



MIMER

Shadowing Manual

Version 7.3

Copyright © 1996 Sysdeco Mimer AB

MIMER version 7.3 Shadowing Manual

November, 1996

Copyright © 1996 Sysdeco Mimer AB.

Published by Sysdeco Mimer AB,

P.O.Box 1713,

S-751 47 Uppsala, Sweden.

Tel +46(0)18-18 50 00.

Fax +46(0)18-18 51 00.

Internet: <http://www.mimer.se>

Produced by Sysdeco Mimer AB, Uppsala, Sweden.

All rights reserved under international copyright conventions.

The contents of this manual may be printed in limited quantities for use at a Mimer installation site. No parts of the manual may be reproduced for sale to a third party.

CONTENTS

1	INTRODUCTION	
1.1	Intended audience.....	1-1
1.2	Organization of this manual.....	1-1
2	GENERAL DESCRIPTION OF DATABANK SHADOWING	
2.1	Databank shadowing	2-1
2.2	Different levels of data protection	2-2
3	SHADOWING MANAGEMENT	
3.1	MIMER Utilities	3-1
3.2	Authorization.....	3-1
3.3	The shadowing utility.....	3-2
3.3.1	Wildcards.....	3-2
3.3.2	Create shadow.....	3-2
3.3.3	Drop shadow.....	3-3
3.3.4	Transform shadow to master	3-4
3.3.5	List shadowing information.....	3-5
3.3.6	Enter and leave a program ident.....	3-6
3.4	Backup on the fly.....	3-7
3.4.1	MIMER backup of a databank	3-7
3.4.2	Host system backup of a databank.....	3-8
3.5	Restore databank.....	3-12
3.6	Handling system databanks	3-13
3.6.1	SYSDB.....	3-13
3.6.2	TRANSDB.....	3-15
3.6.3	LOGDB.....	3-15
3.6.4	SQLDB.....	3-15
3.6.5	If a shadow for SYSDB, TRANSDB or LOGDB is not accessible.....	3-15
3.7	Testing your data protection.....	3-16
3.8	Configuring the system	3-17
3.9	Performance aspects of shadowing	3-17
3.9.1	Tuning	3-18
4	TROUBLESHOOTING	
4.1	Operator error messages.....	4-1
4.2	Miscellaneous problems	4-3

1 INTRODUCTION

MIMER Shadowing makes it possible to create and maintain one or more simultaneously updated copies of a databank. This allows for a higher degree of data availability by giving extra protection from disk crashes, etc. Shadowing also allows for backups to be made while the databanks are in use (so-called 'backup on the fly').

This manual describes the functions and benefits of shadowing, how to use the shadowing system, and how to handle any problems. This manual is a complement to the *MIMER System Management Handbook*.

The SQL commands CREATE SHADOW, DROP SHADOW, ALTER SHADOW and SET SHADOW are described in the *MIMER/SQL Reference Manual*.

Screen display examples in this manual may differ in minor detail from the display in the distribution. The examples are given solely for illustrative purposes.

1.1 Intended audience

The manual is intended for the personnel responsible for the maintenance and backing-up of the database.

1.2 Organization of this manual

Chapter 2 provides a general description of the MIMER shadowing system. The different ways to protect data from being lost are also discussed.

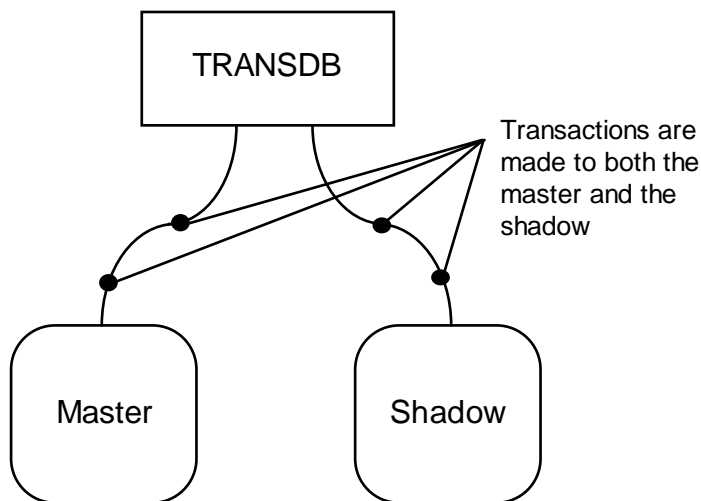
Chapter 3 describes how to use the shadowing system and shows examples of how to create and drop shadows, etc.

Chapter 4 deals with some problem situations and how to handle them.

2 GENERAL DESCRIPTION OF DATABANK SHADOWING

2.1 Databank shadowing

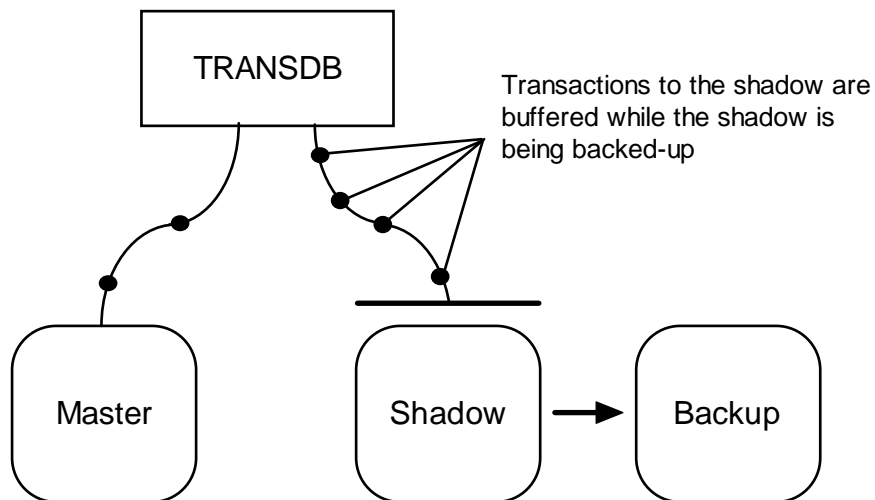
Databank shadowing means working with one or more copies of a databank at the same time. The master is the 'normal' databank file which is accessed for data storage and retrieval. The copies are called shadows. Any changes to the master data is automatically made to the shadows, thus protecting data from a disk crash or other mishap. A databank must have transaction control to be shadowed.



All changes made to a shadowed databank are automatically made to all the shadows for the databank.

If a master databank is lost, a shadow can quickly take the place of the master and operations can be resumed immediately (assuming the shadow is not also damaged in the disk crash). Normally, transforming a shadow to a master is much faster than restoring a databank from a backup copy.

It is recommended to take conventional backups as a supplement to databank shadows, to protect data in the event of a crash that destroys both the master and the shadow. With shadowing, backups can be taken from a shadow without stopping the multi-user system or interrupting work on the databank (so-called 'backup on the fly'). Any changes made to the master during the backup process are buffered in TRANSDB and are written to the shadow when it becomes available again.



Backups can be made without stopping the MIMER system.

Because operations are not interrupted to make backups, and because MIMER databanks are automatically reorganized, true 24 hour-a-day operation is possible.

Databank shadowing is entirely invisible to the application. This means that shadows can be added or removed to already existing applications. No special handling is needed to access tables in a shadowed databank.

2.2 Different levels of data protection


The *MIMER System Management Handbook* describes how data can be protected from loss by using LOGDB and TRANSDB. MIMER shadowing provides an even higher level of protection. Listed below are the different ways in which data can be protected from loss (from the least amount of protection to the highest):

- **All databanks on one disk and no logging**

If a databank is lost with this level of protection, it is only possible to continue operations from the last backup copy (all changes since the last backup was taken are lost). Databanks can be lost due to accidental deletion, disk crashes (which can destroy all files on a disk), etc. This level of protection is not recommended except for 'trash' databanks with unimportant contents.

- **Logging, with LOGDB and TRANSDB on a separate disk from the data**

LOGDB and TRANSDB are vital databanks if the system stops or if any databanks are lost. Because of this, LOGDB and TRANSDB should be stored on a separate disk, as shown in the following figure:



If a databank is lost, it can be restored to the state it had when it was lost by performing the transactions in LOGDB and TRANSDB to a backup copy of the databank (see the *MIMER System Management Handbook*). This may take some time, especially if the databank is large and backup copies are taken at long intervals, because there may be a lot of transactions in LOGDB.

WARNING


If the databank disk and the TRANSDB/LOGDB disk are handled by the same disk controller, the disk controller may cause both disks to crash. If this happens, the databanks can only be restarted from the state of the last backup copy. Therefore, it is advisable to use separate disks with separate disk controllers.

This security level gives a high degree of security and is recommended for databanks containing important data used in a system where the delay before the system is restored after a crash is not critical. To assure this high degree of security the backup copies should always be stored on separate removable media (e.g. tapes).

- **Shadowing, with shadows on a separate disk**

Shadows should always be stored on a separate disk from the masters to protect them from a total disk crash that could destroy both the master and shadow databanks. It is also advisable to use separate disk controllers to assure that a corrupt disk controller does not destroy the disks holding both the masters and the shadows.


If a databank is lost, its shadow can be transformed into a master and operations can be quickly resumed with no loss of data. Since the shadows are updated after the master, and the operations are saved in TRANSDB until the shadow is updated, it is important to have a good copy of TRANSDB when a shadow is transformed. To assure this, it is advisable to shadow TRANSDB. It is strongly recommended that TRANSDB and its shadows are stored on separate disks, as shown in the following diagram.



This arrangement gives a high degree of security and is recommended for databanks containing important data used in a system where it is vital to be able to get the system running again quickly after a disk crash.

- **Shadowing and logging**

Combining shadowing and logging (described in Chapter 6 of the *MIMER System Management Handbook*) gives the highest level of data protection. If logging is not used, the data is not protected if both the master and the shadow disks become corrupted. However, when shadowing is combined with logging (with LOGDB on a third disk) and backups are regularly taken onto separate media (tape, etc.), then data is always protected if any two of the disks crash.



Of course, additional disks can be used, just as long as the databanks that are separated above are not placed onto the same disks. If you only have two disk drives available and all the databanks are shadowed, then logging is of little value. Shadowing of LOGDB will not increase data protection significantly in this configuration.

Different degrees of data security can be used for different databanks, depending on the importance of the data. It is however important that all databanks that depend on each other (foreign keys, etc.) have the same level of protection. Otherwise logical inconsistencies may result if there is a disk crash.

3 SHADOWING MANAGEMENT

3.1 MIMER Utilities

MIMER has a number of utilities to help the system administrator maintain the system. These utilities are described in the *MIMER System Management Handbook*, with the exception of the options that deal with shadowing, which are described in this chapter.

All shadowing operations are handled from the UTIL program or with the SQL commands (CREATE SHADOW, DROP SHADOW, ALTER SHADOW and SET SHADOW; which are described in the *MIMER/SQL Reference Manual*).

3.2 Authorization

Any user can start the utilities program. The separate functions are restricted by special privileges. The following table shows the program ids that a user needs to have execute privilege on to do specific tasks:

<i>Tasks</i>	<i>Program ids</i>		<i>Notes</i>
	MIMER_BR	MIMER_SW	
Create a shadow		X	*
Drop a shadow		X	*
Transform a shadow		X	*
List shadow info.	X	X	**
Backup and restore	X		**
Set shadow offline	X		**
Set shadow online	X		**

Notes:

- * The user must have exclusive use of the databank. This means that no other user can access the databank until the operation is finished.
- ** A user who has EXECUTE privilege on MIMER_SW or MIMER_BR can list shadowing information for any or all databanks. Users can always backup and restore, set shadows offline and online, and list shadowing information for any databanks that they have created.

SYSADM is initially granted EXECUTE with grant option on both the MIMER_BR and MIMER_SW program ids.

3.3 The shadowing utility

This section describes the options in the shadowing utility, which contains functions for the maintenance of a shadowing system. This utility is used to create or drop a shadow, transform a shadow to a master, list information about shadows and databanks, and enter or leave a program ident. The shadowing utility main menu looks like this:

```
-- Shadow utilities --  
  
1. Create shadow  
2. Drop shadow  
3. Transform shadow to master  
4. List shadowing information  
5. Enter program ident  
6. Leave program ident  
  
0. Exit
```

3.3.1 Wildcards

Where noted in the following descriptions, wildcard characters following the SQL standard can be used in databank names. The percent character (%) represents one or more characters, and underscore (_) represents one and only one character. Thus DB% means every databank name beginning with DB, and DB_ specifies all databank names that consist of the letters DB followed by one (and only one) alphanumeric character or blank space. A percent character alone is shorthand for all databanks (excluding TRANSDB).

3.3.2 Create shadow

When the create shadow option is chosen, you are prompted for the name of the databank to be shadowed, the name for the new shadow, and the filename for the new shadow. Databanks can have more than one shadow.

Databanks with NULL option cannot be shadowed, because shadowing requires transaction handling. The databank to be shadowed cannot be used by any other users while a shadow is being created.

Shadow names can be up to 18 characters. The shadow name cannot be the same as the name of the master databank, of any other shadow, or of any shadow that has been transformed to a master.

Creating a shadow for a large databank may take some time, and thus should be carefully planned.

Example (user input is shown in bold):

```
-- SHADOW utilities --

1. Create shadow
2. Drop shadow
3. Transform shadow to master
4. List shadowing information
5. Enter program ident
6. Leave program ident

0. Exit

Select: 1

-- Create shadow --
Databank name           : HOTELDB
Create shadow for databank HOTELDB
Shadow name             : HOTELDB_S
Filename for shadow     : DISK1:[DIRECTORY1]HOTEL_S
Shadow HOTELDB_S created and copied from databank HOTELDB
```

A shadow databank is created on the disk named DISK1 and all tables and indexes in the databank are copied to the shadow.

Shadows can also be created in SQL with the CREATE SHADOW command (described in the *MIMER/SQL Reference Manual*).

3.3.3 Drop shadow

When drop shadow is chosen, you are prompted for the databank name for which a shadow should be dropped. Wildcard characters are allowed in the databank name. A listing of all shadows for the specified databank is displayed. Enter the name of the shadow you want to drop.

When a shadow is dropped, the file where the databank is stored is usually deleted in the operating system (see the machine-specific Users Guide). Otherwise, removing the file requires that it is deleted in the host operating system.

The databank cannot be used by any other users when a shadow is being dropped.

Example (user input is shown in bold):

```

-- SHADOW utilities --

1. Create shadow
2. Drop shadow
3. Transform shadow to master
4. List shadowing information
5. Enter program ident
6. Leave program ident

0. Exit

Select: 2

-- Drop shadow --
Databank name          : HOTELDB
DATABANK              SHADOW          OFFLINE
FILE
=====
HOTELDB                HOTELDB_S        N
DISK1:[DIRECTORY1]HOTEL_S
---
One shadow found
Name of shadow to drop (<CR> = skip) : HOTELDB_S
Shadow HOTELDB_S dropped

```

Shadows can also be dropped in SQL with the `DROP SHADOW` command (described in the *MIMER/SQL Reference Manual*).

3.3.4 Transform shadow to master

If a master is lost, it is possible to continue operations by transforming a shadow to the master. This operation must be done from the Shadowing utilities menu or with the SQL `ALTER SHADOW` statement – it cannot be done in the operating system.

When you choose the Transform shadow to master option from the Shadowing utilities you are prompted for the master databank name (wildcard characters are allowed). After the name is entered, a listing of all shadows for the selected databanks are displayed. Select which shadow to transform by giving the shadow name.

The transform operation only affects the data dictionary. The databank filenames are not changed. The old master databank file is usually deleted in the operating system (see the machine-specific Users Guide). Otherwise, removing the file requires that it is deleted in the host operating system.

The databank cannot be used by any other user when a shadow is being transformed into the master (this is not very likely to happen since this function is normally used when the master has been lost or damaged).

Example (user input is shown in bold):

```

-- SHADOW utilities --

1. Create shadow
2. Drop shadow
3. Transform shadow to master
4. List shadowing information
5. Enter program ident
6. Leave program ident

0. Exit

Select: 3

-- Transform shadow --
Databank name          : HOTELDB
DATABANK              SHADOW          OFFLINE
FILE
=====
HOTELDB              HOTELDB_S          N
DISK1:[DIRECTORY1]HOTEL_S
---
One shadow found
Name of shadow to transform (<CR> = skip) : HOTELDB_S
Shadow HOTELDB_S transformed to master

```

The SQL commands ALTER SHADOW and DROP SHADOW can be used together to transform a shadow to a master.

3.3.5 List shadowing information

The list shadowing information option lists the following information for both master and shadow databanks: databank name, shadow name (if the databank is a shadow), offline status for the shadows (Y = Offline, N = Online) and the filename.

Enter the master databank's name that you want to list the shadowing information for. Wildcard characters are allowed. If a databank is specified, shadowing information for that databank is displayed. If a carriage return (<CR>) is entered without specifying a databank, shadow information for all databanks in the system are displayed. The listing is halted when the screen is full and you can decide whether to continue the listing or not.

Example of list all shadows (user input is shown in bold):

```

-- SHADOW utilities --

1. Create shadow
2. Drop shadow
3. Transform shadow to master
4. List shadowing information
5. Enter program ident
6. Leave program ident

0. Exit

```

```

Select: 4
-- List shadowing information --
Databank name (<CR> => all databanks):
DATABANK          SHADOW          OFFLINE
FILE
=====
SYSDB              -              -
SYSDB7
---
TRANSDB           -              -
TRANSDB
---
LOGDB             -              -
LOGDB
---
SQLDB             -              -
SQL
---
BOOKDB           -              -
BOOKDB
---
BOOKDB            BOOKDB_S        N
DISK1:[DIRECTORY1]BOOKDB_S
---
ROOMSDB          -              -
ROOMSDB
---
HOTELDB          -              -
HOTELDB
---
HOTELDB          HOTELDB_S       N
DISK1:[DIRECTORY1]HOTEL_S
---
Continue <Y>?
7 databanks found
2 shadows found

```

3.3.6 Enter and leave a program ident

These options allow the shadowing responsibilities to be assigned to a program ident rather than a user.

3.4 Backup on the fly

The shadowing system allows for a so-called 'backup on the fly' of a databank to be taken. This means that the backup can be taken while other users are working with the databank. This requires that the databank is shadowed, since a backup on the fly is taken from a shadow (see picture on page 2-2).

First, the flow of transactions to the shadow is stopped by setting the shadow offline. Then a backup is taken of the shadow. Finally, the transaction flow is resumed by setting the shadow online and the pending transactions are executed on the shadow.

Backup on the fly can be taken as a host system backup or as a MIMER backup (both of which are described in Chapter 6 of the *MIMER System Management Handbook*). In either case, the SQL system management functions or the functionality in the MIMER backup and restore utility are used for taking backups.

The Backup/restore menu looks like this:

```
-- Backup/restore utility --
    1. Backup databank
    2. Restore databank
    3. List backups
    4. Drop backups
    5. Shadows offline,online,drop log
    6. Enter program ident
    7. Leave program ident
```

3.4.1 MIMER backup of a databank

If a MIMER backup is used, backup on the fly is automatically done by the SQL command CREATE BACKUP, or the backup and restore utility. The shadow is set offline, the backup is taken and registered in the data dictionary and the shadow is set online again, using the reset log option to drop the log.

If no shadow exists for the databank the backup has to be taken from the databank file. In this case, no other user can access the databank while the backup is being taken.

An example of a MIMER backup on the fly (user input is shown in bold):

```
-- Backup/restore utility --
    1. Backup databank
    2. Restore databank
    3. List backups
    4. Drop backups
    5. Shadows offline,online,drop log
    6. Enter program ident
    7. Leave program ident

    0. Exit

Select: 1
```

```

-- Backup databank --
1. Create full backup and drop log
2. Create backup databank and drop log
3. Create incremental backup and drop log
0. Exit

Select: 1

-- Create full backup --
Databank name          : HOTELDB
Create backup for databank HOTELDB
Incremental backup, file name : HOTELDB_INC
Size                   : 100
Backup databank, file name   : HOTELDB_FULL
Incremental backup HOTELDB_INC created from log for databank HOTELDB
Backup databank HOTELDB_FULL created
The log is reset for databank HOTELDB

```

There may be a delay while synchronizing of shadows occurs, which means that all pending transactions are executed to the shadows to ensure that the shadow state is identical to the master. How long this takes depends on the number of pending shadow server requests (this can be checked with MIMSERV; which is described in the *MIMER System Management Handbook*).

3.4.2 Host system backup of a databank

Shadowed databanks can also be backed-up on the fly with host system backups. If a host system backup is taken, you are responsible for the administration of the backup. You must set the shadow offline to be copied from. Then return to the operating system and copy the shadow file to a backup copy (on disk or tape). Then set the shadow online again, using the reset log option to drop the log records.

Note: The reset log option (used when setting a databank or shadow online) clears the records from LOGDB that were written before the databank or shadow was set offline. These records are cleared so that the log does not contain records covered by a databank backup, thus consistency between the log and the backups is ensured. If the log records are not cleared in this way, all updates will remain recorded in LOGDB.

If transaction handling is used to coordinate modifications to tables in more than one databank, it is important that they are set offline and backed-up simultaneously. Otherwise the backups may be logically inconsistent (this is described in Chapter 6 of the *MIMER System Management Handbook*). This will cause a problem if LOGDB is lost and you need to restore a databank. To prevent this problem, the functionality in the Backup/restore utility can set several databanks offline simultaneously by prompting for databank names until the first empty answer is entered. The same prompt logic is used when setting them online again, although shadows suspended simultaneously do not have to be resumed simultaneously.

When using the SQL system management functions, SET SHADOW may be given a list of shadows but SET DATABANK will only access a single databank name (it is more usual to set the entire database offline for a host backup when shadowing is not used). Consult the *MIMER/SQL Reference Manual* for further information on the SQL system management functions.

If the databank to backup has not been shadowed, you must ensure that you have exclusive use of the databank. This is done by checking with MIMSERV (described in the *MIMER System Management Handbook*) that no databanks are open, and that no shadow servers are working. Then stop the multi-user system to assure that no new users login, and take a backup copy of the databank file. Then enter the *single-user* backup and restore utility and drop the log for that databank. Finally, the multi-user system can be restarted.

It is always advisable to run the DBCHECK utility (described in the *MIMER System Management Handbook*) on a databank (or shadow) before making a backup.

There is a special submenu in the backup and restore utility for setting shadows offline, online and dropping log records:

```
-- Shadows offline,online,drop log --
1. Set shadows offline
2. Set shadows online and drop log
3. Drop entire log
4. List shadowing information
   0. Exit
```

The options in this menu provide the following functions:

1. Set shadows offline

Stops execution of transactions towards shadows for the specified databanks to be able to take a backup on the fly. If there is more than one shadow for the databank, the user is asked which shadow to set offline. All operations on the shadow are buffered in TRANSDB. Wildcard characters are allowed (note that TRANSDB is not included when the wildcard (%) is specified).

2. Set shadows online and drop log

Resumes operations again after a shadow has been offline. This option prompts for the name of the desired databanks and then asks if the log should be dropped for the specified databanks. Even if you do not drop the log, the buffered transactions in TRANSDB are performed, and operations proceed normally. Wildcard characters are allowed (note that TRANSDB is not included when the wildcard (%) is specified).

This option should also be used if a host system copy is taken for a databank that is **not** shadowed. If this option is used on a databank that is not shadowed, the log is dropped and nothing is set online.

3. Drop entire log

This option drops all records for all databanks from LOGDB. This option may only be performed by an ident with EXECUTE privilege on MIMER_BR. The option should be used with great care in production environments since backup protection will be incomplete unless a copy of every databank in the system has just been made. This option does not set any offline shadows online.

4. List shadowing information

This option is described in Section 3.3.5.

Example of a host system backup on the fly on a VAX (user input is shown in bold):

```

-- Backup/restore utility --
  1. Backup databank
  2. Restore databank
  3. List backups
  4. Drop backups
  5. Shadows offline,online,drop log
  6. Enter program ident
  7. Leave program ident
    0. Exit
Select: 5
-- Shadows offline,online,drop log --
  1. Set shadows offline
  2. Set shadows online and drop log
  3. Drop entire log
  4. List shadowing information
    0. Exit
Select: 1
-- Set shadows offline --
Enter list of databanks to set shadows offline for. Terminate with <CR>
Databank name      : BOOKDB
Databank name      : ROOMSDB
Databank name      :
Shadow BOOKDB_S is set offline
Shadow ROOMSDB_S is set offline

-- Shadows offline,online,drop log --
  1. Set shadows offline
  2. Set shadows online and drop log
  3. Drop entire log
  4. List shadowing information
    0. Exit
Select: 0
-- Backup/restore utility --
  1. Backup databank
  2. Restore databank
  3. List backups
  4. Drop backups
  5. Shadows offline,online,drop log
  6. Enter program ident
  7. Leave program ident
    0. Exit
Select: 0
-- MIMER UTILITIES --
  1. Backup/restore
  2. Export/import, SYSxxGEN
  3. SQL statistics
  4. Readlog
  5. Shadowing
    0. Exit
Select: 0
$ RUN MIMEXE7:DBC
MIMER Databank Check Utility
  Version 7.3.1
Databank filename : DISK1:[DIRECTORY1]BOOKDB_S
Result filename   : RESULT1
No errors found

```

```

$ RUN MIMEXE7:DBC
MIMER Databank Check Utility
  Version 7.3.1
Databank filename : DISK1:[DIRECTORY1]ROOMSDB_S
Result filename   : RESULT2
No errors found

$ BACKUP DISK1:[DIRECTORY1]ROOMSDB_S ROOMSCOPY
$ BACKUP DISK1:[DIRECTORY1]BOOKDB_S BOOKCOPY

$ RUN MIMEXE7:UTILM
Username: SYSADM
Password:
      -- MIMER UTILITIES  --
      1. Backup/restore
      2. Export/import, SYSxxGEN
      3. SQL statistics
      4. Readlog
      5. Shadowing
      0. Exit

Select: 1
      -- Backup/restore utility  --
      1. Backup databank
      2. Restore databank
      3. List backups
      4. Drop backups
      5. Shadows offline,online,drop log
      6. Enter program ident
      7. Leave program ident
      0. Exit

Select: 5
      -- Shadows offline,online,drop log --
      1. Set shadows offline
      2. Set shadows online and drop log
      3. Drop entire log
      4. List shadowing information
      0. Exit

Select: 2
      -- Set shadows online and drop log --
Enter list of databanks to set shadows online for. Terminate with <CR>
Databank name      : BOOKDB
Databank name      : ROOMSDB
Databank name      :
Drop log for databanks <Y>? Y
Set online operation initiated for shadow BOOKDB_S
Set online operation initiated for shadow ROOMSDB_S
The log is reset for databank BOOKDB
The log is reset for databank ROOMSDB

```

3.5 Restore databank

Restore databank is intended for databanks without shadows or when both the master and the shadow are lost. The databank must have been defined with the LOG option to be able to restore a databank to the state it was in before it was lost. The restore databank operation can be performed using UTIL or by using the SQL command ALTER DATABANK RESTORE. All shadows for the databank are restored automatically.

When restore databank is chosen, you are prompted for the name of the databank. A listing of shadows for the databank (if any) are displayed in case you forgot that the databank was shadowed. You are then asked if you want to restore from a MIMER backup or a host system backup, or if you want to quit. If a shadow exists, exit the restore function and transform the shadow to the master.

The following example assumes that the master and shadows are lost and you have made a backup, as shown in Section 3.4.1 (user input is shown in bold):

```
-- Backup/restore utility --

1. Backup databank
2. Restore databank
3. List backups
4. Drop backups
5. Shadows offline,online,drop log
6. Enter program ident
7. Leave program ident

0. Exit

Select: 2

-- Restore databank --
Databank name          : HOTELDB
DATABANK SHADOW       : OFFLINE
FILE
=====
HOTELDB HOTELDB_S      N
DISK1:[DIRECTORY1]HOTEL_S
---
One shadow found
Shadow found for this databank
Restoration from MIMER backup (M) or SYSTEM backup (S) or QUIT (Q)?
(M/S/Q): M
Databank HOTELDB created and copied from backup databank HOTELCOPY
Databank HOTELDB restored from log
876 records restored
Shadow HOTELDB_S created and copied from databank HOTELDB
```

3.6 Handling system databanks

The system databanks (SYSDB, TRANSDB, LOGDB, and SQLDB) require special handling in some situations. If a problem occurs with these databanks or their shadows, the only allowed login is by SYSADM to the UTIL program. The utility program will then recognize the problem and help you correct it.

System databanks are handled in the same way as other databanks, with the following exceptions:

- If there is a problem with SYSDB, TRANSDB, SQLDB or LOGDB, new users cannot login. Already active users receives an error message when doing operations that depends on the respective system databank, while other operations still work. The error state is held until MIMER is stopped and the error is corrected.
- If there is a problem with a shadow of SYSDB, TRANSDB, or LOGDB, new users cannot login until the backup and restore utility is used to drop or suspend the shadow (described in Section 3.6.5).
- No users can be connected while a shadow is created, altered or dropped for SYSDB, TRANSDB, or LOGDB.

3.6.1 SYSDB

Because the SYSDB databank holds all the data dictionary information about your database, protecting it with shadowing and/or backups is necessary. Otherwise, the whole database can be unreachable if SYSDB is lost.

Transform a SYSDB shadow to a master

If SYSDB is lost or corrupt, any existing shadow of SYSDB can be transformed into the master to allow MIMER to start again.

In a single-user system, a shadow of SYSDB should be renamed and/or moved to the location where the master SYSDB was. Then start the UTIL program and login as SYSADM. Enter the name of the shadow to be transformed into the master, and exit.

In a multi-user system, the multi-user system should first be started and the shadow's location should be entered for the location of SYSDB. Then start the utility and login as SYSADM. Enter the name of the shadow to be transformed into a master, and exit.

Example (user input is shown in bold):

```

          M I M E R / U T I L
          Username: SYSADM
          Password:
MIMER/DB warning -18013 in function CONNECT
          MIMER/DB started from SYSDB shadow. Transform SYSDB shadow
to master
          with UTIL, or restart system from master SYSDB
          -- Transform shadow --
DATABANK SHADOW          OFFLINE
FILE
=====
SYSDB    SYSDB_S          N
DISK1:[DIRECTORY1]SYSDB_S
---
One shadow found
Name of shadow to transform (<CR> = skip)   : SYSDB_S
Shadow SYSDB_S transformed to master
          -- MIMER UTILITIES  --
          1. Backup/restore
          2. Export/import, SYSxxGEN
          3. SQL statistics
          4. Readlog
          5. Shadowing
          0. Exit

Select: 0

```

Restore SYSDB

If SYSDB is lost and no shadows exists, a backup copy of SYSDB can be restored to allow MIMER to start again (an example of this is given in Chapter 6 of the *MIMER System Management Handbook*).

3.6.2 TRANSDB

Shadowing of TRANSDB assures that you can bring your database up-to-date if your ordinary TRANSDB is lost or damaged.

Transform a TRANSDB shadow to a master

If TRANSDB is lost or corrupt, an existing shadow of TRANSDB can be transformed to the master. Start the UTIL program and login as SYSADM. A message is displayed saying that TRANSDB cannot be opened, and a shadow must be transformed to the master (this is similar to the example in Section 3.6.1). If there are uncompleted transactions they will be completed, as if the "original" TRANSDB was still there.

3.6.3 LOGDB

If some databanks are not shadowed but backup copies of the databanks exist, then a shadow of LOGDB is useful if both an unshadowed databank and the LOGDB master are lost or corrupted.

Transform a LOGDB shadow to a master

If LOGDB is lost or corrupt, an existing shadow of LOGDB can be transformed to the master. Start the UTIL program and login as SYSADM. A message is displayed saying that LOGDB cannot be opened, and a shadow must be transformed to the master (this is similar to the example in Section 3.6.1). If there are transactions not yet written to the log, this will be done automatically.

3.6.4 SQLDB

Shadowing of SQLDB is not allowed because it only contains temporary data. However, SQLDB is required when a user logs on to MIMER. If SQLDB is corrupt or lost, it must be recreated by logging on to the utility program (this automatically recreates SQLDB if the databank is not found).

3.6.5 If a shadow for SYSDB, TRANSDB or LOGDB is not accessible

If a shadow for SYSDB, TRANSDB or LOGDB is not accessible, SYSADM should login to the UTIL program. An error message is given followed by the option to drop the shadow or set it offline. If the shadow is corrupt or missing, it should be dropped. If it is only temporarily unavailable it may be enough to set it offline for a short period of time.

Example (user input is shown in bold):

```

MIMER / UTIL
Username: SYSADM
Password:
MIMER/DB fatal error -16142 in function CONNECT
Cannot open databank SYSDB,
file SYSDB_S not found
Inaccessible shadow encountered. DROP or SET OFFLINE (D/S)? D
-- Drop shadow --
Synchronizing shadows
Shadow SYSDB_S dropped
-- MIMER UTILITIES --
1. Backup/restore
2. Export/import, SYSxxGEN
3. SQL statistics
4. Readlog
5. Shadowing
0. Exit

Select: 0

```

3.7 Testing your data protection

So far, this chapter has described the facilities that are available for the system manager to maintain the system. But how do you know that you are using the functionality in the right way, and that you will be able to get the system running again if something happens? The answer is **practice**. When you have decided on a certain strategy, a disk crash should be simulated and the databanks restored following the chosen strategy. Check that the contents of all tables are the same as before the simulated crash. Do not forget to simulate a crash of the system databanks (SYSDB, TRANSDB, LOGDB, SQLDB) and then restore or recreate them. When you have decided on a successful strategy, then build command files (scripts) that take 'normal' backups, backups on the fly, and creates shadows.

3.8 Configuring the system

In Section 2.2 the various levels of data security are described. When you configure your system, some additional considerations are:

- *How do I divide the work load over several disks to get the best performance?*

TRANSDB and its shadows should preferably be on fast disks (see Section 3.9).

- *How much disk space is required?*

A shadow file occupies the same amount of disk space as the master databank file.

A problem occurs if TRANSDB or LOGDB runs out of disk space. When this happens the system cannot continue. Therefore, it is important to make sure that there is enough space for these databanks.

How large LOGDB grows depends on if LOG option is used and how often backups are taken and drop log is performed, since drop log clears LOGDB and releases space.

How large TRANSDB grows depends on the number of transactions and how fast the shadow servers are able to perform the transactions on the shadows. When the transactions are performed on the shadows, the space in TRANSDB where they were stored is released.

- *How long does a backup on the fly take?*

During a backup on the fly one or more shadows are set offline, and transactions are buffered in TRANSDB. How large TRANSDB grows depends on the amount of time needed to backup the databanks, and the frequency of the transactions on the databanks with offline shadows.

3.9 Performance aspects of shadowing

Performance is not noticeably affected by shadowing in multi-user systems, even though the shadowing system needs more machine resources because more files need to be updated. Applications do not have to wait for the shadows to be updated because this is done by background processes. Actually all updates to the disk, even to master databanks, are performed by background processes (except for updates to TRANSDB and its shadows).

In single-user systems no background process is used. This means that an application has to wait for the shadows to be updated.

3.9.1 Tuning

If the updating of shadows is delayed it will cause TRANSDB to grow. This can happen for several reasons:

- A shadow has been set offline and forgotten. If this happens transactions will be buffered until it is set online again. To check if a shadow is offline use 'list shadowing information' or MIMSERV.
- A shadow is corrupt. Updates on the shadow results in an I/O error, and are buffered in TRANSDB. When this happens the operator is notified by the system. To check it use MIMSERV.
- There are too few shadowing processes to update the shadows, or they get too little machine resources.

Chapter 9 of the *MIMER System Management Handbook* describes how to determine the best number of shadow processes.

4 TROUBLESHOOTING

This chapter lists the error messages that may appear and the solutions to the different problems. Some of these problems have already been presented earlier in the manual, but are combined here for convenience.

When there is a problem, first look for the cause of the error by looking in the multi-user log file (described in the machine specific information). Then, check the following sections for a solution to the problem.

4.1 Operator error messages

There are four error messages that can be sent to the operator of the system. See the machine specific guide for an exact explanation of which users receive these messages. The four error messages and their solutions are:

1) Fatal databank error in MIMER/DB multi-user system

Databank <databank_name> has been disabled

In this case a databank has encountered an error which precludes further transactions to be performed on the databank. A common cause for this can be disk space exhausted. When MIMSERV is run it will display an error status for the databank.

Problem TRANSDB disk space exhausted.

Solution Allocate more disk space or do the following: Investigate why TRANSDB grows. Start more shadow servers, if required. If there are offline shadows, set them online again.

Problem LOGDB disk space exhausted.

Solution Allocate more disk space or do the following: Take new backup copies of all databanks. Remove the LOGDB file. Login to the utility program and create a new LOGDB file. In the future, take backups and drop log more often to prevent LOGDB from growing too large again.

<i>Problem</i>	Disk for master databanks is corrupt.
<i>Solution</i>	Find out which databanks are on the corrupt disk. Transform shadows for corrupt databanks into masters. When possible, create new shadows on another disk (this requires exclusive access to the databank).
<i>Problem</i>	A databank needs to be restored, but LOGDB is lost.
<i>Solution</i>	For this situation there is no satisfactory solution. Use the following emergency procedure to get the system running: Move a backup copy to the location of the databank file. Take new backups of all databanks. Create a new empty LOGDB. Run the DBOPEN utility to open all databanks (this ensures that databank identifications in the new log are up-to-date). Now you can continue from the state of the backup copy. All operations after the backup was taken are lost. If two databanks depend on each other, then both must be restarted from the backups. Note that if data dictionary objects (tables, indexes, etc.) have been created in the databank after the backup was taken, it is advisable to restart all databanks from the backups, since otherwise inconsistency between the databank and data dictionary will occur.
<i>Problem</i>	A master is lost and the shadow is suspended.
<i>Solution</i>	Transform the shadow to a master. When possible, create a new shadow (this requires exclusive access to the databank).

2) Fatal databank error in MIMER/DB multi-user system

Shadow <shadow_name> has been disabled

In this case a shadow has encountered an error. This will cause all transactions performed on the databank to be buffered in TRANSDB.

The first thing that should be done is to set the shadow offline. Then investigate the cause of the error by consulting the multi-user log file. If the error can be corrected the error status is removed if the shadow can be successfully set online, with the preserve log option.

<i>Problem</i>	A shadow has run out of disk space.
<i>Solution</i>	Set the shadow offline and allocate more disk space.
<i>Problem</i>	Disk for shadow databanks is corrupt. Transactions are buffered in TRANSDB, which then grows.
<i>Solution</i>	Drop the shadows on the corrupt disk (this requires exclusive access to the databank). They will only cause TRANSDB to grow. When possible, create new shadows on another disk. Observe that 'create shadow' requires exclusive access to the databank during the copying operation to the shadow so this operation will stop operations to the databank during the creation of a shadow. Plan when new shadows can be created.

3) Unrecoverable error in MIMER/DB

Multi-user system terminated

A fatal error has occurred in the multi-user system. This error is usually due to the lack of space in internal structures in the multi-user system. Please consult the multi-user log file to find out the exact cause of the failure.

4) Unrecoverable internal error in MIMER/DB

Multi-user system terminated

An internal error has been detected in the multi-user system. Store the output from MIMDUMP to a file and then restart the multi-user system. Report the problem to your MIMER service agent.

4.2 Miscellaneous problems

The following problems may also occur:

<i>Problem:</i>	A restore operation was interrupted.
<i>Solution:</i>	Make a new restore from the beginning. To do this you must still have the original backup copy available.
<i>Problem:</i>	Operations appear to stop when you try to set a shadow offline.
<i>Solution:</i>	There is sometimes a delay while synchronization is performed. Wait for the synchronization to be completed (see Chapter 9 of the <i>MIMER System Management Handbook</i>).