

ADABAS D

QUESTIONS AND ANSWERS



**Manual Order Number: ESD611-001ALL**

This document is applicable to ADABAS D Version 6.1.1 PE and to all subsequent releases, unless otherwise indicated in new editions or technical newsletters.

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

Readers' comments are welcomed. Comments may be addressed to the Documentation Department at the address on the back cover.

© April 14, 2025, SOFTWARE AG, Germany & SOFTWARE AG of North America, Inc.

All rights reserved

Printed in the Federal Republic of Germany

The SOFTWARE AG documentation often refers to numerous hardware and software products by their trade names. In most, if not all cases, these designations are claimed as trademarks or registered trademarks by their respective companies.

## 1 General Information

### 1.1 The Name of the DBMS

*What is the database management system called?*

The database management system is called ADABAS D.

### 1.2 Versions

*What is the current version?*

The current version is ADABAS D Version 6.1.1.

### 1.3 The DBMS Supplier

*Who is the supplier of the DBMS?*

The supplier of ADABAS D is SOFTWARE AG.

### 1.4 Language Versions

*Does a German-language version exist?*

User interfaces are available in the languages English and German. Additional languages can easily be implemented. If required, this will be done by the SOFTWARE AG companies in the particular country.

## 1.5 Operating Systems

*Which operating systems can the DBMS be used on?*

ADABAS versions are available for the following operating systems:

- HP-UX
- SUN Solaris
- IBM AIX
- DEC OSF/1
- SNI Sinix
  
- Windows NT
- Windows 95
- Windows 3.1 (clients only)
- OS/2

## 2 Features of the DBMS

### 2.1 Components of the DBMS

*Which components does ADABAS comprise?*

Aside from the database server, the most important individual components are:

- |                        |                       |
|------------------------|-----------------------|
| - Operating tool:      | CONTROL               |
| - Loading tool:        | LOAD                  |
| - Administration tool: | DOMAIN                |
| - End user tool:       | QUERY                 |
| - End user tool:       | EASY                  |
| - End user tool:       | OfficePlus            |
| - Programming tool:    | ADABASIC              |
| - Programming tool:    | C/C++ Precompiler     |
| - Programming tool:    | COBOL Precompiler     |
| - Programming tool:    | CALL Interface (ODBC) |
| - Connectivity tool:   | ODBC driver           |

### CONTROL

*What is CONTROL and what is it used for?*

CONTROL is the menu-driven interface of ADABAS used for database operation (installation and configuration, restart/shutdown, backup and recovery measures, operation monitoring and control, performance monitoring and control).

**LOAD**

*What is LOAD and what is it used for?*

LOAD is the ADABAS component used for exporting and importing database contents and the database catalog.

**DOMAIN**

*What is DOMAIN and what is it used for?*

DOMAIN is the menu-driven interface of ADABAS used for database administration. It provides options for the display, creation, and maintenance of all database objects (tables, users, views, DB procedures, triggers, etc.), as well as information about the usage of these database objects within programs and other database objects. Thus DOMAIN also provides Data Dictionary facilities.

**QUERY**

*What is QUERY and what is it used for?*

QUERY is the interactive SQL interface of ADABAS used to execute any SQL statements. The set of commands of the interactive report generator REPORT which is integrated in QUERY can be used to modify the layout of result tables.

**EASY**

*What is EASY and what is it used for?*

EASY is the end user tool of ADABAS for users without knowledge of SQL. Database requests and modifications are handled according to the QUERY-BY-EXAMPLE PRINCIPLE; i.e., the user can work with ADABAS, guided by menus, without knowledge of SQL.

**OfficePlus**

*What is OfficePlus and what is it used for?*

OfficePlus allows for an optimal integration of ADABAS into Microsoft Office components such as Word, Excel, Access, and Visual Basic. Thus it is possible, e.g., to link a SELECT statement directly to a Word data file or data source or to an Excel spreadsheet and to provide the current SQL data by clicking a button.

**ADABASIC**

*What is ADABASIC and what is it used for?*

ADABASIC is a programming language compatible with Visual Basic. It can be used to create DB procedures, triggers, and DB functions out of the Visual Basic development environment, and to test and operate them on both the client and the server.

## **Precompilers**

*What are precompilers and what are they used for?*

For writing SQL application programs, there are precompilers for the programming languages C/C++ and COBOL. They translate the SQL statements embedded by EXEC SQL into a subroutine interface.

## **CALL Interface**

*What is the CALL Interface and what is it used for?*

In addition to the precompilers, ADABAS provides a CALL Interface to allow SQL applications to be written in programming languages other than C/C++ or COBOL. Some application programmers prefer an SQL-API to a precompiler concept. The SQL-API supported by ADABAS conforms to that of Microsoft ODBC, the only industrial standard in this field. Thus an ODBC-compatible interface is also made available on UNIX systems, for example.

## **ODBC Driver**

*Which Windows end user and development tools does ADABAS support?*

The ODBC driver allows ADABAS to be accessed from any Windows tools with an ODBC interface (e.g., Access, Excel, MS Query, Visual Basic, PowerBuilder, SQLWindows). The ODBC driver is provided in the form of a Windows DLL.



## 2.2 Display-oriented Interfaces

*Which of the cited components are display-oriented interfaces?*

With the exception of the precompilers, the CALL Interface, and the ODBC driver, all the components are display-oriented.

DOMAIN and OfficePlus require a Windows PC as client. QUERY, EASY, and LOAD are provided as both alpha clients and Windows clients.

## 2.3 Interactive SQL

*Does an interactive interface exist for SQL?*

QUERY is the interactive interface to ADABAS which offers the full range of the SQL language.

## 2.4 Report Generator

*Does a report generator exist that can be used to format the results of database queries?*

The report generator REPORT allows the user to format the results either interactively step by step or at a later time in batch mode. REPORT is integrated in QUERY and EASY and can also be called out of programming languages.

## 2.5 SQL

*Which database language is supported and what standards does this language satisfy?*

ADABAS supports the following SQL standards or dialects (SQLMODEs):

ANSI SQL-92 (Entry Level)  
ORACLE V7  
DB2 V3.1  
ADABAS

The applicable SQLMODE can be selected for each individual database session. This means that application programs written for ORACLE or DB2 usually can use ADABAS without having to be modified.

## 2.6 Extended SQL Functions of ADABAS

*Which extended SQL functions does the database system have?*

Extended SQL functions of ADABAS are, for example:

- The primary key concept
- The data types DATE, TIME, TIMESTAMP, BOOLEAN, LONG
- Referential integrity
- Updatable join views
- Outer join
- DB procedures
- Triggers
- DB functions
- Scrollable cursors
- Snapshots

## Referential Integrity

*What does referential integrity mean?*

Referential integrity means the monitoring and maintenance of consistency conditions existing between specific tables. If, e.g., entries of the employees table refer to the department table, it can be ensured that these references do not lead to empty columns, i.e., to non-existent departments. It can also be stipulated that, e.g., departments cannot be deleted as long as there are employee entries referring to them.

## Updatable Join View

*Are views based on several tables updatable?*

If certain conditions are satisfied when join views based on several tables are created, it is possible to perform insert, update, and delete operations on the underlying base tables using this join view.

## Outer Join

*How can rows from several tables be output in a single result set without a join condition being satisfied?*

Unlike "normal" joins, outer joins also include rows in the result set which do not satisfy the join condition. In this case, the corresponding columns are set to NULL values.

## DB Procedure

*How can the database-intensive parts of application systems be transferred to the database server?*

ADABAS allows the definition of DB procedures. These can be executed on both the client or server, as the user chooses. The programming language ADABASIC, which comprises the full Visual Basic language, is provided for this purpose. DB procedures reduce the number of interactions between application and database server and are an important means to performance improvement, especially in client-server configurations. They are also interesting from the maintenance point of view, because centralizing applications logic in a database object means that a modification to the logic of a DB procedure becomes effective in all the applications in which the DB procedure is used.

## Trigger

*Is it possible to initiate the execution of a trigger implicitly?*

While DB procedures have to be called explicitly, thus representing a layer above the normal SQL level, triggers are activated implicitly by means of the SQL statements INSERT, UPDATE, and DELETE. Triggers are special DB procedures which provide a kind of "user exit" for these statements. In addition to the normal effects of these SQL statements, triggers can be used to enable further actions. Triggers are typically used to formulate integrity rules for these modifying statements in a procedural way or to execute derived actions, e.g., for the logging of applications.

## DB Function

*Is it possible to extend the built-in functions of ADABAS?*

DB functions can be defined using the same language elements that are used for DB procedures. Such DB functions can be applied, e.g., in the SELECT statement, in addition to the provided standard built-in functions. They allow data editing or qualification functions specific to a particular application to be put into action on the database server, thus helping to avoid unnecessary data transfer from the server to the client or application.

## Scrollable Cursor

*Is it possible to page back in SELECT results?*

Many SQL systems only support paging forward in SELECT results. Buffering on application level and further SELECT statements are needed to provide also a backward-paging mechanism in output lists of applications systems. ADABAS supports both forward- and backward-paging in SELECT results.

## Snapshot

*Does ADABAS support replicated data?*

Frequently, only parts of data from productive applications systems are needed decentrally, with, e.g., a state of the current day being sufficient. ADABAS supports such a replication of data by the special database object "Snapshot". Snapshots are defined like views, but keep the data defined by them in the form of read-only copies, which are usually stored on another computer. The update of a snapshot is initiated by an explicit REFRESH statement. For simple snapshots (parts of one table), only the modifications are considered; for complex snapshots (joins of tables), the current data is transferred completely.

## 2.7 Precompiler Concept

*How are SQL statements executed in application programs? Are 'application plans' created in this process and stored in the database?*

When an application program is executed, ADABAS - unlike other database systems - compiles each SQL statement upon its first execution (parse phase), storing it as a kind of temporary 'application plan' in the database. To accelerate the processing, the internal representation is accessed for each subsequent execution of the statement (execute phase). This internal representation, however, is only stored for the duration of the application program (i.e., until the end of the session).

This facilitates both the adaptation of the 'application plan' to newly created or dropped indexes and the response to modified user privileges. This technique also enables ADABAS to respond immediately to these types of changes in the database structure during the execution of an application without having to terminate and restart the application.

## 2.8 Data Dictionary

*Which data dictionary functions does ADABAS feature?*

The SQL catalog of ADABAS can be accessed via SELECT statements issued on system tables (views). In addition to these functions, DOMAIN administers all database objects and provides usage records which are maintained implicitly. Especially the usage of database objects within application programs is recorded in this way. Together with the documentation of the database objects and the possibility of including further, user-defined objects, DOMAIN provides the full functionality of a Data Dictionary.

## 2.9 Database Statistics

*Which database statistics are kept?*

Statistics about the database configuration, log status, structure and size of tables, indexes, log, and database can be accessed via special system tables.

*Which monitoring information is available?*

Database monitoring comprises a great number of database events, such as the number of SQL statements issued, the number of logical and physical read and write operations, as well as the hit rate in the different caches.





## 3 Characterization of the DBMS

### 3.1 Database Objects

*Which database objects does ADABAS know?*

ADABAS supports the following database objects:

- Database
- Table
- Synonym
- Index
- View
- Snapshot
- Trigger
- User
- Usergroup
- Privileges
- Domain
- DB procedure
- DB function

### **3.2 The Elements of the Database Language SQL**

*Which elements does the database language SQL consist of?*

The database language SQL comprises the following elements:

DDL	data definition language
DML	data manipulation language
DCL	data control language

#### **The Data Definition Language**

*What is the data definition language?*

The data definition language is used for creating, updating, and dropping database objects: e.g., tables, views, and users.

#### **The Data Manipulation Language**

*What is the data manipulation language?*

The data manipulation language is used for inserting, updating, deleting, and reading the contents of all the database objects.

#### **The Data Control Language**

*What is the data control language?*

The data control language affects connect, manipulation and access authorization, as well as transaction control.

## 4 Operating and Administration

*Which activities are necessary for the operation and administration of ADABAS?*

For the operation of the database, the tool CONTROL is provided. Installation, configuration, database (re)start and shutdown, backup and recovery, all the functions of database operation, can be executed via CONTROL. Three main parameters have to be controlled while operating an ADABAS database: the usage level of the database, the usage level of the log, and the number of concurrent database sessions. An operation for tablespaces (overflow, internal fragmentation) is not required in ADABAS.

For database administration, the Windows-based tool DOMAIN is available. This tool can be used to display, create, and maintain all the database objects. DOMAIN also keeps usage records for the database objects.

### 4.1 Installation and Configuration

*How is the database software installed? How is a database configured? How many parameters must be specified? How long do installation and configuration take?*

Installation and configuration are performed via CONTROL. It takes about 15-20 minutes to install the database software and set the system parameters. There are about 20 parameters to be defined.

### 4.2 Remote Administration

*Can the database system be administered from other nodes within the network?*

For the operation of a remote database, UTILITY is still required. To switch from one database to another, the menu item 'USE SERVERDB' is available in UTILITY. In future, this functionality will be integrated in CONTROL.

### **4.3 Restart**

*How is the DBMS started? How long does this take?*

Prior to starting the DBMS, the code must be loaded and the operating system resources made available. Then database operation is opened via CONTROL. This normally takes 1-2 minutes (i.e., after a proper shutdown). Command procedures can be used for automatic startup.

### **4.4 Shutdown**

*How is the DBMS shut down? How long does this take?*

There are two types of shutdown: slow shutdown refuses new users and waits until all active users have concluded their transactions. Quick shutdown cancels all current sessions.

### **4.5 Access Authorization**

*Who administers and controls the database access authorization? Up to what level can access authorization be granted? Can access authorization granted for foreign objects be passed on to other users?*

Access authorization is administered by the ADABAS kernel. For SELECT and UPDATE statements, access authorization can be granted up to field level. The GRANT ... WITH GRANT OPTION lets users grant access authorization to other users.

Access authorization can be granted to individual users or to a group of users.

## 4.6 User Classes

*Which user classes are there?*

The user classes are SYSDBA, DBA, RESOURCE, and STANDARD.

### The User Class SYSDBA

*What authorization does a user with SYSDBA status have?*

There is one SYSDBA per SERVERDB. The SYSDBA defines users with DBA status.

### The User Class DBA

*What authorization does a user with DBA status have?*

Users with DBA status can create both RESOURCE and STANDARD users, as well as create private data and grant privileges to other users. In addition, DBA users can combine users with identical access authorization to form a usergroup.

### The User Class RESOURCE

*What authorization does a user with RESOURCE status have?*

RESOURCE users can define tables, views, and synonyms of their own, as well as grant privileges.

### **The User Class STANDARD**

*What authorization does a user with STANDARD status have?*

STANDARD users can create views and synonyms. They can only work on database objects for which they have been granted privileges.

### **4.7 Database Objects**

*For which database objects and database operations can privileges be granted?*

Privileges can be defined for tables, table columns, views, indexes, DB procedures, and snapshots.

### **4.8 Accounting**

*Is the use of system resources recorded?*

CONTROL provides access to the accounting table which in turn provides a database administrator with a session-specific record of the CPU time spent, the I/O activity, etc.

### **4.9 Audit**

*Is it possible to determine which user has performed which data modification and when?*

AUDIT uses the log protocol to evaluate modifications of critical data.

## 4.10 Backup/Recovery

### Backup

*In what increments can the database be saved?*

It is possible to back up either the entire database or incremental database modifications (pages), as well as the log or parts of the log (log segments). CONTROL can be used to plan the backup operations in weekly schedules and to initiate them automatically in certain intervals. CONTROL also administers the media required for the backup according to the number of desired backup generations.

### Online Backup

*Does backup require that the DBMS be shut down?*

The database can be backed up parallel to normal processing; i.e., parallel to database modifications. The state of the database at the beginning of the session will be backed up.

### Parallel Backup and Recovery

*How are very large databases backed up and restored?*

The limiting factor for the backup and recovery of very large databases is the read and write speed of the tape devices involved (Video8, DAT). For this reason, ADABAS provides the option to back up to or restore from various tape devices at the same time. Using up to 16 tape devices, the backup and recovery time can be reduced considerably.

## Log

*What logging procedures are there? Is it possible to back up a log?*

There are four logging procedures:

DEMO	This logging procedure is only used to get the database system ready for operation. It is not suitable for production operation.
SINGLE	This logging procedure is designed for configurations with only one devspace. It does not protect against media failures.
NORMAL	This procedure writes a rollback log and a rollforward (archive) log. It requires a minimum of two devspaces. As the devspaces used for the logging should not contain any part of the database, it would be desirable to have at least three devspaces for this logging procedure.
DUAL	For additional protection, this procedure involves keeping two parallel archive logs. The failure of one archive log does not result in the termination of database operation. It requires a minimum of three devspaces.

Logs can be backed up offline and online. The definition of log segments allows completed parts of the log to be backed up. CONTROL provides the AutoSave function for automatic backup.



## MIRRORED DEVSPACES

*Can mirrored devspaces be used for backup?*

ADABAS mirrored devspace operation protects against hardware failures by keeping a copy of every data area. In the event of a disk error, operation can continue, since ADABAS will disregard the defective storage area.

Mirrored devspace operation is an ADABAS feature and independent of the specific operating system in use.

In addition to the mirrored database and log provided by ADABAS, many hardware manufacturers offer either mirror disks on the hardware level or the possibility of connecting RAID periphery. This allows the user to make a choice based on availability requirements.

## Recovery

*How much time does it take to restore a database?*

The time required to restore the database is comparable to that needed to back up the database. Because of the use of parallel backup/recovery, the restore time no longer depends on the database size, but on the capacity of the largest disk within the configuration.

## Recovery of a Previous Database State

*Is it possible to define a cut-off time for recovery when restoring the database?*

The RESTORE LOG function in CONTROL can be used to specify that only transactions completed prior to a specific point in time are to be restored.

### **7-Day Operation**

*Is nonstop, 7-day operation possible with ADABAS?*

Since ADABAS does not require any reorganization and the database and log backup can be performed online, uninterrupted 7-day operation is possible.

### **Resource Control**

*Does the database system have resource control?*

Options for the definition of new users make it possible to control the resource usage of each individual database user. In addition, the component CONTROL supports a user-specific accounting.

### **Security Criteria**

*What security criteria does the database system satisfy?*

The database system satisfies the C2 standard.

### **Consistent Database State**

*Can the savepoints in the log be configured?*

The time interval between two savepoints within the log can be defined in order to restrict the maximum time for a restart.

## 5Transaction Integrity

### 5.1Transaction

*What is a transaction?*

A transaction is a series of operations (SQL statements) on the database with the following characteristics:

Atomicity	The effects of a transaction are either realized completely or not at all.
Consistency	In terms of application logic, the database is in a consistent state both prior to and after a transaction.
Isolation	Concurrent transactions are performed in such a way that they do not affect each other at all or only to the degree defined by the isolation level.
Permanence	The effects of a transaction are lasting; i.e., they survive both the end of the session and the end of database operation and are protected against software and hardware failures.

### 5.2Subtransaction

*What is a subtransaction?*

Subtransactions can be nested within a transaction bounded by 'SUBTRANS BEGIN' and 'SUBTRANS END' or 'SUBTRANS ROLLBACK'. Thus, the effects of a transaction can be partially reset without affecting the lock release.

### **5.3 User Control of Transactions**

*What influence does the user have on transactions?*

The user can specify the beginning and end of both a transaction and its subordinate subtransaction(s) (SUBTRANS BEGIN/SUBTRANS END). The user also has the choice at the end of the transaction of keeping (COMMIT) or undoing (ROLLBACK) data modifications.

### **5.4 Rollback**

*Is there a rollback procedure? How does it work?*

All database modifications are recorded in transaction segments with before/after images in the log. When restarting after a DBMS failure, any completed transactions are redone, if any, and open transactions are rolled back.

### **5.5 Concurrent Access**

*How are different users kept from conflicting with each other? What can be protected for a single user? The entire database, a table, or a single row of a table?*

The problem of concurrent access is dealt with by placing locks on specific rows or tables.

## 5.6 Lock Modes

*What types of locks are available?*

There are read locks (shared mode) and write locks (exclusive mode). Read locks permit several users to read concurrently with each other, but prevent writing. Write locks permit a single user to write, while making the locked data inaccessible to other users.

In addition, ADABAS supports optimistic locks which can be used to determine whether database objects read have been modified in the meantime.

## Isolation Levels

*Which isolation levels are supported?*

ADABAS offers the standard SQL isolation levels 0, 1, 2, and 3. There are some additional isolation levels specific to ADABAS.

## 5.7 Explicit Locking

*Can data objects be locked explicitly?*

Tables, but also rows, can be explicitly locked with the LOCK statement.

## 5.8 Releasing Locks

*Are locks released again?*

Any locks on data objects still existing at the end of a transaction are automatically released. But a lock can also be kept beyond the end of a transaction by explicitly specifying the locked data object at the end of the transaction.

## 5.9 Deadlocks

*Are deadlock conflicts detected?*

Simple deadlocks are detected immediately. More complex deadlock structures, whose detection would involve excessive effort, are solved by a timeout and by cancelling the transaction.

## **6Distribution**

### **6.1SERVERDB**

*What does the term SERVERDB mean?*

A local database is called a SERVERDB.

### **6.2SERVERNODE**

*What is a SERVERNODE?*

The physical node in the network on which a SERVERDB is located is called a SERVERNODE.

### **6.3Distributed Database**

*What are the characteristic features of a distributed database?*

A distributed database consists of several SERVERDBs on one or more SERVERNODEs. Logically, they make up a central database. This means that a user looking for particular data need not know where that data is physically located.

## 6.4 Data Access

*How does data access work in a distributed database?*

From the point of view of the user, data access in a distributed system is obtained, unlike remote data access, without any knowledge of the physical location of the data. The data is simply addressed using its table name and column name.

## 6.5 Two-Phase Commit

*Why is the two-phase commit protocol used in a distributed database?*

In a distributed system, transactions are generally performed on more than one SERVERDB. By limiting the time-frame of possible error states, a two-phase commit ensures that these distributed transactions are realized as atomic units. In the event of an error, this expedites the decision whether to rollback or to repeat the transaction concerned.

## 6.6 Creating a Distributed Database

*How is the infra structure of a distributed database set up?*

The distributed database is set up successively by using CONTROL to introduce each new SERVERDB to the existing distributed system.



## 6.7 Client-Server Support

*How does ADABAS support client-server configurations?*

The complete functionality of ADABAS is available both locally and remotely. The standard connectivity TCP/IP is supported. To reduce client-server communication, ADABAS provides DB procedures, triggers, array statements, and DB functions.



## 7 General Technical Features

### 7.1 Data Types

*Which data types does ADABAS support?*

ADABAS is compatible with SQL application programs whose SQL syntax satisfies the ANSI standard or the DB2 or ORACLE dialect. ADABAS supports the DB2, ORACLE, and ANSI standard data types.

#### Numeric Data Types

*Which numeric data types are supported?*

The following numeric data types are supported:

- |       |  |
|-------|--|
| FIXED | Decimal fixed-point number with a precision of 18 digits.                                      |
| FLOAT | Decimal floating-point number with a precision of 18 digits and a number range of $10^{+63}$ . |

#### Alphanumeric Data Types

*Which alphanumeric data types are supported?*

The following alphanumeric data types are supported:

CHAR/VARCHAR ASCII  
CHAR/VARCHAR EBCDIC  
CHAR/VARCHAR BYTE

with up to 254 characters.

**Date/Time Values**

*Which data types are supported for the date and time values?*

The following date and time data types are supported:

DATE :	YYYYMMDD
TIME :	HHHHMMSS
TIMESTAMP :	YYYYMMDDHHMMSSmmmmμμμ

In addition to an internal format, country-specific representations are supported for date and time values.

**BOOLEAN**

*Are Boolean values supported?*

ADABAS knows the data type BOOLEAN which supports Boolean values and their embedding into programming languages.

**LONG**

*How is unformatted data supported?*

A column of the type LONG can receive unformatted data with a total length up to 2.1 GB.

## Varchar

*How are character strings represented internally?*

CHAR strings up to 30 characters long are represented internally with fixed lengths. CHAR strings with lengths between 31 and 254 characters are represented internally with variable lengths. The internal representations of VARCHAR strings are always variable. Longer character strings up to 2.1 GB in length are represented by the data type LONG.

## 7.2SQLMODE

*How is the distinction between ANSI, DB2, ORACLE, and ADABAS D specified?*

When connecting to the database, the user specifies an SQLMODE (ADABAS, ANSI, DB2, ORACLE).

## 7.3TERMCHAR SET

*What are TERMCHAR SETs used for?*

The ISO-ASCII and EBCDIC character sets include national characters, whose hexadecimal representations vary at certain terminals. To ensure that these characters are displayed correctly, the CONTROL component can be used to define so-called TERMCHAR SETs.

## 7.4 MAPCHAR SET

*What are MAPCHAR SETs used for?*

National characters in ASCII or EBCDIC code generally cause sorting problems, since the sorting order according to the internal code does not meet the users' expectations. Using CONTROL, it is possible to define MAPCHAR SETs in ADABAS which allow a sorting condition to be specified for each special national character (in German, e.g., "oe" for "ö").

## 7.5 Costs of a SELECT Statement

*How can the overhead needed for SELECT statements be calculated?*

The EXPLAIN statement can be used to calculate the expenditure for SELECTs. The applied search strategy is also displayed.

## 7.6 Key Generation

*Is it possible to generate keys automatically?*

The STAMP function or the implicit SYSKEY column can be used to assign network-wide unique, internal keys to rows of tables without user-defined keys.

## 7.7 System Resources

### Space Required by the Database Programs and Documentation

*How much space does the database software require?*

Depending on the hardware platform used, the database software can require up to 60 MB.

### Space Required by the Database

*How much space does an ADABAS database require in its smallest possible configuration?*

In its smallest possible configuration, an ADABAS database requires approximately 50 MB.

### Embedding the Database in the System

*Which processes are started for the database system?*

ADABAS provides a multi-threaded/multi-server process architecture by providing an external and internal tasking which can be configured. In addition to some auxiliary processes and according to the number of CPUs available, ADABAS installs so-called UKPs (user kernel processes). Within these UKPs, internal tasking helps to avoid another operating system process for each application session.

## **7.8Storage Procedures**

### **B\* Trees**

*Which storage procedure does the database system use?*

B\* trees are used as the storage procedure. The data records or index lists are kept in the leave pages; the key information in the index pages. The size of a page is 4 KB. For sequential access, these data pages are chained with the next page; for direct access, they are chained by using a multiple index.

### **No Reorganization**

*How dynamic is storage organization when data is modified?*

The method of storage organization used does not require reorganization. Any data modification is immediately followed by all the necessary, technical reorganizational steps, thus ensuring optimum access behavior. As a rule, a page is used up to about 80%, independent of the modifications made to the pertinent table.

### **Storage Management**

*What happens when a database user is dropped?*

All objects (tables, views, etc.) are dropped and the cleared storage space is released to free space management.



## 7.9 Logging

*What is written to the before and after images?*

INSERT: After image of the row

DELETE: Before image of the row

UPDATE: Field-by-field before and after image (for modified fields only)

## 7.10 Data Storage

*Does ADABAS utilize the UNIX file system or are 'raw devices' processed as well?  
Can data buffered in main memory be lost in the event of a system failure?*

ADABAS uses raw I/O and UNIX files. When UNIX files are used, they must not be subject to UNIX buffering. ADABAS has its own buffer management. Raw devices are preferred from a performance point of view; UNIX files are preferred from an administration point of view. Data consistency is achieved by using the system's own logging function.

## 7.11 Multi-volume Tables

*Is it possible to create tables which are larger than a UNIX file system?*

To distribute the workload, each table is automatically divided among several logical disks. This table striping prevents database hot spots, because all disk mechanics are equally loaded.

## 7.12 Protection

*Which facilities does ADABAS provide for database protection?*

- Access is controlled with passwords, user classes, allocation of data areas and operations, and user-specific modification logs.
- Lock types and lock modes synchronize concurrent processes.
- Transaction commands define consistent database states.

## 7.13 Data Area Failures

*Does the failure of an ADABAS data area terminate database operation?*

Due to ADABAS' totally dynamic and reorganization-free secondary memory organization with its implicit workload distribution among all the data areas available to the database, the data pages belonging to a particular table are generally distributed among all the data areas. A data area failure therefore means the end of database operation.

Considering the disk capacity normally available today, this is not a major drawback for medium-sized to large databases, because the risk is outweighed by the convenience of the secondary memory organization. A data area failure can only present a problem for very large databases because of the recovery time. To deal with this problem, ADABAS offers the following mechanisms:

- a) The use of SMP systems in the high-end field with a failover concept on hardware or operating system level and the use of mirror disks or RAID systems has the effect that, in future, protection against failures will be guaranteed less and less by the DBMS and more and more by the hardware or operating system.

- b) Operation of a distributed configuration with several SERVERDBs on one computer allows failures to be limited to partial databases, thus corresponding to a tablespace concept.
- c) In a failure-tolerant configuration with two (distributed) SERVERDBs on independent computers and with the replication of all tables, database operation will not be affected at all by the failure of one computer.
- d) With mirrored data areas, the database remains fully operable even after a disk failure. The mirroring of data areas does not depend on the characteristics of the underlying operating system.
- e) Parallel backup and recovery done by concurrent use of several tape devices accelerates these procedures considerably and is strictly recommended for large databases. Thus the recovery time no longer depends on the size of the database but only on the capacity of the largest disk and the number of tape devices operated concurrently.
- f) In multi-DB operation, it is possible to run several databases concurrently on one computer - each of these databases with different failure behavior and a separate backup procedure. From an application, it is possible to access up to eight concurrent databases.

Systems offering a recovery of database parts must be checked as to whether they still ensure database consistency. Modifying transactions across several parts of the database (tablespaces) or across referential integrity conditions specified among them cause a database to be inconsistent after a partial recovery. In an actual case of recovery, the operating database usually does not recognize an inconsistency because that requires knowledge of the application.

## **7.14 Optimization**

### **Independent Optimization**

*Which optimization procedures does the database system apply independently?*

- Data compression
- External tasking
- Internal tasking
- I/O optimization
- Access path optimization for local and distributed databases.

### **User-controlled Tuning**

*What additional tuning measures are available to the user?*

- Creating secondary indexes.
- Clustering data using appropriate primary key definitions.
- Dimensioning the caches (data cache, temp cache, catalog cache, converter cache).

## 7.15 Constraints

*What are the constraints on the database?*

Database size:		unlimited
Number of concurrent users:		configurable
Number of tables per database:		unlimited
Number of nodes in a distributed database:	2048	
Table size:		unlimited
Name length:	18	characters
Internal length of a table row:	4047	characters
Length of a LONG column:	2147483647	characters
Columns per table (with key):	255	columns
Columns per table (without key):	254	columns
Number of key columns:	127	columns
Precision of numeric values:	18	digits
Length of alphanumeric columns:	254	characters
Sum of the internal lengths of all key columns:	255	characters
Length of the sort columns in a SELECT:	250	characters
Number of result columns:	254	columns
Number of join tables in a SELECT:	16	tables
Number of join conditions in the WHERE clause of a SELECT:	64	
Number of single indexes per table:	255	
Number of multiple indexes per table:	256	
Number of DEVSPACES:	64	
Statement length:	8240	characters



## Table of Contents

<b>1</b>	<b>General Information.....</b>
1.1	The Name of the DBMS.....
1.2	Versions.....
1.3	The DBMS Supplier.....
1.4	Language Versions.....
1.5	Operating Systems.....
<b>2</b>	<b>Features of the DBMS.....</b>
2.1	Components of the DBMS.....
2.1.1	CONTROL.....
2.1.2	LOAD.....
2.1.3	DOMAIN.....
2.1.4	QUERY.....
2.1.5	EASY.....
2.1.6	OfficePlus.....
2.1.7	ADABASIC.....
2.1.8	Precompilers.....
2.1.9	CALL Interface.....
2.1.10	ODBC Driver.....
2.2	Display-oriented Interfaces.....
2.3	Interactive SQL.....
2.4	Report Generator.....
2.5	SQL.....
2.6	Extended SQL Functions of ADABAS.....
2.6.1	Referential Integrity.....
2.6.2	Updatable Join View.....
2.6.3	Outer Join.....
2.6.4	DB Procedure.....
2.6.5	Trigger.....
2.6.6	DB Function.....
2.6.7	Scrollable Cursor.....
2.6.8	Snapshot.....
2.7	Precompiler Concept.....

2.8	Data Dictionary.....
2.9	Database Statistics.....
<b>3</b>	<b>Characterization of the DBMS.....</b>
3.1	Database Objects.....
3.2	The Elements of the Database Language SQL.....
3.2.1	The Data Definition Language.....
3.2.2	The Data Manipulation Language.....
3.2.3	The Data Control Language.....
<b>4</b>	<b>Operating and Administration.....</b>
4.1	Installation and Configuration.....
4.2	Remote Administration.....
4.3	Restart.....
4.4	Shutdown.....
4.5	Access Authorization.....
4.6	User Classes.....
4.6.1	The User Class SYSDBA.....
4.6.2	The User Class DBA.....
4.6.3	The User Class RESOURCE.....
4.6.4	The User Class STANDARD.....
4.7	Database Objects.....
4.8	Accounting.....
4.9	Audit.....
4.10	Backup/Recovery.....
4.10.1	Backup.....
4.10.1.1	Online Backup.....
4.10.1.2	Parallel Backup and Recovery.....
4.10.2	Log.....
4.10.3	MIRRORED DEVSPACES.....
4.10.4	Recovery.....
4.10.4.1	Recovery of a Previous Database State.....
4.10.5	7-Day Operation.....
4.10.6	Resource Control.....
4.10.7	Security Criteria.....
4.10.8	Consistent Database State.....
<b>5</b>	<b>Transaction Integrity.....</b>



---

5.1	Transaction.....
5.2	Subtransaction.....
5.3	User Control of Transactions.....
5.4	Rollback.....
5.5	Concurrent Access.....
5.6	Lock Modes.....
5.6.1	Isolation Levels.....
5.7	Explicit Locking.....
5.8	Releasing Locks.....
5.9	Deadlocks.....
<b>6</b>	<b>Distribution.....</b>
6.1	SERVERDB.....
6.2	SERVERNODE.....
6.3	Distributed Database.....
6.4	Data Access.....
6.5	Two-Phase Commit.....
6.6	Creating a Distributed Database.....
6.7	Client-Server Support.....
<b>7</b>	<b>General Technical Features.....</b>
7.1	Data Types.....
7.1.1	Numeric Data Types.....
7.1.2	Alphanumeric Data Types.....
7.1.3	Date/Time Values.....
7.1.4	BOOLEAN.....
7.1.5	LONG.....
7.1.6	Varchar.....
7.2	SQLMODE.....
7.3	TERMCHAR SET.....
7.4	MAPCHAR SET.....
7.5	Costs of a SELECT Statement.....
7.6	Key Generation.....
7.7	System Resources.....
7.7.1	Space Required by the Database Programs and Documentation.....
7.7.2	Space Required by the Database.....
7.7.3	Embedding the Database in the System.....
7.8	Storage Procedures.....

7.8.1	B* Trees.....
7.8.2	No Reorganization.....
7.8.3	Storage Management.....
7.9	Logging.....
7.10	Data Storage.....
7.11	Multi-volume Tables.....
7.12	Protection.....
7.13	Data Area Failures.....
7.14	Optimization.....
7.14.1	Independent Optimization.....
7.14.2	User-controlled Tuning.....
7.15	Constraints.....