

Xtract

User's Guide

© PALLAS International Corporation, 1988-2007

1. Introduction

1.1 About XtractTM

XtractTM is a data extractor component used to prepare data for our Lead TimeTM modeling product. Together with model generating facilities incorporated into Lead Time, the goal of Xtract is to allow the user to automatically build a Lead Time model corresponding to the observed workload of a selected set of devices in the user's system. Xtract can process data collected by RMF and compatible monitors, as well as data generated by other system monitors such as ASTEXTM from Computer Associates, MONITOR under VM, or CSIM from Amdahl Corporation.

The flow of information between Xtract and Lead Time is quite simple. Based on user's specifications, Xtract preprocesses the monitor data for the selected devices, and produces a data set that contains the information necessary for Lead Time to build the corresponding model. As a next step, the user transfers the resulting data set to the Personal Computer where it is stored as an ASCII file with the extension .SMF. Such a transfer can be accomplished using standard mainframe to PC communications software, FTP, or terminal emulation board.

Lead Time file handling facilities allow then to read-in such an ASCII file through the Fetch command in the Lead Time File menu. Using the values prepared by Xtract, Lead Time builds a model and initializes several quantities not provided directly by the system monitor. These include the (average) block lengths which are inferred from 'connect times', percentages of no motion seeks and device latency which must correlate with the observed I/O service times, maximum numbers of requests for logical volumes, as well as I/O variability patterns which lead to the observed queueing times.

To reduce the processing time, one Xtract run can create several data sets, each corresponding to a single Control Unit. The net result for the user is a greatly simplified model building and validation process to create a base model of the current workload.

As an added convenience, we provide **Xtgen**, a Windows based interface to Xtract that automatically generates the JCL needed to run Xtract. The current version of Xtgen supports most frequently used Xtract options, viz. the treatment of RMF and IXFP records. Xtgen supports up to 24 control units per Xtract run; each control unit can comprise up to 256 LCUs. We hope that you will find Xtgen to be simple and intuitive to use. All you have to do is to specify the controllers to be treated, for each controller, the desired LCUs and devices, and, for Shared DASD environments, the system ID of the desired system image. Based on your specification, Xtgen generates the necessary JCL and stores it in a file on your PC. *Unless you are planning to use Xtract with GTF traces, ASTEX or MONITOR records, or for extraction at data set level, the on-line help in Xtgen is virtually all the Xtract JCL knowledge you need.* Please read Section 5 of this manual for installation and other important information. Please also consult the Addendum notice at the end of this manual for a brief description of Xtgen dialogs, as well as frequently encountered problems. Additionally, the PALLAS International Web site (www.pallas-international.com) has a section devoted to the latest features and FAQs related to Xtract and Lead Time.

1.2 Processing in XtractTM

To the extent that the data collected by a monitoring facility and processed by Xtract reflects the "true"

activity of the DASD subsystem, it allows the creation of a Lead Time model of DASD subsystem behavior based on realistic workload parameters.

Please keep in mind that any model is an abstraction of reality, taking into account only certain - hopefully, the most important - of its aspects. Even when a base model has been built and calibrated using measured data, no guarantee can be offered that, as parameter values are varied, model results will always match reality.

Xtract can process SMF records created by RMF (or a fully compatible monitor), ASTEX, or the Cache RMF Reporter under MVS/XA, MVS/ESA, z/OS. GTF records also must have been created under MVS/XA or MVS/ESA, z/OS. Records generated by CSIM may additionally pertain to a trace generated under VM/XA SP. Xtract supports also records generated by the MONITOR command of VM/XA SP2 or VM/ESA. Xtract can run under MVS/370, MVS/XA, MVS/ESA, z/OS or VM/CMS. The data processed can have been created under another Operating System (e.g. MONITOR data can be processed under MVS/ESA, or RMF data under VM/CMS).

1.2.1 Processing RMF and compatible records

Xtract cumulates RMF information (SMF records type 73, 74, and 78) for the DASD devices specified by the user. Normally, device information for each Logical Control Unit (LCU) is cumulated into groups, and the group data is written into a sequential data set, to be later transferred to a Personal Computer in a format usable by the Lead Time product. In this Xtract mode, the number of sequential data sets created will be equal to the number of different LCUs for the devices specified by the user.

It is also possible to create a single sequential data set, even if the devices specified span several LCUs. See the description of the LCUGRP keyword in Section 2.3 for more information.

1.2.2 Processing ASTEXTM records

ASTEXTM (or DASDMON) records provide information at the level of individual data sets, and are thus an excellent starting point for a data set level Lead Time model of the configuration under consideration. Xtract automatically distinguishes between ASTEX and DASDMON records based on their internal format.

Analogously to its treatment of RMF records, Xtract cumulates ASTEX information for the DASD devices specified by the user. Normally, device information is cumulated into groups separately for each Contention Group (similar in use to LCUs with RMF records), and the resulting data is written into sequential data sets, one per Contention Group. In this Xtract mode, there will be as many sequential data sets created as there are Contention Groups for the devices specified by the user.

It is also possible to request the creation of a single sequential data set even if the devices processed span several Contention Groups. For more information, see the description of the LCUGRP keyword in Section 2.3.

Each of such sequential data sets created by Xtract can be later transferred to the Personal Computer and given the extension .SMF for use by the Lead Time file handling facilities in the building of a base model.

Note that ASTEX provides considerably more information than simple RMF since it includes data on the activity of individual data sets. However, it does not provide the LCU number or the number of data sets open which are a part of the standard RMF information.

1.2.3. Processing ASTEX with RMF

In this mode, in addition to cumulating ASTEX or DASDMON information about the DASD devices specified by the user, Xtract uses RMF records (type 73, 74, and 78) to collect information regarding Logical Control Units (LCUs), the number of data sets open, as well as the channel configuration and load.

As specified using the LCUGRP keyword (cf. Section 2.3), for each Contention Group or for each LCU, device information is cumulated into groups. The group data is written into sequential data sets, to be later transferred to a Personal Computer in a format usable by the Lead Time product. Again, the number of data sets created is equal to the number of distinct Contention Groups or LCUs for the devices under consideration. The LCUGRP keyword can be used to force the creation of a single sequential data set even if the devices span several Contention Groups or LCUs (see Section 2.3).

1.2.4 Processing CSIM records

Xtract uses the mandatory **UNITS** and the optional **DSNS** data sets created by the Table Dump Facility of CSIM V8R0 or later for the DASD devices specified by the user. Lower versions of CSIM are not supported. Both the UNITS and DSNS data sets must be created by a single CSIM run. Device information is cumulated into groups, and the group data is written as two sequential data sets, to be later transferred to the Personal Computer in a format useable by the Lead Time software.

The first data set (OUT001) is analogous to the data set created by an RMF Xtract run. It describes the I/O activity as observed by CSIM in the trace data, presumably collected in an uncached environment. This data set should be transferred to the Personal Computer as an ASCII file with the extension .SMF, and used for automatic model initialization through the Lead Time Prime facility.

The second data set (SIM001) describes various ratios (e.g. cache hits) and device seek times as predicted by CSIM in a cache simulation run. This data set should be transferred to the Personal Computer as an ASCII file with the extension .MDL. Normally, it should be read by Lead Time following the initialization of the model based on the .SMF file (the OUT001 data set). This can be accomplished using the File Retrieve facility with the Keep option so as to preserve the data obtained from the .SMF file. In this way, the existing initialized model is modified to reflect the cache size and other options specified by the user in the CSIM simulation.

Note that although CSIM provides considerably more information than standard RMF since it includes data on the activity of individual data sets, it does not provide the LCU number, the average number of data sets open, nor the average IOSQ time found in RMF records.

1.2.5 Processing both CSIM and RMF records

In this mode, in addition to processing CSIM information, Xtract uses RMF records (type 73, 74, and 78) to collect data regarding Logical Control Units (LCUs), the average number of data sets open, pending and IOSQ times, as well as channel configuration and load.

Ideally, the selected RMF period should match that of the trace used for CSIM. Xtract will check that the Volume Serial Numbers in the RMF records match those in the CSIM records, as well as that there are no major differences in the workload characteristics for each device (no more than 10%) between both types of records. The pending and IOSQ times reported by RMF will be used only if the match is deemed sufficient.

For each LCU, as specified by the LCUGRP keyword (cf. Section 2.3), the device information is cumulated into two sequential data sets, to be later transferred to the Personal Computer in a format useable by the Lead Time modeling product. Two data sets (OUTnnn and SIMnnn) are created per LCU for the devices under consideration. These data sets should be transferred to the Personal Computer as ASCII files with the extensions .SMF, and .MDL, respectively. Their use is described in Section 1.2.4. Note that the LCUGRP keyword can be employed to force the creation of only two data sets even if the devices span several LCUs.

1.2.6 Processing GTF records

Xtract uses the trace data set created by the Generalized Trace Facility (GTF) under MVS/XA or MVS/ESA for the DASD devices specified by the user. A subsequent section describes the proper parameters to be used for GTF execution. Optionally, Xtract reads the Volume Table of Contents (VTOC) of the specified devices to provide information at the level of individual data sets. Device information is cumulated into groups, and then written as two sequential data sets, to be later transferred to the Personal Computer in a format useable by the Lead Time software package.

The first data set (OUT001) is analogous to the data set created by an RMF or another monitor run. It describes, in the form of appropriate averages, the I/O activity as recorded in the GTF trace. This character data set should be transferred to the Personal Computer as an ASCII file with the extension .SMF, and used for automatic model initialization through the Lead Time Fetch and Prime facility.

The second data set (GTF001) contains the GTF trace after pre-processing in a condensed binary form. This data set should be transferred to the Personal Computer **"as is"** (i.e., without any character translation) and given the extension .GTF. A special facility, called **GTFSIM**, included in the Lead Time package allows to read the .GTF trace file, and to run a simulation of a cached or uncached environment based on the I/O commands in the trace.

A GTFSIM simulation produces estimates of DASD seek times, as well as of various cache statistics such as hit ratios in the form of a Lead Time model file (.MDL).

If the parameter setting used for the GTFSIM simulation corresponds to that under which the GTF trace was collected, the .MDL file should be read-in prior to Lead Time model initialization. If the GTFSIM simulation is used to assess the impact of a different cache or NVS size, or, in general, is run for a parameter setting which does not match that under which the GTF trace was collected, the resulting .MDL file should normally be read by Lead Time following the initialization of the model.

In both cases, the .SMF file (the first data set) must be fetched first. To read in the .MDL file produced by GTFSIM, the File Retrieve facility with the Keep option should be used so as to preserve the data obtained from the .SMF file.

In most cases, for best initialization of the base model, it is desirable to run a GTFSIM simulation with the parameter setting corresponding to the conditions under which the GTF trace was collected so as to complement the data in the OUT001 data set (.SMF file) produced by Xtract. Indeed, certain statistics computed by GTFSIM are not present in this data set.

Note that although GTF provides considerably more information than standard RMF since it allows Xtract to collect data on the activity of individual data sets, it does not provide the LCU number, the number of data sets open, nor the average IOSQ time found in RMF records.

1.2.7 Processing GTF and RMF records

In this mode, in addition to processing GTF information, Xtract uses RMF records (type 73, 74, and 78) to collect information regarding Logical Control Units (LCUs), the average number of data sets open, pending and IOSQ times, as well as channel configuration and load.

Ideally, the selected RMF period should match that of the GTF trace. Xtract will check that there are no major differences in the workload characteristics for each device (no more than 10%) between both sources of information. The pending and IOSQ times reported by RMF will be used only if the match is deemed sufficient. See also the NOMATCH keyword in Section 2.4.

With the inclusion of RMF records, LCU information becomes available. Thus, for each LCU, as specified by the LCUGRP keyword (cf. Section 2.3), the device information is cumulated into two sequential data sets to be later transferred to the Personal Computer in a format useable by the Lead Time modeling package.

Two data sets (OUTnnn and GTFnnn) are created per LCU for the devices under consideration. These data sets should be transferred to the Personal Computer with the extensions .SMF and .GTF, respectively. Their use is described in the preceding section. **Note that the LCUGRP keyword can be used to force the creation of only two data sets even if the devices span several LCUs.**

1.2.8 Processing RMF records and multiple GTF traces

In this mode, multiple GTF traces are processed. They may have been collected on the same system (e.g. to average out DASD activity at different periods, or to represent a planned consolidation on a larger controller) or on different systems (e.g. in Shared DASD configurations).

Ideally, the selected RMF period should match that of the GTF trace. Also, records from all the systems should be present in the RMF data. See also the NOMATCH keyword in Section 2.4.

The user must specify the systems on which the trace was collected using the **TRACEn** statements as in the following example:

```
TRACE1 = SYS1
```

TRACE2 = MVS4
TRACE3 = MVS7.

In the above example, the user specifies that three GTF trace data sets must be processed. The dataset referred to by DDNAME TRACE1 corresponds to an MVS system with System ID SYS1, the one referred to by DDNAME TRACE2 corresponds to System ID MVS4, and the one referred to by DDNAME TRACE3 to an MVS system with System ID MVS7.

Given the intended application to merging of traces, it is difficult to count on the traces being really synchronized. Therefore, Xtract forces the synchronization of traces by assuming that each trace starts at time zero. It is the responsibility of the user to ensure that the traces are started simultaneously if real synchronization is desired. It is also up to the user to select the RMF records which correspond to the traced period(s) using the ST1, ND1, ST2, ND2, STDATE and NDDATE keywords.

Please note that these keywords apply to the RMF (SMF) data, and not the GTF data. The GTF data is always entirely processed. Only RMF records for the specified system IDs (SYS1, MVS4, MVS7 in our example) are processed. Xtract can merge up to eight traces (with DDNAMES TRACE1 through TRACE8) calling whatever sort program is in use in the user's environment. Xtract assumes the program name to be SORT. If the name of the sort program is different from SORT, the user must specify

SORT = sortname

where sortname is the name of the sort program.

With multiple GTF traces, the user must include the necessary statements for sort processing such as

```
//SORTWK01 DD UNIT=SYSDA,SPACE=(CYL,(20))  
//SORTWK02 DD UNIT=SYSDA,SPACE=(CYL,(20))  
//SORTWK03 DD UNIT=SYSDA,SPACE=(CYL,(20))  
//SORTWK04 DD UNIT=SYSDA,SPACE=(CYL,(20))  
//SYSOUT DD SYSOUT=*
```

Please refer to Section 5.6 for an example of JCL control statements and parameters to process multiple GTF traces.

1.2.9 Processing MONITOR records

Xtract processes the following types of MONITOR records:

- Domain 0, Record 9 (Channel Connection Data)
- Domain 1, Record 6 (Configuration Data)
- Domain 6, Record 1 (Vary Online)
- Domain 6, record 3 (Device Activity)
- Domain 6, Record 4 (Cache Activity)
- Domain 6, Record 5 (Attach Device).

The MONITOR file must have been created by the VM/CMS "MONWRITE" command.

Xtract requirements match those of the IBM "VMPRF" program to produce the following reports:

- Cache_DASD_by_Activity
- Channel_Busy
- DASD_by_Activity.

If needed, refer to VM/PRF User's Guide and Reference (SC23-0460-02).

Because MONITOR data does not contain Logical Control Unit (LCU) identification, Xtract creates a virtual LCU for devices accessed via the same Channel Paths, and showing the same type of Control Unit.

If Xtract runs under VM/CMS, "ALLOC = NO" is forced, i.e., the devices are not allocated, and VTOCs are not read.

To treat data produced by the MONITOR command under VM/XA SP2 or VM/ESA, specify

PROCESS = MONITOR.

1.2.10 Dynamic device allocation and VTOCs

To attempt to determine the control unit type and the device type for the devices specified by the user, Xtract dynamically allocates them in order to issue the SENSE ID command. Because MVS rules require that a data set be open on each device, Xtract attempts to open the VTOC data set. Before doing so, Xtract ensures that a UCB exists for the given device address, that it is flagged as a DASD device, and that Volume Serial Number matches that of RMF, ASTEX or DASDMON, or CSIM records. If the access to the VTOCs of some volumes is restricted by products such as RACF or ACF2, Xtract should be given the necessary authority, or "ALLOC = NO" should be specified to bypass the dynamic device allocation.

When processing GTF records, if "ALLOC = YES" is in effect, Xtract additionally reads the VTOCs of the devices under consideration, and builds internal data set tables. Note that Xtract must be able to access the VTOCs to produce information at data set level when processing a GTF trace.

Xtract supports device and control unit type information available in RMF release 4. If this information is available in the records processed, Xtract will not try to allocate the devices to determine their type. Note, however, that VTOC information is read from devices if GTF data is being processed and "ALLOC = YES" is in effect.

1.3 Grouping techniques

Xtract supports four different techniques for device grouping.

1. The number of groups specified by the user (or defaulted to) applies to each String of a given LCU. One output data set (two, in the case of CSIM or GTF) is created for each LCU or Contention Group.
2. The number of groups specified by the user (or defaulted to) applies to each LCU. One output data set (two, in the case of CSIM or GTF) is created for each LCU or contention group.
3. The number of group specified by the user (or defaulted to) applies to all devices included. A single output data set (two, for CSIM or GTF) is created by Xtract.

4. The user specifies explicitly the contents of each group. A single output data set (two, with CSIM or GTF) is created.

When dealing with array control units (such as EMC Symmetrix™ series, IBM ESS, HDS 7700, 9900, STK V2X), it is recommended to generate one group for each logical device. This is accomplished using the specification

GROUPS = 0

Refer to Section 3 for a discussion of device grouping.

1.4 What follows

Section 2 of this manual describes the keywords used to control the execution of Xtract. Section 3 explains how Xtract organizes device information into groups. Examples of control statement are given in Section 4. Section 5 outlines the procedures to be used to install and run Xtract in an MVS/370, MVS/XA, MVS/ESA, z/OS, or VM/CMS Operating System environment. The Appendix lists the condition codes generated by Xtract execution including abnormal end codes, as well as the handling of overflow conditions. It describes also the reports produced by Xtract.

2. The Xtract™ control data set

2.1 Description and syntax

The execution of Xtract is controlled by the SYSIN data set. The parameters in this data set are in pairs of the form

"KEYWORD=OPERAND".

For simplicity of use, defaults are provided for most parameters. For maximum flexibility, over 30 parameters can be specified. In our discussion, we divide the latter into three categories:

- Selection Parameters
- Device and Group Definitions
- Miscellaneous Parameters.

The following rules apply when specifying parameters in the control data set.

- Each keyword must be followed by an equal sign '='. There may be blanks between the keyword and the equal sign.
- The first nonblank character following the equal sign is taken to be the first character of the operand.
- A blank, a comma, or reaching column 73 terminates the operand. The set composed of a keyword and its operand must be contained between columns 1 and 72.
- The logical record length of the SYSIN data set is 80 bytes. Columns 1 through 72 being used for keywords and operands, sequence numbers may be placed in columns 73 to 80, if desired.
- Space permitting, several "KEYWORD=OPERAND" pairs may be placed in one logical record. The keywords and the operands may be specified in any order, except as noted in the following discussion. Comment statements (indicated by an asterisk (*) in column one) may be used.

2.2 Selection parameters

Selection parameters are used to specify which SMF records should be processed. The first seven selection parameters define the days of the week for which SMF data should be selected for processing. The vertical bar '|' is used to denote options only one of which can be chosen.

MONDAY = YES | NO

The MONDAY keyword is used to specify whether SMF records created on Mondays should be processed. This keyword can be abbreviated to "MON". MONDAY = YES is the default.

TUESDAY = YES | NO

The TUESDAY keyword is used to specify whether SMF records created on Tuesdays should be processed. This keyword can be abbreviated to "TUE". TUESDAY = YES is the default.

WEDNESD = YES | NO

The WEDNESD keyword is used to specify whether SMF records created on Wednesdays should be processed. This keyword can be abbreviated to "WED". WEDNESD = YES is the default.

THURSD = YES | NO

The THURSD keyword is used to specify whether SMF records created on Thursdays should be processed. This keyword can be abbreviated to "THU". THURSD = YES is the default.

FRIDAY = YES | NO

The FRIDAY keyword is used to specify if SMF records created on Fridays should be processed. This keyword can be abbreviated to "FRI". FRIDAY = YES is the default.

SATURD = YES | NO

The SATURD keyword is used to specify if SMF records created on Saturdays should be processed. The SATURD keyword can be abbreviated to "SAT". SATURD = NO is the default.

SUNDAY = YES | NO

The SUNDAY keyword is used to specify if SMF records created on Sundays should be processed. This keyword can be abbreviated to "SUN". SUNDAY = NO is the default.

The following two parameters allow you to specify in Julian format the range of dates for which SMF data should be processed.

STDATE = yyyyddd

The STDATE keyword is used to specify the starting day of the date range. The STDATE keyword can be abbreviated to "STD". STDATE = 1988001 is the default and the minimum value for this parameter.

NDDATE = yyyyddd

The NDDATE keyword specifies the ending day of the date range for which SMF records should be processed. The NDDATE keyword can be abbreviated to "NDD". NDDATE = 2010365 is the default and the maximum value for this parameter.

Note that the STDATE and NDDATE keywords define the limits of the date range. SMF data outside this range (if any) present in the input file will not be processed. In our notation, yyyyddd denotes the year and the day of the year.

As an example, specifying

STD=91001, NDD=91365

would cause Xtract to process SMF records generated during the whole year 1991.

The next four parameters are used to define up to two periods of time for which SMF records should be selected for processing. In our notation, hhmm denotes hours and minutes. Hours are taken in 24-hour format.

STTIME1 = hhmm

The STTIME1 keyword specifies the beginning of the first period for which SMF records should be processed. The STTIME1 keyword can be abbreviated to "ST1". STTIME1 = 0000 is the default.

STTIME2 = hhmm

The STTIME2 keyword is used to specify the beginning of the second period for which SMF records should be processed. The STTIME2 keyword can be abbreviated to "ST2". STTIME2 = 0000 is the default.

NDTIME1 = hhmm

The NDTIME1 keyword specifies the end of the first period for which SMF records should be processed (beginning at time STTIME1). The NDTIME1 keyword can be abbreviated to "ND1". NDTIME1 = 2400 is the default.

NDTIME2 = hhmm

The NDTIME2 keyword is used to specify the end of the second period for which SMF records should be processed (beginning at time STTIME2). The NDTIME2 keyword can be abbreviated to "ND2". NDTIME2 = 2400 is the default.

As an example, specifying

ST1=0930, ND1=1145, ST2=1315, ND2=1630

would cause Xtract to process SMF records generated between 9H30 am and 11H45 am, and those generated between 1H15 pm and 4H30 pm.

If the user prefers, he or she can specify a single period.

Since Xtract can process data generated by several system monitors, the PROCESS keyword is provided to select the type of records to be processed.

PROCESS = record type

The following record types are currently accepted: RMF, CARMF, ASTEX, DASDMON, GTF, CSIM, and MONITOR. The use of additional record types (IXFP, SMF42) is described in the addendum.

If "PROCESS = RMF" is specified, Xtract processes only RMF type 73, 74, and 78 records. If "PROCESS = CARMF" is used, Xtract processes additionally RMF Cache Reporter records.

If the specification is "PROCESS = ASTEX" or "PROCESS = DASDMON", Xtract processes ASTEX or DASDMON "Volume" and "Data Set" records. Since Xtract distinguishes automatically between ASTEX and DASDMON records, both specifications can be used interchangeably. RMF type 73, 74, and 78 records are also processed (if present in the input data set) to gather additional information.

If "PROCESS = CSIM" is specified, Xtract processes CSIM "UNITS" (mandatory) and "DSNS" (optional) records, as well as RMF type 73, 74, and 78 records (if present) to collect additional information.

If "PROCESS = GTF" is specified, Xtract processes GTF trace records produced under MVS/XA or MVS/ESA. If "ALLOC = YES" is in effect, the VTOCs of the specified devices are read to correlate I/O operations with individual data sets located on the DASD devices. Optionally, if present, RMF records type 73, 74 and 78 are processed to collect additional information.

If "PROCESS = MONITOR" is specified, Xtract processes records generated by the MONITOR command of VM/XA SP2 or VM/ESA. Note that Xtract can run under MVS/370, MVS/XA, MVS/ESA, z/OS or VM/CMS. The data processed can have been created under another Operating System (e.g. MONITOR data can be processed under MVS/ESA).

Xtract processes the following types of MONITOR records:

- Domain 0, Record 9 (Channel Connection Data)
- Domain 1, Record 6 (Configuration Data)
- Domain 6, Record 1 (Vary Online)
- Domain 6, record 3 (Device Activity)
- Domain 6, Record 4 (Cache Activity)
- Domain 6, Record 5 (Attach Device).

The MONITOR file must have been created by the VM/CMS "MONWRITE" command.

Xtract requirements match those of the IBM "VMPRF" program to produce the following reports:

- Cache_DASD_by_Activity
- Channel_Busy
- DASD_by_Activity.

If needed, refer to VMPRF User's Guide and Reference (SC23-0460-02).

Because MONITOR data does not contain Logical Control Unit (LCU) identification, Xtract creates a virtual LCU for devices accessed via the same Channel Paths, and showing the same type of Control Unit. If Xtract runs under VM/CMS, "ALLOC=NO" is forced, i.e., the devices are not allocated, and VTOCs are not read.

The PROCESS keyword can be abbreviated to "PROC". PROCESS = RMF is the default.

In cases when there is no "standard" SMF record ID for the records created by a monitor, or if the

standard default value is not used, it is necessary to indicate the SMF record ID used. In particular, this is the case when the information processed comes from RMF Cache Reporter, and the default SMF record ID (245) is not used. The SMFREC keyword is then necessary.

SMFREC = number

This keyword specifies the record code (SMFREC#) for the SMF records generated by the monitor. Note that the decimal number specified must be in the range 1 through 255. The SMFREC keyword can be abbreviated to "SMF".

The SMFREC keyword can also be used for ASTEX or DASDMON records if the same SMF record ID is used for both "Volume" and "Data Set" records. If distinct record IDs are used, they should be specified via the VOLREC and DSNREC keywords, respectively.

VOLREC = number

This keyword specifies the record code (SMFREC#) for the SMF "Volume Records" generated by ASTEX or DASDMON. The decimal number specified must be in the range 1 through 255. The SMFREC keyword can be employed if the same record code is used for both "Volume" and "Data Set" records.

DSNREC = number

This keyword is used to specify the record code (SMFREC#) for the SMF "Data Set Records" generated by ASTEX or DASDMON. The decimal number specified must be in the range 1 through 255. If the same record code is used for "Volume" and "Data Set" records, the SMFREC keyword can be employed to specify the common record code.

If "**PROCESS = RMF**" is specified (explicitly or by default), the SMFREC, VOLREC, DSNREC specifications are ignored.

SYSID = xxxx

The SYSID keyword is used to specify the SMF system ID (4 alphanumeric characters) of the RMF records to be processed by Xtract. If SYSID is not specified, RMF records from all systems are processed. If SYSID is used, only the records with the specified ID are processed. *Note that SYSID applies to RMF and RMF Cache Reporter records, not to records created by ASTEX or DASDMON, CSIM, or GTF.*

When processing multiple GTF traces, the TRACEn statements are used to specify the systems on which the traces were collected. The syntax of this statement is

TRACEn = xxxx

An example might be **TRACE2 = MVS4**.

2.3 Device and group definitions

There are several ways to specify the devices for which the SMF data is to be processed by Xtract. Devices can be explicitly and individually selected through their device numbers (keyword DEVICES), or implicitly globally selected by specifying the Logical Control Unit number (keyword LCU) or the Contention Group name (keyword CONTGP) with DASDMON records. In these cases, Xtract will divide each string of devices into groups using the number of groups specified by the user and its own built-in grouping rules (cf. Section 3). Depending on the value of the LCUGRP parameter, Xtract will generate one output data set for each LCU (or Contention Group), or just a single data set (two output data sets are generated in the case of CSIM).

The user can also explicitly define groups of devices (keyword GROUPn) in which case a single output data set will be generated, regardless of string or LCU considerations. Note that devices placed in a single group in Xtract will be treated as a single "device group" in Lead Time. If needed, please refer to Lead Time documentation for more details on the notion of device group.

DEVICES = (cuu,cuu,cuu-cuu,.....cuu)

The DEVICES keyword may be abbreviated to "DEV". This keyword specifies the device numbers of the devices to be processed. The operand consists of a series of 3 or 4 hexadecimal characters. The device numbers must be separated by commas, without intervening blanks. Ranges of consecutive device numbers can be specified as cuu-cuu without intervening blanks. Both types of syntax may be mixed in a single statement. You may specify up to 256 device numbers in one run of Xtract, using as many DEVICES keywords as necessary.

As an example,

DEV = (160-16F,186-188,170)

means that devices 160 through 16F, 186 through 188 as well as device 170 are to be processed.

Except in cases when the "LCU" keyword, the "CONTGP" keyword, or the "GROUPn" keyword are used, the DEVICES keyword must be used to specify device numbers. If the DEVICES keyword is used, then neither the LCU nor CONTGP nor GROUPn keywords may be specified.

LCU = nn | nnn

This keyword specifies the Logical Control Unit Number (LCU) of the devices to be processed. The operand consists of 2 or 3 hexadecimal characters. The LCU number allows the user to select all the devices belonging to a given LCU, without specifying each device number. It is also possible to specify multiple LCU numbers in a single statement

LCU = (nn,nnn,.....,nn,nnn-nnn)

Note that ranges of consecutive LCUs are specified as nnn-nnn without intervening blanks. Both types of

syntax can be mixed in a single statement. If desired, the LCU keyword may be used more than once to process devices belonging to several LCUs.

If the LCU keyword is used, then neither the DEVICES nor CONTGP nor GROUPn keywords may be specified. Note that if the LCU keyword is used when processing ASTEX or DASDMON, GTF, or CSIM records, type 74 records must be present in the SMFIN data set.

CONTGP = name

The CONTGP keyword implicitly defines the devices to be processed through the name of Contention Group to which they belong. The name operand may comprise 1 to 6 characters. If the CONTGP keyword is used, then neither the DEVICES nor LCU nor GROUPn keywords may be specified. Since the notion of contention group is specific to DASDMON records, this keyword should only be used in conjunction with the "PROCESS = ASTEX" or "PROCESS = DASDMON" statement.

STRILEN = NO | 8 | 16 | 32|64

The STRILEN keyword relates to the organization of devices into strings. To create groups within strings, Xtracts needs to know the maximum length of strings on the control unit under consideration.

If STRILEN = NO is specified, Xtract will not attempt to recognize physical strings, and groups will be created at the LCU level (or, depending on the LCUGRP parameter, at the global level) rather than at string level.

If STRILEN=8 is specified, device numbers in each consecutive sequence of 8 hexadecimal numbers (00 through 07, 08 through 0F, etc) will be taken to belong to the same string. The number of groups specified by the user or defaulted to will apply to each string separately.

Similarly, if STRILEN=16, STRILEN=32, or STRILEN=64 is specified, devices in each consecutive sequence of 16 (respectively, 32 or 64) hexadecimal numbers will be assumed to belong to the same physical string. For STRILEN=32, the corresponding sequences are 00 through 1F, 20 through 3F, etc. Again, the specified or defaulted to number of groups will apply to each string.

GROUPS = number

The GROUPS keyword may be abbreviated to "GR". This keyword specifies the number of device groups to be generated. As mentioned before, the number of groups applies to each string, LCU/Contention Group, or to all devices under consideration, depending on the values of the STRILEN and LCUGRP parameters. **GROUPS = 0 is the default resulting in a single volume per group.** The maximum value for GROUPS is defined by the MAXGP parameter. Refer to Section 3 for an explanation of how groups are created.

GROUPn = (cuu,cuu,cuu-cuu,.....cuu)

The GROUPn keyword (where "n" is a number from 1 to MAXGP) may be abbreviated to "GRn". This keyword specifies both the device number of the devices to be processed, and the group to which they

belong. The operand consists of a series of 3 or 4 hexadecimal characters. The device numbers must be separated by commas, without intervening blanks. As for the DEVICES keyword, ranges of device numbers (cuu-cuu) may be included. You may specify up to 256 device numbers in one run of XTRACT, using as many GROUPn keywords as necessary.

If the GROUPn keyword is used, then neither the DEVICES nor LCU nor CONTGP keywords may be specified. Note that when devices are grouped using this keyword Xtract produces a single output data set, OUT001 (two for a CSIM run).

MXDVGP = number

This keyword specifies the maximum number of devices per group. MXDVGP=128 is the default, and the maximum value for MXDVGP is 128. The MAXDVGP keyword must be placed in the input stream before any GROUPS or GROUPn keywords.

MAXGP = number

The MAXGP keyword defines the maximum number of groups which may be generated in a given grouping unit (string, LCU, or globally all devices). MAXGP = 32 is the default. The maximum value for this parameter is 255. The MAXGP keyword must be placed in the input stream before any GROUPS or GROUPn keywords. If the values for MXDVGP and MAXGP conflict, Xtract issues an explanatory message and terminates.

Note that the current version of Lead Time™ requires that this parameter be set to a value such that no string exceeds 64 device groups.

NGRSTR = number

This keyword is used only when groups are constructed by the user (through the GROUPn keyword) or when STRILEN=NO is specified. It then defines the maximum number of groups per string, so that Xtract generates a new string after NGRSTR groups. NGRSTR = 64 is both the default and the maximum value for this parameter.

ZEROS = YES | NO

This keyword defines the treatment of groups for which the observed I/O rate is below 0.005 I/Os per second. Since Xtract keeps only two decimals, the resulting I/O rate for such a group would be 0.00. If included in the .SMF file, such groups would be signaled as having an unacceptable I/O rate by Lead Time when you request the building of a base model. Setting ZEROS = NO lets you exclude such groups from the output data set. ZEROS = NO is the default.

If you wish to keep groups with an I/O rate of 0.00 in the output data set, you should use ZEROS=YES, and explicitly change the I/O rates for such groups within Lead Time after fetching the model, before you request model priming.

When processing records pertaining to data set activity (e.g created by GTF, ASTEX or CSIM), the ZEROS keyword applies also to data sets.

LCUGRP = YES | NO | CNTGP

The LCUGRP keyword is used to help specify how Xtract should generate output data sets. If "LCUGRP = YES" is used, Xtract will generally create one output data set for each Logical Control Unit (LCU). If "LCUGRP = NO" is specified, Xtract will generate a single output data set regardless of the number of LCUs or Contention Groups. For a GTF or CSIM run, pairs of output data sets are always generated.

Note that "LCUGRP = CNTGP" should be used only if "PROCESS = ASTEX" or "PROCESS = DASDMON" is in effect. While "LCUGRP = YES" is generally the default for this parameter, in the case when the monitor data comes from ASTEX or DASDMON and RMF type 74 records are not present in the SMFIN data set, "LCUGRP = CNTGP" becomes the default.

The GPRATE keyword is used to specify a threshold (in I/Os per second) such that devices whose I/O rate is higher or equal to this value are each assigned their own group (i.e., they are alone in this group). All devices whose observed I/O rate was below this threshold value, are then placed in a single "catch all" group. The syntax is

GPRATE = number

As an example, GPRATE=5 specifies that each device with an I/O rate not below 5 I/O per second will be assigned its own group, all other devices will be grouped together in a single group. Clearly, the GPRATE specification should not be used together with another grouping specification (e.g. GROUPS = n).

2.4 Miscellaneous definitions

We describe here the remaining keywords.

LINECNT = number

This keyword is used to specify the number of lines per page for report listings (SYSPRINT). The minimum value for this parameter is 15. The LINECNT keyword can be abbreviated to "LN". LINECNT= 60 is the default.

MAXDSX = number

The MAXDSX keyword specifies the maximum number of data sets handled by Xtract in cases of data set level monitor information (e.g. with "PROCESS = ASTEX"). Xtract uses this value to acquire the main storage used to process data set information. For each data set, 172 bytes of main store are required. If a large MAXDSX is specified, the REGION parameter in the JOB or EXEC JCL statement must be updated accordingly. If "MAXDSX = 0" is specified, Xtract skips the processing of data set records.

If the value of MAXDSX is less than the actual number of data sets for the specified period and dates in

the SMFIN data set, no data sets are processed and a warning is issued by Xtract. If "PROCESS = RMF" is in effect, this parameter is ignored. MAXDSX = 512 is the default.

MAXDSL = number

This keyword specifies the maximum number of data sets to be included in a single output data set, and thus passed to Lead Time. If more than MAXDSL data sets belong to a single output data set, Xtract will make sure that only MAXDSL data sets are included by combining the overflow data sets with the lowest I/O rates in a single data set named "*** VOLUME volser ***", where volser denotes the volume serial number. The MAXDSL parameter is ignored if "PROCESS = RMF" is in effect. MAXDSL = 512 is the default.

With the current version of Lead Time™, this parameter should be set to a value below 10,000.

ALLOC = YES | NO

The ALLOC keyword specifies whether Xtract should attempt to allocate the target devices to determine their type. If GTF trace data are processed, the VTOCs of the target devices are read in order to obtain data set level information. ALLOC = YES is the default. If the access to the VTOCs of some volumes is restricted by products such as RACF or ACF2, Xtract must be given the necessary authority, or "ALLOC = NO" should be used to bypass the dynamic device allocation and VTOC access altogether.

Xtract supports device and control unit type information introduced in RMF release 4. If this information is available in the records processed, Xtract will not try to allocate the devices to determine their type. However, VTOC information is read from devices if GTF data is being processed and "ALLOC=YES" is in effect.

Note that in order to obtain data set level information with a GTF trace, Xtract must be able to read the VTOCs of the devices under consideration, i.e., ALLOC = YES must be in effect.

DSORT = DSN | IOR

This keyword defines how the Data Set Activity report printed by Xtract should be sorted. If "DSORT = DSN" is specified, the report will be sorted by Volume Serial number and by Data Set name within each Volume Serial number. If "DSORT = IOR" is used, the report will be sorted by Volume Serial number and by I/O rate within each Volume Serial number. DSORT = DSN is the default.

The NOMATCH keyword is used to specify allowable mismatch between GTF and RMF data before Xtract issues a warning message. The syntax is illustrated by the following example

NOMATCH = 25

specifies an allowable mismatch of 25% before warnings are issued. NOMATCH = 10 is the implied default.

When processing multiple GTF traces, Xtract can merge up to eight traces (with DDNAMES TRACE1 through TRACE8) calling whatever sort program is in use in the user's environment. Xtract assumes the

program name to be SORT. If the name of the sort program is different from SORT, the user must specify

SORT = sortname

where sortname is the name of the sort program.

3. Creation of groups

3.1 General information

The user is free to specify how the groups are organized, using **GROUPn** keywords. In this case, XTRACT will only ensure that no group contains more devices than specified by the **MXDVGP** parameter.

The user can also let Xtract perform the grouping by simply specifying the desired number of groups (**GROUPS** keyword). Xtract will then use built-in grouping rules described in the following sections.

When dealing with array control units, it is recommended not to "lump" logical devices, i.e., to keep a single logical device per group (see 3.7). This current Xtract default.

3.2 Group naming conventions

With several exceptions, the names given to device groups created by Xtract are formed to include the number of the Logical Control Unit to which the devices of the group belong. Thus, the groups are named GRnnn01, GRnnn02, etc, where nnn denotes the LCU number.

This naming convention is not followed in cases where the LCU is not the basis for device grouping, i.e., if the CNTGP keyword is specified or if "LCUGRP = NO" or "LCUGRP = CNTGP" is used. Then, the groups are simply named GROUP01, GROUP02, etc.

Finally, whenever a group contains a single device the Volume Serial Number of the device is used as the group name.

3.3 Scope of grouping

The scope of grouping is controlled by the **STRILEN** and **LCUGRP** keywords as follows.

If **STRILEN** has a value other than NO (i.e., 8, 16, 32, or 64), and the **LCUGRP** parameter has a value other than NO, the specified or defaulted to number of groups applies to each string of each LCU or Contention Group.

If **STRILEN = NO** is specified, and **LCUGRP** has a value other than NO, the specified or defaulted to number of groups applies to each LCU or Contention Group as a whole.

If STRILEN has a value other than NO, and LCUGRP = NO is used, the number of groups applies to each string of devices irrespective of the LCU or Contention Group to which they belong.

If both STRILEN = NO and LCUGRP = NO are used, the number of groups applies to all devices under consideration in the given Xtract run.

3.4 Creating a single group

If "GROUPS=1" is specified, a single group is created. Any MXDVGP specification is ignored.

3.5 Creating two groups

If "GROUPS=2" is specified, two groups are created. Devices that exhibit an I/O rate of 5 I/Os per second or higher are allocated to the first group. Other devices are allocated to the second group. When this allocation process is completed, Xtract treats possible abnormal conditions:

If the first group contains no device (because all devices exhibited an I/O rate lower than 5 I/Os per second), then the device with the highest I/O rate is placed in the group.

If the second group contains no device (because all devices exhibited an I/O rate of 5 I/Os per second or higher), then the device with the lowest I/O rate is placed in the group.

If the number of devices in the first group exceeds the MXDVGP specification, the excess devices with the lowest I/O rate are moved to the second group.

If the number of devices in the second group exceeds the MXDVGP specification, the excess devices with the highest I/O rate are moved to the first group.

3.6 Creating more than two groups

If "GROUPS=n", where "n" has a value of 3 to MAXGP, is specified, then the device with the highest UTILIZATION (Service time multiplied by I/O rate) is placed in the first group. The device with the next highest utilization is placed in the second group, and so on, until group "n-1" is created. Finally, the remaining devices are allocated to the last group.

Once the allocation process completed, Xtract ensures that the number of devices in the last group does not exceed the MXDVGP specification. If this happens to be the case, the excess devices with the highest utilization are moved to group "n-1". Group "n-1" is then checked for the maximum number of devices, and the same processing is repeated until no group exceeds the MXDVGP specification.

3.7 Creating one group per device

If "GROUPS=0" is specified, Xtract will generate one group for each device. This allows the user to obtain a model at the level of individual devices. ***This is the recommended setting with mapped controllers such as the EMC Symmetrix series, the HDS 7700, 9900, IBM ESS, or STK Iceberg,***

Eclipse, V2X, etc. GROUPS=0 is the current Xtract default.

3.8 Incorrect number of groups

If the number of groups is higher than the number of devices, Xtract reduces the number of groups, and allocates a single device to each group.

If the specified number of groups is too small to accommodate all devices, because of the MXVDGP specification, Xtract increases the number of groups as necessary.

4. Examples of control statements

This Section presents examples of control statements. For each example, the control statements are shown and explained, and the processing performed by Xtract is outlined.

4.1 Example 1

```
LCU=14 GROUPS=4
```

Default operands are in effect for all the unspecified keywords. RMF data for all the DASD devices belonging to Logical Control Unit 14 will be processed. Devices will be allocated into four groups per string. Since a single LCU is processed, Xtract will place the results in a single output data set (OUT001). Xtract will attempt to dynamically allocate all the devices of LCU 14 to determine their type.

4.2 Example 2

```
ST1=0800, ND1=1145, ST2=1300, ND2=1730
```

```
DEV=(150-157)
```

```
DEV=(180,181,182,183,184,185,186,187)
```

```
GROUPS=5
```

Since the STRILEN keyword was not specified, the default value applies. Hence, devices 150 through 157 are taken to belong to a single string, and devices 180 through 187 to another string. Five groups will be allocated for each string. The four devices with the highest utilization will be placed in the first four groups, and the remaining four devices with the lowest utilization will form the fifth group. If both strings belong to the same LCU, a single output data set (OUT001) will be created.

On the other hand, if each string belongs to a different LCU, the two LCUs will be treated separately, and two output data sets (OUT001 and OUT002) will be created.

RMF records created between 8am and 11.45am, and between 1pm and 5.30pm will be processed. Xtract will attempt to dynamically allocate all the specified devices in order to determine their type.

4.3 Example 3

```
ST1=0800
```

```
DEV=(150-157)
```

```
DEV=(180-187)
```

```
GROUPS=2
```

LCUGRP=NO

Since the STRILEN keyword is not specified, the default value applies. Thus, two physical strings will be assumed corresponding to devices 150 through 157, and 180 through 187, respectively. Because "LCUGRP = NO" is specified, the LCU information is not used, and a single output data set (OUT001) will be created. Within each string, the 8 devices will be divided into two groups. GROUP01 will contain devices showing an I/O rate of 5 I/Os per second or higher. GROUP02 will contain devices showing an I/O rate lower than 5 per second. RMF records created before 8am will NOT be processed. Xtract will attempt to dynamically allocate all the devices to determine their type.

4.4 Example 4

GR1=(150-157)

GR1=(158,159,15A,15B)

GR2=(180-187)

ALLOC=NO

The 20 devices will be divided into two groups, as specified by the user. All RMF records, except the ones created on Saturdays and Sundays, will be processed. Xtract will not attempt to dynamically allocate the devices to determine their type, and a single output data set (OUT001) will be created.

4.5 Example 5

TUE=NO,WED=NO,THU=NO,FRI=NO

ST1=1000,ND1=1100

DEV=(160-16F)

DEV=(170-177)

DEV=(370-37F)

DEV=(380-38F)

DEV=(390-39F)

DEV=(870-877)

DEV=(8B0-8B7)

ALLOC=NO

GROUPS=3 LCUGRP=NO STRILEN=16

Because "LCUGRP = NO" is specified, the LCU information is not used, and a single output data set (OUT001) will be generated. Only RMF records created on Monday between 10 and 11am will be processed. Since STRILEN is set to 16, the specified devices are taken to form 7 strings, each divided into 3 groups. Within each string, groups 1 and 2 will contain a single device, and group 3 the remaining devices of the given string. Xtract will not attempt to dynamically allocate the devices.

4.6 Example 6

TUE=NO,WED=NO,THU=NO,FRI=NO

ST1=1000,ND1=1100

DEV=(160-16F)

DEV=(170-177)

DEV=(370-37F)

DEV=(380-38F)

DEV=(390-39F)

DEV=(870-877)

DEV=(8B0-8B7)

ALLOC=NO

GROUPS=9 LCUGRP=NO STRILEN=NO MXDVGP=32

Because LCUGRP=NO is specified, the LCU information is not used, and a single output data set (OUT001) is generated. Only RMF records created on Monday between 10 and 11 am will be processed. The 88 devices will be divided into 9 groups. No group may contain more than 32 devices. Thus, groups 1 through 6 will contain a single device, GROUP007 will contain 18 devices, and both GROUP008 and GROUP009 will contain 32 devices. Xtract will skip the dynamic device allocation.

4.7 Example 7

PROCESS=DASDMON

VOLREC=153,DSNREC=154

CONTGP=PATH01

GROUPS=4 STRILEN=NO

Default values are in effect for all the unspecified keywords. DASDMON (or ASTEX) data for all DASDs belonging to Contention Group PATH01 will be processed. The SMF record ID for "Volume" records is 153, and for "Data Set" records it is 154. Devices will be allocated into four groups. Since a single Contention Group is processed, Xtract will generate a single output data set (OUT001). If RMF type 74 records are present in the SMFIN data set, they will be used to obtain the "number of data sets open" for each device. Xtract will attempt to dynamically allocate all the devices in order to determine their type.

4.8 Example 8

PROCESS=ASTEX

SMFREC=153

LCU=0A3

As always, for all the unspecified parameters, default values are in effect. ASTEX (or DASDMON) data for all the devices belonging to LCU 0A3 will be processed. "Volume" and "Data Set" records have the same SMF record ID of 153. The devices will be allocated into four groups. Xtract will generate a single output data set (OUT001) since only one LCU is processed. RMF type 74 records MUST be present in the SMFIN data set to provide the LCU information. In addition, these records will be used to gather the "number of data sets open" for each device. If RMF type 73 and 78 records are also present in the SMFIN data set, they will be used to determine the channel configuration of LCU 0A3, as well as the channel utilizations. Xtract will try to dynamically allocate all the devices to determine their type.

4.9 Example 9

PROCESS=DASDMON

SMFREC=153

DEV=(150-157)

DEV=(180-187)

LCUGRP=CNTGP

GROUPS=4 STRILEN=NO

DASDMON data for the specified devices will be processed. There will be one output data set created for each Contention Group implied by the device numbers. For each Contention Group, devices will be allocated into four groups. Xtract will attempt to dynamically allocate all the devices to determine their type.

4.10 Example 10

PROCESS=CSIM

LCU=0A3 GROUPS=4

As always, default values are in effect for all unspecified parameters. CSIM data for all the DASD devices belonging to LCU 0A3 will be processed. Devices will be allocated into four groups. Since only one LCU is processed, Xtract will generate a single output data set (OUT001) and a single simulation data set (SIM001). To provide the LCU information, the SMFIN data set must contain at least type 74 RMF records. These records will also be used to gather the "number of data sets open", as well as pending, and IOSQ times for each device. If RMF type 73 and 78 records are also present in the SMFIN data set, they will be used to determine the channel configuration and load of LCU 0A3. Xtract will attempt to dynamically allocate all the devices to determine their type.

4.11 Example 11

PROCESS=CSIM

DEV=(150-157)

DEV=(180-187)

GROUPS=4 STRILEN=NO

CSIM data for the specified devices will be processed. Assuming that the SMFIN data set contains RMF type 74 records to identify LCUs, there will be one output data set and one SIM data set created per LCU implied by the device numbers. For each LCU, devices will be allocated into four groups. Xtract will attempt to dynamically allocate the devices in order to determine their type.

4.12 Example 12

LCU=(000-FFF) GROUPS=4

Default operands are in effect for all the unspecified keywords. RMF data for all DASD devices belonging to Logical Control Units 000 through FFF will be processed. Note that this is an effective way to select all the devices present in the RMF data. Devices will be allocated into four groups per string. One OUTnnn data set will be produced for each LCU in the input data set. Xtract will attempt to dynamically allocate all the devices in order to determine their type.

4.13 Example 13

PROCESS=GTF

LCU=0A3 GROUPS=4

Default operands are in effect for all the unspecified keywords. GTF data for all DASD devices belonging to Logical Control Unit 0A3 will be processed. Devices will be allocated into four groups per string. Since only one LCU is processed, Xtract will generate a single output data set (OUT001) and a single GTF data set (GTF001). To provide the LCU information, the SMFIN data set must contain at least type 74 RMF records. These records will also be used to gather the "number of data sets open" as well as pending and IOSQ times for each device group. If RMF type 73 and 78 records are also present in the SMFIN data set, they will be used to determine the channel configuration and load of LCU 0A3. Xtract will attempt to dynamically allocate all the devices at the specified LCU to determine their type, and to read their VTOCs in order to correlate I/O operations with individual data sets.

4.14 Example 14

PROCESS=GTF

DEV=(150-157)

DEV=(180-187)

STRILEN=NO

GTF data for the specified devices will be processed. Assuming that the SMFIN data set contains RMF type 74 records to identify LCUs, there will one output data set and one GTF data set created per LCU implied by the device numbers. For each LCU, devices will be allocated into individual groups (the default value for GROUPS). Xtract will attempt to dynamically allocate all the devices in order to determine their type.

4.15 Example 15

PROCESS=CARMF,SMFREC=248

SYSID=SYS1

LCU=14 GROUPS=4

As always, for all unspecified parameters, default values are in effect. Both standard RMF and Cache RMF Reporter data will be processed for all the DASD devices belonging to Logical Control Unit 14. Only records created on system "SYS1" will be taken into account. The SMF record ID for Cache RMF Reporter records is not the default 245, but 248 is used instead. Devices will be allocated into four groups per string. Since a single LCU is specified, Xtract will place the results in a single data set (OUT001). Xtract will attempt to dynamically allocate all the devices on LCU 14 to determine their type.

5. Using XtractTM

This section outlines the procedures to be used to install and run Xtract in an MVS/370, MVS/XA, MVS/ESA, z/OS, or VM/CMS Operating System Environment.

Xtract is delivered to the user in one of two forms:

- the object code is supplied on a floppy disk, CD ROM or other medium (referred to as distribution disk),
- load modules are supplied on a magnetic tape.

Sections 5.1 and 5.2 discuss the MVS environment, and section 5.3 is devoted to Xtract under VM/CMS.

5.1 Installing Xtract from disk under MVS

5.1.1 Files on the distribution disk

The distribution disk includes 5 main files, named XTRACT, XTRACT2, XTRACT3, XTRACT4, and XALLOC. Additional modules (some of them may be optional) are provided to treat data created by monitors other than RMF or compatible facilities. For Computer Associates ASTEX (or DASDMON), the additional file is named XTRACT2A. For IBM Corporation RMF Cache Reporter, the additional file is XTRACT2D. For Amdahl Corporation CSIM, the name of the file is XTRACT2E. For GTF, there are four additional files named XTRACT2B, XTRACT5, XVTOC, and XGTF. Other additional modules may be included as needed.

These binary files contain the IBM/390 object code for Xtract. They must be transferred "as is" (binary mode) to the mainframe computer, and link-edited to create the appropriate Xtract load modules. The following sections outline the procedure involved in performing this relatively simple setup.

Additionally, the distribution disk may contain files with sample JCLs for some of the steps. Files other than Xtract object modules are distinguished by the presence of a file extension (other than .OBJ). In some cases, the Xtract object modules may be supplied with the .OBJ file extension.

5.1.2 Creating the Object Library

The following JCL can be used to create the Object Library.

```
//CREATEA JOB .....  
//STEP EXEC PGM=IEFBR14  
//DD DD DSN=xtract.obj,UNIT=disk,VOL=SER=dskvol,  
// SPACE=(TRK,(10,5,20)),DISP=(NEW,CATLG),  
// DCB=(RECFM=FB,LRECL=80,BLKSIZE=400)  
//
```

Substitute the name you want to give to the new library for "xtract.obj", the volume serial number of an existing volume for "dskvol", and a unitname valid for your installation for "disk".

5.1.3 Creating the Load Library

The following JCL can be used to create the Load Library.

```
//CREATEL JOB .....
//STEP EXEC PGM=IEFBR14
//DD DD DSN=xtract.load,UNIT=disk,VOL=SER=dskvol,
// SPACE=(TRK,(30,10,20)),DISP=(NEW,CATLG),
// DCB=(RECFM=U,BLKSIZE=19069)
//
```

Substitute the name you want to give to the new library for "xtract.load", the volume serial number of an existing volume for "dskvol", and a unitname valid for your installation for "disk".

5.1.4 Sending Xtract objects to the mainframe

The Xtract object files on the distribution disk are simple binary files, and they must be transferred "AS IS" (without translation) to the mainframe. This can be accomplished using a modem and any of a number of standard communication software packages. Alternatively, a terminal emulation board can be used if the Personal Computer is connected to the mainframe through such a board. Refer, as needed, to the manual of the corresponding software for details on file transfer between the Personal Computer and the mainframe.

Since each of the object modules is to be stored as a separate member of your data set, you will probably find it necessary to explicitly create empty members before attempting the transfers. Some FTP programs automatically create the appropriate members in the object library as you upload each Xtract object module file.

5.1.5 Link-Editing Xtract

At this point, your OBJECT data set contains as many members as there were object module files on the distribution disk (5, or more if optional modules were included). They should be link-edited in your Load Library.

The following is an example of JCL that can be used to link-edit Xtract.

```
//LXTRACT JOB REGION=1024K
//LKED EXEC PGM=IEWL
//SYSLMOD DD DSN=xtract.load(xtract2),DISP=SHR
//SYSGO DD DSN=xtract.obj(xtract2),DISP=SHR
//SYSPRINT DD SYSOUT=A
```

The second member, named "xtract2", is treated in this JCL example. Lower case letters are used to denote parameter values which are installation dependent or which the user may want to rename.

5.2 Installing Xtract from magnetic tape under MVS

XTRACT, in the form of 5 load modules named XTRACT, XTRACT2, XTRACT3, XTRACT4, and XALLOC has been unloaded from a Partitioned Data Set (PDS) to a Standard Label tape, using the IBM IEBCOPY utility. Additional modules (some optional) may be included to process other monitors than standard RMF. For Computer Associates ASTEX, the load module is named XTRACT2A, and for Amdahl Corporation CSIM, the module name is XTRACT2E.

To reload the load modules, you can tailor a JOB like the ones described in the following sections, or you can develop your own JCL. Refer as needed to the "IBM OS/VS Utilities" manual (GC26-3092), or the "IBM MVS/XA Utilities Extended Architecture" manual (GC26-4018).

5.2.1 Loading Xtract in an existing library

The following JCL can be used to reload XTRACT in an existing library.

```
//RELOAD JOB REGION=2048K
//STEP EXEC PGM=IEBCOPY
//SYSPRINT DD SYSOUT=A
//TAPE DD UNIT=TAPE,DISP=OLD,VOL=SER=volser,DSN=tape.dsn
//DISK DD DISP=SHR,DSN=disk.dsn
//SYSIN DD *
COPY OUTDD=DISK,INDD=((TAPE,R))
//
```

See the cover letter sent with the product tape for the data set name and the volume serial number on this tape. Substitute them for "tape.dsn" and "volser" in the above example. Substitute the name of your library for "disk.dsn". Because the "replace" (R) option is specified to IEBCOPY, if members with the same name existed in the library, they will be replaced.

5.2.2 Loading Xtract in a new library

The following JCL can be used to create a new library and load Xtract in it.

```
//RELOAD JOB REGION=2048K
//STEP EXEC PGM=IEBCOPY
//SYSPRINT DD SYSOUT=A
//TAPE DD UNIT=TAPE,DISP=OLD,VOL=SER=volser,DSN=tape.dsn
//DISK DD DSN=disk.dsn,UNIT=disk,VOL=SER=dskvol,
// DISP=(NEW,CATALOG),
// SPACE=(TRK,(30,10,20)),DCB=(RECFM=U,BLKSIZE=19069)
//SYSIN DD *
COPY OUTDD=DISK,INDD=TAPE
//
```

See the cover letter sent with the product tape for the data set name and the volume serial number on this tape. Substitute them for "tape.dsn" and "volser" in the above example. Substitute the name you want to give to the new library for "disk.dsn", the volume serial number of an existing volume for "dskvol", and a unitname valid for your installation for "disk".

5.3 Xtract under VM/CMS

To install Xtract in a VM/CMS environment, the object files (no file extension or extension .OBJ) on the distribution disk must be transferred **AS IS** to the mainframe. Each of the object modules should be stored as a TEXT file. For example, XTRACT.OBJ should be stored as a VM file named XTRACT TEXT A1.

After transfer, you should have as many TEXT files as there were object module files on the distribution disk. The XTRACT module must now be generated using the following commands:

```
GLOBAL TXTLIB CMSLIB
LOAD XTRACT (CLEAR
GENMOD XTRACT (FROM XTRACT.
```

Note that GENMOD commands should not be used for the other TEXT files. The additional modules are dynamically invoked from the TEXT files during execution.

Under VM/CMS, Xtract must run on a 370 Virtual Machine (24 bit addressing).

A minimum of 4 FILEDEF commands are required:

- The SYSIN FILEDEF defines the control data set. It is fixed block (FB), logical record size 80 characters.
- The SYSPRINT FILEDEF defines the report listing. It is fixed blocked with ASA control characters (FBA), logical record size is 133 bytes.

- The SMFIN FILEDEF defines the input SMF data set, created by the SMF dump program.
- The OUT001 FILEDEF defines the output dataset destined for the Lead Time modelling software. It is fixed blocked (FB), logical record size is 80 bytes.

Depending on the data and treatment options, more than one output data set can be produced. In this case, the FILEDEFs for the additional data sets are OUT002, OUT003, etc up to OUT255.

- If PROCESS = CSIM is in effect, the UNITS FILEDEF defines the UNITS file created by CSIM V8R0 or subsequent releases. The DSNS FILEDEF defines the DSNS file created by CSIM.
- If PROCESS = MONITOR is in effect, the MONITOR FILEDEF defines the file created by the VM/CMS MONWRITE command.
- If PROCESS = GTF is in effect, the TRACE FILEDEF defines the trace data set created by GTF.
- The SIMnnn FILEDEF defines the data set of simulation results corresponding to a given OUTnnn output data set.

When referring to MVS examples in the manual, typically each "DD statement" should be replaced by a FILEDEF command.

Example of Commands:

```
FILEDEF SYSIN DISK XTRACT DATA A1
FILEDEF MONITOR TAP1 NL
FILEDEF SYSPRINT PR
FILEDEF OUT001 DISK LEADTIME VLCU1 A1 (BLKSIZE 80)
FILEDEF OUT002 DISK LEADTIME VLCU2 A1 (BLKSIZE 80)
```

Here, SYSIN control statements are in a CMS file XTRACT DATA A1. The MONITOR data is on unlabeled tape reel. The SYSPRINT listing is directed to the Virtual Machine Printer. Two output files are expected. These will be created as CMS files LEADTIME VLCU1 A1 and LEADTIME VLCU2 A1, respectively.

5.4 Collecting the SMF data

Xtract requires an input file containing SMF records type 74 (RMF DASD Activity), as well as, optionally, type 73 (RMF Channel Path Activity), and type 78 (RMF Monitor I Activity), and/or other monitor records (e.g., ASTEX data set and volume records) created by the SMF dump (IFASMFDP) program. XTRACT will automatically discard other SMF record types. Note that Xtract requires SMF records created on an MVS/XA or MVS/ESA system. MVS/370 records are NOT supported, and should NOT be present in the file.

Refer to "MVS/XA RMF Reference & User's Guide (LC28-1138)" for the RMF parameters to be specified to create type 73, 74, and 78 SMF records.

Refer to "MVS/XA Systems Programming Library - SMF (GC28-1153)" for information about the SMF

dump program.

Refer to "Cache RMF Reporter, Program Description and Operation Manual (SH20-6295) for information about the RMF Cache Reporter.

Refer to your documentation for information about other monitors such as DASDMON or CSIM.

5.5 Collecting the GTF data

For GTF trace treatment, Xtract requires CCW trace records collected on an MVS/XA or MVS/ESA system. The data should be collected over a representative period of time to avoid meaningless results.

GTF, the IBM service aid, is described in "MVS/XA SPL Service Aids (GC28-1159)". CCW records must be collected, although only for I/O interrupts (and not for SSCH). "TIME=YES" must be specified in the PARM field of the EXEC statement used to start GTF, or Xtract will issue an error message and terminate.

The following should be specified to collect GTF/XA or GTF/ESA trace data.

```
TRACE=SSCHP,IOP,CCWP,PCI
```

```
SSCH=IO=(aaa,bbb,ccc,ddd,...)
```

```
SSCH=IO=(xxx,yyy,zzz)
```

```
CCW=(I,CCWN=255,DATA=16,PCITAB=1)
```

Substitute device numbers of the volumes you wish to trace for aaa, bbb, etc.

Note that the operand of the DATA keyword must be set to at least 16 in order for Xtract to operate correctly.

5.6 Typical Job Control Statements

This section explains the Job streams that can be used to run Xtract.

XTRACT requires a minimum of 4 DD statements:

- The SYSIN DD defines the control data set. Its record format is fixed blocked (FB), and its logical record size is 80 bytes.
- The SYSPRINT DD statement defines the report listing. Its record format is fixed, blocked, with ASA control characters (FBA) and its logical record size is 133 bytes.
- The SMFIN DD statement defines the input SMF data set, created by the SMF dump program.
- The OUT001 DD statement defines the output dataset containing the information destined for the Lead Time product. Its format is fixed blocked (FB) and its logical record size is 80 bytes. Depending on the options specified by the user, and on the contents of the SMF records processed, more than one output data sets may be produced. In this case, the DDNAME for the additional data sets are OUT002, OUT003, and so on, up to OUT255.
- If PROCESS = CSIM is in effect, the UNITS DD statement defines the UNITS file created by CSIM V8R0. The DSNS DD statement defines the DSNS file created by CSIM. This latter statement is optional. The SIMnnn DD statement defines the data set of simulation results corresponding to a given OUTnnn output data set.
- If PROCESS = GTF is in effect, the TRACE DD statement defines the trace data set created by GTF. The GTFnnn DD statement defines the condensed GTF data set produced by Xtract corresponding to a given OUTnnn output data set.

The following guidelines may be used to determine the number of output data sets created by Xtract:

Xtract creates a single output data set (OUT001) if :

- a single LCU or Contention Group is specified either explicitly (LCU or CONTGP keywords) or implicitly (DEVICES keywords)
- LCUGRP = NO is specified
- the GRn keyword is used

Xtract creates several output data sets if :

- LCUGRP = YES is in effect. One data set is created for each individual Logical Control Unit. Also, if some of the specified devices were not found in the SMF data, they are assumed to belong to LCU 000, and are grouped in an additional data set.
- LCUGRP = CNTGP is in effect. One data set is created for each individual Contention Group.
- PROCESS = CSIM is in effect. Associated with each output data set OUTnnn, a SIMnnn DD statement should be specified for the corresponding data set of simulation results.

Note that if one or more OUTnnn (or SIMnnn) DD statements are missing from the JCL, depending on the situation, Xtract will not terminate abnormally. Rather, it will write a warning on SYSPRINT, and use the available data sets.

5.6.1 Examples of JCL

```
//XTRACT JOB REGION=1024K
// EXEC PGM=XTRACT
//STEPLIB DD DSN=XTRACT.LOADLIB,DISP=SHR
//SMFIN DD DISP=SHR,DSN=SYS1.SMF.DUMP
//OUT001 DD DISP=SHR,DSN=LEADTIME.DATA.LCU001
//OUT002 DD DISP=SHR,DSN=LEADTIME.DATA.LCU002
//OUT003 DD DISP=SHR,DSN=LEADTIME.DATA.LCU003
//OUT004 DD DISP=SHR,DSN=LEADTIME.DATA.LCU004
//OUT005 DD DISP=SHR,DSN=LEADTIME.DATA.LCU005
//OUT006 DD DISP=SHR,DSN=LEADTIME.DATA.LCU006
//OUT007 DD DISP=SHR,DSN=LEADTIME.DATA.LCU007
//OUT008 DD DISP=SHR,DSN=LEADTIME.DATA.LCU000
//SYSPRINT DD SYSOUT=A
//SYSIN DD *
```

```
ST1=0435,ND1=0435
DEV=(160-16F)
DEV=(170-177)
DEV=(370-37F)
DEV=(380-38F)
DEV=(390-39F)
DEV=(397-39F)
DEV=(870-877)
DEV=(8B0-8B7)
GROUPS=3
//
```

We specified a 1024K region size. This is sufficient for most Xtract runs. Note that if Xtract runs under MVS/XA, MVS/ESA or z/OS, it acquires most of the virtual storage needed for internal tables above the 16 MB line. We assumed that XTRACT had been installed in a cataloged library named XTRACT.LOADLIB.

We show below another example of JCL which illustrates the use of Xtract with the CSIM tool from Amdahl Corporation.

```
//XTRACT JOB REGION=4096K
//STEP1 EXEC PGM=XTRACT
//STEPLIB DD DSN=xtract.load,DISP=SHR
//SYSPRINT DD SYSOUT=A
//UNITS DD DSN=csim8.units,DISP=SHR
//DSNS DD DSN=csim8.dsns,DISP=SHR
//OUT001 DD DISP=SHR,DSN=leadtime.out001
//SIM001 DD DISP=SHR,DSN=leadtime.sim001
//SYSIN DD *
```

```
PROCESS=CSIM
DEV=(383-38A)
DEV=(B80-B87)
MAXDSX=2000
//
```

A larger region of 4096K is specified here to accommodate the treatment of data set level information with potentially a larger number of individual data sets to be processed (MAXDSX).

To treat GTF records, specify PROCESS = GTF. In the JCL, a new TRACE DD statement defines the trace data set created by GTF to be used in the given Xtract run. A WORK DD statement is also needed. An example of JCL follows.

```
//XTRACT JOB REGION=8192K
// EXEC PGM=XTRACT
//STEPLIB DD DSN=XTRACT.LOADLIB,DISP=SHR
//SMFIN DD DISP=SHR,DSN=SYS1.SMF.DUMP
//TRACE DD DISP=SHR,DSN=SYS1.TRACE
//OUT001 DD DISP=SHR,DSN=LEADTIME.DATA
//GTF001 DD DISP=SHR,DSN=LEADTIME.GTF
//WORK DD DSN=&&WK,DISP=(,PASS),SPACE=(CYL,(100))
//SYSPRINT DD SYSOUT=A
//SYSIN DD *
PROCESS=GTF

DEV=(383-38A,B80-B87)
MAXDSX=32000
//
```

A larger region of 8192K is specified to accommodate the treatment of data set level information with a potentially larger number of individual data sets to be processed (MAXDSX).

You will find additional information concerning the proper GTF settings to collect a trace for Xtract treatment in section 5.4.

If your input data set of RMF records contains records pertaining to several systems, you can use the keyword SYSID = SMF_System_Id to select records for the desired system.

When processing multiple GTF traces, the system IDs of the systems to be considered must be specified using the TRACEn statement. A WORK DD statement is also needed. We give below an example of JCL

```
// JOB
//          EXEC PGM=XTRACT
//STEPLIB  DD DSN=MY.LIBRARY,DISP=SHR
//SMFIN    DD DSN=SMF.FROM.SYS1,DISP=SHR
//          DD DSN=SMF.FROM.MVS4,DISP=SHR
//          DD DSN=SMF.FROM.MVS7,DISP=SHR
//TRACE1   DD DSN=GTF.FROM.SYS1,DISP=SHR
//TRACE2   DD DSN=GTF.FROM.MVS4,DISP=SHR
//TRACE3   DD DSN=GTF.FROM.MVS7,DISP=SHR
//OUT001   DD DISP=SHR,DSN=SMF.OUTPUT
//GTF001   DD DISP=SHR,DSN=GTF.OUTPUT
//SORTWK01 DD UNIT=SYSDA,SPACE=(CYL,(20))
//SORTWK02 DD UNIT=SYSDA,SPACE=(CYL,(20))
//SORTWK03 DD UNIT=SYSDA,SPACE=(CYL,(20))
//WORK     DD DSN=&&WK,DISP=(,PASS),SPACE=(CYL,(100))
//SYSPRINT DD SYSOUT=A
//SYSUDUMP DD SYSOUT=A
//SYSOUT   DD SYSOUT=A
//SYSIN    DD *
TRACE1 = SYS1
TRACE2 = MVS4
TRACE3 = MVS7
LCUGRP=NO,NOMATCH=30
PROCESS=GTF,ALLOC=NO
DEVICES=(600-6FF)
//
```

Xtract™ Manual Addendum

This addendum to the Xtract manual describes recent enhancements and changes introduced in Xtract.

Before describing some of the additions to the product, we start by a brief discussion of **frequently encountered problems** when using Xtract.

I. By default, Xtract does not process data generated on Saturdays and Sundays. To change the default, you must specify SATURDAY=YES and/or SUNDAY=YES.

II. When processing GTF information, if you want to work at volume level only, do not specify MAXDSX=0 as it will likely cause Xtract to fail. Use ALLOC=NO instead.

III. A WORK DD statement is needed when using PROCESS=GTF. You need to specify DSN=..., VOLSER=..., UNIT=..., DCB=BLKSIZE=.... Good block sizes are 23476 for a 3380 or 27998 for a 3390.

IV. Using the keywords ST1, ND1, ST2, ND2, Xtract lets you specify the RMF intervals you want to process. For instance, ST1=0900, ND1=1059, ST2=1400, ND2=1459 specifies that you want Xtract to process RMF intervals starting between 9.00 am and 10.59 am, as well as those starting between 2.00 pm and 2.59 pm.

The SMF dump program, IFASMFDP also allows the user to specify periods to be selected. There is one important difference, however. In the case of the SMF dump program, the time refers to the moment the record was written to the SMF data set.

A misunderstanding of this fact may lead to frustrating consequences as in the following example:

RMF intervals are 1 hour long and start on the hour.

The user wants to examine the activity between 2 pm and 3 pm, and specifies this to IFASMFDP.

IFASMFDP dumps all records that were written to the SMF data sets between 2 pm and 3 pm. RMF records written between 2 pm and 3 pm (more exactly, a few seconds after 2 pm) are the ones describing the RMF interval ending at 2 pm, i.e., the 1 pm to 2 pm interval.

Xtract is executed with ST1=1400 and ND1=1459 and does not find any applicable RMF records.

In summary, it is important to remember that Xtract processes all RMF intervals that started during the specified period (or periods), while IFASMFDP dumps the records which were written during the specified period.

We now briefly describe recent Xtract enhancements.

Xtgen JCL generator for Xtract

Xtgen is a utility for the data extraction component of our Lead Time I/O performance modeling software. Its goal is to automate the production of the Job Control Language (JCL) specification for Xtract, the mainframe companion tool for Lead Time. Xtgen allows you to enter through a few simple dialogs the information about the desired data extraction, such as type of controllers, LCUs and device ranges to be included, etc. Based on this information, and additional Operating System data (such as dataset names, account information, etc), Xtgen generates the correct JCL in an ASCII file. All that remains for you to do is to transfer this file to the mainframe as a character file and to use it as the JCL for the desired data extraction with Xtract.

General job parameters dialog

The parameters specified here include:

Record types to process: specifies what record types, in addition to standard RMF, should be processed. For IXFP records, you need also to tell Xtgen the SMF record number used for these records.

Starting date and ending date: dates are specified in the format year (a four digit number) and day of the year (1-366).

Time periods: up to two time periods can be specified. Time is specified as hour (0-23) and minute (0-59).

Exclude records created on: allows to selectively exclude records created on certain days of the week (e.g., weekend).

Select records for a specific system: allows to specify the System Id of the system image to be processed in cases when the input dataset contains records from several system images.

System parameters dialog

The system parameters dialog groups several parameters needed to generate the appropriate Job Control statements for your job, as well as the name of the .jcl file that will contain the JCL generated.

.jcl file to produce: to specify the name of the output file that will be the result of the Xtgen generate command.

Job name: the name of the mainframe job (e.g. T9999) for the JOB statement.

Account information: information for the JOB statement (e.g. (1234567,AB)).

User id: for the JOB statement.

Message class: for the MSGCLASS clause of the JOB statement.

Region: for the REGION clause, in Mbytes.

CPU limits: for the TIME clause of the JOB statement.

Notify user id: for the NOTIFY clause of the JOB statement.

Xtract library dataset name: the name of Xtract load library.

SMIFIN dataset name: the name of the dataset containing the SMF records to be processed.

Use dataset for output: specifies that SYSPRINT will be directed to the corresponding dataset. You need to indicate the format of the corresponding storage unit (3390 or 3380).

Use SYSOUT: use standard system output unit for SYSPRINT.

Controllers dialog

The controllers to be processed are identified through their labels. To include a given controller in the data extraction, check the corresponding check box. To edit the controller properties, click the push button with the controller label.

Controller parameters include:

Name: the name of the Xtract output dataset on the mainframe side corresponding to the given controller.

You can also specify if the given dataset already exists or is a new dataset, as well as the format of the storage device unit (3390 or 3380).

Type: the controller type, to select from a list of controller types recognized by Lead Time.

ID: the ID label that identifies the controller.

Total LCUs: the total number of Logical Control Units defined for this controller that you wish to process through Xtract.

Devices to be processed: you can request that all devices for all LCUs specified by you for this controller be included or you can specify individual device ranges for each LCU included in the data extraction for this controller.

LCU parameters include:

LCU number: the LCU number in hexadecimal (as known to the Operating System) of the given LCU.

Number of device ranges: this parameter is only enabled if you elect to specify the device ranges for each LCU.

Device ranges: these parameters, expressed as hexadecimal number, are only enabled if you elect to specify device ranges for each LCU.

Processing multiple control units in a single Xtract run

The purpose is to handle more easily modern storage controllers that include multiple LCUs and larger configurations with multiple controllers of interest. This is done using the keyword **CNTRL**, and the 3 additional sub-keywords **TYPE** (to define the type of controller), **CLCUS** (to define which LCUs are part of this controller), and **ID** (explained further). This way, there is no need to "merge" multiple LCUs into Lead Time, the entire controller is generated as a single file. For instance:

```
CNTRL TYPE=STKECLIPSE,CLCUS=(27,28,29,2A),ID=STK1
```

Notes:

- Both TYPE and CLCUS keywords are mandatory, and must be specified for each controller. ID is optional.
- The operand of TYPE is up to 16 alphanumeric characters as specified below.
- CLCUS may be specified multiple times, as needed.
- CNTRL is incompatible with LCUGRP=NO, LCUGRP=CONTGP, and user defined groups (i.e. GROUP0=... GROUP1=...)
- A maximum of 256 CNTRL statements may be specified in one XTRACT run.
- A maximum of 256 LCUs may be specified for each controller.
- CLCUS defines the relation between a controller and LCUs, it does NOT specify which LCUs must be handled. The devices to be handled by XTRACT must be specified like in previous releases, using either the LCU keyword or the DEVICE keyword.
- XTRACT will generate one output data set for each controller **and** one output data set for each LCU NOT part of a controller (if any).
- The ID keyword is intended as help for the user to identify the controller. The operand of the ID keyword is up to eight characters, and will be used as DDNAME to write Lead Time parameters for the controller. Consequently, it must follow the same rules as a JCL language DDNAME.
- Controller TYPE values currently supported by Lead Time are:

```
CCUM6          /* cached IBM 3990 Mod 6 */
CCUM3          /* cached IBM 3990 Mod 3 */
```

```

ACM6          /* Amdahl 6690 */
AFCU          /* cached Amdahl 6100 */
CPXM6        /* CPX 6090-6 */
HDSM3        /* cached HDS 7980-3 */
CCUM23       /* cached IBM Mod 23 */
CCUA2        /* cached Amdahl 6880 A2 */
CCUG2        /* cached Amdahl 6880 G2 */
CCUN3C       /* NAS 7880-3C */
CCUSYB       /* cached STC Sybercache */
RCU4         /* uncached 4 paths */
AFCUNC       /* AC 6100 without cache */
RCU          /* uncached CU 2 paths */
SSD          /* solid state disk */
IBM41        /* nonsync IBM 9341 */
IBM43        /* nonsync IBM 9343 */
VKG          /* VIKING */
DMX800       /* EMC DMX800 */
DMX1000      /* EMC DMX1000 */
DMX2000      /* EMC DMX2000 */
DMX3000      /* EMC DMX3000 */
DMX3         /* EMC DMX-3 */
SYM823       /* EMC Symmetrix 8230 */
SYM853       /* EMC Symmetrix 8530 */
SYM883       /* EMC Symmetrix 8830 */
SYM843       /* EMC Symmetrix 8430 */
SYM873       /* EMC Symmetrix 8730 */
SYM570       /* EMC Symmetrix 5700 */
SYM543       /* EMC Symmetrix 5430 */
SYM552       /* EMC Symmetrix 5500-2 */
SYM553       /* EMC Symmetrix 5500-3 */
SYM559       /* EMC Symmetrix 5500-9 */
SYM523       /* EMC Symmetrix 5200-3 */
SYM529       /* EMC Symmetrix 5200-9 */
SYM482       /* EMC Symmetrix 4800-2 */
ICBRG        /* STK Iceberg */
CPX2000      /* CPX Tetragon 2000, HDS 7700 */
HDS7700E     /* HDS 7700-E */
HDS9900      /* HDS 9900 */
HD9900V      /* HDS 9900 V */
HD9900U      /* HDS 9900 U: upgraded version of HDS 9900 */
RAMACAS      /* RAMAC Array Subsystem */
RAMACDA      /* RAMAC DASD Array behind 3990 Mod 6 */
KODIAK       /* STK Kodiak */
RAMELAR      /* RAMAC Electronic Array */
RAMVIRT      /* RAMAC Virtual Array */
INFISP40     /* Encore Infinity SP40 */
STK934       /* STK 9393-4 */
STK932A      /* STK 9393-2 */
STK936       /* STK 9393-6 */
STKECLIPSE   /* STK Eclipse */
STKECLIPS2   /* STK Eclipse-2 */
STKV960      /* STK V960 */
STKV2X9      /* STK V2X */
STKV2XV2     /* STK V2X V2 */

```

```
STKV2XF          /* STK V2X F */
IBM2105          /* IBM Shark */
IBM2105F         /* IBM Shark 1.5 */
IBM800           /* IBM ESS Model 800 */
IBM800T         /* IBM ESS Model 800 Turbo */
```

Example of an Xtract run in CNTRL mode:

```
CNTRL TYPE=STKECLIPSE,CLCUS=(15,16,17,18),ID=STK1
CNTRL TYPE=SYM843,CLCUS=(110,111),ID=EMC1
CNTRL TYPE=HDS9900,CLCUS=(25,26),ID=HDS1
PROCESS=IXFP
SMFREC=250
DEVICES=(0500-05ff)
DEVICES=(7000-733F)
DEVICES=(9100-713F)
DEVICES=(4000-423F)
DEVICES=(a300-a3fF)
ALLOC=NO
ST1=2319,ND1=2323
LN=55
```

In the above run, three "controllers" are defined. DDNAMES STK1, EMC1, and HDS1 must be present in the JCL. Additional DDNAMES OUT001, OUT002... and so on may be needed for each LCU that doesn't belong to any of the three controllers. Note that it would be a good practice to use a CNTRL statement (with an ID keyword) for each physical DASD controller, even if it only includes a single LCU. This is especially true in a multi-system environment.

Support for IXFP SMF records generated for Storage Technology controllers

Xtract will process IXFP SMF records if **PROCESS=IXFP** is specified. You must also tell Xtract the SMF record number through the **SMFREC** keyword. Note that the IXFP keyword implies also PROCESS=RMF so that RMF records are processed accordingly, including type 74 subtype 5 records. Example:

```
PROCESS=IXFP
SMFREC=250
```

Support for RMF type 74 subtype 5 records

Xtract now supports RMF type 74 subtype 5 (replacing cache RMF reporter records). Xtract will automatically process 74-5 records if PROCESS=RMF is specified explicitly or by default. If incompatible processing types are specified (DASDMON, ASTEX, GTF, CARMF, or CSIM), 74-5 records are ignored.

The **maximum number of devices** and groups has been **increased** from 2048 to **9999**.

This change could cause problems in memory constrained environments, because Xtract needs to GETMAIN 560 bytes of memory for each device. Thus, the following approach was adopted:

- 1) The **default** number of devices and groups still is **2048**.
- 2) The **maximum** number of devices and groups is now specified in the JCL PARM field. For example:
`//STEP1 EXEC PGM=XTRACT,PARM=4096`
will allow for a maximum of 4096 devices and groups.
- 3) If a JCL PARM field is not specified, the default of 2048 applies.
- 4) In memory constrained environments, it is possible to decrease the default, also by using the JCL PARM field. For example:
`//STEP1 EXEC PGM=XTRACT,PARM=256`

Note that the current release of Lead Time supports up to **16384** logical volumes for a single control unit.

Support for merging of multiple GTF traces

Xtract now supports the merging of **multiple GTF traces**. These traces may have been collected on the same or on different systems depending on the purpose of the merging. Typically, merged GTF traces are of interest in reconfiguration studies where the workloads from several smaller control units are ported to a single larger controller. Another typical application are Shared DASD configurations where several systems contribute to the total workload on the cached controller.

Other older changes and additions

The other main additions include support for records created by

1. the Cache RMF Reporter utility (IBM Corporation),
2. ASTEX (Computer Associates)
3. the General Trace Facility (GTF, IBM Corporation).
4. the MONITOR command of VM/XA SP2 or VM/ESA.

Additionally, this release of Xtract collects information from **SMF records type 42 subtype 6** (DFSMS data set statistics). It includes two new modules: XTRACT2F and XVTOC2.

As before, Xtract collects information from SMF records type 73, type 74 (subtype 1), and 78 (subtypes 1 and 3). In addition, dataset information is collected from type 42 (subtype 6). Since SMF records 42-6 do not contain device or volume serial numbers, Xtracts reads the VTOCs of the target volumes to determine where the datasets are located.

For best results with type 42 records, it is recommended to specify interval accounting (see the SMF manual GC28-1628), and to synchronize the SMF and RMF intervals.

To use SMF records 42 subtype 6 with Xtract:

1. specify **PROCESS = SMF42**
2. run Xtract on a system that has access to the target devices
3. if VTOCs are under a protection mechanism (ACF2, RACF), give the Xtract job the necessary authorization.

Processing of SMF records created by the [Cache RMF Reporter](#) is very similar to that of standard RMF records. In most cases, all you need to do is to specify PROCESS = CARMF. If you are not using the standard SMF record number (245), you must also tell Xtract the SMF record number through the SMFREC keyword. The advantage of Cache RMF records over standard RMF is, of course, the availability of additional information such as percentages of reads and hit ratios.

ASTEX is the successor of DASDMON. To process SMF records created by ASTEX, you can indifferently use PROCESS = ASTEX or PROCESS = DASDMON since Xtract automatically distinguishes between records created by DASDMON and by ASTEX. Starting with version 3.9, Xtract supports ASTEX Releases 2.1 and 2.5.

Note that Xtract expects the data to be in the format of SMF records and not that of ASTEX or DASDMON database. Use the appropriate utility provided by Computer Associates to create SMF records.

To specify the record code (SMFREC#) for ASTEX SMF "Volume Records" and SMF "Data Set Records", use the new keywords VOLREC = number and DSNREC = number, respectively. If the same record ID is used for both "Volume" and "Data Set" records, you can specify it via the SMFREC keyword.

Records created by the Generalized Trace Facility (GTF) present a more "microscopic" view of system performance, and thus lend themselves to additional treatment such as assessing the expected impact of a given cache or NVS size, or selected aspects of System Managed Storage. When used with the Volume Tables of Contents (VTOCs) for the specified devices, the GTF trace enables you to obtain monitor information at the level of individual data sets.

As a result of processing the GTF trace, Xtract produces two types of output data sets. The first type, OUTnnn, is similar to the data set created by an Xtract run with RMF or another monitor data. It describes the I/O activity as recorded in the GTF trace. This dataset should be transferred to the Personal Computer as an ASCII file with the extension .SMF, and read in through the Lead Time Fetch File facility.

The second type, GTFnnn, contains the GTF trace itself in a condensed binary form. This data set should be transferred to the Personal Computer "as is" (i.e., without any character translation), and given the extension .GTF.

GTFSIM is a utility included with Lead Time to allow you to process this condensed trace under a number of user specified conditions. You can thus obtain cache hit ratios and other cache performance data as a function of the cache and NVS sizes, as well as the treatment of Define Extent flags. For a non cached environment, estimates of device seek times are produced. The GTFSIM utility creates an output file in the .MDL format readable by the Lead Time Retrieve File facility.

If RMF data is present in the input data set, Xtract uses RMF records (type 73, 74, 78) to gather Logical Control Unit and other information. If LCU based grouping is specified, one pair of OUTnnn and GTFnnn data sets is produced per LCU.

To determine the control unit and device types, Xtract dynamically allocates them to issue the SENSE ID command. In conformance with MVS rules, Xtract attempts to open the VTOC data set on each device to be sensed. If the access to VTOCs is restricted by products such as RACF or ACF2, either Xtract should be given the necessary authority or the command ALLOC = NO should be used to suppress dynamic device allocation.

Xtract now supports device and control unit type information available in RMF release 4. If this information is available in the records processed, Xtract will not try to allocate the devices to determine their type. Note, however, that VTOC information is read from devices if GTF data is being processed and "ALLOC=YES" is in effect.

Note that in order to obtain data set level information with a GTF trace, Xtract must be able to read the VTOCs of the devices under consideration, i.e., ALLOC = YES must be in effect.

To treat GTF records, specify PROCESS = GTF. In the JCL, a new TRACE DD statement defines the trace data set created by GTF to be used in the given Xtract run. A WORK DD is required to define a work file used during processing. The size of this work file depends on on the amount of GTF data: 20 3390 cylinder should suffice in most cases. An example of JCL follows.

```

//XTRACT JOB REGION=8192K
// EXEC PGM=XTRACT
//STEPLIB DD DSN=XTRACT.LOADLIB,DISP=SHR
//SMFIN DD DISP=SHR,DSN=SYS1.SMF.DUMP
//TRACE DD DISP=SHR,DSN=SYS1.TRACE
//OUT001 DD DISP=SHR,DSN=LEADTIME.DATA
//GTF001 DD DISP=SHR,DSN=LEADTIME.GTF
//SYSPRINT DD SYSOUT=A
//WORK DD DSN=&&WK,DISP=(,PASS),SPACE=(CYL,(100))
//SYSIN DD *
PROCESS=GTF
DEV=(383-38A,B80-B87)
MAXDSX=32000
//

```

A larger region of 8192K is specified to accommodate the treatment of data set level information with a potentially larger number of individual data sets to be processed (MAXDSX).

You will find additional information concerning the proper GTF settings to collect a trace for Xtract treatment in section 5.4. If your input data set of RMF or Cache RMF records contains records pertaining to several systems, you can use the new keyword **SYSID = SMF_System_Id** to select records for the desired system. If you are a CSIM user (Amdahl Corporation), you may be interested to know that Xtract now supports CSIM Version 12, in addition to Versions 9 and 8.

Xtract supports records generated by the MONITOR command of VM/XA SP2 or VM/ESA. Xtract can run under MVS/370, MVS/XA, MVS/ESA, z/OS or VM/CMS. The data processed can have been created under another Operating System (e.g. MONITOR data can be processed under MVS/ESA, or RMF data under VM/CMS). Xtract processes the following types of MONITOR records:

- Domain 0, Record 9 (Channel Connection Data)
- Domain 1, Record 6 (Configuration Data)
- Domain 6, Record 1 (Vary Online)
- Domain 6, record 3 (Device Activity)
- Domain 6, Record 4 (Cache Activity)
- Domain 6, Record 5 (Attach Device).

The MONITOR file must have been created by the VM/CMS "MONWRITE" command.

Xtract requirements match those of the IBM "VMPRF" program to produce the following reports:

- Cache_DASD_by_Activity
- Channel_Busy
- DASD_by_Activity.

If needed, refer to VMPRF User's Guide and Reference (SC23-0460-02). Because MONITOR data does not contain Logical Control Unit (LCU) identification, Xtract creates a virtual LCU for devices accessed via the same Channel Paths, and showing the same type of Control Unit.

If Xtract runs under VM/CMS, "ALLOC=NO" is forced, i.e., the devices are not allocated, and VTOCs are not read. To treat data produced by the MONITOR command under VM/XA SP2 or VM/ESA, specify **PROCESS = MONITOR**. Refer to Section 5.3 of the manual for information on installing and using Xtract in a VM/CMS environment.

Appendix

A.1 Xtract condition codes

If there is a coding error in the control data set, or a processing error during execution, Xtract is likely to terminate with a final condition code that may help you resolve the error. In most cases, before terminating, Xtract will print an error message on SYSPRINT.

- The following is a list of Xtract condition codes.
- Return code 00 (X'00') : successful execution.
- Return code 02 (X'02') : an overflow condition was detected for a device while cumulating data into tables. Xtract stops processing SMF records as if end-of-file had been encountered. The device number of the device for which the overflow occurred is noted on SYSPRINT, and a second message specifies the amount of SMF data processed (duration). Then, Xtract proceeds with grouping, reporting, and creating output data set(s). If the overflow occurs when processing RMF or compatible data, the results produced by Xtract may be perfectly valid. With DASDMON data, it is recommended to reduce the period for which data is processed, and rerun Xtract.
- Return code 04 (X'04') : see the message on SYSPRINT for an explanation of the error. Either no SMF records were found on SMFIN for the specified period, or there were no records for the specified devices.
- Return code 08 (X'08') : see the message on SYSPRINT for an explanation of the error. Either an error has been detected in the SYSIN control file, or the SMFIN or OUTPUT data sets cannot be opened.
- Return code 16 (X'10') : the SYSIN or the SYSPRINT data set cannot be opened.

If Xtract detects an internal logical error, it may end abnormally with an ABEND code identifying the module in error. If this happens, it is recommended to rerun Xtract adding a SYSUDUMP DD statement, and contact your PALLAS representative for resolution of the problem.

A.2 Report description

Xtract may produce 4 reports: input control statements, DASD activity extracted from RMF or other volume level records (e.g. DASDMON volume records), data set activity (e.g. from DASDMON data set records), and group activity corresponding to the result of Xtract processing.

Control statements from SYSIN

This report is a record-by-record (card-by-card) listing of the control data set, just as it was read from SYSIN. It is the only report produced if there is an error.

If a syntax error or an inconsistency is discovered in the control data set, an error message is printed immediately after the line containing the error. An example might be

```
ERROR -- INVALID DEVICE ADDRESS  
SEE THE USER GUIDE FOR REQUIREMENTS AND RESTRICTIONS
```

Once all control statements have been processed, if an error precludes the successful execution of Xtract, the following message is printed

SYSIN PROCESSING COMPLETED. PROGRAM CANNOT CONTINUE DUE TO ERRORS

and Xtract terminates with condition code 8 (X'08').

DASD activity

This report shows, for each selected device, the group number and the collected RMF, DASDMON, or CSIM data.

Data set activity

This report shows, for data sets allocated on the selected devices, the group number as well as the data collected by the monitor.

Group activity

This report shows, for each group, the number of devices and cumulated RMF, DASDMON, or CSIM data.

1. Introduction	2
1.1 About Xtract	2
1.2 Processing in Xtract	2
1.2.1 Processing RMF and compatible records	3
1.2.2 Processing ASTEX TM records	3
1.2.3. Processing ASTEX with RMF	4
1.2.4 Processing CSIM records	4
1.2.5 Processing both CSIM and RMF records	5
1.2.6 Processing GTF records	5
1.2.7 Processing GTF and RMF records	6
1.2.8 Processing RMF records and multiple GTF traces	6
1.2.9 Processing MONITOR records	7
1.2.10 Dynamic device allocation and VTOCs	8
1.3 Grouping techniques	8
1.4 What follows	9
2. The Xtract TM control data set	9
2.1 Description and syntax	9
2.2 Selection parameters	10
2.3 Device and group definitions	15
2.4 Miscellaneous definitions	18
3. Creation of groups	20
3.1 General information	20
3.2 Group naming conventions	20
3.3 Scope of grouping	20
3.4 Creating a single group	21
3.5 Creating two groups	21
3.6 Creating more than two groups	21
3.7 Creating one group per device	21
3.8 Incorrect number of groups	22
4. Examples of control statements	23
4.1 Example 1	23
4.2 Example 2	23
4.3 Example 3	23
4.4 Example 4	24
4.5 Example 5	24
4.6 Example 6	25
4.7 Example 7	25
4.8 Example 8	26
4.9 Example 9	26
4.10 Example 10	27
4.11 Example 11	27
4.12 Example 12	27
4.13 Example 13	27
4.14 Example 14	28
4.15 Example 15	28
5. Using Xtract TM	29
5.1 Installing Xtract from disk under MVS	29
5.1.1 Files on the distribution disk	29
5.1.2 Creating the Object Library	29
5.1.3 Creating the Load Library	30
5.1.4 Sending Xtract objects to the mainframe	30
5.1.5 Link-Editing Xtract	30
5.2 Installing Xtract from magnetic tape under MVS	31

5.2.1 Loading Xtract in an existing library.....	31
5.2.2 Loading Xtract in a new library.....	32
5.3 Xtract under VM/CMS	32
5.4 Collecting the SMF data.....	33
5.5 Collecting the GTF data	34
5.6 Typical Job Control Statements.....	35
5.6.1 Examples of JCL	36
Addendum	39
Frequently encountered problems	39
Xtgen JCL generator for Xtract.....	40
Processing multiple control units in a single Xtract run.....	41
Support for IXFP SMF records generated for Storage Technology controllers	44
Support for RMF type 74 subtype 5 records	44
Maximum number of devices	44
Multiple GTF traces	44
SMF type 42 subtype 6 records.....	45
ASTEX records	45
GTF records.....	45
VM MONITOR records	47
Appendix	47
A.1 Xtract condition codes	48
A.2 Report description.....	48