

An Oracle and 3PAR White Paper
April 2010

Maintaining High Storage Utilization with Oracle ASM Storage Reclamation Utility and 3PAR Thin Persistence

Introduction	2
Problem: Preserving High Storage Utilization Over Time	2
Solution: Online, Non-Disruptive Space Reclamation for Oracle	3
Overview of ASRU Operation.....	4
When to Use ASRU to Reclaim Storage	5
Example of Using ASRU to Reclaim Storage on 3PAR	6
Summary	8

Introduction

IT Organizations face increasing pressure to deliver high-performance Oracle databases and other mission-critical applications at the lowest possible cost. The old days of throwing disks at storage problems, which resulted in under-utilized resources, is a thing of the past. Deploying Oracle databases with cost-effective thin provisioned storage is an ideal way to achieve high storage efficiency and dramatic storage capacity savings. By boosting storage utilization, thin provisioning drives savings in purchased capacity and associated power and cooling costs.

But, what good are up-front savings if they can't be preserved over time? As common database operations delete or move large amounts of data, storage capacity utilization rates—even in Oracle database environments using thin provisioned storage—are compromised. This technical paper discusses how organizations with Oracle database environments and thin provisioned storage can now capitalize on a new combined Oracle-3PAR solution to maintain high storage utilization over time. Together, Oracle and 3PAR deliver the ability to reclaim unused (but allocated) ASM disk space—simply, quickly, and non-disruptively, as demonstrated by tests conducted by Oracle and 3PAR.

Problem: Preserving High Storage Utilization Over Time

Beginning in 2004, Oracle and 3PAR teamed together to enable high-performance Oracle 10g (and later 11g) databases to be deployed with cost-efficient thin provisioned storage for a dramatic increase in disk utilization, resulting in up to 50% less wasted storage capacity as compared to using Oracle databases with traditional storage arrays.

Enabling technologies for this combined Oracle and 3PAR solution included:

- **Oracle Autoextend** – Oracle has actively supported storage vendors' thin provisioning capabilities with a database feature known as *autoextend*. Autoextend is to databases as thin provisioning is to storage. As applications add new data to database tables and additional space is needed, autoextend will automatically grow the tablespace in size. Accordingly, the storage array would automatically allocate physical storage to support the tablespace growth.
- **Oracle Automatic Storage Management (ASM)** – A purpose-built file system and volume manager integrated in Oracle databases. ASM dramatically simplifies database file management and storage administration. Oracle database deployment with thin provisioned storage was first tested using ASM.
- **3PAR Thin Provisioning** – A green storage technology that dramatically cuts capacity, energy, and related costs while substantially alleviating storage and system administration overhead. Capacity is dedicated and configured autonomically and in small increments from a single, reservationless, comprehensively scalable reservoir, so storage is provisioned automatically, efficiently, and on an as-needed basis.

The announced interoperability was backed with joint testing with results published in a joint Oracle-3PAR white paper published at:

http://www.oracle.com/technology/products/database/asm/pdf/oracle_3par_wp_final_0.pdf

In an ideal world, once the storage for an Oracle database starts thin, it should remain thin. That is, the capacity allocated by the storage array to the thin provisioned volumes supporting the Oracle database should remain in lockstep with the amount of actual written data stored in the Oracle database, so that the capacity savings from thin provisioning are preserved over time. But, that is often not the case. Over the Oracle database lifecycle, the utilization of allocated storage capacity in a thin provisioned volume can decrease as changes are made to the database through common operations such as:

- Dropping of a tablespace or database upon deletion of transient data
- Resizing of an Oracle datafile upon shrinking a tablespace
- Addition of new disks to an ASM disk group to accommodate growth or load balance performance

These changes result in the creation of unused ASM disk space that can build up over time to account for up to 50% of the total storage capacity provisioned to Oracle databases. This space is available for reuse within ASM, but, in the absence of a control communication protocol between applications/file systems and block storage, the storage array is unable to distinguish between capacity associated with deleted data and valid data. Therefore, the unused capacity remains allocated and in use within the storage volume(s) on the storage array. The end result is that the storage utilization falls below desirable levels.

Solution: Online, Non-Disruptive Space Reclamation for Oracle

Building on the original goal of driving high levels of resource utilization, Oracle and 3PAR have partnered again to extend storage efficiency for Oracle database environments. Oracle and 3PAR now offer the ability to improve storage efficiency for Oracle Database 10g and 11g environments by reclaiming unused (but allocated) ASM disk space in thin provisioned environments. This extension of storage efficiency is enabled by two recent innovations:

- **Oracle ASM Storage Reclamation Utility (ASRU)** – Oracle ASRU is a new utility that extends Oracle's support of 3PAR Utility Storage with Thin Provisioning by enabling space reclamation. Oracle ASRU compacts the ASM disks, writes zeroes to the free space, and resizes the ASM disks to original size with a single command, online and non-disruptively.
- **3PAR Thin Persistence** – 3PAR Thin Persistence software detects zero writes and eliminates the capacity associated with free space in thin provisioned volumes—simply, quickly, and without disruption. 3PAR Thin Persistence leverages the unique, built-in, zero-detection capabilities of the 3PAR Gen3 ASIC within all 3PAR InServ Storage Server models with Thin Built In™. Unlike alternative CPU-based zero-detection approaches that are slow and disruptive, 3PAR's revolutionary hardware capability, built right into the Gen3 ASIC, provides an efficient, silicon-based, zero-detection mechanism to identify the unused space—quickly and without performance impact. Subsequently, the virtualization mapping capabilities of 3PAR Thin Engine—built into the 3PAR InForm® Operating System—remap the storage volume without the unnecessary bulk. Together, 3PAR Thin Persistence software and the 3PAR Gen3 ASIC deliver fast, online reclamation of unused storage capacity.

Tests conducted by Oracle and 3PAR demonstrate dramatic gains in storage utilization that can be preserved over time with the use of this new capability. In one test, which began with four databases that occupied 886 GB of storage in a 1-TB Disk Group, deleting two databases and running ASRU recovered 330 GB of allocated but unused storage. In this test, ASRU with 3PAR Thin Persistence was able to reclaim 37% of the original allocated storage, thereby yielding savings in purchased capacity.

This unique Oracle-3PAR solution allows IT organizations to achieve and maintain high levels of utilization for Oracle environments that use thin provisioned storage. This solution can result in significant savings for base capacity and associated costs for power and cooling.

Overview of ASRU Operation

Oracle ASM Storage Reclamation Utility (ASRU) is a stand-alone utility used to reclaim storage in an ASM disk group that was previously allocated but is no longer in use. The ASRU utility, a Perl script, accepts the name of the disk group for which space should be reclaimed. When executed, it writes blocks of zeros to regions on ASM disks where space is currently unallocated. The 3PAR InServ Storage Server, using the zero-detect capability of the Gen3 ASIC, will detect these zero blocks and reclaim any corresponding physical storage.

The administrator invokes the ASRU utility, which operates in three phases:

- **Compaction Phase** – In this phase, ASRU logically resizes the disks downward such that the amount of space in the disk group is at the allocated amount of file space in the disk group, plus a reserve capacity. The default value for the reserve amount is 25 percent; however, the reserve value is a settable option in the utility. The resize operation of the disks is logical to ASM and has no effect on the physical disks. The effect of the resize operation is that file data in the ASM disk group is compressed near the beginning of the disks which is accomplished by an ASM rebalance of the disk group. The utility uses the appropriate database V\$ table to determine the current allocated size of the disk group. The next phase does not begin until the ASM rebalance for the disk group has completed and verified as complete. (Although this phase invokes an ASM rebalance, it does not perform a complete extent relocation operation, just the compaction portion of the rebalance operation. Therefore, it should minimally impact the environment.)
- **Deallocation Phase** – During this phase, ASRU writes zeros above the region where the ASM disks have been resized. The ASRU utility invokes another script called zerofill that does the writing of zeros. It is during this deallocation phase that the zero-detect algorithm within the 3PAR Thin Engine will return the freed storage blocks to the free storage pool.
- **Expansion Phase** – In the final phase, all of the ASM disks will be resized to their original size as determined when ASRU was started. This resize operation is a logical resize of the disks with respect to ASM and does not result in a reorganization of file data in the disk group.

When to Use ASRU to Reclaim Storage

Storage reclamation should be considered after several different types of events:

- Dropping one or more databases
- Dropping one or more tablespaces
- Adding one or more new volumes to an ASM Disk Group, which triggers an ASM rebalance to move a subset of the data from the old volumes to the new volume(s). The storage released from the old volumes is a candidate for reclamation.

To determine whether storage reclamation will be beneficial after one of these operations, it is important to consider the effect of the reserve maintained by ASRU when the utility reduces the size of the disk group during the compaction phase. The temporarily reduced size is equal to the allocated space plus a reserve which allows active databases to grow during the reclamation process; the default reserve is 25% of the allocated storage. Storage reclamation is likely to be beneficial if the amount of allocated physical storage significantly exceeds the amount of storage allocated within ASM plus the reserve.

The amount of physical storage allocated on a 3PAR InServ array can be determined using the 3PAR InForm Operating System's `showvv` command, available from the InForm Command Line Interface (CLI), to show information about the Virtual Volumes (VVs) used by ASM. The usual way of using this command to obtain information related to the effectiveness of Thin Provisioning for a group of volumes matching `oe15.*` is

```
cli% showvv -s oe15.*
```

The `-s` option produces voluminous output, however, so to make the output easier to understand this paper will use more complex options that show just the data columns that are directly relevant to Thin Provisioning:

```
cli% showvv - showcols \  
Name,Usr_Rsvd_MB,Usr_Used_MB,Usr_Used_Perc,Tot_Rsvd_MB,VSize_MB oe15.*  
Name          Usr_Rsvd_MB Usr_Used_MB Usr_Used_Perc Tot_Rsvd_MB VSize_MB  
oe15.1_asm    208896      206433      80.6          209152      256000  
oe15.2_asm    208896      206499      80.6          209152      256000  
oe15.3_asm    208896      206445      80.6          209152      256000  
oe15.4_asm    208896      206443      80.6          209152      256000  
-----  
total         835584      825770          914944      1024000
```

The "Usr_Used_MB" column indicates how many megabytes are actually allocated to user data. In this example, 825,770 megabytes of storage within ASM's volumes has been written.

ASM's view of how much storage is in use can be determined with a SQL query:

```
SQL> select name, state, type, total_mb, free_mb from v$asm_diskgroup where
name = 'LDATA';
```

NAME	STATE	TYPE	TOTAL_MB	FREE_MB
LDATA	MOUNTED	EXTERN	1023984	197986

This example shows 197,986 megabytes of free storage out of 1,023,984 megabytes available or about 19.3%. The difference between these quantities—825,998 megabytes—is how much storage within ASM is in use, i.e., has actual written data.

Example of Using ASRU to Reclaim Storage on 3PAR

To illustrate storage reclamation using Oracle ASRU and the 3PAR InServ Storage Server, a 1-TB ASM disk group was created using four 250-GB Thin Provisioned Virtual Volumes (TPVVs) on the 3PAR InServ array. Zero detection is enabled for the volumes from the InForm CLI as follows:

```
cli% setvv -pol zero_detect oel5.*
```

Four databases were then created, each about 200 GB in size, using 80% of the available storage in the disk group:

```
SQL> select name, state, type, total_mb, free_mb from v$asm_diskgroup where
name = 'LDATA';
```

NAME	STATE	TYPE	TOTAL_MB	FREE_MB
LDATA	MOUNTED	EXTERN	1023984	197986

Here again was the physical storage as seen by the InServ:

```
cli% showvv - showcols \
Name,Usr_Rsvd_MB,Usr_Used_MB,Usr_Used_Perc,Tot_Rsvd_MB,VSize_MB oel5.*
Name      Usr_Rsvd_MB Usr_Used_MB Usr_Used_Perc Tot_Rsvd_MB VSize_MB
oel5.1_asm 208896      206433      80.6          209152      256000
oel5.2_asm 208896      206499      80.6          209152      256000
oel5.3_asm 208896      206445      80.6          209152      256000
oel5.4_asm 208896      206443      80.6          209152      256000
-----
total      835584      825770          914944      1024000
```

After dropping two of the databases, ASM was again queried to verify that the space allocated to the databases had been returned to ASM's free space:

```
SQL> select name, state, type, total_mb, free_mb from v$asm_diskgroup where
name = 'LDATA';
```

NAME	STATE	TYPE	TOTAL_MB	FREE_MB
LDATA	MOUNTED	EXTERN	1023984	610948

Note: DBAs will typically drop database objects (tables, indexes, etc.) or truncate tables in order to free space. Although the truncate operation does free space inside the tablespace, it does not make that space available for reclamation by ASRU. Datafile(s) must be dropped or shrunk in order to reclaim physical space inside the storage array with the ASRU utility.

The physical storage allocation on the InServ (not shown) was unchanged at this point.

The next step was to run ASRU (as the Oracle user, used in installing ASM-Clusterware) to reclaim the storage:

```
# bash ASRU LDATA
Checking the system ...done
Calculating the new sizes of the disks ...done
Writing the data to a file ...done
Resizing the disks...done

/u03/app/oracle/product/11.2.0/grid/perl/bin/perl -I /u03/app/oracle/product/
11.2.0/grid/perl/lib/5.10.0 /home/ora/zerofill 5 /dev/oracleasm/disks/LDATA2
129081 255996 /dev/oracleasm/disks/LDATA3 129070 255996 /dev/oracleasm/disks/
LDATA4 129081 255996 /dev/oracleasm/disks/LDATA1 129068 255996
126928+0 records in
126928+0 records out
133093654528 bytes (133 GB) copied, 2436.45 seconds, 54.6 MB/s
126915+0 records in
126915+0 records out
133080023040 bytes (133 GB) copied, 2511.25 seconds, 53.0 MB/s
126926+0 records in
126926+0 records out
133091557376 bytes (133 GB) copied, 2514.57 seconds, 52.9 MB/s
126915+0 records in
126915+0 records out
133080023040 bytes (133 GB) copied, 2524.14 seconds, 52.7 MB/s

Calculating the new sizes of the disks ...done
Resizing the disks...done
Dropping the file ...done
```

After completion of ASRU, the reclamation of storage was verified on the 3PAR InServ array:

```
cli% showvv - showcols \
Name,Usr_Rsvd_MB,Usr_Used_MB,Usr_Used_Perc,Tot_Rsvd_MB,VSize_MB oe15.*
Name          Usr_Rsvd_MB Usr_Used_MB Usr_Used_Perc Tot_Rsvd_MB VSize_MB
oe15.1_asm    196224     129516      50.6         196608    256000
oe15.2_asm    193792     129462      50.6         194176    256000
oe15.3_asm    196608     129585      50.6         196992    256000
oe15.4_asm    193920     129520      50.6         194304    256000
-----
total         780544     518083      50.6         782080    1024000
```

In the example above, 308 GB of storage was successfully reclaimed, which was 37% of the allocated storage prior to the reclamation process!

Summary

As modern datacenters are constantly asked to do more with less—especially during times of tight IT budgets—deployment of Oracle environments on thin storage is an ideal solution to significantly reduce storage capacity costs. Oracle databases with ASM coupled with 3PAR Thin Provisioning dramatically cuts capacity and related costs while substantially alleviating storage and system administration. Now, for the first time, with Oracle ASRU and 3PAR Thin Persistence, organizations are not only able to achieve high storage utilization upfront, but are also able to maintain it over time—saving up to 50% of the space that is otherwise occupied by allocated but unused data. Oracle ASRU writes zeroes to this unused space while the 3PAR InServ Storage Server with 3PAR Thin Persistence leverages built-in zero-detection capability to intelligently reclaim space while preserving service levels and without disruption or performance impact. With Oracle deployed on 3PAR, achieving and maintaining high storage utilization has never been so simple.

ORACLE®



Maintaining High Storage Utilization with
Oracle ASM Storage Reclamation Utility and
3PAR Thin Persistence
April 2010

Authors: Karl L. Swartz, Sandeep Singh (3PAR)
Jim Williams (Oracle)

Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065
U.S.A.

Worldwide Inquiries:
Phone: +1.650.506.7000
Fax: +1.650.506.7200
oracle.com



Oracle is committed to developing practices and products that help protect the environment

Copyright © 2010, Oracle and/or its affiliates. All rights reserved. This document is provided for information purposes only and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle is a registered trademark of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.