

Demystifying ASM

REQUIRED_MIRROR_FREE_MB and USABLE_FILE_MB

Harald van Breederode
Oracle University
3-JUN-2016

ORACLE

Who Am I

- `#include <std/disclaimer.h>`
- Senior Principal Trainer – Oracle University
- 35 years Unix Experience
- 20 years Oracle DBA Experience
- Oracle8*i*, 9*i*, 10*g*, 11*g* and 12*c* OCP
- Oracle10*g*, Oracle11*g* and Oracle12*c* OCM
- DBA and O/S Certification Exam developer/Reviewer
- DBA and O/S Curriculum developer/Reviewer
- Blog: prutser.wordpress.com
- Visually Impaired (Legally Blind)



ORACLE

Certified Master

Oracle Database 12c
Administrator

ORACLE

Agenda

- Overview
- Automatic Storage Management (ASM)
- ASM Architecture
- ASM Storage Hierarchy
- ASM Disk Group Redundancy
- ASM Disk Group Metadata
- ASM Disk Group Error Handling
- ASM Disk Group Space Administration
- ASM on Exadata Storage
- Live Demo
- Questions & Answers

Overview

The goal of this presentation is to demystify the meaning of:

- `REQUIRED_MIRROR_FREE_MB`
- `USABLE_FILE_MB`

Automatic Storage Management (ASM)

- Integrated LVM and File System for the Oracle RDBMS
- Introduced in Oracle 10g
- **Stripe And Mirror Everything** architecture
- ASM stripes files, not disks
- ASM mirrors files, not disks
- Integral part of Oracle Grid Infrastructure
- Fully cluster aware
- ASM supports OCR and Voting Disks on ASM11gR2+
- Used in many Oracle Engineered Systems
- Foundation for Oracle Cloud File System (ACFS)

ASM Architecture

- ASM Instance (or instances when clustered)
 - Manages ASM metadata
 - Not in client I/O path
- ASM Disk Group
 - Collection of ASM disks
 - Must have one, but can have many
 - Three redundancy levels
- ASM Clients
 - Oracle Database
 - Oracle Clusterware
 - Oracle Cloud File System (ACFS)

ASM Storage Hierarchy

- Physical storage hierarchy
 - ASM Disk Group: Unit of administration
 - ASM Failure Group: Collection of ASM disks that share a path
 - ASM Disk: a physical disk, partition, LUN or Exadata Grid Disk
- Logical storage hierarchy
 - ASM File
 - ASM Extent
 - ASM Allocation Unit (AU)

ASM Disk Group Redundancy

- External Redundancy
 - Redundancy provided outside ASM (if any)
 - No failure groups, just ASM disks
- Normal Redundancy
 - ASM extents are stored in two different failure groups
 - Requires two or more failure groups
 - Can tolerate the loss of one failure group
- High Redundancy
 - ASM extents are stored in three different failure groups
 - Requires three or more failure groups
 - Can tolerate the loss of two failure groups simultaneously

ASM Disk Group Metadata

- V\$ASM_DISKGROUP
 - (...)
 - TYPE
 - TOTAL_MB
 - FREE_MB
 - REQUIRED_MIRROR_FREE_MB
 - USABLE_FILE_MB

- V\$ASM_DISK
 - (...)
 - TOTAL_MB
 - FREE_MB
 - FAILGROUP

ASM Disk Group Metadata (Cont.)

- V\$ASM_ATTRIBUTE
 - (...)
 - DISK_REPAIR_TIMER
 - FAILGROUP_REPAIR_TIMER
- V\$ASM_FILE
 - (...)
 - REDUNDANCY_LOWERED

ASM Disk Group Error Handling

- Single ASM disk failures
 - ASM10g: Failed disk will be dropped immediately
 - ASM11g: Failed disk will be dropped after `DISK_REPAIR_TIMER` expires
 - ASM attempts to restore redundancy
 - Disk Group is forcibly dismounted if too many disks fail

ASM Disk Group Error Handling (Cont.)

- ASM Failure Group failures
 - ASM10g: Failed failure group will be dropped immediately
 - ASM11g: Failed failure group will be dropped after `DISK_REPAIR_TIMER` expires
 - ASM12c: Failed failure group will be dropped after `FAILGROUP_REPAIR_TIMER` expires
 - ASM attempts to restore redundancy
 - Disk Group is forcibly dismounted if too many failure groups fail

ASM Disk Group Space Administration

- `REQUIRED_MIRROR_FREE_MB`
 - Must be available to fully restore redundancy
 - Value depends on disk group redundancy
 - Normal redundancy: size of largest failure group
 - High redundancy: size of two largest failure groups
 - Is NOT enforced!

- `USABLE_FILE_MB`

$$\text{USABLE_FILE_MB} = \frac{\text{FREE_MB} - \text{REQUIRED_MIRROR_FREE_MB}}{\text{REDUNDANCY}}$$

- Amount of usable disk group capacity
- If negative, reduced redundancy will occur after worst case failure

Choosing the number of ASM Failure Groups

- Available hardware should drive this decision
- Single controller: each disk = failure group
- Multiple controllers: all disks on same controller = failure group
- Exadata: all single Storage Server grid disks = failure group

ASM on Exadata Storage

- Single disk failures
 - Failed disk will be dropped immediately
- `REQUIRED_MIRROR_FREE_MB` with *ASM11g*
 - Normal redundancy: size of largest failure group
 - High redundancy: size of two largest failure groups
- `REQUIRED_MIRROR_FREE_MB` with *ASM12c*
 - Normal redundancy: size of largest disk
 - High redundancy: size of two largest disks

Live Demo!

Q U E S T I O N S
&
A N S W E R S

And Finally

Thank you for your kind attention!

For a copy of my demo script email me at:

Harald.van.Breederode@oracle.com

Remember: prutser.wordpress.com