

Adabas® Version 7.4.2

Release Notes

Manual Order Number: ADA742-008IBB

This document applies to the Adabas software package at version 7.4.2 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) June 2003 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 |
| Enhanced SQL Support | 1 |
| Enhanced Support for Multiple-Engine Processors | 1 |
| About this Document | 2 |
| Documentation | 2 |
| Printed Manuals | 2 |
| Online Documentation | 2 |
| 1. ENVIRONMENT | 5 |
| Operating Systems Supported | 5 |
| Support for Previous Versions | 6 |
| Support for 64-Bit Storage | 8 |
| Real Storage | 8 |
| Virtual Storage | 9 |
| Optional ZAPs Dataset Added | 9 |
| Software AG Product Level Compatibility | 9 |
| Adabas and Natural | 9 |
| Entire Net-Work | 9 |
| Adabas SQL Server | 10 |
| Adabas Add-on Products Overview | 10 |
| Adabas Caching Facility | 11 |
| Adabas Cluster Services | 11 |
| Adabas Delta Save Facility | 12 |
| Adabas Fastpath | 12 |
| Adabas Online System | 13 |
| Adabas Parallel Services | 13 |
| Adabas Review | 14 |

Adabas Version 7.4.2 Release Notes

| | |
|--|-----------|
| Adabas SAF Security (OS/390 and z/OS) | 14 |
| Adabas Statistics Facility | 15 |
| Adabas Transaction Manager (OS/390, z/OS, VSE/ESA, and BS2000) | 16 |
| Adabas Vista | 16 |
| Adabas Component Compatibility | 16 |
| Communication Modules | 16 |
| SVC Consolidation | 18 |
| Migration Issues (IBM Platforms) | 19 |
| TP Link Routines (IBM Platforms Only) | 19 |
| Enqueue Conflicts | 20 |
| User Exit Interfaces | 20 |
| Conversion Requirements | 20 |
| New Skeleton Job Control File | 20 |
| Adabas Recovery Aid Compatibility | 21 |
| Migrating Products from Version 7.1 to Version 7.4 | 21 |
| Adabas Dataset Compatibility | 21 |
| Save Datasets | 21 |
| Unload Datasets | 22 |
| ADAORD DD/FILEA Datasets | 22 |
| Sequential Protection Logs | 23 |

| | |
|---|-----------|
| 2. RELEASE NOTES FOR ADABAS VERSION 7.4.2 | 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 (IBM Platforms Only) | 26 |
| File Capacity (BS2000) | 26 |
| Data Integrity Enhancements | 27 |
| Enhanced Logging Performance | 27 |
| Other Performance Enhancements | 27 |
| Enhanced Link Routines | 28 |
| New and Enhanced ADARUN Options | 28 |
| Feature Discussion | 29 |
| NUCID | 29 |
| PLOGRQ=FORCE Option | 29 |
| New Table to Track the Active Nuclei | 29 |
| Switching Between Cluster and Noncluster Mode | 32 |
| Support for the Adabas SQL Server (AQA) | 33 |
| New Utility Function for Use with Fast Copy Products (CE) | 33 |
| Support for Assembler SYSPARM Specification of the Adabas SVC (IBM Platforms) | 34 |
| Messages in Mixed Case (CE) | 35 |
| Automatic Restart Management (ARM) Support | 35 |
| New Buffer Flush Method | 35 |
| New ADARUN Parameters to Include Files in Prefetch/Multifetch | 37 |
| Asynchronous CLOG I/O | 37 |
| Enhancements | 37 |
| CSAs of Replaced SVCs Released (IBM Platforms Only) | 37 |
| Different ISN Returned with Record-related Response Code | 38 |
| ADASAV Processing Change | 38 |
| Number of Blocks Allocated to Work Data Protection Area Increased | 39 |
| Link Routines are UES-enabled by Default | 39 |
| Conversion of High Value in Value Buffer | 40 |
| PRILOG Now Delivered in Both Source and Object Form | 40 |
| Nucleus Processing When All PLOG Datasets are Full | 40 |
| ADACOX (Adabas Conversion Exit) | 40 |
| Conversion of Superdescriptor Search Values | 41 |

Adabas Version 7.4.2 Release Notes

| | |
|---|-----------|
| Removed from Adabas | 41 |
| Support Removed for Alternate RABNs | 41 |
| Mirror Dataset Support Removed | 41 |
| Adabas ESA Option (AEO) No Longer Available | 42 |
| ADAESI No Longer Available | 42 |
| Interlink TCP/IP No Longer Supported | 42 |
| Access Register SVC Under VSE/ESA | 42 |
| Changes to Adabas Online System | 43 |
| Product Interfaces | 43 |
| Installation Requirement | 43 |
| 'In Parallel' Option | 44 |
| ZAP Prefixes for Adabas Online System (AOS) | 44 |
| Updated Adabas Online System Screens | 45 |
| 3. DOCUMENTATION UPDATES | 47 |
| Updates to the Messages and Codes | 48 |
| New ADAM Messages | 48 |
| Updated ADANX1 Message Explanation | 48 |
| New ADAX Messages | 49 |
| New PARM Errors | 49 |
| Response Code 9, Subcodes 80-99 | 50 |
| Response Code 22, Subcode 24 | 50 |
| Response Code 48, Subcodes 25-30 | 50 |
| Response Code 88 | 51 |
| Response Code 219, Subcodes 9-10 | 51 |
| Updated ADARES Messages | 51 |
| New ADARES Messages | 52 |
| New ADASAV Error Message | 53 |
| Updates to User Abend Codes | 54 |
| Updated Messages for ADACOX | 54 |
| New ADACOX Messages | 55 |

Table of Contents

| | |
|--|----|
| Updates to the Adabas Online System Documentation | 56 |
| WORK Status Screen Updated | 56 |
| Additional Option When Adding a New Field | 58 |
| Updated Release Descriptor Screen | 60 |
| Updated Modify File Parameters Screen | 61 |
| Updated Display Cluster Members Screen | 61 |
| Updated Nucleus File Status Screen | 64 |
| Updated Nucleus Status Flags | 65 |
| Updated Cluster Usage Screen | 67 |
| Updated User Table Maintenance Screen | 68 |
| Updates to the Adabas Installation Documentation (OS/390 and z/OS) | 71 |
| ADAGSET Option “ESI” | 71 |
| Adabas Utilities | 72 |
| ADADBS Functional Overview | 72 |
| ADAFRM Formatting Database Components | 72 |
| ADARES Updates | 72 |
| ADARES Backout Functions Information | 72 |
| New ADARES PLCOPY Parameter | 73 |
| Appendix C: Supplied UES Encodings | 73 |
| Updates to the Triggers and Stored Procedures | 75 |
| Updates to the Adabas Operations | 75 |

| | |
|---|-----------|
| APPENDIX A - THE AFPLOOK UTILITY | 77 |
| Overview | 77 |
| How to Run AFPLOOK | 78 |
| AFPLOOK Parameters | 78 |
| Maximum Files | 78 |
| Command/Descriptors per File | 78 |
| Maximum Concurrent Users | 79 |
| Maximum CIDs per User | 79 |
| Maximum Commands Processed | 79 |
| Job Name | 79 |
| Selected Files | 80 |
| Number of Lines per Page | 80 |
| Customer Heading | 80 |
| On/Off Switch | 80 |
| Overriding Parameter Defaults | 81 |
| AFPLOOK Sample File Report | 81 |
| AFPLOOK Potential Optimization Summary | 82 |
| AFPLOOK Sample Command Analysis | 83 |
| AFPLOOK Sample Parameters Report | 84 |
| | |
| APPENDIX B - THE AVILOOK UTILITY | 85 |
| File Maintenance | 86 |
| List AVILOOK Files | 86 |
| Add AVILOOK Files | 88 |
| Maintain AVILOOK Files | 88 |
| | |
| INDEX | 91 |

PREFACE

About this Release

Adabas version 7.4 merges the functionality of the current, generally available Adabas version 7.1 and subsequent versions of Adabas (versions 7.2 and 7.3) that include changes under OS/390 or z/OS only for Adabas Cluster Services and 64-bit virtual support. A few new version 7.4 features have also been added.

All Adabas add-on products have been aligned with the comprehensive, single code base of Adabas version 7.4.

Enhanced SQL Support

Adabas version 7.4 provides support for a closer integration with Adabas SQL Server (AQA) and its underlying database management language (DBML) SQL. At the same time, Adabas SQL Server version 5.1, which is being released in tandem with Adabas version 7.4, has been enhanced to work more closely with the Adabas product family.

Enhanced Support for Multiple-Engine Processors

In conjunction with the release of Adabas version 7.4, Adabas Parallel Services (ASM) is being introduced as a multi-update nucleus replacement for ADASMP. Adabas Parallel Services, which provides nucleus cluster support on a single operating image, is based on the code for Adabas Cluster Services, which supports nucleus clusters across multiple OS/390 or z/OS operating system images in an IBM parallel sysplex environment.

About this Document

These *Release Notes* accompany Adabas version 7.4.2. They contain information about the changes to Adabas since version 7.1.

Chapter 1 contains information about the required environments.

Chapter 2 contains the release notes for Adabas version 7.4.2.

Chapter 3 contains documentation updates for Adabas version 7.4.2.

Appendix A describes the AFPLOOK utility, the Adabas Fastpath command analysis sampler.

Appendix B describes the AVILOOK utility, an Adabas Vista analysis tool.

Documentation

Both printed manuals and online documentation are provided with the Adabas version 7.4.2 release.

Printed Manuals

The following printed manuals are being delivered with the version 7.4.2 release:

| Order No. | Manual |
|------------------|---|
| ADA742-008IBB | <i>Adabas Version 7.4.2 Release Notes</i> (this document) |
| ADA742-010... | <i>Adabas Installation</i> (platform-specific) |

Online Documentation

The documentation for mainframe Adabas Version 7.4.2 is being provided on CD as HTML files with their PDF equivalents. It is also available from the Software AG web site.

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

Installation and Getting Started

- Concepts and Facilities
- Release Notes 7.4.2 (PDF)
- Installation for OS/390 (PDF)
- Installation for BS2000 (PDF)
- Installation for VSE/ESA (PDF)
- DBA Tasks
- Operations

Reference

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

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

For detailed information on changes made to the documentation specifically for the Adabas 7.4.2 release, please see Chapter 3, **Documentation Updates**, in this document.

ENVIRONMENT

The release of Adabas version 7.4.2 for VSE/ESA, z/VM, and BS2000 environments in addition to OS/390 and z/OS is accompanied by

- this document;
- an *Adabas Installation Manual* for each new environment.
- a CD containing Adabas documentation in online format.

These *Release Notes* comprise the version 7.4.1 release notes and all necessary additions and corrections for version 7.4.2.

Note:

The term “dataset” is used as a general term in these release notes. For BS2000 environments, this term indicates a PAM or SAM file, depending on the application.

Operating Systems Supported

Adabas version 7.4.2 is released for the following operating environments:

- OS/390 version 2, release 10
- z/OS version 1, releases 1-4
- z/VM versions 4.2, 4.3, and 4.4
- VSE/ESA version 2, releases 5, 6, and 7
- BS2000 OSD 2.0 and above (3.0 and above if using Universal Encoding Support)

Support for Previous Versions

Software AG plans to provide support for previous Adabas releases and SMs as follows:

| Adabas Product and Release | Release Date | End of Support Date |
|-----------------------------------|---------------------|----------------------------|
| Adabas 7.1.1 | 30 September 1999 | 28 July 2000 |
| Adabas 7.1.2 | 12 April 2000 | 31 October 2002 |
| Adabas 7.1.3 | 24 August 2001 | 31 May 2005 |
| Adabas 7.2.1 | 2 April 2001 | 29 March 2002 |
| Adabas 7.2.2 | 2 November 2001 | 31 May 2005 |
| Adabas 7.4.1 | 2 January 2003 | 30 June 2004 |
| Adabas 7.4.2 | 13 June 2003 | |
| Adabas CICS Interface 7.1.3 | 24 August 2001 | 31 May 2005 |
| Adabas CICS Interface 7.4.1 | 6 December 2002 | 30 June 2004 |
| Adabas CICS Interface 7.4.2 | 13 June 2003 | |
| Adabas IMS Interface 7.1.3 | 24 August 2001 | 31 May 2005 |
| Adabas CICS Interface 7.4.1 | 6 December 2002 | 30 June 2004 |
| Adabas CICS Interface 7.4.2 | 13 June 2003 | |
| Adabas IMS Interface 7.1.3 | 24 August 2001 | 31 May 2005 |
| Adabas IMS Interface 7.4.1 | 6 December 2002 | 30 June 2004 |
| Adabas IMS Interface 7.4.2 | 13 June 2003 | |
| Adabas UTM Interface 7.1.3 | 24 August 2001 | 31 May 2005 |
| Adabas UTM Interface 7.4.2 | 13 June 2003 | |
| Adabas Cluster Services 7.2.1 | 2 April 2001 | 29 March 2002 |
| Adabas Cluster Services 7.2.2 | 27 December 2001 | 31 May 2005 |
| Adabas Cluster Services 7.4.2 | 1 August 2003 | |
| Adabas Native SQL 2.2.1 - MF | 27 March 2002 | 31 May 2005 |
| Adabas SQL Server 1.4.3 - MVS | 2 August 1999 | 31 December 2004 |
| Adabas SQL Server 1.4.3 - VSE | 2 August 1999 | 31 December 2004 |

| Adabas Product and Release | Release Date | End of Support Date |
|-----------------------------------|---------------------|----------------------------|
| AQA SQL Server 5.1.1 - MF | 20 December 2002 | 31 December 2003 |
| AQA SQL Server 5.1.2 - MF | 17 April 2003 | |
| Jadabas 7.1.1 | 9 July 2001 | |
| Adabas Delta Save 7.2.1 | 2 April 2001 | 29 March 2002 |
| Adabas Delta Save 7.2.2 | 27 December 2001 | 31 May 2005 |
| Adabas Delta Save 7.4.1 | 6 December 2002 | 31 December 2004 |
| Adabas Delta Save 7.4.2 | 13 June 2003 | |
| Adabas Fastpath 7.1.3 | 7 September 2001 | 31 May 2005 |
| Adabas Fastpath 7.3.1 | 27 December 2001 | 31 May 2005 |
| Adabas Fastpath 7.4.2 | 19 September 2003 | |
| Adabas Caching Facility 7.3.1 | 23 November 2001 | 31 May 2005 |
| Adabas Caching Facility 7.4.1 | 6 December 2002 | 31 December 2004 |
| Adabas Caching Facility 7.4.2 | 13 June 2003 | |
| Adabas ESI 7.1.3 | 24 August 2001 | 31 May 2004 |
| Adabas SAF Security 7.1.3 | 7 September 2001 | 31 May 2004 |
| Adabas SAF Security 7.3.1 | 27 December 2001 | 31 May 2005 |
| Adabas SAF Security 7.4.2 | 13 June 2003 | |
| Adabas Vista 7.1.3 | 7 September 2001 | 31 May 2005 |
| Adabas Vista 7.3.1 | 21 December 2001 | 31 May 2005 |
| Adabas Vista 7.4.2 | 19 September 2003 | |
| Adabas Transaction Manager 1.2.1 | 1 August 2000 | 31 October 2002 |
| Adabas Transaction Manager 1.2.2 | 27 December 2001 | 31 May 2005 |
| Adabas Transaction Manager 7.4.2 | 19 September 2003 | |
| Adabas Online Services 7.1.3 | 24 August 2001 | 31 May 2005 |
| Adabas Online Services 7.2.1 | 24 August 2001 | 30 September 2003 |
| Adabas Online Services 7.2.2 | 24 August 2001 | 31 May 2005 |
| Adabas Online Services 7.4.1 | 6 December 2002 | 31 December 2004 |

| Adabas Product and Release | Release Date | End of Support Date |
|----------------------------------|------------------|---------------------|
| Adabas Online Services 7.4.2 | 13 June 2003 | |
| Adabas Review 4.2.2 | 1 June 2001 | 31 May 2004 |
| Adabas Review 4.3.1 | 6 June 2002 | 31 December 2004 |
| Adabas Review 4.3.2 | 8 August 2003 | |
| Adabas Statistics Facility 7.1.1 | 30 November 1999 | 31 May 2001 |
| Adabas Statistics Facility 7.1.2 | 1 May 2000 | |
| Adabas DL/I Bridge 2.2.1 | 31 January 1995 | 30 September 1996 |
| Adabas DL/I Bridge 2.2.2 | 5 February 1996 | |
| Adabas VSAM Bridge VSE 3.4.1 | 27 February 1996 | 30 April 2003 |
| Adabas VSAM Bridge OS/390 4.2.1 | 19 June 1998 | 30 April 2003 |
| Adabas VSAM Bridge 5.1.1 | 26 April 2002 | |

Support for 64-Bit Storage

Real Storage

Adabas can now exploit storage occupying real pages above the 2-gigabyte line. This capability allows Adabas I/Os to 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.

Virtual Storage

IBM supports 64-bit virtual storage only for z/OS 1.2 or above.

Software AG provides support for IBM's 64-bit virtual storage with a new product Adabas Caching Facility (ACF; see page 11). Contact your Software AG account representative for more information.

Optional ZAPs Dataset Added

The distributed source library contains a new 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 Level Compatibility

Adabas and Natural

Adabas version 7.4 is downward compatible with Adabas version 7.1 and above.

Natural version 3.1 or above is required to run the Natural (online) components associated with Adabas version 7.4 and its add-ons.

Entire Net-Work

If you use Entire Net-Work with Adabas version 7.4, you must use Entire Net-Work version 5.8.1 or above.

The following Entire Net-Work ZAP must be applied to version 5.8:

| ZAP | Solves a problem with . . . |
|------------|------------------------------------|
|------------|------------------------------------|

| | |
|---------|---|
| WM58028 | Improper re-translation of ASCII client calls to non-mainframe remote targets that do not use the EBCDIC character set. |
|---------|---|

With Adabas version 7.4.2 the “limited” WAL742 library containing platform-dependent modules is also delivered. This library contains the Adabas communication environment (router) used by the Entire Net-Work product. It is optional for users to upgrade their Entire Net-work installation to use this version of the WAL library; however, when using the Adabas Cluster Services or Adabas Parallel Services product, you *must* upgrade the Entire Net-work product in order to utilize the ADASVC and ADALNK / ADALCO components from the WAL742 library and apply the Entire Net-work zap WM58028.

Adabas SQL Server

Adabas SQL Server (AQA; formerly ESQ) provides a standard ANSI/ISO SQL interface for Adabas. It enables SQL access to existing Adabas data structures and the implementation of SQL-based applications.

Adabas version 7.4 is specifically designed to work with Adabas SQL Server (AQA) version 5.1 or above. However, you may also use Adabas version 7.4 with Adabas SQL Server (ESQ) version 1.4 or above.

With Adabas version 7.4.2 the “limited” WAL742 library containing platform-dependent modules is also delivered. This library contains the Adabas communication environment (router) used by AQA. It is optional for users to upgrade their AQA installation to use this version of the WAL library; however, when using the Adabas Cluster Services or Adabas Parallel Services product, you *must* upgrade AQA in order to utilize the ADASVC and ADALNK / ADALCO components from the WAL742 library.

Adabas Add-on Products Overview

Adabas version 7.4 supports add-on products as follows:

| Adabas Product | Compatible Version Level |
|---|--------------------------------------|
| Caching Facility (ACF) | version 7.3 (initial release) or 7.4 |
| Cluster Services (ALS) (OS/390 and z/OS) | version 7.4 |
| Delta Save Facility (ADE) | version 7.4 |
| Fastpath (AFP) | version 7.4 |
| Online System (AOS) | version 7.4 |

| Adabas Product | Compatible Version Level |
|---|--|
| Parallel Services (ASM) | version 7.4 (initial release) Support for Multiprocessing (ADASMP) has been replaced by Adabas Parallel Services for version 7.4; however, you have the option to run ADASMP version 7.1 with the Adabas version 7.4 router. See page 19. |
| Review (REV) | version 4.3 Upgrade Adabas Review to version 4.3 first; then upgrade Adabas to version 7.4. |
| SAF Security (AAF) (OS/390 and z/OS) | version 7.4 |
| Statistics Facility (ASF) | version 7.1; ASF is not supported with ALS or ASM. |
| Transaction Manager (ATM) (OS/390, z/OS, VSE/ESA, BS2000) | version 7.4 |
| Vista (AVI) | version 7.4 |

Adabas Caching Facility

Available on request, Adabas Caching Facility (ACF; formerly Adabas Dynamic Caching) boosts the performance of Adabas. It augments the Adabas buffer pool and so reduces the number of read I/Os to the database. ACF provides buffer space in addition to what is available to the Adabas buffer pool. Also, ACF can be used to specifically cache database blocks belonging to certain files or residing on certain disk volumes.

If you are running z/OS version 1.2 or above, Adabas Caching Facility provides support for 64-bit virtual storage.

Adabas Cluster Services

Adabas Cluster Services (ALS) implements multinucleus, multithread parallel processing and optimizes Adabas in an OS/390 or z/OS parallel sysplex (SYStems comPLEX) environment. Multiple Cluster Services nuclei running in different operating system images can work on the same Adabas database. Each nucleus can perform updates to the database. The nuclei are fully symmetric; they all have the same capabilities.

In addition to the increased throughput that results from parallel processing, Adabas Cluster Services increases database availability during planned or unplanned outages: the database can remain available when a particular operating system image or cluster nucleus requires maintenance or goes down unexpectedly.

To support a cluster environment that includes more than one operating system image, a limited Entire Net-Work library is included as part of Adabas Cluster Services. Entire Net-Work is used to send Adabas commands from users to remote cluster nuclei, if no proper local cluster nucleus is available for command processing. As many as 32 nuclei can run in an Adabas Cluster Services cluster. No changes to application programs are required to utilize Adabas Cluster Services.

Adabas Delta Save Facility

Adabas Delta Save Facility (ADE) offers significant enhancements to ADASAV utility processing by backing up only the changed (delta) portions of Adabas databases.

Adabas Delta Save Facility is intended for Adabas sites with one or more large, heavily updated databases that need to be available most of the time. Especially for sites where the volume of data changed on a day-to-day basis is considerably smaller than the total database volume, Adabas Delta Save Facility provides for

- more frequent saves without interrupting database availability;
- enhanced “24 x 7” operation;
- full offline saving in parallel with the active database; and
- shorter REGENERATE duration during recovery.

Adabas Fastpath

Adabas Fastpath (AFP) reduces CPU consumption and increases throughput for Adabas systems by processing Adabas calls to completion in the client process whenever possible. Database processing is thus avoided and capacity is increased for the whole operation.

Adabas Fastpath optimizes two types of database calls: direct access and sequential access. The results of direct access calls to the database are saved in a cache and repeat calls are satisfied from the cache rather than directly from the database. Read-ahead optimization routines are applied to sequences of commands to reduce redundant activity and thus accelerate the sequence.

Adabas Online System

Adabas Online System (AOS) provides the DBA with an online facility for administration tasks. AOS version 7.4 provides support for nucleus cluster environments, support for SAF Security, and the ability to dynamically modify ADARUN parameters.

AOS is delivered in a separate dataset/library from Adabas. In the basic Adabas delivery, AOS functions as a demonstration system. For full functionality, the contents of the AOS 7.4.2 load library must be copied to the Adabas 7.4.2 load library, or the AOS 7.4.2 load library be concatenated in the STEPLIB of the Adabas nucleus JCL.

In addition, all AOS users (demo or full version) and Pre-dict users must include the load module AOSASM from the Adabas 7.4.2 load library in the link of the Natural nucleus.

Adabas Parallel Services

Adabas Parallel Services (ASM) implements multinucleus, multithread parallel processing and optimizes Adabas in a multiple-engine processor environment on a single operating system image. Up to 31 Adabas nuclei in an Adabas Parallel Services “cluster” are distributed over the multiple engines that are supported by the operating system.

Note:

The precursory product, Adabas Support for Multiprocessing (ADASMP), provided single nucleus update and multinucleus read capability. Adabas Parallel Services (ASM) provides multinucleus update capability.

All nuclei in the cluster access a single physical database simultaneously. The nuclei communicate and cooperate with one another to process the users’ work. Compression, decompression, format buffer translation, sorting, retrieving, searching, and updating operations can all occur in parallel.

Any number of Adabas Parallel Services clusters can operate in the same operating system image under the same or different Routers or SVCs; that is, any number of separate databases can be processed, each with its own Adabas Parallel Services cluster of up to 31 nuclei.

Applications see only one database target; no interface changes are required. Applications still communicate with their intended databases and communicate with an Adabas Parallel Services cluster of nuclei without modification.

In addition to the increased throughput that results from parallel processing, Adabas Parallel Services increases database availability during planned or unplanned outages: the database can remain available when a particular cluster nucleus requires maintenance or goes down unexpectedly.

Adabas Review

Note:

Adabas version 7.4 requires Adabas Review version 4.3. It is important to upgrade Adabas Review to version 4.3 before you upgrade Adabas to version 7.4. See the Review 4.3 Release Notes for how to run Review 4.3 together with earlier versions of Adabas 7.

Adabas Review (REV) provides a set of monitoring, accounting, and reporting tools that enable you to monitor the performance of the Adabas environment and the applications executing within them. Information retrieved about Adabas usage helps you tune application programs to achieve maximum performance with minimal resources.

In addition to the “local” mode with Adabas Review running in the Adabas address space, Adabas Review offers the “hub” mode, a client/server approach to the collection of performance data for Adabas:

- the Adabas Review interface (the client) resides on each Adabas nucleus.
- the Adabas Review hub (the server) resides in its own address space, partition, or region.

New Cluster Services statistics gathering has been introduced to Adabas version 7.4 and Adabas Review version 4.3. Both file-level (CF) caching statistics and Cluster Services locks can be monitored over a period of time and at user defined intervals. Statistics are written to the Review History file and can be retrieved and/or viewed using the power of Review reporting facilities. For additional information refer to the Adabas Review documentation.

Adabas SAF Security (OS/390 and z/OS)

Adabas SAF Security (AAF or ADASAF) is a selectable unit of Adabas that enhances the scope of standard security packages based on the System Authorization Facility (SAF) such as RACF, CA-ACF2, and CA-Top Secret to encompass Adabas resources. It integrates Adabas into a central security repository and enables you to derive maximum benefit from your investment in that repository.

Adabas SAF Security is fully compatible with existing security definitions used by the earlier ADAESI product. ADASAF is described in the *Adabas SAF Security Manual*.

In addition to the features provided by the former ADAESI, the following features are available with the ADASAF license:

- cross-level checking
- protection for transaction data
- passwords and cipher codes
- protection for Adabas Online System (AOS) basic services
- online system
- remote user protection, other than from PC windows

Adabas Statistics Facility

Adabas Statistics Facility (ASF) provides an automated environment for database monitoring.

A “store” program collects database status information during an active nucleus session. The store program is normally scheduled to run at regular intervals (for example once per day) over a period of many weeks or months to collect data that can be statistically evaluated. The store program can also be started by the DBA on an ad hoc basis, using commands in the ASF online menu system.

A set of “evaluation” programs interpret the statistics gathered by the store program and publish summary “evaluation” reports to either the screen or a hardcopy printer. Reports may also be downloaded to a PC using Entire Connection.

Database information can be collected at the start, at the end, and during a nucleus session. The start and end nucleus data, when accumulated over periods of weeks or months, gives an indication of long term database growth and permits projections of future database requirements. The nucleus performance data, such as main memory and pool usage, permits the DBA to analyze and tune the Adabas nucleus parameters.

Adabas Transaction Manager (OS/390, z/OS, VSE/ESA, and BS2000)

Adabas Transaction Manager introduces distributed transaction support to Adabas. Transactions may be distributed over multiple Adabas databases in one or more operating systems (connected by Entire Net-Work). Transactions may also be distributed over non-Adabas DBMSs such as DB2, IMS, and so on. Where transactions are distributed over a non-Adabas DBMS, Adabas Transaction Manager must be configured to inter-operate with other transaction coordinators such as IBM's CICS Syncpoint Manager and/or IBM's Recoverable Resource Management Services. At any time, ATM can account for in-flight transactions, suspect transactions, participating databases, and more.

ATM addresses two basic needs of the enterprise object revolution:

- the need to deliver industrial strength enterprise objects for widespread commercial use in mainstream, critical business systems.
- the need to spread the masses of data that Adabas customers manage more evenly across the computer(s) and organization.

ATM includes an online administration system based on Natural and available through Adabas Online System.

Adabas Vista

Adabas Vista (AVI) is used to partition an Adabas file into multiple, separate files, each containing a part of the original larger whole. Partitioning makes it possible to distribute a file across multiple volumes or computers based on a criterion such as region or date. Adabas Vista is also used to translate Adabas database and file numbers, which allows an application to remain independent of the underlying physical layer.

Adabas Component Compatibility

Communication Modules

The Adabas communication modules include ADALNK, ADAMPM, ADASIP, ADASIR, and the Adabas router/SVC [ADASVC/ADARER (BS2000)].

Software AG recommends that all applications use the ADALNK module from a currently supported Adabas release. Effective with Adabas version 8.1, a new instance of the Adalink will be provided and no source will be provided. All customizations will need to be done through the Before or After Exit of Adalink. The existing version 7.4 Adalinks will still be maintained for upward compatibility; however, they will not be able to support some of the new features from Adabas version 8.

It is recommended to link ADAUSER, rather than ADARUN or ADALNK, to the component application to ease migration of these to new Adabas versions.

Software AG recommends that you install the version 7.4 ADASVC/ADARER (BS2000) and API link routines.

Note:

On BS2000 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.

In a single nucleus (noncluster) environment

- the version 7.4 ADASVC/ADARER (BS2000) is compatible with down-level Adabas databases; that is, you can run nuclei of Adabas version 7.1 or 7.2 with the installed Adabas 7.4 SVC.
- you can run Adabas version 7.4 with version 7.2 or 7.1 ATLASES/ADARERs (BS2000). *However, when using Adabas version 7.4, it is preferable to use the new SVC in order to utilize its full functionality.*

In a cluster environment, the version level of Adabas, the ADASVC/ADARER (BS2000), and the cluster product (Adabas Cluster Services or Adabas Parallel Services) must match.

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 19.

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.

With Adabas 7.4, ADASIP, ADASVC, and ADALNKs can be used with Adabas version 7.1 and Adabas version 7.2 (i.e., they are downward compatible).

Notes:

Adabas version 7.4 is compatible with version 7.2 or 7.1 ATLASES/ADARERs (BS2000). However, when using Adabas version 7.4, it is preferable to use the new SVC in order to utilize its full functionality.

On BS2000 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 ADAMP2, ADARER, SVCCLU and ADAAPSPE(L) into the version 7.1 library in order to enforce this under a mixed version database routing system. The existing ADAMP2 and ADARER modules in the version 7.1 library should be stored.

On BS2000 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 router be loaded into the IDT common memory pool to achieve full functionality.

SVC Consolidation

On OS/390, z/OS, and VSE/ESA 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 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 7.4.2 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 39 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 Dataset Compatibility

Save Datasets

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

Restore Database

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

Restore Files

For restoring files, it is possible

- with Adabas 7.4 to restore files from a save dataset 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 dataset 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 Datasets

- Using version 7.4, it is possible to load files using input unload datasets 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 datasets 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) datasets 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 Datasets

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 datasets with different versions is not possible.

RELEASE NOTES FOR ADABAS VERSION 7.4.2

Summary of Features Added for 7.4

Note:

The term “dataset” is used as a general term in these release notes. For BS2000 environments, this term indicates a PAM or SAM file, depending on the application.

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 OS/390 and 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 OS/390 and z/OS, which means that Adabas I/Os can use 64-bit real addresses.
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 OS/390 or 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 (IBM Platforms Only)

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 dataset 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.

File Capacity (BS2000)

PAM or SAM files larger than 32 gigabytes can be accessed if running on BS2000 OSD 5 or higher.

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 datasets, an installation has more time to react to and resolve any problem in the PLCOPY/CLCOPY process before the last log dataset runs full and the nucleus stalls.
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 without too much drag 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

For the Adabas batch, Com-plete and IMS link routines on OS/390 (z/OS) and for the batch and Com-plete link routines for VSE/ESA, it is no longer necessary to modify the LRVINFO EQU (equate) in the source code.

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 OS/390 and z/OS, the modules affected are ADALNK, ADALNKR, ADALCO and ADALNI. For VSE, the modules are ADALNK, ADALNKR and ADALCO.

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 datasets.
 - PLOGDEV and CLOGDEV to specify the log dataset device type.
 - PLOGSIZE and CLOGSIZE to specify the log dataset size.

These parameters are mutually exclusive with the parameters for establishing dual logging.
2. UEX12 has been added to support multiple dataset 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 MSGCONSO, 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 OS/390 and z/OS operating systems 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) datasets 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 dataset or the dataset 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) datasets, command log (CLOG) datasets, and the Work dataset 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:

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 dataset non-empty and the next nucleus (i.e., the restart nucleus) starts with a different Work, PLOG, or CLOG dataset.

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 datasets 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 datasets 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 dataset names, printing a warning message if the previously recorded PLOGs have not been copied. However, if PLOGRQ=FORCE has been specified, the dataset 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 dataset names into the PPT and overwrites any previously used PLOG or CLOG dataset names that may have been recorded there.

In the case that uncopied PLOG or CLOG datasets remain, they can still be copied with the ADARES utility, by specifying the datasets 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 datasets specified in its JCL. In a noncluster environment, the ADARES PLCOPY and CLCOPY functions always read from the PLOG and CLOG datasets specified in their JCL.

Additionally, the nucleus and ADARES use dynamic allocation (i.e., access the datasets via the names recorded in the PPT, rather than via the JCL) to perform consistency checks if the PPT indicates that the old datasets 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 datasets 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 datasets 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 datasets 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:

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 dataset 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-PLETE 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:

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:

An optional Zap AY742001 available in the ZAPOPT member will set upper case 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 OS/390 and z/OS operating systems 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:

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) datasets can be used (multiple datasets 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 dataset containing important data, the DD/PLOGRn and DD/CLOGRn dataset headers are no longer reset (making the datasets logically empty) in the restore process.

The nucleus will not start if the database has been restored and the dual PLOG datasets 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 datasets 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 dataset 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 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 Datasets are Full

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

Dataset 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.

Changes to Adabas Online System

This section summarizes the changes to the Adabas Online System in version 7.4.2. See Chapter 3, **Documentation Updates**, in these Release Notes for detailed information about Adabas Online System documentation updates.

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, OS/390, 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 dataset containing AOS, Adabas Triggers and Stored Procedures, and Adabas Caching Facility. The delivered ALLINPL file, ADA742.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.

See Chapter 3, **Documentation Updates**, in these Release Notes for Adabas Online System documentation updates related to the ‘In Parallel’ option.

ZAP Prefixes for Adabas Online System (AOS)

The ZAP format is Axvrsn 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 datasets; 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*).

Note:

See Chapter 3, **Documentation Updates**, in these Release Notes for detailed information about Adabas Online System documentation updates.

DOCUMENTATION UPDATES

This chapter includes updates to information published in the following Adabas 7.4 documentation:

- *Adabas Messages and Codes*
- *Adabas Online System*
- *Adabas Installation (OS/390 and z/OS)*
- *Adabas Utilities*
- *Triggers and Stored Procedures*
- *Adabas Operations*

The changes detailed in this chapter are included in the documentation CD distributed with the Adabas 7.4.2 release (for contents of the CD, see section **Online Documentation** in the **Preface**). For documentation that is not yet available online (from this list above, this includes the *Adabas Online System* and the *Triggers and Stored Procedures*), the updates detailed in this chapter will be included in the next documentation release.

Updates to the Messages and Codes

This section describes the changes to the *Messages and Codes* for the 7.4.2 release.

New ADAM Messages

The following new ADAM messages were added:

ADAM88 ddddd Processor(s) = n

Explanation: BS2000 systems only. This is an information message stating how many processors (n) are available on the computer on which the database ddddd is started.

User Action: None.

ADAM94 ddddd Post Module ADAAPSPE loaded

Explanation: BS2000 systems only. This is an information message stating that the Smarts(APS) posting module ADAAPSPE has been loaded into the router. This is an acknowledgement that access to Smarts applications is now enabled.

User Action: None.

Updated ADANX1 Message Explanation

The message explanation for ADANX1 is updated as follows:

**ADANX1 COMMAND cmd COMMAND-ID hex-cid FNR
file-number RESPONSE rsp-code SUBCODE
rsp-subcode FLD field-name TID hex-internal-userid
UID open-userid JOB job-name**

Explanation: Format of the diagnostic information produced by the Adabas PINRSP or PINUES routine. The PINRSP routine will produce only the cmd, hex-cid, rsp-code, and rsp-subcode information.

New ADAX Messages

The following new ADAX messages were added:

ADAX79 dbid Global resource lock on this system is invisible to NUCID nucid on the system sysname

Explanation: Cluster Services nuclei working on the same database synchronize some of their actions using resource locks via Global Resource Serialization (GRS). During session start one nucleus detected that a resource lock it obtained was not effective against the peer nucleus with the NUCID shown, which was active on the system shown. The starting nucleus terminates with parm-error 105.

User Action: Contact your system programmer to ensure that GRS is configured in a way that GRS resource locks are mutually effective against one another on all systems on which you intend to run Cluster Services nuclei.

New PARM Errors

The following new PARM errors were added:

PARM ERROR 105

Explanation: Improper configuration of Global Resource Serialization (GRS). A resource lock acquired by this nucleus was ineffective against a peer nucleus.

User Action: Contact your system programmer to ensure that GRS is configured in a way that GRS resource locks are mutually effective against one another on all systems on which you intend to run Cluster Services nuclei.

PARM ERROR 106

Explanation: When LOCAL=YES is specified, all cluster nuclei for a database must start on the same system.

User Action: If LOCAL=YES is the intended parameter setting, start all cluster nuclei for the database on the same system. Otherwise, change the parameter setting to LOCAL=NO.

Response Code 9, Subcodes 80-99

The following information was added to the description for these subcodes:

80 - 99 CL is invalid because the user has a transaction in progress that is managed by Adabas Transaction Manager. These subcodes are given only when Adabas Transaction Manager is in use. Please see the Adabas Transaction Manager documentation for more information.

Response Code 22, Subcode 24

The following explanation was added for subcode 24 of response code 22 in the *Adabas Messages and Codes*:

24 CL is invalid because the user has a transaction in progress that is managed by Adabas Transaction Manager.

Response Code 48, Subcodes 25-30

The following explanation was added for subcodes 25-30 of response code 48 in the *Adabas Messages and Codes*:

25-30 Reserved for Adabas Transaction Manager. These subcodes are given only when Adabas Transaction Manager is in use. Please see the Adabas Transaction Manager documentation for more information.

Response Code 88

The following explanation was updated for response code 88 in the *Adabas Messages and Codes*:

Explanation: This response code generally occurs when there is a workpool (LWP) shortage. When this response code occurs, refer to the hexadecimal sub-code that appears in the low-order (rightmost) two bytes of the additions 2 field. Following are the subcodes and their meanings.

Response Code 219, Subcodes 9-10

The following explanations were updated for subcodes 9 and 10 of response code 219 in the *Adabas Messages and Codes*:

- 9 The system is currently locked due to a pending Work area 4 overflow.
- 10 The call is not permitted; either ADAEND or an ET-syncpoint is in process for the nucleus.

Updated ADARES Messages

The message text for ADARES Error and Warning messages ERROR-154 through ERROR-157 were updated as follows (“MERGINTI” is now “MERGIN1” and “MERGINTO” is now “MERGIN2”):

ERROR-154 { OPEN | CLOSE } ERROR ON { MERGIN1 | MERGIN2 } FILE

**ERROR-155 READ ERROR ON { MERGIN1 | MERGIN2 } FILE.
SYSTEM ERROR=nnnnnnnn**

**ERROR-156 WRITE ERROR { MERGIN1 | MERGIN2 } FILE.
SYSTEM ERROR=nnnnnnnn**

ERROR-157 INCORRECT { MERGIN1 | MERGIN2 } FILE SUPPLIED

The following error message was *updated*:

**ERROR-051 ERROR OCCURRED DURING EXECUTION OF OPEN dataset/file :
reason**

User Action: The correct JCL must be supplied to the PLCOPY or CLCOPY function. If the PLOGs to be copied are from a cluster database, the NOPPT parameter should be specified.

If no JCL is supplied or the open fails, the following will be displayed: 'The Associator dataset(s) could not be opened. nn - return code explanation.' Processing continues.

New ADARES Messages

The following messages were added for ADARES:

ERROR-170 Parameter SBLKNUM is only allowed if NOPPT

Explanation: The SBLKNUM parameter is only allowed in conjunction with the NOPPT parameter.

**ERROR-171 The PPT is required for a PLCOPY from a cluster database.
This is not possible because the open of the ASSO failed.
Please rerun the PLCOPY with NOPPT specified and the
PLOG DD-statements supplied in the JCL.**

Prior to ERROR - 171, the following should be seen: 'The Associator dataset(s) could not be opened. nn - return code explanation.' Processing continues.

Explanation: The open of the ASSO failed. After checking the PLOGs, it was determined that the PLOG to be copied was from a cluster database. Either the ASSO must be restored or the PLCOPY NOPPT function must be executed to copy off the PLOGs. All pertinent PLOG DD-statement from the cluster should be supplied as input to the NOPPT function.

WARNING - CLOG is from a cluster database. A normal CLCOPY will run since the PPT is not accessible. A subsequent CLOG MERGE must be run to merge the CLOG datasets.

Prior to the warning, the following should be seen: 'The Associator dataset(s) could not be opened. nn - return code explanation. ' Processing continues.

Explanation: The open of the ASSO failed. After checking the CLOGs, it was determined that the CLOG to be copied was from a cluster database. A normal CLCOPY (non-merge) will run. A subsequent CLOG should run to merge the CLOG datasets.

New ADASAV Error Message

The following new ADASAV message was added:

ERROR-143 RESTORE DATABASE NOT PERMITTED. ANOTHER UTILITY HAS EXCLUSIVE CONTROL OF THE DATABASE.

Explanation: ADASAV RESTORE expects to get a Response Code returned that indicates that the database is down. However, ADASAV RESTORE was returned a sub code indicating that another utility currently has exclusive control of the database.

User Action: Ensure the correct nucleus has been accessed; if so, wait for the other utility to run before running the ADASAV RESTORE.

Updates to User Abend Codes

The following user Abend codes were previously undocumented:

| Code | Module | Explanation |
|------|-----------|--|
| 15 | (nucleus) | Workpool too small to execute session autorestart |
| 16 | (nucleus) | DTP=RM: Two-phase commit logic error |
| 17 | (nucleus) | DTP=RM: Work-part-4 overflow |
| 19 | (nucleus) | DTP=RM: Two-phase commit logic error |
| 21 | (nucleus) | I/O error during asynchronous buffer flush (Asso/Data/Work/PLOG) |
| 22 | (nucleus) | I/O error on sequential PLOG and PLOGRQ=YES/FORCE in effect |
| 31 | (nucleus) | Insufficient workpool space for only active command |
| 33 | (nucleus) | Logic error during file number validation |
| 33 | ADARAC | Recovery Aid logic error |
| 34 | utilities | Abnormal termination with dump |
| 37 | (nucleus) | ET post logic error |
| 38 | (nucleus) | Unexpected response code for internal command |
| 40 | (nucleus) | Getmain failure early during session start |
| 40 | ADACOM | Fatal error during ADACOM processing (cluster environments) |
| 48 | (nucleus) | Logic error in ADANCX (cluster environments) |
| 49 | (nucleus) | General internal error (various reasons) |
| 70 | (nucleus) | Parallel Services nucleus canceled by peer nucleus |

Updated Messages for ADACOX

The following existing messages were updated to support the new ADACOX conversion exit:

For Utility Status message **ADAU7A** and Console message **ADAN7A**, the following row was added to the table of ECS functions:

| | |
|----------|---|
| COX LOAD | An error loading ADACOX. ADACOX is loaded if the database is UES-enabled. ADACOX is the Adabas conversion exit for special conversions. |
|----------|---|

New ADACOX Messages

The following new messages were added to support the new ADACOX conversion exit:

COX01I- ADACOX EXIT Vv.r.1 date ACTIVE

Explanation: Informational message showing the version and assembly date of the ADACOX module.

COX02I- ADACOX CONVERTING source TO target

Explanation: Informational message indicating that conversions from source to target will be done by ADACOX.

COX03I- ADACOX source TO target COUNT number AVG time (usec)

Explanation: Informational message at session termination. This message is output only if ADACOX has performed conversions between source and target encodings. *number* shows the number of conversions and time the average time per conversion in microseconds.

Updates to the Adabas Online System Documentation

The following updates have been made to the *Adabas Online System documentation*.

WORK Status Screen Updated

In the **Session Monitoring** chapter, “Resource Utilization” section, “Work Status” subsection, the WORK Status screen was updated to accommodate the display of the Distributed Transaction Processing Area field.

Below is the updated screen capture:

```

15:51:32          ***** A D A B A S  BASIC  SERVICES  *****          2003-04-30
DBID 1955          -  WORK Status  -          PACUW02

                W O R K  Dataset
+-----+
I  Protection Area                500 Blks I
I-----I
I  Intermediate ISN Area          76 Blks I
I-----I
I  Resulting ISN Area            1315 Blks I
I-----I
I  Distributed Transaction Processing Area  500 Blks I
+-----+

PF1----- PF2----- PF3----- PF4----- PF6----- PF7----- PF8----- PF12-----
Help                Exit        DTP-area                Menu

```

If you have DTP=RM in your ADARUN parameter settings, and press PF4 on the new screen, then the new screen DTP Work Area displays (shown below):

```

21:14:47          ***** A D A B A S  BASIC  SERVICES *****          2003-05-14
                    - DTP Work Area -                               PACUK02

          Code      Service
          -----
          D          Display PET-status users
          H          Display heuristically terminated users
          R          Display DTP rabns
          U          Work Part 4 usage
          X          Force heuristic BT/ET
          ?          Help
          .          Exit
          -----

Code ..... _
Selected User ....
Database ID ..... 1955   (WIS1955)

Command ==>
PF1----- PF2----- PF3----- PF4----- PF6----- PF7----- PF8----- PF12-----
Help      Clear UID Exit                               Menu

```

The DTP Work Area menu performs the following functions:

| Function | Action |
|--|--|
| Display PET-status users | Displays the current environment of PET-status users. You can select one or more to display additional information. |
| Display heuristically terminated users | Displays a list of the heuristically terminated users. You can select one or more to display additional information. |
| Display DTP Rabns | Displays a list of allocated RABNs. You can select a user, display additional information, issue BT or ET, and display total RABNs for a user. |
| Work Part 4 Usage | Display counters and information about Work Part 4. |
| Force Heuristic BT/ET | Issues a BT or ET for a selected user. Select the user through the “Display DTP RABNs” menu item. |
| PF2 | Clears the selected user. |


```

16:29:13          ***** A D A B A S  BASIC SERVICES *****      2003-03-04
DBID 1955          - Add New Field(s) -                               PFLCA02

File = 1          (EMPLOYEES)                                       In Parallel .. NO
Enter Password if file is security protected ...

```

| Level | Name | Length | Format | Options | |
|-------|------|--------|--------|---------|---|
| I | I | I | I | I | I |
| I | I | I | I | I | I |
| I | I | I | I | I | I |
| I | I | I | I | I | I |

The option 'In Parallel' allows processing to continue against a file while the function you have selected modifies the file. Specify 'YES' to select this option.

With this option, Adabas performs the function while all applications executing against the file are temporarily quiesced and suspended. With this option set to 'NO' (the default), Adabas requires exclusive file control (that is, no application executing against the file) when performing the function.

The 'Add new field(s)' function corresponds to the Adabas utility function ADADBS NEWFIELD. The equivalent direct command is

ADD FIELD

Updated Release Descriptor Screen

In the **File Maintenance** chapter, “Release Descriptor” section, the screen capture has been updated to accommodate the display of the new ‘In Parallel’ option. The updated screen capture is below:

```
16:17:28          ***** A D A B A S  BASIC  SERVICES  *****      2003-03-04
DBID 1955          -  Release Descriptor  -                          PFLD002
```

```
Descriptor Name .. AA
File Number ..... 1
File Name ..... EMPLOYEES
Password .....
In Parallel ..... NO
```

```
Enter 'RELEASE' to confirm .. _____
```

```
PF1----- PF2----- PF3----- PF4----- PF6----- PF7----- PF8----- PF12-----
Help                Exit                Confirm                Menu
```

Updated Modify File Parameters Screen

In the **File Maintenance** chapter, “Define/Modify FDT” section, “Change File Parameters” subsection, the screen capture has been updated to accommodate the display of the new ‘In Parallel’ option. The updated screen capture is below:

```

14:59:39          ***** A D A B A S  BASIC  SERVICES  *****          2003-03-03
DBID 1955          -  Modify File Parameters  -                          PFLM002

File No. ... 1
File Name .. EMPLOYEES
-----
                                Max. Allocation
ASSO PFAC ..... 10                UI Blks per extent .. 0
DATA PFAC ..... 10                NI Blks per extent .. 0
Max. RECL ..... 4816              DS Blks per extent .. 0

New File Name .. EMPLOYEES_____  ISN Reuse ..... OFF
New File No. ... 1                  with RESET ..... _
                                      in Parallel .... NO
User ISN ..... OFF                 DS Reuse ..... ON_
                                      with RESET ..... ON_
File Password ..                    in Parallel
                                      Mixed DS Device..... OFF
AlphaNum Encoding .. 0              Program Refresh ..... OFF
WideChar Encoding .. 0

PF1----- PF2----- PF3----- PF4----- PF6----- PF7----- PF8----- PF12-----
Help                Exit                Menu

```

Updated Display Cluster Members Screen

From the Session Monitoring menu, a new function **Display cluster members** (option A) produces the following screen:

```

16:21:45          ***** A D A B A S  BASIC SERVICES  *****      2002-07-19
DBID 105          - Display Cluster Members -                      PACA002

```

```
Total number of nuclei in the cluster: 4
```

| I Sel | I Nuc ID | I System ID | I Jobname | I Status | I Available Services | I |
|-------|----------|-------------|------------|------------|----------------------|---|
| I _ | I 1 | I DAEMVS | I ADANUC01 | I Active | I All | I |
| I _ | I 2 | I DAEMVS | I ADANUC02 | I Inactive | I List,Cache | I |
| I _ | I 3 | I DDZMVS | I ADANUC03 | I Active | I All | I |
| I _ | I 4 | I DDZMVS | I ADANUC04 | I Active | I All | I |
| I | I | I | I | I | I | I |
| I | I | I | I | I | I | I |
| I | I | I | I | I | I | I |
| I | I | I | I | I | I | I |
| I | I | I | I | I | I | I |
| I | I | I | I | I | I | I |
| I | I | I | I | I | I | I |
| I | I | I | I | I | I | I |
| I | I | I | I | I | I | I |

```

PF1----- PF2----- PF3----- PF4----- PF6----- PF7----- PF8----- PF12-----
Help                Exit       Refresh                Menu

```

The screen includes a list of nuclei participating in the cluster and information about the current status of each nucleus.

To select a nucleus for additional processing

- Type “S” in the Sel column opposite that nucleus.

To display additional information about a nucleus

- Type “D” in the Sel column opposite that nucleus.

For an Adabas cluster nucleus that has a nonzero nucleus ID, its entry in the parallel participant table (PPT) is displayed in a screen similar to the following:

```

16:21:45          ***** A D A B A S  BASIC SERVICES *****          2002-07-19
DBID 105          - Display PPT Entry -                               PACA002

Nuc ID. . . . . 3 Active Nucleus

Name              Status              DataSet Name
-----
WORK1              SAG.ADABAS.DB105.WORKR1
PLOGR1  Ready to be copied/merged  SAG.ADABAS.DB105.PLOGR1
PLOGR2  Being written by nucleus   SAG.ADABAS.DB105.PLOGR2

PF1----- PF2----- PF3----- PF4----- PF6----- PF7----- PF8----- PF12-----
Help              Exit       Refresh              Menu

```

Updated Nucleus File Status Screen

From the Resource Utilization menu, the **Nucleus file status** (option N) has been added and is the equivalent of the DNFV operator command.

```

16:03:17          ***** A D A B A S  BASIC  SERVICES *****          2002-05-29
DBID  1955          - Nucleus File Status -          PACUN02
NucID 1021

          Locking
          File   NucID  Access count   Update count   State
-----
          24          0             0             0 Access
          25          0             0             0 Access, Update

Last page
PF1----- PF2----- PF3----- PF4----- PF7----- PF8----- PF9----- PF12-----
Help      Repos    Exit    Refresh  -         +         Menu

```

In an Adabas cluster environment, the file may be locked for exclusive use by another cluster nucleus. If this is the case and the file is in the nucleus file status table, the Locking NucID column for the file shows the ID of the nucleus that has exclusive control.

The Access count / Update count fields display the number of access or update users, respectively, that refer to the specified file in their user queue elements (UQEs). These users either have specified the file in an OP command with R-option or are using the file in an as yet incomplete transaction.

A State field indicates when the file is used for access only or for access and update. The State field indicates to what extent a nucleus can use a file on its own. If the requested use exceeds the given state, the nucleus must first communicate with the other nuclei in the cluster in order to upgrade the state.

Updated Nucleus Status Flags

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 datasets; remote and local call distribution; and other current session status information.

```

18:50:16          ***** A D A B A S  BASIC  SERVICES  *****          2002-05-30
DBID 1955          - System Status -          PACUS02
NucID: 1022

          Physical
          Reads          Writes          Call Distribution
          -----
ASSO          370          67  Remote Logical .....          0
DATA          3          18  Remote Physical .....          0
WORK          2          104 Local Logical .....          860
PLOG          67          Local Physical .....          0

Logical Reads .....          349  Logical Reads (binary) ..... 0000015D
Buffer Efficiency ....          0.9  No. of HQEs active .....          0
          No. of UQEs in User Queue ..          2
Format Translations ..          51  No. of CQEs waiting in CQ ..          0
Format Overwrites ....          0
          Total intern. Autorestarts .          0
Throw Backs for ISN ..          0  No. of PLOG switches .....          0
Throw Backs for Space.          0  No. of Bufferflushes .....          18

                                page 1 of 2

PF1----- PF2----- PF3----- PF4----- PF6----- PF7----- PF8----- PF12-----
Help          Exit          Refresh          +          Menu

```

Press PF8 to display an additional screen that indicates if one or more of the following are in progress:

- Online database save running;
- ADAEND in progress;
- Online file save running;
- READONLY/UTIONLY transition;
- READONLY status;
- Update processing suspended;

- ET-sync in progress;
- UTIONLY status; and
- Exclusive-DB-control utility running.

Otherwise, “Adabas operation normal” is displayed.

```
16:47:41          ***** A D A B A S  BASIC  SERVICES *****          2002-05-29
DBID 1955          - System Status -          PACUS02
NucID: 1021
```

Nucleus Status Flags

```
-----
Adabas operation normal
```

page 2 of 2

```
PF1----- PF2----- PF3----- PF4----- PF6----- PF7----- PF8----- PF12-----
Help          Exit      Refresh          +          Menu
```

Updated Cluster Usage Screen

Note:

This option is only active in an Adabas nucleus cluster environment.

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.

The equivalent direct command is

DISPLAY CLUSTERSTATUS

```

16:10:31          ***** A D A B A S  BASIC  SERVICES *****          2002-05-29
                                     - Cluster Usage -                               PACUX02

          Code      Service
          ----      -
          C          Cache statistics
          F          File statistics
          L          Lock statistics
          ?          Help
          .          Exit
          ----      -

Code ..... _
File Number .. 0
Database ID .. 1955   (WIS1955)                NucID .. 1021

Command ==>
PF1----- PF2----- PF3----- PF4----- PF6----- PF10----- PF11----- PF12-----
Help       Exit       Fuse       Flist       Menu

```

Updated User Table Maintenance Screen

Note:

This option is available in Adabas cluster environments only.

A new function has been added to the Session Opercoms menu to support the CLUFREEUSER command. When option V (maintain user table) is selected, the following screen is displayed:

```

16:59:29          ***** A D A B A S  BASIC  SERVICES *****          2002-05-29
                    - User Table Maintenance -                          PACIV02

                Code      Service
                ----      -
                C          Begin CLUFREEUSER process
                ?          Help
                .          Exit
                ----      -

Code ..... _
TNA ..... 0 _____
UID ..... _____
Force ..... _
Global ..... _

Database ID .. 1955 (WIS1955)          NucID .. 1022

Command ==>
PF1----- PF2----- PF3----- PF4----- PF6----- PF7----- PF8----- PF12-----
Help          Exit          Menu

```

The CLUFREEUSER command is only valid in cluster environments. It can be issued against the local nucleus only or, with the Global option, against all active and inactive nuclei in the cluster.

The command is used to delete leftover user table elements (UTEs) in common storage that are no longer associated with user queue elements (UQEs) in a nucleus where

- TNA** is a decimal number specifying the timeout value in seconds.
UTEs that are not used during the time specified may be deleted if other conditions are fulfilled.
If TNA is not specified, UTEs may be deleted without regard to their recent use.
- UID** is a character string or hexadecimal byte string as follows:
- | | |
|---------------------|--|
| ccccccc | —where the argument is 1-8 letters, digits, or embedded ‘-’ signs without surrounding apostrophes. |
| ‘ccccccc’ | —where the argument is 1-8 characters with surrounding apostrophes. |
| X‘xxxxxxxxxxxxxxxx’ | —where the argument is an even number of 2-16 hexadecimal digits enclosed by X‘ ’. |
- A character string must be enclosed in apostrophes if it contains characters other than letter, digits, or embedded ‘-’ signs. If a specified character string is less than 8 characters long, it is implicitly padded with blanks. If a specified hexadecimal string is shorter than 16 hexadecimal digits, it is implicitly padded with binary zeros.
- If the last 8 bytes of a user’s 28-byte communication ID match a specific user ID or user ID prefix, that user’s UTE may be deleted if other conditions are fulfilled.
- If UID not specified, UTEs may be deleted regardless of their user IDs.

- FORCE** Delete leftover UTEs even if the users are due a response code 9, subcode 20. If FORCE is not specified, such UTEs are not deleted. Before using the FORCE parameter, ensure that the users owning the UTEs to be deleted will not expect any of their transactions to remain open. Specify FORCE on this screen by marking the Force field with any character.
- GLOBAL** Delete leftover UTEs throughout the Adabas cluster if they are no longer associated with UQEs and are eligible according to the other specified parameters. Additionally and subject to the other rules, delete leftover UTEs if their assigned nuclei have terminated since their last use. If GLOBAL is not specified, only UTEs assigned to the local nucleus and used since the nucleus start are eligible for deletion. Specify GLOBAL on this screen by marking the Global field with any character.

Updates to the Adabas Installation Documentation (OS/390 and z/OS)

This section describes changes to the *Adabas Installation (OS/390 and z/OS)*.

ADAGSET Option “ESI”

In the **Installing Adabas with TP Monitors** chapter, “Modifying Source Member Defaults (ADAGSET Macro)” section, the option “ESI” has been changed to “SAF”, as follows:

SAF={ NO | YES }

Specify whether the Adabas SAF Security (ADASAF) is used. If you are using ADASAF, you must set SAF=YES.

ADASAF requires the Adabas task-related user exit (TRUE) when running under CICS/ESA 4.1 or above. When SAF=YES and TRUE=YES, the task-related user exit passes the user’s external security ID to Adabas.

If TRUE=YES is not specified in this case, the ADAGSET macro terminates the LNKOLSC, LNKTRUE, or LNKENAB assembly process with an MNOTE and a return code of 16.

TRUE is not required when running ADASAF under CICS/ESA 3.3 or below. The combination SAF=YES and TRUE=NO is valid in such cases.

Adabas Utilities

This section describes changes to the *Adabas Utilities*.

ADADBS Functional Overview

In the **ADADBS: DATABASE SERVICES** chapter, “Functional Overview”, reference to protection logs and command logs has been removed from the ADD function and the DELETE function. (Adding and deleting protection logs and command logs has not yet been implemented in version 7.4.)

The Functional Overview table now contains the following information for ADD and DELETE:

| Function | Action | Page |
|----------|---|------|
| ADD | add a dataset to the Associator or Data Storage | 139 |
| DELETE | delete a file | 151 |

ADAFRM Formatting Database Components

In the **ADAFRM: Format** chapter, “Formatting Database Components”, under “Formatting Modes”, the second paragraph under item number 2 is replaced with the following paragraph:

“This formatting mode is used in combination with the ADADBS INCREASE function for ASSO and DATA. If a greater WORK is needed, then ADADEF NEWWORK should be used.”

ADARES Updates

In the ADARES PLCOPY and CLCOPY utilities sections, all references to “MERGINTI” are now “MERGIN1” and all references to “MERGINTO” are now “MERGIN2”.

ADARES Backout Functions Information

In the **ADARES: Restart** chapter, “BACKOUT Functions”, the following information has been added to the description of BACKOUT DPLOG or MPLOG:

“ADARES BACKOUT DPLOG/MPLOG function is not valid for a cluster database. This is disallowed because a merged PLOG is required in order to perform the BACKOUT.”

New ADARES PLCOPY Parameter

In the **ADARES: Restart** chapter, “ADARES PLCOPY ”, the following new parameter has been added:

SBLKNUM - an optional parameter used only in conjunction with the NOPPT parameter. SBLKNUM allows the user to specify the starting block number for the sequential merge output. This parameter can only be specified in conjunction with NOPPT and only for the PLCOPY function. If this parameter is omitted, there is an attempt to read the PPT and obtain the block number from there. If this read fails, the output will start with block one. To determine the value for this parameter, the user must look at the output from the previous PLCOPY and use the next block number in sequence.

Appendix C: Supplied UES Encodings

The following additions were made to **Appendix C: Supplied UES Encodings**, in the “Coexistent Encodings” section, under “Single-Byte Character Sets”:

| Key | | CS | CP | F/M/S Size | ESID | Fill | Sub | Description |
|-----|-----|----|----|---------------|------|------|-----|---|
| Dec | Hex | | | | | | | |
| 720 | | - | - | F 223 | 2100 | 20 | 1F | Microsoft Windows OEM Codepage 720 (Arabic) |
| 737 | | - | - | F 223 | 2100 | 20 | 1F | Microsoft Windows OEM Codepage 737 (Greek) |
| 775 | | - | - | F 223 | 2100 | 20 | 1F | Microsoft Windows OEM Codepage 775 (Baltic) |
| 852 | | - | - | F 223 | 2100 | 20 | 1F | Microsoft Windows OEM Codepage 852 (cp852_DOSLating2) |
| 857 | | - | - | F 223 | 2100 | 20 | 1F | Microsoft Windows OEM Codepage 857 (Turkish) |
| 862 | | - | - | F 223 | 2100 | 20 | 1F | Microsoft Windows OEM Codepage 862 (Hebrew) |
| 866 | | - | - | F 223 | 2100 | 20 | 1F | Microsoft Windows OEM Codepage 866 (Cyrillic) |

| Dec | Hex | | | Size | | | | |
|------------|------------|------|------|-------------|------|----|----|--|
| 874 | | - | - | F 223 | 2100 | 20 | 1F | Microsoft Windows OEM Codepage 874 (Thai) |
| 1047 | | 697 | - | F 190 | 1100 | 40 | 3F | IBM Latin-1 Open Systems EBCDIC |
| 2084 | | 948 | 878 | F 222 | 4105 | 20 | 1F | KO18-R |
| 2087 | | 813 | 775 | S 222 | 2100 | 20 | 1F | MS-DOS Baltic PC-Data |
| 2258 | | 1408 | 1258 | S212 | 4105 | 20 | 1F | Microsoft Windows-1258 (Vietnam) |
| 3585 | E01 | 646 | - | S 95 | 1100 | 40 | 3F | BS2000 DF.03 International CCSN:EDF03IRV |
| 3586 | E02 | 265 | - | S 95 | 1100 | 40 | 3F | BS2000 DF.03 CCSN:EDF03 Austria, Germany (EBCDIC) - with Dollar |
| 3587 | E03 | - | - | S 95 | 1100 | 40 | 3F | BS2000 DF.03 CCSN:EDF03 Austria, Germany (EBCDIC) - with International Currency Sign |
| 3588 | E04 | 697 | - | F190 | 1100 | 40 | 3F | BS2000 DF.04-DRV LATIN 1, CCSN:EDF04DR Austria, Germany (EBCDIC) |
| 3589 | E05 | 697 | - | F190 | 1100 | 40 | 3F | BS2000 DF.04-1 LATIN 1, CCSN:EDF041IRV International (EBCDIC) |

The following additions were made to **Appendix C: Supplied UCS Encodings**, in the “Coexistent Encodings” section, under “Double- and Multiple-Byte Character Sets”:

| Key | | CS | CP | F/M/S Size | ESID | Fill | Sub | Description |
|------------|------------|-------------|-----------|-------------------|-------------|-------------|------------|---|
| Dec | Hex | | | | | | | |
| 4037 | | 1175 935 | - | S 20360 | 1301 | 40 4040 | 3F FEFE | Same as EDD 937 (Chinese EBCDIC) with code point X'5F' mapped to Unicode U+005E and X'B0' to U+00AC |

Updates to the Triggers and Stored Procedures

This section describes changes to the *Triggers and Stored Procedures* documentation.

Update to the Installation and Operations Chapter

The following subsection was added to the “NATPARAM Considerations” section, under “Special Requirements”:

PROFILE

The PROFILE parameter must not be specified. This parameter causes Adabas calls to be issued before all of the Natural control blocks are initialized. This will cause unpredictable results during the initialization of stored procedures and triggers.

Updates to the Adabas Operations

This section describes changes to the *Adabas Operations* documentation.

Updates to the ADARUN Control Statement Chapter

Addition to the ADARUN Parameters Section

The following note was added under “FORCE : Allow Database ID Table Entry”:

Note:

In an Adabas Cluster Services or Adabas Parallel Services environment, the FORCE parameter applies to the NUCID, rather than the database ID, because a cluster nucleus builds an ID table entry for the NUCID.

Change to the ADARUN Parameters Section

The following information was changed under “NU : Number of User Queue Elements:

Currently, $288 * (NU+4)$ bytes are allocated to the user queue.

is changed to

Currently, $272 * (NU+4)$ bytes are allocated to the user queue.



APPENDIX A - THE AFPLOOK UTILITY

This appendix describes the AFPLOOK utility, which is distributed on the Adabas release tape as part of ADAAFP. It is therefore available to all Adabas customers. In this form, AFPLOOK samples a full database session and prints the results at ADAEND. The sampler can be controlled and viewed online using SYSAFP.

For Fastpath users, the sampler is still relevant. The description of the sampler and its reports provided in this section should be used for reference.

Overview

For Fastpath to operate optimally, the DBA must provide file optimization entries. These entries inform Fastpath of the files, commands, and descriptors that should be optimized. The accuracy of these values depends on an in depth knowledge of the database(s) and its operating environment.

The Adabas Fastpath command analysis sampler (AFPLOOK) assists the DBA in determining where the best results may be expected from Fastpath by reporting on the command constructs that qualify for Fastpath. With this information, the DBA can create the file optimization parameters.

AFPLOOK is set up with certain operational defaults that control the amount of memory used during command analysis by restricting

- the maximum number of files sampled; and
- the number of concurrent users.

If any parameter is exceeded, AFPLOOK tries to ignore the excess while still reporting maximum information. In this way, AFPLOOK audits a general sampling of the database command workload to determine Fastpath optimization parameters. The operational defaults can be modified for site requirements.



How to Run AFPLOOK

The AFPLOOK utility is invoked with:

ADARUN FASTPATH=YES

AFPLOOK Parameters

This section describes the AFPLOOK parameters, which are used to

- define the boundaries of the sample; and
- limit the amount of memory required.

Maximum Files

The maximum number of files to be sampled.

Once the maximum number of files is put in the analysis table, no additional files are sampled; however, additional files show in a command count so that the DBA can judge whether the parameter should be increased for subsequent executions.

Default: 64

Command/Descriptors per File

The maximum number of command/descriptor entries per file.

In conjunction with the Maximum Files parameter, this parameter restricts the amount of memory used. If the maximum entries is reached for a file, the last entry is converted into a general accumulator. Note that only one entry is required for the Adabas command types L1, L2, S8, and S9.

Default: 32



Maximum Concurrent Users

The maximum size for the table of concurrent users.

If all the user areas are being used at one time, a new request is satisfied by releasing the 'oldest usage' user area. The number of times this reusage occurs is noted and printed in the summary. When a sample contains a high percentage of reusage, the parameter should be adjusted.

Default: 100

Maximum CIDs per User

The maximum concurrent Adabas command IDs (CIDs) sampled for each user processed.

In conjunction with the Maximum Concurrent Users parameter, this parameter restricts the amount of memory used. Commands for command IDs that exceed this maximum are ignored and reported as rejected. When a sample contains a high percentage of rejections, the parameter should be adjusted.

Default: 10

Maximum Commands Processed

The maximum number of commands to be sampled.

Default: No limit

Job Name

Used to restrict sampling to a particular job name. Note that one or more asterisks (*) can be used in the job name as a wild card character so that the sample can select all jobs that match the name ignoring the character positions occupied by an asterisk (*).

Default: None



Selected Files

Used to restrict sampling to specific files.

This option may be useful where the maximum files overflowed, or file activity is known and detailed analysis is required.

Default: All files

Number of Lines per Page

The number of lines per page to be used for the report.

Default: 55

Customer Heading

Used to add a 30-character subheading to the report.

This subheading enables the user to label reports from, for example, different databases or divisions within a company.

Default: None

On/Off Switch

Switches sampling on or off.

A site that uses the database continually benefits from this feature. Contact Software AG for more detailed information.

Default: On



Overriding Parameter Defaults

The default AFPLOOK parameters may be temporarily overridden at the start of an AFPLOOK session by using SYSAFP. Such overrides will only take effect for the duration of that particular session. To avoid having to repeat this change for each new AFPLOOK session, an Adabas Fastpath ZAP is available to permanently modify the default parameter values.

For licensed Fastpath users, this ZAP is available on the Fastpath release tape; for non-licensed Fastpath users, this ZAP may be requested from your local Software AG representative.

AFPLOOK Sample File Report

```

-----
FNR CC DESC  DIRECT ACC          RC  SEQUENTIAL  SEQUENCES
-----
 20 L1  --          1
    L2  --                4          4
    L3  CC          1
    L9  AA          1
    L9  BB          2
    L9  CC          2          1          1
    S1  AA          3          3          1
    TOTALS          7                8          21 (18%)
                                EXCLUDED COMMANDS:      2
                                ALREADY PREFETCHED:      3
(UPDATES 2,INSERTS 1,DELETES 1) (MAX.RBL DA 0,SEQ 32)
-----

```

| Column | Explanation |
|------------|---|
| FNR | Adabas file number. |
| CC | Adabas command code. |
| DESC | 'Primary' descriptor for the commands. |
| DIRECT ACC | Maximum direct access commands that can be optimized. |
| RC | Maximum RC commands that can be optimized. |



| Column | Explanation |
|----------------------|---|
| SEQUENTIAL SEQUENCES | Maximum sequence commands that can be optimized. Number of sequences that caused the number of sequential commands. The sequence factor for optimization may be calculated from these two numbers. |

The rightmost number shows the total sampled commands for the file together with the percentage relative to all the sampled file commands. On a large report, a quick glance down these numbers shows the DBA which files are most worthy of detailed attention.

Commands that have been ignored for the file are also listed along with the reason for exclusion.

The final line shows the update commands as well as the maximum record buffer lengths found for direct access or sequential commands that can be optimized.

AFPLOOK Potential Optimization Summary

| ----- POTENTIAL OPTIMIZATION SUMMARY ----- | | | |
|--|------------|----------------------|----------|
| SAMPLED COMMANDS | | MAXIMUM OPTIMIZATION | |
| SAMPLED FILE COMMANDS | 116 (77%) | <----- SEQUENTIAL: | 55 (47%) |
| | | DIRECT ACCESS: | 32 (27%) |
| | | RCS: | 4 (3%) |
| EXCLUDED COMMANDS | 33 (22%) | | |
| TOTALS | 149 (100%) | | 91 (61%) |
| ----- | | | |

This section of the report summarizes the total commands sampled for all files and expresses this as a percentage of all commands seen. Excluded commands are similarly reported.

The maximum optimization numbers show an over-estimation of potential optimization. The sequential commands, direct access, and RC totals are expressed as a percentage of the total sampled file commands, while the total is expressed as a percentage of all commands seen.

These numbers show the **potential** for Fastpath: the actual optimization likely depends on various factors unique to each customer. Contact Software AG for assistance when interpreting samples.



AFPLOOK Sample Command Analysis

```

-----
                          COMMAND ANALYSIS
REJECTED COMMANDS
  MAX. USERS EXCEEDED:           0
  MAX. CIDS EXCEEDED:            0
  MAX. FILES EXCEEDED:          0      0 ( 0%)
EXCLUDED COMMANDS
  BAD COMMANDS:                  4
  NON-FILE COMMANDS:             7
  NON-FILE RCS:                  2
  EXCLUDED FILE COMMANDS:       8
  UPDATE COMMANDS:              4
  ALREADY PREFETCHED:           8      33 ( 22%)
SAMPLED FILE COMMANDS           116 ( 77%)
ALL COMMANDS SEEN               149 (100%)
-----

```

The numbers shown in this section of the report are provided to

- illustrate the type of commands processed by the customer; and
- put the previous section into perspective.

Rejected commands are categorized by users, CIDs, and files exceeded. If the total percentage is high, numbers reported elsewhere may not give an accurate assessment of the situation.

Excluded commands are split into the following categories:

| Category | Explanation |
|------------------------|---|
| Bad commands | Unexpected Adabas response codes. |
| Non-file commands | Commands that cannot be attributed to a file; for example, OP, CL, ET, C1, RE. Plus file commands HI, LF, RI. |
| Non-file RCS | RC all commands plus any RC for which the CID is not stored by AFPLOOK. |
| Excluded file commands | L4, L5, L6, S4, S5. |
| Update commands | A1, A4, E1, E4, N1 |
| Already prefetched | Any command that could qualify for sequential optimization that has prefetch or multifetch already set. |



AFPLOOK Sample Parameters Report

```
-----  
PARAMETERS USED  
MAX. FILES:      64  FILES NEEDED:      5  
  ..MAX. DE:     32  OVERFLOWS:      0  
MAX. USERS:     100  HIGH USERS:      15  
  ..MAX.CID:     10  HIGH CIDS:       4  
* REUSED USER AREA OCCURRENCES:      0  
MAX.RECORDS: NO LIMIT  
-----
```

This section of the report

- shows the important parameters used to produce the report; and
- gives an indication of the parameters needed.



APPENDIX B - THE AVILOOK UTILITY

The Adabas Vista analysis tool AVILOOK provides statistics on the type and quantity of key-based accesses against an Adabas file. You can use these statistics to identify the Adabas files that may benefit from the Adabas Vista partitioning option. AOS must be installed (either the demo version distributed with Adabas beta version or the full version selectable unit at the same level).

The Adabas database where you wish to run AVILOOK and all other system components (link module, router, etc.) must be the current version. The database must be running with the ADARUN VISTA=YES parameter. At normal database termination, the statistics gathered during AVILOOK processing are written to the database job log.

To access the AVILOOK tool

- Select service 4 from the main menu.

```
13:34:34          ***** A D A B A S  VISTA SERVICES 7.4.2  *****          2003-05-05
                               - Avilook -                               V14000M1

                               Code      Service
                               -----
                               1         File Maintenance
                               .         Exit
                               -----
                               Code..:  _

                               Database ID..: _____

                               Nucleus ID..:  _____

                               System Coordinator Node..: _____

Command ==>
Enter-PF1---PF2---PF3---PF4---PF5---PF6---PF7---PF8---PF9---PF10---PF11---PF12---
          Help          Exit                               Menu
```

File Maintenance

List AVILOOK Files

The File Maintenance screen lists the files that are already defined to AVILOOK for a specified database number. The database name is also displayed.

To display the File Maintenance screen from the AVILOOK menu

1. Specify the Adabas database number for the database where you wish to run AVILOOK.
This database and all other other system components (link module, router, etc.) must be Adabas version 7.2.2 or above.

The database must be running with the ADARUN VISTA=YES parameter.

2. (Optional) For a cluster database, specify a nucleus ID.

If the database number you specify is a cluster database, you have the option to specify the ID (that is, the NUCID) of the cluster nucleus you wish to monitor.

```

+-----+
| 13:31:07      Select Cluster Members  2003-05-05 |
|                                                    |
| DBID:   231   |
|                                                    |
|  C   Nuc   C   Nuc   C   Nuc   C   Nuc   |
|  -   1   -   2           |
|                                                    |
| Mark to select |
| Command ==>   |
|                PF3 Exit   PF4 Refr |
+-----+
  
```

You may select the appropriate nucleus from this list.

If you choose not to specify a nucleus ID or you specify a value of 0, you are required to specify the node ID of the local System Coordinator. A window is displayed listing the nuclei that are currently active in the cluster.

If the job within which you are using the Adabas Vista online services is defined to a System Coordinator group, the node ID of the local System Coordinator is automatically set up.

3. Specify service 1.

```

11:14:06      ***** A D A B A S  VISTA SERVICES 7.4.2 *****      2003-05-05
                - AVILOOK  File Maintenance -                          V14100M1

DBID: 231      (TEST-V7-DB)

C  File      Command Limit      Commands      Started      Status
-   12          0          0          Started      Paused
-   2          0          5768      2003-05-05 09:09:20      Active

Mark with (A)ctivate, (P)ause, (R)eset, (S)tatistics, (X)Delete

Command ==>
Enter-PF1---PF2---PF3---PF4---PF5---PF6---PF7---PF8---PF9---PF10--PF11--PF12---
      Help          Exit Refr                          Add          Menu

```

Press PF4 to refresh the command count.



To select a file entry

Enter one of the following options in column C alongside the file entry to be selected:

- a activate; start gathering statistics
- p pause; stop gathering statistics
- r reset statistics to zero
- s display the current statistics for the file
- x delete the file from the file list

See the section **Maintain AVILOOK Files** on page 88 for more information.

Add AVILOOK Files

▶ To add a new AVILOOK file

1. Press PF10 from the AVILOOK File Maintenance screen.

```

+-----+
| 08:11:07          AVILOOK Add File                      2003-05-05 |
|                                                           V14100M2 |
|           File _____ Status (A/P) P      (Active/Paused) |
|                               Command Limit 0 _____ (0=No Limit) |
|                                                           |
|                               PF3 Exit    PF5 Update |
|                                                           |
+-----+

```

2. Define the file number.
3. Indicate whether you want AVILOOK to start gathering statistics immediately (Active status) or you want to define the file now and activate it at a later time (Pause status).
4. (Optional) Predefine the maximum number of commands the active file can process before it automatically reverts to pause status.
5. Press PF5 to add the file.

Maintain AVILOOK Files

▶ To maintain AVILOOK files and file statistics from the AVILOOK file maintenance screen

- Enter the appropriate option (see page 87) in column C next to the file entry.

For example, the option “s” displays the current statistics for a file:

```

16:56:07      ***** A D A B A S   VISTA SERVICES 7.4.2 *****      2003-05-05
              - Avilook File Statistics -                          V14120M1

DBID: 231    (TEST-V7-DB)
File: 2
                                         Started: 2003-05-05 09:09:20
                                         Paused:

      CC   Desc   Command Count      CC   Desc   Command Count
      L3   AA     2836              L3   AA     2836
      S1   AB     1324              S1   AB     1324
      L3   BC      24                L3   BC      24
      L9   S1      26                L9   S1      26

Other Commands not listed above: 1558

Command ==>
Enter-PF1---PF2---PF3---PF4---PF5---PF6---PF7---PF8---PF9---PF10--PF11--PF12---
      Help           Exit Refr                               Menu

```

This screen shows statistics on the command type and quantity of key-based accesses against the file (for example, S1, L3, and L9 commands). The statistics are displayed in descending order by command code (CC) and Adabas two-character field name (Desc).

In the example, an L3 command using the Adabas field AA to sequence by dominates file activity with 4298 accesses. Such a file, where the predominant access is by a single key, may benefit from being partitioned using the Adabas field AA as the Adabas Vista partitioning field.

INDEX

A

- active nuclei, 29
- Adabas, support for earlier releases, 5
- Adabas Caching Facility, 11
- Adabas Delta Save Facility, AOS requirement, 12
- Adabas External Security Interface, 14
- Adabas Fastpath, 12
- Adabas Online System, 13, 43
 - 'In Parallel' option, 44
 - installation requirements, 43
 - ZAP prefixes, 44
- Adabas Parallel Services
 - multiple thread processing, 13
 - parallel processing, 13
- Adabas Review, overview, 14
- Adabas SAF Security, 14
- Adabas SVC, 26, 34
- Adabas Vista, 16
- ADACMP, version compatibility, 22
- ADACOX, 40
- ADAESI, 14
- Adalink, 33
- ADALNKX, 33
- ADALOD, LOAD VERSION parameter, 22
- ADAORD, version compatibility, 22
- ADARAI, compatibility, 21
- ADARUN options, 28
- ADASAF, 14
- ADASAV
 - version compatibility, 21
 - Work dataset reset, 38
- ADASIP, releasing common storage, 37
- ADAULD, version compatibility, 22
- AFPLOOK, 77–84
 - how to run, 78
 - overview, 77

- parameters, 78
 - overriding defaults, 81
- potential optimization summary, 82
- sample command analysis, 83
- sample file report, 81
- sample parameters report, 84
- ZAP for, 81
- Analysis tool, AVILOOK, 85
- AOS, 43
- AQA support, 33
- ARM, 29
 - explanation, 35
- ARMNAME, ADARUN parameter, 35
- asynchronous CLOG I/O, 37
- ASYTVS, ADARUN parameter, 36
- ATM, definition, 16
- Automatic Restart Management, 29, 35
- Automatic Restart Management (ARM), 35
- AVILOOK, 85
 - add files, 88
 - list files, 86
 - maintain files, 88

B

- backup and restore performance, 26
- Buffer flush, 35

C

- CE, 33
- CIDs per user, maximum, 79
- Cluster
 - display usage statistics, using Basic Services, 45, 67
 - maintain user table, 45, 68
- cluster mode vs. noncluster mode, 32
- Command analysis, sampler, 77–84

Adabas Version 7.4.2 Release Notes

Command/descriptors per file, 78
Commands
 excluded, 82
 maximum processed, 79
Communication modules, version compatibility,
 16

D

data integrity enhancements, 27
Database, single physical, defined, 13
DBA function support, 33
DD/JCLIN file, 20
Dynamic allocation, 31

E

end of support, 6
Entire Net-Work, component, 12
Excluded commands, 82

F

fast copy products (CE), 33
features summary, 25
file capacity (BS2000), 26
Files
 maximum, 78
 selected, 80
FMXIO, ADARUN parameter, 36

H

Heading, report, 80

J

Job, name, 79

L

Lines per page, number of, 80
link routines, 28, 39
logging performance, 27

M

messages (mixed case), 35
mixed case messages, 35

N

noncluster mode vs. cluster mode, 32
NUCID, 29
Nucleus, cluster, definition of, 13
Nucleus file, display status, using Basic Ser-
 vices, 45, 64
Nucleus ID, 29

O

On/off switch, 80
Operating systems, supported, 11

P

Parallel Participant Table, 29
performance
 logging, 27
 other enhancements, 27
PLOGRQ=FORCE, 29
PPT, 29
PPT initialization processing, 30
PREFICMD, ADARUN parameter, 37
PREFIFIL, ADARUN parameter, 37
previous version support, 6
PRILOG, 40

R

RC command, 82

Response codes, record-related, 38
restore and backup performance, 26

S

SAF option, ADAGSET macro, 71
Sequential PLOGs, version compatibility, 23
Session, display status, using Basic Services, 45, 65
Sixty-four (64) bit storage support
 real, 8
 virtual, 9
superdescriptor search values, 41
SVC, 26, 34
 and link routines, 37
SYSPARM, 34
System, display status, using Basic Services, 45, 65

T

TRUE option, ADAGSET macro, 71

U

User table, maintain, 45, 68
Users, maximum concurrent, 79

V

value buffer, 40

W

Work dataset, reset by ADASAV, 38

Z

ZAP, for AFPLOOK, 81
ZAP prefixes, Adabas Online System, 44

Adabas Version 7.4.2 Release Notes

Adabas Version 7.4.2 Release Notes