

Adabas

Release Notes

Version 8.2.3

May 2011

This document applies to Adabas Version 8.2.3.

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

Copyright © 1971-2011 Software AG, Darmstadt, Germany and/or Software AG USA, Inc., Reston, VA, United States of America, and/or their licensors.

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).

Table of Contents


1 Adabas 8.2 Release Notes	1
2 Supported Platforms	3
3 Enhancements	5
Extended Address Volume (EAV) Support	8
License Refresh and Verification Support	8
Improved ADASAV Volume Information on z/OS	9
Session Autorestart Progress Messages	9
New Sample COBOL Programs	9
Unique Descriptor Pool Minimum Size Increased	10
SMGT Response Code Handling	10
ADALNK Access to Multiple ID Tables on BS2000	10
System Field Support	11
Partial LOB Field Processing	12
Shared Hold Status	12
SMF Support (z/OS only)	14
Command and Protection Log Alert Messages	15
Printing Multifetch Statistics to a File	16
LOREX DSECT Updates	16
Logically Deleting and Undeleting Descriptors	17
Logically Deleting and Undeleting File Fields	17
Maximum Length of Fixed Point Fields Increased	17
New Date-Time Edit Masks	17
Date-Time Edit Mask Enhancements	18
Time Zone Support	18
Daylight Savings Time Support	19
RI Commands: Change Possibly Affecting Your Application Programs	20
Terminology and Name Changes	20
Lowercase Field Name Enhancement	21
Expanded File Chain Enhancements	21
Putting Files in Read-Only State	21
Database Conversion Improvements	22
Unique Descriptor (UQ) Option Removal	22
Work Area Monitoring	23
PLOG and CLOG Data Set Enhancements	23
Spanned Record Logging	24
New Adabas Event Log Support	25
Response Code 145 (ADARSP145 - Hold Queue Entry) Event Enhancements	25
Multiple CICS Task-Related User Exit (TRUE) Support	26
CICS Installation Changes	26
Performance Enhancements	26
Nucleus 64-Bit Processing Support	28
Adabas Buffer Description (ABD) Structure Changes	28
Large Page Support	28

z/OS SVC Dump Enhancements	29
z/OS SVC Authorization Enhancements	29
z/OS SVC Prelinked to SAF	29
Reporting on Replication-Related Records in Work Part 1	29
Printing Adabas Maintenance and SVC Status Information	30
Replication of Security Definitions	30
Address Converter Extent Handling Enhancement	30
Adabas Delta Save Update: ADASAV MERGE Limitation Lifted	30
New and Changed Utility Functions	31
New and Changed Operator Commands	39
ADARUN Parameter Changes	40
Direct Call Command Enhancements	44
New LGBLSET Macro Parameters	47
4 Dropped Features	49
5 Limitations and Restrictions	51
6 Adabas Data Set Compatibility	55
Importing Files	56
Save Data Sets	56
Unload Data Sets	57
ADAORD DD/FILEA Data Sets	57
Sequential Protection Logs	57
BS2000 Containers Expanded by ADADBS INCREASE	58
7 Applying Zaps	59
8 Software AG Product Compatibility	61
9 End of Maintenance	65
10 Documentation	67
Allowing Active Content in Windows Environments	68
Index	69

1 Adabas 8.2 Release Notes






This document provides a brief summary of the new and changed features included in Adabas 8.2, with links for more information to other areas of the Adabas documentation set.

Because the updates made to Adabas 8 are quite extensive, we recommend that all Adabas 7 (or earlier) users read the *Planning Manual for Adabas 8* before proceeding with the Adabas 8 installation. This planning documentation describes the architectural changes and new features of Adabas 8, as well as any compatibility or migration issues between Adabas 8 and prior versions of Adabas.

 **Important:** Be sure that you apply all supplied Adabas 8 maintenance and concatenate Adabas 8 patch-level libraries (L00*n*), as they are delivered to you. This will ensure that your Adabas 8 code remains up-to-date, supporting all Adabas 8 features as they are enhanced and maintained.

If you are upgrading to this Adabas release from a release prior to the most recent Adabas release (for example, if you are upgrading from Adabas 7.4.4 to Adabas 8.2 and skipping the intermediate upgrade to Adabas 8.1), please read the Release Notes for the releases you are skipping to get a complete understanding of all of the changes implemented in Adabas since you last updated your software. The Release Notes for all currently supported versions of Adabas (on mainframes) are provided in *Previous Release Notes*, elsewhere in this guide.

This document covers the following topics:

 <i>Supported Platforms</i>	Lists the currently supported operating environments for Adabas 8.
 <i>Enhancements</i>	Describes the new and changed features in Adabas 8.2.
 <i>Dropped Features</i>	Lists Adabas features that are no longer supported in Adabas 8.
 <i>Limitations and Restrictions</i>	Lists the limitations and restrictions currently existing in Adabas 8.
 <i>Adabas Data Set Compatibility</i>	Describes the compatibility of Adabas data sets across Adabas releases.



•	<i>Applying Zaps</i>	Describes general information on where to locate and how to apply Adabas zaps.
•	<i>Software AG Product Compatibility</i>	Describes the compatibility of Adabas 8 with other Adabas products.
•	<i>End of Maintenance</i>	Describes how you can determine the end-of-support dates for your Software AG products.
•	<i>Documentation</i>	Describes the documentation provided with this release of Adabas.

2 Supported Platforms

Software AG generally provides support for the operating system versions supported by their respective manufacturers; when an operating system provider stops supporting a version of an operating system, Software AG will stop supporting that operating system version.

Before attempting to install this product, ensure that your host operating system is at the minimum required level. For information on the platform versions supported by Software AG products, access Software AG's Empower web site at <https://empower.softwareag.com>.

Log into Empower. Once you have logged in, you can expand **Products** in the left menu of the web page and select **Product Version Availability** to access the Product Version Availability application. This application allows you to review the supported platforms and releases for any Software AG product.

Product Version Availability		Results per Page	100
Sort by Version	<input checked="" type="radio"/> Descending <input type="radio"/> Ascending		
Time Frame	2010 - 2012 ▼		
Product Name	- ▼		
Product Family	- ▼		
Product Version	[] . [] . []		
Operating System	- ▼		
Operating System Version	[] . []		
Select historical data	<input type="checkbox"/>	 Clear	 Find

Use the fields on this application to filter its results. When you click the **Find** button, a list of the supported Software AG products that meet the filter requirements is shown. You can clear all filter selections using the **Clear** button.



Note: Although it may be technically possible to run a new version of this product on an old operating system, Software AG cannot continue to support operating system versions that are no longer supported by the system's provider. If you have questions about support, or if you plan to install this product on a release, version, or type of operating system other than those described in the Product Version Availability application as described above, consult Software AG technical support to determine whether support is possible, and under what circumstances.

3 Enhancements

▪ Extended Address Volume (EAV) Support	8
▪ License Refresh and Verification Support	8
▪ Improved ADASAV Volume Information on z/OS	9
▪ Session Autorestart Progress Messages	9
▪ New Sample COBOL Programs	9
▪ Unique Descriptor Pool Minimum Size Increased	10
▪ SMGT Response Code Handling	10
▪ ADALNK Access to Multiple ID Tables on BS2000	10
▪ System Field Support	11
▪ Partial LOB Field Processing	12
▪ Shared Hold Status	12
▪ SMF Support (z/OS only)	14
▪ Command and Protection Log Alert Messages	15
▪ Printing Multifetch Statistics to a File	16
▪ LORECX DSECT Updates	16
▪ Logically Deleting and Undeleting Descriptors	17
▪ Logically Deleting and Undeleting File Fields	17
▪ Maximum Length of Fixed Point Fields Increased	17
▪ New Date-Time Edit Masks	17
▪ Date-Time Edit Mask Enhancements	18
▪ Time Zone Support	18
▪ Daylight Savings Time Support	19
▪ RI Commands: Change Possibly Affecting Your Application Programs	20
▪ Terminology and Name Changes	20
▪ Lowercase Field Name Enhancement	21
▪ Expanded File Chain Enhancements	21
▪ Putting Files in Read-Only State	21
▪ Database Conversion Improvements	22
▪ Unique Descriptor (UQ) Option Removal	22
▪ Work Area Monitoring	23
▪ PLOG and CLOG Data Set Enhancements	23
▪ Spanned Record Logging	24

- New Adabas Event Log Support 25
- Response Code 145 (ADARSP145 - Hold Queue Entry) Event Enhancements 25
- Multiple CICS Task-Related User Exit (TRUE) Support 26
- CICS Installation Changes 26
- Performance Enhancements 26
- Nucleus 64-Bit Processing Support 28
- Adabas Buffer Description (ABD) Structure Changes 28
- Large Page Support 28
- z/OS SVC Dump Enhancements 29
- z/OS SVC Authorization Enhancements 29
- z/OS SVC Prelinked to SAF 29
- Reporting on Replication-Related Records in Work Part 1 29
- Printing Adabas Maintenance and SVC Status Information 30
- Replication of Security Definitions 30
- Address Converter Extent Handling Enhancement 30
- Adabas Delta Save Update: ADASAV MERGE Limitation Lifted 30
- New and Changed Utility Functions 31
- New and Changed Operator Commands 39
- ADARUN Parameter Changes 40
- Direct Call Command Enhancements 44
- New LGBLSET Macro Parameters 47

This chapter describes the Adabas 8.2 enhancements.

Category	Enhancements
General 8.2 Enhancements	<i>Extended Address Volume (EAV) Support</i>
	<i>New and Changed Utility Functions</i>
	<i>New and Changed Operator Commands</i>
	<i>ADARUN Parameter Enhancements</i>
	<i>Direct Call Command Enhancements</i>
	<i>New LGBLSET Macro Parameters</i>
	<i>New Sample COBOL Programs</i>
Adabas 8.2 SP3 Enhancements	<i>License Refresh and Verification Support</i>
	<i>Improved ADASAV Volume Information on z/OS</i>
	<i>Session Autorestart Progress Messages</i>
	<i>Unique Descriptor Pool Minimum Size Increased</i>
Adabas 8.2 SP2 Enhancements	<i>SMGT Response Code Handling</i>
	<i>ADALNK Access to Multiple ID Tables on BS2000</i>
	<i>System Field Support</i>
	<i>Partial LOB Field Processing</i>
	<i>Shared Hold Status</i>
	<i>SMF Support (z/OS only)</i>
	<i>Command and Protection Log Alert Messages</i>
	<i>Printing Multifetch Statistics to a File</i>
	<i>LORECX DSECT Updates</i>
	<i>Logically Deleting and Undeleting Descriptors</i>
	<i>Maximum Length of Fixed Point Fields Increased</i>
	<i>Date-Time Edit Mask Enhancements</i>
	<i>Time Zone Support</i>
	<i>Daylight Savings Time Support</i>
<i>RI Commands: Change Possibly Affecting Your Application Programs</i>	
Adabas 8.2 SP1 Enhancements	<i>Logically Deleting and Undeleting File Fields</i>
	<i>New Date-Time Edit Masks</i>
	<i>Lowercase Field Name Enhancement</i>
	<i>Expanded File Chain Enhancements</i>
	<i>Putting Files in Read-Only State</i>
	<i>Database Conversion Improvements</i>
	<i>Unique Descriptor (UQ) Option Removal</i>
<i>Work Area Monitoring</i>	

Category	Enhancements
	<i>PLOG and CLOG Data Set Enhancements</i>
	<i>Spanned Record Logging</i>
	<i>New Adabas Event Log Support</i>
	<i>Response Code 145 (ADARSP145 - Hold Queue Entry) Event Enhancements</i>
	<i>Multiple CICS Task-Related User Exit (TRUE) Support</i>
	<i>CICS Installation Changes</i>
	<i>Performance Enhancements</i>
	<i>Nucleus 64-Bit Processing Support</i>
	<i>Adabas Buffer Description (ABD) Structure Changes</i>
	<i>Large Page Support</i>
	<i>z/OS SVC Dump Enhancements</i>
	<i>z/OS SVC Authorization Enhancements</i>
	<i>z/OS SVC Prelinked to SAF</i>
	<i>Reporting on Replication-Related Records in Work Part 1</i>
	<i>Printing Adabas Maintenance and SVC Status Information</i>
	<i>Replication of Security Definitions</i>
	<i>Address Converter Extent Handling Enhancement</i>
	<i>Adabas Delta Save Update: ADASAV MERGE Limitation Lifted</i>
	<i>Terminology and Name Changes</i>

Extended Address Volume (EAV) Support

Adabas supports the use of extended address volumes (EAV) under z/OS. For a description of the Software AG requirements for this support, review the documentation at the following link: http://www.softwareag.com/corporate/images/Software_AG_ETS_Product_Support_for_Extended_Address_Volumes__EAV_07222010_tcm16-43365.pdf.

License Refresh and Verification Support

You can now refresh the license file currently used by the nucleus. This allows you to overwrite the current license with a new one or force the system to reread and reload its existing license. To support this functionality, Adabas 8.2 SP3 introduces a new operator command: LICREFRESH. For complete information, read *LICREFRESH Command*, in the *Adabas Operations Manual*.

Improved ADASAV Volume Information on z/OS

Adabas 8.2 SP3 includes an enhancement to ADASAV SAVE that provides more accurate volume information in z/OS environments. In previous releases, due to the asynchronous nature of ADASAV I/O processing, a number of blocks written on one volume might have been reported by ADASAV as written on the preceding volume.



Important: To enable this improved ADASAV volume information, apply optional Adabas zap AY823135 and IBM APAR OA33465. If you choose to apply AY823135, make sure you also apply IBM APAR OA33465. Installing AY823135 without IBM APAR OA33465 can cause the ADASAV SAVE backup to become corrupted. An ADASAV RESTORE would then fail (with system abend S002) and the database might then be unusable.

Session Autorestart Progress Messages

While performing a session autorestart after a nucleus failure, Adabas 8.2 SP3 now prints progress messages that indicate how much repair and recovery work has already been done and how much remains to be done. By default, these progress messages occur at the end of each phase of autorestart processing (backward repair, forward repair, autobackout). A new ADARUN parameter, `INTAUTO`, has been provided that allows you to specify the time interval (in seconds) between autorestart progress messages, thus allowing you to request that additional progress messages be printed for each phase of the autorestart process.

For complete information about the `INTAUTO` ADARUN parameter, read *INTAUTO Parameter: Time Interval Between Autorestart Progress Messages*, in the *Adabas Operations Manual*.

New Sample COBOL Programs

Adabas 8.2 SP3 includes two new sample COBOL programs, `COBADA8` and `COBACI8`. These programs are provided to illustrate how to make Adabas ACBX direct calls in a batch/TSO environment (`COBADA8`) and under CICS/TS (`COBACI8`). Their source modules should be compiled using the Enterprise COBOL compiler Version 3 Release 3 or later.

The `COBADA8` member is provided in the `ADAvrs.MVSSRCE` library and the `COBACI8` member is provided in the `ACIvrs.MVSSRCE` library.



Note: These programs are provided "as is" and will not be supported by Software AG.

For more information about these sample programs, read *ACBX Examples*, in the *Adabas Command Reference Guide*.

Unique Descriptor Pool Minimum Size Increased

Starting with Adabas 8.2 SP3, the minimum setting of the unique descriptor pool will be increased from zero to 5000 bytes. This size is established by the ADARUN LDEUQP parameter.



Note: This new functionality is also available in Adabas 8.2 SP2 if you apply zap AN822111.

For complete information, read about the LDEUQP parameter in *LDEUQP Parameter: Length of Unique Descriptor Pool*, in the *Adabas Operations Manual*.

SMGT Response Code Handling

The following changes, made in Adabas 8.2 SP2, affect how SMGT response codes are handled:

- PINRSP will now be loaded automatically at startup if the ADARUN SMGT parameter is set to "YES", similarly to ADAMXY.
- The output from the PIN routines can be directed to DDTRACE1 rather than DDPRINT if ADASMXIT has been updated by the user to specify this and if DDTRACE1 is provided in the nucleus startup JCL.

To activate this support, be sure to apply fix AN822020 and use the ADASMXIT from the ADA822.S001 data set.



Note: If you are using Adabas 8.2 SP3 or a later version of Adabas, applying this fix is unnecessary as the fix has already been applied to the source for you.

Information about this processing and DDTRACE1 are provided for z/OS sessions in *Adabas Session Execution*, in the *Adabas Operations Manual* and in *Installing The Error Handling And Message Buffering Feature*, in the *Adabas z/OS Installation Guide*.

ADALNK Access to Multiple ID Tables on BS2000

In prior versions of Adabas, routing of data from a client to a database on BS2000 systems could only take place over a single ID table defined by the IDTNAME parameter in the following ways:

- As a parameter in the file attached to the link name DDCARD (ADARUN statements);
- As a parameter in the file attached to the link name DDLNKPAR (ADALNK statements); or
- By modifying the parameter ENVNAME in the macro B2CONFIG and assembling the SSFB2C parameter repository module.

In Version 8.2 SP2, it is now possible to route data to specific database IDs to other ID tables, and optionally, to change the database ID to another in the target ID table. For more information, read *ADALNK Access to Multiple ID Tables*, in the *Adabas BS2000 Installation Guide*.

System Field Support

Support is added in Adabas 8.2 SP2 for *system fields*. A system field is a field in an Adabas file whose value is automatically set by the Adabas nucleus when records are inserted or updated on the file. Optionally, you can specify that some system field values only be set when records are inserted. The following types of system fields are currently supported:

- Job name: The job name of the user inserting or updating a record.
- ETID: The eight-byte user ID of the user inserting or updating a record. This is the user ID set in the Additions 1 field of an OP (open) command for the user session.
- Session ID: The 28-byte user ID of the user inserting or updating a record.
- Session user: The last eight bytes of the 28-byte session ID or the user inserting or updating a record.
- Time: The date or date and time at which a record is inserted or updated.

Two new field options have been added to this release to be used when defining your system fields with the ADACMP COMPRESS utility function: SY and CR. The SY option should be used to identify the contents that should be stored in the system field; the CR option should be used to indicate that the system field value should only be set when the record is inserted into the Adabas file (it will not be modified when an update operation occurs). For more information about the SY and CR field options, read *Field Options*, in the *Adabas Utilities Manual*.

In addition, the ADALOD LOAD utility function can now specify the maximum number of occurrences for MU system fields allowed in a file using a new SYFMAXUV parameter. For more information, read *LOAD: Load a File*, in the *Adabas Utilities Manual*.

Finally, a new SYFMAXUV parameter is supported by the ADADBS MODFCB utility function that allows you to modify the maximum number of occurrences for MU system fields allowed in a file. For more information, read *MODFCB: Modify File Parameters*, in the *Adabas Utilities Manual*.

For complete information about system fields, read *System Fields*, in *Adabas Concepts and Facilities Manual*.

Partial LOB Field Processing

Adabas 8.2 SP2 introduces new partial large object field (LOB field) processing. Your application programs can now read or write parts of LOB values, or *LOB segments*. In addition, in a series or read or write requests for a LOB field, your applications can control the position of the next LOB segment within the LOB value. This support introduces the following new features:

- New *LOB segment notation* is provided for you to use in the format buffer of a command to select the specific portion (segment) of the LOB value you want to read or write in the call. For more information about this notation, read *Field Selection Criteria*, in the *Adabas Command Reference Guide*.
- A new command option, *L*, can be specified in Command Option 2 of ACB or ACBX calls for L1/L4 or A1 commands. This command option requests that the current position in the LOB value be tracked in the ISN Lower Limit (ISL) field of the ACBX. In conjunction with this, a new current positioning notation (an asterisk) can now be used in the *LOB segment notation* of a call to indicate that the call should start at the current position (the position where the last call left off) of the LOB value.

These features allow you to read or write the LOB value in segments using multiple calls. The current positioning notation (an asterisk) allows you to easily indicate where the next call should start, without changing the format buffer for every call.

For more information, read *A1 Command: Update Record* and *L1 and L4 Commands: Read / Read and Hold Record*, in the *Adabas Command Reference Guide*.

For complete information about processing partial LOB fields and all of these new features, read *Processing LOB Segments*, in the *Adabas DBA Tasks Manual*.

Shared Hold Status

Adabas 8.2 SP2 introduces a new kind of competitive updating control called *shared hold status*. You can use shared hold status to lock data records in shared mode, rather than in the existing exclusive hold status. This allows your database users to read the same record in parallel transactions, but ensures that no one can update the record concurrently.

Using shared hold status, your users can protect large object values from concurrent updates without locking out other users who may need to read the same LOB value or other LOB values in the same record. It also allows your users to protect the records they read against concurrent updates for specific periods of time:

- For the duration of the read command;
- When the next record in a sequence is read;

- When the user's transaction ends;
- Indefinitely.

This new functionality introduces four new command options for the BT, ET, HI, L4 - L6, RI, and S4 commands, although all four new command options do not necessarily apply to all of these commands. These new options are specified in the command option 3 field of ACBX only calls for the associated command. Each command option specifies a different shared hold time period.

- Option C puts the record in shared hold status for the duration of the read operation. It ensures that the version of the record being read has been committed by the last updater. This option is available for the L4, L5, L6 and S4 commands. For the S4 command, the shared hold remains in place for the duration of the read operation.
- Option Q puts the record in shared hold status until the next record in the read sequence is read or the read sequence or transaction is terminated, whichever happens first. It ensures that the record being read cannot be updated concurrently until the next record in the sequence is read (or the transaction is terminated). This option is available for the L4 (when command option 2 is set to "N"), L5, L6 and S4 commands.
- Option S puts the record in shared hold status until the end of the transaction. It ensures that the record being read cannot be updated concurrently until the transaction is terminated. This option is available for the HI, L4, L5, L6, RI and S4 commands.
- Option H keeps a record in shared hold status indefinitely (until the next ET or BT command). This option is available for the BT and ET commands. Records in shared hold status at the time of the BT or ET command are kept in shared hold status beyond the end of the transaction until another ET or BT command is issued (without this H option or the prefetch or multifetch options). Any records in exclusive control are also changed to shared hold status beyond the end of the transaction.

If the same record is placed in shared hold status more than once (using the C or S options or the Q option for different read sequences), it stays in shared hold status until all of the specified hold lifetimes have expired.

For more information about these command options and their detailed functioning in each command, read about the individual commands in *Commands*, in the *Adabas Command Reference Guide*.

Resource deadlock may occur when users are simultaneously holding records and waiting for records held by other users. Normally, Adabas detects a user deadlock situation and returns a response code to the user who would create the deadlock if placed in wait status. However, the deadlock may not be detected if the users involved are serviced by different nuclei in a cluster. Even in a single nucleus, a deadlock may not be detected if a record involved is being held as a shared resource by more than one user. If a deadlock is not detected, all users involved will wait until the first one of them exceeds the transaction timeout (TT) threshold.

Putting more records in hold status (including shared hold status) may decrease the possible parallelism of transaction processing (and thus, the performance of multiuser applications) and

increase the likelihood of deadlocks between transactions (where each transaction holds a record that the other wants).

Using shared hold status affects the ADARUN NH and NISNHQ parameter settings. Each shared hold request with a different command ID (CID), as well as a (shared or exclusive) hold request without a CID, is counted against the NISNHQ and NH limits. This affects application programs that make use of the new Q option for sequential reads. For example, if a program reads records with the Q option and then updates every record, the shared hold operations from the use of the Q option and the exclusive hold operations from the update commands are counted separately. Such a program might need NISNHQ and NH limits set twice as large as when the Q option is not specified.

For complete information about shared hold status, read *Shared Hold Status*, in the *Adabas Command Reference Guide*.

SMF Support (z/OS only)

For z/OS environments only, support for IBM System Management Facility (SMF) records has been added to Adabas 8.2 SP2. SMF recording is widely used in IBM and ISV software products. Adabas provides support for SMF records in two ways:

1. Adabas nuclei now generate SMF records to record performance and resource consumption. Existing Adabas statistics are the basis for standard Adabas records and a new user exit permits you to include your own data.
2. Adabas nuclei can now be registered with SMF to record CPU consumption in IBM-defined type 89 records.

Seven new ADARUN parameters are introduced as a result of this support:

- The SMF parameter can be used to enable or disable Adabas SMF recording.
- The SMF89 parameter controls whether the Adabas nucleus registers with z/OS for type 89 SMF records.
- The SMFRECNO parameter can be used to define the Adabas SMF record number for user-defined records.
- The SMFINTERVAL parameter can be used to specify whether interval SMF records should be generated and, if so, how often.
- The SMFSUBSYS parameter can be used to associate the SMF records with an IBM or user-defined subsystem.
- The SMFDETAIL parameter can be used to specify the detail sections of Interval and Termination records that should be included (if any).

- The UEXSMF parameter can be used to identify the name of the SMF user exit module that should be used.

For complete information about these ADARUN parameters, read *ADARUN Parameter Directory*, in the *Adabas Operations Manual*.

This enhancement also introduces six new operator commands: SMFRECNO, SMFINTERVAL, SMFSUBSYS, SMFDETAIL, SMFDETAILADD, and SMFDETAILDEL. These commands can be used to alter the comparably named ADARUN parameter settings in a running Adabas session. The SMFDETAILADD and SMFDETAILDEL commands allow you to add and remove specific detail sections of Interval and Termination records from the list of records that should be recorded. These commands can be issued only as operator commands; they cannot be issued in an ADADBS OPERCOM utility function run. For complete information about these commands, read *Operator Commands*, in the *Adabas Operations Manual*.

Finally, this enhancement allows you to provide a user exit that adds a user-specified detail section to the Adabas SMF record. It is a separate load module whose name is provided in the UEXSMF ADARUN parameter. For complete information, read *SMF User Exit*, in the *Adabas User, Hyper-descriptor, Collation Descriptor, and SMF Exits Manual*.

For complete information about Adabas SMF records, read *Adabas SMF Records*, in the *Adabas DBA Tasks Manual*.

Command and Protection Log Alert Messages

Adabas 8.2 SP2 introduces the ability to optionally have Adabas generate alert messages indicating the status of the command log (CLOG) and the protection log (PLOG). The alert messages are written when Adabas detects that a PLOG or CLOG data set has not been copied within a specified time interval. Thereafter, an alert message is written at regular intervals as long as there is a PLOG or CLOG data set waiting to be copied. These alert messages are issued until the PLOG or CLOG data set is copied or until no logs remain to be copied. Your user exits 2 or 12 can be set up so that the messages appear providing a comprehensive trace of activity so you can resolve and problems with the PLOG or CLOG.

This feature helps you prevent the potential problem of Adabas overwriting an older CLOG or PLOG data set without waiting for direction from the user. Ordinarily, when a PLOG or CLOG data set fills up, the Adabas nucleus examines the return code from user exit 2 or user exit 12 to determine whether it should wait before overwriting an older PLOG or CLOG data set with new information. If there is no user exit 2 or 12, Adabas does not wait, but simply overwrites the older log. In prior releases of Adabas, no alert messages were issued before this happened. This new feature provides you with these alert messages. If this LOGWARN parameter is set to a value other than zero, the statuses of the CLOG and PLOG data sets are analyzed at the frequency specified by the parameter, warning about their status and providing you with the information you need to prevent any overwriting of older logs.

The production of CLOG and PLOG alert messages is optional and is determined by the setting of a new LOGWARN ADARUN parameter. A value of zero (0), the default, indicates that no CLOG and PLOG status checks should occur and that no alert messages should be produced. Any other valid LOGWARN value represents the interval (in seconds) at which the PLOG and CLOG status is checked and resulting alert messages are produced. When a value other than zero (0) is specified for LOGWARN, a valid user exit 2 or user exit 12 must also be specified. For complete information about the LOGWARN parameter, read *LOGWARN : CLOG and PLOG Alert Message Interval*, in the *Adabas Operations Manual*.

In addition, a new LOGWARN command has been added that allows you to alter the setting of the LOGWARN ADARUN parameter for a given Adabas session. This command can be run from the operator console or via the ADADBS OPERCOM utility function. For complete information, read *LOGWARN Command*, in the *Adabas Operations Manual*.

Printing Multifetch Statistics to a File

Effective with Adabas 8.2 SP2, multifetch statistics can be printed when a CL command is issued if both of the following conditions are met:

- The new PREFSTDD ADARUN parameter is specified in the startup JCL. This new parameter identifies the job control statement label (DD name) that defines the location to which the multifetch statistics should be written.
- The nucleus startup JCL includes a valid job control statement with the label name identified by the PREFSTDD ADARUN parameter.

For complete information, read *PREFSTDD: Multifetch Statistics DD Name*, in the *Adabas Operations Manual*.

LORECX DSECT Updates

In Adabas 8.2 SP2, the fields LOX1STYP and LORSTYPE in the LORECX DSECT (distributed in the Adabas source library) have been altered so that the CLOGs reflect accurate information describing which search algorithm was used.

Logically Deleting and Undeleting Descriptors

Adabas 8.2 SP2 allows you to logically delete and undelete a descriptor from a database file. Logically deleting a descriptor from a file removes the descriptor from the file, but retains the field data in the database. A logically deleted descriptor cannot be used as a search descriptor. This functionality is provided using two new ADADBS utility functions: DELDE (logically delete a descriptor) and UNDELDE (undelete a logically deleted descriptor). For more information, read *DELDE: Logically Deleting a Descriptor* and *UNDELDE: Undeleting a Logically Deleted Descriptor*, in the *Adabas Utilities Manual*.

Logically Deleting and Undeleting File Fields

Adabas 8.2 SP1 allows you to logically delete and undelete a field from a database file. Logically deleting a field from a file removes the field from the FDT, but retains the field data in the database. Logically undeleting a field that was previously logically deleted reinstates the field in the FDT. This functionality is provided using two new ADADBS utility functions: DELFN (delete a field) and UNDELFN (undelete a field). For more information, read *DELFN: Logically Delete Files* and *UNDELFN: Logically Undelete Files*, in the *Adabas Utilities Manual*.

Maximum Length of Fixed Point Fields Increased

The maximum length of fixed point fields (F format fields) was increased to eight bytes in Adabas 8.2 SP2. Fields defined with format F can now be two, four, or eight bytes long.

For more information, read *Field Length and Data Format*, in the *Adabas Concepts and Facilities Manual*.

New Date-Time Edit Masks

New date-time edit masks are introduced in Adabas 8.2 SP1. For a complete reference of these date-time edit masks, read *Date-Time Edit Mask Reference*, in the *Adabas DBA Tasks Manual*.

- A new DT field option is introduced that allows you to specify the edit mask for a field definition. For more information, read *Date-Time Edit Mask Option DT*, in *Adabas Concepts and Facilities Manual*.

This DT option is supported by ADACMP utility FNDEF field definition statements. For more information, read *DT: Date-Time Edit Mask Field Option*, in the *Adabas Utilities Manual*.

- Date-time edit masks are also supported in the format and search buffers of direct calls. You can add date-time edit mask to fields with formats of B, F, P, or U in a loaded file. For more information, read *Date-Time Edit Mask Processing in Format and Search Buffers*, in the *Adabas Command Reference Guide*.

Date-Time Edit Mask Enhancements

The following date-time edit mask enhancements have been made in Adabas 8.2 SP2:

- The new DT option, introduced in 8.2 SP1, is now supported by the ADADBS NEWFIELD function. In ADADBS NEWFIELD functions, the DT option is supported in FNDEF statements, allowing you to specify a date-time edit mask for a new field.

For more information, read *NEWFIELD: Add New Field*, in the *Adabas Utilities Manual*.

- Changes have been made to the ADADBS CHANGE utility function so you can now change the time zone and the date-time edit mask. For complete information, read *CHANGE: Change Standard Length or Format of a Field*, in the *Adabas Utilities Manual*.
- Changes have been made to the ADADBS ONLINVERT utility function so you can now define a date-time field (with edit mask and time zone) as a descriptor field. For complete information, read *ONLINVERT: Start Online Invert Process*, in the *Adabas Utilities Manual*.

Time Zone Support

Adabas 8.2 SP2 supports variations in time zones. The following updates have been made:

- A new library, the Adabas time zone library, is provided in your installation.

A new TZINFO member is provided in your Adabas source library. It lists the time zone support in Adabas and references members of the new time zone library containing the actual code used for time zone conversions. For more information about the Adabas time zone library, refer to the Adabas installation information for the appropriate system platform. For information about the TZINFO library member, read *Supported Time Zones* in the *Adabas DBA Tasks Manual*.

Adabas uses the time zone data taken from the public-domain **tz database**, (also called the *zoneinfo* or *Olson* database) to create its own time zone library.

- A new TZ parameter can now be specified in the ADACMP COMPRESS and DECOMPRESS utility functions to specify the local time zone. As records are compressed and loaded into the file, the date-time data is converted from the specified time zone to UTC time (Coordinated Universal Time, also known as Greenwich Mean Time); likewise as records are decompressed and read from a file, the date-time data is converted from UTC time to the specified time zone. Date-time field data is always stored in UTC time in an Adabas file. For more information, read

COMPRESS: Compress an Adabas File and *DECOMPRESS: Decompress an Adabas File* , in the *Adabas Utilities Manual*.

In addition, review important information about the [daylight savings time enhancements](#) provided in this release. If a time zone uses daylight savings time, you must be sure to store and retrieve the daylight savings indicator with your date-time data or there will be no way to distinguish date-time values in the hour before the time is switched back to standard time.

- A new TZ (time zone) record buffer option is available for the OP command. You can use this option to set the time zone of an individual user session. If no time zone is specified for a user session, response 55 (ADARSP055) subcode 30 is issued when a field with the TZ option is accessed or updated. For more information about using the TZ record buffer option in an OP command, read about the OP command *Record Buffer*, in the *Adabas Command Reference Guide*.
- A time zone field option, TZ, can now be specified to identify date-time fields that should be presented in the user's local time and stored in UTC time, thus allowing for differences in time zones. For more information, read *Time Zone Option TZ*, in *Adabas Concepts and Facilities Manual*.
- Changes have been made to the ADADBS CHANGE utility so you can now change the time zone and the date-time edit mask for a field. For complete information, read *CHANGE: Change Standard Length or Format of a Field* , in the *Adabas Utilities Manual*.
- Changes have been made to the ADADBS NEWFIELD utility so you can now specify the time zone for a new field. For complete information, read *NEWFIELD: Add New Field* , in the *Adabas Utilities Manual*.
- Changes have been made to the ADADBS ONLINVERT utility function so you can now define a date-time field (with edit mask and time zone) as a descriptor field. For complete information, read *ONLINVERT: Start Online Invert Process* , in the *Adabas Utilities Manual*.



Note: The default number of time zones that can be used in a nucleus session is limited to 8 and the time zone pool size default to 32K. However, Software AG allows you to increase these numbers (refer to the appropriate zap in member ZAPOPT of the distributed source library) so that all time zones can be loaded.

Daylight Savings Time Support

Adabas 8.2 SP2 provides daylight savings time support. If the current time zone uses daylight savings time, Adabas provides the following support:

- A new DST parameter can now be specified in the ADACMP COMPRESS and DECOMPRESS utility functions to indicate that date-time data includes a daylight savings time indicator. If a time zone uses daylight savings time, you must be sure to store and retrieve the daylight savings indicator with your date-time data or there will be no way to distinguish date-time values in the hour before the time is switched back to standard time.

The DST parameter requires that the TZ parameter be set in the same ADACMP run.

For more information, read *COMPRESS: Compress an Adabas File* and *DECOMPRESS: Decompress an Adabas File*, in the *Adabas Utilities Manual*.

- A new daylight savings time indicator, *D*, can be specified in format and search buffers to indicate that a daylight savings adjustment is required for a date-time field with the *TZ* option specified. For complete information, read *Daylight Savings Indicator (D)*, in the *Adabas Command Reference Guide*.

RI Commands: Change Possibly Affecting Your Application Programs

Adabas 8.2 SP2 on mainframes no longer releases records unconditionally. If your application issues an RI command for a record that has been updated in the current transaction, Adabas will now return response code 113 (ADARSP113); if ISN=0 was specified, response code 2 (ADARSP002) is returned. This change in behavior may affect how your application programs are coded.

This change was necessary to accommodate Adabas support for **shared hold status** and **partial large object (LOB) field processing**. It prevents problems that might occur if updated records are released from hold, the transaction is subsequently backed out, and there are competitive updates to the same records. Also, this change ensures that Adabas on mainframe and open systems behave in the same manner regarding RI commands.

Application programs that issue RI commands may be affected, if they perform updates and if they try to release updated records from hold status. Such application programs should be adjusted to either not issue RI commands for updated records or suppress the resulting response code 113 (ADARSP113). In addition, if your installation has Natural installed, you must apply Natural zap NA76045 for downward compatibility, which is available using **Empower**.

For complete information about the RI command, read *RI Command: Release Record*, in the *Adabas Command Reference Guide*.

Terminology and Name Changes

The following terminology and name changes have occurred in Adabas 8.2 SP2:

- UEXITB, also known as user exit B (the before or pre-command user exit), has been renamed LUEXIT1 or link user exit 1 in z/OS and z/VSE environments.
- UEXITA, also known as user exit A (the after or post-command user exit), has been renamed LUEXIT2 or link user exit 2 in z/OS and z/VSE environments.

Lowercase Field Name Enhancement

Adabas 8.2 SP1 introduces the ability to create field names that contain lowercase alphabetic characters. With this change, the maximum number of field definitions for a file has increased from 926 to 3214. Thus, fields WM, wm, wM, and Wm are now all valid field names and now represent four unique fields.



Note: Lowercase fields will not display correctly (they will be converted to uppercase) if you use the ADARUN parameter settings MSGCONSL=UPPER, MSGDRUCK=UPPER, or MSGPRINT=UPPER.

Expanded File Chain Enhancements

With Adabas 8.2 SP1 you can dynamically insert and remove files from an expanded file change using a new ADADBS function, EXPFILE. For more information on this utility function, read *EXPFILE: Insert or Remove Files in an Expanded File Chain*, in the *Adabas Utilities Manual*.

Putting Files in Read-Only State

Adabas 8.2 SP1 allows you place specific database files in read-only state while you maintain them. Your database users will be able to read data from the read-only files while your system administrators perform maintenance on those files. In addition, other files in the database, not in read-only status, can be read and updated as normal.

This support is provided in the form of the following Adabas utility enhancements:

- A new READONLY parameter has been added to the ADASAV RESTONL FILES, ADASAV RESTONL FILES, ADASAV RESTONL GCB, ADASAV RESTORE FILES, ADASAV RESTORE FMOVE, ADASAV RESTORE GCB, ADAORD STORE, and ADALOD LOAD utility functions. READONLY can be used to change the read-only status of a database file.
- A new FILEREADONLY parameter has been added to the ADADBS MODFCB utility function. FILEREADONLY can be used to indicate whether an Adabas database file should be placed in readonly status (where it can be updated only by Adabas utilities) or in normal status (where it can be updated in any normal manner).



Notes:

1. If the `FILEREADONLY` parameter is specified, only the `FILE` parameter (which is required) and the optional `PASSWORD` and `TEST` parameters can be specified in the same `ADADBS MODFCB` run.
2. The `ADADBS DSREUSE`, `ISNREUSE`, `NEWFIELD`, `RELEASE DE`, and `RENAME` utility functions do not need to lock the file for their use; these functions can perform their processing in parallel with active users. This means that you do not need to set a file in read-only status (`FILEREADONLY=YES`) if you will be performing these other `ADADBS` functions only.

For complete information, review the descriptions of the `READONLY` and the `FILEREADONLY` parameter as they are used by each these utility functions in the *Adabas Utilities Manual*.

Database Conversion Improvements

You now no longer need to run multiple `ADACNV` utility runs to convert or revert an Adabas database from one version to another. Adabas 8.2 SP1 includes an updated `ADACNV` utility that will convert or revert a database in a single step.

For complete information, read *ADACNV Utility: Database Conversion*, in the *Adabas Utilities Manual*.

Unique Descriptor (UQ) Option Removal

Adabas 8.2 SP1 introduces the ability to remove the UQ (unique descriptor) option from a descriptor field definition using the `ADADBS CHANGE` utility function. A new option, `NOUQ`, is introduced for the `OPTION` parameter. When you run the utility specifying `OPTION=NOUQ` for a descriptor field, the UQ setting is removed from the descriptor field.

For general information about unique descriptor fields, read *Descriptor Options DE, UQ, and XI*, in *Adabas Concepts and Facilities Manual*. For specific information about the new `NOUQ` option in the `ADADBS CHANGE` utility function, read *CHANGE: Change Standard Length or Format of a Field*, in the *Adabas Utilities Manual*.

Work Area Monitoring

Adabas 8.2 SP1 allows you to monitor the use of parts 1, 2, and 3 of the Work data set. This information allows you to proactively adjust the sizes of these areas and ensure 24-by-7 access to your databases. It also allows you to reduce the risk of receiving response code 9 (ADARSP009), subcode 15 errors while using your Adabas database.

This support is available using the DRES operator command. This support is also available via the ADADBS OPERCOM function. The information is also provided in Adabas Online System and Adabas Manager screens.



Note: The maximum pool value of Work part 1 is derived from the LP parameter. It corresponds to the maximum number of blocks a transaction can spend on Work Part 1 before Adabas decides to back it out.

For more information about this new information, read *DRES Command*, in the *Adabas Operations Manual*.

PLOG and CLOG Data Set Enhancements

Adabas 8.2 SP1 introduces several enhancements to its handling of protection log (PLOG) and command log (CLOG) data sets:

- A new NPLOGBUFFERS ADARUN parameter has been introduced that you can use to increase the number of PLOG protection I/O buffers, allowing multiple PLOG blocks to be written in a single write I/O to the PLOG and improving the performance of applications that were formerly constrained by writes to the PLOG. For more information, read *NPLOGBUFFERS : PLOG Protection I/O Buffer Control*, in the *Adabas Operations Manual*.
- The ADARES CLOG and ADARES PLOG utility functions have been expanded to check the PPT at the end of the CLOG or PLOG run to determine whether any additional CLOG or PLOG data sets need to be copied. If so, the ADARES utility invokes user exit 2 or 12, as appropriate, for each uncopied CLOG or PLOG data set it detects. For complete information, read *CLCOPY: Copy Dual Command Log* and *PLCOPY: Copy Protection Log to Sequential Data Set*, in the *Adabas Utilities Manual*.
- The ADADBS utility introduces four new functions (ADDCLOG, ADDPLOG, DELCLOG, and DELPLOG) that you can use to dynamically add and delete CLOG and PLOG data sets without terminating your current nucleus session. Running these utility functions is invalid when Adabas is running with dual CLOGs or PLOGs.

For complete information, read *ADADBS Utility: Database Services*, in the *Adabas Utilities Manual*.

- You can now log the extended I/O list to the CLOG data sets, for both CLOGLAYOUT=5 and CLOGLAYOUT=8. A new ADARUN parameter, LOGVOLIO, is provided so you can specify this functionality for an Adabas nucleus.

You can also use two new operator commands to dynamically control this logging.

1. Use the new LOGVOLIO operator command on the console (or in an ADADBS OPERCOM utility function run) to dynamically start logging of the extended I/O list. For more information, read *LOGVOLIO Command*, in the *Adabas Operations Manual*.
 2. Use the new NOLOGVOLIO operator command on the console (or in an ADADBS OPERCOM utility function run) to dynamically stop logging the extended I/O list. For more information, read *NOLOGVOLIO Command*, in the *Adabas Operations Manual*.
- A new ADADBS utility function, REACTLOG, now allows you to reactivate command logging in an active nucleus where it had been disabled previously as a result of an I/O error. The cause of the I/O error needs to be corrected before running this utility function or command logging will simply fail again and will not be reactivated.

For complete information, read *REACTLOG: Reactivating Command Logging*, in the *Adabas Utilities Manual*.

- A new SRLOG ADARUN parameter allows you to indicate how spanned records should be logged to the protection logs. Entire records (ALL), modified segments of spanned records (UPD), or entire records up to and including the modified segments (PART) can be logged. The default is UPD. For more information, read *SRLOG : Controlling Spanned Record Logging*, in the *Adabas Operations Manual*.

Spanned Record Logging

Adabas 8.2 SP1 allows you to write complete spanned records to the PLOGs. In the past, spanned records were only partly written to the PLOGs, making them impossible to decompress and use.

- A new SRLOG ADARUN parameter allows you to indicate how spanned records should be logged to the protection logs. Entire records (ALL), modified segments of spanned records (UPD), or entire records up to and including the modified segments (PART) can be logged. The default is UPD. For more information, read *SRLOG : Controlling Spanned Record Logging*, in the *Adabas Operations Manual*.

The SRLOG parameter replaces the PLOGSEGSDS ADARUN parameter; the PLOGSEGSDS parameter has been dropped in this release.

- The ADASEL utility has been enhanced with three new SET GLOBALS parameters to support the decompression of spanned records from the PLOG: SPANREC, MAXLOGRECLLEN, and LWP. Each of these parameters is described in *SET GLOBALS Parameter*, in the *Adabas Utilities Manual*.

In addition, a SPANREC parameter has been added to the ADASEL OUTPUT instruction and can be used in an ADASEL OUTPUT instruction to indicate that new spanned record SELH and SELC output headers should be used for all EXPAN output if there is the possibility that the output decompressed spanned records will exceed the physical record length limitation. DSECTs for the SELH and SELC headers can be found in the Adabas source library. For complete information, read *OUTPUT Instruction*, in the *Adabas Utilities Manual*.

- The ADACDC utility has been enhanced with two new parameters: SPANREC and MAXLOGRELEN. Each of these parameters is described in *Running the Utility*, in the *Adabas Utilities Manual*.

The SPANREC parameter indicates that new spanned record CDCH and CDCN output headers should be used for all CDCOUT output if there is the possibility that the output decompressed spanned records will exceed the physical record length limitation. DSECTs for the CDCH and CDCN headers can be found in the Adabas source library.

New Adabas Event Log Support

Adabas 8.2 SP1 introduces a new Adabas event log that can be used to log database events. In this release, only response code 145 (ADARSP145) events are logged to the Adabas event log. Each entry in the log is currently 128 bytes long, although this may change in future releases.

The size of the Adabas event log for a database is controlled by the INFOBUFFERSIZE ADARUN parameter. In addition, the entries in the Adabas event log can be displayed using a new ADADBS DEVENTLOG utility function.

Response Code 145 (ADARSP145 - Hold Queue Entry) Event Enhancements

Several enhancements have been made in Adabas 8.2 SP1 to support response code 145 (ADARSP145) events. Response code 145 (ADARSP145) occurs when a hold queue entry is not available for an N1 or N2 command or when a command is issued that attempts to put a record on hold but cannot because the record is held by some other user. In this release of Adabas, events triggering response code 145 (ADARSP145) are now logged in the:

- Adabas event log, if an Adabas event log has been allocated. The Adabas event log size and allocation is controlled by the INFOBUFFERSIZE ADARUN parameter for a database.
- command log, if CLOGLAYOUT=8 is specified. A new record type (x'000D') is now produced in a CLOGLAYOUT=8 command log to accommodate response code 145 (ADARSP145) events.

In addition, a special L2 command call can now be issued if you receive a response code 145(ADARSP145). When you execute the L2 command with ISN=0, FNR=-4(X'FFFC), and an appropriate format buffer, you can obtain information about the user currently holding the ISN when

the response code was issued. For complete information on this call, read *Special L2 Command Call for Response 145*, in the *Adabas Command Reference Guide*.

Multiple CICS Task-Related User Exit (TRUE) Support

Adabas 8.2 SP1 supports the installation of multiple CICS task-related user exits (TRUEs) and Adabas application stubs from a single execution of the ADACIC0 installation program. Multiple TRUEs allow your site to tailor different Adabas CICS execution options in the same CICS region with a centralized installation procedure and software. For complete information, read *Installing Adabas with CICS under Adabas 8*, in the *Adabas z/OS Installation Guide*.

CICS Installation Changes

Additional CICS installation changes have occurred in Adabas 8.2 SP1. We no longer deliver ADACIRQ with Adabas on z/OS and z/VSE systems. The CICS input transient data queue is no longer required and the DDLINK input is no longer required. These have been replaced by the Adabas CICS installation options table (member ACIOPT) and the Adabas CICS names module (member ACINAMES). For more information, read *Installing Adabas with CICS under Adabas 8*, in the *Adabas z/OS Installation Guide*.

Performance Enhancements

Various updates have been made in Adabas 8.2 SP1 to ensure it continues to operate at peak performance:

- The performance (CPU consumption) of opening new Adabas user sessions, when there are already thousands or tens of thousands of concurrent users, has been improved.
- The performance of large object (LOB) updating has been improved by eliminating redundant serialization. All nuclei now reduce their use of the unique descriptor pool for LOB updates and customers using LOBs from Adabas 8.1 may be able to reduce the unique descriptor pool size (ADARUN parameter LDEUQP). Adabas Cluster Services and Adabas Parallel Services configurations no longer acquire external unique descriptor locks for LOB updates.
- Design changes have been made to the nucleus so that the blocks of a new address converter (AC) extent are now formatted asynchronously rather than synchronously. This allows a database to continue processing without waiting for all of the new AC extent blocks to be formatted.
- The new V64BIT ADARUN parameter may be used to improve the performance of your Adabas nucleus sessions. V64BIT allows the nucleus to use virtual storage above the 2 gigabyte bar, freeing virtual storage below the bar for other uses.

- For systems running in 64-bit mode (ADARUN parameter V64BIT is set to YES), the new ADARUN parameter `LARGEPAGE` allows the nucleus to use large pages, which may particularly improve the performance of long-running, memory-access intensive applications.
- The new `NPLOGBUFFERS` ADARUN parameter can be used to increase the number of PLOG protection I/O buffers, allowing multiple PLOG blocks to be written in a single write I/O to the PLOG and improving the performance of applications that were formerly constrained by writes to the PLOG.
- The new `NWORK1BUFFERS` ADARUN parameter can be used to increase the number of Work part 1 protection I/O buffers, allowing multiple Work part 1 blocks to be written in a single write I/O to the Work data set and improving the performance of applications that were formerly constrained by writes to Work part 1.
- An internal Software AG Vector Table (SAGVT) is introduced in this release for z/OS systems. This is an internal z/OS CSA structure anchored in a CVT-based field assigned by IBM to Software AG. The vector table is used internally by Adabas to identify the location of an internal SVC table Adabas now uses to store and maintain SVC numbers and their corresponding IDT addresses.

In prior releases, this information was placed in the internal SSCT entry created every time an SVC is installed. To locate the IDT in former releases, Adabas was required to sequentially scan all of the SSCTs. With Adabas 8.2, when an SVC is installed, its IDT information is stored in the SVC table identified by the vector table. To locate the IDT associated with the SVC, Adabas need only scan the internal SVC table, thus eliminating the more time-intensive SSCT scan.

For information on the Adabas SVC, read *Initializing the Adabas Communication Environment in z/OS* in the *Adabas z/OS Installation Guide*.

- Two new `LGBLSET` parameters have been added, `UBTYPE` and `UBSTIME`, that can be used to:
 - obtain a single user block element per CICS task (instead of a user block pool) and
 - set a time interval to clean up locked user blocks.

These new parameters help eliminate abends that can occur when CICS transactions lock the user blocks but do not release them. For more information about the `LGBLSET` macro, read *Modifying Source Member Defaults (LGBLSET Macro) in Version 8*, in the *Adabas z/OS Installation Guide*.

For complete information on the ADARUN parameters mentioned in this section, read *Adabas Initialization (ADARUN Statement)*, in the *Adabas Operations Manual*.

Nucleus 64-Bit Processing Support

Adabas 8.2 SP1 introduces support for 64-bit processing. If your z/OS operating system supports 64-bit virtual storage, you can now request that an Adabas nucleus use this virtual storage above the 2 gigabyte bar. A new ADARUN parameter, *V64BIT*, allows you to indicate whether your Adabas nucleus should use virtual storage above the 2 gigabyte bar. For more information, read *V64BIT : 64-Bit Virtual Storage Option*, in the *Adabas Operations Manual*.

Adabas Buffer Description (ABD) Structure Changes

Updates have been made in Adabas 8.2 SP1 to support 64-bit and ALET-qualified buffers in Adabas calls. The structure of Adabas buffer descriptions (ABDs) has changed to accommodate this support. Two reserved fields (ABDXRSV4 and ABDXRSV5) have been replaced with new DSECT fields ABDXALET and ABDXADRG respectively.

- ABDXALET is now used to store the ALET for the buffer (if ABDXLOC='C'D').
- ABDXADRG is now used to store the 64-bit indirect address pointer (if ABDXLOC='C'I' or 'C'D').

For complete information, read *Adabas Buffer Descriptions (ABDs)*, in the *Adabas Command Reference Guide*.

Large Page Support

Adabas 8.2 SP1 introduces large page support for z/OS systems. If your z/OS operating system supports 64-bit virtual storage and large pages (1 MB pages of real storage), you can request that the Adabas nucleus use large pages. A new ADARUN parameter, *LARGE PAGE*, allows you to indicate whether your Adabas nucleus should use large pages or not. The default is not to use large pages. For more information, read *LARGE PAGE : Large Page Support Option*, in the *Adabas Operations Manual*.

z/OS SVC Dump Enhancements

The z/OS SVC dump support in Adabas 8.2 SP1 has been enhanced to support the changes made by IBM to IEATDUMP in a recent z/OS release. z/OS users running non-authorized should be able to use the transaction dump IEATDUMP service in a similar way to authorized users with SVC-DUMP.

For more information, read *Using the z/OS SVC Dump Facility*, in the *Adabas Operations Manual*.

z/OS SVC Authorization Enhancements

Adabas 8.2 SP1 introduces support for the System Authorization Facility (SAF) in providing authorization for the Adabas 8 SVC in z/OS environments. SAF can now be used in some circumstances to ensure that the Adabas nucleus, as well as other MPM servers (such as Entire Net-Work and the Natural Global Buffer Pool) are authorized to prevent inappropriate use of critical ADASVC functions. There are still some circumstances where APF authorization is required instead. For complete information, read *Authorization Requirements*, in the *Adabas z/OS Installation Guide*.

z/OS SVC Prelinked to SAF

Adabas 8.2 SP1 no longer requires that you link the Adabas SVC (ADASVC) with SAF (SVCSAF) during installation. In Adabas 8.2, the SVC is prelinked to SAF and the SVCSAF WAL library member is no longer supplied.

Reporting on Replication-Related Records in Work Part 1

Adabas 8.2 SP1 provides ADAWRK utility support for replication-related records written to Work part 1. Prior versions of the ADAWRK utility did not provide this support. You can now use the ADAWRK utility to report on replication-related records so you can determine how to handle the data when the database autostart fails and the database will not come up.

For complete information on replication support in the ADAWRK utility, read *ADAWRK Utility: Work Area Recovery Reports*, in the *Adabas Utilities Manual*

Printing Adabas Maintenance and SVC Status Information

A new utility, ADAZIN, is introduced in Adabas 8.2 SP1. ADAZIN can be used to print maintenance information about Adabas load modules and status information about the Adabas SVC on the system in which ADAZIN is run. Names of target load modules and SVC numbers can be specified to the utility to limit the range of the printed report.

For complete information, read *ADAZIN Utility: Print Adabas Maintenance and SVC Information*, in the *Adabas Utilities Manual*.

Replication of Security Definitions

Adabas 8.2 SP1 allows you to replicate security definitions if you have Event Replicator for Adabas 3.2 or later installed. For more information, refer to your Event Replicator for Adabas documentation.

Address Converter Extent Handling Enhancement

With Adabas 8.2 SP1, contiguous address converter (AC) blocks are added to the last existing AC extent, if possible, rather than creating a new AC extent. However, if contiguous blocks are not available, a new AC extent is created.

This enhancement was made in order to reduce the number extents defined in the File Control Block (FCB). At this time, it affects only your Adabas nuclei. The Adabas utilities do not support this update at this time.

Adabas Delta Save Update: ADASAV MERGE Limitation Lifted

Adabas 8.2 SP1 includes internal enhancements that allow Adabas Delta Save to handle overlapping delta save numbers in ADASAV MERGE processing. In prior releases, the merge of two delta save data sets with overlapping delta save numbers would terminate abnormally with the message:

DSF065 Input Full and Delta Save tapes not fitting together

This enhancement allows data sets with DSIDs that share a delta save number to be merged. For example, the save tape with DSID *7/1-5/timestamp1* can now be merged successfully with a save tape with DSID *7/5-9/timestamp2*. The fact that overlapping delta save number 5 exists on both tapes and in both DSIDs can now be handled successfully during the merge.

For complete information about the ADASAV MERGE utility function, refer to your Adabas Delta Save documentation.

New and Changed Utility Functions

Adabas 8.2 introduces the following utility enhancements:

Utility Function	New or Changed?	Enhancement Description
ADACDC	Changed	<p>Spanned records can now be fully decompressed and read from the PLOG. In support of this, two new parameters have been added: SPANREC and MAXLOGRECLLEN. Each of these parameters is described in <i>Running the Utility</i>, in the <i>Adabas Utilities Manual</i>.</p> <p>The SPANREC parameter indicates that new spanned record CDCH and CDCN output headers should be used for all CDCOUT output if there is the possibility that the output decompressed spanned records will exceed the physical record length limitation. DSECTs for the CDCH and CDCN headers can be found in the Adabas source library.</p>
ADACMP	Changed	<p>The ADACMP utility has been updated to support the new date-time edit masks introduced in this version of Adabas. To support these date-time edit masks in field definitions, a new field option, DT, has been added to FNDEF field and group field definition statements. this field option allows you to specify a date-time edit mask for the field.</p>
ADACMP COMPRESS	Changed	<p>The following new parameters are supported in this release:</p> <ul style="list-style-type: none"> ■ A new TZ parameter allows you to specify the local time zone. As records are compressed and loaded into the file, the date-time data is converted from this time zone to UTC time. Date-time field data is always stored in UTC time in an Adabas file. ■ A new DST parameter can now be specified to indicate that date-time data includes a daylight savings time indicator. If a time zone uses daylight savings time, you must be sure to store the daylight savings indicator with your date-time data or there will be no way to distinguish date-time values in the hour before the time is switched back to standard time. <p>The following new field options are supported in this release in FNDEF field definitions:</p>

Utility Function	New or Changed?	Enhancement Description
		<ul style="list-style-type: none"> ■ A new SY option allows you to indicate that the field is a system field and identify the type of system field being defined. ■ A new CR option can be used for system fields (defined with the SY option) to indicate that a system field value should only be set when a record is inserted into a file and not when it is updated. Without the CR option, system field values are set for both record insertions and deletions. The CR option may not be specified for fields defined with the MU option. ■ A new TZ option can now be used to indicated that a date-time field that should be presented in the user's local time and stored in UTC time, allowing for differences in time zones.
ADACMP DECOMPRESS	Changed	<p>The following new parameters are supported in this release:</p> <ul style="list-style-type: none"> ■ A new TZ parameter allows you to specify the local time zone. As records are decompressed and read from the file, the date-time data is converted from UTC to a corresponding time in this time zone. Date-time field data is always stored in UTC time in an Adabas file. ■ In addition, a new DST parameter can now be specified to indicate that date-time data includes a daylight savings time indicator. If a time zone uses daylight savings time, you must be sure to retrieve the daylight savings indicator with your date-time data or there will be no way to distinguish date-time values in the hour before the time is switched back to standard time.
ADACNV	Changed	<p>The ADACNV utility has been updated so that you can now convert and revert Adabas databases between releases in <i>one step</i>, regardless of how many releases have occurred in between the two versions. The version of the earliest Adabas databases that can be converted using this utility is 6.1. Likewise, the earliest version of Adabas database to which you can revert a more current database is 6.1.</p> <p>In addition, during ADACNV processing, all changed RABNs are logged to the Adabas Delta Save area (in case Adabas Delta Save was enabled before the conversion).</p> <p>Finally, it is no longer necessary to make a full backup before and after a conversion; a backup of only the changed database blocks (an Adabas Delta Save ADASAV SAVE DELTA run) is sufficient.</p>
ADADBS ADDCLOG	New	<p>This new ADADBS utility function allows you to dynamically add command log (CLOG) data sets without terminating your current nucleus session. Running this utility function is invalid when Adabas is running with dual CLOGs.</p>
ADADBS ADDPLOG	New	<p>This new ADADBS utility function allows you to dynamically add protection log (PLOG) data sets without terminating your current nucleus session. Running this utility function is invalid when Adabas is running with dual PLOGs.</p>

Utility Function	New or Changed?	Enhancement Description
ADADBS CHANGE	Changed	<p>The following ADADBS CHANGE updates have been made:</p> <ul style="list-style-type: none"> ■ You can now use this utility function to remove the UQ (unique descriptor) option from a descriptor field definition. A new option, NOUQ, has been added for the OPTION parameter of the ADADBS CHANGE function. ■ You can specify an altered Adabas field definition for a field using a new FNDEF parameter. Changes to the time zone and the date-time edit mask are included.
ADADBS DELCLOG	New	This new ADADBS utility function allows you to dynamically delete command log (CLOG) data sets without terminating your current nucleus session. Running this utility function is invalid when Adabas is running with dual CLOGs.
ADADBS DELDE	New	This new ADADBS utility function allows you to logically delete a descriptor from a database file.
ADADBS DELFN	New	This new ADADBS utility function allows you to logically delete a field from a database file.
ADADBS DELPLOG	New	This new ADADBS utility function allows you to dynamically delete protection log (PLOG) data sets without terminating your current nucleus session. Running this utility function is invalid when Adabas is running with dual PLOGs.
ADADBS DSREUSE	Changed	In prior releases, Adabas locked the file associated with ADADBS DSREUSE processing. With this release, enhancements to Adabas have removed the necessity of performing this lock.
ADADBS DEVENTLOG	New	This new ADADBS utility function allows you to display the Adabas event log of a database. In a cluster environment, this function can display the Adabas event log of an individual database or of every database in the cluster.
ADADBS EXPFILE	New	This new utility function can be used to dynamically insert or remove files in expanded file chains.
ADADBS ISNREUSE	Changed	In prior releases, Adabas locked the file associated with ADADBS ISNREUSE processing. With this release, enhancements to Adabas have removed the necessity of performing this lock.
ADADBS MODFCB	Changed	<p>The following updates have been made to the ADADBS MODFCB function:</p> <ul style="list-style-type: none"> ■ A new FILEREADONLY parameter has been added that can be used to indicate whether an Adabas database file should be placed in readonly status (where it can be updated only by Adabas utilities) or in normal status (where it can be updated in any normal manner). <p>Note:</p> <ol style="list-style-type: none"> 1. If the FILEREADONLY parameter is specified, only the FILE parameter (which is required) and the optional PASSWORD and TEST parameters can be specified in the same ADADBS MODFCB run. 2. The ADADBS DSREUSE, ISNREUSE, NEWFIELD, RELEASE DE, and RENAME utility functions do not need to lock the file for their use; these

Utility Function	New or Changed?	Enhancement Description
		<p>functions can perform their processing in parallel with active users. This means that you do not need to set a file in read-only status (FILEREADONLY=YES) if you will be performing these other ADADBS functions only.</p> <ul style="list-style-type: none"> ■ A new SYFMAXUV parameter is supported to specify the maximum number of occurrences for MU system field values that can be stored in a file. ■ In addition in prior releases Adabas locked the file associated with ADADBS MODFCB processing. With this release, enhancements to Adabas have removed the necessity of performing this lock unless the LOBFILE or FILEREADONLY parameters are specified.
ADADBS NEWFIELD	Changed	<p>The DT option is supported in FNDEF statements, allowing you to specify a date-time edit mask and a time zone for a new field. In addition, the system field options SY and CR are supported in FNDEF statements, allowing you to define system fields.</p> <p>In addition, in prior releases Adabas locked the file associated with ADADBS NEWFIELD processing. With this release, enhancements to Adabas have removed the necessity of performing this lock.</p>
ADADBS ONLINVERT	Changed	<p>The FNDEF parameter has been added to the ADADBS ONLINVERT utility function. Using this function you can specify a standard field definition for a new file descriptor. In addition, this new parameter allows you to define a date-time field (with edit mask and time zone) as a descriptor.</p> <p>For complete information, read <i>ONLINVERT: Start Online Invert Process</i>, in the <i>Adabas Utilities Manual</i>.</p>
ADADBS OPERCOM	Changed	<p>You can now use two new operator commands to dynamically control logging of the extended I/O list for both CLOGLAYOUT=5 and CLOGLAYOUT=8.</p> <ol style="list-style-type: none"> 1. Use the new LOGVOLIO operator command to dynamically start logging of the extended I/O list. 2. Use the new NOLOGVOLIO operator command to dynamically stop logging the extended I/O list. <p>For more information, read <i>Operator Commands</i>, in the <i>Adabas Utilities Manual</i>.</p>
ADADBS REACTLOG	New	<p>This new ADADBS utility function allows you to reactivate command logging in an active nucleus where it had been disabled previously as a result of an I/O error. The cause of the I/O error needs to be corrected before running this utility function or command logging will simply fail again and will not be reactivated.</p>
ADADBS RELEASE DESCRIPTOR	Changed	<p>In prior releases, Adabas locked the file associated with ADADBS RELEASE DESCRIPTOR processing. With this release, enhancements to Adabas have removed the necessity of performing this lock.</p>
ADADBS RENAME	Changed	<p>In prior releases, Adabas locked the file associated with ADADBS RENAME processing. With this release, enhancements to Adabas have removed the necessity of performing this lock.</p>

Utility Function	New or Changed?	Enhancement Description
ADADBS UNDELDE	New	This new ADADBS utility function allows you to undelete a logically deleted descriptor from a database file.
ADADBS UNDELFN	New	This new ADADBS utility function allows you to logically undelete a field that had previously been logically deleted from a database file using the ADADBS DELFN function.
ADALOD LOAD	Changed	<p>The following updates have been made to the ADALOD LOAD function:</p> <ul style="list-style-type: none"> ■ A new READONLY parameter has been added that allows you to change the read-only status of a file. ■ A new SYFMAXUV parameter is supported to specify the maximum number of occurrences for MU system field values that can be stored in a file. ■ The following new ADALOD LOAD parameters allow you to manage possible ADALOD LOAD replication errors. These parameters should only be used if you also have Event Replicator for Adabas installed. <ul style="list-style-type: none"> ■ The new RPLINITERROR parameter allows you to indicate whether ADALOD processing should continue if replication initialization errors occur. ■ The new RPLERRORDEACTFILE parameter allows you to indicate whether replication should be deactivated for the file if replication errors occur. <p>Caution: If you elect to have ADALOD LOAD continue its processing after a replication error, you are also responsible for recovering your environment appropriately.</p>
ADAORD STORE	Changed	<p>Four new replication-oriented parameters have been added to the ADAORD STORE function: RPLDSBI, RPLKEY, RPLTARGETID, and RPLUPDATEONLY. These parameters are added to support replication of data during a restore.</p> <p>In addition, a new READONLY parameter has been added that allows you to change the read-only status of the file.</p> <p>Note: These parameters can only be specified if you also have Adabas 8.2 SP2 or later installed.</p>
ADAREP	Changed	<p>You can now produce an XML document containing the ADAREP database status report in XML format. This XML document is structured as defined by an XML schema definition (XSD) file provided with Adabas. You can transfer the XML document to any system (probably Windows or UNIX) using automated tools or user-written programs.</p> <p>The new OUTPUT parameter allows you to request this XML document. A new ACODE parameter allows you to specify the code page used for the XML output.</p> <p>In addition, parallel participant table (PPT) information is now included in the ADAREP report by default. PPT information can be very useful when an abnormal outage has occurred for the database and you need to determine the</p>

Utility Function	New or Changed?	Enhancement Description	
		<p>physical state of the database and its existing checkpoint details. You can use a new NOPPT parameter to indicate that the PPT information should <i>not</i> be included in the ADAREP report.</p>	
ADARES BACKOUT	Changed	<p>The ADARES BACKOUT utility functions have been updated to support the replication of the protection log data set sent to the nucleus. Using the new RPLDATA parameter, you can now activate this replication.</p> <p>In addition, additional statistics that list the number of records backed out for each file in an ADARES BACKOUT run are listed when processing is complete.</p>	
ADARES CLCOPYY	Changed	<p>The ADARES CLCOPYY function has been expanded. It now checks the PPT at the end of the CLCOPYY job to determine whether any additional CLOG data sets need to be copied. If so, the ADARES utility invokes user exit 2 or user exit 12, as appropriate, to accommodate the number of data sets that need copying. For example, if NCLOG=8, once the initial CLCOPYY job completes, the ADARES utility will issue a call to the nucleus to invoke user exit 2 or 12 for each uncopied CLOG data set it detects.</p>	
ADARES PLCOPY	Changed	<p>The ADARES PLCOPYY function has been expanded. It now checks the PPT at the end of the PLCOPYY job to determine whether any additional PLOG data sets need to be copied. If so, the ADARES utility invokes user exit 2 or user exit 12, as appropriate, to accommodate the number of data sets that need copying. For example, if NPLOG=8, once the initial PLCOPYY job completes, the ADARES utility will issue a call to the nucleus to invoke user exit 2 or 12 for each uncopied PLOG data set it detects.</p>	
ADARES REGENERATE	Changed	<p>The ADARES REGENERATE utility functions have been updated to support the replication of the protection log data set sent to the nucleus. Using the new RPLDATA parameter, you can now activate this replication.</p> <p>In addition, additional statistics that list the number of records regenerated for each file in an ADARES REGENERATE run are listed when processing is complete.</p>	
ADASAV RESTONL FILES	Changed	<p>You can now restore a range of files using the ADASAV functions. As before, you can still restore individual or lists of files. To specify a range of files in the FILES or FMOVE parameters of the various ADASAV functions, use a dash (-) between the first and last file number in the range, in the format:</p> <p><i>fnfirst-fnlast.</i></p> <p>Six new replication-oriented parameters have been added to various ADASAV functions: RPLACTIVE, RPLDATA, RPLDSBI, RPLKEY, RPLTARGETID, and RPLUPDATEONLY. These parameters are added to support replication of data during a restore.</p> <p>Note: These parameters can only be specified if you also have Adabas 8.2 SP2 or later installed.</p> <p>In addition, a new READONLY parameter has been added to these functions that allows you to change the original read-only status of the file.</p>	
ADASAV RESTONL FMOVE			
ADASAV RESTONL GCB			
ADASAV RESTORE FILES			
ADASAV RESTORE FMOVE			

Utility Function	New or Changed?	Enhancement Description
ADASAV RESTORE GCB		
ADASAV MERGE	Changed	<p>This release of Adabas includes internal enhancements that allow Adabas Delta Save to handle overlapping delta save numbers in ADASAV MERGE processing. In prior releases, the merge of two delta save data sets with overlapping delta save numbers would terminate abnormally with the message:</p> <p style="background-color: #e0e0e0; padding: 2px;"><code>DSF065 Input Full and Delta Save tapes not fitting together</code></p> <p>This enhancement allows data sets with DSIDs that share a delta save number to be merged. For example, the save tape with DSID <code>7/1-5/timestamp1</code> can now be merged successfully with a save tape with DSID <code>7/5-9/timestamp2</code>. The fact that overlapping delta save number 5 exists on both tapes and in both DSIDs can now be handled successfully during the merge.</p> <p>For complete information about the ADASAV MERGE utility function, refer to your Adabas Delta Save documentation.</p>
ADASEL	Changed	<p>Several new parameters and keywords have been added to the ADASEL utility in this release to enhance its flexibility and functionality:</p> <ul style="list-style-type: none"> ■ A new TEST parameter allows you to test the syntax of your ADASEL utility run without actually processing any PLOG records. ■ A new FDTINPUT parameter allows you to specify an alternate FDT source to be used during ADASEL processing. ■ A new FROM USER clause has been added to the SELECT parameter to allow you to select PLOG records by user or terminal ID, regardless of file number. This new clause is mutually exclusive with the FROM [FILE] and IN [FILE] clauses. <p>In addition, a new NU ADASEL global parameter has been added that allows you to globally limit the number of user or terminal IDs processed by a SELECT parameter. The default is 20. ADASEL will allocate work storage for a run based on the setting of the NU global parameter.</p> <ul style="list-style-type: none"> ■ A new DISPLAY instruction option, ALL, allows you to display all fields of each selected record, including MU and PE group fields and their occurrence counts. ■ Two new DISPLAY instruction options, BEFORE and BOTH, now allow you accurately indicate whether before, after, or before and after images of selected data should be output. BEFORE and BOTH are both optional and are mutually exclusive options; only one of them may be specified in a DISPLAY instruction, but neither of them is required. If BOTH is specified, both the before and after images are output; if BEFORE is specified, only the before images are output; if neither BEFORE nor BOTH is specified, only after images are output. ■ A new DISPLAY instruction option, COUNT, now allows you to display the number of occurrences of an MU field or PE group.

Utility Function	New or Changed?	Enhancement Description
		<p>■ Spanned records can now be fully decompressed and read from the PLOG. In support of this, three new SET GLOBALS parameters have been added: SPANREC, MAXLOGRECLLEN, and LWP. Each of these parameters is described in <i>SET GLOBALS Parameter</i>, in the <i>Adabas Utilities Manual</i>.</p> <p>In addition, a SPANREC parameter has been added to the ADASEL OUTPUT instruction and can be used in an ADASEL OUTPUT instruction to indicate that new spanned record SELH and SELC output headers should be used for all EXPAN output if there is the possibility that the output decompressed spanned records will exceed the physical record length limitation. DSECTS for the SELH and SELC headers can be found in the Adabas source library. For complete information, read <i>OUTPUT Instruction</i>, in the <i>Adabas Utilities Manual</i>.</p> <p>In addition, ADASEL error messages have been enhanced so they are more self-explanatory.</p> <p>In prior releases of Adabas, the IGNORE instruction (that can be specified in the THEN and ELSE clauses of an IF statement) would not actually ignore the before-image of an item. The before-image was output even though IGNORE was specified. In this release of Adabas, neither the before or after-images are output when IGNORE is specified.</p>
ADAWRK	Changed	<p>A new REPLICATION parameter has been added to the ADAWRK utility that provides you with information on replication-related records in Work part 1. This parameter determines whether a new Replication report and new Replication Summary report are produced by the utility and the presence of additional replication statistics on the existing reports.</p> <p>In addition to the replication support, the ADAWRK utility now assigns a transaction sequence number to each transaction it reads from Work part 1, to aid in uniquely identifying and matching the transactions in the different ADAWRK reports</p>
ADAZIN	New	<p>A new utility, ADAZIN, is introduced in this release. ADAZIN can be used to print maintenance information about Adabas load modules and status information about the Adabas SVC on the system in which ADAZIN is run. Names of target load modules and SVC numbers can be specified to the utility to limit the range of the printed report.</p>

New and Changed Operator Commands

Starting with Adabas 8.2, display-type commands from the operator console (such as DSTAT and DTH) will be executed directly if all Adabas threads or Adabas internal command queue elements are in use by other commands. This enhancement allows you to better determine the status of Adabas should a problem occur that impedes normal command processing. Note that this direct execution of display-type commands applies only to commands from the operator console, but not to commands issued by the ADADBS OPERCOM utility function, Adabas Online System, or Adabas Manager. Also, Adabas might be unable to respond to operator commands (issued from the operator console) in certain error situations, such as if Adabas is looping, in a hard wait, or when it is unable to obtain the CPU processing cycles it needs.

The following table lists new and changed operator commands in Adabas 8.2.

Command	New or Changed	Enhancement Description
ALOCKF	Changed	Support for this command, provided for non-cluster environments in prior versions, is now provided in cluster environments.
DRES	Changed	This command now displays information on the high-water marks for Work parts 1, 2, and 3, allowing you to monitor the use of the Work data set.
LICREFRESH	New	This new command allows you to refresh the license file currently used by the nucleus. It allows you to overwrite the current license with a new one or to force the system to reread and reload its existing license.
LOGVOLIO	New	This new command allows you to initiate logging of the extended I/O list to the CLOG data sets, for both CLOGLAYOUT=5 and CLOGLAYOUT=8. This command can be entered on the console and in an ADADBS OPERCOM utility function run.
LOGWARN	New	This new command allows you to alter the setting of the LOGWARN ADARUN parameter for a given Adabas session.
NOLOGVOLIO	New	This new command allows you to stop logging of the extended I/O list to the CLOG data sets, for both CLOGLAYOUT=5 and CLOGLAYOUT=8. This command can be entered on the console and in an ADADBS OPERCOM utility function run.
NWCONNECT	New	This new command allows you to retry establishing the Entire Net-Work DBID target. During nucleus initialization, certain Entire Net-Work errors may have prevented the target from being established, resulting in message ADAM76. Other errors may prevent Entire Net-Work from defining the target when it is started after the nucleus is initialized. Support for the NWCONNECT command in Adabas Cluster Services and Adabas Parallel Services nuclei is provided by Adabas fix AN822079.

Command	New or Changed	Enhancement Description
RALOCKF	Changed	Support for this command, provided for non-cluster environments in prior versions, is now provided in cluster environments.
RALOCKFA	Changed	Support for this command, provided for non-cluster environments in prior versions, is now provided in cluster environments.
SMFDETAIL	New	This new command allows you to alter the complete setting of the SMFDETAIL ADARUN parameter for the running Adabas session.
SMFDETAILADD	New	This new command allows you to add specific detail sections of Interval and Termination records to the list of detail records already specified by the SMFDETAIL ADARUN parameter for the running Adabas session.
SMFDETAILDEL	New	This new command allows you to remove specific detail sections of Interval and Termination records from the list of detail records specified by the SMFDETAIL ADARUN parameter for the running Adabas session.
SMFINTERVAL	New	This new command allows you to override the setting of the SMFINTERVAL ADARUN parameter for the running Adabas session.
SMFRECNO	New	This new command allows you to override the setting of the SMFRECNO ADARUN parameter for the running Adabas session.
SMFSUBSYS	New	This new command allows you to override the setting of the SMFSUBSYS ADARUN parameter for the running Adabas session.

For more information about any Adabas operator command, read *Operator Commands*, in the *Adabas Operations Manual*.

ADARUN Parameter Changes

This section describes the changes made to ADARUN parameters in this release. For more information about any ADARUN parameter, read *Adabas Initialization (ADARUN Statement)*, in the *Adabas Operations Manual*.

The following table summarizes the ADARUN parameter changes in Adabas 8.2.

Parameter	New or Changed	Enhancement Description
CLUCACHEEXTRA	New	Use the new CLUCACHEEXTRA parameter to specify the additional vector size, in blocks, required to allow for the Adabas Caching Facility in the shared Adabas Cluster Services and Adabas Parallel Services cache.

Parameter	New or Changed	Enhancement Description
CLUGROUPNAME	Changed	<p>Support for the CLUGROUPNAME parameter with Adabas Parallel Services is provided in this release. In Adabas Parallel Services, use the CLUGROUPNAME parameter to specify the name of the messaging group.</p> <p>The CLUGROUPNAME parameter is already supported in Adabas Cluster Services; this is the first release in which it is supported for Adabas Parallel Services.</p> <p>The CLUGROUPNAME parameter need not be specified for Adabas Parallel Services except when Adabas Parallel Services is used to recover an Adabas Cluster Services database after a failure that resulted in a remnant DIB entry. In this case, the CLUGROUPNAME parameter for Adabas Parallel Services must be set to the same value as it was for the failed Adabas Cluster Services nuclei.</p>
FORCE	Changed	<p>In this version a data integrity block (DIB) entry will only be removed once the ID Table initialization has been successful. Therefore, you must set IGNDIB and FORCE both to "YES" if either of the following occur:</p> <ul style="list-style-type: none"> ■ You receive a PARM ERROR 26 after parameter settings IGNDIB=NO and FORCE=YES were applied; ■ You receive a PARM ERROR 23 after parameter settings IGNDIB=YES and FORCE=NO were applied.
IGNDIB	Changed	<p>In this version a data integrity block (DIB) entry will only be removed once the ID Table initialization has been successful. Therefore, you must set IGNDIB and FORCE both to "YES" if either of the following occur:</p> <ul style="list-style-type: none"> ■ You receive a PARM ERROR 26 after parameter settings IGNDIB=NO and FORCE=YES were applied; ■ You receive a PARM ERROR 23 after parameter settings IGNDIB=YES and FORCE=NO were applied.
INFOBUFFERSIZE	New	<p>This new parameter allows you to specify the size, in bytes, of the information buffer pool that represents the new Adabas event log for a database. The default is "0" (no buffer pool/event log is allocated). Valid values other than "0" range from 1024 through 2147483647.</p>
INTAUTO	New	<p>This new parameter allows you to specify the time interval (in seconds) between autorestart progress messages. Valid values other than "0" range from 1 through 16777215.</p>
LARGEPAGE	New	<p>This new parameter allows you to indicate whether the Adabas nucleus should use large pages (1 MB pages of real storage above the 2 gigabyte bar). If your z/OS operating system supports 64-bit virtual storage and large pages, you can request that the Adabas nucleus use large pages. The default is NO (not to use large pages).</p>

Parameter	New or Changed	Enhancement Description
LDEUQP	Changed	<p>The following updates have been made to the LDEUQP ADARUN parameter in 8.2:</p> <ul style="list-style-type: none"> ■ Adabas makes an entry in the unique descriptor pool every time a record with a unique descriptor is deleted or altered such that an existing descriptor value is <i>deleted</i> from the index. Adabas Cluster Services and Adabas Parallel Services nuclei will now also make an entry when a record with a unique descriptor is <i>inserted</i> and thus the ADARUN LDEUQP parameter setting may require a larger value. Entries in the unique descriptor pool generally exist for the duration of a transaction. ■ Slight performance improvements have been made in updating large objects (LOBs) by eliminating redundant serialization. All nuclei now reduce their use of the unique descriptor pool for LOB updates and customers using LOBs from Adabas 8.1 may be able to reduce the unique descriptor pool size (ADARUN parameter LDEUQP). Adabas Cluster Services and Adabas Parallel Services configurations no longer acquire external unique descriptor locks for LOB updates. ■ Starting with Adabas 8.2 SP3, the minimum setting of the ADARUN LDEUQP parameter is increased from zero to 5000 bytes.
LFIOP	Changed	<p>The maximum setting for LFIOP is now 18,446,744,073,709,551,615 -- although this is really limited by the available amount of virtual and real memory. In addition, in environments supporting 64-bit virtual storage, the flush I/O pool is now allocated above the 2 gigabyte bar when the V64BIT ADARUN parameter is set to "YES".</p>
LOGV01io	New	<p>This new parameter allows you to indicate whether the extended I/O list should be written to the command log for CLOGLAYOUT=5 and CLOGLAYOUT=8.</p>
LOGWARN	New	<p>This new parameter allows you to specify whether CLOG and PLOG status checking should occur, with resulting alert messages, and, if so, the frequency of the status checking (in seconds). Valid values range from zero (0 - no checking) to 2147483647 seconds. When a value other than zero (0) is specified for LOGWARN, a valid user exit 2 or user exit 12 must also be specified.</p>
LRDP	Changed	<p>With this release of Adabas and the upcoming 8.2 releases of Adabas Cluster Services and Adabas Parallel Services, the default value for LRDP has changed:</p> <ul style="list-style-type: none"> ■ For Adabas Parallel Services 8.2 or later, the default is 0 (zero). ■ For Adabas Cluster Services 8.2 or later, the default is either the value of the ADARUN LFIOP parameter or 10 megabytes, whichever is smaller.
MLWTO	New	<p>Support for the MLWTO parameter is provided in this release of Adabas, in support of the following Adabas product extensions:</p> <ul style="list-style-type: none"> ■ Adabas Cluster Services

Parameter	New or Changed	Enhancement Description
		<ul style="list-style-type: none"> ■ Adabas Parallel Services ■ Event Replicator for Adabas <p>Use this parameter to specify whether multi-line messages that are normally written to the operator console as a series of single-line write-to-operator (WTO) messages should instead appear as multi-line WTO messages. When requested, this parameter will affect the appearance of replication-related and cluster-related messages (messages with prefixes ADAF*, ADAX*, PLI*, and PLX*).</p>
NPLOGBUFFERS	New	This new parameter allows you to specify the number of PLOG protection I/O buffers allocated in the Adabas nucleus session. PLOG protection I/O buffers are only allocated when ADARUN parameter LFIOP is set to a nonzero value. The default is 1.
NT	Changed	The minimum value allowed for the NT parameter was changed from "3" to "4" in this version of Adabas.
NWORK1BUFFERS	New	This new parameter allows you to specify the number of Work part 1 protection I/O buffers allocated in the Adabas nucleus session. Work part 1 protection I/O buffers are only allocated when ADARUN parameter LFIOP is set to a nonzero value. The default is 1.
PAMREQS	New	<p>This new parameter is available in BS2000 environments only. You can use it to set the maximum number of PAM blocks that can be transferred in one PAM macro request. Valid values range from 1 through 1024, with a default of 80.</p> <p>This parameter replaces an optional zap for Adabas in BS2000 environments.</p>
PLOGSEGSDS	Dropped	This parameter has been dropped from this release. It is replaced by the ADARUN SRLOG parameter.
PREFSTDD	New	This new parameter allows you to specify the job control statement label (DD name) in the job stream that defines the location to which multifetch statistics should be written when a CL command is issued.
SMF	New	This new parameter allows you to enable or disable Adabas SMF recording for the nucleus.
SMF89	New	This new parameter controls whether the Adabas nucleus registers with z/OS for type 89 SMF records.
SMFDETAIL	New	This new parameter allows you to specify the specific Interval and Termination detail record sections you want to include in the Adabas SMF records.
SMFINTERVAL	New	This new parameter allows you to specify whether interval SMF records should be generated and how often. A specific interval can be specified or a source from which the interval can be derived.
SMFRECNO	New	This new parameter allows you to specify the Adabas SMF record number for user-defined records.

Parameter	New or Changed	Enhancement Description
SMFSUBSYS	New	This new parameter allows you to associated the SMF records with an IBM or user-defined system.
SRLOG	New	This new parameter allows you to indicate how spanned records should be logged to the protection logs. Entire records (ALL), modified segments of spanned records (UPD), or entire records up to and including the modified segments (PART) can be logged. The default is UPD. The SRLOG parameter replaces the PLOGSEGSDS ADARUN parameter; the PLOGSEGSDS parameter has been dropped in this release.
SWITCHNR	New	This new parameter allows you to set the job switch number that is set when a nucleus or utility exits with a non-zero response code. Valid values range from 1 through 31, with a default of 10. This parameter replaces an optional zap for Adabas in BS2000 environments.
UEXSMF	New	This new parameter allows you to specify the module name of the SMF user exit that should be invoked during Adabas SMF recording.
V64BIT	New	This new parameter allows you to indicate whether the Adabas nucleus should use virtual storage above the 2 gigabyte bar. If your z/OS operating system supports 64-bit virtual storage, you can request that the Adabas nucleus use it also. The default is NO (not to use 64-bit virtual storage).

Direct Call Command Enhancements

The following direct call command enhancements were made in Adabas 8.2:

- Statistics for the internally used V* and Y* commands are now included on the Adabas Online System Command Usage screen as well as in Adabas shutdown statistics. These commands are used internally by Adabas and Adabas add-on products and should not be used in direct calls in your applications. Should you use them, errors will result.
- A new daylight savings time indicator, D, can be specified in format and search buffers to indicate that a daylight savings adjustment is required for a date-time field with the TZ option specified. For complete information, read *Daylight Savings Indicator (D)*, in the *Adabas Command Reference Guide*.
- **Date-time edit masks** are supported in the format and search buffers of direct calls. You can add date-time edit mask to fields with formats of B, F, P, or U in a loaded file. For more information, read *Date-Time Edit Mask Processing in Format and Search Buffers*, in the *Adabas Command Reference Guide*.
- The following table lists new and changed direct call commands in Adabas 8.2.

Command	New or Changed	Enhancement Description
A1	Changed	A new Command Option 2 value, L, can now be specified in the Adabas (ACBX or ACB) control block to indicate that the position within a large object (LOB) field be tracked in the ISN lower limit field. This option is useful when multiple A1 commands are being issued in a series for a LOB field. This option can only be used if a LOB field is specified in the format buffer using LOB segment notation with asterisk notation in the <i>bytenum</i> specification. For more information, read <i>Processing LOB Segments</i> , in the <i>Adabas DBA Tasks Manual</i>
BT	Changed	A new Command Option 3 value, H, can now be specified in the ACBX control block to indicate that a record should be kept in shared hold status indefinitely. For more information about shared hold status, read <i>Shared Hold Status</i> , in the <i>Adabas Command Reference Guide</i> .
CL	Changed	The following updates have been made: <ul style="list-style-type: none"> ■ In ACBX calls only, user CPU consumption values are now returned in eight bytes (instead of four bytes) and are stored in the ACBXISQG field. The unit of measurement is 1/4096 of a microsecond (instead of 1.048576 seconds). ■ You can now request multifetch statistics when a CL command is issued if the PREFSTDD ADARUN parameter is specified in the startup JCL and if the JCL includes a valid job control statement with the label name specified by the PREFSTDD ADARUN parameter.
ET	Changed	A new Command Option 3 value, H, can now be specified in the ACBX control block to indicate that a record should be kept in shared hold status indefinitely. For more information about shared hold status, read <i>Shared Hold Status</i> , in the <i>Adabas Command Reference Guide</i> .
HI	Changed	A new Command Option 3 value, S, can now be specified in the ACBX control block to indicate that a record should be kept in shared hold status until the end of the transaction. For more information about shared hold status, read <i>Shared Hold Status</i> , in the <i>Adabas Command Reference Guide</i> .
L1	Changed	A new Command Option 2 value, L, can now be specified in the Adabas (ACBX or ACB) control block to indicate that the position within a large object (LOB) field be tracked in the ISN lower limit field. This option is useful when multiple L1 commands are being issued in a series for a LOB field. This option can only be used if a LOB field is specified in the format buffer using LOB segment notation with asterisk notation in the <i>bytenum</i> specification. For more information, read <i>Processing LOB Segments</i> , in the <i>Adabas DBA Tasks Manual</i>
L4	Changed	Several new Command Option 3 values, C, Q and S, can now be specified in the ACBX control block to indicate that a record should be kept in shared hold status for varying time periods, depending on the option. Option C puts the record in shared hold status for the duration of the read operation; option Q puts the record in shared hold status until the next record in the read sequence is read or the read sequence or transaction is terminated; option S puts the record in shared hold status until the end of the transaction. For more information about shared hold status, read <i>Shared Hold Status</i> , in the <i>Adabas Command Reference Guide</i> .

Command	New or Changed	Enhancement Description
		<p>A new Command Option 2 value, L, can now be specified in the Adabas (ACBX or ACB) control block to indicate that the position within a large object (LOB) field be tracked in the ISN lower limit field. This option is useful when multiple L1 commands are being issued in a series for a LOB field. This option can only be used if a LOB field is specified in the format buffer using LOB segment notation with asterisk notation in the <i>bytenum</i> specification. For more information, read <i>Processing LOB Segments</i>, in the <i>Adabas DBA Tasks Manual</i></p>
L5	Changed	<p>Several new Command Option 3 values, C, Q and S, can now be specified in the ACBX control block to indicate that a record should be kept in shared hold status for varying time periods, depending on the option. Option C puts the record in shared hold status for the duration of the read operation; option Q puts the record in shared hold status until the next record in the read sequence is read or the read sequence or transaction is terminated; option S puts the record in shared hold status until the end of the transaction. For more information about shared hold status, read <i>Shared Hold Status</i>, in the <i>Adabas Command Reference Guide</i>.</p>
L6	Changed	<p>Several new Command Option 3 values, C, Q and S, can now be specified in the ACBX control block to indicate that a record should be kept in shared hold status for varying time periods, depending on the option. Option C puts the record in shared hold status for the duration of the read operation; option Q puts the record in shared hold status until the next record in the read sequence is read or the read sequence or transaction is terminated; option S puts the record in shared hold status until the end of the transaction. For more information about shared hold status, read <i>Shared Hold Status</i>, in the <i>Adabas Command Reference Guide</i>.</p>
LF	Changed	<p>A new Command Option 2 setting of "X" is supported for this command. This setting produces a record buffer containing extended field information, including collation descriptor, subfield, superfield, subdescriptor, superdescriptor, hyperdescriptor, and phonetic descriptor information with extended date and time formats.</p>
OP	Changed	<p>A new TZ (time zone) record buffer option is available for the OP command. You can use this option to set the time zone of an individual user session. If no time zone is specified for a user session, response 55 (ADARSP055) subcode 30 is issued when a field with the TZ option is accessed or updated. For more information about using the TZ record buffer option in an OP command, read about the OP command <i>Record Buffer</i>, in the <i>Adabas Command Reference Guide</i>.</p>
RI	Changed	<p>Adabas on mainframes no longer releases records unconditionally. If your application issues an RI command for a record that has been updated in the current transaction, Adabas will now return response code 113 (ADARSP113) or, if ISN=0 was specified, response code 2 (ADARSP002) is returned. This change in behavior may affect how your application programs are coded. For more information, read RI Commands: Change Possibly Affecting Your Application Programs, elsewhere in this section.</p> <p>In addition a new Command Option 3 value, S, can now be specified in the ACBX control block to indicate that a record should be kept in shared hold status until</p>

Command	New or Changed	Enhancement Description
		the end of the transaction. For more information about shared hold status, read <i>Shared Hold Status</i> , in the <i>Adabas Command Reference Guide</i> .
S4	Changed	Several new Command Option 3 values, C, Q and S, can now be specified in the ACBX control block to indicate that a record should be kept in shared hold status for varying time periods, depending on the option. Option C puts the record in shared hold status for the duration of the read operation; option Q puts the record in shared hold status until the next record in the read sequence is read or the read sequence or transaction is terminated; option S puts the record in shared hold status until the end of the transaction. For more information about shared hold status, read <i>Shared Hold Status</i> , in the <i>Adabas Command Reference Guide</i> .

For more information about direct call commands, read *Commands*, in the *Adabas Command Reference Guide*.

New LGBLSET Macro Parameters

The following new LGBLSET parameters were added in Adabas 8.2 SP2.

- LUIDX can be used to indicate whether or not a CICS link routine user ID user exit should be invoked.
- LUIXNAM can be used to specify the name of the CICS link routine user ID user exit program that should be invoked.
- UBTYPE can be used to obtain a single user block element per CICS task (instead of a user block pool). This new parameter helps eliminate abends that can occur when CICS transactions lock the user blocks but do not release them.
- UBSTIME can be used to set a time interval to clean up locked user blocks. This new parameter helps eliminate abends that can occur when CICS transactions lock the user blocks but do not release them.

For more information about the LGBLSET macro, read *Modifying Source Member Defaults (LGBLSET Macro) in Version 8*, in the *Adabas z/OS Installation Guide*.

4 Dropped Features

The following features have been dropped from Adabas in version 8.2:

- On z/VSE systems, device types 2512 and 3512 have been discontinued, because the definitions were incompatible with the physical characteristics of the SCSI devices.

Device type 1512 as supplied in Adabas 8.1 SP3 and 8.1 SP4 was also discontinued. However, device type 1512 has been redefined with changed track and cylinder sizes. The track and cylinder sizes now in use are compatible with those introduced with zap AD813035 (for Adabas 8.1 SP3) or AD814005 (for Adabas 8.1 SP4).

For more information, read *Device and File Considerations*, in the *Adabas z/VSE Installation Guide*.

- This release no longer requires that you link the Adabas SVC (ADASVC) with SAF (SVCSAF) during installation. In Adabas 8.2, the SVC is automatically linked to SAF and the SVCSAF WAL library member is no longer supplied.
- The NONUC parameter in the ADARES COPY utility function is deactivated in Adabas 8.2. You can still specify the parameter, but it is ignored.
- We no longer deliver ADACIRQ with Adabas on z/OS and z/VSE systems and DDLINK input is no longer required in CICS installations. For more information, read [CICS Installation Changes](#), elsewhere in this guide.
- The ADARUN PLOGSEGSDS parameter has been dropped and has been replaced with the ADARUN SRLOG parameter. For more information, read *SRLOG Parameter: Controlling Spanned Record Logging*, in the *Adabas Operations Manual*.

5

Limitations and Restrictions

The following limitations and restrictions exist in Adabas 8. Enhancements to resolve these limitations may be considered in a future release.

1. The following restrictions and limitations apply to large object (LB) fields in this release:
 - At this time, character conversion of LB field values from one code page to another is not supported. This functionality may be considered in a future release.
 - Some utility parameters are not supported for files containing LB fields. For more information, refer to the documentation for the utility in *Adabas Utilities Manual* and to the utility limitations and restrictions, provided later in this chapter.
 - At this time, large object (LB) fields can be define only with format A.
2. The new format buffer length indicator is only supported for LA and LB fields. Future versions of Adabas will consider supporting the specification of the length indicator for other fields too. For more information about the format buffer length indicator, read *Length Indicator (L)*, in the *Planning Manual for Adabas 8*.
3. At this time, when specifying an Adabas buffer description (ABD), you must specify the same value for the number of bytes sent to Adabas (ABDXSEND field) as you specify for the maximum buffer size (ABDXSIZE field).
4. The prefetch feature is not supported in ACBX interface direct calls -- it will not support ACBX calls with multiple buffers; you should use the multifetch feature instead. However, the prefetch feature still supports ACB interface direct calls.
5. At this time, system files do not support spanned records or the extended MU and PE field counts.
6. At this time, fields defined with the NB option must also be defined with either the NU or NC option.
7. The following restrictions and limitations apply to spanned records in this release:
 - The ADACDC utility does not support spanned records at this time.

- The ADAULD utility does not support spanned records on ADAULD SAVETAPE runs.
- At this time, ADAM files do not support spanned records.
- System files do not support spanned records at this time.
- The number of records that comprise a spanned record is limited. The Adabas nucleus allows up to five physical records (one primary record and four secondary records) in a spanned record. If you need more space, try relocating the Data Storage of the file to a different device type with a larger block size.

For more information about spanned record support in Adabas 8, read *Spanned Record Support*, in the *Adabas Concepts and Facilities Manual*.

8. At this time, Adabas Review Pulse reports do not support ADARUN CLOGLAYOUT=8.
9. The following table lists restrictions and limitations of the Adabas 8 utilities:

Utility	Restrictions or Limitations
ADACDC	At this time, ADACDC cannot process spanned records. When spanned records are updated, only the changed records are written to the PLOG and WORK data sets. Since ADACDC presents the changes of the record in external form, it needs to decompress the logical record starting from the beginning. However, if some early piece is missing, the decompression will not work.
ADACMP	At this time, LB fields cannot be specified in the FORMAT parameter for either ADACMP COMPRESS or ADACMP DECOMPRESS.
ADACNV	<p>If a database makes use of any of the following extended features of Adabas 8.2, ADACNV will not allow you to REVERT the database to a version prior to Adabas 8.2:</p> <ul style="list-style-type: none"> ■ UQDE option of one or more fields of a file was removed. <p>If a database makes use of any of the following extended features of Adabas 8, ADACNV will not allow you to REVERT the database to a version prior to Adabas 8:</p> <ul style="list-style-type: none"> ■ More than five ASSO, DATA, or DSST extents ■ More than five file extents ■ Files that allow spanned records ■ Files that allow more than 191 MU and PE occurrences ■ Files that make use of LB fields ■ Files with fields that have the NB (no blank compression) option ■ System files with two-byte file numbers. <p>If you want to complete the backward conversion (reversion), you must first remove any file with these new features from the Adabas database.</p> <p>The use of the following other new features provided in Adabas 8 do not prevent backward conversion to Adabas 7.4, but, of course, the new features cannot be used in Adabas 7.4:</p> <ul style="list-style-type: none"> ■ Adabas commands issued via the ACBX interface (for example, with long or segmented buffers)

Utility	Restrictions or Limitations
	<ul style="list-style-type: none"> ■ Commands using the new format buffer features (for example, the length indicator). <p>Note: There may be files in a database that are not loaded but that have a Field Definition Table (FDT) stored in the FDT blocks. If ADACNV encounters such FDTs while converting a database to Version 8, they are deleted as part of its cleanup processing.</p>
ADALOD	<p>The MIXDSDEV parameter is not supported in an ADALOD LOAD run if the file you are loading is LOB file or contains spanned records or extended MU or PE fields. Spanned records are based on the data storage block size. If you use a mixture of device types (as allowed by the MIXDSDEV parameter), your spanned records would end up with different record lengths, which is not allowed.</p> <p>You cannot use the DDISN or DELISN parameters in an ADALOD UPDATE function to delete records in a <i>LOB file</i>. However, you can use these parameters against a <i>base file</i> of a <i>LOB group</i> unless one or more of the records to be deleted contain references to LOB values longer than 253 bytes which are stored in the <i>LOB file</i> (ADALOD will terminate with an error if such a request is made).</p>
ADAULD	<p>For this release, you cannot unload from a SAVETAPE if the file you are trying to unload contains spanned records.</p>

6 Adabas Data Set Compatibility

▪ Importing Files	56
▪ Save Data Sets	56
▪ Unload Data Sets	57
▪ ADAORD DD/FILEA Data Sets	57
▪ Sequential Protection Logs	57
▪ BS2000 Containers Expanded by ADADBS INCREASE	58

Generally, Adabas utilities accept sequential input data sets that were produced as output data sets by utilities of the same version. Utilities of Adabas 8 also accept input data sets produced by utilities of versions prior to Adabas 8. However, utilities for versions prior to Adabas 8 cannot generally work with input data sets produced by Adabas 8 utilities.

Importing Files

A file can be imported (loaded, stored, or restored) into a database running with an earlier Adabas version as long as it does not use features that are supported only in a later Adabas version.

Save Data Sets

Generally, restoring a whole database is possible only with the same Adabas release used for creating the save data set. Restoring individual files is possible with the same or any later Adabas release used for creating the save data set.

Using the ADASAV utility of Adabas 8, you can restore files into an Adabas 8 database from a database save or file save data set created with the ADASAV utility from any prior Adabas version (for example, Adabas 7). Likewise, using the ADAREP or ADAULD utilities of Adabas 8, you can print a report or unload a file from a database save or file save data set created with the ADASAV utility from any prior Adabas version. However, you can only restore files from an Adabas 8 database save or file save data set into a database running with Adabas 7.4 if you have applied one of the following zaps to your Adabas 7 installation, as appropriate:

- For Adabas 7.4.4 installations, apply zap AU744032.
- For Adabas 7.4.3 installations, apply zap AU743099.
- For Adabas 7.4.2 installations, apply zap AU742186.

You cannot use an ADAREP or ADAULD utility from an Adabas version prior to Adabas 8 to print a report or unload a file from a save data set created by the ADASAV utility of Adabas 8. If you are using Adabas Delta Save, you cannot merge together delta save data sets created by different versions of the ADASAV utility. After the database has been converted to version 8, the next save operation must be a full save.

Unload Data Sets

Using the ADALOD utility of Adabas 8, you can load a file from an unload data set created using the ADAULD or ADACMP utilities from any prior version of Adabas. However, you can only use the ADALOD or ADACMP utilities from Adabas 7 to load or decompress a file from an unload data set created by the Adabas 8 ADAULD or ADACMP utilities if you have applied one of the following zaps to your Adabas 7 installation, as appropriate:

- For Adabas 7.4.4 installations, apply zap AU744033.
- For Adabas 7.4.3 installations, apply zap AU743100.
- For Adabas 7.4.2 installations, apply zap AU742187.

ADAORD DD/FILEA Data Sets

You can restructure databases and files from an Adabas version prior to Adabas 8 and store them in an Adabas 8 database using ADAORD STORE. However, you cannot store the restructured output of an Adabas 8 database or file in a database running with any prior Adabas version (for example, Adabas 7). If you attempt this, the following warning will be generated and ADAORD will end with a CC=4:

```
*** Warning: The input dataset is from V8 and will not be processed
```

Sequential Protection Logs

Any sequential protection log (PLOG) used for a REGENERATE or BACKOUT function under Adabas 8 must have been created with Adabas 8. Using a sequential PLOG created with Adabas 8 for a REGENERATE or BACKOUT function under any prior Adabas 8 version is not supported.

You cannot run the ADASEL utility of Adabas 8 to select data from a sequential PLOG created by any Adabas version prior to Adabas 8. In addition, you cannot run the ADASEL utility of a prior Adabas version (for example, Adabas 7) to select data from a PLOG created by Adabas 8.

You also cannot run the ADARES COPY utility of Adabas 8 to copy data from a sequential PLOG created by any Adabas version prior to Adabas 8. In addition, you cannot run the ADARES COPY utility of a prior Adabas version (for example, Adabas 7) to copy data from a PLOG created by Adabas 8.

BS2000 Containers Expanded by ADADBS INCREASE

If the database had previously been expanded using the ADADBS INCREASE utility followed by an ADAFRM FROMRABN utility in a previous version, we recommend that you check the size of that database container.


1. Run an ADAREP utility report to determine the highest RABNs in the container.
2. Calculate the highest PAM page in the container using the RABN STD block size as shown in the device table for that device type in *Device and File Considerations*(in the *Adabas BS2000 Installation Guide*) as follows:

```
Highest PAM Page = (Highest RABN) * (RABN STD Block Size)
```

3. Check this value with the highest physical PAM page output from the following command:
`/SHOW-FILE-ATTRIBUTES container-name,ALL` . The value in the HIGH-US-PA field contains the highest physical PAM page (this will have to be done cumulatively for multiple containers).

If the highest PAM page is less than the highest PAM pages shown in the ADAREP output, accessing the highest RABNs will result in a DMS0922 I/O error.

7 Applying Zaps

 **Important:** Be sure that you apply all supplied Adabas maintenance and concatenate Adabas patch-level libraries (L00*n*), as they are delivered to you. This will ensure that your Adabas code remains up-to-date, supporting all Adabas features as they are enhanced and maintained. The latest zaps for this product are available in the Knowledge Center in Software AG's Empower (<https://empower.softwareag.com>) web site.


In general, zaps for Adabas components (such as Adabas nuclei, the Adabas router, Adabas utilities) can be applied and made active one component at a time.

- Adabas utility zaps should be applied to the load library. The utility can then be run or rerun to make use of the zap.
- Adabas nucleus zaps should first be applied to the load library. Then the nucleus should be stopped and restarted to activate the zap.
- Adabas router (on z/OS, ADASVC) zaps should first be applied to the load library. Then all Adabas nuclei and other MPM servers running on the router should be stopped and the router should be reinstalled to activate the zap. Finally, the Adabas nuclei and MPM servers should be restarted.

Finally, the distributed source library contains member ZAPOPT, which lists some optional zaps that you may choose to apply for the activation or deactivation of various features and optional user settings of Adabas. A ZAPOPT member will be included with each SM level distribution.

8 Software AG Product Compatibility

The following table describes Adabas 8 compatibility with other mainframe Adabas products, including prior releases of Adabas itself. You may need to upgrade your installation of the software if your existing release is not listed.

 **Note:** Any exceptions to the product compatibility described here will be covered in the documentation for the specific product.

Product	Compatible Version Levels and Notes
Adabas (ADA)	<p>The version used of the Adabas SVC must be the same as or greater than the version of any Adabas database used in your Adabas environment. For example, the Adabas 8.2 SVC can be run in the same environment with Adabas 7.4, 8.1, or 8.2 databases. However, an Adabas 8.2 database cannot run in the same environment with an Adabas 8.1 SVC.</p> <p>For any given database (on disk), the Adabas nucleus and utilities of the same version and release level as the database must be used. If you need to convert a database to a higher version or release level, or revert it to a lower version or release level, the ADACNV utility of the higher level must be used.</p> <p>The Adabas link (Adalink) routines can be used across versions. For example, Adabas 8 link routines can be used to issue calls to Adabas 7 databases. Unless the new ACBX direct call interface is used, Adabas 7.4 link routines can be used to issue calls to Adabas 8 databases. Software AG recommends that you use the Adabas 8 link routines for all programs that issue Adabas direct calls.</p>
Adabas Bridge for DL/I (ADL)	Version 2.3 SP1 and above supports Adabas 8.
Adabas Bridge for VSAM (AVB)	Version 5.1 SP1 releases support Adabas 8 databases that do not make use of the expanded features (for example, spanned records, increased limits, or large object fields) available in Adabas 8.

Product	Compatible Version Levels and Notes
Adabas Caching Facility (ACF)	<p>Version 8.2 and above fully support all Adabas 8.2 databases and expanded features.</p> <p>Adabas Caching Facility 8.1 must be run with Adabas 8. Likewise, Adabas Caching Facility 7.4 must be run with Adabas 7.4. You cannot run Adabas Caching Facility 7.4 with Adabas 8.</p>
Adabas CICS Interface (ACI)	Version 8.2 and above fully supports all Adabas 8.2 databases and expanded features.
Adabas Cluster Services (ALS)	Version 8.2 SP3 supports Adabas 8.2 databases and requires the Adabas 8.2 SP3 load library, with appropriate Adabas 8.2 SP3 zaps applied.
Adabas Delta Save (ADE)	<p>Version 8.2 and above fully support all Adabas 8.2 databases and expanded features.</p> <p>Adabas Delta Save 8.1 must be run with Adabas 8.1. Likewise, Adabas Delta Save 7.4 must be run with Adabas 7.4. You cannot run Adabas Delta Save 7.4 with Adabas 8.1.</p>
Adabas Fastpath (AFP)	<p>Version 8.1 and above fully support all Adabas 8.1 databases and expanded features.</p> <p>Version 8.1 and above fully support Adabas 8.2 databases including any features that had been introduced by Adabas 8.1 (no new Adabas 8.2 features are supported), but this support requires additional installation steps and applied zaps. For more information, refer to the instructions provided in Empower. Once you have logged into Empower, select Knowledge Center and then Technical Papers. The paper is entitled "Using 8.1.2 ATM, AFP, AVI, AAF, COR with ADA 8.2 SP2 in z/OS".</p>
Adabas IMS Interface (AII)	Version 8.2 and above fully support all Adabas 8.2 databases and expanded features.
Adabas Manager (AMA)	<p>Version 7.5 and 7.6 releases support Adabas 7 and 8.1 databases. Version 7.7 fully supports Adabas 8.2 databases and expanded features.</p> <p>Note: Version 7.7 supports Adabas Cluster Services 8.2 and Adabas Parallel Services 8.2, but does not display session-related information from a specific nucleus (NUCID) in a cluster.</p>
Adabas Native SQL (SQL)	Version 2.4 SP1 supports Adabas 8 databases that do not make use of the expanded features (for example, spanned records, increased limits, or large object fields) available in Adabas 8. Full support for Adabas 8 is planned.
Adabas Online System (AOS)	Version 8.2 and above fully support all Adabas 8.2 databases and expanded features. You can also access Adabas 7.4 databases using Adabas Online System Version 8, but you cannot access Adabas 8 databases using Adabas Online System 7.4.
Adabas Parallel Services (ASM)	Version 8.2 SP3 supports Adabas 8.2 databases and requires the Adabas 8.2 SP3 load library, with appropriate Adabas 8.2 SP3 zaps applied.

Product	Compatible Version Levels and Notes
Adabas Review (REV)	<p>Version 4.5.1 and above support Adabas 8.2 databases and expanded features.</p> <p>Version 4.4 SP1 supports Adabas 8.1 databases and expanded features.</p> <p>For a complete description of the compatibility between Adabas Review and Adabas 8, read the section entitled <i>Adabas Review Version Compatibility</i> in the Adabas Review Release Notes.</p>
Adabas SAF Security z/OS (AAF)	<p>Version 8.1 and above fully support all Adabas 8.1 databases and expanded features.</p> <p>Version 8.1 and above fully support Adabas 8.2 databases including any features that had been introduced by Adabas 8.1 (no new Adabas 8.2 features are supported), but this support requires additional installation steps and applied zaps. For more information, refer to the instructions provided in Empower. Once you have logged into Empower, select Knowledge Center and then Technical Papers. The paper is entitled "Using 8.1.2 ATM, AFP, AVI, AAF, COR with ADA 8.2 SP2 in z/OS".</p>
Adabas Statistics Facility (ASF)	<p>Version 8.1 and above fully support all Adabas 8 databases and expanded features. It also supports Adabas 7 databases.</p>
Adabas SQL Gateway (ACE)	<p>Version 6.6 and above support Adabas 8 databases that do not make use of the expanded features (for example, spanned records, increased limits, or large object fields) available in Adabas 8. Full support for Adabas 8 databases and expanded features is planned.</p>
Adabas SQL Server (ESQ)	<p>Version 1.4 SP3 supports Adabas 8 databases that do not make use of the expanded features (for example, spanned records, increased limits, or large object fields) available in Adabas 8.</p>
Adabas System Coordinator (COR)	<p>Version 8.1 and above fully support all Adabas 8.1 databases and expanded features.</p> <p>Version 8.1 and above fully support Adabas 8.2 databases including any features that had been introduced by Adabas 8.1 (no new Adabas 8.2 features are supported), but this support requires additional installation steps and applied zaps. For more information, refer to the instructions provided in Empower. Once you have logged into Empower, select Knowledge Center and then Technical Papers. The paper is entitled "Using 8.1.2 ATM, AFP, AVI, AAF, COR with ADA 8.2 SP2 in z/OS".</p>
Adabas Text Retrieval (TRS)	<p>Version 2.1 SP4 works with Adabas 8.2 SP2 when the Adabas Text Retrieval 2.1 SP4 hyperexit TRSHEX12 is enabled to run with the Version 8 interface by applying zaps TR21454 and TR21455. This hyperexit will then only operate with Adabas Version 8. If you then want to run Adabas Text Retrieval 2.1 SP4 with Adabas Version 7 again, you must either undo the zaps or use a copy of the hyperexit where zaps TR21454 and TR21455 are not applied. If you do use TR21454 and TR21455, the following additional Adabas Text Retrieval fixes must be applied as prerequisite zaps: TR21420, TR21421, TR21422, TR21423 and TR21424. These zaps can be found in Empower.</p>

Product	Compatible Version Levels and Notes
	It is not necessary to use the Adabas Hyperexit Stub in conjunction with Adabas Text Retrieval.
Adabas Transaction Manager (ATM)	<p>Version 8.1 and above fully support all Adabas 8.1 databases and expanded features.</p> <p>Version 8.1 and above fully support Adabas 8.2 databases including any features that had been introduced by Adabas 8.1 (no new Adabas 8.2 features are supported), but this support requires additional installation steps and applied zaps. For more information, refer to the instructions provided in Empower. Once you have logged into Empower, select Knowledge Center and then Technical Papers. The paper is entitled "Using 8.1.2 ATM, AFP, AVI, AAF, COR with ADA 8.2 SP2 in z/OS".</p>
Adabas Vista (AVI)	<p>Version 8.1 and above fully support all Adabas 8.1 databases and expanded features.</p> <p>Version 8.1 and above fully support Adabas 8.2 databases including any features that had been introduced by Adabas 8.1 (no new Adabas 8.2 features are supported), but this support requires additional installation steps and applied zaps. For more information, refer to the instructions provided in Empower. Once you have logged into Empower, select Knowledge Center and then Technical Papers. The paper is entitled "Using 8.1.2 ATM, AFP, AVI, AAF, COR with ADA 8.2 SP2 in z/OS".</p>
Data Archiving for Adabas (ADR)	Version 1.1 supports Adabas 8.1 or above in z/OS environments.
Entire Net-Work (WCP)	Version 6.1 and above fully support Adabas 8 databases and expanded features, as well as ACBX interface direct calls.
Entire System Server	Entire System Server Version 3.4 fully supports Adabas 8 databases and expanded features.
Event Replicator for Adabas (ARK)	Version 3.3 SP2 supports Adabas 8.2 SP2 and 8.2 SP3. Version 3.2 SP1 with additional load library ARF321.L001 supports Adabas 8.2 SP2 and 8.2 SP3. A z/OS Event Replicator Server can also be an Adabas 8.2 SP2 or later nucleus (server). For more information, refer to your Event Replicator for Adabas documentation.
Natural (NAT)	Version 4.2, with Entire System Server 3.3 fully supports Adabas 8.1 databases and expanded features. Please consult the applicable Natural Release Notes for information regarding Natural support of the new Adabas 8.2 features that enhance the Adabas direct call interface.
Predict	Version 4.5 fully supports Adabas 8.1 databases and expanded features. Please consult the applicable Predict Release Notes for information regarding Predict support of the new Adabas 8.2 features that enhance the Adabas direct call interface.
EntireX Communicator/webMethods EntireX	Version 7.2 SP1 and later fully support the Adabas 8 SVC and the use of an Adabas 8 database as a persistent store.

9 End of Maintenance

For information on how long a product is supported by Software AG, access Software AG's Empower web site at <https://empower.softwareag.com>.

Log into Empower. Once you have logged in, you can expand **Products** in the left menu of the web page and select **Product Version Availability** to access the Product Version Availability application. This application allows you to review support information for specific products and releases.

10 Documentation

- Allowing Active Content in Windows Environments 68

The documentation for this product is new with this release. When additional updated versions of the documentation are created, you can review them by linking to the Software AG documentation web site: <http://documentation.softwareag.com/>. If you have an Empower account, updated and past versions of the documentation can also be reviewed and downloaded by linking to the Software AG Empower web site: <https://empower.softwareag.com>. If you do not have an Empower user ID and password yet, you will find instructions for registering on this site (free for customers with maintenance contracts).

This documentation includes:

- online HTML topics describing all aspects of the product;
- Adobe Acrobat Reader Portable Document Format (PDF) files created from the HTML topics;
- Adobe Acrobat Reader Portable Document Format (PDF) files for a series of manuals created from the HTML topics.

No hard-copy documentation is provided, but you can print the PDF and HTML files on your local printer.

Allowing Active Content in Windows Environments

With Service Pack 2 (SP2) for Windows XP and Service Pack 1 (SP1) for Server 2003 and later Windows releases, Microsoft introduced a range of powerful new security features that restrict active content that runs locally on your computer. Active content includes ActiveX controls, Java applets, and JavaScript. Software AG's documentation web pages contain some JavaScript, and the SEARCH, INDEX and CONTENTS capabilities are implemented as Java applets. As a result, when viewing documentation web pages that reside on your PC using Internet Explorer and Mozilla Firefox under Windows XP SP2, note that active content is blocked. You must explicitly and repeatedly allow active content if you want to make use of the documentation's full navigation features. Note that this behavior is only observed when reading web pages installed locally on your PC, including those on CD in the PC's CD-ROM drive.

The active content for which Software AG is responsible, that is, the JavaScript code in our HTML documentation pages, will not harm your computers. The risk in using the navigation applets is negligible: Software AG has received no reports from users concerning any harm caused to a computer by the applets. We therefore suggest that when reading Software AG documentation in a local context, you should allow active content via the Security settings in the browser (with Internet Explorer, usually found under Tools > Internet Options > Advanced).

Full details of alternatives can be found on the home page of the suppliers of the navigation applets: <http://www.phdccc.com/xpsp2.htm>.

Index

A

applying zaps, 59

D

documentation
 allowing active content in Windows, 68
 obtaining updates, 67

E

Empower
 documentation, 67
 end-of-maintenance dates, 65
 platform support, 3
end-of-maintenance dates, 65

P

platform support, 3
product support
 end-of-maintenance dates, 65
 obtaining updated documentation, 67
 supported platforms, 3

R

response code 145 (ADARSP145)
 logging events, 25

S

spanned records
 logging, 24
support
 end-of-maintenance dates, 65
 obtaining updated documentation, 67
 platforms supported, 3
support dates, 65
supported platforms, 3

V

V* commands, 44

Y

Y* commands, 44

Z

zaps, 59

