

Adabas® Version 7.4

Release Notes

Manual Order Number: ADA744-008IBB

This document applies to the Adabas software package at version 7.4 and to all subsequent versions, unless otherwise indicated in new editions or technical newsletters.

Specifications contained herein are subject to change, and these changes will be reported in subsequent revisions or editions.

Readers' comments are welcomed. Comments may be addressed to the Documentation Department at the address on the back cover or to the following e-mail address:

Documentation@softwareag.com

Copyright (c) November 2007 by Software AG. All rights reserved.

The name Software AG and/or all Software AG product names are either trademarks or registered trademarks of Software AG. Other company and product names mentioned herein may be trademarks of their respective owners.

TABLE OF CONTENTS

PREFACE	1
About this Release	1
About this Document	1
Documentation	1
Printed Manuals	1
Online Documentation	2
ENVIRONMENT	3
Operating Systems Supported	3
Support for Prior Versions	4
Optional ZAPs Data Set	6
Software AG Product Compatibility	7
Adabas Component Compatibility	9
Communication Modules	9
SVC Consolidation	12
Migration Issues (IBM Platforms)	12
TP Link Routines (IBM Platforms Only)	13
Enqueue Conflicts	13
User Exit Interfaces	14
Conversion Requirements	14
New Skeleton Job Control File	14
Adabas Recovery Aid Compatibility	14
Migrating Products from Version 7.1 to Version 7.4	15
Adabas Data Set Compatibility	15
Save Data Sets	15
Unload Data Sets	16
ADAORD DD/FILEA Data Sets	16
Sequential Protection Logs	16

Adabas Version 7.4 Release Notes

ADABAS VERSION 7.4.4 RELEASE INFORMATION	17
Enhancements	17
New ADARUN MXCANCEL Parameter	17
New ADARUN PGFIX Parameter	18
Large Format File Support	18
ADACDC Utility Updates	19
VSE and VM FBA Device Support Updated	19
Adding UES Support to VM Databases	19
VM Installation Procedure Changes	20
Event Replicator for Adabas	20
Changes in Future Versions	21
Removed from Adabas	21
ADALNA5 Module Dropped	21
ADABAS VERSION 7.4.3 RELEASE INFORMATION	23
Enhancements	23
Enhanced Interfaces between Products	23
Adabas Parallel Services 7.5.1 Support	23
CICS Link Routine Updates	23
ADARUN FORCE Parameter Changes	24
ADABAS VERSION 7.4.2 RELEASE INFORMATION	25
Summary of Features Added for 7.4	25
Enhanced Interfaces between Products	25
Enhanced Support for Operating System Features	25
More SVC Flexibility (IBM Platforms Only)	26
Enhanced Backup and Restore Performance	26
File Capacity (BS2000/OSD)	27
Data Integrity Enhancements	27
Enhanced Logging Performance	28
Other Performance Enhancements	28
Enhanced Link Routines	29
New and Enhanced ADARUN Options	29

Table of Contents

Feature Discussion	30
NUCID	30
PLOGRQ=FORCE Option	30
New Table to Track the Active Nuclei	31
Switching Between Cluster and Noncluster Mode	33
Support for the Adabas SQL Server (AQA)	34
New Utility Function for Use with Fast Copy Products (CE)	34
Support for Assembler SYSPARM Specification of the Adabas SVC (IBM Platforms)	35
Messages in Mixed Case (CE)	36
Automatic Restart Management (ARM) Support	36
New Buffer Flush Method	36
New ADARUN Parameters to Include Files in Prefetch/Multifetch	38
Asynchronous CLOG I/O	38
Enhancements	38
CSAs of Replaced SVCs Released (IBM Platforms Only)	38
Different ISN Returned with Record-related Response Code	39
ADASAV Processing Change	39
Number of Blocks Allocated to Work Data Protection Area Increased	40
Link Routines are UES-enabled by Default	40
Conversion of High Value in Value Buffer	41
PRILOG Now Delivered in Both Source and Object Form	41
Nucleus Processing When All PLOG Data Sets are Full	41
ADACOX (Adabas Conversion Exit)	41
Conversion of Superdescriptor Search Values	42
Removed from Adabas	42
Support Removed for Alternate RABNs	42
Mirror Data Set Support Removed	42
Adabas ESA Option (AEO) No Longer Available	43
ADAESI No Longer Available	43
Interlink TCP/IP No Longer Supported	43
Access Register SVC Under VSE/ESA	43
Dropped Support for the Adabas CICS Macro-level Link Routine (ADALNC)	44
Dropped Support for DSDEV Subparameter	44
Changes to Adabas Online System	44
Product Interfaces	44
Installation Requirement	45

Adabas Version 7.4 Release Notes

‘In Parallel’ Option	45
ZAP Prefixes for Adabas Online System (AOS)	46
Updated Adabas Online System Screens	46
INDEX	49

PREFACE

About this Release

Adabas version 7.4.4 is a maintenance release. It includes fixes and documentation updates made since the release of Adabas 7.4.3.

About this Document

These *Release Notes* accompany Adabas version 7.4.4. They contain information about the changes to Adabas since version 7.4.3.

Chapter 1 contains information about required and compatible environments.

Chapter 2 describes the changes made to Adabas for 7.4.4.

Chapter 3 describes the changes made to Adabas for 7.4.3.

Chapter 4 describes the changes made to Adabas for 7.4.2 and prior 7.4 releases.

Documentation

One printed manual is provided with the Adabas version 7.4.4 release. All other documentation is maintained online. To review past versions of the documentation, link to the Software AG Web site: <http://servline24.softwareag.com/public/>. The Adabas documentation at this site will be updated as required.

Printed Manuals

The following printed manual is delivered with the version 7.4.4 release:

Order No.	Manual
ADA744-008IBB	<i>Adabas Version 7.4.4 Release Notes</i> (this document)

Online Documentation

The documentation for mainframe Adabas Version 7.4.4 is being provided on CD as HTML topics with PDFs generated from the HTML pages.

This online release of the Adabas mainframe documentation covers the following areas:

Installation and Getting Started

- Concepts and Facilities
- Release Notes 7.4.4 (PDF)
- Installation for z/OS and OS IV/F4
- Installation for BS2000/OSD
- Installation for z/VM
- Installation for z/VSE
- DBA Tasks
- Operations

Reference

- Command Reference
- Utilities
- Messages and Codes
- User Exits
- Command Log Formats
- Supplied UES Encodings
- Triggers and Stored Procedures

In addition, documentation for all Adabas add-ons is provided on the CD.

Customer feedback on the usability and quality of this documentation is welcomed.

ENVIRONMENT

This chapter describes:

- operating systems supported for this release
- support for prior versions of Adabas and its add-on products
- the optional ZAPs data set
- Software AG products that are compatible with this release of Adabas
- Adabas component compatibility with this release of Adabas, including migration considerations
- Adabas data set compatibility with this release of Adabas, including migration considerations

Operating Systems Supported

Adabas version 7.4.4 is released for the following operating environments:

- z/OS (MVS): All versions currently supported by IBM.
- z/VM: All versions currently supported by IBM.
- z/VSE: All versions currently supported by IBM.
- BS2000/OSD: For information on the BS2000 OSD versions supported by this release of Adabas, access Software AG's ServLine24 web site at <http://servline24.softwar eag.com/public>. In the left menu of this web page, expand My ServLine24 and log into Secured Services. Once you have logged in, you can expand Products in the left menu of the web page and select Product Roadmaps to access the Product Version Roadmaps application. This application allows you to review platform support information for specific Software AG products and releases.
- MSP: All versions currently supported by Fujitsu
- VOS: All versions currently supported by Hitachi.

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

Although it may be technically possible to run a new version of Adabas 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 Adabas on a release, version, or type of operating system other than those listed above, consult your Adabas technical support facility to determine whether support is possible, and under what circumstances.

Support for Prior Versions

Software AG plans to provide support for prior Adabas releases as follows:

Adabas Product and Release	Release Date	End-of-Support Date
Adabas 7.1.3	24 August 2001	31 December 2006
Adabas 7.2.2	2 November 2001	31 May 2005 (controlled release only)
Adabas 7.4.2	13 June 2003	
Adabas 7.4.3	26 November 2004	
Adabas 7.4.4	09 December 2005	
Adabas Bridge for DL/I 2.2.2	5 February 1996	
Adabas Bridge for VSAM 5.1.1	26 April 2002	
Adabas Caching Facility 7.3.1	23 November 2001	31 May 2005
Adabas Caching Facility 7.4.2	13 June 2003	
Adabas CICS Interface 7.1.3	24 August 2001	31 December 2006
Adabas CICS Interface 7.4.2	13 June 2003	
Adabas CICS Interface 7.4.3	26 November 2004	
Adabas CICS Interface 7.4.4	09 December 2005	
Adabas Cluster Services 7.2.2	27 December 2001	31 May 2005 (controlled release only)
Adabas Cluster Services 7.4.2	1 August 2003	
Adabas Delta Save 7.1.2	30 April 2000	31 December 2006

Adabas Product and Release	Release Date	End-of-Support Date
Adabas Delta Save 7.2.2	27 December 2001	31 May 2005 (controlled release only)
Adabas Delta Save 7.4.2	13 June 2003	
Adabas Delta Save 7.4.3	26 November 2004	
Adabas Fastpath 7.1.3	7 September 2001	31 December 2006
Adabas Fastpath 7.3.1	27 December 2001	31 May 2005
Adabas Fastpath 7.4.2	19 September 2003	
Adabas IMS Interface 7.1.3	24 August 2001	31 December 2006
Adabas IMS Interface 7.4.2	13 June 2003	
Adabas IMS Interface 7.4.3	26 November 2004	
Adabas IMS Interface 7.4.4	09 December 2005	
Adabas Native SQL 2.2.1 – MF	27 March 2002	28 February 2006
Adabas Native SQL 2.3.1 – MF	25 February 2005	
Adabas Online Services 7.1.3	24 August 2001	31 December 2006
Adabas Online Services 7.2.2	24 August 2001	31 May 2005 (controlled release only)
Adabas Online Services 7.4.2	13 June 2003	
Adabas Online Services 7.4.3	26 November 2004	
Adabas Online Services 7.4.4	09 December 2005	
Adabas Parallel Services 7.4.2	5 August 2003	30 November 2005
Adabas Parallel Services 7.5.1	5 April 2004	
Adabas Review 4.3.1	6 June 2002	31 December 2004
Adabas Review 4.3.2	8 August 2003	
Adabas SAF Security 7.1.3	7 September 2001	31 December 2006
Adabas SAF Security 7.3.1	27 December 2001	31 May 2005
Adabas SAF Security 7.4.2	13 June 2003	
Adabas SQL Gateway 6.1.1	28 May 2004	30 September 2005
Adabas SQL Gateway 6.2.1	15 September 2004	31 December 2005

Adabas Product and Release	Release Date	End-of-Support Date
Adabas SQL Gateway 6.3.1	20 December 2005	29 September 2006
Adabas SQL Gateway 6.4.1	12 May 2005	
Adabas SQL Server 1.4.3 – MVS	2 August 1999	31 December 2007
Adabas SQL Server 1.4.3 – VSE	2 August 1999	31 December 2007
Adabas Statistics Facility 7.1.2	1 May 2000	
Adabas Text Retrieval 2.1.5	2 August 2002	
Adabas Transaction Manager 1.2.2	27 December 2001	31 May 2005
Adabas Transaction Manager 7.4.2	19 September 2003	30 November 2006
Adabas Transaction Manager 7.5.1	29 April 2005	
Adabas UTM Interface 7.1.3	24 August 2001	31 December 2006
Adabas UTM Interface 7.4.2	13 June 2003	
Adabas UTM Interface 7.4.3	26 November 2004	
Adabas Vista 7.1.3	7 September 2001	31 December 2006
Adabas Vista 7.3.1	21 December 2001	31 May 2005
Adabas Vista 7.4.2	19 September 2003	
Event Replicator for Adabas 2.1.1	30 June 2005	29 June 2007
Event Replicator for Adabas 2.2.1	23 December 2005	
Jadabas 7.1.1	9 July 2001	31 May 2006
Jadabas 7.4.2	4 February 2005	

Optional ZAPs Data Set

The distributed source library contains a member ZAPOPT, which lists zaps that may optionally be applied for the activation or deactivation of various features of Adabas. A ZAPOPT member will be included with each SM level distribution.

Software AG Product Compatibility

The following table describes Adabas 7.4.4 compatibility with other Adabas products, including prior releases of Adabas itself. You may need to upgrade your installation of this software if your existing release is not listed.

Adabas Product	Compatible Version Levels and Notes
Adabas (ADA)	Version 7.1 and above
Adabas Caching Facility (ACF)	Version 7.3 (initial release) or 7.4
Adabas Cluster Services (ALS) (z/OS)	Version 7.4
Adabas Delta Save Facility (ADE)	Version 7.4
Adabas Fastpath (AFP)	Version 7.4.2 If you install the Adabas 7.4.4 INPL, you must reapply the Adabas Fastpath base INPL and the most recent Adabas Fastpath INPL update. This is because the Adabas 7.4.4 INPL file includes a basic "demo" level which is not at the fully licensed level of functionality.
Adabas IMS/TM Interface	Version 7.4.4
Adabas Online System (AOS)	Version 7.4.4 In the initial Adabas delivery, AOS functions as a demonstration system. For full functionality, the contents of the AOS 7.4.4 load library must be copied to the Adabas 7.4.4 load library, or the AOS 7.4.4 load library must be concatenated in the STEPLIB of the Adabas nucleus JCL. In addition, all AOS users (demo or full version) and Predict users must include the load module AOSASM from the Adabas 7.4.4 load library in the link of the Natural nucleus.
Adabas Parallel Services (ASM)	Version 7.4.2 z/OS and z/VSE only Version 7.5.1

Adabas Product	Compatible Version Levels and Notes
Adabas Review (REV)	Version 4.3 Upgrade Adabas Review to version 4.3 first; then upgrade Adabas to version 7.4. See the Adabas Review 4.3 Release Notes for information on how to run Review 4.3 with versions of Adabas prior to 7.4.
Adabas SAF Security (AAF) (z/OS)	Version 7.4.2 (Adabas SAF Security must be upgraded to 7.4.2 for Adabas 7.4 databases). If you install the Adabas 7.4.4 INPL, you must reapply the Adabas SAF Security base INPL and the most recent Adabas SAF Security INPL update. This is because the Adabas 7.4.4 INPL file includes a basic "demo" level which is not at the fully licensed level of functionality.
Adabas Statistics Facility (ASF)	Version 7.1; Adabas Statistics Facility is not supported with Adabas Cluster Services or Adabas Parallel Services.
Adabas System Coordinator (COR)	Version 7.4.2 If you install the Adabas 7.4.4 INPL, you must reapply the Adabas System Coordinator base INPL and the most recent Adabas System Coordinator INPL update. This is because the Adabas 7.4.4 INPL file includes a basic "demo" level which is not at the fully licensed level of functionality.
Adabas Transaction Manager (ATM) (z/OS, z/VSE, BS2000/OSD)	Version 7.4.2 and 7.5 If you install the Adabas 7.4.4 INPL, you must reapply the Adabas Transaction Manager base INPL and the most recent Adabas Transaction Manager INPL update. This is because the Adabas 7.4.4 INPL file includes a basic "demo" level which is not at the fully licensed level of functionality.
Adabas Vista (AVI)	Version 7.4 If you install the Adabas 7.4.4 INPL, you must reapply the Adabas Vista base INPL and the most recent Adabas Vista INPL update. This is because the Adabas 7.4.4 INPL file includes a basic "demo" level which is not at the fully licensed level of functionality.

Adabas Product	Compatible Version Levels and Notes
Entire Net–Work (WCP)	Version 5.8.1 or above is required. In addition, ZAP WM58028 must be applied if Entire Net–Work 5.8 is used. This ZAP solves a problem with the improper retranslation of ASCII client calls to non–mainframe remote targets that do not use the EBCDIC character set. If Entire Net–Work 5.9.1 is used, this ZAP is not necessary.
Event Replicator for Adabas (ARK)	Version 2.1.1
Jadabas (AVA)	Version 7.4 Version 7.1 classes can still be used. Entire Net–Work 5.8.1 or higher is required to access mainframe databases. To access databases on UNIX, Linux, and Windows, Entire Net–Work 2 can be used, however we recommend the use of Entire Net–Work 7.2.2.
Natural (NAT)	Version 4.1 or above is required to run the Natural (online) components associated with Adabas version 7.4.4 and its add-ons.

Adabas Component Compatibility

Communication Modules

The Adabas communication modules include ADALNK, ADAMPM, ADASIP, ADASIR, and the Adabas SVC/router [ADASVC/ADARER (BS2000/OSD)]. This section describes the compatibility issues for these modules.

Warning:

Failure to comply with these compatibility restrictions may lead to abnormal program termination of communications modules and may require a complete reinitialization of the environment.

Module	Compatibility Notes
--------	---------------------

Module	Compatibility Notes
ADALNK	<p>Software AG recommends that all applications use the ADALNK module from a currently supported Adabas release. <i>It is recommended that you link ADAUSER, rather than ADARUN or ADALNK, to the component application for easier migration to new Adabas versions.</i></p> <p>The ADALNK module can be used with Adabas version 7.1 and Adabas version 7.2 (i.e., it is downward compatible). However, you should read <i>Planned Changes in Future Versions</i> on page 21 for more information about the future of ADALNK.</p>
ADASIP	<p>The ADASIP module can be used with Adabas version 7.1 and Adabas version 7.2 (i.e., it is downward compatible).</p>

Module	Compatibility Notes
ADASVC/ADARER	<p data-bbox="542 286 1228 385">Software AG recommends that you install version 7.4 ADASVC or ADARER (BS2000/OSD) and API link routines.</p> <p data-bbox="542 394 1228 600"><i>Note:</i> <i>On BS2000/OSD systems, an attempt to start a V74 nucleus on an ID Table containing a router(ADARER) on a lower version will result in an error message. This is because new functionality for V74 routing is not available in the old version. Older version databases can, however, run with the new IDT/Router.</i></p> <p data-bbox="542 609 1228 645">In a single nucleus (noncluster) environment:</p> <ul data-bbox="606 654 1228 940" style="list-style-type: none"><li data-bbox="606 654 1228 797">• version 7.4 ADASVC/ADARER is compatible with lower-level Adabas databases; that is, you can run nuclei of Adabas version 7.1 or 7.2 with the installed Adabas 7.4 SVC.<li data-bbox="606 806 1228 940">• you can run Adabas version 7.4 with version 7.2 or 7.1 ADASVC/ADARERs. <i>However, when using Adabas version 7.4, it is preferable to use the new SVC to take advantage of its full functionality.</i> <p data-bbox="542 949 1228 1048">In a cluster environment, the version level of Adabas, the ADASVC/ADARER, and the cluster product (Adabas Cluster Services or Adabas Parallel Services) must match.</p> <p data-bbox="542 1057 1228 1236">If you want to continue using Adabas Support for Multiprocessing, you may run the version 7.1 ADASMP module using an Adabas version 7.4 SVC/router; however, you must also use the Adabas version 7.4.2 SMPSVC linked into the version 7.4 ADASVC/router. See the section Migration Issues on page 12.</p> <p data-bbox="542 1245 1228 1306">The ADASVC module can be used with Adabas version 7.1 and Adabas version 7.2 (i.e., it is downward compatible).</p>

Notes:

On BS2000/OSD systems, version 7.4 systems cannot function with version 7.1 Routers. It is necessary to establish an IDT common memory pool with the V74 Router. It is recommended that you copy the modules ADAMPM, ADAMP2, ADARER, SVCCLU(all R) and ADAAPSPE(L) into the version 7.1 library in order to enforce this under a mixed version database routing system. The existing ADAMPM, ADAMP2 and ADARER modules in the version 7.1 library should be saved in case the previous contents of the library are needed later.

On BS2000/OSD systems, the first application requiring Adabas routing services ADAUSER/ADALNK/ADARER determines from which Adabas library the router ADARER is loaded into the ID table. Thus it is important to ensure that the highest version of the router be loaded into the IDT common memory pool to achieve full functionality.

SVC Consolidation

On z/OS and z/VSE platforms:

Previously, users needed to link special SVC extensions such as SMPSVC or SVCCLU to the ADASVC to enable Adabas SMP or Adabas Cluster Services, respectively. With Adabas version 7.4, the SVC extension SVCCLU is prelinked to ADASVC and delivered only on the base Adabas load library. This means that there is now only one official instance of the SVC. To correct errors found in this code, it is only necessary to apply Adabas ZAPs with AO and AI prefixes.

On BS2000/OSD Platforms:

SVCCLU will be automatically loaded into the IDT common memory pool when a V74 nucleus is the first to use that environment.

On VSE/ESA Platforms:

The Access Register SVC has been eliminated for VSE/ESA.

Migration Issues (IBM Platforms)

Because of the SVC consolidation, there is now an incompatibility between the SVC needed for Adabas SMP 7.1 and the prelinked ADASVC used for Adabas Cluster Services 7.4 and Adabas Parallel Services 7.4.

If you want to use Adabas SMP 7.1 rather than or in addition to Adabas Parallel Services 7.4, there are two methods you can use to resolve this incompatibility:

- you can use the Migration Tool documented in the *Adabas Installation Manual (OS/390)* if you want to access two different databases using two different SVCs concurrently from the same client.
- you can link the SMPSVC from the Adabas 7.4 library to the (prelinked) ADASVC of 7.4 if you want to use a single SVC to run, for example, one database under Parallel Services 7.4 and another database under Adabas SMP 7.1.

TP Link Routines (IBM Platforms Only)

You can run application programs interchangeably with the version 7.1, 7.2, or 7.4 link routines against either a version 7.1, 7.2, or 7.4 nucleus. Therefore, if you install the version 7.4 link routines and then decide to convert the database back to Adabas version 7.1 or 7.2, you can continue to run the application programs with the version 7.4 link routines.

When installing the version 7.4 SVC with the version 7.4 ADASIP, Software AG recommends that you also use the version 7.4 link routines.

UES-Enabled Link Routines

For Adabas version 7.4, UES is enabled by default for the batch/TSO, Complete, and IMS link routines. It is **not** necessary to disable UES support. Applications that do not require UES translation continue to work properly even when the UES components are linked with the Adabas link routines. See the *Adabas Installation Manual (OS/390)* and the *Adabas Installation Manual (VSE/ESA)* and page 40 in these *Release Notes* for more information.

However, if for some reason you feel it necessary to disable UES support in the Adabas link routines, a procedure is provided to do so.

Enqueue Conflicts

During initialization, PLOG and CLOG processing, and termination, Adabas version 7.4 uses system enqueues to coordinate certain activities. It uses the DBID as an identifier for the enqueue names.

Users who run different Adabas environments that use the same DBID within a sysplex but use different SVCs and/or LPARs may encounter situations where enqueue conflicts occur.

For example, in a situation where:

- production and test are run within the same sysplex;
- test DBID 199 experiences a problem during initialization causing it to hang while it is holding enqueue SSE00199; and
- an attempt is made to start production database 199 on a different LPAR within the same sysplex;

it will not initialize until test database 199 releases SSE00199.

Under normal circumstances, enqueue conflicts are resolved before there is any impact.

User Exit Interfaces

The interfaces to user exits are fully compatible between version 7.1, 7.2, and version 7.4; that is, existing version 7.1 or 7.2 user exits can be used without modification with version 7.4.

Conversion Requirements

The minimum requirement when converting Adabas version 7.1 (or 7.2) to version 7.4 is to replace the load library used for the nucleus and utility tasks/jobs and then run the Adabas version 7.4 conversion utility ADACNV to convert to Adabas version 7.4.

When you want to convert your test database back to version 7.1 (or 7.2), you can run ADACNV with the TOVERS=71 (or TOVERS=72) parameter setting.

New Skeleton Job Control File

Skeleton job control is contained in the DD/JCLIN file and is read as input to the RECOVER function of the Adabas Recovery Aid (ADARAI).

With Adabas version 7.4, the DD/JCLIN file has changed; Adabas version 7.4 is not compatible with the DD/JCLIN file delivered with earlier versions of Adabas. If an older version of the DD/JCLIN file is used, you will need to remove any DDRLOGM1 statements.

Adabas Recovery Aid Compatibility

The ADARAI LIST function supports Adabas version 7 and above RLOGs; Adabas version 6 RLOGs are not supported.

Migrating Products from Version 7.1 to Version 7.4

For customers needing to migrate a set of products from version 7.1 to version 7.4, we recommend migrating the add-on products first.

Adabas Data Set Compatibility

Save Data Sets

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

Restore Database

It is possible with Adabas 7.4 to restore a database from a save data set created by Adabas 7.2 or 7.4. It is also possible to use Adabas 7.2 to restore a database from a save data set created by Adabas 7.4.

Restore Files

For restoring files, it is possible

- with Adabas 7.4 to restore files from a save data set created by Adabas 7.2, 7.1, 6.2, 6.1, 5.3, 5.2, or 5.1 by using the RESTORE or RESTONL functions with the FILE or FMOVE parameters.
- with Adabas 7.2 or 7.1 (including zaps AU713084/AU713085) to restore files from a save data set created by Adabas 7.4 by using the RESTORE or RESTONL functions with the FILE or FMOVE parameters.

Considerations when Restoring Files from V5 to V7.4

- The RESTORE FILE function discards the unused RABN chains that may be present for the normal index or upper index. This makes all blocks of these chains “unreachable index blocks” as reported by the ADAICK ICHECK function in WARNING-163. These blocks will not be reused until they are reordered by the ADAORD REORFASSO or other reorder functions. The RESTORE FMOVE function does not discard the unused RABN chains, but rather transforms them to the new version 6 and above structure.

- If the database contains different device types for Data Storage and Work, restoring might be difficult if the Data Storage blocksize is larger than the Work blocksize. ADASAV may reject the restore because the maximum compressed record length of the file exceeds the length allowed by the Work blocksize. This is due to the increase in the length of protection record headers in version 6 and above. To restore the file in this case, a new Work device type must be installed using the ADADEF NEWWORK function.

Unload Data Sets

- Using version 7.4, it is possible to load files using input unload data sets created by ADAULD or ADACMP of Adabas 7.2, 7.1, 6.2, 6.1, 5.3, 5.2, or 5.1.
- Using version 7.2, 7.1, 6.2, 6.1, 5.3.4, or 5.3.3, it is possible to load files using input unload data sets created by Adabas version 7.4. This is **not** possible when using versions of Adabas prior to version 5.3.3.

When loading version 7.4 unload (ADACMP) data sets into version 5.3 databases using ADALOD, specify the input data format with the VERSION=6 parameter (version 5 does not recognize VERSION=7). The default value for this parameter is 5.

ADAORD DD/FILEA Data Sets

It is possible to run a RESTRUCTURE under

- version 7.2, 7.1, 6.2, 6.1, 5.3, 5.2, or 5.1 and a STORE under version 7.4.
- version 7.4 and a STORE under version 7.2, 7.1, 6.2, or 6.1.

It is **not** possible to run a RESTRUCTURE under version 7.4 and a STORE under version 5.

Sequential Protection Logs

Any sequential protection log (PLOG) used for a REGENERATE or BACKOUT under version 7.4 must have been created using version 7.2 or 7.4. PLOGs created using version 7.1 or earlier are incompatible with PLOGs created using version 7.2 or 7.4.

Only the ADASEL utility of versions 7.2 and above is capable of reading and processing PLOGs from both 7.4 and versions 7.2, 7.1, 6.2, 6.1, 5.3, 5.2, or 5.1. Concatenation of PLOG data sets with different versions is not possible.

ADABAS VERSION 7.4.4 RELEASE INFORMATION

These *Release Notes* accompany Adabas version 7.4.4. They contain information about the changes to Adabas since version 7.4.3.

Enhancements

Adabas version 7.4.4 is largely a maintenance release. It includes fixes and documentation updates made since Adabas version 7.4.3 was released.

The updated documentation can be reviewed on the Adabas documentation CD. This documentation can also be viewed on the Servline24 online support Web site <http://servline24.softwareag.com/public/>.

New ADARUN MXCANCEL Parameter

A new ADARUN parameter, MXCANCEL, is included in this release. This parameter is only valid for use with Adabas Parallel Services and Adabas Cluster Services. It allows you to specify the maximum time a nucleus must wait for a canceled peer nucleus to terminate before the original nucleus terminates. If no value is specified for the MXCANCEL parameter, the value of the MXMSG parameter is used.

A slight change to the MXMSG parameter is associated with the new MXCANCEL parameter in this release. The MXMSG parameter used to specify the length of time a nucleus should wait for:

- An internucleus command to be processed by a peer nucleus before the sending nucleus issues a cancel request against the unresponsive peer. With Adabas 7.4.4, the MXMSG parameter continues to specify this time-out period.
- The canceled peer nucleus to terminate before the sending nucleus terminates itself. With Adabas 7.4.4, the MXMSG parameter is no longer used to specify this time-out period. Instead, the MXCANCEL parameter now specifies this value.

For complete information on these and all ADARUN parameters, refer to your *Adabas Operations* online documentation.

New ADARUN PGFIX Parameter

A new ADARUN parameter, PGFIX, is included in this release. This parameter is only valid for z/OS users using EXCPVR (for example, when running the Adabas nucleus or utility jobs APF-authorized). It allows you to control how page-fixing occurs when EXCPVR is in use. Valid values for the PGFIX parameter are YES and NO; the default is NO. When PGFIX=NO is specified, pages containing the I/O control blocks are fixed only for the duration of the I/O processing. When PGFIX=YES is specified, pages containing the I/O control blocks are fixed for the duration of the job.

For complete information on these and all ADARUN parameters, refer to your *Adabas Operations* online documentation.

Large Format File Support

This release supports large format files introduced with IBM z/OS 1.7. Large format file support is also known as large format sequential file support. This support allows a data set to occupy more than 65,535 tracks per volume.

With Adabas large format file support, the extents of an Adabas direct access file (ASSO, DATA, WORK, TEMP, SORT, PLOG, CLOG, RLOG, or DSIM data sets) can occupy an entire volume. As before, the data set can have up to 16 extents on a volume and can span up to 59 volumes.

In addition to supporting large format files for direct access data sets, Adabas sequential data sets (typically those read or written by the Adabas utilities) can also take advantage of IBM large format file features when running on z/OS 1.7 or higher. If the IBM PARMLIB parameter, BLOCKTOKENSIZE(REQUIRE), is specified, Adabas sequential files using QSAM will support large format files on z/OS 1.7 systems.

If you need to allocate a new Adabas data set (direct access or sequential) where the space on any single volume exceeds 65,535 tracks, you must be using a z/OS 1.7 system and the DSNTYPE=LARGE parameter must be specified in the JCL DD statements used to allocate the data set. Once a large format file is allocated, the DSNTYPE=LARGE specification is no longer necessary in JCL for the file.

Warning:

If you choose to use large format files, these files will not be transportable to systems where support is not available. This might include a site being used as a backup facility for disaster recovery.

ADACDC Utility Updates

Two updates are provided in this release for the ADACDC utility.

- A new parameter, ISN, has been added to the ADACDC utility. This parameter alters ADACDC processing so that it will *not* consolidate delete and insert transactions for the same ISN, creating a single update transaction for the ISN. Instead, each delete and insert transaction is recorded in the primary output file (CDCOUT) individually. So, when you use the ADACDC utility with the ISN parameter, the number of records produced in the primary output file will increase, possibly dramatically.
- There are now some additional sorts that must be performed if you run ADACDC for a file for which ISNs can be reused (ADADBS ISNREUSE is specified). In this case, a special sort must be performed on the ADACDC output to ensure that the PLOG records are processed properly by the application to which you apply the ADACDC output. In addition, the transaction file must also be sorted so it is in correct PLOG format for any additional ADACDC runs you perform. To support these sorts, the structure of the primary output record prefix (CDCO DSECT) has been altered.

For complete information on the ADACDC utility, read the online *Adabas Utilities* documentation.

VSE and VM FBA Device Support Updated

This release includes updated FBA device support for Adabas z/VSE and z/VM database files. Before you can use this support on z/VM systems, you must be sure to add the APSV272 TXTLIB to the SETTXTLB EXEC before the ADAV744 member name.

For more information, read about FBA device support in the *Adabas Installation for z/VSE* and the *Adabas Installation for z/VM* online documentation.

Adding UES Support to VM Databases

On VM systems, you can now add Universal Encoding Support (UES) to an existing Adabas VM database. In other words, you can convert a non-UES-enabled VM database to a UES-enabled one. Before you can use this support on z/VM systems, you must be sure to add the APSV272 TXTLIB to the SETTXTLB EXEC before the ADAV744 member name.

For more information, read about UES support in the *Adabas Installation for z/VM* online documentation.

VM Installation Procedure Changes

An easier installation routine is now supported for Adabas on z/VM. Instead of using TAPPDS format, TAPE LOAD format is now used. For more information, read the *Adabas Installation for z/VM* online documentation.

Event Replicator for Adabas

While this is not directly related to Adabas 7.4.4, the Event Replicator for Adabas (ARK) has been added to the Adabas mainframe product suite this year. The documentation for this new product is included on the Adabas documentation CD.

The Event Replicator for Adabas is composed of a family of Software AG products. The basic product can be used to monitor data modifications in an Adabas database and replicate the modified data to a target application through a messaging system, such as EntireX Communicator or IBM MQSeries (also known as WebSphere MQ).

The other products that comprise the Event Replicator for Adabas family are:

- Event Replicator Target Adapter, which can be used to transform and apply replicated data to a relational database, such as Oracle, DB2, Microsoft SQL Server, MySQL, or Sybase.
- Event Replicator Administration, which is a web-based graphical user interface (GUI) you can use to perform administrative tasks for the Event Replicator.
- Entire Net-Work Administration and Entire Net-Work Client, which are prerequisites to using Event Replicator Administration to perform Event Replicator administration tasks.

Finally, updates to Adabas Online System have been made to support Event Replicator for Adabas. These are described in the Event Replicator for Adabas documentation. For complete information regarding Event Replicator for Adabas, read the online Event Replicator for Adabas documentation.

Changes in Future Versions

This section describes planned changes in the support of Adabas features and modules.

- On z/OS systems, it is currently possible to define the Adabas container files (Asso, Data, Work, PLOG, RLOG, CLOG, Temp, Sort) as VSAM relative record or linear data sets. This VSAM support for Adabas containers on z/OS systems will be removed in Adabas version 8.
- Effective with Adabas version 8, new ADALNK routines will be provided as executable modules only. Source will no longer be provided. All customization will need to be done through the before or after exits of ADALNK. The existing version 7.4 ADALNKs will still be maintained for upward compatibility; however, they will not be able to support some of the new features from Adabas version 8.

Removed from Adabas

ADALNA5 Module Dropped

The ADALNA5 source module has been dropped from the Adabas 7.4.4 source and WAL 7.4.4 source libraries. The ADALNA5 was provided for versions of Entire System Server (XCOM or Natural Process) prior to ESS Version 3.2.1. Versions of Entire System Server prior to 3.2.1 are no longer supported by Software AG. Entire System Server Version 3.2.1 and higher use the ADALNKR module to handle client communications with the Adabas router, so there is no longer any need for the ADALNA5 source module.

ADABAS VERSION 7.4.3 RELEASE INFORMATION

These *Release Notes* accompany Adabas version 7.4.3. They contain information about the changes to Adabas since version 7.4.2.

Enhancements

Adabas version 7.4.3 is a maintenance release. It includes fixes and documentation updates made since Adabas version 7.4.2 was released.

The updated documentation can be reviewed on the Adabas documentation CD. This documentation can also be viewed on the Servline24 online support Web site <http://servline24.softwareag.com/public/>.

Enhanced Interfaces between Products

Adabas SQL Gateway (ACE) supports an ODBC interface to Adabas. ACE replaces the earlier Adabas SQL Server (ESQ) product and Adabas SQL Server (AQA).

Adabas Parallel Services 7.5.1 Support

Data set FRZ.ADA743.MVSLX01, provided with Adabas 7.4.3, includes all maintenance currently available to use Adabas Parallel Services version 7.5.1. This data set supercedes the ADA742.LX01 data set referenced in the Adabas Parallel Services Release Notes.

CICS Link Routine Updates

- The default value for the EXECKEY keyword in member DEFADAC for the Adabas link routine (LNKOLSC/LNKOLM) has been changed from "USER" to "CICS". For most customers, particularly those using the LNCSTUB performance enhancement module, CICS-key is the recommended execution key for the Adabas link routine.
- The Adabas enhanced installation module, LNKENAB/ADAENAB, has been updated to load the Adabas link routine with the HOLD option to enforce the residency requirement for the module. If the load of the Adabas command-level link routine module fails, message ADAK069 is produced by the Adabas enhanced installation program.

The Adabas enhanced installation module will also release the CICS command-level link routine and disable the Adabas TRUE (Task-Related User Exit) if a failure occurs during the execution of the enhanced installation program. Two new messages, ADAK070 and ADAK071, are generated by the installation program if an error occurs.

ADARUN FORCE Parameter Changes

If you have Adabas 7.4.2 (with fix AN742281 applied) or Adabas 7.4.3 (with fix AN743052 applied), when an Adabas Parallel Services nucleus starts, the PLXNUC entries in the PLXCB structure are also scanned to ensure an active entry for the same nucleus does not exist. In addition, you can use the ADARUN FORCE parameter to indicate whether or not the Adabas Parallel Services nucleus is able to overwrite an existing PLXNUC entry should a duplicate PLXNUC entry be found. For more information, read about the ADARUN FORCE parameter in the Adabas online documentation.

ADABAS VERSION 7.4.2 RELEASE INFORMATION

Summary of Features Added for 7.4

Note: The SUSPEND/RESUME functionality is available on BS2000 platforms also.

Enhanced Interfaces between Products

1. The Adabas nucleus and Adabas add-on products are more easily integrated with improved interfaces and security checks on particular activities or functions.
2. Adabas supports cluster environments using Adabas Cluster Services (ALS) or Adabas Parallel Services (ASM). Additionally, Adabas Support for Multiprocessing (ADASMP) version 7.1 can be used with the Adabas version 7.4 router when the SMPSVC module from the Adabas 7.4.2 load library is linked with the version 7.4 ADASVC module. ALS is available for z/OS.
3. Adabas SQL Server (AQA) supports an ODBC interface to Adabas. AQA replaces the earlier Adabas SQL Server (ESQ) product.

Enhanced Support for Operating System Features

1. Adabas exploits storage occupying real pages above the 2-gigabyte line under z/OS, which means that Adabas I/Os can use 64-bit real addresses.

Support for 64-bit real storage is available whether you are running APF-authorized (using EXCP VR) or not (using EXCP). The run mode is indicated in the ADAI65 message:

ADAI65 EXCPVR IS {BEING | NOT BEING} USED FOR THIS RUN IN ESA64 MODE

Support for 64-bit real storage requires either

- OS/390 R10 in ARCHLEVEL=2 (that is, z/architecture mode); or
- z/OS 1.2 or above

—on a processor of the IBM 2064 family with an LPAR greater than 2 gigabytes for real storage allocation.

- The ADAI65 message appears in the startup messages only if these requirements are met.
2. Adabas Caching Facility (ACF) can be used to exploit 64-bit virtual storage under z/OS 1.2 or above.
 3. An Adabas nucleus under z/OS supports the use of the automatic restart management (ARM) feature for automatic restarting.

More SVC Flexibility (IBM Platforms Only)

1. Since ADASVC has been made downward compatible, you can run a version 7.1 ADASMP with a version 7.4 ADASVC.
2. Adabas releases the common storage areas (CSAs) associated with an existing SVC when running the version 7.4 ADASIP to replace that SVC with a new one.

Enhanced Backup and Restore Performance

1. Adabas exploits IBM's large block (more than 32,760 bytes) support for sequential access methods BSAM and QSAM under OS/390 version 2 release 10 and above. ADAIOR supports tape drives with a block size of up to 256K for 3590 devices and 64K for 3490/3490E devices. This support can provide performance benefits for any utility writing to tape; for example, ADASAV.

Users must ensure that they have applied the PTF for their environment that fixes IBM APAR OW55220. Without this fix, ADARES BACKOUT from a tape file written with large block support will fail.

Warning:

If you choose to write tape files with large block sizes (for example, for database backups), these files will not be transportable to systems where support for large blocks is not available. This might include a site being used as a backup facility for disaster recovery.

2. (CE) Adabas provides a new ADADBS utility function SUSPEND/RESUME to quiesce database update processing and then later resume normal processing. This function facilitates the use of non-Software AG DASD fast data set copy products such as IBM's Fastcopy and Storage Tek's Snapshot. Software AG does not recommend this technology as a substitute for normal backups using ADASAV.

Note: The SUSPEND/RESUME functionality is available on BS2000 platforms also.

File Capacity (BS2000/OSD)

From BS2000/OSD 5 onwards, you can allocate PAM or SAM files up to 4 terabytes(4000 gigabytes) in size. These files, which are larger than 32 gigabytes, can be allocated under the following conditions:

- The operating system must be OSD 5 or higher.
- The volume where the file is to be allocated must be non-Pamkey (NK2 or NK4).

Note: The SUSPEND/RESUME functionality is available on BS2000 platforms also. If it is an NK4 Disk, the default database device 2000 cannot be used because its containers have odd standard block sizes. Only devices 2002,2008,2010,2200,2201 and 2202 may be used on these volumes.

- The pubset where the file is to be allocated must allow large files by means of the SET-PUBSET-ATTRIBUTES command.
- Adabas container files are allocated during the ADAFRM utility. The container can be established using the following link command:

```
/ADD-FILE-LINK <link_name>,<file_name>,BUF-LEN=STD(<std_block_size>),-  
/ OPEN-MODE=OUTIN,EXCEED-32GB=*ALLOWED
```

Once established, it is no longer necessary to state "EXCEED-32GB=*ALLOWED". Stating one of the following will be enough:

```
/ADD-FILE-LINK <link_name>,<file_name>  
/SET-FILE-LINK <link_name>,<file_name>
```

Data Integrity Enhancements

1. Adabas now detects and reports (with response code 175) an inconsistency between the index and data storage value for a descriptor read by an L3/L6 command or an S1 command with search criterion 'descriptor=value' and a non-empty format buffer. This inconsistency is also reported by ADAVAL.
2. Changes to the RC command make it possible for Adabas to release formats (command option "D") and global formats (command option "E") for a given file number and now also descriptor name specified in the additions 1 field. These changes apply only to formats created using L3/L6 commands or S1 commands with search criterion "descriptor=value" and non-empty format buffer. The changes are based on a 12-byte format identifier that comprises an 8-byte format ID followed by a 2-byte file number and a 2-byte descriptor name.

3. With Adabas 7.4 you can alternately run cluster and noncluster nucleus sessions on a database. Switching from noncluster mode to cluster mode (and vice versa) is only possible if the previous session ended normally (that is, not after abnormal termination).

Enhanced Logging Performance

1. Adabas now supports eight (8) rather than two (2) PLOGs/CLOGs per nucleus. With more than two PLOG/CLOG data sets, an installation has more time to react to and resolve any problem in the PLCOPY/CLCOPY process before the last log data set becomes full.
2. Adabas performs CLOG I/Os asynchronously; i.e., the nucleus can continue processing while command logs are being written to disk. Additionally, the I/Os now comprise multiple blocks. These asynchronous multi-block I/Os improve the CLOG performance and extend the number of scenarios where the use of command logging is feasible with limited effect on the overall performance of the nucleus.

Other Performance Enhancements

Adabas 7.4 makes it possible to start multiple I/Os per volume in parallel to decrease buffer flush duration. This feature improves performance when the update load on the database is so high that the buffer flush becomes the bottleneck.

Enhanced Link Routines

You no longer need to modify the LRVINFO EQU (equate) in the source code for:

- Batch, Com-plete, and IMS link routines on z/OS
- Batch and Com-plete link routines on VSE/ESA
- Batch and UTM link routines on BS2000.

The default value is always set to 256. If the Review exit is link-edited with these link routines, the buffer associated with the UB (User Block) will include the 256-byte area required by the Review exit. If the Review exit is not link-edited with these link routines, then the 256-byte area will not be included in the buffer area associated with the UB. If user information is to be passed to UEXITB and or UEXITA, the LUINFO EQU (equate) value must still be provided in the source module before assembling and linking it.

For z/OS, the modules affected are ADALNK, ADALNKR, ADALCO and ADALNI. For VSE, the modules are ADALNK, ADALNKR and ADALCO. For BS2000, the module affected is ADALNK.

New and Enhanced ADARUN Options

1. The following parameters have been added to support up to eight PLOGs/CLOGs per nucleus:
 - NPLOG and NCLOG to specify the number of log data sets.
 - PLOGDEV and CLOGDEV to specify the log data set device type.
 - PLOGSIZE and CLOGSIZE to specify the log data set size.

These parameters are mutually exclusive with the parameters for establishing dual logging.

2. UEX12 has been added to support multiple data set PLOG and CLOG processing. It is mutually exclusive with UEX2, which supports only dual log processing. If NCLOG/NPLOG are being used, then UEX12 must be used *instead of* UEX2.
3. Most informational and error messages are now displayed in mixed upper/lowercase characters. The new ADARUN parameters MSGCONSL, MSGDRUCK, and MSGPRINT are available to allow translation of the messages to all uppercase for environments that use character sets in which the lowercase character codes are assigned to other characters.
4. You can now specify the files and commands to be included for prefetch/multifetch using the new ADARUN parameters PREFIFIL and PREFICMD. These commands are mutually exclusive with the commands PREFXFIL/PREFXCMD, which are used to exclude files and commands.

5. You can now increase Work part 1 (LP parameter) to a virtually unlimited size to allow more and longer incomplete transactions to be active in parallel.
6. The meaning of the FMXIO parameter now changes based on the value of ADARUN ASYTVS to accommodate the new multiple-I/O-per-volume buffer flush method.
7. You can identify with the ADARUN ARMNAME parameter the element in the automatic restart management (ARM) 'policy' that is to be activated. Adabas utilizes this feature of the z/OS operating system in both cluster and noncluster environments to automatically restart a nucleus when it unintentionally ABENDs.
8. The minimum value for LWP is now $NT * 25000$ or 80000, whichever is higher.

Feature Discussion

NUCID

Adabas Cluster Services and Adabas Parallel Services provide for a cluster of Adabas nuclei all working on the same database.

The nucleus ID (NUCID) is used to identify an individual nucleus as a member of a cluster. Standalone nuclei, which are not members of a cluster, have a NUCID of zero (the default). A value greater than zero identifies a nucleus as an Adabas cluster nucleus.

PLOGRQ=FORCE Option

Adabas 7.4 offers the new setting 'FORCE' for the PLOGRQ ADARUN parameter. The purpose of this option is to enforce that all database updates by the nucleus are backed by data on the protection log (PLOG) data sets and that no PLOGs are overwritten before they have been copied.

If PLOGRQ=FORCE has been specified, any attempt to start an Adabas update nucleus without UEX2 or UEX12, or without a protection log, will cause an initialization error. Furthermore, an error occurs if the nucleus gets into a position where it would have to overwrite the PPT entry of an uncopied PLOG data set or the data set itself.

The default setting of PLOGRQ is unchanged, in order to avoid impacting databases which do not run with UEX2/UEX12 all the time.

New Table to Track the Active Nuclei

A new table in the Associator tracks the active Adabas nucleus (or nuclei in a clustered environment) that manages a particular database. When the nucleus becomes active, it registers itself in the parallel participant table (PPT).

The PPT comprises 32 blocks: one per possible nucleus in a cluster. A single (noncluster) nucleus always occupies the first block of the table and the NUCID is always zero (0).

The PPT is created by ADADEF for a new database or by ADACNV when an existing database is converted to version 7.4. Its location is stored in the GCB. PPT entries can be modified by the nucleus or by the ADARES PLCOPY/CLCOPY functions. The new operator command DPPT displays the PPT contents. The new ADAICK PPTPRINT function prints out the PPT contents in interpreted and dump format.

If the nucleus terminates normally, the PPT entry is retained and marked as inactive. If the nucleus terminates abnormally, its entry in the PPT remains unchanged until the recovery process is complete.

Information about protection log (PLOG) data sets, command log (CLOG) data sets, and the Work data set of a nucleus is logged in its PPT entry. As each PLOG and CLOG becomes active, it is registered in the PPT. Once all PLOGs or CLOGs have been copied after session end, the PPT entry is marked accordingly. ADARES and the nucleus both check the PPT.

Note: The SUSPEND/RESUME functionality is available on BS2000 platforms also. Sequential PLOGs and CLOGs are neither recorded nor maintained in the PPT.

The PPT is used in the following situations:

- when ADARES PLCOPY or CLCOPY is copying all active protection logs or command logs, respectively. ADARES only pulls the PLOG and CLOG information from the PPT if the PLCOPY or CLCOPY is for a cluster database; otherwise, it will use the JCL provided in the PLCOPY/CLCOPY job, as it did in previous versions of Adabas.
- for warning when a nucleus session leaves the Work, PLOG, or CLOG data set non-empty and the next nucleus (i.e., the restart nucleus) starts with a different Work, PLOG, or CLOG data set.

PPT Initialization Processing

The parallel participant table (PPT) in the Associator exists for both cluster and noncluster nuclei. It is used to determine which PLOGs, CLOGs, and Work data sets have been last used by each nucleus. The following discussion applies to the session start of noncluster nuclei.

The starting nucleus uses the PPT to determine if any PLOGs or CLOGs still contain uncopied data from a previous session. If they do, the nucleus calls UEX2 or UEX12, if specified, to submit ADARES PLCOPY or CLCOPY jobs that will copy these PLOGs and CLOGs (or merge them, if they were written by cluster nuclei). If the starting nucleus shares PLOG data sets with a previously run cluster nucleus and one of the PLOGs contains uncopied data, the nucleus either reports an error (if PLOGRQ=FORCE) or marks the PLOG as empty after giving UEX2/12 a chance to copy off the protection data (otherwise).

If the first block of the PPT has been used by a cluster nucleus, the starting noncluster nucleus gives UEX2/12 a chance to copy off any uncopied PLOGs and then overwrites the PPT entry with its own PLOG and/or CLOG data set names, printing a warning message if the previously recorded PLOGs have not been copied. However, if PLOGRQ=FORCE has been specified, the data set names of the uncopied PLOGs cannot be overwritten in the PPT and the session start fails.

In any case, once the starting nucleus prepares its own PLOGs and CLOGs (specified in its JCL) for logging, it inserts their data set names into the PPT and overwrites any previously used PLOG or CLOG data set names that may have been recorded there.

In the case that uncopied PLOG or CLOG data sets remain, they can still be copied with the ADARES utility, by specifying the data sets in the ADARES JCL. Uncopied PLOGs from a cluster session can be copied and merged using the NOPPT option. Uncopied CLOGs from a cluster session can be copied and merged using the CLOG MERGE function.

Dynamic Allocation

The nucleus always writes its protection and command log data to the Work, PLOG, and CLOG data sets specified in its JCL. In a noncluster environment, the ADARES PLCOPY and CLCOPY functions always read from the PLOG and CLOG data sets specified in their JCL.

Additionally, the nucleus and ADARES use dynamic allocation (i.e., access the data sets via the names recorded in the PPT, rather than via the JCL) to perform consistency checks if the PPT indicates that the old data sets still contain needed information.

Switching Between Cluster and Noncluster Mode

Switching from cluster to noncluster mode (or vice versa) is possible only after normal termination. A starting nucleus checks in the PPT whether the previous session ended abnormally with a pending autorestart. If this is the case and the previous nucleus ran in the same mode as the starting nucleus, the session autorestart logic will be executed. If the previous nucleus ran in a different mode than the starting nucleus, the session start will terminate with an error.

The following sections illustrate a few scenarios where a noncluster nucleus starts after the normal termination of a cluster. PLOGRQ is not set to FORCE. These scenarios apply to two PLOGs as well as up to eight PLOGs.

Scenario 1

The previous session was cluster mode, there are remaining PLOGs to be copied, there is no UEX2/12 in use, and the PLOG data sets are different from what was used in the previous session.

The results of this scenario are as follows:

- The information in the current PPT entries remain, and the new entry is written.
- Initialization continues.

Scenario 2

The previous session was cluster mode, there are remaining PLOGs to be copied, there is no UEX2/12 in use, and the PLOG data sets are the same as what was used in the previous session.

The results of this scenario are as follows:

- A warning that the PLOG is being overwritten will occur and the PLOG flag in the previously used PPT block will be reset or the PPT entry will be overwritten (whichever is appropriate).
- Initialization continues.

Scenario 3

The previous session was cluster mode, there are remaining PLOGs to be copied, UEX2/12 is in use, and the PLOG data sets are different from what was used in the previous session.

The results of this scenario are as follows:

- UEX2/12 is called to submit a PLCOPY job that will copy and merge the PLOGs. The information in the current PPT entries remains, and the new entry is written.
- Initialization continues.

Note: The SUSPEND/RESUME functionality is available on BS2000 platforms also. If ADARES detects that there is data to be copied both from a cluster nuclei and from a noncluster nucleus (different PLOGs), it will copy the oldest data first.

Support for the Adabas SQL Server (AQA)

With Adabas version 7.4, SQL support has been more tightly integrated into Adabas.

Adalink ADALNKX

A special version of Adalink ADALNKX that supports XTS has been created to optimize communication between Adabas and Adabas SQL Server (AQA) address spaces. This Adalink supports Adabas add-ons such as Adabas Vista and Adabas Fastpath as well.

Delivered in the Adabas load library, ADALNKX must be part of the STEPLIB(BLSLIBnn) used for the Adabas SQL Server (AQA) job. See the Adabas SQL Server (AQA) documentation for more information.

Support for DBA Functions

Adabas version 7.4 supports new DBA functions used by Adabas SQL Server:

- Create Table (i.e., define FDT, find next file, define file) and Drop Table
- Count the number of records in a file
- Create Index and Drop Index (i.e., invert or release a descriptor)
- Alter Table Add Column (add field) and Alter Column (modify field length, format)

New Utility Function for Use with Fast Copy Products (CE)

Adabas version 7.4 provides a new ADADBS utility function SUSPEND/RESUME to quiesce database update processing and then later resume normal processing. Quiescing update processing means that once users have reached ET status (end of transaction), new update commands are delayed until normal update processing resumes. No commands are lost or rejected.

This function facilitates the use of third-party fast data set copy products such as SnapShot, FlashCopy, One Point Copy, or SnapView. Software AG does not recommend this technology as a substitute for normal backups using ADASAV.

Parameters of the new utility function set time limits so that

- running transactions have time to finish before being forcibly backed out; and
- normal processing resumes in a timely manner. If the RESUME function is not issued within the time limit, database update processing resumes automatically once the specified time period elapses.

A checkpoint is issued when the nucleus successfully quiesces update processing and another when it resumes normal processing.

Support for Assembler SYSPARM Specification of the Adabas SVC (IBM Platforms)

The Adabas 7.4 batch, COM-LETE and IMS link routines now support the specification of the default Adabas SVC number through the SYSPARM parameter specified on the PARM statement of the assembly step for the source modules. The modules that support this facility are ADALNK, ADALNKR, ADALNA5, ADALCO, and ADALNI.

Previous versions of these modules required editing the source member to change the SVCNR EQU to override the default value of 249 for the z/OS Adabas SVC number. Now the value may be specified in the assembly JCL on the SYSPARM assembler step PARM statement. The value provided will be checked by conditional assembly statements to ensure that it is a valid decimal number in the range 200–255.

If an improper SYSPARM value is provided, an MNOTE is generated and the assembly step will terminate with condition code 16. If no value is coded on the SYSPARM assembly step PARM statement, the value coded on the SVCNR EQU statement in the source module will be used. This permits either the specification of the default Adabas SVC number by editing the source member and then assembling it, or through the SYSPARM JCL parameter.

Note: The SUSPEND/RESUME functionality is available on BS2000 platforms also. It is still necessary to modify the supplied source members to set other default values such as LOGID or LNUINFO. Refer to the Adabas Installation manual for more information.

Messages in Mixed Case (CE)

Adabas version 7.4 messages are displayed in mixed upper-/lowercase. If you want to display one or more types of these messages in uppercase, you can use the new ADARUN parameter MSGCONSL for console messages, MSGDRUCK for messages sent to DD/DRUCK, and MSGPRINT for messages sent to DD/PRINT. The first line in DDPRINT will not be influenced by the value of parameter MSGPRINT.

Note: The SUSPEND/RESUME functionality is available on BS2000 platforms also. An optional Zap AY744001 available in the ZAPOPT member will set uppercase as the default for console messages, DDPRINT and DDRUCK, no matter what is set for the ADARUN parameters.

Automatic Restart Management (ARM) Support

Automatic restart management (ARM) is a feature of the z/OS operating system that is installed and maintained by the system programmer. ARM is used for automatically restarting failed servers according to the rules provided to it. Adabas utilizes this feature to automatically restart a nucleus when it ABENDs. Automatic restart is suppressed when the ABEND is intentional; for example, when it results from a parameter error.

ARM can be used for Adabas nuclei in both cluster and noncluster environments.

The ADARUN parameter ARMNAME is used to identify the element in the ARM ‘policy’ that is to be activated. Each element specifies when, where, and how often an automatic restart is to be attempted.

If an ARM policy has not been defined, the ARMNAME parameter has no effect.

New Buffer Flush Method

Adabas now offers three different methods for the asynchronous buffer flush:

1. The buffer flush activated with ADARUN ASYTVS=NO is asynchronous in ascending RABN sequence without regard to the distribution of the blocks over volumes. This was the only buffer flush method available prior to Adabas version 7.1.
2. The buffer flush activated with ADARUN ASYTVS=YES and FMXIO=1 is asynchronous in parallel over all database volumes with one I/O per volume. This method was added with Adabas version 7.1.

3. The buffer flush activated with ADARUN ASYTVS=YES and FMXIO greater than one is asynchronous in parallel over all database volumes with multiple parallel I/Os per volume for PAV (Parallel Access Volume) devices. This method is available for Adabas version 7.2.2 and above.

This new buffer flush method (3 above) reduces the duration of buffer flushes. It improves performance when the update load on the database is so high that the buffer flush becomes the bottleneck.

With method 3, a buffer flush initially starts a predetermined number of I/Os on each volume and then starts a new one once another I/O on the same volume finishes. This occurs independently on each volume.

The previously used buffer flush method (2), which remains available, starts only one I/O per volume at a time.

Meaning of ADARUN FMXIO Parameter Changed

The meaning of the FMXIO parameter has changed for the new buffer flush method.

When ASYTVS=YES (buffer flushes occur by volume), FMXIO now specifies the number of I/Os to be started in parallel **on each volume**. The minimum and default number is 1; the maximum number is 16. If you specify a number greater than 16, it is reduced to 16 without returning a message.

Note: The SUSPEND/RESUME functionality is available on BS2000 platforms also. If ASYTVS=YES is in effect in your current ADARUN parameters and you have specified the FMXIO parameter, you need to consider setting a new value of FMXIO according to its new meaning.

When ASYTVS=NO (perform buffer flushes in ascending RABN sequence, without regard to the distribution of the blocks over volumes), FMXIO continues to specify the number of I/Os to be started in parallel overall. The minimum value is 1, the maximum value is 100, and the default value is 60.

Dynamically Modifying the FMXIO Parameter Setting

The setting of FMXIO can be modified dynamically using the FMXIO=nn command from the operator console or the Modify Parameter function of Adabas Online System.

New ADARUN Parameters to Include Files in Prefetch/Multifetch

The new ADARUN parameters PREFIFIL/PREFICMD make it possible to “include” lists of files in prefetching/multifetching.

Along with the current parameters PREFXFIL/PREFXCMD for “excluding” lists of files in prefetching/multifetching, it is now possible to include or exclude lists of files depending on which list is shorter.

Only one set of parameters may be used at a time.

Asynchronous CLOG I/O

Because it is not necessary to synchronize the writing of command logs to disk with any other activity in the nucleus, this I/O is now done asynchronously. Performance gains are realized because the nucleus continues processing while the I/O is in progress.

Additional performance gains are realized by support for multiblock I/Os. The CLOG provides an audit trail and can be used for debugging and for monitoring the use of resources. Single, dual, or multiple (2–8) data sets can be used (multiple data sets are recommended).

Enhancements

CSAs of Replaced SVCs Released (IBM Platforms Only)

When you run ADASIP to replace an existing Adabas SVC with a new one, ADASIP now releases almost all common storage areas (CSAs) associated with the SVC being replaced, including the SVC code and the ID table. The CSAs released are identified in the new ADAS20 message.

Releasing the CSAs means that this finite resource remains available for other uses in the system.

The only common storage that remains is the first eight (8) bytes of the ID table in which an “eye-catcher” is changed to indicate the released ID table. The version 7.2.2 and above link routines use this eye-catcher to verify that the ID table address they saved from earlier calls is still valid.

Earlier versions of the link routines do not check this eye-catcher in the ID table. If you dynamically install the version 7.4.2 SVC but retain earlier version link routines, you may encounter a program check (S0C4) depending on whether the older link routine has saved and reuses the old address of the ID table.

If you install the new version 7.4.2 SVC and you use it from Complete or CICS, do one of the following:

- install the new SVC permanently, not dynamically “over” an existing SVC instance; or
- install the version 7.4.2 link routine in the TP monitor to begin with; or
- bounce the TP monitor after installing the new SVC to force the existing, earlier version link routine to obtain the new ID table address.

Using the new version SVC in batch presents no special problems for the Adabas link routines.

Different ISN Returned with Record-related Response Code

For record-related nonzero response codes such as 113 or 177, Adabas now returns to the application the ISN of the record that caused the response code. For nonrecord-related response codes such as 3 or 17, Adabas continues as before to return the ISN provided by the application.

This change is intended to alleviate any confusion that may have arisen when a record-related response code returned the ISN provided by the application rather than the ISN of the record that caused the response code.

ADASAV Processing Change

For the following ADASAV functions:

RESTONL (database),
RESTONL GCB;
RESTORE (database),
RESTORE GCB

—in order to avoid accidentally overwriting a log data set containing important data, the DD/PLOGRn and DD/CLOGRn data set headers are no longer reset (making the data sets logically empty) in the restore process.

The nucleus will not start if the database has been restored and the dual PLOG data sets are not empty. This can happen, for example, if you restore a test database that you ran with dual PLOGs but without UEX2. If the nucleus refuses to start because the dual PLOG data sets are not empty, you must run either ADARES PLCOPY to copy off the latest protection data or ADAFRM PLOGFRM to clear the PLOG contents.

Number of Blocks Allocated to Work Data Protection Area Increased

The number of blocks allocated to the Work data protection area is no longer limited to 65535 blocks; it is now only limited by the size of the Work data set minus the sizes of Work parts 2, 3, and 4. The number of blocks allocated is specified with the LP parameter. See the *Adabas Operations Manual* for more information.

Link Routines are UES-enabled by Default

The Adabas version 7.4 link routines for batch/TSO, Com-plete, and IMS/DC are UES-enabled by default.

Load Modules

The load modules ADALNK, ADALNKR, and ADALCO in the Adabas load library ADA74x.MVSLOAD (ADALNK in ADA74x.MOD on BS2000) have been linked with LNKUES and the default translation tables ASC2EBC and EBC2ASC. The load module LNKUESL has been dropped from the ADA74x.MVSLOAD (ADALNK in ADA74x.MOD on BS2000) library since the link jobs now specify the LNKUES module and the translation tables separately.

Note: The SUSPEND/RESUME functionality is available on BS2000 platforms also. The use of UES-enabled link routines should be transparent to applications, including applications that do not require the UES conversion of Adabas calls.

Source Modules

The source modules ADALNK, ADALNKR, ADALCO, and ADALNI in the source libraries ADA74x.MVSSRCE and ADALNK in ADA74x.SRC (on BS2000) are set to enable UES by default. This means that the &UES Boolean assembly variable is set to 1; the statement to set the variable to 0 has been commented out. The setting of other Boolean variables and equates such as the SVC number and database ID remains unchanged.

Job Members (IBM Platforms Only)

The members ALNKLCO, ALNKLNK, and ALNKLNKR in the ADA74x.MVSJOBS library have been modified to assemble and link the ADALCO, ADALNK, and ADALNKR modules, respectively, with the UES components into the Adabas load library. The following step assembles the two default translation tables, ASC2EBC and EBC2ASC, into the Adabas load library. Finally, the link routine is linked with LNKUES, ASC2EBC, and EBC2ASC and the resulting load module is put in a “user” load library.

Conversion of High Value in Value Buffer

When UES=YES and Alpha (or Wide) conversion is used, all FROM-TO Search/Logical Read Criteria are handled in such a way that in the TO criterion the high value characters at the value end are preserved when converted into the internal search value and are excluded from value conversion.

This prevents a problem that may have occurred when performing searches using the S operator with UES=YES, i.e., that the X'FF' would not always remain the X'FF' in the converted value. For more information, see the *Adabas DBA Reference Manual*.

PRILOG Now Delivered in Both Source and Object Form

The PRILOG program used in conjunction with user exit 4 is now provided in both source and object form. The object form in the load library can be used without modification.

Nucleus Processing When All PLOG Data Sets are Full

Prior to Adabas 7.2, the entire Adabas nucleus went into a wait status when both PLOG data sets were full. No commands were processed by the Adabas nucleus until one of the PLOG data sets was copied by ADARES PLCOPY. With Adabas 7.2 and above, the nucleus continues to process access and operator commands when all PLOG data sets are full.

ADACOX (Adabas Conversion Exit)

With Adabas 7.4.2 the new conversion exit ADACOX is available. With UES-enabled databases, ADACOX supports “context-sensitive” conversion between Windows-1256 and EBCDIC Arabic or EBCDIC Farsi code page.

While Arabic characters are ‘unshaped’ in Windows–1256, the supported EBCDIC encodings use ‘shaped’ forms, depending on previous or following characters. In addition, for certain consecutive characters the combined form is used (for example, LAM–ALEF ligature).

Following is some information about ADACOX implementation in Adabas 7.4.2:

- Currently, no support is included for conversion from logical to visual order and vice versa, symmetric character conversion.
- The conversion exit will always be loaded for UES enabled databases.
- When a new conversion between two encodings is first used, the exit is queried as to whether it supports the conversion. If it does, the exit will be called for any such conversion; otherwise, Adabas and/or Entire Conversion Services will do the conversion. Still, the conversions defined in ADACOX need to be backed by corresponding ECS objects. For example, for the conversion of 420 to 1256, the character set properties are determined by the ECS objects.

Conversion of Superdescriptor Search Values

In alignment with Adabas Open/Systems, the conversion of the superdescriptor parts is performed according to its parent’s format.

Previously, the superdescriptor format was used for conversion. For example, this would have converted binary data using ASCII–EBCDIC conversion with superdescriptors composed of alpha and binary fields.

Removed from Adabas

Support Removed for Alternate RABNs

Alternate RABNs, which were recorded in the Block of Unreadable Blocks (BUB), are no longer supported.

Mirror Data Set Support Removed

Data set mirroring, including PLOG mirroring and RLOG mirroring, is no longer supported.

Adabas ESA Option (AEO) No Longer Available

The Adabas ESA Option is no longer available as a product. Its components Adabas Dynamic Caching and Adabas Support for Multiprocessing (ADASMP) have been replaced by the selectable units Adabas Caching Facility (ACF) and Adabas Parallel Services (ASM), respectively.

The Adabas version 7.4 SVC/router allows you to continue to use the version 7.1 ADASMP as needed.

For VSE, additional functions now available without the Adabas ESA Option include support for buffers above the 16-megabyte line and dynamic partitions.

ADAESI No Longer Available

The Adabas External Security Interface (ADAESI) is no longer available. It has been replaced by Adabas SAF Security (AAF), which provides ADAESI functionality as part of its service.

Adabas SAF Security is a selectable unit. Special consideration will be given to those who already use ADAESI. AAF is available for OS/390 only.

Interlink TCP/IP No Longer Supported

The Adabas TCP/IP direct link supports the application programming interface (API) for the IBM TCP/IP stacks (HPS, OES) only. The API for the Interlink TCP/IP stack is no longer supported.

Access Register SVC Under VSE/ESA

The Adabas SVC under VSE has been redesigned in Adabas 7.4.2 to take advantage of new technology that will significantly enhance system performance. Therefore, as announced in the Adabas version 7.1.3 Release Notes, support for the access register SVC has been removed in Adabas 7.4.2.

Dropped Support for the Adabas CICS Macro-level Link Routine (ADALNC)

IBM has dropped CICS macro-level support with their latest releases of VSE. For this reason, Software AG will no longer distribute the CICS macro-level link routine, ADALNC, in the VSE Adabas or WAL sublibraries. Customers still running CICS/VSE 2.3 or earlier may want to retain a copy of ADALNC from the Adabas 7.1.3 or WAL 7.1.3 sublibraries for use in these older releases of CICS/VSE.

Dropped Support for DSDEV Subparameter

When decompressing files using the ADACMP DECOMPRESS function, you used to be able to use the DSDEV subparameter to specify a device type for a file identified by the INFILE parameter. Support for the DSDEV subparameter has been dropped in 7.4 because Adabas can now automatically determine the device type when the data set is opened.

Changes to Adabas Online System

This section summarizes the changes to the Adabas Online System in version 7.4.2.

Product Interfaces

Depending on the platforms supported, demo versions of Adabas Vista (AVI), Adabas Fastpath (AFP), Adabas SAF Security (AAF), and Adabas Transaction Manager (ATM) are automatically installed when you install either the demo or full version of Adabas Online system (AOS). AAF is available for OS/390 only; ATM is available for z/OS, VSE/ESA, and BS2000.

AOS version 7.4 requires Natural version 3.1 or above.

If you do not upgrade or install Natural with this Adabas release, you will need to relink the Adabas version 7.4 AOSASM module with your Natural in order for AOS 7.4 to function correctly. This AOSASM module is downward compatible with previous versions.

Installation Requirement

AOS 7.4 is distributed via a single ALLINPL data set containing AOS, Adabas Triggers and Stored Procedures, and Adabas Caching Facility. The delivered ALLINPL file, ADA743.ALLINPL, contains the Adabas Online System (AOS) base and also functions as a demo.

To use the full version of AOS, you must copy the contents of the AOS version 7.4 load library to the base Adabas version 7.4 load library or concatenate the AOS library to the base Adabas library. When allocated to the “running” Adabas environment, this completes the AOS installation to make it a licensed copy.

The full version AOS selectable unit is required for the triggers and stored procedures facility.

If you try to run AOS version 7.1 against Adabas version 7.4, AOS will function as a demo.

If you try to run AOS version 7.4 with an earlier version of ADASVC, AOS will function as a demo. For ADASVC 7.1.2 and 7.1.3, special ZAPs can be requested from Software AG Customer Support.

If you try to run AOS version 7.4 with an earlier version of Adabas, AOS will function as a demonstration system, unless the following Adabas Zaps are applied, as follows:

Adabas Version	Zap Number(s)
7.1.2	A\$712099, A\$712063 *
7.1.3	A\$713099, A\$713031 *
7.2.2	AY722099

* ZAPs A\$712063 or A\$713031 update the ADASVC version 7.1.2 or version 7.1.3.

Beginning with ADASVC version 7.2.2, no zap for the ADASVC is required.

‘In Parallel’ Option

The option ‘In Parallel’ has been added to the Adabas Online System. This new option allows processing to continue against a file while the function you have selected modifies the file. The ‘In Parallel’ option is available in Adabas Online System when adding a new field, releasing a descriptor, and modifying file parameters.

When using the 'In Parallel' option, file integrity is always maintained. Maintaining this integrity, however, could cause unexpected behavior in the applications that are executing against the file. For this reason, you should be aware of the impact on your operating environment and the consequences of performing functions against a file before using this option.

For example, if the 'In Parallel' option is active and you add a new field or release a descriptor, any application running against the file containing the new field might experience a performance slowdown; in some instances, any application running against the file containing the affected field might become unstable, and could even terminate abnormally.

ZAP Prefixes for Adabas Online System (AOS)

The ZAP format is Axvrsnnn where "x" identifies the component affected by the ZAP. The following ZAP values are used for the AOS Natural components:

- AAvrsnnn is a ZAP that affects the Natural components of the AOS product; and
- AZvrsnnn is a ZAP that affects the Natural components of the demo version of AOS.

ZAPs that make changes reflected in AOS may not affect AOS itself, but rather some part of the underlying Adabas. For example:

- ANvrsnnn is a ZAP that affects the Adabas nucleus; or
- AUvrsnnn is a ZAP that affects an Adabas utility.

Updated Adabas Online System Screens

The following screen additions and changes were made to Adabas Online System in version 7.4.2:

- From the Session Monitoring menu, a new function **Display cluster members** (option A) produces the Display Cluster Members screen. The screen includes a list of nuclei participating in the cluster and information about the current status of each nucleus.
- From the Resource Utilization menu, the **Nucleus file status** (option N) has been added and is the equivalent of the DNFV operator command.
- From the Resource Utilization menu, a second screen has been added to the **System status** (option S), which displays I/O counts for the ASSO, DATA, WORK, and PLOG data sets; remote and local call distribution; and other current session status information.

- From the Resource Utilization menu, **Cluster usage** (option X) displays nucleus cluster statistics that are equivalent of those displayed using the DXCACHE, DXLOCK, and DXFILE operator commands. See the *Adabas Cluster Services Manual* or the *Adabas Parallel Services Manual* for more information (*available in Adabas cluster environments only*).
- A new function has been added to the Session Opercoms menu to support the CLUFREEUSER command (*available in Adabas cluster environments only*).

INDEX

A

- active nuclei, 31
- Adabas, support for earlier releases, 3
- Adabas Online System, 44
 - 'In Parallel' option, 45
 - installation requirements, 45
 - ZAP prefixes, 46
- Adabas SVC, 26, 35
- ADACDC utility, 19
 - ISN parameter, 19
- ADACMP, version compatibility, 16
- ADACOX, 41
- Adalink, 34
- ADALNA5 source module, 21
- ADALNKX, 34
- ADALOD, LOAD VERSION parameter, 16
- ADAORD, version compatibility, 16
- ADARAI, compatibility, 14
- ADARUN FORCE parameter, 24
- ADARUN options, 29
- ADARUN parameters
 - MXCANCEL, 17
 - MXMSG, 17
 - PGFIX, 18
- ADASAV
 - version compatibility, 15
 - Work data set reset, 39
- ADASIP, releasing common storage, 38
- ADAULD, version compatibility, 16
- AOS, 44
- AQA support, 34
- ARM, 30
 - explanation, 36
- ARMNAME, ADARUN parameter, 36
- asynchronous CLOG I/O, 38
- ASYTVS, ADARUN parameter, 37
- Automatic Restart Management, 30, 36

Automatic Restart Management (ARM), 36

B

- backup and restore performance, 26
- BLOCKTOKENSIZE parameter, 18
- Buffer flush, 36

C

- CE, 34
- Cluster
 - display usage statistics, using Basic Services, 47
 - maintain user table, 47
- cluster mode vs. noncluster mode, 33
- Communication modules, version compatibility, 9

D

- data integrity enhancements, 27
- DBA function support, 34
- DD/JCLIN file, 14
- device support, FBA devices, 19
- DSNTYPE parameter, 18
- Dynamic allocation, 32

E

- end of support, 4
- Event Replicator for Adabas, 20
- EXCPVR processing under z/OS, 18

F

- fast copy products (CE), 34
- FBA device support, 19

Adabas Version 7.4 Release Notes

features summary, 25
file capacity (BS2000), 27
files, large format file support, 18
FMXIO, ADARUN parameter, 37

I

installation changes, z/VM, 20
ISN parameter, 19

L

large format file support, 18
link routines, 29, 40
logging performance, 28

M

messages (mixed case), 36
mixed case messages, 36
MXCANCEL parameter, 17
MXMSG parameter, 17

N

noncluster mode vs. cluster mode, 33
NUCID, 30
Nucleus file, display status, using Basic Services, 46
Nucleus ID, 30

P

page fixing, 18
Parallel Participant Table, 30, 31
performance
 logging, 28
 other enhancements, 28
PGFIX parameter, 18
PLOGRQ=FORCE, 30

PPT, 30, 31
PPT initialization processing, 31
PREFICMD, ADARUN parameter, 38
PREFIFIL, ADARUN parameter, 38
previous version support, 4
PRILOG, 41

R

Response codes, record-related, 39
restore and backup performance, 26

S

Sequential PLOGs, version compatibility, 16
Session, display status, using Basic Services, 46
Sixty-four (64) bit storage support
 real, 25
 virtual, 26
superdescriptor search values, 42
SVC, 26, 35
 and link routines, 38
SYSPARM, 35
System, display status, using Basic Services, 46

U

UES, see Universal Encoding Support (UES), 19
Universal Encoding Support (UES, for VM data-bases, 19
User table, maintain, 47
Utilities, ADACDC updates, 19

V

value buffer, 41
VSAM support for z/OS Adabas containers, 21

W

Work data set, reset by ADASAV, 39

Z

- z/VM installation procedure changes, 20
- ZAP prefixes, Adabas Online System, 46

Adabas Version 7.4 Release Notes

Adabas Version 7.4 Release Notes