

Adabas

User, Hyperdescriptor, and Collation Descriptor Exits

Version 8.1.4

June 2014

This document applies to Adabas Version 8.1.4.

Specifications contained herein are subject to change and these changes will be reported in subsequent release notes or new editions.

Copyright © 1971-2014 Software AG, Darmstadt, Germany and/or Software AG USA, Inc., Reston, VA, USA, and/or its subsidiaries and/or its affiliates and/or their licensors..

The name Software AG and all Software AG product names are either trademarks or registered trademarks of Software AG and/or Software AG USA, Inc. and/or its subsidiaries and/or its affiliates and/or their licensors. Other company and product names mentioned herein may be trademarks of their respective owners.

Detailed information on trademarks and patents owned by Software AG and/or its subsidiaries is located at <http://documentation.softwareag.com/legal/>.

Use of this software is subject to adherence to Software AG's licensing conditions and terms. These terms are part of the product documentation, located at <http://documentation.softwareag.com/legal/> and/or in the root installation directory of the licensed product(s).

This software may include portions of third-party products. For third-party copyright notices and license terms, please refer to "License Texts, Copyright Notices and Disclaimers of Third-Party Products". This document is part of the product documentation, located at <http://documentation.softwareag.com/legal/> and/or in the root installation directory of the licensed product(s).

Document ID: ADAMF-EXITS-814-20140626

Table of Contents

Preface	v
1 Conventions	1
2 User Exit 1 (General Processing)	3
3 User Exit 2 (Dual Log Processing)	5
User Exit 2 Calling Sequence	7
Input Parameters	9
Output Parameter	11
BS2000 Options for Invoking User Exit 2	11
4 User Exit 3 (User-Defined Phonetic Processing)	13
Input Parameters	14
5 User Exit 4 (User-Generated Log Data)	15
Command Log Format	16
6 User Exit 5 (Adabas Review Hub Event Handler)	21
Input Parameters	22
Output Parameters	23
7 User Exit 6 (User Processing Before Data Compression)	25
ADACMP Header Processing	28
Input Parameters	29
Output Parameters	29
8 User Exit 8 (Operator Interface)	31
Input Parameters	33
9 User Exit 9 (ADAULD)	35
10 User Exit 11 (General Processing)	37
Input and Output Parameters	39
11 User Exit 12 (Multiple Data Set Log Processing)	41
User Exit 12 Calling Sequence	43
User Exit Interface	44
Output Parameter	47
Activating the Sample User Exit	47
12 Hyperdescriptor Exits 01 - 31	51
Main Parameter Area	53
Input Parameter Area (Pointed to by Third Parameter Address)	54
Output Parameter Area	56
Null Value Option	57
Hyperexit Initialization Call	58
Hyperexit Stub	58
13 Collation Descriptor Exits 01 - 08	61
Collation Descriptor Exit Interface	62
Index	65

Preface

This document refers to the user exits activated by the ADARUN parameters UEXn, HEXnn, and CDXnn (see the *Adabas Operations* documentation for descriptions of the ADARUN parameters).

The user exits documented in this document are as follows:

User Exit	ADARUN	Use
User Exit 1	UEX1	Command processing (Adabas nucleus) -- being retired.
User Exit 2	UEX2	Dual log processing
User Exit 3	UEX3	User-defined phonetic processing
User Exit 4	UEX4	User-generated log data
User Exit 5	UEX5	Adabas Review hub event handler
User Exit 6	UEX6	Data compression (ADACMP)
User Exit 8	UEX8	Operator interface
User Exit 9	UEX9	Data unload (ADAULD)
User Exit 11	UEX11	Command processing (Adabas nucleus) -- new.
User Exit 12	UEX12	Multiple log processing

Hyperdescriptor Exit	ADARUN	Use
1..31	HEX01 .. HEX31	User-supplied algorithm to create hyperindex values . . .

Collation Descriptor Exit	ADARUN	Use
1..8	CDX01 .. CDX08	User-supplied algorithm to encode and decode values for the corresponding collation descriptors . .

Other Exits Supported by Adabas

Other user exits supported by Adabas include the following:

Entry Name	Use
ADACDCUX	Allows you to obtain control at strategic points during ADACDC utility processing. See the <i>Adabas Utilities</i> documentation.
ADACSHUX	Allows you to obtain control at strategic points during Adabas Caching Facility processing. See the <i>Adabas Caching Facility</i> documentation.
ADASMxit	Allows you to supply parameters to a PIN routine or examine a condition when it is encountered before the PIN routine is invoked so that recovery actions other than those

Entry Name	Use
	provided by Adabas can be implemented. See <i>Adabas Online System Demo Version</i> in the DBA Tasks documentation..
DSFEX1	Automatically submits the necessary job to prevent overflow of the DLOG area. See the <i>Adabas Delta Save Facility</i> documentation.
UEXITA	Linked with Adalink for Adabas 7: receives control after a command is processed by a target, the router, or Adalink itself. See the <i>Adabas Installation</i> documentation .
UEXITB	Linked with Adalink for Adabas 7: receives control before a command is passed to a target with the router 04 call. See the <i>Adabas Installation</i> documentation .
UEXIT1	Linked with Adalink for Adabas 8: receives control before a command is passed to a target with the router 04 call. See the <i>Adabas Installation</i> documentation .
UEXIT2	Linked with Adalink for Adabas 8: receives control after a command is processed by a target, the router, or Adalink itself. See the <i>Adabas Installation</i> documentation .
UEXRAI	Allows you to change automatically generated ADARAI RECOVER JCL before it is written to DDJCLOUT. See the <i>Adabas Utilities</i> documentation.

1 Conventions

Notation *vrs*, *vr*, or *v*: When used in this documentation, the notation *vrs* or *vr* stands for the relevant version of a product. For further information on product versions, see *version* in the *Glossary*.

2 User Exit 1 (General Processing)

With the introduction of user exit 11, support for user exit 1 is dropped. However, to ease the migration, a sample user exit, UEX11UX1, is supplied that you can insert in front of your existing user exit 1 to have it invoked as user exit 11. This sample will only work for direct calls made using the ACB direct call interface; it will not work for direct calls made using the Adabas 8 ACBX direct call interface. The exit is still subject to exit 11 constraints, as described in [User Exit 11](#), elsewhere in this guide. In particular, changes are allowed only to the file number (CQXFNR), Additions 2 (ACBADD2), Additions 3 (ACBADD3), and user area (ACBUSER) fields. The nucleus will ignore changes in any other ACB field and all other changes to the CQX. Please refer to comments in the sample user exit for more details, including how to link it with an existing user exit 1. Adabas nucleus support for this transition aid will be withdrawn in a future Adabas release.

3 User Exit 2 (Dual Log Processing)

- User Exit 2 Calling Sequence 7
- Input Parameters 9
- Output Parameter 11
- BS2000 Options for Invoking User Exit 2 11

This user exit is given control by the Adabas nucleus during a switch from one dual log to the alternate dual log for the purpose of copying the log before it is reused by Adabas. This switch occurs only if dual data protection logging or dual command logging is in effect for the session.

The user exit routine must invoke a procedure whereby the appropriate function of the ADARES utility (CLCOPY or PLCOPY) is executed.

User exit 2 is invoked:

- during Adabas nucleus startup if a PLOG/CLOG has to be copied;
- whenever a dual command or dual protection log switch occurs between two log data sets;
- during Adabas nucleus shutdown.

The user exit is provided with information about the status of the dual log data sets.

The user exit can decide which action is to be taken:


- Ignore the call;
- Submit a job to copy the log data set just filled up (ADARES utility);
- Wait for completion of the copy job just submitted.

If the data set to be overwritten contains data, console message ADAN46 Function not executable is issued.

An example of user exit 2 is supplied with the Adabas installation procedure. Refer to the *Adabas Installation* documentation for more information.


The call to the user exit is made using a standard BASR 14,15 Assembler instruction. All registers must be saved when control is received and restored immediately prior to returning control to Adabas. Register 15 contains an action code as described in *Output Parameter*, elsewhere in this chapter.

User Exit 2 Calling Sequence



Dual Log Processing Flow

Input Parameters



Dual Log Processing User Exit (2) Parameters

The input parameters for the address list are as follows:

Parameter	A fullword address of . . .
0 (R1)	the C/PLOG indicators and flag 1/2.
4 (R1)	the four-byte timer 1 field.
8 (R1)	the four-byte timer 2 field.
12 (R1)	the current session's PLOG number, followed by the database ID.
16 (R1)	a four-byte area where the first two bytes contain the number of PLOG1, and the second two bytes hold the number of PLOG2.

Other input parameters are explained in the following table:

Parameter	Usage								
Flag 1	Status flags for DD/PLOGR1 and DD/CLOGR1; and								
Flag 2	<p>Status flags for DD/PLOGR2 and DD/CLOGR2:</p> <table border="1"> <tr> <td>B'1...' :</td> <td>Data set being written by nucleus</td> </tr> <tr> <td>B'.1..' :</td> <td>Data set has been completed by nucleus</td> </tr> <tr> <td>B'.11.' :</td> <td>Being copied by ADARES</td> </tr> <tr> <td>B'0000 0000' :</td> <td>Data set is empty (or copied) and reusable for the nucleus.</td> </tr> </table> <p>All other flag 1/2 field values are reserved. For DD/CLOGR1/2 only: X'08' for CLOGLAYOUT=5. Flag 1/2 bit settings can be combined (X'40' and X'20' as X'60', for example).</p> <p>If OPENOUT is specified, these flags are set after OPEN is issued for the output data set; otherwise, the flags are set before the OPEN is issued.</p>	B'1...' :	Data set being written by nucleus	B'.1..' :	Data set has been completed by nucleus	B'.11.' :	Being copied by ADARES	B'0000 0000' :	Data set is empty (or copied) and reusable for the nucleus.
B'1...' :	Data set being written by nucleus								
B'.1..' :	Data set has been completed by nucleus								
B'.11.' :	Being copied by ADARES								
B'0000 0000' :	Data set is empty (or copied) and reusable for the nucleus.								
Session Status	<p>Contains information about the status of the nucleus when the exit was called:</p> <table border="1"> <tr> <td>X'S'</td> <td>Called during nucleus session startup.</td> </tr> <tr> <td>X'T'</td> <td>Called while terminating the nucleus session.</td> </tr> <tr> <td>X'W'</td> <td>Called following a dual protection log switch.</td> </tr> </table>	X'S'	Called during nucleus session startup.	X'T'	Called while terminating the nucleus session.	X'W'	Called following a dual protection log switch.		
X'S'	Called during nucleus session startup.								
X'T'	Called while terminating the nucleus session.								
X'W'	Called following a dual protection log switch.								
TIMERn	Time-stamp (highest four bytes of a STCK instruction) for the time the first block of the log data set has been written. TIMER1 for DD/PLOGR1 and DD/CLOGR1, and TIMER2 for DD/PLOGR2 and DD/CLOGR2								
PLOG	Current session protection log number (two bytes). This value is set for PLOG only; the field contains X'00' for CLOG.								
DBID	Database ID (two bytes).								
PLOG1/2	Two 2-byte PLOG numbers found on PLOG 1 and PLOG 2. If the previous nucleus session ended abnormally, these four bytes contain that session's PLOGNUM value, which can be used in the initial user exit 2 call to copy that session's PLOG. During any subsequent session, these bytes contain the current PLOGNUM value. If the preceding session ends abnormally, these four bytes contain the ended session's PLOG numbers during the nucleus start phase. This PLOG information is needed during the start phase to assign the correct PLOG numbers to the PLOG areas to be copied. During subsequent exit calls, the current PLOG values are in these fields.								

Output Parameter

Parameter	Usage
R15 = 0	Nucleus continues processing.
R15 > 0	R15 is treated as the number of seconds to wait before calling user exit 2 again. During this time, the nucleus is in a "hard" wait. No commands are processed during the wait.

BS2000 Options for Invoking User Exit 2

When using user exit 2 with BS2000 systems, the name of the ADARES CLCOPY or PLCOPY job is no longer "hard-coded" in the user exit as in previous releases. Therefore, there are now two ways of specifying the /ENTER job:

Method 1

Create a job variable containing the complete "enter job" command. For example:

```
/DCLJV  ENTER.ADARES.PLCOPY, LINK=*DDJBPLC
/SETJV  *DDJBxLC,C 'ENTER JOB. ENTER.ADARES.PLCOPY, ST=IMM'
```

Note that all operands of the /ENTER command may be used. The presence of the /DCLJV statement overrides the other possibilities. If necessary, the content of the job variables can be changed during a nucleus session.

Use the link names "*DDJBPLC" for assigning the ADARES PLCOPY job, and "*DDJBCLC" for assigning the ADARES CLCOPY job.

Method 2

Omit the JV specification. The user exit 2 will then issue the following command:

```
'ENTER  RES.E.xLCO'
```



Note: For BS2000 systems, user exit 2 is delivered as a source element only.

4 User Exit 3 (User-Defined Phonetic Processing)


- Input Parameters 14

This user exit may be used to perform user-defined phonetic processing. It is given control by the ADACMP utility or the Adabas nucleus whenever phonetic processing is required.

The user exit must develop a three-byte phonetic key using the value supplied. The address of the resulting phonetic key must be placed at 8(R1) before control is returned.

Input Parameters

Register 1 contains the address of the following parameter list:



User-Defined Phonetization User Exit (3) Parameters

Parameter	A fullword address of . . .
0(R1)	the four-byte length for the value to be phonetically processed.
4(R1)	the address of the value to be phonetically processed.
8(R1)	a three-byte location to contain the phonetic key. This address is set to zero before the user exit and must be set to the actual address during the user exit.

The call to the user exit is made using a standard BASR 14,15 assembler instruction. All registers must be saved when control is received and restored immediately prior to returning control to Adabas. The content of R15 is ignored.

5 User Exit 4 (User-Generated Log Data)

- Command Log Format 16

User exit 4 is called immediately before an Adabas command log record is to be written. It may be used to generate any required user log data (SMF records) special statistics, or to suppress writing a log record.



Note: User exit 4 is still called even if ADARUN LOGGING=NO. The only way to disable user exit 4 is to remove the ADARUN UEX4 parameter from the Adabas run.


Command Log Format

Adabas supports two different command log formats. The ADARUN CLOGLAYOUT parameter determines which format is used:

- CLOGLAYOUT=5 (the default) is supported only in Adabas versions 5.2 and above.
- CLOGLAYOUT=8 specifies the new format, which is supported only in Adabas versions 8 and above.

Both formats are described in *Command Log Formats*, in *Adabas DBA Tasks Manual*.

Ensure that your user exit and command log evaluation programs recognize the format in use before switching to it.



User-Generated Log Data User Exit (4) Parameters

¹*Address of ACB Copy:* This address should be set to zero if the command is using an ACBX interface direct call.

²*Address of the first ABD:* The Adabas buffer descriptions (ABDs) are in a contiguous array. For complete information about locating ABDs in this array, read *Locating the Correct ABD*, elsewhere in this section.

Parameter	Address of . . .
0(R1)	a byte containing a logging action code. This byte contains: <ul style="list-style-type: none"> ■ +00 -- action code to log the record upon each call. If changed to a nonzero value, this record will not be written to DDLOG. ■ +01 -- reserved for future use ■ +02 -- two-byte database ID.
4(R1)	the record to be logged. This address is zero if the exit is called at the end of the nucleus session.
8(R1)	the end of the Adabas I/O area. This address is zero if the exit is called at the end of the nucleus session.
12(R1)	the command queue element (CQX). This address is zero if the exit is called at the end of the nucleus session.

The record to be logged may be modified by the user exit. The record's address in 4(R1) may also be modified. The logging action code must always be specified before returning to the Adabas nucleus.



Caution: When modifying the record, do not exceed the end address of the ADALOG I/O area contained in 8(R1).

Locating the Correct ABD

Internally, Adabas 8 only uses extended Adabas control blocks (ACBX) and Adabas buffer descriptions (ABDs). Direct calls made using the classic Adabas control block (ACB) and buffer definitions have their data structures converted to ACBX calls and ABDs by ADASVC before the nucleus sees the call. Thus, the protocol for locating and accessing buffers in user exits, such as this one, has changed as of Adabas 8.

The Adabas buffer descriptions (ABDs) are now in a contiguous array. However, the internal representation of the ABD may not have the same length as the base ABD, as defined by the value of the ABDXQLL symbol in the ADABDX DSECT, although the first ABDXQLL bytes continue to be mapped by ADABDX. This means that you should not use the ABDXQLL value in the ADABDX DSECT to locate the next ABD in the ABD array. Instead, you should use the value of the two-byte ABDXLEN field at offset +x'00' of the ABD to determine the end of that ABD and the start of the next ABD in the array. Do not assume that all internal ABD representations have the same length: each must be located in turn by applying its predecessor's ABDXLEN value.

In addition, the order of the ABDs is not defined and may change over time or from command to command, although within the array all ABDs of a given type (format buffer, record buffer, etc.) are contiguous. There will be an ABD for every buffer provided by the user that is documented as an input or output buffer for the specific command. There may also be additional buffers created by other components. When there are multiple instances of format, record and (optional) multifetch buffers, they are related based on their position: the first format buffer is associated with the first record (and optional multifetch) buffer, the second with the second, and so forth. If the caller

provides an unequal number of format, record and (optional) multifetch buffers, dummy descriptors with a zero buffer length are created to bring about equal quantities. When multifetch is used with a classic ACB call, certain commands (L1/2/3/4/9) will have their ISN buffer converted into a multifetch buffer. Here are some examples:

- If a caller (using either an ACB or ACBX call) issues an OP command and provides a record buffer and search buffer, the array of ABDs will have one record buffer ABD and one dummy format buffer ABD (to satisfy the internal requirement that there be equal numbers of format and record buffers). There is no ABD for the search buffer because that is not a documented input or output buffer for the OP command.
- If a caller uses an ACBX call to issue an L1 command and provides two format buffers and three record buffers, the array of ABDs will have three record ABDs and three format ABDs, the last one of which is a dummy format ABD. The first record buffer is associated with the first format buffer; the second record buffer is associated with the second format buffer; and the third record buffer is associated with the third (dummy) format buffer.
- Suppose a caller uses an ACB call to issue an L3 command with Command Option 1 set to "M" (multifetch) and Command Option 2 set to "A" (ascending retrieval from a specified value). In addition, the caller provides a format buffer, a record buffer, an ISN buffer, a search buffer and a value buffer. In this case, the array of ABDs will have one format buffer ABD, one record buffer ABD, one multifetch buffer ABD, one search buffer ABD, and one value buffer ABD. The caller's ISN buffer will have been converted to a multifetch buffer.

6 User Exit 5 (Adabas Review Hub Event Handler)

- Input Parameters 22
- Output Parameters 23

User exit 5 is called by the Adabas nucleus when an *event* occurs with the Adabas Review hub. User exit 5 must be specified in ADARUN parameter UEX5 in the Adabas nucleus startup job. An event is defined as:

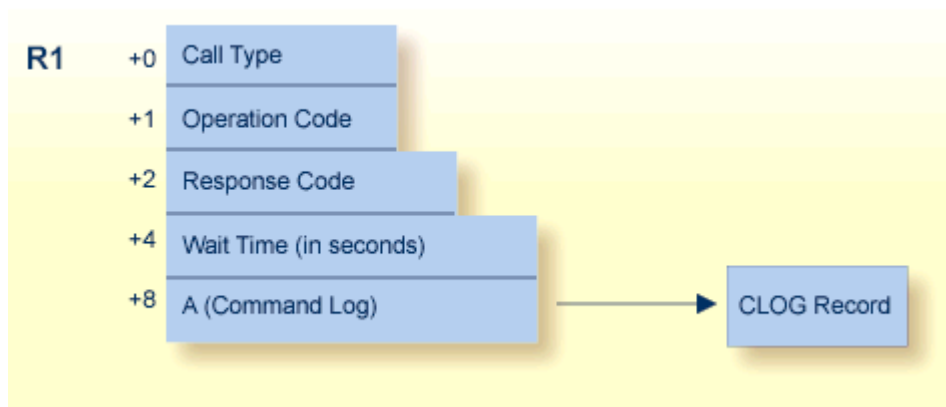
- a connection made with the Adabas Review hub during Adabas session open;
- a connection ended with the Adabas Review hub during Adabas session close; or
- a non-zero return code received from the send operation for a command log record. When buffering is active, this return code is provided once for a whole buffer and it is possible that only parts of the buffer are not transferred correctly.

The exit is invoked with AMODE=31 and should return control in the same state.

The exit is required to process logging errors. It determines how the failure is handled. The parts of the buffer that were not logged and the response code received from the Adabas Review hub logging request are provided to assist in making the determination.

Input Parameters

On entry, the register 1 points to the following parameter list:



Parameter	Usage						
0(R1)	Exit call indication. The value of this byte can be: <table border="1" style="width: 100%; border-collapse: collapse;"> <tr> <td style="width: 10%;">O</td> <td>connection with Adabas Review hub opened;</td> </tr> <tr> <td>C</td> <td>connection with Adabas Review hub closed; or</td> </tr> <tr> <td>L</td> <td>sending logging error to Adabas Review hub.</td> </tr> </table>	O	connection with Adabas Review hub opened;	C	connection with Adabas Review hub closed; or	L	sending logging error to Adabas Review hub.
O	connection with Adabas Review hub opened;						
C	connection with Adabas Review hub closed; or						
L	sending logging error to Adabas Review hub.						
1(R1)	Action to handle a logging error (ignored for open and close). The exit must provide one of the following values for this field in the parameter list for a logging error: <table border="1" style="width: 100%; border-collapse: collapse;"> <tr> <td style="width: 10%;">W</td> <td>wait a specified time and then retry;</td> </tr> </table>	W	wait a specified time and then retry;				
W	wait a specified time and then retry;						

Parameter	Usage	
	R	retry logging operation immediately; or
	I	ignore the logging failure and continue without consequence.
2(R1)	Response code for logging errors. This response code is the same as the Adabas response code found in the <i>Adabas Messages and Codes</i> .	
4(R1)	Fullword where the exit must provide a wait time (in seconds) for the logging failures that are to be retried after waiting.	
8(R1)	Address of the command log record that the Adabas nucleus was attempting to send to the Adabas Review hub.	

Other Register Values at Entry

R13	save area of calling Adabas nucleus routine
R14	return address in Adabas nucleus
R15	entry point address for exit

Output Parameters

- For logging errors, the exit is required to set a value in the 'operation' field. If the wait value (W) is chosen, the exit is also required to provide a non-zero time value.
- Register 15 should be set to zero. All other registers should be returned intact.

7 User Exit 6 (User Processing Before Data Compression)

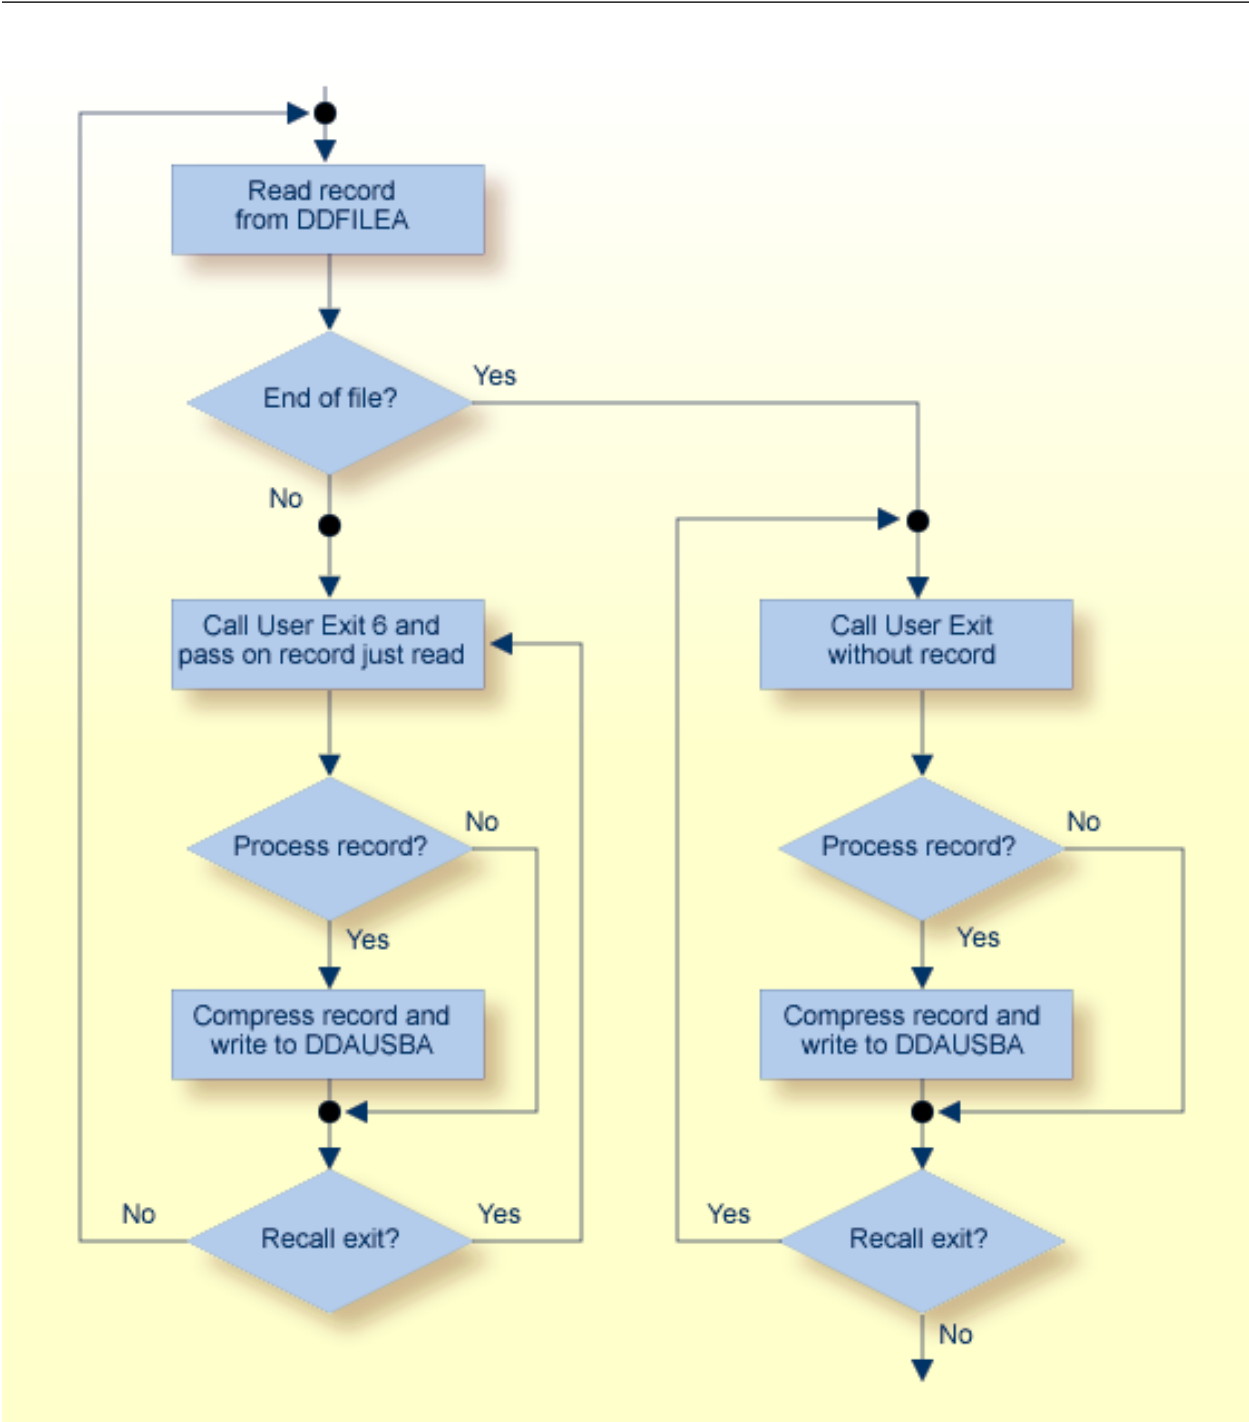
- ADACMP Header Processing 28
- Input Parameters 29
- Output Parameters 29

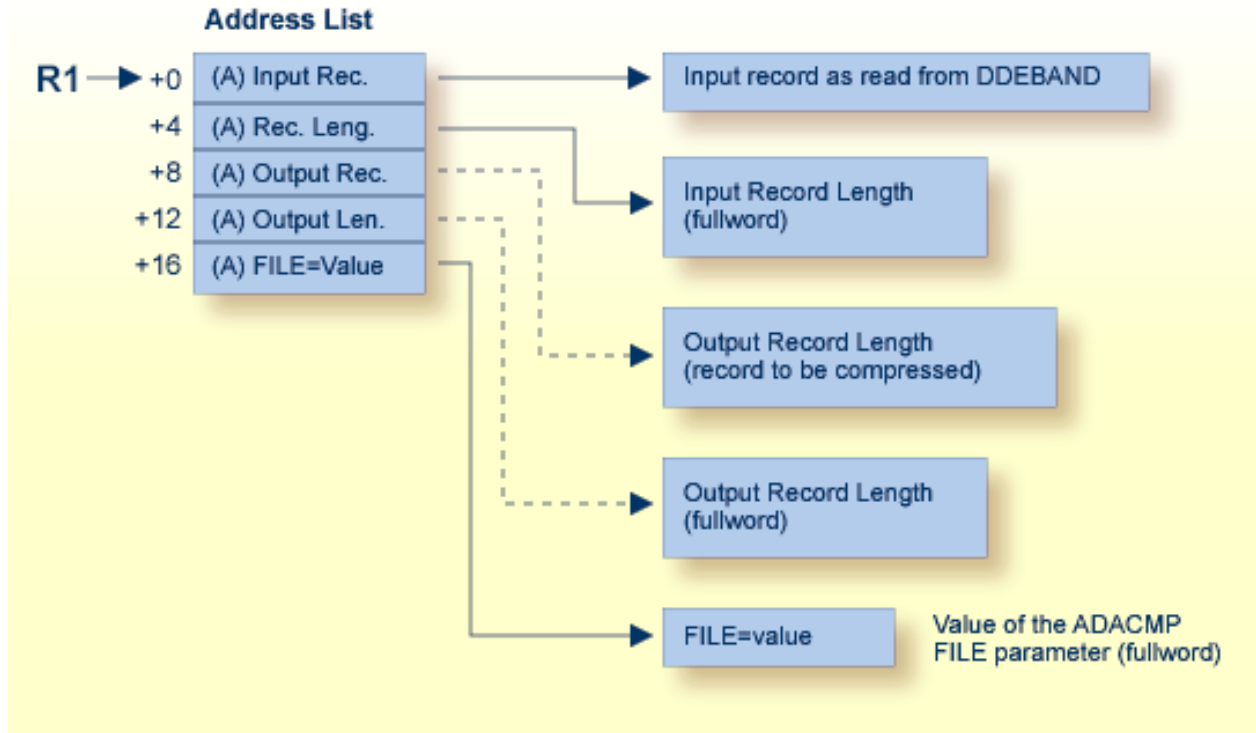
This user exit can be used to perform user processing on a record before it is processed by the ADACMP COMPRESS utility. It can also be used to control the sequence and contents of the decompressed records that are output from the ADACMP DECOMPRESS utility; when used in this way, the user exit controls which decompressed records ADACMP writes to the DDAUSBA data set.

The user exit program can be written in Assembler language or COBOL. The ADACMP utility job must specify the ADARUN UEX6=*program parameter*, where *program* is the name of the user program and *parameter* is a parameter passed to that program.

User exit 6 is called by the ADACMP COMPRESS utility function immediately after one of the following occurs so that it can append records to the input:

- A record has been read from DDEBAND.
- An end-of-file condition has occurred on DDEBAND.





ADACMP User Processing User Exit (6) Parameters

ADACMP Header Processing

When ADACMP is run with the parameter HEADER=YES, all input records for ADACMP COMPRESS and output records for ADACMP DECOMPRESS are preceded by 32-byte ADAH and ADAC headers that describe the grouping of physical records into logical records that may be larger than 32 KB. DSECTs for the ADAH and ADAC headers can be found in members ADAH and ADAC in the distributed Adabas 8 SRCE library. These headers identify how one logical record containing uncompressed data is composed of one or more physical records.

When ADACMP is run with user exit 6 and HEADER=YES, ADACMP will pass each physical record to user exit 6. The user exit application may need to use these headers to determine the relationship between the physical and logical records.

For more information about the ADACMP headers and record segmenting, read *Segmented Record Considerations*, in *Adabas Utilities Manual*.

Input Parameters

Parameter	Usage
0 (R1)	Address of an input record. The length field preceding the variable record is skipped. The address is of a fullword containing -1 (X'FFFF FFFF') if the user exit is called after ADACMP detects end-of-file in DD/EBAND.
4 (R1)	Address of the field containing the input record length. For fixed records, this is a logical record length. For variable records, this is the length of the actual data only (excluding the length field itself). The address points to a fullword containing minus 1 (X'FFFFFFFF') if the user exit is called after ADACMP detects end-of-file in DD/EBAND.
8 (R1)	Contains binary zeros on entry to the user exit (see Output Parameters).
12 (R1)	Contains binary zeros on entry to the user exit (see Output Parameters).
16 (R1)	Address of the FILE parameter value specified by the ADACMP COMPRESS utility job. The address is in the rightmost/low-order two bytes. The location and content of this fullword must remain unchanged during the time of the user exit. If ADACMP COMPRESS did not specify the FILE parameter, the fullword is X'00000000'.

Output Parameters

Parameter	Usage
8 (R1)	Address of the user exit output record. This record will be used as input to the ADACMP compression algorithm. The address of this record must be placed into 8 (R1) each time the user exit is called. If this field contains binary zeros on return, ADACMP will ignore the input record and will continue processing.
12 (R1)	Address of a 4-byte field containing the length of the returned record. The address of this field must be placed into 12 (R1) each time the exit is called. If this field contains binary zeros on return, ADACMP will ignore the record and will continue processing. Though the length field pointed to by 12 (R1) has a length of 4 bytes, only the low-order/rightmost halfword is used (bytes 3 and 4). If byte 2 contains a X'01' on return, the exit is recalled before the next record is read from DDEBAND. This enables the user to return more than one record to ADACMP for each record read from DD/EBAND.

8

User Exit 8 (Operator Interface)

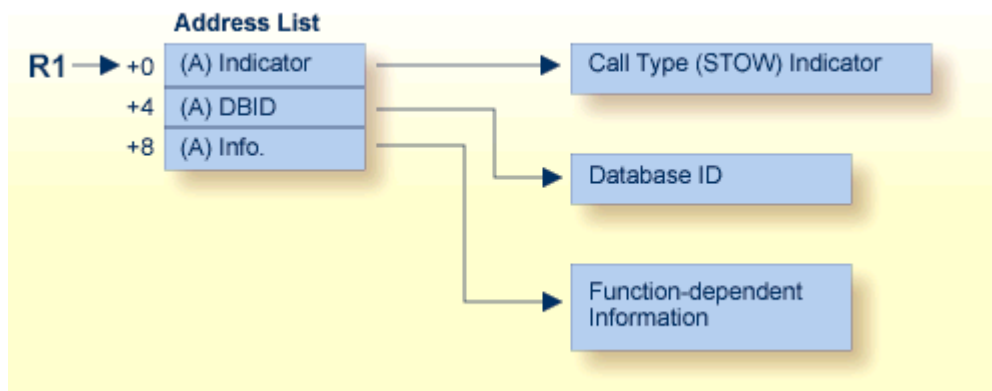
- Input Parameters 33

This user exit receives control from the Adabas nucleus whenever the nucleus starts or stops, or whenever the nucleus or an Adabas utility receives a message from or sends a message to the operator. User exit 8 can be used to provide specific instructions to the operator

- when the nucleus starts and (normally) stops operation;
- as added information when Adabas sends console messages to the operator;
- to confirm commands entered by the operator.

User exit 8 is invoked

- (MODE=MULTI only) after Adabas startup, as soon as the nucleus is able to answer calls from user programs. At this point, the nucleus is now active.
- immediately after the Adabas nucleus or utility issues a console operator message. The user exit call is in addition to the standard message processing; the message itself cannot be changed during the user exit.
- after the Adabas nucleus or utility receives an operator command. The exit is called before the command is actually processed, and can reject or replace the command. The command cannot be modified in its original area.
- before a normal Adabas nucleus stop. At this point, the nucleus is no longer active; any more nucleus calls result in response code 148. This exit is not called if the nucleus ends abnormally.



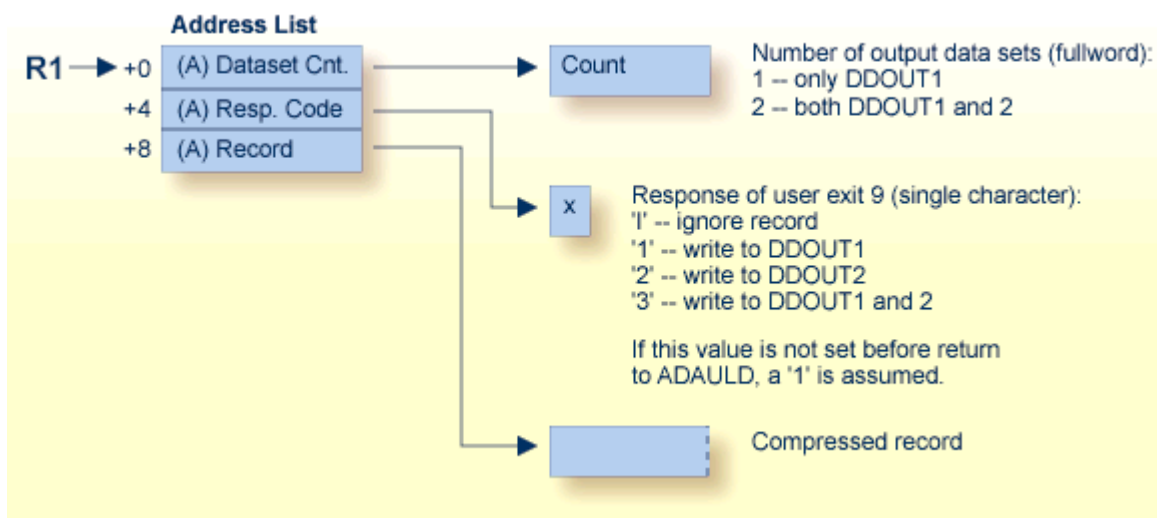
Operator Interface User Exit (8) Parameters

Input Parameters

Parameter	Usage								
0 (R1)	Address of the byte containing the call type ("STOW") indicator: <table border="1" data-bbox="373 441 1396 630"> <tr> <td>S</td> <td>called at nucleus start</td> </tr> <tr> <td>T</td> <td>called at normal nucleus termination</td> </tr> <tr> <td>O</td> <td>called with an operator message to the nucleus/utility</td> </tr> <tr> <td>W</td> <td>called with a nucleus/utility message to the operator</td> </tr> </table>	S	called at nucleus start	T	called at normal nucleus termination	O	called with an operator message to the nucleus/utility	W	called with a nucleus/utility message to the operator
S	called at nucleus start								
T	called at normal nucleus termination								
O	called with an operator message to the nucleus/utility								
W	called with a nucleus/utility message to the operator								
4 (R1)	Address of the fullword containing the database ID.								
8 (R1)	Address of variable-length message-related information for "O" and "W" type calls. The information at this address has the following format: Call format: <table border="1" data-bbox="373 756 1477 1155"> <tr> <td>O</td> <td>The one-byte message length, followed by the alphanumeric message. The length excludes the length byte itself. If the message is to be changed, location 8(R1) must point to the new message on return. This message is structured as described above. If the message is to be suppressed, location 8(R1) must point to a two-byte field containing X'0140'.</td> </tr> <tr> <td>W</td> <td>8(R1) points to the message, which has the following structure: DC H'message-length+4' DC H'0' DC C'message-text' ↵</td> </tr> </table>	O	The one-byte message length, followed by the alphanumeric message. The length excludes the length byte itself. If the message is to be changed, location 8(R1) must point to the new message on return. This message is structured as described above. If the message is to be suppressed, location 8(R1) must point to a two-byte field containing X'0140'.	W	8(R1) points to the message, which has the following structure: DC H'message-length+4' DC H'0' DC C'message-text' ↵				
O	The one-byte message length, followed by the alphanumeric message. The length excludes the length byte itself. If the message is to be changed, location 8(R1) must point to the new message on return. This message is structured as described above. If the message is to be suppressed, location 8(R1) must point to a two-byte field containing X'0140'.								
W	8(R1) points to the message, which has the following structure: DC H'message-length+4' DC H'0' DC C'message-text' ↵								

9 User Exit 9 (ADAULD)

User exit 9 is called by ADAULD whenever a compressed record is ready to be written. The user exit decides whether a record is written to DD/OUT1, DD/OUT2, both, or neither.



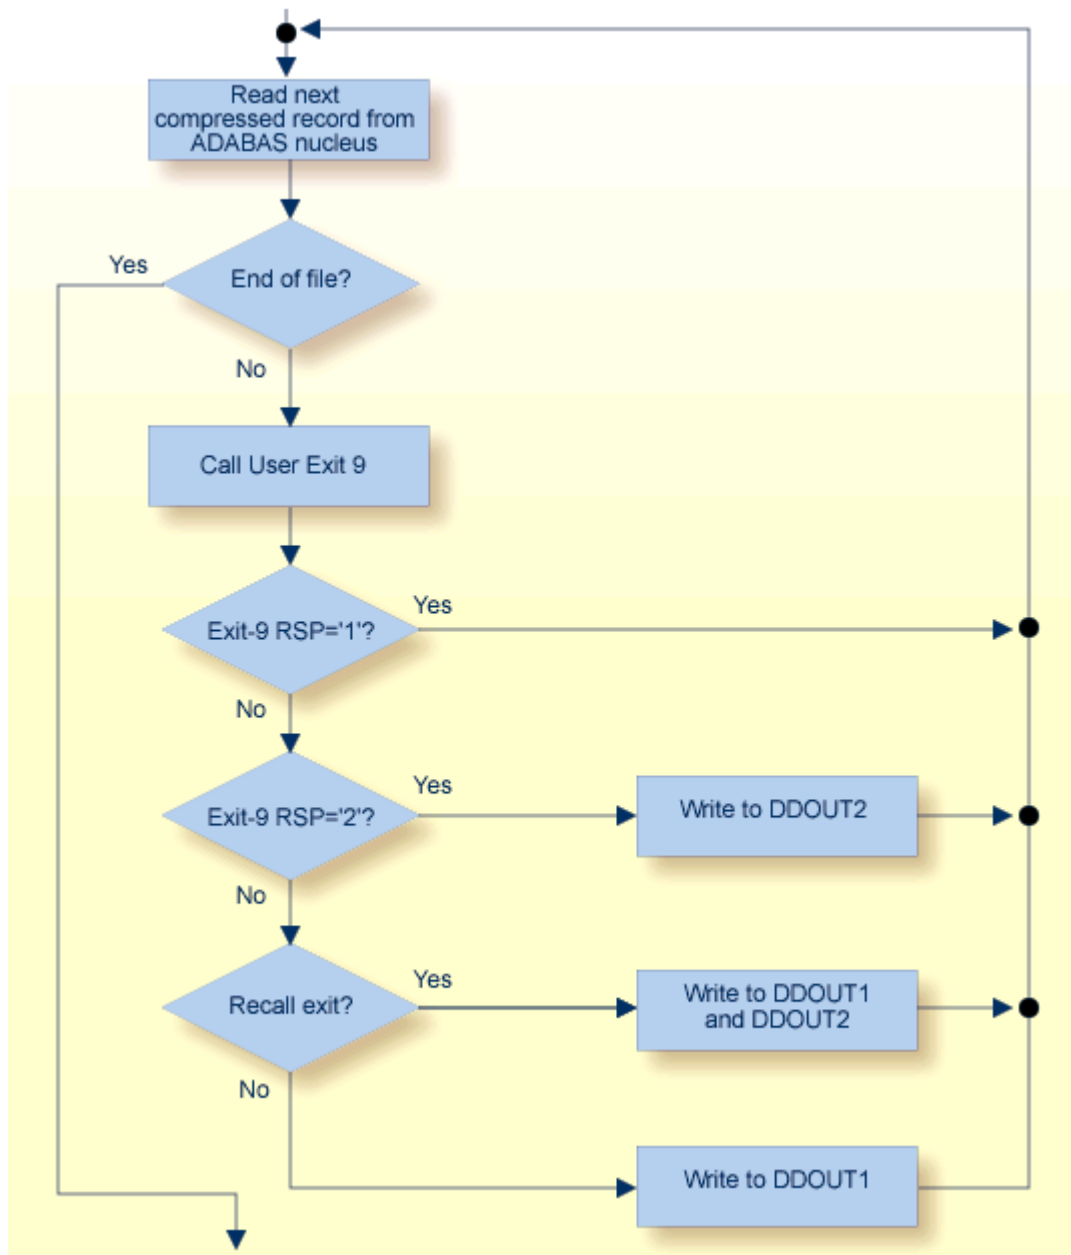
ADAULD User Exit (9) Parameters

 **Notes:**

1. DDOUT1 & 2 must have the same block size, or an ADAULD error occurs.

The compressed record pointed to by the third address has the following structure:

AL2 (L1) total length (inclusive)
 AL2 (L2) record length (inclusive)
 AL4 (ISN)
 XL (L2 - 8) '...compressed fields...'
 XL (L1 - L2 - 2) '...DVT entries...'



User Exit 9 Output Control Flow

The call to the user exit is made via a standard BASR 14,15 assembler instruction. All registers must be saved when control is received and restored immediately prior to returning control to ADAULD. The content of R15 is ignored.

10 User Exit 11 (General Processing)

- Input and Output Parameters 39

This user exit is given control by Adabas immediately after a command is received by the Adabas nucleus. The command itself has yet to be processed except for the determination of the type of command (simple access, complex access, update).

One of the most common applications of this user exit is to insert a security password or a cipher code into the ACBX.

This user exit functionality largely matches that of the classic user exit 1, except for the fact that edited copies of the CQX and ACBX data structures are used during user exit 11 processing, rather than the actual structures used by user exit 1. In addition, support for user exit 1 is dropped in Adabas 8 (or later).

Only certain fields in the ACBX may be changed by the exit: ACBXFNR (file number), ACBXADD2 (Additions 2), ACBXADD3 (Additions 3), and ACBXUSER (user area). The nucleus will ignore changes in any other ACBX fields and all changes to the CQX. DSECT EX11PARM maps the user exit 11 parameter list. In addition, a sample user exit 11 skeleton called UEX11 is provided. Both the DSECT and the user exit skeleton are provided in the Adabas source library.


The call to the user exit is made using a standard BASR 14,15 assembler instruction. Register 1 contains the address of a parameter list. All registers must be saved when control is received and restored immediately prior to returning control to Adabas, with the exception of Register 15 which contains the return code. A non-zero value means that the command should not be executed and returns response code 22.



Notes:

1. All user exits must return the same program status word (PSW) fields to the calling program that were active when the user exit was called. This applies in particular to the addressing mode (AMODE), program mask, problem state flag, PSW key, and address space control setting. The condition code need not be preserved. If any of these PSW fields is changed by the user exit, one way to ensure that their previous values are returned is to envelope the code where the change is in effect with a pair of the BAKR ... PR instructions.
2. The file number specified in the original ACBX for the call cannot be changed. If it is necessary to change a file number with the user exit, change the field ACBXFNR field in the copy of the ACBX that is made for the exit.
3. The command code field in the Adabas control block cannot be changed; a response code is returned if you attempt to do so.
4. The length of an Adabas buffer in any Adabas buffer description (ABD) used by the call cannot be changed.

Input and Output Parameters



General Processing User Exit (1) Parameters

¹*Indicators:* Before calling user exit 11, the fullword indicator areas are set to zero.

²*Parmlist length:* The EX11PARM parameter list length is at least 28 bytes.

³*Address of ACB Copy:* This address should be set to zero if the command is using an ACBX interface direct call.

⁴*Address of the first ABD:* The Adabas buffer descriptions (ABDs) are in a contiguous array. For complete information about locating ABDs in this array, read *Locating the Correct ABD*, next in this section.

Locating the Correct ABD

Internally, Adabas 8 only uses extended Adabas control blocks (ACBX) and Adabas buffer descriptions (ABDs). Direct calls made using the classic Adabas control block (ACB) and buffer definitions have their data structures converted to ACBX calls and ABDs by ADASVC before the nucleus sees the call. Thus, the protocol for locating and accessing buffers in user exits, such as this one, has changed as of Adabas 8.

The Adabas buffer descriptions (ABDs) are now in a contiguous array. However, the internal representation of the ABD may not have the same length as the base ABD, as defined by the value of the ABDXQLL symbol in the ADABDX DSECT, although the first ABDXQLL bytes continue

to be mapped by ADABDX. This means that you should not use the ABDXQLL value in the AD-ABDX DSECT to locate the next ABD in the ABD array. Instead, you should use the value of the two-byte ABDXLEN field at offset +x'00' of the ABD to determine the end of that ABD and the start of the next ABD in the array. Do not assume that all internal ABD representations have the same length: each must be located in turn by applying its predecessor's ABDXLEN value.

In addition, the order of the ABDs is not defined and may change over time or from command to command, although within the array all ABDs of a given type (format buffer, record buffer, etc.) are contiguous. There will be an ABD for every buffer provided by the user that is documented as an input or output buffer for the specific command. There may also be additional buffers created by other components. When there are multiple instances of format, record and (optional) multifetch buffers, they are related based on their position: the first format buffer is associated with the first record (and optional multifetch) buffer, the second with the second, and so forth. If the caller provides an unequal number of format, record and (optional) multifetch buffers, dummy descriptors with a zero buffer length are created to bring about equal quantities. When multifetch is used with a classic ACB call, certain commands (L1/2/3/4/9) will have their ISN buffer converted into a multifetch buffer. Here are some examples:

- If a caller (using either an ACB or ACBX call) issues an OP command and provides a record buffer and search buffer, the array of ABDs will have one record buffer ABD and one dummy format buffer ABD (to satisfy the internal requirement that there be equal numbers of format and record buffers). There is no ABD for the search buffer because that is not a documented input or output buffer for the OP command.
- If a caller uses an ACBX call to issue an L1 command and provides two format buffers and three record buffers, the array of ABDs will have three record ABDs and three format ABDs, the last one of which is a dummy format ABD. The first record buffer is associated with the first format buffer; the second record buffer is associated with the second format buffer; and the third record buffer is associated with the third (dummy) format buffer.
- Suppose a caller uses an ACB call to issue an L3 command with Command Option 1 set to "M" (multifetch) and Command Option 2 set to "A" (ascending retrieval from a specified value). In addition, the caller provides a format buffer, a record buffer, an ISN buffer, a search buffer and a value buffer. In this case, the array of ABDs will have one format buffer ABD, one record buffer ABD, one multifetch buffer ABD, one search buffer ABD, and one value buffer ABD. The caller's ISN buffer will have been converted to a multifetch buffer.

11 User Exit 12 (Multiple Data Set Log Processing)

- User Exit 12 Calling Sequence 43
- User Exit Interface 44
- Output Parameter 47
- Activating the Sample User Exit 47



Note: UEX2 and UEX12 are mutually exclusive for an Adabas nucleus session: only one can be specified.

This user exit is given control by the Adabas nucleus during a switch from one multiple log data set to another for the purpose of copying the log data set before it is reused by Adabas. This switch occurs only if multiple data set data protection logging and/or multiple data set command logging is in effect for the session.

The user exit routine is designed to invoke a procedure that will execute the appropriate function (CLCOPY or PLCOPY) of the ADARES utility.

User exit 12 is invoked

- during Adabas nucleus startup if a multiple PLOG/CLOG data set has to be copied;
- whenever a switch to another log data set occurs;
- during Adabas nucleus shutdown.

The user exit is provided with information about the type of log (PLOG or CLOG) and the status of the multiple log data sets.

The user exit can decide which action is to be taken:

- Ignore the call and allow Adabas to proceed;
- Submit a job to copy and mark as empty the log data set just filled (ADARES utility);
- Direct Adabas to wait for a specified interval, then call the user exit again with updated PLOG/CLOG data set status information. During the wait interval, no commands that may produce log records for the log type being processed are allowed to proceed.

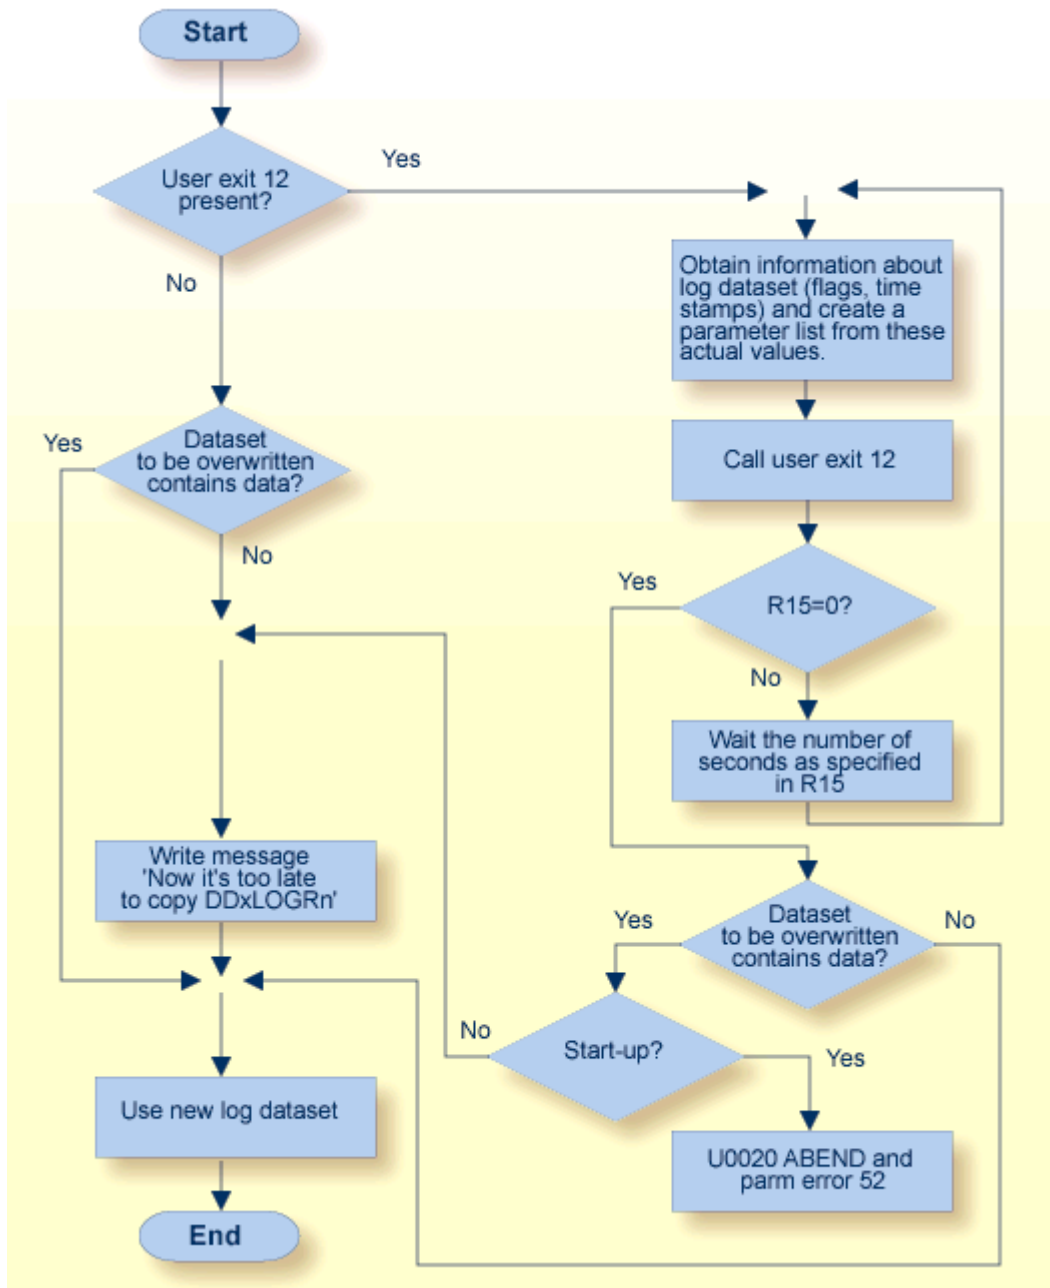
An example of user exit 12 is supplied with the Adabas installation procedure. Refer to the *Adabas Installation* documentation for more information.

The call to the user exit is made using a standard BASSM R14,R15 Assembler instruction. All registers must be saved when control is received and restored immediately prior to returning control to Adabas. Register 15 contains an action code as described in the user exit 2 section *Output Parameter*, elsewhere in this guide.



Note: User exit 12 must return the same AMODE value to the calling program that was active when user exit 12 was called. The recommended Assembler instruction to return is BSM 0,R14.

User Exit 12 Calling Sequence




Multiple Log Processing Flow

User Exit Interface

Optionally, the user exit may initialize its operation. It may store any value in field EX12USER of the EX12PARM parameter block to keep track of its resources. This field is considered as "owned" by the user exit and is supplied again for all subsequent executions of the exit. It is set to zero when the exit is first called and is not modified by Adabas thereafter.

The user exit is called again during termination to do any necessary finishing or 'cleanup' work.

Parameters



User Exit 12 Parameters

DSECT of the EX12PARM Parameter Block

```

MACRO
EX12PARM
*****
.*
.* Name          Ex12Parm                               .*
.*
.* Component     ADABAS User Exit                       .*
.*
.* Function      Parameter list for User Exit 12        .*
.*               (replacement for User Exit 2 for use when there are .*
.*               more than 2 PLOGs or CLOGs)            .*
.*
.* Parameters    None                                    .*
.*
.* Restrictions  None                                    .*
.*
.* Notes        None                                    .*
.*
*****
.*
EX12PARM DSECT ,           User Exit 12 Parameter List
.*
EX12HDR  DS    0F          Common header section
.*
EX12USER DS    F          Reserved for the user. This field +
                          is initialized to zero before the +
                          exit is called the first time, and +
                          will not be altered by ADABAS after +
                          that. It can be used to maintain +
                          information across invocations.
.*
EX12LOGT DS    X          Log type
EX12PLOG EQU    C'P'      PLOG
EX12CLOG EQU    C'C'      CLOG
.*

```

User Exit 12 (Multiple Data Set Log Processing)

```

EX12TYPE DS      X      Call type
EX12TBEG EQU     C'S'   Nucleus start
EX12TSW  EQU     C'W'   Log switch
EX12TEND EQU     C'T'   Nucleus termination
          DS      XL2    Reserved
*
EX12NLOG DS      F      Number of logs
EX12DBID DS      F      Database ID
EX12NUCI DS      F      Nucleus ID
*
EX12PLGN DS      F      Current session PLOG number      +
                       (zero for CLOGs)
EX12NCMP DS      F      Log just completed
EX12STAT DS      X      Flags of next log in sequence
EX12WNUC EQU     X'80'  Being written by the nucleus
EX12FULL EQU     X'40'  Completed by the nucleus
EX12RES  EQU     X'20'  Being copied by ADARES
EX12CL5  EQU     X'08'  CLOGLAYOUT=5
EX12UNUS EQU     X'00'  Unused and/or copied
          DS      XL3    Reserved
          DS      4F     Reserved
*
EX12HDRL EQU     *-EX12HDR Length of header section
*
EX12ENT  DSECT   ,
EX12LOG  DS      0F     Start of individual log data set      +
                       entries. This section is repeated      +
                       for the number of logs specified in      +
                       field EX12NLOG
*
EX12LTIM DS      XL8    Time stamp of write to log data set
EX12LNUM DS      F      Number of log data set
EX12LFLG DS      X      Flags (mapped as in EX12STAT)
          DS      XL3    Reserved
          DS      4F     Reserved
*
EX12LOGL EQU     *-EX12LOG Length of a log data set entry
*
*
          MEND

```

Output Parameter

Parameter	Usage
R15 = 0	Nucleus continues processing.
R15 > 0	R15 is treated as the number of seconds to wait before calling user exit 12 again with updated status for all log data sets. During this time, no commands that may create log entries are processed.

Activating the Sample User Exit

The sample user exit is written in Assembler language. It performs the following functions:

- Issues a message identifying the reason and the type of log for which it was called.
- Issues a message with the status and timestamp of all log data sets that are not empty.
- If any log data set is full and at least one log data set has a status that is different from the last time the exit was called, the exit reads 80-byte records from an input file and writes them to an output file. It replaces all occurrences of the character "?" with either "P" or "C", depending on whether the exit was invoked to process a PLOG or CLOG event, respectively. This allows the input file to accommodate an event for either log type. Normally, the input file contains job control statements and the output file is directed to a job execution queue.
- If at least one log data set is not full, the exit returns to the caller with R15 zero, which allows Adabas to proceed.
- If all log data sets are full, the exit returns to the caller with R15 nonzero, which directs Adabas to wait for the number of seconds in R15, then call the exit again with an updated status of all log data sets. The default delay time is 30 seconds.
 - [Activating in z/OS](#)
 - [Activating in VSE](#)

- [Activating in BS2000](#)

Activating in z/OS

The sample user exit UX12SAMP is delivered on z/OS as source and as a load module that can be used without change or reassembly.

The source and load forms of the user exit are delivered in the Adabas source and load libraries, respectively. The job to assemble the user exit UX12ASML is located in the Adabas jobs library. The jobs library also contains a sample job UX12CJOB to be customized and submitted by the user exit that invokes the ADARES utility PLCOPY or CLCOPY function.

Activate the sample user exit as follows:

1. In addition to ADARUN NxLOG={2-8}, specify ADARUN UEX12=UX12SAMP for the Adabas nucleus.
2. Supply the job control model that the user exit is supposed to submit under the DDNAME COPYJOB.

Provide the following DD statement:

```
//INTRDR2 DD SYSOUT=(*,INTRDR)
```



Note: The same DD statement is used by the sample user exit 2 or 12 for submitting PLCOPY or CLCOPY jobs.

Activating in VSE

The sample user exit UX12SAMP is delivered on VSE as source and as a phase that can be used without change or reassembly.

The source A.book and phase are delivered in the Adabas VSE sublibrary. The job to assemble the user exit UX12ASML.X is also provided in the Adabas sublibrary. The sample job UX12CJOB.X is to be customized and submitted by the user exit that invokes the ADARES utility PLCOPY or CLCOPY function.

Activate the sample user exit as follows:

1. In addition to ADARUN NxLOG={2-8}, specify ADARUN UEX12=UX12SAMP for the Adabas nucleus.
2. Supply the job control model that the user exit is supposed to submit under the DLBL COPYJOB.

Activating in BS2000

The sample user exit UEX12 is delivered on BS2000 as source and can be assembled using the sample job ASMB2.as as an object module that can be used without change or reassembly.

To activate the sample user exit:

1. In addition to ADARUN NxLOG={2-8}, specify ADARUN UEX12=UEX12 for the Adabas nucleus.
2. Use one of the following methods to set up the Adabas nucleus job appropriately:
 - [Method 1 - With Job Variables](#)
 - [Method 2 - Without Job Variables](#)

Method 1 - With Job Variables

If the machine can process job variables, place the following statement in the Adabas nucleus job:

```
/DEL-JV ENTER.ADARES.xLCOPY
/SET-JOB-STEP
/CRE-JV ENTER.ADABAS.xLCOPY
/SET-JV-LINK *DDJBx12,ENTER.ADABAS.xLCOPY
/MOD-JV LINK(*DDJBx12),-
/ SET-VAL='ENT-PROC (src1ib,ux12cjbx,DB=dbid,J-N=X12COPY'
```

or, in ISP format:

```
/DCLJV ENTER.ADARES.xLCOPY, LINK=*DDJBx12
/SETJV *DDJBx12,C'ENTER-PROC(src1ib,ux12cjbx),(DB=dbid),J-N=X12COPY'
```

- x* is either "P" for a PLCOPY job or "C", for a CLCOPY job
- src1ib* is the library where the UX12CJBx is stored
- ux12cjbx* is the element name of the respective procedure job.
- dbid* is the database ID of this nucleus.

Method 2 - Without Job Variables

If the machine cannot process job variables, place the following statement in the Adabas nucleus job:


```
ENTER-PROC (ADABAS.SRC,UX12CJBx),(DBID=dbid),J-N=X12xCOPY
```

Be sure to set *x* to "P" for a PLCOPY job or "C" for a CLCOPY job. Also set *dbid* to the calling DBID (UX12CJBC and UX12CJBP examples can be taken from the ADARESCC and ADARESCP examples respectively).

12 Hyperdescriptor Exits 01 - 31

- Main Parameter Area 53
- Input Parameter Area (Pointed to by Third Parameter Address) 54
- Output Parameter Area 56
- Null Value Option 57
- Hyperexit Initialization Call 58
- Hyperexit Stub 58

The hyperdescriptor exits (hyperexits) 1 through 31 (HEX01...HEX31) are required to define the algorithm for user-supplied descriptor values (see the *Adabas Utilities Manual* documentation). A hyperexit is called by the ADACMP utility or the Adabas nucleus whenever a hyperdescriptor value is to be generated. ADACMP always uses the hyperdescriptor exit specified in its own ADARUN statement. When the ADAINV utility specifies a hyperdescriptor exit, the exit used is the one specified in the Adabas nucleus' ADARUN statement.



Hyperdescriptor Exit Parameters




Notes:

1. Hyperexits must return the same AMODE value to the calling program that was active when the hyperexit was called.
2. If Adabas 8 (or later) is installed and your hyperexit has not been updated to use the new parameter list, the Hyperexit Stub can be linked to your hyperexit to provide the necessary parameter list changes and input parameter conversion. The Hyperexit Stub does not provide extended MU or PE support. For more information, read [HyperExit Stub](#), elsewhere in this chapter.
3. An initialization call is made to each loaded hyperexit during Adabas nucleus or ADACMP startup. For more details, read [Hyperexit Initialization Call](#), elsewhere in this chapter.

Main Parameter Area

Parameter	Content
0 (R1)	Reserved (must not be changed)
4 (R1)	Fullword of zeros (must not be changed)
8 (R1)	Address of the beginning of the input parameter area.
12 (R1)	Address of the beginning of the output parameter area. This address must be inserted by the user-written program. An output parameter area must always be returned by the user hyperexit. If no values are to be returned, the address will point to a Header Element with a total length that indicates no Value Elements exist.

Input Parameter Area (Pointed to by Third Parameter Address)



Header Element Fields

LL	Total length of the input parameter area, including this length field
FNR	File number
ISN	ISN assigned to the record
HN	Name of the hyperdescriptor
F	Flag byte: <ul style="list-style-type: none"> ■ X'02' indicates file with extended MU or PE fields ■ X'80' indicates initialization call

Parent Element Fields

FN	Name of the parent field
L	Length of the value pointed to by VALADDR if the parent field is defined with the FI option.
I	Four-byte periodic group index of the parent field. If the parent field is not part of a PE group, these bytes contain zeros.
VALADDR	Address of the value of the parent field. The format of the value depends on the options of the fields. If the parent field is defined with the NU (null value suppression) option and the value for this field is suppressed, no input parameter element is created.

The following examples show formats for the value pointed to by VALADDR for parent fields with combinations of the FI (fixed storage) and MU (multiple-value) options:


Fields without FI and without MU option:




Fields with FI and without MU option:




Fields without FI and with MU option:



Fields with FI and MU option:



Fields without FI and with MU option when extended count is present:



where:


- C A one-byte value representing the MU count. If the MY value is for a file defined with extended MU or PE fields, an extended count may be present. For more details on the extended count, read *Identifying MU and PE Occurrences Greater Than 191 in Compressed Records*, found in the ADACMP documentation in *Adabas Utilities Manual*.
- *L A hexadecimal value length, including this one- or two-byte length value. For lengths from 1 through 127, only a single byte is required. For lengths ranging 128 to 255, two bytes are needed: the first byte is set to X'80', and the second byte is set to the actual length value (see the following example table):

Length	Byte 1	Byte 2
L=127:	x'7F'	(x'80')
L=128:	x'80'	x'80'
L=255:	x'80'	x'FF'

Output Parameter Area

This area must be allocated and filled within the hyperdescriptor user exit. The address of this area must be placed into the second position of the main parameter area.

This area consists of a 8-byte header followed by the generated hyperdescriptor values in compressed format.



Header Element

LL	Total length of the output parameter area, including this length field. If no values are returned, the total length is set to the length of the Header Element.
00	Reserved space. This must be set to zeros.
RC	Return code. The hyperexit may set a non-zero value here to indicate the call is rejected; a value of "16" is recommended. If this field is non-zero, the call will fail with response code 79.
ISN	The ISN to be assigned to the descriptor values. If the original ISN is to be changed, the new ISN must be inserted here. If these four bytes contain zero on return to the Adabas nucleus, the original ISN is used. This is a four-byte binary value.



Note: If the hyperexit returns an ISN in the ISN field of the header element, the file must be defined with USERISN=YES to prevent ISN reassignment when the file is later reloaded.

Value Elements

L	Length of the following value, including this length byte. The maximum length depends on the format in use for the hyperdescriptor.
Value	The descriptor value to be inserted into the index. The value must follow the rules in effect for the format assigned to this hyperdescriptor. If the hyperdescriptor is defined with the PE option, one byte containing the one-byte PE index must immediately follow the value and be included in length L. If the hyperdescriptor defined with the PE option is for a file defined with extended MU or PE fields, two bytes containing the two-byte PE index must immediately follow the value and be included in length L. The nucleus checks values of packed or numeric format for validity. Valid signs for packed fields are A,C,E,F (positive) and B,D (negative). The nucleus changes all signs to F or D.

Examples:

L	Value	Notes
04	R E D	
06	B L U E02	where X'02' is a PE index
03	123F	packed 123
04	123F01	packed 123 in PE group with index 1
07	B L U E0002	where X'0002' is a PE index for a file defined with extended MU or PE fields
05	123F010A	packed 123 in extended PE group with index 266

Null Value Option

The NU (null value) option is possible for the hyperdescriptor or parent fields. The possible combinations are as follows:

- The hyperdescriptor is not NU:
 - The parent field is not NU and the value is null, the hyperexit is called and the null value is passed.
 - The parent field is NU and the value is null, the hyperexit is called and no input parameter element is created for this parent field.
 - All parent fields are NU and all values are null, the hyperexit is called and no input parameter element is created for any parent field.
- The hyperdescriptor is NU:

- The parent field is not NU and the value is null, the hyperexit is called and the null value is passed.
- The parent field is NU and its value is null, the hyperexit is called and no input parameter element is created for this parent field.
- All parent fields are NU and all values are null, the hyperexit is not called.

Hyperexit Initialization Call

During Adabas nucleus or ADACMP startup, each loaded hyperexit is called with an initialization call. The main parameter area must be used as documented. The third parameter address will point to an input parameter area with a header length indicating that no values follow. The flag byte will be set to x'80' to indicate the initialization call. Upon return, the hyperexit must set the fourth parameter address to an output parameter area with a header length indicating that no values are returned.

Hyperexit Stub

The Hyperexit Stub is provided to allow earlier hyperexits to use the Adabas 8 parameter list without change. The Hyperexit Stub is intended as a temporary solution for those customers who do not wish to immediately update their hyperexits to use the new parameter areas. The Hyperexit Stub will not function for files that are defined with extended MU or PE fields; a response code will be given when the Hyperexit Stub is called for such files. Hyperexits linked with the Hyperexit Stub may be used with earlier versions of Adabas, however, the Hyperexit Stub must not be used with hyperexits that use the Adabas 8 parameters.

Sample job LNKHEX8 in the JOBS data set provides an example for linking the Hyperexit Stub to your hyperexit.

In z/VM environments, use the following procedure to build your hyperexit with the hyperexit stub:

1. Edit the file INPUT TEXT provided on the distribution tape. On the third line (NAME HEXn-nMOD), change the name to your hyperexit. The same name should be set in the ADARUN HEXnn parameter. Do not change any other lines. Save the file.


```
INCLUDE LIBDEF(HEXV87,HEXOLD)
INCLUDE TXTDEF
NAME HEXnnMOD
```

2. Generate the hyperexit with the following commands:
 - FILEDEF LIBDEF DISK ADAV813 TXTLIB *x* (where *x* is the filemode containing ADAV813 TXTLIB)
 - FILEDEF TXTDEF DISK *HEXnn* TEXT *x* (where *HEXnn* is your hyperexit, and *x* is the filemode where it resides)
 - LKED INPUT (LIBE *HEXnn*) (where *HEXnn* is the name of the new LOADLIB containing the hyperexit)
3. Before starting the Adabas nucleus, specify GLOBAL LOADLIB *HEXnn*.

13 Collation Descriptor Exits 01 - 08

- Collation Descriptor Exit Interface 62

The collation descriptor exits 1 through 8 (CDX01 through CDX08) are used for encoding and decoding values for the corresponding collation descriptors.

A collation descriptor may be defined for a field with alphanumeric or wide format. Its values are stored in the index, not in the record itself. The number of the collation descriptor exit used to derive the values is associated with the collation descriptor.

A sample collation descriptor exit CDXE2A is provided in the Adabas source data set. It converts EBCDIC to ASCII for the encoding function and the reverse (ASCII to EBCDIC) for the decoding function.

The Collation Exit implements three function entry points which are called on the following events:

INITIALIZE function

- nucleus session start
- utility initialization when collation exits have been defined (ADARUN parameters)

ENCODE function

- update/insert/delete of the parent's value (Nucleus)
- Search specifying the collation descriptor with the search value (Nucleus)
- compression of a record (ADACMP)

DECODE function

- Read Index (L9) by Collation DE, only if the exit supports the DECODE function (Nucleus)

Collation Descriptor Exit Interface


The collation descriptor exit interface is defined in the CDXPARAM DSECT in the Adabas source data set. The interface has three functions:

- initialization
- encoding
- decoding (optional)

Initialization Parameters

R1 points to a list of addresses that point to five storage areas of the caller. The collation descriptor exit must set the five areas as follows:

CDXSPC	default space character; a maximum of 4 bytes
CDXSPCL	fullword containing the size of the space character
CDXENC	address of encoding function
CDXDEC	address of decoding function If the returned address is zero, decoding is not supported. The collation descriptor cannot then be used for L9 processing.
CDXVER	address of zero-byte delimited version string




Encoding/Decoding Parameters

R1 points to a list of five fullword-sized parameters:

CDXIA	address of the input string
CDXIL	length of the input string
CDXOA	address of the output string
CDXOL	size of the output area
CDXARL	address of the length of the returned output string

The collation descriptor exit stores

- the output string in the area at the address specified by CDXOA; and
- the input string length in the fullword at the address specified by CDXARL.



Index

A

Adabas Review
 hub event handler user exit 5, 21
ADAULD utility
 user exit 9 processing, 35

B

BS2000
 options for user exit 2, 11

C

cipher code
 program to insert in ACBX, 38
collation descriptor exits, 61
command log
 dual data set user exit, 6
command log (CLOG)
 format, 16
 multiple data set user exit, 41
 user exit 4 processing, 15
command processing user exit, 3, 37

D

data compression
 user exit 6 processing before, 25
dual log processing, 5

E

exit 1, 3
exit 11, 37
exit 12, 41
exit 2, 5
exit 3, 13
exit 4, 15
exit 5, 21
exit 6, 25
exit 8, 31
exit 9, 35

H

hyperdescriptor
 user exit, 51

description, v

L

logs
 dual data set user exit 2 processing, 5
 multiple data set user exit, 41

O

operator
 program to provide instructions to, 32
operator interface user exit, 31

P

password
 program to insert in ACBX, 38
phonetic processing
 processing with user exit 3, 13
protection log
 dual data set user exit, 6
protection log (PLOG)
 multiple data set user exit, 41

U

user
 log data processing, 16
user exit
 exit 12, 44
user exits
 BS2000 user exit 2 options, 11
 descriptions, v
 exit 1, 3
 exit 11, 37
 exit 12, 41
 exit 2, 5
 exit 3, 13
 exit 4, 15
 exit 5, 21
 exit 6, 25
 exit 8, 31
 exit 9, 35
 overview, v

