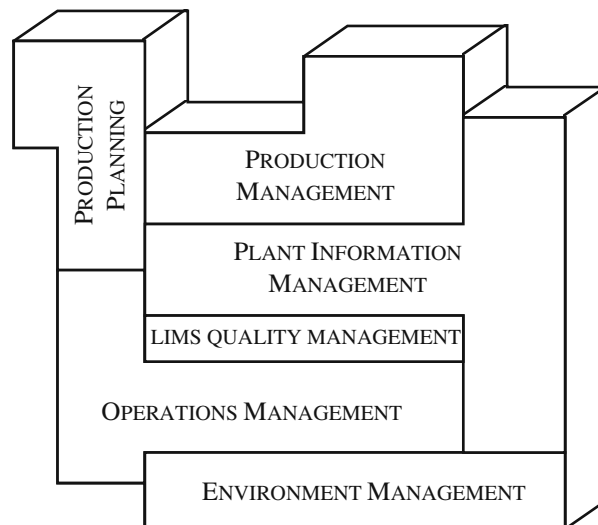


PHD/Oracle Backup Strategies



Copyright, Notices, and Trademarks

Honeywell IAC courseware is subject to change without notice.

This courseware is copyrighted and all rights are reserved by Honeywell Inc. These materials are intended for use solely in conjunction with Honeywell products. The materials comprising the courseware may not, in whole or in part, be copied, photocopied, reproduced, translated or reduced to any electronic medium or machine readable form without the prior, express written consent of Honeywell Inc.

No Warranty: Honeywell Automation College provides no warranties respecting the program, documentation or services provided here under, either expressed or implied, including but not limited to any warranty of design, merchantability, or fitness for a particular purpose, even if Honeywell Automation College has been informed of such purpose. No agent of Honeywell Automation College is authorized to alter or exceed the warranty obligations of Honeywell Automation College as set forth in this agreement.

TotalPlant is a US registered trademarks of Honeywell Inc.
Other brand or product names are trademarks of their respective owners.

Honeywell
Industrial Automation and Control
Automation College
2820 West Kelton Lane
Phoenix, AZ USA 85053
1-800-852-3211

PHD/Oracle Backup Strategies

The PHD Administrator should know how to perform these types of backup strategies around PHD and Oracle:

- Full Backups
- Incremental Backups
- Online Backups
- Offline Backups

The following are guidelines for developing backup strategies. Use the guidelines as a starting point. Your final strategy should take care of specific site/project requirements and hardware limitations.

Each site should document their specific backup procedure.

Each site's PHD administrator should go to Oracle training, especially if the site uses Oracle for more than just the PHD point configuration. There are more database administration issues that need to be handled, in addition to backups.

Backup Types

Full Backups—Include everything on the hard disk.

Incremental backups

Online Backups

Offline Backups

You should get Full Backups of *all* archives periodically (daily or weekly). If your Full Backups are weekly, you should get Incremental Backups daily.

Script files should be on every backup. Script files include cold_start.cmd, interfaces.dat, and phd_calculations.phd.

Be sure to backup script files before a release migration, in case something goes wrong during the migration process.

Backup of Archives

Whenever a new scan-archive file becomes the active one, the previous one is made inactive and may be backed up. To reduce the risk of losing any data, backup the previous active archive file and the current active archive file.

To backup of the *previous* archive file, just perform a file backup.

If you decide to backup the scan-archive file as soon as it is made inactive, you may choose to have PHD call an archiving procedure at the time it creates a new archive. The archiving procedure is a command procedure normally used to delete the oldest scan file in the archive set; however, you can use it for anything. You can modify the procedure to grab the previous archive and copy it to a safe area until the standard backup can run.

If you follow a standard backup regime (daily), you do not necessarily have to backup the scan-archive file as soon as it is made inactive. The potential exists for losing a maximum of one day's worth of data; however, it is unlikely and would be the result of some catastrophic failure.

Another consideration for backups of the previous archive file is if your end users do manual entries to archives other than the current archive, there is a potential for backup problems.

To backup the *current active* archive file, do the following:

1. Pause the Constore process: PHDMAN> SET STORE:ENABLE 0

(Stopping the Constore process prevents it from modifying the scan-archive file. This guarantees that the index file will be kept synchronized with data file. The only way synchronization can be guaranteed is if nothing is writing to the files at the time the backup is made.)

2. Since Constore may have started writing out data already, wait n seconds, where n is the Constore interval (most sites leave the interval at the default of 300 = every 5 minutes).
3. Copy the current archive file to another place on the hard disk.
Note: You must allocate space on the hard disk for this purpose.
4. Resume the Constore process: PHDMAN> SET STORE:ENABLE 1
5. Backup the copy of the current archive file to tape or other media.

Considerations

The minimum recommendation for the default Data Queue size (DEF_QMINUTE parameter) is twice the store interval plus the amount of time it takes to backup the archive.

When PHD storing is stopped (constore inactive), the data queues are not being emptied and there is a potential for data to be lost. **You should carefully determine the required size for your data queues**, as the size determines the buffer time.

Generally, the data queues are set to 15 minutes by default in the phdparams.cmd file. (If the DEF_QMINUTE is not set in the phdParams.cmd file, then the true default is 30 minutes.)

Example:

Store Interval = 600 seconds (10 minutes)

Time required to copy archive = 5 minutes

DEF_QMINUTES should be 25 minutes.

Example:

DEF_QMINUTE = 15 minutes.

Store Interval = 300 seconds (5 minutes)

If you stop Constore and wait 5 minutes, you have 10 minutes to complete the backup and get Constore started again.

Depending on the size of your archive files and speed of your archive device (tape drive?), this may or may not be a problem.

A tape backup may not be possible during the buffer time; therefore (to prevent queue overwrites and lost data), ensure that there is disk space to copy the archive files to a temp directory or network drive on hard disk first. Backup the files to the archive device from there. This is a good, quick way to “snapshot” the files so you can set Constore active again.

Backup of Oracle

Oracle supports online and offline backups.

Online backups are called 'archiving logging' and allow you to maintain the backup up to the current open transaction.

Offline backups tend to be done periodically (recommendation is daily). Using an offline backup has the potential, on catastrophic failure, to lose data from now to your last backup. Even PHD-only implementations may need more frequent backups (such as for tag synch records and event journal records).

Backing up Oracle is not trivial; typically, a rigorous procedure includes

Archive Logging—

Exports—Full export dump files also contain structural information (in logical form) that can be recovered. Exports (for the typical size that databases grow to) allow any one piece to be recovered.

Image Backups—Image backups are all or nothing when it comes to recovery.

Honeywell recommends a full Oracle backup that captures data, structure, security, etc. The Oracle backup utility is used to capture full oracle backups. Backup of the physical database files is not necessary. Full backups require that all users and background processes be disconnected from Oracle.

Because the PHD oracle database does not change continuously, Honeywell set up the database to run in the NOARCHIVELOG mode, giving protection against loss of the server's memory (instance failure), but not against hard disk failures.

To reduce the impact of a hard disk failure, do a full backup and export.

Full Backup (structure + content)—Offline Backup

The full backup is an off-line procedure. This means that the database should be stopped before a backup can be taken. On NT, Oracle includes an NT backup manager utility that can do this task automatically for you: stopping the database, backing-up, and restarting the database.

Here is a method to find the files that make up the database.

Run the following SQL statements:

select * from v\$datafile (returns the datafiles)

select * from v\$logfile (returns the redo log files)

And backup as well:

initprd.ora

(inside this file, the path for the control files can be found)

control_files

Export (content only)—Online Backup

The Oracle export utility is an on-line utility that allows you to dump the content of the database into one file (Oracle format). This file can be used to recover through the import utility. Note that even though this utility works on-line, care should be taken that nobody modifies the database at the time the export is done.

To recover from a hard disk crash, the entire Oracle database should be rebuilt. This means that the structure of the database must be recovered. This can be done starting from the Honeywell installation CDROM or from a full backup. Once the structure is recovered, data can be restored through the import utility.

Export command example: (from the Command Prompt)

```
EXP73 TOTALPLANT/<password>@TOTALPLANT FULL=Y FILE=BKUP.DMP  
BUFFER=2000000
```

(This command can be scheduled.)

Example PHD Backup Procedures

The following procedures were created for a specific customer project and were supplied to Automation College by Diane Pedersen, Honeywell HiSpec, Belgium.

Background

Backup directly to tape can take a very long time, and the customer wanted to leave the Oracle database in "backup" mode for the shortest time possible. (The database has been put in ARCHIVELOG mode, so on-line backups are possible.) Since they have the Event Journal application, their Oracle database is large and is constantly getting a lot of new records inserted into it. For the PHD part, the STORE flag is set to DISABLE during the time that the active PHD archive file is being backed up.

To make sure that the backups take as little time as possible, a daily backup to disk is done. The NT Schedule service is used to schedule 2 command files, one for PHD and one for Oracle, to run in the middle of the night. These procedures do the backups to disk, and then the Seagate Backup Exec software (included in the delivery of the Dell machine) is used to copy the backed up files from disk to tape every night after the "disk" backup runs. (Of course this assumes there is lots of free disk space to work with.)

The Oracle backup takes all the database, archive and log files, plus the database control file. The PHD backup takes only the two most recent PHD archive files and their associated index files. The customer has about 45GB of disk space for the PHD archive files, so it is not feasible to backup everything every day. The customer has put in place an additional procedure to make sure that all the backed up PHD archive files get put into long-term storage.

The overhead required to use this method:

- *Must run Oracle in ARCHIVELOG mode,*
- *Need to manage the Oracle log files separately,*
- *Need extra disk space to hold the backups, and*
- *Need a third party package to copy the files to tape, etc.*

10.3 Backup Procedures

The following sections describe the backup procedures for the PHD Shadow Server. This includes a daily backup of the Oracle database and the most recent PHD archive files. A manually initiated full backup is also available. The backups are implemented through a combination of some custom programs, and through the configuration of the Seagate Backup Exec software that was provided by Dell with the PHD Shadow Server machine.

Daily backups of the PHD Nodes are not included, since the PHD Nodes contain no Oracle database and the PHD archive data on these nodes is duplicated on the PHD Shadow Server machine. These nodes do, however, contain some PHD configuration information, and should be backed up once after installation and then again if a significant configuration change is made.

10.3.1 Backup Strategy for PHD Shadow Server

The PHD Shadow Server is a process historian which stores large amounts of data, both in its Oracle database and its PHD archive files. Because of the amount of data stored on this machine, it is not feasible to do a full backup of the machine on a daily basis. Therefore, custom routines have been developed to copy only the most important files to a backup directory on the disk, then the files are uploaded from this backup directory to tape using the Seagate Backup Exec software.

The normal backup procedure at Customer Site A is to perform a daily backup to a separate tape every night of the workweek. The tape that is inserted on Friday is overwritten on Saturday and Sunday nights. Once per month, the backup tape is taken and put into long-term storage.

The above procedure will work fine for the Oracle database backup, but a slightly different procedure is needed for the PHD archive files.

The PHD Shadow Server is configured to hold a maximum of 48 archives. Each PHD “archive” actually contains two files, a .DAT file that holds the actual process data, and an .IDX file containing the information needed for PHD to access the data in the .DAT file. The size of the .DAT file at Customer Site A is configured to be a maximum of 500 MB. Each time an archive fills up, PHD automatically creates another archive, up to the configured limit. After that, the oldest archive is overwritten each time a new archive is created.

The daily backup of the PHD archives includes only the two most recent archives. The current archive is backed up because it contains the most recent process information. Since it is not possible to predict exactly when an archive will fill up and a new one will be created, the second most recent file is also backed up daily.

The backup strategy must also ensure that the older archives are backed up and put in long-term storage. The standard practice of putting one tape per month into long-term storage should result in all archive files being placed in long-term storage, since each archive holds approximately one month of data, and there are always two archives on each tape. This practice, however, makes it difficult to find old archives in the event that they need to be restored to disk. Each tape has two archives on it, and you must work backwards to find the tape containing the most recent version of each archive. Depending on how much data is stored, data compression, etc., there may be more than one tape with the same two archives on it.

It is, therefore, recommended to make a separate set of tapes containing only the PHD archives. This is a manual process that can be performed, for example, once every 3 months. If there is a need to recover archives less than 3 months old, they can be individually restored from the daily backup tapes that were put into long-term storage (one tape for each month). All archives older than 3 months can be found all together on the PHD Archive backup tapes. A longer interval between "full archive backups" can also be used (that is, 6 months or 1 year), but if the on-line archives are lost, it will require more work to find the correct files to restore from the individual daily backup tapes.

10.3.2 Oracle Database Backup & Restore

10.3.2.1 Oracle Database Backup Procedure

An automatic backup procedure for the Oracle database at Customer Site A has been developed. This backup procedure resides in the *F:\Backup* directory on the PHD Shadow Server machine. When executed, it puts the Uniformance Oracle database into "backup mode", then copies the tablespaces and control files to the *F:\Backup\Oracle* and *F:\Backup\Oracle\Archive* directories, then puts the database back into the normal mode of operation.

The backup procedure consists of one NT script file and two SQL script files:

- *Oracle_Backup.cmd* (main command procedure)
- *Start_Oracle_Backup.sql* (called by *Oracle_Backup.cmd*)
- *End_Oracle_Backup.sql* (called by *Oracle_Backup.cmd*)

The main procedure *Oracle_Backup.cmd* file is scheduled to run once every day in background mode using the standard NT Schedule service. The NT Schedule Service must be configured to run in a special NT account that has the Administrator, ORA_DBA, and all PHD privileges. The NT user account named "Schedule" has been created on the PHD Shadow Server machine for this purpose.

To view all scheduled jobs:

- Open a DOS window on the PHD Shadow Server
- Enter the command AT. The job IDs are shown in the first column of the display, the scheduling information appears in the next two columns, and the last column shows the command that is executed to start the backup.

To delete the job from the queue, and thus stop the automatic backup procedure:

- View all scheduled jobs as described above, then find the ID number associated with the *Oracle_Backup.cmd* job
- At the DOS prompt, enter the command AT ID /DEL, where *ID* is the number of the job

To add the backup procedure to the queue, and thus schedule the backup procedure for automatic execution:

- At the DOS prompt, enter the command AT hh:mm /every:M,T,W,Th,F,S,Su F:\Backup\Oracle_Backup.cmd, where *hh:mm* is the time (hour and minute in 24 hour format) at which to execute the backup procedure. The current *hh:mm* setting is 01:30.

10.3.2.2 Oracle Database Restore Procedure

In the event of an Oracle database crash, the standard Oracle recovery procedures should be followed.

10.3.3 PHD Archives Backup & Restore

10.3.3.1 PHD Archives Backup Procedure

An automatic backup procedure for the two most recent PHD SCAN archive files has been developed for Customer Site A. Because PHD archive files are large and the number of these files grows over time, only the two most recent SCAN archive files (and their associated index files) are included in the automatic backup procedure. This ensures that all new data over a 24 hour period is archived, even when one archive file fills up and another is automatically created by PHD. The size of the archive files is sufficiently large so that no more than one new archive file will be created in any 24 hour period. The older archive files that are not included in the automatic backup should be archived separately and put in long term storage.

This backup procedure resides in the *F:\Backup* directory on the PHD Shadow Server machine. When executed, it runs the FILEMOD.EXE program to get the names of the two most recent archive files and their associated index files (maximum of four files). It then disables PHD from writing to the active SCAN archive file, copies the files to the *F:\Backup\PHD* directory, then re-enables the PHD write to the active SCAN archive.

The backup procedure consists of one NT script file and one executable:

- *Phd_Backup.cmd* (main command procedure)
- *Filemod.exe* (called by *Phd_Backup.cmd*)

The main procedure *Phd_Backup.cmd* file is scheduled to run once every day in background mode using the standard NT Schedule service. The NT Schedule Service must be configured to run in a special NT account that has the Administrator, ORA_DBA, and all PHD privileges. The NT user account named "Schedule" has been created on the PHD Shadow Server machine for this purpose.

To view all scheduled jobs:

- Open a DOS window on the PHD Shadow Server
- Enter the command AT. The job IDs are shown in the first column of the display, the scheduling information appears in the next two columns, and the last column shows the command that is executed to start the backup.

To delete the job from the queue, and thus stop the automatic backup procedure:

- View all scheduled jobs as described above, then find the ID number associated with the *Phd_Backup.cmd* job
- At the DOS prompt, enter the command AT ID /DEL, where *ID* is the number of the job


To add the backup procedure to the queue, and thus schedule the backup procedure for automatic execution:

- At the DOS prompt, enter the command AT hh:mm /every:M,T,W,Th,F,S,Su F:\Backup\Phd_Backup.cmd, where *hh:mm* is the time (hour and minute in 24 hour format) at which to execute the backup procedure. The current *hh:mm* setting is 01:00.

10.3.3.2 PHD Archives Restore Procedure

The logical archives, which are archived through the backup procedures, can be restored to the *F:\PHDServer\Archive* directory on the PHD Shadow Server. This procedure is used to bring data back on-line from a period that is no longer available.

When restoring old archive files, both the .DAT and the .IDX files must be copied to the *F:\PHDServer\Archive* directory. The logical archive named RECOVER is dedicated to restore data into and is not used in the normal archive shift rotation.

 **NOTE:** Recovered archive files should **not** be restored to the SCAN logical archive because they will be automatically deleted by the system when the maximum numbers of archive files exist and a new SCAN archive is created.

To create the RECOVER logical archive (this only needs to be done once):

- Open a DOS window on the PHD Shadow Server
- Enter the PHDMAN environment by entering the command PHDMAN
- Create the RECOVER logical archive by entering the command DEFINE LARC RECOVER
- Check the status of the logical archives by entering the command SHOW LARC

To connect an archive file to the RECOVER logical archive:

- Use the Windows Explorer to copy the old archive files to the *F:\PHDServer\Archive* directory (for example, SCAN00001.DAT and SCAN00001.IDX)
- Open a DOS window on the PHD Shadow Server
- Enter the PHDMAN environment by entering the command PHDMAN
- Connect the old archive to the RECOVER logical archive by entering the command CONNECT <arcname> RECOVER, where <arcname> is the name of the archive file (for example, SCAN00001)
- Check the status of the archive files by entering the command SHOW ARC

To disconnect an archive file from the RECOVER logical archive:

- Open a DOS window on the PHD Shadow Server
- Enter the PHDMAN environment by entering the command PHDMAN
- Disconnect the old archive from the RECOVER logical archive by entering the command DISCONNECT <arcname> RECOVER, where <arcname> is the name of the archive file (file (for example, SCAN00001)
- Check the status of the archive files by entering the command SHOW ARC

10.3.4 Daily Backup to Tape

The above procedures described the daily backups of the PHD Shadow Server that are done to disk. After the disk backup is complete, the files are copied from the backup directory (*F:\Backup*) to the tape. A disk backup is performed before the tape backup for the following reasons:

- PHD may not write data into the current archive while it is being backed up. During the backup, PHD must keep all new data in memory buffers. For Customer Site A, approximately 20 minutes of data can be stored in memory, and if the backup takes longer than that, data can be lost. Backing up two archives to disk takes only a fraction as long as backing up directly to tape.
- The Oracle database must be put in “backup” mode during the backup. Backing up to disk reduces the time that the database must be in this mode.
- The backup will not be lost if there is an error writing to the tape, the tape is full, or there is no tape mounted in the drive. In this case, the backup can be taken manually by copying all files and subdirectories from the *F:\Backup* directory onto the tape.

10.3.4.1 Seagate Software Backup Exec Configuration

A software product called *Backup Exec* is used to schedule and run the daily backup to tape. This software can be accessed directly from the shadow server, via the NT Start Menu:

Start → Programs → Seagate Backup Exec → Seagate Backup Exec

The main window appears with a number of tabs at the bottom to toggle among the different windows. The windows that are important for this backup procedure are:

- **Job Monitor** – Shows the job status, completion status, start time, elapsed time of the job. Double-click on a job to see more details about a job in progress. The system has been configured to keep history for up to 31 jobs. Jobs older than this are automatically purged by the system.
- **Job Definition** – Shows a list of the configured jobs. Double-click on a job to see the detailed job configuration.
- **Alerts** – Shows any error messages, warnings or other user messages relating to processing of the jobs.

Log files are generated during the processing of each job. The log files show summary information about the job execution, plus a list of the directories that have been processed. An option exists to send more detailed information to the log files if desired. The log files are named *BEXnn.TXT*, and are located in the following directory on the PHD Shadow Server:

C:\Program Files\Seagate Software\Backup Exec\NT\Data

One job has been configured to perform the daily backup. This information can be viewed by performing the following steps from the Backup Exec menu:

Jobs → Setup → Select PHD Shadow Server Daily → Select the *Edit* button

The configuration details are as follows:

Tab / Dialog	Parameter	Setting
General	Job Name	PHD Shadow Server Daily
	Write Option	Overwrite
	Media Name	Media created on mm/dd/yy hh:mm:ss am/pm
	Backup Set Description	PHD Shadow Server Backup (Oracle & PHD)
	Backup Method	NORMAL – Backup Files – Reset Archive Bit
	Device	ARCHIVE 0
	Media Set	Media Set 1
	Verify after backup completes	Enabled
Advanced	Compression	None
	Backup Open Files	Never
	Files To Back Up	F:\Backup*. F:\Backup\Oracle*. F:\Backup\Oracle\Archive*. F:\Backup\PHD*.*
Schedule	Execution Frequency	Once every 24 hours
	Execution Time	02h00

A second job has been configured to perform a full system backup. This job is not scheduled, but can be demanded manually from the “Job Setup” dialog box. The configuration details are as follows:

Tab / Dialog	Parameter	Setting
General	Job Name	PHD Shadow Server Full
	Write Option	Overwrite
	Media Name	Media created on mm/dd/yy hh:mm:ss am/pm
	Backup Set Description	PHD Shadow Server Full Backup (All Drives)
	Backup Method	NORMAL – Backup Files – Reset Archive Bit
	Device	ARCHIVE 0
	Media Set	Media Set 1
	Verify after backup completes	Enabled
Advanced	Compression	None
	Backup Open Files	Yes
	Files To Back Up	All Files on Drive C: All Files on Drive E: All Files on Drive F:
Schedule	Execution Frequency	None
	Execution Time	On Demand

A full set of documentation for the Backup Exec software is installed on the PHD Shadow Server, and can be accessed via the NT Start Menu:

Start → Programs → Seagate Backup Exec → Seagate Backup Exec Documentation

The following are the Oracle and PHD command files associated with the previous example procedures.

```

REM -----
REM -- CUSTOMER   : Customer Name
REM -- PROJECT NO. : #####
REM --
REM -- FILE NAME   : ORACLE_BACKUP.CMD
REM -- DESCRIPTION : Oracle Backup Procedure for PHD Database
REM --
REM -- This procedure calls the START_ORACLE_BACKUP.SQL script file
REM -- to set the tablespaces to backup mode. It then uses DOS
REM -- commands to backup the Oracle database and archive files, and
REM -- finally calls the END_ORACLE_BACKUP.SQL script file to reset
REM -- the tablespaces back to the normal mode of operation and to
REM -- backup the PRD database control file.
REM --
REM -- This procedure is scheduled to run in the background once per
REM -- day using the NT "AT" command. NOTE: The NT Schedule Service
REM -- must be configured to run in a special account that has the
REM -- Administrator, ORA_DBA, and all PHD privileges. The user
REM -- account named "Schedule" has been created for this purpose.
REM --
REM -- Honeywell Hi-Spec Solutions
REM -- Brussels, Belgium
REM -- Copyright 1998
REM --
REM -- 13-Jul-98      : Initial Version
REM -- 07-Sep-98      : Rev 1.
REM --              Update password for SQL*Plus login
REM -----

```

```

REM -- Make sure that environment variable definitions remain local
REM -- to this command file.

```

```

SETLOCAL

```

```

REM -- Define the Oracle directory where SQL*Plus is located.

```

```

SET OraDir=E:\ORANT\BIN

```

```

REM -- Define the home directory where this command file will execute.

```

```

SET HomeDir=F:\BACKUP

```

```

REM -- Define the source directory where the database files are
REM -- located.

```

```

SET DbaseDir=E:\ORANT\DATABASE

```

```

REM -- Define the backup directory that the database files will be
REM -- copied to.

```

```

SET BackDbaseDir=F:\BACKUP\ORACLE

```

```

REM -- Define the source directory where the archive files are located.

```

```

SET ArcDir=E:\ORANT\DATABASE\ARCHIVE

REM -- Define the backup directory that the archive files will be
REM -- copied to.

SET BackArcDir=F:\BACKUP\ORACLE\ARCHIVE

REM -- Delete all files from the backup directories before starting
REM -- the new backup.

DEL /Q %BackDbaseDir%\*. *
DEL /Q %BackArcDir%\*. *

REM -- Connect to Oracle database and set tablespaces to backup mode.

%OraDir%\PLUS33 TOTALPLANT/TOTALPLANT@TOTALPLANT
@%HomeDir%\START_ORACLE_BACKUP.SQL

REM -- Copy the Oracle database files to the backup directory.

COPY %DbaseDir%\IP_DATA1.DBS %BackDbaseDir%\*. *
COPY %DbaseDir%\IP_IDX1.DBS %BackDbaseDir%\*. *
COPY %DbaseDir%\USERS1.DBS %BackDbaseDir%\*. *
COPY %DbaseDir%\TEMP1.DBS %BackDbaseDir%\*. *
COPY %DbaseDir%\RBS1.DBS %BackDbaseDir%\*. *

REM -- Copy the Oracle log files to the backup directory.

COPY %DbaseDir%\*.ORA %BackDbaseDir%\*. *

REM -- Copy the Oracle archive files to the backup directory.

COPY %ArcDir%\*. * %BackArcDir%\*. *

REM -- Connect to Oracle database, set tablespaces to normal mode
REM -- and backup control file.

%OraDir%\PLUS33 TOTALPLANT/TOTALPLANT@TOTALPLANT
@%HomeDir%\END_ORACLE_BACKUP.SQL

REM -- End of locally defined environment variables.

ENDLOCAL

```

```

-----
-- CUSTOMER      : Customer Name
-- PROJECT NO.   : #####
--
-- FILE NAME     : START_ORACLE_BACKUP.SQL
-- DESCRIPTION   : Oracle Backup Procedure for PHD Database
--
-- This procedure sets the tablespaces for the PRD database to
-- backup mode. This procedure is called from the ORACLE_BACKUP.CMD
-- command procedure.
--
-- Honeywell Hi-Spec Solutions
-- Brussels, Belgium
-- Copyright 1998
--
-- 13-Jul-98      : Initial Version
-----

-- Set tablespaces to backup mode

ALTER TABLESPACE SYSTEM      BEGIN BACKUP;
ALTER TABLESPACE IP_DATA_TS  BEGIN BACKUP;
ALTER TABLESPACE IP_IDX_TS   BEGIN BACKUP;
ALTER TABLESPACE USERS_TS    BEGIN BACKUP;
ALTER TABLESPACE TEMP_TS     BEGIN BACKUP;
ALTER TABLESPACE RBS_TS      BEGIN BACKUP;

-- Exit back to DOS

EXIT;

```

```

-----
-- CUSTOMER      : Customer Name
-- PROJECT NO.   : #####
--
-- FILE NAME     : END_ORACLE_BACKUP.SQL
-- DESCRIPTION    : Oracle Backup Procedure for PHD Database
--
-- This procedure sets the tablespaces back to the normal mode of
-- operation and backs up the control file for the PRD database.
-- This procedure is called from the ORACLE_BACKUP.CMD command
-- procedure.
--
-- Honeywell Hi-Spec Solutions
-- Brussels, Belgium
-- Copyright 1998
--
-- 13-Jul-98      : Initial Version
-----

-- Set tablespaces back to normal mode

ALTER TABLESPACE SYSTEM      END BACKUP;
ALTER TABLESPACE IP_DATA_TS  END BACKUP;
ALTER TABLESPACE IP_IDX_TS   END BACKUP;
ALTER TABLESPACE USERS_TS    END BACKUP;
ALTER TABLESPACE TEMP_TS     END BACKUP;
ALTER TABLESPACE RBS_TS      END BACKUP;

-- Backup control file

ALTER DATABASE BACKUP CONTROLFILE TO 'F:\BACKUP\ORACLE\CTL1PRD.ORA'
REUSE;

-- Exit back to DOS

EXIT;

```

```

REM -----
REM -- CUSTOMER      : Customer Name
REM -- PROJECT NO.   : #####
REM --
REM -- FILE NAME      : PHD_BACKUP.CMD
REM -- DESCRIPTION    : PHD Archive Backup Procedure
REM --
REM -- This command procedure backs up the 2 most recent SCAN archive
REM -- files from the PHD archive directory, along with their associated
REM -- index files. To find out which are the 2 most recent archives,
REM -- a program called FILEMOD.EXE is executed. This program creates
REM -- a temporary command file that is used to copy the files to the
REM -- backup directory.
REM --
REM -- This procedure is scheduled to run in the background once per
REM -- day using the NT "AT" command. NOTE: The NT Schedule Service
REM -- must be configured to run in a special account that has the
REM -- Administrator, ORA_DBA, and all PHD privileges. The user
REM -- account named "Schedule" has been created for this purpose.
REM --
REM -- Honeywell Hi-Spec Solutions
REM -- Brussels, Belgium
REM -- Copyright 1998
REM --
REM -- 13-Jul-98      : Initial Version
REM -- 07-Sep-98      : Rev 1.
REM --
REM --                                Update DIR command to sort files by creation
REM --                                date and not by modification date
REM -----

REM -- Make sure that environment variable definitions remain local
REM -- to this command file.

SETLOCAL

REM -- Define the home directory where this command file will execute.

SET HomeDir=F:\BACKUP

REM -- Define the drive and directory where the PHD archive files
REM -- reside.

SET ArcDrive=F:
SET ArcDir=F:\PHDSERVER\ARCHIVE

REM -- Define the backup directory where the files will be copied.

SET BackDir=F:\BACKUP\PHD

REM -- Delete old archive files from the "Backup" directory.

DEL /Q %BackDir%\*. *

REM -- Make a list of all the files in the archive directory and save
REM -- it in the "Home" directory. This is a temporary file which is
REM -- deleted at the end of this script.

DIR /b /a:-d /o:d /t:c %ArcDir% > %HomeDir%\FILELIST.TXT

REM -- Turn off PHD archiving for a minute so that the files can be
REM -- copied without interference from PHD.

```

```

%PHD_PATH%\PHDMAN SET STORE:ENABLE 0
%PHD_PATH%\SLEEP 10

REM -- Run the program FILEMOD.EXE to take the last two entries in
REM -- the file created above and create a command file that copies
REM -- these two files to the "Backup" directory. The program
REM -- FILEMOD.EXE executes from the "Home" directory. The file
REM -- specification of the second argument must be the path to the
REM -- "Backup" directory. The command file is a temporary file which
REM -- is deleted at the end of this script.

%HomeDir%\FILEMOD.EXE %HomeDir%\FILELIST.TXT %BackDir%\FILELIST.CMD

REM -- Set the default drive for the command file, and go to that
REM -- directory. The command file runs from the "Archive" directory.

%ArcDrive%
CD %ArcDir%

REM -- Run the command file that copies the last two archive files and
REM -- the last two index files to the "Backup" directory.

CALL %BackDir%\FILELIST.CMD

REM -- Turn PHD archiving back on.

%PHD_PATH%\PHDMAN SET STORE:ENABLE 1

REM -- Delete the temporary files.

DEL %HomeDir%\FILELIST.TXT
DEL %BackDir%\FILELIST.CMD

REM -- End of locally defined environment variables.

ENDLOCAL

```

```

/* filemod.c ...generate a command procedure for copying latest
   2 archive files; needs to be run in connection
   with arcback.cmd command procedure -- which yields the
   text file called filelist.txt which has the
   most recent archive and its corresponding index
   in the last 2 lines */
#include <stdio.h>
#include <stdlib.h>
#include <time.h>
#include <string.h>
#include <windows.h>
main (int argc, char *p_argv[]){
char line[200],input_file_name[80], output_file_name[80], out_dir[80];
FILE *file_in, *file_out;
int numlines = 0;
int i = 0;
int n = 12;
int nn = 0;
int nnn = 1;
int n4 = 0;
int j = 0;
int n5 = 0;
{
if ( argc < 2 )
printf( "usage: filemod <infilename> <outfilename> \n");
strcpy( input_file_name, p_argv[1] );
strcpy( output_file_name, p_argv[2] );
/* copy filespec on command file to use as destination directory for
backup files */
nn = strlen(output_file_name);
nnn = nn - n + 2;
strncat(out_dir,output_file_name,nnn);
out_dir[0] = ' ';
out_dir[1] = ' ';
out_dir[2] = ' ';
/* read input file */
if (!(file_in = fopen (input_file_name,"r")))
{
printf("Error: unable to open input file %s\n", input_file_name);
return -1;
}

/* read the file to get the number of records /tags */
while (TRUE)
{
if (fgets (line,100,file_in) == NULL )
break;
/* count number of line in input file */
numlines++; }
/* open output file */
if (!(file_out = fopen (output_file_name,"w")))
{
printf("Error: unable to open input file %s\n", output_file_name);
return -1;
}
rewind( file_in );
/* the last 2 lines of the input file contain the names for the
current active archive and the associated index file; keep
these and discard the preceeding lines */
for( i = 0; i < numlines; i++)
{

```



```

        if (fgets (line,100,file_in) == NULL )
            break;
/* this line determines how many files are kept - to
keep the current active archive and its index this is -3;
to keep the last two and their indices this is -5*/
    if (i > numlines - 5)
    {
        if (strlen(line) > 0)
        {
            n4 = strlen(line);
            line[n4] = ' ';
            n4 = n4 -1;
            line[n4] = '\\0 ';
            strncat(line, out_dir,nnn);
            fprintf( file_out, "copy %s\\n",line);
            n5 = n4 + nnn +3;
            for (j=1; j < n5; j++)
                {line[j] = '\\0';}
            line[0] = '\\0';
        }
    }
}
fclose( file_in );
fclose( file_out );
exit(0);
}
}

```

Honeywell

Helping You Manage Your World