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

Draško Tomić, Ph.D.
Date: 19/10/2011

Problem: Preserving High Storage Utilization Over Time

Beginning in 2004, Oracle and 3PAR teamed together to enable high-performance Oracle 10*g (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

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.*

Enabling technologies for this combined Oracle and 3PAR solution

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.

Enabling technologies for this combined Oracle and 3PAR solution

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.

In an ideal world, once the storage for an Oracle database starts thin, tt should remain thin.

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

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.

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.

cli% **showvv - showcols **

3PAR Hardware Architecture: Scale without Tradeoffs

3PAR InSpire® F-Series Architecture

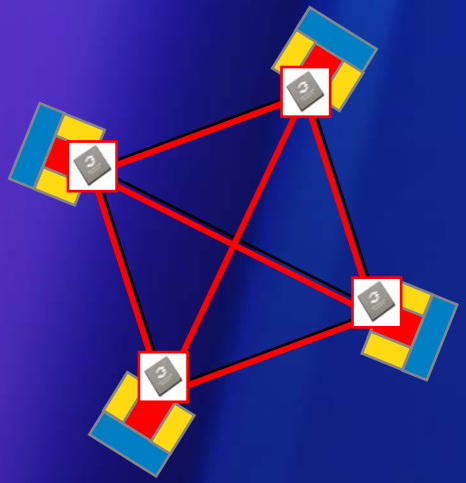

F200,

F400

A finely, massively, and automatically load balanced cluster

3PAR InSpire® T-Series Architecture

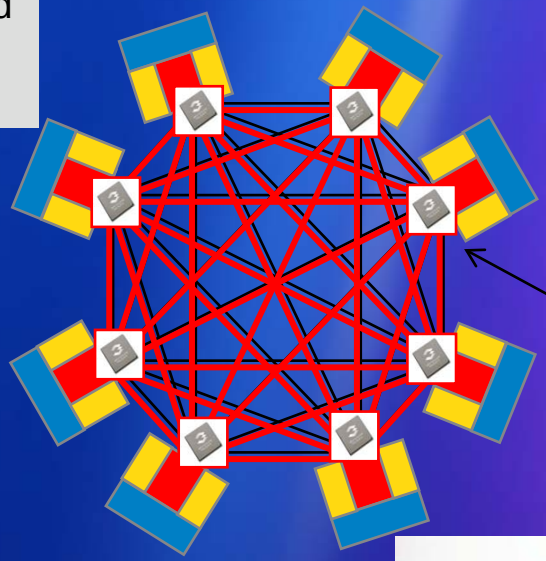

T400,

T800

Legend

- Host Connectivity
- Data Cache
- Disk Connectivity
- Passive Backplane

3PAR Hardware Architecture: Scale without Tradeoffs

Traditional Controllers

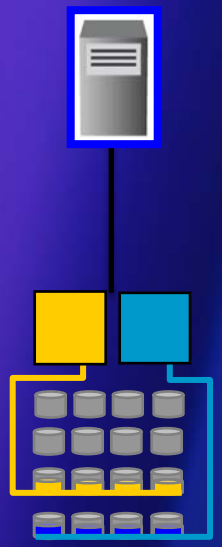

- Each volume is active on only one controller – really “Active-Passive”
- Volume is restricted to drives behind a single controller
- Manual planning and load balancing for each controller
- Cache not necessarily shared

Mesh-Active Controllers

- Each volume is active on all controllers – true “Active-Active”
- Volume is evenly spread across all resources
  - drives
  - controllers
  - cache
  - IO
- Autonomically provisioned
- Cache Coherent

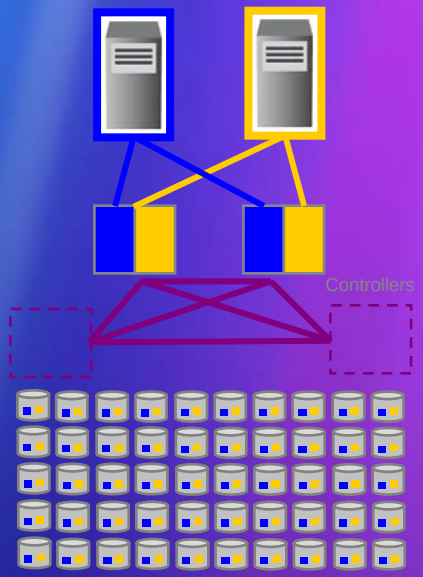

RAID efficiency & resiliency without tradeoffs

Fast RAID 5 and 6

Key enablers:

- Massive parallelization
- 3PAR ASIC for XOR calculations

Tested by: ORACLE

RAID ORGANIZATION SENSITIVITY

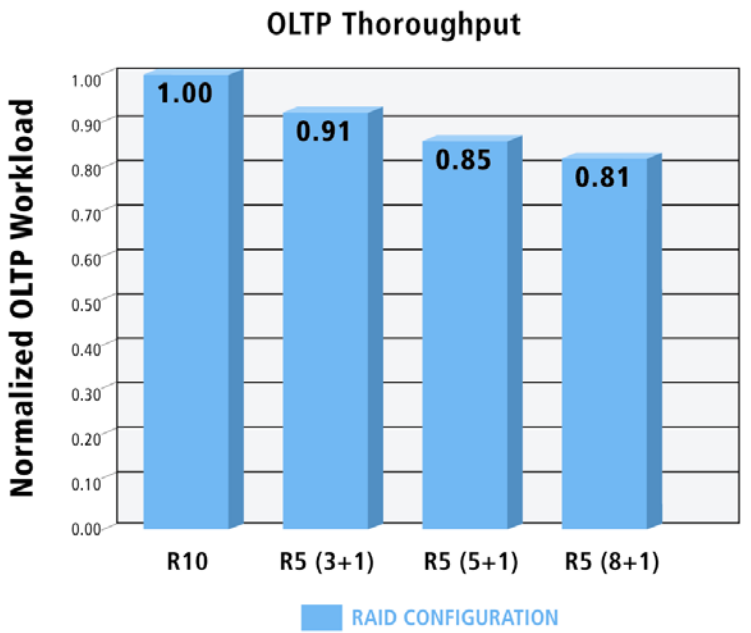

RAID CONFIGURATION SENSITIVITY

	R10	R5 (3+1)	R5 (5+1)	R5 (8+1)
Performance	100%	91%	85%	81%
Space Requirement	200%	133%	120%	112.5%

Test setup included four 100 GB virtual volumes connected via four HBAs to the server. The virtual volumes were statistically balanced across available HBAs.

Mixed Workload:
SQL Server Tests from MSFT

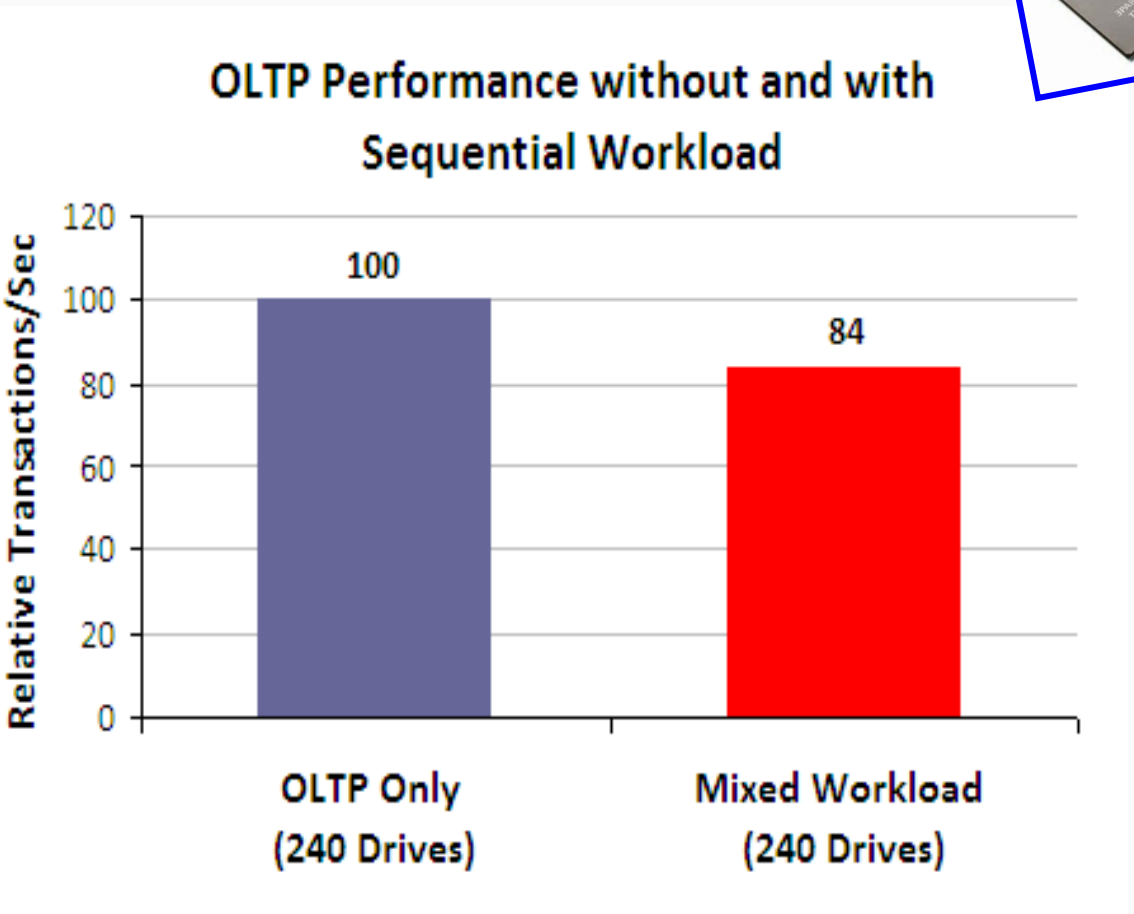

Test Configuration

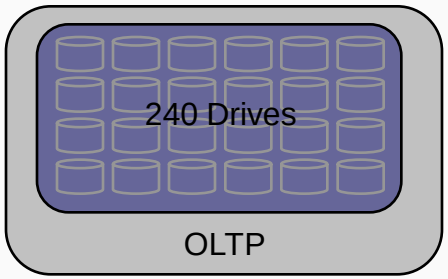

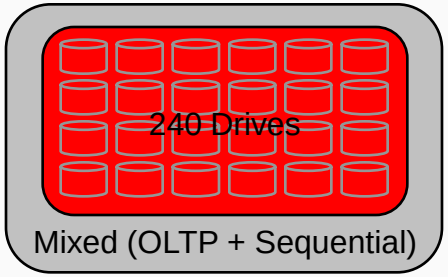

Overview 3PAR Thin Technologies

Thin Conversion	– Rapid migration from Fat to Thin allows for easy conversion from legacy storage environments at line speeds – Preserve service levels without disruption or performance impact during migrations
Thin Persistence	– Reduce storage capacity costs by allowing “thin” volumes to stay “thin”
Thin Copy Reclamation	– Reduce storage capacity costs by reclaiming storage capacity that is deleted but still “locked” in allocated space

Get Thin: 3PAR Thin Conversion

- Thin your online SAN storage up to 75%
- A practical and effective solution to eliminate costs associated with:
  - Storage arrays and capacity
  - Software licensing and support
  - Power, cooling, and floor space
- Unique 3PAR Gen3 ASIC with built-in zero detection delivers:
  - Simplicity and speed – eliminate the time & complexity of getting thin
  - Choice - open and heterogeneous migrations for any-to-3PAR migrations
  - Preserved service levels – high performance during migrations

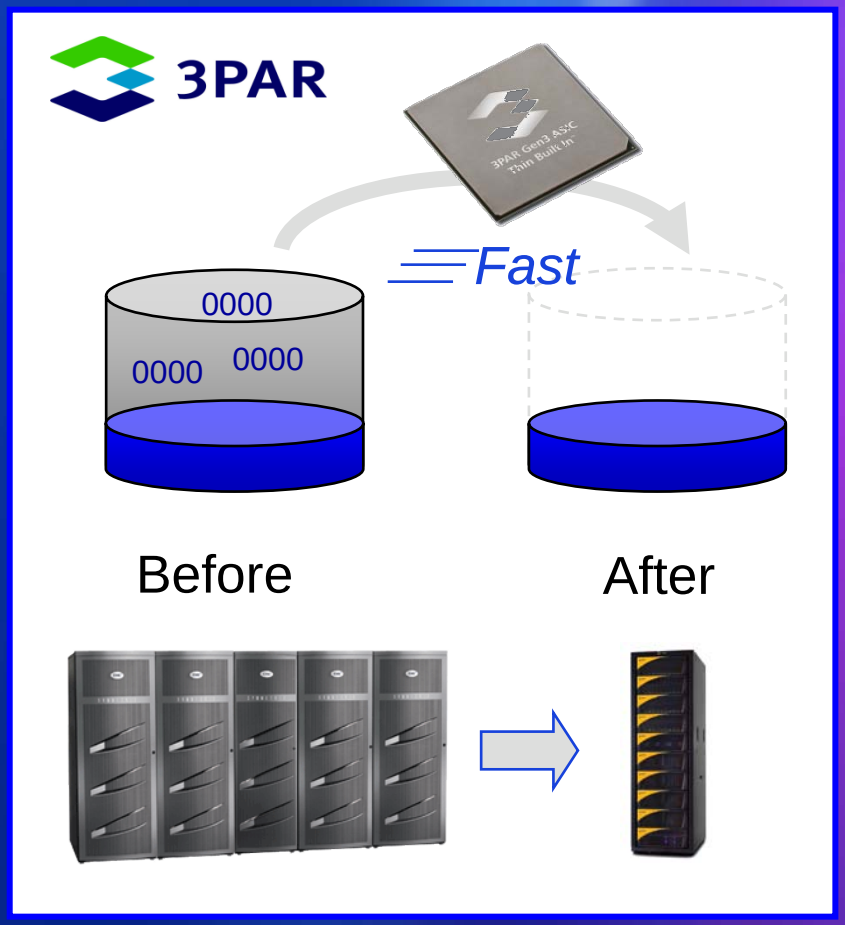

GEN3 ASIC CIRCUIT DIAGRAM

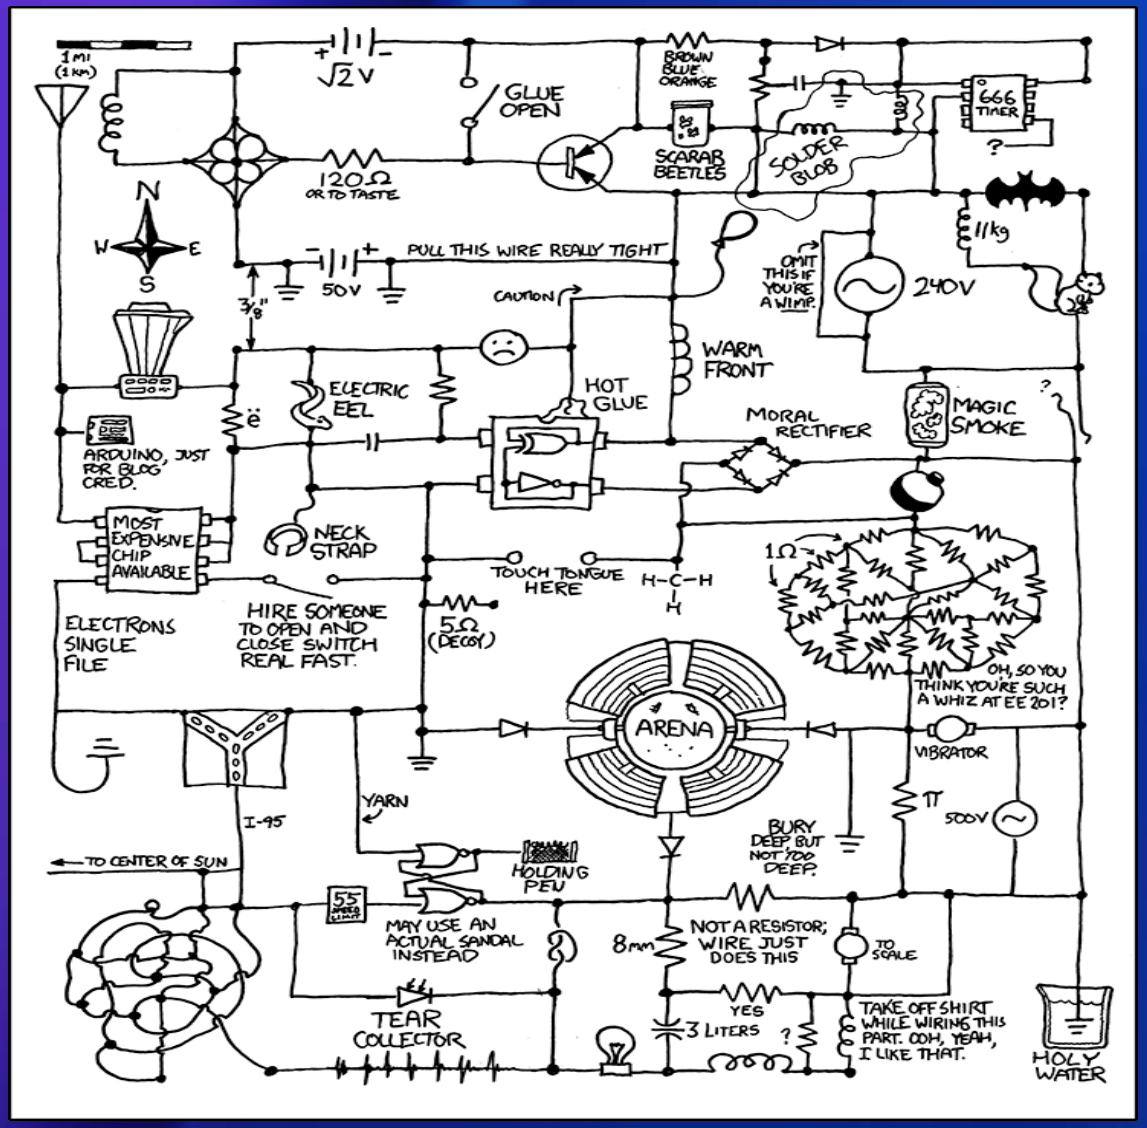

HP Education Services – 3PAR

HP 3PAR Managing 3PAR Disk Arrays

Category: Storage
Course length: 3 days
Level: Beginner/Intermediate

Delivery method: classroom and RAIL (Remotely Assisted Instructional Learning = combines the best of classroom training with the best of online training: without the travel and with reduced ‘out-of-office’ time.
You attend our best-in-class instructor-led training from wherever you are: home, office, anywhere)

This 3 day IDL overviews 3PAR hardware and architecture along with giving administrators insight in to the constructs within the 3PAR array family. Students will explore Virtual Volumes and Thin Provisioning along with performing other administrative tasks via both GUI and CLI methods. The class discusses local and remote data copies along with reporting, scheduling and space management.

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.