

Adabas Version 7.4.1

Release Notes

Manual Order Number: ADA741-008ZOS

This document applies to Adabas Version 7.4.1 and to all subsequent releases.

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

© December 2002, Software AG

All rights reserved

Printed in the Federal Republic of Germany

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

TABLE OF CONTENTS

PREFACE	1
About this Release	1
Enhanced SQL Support	1
Single Code Base	1
Enhanced Support for Multiple-Engine Processors	1
About this Document	2
Documentation	2
Printed Manuals	2
Online Documentation	3
Coordinated Releases	4
1. ENVIRONMENT	5
Operating Systems Supported	5
Support for 64-Bit Storage	5
Real Storage	5
Virtual Storage	6
Optional ZAPs Dataset Added	6
Software AG Product Level Compatibility	6
Adabas and Natural	6
Entire Net-Work	6
Entire Net-Work X-Hub	7
Adabas SQL Server	7
Adabas Add-on Products Overview	7
Adabas Caching Facility	8
Adabas Cluster Services	8
Adabas Delta Save Facility	9
Adabas Fastpath	9
Adabas Manager	9
Adabas Online System	10
Adabas Parallel Services	10
Adabas Review	11

Adabas Version 7.4.1 Release Notes

Adabas SAF Security	11
Adabas Statistics Facility	12
Adabas Transaction Manager	13
Adabas Vista	13
Adabas Component Compatibility	14
Communication Modules	14
Enqueue Conflicts	15
SVC Consolidation	16
Migration Issues	16
TP Link Routines	17
User Exit Interfaces	17
Conversion Requirements	17
New Skeleton Job Control File	18
Adabas Recovery Aid Compatibility	18
Save Datasets	18
Unload Datasets	19
ADAORD DD/FILEA Datasets	19
Migrating Products from Version 7.1 to Version 7.4	20
Sequential Protection Logs	20

2. RELEASE NOTES FOR ADABAS VERSION 7.4.1 **21**

Summary of Features Added for 7.4	21
Enhanced Interfaces between Products	21
Enhanced Support for Operating System Features	21
More SVC Flexibility	21
Enhanced Backup and Restore Performance	22
Data Integrity Enhancement	22
Enhanced Logging Performance	23
Other Performance Enhancements	23
New and Enhanced ADARUN Options	23
Feature Discussion	24
NUCID	24
New Table to Track the Active Nuclei	24
Support for the Adabas SQL Server (AQA)	26
New Utility Function for Use with Fast Copy Products (CE)	27
Support for Assembler SYSPARM Specification of the Adabas SVC	27

Table of Contents

Messages in Mixed Case (CE)	28
Automatic Restart Management (ARM) Support	28
New Buffer Flush Method	29
New ADARUN Parameters to Include Files in Prefetch/Multifetch	30
Asynchronous CLOG I/O	30
Enhancements	31
CSAs of Replaced SVCs Released	31
Different ISN Returned with Record-related Response Code	31
ADASAV Processing Change	32
Number of Blocks Allocated to Work Data Protection Area Increased	32
Link Routines are UES-enabled by Default	32
Conversion of High Value in Value Buffer	33
PRILOG Now Delivered in Both Source and Object Form	33
Removed from Adabas	34
BUB Support Removed	34
Mirror Dataset Support Removed	34
Adabas ESA Option (AEO) No Longer Available	34
ADAESI No Longer Available	34
Interlink TCP/IP No Longer Supported	34
Changes to Adabas Online System	35
Product Interfaces	35
Installation Requirement	35
ZAP Prefixes for Adabas Online System (AOS)	36
Display Cluster Members	36
Nucleus File Status	39
Nucleus Status Flags	40
Cluster Usage	42
Maintain the User Table	51
APPENDIX A – THE AFPLOOK UTILITY	55
Overview	55
How to Run AFPLOOK	56
AFPLOOK Parameters	56
Maximum Files	56
Command/Descriptors per File	56
Maximum Concurrent Users	57

Adabas Version 7.4.1 Release Notes

Maximum CIDs per User	57
Maximum Commands Processed	57
Job Name	57
Selected Files	58
Number of Lines per Page	58
Customer Heading	58
On/Off Switch	58
Overriding Parameter Defaults	59
Adabas Version 7.1	59
Adabas Version 7.2.2	60
AFPLOOK Sample File Report	61
AFPLOOK Potential Optimization Summary	62
AFPLOOK Sample Command Analysis	63
AFPLOOK Sample Parameters Report	64
APPENDIX B – THE AVILOOK UTILITY	65
File Maintenance	66
List AVILOOK Files	66
Add AVILOOK Files	69
Maintain AVILOOK Files	70
INDEX	71

PREFACE

About this Release

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.

Single Code Base

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 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-nucleus update 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 7.4.1. They contain information about the changes to Adabas since version 7.1.3.

Chapter 1 contains information about the required environments.

Chapter 2 contains the release notes for Adabas version 7.4.1.

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 7.4.1 release.

Printed Manuals

The following printed, bound manuals are being delivered with the version 7.4.1 release:

Order No.	Manual
ADA741-006IBB	Adabas Concepts and Facilities
ADA741-008IBB	Adabas Version 7.4.1 Release Notes (this document)
ADA741-010...	Adabas Installation Manual (platform-specific)
ADA741-030IBB	Adabas DBA Reference Manual
ADA741-032IBB	Adabas Triggers and Stored Procedures Manual
ADA741-050IBB	Adabas Command Reference Manual
ADA741-060IBB	Adabas Messages and Codes
ADA741-080IBB	Adabas Utilities Manual, Volume 1
ADA741-081IBB	Adabas Utilities Manual, Volume 2
ADA741-090IBB	Adabas Security Manual (available only to authorized personnel)
ADA741-110IBB	Adabas Operations Manual

Printed, bound manuals are also being delivered for several add-on products:

Order No.	Manual
ACF741-008ZOS	Adabas Caching Facility Version 7.4.1 Release Notes
ACF741-030ZOS	Adabas Caching Facility Manual
ADE741-030IBB	Adabas Delta Save Facility Manual
AOS741-030IBB	Adabas Online System Manual

Online Documentation

For the first time, the mainframe Adabas documentation is being provided on CD as HTML online files with their PDF equivalents. This documentation is also available from the Software AG web site for product support.

The initial online release provides some, but not all, of the Adabas 7.4.1 documentation. It includes the following sections:

- Concepts and Facilities
- DBA Tasks
- Command Reference
- Messages and Codes
- User Exits
- Command Log Formats
- Supplied UES Encodings

Customer feedback on the usability and quality of this documentation are welcomed. Beginning with Adabas version 8, the entire product documentation set will be provided in the online format.

Coordinated Releases

The following Adabas version 7.4 add-on products and co-products will be released separately:

Code	Product
ALS	Adabas Cluster Services version 7.4
ASM	Adabas Parallel Services version 7.4
AQA	Adabas SQL Server version 5.1
AUT	Adabas UTM Interface version 7.4
REV	Adabas Review version 4.3
AAF	Adabas SAF Security version 7.4
AFP	Adabas Fastpath version 7.4
AMA	Adabas Manager version 7.4
ASF	Adabas Statistics Facility version 7.4
ATM	Adabas Transaction Manager version 7.4
AVA	Jadabas version 7.4
AVI	Adabas Vista version 7.4
COR	System Coordinator for Adabas Options version 7.4

ENVIRONMENT

Operating Systems Supported

Adabas version 7.4.1 is released for the following operating environments:

- OS/390 version 2, releases 6–10
- z/OS version 1, releases 1–2

Other operating environments supported by mainframe Adabas will be released at the version 7.4 SM 2 level.

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 8). Contact your Software AG account representative for more information.

A demo of Adabas Caching Facility is delivered in the ADA741.ALLINPL file.

Optional ZAPs Dataset Added

The distributed source library contains a new member ZAPOPT, which contains ZAPs that may optionally be applied. 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. You will need to update the Entire Net-Work ADALNK to use the version 7.4.1 Adabas router (SVC).

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.

Entire Net-Work X-Hub

If you use Entire Net-Work X-Hub with Adabas version 7.4, you must use Entire Net-Work X-Hub version 1.2.1 or above. You will need to update the Entire Net-Work X-Hub ADALCO to use the version 7.4.1 Adabas router (SVC).

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.

Adabas Add-on Products Overview

Adabas version 7.4.1 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)	version 7.4
Delta Save Facility (ADE)	version 7.4
Fastpath (AFP)	version 7.1, 7.3, or above
Manager (AMA)	version 1.1.1 or above
Online System (AOS)	version 7.4
Parallel Services (ASM)	version 7.4 (initial release)
Review (REV)	version 4.3 Upgrade Adabas Review to version 4.3 first; then upgrade Adabas to version 7.4.
SAF Security (AAF)	version 7.1, 7.3, or above
Statistics Facility (ASF)	version 7.1
Support for Multiprocessing (ADASMP)	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 16.

Adabas Product	Compatible Version Level
Transaction Manager (ATM)	version 1.2.2.
Vista (AVI)	version 7.1 or above

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.

Note:

Write IOs are always issued to maintain the integrity of the database.

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

Note:

A demo of Adabas Caching Facility version 7.4.1 is delivered in the ADA741.ALLINPL file.

Adabas Cluster Services

Adabas Cluster Services (ALS; formerly Adaplex+) implements multinucleus, multithread parallel processing and optimizes Adabas in an OS/390 or z/OS parallel sysplex (SYStems comPLEX) environment.

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 Entire Net-Work are required to accommodate Adabas Cluster Services.

Adabas Delta Save Facility

Adabas Delta Save Facility (ADE) offers significant enhancements to ADASAV utility processing by backing up and restoring 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 Manager

Adabas Manager (AMA) is a web-based user interface to Adabas on OS/390 or z/OS. It is used to analyze and control the operation of Adabas and the related environment: monitor system performance, administer databases, track operations and performance statistics, and analyze trends based on statistics over time. It operates under the Software AG System Management Hub, which is used to centrally manage a variety of Software AG products. The user web interface is a Javascript-enabled browser (Netscape or Internet Explorer).

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.

The delivered product contains the AOS base and also functions as a demo. See the section **Installation Requirements** on page for details.

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 user’s 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 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; however, you can optionally run Adabas Review version 4.3 with earlier versions of Adabas 7. If this is a requirement, contact your Adabas technical support representative. It is important to upgrade Adabas Review to version 4.3 before you upgrade Adabas to version 7.4.

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

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.

ASF uses approximately 450 data fields for monitoring Adabas databases and files. The data fields represent aspects of an Adabas database such as disk and buffer usage, thread usage, database load, ADARUN parameters, pool usage, and frequency of use of particular Adabas commands.

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

Adabas Transaction Manager (ATM) is a server for coordinating “two-phase commit processing” in distributed Adabas environments. It manages “global” transactions that are distributed across multiple Adabas databases by coordinating changes to the databases in a seamless, integrated way.

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).

Software AG recommends that all applications use the ADALNK module from a currently supported Adabas release. Effective with Adabas version 8.1, user applications using the ADALNK module of an Adabas release that is no longer supported may not function correctly.

Software AG recommends that you install the version 7.4 ADASVC and API link routines.

In a single nucleus (noncluster) environment

- the version 7.4 ADASVC is compatible with down-level Adabas databases; that is, even if you install the new SVC, you can run nuclei of Adabas version 7.1 or 7.2 with the installed Adabas 7.4 SVC.

Note:

Version 7.4 ADASVC with version 7.1.2 nucleus: a problem with DDPRINT displaying incorrect information behind the ZAP table of the ADASVC was corrected with ZAP AI712023.

- you can run Adabas version 7.4 with version 7.2 or 7.1 ADASVCs; that is, you can but need not install a new SVC.

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

For those who want to continue using Adabas Support for Multiprocessing, you may run the version 7.1 ADASMP module with the Adabas version 7.4 SVC/router if the related version 7.1 SMPSVC is linked into the version 7.4 ADASVC. See the section **Migration Issues** on page 16.

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.

In a noncluster environment when no new functionality is required, the following matrix shows which components are compatible with version 7.1 and 7.2 where ‘+’ is compatible; ‘-’ is incompatible; and ‘0’ is unrelated. The information in this matrix can be understood as follows:

- ADASVC, ADASIP, and ADASIR build one group.
- ADALNK constitutes a group by itself.
- ADANUC, ADAMPM, and ADAIOR build one group.
- Within each group, modules must be of the same version.
- Between groups, versions 7.1 and 7.2 or 7.4 can be mixed.

Note:

In all cases, the Adabas 7.4 SVC must be installed using the Adabas 7.4 SIP/SIR.

Version 7.4	Version 7.1 or 7.2					
	ADASIP/SIR	ADASVC	ADAMPM	ADALNK	ADAIOR	ADANUC
ADASIP/SIR		-	0	0	0	0
ADASVC	-		+	+	0	+
ADAMPM	0	+		0	0	-
ADALNK	0	+	0		0	+
ADAIOR	0	0	0	0		-
ADANUC	0	+	-	+	-	

Note:

Although the version 7.1 and 7.2 ADALNK routines can be used with the version 7.4 ADASVC, Software AG recommends using the version 7.4 ADALNK routines with the version 7.4 ADSVC.

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.

SVC Consolidation

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.

Migration Issues

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

If you want to use Adabas SMP rather than or in addition to Adabas Parallel Services, 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 to the (prelinked) ADASVC if you want to use a single SVC to run, for example, one database under Parallel Services and another database under Adabas SMP.

TP Link Routines

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, Com-plete, 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 chapter **Connecting UES-Enabled Databases** in the *Adabas Installation Manual (OS/390)* and page 32 in these *Release Notes*.

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.

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.

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 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.

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.

When migrating Adabas SAF Security (AAF), Adabas Fastpath (AFP), Adabas Vista (AVI), or System Coordinator for Adabas Options (COR), use the ADA731.MVSL001 dataset

The ADA731.MVSL001 library should be placed ahead of the ADA71x.MVSLOAD library in the JCL until the base Adabas product is also upgraded to version 7.4.

Warning:

The ADA731.MVSL001 library dataset must not be used with Adabas version 7.4.

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.

RELEASE NOTES FOR ADABAS VERSION 7.4.1

Summary of Features Added for 7.4

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 version 7.1 SMPSVC module is linked with the version 7.4 ADASVC module.
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 IBM's automatic restart management (ARM) product for automatic restarting.

More SVC Flexibility

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

Enhanced Backup and Restore Performance

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

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.

Data Integrity Enhancement

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 new additions 1 field. These changes apply only to formats created using the L3/L6 command to ensure that correct data are returned in an environment where Smith/Jones problems are possible. 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. Physical datasets for command and protection logging can be dynamically added and removed to increase availability by removing limits.
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.

Other Performance Enhancements

1. Adabas 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.
2. Adabas batch, COM-LETE and IMS link routines 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.

New and Enhanced ADARUN Options

1. The following parameters have been added to support 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.
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.

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 only one block of the table and the NUCID is always zero (0).

The PPT is created by ADADEF for a new database or by ADACNV for an existing database. Its address is stored in the GCB. The PPT entry can be modified by the nucleus or by ADARES.

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 the PLOG or CLOG is copied, its PPT entry is updated with the information. 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

- when the nucleus is autorestarting to warn about changes in the Work dataset.
- when ADARES PLCOPY or CLCOPY is copying all active protection logs or command logs, respectively.
- 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), which exists for both cluster and noncluster nuclei, is used to determine if any PLOGs or CLOGs still need to be copied from previous sessions. If the PPT indicates that PLOGs or CLOGs remain to be copied, the PLOG or CLOG datasets are read and if necessary, UEX2 or UEX12, as appropriate, is invoked.

A noncluster nucleus checks whether the previous session was a nucleus cluster session and has a pending autorestart. After a noncluster nucleus failure, the nucleus cluster is not allowed to start.

If PLOGs or CLOGs from a previous nucleus cluster session remain to be copied and a UEX2 or UEX12 exists, it will be invoked. This should invoke ADARES merge or PLCOPY/CLCOPY as required. A noncluster nucleus can overwrite block 1 of the PPT only when PLOGs from previous sessions have been processed to completion if PLOGRQ=FORCE is specified. If PLOGRQ=FORCE is not specified, a warning will be given and the entry will be over-written.

UEX2 or UEX12 controls the copy/merge process in conjunction with the PLOGRQ=FORCE parameter. UEX2 or UEX12 is required when PLOGRQ=FORCE is specified. If there is no UEX2 or UEX12, the PLOG and PPT entry are overwritten.

If an uncopied PLOG is detected that does not match that specified in the last session, and PLOGRQ=FORCE is specified, a parameter error occurs. If PLOGRQ=FORCE has not been specified, a warning will be given and the entry in the PPT will be overwritten. The NOPPT function can still be used to copy these PLOGs and the CLOG MERGE function can be used to merge the CLOGs after they have been copied. If the PLOG and CLOG are different from the previous session but they have been copied, the PPT entry is overwritten and the nucleus starts.

Log Copy Processing

In a cluster environment, the PLCOPY and CLCOPY processes access the parallel participant table (PPT) to determine the PLOGs or CLOGs to copy.

Dynamic Allocation

In a noncluster environment, the nucleus and ADARES use all of the WORK, PLOG, and CLOG datasets listed in the JCL. They use dynamic allocation to perform consistency checks only if the dataset names are changed and the PPT indicates that the old ones still contain needed information.

Support for the Adabas SQL Server (AQA)

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

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 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

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

Previous versions of these modules required editing the source member to change the SVCNR EQU to override the default value of 249 for the z/OS Adabas SVC number. Now the value may be specified in the assembly JCL on the SYSPARM assembler step PARM statement. The value provided will be checked by conditional assembly statements to ensure that it is a valid decimal number in the range 200–255. If an improper SYSPARM value is provided, an MNOTE is generated and the assembly step will terminate with condition code 16. If no value is coded on the SYSPARM assembly step PARM statement, the value coded on the SVCNR EQU statement in the source module will be used. This permits either the specification of the default Adabas SVC number by editing the source member and then assembling it, or through the SYSPARM JCL parameter.

Note:

It is still necessary to modify the supplied source members to set other default values such as LOGID, LUINFO, LRINFO, etc. 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.

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 (buffer flushes occur in ascending RABN sequence without regard to the distribution of the blocks over volumes), the minimum, default, and maximum values of FMXIO continue to be 1, 60, and 100, respectively.

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

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 7.4.1 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 7.4.1 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 7.4.1 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 DDPLOGnn and DDCLOGnn 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 have been linked with LNKUES and the default translation tables ASC2EBC and EBC2ASC. The load module LNKUESL has been dropped from the ADA74x.MVSLOAD 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 AII74x.MVSSRCE 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

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.

Removed from Adabas

BUB Support Removed

The Block of Unreadable Blocks (BUB) is 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.

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.

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.

Changes to Adabas Online System

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).

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, ADA741.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 with the REPLACE option. When allocated to the 'running' Adabas environment, this completes the AOS 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 Adabas, AOS will function as a demo unless an Adabas ZAP is applied, as follows:

Adabas Version	ZAP Number
7.1.2	A\$712099
7.1.3	A\$713099
7.2.2	AY722099

ZAP Prefixes for Adabas Online System (AOS)

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

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

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

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

Display Cluster Members

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

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 *****          2000-07-19
DBID 105          - Display PPT Entry -          PACA002

Nuc ID. . . 3 Active Nucleus

Name              Status              DataSet Name
-----
WORK1
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
```

Nucleus File Status

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  Access
25          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 update 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.

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          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
```

Cluster Usage

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

```

Cache Statistics

Choosing **cache statistics** (option C) from the Cluster Usage menu displays the following menu:

```
16:14:23          ***** A D A B A S  BASIC  SERVICES *****          2002-05-29
                                     - Cache Statistics -                               PACUX12

Code   Service
----   -
  K   Cast-out / Directory
  P   Publishing requests
  X   Individual cache blocks
  .   Exit
  ?   Help
----   -

Code .....
Database ID .. 1955   (WIS1955)                NucID .. 1021

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

Cast-out / Directory

Choosing **cast-out / directory** (option K) from the Cache Statistics menu display the following:

```

16:14:23          ***** A D A B A S  BASIC  SERVICES *****          2002-05-29
DBID 1955          - Cast-out / Directory -          PACUX12
NucID 1021

      Cast-out Directory Reads          Directory Reads
      -----
Total .....          28          Total .....          5
  Sync .....          1          Sync .....          1
  Async ....          27          Async ....          4

Unlock Cast-out Calls
-----
Total .....          28
  Sync .....          1
  Async ....          27

PF1----- PF2----- PF3----- PF4----- PF7----- PF8----- PF9----- PF12-----
Help          Exit      Refresh          Detail      Menu
  
```

Counters have a multiplier column with the following values:

Value The total shown is in . . .

blank	(factor of 1)
K	kilo (factor of 1,000)
M	mega (factor of 1,000,000)
G	giga (factor of 1,000,000,000)

If a number has a multiplier shown, it has been divided by the multiplier, showing the significant digits to 9 places with no decimal point.

Press PF9 to see the entire value. This value is the exact count up to 20 digits in length.

Publishing Requests

Choosing **publishing requests** (option P) from the Cache Statistics menu display the following:

```
16:26:21          ***** A D A B A S  BASIC  SERVICES *****          2002-05-29
DBID  1955          - Publishing Requests -          PACUX12
NucID 1021

Publishing Request Category
-----
Update sync .....          34

BT or CL or ET ....          162

Redo threshold ....          2

Full bufferpool ...          0

All blocks .....          84

Specific RABN .....          0

File DS blocks ....          4

PF1----- PF2----- PF3----- PF4----- PF7----- PF8----- PF9----- PF12-----
Help          Exit      Refresh          Detail      Menu
```

All Cache Blocks

Choosing **all cache blocks** (option X) from the Cache Statistics menu display the following:

```

16:27:05          ***** A D A B A S  BASIC  SERVICES *****          2002-05-29
DBID  1955          - All Cache Blocks -          PACUX12
NucID 1021

  Reads                                     Writes
-----
Total .....                167                Total .....                38,176
  Sync .....                 24                Sync .....                15,148
  Async .....                143                Async .....                23,028

  In cache .....              49                Written .....              38,176
  Not in cache ..             118               Not written .....          0
  Struc. full ...              0                Struc. full .....          0

Cast-out Reads                               Other
-----
Total .....                212                Validates .....           187,677
  Sync .....                 212                Invalid .....              43
  Async .....                  0                Deletes .....              0
                                                Timeouts .....             0
                                                Redo processes .....        0

PF1----- PF2----- PF3----- PF4----- PF7----- PF8----- PF9----- PF12-----
Help   Repos   Exit    Refresh  PrevBlk  NxtBlk  Detail   Menu
  
```

Use PF7 and PF8 to scroll through the cache blocks; use PF2 to reposition.

Statistics are displayed for the following:

- All cache blocks;
- Address converter (AC) cache blocks;
- Data Storage (DS) cache blocks;
- Data Storage space table (DSST) cache blocks;
- File control block (FCB) cache blocks;
- Normal index (NI) cache blocks; and
- Upper index (UI) cache blocks.

Press PF9 from the above screen to display the following detail screen:

```

16:27:05          ***** A D A B A S  BASIC  SERVICES *****          2002-05-29
DBID  1955          - All Cache Blocks -          PACUX12
NucID 1021
Reads                                Writes
-----
Total .....                167 Total .....                38,176
  Sync .....                 24  Sync .....                15,148
  Async .....                143 Async .....                23,028

  In cache..                 49  Written ...                38,176
  Not in ...                 118 Not writ ..                 0
  Stru.full.                  0   Stru.full .                 0

Cast-out Reads                      Other
-----
Total .....                212 Validates ...            187,677
  Sync .....                212 Invalid ...                43
  Async .....                 0 Deletes .....                0
                                   Timeouts ..                0
                                   Redo procs ..                0

```

Press Enter to continue

File Statistics

Choosing **file statistics** (option F) from the Cluster Usage menu for file 25 displays the following menu:

```

16:37:02          ***** A D A B A S  BASIC  SERVICES *****          2002-05-29
DBID  1955          - File 25 Statistics -          PACUX22
NucID 1021

  Reads                                     Writes
-----
Total .....          67          Total .....          20,157
  Sync .....          0          Sync .....          7,583
  Async .....         67          Async .....          12,574

  In cache .....          0          Written .....          20,157
  Not in cache ..         67          Not written .....          0
  Struc. full ...          0          Struc. full .....          0

Cast-out Reads                             Other
-----
Total .....          78          Validates .....          79,248
  Sync .....          78          Invalid .....          0
  Async .....          0          Deletes .....          0
                                         Timeouts .....          0
                                         Redo processes .....          0

PF1----- PF2----- PF3----- PF4----- PF7----- PF8----- PF9----- PF12-----
Help      Repos      Exit      Refresh          Detail      Menu

```

Lock Statistics

Choosing **lock statistics** (option L) from the Cluster Usage menu displays the following menu:

```

16:38:16          ***** A D A B A S  BASIC  SERVICES *****          2002-05-29
                                - Lock Statistics -                               PACUX32

Code  Service                      Code  Service
-----
A    Buffer flush lock              I    Global update command sync lock
B    Checkpoint lock              J    Hold ISN lock
C    DSF lock                     K    New-Data-RABN lock
D    ETID lock                    L    Online save lock
E    File-lock-table lock         M    Parameter lock
F    FST lock                     N    Recovery lock
G    GCB lock                     O    RLOG lock
H    Global ET sync lock         P    Security lock
.    Exit                         Q    Spats lock
?    Help                         R    Unique descriptor lock
-----

Code ..... _
Database ID .. 1955   (WIS1955)           NucID .. 1021

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

```

Each of the options on the Lock Statistics menu displays statistics for a particular lock. For each lock, the screen displays obtain and release information about the various types of that lock that are currently in use by a cluster nucleus:

- The system may obtain locks conditionally or unconditionally, synchronously or asynchronously. A conditional request for a lock may be granted or rejected.
- Releases may be issued synchronously or asynchronously.

Hold ISN Lock

Choosing **hold ISN lock** (option J) from the Lock Statistics menu displays the following:

16:38:16	***** A D A B A S BASIC SERVICES *****	2002-05-29					
DBID 1955	- Hold ISN Lock -	PACUX32					
NucID 1021							
Obtains		Releases					
-----		-----					
Conditional	16,017	Issued	16,017				
Granted	16,017	Sync	15,971				
Rejected	0	Async	46				
Unconditional ..	0						
Sync	158						
Async	15,859						
PF1-----	PF2-----	PF3-----	PF4-----	PF6-----	PF7-----	PF8-----	PF12-----
Help	Repos	Exit	Refresh		PrevLok	NxtLok	Menu

Use PF7 and PF8 to scroll through the locks; use PF2 to reposition.

Maintain the User Table

Note:

This option is available in Adabas nucleus 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.



APPENDIX A – THE AFPLOOK UTILITY

This appendix describes the AFPLOOK utility, which is distributed on the Adabas 7.1 release tape as part of ADAAFP and on the Adabas version 7.2.2 beta test tape as part of AFPADA. It is therefore available to all Adabas customers before any Fastpath decision is made. 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 utility AFPLOOK is included with Adabas version 7.1 as ADAAFP and with Adabas beta version 7.2.2 as AFPADA. It 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

Full-function Fastpath users may wish to zap the AFPLOOK parameters to avoid the need to change them each time AFPLOOK is started using SYSAFP.

Especially for this purpose, an option mentioned in the section **Starting an AFPLOOK Session** in the *Adabas Fastpath Manual* makes it possible to **mark to use zapped parameters**. In this case the ZAP in the following sections applies.

Adabas Version 7.1

To override the parameter defaults, apply a ZAP to AFPLOOK similar to the following:

```

NAME ADAAFP AFPPCLK
VER 005C 0040           Maximum files
VER 005E 0020           Commands/descriptors per file
VER 0060 0064           Maximum concurrent users
VER 0062 000A           CIDs per user
VER 0064 7FFFFFFF       Maximum commands
VER 0068 40404040404040 Job name
VER 0070 0037           Lines per page
VER 0072 40404040404040 Customer subheading (30 characters)
VER 0090 FF             On/off switch (ON)
VER 0091 00000000000000 Selected files (2-byte hexadecimal)
REP 005C 0040
REP 005E 0020
REP 0060 0064
REP 0062 000A
REP 0064 7FFFFFFF
REP 0068 40404040404040
REP 0070 0037
REP 0072 40404040404040
REP 0090 FF
REP 0091 000700FE8FFF   Select files 7 254 2303

```



Adabas Version 7.2.2

To override the parameter defaults, apply a ZAP to AFPLOOK similar to the following:

```
NAME AFPADA AFPPCLK
VER 00A8 0040           Maximum files
VER 00AA 0020           Commands/descriptors per file
VER 00AC 0064           Maximum concurrent users
VER 00AE 000A           CIDs per user
VER 00B0 7FFFFFFF       Maximum commands
VER 00B4 40404040404040 Job name
VER 00BC 0037           Lines per page
VER 00BE 40404040404040 Customer subheading (30 characters)
VER 00DC FF             On/off switch (ON)
VER 00DD 00000000000000 Selected files (2-byte hexadecimal)
REP 00A8 0040
REP 00AA 0020
REP 00AC 0064
REP 00AE 000A
REP 00B0 7FFFFFFF
REP 00B4 40404040404040
REP 00BC 0037
REP 00BE 40404040404040
REP 00DC FF
REP 00DD 000700FE8FFF       Select files 7 254 2303
```



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 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.
SEQUENTIAL	Maximum sequence commands that can be optimized.
SEQUENCES	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 7.2.2 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 version 7.2.2 or above. 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 7.3.1 *****      2001-11-08
                                     - 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 PLXID or SMPID) of the cluster nucleus you wish to monitor.

```

+-----+
| 13:31:07      Select Cluster Members      2001-02-09 |
|                                                    |
| 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 7.3.1 *****      2001-09-07
                - AVILOOK  File Maintenance -                          V14100M1

DBID: 231      (TEST-V7-DB)

C   File      Command Limit      Commands      Started      Status
-   12         0                0              0              Paused
-   2          0                5768          1999-09-24 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 70 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                      1999-09-24 |
|                                     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 68) 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 7.3.1 *****      2001-09-07
                - Avilook File Statistics -                          V14120M1

DBID: 231      (TEST-V7-DB)
File: 2                                               Started: 1999-09-24 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

- Adabas, support for earlier releases, 5
- Adabas Caching Facility, 8
- Adabas Delta Save Facility, AOS requirement, 9
- Adabas External Security Interface, 11
- Adabas Fastpath, 9
- Adabas Online System, 10
 - installation requirements, 35
 - ZAP prefixes, 36
- Adabas Parallel Services
 - multiple thread processing, 10
 - parallel processing, 10
- Adabas Review, overview, 11
- Adabas SAF Security, 11
- Adabas Vista, 13
- ADACMP, version compatibility, 19
- ADAESI, 11
- ADALOD, LOAD VERSION parameter, 19
- ADAORD, version compatibility, 19
- ADARAI, compatibility, 18
- ADASAF, 11
- ADASAV
 - version compatibility, 18
 - Work dataset reset, 32
- ADASIP, releasing common storage, 31
- ADAULD, version compatibility, 19
- AFPLOOK, 55–64
 - how to run, 56
 - overview, 55
 - parameters, 56
 - overriding defaults, 59
 - potential optimization summary, 62
 - sample command analysis, 63
 - sample file report, 61

- sample parameters report, 64
 - ZAP for, 59
- Analysis tool, AVILOOK, 65
- ARM, 24
 - explanation, 28
- ARMNAME, ADARUN parameter, 28
- ASYTVS, ADARUN parameter, 29
- ATM, definition, 13
- Automatic Restart Management, 24, 28
- AVILOOK, 65
 - add files, 69
 - list files, 66
 - maintain files, 70

B

- Buffer flush, 29

C

- CIDs per user, maximum, 57
- CLCOPY, 25, 26
- Cluster
 - display usage statistics, using Basic Services, 42
 - maintain user table, 51
- Command analysis, sampler, 55–64
- Command/descriptors per file, 56
- Commands
 - excluded, 62
 - maximum processed, 57
- Communication modules, version compatibility, 14

Adabas Version 7.4.1 Release Notes

D

data field, 12
Database, single physical, defined, 10
DD/JCLIN file, 18
Dynamic allocation, 26

E

Entire Net-Work, component, 8
Excluded commands, 62

F

Files
 maximum, 56
 selected, 58
FMXIO, ADARUN parameter, 29

H

Heading, report, 58

J

Job, name, 57

L

Lines per page, number of, 58
Log copy process, 26

N

Nucleus, cluster, definition of, 10

Nucleus file, display status, using Basic Services, 39
Nucleus ID, 24

O

On/off switch, 58
Operating systems, supported, 8

P

Parallel Participant Table, 24
PLCOPY, 25, 26
PPT, 24
PREFICMD, ADARUN parameter, 30
PREFIFIL, ADARUN parameter, 30

R

RC command, 62
Response codes, record-related, 31

S

Sequential PLOGs, version compatibility, 20
Session, display status, using Basic Services, 40
Sixty-four (64) bit storage support
 real, 5
 virtual, 6
SVC, and link routines, 31
System, display status, using Basic Services, 40

U

User table, maintain, 51
Users, maximum concurrent, 57

W

Work dataset, reset by ADASAV, 32

Z

ZAP, for AFPLOOK, 59

ZAP prefixes, Adabas Online System, 36

Adabas Version 7.4.1 Release Notes

Adabas Version 7.4.1 Release Notes