



NetApp™
Go further, faster

Technical Report

NetApp ILM for Oracle Database v1.1 with Oracle ILM Assistant

Padmanabhan Sadagopan, NetApp
September 2009 | TR-3781

TABLE OF CONTENTS

1	INTRODUCTION	3
1.1	INTENDED AUDIENCE AND ASSUMPTIONS	3
2	ARCHITECTURE	4
2.1	ORACLE DATABASE	4
2.2	ORACLE ILM ASSISTANT	5
3	INSTALLATION	5
3.1	PREREQUISITES	5
3.2	INSTALLING ORACLE DATABASE	5
3.3	INSTALLING ORACLE APPLICATION EXPRESS	6
3.4	CONFIGURING ORACLE APPLICATION EXPRESS USING EMBEDDED PL/SQL GATEWAY	6
3.5	INSTALLING ORACLE ILM ASSISTANT	7
4	NETAPP ILM INSTALLATION	7
5	NETAPP ILM CONFIGURATION	11
5.1	PREREQUISITES	11
5.2	CONFIGURING THE STORAGE SYSTEM	11
5.3	CONFIGURING THE DATABASE HOST	14
5.4	CONFIGURING THE NETAPP ILM EXTENSION	14
6	IMPLEMENTING ILM FOR YOUR ENVIRONMENT	17
6.1	DEFINING DATA CLASSES (USING ORACLE PARTITIONING OPTION)	17
6.2	CREATING STORAGE TIERS FOR THE DATA CLASSES	20
6.3	CREATING DATA ACCESS AND MIGRATION POLICIES	25
6.4	DEFINING AND ENFORCING COMPLIANCE POLICIES	26
6.5	VERIFYING THE RETENTION DATE	28
7	SUMMARY	28
	APPENDIXES	29
	APPENDIX A: REFERENCES	29
	APPENDIX B: SCRIPTS	29
	APPENDIX C: TROUBLESHOOTING TIPS	30
	APPENDIX D: NETAPP ILM SQL SCRIPT FILES	32
	ACKNOWLEDGEMENTS	33
	DISCLAIMER	33

1 INTRODUCTION

Information lifecycle management (ILM) is set of strategies and rules that define what happens to data during its lifetime. Government regulations and guidelines, such as Sarbanes-Oxley (SOX), HIPAA, and the European Data Privacy Directive in the European Union, are key driving forces in how and why data is being retained, as they are now requiring organizations to retain and control information for very long periods of time.

The basic idea of an ILM is to develop a hierarchy of storage based on storage cost and data activity, for instance:

- Retain highly accessed "current" data on a high-performance disk
- Archive low-activity historical data to cheaper disk

NetApp® ILM for Oracle® Database is an archival and compliance solution from NetApp that enables you to easily create tiered storage for data stored in the Oracle Database and apply granular controls to meet regulatory and corporate retention policies. With complete support for NetApp SnapLock® (write once, read many (WORM) storage), you can better comply with applicable government and industry regulations and keep data more secure, in unalterable and immediately accessible storage media.

The combination of Oracle ILM Assistant and NetApp ILM for Oracle Database enables DBAs to create partitions and tier storage across the NetApp unified storage architecture. Data is automatically migrated between tiers based upon service requirements and retention policies without navigating complex and heterogeneous storage environments.

Key objectives of this solution are to store vast quantities of data, for the lowest possible cost; and to meet the new regulatory requirements for data retention and protection.

1.1 INTENDED AUDIENCE AND ASSUMPTIONS

This document is for system administrators, database administrators, and personnel responsible for implementing IT infrastructure solutions.

This document assumes familiarity with Oracle Database concepts, PL/SQL, and Data ONTAP® features such as SnapLock and SnapMirror®.

2 ARCHITECTURE

Figure 1 depicts the high-level implementation of the NetApp ILM for Oracle Database.

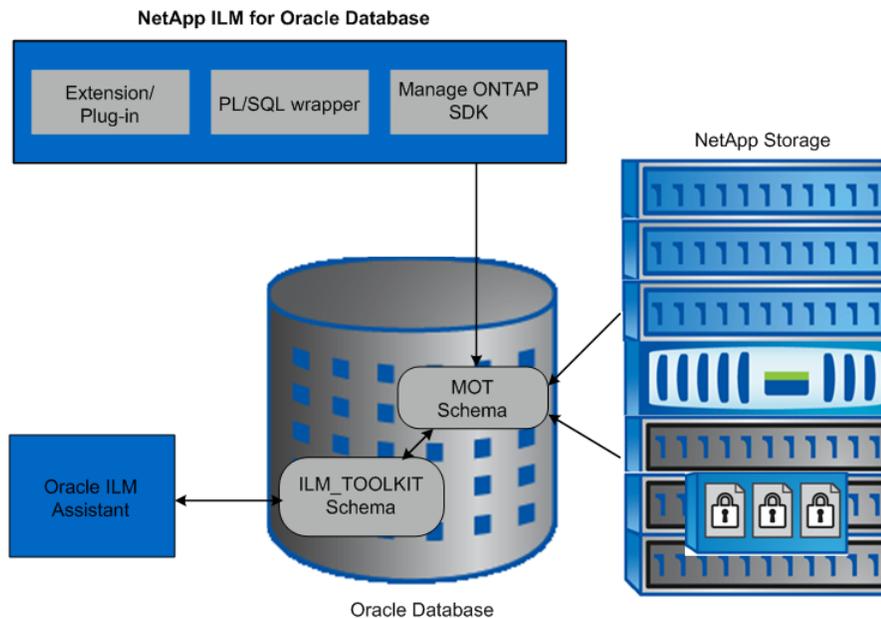


Figure 1) Architecture.

The three key components of NetApp ILM for Oracle Database are:

- **Extension/Plug-in:** This is a PL/SQL program that provides information such as directory paths of the NetApp ILM volume, special attributes of the NetApp ILM volume (either SnapLock or FlexVol®), and other details about the storage system to Oracle ILM Assistant. It creates the db objects, grants necessary privileges between the schema objects, and validates an added SnapLock (WORM) device.
- **PL/SQL Wrapper:** It is a combination of a Java™ program and PL/SQL package (called `DBMS_ONTAP`), taking advantage of Oracle JVM to load and publish the Java class in the database. This program interacts with the Manage ONTAP® SDK to acquire storage system information.
- **Manage ONTAP SDK:** It provides infrastructure to develop applications to monitor and manage the NetApp storage system. The Java API is used to develop this solution. `Manageontap.jar (v3.0)` is used for this solution.

2.1 ORACLE DATABASE

Oracle Database simplifies data management as it offers the advantage of storing different data formats in one location.

The following unique features make Oracle Database an ideal platform to implement an information lifecycle management policy, which makes it very easy to implement an ILM solution:

- Fine-grained: Managing data at individual row level
- Application transparency: Data classification is transparent
- Low cost: Uses low-cost storage to reduce costs
- Enforceable compliance policies

Oracle partitioning option allows data stored in a table to be physically segmented according to a data value. Partitioning data is completely transparent to the application, providing an easy way to distribute data across appropriate storage tiers depending on its value, while keeping the data online and stored on the most cost-effective device. For example, all of last year's orders can be stored in one or more partitions on a low-cost storage tier, while all the current year's orders are stored on the high-performance storage tier.

In Oracle Database 11g, ILM is enhanced to incorporate Oracle partitioning (and the all-important "exchange partition" syntax) to allow easy data movement at the partition level. As highly active data "ages out" and becomes low activity, ILM allows the DBA to easily "move" it away from the high-cost storage onto cheaper media. ILM also uses other Oracle Database 11g features such as Oracle Advanced Data compression and other partitioning methods.

2.2 ORACLE ILM ASSISTANT

The Oracle ILM Assistant is a GUI-based tool for managing your ILM environment. It provides the ability to create lifecycle definitions, which are assigned to tables in the database. Then based on the lifecycle policy, the ILM Assistant advises when to move, archive, or delete data. It also illustrates the storage requirements and cost savings associated with moving the data. Other capabilities include the ability to show how to partition a table based on a lifecycle definition and to simulate the events on a table if it were partitioned.

To assist with managing compliance issues, the ILM Assistant displays all VPD and FGA policies that are defined. It also displays and queries all audit records and generates and compares digital signatures.

Oracle ILM Assistant requires Oracle Database 9i or greater, and Oracle Application Express (formerly HTML DB) installed in the database where the tables to be managed by the ILM Assistant reside.

Initially, only tables partitioned on a date are eligible, and the ILM Assistant does not make any physical changes to the database. Instead it generates scripts so that you can perform the tasks when you are ready.

3 INSTALLATION

This section describes the procedure to install NetApp ILM v1.1 for Oracle Database with Oracle Database 11g R1 (11.1.0.6) or higher.

3.1 PREREQUISITES

GENERAL PREREQUISITES

- Check if any of the Oracle Database version 9i/10g/11g is installed with Oracle JVM option enabled.
- Check if Oracle Application Express (HTML DB) release 2.2 or higher is installed.
- Check if Oracle Application Express is configured with Oracle HTTP Server or embedded PL/SQL gateway to run it from the Web.
- NetApp storage system; check if Data ONTAP 7.2.x or higher is installed and SnapLock feature is available.
- SnapLock Compliance (SLC) requires all NetApp storage systems with the release of Data ONTAP 7.0.7, 7.1.3 or 7.2.5.

TEST SETUP

The following software components were used in our test setup:

- On the storage system, Data ONTAP 7.2.3 with SnapLock feature enabled
- Oracle Database 11g
Note: Any database version from 9i/10g/11g can be used. The NFS protocol is preferred as it enables using the SnapLock features.
- Oracle Application Express release 3.1.1

3.2 INSTALLING ORACLE DATABASE

Install Oracle Database 11g with Oracle JVM option enabled.

3.3 INSTALLING ORACLE APPLICATION EXPRESS

Oracle Application Express release 3.1.1 is automatically installed when installing Oracle Database 11g.

If you still need to install Oracle Application Express, do as follows; otherwise, go to section 2.4:

1. Log in as `sysdba` and execute the `apexins` script located in the `$ORACLE_HOME/apex/install` directory:

```
sqlplus / as sysdba
SQL> @apexins users users TEMP /i/
```

2. To change or set the password for the APEX user, run the `apxxepwd` script using the password as "passed." For example: `@apxxepwd.sql admin`.

3.4 CONFIGURING ORACLE APPLICATION EXPRESS USING EMBEDDED PL/SQL GATEWAY

Oracle Application Express should be configured with either Oracle HTTP Server or embedded PL/SQL gateway to enable it to run from the Web.

The embedded PL/SQL gateway provides the Oracle Database with a Web server and also the necessary infrastructure to create dynamic applications. The embedded PL/SQL gateway runs in the Oracle XML DB HTTP server in the Oracle Database and includes the core features of `mod_plsql`.

Although the embedded PL/SQL gateway installs with the Oracle Database 11g, you must configure it before you can use it with Oracle Application Express. To do this, run a configuration file and unlock the ANONYMOUS account.

```
$ cd /apex/
$ sqlplus / as sysdba
$ @apex_epg_config
SQL> alter user anonymous account unlock;
```

VERIFYING THE ORACLE XML DB HTTP SERVER PORT

```
SQL> SELECT DBMS_XDB.GETHTTPPORT FROM DUAL;
```

If the port number appears as 0, it indicates that the Oracle XML DB HTTP Server is disabled.

DISABLING THE ORACLE XML DB HTTP SERVER

Log in as `sysdba` and enter:

```
SQL> EXEC DBMS_XDB.SETHTTPPORT (<port>);
```

For example:

```
SQL> EXEC DBMS_XDB.SETHTTPPORT (0);
```

ENABLING THE ORACLE XML DB HTTP SERVER

Log in as `sysdba` and enter:

```
SQL> EXEC DBMS_XDB.SETHTTPPORT (<port >);
```

For example:

```
SQL> EXEC DBMS_XDB.SETHTTPPORT (8080);
```

Where `<port>` is the port number, which will be used in the URL to invoke the Oracle ILM Assistant.

3.5 INSTALLING ORACLE ILM ASSISTANT

To install the Oracle ILM Assistant, do as follows:

1. Download the Oracle ILM Assistant v1.3 or higher from OTN.

www.oracle.com/technology/software/deploy/ilm/index.html

Install and configure Oracle ILM Assistant by referring to the *ILMA Installation Guide*, listed in [Appendix A: References](#).

```
$sqlplus "sys/<password> as sysdba"  
SQL>@ilma_install apex_password tablespace connect  
SQL>Exit
```

Where:

- *apex_password* is the password for the Application Express database user.
- *tablespace* is the tablespace in which to store the ILM Assistant database objects created by this installation.
- *connect* is the Oracle Net connect string to the database. If this is a local install, use none or NONE.

2. Invoke the Oracle ILM Assistant using a URL similar to the following:

<http://:<hostname>:<port number>/apex/f?p=737677:1>

For example:

<http://ibmx345-sv103.iop.eng.netapp.com:9090/apex/f?p=737677:1>

3. Log in to the Oracle ILM Assistant:

Oracle ILM Assistant expects the database user created with the correct privilege to log in and set up the ILM for your environment.

4 NETAPP ILM INSTALLATION

1. Log in to the database server as "Oracle" or using the privileges set when installing the Oracle privileges who installed Oracle Database.
2. Change directory to the place wherever you want to unzip the files.
3. Download NetApp ILM for Oracle Database v1.1 from [NOW](#)™.
4. Create a separate directory for NetApp_ILM (for example, NetApp_ILM).
5. Unzip the download (NetApp_ILM.zip) file to the directory created in step 3.
6. Unzip will create a directory "NetApp_ILM" and download the files.
7. Change directory to the NetApp_ILM.
8. Make sure ORACLE_SID and ORACLE_HOME set correctly before starting the NetApp ILM for Oracle Database Installation.
9. Make sure you have the default tablespace and temporary tablespace name as you may be prompted for it during the installation.
10. Log in to Oracle Database as SYSDBA and execute:

```
run_install.sql
```

The installation is a two-step process. First, NetApp ILM for Oracle Database is installed, and then, NetApp ILM Extension is installed.

NetApp ILM Extension requires Oracle ILM Assistant (ILMA) runtime PL/SQL package, so installation of an extension can only be performed after a successful ILM Assistant installation.

```
oracle@ibmx345-sv101 NETAPP_ILM_v11]$ sqlplus "/as sysdba"
```

```
SQL*Plus: Release 11.1.0.6.0 - Production on Thu Mar 12 10:38:59 2009  
Copyright (c) 1982, 2007, Oracle. All rights reserved.
```

```
Connected to:
```

```
Oracle Database 11g Enterprise Edition Release 11.1.0.6.0 - Production  
With the Partitioning, OLAP, Data Mining and Real Application Testing  
options
```

```
SQL> @run_install
```

```
*****
```

```
NetApp ILM for Oracle Database v1.1 installation
```

```
*****
```

```
Step 1 - MOT schema creation
```

```
This step will create a new schema called MOT.
```

```
You will be asked to provide the MOT schema password,
```

```
DEFAULT tablespace and TEMPORAY tablespace detail.
```

```
*****
```

```
Press Enter to Continue with the installation
```

```
Enter MOT Schema Password : <<Enter the password – it will not be displayed>>
```

```
Enter default tablespace for MOT :MOT_TBS <<Provide the default tablespace name  
>>
```

```
Enter temporary tablespace for MOT :TEMP <<Provide the temporary tablespace>>
```

```
PL/SQL procedure successfully completed.
```

```
*****
```

```
Step 2 - Load the manageontap.jar and DBMSWrapper.class into MOT  
schema.
```

```
Note down the directory name where you unzipped these files
```

```
You must provide the fully qualified path name of these files
```

```
For example: </u01/NetApp_ILM/source_home>/manageontap.jar
```

```
For example: </u01/NetApp_ILM/source_home>/DBMSWrapper.class
```

```
*****
```

```
Press Enter to Continue with the installation
```

```
Enter fully qualified name of manageontap.jar
```

```
:/u03/UAT/NETAPP_ILM_v11/manageontap.jar
```

Enter fully qualified name of DBMSWrapper.class
:/u03/UAT/NETAPP_ILM_v11/DBMSWrapper.class

PL/SQL procedure successfully completed.

.....

DBMS_JAVA.LONGNAME(OBJECT_NAME)	OBJECT_TYPE	STATUS
-----	-----	-----
netapp/manage/NaSshHost	JAVA CLASS	INVALID
netapp/manage/NaSshHost\$1	JAVA CLASS	INVALID

You can IGNORE the above INVALID status warning....

It is OK to have netapp/manage/NaSshHost and netapp/manage/NaSshHost\$1 objects are in status INVALID

"Press Enter to Continue with the installation"

Step 3 - Create DBMS_ONTAP package under MOT schema

Trigger created.

Trigger altered.

Package created.

Package body created.

No errors.

Package body altered.

OBJECT_NAME	OBJECT_TYPE	STATUS
-----	-----	-----
DBMS_ONTAP	PACKAGE	VALID
DBMS_ONTAP	PACKAGE BODY	VALID

"Press Enter to Continue with the installation"

NetApp ILM for Oracle Database v1.1 Installation Completed Successfully

NetApp ILM Extension Installation

Oracle ILM Assistant v1.3 must be installed before proceeding further.

Press Enter to Continue with the installation

PL/SQL procedure successfully completed.

Grant succeeded.

Grant succeeded.

No errors.

SP2-0272: escape character cannot be alphanumeric or whitespace

Type body created.

No errors.

Type body created.

NetApp ILM Extension Installation Completed

11. Check the netapp_ilm_install.log file created in the current directory for any errors.

5 NETAPP ILM CONFIGURATION

This section describes the procedure to configure NetApp ILM for Oracle Database.

5.1 PREREQUISITES

- Make sure that the Oracle ILM Assistant is installed correctly and that you are able to invoke the URL similar to the following:

`http://<hostname>:<port number>/apex/f?p=737677:1`

For example:

<http://ibmx345-svl03.iop.eng.netapp.com:9090/apex/f?p=737677:1>

- Oracle Application Express should be configured with Oracle HTTP Server or embedded PL/SQL gateway to run it from the Web. For the configuration procedure, see http://download.oracle.com/docs/cd/B32472_01/doc/install.300/b32468/toc.htm.

5.2 CONFIGURING THE STORAGE SYSTEM

SnapLock is NetApp's implementation of high-performance disk-based magnetic WORM storage.

The primary objective of this Data ONTAP feature is to provide secure, storage-enforced data retention functionality via open file protocols such as CIFS and NFS while leveraging existing NetApp technologies to the greatest degree possible.

SnapLock is available in two versions: SnapLock Compliance and SnapLock Enterprise. Make sure that your version of Data ONTAP supports both versions.

You can use the NetApp FilerView® GUI—a Web based tool, or storage system command line to perform the following steps:

1. License the SnapLock feature.

Check that the SnapLock Enterprise or SnapLock Compliance license is added correctly.

Use the license command at the storage system to license the correct SnapLock version.

```
Storage System > license add <snaplock_license_code>
```

2. Create SnapLock specific aggregates:

SnapLock is activated at the aggregate level. Therefore, it can be turned on per traditional volume, but FlexVol volumes inherit SnapLock attributes from their aggregate. If SnapLock is enabled on an aggregate, any FlexVol volumes created on the aggregate will be SnapLock enabled.

```
FilerView > Aggregates > Add > <aggregate name>
```

Or

```
Storage System> aggr create SL_ENT_AGGR -L enterprise 9
```

Following is a sample output from FilerView:

Name:	SL_ENT_AGGR		
Type:	Aggregate	Root Aggregate?	-
Status:	online,raid_dp	Raid Size:	16
Used Capacity:	144 KB	Checksums:	block
% Used:	0%		
Total Capacity:	397 GB	Number of Disks:	9
Number of Files:	98		
Max Files:	31.1 k	Double Parity?	✓
Snaplock type:	enterprise		

3. Check the status of the aggregate created previously:

```
Storage System> aggr status
```

```
Aggr State          atus          Options
SL_ENT_AGGR online   raid_dp, aggr snaplock_enterprise
aggr0 online        aid_dp, aggr  root
```

4. Create a FlexVol volume in the SnapLock aggregate using one of the following methods:

1. Enter:

```
FilerView > Volumes > Add > <volume_name>
```

Name:	SL_ENT_2YEARS		
Type:	Flexible	Root Volume?	-
FlexClone ?	-	Containing Aggregate:	SL_ENT_AGGR
Status:	online,raid_dp		
Used Capacity:	96 KB	Space Guarantee	volume
% Used:	0%	Language:	undefined
Total Capacity:	1.6 GB	Total Size:	2 GB
Number of Files:	102	Max Directory Size:	81.8 MB
Max Files:	77.6 k		
Snaplock type:	enterprise	Snaplock minimum period:?	0y
Snaplock maximum period:?	2y	Snaplock default period:?	min

2. Enter:

```
Storage System> vol create SL_ENT_2YEARS SL_ENT_AGGR 2G
```

```
Creation of volume 'SL_ENT_2YEARS' with size 2g on containing aggregate
'SL_ENT_AGGR' has completed.
```

```
Storage System> vol status
```

```
Volume State  Status          Options
vol0 online   raid_dp, flex  root
SL_ENT_2YEARS online   raid_dp, flex  no_atime_update=on,
                                     snaplock_enterprise
```

5. Set a retention period for the SnapLock volume as follows:

Each SnapLock traditional volume or FlexVol volume can have the following three retention periods, which are set using the `vol options` command:

- **Maximum:** This is the longest retention period that Data ONTAP allows to be set on a file, regardless of what the application attempts to overload the last access time attribute for a file with. This protects against situations such as an application being set with a longer retention period than the business requires.
- **Minimum:** This is the shortest retention period that Data ONTAP allows to be set on a file, regardless of what the application attempts to overload the last access time attribute for a file with. This protects against files having too short a retention period to meet regulatory compliance.
- **Default:** This is the retention period for a record or file if the last access time attribute was not overloaded prior to putting a file into a SnapLock WORM state. Default retention dates can be any value between the minimum and maximum retention period settings.

By default, the maximum retention period is 30 years, the minimum retention period is zero years, and the default is set to equal the minimum retention period.

To change any of the retention values, use the `vol options` command or enter:

```
FilerView > Volumes > Manage > <volume_name> > Modify
Storage System> vol options volume_name snaplock_minimum_period 0y
Storage System> vol options volume_name snaplock_maximum_period 2y
```

6. Export NFS for the SnapLock volume created in the previous steps by entering:

```
FilerView > NFS > Manage Exports > <volume_name>
```

Path	Options
/vol/SL_ENT_2YEARS	Anonymous User ID=0 Read-Write Access (All Hosts) Security (sys)

7. Initialize the compliance clock:

The compliance clock is a software-driven secure clock that runs independently of the regular system clock and persistently stores state on disk. It provides a secure time base, and all WORM delete operations are based on this clock. It can only be initialized once and is tamper resistant.

Based on the clock, WORM files and SnapLock compliant volumes can be deleted once their retention period has expired.

To initialize the compliance clock, enter:

```
date -c initialize
```

Use this option only if one you have installed a SnapLock Compliance or SnapLock Enterprise license. This option initializes the compliance clock to the current value of the system clock.

Note: Make sure that system clock is appropriately set before running the above command, as the compliance clock may only be set once; there is no mechanism for resetting the compliance clock.

For more information, see the NetApp storage system specific documentation:

<http://now.netapp.com/NOW/knowledge/docs/ontap/rel7251/html/ontap/onlinebk/7snaplo2.htm>.

NetApp FilerView can be invoked by using the URL similar to the following:

```
http://<storage system name>/na_admin
```

For example:

```
http://10.61.162.12/na_admin/
```

5.3 CONFIGURING THE DATABASE HOST

Make sure you have root privileges to create directories and the operating system user such as "oracle" with group "dba" with the right privilege which you used to create the Oracle Database.

The following example describes how to create a physical directory in the database host and NFS mount them in Linux®. Make appropriate changes when installing on Solaris™ and UNIX®.

1. Create a physical directory on the database hosts to mount the volume created in the storage system. Change the ownership of the directory to the OS user (in this case oracle).

For example:

```
# mkdir -p /u02/spectrevol/sl_2to20
# chown -R oracle:dba /u02/spectrevol/sl_2to20
# mount /u02/spectrevol/sl_2to20
# ls -ld /u02/spectrevol/sl_2to20
drwxr-xr-x 3 root root 4096 Jul 21 2008 /u02/spectrevol/sl_2to20
# chown -R oracle:dba sl_2to20/
chown: changing ownership of `sl_2to20/.snapshot': Read-only file system
# ls -ltr
total 4
drwxr-xr-x 3 oracle dba 4096 Jul 21 2008 sl_2to20
```

2. Use the `mount` command to mount the volume to the physical directory or add entries to the `/etc/fstab` file (in the Linux OS).

For example:

```
10.61.162.12:/vol/sl_2to20 /u02/spectrevol/sl_2to20 nfs
hard,nointr,proto=tcp,suid,rw,bg,vers=3,rsize=32768,wsiz=32768,actimeo=0,ti
meo=600
```

3. Verify the NFS volumes mounted from the hosts.

For example:

```
$mount
10.61.162.12:/vol/sl_2to20 on /u02/spectrevol/sl_2to20 type nfs
(rw,hard,nointr,proto=tcp,bg,nfsvers=3,rsize=32768,wsiz=32768,actimeo=0,tim
eo=600,addr=10.61.162.12)
```

5.4 CONFIGURING THE NETAPP ILM EXTENSION

This section describes how to create a database user and assign credentials to this user to access the Oracle ILM Assistant.

To configure the NetApp ILM Extension, every SnapLock volume created and NFS mounted should be added as a file system to the Oracle ILM Assistant.

1. Log in as `sysdba` and grant execute privilege on `DBMS_ONTAP` to the Oracle ILM Assistant.

For example:

```
SQL> grant execute on mot.dbms_ontap to ilm_toolkit ;
```

2. Create a database user (for example, `netapp_ilm_demo`) to log into the Oracle ILM Assistant GUI.

For example:

```
SQL> CREATE USER "NETAPP_ILM_DEMO" PROFILE "DEFAULT" IDENTIFIED BY "*****"
DEFAULT TABLESPACE "NETAPP_ILM_DEMO" TEMPORARY TABLESPACE "TEMP" ACCOUNT
UNLOCK
GRANT "CONNECT" TO "NETAPP_ILM_DEMO"
```

3. Assign privileges to the database user created.

Log in as sysdba and run the `grant2run_ilma.sql` script. Pass the database user name created above as an argument.

For example:

```
SQL > @grant2run_ilma.sql <netapp_ilm_demo>
```

4. Add the file system to the Oracle ILM Assistant.

Every SnapLock volume created on the storage system that is mounted on the database host should be added as a file system to Oracle ILM Assistant with the required credentials.

For every file system added, the Oracle ILM Assistant expects a directory object created with the directory path same as the physical path that is mounted to the SnapLock volume. This step is automated so that the directory object appears in the output when running the `netapp_add_fs.sql` script.

Log in as MOT user or sysdba privilege and run the `ilma_netapp_add.sql` script.

For example:

```
SQL> @netapp_add_fs
Enter Storage System Name      :10.61.162.12
Enter Volume/Device Path      :/vol/sl_2to20
Enter Storage System Username  :root
Enter Storage System Password  :
Enter File System Display Name  :SL_ENTERPRISE_2TO20
Enter File System Unique ID    :4444
Enter Physical path from the Host :/u02/spectrevol/sl_2to20
PL/SQL procedure successfully completed.
Device ID Display Name
-----
      4444 SL_ENTERPRISE_2TO20
DIRECTORY_NAME      DIRECTORY_PATH
-----
NETAPP_SNAPE_4444   /u02/spectrevol/sl_2to20
SQL>
```

Where:

File system display name: The name of the file system displayed in Oracle ILM Assistant.

File system unique ID: The unique ID to identify the file system.

Physical path from the host: The physical directory created and mounted to the SnapLock volume. Oracle ILM Assistant expects data files for this volume to be created in this directory.

The device ID and display name output shows that volume is successfully added as a file system to Oracle ILM Assistant. It also internally creates a directory structure where the data files for this volume are to be created.

5. Create a tablespace template for each file system added.

At least one tablespace must exist on the managed file system to be usable by the ILM Assistant. The purpose of the tablespace is to be a template for new tablespaces targeted to receive migrated database partitions. The new tablespaces inherit every attribute of the template, except the file name, which is fabricated by the ILM Assistant.

Template tablespaces must not contain volume paths associated with multiple file systems. If this occurs, the ILM Assistant will ignore the tablespace and possibly the entire file system.

Make sure that the data file location is the physical directory created and mounted as shown in step 4.

For example:

```
SQL> CREATE SMALLFILE TABLESPACE "SL_2TO20_TPLATE" DATAFILE
'/u02/spectrevol/sl_2to20/sl_2to20_tplate01.dbf' SIZE 50M REUSE LOGGING
EXTENT MANAGEMENT LOCAL SEGMENT SPACE MANAGEMENT AUTO DEFAULT NOCOMPRESS
```

6. Verify whether you can log into Oracle ILM Assistant GUI using the above URL and use the database user created in previous steps.

For example: <http://ibmx345-svl03.iop.eng.netapp.com:9090/apex/f?p=737677:1>

Database username: <netapp_ilm_demo>

Database password: <XXXXXX>

7. Set the retention date.

Note: The retention date on a particular database is set when implementing the ILM for your environment as described in Section 6.

A retention date is the date after which Data ONTAP allows deletion of a WORM file on a SnapLock volume. The retention date for WORM records on a SnapLock volume is stored in the last access timestamp of the record file metadata.

To set a retention date for a WORM record, the application must explicitly set the file last access time to the desired retention date before setting the file to read-only and engaging the WORM commit operation. Once committed to the WORM state, the access time of the file is immutable, with the exception of extending the record retention period.

Oracle ILM Assistant will provide the script to set the retention date for a particular tablespace, based on the event detection. That script internally calls the NetApp ILM Extension to set the retention date correctly on the SnapLock volumes on the NetApp storage system.

Sample script:

```
begin
ilm_toolkit.ilm_toolkit.set_retention ( 'SL_2TO20_TPLATE4',
to_date('10/11/2015', 'MM/DD/YYYY') );
end;
/
```

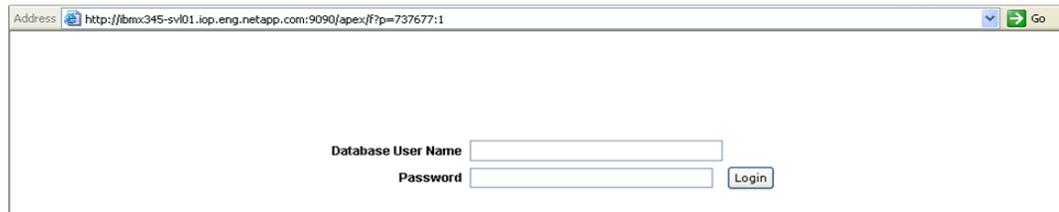
6 IMPLEMENTING ILM FOR YOUR ENVIRONMENT

Following are the steps to implement ILM:

1. Define data classes (see [section 6.1](#)).
2. Create storage tiers and define lifecycle for data classes (see [section 6.2](#)).
3. Create data access and migration policies (see [section 6.3](#)).
4. Define and enforce compliance policies (see [section 6.4](#)).

To implement ILM, log into Oracle ILM Assistant (OILMA) using a URL as follows:

`http://<host_name>:<port#>/apex/f?p=737677:1`



Lifecycle Setup

If this is the first time that you have used the ILM Assistant, then it is here where you specify exactly how the data is to be managed by the ILM Assistant. You must complete the following steps before the ILM Assistant is able to give advice on data placement.

1. Define the [logical storage tiers](#).
2. Define the [lifecycle definitions](#).
3. [Select tables](#) to be managed by the lifecycle definitions.

6.1 DEFINING DATA CLASSES (USING ORACLE PARTITIONING OPTION)

Data can be classified based on frequency of usage. Storage tiers can then be created for the identified data classes. During the lifetime of the data, it will then migrate between the data classes, and access to it will be controlled. Eventually, the data can be archived or remain in the database forever.

Defining the data class helps you separate the data physically. Once the data has been classified, you can map them with a data attribute. Oracle table partition enables you to separate data by a data attribute, as Oracle partitions are transparent to the applications.

For example, a sample table named `SL_EMP` with partition by date range is created as follows:

```
CREATE TABLE "ILMTEST"."SL_EMP"  
( "NAME" VARCHAR2(25) NOT NULL , "DATE" DATE NOT NULL , "SALARY" NUMBER)  
TABLESPACE "NETAPP_TEST"  
PARTITION BY RANGE ("DATE")  
(PARTITION "SL_EMP_P1" VALUES LESS THAN (TO_DATE(' 1995-01-01 00:00:00',  
'SYYYY-MM-DD HH24:MI:SS')) TABLESPACE "NETAPP_TEST"  
PARTITION "SL_EMP_P2" VALUES LESS THAN (TO_DATE(' 2009-01-01 00:00:00',  
'SYYYY-MM-DD HH24:MI:SS')) TABLESPACE "NETAPP_TEST"
```

The SL_EMP Oracle table created above with partition by date, qualifies as a table to be migrated from one partition to other partition when it falls within the date range.

You can check the details of this table by clicking Life Cycle Setup > Life Cycle Table.

ORACLE ILM Assistant [Skip navigation](#) [Print](#) [Logout](#)

User: ILMTEST **Lifecycle Setup** Lifecycle Management Reports Compliance & Security Help

[Logical Storage Tiers](#) [Lifecycle Definitions](#) [Lifecycle Tables](#) [Preferences](#)

Lifecycle Tables

Shows all tables within the view of the ILM Assistant. Candidate tables can be managed by the ILM Assistant. For tables having no partitioning, storage cost savings and storage tier migration can be modeled using a simulated partitioning strategy. Last table list refresh: 30-Oct-2008 16:59 ⓘ

[Tasks & Options](#)
[Refresh table list](#)
[Show filter options](#)

Show All Managed Tables Simulated Tables Candidate Tables Ineligible Tables

	Table Owner	Table Name	Storage Size (GB) ▼	Data Reads	Data Writes	Lifecycle Definition	Lifecycle Status	Table Partitioning	Cost Savings	Placement Map	
	ILM_DEMO	ALLORDERS	453	0	0	General Orders	Simulated	Simulated	\$20,354	✓	🗑️
	ILM_DEMO	PART_ORDERS	0	256	21	General Orders	Managed	Implemented	\$18	✓	🗑️
	ILM_DEMO	CC_TRAN	0	176	17	Audit Data	Managed	Implemented	\$17	✓	🗑️
	ILMTEST	CP_CC_TRAN	0	0	0	my_life_cycle	Managed	Implemented	\$21	✓	🗑️
	ILM_DEMO	USER_AUDIT	0	64	12	Audit Data	Managed	Implemented	\$12	✓	🗑️
	ILMTEST	SL_EMP	0	17	5	my_life_cycle	Managed	Implemented	\$6	✓	🗑️
	ILMTEST	CP_SALGRADE	0	1	0		Candidate	Unavailable			
	ILMTEST	CP_DEPT	0	1	0		Candidate	Unavailable			
	ILMTEST	CP_BONUS	0	1	0		Candidate	Unavailable			
	ILMTEST	CP_ALLORDERS	0	1	0		Candidate	Available			

row(s) 1 - 10 of 14 ▶

Figure 2) Lifecycle tables.

The Lifecycle Tables page (Figure 2) identifies tables that can be managed by the ILM Assistant, and allows mapping these tables to a lifecycle definition. A database may contain many tables, only some of which you can consider as candidates for ILM. A table is automatically eligible if it is range partitioned on a date column, and then when it is associated with a lifecycle definition, the ILM Assistant can manage its data. For tables having no partitioning, storage cost savings and storage tier migration can be modeled using a simulated partitioning strategy. This page is the final step in mapping a table to a lifecycle definition so the ILM Assistant can report when it is time to move the data.

Though the Lifecycle Tables page view shows all accessible tables, the ILM Assistant might not be able to manage every table. In this case, the table will be marked as ineligible, and a link will be provided to explain the exception. Some examples of ineligible tables:

- Tables partitioned on non-date columns
- Tables partitioned using a partition type other than range

For tables that are not partitioned, the ILM Assistant now permits an alternative placement approach that does not require partitioning on a date. If a table is small or has incompatible attributes, the user may manage the table by allowing it to be placed as a whole object. Just as one does with a partition candidate, the table will be associated with a lifecycle. However, instead of requiring the table to have a date column, the Assistant will prompt the user for a specific date value. That date value will then be used by the Assistant to evaluate correct placement.

The display for lifecycle tables can be customized to display managed, simulated, candidate, and ineligible tables.

To begin lifecycle management of a table:

1. Click Candidate.

If the table is already partitioned, the [Manage Lifecycle Table](#) page appears. In this page, you can set up an association with a lifecycle definition, as shown in step 2.

If the candidate is not yet partitioned, the [Partition Simulation](#) page appears. In this page, you can set up a full simulation, as shown in step 3.

2. In [Manage Lifecycle Table](#) page, select a lifecycle for the table and preview the storage cost benefits and the projected lifecycle stage effect on the table.

Review the changes and click OK to return to the Lifecycle Tables page or click Cancel to choose another lifecycle definition.

3. In the [Partition Simulation](#) page, select a lifecycle for the table and preview the storage cost benefits and the projected lifecycle stage effect on the table.

Review the changes. Click OK to return to the Lifecycle Tables page or click Cancel to choose another lifecycle definition.

The table is now eligible for full lifecycle management in simulation mode.

The difference between a managed table and a simulated table is that a managed table contains actual partitions, and a simulated table contains fake partitioning data. All reports and event detection work with both types of lifecycle tables. However, any table upon which partitioning is being simulated will only be seen as being partitioned from within the ILM Assistant. All other tools will continue to see it as a non-partitioned table.

6.2 CREATING STORAGE TIERS FOR THE DATA CLASSES

You can create separate storage tiers for high-performance, low-cost, and archival purposes.

High-performance storage tier can be using FC disks and low-cost storage tier may be using SATA disk and the same for online/offline archival disk.

A logical storage tier is a name given to a logical group of storage devices, and typically all disks of the same type will be identified by that name. For example, the group called High Performance could refer to all the high-performance disks. Any number of logical storage tiers may be defined, and the devices are identified by the assigned tablespaces that reside upon them.

Figure 3 shows an example of a logical storage tier.

The screenshot shows the Oracle ILM Assistant interface. At the top, the user is identified as 'User: ILMTEST'. The navigation bar includes 'Lifecycle Setup', 'Lifecycle Management', 'Reports', 'Compliance & Security', and 'Help'. Under 'Lifecycle Setup', 'Logical Storage Tiers' is selected. Below the navigation bar, there is a section titled 'Logical Storage Tiers' with a description: 'A logical storage tier is the name given to a pool of tablespaces that all reside on the same class of storage device.' To the right of this description is a 'Tasks & Options' box with links for 'New storage tier' and 'Show filter options'. Below this is a table with the following data:

	Storage Tier Name	File System		Description	Cost Per GB ▼	In Use	
	NetApp_High_Cost	Conventional		NetApp FC disk based high cost tier	\$125.00	✓	
	SL_ENT_5YEAR	SL_ENT_MAX5YEAR		Archive data older than 5 years (SnapLock Enterprise)	\$70.00	✓	
	SNAPLOCK_ENTERPRISE_0-30Y	PAD_ARCHIVE			\$50.00	✓	
	High Performance	Conventional		Very high performance, used for current, active data	\$50.00	✓	
	NetApp_Low_Cost	Conventional		SATA disk based Low cost tier	\$50.00	✓	
	Low Cost	Conventional		Lower Costs ATA discs, used for older data	\$12.00	✓	
	Online Archive	Conventional		Used for data older than 2 years	\$4.00	✓	

At the bottom right of the table area, it says '1 - 7'.

Figure 3) Logical storage tiers.

Cost per GB

The Cost per GB value must be greater than zero. This value is used by the ILM Assistant to project storage costs when data is mapped to the tier. It is recommended to enter a value that represents a reasonably accurate cost of storing data on the tier. This would include the physical purchase price of a device; however, you may also want to consider other associated costs such as maintenance and running costs.

File System

A storage tier must be associated with a particular file system. By default, the Oracle Database understands the conventional system that is implemented by the operating system. Additional file systems can be defined within the ILM Assistant to support special devices that save energy or protect data. If multiple file system types are found, then the user may associate the storage with the desired file system; otherwise, the ILM Assistant assumes all devices are conventional systems.

Data can be classified based on frequency of usage. Storage tiers can then be created for the identified data classes. During the lifetime of the data, it will then migrate between the data classes, and access to it will be controlled. Eventually, the data can be archived or remain in the database forever.

When creating a logical storage tier, every SnapLock volume created as a file system will be available as a list of values to assign to a particular storage tier.

A storage tier can only be associated with one type of file system. The file system also determines which tablespace objects are assigned to the storage tier. There can be more than one file system available for the storage tier usage.

The conventional system is the default file system used by the Oracle Database. It permits standard read or write operations.

The other file systems that are created as part of NetApp specific extension configuration are initially read-write file systems. After the migration of the data to be protected, the file system becomes read-only.

For the more details, see the section “Logical Storage Tier” in the *Oracle ILM Assistant User Guide*.

CREATING A STORAGE TIER

To create a storage tier, do as follows:

1. On the Oracle ILM page, click Lifecycle Setup > Logical Storage Tiers.
2. Under Tasks & Options, click New storage tier (see Figure 3).

The Manage Logical Storage Tier page appears.

The screenshot shows the 'Manage Logical Storage Tier' page in the Oracle ILM Assistant. The page title is 'Manage Logical Storage Tier' and it includes 'Cancel', 'Apply', and 'OK' buttons. The main content area has the following fields:

- Name:** NetApp_High_Cost
- Description:** NetApp FC disk based high cost tier
- Cost per GB:** 125
- File System:** Conventional (with a help icon and a note: 'The default file system used by an Oracle database')

Below the configuration fields is the 'Assigned Tablespaces' section, which includes a 'Go' button and a table of assigned tablespaces.

Tablespace Name	Lifecycle Status	Tablespace Status	Read-Write Preferred	Read-Only Preferred	Secondary Tablespace
NETAPP_HIGH_COST	Managed	ONLINE	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

The table shows 1 of 1 tablespace assigned.

Figure 4) Manage logical storage tier.

3. In the Manage Logical Storage Tier page, enter a name, description, cost per GB, and preferred tablespaces for the storage tier. Click OK to apply the changes.
4. Click Create and return to the Logical Storage Tiers page.

DELETING A STORAGE TIER

To delete a logical storage tier, in the Logical Storage Tiers page, click the Delete icon . Click OK when prompted to confirm deletion.

MODIFYING A STORAGE TIER

Figure 4 shows the logical storage tier named “NETAPP_HIGH_COST” with its defined name, description, and cost per GB.

Each storage tier has a set of assigned tablespaces. A tablespace can only be assigned one type of tablespace:

- **Read-Write Preferred:** If read-write data is migrated to the tier, then the read-write preferred tablespace is required. If the storage tier accepts read-only data, a read-only preferred tablespace can be assigned. In addition to the preferred tablespaces, one or more secondary tablespaces may be assigned to the tier.
- **Read-Only Preferred:** A read-only preferred tablespace may not necessarily receive an object. If the tablespace already contains an object or if the tablespace resides on a protected file system device, the

ILM Assistant indicates that a new tablespace will be created using attribute settings from the preferred read-only tablespace. If the recommendation is implemented, the newly cloned tablespace will automatically become a secondary tablespace for the storage tier.

- Secondary Tablespace:** Secondary tablespaces are typically located in the same location as the read-write preferred tablespace for the storage tier. As the ILM Assistant only supports a single preferred tablespace, any read-write data that must reside on the tier would generate a data movement event to relocate the data to the read-write preferred tablespace. To avoid unnecessary move events, the ILM Assistant will allow existing data to remain on a secondary tablespace for the storage tier. The ILM Assistant will not presently use a secondary tablespace as the target location for data movement.

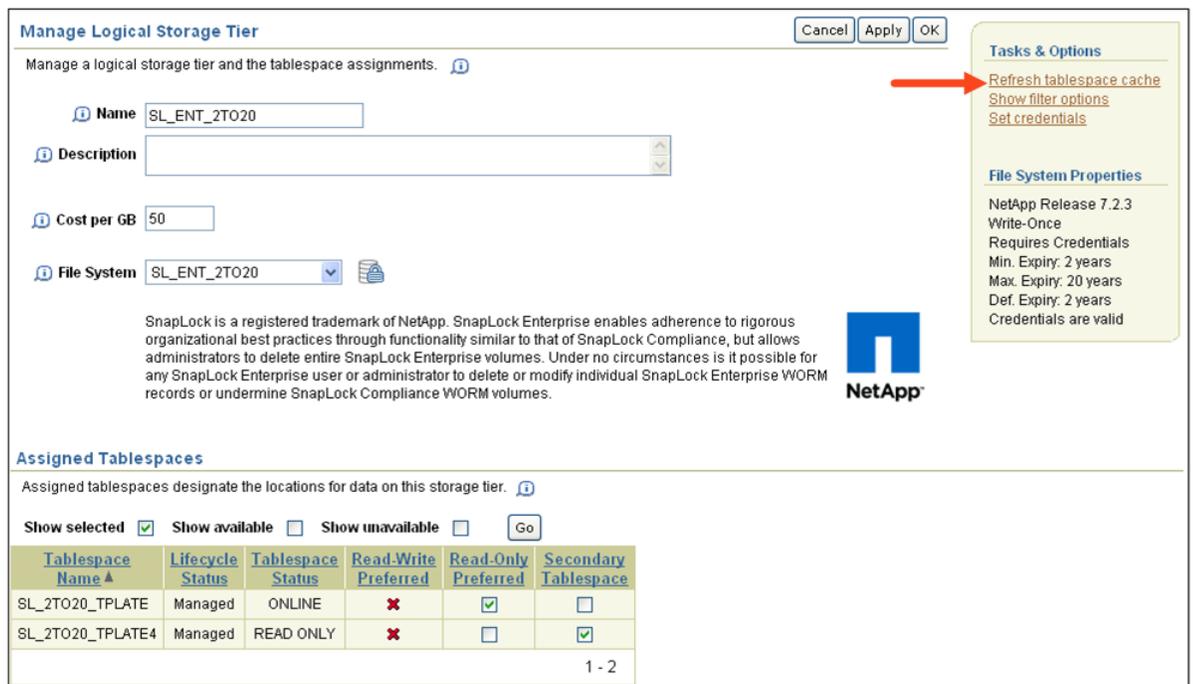
To modify storage tier attributes and tablespace assignment, do as follows:

- In the Logical Storage Tiers page, click the Edit icon . The Manage Logical Storage Tier page appears (Figure 4).
- Make the required changes and click OK.

REFRESHING THE TABLESPACE

To handle large installations, an internal cache of known tablespaces is maintained by the ILM Assistant. It is simply a private copy of tablespace information to provide better performance within the Assistant. Typically, the ILM Assistant will keep the cache up to date; however, it is possible that it may miss newly created tablespaces, especially if they are created outside of the ILM Assistant domain.

In this case, to manually fresh and update the cache, under Tasks & Options, click the Refresh tablespace cache (see Figure 5).



Manage Logical Storage Tier Cancel Apply OK

Manage a logical storage tier and the tablespace assignments. ⓘ

Name

Description

Cost per GB

File System 

SnapLock is a registered trademark of NetApp. SnapLock Enterprise enables adherence to rigorous organizational best practices through functionality similar to that of SnapLock Compliance, but allows administrators to delete entire SnapLock Enterprise volumes. Under no circumstances is it possible for any SnapLock Enterprise user or administrator to delete or modify individual SnapLock Enterprise WORM records or undermine SnapLock Compliance WORM volumes.

NetApp

Tasks & Options

- [Refresh tablespace cache](#)
- [Show filter options](#)
- [Set credentials](#)

File System Properties

- NetApp Release 7.2.3
- Write-Once
- Requires Credentials
- Min. Expiry: 2 years
- Max. Expiry: 20 years
- Def. Expiry: 2 years
- Credentials are valid

Assigned Tablespaces

Assigned tablespaces designate the locations for data on this storage tier. ⓘ

Show selected Show available Show unavailable Go

Tablespace Name ▲	Lifecycle Status	Tablespace Status	Read-Write Preferred	Read-Only Preferred	Secondary Tablespace
SL_2T020_TPLATE	Managed	ONLINE	✘	<input checked="" type="checkbox"/>	<input type="checkbox"/>
SL_2T020_TPLATE4	Managed	READ ONLY	✘	<input type="checkbox"/>	<input checked="" type="checkbox"/>

1 - 2

Figure 5) Refresh tablespace cache.

DEFINING THE LIFECYCLE

To define the lifecycle, do as follows:

1. On the Oracle ILM page, click Lifecycle Setup > Lifecycle Definitions.
The Manage Lifecycle Definition page appears.

User: ILMTEST | Lifecycle Setup | Lifecycle Management | Reports | Compliance & Security | Help

Logical Storage Tiers | **Lifecycle Definitions** | Lifecycle Tables | Preferences

Manage Lifecycle Definition [Cancel] [Apply] [OK]

Manage a lifecycle definition. ⓘ

Lifecycle Name:

Description:

Fiscal Start for this Lifecycle January 1

Fiscal Quarters
 Q1: January 1 - March 31
 Q2: April 1 - June 30
 Q3: July 1 - September 30
 Q4: October 1 - December 31

Total Retention 10 Years 1 Month 10 Days

Stage Details | Graphical View [New Stage]

Each stage describes a destination and how long data will reside there during its lifetime. ⓘ

	Stage Name	Description	Action	Tier Name	Attributes	Retention Period	Stage Start Date	Stage End Date
Current Stage	my_high_perf_stage	current	Remain Online	my_high_perf	None	Current 10 Days	30-Sep-2008	09-Oct-2008
	my_lowcost_stage	second stage	Remain Online	Low Cost	None	1 Month	31-Aug-2008	29-Sep-2008
	my_final_sl_stage		Remain Online	SL_ENT_2T020	Read-Only	10 Years	31-Aug-1998	30-Aug-2008
Final Stage	Final_purge_stage	third and final stage	Purge		None			30-Aug-1998

Figure 6) Manage lifecycle definition.

Make sure that the `SL_EMP` table created in section 5.1 is assigned to the Lifecycle Definitions created in this step.

2. Enter a lifecycle name and description and click OK.
3. You might have to create stages for Lifecycle definition created in the previous step by clicking on 'New Stage'.

DEFINING THE STAGES

A lifecycle definition comprises a number of stages that describe what happens to data during its lifetime. Lifecycle stages are initially created in reverse time order, that is, working backwards in time from the current date.

Therefore, the current stage for new data is always created on the **Create Lifecycle Definition** page and then additional stages are defined in the **Create Stage** screen. Every stage must have a unique name and an optional description can be supplied. If the stage is not the final stage, you must specify how long the data is to remain in this stage and define any stage attributes such as whether the data should be compressed or set to read-only. Note that it is only possible to specify a read-only stage if a preferred read-only tablespace has been defined for the logical storage tier for this stage.

Click **Create and Add Another** to save this stage and add another stage, or click **Create** if this is the last stage. As the stages are defined a report detailing every stage defined for this lifecycle is shown in the **Stage Details** report.

The current stage represents the present time but can span any length of time. A lifecycle can only have one current stage.

The final stage is required as it describes what happens when data reaches its end-of-life. A lifecycle can only have one final stage and is automatically created if the user does not create one. Possible actions that can be set are:

- Purge the data
- Archive the data offline
- Allow the data to remain online

To define the stages for the lifecycle, do as follows:

1. On the Oracle ILM page, click Lifecycle Setup > Lifecycle Definitions > New Stage.

Manage Stage for Lifecycle Definition my_life_cycle Back to Lifecycle Definition Apply OK

Manage a stage for a lifecycle definition. ⓘ

ⓘ Stage Name

ⓘ Description

ⓘ Stage Immediately Precedes Stage

ⓘ Retain Data for

ⓘ Logical Storage Tier ⓘ Stage Attributes

 Data that resides on this stage will be protected from change. Once the lock has been applied, it cannot be unlocked until the retention period expires or canceled by an authorized user.

Stage Reference

This is a reference of all existing stages for the lifecycle definition my_life_cycle. Current dates are derived from retention and change with each event scan. ⓘ

	Stage Name	Description	Action	Tier Name	Attributes	Retention Period	Stage Start Date	Stage End Date
Current Stage	my_high_perf_stage	current	Remain Online	my_high_perf	None	Current 10 Days	30-Sep-2008	09-Oct-2008
	my_lowcost_stage	second stage	Remain Online	Low Cost	None	1 Month	31-Aug-2008	29-Sep-2008
	my_final_sl_stage		Remain Online	SL_ENT_2TO20	Read-Only	10 Years	31-Aug-1998	30-Aug-2008
Final Stage	Final_purge_stage	third and final stage	Purge		None			30-Aug-1998

Figure 7) Manage stage for lifecycle definition.

2. Enter the appropriate fields and click OK.

6.3 CREATING DATA ACCESS AND MIGRATION POLICIES

Access policies determine data visibility and are transparent to the applications. Only special users or operations access historical data, as hiding historical data speeds up data scans and maintenance. OILMA advises when it is time to move data, and recommends when to move data between storage tiers as access patterns change.

Based on the enterprise requirements that specify what data needs to be archived and how often, enterprises can define automatic policies to migrate the data from a read-write storage tier to a WORM storage tier.

SCANNING LIFECYCLE EVENTS

Lifecycle event scanning determines when it is time to move data across storage tiers or whether it is ready for archiving or deletion.

Events are organized in groups called recommendations. This is necessary to guarantee that related events are implemented or rejected together. Therefore, you must accept or reject at the recommendation level rather than at the individual event level.

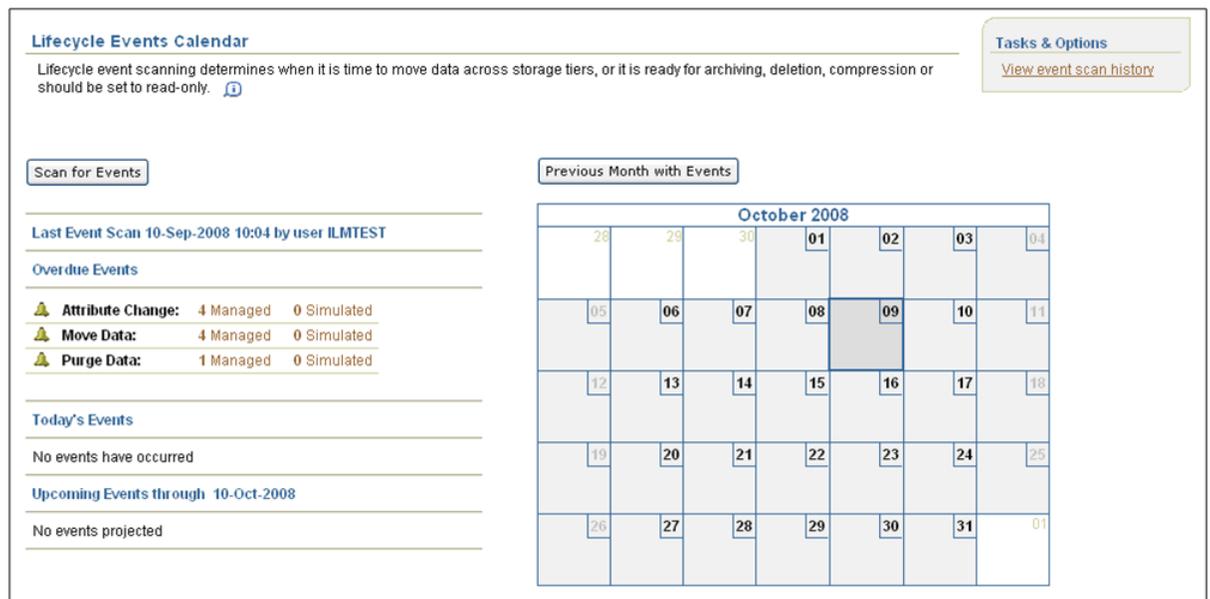


Figure 8) Lifecycle events calendar.

RECOMMENDATIONS FOR LIFECYCLE EVENT

Lifecycle recommendations are implemented using Oracle SQL scripts. By rule, only recommendations that are accepted will be implemented; however, the user may further filter the implementation script by selecting specific tables and target execution date at script generation time.

A script can be previewed on the screen, saved to a file, or even dropped onto a queue for automatic execution. When implementing recommendations within the ILM Assistant, the script actions are sent to the Oracle scheduler system for execution. The specified event date is used as the execution date and time for the submitted job.

To view the lifecycle events, on the Oracle ILM page, click Lifecycle Management > Lifecycle Events.

Figure 9 shows an example of the lifecycles events for the time period selected.

Lifecycle Events for 09-Sep-2008 [Back to Calendar](#)

Summary of outstanding lifecycle events. ⓘ

Last scan date 10-Sep-2008 10:04 **Recommendations accepted** 5
Last event date 10-Oct-2008 **Recommendations rejected** 0
Last scanned by ILMTEST

Tasks & Options
[View implementation status](#)
[Implement accepted recommendations](#)
[Accept selected recommendations](#)
[Reject selected recommendations](#)
[Show filter options](#)

Events for Managed Tables | [Events for Simulated Tables](#)

Table Name: ILMTEST.CP_CC_TRAN

<input type="checkbox"/>	Event Status	Recm ID	Recommended Action	Partition Name	Current Tier	Recommended Tier	Cost Savings	Event Date	Implement Status	Event Details
<input type="checkbox"/>	Accepted	21	Purge Data	VERY_OLD_CC_TRANS	Online Archive		\$0	Past	Ready	
<input type="checkbox"/>	Accepted	22	Move Partition	OLD_CC_TRANS	Online Archive	SL_ENT_2TO20	-\$3	Past	Ready	
<input type="checkbox"/>	Accepted	22	Set Retention	OLD_CC_TRANS	Online Archive	SL_ENT_2TO20	\$0	Past	Ready	
<input type="checkbox"/>	Accepted	23	Move Partition	ALL_YEAR_CC_TRANS	Low Cost	SL_ENT_2TO20	-\$2	Past	Ready	
<input type="checkbox"/>	Accepted	23	Set Retention	ALL_YEAR_CC_TRANS	Low Cost	SL_ENT_2TO20	\$0	Past	Ready	
<input type="checkbox"/>	Accepted	24	Move Partition	THIS_YEAR_CC_TRANS	Low Cost	SL_ENT_2TO20	-\$2	Past	Ready	
<input type="checkbox"/>	Accepted	24	Set Retention	THIS_YEAR_CC_TRANS	Low Cost	SL_ENT_2TO20	\$0	Past	Ready	
<input type="checkbox"/>	Accepted	25	Move Partition	NEW_CC_TRAN	High Performance	SL_ENT_2TO20	\$0	Past	Ready	
<input type="checkbox"/>	Accepted	25	Set Retention	NEW_CC_TRAN	High Performance	SL_ENT_2TO20	\$0	Past	Ready	

1 - 9

Figure 9) Lifecycle events.

6.4 DEFINING AND ENFORCING COMPLIANCE POLICIES

Figure 10 shows the implemented compliance and security policies. You may have to assign the policies, when defining the lifecycle definition.

To create a new policy, click Policy Notes > New Policy.

ORACLE ILM Assistant [Skip navigation](#) [Print](#) [Logout](#)

User: ILMTEST Lifecycle Setup | Lifecycle Management | Reports | **Compliance & Security** | Help

Current Status | Immutability | Privacy & Security | Auditing | **Policy Notes**

Overview
 Displays the current status of functionality implemented to ensure security and regulatory compliance. ⓘ

Data Retention Policies

Retention Policies Defined: Yes
 Number of Lifecycle Definitions: 3
 Number of Lifecycle Tables: 6

Database Security

Column: No
 Transparent Data Encryption: No
 Oracle Wallet: No

Virtual Private Database Policies

VPD Policies in Use: Yes
 Number of VPD Policies: 1

Auditing

Database: On
 Fine-Grained Auditing: On
 Number of Audit Policies: 5

Immutability

Number of Result Sets: 6
 Number of Signatures: 6
 Last Generated: 16-Oct-2006 18:01
 Last Comparison: 16-Oct-2006 17:53

User: ILMTEST Language: en-us Copyright © 2006, 2007, Oracle. All rights reserved.

Figure 10) Implemented compliance and security settings.

ENFORCING COMPLIANCE POLICIES: RETENTION

Lifecycle event details include information specific the event, including recommended tablespace, compression attributes, read-write attributes, retention settings for protected data, and estimated cost savings. Additionally, a link is provided to view the SQL statement that would be used to implement the recommendation.

To implement the recommendations, run the SQL script.

The screenshot displays the Oracle Lifecycle Management interface. At the top, the user is identified as 'User: ILMTEST'. The navigation menu includes 'Lifecycle Setup', 'Lifecycle Management', 'Reports', 'Compliance & Security', and 'Help'. The current view is 'Lifecycle Events', with a sub-tab for 'Lifecycle Events Calendar'. The main content area is titled 'Lifecycle Event Details' and includes a 'Back to Lifecycle Events' button. Below the title, there is a link for 'Detailed information on lifecycle event.' and a small icon. The event details are presented in two columns. The left column lists: Managed Table Name (ILMTEST.CP_CC_TRAN), Action (Move Partition), Partition Name (OLD_CC_TRANS), Current Storage Tier (Online Archive), Current Tablespace (ILMDEMO_10), Recommended Storage Tier (SL_ENT_2TO20 (SL_ENT_2TO20)), Recommended Tablespace (SL_2TO20_TPLATE1 (Cloned from SL_2TO20_TPLATE)), and Recommended Attributes (Read-Only). The right column lists: Lifecycle Definition (my_life_cycle), Event Date (Past), Event Status (Accepted and ready to implement), Cost Savings (-\$3), and Explanation (Partition OLD_CC_TRANS needs to be moved to a new storage tier. A new tablespace must be created because the target location is a worm device. Data will be set to read-only, and compression will be disabled). Below this, there is a section for 'Event SQL' with a dropdown arrow. The text states: 'The following SQL operations are used to implement the current event.' followed by a SQL script:

```
begin
  ilm_toolkit.ilm_toolkit.set_tablespace('ILMDEMO_10','READ WRITE');
end;
/
create tablespace "SL_2TO20_TPLATE1"
  datafile '/u02/spectrevol/sl_2to20/sl_2to20_tplate0111.dbf' size 50m
  blocksize 8192;
begin
  ilm_toolkit.ilm_toolkit.add_clone_tablespace(169,'SL_2TO20_TPLATE1','SL_2TO20_TPLATE',TRUE);
end;
```

Figure 11) Detailed information on lifecycle event.

6.5 VERIFYING THE RETENTION DATE

You can verify the data file associated with a tablespace from DB host, and check if the file is read-only and if the retention date is set, by logging in to the database host and entering:

```
$pwd
/u02/spectrevol/sl_2to20
$ ls -lt
total 102636
-rw-r----- 1 oracle dba 52436992 Oct 31 2008 sl_2to20_tplate01.dbf
-r--r----- 1 oracle dba 52436992 Aug 29 22:03 sl_2to20_tplate0141.dbf
```

User: ILMTEST

Storage Costs | Storage Tiers | Tables | Partitions | Data Usage | Lifecycle Retention | Data Protection | Lifecycle Setup | Lifecycle Management | Reports | Compliance & Security | Help

Partitions by Table | by Storage Tier

Displays partitions which are managed by the ILM Assistant. ⓘ

Tasks & Options
Show filter options

Managed Tables | Simulated Tables

Table Name: ILMTEST.CP_CC_TRAN Partition Key: TRAN_DT Lifecycle Definition: my_life_cycle

Partition Name	Low Key Value	High Key Value	Tablespace Name	Read Only	Compression	Stage Name	Logical Storage Tier	Expiration Date	Storage Size (GB)	Cost Per GB	Storage Cost
NEW_CC_TRAN	01-Sep-2005	30-Sep-2005	ILMDEMO_13	No	Disabled	my_final_sl_stage	High Performance		0	\$50.00	\$3
RECENT_CC_TRANS	01-Jul-2005	31-Aug-2005	SL_2T020_TPLATE4	Yes	Disabled	my_final_sl_stage	SL_ENT_2T020	11-Oct-2015	0	\$50.00	\$3
THIS_YEAR_CC_TRANS	01-Oct-2004	30-Jun-2005	ILMDEMO_11	No	Disabled	my_final_sl_stage	Low Cost		0	\$12.00	\$1
ALL_YEAR_CC_TRANS	01-Oct-1998	30-Sep-2004	ILMDEMO_11	No	Disabled	my_final_sl_stage	Low Cost		0	\$12.00	\$1
OLD_CC_TRANS	01-Oct-1996	30-Sep-1998	ILMDEMO_10	No	Disabled	my_final_sl_stage	Online Archive		0	\$4.00	\$0
VERY_OLD_CC_TRANS		30-Sep-1996	ILMDEMO_10	No	Disabled	Final_purge_stage	Online Archive		0	\$4.00	\$0
Totals for table:									0		\$8

Table Name: ILM_DEMO.CC_TRAN Partition Key: TRAN_DT Lifecycle Definition: Audit Data

Partition Name	Low Key Value	High Key Value	Tablespace Name	Read Only	Compression	Stage Name	Logical Storage Tier	Expiration Date	Storage Size (GB)	Cost Per GB	Storage Cost
NEW_CC_TRAN	01-Sep-2005	30-Sep-2005	ILMDEMO_13	No	Disabled	old_audit	High Performance		0	\$50.00	\$3
RECENT_CC_TRANS	01-Jul-2005	31-Aug-2005	ILMDEMO_12	No	Disabled	old_audit	Low Cost		0	\$12.00	\$1
THIS_YEAR_CC_TRANS	01-Oct-2004	30-Jun-2005	ILMDEMO_11	No	Disabled	old_audit	Low Cost		0	\$12.00	\$1

Figure 12) Verifying the retention date.

Now you can observe that data has been migrated to WORM tier (SnapLock volume) and the retention date (expiration date) has been set, so that the data is protected until that period.

This NetApp ILM for Oracle Database product helps customer to implement 'Tiered Storage' architecture in their environment and also helps them meet the regulatory requirements.

7 SUMMARY

NetApp ILM for the Oracle Database is an effective and simple solution that helps to easily manage data through its lifecycle. It helps in defining data archival policies, adhering to defined security policies, and improving space efficiency and reducing costs. This document intends to provide detailed instructions on the how to install, configure and effectively implement NetApp ILM for your environment.

In case of errors, omissions, or feedback, email padmanabhan.sadagopan@netapp.com.

APPENDICES

APPENDIX A: REFERENCES

- http://now.netapp.com/knowledge/docs/NetApp_ILM/1.1/pdfs/NetApp_ILM_install_config_guide.pdf
- For information about the NetApp specific extension for Oracle ILM Assistant, see the ILM Assistant SDK Guide Version 1.3.
- www.netapp.com/us/solutions/applications/oracle/oracle-archive-compliance.html
- <http://media.netapp.com/documents/ds2862-oracle-archive.pdf>
- www.oracle.com/technology/software/deploy/ilm/index.html
- www.oracle.com/technology/software/deploy/ilm/ilma_users_guide_0.html

APPENDIX B: SCRIPTS

Table 1 lists the scripts that are invoked from `run_install.sql`.

Table 1) Scripts.

Script	Description
<code>create_mot_user1.sql</code>	Creates a MOT schema. It requires the MOT schema password, default tablespace, and temporary tablespace.
<code>grants_to_mot.sql</code>	Assigns necessary grants to MOT.
<code>load_mot_java.sql</code>	Loads the <code>manageontap.jar</code> and <code>DBMSWrapper.class</code> files to the Oracle Database.
<code>compile_java_class.sql</code>	Compiles the Java classes.
<code>verify_java_class.sql</code>	Verifies the status of the Java classes. Check whether the status is "VALID"; otherwise recompile.
<code>create_dbms_ontap_pkg_head.sql</code>	Creates the DBMS_ONTAP package specification in MOT schema.
<code>create_dbms_ontap_pkg_body.plb</code>	Creates the DBMS_ONTAP package body in MOT schema.
<code>verify_dbms_ontap_pkg.sql</code>	Verifies the status of the DBMS_ONTAP package.
<code>netapp_ilma_install_head.sql</code>	NetApp specific extension to Oracle ILM Assistant: Creates the object type instantiating from ILM_TOOLKIT schema.
<code>netapp_ilma_install_body.plb</code>	NetApp specific extension to Oracle ILM Assistant: Creates the object type body instantiating from ILM_TOOLKIT schema.

APPENDIX C: TROUBLESHOOTING TIPS

ISSUE

How to delete a file system that was added using the “netapp_add_fs.sql” script.

SOLUTION

To remove an incorrectly added file system, do as follows:

1. Execute the `ilm_toolkit.ilm_toolkit.delete_file_system` package by logging in as `sysdba`. Use the file system unique ID as a parameter.

For example:

```
SQL> exec ilm_toolkit.ilm_toolkit.delete_file_system (<file system id>);
```

```
SQL> exec ilm_toolkit.ilm_toolkit.delete_file_system (4444);
```

PL/SQL procedure successfully completed.

2. Drop the directories manually by entering:

```
SQL> drop directory <Directory Name>
```

```
SQL> drop directory NETAPP_SNAPE_4444
```

ISSUE

The following error appears when trying to run “netapp_add_fs” to add a SnapLock volume as a file system:

```
SQL> @netapp_add_fs
```

```
l_obj ilm_toolkit.ilm$snap_t;
```

```
*
```

ERROR at line 3:

ORA-06550: line 3, column 9:

PLS-00201: identifier 'ILM_TOOLKIT.ILM\$SNAP_T' must be declared

ORA-06550: line 3, column 9:

PL/SQL: Item ignored

ORA-06550: line 75, column 1:

PLS-00320: the declaration of the type of this expression is incomplete or malformed

SOLUTION

Make sure NetApp ILM extension installation has completed successfully as it creates the `ilm$snap_t`. You can rerun the `run_install.sql` script or only the following two files:

```
@@netapp_ilma_install_head.sql
```

```
@@netapp_ilma_install_body.plb
```

ISSUE

During re-install, during the NetApp ILM Extension installation the following ORACLE errors appear. But, the install completes successfully,

```
NetApp ILM Extension Installation
```

```
Oracle ILM Assistant v1.3 must be installed before proceeding further.
```

```
*****
```

```
Press Enter to Continue with the installation
```

```
PL/SQL procedure successfully completed.
```

```
Grant succeeded.
```

Grant succeeded.

```
DROP TYPE ilm_toolkit.ilm$_snap_t validate
```

*

ERROR at line 1:

ORA-30746: error occurred while trying to drop column "DEVICE_TYPE" in table
"ILM_TOOLKIT.ILM\$_FILE_SYSTEMS"

ORA-02303: cannot drop or replace a type with type or table dependents

```
CREATE OR REPLACE TYPE ilm_toolkit.ilm$_snap_t
```

*

ERROR at line 1:

ORA-02303: cannot drop or replace a type with type or table dependents

No errors.

SP2-0272: escape character cannot be alphanumeric or whitespace

Type body created.

No errors.

Type body created.

NetApp ILM Extension Installation Completed

SOLUTION

This is normal behavior as the `ilm_toolkit.ilm$_snap_t` Oracle object type is already created, and there are file systems based on it.

Before you drop the type, delete the dependents; in this case, a file system. Use the `ilm_toolkit.ilm_toolkit.delete_file_system` to delete dependents.

APPENDIX D: NETAPP ILM SQL SCRIPT FILES

Using the following SQL script files, you can directly connect to the database and get the NetApp storage system specific information:

DataFileToVol.sql

PhyToVol.sql

VolToPhyPath.sql

getretention.sql

setretention.sql

getsysteminfo.sql

getslvolinfo.sql

getslvoloptioninfo.sql

getontapver.sql

SAMPLE SQL File: (Get Storage System information from Oracle Database)

You can create your own SQL file as shown below to get NetApp storage system related information from your Oracle Database:

```
SELECT DBMS_ONTAP.getSLVolOptionInfo(
'&&Filer_IP_Address', '&&UserName_of_Filer', '&&Password_of_Filer',
'&&AbsolutePath_of_DataFileName' , '&&attribute', '&&attribute1') "Option"
FROM DUAL ;

SELECT
DBMS_ONTAP.getSystemInfo('&&Filer_IP_Address', '&&UserName_of_Filer', '&&Password_
of_Filer', '&&attribute') "System Information"
FROM DUAL ;
```

The following is a sample output from running getslvoloptioninfo.sql from the SQL prompt:

```
SQL> @getslvoloptioninfo

"-----This file will find SnapLock Volume options like - like
'snaplock_maximum_period' an so on -----"

"----- Please enter YOUR Filer Details -----"
"
"-----Filer_IP_Address : <Enter IP Address> -----"
"-----UserName_of_Filer : <Enter User name> -----"
"-----Password_of_Filer : <Enter Password> -----"
"-----Path_DataFile_or_Device : <Either enter the device name OR complete data file name>
--"
"-----attribute : <Can be any of the following—'snaplock_maximum_period',
'snaplock_minimum_period' and so on> -----"
"-----attribute1 : Enter as 'value' if only value is required; enter it as
'metric' if only metric is required; enter it as 'null' if both value and metric
is required --"
```

"-----attribute2 : Enter as 'DataFile' if a datafile name is passed. Other
wise enter 'null' ---"

Enter value for filer_ip_address: **10.x1.ly2.12**

Enter value for username_of_filer: **root**

Enter value for password_of_filer: **XXX**

Enter value for path_datafile_or_device: **/vol/sl_2to20**

Enter value for attribute: **snaplock_default_period**

Enter value for attribute1: **null**

Enter value for attribute2: **null**

Option

min

SQL>

ACKNOWLEDGEMENTS

Niranjan Mohapatra

Preetom Goswami

Uday Shet

Lynne Thieme

Steven Schuettinger

Esther Smitha

DISCLAIMER

NetApp provides no representations or warranties regarding the accuracy, reliability or serviceability of any information or recommendations provided in this publication, or with respect to any results that may be obtained by the use of the information or observance of any recommendations provided herein. The information in this document is distributed AS IS, and the use of this information or the implementation of any recommendations or techniques herein is a customer's responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. This document and the information contained herein must be used solely in connection with the NetApp products discussed in this document.