

Adabas Manager

Adabas Manager Documentation

Version 7.5.1

April 2008

Adabas Manager

This document applies to Adabas Manager Version 7.5.1 and to all subsequent releases.

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

Copyright © Software AG 2008. All rights reserved.

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

Table of Contents

| | |
|---|----|
| 1 Adabas Manager Documentation | 1 |
| 2 Introduction | 3 |
| 3 Working with Adabas Manager | 5 |
| 4 Accessing the Adabas Manager Area | 7 |
| 5 Adabas Manager Window Layout | 9 |
| 6 Adabas Manager Session and Cache Administration | 11 |
| 7 Leaving Adabas Manager | 13 |
| 8 Adabas Manager Security | 15 |
| 9 SAF Security Support | 17 |
| 10 Working with a Protected Nucleus | 19 |
| Accessing a Protected Nucleus | 20 |
| Terminating a Session with a Protected Nucleus | 20 |
| Passwords | 21 |
| Non-activity Timeout | 21 |
| 11 Databases | 23 |
| 12 Listing Databases | 25 |
| 13 Selecting a Database | 27 |
| Changing the Columns in the Database Overview Display | 29 |
| 14 Databases that Cannot Be Reached | 33 |
| 15 Stopping a Database | 35 |
| 16 Registering or Unregistering a Database | 37 |
| Registering Databases | 38 |
| Unregistering Databases | 38 |
| 17 Reviewing AdabasDatabase DIBs | 41 |
| 18 Checkpoints | 43 |
| Types of Adabas Users | 45 |
| Checkpoint Names Starting with 'SYN' | 45 |
| Checkpoint Numbers for 'SYN' Names | 46 |
| 19 Database Files | 49 |
| 20 Listing Database Files | 51 |
| Displaying the List of Files | 52 |
| Changing the Columns in the File List Display | 52 |
| 21 Reviewing File Properties | 55 |
| 22 Displaying Files with a Critical Number of Extents (Adabas Version 7 only) | 59 |
| 23 Files with Extent List Saturation (Adabas Version 8 only) | 61 |
| 24 Field Definition Table (FDT) | 63 |
| Changing Field Parameters | 65 |
| Releasing Descriptors | 66 |
| Predict Parameters | 66 |
| Adding New Fields in the Field Definition Table (FDT) | 67 |
| 25 Modifying File Parameters | 69 |
| Changing Padding Factors, Maximum Allocation, and Record Length | 70 |
| Selecting File Options | 71 |

| | |
|--|-----|
| Record Spanning, MUPEX, LOB Fields/File (Adabas V8 only) | 71 |
| 26 Allocating/Deallocating File Space | 73 |
| 27 Modifying Files | 75 |
| Locking/Unlocking Files | 76 |
| Reordering Files | 76 |
| Deleting Files | 77 |
| Renumbering Files | 77 |
| Renaming Files | 77 |
| Uncoupling Files | 78 |
| Inverting Files | 78 |
| 28 Maintaining Expanded Files | 81 |
| 29 Creating a New File | 83 |
| 30 Creating a New LOB File (Adabas Version 8 only) | 89 |
| 31 Session Monitoring | 91 |
| 32 Active Parameters | 93 |
| 33 Disk Storage | 97 |
| Adding a Data Storage or Associator Dataset | 98 |
| Modifying a Data Storage or Associator Dataset | 99 |
| Recovering Unused Storage | 100 |
| Reviewing Volume Distribution | 100 |
| 34 Database Maintenance Levels | 101 |
| 35 System Status | 103 |
| Reviewing General System Status of Databases | 104 |
| Resetting the System Status | 106 |
| CLOG/PLOG Statistics | 107 |
| Reviewing Command Queue | 108 |
| Reviewing Command Usage | 108 |
| Reviewing File Usage | 110 |
| Reviewing High Water Marks | 110 |
| Reviewing Hold Queue | 111 |
| Installed Products | 112 |
| Reviewing Nucleus File Status | 112 |
| Reviewing Online Status | 113 |
| Reviewing PPT (Parallel Participant Table) Status | 114 |
| Reviewing Resource Statistics | 114 |
| Reviewing Thread Status | 118 |
| User Profiles | 118 |
| User Queue and User Queue Elements (UQEs) for Databases | 120 |
| Reviewing Work Status | 126 |
| Reviewing Work Pool Usage | 126 |
| 36 Space Calculation | 127 |
| Estimating Associator Space | 129 |
| Estimating Data Storage Space | 129 |
| Estimating Space for the DD/FILEA Sequential Dataset | 130 |
| Estimating the Sort Dataset | 130 |

Estimating the Temp Dataset 131

Estimating the Work Dataset 132

Index 133

1 Adabas Manager Documentation

This document tells you how to install, configure, and use the Adabas Manager version 7.5.1, a web-based graphical user interface (GUI) you can use to perform administrative tasks for Adabas.

It 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 7.5.1 provides

- browser-based Adabas administration and monitoring;
- 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. |
| • Installing Adabas Manager | describes the prerequisites and the environment setup required to install and run Adabas Manager; describes the Adabas Manager installation procedure. |
| • 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. |
| • Databases | introduces database services: listing databases, registering and unregistering databases, and selection of a database. |
| • 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. |

2 Introduction

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.

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); and
- control logging for the Adabas Delta Save Facility

Enterprise Edition and Standard Edition

The Adabas Manager Enterprise Edition supports the following functionality, which is not available in the Standard Edition:

- 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

3 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 Session and Cache Administration**
- **Leaving Adabas Manager**

4 Accessing the Adabas Manager Area

Logging in to Adabas Manager involves first logging in 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 SMH:**

- 1 During installation, an entry for Adabas Manager has been created in your Start Menu. Select **Software AG Adabas Manager 7.5.1 -> Administration** from your program list.

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.

► **To enter your Adabas Manager license file:**

- 1 Select Adabas Manager in tree-view.

The Adabas Manager **Welcome** panel is displayed in detail-view.

- 2 Click on the **License** button at the bottom of the page.


The **License Files** panel is displayed in detail-view.

- 3 Click on the **Modify** button below the table or **Cancel** to return to the **Welcome** panel.

The table in detail-view now allows you to either enter the path for your Adabas Manager license file or browse to the location of the file.

- 4 Click **OK** or **Cancel** to return to the **Welcome** panel.



Note: The Adabas Manager Documentation has been installed with the product. You can either view it by selecting **Software AG Adabas Manager 7.5.1 -> Documentation -> Overview** from your program list, or call up the appropriate page by clicking on the **Help** button () in the top right corner of your Adabas Manager panel.

5

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 on) 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 on the plus sign (+) to the left of the high-level entry.



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.

Once you have highlighted an item in the tree, context-specific commands are listed in the **Commands** dropdown list in the screen title bar.

You can also see these commands by right-clicking on the highlighted menu item. Only this option is described in the documentation, but both options are always available.

All items in the tree show a **Refresh** option in their dropdown 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 on it (or view the **Commands** dropdown in the screen title bar).
- 2 Select **Refresh** on the dropdown menu.

The information displayed in the detail-view is updated.

6

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 on it.
- 2 Select **Display Cache** on the dropdown 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 on it.
- 2 Select **Clear Cache** on the dropdown 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/7.5.1/* 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.

7 Leaving Adabas Manager

▶ **To leave Adabas Manager:**

- 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.

8 Adabas Manager Security

Adabas Manager version 7.5.1 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](#).

The information is organized in the following topic:

- **Adabas SAF Security Support**
- **Working with a Protected Nucleus**

9 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*.

10

Working with a Protected Nucleus

| | |
|--|----|
| ■ Accessing a Protected Nucleus | 20 |
| ■ Terminating a Session with a Protected Nucleus | 20 |
| ■ Passwords | 21 |
| ■ Non-activity Timeout | 21 |

The information is organized under the following headings:

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.

11 Databases

The information is organized under the following headings:

- **Listing Databases**
- **Selecting a Database**
- **Databases that Cannot Be Reached**
- **Stopping a Database**
- **Registering or Unregistering a Database**
- **Reviewing Adabas Database DIBs**

12

Listing Databases

▶ To list the databases that have been registered to Adabas Manager:

- Once you have accessed the Adabas Manager area of SMH, expand the area and then select and expand **Adabas Databases**.
 - The databases and their IDs are listed in tree-view under **Adabas Databases**. In addition, the database IDs, names, Adabas versions, and required logins are shown in detail-view.

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

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

13

Selecting a Database

- Changing the Columns in the Database Overview Display 29

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 the database is protected or not supported.
- Red indicates that the database is inactive.

When **Adabas Databases** is highlighted in the tree-view, 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 databases a login with password is required.



Note: The databases can be accessed from this table by marking the appropriate check boxes and clicking **Login**. For protected databases, 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.

▶ 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** If the 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 database, the dropdown menu will offer you two additional menu items; **Logout** will terminate the session with the database, **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. |

| Property | Description |
|----------------------------------|--|
| 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. |
| Event Replication | Indicates whether or not data replication is active for an Adabas 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 on the **Properties** button ()

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 on 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 on 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 the **Up** and **Down** buttons 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 the **Save** button to save the column setting definition and return to the main database overview display, which will show the column settings you selected.

Or:

Click on the **Cancel** button 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 on the **Properties** button ()

The **Database Overview** panel appears.

- 2 Click on the **Restore Default** button.

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

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

14

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 the **Retry** button will reach the database only if the error has been corrected.

15

Stopping a Database

► **To stop an Adabas database:**

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

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

- Normal session termination (ADAEND)
 - Cancel session immediately (CANCEL)
 - Stop session (HALT)
- 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 on the **Cancel** button to continue working with this database.

16

Registering or Unregistering a Database

| | |
|---------------------------------|----|
| ■ Registering Databases | 38 |
| ■ Unregistering Databases | 38 |

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

This chapter is organized as follows:

Registering Databases

► To register a database to Adabas Manager:

- 1 Select an Adabas database in tree-view.
- 2 Right-click on the database name in tree-view.
- 3 Select **Register DB** on the dropdown menu.

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

- 4 Enter the ID and name of the database to be registered.
- 5 Click the **Register** button to complete the registration, or the **Cancel** button to quit the registration window without registering the database. Press the **Help** button 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.

- 6 Click on the **Close** button 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 on the database name in tree-view.
- 3 If you are unregistering an Adabas database, select **Unregister DB** on the dropdown menu.

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

- 4 Click the **Unregister** button to complete the unregistration, or the **Cancel** button to quit the unregistration window without unregistering the database. Press the **Help** button 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 on the **Close** button in detail-view to complete the transaction.

17

Reviewing AdabasDatabase DIBs

You can review the DIBs (Data Integrity Blocks) of Adabas databases as follows:

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

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

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

- The name of the service (ADANUC)
- The name of the job that started the Adabas database
- The date and time the Adabas database was started
- The type of database
- 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 on the database name in tree-view.
- 3 Select **Display DIB** on the dropdown 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 on the database name in tree-view.
- 3 Select **Display DIB** on the dropdown menu.

The DIB information appears in detail-view.

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

18

Checkpoints

| | |
|--|----|
| ■ Types of Adabas Users | 45 |
| ■ Checkpoint Names Starting with 'SYN' | 45 |
| ■ Checkpoint Numbers for 'SYN' Names | 46 |

This chapter is organized as follows:

► **To review the checkpoints for an Adabas database:**

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

The checkpoints for the database are displayed in detail-view.

The table columns are described as follows:

| Column | Explanation |
|----------------------|---|
| Name | The checkpoint name; that is, the checkpoint identifier. In the case of a user nonsynchronized checkpoint, the name is supplied by the user program. Checkpoint names starting with 'SYN' are reserved for the Adabas nucleus and utilities and are described in the section Checkpoint Names Starting with 'SYN' . |
| Type | The checkpoint number. For checkpoint names starting with 'SYN', the checkpoint numbers are provided in the section Checkpoint Numbers for 'SYN' Names . |
| Date | The date and time the checkpoint was taken. |
| PLOG No. | The PLOG that the nucleus is currently writing. |
| BlockNo. | The block number of the data protection log in which the checkpoint was written. |
| VOLSER | Either the number of the sequential protection (DD/SIBA) log or 'DUAL' if dual logging is being used. |
| User | The type of Adabas user that set the checkpoint. These types are described in the section Types of Adabas Users . |
| JobName | The name of the job that created the checkpoint. |
| Extended Information | |

For more information about checkpoints, see the section *ADAREP* in the *Adabas Utilities* documentation.

► **To delete the checkpoints:**

- 1 Right-click on **Checkpoints** in tree-view.
- 2 Select **Delete Checkpoints** on the dropdown menu.

The **Delete Checkpoints** panel appears in detail-view.

- 3 The available options are:

| | |
|--|--|
| Delete Checkpoints from (and older): | Remove all checkpoint file entries up to the date you specify in the field. |
| Delete Checkpoints older than XX days: | Remove all checkpoint file entries older than the number of days you specify in the field. |

Enter a date or a value for the number of days and click **Delete** to confirm.

The **Delete Checkpoints Success** panel is displayed.

- 4 Click **OK** to return to the **Checkpoints** display panel.

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 | 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 | 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 | 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 | A checkpoint made at the beginning of online ADASAV execution (SAVE database function). |
| SYN2 | A checkpoint made at the end of online ADASAV execution (SAVE database function). |
| SYN4 | A checkpoint made at the beginning of online ADASAV execution (SAVE files operation). |
| SYN5 | A checkpoint made at the end of online ADASAV execution (SAVE files operation). |

Checkpoint Numbers for 'SYN' Names

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 |

| Type | Name | Origin | Description |
|------|------|--------|--|
| 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 |
| 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 | PLCOPY function successfully completed |
| 28 | SYNV | ADARES | PLCOPY dual protection log |
| 29 | SYNV | ADARES | CLCOPY dual command log |
| 2A | SYNP | ADARES | PLCOPY MERGE function successfully completed |
| 2A | SYNV | ADARES | PLCOPY 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 |

| Type | Name | Origin | Description |
|------|------|--------|---|
| 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 |
| 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 |

19 Database Files

The information is organized under the following headings:

- Listing Database Files
- Reviewing File Properties
- Displaying Files with a Critical Number of Extents
- Files with Extent List Saturation (Adabas version 8 only)
- Field Definition Table (FDT)
- Modifying File Parameters
- Allocating/Deallocating File Space
- Modifying Files
- Maintaining Expanded Files
- Creating a New File
- Creating a New LOB File (Adabas version 8 only)

20

Listing Database Files

| | |
|---|----|
| ■ Displaying the List of Files | 52 |
| ■ Changing the Columns in the File List Display | 52 |

This chapter is organized as follows:

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 on the **Go to** button.

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 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 on the **Properties** button ().

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 on 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 on 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 the **Up** and **Down** buttons 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 the **Save** or **Save As** button to save the column setting definition.

Or:

Click on the **Cancel** button 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 dropdown list on the upper right side of the file list display.

21

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.

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. |
| 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. |
| 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). |

| Property | Description |
|--|--|
| 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 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 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. |
| has LOB Fields (Adabas version 8 only) | Indicates whether the file contains one or more LB fields. |
| LOB File (Adabas version 8 only) | Specifies the file number of the associated LOB file. |
| Base File (Adabas version 8 LOB file only) | Specifies the file number of the associated base file. |

| Property | Description |
|---|--|
| 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). |
| User 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. |
| Replication Target ID (RPLTARGETID) | The Event Replicator target ID for the Adabas file to be replicated (if replication is active). |
| Replication Primary Key (RPLKEY) | The primary key for replication (if replication is active). |
| Replication Active | Indicates whether or not data replication is active for the Adabas database. |
| Replication of Before Images of Updates (RPLDSBI) | Indicates whether to activate the collection of before images of data storage during an update command to the file (if replication is active). |

22

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 on it.
- 3 Select **Files with a Critical Number of Extents** on the dropdown 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 dropdown 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.

23

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 on it.
- 3 Select **Files with Extent List Saturation** on the dropdown 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 dropdown menu.

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

24 Field Definition Table (FDT)

| | |
|---|----|
| ■ Changing Field Parameters | 65 |
| ■ Releasing Descriptors | 66 |
| ■ Predict Parameters | 66 |
| ■ Adding New Fields in the Field Definition Table (FDT) | 67 |

This chapter is organized as follows:

► **To review the Field Definition Table (FDT) for a file:**

- 1 List the files in a database, as described in [Listing Database Files](#).
- 2 If an icon appears in the Fields column of the list of files table in detail-view, click on the icon.

The FDT for the file appears in detail-view. The table provides the following information:

| Field | Explanation | |
|--------------|---------------------------------|--|
| Level | the field level | |
| Name | the field name | |
| Length | the field length | |
| Format | the data type of the field: | |
| | A | alphanumeric |
| | B | binary |
| | F | fixed point |
| | P | packed decimal |
| | G | floating point |
| | U | unpacked decimal |
| | W | wide-character |
| Option | DE | descriptor |
| | FI | fixed storage |
| | LA | long alphanumeric |
| | LB | large object field |
| | MU | multiple-value field |
| | NC | null/not counted |
| | NN | null not allowed |
| | NU | null value suppression |
| | NV | not converted (alpha and wide-character fields) |
| | PE | periodic group |
| | UQ | unique descriptor |
| | XI | index (occurrence) number excluded from UQ in PE |
| Predict Name | equivalent Predict name, if any | |

When available, the Special Descriptor Table (SDT) is also displayed. It shows field information about all sub-/super-/hyperdescriptors, collation descriptors, phonetic descriptors, and sub-/superfields for the file.

The table provides the following information:

| Column | Explanation | |
|--------------|---|--|
| Type | SUB | subfield/subdescriptor |
| | SUPER | superfield/superdescriptor |
| | PHON | phonetic descriptor |
| | HYPER | hyperdescriptor |
| | COL | collation descriptor |
| Name | the field name | |
| Length | the field length | |
| Option | DE | descriptor field |
| | FI | fixed point |
| | LA | long alphanumeric |
| | MU | multiple-value field |
| | NC | null/not counted |
| | NN | null not allowed |
| | NU | null value suppression |
| | NV | not converted (alpha and wide-character fields) |
| | PE | periodic group |
| | UQ | unique descriptor |
| | XI | index (occurrence) number excluded from UQ in PE |
| Parent | the name of the parent field for this descriptor. | |
| Predict Name | equivalent Predict name, if any | |

Changing Field Parameters

► To change field parameters:

- 1 In the Field Definition Table (FDT), click on 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. 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.

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.

▶ To release a descriptor:

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

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

- 2 Click on **Release Descriptor** to release the descriptor and return to the FDT.
- 3 If a Special Descriptor Table (SDT) is shown, click on the name of the descriptor you want to release.

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

- 4 Click on **Release Descriptor**.
- 5 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.


Predict Parameters

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 on 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 on the **Properties** button ()

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 on **Add**.

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.

25

Modifying File Parameters

| | |
|---|----|
| ■ Changing Padding Factors, Maximum Allocation, and Record Length | 70 |
| ■ Selecting File Options | 71 |
| ■ Record Spanning, MUPEX, LOB Fields/File (Adabas V8 only) | 71 |

► **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 on **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 version 8 databases, the following additional parameters may be specified:

- Record Spanning
- MUPEX
- Has LOB Fields
- LOB File

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

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 V8 only)

Record Spanning

With Adabas 8, 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 on 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

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 (LB) Fields* in the *Adabas 8 for Mainframes* documentation.

26 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 on the icon.

The space allocation for the file appears in detail-view.
- 3 Click on **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.

27

Modifying Files

| | |
|---------------------------------|----|
| ▪ Locking/Unlocking Files | 76 |
| ▪ Reordering Files | 76 |
| ▪ Deleting Files | 77 |
| ▪ Renumbering Files | 77 |
| ▪ Renaming Files | 77 |
| ▪ Uncoupling Files | 78 |
| ▪ Inverting Files | 78 |

The information is organized under the following headings:

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)'  
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.

28

Maintaining Expanded Files

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 on it.
- 3 Select **Maintain Expanded Files** on the dropdown 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 on the **Insert** button.

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 on it.
- 3 Select **Maintain Expanded Files** on the dropdown 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 on the **Remove** button.

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

29 Creating a New File

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).

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 on 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 on it.
- 3 Select **New File** on the dropdown 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 on 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 on the **Find** button 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 on the **Find** button, 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. 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.

- 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.

- 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 on 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 on 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 |

| Property | Description |
|------------------------------|--|
| | 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 |
| | 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.

30

Creating a New LOB File (Adabas Version 8 only)

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 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 on it.
- 3 Select **New LOB File** on the dropdown 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 on the **Find** button 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 on the **Find** button, 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.

31

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**
- **Disk Storage**
- **Database Maintenance Levels**
- **System Status**

32

Active Parameters

► To review the ADARUN 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 ADARUN parameters for the database are displayed in tables in detail-view:

- Pools and Queues
- Time Limits
- Logging
- Options

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

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 Value Buffer | Enable value buffer logging. |
| Log ISN Buffer | Enable ISN buffer logging. |
| Log I/O Activity | Enable I/O activity logging. |

| Parameter | Action |
|--|--|
| 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). |
| Event Replication | Enable data replication. |

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

- 1 Click the **Modify** command button.

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
- Max 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 Value Buffer
- Log ISN Buffer
- Log I/O Activity
- Log User Exit B Data
- Merge Cluster Command Logs automatically

Options (on/off)

The following options can be checked:

- Log AOS/DBS Update
- Async by Vol-Ser
- 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.

33

Disk Storage

| | |
|--|-----|
| ■ Adding a Data Storage or Associator Dataset | 98 |
| ■ Modifying a Data Storage or Associator Dataset | 99 |
| ■ Recovering Unused Storage | 100 |
| ■ Reviewing Volume Distribution | 100 |

► **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.

In detail-view, a selection box allows you to choose between a **Containers**, **Unused Storage**, and **Volume Distribution** display. The default is set to **Containers**.

► **To view RABNs for a database:**

- 1 Select an Adabas database in tree-view and expand it.
- 2 Select **Disk Storage** in tree-view and right-click on it.
- 3 Select **Display ASSO/DATA Block (RABN)** on the dropdown 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.

The information is organized under the following headings:

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 selection box 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 on **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.

- 5 Allocate the dataset with the operating system and format the additional space using the ADAFRM utility.
- 6 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

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 selection box 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 on **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.

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 selection box 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 on **Recover** to recover unused storage space.

Reviewing Volume Distribution

► **To review the volume distribution for a database:**

- 1 Select **Disk Storage** in the expanded tree-view of the selected database.
- 2 Select **Volume Distribution** in the selection box 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 Database Maintenance Levels

► To review maintenance levels:

- 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 | 104 |
| ■ Resetting the System Status | 106 |
| ■ CLOG/PLOG Statistics | 107 |
| ■ Reviewing Command Queue | 108 |
| ■ Reviewing Command Usage | 108 |
| ■ Reviewing File Usage | 110 |
| ■ Reviewing High Water Marks | 110 |
| ■ Reviewing Hold Queue | 111 |
| ■ Installed Products | 112 |
| ■ Reviewing Nucleus File Status | 112 |
| ■ Reviewing Online Status | 113 |
| ■ Reviewing PPT (Parallel Participant Table) Status | 114 |
| ■ Reviewing Resource Statistics | 114 |
| ■ Reviewing Thread Status | 118 |
| ■ User Profiles | 118 |
| ■ User Queue and User Queue Elements (UQEs) for Databases | 120 |
| ■ Reviewing Work Status | 126 |
| ■ Reviewing Work Pool Usage | 126 |

Adabas Manager System Status provides a comprehensive overview of Adabas operation.

This chapter covers the following topics:

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 three tables, 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 table 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 table 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). |

| Property | Explanation |
|-----------------------|--|
| 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 table is a high-water mark display. It is the third table from the top of the **System Status** panel. 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, replication, 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

You have the following options for resetting the system status:

▶ 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 on it.
- 3 Select **Reset Pool** on the dropdown 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 on it.
- 3 Select **Reset Counter** on the dropdown 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 on it.
- 3 Select **Reset All** on the dropdown 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** buttons 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 |

► **To reset the command usage status display:**

- 1 Right-click on **Command Usage** under **System Status**.
- 2 Select **Reset** on the dropdown menu.

The counters in the command usage status table are reset to zero.



Note: Selecting **Reset All** on the dropdown menu will reset all counters in the **System Status** tree to zero.

Reviewing File Usage

▶ **To review file usage:**

- 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:**

- 1 Right-click on **File Usage** under **System Status**.
- 2 Select **Reset** on the dropdown menu.

The counters in the file usage status table are reset to zero.



Note: Selecting **Reset All** on the dropdown menu will reset all counters in the **System Status** tree to zero.

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).

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 the **Stop User** button 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 on 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 on the **OK** button to save the filters or click on the **Cancel** button 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 the **Stop User** button.

► 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 on 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 on the **OK** button to save the filters or click on the **Cancel** button 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 the 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. |

| Property | Explanation |
|--------------|--|
| 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

► To review the online status of the 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 multiclient 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:**

- 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 on 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 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 on the **Start General Statistics** button.

The **General Statistics** collection is started.

- 4 After the specified duration, click on the **Read General Statistics** button.

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 |
| 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 |

| Command | Description |
|----------|------------------------|
| 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 on the **Back** button to return to the **Resource Statistics** panel.

► **To obtain statistics for a single 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 on the **Select** button.

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 on the **Start User Statistics** button.
- 6 After the specified duration, click on the **Read User Statistics** button.

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 |
| 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 on the **Back** button 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 on **Threads** under **System Status**.
- 2 Select **Reset** on the dropdown menu.

The counters in the thread status table are reset to zero.



Note: Selecting **Reset All** on the dropdown 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 multiclient support.

▶ **To review the 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 on 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 on the **Modify** button 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 on the user ID you want to delete.
The profile table for the user is displayed.
- 2 Click on the **Delete** button 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 user profile:**

- 1 Select an Adabas database in tree-view and expand it.
- 2 Select **System Status** in tree-view and right-click on it.
- 3 Select **New Profile** on the dropdown 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 on the **OK** button 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. This section covers the following topics:

- [User Queue](#)
- [User Queue Element \(UQE\)](#)
- [Filtering the User Queue Display](#)

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: | |
| | 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: | |
| | ET | The user is in "End Transaction" state. |
| | T-OUT | The user has been timed out. |
| | TRANS | The user is in "Open Transaction" state. |

| Statistics | Description |
|---------------|---|
| 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 the **Stop User** button 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 on the **Yes** button to confirm or click on the **No** button 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 on **User Queue** under the expanded **System Status**.
- 2 Select **Stop Users Using File** on the dropdown 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 on the **OK** button to confirm or click on the **Cancel** button to 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 on **User Queue** under the expanded **System Status**.
- 2 Select **Stop Inactive Users** on the dropdown 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 on the **OK** button to stop the users or click on the **Cancel** button to 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 on **User Queue** under the expanded **System Status**.
- 2 Select **Stop Users by Jobname** on the dropdown 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 on the **OK** button to stop the users or click on the **Cancel** button to 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 on 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: <table> <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> <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). | | | | | | | | | | | | |

| Statistics | Description |
|--------------------------------|---|
| 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) | The 28-byte hexadecimal communication ID of the user. 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. |
| 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. |
| 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 the **Back** button 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 the **Stop User** button.

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 on 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 on the **OK** button to save the filters or click on the **Cancel** button 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 the **Stop User** button.

▶ **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 on 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 on the **OK** button to save the filters or click on the **Cancel** button 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.

36

Space Calculation

| | |
|--|-----|
| ■ Estimating Associator Space | 129 |
| ■ Estimating Data Storage Space | 129 |
| ■ Estimating Space for the DD/FILEA Sequential Dataset | 130 |
| ■ Estimating the Sort Dataset | 130 |
| ■ Estimating the Temp Dataset | 131 |
| ■ Estimating the Work Dataset | 132 |

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 selection box 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

This chapter is organized as follows:

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 selection box.
- 2 Enter your requirements in the table and click the **OK** button 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 selection box.
- 2 Enter your requirements in the table and click the **OK** button 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 selection box.
- 2 Enter your requirements in the table and click the **OK** button 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 selection box.
- 2 Enter your requirements in the table and click the **OK** button 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 selection box.
- 2 Enter your requirements in the table and supply the password (if required). Click the **OK** button 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 selection box.
- 2 Enter your requirements in the table and supply the password (if required). Click the **OK** button 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 selection box.
- 2 Enter your requirements in the table and supply the password (if required). Click the **OK** button 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 selection box.
- 2 Enter your requirements in the table and supply the password (if required). Click the **OK** button 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 selection box.
- 2 Enter your requirements in the table and supply the password (if required). Click the **OK** button 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 selection box.
- 2 Enter your requirements in the table and supply the password (if required). Click the **OK** button 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 selection box.
- 2 Enter your requirements in the table and click the **OK** button to calculate.

Index

A

- active parameters, 93
- ADAINV sort, 130
- ADAINV temp storage, 131
- ADALOD DELETE temp storage, 132
- ADALOD LOAD sort, 130
- ADALOD LOAD/UPDATE temp storage, 131
- ADALOD UPDATE sort, 131
- adding new fields
 - Field Definition Table (FDT), 67
- address converter space, 129
- allocating/deallocating file space, 73
- associator space, 129

C

- changing field parameters, 65
- changing the columns in the database overview display, 29
- checkpoints, 43
- CLOG statistics, 107
- CLOG/PLOG
 - forcing, 108
- columns
 - changing in file list display, 52
- command queue, 108
- command usage, 108
- creating new file, 83
- creating new LOB file (Adabas V8 only), 89

D

- data storage space, 129
- database files, 49
- database overview display
 - changing the columns, 29
- database system status, 103
 - CLOG/PLOG statistics, 107
 - command queue, 108
 - command usage, 108
 - file usage, 110
 - high water marks, 110
 - hold queue, 111
 - installed products, 112
 - nucleus file status, 112
 - online status, 113
 - parallel participant table (PPT), 114
 - reset, 106

- resource statistics, 114
- thread status, 118
- user profiles, 118
 - new user profiles, 119
- user queue, 120
- user queue elements, 120
- work pool, 126
- work status, 126
- databases
 - active parameters, 93
 - expanded files
 - maintenance, 81
 - FDT, 63
 - files with a critical number of extends (Adabas V7 only), 59
 - files with extent list saturation (Adabas V8 only), 61
 - listing, 25
 - listing files, 51
 - registering, 37
 - reviewing Adabas database DIBs, 41
 - reviewing file properties, 55
 - reviewing maintenance levels, 101
 - startup parameters, 93
 - system status, 103
 - unregistering, 37
- dataset
 - add, 98
 - increase/decrease, 99
- DD/FILEA space, 130
- deleting files, 76, 77
- disk storage
 - add dataset, 98
 - increase/decrease dataset, 99
 - unused storage, 100
 - volume distribution, 100

E

- expanded files, 81
- extent list saturation (Adabas V8 only)
 - database files, 61
- extents
 - database files (Adabas V7 only), 59

F

- FDT (see Field Definition Table (FDT))
- Field Definition Table (FDT), 63
 - adding new fields, 67
- field parameters
 - changing, 65

- releasing a descriptor, 66
- file
 - creating, 83
- file list display
 - changing columns, 52
- file modification, 75
- file parameters
 - modifying, 69
- file space
 - allocating file space, 73
 - deallocating file space, 73
- file usage, 110
- filtering
 - hold queue display, 111
 - user queue display, 125
- forcing CLOG/PLOG, 108

H

- high water marks, 110
- hold queue, 111
- hold queue display
 - filtering, 111

I

- installed products, 112
- inverting files, 78

L

- listing
 - database files, 51
 - databases, 25
- LOB fields (Adabas V8 only), 71
- LOB file (Adabas V8 only), 71
 - creating, 89
- locking files, 76
- logging in
 - SMH, 7

M

- maintenance levels, 101
- modifying file parameters, 69
- modifying files, 75
- MUPEX (Adabas V8 only), 71

N

- new file, 83
- new user profiles, 119
- normal/upper index space, 129
- nucleus file status, 112

O

- online status, 113

P

- parallel participant table (PPT), 114
- parameters

- Predict, 66
- PLOG statistics, 107
- PPT (see parallel participant table)
- Predict parameters, 66

R

- Record spanning (Adabas V8 only), 71
- registering
 - databases, 37
- releasing a descriptor, 66
- renaming files, 77
- renumbering files, 77
- reordering files, 76
- reset system status, 106
- resource statistics, 114
- reviewing
 - Adabas database DIBs, 41
 - database file FDT, 63
 - database file properties, 55
 - database maintenance levels, 101

S

- selecting a database, 27
- sort space, 130
- space calculation, 127
 - ADAINV sort, 130
 - ADAINV temp storage, 131
 - ADALOD DELETE temp storage, 132
 - ADALOD LOAD sort, 130
 - ADALOD LOAD/UPDATE temp storage, 131
 - ADALOD UPDATE sort, 131
 - address converter, 129
 - associator, 129
 - data storage, 129
 - DD/FILEA, 130
 - normal/upper index, 129
 - sort, 130
 - temp storage, 131
 - work dataset, 132
- stopping a database, 35
- System Management Hub
 - logging in, 7

T

- temp storage space, 131
- thread status, 118

U

- uncoupling files, 78
- unlocking files, 76
- unregistering
 - databases, 37
- unused storage
 - disk storage, 100
- user profiles, 118
- user queue, 120
- user queue display
 - filtering, 125
- user queue elements, 120

V

volume distribution
 reviewing, 100

W

work dataset space, 132
work pool usage, 126
work status, 126

