StorageTek

Expert

Performance

Reporter

(ExPR)

MVS User's Guide

Version 1 Release 1

Contents

Contents   **5**

Contents   **7**

Preface

Overview

This document describes how to install, configure, and use the MVS component of the StorageTek Expert Performance Reporter (ExPR) software product.

Optional Windows Help System

Reference information in this MVS User's Guide is also available as a Windows-based Help system on the PC. The help system includes help for CONFIG and SYSIN statement syntax and ExPR error messages.

The ExPR MVS Help system can be accessed from the ExPR PC main window or from the ExPR PC Help system.

Audience

The audience for this document includes MVS system programmers and capacity planners who will install, configure, and use the ExPR MVS component, as well as StorageTek support personnel who support ExPR customers in these efforts.

Organization

This document is organized as follows:

- Chapter 1, *Introduction*, describes the ExPR software product, including what ExPR does and how it works.

- Chapter 2, *Installation and Customization*, describes a set of sequential tasks to configure the ExPR system for your site after the initial SMP/E installation is completed.

- Chapter 3, *Managing the ExPR Database*, describes the database update process and also how to verify and consolidate the database.

- Chapter 4, *Exception Processing and Reports*, describes how to configure and use the exception processing feature of ExPR to define exception thresholds and generate mainframe tabular exception reports.

- Chapter 5, *Jobname Workload Groups*, describes how to configure and use the jobname workload groups features of ExPR to logically map Nearline activity by specified jobnames for use in ExPR reports.

- Chapter 6, *Tape Catalog Processing and Reports*, describes how to configure and use the tape catalog processing feature of ExPR to collect information from your site's tape catalog system and generate tape cartridge-based reports. This chapter also describes how to use the dataset workload group feature in conjunction with tape catalog processing to logically map tape cartridge activity by specified datasets for use in ExPR reports.

- Chapter 7, *Mainframe Performance Reports*, describes how to generate and interpret mainframe tabular performance reports.

- Chapter 8, *File Transfer*, describes the process of extracting data from the mainframe database and transferring it to a PC for further processing in the ExPR PC environment.

- Chapter 9, *API*, describes the ExPR Application Programming Interface (API) and how it is used to create customized mainframe tabular reports.

- Appendix A, *CONFIG Parameters*, lists and describes ExPR configuration (CONFIG DD) parameter statements.

- Appendix B, *SYSIN Parameters*, lists and describes ExPR runtime (SYSIN DD) parameter statements.

- Appendix C, *Messages*, lists messages that are generated by ExPR and provides possible solutions to problems that are indicated by these messages.

- Appendix D, *Database Layout*, provides supplemental technical information about the ExPR mainframe database.

- Appendix E, *ExPR Version 1 Release 1 Modifications*, describes modifications and enhancements to ExPR that are introduced in this release of the product.

Related Documentation

- *ExPR MVS Installation Notes* - describes how to SMP/E install the ExPR MVS software on the MVS host system.

- *ExPR PC User's Guide* - describes the ExPR software product, including instructions for installing the software in the Windows environment, building and maintaining the PC database of downloaded information, and selecting reports for graphical display.

- *IBM System Management Facility* (GC28-1628-nn), describes the SMF record collection process.

Conventions

The following conventions are used in this document:

- Parameter keywords are shown in **boldface** type

- Variable information you enter is shown in *italics*

- A bar (|) separates mutually exclusive choices in command strings

- Parentheses and single quotation marks must be entered as shown in command strings

StorageTek Support

To report problems and receive support from StorageTek Software Support Representatives (SSRs), refer to the separate document, *How to Request Help or Software Enhancements*.

Chapter 1

Introduction

Overview

This chapter describes the StorageTek ExPR software product, including what ExPR does and how it works.

What ExPR Does

ExPR collects performance data and generates reports about StorageTek Nearline Automated Cartridge System (ACS) status and performance. ExPR reports are available in tabular format in the MVS mainframe environment and in graphical or tabular format in the PC Windows environment.

How ExPR Works

ExPR software is packaged as two unique software systems, ExPR MVS and ExPR PC:

- ExPR MVS, which resides on an MVS host system, builds and maintains a database of historical performance data that it collects from the Nearline system, from the MVS operating system, and optionally from the site's

tape management system. Tabular performance and exception reports are generated directly from this database for display in the MVS environment. For ExPR MVS, all processes are controlled by a batch task interface. ExPR MVS is described in this document.

- ExPR PC, which resides on one or more PCs, is a Windows application that manages user-customized subsets of the mainframe database on the PC. Data that is transferred from the mainframe database can be queried and displayed in graphical or tabular performance, exception, and quick summary reports. Data can also be ported to a Microsoft Excel-compatible spreadsheet and other external applications for further processing. ExPR PC is described in the *ExPR PC User's Guide*.

ExPR Mainframe Databases

ExPR MVS collects information relating to Nearline performance and inputs it to an ExPR mainframe database. This process is controlled by user-specified parameters.

The ExPR mainframe database is a VSAM KSDS file. Data can be stored as a single database or as multiple databases (one for each ACS, in multiple ACS configurations, etc.), depending on the size and complexity of the installation and its reporting requirements.

A single database can be maintained containing data related to all ACS environments, or multiple databases can be maintained to reflect the reporting needs of various groups within the organization. Databases can be defined for each

individual ACS or a logical group of ACSs determined by
application, host system attachment, or geographic location.

Data from more than one MVS system can also be input to a
single ExPR database. This is useful where multiple MVS
systems share a single Nearline ACS environment.

The ExPR database update and maintenance process is
described in chapter 3, *Managing the ExPR Database*.

ExPR PC Databases

A subset of the mainframe database is maintained as a
relational database (or databases) on each PC where ExPR PC
software is installed. For each PC user, a batch job process is
executed at user-specified intervals on the mainframe,
requesting that selected data be extracted to the mainframe
extract file for transfer to the PC. Once transferred, the PC
user inputs the file to the PC database. Graphical
performance reports are generated from this database. The
ExPR PC database is described in the *ExPR PC User's Guide*.

ExPR Data Flow

Figure 1-1 provides a simplified overview of how data flows
through the ExPR system. Information is collected from
various sources (PGMI, tape catalog, SMF, RMF, MVS
SYSLOG, etc.) and input to the ExPR mainframe database.
Collected data can then be processed for mainframe tabular
reports or transferred to a PC for graphical display and/or
porting to a Microsoft Excel-compatible spreadsheet and
other external applications.
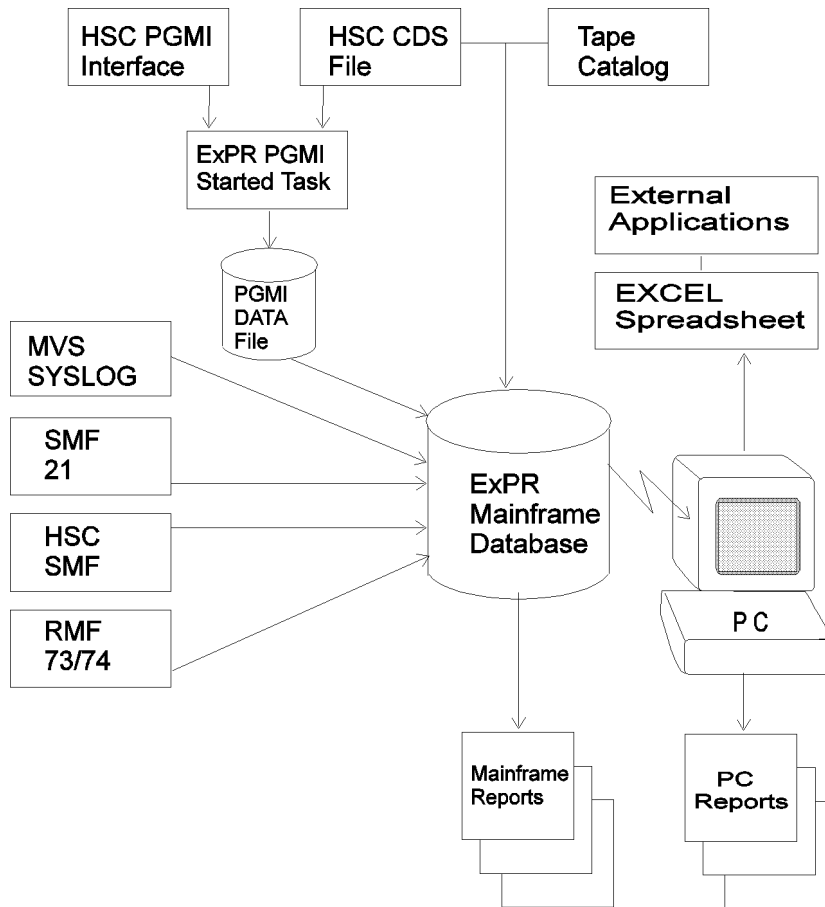
Figure 1-1: ExPR Data Flow

　　　　ExPR MVS User's Guide

ExPR Reporting Features

Reports can be generated to track the capacity, utilization, and performance of Nearline configurations ranging from a single LSM to complex multi-ACS environments utilizing mixed transport types.

The following types of reports are available:

- Tabular performance, exception analysis, and tape catalog reports (host and PC)

- Graphical performance, exception analysis, and tape catalog reports (PC)

- User-customizable tabular reports through the supplied API (host)

Report Granularity

ExPR provides mechanisms for isolating precise data in your reports. Reports can include data for all or specific MVS systems, ACSs, LSMs, and device types. Information can be displayed on the mainframe as hourly detail and daily summary formats. If greater granularity is required, the PC component supports dynamic summarization in hourly, daily, weekly, monthly, quarterly, or yearly formats.

Exception Processing

ExPR exception processing monitors user-specified exception thresholds, which represent the point at which an event is considered to be an exception, and generates a report when

thresholds have been exceeded, such as when the number of mounts or mount response time exceeds the specified value.

Thresholds can be set for scratch mounts, non-scratch mounts, scratch mount response time, non-scratch mount response time, enters, ejects, passthroughs, control unit load, number of drives in use, available scratch, available free cells, and maximum allocation recovery time.

Exception reports may be run in tabular hardcopy format on the mainframe. Input is user report requests that are run as a batch job process against exception data in the mainframe database.

Mainframe exception reports include:

- The THRESHOLD EXCEPTIONS report, which compares the thresholds you have defined against the database records over a selected period of time and highlights those fields which exceed a threshold.

- The SMF EXCEPTION EVENT report, which compares the thresholds defined for mount response time with the mount events as they were recorded.

Graphical or tabular exception reports can also be generated on the PC. These reports can be run against the mainframe thresholds or against user-specified local thresholds.

Exception processing is described in chapter 4, *Exception Processing and Reports*. PC exception reports are described in the *ExPR PC User's Guide*.

Jobname and Dataset Workload Groups

Site-specified jobname and dataset workload groups provide mechanisms for logically mapping activity from specified jobnames or datasets into groups against which reports can be generated, increasing the ExPR reporting capability by further defining Nearline activities within your organization.

- Jobname workload groups map Nearline activity from specified jobnames into logical groups, typically defining activities generated by specific departments or functions. Jobname workload groups are described in chapter 5, *Jobname Workload Groups*.

- Dataset workload groups, used in conjunction with ExPR tape catalog processing, logically map the contents of a library based on dataset names, typically defining critical applications or system components. Dataset workload groups are described in chapter 6, *Tape Catalog Processing and Reports*.

Note: Jobname workload groups and dataset workload groups are completely separate ExPR features; there is no direct correlation between the reports that are generated by either type of workload group.

Tape Catalog Processing

Tape catalog processing is an optional ExPR feature that provides a comprehensive set of snapshot and historical comparison reports relating to the age and utilization of cartridges in the library. You can feed your entire tape catalog into ExPR or allow ExPR to filter only Nearline entries. Mixed media and transport device types are

supported. Output is in a tabular hardcopy format on the mainframe and in a graphical or tabular format on the PC.

Mainframe tape catalog reports include the following:

- The tape processing update Aging report lists, for each LSM and for each user-defined dataset workload group, the number of cartridges within user-specified age bands and their average ages, the number of cartridges containing multi-volume datasets, and the number of cartridges containing multiple datasets.

- The tape processing update Utilization report lists, for each LSM and for each user-defined dataset workload group, the number of cartridges within each tape utilization percentage band.

- The Volume Details report lists, for each volume, the primary dataset name (or optionally all dataset names), volume sequence, number of megabytes, ACS/LSM location, number of datasets, estimated utilization percentage, last-reference date, number of accesses, device type and cartridge length, scratch status, and the name, file sequence, block count, block size, record size, and record format of each dataset.

- The Volume Summary report provides tape catalog volume contents information summarized for each ACS, each LSM, each dataset workload group, each device type, each media type, and each defined tape length.

- The Tape Catalog History report provides comparative activity analysis between two tape catalog images by listing volumes that have changed status.

The tape catalog feature is described in chapter 6, *Tape Catalog Processing and Reports*.

Mainframe Performance Reports

Mainframe performance reports are run as a batch job process against the mainframe database. Output is in a tabular hardcopy format. Each report presents data sorted by hour for a specified date range on all or selected ACSs and LSMs.

Mainframe performance reports include the following:

- The MOUNTS report lists the number of Nearline mounts, average mount time, and maximum mount time. Separate tallies are reported for scratch and non-scratch volumes. Total global (including non-Nearline) mounts are also reported separately.

- The MOUNTS-DETAIL report provides a breakdown of the mount response time components for both scratch and non-scratch mounts.

- The CONTENTS report lists cartridge movements (enter/eject count and number of passthroughs) and contents information (scratch count, free cells, total cells).

- The UTILIZATION report lists the percentage of time the robotics system was in use and the percentage of time that drives were concurrently in use.

- The CU-BUSY report produces control unit information, including the number of drives with a disconnect time exceeding the user-defined threshold.

- The PATH-BUSY report produces information about user-defined channel group utilization.

- The TAPE-ERRORS report provide listings of the number of temporary and permanent read/write errors that occur. The report also provides an audit trail of which drive and volume the media errors occurred on.

- The ALLOC-REC report provides information about allocation recovery tasks on JES2 systems, including the number of events and the delay times.

- The DEMAND-ENTERS report provides information about demand enters that occurred during the reporting period.

These reports are described in chapter 7, *Mainframe Performance Reports*.

Graphical Reports

Graphical performance and exception reports are generated on the PC using data that is extracted from the mainframe database and transferred to the PC. Data that is used for generating graphical reports can also be displayed in tabular format on the PC. These reports are described in the *ExPR PC User's Guide*. The process of transferring data from the mainframe is described in chapter 8, *File Transfer to PC*.

User-Customizable Reports

The supplied ExPR Application Programming Interface (API) provides a mechanism for writing custom reports against the

ExPR mainframe database. These reports can be developed by your technical staff. The API is described in chapter 9, *ExPR API*.

Chapter 2

Installation and Customization

Overview

This chapter describes the sequential tasks you will perform after the basic SMP/E software installation is complete. The SMP/E software installation process is described in the separate document, *ExPR MVS Installation Notes*.

IMPORTANT

The ExPR installation and customization process includes the establishment of a started task in your MVS startup parameters. Depending on your site rules, you may need to schedule this procedure in advance with your system administrator.

Before you install the software, you should review this chapter to become familiar with ExPR site planning issues and system configuration processes to assess their impact at your site.

MVS tailoring may be required to define how ExPR fits into your mainframe environment. Specifically, changes to SYS1.PARMLIB and SYS1.PROCLIB may be required to reflect site parameters and how ExPR interacts with other products in your MVS operating environment.

Skills Required

The person who installs and configures ExPR should have experience with the following:

- JCL
- SMF/RMF post-processing
- MVS SYSLOG external writer
- SYS1.PARMLIB members
- VSAM
- SMP/E
- IBM Utilities

Control Statements

Two types of control statements, CONFIG and SYSIN, are used to configure and operate the ExPR MVS software. These statements can be interchanged (i.e., a CONFIG statement could be specified in the SYSIN jobstream). However, the following rules have been applied to the sample decks supplied with the product:

- CONFIG statements are used for static settings, such as defining the Nearline configuration
- SYSIN statements are used for day-to-day requests, such as generating reports

Information that is specific to your installation, such as MVS host names and ACS, LSM, and tape drive device IDs, can be collected automatically by ExPR from the HSC control dataset (CDS) and input to the configuration process by the ExPR automatic Config generator feature. However, ExPR cannot automatically collect all site data (for example, device types, channel paths, LSM types, and optional ACS and LSM names), so some manual entry is required even when the automatic Config generation feature is used.

When interpreting parameters, ExPR looks first to the CDSCOPY dataset to see if automatically generated data is available, then to the CONFIG dataset, and finally to the SYSIN dataset.

The automatic Config generator feature is described in more detail in *Task 7: Review ExPR Parameters and Configuration Options* in this chapter. At that point in the installation process, you will make a decision to use the ExPR automatic Config generator feature or input all site-specific information manually.

CONFIG and SYSIN control statements are described contextually throughout this user guide and individually in appendix A, *CONFIG Statements* and appendix B, *SYSIN Statements*. Sample control cards are also provided throughout to assist you in the ExPR configuration process.

Configuration Task Summary

The ExPR MVS configuration process is organized as a series of sequential tasks you will perform.

- The initial tasks provide you with technical information about ExPR and help you make site planning decisions.

- The remaining tasks involve copying sample control cards from the supplied SAMPLIB, customizing them where appropriate to substitute site-specific data, and inserting them in your batch jobstream. The ExPR automatic Config generator feature may be used to automatically supply some of this data.

Note: The configuration tasks described in this chapter will build a basic functional system for your MVS environment. Additional customization is required if you wish to implement optional features, including exception processing, jobname workload groups, and tape catalog processing. These features are described in chapters 4, 5, and 6 respectively.

The table below summarizes the tasks you will perform to configure the MVS portion of the ExPR software for your site.

ExPR Configuration Task Summary

Task	Description
Task 1: Review Data Collection Procedures	This task provides a necessary technical introduction to the types of data ExPR will collect, and reviews several housekeeping issues that should be considered when inserting the ExPR data collection process into your MVS batch jobstream.
Task 2: Space Planning	This task helps you determine the probable DASD requirements at your site for ExPR datasets and the ExPR database.

Task 3: Review Hardware and Software Prerequisites	This task helps you ensure that all required hardware and software components are in place before you start configuring ExPR.
Task 4: Define ExPR Datasets and Tailor the EXPRPROC Procedure	This task allocates the datasets that are required by ExPR and initializes the PGMI collection file. In this task, you will customize sample control cards with site data you calculated in *Task 2, Space Planning*. You will also customize the standard EXPRPROC procedure and place it in an MVS PROCLIB.
Task 5: Define ExPR Database(s)	This task helps you determine if a single or multiple ExPR databases are suitable for your site's reporting requirements, and then defines the database or databases. In this task, you will customize sample control cards with site data and database information you calculated in *Task 2, Space Planning*.
Task 6: Initialize the ExPR Database	This task initializes the database(s) that were defined in the previous task. In this task, you will run the sample initialize on the database(s) that were defined in *Task 5, Define the ExPR Database(s)*.
Task 7: Review ExPR Parameters and Configuration Options	This task reviews the @-prefixed SAMPLIB members and the CONFIG and SYSIN parameter-generation options you will use to establish ExPR functions for your site. This task also helps you determine whether to use the ExPR automatic Config generator feature or input site-specific information manually.

Chapter 2: Installation and Customization   **29**

Task 8: Define the Nearline Configuration to ExPR	This task identifies your Nearline configuration to ExPR, including ACSs and LSMs, tape drive devices and device types, and channel paths and channel path groups.
Task 9: Install the ExPR Started Task	This task defines the ExPR PGMI started task and establishes it in your MVS jobstream. In this task, you will copy sample control cards to your selected PROCLIB library and customize them to define started task runtime parameters for your site.
Task 10: Specify HSC SMF Number to ExPR	This task identifies the HSC SMF number for your site to ExPR. In this task, you will customize sample control cards with the HSC SMF number for your site.
Task 11: Define the Control Dataset to ExPR	This task identifies the name of the HSC Control Dataset (CDS) that ExPR is to reference. In this task, you will customize sample control cards with the HSC CDS name that ExPR is to reference.
Task 12: Create the ExPR CONFIG Dataset	This task copies the customized sample control cards from previous tasks into the live CONFIG dataset. In this task, you will customize parameters provided in the sample control cards to copy your edited SAMPLIB entries into the live CONFIG dataset.
Task 13: Installation Verification	This task tests the basic functionality of the completed ExPR installation to ensure that all required processes have been satisfactorily performed.

IMPORTANT: Copy Your SAMPLIB First!

It is highly recommended that you make a copy of the ExPR SAMPLIB files and retain a clean master copy of the files as they were installed from the distribution tape. This way, if your initial tailoring efforts do not produce the desired results, you will still have a clean copy of the SAMPLIB files available.

Task 1: Review Data Collection Procedures

This task provides a necessary technical introduction to the types of data ExPR will collect, and reviews several housekeeping issues that should be considered when inserting the ExPR data collection process into your MVS batch jobstream.

Data Collection Sources

ExPR collects data relating to Nearline performance from the following sources:

- MVS SMF records

- MVS RMF records

- HSC SMF records

- MVS system log (console output)

- HSC PGMI records (started task)

- HSC CDS (started task)

- Tape catalog data (optional process)

MVS SMF Processing

ExPR uses standard MVS SMF type 21 records to extract external tape statistics and media error information. Input datasets containing the SMF records are identified in the JCL, with ExPR parameters providing controls for date range.

MVS RMF Processing

RMF records are used to identify 4400 channel and control unit activity and SSCH (start subchannel) counts. Configuration information provided during installation is cross-referenced to identify the specific 4400 components. Record types 73 and 74 are used. Input datasets containing the RMF records are identified in the JCL, with ExPR parameters providing controls for date range and configuration information.

HSC SMF Processing

The HSC User SMF record provides the primary input source for ExPR. Data is classified by ACS, LSM, device type, and by user-defined jobname workload groups. Subtype 4 and subtype 7 records are used. The input dataset containing HSC SMF records is identified in the JCL, with ExPR parameters providing controls for date range, SMF number, and workload group identification.

MVS System Log

Some events that are useful in diagnosing performance problems in a Nearline environment are not available through SMF-based reporting. An example of this is JES allocation recovery. ExPR uses the external writer SYSLOG file

generated by MVS from console messages for this information.

HSC PGMI Processing

The HSC Programmatic Interface (PGMI) is referenced through the EXPR started task that reads the HSC CDS, issues regular PGMI calls to the HSC, and directs its output to a collection file for input to the ExPR database. Collected data includes cell count statistics, which are not available through HSC SMF records.

Tape Catalog/HSC CDS Processing

Selected fields may be extracted from your CA-1 or TLMS tape catalog. This catalog is first cross-referenced with the HSC CDS, which contains an inventory of cartridges residing in the library. (Optionally, you may specify that the entire tape catalog be read without cross-reference against the CDS.)

Processing would normally take place on a daily "snapshot" basis to minimize the overhead of reading both the CDS and tape catalog. However, the ExPR database structure would allow you to record the tape catalog aging and utilization data on a more frequent basis if required.

Direct reads of the CDS are also undertaken through the ExPR started task to extract scratch status information.

Tape catalog processing is optional. If you do not wish to collect tape catalog data (or if this release of ExPR does not support your tape catalog system), ExPR can be used without collecting these records.

Review SYS1.PARMLIB Member SMFPRMnn

A general review of SMF record collection is recommended. You should look at member SMFPRMnn in SYS1.PARMLIB and refer to the IBM SMF manual, *System Management Facility* (GC28-1628-nn).

The MVS SMF, HSC SMF, and MVS RMF record types described above are all derived from the MVS SMF. RMF records are in fact SMF records in the range 70 through 79, and the HSC data is written to a user-defined SMF record.

The HSC SMF is usually assigned 255. However, this may not be the case and should be verified with the appropriate systems programmer.

The following SMF record types and subtypes are required by ExPR:

- SMF type 21
- RMF type 73
- RMF type 74
- SMF user record for HSC

SMF Housekeeping Procedures

Current housekeeping procedures should be examined to identify the best place and time to run the ExPR data collection and update procedures.

SMF, RMF, and the HSC SMF records are collected by MVS into a single system file named SYS1.MANx. This file will fill up on a regular basis and be automatically switched to the next file in sequence. Procedures will be established for downloading these files on a regular basis before the sequence

of files are all full. If the files are all full, data collection ceases.

The recommended approach is to extract the required record types using an IBM supplied SMF utility called IFASMFDP. Information can be found in the IBM document *System Management Facility* (GC28-1628-nn). In most installations similar processing takes place for other products. Where possible, ExPR processing should be integrated at this processing point to avoid reading the main SMF/RMF collection file unnecessarily.

Note: The installation download procedure usually moves the data to a daily GDG file which in turn will be moved into a weekly and then monthly GDG cycle. Theoretically, this entire system-wide file could be input to ExPR and only the required record types would be selected for processing by the product. However, this is very inefficient and would normally involve a huge number of cartridge mounts.

Task 2: Space Planning

This task helps you determine the probable DASD requirements at your site for ExPR datasets, including the ExPR database. These requirements are in addition to target and distribution libraries already allocated during the basic SMP/E installation.

ExPR Datasets

The following files are created by ExPR MVS:

- DATABASE (ExPR mainframe database file)
- FTEXTR (File transfer extract file)

- FTCNTL (File transfer control file)
- CONFIG (ExPR parameter file)
- CDSCONF (Generated configuration statements from CDS)
- PGMIDATA (Collection file for PGMI data)
- EXPORT (Consolidation file for exported (old) records)
- TMCHIST (CDS/tape catalog volume history file)
- TLMSRPT (report file for CA-TLMS installations only)

Additionally, a number of temporary files are allocated during various update processes. These are allocated to SYSDA; it is assumed that sufficient temporary DASD is available.

DASD space estimates for ExPR files are described in the sections below. Estimates are based on 3390 device types.

DATABASE

DASD requirements for the ExPR mainframe database or databases depend on the size of the Nearline system being reported on and other factors. The calculations below can be used to estimate ExPR mainframe database size requirements. Refer also to appendix D, *Database Layout*, for additional information about the contents of each record type.

Note: A cylinder of 3380 will hold approximately 700K of ExPR database, and a 3390 cylinder will hold approximately 800K. If, at this point, you are uncertain about how much data you will collect, you could use the supplied model defined in SAMPLIB member DBDEF without changing the allocation size.

Database Size Calculation Worksheet

RECORD TYPE 0   HSC SMF Hourly Mount Statistics

(# hours)* x (# LSMs) x (1 + # of workload groups) x (record size)**

_____ x _____ x _____ x <u>145</u> = _____

RECORD TYPE 1   Hourly Control Unit Activity

(# hours)* x (# LSMs) x (record size)**

_____ x _____ x <u>66</u> = _____

RECORD TYPE 2   Hourly Tape Error Statistics

(# hours)* x (record size)**

_____ x <u>64</u> = _____

RECORD TYPE 3   Hourly Global Mounts

(# hours)* x (record size)**

_____ x <u>24</u> = _____

 * The period of time collected data will be kept in the database (i.e., six months is 4368 hours, one month is 720 hours).

** These record sizes include the 16-byte VSAM key at the start of each record.

RECORD TYPE 4   LSM Volume Aging Statistics

                       (#days)* x (#LSMs) x (record size)**

                       _____ x _____ x <u>84</u> = _____

RECORD TYPE 5   Dataset Group Utilization Statistics

                       (#days)* x (#dataset workload groups) x (record size)**

                       _____ x _____ x <u>68</u> = _____

RECORD TYPE 6   Dataset Group Aging Statistics

                       (#days)* x (#dataset workload groups) x (record size)**

                       _____ x _____ x <u>68</u> = _____

RECORD TYPE 7   LSM Volume Utilization Statistics

                       (#days)* x (#LSMs) x (record size)**

                       _____ x _____ x <u>84</u> = _____

   * The period of time collected data will be kept in the database (i.e., six months is 182 days, one month is 30 days).

 ** These record sizes include the 16-byte VSAM key at the start of each record.

RECORD TYPE 8   Hourly Scratch and Free Cell Statistics

(#hours)* x (#LSMs) x (record size)**

_____ x _____ x <u>48</u> = _____

RECORD TYPE 9   Hourly Channel Group Statistics

(#hours)* x (#groups) x (record size)**

_____ x _____ x <u>24</u> = _____

RECORD TYPE 10   Hourly Allocation Recovery Per ACS/LSM

(#hours)* x (#LSMs) x (record size)**

_____ x _____ x <u>56</u> = _____

RECORD TYPE 11   Tape Error Event Detail

(#days)* x (estimated tape errors per day) x (record size)**

_____ x _____ x <u>59</u> = _____

  * The period of time collected data will be kept in the database (i.e., six months is 4368 hours, one month is 720 hours, six months is 182 days, one month is 30 days).

 ** These record sizes include the 16-byte VSAM key at the start of each record.

Chapter 2: Installation and Customization   **39**

RECORD TYPE 12   Demand Enter Event Detail

(#days)* x (estimated demand enters per day) x (record size)**

_____ x _____ x <u>39</u> = _____

RECORD TYPE 13   Hourly Allocation Recovery For Non-ACS Devices

(#hours)* x (#LSMs) x (record size)**

_____ x _____ x <u>56</u> = _____

 * The period of time collected data will be kept in the database (i.e., six months is 4368 hours, one month is 720 hours, six months is 182 days, one month is 30 days).

 ** These record sizes include the 16-byte VSAM key at the start of each record.

Total of Database Record Calculations

RECORD TYPE 0: _____

RECORD TYPE 1: _____

RECORD TYPE 2: _____

RECORD TYPE 3: _____

RECORD TYPE 4: _____

RECORD TYPE 5: _____

RECORD TYPE 6: _____

RECORD TYPE 7: _____

RECORD TYPE 8: _____

RECORD TYPE 9: _____

RECORD TYPE 10: _____

RECORD TYPE 11: _____

RECORD TYPE 12: _____

RECORD TYPE 13: _____

=====================

Sum of Calculations for All Database Record Types: _____

FTEXTR

The FTEXTR file contains information that is extracted from the ExPR database and formatted for user file transfer to a PC. The size of the file is dependent on the period of data to be extracted and the size of the database.

This file is formatted as a comma-separated ASCII file, and as such may be up to 50% larger than the corresponding database segment where its information is extracted from.

The FTEXTR file can be deleted between extract runs if necessary due to DASD space considerations. However, successful updating of the PC database should be verified first if extract files are to be deleted.

FTCNTL

The FTCNTL file contains file transfer control information, specifically the last-extract date stamp. This file requires minimal DASD; estimate one track.

CONFIG

The CONFIG file contains configuration parameters ExPR requires to perform database build and update procedures, and to collect appropriate run data. This file requires minimal DASD; estimate two tracks.

CDSCONF

This file is required by the automatic configuration-generation routines. Allocate at least five tracks for this file if you are using the automatic configuration feature.

PGMIDATA

The PGMIDATA file contains PGMI data collected by the ExPR started task. The size of the file depends on the number of days between database updates (up to the maximum number of days set by the ExPR oldest record parameter MAXAGE). Estimate two tracks per day for average-to-medium-size sites with up to five LSMs.

EXPORT

The EXPORT dataset is used to hold old database records that are not required for immediate access and can therefore be archived in some manner. The format of this dataset is:

- DSORG   PS

- RECFM   VB

- LRECL   500

- BLKSIZE   504 or greater

The EXPORT file can be imported back into the ExPR KSDS database with the IDCAMS REPRO command. The size of the dataset can only be determined by knowing how many records are to be exported.

TMCHIST

This file is created each time you run the TAPECAT UPDATE function. It contains information on each volume in the CDS/tape catalog. Task 4 in this chapter defines a GDG for TMCHIST.

TLMSRPT

This file is only required by CA-TLMS installations. It is an FBA/133-byte file. You should block this dataset with the maximum one-half track blocksize. As a guide to sizing this file, a 380-page TLMS report requires 90 tracks (3390) when one-half track blocking is used.

Total ExPR DASD Requirement

DATABASE:	_____
FTEXTR:	_____
FTCNTL:	_____
CONFIG:	_____
CDSCONF:	_____
PGMIDATA:	_____
EXPORT:	_____
TMCHIST	_____
TLMSRPT:	_____
Total ExPR DASD Space Requirement:	_____

Temporary SYSDA Work Files

The EXPRPROC member includes work DDs for four pairs of
temporary files which are used by the TAPECAT function
and its called sorts. Larger installations may need to increase
the space allocation for these files, which should always be
blocked at the maximum half-track capacity for your DASD
type. CDSTEMP and CDSTEMPX are also used as workfiles
by other ExPR functions.

The table below describes the characteristics and usage of
these files. Records-per-track are for 3380/3390, assuming
maximum half-track blocking.

Temporary SYSDA Work Files

File	Type / Size	Description	Records-per-Track
CDSTEMP CDSTEMPX	FB 40 bytes	One record per volume defined in the HSC CDS.	1200 (3380) 1400 (3390)
TMCTEMP TMCTEMPX	VBS 167 bytes. LRECL must be 32756.	One record per volume in the tape catalog, plus 57 bytes per secondary dataset in the tape catalog.	210 (3380) 250 (3390)
TMCDSN TMCDSNX	FB 170 bytes	One record per secondary dataset in the tape catalog.	270 (3380) 320 (3390)
TMCVOL TMCVOLX	FB 340 bytes	One record per volume in the tape catalog.	136 (3380) 163 (3390)
TMCHISTX	FB 110 bytes	One record per dataset matched between the tape catalog and the HSC CDS. Used in conjunction with the TMCHIST GDG file.	416 (3380) 510 (3390)

Task 3: Review Hardware and Software Prerequisites

This task helps you ensure that all required hardware and software components are in place before you start configuring ExPR.

Generally, the hardware prerequisites for the MVS component of ExPR are very basic. The software must run on an MVS/XA or MVS/ESA system. The MVS system does not necessarily need to be the system with the Nearline configuration attached, except for the PGMI task, which must run on the HSC-resident system. Additionally, automatic Config would require read access to the CDS dataset. All other hardware conditions are satisfied by the presence of a Nearline configuration.

Required software for the MVS environment includes the MVS operating system, StorageTek Host Software Component (HSC) software, a SORT product, RMF, and optionally a tape management system. Refer to the separate document *ExPR MVS Installation Notes* for a list of the supported or required versions of these products for this release of ExPR.

Task 4: Define ExPR Datasets and Tailor EXPRPROC Procedure

This task allocates the datasets that are required by ExPR and initializes the PGMI collection file. In this task, you will customize sample control cards with site data you calculated in *Task 2, Space Planning*.

Data Collection and Parameter Files

ExPR requires the following data collection and parameter files:

- File transfer extract dataset
- File transfer control dataset
- CONFIG dataset
- PGMI collection file (VSAM KSDS)
- ExPR consolidation EXPORT file
- TMCHIST pattern DSCB (for non-SMS installations only)
- TMCHIST GDG definition (IDCAMS)
- Optional TLMSRPT file for CA-TLMS users
- Optional CDSCONF automatic Config generation file

```
ACTION:

⇒  Refer to Task 2: Space Planning in this chapter to
   determine appropriate sizes for ExPR datasets.

⇒  Allocate required datasets using SAMPLIB member
   DSNDEF.  Non-SMS installations must be sure that
   TMCHIST pattern DSCB is allocated to the correct
   catalog volume.  SMS installations must remove this DD
   statement.

⇒  Initialize the PGMI collection file, also using DSNDEF.

⇒  CA-TLMS users must amend DSNDEF to allocate the
   TLMSRPT file.

⇒  Tailor SAMPLIB member EXPRPROC with your dataset
   names; catalog to your MVS PROCLIB for subsequent
   use by EXPRJOB.  Be sure that the correct HSC steplib is
   included, and be sure to include the CAPS command
   when tailoring SAMPLIB members.
```

Note: Generally, from this point on, you will be using the
tailored procedure EXPRPROC and the supplied job
EXPRJOB when running the ExPR batch functions. These
will be used in conjunction with the various SAMPLIB
members.

Task 5: Define the ExPR Database(s)

This task helps you determine if a single or multiple ExPR
databases are suitable for your site's reporting requirements,
and then defines the database or databases. In this task, you
will customize sample control cards with site data you

calculated in *Task 2, Space Planning*. You will also provide database names in this task.

Database Overview

A single ExPR mainframe database can be maintained that contains data relating to all ACS environments, or multiple mainframe databases can be defined. How you define the ExPR MVS database for your site depends on a number of factors, including the size of the Nearline configuration and the number of geographic locations involved. As a general rule, if in doubt, create a single database. This is especially true for small to medium installations with five or fewer LSMs.

For larger installations, it may be desirable to establish a unique database for each individual ACS, or for a logical group of ACSs determined by application, host system attachment, or geographic location. In multi-site installations, the logistics of updating the ExPR database may be streamlined by establishing a database for each site. Similarly, if different ACSs in the organization have unique dynamics and it is appropriate to create separate CONFIG files to uniquely handle ExPR features such as exception processing, separate databases should be established.

Database File Structure

The ExPR database is a VSAM KSDS file. It contains a number of different record types; see appendix D, *Database Layout* for specifics. Duplicate data handling and update of existing records are supported.

Task 6: Initialize the ExPR Database

This task initializes the database(s) that were defined in the previous task. In this task, you will run the sample initialize with database name(s) that were supplied in *Task 5, Define the ExPR Database(s)*.

The mainframe database must be initialized as a VSAM KSDS file before data collection can begin.

ACTION:

⇒ Initialize the database using SAMPLIB members DBINIT and EXPRJOB.

⇒ Refer to appendix B, *SYSIN Parameters* for information about runtime parameters in the SAMPLIB.

⇒ Be sure to include the CAPS command when tailoring SAMPLIB members.

Note: You must add the DD DUMMY for EXPR.CONFIG, as noted in DBINIT. Remove it after the initialize function.

Task 7: Review ExPR Parameters and Configuration Options

The ExPR environment and configuration are established for your site using the SAMPLIB members that are prefixed "@". The "@" SAMPLIB members contain parameters that will be built into the ExPR CONFIG file. These members and their functions are described in the table below.

"@" SAMPLIB Member	Description
@CONFIG	Defines hardware specification
@HSC	Defines SMF number
@PGMI	Defines started task controls
@CDS	Defines HSC CDS and DDname parameters
@THRSHLD	Defines exception threshold values
@MAP	Defines jobname workload groups
@TAPCAT	Defines tape processing controls
@DSMAP	Defines dataset workload groups

Note: Some of the site configuration information in @CONFIG and @HSC can be supplied automatically by ExPR from the HSC CDS if the automatic Config generator feature is used. This feature is described below.

Before proceeding with the remaining tasks in this chapter, you should determine if you will use the ExPR automatic Config generator feature or input your site data manually. You should also review the "@" SAMPLIB members to become familiar with their contents. Examples of fully completed configurations are in SAMPLIB members CONFSMPA (for automatically generated configurations) and CONFSMPM (for manually generated configurations).

You should also review the following appendixes in this guide:

- Appendix A, *CONFIG Parameters.* Configuration parameters that are fixed values (or unlikely to change on a run-to-run basis) are specified by the CONFIG DD dataset.

- Appendix B, *SYSIN Parameters.* Runtime parameters that are specific to each job are specified by the SYSIN DD dataset.

Note: It is important to understand that ExPR does not execute *from* the SAMPLIB library, and that only those parameters copied *to* the CONFIG file or automatically generated by ExPR will become effective.

Determining Whether to Use the Automatic Config Generator

There are two available methods for building the ExPR CONFIG dataset, manually or with assistance from the ExPR automatic Config generator feature.

The automatic Config generator feature, which is activated by default, collects your site's hardware configuration information from the HSC CDS and inputs it to the CONFIG process. However, not all site data can be collected automatically, so smaller sites (with five or fewer LSMs) may in fact find it more efficient to just enter the data manually using ExPR CONFIG parameter statements.

Data that is automatically collected is written to the CDSCONF dataset. When ExPR interprets parameters, it looks first to the CDSCONF dataset and then to the CONFIG dataset and finally to the SYSIN DD.

How the Automatic Config Generator Works

The automatic Config generator feature is invoked automatically by ExPR when any ExPR function is run. No user command specification is required.

This function is controlled by JCL in EXPRPROC:

- If the CDSCONF DD statement is present, the automatic CONFIG statements will be written and used when ExPR looks for CONFIG parameters.

- If the CDSCONF DD statement is commented out, no automatic Config statements can be written. ExPR will not return an error for this condition. It will look instead to the CONFIG dataset for site configuration data.

Input to the automatic Config function is from the HSC CDS that is specified in the CDSCOPY DD statement, also found in EXPRPROC. This can be either the live CDS or a copy you have made. Remember that ExPR will look at this file and attempt to update the CDSCONF file every time it runs (this represents a very low overhead) so this should be the live CDS dataset if you want all changes to your CDS to be immediately reflected in ExPR reporting. If you wish to control when CDS changes are seen in ExPR reports (this can be important in maintaining consistency for historical reporting), then use a copy of the CDS.

Determining Channel Path Group Requirements

ExPR reports provide information about channel path activity related to your Nearline configuration. To use this feature effectively, you need to map your channel paths into logical groups that reflect how they are organized in your

Chapter 2: Installation and Customization   **53**

configuration. Each channel path group may contain one or more channel paths (and/or a range of channel paths) whose activity should be reported on as a single unit. These groups are defined to ExPR with the CONFIG DD statement CHPATH in the next task.

Determining Device-Type Requirements

During SMF/RMF processing, ExPR will collect mount counts, mount times, mount robotics breakdown, and drive utilization statistics on a per LSM, per hour basis. All devices attached to a given LSM will be accounted for in a single hourly record.

Until recently, this has been adequate for most customers. However, it is now possible to have mixed devices on a single LSM (4480/4490 and 9490). Your site may wish to analyze the performance and utilization of the newer faster devices separately from the main device pool. SD-3 Redwood devices and HSC 2.0.1 further raises the importance of device-type reporting and maximization of your hardware investment.

ExPR addresses this requirement with the following features and control statements:

- DEVICE-TYPE DEFAULT allows you to define the majority device type across all LSMs.

- DEVICE TYPE assigns a device type to a specific device on a single MVS host.

- APPEND DEVICE TYPE adds a device type to a previously defined device (i.e., from Auto-Config).

- REPORT-OPTIONS causes the MOUNTS, MOUNTS-DETAIL, and UTILIZATION reports to print summary lines for individual device types.

If you have a single device type in your configuration, or if you do not wish to break down the device-type statistics, then it is not necessary to specify any of the above parameters.

If you choose to use device-type analysis, it is important to understand the order of processing:

1. At the commencement of ExPR control statement processing, the default device type is set to NONE.

2. A DEVICE-TYPE DEFAULT statement, when encountered, will set the new default device type for all MVS systems, until another DEVICE-TYPE DEFAULT statement is processed.

3. When a DEVICE statement without a TYPE parameter is processed, it will receive the current DEVICE-TYPE DEFAULT setting.

4. When all AUTO-Config, CONFIG, and SYSIN statements have been processed, the final DEVICE-TYPE DEFAULT setting will be applied to all devices in all MVS systems that do not have a specific device type assigned by DEVICE TYPE.

5. An APPEND DEVICE TYPE statement will always over-ride any previous setting for a specific device.

6. A summary of the final device-type settings will be printed if PARM='OPTIONS(+A)' is specified on the ExPR EXEC statement.

Large users of mixed device types may choose to take a copy of the generated CDSCONF auto-Config file and add the device-type values to it. This could then be merged into the CONFIG file, and the automatic Config feature could be disabled by removing DD CDSCONF. In these circumstances, you must remember to reactivate automatic Config if the Nearline configuration is amended.

Task 8: Define Nearline Configuration to ExPR

This task identifies your Nearline configuration to ExPR. Instructions are included here for performing this task with the automatic Config generator feature or entering the data manually.

ExPR requires the following information about your site:

- For each MVS host system, the host system identifier

- For each ACS, the decimal ACS identifier found in HSC SMF records, and optionally the ACS text name to appear in report headings

- For each LSM in each ACS, the decimal LSM identifier and optionally the LSM type (4400, Wolfcreek, etc.) and the LSM text name to appear in report headings

- For each LSM in each ACS, the hexadecimal device address for each tape drive and its device type

- Channel path IDs and channel path group definitions to logically group channel path IDs for reporting purposes.

Using the ExPR Automatic Config Generation Feature

Note: Skip this subtask if you intend to enter site information manually.

When the automatic Config generation feature is invoked (the default), ExPR will automatically collect the following site data from the HSC CDS and input it to the CDSCONF dataset:

- The HSC SMF record type number

- The MVS host system identifier for each host

- The ACS identifier for each ACS

- The LSM identifier for each LSM

- The device address for each tape drive

The following site data *cannot* be collected automatically and therefore must still be entered manually in the CONFIG DD dataset:

- The channel path ID and optional channel path group for each channel path

- The optional LSM type (4400, Wolfcreek, etc.) for each LSM

- The optional ACS, LSM, and customer names that appear in report headings

- The device type (4480, Redwood, etc.) for each tape drive

The optional customer name, ACS name, ACS type, and LSM name appear in report headers for ease of identification;

Chapter 2: Installation and Customization   **57**

these optional parameters are not used by ExPR for any direct processing.

Parameters that identify device types, channel paths, company name, and ACS/LSM names are shown in the table below.

Site Data	CONFIG DD Statements
Channel Paths	CHPATH ID(*xx-xx*) GROUP(*nnn*) SYSID(*mvshost*)
Company Name	CUSTOMER NAME('*text*')
ACS Names	APPEND ACS ID(*aaa*) NAME('*text*')
LSM Names	APPEND LSM ID(*aaa ll*) NAME('*text*') TYPE('*t*')
Device Types	DEVICE-TYPE DEFAULT(*tttt*)
Tape Drives	APPEND DEVICE ID(*dddd*) LSM(*aaa ll*) TYPE(*tttt*) SYSID(*mvshost*)

SAMPLIB member @CONFIG contains sample control cards for each of these functions. You must comment out those statements that do not apply to your installation.

> **ACTION:**
>
> ⇒ No action is required to generate the CDSCONF dataset. The automatic Config feature is invoked automatically when any ExPR function is run.
>
> ⇒ Define Nearline configuration information for your site that cannot be collected automatically using the following CONFIG statements in the SAMPLIB member @CONFIG: APPEND, CUSTOMER, CHPATH, and DEVICE-TYPE. Refer to appendix A, *CONFIG Parameters* for information about these parameters.
>
> ⇒ Be sure to include the CAPS command when tailoring SAMPLIB members.
>
> ⇒ Print the SYSOUT listing from the auto-Config and CONFIG files as described in chapter 7, *Mainframe Performance Reports* under *Printing SYSOUT Information.*

Entering CONFIG Data Manually

Note: Skip this subtask if you intend to use the automatic Config feature.

If you are not using the automatic Config generation feature, you must:

1. Comment out the CDSCONF DD statement in EXPRPROC to prevent ExPR from reading CONFIG statements from this source.

2. Customize sample control cards in SAMPLIB member @CONFIG with site data about your installation's configuration.

Parameters that identify each MVS host, ACS, LSM, tape drive device, and channel path, as shown in the table below.

Site Data	CONFIG DD Statements
MVS Hosts	HOST SYSID(*mvshost*)
ACSs	ACS ID(*aaa*) NAME('*text*')
LSMs	LSM ID(*aaa ll*) NAME('*text*') TYPE('*t*')
Tape Drives	DEVICE ID(*dddd*) LSM(*aaa ll*) TYPE (*tttt*) SYSID(*mvshost*)
Device Types	DEVICE-TYPE DEFAULT(*tttt*)
Channel Paths	CHPATH ID(*xx-xx*) GROUP(*nnn*) SYSID(*mvshost*)
Company Name	CUSTOMER NAME('*text*')

SAMPLIB member @CONFIG contains sample control cards for each of these functions.

```
┌─────────────────────────────────────────────────┐
│ ACTION:                                         │
│                                                 │
│ ⇒  Comment out the CDSCONF DD statement in      │
│    EXPRPROC to prevent ExPR from reading CONFIG │
│    statements from this source.                 │
│                                                 │
│ ⇒  Define Nearline configuration information for your site │
│    to ExPR using SAMPLIB member @CONFIG.        │
│                                                 │
│ ⇒  Refer to appendix A, CONFIG Parameters for   │
│    information about the CONFIG dataset parameters. │
│                                                 │
│ ⇒  Be sure to include the CAPS command when tailoring │
│    SAMPLIB members.                             │
│                                                 │
│ ⇒  Print the SYSOUT listing from CONFIG DD, as  │
│    described in chapter 7, Mainframe Performance Reports. │
└─────────────────────────────────────────────────┘
```

Task 9: Install the ExPR Started Task

This task defines the EXPRPGMI started task and establishes
it in your MVS startup parameters. In this task, you will copy
sample control cards to your selected PROCLIB library and
customize them to define started task runtime parameters for
your site.

Started Task Overview

The EXPRPGMI started task is required to collect data for
scratch count and free cell information on an hourly basis.
This task is required because these statistics are not provided

in the HSC SMF record. The started task, EXPRPGMI, is included in the ExPR SAMPLIB.

EXPRPGMI issues regular PGMI calls to the HSC and also reads the HSC CDS. Output is then written to a collection file, EXPR.PGMIDATA, which is subsequently input to the database update process as a separate step. The size of the collection file is controlled by applying an age limit to the data contained in it.

Housekeeping procedures are performed automatically by the started task to maintain the collection file at the correct size. Record keys are checked for age each time the task is started and also after database update processing.

Started Task Parameters

The started task is run from the SYS1.PROCLIB library (or a site-specified concatenated PROCLIB). Optionally, it can be started automatically at IPL by a start command in SYS1.PARMLIB member IEACMDnn (for example, 'COM=START EXPRPGMI').

The CONFIG statement PGMI defines the frequency that PGMI calls are made and the age of the data that is retained in the collection file:

- PGMI call frequency: An EVENT parameter controls the rate at which records are generated by ExPR (every minute, every 15 minutes, or every hour).

- Age specification: The started task checks records in the collection file against the age parameter (MAXAGE) and deletes outdated records. Thus, the age limit should cover the longest period between database update runs (seven days is recommended). This ensures that records are retained in the collection file long enough to be included in the database update process without requiring excess DASD allocation.

Data that is written to EXPR.PGMIDATA is selected and input for database update processing by a single process.

Notes:

1. The EXPR.PGMIDATA collection dataset referenced by DDname PGMIDATA must be defined and initialized prior to running the started task. This process was described in *Task 4: Define ExPR Datasets and Tailor EXPRPROC Procedure* in this chapter.

2. The EXPRPGMI started task references HSC module SLSXCAL. You must ensure that the HSC load library where this module resides is either included in the link list or is specified on the STEPLIB DD statement for the EXPRPGMI started task.

3. The SLSXCAL module must be the same version/release as the executing HSC started task.

> **ACTION:**
>
> ⇒ Copy the started task EXPRPGMI from SAMPLIB to your selected PROCLIB library.
>
> ⇒ Establish started task parameters for your site using SAMPLIB member @PGMI.
>
> ⇒ Refer to appendix A, *CONFIG Parameters* for information about CONFIG dataset parameters, and to appendix B, *SYSIN Parameters* for information about runtime parameters.
>
> ⇒ Be sure to include the CAPS command when tailoring SAMPLIB members.

Task 10: Specify HSC SMF Number to ExPR

This task identifies the HSC SMF number for your site to ExPR. In this task, you will customize sample control cards with the HSC SMF number for your site.

Note: If you are using the ExPR automatic configuration generator feature, the HSC SMF number is supplied automatically and you do not need to perform this task. However, you must comment out the statement in SAMPLIB member @HSC.

The HSC SMF number is usually assigned 255. However, it is a site-variable number and thus should be verified with the appropriate systems programmer.

Task 11: Define the Control Dataset to ExPR

This task identifies the DDname of the HSC Control Dataset
(CDS) that ExPR is to reference. In this task, you will
customize sample control cards by specifying the DDname
that points to the CDS. Typically, this would be one of the
three images of the CDS supported by HSC. ExPR's default
value is the primary image, CDSPRIM, which would
normally point to the primary CDS dataset.

Task 12: Create the ExPR CONFIG Dataset

This task copies the customized sample configuration options
from previous tasks into the live CONFIG dataset. In this
task, you will customize parameters provided in the sample
control cards to copy your edited SAMPLIB entries into the
live CONFIG dataset. This task is essential for a successful

installation. ExPR does not execute *from* the SAMPLIB
library, so only those parameters copied *to* the CONFIG file
(or generated automatically by ExPR) will become effective.

Note: Before you perform this task, you may want to first set
up the optional exception processing, jobname workload
group, and tape catalog processing features described in
chapters 4, 5, and 6 respectively. It does not matter when or
if you set up those optional features, but if you do, you must
return to this task to copy any SAMPLIB members you
customize into the live CONFIG dataset. There are four
SAMPLIB members relating to these optional features -
@DSMAP, @MAP, @TAPCAT and @THRSHLD. The
sample statements are commented out, so it is not necessary
to edit these members if you do not wish to utilize the
associated features of ExPR.

ACTION:

⇒   Create the CONFIG dataset for your site using SAMPLIB
    member CONFIG.

⇒   Be sure to include the CAPS command when tailoring
    SAMPLIB members.

Note: Examples of fully completed configurations can be
found in SAMPLIB members CONFSMPA (for automatically
generated configurations) and CONFSMPM (for manually
generated configurations).

Task 13: Installation Verification

At this point, the basic ExPR system is installed and configured for use in the MVS environment at your site. The best way to test and verify your installation is to update the database with your site data and begin generating mainframe tabular reports.

> **ACTION:**
>
> ⇒ Build the database using SAMPLIB member SMFUPDAT. Refer to chapter 3, *Managing the ExPR Database* for more information about the database update process.
>
> ⇒ Generate mainframe performance reports using SAMPLIB members MFREPS and EXPRJOB. Refer to chapter 7, *Mainframe Performance Reports* for more information about the mainframe reporting process.
>
> ⇒ Be sure to include the CAPS command when tailoring SAMPLIB members.
>
> ⇒ Refer to chapters 4, 5, and 6 respectively for information about configuring the optional exception processing, jobname workload groups, and tape catalog processing features.

Chapter 3

Managing the ExPR Database

Overview

This chapter describes required steps for selecting data to update the database with, various database processes, and maintenance procedures required to maintain the ExPR database, including verification and consolidation. Examples are included for reference.

Note: Before proceeding with this chapter, be sure you have successfully completed all tasks in chapter 2, *Installation and Customization*.

Input Data Review

ExPR allows data input from a number of MVS and HSC sources. These are:

- SMF type 21 (tape error statistics)

- RMF (Nearline channel and CU statistics and SSCH counts)

- HSC SMF (Nearline activity)

- PGMI (Nearline cell information)

- CDS (scratch and free cell statistics)
- Tape Catalog (tape usage statistics)
- MVS SYSLOG (console messages for allocation recovery analysis)

Not all data sources follow the same selection path. Figure 3-1 below illustrates the various update paths that are taken by different data sources.

- SMF type data requires pre-selection and sorting

- PGMI and CDS data is first written to an intermediate collection file

- Tape Catalog data is updated directly to the ExPR database

- MVS SYSLOG data is updated directly to the ExPR database
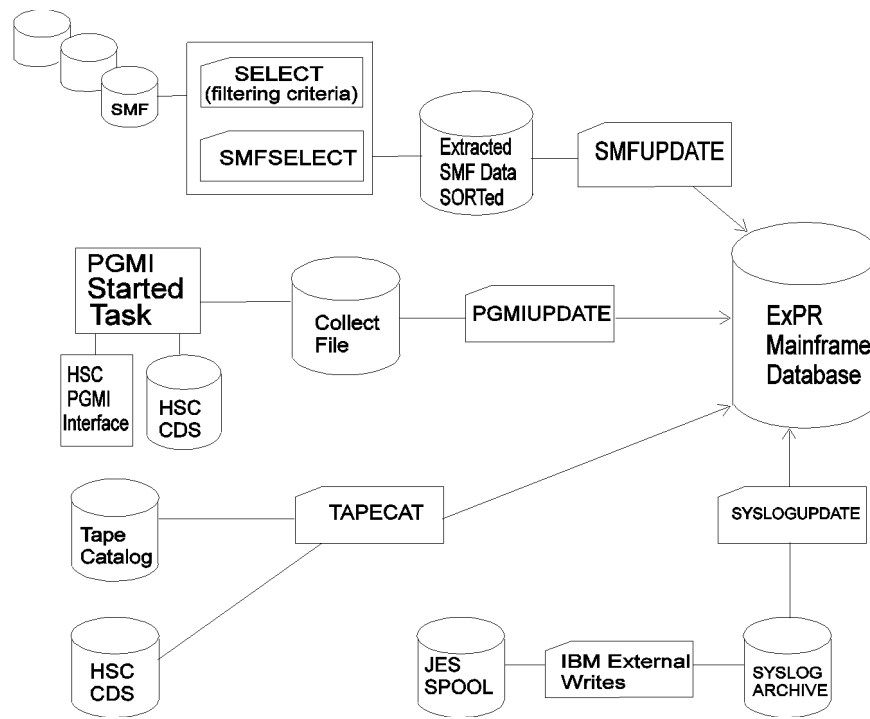
Figure 3-1: ExPR Data Update Paths

Database Processes

The database should be updated as part of the same housekeeping procedure as SMF, RMF, and HSC SMF processing. The database update process can be selective with ExPR parameters providing controls for the record types to be input. Update (new date ranges) and replace (overwrite existing records) processing are supported.

Each of the ExPR SYSIN parameters that relate to database processing are described below. The syntax for these parameters is described in appendix B, *SYSIN Parameters*.

Note: If you are processing data that is more than 30 days old, you must include a PERIOD FROM/TO range; the default is from 30 days previous to the current date if a period is not specified.

SELECT

SELECT defines filtering criteria for the SMF record types to be extracted from the general SMF housekeeping file maintained by the installation. SELECT subparameters include HSC, SMF21, and RMF. This parameter should be used when not all SMF type records handled by ExPR are required. The SELECT command must precede the SMFSELECT function.

PERFORM SMFSELECT

Extracts the required ExPR SMF record types (subject to any SELECT filtering criteria) and writes them to a sequential file for input to the SORT utility.

PERFORM SMFUPDATE

Updates the ExPR VSAM database file with the sorted output from SMFSELECT processing.

PERFORM PGMIUPDATE

Updates the ExPR database with data records that are held in the collection file that is written to by the PGMI started task processes.

PERFORM SYSLOGUPDATE

Update the ExPR database with console message records from the MVS SYSLOG. Selective message numbers are scanned and the date and timestamp associated with each message is used to build a record containing the number of allocation recovery events and the duration of each event.

PERFORM TAPECAT

Performs tape catalog processing and updates the ExPR database with information extracted from the tape catalog and CDS. This information relates to tape ages and utilization within the ACS. Tape catalog processing is detailed in chapter 6, *Tape Catalog Processing and Reports*.

PERFORM CONVERTDB

Allows users of ExPR 1.0 to convert their existing database for use with ExPR 1.1. Major database structure changes have occurred in ExPR 1.1; consequently, it is not possible to run the ExPR 1.1 software against the older database.

> **ACTION:**
>
> ⇒ Update the database using SAMPLIB members
>   SMFUPDAT, PGMIUPDT, SYSLOG, and TAPECAT1
>   or TAPECATT.
>
> ⇒ Refer to appendix B, *SYSIN Parameters* for information
>   about runtime parameters that appear in the SAMPLIB.
>
> ⇒ Refer to chapter 6, *Tape Catalog Processing and
>   Reports*, for more detail about the TAPECAT function.
>
> ⇒ Be sure to include the CAPS command when tailoring
>   SAMPLIB members.

SYSIN Examples

The following examples illustrate how SYSIN parameters
described above are used to update the database.

Example 1: Select HSC SMF Records Only and Update the Database

```
 PERIOD FROM(yyyyddd) TO(yyyyddd) ;
INPUT   DDNAME(INDATA1 INDATA2) ;
SELECT TYPE(HSC) SYSID(mvshost);
PERFORM
SMFSELECT ;

              --- SORT step goes here ---

PERIOD FROM(yyyyddd) TO(yyyyddd) ;
INPUT   DDNAME(SORTDAT)  ;
PERFORM
SMFUPDATE  ;
```

Note: If all SMF data is required, do not specify a SELECT statement. If SMF data from a specific system is required, specify SELECT SYSID without the TYPE(HSC) parameter.

The SMFUPDATE function can be replaced by the SMFREPORT function if you want to analyze SMF data for exception events without updating the ExPR database. See chapter 4, *Exception Processing and Reports* for further details and a sample of the SMF Exception Events report.

All SMF type data needs to be sorted before being input to the database update process. If this is done already at your installation, the sort steps provided in the control card examples found in the ExPR SAMPLIB can be omitted. In the above example of SYSIN streams, a sort step would be performed between SMFSELECT and SMFUPDATE processing.

Example 2: Select PGMI Data and Update the Database

```
PERIOD FROM(yyyyddd) TO(yyyyddd) ;
PERFORM
PGMIUPDATE ;
```

Data that is collected by the EXPRPGMI started task is written to the PGMIDATA collection file. The database is then updated with this collected data by the SYSIN parameter PERFORM PGMIUPDATE.

A report is generated automatically by this process. The report (titled "Update Database from STC Records") displays totals for records read and records added to the database by the PGMIUPDATE process. Additionally, hourly averages are printed for total cells, free cells, scratches, and cleaning cartridges.

The report shown below is generated during PGMIUPDATE processing.

1995.353   S T O R A G E T E K E X P R 1 . 1   XYZ COMPANY   PAGE   1
17:03:11

UPDATE DATABASE FROM STC RECORDS

| | | | | NO OF INTER | TOTAL | FREE | - - A V E R A G E S - - SCRAT | CLEAN |
DATE	HOUR	SYS	ACSLSM	VALS	CELLS	CELLS	CHES	CARTS
1994.340	01.00	CPUA	000 00	7	5800	256	384	8
1994.340	01.00	CPUA	000 01	7	5800	256	384	8
1994.340	01.00	CPUA	000 02	7	5800	256	384	8
1994.340	12.00	CPUA	000 00	4	5800	256	384	8
1994.340	12.00	CPUA	000 01	4	5800	256	384	8
1994.340	12.00	CPUA	000 02	4	5800	256	384	8
1994.340	13.00	CPUA	000 00	4	5800	304	640	8
1994.340	13.00	CPUA	000 01	4	5800	304	640	8
1994.340	13.00	CPUA	000 02	4	5800	304	640	8
1994.340	14.00	CPUA	000 00	4	5800	243	291	6
1994.340	14.00	CPUA	000 01	4	5800	243	291	6
1994.340	14.00	CPUA	000 02	4	5800	243	291	6
1994.340	15.00	CPUA	000 00	4	5800	273	801	6
1994.340	15.00	CPUA	000 01	4	5800	273	801	6
1994.340	15.00	CPUA	000 02	4	5800	273	801	6
1994.340	16.00	CPUA	000 00	4	5800	200	255	6
1994.340	16.00	CPUA	000 01	4	5800	200	255	6
1994.340	16.00	CPUA	000 02	4	5800	200	255	6
1994.340	17.00	CPUA	000 00	4	5800	153	528	6
1994.340	17.00	CPUA	000 01	4	5800	153	528	6
1994.340	17.00	CPUA	000 02	4	5800	153	528	6
1994.341	16.00	CPUA	000 00	3	0	0	242	9
1994.341	16.00	CPUA	000 01	3	0	60	306	8
1994.341	16.00	CPUA	000 02	3	0	170	332	4
1994.341	17.00	CPUA	000 00	4	0	0	242	9
1994.341	17.00	CPUA	000 01	3	0	60	306	8
1994.341	17.00	CPUA	000 02	3	0	170	332	4
1994.341	18.00	CPUA	000 00	12	0	0	242	9
1994.341	18.00	CPUA	000 01	12	0	60	306	8
1994.341	18.00	CPUA	000 02	12	0	170	332	4
1994.342	10.00	CPUA	000 00	3	0	0	242	9
1994.342	10.00	CPUA	000 01	3	0	60	306	8
1994.342	10.00	CPUA	000 02	3	0	170	332	4
1994.342	11.00	CPUA	000 00	60	0	0	242	9
1994.342	11.00	CPUA	000 01	60	0	60	306	8
1994.342	11.00	CPUA	000 02	60	0	170	332	4
1994.342	12.00	CPUA	000 02	18	0	170	332	4
1994.342	12.00	CPUA	000 01	18	0	60	306	8
1994.342	12.00	CPUA	000 00	18	0	0	242	9

COLLECTION FILE INPUT:
 RECORDS READ...............................511
DATABASE OUTPUT ANALYSIS:
 RECORDS ADDED.................................0
 RECORDS REPLACED........................39
 UPDATE FROM STC DATA COMPLETED

Chapter 3: Managing the ExPR Database   **77**

Example 3: Select MVS SYSLOG Data and Update the Database

The SYSLOGUPDATE control statement below is for a JES2 system. JES3 users should refer to the next page for a JES3 example.

```
//EXPR.SYSLOG DD DISP=SHR,DSN=user.syslog
//EXPR.SYSIN DD *
  PERFORM SYSLOGUPDATE ;
 /*
```

PERFORM SYSLOGUPDATE selects SYSLOG data and updates the database. The SYSLOG input is located by DDNAME SYSLOG. This is a fixed value and should not be changed. Also, there is no PERIOD card required with this function; the full XWTR dataset will be processed by ExPR.

The SYSLOGUPDATE process writes hourly summary records to the database for each LSM and additionally for all non-Nearline activity (manual tape devices and other device types which invoked MVS allocation recovery). Allocation recovery is a single-thread function of MVS. It is therefore important to understand how much throughput delay is attributable to it.

The report shown below is generated during SYSLOGUPDATE processing.

```
               MVS ALLOCATION/RECOVERY UPDATE/REPORT FOR SYSTEM CPUA      DATE RANGE: 1993001 TO 1995365
---
JESJOBNO JOBNAME STEPNAME DDNAME   VOLSER  DATE  STRTTIME END-TIME DURATION UNIT ACS LSM ------ COMMENTS ----

JOB01121  CPF002   DFDSS    TAPE      PRIVAT 1994.341 10.31.44 10.31.56 00:00:12  393  000  01
JOB01127  CPF001   DFDSS    TAPE      PRIVAT 1994.341 10.34.39 10.34.44 00:00:05  394  000  01
JOB01163  N51T001A S1       SYSUT1    PP2V3C 1994.341 10.40.32 10.41.04 00:00:32  393  000  01
JOB01129  CPF001   DFDSS    TAPE      PRIVAT 1994.341 10.47.42 10.48.24 00:00:42  395  000  01
JOB00895  SMFSAVT  INIT     TAPEUNIT  SCRTCH 1994.341 10.50.01 10.50.12 00:00:11  394  000  01
JOB02862  B30REST  S010     INDD1     BRM006 1994.341 11.03.55 11.06.01 00:02:06  394  000  01
JOB02521  B30T005S EXPR     INDATA1   OTS403 1994.341 11.07.49 11.08.14 00:00:25                    OPERATOR REPLIED CANCEL
JOB02863  B30REST  S010     INDD1     BRM007 1994.341 11.25.56 11.29.35 00:03:39  394  000  01
JOB03451  CPF002   DFDSS    TAPE      PRIVAT 1994.341 11.30.23 11.32.01 00:01:38  391  000  00
JOB02591  B30T005S EXPR     INDATA1   OTS403 1994.341 11.32.11 11.32.50 00:00:39  392  000  00
JOB01638  CPF002   DFDSS    TAPE      PRIVAT 1994.341 12.17.27 12.17.35 00:00:08  395  000  01
JOB01373  CPFRS1   DFDSS    TAPE      PRIVAT 1994.341 13.02.47 13.02.53 00:00:06  397  000  02
JOB02197  B30T005C GENER    SYSUT1    CMF001 1994.341 13.49.59 13.51.06 00:01:07  393  000  01
JOB03041  B30COPY  S1       IN        PA0001 1994.341 14.23.18 14.28.53 00:05:35                    OPERATOR REPLIED WAIT
JOB05143  CPF001   DFDSS    TAPE      PRIVAT 1994.341 14.29.10 14.29.17 00:00:07  392  000  00
JOB05145  CPFRS1   DFDSS    TAPE      PRIVAT 1994.341 14.29.27 14.29.32 00:00:05  393  000  01
JOB03041  B30COPY  S1       IN        PA0001 1994.341 14.28.53 14.41.58 00:13:05  396  000  02 MOUNT SERVICED AFTER WAIT
 .
---
```

JES3 Support

The SYSLOGUPDATE example above is for a JES2 system.
JES3 users should specify the following control statement:

```
//EXPR.SYSLOG DD DISP=SHR,DSN=user.syslog
//EXPR.SYSIN DD *
 PERFORM SYSLOGUPDATE JES3(x=yyyy ...) ;
/*
```

The additional parameter JES3 lists the system-ids to be
extracted from the GLOBAL JES3 console file. In the above
example, for each system, x represents the message origin
prefix in the JES3 complex and $yyyy$ represents the

Chapter 3: Managing the ExPR Database   **79**

MVS/JES3 system-id associated with each system (normally the SMF-id). For example, if "1" is the message prefix for "SY1X" and "2" is the message prefix for "SY2X", you would code 1=SY1X 2=SY2X.

When multiple JES3 systems are specified, a separate allocation recovery report is generated for each system. The JES3 console file is read via the DDname SYSLOG, the same as for JES2. It is not possible to mix JES2 and JES3 SYSLOGUPDATE functions in a single execution of ExPR.

In addition to MVS allocation/recovery, ExPR will also report on JES3 mount setup processing as recorded by messages IAT5210 and IAT5918.

RECFM Support

Although the SYSLOGUPDATE function was designed to read the output file from the IBM external writer, it is recognized that many installations use other utilities or packages to manage SYSLOG output. These utilities may not produce a VBA/VBM file like the external writer.

If your SYSLOG file is not VBA/VBM, you can specify the RECFM parameter on the SYSLOGUPDATE statement, as shown below. Valid RECFMs are VBA, VBM, FBA, FBM, FB, and VB. (The external write record format for JES2 is VBA; for JES3 it is VB.)

```
//EXPR.SYSLOG DD DISP=SHR,DSN=user.syslog
//EXPR.SYSIN DD *
  PERFORM SYSLOGUPDATE JES3(x=yyyy ...) RECFM(valid recfm) ;
/*
```

Example 4: Select Tape Catalog Data and Update the Database

```
CDS DDNAME(CDS) ;
PERFORM
TAPECAT OPTION(UPDATE) SYSID(mvshost);
```

PERFORM TAPECAT OPTION(UPDATE) updates the database using tape catalog information if the tape catalog processing feature has been implemented for your site with the @TAPCAT SAMPLIB member. Tape catalog processing and the generated reports are described in chapter 6, *Tape Catalog Processing and Reports*.

UPDATE scans the CDS and tape catalog, calculates age and utilization statistics, updates the database file, creates the tape catalog history file, and generates basic age and usage summary reports. Examples of the control cards to run this task are in SAMPLIB members TAPECAT1 for CA-1 and TAPECATT for CA-TLMS.

Due to the potential overhead caused by reading both the CDS and tape catalog, it is recommended that this task only be run once per day. This level of recording should be adequate for the reports generated from this source. However, the ExPR data key structure would permit the recording of tape catalog information on a more frequent basis if required.

Additionally, this task should be run at the same time each day to avoid data being distorted by any workload variations, such as enter/eject activity.

When running TAPECAT UPDATE to create a TMCHIST file (the default option), be sure to un-comment the TMCHIST DD in EXPRJOB. Do not un-comment this DD for other ExPR functions or when using option NOHIST; you will create empty generations of TMCHIST files.

Example 5: Converting an ExPR 1.0 Database

```
//EXPR.OLDDB      DD   DISP=SHR,DSN=unloaded.EXPR10.database
//EXPR.SYSIN      DD  *
  PERFORM CONVERTDB SYSID(mvshost) ;
/*
```

PERFORM CONVERTDB will read an unloaded sequential copy of the ExPR 1.0 database and write the data directly into the new ExPR 1.1 database. The input file is located by DDNAME OLDDB and should have been previously created using the IDCAMS REPRO function as follows:

```
//DATABASE  DD  DISP=SHR,DSN=EXPR10.database
//UNLOAD     DD  DISP=SHR,DSN=unloaded.EXPR10.database,
               DCB=(LRECL=500,BLKSIZE=nnn,RECFM=VB)
//SYSIN       DD  *
  REPRO       INFILE(DATABASE)  OUTFILE(UNLOAD)
/*
```

ExPR 1.0 data can be loaded into a newly defined (and initialized) 1.1 database or it can be merged into an existing database with newer 1.1 records. Once converted, there is no differentiation of 1.0 and 1.1 data.

When converting the 1.0 database, the SYSID(*mvshost*) parameter must be specified. If you have several 1.0 databases, each representing a different MVS host, the conversion process of unloading and running CONVERTDB can be repeated. Each time, you would change the SYSID parameter, thereby loading all systems into a single ExPR 1.1 database.

The CONVERTDB process produces a small report of records read and written. Due to the database structure changes, the input and output record counts do not always correlate one-to-one:

- Record types 0, 3, 4, 5, 6, 7 ,8, and 9 are processed on a one-to-one basis.

- Record type 2 was previously created one per hour, per LSM; however in ExPR 1.1, this is now one per hour for the entire system.

- Record type 1 is not converted at all.

Record type 1 data can be reproduced by running the SMFSELECT/SMFUPDATE process against your old RMF data. This should be run with a SELECT TYPE(RMF) statement.

Note: An alternative to database conversion would be to maintain the previous release of the mainframe and PC components of ExPR. When the new release has collected sufficient data (for example, two or three months), you could then purge the old product and data.

Multiple Database Support

More than one ExPR database can be maintained. To update a second database, the DATABASE DD statement must point to the appropriate database, and the SELECT statement must identify those systems and ACSs to be input to this database.

Consolidate the Database

The database consolidation procedure allows you to control the size of your mainframe database by moving date ranges of records from the database to an external export file created through JCL or TSO facilities. The export file is a variable blocked QSAM file with a maximum record length of 500 bytes. It can be deleted or archived following the database consolidation.

Two functions are supported, EXPORT and PURGE:

- EXPORT copies database periods to the QSAM file

- PURGE deletes the selected periods from the database

EXPORT and PURGE appear as PERFORM options in the SYSIN runtime parameters. EXPORT and PURGE can be used alone or in combination. When used in combination, PURGE is not executed if the EXPORT command fails.

Reorganizing the Database

At some point, you may wish to reorganize your ExPR database, particularly in larger installations, because of CI/CA splits and secondary extent allocations. This can be achieved with the ExPR EXPORT command and IDCAMS DELETE/DEFINE and REPRO.

The ExPR EXPORT file is a standard VB file which could be reloaded to the main database at any time. When using ExPR EXPORT as part of a database reorganization, be sure to specify a PERIOD FROM/TO that ensures all data is exported before the DELETE/DEFINE, such as PERIOD FROM(1990000) TO(1999365).

Chapter 4

Exception Processing and Reports

Overview

This chapter describes the ExPR exception processing feature and the exception-based reports that can be generated in the mainframe environment.

Process Description

ExPR exception processing monitors user-specified exception thresholds, which represent the point at which an event is considered to be an exception. Exception processing generates a report when thresholds have been exceeded, such as when the number of mounts or mount response time exceeds the specified value.

Thresholds can be set for scratch mounts, non-scratch mounts, scratch mount response time, non-scratch mount response time, enters, ejects, passthroughs, control unit load, number of drives in use, available scratch, available free cells, and maximum allocation recovery time.

Exception reports are produced in tabular hardcopy format on the mainframe. These reports include an event report that lists all SMF exception events and a summary report that provides hourly

threshold exception totals by exception type. Graphical or tabular exception reports can also be generated on the PC. These reports can be run against the mainframe thresholds or against user-specified local thresholds.

Data collected from HSC SMF records and MVS SMF/RMF records is eligible for SMF exception processing. Processing of exceptions takes place automatically during the SYSIN PERFORM SMFUPDATE database update operation or on demand when the SYSIN PERFORM SMFREPORT parameter is specified in the control cards.

Eligible fields are checked against the user-defined thresholds to determine if an exception event has been triggered. If so, an exception event is printed for subsequent review and analysis.

The Threshold Exceptions report reads the ExPR database and prints an hourly summary of the collected data. The current threshold settings are also listed. Database fields exceeding those thresholds are highlighted to draw your attention to them.

Exception Types

All exception values are specified at the LSM level. There are two exception types:

- EVENT: Where an individual event triggers the exception, such as mount time

- INTERVAL: Where the number of events over a recorded period triggers the exception, such as mounts

Default values can apply to event exceptions until some analysis has been done on the reports generated by ExPR. For interval exceptions, some knowledge of activity volume is required. If you are unsure of what these values should be, specify a high value to

avoid excessive exception recording until further analysis of real data has been completed.

Note: As a general rule, avoid frequent adjustment of exception values. The database exception records provide historical perspectives on exceptions and are based on the percentage of total records exceeding the threshold. Therefore, this figure will be distorted if the original threshold is not consistent.

Defining Threshold Values for Your Site

Exception threshold values are defined for your site by customizing sample control cards in the SAMPLIB library. Exception thresholds should be set at a level where you are prepared to take some follow-up action when they are exceeded. This may, for example, be a service level that your installation has implemented.

If there are no service levels in place at your site, then the threshold should be set at a level where you would wish to spend time performing further analysis to identify the reason behind the threshold being exceeded. If thresholds are set too low, there is a danger they will be ignored since this event will become or seem to be the norm. Also, in larger installations, the volume of exceptions being generated could become unmanageable.

Threshold Values

Exception threshold values are set with the CONFIG control statement THRESHOLD. Subparameters for this statement are listed in the table below.

Threshold Name	Description
SCRMOUNTCOUNT	Scratch mount count
SCRMOUNTTIME	Scratch mount time in seconds
NONSCRMOUNTCOUNT	Non-scratch mount count
NONSCRMOUNTTIME	Non-scratch mount time in seconds
ENTERS	Enter count
EJECTS	Eject count
PASSTHRUS	Passthrough count
DRIVESINUSE	Number of drives in use within an LSM
CULOAD	Control unit load percentage threshold
AVAILSCRATCH	Minimum number of available scratch cartridges
AVAILCELLS	Minimum number of available free cells
MAXALLOCREC	Maximum time in seconds for an allocation recovery to be serviced

ACTION:

⇒ Define exception thresholds for your site using SAMPLIB member @THRSHLD.

⇒ Refer to appendix A, *CONFIG Parameters* for information about the THRESHOLDS configuration parameter.

⇒ Be sure to include the CAPS command when tailoring SAMPLIB members.

⇒ Add the customized SAMPLIB member to the CONFIG dataset using SAMPLIB member CONFIG. ExPR does not execute *from* the SAMPLIB library, so only those parameters copied *to* the CONFIG file or automatically generated by ExPR will become effective.

Exception Report Generation

Tabular hardcopy reports are generated against the exceptions created by exception processing. These reports and how to use them are described below.

Two mainframe exception reports are generated; the SMF Exception Event report and the Threshold Exceptions report. These reports are shown in the samples on the following pages.

In addition, exceptions for allocation recovery and percent disconnect for tape devices are shown on the ALLOC-REC and CU-BUSY reports. These reports are described in chapter 7, *Mainframe Performance Reports*.

The Thresholds Exception Summary report is generated by the following SYSIN statement:

- PERFORM REPORT NAME(THRESHOLDS)

The SMF Exception Event report is generated automatically as part of the SMFUPDATE function, or by the SMFREPORT function:

- PERFORM SMFUPDATE reads the SMF data, updates the ExPR mainframe database, and produces the SMF Exception Event report as a part of the normal database update operation.

- PERFORM SMFREPORT reads the SMF data and prints the SMF Exception Event report without updating the database.

ACTION:

$\Rightarrow$ Generate the Threshold Exceptions report using SAMPLIB member EXCPREP.

$\Rightarrow$ To run the SMF Exceptions Event report (SMFREPORT), refer to SAMPLIB member SMFUPDAT. Also refer to *SYSIN Examples* in chapter 3, *Managing the ExPR Database*.

$\Rightarrow$ Refer to appendix B, *SYSIN Parameters* for information about runtime parameters that appear in the SAMPLIB.

$\Rightarrow$ Be sure to include the CAPS command when tailoring SAMPLIB members.

Note: You can also produce exception reports in a graphical or tabular format on a PC, and each ExPR PC user can set local exception thresholds to override the mainframe defaults. Refer to the *ExPR PC User's Guide* for information about PC-based exception reports.

Threshold Exceptions Report

THRESHOLD EXCEPTIONS REPORT FOR SYSTEM CPUA   DATE RANGE: 1993001 TO 1995365
ACS: 000 NEARLINE ACS 000
LSM: 02 NEARLINE LSM 000 02

THRESHOLD VALUES ===> DATE HOUR	MNT CNT SCRATCH (45)	MNT CNT NON-SCR (55)	MNT TIME SCRATCH (50)	MNT TIME NON-SCR (60)	ENTERS (25)	EJECTS (30)	PASSTHRUS (35)	MAX USED DRIVES (6)	AVAILABLE SCRATCHES (200)	AVAILABLE CELLS (200)	MAX ALLOC RECOVERY (300)
1994340 6	4	2	17	20	0	0	0	=> 16 <=	384	256	
7	17	3	26	34	0	0	0	=> 16 <=	384	256	
8	8	3	23	18	0	0	2	=> 16 <=	384	256	
9	15	11	18	25	0	0	1	=> 16 <=	384	256	
10	15	8	21	22	0	0	1	=> 16 <=	384	256	
11	9	10	17	39	0	0	6	=> 16 <=	384	256	
12	0	1	0	17	0	0	0	=> 16 <=	384	256	
17	0	1	0	17	0	0	0	=> 16 <=	528	=> 153 <=	
1994341 6	9	3	28	30	0	=> 131 <=	2	=> 16 <=	528	=> 153 <=	
7	7	4	26	38	0	7	1	=> 16 <=	528	=> 153 <=	
8	11	0	21	0	0	0	0	=> 16 <=	528	=> 153 <=	
9	11	1	27	28	0	0	1	=> 16 <=	528	=> 153 <=	
10	8	3	24	34	0	0	1	=> 16 <=	528	=> 153 <=	
11	6	8	38	38	0	0	4	=> 16 <=	528	=> 153 <=	
12	12	5	28	35	0	0	4	=> 16 <=	528	=> 153 <=	
13	10	6	33	20	0	0	0	=> 16 <=	528	=> 153 <=	
14	5	11	22	18	21	0	0	=> 16 <=	528	=> 153 <=	
15	4	7	33	27	0	0	3	=> 16 <=	528	=> 153 <=	
16	4	5	30	30	0	0	2	=> 16 <=	332	=> 170 <=	
17	3	3	32	58	0	0	1	=> 16 <=	332	=> 170 <=	
18	3	3	=> 51 <=	28	0	0	0	=> 16 <=	332	=> 170 <=	
19	15	1	32	19	0	0	1	=> 16 <=	332	=> 170 <=	
20	15	4	33	30	0	0	1	=> 16 <=	332	=> 170 <=	
21	15	4	35	30	0	0	1	=> 16 <=	332	=> 170 <=	
22	14	1	34	17	0	0	4	=> 16 <=	332	=> 170 <=	
23	16	0	27	0	0	0	1	=> 16 <=	332	=> 170 <=	
1994342 0	19	1	23	58	0	0	0	=> 16 <=	332	=> 170 <=	
1	8	3	38	35	0	0	3	=> 16 <=	332	=> 170 <=	

.
.
.

The Threshold Exceptions report is run against the ExPR
mainframe database. It compares the thresholds you have
specified against the database records over a selected period of
time. The report highlights those fields that exceed the threshold,
providing you with information to use as a starting point when
performing detailed analysis on any problem areas.

Chapter 4: Exception Processing and Reports   **93**

Reports are generated per-LSM and show only hours where at least one field has exceeded its threshold.

If a threshold is exceeded, all the other fields for that hour are displayed even if they are within threshold. This enables a basic analysis of related data to be completed.

If data is not present, the field is filled with blanks.

Where exceptions are present, they are highlighted with => and <= characters on either side of the field.

Beneath each exception field heading, the exception threshold specified in the CONFIG file is displayed.

SMF Exception Event Report

SMF UPDATE - SMF EXCEPTION EVENTS

DATE	TIME	SYS	DRIVE	TYPE	ACS	LSM	EVENT	OBSERVATION	THRESHOLD	VARIATION
1994.341	13:01:40	CPUA	0C25	4490	000	02	UNEXPECTED DISMOUNT			
1994.341	13:01:46	CPUA	0C09	9490	000	00	UNEXPECTED DISMOUNT			
1994.341	13:06:59	CPUA	0C00	4490	000	02	UNEXPECTED MOUNT			
1994.341	13:07:58	CPUA	0C0A	9490	000	00	UNEXPECTED DISMOUNT			
1994.341	13:09:03	CPUA	0C0B	9490	000	00	NON-SCRATCH MOUNT	76	60	126% ------>
1994.341	13:28:49	CPUA	0C02	9490	000	01	NON-SCRATCH MOUNT	67	60	111% ----->
1994.341	13:30:34	CPUA	0C09	9490	000	00	NON-SCRATCH MOUNT	70	60	116% ----->
1994.341	13:33:43	CPUA	02B0		000	02	DEMAND ENTER - VOLSER=030471 JOB=MAINT13C			
1994.341	13:33:43	CPUA	02B0		000	02	NON-SCRATCH MOUNT	103	60	171% -------->
1994.341	13:36:12	CPUA	02B0		000	02	DEMAND ENTER - VOLSER=030470 JOB=MAINT13C			
1994.341	13:38:15	CPUA	0C27	9490	000	02	SCRATCH MOUNT	63	50	126% ------>
1994.341	13:45:03	CPUA	0C00	4490	000	01	HIGH CU DISCONNECT	92%	80%	+12%
1994.341	13:46:29	CPUA	0C0B	9490	000	01	DRIVESINUSE EXCEEDED	7	6	116% ----->
1994.341	13:59:59	CPUA			000	01	EXCESSIVE SCR MOUNTS	40	20	200% --------->
1994.341	14:02:31	CPUA	0C2A	4480	000	01	NON-SCRATCH MOUNT	80	60	133% ------>
1994.341	14:30:03	CPUA	0C2A	4480	000	01	HIGH CU DISCONNECT	81%	80%	+ 1%
1994.341	14:45:04	CPUA	0C25	4490	000	02	HIGH CU DISCONNECT	96%	80%	+16%
1994.341	14:47:47	CPUA	0C00	4490	000	01	DEMAND ENTER - VOLSER=104045 JOB=PRDACNTS			
1994.341	14:47:53	CPUA	0C01	4490	000	01	DEMAND ENTER - VOLSER=102890 JOB=PRDACNTS			
1994.341	14:48:01	CPUA	0C02	9490	000	01	DEMAND ENTER - VOLSER=102914 JOB=PRDACNTS			
1994.341	14:48:13	CPUA	0C03	9490	000	01	DEMAND ENTER - VOLSER=102916 JOB=PRDACNTS			
1994.341	14:48:23	CPUA	0C00	4490	000	01	DEMAND ENTER - VOLSER=102635 JOB=PRDACNTS			
1994.341	14:48:35	CPUA	0C01	4490	000	01	DEMAND ENTER - VOLSER=107178 JOB=PRDACNTS			
1994.341	14:48:36	CPUA	0C02	9490	000	01	DEMAND ENTER - VOLSER=102645 JOB=PRDACNTS			
1994.341	14:48:41	CPUA	0C03	9490	000	01	DEMAND ENTER - VOLSER=000006 JOB=PRDACNTS			
1994.341	14:50:24	CPUA	0C27	9490	000	02	PERM I/O ERROR VOLSER=013539			
1994.341	14:51:34	CPUA	0C00	4490	000	01	PERM I/O ERROR VOLSER=013539			
1994.341	14:52:45	CPUA	0C02	9490	000	01	NON-SCRATCH MOUNT	69	60	115% ----->
1994.341	14:52:54	CPUA	0C0B	9490	000	00	DRIVESINUSE EXCEEDED	7	6	116% ----->
1994.341	14:55:23	CPUA	0C02	9490	000	01	PERM I/O ERROR VOLSER=017003			
1994.341	14:59:17	CPUA	0C22	4490	000	00	SCRATCH MOUNT	54	50	108% ----->
1994.341	14:59:50	CPUA	0C2A	4480	000	01	SCRATCH MOUNT	67	50	134% ------>

The SMF Exception Event report is generated automatically by PERFORM SMFUPDATE during database update processing, and on demand by PERFORM SMFREPORT. The report compares the threshold defined for mount response time for scratch and non-scratch volumes with the mount events as they were recorded. No summarization of this data has taken place, so reporting is at the lowest possible level.

Possible exceptions from SMF data are listed below:

Chapter 4: Exception Processing and Reports   **95**

Exception	Description
SCRATCH MOUNT	Scratch mount time exceeded the threshold.
NON-SCRATCH MOUNT	Non-scratch mount time exceeded the threshold.
EXCESSIVE SCR MOUNTS	The scratch mount count threshold has been exceeded for this LSM/hour.
EXCESSIVE NSC MOUNTS	The non-scratch mount count threshold has been exceeded for this LSM/hour.
UNEXPECTED MOUNT	A mount occurred for an already mounted drive.
UNEXPECTED DISMOUNT	A dismount occurred for an already free drive.
DRIVESINUSE EXCEEDED	The number of active drives was exceeded for this LSM.
HIGH CU DISCONNECT	The tape control disconnect time exceeded the threshold.
PERM I/O ERROR VOL=*xxxxxx*	The identified volume had a permanent read or write error.
TEMP I/O ERROR VOL=*xxxxxx*	The identified volume had a temporary read or write error.
DEMAND ENTER - VOL=*xxxxxx* JOB=*xxxxxx*	A mount was issued against a Nearline drive for a volume that was not in the library. The operator had to load the cartridge into a CAP to satisfy the mount.
HSC COUNTER OVERFLOW	An HSC LMU statistics counter has overflowed and been reset to zero by ExPR. This prevents distortion of mount-time breakdown values.

An unexpected dismount can occur at the start of the SMF data if the corresponding mount preceded the start of the SMF file.

A returned RC=4 from the report indicates a possible break in SMF recording, resulting in a mount-mount sequence with no intervening dismount, or a dismount with no preceding mount. If this occurs, check that SMF input data is complete and continuous.

A returned RC=8 from any ExPR report request indicates that the database did not contain records to match your selection criteria (date range, system, or ACS/LSM ids).

Note: The SMFREPORT/SMFUPDATE functions will return RC=8 if the input file is not in ascending date/time sequence. Ensure that the SMFSELECT and SORT steps are run correctly.

Chapter 5

Jobname Workload Groups

Overview

This chapter describes the optional jobname workload groups feature. Jobname workload groups map Nearline activity from specified jobnames into logical groups, typically defining activities generated by specific departments or functions. Reports can be generated for these groups, increasing the ExPR reporting capability by further defining Nearline activities within your organization.

Defining Jobname Workload Groups

A workload group can comprise one or more jobnames. Groups are specified by including one or more jobnames under a user-selected group name. Masks and wildcard characters are used to select a range of jobnames for a group.

The following masks can be used when constructing ExPR jobname workload groups:

? or %	Represents a wildcard character, one for one. Blanks are not allowed.
*	Represents a wild prefix or suffix (0-8 characters). Can be used as either the prefix or suffix, or both.

For example:

JT100*	Any job starting JT100.
J*100ABC	Jobs such as JT100ABC or JD100ABC
J????ABC	Jobs starting with J and ending with ABC, with any alphanumeric characters in positions 2-5.

Workload groups are defined with the CONFIG control statement MAP. Identifying names can be defined for jobname workload groups with the CONFIG control statement WORKLOAD.

Note: Jobname workload groups will significantly increase the size of the ExPR database, so some thought is required at specification time. Refer to *Task 2: Space Planning* in chapter 2, *Installation and Customization Tasks*.

```
ACTION:

⇒  Define jobname workload groups for your site using
   SAMPLIB member @MAP.

⇒  Refer to appendix A, CONFIG Parameters for
   information about the MAP and WORKLOAD
   configuration parameters.

⇒  Be sure to include the CAPS command when tailoring
   SAMPLIB members.

⇒  Add the customized SAMPLIB member to the CONFIG
   dataset using SAMPLIB member CONFIG.  ExPR does
   not execute from the SAMPLIB library, so only those
   parameters copied to the CONFIG file or automatically
   generated by ExPR will become effective.
```

Jobname Workload Group Reports

Jobname workload group reporting is available primarily
through the ExPR PC component. Reports for mount counts
and mount response time can be generated by jobname
workload group in that environment. Refer to the *ExPR PC
User's Guide* for information about these reports. The
mainframe REPORT NAME(MOUNTS) report can
optionally print jobname workload group information when
REPORT-OPTIONS(WORKLOADS) is specified.

Chapter 6

Tape Catalog Processing and Reports

Overview

This chapter describes the optional ExPR tape catalog processing function. If you use this function, ExPR will process your tape catalog as an input source to the database update process that takes place on the mainframe. Database records are written that can be later reported against as mainframe tabular reports (or as PC graphical reports when the file is transferred to the ExPR PC component).

Tape catalog processing makes another level of Nearline reporting available to you. Other input sources look at how the various hardware components are performing and identify processing trends and potential bottlenecks. With tape catalog database records, you can also account for who is occupying the slots within each library and also how efficiently the tape is being used. Additionally, you can produce historical comparisons of your tape catalog's contents.

Note: Important tailoring steps are required to use this feature of ExPR. Tailoring steps you must complete are described in this chapter; syntax for these steps is described in appendix A, *CONFIG Parameters*.

Supported Environments

ExPR currently supports the CA-1 and CA-TLMS tape management systems. Tape catalog users with products other than CA-1 or CA-TLMS can also write an interface to the ExPR Custom Tape Management Interface (CTMI) architecture. Contact your support representative for technical specifications.

It is still possible to use the tape processing facilities of ExPR even if your installation does not have a tape management catalog, or if ExPR does not support your tape management system. Refer to "Running Without a Tape Catalog" at the end of this chapter.

Note: CA-TLMS users should refer to the end of this chapter for additional considerations about setting up ExPR tape catalog processing for use with CA-TLMS.

How ExPR Tape Processing Works

ExPR requires two input sources to build its tape catalog-related database fields:

- Installation Tape Catalog

- HSC Control Dataset

ExPR reads the control dataset (CDS) and builds a list of cartridge volumes in the libraries. This list is referenced against the tape catalog and the appropriate fields are extracted. Optionally, information about all cartridge volumes in the tape catalog can be extracted.

ExPR tape catalog processing analyzes the collected data and generates two summary reports automatically. These reports analyze tape aging and utilization statistics at the ACS/LSM and dataset

workload group level. Depending on the submitted SYSIN control statements, ExPR also writes database records types 4, 5, 6, and 7 to its mainframe database (see appendix D, *Database Layout* for information about individual record types) and writes one record per volume to a TMCHIST history file. This file is a GDG dataset created each time the tape processing update function is run, allowing comparative analysis on a day-to-day basis with the tape processing history reports.

Unlike the other ExPR data sources, the tape catalog needs to be read only once a day under normal circumstances. This is to reduce the overhead of processing such a large file, and also because the profile of information contained within the tape catalog should not typically change significantly during each 24-hour cycle. If this is not the case at your installation, you can run this process more frequently and get ExPR to write additional database records.

Figure 6-1 illustrates how ExPR processes tape catalog-based input.

_____

_____
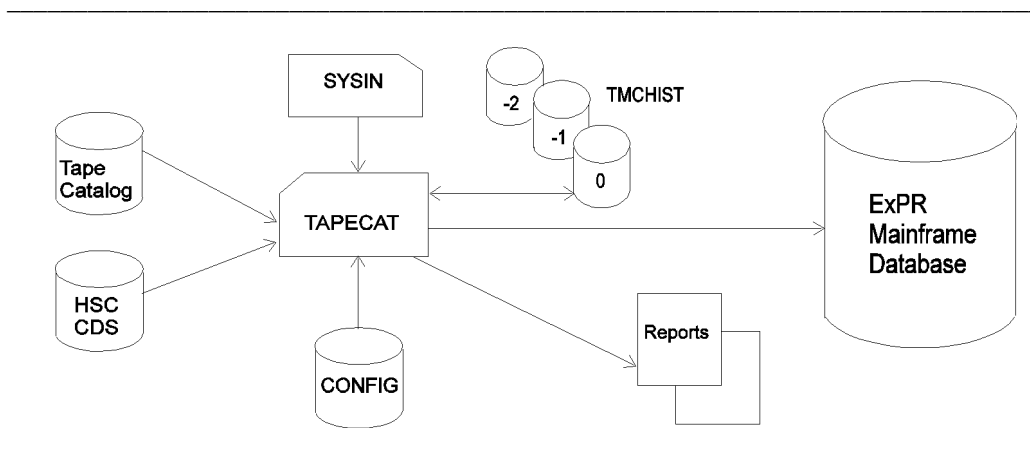
Figure 6-1: Tape Processing Data Flow

What Information Is Generated

Tape processing gives access to information that you can use to determine which users or applications are using the Nearline libraries and how efficiently they are using them. ExPR has a number of unique analysis features to help isolate tape catalog usage and trends, including:

- Dataset workload group analysis

- Tape utilization estimation analysis

- User-specified age band analysis

- Comparative historical analysis

- Media type analysis

These analysis features are reflected in tape catalog reports that are generated at the mainframe in a tabular hardcopy format. The following types of mainframe tape catalog reports are available:

- The tape update Aging summary report lists, for each LSM and for each user-defined dataset workload group, the number of cartridges within user-specified age bands and their average ages, the number of cartridges containing multi-volume datasets, and the number of cartridges containing multiple datasets.

- The tape update Utilization report lists, for each LSM and for each user-defined dataset workload group, the number of cartridges within each tape utilization percentage band.

- The Volume Details report lists, for each volume, the volume sequence, number of megabytes, ACS/LSM location, number of datasets, estimated utilization percentage, last-reference date, number of accesses, device type and cartridge length, scratch

status, and the name, file sequence, block count, block size, record size, and record format of each dataset.

- The Volume Summary report provides tape catalog volume contents information summarized for each ACS, each LSM, each dataset workload group, each device type, each media type, and each defined tape length.

- The Tape Catalog History report provides comparative activity analysis between two tape catalog images by listing volumes that have changed status.

Defining Tape Catalog Processing for Your Site

This task provides information about specifying the CONFIG parameters ExPR will require to understand what tape management system is in place at your installation and what runtime options you wish to apply.

Tape catalog processing is established with the @TAPCAT member of the ExPR SAMPLIB. This member contains three statements:

- TMS tells ExPR which of its tape management modules to call and which DD statement identifies the tape catalog. The supplied EXPRPROC procedure requires a DDname of TMC (the default).

- AGEBANDS specify a series of two to four numeric values, each of which represents a number of days. These are in a range, low to high, and are used to provide a distribution report of how old the cartridges are in the libraries.

- TAPEDEF defines ranges of cartridges based on length, Redwood capacity, or desired exclusion from tape catalog processing.

Defining Age Bands

Age bands are used to classify last-referenced date information into concise groupings that show the distribution of data across a dataset workload group, LSM, or entire ACS. They are specified in the ExPR CONFIG file with the AGEBANDS control statement.

You can specify up to four age bands on the AGEBANDS command. For example, AGEBANDS(365 183 91 31) would summarize the ACS-resident cartridges into age groupings of up to one month, three

months, six months, and one year since last being referenced/accessed by the host system. A fifth "catchall" band will account for cartridges older than the oldest ageband.

Tape Utilization and Capacity Estimation

The ExPR tape utilization feature provides an estimation of the length of tape that has been used for groups of cartridges, defined by dataset workload groups, or for an LSM. It should be stressed that this is an estimation tool only and many variable factors exist that can affect the accuracy of the information recorded.

ExPR makes utilization estimates based on cartridge tape lengths and capacities (either default lengths or lengths you specify in the CONFIG dataset with the TAPEDEF control statement).

If you require very precise information, it is recommended that you benchmark this feature against known cartridge capacities and data profiles before making any external processing decisions. ExPR takes into account factors such as HSM, variable and unblocked files, IDRC, hardware type, and tape lengths. However, in some cases, these rules are applied at a global level. For example, the IDRC compression ratio for 3480- and 3490-type devices is fixed at 2:1 against non-IDRC devices.

The percentage utilization calculation is not a measurement of blocking efficiency.

ExPR will also calculate the amount of data written to a tape in megabytes. Once again, the device type, recording technique, and length of tape is taken into account. An average blocksize is calculated from catalog information about all datasets on a given volume. This average blocksize and the number of blocks is used to determine if a volume's theoretical capacity has been exceeded. If

the megabytes value exceeds 80 percent of the theoretical capacity, the volume's capacity will be capped. This reduces distortion caused by incorrectly stating cartridge lengths or, for multi-volume files, excessively large block counts stored in the tape catalog. Message XPR0097W will advise you how many volumes have been capped.

Defining Tape Ranges

Ranges of tape cartridges can be defined by volser number using the CONFIG TAPEDEF control statement. Ranges are typically defined based on length, Redwood capacity, or desired exclusion from tape catalog processing. Each of these uses is described in this section.

Defining Tape Ranges Based on Length

Since cartridges are usually purchased in blocks of volser numbers, it should be possible to match ranges of volsers to a given tape length. By supplying this information to ExPR, you can increase the accuracy of tape utilization estimates in ExPR tape catalog reports.

Generally, cartridges come in four lengths - 550, 685, 800, and 1100 feet. Manufacturers have each produced variations in cartridge lengths, particularly in the 550 foot "standard length" range (i.e., 541/548/565/575). It is strongly recommended that all "standard length" cartridges be grouped together under a common length classification of 550 feet.

The sample CONFIG TAPEDEF control statements coded below would define two ranges of volsers with a unique tape length for each range:

```
TAPEDEF LOVOL(000000) HIVOL(010000) LENGTH(550) ;
TAPEDEF LOVOL(010001) HIVOL(015000) LENGTH(1100) ;
```

In this example, volsers 000000-010000 are defined as 550 feet in length and volsers 010001-015000 are defined as 1100 feet. During tape catalog processing, the volser will be used to determine the length of each cartridge. This value will be used with the tape catalog device type density value to estimate the capacity and percentage utilization of the cartridge.

Alternatively, if you are not particularly interested in the volume utilization statistics, you can define a standard length for each type of cartridge. The sample CONFIG TAPEDEF control statement coded below would define standard lengths for all 3480 and 3490 cartridges:

```
TAPEDEF CART3480(550)  CART3490(1100) ;
```

In this example, all 3480 cartridges are defined as 550 feet and all 3490 cartridges are defined as 1100 feet. During tape catalog processing, the cartridge type will be used to determine the length of each cartridge. This value will be used with the tape catalog device type density value to estimate the capacity and percentage utilization of the cartridge. The estimation provided by this general assessment of cartridge lengths will be somewhat less accurate than in the previous example where cartridge lengths were more precisely defined.

You can define as many TAPEDEF ranges as you wish, there is no limit. The default 3480/3490 length is assumed for volumes that are not covered by a specific range.

Defining Tape Ranges Based on Redwood Capacity

HSC 2.0.1 introduces support for the helical scan Redwood devices and three capacities of cartridges, 10Gb, 25Gb and 50Gb. ExPR supports Redwood through the VARMEDIA flag in the DVAR CDS records. The following VARMEDIA values are supported:

VARMEDIA Value	Description
X'00' or C'1'	Tape is 3480/3490 type. Use TAPEDEF definitions to determine capacity and length.
C'A'	Redwood 10Gb cartridge
C'B'	Redwood 25Gb cartridge
C'C'	Redwood 50Gb cartridge
C'E'	3490E extended length cartridge (1100 ft.). This takes precedence over any TAPEDEF range that may include this cartridge.

If the CDS DVAR record returns a VARMEDIA value of A,B,C or E, then that device/media type takes precedence over TAPEDEF definitions. However, if you inadvertently define a range of 3480/3490 type cartridges as Redwood, then the TAPEDEF definition will take precedence.

Under normal circumstances you should not need to define ranges of Redwood cartridges, though you may wish to for purposes of clarity and documentation. The sample CONFIG TAPEDEF control statements coded below would define ranges of volsers for each type of Redwood cartridge:

```
TAPEDEF LOVOL(100000) HIVOL(100499) OPTION(RW10) ;
TAPEDEF LOVOL(100500) HIVOL(100999) OPTION(RW25) ;
TAPEDEF LOVOL(101000) HIVOL(101499) OPTION(RW50) ;
```

One example of how you might use this feature is to define a range of 3490 cartridges as Redwood cartridges, sort them by dataset, and analyze the results to help determine the impact Redwood device types might have on your library capacities and utilization. If you set up a scenario like this, be sure to include the NODBUP and NOHIST options in the PERFORM TAPECAT OPTION(UPDATE) control card so as to not store these "what if" results in your database and history files.

Defining Tape Ranges to EXCLUDE from Processing

ExPR also allows the exclusion of a range or ranges of tape cartridges from tape catalog processing. This option is intended for sites with multiple MVS systems and sites that want to analyze a subset of the complete ACS tape catalog.

An example of excluding a range of tapes in a multiple shared MVS/ACS could be a situation where an ACS holds volumes 000000-010000 and 100000-106000, with the first range being production cartridges in a production MVS system and tape catalog, and the second range belonging to a test/development MVS on another LPAR/CPU with its own tape catalog. TAPECAT UPDATE processing on the production system would issue 6000 occurrences of the XPR0080W NOT IN TAPE CATALOG message. This could be prevented by the following TAPEDEF statement:

```
TAPEDEF LOVOL(100000) HIVOL(106000) OPTION(EXCLUDE) ;
```

This example would exclude volsers 100000-106000 from tape catalog processing. The exclusion filter is applied to records retrieved from the tape catalog and also to those retrieved from the HSC CDS.

Another example of how the EXCLUDE option could be used is to generate snapshot statistics of a subset of the tape catalog. This subset could be defined by the following TAPEDEF statement:

```
TAPEDEF LOVOL(050000) HIVOL(999999) OPTION(EXCLUDE) ;
```

This example would exclude volsers 050000-999999 from tape catalog processing, effectively creating a subset of the tape catalog composed of cartridges that were not excluded. Such ad-hoc statistics could be generated with the following TAPECAT statement:

```
PERFORM TAPECAT OPTION(UPDATE NOHIST NODBUP) ;
```

This example would process all but the excluded data and would also preclude ExPR from updating the history and database files with the results of this processing.

Refer to appendix B, *SYSIN Parameters* for more information about the TAPECAT control statement and the options that are used in this example.

Defining Dataset Workload Groups

Dataset workload groups, used in conjunction with ExPR tape catalog processing, provide a mechanism for logically mapping the contents of a library to particular systems or applications based on dataset names. A workload group typically defines a critical application or system component, such as HSM. While optional, these workload groups can greatly enhance the level of detail of your Nearline reporting. A dataset workload group records statistics for cartridges

whose primary dataset name matches a set of masks defined by you in the ExPR CONFIG file.

If dataset workload groups are to be used, you can define one or many groups to allow ExPR to report at an appropriate level of detail. However, unless clearly defined naming standards or existing workload specifications are in place at your site, you should not define overly complex groups.

Note: Cartridges may appear in more than one dataset workload group. A duplicate entry report is automatically generated that lists cartridge volumes that are in more than one group.

Dataset workload groups are specified by including one or more datasets under a user-selected group name. The mask that defines the group includes dataset qualifiers and wildcard characters. The rules and syntax for specifying workload groups are detailed below.

The following masks can be used when constructing ExPR dataset workload groups:

? or %	Represents a wildcard character, one for one. Blanks are not allowed.
*	Represents a single node wild prefix or suffix (1-8 characters). Can be used as either the prefix or suffix, or both.
**	Represents a wild single or multiple node up to 21 characters long. Can be used as a prefix or suffix (but not both) or on its own.

For example:

*.PROD?	SYS1.PROD1, SYS1.PROD2, and IPO1.PROD3 but not ACCOUNTS.DATA.PROD4 or IP01.SYSTEM.PROD3
SYS1.**.TEST	SYS1.LINKLIB.TEST and SYS1.LINKLIB.MVS1.TEST but not SYS2.LINKLIB.TEST or SYS1.LINKLIB.MVS1.PROD3

SYS1.BACKUP.*	SYS1.BACKUP.DAILY and SYS1.BACKUP.WEEKLY but not SYS1.BACKUP.DAILY.MONDAY or SYS2.BACKUP.WEEKLY
SYS2.BACKUP.**	SYS2.BACKUP.DAILY and SYS2.BACKUP.DAILY.MONDAY but not SYS1.BACKUP.DAILY or SYS2.DAILY.MONDAY

Dataset workload groups are defined and mapped by the CONFIG control statements DSGROUP and DSMAP. DSGROUP defines the group and DSMAP maps datasets to the group. The CLASS subparameter in each statement links them together.

Using the previous examples, four unique dataset workload groups could be established with the following CONFIG control statements:

```
DSGROUP CLASS(1) NAME('ALLPRODX') ;
DSMAP CLASS(1) DSN(*.PROD?) ACS(0) ;
DSGROUP CLASS(2) NAME('S1ALLTEST') ;
DSMAP CLASS(2) DSN(SYS1.**.TEST) ACS(0) ;
DSGROUP CLASS(3) NAME('S1BACKUP') ;
DSMAP CLASS(3) DSN(SYS1.BACKUP.*) ACS(0) ;
DSGROUP CLASS(4) NAME('S2BAKALL') ;
DSMAP CLASS(4) DSN(SYS2.BACKUP.**) ACS(0) ;
```

If you have not previously classified applications that use the Nearline system in a way that can be easily mapped to dataset names, it is recommended that you begin by specifying some simple and fairly broad groups (for example, HSM or other archive/backup workloads). Care should be taken when specifying complex masks to ensure that a dataset only qualifies for its intended group or groups.

> **ACTION:**
>
> ⇒ Define and map dataset workload groups for your site using SAMPLIB member @DSMAP.
>
> ⇒ Refer to appendix A, *CONFIG Parameters* for information about the DSMAP and DSGROUP configuration parameters.
>
> ⇒ Be sure to include the CAPS command when tailoring SAMPLIB members.
>
> ⇒ Add the customized SAMPLIB member to the CONFIG dataset using SAMPLIB member CONFIG. ExPR does not execute *from* the SAMPLIB library, so only those parameters copied *to* the CONFIG file or automatically generated by ExPR will become effective.

Special Considerations for the TMCHIST File

During installation and customization of ExPR in chapter 2, SAMPLIB member DSNDEF created a generation data group for the TAPECAT TMCHIST file. This file holds a record per volume in the CDS (and optionally the tape catalog) and is a snapshot of the tape library at the time the TAPECAT UPDATE function is run.

The optional TMCHIST DD statement in EXPRJOB should only be set to create generation (+1) when running the TAPECAT UPDATE function; otherwise you will create empty generations of TMCHIST whenever you run EXPRJOB.

If the TAPECAT reporting options REPORT or SUMMARY are requested in the same run as UPDATE, they will read the newly created (+1) generation. However, when running TAPECAT REPORT or SUMMARY standalone, the TMCHIST DD should specify (0) or a previous generation (-1, -2, etc.).

The TAPECAT HISTORY function requires two generations of TMCHIST for comparison purposes. These are read via DD statements TMCHIST1 and TMCHIST2 and should always specify previous generations (0, -1, -2, etc.).

Note: If your installation wishes to read the TMCHIST file for your own processing requirements, SAMPLIB member EXPRHIST contains an assembler record layout. Appendix D, *Database and TMCHIST Layouts*, also describes the TMCHIST record.

Processing Tape Catalog Data

Tape catalog processing is controlled and initiated by mainframe batch tasks that scan the CDS and tape catalog, calculate age and utilization statistics, update the database file, create the history file, generate basic age and usage summary reports, and produce optional reports. Sample SYSIN decks are provided below.

Note: If you use CA-TLMS as your tape management system, refer to the section *CA-TLMS Considerations* at the end of this chapter.

ACTION:

⇒ Refer to appendix B, SYSIN Parameters for information about the PERFORM TAPECAT runtime parameters.

⇒ Review SAMPLIB members TAPECAT1 and TAPECATT for examples of PERFORM TAPECAT.

Example 1: Process Tape Catalog and Update Database and History Files

This example will read the CDS and tape catalog, calculate tape utilization estimations, and calculate age statistics based on last-reference date and data profile, such as stacked and multi-volume statistics. The database file is updated and the history file is created in this example, and two basic age and usage summary reports are generated listing each ACS/LSM or dataset group that was processed (ExPR messages XPR0091I and XPR0092I).

```
//SYSINDD *
  PERFORM TAPECAT OPTION(UPDATE) SYSID(mvshost) ;
```

Do not forget to un-comment the TMCHIST DD in EXPRJOB for TAPECAT UPDATE.

Note: The SYSID parameter is required; an error message will be issued if it is omitted. In a multi-MVS shared tape catalog/CDS environment, you should specify the system-id of the main system. The TAPECAT database records are written to the database with the SYSID value in the record key.

Example 2: Process Tape Catalog Without Updating Database and History Files

This example will perform all of the calculations and generate the two summary reports as in the previous example, but will not update the database file and will not create the history file. This example is useful for generating snapshot statistics for subsets of the tape catalog when the CONFIG statement TAPEDEF EXCLUDE is specified, as

was previously described in this chapter in the section *Defining Tape Ranges to EXCLUDE from Processing*.

```
//SYSIN  DD *
  PERFORM TAPECAT OPTION(UPDATE NOHIST NODBUP) SYSID(mvshost) ;
```

When using option NOHIST, you should *not* un-comment the TMCHIST DD in EXPRJOB. Doing so will create an empty generation of the TMCHIST file.

Example 3: Process Tape Catalog Including All Volumes

By default, ExPR processes just those volumes in the tape catalog that are listed in the HSC CDS. However, the SYSIN control statement TAPECAT OPTION includes a subparameter option, FULLCAT, which includes all volumes in the tape management catalog for ExPR tape catalog processing. Use of this subparameter does not affect the calculation of age and utilization data for the ACSs, LSMs, and dataset workload groups. However, it will cause additional records to be written to the history file for subsequent comparison, sorting, and reporting (unless NOHIST is also specified).

```
//SYSIN  DD *
  PERFORM TAPECAT OPTION(UPDATE FULLCAT) SYSID(mvshost) ;
```

Do not forget to un-comment the TMCHIST DD in EXPRJOB for TAPECAT UPDATE.

Generating Tape Catalog Reports

Tape catalog reports are produced by the following batch processes:

- TAPECAT OPTION(UPDATE) produces the tape update processing Aging and Utilization reports as part of the ExPR tape catalog processing function.

- TAPECAT OPTION(REPORT BY*xxxx*) produces the Nearline Volume Detail report.

- TAPECAT OPTION(SUMMARY) produces the Nearline Volume Summary report.

- TAPECAT OPTION(HISTORY) produces the History comparison report.

Each of these reports is described on the following pages.

TAPECAT Update Processing Aging and Utilization Summary

```
----------------------------------------------------------------------------------------------------------------------------
1994.340      S T O R A G E T E K   E X P R  1 . 1     XYZ COMPANY                                              PAGE  1
17:32:20
                                      TAPE CATALOG PROCESSING FOR SYSTEM CPUA
----------------------------------------------------------------------------------------------------------------------------
XPR0709I HSC / SLSXCAL LEVEL IS 2.0.0
XPR0080W Volser ABS001 not in Tape Catalog
XPR0080W Volser FAT001 not in Tape Catalog
XPR0087W VOLUME 000020 DSN=IDMS.LIBRARY.WEEKLY.BACKUP.G0321V00        SELECTED IN DATASET-GROUP   1 AND  1026
XPR0087W VOLUME 000105 DSN=TEST.OLAS.IBCATB.BACKUP.G0079V00           SELECTED IN DATASET-GROUP   1 AND   2
XPR0082W Volser 004954: Catalog says Scratch, CDS says not
XPR0080W Volser 999995 not in Tape Catalog


----------------------------------------------------------------------------------------------------------------------------
1994.340      S T O R A G E T E K   E X P R  1 . 1     XYZ COMPANY                                              PAGE  1
17:32:20
                                      TAPE CATALOG PROCESSING FOR SYSTEM CPUA
----------------------------------------------------------------------------------------------------------------------------
XPR0092I      OLDEST # MULTI  # MULTI (9999- 457 DAYS)  ( 456- 366 DAYS)   ( 365- 275 DAYS)   ( 274- 184 DAYS)   ( 183-  1 DAYS)
ACS LSM      IN DAYS VOLUMES  FILES (# VOLS/AVG AGE)(# VOLS/AVG AGE)(# VOLS/AVG AGE)(# VOLS/AVG AGE)(# VOLS/AVG AGE)
 0    0       322    1393    1147    0      0       0       0       26     304      51     216     5409     5
 0    1       314    1123     998    0      0       0       0       19     299      75     212     5265     5
 0    2       315    1317    1049    0      0       0       0       42     305     103     216     5264     5


XPR0092I DATA- OLDEST # MULTI  # MULTI (9999- 457 DAYS)  ( 456- 366 DAYS)   ( 365- 275 DAYS)   ( 274- 184 DAYS)   ( 183-  1 DAYS)
ACS LSM GROUP IN DAYS VOLUMES  FILES (# VOLS/AVG AGE)(# VOLS/AVG AGE)(# VOLS/AVG AGE)(# VOLS/AVG AGE)(# VOLS/AVG AGE)
ALL  1      290    3705    3037    0      0       0       0       21     285     136     227     7429     4
 0   1   2   32      38      39    0      0       0       0        0       0       0       0      106     3
ALL ALL 1026 290     618     132    0      0       0       0       12     290      12     260      779     0
 0   2  1027  0       8       6    0      0       0       0        0       0       0       0       10     0
 0  ALL 1028  0       0       0    0      0       0       0        0       0       0       0       10     0
 0  ALL 1029  0       0       0    0      0       0       0        0       0       0       0       10     0

XPR0091I      TOTAL UTILIZED AVERAGE - - - - - - - PERCENTAGE UTILIZED BANDS (NO OF VOLUMES) - - - - - - - - -
ACS LSM      VOLUMES VOLUMES % UTIL  00-10%  11-20%  21-30%  31-40%  41-50%  51-60%  61-70%  71-80%  81-90%  91-100%
 0    0       5728    4141    41    1189    317     679     577     130     62      40      13      15     1119
 0    1       5665    3666    38    1142    321     591     501     108     55      28      19       5      896
 0    2       5741    3853    40    1241    328     466     532     123     59      19      12      13     1060


XPR0091I DATA- TOTAL UTILIZED AVERAGE - - - - - - - PERCENTAGE UTILIZED BANDS (NO OF VOLUMES) - - - - - - - - -
ACS LSM GROUP VOLUMES VOLUMES % UTIL  00-10%  11-20%  21-30%  31-40%  41-50%  51-60%  61-70%  71-80%  81-90%  91-100%
 0  ALL  1     7586    7586    44    3099    747     449     143      88      64      11       6      17     2962
 0   1   2      106     106    34      39     29       2       3       2       4       0       0       2       25
ALL ALL 1026    803     803    75     128     15      26      34      16      11       1       3      12      557
 0   2  1027     10      10    45       3      3       0       0       0       0       0       0       0        4
 0  ALL 1028     10      10    10       7      0       3       0       0       0       0       0       0        0
 0  ALL 1029     10      10    10       7      0       3       0       0       0       0       0       0        0

XPR0093W AT LEAST ONE LSM OR DATASET-GROUP HAS NO VOLUMES
XPR0094I   4910 CDS VOLUMES NOT IN TAPE CATALOG
XPR0095I      3 VOLUMES HAVE A CONFLICTING SCRATCH STATUS
XPR0097W    254 VOLUMES HAVE HAD THEIR MEGA-BYTES CAPACITY CAPPED
----------------------------------------------------------------------------------------------------------------------------
```

The tape update processing Aging and Utilization summary reports are generated automatically by the ExPR tape catalog processing function TAPECAT OPTION(UPDATE). Processing anomalies between the tape catalog and CDS are also printed as XPR messages with these reports.

- The Aging summary report (XPR0092I) lists, for each LSM and again for each user-defined dataset workload group, the number of non-scratch cartridges within user-specified age bands and their average ages, the number of cartridges containing multi-volume datasets, and the number of cartridges containing multiple datasets.

- The Utilization summary report (XPR0091I) lists, for each LSM and again for each user-defined dataset workload group, the number of non-scratch cartridges within each tape utilization percentage band.

Fields in these reports are as follows:

- ACS: The ACS being reported.

- LSM: The LSM being reported.

- DSGRP: The dataset workload group being reported.

- OLDEST IN DAYS: The highest number of days since the last-reference date, i.e., the oldest tape within the LSM or dataset workload group.

- # MULTI VOLUMES: The number of multi-volume cartridges (more than one cartridge per file) held within the LSM or dataset workload group.

- # MULTI FILES: The number of multi-file cartridges (more than one file per cartridge) held within the LSM or dataset workload group.

- # VOLS/AVG AGE (1st occurrence): The number of non-scratch cartridges (#VOLS) that are older than the highest age band and the average number of days since their last-reference date (AVG AGE) within the LSM or dataset workload group. The age band range (in days) is printed above the occurrence.

- # VOLS/AVG AGE (2nd through 5th occurrences): For each age band, the number of non-scratch cartridges (# VOLS) that are within the specified age band and the average number of days since their last reference (AVG AGE). The age band range (in days) is printed above each occurrence.

- UTILIZED VOLUMES: The total number of entries for which utilization calculations have been performed. For LSM, this is the number of non-scratch cartridges matched in the tape catalog and HSC CDS. For dataset workload groups, this is the number of cartridges in the group.

- AVERAGE % UTIL: The LSM or group average of how much of the cartridges were utilized, expressed as a percentage of the length.

- *xx-xx*%: A series of percentage utilization bands showing the number of cartridges in each band.

TAPECAT Nearline Volume Details Report

NEARLINE VOLUME DETAILS REPORT   (AS AT 1994.340 17:32:20)

SORTED BY - VOLUME SERIAL NO

VOLUME SERIAL	-------------- DATASET NAME --------------	VOL SEQ	DS NS	PER CNT	MEGA-BYTES	LAST REF/ USED DATE	ACS LSM	USAGE COUNT	DEVICE-TYPE/MEDIA	SCRA TCH	FIL SEQ	BLOCK COUNT	BLOCK SIZE	REC LRECL FM
000002	HSM.HMIGTAPE.DATASET	1	1	47	646	1995100	000 01	100	3490I 1100	1		41369	16384	16384 F
000003	HSM.HMIGTAPE.DATASET	1	1	28	395	1995120	000 00	394	3490I 1100	1		25299	16384	16384 F
000004	HSM.HMIGTAPE.DATASET	1	1	37	513	1995110	000 01	301	3490I 1100	1		32812	16384	16384 F
000005	HSM.HMIGTAPE.DATASET	1	1	35	491	1995123	000 00	630	3490I 1100	1		31415	16384	16384 F
000010	PROD.IB.IBCA.IC1.DTH3DDA.DTH3	6	5	100	778	1995090	000 02	535	3490I 1100	1		43482	32760	28014 VBS
000014	HSM.BACKTAPE.DATASET	1	1	31	434	1995089	000 02	246	3490I 1100	1		27761	16384	16384 F
000017	PROD.OB.RPDS.BACKUP.G4631V00	1	3	0	0	1995111	000 02	345	3490I 1100	1		0	32720	80 FB
000018	PROD.IB.IBCA.IC1.DTPOLDA	8	6	100	804	1995083	000 02	242	3490I 1100	1		44973	32760	28014 VBS
000019	PROD.IB.IBCA.IC1.DTACDDA	1	3	100	886	1995101	000 02	389	3490I 1100	1		2776	32760	28014 VBS
000020	IDMS.LIBRARY.WEEKLY.BACKUP	1	1	7	102	1995090	000 00	300	3490I 1100	1		6551	0	0 UN
000021	PROD.GO.GLP.DGLPP001.SGLPP01	1	20	11	210	1995101	000 02	269	3490I 1100	1		1	28672	4096 FB
000030	HSM.HMIGTAPE.DATASET	1	1	30	417	1995112	000 01	423	3490I 1100	1		26685	16384	16384 F
000033	PROD.PU.PU001.G1405V00	1	1	2	50	1995080	000 01	375	3490I 1100	1		1588	32760	32756 VBS
000034	PROD.HK.DB2E.SYSDDF.IMAGE	1	16	2	38	1995108	000 01	212	3490I 1100	1		1	28672	4096 FB
000035	HSM.BACKTAPE.DATASET	1	1	31	427	1995047	000 02	236	3490I 1100	1		27339	16384	16384 F
000036	HSM.HMIGTAPE.DATASET	1	1	37	509	1995123	000 00	419	3490I 1100	1		32545	16384	16384 F
000038	HSM.BACKTAPE.DATASET	1	1	38	523	1994159	000 02	424	3490I 1100	1		33486	16384	16384 F
000044	XYZ.SAR.SARTAPE.T0012447	1	1	0	0	1995099	000 00	391	3490I 1100	1		0	0	0 UN
000047	PROD.OB.OB001.G0594V00	1	1	10	202	1995123	000 00	225	3490I 1100	1		6495	32640	128 FB
000048	PROD.IB.IBCA.IC1.DTRNDDA	9	2	100	198	1995083	000 02	127	3490I 1100	1		2962	32760	28014 VBS
000049	HSM.HMIGTAPE.DATASET	1	1	32	448	1995123	000 01	399	3490I 1100	1		28659	16384	16384 F
000050	HSM.HMIGTAPE.DATASET	1	1	28	391	1995123	000 00	245	3490I 1100	1		25030	16384	16384 F
000051	PROD.VSAM.STATS.BACKUP	1	1	1	36	1995123	000 01	517	3490I 1100	1		1158	32760	32756 VBS
000052	HSM.BACKTAPE.DATASET	1	1	31	437	1995037	000 02	267	3490I 1100	1		27966	16384	16384 F
000059	PROD.OG.DEPUBDA.IC1.G0087V00	1	5	100	1202	1995116	000 00	424	3490I 1100	1		1	32760	4105 VBS
000062	HSM.HMIGTAPE.DATASET	1	1	36	501	1995123	000 00	725	3490I 1100	1		32086	16384	16384 F
000068	HSM.HMIGTAPE.DATASET	1	1	28	395	1995120	000 00	397	3490I 1100	1		25293	16384	16384 F
000069	PROD.IB.IBCA.IC1.DTH3DDA	5	1	100	827	1995106	000 01	119	3490I 1100	1		46250	32760	28014 VBS
000071	PROD.PR.PR301.PRE.BATCH	1	1	7	134	1995115	000 00	209	3490I 1100	1		4295	32760	32756 VBS
000072	HSM.HMIGTAPE.DATASET	1	1	30	418	1995116	000 01	270	3490I 1100	1		26731	16384	16384 F

The Nearline Volume Details report lists, for each volume, the
primary dataset name (or optionally all primary and secondary
dataset names), volume sequence, number of datasets, percent used,
number of megabytes, last reference date, ACS/LSM location,
number of accesses, device type and cartridge length, scratch status,

Chapter 6: Tape Catalog Processing and Reports   **125**

and the file sequence, block count, block size, record size, and record format of each listed dataset.

This report is produced by the SYSIN control statement PERFORM TAPECAT OPTION(REPORT BY*xxxx*). Refer to appendix B, *SYSIN Parameters* for syntax information and additional reporting options.

Fields in this report are as follows:

- VOLUME SERIAL: The volser of the volume cartridge.

- DATASET NAME: The primary dataset name on the volume (and optionally all secondary datasets).

- VOL SEQ: The volume sequence within a multi-volume stack.

- DSNS: The number of datasets on the volume.

- PER CNT: The approximate percentage used of the volume. This is a percentage of the length of the cartridge that has been covered with data blocks and inter-block gaps.

- MEGA-BYTES: The approximate amount of data written to the volume. A "+" symbol beside this value indicates that the capacity has been capped.

- LAST REF/USED DATE: The last date the volume was last accessed.

- ACS LSM: The location of the volume within the Nearline library.

- USAGE COUNT: The number of times the volume has been accessed.

- DEVICE-TYPE/MEDIA: The creating device type and cartridge length. 3480I/3490I indicates that IDRC/ICRC compression was active for this cartridge.

- SCRATCH: The volume's scratch status. Blank indicates that the cartridge is not in scratch status. YES indicates the cartridge is a scratch. ??? indicates that message XPR0081W or XPR0082W was issued for this cartridge.

- FIL SEQ: The file sequence on a multi-file cartridge.

- BLOCK COUNT: The block count of the dataset.

- BLOCK SIZE: The block size of the dataset.

- LRECL: The record size of the dataset.

- RECFM: The record format of the dataset.

TAPECAT Nearline Volume Summary Report

NEARLINE VOLUME SUMMARY REPORT   (AS AT 1994.340 17:32:20)

PART 1 - VOLUME AGING AND TOTALS

ACS	LSM	DATA-GROUP	DEVICE-TYPE/MEDIA	TOTAL VOLUMES	IN-USE VOLUMES	SCRA-TCHES	UTILIZED VOLUMES	NOT IN CATALOG	MULTI VOLUMES	MULTI FILES	OLDEST DAYS	AVG DAYS	AVG MEGS	HIGH MEGS	CAPPED VOLUMES	TOTAL DSNS	MAX DSNS
000	00		ALL MEDIA	3134	2892	242	1547	1403	33	50	322	9	462	2275	3	2062	84
000	00		3480	1405	1345	60	0	1403	0	0	322	10	127	219	1	2	1
000	00		3480 0550	1405	1345	60	0	1403	0	0	322	10	127	219	1	2	1
000	00		3490	38	35	3	35	0	0	35	0	0	293	387	0	73	2
000	00		3490 1100	38	35	3	35	0	0	35	0	0	293	387	0	73	2
000	00		3490IDRC	1691	1512	179	1512	0	33	15	315	8	466	2275	2	1987	84
000	00		3490I 1100	1691	1512	179	1512	0	33	15	315	8	466	2275	2	1987	84
000	00	1	ALL MEDIA	2568	2568	0	2568	0	1339	1097	279	7	527	2257	84	19016	181
000	00	1	3480	9	9	0	9	0	8	0	161	45	64	157	0	9	1
000	00	1	3480 0550	9	9	0	9	0	8	0	161	45	64	157	0	9	1
000	00	1	3490IDRC	2559	2559	0	2559	0	1331	1097	279	7	529	2257	84	19007	181
000	00	1	3490I 1100	2559	2559	0	2559	0	1331	1097	279	7	529	2257	84	19007	181
000	00	1026	ALL MEDIA	23	23	0	23	0	21	0	0	0	744	934	0	23	1
000	00	1026	3490IDRC	23	23	0	23	0	21	0	0	0	744	934	0	23	1
000	00	1026	3490I 1100	23	23	0	23	0	21	0	0	0	744	934	0	23	1

.
.
.

NEARLINE VOLUME SUMMARY REPORT   (AS AT 1994.340 17:32:20)

PART 2 - VOLUME UTILIZATION

ACS	LSM	DATA-GROUP	DEVICE-TYPE/MEDIA	TOTAL VOLUMES	UTILIZED VOLUMES	AVG% UTIL	00-10%	11-20%	21-30%	31-40%	41-50%	51-60%	61-70%	71-80%	81-90%	91-100%
000	00		ALL MEDIA	3134	1547	30	193	80	527	532	92	45	32	10	6	30
000	00		3480	1405	0	0	0	0	0	0	0	0	0	0	0
000	00		3480 0550	1405	0	0	0	0	0	0	0	0	0	0	0
000	00		3490	38	35	40	2	3	3	3	24	0	0	0	0
000	00		3490 1100	38	35	40	2	3	3	3	24	0	0	0	0
000	00		3490IDRC	1691	1512	30	191	77	524	529	68	45	32	10	6	30
000	00		3490I 1100	1691	1512	30	191	77	524	529	68	45	32	10	6	30
000	00	1	ALL MEDIA	2568	2568	47	992	237	150	45	38	17	8	3	9	1069
000	00	1	3480	9	9	88	0	0	1	0	0	0	0	1	0	7
000	00	1	3480 0550	9	9	88	0	0	1	0	0	0	0	1	0	7
000	00	1	3490IDRC	2559	2559	47	992	237	149	45	38	17	8	2	9	1062
000	00	1	3490I 1100	2559	2559	47	992	237	149	45	38	17	8	2	9	1062
000	00	1026	ALL MEDIA	23	23	88	2	0	1	0	0	0	0	0	0	20
000	00	1026	3490IDRC	23	23	88	2	0	1	0	0	0	0	0	0	20

.
.
.

The Nearline Volume summary report provides tape catalog volume contents information summarized for each ACS, each LSM, each dataset workload group, each device type, each media type, and each defined tape length.

This report is produced by the SYSIN control statement PERFORM TAPECAT OPTION(SUMMARY).   Refer to appendix B, *SYSIN Parameters* for syntax information and additional reporting options.

The report is presented in two parts (due to page width constraints), volume aging and totals, and volume utilization.

Fields in this report are as follows:

- ACS: The ACS being reported.

- LSM: The LSM being reported.

- DATA-GROUP: The dataset workload group being reported.

- DEVICE-TYPE/MEDIA: The device type and cartridge length being reported.

- TOTAL VOLUMES: The total of all volumes resident within this LSM.

- IN-USE VOLUMES: The volumes not in scratch status.

- SCRATCHES: The volumes available as scratches

- UTILIZED VOLUMES: The volumes for which a utilization calculation was performed.

- NOT IN CATALOG: The number of volumes that were not in the tape catalog.

- MULTI-VOLUMES: The volumes that are part of a multi-volume stack.

- MULTI-FILES: The volumes with more than one dataset.

- OLDEST DAYS: The number of days since the oldest volume was last referenced.

- AVG DAYS: The average number of days since these volumes were last referenced.

- AVG MEGS: The average megabytes of data on these volumes.

- HIGH MEGS: The highest number of megabytes on any volume.

- CAPPED VOLUMES: The number of volumes to which message XPR0097W applied.

- TOTAL DSNS: The total of all primary and secondary datasets.

- MAX DSNS: The maximum number of datasets on any volume.

- AVG % UTIL: The average percentage utilization for UTILIZED VOLUMES.

- 00-10 / 91-100%: The ten percentage bands showing the spread of cartridge utilization.

TAPECAT Tape Catalog History Reports

The Nearline and tape catalog history report provides comparative activity analysis between two tape catalog images by listing volumes that have changed status and summarizing activity with the library and tape catalog.

The HISTORY report reads two generations of the TMCHIST GDG, as shown below.

```
//STEPABC  EXEC  EXPRPROC
//EXPR.TMCHIST1  DD  DSN=user.name.SPR110.TMCHIST(-1),
//               DISP=SHR,DCB=(BUFNO=20)
//EXPR.TMCHIST2  DD  DSN=user.name.SPR110.TMCHIST(0),
//               DISP=SHR,DCB=(BUFNO=20)
//EXPR.SYSIN      DD  *
 PERFORM  TAPECAT  OPTION(HISTORY) ;
```

Note: Always ensure TMCHIST1 points at the older dataset.

The HISTORY report lists those volumes that have changed status between the two runs of TAPECAT UPDATE that generated the TMCHIST files. These changes can be:

- a volume being added to the TMC or deleted from the TMC

- a volume being entered into the Nearline library or ejected from the Nearline library

- a volume becoming a scratch or non-scratch

- a volume recording mode/density changing (18/36 track)

Additionally, history file comparison statistics are listed at the end of the report.

```
----------------------------------------------------------------------------------------------------------------------------------

1996.015    S T O R A G E T E K  E X P R  1 . 1    XYZ COMPANY                                    PAGE  1
11:14:33
                        NEARLINE VOLUME HISTORY REPORT
                        -------------------------------------------------
    (1ST FILE CREATED: 1996.011 AT 08:02:12  2ND FILE CREATED: 1996.012 AT 08:03:24)

    VOLUME     CDS/ACS     TMC       CHANGED   BECOME    BECOME    VOLUME
    SERIAL     STATUS      STATUS    DENSITY   SCRTCH    NONSCR    RE-USED
    ----------------------------------------------------------------------------------------------

    000001                           YES       YES       YES
    000002     ENTERED               YES       YES       YES
    000005                           YES       YES       YES
    000006     EJECTED               YES       YES       YES
    000008                           YES       YES       YES      YES
    000010                           YES       YES       YES
    000011     EJECTED     DELETED
    000014                 DELETED
    000016                 DELETED
    000017     ENTERED               YES       YES       YES      YES
    000019     ENTERED               YES       YES       YES      YES
    000020                           YES       YES       YES
    000021     ENTERED     ADDED
    000023     EJECTED               YES       YES       YES
    000025                 ADDED
    000027     ENTERED     ADDED


    HISTORY FILE COMPARISON STATISTICS :
    VOLUMES ENTERED.....................................  70
    VOLUMES EJECTED.....................................  12
    VOLUMES MOVED BETWEEN ACS/LSM'S.............  356
    VOLUMES ADDED TO THE TMC......................  50
    VOLUMES DELETED FROM THE TMC...................  11
    VOLUMES WITH CHANGED DENSITY/MODE........  22
    VOLUMES THAT HAVE BECOME SCRATCH...........  514
    VOLUMES THAT HAVE BECOME NON-SCRATCH.  558
    VOLUMES THAT HAVE BEEN RE-USED..................  14

    VOLUMES THAT HAVE BEEN READ AS INPUT......  381
    TOTAL NO OF INPUT OPENS.......................  435

    TMCHIST1 VOLUMES READ......................................  45935
    TMCHIST1 VOLUMES IN LIBRARY............................  15993
    TMCHIST1 SCRATCHES IN LIBRARY.......................  1265
    TMCHIST1 VOLUMES IN TMC.....................................  41250
    TMCHIST2 VOLUMES READ......................................  45993
    TMCHIST2 VOLUMES IN LIBRARY............................  16051
    TMCHIST2 SCRATCHES IN LIBRARY........................  1228
    TMCHIST2 VOLUMES IN TMC.....................................  41289


----------------------------------------------------------------------------------------------------------------------------------
```

CA-TLMS Considerations

Installations that use the CA-TLMS tape management system must run an extra job before the ExPR TAPECAT UPDATE function. This job is a TLMS report run, as documented in the ExPR SAMPLIB members TLMS53 and TLMS54. Users of CA-TLMS 5.3 (or earlier) must use job TLMS53 and users of CA-TLMS 5.4 (or later) must run job TLMS54. This is due to differences in the TLMS018 user report definitions.

The ExPR/TLMS interface module does not directly read the tape catalog (as with CA-1), but extracts the required information from TLMS reports. These reports are read from the TLMSRPT dataset by the TLMS interface module. SAMPLIB member TAPECATT is an example of running the TAPECAT function for CA-TLMS installations.

It is recommended that the TLMS report step be run immediately before the ExPR TAPECAT UPDATE function. This will ensure that the tape catalog and CDS are in synchronization.

Running Without a Tape Catalog

If a tape catalog is not available to ExPR, limited information can be extracted from the CDS alone. This will include volume aging statistics only, and the TMCHIST file will be empty. To operate without a tape catalog, specify the following TMS statement in the CONFIG file.

```
TMS MODNAME(IEFBR14) PARM(EXPRFORM) ;
```

Chapter 7

Mainframe Performance Reports

Overview

This chapter describes mainframe tabular performance reports that are produced by ExPR.

Note: You can also generate performance reports in a graphical format on a PC with ExPR. Graphical reports are described in the *ExPR PC User's Guide*.

Generating Reports

Mainframe performance reports are generated by running user report requests and selection parameters against the ExPR mainframe database as a batch task. ExPR mainframe performance reports are formatted as tabular hardcopy output.

Mainframe performance reports provide a simple and quick mechanism for report generation without the need for advanced graphics facilities. If more sophisticated reporting is required, use the ExPR PC graphical reports. If more flexible mainframe reports are required, refer to the ExPR API, which is described in chapter 9, *ExPR API*.

The procedure for generating mainframe reports can be considered in two parts: specifying selection criteria and running the batch task to build the report. Report availability is subject to the raw data being made available to ExPR during update processing.

Mainframe performance reports include the following:

- The MOUNTS report lists the number of Nearline mounts, average mount time, and maximum mount time. Separate tallies are reported for scratch and non-scratch volumes. Total global (including non-Nearline) mounts are also reported separately.

- The MOUNTS-DETAIL report provides a breakdown of the mount response time components for both scratch and non-scratch mounts.

- The CONTENTS report lists cartridge movements (enter/eject count and number of passthroughs) and contents information (scratch count, free cells, total cells).

- The UTILIZATION report lists the percentage of time the robotics system was in use and the percentage of time that drives were concurrently in use.

- The CU-BUSY report produces control unit information, including the number of drives with a disconnect time exceeding the user-defined threshold.

- The PATH-BUSY report produces information about user-defined channel group utilization.

- The TAPE-ERRORS report provides listings of the number of temporary and permanent read/write errors that occurred. The report also provides an audit trail of which drive and volume the media errors occurred on.

- The ALLOC-REC report provides information about allocation recovery tasks on JES2 systems, including the number of events and the delay times.

- The DEMAND-ENTERS report provides information about demand enters that occurred during the reporting period.

Report samples and information field summaries follow in this chapter. Mainframe performance reports are generated by the SYSIN control statement PERFORM REPORT NAME(*reportname*).

ACTION:

⇒ Generate mainframe performance reports using SAMPLIB member MFREPS.

⇒ Refer to appendix B, *SYSIN Parameters* for information about runtime parameters that appear in the SAMPLIB.

Printing SYSOUT Information

By default, the SYSIN controls that were used to generate performance reports are printed with the report. Three other SYSOUT listings are not printed automatically; these are the Configuration listing, the trace Analysis listing, and auto-CONFIG Generated statements listing. These listings can be produced along with or in place of the SYSIN controls listing with the OPTIONS control card.

The OPTIONS control card has the following syntax:

+	Prefix to include an output type
-	Prefix to exclude an output type
S	Runtime statements (from SYSIN DD)
C	Configuration statements (from CONFIG DD)
G	Generated statements (from CDSCOPY DD)
A	Analysis report (trace use only)

The default settings for these parameters includes SYSIN but excludes Configuration , Generated, and Analysis listing. If this is acceptable, it is not necessary to code the PARM statement.

For example, the default

```
//STEPn EXEC PGM=EXPRPROC,PARM='OPTIONS(+S-CGA)'
```

lists just SYSIN runtime statements, while the coded statement

```
//STEPn EXEC PGM=EXPRPROC,PARM='OPTIONS(-S+CGA)'
```

would list everything except the SYSIN statements.

Order of SELECT Statement Processing

ExPR processes the complete generated AUTO-CONFIG, CONFIG, and SYSIN files before acting on any update or report requests.

The reports are then generated in the same order as the PERFORM REPORT statements, but will be produced subject to the final status of the SELECT SYSID criteria.

Thus, the following example would result in three reports for system CPUA:

```
SELECT SYSID(MVSA) ;
PERFORM REPORT NAME(MOUNTS) ;
PERFORM REPORT NAME(CU-BUSY) ;
SELECT SYSID(CPUA) ;
PERFORM REPORT NAME(ALLOC-REC) ;
/*
```

Selection by ACS/LSM-ids is cumulative, however. All requested reports will process all selected ACS/LSM-ids regardless of the order of SELECT and REPORT statements.

A returned RC=8 from any ExPR report request indicates that the database did not contain records to match your selection criteria.

MOUNTS Report

```
                        MOUNTS REPORT FOR SYSTEM CPUA          DATE RANGE: 1993001 TO 1995365
                        ACS: 000 NEARLINE ACS 000
                        LSM: 02 NEARLINE LSM 000 02
```

---PERIOD---		DEVICE-TYPE/	-----------------------------NEARLINE MOUNTS-----------------------------						---GLOBAL MOUNTS--	
DATE	HOUR	WORKLOAD	NUMBER OF MOUNTS		AVERAGE MOUNT TIME		MAXIMUM MOUNT TIME		NUMBER OF MOUNTS	
			SCRATCH	NON-SCR	SCRATCH	NON-SCR	SCRATCH	NON-SCR	SCRATCH	NON-SCR
1994341	00	NO DATA								
	01	"								
	02	"								
	03	"								
	04	"								
	05	"								
	06	ALL DEVICES	9	3	28	30	60	41	9	7
	07	ALL DEVICES	7	4	26	38	36	51	20	8
	08	ALL DEVICES	11	0	21	0	32	0	21	15
	09	ALL DEVICES	11	1	27	28	61	28	14	47
	10	ALL DEVICES	8	3	24	34	40	42	15	69
	11	ALL DEVICES	6	8	38	38	80	80	7	52
	12	ALL DEVICES	12	5	28	35	41	48	7	40
	13	ALL DEVICES	10	6	33	20	63	39	12	46
	14	ALL DEVICES	5	11	22	18	46	23	3	31
	15	ALL DEVICES	4	7	33	27	38	59	8	65
	16	ALL DEVICES	4	5	30	30	36	41	7	67
	17	ALL DEVICES	3	3	32	58	33	86	8	39
	18	ALL DEVICES	3	3	51	28	75	34	38	34
	19	ALL DEVICES	15	1	32	19	54	19	7	7
	20	ALL DEVICES	15	4	33	30	81	45	39	6
	21	ALL DEVICES	15	4	35	30	72	46	20	6
	22	ALL DEVICES	14	1	34	17	70	17	24	7
	23	ALL DEVICES	16	0	27	0	42	0	7	5

The MOUNTS report is generated per-LSM and shows one day of data per page. The report title section shows the ACS/LSM being reported on, with any user annotation from the CONFIG file and the user-selected date range.

This report is produced by the SYSIN control statement PERFORM REPORT NAME(MOUNTS). Refer to appendix B, *SYSIN Parameters* for syntax information.

The MOUNTS report provides the following types of information:

- Nearline Mounts: Mount event statistics within the Nearline environment, including the number of mounts, the average mount time (in seconds), and the highest recorded time to mount a volume over the interval. Each report is further categorized by scratch and non-scratch volumes.

- Global Mounts: Mount events (including those outside the Nearline environment), showing the number of scratch and non-scratch mounts for this MVS system.

Mounts that exceed the user-defined thresholds are listed individually on the reports generated by SMFUPDATE or SMFREPORT.

The MOUNTS report can produce additional summary lines for individual device-types and user-defined workload groups. This facility is controlled by the REPORT-OPTIONS statement. Refer to appendix B, *SYSIN Parameters* for syntax information.

Note: It is possible that the Nearline scratch and non-scratch mount counts may exceed the global system mounts for a given hour. This is due to the different data sources used by ExPR. The Nearline counts are derived from HSC SMF records (written at the time of the mount event), whereas the global values are extracted from SMF 21 dismount records and HSC SMF records.

MOUNTS-DETAIL Report

```
MOUNTS-DETAIL REPORT FOR SYSTEM CPUA                          DATE RANGE: 1994340 TO 1994342
ACS: 000 NEARLINE ACS 000
LSM: 02 NEARLINE LSM 000 02
```

			<--- SCRATCH MOUNT BREAKDOWN DETAILS --->											<--- NON-SCRATCH MOUNT BREAKDOWN DETAILS --->										
DATE HOUR	WORKLOAD	NO MNTS	AVG TIME	TOTL TIME	HSC Q	DRIVE WAIT	DEST ROB	DEST ROBQ	OTHR ROB	OTHR ROBQ	PTHR ROB	PTHR ROBQ	NO MNTS	AVG TIME	TOTL TIME	HSC Q	DRIVE WAIT	DEST ROB	DEST ROBQ	OTHR ROB	OTHR ROBQ	PTHR ROB	PTHR ROBQ
1994341 00	NO DATA																						
01	"																						
02	"																						
03	"																						
04	"																						
05	"																						
06	ALL DEVICES	9	28	252	86	78	81	7	0	0	0	0	3	30	90	32	27	27	4	0	0	0	0
07	ALL DEVICES	7	26	182	62	59	59	2	0	0	0	0	4	38	152	75	34	34	0	4	3	2	0
08	ALL DEVICES	11	21	231	30	93	108	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
09	ALL DEVICES	11	27	297	109	94	94	0	0	0	0	0	1	28	28	13	8	7	0	0	0	0	0
10	ALL DEVICES	8	24	192	45	76	71	0	0	0	0	0	3	34	102	49	25	28	0	0	0	0	0
11	ALL DEVICES	6	38	228	123	50	55	0	0	0	0	0	8	38	304	99	75	82	0	31	5	11	1
12	ALL DEVICES	12	28	336	90	105	124	17	0	0	0	0	5	35	175	80	42	45	8	0	0	0	0
13	ALL DEVICES	10	33	330	137	85	103	5	0	0	0	0	8	34	272	129	65	59	19	0	0	0	0
14	ALL DEVICES	5	22	110	23	42	37	8	0	0	0	0	11	18	198	10	93	95	0	0	0	0	0
15	ALL DEVICES	4	33	132	63	34	35	0	0	0	0	0	7	27	189	35	59	71	0	12	11	1	0
16	ALL DEVICES	4	30	120	54	34	32	0	0	0	0	0	5	30	150	59	42	49	0	0	0	0	0
17	ALL DEVICES	3	32	96	50	25	21	0	0	0	0	0	3	58	174	120	25	29	0	0	0	0	0
18	ALL DEVICES	3	51	153	94	25	34	0	0	0	0	0	3	28	84	38	25	21	0	0	0	0	0
19	ALL DEVICES	15	32	480	184	130	138	28	0	0	0	0	1	19	19	1	8	10	0	0	0	0	0
20	ALL DEVICES	15	33	495	208	127	160	0	0	0	0	0	4	30	120	43	34	34	9	0	0	0	0
21	ALL DEVICES	15	35	525	264	127	121	13	0	0	0	0	4	30	120	36	34	36	0	12	0	2	0
22	ALL DEVICES	14	34	476	242	119	112	3	0	0	0	0	1	17	17	1	8	8	0	0	0	0	0
23	ALL DEVICES	16	27	432	105	137	147	43	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

The MOUNTS-DETAIL report provides a breakdown of the statistics produced in the MOUNTS report, particularly the components that make up the total response time of a Nearline mount request.

This report is produced by the SYSIN control statement PERFORM REPORT NAME(MOUNTS-DETAIL). Refer to appendix B, *SYSIN Parameters* for syntax information.

The report is in two sections, scratch and non-scratch. For each section, the MOUNTS-DETAIL report provides the following types of information on an hourly basis:

- The number of mounts requested.

- The average time in seconds to service a mount.

- The total time in seconds spent servicing mounts.

- HSC queue time: The time spent while HSC processed other events.

- Drive Wait: The duration of tape drive threading/loading.

- Dest Robotics: The time the robotics arm was busy/moving.

- Dest Robotics Queue: The time awaiting the robotics arm to become free.

- Other Robotics: The time robotics arm in the adjacent LSM was busy.

- Other Robotics Queue: The time awaiting the robotics arm in the adjacent LSM.

- Passthru Robotics: The time cartridges were traveling between LSMs.

- Passthru Robotics Queue: The time awaiting the passthrough port to become free.

Destination Robotics refers to the robotics arm of the LSM that is attached to the allocated drive.

Other Robotics refers to the robotics activities of all other LSMs involved in moving the cartridge from its cell to the destination LSM.

The MOUNTS-DETAIL report can produce additional summary lines for individual device-types. This facility is controlled by the REPORT-OPTIONS statement. Refer to appendix B, *SYSIN Parameters* for syntax information.

UTILIZATION Report

UTILIZATION MOUNTS REPORT FOR SYSTEM CPUA   DATE RANGE: 1993001 TO 1995365
ACS: 000 NEARLINE ACS 000
LSM: 02 NEARLINE LSM 000 02

DATE	HOUR	DEVICE-TYPE/ WORKLOAD	DRIVE UTIL% SCR	DRIVE UTIL% NON-SCR	ROBOTICS UTIL %	\--------- PERCENTAGE OF TIME DRIVES WERE IN USE \------------ NO OF IN USE DRIVES																	THRESHOLD DRIVESINUSE(06)
						00	01	02	03	04	05	06	07	08	09	10	11	12	13	14	15	16	
1994341	00	NO DATA																					
	01	"																					
	02	"																					
	03	"																					
	04	"																					
	05	"																					
	06	ALL DEVICES	68	19	38	0	1	6	24	69	0	0	0	0	0	0	0	0	0	0	0	0	
	07	ALL DEVICES	71	22	6	0	0	7	7	86	0	0	0	0	0	0	0	0	0	0	0	0	
	08	ALL DEVICES	70	0	3	0	1	36	40	23	0	0	0	0	0	0	0	0	0	0	0	0	
	09	ALL DEVICES	95	2	5	0	0	0	8	92	0	0	0	0	0	0	0	0	0	0	0	0	
	10	ALL DEVICES	84	11	6	0	0	1	12	87	0	0	0	0	0	0	0	0	0	0	0	0	
	11	ALL DEVICES	44	37	7	0	2	15	32	51	0	0	0	0	0	0	0	0	0	0	0	0	
	12	ALL DEVICES	60	22	9	0	1	20	23	56	0	0	0	0	0	0	0	0	0	0	0	0	
	13	ALL DEVICES	37	10	9	0	41	30	25	4	0	0	0	0	0	0	0	0	0	0	0	0	
	14	ALL DEVICES	18	15	8	18	47	22	10	3	0	0	0	0	0	0	0	0	0	0	0	0	
	15	ALL DEVICES	23	13	7	1	59	30	10	0	0	0	0	0	0	0	0	0	0	0	0	0	
	16	ALL DEVICES	24	17	4	0	29	71	0	0	0	0	0	0	0	0	0	0	0	0	0	0	
	17	ALL DEVICES	24	24	4	0	4	96	0	0	0	0	0	0	0	0	0	0	0	0	0	0	
	18	ALL DEVICES	20	11	6	15	39	46	0	0	0	0	0	0	0	0	0	0	0	0	0	0	
	19	ALL DEVICES	55	2	12	3	23	8	66	0	0	0	0	0	0	0	0	0	0	0	0	0	
	20	ALL DEVICES	72	7	9	0	0	10	58	32	0	0	0	0	0	0	0	0	0	0	0	0	
	21	ALL DEVICES	79	11	9	0	0	1	30	69	0	0	0	0	0	0	0	0	0	0	0	0	
	22	ALL DEVICES	88	3	11	0	0	6	13	81	0	0	0	0	0	0	0	0	0	0	0	0	
	23	ALL DEVICES	85	0	7	0	4	8	24	64	0	0	0	0	0	0	0	0	0	0	0	0	

The UTILIZATION report is generated per-LSM and shows one day of data per page. The report title section shows the ACS/LSM being reported on, with any user annotation from the CONFIG file and the user-selected date range.

This report is produced by the SYSIN control statement PERFORM REPORT NAME(UTILIZATION). Refer to appendix B, *SYSIN Parameters* for syntax information.

The UTILIZATION report provides the following types of information:

Drive Util%: A percentage breakdown between scratch and non-scratch volumes mounted on the drives.

Robotics: The percentage of time over an interval the robotics arm was performing some kind of action.

Drives in Use: The percentage of time a number of drives were concurrently in use.

EXCEEDED: This messages is printed in the right-most column when the user-defined threshold DRIVESINUSE is exceeded.

The UTILIZATION report can produce additional summary lines for individual device-types. This facility is controlled by the REPORT-OPTIONS statement. Refer to appendix B, *SYSIN Parameters* for syntax information.

CONTENTS Report

```
-----------------------------------------------------------------------------------------------------------------------------------
1994.340      S T O R A G E T E K  E X P R  1 . 1      XYZ COMPANY                                              PAGE 1
17:52:29
-----------------------------------------------------------------------------------------------------------------------------------
                                    CONTENTS REPORT FOR SYSTEM CPUA            DATE RANGE: 1993001 TO 1995365
                                    ACS: 000 NEARLINE ACS 000
                                    LSM: 02 NEARLINE LSM 000 02
-----------------------------------------------------------------------------------------------------------------------------------
---PERIOD----          -------CARTRIDGE MOVEMENTS----------   -----------------------LIBRARY CONTENTS-------------------
DATE   HOUR                                        DEMAND      SCRATCH COUNTS   --FREE CELLS---  TOTAL CELLS
                       PASSTHRUS EJECTS ENTERS ENTERS   MIN AVG MAX     MIN AVG MAX

1994341  00    NO DATA
         01    "
         02    "
         03    "
         04    "
         05    "
         06    PARTIAL      2       0      131      3
         07    "            1       0       7       0
         08    "            0       0       0       0
         09    "            1       0       0       0
         10    "            1       0       0       0
         11    "            4       0       0       0
         12    "            4       0       0       0
         13    "            0       0       0       0
         14    "            0      21       0       0
         15    "            3       0       0       0
         16                 2       0       0       0    332  332 332   170 170  170    5400
         17                 1       0       0       0    332  332 332   170 170  170    5400
         18                 0       0       0       0    332  332 332   170 170  170    5400
         19    PARTIAL      1       0       0       0
         20    "            1       0       0       0
         21    "            1       0       0       0
         22    "            4       0       0       0
         23    "            1       0       0       0
-----------------------------------------------------------------------------------------------------------------------------------
```

The CONTENTS report is generated per-LSM and shows one day of data per page. The report title section shows the ACS/LSM reported on, with any user annotation from the CONFIG file and the user-selected date range.

This report is produced by the SYSIN control statement PERFORM REPORT NAME(CONTENTS). Refer to appendix B, *SYSIN Parameters* for syntax information.

The CONTENTS report provides the following types of information:

- Cartridge Movements: Statistics on the movement of cartridges in, out, and around LSMs, including enters, ejects, and passthroughs.

- Library Contents: LSM cell information, including scratch count, free cells, and total cells.

CU-BUSY Report

```
                        CONTROL UNIT BUSY REPORT FOR SYSTEM CPUA                   DATE RANGE: 1994340 TO 1994342
                        CONTROL UNIT: 0C00 (LCU#  54)
                        DEVICES: 0C00 THRU 0C0F

ACS+LSM  ====>     00001   00001 00001 00001                    00000 00000   00000 00000
DEVICE-TYPE ====>
DEVICE-ADDR ====> -CU-  0C00    0C01 0C02 0C03 0C04 0C05   0C06 0C07 0C08 0C09   0C0A 0C0B 0C0C 0C0D 0C0E 0C0F
          TOTAL <--------------------- CONTROL UNIT AND INDIVIDUAL DEVICE PERCENTAGE DISCONNECTED --------------------->THRESHOLD
DATE HOUR  SSCHS  AV HI AV HI AV HI AV HI AV HI AV HI AV HI AV HI AV HI AV HI AV HI AV HI AV HI AV HI AV HI AV HI CULOAD(80)

1994341 00    NO DATA
        01     "
        02     "
        03     "
        04     "
        05     "
        06   295508  43 94   60 74 30 39 49 61 47 73   0 0 0 0 0 0 0 0   87 94 24 31   35 44 44 49 0 0 0 0 0 0 0 0  *EXCEEDED*
        07   394380  44 85   68 77 46 51 81 85 47 85   0 0 0 0 0 0 0 0   41 50 37 49   39 47 57 59 0 0 0 0 0 0 0 0  *EXCEEDED*
        08   223900  35 81   54 62 45 48 27 81 27 81   0 0 0 0 0 0 0 0   28 62 21 33   34 35 30 47 0 0 0 0 0 0 0 0  *EXCEEDED*
        09   138741  18 67   12 60 52 55 1  1  1  1    0 0 0 0 0 0 0 0   39 61 6  58   4  55 39 67 0 0 0 0 0 0 0 0
        10   186768  33 67   7  26 61 66 48 67 48 67   0 0 0 0 0 0 0 0   51 60 21 35   31 62 24 63 0 0 0 0 0 0 0 0
        11   219755  26 66   22 56 58 59 25 66 25 66   0 0 0 0 0 0 0 0   9  40 0  0    28 40 29 37 0 0 0 0 0 0 0 0
        12   302194  14 79   19 79 44 58 1  3  1  3    0 0 0 0 0 0 0 0   8  26 59 59   16 20 21 34 0 0 0 0 0 0 0 0
        13   55157   17 95   26 92 60 62 6  69 6  69   0 0 0 0 0 0 0 0   35 49 30 95   5  58 57 63 0 0 0 0 0 0 0 0  *EXCEEDED*
        14   31351   19 66   23 66 15 54 28 60 28 60   0 0 0 0 0 0 0 0   4  58 57 62   3  57 29 59 0 0 0 0 0 0 0 0
        15   51265   24 70   4  36 3  11 51 58 51 58   0 0 0 0 0 0 0 0   41 55 52 57   14 70 23 63 0 0 0 0 0 0 0 0
        16   24142   40 77   51 77 0  0  54 61 54 61   0 0 0 0 0 0 0 0   0  0  34 62   63 64 20 58 0 0 0 0 0 0 0 0
        17   27082   23 85   40 85 0  0  5  12 5  12   0 0 0 0 0 0 0 0   0  0  15 57   10 54 9  58 0 0 0 0 0 0 0 0  *EXCEEDED*
        18   110018  22 85   75 85 0  0  57 64 0  0    0 0 0 0 0 0 0 0   0  0  8  79   28 53 12 20 0 0 0 0 0 0 0 0  *EXCEEDED*
        19   151389  51 91   76 85 0  0  86 91 30 40   0 0 0 0 0 0 0 0   86 86 41 57   30 86 50 67 0 0 0 0 0 0 0 0  *EXCEEDED*
        20   309177  51 85   72 84 48 56 60 85 29 38   0 0 0 0 0 0 0 0   48 56 0  0    0  0  61 74 0 0 0 0 0 0 0 0  *EXCEEDED*
        21   170785  54 92   71 75 39 56 76 92 57 88   0 0 0 0 0 0 0 0   39 56 0  0    48 66 70 78 0 0 0 0 0 0 0 0  *EXCEEDED*
        22   186034  43 91   58 71 23 31 83 91 73 78   0 0 0 0 0 0 0 0   23 31 0  0    48 74 60 75 0 0 0 0 0 0 0 0  *EXCEEDED*
        23   153333  20 44   0  0  0  0  33 44 26 33   0 0 0 0 0 0 0 0   0  0  0  0    10 23 8  11 0 0 0 0 0 0 0 0
```

The CU-BUSY report is generated for each tape control unit defined in the ExPR configuration. The headings identify the address of each control unit and its logical control unit number (LCU) from the system IOCP/IOCDS.

This report is produced by the SYSIN control statement PERFORM REPORT NAME(CU-BUSY). Refer to appendix B, *SYSIN Parameters* for syntax information.

The CU-BUSY report provides the following types of information:

- Total SSCHs: The total start subchannel commands (physical SIOs) to devices attached to this control unit.

- CU Average: The average percentage of disconnected time for this control unit.

- CU Highest: The highest percentage of disconnected time of any device attached to this control unit.

- Device Average: The average percentage of disconnected time for each device attached (0-F).

- Device Highest: The highest percentage of disconnected time for each device attached (0-F).

- CULOAD Threshold: The user-defined control unit loading threshold. *EXCEEDED* is printed to indicate that this value has been surpassed. (Excessive control unit busy is also listed as an exception in the SMFUPDATE report.)

The percentage disconnect time is collected by RMF on a per-device basis. Disconnected time is defined as the time spent by the device waiting to transfer data across the already busy control unit interface.

Note: When the devices are 9490 or SD-3 (Timberline or Redwood), each device actually has its own integrated control unit. In this case, you should treat each pair of device columns within the CU-BUSY report as an individual control unit. The SSCH count will be for all devices/control units. When you review the data on the PC, you should also select the "device-specific" option.

PATH-BUSY Report

CHANNEL PATH GROUP BUSY REPORT FOR SYSTEM CPUA   DATE RANGE: 1994340 TO 1994342

DATE	HOUR		CPG 0 %	CPG 1 %	CPG 2 %	CPG %	CPG %	CPG %	CPG %	CPG %	CPG %	CPG %	CPG %	CPG %	CPG %	CPG %	CPG %	CPG %	CPG %	CPG %	CPG %	CPG %	CPG %	CPG %	CPG %	CPG %
1994341	00	NO DATA																									
	01	"																									
	02	"																									
	03	"																									
	04	"																									
	05	"																									
	06		6	2	4																						
	07		6	2	5																						
	08		6	2	5																						
	09		6	2	5																						
	10		6	2	5																						
	11		6	2	5																						
	12		6	2	5																						
	13		6	2	4																						
	14		6	2	4																						
	15		6	2	4																						
	16		6	2	4																						
	17		6	2	4																						
	18		6	2	4																						
	19		6	2	4																						
	20		6	2	4																						
	21		6	2	4																						
	22		6	2	4																						
	23		6	2	4																						

The PATH-BUSY report produces information about user-defined channel group utilization. This report is produced by the SYSIN control statement PERFORM REPORT NAME(PATH-BUSY). Refer to appendix B, *SYSIN Parameters* for syntax information.

The report is generated as one page per day for each system defined. Up to 30 user-defined channel path groups are listed across the page. The user-defined group numbers (0-255) are inserted in the column headings. The percentage busy time is printed under each defined group. This percentage is derived by totaling all RMF samples for all channels in a specific group and all RMF busy samples within the group.

TAPE-ERRORS Report

GLOBAL SUMMARY OF TAPE ERRORS FOR SYSTEM CPUA   DATE RANGE: 1993001 TO 1995365

			<---TOTALS FOR ACS VOLUMES AND DRIVES--->						<-TOTALS FOR NON-ACS VOLUMES AND DRIVES->					
			<PERM ERRORS>		<TEMP ERRORS>		<NO OF VOLS>		<PERM ERRORS>		<TEMP ERRORS>		<NO OF VOLS>	
-DATE-	HOUR	SYSTEM	READ	WRITE	READ	WRITE	PERM	TEMP	READ	WRITE	READ	WRITE	PERM	TEMP
1994.341	22.00	CPUA	0	2	0	0	2	0	0	0	0	0	0	0
1994.342	00.00	CPUA	1	0	0	1	0	0	0	0	0	0	0	0

.
.

GLOBAL DETAILS OF TAPE ERRORS FOR SYSTEM CPUA   DATE RANGE: 1993001 TO 1995365

					UNIT	VOLUME	< PERM ERRORS >		< TEMP ERRORS >		NO OF	<DATA TRANSFERRED>	
-DATE-	HOUR	SYSTEM	ACS	LSM	ADDR	SERIAL	READ	WRITE	READ	WRITE	SSCH'S	READ	WRITTEN
1994.341	22.00	CPUA	000	02	0C27	013539	0	1	0	0	1207	12K	965M
1994.341	22.00	CPUA	000	01	0C00	013539	0	1	0	0	191	4K	28M
1994.342	00.00	CPUA	000	01	0C02	017003	0	1	0	0	275	268K	35M

.
.

TAPE-ERRORS reports provide summary and detail listings of the number of temporary and permanent read/write errors that occur; the detail report also provides an audit trail of which tape drive and cartridge volume the media errors occurred on.

The reports are produced by the SYSIN control statement PERFORM REPORT NAME(TAPE-ERRORS). Refer to appendix B, *SYSIN Parameters* for syntax information.

The Global Summary report has two sections, one for Nearline volumes and devices and another for non-Nearline volumes and devices. The following information is listed for each section:

- The number of permanent read errors recorded.

Chapter 7: Mainframe Performance Reports   **151**

- The number of permanent write errors recorded.

- The number of temporary read errors recorded.

- The number of temporary write errors recorded.

- The number of volumes that had permanent errors.

- The number of volumes that had temporary errors.

The Global Details report lists one record per drive/volume that had I/O errors associated with it. For each drive/volume, the following information is printed:

- The ACS and LSM that owned the drive (these are blank for non-Nearline devices).

- The device address of the unit involved.

- The volser of the cartridge.

- Counts of permanent/temporary read and write errors.

- The SSCH count (number of physical I/Os to the volume).

- The quantity of data read and written to the volume.

The SSCH count taken with the data transferred can give a ratio of errors to I/O activity. For example, one temporary error on a volume of 800Mb with 20,000 physical I/Os is not as important as a permanent error while writing tape labels.

Note: Tape errors are also listed as exceptions by the SMFUPDATE and SMFREPORT functions.

ALLOC-REC Report

MVS ALLOCATION RECOVERY SUMMARY FOR SYSTEM CPUA   DATE RANGE: 1993001 TO 1995365

					<--- ACS/LSM SPECIFIC --->					<--- GLOBAL/COMPLETE SYSTEM --->			
-DATE-	HOUR	SYSTEM	ACS	LSM	NO OF OPERATOR REPLIES	TOTAL OPERATOR REPLY TIME	LONGEST OPERATOR REPLY	NO OF WAITS SERVICED	LONGEST WAIT SERVICED	NO OF CANCEL REPLIES	NO OF WAIT REPLIES	HIGHEST NO WAITS PENDING	TIME ANY WAIT WAS PENDING	TIME ANY REPLY WAS PENDING
1994.341	11.00	CPUA	000	02	0	00.00.00	00.00.00			1				00.08.23
1994.341	13.00	CPUA	000	02	1	00.00.06	00.00.06							00.01.12
1994.341	14.00	CPUA	000	02	1	00.00.03	00.00.03	1	00.13.05		1	1	00.13.05	00.05.55
1994.341	15.00	CPUA	000	02	1	00.00.06	00.00.06			1				00.01.00
1994.341	16.00	CPUA	000	02	0	00.00.00	00.00.00			1				00.00.37
1994.341	17.00	CPUA	000	02	0	00.00.00	00.00.00			1				00.00.52
1994.342	14.00	CPUA	000	02	1	00.00.07	00.00.07							00.00.07
1994.342	18.00	CPUA	000	02	0	00.00.00	00.00.00			1				00.04.05
1994.343	10.00	CPUA	000	02	0	00.00.00	00.00.00			1				00.00.07
1994.343	14.00	CPUA	000	02	0	00.00.00	00.00.00			1				00.00.19
1994.344	10.00	CPUA	000	02	1	00.02.12	00.02.12				1	1	00.21.32	00.02.12
1994.344	11.00	CPUA	000	02	0	00.00.00	00.00.00			3				00.01.02
1994.344	12.00	CPUA	000	02	0	00.00.00	00.00.00			3				00.00.51
1994.344	13.00	CPUA	000	02	1	00.00.16	00.00.16							00.00.16

.
.
.

The ALLOC-REC allocation recovery summary report is a per-system per-LSM report of data collected from the SYSLOGUPDATE function. The report only prints those hours with non-zero data recorded.

Information is also produced for non-Nearline mounts that involved allocation recovery; these manual tape drives are listed with a blank ACS/LSM-id.

This report is produced by the SYSIN control statement PERFORM REPORT NAME(ALLOC-REC). Refer to appendix B, *SYSIN Parameters* for syntax information.

Note: A second report, the MVS Allocation Recovery Update report, is generated automatically during SYSLOGUPDATE processing. That report was previously described in chapter 3, *Managing the ExPR Database.*

Chapter 7: Mainframe Performance Reports   **153**

The ALLOC-REC report contains two sections. The first relates to the specific LSM and the second is for the whole MVS system and contains global information.

The first part of the report provides the following types of information:

- The date, hour, system-id, and ACS/LSM-id.

- Operator reply measurements, including the number of replies that allocated devices on this LSM, the total time spent awaiting a reply, and the longest operator reply.

- System throughput measurements, including the number of outstanding WAITs serviced and the longest WAIT serviced by devices attached to each LSM.

These reply and wait times are recorded in the hour in which the event completed. They can therefore exceed one hour and may have been outstanding for many hours previously.

The second part of each line relates to global system-wide measurements, including:

- The number of CANCEL replies given by the operator.

- The number of WAIT replies given by the operator.

- The highest number of WAITs outstanding at any time (i.e. between the operator replying WAIT and the system allocating a drive and issuing a mount message).

- The time any WAIT or any REPLY was outstanding.

These last two measurements will never exceed one hour. Outstanding events that span several hours will have the time spread across those hours. These fields can be plotted on the ExPR PC component as a measure of operator and system performance.

When an operator replies WAIT or CANCEL to an allocation recovery, there is no associated LSM. Therefore, the operator reply times for WAIT and CANCEL are recorded in the non-Nearline portion of the report.

Whenever an operator reply or a WAITed mount exceeded the user-defined MAXALLOCREC threshold, this is noted in the rightmost column of the report by the *EXCEEDED* message.

DEMAND-ENTERS Report

GLOBAL DETAILS OF DEMAND ENTERS FOR SYSTEM CPUA   DATE RANGE: 1994340 TO 1994342

--DATE--	HOUR	SYSTEM	ACS	LSM	UNIT ADDR	VOLUME SERIAL	JOBNAME
1994.341	13.00	CPUA	000	02	02B0	030471	MAINT13C
1994.341	13.00	CPUA	000	02	02B0	030470	MAINT13C
1994.341	14.00	CPUA	000	01	0C00	104045	PRDACNTS
1994.341	14.00	CPUA	000	01	0C01	102890	PRDACNTS
1994.341	14.00	CPUA	000	01	0C02	102914	PRDACNTS
1994.341	14.00	CPUA	000	01	0C03	102916	PRDACNTS
1994.341	14.00	CPUA	000	01	0C00	102635	PRDACNTS
1994.341	14.00	CPUA	000	01	0C01	107178	PRDACNTS
1994.341	14.00	CPUA	000	01	0C02	102645	PRDACNTS
1994.341	14.00	CPUA	000	01	0C03	000006	PRDACNTS

The DEMAND-ENTERS report lists any demand enters that occurred during the reporting period.

This report is produced by the SYSIN control statement PERFORM REPORT NAME(DEMAND-ENTERS). Refer to appendix B, *SYSIN Parameters* for syntax information.

The DEMAND-ENTERS report provides the following types of information:

- The date and hour in which the demand enter occurred.

- The MVS system, ACS, LSM, and drive involved in the demand enter.

- The volser and jobname that requested the mount.

Demand enters should ideally be rare events and are therefore also listed as exceptions in the SMFUPDATE report. The same mount may also be listed as a mount-time threshold exception due to the manual intervention required.

Chapter 8

File Transfer

Overview

This chapter describes the process of transferring user-extracted data from the ExPR MVS environment to a PC, and the administration that is required for each PC user.

Process Description

User data extraction processing extracts selected records from the mainframe database and formats them to be readable by the ExPR PC component. Records are selected based on date and written to a comma-separated ASCII file that is formatted for file transfer processing to a PC.

Each ExPR user can define their own criteria for selecting records without impacting other users of the system. Control of this process is achieved by using a unique extract dataset for each user or group (DDname EXTRACT) and similarly a unique dataset which contains a single record with a last extract date stamp (DDname LASTEXTR).

The LASTEXTR dataset, referenced through the user's JCL, determines where subsequent database extracts for that user

should begin. Use of the last-extract flag is optional; users can manually override the value by editing the dataset. If this file is not specified, default processing will extract the complete database file.

Note: ExPR maintains multiple last-extract flags. These are unique to particular users and can avoid multiple mainframe databases having to be maintained if a significantly different view or frequency of update is required by different PC users.

- Each user maintains a control and extract file unique to them. These are identified and specified by the user in the JCL. ExPR will use the file associated with the corresponding DD name. These are as follows:

- LASTEXTR: Last extract date specification

RECFM=FB, LRECL=80, BLKSIZE=80, DSORG=PS

- EXTRACT: Extract file containing extracted database records

    RECFM=VB,LRECL=500,BLKSIZE=504 (or greater - normally half-track)

The LASTEXTR dataset is updated automatically by the extract processing when the EXTRACT statement is set as OPTION(FLAG).

The file transfer dataset is sized to match the largest period between file transfer executions. Each update to this file is a replace operation and includes all data with a date higher than the last-update flag.

Some consideration is required to control the specification of user control and extract datasets. In particular, a default ExPR processing option is recommended with additional user

datasets only being used where the default option does not satisfy the requested data profile.

Selecting the Extraction Period

The table below shows how the FROM subparameter of the SYSIN parameter PERIOD is used to control data extraction for individual users. This parameter appears in the control card for each user. (Complete syntax for ExPR SYSIN parameters is described in appendix B, *SYSIN Parameters*.)

PERIOD Subparameter	Description
FROM(*yyyyddd*)	Extracts data from the specified date to the present
FROM(*yyyyddd*) TO(*yyyyddd*)	Extracts data from the specified range of dates
FROM(LASTEXTRACT)	Extracts data from the last extraction date to the present

Note: The SAMPLIB supplied with ExPR suggests using FROM(*yyyyddd*). This is to set up the initial extraction of data. Subsequently, you will probably want to change each user's JCL to LASTEXTRACT.

Transferring Extracted Data to a PC

File transfer operations are initiated from the PC. Files that
are transferred from ExPR MVS to a PC are used as the input
source to a PC database build process. This database is
controlled and queried by the ExPR PC component to build
graphical reports.

Note: The actual mainframe-to-PC file transfer process is
external to ExPR. There are no specification requirements
for this software; you can use the file transfer product you
have in place to transfer extracted data to a PC.

File Transfer Recovery Procedures

Within the file transfer processing cycle, there is a possibility
for the extract file to be deleted or overwritten at the
mainframe before a successful update has been completed on
the PC. This situation can be easily rectified by two methods:
the "user.name.LASTEXTR" dataset can be manually edited
with a different date and time value or an explicit date can be
specified for the extract criteria.

Chapter 9
ExPR API

Overview

This chapter provides information about developing customized mainframe reporting applications using the ExPR Application Programming Interface (API).

Using the API

ExPR mainframe reports have been developed to cover most common user information needs. However, it is understood that these reports may not adequately cover all requirements for all users. The ExPR API has been developed to assist you in developing reporting formats that meet your specifications.

The ExPR API provides consistent access to all database fields for user-written modules that create customized mainframe reports. The advantage of using the API for customized report development is that it provides a consistent development format that will remain valid across releases of the product even if the database format changes.

Specifically, the API provides the user module with an interface containing values which may be useful in creating a

report, and also the addresses of routines that it may call to perform certain functions (such as getting database records or writing report lines).

The user module to be called is defined to ExPR as follows:

```
PERFORM REPORT MODNAME(usermod) ;

or

PERFORM REPORT MODNAME (usermod) GROUP ;
```

The presence or absence of the GROUP parameter influences the API as follows:

If GROUP is specified, the user module is only presented with records for a single LSM as it calls for database records. When it returns control to ExPR, it will have seen only records for a single LSM. ExPR will then invoke the user module again for the next LSM, and so on until all LSMs have been processed.

If GROUP is not specified, the database records are returned in their natural sequence on the database, so any given interval records for each LSM are in turn returned to the caller.

ExPR cannot know which is the appropriate method to use, so it is the user's responsibility to ensure that the GROUP parameter is correctly coded.

You should be aware that not all ExPR database records are keyed by ACS/LSM-id (for example, demand enters, tape errors, and global allocation recovery records). The ExPR API will always pass such records across regardless of the specification of the GROUP parameter or the SELECT statement, such as SELECT LSM(001 02). A system-based SELECT statement will however cause only those records for that system to be passed to the user module, such as SELECT SYSID(CPUA).

You must also be aware that device-type and workload-specific type-0 records will be passed across to your report module, in addition to the global record per LSM/per hour.

All record types are subject to filtering against the PERIOD FROM/TO criteria.

The following record types are not keyed by ACS/LSM-id and will always be passed across:

Type 01 - Control unit and device activity

Type 02 - Global tape volume error counters

Type 03 - Global tape mount counters

Type 05 - Tape dataset group utilization information

Type 06 - Tape dataset group aging information

Type 09 - Channel path group data

Type 11 - Tape error detail event

Type 12 - Demand enter detail event

Type 13 - Allocation recovery global statistics

The ExPR database assembler DSECT, in SAMPLIB member
EXPRDB, gives details of the key values used for each record
type.

ACTION:

⇒ For examples of how to code the API and the conventions
used, refer to SAMPLIB members APIDEMO,
EXPRAPI, and EXPRDB.

⇒ Refer to appendix B, *SYSIN Parameters* for information
about runtime parameters that appear in the SAMPLIB.

Appendix A

CONFIG Parameters

This appendix lists CONFIG DD parameters for ExPR. These parameters are used for relatively static information that will be common to all ExPR jobs, such as hardware and software specifications.

Note: CONFIG parameters can also be specified in the SYSIN and vice versa.

The following conventions are used in the tables below:

- Parameter keywords are shown in **boldface** type

- Parentheses and single quotation marks must be entered as shown

- Variable information you enter is shown in *italics*

- A bar (|) separates mutually exclusive choices

Note: A semicolon (;) delimiter is required at the end of each parameter statement. For example:

ACS ID(001) NAME('*acsnew*')**;**

LSM ID(001 01) NAME('*lsmwolf*') TYPE('3')**;**

CONFIG DD Statements

Hardware Configuration Parameters	Description
HOST SYSID(*mvshost*) ;	HOST defines an MVS host system to ExPR. SYSID identifies the MVS host system name (1-4 characters). **Note:** All HOST statements must appear before any ACS or LSM statements in the ExPR CONFIG file.

ACS
 	
   ID(*aaa*)
 	
   NAME('*text*') ;
	ACS defines an ACS to ExPR.
	
	ID identifies a decimal identifier (000-255) that identifies the ACS. This identifier is part of the site hardware configuration data.
	
	NAME defines an optional user-defined textual identifier (1-60 characters) for the ACS.

LSM ID(*aaa ll*)	LSM defines an LSM to ExPR. ID identifies a decimal identifier that is the associated ACS number (000-255) in the first three positions, a space, and the individual LSM number (0-15) in the last two positions. These numbers in combination identify

	the LSM.
NAME('*text*')	NAME defines an optional user-defined textual identifier (1-60 characters) for the LSM.
TYPE('*t*') ;	TYPE identifies the Nearline system, where 1=4400, 2=Powderhorn, 3=Wolfcreek, 4=Timberwolf, and 5=unspecified. This value is not used by ExPR processing at present; however, you may wish to specify it for documentation purposes.
DEVICE	DEVICE defines Nearline tape drive devices to ExPR.
ID(*dddd*)	ID identifies the device address in the hex format 0000-FFFF.
LSM(*aaa ll*)	LSM identifies the numerical LSM identifier that identifies where the device is located.
TYPE(*device-code*)	TYPE optionally specifies a device-type code (as documented under the DEVICE-TYPE statement).
SYSID(*mvshost*) ;	SYSID identifies the MVS host system associated with the device.
DEVICE-TYPE	DEVICE-TYPE specifies optional device-type codes that cause the ExPR SMFUPDATE function to generate overall statistics for all

	devices per LSM, and additionally statistics for each device-type subgroup.
DEFAULT(*device-code*) ;	The values listed below are valid for the DEVICE-TYPE and DEVICE statements: NONE - no specific device-type assigned (the default) 4480 - the device is an 18-track 4480 (3480-compatible) 4490 - the device is a 36-track 4490 (3490-compatible) 9490 - 36-track 9490 Silverton (3490-compatible) SD-3 - the device is a Redwood Helical Scan device

APPEND
 	
  [ACS ID(*aaa*) NAME('*text*') |
 	
  LSM ID(*aaa ll*) NAME('*text*') TYPE('*t*') |
 	
  DEVICE ID(*ccuu*) LSM(*aaa ll*) TYPE(*device code*) SYSID(*mvshost*)] ;
	APPEND allows you to identify ACS and LSM site names and device types to appear in report headers for automatically generated configurations.
	
	APPEND DEVICE permits addition of a device-type code as documented under the DEVICE-TYPE statement.

CUSTOMER NAME('*text*') ;	CUSTOMER allows you to identify your company or site name to appear in report headers.

CHPATH	CHPATH allows you to logically group channel path IDs associated with the Nearline configuration.
ID(*xx-xx*)	ID identifies a single channel path ID (*xx*) or a range of paths (*xx-xx*) specified in the hex format 00-FF.
GROUP(*nnn*)	GROUP is the common reporting identifier (decimal) for the channel paths specified by the ID parameter. **Note:** Multiple statements with the same GROUP value are supported if you wish to include multiple IDs (such as 10, 20, 30, 35-40) in a single group.
SYSID(*mvshost*) ;	SYSID identifies the MVS host system associated with the channel path.

Software Configuration Parameters	Description
HSC	HSC defines site-specific HSC characteristics to ExPR.
SMFTYPE(*nnn*) ;	SMFTYPE identifies the HSC SMF record type (default 255).

Dataset Usage Parameters	Description
CDS	CDS allows variation of the DD name for the CDS dataset. The default is CDSPRIM.
DDNAME(*dddddd*) ;	DDNAME defines the variation.

Tape Catalog Processing Parameters	Description
DSGROUP	DSGROUP defines a dataset workload group and gives a meaningful name to it.
CLASS(*nnnnn*)	CLASS defines a unique dataset identifier (1-32767) for the dataset workload group.
NAME('*cccccccc*') ;	NAME is a user-defined name to associate with this dataset class identifier.
DSMAP	DSMAP maps a dataset name read from the tape catalog to a dataset class identifier. Multiple DSMAP statements can associate many dataset name masks into a single DSGROUP.
CLASS(*nnnnn*)	CLASS identifies the dataset class (1-32767) that corresponds to the DSGROUP CLASS parameter.
DSN(*dataset_name_mask*)	DSN defines the dataset or dataset mask to include in the

	work group.
ACS(*aaa*)	ACS identifies an ACS to limit the workload group matching.
LSM(*ll*) ;	LSM identifies an LSM to limit the workload group matching.
AGEBANDS(*nnn nnn nnn nnn*) ;	AGEBANDS provides a list of up to four values, representing a number of days, used in tape catalog reporting of cartridge age distribution. These values are sorted into descending order.
TMS	TMS identifies the external module that will read the tape catalog and return selected information.
MODNAME(*extrnlmod*)	MODNAME identifies the external module. You must use the ExPR-supplied modules for TAPECAT processing. The default is XPRCA1X, for CA-1. CA-TLMS users should specify MODNAME(XPRTLMX). If your installation does not have a tape management catalog, specify MODNAME(IEFBR14).
DDNAME(*ddddddd*)	DDNAME identifies the DD file from which the tape catalog will be read. The default is TMC.
PARM(*xxxxxxxxxx*) ;	PARM specifies up to 50 bytes of data that will be passed to the "Clemson-type" interface module

	specified by MODNAME. For ExPR TAPECAT processing, you must specify PARM(EXPRFORM).
TAPEDEF CART3480(*nnnn*) CART3490(*nnnn*) ;	TAPEDEF defines the tape length in feet for all 3480 or 3490 cartridges, or for a range of cartridges. It can also be used to identify ranges of Redwood cartridges by capacity, and to define ranges of cartridges to exclude from tape catalog processing. CART3480 and CART3490 define standard tape lengths in feet for all 3480 and 3490 cartridges. Either or both can be specified in a single TAPEDEF statement. The tape length (*nnnn*) can be any value between 1 and 2000 feet. The default value is 550 feet. **Note:** CART3480 and CART3490 are standalone subparameters; that is, they are not allowed in the same TAPEDEF statement with subparameters listed below that define ranges of cartridges.
--	---
TAPEDEF LOVOL(*volser*) HIVOL(*volser*) LENGTH(*nnnn*) [OPTION([EXCLUDE | RW10 | RW25 | RW50])] ;	LOVOL identifies the lowest-numbered volser to include in the range. HIVOL identifies the highest-

	numbered volser to include in the range.
	LENGTH identifies the tape length in feet for volsers in the defined range, where (*nnnn*) is any value between 1 and 2000 feet. LENGTH is ignored if OPTION is specified.
	OPTION(EXCLUDE) indicates the defined range is to be excluded from tape catalog processing.
	OPTION(RW10) indicates the defined range is composed of Redwood 10Gb cartridges.
	OPTION(RW25) indicates the defined range is composed of Redwood 25Gb cartridges.
	OPTION(RW50) indicates the defined range is composed of Redwood 50Gb cartridges.

Started Task Parameters	Description
PGMI	PGMI defines the HSC PGMI as the started task.
MODNAME(*extrnlmod*)	MODNAME defines an external module to be loaded and then called periodically at intervals indicated by the EVENTS operand. The default is a built-in

	function, so the operand is optional. (For normal processing, this parameter is not required.)
EVENTS(*n*)	EVENTS defines intervals for calling the external module defined in MODNAME. May be set to H, Q, or M, indicating every hour (H), every quarter hour (Q), or every minute (M) respectively. Default is Q.
MAXAGE(*aaa*) ;	MAXAGE defines the oldest data to be retained on the collection file to which the started task writes, as a number of days in the range 0-999.

Job Workload Class Parameters	Description
MAP	MAP maps jobnames to workload groups. Multiple MAP statements can associate many jobname masks into a single workload group.
JOB(*cccccccc*)	JOB defines a mask to be matched against a job name, including wild cards ('*' and '?').
WORKLOAD(*a*) ;	WORKLOAD specifies the associated workload group as a single alphanumeric character (A-Z or 0-9).

WORKLOAD	WORKLOAD complements the MAP statement to provide a descriptive name for the workload group.
ID(*a*)	ID identifies the workload group being named (corresponding to the WORKLOAD subparameter in MAP (A-Z or 0-9).
NAME(*xxxxxxxx*) ;	NAME is a description to identify those jobs that are associated with this workload group.

Exception Threshold Parameters	Description
THRESHOLD	THRESHOLD sets exception threshold values to be used by ExPR exception processing. Threshold values can be defined for the following:
SCRMOUNTTIME(sss)	SCRMOUNTTIME = number of seconds to mount scratch volumes
NONSCRMOUNTTIME(sss)	NONSCRMOUNTTIME = number of seconds to mount non-scratch volumes
SCRMOUNTCOUNT(*nnnn*)	SCRMOUNTCOUNT = number of scratch mounts per LSM hour
NONSCRMOUNTCOUNT(*nnnn*)	NONSCRMOUNTCOUNT = number of non-scratch mounts per LSM hour

Appendix A: CONFIG Parameters   **175**

ENTERS(*nnnn*)	ENTERS = number of enters per LSM hour
EJECTS(*nnnn*)	EJECTS = number of ejects per LSM hour
PASSTHRUS(*nnnn*)	PASSTHRUS = number of passthroughs per LSM hour
CULOAD(*nnn*)	CULOAD = control unit percentage disconnected threshold
AVAILSCRATCH(*nnnn*)	AVAILSCRATCH = minimum number of available scratch cartridges
AVAILCELLS(*nnnn*)	AVAILCELLS = minimum number of available cells
MAXALLOCREC(*ssss*)	MAXALLOCREC = maximum time in seconds for an allocation recovery to be serviced
DRIVESINUSE(*nn*)	DRIVESINUSE = the number of drives in use per LSM at any time

Appendix B

SYSIN Parameters

This appendix lists SYSIN DD parameters for ExPR. These parameters are used for information that is specific to each job, and indicate the function to be performed within that execution of ExPR.

Note: SYSIN parameters can also be specified in the CONFIG and vice versa.

The following conventions are used in the tables below:

- Parameter keywords are shown in **boldface** type

- Parentheses and single quotation marks must be entered as shown

- Variable information you enter is shown in *italics*

- A bar (|) separates mutually exclusive choices

Note: A semicolon (;) delimiter is required at the end of each parameter statement. For example:

> PERIOD FROM(LASTEXTRACT)**;**

> PERFORM EXTRACT**;**

SYSIN DD Statements

Input SMF Source Parameters	Description
INPUT	INPUT allows SMF data input from multiple DD names.
DDNAME(*dd1 dd2 ... ddN*) ;	DDNAME identifies the DD names to accept input from.

Record Selection Parameters	Description
PERIOD	PERIOD selects inclusive date range to process; if not specified, the default is *from* 30 days previous *to* the current date.
FROM(*yyyyddd*) TO(*yyyyddd*) ;	FROM and TO define the range, in the format "yearday" (where January 1, 1995 is 1995001). The default is 30 days previous to the current date.
FROM(LASTEXTRACT) ;	FROM(LASTEXTRACT) indicates that the value is to be obtained from the LASTEXTR dataset.
SELECT	SELECT selects a subset of the data in the date range.
ACS(*aaa*)	ACS indicates all LSMs in the identified ACS are selected for report generation.
LSM(*aaa ll*)	LSM indicates just the identified

	LSM is selected for report generation.
TYPE(HSC SMF21 RMF) ;	TYPE is meaningful only if the input source is SMF data, and allows selective processing of a subset of the SMF input.
SYSID(*mvshost*)	The optional parameter SYSID identifies the MVS host system to be selected for either general report generation or SMF input filtering during SMFREPORT or SMFUPDATE processing. Only one system-id can be selected; generic selection is not supported. If SYSID is not specified, then data for all systems that have been defined by a HOST statement will be selected.

Report Detail Request Parameters	Description
REPORT-OPTIONS (DEVICE-TYPES WORKLOADS) ;	REPORT-OPTIONS controls the level of detail reporting generated by the MOUNTS, MOUNTS-DETAIL, and UTILIZATION reports. DEVICE-TYPES requests that in addition to the general ALL DEVICES summary line, a summary is also printed for each device-type within an LSM. WORKLOADS only applies to

| | the MOUNTS report; it requests a summary line for each user-defined workload group. |

Action Request Parameters	Description
PERFORM	PERFORM specifies an action to be performed in this execution of ExPR and is followed by an action request (e.g., PERFORM INITIALIZE ;).
CONVERTDB SYSID(*mvshost*) ;	CONVERTDB converts ExPR database files from previous software versions to the current software version's format. The SYSID parameter is required.
INITIALIZE ;	INITIALIZE initializes a newly-defined database.
SMFSELECT ;	SMFSELECT selects SMF records from input dataset(s) and writes them to SMFOUT ready for sorting.
SMFUPDATE DUPKEY(*option*) ;	SMFUPDATE selects SMF records from a sorted input dataset, creates database records, and produces an Exception Event report. DUPKEY *option* is IGNORE, REPLACE, or CANCEL. The default is REPLACE.
PGMIUPDATE ;	PGMIUPDATE selects records from the PGMIDATA dataset

	written by the ExPR started task and creates database records.
EXTRACT OPTION(FLAG|NOFLAG);	EXTRACT extracts data from the ExPR database for transfer to PC.
	OPTION indicates whether the last extract value flag is to be updated, either FLAG or NOFLAG. The default is NOFLAG.
PGMI ;	PGMI performs started task processing (periodic data collection from CDS and HSC PGMI).
EXPORT ;	EXPORT selects data from the ExPR database and writes to a QSAM dataset for archiving.
PURGE ;	PURGE removes the selected range of records from the ExPR database.
SYSLOGUPDATE [JES3(*x*=*yyyy* ...)] [RECFM(*valid recfm*)] ;	SYSLOGUPDATE requests allocation/recovery analysis of the JES console log.
	JES3 defines those systems in a JES3 complex that are to be extracted. For each system, *x* represents the message origin prefix and *yyyy* represents the associated MVS system-id (normally the SMF-id) . Do not specify the JES3 parameter if you are processing JES2 console logs.

Appendix B: SYSIN Parameters   **181**

	RECFM - this optional parameter is valid for JES2 or JES3 installations. The IBM external writer produces a VBA file of JES2 SYSLOG messages and VB for JES3 messages. However, many sites use other utilities to manage their SYSLOG, and these utilities may not produce their files in VBA format. The valid RECFMs are VBA, VBM, FBA, FBM, FB, and VB. The default is VBA/VBM.

Report Request Action Parameters	Description
PERFORM	PERFORM specifies an action to be performed in this execution of ExPR and is followed by an action request (e.g., PERFORM SMFREPORT ;).
SMFREPORT ;	SMFREPORT selects SMF records from a sorted input dataset and produces an Exception Event report, the same as SMFUPDATE, but does not create database records.
REPORT NAME(*predefrep*/*extrnlmod*) {GROUP|NOGROUP} ;	REPORT generates ExPR mainframe reports. Reporting is performed through an API to allow local extension of the

For example:	reporting facilities.
PERFORM REPORT NAME(MOUNTS) ;	NAME identifies the report source, either predefined or from an external module.
	predefreport represents one of the following values:
	MOUNTS - produce the mounts report
	CONTENTS - produce the contents report
	UTILIZATION - produce the utilization report
	MOUNTS-DETAIL - produce the mounts detail report
	ALLOC-REC - produce the allocation recovery report
	TAPE-ERRORS - produce the tape errors reports
	THRESHOLDS - produce the threshold exceptions report
	CU-BUSY - produce the control unit busy report
	PATH-BUSY - produce the channel group utilization report
	DEMAND-ENTERS - produce the demand errors report

	extrnlmod represents a load module to which ExPR links, and that module obtains data and writes the report by calling back through the API.
	GROUP indicates that the report is 'per LSM' and only records for a single LSM should be passed during each link. It is optional and is not needed for the predefined reports.
	NOGROUP is the default for user-written reports using the API. If specified, all records for an ACS will be passed in a single invocation of the report module, the opposite of GROUP.

Tape Catalog Action Request Parameters	Description
PERFORM TAPECAT OPTION(*subparameter*)	PERFORM requests an action to be performed.
	TAPECAT indicates the tape catalog function.
	OPTION(*subparameter*) identifies the TAPECAT function to perform, either UPDATE, REPORT, SUMMARY, or HISTORY. Only one of these four options may be specified in each TAPECAT control

	statement. OPTION must be immediately followed by one of the four keywords, which is set in parentheses.
TAPECAT OPTION(*subparameter*) strings are listed below:	
--	--------------------------------------
(UPDATE [NOHIST FULLCAT NODBUP]) SYSID(*mvshost*) For example: PERFORM TAPECAT OPTION(UPDATE NOHIST) SYSID(MVS1) ;	UPDATE reads the HSC CDS and tape catalog. All volumes in the tape catalog that are also in the HSC CDS are processed. The database file is updated and the history file is created. Summary aging and utilization reports are produced. NOHIST - This UPDATE option performs tape catalog processing but does not create the history file. NODBUP - This UPDATE option performs tape catalog processing but does not update the database. FULLCAT - This UPDATE option processes all volumes in the tape catalog in addition to volumes found in the HSC CDS (i.e., non-ACS volumes are included).
SYSID(*mvshost*)	SYSID identifies the MVS host system associated with the tape catalog. This parameter is required for UPDATE processing. In a multi-MVS/shared system, the main system-id should be

Appendix B: SYSIN Parameters   185

	specified.
--	---------------------------------------
(REPORT [ALLDSNS FULLCAT] BY*xxxx*...) For example: PERFORM TAPECAT OPTION(REPORT BYDSN) ;	REPORT produces the Nearline Volume Details report from the history file created by UPDATE processing. The required BY*xxxx* parameter specifies the sorting sequence. A summary line is printed for each primary dataset that was written from the HSC CDS and tape catalog, sorted into the appropriate sequence. ALLDSNS - This REPORT option prints a summary line for all primary and secondary datasets. The default is primary only. FULLCAT - This REPORT option prints a summary line for all non-ACS volumes (but only if FULLCAT was also specified as an UPDATE option at the time the history file was created).
Note: Multiple BY*xxxx* options can be specified on a single statement to generate multiple reports.	BY*xxxx* - This required REPORT parameter specifies the sorting sequence(s) for the report: BYVOL - by volser, file sequence BYDSN - by dataset name, volume sequence

	BYREF - by reference date, dataset name
	BYUTIL - by percentage/Mb, dataset name
	BYUSES - by use count, reference date
	BYMEGS - by Mb, dataset name
	BYACS - by ACS/LSM, volser/file sequence
	BYMEDIA - by media type, volser/file sequence

TAPECAT SUMMARY AND HISTORY REPORTS	Description
PERFORM TAPECAT OPTION(SUMMARY) ;	SUMMARY produces the Nearline Volume Summary report from the history file created by TAPECAT UPDATE processing. Volume statistics and utilization estimates are summarized for each ACS, LSM, dataset workload group, and media type (device type and cart length). Statistics are also summarized separately for non-ACS volumes if FULLCAT was specified during the UPDATE processing that created the history file.

--	---
PERFORM TAPECAT OPTION(HISTORY) ;	HISTORY produces comparison reports for two history files, listing volumes that have changed status and a summary of activity within the Nearline library and tape catalog.

Appendix C

Messages

Overview

This appendix lists error, warning, and informational messages that may be generated by ExPR.

Message Format

ExPR messages are identified by an eight-character message identifier in the format XPRnnnn%, where:

XPR   Identifies the message as an ExPR message

nnnn   The message number

%   Indicates if the message is for an error (E), a warning (W), or is for informational (I) purposes

ExPR Messages

Messages are listed below in numerical sequence. Where appropriate, information is included to help you determine why the message was produced and steps you should take to correct the problem.

Appendix C: Messages   **189**

Message ID	Description
XPR0001I	Error opening SYSIN This informational message indicates SYSIN is either missing or not pointing to a valid dataset (QSAM, 80 bytes fixed). Check the console log for more information.
XPR0002I	Error opening CONFIG This informational message indicates CONFIG is either missing or not pointing to a valid dataset (QSAM, 80 bytes fixed). Check the console log for more information.
XPR0003E	Unclosed quotes This error message indicates a syntax error in CONFIG or SYSIN parameters. Quotation marks must be entered as shown in the parameter tables. Check appendix A and B for correct syntax.
XPR0004E	Unmatched parenthesis This error message indicates a syntax error in CONFIG or SYSIN parameters. Parentheses must be entered as shown in the parameter tables. Check appendix A and B for correct syntax.
XPR0005E	No Delimiter (*/) on Final Statement This error message indicates a syntax error in CONFIG or SYSIN parameters. A */ delimiter is required after each block of comments (/*...comments...*/).
XPR0006E	No delimiter (semicolon) on final statement This error message indicates a syntax error in

	CONFIG or SYSIN parameters. A semicolon delimiter is required after each keyword statement. Check appendix A and B for correct syntax.
XPR0007I	Parameter File Processing Completed Successfully. This informational message indicates that ExPR has read and validated all CONFIG and SYSIN control statements, and no errors were found. Requested functions are then performed.
XPR0008E	Parameter file processing ended, errors detected This error message indicates a syntax error in CONFIG or SYSIN parameters. This message is issued with another message that indicates a probable cause. Check appendix A and B for correct syntax.
XPR0009I	Last Statement Processed was __ / Parameter __ - Statement no __ This informational message is a supplementary diagnostic message to assist in determining which ExPR AUTO-CONFIG, CONFIG, or SYSIN statement is in error. The statement number value is the sequence number of the erroneous statement within the complete AUTO-CONFIG, CONFIG, SYSIN sequence (excluding comment statements).
XPR0010E	Parameter __ unrecognized This error message indicates a syntax error in CONFIG or SYSIN parameters, possibly a missing semicolon delimiter at the end of the previous statement. Check appendix A and B for correct syntax.
XPR0011E	Expecting keyword before __ This error message indicates a syntax error in CONFIG or SYSIN parameters, possibly an open parenthesis with no preceding keyword caused by a misplaced semicolon delimiter, or by a missing

	keyword subparameter. Check appendix A and B for correct syntax.
XPR0012E	Too many parameters specified The INPUT DD(....,....,....) command has too many DDnames specified for SMF input data. Check appendix B for correct syntax.
XPR0013E	__ does not accept value operands This error message indicates a syntax error in CONFIG or SYSIN parameters, possibly a missing keyword or a misplaced parenthetical entry. Check appendix A and B for correct syntax.
XPR0014E	Expected numeric value, found __ This error message indicates a syntax error in CONFIG or SYSIN parameters, where a numeric value was expected but a non-numeric value was found. Check appendix A and B for correct syntax.
XPR0015E	Cannot process numeric value, __ too large This error message indicates a syntax error in CONFIG or SYSIN parameters, where more than the allowable number of digits were entered. Check appendix A and B for correct syntax.
XPR0016E	ACS __ invalid, outside range 0-255 This error message indicates a syntax error in CONFIG or SYSIN parameters, where an invalid ACS id was entered. Check appendix A and B for correct syntax.
XPR0017E	Expected date as yyyyddd, found __ This error message indicates a syntax error in SYSIN parameters, where a date was entered that was not in the expected format. Check appendix B for correct syntax.
XPR0018E	Value __ is beyond the valid range / Invalid - the valid

	syntax is __
	This error message indicates a value error in a CONFIG or SYSIN statement. This is normally a numeric value which has been specified as too low or too high. The acceptable syntax of the statement/parameter is also listed. Check appendix A or B for further clarification.
XPR0020E	Device addr __ invalid, outside range 0000-FFFF This error message indicates a syntax error in CONFIG parameters, where an invalid device address was entered. Check appendix A for correct syntax.
XPR0021E	Channel path __ invalid, outside range 00-FF This error message indicates a syntax error in CONFIG parameters, where an invalid channel path was entered. Check appendix A for correct syntax.
XPR0022E	Incomplete parameter set for this statement This error message indicates a syntax error in CONFIG or SYSIN parameters, where an incomplete set of the required or minimum parameters was specified. Check appendix A and B for correct syntax.
XPR0023E	Open error on SMFOUT This error message indicates SMFOUT is either missing or not pointing at a valid dataset. Check the console log for more information.
XPR0024E	Error on SMFOUT, record __, length __ This error message indicates a write error occurred while processing SMFOUT. Check the space allocation for the SMFOUT dataset.
XPR0025E	Dataset mask format error: __ This error message identifies an invalid dataset workload group mask specification. Refer to chapter 6

	for valid mask formats.
	The valid error codes are:
	4 - illegal character in DSN
	8 - a DSN node was >8 or zero characters
	12 - DSN ends in a period/full-stop or has trailing non-blanks
	16 - Multiple wildcards (**) specified
XPR0026E	Open error on __
	This error message indicates the named DCB is invalid, which usually means the DDNAME is spelled incorrectly or allocated to the wrong dataset. Check the console log for more information.
XPR0027E	Max age __ invalid, outside range 1-999
	This error message indicates the PGMI MAXAGE value is outside the absolute range. Enter a valid number.
XPR0028E	No report name specified
	This error message indicates a PERFORM REPORT parameter is missing either the NAME subparameter or a report name to generate. Check appendix B for correct syntax.
XPR0029E	Workload mask format error in __: __
	This error message identifies an invalid workload mask. Refer to chapter 4 for valid mask formats.
XPR0030E	Internal Tables Full - Too Many Statements or Parameters
	This error message indicates that, during user statement processing, the total of all AUTO-CONFIG, CONFIG, and SYSIN statements (excluding comment statements) has caused the internal tables to fill up. Review the quantity of input statements and remove

	any defunct system definitions or devices that no longer exist. It may be necessary to copy the AUTO-CONFIG file (CDSCONF) into your CONFIG file and edit it to reduce the size. You can then deactivate AUTO-CONFIG by removing the CDSCONF DD statement.
XPR0031E	Database initialization SHOWCB failed This error message indicates an internal VSAM file handling error occurred during database initialization. Contact Customer Services.
XPR0032E	Database initialization failed: missing DD statement This error message indicates a VSAM file open error occurred during database initialization. Check for the correct DDNAME.
XPR0033E	Database initialization failed: open error This error message indicates database initialization failed because the database initialization module could not open the file. Check the console log for more information.
XPR0034E	Database initialization failed: not a KSDS This error message indicates the database initialization failed because the DDNAME does not point to a valid dataset. Check the console log for more information.
XPR0035E	Database initialization failed: close error This error message indicates database initialization failed because the database initialization module could not close the database correctly. Check the console log for more information.
XPR0036E	Database initialization failed: Open Error This error message indicates that the second open of the database has failed. Check the console log for more information.

| XPR0037E | Database initialization MODCB failed

This error message indicates an internal error occurred during database initialization. Contact Customer Services. |
|---|---|
| XPR0038E | Database initialization failed: database not empty

This error message indicates the database initialization failed because the database has already been initialized. |
| XPR0039E | Database initialization failed: couldn't write control record

This error message indicates an internal error occurred during database initialization. Check that the VSAM definition is consistent with ExPR requirements (i.e., not below minimum record length, etc.). |
| XPR0040I | Database Initialized Successfully.

This information message indicates the requested PERFORM INITIALIZE function has completed successfully. |
| XPR0041E | Database initialization ERASE failed

This error message indicates an internal error occurred during database initialization. Contact Customer Services. |
| XPR0042E | Database initialization GET failed

This error message indicates an internal error occurred during database initialization. Check that the VSAM definition is consistent with ExPR requirements (i.e., not below minimum record length, etc.). |
| XPR0043E | GETMAIN Error - Insufficient Storage - Ensure REGION=4M or Greater - Abend S0C1 Follows

This error message indicates the ExPR storage management module has received a non-zero return |

	code from a GETMAIN request. This is probably due to insufficient region size. Increase the region size and retry the job. An abend S0C1 is forced immediately after this message for the purpose of debugging. R7 contains the required storage length and R8 contains the R15 return code from the GETMAIN macro. A value of hex 'ABC' in R8 indicates an internal ExPR GETMAIN request error.
XPR0044E	FREEMAIN Error Code nn - Abend S0C1 Follows This error message indicates the ExPR storage management module has detected an error while processing a FREEMAIN request. The possible error codes are: 01 - the leading storage eye-catcher has been corrupted - part 1. 02 - the leading storage eye-catcher has been corrupted - part 2. 03 - the storage area trailer area has been overwritten. 04 - the FREEMAIN macro gave a non-zero return code. An Abend S0C1 is forced immediately after this message for the purpose of debugging. R6 points at the area to be freed, R7 contains the length of the area, and R8 contains the R15 return code from the FREEMAIN macro.
XPR0050E	Input SMF records out of sequence __ follows __ This error message indicates that records were not updated in ascending sequence during database update. Be sure to include the select and sort steps in your JCL for updating the database with SMF records, as described in chapter 2.
XPR0051E	Error on EXTRACT, record __ length __ invalid

	This error message indicates an internal error occurred while writing extract records. Contact Customer Services.
XPR0053E	Error opening EXTRACT

This error message indicates that the EXTRACT dataset is missing, which usually means the DDNAME or dataset type is incorrect. Check the console log for more information. |
| XPR0060E | Couldn't load external module __

This error message indicates a syntax error in SYSIN parameters, where the named predefined report or external module is invalid or spelled incorrectly. Refer to appendix B for correct syntax. |
| XPR0071E | Error writing to __ R15=__

This error message indicates an unrecoverable error was detected while trying to write to the database. Check the console log for more information. |
| XPR0072E | Duplicate key encountered, DUPKEY(CANCEL) specified

This error message indicates that, during database update, when CANCEL was specified, a duplicate key was encountered. This could be caused by an attempt to update the wrong database or if the update was being re-run. |
| XPR0073E | Duplicate key encountered, get-for-update failed

This error message indicates an internal VSAM error occurred. Call Customer Services. |
| XPR0074E | Duplicate key encountered, replace failed

This error message indicates an internal VSAM error occurred. Call Customer Services. |
| XPR0075E | Input records out of sequence ___ follows ___ |

	During the PGMIUPDATE function, ExPR found that the records in the PGMIDATA file were not in ascending date/time sequence. As PGMIDATA is a KSDS file, this should not occur. Run IDCAMS VERIFY against PGMIDATA and retry the PGMIUPDATE function.
XPR0076E	CDS relative record __ id was not __

This error message indicates an internal error caused by an unexpected response from the CDS or possibly by a record format version incompatibility. Call Customer Services. |
| XPR0077E | Process aborted due to foregoing error

This error message indicates another error caused the process to stop. A separate error message will be generated for that error. |
| XPR0078E | Process aborted due to error in sort

This error message indicates that an internal sort process failed. Check the console log for more information. |
| XPR0079W | No Tape Catalog data exists

This warning message indicates that an attempt was made to read tape catalog data but no data was found. This is possibly caused by an empty tape catalog. |
| XPR0080W | Volser __ not in Tape Catalog

This warning message indicates that the CDS knows of the named cartridge but the tape catalog does not. This message indicates a possible HSC synchronization problem. |
| XPR0081W | Volser __: CDS says Scratch, Catalog says not

This warning message indicates that the CDS knows the named cartridge as a scratch volume but the tape |

	catalog does not. This message indicates a possible HSC synchronization problem.
XPR0082W	Volser __: Catalog says Scratch, CDS says not

This warning message indicates that the CDS knows the named cartridge as a non-scratch volume but the tape catalog does not. This message indicates a possible HSC synchronization problem. |
| XPR0083W | Volser __ in transit or location unknown

This warning message indicates that the CDS has a record indicating the named cartridge exists but does not know its location. This message indicates a possible HSC synchronization problem. |
| XPR0084W | Volser __ location invalid - LSM __, __

This warning message indicates that the CDS has a record indicating the named cartridge's location but the LSM does not exist. This message indicates a possible HSC synchronization problem. |
| XPR0085E | Cleaning Cartridge Prefix record absent from CDS

This error message indicates that the CDS does not contain a record which indicates a cleaning cartridge prefix, possibly caused by an invalid CDS. |
| XPR0086E | Volser __: Problem with DSNAME - Reason __

This error message indicates an internal error caused by a dataset name on the named cartridge that does not match dataset group masks. Refer to chapter 6 for valid mask formats.

Valid reason codes are:

8 - invalid DSN node length >8 or zero
12 - position 44 of DSN was a period/full-stop.

This will occur against CA-TLMS "crash-protected" |

	tapes and other special status volumes within CA-1 and CA-TLMS.
XPR0087W	Volume __ DSN=___ Selected in dataset group __ and ___ This warning message indicates that dataset group and map criteria have selected the named volume in more than one dataset workload group. You should be aware of this when comparing or combining the results from dataset workload group output.
XPR0091I	This message is a heading for the utilization summary produced by the TAPECAT UPDATE function. Each line printed represents an LSM silo or a dataset group.
XPR0092I	This message is a heading for the aging summary produced by the TAPECAT UPDATE function. Each line printed represents an LSM silo or a dataset group.
XPR0093W	At Least One LSM or dataset group has No Volumes During the preceding XPR0091/XPR0092 summary, an LSM silo or dataset group had no volumes matched/selected against it. For a dataset group, this could merely be an error in the dataset name masks. If an LSM has no volumes/cartridges, this would require further investigation.
XPR0094I	_____ CDS Volumes Not in Tape Catalog This is a total of message XPR0080 occurrences.
XPR0095I	_____ Volumes Have a Conflicting Scratch Status This is a total of the occurrences of messages XPR0081 and XPR0082.
XPR0096E	'PARM (EXPRFORM)' required for ExPR Tape Catalog Interface Module For ExPR TAPECAT processing, you must specify TMS PARM(EXPRFORM).
XPR0097W	___ Volumes have had their Megabytes Capacity Capped.

	This warning message is a count of volumes where the tape catalog values exceed 80% of the theoretical maximum. These volumes have a plus symbol printed against the megabytes column of the Volume Details report. During ExPR TAPECAT UPDATE processing, the tape catalog information is used to calculate how much data is stored on a cartridge. This is compared against a theoretical maximum capacity based on recording technique and cartridge length.
XPR0098E	Invalid Record Type Returned by External TMC Interface Module This error message indicates the external "Clemson" module, as defined by TMS MODNAME(....) PARM(EXPRFORM), has failed to pass back an extended format record for TAPECAT UPDATE processing. Currently only ExPR-supplied modules support the extended format and parameter PARM(EXPRFORM). If this message occurs with an ExPR-supplied module, call Customer Services.
XPR0100W	DSN Count Mismatch: Volume _____ says _____ DSNS - _____ DSNS Found The XPRCA1X CA-1 catalog interface module has detected a discrepancy in the catalog. The DSNB record count in a volumes base record does not match the number of DSNBs found for that volume. This may be due to in-flight activity while the TAPECAT function was running.
XPR0101E	ERROR-VOL _____, DSNB _____ Requested, DSNB _____ RETRIEVED During CA-1 catalog processing, the XPRCA1X module has attempted to retrieve a DSNB associated with the stated volume. However, a different DSNB was passed back by the DA logic module. This could be an error in XPRCA1X or a corruption of the

	catalog pointers.
XPR0103E	Tape Catalog Does Not Have Exactly Three Control Records

The first three records of the CA-1 catalog are control records with the identifiers TMSCTL#1, 2 & 3. XPRCA1X has failed to find these and cannot process the catalog. The catalog may be corrupted. |
| XPR0104E | Called Sort of Tape Catalog Failed

XPRCA1X invokes the SORT utility to sort tape catalog DSNBs and optionally the volume base records. One of these sorts has filed. Review the SYSPRINT/SYSOUT to determine the reason for the failure. Each sort has a unique identifier - DSNS or VOLS. |
| XPR0105I | Catalog Volumes Out of Sequence - Sort will be Called

While retrieving the CA-1 catalog volume base records, XPRCA1X has detected that they are not in ascending order - this is not an error. This will occur if your installation has non-numeric volume serial numbers and uses the CA-1 exits TMSUXOE and TMSUXOU. |
| XPR0106W | Unowned/Orphaned DSNB Found: VOL=_____ DSN = dataset name _____

The XPRCA1X CA-1 interface has found a DSNB record for which there is no matching volume base record. This may be due to in-flight activity while the TAPECAT function was running, alternatively there may be a corruption of the tape catalog. |
| XPR0107I | XPRCA1X Run Stats - ____ Vols, ____DSNBS, ____ DA READS

XPRCA1X has completed processing and is returning control to the main ExPR module. The statistics |

	represent the number of base volume records and secondary dataset DSNBs read from the catalog. The DA READ count is associated with the processing of the multi-file/multi-volume tape stacks.
XPR0108W	Volume ___ has ___ Datasets - Truncated at ___

This warning message indicates that ExPR has not been able to collect all secondary dataset information on the named volume because the volume has more secondary datasets than can be accommodated in a 32756-byte record. MVS restricts the maximum VB record length to 32756 bytes, so ExPR truncates the record at 572 datasets. The dataset count for the volume will include the additional secondary datasets; however, they will not be included in the calculations of percentage used and megabytes. The extra datasets will also not be written to the TMCHIST file for subsequent reporting by the TAPECAT OPTION(REPORT ALLDSNS) function. |
| XPR0120I | TLMS Control: ___ TLMS report options ___

The CA-TLMS interface XPRTLMX will print several of these informational messages to show the control cards used for the TLMS003 and TLMS018 reports - this is for audit trail purposes. |
| XPR0121E | Input Report is not Recognizable TLMS Format - Reason =

XPRTLMX has detected an error in the processing/extraction from the TLMS report file. The reason can be:

(01) - unable to find the TLMS001 control page.

(02) - unable to find the CA-DYNAM/TLMS eye-catcher and level/service indicators.

(03) - the reports date format is not MM/DD/YY or |

	DD/MM/YY. (04) - while processing the TLMS003 report a discrepancy was discovered - this is probably insufficient columns of data. (05) - while processing the TLMS018 report a discrepancy was discovered - this is probably insufficient columns of data. Check that the TLMS018 control card was correctly specified. (06) - no records were extracted from the TLMS003 and TLMS018 reports.
XPR0124E	Called Sort of TLMS Catalog Volumes Failed The called SORT from XPRTLMX gave a non-zero return-code. Review the SYSPRINT/SYSOUT to determine the reason for the failure.
XPR0127I	XPRTLMX Run Stats - ___ Scratches, ____ Active Volumes, _____ Datasets, ____ Crash-Protected Volumes XPRTLMX has completed processing and is returning control to the main EXPR module. These control counters are self-explanatory.
XPR0700I	variable text This message number is used to prefix messages passed to EXPR by external tape catalog interfaces, such as Clemson.
XPR0701E	I/O Error: \<SYNAD DATA\> This error message indicates an I/O error while reading or writing a file. Check that you have enough space allocated for the file and that the DDname is correct.. If the problem is not due to the above, call Customer Services.
XPR702E	CDS Relative Record _____ ID was not

	_____ This error message indicates an inconsistency was found while reading the CDS for record type, possibly caused by an invalid CDS.
XPR0703E	Error Loading SLSXCAL This error message indicates an interface problem while attempting to load the HSC interface. Check that HSC is active.
XPR0704E	Return Code R15=_____ From Programmatic Interface This error message indicates an invalid return code from HSC. Check that HSC is active.
XPR0705E	Return Code SLXCMDRC=___ From Programmatic Interface This error message indicates an invalid return code from HSC. Check that HSC is active
XPR0706E	Inconsistent Data Returned From Programmatic Interface This error message indicates that HSC may not be properly installed.
XPR0707E	Programmatic Interface Response Conflicts With Control Dataset This error message indicates that the CDS and programmatic interface do not agree about the number of LSMs in the ACS. Check that the CDS is valid and that has not been updated since the PGMI was started.
XPR0708E	Unable to Open DDNAME _____ This error message indicates an error occurred while opening a DCB dataset, possibly caused by a missing DDname. Check the console log for more

	information.
XPR0709I	HSC /SLSXCAL Level is v.r.m. (assumed as default)
	The ExPR PGMI interface module has determined the level of HSC installed by scanning the SLSXCAL module. If the additional text ASSUMED AS DEFAULT appears, the level of HSC could not be determined and 1.2 is assumed.
XPR0710E	XPRQAC1 Incorrectly Invoked
	This error message indicates a logic error in the EXPR PGMI interface routines. No configuration or ACS/LSM information has been extracted from the CDS or PGMI.

Appendix D

Database and TMCHIST Layout

The ExPR mainframe database is a VSAM KSDS file. It has a 16-byte key, which should be added to the length of each record.

EXPR DATABASE KEY SECTION

FIELD	LENGTH	DESCRIPTION
DBDATE	4	Date - YYYYDDD
DBHOUR	1	HOUR - x'00' - x'17'
DBSYS	4	MVS SYSTEM-ID
DBACS	1	ACS Number
DBLSM	1	LSM Number
DBDRVTYP	1	Device-Type
--------	1	reserved
DBRECTYP	1	Record Type
DBWKLOAD	1	Workload ID
DBMINUTE	1	Minutes of Hour.
TOTAL	16	

RECORD TYPE 0: HSC SMF HOURLY MOUNT STATISTICS

FIELD	LENGTH	DESCRIPTION
DBMCSC	4	Mount count scratch
DBMCNS	4	Mount count non scratch
DBMCNJ	4	Mount count non-job mounts
DBMTSC	4	Mount time scratch
DBMTNS	4	Mount time non scratch
DBMTSCM	4	Maximum mount time scratch
DBMTNSM	4	Maximum mount time non scratch
DBDUSC	4	Drive utilization scratch
DBDUNS	4	Drive utilization non scratch
DBDENTER	4	Demand enters count
DBENTER	4	Enters count
DBEJECT	4	Ejects count
DBPTHRU	4	Passthroughs count
DBROBOT	4	Robotics utilization
DBDWT	4	Drive wait time - scratch
DBDRO	4	Destination robotics time - scratch
DBDRQ	4	Destination robotics queue time - scratch
DBORO	4	Other robotics time - scratch
DBORQ	4	Other robotics queue time - scratch

DBPRO	4	Passthrough robotics time - scratch
DBPRQ	4	Passthrough robotics queue time - scratch
DBDWTNS	4	Drive wait time - non-scratch
DBDRONS	4	Destination robotics time - non-scratch
DBDRQNS	4	Destination robotics queue time - non-scratch
DBORONS	4	Other robotics time - non-scratch
DBORQNS	4	Other robotics queue time - non-scratch
DBPRONS	4	Passthrough robotics time - non-scratch
DBPRQNS	4	Passthrough robotics queue time - non-scratch
DBCMPX	17	Concurrent drive array (0-16 in-use)
TOTAL	**129**	(excluding key area length)

RECORD TYPE 1: HOURLY CONTROL UNIT ACTIVITY

FIELD	LENGTH	DESCRIPTION
DB1SSCH	4	Total SSCHs thru CU
DB1CUB	4	CU busy time
DB1LCU	4	Logical CU number
DB1TABL	0	Table of 17x2 two-byte entries
DB1CUAV	1	CU - Average % disconnected
DB1CUH1	1	CU - Highest % disconnected
DB1DVAV	1	Device xx0 - Average % disconnected

DB1DVH1	1	Device xx0 - Highest % disconnected
	1	Device xx1 - Average % disconnected
	1	Device xx1 - Highest % disconnected
	28	Devices xx2-xxF Avg and High % disconnected
DB1ADRT	4	Base address of CU x 'ccuu' (in character)
TOTAL	**50**	(excluding key area length)

RECORD TYPE 2: HOURLY TAPE ERROR STATISTICS

FIELD	LENGTH	DESCRIPTION
DB2PREA	4	Number of permanent read errors - ACS
DB2PWEA	4	Number of permanent write errors - ACS
DB2TREA	4	Number of temporary read errors - ACS
DB2TWEA	4	Number of temporary write errors - ACS
DB2PNOA	4	Number of volumes with perm errors - ACS
DB2TNOA	4	Number of volumes with temp errors - ACS
DB2PREN	4	Number of permanent read errors - non-ACS
DB2PWEN	4	Number of permanent write errors - non-ACS
DB2TREN	4	Number of temporary read errors - non-ACS
DB2TWEN	4	Number of temporary write errors - non-ACS
DB2PNON	4	Number of volumes with perm errors - non-ACS

DB2PTNON	4	Number of volumes with temp errors - non-ACS
TOTAL	48	(excluding key area length)

RECORD TYPE 3: HOURLY GLOBAL MOUNTS

FIELD	LENGTH	DESCRIPTION
DBGMNTS	4	Global scratch mounts
DBGMNTN	4	Global non-scratch mounts
TOTAL	**8**	(excluding key area length)

RECORD TYPE 4: LSM VOLUME AGING STATISTICS

FIELD	LENGTH	DESCRIPTION
DB4NCEL	4	Total number of cells
DB4NFRE	4	Number of free cells
DB4NCLN	4	Number of cleaning cartridges
DB4NSCR	4	Number of scratch tapes
DB4OLDS	4	Longest days since last mount
DB4MVOL	4	Number of volumes multi volume
DB4MDS	4	Number of volumes multi dataset
DB4BNDS	0	Array of age band data x 5
DB4NVOL	4	Band x - number of volumes
DB4AGET	4	Band x - sum of ages of volumes

Appendix D: Database Layout   **213**

-----	4	Band A - number of volumes
-----	4	Band A - sum of ages of volumes
-----	4	Band B - number of volumes
-----	4	Band B - sum of ages of volumes
-----	4	Band C - number of volumes
-----	4	Band C - sum of ages of volumes
-----	4	Band D - number of volumes
-----	4	Band D - sum of ages of volumes
TOTAL	**68**	(excluding key area length)

RECORD TYPE 5: DATASET GROUP UTILIZATION STATISTICS

FIELD	LENGTH	DESCRIPTION
DB5TTAPE	4	Total # of tapes
DB5AVGPC	4	Average % utilization
DB5T010	4	Total tapes in range 00 - 10%
DB5T020	4	Total tapes in range 11 - 20%
DB5T030	4	Total tapes in range 21 - 30%
DB5T040	4	Total tapes in range 31 - 40%
DB5T050	4	Total tapes in range 41 - 50%
DB5T060	4	Total tapes in range 51 - 60%
DB5T070	4	Total tapes in range 61 - 70%
DB5T080	4	Total tapes in range 71 - 80%

FIELD	LENGTH	DESCRIPTION
DB5T090	4	Total tapes in range 81 - 90%
DB5T100	4	Total tapes in range 91 - 100%
----	4	Spare/unused
TOTAL	**52**	(excluding key area length)

RECORD TYPE 6: DATASET GROUP AGING STATISTICS

FIELD	LENGTH	DESCRIPTION
DB6OLDS	4	Longest days since last mount
DB6MVOL	4	Number of multi-volume volumes
DB6MDS	4	Number of multi-dataset volumes
DB6BNDS	0	Array of age band data x 5
DB6NVOL	4	Band x - Number of volumes
DB6AGET	4	Band x - Sum of ages of volumes
-----	4	Band A - number of volumes
-----	4	Band A - sum of ages of volumes
-----	4	Band B - number of volumes
-----	4	Band B - sum of ages of volumes
-----	4	Band C - number of volumes
-----	4	Band C - sum of ages of volumes
-----	4	Band D - number of volumes
-----	4	Band D - sum of ages of volumes
TOTAL	**52**	(excluding key area length)

RECORD TYPE 7: LSM VOLUME UTILIZATION STATISTICS

FIELD	LENGTH	DESCRIPTION
DB7TTAPE	4	Total # of tapes
DB7AVGPC	4	Average % utilization
DB7T010	4	Total tapes in range 00 - 10%
DB7T020	4	Total tapes in range 11 - 20%
DB7T030	4	Total tapes in range 21 - 30%
DB7T040	4	Total tapes in range 31 - 40%
DB7T050	4	Total tapes in range 41 - 50%
DB7T060	4	Total tapes in range 51 - 60%
DB7T070	4	Total tapes in range 61 - 70%
DB7T080	4	Total tapes in range 71 - 80%
DB7T090	4	Total tapes in range 81 - 90%
DB7T100	4	Total tapes in range 91 - 100%
----	20	Spare/Unused
TOTAL	**68**	(excluding key area length)

RECORD TYPE 8: HOURLY AVERAGE SCRATCH & FREE CELL STATISTICS

FIELD	LENGTH	DESCRIPTION
DB8NCEL	4	Number of cells
DB8NFRELO	4	Minimum free cells

DB8NFRE	4	Average free cells
DB8NFREHI	4	Maximum free cells
DB8NSCRLO	4	Minimum available scratch
DB8NSCR	4	Average available scratch
DB8NSCRHI	4	Maximum available scratch
DB8NCLN	4	Number of cleaning cartridges
TOTAL	**32**	(excluding key area length)

RECORD TYPE 9: HOURLY CHANNEL GROUP STATISTICS

FIELD	LENGTH	DESCRIPTION
DBCHGBSY	4	Channel group busy samples
DBCHGSMP	4	Total channel path samples
TOTAL	**8**	(excluding key area length)

Appendix D: Database Layout   **217**

RECORD TYPE 10: HOURLY ALLOCATION RECOVERY PER ACS/LSM

FIELD	LENGTH	DESCRIPTION
DB10TOT	4	Total operator reply time (in seconds)
DB10REP	4	Total number of operator replies
DB10LNG	4	Longest operator reply (in seconds)
DB10SVCD	4	Number of WAITs serviced this hour
DB10SVCL	4	Longest WAIT serviced (in seconds)
DB10WAIT	4	Global - number of WAIT replies
DB10NCAN	4	Global - number of CANCEL replies
DB10HIWT	4	Global - highest number of outstanding WAITs
DB10ANYW	4	Global - time with any WAIT outstanding (in seconds)
DB10ANYR	4	Global - time with any reply outstanding (in seconds)
TOTAL	**40**	(excluding key area length)

RECORD TYPE 11: TAPE ERROR EVENT DETAIL

FIELD	LENGTH	DESCRIPTION
DB11PRD	4	Permanent read error count
DB11TRD	4	Temporary read error count
DB11PWR	4	Permanent write error count
DB11TWR	4	Temporary write error count
DB11SSCH	4	Physical I/O count (start subchannel)

DB114KRD	4	Multiples of 4K read from device
DB114KWR	4	Multiples of 4K written to device
DB11ACS	3	ACS number (in character)
DB11LSM	2	LSM number (in character)
DB11VOL	6	Volume serial number
DB11CUU	4	Device address (in character)
TOTAL	**43**	(excluding key area length)

RECORD TYPE 12: DEMAND ENTER EVENT DETAIL

FIELD	**LENGTH**	**DESCRIPTION**
DB12ACS	3	ACS number (in character)
DB12LSM	2	LSM number (in character)
DB12VOL	6	Volume serial number
DB12CUU	4	Device address (in character)
DB12JOB	8	Jobname
TOTAL	**23**	(excluding key area length)

RECORD TYPE 13: HOURLY ALLOCATION RECOVERY FOR NON-ACS DEVICES

FIELD	LENGTH	DESCRIPTION
DB13TOT	4	Total operator reply time (in seconds)
DB13REP	4	Total number of operator replies
DB13LNG	4	Longest operator reply (in seconds)
DB13SVCD	4	Number of WAITs serviced this hour
DB13SVCL	4	Longest WAIT serviced (in seconds)
DB13WAIT	4	Global - number of WAIT replies
DB13NCAN	4	Global - number of CANCEL replies
DB13HIWT	4	Global - highest number of outstanding WAITs
DB13ANYW	4	Global - time with any WAIT outstanding (in seconds)
DB13ANYR	4	Global - time with any reply outstanding (in seconds)
TOTAL	**40**	(excluding key area length)

**TMCHIST: CDS/TMC VOLUME AND DATASET HISTORY
RECORD**

FIELD	LENGTH	TYPE	DESCRIPTION
ORF_USE	4	Binary	Volume usage count
ORF_BLKS	4	Binary	File block size
ORF_BLKC	4	Binary	File block count
ORF_LREC	4	Binary	File record length
ORF_VSEQ	2	Binary	Volume sequence within multi-volume set
ORF_DSNS	2	Binary	Volume number of DSN=
ORF_VOL	6	Char	Volume serial
ORF_FSEQ	2	Binary	File sequence within volume
ORF_ACS	3	Char	ACS location of volume
ORF_LSM	2	Char	LSM location of volume
ORF_DSN	44	Char	File dataset name
ORF_REF	12	Char	Volume last use/reference date
ORF_UTIL	3	Char	Volume percentage utilized
ORF_MEGS	5	Char	Volume megabytes capacity
ORF_MEGX	1	Char	+ = megabytes capacity capped
ORF_SCR	1	Char	Volume scratch status
ORF_RECF	3	Char	File record format
ORF_DEN	1	Char	Volume creation device-type density
ORF_CLEN	2	Binary	Volume cat length

ORF_AGE	2	Binary	Volume age since last use/reference
ORF_DSG	2	Binary	Primary dataset group number
ORF_MVOL	1	Char	Y = Part of multi-volume sets
TOTAL	**110**		

A TMCHIST record exists for each primary and secondary dataset on a volume. SAMPLIB member EXPRHIST contains an assembler layout.

Appendix E

ExPR Version 1 Release 1 Modifications

ExPR Version 1 Release 1 introduces the following new or revised functions:

Automate ExPR CONFIG Update Procedure - Provides automatic collection of site data from the HSC control dataset, eliminating the need for manually specifying device addresses and IDs for each ACS, LSM, and device during the installation and configuration process. Additionally, subsequent changes to the configuration are automatically recognized by ExPR.

Tape Catalog Processing Enhancements - Extends tape catalog reporting by providing cartridge access reporting, summary reporting by device type or length, tape range specification with length/include/exclude options, Redwood support, extended SORT key options, splitting of update and reporting functions to allow report generating without affecting the status of the database, and historical comparisons of volumes.

Allocation Recovery Analysis - Extends mount analysis to fully account for mount delays caused by insufficient drives or operator performance.

Mount Breakdown Report - Extends mount analysis to provide a detailed breakdown of mount time components, including delay time.

Mount Analysis by Drive Type - Extends mount analysis to differentiate between device types for mount count and mount response time.

PC Tabular Reporting - Extends PC graphical reporting options to display the numeric values behind the graphical presentation.

PC Multiple Print Option - Extends PC graphical reporting options to print either the current graph or all displayed graphs.

PC "Save As" Option - extends PC graphical reporting options to save either the current graph or all displayed graphs.

PC Database Purge - Allows records beyond a certain age to be purged from the PC database.

PC Utility Functions Grouping - Combines database utility functions (repair, compress, purge, scan, etc.) under a single icon.

Dynamic Summarization of PC Database - Extends PC graphical reporting options to include hourly, daily, weekly, monthly, quarterly, and yearly summary options.

Microsoft Excel Interface - Extends PC graphical reporting options to include copying graph data into an Excel-compatible spreadsheet through the clipboard feature.

Channel and Control Unit Reporting Enhancements - Extends control unit reports to include control unit saturation

and subsystem loading, and extends channel reporting with the PATH-BUSY report.

Tape Error Detail Report - Extends media error reporting to include the volser and drive associated with each media error, providing an audit trail of bad volumes or drives.

Demand Enter Detail Report - Extends the current enter and eject counts to highlight enters required to immediately service a mount request. An audit trail report additionally gives details of the volumes and jobs involved.

Multiple CPU Support - Extends MVS system-level reporting by grouping data by SMF ID to facilitate reporting of ACS/LSM utilization by individual MVS system in a shared environment, providing a base for future Sysplex data collection.

Exception Reporting Enhancements - Extends exception reporting to include RMF-based channel and control unit data, JES2 and JES3 allocation recovery, and available free cell and scratch information.

Database Verification and Scan Report Modifications - These reports are reformatted to reflect changes in database layout.

Database Conversion Facility - Enables ExPR 1.0 installations to convert existing databases at the host and PC level.

Index

—E—

—F—

—G—

—H—

—I—

—J—