



User Manual

Software Version E.01.05 for the HP 3000



COPYRIGHT

Copyright © 1998-1999 Lund Performance Solutions. All rights reserved.

Shadow D/R (Data Replicator) User Manual

This manual, as well as the software described in it, is furnished under license and may only be used or copied in accordance with the terms of such license. The information in this manual is furnished for informational use only, is subject to change without notice, and should not be construed as a commitment by Lund Performance Solutions.

Except as permitted by such license, no part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, recording, or otherwise, without the prior written permission of Lund Performance Solutions.

Lund Performance Solutions
240 2nd Avenue SW
Albany, OR 97321
USA

TRADEMARKS

De-Frag/X, Developers Toolbox, Forecast/3000 Capacity Planner, Forecast/9000 Capacity Planner, Intact D/R, Performance Gallery, Performance Gallery Gold, SOS/3000 Performance Advisor, SOS/9000 Performance Advisor, Q-Xcelerator Resource Manager, **Shadow D/R**, System Managers Toolbox are trademarks owned by Lund Performance Solutions. Other brand and product names are trademarks or registered trademarks of their respective holders and should be noted as such.

Printed in the United States of America

For Technical Support, please contact:

Lund Performance Solutions

Phone: (541) 926 3800

Fax: (541) 926 7723

E-mail: support@lund.com

Web Home Page: <http://www.lund.com/>

Table of Contents

1 Introduction	1
Shadow Data Replicator	1
Advantages of Shadow D/R	1
Technical Support	1
LPS Technical Support Team	1
Shadow D/R User Manual	2
2 System Description	4
System Components	4
Primary and Secondary Systems and Processes	4
Shadow D/R Files	5
Shadow D/R Reports	5
System Design	6
System Operations	7
Database Replication	7
Offline Backup	9
Load Balancing	13
3 System Installation	15
Starting a New Log Cycle	15
On the Primary System	15
On the Secondary System	15
Loading the Shadow D/R Files	15
4 System Configuration, Setup and Startup	16
System Configuration	16
Understanding the System Configuration File	16
Configuring the Primary System	19
Configuring the Secondary System	20
Configuration Error Messages	21
System Setup	21
Setting Up Databases and Logging (Primary System)	21
Installing Files (Secondary System)	22
Setting Shadow D/R System Parameters	23
Setting Switch Environments (Secondary System)	23
System Startup and Shutdown	24
Startup	24
Shutdown	24

5 Normal Operations	26
Backing Up the Database	26
Starting a New User Logging Cycle	26
With Automatic CHANGELOG	27
Without Automatic CHANGELOG	28
6 DS Line Failure and Recovery	29
DS Line Failure Overview	29
DS Line Recovery Procedures	29
Recovering One of Several DS Lines	29
Recovering All DS Lines	30
7 Secondary Failure and Recovery	31
Secondary System Failure Overview	31
Assessing a Secondary Failure	32
Enabling Rollback for Easier Recovery	33
Recovering the Secondary System	33
Recovery Procedure for a Minor Secondary System Failure	33
Recovery Procedure for a Major Secondary System Failure	33
8 Primary Failure and Recovery	34
Primary System Failure Overview	34
Minor Primary System Failure	34
Major Primary System Failure	34
Recovering from a Major Primary Failure	35
Fast Recovery Procedure	35
Fast Recovery Procedure for Non-CHANGELOG Environments	36
Slow Recovery Procedure	38
Slow Recovery Procedure for Non-CHANGELOG Environments	38
Returning Systems to Their Original State	39
Procedure for Systems with CHANGELOG Environments	39
Procedure for Systems without CHANGELOG Environments	40
Example of a Primary System Failure and Recovery	41
Operating Under Normal Conditions	41
Switching Roles After a Primary System Failure	41
Updating the New Secondary System	42
Returning the Systems to Their Original Roles	42
9 Data Capture	44
Overview of the Data Capture Process	44
MPE User Logging Facility	45

MPE Logging Intrinsic	46
MPE Logging Commands	47
TurboIMAGE Database Logging Facility	47
TurboIMAGE Intrinsic	48
TurboIMAGE Logging-specific Intrinsic	48
Shadow D/R Data Capture Mechanism	49
Data Capture Intrinsic	49
WAIT Time	49
Timeout Implementation	49
Wait Time Selection	49
Log Record Blocks	50
Log File Reads	50
Record Transmission	50
OUTSTANDING Log Records	50
Activities of the Tape Process	50
Tape Process Activation	51
Tape Process Deactivation	51
Tape Process Creation	51
Tape File Creation	51
CHANGELOG Records	51
User Logging and CHANGELOG	51
CHANGELOG Record Processing	52
The Heartbeat	52

10 Data Transport 53

Network Links	53
Multiple Data Communication Links	53
Blocked Data	54
Advantages of Blocking Data	54
Disadvantages of Blocking Data	54
Transmission Frequency	54
Transmission Algorithm	54
Role of the SP	55

11 Data Posting 56

Receiving and Validating Data	56
Receiving and Executing Commands	56
Command Requests	57
Receiving the Replicate Log File	58
Keeping Track of Database Users	58
Updating the Replicate Database(s)	58
Database Integrity Checks	60
Maintaining System Restartability	60

Secondary System Recovery Enabled Feature	60
SAVE File Updates	61
System Restart	61
Switch Execution and Transaction Rollback	62

12 Security 64

MODE=1 Security Solution	64
File System Security Solution	64
TurboIMAGE Password Security Solution	65
Set Subsystem=Real Solution	65

13 Troubleshooting 66

Overview	66
Diagnostic Tools	66
Diagnostic Commands	66
Diagnostic Messages	66
Error Messages	69
Frequently Asked Questions	69

A Commands 72

Overview	72
Shadow D/R Commands	73
BACKUP Command	73
CONFIG Command	74
EXIT Command	75
HELP Command	75
KILL Command	76
RECOVER Command	77
RELEASE Command	78
RESUME Command	79
START Command	80
STATUS Command	81
STOP Command	83
SWITCH Command	84
SYSTEM Command	85
USE Command	85
VERIFY Command	86
VERSION Command	87

B Messages 88

Configuration Messages	88
Manager Process (MP) Messages	90
Communications Process (CP) Messages	95

Recovery Process (RP) Messages 99

Function Messages 106

Quit Messages..... 111

List of Figures 114

Index 116

1 Introduction

Shadow Data Replicator

Welcome to **Shadow D/R** (Data Replicator), a unique performance tool from Lund Performance Solutions. **Shadow D/R** is a transport mechanism which copies TurboIMAGE databases from a primary system to replicate databases on one or more secondary systems. Each secondary database exists as a “shadow” of the primary database, that is, each secondary database is automatically updated when changes or additions are made to the primary database.

With **Shadow D/R** you can:

- Replicate the database on one primary to multiple secondary systems.
- Replicate the databases on multiple primary systems to one secondary system.
- Replicate databases between two systems, each mirroring the databases of the other.

You do not need to modify any application software to use **Shadow D/R** and the product is not seen by application users.

Advantages of Shadow D/R

High Availability	Shadow D/R lets you maintain production during a primary system failure. When a primary system fails, the Shadow D/R operator can switch systems, so the secondary system becomes the temporary primary system.
Application Availability	Shadow D/R 's transport mechanism lets you do a full backup of data files to the secondary system without interrupting production on the primary system. While the backup is being done on the secondary, shadowing is temporarily interrupted. However, the database on the primary system continues to record production. When the backup is complete, shadowing is resumed and the secondary database is brought up to date.
Load Balancing Ability	Shadow D/R allows inquiry-only access to the replicate TurboIMAGE databases on the secondary system. This means you can move your read-only users from the primary system to the secondary system, thereby reducing the primary system workload.

Technical Support

LPS Technical Support Team

Lund Performance Solutions (LPS) will answer questions and assist you with problems you may have regarding **Shadow D/R**.

LPS Phone Number	Our Technical Support Team can be reached at (541) 926-3800 , Monday through Friday during the hours of 8:00 A.M., to 5:00 P.M., Pacific time, excluding holidays (New Year's Day, Memorial Day, Independence Day, Labor Day, Thanksgiving Day and Christmas Day).
LPS WWW Site URL	We invite you to visit the Lund Performance Solutions World Wide Web site. You can visit us at the following location: http://www.lund.com/ .
LPS Email Address	Send electronic mail to support@lund.com .

**LPS Postal
Address**

Send mail to **Lund Performance Solutions, 240 2nd Ave SW, Albany OR 97321.**

Shadow D/R User Manual

The *Shadow D/R* User Manual contains detailed information about the system structure, operation and commands. It is designed to be used as a reference tool in your everyday work with *Shadow D/R*.

**Corrections
and
Comments**

Lund Performance Solutions has made every effort to ensure the accuracy of this manual. If you have comments, questions or ideas regarding this document, please send them to us.

**Overview of
Contents**

The following is a summary of this manual's contents:

Chapter 1 Introduction

Chapter 1 briefly describes the features and advantages of *Shadow D/R* and provides information about technical support and the user manual.

Chapter 2 System Description

Chapter 2 describes and illustrates the components of the *Shadow D/R* system and provides an overview of the system's design and operations. (See Chapters 9, 10 and 11 for a more detailed explanation of *Shadow D/R*'s design.)

Chapter 3 System Installation

Chapter 3 lists the step-by-step instructions to install the *Shadow D/R* software.

Chapter 4 System Configuration, Setup and Startup

Chapter 4 provides all the information necessary to configure, setup, startup and shut down the *Shadow D/R* system.

Chapter 5 Normal Operation

Chapter 5 gives step-by-step instructions to perform a database backup and start a new user logging cycle.

Chapter 6 DS Line Failure and Recovery

Chapter 6 explains how *Shadow D/R* handles DS line failures and provides instructions to recover from one or more DS line failures.

Chapter 7 Secondary Failure and Recovery

Chapter 7 explains how *Shadow D/R* handles secondary system failures and provides instructions to recover from a secondary system failure.

Chapter 8 Primary Failure and Recovery

Chapter 8 explains how *Shadow D/R* handles primary system failures and provides instructions to recover from a primary system failure.

Chapter 9 Data Capture

Chapter 9 provides both an overview and a detailed explanation of the data capture process.

Chapter 10 Data Transport

Chapter 10 provides both an overview and a detailed explanation of the data transport process.

Chapter 11 Data Posting

Chapter 11 provides both an overview and a detailed explanation of the data posting process.

Chapter 12 Security

Chapter 12 describes *Shadow D/R* Mode = 1 security solution and also security solutions associated with files systems, TurboIMAGE passwords and set sub-systems.

Chapter 13 Troubleshooting

Chapter 13 includes information about diagnostic commands and messages and gives answers to some frequently asked questions.

Appendix A Commands

Appendix A contains information about each *Shadow D/R* command.

Appendix B Messages

Appendix B lists and explains all *Shadow D/R* messages.

2 System Description


System Components

The information in this section describes the components of the *Shadow D/R* system and illustrates their interconnections during normal system operation.

Primary and Secondary Systems and Processes

The terms, “primary system” and “secondary system,” refer to the function and the direction of the data replication, not the actual mechanisms. For example, in some environments where data replication is directed both ways, the same system can be considered both a primary system for one database and a secondary system for another database.

- CP** The communications process (CP) job transfers files from the primary to the secondary machine. If you are Shadowing to more than one secondary system, there is one CP per secondary.
- MP** The manager process (MP) job is the interface used to configure and manage the *Shadow D/R* application.
- RP** The recovery process (RP) job exists on the secondary system and performs the actual updates to the database.



System Components

- SP** The secondary communications process (SP) job is a son process of the CP.

System Description

TP The tape process (TP) job creates tapes when needed to resolve system failures or communications problems.

Shadow D/R Files

Configuration File The configuration file stores specific information about *Shadow D/R*'s configuration (refer to "System Configuration" on page 16).

Message Files All processes in the *Shadow D/R* system communicate with other processes in the system by means of message files; one for sending and one for receiving. The names of these files are specified in the configuration file.

- The CP opens the "command" file for reading only and the "response" file for writing only.
- The MP/TP process opens the "command" file for writing only and the "response" file for reading only.

This allows a secure route for information exchanged between the cooperating processes of the *Shadow D/R* system.

SAVE Files The SAVE file exists on the secondary system only. It is used by the *Shadow D/R* system to store information about the following records:

- Records transported to the replicate log file.
- Records processed against the replicate database on the secondary system.

In the event of a failure on the secondary system, *Shadow D/R* can reconstruct the database by using the database and SAVE files from the backup tape and the current log file from the primary system.

NOTE: The SAVE file indicates to the RP where to begin processing in the current log file. For this reason it is a good idea to backup the SAVE file at the same time you do a regular secondary system backup of the database and log files.

NOTE: The SAVE file is paired with the log file on the primary system. Therefore, if a new log cycle is started on the primary, a new SAVE file must be used on the secondary system. Errors will occur if an old SAVE file is used to reconstruct a database from records stored in a new log file.

File Names Names used for the following files must not begin with the letters x, y or z:

- Message files on the primary and secondary systems.
- Log files on the primary and secondary systems.
- The SAVE file on the secondary system.

The file names must be unique in at least one position other than the first letter. For example:

Acceptable File Names	Unacceptable File Names
MESSAGEA	AMESSAGE
MESSAGERBR	BMESSAGE
A1MESS	YMESS
A2MESS	ZMESS

Shadow D/R Reports

The *Shadow D/R* system provides a few reports, since it is primarily concerned with data transportation.

**STATUS
Reports**

The STATUS command produces information about the *Shadow D/R* system, but it is intended for online use by the *Shadow D/R* operator.

**RECOVER
Reports**

The RECOVER command produces a report when it brings a backup database up to date using a current log file. The report shows the last complete transaction for each user of the database and the last incomplete transaction for each user, if there was one.

**SWITCH
Reports**

The report generated by the SWITCH command is similar to the report produced by the RECOVER command.

The SWITCH report contains the following information:

- The user name.
- The process used.
- The database name.
- The process used.
- The database name.
- The log record number.

The logical device number for each user for every transaction reported.

System Design


All modifications to the database are recorded by the TurboIMAGE logging facility by way of the MPE user logging facility. This data is recorded in a disc resident log file for *Shadow D/R* databases.

All communications between the user and the *Shadow D/R* program are accomplished through a manager process (MP). The MP is invoked every time the user issues a command to the *Shadow D/R* system (usually done interactively through a terminal session).

Once *Shadow D/R* is activated, the communications process (CP) remains active as the controlling process for the system. (It is the program that provides control of the other components on both the primary and secondary systems.) The major function of the CP is to monitor the log file created by the MPE user logging process and pick up any data written to the log file. The CP transports this data to a slave program that resides on the secondary system, using the facilities of NS or DS.

NOTE: If the communication link between the systems breaks, the master communications process (CP) has the option to temporarily store the untransmitted data to a tape on the primary system. The tape process serves as an additional level of security in case the primary system fails before the communication line is repaired.

The data from the primary system is received by the secondary communications process (SP), a slave process owned by the CP on the primary system. When multiple DS lines are used, one of these processes is active for each DS line.



System Flow

Data from the CP is quickly passed to the recovery process (RP) on the secondary system. In addition to updating the databases on the secondary system, the RP maintains a replicate copy of the log file from the primary system. This replicate log file can be used to resolve a catastrophic failure.

At any time, the data flow between the systems can be suspended to allow the copied databases to reach a consistent state. Then, a backup can be made of the secondary system while processing continues on the primary system. Once the backup is complete, the data flow can be resumed and the secondary system will attempt to catch up as quickly as possible.


System Operations

Database Replication

A database which is to be Shadowed must be set up for logging. Transactions which alter the database cause log records to be automatically written to the log file, which **Shadow D/R** copies from the primary system to a replicate log file on the secondary system. The replicate log file is then applied to the replicate database.

As the following diagram illustrates, the **Shadow D/R** software on the primary system operates in isolation from the user applications on that system.

In fact, the only person on the primary system who needs to be aware of *Shadow D/R* at all is the operator who starts the *Shadow D/R* system.



Database Replication

The diagram shows the following process:

- 1 A user modifies the database.
- 2 The log records about that alteration are written to a log file.
- 3 Log records are read by the communications process (CP) and sent via a data communications line to the secondary communications process (SP) on the secondary system.
- 4 The SP passes the records on to the recovery process (RP), the RP posts these records to the replicate log file and processes the transactions against the replicate database maintained on the secondary system.

Data Transfer Frequency

The CP regularly checks the primary log file to see if it has received any new records. It buffers any records that it finds and transports them, one block at a time, at the time interval specified in the configuration file or when the buffer is full. This frequency insures that the *Shadow D/R* system has a very low impact on the performance of the primary system.

Data Reliability

The SP informs the CP on the primary system when it receives the data correctly. This is done via sequential record numbers used by both the SP and the CP to verify all records were sent.

If, for any reason, the SP on the secondary system receives the records but cannot reply to the CP, the CP will transmit the same data records again. If the SP is able to reply after the second transmission, it will compare the record numbers from the first and second transmissions and discard the replicate records from the first transmission. The SP knows whether or not it successfully informed the CP that the data was received. This tight end-to-end coupling

System Description

of the communication link allows the *Shadow D/R* software to guarantee that data is transmitted successfully.

For further reliability, the SP passes the data to the recovery process (instead of passing it directly to the log file). This additional process is needed because the SP is vulnerable to DS line failures. If, for example, a line failure caused the SP to be aborted by the MPE operating system on the secondary machine while the SP was updating the replicate log file, data could be lost. However, since the records are passed to the RP before the SP reports back to the CP that the data was not successfully received, even if the communication line fails and the SP aborts, no data is lost. The SP is best conceived as an extension of the CP, but resident on the secondary system.

"Heartbeat" A "heartbeat" status request is sent by the CP on the primary system to the RP on the secondary system. This enables the CP to determine the current status of the RP. For example, the RP can inform the CP via the heartbeat that the secondary database has been released for backup. The RP expects to be asked for the heartbeat status and, likewise, can determine whether the CP is still operating or not.


If one process misses a heartbeat, the other processes will send a message to the console. If both processes are active and the communication link is broken, both processes will miss heartbeats and both will report to their respective consoles that something is wrong with the *Shadow D/R* system.

Offline Backup

To ensure that users have continual access to their data on the primary system, perform backups of the secondary system instead of the primary. Although it is true that the secondary system is already a backup, the backup of the replicate database is necessary for recovering from a secondary system failure. The backup can also be archived.

Before an Offline Backup

To prepare for an offline backup, the following takes place




Offline Backup (before)

- 1 The **Shadow D/R** operator on the primary system uses the MP (manager process) to issue a backup request.
- 2 The MP process informs the CP of the backup request. The CP informs the SP, then the SP relays the message to the RP.
- 3 The RP acknowledges the backup request by changing its internal backup flag from "No" to "Pending."

Once the systems are prepared, the **Shadow D/R** operator can use the MP process to get a status report.

During an Offline Backup


The backup does not begin until the RP finds a quiet period (a point when all processes to modify the primary database are not active) (see "Waiting for a Quiet Period" on page 12). When the RP finds a quiet period, it changes its internal marker from "Yes" to "No" and releases the replicate database and log file for backup. The CP on the primary system is notified that the RP will not accept new records. The CP stops sending records to the secondary system, instead allowing the records to accumulate in the log file.



Offline Backup (during)

After an Offline Backup

The CP will begin to send records to the secondary system when the [Shadow D/R](#) operator tells the CP, via the MP, that the backup is done and log records can be sent to the secondary. The replicate database on the secondary system is updated by the RP. All outstanding records in the primary log file are recorded in the secondary log file and database until the secondary system is current with the primary system.



Offline Backup (after)

This catch-up process occurs automatically at the **Shadow D/R** system's normal, high speed. Depending upon the length of time it takes to do the backup, the secondary can be back in step in a relatively short time.

Waiting for a Quiet Period

An offline backup cannot start until **Shadow D/R** can insure transaction integrity in the replicate database. This integrity cannot be insured until after all database transactions are complete.

NOTE: Transactions to a database can be either single or multiple.

- Single transactions are complete in themselves, involving only one modification to the database. For example, deleting a record from a data set is a single transaction.
- Multiple transactions involve several modifications to the database and cannot be interrupted without leaving the database in an inconsistent state. For example, adding a sales record and updating a total sales counter is a multiple transaction, since two modifications are made to the database. If the database is backed up after the sales record is updated, but before the total sales counter is increased, the backup database is inconsistent. To avoid this problem, multiple transactions are marked with BEGIN and END intrinsics in the log file.

The occasions when the system's integrity is secure and all transactions are complete are referred to as "quiet periods."

System Description

The following diagram shows the difference between a busy period and a quiet period.



Quiet Period vs. Busy Period


The recovery process (RP) on the secondary system waits until it finds a quiet period in the stream of records coming to it from the CP. If the RP allowed a backup at the busy point (shown in the example above), both transactions would be incomplete and the resulting backup database would be inconsistent. A backup taken during the quiet period will result in a consistent backup, because at this point, all transactions have completed.

Load Balancing

"Load balancing" refers to using the secondary database in read mode for inquiry tasks while keeping write mode access on the primary system. Load balancing is beneficial for the following reasons:

- The processing load on the primary system is reduced.
- System performance in general is improved.
- Response time is improved.

Load balancing is accomplished by replicating the primary database onto multiple secondary systems. Distributed Shadowing allows massive inquiry applications to access the same database on different machines. *Shadow D/R* provides a means of circumventing the limitation on the number of terminals connected to the CPU. For instance, if an inquiry application requires 300 terminals with read access to a database, *Shadow D/R* can replicate the database on three secondary systems, each of which supports 100 users (see the following diagram).



3 System Installation

Starting a New Log Cycle

NOTE: Version E.01.03 of *Shadow D/R* has changed the size of the SAVE file. If you are currently using version E.01.02 (or earlier), you must begin a new log cycle as part of the installation procedure.

On the Primary System

- 1 Use DBUTIL to disable the database(s) for access.
- 2 Wait for *Shadow D/R* to synchronize the secondary database(s).
Outstanding=0 and DATABASE Opens=0 in STATUS display.
- 3 Stop *Shadow D/R*.
- 4 Reset logid:
:LOG logid, STOP
:PURGE LOG@
:BUILD LOG001; code=LOG; disc=1000, 8, 1
:ALTLOG logid; LOG=LOG001, DISC; AUTO
:LOG logid, RESTART
- 5 Use DBUTIL to enable the database(s) for access.

On the Secondary System

- 1 Purge the **SAVE** file.
- 2 Backup the database.
- 3 Rebuild the log file:
:PURGE LOG@
:BUILD LOG001; code=log; disc=1000, 8, 1

Loading the Shadow D/R Files

- 1 Logon as a user with SM capability and restore *Shadow D/R* files from tape:
:FILE T; DEV=TAPE
:RESTORE *T; @.PUB.LPS; CREATE; SHOW; OLDDATE
- 2 Stream the job, LPSINST.PUB.LPS.

NOTE: You may need to reply to the request on the console for the RESTORE command. Perform this procedure on both the primary and secondary systems.

4 System Configuration, Setup and Startup

System Configuration

Understanding the System Configuration File

The configuration file (config file) contains information regarding the environment for a *Shadow D/R* system. One configuration file is required for each primary system. An example of a config file is shown below.

BEGIN <<Sample Shadow configuration file>>

LOCAL=!*JOB shadow,mgr/mgrpass.aaa.ggg;hipri;outclass=lp,1,1:*

COMMANDS = **shadowp1.gg.aa:** <<pipe to CP>>

RESPONSES = **shadowp2.gg.aa:** <<pipe from CP>>

SAVEFILR = **shadowsv.gg.aa:** <<saves environment>>

USE logproc WITH databasename.ggg.aaa/pass:

WAIT = **6000:** <<milliseconds>>

MAXRECS = **100:** <<before requesting tape>>

DSLINE = **75, 70:** <<2 DS lines>>

MODE = **4:** <<MODE=4 is recommended>>

REMOTE = **hello shadows,mgr/mgrpass.aaa.ggg;pri=bs:**

LOG = **logfl.gg.aa:**

TAPEDEV = **tape:** <<tape device class>>

END. <<of configuration>>

Shadow D/R Configuration File

The configuration file is an ASCII file which can be created or modified with any text editor. It can be assigned any name, however the name must not begin with the letters X, Y or Z. The example (above) shows a configuration file named SHADCONF installed in PUB.SHADOW during the installation of *Shadow D/R*.

NOTE: If you make any changes to the configuration file, you should start a new logging cycle to ensure the changes are taken into account.

Comments

Comments may be inserted anywhere in the config file that space is allowed. The comment is started with "<<," and ended with ">>." For example,

<<comment>>

Parameters Parameters can be given in any order. They can be in upper or lower case. Parameter specifications are written one per line and they must end with a colon (:). File records can be up to 132 bytes wide.

BEGIN <<Sample Shadow configuration file>>

The BEGIN keyword at the top of the example file indicates the start of the configuration parameters within the configuration file. Any text prior to this is treated as one or more comments.

LOCAL !JOB shadow,mgr/mgrpass.aaa,ggg;hipri;outclass=lp,1,1:

This parameter is the jobcard used to do the following:

- Log on to the primary system for the batch job that runs the CP.
- Log on to the secondary system to run the RP.

The account and group should be on the same line as the log file. The account must have MAXPRI-BS and the user and account must have CAP=LG,OP added as capabilities.

COMMANDS=shadowp1.gg.aa:

For the primary system, the COMMAND parameter specifies the name of the message file that will be used by MP to send instructions to the CP.

For the secondary system, this specifies the name of the data and message files the SP will use to transfer data and communicate with the RP.

In both cases, the *Shadow D/R* software will handle the building and maintenance of this file. This file must have a name which does not begin with the letter x, y or z, and it must reside in the same group and account as the database and configuration file.

RESPONSES=shadowp2.gg.aa:

For the primary system, the RESPONSES parameter specifies the name of the message file that will be used by the CP to return responses to the MP.

For the secondary system, this specifies the name of the message file that the RP will use to return responses back to the SP.

In both cases, *Shadow D/R* will handle the building and maintenance of this file. This file must have a name which does not begin with the letter x, y or z, and it must reside in the same group and account as the database and configuration file.

SAVEFILE=shadowsv.gg.aa:

The SAVEFILE parameter specifies the name of a file that the *Shadow D/R* software will build, maintain and use to store the current environment information for the systems. This file must have a name which does not begin with the letter x, y or z, and it must reside in the same group and account as the database and configuration file. (For information on SAVE files, see "SAVE Files" on page 5.)

USE logproc WITH databasename.ggg.aaa/pass:

This parameter specifies the name of the logging process that is to be monitored. In addition, it specifies the databases that should be Shadowed using that log process and a password that can be used to get "creator" capabilities for each of the databases. You can configure up to 12 databases, all of which must log to the same logid.

WAIT=6000:

The WAIT parameter specifies the maximum time in milliseconds that *Shadow D/R* should wait before examining the log file on the primary system. This is used when the *Shadow D/R* system finds no more log records to send to the secondary, so it pauses because it has nothing to do. By specifying a short period of time, you minimize the data that might be lost during a failure of the system, but at the same time, you may cause the *Shadow D/R* overhead to be higher than necessary.

MAXRECS=100:

The MAXRECS parameter is the taping threshold that specifies how far the *Shadow D/R* software is allowed to get behind in the log file before the operator is requested to provide a tape for storing the records.

The maximum allowable number for MAXRECS is 9999.

There are two reserved words which can be used with the MAXRECS parameter: NEVER and ALWAYS.

- If MAXRECS=NEVER, the *Shadow D/R* program will never write records to tape.
- If MAXRECS=ALWAYS, all transactions will be sent to tape as well as across the DS line to the secondary system.

Setting the Taping Threshold Value

The taping threshold is specified by the *Shadow D/R* operator in the configuration file. The value of the taping threshold represents a compromise between convenience and risk.

- If the mark is too low, the *Shadow D/R* operator may have to mount and dismount a tape every few minutes when the TP process alternately falls behind and catches up as the processing load on the primary system fluctuates.
- On the other hand, the taping threshold is a measure of the number of records which might be lost forever if the secondary and primary systems fail just before the taping threshold is reached. In this situation, there is no backup anywhere of the latest records in the primary log file.

DSLINE=75;70:

The DSLINE parameter specifies which DS lines will be available for the *Shadow D/R* program to use. This can be specified as either a logical device numbers or device class names for DS devices, or as node names for LANIC devices. You can have a maximum of 4 DS lines configured (one of which may be LANIC). You can have more than one DS line per secondary

MODE=4:

The MODE parameter should be left at MODE=4 (default) when possible.

If you absolutely require the secondary database be opened in MODE 1 because of existing applications in MODE 5, this parameter can be used to change the default open mode (4).

A problem may arise with locking contention on the secondary database. The RP uses DATASET level locking. If there are no other processes doing some locking, this can impede the RP. Security of the secondary database is also a potential problem when this parameter is used. Please refer to "Security" on page 64.

You may use LOCK or NOLOCK instead of MODE numbers.

REMOTE=Hello shadowsp,mgr/mgrpass.aa,ggg;pri=bs:

The REMOTE parameter specifies the logon command that will be used by any *Shadow D/R* components that must log on the secondary system. The account and group should be the same as that for the log file. This account must have MAXPRI=BS and the account and user must have CAP=LG,OP as capabilities.

LOG=logfl.gg.aa:

The LOG parameter specifies the name of the disc file to which the logging process on the primary system will write. It is also used to access a copy of this log file on the secondary system. This file must have a name which does not begin with the letter x, y or z, and it must reside in the same group and account as the database and configuration file.

TAPEDEV=tape:

The TAPEDEV parameter supplies the device class name to be used by *Shadow D/R* on the primary system when it finds that the transmission process is running behind and must, therefore, copy the outstanding data to tape.

END. <<configuration>>

The END line defines the end to the configuration parameters within a configuration file.

Wherever passwords are used in the system, they must be included in the configuration file. If the logging procedure has a password, for example the "USE..." line will have to begin "USE logproc/logpass WITH..."

Configuring the Primary System

Please read "System Description" on page 4 and "Understanding the System Configuration File" on page 16 before attempting to configure the *Shadow D/R* program.

To configure the primary system, do the following:

- 1 On the primary system, create a configuration file or edit the example SHAD-CONF.PUB.SHADOW (see "Shadow D/R Configuration File" on page 16).
- 2 Run the Manager Process (MP), you must have either SM or PM capability, and verify the configuration file:

:RUN SHADOWMP.PUB.SYS

VERIFY configfilename

The configuration file is ready when no errors are reported by the VERIFY command. If there are errors, correct the configuration file and verify again. (See the following configuration file example.)


```
:RUN shadowmp.pub.sys
```

```
V.UU.FF Shadow Manager Program
```

```
(C) COPYRIGHT Carolian Systems International Inc. 1984
```

```
>verify config.shadow.sys <<test-bed Shadow Config File>>
```

```
BEGIN<<Sample Shadow configuration file>>
```

```
LOCAL=!!JOB shadow,mgr/mgrpass.aaa.ggg;hipri;outclass=lp,1,1:
```

```
COMMANDS      = shadowp1.gg.aa:      <<pipe to CP>>
```

```
RESPONSES     = shadowp2.gg.aa:      <<pipe from CP>>
```

```
SAVEFILR      = shadowsv.gg.aa:      <<saves environment>>
```

```
USE logproc WITH databasename.ggg.aaa/pass:
```

```
WAIT          = 6000:                <<milliseconds>>
```

```
MAXRECS       = 100:                <<before requesting tape>>
```

```
DSLIN         = 75, 70:             <<2 DS lines>>
```

```
MODE          = 4:                  <<MODE=4 is recommended>>
```

```
REMOTE        = hello shadows,mgr/mgrpass.aaa.ggg;pri=bs:
```

```
LOG           = logfl.gg.aa:
```

```
TAPEDEV       = tape:              <<tape device class>>
```

```
END. <<of configuration>>
```

```
no errors found in Shadow configuration file. (MPMSG 81)
```

```
>exit
```

```
END OF PROGRAM
```

Verification of Configuration File

- 3 After verifying the system configuration, secure the file to restrict access from other users. The following command will restrict access to only the creator:

```
:ALTSEC filename; (R,A,L,W,X:CR)
```

Configuring the Secondary System

To set up the secondary database for recovery, do the following;

NOTE: The log files created by the user logging process on the secondary as a result of this particular rollback process will have no DBBEGIN's or DBEND's and will show that every record came from the same user (the RP).

- 1 Shut *Shadow D/R* processing down:

```
:RUN SHADOWMP.PUB.SYS  
>SYSTEM config  
>STOP
```

- 2 Ensure there is not access to the secondary database:

```
:RUN DBUTIL.PUB.SYS  
>SHOW databasename USERS  
>EXIT
```

- 3 On the secondary system, build a log file using automatic CHANGELOG (if there is currently a Shadow log file in this group, be sure to use a different name):

```
:BUILD logf1001.group.account;CODE=LOG;DISC=10000,10,10
```

NOTE: This log file is a TurboIMAGE byproduct from enabling rollback recovery for your database. Shadow does not make any use of the log file generated by the logging process on the secondary, and inactive files may be purged at the operator's convenience. To help facilitate this process, use CHANGELOG.

- 4 Obtain a logid:

```
:GETLOG logproc;LOG=logf1001.group.account,DISC;AUTO
```

- 5 Start *Shadow D/R*:

```
:RUN SHADOWMP.PUB.SYS  
>SYSTEM config  
>START
```

When *Shadow D/R* is restarted, the Recovery Process (RP) will determine whether rollback is enabled on the secondary database by checking the root file. If rollback is enabled, all I/O to SAVE files will complete. The message, "DB's enabled with ROLLBACK - physical I/O to SAVE," should appear in the RP joblisting spool file, indicating that rollback is enabled.

Configuration Error Messages

Please refer to "Configuration Messages" on page 88.

System Setup

Setting Up Databases and Logging (Primary System)

To set up databases and logging on the primary system, do the following:

- 1 Ensure no users are accessing the database on the primary system:

```
:RUN DBUTIL.PUB.SYS  
>SHOW databasename USERS
```

- 2 Disable the database for access.

```
>DISABLE databasename FOR ACCESS  
>EXIT
```

- 3 If you are currently logging, do the following:
 - a At the system console, turn off logging and alter the log procedure file declaration.
:LOG logprocedurename, STOP
:ALTLOG logprocedurename;LOG=logfilename.group.account
 - b FCOPY or RENAME the pre-existing log file to the group.account used in the LOCAL parameter of the configuration file.
- 4 If you are not currently logging, get a log process called LOGPROC.
:HELLO mgr.account,group
:BUILD logfl.group.account;CODE=LOG;DISC=10000,10,10
:GETLOG logproc;LOG=logfl.group.account,DISC
 - a Use the GETLOG AUTO option for automatic CHANGELOG environments. If AUTO is not specified, the default, NOAUTO, is in effect.
 - b In the :BUILD command, the parameters suggested in the DISC portion are not necessarily best for every system. The correct values depend entirely on the rate at which transactions are applied to the database and the frequency of log cycles.
 - c The group.account should be the same as used in the LOCAL parameter of the configuration file.
- 5 Store all files from the group and account (where the databases reside) to tape:
:STORE @.ggg.aaa;*t;SHOW
 This command stores a copy of the configuration, log and database files on tape. This tape can be used later to set up the secondary system.
- 6 Start logging at the system console (for an existing log file, use the RESTART option).
:LOG logproc, START
- 7 Enable the database on the primary system for logging:
:RUN DBUTIL.PUB.SYS
>SET databasename LOGID=logproc
>ENABLE databasename FOR ACCESS
>ENABLE databasename FOR LOGGING
>EXIT
 Ignore the error message, "Login is being enabled when no preceding DBSTORE was done," as the MPE STORE process completes.

Installing Files (Secondary System)

To install files onto the secondary system, do the following:

- 1 Log on to the secondary system and restore the files from the tape created in step 5 of "Setting Up Databases and Logging (Primary System)" on page 21." If the "ggg.aaa" group and account do not exist on the secondary system, they must be created to match the structure on the primary exactly. All capabilities and passwords must be identical on the two systems.
:HELLO MGR.aaa.ggg
:FILE;DEV=TAPE
:RESTORE*T;@.@;SHOW

This will restore the database files, the configuration file and the log file to the "ggg.aaa" group and account.

- 2 Enable the secondary database for access and disable it for logging.

:RUN DBUTIL.PUB.SYS

>ENABLE databasename FOR ACCESS

>DISABLE databasename FOR LOGGING

>EXIT

- 3 If you are Shadowing more than one database, make sure that the RP on the secondary system can write to each database that is not in the "ggg.aaa" group and account. There are two ways to do this:

- Give the user "mgr.aaa" System Manager (SM) capability, or
- Use DBUTIL on the secondary system to release each database that is not in the "ggg.aaa" location.

This release command maintains normal TurboIMAGE security.

Setting Shadow D/R System Parameters

To set system parameters for *Shadow D/R*, do the following:

- 1 Increase the number of jobs allowed on each system by 1. This will ensure that one extra job (streamed by *Shadow D/R*) can run on the primary and secondary systems.
- 2 Increase the session limit by the number of DS lines to be used.
- 3 Set the Communications Subsystem Configuration for the INP connect time-out to a short value (30 for hardwired and as short as possible for dial connections consistent with dial facilities is recommended).

This step is particularly important if you use *Shadow D/R* as a high-availability business solution. It provides an early warning to operators when a remote system is experiencing problems.

- 4 Determine the maximum number of users that will need to access the log file and decide if the current system configuration needs adjustment. View the SYSDUMP/SYSGEN#OF USERS PER LOGGING PROCESS in the system tables dialog.

Setting Switch Environments (Secondary System)

Use the following checklist when preparing the secondary system for a smooth switch and transition to primary system status:

- 1 Ensure the logging process for the secondary system is identical to the primary system using GETLOG.
- 2 Copy the confirmation file created on the primary system to the secondary system.
 - a Ensure the user identified in the file has LG and OP capabilities.
 - b If the account on the primary system has MAXPRI=BS, then the secondary system should have the same setting.
 - c Modify the DSLINE parameter to identify the current primary system for shadowing in the new direction after a switch.
- 3 Ensure the MAXIMUM USERS PER LOG PROCESS system parameter is adequate on the secondary system.

- 4 Ensure the density on the backup tapes is compatible between machines.
 - a A logon UDC can be set for the configuration file logon users or account.
 - b The MPE file equation, FILESHADTAPE;DEV=TAPE;DEN=1600, will redirect the TP output file.

System Startup and Shutdown

Startup

Once the *Shadow D/R* system has been installed successfully and configured, startup of the system is straightforward. The *Shadow D/R* operator issues the START command from the manager process (MP) on the primary system and the *Shadow D/R* software does the rest, as follows:

- 1 The MP creates the communications process (CP).
- 2 The CP checks the configuration file for existing DS line.
- 3 The CP logs on to the secondary system to which each DS line is connected and creates a secondary communications process (SP) as an extension of itself.
- 4 The SP (or, if more than one SP is created, the initial SP) sets up the recovery process (RP).
- 5 The RP obtains control of the replicate database and log file.

To start the *Shadow D/R* system, do the following:

- 1 Run the MP on the primary system:


```
:RUN SHADOWMP.PUB.SYS
>SYSTEM configfilename
>START
```
- 2 Observe the following response by *Shadow D/R*:


```
#Jnnn ("nnn" is the number of the job Shadow D/R is running.)
The Shadow system was started successfully. (MPWARN 55)
>
```

NOTE: If logging has not been started, you will see the following response:

COPENLOG failure #3

This message indicates that logging is not set up correctly. If you receive this message, verify that logging is properly enabled.

Shutdown

The *Shadow D/R* operator can stop the Shadowing process using the standard procedure or abruptly, using the KILL command.

Standard Shutdown

The *Shadow D/R* operator issues the STOP command to the MP on the primary system. The MP informs the CP to wait for a quiet period to stop transporting log records (see "Waiting for a Quiet Period" on page 12). This leaves the database in a logically consistent state.

Abrupt Shutdown

The *Shadow D/R* operator issues the KILL command to the MP on the primary system. The MP informs the CP to stop transporting records immediately. This may leave the secondary database in an inconsistent state. The operator can then restart *Shadow D/R* and any partial transactions will be completed by the *Shadow D/R* program.

5 Normal Operations

Backing Up the Database

It is a good idea to make periodic backups of the database on the secondary system; at least once during each logging cycle on the primary system is recommended. Since the SAVE file maintains the check points for restarting *Shadow D/R* in the log file, it is important to always treat SAVE files and log files in pairs in regards to backup. However, only a database backup is required at the start of a new log cycle.

To backup the database, log file and SAVE file, do the following:

- 1 Issue a BACKUP command on the primary system via the Manager Process (MP).
- 2 When the STATUS command shows that the backup status is "Now" (see the following figure) inform the operator on the secondary system that the *Shadow D/R* system is in backup mode.

>STATUS									
=====			Shadow		STATUS		=====		
					1:34:57				
*PRIMARY					SECONDARY				
LOGFILE	record#	=	4675 (001)		LOGFILE	record#	=	4623 (001)	
	time	=	1:34:54			time	=	1:30:36	
	outstanding	=	52 recs		FLAGS	Backup :	Now		
	max ever	=	4675 recs						
SENT	logical	=	4875 recs		DATABASE	opens	=	3 paths	
	physical	=	198 blocks			transact	=	2 outst	
						Open mode	=	1 LOCK	
	[current status]					[current status]			
LINES OPEN:SYSTEMB									
=====									

STATUS Command (status of backup on the secondary system is "Now")

- 3 When *Shadow D/R* releases the database on the secondary system, do a complete backup of the log file, database and SAVE file on the secondary system.
- 4 When the backup is complete, issue the RESUME command on the primary system.

Starting a New User Logging Cycle

Periodically, it will be necessary to start a new log cycle on the primary system. With the *Shadow D/R* system in operation, take care to keep the SAVE file on the secondary system in step with the log file on the primary system.

With Automatic CHANGELOG

To create a new logging cycle on a system with automatic CHANGELOG, do the following:

On the Primary System

- 1 Prevent users from accessing the primary database.
- 2 When no users are accessing the database, use DBUTIL to disable any further access.
- 3 Issue the CHANGELOG command.
- 4 Ensure the databases are up-to-date and synchronized by checking the STATUS display. The "LOGFILE record#" values for the primary and the secondary should be equal and the "DATABASE opens" parameter should indicate that no users are accessing the secondary database (shows 0 paths, as in the following figure). When this state is achieved, use the STOP command to shutdown *Shadow D/R* (the user logging will record transactions).

>STATUS											
=====			Shadow		STATUS		=====				
					1:34:57						
			*PRIMARY			SECONDARY					
LOGFILE	record#	=	4675 (001)			LOGFILE	record#	=	4675 (001)		
	time	=	1:34:54				time	=	1:34:54		
	outstanding	=	52 recs			FLAGS	Shadow Mode				
	max ever	=	4675 recs								
SENT	logical	=	4875 recs			DATABASE	opens	=	0 paths		
	physical	=	198 blocks				transact	=	0 outst		
							Open mode	=	1 LOCK		
	[current status]						[current status]				
LINES OPEN:SYSTEMB											
=====											

STATUS Display (record#'s are equal for both systems)

- 5 Use DBUTIL to enable the database for access.
- 6 Allow users to access the database on the primary system again.

On the Secondary System

- 7 Backup the database on the secondary system along with the SAVE file, but not the log file.

On the Primary System (Again)

- 8 Start *Shadow D/R* when the secondary backup is complete.

A store of the SAVE file is necessary for automatic CHANGELOG environments. In the event of a secondary system failure, the SAVE file will indicate where to begin in the log file set with the appropriate log file sequence number.

Without Automatic CHANGELOG

To create a new logging cycle on a system without automatic CHANGELOG, do the following:

On the Primary System

- 1 Prevent users from accessing the primary database.
- 2 When no users are accessing the database, use DBUTIL to disable any further access.
- 3 Ensure the databases are up-to-date and synchronized by checking the STATUS display.
 - a The "LOGFILE record# =" values fro the primary and secondary systems should be equal.
 - b The "DATABASE opens =" parameter should show that no users are accessing the secondary database with 0 paths.

When this state is achieved, the STOP command can be issued to shutdown *Shadow D/R*.

On the Secondary System

- 4 Purge the SAVE file.
- 5 Backup the secondary database only, not the SAVE or log files.
- 6 Purge and rebuild the log file.

On the Primary System (Again)

- 7 Issue the MPE command:
:LOG logid, STOP
- 8 Purge or rename the log file and build a new one. (Refer to "Setting Up Databases and Logging (Primary System)" on page 21 for instructions.)
- 9 Issue the MPE command:
:LOG logid, START
- 10 Use DBUTIL to enable the database for access.
- 11 Allow users to access the database.
- 12 Start *Shadow D/R* when the secondary backup is done.

6 DS Line Failure and Recovery

DS Line Failure Overview

In the event of a DS line failure, *Shadow D/R* simply distributes the transmission load over the remaining DS lines. If the last DS line becomes unavailable, *Shadow D/R* will direct activities on the primary and secondary systems until the DS lines are restored.

The following activity takes place on the primary system:

- 1 The CP on the primary system informs the *Shadow D/R* operator of the DS line failure.
- 2 The CP begins to monitor the number of records added to the primary log file (records that cannot be transmitted) as users continue to update the database.

If the number of records in the primary log file exceeds the taping threshold (see "MAXRECS=100:" on page 18), a backup tape will be produced (see "Secondary Failure and Recovery" on page 31).

The following occurs on the secondary system:

- 1 MPE aborts the SP, in most instances.
- 2 The RP detects the missing "heartbeats" and informs the *Shadow D/R* operator via the secondary system that the primary is no longer responding.

The operator is informed the first two times that a heartbeat is missed on either system and also periodically afterward. This activity remains in force until at least one DS line returns to *Shadow D/R* service.

When the first DS line is restored for *Shadow D/R* system use, the following occurs:

- 1 The CP logs on to the secondary system and creates a new SP.
- 2 The new SP checks to see if an RP already exists. If the RP is not there, the SP creates one.
- 3 The SP asks the RP for the last record number processed to the secondary log file and applied to the secondary database. This record number is returned to the CP on the primary system and Shadowing begins with the record immediately following that record number in the log file.
- 4 Transmission of records continues at high speed until the secondary database and log file are back in step with the primary system. Any incomplete transactions are completed automatically.

DS Line Recovery Procedures

DS line failures are the most straightforward type of failure to resolve, because the *Shadow D/R* system remains active on both the primary and secondary systems.

Recovering One of Several DS Lines

If multiple DS lines are used and one line shuts down, the other lines will continue transmission. A warning message will display on the system console to inform the operator of the downed line.

To recover the DS line, do the following:

- 1 Determine why the DS line failed and attempt to resolve the problem.
- 2 Issue a SYSTEM command specifying the configuration file for the downed DS line.
- 3 Once the line has been restored, run SHADOWMP and issue a USE command, specifying the DS line number. *Shadow D/R* will attempt to re-establish communications.

Recovering All DS Lines

If all DS lines configured for *Shadow D/R* are lost, the *Shadow D/R* operator will receive a warning message on the system console.

To recover all DS lines, do the following:

- 1 Repeat the DS line procedure, "Recovering One of Several DS Lines" on page 29, for each of the downed lines.
- 2 If the CP determines the count of untransmitted records exceeds the taping threshold, it will prompt the *Shadow D/R* operator to mount a tape.
 - a Mount a tape on the tape drive.
 - b Reply to the tape request at the system console.Log records will be sent to tape, starting with the first record not previously transmitted.
- 3 Repeat the DS line recovery procedure for each of the downed DS lines.

After the CP transports enough records to tape and the number of records **not** sent falls below the taping threshold, it will release the tape. Once the CP's transmissions are up-to-date, the tape is no longer needed by the *Shadow D/R* system.


7 Secondary Failure and Recovery

Secondary System Failure Overview

If the secondary system fails, the CP on the primary system will not receive confirmation of its last transmission to the SP. To the CP, there is no difference between this situation and a simple DS line failure.

The following activities take place after a secondary system failure:

- 1 The CP on the primary system informs the **Shadow D/R** operator of the secondary failure.
- 2 The CP begins to monitor the number of records added to the primary log file (records that cannot be transmitted) as users continue to update the database.
- 3 If the number of records in the primary log file exceeds the taping threshold (see "MAXRECS=100:" on page 18), the CP notifies the tape process (TP).
- 4 The TP prompts the **Shadow D/R** operator to mount a tape. When the tape request is answered, the TP reads directly from the log file and writes the records to tape as shown in the following diagram.



Tape Process

All log file records not previously sent to the secondary system are written to tape. This ensures that some form of Shadowing occurs even when the secondary system is inoperable.

When the secondary system is returned to service, the normal catch up process operates automatically to bring the secondary system up to date with the primary system.

The tape produced by the TP is used if the primary system fails while the secondary system is down. If both systems fail at the same time, the tape is used to bring the secondary system up-to-date without the primary, so that processing can continue (see "Primary Failure and Recovery" on page 34).

Assessing a Secondary Failure

If you are running MPE V with disc caching or MPE XL, there is a remote chance that *Shadow D/R* will restart successfully (with no errors), even though the secondary database, log and SAVE files are no longer synchronized. Enabling the secondary failure protection mechanism will insure secondary database integrity on both MPE V and MPE XL systems and allow you to use the minor failure recovery procedure after every secondary system failure. The performance penalty and extra operational requirements make this feature unnecessary for most *Shadow D/R* environments.

To help evaluate the necessity for use of this feature, review “Configuring the Secondary System” on page 20 and “MODE=1 Security Solution” on page 64.

Secondary system failures can be minor or major in the same way primary system failures are with regard to whether the database is intact or corrupt.

- If there was no activity to the secondary database, use the following minor failure recovery procedure.
- If the STATUS display shows there was activity generated by the RP to the database, use the major failure recovery procedure.

If the RP was updating the database at the time of failure, there is really no conclusive way to determine if the database at the time was intact. If the secondary fails, you should verify the RP was not updating the database at the time of the failure. This can only be accomplished by confirming the secondary was current with the primary system before the failure occurred.

- The last STATUS display should show the outstanding records are equal to zero.
- The primary and secondary log file record counts are equal.
- No outstanding open transactions are indicated at the bottom of the secondary side of the display.

The following example shows that the secondary system is current with the primary system.

>STATUS					
=====		Shadow	STATUS	=====	
			1:34:57		
*PRIMARY			SECONDARY		
LOGFILE	record#	= 4675 (001)	LOGFILE	record#	= 4675 (001)
	time	= 1:34:54		time	= 1:34:54
	outstanding	= 52 recs	FLAGS	Shadow Mode	
	max ever	= 4675 recs			
SENT	logical	= 4875 recs	DATABASE	opens	= 0 paths
	physical	= 198 blocks		transact	= 0 outst
				Open mode	= 1 LOCK
	[current status]			[current status]	
=====					
LINES OPEN:SYSTEMB					
=====					

STATUS Display (secondary and primary systems are current)

Enabling Rollback for Easier Recovery

Protecting data for easier recovery in the event of a secondary system failure is optional. Enabling rollback on the secondary database guarantees the logical integrity of the secondary databases after a secondary system failure, so you can simply restart the *Shadow D/R* software and the program will resume data replication from the point of failure.

This recovery feature is useful if the following is true:

- Secondary database activity is low (the amount of data replicated in the secondary is relatively small).
- The integrity of the databases is guaranteed in the event of a secondary failure.
- Secondary system failures are frequent.

To successfully recover *Shadow D/R* from secondary system failures, the following requirements must be met:

- All database, SAVE and log files must be logically intact and up-to-date after a system failure or halt.
- The SAVE file must reflect all the most recent upgrades to that database.

This requires you to enable the secondary database for rollback recovery, which insures this recovery state and eliminates the need for secondary database restoration. However, it can also affect system performance on the secondary system. (See "MODE=1 Security Solution" on page 64.)

This procedure enables the secondary database for rollback recovery and, thereby, ensures physical I/O by TurboIMAGE to the databases. If the *Shadow D/R* program detects that rollback is enabled, it will make sure that physical I/O is performed for all writes to the SAVE file (MPE V and MPE XL).

- A "warm start" procedure must be used to recover an MPE V system.
- A "start recovery" procedure must be used to recover an MPE XL system.

To set up the secondary database for recovery, follow the instructions in "Configuring the Secondary System" on page 20.

Recovering the Secondary System

Recovery Procedure for a Minor Secondary System Failure

A system failure is considered to be minor if a remote process on the secondary system was mistakenly aborted or if the database on the secondary system is known to be intact. In situations where an RP or an SP job was mistakenly aborted, *Shadow D/R* will recover whether disc caching is enabled or not.

Recovery Procedure for a Major Secondary System Failure

- 1 Restore the database files and SAVE file from the last backup of the secondary system.
 - If the backup was made at the beginning of the cycle, purge or rename the SAVE file.
 - If the backup was made during a log cycle or if you are using automatic CHANGELOG, preserve the SAVE file.
- 2 On the primary system, issue the USE command to enable the DS lines and start *Shadow D/R* again.

8 Primary Failure and Recovery

Primary System Failure Overview

A primary system failure can be an operating system failure or a failure of some component of the computer system, such as a disc drive. The severity of the problem will dictate the procedure needed to recover.

Minor Primary System Failure

A system failure is considered minor if the database remains intact and suffers no damage as a result of the failure. This is usually the case if no one was updating or modifying the database when it failed.

To recover from a minor system failure, do the following:

- 1 Restart MPE.
- 2 Resume logging by issuing the command:
:LOG logid,RESTART
- 3 Resume *Shadow D/R* operation by running SHADOWMP and issuing a START command.
- 4 Allow users access to the database.

Major Primary System Failure

In typical environments, you can easily assess damage by using the DBRECOV "control stats" option. This reports database activity from the log file. If the DBRECOV report indicates read or write access to the database at the time of the system failure, the failure would be considered major, not minor.

When the primary system incurs a major system failure, the SP on the secondary system is usually aborted by MPE. The RP detects the missing "heartbeats" (see "The Heartbeat" on page 52) and informs the *Shadow D/R* operator. *Shadow D/R* will repeat this message until the situation is corrected.

The *Shadow D/R* operator can "switch" systems, so the secondary system acts as the primary system:

- The operator of the secondary system issues a SWITCH command using the manager process (MP) on the secondary system.
- The RP looks for a tape produced on the primary system before it failed. (A tape is produced only if the contents of the log file on the primary system exceeded the maximum taping threshold value.) The RP uses this tape to update the secondary database and log file to the point when the primary system failed.
- When the RP completes the update process, the *Shadow D/R* operator can start a new logging cycle with or without a database backup (depending on time constraints). All primary system activity is switched to the secondary system where processing can continue.
- The *Shadow D/R* operator runs the MP and issues the START command. This creates a CP process on the new primary (original secondary) system.
- The new CP shortly detects the failed DS lines or the missing responses from the secondary system (because the new secondary system is still down) and initiates the tape process on the new primary system.

- When the maximum taping threshold is exceeded, the new TP will write records directly from the log file to tape.

The **Shadow D/R** software recognizes the original secondary system as the new primary system in every respect. The users can simply switch their terminals to the second machine and resume processing as soon as the **Shadow D/R** operator tells them the system is ready.

After the new secondary system is restored, the system can remain as it is, with the roles of the original primary and secondary systems switched. The **Shadow D/R** software will automatically update the secondary system.

The systems can be returned to their original primary/secondary roles. The **Shadow D/R** operator on the new primary system issues a STOP command and the roles are reversed, provided all transactions are sent from the new primary system to the new secondary system and no new transactions are applied to the database. (Please refer to "Returning Systems to Their Original State" on page 39.)

Recovering from a Major Primary Failure

Choose one of the following recovery procedures, depending on how quickly the secondary system must be available to users:

- | | |
|-----------------|--|
| Fast Recovery | The fast recovery procedure allows the quickest access to the database by users, but the recovery implications are greater should a primary system failure occur after the SWITCH and before a secondary database backup.

Systems with CHANGELOG should use the procedure on page 35.
Systems without CHANGELOG should use the procedure on page 36. |
| Secure Recovery | If more time is available, follow the procedure for a more secure recovery.

Systems with CHANGELOG should use the procedure on page 38.
Systems without CHANGELOG should use the procedure on page 38. |

Fast Recovery Procedure

In following recovery procedure, "System A" refers to the original primary **Shadow D/R** system and "System B" refers to the original secondary system.

- | | |
|-----------------|---|
| System B | <ol style="list-style-type: none">1 Run SHADOWMP and issue SYSTEM and SWITCH commands.
This will cause the Shadow D/R program to terminate and produce listings of the status of each database user. You may apply any transactions stored to tape (from primary overflow backup) at this time.2 The SWITCH command renames the log file set and the SAVE file:<ul style="list-style-type: none">• The log file set is renamed "zog####," where "####" represents the sequence number. For example, the log file set, "log001 to log008," would be renamed "zog001 to zog008."• The SAVE file is renamed "ZAVE filename."Save both of these files to tape. |
|-----------------|---|

NOTE: If this operation is likely to take too much time, it can be performed later, after allowing users access to the database.

- 3 Issue a RELLOG command and then a GETLOG command to reset the log file sequence to "001." Build a log file with the sequence number "001."
- 4 Start logging and ensure that all databases that are part of the *Shadow D/R* system are enabled for logging with DBUTIL.
- 5 Allow users access to the database. If the zog#### and ZAVE files were not stored earlier, store them to tape.
- 6 Start *Shadow D/R* by running SHADOWMP and using the START command (verify the DS lines are not open or available in case the primary system has become operational in the meantime).

A backup of the transactions created on the new primary will be created if you have set an appropriate MAXRECS taping threshold. Since System A (the original primary system) is down, a tape request will appear at the console. The tape will store records which cannot be sent to the original primary.

Mount this tape and reply to the system console tape request.

System A

- 7 Restore the database files and the SAVE file from the last backup of the secondary system. For automatic CHANGELOG environments, the SAVE file indicates the log file sequence number and log record number at the point at which the database was locked up. You will need the SAVE file, even if the backup was from the start of a new log cycle.
- 8 Rename the last log file in the log file set to any other name.
- 9 Restore the zog#### file with the renamed sequence number.
- 10 Rename the restored zog#### file back to its original log#### name.
- 11 Run the SHADOWMP program and user the SYSTEM command to identify the current configuration file.
- 12 Issue the RECOVER command on the MP on System A.
- 13 Answer "No" to the MP's tape request. (You are saying that you do not have a *Shadow D/R* tape of unsent log file records to be added to the log file on the primary system.)
 - The MP will bring the database and log file up-to-date and produce a report to be examined by the *Shadow D/R* operator showing the last complete transaction for each user of the database and any incomplete transactions which it removed from the log file.
 - The MP will rename the log and SAVE files to zog#### and ZAVE again. The backup database is now in exactly the same state as System B before the users began entering transactions.
- 14 Build a new log file identical to the log file built on System B.

Both Systems

- 15 Up the DS lines on both the primary and secondary systems.
- 16 Issue the USE command on System B to start *Shadow D/R* activity again.

Shadowing will occur now automatically, but in reverse of the original direction.

Fast Recovery Procedure for Non-CHANGELOG Environments

In following recovery procedure, "System A" refers to the original primary *Shadow D/R* system and "System B" refers to the original secondary system.

System B

- 1 Run SHADOWMP and issue SYSTEM and SWITCH commands. This will cause the *Shadow D/R* program that is currently running to terminate and produce a list of the status of each database user.

- 2 The SWITCH command renames the log file set and the SAVE file:
 - The log file set is renamed "zog####," where "####" represents the sequence number. For example, the log file set, "log001 to log008," would be renamed "zog001 to zog008."
 - The SAVE file is renamed "ZAVE filename."

Store the database files and the zog#### and ZAVE files to tape.

NOTE: If this step will take too long, it can be performed later, after allowing users access to the database system.

- 3 Build a new log file.
- 4 Start logging (verify that all database that are part of the *Shadow D/R* system are enabled for logging).
- 5 Allow users access to the database. If the zog#### and ZAVE files were not stored earlier, store them to tape now.
- 6 Start *Shadow D/R* by running SHADOWMP and using the START command (verify the DS lines are not open or available in case the primary system has become operational in the meantime).

A backup of the transactions created on the new primary will be created if you have set an appropriate MAXRECS taping threshold. Since System A (the original primary system) is down, a tape request will appear at the console. The tape will store records which cannot be sent to the original primary.

Mount this tape and reply to the system console tape request.

- 7 Once System A is up, restore only the database files and the SAVE file from the last backup of the secondary system.
 - If the backup was made at the beginning of a log cycle, purge or rename the SAVE file.
 - If the backup was made during a log cycle, preserve the SAVE file.
- 8 Rename the old file to any other name.
- 9 Restore the zog file from tape.
- 10 Rename the zog file its original log file name.
- 11 Run SHADOWMP and use the SYSTEM command to identify the current configuration file.
- 12 Issue the RECOVER command.
- 13 Answer "No" to the tape request. (You are saying that you do not have a *Shadow D/R* tape of unsent log file records to be added to the log file on the primary system.)

The MP will bring the database and log file up-to-date and produce a report to be examined by the *Shadow D/R* operator showing the last complete transaction for each user of the database and any incomplete transactions which it removed from the log file.

The MP will rename the log and SAVE files to zog#### and ZAVE again. The backup database is now in exactly the same state as System B before the users began entering transactions.
- 14 Build a new log file identical to the log file built on System B.

Both Systems

- 15 Up the DS lines on both the primary and secondary systems.
- 16 Issue the USE command on System B to start *Shadow D/R* activity again. Shadowing will now occur automatically, but in reverse of the original direction.

Slow Recovery Procedure

If your system is not equipped with CHANGELOG, please use the procedure, "Slow Recovery Procedure for Non-CHANGELOG Environments" on page 38.

In following recovery procedure, "System A" refers to the original primary *Shadow D/R* system and "System B" refers to the original secondary system.

- | | |
|---------------------|--|
| System B | <ol style="list-style-type: none"> 1 Run SHADOWMP and issue SYSTEM and SWITCH commands. This will cause the <i>Shadow D/R</i> program that is currently running to terminate and produce a list of the status of each database user. 2 The SWITCH command renames the log file set and the SAVE file: <ul style="list-style-type: none"> • The log file set is renamed "zog####," where "####" represents the sequence number. For example, the log file set, "log001 to log008," would be renamed "zog001 to zog008." • The SAVE file is renamed "ZAVE filename." <p>Store the database files and the zog#### and ZAVE files to tape.</p> 3 Issue a RELLOG command and a GETLOG command to reset the log file sequence to "001." 4 Start logging. Verify that all database that are part of the <i>Shadow D/R</i> system are enabled and store them to tape now. 5 Allow users access to the database. If the zog#### and ZAVE files were not stored earlier, store them to tape now. 6 Start <i>Shadow D/R</i> by running SHADOWMP and using the START command (verify the DS lines are not open or available in case the primary system has become operational in the meantime). <p>A backup of the transactions created on the new primary will be created if you have set an appropriate MAXRECS taping threshold. Since System A (the original primary system) is down, a tape request will appear at the console. The tape will store records which cannot be sent to System A.</p> <p>Mount this tape and reply to the system console tape request.</p> |
| System A | <ol style="list-style-type: none"> 7 Restore only the database files from the secondary backup tape that was made after the SWITCH command. 8 Build a new log file identical to the log file built on System B. |
| Both Systems | <ol style="list-style-type: none"> 9 Up the DS lines on both the primary and secondary systems. 10 Issue the USE command on System B to start <i>Shadow D/R</i> activity again. Shadowing will now occur automatically, but in reverse of the original direction. |

Slow Recovery Procedure for Non-CHANGELOG Environments

In following recovery procedure, "System A" refers to the original primary *Shadow D/R* system and "System B" refers to the original secondary system.

- | | |
|-----------------|---|
| System B | <ol style="list-style-type: none"> 1 Run SHADOWMP and issue SYSTEM and SWITCH commands. This will cause the <i>Shadow D/R</i> program that is currently running to terminate and produce a list of the status of each database user. |
|-----------------|---|

- 2 The SWITCH command renames the log file set and the SAVE file:
 - The log file set is renamed "zog####," where "####" represents the sequence number. For example, the log file set, "log001 to log008," would be renamed "zog001 to zog008."
 - The SAVE file is renamed "ZAVE filename."Store the database files and the zog#### and ZAVE files to tape.
- 3 Build a new log file identical to the log file built on System A.
- 4 Start logging (verify that all database that are part of the **Shadow D/R** system are enabled and store them to tape now).
- 5 Allow users access to the database. If the zog#### and ZAVE files were not stored earlier, store them to tape now.
- 6 Start **Shadow D/R** by running SHADOWMP and using the START command (verify the DS lines are not open or available in case the primary system has become operational in the meantime).

A backup of the transactions created on the new primary will be created if you have set an appropriate MAXRECS taping threshold. Since System A (the original primary system) is down, a tape request will appear at the console. The tape will store records which cannot be sent to System A.

Mount this tape and reply to the system console tape request.
- 7 Once System A is up, restore only the database files from the secondary backup tape that was made after the SWITCH command.
- 8 Build a new log file identical to the log file built on System B.

Both Systems

- 9 Up the DS lines on both the primary and secondary systems.
- 10 Issue the USE command on System B to start **Shadow D/R** activity again. Shadowing will now occur automatically, but in reverse of the original direction.

Returning Systems to Their Original State

After completing one of the failure recovery procedures, you may want to return the database systems to their original roles as primary and secondary, that is, with **Shadow D/R** occurring in the direction used before the SWITCH command.

Procedure for Systems with CHANGELOG Environments

The following instructions are for systems with automatic CHANGELOG environments. For non-CHANGELOG environments, use the instructions on page 40.

- 1 Stop all database activity on System B (the current primary system/original secondary system).
- 2 Verify the databases are up-to-date and synchronized by checking the status display.
 - The "LOGFILE record#=" values for both systems should be equal.
 - The "DATABASE opens=" parameter should show that no users are accessing the secondary database (there are 0 paths).

>STATUS											
=====				Shadow		STATUS		=====			
						1:34:57					
*PRIMARY					SECONDARY						
LOGFILE	record#	=	4675 (001)			LOGFILE	record#	=	4675 (001)		
	time	=	1:34:54				time	=	1:34:54		
	outstanding	=	52 recs				FLAGS	Shadow Mode			
	max ever	=	4675 recs								
SENT	logical	=	4875 recs			DATABASE	opens	=	0 paths		
	physical	=	198 blocks				transact	=	0 outst		
					Open mode		=	1 LOCK			
[current status]					[current status]						
LINES OPEN:SYSTEMB											
=====											

STATUS Display (paths at 0)

- 3 STOP *Shadow D/R*.
- 4 Stop the logging process.
- 5 Disable the System B database for logging.
- 6 Store the log file to tape.
- 7 Restore this log file to System A.
- 8 Do a RELLOG and GETLOG if the log file sequence number needs to be reset to sequence 001 on System A.
- 9 Build a new log file with the sequence number 001 on both systems.
- 10 Purge the SAVE file on System A.
- 11 Do a backup of the secondary system to keep log files and databases in sequence.
- 12 Start *Shadow D/R*.

Procedure for Systems without CHANGELOG Environments

- 1 Stop all database activity on System B (the current primary system/original secondary system).
- 2 Verify the databases are up-to-date and synchronized by checking the STATUS display.
 - The "LOGFILE record#=" values for both systems should be equal.
 - The "DATABASE opens=" parameter should show that no users are accessing the secondary database (there are 0 paths).
- 3 STOP *Shadow D/R*.
- 4 Stop the logging process.
- 5 Disable the System B database for logging.
- 6 Store the log file to tape.

- 7 Restore this log file to System A.
- 8 Store the SAVE file from System A to tape.
- 9 Restore the SAVE file to System B.
- 10 Enable the database for logging on System A.
- 11 Restart system logging on System A.
- 12 Allow users access to the database on System A.
- 13 Start *Shadow D/R*. The original primary is the primary system.


Example of a Primary System Failure and Recovery

The following diagrams illustrate the states of the database, log file and SAVE file during a fast recovery procedure on a primary system. The arrows indicate the direction of data transport between the primary and secondary database systems.

In following illustrations, "A" indicates the original primary *Shadow D/R* system and "B" indicates the original secondary system.


Operating Under Normal Conditions

The following diagram shows the *Shadow D/R* system operating under normal conditions. A backup of the database, log file and SAVE file was made recently, during a log cycle.



Switching Roles After a Primary System Failure

The next diagram shows the database systems after a primary system failure.



The following events take place after a primary system failure:


- System B**
 - The operator on system B issues a SWITCH command. The SWITCH operation processes the remaining log transactions from the database and produces a report.
 - The previous log file and SAVE file are renamed to files beginning with the letter "z."
 - A new log is built, user logging is started and users are allowed to modify the database.
- System A**
 - The "ZOG" file is restored and it is renamed its original name, "LOG" file.
 - A RECOVER command is issued to bring the database from a backup state to the state at the point of failure.
 - The restored log file, LOG', is used to recover transactions that occurred after backup (indicated by the checkpoint in SAVE) and before the failure (to the end of LOG').
 - When LOG' is processed completely, incomplete transactions are backed out and a report is produced.

Updating the New Secondary System

The following events take place to update the new secondary system (System A):

- When System A is recovered, the log and SAVE files are again renamed to files beginning with the letter "z."
- A new log file is created to enable System A's transition to primary status.
- Upon the initial startup of **Shadow D/R** in the new direction, the RP checks the SAVE file checkpoint records to determine where to begin processing in the primary log file.
- Since there is no SAVE file, the RP tells the master CP to begin data capture starting with the first log file record. A SAVE file is created by the RP on the secondary system at this time.
- The secondary system will catch up by processing the transactions which occurred during the period of recovery on System A (LOG' plus any new transactions that occurred since System B became the new primary system).

The following diagram illustrates this process:



Returning the Systems to Their Original Roles

Before the systems are returned to their original roles, the operator verifies (using the STATUS command) that all records have been sent and updated on the secondary system (System A). Checking the parameters in the STATUS display is the only way to ensure database activity is completed on both systems.


NOTE: A STOP command only looks for a quiet period shutting down *Shadow D/R*. You will lose synchronization of our primary and secondary databases if the STOP occurred before all database activity completed.

The operator restores the log file from System B onto System A to keep the log files synchronized. This is necessary because a few CLOSELOG records were added to the System B log file when the logging process was stopped.

The SAVE file is also moved to the new secondary (system A). It is highly recommended that the System B database backup is made at this time.

Logging is restarted on System A, users are allowed access to the database and *Shadow D/R* is started (Shadowing in the opposite direction).

The following diagram illustrates the results of this process. Notice the original primary system (System A) is operating as the primary system again.



9 Data Capture

This section focuses on the mechanisms which operate on the primary system, specifically, how log records are captured from the primary log file and prepared for data transport to the secondary.

Overview of the Data Capture Process


When a user calls any of the TurboIMAGE intrinsics that might need logging, a test is made to see if the database is enabled for logging. If logging is enabled in addition to modifications to the data with the database, the pertinent data is written to the log file using the available MPE logging intrinsics.

All of the interfacing to the log file data under normal circumstances is performed by the CP. When *Shadow D/R* is started, it does the following:

- Determines the name of the logging ID as well as the current log file name.
- Uses the logging identifier name to access the logging process using the MPE intrinsic, OPENLOG.
- Opens the disc log file using a standard FOPEN intrinsic, so the file may be accessed by the MPE file system intrinsics as needed.
- Periodically checks the status of the logging process using LOGSTATUS to see if any log file records have been added which have not been seen by *Shadow D/R*.

(If there are additional log files records, FLUSHLOG is used to ensure the records are physically posted to the log file on disc.)

- Reads the necessary records from the disc log file using the MPE file system intrinsics.
- Buffers these records within the SHADOWCP process and transfers them to the secondary system (refer to the low diagram on page 45).



Data Capture

MPE User Logging Facility

The user logging system is a facility within the MPE operating system that provides a mechanism for users to deposit information into a centralized storage facility. The nature of this information is completely dependent upon the use of this facility.

The method of physical storage can be disc, tape or cartridge tape. The actual medium used is not readily apparent to the user of the system.

A particular logging facility is assigned a logging identifier by the original creator. This must be a unique name since it is used by all subsequent commands and intrinsics to communicate with the particular logging entity. The management of the facility is usually part of the function of the system operator, although any given logging identifier appears to belong to its original creator.

The following MPE commands are used to create the logging identifier:

- GETLOG creates a logging identifier.
- RELOG purges a logging identifier from the system.
- LOG manages the activation and deactivation of the actual logging process.

In order to manage the logging facility or use it, the user must be assigned the LG capability. Once activated, a logging identifier becomes a separate process within MPE and the logging-related commands and intrinsics communicate with this process through assigned extra data segments and a defined protocol.

Within MPE there is a data segment that keeps management information about all current logging identifiers. The maximum number of user logging processes is a SYSDUMP configuration parameter, since the data segment must be sized accordingly. Whenever a particular logging identifier is activated, a process is created and a data segment is assigned which can be used for communications, buffering and control of the logging process. The process and data segment are identified in the global logging data segment in the entry for the particular logging identifier. As additional users attach themselves to the logging process, an entry is made in the individual log process data segment to keep track of them. Therefore, the maximum number of users per logging process is also a configuration parameter and is used to size individual log process extra data segments.

MPE Logging Intrinsic

There are a number of intrinsics within MPE which make it fairly simple to operate and use the logging facilities. With the exception of the OPENLOG intrinsic which returns the unique user index, all other intrinsics require the unique user index as a parameter to identify the user to the logging process. The intrinsics are as follows:

OPENLOG

```
OPENLOG (index, logid, password, mode, status);
```

The OPENLOG intrinsic initiates access of a particular logging process for the caller. It returns the index which is employed as the user's identifier for calls to the other logging intrinsics.

WRITELOG

WRITELOG (index, data, length, mode, status);

The WRITELOG intrinsic writes the data contained in the data array into the log file. Since the logging record is fixed in length at 128 words, it is quite possible that the data will not fit into a single record. In this case, the data is split into multiple records that are guaranteed to be adjacent in the log file.

LOGSTATUS

LOGSTATUS (index, loginfo, status);

The LOGSTATUS intrinsic provides a simple interface for the user to determine the state of the particular logging process. It includes file size information, current space used and available in the file as well as a count of current users of the logging process.

FLUSHLOG

FLUSHLOG (index, status);

The FLUSHLOG intrinsic provides a means of ensuring that all data logically written to the log file is physically posted to the disc media. *Shadow D/R* uses this intrinsic to ensure the disc copy is up to date prior to reading the data. As a side effect, when activity to the logging process is slow, the *Shadow D/R* software will ensure that only a specified period of time elapses before the data partially filling a buffer is flushed to disc.

CLOSELOG

CLOSELOG (index, mode, status)

The CLOSELOG intrinsic breaks the communication between the user and the particular logging process.

LOGINFO

```
LOGINFO (index, status      [,item1, val1]
                                [,item2, val2]
                                [,item3, val3]
                                [,item4, val4]
```

The LOGINFO intrinsic is used to obtain information from the user logging data segments. Up to four different types of information can be accessed in one cell. Information available includes log file size, remaining space in the log file, current and previous log file names (and sequence numbers), and total records written in the current file or the whole set.

MPE Logging Commands

In order to manage the user logging facilities, MPE has several commands that can be used to provide maintenance functions, as follows:

GETLOG

:GETLOG logid;LOG=log file, {DISC/TAPE/SDISC/CTAPE};[PASS=password]

The GETLOG command is used to create a new logging identifier within the system. The logid must be unique within the system. A warning is printed if the AUTO option is used and the log file created does not end in 001.

SHOWLOG-STATUS

:SHOWLOGSTATUS

The SHOWLOGSTATUS command is used to find out what logging identifiers have been created for the system and the current status of each logging identifier.

ALTLOG

**ALTLOG logid;LOG=log file, {DISC/TAPE/SDISC/CTAPE} [PASS=password]
[;{AUTO/NOAUTO}]**

The ALTLOG command allows you to alter the characteristics of a particular logging identifier when it is not active.

LOG

**LOG logid [,START
 [,STOP]
 [,RESTART]**

The LOG operator command provides the mechanism for activating and deactivating a logging process specified by a particular logging identifier.

RELLOG

:RELLOG logid

The RELLOG command provides the mechanism to delete a logging identifier from the system permanently.

CHANGELOG

:CHANGELOG logid [DEF=device]

The CHANGLOG command closes the current log file and opens a new file with the next sequence number.

TurboIMAGE Database Logging Facility

The TurboIMAGE transaction logging facility is a TurboIMAGE feature that uses the MPE user logging facility for its implementation. As such, the TurboIMAGE subsystem is simply a user of the base MPE intrinsics.

The TurboIMAGE database environment provides several specialized intrinsics for the user as well as the built-in logging intrinsics required to reflect a particular database activity. Because of this, logging of database transactions can be done into the same logging identifier that other database users are sharing.

The management of opening, closing, writing to and flushing the logging identifier is transparent to the user application code although logging-related TurboIMAGE intrinsics can be used to advantage.

In order to activate user logging for a database, the creator of the database must do the following:

```
:RUN DBUTIL.PUB.SYS  
>SET databasename LOGID=logid  
>ENABLE databasename FOR LOGGING  
>EXIT
```

TurboIMAGE will ensure that whenever there is a current user of the database, logging must be active for the logid, and all relevant database transactions will be reflected into the specified user logging system.

TurboIMAGE Intrinsics

When logging has been set up for a database, all TurboIMAGE intrinsics that modify the contents of the database will get logged.

DBOPEN	Wherever the database is opened, the information from the intrinsic will be logged.
DBPUT	Wherever new data is added to a dataset, all information required to reproduce the intrinsic is logged.
DBUPDATE	Wherever the contents of the existing data entry is changed, both the original contents as well as the new contents are logged.
DBDELETE	When an existing entry is deleted, the data entry from the entry is logged.
DBCLOSE	Whenever the database is closed, this is recorded in the log file.

TurboIMAGE Logging-specific Intrinsics

In addition to the automatic logging of the activity generated by the normal TurboIMAGE intrinsics, there are several intrinsics which are part of TurboIMAGE that relate only to logging. In the event that logging is not enabled, the intrinsics are ignored. These intrinsics and their functions are as follows:

DBBEGIN	The DBBEGIN intrinsic provides an indication in the log file that this is the beginning of a series of entries for this user that are part of one logical transaction. In the event that the log file is used for failure recovery, the recovery software can use this marker to help identify a single logical transaction that should either all be recovered or none of it recovered.
DBDEMO	The DBDEMO intrinsic provides a mechanism for the user to insert commentary into the log file as appropriate. This commentary has no bearing on any subsequent recovery activities.
DBEND	The DBEND intrinsic provides an indication in the log file that a logical transactions started by a DBBEGIN is now complete. This can be used in recovery logic to identify the ending of a single logical transaction.

Shadow D/R Data Capture Mechanism

The user logging facility provides a reliable and stable method of collecting information reflecting changes being processed against a database. Because of this, user logging was selected as the interface by which *Shadow D/R* initially captures this data from the primary system. Because no changes were made to TurboIMAGE code, all legitimate uses of TurboIMAGE databases through the TurboIMAGE intrinsics will be subject to this data capture.

NOTE: The use of some tools to manipulate the database do not reflect their changes in the log files which *Shadow D/R* relies on as its source of input. This means that the use of such products as ADAGER<DBCHANGE, DBUITL's erase function, and other tools will violate the integrity of the *Shadow D/R* environment if not used very carefully. Either the primary database should be restored to the secondary after any manipulation by any of these tools, or any of the changes done on the primary by these tools should identically be done to the secondary database to maintain consistency on both systems.

Data Capture Intrinsics

- | | |
|------------------|--|
| FOPEN | The <i>Shadow D/R</i> program opens the log file using the FOPEN intrinsic. |
| FREADDIR | Once the log file is open, data in the log file is read using relative record numbers with the FREADDIR intrinsic. |
| LOGSTATUS | Prior to all read operations, the LOGSTATUS intrinsic is used to determine how many records are active in the log file. |
| FLUSHLOG | If necessary, the FLUSHLOG intrinsic is used to ensure the memory resident buffers are physically written to disc prior to the FREADDIR intrinsic. |

WAIT Time

Whenever the *Shadow D/R* CP has no log records to send to the secondary system and no MP commands to process, it waits on an FREAD on the command message file to the MP. When this read completes, another check is made to see if any log records have been posted to the log file.

Timeout Implementation

The timeout is implemented by calling FCONTROL on the command message file from the *Shadow D/R* MP (whether the MP is there or not). This means that the timeout may actually only be specified in seconds, not milliseconds, as is implied by the WAIT parameter in the configuration file. In fact, any value supplied will be rounded to the next higher or lower integer for use in the FCONTROL intrinsic. If the value supplied is less than 500 milliseconds, resulting in a round down to 0, the value is flagged as an error.

Wait Time Selection

The choice of WAIT time is dependent on the load you want placed on the system and the maximum time that the secondary system may be behind the primary system.

Actually, the choice is rather limited, as the maximum WAIT time is 10 seconds. However, selecting a low time will cause a call to the LOGSTATUS intrinsic more frequently and, if the rate of writing to the log file is low, transmission of many buffers to the secondary system, each containing few log records. This makes poor use of the existing communications bandwidth and causes more primary system overhead. In any case, the *Shadow D/R* CP will never transmit to the secondary for more than 10 seconds at a time without checking for commands which the *Shadow D/R* MP has sent to be executed.

Log Record Blocks

Log File Reads

All log reads are done one block at a time, 32 log records maximum. If the WAIT parameter is small or the logging activity is low, the likelihood of reading fewer than 32 log records each physical read is high. When the interval is longer, or the activity is higher, the efficiency of each block read from the log file will rise, guaranteeing that as the need increases, the efficiency increases.

Record Transmission

When the log file has no more records to be read, the log records assembled in the transmission buffer are sent to the secondary system. If there are more records to be read from the log file, the transmissions buffer is sent when it contains 24 log records. Again, this means that if the WAIT interval is short or the log activity is long, there is a greater chance that the transmission buffer will be full.

OUTSTANDING Log Records

The OUTSTANDING value found in the STATUS display indicates the number of log records not yet read and transmitted to the secondary system.

This value is controlled by the rate of log activity, the frequency of transmission and the ability of both the DS line and the secondary system to keep up to the primary transaction load. As long as the OUTSTANDING value is greater than 32, all log file read should be full blocks and as long as the OUTSTANDING is greater than 24, the transmission buffers should all be the maximum of 24 log records each.

Activities of the Tape Process

Whenever there are any logging records that have not been transmitted successfully to the secondary system, there is a backlog. In reality, there will almost always be a backlog on a system that has an active database. The backlog is only a concern when more than a few records have accumulated. No matter how few records are backlogged, in any are in this state, they will be lost in the event of a failure of the primary system. To minimize the impact of such an occurrence, *Shadow D/R* provides a secondary method of storing the data off the primary system to tape where it is isolated from primary system failures.

Tape Process Activation

The communication process (CP) component of *Shadow D/R* is continually monitoring the status of the log file. At the same time, it always knows what the last file record successfully transmitted to the secondary system was. By comparing these pieces of information, the backlog can always be determined. Whenever a backlog exists, the number of records backlogged is compared to a configured parameter for the particular *Shadow D/R* environment. If the current backlog exceeds this specified threshold, the TP is activated on the primary system. This process independently monitors the status of the log file and writes new log file records out to tape as long as it is active.

Tape Process Deactivation

When the CP detects the backlog has fallen below the value configured for the particular environment, it communicates with the tape process again and tells it to stop copying the data to tape. At this point, the tape is rewound and unloaded. The drive is now released and made available to the system again.

Tape Process Creation

The TP is a son process of the CP and is the same program file started with the entry point, "TAPE." It independently and synchronously accesses the log file, flushing and reading the log blocks are written to tape. If there is no log records available to fill the current tape output buffer, the TP waits three seconds by reading the message file from the CP with a three second timed read. The TP is created at *Shadow D/R* system startup by the CP.

Tape File Creation

Operator interaction is required for mounting tape reels. Any number of reels may be created in a set of *Shadow D/R* log tapes. Each tape has a header record and a trailer record indicates whether any tapes follow in the set and what the next log record number should be. All records but the header record are 4096 words long, blocked into 32x128 word records. The trailer record, as the header record, is 40 words in length, but uses one block (4096 on tape due to padding of the block by MPE).

CHANGELOG Records

User Logging and CHANGELOG

With TurboIMAGE automatic CHANGELOG, the concept of the logging cycle has changed so that it is no longer necessary to stop the user logging process and the user's applications to switch to a new log file. When a log file fills up, a new log file is built, with a related name, and logging continues to this new file. It is possible to have up to 999 log files in a set of log files. As soon as one log file is filled, it may be stored to tape and then purged, if the space is needed.

Each time a file switch is requested, a final log record is written to the log file to indicate that the log records continue in a new log file. The name and type of the new log file are found in this record. In addition, the first record in the new log file is a record pointing to the previous log file in the set.

These records are used by *Shadow D/R* to ensure that records are read out of the correct log file on the primary system and put into the correct file on the secondary system. *Shadow D/R* tracks log file switches as they are done by the user logging subsystem, or as it reads the CHANGELOG records from the log files.

CHANGELOG Record Processing

There are two types of CHANGELOG records: type 12 at the beginning of a log file and type 13 at the end of a log file. The type 12 CHANGELOG record is simply passed to the secondary system, where it is written to the log file copy. Type 13 log records cause the *Shadow D/R* CP to close the current log file and open the new log file indicated in the CHANGELOG record, after first checking that the new log file is a disc log file (and not tape). The CP waits until the new log file has been built and is in use, if necessary. At the secondary, the same type 13 log record causes the current log file to be closed and the one indicated in the log record to be built with the same file size as the original previous log file. The CHANGELOG record is then written to the second user label of the SAVE file for use in START commands to indicate the current log file to be accessed. Upon initial startup of the RP when there is as yet no SAVE file, the RP places a dummy type 12 log record containing the name of the current (first in the set) log file.

The Heartbeat

A special status request is sent from the CP to the RP via the SP every 30 seconds to 40 seconds to provide a positive indication to the RP that the primary system is still up and functioning correctly. The RP expects to receive this heartbeat signal at a minimum of every 60 seconds. This allows a buffer of 20 seconds for busy systems or DS lines. If the RP misses a single heartbeat, it issues a message but does nothing further. At least two successive heartbeats must be status from primary to secondary systems. This status is used as the basis for the STATUS command on the secondary system and explains why the STATUS display on the secondary system appears to be slightly behind the CP STATUS display. The RP follows the following pattern for missing heartbeats:

- First miss, warning to the operator.
- Second miss, warning plus a message indicating potential trouble.

If a new heartbeat arrives after the warning, the RP resets and carries on.

10 Data Transport

This chapter explains the data communication between the CP on the primary and the SP on the secondary. *Shadow D/R* will use either RPM and NetIPC or DS PTOP intrinsics, depending on which intrinsics are used on your systems. The code logic for either data communication implementation for using multiple communication links, blocking data, transmission frequency and the transmission algorithm is the same.

Network Links

Shadow D/R must choose one of the following data communication services:

- NS/3000 for LAN.
- DS/3000.

Shadow D/R was originally implemented for the DS/3000 system using PTOP services. The process-to-process communication for NS/3000 Network Inter-process Communication (NetIPC) and Remote Process Management (RPM) is more flexible than PTOP, because the relationship between NetIPC processes is peer-to-peer rather than master-to-slave. However, the inherent design of *Shadow D/R* is a master-slave relationship. The NetIPC is implemented in the latest version of *Shadow D/R* and basically simulates DS PTOP design.

To implement NetIPC calls into the design of *Shadow D/R*, the DS PTOP calls were replaced with procedures which contain logic to switch to NetIPC or PTOP, whichever is being used. Hence, the procedure names: MMY'POPEN, MY'ACCEPT AND MY'REJECT replaced DS PTOP CALLS to POPEN, ACCEPT AND REJECT for EXAMPLE.

Using NetIPC and RPM allows *Shadow D/R* to run in compatibility mode on MPE XL systems to support MPE XL-MPE V mixed configurations. One version of *Shadow D/R* is still required for all systems in the *Shadow D/R* network. So, whether you are using point-to-point, X.25 or LAN, *Shadow D/R* is able to determine the appropriate intrinsics to call. This happens during the startup routine.

The transmission algorithm (see "Transmission Algorithm" on page 54) describes the logic used to determine which network software to use.

When the CP is started by the MP, a procedure opens all the necessary files and calls the OPEN' LINES procedure described in the transmission algorithm (see page 54).

The subsequent successful or unsuccessful call to PCCreate from MP'POPEN determines if NetIPC and RPM NS/3000 are installed on the system. If MY'OPEN indicates DS PTOP is to be used, a remote process with entry point "SLAVEDS" using DS PTOP will be created. Otherwise, a remote process with entry point, "SLAVE," will be created for NetIPC-RPM data communications. LINE'TYPE is a flag for procedures to determine whether to call NetIPC or the DS PTOP intrinsics. Since the data communication type is known at process creation of the secondary CP, the LINE'TYPE flag is set immediately after the corresponding entry point for the secondary CP.

Multiple Data Communication Links

If you have more than one data communication link and have configured them in the configuration file, review the following items.

Due to the single-threaded nature of the data communication facilities and the design of the transport mechanism, the throughput advantages from multiple lines is only minor. Adding more communication lines to the configuration file will not significantly increase data transport throughput. However, adding more lines will reduce the impact if you have a lot of traffic from non-*Shadow D/R* applications also using the same lines.

By providing multiple paths between the computers, you reduce the effects of single line failure since the other lines can continue processing. This can be a real advantage if the failure involves HP3000 hardware or modem hardware. If the failure is in the linkage provided by common carrier facilities, then both lines between the same two points over the same media will likely fail. In this case, supplying multiple paths would have no advantage.

Blocked Data

Advantages of Blocking Data

If you block the data transmitted you cause less overhead, since the code required to be executed to perform the actual transmission will be run less often and therefore cause less overhead on the system.

The throughput on the data communication linked will also be higher if the data is blocked. Just as with disc subsystems, the limitation on throughput is very severely limited to a fixed number of transmissions per second. By blocking, we move more data during our fixed number of transmissions and therefore optimize the throughput. Of course, on a slow line the actual transmission time will become significant; but as line speed increases, the benefits of blocking provide clearly superior line throughput rates.

Disadvantages of Blocking Data

The only real drawback to blocking the data is the amount of data lost within the buffers when the primary system fails. The larger the blocking factor, the higher the number of lost records is likely to be. This is compensated for by the *Shadow D/R* facility that allows you to specify that any data not sent after a specified period of time will be sent in a partial block when the timer elapses.

Transmission Frequency

Whenever a data buffer being accumulated within the CP becomes full, that buffer of data is transmitted to the secondary system. You may recall that this buffer size is hardcoded at 24 records.

Within the *Shadow D/R* configuration, a facility exists to specify the maximum time that should elapse between one physical transmission and the next (if there is something to transfer). This provides a mechanism to protect the primary system from the condition where several log file records are buffered up and then the database modification activity stops for a while, leaving a partially filled buffer un-transmitted.

Transmission Algorithm

The algorithm below uses a WAIT time of 10 seconds and ensures no missed heartbeats and that all commands from the MP get executed.

Logic Flow for Data Transmission:

Main execution loop:

1Heartbeat request.

2Command execution.

3Data transmission.

Must have at least one outstanding log record to transmit.

Place records into transmission buffer until full or outstanding is 0.

Flush buffer to secondary system in either of above cases.

If outstanding is 0, repeat main execution loop.

If more log records exist, check if 10 seconds has passed.

After 10 second transmission, repeat the main execution loop.

Role of the SP

The SP deblocks the transmission buffer and passes data into a message file. This message file is often referred to as the data "pipe." Prior to an ACCEPT (MY"ACCEPT if NetIPC is used) of the data, disc file space is checked. Commands and status are returned via the PCONTROL intrinsic or MY'CONTROL routine. The SP also monitors any problems with the RP and will report back to the CP.

11 Data Posting

The majority of work performed by *Shadow D/R* is done by the RP on the secondary system. An understanding of the order in which the RP file updates occur is very important.

The procedure is:

- 1 The RP pulls data from the data pipe which is filled by the secondary communications process (SP).
- 2 The replicate log file is updated with the current log records.
- 3 The SAVE file checkpoint record 0 is updated to reflect the current log file record number.
- 4 Continuation records are assembled to create the records for the database call.
- 5 The TurboIMAGE intrinsic which corresponds to the assembled log record is used to update the database. If the database is successfully updated, the SAVE file will be updated only according to certain conditions, which are described later in this chapter.

Also on the secondary, the SP builds all message files, deblocks the log records for sending to the RP and also helps to self-monitor system status.

Receiving and Validating Data

The data from the remote system is received by FREADING a message file specified by the "COMMANDS=" record in the *Shadow D/R* configuration file. This file, in fact, contains only user logging records, not commands. The message file is open with extended wait and has no read timeout. Record format of data is exactly the user logging record format as found in the log file in use on the main system.

The RP uses the knowledge of the log file record number last received to validate the incoming record number. This must be consecutive. If the record received has a record number preceding the expected number, the local log file is read and the received record is compared to the existing record. If the records are identical, it is assumed that the record was inadvertently retransmitted and nothing is done. If the record received has a higher record number than the expected record, it is assumed that records have been missed and/or the log file being used for transmission does not match the log file and SAVE file on the local *Shadow D/R* system. This would be a serious error and the RP would terminate immediately.

Receiving and Executing Commands

Commands and status requests are received from the primary system by FREADING a message file specified by replacing the first character of the data file name with an "X." The data file name is specified in the "COMMANDS=" record in the *Shadow D/R* configuration file. The message file is opened with extended wait and a 60 second read timer. NO-WAIT reads are used to allow receiving either commands or data at any time. There may be many outstanding data records at any time, but only one command at a time is allowed.

Read completion of the command file takes two forms under normal conditions:

- Data received.
- Timeout occurred.

If the completion indicates a file error, the RP will terminate (currently with a QUIT 300).

A read timeout indicates a missed heartbeat status request from the main system. A flag is set to indicate a missed heartbeat and a message is sent to the console. If the missed heartbeat flag is already set, this is not the first timeout and second flag is set to suppress additional messages to the console, and a message is generated to indicate that the remote primary is no longer communicating. It is then the operator's responsibility to determine whether the data communication line(s) are down or the primary has failed.

The normal data and command read loop is resumed to allow continuing the RP function as soon as the other parts of *Shadow D/R* are again operational. As soon as any command, heartbeat status or otherwise, is received from the primary system, the missed heartbeat flags are reset.

A normal read completion (with error) on the command file will cause the request to be processed. The request may have two sources: the SP or a local (to the RP) copy of the *Shadow D/R* manager process (MP). Only those commands associated with SWITCH or RECOVER are allowed from a local MP process. Remote commands originating either from the slave CP or the remote MP are used to start up and control the RP or ask for status.

In the description of command requests below, the req# is the request number which is returned in Status Array (0) in CP'REQUEST.

Command Requests

From the MP

WHO'R'U	Request for identification (req# 0). Set response bit, reply "RP."
OPENBASE	Open all requested database for recovery (req# 2). Report success or failure of DBOPEN call.
SWITCH	Switch to backup databases (req#8). Roll out incomplete transactions, print report. Close all databases, rename log and SAVE files.
RECOVER	Recovery from log file to database(s) (req# 11). Use SAVE file as start point, process to end of log file and process SWITCH code.

From the SP

GET'PARMS	Return current environment of RP (req# 0). Reply with the last log record received. (Currently calls QUIT 232 if can't read log.)
SHUTDOWN	Stop RP processing and terminate KILL request (req# 2). Close all databases, log and SAVE files.
OPENBASE	Open requested database, report errors (req# 2).

GETSTATS	Return status information for status request (req# 3). Return DBOPEN count, max user count, outstanding BEGIN's and status flags (STOP/BACKUP/SUSPENDED). Return user and database information for BEGIN's outstanding (currently calls QUIT 290 if response fails). Used by heartbeat also.
HEARTBEAT	Keep-alive message from <i>Shadow D/R</i> communications process (req# 5). Return status flags to caller.
RESUME	Startup recovery processing after a backup (req# 6). Clear Stop and Backup request flags. Re-open databases, log and SAVE files.
RECOVER	Not currently in use (req# 7).
BACKUP	Prepare for backup of databases (req# 8). Set Backup request flag and look for backup log record. When backup logrec received, find quiet period and close all databases, log and SAVE files.

Receiving the Replicate Log File

Log file records are received from the SP via the data pipe. They are an exact image of primary log file records. FWRITEDIR uses the log record number minus the base for the current log file to write the image to the exact same relative location in the log file as the primary log file.

There is sequence number checking on all log file records from the SP to ensure data is received in sequential order.

Keeping Track of Database Users

In order to keep track of database accessories and their activities, the RP keeps a data structure in the DL area of its stack. This data area contains information such as outstanding DBBEGIN's, user logging logid and position in the log file of the DBOPEN and last data records. The RP SAVE file contains this DL area data and is kept updated by Nobuf waited writes from the RP at periodic intervals.

Updating the Replicate Database(s)

Target databases obtained from the configuration file are put into a table. Log records for other databases or other non-TurboIMAGE data that are contained in the same log file are ignored.

All databases are opened at the start using MODE 1 or 4. Standard TurboIMAGE intrinsics are used for updating. A DBLOCK or DBUNLOCK is around each call at the set level if DBOPEN mode 1 is used.

TurboIMAGE E Intrinsic

Each TurboIMAGE intrinsic which modifies a database is executed with the user logging record information supplied. In addition, DBOPEN, DBCLOSE, DBBEGIN and DBEND log records cause RP table entries to be made, deleted or updated. Following, is a summary of the actions performed by the RP for each type:

DBOPEN	Look up the database in the list of Shadowed databases kept in an RP table and, if present, create a user entry for this access path. Store the DBOPEN log record number, the logid, the database index (to the RP database table) and the user's logical device number (meaningful only for a session). Increment the database user count.
DBCLOSE	If there is an entry in the user table for this logid, delete the entry and decrement the database user count.
DBEND	Check for logid in the user table. If an entry exists, set the BEGIN flag to FALSE, update the completion record of the transaction to the DBEND record number. Decrement the transaction (begin) counter. If the transaction counter is zero and STOP is pending, close the data message file from the slave CP.
DBUPDATE	Check for logid in user table. If an entry exists, apply the data to the correct database. Create a DBGET buffer at the top-of-stack and do a mode 7 (calculated) get for a master or a mode 4 (directed) get for a detail dataset. (Currently, call QUIT 214 on get error.) Compare the current data in the entry with the BEFORE data in the log record. (Currently, call QUIT 216 on FALSE compare.) Apply a DBUPDATE against the found record, using the AFTER data in the log record. (Currently, call QUIT 218 on update error.) If a transaction is in progress, update the current record processed field in the user entry, otherwise update the start of transaction field.
DBPUT	Check for logid in user table. If an entry exists, do a DBPUT using the item list and data provided. (Currently call QUIT 220 on put error.) The TurboIMAGE record number found in the user logging record and the TurboIMAGE record number returned in the status array are compared to ensure that the record went into the same location. (Currently, call QUIT 222 on not equal.) If a transaction is in progress update the current record processed field in the user entry, otherwise update the start of the transaction field.
DBDELETE	Check for logid in user table. If any entry exists, remove the entry from the correct database. Create a DBGET buffer at the top-of-stack and do a mode 7 (calculated) get for a master or a mode 4 (directed) get for a detail dataset. (Currently, call QUIT 224 on get error.) Compare the current data in the entry with the delete data in the log record. (Currently, call QUIT 226 on FALSE compare.) Apply a DBDELETE against the found record. (Currently call QUIT 228 on update error.) If a transaction is in progress, update the current record processed field in the user entry, otherwise update the start of the transaction field.

DBMEMO Check for logid in user table.

Call DBMEMO with data buffer supplied. (Currently call QUIT 230 on memo error.)

If a transaction is in progress, update the current record processed field in the user entry, otherwise update the start of the transaction field.

Each intrinsic above will decrement a SAVE file update counter. When the counter reaches 0, the updated entry will be posted to record 1 and the tables area of the SAVE file. More information about the SAVE file will be discussed later in this chapter.

Database Integrity Checks

A DBGET, mode 7 for a master or mode 4 for detail sets, will precede a DBDELETE/DBUPDATE to compare the contents log file "before" image with the current database contents. The "before" image must match before completing the DBDELETE/DBUPDATE. A QUIT 216 is called if the values are not equal. For DBPUT's, the TurboIMAGE record number is found in the user logging record. This record number is compared to the TurboIMAGE record number returned in the status array to ensure DBPUT will occur in the exact same location. A QUIT 222 is called if the numbers do not equal.

Maintaining System Restartability

The SAVE file serves the purpose of allowing a complete restart of the *Shadow D/R* program at the point at which processing left off. However, the conditions under which *Shadow D/R* is able to restart makes it necessary to stress that not all successful system restarts are correct. The Shadow processing environments on MPE XL systems (when not linked into Transaction Manager/XL) or MPE V system using disc caching allows the small probability for some updates to the database to not physically occur and yet have those updates accounted for during table recovery. In this situation, *Shadow D/R* will restart successfully even though the SAVE is not consistent with the database. Table recovery refers to the action of bringing the user and transactions tables up to date from the log record in the BEFORE checkpoint and the last log record recorded in the TABLE checkpoint. This is done because at the time of failure, the last update to the SAVE file may have occurred by as many as 1K log records ago.

Some examples of the conditions in which *Shadow D/R* would restart correctly are:

- Accidental RP job aborts by the operator.
- When the ROLLBACK feature is enabled (explained later in this chapter).
- Secondary system failures in which it can be assured that there was no RP activity to the secondary database at the time of the failure.

These will be known as "*Shadow D/R* automatic recover conditions."

If the secondary system fails while RP activity to the database is occurring and the rollback recovery feature is not enabled, then the appropriate recovery operation would be a secondary database restore.

Secondary System Recovery Enabled Feature

The purpose of enabling the secondary database for rollback recovery is to ensure physical I/O by TurboIMAGE to the databases.

At startup, **Shadow D/R** determines if TurboIMAGE rollback is enabled for any databases in the configuration file by reading the root file for that database. If **Shadow D/R** has determined that rollback is enabled, **Shadow D/R** will also ensure that physical I/O is performed for all writes to the SAVE file by using an FSETMODE 3. This method protects against secondary failures on both MPE V (with disc caching on in particular) and MPE XL systems. The secondary system must be recovered using a "warmstart" for MPE V systems and a "start recovery" for MPE XL systems. Enabling ROLLBACK on the secondary database will now guarantee logical integrity of secondary databases and synchronization with the recovery-SAVE file on both MPE XL and MPE V systems. With this feature, after recovering the secondary system from failure, you simply restart **Shadow D/R** and the software will resume replicating from the point of failure.

SAVE File Updates

Under certain conditions, the SAVE file serves the purpose of allowing a complete restart of the **Shadow D/R** RP program at the point at which processing left off. Two of the data areas in the SAVE file are related to the log file and one is related to the RP DL-area tables.

A much lower frequency of updates to the SAVE file by the recovery process has increased **Shadow D/R** system throughput starting with version A.00.07.

The RP posts the current WRITELOG record to the BEFORE check point record of the SAVE file (record 0) during log record processing as before. However, the second check point (previously known as the AFTER record or SAVE file record 1) write will be reduced to one every 1K WRITELOG records during processing.

There are other situations which require posting of the BEFORE and TABLE checkpoint records. During normal termination of log record processing (STOP, KILL and BACKUP commands) both the BEFORE and TABLE check point records are posted. Both check point records are also posted upon reading the next header record after a CHANGELOG record (log record type 13), as well as at the high water mark of RECOVER and SWITCH commands.

The DATABASE TABLE contains a list of each database from the configuration file and associates with it an index number and other information. Therefore, if **Shadow D/R** is restarted with a different configuration file which has the database listed in a different order, or added or deleted a database from the previous configuration, the RP will no longer abort with a configuration file error.

The LOG EVENT TABLE keeps track of logging related records such as SWITCHLOG's as they are processed by the **Shadow D/R** transport mechanism. The RP EVENTABLE does the same for RP related events such as STOP, START and RESTART. Every event recorded in each table in the SAVE file has entries for the log file record number that the RP last completed, and the date and time that log record was processed on the primary and secondary systems.

The STATUS TABLE includes a SAVE file version stamp so that the RP can determine at startup whether the SAVE file needs to be converted to the new format or not.

System Restart

This section describes how **Shadow D/R** is able to restart processing after a **Shadow D/R** automatic recover condition.

At system startup, the RP compares the BEFORE and TABLE checkpoint records. If the BEFORE record is older than the TABLE checkpoint record, the RP will terminate indicating a FATAL error and request a restore of the databases. If the TABLE record is older than the BEFORE record, the RP jumps back in the log file beginning at the sequence number indicated by the BEFORE record to recover the USER table.

The RP then posts the updated USER table to the SAVE FILE when the BEFORE record is reached. Then the BEFORE record is actively checked as outlined in the paragraph below. Finally, if the BEFORE and TABLE records are the same, only the BEFORE record is actively checked. The above three situations correspond to FATAL error conditions, abnormal RP termination, and normal termination of the RP.

During active checking of the BEFORE check point record, the RP attempts to verify that the last intrinsic. If the BEFORE checkpoint record is an UPDATE, PUT or DELETE IMAGE WRITELOG record corresponding to one of the Shadowed databases, the RP attempts a DBGET to obtain the current data and decide if the operation completed successfully. If the data is unrecognizable, the RP terminates indicating a FATAL error and the requests that the database be restored. All other passive (non-WRITELOG) BEFORE record types are not checked by the RP.

If the checkpoint records in the SAVE files are not synchronized (loss of synchronization is usually the result of a KILL operation, operator abort, system failure or other abnormal termination) the following message will display on the secondary console indicating the range of log file records used to update RP's user table.

RECOVER SAVE FILES TABLES from record mmm to nnn (RPMMSG 82)

Since *Shadow D/R* does not update the SAVE file from the RP's tables after each WRITELOG record, SAVE file tables recovery is required to synchronize the RP and the SAVE file tables to the point of the system interruption. *Shadow D/R* will then check that the last IMAGE operation completed successfully. When the RP restarts, the following message will appear on the secondary system's console and in the RP's job listing.

ACTIVELY CHECK last log record nnn (RPMMSG 83)

During table recovery and active checking of the BEFORE record, the RP's commands are restricted so the BACKUP, STOP and SWITCH commands are not allowed until normal log record processing is initiated.

Switch Execution and Transaction Rollback

Upon a SWITCH command with no options being received by the RP, the command file remains open only for commands compatible with its current mode. If the RP is suspended for BACKUP, a RESUME is executed to reopen all databases. Any log records in the message file from the SP are read and processed in the normal manner. When all the records from the SP are processed, if the operator issuing the SWITCH command indicated that there were one or more *Shadow D/R* tapes, the tape(s) will be read, ignoring any replicate log records and processed into the database(s) and user calls.

For each DBOPEN path that a user has (entry in the user table), the last complete transaction will be printed, complete with user name, program name, database name, dataset name and intrinsics called, with the entire data buffer.

For any outstanding incomplete transactions, the action of the intrinsics is reversed and the complete information for each intrinsic is listed as it is for complete transactions. Data is removed in the reverse chronological order of original application. What follows is the action performed for the three intrinsics involved in transaction rollout:

- | | |
|----------|---|
| DBUPDATE | Performed exactly as the original DBUPDATE, except the BEFORE data is used for the update and no data comparison is done for the existing old data. No user table maintenance or log or SAVE file updates are done. |
| DBDELETE | A DBPUT call is made with the data buffer from the delete log record, using "@" as the item list. Except for skipping the data comparison, this is identical to a DBPUT call in forward mode. |

DBPUT A DBDELETE call is made. Again, except for the data comparison, this call is identical to a DBDELETE call in forward mode.

After all rollouts have been completed, and the report printed, the log file and SAVE file are renamed by replacing the first character in the file name with a "z." This prevents inadvertent restarting of *Shadow D/R* using this log and SAVE file. The assumption is made that the database is no longer identical to the database(s) on the primary HP3000. A new log cycle with a new log file should be started on the secondary system when users begin to use it again.

There is a base level lock done before transaction rollback. Users must ensure there is no access to the database prior to a SWITCH operation.

12 Security

MODE=1 Security Solution

Shadow D/R expects the database(s) on the secondary system to be identical to the database(s) on the primary system at the beginning of a TurboIMAGE log cycle. After this point, *Shadow D/R* must be the only process that is allowed to update the database(s) on the secondary system. If another process is allowed to update them, the *Shadow D/R* process will detect the change and will terminate, requiring the database(s) to be brought back to an identical state before *Shadow D/R* can be restarted. In most cases, each database will have to be stored from the primary system and restored to the secondary system(s) which can take anywhere from a few minutes to a few days, depending on the size of your database(s).

The users on the secondary system should access the databases in a read mode to protect against updating the database on the secondary system. *Shadow D/R* currently opens the database(s) on the secondary system in TurboIMAGE mode 4, which does not allow concurrent write access, but does allow the concurrent read access using TurboIMAGE mode 6.

This implementation does not take into account that most of the database applications in use on HP3000's today open database(s) in either mode 1 or mode 5. There are some HP products in addition to some third-party products that require a TurboIMAGE open mode of 5. This can represent a significant problem if you require the use of these products on your secondary system(s).

A short term solution to this problem has been provided in *Shadow D/R*, which is to specify a "MODE=1:" statement in the *Shadow D/R* configuration file. Adding this statement to the configuration will cause *Shadow D/R* to open the database(s) on the secondary system with TurboIMAGE mode 1 rather than the default open mode of 4. Although this solution solves the problem of using products that require TurboIMAGE open modes 1 and 5, it also provides a way to update the database(s) on the secondary system with the problems described above.

There are two methods you can use to prevent users from updating the database(s) on the secondary system while maintaining the flexibility of having *Shadow D/R* open these database(s) with TurboIMAGE open mode 1. The first method deals with MPE file system security while the second deals with Image password security.

File System Security Solution

Write access to the database can be controlled by using the normal MPE file system security. A user with write access to the group would need to be created in the account where the database resides for the *Shadow D/R* process. All other users would need to be allowed read access only to the group and account where the database resides. Establishing MPE file system security like this would enable users to open the database in any compatible mode, but would prevent them from updating the database in the event that the *Shadow D/R* process stops accessing the database, such as when a BACKUP is requested.

The MPE file system security work around could require some additional planning in the event of a primary system failure. The file system security may need to be changed in order to make the secondary system a new primary system.

TurboIMAGE Password Security Solution

The second method of working around this mode 1/5 problem involves the use of TurboIMAGE password security. If the database has a password which only allows read access to the database and another password which allows write access to the entire database, the passwords could be manipulated to ensure that the password the database user used would only allow read access on the secondary system.

If the database on the primary system had password A assigned to class number 1, allowing read only access, and password B assigned to class number 2, allowing write access, could be changed to some new value which is only given to *Shadow D/R*. This would prevent anyone from modifying the database except the database creator and the *Shadow D/R* process. The password could then be changed back if the primary system failed and a switch was made to the secondary system. The passwords can be easily changed by the database creator using DBUTIL's SET command. Different variations of this work around may be used with more complex database security.

Set Subsystem=Real Solution

This solution will help to prevent QUERY subsystem access to the secondary database by any other users except the database creator. It does not prevent programmatic update by an application.

In situations where the secondary database modifications are suspected to emanate from QUERY updates, the DBUTIL SET command can allow only read access to the database by subsystems.

13 Troubleshooting

Overview

The complexity of the *Shadow D/R* processing environment makes solving most problems highly investigative in nature. Using the *Shadow D/R* recovery procedures to recover from a problem will not help determine the cause of a problem. Many *Shadow D/R* problems will stem from a cause which occurred much earlier, such as:

- System failures.
- Execution of non-TurboIMAGE tools without a new log cycle.
- Improper operations.
- Job stream aborts.

This chapter discusses the knowledge and tools required to isolate or resolve *Shadow D/R* problems.

Diagnostic Tools

Diagnostic Commands

TSTART Command	The TSTART command is available for complete tracking of the START command operation. The name, "TSTART," stands for "Trace START." The command sequence for the TSTART command is nearly the same as for the START command. <pre>>SYSTEM MCONFIGS >TSTART</pre> Please contact the Technical Support Team at Lund Performance Solutions if you want assistance with this command.
VERIFY Command	VERIFY will check syntax and completeness of the parameters in the configuration file. It does not check items such as incorrect node names or non-existent log files.
VERSION Command	<i>Shadow D/R</i> will not come up if the V.UU.FF levels are mismatch. VERSION will display this information for the local machine.

Diagnostic Messages

The messages described in this section are produced by *Shadow D/R* to inform the user about the state of the system components, during the execution of the **START** command and throughout *Shadow D/R* operation.

Each diagnostic message may be followed by one or more error messages (see "Error Messages" on page 69). If a diagnostic message is not followed by an error message, then it is merely informing the operator of the usual course of events, usually during the execution of the **START** command. These diagnostic messages are designed to inform the operator where events are occurring and which system processes are involved.

Requests from the MP to the CP

The following table lists and describes the requests from the MP to the CP:

Message	Description
CP/RP ident request	The MP is asking the RP or the CP to identify itself.
CP to receive config	The primary CP is about to receive the configuration file information from the MP.
CP to send config	The primary CP is requested to send the configuration file to the MP.
KILL request to CP	The primary CP received a KILL command from the MP.
STOP request to CP	The primary CP received a STOP command from the MP.
STATUS request to CP	The primary CP received a STATUS command from the MP.
BACKUP request to CP	The primary CP received a BACKUP command from the MP.
RESUME request to CP	The primary CP received a RESUME command from the MP.
USE request to CP	The primary CP received a USE command from the MP.
RELEASE request to CP	The primary CP received a RELEASE command from the MP.
VERSION request to CP	Request from the MP to the primary CP for the full program version of the CP and RP.

Requests from the MP to the RP

The following table lists and describes the requests from the MP to the RP:

Data Item	Description
Identify request to RP	The MP is asking the RP to identify itself.
Unused RP request	This message is reserved for later use.
RP open databases	The RP is requested to open databases on the secondary system.
RP shutdown (error)	The RP encountered errors when opening the databases on the secondary system and is requested to shut down.
RP STOP (unused)	This message is not used in the current version.
Request RP status	The RP received a STATUS request from the MP on the secondary system.
RP BACKUP (unused)	This message is not used in the current version.
RP RESUME (unused)	This message is not used in the current version.
SWITCH request to RP	The RP received a SWITCH request from the MP on the secondary system.
RP USE (unused)	This message is not used in the current version.
RP RELEASE (unused)	This message is not used in the current version.
Stand-alone RP recov	The RP received a stand-alone RECOVER command from the MP.
VERSION request, RP	The secondary CP requests the full program version from the RP.

Requests from the CP to the SP

The following table lists and describes the requests from the CP to the SP:

Data Item	Description
CP send cnfg to SP	The CP sent the configuration file information to the SP.
Open secondary msgfiles	The SP is requested to open the message pipe files to the RP on the secondary system.
Create RP job	The SP is requested to create the RP job on the secondary system.
Terminate remote	The CP on the primary system is about to KILL an SP associated with a DS line.
Terminate remote/RP	The CP is about to KILL an SP and the RP on the secondary system.

Data Item	Description
Secondary to open DB's	The secondary databases are about to be opened via a request to the RP.
Get BEGIN's from RP	The RP has been requested to send BEGIN records from the log file on the secondary system which belong to DBOPEN's for users with outstanding transactions.
Tell RP to STOP	The SP has been told to STOP by the CP on the primary system.
Tell RP to RESUME	The SP has been told by the CP on the primary system to cause the RP to RESUME processing.
Unused SP request	This message is not in use in this version.
Heartbeat to RP	The SP is requested to send a heartbeat to the RP on the secondary.
Tell RP to BACKUP	The SP has received a BACKUP request from the CP on the primary system and is requested to pass it on to the RP on the secondary.
VERSION request, SP	Request from the primary CP to secondary CP for full program version of CP and RP.

Requests from the SP to the RP

The following table lists and describes the requests from the SP to the RP:

Data Item	Description
RP GET SAVE logrec	The RP is about to get record 1 from the SAVE file on the secondary system to determine the log file restart record number.
RP terminate request	The RP has received a request to shut down.
RP to open Bases	The RP is requested to open the databases on the secondary system.
RP return status	The RP is requested to return STATUS information to the CP's on the primary and secondary systems.
RP orderly shutdown	The RP is about to start a controlled shutdown of the databases on the secondary system.
Heartbeat request to RP	The RP received a request to send a heartbeat to the CP's on the primary and secondary systems.
RP RESUME for backup	The RP received a request to RESUME after a BACKUP.
BACKUP request to RP	The RP has received a request to look for a backup log record and then wait for a quiet period before releasing the SAVE, log and data-base files.

CP States

The following table lists and describes the CP states:

Data Item	Description
REMOTE Hello	The CP is about to try a logon on the secondary system in order to create an SP.
Master POPEN	The primary CP is about to create the secondary CP.
CP Open log file	The CP is about to open the log file on the primary system.
CP OPENLOG intrinsic	The CP is about to use the OPENLOG intrinsic on the primary system.
Create tape process	The CP is about to create a tape process to handle records in excess of the taping threshold on the primary system.
CP files/lines open	The CP is reporting success in opening files and communications lines.
CP check tape	The CP is about to check the taping threshold to determine if it needs to do anything related to tape processing.
Send data block	Function code for request to send a block of data (none of the others).

Error Messages

This section provides a structured methodology for analyzing SHADOWRP errors to determine the cause of the abort. Familiarity and understanding of TurboIMAGE, QUERY, FCOPY, MPE User Logging, MPE Stack Dump, and the contents of this manual is required to successfully understand and complete this section.

(For a complete list of *Shadow D/R* messages, see Appendix B, "Messages" on page 88.)

Types of Shadow Errors

There are three main types of *Shadow D/R* errors:

- CP Errors This type of error comes from SHADOWCP, which is a job running on the primary system. Check the STDLIST of SHADOWCP job and see "Error Messages" in this manual for a list of diagnostic messages.
- MP Errors SHADOWMAP is the user interface to the *Shadow D/R* system. It is invoked by running SHADOWMP.PUB.SYS. See "Error Messages" for a list of diagnostic messages (self-explanatory).
- RP Errors SHADOWRP is the remote job running on the secondary system. The job name is also contained on the LOCAL= line in the configuration file. The STDLINS should contain the exact error message and a stack dump. This is usually the most difficult type of error to analyze.

TurboIMAGE Intrinsic Errors

TurboIMAGE intrinsic errors are not only the most reported *Shadow D/R* errors, but also the most demanding error to diagnose. They usually result from one of the following situations:

- Databases were not equal at the start of the logging cycle.
- Databases on the secondary system are modified by process other than the SHADOWRP.
- Improper recovery procedure after a system failure.
- Use of non-TurboIMAGE logging tools such as ADAGER, DBUTIL erase, etc.
- Log file corruption.
- Undetected hardware problem, e.g., bad memory board.
- Software (TurboIMAGE, MPE, *Shadow D/R*) bugs.

In order to pinpoint the cause, all of the following data must be present. However, because the error does not occur until the record is needed again, the data is usually long gone.

- STDLIST of aborted SHADOWRP.
- *Shadow D/R* CONFIG file from the primary system.
- SAVE file from the secondary system.
- Log files from BOTH systems.
- Copy of databases at the beginning of log cycle on primary (not essential, but required to determine if operational problems exist.)
- Copy of databases at the beginning of log cycle on the secondary system.

Frequently Asked Questions

Can a Shadow D/R environment be configured which Shadows separate databases in both directions? If yes, what special considerations are there for this type of environment?

Shadow D/R is fully supported for this type of environment. There are no special considerations in addition to the information presented in this manual. Recovery procedures are more com-

plex, however, because procedures must be followed for both primary and secondary failures for each machine.

Are there some primary failure recovery procedure considerations for users who are using Shadow D/R for only load balancing and never plan to implement a SWITCH?

For primary system failures, you have two options depending on the recovery which is most appropriate. Some people will use DBRECOV roll-forwarding recovery to recover the primary databases. If you do this, you should ensure that the database backup which is used for recovery was made when the primary and secondary databases were synchronized. This means that the backup was done on either system at the start of new log cycle. You must do a secondary store of the recovered database before starting *Shadow D/R* again.

Another option is to restore the last backup, copy the log file from secondary and use the SHADOWWP RECOVER with rollback option. Use the SAVE if the backup was made during the log cycle. The recover will process all the transactions that made it to the secondary and rollout any incomplete transactions. Do a RECOVER with rollback on the secondary as well, to synchronize the database and log files.

What can be done to flag people modifying the secondary database if the MODE=1 option is used?

There are a few suggestions contained in Appendix D of this manual for security provisions. Another method to isolate the culprit applications is to temporarily initiate logging on the secondary database. The DBRECOV statistics reporting function can be used to locate modify intrinsics from applications other than the RP. This is a tool which will reduce secondary performance and should only be used to troubleshoot this situation.

Why do we copy the log file over to the new primary when switching systems back to their original roles? They appear to be equal.

For recording purposes, the log file is copied to the new primary because of the CLOSELOG records added to the end when logging was stopped.

What extra considerations are there for a SWITCH operation if you are Shadowing to several different secondary system and the primary system fails?

Doing a SWITCH for this type of environment has ramifications regarding which system to make the new primary and synchronization of the remaining secondary systems with the primary. Synchronizing the remaining system with the new primary will in most situations require a complete database restore. Because of them, any ramifications of this configuration, you should immediately contact Lund Performance Solutions to help arrange optimal operations procedures to implement this configuration smoothly.

If I am Shadowing to several systems and the overflow mark has been reached for each secondary system, will there be a tape process for each configuration file? Should I have three tape drives?

Yes, there will be a tape process (TP) for each overflow mark that has been reached. However, since you are Shadowing the same log file to all other systems, only the transactions from one process are necessary. You can REPLY 0 to the others to stop the requests. You may need more than one tape drive if you plan to send these tapes to remote sites.

Why do we need MAXPRI=BS for the account in which the database(s) and log file resides?

In some *Shadow D/R* operations it is extremely important for the secondary to stay current with the primary. In these environments, the RP should be allowed a higher priority queue than user jobs and sessions. If your use of *Shadow D/R* is only load-balancing, you can remove the PRI=BS in the logon parameters in the configuration file. This means that the secondary has a much greater probability of falling behind the primary and the heartbeat mechanism may report missing beats because of the increased contention for CPU time.

What should be considered if I want to fit primary database backups into my operations instead of secondary backup?

Since you do not have the benefit of the SAVE file to mark the backup point in the log file, you need to ensure that primary and secondary databases are equal at the time of the backup. This procedure is outlined in the "Using Shadow for normal operations" chapter of this manual. It describes using the MP STATUS command to ensure that the secondary is current with the primary. When a recovery is needed, the SAVE file from the secondary, stored at the time of the primary backup, will still be needed with the database backup on the primary system.

What considerations are there for the use of TurboIMAGE ILR in conjunction with Shadow D/R systems?

If ILR is enabled on the secondary database, there is a performance penalty. However, you will benefit with ILR enabled by decreasing the probability of having to do secondary major failure recovery procedures. ILR on the primary depends on the local application design.

A Commands

Overview

In order to control the *Shadow D/R* environment, you must run the Manager Process (MP) SHADOWMP.PUB.SYS. Once this process is alive, it accepts commands from the user and performs the requested functions. This is the only interface between the user and the *Shadow D/R* system.

Each configured *Shadow D/R* system allows only one user at a time to run the MP for each configuration file. (More than one configuration file is used if you are Shadowing to more than one secondary system.)

Commands which act on the configuration file can be used with the specific name of the configuration file, or the SYSTEM parameter can be used to set the default configuration file for all commands. For example, you can enter either:

```
>SYSTEM config 1           or           >START config1
>START                     >START config2
>SYSTEM config 2
>START
```

and get the same results.

All commands except for SWITCH are valid on the primary system and only the SYSTEM, STATUS, SWITCH, VERIFY and VERSION commands are valid on the secondary system.

The following lists describe the different types of *Shadow D/R* commands.

Commands that Manage Configuration Files:

SYSTEM	See "SYSTEM Command" on page 85.
VERIFY	See "VERIFY Command" on page 86.

Commands that Start and Stop Shadow D/R

START	See "START Command" on page 80.
STOP	See "STOP Command" on page 83.
KILL	See "HELP Command" on page 75.
BACKUP	See "BACKUP Command" on page 73.
RESUME	See "RESUME Command" on page 79.
USE	See "USE Command" on page 85.
RELEASE	See "RELEASE Command" on page 78.
SWITCH	See "SWITCH Command" on page 84.
RECOVER	See "RECOVER Command" on page 77.

Commands to Request Status

CONFIG	See "BACKUP Command" on page 74.
STATUS	See "STATUS Command" on page 81.

Miscellaneous Commands

EXIT See "EXIT Command" on page 75.
HELP See "HELP Command" on page 75.
VERSION See "VERSION Command" on page 87.

UDC's

In order to allow UDC's, a command will be accepted using the INFO parameter of the RUN command. In this case, the single command is executed and the process is then terminated. For example:

START

OPTION LIST

RUN SHADOWMP.PUB.SYS;INFO="START CONFIG"

STOP

OPTION LIST

RUN SHADOWMP.PUB.SYS;INFO="STOP CONFIG"

Shadow D/R Commands

BACKUP Command

The BACKUP command suspends the RP and releases the database and files.

Syntax **>BACKUP [configfile]**

Parameter **configfile**

The configfile parameter specifies the name of the configuration file of the database to be backed up.

Function The BACKUP command notifies the *Shadow D/R* system that you wish to make a backup of the secondary database.

- Shadow will search for a quiet period where no multi-step transactions are in progress on the secondary database. During this search, a STATUS request will show that a backup is pending.
- When *Shadow D/R* finds a quiet period, the word "Now" appears on the next STATUS display to inform you that the databases have been released for storing.
- The *Shadow D/R* software on the secondary system will close the database, allowing the backup to take place. At this point, all databases and files specified in the SAVE and LOG lines in the active configuration file are stored to magnetic tape.
- Once the backup is complete, the RESUME command, issued on the primary system, will restart the processing on the secondary system.
- *Shadow D/R* automatically retrieves any queued transactions from the primary system, bringing the secondary system up to date.

Since the **Shadow D/R** recovery program requires the BACKUP request to be posted in the log file on the primary system, and the log record is the signal to the secondary system to look for a gap, a moment may pass before the command is acted upon. The STATUS command will show "BACKUP Pending" when the secondary system sees the request. When the secondary system finds a gap in the transaction stream, the STATUS display will read "Backup: Now" and the secondary database becomes available for backup.

Example

```

>BACKUP
Backup req. to RP: Backup pending on outstanding transactions. (RQMSG 38)
>STATUS
=====          Shadow          STATUS          =====
                                1:34:57

          *PRIMARY          SECONDARY
LOGFILE  record#    =  4675 (001)    LOGFILE  record#    =  4623 (001)
          time      =  1:34:54          time      =  1:30:36

          outstanding =  52 recs          FLAGS      Backup: Pending
          max ever    =  4675 recs

SENT      logical    =  4875 recs          DATABASE opens      =  3 paths
          physical   =  198 blocks          transact      =  2 outst
                                           Open mode     =  1 LOCK

          [current status]                [current status]

=====
LINES OPEN:SYSTEMB
=====

```

BACKUP Command

CONFIG Command

The CONFIG command prints the current configuration.

Syntax

>CONFIG

Function

The CONFIG command is used to get information about the current configuration. This command must be preceded by the SYSTEM command to establish the configuration file name.

The CONFIG command shows the information contained in the file. Unlike the VERIFY command, it does not check the file for validity.

NOTE: If you do not have SM, OP or PM, the passwords and DSLINE display as asterisks.

Commands

Example

```
>CONFIG

Shadow configuration came from: SHADCONF.GROUP.ACCOUNT

Local login:          LJOB shadjob,MG.account/****,group;pri=bs;hipri
Remote login:         HELLO shadsess,mgr.account/****,group;pri=bs;hipri

Commands file is:     p1.group.account
Responses file is:    p2.group.account
Log file name is:     logfl.group.account
Save file name is:    sv.group.account

Logging process is:   LOGPROC          Logging password is:
  With Databases:     DBTST1.GROUP.ACCOUNT/*****,
                     DBTST2.GROUP.ACCOUNT/*****,
                     DBTST3.GROUP.ACCOUNT/*****,

In LOCK mode

Wait time interval is: 10000
Max outstanding records: 100
DS lines specified:    ***
TAPEDEV specification: TAPE

This system was turned on: THUR, OCT 1, 1998, 2:40PM
>
```

CONFIG Command

EXIT Command

The EXIT command exits the MP and returns to MPE.

Syntax >EXIT

Function The EXIT command exits MP and returns to MPE. This does not affect the running of the *Shadow D/R* program that is maintaining the log files and databases.

HELP Command

The HELP command displays a one-line description of all commands.

Syntax >HELP

Function The HELP command displays information about each of the *Shadow D/R* commands.

Example

>HELP

BACKUP	Releases secondary system databases and system files for backup.
CONFIG	Displays the current configuration in effect.
EXIT	Terminates the Manager Program (but not Shadow)
HELP	Displays this list.
KILL	Causes an immediate shutdown of the Shadow system.
RECOVER	Initiates stand-alone database recovery from a log file.
RELEASE	Causes Shadow to release a DS line from active use.
RESUME	Resumes normal operation after doing a backup.
START	Initiates Silhouette processing between systems.
STATUS	Displays current system status information.
STOP	Causes an orderly shutdown of the Shadow system.
SWITCH	Initiates a switch from the primary to the secondary.
SYSTEM	Defines the current configuration file name.
USE	Requests that a DS line be placed in active service.
VERIFY	Validates the syntax of a configuration file.
VERSION	Displays the current software version and fix level.
:command	Issues the command following the colon to MPE.
>	

HELP Command

KILL Command

The KILL command terminates the *Shadow D/R* program immediately.

Syntax **>KILL [configfile]**

Parameters **configfile**

The configfile parameter specifies the configuration file you want *Shadow D/R* to terminate.

Function The KILL command operates on the supplied configuration file. If no configuration file is supplied, the current configuration file is terminated. The *Shadow D/R* program will terminate the replication process immediately without waiting to transfer any queued transactions from the primary to the secondary systems. The database on the secondary system, therefore, may contain partial transactions.

NOTE: A KILL or STOP may not take effect immediately if *Shadow D/R* exceeded the taping threshold. *Shadow D/R* may have a tape request pending for a tape device which is set for auto-reply and it will not terminate until a tape has been mounted.

Commands

Example

```
>SYSTEM siconfig
This Shadow system is currently active. (MPWARN 51)
>KILL
Terminate remote: Successful execution of function / request (RQMSG 0)
>
```

KILL Command

RECOVER Command

The RECOVER command applies transactions from a log file to the database.

Syntax

```
>RECOVER [configfile] [,optionword]
```

Parameters

configfile

The configfile parameter specifies the configuration file name. It must be used if the RECOVER command is issued via the INFO option of the RUN command in a UDC or job.

optionword

The optionword parameter specifies the RECOVER option. The options are as follows:

- | | |
|------------|--|
| report | The report option tells the <i>Shadow D/R</i> program to read the log file records and produce a normal recovery report without making any changes to the database. In other words, this option generates a report describing what the database would look like if a normal recovery is performed. |
| norollback | The norollback option tells the <i>Shadow D/R</i> program to perform a normal recovery without removing incomplete transactions from the database. A database recorded in this fashion may be internally inconsistent, but is guaranteed to contain every transaction applied to the database prior to the system failure. |
| norollfwd | The norollfwd option tells the <i>Shadow D/R</i> program to process the log file records without applying them to the database. Incomplete transactions will be removed as they are in a normal recovery. This option is used to clean up a log file which is known to contain incomplete transactions. |

Function

The RECOVER command is used to rebuild a database from a database backup.

In an interactive session, the *Shadow D/R* operator could give the configuration file name via the SYSTEM command. If the RECOVER command is issued with no optionword parameter, the *Shadow D/R* program will perform the equivalent of the SWITCH command including the normal SWITCH command report. The database will be rebuilt and all incomplete transactions will be removed. In order to insure recovery is correct using this method, either provide an appropriate SAVE file or make sure the log file contains only records **not** already applied to the database.

The RECOVER command checks the configuration file for the name of the SAVE file for the system. If the SAVE file exists, the RECOVER command will use the SAVE file record numbers to decide which records in the log file to process. If the SAVE file does not exist, the RECOVER command will process all records in the log file.

Example

```

V.UU.FF Shadow Manager Program
(C) COPYRIGHT 1998-1999 Lund Performance Solutions
>SYSTEM siconfig
>RECOVER
Do you wish to mount a Shadow transaction tape (y/n)?
N
      #J5
The recovery was started successfully for SICONFIG. (MPMSG 67)
>

```

RECOVER Command

RELEASE Command

The RELEASE command takes a communications line out of service.

Syntax **>RELEASE dsline**

Parameters **dsline**

The dsline parameter specifies the number or name of the DS line you wish to release.

Function The RELEASE command is used to tell the *Shadow D/R* program to stop using a particular DS line.

The *Shadow D/R* program completes any outstanding processing on the specified DS line and then stops using it. This command is helpful if spare lines or synchronous hardware usually used for *Shadow D/R* are periodically required for some other data communications sub-system.

Commands

Example

```
>RELEASE series42
Release req. to CP: There are no remaining DS lines available. (RQWARN 9)
>STATUS

===== Shadow STATUS =====
                        1:28:15

      *PRIMARY                                SECONDARY
-----
LOGFILE record# = 4675 (001) LOGFILE record# = 4675 (001)
      time      = 1:13:15      time      = 1:13:15

      outstanding = 0 recs      FLAGS      Shadow Mode
      max ever    = 4675 recs

SENT logical      = 4875 recs   DATABASE opens      = 0 paths
      physical    = 197 blocks  transact       = 0 outst
                                   Open mode      = 1 LOCK

      [current status]         [no DS line, old status]

-----
LINES OPEN:SYSTEMB
=====
```

RELEASE Command

RESUME Command

The RESUME command restarts the RP and opens the database and files.

Syntax >RESUME

Function The RESUME command is used to re-initiate Shadowing after a backup is completed. This command causes the *Shadow D/R* program on the secondary system to be reactivated and causes Shadowing to continue from the point where it left off when the BACKUP command was issued.

NOTE: Issuing the RESUME command while a STOP or BACKUP command is waiting for a quiet period will cancel the RESUME command.

Example

```
>BACKUP
Backup req. to CP: A backup request is already pending.
>RESUME
Tell RP to resume: Successful execution of function/request.
>STATUS

===== Shadow STATUS =====
                               1:34:57

*PRIMARY                                SECONDARY
-----
LOGFILE record# = 4675 (001) LOGFILE record# = 4675 (001)
      time      = 1:34:54      time      = 1:34:54

      outstanding = 0 recs      FLAGS      Shadow Mode
      max ever    = 4675 recs

SENT logical      = 4875 recs      DATABASE opens      = 0 paths
      physical    = 198 blocks      transact      = 0 outst
                                     Open mode      = 1 LOCK

      [current status]            [current status]
```

RESUME Command

START Command

The START command creates all components of *Shadow D/R*.

Syntax **>START [configfile]**

Parameters **configfile**

The configfile parameter will cause the *Shadow D/R* system to begin processing based on the specifications in the current or supplied configuration. If the START command follows the creation of a new log cycle, make sure the SAVE file is purged from the secondary system and backup the database on that system. (See “Starting a New User Logging Cycle” on page 26.)

Commands

Example

```
>START
Please be patient, this may take a minute or two to execute. (MPMSG 80)
#165
CP files/lines open: Number of DS lines opened = 1. (RQMSG 34)
The Shadow system was started successfully. (MPMSG 55)
>
```

START Command

STATUS Command

The STATUS command prints the status of *Shadow D/R*.

Syntax >STATUS [configfile]

Parameters configfile

The configfile parameter specifies the configuration file for which a status report is to be generated.

Function The STATUS command provides a method of inquiring about the current status of a *Shadow D/R* environment.

Example

```
>STATUS
===== Shadow STATUS =====
                               1:34:57

      *PRIMARY
-----
LOGFILE record# = 4675 (001)
        time    = 1:34:54
        outstanding = 52 recs
        max ever  = 4674 recs
SENT    logical  = 4875 recs
        physical = 198 blocks
        [current status]

      SECONDARY
-----
LOGFILE record# = 4623 (001)
        time    = 1:30:36
        FLAGS    Backup: Now
        DATABASE opens      = 3 paths
        transact  = 2 outst
        Open mode = 1 LOCK
        [current status]

=====
LINES OPEN:SYSTEMB
=====
```

STATUS Command

Note the following about the STATUS Command Example (above):

- The log file record numbers and the log record timestamp of the last record processed on both systems.

The **Shadow D/R** operator can use this information to determine how far the secondary system is behind the primary system. Log record timestamps for multiple record transactions are not reported until all records for the transactions are received on the secondary.

The primary log file record number specifies the last record sent to the SP and queued in the data pipe for RP processing. The secondary log file record number represents the last record processed to the database by the RP.

- The outstanding number of records in the primary system specifies the number of records which have been logged, but are not yet collected by the CP. These records are vulnerable to a system failure. The highest number of records outstanding since **Shadow D/R** was last started is displayed in "max ever."

If the primary system falls behind in its transmissions and the taping threshold is exceeded, or if the MAXRECS parameter in the configuration file is set to ALWAYS, an asterisk will display beside the entry for outstanding records. This asterisk means that a request has been sent to the system console.

- On the secondary side of the STATUS display, the FLAGS line shows the current state of the RP.

The possible FLAGS are:

Recovering Tables	This flag may be noticed during restart. It indicates the RP is updating its recovery tables from the data in the log file, or it is actively check the BEFORE check point record.
	This flag is incorrect if the current STATUS display shows that no RP is present (displayed in brackets on the lower right side of the screen). This means that the RP could not restart. The job listing should be checked for further information. The probable cause is that the active checking step uncovered a FATAL error.
Shadow Mode	This flag indicates the table recovery and active checking is complete and normal log record processing is under way.
Stand Alone Mode	This flag indicates the table recovery and active checking (done by RECOVER) is complete and normal log processing is under way.
Switch Mode	This flag shows the RP is operating in SWITCH mode. This information is available until the end of the data pipe is reached, but it is not available while the optional tape transactions are being processed.

- The SENT lines show the number of records transported since the last START. The number of blocks (24 records if the buffer is full) when compared to the number of records sent will help to determine the efficiency of the data transport mechanism.
- The DATABASE line shows the current database open mode and other information about open transactions and paths into the database(s). If the default open mode is used, "4 NOLOCK" means there are no locks on the database(s) and the open mode is four.
- LINES OPEN refers to the lines open on the primary system. This flag remains constant, even if you are looking at the STATUS from the secondary system. If no lines are open, there is a problem with the communication line.
- CURRENT STATUS indicates the system is active.

STOP Command

The STOP command terminates *Shadow D/R* when gracefully when all logical transactions are complete.

Syntax >STOP [configfile]

Parameters configfile

The configfile parameter specifies the configuration file for which *Shadow D/R* will stop processing.

Function The STOP command is used to stop Shadowing for the current configuration.

The *Shadow D/R* program will wait until a quiet period can be found before shutting the system down in an orderly fashion. Note the word "pending" in the example. Once the STOP is successfully completed, the *Shadow D/R* operator will not be able to get a STATUS display, because the configuration will no longer be active.

NOTE: Issuing a RESUME command while the STOP command is waiting for a quiet period will cancel the STOP command. (If *Shadow D/R* ever exceeded the taping threshold and has a tape request pending for a tape device, it will not terminate after a STOP or KILL command until a tape has been mounted on the device.)

Example

```
>STOP
Terminate the remote: Successful execution of function/request. (RQMSG 0)
>STATUS

===== Shadow STATUS =====
                               1:34:57

               *PRIMARY                SECONDARY
-----
LOGFILE record# = 4675 (001) LOGFILE record# = 4623 (001)
      time      = 1:34:54      time      = 1:30:36

      outstanding = 52 recs      FLAGS      Backup: Now
      max ever    = 4675 recs

SENT logical      = 4875 recs      DATABASE opens      = 3 paths
      physical    = 198 blocks      transact      = 2 outst
                                     Open mode      = 1 LOCK

      [current status]            [current status]

-----
LINES OPEN:SYSTEMB
=====
```

STOP Command

SWITCH Command

The SWITCH command changes the role of a secondary system to that of a primary system.

Syntax >SWITCH [,optionword]

Parameters optionword

The optionword parameter specifies the RECOVER option. The options are as follows:

- report The report option tells the *Shadow D/R* program to read the log file records and produce a normal recovery report without making any changes to the database. In other words, this option generates a report describing what the database would look like if a normal recovery is performed.
- norollback The norollback option tells the *Shadow D/R* program to perform a normal recovery without removing incomplete transactions from the database. A database recorded in this fashion may be internally inconsistent, but is guaranteed to contain every transaction applied to the database prior to the system failure.
- norollfwd The norollfwd option tells the *Shadow D/R* program to process the log file records without applying them to the database. Incomplete transactions will be removed as they are in a normal recovery. This option is used to clean up a log file which is known to contain incomplete transactions.

Function The SWITCH command is used to tell *Shadow D/R* to make the secondary system the new primary system.

If the SWITCH command is issued with no optionword parameter, any transactions already sent to the secondary system but not yet processed will be processed. The database will then be updated to a consistent state, with incomplete transactions removed from the database as required. A report will be printed showing the last complete transaction for each user of the database and any incomplete transactions the database may have contained before they were removed by *Shadow D/R*.

NOTE: Once the SWITCH command is issued, the databases are no longer synchronized and may not be used for Shadowing until they are copied to the new primary system.

Example

```
V.UU.FF Shadow Manager Program
(C) COPYRIGHT 1998-1999 Lund Performance Solutions
>SYSTEM siconfig
This Shadow system is currently active. (MPWARN 51)
>SWITCH
Do you wish to mount a Shadow transaction tape (Y/N)?
N
SWITCH req. to RP: Successful execution of function/request. (RQMSG 0)
>
```

SWITCH Command

Commands

SYSTEM Command

The SYSTEM command specifies a current configuration file.

Syntax >SYSTEM configfile

Parameters configfile

The configfile parameter specifies the configuration file to be the default file for all *Shadow D/R* commands.

Function The SYSTEM command is used to set the current configuration. This configuration will be the default for all subsequent commands until you either override it with a new specification or exit MP.

Example

```
>SYSTEM siconfig  
>
```

SYSTEM Command

USE Command

The USE command puts a communications line out of service.

Syntax >USE dsline

Parameters dsline

The dsline parameter specifies the DS line number or the name of the communications line that *Shadow D/R* will use.

Function The USE command instructs the *Shadow D/R* program to use the specified DS line for the current configuration (a DS line that was unavailable is now re-established). This command is also used to give back a line that was removed by the RELEASE command (the line must be specified in the configuration file).

Example

```

>USE
Please be patient. This may take a minute or two to execute. (MPMSG 80)
USE req. to CP: Successful execution of function/request. (RQMSG 0)
>STATUS

===== Shadow STATUS =====
                        1:34:57

*PRIMARY
LOGFILE record# = 4675 (001)
        time    = 1:34:54
        outstanding = 0 recs
        max ever   = 4675 recs
SENT    logical  = 4875 recs
        physical = 198 blocks
        [current status]

SECONDARY
LOGFILE record# = 4675 (001)
        time    = 1:34:54
        FLAGS   Shadow Mode
        DATABASE opens      = 0 paths
        transact = 0 outst
        Open mode = 1 LOCK
        [current status]

=====
LINES OPEN:SERIRES42
=====

```

USE Command

VERIFY Command

The VERIFY command validates the statements in a configuration file without starting a *Shadow D/R* process.

Syntax >VERIFY configfile

Parameters configfile

The configfile parameter specifies the configuration file to verify.

Function The VERIFY command provides a method of validating a *Shadow D/R* configuration file. It takes longer than a CONFIG command to execute, so use the CONFIG command to get a quick reminder of the configuration file contents. Use the VERIFY command to see how *Shadow D/R* will interpret a configuration file, without actually activating the configuration.

The user must have SM, OP or PM to see the file contents. If not, the MPMSG 81 is displayed.

Commands

Example

```
:RUN shadowmp.pub.sys

V.UU.FF Shadow Manager Program
(C) COPYRIGHT 1998-1999 Lund Performance Solutions
>system config.shadowtest.sys
    <<Shadow Configuration File>>

BEGIN
    LOCAL=!JOB shadjob,MGR.account/suit,group;pri=bs;hipri:
    COMMANDS      = p1.group.account:          <<Pipe to CP>>
    RESPONSES     = p2.group.account:          <<Pipe from CP>>
    SAVEFILE      = sv.group.account:          <<Saves environment>>
    USE logproc WITH dbtst1.group.account/MGR, <<all three same>>
                  dbtst2.group.account/MGR,
                  dbtst3.group.account/MGR:
    WAIT          = 1000:                      <<ten seconds>>
    MAXRECS       = 0100:                      <<before requesting tape>>
    DSLINE        = series 42:                 <<1 DS line>>
    MODE          = 1:
END. <<of configuration>>
No errors found in Shadow configuration file. (MPMSG 81)
>
```

VERIFY Command

VERSION Command

The VERSION command displays the current software version.

Syntax

>VERSION

Function

The VERSION command provides a method of asking *Shadow D/R* to report on the version of the active software. This command is useful to verify software revision numbers while contacting the Technical Support Team at Lund Performance Solutions.

Example

In the following example, "V.UU.FF" is used to represent the correct version of your software.

```
>VERSION
ShadowMP
V.UU.FF
>
```

VERSION Command

B Messages

Configuration Messages

The following messages result primarily from errors in syntax:

CONFIGERR 0	Commandname command expects a file name.
CONFIGERR 1	Configuration file READ error.
CONFIGERR 2	Unknown keyword or symbol.
CONFIGERR 3	Found more than one BEGIN keyword in configuration file.
CONFIGERR 4	Missing BEGIN keyword in configuration.
CONFIGERR 5	Missing "=" after keyword.
CONFIGERR 6	Expected colon ":" as a statement terminator.
CONFIGERR 7	Expected a "!" to precede JOB in jobcard.
CONFIGERR 8	Expected "!JOB" to start Shadow jobcard.
CONFIGERR 9	Missing COMMANDS file name.
CONFIGERR 10	Missing RESPONSE file name.
CONFIGERR 11	Expected HELLO to start remote logon.
CONFIGERR 12	Missing time for WAIT parameter.
CONFIGERR 13	Expected wait time between 1000 and 32000 msecs.
CONFIGERR 14	Missing number for MAXRECS parameter.
CONFIGERR 15	Expected MAXRECS 0,1-999, NEVER or ALWAYS.

Messages

CONFIGERR 16	Missing device class or ldev for TAPEDEV parameter.
CONFIGERR 17	Invalid ldev or class specified for TAPEDEV parameter.
CONFIGERR 18	Missing logging identifier following USE keyword.
CONFIGERR 19	Invalid logging identifier specified.
CONFIGERR 20	Invalid logging identifier password.
CONFIGERR 21	Expected database name following WITH keyword.
CONFIGERR 22	Database name invalid or not fully qualified.
CONFIGERR 23	Expected a password following database name.
CONFIGERR 24	Invalid database password specified.
CONFIGERR 25	Expected a comma "," separating database names.
CONFIGERR 26	Expected a "/" to precede the password.
CONFIGERR 27	Missing WITH for USE keyword.
CONFIGERR 28	Too many databases specified. Maximum is 12.
CONFIGERR 29	Missing "END." in configuration file.
CONFIGERR 30	Too many DS lines were specified. Maximum is 4.
CONFIGERR 31	Expected at least one DS line name or ldev.
CONFIGERR 32	Invalid DS line name or ldev specified.
CONFIGERR 33	Missing LOCAL keyword in configuration file.
CONFIGERR 34	Missing COMMANDS keyword in configuration file.

CONFIGERR 35	Missing RESPONSES keyword in configuration file.
CONFIGERR 36	Missing USE keyword in configuration file.
CONFIGERR 37	Missing WAIT keyword in configuration file.
CONFIGERR 38	Missing MAXRECS keyword in configuration file.
CONFIGERR 39	Missing DSLINE keyword in configuration file.
CONFIGERR 40	Missing REMOTE keyword in configuration file.
CONFIGERR 41	Missing TAPEDEV keyword in configuration file.
CONFIGERR 47	Missing log file name in configuration file.
CONFIGERR 48	Invalid log file name specified in configuration file.
CONFIGERR 49	Missing LOG keyword in configuration file.
CONFIGERR 56	Missing SAVEFILE name in configuration file.
CONFIGERR 57	Mising SAVEFILE keyword in configuration file.
CONFIGERR 58	Duplicate configuration statement found, not allowed.
CONFIGERR 59	Logon image must be 100 characters or less.
CONFIGERR 60	Bad Recovery Program DBOPEN mode. Only 0, 1 or LOCK and 4 or NOLOCK are valid.

Manager Process (MP) Messages

The following messages are associated with the Manager Process (MP):

MPERR 4	Unknown Shadow command. <i>Shadow D/R</i> did not recognize the command specified. Check the spelling of the command and enter it again.
---------	---

MPERR 5	The commandname command has not yet been implemented. The command you chose does not work in this version of <i>Shadow D/R</i>. Use a valid <i>Shadow D/R</i> command (see “Shadow D/R Commands” on page 73).
MPERR 6	MPE command error error,param. This message shows the MPE COMMAND intrinsic returns "error" and "param." Call the system manager.
MPERR 7	The commandname command expects a file name. Specify a file name with the command.
MPERR 8	The commandname command expects a DS line number or class. Specify a DS line number or class following the command.
MPERR 10	File System Error nnn opening configuration file filename. An MPE file system error was encountered in the FOPEN of the specified configuration file. Call the system manager.
MPERR 11	File System Error nnn opening RESPONSES file. An MPE file system error was encountered in the FOPEN of the specified configuration file. Call the system manager.
MPEWARN 39	There are no DS lines available. All DS lines are unavailable. <i>Shadow D/R</i> will continue trying until a DS line becomes available.
MPERR 40	DS line dsline cannot be found in the configuration file. Specify a different DS line or modify the DS line parameter in the configuration file.
MPWARN 42	No configuration is currently active. Enter START and specify the configuration file.
MPWARN 43	This configuration is NOT currently active. Enter START and specify the configuration file.
MPWARN 44	The CP is present but does not respond. The CP is busy elsewhere and did not respond. This situation normally corrects itself.

MPWARN 45	This system is currently being controlled by another Shadow Manager Program. You are not allowed to run more than one MP per configuration. Wait until the other user exits the MP and try again.
MPERR 46	File System Error opening Commands file to the CP job. An MPE file system error occurred in the FOPEN of the COMMANDS file. Call the system manager.
MPERR 50	The CP or RP process does not respond or is not present. This error is due to a DS line error, a connect timeout, or because the CP job could not start. Call the Shadow D/R manager or system manager.
MPWARN 51	This Shadow system is currently active. The MP referred to a configuration file that was already started. You cannot specify an active configuration.
MPWARN 52	This Shadow system was started with a different configuration file. Use the CONFIG command. The MP is attempting to control the system with a configuration file other than the one used to start it. Use the configuration file that was used to start the system.
MPERR 53	The CP or RP did not supply the expected response to the Manager Program. Try again. A previous error may have caused an abort, leaving corrupt messages in the Commands or Responses pipe files.
MPERR 54	The specified system was NOT started. Check STDLIST or the CP job for any errors and try again. This error is displayed with many errors on the terminal screen. Check the \$STDLIST for any errors which may not have appeared on the terminal screen.
MPMSG 55	The Shadow system was started successfully. This message is displayed after a successful START has occurred. No action is necessary.
MPERR 58	The commandname command is invalid on an active Shadow recovery system. The specified command is not permitted on the secondary system. Determine if the command used is proper for the current mode.
MPWARN 59	This system is NOT suspended for BACKUP. Any pending requests for either STOP or BACKUP were cleared. This message displays after a successful RESUME, when the RP has not yet released the databases for backup. No action is necessary.

MPERR 60	The commandname command is invalid on the main system. The specified command is not permitted on the primary system. Refer to “Shadow D/R Commands” on page 73 for a list of permitted commands.
MPERR 62	Cannot RECOVER an active Shadow system. If databases are active, they do not need to be recovered. Do not attempt to recover an active database.
MPERR 63	The recovery job was not started. This message may display because the RP could not be started, the databases could not be opened, or the log file does not exist. If the RP was started, examine the job stream for information.
MPERR 64	There is no answer from the recovery job. The RP started but did not complete, possibly opening pipe files. Call the system manager.
MPERR 65	There is no answer from the recovery job. All the databases specified in the configuration file could not be opened. Check the RP job \$STDLIST to determine the problem.
MPERR 66	The recovery job could not continue. The RP aborted because of an error. Check the job stream for information.
MPMSG 67	The recovery was started successfully for filename. This message displays when a stand-alone recovery has been started successfully. It identifies the configuration file name. If the file name is correct, no action is necessary.
MPERR 68	Bad option specified. The option supplied with the command is invalid, incorrect or misspelled.
MPERR 69	The recovery job was not scheduled. The STREAM command for the recovery job failed. Verify the STREAM command is correct and try again.
MPWARN 70	(filename) is not enabled for logging. <i>Shadow D/R</i> will not work unless the database is enabled for logging. Enable it.
MPWARN 71	(filename) is not using the correct logid. The logid for the configuration file and root do not match. They must match.

MPERR 73	File System Error opening root file for filename. An MPE file system error occurred in the FOPEN for the root file of the specified database. Call the system manager.
MPERR 74	File System Error nnn reading root file for filename. An MPE file system error occurred in an FREAD of the root file for the specified database. Call the system manager.
MPERR 75	Stream line open failure# FSERR#nnn. An MPE error occurred in FOPEN for the job stream file on the primary system. Call the system manager.
MPERR 76	Stream of CP job failed. Error = nnn. An MPE error occurred in the COMMAND intrinsic to steam the CP job. Call the system manager.
MPERR 77	Filename is enabled for Shadow recovery. This database may not be used as a primary system. Not currently in use. The user is trying to use a secondary database as a primary database without entering a START command first.
MPWARN 78	This Shadow system is stopping or something has stopped. This message is generated in response to any command entered after a STOP command. To continue, enter a START command.
MPERR 79	There is not a recovery process on the secondary system. Examine the error message(s) above, or check the job listing on the secondary printer for information. The RP has logged off due to abnormal conditions. Check the job listing on the secondary printer for additional information.
MPERR 80	Please be patient, this may take a minute or two to execute. This message informs the user that a wait of a few minutes may be normal for both the START and USE commands. Wait for the command to execute, then continue.
MPERR 81	No errors found in the Shadow configuration file. This message is generated in response to a VERIFY command when there are no errors. No action is required.
MPERR 85	The Shadow program versions are mismatched. n.nn.nn versus n.nn.nn. Some combination of the CP, RP or MP has unmatched program version. The first version given is for the MP, then either the RP or CP. Ensure that the program version matches for all three components.

MPERR 86	Cannot start Shadow subsystem. Verify installation. <i>Shadow D/R</i> is not properly installed on the system. Install the product again to ensure all versions match.
MPWARN 87	Warning from the Stream command starting CP. CIWARN = nnn. A non-fatal error in LOCAL definition in configuration file. Examine CI errors and the configuration file to determine the cause of the problem.
MPWARN 88	No configuration has been specified. Use of the STATUS command before specifying a configuration file name with either the CONFIG, SYSTEM or STATUS commands will result in this warning. Specify a configuration file name with any of these commands.

Communications Process (CP) Messages

The following messages are associated with the Communications Process (CP):

CPERR 1	PTOP error on internal line (line), DSLINE (dsline), File System or DS Error = nnn. This is a failure in the program-to-program communication intrinsic. Call the system manager.
CPERR 2	Associated DS secondary terminated. After a PTOP error, a PCLOSE and DS line close were issued. Check the communications link to the secondary system.
CPERR 11	Error writing to RESPONSES file. n MPE file system error occurred in an FWRITE to the RESPONSES file to the MP. Call the system manager.
CPWARN 12	Shadow terminated by STOP/KILL. This warning is generated after a successful STOP or KILL. No action is required.
CPERR 13	Control request rejected by remote. Request = nnn. A control request to the secondary CP was rejected. Examine the \$STDLIST for information.
CPERR 14	No Recovery Process on remote system after a secondary system failure or an aborted RP job. Wait for the RP to be created again. Issue the STATUS command to see if the secondary flag is "Shadow Mode."

CPERR 15	No response received from the remote recovery process. A response was expected from the RP, but was not received before the timeout occurred. Check that the secondary <i>Shadow D/R</i> job is running.
CPWARN 16	Stopping... This message is generated when the STOP command is in progress and waiting for a quiet period. No action is required.
CPERR 17	Couldn't create the remote Recovery Process. The job stream of the RP job failed.
CPERR 18	Remote command already in progress. Internal error. The command pipe file already contained a command. Investigate the secondary for reasons the RP is impeded. The system may just be slow temporarily.
CPERR 19	WRITELOG error=nnn. An MPE error occurred while writing to the log file on the secondary system. Call the system manager.
CPERR 20	CREATE error=nnn. An MPE error occurred in the FCREATE of the TP. Call the system manager.
CPERR 21	File System Error reading the tape response. The system cannot communicate with the TP. Call the system manager.
CPERR 22	Tape Process did not respond OK. The system cannot communicate with the TP. Call the system manager.
CPERR 23	POPEN on line (line) failed. Error=nnn. An MPE error occurred in the PTOPI communications system. Call the system manager.
CPERR 24	Error sending configuration to secondary. This error is probably due to a DS line problem. Call the system manager.
CPERR 25	Error opening secondary pipes. An MPE file system error occurred in the FOPEN of the pipe files on the secondary system. Call the system manager.

CPERR 26	Secondary couldn't open the pipes. An MPE file system error occurred in the FOPEN of the pipe files on the secondary system. Call the system manager.
CPERR 27	Secondary couldn't create SRP. The secondary system could not stream the internal job to run RP. Check the jobfence and limits.
CPERR 28	SRP couldn't open the database. The RP on the secondary system could not open the databases. Check the databases access and security.
CPERR 29	Logfile open failed. FSERRnnn. An MPE file system error occurred in the FOPEN of the log file on the secondary system. Call the system manager.
CPERR 30	OPENLOG failure nnn. An MPE error occurred involving the OPENLOG intrinsic. Call the system manager.
CPERR 31	Error creating Tape Process nnn. An MPE error occurred involving the CREATE intrinsic. Call the system manager.
CPERR 32	Couldn't open any DSLINE or slave error. This message is generated as a result of a DS line error or a remote secondary error. Call the system manager.
CPERR 35	Couldn't request tape start. The message file write failed to TP. Call the system manager.
CPERR 36	Couldn't request a tape stop. The message file write failed to TP. Call the system manager.
CPERR 37	Error reading start record from SRP. An MPE file system error occurred in the FREAD of the SAVE file. Call the system manager.
CPERR 39	...creating a remote Recovery Process. This message displays when restarting after a secondary system failure or abort RP job. No action is required.

CPERR 41	TCOP open error. An MPE file system error occurred in the FOPEN of the TCOP file in the TP. Call the system manager.
CPERR 42	FROMCP open error. An MPE file system error occurred in the FOPEN of the FROMCP file in the TP. Call the system manager.
CPERR 43	SP: File System or DS error nnn occurred on GET. There is a secondary GET error in the PTOP system. Call the system manager.
CPERR 44	Error reading logfile. An MPE error occurred when the CP tried to FREAD the log file. Call the system manager.
TPERR 44	Error reading logfile. An MPE error occurred when the tape handling process tried to FREAD the log file. Call the system manager.
CPERR 46	Dsline (dsline) open error=nnn returned by DS. The line number in error is a 0-relative index into list of lines. See your DS manual for the specific error.
CPMSG 47	End of Shadow. This message indicates the end of <i>Shadow D/R</i>. No action is required.
CPERR 48	Only U-Mit Logging allows a logfile sequence number. Sequence number nnn returned by SHADOWRP. This message only displays when the user backdates MPE on the primary system to a version prior to U-MIT. Do not use AUTO CHANGELOG feature.
CPERR 49	LOGINFO intrinsic call failed while opening logfile. User logging error nnn by intrinsic. The log file could not be opened. Check the error provided in the message.
CPERR 50	Attempt to START with an out-of-sequence logfile. Secondary logfile sequence number is nnn. A new log cycle was started with an old SAVE file. Purge the old SAVE file and restart.
CPERR 51	CHANGELOG enabled and logfile doesn't end in 3 digits. The log file name is incompatible with the U-MIT CHANGELOG requirements. Rename the file to be compatible with U-MIT CHANGELOG requirements.

Messages

CPERR 52	START with non-identical/unavailable secondary logfile. The secondary log file was either purged or changed while <i>Shadow D/R</i> was not running. Check to make sure the specified log file exists on the secondary system.
CPERR 53	Tape Process OPENLOG failure. Logging error=nnn. The internal TP encountered an error opening the log process. Check the error number to determine the problem.
CPERR 54	Tape Process Logfile open failed. Cannot write to tape. The TP couldn't open the log file. Check the \$STDLIST for SHADOWCP and correct the problem.
CPWARN 55	Warning from the REMOTE HELLO command. CIWARN=nnn. A non-fatal error occurred involving the REMOTE definition in the configuration file. Examine CI errors and the configuration file to determine the cause.
CPWARN 56	Warning from the STREAM command starting RP. CIWARN=nnn. A non-fatal error occurred involving the LOCAL definition in the configuration file. Examine CI errors and the configuration file to determine the cause.
TPERR 57	Tape process failed to open tape device. The tape device on the primary system is logically down. Investigate the status of the tape drives.

Recovery Process (RP) Messages

The following messages are associated with the Recovery Process (RP):

RPERR 1	Error writing to SAVE file. An MPE file system error occurred in the FWRITE to the SAVE file on the secondary system. Call the system manager.
RPERR 2	Couldn't open the input pipe file. An MPE file system error occurred in the FOPEN of the input file on the secondary system. Call the system manager.
RPERR 3	Error setting the extended wait. Internal file system error. Call the system manager.
RPERR 4	Couldn't open the command pipe file. An MPE file system error occurred in the FOPEN of the command pipe file on the secondary system. Call the system manager.

RPERR 5	Error setting command wait. Internal file system error. Call the system manager.
RPERR 6	Error setting command timeout. Internal file system error. Call the system manager.
RPERR 7	Couldn't open output pipe file. An MPE file system error occurred in the FOPEN of the output pipe file on the secondary system. Call the system manager.
RPERR 8	Couldn't open the SAVE file. An MPE file system error occurred in the FOPEN of the SAVE file on the secondary system. Cannot restart the <i>Shadow D/R</i> subsystem. Call the system manager.
RPERR 9	Couldn't close the SAVE file. An MPE file system error occurred in the FCLOSE of the SAVE file on the secondary system. Call the system manager.
RPERR 10	Error writing to the SAVE file. An MPE file system error occurred in an FWRITE to the SAVE file on the secondary system. Call the system manager.
RPERR 11	Error reading the SAVE file. An MPE file system error occurred in an FREAD to the SAVE file on the secondary system. Call the system manager.
RPERR 12	Checkpoint (SAVE) file indicated potentially bad database. The first work records in the SAVE file are different, indicating a DB call failure or critical system failure during a DB update. Use the recovery procedure for a major failure on the secondary system.
RPERR 13	Restore Shadow databases, checkpoint (SAVE) file and logfiles. A failure occurred on the secondary system and <i>Shadow D/R</i> requires files to be restored. The database could be corrupt. Restore the databases, SAVE file and log files.
RPERR 14	Error opening the logfile. The RP could not open the log file on the secondary system. Call the system manager.

RPERR 15	Error closing the logfile. An MPE file system error occurred in an FCLOSE to the log file on the secondary system. Call the system manager.
RPERR 16	Error writing to the logfile. An MPE file system error occurred in an FWRITE to the log file on the secondary system. Call the system manager.
RPERR 17	Error reading the logfile. An MPE file system error occurred in an FREAD of the log file on the secondary system. Call the system manager.
RPERR 18	Error issuing backup tape equate. An MPE COMMAND intrinsic error occurred issuing an equate for the backup tape on the secondary system. Call the system manager.
RPWARN 22	These transactions have been removed from the Shadow database. This system must now be used as the MAIN system. The databases are no longer identical. You must ensure that the Databases are identical before starting the Shadow system again. This warning displays after a successful SWITCH on the secondary system. No action is required.
RPMSG 23	There were no incomplete Shadow transactions. An MPE file system error occurred in an FREAD of the log file on the secondary system. Call the system manager.
RPERR 24	The DBUPDATE data doesn't compare. When processing a DB update request from the primary system, the "before change" data was not the same. Check for secondary database modify occurrence. Restore database(s), SAVE file and log file.
RPERR 25	DBPUT record number error. Old=nnn New=nnn. Detail data set record numbers do not match. Databases are not the same. Check for secondary database modify occurrence. Restore database(s), SAVE file and log file.
RPERR 28	Error writing to secondary system. The RP failed in writing back to the CP via the message pipe on the secondary system. Call the system manager.
RPMSG 29	RESTART requested at logfile record nnn. Appears at the console in response to a START or USE command when RP already existed. No action is required.

RPMSG 30	***Shadow requested shutdown.*** This message appears after a KILL or STOP command when there is no outstanding BEGIN. No action is required.
RPMSG 31	Shadow recovery process shutdown. This message appears after a KILL or STOP command when there is no outstanding BEGIN. No action is required.
RPMSG 32	Database opened filename. This message is issued after a successful DB open requested by a MP START or RESUME command. No action is required.
RPERR 33	Unable to re-open SAVE file or logfile. The RESUME or START command is not able to open the SAVE or log file via RP. Call the system manager.
RPMSG 34	RESUME request executed. This message appears after a RESUME command is entered. No action is required.
RPMSG 35	Starting stand-alone recover. This message appears after a RECOVER command. No action is required.
RPERR 36	(No rollout of incomplete transactions.) This message appears after a successful RECOVER command with the NOROLL-OUT or REPORT options. No actions is required.
RPERR 37	(No new transactions added.) This message appears after a successful RECOVER command with the NOROLL-OUT or REPORT options. No actions is required.
RPERR 38	Starting recover. This message appears after a RECOVER command from the MP. No action is required.
RPERR 39	Unable to open SWITCH log nnn. An MPE file system error occurred in the FOPEN of the SAVE file on the secondary system. Call the system manager.
RPERR 40	Unable to open SWITCH SAVE file filename. An MPE file system error occurred in the FOPEN of the SAVE file on the secondary system. Call the system manager.

RPERR 41	Unable to read SAVE file label. An MPE file system error occurred in the FREADLABEL of the SAVE file. Call the system manager.
RPERR 42	SAVE file does not correspond to a backup quiet point. This message appears after the FREADLABEL, indicating a backup, though there are still outstanding transactions. Call the system manger.
RPERR 43	Recovery log record number nnn. This message provides the log file record number in which the error was found. No action is required.
RPERR 44	Finished recover. This message appears after a successful RECOVER. No action is required.
RPERR 45	FREDADDIR error on log file. An MPE file system error occurred in an FREDADDIR of the log file on the secondary system. Call the system manager.
RPERR 46	Tape operation failed. FSERR nnn. An MPE file system error occurred in reading the transaction tape during a SWITCH or RECOVER. Call the system manager.
RPERR 47	Record out of sequence. Expected nnn. Got nnn. This message appears when a log record is received out of sequence. Call the system manager.
RPERR 48	Tape is Shadow nnn. This message provides the tape sequence number. No action is required.
RPERR 49	Wrong tape set? Missing records. This message appears when an incorrect tape is mounted during a SWITCH or RECOVER. Mount the correct tape and try again.
RPWARN 50	End of Shadow tape records. This warning displays when there is no trailer on the Shadow tape. This is a normal occurrence when the primary system was failed, no action is required.
RPMSG 51	Shadow tapes processed OK. This message appears after a SWITCH or RECOVER is finished. No action is required.

RPERR 52	Cannot purge filename, rename did not take place. An MPE COMMAND intrinsic error occurred in renaming the log and SAVE files after a SWITCH or RECOVER. Call the system manager.
RPERR 53	***Manager requested a SWITCH.*** This message displays after a SWITCH request on the secondary system. No action is required.
RPERR 54	End of stand-alone recovery. This message appears after a successful stand-alone RECOVER. No action is required.
RPERR 55	Command file FSERR nnn. An MPE file system error occurred in reading a command from the CP on the secondary system. Call the system manager.
RPERR 56	Missed a heartbeat. This message appears after a missed heartbeat from the primary system is detected. Check the primary system status.
RPERR 57	Recovery process no longer communicating to main system. This message is generated after two consecutive missed heartbeats from the primary system are detected. Check the primary system status.
RPMSG 59	End of stand-alone recovery. (Now rollout of incomplete transactions.) This message displays after a RECOVER with the NOROLLFWD or REPORT options. No action is required.
RPMSG 60	End of stand-alone recovery. (No new transactions added.) This message displays after a RECOVER with the NOROLLFWD or REPORT options. No action is required.
RPMSG 69	Not Shadow log tape. Mount another? (Y/N) This is not recognized as a Shadow log tape. Determine whether or not this is the correct tape.
RPERR 70	Database not opened filename. The RP could not open the database. Check the \$STDLIST to determine the TurboIMAGE error.

RPERR 71	The SAVE file on the secondary is out-dated. <i>Shadow D/R</i> is using an old SAVE file with a new version of <i>Shadow D/R</i>. Start <i>Shadow D/R</i> with a new log cycle.
RPERR 72	START request at logfile record nnn. This message appears at the console in response to a START or USE command. Similar to RPERR 29, except this is a START with no RP on the secondary. No action is required.
RPERR 73	Bad log file sequence number encountered: nnn. The user started a new log cycle incorrectly or changed the log file name and caused the RP log file name to be out of sequence. Synchronize the primary and secondary log file sequence numbers.
RPMSG 74	Logfile change from logid to logid. This message is sent to the console to show a change in log files (U-MIT CHANGELOG console message). No action is required.
RPMSG 75	Logfile logid opened. The specified log file has been opened. This is usually the result of the CHANGELOG. No action is required.
RPMSG 76	Failed to open logfile logid. The specified log file could not be opened (console message). No action is required.
RPMSG 77	SWITCH logfile logid not present. Available? This is an operator message to supply a log file for a SWITCH request. This occurs when a previous log file is required for rollback. No action is required.
RPERR 78	SWITCH operator request for new logfile failed. This message is located on the \$STDLIST of a recovery job if the operator request failed. Call the system manager.
RPERR 79	SWITCH command aborted due to errors. This message is located on the \$STDLIST of a recovery job if the operator request failed. Call the system manager.
RPERR 80	Logfile "logfilename" is not available. During rollback, a previous log file was required, but not found. Restore the backup copy of the required log file and respond to the associated console request.

RPWARN 81	Rename logfile to "logfile-name." RP couldn't. Rollback was completed successfully, however, the renaming of the log file to "logfile-name" was not possible. Rename the log file to "zog" plus the filename manually under MPE.
RPMSG 82	Recover SAVE file tables from log record nnn to nnn. A previous run of RP terminated abnormally. As a result, the SAVE file tables are out of date. The RP will update the SAVE file. No action is required.
RPMSG 83	Actively check last record nnn. The RP restart code will check the last log record to verify the database I/O is completed. No action is required.

Function Messages

The following are function messages associated with the *Shadow D/R* program:

RQWARN 1	A backup request is already pending. An outstanding backup transmission is waiting to complete. Do not request a backup when there is one already pending.
RQERR 2	Bad checksum = different configuration used. The wrong configuration file was specified. Supply the correct name of the configuration file.
RQERR 3	Recovery Process internal trans'tab broken. The RP has a bad internal table. Call the system manager.
RQERR 4	Not all requested data bases were opened. The RP could not open all the databases specified in the configuration file. Call the system manager.
RQERR 5	Error accessing the Command msg file. There was a pipe system error in accessing the COMMANDS pipe file. Call the system manager.
RQERR 6	Data compare error between the two logfiles. The RP detected a database inconsistency. Determine the cause of the inconsistency, then synchronize the databases, log files and SAVE file before restarting <i>Shadow D/R</i>.

RQERR 7	Error accessing the data msg file. The secondary system or RP had a file system error accessing the data pipe file. Call the system manager.
RQERR 8	DS PTOP error on Master program (CP). The CP encountered a PTOP error. Call the system manager.
RQWARN 9	There are no remaining DS lines available. The last DS line is out of use. When a DS line becomes available, try again.
RQWARN 10	No Recovery Process on the secondary computer. There was no reply to a heartbeat. Check for the RP job stream on the secondary system. Determine the cause of the error in the \$STDLIST spoolfile.
RQERR 11	IMAGE database intrinsic error on secondary. The RP got a DB status not equal to zero. See the RP job stream for information.
RQMSG 12	The requested DS line is already in use. The USE command cannot open a line that is already open. No action is required.
RQWARN 13	The requested command is not allowed here. The SWITCH command is not permitted on the primary system (it is almost the only command allowed on the secondary system). Only enter the SWITCH command on the secondary system.
RQERR 14	Error accessing the User Logging file. There was a logging system error. Call the system manager.
RQERR 15	Error executing the OPENLOG intrinsic. The CP could not open the log file. Call the system manager.
RQERR 16	Internal error getting logging status. A log status request from the CP failed. Call the system manager.
RQMSG 17	There is no backup in progress at this time. A RESUME command was given when no BACKUP or STOP was pending. No action is required.

RQERR 18	No response to command/request issued. A request was not answered either by the RP or by the CP on either system. Check the <i>Shadow D/R</i> system status.
RQMSG 19	Recovery data pipe nnn has records remaining. A diagnostic which under the current implementation is not printed. Used for internal module communication. No action is required.
RQMSG 20	The requested DS line is not in use. The RELEASE command cannot operate on a line which is not active. Resume DS line operation before issuing this command.
RQERR 21	A request is already in process. The command pipe file is not empty (internal error). Call the system manager.
RQERR 22	Error issuing POPEN intrinsic to start secondary CP. <i>Shadow D/R</i> cannot open the secondary CP. See the job stream for information.
RQERR 23	Error in PTOP master call. There was a PTOP error in the CP on the primary system. See the job stream listing for information.
RQERR 24	Error in issuing REMOTE HELLO command. The COMMAND intrinsic failed. Check the jobfence and session limit on the secondary system.
RQERR 25	Error accessing the RESPONSES file. There was a file system error reading from or writing to the RESPONSES file. Call the system manager.
RQERR 26	Error while restarting Shadow DS line. The RESUME command failed. See the RP and CP job stream listing for information.
RQERR 27	Error accessing the SAVE file on secondary. Could not open or read the SAVE file. Call the system manager.
RQERR 28	Record number mismatch in database or logfile. A DBPUT on a detail record detected a database mismatch. Call the system manager.

RQERR 29	Unsuccessful attempt to start up Shadow. A START command failed. Reasons for the failure precede this message line. Call the system manager.
RQERR 30	Should not stream the file to start CP or RP. An MPE COMMAND intrinsic failed. Call the system manager.
RQERR 31	Error writing to \$OLDPASS for stream file. An MPE error occurred in constructing the stream file text. Call the system manager.
RQERR 32	Unable to communicate with Tape Process. A file system error occurred while accessing the TP. Call the system manager.
RQERR 33	WRITELOG intrinsic returned an error. There was a logging system error. Call the system manager.
RQMSG 34	Number of DS lines opened = nnn. This specifies the number of DS lines open that were successful. No action is required.
RQERR 35	No DS lines could be opened. No DS line opens were successful. Check the status of the communication lines.
RQMSG 36	Secondary system was not suspended for backup. The RESUME command was issued when no backup was pending or completed. Review operational procedures and issue a STATUS command to check the flags on the secondary.
RQMSG 37	Stop pending on transaction(s) in progress. The secondary system is looking for a quiet period to begin backup. No action is required.
RQMSG 38	Backup pending on outstanding transactions. The secondary system is looking for a quiet period to begin backup. No action is required.
RQERR 39	SAVE file indicates bad database. The secondary database is corrupt. See RPERR 12. Follow secondary system major failure recovery procedures.

RQERR 40	Incompatible SAVE file version. The SAVE file being used on the secondary is from an older version of Shadow. When updating to A.00.07 (or later) from any earlier versions, start a new log cycle.
RQERR 41	Target process QUIT, see parameter. This indicates an internal error (encountered a fatal error). Check the listing for previous messages.
RQERR 42	U-MIT logfile sequence no. on non U-MIT. The user backdated MPE without purging the SAVE file. Start a new log cycle when backdating.
RQERR 43	LOGINFO call failure while opening logfile. The log file open could not be completed. Check the listing to determine the error.
RQERR 44	Out-of-sequence logfile sequence number. The log cycle was incorrectly started, or the SAVE file was not changed with the new log cycle, creating a restart sequence number greater than the current sequence number. Purge the old SAVE file and restart. (See CPERR 50.)
RQERR 45	Logfile name doesn't end in 3 digits. CHANGELOG requires special file names. Check your TurboIMAGE manual for user logging information and verify the configuration has the correct log file name.
RQERR 46	Out-of-sequence log record or file. The log files on the primary and secondary are not identical, or the starting log record number is greater than the current maximum log record number. Check operational procedures when starting a new log cycle for the secondary.
RQERR 47	Logfile open not successful. A new log cycle was started with an old SAVE file. Purge the old SAVE file and restart. (See CPERR 50.)
RQWARN 48	REMOTE HELLO warning: CIWARN = nnn. A non-fatal error occurred in the REMOTE definition in the configuration file. Examine CI errors and the configuration file to determine the cause.
RQWARN 49	Command not allowed - check status. While RP is recovering its tables, certain MP commands are not allowed. See RPMSG 82. No caution is required. Allow operation to complete.

Messages

RQWARN 50	Converting old SAVEFILE version. A previous or backup SAVE file corresponds to the old Shadow version. This requires a conversion of formats. Old SAVE file is renamed to ZAVE file. No action is required.
CPERR 70	Parameter returned by above error: nnn. This message informs the user of the error number. Refer to preceding error messages output by the system.
RPERR 75	Parameter returned by above error: nnn. Returns an error diagnostic for any secondary CP error on a request from the RP. Refer to the preceding error message output by the system.

Quit Messages

The following messages are associated with the Quit process:

QUIT 1	Can't build or access MSG file to Tape Process. An MPE error occurred while building or accessing the message pipe file. Investigate the ability to create files.
QUIT 2	Can't build or access MSG file from Tape Process. An MPE error occurred while building or accessing the message pipe file. Investigate the ability to create files.
QUIT 3	CP log file read failure. An MPE error occurred on either system when the CP tried to read the log file. Determine why reading the current log file is not possible.
QUIT 4	FLUSHLOG call failed. An MPE error occurred while executing the FLUSHLOG intrinsic. There is an internal MPE problem. Contact technical support at Hewlett-Packard.
QUIT 5	Tape Process tape write failure. A file system error occurred when the TP tried to write a tape. There is a possible tape drive or media problem. Contact technical support at Hewlett-Packard.
QUIT 6	Tape Process couldn't get DL area buffer. An MPE error occurred when the CP tried to reserve DL space. Contact technical support at Hewlett-Packard.

QUIT 7	CP couldn't get DL area Send Buffer. An MPE error occurred when the CP tried to reserve DL space. Contact technical support at Hewlett-Packard.
QUIT 8	CP couldn't get DL area Log Buffer. An MPE error occurred when the CP tried to reserve DL space. Contact technical support at Hewlett-Packard.
QUIT 9	CP had unexpected file error reading MP command. An MPE error occurred when the CP tried to read the message pipe file from the MP. Contact technical support at Hewlett-Packard.
QUIT 10	RP couldn't update SAVE file. A file system error occurred while writing to the SAVE file. Contact technical support at Hewlett-Packard.
QUIT 11	RP can't continue processing due to above error. This message is generated after any fatal error. The cause of the error will always precede this message. Do not purge the SAVE file, it is useful for diagnostics. Contact the technical support team at Lund Performance Solutions.
QUIT 12	RP user/open table overflow for DBOPEN. This message is generated after an internal error. Contact the technical support team at Lund Performance Solutions. Do not purge the SAVE file, it is useful for diagnostics.
QUIT 13	RP transaction table overflow for DBBEGIN. This message is generated after an internal error. Contact the technical support team at Lund Performance Solutions. Do not purge the SAVE file, it is useful for diagnostics.
QUIT 14	RP user/open table overflow for DBBEGIN. This message is generated after an internal error. Contact the technical support team at Lund Performance Solutions. Do not purge the SAVE file, it is useful for diagnostics.
QUIT 15	RP corrupt user/open table. Run SHADOWMP on SAVE file. Save output for Lund Performance Solutions analysis. This message is generated after an internal error. Contact the technical support team at Lund Performance Solutions. Do not purge the SAVE file, it is useful for diagnostics.
QUIT 16	RP failed database intrinsic call, can't continue. A database intrinsic returned an unexpected error. DBEXPLAIN output precedes this message. Investigate the reason for failure in the DBEXPLAIN message.

QUIT 17	RP couldn't re-open databases after backup request. DBOPEN failed. DBEXPLAIN text precedes this message. Investigate the reason for open failure in the DBEXPLAIN message.
QUIT 18	RP couldn't expand DL area for program tables. This message is generated after an internal error. Contact the technical support team at Lund Performance Solutions.
QUIT 19	CP logging status request failed. An internal error occurred while accessing logging tables. Contact the technical support team at Lund Performance Solutions.
QUIT 20	CP close of logfile failed during CHANGELOG. A file system error occurred while accessing the next log file during the CHANGELOG processing. Contact the technical support team at Lund Performance Solutions.
QUIT 22	RP determined next logfile not DISC logfile. CHANGELOG was done specifying the next log file is on a serial device. Not allowed. Contact the technical support team at Lund Performance Solutions.
QUIT 23	CP couldn't open the logfile. During CHANGELOG, the next file in sequence couldn't be accessed. An error message precedes this message. Contact the technical support team at Lund Performance Solutions.
QUIT 24	SHADOWRP unable to change DB's. Run SHADOWDP on SAVEFILE. When outstanding transactions exist, databases cannot be added or deleted from a configuration file in order to restart <i>Shadow D/R</i>. Run SHADOWDP on the SAVE file. Contact the technical support team at Lund Performance Solutions.

List of Figures

System Description

System Components	4
System Flow	7
Database Replication	8
Offline Backup (before)	10
Offline Backup (during)	11
Offline Backup (after)	12
Quiet Period vs. Busy Period	13

System Configuration, Setup and Startup

Shadow D/R Configuration File	16
Verification of Configuration File	20

Normal Operations

STATUS Command (status of backup on the secondary system is "Now")	26
STATUS Display (record#'s are equal for both systems)	27

Secondary Failure and Recovery

Tape Process	31
STATUS Display (secondary and primary systems are current)	32

Primary Failure and Recovery

STATUS Display (paths at 0)	40
---------------------------------------	----

Data Capture

Data Capture	45
------------------------	----

Commands

BACKUP Command	74
CONFIG Command	75
HELP Command	76
KILL Command	77
RECOVER Command	78
RELEASE Command	79
RESUME Command	80
START Command	81
STATUS Command	81
STOP Command	83

List of Figures

SWITCH Command84

SYSTEM Command.....85

USE Command86

VERIFY Command.....87

VERSION Command87

Index

C

CHANGELOG, 51, 52

Command Requests, 57

from the MP

OPENBASE, 57

RECOVER, 57

SWITCH, 57

WHO'R'U, 57

from the SP

BACKUP, 58

GETPARMS, 57

GETSTATS, 58

HEARTBEAT, 58

OPENBASE, 57

RECOVER, 58

RESUME, 58

SHUTDOWN, 57

Commands, 72

ALTLOG, 47

BACKUP, 73

CHANGELOG, 47

CONFIG, 74

EXIT, 75

GETLOG, 47

HELP, 75

KILL, 76

LOG, 47

RECOVER, 77

RELEASE, 78

RELLOG, 47

RESUME, 79

SHOWLOGSTATUS, 47

START, 80

STATUS, 81

STOP, 83

SWITCH, 84

SYSTEM, 85

TSTART, 66

USE, 85

VERIFY, 66, 86

VERSION, 87

Communications Process (CP), 4, 68, 69, 95

Configuration, 16

configuration file, 5, 16, 17

configuration messages, 88

configuring the primary system, 19

configuring the secondary system, 20

Customer Support (see Technical Support)

D

Data Capture

about log record blocks

log file reads, 50

outstanding log records, 50

record transmission, 50

about the "heartbeat", 52

about the data capture mechanism, 49

about the WAIT= parameter

timeout implementation, 49

wait time selection, 49

activities of the tape process, 50

enabling the database logging facility, 47

enabling the user logging facility, 45

CHANGELOG records, 51

Data Posting

receiving and executing commands, 56, 57

receiving and validating data, 56

receiving the duplicate log file, 58

tracking database users, 58

updating the duplicate database(s), 58, 59

Data Reliability, 8

Data Transport

data blocks, 54

data transmission frequency, 8, 54

multiple data communication links, 53

network links, 53

Database Duplication

about the "heartbeat" status request, 9

data reliability, 8

data transmission frequency, 8

Diagnostics

diagnostic commands, 66

diagnostic messages, 66

DS Lines

- failure, 29
- recovery, 29, 30

F

Files

- acceptable file names, 5
- configuration files, 5, 16
- log files, 50
- message files, 5
- SAVE files, 5, 17
- Shadow D/R files, 15
- tape files, 51

Frequently Asked Questions, 69

H

Heartbeat, 9, 52, 58

I

Installation

- loading the Shadow D/R files, 15
- starting a new log cycle, 15

Intrinsics

- MPE
 - CLOSELOG, 46
 - FLUSHLOG, 46
 - LOGINFO, 46
 - LOGSTATUS, 46
 - OPENLOG, 46
 - WRITELOG, 46
- Shadow D/R
 - FLUSHLOG, 49
 - FOPEN, 49
 - FREADDIR, 49
 - LOGSTATUS, 49
- TurboIMAGE
 - DBBEGIN, 48
 - DBCLOSE, 48, 59
 - DBDELETE, 48, 59
 - DBDEMO, 48
 - DBEND, 48, 59
 - DBMEMO, 60
 - DBOPEN, 48, 59

- DBPUT, 48, 59
- DBRECOV, 70
- DBUPDATE, 48, 59

L

Log Cycles, 15, 26

Lund Performance Solutions

- business hours, 1
- electronic mail address, 1
- land mail address, 2
- phone number, 1
- technical support team, 1
- world wide web site, 1

M

Manager Process (MP), 4, 57, 69, 90

Messages, 88

- communications process messages, 95
- configuration messages, 88
- diagnostic messages, 66
 - CP states, 68
 - requests from the CP to the SP, 67
 - requests from the MP to the CP, 67
 - requests from the MP to the RP, 67
 - requests from the SP to the RP, 68
- error messages
 - communications process errors, 69
 - manager process errors, 69
 - recovery process errors, 69
 - TurboIMAGE intrinsic errors, 69
- manager process messages, 90
- quit messages, 111
- recovery process messages, 99

MPE

- user logging facility, 45
 - logging commands, 47
 - logging intrinsics, 46

N

Normal Operations

- backing up the database, 26
- starting a new log cycle
 - with automatic CHANGELOG, 27
 - without automatic CHANGELOG, 28

O

Offline Backup

- waiting for a quiet period, 12
- what happens after an offline backup, 11
- what happens before an offline backup, 9
- what happens during an offline backup, 10

P

Parameters

- BEGIN, 17
- COMMANDS, 17
- DSLIME, 18
- END, 19
- LOCAL, 17
- LOG, 19
- MAXRECS, 18
- MODE, 18
- MODE=, 64, 70
- REMOTE, 19
- RESPONSES, 17
- SAVEFILE, 17
- TAPEDEV, 19
- USE, 17
- WAIT, 18
- WAIT=, 49

Primary System

- configuration, 19
- failure and recovery
 - about primary system failure, 34
 - example of a primary system recovery, 41
 - recovering from a major failure, 35
 - recovering from a minor failure, 34
 - returning systems to their original states, 39
- setup, 21

Processes

- communications process, 4, 68, 95
- manager process, 4, 57, 90
- recovery process, 4, 99
- secondary communications process, 55, 57
- tape process, 5, 31, 50

Q

Quiet Periods, 12

R

Recovery Process (RP), 4, 69, 99

Reports

- RECOVER reports, 6
- STATUS reports, 6
- SWITCH reports, 6

S

Secondary Communications Process (SP), 55, 57

Secondary System

- configuration, 20
- failure and recovery
 - about secondary system failure, 31
 - assessing a secondary failure, 32
 - enabling a database rollback, 33
 - recovering from a major failure, 33
 - recovering from a minor failure, 33
- setup, 22, 23

Setup

- installing files, 22
- setting switch environments, 23
- setting system parameters, 23
- setting up databases and logging, 21

Shadow D/R

- errors, 69
- intrinsic, 49
- operations
 - backing up the database, 26
 - balancing the data load, 13
 - capturing data, 44
 - duplicating the database, 7
 - performing an offline backup, 9
 - posting data, 56
 - starting a new log cycle, 26
 - transporting data, 53
- product features and advantages, 1
- reports, 5
- security solutions, 64
- system commands, 72
- system components, 4
- system configuration, 16
- system description, 4
- system design, 6
- system failure and recovery

Index

- recovering failed DS lines, 29
- recovering the primary system, 34
- recovering the secondary system, 31
- system files, 5
- system messages, 88
 - configuration messages, 88
 - CP messages, 95
 - MP messages, 90
 - quit messages, 111
 - RP messages, 99
- system parameters, 17
- system processes, 4
- system security, 64
- system setup, 16, 21
- system shutdown, 24
- system startup, 16, 24
- system troubleshooting
 - diagnostic tools, 66
 - frequently asked questions, 69
- systems
 - primary, 4
 - secondary, 4

Shutdown, 24

- performing a standard shutdown, 24
- performing an abrupt shutdown, 25

Startup, 24

T

Tape Process (TP), 5, 31, 50

- activation, 51
- creation, 51
- deactivation, 51

Taping Threshold, 18

Technical Support

- LPS business hours, 1
- LPS electronic mail address, 1
- LPS land mail address, 2
- LPS phone number, 1
- LPS technical support team, 1
- LPS world wide web site, 1

TurboIMAGE

- database logging facility, 47
- errors, 69
- intrinsic, 48, 59, 70

- passwords, 65
- user logging facility, 51

U

User Manual

- corrections and comments, 2
- overview of contents, 2