

Adabas Manager

Adabas Manager Documentation

Version 7.7.3

March 2013

Adabas Manager

This document applies to Adabas Manager Version 7.7.3.

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

Copyright © 2013 Software AG, Darmstadt, Germany and/or Software AG USA, Inc., Reston, VA, United States of America, and/or their licensors.

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

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

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

Document ID: AMA-MANUAL-773-20130315

Table of Contents

Preface	vii
1 Introduction	1
What You Can Do with Adabas Manager	2
Workbench Functionality	4
I Working with Adabas Manager	9
2 Accessing the Adabas Manager Area	11
3 Adabas Manager Window Layout	13
4 Adabas Manager Configuration	15
5 Adabas Manager Session and Cache Administration	17
6 Trace Version	19
7 Leaving Adabas Manager	21
II Adabas Manager Security	23
8 SAF Security Support	25
9 Working with a Protected Nucleus	27
Accessing a Protected Nucleus	48
Terminating a Session with a Protected Nucleus	28
Passwords	29
Non-activity Timeout	29
III Adabas Databases	31
10 Listing Databases	33
Registering or Unregistering Mainframe Databases	35
Adding Existing Open Systems Databases	36
Creating New Open Systems Databases	37
Deleting Open Systems Databases	39
11 Starting and Stopping Databases	41
12 Mainframe Databases that Cannot Be Reached	43
13 Selecting Databases	45
Selecting Mainframe Databases	46
Selecting Open Systems Databases	49
14 Reviewing Adabas Mainframe Database DIBs	53
15 Reviewing Adabas Open Systems Database UCBs	55
16 Database-related Long Running Utilities for Open Systems	57
Backup (ADABCK)	58
Create WORK (ADD_CONTAINER)	59
Create SORT (ADD_CONTAINER)	59
Create TEMP (ADD_CONTAINER)	60
Error Log (ADAERR)	60
Display Descriptors (ADAFIN)	60
File Usage (ADAFIN)	61
List Protection Log (ADAPLP)	61
Regenerate (ADAREC)	61
Reorder (ADAORD)	62
Restore (ADABCK)	63

Unused Storage (ADAREP)	64
Show Layout (ADAREP)	64
Verify (ADAVFY)	64
IV Checkpoints	67
17 Checkpoints	69
Types of Adabas Users	70
Checkpoint Names Starting with 'SYN'	70
Checkpoint Numbers for 'SYN' Names (Mainframe Databases Only)	71
V Database Files	75
18 Listing Database Files	77
Displaying the List of Files	78
Changing the Columns in the File List Display	78
Displaying Locked Files	80
19 Reviewing File Properties	81
20 Displaying Files with a Critical Number of Extents (Adabas Version 7 only)	85
21 Files with Extent List Saturation (Adabas Version 8 only)	87
22 Field Definition Table (FDT)	89
Referential Integrity Constraints	90
Changing Field Parameters	91
Logically Deleting/Undeleting Fields	91
Descriptors	92
Predict Parameters for Mainframe Databases	94
Adding New Fields in the Field Definition Table (FDT)	94
23 Modifying File Parameters	95
Changing Padding Factors, Maximum Allocation, and Record Length	96
Selecting File Options	97
Record Spanning, MUPEX, LOB Fields/File (Adabas Version 8 and Open Systems Only)	97
24 Allocating/Deallocating File Space	101
25 Modifying Files	103
Locking/Unlocking Files	104
Reordering Files	104
Deleting Files	105
Renumbering Files	105
Renaming Files	105
Uncoupling Files	106
Inverting Files	106
26 Maintaining Expanded Files (Mainframes Only)	109
27 Creating New Files	111
28 Creating New LOB Files (Adabas for Mainframes Version 8 and Open Systems Only)	119
29 File-related Long Running Utilities for Open Systems	121
Backup (ADABCK)	122
Restore (ADABCK)	123

Load (ADAMUP)	124
Unload (ADAULD)	125
Compress (ADACMP)	127
Decompress (ADADCU)	129
Import (ADAORD)	131
Export (ADAORD)	132
Verify (ADAVFY)	133
Invert Descriptor (ADAINV)	134
Reinvert Descriptor	135
Verify Descriptor (ADAINV)	135
Descriptor Summary (ADAINV)	136
Display Descriptors	136
File Usage	136
VI Session Monitoring	137
30 Active Parameters	139
Active Parameters for Mainframe Databases	140
Active Parameters for Open Systems Databases	143
31 Caching Facility	151
Defining a File Cache	152
Maintaining Cache Parameters	153
Cache Session Summary	154
32 Delta Save Facility	155
Displaying Files Modified by Utilities	157
Displaying Size Estimate for Delta	157
33 Disk Storage	159
Adding a Data Storage or Associator Dataset	160
Modifying a Data Storage or Associator Dataset for Mainframe Databases	161
Modifying a Data Storage or Associator Dataset for Open Systems Databases	162
Recovering Unused Storage	163
Reviewing Volume Distribution for a Mainframe Database	163
34 Mainframe Database Maintenance Levels	165
35 System Status	167
Reviewing General System Status of Databases	168
Resetting the System Status for Mainframe Databases	170
CLOG/PLOG Statistics	171
Reviewing Command Queue	172
Reviewing Command Usage	172
Reviewing File Usage	174
Reviewing High Water Marks	174
Reviewing Hold Queue	175
Installed Products	176
Reviewing Nucleus File Status	176
Reviewing Online Status	177

Reviewing PPT (Parallel Participant Table) Status	178
Reviewing Resource Statistics	179
Reviewing Thread Status	182
User Profiles	183
User Queue and User Queue Elements (UQEs) for Databases	184
Reviewing Work Status	190
Reviewing Work Pool Usage	191
VII Space Calculation for Mainframe Databases	193
36 Space Calculation for Mainframe Databases	195
Estimating Associator Space	197
Estimating Data Storage Space	197
Estimating Space for the DD/FILEA Sequential Dataset	198
Estimating the Sort Dataset	198
Estimating the Temp Dataset	199
Estimating the Work Dataset	200
Index	201

Preface

This document tells you how to install, configure, and use the Adabas Manager version 7.7.3, a web-based graphical user interface (GUI) you can use to perform administrative tasks for Adabas.

Adabas Manager provides browser-based administration and monitoring of Adabas on mainframes as well as on open systems platforms; it offers the functionality previously available with the DBA Workbench tool.

Adabas Manager runs in the System Management Hub (SMH), the standard, GUI-based, central point of administration for Software AG's products. The System Management Hub itself runs in a standard web browser.

Adabas Manager version 7.7.3 also serves as the framework for additional components of Adabas and related products.

Who Should Read this Documentation

The success of a database environment depends on central control of database design, implementation, and use. This central control and coordination is the role of the database administrator (DBA).

The DBA also needs

- administrative skill to set up and enforce the standards and procedures for using the database;
- technical ability to understand the factors governing hardware performance, with considerable knowledge both of the operating system software and the DBMS being used;
- a thorough knowledge of existing and future applications; and
- skills to produce an efficient database design that meets the application requirements.

The DBA using Adabas Manager to monitor and change aspects of an Adabas database while an Adabas session is active should have in-depth understanding of the database operation and the service it provides to the organization.

Contents of the Documentation

The Adabas Manager documentation is organized in the following parts:

Introduction	describes the functionality of Adabas Manager and introduces the new functionality for administering open systems databases in addition to Adabas mainframe databases.
Working with Adabas Manager	explains how to access the Adabas Manager area and the Adabas Manager window layout.
Security	describes the Adabas Manager security functionality.
Working with Adabas Databases	introduces database services: listing databases, registering and unregistering mainframe databases, adding and existing or creating a new open systems database, and selection of databases.
Checkpoints	lists checkpoints.
Database Files	lists files, displays for a selected file its file control block (FCB) and file space; allows you to add new files.
Session Monitoring	presents the Adabas session monitoring functions, which allow you to control and manage major Adabas resources. These functions are most useful when analyzing system performance or seeking the cause of performance problems.
Space Calculation	introduces a planning tool for adding new components or recalculating existing space requirements.

1 Introduction

- What You Can Do with Adabas Manager 2
- Workbench Functionality 4

This section of the documentation describes the functionality of Adabas Manager and introduces the new functionality for administering open systems databases in addition to Adabas mainframe databases.

What You Can Do with Adabas Manager

Adabas Manager replaces the Basic Services component of the Adabas Online System and includes functions comparable to the Adabas operator commands and utilities.

Version 7.7 of Adabas Manager also provides the full functionality of the DBA Workbench tool and supports the administration and monitoring of Adabas databases on mainframes as well as on open systems platforms.

Adabas Manager 7.7 also supports Adabas Caching Facility and Adabas Delta Save Facility.

Adabas Manager allows the DBA to interactively monitor and change aspects of an Adabas database while an Adabas session is active. Its web-based graphical user interface (GUI) allows the DBA to easily view resource status and user queues; display and revise space allocation; change file and database parameters; define a new file online; and stop a selected user or current Adabas session.

DBA tasks directly affecting the database (designing, selecting tuning aids, predicting changes, deciding and enforcing database rules, organizing and administering file conversions, measuring database use, and tuning for best database performance) can all be done using Adabas Manager.

For analyzing performance and monitoring database operation, Adabas Manager allows you to see the system from the viewpoint of either a user or a particular system resource. For example, you can perform the following user-related tasks:

- monitor the operation of a particular user or all users
- check user and command queue status
- change user priority
- stop a user from using a particular file, or stop all inactive users

Some of the resource-related tasks you can perform are

- view and update nucleus parameters
- monitor command and file usage and system performance information
- list file layout and extent status; and
- list file distribution of the database by VOLSER

For controlling the overall Adabas session, Adabas Manager provides you with ways to

- change file parameters
- define new files and rename, renumber, refresh, uncouple, or delete files
- create new FDTs, fields, and change field lengths
- switch ISN and Data Storage reuse on and off by file, and release descriptors
- terminate an Adabas nucleus session (ADAEND, HALT, CANCEL)
- enable better use of the available operating system facilities through the Adabas Caching Facility, in that valuable virtual memory resources are not monopolized.
- enhance ADASAV utility processing by backing up and restoring only the changed (delta) portions of Adabas databases through the Adabas Delta Save Facility.

Enterprise Edition and Standard Edition

The Adabas Manager Enterprise Edition supports a range of functionality which is not available in the Standard Edition for mainframe databases; the functionality for open systems databases is *not* restricted in this way (this may change in future versions of Adabas Manager).

For mainframe databases, the following functionality is supported by the Adabas Manager Enterprise Edition only:

- **Caching Facility:**
enables better use of the available operating system facilities, in that valuable virtual memory resources are not monopolized.
- **Delta Save Facility:**
enhances ADASAV utility processing by backing up and restoring only the changed (delta) portions of Adabas databases.
- **Checkpoints:**
allows the administrator to review the checkpoint list for an Adabas database and to delete checkpoints.
- **Disk Storage:**
allows the administrator to modify or add Associator and Data Storage datasets and to recover unused storage.
- **Space Calculation:**
is a planning tool for adding new components or recalculating existing space requirements.
- **System Status (extended):**
provides a comprehensive overview of Adabas operation:

CLOG/PLOG Statistics

Command Queue

Command Usage

File Usage


High Water Marks

Nucleus File Status


Online Status
 PPT (Parallel Participant Table)
 Resource Statistics
 User Profiles
 User Queue
 Work Status
 Work Pool Usage

Workbench Functionality

Adabas Manager version 7.7 is able to administer and monitor Adabas databases on open systems platforms. It replaces the DBA Workbench tool; Adabas Manager offers the same functionality previously available through the workbench and has been designed to handle different versions of Adabas in parallel. Adabas 6.1.11 for Open Systems is the minimum version supported.

 **Important:** Although it is possible to run Adabas Manager and the DBA Workbench in parallel, Software AG does not recommend this under any circumstances. In cases where a parallel use cannot be avoided, the DBA Workbench must never be active while Adabas Manager is running.

The following tables explain the DBA Workbench functions which can now be executed with Adabas Manager and refers to the information in this documentation.

 **Note:** Where a target can be reached by right-clicking on a node, the character † is displayed; for example, "DB nnn † Nucleus Log" asks you to right-click on the database name and choose **Nucleus Log** from the dropdown menu.

Report

DBA Workbench	Adabas Manager	Documentation
General Information		
Contents	DB nnn >Database Files > <i>File Name</i>	Reviewing File Properties in <i>Database Files</i>
Summary	DB nnn and DB nnn > Disk Storage > Containers	Listing Databases in <i>Adabas Databases</i> and Disk Storage in <i>Session Monitoring</i>
Layout	DB nnn > Utilities > Show Layout	Show Layout in <i>Database-related Long Running Utilities for Open Systems</i>
Free Space	DB nnn > Utilities > Unused Storage	Unused Storage in <i>Database-related Long Running Utilities for Open Systems</i>
Nucleus Log	DB nnn † Nucleus Log	Nucleus Log in <i>Listing Databases</i>
File Information		

DBA Workbench	Adabas Manager	Documentation
Status	DB nnn > Database Files > <i>File Name</i> and DB nnn > Disk Storage	Reviewing File Properties in <i>Database Files and Disk Storage in Session Monitoring</i>
Descriptor	DB nnn > Utilities > Display Descriptors	
Usage	DB nnn > Utilities > File Usage	File Usage in <i>Database-related Long Running Utilities for Open Systems</i>
Session		
Command Queue	DB nnn > System Status > Command Queue	Reviewing Command Queue in <i>System Status</i>
Full User Queue	DB nnn > System Status > User Queue	User Queue and User Queue Elements (UQEs) for Databases in <i>System Status</i>
Hold Queue	DB nnn > System Status > Hold Queue	Reviewing Hold Queue in <i>System Status</i>
Internal Command Queue	DB nnn > System Status > Command Queue Internal CMDQ	Reviewing Command Queue in <i>System Status</i>
Thread Table	DB nnn > System Status > Threads	Reviewing Thread Status in <i>System Status</i>
User Files	DB nnn > System Status > User Queue	User Queue and User Queue Elements (UQEs) for Databases in <i>System Status</i>
User Queue	DB nnn > System Status > User Queue	User Queue and User Queue Elements (UQEs) for Databases in <i>System Status</i>
User Time Limits	DB nnn > System Status > User Queue	User Queue and User Queue Elements (UQEs) for Databases in <i>System Status</i>
Utility Communication Block	DB nnn Utility Communication Block	Reviewing Adabas Open Systems Database UCBs in <i>System Status</i>
Session Parameters		
Dynamic	DB nnn > Active Parameters	Active Parameters in <i>Session Monitoring</i>
Static	DB nnn > Active Parameters	Active Parameters in <i>Session Monitoring</i>
Statistics		
High Water Marks	DB nnn > System Status > High Water Marks	Reviewing High Water Marks in <i>System Status</i>
Command	DB nnn > System Status > Command Usage	Reviewing Command Usage in <i>System Status</i>
Activities	DB nnn > System Status	System Status in <i>Session Monitoring</i>
File IO	DB nnn > System Status > File Usage	Reviewing File Usage in <i>System Status</i>
Set List of Files	DB nnn > System Status > File Usage > Define Files	Reviewing File Usage in <i>System Status</i>
Buffer Pool	DB nnn > System Status	System Status in <i>Session Monitoring</i>
Format Pool	DB nnn > System Status	System Status in <i>Session Monitoring</i>
Protection Log	DB nnn > System Status > CLOG/PLOG Statistics	CLOG/PLOG Statistics in <i>System Status</i>

DBA Workbench	Adabas Manager	Documentation
Checkpoints	DB nnn > Checkpoints	Checkpoints
Command Log	DB nnn > System Status > CLOG/PLOG Statistics	CLOG/PLOG Statistics in <i>System Status</i>
Error Log	DB nnn > Utilities > Error Log	Error Log in <i>Database-related Long Running Utilities for Open Systems</i>
List Protection Log	DB nnn > Utilities > List Protection Log	List Protection Log in <i>Database-related Long Running Utilities for Open Systems</i>

Profile

DBA Workbench	Adabas Manager	Documentation
Containers	DB nnn > Disk Storage	Disk Storage in <i>Session Monitoring</i>
Nucleus Parameters	DB nnn > Active Parameters	Active Parameters in <i>Session Monitoring</i>
Start as Service	DB nnn Enable Start as Service	Starting and Stopping Databases in <i>Adabas Databases</i>
Environment Extensions	DB nnn > Active Parameters Environment Extensions	Environment Extensions in <i>Active Parameters</i>
Operations		
Node Parameters	DB nnn > Active Parameters Node Parameters	Node Parameters in <i>Active Parameters</i>
Database Actions	DB nnn > Active Parameters Database Actions	Database Actions in <i>Active Parameters</i>
Action Users	DB nnn > Active Parameters Action Users	Action Users in <i>Active Parameters</i>
Add Existing Database	Adabas Databases Add Existing DB	Add Existing Open Systems Databases in <i>Listing Databases</i>

File

DBA Workbench	Adabas Manager	Documentation
New	DB nnn > Database Files New File	Creating New Files in <i>Database Files</i>
Delete	DB nnn > Database Files > Delete	Deleting Files in <i>Modifying Files</i>
Refresh	DB nnn > Database Files > Refresh	
Add LoB File	DB nnn > Database Files New LOB File	Creating New LOB Files in <i>Database Files</i>
Backup	DB nnn > Database Files > Utilities > Backup (dropdown list)	Backup in <i>File-related Long Running Utilities for Open Systems</i>
Restore	DB nnn > Database Files > Utilities > Restore (dropdown list)	Restore in <i>File-related Long Running Utilities for Open Systems</i>
Load	DB nnn > Database Files > Utilities > Load (dropdown list)	Load in <i>File-related Long Running Utilities for Open Systems</i>

DBA Workbench	Adabas Manager	Documentation
Unload	DB nnn > Database Files > Utilities > Unload (dropdown list)	Unload in <i>File-related Long Running Utilities for Open Systems</i>
Compress	DB nnn > Database Files > Utilities > Compress (dropdown list)	Compress in <i>File-related Long Running Utilities for Open Systems</i>
Decompress	DB nnn > Database Files > Utilities > Decompress (dropdown list)	Decompress in <i>File-related Long Running Utilities for Open Systems</i>
Export	DB nnn > Database Files > Utilities > Export (dropdown list)	Export in <i>File-related Long Running Utilities for Open Systems</i>
Import	DB nnn > Database Files > Utilities > Import (dropdown list)	Import in <i>File-related Long Running Utilities for Open Systems</i>
Verify	DB nnn > Database Files > Utilities > Verify (dropdown list)	Verify in <i>File-related Long Running Utilities for Open Systems</i>

Administrate (File Level)

DBA Workbench	Adabas Manager	Documentation
Allocate	DB nnn > Database Files > Alloc. > Allocate	Allocating/Deallocating File Space in <i>Database Files</i>
Deallocate	DB nnn > Database Files > Alloc. > Deallocate	Allocating/Deallocating File Space in <i>Database Files</i>
Reusage		
Data Storage	DB nnn > Database Files > <i>File Name</i> > Modify	Modifying Files in <i>Database Files</i>
ISNs	DB nnn > Database Files > <i>File Name</i> > Modify	Modifying Files in <i>Database Files</i>
Lock	DB nnn > Database Files > Lock	Locking/Unlocking Files in <i>Modifying Files</i>
Unlock	DB nnn > Database Files > Unlock	Locking/Unlocking Files in <i>Modifying Files</i>
Enable PGM_REFRESH	DB nnn > Database Files New File > Enable PGM_REFRESH	Creating New Files in <i>Database Files</i>
Disable PGM_REFRESH	DB nnn > Database Files > <i>File Name</i> > Modify	Modifying Files in <i>Database Files</i>
Modify FDT	DB nnn > Database Files > <i>Fields</i> > Add	Adding New Fields in the Field Definition Table (FDT) in <i>Field Definition Table (FDT)</i>
Rename	DB nnn > Database Files > Rename	Renaming Files in <i>Modifying Files</i>
ReNUMBER	DB nnn > Database Files > ReNUMBER	Renumbering Files in <i>Modifying Files</i>

Descriptor (File Level)

DBA Workbench	Adabas Manager	Documentation
Invert	DB nnn > Database Files > Utilities > Invert Descriptor (dropdown list)	Invert Descriptor in <i>File-related Long Running Utilities for Open Systems</i>
Reinvert	DB nnn > Database Files > Utilities > Reinvert Descriptor (dropdown list)	Reinvert Descriptor in <i>File-related Long Running Utilities for Open Systems</i>
Release	DB nnn > Database Files > <i>Fields</i> > Name Release Descriptor	Setting and Releasing Descriptors in <i>Field Definition Table (FDT)</i>
Unique		
Set	DB nnn > Database Files > <i>Fields</i> > Name Set Unique Descriptor	Setting and Releasing Descriptors in <i>Field Definition Table (FDT)</i>
Reset	DB nnn > Database Files > <i>Fields</i> > Name Reset Unique Descriptor	Setting and Releasing Descriptors in <i>Field Definition Table (FDT)</i>
Verify	DB nnn > Database Files > Utilities > Verify Descriptor (dropdown list)	Verify Descriptor in <i>File-related Long Running Utilities for Open Systems</i>
Summary	DB nnn > Database Files > Utilities > Descriptor Summary (dropdown list)	Descriptor Summary in <i>File-related Long Running Utilities for Open Systems</i>

Report (File Level)

DBA Workbench	Adabas Manager	Documentation
Display FDT	DB nnn > Database Files > <i>Fields</i>	Field Definition Table (FDT) in <i>Database Files</i>
File Information		
Descriptor	DB nnn > Database Files > Utilities > Display Descriptors (dropdown list)	Display Descriptors in <i>File-related Long Running Utilities for Open Systems</i>
Usage	DB nnn > System Status > File Usage	Reviewing File Usage in <i>System Status</i>

I Working with Adabas Manager

This chapter explains the layout of the browser window through which System Management Hub and Adabas Manager are viewed. It takes you from system login to the main window for Adabas Manager.

The information is organized under the following headings:

Accessing the Adabas Manager Area

Adabas Manager Window Layout

Adabas Manager Configuration

Adabas Manager Session and Cache Administration

Trace Version

Leaving Adabas Manager

2 Accessing the Adabas Manager Area

Logging in to Adabas Manager involves first logging on to System Management Hub, then logging in to a particular host, and then selecting the Adabas Manager application.

▶ **To access the Adabas Manager area of System Management Hub:**

- 1 Select **Software AG Adabas Manager 7.7.3 -> Web Interface** from the program list in your start menu.

Or:

From a web browser, logon to System Management Hub at <http://<hostname>:49981/smh/login.htm>.

The logon screen for the System Management Hub (SMH) appears.

- 2 Logon to the System Management Hub...

The System Management Hub main panel appears.

- 3 Select the name of the managed host on which Adabas Manager is installed.
- 4 Expand the tree-view frame for the managed host by clicking on the plus sign (+) to the left of its name.
- 5 Expand the tree-view for Adabas Manager.


Deactivating/Activating Components

When you select Adabas Manager in tree-view, the Welcome to Software AG's Adabas Manager panel is displayed in detail-view. This table lists the components which have been installed with or are available to Adabas Manager. You can deactivate functionality you do not require by clicking **Deactivate**; you can reactivate the functionality any time by clicking **Activate**.



Important: When Adabas Manager is started for the the first time after installation, the system checks for the required version of Adabas for Open Systems. If Adabas for Open Systems has not been installed, the associated Adabas Manager functionality is automatically *deactivated*. Once Adabas for Open Systems has been installed, this functionality must be reactivated *manually*.

Documentation

The Adabas Manager documentation has been installed with the product. You can either view it by selecting **Software AG Adabas Manager 7.7.3 -> Documentation** from your program list, or call up the appropriate page by clicking **Help** () in the top right corner of your Adabas Manager panel.

3 Adabas Manager Window Layout

Each Adabas Manager window is divided into two scrollable parts:

- a tree-view on the left side of the screen
- detail information (detail-view) on the right side of the screen

You can navigate in the tree-view as you would in any directory tree. When you highlight (click) an item in the tree, the detailed information about that entry is displayed in detail-view. To expand items under a high-level entry in the tree, click the plus sign (+) to the left of the high-level entry.

Adabas databases on mainframes and on open systems platforms of any supported version are visible in the same tree-view. By selecting **Adabas Databases** in the tree-view, the number of databases of each type is displayed in the detail-view.

 **Important:** When Adabas Manager is started for the the first time after installation, the system checks for the required version of Adabas for Open Systems. If Adabas for Open Systems has not been installed, the associated Adabas Manager functionality is automatically *deactivated*. Once Adabas for Open Systems has been installed, this functionality must be reactivated *manually*.

The databases in the view can be grouped by by platform or separate sections and sorted by ID or name.

▶ **To group or sort your Adabas databases:**

- 1 Select the Adabas Manager item in the tree-view and expand it
- 2 Right-click **Adabas Databases** on the drop-down menu and select **Sort Order**.
- 3 In the detail-view panel, mark your preferences and click **OK** to save the grouping/sorting order or **Cancel**.

The tree-view is stored separately for each individual user.



Notes:

1. Tree items you are not authorized to view for security reasons are not displayed; where menus are available for viewing but may not be modified, the respective buttons are grayed out in the detail-view panels.
2. If Adabas Caching and Adabas Delta Save Facility are not installed on your system, their icons will not be displayed in your Adabas Manager tree-view.

Once you have highlighted an item in the tree, context-specific commands are listed in the drop-down menu.

You can personalize your work area by adding views to or removing them from your **System Management View**.

▶ **To personalize your work area:**

- 1 Right-click the item in the tree you want to add to the **System Management View** and select **Add to View** in the drop-down menu.

The item is added to your **System Management View**.

- 2 Click **System Management View** at the bottom of the tree-view area. The tree-view will collapse to show only the items you have added to your View.
- 3 To remove an item from your View, right-click it in your **System Management View** and select **Remove from View**.

The item is removed from the view.



Note: To return to the general work area, click **System Management** at the top of the tree-view.

All items in the tree show a **Refresh** option in their drop-down list.

▶ **To refresh any item in the tree:**

- 1 Select the item in the tree for which you want to view information and right-click it.
- 2 Select **Refresh** on the drop-down menu.

The information displayed in the detail-view is updated, i.e., you return to the main overview display of a function.

4 Adabas Manager Configuration

The Adabas Manager **Configuration** panel allows you to enter your license key for the Adabas Manager Enterprise Edition on mainframes and deactivate/activate the license for your version of Adabas on Open Systems; you also have to define the location for large output files of Adabas Manager's *Long Running Utilities*.

▶ **To view the Adabas Manager Configuration panel:**

- 1 Select **Adabas Manager** in tree-view and right-click it.
- 2 Select **Configuration** in the menu.

The Adabas Manager **Configuration** panel is displayed.

▶ **To enter your Adabas Manager mainframe license key file:**

- 1 Display the Adabas Manager **Configuration** panel.

The current license file name is displayed. If the license key field is empty, no license file is active and Adabas Manager is running in Standard Edition mode.



Note: Click **Display** in the **Configuration** panel to view a current license file text.

- 2 Click **Modify** to browse to the new license file and to activate it.
- 3 After selecting the new license key file click **OK** to activate it or **Cancel** to return to the **Configuration** panel.

▶ **To deactivate your Adabas for Open Systems license:**

- 1 Display the Adabas Manager **Configuration** panel.
- 2 Press **Deactivate** for the version you wish to deactivate.

The **Deactivate License** panel is displayed.



Caution: If you proceed with the deactivation, the license key will be deleted; before you confirm the deactivation, ensure that you have a copy of the license key in a safe location or you will not be able to activate your version of Adabas for Open Systems again.

- 3 Click **OK** or **Cancel** to return to the **Configuration** panel.

▶ **To activate your Adabas for Open Systems license:**

- 1 Display the Adabas Manager **Configuration** panel.
- 2 Press **Activate** for the version you wish to activate.

The **Activate License** panel is displayed.

- 3 Browse to the location where your license is stored.
- 4 Click **OK** or **Cancel** to return to the **Configuration** panel.

The System Management Hub job monitor receives the output of Adabas Manager's *Long Running Utilities*. In case of very large reports, Adabas Manager only sends the first and last lines to the job monitor; the complete output is written to the location defined in this panel. You can also modify the number of lines that should be displayed by the job monitor.

▶ **To modify the Job Monitor parameters:**

- 1 Display the Adabas Manager **Configuration** panel.
- 2 Press **Modify** and browse to a location where you want the output of the long running utilities to be written.
- 3 Enter the number of lines to be displayed by the job monitor (500 is the maximum value).
- 4 Click **OK** or **Cancel** to return to the **Configuration** panel.

5

Adabas Manager Session and Cache Administration

With Adabas Manager, several users can work independently and in parallel with all registered databases on a managed host. Adabas Manager stores frequently used data such as for example database GCB (general control block) information in the main memory of the Windows machine which displays the host (CS layer server).

The caching for each registered user (local objects) and for user-independent (global) objects is kept separate. A local object for example would be information on files that are "under construction", i.e., information from an incomplete file creation process, whereas GCB information of the Adabas databases is viewed as global objects.

After logout or after the session has been ended by closing the browser window, cached objects are no longer available.

To avoid a cache overflow, unused objects are deleted from the cache after a set time limit. A local object is regarded as *unused* if its user has been inactive for longer than 3600 seconds; a global object is considered *unused* if its period of inactivity is longer than 7200 seconds.

The cache status can be reviewed and cleared manually.

▶ **To display the cache status of an Adabas Manager session:**

- 1 Select the Adabas Manager item in the tree-view and right-click it.
- 2 Select **Display Cache** on the drop-down menu.

The cache information is displayed in the detail-view.

▶ **To clear the cache of an Adabas Manager session:**


- 1 Select the Adabas Manager item in the tree-view and right-click it.
- 2 Select **Clear Cache** on the drop-down menu.

The cache for all users using the same C/S layer server is cleared and all SAF-relevant data is lost. This means that after a cache clear, all users sharing a CS/layer server must logon again.

Internal settings of the cache and timeout limits can be set in the *HKLM/SOFTWARE/Software AG/System Management Hub/Products/Adabas Manager/Versions/7.7.3/* registry; the following properties may be modified:

Sess_Mem_Block	cache size in bytes	minimum: 1000000
Sess_Mem_Max_Global_Idle	maximum lifetime of global objects in seconds	minimum: 3600
Sess_Mem_Max_Local_Idle	maximum lifetime of local objects in seconds	minimum: 3600

After modification of a setting in the registry, the cache must be cleared with **Clear Cache**; a new cache is set up automatically.

-  **Caution:** Please note that serious problems might occur if you modify the registry incorrectly; make sure to exercise great care when working in the registry and only modify the settings listed in this table.

6 Trace Version

In case of a problem it may become necessary to produce a trace output for the Software AG support team. The trace output is created during normal Adabas Manager function:

- On *UNIX platforms* it is written to the folder **trace** in the Adabas Manager installation directory.
- On *Windows platforms* it is written to the folder **trace** in the application data directory, which contains application data for all users. Typical paths are:
 - C:\Documents and Settings\All Users\Application Data\Software AG\Adabas Manager 7.7.3\trace
 - C:\ProgramData\Software AG\Adabas Manager 7.7.3\trace



Note: The trace mode can only be activated if the trace version of the binaries was installed during the installation of Adabas Manager.

▶ To activate the trace mode:

- 1 Expand the tree-view frame for the managed host by clicking on the plus sign (+) to the left of its name.
- 2 Select Adabas Manager in tree-view and right-click it.
- 3 Select **Activate Trace Version** on the drop-down menu.
- 4 Carry out the function which causes the problem.

A trace report is written , which you can supply to your Software AG support team.

▶ To deactivate the trace mode:

- 1 Expand the tree-view frame for the managed host by clicking on the plus sign (+) to the left of its name.
- 2 Select Adabas Manager in tree-view and right-click it.

- 3 Select **Deactivate Trace Version** on the drop-down menu.

The trace version is deactivated.

7 Leaving Adabas Manager

▶ **To leave Adabas Manager:**

- Log off from the System Management Hub.

Or:

Close the Browser window.

The System Management Hub window and the Adabas Manager area are closed.

For new sessions with SAF-protected nuclei the user must logon again.

II Adabas Manager Security

Adabas Manager version 7.7 provides a security system to control the resources available to the users who administrate the Adabas databases:

- access control by user ID and password (authentication)
- access protection of the individual objects of an Adabas database (authorization)

Selecting **Adabas Databases** in the Adabas Manager tree will display an **Overview of Adabas Databases** in detail-view. The available database names and IDs are listed with their Adabas version and the information where logon is required.



Note: For protected databases, tree items you are not authorized to view are not displayed; where menus are available for viewing but may not be modified, the respective buttons are grayed out or not available (depending on the browser you use) in the detail-view panels.

Adabas Manager also supports simultaneous logon to more than one nucleus, which is useful in configurations where several nuclei run on the same server. It transfers user IDs and passwords (including new passwords) to all control instances. Adabas Manager neither restricts nor extends the user data rules which apply in the data center.

You can logon to Adabas databases by marking the appropriate check boxes in the table and clicking **Login**. For detailed information on how to logon to Adabas databases with Adabas Manager, see [Selecting a Database](#).

8 SAF Security Support

Adabas Manager can administrate a nucleus protected by Adabas SAF Security and processes for example the RACF profiles defined for Adabas Online System.

For a detailed table of the Adabas Manager tree items and the RACF/SAF resources used for their protection, see *RACF and SAF Resources*.

9 Working with a Protected Nucleus

- Accessing a Protected Nucleus 48
- Terminating a Session with a Protected Nucleus 28
- Passwords 29
- Non-activity Timeout 29

Accessing a Protected Nucleus

The user has to supply user ID and password when attempting to access a protected nucleus. The ID and password supplied for System Management Hub (SMH) are not applied here.

If the nucleus is protected by Adabas SAF Security, Adabas Manager will verify the user ID and password defined in the mainframe security system.

You are prompted for login

- if you are accessing the protected nucleus for the first time in the session
- after a non-activity timeout
- if the SMH session has been terminated (by logoff, timeout or network failure)



Note: Tree items you are not authorized to view for security reasons are not displayed; where menus are available for viewing but may not be modified, the respective buttons are grayed out in the detail-view panels.

Terminating a Session with a Protected Nucleus

The session with a protected nucleus is terminated

- when the browser window is closed
- when another URL is entered in the browser
- when the System Management Hub session is closed
- when the System Management Hub non-activity timeout limit is reached
- when the Adabas Manager non-activity timeout limit is reached
- by clearing the Adabas Manager session cache
- by logout
- by a new login

Passwords

Adabas Manager does its own user authentication check independent of the credentials supplied for System Management Hub (SMH) logon. The authentication is specific for each nucleus, but you can simultaneously logon to several with one password.

On expiry of a user password, the databases protected by Adabas SAF Security will request that you set a new password. User IDs and passwords in Adabas Manager are restricted to 8 characters. Other than that, Adabas Manager will check that a new password is being entered and that it is not the same as the previous one.

Non-activity Timeout

The duration of an Adabas Manager session with a protected nucleus can be determined by setting the `non_activity_time` value in the registry; this will be the priority timeout value and apply to all databases.

If `non_activity_time` has not been set in the registry, the ADARUN parameter `TNAE` will set the timeout for the session with a database.

III Adabas Databases

This chapter introduces the Adabas Manager database services, i.e., listing databases, registering and unregistering mainframe databases, adding and existing or creating a new open systems database, and selection of databases.

Adabas Manager can administer the following databases:

Mainframe

The remote databases must be UES-enabled and accessible by Entire Net-Work Client and Entire Net-Work Administration. Target databases on different mainframe servers are displayed under the same managed host in the System Management Hub environment independently of their operating system.

A maximum of 512 mainframe databases can be registered with Adabas Manager.

Open Systems

Only local open systems databases are displayed under the same managed host. For open systems databases with different operating systems, System Management Hub and Adabas Manager must be installed on all hosts holding the databases:

1. Install System Management Hub and Adabas Manager on the hosts.
2. Add the host by right-clicking **Managed Hosts** and choosing **Add Host** in the dropdown menu.

Open systems databases are always displayed under their host (they cannot be grouped under a single **Adabas Databases** entry).

The information is organized under the following headings:

[Listing Databases](#)

[Starting and Stopping Databases](#)

[Mainframe Databases that Cannot Be Reached](#)

[Selecting Databases](#)

[Reviewing Adabas Database DIBs](#)

[Reviewing Adabas Open Systems Database UCBs](#)

[Database-related Long Running Utilities for Open Systems](#)

10 Listing Databases

- Registering or Unregistering Mainframe Databases 35
- Adding Existing Open Systems Databases 36
- Creating New Open Systems Databases 37
- Deleting Open Systems Databases 39

▶ **To list the databases that are available to Adabas Manager:**

- Once you have accessed the Adabas Manager area of SMH, expand the area and then select **Adabas Databases**.

If you have not grouped your databases, the **Overview of Adabas Databases** in detail-view lists *all* databases that can be administered with Adabas Manager (mainframe and open systems databases).

Or:

If you have grouped your databases by section, select and expand either **Mainframe** or **Open Systems**.

The **Overview of Adabas Databases** panel is displayed in detail-view; the table shows database IDs, names, platforms and Adabas versions.

Database entries listed in the tree-view and in the overview tables are preceded by a stoplight icon:

- Green indicates that the database is active.
- Yellow indicates that this *mainframe* database is protected.
- Red indicates that the database is inactive.

Protected Adabas mainframe databases can be accessed from both these tables; mark the checkbox for the database, enter your user ID and password in the resulting **Security Login** window, and click **Login**, or **Cancel** to return to the overview of databases display.

In the table displaying Adabas open systems databases only, you can customize the information in the table by right-clicking on any of the table rows; the drop-down menu lets you select which columns you want to display and you can also present the information in different charts (bar, pie, line and points).

Nucleus Log

For more information on the open systems database you have chosen to work with, the contents of the nucleus log file of the last session of the selected database can be displayed. This contains information about the settings of the nucleus parameters when the current session was started, together with the version of the nucleus, and the number of the current or most recent session. The nucleus log file must exist.

▶ **To display the nucleus log of a database:**

- Right-click on the database name in tree-view and choose Nucleus Log.

The nucleus log information is displayed in detail-view.

The information in this chapter is organized under the following headings:

Registering or Unregistering Mainframe Databases

Mainframe databases must be registered with Adabas Manager. You can register a new database or unregister a currently registered database.

A maximum of 512 mainframe databases can be registered with Adabas Manager.

The information is organized under the following headings:

Registering Databases

▶ To register a database to Adabas Manager:

- 1 Right-click **Adabas Databases** in tree-view.



Note: If you have grouped your databases by section, you can also right-click **Mainframe**.

- 2 Select **Register DB** on the drop-down menu.

The **Database Registration** panel appears in detail-view.

- 3 Enter the ID and name of the database to be registered.
- 4 Click **Register** to complete the registration, or **Cancel** to quit the registration window without registering the database. Click **Help** for more information about registering a database.



Note: If the database is currently unreachable, you will be prompted to verify that you still want to register it.

If you clicked **Register**, a message is eventually returned indicating that the database has been successfully registered. The database appears as an option in tree-view.

- 5 Click **Close** in detail-view to complete the transaction.

Unregistering Databases

► **To unregister a database to Adabas Manager:**

- 1 Select an Adabas database in tree-view.
- 2 Right-click the database name in tree-view.
- 3 If you are unregistering an Adabas database, select **Unregister DB** on the drop-down menu.

The **Unregistration of Database** panel appears in detail-view.

- 4 Click **Unregister** to complete the unregistration, or **Cancel** to quit the unregistration window without unregistering the database. Click **Help** for more information about unregistering a database.

A message is returned indicating that the database has been successfully unregistered. The database is removed as an option in tree-view.

- 5 Click **Close** in detail-view to complete the transaction.

Adding Existing Open Systems Databases

Existing open systems databases can be made available to Adabas Manager.

► **To add an existing open systems database:**

- 1 Right-click **Adabas Databases** in tree-view.



Note: If you have grouped your databases by section, you can also right-click **Open Systems**.

- 2 Select **Add Existing DB** on the drop-down menu.

The **Add Existing Database** panel appears in detail-view.

- 3 Enter the DBID of the database you wish to add.
- 4 Click **OK** to add the database, or **Cancel** to quit without adding the database.

The job status is displayed in detail-view and will inform you once your database has been created successfully.

- 5 Click **Back** to return to the **Overview of Adabas Databases - Open Systems**.




Note: You can change the name of a database by right-clicking it in tree-view and choosing **Rename Database**. Enter the new database name in the resulting **Rename Database** panel in detail-view and confirm.

Creating New Open Systems Databases

▶ **To create a new open systems database:**


- 1 Right-click **Adabas Databases** in tree-view.

 **Note:** If you have grouped your databases by section, you can also right-click **Open Systems**.

- 2 Select **Create DB** on the drop-down menu.

The **Create New Database** panel appears in detail-view.

- 3 Enter the database specifications according to your requirements.

 **Note:** For information on open systems databases see section *Database Design* in the *Adabas for Open Systems* documentation.

1. Enter the DBID number in the **DBID** box; if this DBID already exists an error message will be displayed. A valid DBID is any number between 1 and 255 that has not already been assigned to another database.
2. Enter the name of the database in the **Name** box. A maximum of 16 characters may be specified; the database name must not contain blanks or special characters. Alternatively, you may accept the default name.
3. If different versions of Adabas are installed on your system, you may choose the version with which you want to create the database; by default the highest installed version of Adabas is selected.

4. Database Files

The default values of the database containers (Block size, Container size and Location) are displayed. Change the values according to your requirements.

The default block size that is to be used for the Associator container file (ASSO) is displayed in the ASSO1 **ASSO1 Blocksize** box. If you do not specify a block size, the default will be used.

 **Notes:**

- a. Ensure that the ASSO block size is smaller than the WORK block size.
- b. The ASSO2 container is generated by default to support Large Descriptor Values (i.e. Descriptor Values up to 1144 bytes), which requires a block size greater or equal to 16 KB. If you do not plan to use Large Descriptor Values, set either the ASSO2 block size or the ASSO2 container size to 0 and only the ASSO1 container will be generated. If you

later plan to use Large Descriptor Values a new container with larger block size can be added.

The amount of space to be assigned to the Associator container file is displayed in the **ASSO Size** box. You can accept the default value or change it according to your requirements.

The full specification of the Associator container file is displayed in the **Location** box. Either leave the specification displayed in the box unchanged if you want to accept it, or enter the specification directly in the box, or browse to another file.

5. The default block size that is to be used for the Data Storage container file (DATA) is displayed in the **DATA1 Blocksize** text box. If you do not specify a block size, the default will be used.

The amount of space to be assigned to the Data Storage container file is displayed in the **DATA1 Size** box.

The full specification of the Data Storage container file is displayed in the **Location** box. Either leave the specification displayed in the box unchanged if you want to accept it, or enter the specification directly in the box, or use browse to another file.

6. The default block size that is to be used for the WORK container file is displayed in the **WORK1 Blocksize** box. If you do not specify a block size, the default will be used.

The amount of space to be assigned to the WORK container file is displayed in the **WORK1 Size** box.

The full specification of the WORK Storage container file is displayed in the **Location** box. Either leave the specification displayed in the box unchanged if you want to accept it, or enter the specification directly in the box, or use browse to another file.

7. Miscellaneous

If the **Load Demofiles** box is checked there will be five demo files loaded into the new database:

File 9: Employees

File 11: EMP-NAT

File 12: Vehicles

File 13: Misc File

14: Lobfile of 9

The system files (checkpoint file, security file, and user data file) are automatically loaded with the default file numbers (checkpoint=1, security=2, user data=3) when you create a new database. Change the value in each box if you want to use other file numbers.

- 4 Click **OK** to create the database, or **Cancel** to quit without creating the database.

The job status is displayed in detail-view and will inform you once your database has been created successfully.

- 5 Click **Back** to return to the **Overview of Adabas Databases - Open Systems**.

Deleting Open Systems Databases

▶ To delete an Adabas open systems database from Adabas Manager:

- 1 Select an Adabas database in tree-view.
- 2 Right-click the database name in tree-view.
- 3 Select **Delete Database** on the drop-down menu.



Note: Only inactive databases can be deleted. Once a database has been deleted, it cannot be recovered.

The **Delete Database** *name* is displayed in detail-view and you are requested to confirm your delete action.

- 4 Click **OK** to confirm and return to the **Overview of Adabas Databases - Open Systems** display, or **Cancel** to return to the overview display of your database.

11 Starting and Stopping Databases

▶ To start an inactive Adabas open systems database:

- Right-click the name of the inactive database you want to start in tree-view and choose **Start Database** in the dropdown menu.

The detail-view panel displays a message that the database has been started successfully. Confirm and return to the **Database Overview** information.



Note: You can specify that the database is started when the operating system starts. To do this, right-click the database name and choose **Enable Start as Service** in the dropdown menu.

▶ To stop an Adabas database:

- 1 Select the Adabas database in tree-view.
- 2 Right-click the database name in tree-view.
- 3 Select **Stop Database** on the drop-down menu.

The session termination selection is displayed in detail-view. You have the following choices:

- Normal session termination (Shutdown)
- Cancel session immediately (Cancel)
- Stop session (Abort)



Important: Only open systems databases can be started with Adabas Manager; a main-frame database which has been stopped must be restarted by the DBA.

- 4 Choose how you want to stop the database by clicking in the appropriate check box for the user until a check mark appears and click **OK** to confirm.

Or:

Click **Cancel** to continue working with this database.

If you have chosen to stop the database, the shutdown information is displayed in detail-view and you are asked to confirm the shutdown.

- 5 Click **OK** to confirm the shutdown.

The database information is displayed in detail-view, with the target marked as inactive.



Note: If you have logged in to protected mainframe databases, the drop-down menu will offer you an additional menu item; **Logout All** terminate all sessions with protected databases.

12 Mainframe Databases that Cannot Be Reached

If you cannot access the database you have selected, it may be for one of the following reasons:

- the database is not running
- network problems prevent the database from being reached
- the database is not UES-enabled
- the database you have selected is not an Adabas mainframe database version 7.4.2 (or higher)
- mandatory ZAPs have not been applied to the database load library

In all of these cases a message is displayed with the information that the target cannot be reached.



Note: If you select a database which is not reachable due to network problems, the message "Target not responding (network timeout)" is displayed. Please ensure that all networks are functioning, i.e., take action to correct any possible errors. Clicking **Retry** will reach the database only if the error has been corrected.

13

Selecting Databases

- Selecting Mainframe Databases 46
- Selecting Open Systems Databases 49

When you select **Adabas Databases** to manage, you are first presented with a list of databases to choose from.

Database entries listed in the tree-view are preceded by a stoplight icon:

- Green indicates that the database is active.
- Yellow indicates that this mainframe database is protected or not supported.
- Red indicates that the database is inactive.

When **Adabas Databases** is highlighted in tree-view and you have not grouped your databases, a table is displayed in the detail window providing in addition to the stoplight icon, the database ID, database name, and Adabas version of each database registered to Adabas Manager, and for which mainframe databases a login with password is required.

If you have grouped your databases, click **Mainframe** or **Open Systems** to display the databases information.



Note: Protected mainframe databases can be accessed from both these tables; mark the checkbox for the database, enter your user ID and password in the resulting **Security Login** window, and click **Login**, or **Cancel** to return to the overview of databases display.

If a database is currently not reachable due to network problems, it will also be preceded by a red stoplight icon; the information "Target not responding" is displayed.

The information is organized under the following headings:

Selecting Mainframe Databases

► **To select an Adabas mainframe database to work with:**

- 1 Select and expand **Adabas Databases** in tree-view.
- 2 Select the database entry in tree-view under **Adabas Databases** If the mainframe database you want to access is preceded by a yellow stoplight sign, i.e., is protected, a Security Login window is displayed in detail view; enter your user ID and password and click **Login**, or **Cancel** to see the overview of databases display.



Note: If you have logged in to a protected mainframe database, the drop-down menu will offer you an additional menu item; **Display Authorizations** will display the authorizations you have with this particular database.

Once a database is selected, the following property default information about that database appears in detail-view:

Property	Description
Database ID	The number of the database.
Database Name	The name of the database.
Database Load Date	The date and time the database was loaded.
Maximum Number of Files	The maximum number of files permitted for the database.
Number of Files Loaded	The total number of files currently loaded in the database.
Lowest File Loaded	The lowest file number currently in use in the database.
Highest File Loaded	The highest file number currently in use in the database.
System File(s)	The number of Adabas system files allocated to the database.
Universal Encoding Support (UES)	Indicates whether Universal Encoding Support (UES) is active or inactive for the database.
Current Log Tape Number	The number of the most recent data protection log tape for the database.
RABN Size (RABNSIZE)	Indicates whether 3- or 4-byte RABNs are being used for the database.
Database Version	The version of the database.


Changing the Columns in the Database Overview Display

You can change the columns shown in the **Database Overview** display. Specifically, you can indicate whether the following information appears in the database overview display table:

Asso used	percentage of the Associator area used by the file
Data used	percentage of the Data storage area used by the file
Work used	percentage of the Work storage area used by the file
ASSO	Associator
DATA	Data storage
WORK	Work dataset
PLOG Statistics	Protection Log statistics
Logical Reads	Number of logical reads
Buffer Efficiency	Percentage of buffer efficiency
Format Translations	Number of format translations
Format Overwrites	Number of format overwrites
Throwbacks	Number of throwbacks
Thread Switches	Number of thread switches
Number of Calls	Number of calls
Internal Autorestarts	Number of internal autorestarts
PLOG Switches	Number of PLOG switches

Buffer Flushes	Number of buffer flushes
Attached Buffer	Percentage of attached buffer space used
Command Queue	Percentage of the command queue used
Format Pool	Percentage of the format pool space used
Hold Queue	Percentage of the hold queue space used
ISN-List Table	Percentage of the ISN list table queue space used
Seq. Cmd Table	Percentage of the seq. command queue space used
User Queue	Percentage of the user queue space used
Unique Descriptor Pool	Percentage of the unique descriptor pool queue space used
Security Pool	Percentage of the security pool space used
User File List Pool	Percentage of the user file list pool space used
DTP Transaction ID Pool	Percentage of the DTP transaction ID pool space used
Work Pool	Percentage of the work pool space used
Redo Pool	Percentage of the redo pool space used

► **To change the columns in the database overview display:**

- 1 In the upper left corner of the database overview display, click **Properties** (.

The **Database Overview** panel appears.

- 2 On the panel, select columns for the database overview display by moving them into and out of the **Selected Columns** list. To move a column into the **Selected Columns** list, select the column name in the **Available Columns** list and click the right arrow button. To move a column out of the **Selected Columns** list, select the column name in the **Selected Columns** list and then click the left arrow button.
- 3 Once you have selected the columns for the database overview display, select the order in which they should appear in the database overview display table. Use **Up** and **Down** to move the column headings up and down in the **Selected Columns** list.


A number of separators is available, allowing you to partition the tables as you wish.

- 4 When the columns you want are listed in the order you want in the **Selected Columns** list, click **Save** to save the column setting definition and return to the main database overview display, which will show the column settings you selected.

Or:

Click **Cancel** to cancel the column setting definition.

▶ **To restore the default database overview display:**

- 1 In the upper left corner of the database overview display, click **Properties** ().
The **Database Overview** panel appears.
- 2 Click **Restore Default**.
The non-default columns you have previously selected will move immediately from the **Selected Columns** list into the **Available Columns** list.
- 3 Click **Save** to save the column setting definition and return to the main database overview display, which will show the default column settings again.

Selecting Open Systems Databases

When you select **Adabas Databases** to manage, you are first presented with a list of databases to choose from.

Database entries listed in the tree-view are preceded by a stoplight icon:

- Green indicates that the database is active.
- Red indicates that the database is inactive.

When **Adabas Databases** is highlighted in tree-view and you have not grouped your databases, a table is displayed in the detail window providing in addition to the stoplight icon, the database ID, database name, and Adabas version of each database that can be administered with Adabas Manager, and for which databases a login with password is required.

If you have grouped your databases, click **Open Systems** to display the databases information.

▶ **To select an Adabas database to work with:**


- 1 Select and expand **Adabas Databases** in tree-view.
- 2 Select the database entry in tree-view under **Adabas Databases**.

Once a database is selected, the following property default information about that database appears in detail-view:

Property	Description
Database ID	The number of the database.
Database Name	The name of the database.
Database Load Date	The date and time the database was loaded.
Highest File Loaded	The highest file number currently in use in the database.
Lowest File Loaded	The lowest file number currently in use in the database.
Number of Files Loaded	The total number of files currently loaded in the database.
System File(s)	The number of Adabas system files allocated to the database.
Current PLOG	The number of the most recent data protection log for the database.

You can change the columns shown in the **Database Overview** display to include the database version and the maximum number of files.

▶ **To change the columns in the database overview display:**

- 1 In the upper left corner of the database overview display, click **Properties** ().

The **Database Overview** panel appears.

- 2 On the panel, select columns for the database overview display by moving them into and out of the **Selected Columns** list. To move a column into the **Selected Columns** list, select the column name in the **Available Columns** list and click the right arrow button. To move a column out of the **Selected Columns** list, select the column name in the **Selected Columns** list and then click the left arrow button.
- 3 Once you have selected the columns for the database overview display, select the order in which they should appear in the database overview display table. Use **Up** and **Down** to move the column headings up and down in the **Selected Columns** list.

A number of separators is available, allowing you to partition the tables as you wish.

- 4 When the columns you want are listed in the order you want in the **Selected Columns** list, click **Save** to save the column setting definition and return to the main database overview display, which will show the column settings you selected.

Or:

Click **Cancel** to cancel the column setting definition.

▶ **To restore the default database overview display:**

- 1 In the upper left corner of the database overview display, click **Properties** ().

The **Database Overview** panel appears.

2 Click **Restore Default**.

The non-default columns you have previously selected will move immediately from the **Selected Columns** list into the **Available Columns** list.

3 Click **Save** to save the column setting definition and return to the main database overview display, which will show the default column settings again.

14

Reviewing Adabas Mainframe Database DIBs

The data integrity block (DIB) comprises entries for each Adabas utility currently operating, describing the resources each utility is using.

You can review the DIB of an Adabas mainframe database as follows:

▶ **To review the DIB of an Adabas database:**

- 1 Select an Adabas database in tree-view.
- 2 Right-click the database name in tree-view.
- 3 Select **Display DIB** on the drop-down menu.

The DIB information appears in detail-view. It includes the following information about the Adabas database:

- the name of the service
- the name of the job
- the date and time the job was started
- the type of job
- the number of locked files in the database

An Adabas utility removes its entry from the DIB when it terminates normally. If a utility terminates abnormally (for example, the job is cancelled by the operator), the files used by that utility remain "in use". Any such files may be released with the **Reset DIB** function.

▶ **To reset the DIB of an Adabas database:**

- 1 Select an Adabas database in tree-view.

- 2 Right-click the database name in tree-view.
- 3 Select **Display DIB** on the drop-down menu.

The DIB information appears in detail-view.

- 4 Click **Reset DIB** to return to the DIB information panel.

When an ADASAV online dump operation has ABENDED, the online dump status can be reset.

▶ **To reset the online dump status:**

- 1 Select an Adabas database in tree-view.
- 2 Right-click the database name in tree-view.
- 3 Select **Display DIB** on the drop-down menu.

The DIB information appears in detail-view.

- 4 Click **Reset Online Dump Status** to return to the DIB information panel.

15

Reviewing Adabas Open Systems Database UCBs

The utility communication block (UCB) comprises entries for each Adabas utility currently operating, describing the resources each utility is using.

You can review the Utility Communication Block of Adabas open systems databases as follows:

▶ **To review the Utility Communication Block of an Adabas database:**

- 1 Select an Adabas database in tree-view.
- 2 Right-click the database name in tree-view.
- 3 Select **Utility Communication Block** on the drop-down menu.

The Utility Communication Block panel appears in detail-view. It includes the following information about the Adabas database:

- the entry ID
- the utility
- the mode
- the number of locked files

An Adabas utility removes its entry from the UCB when it terminates normally. If a utility terminates abnormally (for example, the job is cancelled by the operator), the files used by that utility remain "in use". Any such files may be released with the **Reset UCB** function.

▶ **To reset the UCB of an Adabas database:**

- 1 Select an Adabas database in tree-view.
- 2 Right-click the database name in tree-view.

- 3 Select **Utility Communication Block** on the drop-down menu.

The UCB information appears in detail-view.

- 4 Click **Reset UCB** to return to the UCB information panel.

16 Database-related Long Running Utilities for Open Systems

▪ Backup (ADABCK)	58
▪ Create WORK (ADD_CONTAINER)	59
▪ Create SORT (ADD_CONTAINER)	59
▪ Create TEMP (ADD_CONTAINER)	60
▪ Error Log (ADAERR)	60
▪ Display Descriptors (ADAFIN)	60
▪ File Usage (ADAFIN)	61
▪ List Protection Log (ADAPLP)	61
▪ Regenerate (ADAREC)	61
▪ Reorder (ADAORD)	62
▪ Restore (ADABCK)	63
▪ Unused Storage (ADAREP)	64
▪ Show Layout (ADAREP)	64
▪ Verify (ADAVFY)	64

Adabas Manager has two different ways of working with utilities:

- For utilities with a short processing time the job status is displayed in the detail-view. The user has to wait until the action is completed.
- For the so-called *Long Running Utilities* the job status is written to the System Management Hub Job Monitor. The utility runs in the background and leaves the user free to carry out other actions.

For detailed information on the SMH Job Monitor see *Monitoring Events and Jobs* in the *System Management Hub* documentation. For detailed information on Adabas utilities see the *Adabas for Open Systems Utilities* documentation.

The functions *Create Database* and *CLOG/PLOG Statistics* are designated long running utilities.

A number of other long running utilities have been grouped as **Utilities**.

▶ **To display database-related long-running utilities:**

- 1 Expand the open systems databases and click **Utilities** for the database you are working with.
The utilities are listed in detail-view.
- 2 Mark the utility you want to start and click **OK**.
- 3 In the detail-view panel, specify your options and the location of the output file and click **OK**

The utilities are described in more detail in the following sections:

Backup (ADABCK)

Options:

1. Allow parallel updates (DUMP)
Parallel updates are permitted on the files to be backed up while the backup is in progress. If this option is disabled, only ACC users (access only) can access the files to be backed up while the backup is in progress (EXU_DUMP).
2. Enable **Create New Protection Log** if you want to close the current protection log file and create a new log file after executing the backup.
3. If you have allowed parallel updates, enter a value to define the time (in seconds) that the system waits for ET-logic (end transaction) users to come to ET status at the end of the backup in the **Time To Wait For ET Status** field. If you do not enter a value, the value currently in effect for the database nucleus is taken.

Output:

1. Enable/Disable **Use Parallel Output Devices**.
2. The specification of the backup save set is displayed in the **File x** box. Either leave the specification displayed in the box unchanged if you want to accept it, or enter the specification directly. You can also browse to another file.
3. Enable **Create Two Physical Copies** check box if you want to create two copies of the save set.
4. Enable/Disable **Replace Existing Files**.
5. Click **OK** to start the utility or **Cancel**.

Create WORK (ADD_CONTAINER)

Options:

1. The default block size that is to be used for the new WORK container file is displayed in the **Block Size** box. To change it, enter the required value.
2. The amount of space to be assigned to the new WORK container file is displayed in the **Size** box. To change it, enter the required value.
3. Either leave the **Path** specification displayed in the box unchanged if you want to accept it, or enter the specification directly. You can also browse to another file.
4. Enable/Disable **Replace Existing Files**.
5. Click **OK** to start the utility or **Cancel**.

Create SORT (ADD_CONTAINER)

Options:

1. The amount of space to be assigned to the SORT container file is displayed in the **Size** box. To change it, enter the required value.
2. Either leave the **Path** specification displayed in the box unchanged if you want to accept it, or enter the specification directly. You can also browse to another file.
3. Enable/Disable **Replace Existing Files**.
4. Click **OK** to start the utility or **Cancel**.

Create TEMP (ADD_CONTAINER)

Options:

1. The amount of space to be assigned to the TEMP container file is displayed in the **Size** box. To change it, enter the required value.
2. Either leave the **Path** specification displayed in the box unchanged if you want to accept it, or enter the specification directly. You can also browse to another file.
3. Enable/Disable **Replace Existing Files**.
4. Click **OK** to start the utility or **Cancel**.

Error Log (ADAERR)

Option:

1. Either enter the specification for the error log input file directly or browse to another file.
2. Click **OK** to start the utility or **Cancel**.

Display Descriptors (ADAFIN)

Options:

1. Choose **All Descriptor** to create a report about all of the descriptors in all of the files in the database or enter the names of the descriptor you want to display.
2. Select the **Histogram** box to include a graphical overview of the descriptor-value lengths in the report.
3. Click **OK** to start the utility or **Cancel**.

File Usage (ADAFIN)

Create a report about the percentage to which the blocks in the database are used. The report also includes information about the number of allocated, used and free blocks, together with information about the length of the records in the files.

Option:

1. Select the data structure for which you want to create the report by selecting the appropriate check box.
2. Click **OK** to start the utility or **Cancel**.

List Protection Log (ADAPLP)

The function is used to list the contents of a protection log. The information displayed includes the protection log number, its creation date, information about the checkpoints, and the number of modifications for each file. The List Protection Log function also checks the internal structure of the protection log to ensure that it is consistent.

Option:

1. Either leave the specification displayed in the box unchanged if you want to accept it, or enter the specification directly. You can also browse to another file.
2. Click **OK** to start the utility or **Cancel**.

Regenerate (ADAREC)

The Regenerate function is used to recover a database. The database to be recovered must be active. Some functions however, (List, Close, Copy) can be executed when the database is inactive.

Data protection information, in the form of 'before', 'after' and 'delta' images of all updated records, is written to the protection log during each Adabas session provided that the nucleus parameter PLOG is set. This information is used to reapply the updates.

Protection Log:

1. The specification of the protection log file that is to be used for the regenerate is displayed in the **Input File** box. Either leave the specification displayed in the box unchanged if you want to accept it, or enter the specification directly. You can also browse to another file.

2. Either accept the number of the protection log that is displayed in the **Log Number** box, or enter the number of the protection log that is to be used to regenerate the database. Use **Find** to display information about the protection log if it is in a raw disk section.
3. Either leave the specification displayed in the **Output Errors File** box unchanged if you want to accept it, or enter the specification directly. You can also browse to another file.
4. Enable/Disable **Replace Existing Files**.

Action on Error:

1. Use the options **Exclude File with Error** and **Abort Regeneration** to select the action to be taken if non-fatal errors are detected during processing.

If **Exclude File with Error** is enabled, the file in question is excluded from the regenerate if Data Storage errors occur (nucleus response codes 17, 49, 75, 77 and 79).

If **Abort Regeneration** is enabled, the regenerate is abnormally terminated if errors are detected.

Miscellaneous:

1. Select **Check Before Image Consistency** to check the consistency of the before images in the protection log against the data in the database.

If you have not selected **Check Before Image Consistency**, it is strongly recommended that you select **Log Inconsistencies**. If any inconsistencies are detected, they will be written to the error log, which is then displayed automatically, or can be redisplayed using the **Error Log** function.



Important: If you do not select **Check Before Image Consistency** or **Log Inconsistencies**, the regenerate process may result in an inconsistent (i.e. invalid) database.

2. Click **OK** to start the utility or **Cancel**.

Reorder (ADAORD)

Reordering a database rearranges the database's global areas, and eliminates fragmentation of the files' Address Converter, Data Storage and Index extents by physically changing their placement. The Reorder function consists of implicit Export and Import functions.

Option:

1. The specification of the temporary data file (ORDEXP) that is used during the reorder is displayed in the **Export Data** box. Either leave the specification that is displayed in the box unchanged if you want to accept it, or enter the specification directly. You can also browse to another file.
2. Click **OK** to start the utility or **Cancel**.

Restore (ADABCK)

The Restore function is used to restore a database. The database to be restored must be inactive.

Options:

1. If **Allow for Allocation** is specified, the utility reallocates all files to be overlaid or the specified subset rather than attempting to restore them in the same block ranges as in the backup. Using this option reduces the number of file extents as much as possible.
2. If you want to restore the save set of one database into another database with a different DBID, select the **Ignore DBID in Backup File** check box, for example to restore the save set 'BCK001.004' (from the database with the DBID 4) into the database with the DBID 12.
3. Select the **Format Associator** and/or **Format Data Storage** check box if you want to format the blocks of the Associator and/or Data Storage. Formatting increases the time required for the restore to complete. If **Allow Relocation** is specified, ADABCK reallocates all files to be overlaid or the specified subset rather than attempting to restore them in the same block ranges as in the backup. Using this option reduces the number of file extents as much as possible.

Input:

1. The specification(s) of the save set(s) to be used for the restore is/are displayed in the **File x** box(es). Either leave the specification that is displayed in the box unchanged if you want to accept it, or enter the specification directly. You can also browse to another file.

Additional Options:

1. It is possible to see whether the save set can be used to restore a database by choosing **Read Check**. This function checks the readability (absence of parity errors) and completeness of the ADABAS backup copy.
2. You can display information about the backup save set by selecting **Report**.

Type:

The options **Contents**, **Summary**, and **Files** are available.

Contents displays a list of files (names and numbers) in an ADABAS backup save set, together with the name and number of the database.

Summary displays general information (name and number of the database, number of files loaded, maximum number of files, system file numbers), and the physical layout of the database.

Files displays status information for the files in an ADABAS backup save set.

3. Click **Restore** to start the utility or **Cancel**.



Note: An interrupted restore of a database must be rerun from the beginning.

Unused Storage (ADAREP)

This report shows the amount and location of unused space available for Associator and Data Storage.

- Mark **Unused Storage** and click **OK** to create the report.

Show Layout (ADAREP)

This report displays a summary of the use of the logical blocks in the ASSO and DATA, and reports lost blocks.

- Mark **Show Layout** and click **OK** to create the report.

Verify (ADAVFY)

Verify is used to verify the internal consistency of a database. The general control block (GCB) is validated together with each file control block (FCB) and each field definition table (FDT) of the loaded files. You can verify the Associator (Address Converter and Index) and/or the Data Storage.

Running the **Verify** function with an active database may lead to errors being reported; this is because updates may be made while the function is processing, and these updates will only be reflected in the nucleus buffer pool.

Verify Associator:

1. Select **File Control Block** to verify the file control blocks together with the field definition tables for all of the files in the database.
2. Select **Address Converter** to verify the Address Converter; this function validates from the Address Converter to the Data Storage, and checks to see that records can be found in the specified Data Storage.

Choose the extent of your output in the dropdown list.

3. Select the **Index** check box if you want to verify the index.

The dropdown list contains a list of the parts of the index for which information can be obtained; select the part that you want to be dumped.

4. The **RABN to dump from** boxes are used to specify a range of RABNs (Relative ADABAS Block Numbers) in the index. Only the RABNs specified will be dumped.

Verify Data:

1. Select **Data** to verify the Data Storage of the files in the database. This validates from the Address Converter to the Data Storage, and from the Data Storage to the Address Converter.

Record validates the Data Storage and check the structure of each record.

Choose the extent of your output in the dropdown list.

2. Enable **Field** to verify the record structure and validate the contents of packed, unpacked and floating point fields.

Miscellaneous:

1. Select **Search Lost Blocks** to search for lost RABNs in the database. The Dump Physical Structure option dumps the layout of the database.
2. Enter the maximum number of errors to be reported before the Verify function terminates.
3. Click **OK** to start the utility or **Cancel**.

IV

Checkpoints

17 Checkpoints

- Types of Adabas Users 70
- Checkpoint Names Starting with 'SYN' 70
- Checkpoint Numbers for 'SYN' Names (Mainframe Databases Only) 71

Types of Adabas Users

The following user types may appear in the list of checkpoints:

ACC	Access-only user
ET	ET logic user
EXF	Exclusive file control user or utility (privileged user)
EXU	Exclusive file update user or utility
UTI	Utility update control utility (privileged user)
UTS (mainframe only)	Online ADASAV SAVE file (privileged user)

Checkpoint Names Starting with 'SYN'

Checkpoint names starting with 'SYN' are reserved for the Adabas nucleus and utilities:

SYNC	A synchronized checkpoint made during nucleus initialization, including the status of the ADARUN IGNDIB and FORCE parameters.
SYNF	A checkpoint taken by a user program or utility that requires exclusive (EXF) control of one or more files.
SYNP	A checkpoint from a utility that requires privileged control. Such a utility can perform updating without using the Adabas nucleus.
SYNS (mainframe only)	A checkpoint from Adabas Online System (SYSAOS) or ADADBS with three exceptions from the nucleus. The function identified by this checkpoint is implemented without user intervention during regeneration. Exceptions include a second SYNS 5B recorded at the end of a nucleus session; SYNS 60 recorded at an interval specified by the ADARUN INTNAS parameter; and SYNS 61 recorded when more space is allocated for a file.
SYNV (mainframe only)	Indicates that a volume ID changed during sequential write to a dataset that is being closed.
SYNX	A checkpoint from a utility requiring exclusive control (EXU) of one or more files.
SYN1 (mainframe only)	A checkpoint made at the beginning of online ADASAV execution (SAVE database function).
SYN2 (mainframe only)	A checkpoint made at the end of online ADASAV execution (SAVE database function).
SYN4 (mainframe only)	A checkpoint made at the beginning of online ADASAV execution (SAVE files operation).
SYN5 (mainframe only)	A checkpoint made at the end of online ADASAV execution (SAVE files operation).

Checkpoint Numbers for 'SYN' Names (Mainframe Databases Only)

The following table describes the checkpoints written by an Adabas nucleus or utility:

Type	Name	Origin	Description
01	SYNC	ADANUC	Written by a nucleus at the start of a nucleus session
01	SYNF	user/utility	User/utility session OPEN with files used in EXF mode
01	SYNX	EXU user	Open checkpoint for an EXU user
02	SYNV	ADANUC	VOLSER entry; written at volume switch on DD/SIBA and at the end of the session if sequential logging is used
03	SYNF	user/utility	Close checkpoint for an EXF user
03	SYNX	EXU	Close checkpoint for an EXU user
05	SYNP	ADASAV	Start of SAVE file(s) operation
06	SYNP	ADASAV	Start of SAVE database operation
07	SYNP	ADASAV	End of RESTORE file(s) operation
08	SYNP	ADASAV	End of RESTPLOG operation
09	SYNV	ADASAV	SAVE file(s), VOLSER entry; written at volume switch on DD/SAVE or at the end of a SAVE operation
0A	SYNV	ADASAV	SAVE database, VOLSER entry; written at volume switch on DD/SAVE and at the end of a SAVE operation
0B	SYNP	ADASAV	End of SAVE DELTA operation
0C	SYNP	ADASAV	End of RESTORE operation
0D	SYNP	ADASAV	End of MERGE operation
0E	SYNV	ADASAV	SAVE DELTA, VOLSER entry
0F	SYNV	ADASAV	MERGE, VOLSER entry
10	SYNP	ADAINV	COUPLE files
11	SYNP	ADAINV	INVERT field(s)
15	SYNP	ADAORD	REORDER Associator database
16	SYNP	ADAORD	REORDER Data Storage database
17	SYNP	ADAORD	REORDER database
18	SYNP	ADAORD	REORDER Associator file
19	SYNP	ADAORD	REORDER Data Storage file
1A	SYNP	ADAORD	REORDER file
1B	SYNP	ADAORD	STORE
1C	SYNP	ADAORD	RESTRUCTURE
1D	SYNP	ADADEF	DEFINE NEWWORK
1E	SYNP	ADADEF	MODIFY default character encodings

Checkpoints

Type	Name	Origin	Description
22	SYNX	ADARES	REGENERATE file
23	SYNX	ADARES	BACKOUT file
24	SYNX	ADARES	REGENERATE all; CPEXLIST lists excluded files
25	SYNX	ADARES	BACKOUT all; CPEXLIST lists excluded files
26	SYNP	ADARES	REPAIR Data Storage
27	SYNV	ADARES	COPY sequential protection log
28	SYNP	ADARES	PLCOPYY function successfully completed
28	SYNV	ADARES	PLCOPYY dual protection log
29	SYNV	ADARES	CLCOPYY dual command log
2A	SYNP	ADARES	PLCOPYY MERGE function successfully completed
2A	SYNV	ADARES	PLCOPYY MERGE dual protection log
2B	SYNP	ADARES	CLOG MERGE function successfully completed
2B	SYNV	ADARES	CLOG MERGE dual command log
30	SYNP	ADALOD	LOAD file
31	SYNP	ADALOD	Mass update
35	SYNX	ADAULD	Unload file
3F	SYNP	ADAZAP	VERIFY - REPLACE successfully completed
40	SYNS	SYSAOS	Add extent
41	SYNS	SYSAOS	CHANGE default field length
42	SYNS	SYSAOS	DECREASE database size
44	SYNS	SYSAOS	Delete file
45	SYNS	SYSAOS	INCREASE database size
47	SYNS	SYSAOS	RECOVER space
48	SYNS	SYSAOS	Refresh file
49	SYNS	SYSAOS	Remove component file from expanded file chain
4A	SYNS	SYSAOS	Release descriptor
4B	SYNS	SYSAOS	RENAME file
4C	SYNS	SYSAOS	RENUMBER file
4D	SYNS	SYSAOS	RESET DIB
4E	SYNS	SYSAOS	Reuse ISN
4F	SYNS	SYSAOS	Reuse Data Storage
50	SYNS	SYSAOS	UNCOUPLE file
51	SYNS	SYSAOS	ALLOCATE file extent
52	SYNS	SYSAOS	DEALLOCATE file extent
53	SYNS	SYSAOS	Delete checkpoint
54	SYNS	SYSAOS	Set user priority

Type	Name	Origin	Description
55	SYNS	SYSAOS	Modify FCB
57	SYNS	SYSAOS	DEFINE file
58	SYNS	SYSAOS	Write FDT
59	SYNS	SYSAOS	DEFINE new field
5B	SYNS	ADADBS	Write refreshed statistics (some or all per user request)
5B	SYNS	ADANUC	Write (all) statistics at end of nucleus session
5B	SYNS	ADARES	Write refreshed statistics (command usage, file usage, DRES, DSTAT, thread usage)
5C	SYNS	SYSAOS	Change default field format
5D	SYNS	SYSAOS	Change file encoding
60	SYNS	ADANUC	Nucleus statistic checkpoint
61	SYNS	ADANUC	Allocate file space
64	SYNS	ADASCR	Protect files
65	SYNS	ADASCR	Protect fields
66	SYNS	SYSAOS	Link component file into expanded file chain
68	SYNS	SYSAOS	Set USERISN on/off.
69	SYNS	SYSAOS	Set MIXDSDEV on/off
6A	SYNS	SYSAOS	Install Delta Save DLOG area
6B	SYNS	SYSAOS	Change Delta Save DLOG area
6C	SYNS	SYSAOS	Remove Delta save DLOG area
6F	SYNS	SYSAOS	Online process initiated
70	SYNS	SYSAOS	Online invert process
71	SYNS	SYSAOS	Online reorder process
73	SYNC	ADANUC	Nucleus (nuclei) successfully quiesced
74	SYNC	ADANUC	Nucleus (nuclei) have resumed normal processing
75	SYNS	ADANUC	Delete heuri-user-entry after Response 72 was detected during nucleus startup
76	SYNS	ADANUC	Delete heuri-user-entry after Response 72 was detected during nucleus session
77	SYNS	ADADBS	Enable spanned record support
78	SYNS	ADADBS	Enable or disable extended MU or PE fields

V Database Files

The information is organized under the following headings:

Listing Database Files

Reviewing File Properties

Displaying Files with a Critical Number of Extents (Adabas for Mainframes Version 7 Only)

Files with Extent List Saturation (Adabas for Mainframes Version 8 Only)

Field Definition Table (FDT)

Modifying File Parameters

Allocating/Deallocating File Space

Modifying Files

Maintaining Expanded Files (Mainframes Only)

Creating New Files

Creating New LOB Files (Adabas for Mainframes Version 8 and Open Systems only)

File-related Long Running Utilities for Open Systems

18

Listing Database Files

- Displaying the List of Files 78
- Changing the Columns in the File List Display 78
- Displaying Locked Files 80

Displaying the List of Files

▶ **To list the files associated with a database:**

- 1 Select an Adabas database in tree-view and expand it.
- 2 Select **Database Files** in tree-view.

The files are listed in a file list display table in detail-view. By default, the following information about each file is provided:

- the file number
- the file name

Optionally, additional information about each file can also be listed, if you change the columns in the file list display, as described in *Changing the Columns in the File List Display*.

Links are also available for each file listed from which you can review the file properties and the Field Definition Table (FDT) for the file.



Note: To view a specific file which may be displayed on another page of the file list, enter the file number in the text box above the table and click **Go to**.

Changing the Columns in the File List Display

You can change the columns shown in the file list display. Specifically, you can indicate whether the following information appears in the file list display table:

- the percentage of the Associator area used by the file (Asso Used %)
- the percentage of the Data storage area used by the file (Data Used %)
- the lowest ISN that can be assigned to a record in the file (Min ISN)
- the highest ISN that can be assigned to a record in the file (Max ISN)
- the highest ISN currently assigned to a record in the file (Top ISN)
- the padding factor for the Associator (ASSO Padding Factor)
- the padding factor for Data Storage (DATA Padding Factor)
- the date the file was loaded
- the date of the last update
- the number of records loaded



Note: Depending on the number of records, the determination of this value may cause longer response times.

- the number of updates since the file was last loaded
- the extent list saturation (% , for Adabas for mainframes version 8 databases only)
- the number of logical extents currently assigned by Associator: normal index (NI Extent), upper index (UI Extent), address converter (AC Extent)
- the number of logical extents currently assigned by Data Storage (DATA Extent)

► **To change the columns in the file list display:**

- 1 Access the file list display, as described in *Displaying the List of Files*, earlier in this section.
- 2 In the upper left corner of the display, click **Properties** (.

The **File Properties** panel appears.

- 3 On the **File Properties** panel, select columns for the file list display by moving them into and out of the **Selected Columns** list. To move a column into the **Selected Columns** list, select the column name in the **Available Columns** list and click the right arrow button. To move a column out of the **Selected Columns** list, select the column name in the **Selected Columns** list and then click the left arrow button.
- 4 Once you have selected the columns for the file list display, select the order in which they should appear in the file list display table. The topmost column in the **Selected Columns** list will appear immediately to the right of the Fields column in the table. Use **Up** and **Down** to move the column headings up and down in the **Selected Columns** list.
- 5 When the columns you want are listed in the order you want in the **Selected Columns** list, click **Save** or **Save As** to save the column setting definition.

Or:

Click **Cancel** to cancel the column setting definition.

If you click **Save** or **Save As**, you are prompted for the unique name of the column setting definition. Continue with the next step.

If you click **Cancel**, the column setting definition is cancelled and you can ignore the remaining steps in this section.

- 6 Type in a unique name for the column setting definition and click **OK** to save the definition or **Cancel** to cancel the definition.
- 7 You can also rename or delete an existing column setting definition. Select the definition name in the selection box and click **Delete** or **Rename**.

If you click **Delete** the column setting definition will be deleted.

If you click **Rename**, you will be prompted to enter a new name for the definition. Enter the name you have chosen and click **OK**.

When you refresh the file list display, the column settings you selected will appear. The selected column settings appear in all file list displays for all databases registered until you change the column settings with a new definition.

If more than one column setting definition is created, you can select the one you want using the drop-down list on the upper right side of the file list display.

Displaying Locked Files

▶ To display locked files of a database:

- 1 Select an Adabas database in tree-view and expand it.
- 2 Right-click **Database Files** in tree-view and choose **Locked files** in the dropdown menu.

The locked files of the databases are displayed in detail.view. You can unlock a file by marking it and clicking **Unlock**.


19

Reviewing File Properties

▶ **To review the properties for a file:**

- 1 List the files in a database, as described in [Listing Database Files](#).
- 2 Select the name of a file in the table of files.



Note: The  button above the table in detail-view allows switching from regular layout of the tables to expert layout. Expert layout will also list Adabas nucleus parameters ADARUN for mainframe databases and ADANUC for open systems databases.

The properties for the file appears in detail-view. Please note that, depending on the file type, not all properties listed in the following table may be displayed:

Property	Description
Database ID (DBID)	The ID of the database in which the file resides.
File Number (FILE)	The number of the file.
File Name (NAME)	The name of the file.
Read Only	Shows whether the file is in read-only mode or not.
Records Loaded	The number of records currently contained in the file.
Number of Updates	The total number of blocks in the file that have been updated since the file was last loaded.
Commands during Nucleus Session	The number of operator commands entered during the session.
Date Loaded	The date and time the file was last loaded into the database.
Date of Last Update	The date and time the file was last updated.
Lowest ISN to be Allocated (MINISN)	The lowest ISN that can be assigned to a record in the file ("0" indicates that the parameter is not set; the default then is "1"). MINISN is required for expanded files.
TOP ISN	The highest ISN currently used in the file.

Property	Description
Highest ISN to be Allocated (MAXISN)	The highest ISN that can be assigned to a record in the file.
ASSO Padding Factor (ASSOPFAC)	The padding factor (percentage of each block) set for the ASSO dataset (the default is 10).
DATA Padding Factor (DATAPFAC)	The padding factor (percentage of each block) set for the DATA dataset (the default is 10).
ADAM File (ADAMFILE)	Whether the file is loaded using the ADAM option ("NO" indicates that the parameter is not specified).
Expanded File Component	Whether the file is part of an expanded file.
Ciphered File	Whether the file is encrypted (see the optional ADACMP COMPRESS CODE parameter; "NO" indicates that the parameter is not specified).
Coupled File	Whether the file is physically coupled to another file (see the ADAINV COUPLE function for more information).
DATA Space Reusage (DSREUSE)	Whether Data Storage space can be reused (the default is "YES").
Compress File Index (INDEXCOMPRESSION)	Whether the index for the file is loaded in compressed form (the default depends on the input file).
ISN Reusage (ISNREUSE)	Whether a freed ISN can be reused for a new record (the default is "NO").
Size of ISN (ISNSIZE)	Whether 3- or 4-byte ISNs are used for the file (the default is 3).
Owner ID Length for Multiclient File (LOWNERID)	The length of the internal owner ID for multiclient files ("0" indicates that it is not a multiclient file).
Maximum Blocks per DS Extent (MAXDS)	The number of blocks allowed per secondary Data Storage extent ("0" indicates that the parameter is not set; the default then is no limit).
Maximum Blocks per NI Extent (MAXNI)	The number of blocks allowed per secondary Associator normal index extent ("0" indicates that the parameter is not set; the default then is no limit).
Maximum Blocks per UI Extent (MAXUI)	The number of blocks allowed per secondary Associator upper index extent ("0" indicates that the parameter is not set; the default then is no limit).
Maximum Compressed Record Length (MAXRECL)	The maximum compressed record length (in bytes) allowed for the file (the default is the maximum possible size).
Mixed Device Types for DS Extents Allowed (MIXDSDEV)	Whether secondary Data Storage extents may be allocated on different device types and therefore with different block lengths ("NO" indicates that the parameter is not specified).
Record Spanning (Adabas version 8 and open systems only)	Whether record spanning is activated for the file.
MUPEX (Adabas version 8 only)	Whether the file has been established with the extended MU or PE limits available with Adabas version 8.

Property	Description
Has LOB Fields (Adabas version 8 and open systems only)	Indicates whether the file contains one or more LB fields.
LOB File (Adabas version 8 and open systems only)	Specifies the file number of the associated LOB file.
Base File (Adabas version 8 and open systems LOB file only)	Specifies the file number of the associated base file.
No AC Extensions (NOACEXTENSION)	Whether the MAXISN value for the file can be increased ("NO" indicates that the parameter is not specified; "YES" is required for expanded files).
Program Refresh Allowed (PGMREFRESH)	Whether a user program is allowed to refresh the file using an E1 call to Adabas when the file is loaded (the default is "NO").
Requires Universal Encoding Database	Whether the file has settings that require a UES-enabled database (see the ADACMP COMPRESS parameters FACODE, FWCODE, FUWCODE, UACODE, UWCODE, and UARC).
Alpha Num Encoding	
Wide Char Encoding	
User ISN (USERISN)	Whether ISN numbers are assigned by the user ("YES") rather than by Adabas ("NO"). The default depends on the input file.
Requires Adabas Version (or Higher)	The version of Adabas required to process the database and its files.
Highest Index Level	The highest index level currently active for the file.
FDT Length	The size (in blocks) of the file's field definition table (FDT).
Online Invert Running	Whether an online invert process is currently running in the background against this file. See the ADADBS ONLINVERT function for more information. Since such a background process runs as an ET logic job, 'yes' indicates that EXU/EXF users and utilities are prohibited from using the file, but other users can use it.

20

Displaying Files with a Critical Number of Extents (Adabas Version 7 only)

Adabas files can have five extents. You can display a list of the files in the database that are at or near the limit of five extents, i.e., near the "critical" number of extents.

▶ **To display files with a critical number of extents in a database:**

- 1 Select an Adabas database in tree-view and expand it.
- 2 Select **Database Files** in tree-view and right-click it.
- 3 Select **Files with a Critical Number of Extents** on the drop-down menu.

The **Files with a Critical Number of Extents** panel appears in detail-view.

A table displays the total number of files, the number of files that have already been processed, and once the operation has been completed, the number of files with critical extents.



Note: You can cancel the operation at any time by selecting **Cancel 'Files with a Critical Number of Extents'** on the drop-down menu.

The exact extent count is listed by table type (AC, DATA, NI, or UI).

If no files with four or five extents exist in the database, a message will be displayed in the **Files with a Critical Number of Extents** panel.

21 Files with Extent List Saturation (Adabas Version 8 only)

The limit of five logical file extents for each Adabas file extent type has been lifted. The maximum number of logical file extents that can be defined is limited only in that the extent information of all address converters, Data Storage, normal index, and upper index extents for the file must fit into the file control block (FCB). (The extent information is stored in a variable section of the FCB.) For example, on a standard 3390 device type, a file could have more than 40 extents of each type (or there could be more of one type if there are less for another). For more information about Adabas 8 logical extents, see *Adabas Logical Extents* in the *Adabas 8 for Mainframes* documentation.

For Adabas version 8 databases, the extent list saturation of the files can optionally be displayed in the **List Files** detail-view. For detailed information see [Changing the Columns in the File List Display](#) in *Listing Database Files*.

▶ **To view files with extent list saturation:**

- 1 Select an Adabas version 8 database in tree-view and expand it.
- 2 Select **Database Files** in tree-view and right-click it.
- 3 Select **Files with Extent List Saturation** on the drop-down menu.

The **Files with Extent List Saturation** panel appears in detail-view.



Note: You can cancel the operation at any time by selecting **Cancel 'Files with Extent List Saturation'** on the drop-down menu.

If no files with critical saturation exist in the database, a message will be displayed in the **Files with Extent List Saturation** panel.

22

Field Definition Table (FDT)

- Referential Integrity Constraints 90
- Changing Field Parameters 91
- Logically Deleting/Undeleting Fields 91
- Descriptors 92
- Predict Parameters for Mainframe Databases 94
- Adding New Fields in the Field Definition Table (FDT) 94

Referential Integrity Constraints

A referential constraint ensures referential integrity between two keys. Keys can be descriptors, superdescriptors or ISNs. Referential integrity means that for every value in a descriptor called “foreign key”, there must be a value in a descriptor called “primary key” in a primary file. The primary key must be defined as unique and the options NC and NN must be set. For the foreign key, the option NC must be set. A pair of primary and foreign key must have the same format. ISNs can be used as primary keys and the corresponding foreign key must be binary. If primary and foreign keys are superdescriptors, then

- The corresponding key must be a superdescriptor;
- The superdescriptors must have the same number of parent fields;
- The corresponding parent fields must have the same format;
- No parent field may occur twice in a superdescriptor;
- No parent field may occur in more than one foreign key;
- All parent fields must have the option NC;
- All parent fields of the primary key must have the option NN;
- The primary key superdescriptor must be unique.

You can specify a referential action that is executed on the foreign key record if a primary key value is modified.

The referential constraint is added to the file to which the foreign key belongs.



Note: The setting of referential constraints is a designated long running utility because all descriptor entries must be checked to ensure that they do not violate the new constraint.

See also *Referential Constraints in FDT Record Structure* in the *Adabas Administration* documentation for open systems.

▶ To add a constraint:

- 1 In the Field Definition Table (FDT), press **Add Constraint**.
- 2 In the resulting **Referential Constraints** panel, enter the required specifications and click **OK** to save your entries or **Cancel** to return to the FDT.

Changing Field Parameters

Note that Adabas for Open Systems now supports referential integrity constraints - this means that any field in a file which is defined as a foreign key can only contain values from a primary file's primary key. This prevents you from defining records in a file containing a foreign key value that is not contained in the corresponding primary file as a primary key value (for more information see [Referential Integrity Constraints](#) and also *Referential Constraints in FDT Record Structure* in the *Adabas Administration* documentation for open systems).

▶ **To change field parameters:**

- 1 In the Field Definition Table (FDT), click the name of the field you want to change the parameters for.

The **Change Field Parameters** panel is displayed.

- 2 To modify the standard length of an Adabas field, update the **Length** value as required.

For mainframe databases you can also select the **LA** option to change a normal alphanumeric (A) field to a long-alpha (LA) field. If the file is protected, supply the password.

- 3 Click **OK** to save your entries or **Cancel** to return to the file list display.

Logically Deleting/Undeleting Fields

With Adabas Manager version 7.7.3 you can logically delete and undelete a field from a database file. Logically deleting a field from a file removes the field from the FDT, but retains the field data in the database. Logically undeleting a field that was previously logically deleted reinstates the field in the FDT.



Note: Logically undeleting a field is only possible for mainframe databases.

▶ **To delete a field:**

- 1 In the Field Definition Table (FDT), click the name of the field you want to delete.

The **Change Field Parameters** panel is displayed.

- 2 Click **Delete Field** or **Cancel** to return to the file list display.

The field is displayed in the FDT with the option **deleted**.



Note: The parent field of a descriptor cannot be deleted; attempting to delete such a field will result in an error message.

Deleting a field in the Field Definition Table (FDT) only logically deletes the field; the field can be reactivated at any time.

▶ **To reactivate a logically deleted field (mainframe databases only):**

- 1 In the Field Definition Table (FDT), click the name of the deleted field you want to reactivate.

The **Change Field Parameters** panel is displayed.

- 2 Click **Undo Field Deletion** or **Cancel** to return to the file list display.

The field is displayed in the FDT without the option **deleted**.

Descriptors

Setting and Releasing Descriptors

Field names listed in the Field Definition Table (FDT) with an option "DE" are descriptors. You can remove a descriptor by freeing the specified field's inverted list in the Associator.

The functions **Release Descriptor**, **Set Unique Descriptor** and **Reset Unique Descriptor** are designated long running utilities.

▶ **To release a descriptor:**

- 1 In the Field Definition Table (FDT), click the name of the field you want to release the descriptor for.

Or:

If a Special Descriptor Table (SDT) is shown, click the icon in the **Parameters** column of the descriptor you want to release.

The **Change Field Parameters** panel is displayed.

- 2 Click **Release Descriptor**.
- 3 In the resulting panel, confirm that you want to release the descriptor by clicking on **OK**, or **Cancel**. Both options will return you to the FDT display.

▶ **To set a unique descriptor for open systems databases:**

- 1 In the Field Definition Table (FDT), click the name of the field you want to set the descriptor for.

Or:

If a Special Descriptor Table (SDT) is shown, click the icon in the **Parameters** column of the descriptor you want to set to unique.

The **Change Field Parameters** panel is displayed.

- 2 Click **Set Unique Descriptor**.
- 3 In the resulting panel, confirm that you want to set the descriptor by clicking on **OK**, or **Cancel**. Both options will return you to the FDT display.

▶ **To reset a unique descriptor for open systems databases:**

- 1 In the Field Definition Table (FDT), click the name of the field you want to reset the descriptor for.

Or:

If a Special Descriptor Table (SDT) is shown, click the icon in the **Parameters** column of the descriptor you want to reset.

The **Change Field Parameters** panel is displayed.

- 2 Click **Reset Unique Descriptor**.
- 3 In the resulting panel, confirm that you want to reset the descriptor by clicking on **OK**, or **Cancel**. Both options will return you to the FDT display.

Logically Deleting Descriptors (Adabas version 8.2 only)

Field names listed in the Field Definition Table (FDT) with an option "DE" are descriptors; this option can be logically removed from a field.

▶ **To logically delete a descriptor:**

- 1 In the Field Definition Table (FDT), click the name of the field you want to remove the descriptor from.

The **Change Field Parameters** panel is displayed.

- 2 Click **Logically Delete** to remove the descriptor and return to the FDT.

Predict Parameters for Mainframe Databases

The following Predict parameters must be specified in the FDT to enable the display of the Predict field name of a file.

▶ **To set the Predict parameters:**

- 1 In the Field Definition Table (FDT), click the name of the field you want to set the parameters for.

The **Change Field Parameters** panel is displayed.

- 2 In the upper left corner of the FDT display, click **Properties** .

The **Predict Parameters** panel appears.

- 3 Enter the following parameters:

FDIC DBID	The ID of the database containing the Predict repository.
FDIC FNR	The file number of the FDIC file.

- 4 Click **OK** to save your entries or **Cancel** to return to the file list display.

Adding New Fields in the Field Definition Table (FDT)

▶ **To add new fields in the Field Definition Table (FDT):**

- 1 In the Field Definition Table detail-view, click **Add Field**.

The **Add New Fields** panel for the file is displayed.

- 2 Specify the new fields you want to add to the file.
- 3 Click **New** to add more rows to the table if necessary.
- 4 Click **OK** to save your modifications or **Cancel** to return to the previous panel.

23

Modifying File Parameters

- Changing Padding Factors, Maximum Allocation, and Record Length 96
- Selecting File Options 97
- Record Spanning, MUPEX, LOB Fields/File (Adabas Version 8 and Open Systems Only) 97

► **To modify parameters for a file:**

- 1 List the files in a database, as described in *Listing Database Files*.
- 2 Select the file for which you want to modify the parameters by clicking on the filename.

The property table for the file appears in detail-view.

- 3 Click **Modify**.

A property table for the file appears in detail-view, which allows you to modify the padding factor, the maximum compressed record length, file number, file name, extent allocation for NI/UI/AC/DS, ISN reusage, and DS reusage.

- 4 Modify the parameter settings according to your requirements and click **OK** to save your modifications and return to the previous panel.

Or:

For Adabas for mainframes version 8 and/or open systems databases, the following additional parameters may be specified:

- Record Spanning
- MUPEX (Adabas for mainframes version 8 only)
- Has LOB Fields
- LOB File

For detailed information on record spanning, MUPEX, LOB fields and files, see the *Adabas for Mainframes* and the *Adabas for Open Systems* documentation.

Changing Padding Factors, Maximum Allocation, and Record Length

You can use the Modify File Parameters function to change the Associator and Data Storage padding factors for the file, the maximum number of blocks that can be allocated (for a new Data Storage, normal index, or upper index extent), and the maximum compressed record length allowed.

The "padding factor" is the percentage (%) of each Associator or Data Storage block that is reserved; that is, not loaded. This area is used to create new records later. The range is from 3 to 90 percent. The factor size allocated should depend on the amount of updating that is expected. The number of bytes left in the Associator after padding must exceed the largest descriptor value by at least 10.

Selecting File Options

With Adabas version 8.2, you have the option to put files in read-only state; if you want a file to be accessed only in read-only mode, mark the **Read Only** checkbox in this table. This is useful if you need to maintain this file while the rest of the database is up.

You can also turn off or on several file options in this table:

- Data Space Reusage ...with RESET / in Parallel
- ISN Reusage ...with RESET / in Parallel
- Mixed Device Types for DS Extents Allowed
- User Refresh Allowed
- User ISN

ISN Reusage and Data Storage Reusage determine whether ISNs and Data Storage blocks for deleted records are reused as new records are added to the file.

When setting either of these two options to "ON", you can also set the RESET option "ON" to start the search for an unused ISN and/or Data Storage block at the beginning of the file.

Record Spanning, MUPEX, LOB Fields/File (Adabas Version 8 and Open Systems Only)

Record Spanning

With Adabas 8 and Adabas for open systems, records can be spanned in a database. In the database, the logical record is split into a number of physical records, each part fitting into a single Data Storage (DS) block. The resulting physical records are each assigned individual ISNs. The first physical record is called the primary record and contains the beginning of the compressed record and is assigned a primary ISN. The remaining physical records are called secondary records and contain the rest of the data of the logical record. Secondary records are assigned secondary ISNs. These ISNs do not affect the user ISNs assigned when using the N2 command or the ISNs used when using the I option of the L1 command. If spanned records are used, a secondary address converter is used to map the secondary ISNs to the RABNs of the Data Storage blocks where the secondary records are stored.

A spanned record is comprised of one primary record and one or more secondary records. However, the number of segments in a spanned record is limited. The Adabas nucleus allows up to five physical records (one primary record and four secondary records) in a spanned record.

Spanned records are not directly visible to application programs. Applications always address spanned records via the primary ISN.

Spanned records are also supported in expanded Adabas files and in multi-client files.

▶ To count spanned records:

- 1 If the parameter has been specified, click **Record Spanning** in the file property table to count the spanned records.

The Count Spanned Records table is displayed in detail-view.

- 2 Set the maximum running time for the count (default is 60 seconds) and supply the password if required.
- 3 Click **OK** to start the count.

MUPEX (Adabas for Mainframes Version 8 Only)

With Adabas version 8 the number of occurrences of each MU field or each PE group in a record has been increased from 191 to about 65,534. However, the actual number of occurrences is limited to the size of the data-block size, the device type and the file type (spanned or not spanned). All MU fields and PE groups and other fields must fit into one compressed record. If you are using spanned records (introduced with Adabas 8), more MU fields and PE groups can be stored.

In addition, subdescriptors and superdescriptor definitions can affect the number of MU fields or PE groups in the record. For example, if a superdescriptor is created as a combination of a PE group and one or more MU fields and the number of occurrences is high, performance and resource problems can occur.



Note: Excessive use of extended MU and PE fields might cause performance and resource problems. These can result in a work storage overflow, resulting in Response code 9. If this should happen, increase the ADARUN LP size for the database.

The use of more than 191 MU or PE fields in a record must be explicitly allowed for a file (it is not allowed by default).

LOB Fields/Files

Adabas 8 introduces large object fields (LB fields). These are fields that may contain much more data than the 253 bytes of normal alphanumeric fields or the 16,381 bytes of LA fields. Fields containing large objects are defined with the new LB option. The theoretical maximum size of the value of an LB field is just short of 2 GB; practical usable sizes are smaller.

Adabas version 8 stores large object field values in a separate file, called a LOB file, that is tightly associated with the file containing the LB fields, which is called the base file. Behind the scenes, Adabas will store LB field values (except for very short ones) in the LOB file, but this is transparent

to your application. Commands in application programs should always be directed against the base file; application programs need neither know nor care about the existence of a LOB file.

For more information, see section [Creating a New LOB File](#), and *Getting Started with Large Object (LOB) Fields* in the *Adabas 8 for Mainframes* documentation and the *Adabas for Open Systems* documentation.

24

Allocating/Deallocating File Space

▶ **To allocate or deallocate space for a file:**

- 1 List the files in a database, as described in *Listing Database Files*.
- 2 If an icon appears in the **Alloc.** column of the list of files table in detail-view, click the icon.

The space allocation for the file appears in detail-view.

- 3 Click **Allocate** or **Deallocate**.

The Allocate Space or the Deallocate Space table is displayed in addition to the existing space allocation table.

- 4 You can create or deallocate extents for address converter, data storage, normal index and upper index of a file. You can specify the allocation in blocks or in cylinders, a starting relative Adabas block number (STARTRABN), and (if you are allocating extents) the device where allocated space is to be located. If the file is protected, supply the password.
- 5 Click **OK** to save your modifications or **Cancel** to return to the previous panel.

25

Modifying Files

- Locking/Unlocking Files 104
- Reordering Files 104
- Deleting Files 105
- Renumbering Files 105
- Renaming Files 105
- Uncoupling Files 106
- Inverting Files 106

Locking/Unlocking Files

▶ **To lock/unlock files:**

- 1 List the files in a database, as described in *Listing Database Files*.
- 2 The files are listed in a file list display table in detail-view; select the file you want to lock or unlock by clicking in the check box until a check mark appears.
- 3 Click **Lock** or **Unlock**.
- 4 Click **OK** to confirm or **Cancel**, and return to the file list.

Reordering Files

▶ **To reorder files:**

- 1 List the files in a database, as described in *Listing Database Files*.
- 2 The files are listed in a file list display table in detail-view; select the file you want to reorder by clicking in the check box until a check mark appears.
- 3 Click **Reorder**.
- 4 In the resulting **Reorder Files** panel, select the preferred reorder type in the selection box.

You can reorder:

- the entire file,
- the Associator for the file only,
- the Data Storage for the file only.

When reordering Data Storage for a file, you can specify a sort sequence. The default is physical sequence. Other possible options include "ISN" if the file is to be sorted in ISN order, or the two character descriptor name to sort the file according to the value of the specified descriptor.

- 5 Supply the password and click **OK** to confirm or **Cancel**, and return to the file list.

Deleting Files

▶ **To delete files:**

- 1 List the files in a database, as described in *Listing Database Files*.
- 2 The files are listed in a file list display table in detail-view; select the file you want to delete by clicking in the check box until a check mark appears.
- 3 Click **Delete**.
- 4 In the resulting **Delete Files** panel, supply the password and click **OK** to confirm or **Cancel**, and return to the file list.

If a file to be deleted is a coupled file, it must first be uncoupled.

Renumbering Files

▶ **To renumber files:**

- 1 List the files in a database, as described in *Listing Database Files*.
- 2 The files are listed in a file list display table in detail-view; select the file you want to renumber by clicking in the check box until a check mark appears.
- 3 Click **Renumber**.



Note: Password-protected files cannot be renumbered.

- 4 In the resulting **Renumber Files** panel, modify the file number according to your requirements and click **OK** to confirm or **Cancel**, and return to the file list.

Renaming Files

▶ **To rename files:**

- 1 List the files in a database, as described in *Listing Database Files*.
- 2 The files are listed in a file list display table in detail-view; select the file you want to delete by clicking in the check box until a check mark appears.
- 3 Click **Rename**.

- 4 In the resulting **Rename Files** panel and enter the new file name.



Note: Password-protected files cannot be renamed.

- 5 Click **OK** to confirm or **Cancel**, and return to the file list.

Uncoupling Files

The uncoupling option is used to remove the physical coupling between the specified files by erasing the coupling inverted lists from each file's Associator. No change is made to the field definition tables (FDTs) or descriptors for the specified files. This option must be executed before either of the specified files is deleted.

▶ To uncouple files:

- 1 List the files in a database, as described in *Listing Database Files*.
- 2 The files are listed in a file list display table in detail-view; select the file you want to delete by clicking in the check box until a check mark appears.
- 3 Click **Uncouple**.
- 4 In the resulting **Uncouple Files** panel, enter number of the file to be uncoupled and supply the password.
- 5 Click **OK** to confirm or **Cancel**, and return to the file list.

Inverting Files

▶ To invert files:

- 1 List the files in a database, as described in *Listing Database Files*.
- 2 The files are listed in a file list display table in detail-view; select the file you want to delete by clicking in the check box until a check mark appears.
- 3 Click **Invert**.
- 4 In the resulting **Invert Files** panel, enter the definition for the type of descriptor to be inverted and supply the password.

Examples:

```
FIELD=' AA '                HYPDE=' 01,HD,20,A,NU,MU=AA,AB '  
SUBDE=' SB=AA(1,5) '        PHONDE=' PH(AA) '  
SUPDE=' SP=AA(1,5),BB(1,2),CC(3,5) ' COLDE=' 1,CD=AA '
```

For detailed information on the syntax, see *ONLINVERT: Start Online Invert* in the *Adabas Utilities* documentation.



Note: Only one descriptor can be specified per process. Only one invert process can be active for a file at any time.

- 5 Click **OK** to confirm or **Cancel**, and return to the file list.

26

Maintaining Expanded Files (Mainframes Only)

This function corresponds to the utility functions ADALOD LOAD and ADALOD UPDATE.

▶ **To insert an existing file into an expanded file chain:**

- 1 Select an Adabas database in tree-view and expand it.
- 2 Select **Database Files** in tree-view and right-click it.
- 3 Select **Maintain Expanded Files** on the drop-down menu.

The **Expanded File Maintenance** panel appears in detail-view.

- 4 Specify the number of the file to be inserted.
- 5 Specify the number of the master file of the expanded file chain into which the file is being inserted.



Note: If the file to be inserted is to be the first (master) file of the chain, both **File Number** and **Master File Number** fields must be set to the same value.

- 6 Supply the password if required and click **Insert**.

A message will be displayed that the function has been completed successfully.

▶ **To remove a component file from an expanded file chain:**

- 1 Select an Adabas database in tree-view and expand it.
- 2 Select **Database Files** in tree-view and right-click it.
- 3 Select **Maintain Expanded Files** on the drop-down menu.

The **Expanded File Maintenance** panel appears in detail-view.

- 4 Specify only the number of the file to be removed from the chain.

- 5 Supply the password if required and click **Remove**.

A message will be displayed that the function has been completed successfully.

27

Creating New Files

The **New File** functionality of the Adabas Manager allows the creation of simple Field Definition Tables (FDT). There is no limitation in the number of fields you can create or to their complexity; however, for the creation of extensive and complex FDTs, the use of the ADACMP utility is strongly recommended (for details see the section *ADACMP: Compress-Decompress* in the *Adabas Utilities* documentation).

For open systems databases you also have the option to import an FDT (Page 2 of the **New File** process).

Once you have started the creation of a new FDT, its status will be stored even after a cancellation or during inactivity of the user until a timeout occurs (for timeout settings, see *Adabas Manager Session and Cache Administration*). It is recommended though that you carry out the file creation process without interruption; once you click the Adabas Manager entry in the tree-view or leave your session, any incomplete file creation process is terminated.

▶ **To create a new file for a database:**

- 1 Select an Adabas database in tree-view and expand it.
- 2 Select **Database Files** in tree-view and right-click it.
- 3 Select **New File** on the drop-down menu.

The **New File** panel appears in detail-view.



Note: During the following file creation procedure, clicking **Back** will return you to the previous page of the **New File** (i.e., for modifications or additional entries); **Cancel** will cancel the file creation and return to the file list display (after cancellation, the system will store the settings you have entered for the new file until you leave your session, a timeout occurs or you click the Adabas Manager entry in the tree-view).

- 4 Enter the following parameter values for the new file:

- File Number
- File Name (optional)
- MAXISN
- Data Storage Size
- Normal Index Size
- Upper Index Size



Note: If you click **Find** next to the **File Number** text box, the next free file number will be displayed. If you enter a file number that already exists and click **Find**, the next free file number after the number you entered will be displayed.

- 5 Click **Next** to continue with the file creation.

Page 2 of the **New File** appears in detail-view.

- 6 Define the File Descriptor Table for the new file. Note that for open systems databases you have the option to import an FDT by specifying the path to the file.

If you are not importing an FDT, complete the following entries:

- Level

a one- or two-digit number in the range 01-07 (the leading zero is optional) used in conjunction with field grouping.

- Name

the name to be assigned to the field (or group). Names must be unique within a file. The name must be two characters long: the first character must be alphabetic; the second character can be either alphabetic or numeric. No special characters are permitted.

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

Adabas version 6.2 for open systems also supports lower case field names; however, if the database you are creating the new file for is a version 6.1 database, any lower case entry will automatically be converted to upper case characters.

- Length

the length of the field (expressed in bytes).

- Select the format in the selection box.

Check the options you want to set; enter the position of the field if required.

Date-Time Option DT

Adabas versions 8.2 and 6.2 introduce a new DT field option which allows you to specify the edit mask for a field definition. This option DT assigns a date-time edit mask to a binary, fixed point, packed decimal, or unpacked decimal field. This option cannot be specified for fields of other formats.

The timestamp function is set by checking the TZ box and selecting the required value from the list.

Valid values are described in the following table. It also shows the required minimum field lengths for the different formats of fields that can specify the DT option; the length of the field must be large enough to store the date-time values.



Note: In the table, "YYYY" represents the 4-digit year (1-9999), "MM" represents the 2-digit month (1-12), "DD" represents the 2-digit day of the month (1-31), "HH" represents the 2-digit hour (0-23), "II" represents the 2-digit minute within the hour (0-59), "SS" represents the 2-digit second within the minute (0-59), and "XXXXXX" represents the 6-digit microsecond within the second.

	Description	Minimum Field Length for Field Format			
		B	F	P	U
DATE	The date field is in the format Z'YYYYMMDD'.	4	4	5	8
TIME	The time field is in the format Z'HIISS'.	3	4	4	6
DATETIME	The date and time field is in the format Z'YYYYMMDDHIISS'.	6	8	8	14
TIMESTAMP	The date and time field is in the format Z'YYYYMMDDHIISSXXXXXX', with microsecond precision.	9	-	11	20
NATTIME	The time field is in Natural T format (tenths of seconds since year zero).	6	8	7	13
NATDATE	The date field is in Natural D format (days since year zero).	3	4	4	7
UNIXTIME	The time field is in UNIX time_t type format (seconds since January 1, 1970).	4	4	6	10
XTIMESTAMP	The date and time field is in UNIX timestamp format, with microsecond precision, since January 1, 1970 (UNIXTIME * 60**6 + microseconds).	8	8	10	18

Time zone Option TZ

If this option is specified, the date-time field value is presented in the user's local time and stored in UTC time, allowing for differences in time zones.

Option CR

Use the CR option to indicate that the system field value should only be set when a record is inserted and not when it is updated. The CR field option can only be specified for fields defined with the SY field option, but cannot be specified for an MU field.

For more information see *Field Definition Statements* in the *Adabas Utilities* documentation for mainframes and *FDT Record Structure* in the *Adabas Administration* documentation for open systems.

- 7 Click **Add Field** to add the field to the Field Definition Table of the new file.



Note: Before the field is added to the FDT, Adabas Manager performs a check on all entries. Field names and their syntax and length are checked as well as any dependencies between fields and between field attributes. Existing Special Descriptors are checked, so that for example a field which is parent of a Special Descriptor cannot be deleted. Group fields are initially defined with all properties. During generation of the FDT, group fields are recognized and properties not required (format/length) are not used. In case of invalid values, the resulting error messages will show you how to modify your entries so that the field can be added to the FDT. For detailed information on the field definitions see the section *ADACMP: Compress-Decompress* in the *Adabas Utilities* documentation.

- 8 Specify all required fields and add them to the FDT.

Your entries are displayed in the table below the entry fields; to remove a field you created, click its check box and then click **Remove**.

To modify a field, select the field by clicking in the appropriate check box until a check mark appears and click **Edit**.

- 9 After defining all fields, click **Next** to continue with the file creation.



Note: Clicking **Finish** if the definition of special descriptors is not required will create the new file.

Page 3 of the **New File** appears in detail-view.

- 10 Define the Special Descriptor Table for the new file. Select the descriptor type in the selection box and complete the following entries:

- **Number**

only available for hyperdescriptor and collation descriptor; the user exit number to be assigned to the descriptor. The Adabas nucleus uses this number to determine the descriptor user exit to be called.

- **Name**

the name to be used for the descriptor/field. The naming conventions are identical to those for Adabas field names.

- Length

only available for hyperdescriptors; the default length of the hyperdescriptor.

- Select the format in the selection box.

- Parent

the name of the field from which the super-/subfield or super-/subdescriptor is to be derived. For collation descriptors and hyperdescriptors, an elementary field.

The following table shows some examples of the parent field entry syntax required by Adabas Manager in comparison to the ADACMP syntax:

Adabas Manager	ADACMP
AA,AB	HYPDE='01,HD,20,A,NU,MU=AA,AB'
AA(1,5)	SUBDE='SB=AA(1,5)'
AA	PHONDE='PH(AA)'
AA(1,5),BB(1,2),CC(3,5)	SUPDE='SP=AA(1,5),BB(1,2),CC(3,5)'
AA	COLDE='1,CD=AA'

Check the options you want to set; enter the position of the field if required.

- 11 Click **Add Descriptor** to add the descriptor to the Special Descriptor Table.



Note: Before the field is added to the Special Descriptor Table, Adabas Manager performs a check on all entries. In case of invalid values, the resulting error messages will show you how to modify your entries so that the field can be added to the FDT. For detailed information on the field definitions see the section *ADACMP: Compress-Decompress* in the *Adabas Utilities* documentation.

- 12 Repeat the action for each descriptor you want to create.

The result of your modifications is shown in a table below the entry fields; to remove a descriptor you created, click its check box and then click **Remove**.

To modify a descriptor, select the field by clicking in the appropriate check box until a check mark appears and click **Edit**.

- 13 After defining all required descriptors, click **Next** to continue with the file creation.



Note: Clicking **Finish** if the setting of optional parameters is not required will create the new file.

Page 4 of the **New File** appears in detail-view.

14 The following optional parameters can be set:

Property	Description
ACRABN	The address converter RABN space.
DSRABN	The data storage RABN space.
NIRABN	The normal index RABN space.
UIRABN	The upper index RABN space.
ASSO Padding	The padding factor (percentage of each block) set for the ASSO dataset (the default is 10).
DATA Padding	The padding factor (percentage of each block) set for the DATA dataset (the default is 10).
Max. Blocks of DS Extents	The number of blocks allowed per secondary Data Storage extent ("0" indicates that the parameter is not set; the default then is no limit).
Max. Blocks of NI Extents	The number of blocks allowed per secondary Associator normal index extent ("0" indicates that the parameter is not set; the default then is no limit).
Max. Blocks of UI Extents	The number of blocks allowed per secondary Associator upper index extent ("0" indicates that the parameter is not set; the default then is no limit).
Maximal Compressed Record Length	The maximum compressed record length (in bytes) allowed for the file (the default is the maximum possible size). You can also set
	Index Compression if the index for the file can be loaded in compressed form (the default depends on the input file).
	No AC Extension if the MAXISN value for the file can be increased ("YES" is required for expanded files).
	Program Refresh if a user program is allowed to refresh the file using an E1 call to Adabas when the file is loaded (the default is "NO").
Alpha Code	Current file encoding set for alphanumeric fields in the file.
Wide Code	Current file encoding set for wide-character fields in the file.
User Wide Code	Current user encoding set for wide-character fields in the file.
MINISN	The lowest ISN that can be assigned to a record in the file ("0" indicates that the parameter is not set; the default then is "1"). MINISN is required for expanded files.
ISN Size	Whether 3- or 4-byte ISNs are used for the file (the default is 3). You can also set
	ISN Reuse if a freed ISN can be reused for a new record (the default is "NO").
	DS Reuse if Data Storage space can be reused (the default is "YES").
Data Device	The device containing the Associator, Data Storage, or Work component. You can also set

Property	Description
	MIXDSDEV if secondary Data Storage extents may be allocated on different device types and therefore with different block lengths.
Anchor Fnr	Whether the file is part of an expanded file. You can also set Ciphering if the file is to be encrypted.
Multi Client Owner-ID Length	The length of the internal owner ID for multiclient files (if Multi Client Support has been set).
ADAMDE (Field/ISN)	Specifies the field to be used as the ADAM key (if the ADAM option has been selected for the file).
ADAPARM	The number of bits to be truncated from the ADAM descriptor value before it is used as input to the ADAM randomizing algorithm (if the ADAM option has been selected for the file). A value in the range 1-255 may be specified. If this parameter is omitted, a value of 0 bits (no truncation) will be used.
ADAM Overflow	The size of the Data Storage area to be used for ADAM file overflow (if the ADAM option has been selected for the file).

- 15 Specify any parameters you may wish to set and click **Finish** to create the new file.

28

Creating New LOB Files (Adabas for Mainframes Version 8 and Open Systems Only)

Large object fields (LB fields) are available with Adabas for mainframes version 8 and open systems. These are fields that may contain much more data than the 253 bytes of normal alphanumeric fields or the 16,381 bytes of LA fields. Fields containing large objects are defined with the new LB option. The theoretical maximum size of the value of an LB field is just short of 2 GB; practical usable sizes are smaller.

Adabas stores large object field values in a separate file, called a LOB file, that is tightly associated with the file containing the LB fields, which is called the base file. Behind the scenes, Adabas will store LB field values (except for very short ones) in the LOB file, but this is transparent to your application. Commands in application programs should always be directed against the base file; application programs need neither know nor care about the existence of a LOB file.

The **New LOB File** functionality of the Adabas Manager allows the creation of new LOB files.

▶ **To create a new LOB file for a database:**

- 1 Select an Adabas database in tree-view and expand it.
- 2 Select **Database Files** in tree-view and right-click it.
- 3 Select **New LOB File** on the drop-down menu.

The **Create LOB File** panel appears in detail-view.

- 4 Enter the following parameter values for the new file:
 - File Number
 - File Name (optional)
 - Base File (the file containing the large object (LB) fields)
 - MAXISN

- Data Storage Size
- Normal Index Size
- Upper Index Size



Note: If you click **Find** next to the **File Number** text box, the next free file number will be displayed. If you enter a file number that already exists and click **Find**, the next free file number after the number you entered will be displayed.

5 The following optional parameters can be set:

Property	Description
ACRABN	The address converter RABN space.
DSRABN	The data storage RABN space.
NIRABN	The normal index RABN space.
UIRABN	The upper index RABN space.
ASSO Padding	The padding factor (percentage of each block) set for the ASSO dataset (the default is 10).
DATA Padding	The padding factor (percentage of each block) set for the DATA dataset (the default is 10).
Max. Blocks of DS Extents	The number of blocks allowed per secondary Data Storage extent ("0" indicates that the parameter is not set; the default then is no limit).
Max. Blocks of NI Extents	The number of blocks allowed per secondary Associator normal index extent ("0" indicates that the parameter is not set; the default then is no limit).
Max. Blocks of UI Extents	The number of blocks allowed per secondary Associator upper index extent ("0" indicates that the parameter is not set; the default then is no limit).
Data Device	The device containing the Associator, Data Storage, or Work component.

You can also set:

Index Compression	if the index for the file can be loaded in compressed form (the default depends on the input file).
ISN Reuse	if a freed ISN can be reused for a new record (the default is "NO").
DS Reuse	if Data Storage space can be reused (the default is "YES").

6 Specify any parameters you may wish to set and click **Finish** to create the new LOB file.

29

File-related Long Running Utilities for Open Systems

▪ Backup (ADABCK)	122
▪ Restore (ADABCK)	123
▪ Load (ADAMUP)	124
▪ Unload (ADAULD)	125
▪ Compress (ADACMP)	127
▪ Decompress (ADADCU)	129
▪ Import (ADAORD)	131
▪ Export (ADAORD)	132
▪ Verify (ADAVFY)	133
▪ Invert Descriptor (ADAINV)	134
▪ Reinvert Descriptor	135
▪ Verify Descriptor (ADAINV)	135
▪ Descriptor Summary (ADAINV)	136
▪ Display Descriptors	136
▪ File Usage	136

Adabas Manager has two different ways of working with utilities:

- For utilities with a short processing time the job status is displayed in the detail-view. The user has to wait until the action is completed.
- For the so-called *Long Running Utilities* the job status is written to the System Management Hub Job Monitor. The utility runs in the background and leaves the user free to carry out other actions.

For detailed information on the SMH Job Monitor see *Monitoring Events and Jobs* in the *System Management Hub* documentation. For detailed information on Adabas utilities see the *Adabas for Open Systems Utilities* documentation.

The functions **Release Descriptor**, **Set Unique Descriptor**, **Reset Unique Descriptor** and **Add Constraint** are designated long running utilities (see [Setting and Releasing Descriptors](#)).

A number of other long running utilities have been grouped in a dropdown list in the **Database Files** detail-view.

▶ **To display the file-related utilities available with Adabas Manager:**

- 1 Select **Database Files** in tree-view.

The files are listed in a file list display table in detail-view. The available long running utilities are listed in a dropdown list at the bottom of the panel.

- 2 Mark the file and select the utility you want to run in the list. Click **Execute**.
- 3 In the detail-view panel, specify your options and the location of the output file and click **OK**.

The utilities are described in more detail in the following sections:

Backup (ADABCK)

Dumping only selected files allows a controlled backup of certain parts of the database in cases where backing up the complete database is unnecessary.

Options:

1. Allow parallel updates (DUMP)
Parallel updates are permitted on the files to be backed up while the backup is in progress. If this option is disabled, only ACC users (access only) can access the files to be backed up while the backup is in progress (EXU_DUMP).
2. Enable **Create New Protection Log** if you want to close the current protection log file and create a new log file after executing the backup.

3. If you have allowed parallel updates, enter a value to define the time (in seconds) that the system waits for ET-logic (end transaction) users to come to ET status at the end of the backup in the **Time To Wait For ET Status** field. If you do not enter a value, the value currently in effect for the database nucleus is taken.

Output:

1. Enable/Disable **Use Parallel Output Devices**.
2. The specification of the backup save set is displayed in the **File x** box. Either leave the specification displayed in the box unchanged if you want to accept it, or enter the specification directly. You can also browse to another file.
3. Enable **Create Two Physical Copies** check box if you want to create two copies of the save set.
4. Enable/Disable **Replace Existing Files**.
5. Click **OK** to start the utility or **Cancel**.

Restore (ADABCK)

A backup copy can be used to restore/overlay selected files if single files are corrupted. When restoring/overlying files, the RABNs required by the files must be available.

Options:

1. If **Allow for Allocation** is specified, the utility reallocates all files to be overlayed or the specified subset rather than attempting to restore them in the same block ranges as in the backup. Using this option reduces the number of file extents as much as possible.
2. If you want to restore the save set of one database into another database with a different DBID, select the **Ignore DBID in Backup File** check box, for example to restore the save set 'BCK001.004' (from the database with the DBID 4) into the database with the DBID 12.
3. Select the **Format Associator** and/or **Format Data Storage** check box if you want to format the blocks of the Associator and/or Data Storage. Formatting increases the time required for the restore to complete. If **Allow Relocation** is specified, ADABCK reallocates all files to be overlayed or the specified subset rather than attempting to restore them in the same block ranges as in the backup. Using this option reduces the number of file extents as much as possible.

Input:

1. The specification(s) of the save set(s) to be used for the restore is/are displayed in the **File x** box(es). Either leave the specification that is displayed in the box unchanged if you want to accept it, or enter the specification directly. You can also browse to another file.

Additional Options:

1. It is possible to see whether the save set can be used to restore a database by choosing **Read Check**. This function checks the readability (absence of parity errors) and completeness of the ADABAS backup copy.
2. You can display information about the backup save set by selecting **Report**.

Type:

The options **Contents**, **Summary**, and **Files** are available.

Contents displays a list of files (names and numbers) in an ADABAS backup save set, together with the name and number of the database.

Summary displays general information (name and number of the database, number of files loaded, maximum number of files, system file numbers), and the physical layout of the database.

Files displays status information for the files in an ADABAS backup save set.

3. Click **Restore** to start the utility or **Cancel**.



Note: An interrupted restore/overlay of one or more files will result in lost RABNs. The lost RABNs can be recovered by using the Recover Lost Blocks function from the Administrate menu.

Load (ADAMUP)

Load is used to load compressed data into an existing file in a database. The output of the Compress or Unload functions can be used as input. Input files that do not contain descriptor value tables (produced using the 'Do not Unload Descriptor' check box of the Unload function) can also be processed if the file to be updated does not contain any descriptors.

The selected database can be either active or inactive.



Note: The selected database must be inactive if you are updating the checkpoint or security file.

Input:

1. The specifications of the input files that contain the compressed data and the descriptor values are displayed in the **Data** and **Descriptors** boxes. Either leave the specification that is displayed in the box unchanged if you want to accept it, or enter the specification directly. You can also browse to another file.

Options:

1. If you want to skip records on the input file before starting to load data, enter the number of records to be skipped in the **Number of Records to Skip** box.
2. If you do not want to load all of the records on the input file, enter the number of records to be added in the **Max. Number of Records to Add** box. The Load function loads all of the records on the input file if you do not enter anything in this text box, and have not specified that records are to be skipped.
3. Enter the size of the workpool in the **Size of Workpool** box.
4. Select the **Load from Single File** check box if the input file was produced by the Compress or Unload functions with the **Compress to Single File** or **Unload to Single File** check boxes selected.
5. Select the **User ISNs** check box if you want to control ISN assignment for each record that is added to the database. Each ISN provided must be a four byte binary number immediately preceding each data record; within the current limit (MAXISN) for the file; unique within the specified file.
6. Select the **Allow for Other Descriptors** check box to continue processing if a field that is defined as a descriptor in the input data is not a descriptor in the FDT of the file being updated. Otherwise the Load function terminates and returns an error message.
7. Select the **Continue on UNIQUE Conflicts** check box to continue processing if conflicting values for a unique descriptor are detected. In this case, conflicting ISNs are written to the error log and the unique status of the descriptor in question is removed. Otherwise, the Load function terminates and returns an error message.

Miscellaneous:

1. The specification of the file that is to be used as the TEMP container is displayed in the **Temp Data** box. Either leave the specification that is displayed in the box unchanged if you want to accept it, or enter the specification directly. You can also browse to another file.
2. Click **OK** to start the utility or **Cancel**.

Unload (ADAULD)

Unload is used to retrieve records from a file in a database and write them to a sequential file. The main reasons for unloading a file are to change the space allocation, to reduce the number of logical extents assigned to the index, Address Converter or Data Storage, or to change the padding factor. In this case, the file has to be unloaded, deleted and then reloaded. Alternatively, the Unload function can be used to produce one or more test files that all contain the same data; this means that the file has to be unloaded and then reloaded using a different file number.



Note: The selected database must be inactive if you are unloading the checkpoint or security file.

Output:

1. The specifications of the output files that will contain the compressed data and the descriptor values are displayed in the **Data and Descriptors** boxes. Either leave the specification that is displayed in the box unchanged if you want to accept it, or enter the specification directly. You can also browse to another file.
2. Enable/Disable **Replace Existing Files**.

Options:

1. If you want to skip records before starting to unload data, enter the number of records to be skipped in the **Number of Records to Skip** box.
2. If you do not want to unload all of the records, enter the number of records to be unloaded in the **Max. Number of Records to Unload** box. The Unload function unloads all of the records if you do not enter anything in this text box, and have not specified that records are to be skipped.
3. **Sort Sequence:**
Select the sequence in which the records are to be unloaded: by physical sequence, by ISN or by Descriptor and enter the name of a descriptor, or sub- or superdescriptor, or the number of the ISN at which you want to start unloading (optional). If the specified ISN does not exist, unloading starts at the next highest ISN found.
4. Select the **Unload to a Single File** check box if you want to write the data information and the descriptor information to a single file.
5. Select the **Do not Unload Descriptor** check box if you want to omit the descriptor values that are used to build the index. You must not select this check box if you intend to use the output of the Unload function as input for the Load function. The **Unload to a Single File** check box and the **Do not Unload Descriptor** check box are mutually exclusive.
6. Select the **Decompress Unloaded Data** check box if you want to decompress the data immediately after unloading.
7. Click **OK** to start the utility or **Cancel**.



Note: The Unload function does not have a restart capability. An abnormally terminated unload must be rerun from the beginning.

Compress (ADACMP)

Compress is used to compress raw user data into a form which can be used by the Load function. The input data must be in a sequential file. Each field in the input record is compressed. Compression consists of removing trailing blanks from alphanumeric fields and leading zeros from numeric fields. Unpacked and packed fields are checked for correct data. Fields that are defined with fixed length are not compressed.

This function also covers format conversions for data that originally comes from mainframes.

The selected database can be either active or inactive.

Input:

1. The specifications of the input files that contain the raw data and the field definition table are displayed in the **Data** and **FDT** boxes. Either leave the specification that is displayed in the box unchanged if you want to accept it, or enter the specification directly. You can also browse to another file.
2. Specify the type of record separation used in the input file by selecting the corresponding structure from the **Record Structure** dropdown list:
 - **Excl.Len.Prefix**
The records are separated by a two-byte exclusive length field.
 - **Incl.Len.Prefix**
The records are separated by a two-byte inclusive length field.
 - **ASCII-File**
The records are separated by a new-line character.
 - **Host:Var.Blocked (Tape)**
The records are in variable blocked format.
 - **BS2000 (FTP)**
Format conversion BS2000 to UNIX or PC systems.
 - **IBM-Host (FTP), VAX**
Format conversion IBM (MVS, CMS) to UNIX or PC systems.
 - **RDW**
Format conversion of data that has been transferred from an IBM host using the FTP site rdw option.
3. Select the **User ISNs** check box if you want to assign the ISNs for the records in the input file. In this case you must provide the ISN to be assigned to each record as a four-byte binary number immediately preceding each record. ISNs can be assigned in any order and must be unique for the file. The ISN must not exceed the maximum number of records specified for the file. If you do not select this check box, the ISNs are assigned by the system.



Note: User ISNs cannot be used in conjunction with the ASCII-File record structure.

4. Select the **Null Value** check box if you are compressing data according to the standard FDT and the status values of the NC option fields are provided in the input data.
5. Select the **MU/PE count length**.
6. If you want to skip records before starting to compress data, enter the number of records to be skipped in the **Number of Records to Skip** box.
7. If you do not want to compress all of the records, enter the number of records to be compressed in the **Max. Number of Records to Process** text box. The Compress function compresses all of the records if you do not enter anything in this text box, and have not specified that records are to be skipped.
8. If you enter a character in the **Separator** text box, the fields in the raw data records must be separated by the character you specify. In this case, the same fields in different records can be of different lengths. If you specify a format buffer by choosing Fields, the order in which the field names are selected must correspond to the order in which the fields are specified in the FDT.
9. Enter the **Character Set** which specifies the default encoding used in the decompressed file based on the encoding names listed at <http://www.iana.org/assignments/character-sets>.

Temporary Data:

1. The file specified in the **Temporary Data** box is used to store the output of the format conversion step before the actual compression step starts if you have selected BS2000 or IBM-Host (FTP) as **Record Structure**. Either leave the specification that is displayed in the box unchanged if you want to accept it, or enter the specification directly. You can also browse to another file.

Output:

1. The specifications of the output files that will contain the compressed data and the descriptor values are displayed in the **Data** and **DVT** boxes. Either leave the specification that is displayed in the box unchanged if you want to accept it, or enter the specification directly. You can also browse to another file.

If you select **Compress to Single File**, the data and the descriptor information are written to a single file, and the DVT text box is disabled.

2. Select the **Load Compressed Data** check box if you want to load the compressed data immediately after compression. After choosing OK, the **Load** dialog box will be displayed before the Compress function starts.
3. Enable/Disable **Replace Existing Files**.

Source Architecture:

1. Specify the source architecture (character set, floating point format, and byte order) of the input records by selecting the corresponding entries from the **Source Architecture** dropdown lists.
2. Choose **Fields** if you want to specify a subset of the fields defined in the FDT. In this case, the input records do not have to contain all of the fields that are defined in the FDT. Enter the desired format buffer in the dialog box.
3. Click **OK** to start the utility or **Cancel**.



Note: The Compress function does not have a restart capability. An abnormally terminated compress must be rerun from the beginning. The Compress function does not change the database, therefore no considerations need to be made concerning the status of the database before a compress is restarted.

Decompress (ADADCU)

Decompress is used to decompress records produced by the Compress and Unload functions. The output produced can be used as input for programs that use standard file access methods. It can also be used as input for the Compress function once you have made any changes to the data structure or to the data definitions.

You can decompress either complete records, or selected fields within records. If you decompress complete records, any multiple-value fields or periodic groups are preceded by a one-byte count field. If you only decompress selected fields, you can change the length and formats of the fields, and allocate space for the subsequent addition of new fields by using a literal element or a blank element.

The selected database can be either active or inactive.

Input:

1. The specification of the input file that contains the compressed data is displayed in the **Data** box. Either leave the specification that is displayed in the box unchanged if you want to accept it, or enter the specification directly. You can also browse to another file.
2. If you want to skip records before starting to decompress data, enter the number of records to be skipped in the **Number of Records to Skip** box.
3. If you do not want to decompress all of the records, enter the number of records to be decompressed in the **Max. Number of Records to Process** text box. The Decompress function decompresses all of the records if you do not enter anything in this text box, and have not specified that records are to be skipped.

Output:

1. The specifications of the output files that will contain the decompressed data and the field definition table are displayed in the **Data** and **FDT** boxes. Either leave the specification that is displayed in the box unchanged if you want to accept it, or enter the specification directly. You can also browse to another file.

The **FDT** box is only enabled if you select the **Dump FDT** box.

2. Enable/Disable **Replace Existing Files**.
3. Enable/Disable **Dump FDT**.
4. Select the **Null Value** check box if you want to decompress the records according to the standard FDT, and the records contain NC option fields with their status values.
5. Select the **MU/PE count length**.
6. Select the **User ISNs** check box if you want the ISN to be output together with each decompressed record. You should do this if you want to subsequently reload the file with the same ISNs.



Note: User ISNs cannot be used in conjunction with the ASCII-File record structure.

7. Select the **Truncation of Alphanumeric Fields** check box if you want to enable the truncation of alphanumeric field values. If you do not select this check box, all of the records which contain truncated alphanumeric field values are written to the error log.
8. Specify the type of record separation to be used in the output file by selecting the corresponding structure from the **Record Structure** dropdown list. The following structures are supported: .

■ **Excl.Len.Prefix**

The records are separated by a two-byte exclusive length field. There is no separator character, and there are no restrictions on the use of this format.

■ **Incl.Len.Prefix**

The records are separated by a two-byte inclusive length field. There is no separator character, and there are no restrictions on the use of this format.

■ **ASCII-File**

The records are separated by a new-line character. This format can only be specified if the field values of the output do not contain the new-line character (i.e., if the record only contains packed, unpacked, and alphanumeric fields, and if the alphanumeric fields only contain printable characters).

■ **Host:Var.Blocked (Tape)**

The records are stored in blocks. Each record begins with an inclusive four-byte length field.

■ **BS2000 (FTP)**

Format conversion BS2000 to UNIX or PC systems.

- **IBM-Host (FTP), VAX**

Format conversion IBM (MVS, CMS) to UNIX or PC systems.

- **RDW**

Format conversion of data that has been transferred from an IBM host using the FTP site rdw option.

9. Enter the **Character Set** which specifies the default encoding used in the decompressed file based on the encoding names listed at <http://www.iana.org/assignments/character-sets>.

Target Architecture:

1. Specify the target architecture (character set, floating point format, and byte order) of the output records by selecting the corresponding entries from the **Target Architecture** dropdown lists.
2. Choose **Fields** if you want to define a new structure for the output records. The fields to be decompressed are specified by their field name, format and length. Enter the format buffer that you want in the dialog box that appears.
3. Click **OK** to start the utility or **Cancel**.



Note: The Decompress function does not have a restart capability. An abnormally terminated decompress must be rerun from the beginning. The Decompress function does not change the database, therefore no considerations need to be made concerning the status of the database before a decompress is restarted.

Import (ADAORD)

Import imports one or more files into a database, using the data in the sequential file (ORDEXP) produced when exporting files or a database. The files to be imported may not already be loaded in the database and must not exceed the maximum file number defined for the database.



Note: By default, the Import function controls the placement of files and the allocation quantities. The options and parameters that can be used to overwrite the default settings can only be used if you are importing a single file.

If you are importing a system file (checkpoint or security file), the selected database must be inactive.

Options:

1. Enter the number(s) of the file(s) that you want to import in the **File Number(s)** box. The file numbers specified must not exceed the maximum file number defined for the database in question.
2. Enter a new value for the maximum ISN for the file in the **Maximum ISN** box. The Import function uses this number to determine the new amount of space to be assigned to the file's

Address Converter. If you do not enter a value, the maximum ISN in effect for the file when it was exported is used.

3. Enter a new file number for the file to be imported in the **New File Nr.** box if you want the file to have a different file number to the one it was exported with.
4. Enter a new lobfile number for the LOB file to be imported in the **New LOB File Nr.** text box if you want the LOB file to have a different file number to the one it was exported with.
5. Enter the **LOB Size**.
6. The specification of the input file that is to be used is displayed in the **Input** box. Either leave the specification that is displayed in the box unchanged if you want to accept it, or enter the specification directly. You can also browse to another file.

Structure:

1. Enter the new amount of space to be allocated to the structures in the **Size** boxes. If you want to allocate a structure to a given disk section or to a specific container extent, enter the starting RABN in the **Start RABN** box. If you do not enter any values, the Import function assigns the starting RABNs and calculates the respective sizes.

Padding:

1. Enter new values in the **Padding** boxes if you want to change the padding factors for the Associator and Data Storage. Specify a value in the range 0 - 10 if you expect little or no descriptor updating (ASSO), or little or no record expansion (DATA). You should specify a value in the range 10 - 50 if you expect a large amount of descriptor updating/record expansion.
2. You can display information about the input file by choosing **Report**. The information displayed includes the time and date when the files were exported, the name and number of the database from which they were exported, the names and numbers of the exported files, and the dates on which they were loaded.
3. Click **OK** to start the utility or **Cancel**.



Note: An abnormally-terminated import of one or more files will result in lost RABNs for the last file being imported.

Export (ADAORD)

Export exports (copies) one or more files from the database to a sequential output file (ORDEXP). The Data Storage of each file is copied, together with the information that is required to reestablish its index.

If you are exporting a system file (checkpoint or security file), the selected database must be inactive.

Options:

1. The specification of the data file (ORDEXP) that is produced is displayed in the **Export Data** box. Either leave the specification that is displayed in the box unchanged if you want to accept it, or enter the specification directly. You can also browse to another file.

2. Enable/Disable **Replace Existing Files**.

3. Sort Sequence:

Select the sequence in which the records are to be exported: by physical sequence, by ISN or by Descriptor and enter the name of a descriptor, or sub- or superdescriptor. The descriptor specified must not be a multiple-value field, be contained in or derived from a periodic group or refer to a field defined with null value suppression.



Note: You can only sort by descriptor if only one file has been selected.

4. Click **OK** to start the utility or **Cancel**.



Note: The Export function does not have a restart capability. An abnormally terminated export must be rerun from the beginning.

Verify (ADAVFY)

Verify is used to verify the internal consistency of a file or selected files. The general control block (GCB) is validated together with the file control block (FCB) and the field definition table (FDT) of the selected file. You can verify the Associator (Address Converter and index) and/or the Data Storage. Running the Verify function with an active database may lead to errors being reported; this is because updates may be made while the function is processing, and these updates will only be reflected in the nucleus buffer pool.

Running the Verify function with an active database may lead to errors being reported; this is because updates may be made while the function is processing, and these updates will only be reflected in the nucleus buffer pool.

Verify Associator:

1. Select **File Control Block** to verify the the file control block together with the field definition table for the selected file.
2. Select **Address Converter** to verify the Address Converter; this function checks the existence of a record within the Data Storage for an existing ISN in the Address Converter.

Choose the extent of your output in the dropdown list.

3. Select the **Index** check box if you want to verify the index.

Choose the extent of your output in the dropdown list.

4. The **RABN to dump from** boxes are used to specify a range of RABNs (Relative ADABAS Block Numbers) in the index. Only the RABNs specified will be dumped.

Verify Data:

1. Select **Data** to verify the Data Storage of the selected file. This function checks the existence of an ISN within the Address Converter for an existing record in the Data Storage.

Record validates the Data Storage and checks the structure of each record.

Choose the extent of your output in the dropdown list.

2. Enable **Field** to verify the record structure and validate the contents of packed, unpacked and floating point fields.

Miscellaneous:

1. Select **Search Lost Blocks** to search for lost RABNs in the selected file. The Dump Physical Structure option dumps the layout of the database.
2. Enter the maximum number of errors to be reported before the Verify function terminates.
3. Click **OK** to start the utility or **Cancel**.

Invert Descriptor (ADAINV)

Invert is used to create new elementary, sub-, super-, hyper- and phonetic descriptors at any time after a file has been initially loaded.

Options:

1. Specify the **Size of Workpool**.
2. Enable/disable the **Continue on Unique Conflicts** box.

Field Options:

1. Enter the name of the field that you want to invert in the **Short Name** box.

The short name must either be the name of an existing elementary field for which an elementary descriptor is to be defined, or a new name that has not been used in the FDT yet in the case of sub-, super-, hyper-, and phonetic descriptors.

2. Use the **Descriptor** dropdown list and the **Unique** check box (if required) to specify the type of descriptor.
3. Click **Add** to add the new descriptor to the field list.

4. Click **OK** when you have finished defining all of the new descriptors, or **Cancel**.



Note: Defining a new descriptor leads to modifications in the Field Definition Table (FDT) and results in the creation of an inverted list. New Main Index and Normal Index blocks are required to store the inverted list entries for the new descriptor. The new Normal Index and Main Index are built on a descriptor by descriptor basis. During this process, the linking entries are still missing in the Upper Index and none of the new inverted lists can be accessed. When the Normal Index and Main Index of all the new descriptors have been built, the FDT is updated and corresponding entries are added to the Upper Index.

Reinvert Descriptor

Reinvert is used to implicitly release and subsequently invert one or more descriptors.

Options:

1. Select the **Continue on Unique Conflicts** check box to continue processing if duplicate values are found for a unique descriptor. In this case, the unique status of the descriptor in question is removed, and processing continues. Otherwise the Reinvert function terminates and returns an error message if duplicate values are detected.
2. Select the **Format** box if you want to format index blocks that are no longer used but which are still allocated to the selected file; this is because the new index created when a descriptor is re-inverted is generally smaller than the old index.
3. Specify the **Size of Workpool**.
4. Select the descriptors.
5. Click **OK**, or **Cancel**.

Verify Descriptor (ADAINV)

Verify Descriptor checks the integrity of the inverted lists of the specified file.

Options:

1. Specify the **Size of Workpool**.
2. In the **Max. No. of Errors** box, enter the number of errors that have to be detected before the Verify function terminates. If you do not enter anything in this box, the default value is used.
3. Select the descriptors.
4. Click **OK**, or **Cancel**.

Descriptor Summary (ADAINV)

Verify Descriptor displays the descriptor space summary for specified descriptors and the sizes required to process these descriptors.

Options:

1. Select the **Full Summary** box if you want to display each descriptor together with the sizes required for that descriptor.
2. Select the descriptors.
3. Click **OK**, or **Cancel**.



Note: Inverted lists are maintained for each elementary, sub-, super-, hyper- and phonetic descriptor defined within a file. In order to guarantee their integrity, it must be ensured that each ISN in the inverted list is associated with an existing data record and that this data record is the correct one; that each record in the data storage is represented in the inverted list by its ISN and the descriptor value entries generated. When verifying a descriptor, the Verify function simulates loading of the normal index and matches the output from the sort against the content of the inverted list. This checks both of the points mentioned above in one run and detects uniqueness conflicts. All inconsistencies found will be reported. The file remains unchanged.

Display Descriptors

Options:

1. Select **All Descriptors** or enter the names of the descriptors you want to display.
2. Enable/disable Histogram.
3. Click **OK**, or **Cancel**.

File Usage

Options:

1. Specify **Data Storage**, **Normal Index** or **Upper Index**.
2. Click **OK** to run the utility, or **Cancel**.

VI

Session Monitoring

This part of the Adabas Manager documentation presents the Adabas session monitoring functions, which allow you to control and manage major Adabas resources. These functions are most useful when analyzing system performance or seeking the cause of performance problems.

The information is organized under the following headings:

Active Parameters

Cache Facility

Delta Save Facility

Disk Storage


Mainframe Database Maintenance Levels

System Status

30 Active Parameters

- Active Parameters for Mainframe Databases 140
- Active Parameters for Open Systems Databases 143



Note: The  button above the table in detail-view allows switching from regular layout of the tables to expert layout. Expert layout will also list Adabas nucleus parameters ADARUN for mainframe databases and ADANUC for open systems databases.

▶ **To review the parameter settings for an Adabas database:**

- 1 Select an Adabas database in tree-view and expand it.
- 2 Select **Active Parameters** in tree-view.

The **Active Parameters** for the database are displayed in tables in detail-view:

- Pools and Queues
- Time Limits
- Logging
- Options
- User Exits to be Used (for open systems databases only)



Note: The sections in the table can be collapsed and expanded by clicking the arrow in the section heading.

Active Parameters for Mainframe Databases

When Active Parameters is selected in tree-view, the following table is displayed in detail-view:

- Pools and Queues
- Time Limits
- Logging
- Options

The meaning of ON/OFF options in the **Logging** section of the active parameters panel are explained in the following table:

Parameter	Action
Command Logging	Enable command logging.
Log Adabas Control Block	Enable Adabas control block logging.
Log Format Buffer	Enable format buffer logging.
Log Record Buffer	Enable record buffer logging.
Log Search Buffer	Enable search buffer logging.

Parameter	Action
Log Value Buffer	Enable value buffer logging.
Log ISN Buffer	Enable ISN buffer logging.
Log I/O Activity	Enable I/O activity logging.
Log User Exit B Data	Log User Exit B data (only available if CLOGLAYOUT=5 (the default) is also used).
Merge Cluster Command Logs automatically	Enable automatic merging of cluster command logs.
Protection Log Required	Indicates whether a protection log is required for this session (the default is On).

The meaning of ON/OFF options in the **Options** section of the active parameters panel are explained in the following table:

Parameter	Action
Log AOS/DBS Update	Log commands that update the active nucleus coming from AOS or ADADBS.
Batch Support	Enable batch support.
Async. by Vol-Ser	Flush buffers asynchronously based on volume serial number.
Cache Facility	Enable support for Dynamic Caching.
Delta Save Facility	Enable support for the Delta Save Facility.
Resource Manager	Enable distributed transaction processing as a resource manager with other resource managers (databases) coordinated by a transaction manager.
Transaction Manager	Enable distributed transaction processing as a transaction manager.
Ignore DIB Entry	Ignore and delete an existing data integrity block (DIB) entry in the Associator and initiate a new session.
Ignore Work Part 4	Ignore data related to the recovery of previously incomplete transactions that resides in the two-phase commit area (Work part 4) and start the nucleus anyway. For emergency use only: contact your Software AG technical support center before using.
LOCAL Nucleus	Restrict a nucleus to local use only (no connection to other Entire Net-Work nodes).
Non-DE Search	Restrict searches to descriptors only in search buffers (no searches on nondescriptors).
OPEN Required	Require an OPEN command as the first command of a user session.
Read-Only Session	Restrict a database session to reading data only (no updates allowed).
Ext. Error Recovery	Enable extended error handling and message buffering.
Triggers and Procedures	Enable use of triggers and stored procedures.
UTI-only Session	Restrict a database session to running utilities only.
TCP/IP Access Control	Enable a direct TCP/IP link to the Adabas nucleus (UES-enabled databases only).

► **To modify the parameter settings for an Adabas mainframe database:**

1 Click **Modify**.

The parameter tables are displayed. Where settings can be modified, text boxes allow you to enter the required values; where options can be selected, check boxes can be marked.

2 The following parameter values can be modified:

Pools and Queues

- Sort Area
- Int. User Buffer
- Parallel LFIOP I/O
- Hold Queue Limit
- Active Command IDs per User
- ISN per TBI Element

Time Limits

- Command Time
- Bufferflush Duration
- Time Limit Sx-Commands
- Nonactivity ACC-User
- Nonactivity ET-User
- Nonactivity EXU-User
- Transaction Time

Logging (on/off)

The following parameters can be set:

- Command Logging
- Log Adabas Control Block
- Log Format Buffer
- Log Record Buffer
- Log Search Buffer
- Log Value Buffer
- Log ISN Buffer
- Log I/O Activity

- Log User Exit B Data
- Merge Cluster Command Logs automatically
- Protection Log Required (the default is **On**)

Options (on/off)

The following options can be checked:

- Log AOS/DBS Update
- Batch Support
- Async by Vol-Ser (the default is **On**)
- Read-Only Session
- UTI-Only Session

For detailed information about ADARUN parameters and their possible values, see *ADARUN Control Statement* in the *Adabas Operations* documentation.

- 3 Modify the settings in the table as required and click **OK** to confirm your changes and to return to the **Active Parameters** display.

Active Parameters for Open Systems Databases

When Active Parameters is selected in tree-view, the following table is displayed in detail-view:

- Pools and Queues
- Time Limits
- Logging
- Options
- User Exits to be Used



Note: A dropdown menu allows the user to alternate the display between the current and the default parameters.

The meaning of ON/OFF options in the **Logging** section of the active parameters panel are explained in the following table:

Parameter	Action
Log Adabas Control Block	Enable Adabas control block logging.
Log Format Buffer	Enable format buffer logging.
Log Record Buffer	Enable record buffer logging.
Log Search Buffer	Enable search buffer logging.
Log Value Buffer	Enable value buffer logging.
Log ISN Buffer	Enable ISN buffer logging.
Log I/O Activity	Enable I/O activity logging.
Log Buffer Description	Enable buffer description logging.
Command Log Layout	New CLOG Layout
Protection Log required	Indicates whether a protection log is required for this session (the default is On).

The meaning of ON/OFF options in the **Options** section of the active parameters panel are explained in the following table:

Parameter	Action
Truncation	Enable truncation of alphanumeric field values.
Utilities Only	If this is enabled, all non-utility calls are rejected and the DBA has exclusive control over all database files.
Local Utilities	If <code>Local Utilities</code> is enabled, remote utilities cannot be executed.
Open Required	If <code>Open Required</code> is enabled, an open (OP) command must be issued as the first command of a user session. This option should be set if <code>lnk_set_adabas_id</code> is used when calling Adabas from application servers, and also when using Net-Work, otherwise in these cases Adabas cannot guarantee transaction integrity following an ADANUC restart.
Fault Tolerant AR	When <code>Fault Tolerant AR</code> is enabled, the nucleus behaviour in the event of an Adabas error during autorestart can be controlled. If an error occurs for a given file, a detailed entry is made in the nucleus log, but the autorestart continues. When the autorestart completes, the DBA can restore and regenerate the file by using <code>ADABCK RESTORE</code> and <code>ADAREC REGENERATE</code> for the affected file. If however the error occurred for the index of a file, it is sufficient to rebuild the file's index by using the <code>REINVERT</code> function of <code>ADAINV</code> . If <code>Fault Tolerant AR</code> is not selected, Adabas aborts an autorestart if an error is detected, and this can cause the database to be in an inconsistent state. Note: The nucleus log can be displayed by right-clicking on the database name and choosing Nucleus Log in the dropdown menu.
Autorestart Only	The <code>AUTORESTART_ONLY</code> keyword shuts down the nucleus immediately after its startup sequence has completed. If an autorestart is pending, the autorestart will be performed. No user commands or utility calls will be accepted by the nucleus.
Read Only	The <code>READONLY</code> option causes ADANUC to run in read-only mode. Refer to the <i>Adabas for Open Systems Administration</i> documentation for more details.

Parameter	Action
XA	The keyword XA indicates that the server will support distributed transaction processing according to the X/Open XA specification. If the Adabas XA interface is to be used by an application, OPTION=XA must be used. See <i>XA Support</i> in the <i>Adabas for Open Systems Administration</i> documentation for further information.
Auto Expand	Enable auto expand when the database becomes full. You can specify reserved and temporary locations for database auto expand by right-clicking the database name and selecting Locations in the dropdown menu. For detailed information see the <i>Adabas Extended Operation</i> documentation.
Write Before Image to PLOG	Enable Write Before Image to PLOG.
Behaviour on Conflict During AUTORESTART	If Continue is enabled, the autorestart will continue even in case of a conflict (the default is Abort).

For more information on the nucleus parameters see the *Adabas Open Systems Utilities* documentation.

► **To modify the parameter settings for an Adabas open systems database:**

1 Click **Modify**.

The parameter tables are displayed. Where settings can be modified, text boxes allow you to enter the required values; where options can be selected, check boxes can be marked.

2 The following parameter values can be modified:

Pools and Queues

- Number of Threads
- User Queue
- Hold Queue Limit
- Group Commits Limit
- Number of Client Threads Elements
- Max. Usage of Buffer Pool Before Flushed to Disk
- Attached Buffer Length
- Attached Buffer Length Extended
- Buffer Pool
- Work Pool
- XA Area Limit
- Parallel I/O Request Limit

- Parallel Records Read Limit
- Parallel Blocks Read Limit
- Parallel Total Read Limit

Time Limits

- Nonactivity ACC-User
- Nonactivity ET-User
- Nonactivity EU-User
- Transaction Time

Logging (on/off)

The following parameters can be set:

- Log Adabas Control Block
- Log Format Buffer
- Log Record Buffer
- Log Search Buffer
- Log Value Buffer
- Log ISN Buffer
- Log Fault Tolerant Auto Activity
- Log I/O Activity
- Log Buffer Description
- Maximum Command Log Size
- Command Log Layout (New CLOG Layout)
- Protection Log Required (the default is **On**)

Options (on/off)

The following options can be checked:

- Utilities Only
- Local Utilities
- Open Required
- Fault Tolerant AR
- Autorestart Only
- Read Only

- XA
- Auto Expand
- Write Before Image to PLOG (the default is **On**)
- Behaviour on Conflict During AUTORESTART (the default is **Abort**)

User Exits to be Used

User exits can be activated or deactivated as required.

- 3 Modify the settings in the table as required and click **OK** to confirm your changes and to return to the **Active Parameters** display.

Environment Extensions

▶ To display the environment extensions for an Adabas database:

- 1 Select an Adabas database in tree-view and expand it.
- 2 Right-click **Active Parameters** in tree-view and choose **Environment Extensions** in the dropdown menu.

The **Environment Extensions** for the database are displayed in detail-view.

- 3 To add extensions, click **Add**.
- 4 Enter name and value of the extensions as required and click **OK** to confirm your settings.

Action Users

▶ To add an action user:

- 1 Select an Adabas database in tree-view and expand it.
- 2 Right-click **Active Parameters** in tree-view and choose **Action Users Extensions** in the dropdown menu.

The **Add Action User** panel is displayed in detail-view.

- 3 Enter the user name and click **OK** to confirm your settings.

Node Parameters

▶ **To display the node parameters for an Adabas database:**

- 1 Select an Adabas database in tree-view and expand it.
- 2 Right-click **Active Parameters** in tree-view and choose **Node Parameters** in the dropdown menu.

The **Node Parameters** for the database are displayed in detail-view:

- Adabas Operations
 - Logging
 - Archive Log Files
 - Alerts
- 3 In the table you can activate/deactivate and enable/disable these parameters according to your requirements and browse to the locations where the reports should be written to.

Click **Save** to save your new settings.

Database Actions

▶ **To display the database actions for an Adabas database:**

- 1 Select an Adabas database in tree-view and expand it.
- 2 Right-click **Active Parameters** in tree-view and choose **Database Actions** in the dropdown menu.

The **Database Actions** for the database are displayed in detail-view:

- Nucleus Startup
- Nucleus Termination
- Delete Checkpoints
- Increase Size of ASSO
- Increase Size of DATA
- Increase Nucleus Parameter LBP
- Increase Nucleus Parameter LS
- Increase Nucleus Parameter LWP
- Increase Nucleus Parameter NH
- Increase Nucleus Parameter NT

- Increase Nucleus Parameter NU
 - Observe Number of Offline Checkpoint
 - Recover Lost Blocks Required
 - Reorder File
 - Backup Database
 - Nucleus Termination Using 'adastop'
- 3 In the table you can enable or disable these parameters according to your requirements and browse to scripts you may wish to use.

Click **Save** to save your new settings.

31 Caching Facility

▪ Defining a File Cache	152
▪ Maintaining Cache Parameters	153
▪ Cache Session Summary	154

The Adabas Caching Facility is used to define caching parameter settings and also to monitor caching operations. For detailed information, see the *Caching Facility* section in the *Adabas for Mainframes* documentation.

When the Adabas Caching Facility is installed on your system, the Adabas Manager provides easy access to the optional caching facilities.



Note: If the Adabas Caching Facility is not installed on your system, the icon will not be displayed in your Adabas Manager tree-view.

▶ **To view the files with caching for a database:**

- 1 Select an Adabas database in tree-view and expand it.
- 2 Select **Caching Facility** in tree-view.

In detail-view, a table displays file names and numbers of the files with caching as well as their cache type and status.

To change the status of a file, select the file by checking the appropriate box in the table and click **Enable**, **Disable** or **Delete**.

Defining a File Cache

▶ **To define a file cache:**

- 1 Select **Caching Facility** in the tree-view and right-click it.
- 2 Select **Define File Cache** on the drop-down menu.

In detail-view, the **Define File Cache** table is displayed and allows you to define one or more files to be cached.



Note: Files already cached are displayed in a second table together with their cache type and status.

- 3 The following specifications can be set in the table:

■ **Cache Type**

Files may be cached in **ASSO**, **ASSO and DATA**, and **DATA** space.

- **Memory Type**

The memory type can be **Extended**, **Hiperspace**, **Dataspace**, or **Virtual 64**. The choices available depend on the operating system configuration under which Adabas is running. Virtual 64 is only available for systems running z/OS version 1.2 or above on the appropriate hardware.

- **File Number From... To**

Specifying the start file number and the end file number defines a file range to be cached in the defined cache space. If only one file is to be cached, it is not necessary to specify an ending file number.

- **Class of Service**

Allows you to define the percentage of the file you want to be cached (100%, 75%, 50%, 25%, or 10%).

Enter your specifications in the table and click **OK** or **CANCEL**.

The files you have successfully cached are displayed in a second table together with their cache type and status, i.e, added to the table where cached files already existed.

Maintaining Cache Parameters

▶ **To maintain cache parameters:**

- 1 Select **Caching Facility** in the tree-view and right-click it.
- 2 Select **Maintain Cache Parameters** on the drop-down menu.

In detail-view, the **Maintain Cache Parameters** table displays the current setting of the Adabas Caching (ADARUN) parameter values and allows you to change them.

- 3 The following specifications can be set in the table:

- Display Cache Space Statistics - Yes or No
- Minimum Buffer Efficiency
- Elapsed Time (GETMAIN)
- Time of Inactivity
- Maximum Number Permitted
- Maximum ASSO Cache Size
- Maximum DATA Cache Size
- WORK Memory Type
- Percent WORK2 RABNs

- Percent WORK3 RABNs
- **Read Ahead Caching**
 - Number of Buffers
 - Command Types Excluded



Note: Values for Max. ASSO Cache Size and Max. DATA Cache Size represent bytes, but may be marked with a K (1 kilobyte=1024 bytes), M (1 megabyte=1,048,576 bytes), or G (1 gigabyte=1,073,741,824 bytes) indicator when the size of the value requires it.

For more information about these parameters, see *Adabas Caching Parameters* in the *Caching Facility* section in the *Adabas for Mainframes* documentation.

Enter your specifications in the table and click **OK** or **CANCEL**.

Cache Session Summary

▶ **To view the cache session summary:**

- 1 Select **Caching Facility** in the tree-view and right-click it.
- 2 Select **Cache Session Summary** on the drop-down menu.

In detail-view, the **Cache Session Summary** table displays a summary of information for the Adabas Caching session including cache statistics, read-ahead statistics, and allocation statistics for active and inactive cache spaces.

32

Delta Save Facility

- Displaying Files Modified by Utilities 157
- Displaying Size Estimate for Delta 157

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



Note: If the Adabas Delta Save Facility is not installed on your system, the icon will not be displayed in your Adabas Manager tree-view.

▶ **To view the Delta Save information:**

- 1 Select an Adabas database in tree-view and expand it.
- 2 Select **Delta Save Facility** in tree-view.

In detail-view, a table displays the Delta Save Facility status as well as the following allocation information:

- Number of Full Saves
- Last Full Save Date
- Number of Delta Saves
- Last Delta Save
- DLOG Area Size, which is the number of Associator blocks to be allocated. The area must be large enough to hold all blocks changed between two consecutive Delta Save operations. About 5 bytes of DLOG area storage are needed for every database write operation. To calculate the size of the DLOG area, see *Calculating the DLOG Area Size* in the *Delta Save Facility* section of the *Adabas for Mainframes* documentation.
- DLOG Area Usage (%)
- DLOG Area Location (Start RABN), which is the first RABN to be allocated.
- Device Type, the device type to be used for DLOG area allocation.

To disable the Delta Save Facility, click **Disable**.



Note: When DSF logging is disabled, you may no longer perform Delta Save operations. To enable DSF logging again, you will need to do a full database save using ADASAV SAVE. For that reason you will be prompted to confirm that you want to disable DSF logging.

Displaying Files Modified by Utilities

▶ **To display files modified by utilities:**

- 1 Select Delta Save Facility in the tree-view and right-click it.
- 2 Select **Files Modified by Utilities** on the drop-down menu.

In detail-view, the **Delta Save Facility - Files Modified** table displays the files modified by utilities, indicating the name and number of the files changed as well as the following information:

- total blocks changed in all logged areas
- total AC/DS blocks used, changed, and the percentage of the used that were changed
- total index blocks used, changed, and the percentage of the used that were changed

Separate totals for the number of DS/AC blocks and index blocks changed for all files listed are displayed in a second table.

Displaying Size Estimate for Delta

▶ **To display the size estimate for delta:**

- 1 Select Delta Save Facility in the tree-view and right-click it.
- 2 Select **Size Estimate for Delta** on the drop-down menu.

In detail-view, the **Delta Save Facility - Estimated Number of Changed Blocks** table displays an estimate of the number of database blocks that would be saved in a Delta Save operation at the current time.



Note: The estimate is not precise: the actual size may be larger or smaller.

The fixed area includes the blocks for the first 30 RABNs, all FDTs and all FCBs. Separate totals are provided for the number of index blocks as opposed to the number Data Storage and address converter blocks changed by utilities. The calculated number of blocks for the DLOG area is a rounded number. The total of all blocks to be saved is approximate to the nearest 100.

33

Disk Storage

- Adding a Data Storage or Associator Dataset 160
- Modifying a Data Storage or Associator Dataset for Mainframe Databases 161
- Modifying a Data Storage or Associator Dataset for Open Systems Databases 162
- Recovering Unused Storage 163
- Reviewing Volume Distribution for a Mainframe Database 163

▶ **To view the disk storage information for a database:**

- 1 Select an Adabas database in tree-view and expand it.
- 2 Select **Disk Storage** in tree-view.

The **Containers** of the database are displayed in detail-view.



Note: For mainframe databases, a dropdown list in detail-view allows you to choose between a **Containers**, **Unused Storage**, and **Volume Distribution** display. The default is set to **Containers**.

▶ **To view RABNs for a mainframe database:**

- 1 Select an Adabas database in tree-view and expand it.
- 2 Select **Disk Storage** in tree-view and right-click it.
- 3 Select **Display ASSO/DATA Block (RABN)** on the drop-down menu.

The **ASSO/DATA Block** information is displayed in detail-view

- 4 Select the RABN type (ASSO or DATA) in the selection box and enter the RABN number in the text box.
- 5 Enter the offset and length values to determine starting point and length of the RABN display.

The specified RABN block from the Associator or Data Storage in hexadecimal format is displayed. Both the hexadecimal data and its alphanumeric equivalent are displayed. If the block is not assigned, zeros are displayed.

Adding a Data Storage or Associator Dataset

The Associator and Data Storage dataset sizes must be added separately. It is not possible to add both with a single operation.

▶ **To add a DATA or ASSO dataset:**

Backup the database using the ADASAV utility. This step is optional, but recommended.

- 1 Select **Disk Storage** in the expanded tree-view of the selected database.

The default setting in the of the detail-view is **Containers**. Size information, RABN range and device type of each extent are displayed in a table in detail-view.

- 2 Click **Add**.

- 3 In the resulting detail-view panel, enter the required parameters to add either a DATA or ASSO dataset.
- 4 Click **OK** to save your modifications or **Cancel** to return to the previous panel.



Caution: After an *Add* operation has completed, the nucleus session will be automatically terminated to allow for the necessary formatting of the Associator or Data Storage. If the new dataset is not formatted as required, an I/O error will occur once the nucleus tries to use it and the nucleus will ABEND.

▶ **The following additional two steps are required for mainframe databases:**

- 1 Allocate the dataset with the operating system and format the additional space using the ADAFRM utility.
- 2 Add necessary JCL/JCS to all Adabas nucleus and Adabas utility execution procedures.

For more information, please refer to *ADD: Add Dataset* in section *ADADBS: Database Services* in the *Adabas Utilities* documentation.

Modifying a Data Storage or Associator Dataset for Mainframe Databases

The Associator and Data Storage dataset sizes must be modified separately. It is not possible to modify both with a single operation.

▶ **To increase or decrease a DATA or ASSO dataset:**

Backup the database using the ADASAV utility. This step is optional, but recommended.

- 1 Select **Disk Storage** in the expanded tree-view of the selected database.

The default setting in the dropdown list of the detail-view is **Containers**. Allocation, size information, RABN range and device type of each extent are displayed in a table in detail-view.

- 2 Click **Increase** or **Decrease**.
- 3 In the resulting detail-view panel, enter the required parameters to increase/decrease either a DATA or ASSO dataset.
- 4 Click **OK** to increase/decrease the dataset or **Cancel** to return to the previous panel.



Caution: Increasing a DATA or ASSO dataset will result in the logical and physical sizes of the database being different.

After an *Increase* operation has completed, the nucleus session will be automatically terminated to allow for the necessary formatting of the Associator or Data Storage. If

the additional space is not formatted as required, an I/O error will occur once the nucleus tries to use it and the nucleus will ABEND.

- 5 Format the additional space using the ADAFRM utility.

For more information, please refer to *INCREASE: Increase Associator Data Storage* in section *ADADBS: Database Services* in the *Adabas Utilities* documentation.

Modifying a Data Storage or Associator Dataset for Open Systems Databases



Note: For detailed information please refer to the *Adabas for Open Systems* documentation.

▶ To add a new extent to an existing dataset:

- 1 Select **Disk Storage** in the expanded tree-view of the selected database.

The **Containers** of the database are displayed in detail-view.

- 2 Click **Links**.

- 3 In the detail-view panel, mark the dataset you want to modify and click **Add**.

The **Add Container** panel is displayed in detail-view.

- 4 Select the file type you want to add and enter the path. Click **OK** to confirm.

- 5 Format the additional space using the ADAFRM utility.

▶ To modify container settings:

- 1 Select **Disk Storage** in the expanded tree-view of the selected database.

The **Containers** of the database are displayed in detail-view.

- 2 Click **Links**.

- 3 In the detail-view panel, mark the dataset you want to modify and click **Modify**.

The **Modify Container Settings** panel is displayed in detail-view.

- 4 Edit the container settings according to your requirements. Click **OK** to confirm.

▶ To remove a container

- 1 Select **Disk Storage** in the expanded tree-view of the selected database.

The **Containers** of the database are displayed in detail-view.

- 2 Click **Links**.
- 3 In the detail-view panel, mark the dataset you want to modify and click **Remove**.

The **Remove Container** panel is displayed in detail-view.

- 4 Confirm that you want to remove the container you have marked by clicking **OK**.



Important: Please note that removing the last Associator or Data Storage container will make the database unusable.

Recovering Unused Storage

Space allocated for utility operations that ended abnormally remains unavailable unless it is intentionally recovered.

▶ To recover unused storage:

- 1 Select **Disk Storage** in the expanded tree-view of the selected database.
- 2 Select **Unused Storage** in the dropdown list in the detail-view panel.

The resulting table shows the device type where the unused blocks are located, the number of unused storage blocks and cylinders, and the range of unused block numbers.

- 3 Click **Recover** to recover unused storage space.

Reviewing Volume Distribution for a Mainframe Database

▶ To review the volume distribution for a mainframe database:

- 1 Select **Disk Storage** in the expanded tree-view of the selected database.
- 2 Select **Volume Distribution** in the dropdown list in the detail-view panel.

The **Volume Distribution** panel is displayed in detail-view. The file extents of the volume are shown indicating the RABN range and I/O count in each file.

34 Mainframe Database Maintenance Levels

▶ **To review maintenance levels for mainframe databases:**

- 1 Select an Adabas database in tree-view and expand it.
- 2 Select **Maintenance**.

The maintenance levels for the database are displayed in a table in detail-view.

The table displayed provides maintenance information about the loaded Adabas nucleus modules:

Information Type	Description
Module	The name of a nucleus module.
CSECT	The name of the CSECT associated with a nucleus module.
Date	The date of the last assembly.
Version	The Adabas version of the module.
Base	The base for ZAP numbering (corresponds with the SM level of the module).
Zaps	A list of zaps that have been applied to the module.

35 System Status

▪ Reviewing General System Status of Databases	168
▪ Resetting the System Status for Mainframe Databases	170
▪ CLOG/PLOG Statistics	171
▪ Reviewing Command Queue	172
▪ Reviewing Command Usage	172
▪ Reviewing File Usage	174
▪ Reviewing High Water Marks	174
▪ Reviewing Hold Queue	175
▪ Installed Products	176
▪ Reviewing Nucleus File Status	176
▪ Reviewing Online Status	177
▪ Reviewing PPT (Parallel Participant Table) Status	178
▪ Reviewing Resource Statistics	179
▪ Reviewing Thread Status	182
▪ User Profiles	183
▪ User Queue and User Queue Elements (UQEs) for Databases	184
▪ Reviewing Work Status	190
▪ Reviewing Work Pool Usage	191

Adabas Manager System Status provides a comprehensive overview of Adabas operation.



Note: A number of information displays are only available for mainframe databases. These are **Installed Products**, **Nucleus File Status**, **Online Status**, **PPT**, **Resource Statistics**, **User Profiles**, **Work Status** and **Work Pool Usage**.

Reviewing General System Status of Databases

► To review the general system status of an Adabas database:

- 1 Select an Adabas database in tree-view and expand it.
- 2 Select **System Status** in tree-view.

The system status information for the database is displayed in tables on the **System Status** panel in detail-view.

Read about *Database Monitoring and Tuning* in the *Adabas DBA Tasks* documentation for detailed information about system status statistics.

The system status information for a database is provided in sections, as described in this section:

- [Physical I/O Table](#)
- [General Statistics Table](#)
- [Pool and Queue Statistics Table](#)

Physical I/O Table

The physical I/O table is the first section on the **System Status** panel. It displays the number of physical reads and writes executed during the session to the Associator (ASSO), Data Storage (DATA), and Work (WORK) data sets. In addition, it displays the number of physical writes during the session to the data protection log (PLOG).

General Statistics Table

The general statistics table is the second section from the top of the **System Status** panel. It lists the following information about the database:

Property	Explanation
Logical Reads	The number of logical reads issued for the buffer pool.
Buffer Efficiency	The number of logical reads divided by the number of ASSO and DATA reads. The higher the value for buffer efficiency, the more efficient is buffer pool usage. If the value is less than 10, you may want to increase the size of the Adabas buffer pool (see the ADARUN LBP parameter).
Format Translations	Adabas read and update commands require a format buffer that specifies the fields to be read or updated. Adabas interprets the format buffer and translates it into an internal form and stores it in an internal format buffer pool. The internal format buffer is identified by a combination of user ID/ command ID. Format buffer translation is CPU-intensive.
Format Overwrites	When the internal format buffer pool becomes full, an existing entry must be overwritten to accommodate a new entry. If you have an excessive number of format overwrites, you may need to increase the size of the internal format buffer pool (see the ADARUN LFP parameter).
Throwbacks	The number of times a command could not be executed because the Adabas nucleus was waiting for an available ISN or Adabas work pool space. In these cases, the command is thrown back into the command queue for processing at a later point in time.
Thread Switches	The number of times in this session that a thread was automatically switched by Adabas during the processing of a single command. A thread switch occurs if I/O activity suspends command processing in an active thread. Ideally, each command should be processed to completion without a thread switch. The count of thread switches is thus an indication of the balance between the size of the I/O buffer and the number of threads allocated.
Number of Calls	Total number of calls issued to Adabas during the session.
Internal Autorestarts	The number of autorestarts performed during the session.
PLOG Switches	Applies only when dual protection logging is in effect; when one dataset is full, Adabas switches automatically to a second dataset and continues writing protection log records. In the meantime, data in the first dataset is copied off using ADARES PLCOPY to a sequential file, either manually or under the control of user exit 2, so that the dataset can be safely overwritten when the second dataset becomes full and Adabas switches automatically back to the first dataset. The system is most efficient when sufficient time is available to copy a full protection log off to the sequential dataset before the other protection log dataset becomes full. Protection logs are sized using the DUALPLD parameter to specify the device and thus the block size and the DUALPLS parameter to specify the number of blocks allowed in a log.
Buffer Flushes	The number of times the buffer pool has been flushed to Data Storage during the session. The buffer pool contains the most frequently used ASSO and DATA blocks and its purpose is to minimize physical I/O activity. The buffer pool should be as large as possible without causing undue paging by the operating system.

Pool and Queue Statistics Table

The pool and queue statistics section of the table is a high-water mark display. It shows the usage statistics and graphs the highest percentage used of selected pools and queues in the current session, as well as the date and time when the high point was reached. Values are displayed for the:

- user, command, and hold queues
- ISN list and sequential command tables
- format, redo, security, user file list, DTP transaction ID, unique descriptor, and work pools
- attached buffers.

In addition, the units and defined settings of the ADARUN parameters associated with establishing the sizes of these queues and pools is also listed in the table.

The high-water values are a good starting point when looking for a problem with limited buffer, pool, or queue space, or if you are looking for unused storage resources.

Resetting the System Status for Mainframe Databases

You have the following options for resetting the system status for mainframe databases:

▶ **To reset the pool and queue statistics display:**

- 1 Select an Adabas database in tree-view and expand it.
- 2 Select **System Status** in tree-view and right-click it.
- 3 Select **Reset Pool** on the drop-down menu.

The counters in the pool and queue statistics table are reset to zero.

▶ **To reset the counter:**

- 1 Select an Adabas database in tree-view and expand it.
- 2 Select **System Status** in tree-view and right-click it.
- 3 Select **Reset Counter** on the drop-down menu.

The counters in the general statistics table are reset to zero.

▶ **To reset all:**

- 1 Select an Adabas database in tree-view and expand it.

- 2 Select **System Status** in tree-view and right-click it.
- 3 Select **Reset All** on the drop-down menu.

All counters in the **System Status** tree are reset to zero.

CLOG/PLOG Statistics

▶ **To review the CLOG/PLOG statistics:**

- 1 Select an Adabas database in tree-view and expand it.
- 2 Select **CLOG/PLOG Statistics** in tree-view under the expanded **System Status**.

Two tables are displayed in the detail-view panel; in the first table the following command log properties are shown:

- the name of the CLOG that the nucleus is currently writing
- the size of a CLOG area in blocks
- the number of the last block written
- the percentage of the last block that was written (in graph form)
- the number of dual log switches since the nucleus started; and
- the date and time of the last CLOG switch

In the second table the following protection log properties are shown:

- the name of the PLOG that the nucleus is currently writing
- the size of a PLOG area in blocks
- the number of the last block written
- the percentage of the last block that was written (in graph form)
- the number of dual log switches since the nucleus started; and
- the date and time of the last PLOG switch

Forcing CLOG/PLOG

Using this option, you can immediately switch (by forcing an end-of-file) between dual or multiple command log (CLOG) or protection log (PLOG) files. Switching (that is, "toggling") changes from one CLOG or PLOG file to another.

▶ **To force a CLOG/PLOG switch:**

- 1 Select an Adabas database in tree-view and expand it.
- 2 Select **CLOG/PLOG Statistics** in tree-view under the expanded **System Status**.

Two tables are displayed in the detail-view panel, showing the command/protection log properties. If dual logging is active, **Force CLOG Switch** and **Force PLOG Switch** are displayed.

- 3 Click the action you want to perform.

Reviewing Command Queue

▶ **To review the command queue:**

- 1 Select an Adabas database in tree-view and expand it.
- 2 Select **Command Queue** in tree-view under the expanded **System Status**.

A list of the current commands in the queue and their status are listed.

Reviewing Command Usage

▶ **To review command usage:**

- 1 Select an Adabas database in tree-view and expand it.
- 2 Select **Command Usage** in tree-view under the expanded **System Status**.

A table is displayed in detail-view showing the total and average execution time of each Adabas command type issued during the current session and processed by the Adabas nucleus, and the total of all Adabas commands issued.

The following commands are listed:

Command	Description
A1/4	Update Existing Records
BT	Backout Transactions
CL	Close User Session
ET	End Transaction
E1/2	Delete Records
L1/4	Read by ISN
L2/5	Read Physical
L3/6	Read by Descriptor
L9	Read Descriptor
LF	Read Field Definition
N1/2	Add Records
OP	Open User Session
UC	Utility Commands
RC	Release Command id's
RE	Read ET User Data
REST	Other Commands
S1/4	Find Records
S2	Find records sorted
S5	Find Coupled ISN
S8	Process ISN List
S9	Sort ISN List
YCAL	Internal Communication
V1 to V4	Internal Communication
UCAL	Utility Commands

With Adabas version 8.2, statistics for the internally-used Y* and V* commands are also included in the **Command Usage** display. These commands are used internally by Adabas and Adabas add-on products and should not be used in direct calls in your applications because their use will result in errors.

► **To reset the command usage status display:**

- 1 Right-click **Command Usage** under **System Status**.
- 2 Select **Reset** on the drop-down menu.

The counters in the command usage status table are reset to zero.



Note: Selecting **Reset All** on the drop-down menu will reset all counters in the **System Status** tree to zero.

Reviewing File Usage

▶ **To review file usage for mainframe databases:**

- 1 Select an Adabas database in tree-view and expand it.
- 2 Select **File Usage** in tree-view under the expanded **System Status**.

A table is displayed in detail-view showing the total number of accesses for all files of the database used during the current session is displayed.

▶ **To reset the file usage status display for mainframe databases:**

- 1 Right-click **File Usage** under **System Status**.
- 2 Select **Reset** on the drop-down menu.

The counters in the file usage status table are reset to zero.



Note: Selecting **Reset All** on the drop-down menu will reset all counters in the **System Status** tree to zero.

▶ **To review file usage for open systems databases:**

- 1 Select an Adabas database in tree-view and expand it.
- 2 Select **File Usage** in tree-view under the expanded **System Status**.
- 3 In detail-view, click **Define Files**.
- 4 Enter your selection of files in the **Enter File List** box and click **OK**.

A table is displayed in detail-view showing the total number of accesses for all files of the database used during the current session is displayed.

Reviewing High Water Marks

▶ **To review high water marks:**

- 1 Select an Adabas database in tree-view and expand it.
- 2 Select **High Water Marks** in tree-view under the expanded **System Status**.

A table is displayed in detail-view showing the maximum percent of use of selected pools and queues in the current session, and the date and time when the high point was reached.

Values are displayed for the user, command, and hold queues; the ISN list and sequential command tables; the format, security, and work pools; and the attached buffers (NAB).

With Adabas version 8.2, the **High Water Marks** display includes data for Work parts 1, 2, and 3. This allows you to monitor the use of the Work data set.

These values are a good starting point when looking for a problem with limited buffer, pool, or queue space, or if you are looking for unused storage resources.

Reviewing Hold Queue

▶ To review the hold queue:

- 1 Select an Adabas database in tree-view and expand it.
- 2 Select **Hold Queue** in tree-view under the expanded **System Status**.

A list of ISNs currently in hold status is displayed along with the file in which it occurs, the related TID, CPU-ID, VM-ID and OPSYS-ID.

▶ To stop a user:


- 1 Click in the check box for the user you want to stop until a check mark appears.
- 2 Click **Stop User** to confirm your selection.

Filtering the Hold Queue Display

You can filter the hold queue display by file number or name and User ID.

Once you have specified a filter for the hold queue display, it remains in effect for all databases registered until you turn it off.

▶ To filter the hold queue display:

- 1 Access the hold queue display and in the upper left corner of the display, click the **Filter** button ().

The **Hold Queue** filter panel appears.

- 2 Enter the file name or number and/or the TID you want to filter the hold queue for. You can use an asterisk (*) as a wildcard character.


- 3 Verify that the **Turn Off Filter** check box does not have a check mark in it. If it does, click in the check box until the check mark is gone.
- 4 When all filters have been specified, click **OK** to save the filters or click **Cancel** to cancel the filter settings and revert to the prior settings (if any).

When you refresh the hold queue display, the filters are applied to the hold queue display.

▶ **To stop a user:**

- Select the user you want to stop from the resulting list and click **Stop User**.

▶ **To clear all filters on the hold queue display:**

- 1 Access the hold queue display and in the upper left corner of the display, click the **Filter** button ().

The **Hold Queue** filter panel appears.

- 2 Click in the **Turn Off Filter** check box until a check mark appears.
- 3 Click **OK** to save the filters or click **Cancel** to cancel the filter settings and revert to the prior settings (if any).

Installed Products

▶ **To review the installed products associated with the database:**

- 1 Select an Adabas database in tree-view and expand it.
- 2 Select **Installed Products** in tree-view under the expanded **System Status**.

The products associated with the database are displayed in detail-view.

Reviewing Nucleus File Status

▶ **To review the nucleus file status of a mainframe database:**

- 1 Select an Adabas database in tree-view and expand it.
- 2 Select **Nucleus File Status** in tree-view under the expanded **System Status**.

The nucleus file status of the database is displayed in a table in detail-view.

The following information is listed:

Property	Explanation
File Number	The Adabas file number.
Locking Nucleus ID	In an Adabas cluster environment, the file may be locked for exclusive use by another cluster nucleus. If this is the case and the file is in the nucleus file status table, this column for the file shows the ID of the nucleus that has exclusive control.
Access Count	Display the number of access 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.
Update Count	Display the number of 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.
Status	Indicates when the file is used for access only or for access and update. The Status column field indicates to what extent a nucleus can use a file on its own. If the requested use exceeds the given state, the nucleus must first communicate with the other nuclei in the cluster in order to upgrade the state.

Reviewing Online Status

The **Online Status** panel displays the online utility processes (inverting files, reordering files) currently in the database nucleus. For details on how to start the online utilities, see sections [Inverting Files](#) and [Reordering Files](#).

► To review the online status of a mainframe database:

- 1 Select an Adabas database in tree-view and expand it.
- 2 Select **Online Status** in tree-view under the expanded **System Status**.

The online status of the database is displayed in a table in detail-view.

The following information is listed:

Property	Explanation
ET-ID	The logical user ID assigned by the ADALOD LOAD parameter ETID (which defines owner IDs to all records being loaded into a multicient file).
Process Type	The current process type.
Status	The current status.
File Number	The Adabas file number.
Reorder Type	The reorder type set for the file.
Current RABN	The RABN assigned for the file.
Current ISN	The ISN allocated for the file.
Number of Throwbacks	The number of commands in the command queue for processing at a later point in time.

Reviewing PPT (Parallel Participant Table) Status

The parallel participant table (PPT) located in a database's Associator tracks all active Adabas nuclei in the cluster and displays additional information about a nucleus.

The information about each Adabas nucleus maintained in the PPT is used when ADARES PLCOPY is copying and merging all active protection logs in the cluster. It is also used during nucleus initialization to ensure that the dataset information required for autorestart and PLOG merge is available:

- If the previously used Work dataset and PLOG datasets are already in use by another nucleus, the nucleus will not start.
- If a different Work dataset is provided and the previously used Work dataset contains autorestart information, the nucleus will start but a warning will be printed.
- If different PLOGs or no PLOGs are provided and the previous session had PLOGs still to be copied, the nucleus will not start until those PLOGs have been copied/merged if PLOGRQ=FORCE is specified. If PLOGRQ=FORCE is not specified, the nucleus will start but a warning will be printed.

▶ To review the PPT of a mainframe database:

- 1 Select an Adabas database in tree-view and expand it.
- 2 Select **PPT** in tree-view under the expanded **System Status**.

The table in detail-view displays a list of nuclei participating in the cluster and information about the current status of each nucleus.

- 3 Click a nucleus ID in the table to display additional information about a nucleus.

For an Adabas cluster nucleus that has a nonzero nucleus ID, its entry in the PPT is displayed in the detail-view panel.

Reviewing Resource Statistics

Resource statistics provide information about file and command use either for all currently active users (general statistics) or for a single user (user statistics). The statistics are collected by starting a sampling period for which you have specified a time period in seconds (duration).

► To obtain statistics for all currently active mainframe users:

- 1 Select an Adabas database in tree-view and expand it.
- 2 Select **Resource Statistics** in tree-view under the expanded **System Status**.

The **Resource Statistics** panel is displayed in detail-view.

- 3 Specify the duration period in seconds (default is 60 seconds) and click **Start General Statistics**.

The **General Statistics** collection is started.

- 4 After the specified duration, click **Read General Statistics**.

The following command usage information for all currently active users is displayed in the detail-view panel:

- the total number of commands
- the command type
- the number of commands
- the number of commands (%)
- the average duration (milliseconds)

The following commands are listed:

Command	Description
A1/4	Update Existing Records
BT	Backout Transactions
CL	Close User Session
ET	End Transaction
E1/2	Delete Records
L1/4	Read by ISN
L2/5	Read Physical

Command	Description
L3/6	Read by Descriptor
L9	Read Descriptor
LF	Read Field Definition
N1/2	Add Records
OP	Open User Session
UC	Utility Commands
RC	Release Command id's
RE	Read ET User Data
REST	Other Commands
S1/4	Find Records
S2	Find records sorted
S5	Find Coupled ISN
S8	Process ISN List
S9	Sort ISN List
YCAL	Internal Communication
V1 to V4	Internal Communication
UCAL	Utility Commands

- 5 Choose **Files** in the selection box in this panel to display the file usage information for all currently active users.

The resulting table provides the following file usage details:

- the file number
- the number of commands
- the number of commands (%)

Click **Back** to return to the **Resource Statistics** panel.

► **To obtain statistics for a single mainframe user:**

- 1 Select an Adabas database in tree-view and expand it.
- 2 Select **Resource Statistics** in tree-view under the expanded **System Status**.

The **Resource Statistics** panel is displayed in detail-view.

- 3 Click **Select**.

The **User Queue** panel is displayed in detail-view.

- 4 Select a user ID to retrieve statistics for that user. This will return you to the **Resource Statistics** panel.
- 5 Specify the duration period in seconds (default is 60 seconds) and click **Start User Statistics**.
- 6 After the specified duration, click **Read User Statistics**.

The following command usage information for the user you selected is displayed in the detail-view panel:

- the total number of commands
- the command type
- the number of commands
- the number of commands (%)
- the average duration (milliseconds)

The following commands are listed:

Command	Description
A1/4	Update Existing Records
BT	Backout Transactions
CL	Close User Session
ET	End Transaction
E1/2	Delete Records
L1/4	Read by ISN
L2/5	Read Physical
L3/6	Read by Descriptor
L9	Read Descriptor
LF	Read Field Definition
N1/2	Add Records
OP	Open User Session
UC	Utility Commands
RC	Release Command id's
RE	Read ET User Data
REST	Other Commands
S1/4	Find Records
S2	Find records sorted
S5	Find Coupled ISN
S8	Process ISN List
S9	Sort ISN List

Command	Description
YCAL	Internal Communication
V1 to V4	Internal Communication
UCAL	Utility Commands

- 7 Choose **Files** in the selection box in this panel to display the file usage information for this user.

The resulting table provides the following file usage details:

- the file number
- the number of commands
- the number of commands (%)

Click **Back** to return to the **Resource Statistics** panel.

Reviewing Thread Status

▶ To review threads:

- 1 Select an Adabas database in tree-view and expand it.
- 2 Select **Threads** in tree-view under the expanded **System Status**.

The threads are listed along with their status (active or not active), the command type currently in process in each active thread, and the number of commands processed by each thread in the current session.

▶ To reset the thread status display:

- 1 Right-click **Threads** under **System Status**.
- 2 Select **Reset** on the drop-down menu.

The counters in the thread status table are reset to zero.



Note: Selecting **Reset All** on the drop-down menu will reset all counters in the **System Status** tree to zero.

User Profiles

Adabas allows you to retain user-related information from session to session in a user profile table. For each user, you can maintain

- a user priority to add "weight" to the normal, built-in priorities of Adabas commands issued by a specific user when they contend with other commands for Adabas database priority. The effect is to change the user's database access priority.
- nonactivity timeout values for access-only users (TNAA), ET logic users (TNAE), and EXU users (TNAX).
- transaction time limits for ET Logic users (TT).
- a time limit for executing a database query (Sx) command (TLSCMD).
- the number of ISNs allowed per TBI element (NSISN).
- the number of records that can be placed in hold status at one time (NISNHQ).
- the number of active command IDs allowed (NOCID).
- an owner ID for multicient support.

▶ To review the mainframe user profiles:

- 1 Select an Adabas database in tree-view and expand it.
- 2 Select **User Profiles** in tree-view under the expanded **System Status**.

The user profiles are displayed in a table.

▶ To modify a user profile:

- 1 In the **User Profiles** table displayed in detail-view, click the user ID for which you want to modify the profile.

The profile table for the user is displayed.
- 2 In the resulting profile table, enter the values you require in the text boxes.
- 3 Click **Modify** to confirm your changes or **Cancel** to return to the **User Profiles** table without modifying the profile.

▶ To delete a user profile:

- 1 In the **User Profiles** table displayed in detail-view, click the user ID you want to delete.

The profile table for the user is displayed.

- 2 Click **Delete** to delete the user ID or **Cancel** to return to the **User Profiles** table without deleting the profile.

Creating New User Profiles

In addition to modifying or deleting existing profiles, you can also add new user profiles.

▶ To create a new mainframe user profile:

- 1 Select an Adabas database in tree-view and expand it.
- 2 Select **System Status** in tree-view and right-click it.
- 3 Select **New Profile** on the drop-down menu.

A user profile table appears in detail-view.

- 4 Enter the user ID for the new profile and allocate the values you require in the text boxes (as default the boxes show the active parameter values).
- 5 Click **OK** to confirm the new user profile or **Cancel** and return to the **User Profiles** table.

The new user profile is now displayed in the table.

User Queue and User Queue Elements (UQEs) for Databases

Each user who calls Adabas is assigned a user queue element (UQE). The user queue stores the UQEs for all users of a given database.

User Queue

▶ To review the user queue for a database:

- 1 Select an Adabas database in tree-view and expand it.
- 2 Select **User Queue** in tree-view under the expanded **System Status**.

The user queue for the database appears on the **User Queue** panel in detail-view.

The user queue information is displayed in a table that provides the following information:

Statistics	Description												
TID	The unique internal identifier assigned to the user.												
ET-ID	Logical user ID assigned by the ADALOD LOAD parameter ETID (which defines owner IDs to all records being loaded into a multiclient file).												
Job Name	The job name associated with the user.												
User Type	The user type. Possible values include: <table border="1"> <tr> <td>ACC</td> <td>access-only user</td> </tr> <tr> <td>AOS</td> <td>AOS/AMA user</td> </tr> <tr> <td>ET</td> <td>ET user</td> </tr> <tr> <td>EXF</td> <td>exclusive file control access user</td> </tr> <tr> <td>EXU</td> <td>exclusive file update user</td> </tr> <tr> <td>UTI</td> <td>utility update control user</td> </tr> </table>	ACC	access-only user	AOS	AOS/AMA user	ET	ET user	EXF	exclusive file control access user	EXU	exclusive file update user	UTI	utility update control user
ACC	access-only user												
AOS	AOS/AMA user												
ET	ET user												
EXF	exclusive file control access user												
EXU	exclusive file update user												
UTI	utility update control user												
Status	The user status. Possible values include: <table border="1"> <tr> <td>ET</td> <td>The user is in "End Transaction" state.</td> </tr> <tr> <td>T-OUT</td> <td>The user has been timed out.</td> </tr> <tr> <td>TRANS</td> <td>The user is in "Open Transaction" state.</td> </tr> </table>	ET	The user is in "End Transaction" state.	T-OUT	The user has been timed out.	TRANS	The user is in "Open Transaction" state.						
ET	The user is in "End Transaction" state.												
T-OUT	The user has been timed out.												
TRANS	The user is in "Open Transaction" state.												
Last Activity	Period of inactivity (in seconds).												
Files	The file numbers of the Adabas files currently in use by the user. Note that no file list is provided for ET users.												

You can stop a specific user, all users from a specific job, all users using a specific file, or all inactive users. Any open transactions of the stopped users are backed out. You can also request that the stopped users be deleted. Note that EXF and UTI users are not stopped or deleted.

► **To stop specific users:**

- 1 Click in the check boxes for all users you want to stop until check marks appear.
- 2 Click **Stop User** to confirm your selection.



Caution: If Adabas is running with ADARUN OPENRQ=NO (specifying that users are not required to issue an OP as the first command of the session), stop specific users only if you are certain that the users to be deleted are no longer active. If a user with an open transaction is deleted, but then returns (by sending a command), no indication is given about the transaction backout. If the user continues the transaction, logical inconsistencies in the database could occur.

- 3 Click **Yes** to confirm or **No** to cancel. Both will return you to the previous panel.

The nucleus backs out any open transactions of the selected users and deletes them (purges the user queue elements).

▶ **To stop users using a specific file number:**

- 1 Right-click **User Queue** under the expanded **System Status**.
- 2 Select **Stop Users Using File** on the drop-down menu.

The **Stop Users using File** panel is displayed in detail-view.

- 3 Enter the file number for which you want users stopped.
- 4 If you want the stopped users deleted (if you want to purge their user queue elements), check the **Purge User Queue Elements** check box.



Caution: If Adabas is running with ADARUN OPENRQ=NO (specifying that users are not required to issue an OP as the first command of the session), stop users with **Purge User Queue Elements** checked only if you are certain that the users to be deleted are no longer active. If a user with an open transaction is deleted, but then returns (by sending a command), no indication is given about the transaction backout. If the user continues the transaction, logical inconsistencies in the database could occur.

- 5 Click **OK** to confirm or **Cancel**. Both will return you to the previous panel.

The nucleus backs out any open transactions of the affected users. If **Purge User Queue Elements** is checked, the nucleus also deletes the users (purges their user queue elements).

▶ **To stop inactive users:**

- 1 Right-click **User Queue** under the expanded **System Status**.
- 2 Select **Stop Inactive Users** on the drop-down menu.

The **Stop Inactive Users** panel is displayed in detail-view.

- 3 Enter the inactivity time interval (in seconds). Users who have not executed a command during this time limit will be stopped.
- 4 If you want the stopped users deleted (if you want to purge their user queue elements), check the **Purge User Queue Elements** check box.



Caution: If Adabas is running with ADARUN OPENRQ=NO (specifying that users are not required to issue an OP as the first command of the session), stop users with **Purge User Queue Elements** checked only if you are certain that the users to be deleted are no longer active. If a user with an open transaction is deleted, but then returns (by sending a command), no indication is given about the transaction backout. If the user continues the transaction, logical inconsistencies in the database could occur.

- 5 Click **OK** to stop the users or click **Cancel**. Both will return you to the previous panel.

The nucleus backs out any open transactions of the affected users. If **Purge User Queue Elements** is checked, the nucleus also deletes the users (purges their user queue elements).

▶ **To stop users by jobname:**

- 1 Right-click **User Queue** under the expanded **System Status**.
- 2 Select **Stop Users by Jobname** on the drop-down menu.

The **Stop Users by Jobname** panel is displayed in detail-view.

- 3 Enter the name of the job whose users you wish to stop.



Caution: If Adabas is running with ADARUN OPENRQ=NO (specifying that users are not required to issue an OP as the first command of the session), stop users by jobname only if you are certain that the users to be deleted are no longer active. If a user with an open transaction is deleted, but then returns (by sending a command), no indication is given about the transaction backout. If the user continues the transaction, logical inconsistencies in the database could occur.

- 4 Click **OK** to stop the users or **Cancel**. Both will return you to the previous panel.

The nucleus backs out any open transactions of the affected users and deletes them (purges their user queue elements).

User Queue Element (UQE)

▶ **To review the user queue element for a specific user:**

- 1 Select an Adabas database in tree-view and expand it.
- 2 Select **User Queue** under the expanded **System Status**.

The user queue for the database appears on the **User Queue** panel in detail-view.

- 3 On the **User Queue** panel in detail-view, click the TID of the user whose user queue element you want to review.


The UQE information is displayed in several tables that provide the following information.

Statistics	Description
TID	The unique internal identifier assigned to the user.
User ID	The Adabas-assigned user ID, if available.
Job Name	The job name associated with the user.
User Type	The user type. Possible values include:
	ACC access-only user

Statistics	Description										
	<table border="1"> <tr> <td>AOS</td> <td>AOS/AMA user</td> </tr> <tr> <td>ET</td> <td>ET user</td> </tr> <tr> <td>EXF</td> <td>exclusive file control access user</td> </tr> <tr> <td>EXU</td> <td>exclusive file update user</td> </tr> <tr> <td>UTI</td> <td>utility update control user</td> </tr> </table>	AOS	AOS/AMA user	ET	ET user	EXF	exclusive file control access user	EXU	exclusive file update user	UTI	utility update control user
AOS	AOS/AMA user										
ET	ET user										
EXF	exclusive file control access user										
EXU	exclusive file update user										
UTI	utility update control user										
Status	<p>The user status. Possible values include:</p> <table border="1"> <tr> <td>ET</td> <td>The user is in "End Transaction" state.</td> </tr> <tr> <td>T-OUT</td> <td>The user has been timed out.</td> </tr> <tr> <td>TRANS</td> <td>The user is in "Open Transaction" state.</td> </tr> </table>	ET	The user is in "End Transaction" state.	T-OUT	The user has been timed out.	TRANS	The user is in "Open Transaction" state.				
ET	The user is in "End Transaction" state.										
T-OUT	The user has been timed out.										
TRANS	The user is in "Open Transaction" state.										
Last Activity	The date and time activity last occurred by the user.										
CPU ID	The CPU ID of the user (a component of the Communication ID).										
VM ID	The VM ID of the user (a component of the Communication ID).										
OPSYS ID	The operating system ID of the user (a component of the Communication ID).										
User ID	The user ID (a component of the Communication ID).										
Communication ID (hexadecimal)	<p>The 28-byte hexadecimal communication ID of the user.</p> <p>Note: CPU ID, VM ID, OPSYS ID and User ID are all components of the <i>Communication ID</i> and guarantee its network-wide uniqueness even in heterogeneous environments. The individual IDs are dependent on the operating system the <i>Adabas</i> call has been issued from.</p>										
Hold Queue Limit	The maximum number of held ISNs the hold queue can contain from one user. For detailed information about the NISNHQ parameter, see the <i>ADARUN Control Statement</i> in the <i>Adabas Operations</i> documentation.										
Max. parallel CIDs per User	The total number of CIDs (command IDs) allowed per user, and therefore the maximum number of table elements the user has available at any point in time. For detailed information about the NQCID parameter, see the <i>ADARUN Control Statement</i> in the <i>Adabas Operations</i> documentation.										
Max. ISNs per TBI Element	The maximum number of ISNs kept in a TBI (table of ISNs) element in memory (list of resulting ISN lists). For detailed information about the NSISN parameter, see the <i>ADARUN Control Statement</i> in the <i>Adabas Operations</i> documentation, .										
Max. Time of Nonactivity	The maximum length of time, in seconds, that the user can be inactive before inactive processing occurs. For more information about Adabas timeouts, refer to your Adabas documentation.										
Time Limit for Sx Commands	The maximum length of time, in seconds, permitted for Sx commands by this user.										
No. of ISNs currently held	The number of ISNs currently held by the user.										
No. of CIDs currently in use	The number of CIDs (command IDs) currently in use. For detailed information about the NQCID parameter, see the <i>ADARUN Control Statement</i> in the <i>Adabas Operations</i> documentation.										

Statistics	Description
No. of Calls	The number of calls issued by the user during the user session.
No. of I/Os	The number of I/Os initiated by the user during the user session.
Priority from ET/CP File	The priority of the user.
Start Session	The start time of the user session.
Start Transaction	The start time of the transaction.
File List	The file numbers of the files currently in use by the user. Note that no file list is provided for ET users.

You can return to the **User Queue** panel by clicking **Back** or you can stop this user. When users are stopped, they are deleted. Any open transactions of the stopped/deleted users are backed out. Note that EXF and UTI users are not stopped or deleted.

 **Caution:** If Adabas is running with ADARUN OPENRQ=NO (specifying that users are not required to issue an OP as the first command of the session), stop them only if you are certain that the users to be deleted (stopped) are no longer active. If a user with an open transaction is deleted, but then returns (by sending a command), no indication is given about the transaction backout. If the user continues the transaction, logical inconsistencies in the database could occur.

▶ **To stop this user:**

- Click **Stop User**.

Filtering the User Queue Display

You can filter the user queue display by last activity (elapsed time in seconds), ET-ID, user type, job name and user ID.

Once you have specified a filter for the user queue display, it remains in effect for all databases registered until you turn it off.

▶ **To filter the user queue display:**

- 1 Access the user queue display and in the upper left corner of the display, click the **Filter** button ().

The **User Queue** filter panel appears.

- 2 Enter the values you want to filter the user queue for. You can use an asterisk (*) as a wildcard character.
- 3 Verify that the **Turn Off Filter** check box does not have a check mark in it. If it does, click in the check box until the check mark is gone.

- 4 When all filters have been specified, click **OK** to save the filters or **Cancel** to cancel the filter settings and revert to the prior settings (if any).

When you refresh the user queue display, the filters are applied to the file list display.

▶ **To stop a user:**

- Select the user you want to stop from the resulting list and click **Stop User**.

▶ **To clear all filters on the user queue display:**

- 1 Access the user queue display and in the upper left corner of the display, click the **Filter** button ()

The **Hold Queue** filter panel appears.

- 2 Click in the **Turn Off Filter** check box until a check mark appears.
- 3 Click **OK** to save the filters or **Cancel** to cancel the filter settings and revert to the prior settings (if any).

Reviewing Work Status

▶ **To review the work status:**

- 1 Select an Adabas database in tree-view and expand it.
- 2 Select **Work Status** under the expanded **System Status**.

Work Status displays the WORK area sizes (in blocks) for the

- data protection area (Work part 1; ADARUN LP parameter)
- area used for intermediate ISN lists (Work part 2; ADARUN LWKP2 parameter)
- area used for resulting ISN lists (Work part 3)
- distributed transaction processing area (Work part 4; only available if Adabas Transaction Manager is installed on a separate nucleus in your system)

Reviewing Work Pool Usage

▶ **To review the work pool:**

- 1 Select an Adabas database in tree-view and expand it.
- 2 Select **Work Pool Usage** under the expanded **System Status**.

Work pool statistics are displayed including the total length set in the ADARUN parameter LWP, the length of the used and unused parts of the pool, and the length of the longest single unused part. These numbers can be used to tune the work pool length for the next session.

VII

Space Calculation for Mainframe Databases

36

Space Calculation for Mainframe Databases

- Estimating Associator Space 197
- Estimating Data Storage Space 197
- Estimating Space for the DD/FILEA Sequential Dataset 198
- Estimating the Sort Dataset 198
- Estimating the Temp Dataset 199
- Estimating the Work Dataset 200

The space calculation function is a planning tool for adding new components or recalculating existing space requirements. Each calculation provides a block or cylinder estimate according to information you provide. In general, you must provide

- the maximum estimated record count
- average number of DE values
- average descriptor, compressed record, or normal record length
- estimated padding factor
- device type where the Adabas component being estimated resides

In many cases, the results are "best guess" estimates; other than a device type, no defaults are assumed. Because no values are actually changed by the Space Calculation function, unrealistic estimates cause no harm.

▶ **To use the space calculation function:**

- 1 Select an Adabas database in tree-view and expand it.
- 2 Select **Space Calculation** in tree-view.

In detail-view, a dropdown list allows you to choose between the following components:

- ASSO - Address Converter
- ASSO - Normal/Upper Index
- DATA
- DDFILEA
- SORT - ADAINV
- SORT - ADALOD load
- SORT - ADALOD update
- TEMP - ADAINV
- TEMP - ADALOD load/update
- TEMP - ADALOD delete
- WORK

Estimating Associator Space

You can calculate two Associator component values: the address converter (AC) space, or the normal (NI) and upper (UI) index space.

Address Converter Space

Address converter space is based on the device type and the estimated number of records in the related Data Storage file.

▶ **To calculate address converter space:**

- 1 In the **Space Calculation** detail-view, select **ASSO - Address Converter** in the dropdown list.
- 2 Enter your requirements in the table and click **OK** to calculate.

Normal/Upper Index Space

The index values are calculated for a single descriptor, requiring you to estimate such things as the average descriptor length, the number of multiple descriptors you expect to have, the total number of unique descriptor values for that field, an Associator padding factor, and a device type if other than the default.

▶ **To calculate normal/upper index space:**

- 1 In the **Space Calculation** detail-view, select **ASSO - Normal/Upper Index** in the dropdown list.
- 2 Enter your requirements in the table and click **OK** to calculate.

Estimating Data Storage Space

The DATA storage calculation is based on values you provide for estimated maximum record count, the average length of a compressed record, a Data Storage padding factor, and device type. Results are specified in both blocks and cylinders.

▶ **To calculate data storage space:**

- 1 In the **Space Calculation** detail-view, select **DATA** in the dropdown list.
- 2 Enter your requirements in the table and click **OK** to calculate.

Estimating Space for the DD/FILEA Sequential Dataset

The space required for the DD/FILEA sequential dataset is calculated when it is used with the ADAORD utility. (The dataset is also used with the ADALOD utility.)

▶ **To calculate space for the DD/FILEA sequential dataset:**

- 1 In the **Space Calculation** detail-view, select **DDFILEA** in the dropdown list.
- 2 Enter your requirements in the table and click **OK** to calculate.

Estimating the Sort Dataset

The storage needed on SORT for the utility function chosen is calculated.

The topic is organized as follows:

- [ADAINV Sort](#)
- [ADALOD LOAD Sort](#)
- [ADALOD UPDATE Sort](#)

ADAINV Sort

▶ **To calculate the space needed on SORT for the ADAINV utility function:**

- 1 In the **Space Calculation** detail-view, select **SORT - ADAINV** in the dropdown list.
- 2 Enter your requirements in the table and supply the password (if required). Click **OK** to calculate.

ADALOD LOAD Sort

For the ADALOD LOAD calculation, the default number of records is MAXISN rather than TOPISN as it is for the ADAINV function.

▶ **To calculate the space needed on SORT for the ADALOD LOAD utility function:**

- 1 In the **Space Calculation** detail-view, select **SORT - ADALOD Load** in the dropdown list.
- 2 Enter your requirements in the table and supply the password (if required). Click **OK** to calculate.

ADALOD UPDATE Sort

For the ADALOD UPDATE calculation, the default number of records is 0.

▶ **To calculate the space needed on SORT for the ADALOD UPDATE utility function:**

- 1 In the **Space Calculation** detail-view, select **SORT - ADALOD Update** in the dropdown list.
- 2 Enter your requirements in the table and supply the password (if required). Click **OK** to calculate.

Estimating the Temp Dataset

The storage needed on TEMP for the utility function chosen is calculated.

The topic is organized as follows:

- [ADAINV Temp](#)
- [ADALOD LOAD/UPDATE Temp](#)
- [ADALOD DELETE Temp](#)

ADAINV Temp

▶ **To calculate the space needed on TEMP storage for the ADAINV utility function:**

- 1 In the **Space Calculation** detail-view, select **TEMP - ADAINV** in the dropdown list.
- 2 Enter your requirements in the table and supply the password (if required). Click **OK** to calculate.

ADALOD LOAD/UPDATE Temp

▶ **To calculate the space needed on TEMP storage for the ADALOD LOAD/UPDATE utility function:**

- 1 In the **Space Calculation** detail-view, select **TEMP - ADALOD Load/Update** in the dropdown list.
- 2 Enter your requirements in the table and supply the password (if required). Click **OK** to calculate.

ADALOD DELETE Temp

▶ To calculate the space needed on SORT for the ADALOD DELETE utility function:

- 1 In the **Space Calculation** detail-view, select **TEMP - ADALOD Delete** in the dropdown list.
- 2 Enter your requirements in the table and supply the password (if required). Click **OK** to calculate.

Estimating the Work Dataset

The Work dataset requires the most estimating. Although many initial values may be arbitrary, keep a record of them to ensure that subsequent tuning of the Work parameters has a realistic basis. Results comprise block estimates for the three parts of the Work area. A total of these values in blocks and cylinders is also provided.

▶ To calculate work dataset space:

- 1 In the **Space Calculation** detail-view, select **WORK** in the dropdown list.
- 2 Enter your requirements in the table and click **OK** to calculate.

Index

A

- access a protected nucleus, 28
- access the Adabas Manager area, 11
- active parameters, 139
- Adabas Manager
 - area access, 11
 - cache and session administration, 17
 - configuration, 15
 - session and cache administration, 17
 - window layout, 13
- ADAINV sort, 198
- ADAINV temp storage, 199
- ADALOD DELETE temp storage, 200
- ADALOD LOAD sort, 198
- ADALOD LOAD/UPDATE temp storage, 199
- ADALOD UPDATE sort, 199
- add existing
 - open systems databases, 36
- add new fields
 - Field Definition Table (FDT), 94
- address converter space, 197
- allocate/deallocate file space, 101
- associator space, 197

C

- cache and session administration, 17
- cache Facility
 - display files modified by utilities, 157
- cache parameters, 153
- cache session summary, 154
- Caching Facility, 151
 - cache session summary, 154
 - define a file cache, 152
 - maintain cache parameters, 153
- change field parameters, 91
- change the columns in the mainframe database overview display, 47
- checkpoints, 69
- CLOG statistics, 171
- CLOG/PLOG
 - force, 172
- columns
 - change in file list display, 78
- command queue, 172
- command usage, 172
- configuration, 15
- CR option, 114

- create
 - open systems databases, 37
- create new LOB file (Adabas V8 only), 119
- creating new file, 111

D

- data storage space, 197
- database
 - open systems
 - rename, 36
- database system status, 167
 - CLOG/PLOG statistics, 171
 - command queue, 172
 - command usage, 172
 - file usage, 174
 - high water marks, 174
 - hold queue, 175
 - installed products, 176
 - nucleus file status, 176
 - online status, 177
 - parallel participant table (PPT), 178
 - reset, 170
 - resource statistics, 179
 - thread status, 182
 - user profiles, 183
 - new user profiles, 184
 - user queue, 184
 - user queue elements, 184
 - work pool, 191
 - work status, 190
- databases
 - active parameters, 139
 - add existing
 - open systems databases, 36
 - expanded files
 - maintenance, 109
 - FDT, 89
 - files with a critical number of extends (Adabas V7 only), 85
 - files with extent list saturation (Adabas V8 only), 87
 - list, 33
 - list files, 77
 - open systems
 - create, 37
 - delete, 39
 - register mainframe databases, 35
 - review file properties, 81
 - review mainframe database DIBs, 53
 - review maintenance levels for mainframe databases, 165

- review open system database Utility Communication Block, 55
- select, 45
- start
 - databases, 41
- startup parameters, 139
- stop
 - databases, 41
- system status, 167
- unregister, 35
- unregister mainframe databases, 35
- dataset
 - add, 160
 - increase/decrease for a mainframe database, 161
 - modify for an open systems database, 162
- date-time option DT, 113
- DBA Workbench functionality, 4
- DD/FILEA space, 198
- define a file cache, 152
- delete
 - open systems databases, 39
- delete files, 104-105
- deleting/undeleting fields, 91
- Delta Save, 157
- Delta Save Facility, 155
- descriptors, 92
- disk storage, 159
 - add dataset, 160
 - increase/decrease dataset for a mainframe database, 161
 - modify dataset for an open systems database, 162
 - unused storage, 163
 - volume distribution for mainframe databases, 163

E

- expanded files, 109
- extent list saturation (Adabas V8 only)
 - database files, 87
- extents
 - database files (Adabas V7 only), 85

F

- FDT (see Field Definition Table (FDT))
 - referential integrity constraints, 90
- Field Definition Table (FDT), 89
 - add new fields, 94
- field parameters
 - change, 91
 - logically delete descriptors, 92
 - set and release descriptors, 92
- fields
 - delete/undelete, 91
- file
 - create, 111
- file cache, 152
- file list display
 - change columns, 78
- file modification, 103
- file parameters
 - modify, 95
- file space
 - allocate file space, 101
 - deallocate file space, 101

- file usage, 174
- filter
 - hold queue display, 175
 - user queue display, 189
- force CLOG/PLOG, 172

H

- high water marks, 174
- hold queue, 175
- hold queue display
 - filter, 175

I

- installed products, 176
- invert files, 106

L

- list
 - database files, 77
 - databases, 33
- LOB fields (Adabas V8 and open systems only), 97
- LOB file (Adabas V8 and open systems only), 97
- LOB file (Adabas V8 only)
 - create, 119
- lock files, 104
- logging in
 - SMH, 11
- logically delete descriptors, 92

M

- mainframe database overview display
 - change the columns, 47
- maintain cache parameters, 153
- maintenance levels, 165
- modify file parameters, 95
- modify files, 103
- MUPEX (Adabas V8 only), 97

N

- new file, 111
- new user profiles, 184
- non-activity timeout, 29
- normal/upper index space, 197
- nucleus file status, 176

O

- online status, 177

P

- parallel participant table (PPT), 178
- parameters
 - Predict, 94
- passwords, 29
- PLOG statistics, 171
- PPT (see parallel participant table)
- Predict parameters, 94

protected nucleus
 access, 28
 terminate a session, 28

R

Record spanning (Adabas V8 and open systems only), 97
 referential integrity constraints, 90
 register
 mainframe databases, 35
 release and set descriptors, 92
 rename files, 105
 rename open systems database, 36
 renumber files, 105
 reorder files, 104
 reset system status, 170
 resource statistics, 179
 review
 database file FDT, 89
 database file properties, 81
 mainframe database DIBs, 53
 mainframe database maintenance levels, 165
 open systems database Utility Communication Block, 55

S

SAF security support, 25
 select
 databases, 45
 mainframe databases, 46
 open systems databases, 49
 select a database
 mainframe, 46
 open systems, 49
 session and cache administration, 17
 set and release descriptors, 92
 sort space, 198
 space calculation
 ADAINV sort, 198
 ADAINV temp storage, 199
 ADALOD DELETE temp storage, 200
 ADALOD LOAD sort, 198
 ADALOD LOAD/UPDATE temp storage, 199
 ADALOD UPDATE sort, 199
 address converter, 197
 associator, 197
 data storage, 197
 DD/FILEA, 198
 normal/upper index, 197
 sort, 198
 temp storage, 199
 work dataset, 200
 space calculation for mainframe databases, 195
 starting
 databases, 41
 stop
 databases, 41
 System Management Hub
 logging in, 11

T

temp storage space, 199
 terminate a session with a protected nucleus, 28

thread status, 182
 time zone
 TZ option, 113
 time zone option TZ, 113
 timestamp
 DT option, 113
 trace version, 19

U

uncouple files, 106
 unlock files, 104
 unregistering
 mainframe databases, 35
 unused storage
 disk storage, 163
 user profiles, 183
 user queue, 184
 user queue display
 filter, 189
 user queue elements, 184

V

volume distribution
 review for mainframe databases, 163

W

window layout, 13
 work dataset space, 200
 work pool usage, 191
 work status, 190

