log should be increased.
- Tracks in use (maximum since the last IPL)
This count indicates the amount of recoverable data on the recovery log. Because the recovery log is processed as a wraparound file, there is a danger that the log manager will overwrite valid recoverable data. The TPF system prevents this by keeping track of the number of tracks in use and causing a catastrophic dump when the count reaches a maximum value. The maximum value is defined as the total number of tracks allocated minus the number of tracks reserved for recovery. You want to keep the total number of tracks at a high enough level to ensure that the recovery log will not become full. For example, if you have a system with 400 tracks minus the 50 tracks that are kept as a recovery buffer, and the number of tracks in use is 343, you want to increase the number of log records that you have allocated.
- Commit scope buffers per commit scope
 - User-specified maximum
This number will be 0 if you do not specify a maximum number. (0 means that the maximum number is unlimited.) If the commit scope gets to the maximum size, the TPF system dumps and the ECB exits.
 - Maximum since the last IPL
This is the largest size commit scope that has existed in the TPF system since the last IPL. Use this number to see how close you are getting to the maximum value.
- WLOGC blocked conditions (per second)
A blocked condition indicates that the log manager track buffer area is not large enough to hold all of the log write requests. During these blocked conditions, the TPF system suspends any ECBs that are in the process of committing data and forces them to wait until buffers are free. Buffers become free when they have been successfully written out to file. A high amount of blocked conditions could indicate a performance problem. If you have a high number of blocked conditions per second, you need to define more recovery log buffers.

- Recovery log track writes (per second)
- Commits (per second)
- Rollbacks (per second)
- Commit buffers in use in VFA (381-, 1055-, and 4095-byte).

TPF Shutdown Conditions

As noted previously, control parameters are used to insure that maximum utilization of system resources can be realized, while at the same time protecting against irrecoverable conditions caused by inordinate peak loads. For instance, there is little to be gained and much to be lost by activating a message from the input list when the system activity (indicated by the ECB count) is already very high or the number of available working storage blocks is extremely low.

The control parameters are user-specified and will vary from system to system. Furthermore, there is no exact or scientific method for establishing the value of the various parameters. The System Summary Report is a valuable tool to assist in this effort because the report lists the user-assigned value of each parameter and calculates the number of data collection samples where shutdown occurred because the limits of the particular parameter were exceeded.

System Pools Summary Report

This report lists all active pool sections for the reduced subsystem. The dispensed addresses are shown for each active section. The number of return hits per second is shown for short-term pool sections only. The totals for all short-term sections are listed at the end of the report. The report also lists the number of times that reorder did not occur before active buffer depletion. The System Pools Summary Report is printed as part of SYSTEM reduction only (see Figure 6).

SYSTEM POOLS SUMMARY 30 OBSERVATIONS		BSS	SUBSYSTEM	18 APR	09:14:33	SYSTEM	PAGE	7
POOL SECTION	SET SIZE	REORDER TIME(MIN)	DISPENSED ADDRS/SEC	DISPENSED ADDRS/MSG	STP RETURN HITS/SEC			
DEVA-SST	2	0.00	0.60	1.61	0.31			
DEVA-SDP	2	0.00	0.40	1.09	N/A			
DEVA-LST	2	0.00	0.57	1.52	0.00			
DEVA-LDP	1	0.00	0.00	0.00	N/A			
DEVA-4ST	2	3.16	0.89	2.40	0.00			
DEVA-4DP	1	0.00	1.46	3.92	N/A			
DEVB-4ST	1	0.00	0.00	0.00	0.00			
DEVB-4DP	1	0.00	0.00	0.00	N/A			
DEVA-4D6	2	25.03	1.24	3.31	N/A			
DEVB-4D6	2	21.80	1.77	4.75	N/A			
SHORT TERM POOL TOTALS (ALL SECTIONS)								
TOTAL ADDRESSES DISPENSED -			1250					
TOTAL ADDRESSES RETURNED -			180					
REORDER DID NOT COMPLETE BEFORE ACTIVE BUFFER DEPLETION DURING 0 INTERVALS								

Figure 6. System Pools Summary Report

TPF Logical Record Cache Summary Report

This report includes the following information about the logical record cache to help you determine the most effective cache size for your installation:

- Cache Size
The number of entries in the cache.
- Timeout Value

The amount of time, in seconds, that an entry resides in the cache before it is removed.

- Read Calls Per Second

The number of calls to read an entry from the cache.

- Read Misses Per Second

The number of calls that were not satisfied by the cache because the entry either was not in the cache or had timed out and was not valid.

- Castouts Per Second

The number of valid entries that had to be removed from the cache to make room for a new entry.

- Hit Ratio

The ratio of read calls to read misses.

- Update Invalidates Per Second

The number of times the cache will request the coupling facility (CF) to invalidate the cache entries for other processors for the same hash name because of an updateCacheEntry call by the application. This count will be zero for processor unique caches.

- Read Buffer Invalidates Per Second

The number of times a CF notifies an application to invalidate an entry in the cache. This count will be zero for processor unique caches.

- Duplicate Hash Refused

The number of times a cache entry was not added to the cache because its 16-byte name was the same as an entry already in the cache.

The following shows an example of a TPF Logical Record Cache Summary report.

TPF LOGICAL RECORD CACHE SUMMARY									
11 OBSERVATIONS									
CACHE NAME	CACHE SIZE (ENTRIES)	TIMEOUT VALUE (SECONDS)	READ CALLS PER SECOND	READ MISSES PER SECOND	CASTOUTS PER SECOND	HIT RATIO	UPDATE INVALIDATES PER SECOND	READ BUFFER INVALIDATES PER SECOND	DUPLICATE HASH REFUSED
IDNSHOSTADDR	10	3600	0.00	0.00	0.00	0.00%	0.00	0.00	0
IDNSHOSTNAME	10	3600	0.00	0.00	0.00	0.00%	0.00	0.00	0
MAIL_CACHES	10	3600	0.00	0.00	0.00	0.00%	0.00	0.00	0
TPF_FS_DIR	200	60	0.35	0.24	0.00	30.76%	0.00	0.00	0
TPF_FS_INODE	200	60	0.53	0.01	0.00	98.30%	0.11	0.01	0
CASH	16	3	683.22	1.20	0.00	99.82%	1.04	1.17	0
MONEY	16	3	970.67	2.72	0.00	99.71%	2.97	2.65	0
GOLDBARS	16	3	701.56	0.88	0.00	99.87%	1.57	0.86	0
INGOTS	16	3	1201.32	2.50	0.00	99.84%	1.19	0.82	0
DOLLARS	16	3	664.93	1.00	0.00	99.86%	1.29	1.03	0
ANDCENTS	16	3	769.09	1.04	0.00	99.84%	0.92	1.31	0
PENNIES	16	3	850.94	1.33	0.00	99.81%	2.51	2.15	0
CACH2	16	3	1157.65	2.19	0.00	99.86%	1.47	0.96	0
CACH4	16	3	734.97	0.97	0.00	99.70%	1.49	1.20	0
CACH6	16	3	411.13	1.21	0.00	99.84%	1.19	0.82	0
LOCALC	16	3	0.00	0.00	0.00	0.00%	0.00	0.00	0

Figure 7. TPF Logical Record Cache Summary Report

Coupling Facility Data Collection Reports

Data collection for coupling facility (CF) support, CF record lock support, and logical record cache support provides data to effectively manage the performance of a CF, the CF list structures, and the CF cache structures. The following reports are provided:

- Coupling Facility Usage Summary

- Coupling Facility Structure Summary
- Coupling Facility Locking Summary
- Coupling Facility Caching Summary.

See Figure 8 for an example of the Coupling Facility Usage Summary report. This report lists all the CFs defined in the processor configuration.

The Coupling Facility Usage Summary report provides:

- CF Name
The coupling facility (CF) name, which is a 5- to 8-character alphanumeric name that begins with an alphabetic character and is used by the TPF system to identify a CF.
- Storage in Use
The maximum storage in use, both in megabytes (MB) and as a percentage of total CF storage.
- Chgd?
A Y in this column indicates that the amount of storage in use on this CF has changed during the data collection run. Otherwise, this column is blank.
- Utilization
The CF processor usage, which is calculated as B/B + W, where B represents busy time and W represents the wait time.
- Repl?
A Y in this column indicates that this CF was deleted and another CF with the same name was added during the data collection run. Otherwise, this column is blank.

COUPLING FACILITY USAGE SUMMARY					14 SEP	14.19.15	SYSTEM	PAGE	8
148 OBSERVATIONS									
CF NAME	STORAGE IN USE		CHGD?	UTILIZATION		REPL?			
	MB	PERCENT		MEAN	MAX				
TESTCF1	5.750	24.46%		2.65%	3.23%				
TESTCF2	6.500	27.65%		0.01%	0.33%				

Figure 8. Coupling Facility Usage Summary Report

See Figure 9 on page 31 for an example of the Coupling Facility Structure Summary report. This report lists all the CF structures sorted by the CF names on which the CF structures reside. This particular report shows multiple CFs and many CF structures.

The Coupling Facility Structure Summary report provides:

- Structure Name
The name of the CF structure, organized by CF name.
- Structure Size
The size of the CF structure both in megabytes (MB) and as a percentage of total CF storage.
- Reqs/Sec
The average number of CF requests issued for each second.
- Service Time Usec
The average time, in microseconds, to process a request.

- Queue Time Usec
The average time, in microseconds, that a request waits in a queue before being started.
- Repl?
A Y in this column indicates that a CF structure was deleted and then reallocated during a data collection run. Otherwise, this column is blank.
- CF
The coupling facility (CF) name, which is a 5- to 8-character alphanumeric name that begins with an alphabetic character and is used by the TPF system to identify a CF.

COUPLING FACILITY STRUCTURE SUMMARY				14 SEP	14.19.15	
11 OBSERVATIONS						
STRUCTURE NAME	STRUCTURE SIZE		REQS/SEC	SERVICE TIME	QUEUE TIME	REPL?
	MB	PERCENT		USEC		
CF: FORTKNOX						
ANDCENTS	0.250	1.28%	2.40	9411	1520	
CASH	0.250	1.28%	2.31	14067	991	
DOLLARS	0.250	1.28%	2.25	20184	2251	
GOLDBARS	0.250	1.28%	2.50	9282	6744	
INGOTS	0.250	1.28%	5.54	11443	863	
ITPFLK1_TU0001	5.000	25.64%	55.79	10884	459	
ITPFLK2_TU0001	1.000	5.12%	67.65	6802	393	
MONEY	0.250	1.28%	5.79	11229	553	
PENNIES	0.250	1.28%	2.33	8262	742	
TPF_FS_DIR	0.250	1.28%	0.10	104418	2196	
CF: PIGGYBNK						
CACH2	0.250	0.48%	4.80	13842	4668	
CACH4	0.250	0.48%	2.51	5637	2558	
CACH6	0.250	0.48%	2.80	8664	689	
ITPFLK1_TU0001	9.250	17.96%	23.12	8822	3060	
ITPFLK2_TU0001	1.000	1.94%	23.44	8771	1133	
TPF_FS_INODE	0.250	0.48%	0.22	10635	15572	

Figure 9. Coupling Facility Structure Summary Report

See Figure 10 on page 32 for an example of the Coupling Facility Locking Summary report. This report provides the mean and maximum values for each variable.

The Coupling Facility Locking Summary report provides:

- CF Name
The coupling facility (CF) name, which is a 5- to 8- character alphanumeric name that begins with an alphabetic character and is used by the TPF system to identify a CF.
- Modules
The number of prime modules that are active on the CF and that use the CF for locking at the start of data collection.
- Chgd?
A Y in this column indicates that the number of modules using this CF for locking changed during the data collection run. Otherwise, this column is blank.
- Operations/Request
The average (mean) and maximum number of lock operations performed for each CF request.
- Lists
The number of lists allocated in the locking structure on this CF.

- List Depth
The average (mean) and maximum number of locks (list entries) for a randomly selected set of lists.
- Locks Held
The average (mean) and maximum number of locks currently held in the CF.
- Locks/List
The average (mean) number of locks held in the CF divided by the total number of lists allocated in the locking structure.
- Repl?
A Y in this column indicates that this CF was deleted and another CF with the same name was added during the data collection run. Otherwise, this column is blank.

COUPLING FACILITY LOCKING SUMMARY										14 SEP	14:19:15
11 OBSERVATIONS											
CF NAME	MODULES	CHGD?	OPERATIONS/REQUEST		LISTS	LIST DEPTH		LOCKS HELD		LOCKS/LIST MEAN	REPL?
			MEAN	MAX		MEAN	MAX	MEAN	MAX		
FORTKNOX	14		1.07	10	4929	0.11	8	1223.27	1666	0.25	
PIGGYBNK	9		1.18	11	8427	0.11	4	1003.45	1352	0.12	

Figure 10. Coupling Facility Locking Summary Report

See Figure 11 for an example of the Coupling Facility Caching Summary report. This report provides the cache size and castout values for each CF cache.

The Coupling Facility Caching Summary report provides:

- Cache Name
The cache name, which is a 4- to 12-character alphanumeric name that begins with an alphabetic character and is used by the TPF system to identify a CF cache.
- Cache Size (Entries)
The number of entries in the cache.
- Castouts Per Second
The number of valid entries that had to be removed from the cache to make room for a new entry.
- CF
The coupling facility (CF) name, which is a 5- to 8-character alphanumeric name that begins with an alphabetic character and is used by the TPF system to identify a CF.

COUPLING FACILITY CACHING SUMMARY				27 JUN	07:03:42	SYSTEM	PAGE	12
18 OBSERVATIONS								
CACHE NAME	CACHE SIZE (ENTRIES)	CASTOUTS PER SECOND						
CF: CFSEARS								
TPF_FS_DIR	400	0.00						
TPF_FS_INODE	300	0.00						

Figure 11. Coupling Facility Caching Summary Report

TPF Internet Mail Server Summary Report

This report includes the following information about the TPF Internet mail server to help you determine the most effective configuration for your installation:

- Messages Per Second

Indicates the number of messages per second that were processed by the TPF Internet mail server. This entry shows the number of messages that were received, sent to local and remote clients, and not delivered (bounced).

- Characters Per Message

Indicates the number of characters per message that were received and sent by the TPF Internet mail server.

Note: Internally, the number of characters per message is calculated in units of 1000 to handle larger numbers; therefore, at low message rates or fewer characters per message there will be some inaccuracies. In production systems running normal workloads, any inaccuracy in the characters per message numbers should not be significant.

- Queue Length (Blocks)

Indicates the amount of mail (in blocks) in the active and deferred queues that is currently waiting to be delivered. Each block contains from 1 to 50 pieces of mail.

- Delivery Managers

Indicates the following for the local, remote, and deferred delivery managers:

- Mean Number Active

The number of ECBs that are currently sending mail at the same time.

- Maximum Allowed

The maximum number of ECBs that are allowed to send mail at the same time.

If the queue length of the active queue is too large for your environment or if the queue length increases over time, there are not enough delivery managers defined for the local or remote active queue. Similarly, if the queue length of the deferred queue is too large or increases, there are not enough delivery managers defined for the deferred queue. The number of delivery managers is defined in the TPF configuration file (`/etc/tpf_mail.conf`). See *TPF Transmission Control Protocol/Internet Protocol* for more information about the TPF configuration file.

The following shows an example of a TPF Internet mail server summary report.

TPF INTERNET MAIL SERVER SUMMARY					16 FEB	10:24:36	SYSTEM	PAGE	11
300 OBSERVATIONS									
	IN	LOCAL	OUT	REMOTE	BOUNCED				
MESSAGES PER SECOND	4.77	0.63		0.00	0.09				
CHARACTERS PER MESSAGE	602			0					
	ACTIVE	DEFERRED							
QUEUE LENGTH (BLOCKS)	80.62	176.12							
DELIVERY MANAGERS	LOCAL	REMOTE		DEFERRED					
MEAN NUMBER ACTIVE	9.37	7.62		0.00					
MAXIMUM ALLOWED	10.00	10.00		10.00					

Figure 12. TPF Internet Mail Server Summary Report

TCP/IP Weighted Input Messages by Application Report

This report includes information about the number of TCP/IP weighted input messages for each TCP/IP native stack application for which there is activity. The TCP/IP weighted message by application report provides the following information:

- **Application**
The name of the application. The application must be defined in the TCP/IP network services database. If an application is not defined in the database, the number of messages is shown in the OTHER category. See *TPF Transmission Control Protocol/Internet Protocol* for more information about defining applications in the TCP/IP network services database.
- **Port**
The port associated with this application.
- **Weight**
The weighting factor used when counting the input messages for this application. This weighting factor can be defined for each application in the TCP/IP network services database by specifying the weight parameter. If a weighting factor is not defined for the application, this column contains asterisks (***) . See *TPF Transmission Control Protocol/Internet Protocol* for more information about defining applications in the TCP/IP network services database and how message counts are incremented for TCP/IP applications.
- **Weighted Messages**
The number of weighted messages received by this application.
- **Weighted Msgs/Sec**
The number of weighted messages per second received by this application.
- **Percent of Total**
The percentage of the total number of TCP/IP messages received in the TPF system.
- **Cumulative Percent**
The cumulative percentage of the total number of TCP/IP messages received in the TPF system.

The information in this report is sorted in descending order by activity; that is, the application with the highest activity is shown at the top of the report.

The following shows an example of a TCP/IP weighted input messages by application report.

TCP/IP WEIGHTED MESSAGES BY APPLICATION							04 FEB	14:44:57	SYSTEM	PAGE	9
APPLICATION	PORT	WEIGHT	WEIGHTED MESSAGES	WEIGHTED MSGS/SEC	PERCENT OF TOTAL	CUMULATIVE PERCENT					
TEST-9981	9981	***	278463	153.36	17.07%	17.07%					
TEST-9980	9980	***	278304	153.27	17.06%	34.14%					
TEST-9982	9982	***	278138	153.18	17.05%	51.20%					
TEST-9984	9984	***	278041	153.12	17.05%	68.25%					
TEST-9983	9983	***	258880	142.57	15.87%	84.13%					
TEST-9985	9985	***	258382	142.30	15.84%	99.97%					
OTHER		***	169	0.09	0.01%	99.98%					
RIP	520	50	65	0.04	0.00%	99.99%					
FTP-DATA	20	100	58	0.03	0.00%	99.99%					
TFTP	69	100	42	0.02	0.00%	100.00%					
TOTAL			1630542	897.98	100.00%	100.00%					

Figure 13. TCP/IP Weighted Input Messages by Application Report

ECB Frame Usage Summary Report

The number of frames held by exiting ECBs is reported in histogram form.

TPF ECB FRAME USAGE SUMMARY REPORT			SYSTEM WIDE	27 FEB	14:18:27	SYSTEM	PAGE 7
MEAN TOTAL FRAMES USED DURING ECB LIFETIME			1 FRAMES				
CLASS UPPER LIMIT	FREQUENCY OBSERVED	PERCENT OF TOTAL	FREQUENCY DIAGRAM (SCALE = 968/1)				
0	77362	58.58%	*****				
1	428	0.32%	*				
2	158	0.12%	*				
3	19900	15.07%	*****				
4	34022	25.76%	*****				
5	41	0.03%	*				
6	40	0.03%	*				
7	19	0.01%	*				
8	0	0.00%					
9	41	0.03%	*				
10	0	0.00%					
11	0	0.00%					
12	0	0.00%					
13	0	0.00%					
14	0	0.00%					
15	0	0.00%					
16	0	0.00%					
17	0	0.00%					
18	0	0.00%					
19	0	0.00%					
20	0	0.00%					
25	20	0.02%	*				
30	20	0.02%	*				
35	0	0.00%					
40	1	0.00%	*				
45	0	0.00%					
50	0	0.00%					
60	0	0.00%					
70	8	0.01%	*				
80	0	0.00%					
90	0	0.00%					
100	0	0.00%					
120	0	0.00%					
140	0	0.00%					
160	0	0.00%					
180	0	0.00%					
200	0	0.00%					
220	0	0.00%					
240	0	0.00%					
> 240	0	0.00%					

Figure 14. ECB Frame Usage Summary Report

ECB Heap Area Usage Summary Report

The number of heap frames held by exiting ECBs is reported in histogram form.

TPF ECB HEAP AREA USAGE SUMMARY REPORT			SYSTEM WIDE	27 FEB	14:18:27	SYSTEM	PAGE	8
MEAN FRAMES USED FOR HEAP STORAGE DURING ECB LIFETIME			0 FRAMES					
MAXIMUM FRAMES USED FOR HEAP STORAGE BY AN ECB			58 FRAMES					
CLASS	UPPER LIMIT	FREQUENCY OBSERVED	PERCENT OF TOTAL	FREQUENCY DIAGRAM (SCALE = 1650/1)				
0		131970	99.93%	*****				
1		82	0.06%	*				
2		0	0.00%					
3		0	0.00%					
4		0	0.00%					
5		0	0.00%					
6		0	0.00%					
7		0	0.00%					
8		0	0.00%					
9		0	0.00%					
10		0	0.00%					
11		0	0.00%					
12		0	0.00%					
13		0	0.00%					
14		0	0.00%					
15		0	0.00%					
16		0	0.00%					
17		0	0.00%					
18		0	0.00%					
19		0	0.00%					
20		0	0.00%					
25		0	0.00%					
30		0	0.00%					
35		0	0.00%					
40		0	0.00%					
45		0	0.00%					
50		0	0.00%					
60		8	0.01%	*				
70		0	0.00%					
80		0	0.00%					
90		0	0.00%					
100		0	0.00%					
120		0	0.00%					
140		0	0.00%					
160		0	0.00%					
180		0	0.00%					
200		0	0.00%					
220		0	0.00%					
240		0	0.00%					
> 240		0	0.00%					

Figure 15. Heap Area Usage Report

MPIF Configuration Report

This report contains static Multi-Processor Interconnect Facility (MPIF) information. It can be used to determine the general MPIF environment specifications. MPIF-related reports are omitted when MPIF is not active.

MPIF CONFIGURATION	SYSTEM WIDE	27 FEB	14:18:27	SYSTEM	PAGE	9
CRITICAL ACTIVE ECB BLOCKS	:	0				
NUMBER OF PROCESSORS	:	10				
NUMBER OF GLOBAL USERS	:	100				
NUMBER OF RESIDENT USERS	:	24				
NUMBER OF CONNECTIONS	:	90				
NUMBER OF PATHS	:	23				
NUMBER OF PATH ACTIVATION NOTIFICATION	:	90				
NUMBER OF DIRECTORY NOTIFICATION	:	90				
NUMBER OF CLASSES	:	3				
BUFFER SIZE	:	2265088				
INTERFACE VERSION NUMBER	:	3				
NAME OF THE SYSTEM	:	CPUB.JM				
CONNECTION TIMEOUT INTERVAL	:	80				
1ST LEVEL PATH TIMEOUT INTERVAL	:	81				
2ND LEVEL PATH TIMEOUT INTERVAL	:	82				
PDT OUTPUT QUEUE DEPTH	:	10				
NUMBER OF LINKS BETWEEN PROCESSORS	:	10				

Figure 16. MPIF Configuration Report

Interprocessor Communication MPIF Summary Report

This report is generated when MPIF IPC is active. The report gives information about path activity between the origin processor and other processors, and includes mean values per second for the following:

- Total SIPCC items sent
- Total SIPCC items received
- Total SIPCC items returned
- Transmits failed with return
- Transmits failed without return.

INTERPROCESSOR COMMUNICATION MPIF SUMMARY						
ORIGIN PROCESSOR TPF CPU-ID = B						
DEST TPF	TOT SIPCC	TOT SIPCC	TOT SIPCC	XMIT FAIL	XMIT FAIL	
CPU-ID	ITEMS SENT	RECEIVED	RETURN	RETURN	NO RETURN	
B	0.00	0.00	0.00	0.00	0.00	0.00
C	0.00	0.00	0.00	0.00	0.00	0.00
D	0.00	0.00	0.00	0.00	0.00	0.00
E	0.00	0.00	0.00	0.00	0.00	0.00
Z	0.00	0.00	0.00	0.00	0.00	0.00
0	0.00	0.00	0.00	0.00	0.00	0.00

NOTE: COUNTS ARE FROM ORIGIN CPU VIEW POINT
 IE: SENT TO DESTINATION CPU OR RECEIVED FROM DESTINATION CPU
 NOTE: ALL VALUES SHOWN ARE TOTALS PER SECOND

Figure 17. Interprocessor Communication MPIF Report

MPIF Path Activity Report

Performance data for the MPIF paths are contained in the report. This data can be used to detect if a bottleneck exists on a particular path and help to determine the cause. The optional plot reports can help pinpoint heavy utilization periods.

The report is printed in order of path class, and in addition to class, path, and device names, it contains: message rate (messages per second), the average message size, the reads and writes per second, the number of requests on the pending queue, and the number of queue overruns. See Figure 18 for the format of the report.

MPIF PATH ACTIVITY REPORT										27 FEB	14:18:27	SYSTEM	PAGE	10
CTCNAME	PATHNAME	CL	MSGRATE (/SEC)	MSGSIZE (AVG)	READS (/SEC)	WRITES (/SEC)	QUEUED (AVG)	OVERRUNS (TOTAL)						
	CPUB.A1	A	0.00	0.00	0.00	0.00	0	0						
	CPUC.A1	A	0.00	0.00	0.00	0.00	0	0						
	CPUD.A1	A	0.00	0.00	0.00	0.00	0	0						
	CPUE.A1	A	0.00	0.00	0.00	0.00	0	0						
	CPUZ.A1	A	0.00	0.00	0.00	0.00	0	0						
	CPU0.A1	A	0.00	0.00	0.00	0.00	0	0						
	CPU1.A1	A	0.00	0.00	0.00	0.00	0	0						
	CPU2.A1	A	0.00	0.00	0.00	0.00	0	0						
	CPUB.B1	B	0.00	0.00	0.00	0.00	0	0						
	CPUC.B1	B	0.00	0.00	0.00	0.00	0	0						
	CPUD.B1	B	0.00	0.00	0.00	0.00	0	0						
	CPUE.B1	B	0.00	0.00	0.00	0.00	0	0						
	CPUZ.B1	B	0.00	0.00	0.00	0.00	0	0						
	CPU0.B1	B	0.00	0.00	0.00	0.00	0	0						
	CPU1.B1	B	0.00	0.00	0.00	0.00	0	0						
	CPU2.B1	B	0.00	0.00	0.00	0.00	0	0						
	CPUC.C1	C	0.00	0.00	0.00	0.00	0	0						
	CPUD.C1	C	0.00	0.00	0.00	0.00	0	0						
	CPUE.C1	C	0.00	0.00	0.00	0.00	0	0						
	CPUZ.C1	C	0.00	0.00	0.00	0.00	0	0						
	CPU0.C1	C	0.00	0.00	0.00	0.00	0	0						
	CPU1.C1	C	0.00	0.00	0.00	0.00	0	0						
	CPU2.C1	C	0.00	0.00	0.00	0.00	0	0						

Figure 18. MPIF Path Activity Report

Frequency Distribution Reports

Frequency distribution reports can be optioned for all the parameters in the system collector (see Figure 19). Actually, the distribution reports for almost all collected parameters can be obtained by placing the DISTRIBUTION keyword in the option field of any option card.

CREATED ENTRIES PER SECOND		BSS		SUBSYSTEM		
40 OBSERVATIONS		MIN =	1.063	VARIANCE		
		MAX =	16.059	STANDARD D		
		MEAN =	6.007	COEF. OF V		
CLASS	FREQUENCY	PERCENTAGES			MULTIPLE	STD.
UPPER LIMIT	OBSERVED	CLASS	ACCUM	REMAIN	OF MEAN	DEVI
1.85	1	2.50	2.5	97.5	0.30	-0.78
2.64	8	20.00	22.5	77.5	0.43	-0.63
3.43	14	35.00	57.5	42.5	0.57	-0.48
4.22	4	10.00	67.5	32.5	0.70	-0.33
5.01	3	7.50	75.0	25.0	0.83	-0.18
5.80	0	0.00	75.0	25.0	0.96	-0.03
6.59	0	0.00	75.0	25.0	1.09	0.11
7.38	0	0.00	75.0	25.0	1.22	0.26
8.17	0	0.00	75.0	25.0	1.36	0.40
8.96	0	0.00	75.0	25.0	1.49	0.55
9.75	0	0.00	75.0	25.0	1.62	0.70
10.54	0	0.00	75.0	25.0	1.75	0.85
11.33	0	0.00	75.0	25.0	1.88	1.00
12.12	0	0.00	75.0	25.0	2.01	1.15
12.91	0	0.00	75.0	25.0	2.14	1.30
13.70	1	2.50	77.5	22.5	2.28	1.45
14.49	2	5.00	82.5	17.5	2.41	1.60
15.28	2	5.00	87.5	12.5	2.54	1.75
16.07	5	12.50	100.0	0.0	2.67	1.90
16.86	0	0.00	100.0	0.0	2.80	2.05

Figure 19. Frequency Distribution Report

These distribution reports are used to investigate any unusual data such as extreme maximum or minimum values found in the summary. Wide variances indicate that the processing flow through the system is not smooth. Examination of the class

limits and observed frequencies can verify that the system is experiencing wide swings in either input or utilization of one or more of its resources. When such a situation occurs, use caution when comparing one variable to another in order to distinguish cause from effect. Peaks occurring in one variable usually impact some other resource later in the processing cycle. Note that these reports are concerned with frequency distribution only; no chronology of samples is set up.

Plot Reports

To distinguish the cause from effect mentioned previously, a chronological listing is necessary. In addition to the frequency distributions produced, the plot reports show each variable chronologically by interval over the life of the collection (see Figure 20). A separate page is produced for each 100 intervals of the collection period. The clock time associated with the start of each interval is shown along the abscissa.

Variables may now be compared. For example, a very high working storage utilization might have caused polling to be suspended; a high input rate and a long input list appear after polling is resumed. Depending on the order of occurrence, these same variables might indicate an entirely different situation. An extreme peak in high-speed messages can result in a long input list and high core utilization. Any regularity in this type of peaking suggests an adjustment of system parameters that control polling frequency. Messages are allowed to queue in the terminal interchange buffers before being polled into the system.

The proper collection mode, period, and interval must be chosen for the type of analysis suggested previously. Fluctuations may be hidden completely by the smoothing effect of long intervals.

The plot reports are available for all the same parameters as the distribution reports.




Figure 20. Plot Report

Data Reduction Limitations Report

The limitations imposed on the reduction by the PL/I pre-compiler options appear near the end of the reduction report. A sample of this page is shown in Figure 21.

These limits are set by you to match the online system or to restrict the types of data to be reduced. For instance, two versions of the data reduction package could be maintained in the system library, one that handles only system collector data, and the other that reduces data from all collectors. The system collector version would be used more often and would run in a smaller memory partition.

```

THE PRECOMPILER OPTIONS HAVE GIVEN THE TPF 4.1 HPO SYSTEM      REDUCTION PROGRA
COLLECTIONS MADE BY COLLECTOR RELEASES PRIOR TO TPF VERSION 4.1 CANNOT BE REDUCE
THE FOLLOWING COLLECTOR'S OUTPUT CAN BE REDUCED:

                                SYSTEM
                                FILE
                                PROGRAM
                                MESSAGE

THESE PROGRAMS CAN PRODUCE BOTH SUMMARIES AND DETAIL REPORTS OF THEIR PARAMETERS
IN THE SUMMARIES AND DETAIL REPORTS THE FOLLOWING LIMITATIONS ARE IN EFFECT:
THE MAXIMUM NUMBER OF COLLECTION INTERVALS ALLOWED IS      300
THE LOW SPEED MESSAGE WEIGHTING FACTOR IS      5
THE ROUTED MESSAGE WEIGHTING FACTOR IS      0.3
THE MAXIMUM NUMBER OF RANDOM FILES IN THE SYSTEM IS      132
THE MAXIMUM NUMBER OF RECORD ID'S IN THE SYSTEM IS      100
THE MAXIMUM NUMBER OF TAPE UNITS IN THE SYSTEM IS      64
THE MAXIMUM NUMBER OF PROGRAMS IN THE SYSTEM IS      600
THE MAXIMUM NUMBER OF LINES IN THE SYSTEM IS      244
THE MAXIMUM NUMBER OF SNA LINES IN THE SYSTEM IS      255
THE MAXIMUM NUMBER OF NODENAMES IN THE SYSTEM IS      2750
THE MAXIMUM NUMBER OF BSC STATIONS IN THE SYSTEM IS      152
THE MAXIMUM NUMBER OF INTERCHANGES IN THE SYSTEM IS      220
THE MAXIMUM NUMBER OF SUB SYSTEMS IN THE SYSTEM IS      4
THE MAXIMUM NUMBER OF L/C CPU'S IN THE SYSTEM IS      6
THE MAXIMUM NUMBER OF CITIES IN THE SYSTEM IS      35
THE MAXIMUM NUMBER OF TERMINALS IN ONE CITY IS      700
THE MAXIMUM NUMBER OF INTERCHANGES IN ONE CITY IS      30
THE MAXIMUM NUMBER OF APPLICATIONS IN THE NETWORK IS      164

ALL VALID APPLICATION NAMES ARE SHOWN IN THE APPLICATION SUMMARY REPORT.

IF THESE LIMITS ARE EXCEEDED UNDETERMINED ERRORS WILL TAKE PLACE.
THEREFORE CHECK THESE PARAMETERS WITH THE ONLINE SYSTEM AFTER EACH CHANGE.

THE FOLLOWING MESSAGES ARE ERRORS FOUND BY THE REDUCTION PROGRAM, THE OPERATING
EXPLANATION OF THESE MESSAGES MAY BE FOUND IN THE SYSTEM PERFORMANCE AND MEASURE
THE PL/1 PROGRAMMERS GUIDE, THE OS MESSAGE GUIDE, OR THE OS SORT GUIDE.
*****

```

Figure 21. Data Reduction Limitations Report

File Data Analysis

Reports explained under this category include:

- Various file summary reports:
 - Virtual File Access (VFA) Summary Report
 - File Access Summary Report
 - Random File Access Summary Report
 - Channel Path Activity reports
 - Sequential File Access Summary Report
- File Accesses per Record ID Report
- File Accesses per Cylinder Report
- File Comparison Plot Report
- Storage Control reports.

These are each described in this section.

File Summary Reports

These reports summarize all I/O activity except communication I/O over the life of the collection. The total number of DASD accesses, by device, are subdivided into data writes, data reads, and program reads; only writes and character counts are

shown for tapes. All access data is on a per second basis except mean device queue, which is a mean of all interval snapshots taken during the collection. Tape and DASD accesses are totaled by control unit and channel as well as by system.

- The File Access Summary provides
 - Mean Data Reads Per Weighted Message
This number shows average data reads per weighted message, regardless of whether a physical I/O access is required.
 - Mean Data Writes Per Weighted Message
This number shows average data writes per weighted message, including both immediate and delayed file requests. A write request to a duplicated record is counted once, although at least two physical write operations are ultimately initiated.
 - Mean Program Reads Per Weighted Message
This number shows average ENTER requests to VFA per weighted message.

FILE ACCESS SUMMARY	SYSTEM WIDE	27 FEB	14:18:27	FILE	PAGE	1
20 OBSERVATIONS						
	MEAN DATA READS PER WEIGHTED MSG	MEAN DATA WRITES PER WEIGHTED MSG	MEAN PROGRAM READS PER WEIGHTED MSG	TOTAL ACCESSES PER WEIGHTED MSG		
VFA ACCESSES	0.00	0.08	0.00	0.08		
IMMEDIATE VFA FILES		0.08				
TAPE ACCESSES	0.00	0.01	0.00	0.01		
3390 ACCESSES	14.77	3.66	0.00	18.43		
WEIGHTED MESSAGE RATE = (H.S.MSG IN - H.S.MSG ROUTED) + (L.S. * FACT) + (H.S.MSG ROUTED * FACT2)						
= (89.05 - 0.00) + (0.00 * 5.00) + (0.00 * 0.30)						
= 89.05						

Figure 22. File Access Summary Report

More specific information on program reads is found in DASD device statistics in the Record ID Report.

- The VFA summary reports can be used to determine:
 - The efficiency of VFA in reducing the number of physical I/O operations.
When the number of physical I/O operations are excessive there are two tuning possibilities: either expand the VFA area or eliminate the less frequently used record groups from VFA candidacy. (File accesses per record ID should be observed.)
 - VFA Efficiency Statistics.
This report provides systemwide and basic subsystem counts per second of program reads, data reads, FINDs, various kinds of files, and VFA synchronization data.

VFA SUMMARY	SYSTEM WIDE	2 20 APR	08:55:07	FILE	PAGE	2
20 OBSERVATIONS						
VFA EFFICIENCY STATISTICS						
	SYSTEM WIDE	BSS	SUBSYSTEM			
PROGRAM READS	0.00 PER SECOND	0.00	PER SECOND			
DATA READS	34.67 PER SECOND	34.67	PER SECOND			
FINDS (WITH I/O)	35.46 PER SECOND	35.46	PER SECOND			
CANDIDATE FILES	54.00 PER SECOND	54.00	PER SECOND			
FORCE FILES	3.34 PER SECOND	3.34	PER SECOND			
NON-CANDIDATE FILES	1.17 PER SECOND	1.17	PER SECOND			
FILE IMMEDIATES	48.78 PER SECOND	48.78	PER SECOND			
BUFFER UNAVAILABLE	0.00 PER SECOND	0.00	PER SECOND			
381 BUFFER USAGE	1.50 PER SECOND	1.50	PER SECOND			
1055 BUFFER USAGE	1.77 PER SECOND	1.77	PER SECOND			
4K BUFFER USAGE	37.77 PER SECOND	37.77	PER SECOND			
VFA-S LOCK HITS	7.71 PER SECOND	7.71	PER SECOND			
VFA-X LOCK HITS	66.68 PER SECOND	66.68	PER SECOND			
VFA SYNC MISSES	40.74 PER SECOND	40.74	PER SECOND			
VFA SYNC CONTENTIONS	0.00 PER SECOND	0.00	PER SECOND			
RHT CONTENTIONS	0.00 PER SECOND	0.00	PER SECOND			
CACHED RHT HITS	7.77 PER SECOND	7.77	PER SECOND			
VFA SYNC AGE OUTS	0.00 PER SECOND	0.00	PER SECOND			
SIZE OF VFA	6582272 BYTES	(6.6 MEGABYTES)		

Figure 23. VFA Summary Report

- Random File Access Summary Report

This report provides two types of VFA statistics:

- TOT, which gives accesses per second instead of per weighted message. TOT includes all accesses to VFA (including WIO), and
- WIO under device type. WIO includes only those devices that require physical I/O.

TOT refers to total reads (physical and logical). WIO includes only physical reads and is not broken up into program reads and data reads.

The report provides:

- Data Writes

This is the number of physical writes per second to DASD devices, including immediate file requests and page-out file operations that occur when VFA space must be freed for read requests (data reads).

- Data Reads

This is the number of physical reads per second from DASD devices, including physical data reads and physical program reads that result from no-hit in the VFA area.

- Device Queues

This is the maximum and minimum number of outstanding I/O requests in the queue during the data collection period. The FQ record (file queue counts for DASD devices) is used for this report. The FF record (device queues per interval) is used for plotting reports. Therefore, minimum and maximum values printed on the two different reports may not be identical. These physical reads and writes are also counted in the appropriate DASD device reports.

By using this information, you can assess whether:

- The ratios of reads, writes, and program reads are reasonable. The read/write ratio varies depending on how many record types are duplicated; the access counts include duplicate writes to file. The ratio of program reads to total access depends on how much high-speed memory is available for program allocation.

- Accesses are well distributed across all devices, control units, and channels. Abnormal accesses to a single device may be attributable to software, record allocation, message mix, or hardware errors. Access distribution data can also be used to develop operating procedures by which system operators can obtain a proper mix of prime, duplicate, and keypoint modules in control units so that activity is distributed properly across the file subsystem.
- Device, control unit, and channel utilizations are in acceptable limits.
- Channel path usage is evenly distributed across all enabled channel path IDs for a given SDA.

RANDOM FILE ACCESS SUMMARY				BSS		SUBSYSTEM		27 FEB		14:18:27		FILE	PAGE	3
20 OBSERVATIONS														
HDWR	S.M.	SUB	DEVICE	MEAN ACCESSES PER SECOND				SSCH(S)		DEVICE QUEUES		SERVICE		
ADDR	NO.	SYS	TYPE	DATA	DATA	PROG	TOTAL	LOCK	UNLOCK	MAX.	MEAN	LOCK	FDCTC	TIME
				WRITES	READS	READS		ONLY	ONLY			CONTENTION	COMPLETE	(MILS)
VFA	ALL	TOT		7.30	0.00	0.00	7.30	----	----	----	----	----	----	----
VFA	ALL	WIO		7.30	0.00	----	7.30	----	----	----	----	----	----	----
VFA	BSS	TOT		7.30	0.00	0.00	7.30	----	----	----	----	----	----	----
VFA	BSS	WIO		7.30	0.00	----	7.30	----	----	----	----	----	----	----
CHANNEL 14														
CONTROL UNIT 1														
141C	0066	BSS	3390	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
141D	0067	BSS	3390	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
141E	004E	BSS	3390	18.75	66.77	0.00	85.52	0.86	1.47	1.00	0.04	0.00	0.00	1.33
141F	004F	BSS	3390	18.68	66.45	0.00	85.13	0.78	1.28	1.00	0.04	0.00	0.00	1.22
TOTAL CONTROL UNIT 1				37.43	133.22	0.00	170.65	1.64	2.75	----	0.08	0.00	0.00	----
TOTAL CHANNEL 14				37.43	133.22	0.00	170.65	1.64	2.75	----	0.08	0.00	0.00	----
CHANNEL 15														
CONTROL UNIT 1														
151C	0050	BSS	3390	19.01	68.26	0.00	87.27	0.69	1.32	1.00	0.04	0.00	0.00	1.33
151D	005B	BSS	3390	26.63	130.06	0.00	156.69	0.00	1.50	1.00	0.04	0.00	0.00	1.21
TOTAL CONTROL UNIT 1				45.64	198.32	0.00	243.96	0.69	2.82	----	0.08	0.00	0.00	----
TOTAL CHANNEL 15				45.64	198.32	0.00	243.96	0.69	2.82	----	0.08	0.00	0.00	----
:														
CHANNEL 21														
CONTROL UNIT 1														
211C	004A	BSS	3390	18.69	67.19	0.00	85.88	0.70	1.24	1.00	0.04	0.00	0.00	1.27
211D	004B	BSS	3390	18.60	65.91	0.00	84.51	0.78	1.24	0.00	0.00	0.00	0.00	1.20
211E	0068	BSS	3390	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
TOTAL CONTROL UNIT 1				37.29	133.10	0.00	170.39	1.48	2.48	----	0.04	0.00	0.00	----
TOTAL CHANNEL 21				37.29	133.10	0.00	170.39	1.48	2.48	----	0.04	0.00	0.00	----
TOTAL SUBSYSTEM				326.15	1315.71	0.00	1641.86	7.30	20.02	----	0.49	0.00	0.00	----
TOTAL DEVICES REPORTED = 27														
NOTE: TOTAL ACCESSES ON THE RANDOM FILE ACCESS SUMMARY MAY DIFFER FROM OTHER REPORTS.														
1) THE TOTAL ACCESSES ON DETAIL REPORTS = THE TOTAL ACCESSES LISTED ON THIS REPORT + THE FDCTC COUNT LISTED ON THIS REPORT.														
2) THE TOTAL ACCESSES ON CACHE REPORTS = THE TOTAL ACCESSES LISTED ON THIS REPORT + (MULTIPLIER * THE FDCTC COUNT),														
WHERE THE MULTIPLIER IS EQUAL TO THE MEAN NUMBER OF I/O OPERATIONS GENERATED BY AN FDCTC.														

Figure 24. Random File Access Summary Report

The VFA efficiency summary, file access summary, and the random file access report provide a fairly clear picture of VFA activity in main and secondary storage.

Channel Path Reports

The Channel Path Activity reports consists of:

- Channel Path ID (CHPID) Status Exception Report

These are produced when the FILE option is used. See Figure 25.

If a path is logically activated (that is, it exists), it is shown as either ENABLED if the path is available or DISABLED if the path is not available. Only CHPIDs whose status was DISABLED at the beginning or the end of data collection are included in this report. If all CHPIDs are enabled during data collection, no report is produced, and a note is added at the end of the Path Activity Per Channel Path ID Report stating that all CHPIDs were active.

CHANNEL PATH ID STATUS		SYSTEM WIDE	
43 OBSERVATIONS			
CHANNEL PATH ID STATUS EXCEPTION REPORT			
SDA	CHPID	STATUS(START)	STATUS(END)
0123	01	DISABLED	DISABLED
0142	01	DISABLED	DISABLED

Figure 25. Channel Path ID Status Exception Report

- Path Activity per Channel Path ID Report.

This is produced whenever the FILE option is used. See Figure 26.

The following information is reported for each CHPID:

- The number of accesses for that CHPID
- The percentage of the total number of accesses that this represents
- The mean number of accesses per second.

PATH ACTIVITY PER CHANNEL PATH ID			
CHPID	TOTAL ACCESSES	ACCESSES%	MEAN ACCESSES (PER SEC)
25	11	78	3.25
65	3	21	0.77
TOTAL MEAN ACCESSES FOR ALL CHPIDS =			4.02
NOTE: ALL THE CHANNEL PATH ID'S WERE ENABLED DURING DATA COLLECTION			

Figure 26. Path Activity per Channel Path ID Report

- SDA Activity Exception Report

This report is used to analyze discrepancies in channel path activity determined by the SDA percentage cutoff value (SDA_FAC) defined in JPC0. See Figure 27 on page 46.

A report is produced only if the PATHACTIVITY suboption is specified and if the relative difference between access percentages (calculated as the ratio of maximum-minimum/minimum) of any CHPIDs exceeds the percentage specified by SDA_FAC.

SDA ACTIVITY SUMMARY		SYSTEM WIDE		
40 OBSERVATIONS				
SDA ACTIVITY EXCEPTION REPORT				
SDA	CHPID	ACCESSES%	STATUS(START)	STATUS(END)
04E0	25	78	ENABLED	ENABLED
	65	21	ENABLED	ENABLED

Figure 27. SDA Activity Exception Report

Sequential File Access Report

This report (see Figure 28) provides the same kind of information as the Random File Access Report except that sequential devices (tapes) are focused on. The number of tape writes per second, the number of characters written to tape per second, and the device queue are all reported for the various channels and control units in use.

SEQUENTIAL FILE ACCESS SUMMARY			SYSTEM WIDE			27 FEB	14:18:27	FILE	PAGE	10
20 OBSERVATIONS										
HDWR. ADDR.	TAPE NAME	SUBSYSTEM NAME	TAPE WRITES PER SECOND			CHAR TO TAPE PER SECOND			DEVICE QUEUE	
			MAX.	MIN.	MEAN	MAX.	MIN.	MEAN	MAX.	MEAN
CHANNEL F										
CONTROL UNIT 1										
F11	RTA	BSS	0.00	0.00	0.00	0	0	0.00	0	0
F12	RTC	BSS	1.20	1.19	1.20	4945	4928	4935.20	0	0
TOTAL CONTROL UNIT 1			-----	-----	1.20	-----	-----	4935.00	-----	-----
TOTAL CHANNEL F			-----	-----	1.20	-----	-----	4935.00	-----	-----
TOTAL SYSTEM			-----	-----	1.20	-----	-----	4935.00	-----	-----

Figure 28. Sequential File Access Report

Record ID Reports

The Record ID report (see Figure 29) can be produced for the system or by module number. The IDs are listed in alphabetic order after the two characters that represent program accesses. The level of activity against each record ID can be readily observed. The most active record IDs should always be allocated to cylinders near the midpoint of movable head files to reduce seek time. Unusual shifts in record ID activity can often be traced to a change in the input message mix.

FILE ACCESSES PER RECORD ID 20 OBSERVATIONS				BSS	SUBSYSTEM	27 FEB	14:18:27	FILE	PAGE	202	
RECORD I.D.	SYNC	READS	WRITES IMMEDIATE	VFA SUMMARY ACCESS COUNTS			CFWD	RETENTIVE	BYPASS	PERCENT OF TOTAL	ACCESSES PER SECOND
				DELAYED	DFW	CFWS					
X'00E0'		0	1							0.06	0.50
CD,X'C3C4'		0	1							0.06	0.50
E5,X'C5F5'		11	19							1.83	29.51
E6,X'C5F6'		19	10							1.77	28.51
E7,X'C5F7'		18	15							2.01	32.51
F6,X'C6F6'		11	12							1.40	22.51
GL,X'C7D3'		9	15							1.46	23.51
GS,X'C7E2'		1222	161							84.28	1383.11
G7,X'C7F7'		17	10							1.65	26.51
LD,X'D3C4'		18	13							1.89	30.51
OM,X'D6D4'		0	2							0.12	1.50
SD,X'E2C4'		19	7							1.58	25.51
YC,X'E8C3'		0	13							0.79	12.51
YD,X'E8C4'		0	18							1.10	17.51
TOTALS**		1344	297							99.99	1641.22

LEGEND: 'PROGRAM' I.D. = FILE ACCESSES DUE TO ENTERS, AND 'OVERFLOW' I.D. = OVERFLOW FILE ACCESSES DUE TO FINDS/FILES
ALL OTHER I.D.'S PRINTED AS XX,X'XXXX', BOTH CHARACTER AND HEX REPRESENTATIONS
NOTE: DFW = DASD FAST WRITE, CFWS = CACHE FAST WRITE SIMPLEX, CFWD = CACHE FAST WRITE DUPLEX

Figure 29. Record ID Report

File Accesses per Cylinder Report

The Cylinder Analysis Report (see Figure 30) can be produced for SYS, sum, or individual module numbers. SYS produces a report for each DASD device in the system, which can be correlated with the Record ID report to determine whether or not the most active records have been grouped around the center of the module. Arm movement analysis is included in this report. Arm movement of a particular module is calculated by using the following formula:

$$E(X) = P(i) * P(j) * |i-j| \text{ for every } i,j \text{ from } 0 \text{ to } N$$

where:

- P(i) is the probability of being on cylinder i.
- N is the maximum number of cylinders on a module.
- X is a random variable for cylinders moved.
- i and j are cylinder locations on a module.

FILE ACCESSES PER CYLINDER 40 OBSERVATIONS		BSS	SUBSYSTEM
		DEVICE 04E0	MODULE
CYLINDER	ACCESSES		
13	2		
34	1		
47	2		
50	2		
54	1		
61	2		
65	1		
66	1		
147	2		
TOTAL	14		

THE EXPECTED ARM MOVEMENT IS : 37.97 CYLINDERS

Figure 30. Cylinder Analysis Report

File Comparison Plot Report

This report provides a way of comparing the total file access per second for each device type. It provides a concise summary in plot format of the average accesses

per second for all devices in the system by device type. There is a section in this report for DASD devices. System load balancing is more feasible with this information.




Figure 31. File Comparison Plot

Storage Control Reports

The Storage Control reports are primarily concerned with the record caching performed by certain 3880 and 3990 storage controls. The 3880 storage controls are reported on from the point of view of the individual storage director, while the 3990 storage controls are reported on from the point of view of the individual cache subsystem. This is due to differences in storage control organization: each 3880 storage director has its own record cache, but each 3990 cache subsystem has a record cache that is shared by the storage directors in the cache subsystem.

- 3880 Storage Control reports

There are three 3880 storage control reports with essentially the same format. These reports differ with respect to degree of summarization: one contains totals for all of the 3880 storage directors, another contains subtotals for a specific 3880 storage director, and the third contains statistics for a specific DASD backing a specific 3880 storage director.

The data for the 3880 storage control reports is collected from the results of the Sense Subsystem Counts command.

- 3880 Storage Control Report: Totals

This report contains totals for all of the active 3880 storage directors (see Figure 32).

3880 STORAGE CONTROL REPORT: TOTALSCATEGORY		COUNT					
-----	-----	-----					
BUFFER-SUBMODE DATA-TRANSFER-CHAINS	3616						
BUFFER-SUBMODE BYTES-TRANSFERRED	7568						
BUFFER-SUBMODE WAITS	20464						
		381-BYTE SLOTS	1055-BYTE SLOTS	4095			
		-----	-----	-----			
		COUNT	RATE (*)	COUNT	RATE (*)	COUN	
		-----	-----	-----			
NORMAL CACHE REPLACEMENT:		READ REQUESTS	3088	10.22	7184	23.79	
READ HITS	2832	9.38	6928	22.94			
READ HIT-PERCENTAGE		91.71		96.44			
WRITE REQUESTS	2576	8.53	6672	22.09			
WRITE HITS	2320	7.68	6416	21.24			
WRITE HIT-PERCENTAGE		90.06		96.16			
READS-PER-WRITE RATIO		1.20		1.08			
ALL REQUESTS	5664	18.75	13856	45.88	1		
NONRETENTIVE ACCESS:		READ REQUESTS	2064	6.83	6160	20.40	
READ HITS	1808	5.99	5904	19.55			
READ HIT-PERCENTAGE		87.60		95.84			
WRITE REQUESTS	1552	8.53	5648	22.09			
WRITE HITS	1296	4.29	5392	17.85			
WRITE HIT-PERCENTAGE		83.51		95.47			
READS-PER-WRITE RATIO		1.20		1.08			
ALL REQUESTS	3616	11.97	11808	39.10	1		
OTHER:							
INHIBIT-CACHE-LOADING OPERATIONS	1040	3.44	5136	17.01			
RECORD DISCARDS	784	2.60	4880	16.16			
INTERNAL-DASD UPDATE-OPERATIONS	528	1.75	4624	15.31			
INTERNAL-DASD RECORD-UPDATES	272	0.90	4368	14.46			
=====							
NOTE: THE RATE COLUMN ALSO CONTAINS PERCENTAGES AND RATIOS							
NOTE: THESE TOTALS ARE FOR 16 DASD BACKING 4 STORAGE DIRECTORS							

Figure 32. 3880 Storage Control Report: Totals

– 3880 Storage Control Report: Subtotals for SDID nn

This report contains subtotals for a specific active 3880 storage director (see Figure 33).

3880 STORAGE CONTROL REPORT: SUBTOTALS FOR SDID 04					
CATEGORY	COUNT				

BUFFER-SUBMODE DATA-TRANSFER-CHAINS	904				
BUFFER-SUBMODE BYTES-TRANSFERRED	1892				
BUFFER-SUBMODE WAITS	5116				
		381-BYTE SLOTS	1055-BYTE SLOTS	4095	
		-----	-----	-----	
		COUNT	RATE (*)	COUNT	RATE (*)
		-----	-----	-----	-----
NORMAL CACHE REPLACEMENT:					
READ REQUESTS	772	2.56	1796	5.95	
READ HITS	708	2.34	1732	5.73	
READ HIT-PERCENTAGE		91.71		96.44	
WRITE REQUESTS	644	2.13	1668	5.52	
WRITE HITS	580	1.92	1604	5.31	
WRITE HIT-PERCENTAGE		90.06		96.16	
READS-PER-WRITE RATIO		1.20		1.08	
ALL REQUESTS	1416	4.69	3464	11.47	
NONRETENTIVE ACCESS:					
READ REQUESTS	516	1.71	1540	5.10	
READ HITS	452	1.50	1476	4.89	
READ HIT-PERCENTAGE		87.60		95.84	
WRITE REQUESTS	388	2.13	1412	5.52	
WRITE HITS	324	1.07	1348	4.46	
WRITE HIT-PERCENTAGE		83.51		95.47	
READS-PER-WRITE RATIO		1.20		1.08	
ALL REQUESTS	904	2.99	2952	9.77	
OTHER:					
INHIBIT-CACHE-LOADING OPERATIONS	260	0.86	1284	4.25	
RECORD DISCARDS	196	0.65	1220	4.04	
INTERNAL-DASD UPDATE-OPERATIONS	132	0.44	1156	3.83	
INTERNAL-DASD RECORD-UPDATES	68	0.22	1092	3.62	
=====					
NOTE: THE RATE COLUMN ALSO CONTAINS PERCENTAGES AND RATIOS					
NOTE: THESE SUBTOTALS ARE FOR 4 DASD					

Figure 33. 3880 Storage Control Report: Subtotals for SDID nn

- 3880 Storage Control Report for SDA mmmm, SDID nn
This report contains statistics for a specific active DASD backing a specific 3880 storage director (see Figure 34).

3880 STORAGE CONTROL REPORT FOR SDA 0028, SDID 04					
CATEGORY	COUNT				

BUFFER-SUBMODE DATA-TRANSFER-CHAINS	226				
BUFFER-SUBMODE BYTES-TRANSFERRED	473				
BUFFER-SUBMODE WAITS	1279				
	381-BYTE SLOTS		1055-BYTE SLOTS		4095
	COUNT	RATE (*)	COUNT	RATE (*)	COUN

NORMAL CACHE REPLACEMENT:					
READ REQUESTS	193	0.64	449	1.49	
READ HITS	177	0.59	433	1.43	
READ HIT-PERCENTAGE		91.71		96.44	
WRITE REQUESTS	161	0.53	417	1.38	
WRITE HITS	145	0.48	401	1.33	
WRITE HIT-PERCENTAGE		90.06		96.16	
READS-PER-WRITE RATIO		1.20		1.08	
ALL REQUESTS	354	1.17	866	2.87	
NONRETENTIVE ACCESS:					
READ REQUESTS	129	0.43	385	1.27	
READ HITS	113	0.37	369	1.22	
READ HIT-PERCENTAGE		87.60		95.84	
WRITE REQUESTS	97	0.53	353	1.38	
WRITE HITS	81	0.27	337	1.12	
WRITE HIT-PERCENTAGE		83.51		95.47	
READS-PER-WRITE RATIO		1.20		1.08	
ALL REQUESTS	226	0.75	738	2.44	
OTHER:					
INHIBIT-CACHE-LOADING OPERATIONS	65	0.21	321	1.06	
RECORD DISCARDS	49	0.16	305	1.01	
INTERNAL-DASD UPDATE-OPERATIONS	33	0.11	289	0.96	
INTERNAL-DASD RECORD-UPDATES	17	0.06	273	0.90	
=====					
NOTE: SYMBOLIC DEVICE ADDRESS 0028 IS SYMBOLIC MODULE NUMBER 000D IN TPF SUBSYS					
NOTE: THE RATE COLUMN ALSO CONTAINS PERCENTAGES AND RATIOS					

Figure 34. 3880 Storage Control Report for SDA mmmm, SDID nn

- 3990 Storage Control reports

There are five 3990 storage control reports in two categories: full reports and summary reports.

There are three full reports, each several pages long. The first few pages of the full reports all have essentially the same format. They only differ in what they summarize: one contains totals for all of the 3990 cache subsystems, another contains subtotals for a specific 3990 cache subsystem, and the third contains statistics for a specific DASD backing a specific 3990 cache subsystem. The subsequent pages of the reports in this category are composed of different groups of the following types of information:

- Cache capacity information
- Cache shape information
- Cache status information
- Nonvolatile-storage capacity information
- Nonvolatile-storage status information
- DASD status information
- Locking information (when the concurrency filter lock facility (CFLF) is supported).

The two summary reports are about one page long. Both reports have essentially the same format. They also differ to the degree they summarize: one contains a column of totals for all of the 3990 cache subsystems, and a column of subtotals for each 3990 cache subsystem; the other contains a column of totals for all of

the 3990 cache subsystems, a column of subtotals for a specific 3990 cache subsystem, and a column of statistics for each of the DASDs backing the 3990 cache subsystem.

The data for the 3990 Storage Control reports are collected from the results of the Sense Subsystem Status, Read Allocation Data, and Read Subsystem Statistics commands.

3990 Storage Control Report: Totals

This report contains totals for all of the active 3990 cache subsystems (see Figure 35 and Figure 36).

3990 STORAGE CONTROL REPORT: TOTALS (PART 1)									
CATEGORY	TOTALS		381-BYTE SLOTS		COMPLEXWIDE 1055-BYTE SLOTS		27 FEB 14:18:27 4096-BYTE SLOTS		FILE PAGE 205
	COUNT	RATE (*)	COUNT	RATE (*)	COUNT	RATE (*)	COUNT	RATE (*)	
READS PER WRITE		4.02		0.00		4.02		0.00	
READS	603784	1315.43	0	0.00	603784	1315.43	0	0.00	
CFW READS	603784	1315.43	0	0.00	603784	1315.43	0	0.00	
NORMAL READS	0	0.00	0	0.00	0	0.00	0	0.00	
READ HITS	603784	1315.43	0	0.00	603784	1315.43	0	0.00	
CFW READ HITS	603784	1315.43	0	0.00	603784	1315.43	0	0.00	
NORMAL READ HITS	0	0.00	0	0.00	0	0.00	0	0.00	
READ HIT PERCENTAGE		100.00		0.00		100.00		0.00	
CFW READ HITS		100.00		0.00		100.00		0.00	
NORMAL READ HITS		0.00		0.00		0.00		0.00	
WRITES	150041	326.89	0	0.00	150041	326.89	0	0.00	
CFW WRITES	128558	280.08	0	0.00	128558	280.08	0	0.00	
NORMAL WRITES	21483	46.80	0	0.00	21483	46.80	0	0.00	
RETENTIVE WRITES	0	0.00	0	0.00	0	0.00	0	0.00	
DFW WRITES	21483	46.80	0	0.00	21483	46.80	0	0.00	
DFW BYPASSES	0	0.00	0	0.00	0	0.00	0	0.00	
WRITE HITS	150000	326.80	0	0.00	150000	326.80	0	0.00	
CFW WRITE HITS	128558	280.08	0	0.00	128558	280.08	0	0.00	
HIT PERCENTAGE		100.00		0.00		100.00		0.00	
DFW WRITE HITS (NVS)	21442	46.71	0	0.00	21442	46.71	0	0.00	
NVS HIT PERCENTAGE		99.81		0.00		99.81		0.00	
READS AND WRITES	753825	1642.32	0	0.00	753825	1642.32	0	0.00	
DESTAGE OPERATIONS	0	0.00	0	0.00	0	0.00	0	0.00	
DESTAGE PERCENTAGE		0.00		0.00		0.00		0.00	
CAPTURE CHAINS	0	0.00							
CAPTURE CHAIN HITS	0	0.00							
HIT PERCENTAGE		0.00							
RESTORE CHAINS	0	0.00							
TOTAL BYPASS CACHE CHAINS	0	0.00							
BAD READS AND WRITES	0	0.00							
=====									
NOTE: THE RATE COLUMN ALSO CONTAINS PERCENTAGES AND RATIOS									
NOTE: THESE TOTALS ARE FOR 27 DASD BACKING 8 CACHE SUBSYSTEMS									

Figure 35. 3990 Storage Control Report: Totals (Part 1)

3990 STORAGE CONTROL REPORT: TOTALS (PART 2)				COMPLEXWIDE	27 FEB	14:18:27	FILE	PAGE	206
CACHE SHAPE	CURRENT IML	NEXT IML							

SLOTS	SIZE	PERCENT	SIZE	PERCENT					

FULL TRACK		25.00		25.00					
RECORD 1	381	25.00	381	25.00					
RECORD 2	1055	25.00	1055	25.00					
RECORD 3	4096	25.00	4096	25.00					

LOCKING									
CATEGORY	CURRENT IML	NEXT IML	SOE	EOE					

LOCKS ALLOCATED	131072	0							
LOCKS IN USE				5	4				
PERCENTAGE IN USE				*,**	*,**				
RAN OUT OF LOCKS?									
NO									

NOTE: SOE = START OF (DATA COLLECTION) EPISODE, EOE = END OF EPISODE									

Figure 36. 3990 Storage Control Report: Totals (Part 2)

3990 Storage Control Report: Subtotals for Cache Subsystem nnnn

This report contains subtotals for a specific active 3990 cache subsystem (see Figure 37, Figure 38 and Figure 39).

3990 STORAGE CONTROL REPORT: SUBTOTALS FOR CACHE SUBSYSTEM 6207 (PART 1)									
COMPLEXWIDE 27 FEB 14:18:27 FILE PAGE 207									
TOTALS		381-BYTE SLOTS		1055-BYTE SLOTS		4096-BYTE SLOTS			
CATEGORY	COUNT	RATE (*)	COUNT	RATE (*)	COUNT	RATE (*)	COUNT	RATE (*)	
READS PER WRITE		3.52		0.00		3.52		0.00	
READS	59687	130.04	0	0.00	59687	130.04	0	0.00	
CFW READS	59687	130.04	0	0.00	59687	130.04	0	0.00	
NORMAL READS	0	0.00	0	0.00	0	0.00	0	0.00	
READ HITS	59687	130.04	0	0.00	59687	130.04	0	0.00	
CFW READ HITS	59687	130.04	0	0.00	59687	130.04	0	0.00	
NORMAL READ HITS	0	0.00	0	0.00	0	0.00	0	0.00	
READ HIT PERCENTAGE		100.00		0.00		100.00		0.00	
CFW READ HITS		100.00		0.00		100.00		0.00	
NORMAL READ HITS		0.00		0.00		0.00		0.00	
WRITES	16955	36.94	0	0.00	16955	36.94	0	0.00	
CFW WRITES	13221	28.80	0	0.00	13221	28.80	0	0.00	
NORMAL WRITES	3734	8.13	0	0.00	3734	8.13	0	0.00	
RETENTIVE WRITES	0	0.00	0	0.00	0	0.00	0	0.00	
DFW WRITES	3734	8.13	0	0.00	3734	8.13	0	0.00	
DFW BYPASSES	0	0.00	0	0.00	0	0.00	0	0.00	
WRITE HITS	16955	36.94	0	0.00	16955	36.94	0	0.00	
CFW WRITE HITS	13221	28.80	0	0.00	13221	28.80	0	0.00	
HIT PERCENTAGE		100.00		0.00		100.00		0.00	
DFW WRITE HITS (NVS)	3734	8.13	0	0.00	3734	8.13	0	0.00	
NVS HIT PERCENTAGE		100.00		0.00		100.00		0.00	
READS AND WRITES	76642	166.98	0	0.00	76642	166.98	0	0.00	
DESTAGE OPERATIONS	0	0.00	0	0.00	0	0.00	0	0.00	
DESTAGE PERCENTAGE		0.00		0.00		0.00		0.00	
CAPTURE CHAINS	0	0.00							
CAPTURE CHAIN HITS	0	0.00							
HIT PERCENTAGE		0.00							
RESTORE CHAINS	0	0.00							
TOTAL BYPASS CACHE CHAINS	0	0.00							
BAD READS AND WRITES	0	0.00							

=====

NOTE: THE RATE COLUMN ALSO CONTAINS PERCENTAGES AND RATIOS
NOTE: THESE SUBTOTALS ARE FOR 3 DASD

Figure 37. 3990 Storage Control Report: Subtotals for Cache Subsystem nnnn (Part 1)

CATEGORY	SOE	EOE	PROTOTYPE
CACHING ACTIVE	YES	YES	000.
CACHE COMING ONLINE	NO	NO	001.
CACHING FAILED	NO	NO	010.
UNDEFINED	NO	NO	011.
CACHING DEACTIVATED	NO	NO	100.
CACHING SUSPENDED	NO	NO	101.
DESTAGING IN PROGRESS	NO	NO	110.
DESTAGING FAILED	NO	NO	111.
CACHE DISABLED FOR MAINTENANCE	NO	NO	...1
UNDEFINED	NO	NO	... 1...
UNDEFINED	NO	NO1..
IML DEVICE UNAVAILABLE	NO	NO1.
CACHE-FAST-WRITE DEACTIVATED	NO	NO1

NVS STATUS

CATEGORY	SOE	EOE	PROTOTYPE
NVS ACTIVE	YES	YES	00..
NVS FAILED	NO	NO	01..
NVS UNAVAILABLE	NO	NO	10..
NVS PENDING	NO	NO	11..
UNDEFINED	NO	NO	.1.
NVS DISABLED FOR MAINTENANCE	NO	NO	...1
NVS PENDING DUE TO ERROR	NO	NO	... 1...
UNDEFINED	NO	NO1..
UNDEFINED	NO	NO1.
UNDEFINED	NO	NO1

CACHE CAPACITY		START OF EPISODE		END OF EPISODE	
CATEGORY	IML AMOUNT	AMOUNT	PERCENT	AMOUNT	PERCENT
CONFIGURED	1073741824				
AVAILABLE		1069498368	99.60	1069498368	99.60
PINNED		0	0.00	0	0.00
OFFLINE		0	0.00	0	0.00

NVS CAPACITY		START OF EPISODE		END OF EPISODE	
CATEGORY	IML AMOUNT	AMOUNT	PERCENT	AMOUNT	PERCENT
CONFIGURED	16777216				
PINNED		0	0.00	0	0.00

NOTE: SOE = START OF (DATA COLLECTION) EPISODE, EOE = END OF EPISODE

Figure 38. 3990 Storage Control Report: Subtotals for Cache Subsystem mmmm (Part 2)

3990 STORAGE CONTROL REPORT: SUBTOTALS FOR CACHE SUBSYSTEM 6207 (PART 3)				COMPLEXWIDE	27 FEB	14:18:27	FILE	PAGE	209
CACHE SHAPE	CURRENT IML	NEXT IML							
SLOTS	SIZE	PERCENT	SIZE	PERCENT					
FULL TRACK		25.00		25.00					
RECORD 1	381	25.00	381	25.00					
RECORD 2	1055	25.00	1055	25.00					
RECORD 3	4096	25.00	4096	25.00					
LOCKING									
CATEGORY	CURRENT IML	NEXT IML	SOE	EOE	HWM				
LOCKS ALLOCATED	16384	0							
LOCKS IN USE			1	0	0				
PERCENTAGE IN USE			0.01	0.00	0.00				
RAN OUT OF LOCKS?	NO								
=====									
NOTE: SOE = START OF (DATA COLLECTION) EPISODE, EOE = END OF EPISODE, HWM = HIGHWATER MARK DURING EPISODE									

Figure 39. 3990 Storage Control Report: Subtotals for Cache Subsystem mmmm (Part 3)

3990 Storage Control Report for SDA mmmm, Cache Subsystem nnnn

This report contains subtotals for a specific SDA on an active 3990 cache subsystem (see Figure 40 and Figure 41 on page 58)

3990 STORAGE CONTROL REPORT FOR SDA 0231, CACHE SUBSYSTEM 00FE (PART 1)					
CATEGORY	TOTALS		381-BYTE SLOTS		1055-BY
	COUNT	RATE (*)	COUNT	RATE (*)	COUNT
READS PER WRITE		1.00		0.69	
READS	24733	83.00	206	0.69	
CFW READS	24733	83.00	206	0.69	
NORMAL READS	0	0.00	0	0.00	
READ HITS	24721	82.96	206	0.69	
CFW READ HITS	24721	82.96	206	0.69	
NORMAL READ HITS	0	0.00	0	0.00	
READ HIT PERCENTAGE		99.95		100.00	
CFW READ HITS		99.95		100.00	
NORMAL READ HITS		0.00		0.00	
WRITES	24793	83.20	297	1.00	
CFW WRITES	24432	81.99	0	0.00	
NORMAL WRITES	361	1.21	297	1.00	
RETENTIVE WRITES	361	1.21	297	1.00	
DFW WRITES	0	0.00	0	0.00	
DFW BYPASSES	0	0.00	0	0.00	
WRITE HITS	24432	81.99	0	0.00	
CFW WRITE HITS	24432	81.99	0	0.00	
HIT PERCENTAGE		100.00		0.00	
DFW WRITE HITS (NVS)	0	0.00	0	0.00	
NVS HIT PERCENTAGE		0.00		0.00	
READS AND WRITES	49526	166.19	503	1.69	
DESTAGE OPERATIONS	0	0.00	0	0.00	
DESTAGE PERCENTAGE		0.00		0.00	
CAPTURE CHAINS	0	0.00			
CAPTURE CHAIN HITS	0	0.00			
HIT PERCENTAGE		0.00			
RESTORE CHAINS	0	0.00			
TOTAL BYPASS CACHE CHAINS	0	0.00			
BAD READS AND WRITES	0	0.00			
=====					
NOTE: SYMBOLIC DEVICE ADDRESS 0231 IS SYMBOLIC MODULE NUMBER 0047 IN TPF SUBSYS					
NOTE: THE RATE COLUMN ALSO CONTAINS PERCENTAGES AND RATIOS					

Figure 40. 3990 Storage Control Report for SDA nnnn, Cache Subsystem mmmm (Part 1)

3990 STORAGE CONTROL REPORT FOR SDA 0231, CACHE SUBSYSTEM 00FE (PART 2)

DASD STATUS

CATEGORY	SOE	EOE	PROTOTYPE
CACHING ACTIVE	YES	YES	00..
UNDEFINED	NO	NO	01..
DEACTIVATE PENDING	NO	NO	10..
CACHING DEACTIVATED	NO	NO	11..
DASD-FAST-WRITE ALLOWED	YES	YES	..00
UNDEFINED	NO	NO	..01
DEACTIVATE PENDING	NO	NO	..10
DASD-FAST-WRITE DEACTIVATED	NO	NO	..11
PRIMARY OF A DUPLEX PAIR	NO	NO 1... .
SECONDARY OF A DUPLEX PAIR	NO	NO1... .
DUPLEX PAIR SUSPENDED	NO	NO **1. .
DUPLEX PAIR AVAILABLE	NO	NO **00 .
DUPLEX PAIR PENDING	NO	NO **01 .
PINNED DATA EXISTS	NO	NO 1

NOTE: SYMBOLIC DEVICE ADDRESS 0231 IS SYMBOLIC MODULE NUMBER 0047 IN TPF SUBSYS
 NOTE: SOE = START OF (DATA COLLECTION) EPISODE, EOE = END OF EPISODE

Figure 41. 3990 Storage Control Report for SDA nnnn, Cache Subsystem mmmm (Part 2)

3990 Storage Control Cache Summary Report

This report contains summarized totals for all of the 3990 cache subsystems and summarized subtotals for each of the active 3990 cache subsystems (see Figure 42).

3990 STORAGE CONTROL CACHE SUMMARY REPORT					COMPLEXWIDE		27 FEB	14:18:27	FILE	PAGE	203
CATEGORY	TOTALS	SSID 6207	SSID 6208	SSID 6209	SSID 620B	SSID 620C	SSID 620D	SSID 620E			
NUMBER OF CACHE SUBSYSTEMS	8										
NUMBER OF DASD	27	3	3	4	4	3	4	4			
NUMBER OF LOCKS ALLOCATED	131072	16384	16384	16384	16384	16384	16384	16384			
RAN OUT OF LOCKS?											
NO	NO	NO	NO	NO	NO	NO	NO	NO			
SOE: CACHE CONFIGURED CAPACITY		1073741824	536870912	536870912	1073741824	536870912	536870912	536870912			
PERCENTAGE PINNED		0.00	0.00	0.00	0.00	0.00	0.00	0.00			
PERCENTAGE OFFLINE		0.00	0.00	0.00	0.00	0.00	0.00	0.00			
NVS CONFIGURED CAPACITY		16777216	16777216	16777216	16777216	16777216	16777216	16777216			
PERCENTAGE PINNED		0.00	0.00	0.00	0.00	0.00	0.00	0.00			
LOCKS IN USE	5	1	2	0	2	0	0	0			
PERCENTAGE IN USE	*.**	0.01	0.01	0.00	0.01	0.00	0.00	0.00			
EOE: CACHE CONFIGURED CAPACITY		1073741824	536870912	536870912	1073741824	536870912	536870912	536870912			
PERCENTAGE PINNED		0.00	0.00	0.00	0.00	0.00	0.00	0.00			
PERCENTAGE OFFLINE		0.00	0.00	0.00	0.00	0.00	0.00	0.00			
NVS CONFIGURED CAPACITY		16777216	16777216	16777216	16777216	16777216	16777216	16777216			
PERCENTAGE PINNED		0.00	0.00	0.00	0.00	0.00	0.00	0.00			
LOCKS IN USE	4	0	0	0	2	0	0	0			
PERCENTAGE IN USE	*.**	0.00	0.00	0.00	0.01	0.00	0.00	0.00			
HWM: LOCKS IN USE		0	0	0	0	0	0	0			
PERCENTAGE IN USE		0.00	0.00	0.00	0.00	0.00	0.00	0.00			
READS PER WRITE	4.02	3.52	3.56	3.55	4.62	4.33	3.51	4.36			
READS	1315.43	130.04	133.45	133.33	260.57	130.04	132.65	197.56			
CFW READS	1315.43	130.04	133.45	133.33	260.57	130.04	132.65	197.56			
NORMAL READS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00			
READ HITS	1315.43	130.04	133.45	133.33	260.57	130.04	132.65	197.56			
CFW READ HITS	1315.43	130.04	133.45	133.33	260.57	130.04	132.65	197.56			
NORMAL READ HITS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00			
READ HIT PERCENTAGE	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00			
CFW READ HITS	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00			
NORMAL READ HITS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00			
WRITES	326.89	36.94	37.52	37.52	56.40	30.04	37.79	45.35			
CFW WRITES	280.08	28.80	29.64	29.64	52.93	26.85	29.91	41.25			
NORMAL WRITES	46.80	8.13	7.88	7.87	3.47	3.19	7.88	4.10			
RETENTIVE WRITES	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00			
DFW WRITES	46.80	8.13	7.88	7.87	3.47	3.19	7.88	4.10			
DFW BYPASSES	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00			
WRITE HITS	326.80	36.94	37.52	37.51	56.38	30.03	37.78	45.33			
CFW WRITE HITS	280.08	28.80	29.64	29.64	52.93	26.85	29.91	41.25			
HIT PERCENTAGE	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00			
DFW WRITE HITS (NVS)	46.71	8.13	7.88	7.87	3.45	3.18	7.86	4.08			
NVS HIT PERCENTAGE	99.81	100.00	99.97	99.97	99.37	99.52	99.83	99.47			
READS AND WRITES	1642.32	166.98	170.98	170.84	316.98	160.09	170.44	242.91			
DESTAGE OPERATIONS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00			
DESTAGE PERCENTAGE	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00			
CAPTURE CHAINS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00			
CAPTURE CHAIN HITS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00			
HIT PERCENTAGE	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00			
RESTORE CHAINS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00			
TOTAL BYPASS CACHE CHAINS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00			
BAD READS AND WRITES	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00			

NOTE: SOE = START OF (DATA COLLECTION) EPISODE, EOE = END OF EPISODE, HWM = HIGHWATER MARK DURING EPISODE
NOTE: CFW = CACHE FAST WRITE, DFW = DASD FAST WRITE, NVS = NONVOLATILE STORAGE

Figure 42. 3990 Storage Control Cache Summary Report

3990 Storage Control Device Summary Report for Cache Subsystem

This report contains summarized totals for all 3990 cache subsystems, summarized subtotals for a specific active 3990 cache subsystem, and summarized statistics for each of the active DASDs backing the 3990 cache subsystem (see Figure 43).

3990 STORAGE CONTROL DEVICE SUMMARY REPORT FOR CACHE SUBSYSTEM 6207					COMPLEXWIDE	27 FEB	14:18:27	FILE	PAGE	231
CATEGORY	TOTALS	SSID 6207	SDA 201C	SDA 201D	SDA 201E					
NUMBER OF CACHE SUBSYSTEMS	8									
NUMBER OF DASD	27	3								
TPF SUBSYSTEM		BSS	BSS	BSS						
NUMBER OF LOCKS ALLOCATED	131072	16384								
RAN OUT OF LOCKS?										
NO	NO									
SOE: CACHE CONFIGURED CAPACITY		1073741824								
PERCENTAGE PINNED		0.00								
PERCENTAGE OFFLINE		0.00								
NVS CONFIGURED CAPACITY		16777216								
PERCENTAGE PINNED		0.00								
LOCKS IN USE	5	1								
PERCENTAGE IN USE	*.**	0.01								
EOE: CACHE CONFIGURED CAPACITY		1073741824								
PERCENTAGE PINNED		0.00								
PERCENTAGE OFFLINE		0.00								
NVS CONFIGURED CAPACITY		16777216								
PERCENTAGE PINNED		0.00								
LOCKS IN USE	4	0								
PERCENTAGE IN USE	*.**	0.00								
HWM: LOCKS IN USE		0								
PERCENTAGE IN USE		0.00								
READS PER WRITE	4.02	3.52	3.51	0.00	3.53					
READS	1315.43	130.04	63.71	0.00	66.32					
CFW READS	1315.43	130.04	63.71	0.00	66.32					
NORMAL READS	0.00	0.00	0.00	0.00	0.00					
READ HITS	1315.43	130.04	63.71	0.00	66.32					
CFW READ HITS	1315.43	130.04	63.71	0.00	66.32					
NORMAL READ HITS	0.00	0.00	0.00	0.00	0.00					
READ HIT PERCENTAGE	100.00	100.00	100.00	0.00	100.00					
CFW READ HITS	100.00	100.00	100.00	0.00	100.00					
NORMAL READ HITS	0.00	0.00	0.00	0.00	0.00					
WRITES	326.89	36.94	18.16	0.00	18.78					
CFW WRITES	280.08	28.80	14.18	0.00	14.63					
NORMAL WRITES	46.80	8.13	3.98	0.00	4.16					
RETENTIVE WRITES	0.00	0.00	0.00	0.00	0.00					
DFW WRITES	46.80	8.13	3.98	0.00	4.16					
DFW BYPASSES	0.00	0.00	0.00	0.00	0.00					
WRITE HITS	326.80	36.94	18.16	0.00	18.78					
CFW WRITE HITS	280.08	28.80	14.18	0.00	14.63					
HIT PERCENTAGE	100.00	100.00	100.00	0.00	100.00					
DFW WRITE HITS (NVS)	46.71	8.13	3.98	0.00	4.16					
NVS HIT PERCENTAGE	99.81	100.00	100.00	0.00	100.00					
READS AND WRITES	1642.32	166.98	81.87	0.00	85.10					
DESTAGE OPERATIONS	0.00	0.00	0.00	0.00	0.00					
DESTAGE PERCENTAGE	0.00	0.00	0.00	0.00	0.00					
CAPTURE CHAINS	0.00	0.00	0.00	0.00	0.00					
CAPTURE CHAIN HITS	0.00	0.00	0.00	0.00	0.00					
HIT PERCENTAGE	0.00	0.00	0.00	0.00	0.00					
RESTORE CHAINS	0.00	0.00	0.00	0.00	0.00					
TOTAL BYPASS CACHE CHAINS	0.00	0.00	0.00	0.00	0.00					
BAD READS AND WRITES	0.00	0.00	0.00	0.00	0.00					

NOTE: SOE = START OF (DATA COLLECTION) EPISODE, EOE = END OF EPISODE, HWM = HIGHWATER MARK DURING EPISODE
NOTE: CFW = CACHE FAST WRITE, DFW = DASD FAST WRITE, NVS = NONVOLATILE STORAGE

Figure 43. 3990 Storage Control Device Summary Report for Cache Subsystem mmmm

DP Record Anomaly Log

There is an error log that reports any anomalies detected in the data collected for the storage control reports (see Figure 44).

```

***** START OF DP RECORD ANOMALY LOG
***** START OF DP RECORD ANOMALY LOG
***** START OF DP RECORD ANOMALY LOG
THE FOLLOWING DASD HAD DUPLICATE START-OF-EPIISODE DATA:
 0003 0004 0005
THE FOLLOWING DASD HAD DUPLICATE END-OF-EPIISODE DATA:
 0004 0005 0006 0007 0008 0013 0015 0062
THE FOLLOWING DASD WERE APPARENTLY ONLINE AT THE END OF DATA COLLECTION BUT OFFL
 0006 0008 0013 0015 0062
THE FOLLOWING DASD WERE APPARENTLY ONLINE AT THE START OF DATA COLLECTION BUT OF
 0003
THE FOLLOWING DASD WERE APPARENTLY RESET DURING DATA COLLECTION:
 0004 0005 0007 000D 001D 002B 0033 0037 0038 003C 0046 004F 0051
 00CB 00DE 00E0 00E2 00F2 00F7
THE FOLLOWING DASD HAD DIFFERENT START- AND END-OF-EPIISODE CONFIGURED-CACHE-CAPA
 0004 0005 0007
THE FOLLOWING DASD HAD DIFFERENT START- AND END-OF-EPIISODE CONFIGURED-NONVOLATIL
 0004 0007
THE FOLLOWING DASD HAD DIFFERENT START- AND END-OF-EPIISODE RECORD SIZES:
 0004 0005 0007
THE FOLLOWING DASD HAD DIFFERENT START- AND END-OF-EPIISODE TRACK/RECORD PERCENTA
 0004 0005 0007
THE FOLLOWING DASD WERE BACKING INCONSISTENT STORAGE CONTROLS:
 0003 0004 0005 0006 0008 0013 0015 0062
THE FOLLOWING DASD HAD DIFFERENT START- AND END-OF-EPIISODE LOCKING CHARACTERISTI
 0074 0075 0077
***** END OF DP RECORD ANOMALY LOG
***** END OF DP RECORD ANOMALY LOG
***** END OF DP RECORD ANOMALY LOG

```

Figure 44. DP Record Anomaly Log

See the publications listed in “Related Information” on page xiv for definitions of the terms used in these reports and for more information about the 3990 cache control units.

Program Data Analysis

Program Reports

A TPF system analyst has a number of uses for the program reports:

- One use for the program reports is to track program usage over time as an indirect gauge of the growth of the load on the system.

Every Enter does not represent an identical load on the system. Different entry points may have different instruction path lengths. Such differences tend to cancel one another out when aggregates of programs are considered instead of individual programs.

To assist the study of aggregates of programs, the OPTIONS file allows definition of packages of programs through program naming patterns.

- Another use of the program reports is to compare packages of programs. The purpose of this comparative study is to roughly gauge the relative share of packages on system load. Alternatively, differences in the design and implementation of the respective packages can be roughly gauged.

For example, the TPF analyst may want to investigate the reasons underlying an observation that a package of programs has a comparatively high mean nesting level but a comparatively low percentage of ENTNCs and ENTDCs in favor of ENTRCs. One explanation may be that classical structured programming guidelines were followed too literally for the TPF environment. A high mean nesting level but a low percentage of ENTNCs or ENTDCs can mean that the package is holding on to more nested system resources than necessary and that the package is placing an unnecessarily large load on the system.

The TPF analyst can use the program reports to look for significant differences between packages and between subsystems, and then try to account for those differences.

Program Names and Enters Report

The Program Names and Enters Report shows which of the programs of a subsystem actually received control while the program data was being collected (see Figure 45).

PROGRAM NAMES AND ENTERS		PROGRAM NAMES AND ENTERS		PROGRAM NAMES AND ENTERS		BASIC	SUB	SY
PROGRAM NAMES AND ENTERS		PROGRAM NAMES AND ENTERS		PROGRAM NAMES AND ENTERS		BSS	SUBSYSTEM	
PROG	CAT	#ENTERS	PROG	CAT	#ENTERS	PROG	CAT	#ENTE
====	---	-----	====	---	-----	====	---	-----
CEL5	F31	0						
CLEN	F31	1						
CLER	F31	0						
CLEZ	F31	18						
COAB	C31	1						
COLI	F31	0						
CSA0	C31	0						
CSA7	C31	2						
CSHA	C31	0						
CSHB	C31	2						
CSHC	C31	1						
CSHD	C31	2						
CVIQ	C31	1						
CYC3	C31	1						
JCD2	F31	1						
JCF0	F31	300						
XLDD	C31	1						
XLII	C31	1						

Figure 45. Program Names and Enters Report

The report identifies the programs that received control as a result of Enter or Back events. An Enter event is the execution of an ENTRC, ENTNC, or ENTDC macro. A Back event is the execution of a BACKC macro.

Program Details Report

The Program Details Report provides a detailed view of the Enter/Back activity of those subsystem programs that actually received control (see Figure 46).

PROGRAM DETAILS								BSS	SUBSYSTEM	
PROG CAT	#EVENTS	#ENTERS	#RETURNS	%DLMFC	%ENTRC	%ENTNC	%ENTDC	%DLMR	%BACKC	%F
====	----	-----	-----	-----	-----	-----	-----	-----	-----	-----
ARPK C31	74	74	0	0.00	0.00	100.00	0.00	0.00	0.00	
ARPT C31	146	71	75	0.00	0.00	100.00	0.00	0.00	100.00	
CAPB C31	4	4	0	0.00	0.00	100.00	0.00	0.00	0.00	
CAPP C31	17	8	9	0.00	0.00	100.00	0.00	0.00	100.00	
CAPQ C31	9	9	0	0.00	100.00	0.00	0.00	0.00	0.00	
CAPR C31	9	9	0	0.00	0.00	100.00	0.00	0.00	0.00	
CBC3 C31	2	2	0	0.00	0.00	100.00	0.00	0.00	0.00	
CBLO C31	11	5	6	0.00	0.00	100.00	0.00	0.00	100.00	
CBPK C31	2	2	0	0.00	100.00	0.00	0.00	0.00	0.00	
CBW0 C31	18	18	0	0.00	33.33	66.67	0.00	0.00	0.00	
CDTF C31	20	20	0	0.00	0.00	100.00	0.00	0.00	0.00	
CELM F31	72	72	0	0.00	100.00	0.00	0.00	0.00	0.00	
CELT F31	178	178	0	0.00	100.00	0.00	0.00	0.00	0.00	
CELV F31	177	76	101	0.00	100.00	0.00	0.00	0.00	100.00	
CELW F31	250	26	224	0.00	100.00	0.00	0.00	0.00	100.00	
CEL5 F31	418	23	395	0.00	0.00	100.00	0.00	0.00	100.00	
CHKA C31	4	4	0	0.00	100.00	0.00	0.00	0.00	0.00	
CIAA C31	4	4	0	0.00	0.00	100.00	0.00	0.00	0.00	
CINN C31	4	4	0	0.00	0.00	100.00	0.00	0.00	0.00	
CLDA F31	26	26	0	0.00	100.00	0.00	0.00	0.00	0.00	
CLEG F31	26	26	0	0.00	100.00	0.00	0.00	0.00	0.00	
CLEK F31	48	48	0	0.00	100.00	0.00	0.00	0.00	0.00	
CLEM F31	1667	705	962	0.00	100.00	0.00	0.00	0.00	100.00	
CLEN F31	257	257	0	0.00	100.00	0.00	0.00	0.00	0.00	
CLER F31	2236	26	2210	0.00	100.00	0.00	0.00	0.00	100.00	
CLES F31	62	62	0	0.00	100.00	0.00	0.00	0.00	0.00	
CLET F31	20	10	10	0.00	100.00	0.00	0.00	0.00	100.00	2
:										
:										
CYYM C31	74	74	0	0.00	100.00	0.00	0.00	0.00	0.00	
JCD2 F31	120	40	80	0.00	0.00	100.00	0.00	0.00	100.00	
JCD3 F31	40	40	0	0.00	100.00	0.00	0.00	0.00	0.00	1
JCP0 F31	40	40	0	0.00	100.00	0.00	0.00	0.00	0.00	
:										
:										
UOP1 C31	36	36	0	0.00	100.00	0.00	0.00	0.00	0.00	
WGR1 C31	3	3	0	0.00	100.00	0.00	0.00	0.00	0.00	
XLAA C31	32	32	0	0.00	0.00	100.00	0.00	0.00	0.00	
XLII C31	35	35	0	0.00	0.00	100.00	0.00	0.00	0.00	
XLJJ C31	40	40	0	0.00	0.00	100.00	0.00	0.00	0.00	
TOTALS:	75274	45265	30009	13.46	62.16	17.38	7.00	23.84	76.16	

Figure 46. Program Details Report

When no packages of programs are defined in the OPTIONS file, the report generated encompasses all of the programs in the subsystem being reduced.

Cumulative Program Enters Report

The Cumulative Program Enters report provides a concise view of the programs that received control most often as a result of Enter events (see Figure 47). The report stops identifying programs when the cumulative percentage cutoff is exceeded. The default cutoff is 90%.

When no packages of programs are defined in the OPTIONS file, the report generated encompasses all of the programs in the subsystem being reduced.

CUMULATIVE PROGRAM ENTERS FOR PACKAGE: XXXXXXXX					BASIC SUB SY
CUMULATIVE PROGRAM ENTERS					BSS SUBSYSTEM
PROG CAT	#ENTERS	ENTERS RATE	ESREL%	ESCUM%	
----	----	-----	-----	-----	*****
JCF0 F31	300	104.91	90.36	90.36	
CLEZ F31	18	6.13	5.42	95.78	
CSA7 C31	2	0.53	0.60	96.38	
CSHB C31	2	0.53	0.60	96.99	
CSHD C31	2	0.53	0.60	97.59	
CLEN F31	1	0.18	0.30	97.89	
COAB C31	1	0.18	0.30	98.19	
CSHC C31	1	0.18	0.30	98.49	
CVIQ C31	1	0.18	0.30	98.79	
CYC3 C31	1	0.18	0.30	99.10	
JCD2 F31	1	0.18	0.30	99.40	
XLDD C31	1	0.18	0.30	99.70	
XLII C31	1	0.18	0.30	100.00	
CEL5 F31	0	0.00	0.00	100.00	
CLER F31	0	0.00	0.00	100.00	
COLI F31	0	0.00	0.00	100.00	
CSA0 C31	0	0.00	0.00	100.00	
CSHA C31	0	0.00	0.00	100.00	
TOTALS:	332	116.12			

Figure 47. Cumulative Program Enters Report

Cumulative Program On-File Enters Report

The Cumulative Program On-File Enters report provides a concise view of the programs that received control most often as a result of Enter events that caused program retrieval references to VFA or to DASD (see Figure 48). The report stops identifying programs when the cumulative percentage cutoff is exceeded. The default cutoff is 90%.

When no packages of programs were defined in the OPTIONS file, the report is generated once for all of the programs in the subsystem being reduced.

CUMULATIVE PROGRAM ON-FILE ENTERS					BSS SUBSYSTEM
PROG CAT	#FILE	FILE RATE	FSREL%	FSCUM%	
----	----	-----	-----	-----	*****
CEL5 F31	0	0.00	0.00	0.00	
CLEN F31	0	0.00	0.00	0.00	
CLER F31	0	0.00	0.00	0.00	
CLEZ F31	0	0.00	0.00	0.00	
COAB C31	0	0.00	0.00	0.00	
COLI F31	0	0.00	0.00	0.00	
CSA0 C31	0	0.00	0.00	0.00	
CSA7 C31	0	0.00	0.00	0.00	
CSHA C31	0	0.00	0.00	0.00	
CSHB C31	0	0.00	0.00	0.00	
CSHC C31	0	0.00	0.00	0.00	
CSHD C31	0	0.00	0.00	0.00	
CVIQ C31	0	0.00	0.00	0.00	
CYC3 C31	0	0.00	0.00	0.00	
JCD2 F31	0	0.00	0.00	0.00	
JCF0 F31	0	0.00	0.00	0.00	
XLDD C31	0	0.00	0.00	0.00	
XLII C31	0	0.00	0.00	0.00	
TOTALS:	0	0.00			

Figure 48. Cumulative Program On-File Enters Report

Package Details Report

The Package Details report provides a detailed view of the Enter/Back activity of the packages of programs defined in the OPTIONS file for the subsystem (see Figure 49).

PACKAGE DETAILS								BSS	SUBSYSTEM	
PACKAGE	#EVENTS	#ENTERS	#RETURNS	%DLMFC	%ENTRC	%ENTNC	%ENTDC	%DLMR	%BACKC	%F
A	47	31	16	0.00	0.00	100.00	0.00	0.00	100.00	
B	0	0	0	0.00	0.00	0.00	0.00	0.00	0.00	
C	7098	3908	3190	0.00	66.63	33.29	0.08	0.00	100.00	
D	0	0	0	0.00	0.00	0.00	0.00	0.00	0.00	
E	0	0	0	0.00	0.00	0.00	0.00	0.00	0.00	
F	0	0	0	0.00	0.00	0.00	0.00	0.00	0.00	
G	0	0	0	0.00	0.00	0.00	0.00	0.00	0.00	
H	0	0	0	0.00	0.00	0.00	0.00	0.00	0.00	
I	0	0	0	0.00	0.00	0.00	0.00	0.00	0.00	
J	97	52	45	0.00	88.46	11.54	0.00	0.00	100.00	
K	0	0	0	0.00	0.00	0.00	0.00	0.00	0.00	
L	0	0	0	0.00	0.00	0.00	0.00	0.00	0.00	
M	0	0	0	0.00	0.00	0.00	0.00	0.00	0.00	
N	0	0	0	0.00	0.00	0.00	0.00	0.00	0.00	
O	0	0	0	0.00	0.00	0.00	0.00	0.00	0.00	
P	0	0	0	0.00	0.00	0.00	0.00	0.00	0.00	
Q	244543	157041	87502	19.87	53.63	19.48	9.76	35.66	64.34	
R	79803	57238	22565	22.38	53.08	6.69	15.71	56.76	43.24	
S	15584	7882	7702	9.63	89.71	3.39	0.00	9.85	90.14	
T	0	0	0	0.00	0.00	0.00	0.00	0.00	0.00	
U	200	200	0	0.00	100.00	0.00	0.00	0.00	0.00	
V	0	0	0	0.00	0.00	0.00	0.00	0.00	0.00	
W	0	0	0	0.00	0.00	0.00	0.00	0.00	0.00	
X	20	20	0	0.00	0.00	100.00	0.00	0.00	0.00	
Y	0	0	0	0.00	0.00	0.00	0.00	0.00	0.00	
Z	0	0	0	0.00	0.00	0.00	0.00	0.00	0.00	

Figure 49. Package Details Report

Cumulative Package Enters Report

The Cumulative Package Enters report provides a concise view of the packages that received control most often as a result of Enter events (see Figure 50).

CUMULATIVE PACKAGE	PACKAGE #ENTERS	ENTERS RATE	ESREL%	ESUM%	BSS	SUBSYSTEM
J	301	105.26	90.66	90.66		
C	29	9.98	8.73	99.40		
X	2	0.53	0.60	100.00		
A	0	0.00	0.00	100.00		
B	0	0.00	0.00	100.00		
D	0	0.00	0.00	100.00		
E	0	0.00	0.00	100.00		
F	0	0.00	0.00	100.00		
G	0	0.00	0.00	100.00		
H	0	0.00	0.00	100.00		
I	0	0.00	0.00	100.00		
K	0	0.00	0.00	100.00		
L	0	0.00	0.00	100.00		
M	0	0.00	0.00	100.00		
N	0	0.00	0.00	100.00		
O	0	0.00	0.00	100.00		
P	0	0.00	0.00	100.00		
Q	0	0.00	0.00	100.00		
R	0	0.00	0.00	100.00		
S	0	0.00	0.00	100.00		
T	0	0.00	0.00	100.00		
U	0	0.00	0.00	100.00		
V	0	0.00	0.00	100.00		
W	0	0.00	0.00	100.00		
Y	0	0.00	0.00	100.00		
Z	0	0.00	0.00	100.00		

Figure 50. Cumulative Package Enters Report

Cumulative Package On-File Enters Report

The Cumulative Package On-File Enters report provides a concise view of the packages that received control most often as a result of Enter events that caused program-retrieval references to VFA or to DASD (see Figure 51).

CUMULATIVE PACKAGE	PACKAGE #FILE	ON-FILE FILE RATE	ENTERS FSREL%	FSCUM%	BSS	SUBSYSTEM
-----	=====	=====	=====	*****		
A	0	0.00	0.00	0.00		
B	0	0.00	0.00	0.00		
C	0	0.00	0.00	0.00		
D	0	0.00	0.00	0.00		
E	0	0.00	0.00	0.00		
F	0	0.00	0.00	0.00		
G	0	0.00	0.00	0.00		
H	0	0.00	0.00	0.00		
I	0	0.00	0.00	0.00		
J	0	0.00	0.00	0.00		
K	0	0.00	0.00	0.00		
L	0	0.00	0.00	0.00		
M	0	0.00	0.00	0.00		
N	0	0.00	0.00	0.00		
O	0	0.00	0.00	0.00		
P	0	0.00	0.00	0.00		
Q	0	0.00	0.00	0.00		
R	0	0.00	0.00	0.00		
S	0	0.00	0.00	0.00		
T	0	0.00	0.00	0.00		
U	0	0.00	0.00	0.00		
V	0	0.00	0.00	0.00		
W	0	0.00	0.00	0.00		
X	0	0.00	0.00	0.00		
Y	0	0.00	0.00	0.00		
Z	0	0.00	0.00	0.00		

Figure 51. Cumulative Package On-File Enters Report

Program Legend Report

The Program Legend report defines the column headings used in most of the program reports (see Figure 52).

PROGRAM LEGEND COLUMN	DESCRIPTION	BSS	SUBSYSTEM
=====	-----		
#COME	THE NUMBER OF ENTERS THAT HAD TO AWAIT INCOMING I/O IN ORDER TO RE		
#ENTERS	THE NUMBER OF GENERIC ENTERS		
#EVENTS	THE NUMBER OF GENERIC ENTERS OR GENERIC RETURNS		
#FILE	THE NUMBER OF ENTERS THAT HAD TO ACCESS THE VFA OR DASD IN ORDER T		
#IN	THE NUMBER OF ENTERS THAT FOUND THE PROGRAM BEING ENTERED ALREADY		
#OTHER	THE NUMBER OF ENTERS THAT DID NOT HAVE TO ACCESS THE VFA OR DASD I		
#RETURNS	THE NUMBER OF GENERIC RETURNS		
%BACKCS	THE PERCENTAGE OF RETURNS THAT WERE BACKCS		
%COME	THE PERCENTAGE OF ENTERS THAT HAD TO AWAIT INCOMING I/O IN ORDER T		
%DLMFC	THE PERCENTAGE OF ENTERS THAT WERE DLM EXTERNAL FUNCTION CALLS		
%DLMR	THE PERCENTAGE OF RETURNS THAT WERE DLM RETURNS		
%ENTDC	THE PERCENTAGE OF ENTERS THAT WERE ENTDCS		
%ENTNC	THE PERCENTAGE OF ENTERS THAT WERE ENTNCS		
%ENTRC	THE PERCENTAGE OF ENTERS THAT WERE ENTRCS		
%FILE	THE PERCENTAGE OF ENTERS THAT HAD TO ACCESS THE VFA OR DASD IN ORD		
%I1 %I2 ...	THE PERCENTAGE OF GENERIC ENTERS OR GENERIC RETURNS THAT OCCURRED		
%IN	THE PERCENTAGE OF ENTERS THAT FOUND THE PROGRAM BEING ENTERED ALRE		
%OTHER	THE PERCENTAGE OF ENTERS THAT DID NOT HAVE TO ACCESS THE VFA OR DA		
CAT	THE PROGRAM CATEGORIES:		
	*** - INDETERMINATE PROGRAMS		
	C24 - CORE RESIDENT, 24-BIT PROGRAMS		
	C31 - CORE RESIDENT, 31-BIT PROGRAMS		
	F24 - FILE RESIDENT, 24-BIT PROGRAMS		
	F31 - FILE RESIDENT, 31-BIT PROGRAMS		
	X24 - CORE RESIDENT CONVERTED TO FILE RESIDENT, 24-BIT PROGRAMS		
	X31 - CORE RESIDENT CONVERTED TO FILE RESIDENT, 31-BIT PROGRAMS		
ENTERS RATE	THE NUMBER OF GENERIC ENTERS PER SECOND		
ESCUM%	THE CUMULATIVE PERCENTAGE OF THE SUBSYSTEM'S ENTERS		
ESREL%	THE RELATIVE PERCENTAGE OF THE SUBSYSTEM'S ENTERS		
FILE RATE	THE NUMBER OF ENTERS, THAT HAD TO ACCESS THE VFA OR DASD IN ORDER		
FSCUM%	THE CUMULATIVE PERCENTAGE OF THE SUBSYSTEMS'S ENTERS THAT HAD TO A		
FSREL%	THE RELATIVE PERCENTAGE OF THE SUBSYSTEMS'S ENTERS THAT HAD TO AC		
PACKAGE	THE NAME OF THE PACKAGE OF PROGRAMS		
PROG	THE NAME (OR NUMBER) OF THE PROGRAM ('OTHER' IS A SPECIAL PROGRAM-		
SUBSYS	THE NAME OF THE SUBSYSTEM		

Figure 52. Program Legend Report

Weighted-Message Rate Report

The Weighted-Message Rate Report (see Figure 53) shows you how many Enter events occurred per second, per weighted message.

WEIGHTED-MESSAGE RATE	BSS	SUBSYSTEM	04 FEB	14:44:57	PROGRAM	PAGE	2
THERE WERE	6.12	ENTERS PER WEIGHTED MESSAGE.					
#WEIGHTED-MESSAGES =	(#HIGHSPEED - #ROUTED) + (5.00 * #LOWSPEED) + (0.30 * #ROUTED) + #TCP/IP	=	897.92	WEIGHTED-MESSAGES.	

Figure 53. Weighted-Message Rate Report

Message Data Analysis

Message Summary Report

Message Summary reports are produced according to your options (see Figure 54). The system-wide message summary report presents basic local network line

statistics calculated from captured TPF counters. The application-based message summary report presents local network line statistics for each specified application, calculated from intercepted input/output messages. This report is related only to the high-speed local network lines including SDLC cross-domain links, and BSC/SLC links (except for SDLC and ALC lines of cross-domain resources which are reported if the associated NCP is shared by other CPUs).

The systemwide message summary is the first part of the Message Summary Report. Each pseudo line is reported by the actual line, that is, by the name given to the physical line. For the Network Extension Facility (NEF), the SNA node name given to each ALC line is used instead of the pseudo ALC lines. The Message Summary by Line Report may not be meaningful for a NEF network. The application based message summary does not report message-routing traffic to and from SDLC cross-domain links, except for application-to-application messages which are reported in the CPU section of the Message Summary Report. Messages of cross-domain sessions — those between applications of the TPF system and the cross-domain resources (including primary and secondary applications LUs) — are reported in the associated cross-domain link section. Because NEF resources are considered cross-domain resources, NEF messages are reported in the cross-domain link section.

The message rates are calculated from messages recorded during the message collector interval. Comparing these values to those reported in the System Summary Report should provide close correlation for the means.

The high-speed output message queue refers to a snapshot sample of the total number of messages which are in queue for output to the 2703, (3705EP), and 3705 NCP lines.

Inordinately high line queues are indications of line problems or overloading and should be investigated further. Line loading or utilization can be readily calculated by using 400 and 800 characters per record to represent 100% utilization of a 2400 and 4800 baud line, respectively.

MESSAGE SUMMARY		SYSTEM WIDE			
LINE NO/ NODENAME	TYPE	INPUT MSG/SEC	MEAN INPUT MSG LENGTH	INPUT CHARS/SEC	OUTPUT MSGS /INPUT MSGS
3E	BYSC PP	0.00	0.00	0.00	0.00
3F	BYSC PP	0.00	0.00	0.00	0.00
40	BYSC PP	0.00	0.00	0.00	0.00
41	BYSC MP	0.00	0.00	0.00	0.00
	STB1	0.00	0.00	0.00	0.00
	STB3	0.00	0.00	0.00	0.00
	ASSS	0.00	0.00	0.00	0.00
	BIRD	0.00	0.00	0.00	0.00
	COWS	0.00	0.00	0.00	0.00
	CROW	0.00	0.00	0.00	0.00
	DOGS	0.00	0.00	0.00	0.00
	EELS	0.00	0.00	0.00	0.00
	FAWN	0.00	0.00	0.00	0.00
	GOAT	0.00	0.00	0.00	0.00
	LAMB	0.00	0.00	0.00	0.00
	OWLS	0.00	0.00	0.00	0.00
	SEAL	0.00	0.00	0.00	0.00
42	BYSC MP	0.00	0.00	0.00	0.00
	STB4	0.00	0.00	0.00	0.00
:					
49	BYSC PP	0.00	0.00	0.00	0.00
4A	BYSC MP	0.00	0.00	0.00	0.00
	STD5	0.00	0.00	0.00	0.00
	STB2	0.00	0.00	0.00	0.00
	ASSS	0.00	0.00	0.00	0.00
	BIRD	0.00	0.00	0.00	0.00
	COWS	0.00	0.00	0.00	0.00
	CROW	0.00	0.00	0.00	0.00
	DOGS	0.00	0.00	0.00	0.00
	EELS	0.00	0.00	0.00	0.00
	FAWN	0.00	0.00	0.00	0.00
	GOAT	0.00	0.00	0.00	0.00
	LAMB	0.00	0.00	0.00	0.00
	OWLS	0.00	0.00	0.00	0.00
	SEAL	0.00	0.00	0.00	0.00
4B	BYSC PP	0.00	0.00	0.00	0.00
4C	BYSC PP	0.00	0.00	0.00	0.00
4D	BYSC PP	0.00	0.00	0.00	0.00
	LINK 1	0.00	0.00	0.00	0.00
34	AI	0.00	0.00	0.00	0.00
	LINK 2	0.00	0.00	0.00	0.00
35	AI	0.00	0.00	0.00	0.00
36	AI	0.00	0.00	0.00	0.00
37	AI	0.00	0.00	0.00	0.00
38	AI	0.00	0.00	0.00	0.00
	LINK 3	0.00	0.00	0.00	0.00
39	AI	0.00	0.00	0.00	0.00
3A	AI	0.00	0.00	0.00	0.00
3B	AI	0.00	0.00	0.00	0.00
3C	AI	0.00	0.00	0.00	0.00

Figure 54. Message Summary Report

Action Code Summary Reports

System and File Summaries include only the number of messages on a per second basis. To establish a base from which data collection comparisons can be made, the number of input messages must be related to a message mix: that is, messages grouped by user defined transaction category.

When the message collector is run in continuous mode, the mean output lengths and mean existence time of each action group is also available. The existence time is from the time a message is placed on the input list until the SEND macro is issued. The sample reports represent the total system. The Action Code Summary (see Figure 55) may also be selected by line, city, or particular terminal interchange.

Because NEF terminal addresses (LEIDs have no relationship to lines, cities, or terminal interchanges), the Action Code Summary Report based on line, city, or terminal interchange is not meaningful.

ACTION CODE SUMMARY					
AAAA ACTION GROUP	TOTAL INPUT MESSAGES	% OF TOTAL	MSGs/ SEC	MEA LE	
AM	45	100.00	0.06	2	
TOTALS	45	100.00	0.06	2	
UNMATCHED OUTPUT	1				
% OF LINE'S TOTAL OUTPUT	2.50				
MEAN OUTPUT LENGTH	24.00				

'MEAN OUTPUT LENGTH' REFERS TO OUTPUT MESSAGES AS SEEN BY THE AGENT (MAY BE COM

Figure 55. Action Code Summary Report

Application Summary and Detail Reports

Two reports related to applications are produced when the TPF system is generated with the Message Routing Facility. The Application Summary report is included in the System Summary (see Figure 5) section and is always produced. The Application Detail Report, when selected in the option cards, is produced to provide additional information (see Figure 56).

The Application Summary Report is used to analyze messages from the cross-domain network (whether they are through cross-domain SLU-PLU session or via Message Routing facility) to applications in the TPF system, and messages from local-domain resources to applications in other CPUs (through the message routing facility). The Application Detail Report breaks down, by each application in the network, the number of input messages for each application, the percentage of the total number for the system, the number of messages per second, and the mean input message length. The Application Detail report is printed only if the message routing facility is system-generated and message collection was active. If message collection was done in continuous mode, this report provides mean existence time for each application. In addition to local-domain messages, this report reflects cross-domain session input messages. System application (SA) records are used as input to this report.

APPLICATION	TOTAL INPUT MESSAGES	SYSTEM WIDE % OF TOTAL	10 MEAS/SEC	MEAS/LE
AAAA				
BALU				
BBBB				
B320				
CCCC				
CISB				
CLGA				
CLGB				
CLGC				
CLGD				
CXFT				
CXJA				
CXJB				
DDDD				
ECHO				
EEEE				
EEE1				
FC14				
FFFF				
FFF1				
GAMB				
GGGG				
GGG1				
GMC1				
GMC2	27	12.75		0.17
GMC3	15	7.12		0.09
GMT1	14	6.62		0.08
GMT2				
GMT3				
GMT4	13	6.12		0.08
GMT5				
GMT6				
GMT7				
GMT8				
GMZ1	61	29.00		0.39
GMZ2	80	38.00		0.51
GMZ3				
GMZ4				
HHHH				
HHH1				
:				
VVVV				
WWW				
XXXX				
YYYY				
ZZZZ				
32NA				
32NB				
32NC				
32ND				
TOTALS	210	100.00		1.34

Figure 56. Application Detail Report

City Summary Report

The City Summary Report (see Figure 57) lists input and output messages by cities served by the reservations network. Activity on remote terminals that are served by local line drops off a terminal interchange is reported as part of the traffic originating from the city code assigned to the terminal interchange. This report is also broken down by optioned application. This report includes messages of pseudo 3270 terminals and terminals of the Network Extension Facility (NEF) in the cross-domain network.

CITY SUMMARY		SYSTEM WIDE			03 APR	10:36:26	MESSAGE
		INPUT MESSAGES			OUTPUT MESSAGE		
CITY	MSG COUNT	% OF TOTAL	MSGs/ SEC	MEAN INP LENGTH	MSG COUNT	% OF TOTAL	MSGs/ SEC
LAX	9	100.00	0.06	7.000	9	100.00	0.06
SLC							
STL							
DEN							
ORD							
PHX							
COS							
APPLICATION RES0							
		INPUT MESSAGES			OUTPUT MESSAGES		
CITY	MSG COUNT	% OF TOTAL	MSGs/ SEC	MEAN INP LENGTH	MSG COUNT	% OF TOTAL	MSGs/ SEC
LAX	125	26.65	0.06	15.00	210	35.59	0.10
COS							
DEN	227	48.40	0.11	11.18	227	38.47	0.11
ABO							
ORD							
PHX							
MKC							
EIP							
SLC							
STL							
TUS							

Figure 57. City Summary Report

Terminal Activity Report

The Terminal Activity Report (see Figure 58) is optioned by city and application. Individual terminal activity for all terminals associated with a city code is grouped in a convenient one-page report. The total counts of input and output messages by LNIA are also shown. At the bottom of the page the totals of message counts for the city are listed along with the percentages of the total system traffic that the counts represent. This report includes terminals of the Network Extension Facility (NEF) in the cross-domain network. Note that pseudo 3270 terminals are reported in the Logical Unit (LU) Activity Report. The Terminal Activity Report optioned by city is not meaningful for NEF terminal networks.

TERMINAL ACTIVITY		TPF/DSD Development				22 May 13:21:39		MESSAGEHOU	
	LNIA	0402	LNIA	0503	LNIA	0906	LNIA	0B01	
TA	INP	OUT	INP	OUT	INP	OUT	INP	OUT	
02			6	1					
10	2	1	3	3	4	3			
11	4	3	9	10					
12	1	1			1	1			
13	11	11	1	1			4	4	
14					1	1	6	6	
15									
16	9	10	10	10			3	3	
17	16	15	6	6	1	1			
18	2	2	3	4					
19	3	3			1	1			
1A	10	9	1	1	3		6	7	
1B							1	1	
20	2	2					3	3	
22			4	4					
23	6	7							
25	8	9							
26			1	1					
28	5	5					14	14	
2A							7	7	
TOTAL	79	78	44	41	12	8	44	45	
TOTAL INP MESSAGES = 179									
PERCENT OF TOTAL INP= 8.16									
TOTAL OUTP MESSAGES = 172									
PERCENT OF TOTAL OUT= 7.65									

Figure 58. Terminal Activity Report (Part 1 of 2)

TERMINAL ACTIVITY	TPF Development		18 Oct	14:16:02	MESSAGE	
LAX						
RES0	LNIA 0201	LNIA 0202	LNIA 0203	LNIA 0204	LNIA 0205	LNIA 0206
TA	INP	OUT	INP	OUT	INP	OUT
TOTAL	ALL VALUES ARE ZERO					
TERMINAL ACTIVITY	TPF Development		18 Oct	14:16:02	MESSAGE	
(CONT'D)						
	LNIA 0301	LNIA 0302	LNIA 0303	LNIA 0304	LNIA 0305	LNIA 0306
TA	INP	OUT	INP	OUT	INP	OUT
02			1			
05						
06	13	20				
TOTAL	13	20	1			
TERMINAL ACTIVITY	TPF Development		18 Oct	14:16:02	MESSAGE	
(CONT'D)						
	LNIA 0309	LNIA 030B	LNIA 0701	LNIA 0708	LNIA 0709	LNIA 0804
TA	INP	OUT	INP	OUT	INP	OUT
TOTAL	ALL VALUES ARE ZERO					
TERMINAL ACTIVITY	TPF Development		18 Oct 14	:16:02	MESSAGE	
(CONT'D)						
	LNIA 0C01	LNIA 0C02	LNIA 0C03	LNIA 0C04	LNIA 1804	
TA	INP	OUT	INP	OUT	INP	OUT
TOTAL	ALL VALUES ARE ZERO					
TOTAL INP MESSAGES =	125					
% OF INP OF THIS CITY	81.16					
PERCENT OF TOTAL INP=	26.65					
TOTAL OUTP MESSAGES =	210					
% OF OUT OF THIS CITY	87.86					
PERCENT OF TOTAL OUT=	35.59					

Figure 58. Terminal Activity Report (Part 2 of 2)

Logical Unit Activity Report

The Logical Unit (LU) Activity Report is optioned by SNA line name (node name) and application (see Figure 59). The individual activity for all logical units associated with a line is grouped by cluster controller and its related logical units on this report. The total counts of input and output messages by cluster are also shown. At the bottom of the page the totals of message counts for the line are listed along with the percentages of the total system-wide SNA traffic that the counts represent. Pseudo 3270 terminals are included in this report. To obtain activities of cross-domain logical units, specify the proper cross-domain link in the option cards.

LU ACTIVITY		TPF Development		03 APR		12:16:32		MESSAGE			
WILDCARD SPECIFIED: LN2C*											
LU NAME	INP	OUT	LU NAME	INP	OUT	LU NAME	INP	OUT	LU NAME	INP	OUT LU
LN2C1411	0	0	LN2C1412	0	0	LN2C1413	37	31	LN2C1414	0	0 LN2C1
LN2C1417	0	0	LN2C1418	0	0	LN2C1421	0	0	LN2C1422	0	0 LN2C1
LN2C1425	40	60	LN2C1426	0	0	LN2C1427	0	0	LN2C1428	0	0 LN2C1
LN2C1433	0	0	LN2C1434	0	0	LN2C1435	0	0	LN2C1436	0	0 LN2C1
LN2C1440	0	0	LN2C1441	62	48	LN2C1442	0	0	LN2C1443	0	0 LN2C1
LN2C1446	0	0	LN2C1447	0	0	LN2C1448	0	0	LN2C1450	0	0 LN2C1
LN2C1453	0	0	LN2C1454	0	0	LN2C1455	0	0	LN2C1456	0	0 LN2C1
LN2C1460	0	0	LN2C1461	0	0	LN2C1462	0	0	LN2C1463	0	0 LN2C1
LN2C1466	0	0	LN2C1467	0	0	LN2C1468	0	0			
TOTAL INP MESSAGES = 139											
TOTAL OUTP MESSAGES = 139											

Figure 59. Logical Unit Activity Report

Logical Unit Message Stream Report

This report is optioned by SNA line name (node name) (see Figure 60). Only one line is permitted for each run against the reduction. If both the SNA line and city Message Stream reports are optioned, only one will be generated while the other is ignored. The normal report will print the total text of the input messages and as many as 65 characters of the output messages unless the suboption to print the entire output message has been chosen. The report is arranged by nodename so that all messages from a specific LU are listed in sequence before beginning the printout of messages from the next LU. This report, when optioned by city, may not be meaningful for NEF terminal networks.

To obtain the message stream of pseudo 3270 terminals, a SNA line node name must be chosen.

The TOD clock is listed for each message so that the time between successive message entries from an LU can be analyzed. The existence time approximates the length of time required to process the message in the CPU. The existence time is not printed for some messages when collection is made in sampling mode because input or output messages may not be available at the beginning and end of intervals. This section explains the variables that are affected by the online system environment and by the collectors' data recording methods, and offline by the unique constants of the reduction programs, assumptions, and data manipulation techniques. The aim of this section is to provide those details of collection and reduction that may not be readily apparent to the user, but are necessary for a serious analysis of one's system based on the System Performance reports. The I-stream number will be reported for output messages.

Note: A logical unit (LU) resource is included in the LU Message Stream Report only if a resource vector table (RVT) entry exists for that LU resource when data collection is started. If an RVT entry is created for an LU resource after data collection is started, that LU resource is not included in this report.

MESSAGE STREAM BY AGENT			TPF Development	MESSAGE	
LN	IA	TA	MESSAGE TYPE	MESSAGE TEXT	TIM
06	02	10			SIN
			INPUT	*R	2
			OUTPUT	/\$ 1. 1PPAAAAA +	
				1 CO 165Y 23OCT ABQELP SS1 327P 413P	
				FONE-ABQABQ0000406	
				TLT-TL0	
				HA FAX-OSI HAFAX	
				GEN FAX-OSI FAX	
				▶+	
			INPUT	0A	
			OUTPUT	/F 2 ARNK +	
				▶+	
			INPUT	0166Y19NOVDENMKCNN1	
			OUTPUT	/G 3 166Y 19NOV DENMKC SS1 830P 952P +	
				▶+	
			INPUT	0169Y19DECMKCDENNN1	
			OUTPUT	/ +	
				RPT ENT	
				▶+	
			INPUT	0A	
			OUTPUT	/H 4 ARNK +	
				▶+	
			INPUT	0173Y18JANDALABQNN1	
			OUTPUT	/ +	
				INVLD DATE	
				▶+	
			INPUT	0174Y17FEBABQDALNN1	
			OUTPUT	/ +	
				INVLD DATE	
				▶+	
			INPUT	*1	
			OUTPUT	/\$ 1 CO 165Y 23OCT ABQELP SS1 327P 413 +	
				2 ARNK	
				3 CO 166Y 19NOV DENMKC SS1 830P 952P	
				4 ARNK	
				▶+	

Figure 61. Message Stream Report by Agent

Distributed Data Management (DDM) Message Stream Report

This report is printed if the SNA STREAM DDM option is requested (see Figure 62). One report is printed for each SLU (that is, node) found to be sending or receiving DDM message traffic. The output DDM messages and the input DDM messages are listed in chronological order. The remote response time is calculated for any input message that immediately follows an output message. It should be noted that DDM messages are paired as output/input message pairs and that there may be multiple responses (input messages) to a single request (output message).

See *TPF Application Requester User's Guide* for more information about interpreting the data collection report for the TPF Application Requester (TPFAR) feature.

Note: A logical unit (LU) resource is included in the DDM Message Stream Report only if a resource vector table (RVT) entry exists for that LU resource when

data collection is started. If an RVT entry is created for an LU resource after data collection is started, that LU resource is not included in this report.

DDM MESSAGE NODENAME	STREAM REPORT MESSAGE TYPE	MESSAGE TEXT	SYSTEM WIDE
CTCCC003			
SESSION 1	CONVERSATION 1		
	OUTPUT	0048D041 00000042 10410014 14041403 00031444 00032407 0003240F 00030004 115D0008 1147D8E3 D7C60015 116D40C4 C1D5C2E4 D9E8E3D7 C6D5C5E3 4040C300 09115AE3 D7C6F3F1 0049D001 00010043 20010016 2110C4C2 F2F3E3E2 E3404040 40404040 40404040 0006210F 2407000D 002FD8E3 C4E2D8D3 F3F7F000 0A003500 06119C01 F4000C11 2EE3D7C6 F0F3F0F1 F0	
	INPUT	0056D043 00000050 14430010 115EC4C2 C1C1F0F2 F7C3F5F5 C4F80014 14041403 00031444 00032407 0003240F 00030008 1147D8C4 C2F20014 116DC4C2 F2F3E3E2 E3404040 40404040 4040000C 115AC4E2 D5F0F2F0 F3F0004C D0020001 00462201 00061149 0000000D 002FD8E3 C4E2D8D3 F3F7F000 0C112EC4 E2D5F0F2 F0F3F000 0A003500 06119C01 F4001921 35E3D7C6 D5C5E34B C3E3C3C3 C3F0F0F3 A3F87E4A 5554	
	OUTPUT	0069D051 00020063 200B0016 2110C4C2 F2F3E3E2 E3404040 40404040 40404040 00442113 C4C2F2F3 E3E2E340 40404040 40404040 4040E3D7 C6D5C5E3 6DC2E2E2 6DC8D7D5 40404040 D8E7D5F0 40404040 40404040 40404040 40401477 8BE210BD 10100001 00052111 F00013D0 53000200 0D002FD8 E3C4E2D8 D3F3F7F0 0010D003 0002000A 00350006 119C01F4	
	INPUT	000BD003 00020005 2408FF	
	OUTPUT	0069D051 00030063 200B0016 2110C4C2 F2F3E3E2 E3404040 40404040 40404040 00442113 C4C2F2F3 E3E2E340 40404040 40404040 4040E3D7 C6D5C5E3 6DC2E2E2 6DC8D7D5 40404040 D8E7D5F0 40404040 40404040 40404040 40401477 8BE210BD 10100002 00052111 F00013D0 53000300 0D002FD8 E3C4E2D8 D3F3F7F0 0010D003 0003000A 00350006 119C01F4	
	INPUT	000BD003 00030005 2408FF	

Figure 62. Distributed Data Management (DDM) Message Stream Report

TCP/IP Message Summary Report

This report includes information about the number of messages, bytes, and packets that are sent and received by each TCP/IP native stack application. The TCP/IP message summary report includes the following information:

- **Application**
The name of the application. The application must be defined in the TCP/IP network services database. If an application is not defined in the database, the number of messages is shown in the OTHER category. See *TPF Transmission Control Protocol/Internet Protocol* for more information about defining applications in the TCP/IP network services database.
- **Port**
The port associated with this application.
- **Weight**
The weighting factor used when counting the input messages for this application. This weighting factor can be defined for each application in the TCP/IP network services database by specifying the weight parameter. If a weighting factor is not

defined for the application, this column contains asterisks (***) . See *TPF Transmission Control Protocol/Internet Protocol* for more information about defining applications in the TCP/IP network services database and how message counts are incremented for TCP/IP applications.

- **Input Msgs/Sec**
The number of messages that were received by this application each second. These are raw message counts; that is, they are not weighted.
- **Input Pkts/Sec**
The number of packets containing data that were received by this application each second.
- **Input Bytes/Sec**
The number of bytes that were received by this application each second.
- **Output Msgs/Sec**
The number of messages that were sent by this application each second. These are raw message counts; that is, they are not weighted.
- **Output Pkts/Sec**
The number of packets containing data that were sent by this application each second.
- **Output Bytes/Sec**
The number of bytes that were sent by this application each second.

The information in this report is shown in descending order by input messages per second; that is, the application with the highest number of input messages per second is shown at the top.

The following shows an example of a TCP/IP message summary report.

TCP/IP MESSAGE SUMMARY			04 FEB	14:44:57	MESSAGE	PAGE	1	
APPLICATION	PORT	WEIGHT	MSG/SEC	INPUT PKTS/SEC	BYTES/SEC	MSG/SEC	OUTPUT PKTS/SEC	BYTES/SEC
TEST-9981	9981	***	149.60	149.60	149598	149.60	149.60	149598
TEST-9980	9980	***	149.49	149.49	149492	149.49	149.49	149492
TEST-9984	9984	***	149.45	149.45	149451	149.45	149.45	149451
TEST-9982	9982	***	149.44	149.44	149436	149.44	149.44	149436
TEST-9983	9983	***	139.69	139.69	139686	139.69	139.69	139686
TEST-9985	9985	***	139.66	139.66	139658	139.66	139.66	139658
OTHER		***	0.09	0.09	39	0.11	0.11	0
RIP	520	50	0.07	0.07	4	0.07	0.07	1
FTP-DATA	20	100	0.04	8.76	11883	0.00	1.39	0
TFTP	69	100	0.02	0.02	0	0.00	0.00	0
TOTAL			877.55	886.28	889251	877.50	878.89	877326

Figure 63. TCP/IP Message Summary Report

TPF MQSeries Data Collection Reports

MQSeries data collection provides data to effectively manage the performance of both channels and queues.

The TPF system collects data for each channel defined in MQSeries, including both sender and receiver channels. The channel data consists of the number of messages and the number of bytes being processed by each channel. This is for data arriving at the TPF system and for data being sent by the TPF system to remote MQSeries servers. The report provides mean message and data rates as well as the maximum message and data rates. See Figure 64 on page 80 for an example of the TPF MQSeries message channel report.

The TPF system also collects data showing activity for the queue of each type defined. The report includes the following information:

- The average number of times per second that a queue is opened and closed
- The maximum number of times that a queue is opened and closed
- The average number of times that messages are added to each queue
- The maximum number of times that messages are added to each queue.

Note: The number of times the queue is opened and closed is not shown for transmission queues.

See Figure 65 on page 81 for an example of the TPF MQSeries message queue report.

```

TPF MQ*SERIES MESSAGE - CHANNEL REPORT          04 MAR      14:29:43      MESSAGE PAGE 1
BSS SUBSYSTEM
=====
MESSAGE CHANNELS FOR LOCAL QUEUE MANAGER: TPFQM          , 5 OBSERVATIONS, 24 SECONDS, CPU B
=====
MESSAGE CHANNEL NAME  TYPE          MESSAGE RATE      TOTAL DATA RATE (BYTES)  INCOMING DATA RATE      OUTGOING DATA RATE
MEAN                 MAXIMUM          MEAN              MAXIMUM                 MEAN                     MAXIMUM
-----
CHL1                  SERVER         3.72             18.61                  3833.13                 19164.87
CHL3                  SENDER        2.24             10.80                  1198.93                 5771.55
MESSAGE CHANNEL TOTALS          5.96 *****          5032.05 *****          2027.85 *****          3004.20 *****
TPF MQ*SERIES MESSAGE - CHANNEL REPORT          04 MAR      14:29:43      MESSAGE PAGE 2
BSS SUBSYSTEM
=====
MESSAGE CHANNELS FOR LOCAL QUEUE MANAGER: TPFQM          , 5 OBSERVATIONS, 24 SECONDS, CPU B
=====
MESSAGE CHANNEL NAME  TYPE          NEGOTIATED      BATCH SIZE
MEAN                 BATCH SIZE      MEAN             MAXIMUM
-----
CHL3                  SENDER         10              1.50                 5
MESSAGE CHANNEL TOTALS          *****          1.50 *****
TPF MQ*SERIES MESSAGE - CHANNEL REPORT          04 MAR      14:29:43      MESSAGE PAGE 3
BSS SUBSYSTEM
=====
MESSAGE CHANNELS FOR LOCAL QUEUE MANAGER: TPFQM          , 5 OBSERVATIONS, 24 SECONDS, CPU B
=====
MESSAGE CHANNEL NAME  TYPE          MEAN ACTIVE      HIGH WATER MARK
MEAN                 INSTANCES      INSTANCES
-----
CHL1                  SERVER         0.25             2
MESSAGE CHANNEL TOTALS          0.25             *****

```

Figure 64. TPF MQSeries Message Channel Report

```

TPF M*Q*SERIES MESSAGE - QUEUE REPORT          04 MAR      14:29:43      MESSAGE PAGE 1
BSS  SUBSYSTEM
=====
MESSAGE QUEUES FOR LOCAL QUEUE MANAGER: TPFQM          , 5 OBSERVATIONS, 24 SECONDS, CPU B
=====
MESSAGE QUEUE NAME: DEAD.LETTER.QUEUE          , MESSAGE QUEUE TYPE: LOCAL
      OPEN QUEUE RATE:          0.00 MEAN,          0.00 MAXIMUM OBSERVED OPENS PER SECOND
      CLOSE QUEUE RATE:         0.00 MEAN,          0.00 MAXIMUM OBSERVED CLOSES PER SECOND
      ADD PERSISTENT MESSAGE RATE: 0.00 MEAN,          0.00 MAXIMUM OBSERVED PERSISTENT MESSAGES ADDED PER SECOND
      ADD NONPERSISTENT MESSAGE RATE: 0.00 MEAN,          0.00 MAXIMUM OBSERVED NONPERSISTENT MESSAGES ADDED PER SECOND
      ADDED PERSISTENT MESSAGE LENGTH: 0.00 MEAN,          0.00 MAXIMUM OBSERVED BYTES PER ADDED PERSISTENT MESSAGE
      ADDED NONPERSISTENT MESSAGE LENGTH: 0.00 MEAN,          0.00 MAXIMUM OBSERVED BYTES PER ADDED NONPERSISTENT MESSAGE
      SWEEP COUNT:              0 TOTAL OBSERVED SWEEPS DURING DATA COLLECTION RUN
.....
MESSAGE QUEUE NAME: Q1          , MESSAGE QUEUE TYPE: LOCAL
      OPEN QUEUE RATE:          0.00 MEAN,          0.00 MAXIMUM OBSERVED OPENS PER SECOND
      CLOSE QUEUE RATE:         0.04 MEAN,          0.20 MAXIMUM OBSERVED CLOSES PER SECOND
      ADD PERSISTENT MESSAGE RATE: 3.00 MEAN,          15.01 MAXIMUM OBSERVED PERSISTENT MESSAGES ADDED PER SECOND
      ADD NONPERSISTENT MESSAGE RATE: 0.64 MEAN,          3.20 MAXIMUM OBSERVED NONPERSISTENT MESSAGES ADDED PER SECOND
      ADDED PERSISTENT MESSAGE LENGTH: 50.00 MEAN,          50.00 MAXIMUM OBSERVED BYTES PER ADDED PERSISTENT MESSAGE
      ADDED NONPERSISTENT MESSAGE LENGTH: 50.00 MEAN,          50.00 MAXIMUM OBSERVED BYTES PER ADDED NONPERSISTENT MESSAGE
      SWEEP COUNT:              5 TOTAL OBSERVED SWEEPS DURING DATA COLLECTION RUN
.....
MESSAGE QUEUE NAME: SPECIAL.RECOVERY.QUEUE          , MESSAGE QUEUE TYPE: LOCAL
      OPEN QUEUE RATE:          0.00 MEAN,          0.00 MAXIMUM OBSERVED OPENS PER SECOND
      CLOSE QUEUE RATE:         0.00 MEAN,          0.00 MAXIMUM OBSERVED CLOSES PER SECOND
      ADD PERSISTENT MESSAGE RATE: 0.00 MEAN,          0.00 MAXIMUM OBSERVED PERSISTENT MESSAGES ADDED PER SECOND
      ADD NONPERSISTENT MESSAGE RATE: 0.00 MEAN,          0.00 MAXIMUM OBSERVED NONPERSISTENT MESSAGES ADDED PER SECOND
      ADDED PERSISTENT MESSAGE LENGTH: 0.00 MEAN,          0.00 MAXIMUM OBSERVED BYTES PER ADDED PERSISTENT MESSAGE
      ADDED NONPERSISTENT MESSAGE LENGTH: 0.00 MEAN,          0.00 MAXIMUM OBSERVED BYTES PER ADDED NONPERSISTENT MESSAGE
      SWEEP COUNT:              0 TOTAL OBSERVED SWEEPS DURING DATA COLLECTION RUN
.....
MESSAGE QUEUE NAME: SYSTEM.TEMPORARY.QUEUE          , MESSAGE QUEUE TYPE: LOCAL
      OPEN QUEUE RATE:          0.00 MEAN,          0.00 MAXIMUM OBSERVED OPENS PER SECOND
      CLOSE QUEUE RATE:         0.00 MEAN,          0.00 MAXIMUM OBSERVED CLOSES PER SECOND
      ADD PERSISTENT MESSAGE RATE: 0.00 MEAN,          0.00 MAXIMUM OBSERVED PERSISTENT MESSAGES ADDED PER SECOND
      ADD NONPERSISTENT MESSAGE RATE: 0.00 MEAN,          0.00 MAXIMUM OBSERVED NONPERSISTENT MESSAGES ADDED PER SECOND
      ADDED PERSISTENT MESSAGE LENGTH: 0.00 MEAN,          0.00 MAXIMUM OBSERVED BYTES PER ADDED PERSISTENT MESSAGE
      ADDED NONPERSISTENT MESSAGE LENGTH: 0.00 MEAN,          0.00 MAXIMUM OBSERVED BYTES PER ADDED NONPERSISTENT MESSAGE
      SWEEP COUNT:              0 TOTAL OBSERVED SWEEPS DURING DATA COLLECTION RUN

```

Figure 65. TPF MQSeries Message Queue Report (Part 1 of 2)

BSS SUBSYSTEM

=====

MESSAGE QUEUES FOR LOCAL QUEUE MANAGER: TPFQM , 5 OBSERVATIONS, 24 SECONDS, CPU B

=====

MESSAGE QUEUE NAME: XQ1 , MESSAGE QUEUE TYPE: TRANSMISSION

 ADD PERSISTENT MESSAGE RATE: 0.00 MEAN, 0.00 MAXIMUM OBSERVED PERSISTENT MESSAGES ADDED PER SECOND

ADD NONPERSISTENT MESSAGE RATE: 3.00 MEAN, 15.00 MAXIMUM OBSERVED NONPERSISTENT MESSAGES ADDED PER SECOND

 ADDED PERSISTENT MESSAGE LENGTH: 0.00 MEAN, 0.00 MAXIMUM OBSERVED BYTES PER ADDED PERSISTENT MESSAGE

ADDED NONPERSISTENT MESSAGE LENGTH: 50.00 MEAN, 50.00 MAXIMUM OBSERVED BYTES PER ADDED NONPERSISTENT MESSAGE

 SWEEP COUNT: 0 TOTAL OBSERVED SWEEPS DURING DATA COLLECTION RUN

.....

MESSAGE QUEUE T O T A L S

OPEN QUEUE RATE: 0.04 TOTAL MEAN OBSERVED OPENS PER SECOND

CLOSE QUEUE RATE: 0.04 TOTAL MEAN OBSERVED CLOSES PER SECOND

ADD PERSISTENT MESSAGE RATE: 3.00 TOTAL MEAN OBSERVED PERSISTENT MESSAGES ADDED PER SECOND

ADD NONPERSISTENT MESSAGE RATE: 3.64 TOTAL MEAN OBSERVED NONPERSISTENT MESSAGES ADDED PER SECOND

 ADDED PERSISTENT MESSAGE LENGTH: 50.00 TOTAL MEAN OBSERVED BYTES PER ADDED PERSISTENT MESSAGE

ADDED NONPERSISTENT MESSAGE LENGTH: 50.00 TOTAL MEAN OBSERVED BYTES PER ADDED NONPERSISTENT MESSAGE

 SWEEP COUNT: 5 TOTAL OBSERVED SWEEPS DURING DATA COLLECTION RUN

Figure 65. TPF MQSeries Message Queue Report (Part 2 of 2)

Details of Collection/Reduction

This section describes the environment, procedures and effects of collection and reduction elements.

Details of System Collection

Working Storage Usage

When JCS0 records the number of storage blocks in use, the count includes those blocks used by data collection.

Time Stamp

Once activated, JCS0 collects specified data and then time stamps the data record just before passing the record to the data collection tape. During the continuous mode, JCS0 waits for JCF0 to collect the file data prior to collecting the system data and recording the time stamp.

Details of System Reduction

Total Working Storage Assigned

Method: When computing this value the physical size of each block is used rather than the programmer's usable size.

Result: The value printed will be slightly larger and will be a true value of the amount of storage used for pools.

Milliseconds Per Weighted Message

Method: The overhead of data collection is not removed from this value.

Result: The value will be inflated very slightly.

Details of File Collection

File Intercepts

All find and file macro calls are intercepted and stored on the RTC tape, except FDCTC macro calls.

File Queues

Environment: Each time CDC1 intercepts a file macro, the queue for the file module to be accessed is recorded in the FF Record. In addition, the queues for all file modules are recorded at the beginning of each interval in the FQ record.

Effect: File queue information will have a better statistical sample.

Details of File Reduction

Random File Access Summary - Total Accesses Per Second

Data write, data read, and program read counts are accumulated each interval, divided by seconds, and added together to equal total accesses per second during a given interval. The total accesses are then accumulated for the mean. At the end of the tape, the accumulation for the mean is divided by the number of intervals.

Random File Access Summary - Device Queues Mean

Device queue counts are sampled every access, accumulated each interval, divided by samples, and accumulated for the mean. At the end of the tape the accumulation for the mean is divided by the number of intervals.

Random File Access Summary - Total System

Method: VFA accesses are not considered to be random file accesses.
DASD accesses are considered to be random file accesses.

Result: The total system line excludes VFA accesses and includes DASD accesses.

Sequential File Access Summary - YYY Module

A symbolic tape name of YYY denotes that a specified tape drive was inactive at the start of the collection, and that an unspecified tape was mounted on that tape drive or varied online during the collection.

File Accesses Per Record ID Report

Method: Program access IDs are collected as non-printable EBCDIC codes.

Result: An ID of '--' indicates that a program access ID has been intercepted.

Details of Program Reduction

The Identification of Individual Programs

Environment: Individual programs are identified by name in four reports: the Program Names and Enters Report, the Program Details Report, the Cumulative Program Enters Report, and the Cumulative Program On-File Enters Report. There are two avoidable situations where identification is difficult.

Effect: The first situation occurs when the PGMS parameter causes too few rows to be allocated in the PRTABLE and PRISN tables.

The PGMS parameter defines the number of programs requiring data reduction. It does this by setting the number of rows in the PRTABLE and PRISN tables. Each file-resident program generally requires 2 rows in each of those tables, identified by ordinal number in one row, and by name in the other row.

When there are no more rows in PRTABLE and PRISN, an auxiliary row in each of those tables is used to keep track of the overflow information. The information in those auxiliary rows is reported under the pseudo program name of OTHER.

When OTHER appears on a report, the rest of the information in the report, the information ostensibly about individual programs, may be misleading. For example, the data collected for a particular file-resident program that requires two rows to reduce properly may have been reduced with just one row. The information reported about one or more file-resident program may not be complete.

The System of Graphic Cues Used in the Program Reports

A system of graphic cues is used in the program reports to make them easier to read.

- Strings of =’s are used to underline the report columns on which the report rows are ordered.
- Strings of *’s are used to underline the report columns on which the report rows are ordered and that contain cumulative percentages.
- Strings of -’s are otherwise used to underline report columns.
- *’s (for example, * or *.*) are used instead of 0’s (e.g., 0 or 0.00) to denote those rates or percentages that are greater than zero but that would otherwise round down to zero. 0’s are reserved for true zero.
- 100+++ is used to denote a cumulative percentage greater than 100. Such a cumulative percentage can occur when dealing with packages of programs defined in the OPTIONS file, inasmuch as it is possible to define overlapping packages.

The Categorization of Programs

Program reduction determines the category of a program largely on the basis of the Program Allocation Table (IDSPAT). The categories are either file resident (FR) or core resident (CR).

Details of Message Collection

Time Stamp - High-Speed Input Message

Environment: High-Speed input messages are time stamped when the ECB is created; this four-byte time stamp is still in the ECB when intercepted by CDC1.

Effect: It is not necessary, therefore, for CDC1 to time stamp an input message before recording it on the RTC tape.

Time Stamp - High Speed Output Message

Environment: There is no time stamp in an output message block when intercepted by CDC1.

Effect: After intercepting the output message block, CDC1 time stamps it. The time stamp is four bytes long because CDC1 records bits 20 through 51 of the 64-bit time-of-day (TOD) clock.

Details of Message Reduction

All Reports

As noted on the Message Summary Report, the phrase *output messages* refers to output message segments, not output messages as viewed by the terminal

operator. It may take three message segments to fill up a CRT screen with what appears, to the terminal operator, to be one output message.

Action Code Summary, Application Summary - Mean Existence Time

For a continuous mode run only, an input message followed by an output message with the same LNIATA or resource ID (for SNA devices) is considered to be a message pair; its existence time is calculated by subtracting time stamps (except for line 01 because there is no time stamp on the input message).

Action Code Summary - Mean Output Length

Method: When an input message is followed by more than one output message with the same LNIATA, it is assumed this is an output message comprised of more than one segment. The character count field in the additional segments is added to that of the first output message segment.

Result: The total mean output length may not correspond to the mean output length figures seen in other message reports because this is the only report where output message segment lengths of one logical message are added together. This figure then is the mean output message length as seen by the terminal operator, not by the line over which the segments were transmitted.

Action Code Summary, Application Detail - Unmatched Output

When an output message has not been preceded by an input message with the same LNIATA or resource ID (for SNA devices), it is considered unmatched. Its message type cannot be determined, nor can an existence time be calculated.

City Summary

Method: Because line 01 does not have a city code in the DA record, data from messages on this line is not included in City Summary totals.

Result: City Summary totals may not correspond with Message Summary and Action Code Summary totals which include line 01 data.

Terminal Activity

Only those terminals for which activity was intercepted are reported.

Plot/Dist Option - Message

Those variables which have the same value every interval are not plotted.

Weighted Message Rate

Environment: The weighted message rate is obtained from the data reduction driver.

Effect: The message rate is an average over the life of collection, not the rate for just the file collector intervals.

If the weighted message rate is insignificant, the File Access Summary Report is not printed.

Message Stream - Output Messages

- Method: If the optioned city code is followed by an A, a new print line will be advanced to each time a carriage return is encountered in the output message text.
- Result: The output message will be printed in the format seen by the terminal operator. Otherwise, as many as the first 60 characters only are printed unformatted as they appear in the message block.

Message Summary - Max/Min

The maximum and minimum values of those parameters which are not reported on a per second basis are actual observed values. They are not the maximum mean or minimum mean calculated for an interval, nor do they necessarily occur in the same interval.

Use of DA Records

There are two methods of searching through the DA records in order to match a city with its LNIAs:

- Method 1: If the City Summary is optioned, the records are searched by LNIA. Each time a new LNIA is encountered, it is matched with the city code name of the first terminal on that LNIA - unless the device is a CRT on a 2962 (Remote Display Interchange) in which case the next terminal's city code is used. The net result is a table of city code names and their associated LNIAs.
- For all subsequent options by city, this created table is searched for the appropriate LNIAs on the optioned city, rather than researching all the DA records for the necessary information.
- Method 2: If the City Summary is not optioned, the DA records are searched by each city requested for the option being processed. Each unique LNIA associated with the optioned city in the DA record is entered in that city's table unless it is in another city's table.
- Result: Message counts for all terminals on a LNIA will be attributed to the city with which the LNIA has been matched. Occasionally, however, all terminals on one LNIA may not have the same city code name in the DA records. For example, a terminal in the back room may have a different city code name from those terminals used by the reservation agents. The message counts for all these terminals will still be attributed to the one city matched with the LNIA as stated previously.

Limitations of Message Stream Reports

A logical unit (LU) resource is included in the LU message stream, message stream, and distributed data management (DDM) message stream reports only if a resource vector table (RVT) entry exists for that LU resource when data collection is started. If an RVT entry is created for an LU resource after data collection is started, that LU resource is not included in the Message Stream reports.

Collection Details Of Driver

Message Counters

- Environment: At both the start and end of collection, the driver records the

non-keypointed control program message counters (high speed processed, low speed, high speed routed, and created entries message counts), and the system message counters (program entry counters maintained by UII).

Effect: Because these are non-keypointed counters, the driver will record zero message counts if data collection is abnormally stopped because of a system failure.

One-Time Driver Records

Environment: At the start of collection, the driver records on the RTC tape contain various configuration defining records.

Effect: Provides reduction programs with data needed to reference configuration dependent arrays.

Additional Data Collected

Data describing entry control block (ECB) resource requirements is collected to allow the development of other analyses. This data is placed in ECB fields. It must be accessed by customer applications using the EXI user exit. See the *TPF System Installation Support Reference* for more information. The data resides in the ECB fields listed in Table 2.

Table 2. ECB Fields Containing Additional Data

ECB Field	Description
CE1CTRS	The 40-byte area including all the following fields.
CE1DSTMP	An 8-byte field containing the time-of-date (TOD) clock value when the ECB was last dispatched.
CE1EXTIM	An 8-byte field containing the TOD clock value when the ECB was exited. This is set immediately prior to calling the EXI user exit.
CE1ISTIM	An 8-byte field containing the accumulated time the ECB has run on any I-stream. The data is in TOD clock format. It is updated whenever the ECB gives up control or when it exits.
CE1FINDS	A 4-byte field containing the number of Find requests made by the ECB. When the ECB starts a FINDC, FINHC, FINSC, FINWC, FIWHC, or FNSPC macro, the field is updated.
CE1FILES	A 4-byte field containing the number of File requests made by the ECB. When the ECB starts a FILEC, FILNC, FILSC, FILUC, or FLSPC macro, this field is updated.
CE1GETFCS	A 4-byte field containing the number of GETFC requests made by the ECB.
CE1USRID	A 4-byte field initialized to zero at ECB creation available for application programs to store information describing the ECB to facilitate later analysis. This field is not modified by the TPF system.

The duration of the ECB can be determined by subtracting its creation time (found in ECB field CE2TIME) from the exit time (CE1EXTIM). Subtracting the accumulated time the ECB ran from the ECB's duration yields the amount of time the ECB spent waiting to process.

Installation Details

To install the data reduction function:

1. Update the source statements to tailor the function to the installation.
2. Run SIP Stage I and II according to information provided in *TPF System Generation*. This includes compiling and link-editing the data reduction programs.

To optimize the data reduction program for each installation several statements are required to be changed before compilation can begin. These statements are PL/I preprocessor statements, PL/I statements, and linkedit control statements.

This section describes these options for compiling, link-editing, and running the function, as well as input, output, and preprocessor information.

Compilation Phase (Using PL/I Optimizing Compiler)

For the compilation phase, the following minimum options are required:

- MACRO
- MARGINS (2,72,1)
- NOCOMPILE (S)
- OBJECT
- NODECK
- OPTIMIZE (TIME) - (OPT(2))
- SOURCE
- GOSTMT or STMT
- OFFSET

The following options are desirable: AGGREGATE, ATTRIBUTES, ESD, NEST, OPTIONS, and XREF.

The specified options can be found in the PARM field on the EXEC card for the compiler.

A definition of each option can be found in *OS PL/I Optimizing Compiler Programmer's Guide*.

Preprocessor Statements

The following preprocessor statements should be checked before compiling begins. All source modules should use the same settings; otherwise execution errors will occur. Note that where the term allocated is used reference is made to the number assembled into the control program, not the number currently in use. These statements are found in ACP.SYMACRO.RELv (JPC0). All preprocessor statements are identified by a percent sign (%) preceding the statement. See the listing of compiler labels that follows this section.

Table 3. Preprocessor statements

Statement	Shipped Values	Description
BSC	1	This is set to 1 if BSC lines are system generated in the host CPU.
BSCMP	1	This is set to 1 if BSC multipoint lines are system generated in the host CPU.

Table 3. Preprocessor statements (continued)

Statement	Shipped Values	Description
BSCSCT	152	This is set to the maximum number of BSC multipoint station slots system generated in the host CPU.
CIT	35	This is set to the number of cities that have line interchanges.
CP	6	This is the number of loosely coupled CPUs in the system.
CSDMAX	200	This is the number of 3990 subsystem devices.
CSMAX	20	This is set to the number of 3990 cache subsystems.
DASDMOD	3999	This is the maximum number of symbolic modules.
DCMPIF	1	This is set to 1 if MPIF is system generated in the host CPU.
DETAIL	1	If this variable is set to one, code is generated to reduce detail reports (for example, PLOT, DIST, STREAM), as well as summaries. Otherwise, only summaries are produced.
FILE	1	If this variable is set to 1, code is generated to reduce file collector output.
IAC		This is the maximum number of interchange addresses per city.
IAS	220	This is set to the maximum number of Unique Line Interchange city codes in the high-speed line system.
INTERVAL	300	This is set to the maximum number of sampling intervals to be processed.
ISTREAM	8	This is set to the maximum number of I-streams in the system.
LN	244	This is set to the maximum number of high-speed lines allocated to the total system (2703 and 3705EP, plus the number of BSC lines and stations, pseudo lines, links, and AI lines defined, plus one).
MAPPL	\$\$\$	The maximum number of applications in the network (see RC11T).
MESSAGE	1	If this variable is set to 1, code is generated to reduce message collector output.
MLSIZE	2928	This is set to 12 times the value of LN.
MXRECID	100	This is set to the maximum number of Unique Record IDs to be processed.
NPGM	35	The number of program names in the SM record (the maximum is 63).
PGMS	600	This is set to the maximum number of programs expected to be observed during a collection.
PROGRAM	1	If this variable is set to 1, code is generated to reduce program collector output.
RANDOM	132	This is set to the total number of DASD modules allocated in the system.
SDA_FAC	25	Cutoff percentage for Channel Path ID Activity Exception Report.
SDDMAX	200	This is set to the number of DASDs backing 3990 storage directors.
SDMAX	20	This is set to the number of 3880 storage directors.
SLC	1	This is set to 1 if synchronous link control (SLC) support is system generated in the host CPU.
SNA	1	This is set to 1 if SNA support is system generated in the host CPU.
SNALNE	255	This is set to the maximum number of SNA NCP/ALS devices system generated in the host CPU plus one.
SNODES	2750	This is set to the maximum number of SNA nodes system generated in the host CPU. This value must not exceed 98301.
SNSIZE	4080	This is set to 20 times the value of SNALNE.
SS	4	This is the number of subsystems.
SSU	8	This is the number of subsystem users.

Table 3. Preprocessor statements (continued)

Statement	Shipped Values	Description
SYSTEM	1	If this variable is set to 1, code is generated to reduce system collector output.
TAPES	64	This is set to the number of tape entries allocated to the tape status table.
TMC	700	This is set to the maximum number of terminals in 1 city.
USERNAM	TPF 4.1 HPO SYSTEM	This variable is initialized to 25 characters and should contain the name of your system.
VFA	1	This is always set to one.
WT	5	This is set to the weighting factor for low-speed messages.
WT2	3	This is set to the weighing factor for routed messages and coded as 3 in the DATACO macro. SKJPC0 then converts this to a 0.3, which is the value used by the data collection reporting programs for the weighted message rate calculation.

Variables

JRA1

The following variables need to be initialized with the proper constants. These variables are found in ACP.SRCE.OL.RELvv (JRA1XX).

DSUTLIST	System utilities
DAUTLIST	Application utilities
DNMTBLC	This array contains 4-character names that should be made the same as the system message counter table in the online system (SM record).
DMSGDES	This array contains 30-character descriptions of the functions of the program listed in DNMTBLC. The order of DNMTBLC and DMSGDES must be alike.

JRM2 and JRM4

Action Code Changes/Additions

The first byte of the action code can be 1 of the following characters: A-Z, 0-9. The second byte of the action code can be 1 of these characters: A-Z, 0-9, =, *, /, -, or blank. Any changes in these action code groups require changes in JRM2 and JRM4.

Linkedit Phase

The linkedit control cards do not need to be changed unless you change the names of the procedures, alter the overlay structure, or add new procedures. The linkedit is punched out through SIP Stage I. These statements call modules in from the library (INCLUDE).

Take care altering the order of these cards. The link editor must receive the module with the initial attribute for an external variable before any other reference to that variable. Therefore, the sequence of the include statements is very important.

Use the following options for the linkedit phase:

- LET
- XREF

- LIST
- DCBS.

Execution Phase

See *TPF Operations* for more information about running the DATAREAD load module and the options to use.

Input

The input for installation of the data reduction function consists of 19 data sets of source statements, and 1 file of preprocessor statements (JPC0). A group of intermediate files (object library) is used to store the object code between compilation and linkedit.

Files

The system performance programs reside in 20 data set members. The unpacking procedure described in the *TPF Program Directory* creates the following libraries:

- ACP.SYMACRO.RELv
- ACP.SRCE.OL.RELv
- ACP.OBJ.RELv
- ACP.LINK.RELv

The following tasks are one way to install the data reduction system into an IBM MVS environment. There is no one way to do so. The one presented is one known to work.

1. Print the following datasets:
 - ACP.SRCE.OL.RELv(JRA1xx)
 - ACP.SRCE.OL.RELv(JRM2xx)
 - ACP.SRCE.OL.RELv(JRM4xx)
 - ACP.SYMACRO.RELv(JPC0xx).

You now have a listing of the datasets to be altered to install the data reduction package.

2. Alter the statements that need to be changed to fit your installation. This can be done using the standard library update procedure.
3. After the updates have been made, SIP Stages I and II must be run to provide the executable DATAREAD module. The SIP Stage I input contains all the parameters required to initialize macro JPC0.
4. Verify that all the segments were compiled without severe errors and the link-edit step has no errors.
5. Data reduction is now installed. See *TPF Operations* for more information about system performance and running the program.

The macro library ACP.SYMACRO.RELv has member JPC0.

The source library ACP.SRCE.OL.RELv and the object library ACP.OBJ.RELv have the following members: JRA0xx, JRA1xx, JRA2xx, JRA3xx, JRA4xx, JRA5xx, JRA6xx, JRF1xx, JRF4xx, JRF5xx, JRF6xx, JRM1xx, JRM2xx, JRM4xx, JRM5xx, JRP1xx, JRP3xx, JRS1xx, JRS3xx

The library ACP.LINK.RELv has 1 member: DATAREAD.

Notes:

1. xx represents the current version number of the module.
2. vv represents the current release number of TPF.

These data sets contain all the statements needed to install the data reduction function.

Output

The output from the compilation and link-edit phases will be standard PL/I and link-edit listings. See *PL/I Optimizing Compiler: Language Reference Manual* and *Linkage Editor Manual* for descriptions.

Error Messages

Each compile step starts with a terminal error. This error is caused by the PROCESS card being the first input statement. This condition causes no problems. JRA0 is the only procedure with the option MAIN, therefore all other procedures will have a warning message. Several errors and warnings will be produced. These will be because of:

- A variable being defined on another variable of different data characteristics.
- Failure to initialize all elements of an array.

Data conversion errors are expected and will not cause execution problems. Any other errors should be investigated.

The options used to generate Data Reduction reports and sample JCL for running data reduction are in *TPF Operations*.

Index

Numerics

3880 storage control report complexwide 49
3990 storage control report complexwide 52

A

access per CHPID 45
accesses per second 43, 45
action code summary 70, 86
 mean output length 86
 report 69
active queue 33
active subsystem pools 28
activity, system 25
Airlines Line Control (ALC) 2
analysis
 cylinder 47
 data reduction 6
 file data 41
 program data 61
 system data 11
 system performance 11
Anomaly Log, DP Record 61
application detail 86
application detail report 71
application summary 86
application summary and detail reports 70

B

back event 62
blocked conditions 26

C

cache
 logical record cache summary report 29
channel path id status 45
CHPID accesses percentage 45
CHPID reports. 45
city summary 86
 message 72
 report 71
collection
 interval 24
 message 85
collectors 4
collectors chosen 13
commit data 26
compiler
 labels 91
 PL/I optimizing 91
conditions, shutdown 28
configuration
 working storage 26
continuous data collection 9
continuous mode in data collection 6

Coupling Facility Caching Summary report 29
Coupling Facility Locking Summary report 29
Coupling Facility Structure Summary report 29
Coupling Facility Usage Summary report 29
CPU lists 24
cylinder analysis report 47

D

data
 collection 4
 manipulation equations 26
 reduction analysis 6
 threads support 25
data collection
 3880 storage control report 49
 3880 Storage Control Report for SDA mmmm, SDID
 nn 51
 3880 Storage Control Report: Subtotals for SDID
 nn 50
 3880 Storage Control Report: Totals 49
 3990 Storage Control Cache Summary Report 58
 3990 storage control report 52
 3990 Storage Control Report for SDA mmmm, Cache
 Subsystem nnnn 57
 3990 Storage Control Report: Subtotals for Cache
 Subsystem nnnn 53
 3990 Storage Control Report: Totals 52
 application detail report 71
 channel path activity 46
 channel path id status 45
 city summary 72
 collectors chosen 13
 Coupling Facility Caching Summary report 29
 Coupling Facility Locking Summary report 29
 Coupling Facility Structure Summary report 29
 Coupling Facility Usage Summary report 29
 cylinder analysis 47
 cylinder analysis report 47
 data reduction limitations report 41
 DDM message stream report 78
 DP Record Anomaly Log 61
 ECB frame usage summary report 35
 expected arm movement 47
 file accesses per cylinder 47
 file accesses per record id 46
 file comparison 47, 48
 file comparison plot 48
 file data analysis 41
 frequency distribution 38
 heap area usage report 35
 input messages per second 16
 instantaneous activity 17
 logical record cache summary report 29
 LU activity
 See logical unit
 message stream by agent 77
 message stream by lu 76

- data collection (*continued*)
 - message summary 69
 - MPIF path activity report 37
 - nonretentive access: 49
 - normal cache replacement 49
 - package collection details 65
 - package enters 65
 - package on-file enters 66
 - plot 40
 - processor utilization 16
 - program details for a package 63
 - program enters for a package 64
 - program legend 67
 - program names and enters 62
 - program on-file enters for a package 64
 - program report 61
 - pushbutton application summary report 14
 - random file access summary 44
 - record ID file accesses report 46
 - record-caching storage controls 48
 - report 16
 - sequential file access summary 46
 - shutdown conditions 16
 - system 11
 - system pools summary 28
 - TCP/IP message summary report 78
 - TCP/IP weighted input messages by application 34
 - terminal activity 73
 - threads support data 17
 - TPF Internet mail server summary report 33
 - TPF MQSeries message channel report 79
 - TPF MQSeries message queue report 79
 - weighted-message rate 67
 - working storage utilization 16
- data collection basics 1
- data collection/reduction overview 9
- data reduction limitations report 40, 41
- DDM message stream report 78
- deferred queue 33
- distributed data management (DDM) message stream report 77
- DP Record Anomaly Log 61
- driver collection/reduction 87

E

- e-mail
 - TPF Internet mail server summary report 33
- ECB classifications
 - low-priority 25
- ECB classifications, low-priority 25
- ECB Frame Usage Summary Report 35
- electronic mail (e-mail)
 - TPF Internet mail server summary report 33
- enter event 62
- enters per second per message 67
- environment
 - report 11
- equations, data manipulation 26
- error messages 95

- examples
 - file access report 42
 - system report 13, 16
- execution phase 94
- expected arm movement 47

F

- file
 - 3880 Storage Control Report for SDA mmmm, SDID nn 51
 - 3880 Storage Control Report: Subtotals for SDID nn 50
 - 3880 Storage Control Report: Totals 49
 - 3990 Storage Control Cache Summary Report 58
 - 3990 Storage Control Device Summary Report for Cache Subsystem nnnn 59
 - 3990 Storage Control Report for SDA mmmm, Cache Subsystem nnnn 57
 - 3990 Storage Control Report: Subtotals for Cache Subsystem nnnn 53
 - 3990 Storage Control Report: Totals 52
 - access report 42
 - collection/reduction 83
 - comparison plot 48
 - comparison report 47
 - DP Record Anomaly Log 61
 - library 94
 - storage control reports 48
- file accesses per cylinder 47
- file accesses per record id 46
- file accesses per record id report 84
- file comparison plot 48
- file queues 83
- frame usage report 35
- frequency distribution report 38
- functions, collection 1

H

- heap area usage report 35
- high-speed
 - input message, time stamp 85
 - output message, time stamp 85

I

- initiation, variable 93
- input message, high-speed 85
- input, installation 94
- installation
 - data collection/reduction programs 7
 - input 94
- installation details 91
- instantaneous activity 25
- intercept 83
- intercept, file 83
- Internet mail
 - TPF Internet mail server summary report 33
- interprocessor communication MPIF report 37

interval, collection 24

J

JC data collection 1
JRA1 93
JRM2 93
JRM4 93

L

linkedit phase 93
lists, CPU 24
LNIATA 2
logical end identifiers 3
logical record cache summary report 29
logical unit 75
 activity 75
 activity report 74
 SNA 3
low priority ECB classifications 25
LU activity
 See logical unit

M

mail queue
 active 33
 deferred 33
maximum pages used 35
mean existence time 86
mean pages used 35
message
 collection 85
 error 95
 high-speed 22
 input 22
 processing 2
 routing facility 3
 TCP/IP 22, 34, 78
 weighted 2, 22, 83
 weighted rate 86
message data analysis 67
message rate, weighted 67
message stream
 by agent 77
 by lu 76
 DDM report 78
 logical unit report 75
 output messages 87
 report 76
 reports, limitations of 87
 TPF MQSeries message channel report 79
 TPF MQSeries message queue report 79
message summary 69
 max/min 87
message summary report 67
messages per second performance 22
MPIF 37
MPIF path activity report 37
MQSeries message channel report 79

MQSeries message queue report 79
multiple database function (MDBF) 1

N

NEF 3
network control program (NCP) 3
network extension facility 3
nonretentive access: 49
normal cache replacement 49

O

offline data reduction 1
one-time driver records 88
online data collection 1
options
 compilation phase 91
 execution phase 94
 linkedit phase 93
output message
 high-speed 85
overview
 data collection/reduction 9
 performance 1

P

package
 collection details 65
 details report 64
 enters 65
 enters report 65
 on-file enters 66
 on-file enters report 66
pages allocated for ECB 35
path activity per CHPID report 45, 46
path, MPIF 37
percentage access per CHPID 45
performance basics 1
performance overview 1
PL/I optimizing compiler 91
plot report 39, 40
plot/dist option - message 86
preprocessor statements 91
processing, message 2
processor utilization 23
program
 collection/reduction 84
 data analysis 61
 installing data collection/reduction 7
program details
 for a package 63
 report 62
program enters
 for a package 64
 report 63
program legend 67
program legend report 66
program names and enters 62
program names and enters report 62

- program on-file enters for a package 64
- program on-file enters report 64
- pushbutton application
 - summary report 14

Q

- queue, file 83

R

- random file access
 - device queues mean 84
 - summary 44
 - total accesses per second 84
 - total system 84
- Record Anomaly Log 61
- record ID file accesses report 46
- reports
 - 3990 Storage Control Device Summary Report for Cache Subsystem nnnn 59
 - Coupling Facility Caching Summary report 29
 - Coupling Facility Locking Summary report 29
 - Coupling Facility Structure Summary report 29
 - Coupling Facility Usage Summary report 29
 - cylinder analysis 47
 - file 42
 - logical record cache summary report 29
 - MQSeries message channel report 79
 - MQSeries message queue report 79
 - plot 39
 - system 13
 - TPF Internet mail server summary report 33
- rollback data 26

S

- sampling mode in data collection 6
- SDA activity exception report 46
- sequential file access 84
- sequential file access summary 46
- shutdown 28
- shutdown conditions 15, 16
- SNA logical unit 3
- stamp, time 83
- subsystem pools 28
- subsystems 1
- Synchronous Data Link Control (SDLC) 3
- system
 - activity 25
 - collection, details 83
 - collection/reduction 83
 - data analysis 11
 - performance analysis 11
 - pools summary 28
 - reduction 83
 - report 11
- system control program 3
- system pools 28

T

- TCP/IP message summary report 78
- TCP/IP weighted input messages by application 34
- techniques
 - data collection 4
 - data reduction 6
- terminal activity 73, 86
- terminal activity report 72
- threads support 25
- time stamp
 - general 83
 - high-speed input message 85
 - high-speed output message 85
- total working storage 83
- TPF Internet mail server summary report 33
- TPF logical record cache summary report 29
- TPF MQSeries message channel report 79
- TPF MQSeries message queue report 79
- TPF transaction services data 26

U

- unmatched output 86
- use of da records 87
- utilization
 - processor 23
 - resource 23

V

- variable initiation 93

W

- weighted message 22
- weighted message milliseconds 83
- weighted-message rate
 - report 67
- weighting factor 23
- working storage
 - configuration 26
 - total 83
 - using 83
- working storage utilization 16



File Number: S370/30XX-34
Program Number: 5748-T14



Printed in the United States of America
on recycled paper containing 10%
recovered post-consumer fiber.

SH31-0170-12

