Key Technical Features of Oracle RAC 12c

Ian Cookson
Product Manager – Oracle Clusterware
Oct 25, 2017

Safe Harbor Statement

The following is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described for Oracle's products remains at the sole discretion of Oracle.

Key Technical Features of Oracle RAC 12c

- **Applied Machine Learning**
  - Autonomous Health Framework (AHF); specifically Cluster Health Advisor – “Sherlock” for the cloud
- **Smart Reconfiguration**
  - Node Weighting lets most of the workload survive; Recovery Buddy
- **Massive Scaling**
  - Pluggable Database Isolation; Service-oriented Buffer Cache Access; Flex Cluster
- **Fleet Management**
  - Cluster Domains and the Domain Services Cluster; Rapid Home Provisioning
- **Database-Oriented Storage Management**
  - ASM Flex Diskgroups

Key Technical Features of Oracle RAC 12c

- **Applied Machine Learning**
  - *Autonomous Health Framework (AHF); specifically Cluster Health Advisor – "Sherlock" for the cloud*
- **Smart Reconfiguration**
  - *Node Weighting lets most of the workload survive*; Recovery Buddy
- **Massive Scaling**
  - Pluggable Database Isolation; Service-oriented Buffer Cache Access; Flex Cluster
- **Fleet Management**
  - *Cluster Domains and the Domain Services Cluster*; Rapid Home Provisioning
- **Database-Oriented Storage Management**
  - ASM Flex Diskgroups

Oracle Autonomous Health Framework

Powered by Applied Machine Learning

- Integrates next generation tools running 24/7
- Discovers Potential Issues and Notifies with Corrective Actions
- Speeds Issue Diagnosis and Recovery
- Preserves Database and Server Availability and Performance
- Autonomously Monitors and Manages resources to maintain SLAs

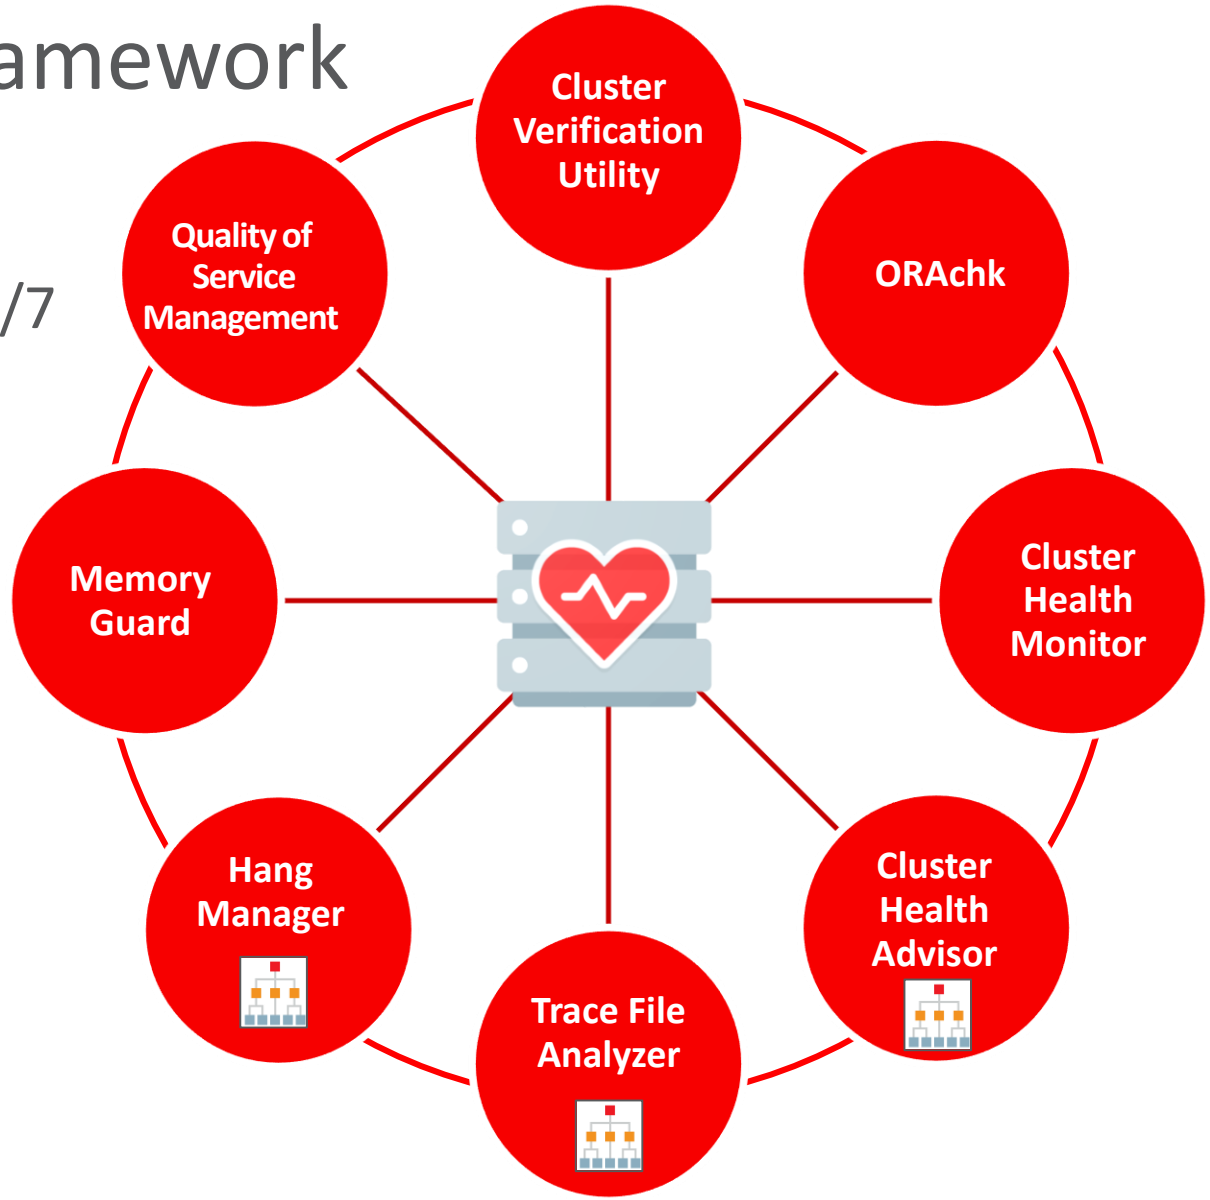

5

Cluster Health Advisor (CHA) Architecture Overview

- Monitors in real-time Oracle database* systems and their hosts
- Detects early impending as well as ongoing system faults
- Diagnoses and identifies the most likely root causes
- Provides targeted actions for prevention or escalation of DB/server problems
- Generates relevant alerts and notifications for rapid response

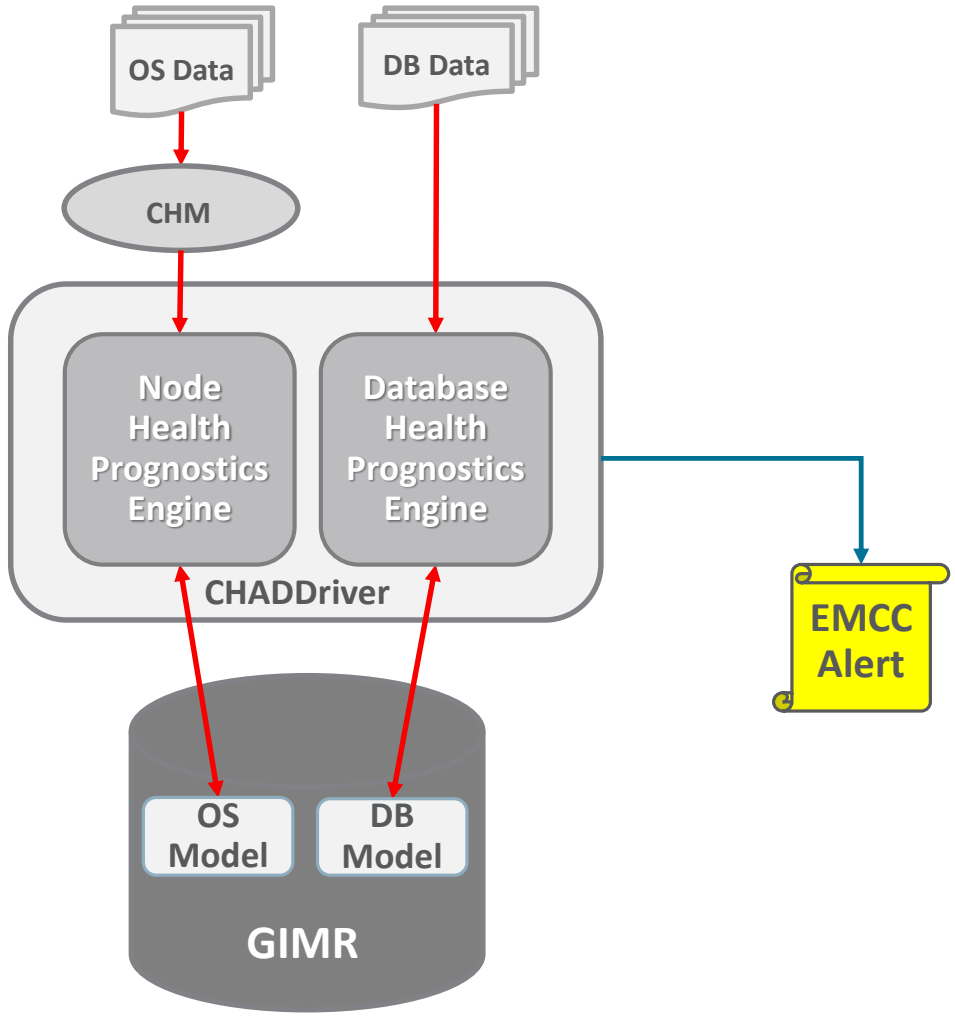

*Oracle RAC/R1N databases only

Incident Manager

Page Refreshed **Sep 29, 2017 12:08:13 PM PDT**

ASM Cluster-wide Disk Utilization on Host rwsbi06 Database/Cluster rwsbi0508-mb2 Instance . The Cluster Health Advisor (CHA) detected slower tl...

Open in new tab

Unassigned, Not acknowledged

Incident Details

ID	766
Metric	Alert Level
Metric Group	CHA Alerts
Key	CHA_INCIDENT_STATE_CHANGE_CLUSTERWARE_rwsbi0508-mb2__CHA_...
Target	rwsbi0508-mb2 (Cluster)
Incident Created	Sep 29, 2017 7:06:45 PM GMT
Last Updated	Sep 29, 2017 7:06:45 PM GMT
Summary	ASM Cluster-wide Disk Utilization on Host rwsbi06 Database/Cluster rwsbi0508-mb2 Instance . The Cluster Health Advisor (CHA) detected slower than expected disk performance because the high disk I/O demand from the other servers increased the utilization of the shared disks. Review the CHA findings and corrective actions from the other servers and database instances in the cluster for IO issues. Add disks to the database disk groups.
Internal Event Name	cha_alerts:cha_alert_level
Event Type	Metric Alert
Category	Unclassified

Show internal values for attributes ...

Metric Data

Critical Threshold	Not Applicable
Warning Threshold	Not Applicable
Number of Occurrences	0
Last Known Value	Critical

Tracking

Acknowledge | Add Comment ... | Manage ... | More

Escalated	No	Owner	-
Priority	None	Acknowledged	No
Status	New		

Last Comment Incident created by rule (Name = Incident management rule set for all targets, Create incident for critical metric alerts [System generated rule]).: on Sep 29, 2017 7:06:45 PM GMT

This incident will be automatically cleared when the underlying issue is resolved.

Guided Resolution

Diagnostics	Actions	Corrective Actions
Problem Analysis	Edit Thresholds	No corrective action defined.
View Metric Help		Add corrective action

Cluster Health Advisor

Problem	The degradation is caused by a higher than expected utilization of shared storage devices for this database. No evidence of significant increase in I/O demand on the local node.
Confidence	95.17%
Action	Validate whether there is increase in I/O demand on other nodes than the local and find I/O intensive SQL . Add more disks to disk group or move database to faster disks.

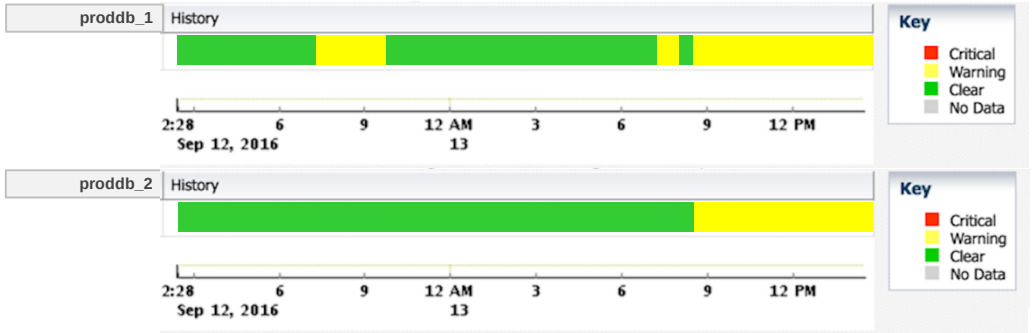

Rapid Recovery with TFA

Smart Collection with TFA Collector

- Always on
- Collects comprehensive first failure diagnostics on each node
- Filters and packages relevant diagnostic data using Applied ML model
- Automatically notifies DBAs and Sys Admins of events
- Transfers data to centralized storage for detailed analysis with TFA Receiver
- Optionally allows quick issue resolution with Oracle Support

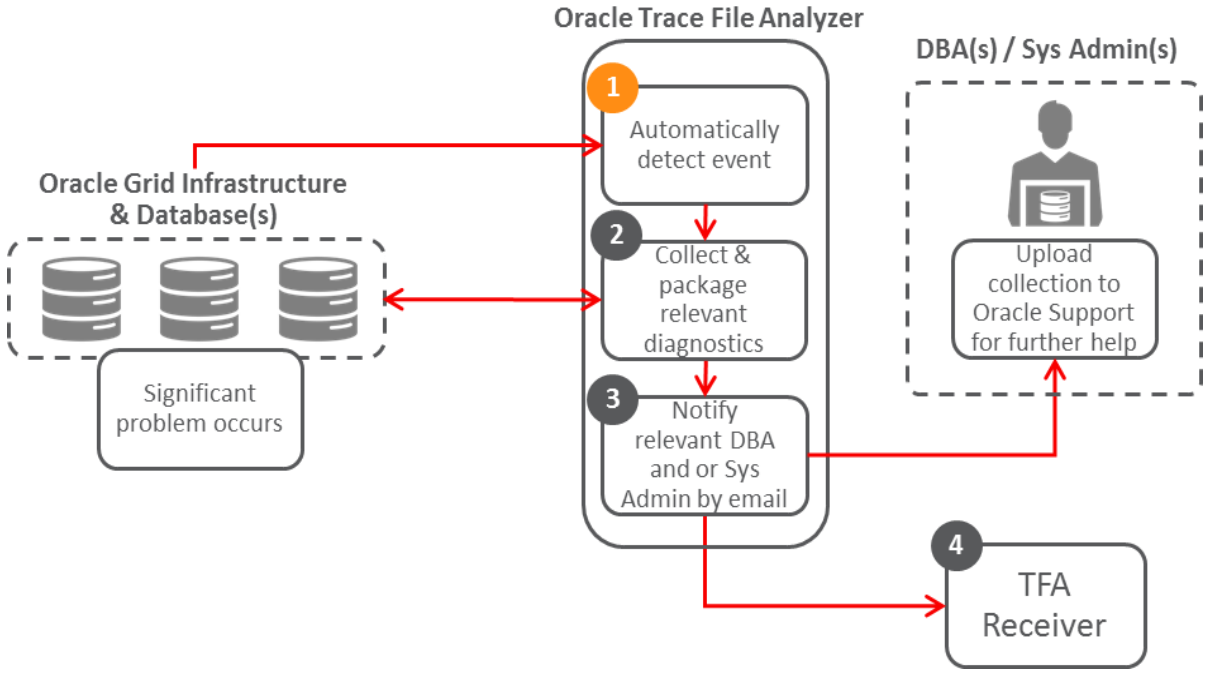

Oracle 12c Hang Manager

Autonomously Preserves Database Availability and Performance

- Always on - Enabled by default
- Reliably detects database hangs and deadlocks
- Autonomously resolves them
- Supports QoS Performance Classes, Ranks and Policies to maintain SLAs
- Logs all detections and resolutions
- New SQL interface to configure sensitivity (Normal/High) and trace file sizes

NEW IN 12.2

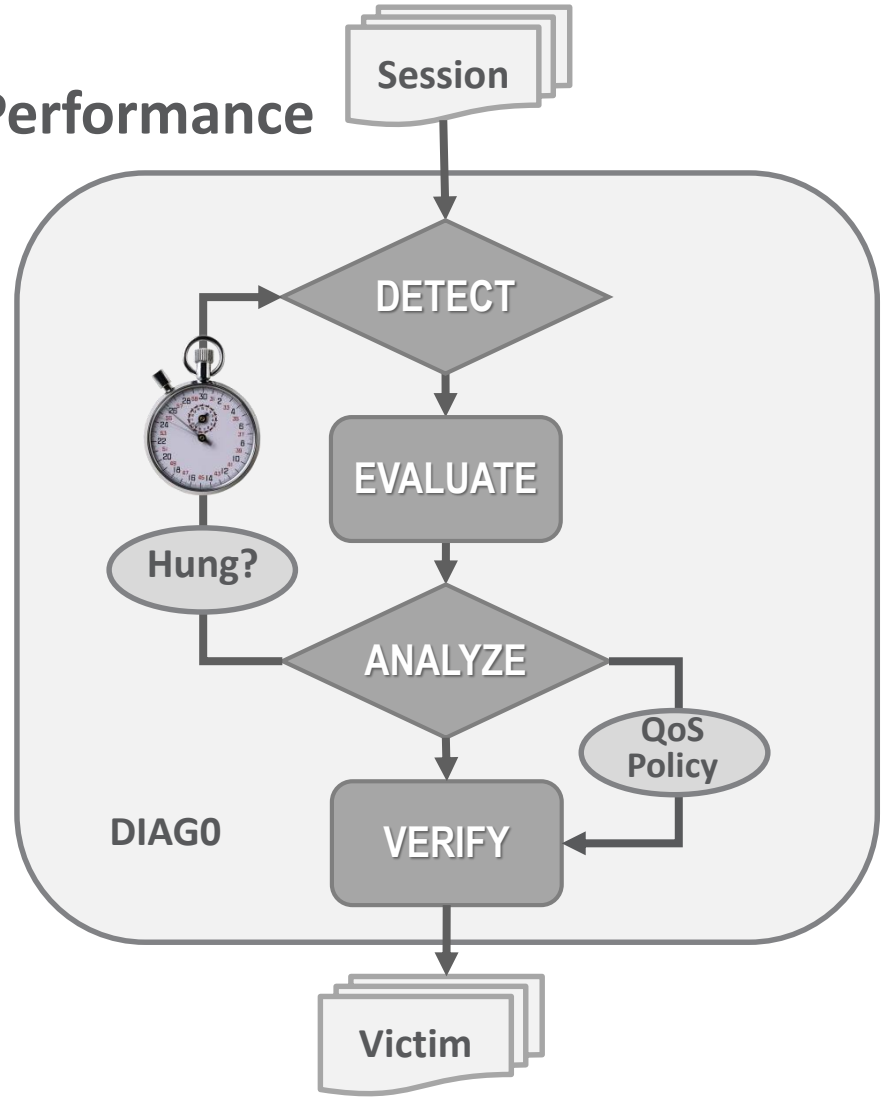

Key Technical Features of Oracle RAC 12c

- **Applied Machine Learning**
  - *Autonomous Health Framework (AHF); specifically Cluster Health Advisor – "Sherlock" for the cloud*
- **Smart Reconfiguration**
  - *Node Weighting lets most of the workload survive*; Recovery Buddy
- **Massive Scaling**
  - Pluggable Database Isolation; Service-oriented Buffer Cache Access; Flex Cluster
- **Fleet Management**
  - *Cluster Domains and the Domain Services Cluster*; Rapid Home Provisioning
- **Database-Oriented Storage Management**
  - ASM Flex Diskgroups

Split Brain – What Does It Mean for Oracle Clusterware?

"a condition in which Oracle Clusterware believes that there is a communication failure between nodes"

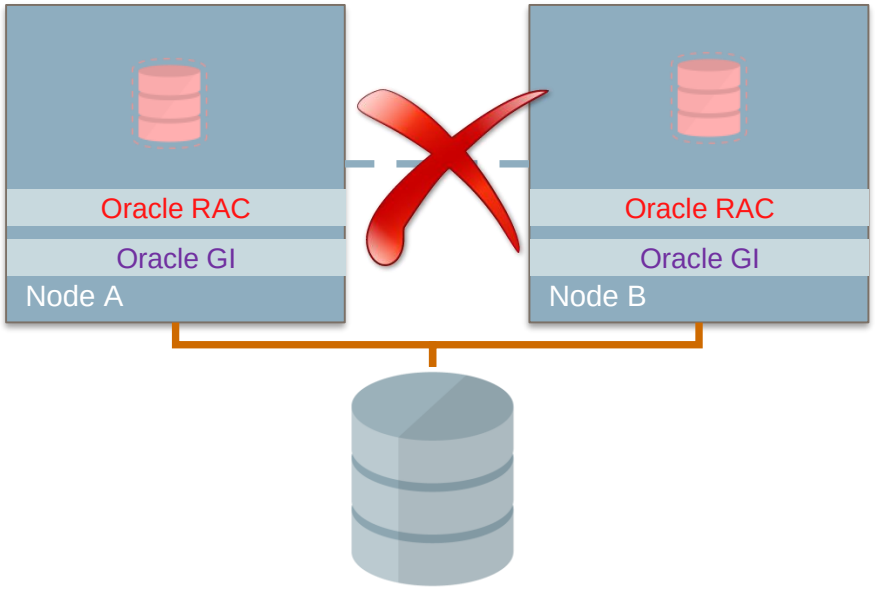

Node Eviction Basics

http://www.slideshare.net/MarkusMichalewicz/oracle-clusterware-node-management-and-voting-disks

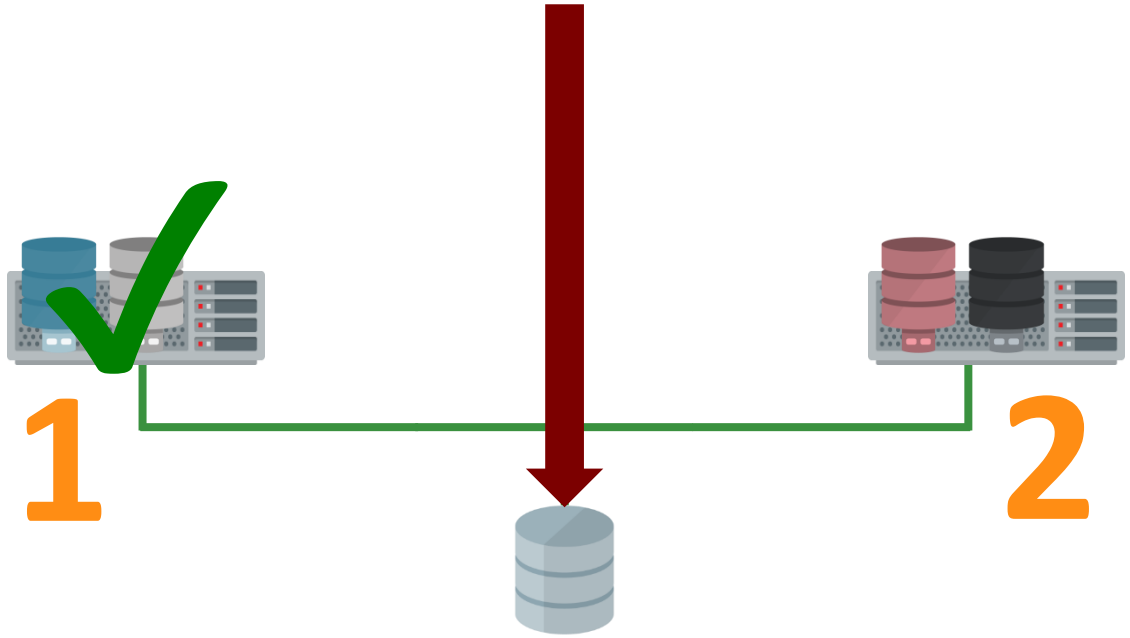

- Pre-12.2, node eviction follows a rather simplistic algorithm
  - Example in a 2-node cluster: The node with the lowest node number survives.
- Customers must not base their application logic on which node survives the split brain.
  - As this may(!) change in future releases

Node Weighting in Oracle RAC 12*c* Release 2

Idea: If *Everything is equal*, let the majority of work survive

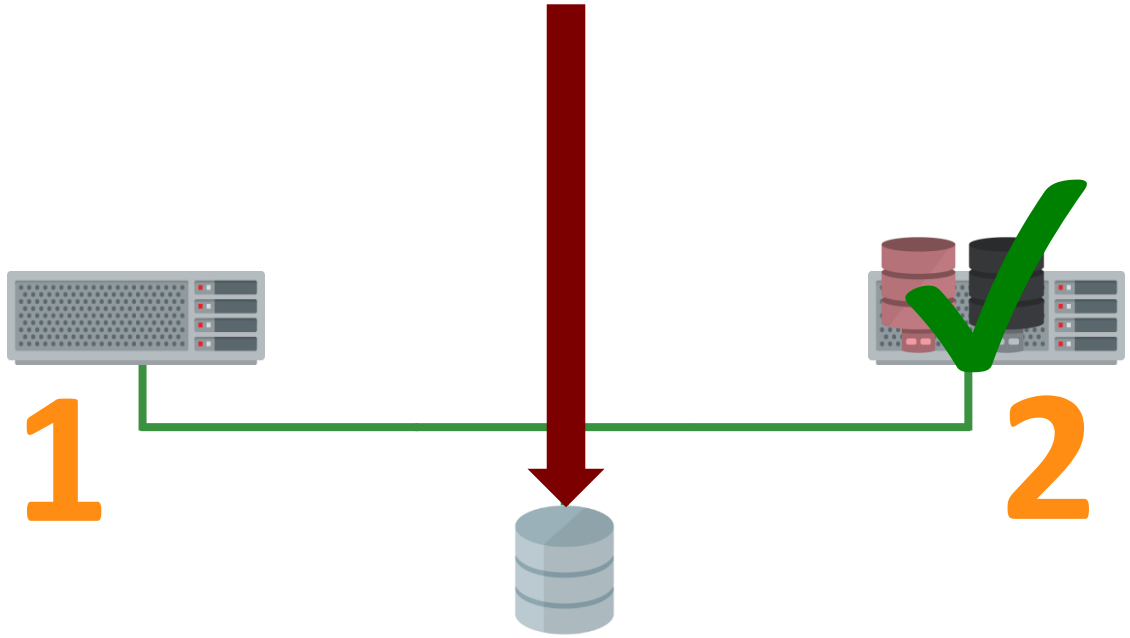

- Node Weighting is a new feature that considers the workload hosted in the cluster during fencing
- The idea is to let the majority of work survive, if *everything else is equal*
  - Example: In a 2-node cluster, the node hosting the majority of services (at fencing time) is meant to survive

Headache-free Split Brain Resolution

Split Brain Resolution in Oracle Clusterware 12c Rel 2
If Everything Else is Equal…

1. Customer can designate which server(s) and resource(s) are **critical**
2. Clusterware will evaluate cluster resources on **implied workload**
3. Cluster cohort containing the **lowest cluster node number**

Let’s Define “Equal”

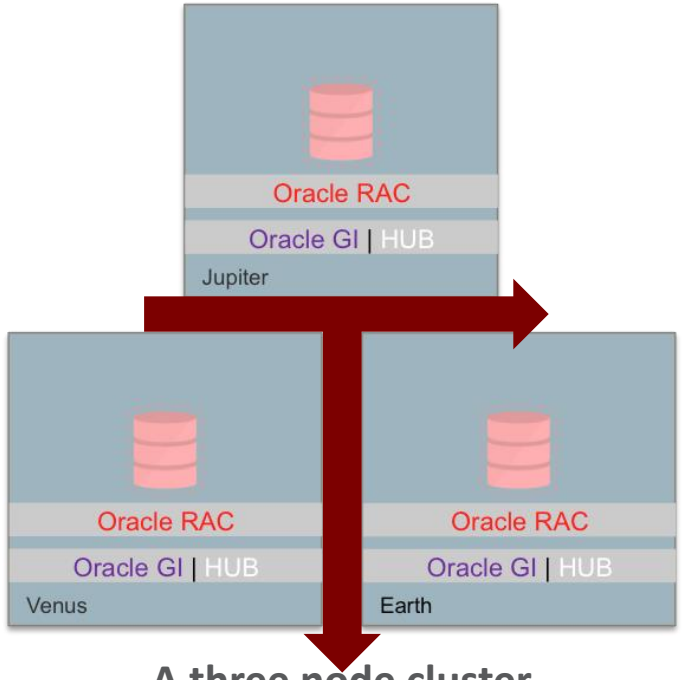

A three node cluster
will benefit from “Node Weighting”, if three equally sized sub-clusters are built as s result of the failure, since two differently sized sub-clusters are not equal.

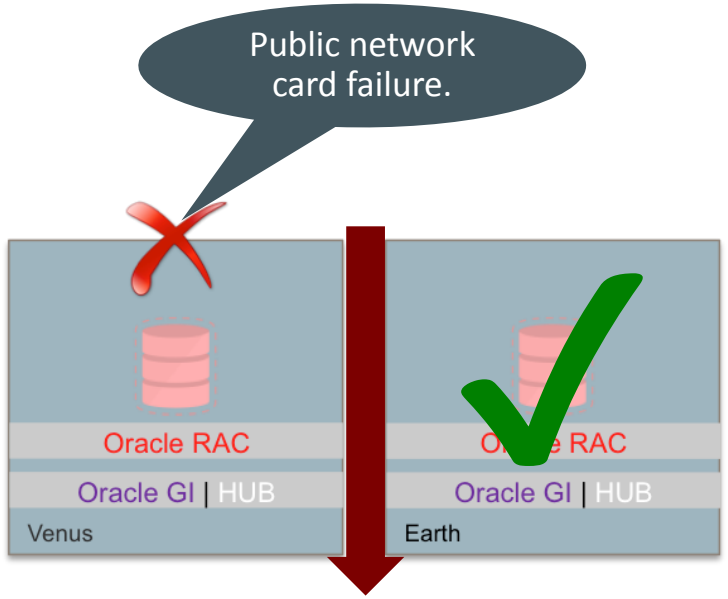

Secondary failure consideration can influence which node survives. Secondary failure consideration will be enhanced successively.

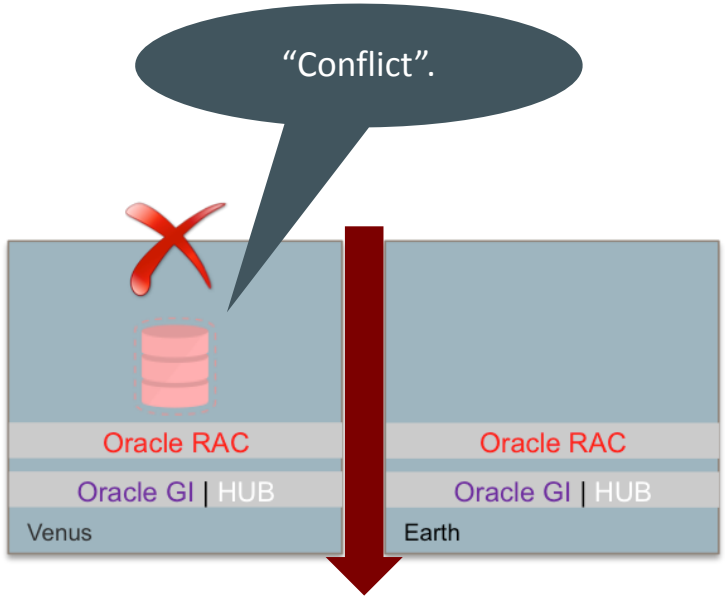

A fallback scheme
is applied if considerations do not lead to an actionable outcome.

CSS_CRITICAL – Fencing with Manual Override

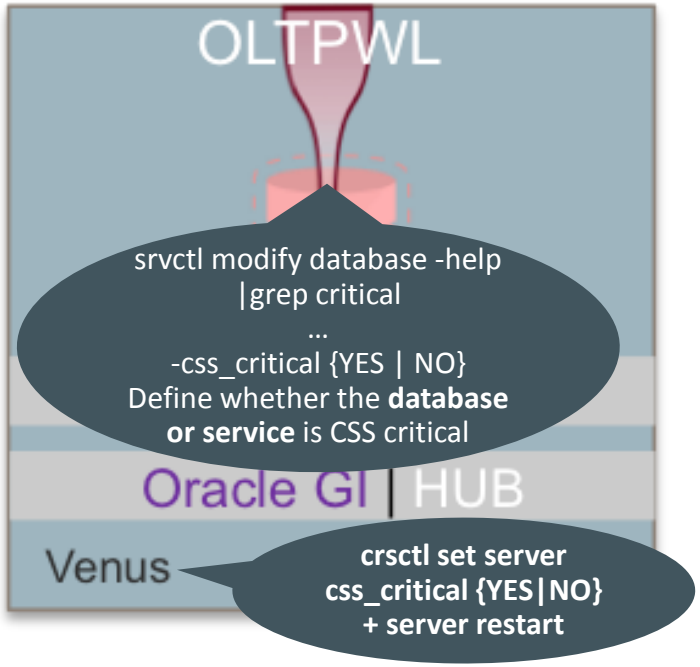

CSS_CRITICAL
can be set on various levels / components to mark them as "critical" so that the cluster will try to preserve them in case of a failure.

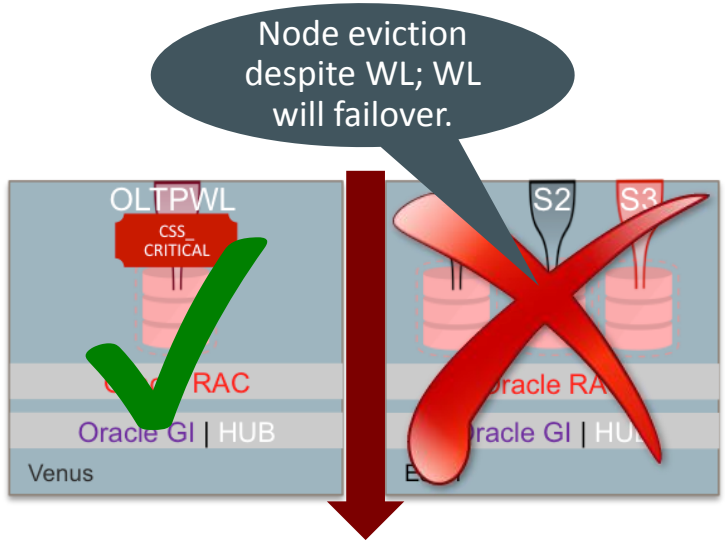

CSS_CRITICAL will be honored
if no other technical reason prohibits survival of the node which has at least one critical component at the time of failure.

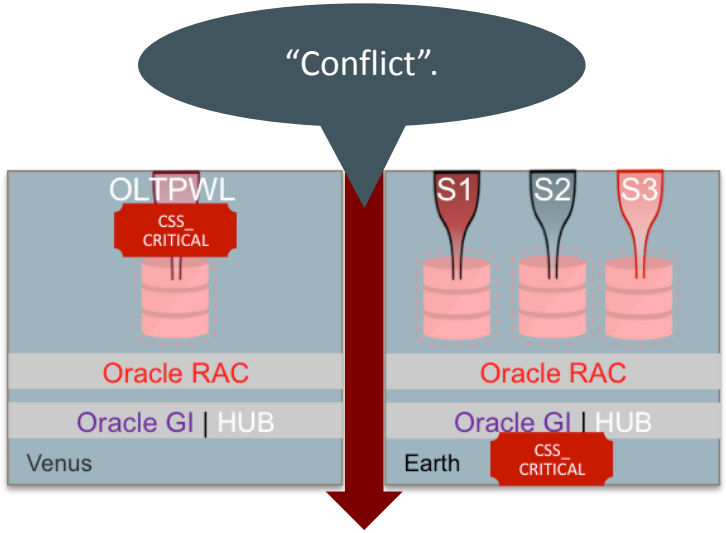

A fallback scheme is applied if CSS_CRITICAL settings do not lead to an actionable outcome.

Key Technical Features of Oracle RAC 12c

- **Applied Machine Learning**
  - *Autonomous Health Framework (AHF); specifically Cluster Health Advisor – "Sherlock" for the cloud*
- **Smart Reconfiguration**
  - *Node Weighting lets most of the workload survive*; Recovery Buddy
- **Massive Scaling**
  - Pluggable Database Isolation; Service-oriented Buffer Cache Access; Flex Cluster
- **Fleet Management**
  - *Cluster Domains and the Domain Services Cluster*; Rapid Home Provisioning
- **Database-Oriented Storage Management**
  - ASM Flex Diskgroups

Fleet Management

- More efficient and intelligent management ***of large cluster estates***
- Based on Cluster Domains and the Domain Services Cluster
- Optimization:
  - Oracle's Autonomous Health Framework (AHF) powered by Oracle Machine Learning
  - Rapid Home Provisioning

Oracle RAC 12*c* Rel. 2 Cluster Domain

Centralized Management for Cluster Estates "too big to manage" otherwise

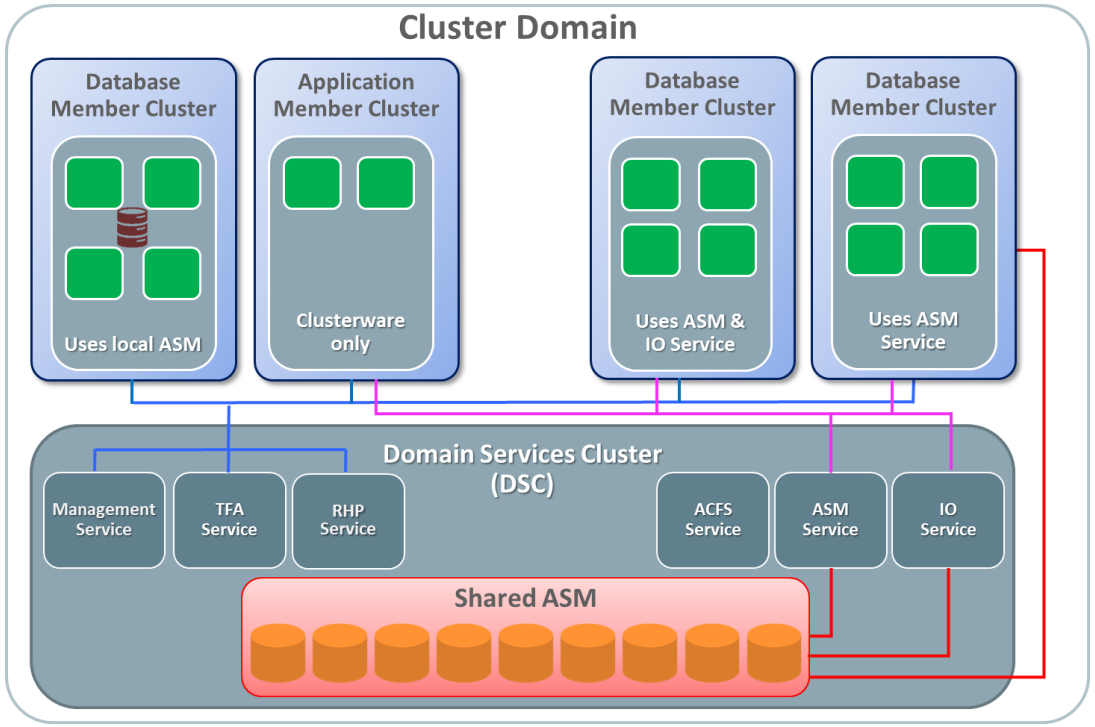

- Simplified Management
  - Fleet Management for installation, update, patching and maintenance
- Reduced Local Overhead
  - Member Clusters benefit from the consolidation of common services on the Domain Services Cluster
- Improved IO Performance
  - Utilizing consolidated shared storage

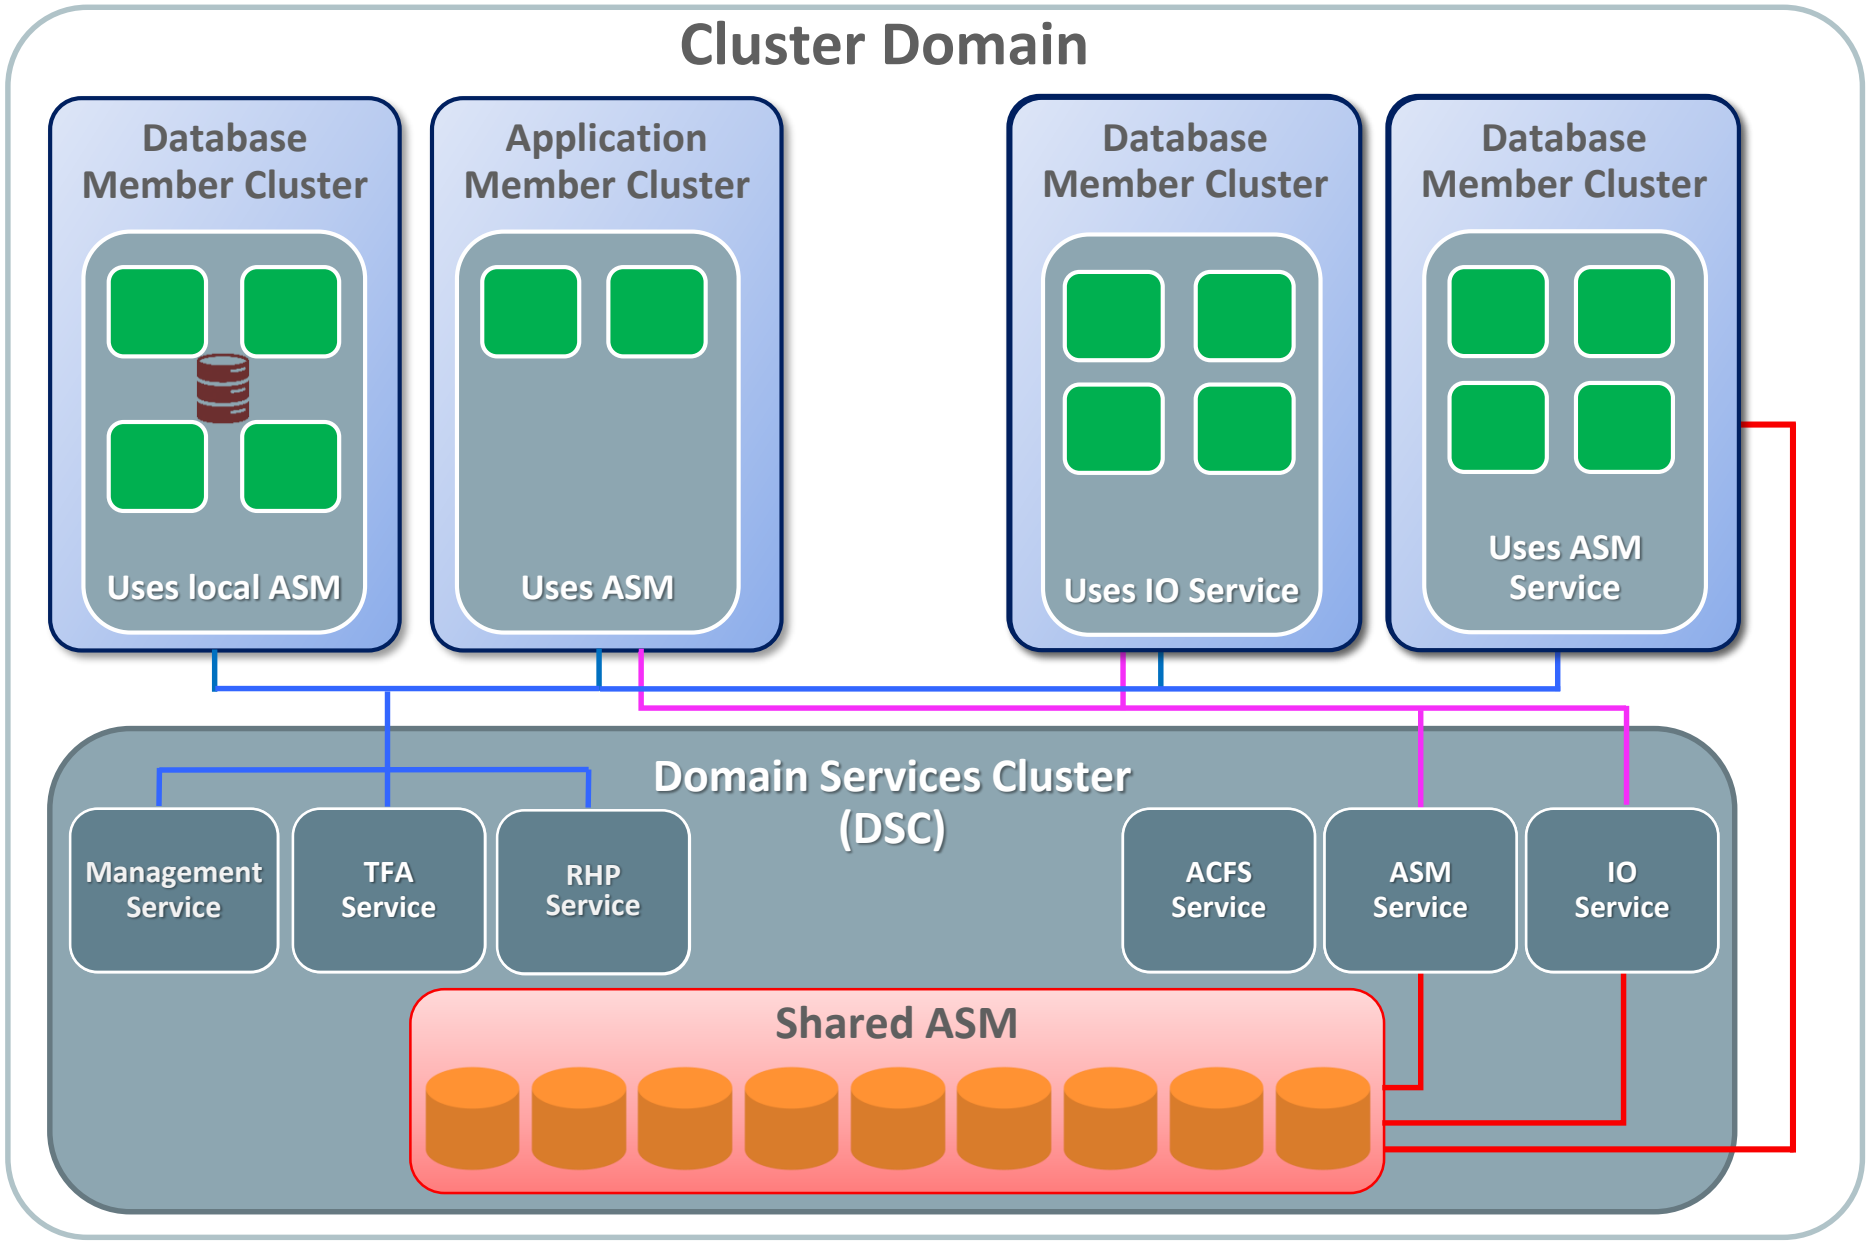

The DSC Management Service

Applied Machine Learning for Database Diagnostics

- Efficient diagnosis using Machine Learning
- Automatically performs corrective actions to prevent possible issues
- Provides simple alerts & recommendations for issues that require manual intervention

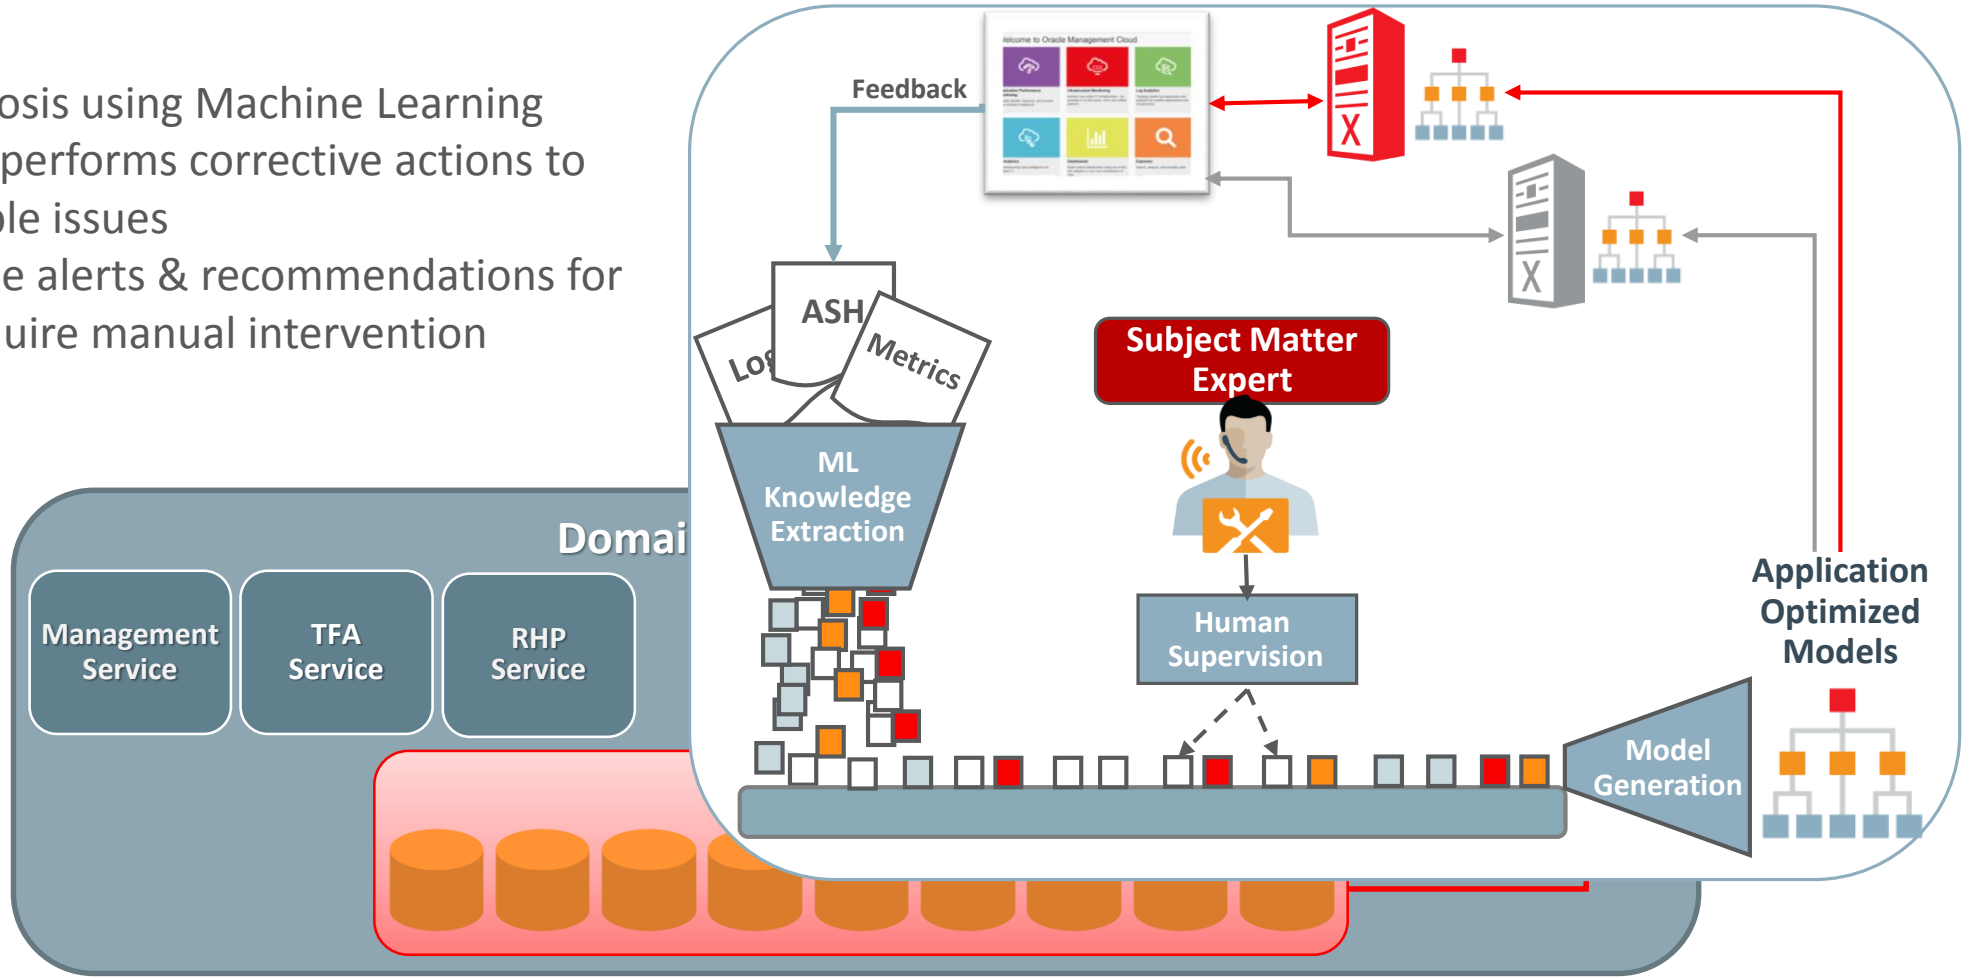

Rapid Home Provisioning Service

Fleet Management for On-Premise Deployments

- Provision new pools onto base machines
- DB and GI: provision, scale, patch, upgrade
- Custom workflow framework
- Notification model
- Audit capabilities

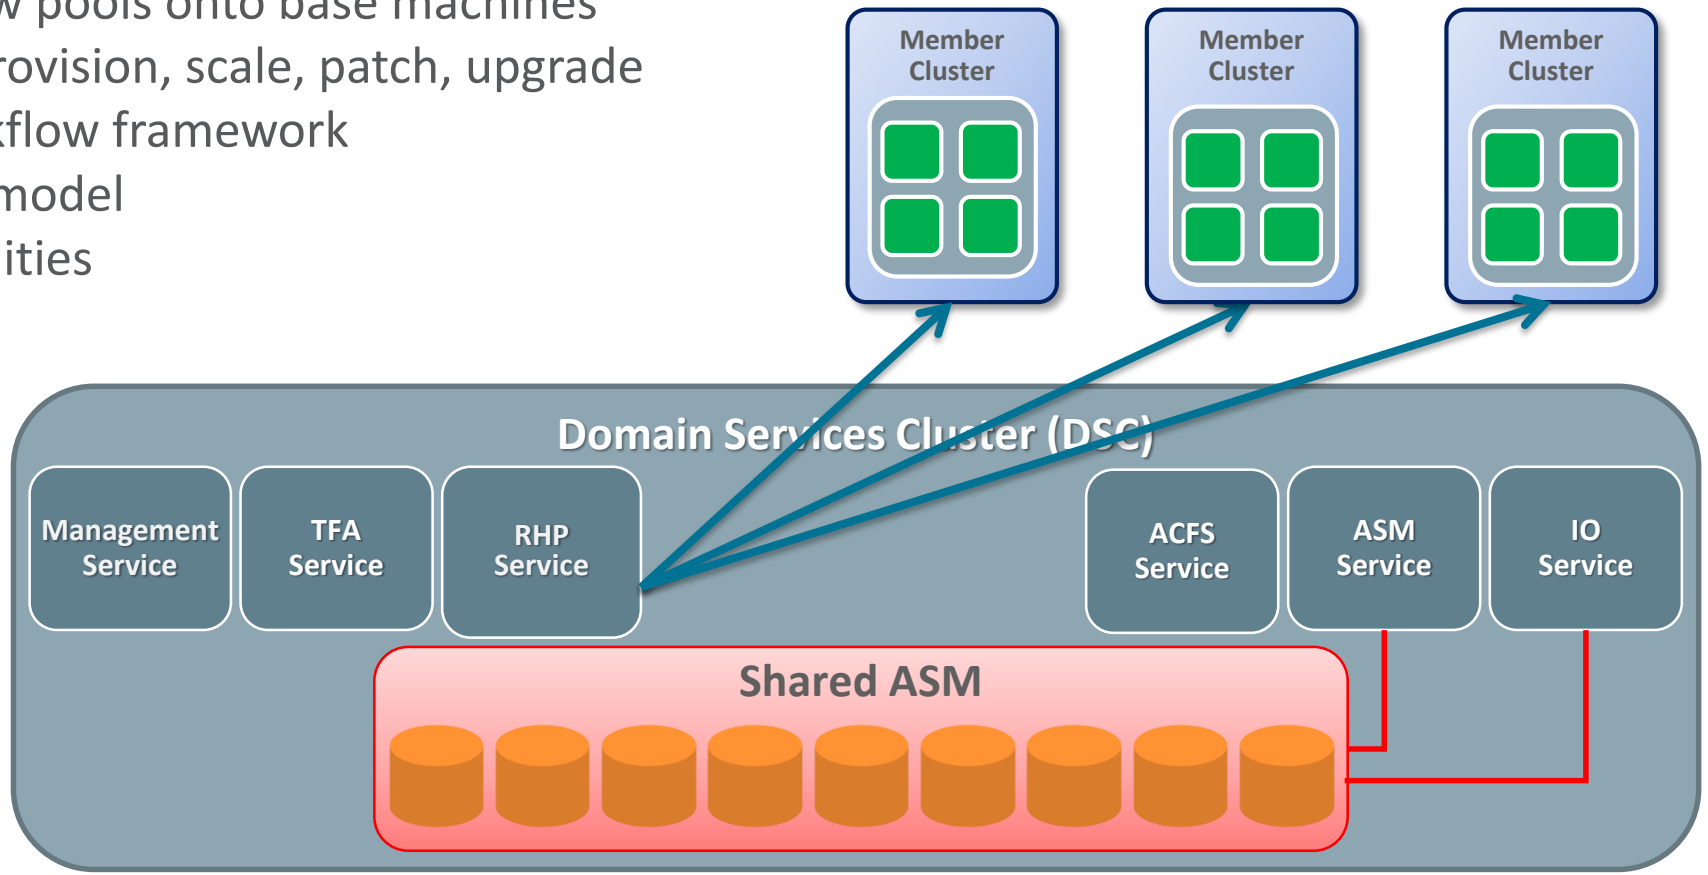

Rapid Home Provisioning Service

Fleet Management in the Oracle Cloud

- Provision new pools onto base machines
- DB and GI: provision, scale, patch, upgrade
- Custom workflow framework
- Notification model
- Audit capabilities

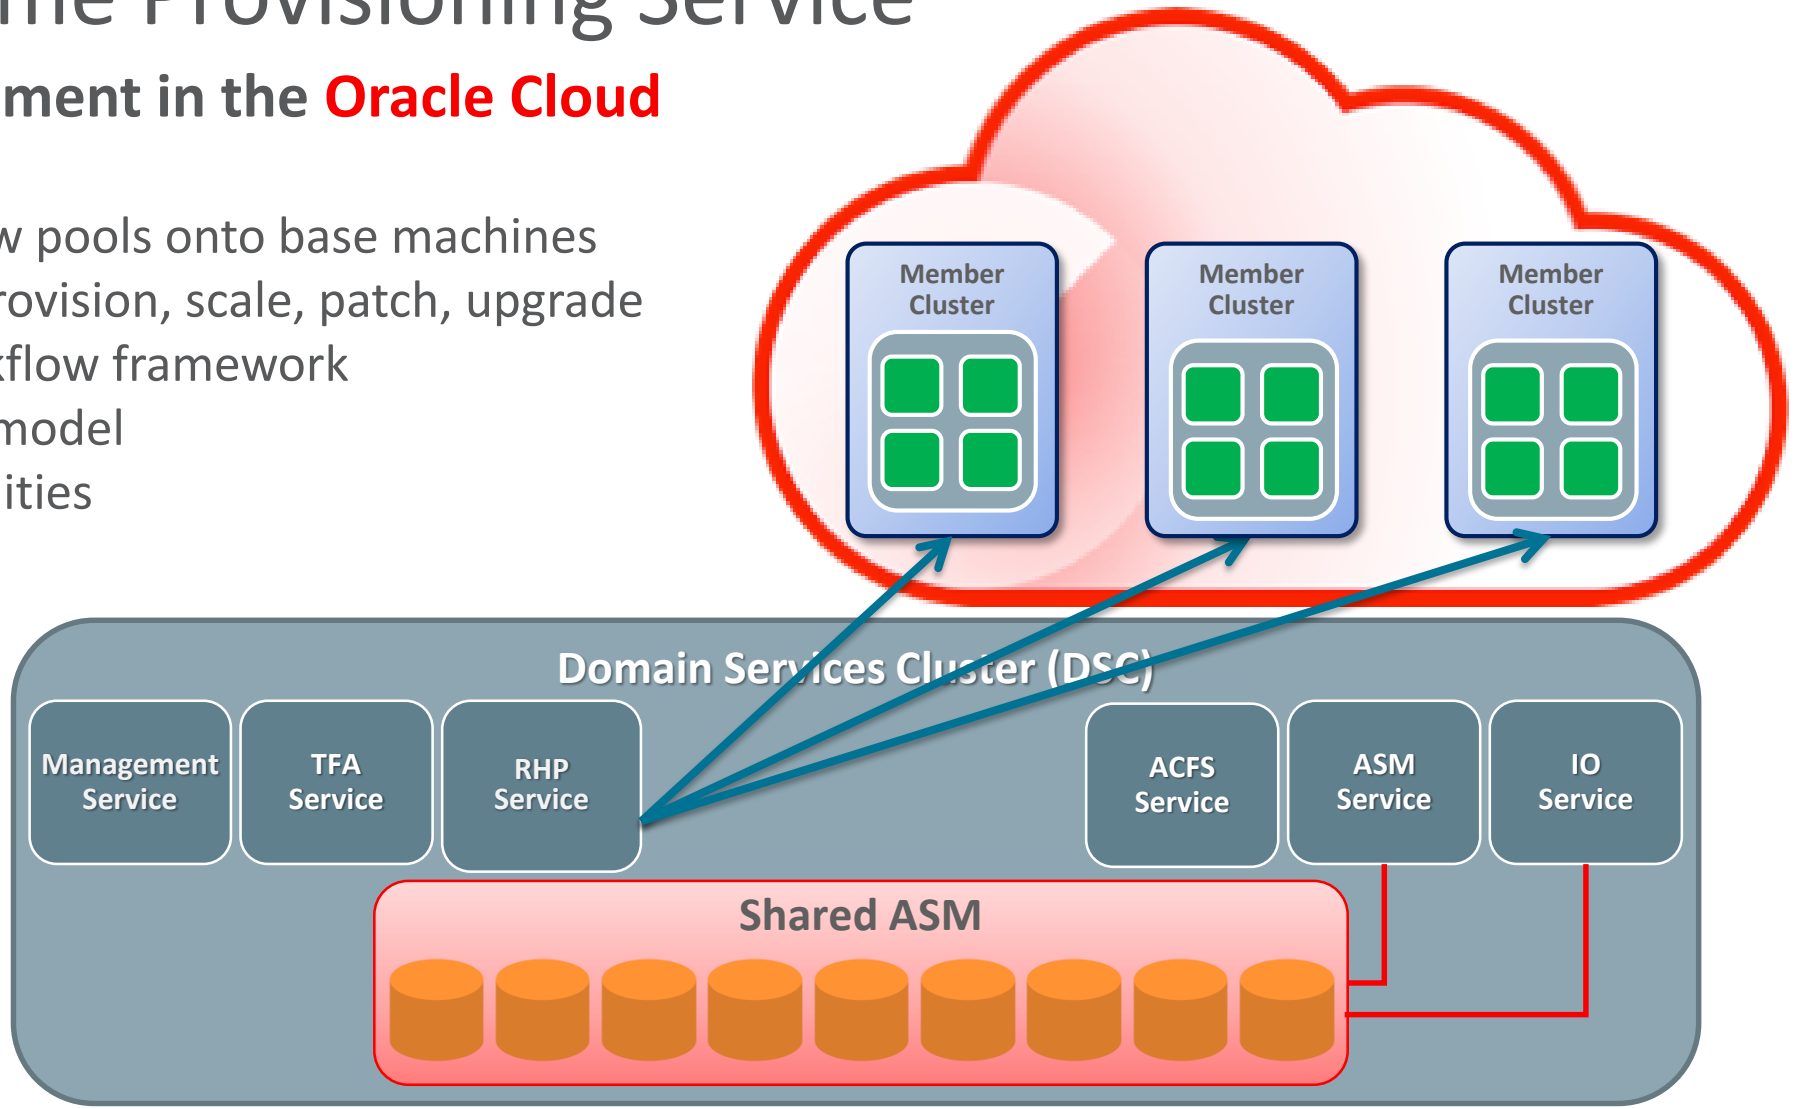

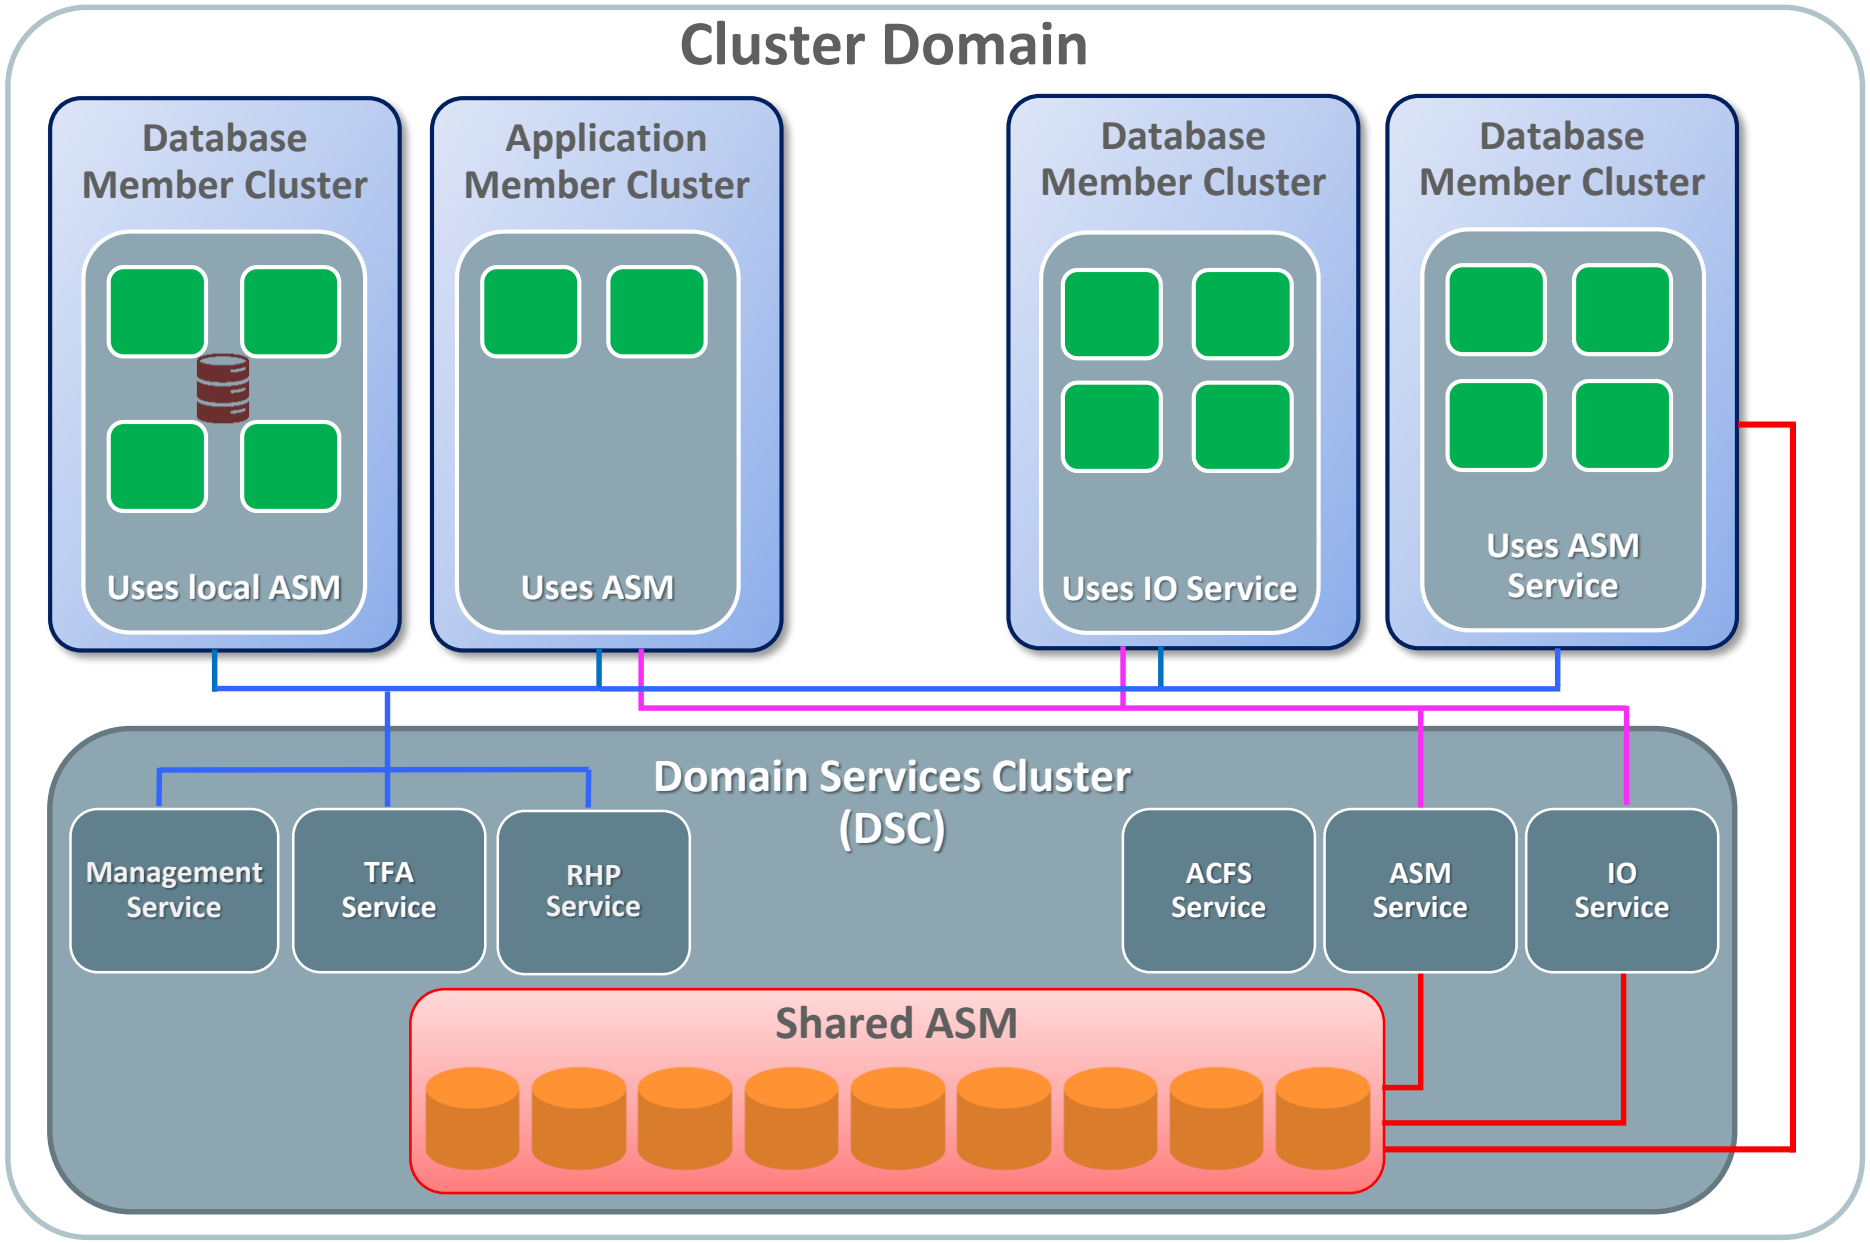

Database Member Cluster with Local ASM

Standalone isolation with reduced local overhead

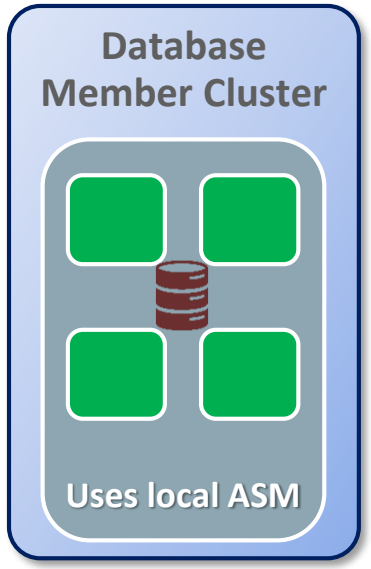

- For databases requiring
  - Full Isolation and performance stability
  - That can benefit from the centralized Management Service on the DSC
- Particularly suitable for unpredictable workloads, or highly variable workloads
- Examples include
  - Business Intelligence and Analytics systems
  - Batch processing systems
  - Response-critical, user-facing systems

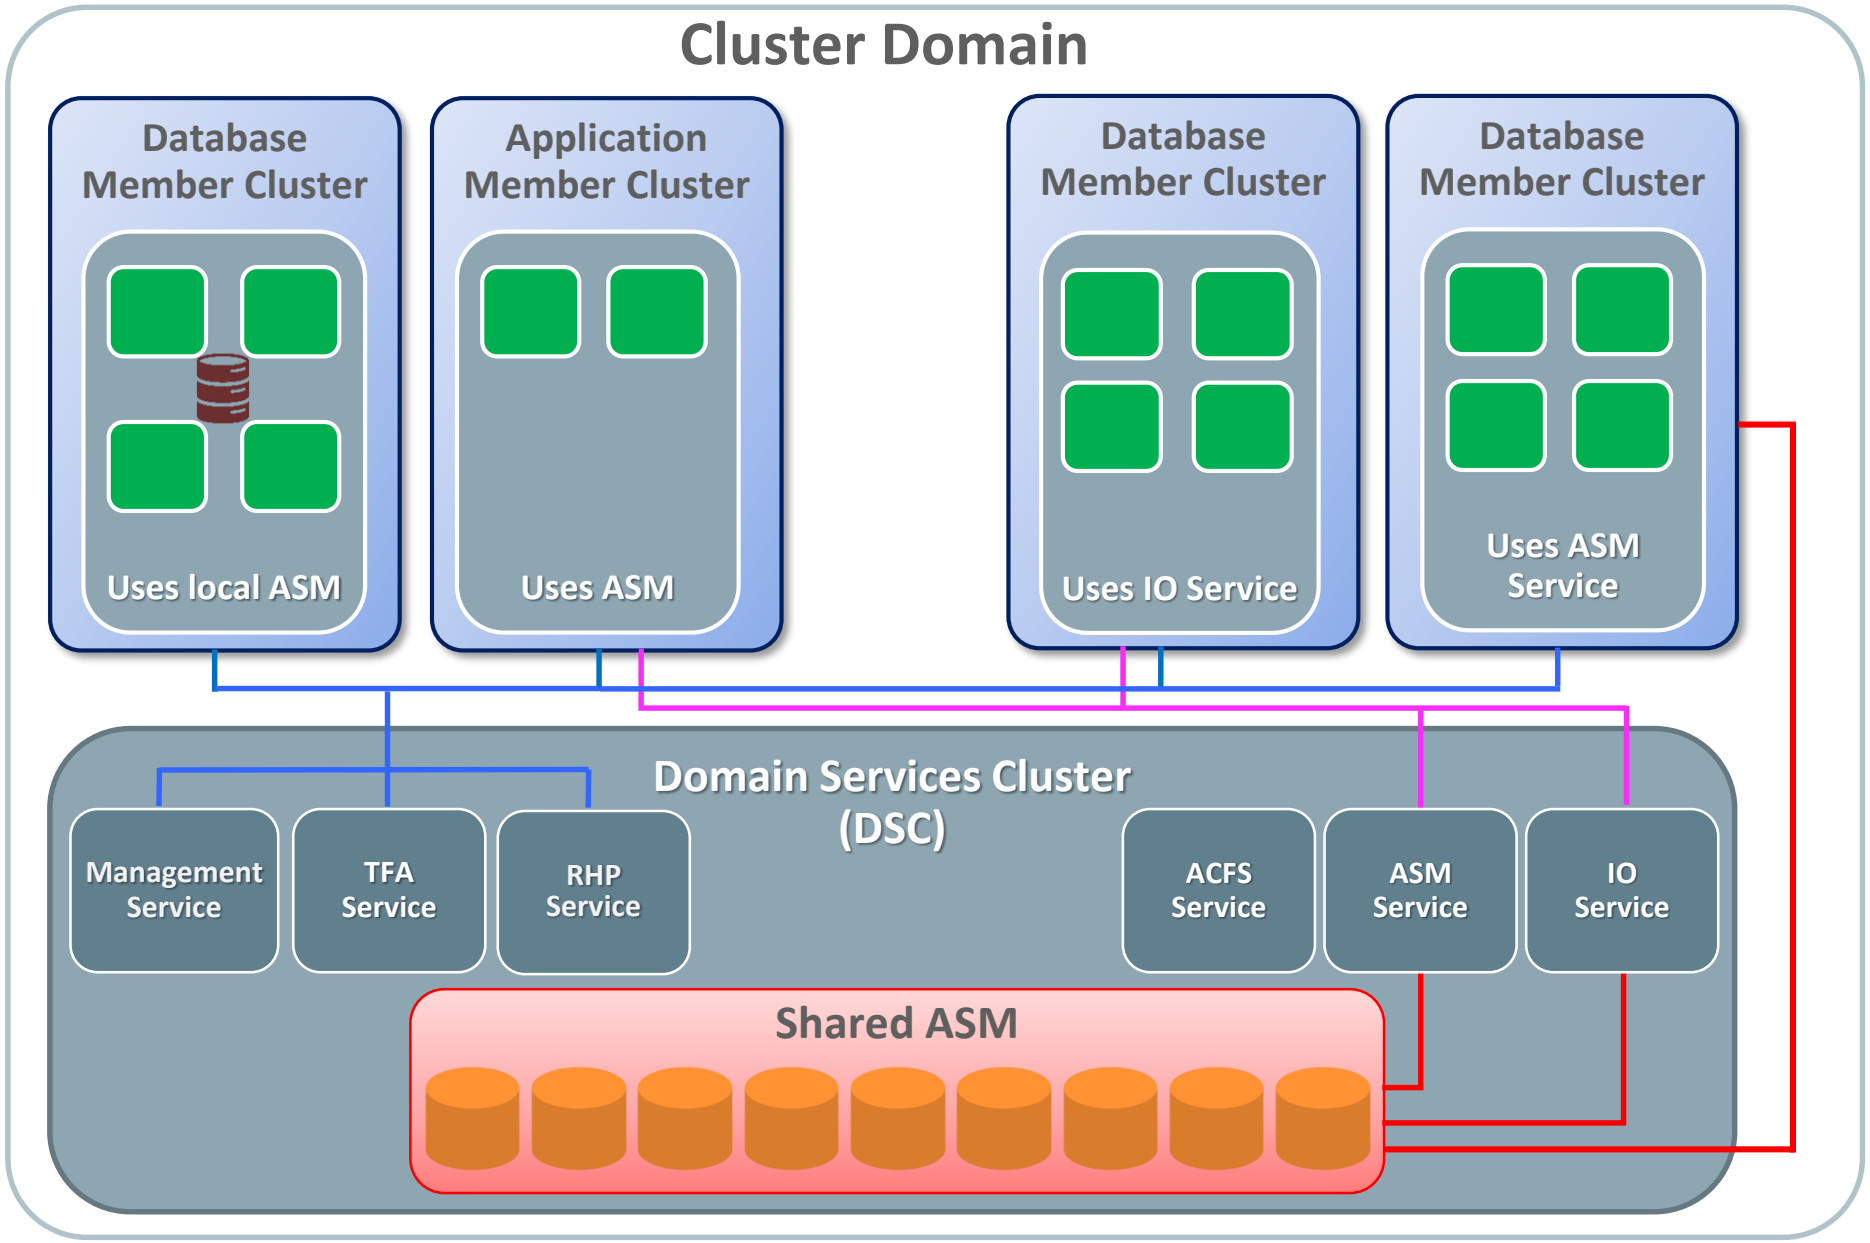

Database Member Cluster Using ASM Service

Standalone isolation benefitting from consolidated shared storage

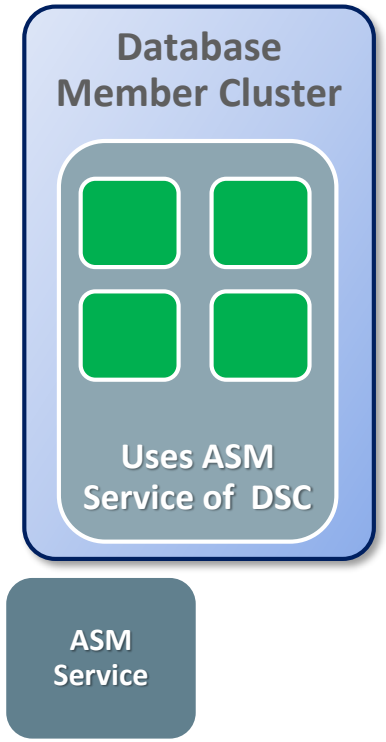

- For databases requiring
  - Isolation and performance stability
  - That can benefit from the centralized Management Service on the DSC
  - And the centralized ASM Storage Management Service on the DSC
- Best suited for workloads for which IO stability is important, but benefit from the centralized ASM Services on the DSC
- Examples include
  - OLTP systems
  - Reporting systems

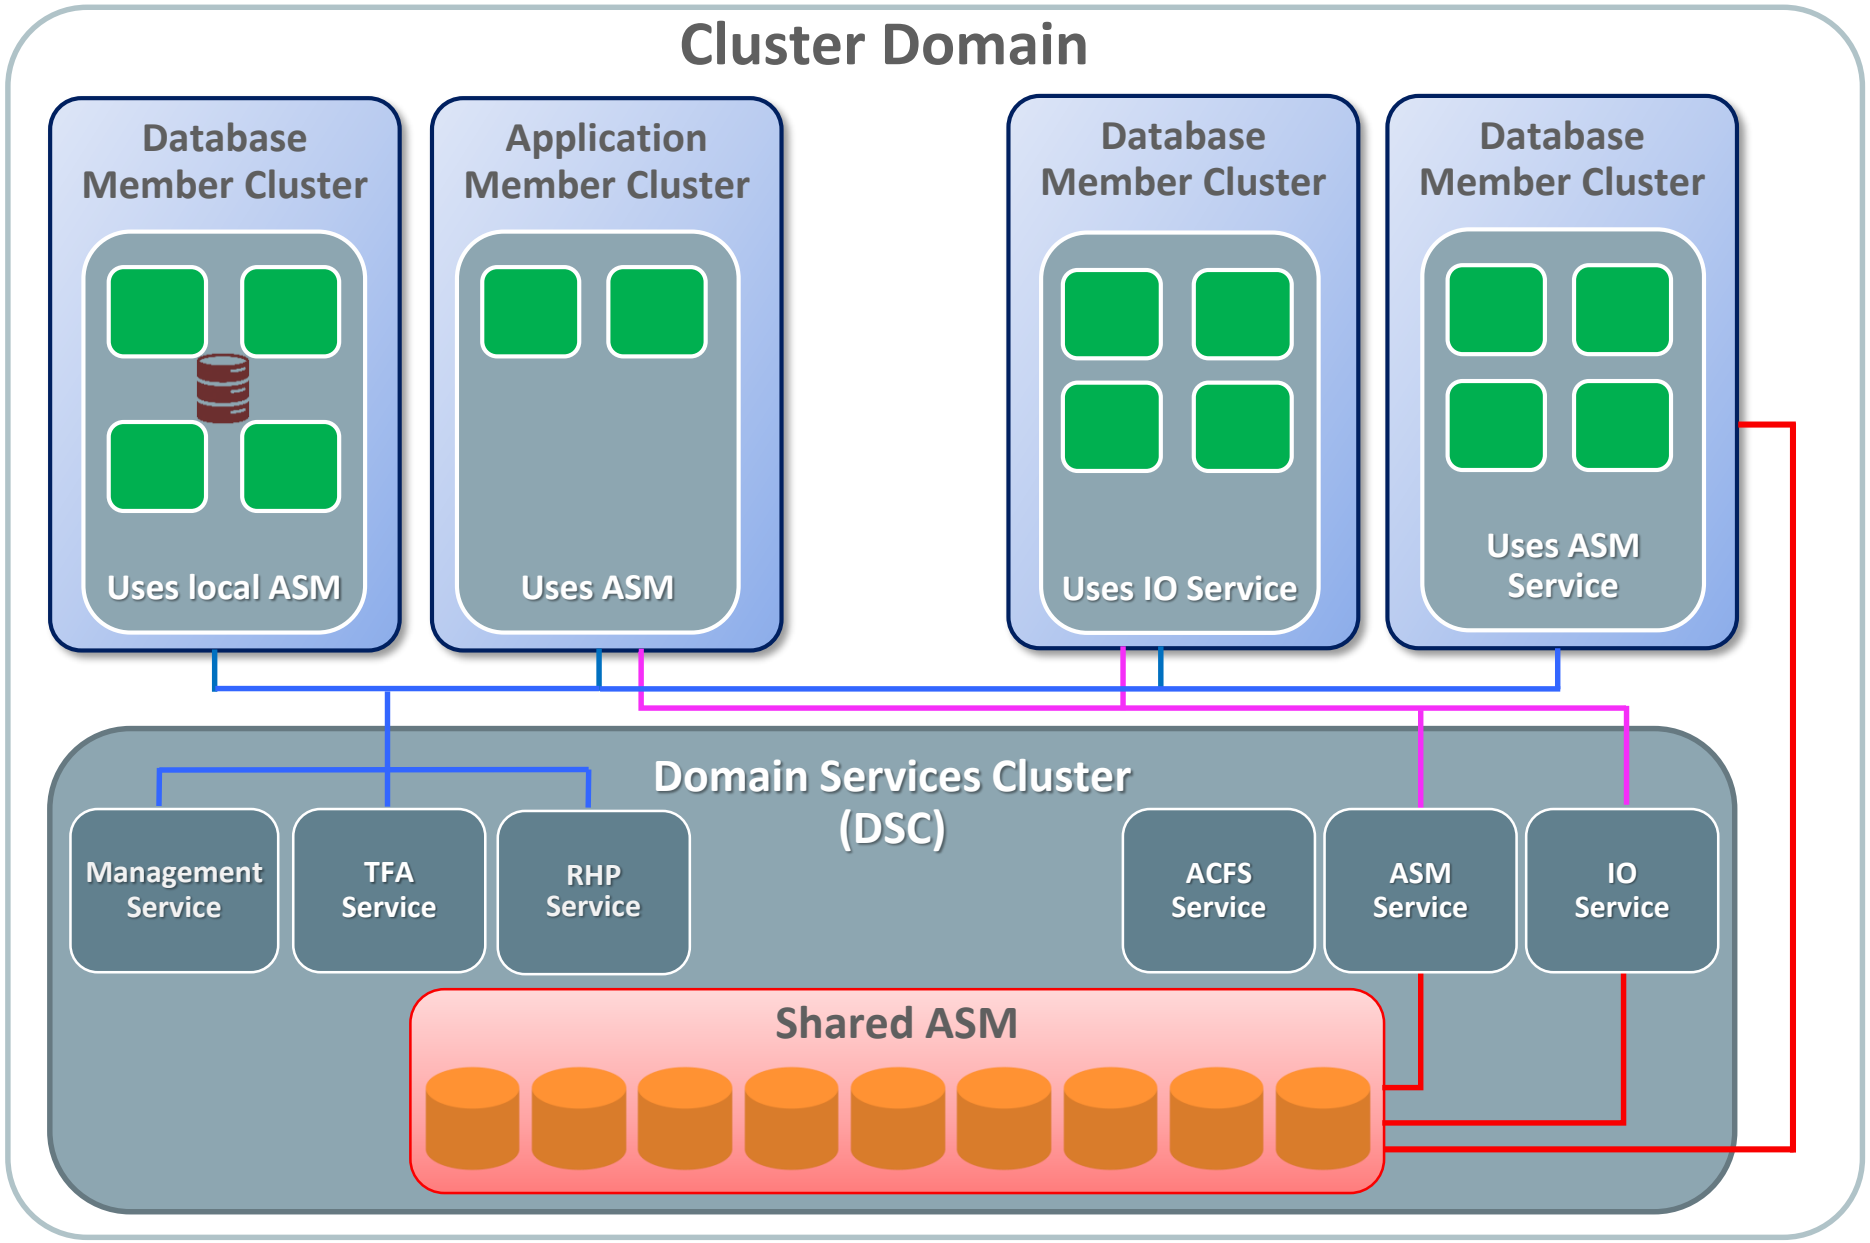

Database Member Cluster Using the IO Service

Consolidation at its best utilizing full resource sharing

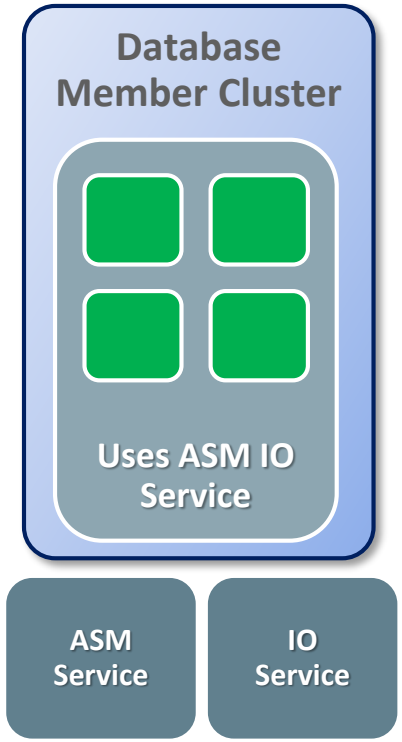

- For databases
  - That need to be deployed quickly
  - or cloned or duplicated frequently
- Ideal for databases that can allow for IO path sharing with other Member Clusters, while maintaining control of local cluster processing, workload and resources
- Examples include
  - Test, integration, development systems
    - For which maintaining additional hardware (e.g. SAN networks and storage) is inefficient

Summary

Oracle is Listening to Customers

- Reducing the Management Overhead
  - Applied Machine Learning
    - AHF, TFA, Hang Manager
  - Fleet Management
    - Cluster Domains, RHP
- Making Automatic & Autonomous Decisions Logically
  - Predictive Problem Analysis and Warnings
  - Automatic Hang Detection and Resolution
  - Smart Reconfiguration

Integrated Cloud

Applications & Platform Services