18c ORACLE Database

Oracle Grid Infrastructure – Cluster Domains

ORACLE WHITE PAPER | FEBRUARY 2018

ORACLE

Table of Contents

Introduction

Cluster Domains were introduced in Oracle Database 12*c* Release 2 as a cluster architecture designed to reduce management overhead by centralizing and consolidating storage management and other common services for groups of otherwise independent Oracle Clusters. While the Standalone Cluster configuration (the only cluster architecture available prior to Oracle Grid Infrastructure 12*c* Release 2) is still supported, there are significant benefits to be realized in moving to Cluster Domains, particularly for larger cluster estates in which the management needs of the many clusters are ever-increasing.

Oracle's Cluster Domain architecture enables simpler, easier deployments, reduced storage management effort and performance gains for I/O operations. Adoption of the new Cluster Domain architecture can be via new installations through the Oracle Universal Installer, or via upgrades of established clusters.

New with Oracle Database 18c is the addition of an automated conversion from the Standalone Cluster to a Member Cluster within a Cluster Domain. While in Oracle Database 12c Release 2 there was only the option to create Member Clusters as part of a fresh deployment, the new conversion capability enables customers' current Standalone Clusters to benefit from the management improvements available with the Cluster Domain architecture.

In addition, a new service hosted on the Domain Services Cluster has been introduced to enable the hosting of local ACFS file systems on Member Clusters that are configured to use remote ASM-managed storage. This includes the full ACFS functionality to the Member Clusters, without having to manage NFS mounts to remote storage.

Clustering with Oracle Clusterware

Oracle Clusterware enables the clustering of otherwise independent servers so that they co-operate as a single system. As a cluster, these servers then provide the integrated foundation which Oracle Real Application Cluster (RAC) Databases and user applications can leverage for high availability and scalability.

The cluster of servers is coordinated via Oracle Clusterware, with cluster resources made available as required in support of the high availability requirements of the Oracle RAC Databases and applications running on the cluster, on one or more of the clustered servers, or nodes. Introduced in Oracle 10g Release 1, Oracle Clusterware has evolved and broadened its capabilities to meet the demand for a more versatile and more capable infrastructure.

With Oracle Grid Infrastructure 12c Release 2[1], new cluster architectures were made available, all based upon the Flex Cluster architecture that was introduced with Oracle Grid Infrastructure 12c Release 1. Flex Clusters were designed to enable customers to scale their clusters beyond established norms in support of mixed workloads, with specific support for Oracle RAC Database instances and applications in the same cluster.

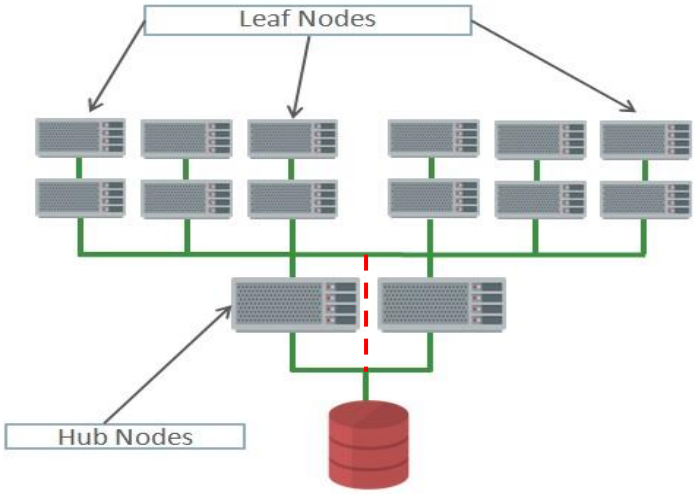

Figure 1: Flex Cluster Architecture

As Oracle Clusters have been accepted and deployed to advantage for a broad range of customer database and application requirements, the demand has risen to both simplify the management of multi-cluster environments and to consolidate them under a reduced management framework. This is immediately apparent to customers who have multiple clusters deployed, as the operational demands have increased with the deployment of each new cluster.

Then, with Oracle Grid Infrastructure 12c Release 2, the customer has options in their cluster deployment, to stay with the previously available cluster architecture (that has been well proven over the last decade) or to adopt the new multi-cluster architectures of the Cluster Domain, especially if managing a bigger estate of Oracle Clusters.

1 Oracle Grid Infrastructure was introduced with Oracle 11g Release 1 as a composite of Oracle Clusterware and Oracle Automatic Storage Management (ASM).

Oracle Grid Infrastructure Architectures

There are now two architectures for deploying clusters using Oracle Grid Infrastructure. These are the Standalone Cluster and the new Cluster Domain. The clustering software for these deployments is identical, differing only how they are configured during deployment.

Standalone Cluster

The Standalone Cluster consists of one or more cluster nodes configured with locally available shared storage, a private interconnect, local instances of Automatic Storage Management (ASM) for managing that shared storage, and the GI Management Repository (GIMR) for storing cluster health and diagnostic information.

This is essentially the same as the Standard Cluster of previous releases. It offers the highest degree of workload, storage and resource isolation possible for a highly available, scalable clustered deployment. Upgrades from previous GI versions will automatically be Standalone Clusters.

They are best suited for supporting databases with unpredictable or highly variable workloads. Essentially, those databases with workloads that are not good candidates for consolidation and resource sharing. Candidates for such criteria would include Business Intelligence and Analytics, or Batch Processing systems.

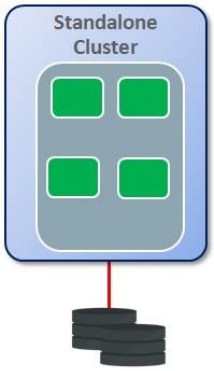

Figure 2: Standalone Cluster

Cluster Domain

A Cluster Domain is actually a grouping of clusters. A Cluster Domain consists of a single Domain Services Cluster and a number of Member Clusters (hosting applications or databases) that utilize services offered on the Domain Services Cluster. Centralized and consolidated services are hosted by the Domain Services Cluster, and consumed by the Member Clusters that are registered with that Domain Services Cluster.

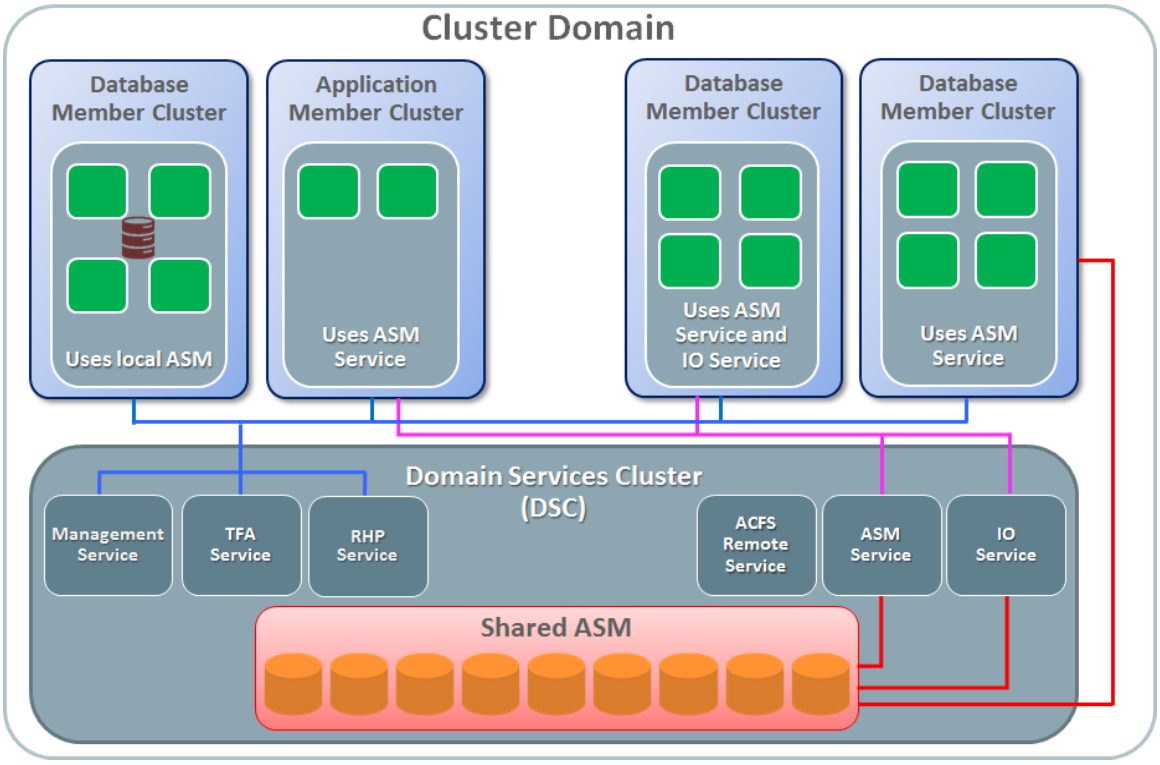

Figure 3: Cluster Domain

Domain Services Cluster

The Domain Services Cluster is the heart of the Cluster Domain, as it is configured to provide the services that will be utilized by the various Member Clusters within the Cluster Domain. As per the name, it is a cluster itself, thus providing the required high availability and scalability for the provisioned services.

The services available consist of:

- Management Service,
- Storage Services (ACFS Service, ASM Service and IO Service),
- Trace File Analyzer Service, and
- Rapid Home Provisioning Service.

Member Clusters

The Member Clusters within a Cluster Domain are deployed to host databases and applications. They can be configured purely for applications, in which case they lack the support for databases, purely for databases, or for a mix of databases and applications. Depending upon the centralized services to which they subscribe, they will be configured as one of four constructs:

- Application Member Cluster – cannot support database instances, but runs with a lighter memory footprint
- Database Member Cluster – configured with local ASM and storage

- Database Member Cluster – configured with directly attached shared storage that is managed through ASM on the Domain Services Cluster, or
- Database Member Cluster – configured with indirect access to the shared storage managed and presented by ASM on the Domain Services Cluster.

Database Member Clusters are designed to host databases, but may also host applications or programs that benefit from direct local access to the local database instances (for example, Oracle GoldenGate). For a DBA managing their database and instances, the database management on any of the Database Member Cluster constructs will appear to be identical. It is only at the cluster deployment and administration levels that there are differences.

Domain Services Cluster

At the heart of the Oracle Cluster Domain, the Domain Services Cluster, or DSC, must be deployed first and configured with the services required for subsequent Member Cluster deployments. Shared storage is attached to the DSC, which is managed by Oracle ASM and configured for Member Cluster access.

Being a cluster itself, the Domain Services Cluster also has its management repository stored in the centralized Grid Infrastructure Management Repository hosted on the DSC.

Management Service

The GIMR hosted on the Domain Services Cluster is configured as a Multitenant Single-Instance Database (CDB) in which each of the Member Clusters and the DSC stores its management repository as an individual PDB. Cluster health and diagnostic information is written to the GIMR by background processes (daemons) on each of the Member Clusters over SQL*Net. Data security and isolation are carefully maintained by deploying separate PDB's for each Member Cluster. Accesses to a cluster-specific data for diagnostic and analysis purposes must be from the cluster itself, again over SQL*Net, against that cluster's Management Database in the GIMR on the DSC (the only exception to this is for Oracle Enterprise Manager, as it references this information across the Public Network, much as it does for any other target).

The centralized GIMR is host to cluster health and diagnostic information for all the clusters in the Cluster Domain. As such, it is accessed by the client applications of the Autonomous Health Framework (AHF), the Trace File Analyzer (TFA) facility and Rapid Home Provisioning (RHP) Server across the Cluster Domain. Thus, it acts in support of the DSC's role as the management hub.

Storage Services

By centralizing the Oracle ASM storage management, Database Member Clusters can be deployed without configuring ASM to run directly on those clusters. Oracle ASM storage management is consolidated onto the DSC. In addition, the ASM instances on the DSC will now manage the consolidated storage pool for multiple Member Clusters.

Where previously, with ASM instances configured locally to manage the storage for each cluster, there would be a single, centralized ASM deployment tasked with managing all the disks in support of the entire Cluster Domain. This increases the pool of shared disks over which ASM can distribute shared storage for each cluster, resulting in potential I/O performance gains. It also allows ASM to manage storage changes more efficiently. In addition, ASM's new Flex Disk Groups feature will maintain the security and isolation of data and data files on a database basis.

In addition, the ACFS Remote Service has been introduced with Oracle Database 18c to enable the deployment of ACFS file systems directly on Member Clusters, yet still using the remote ASM storage management services hosted on the Domain Services Cluster. The increased flexibility of such deployments will now enable the use and optimization of local shared storage on Member Clusters without having to configure NFS mounts.

The ACFS Remote service will access storage via iSCi connections over the network, as designated during deployment. This will provide the best combination of performance and reliability (relying upon how the network multi-pathing and bonding has been configured).

Accessing Centralized Storage

Shared storage service access from the Database Member Clusters may be via direct paths to the disk storage (i.e. the disk storage is mounted to both the DSC and to the Database Member Cluster) or via an indirect I/O path (without disk storage mounted on the Database Member Cluster).

Using locally attached shared storage (that is managed from the DSC) has all the benefits of local shared storage without the overhead of running locally configured ASM instances. So, instead of registering with local ASM instances, the database instances on the Member Cluster register with the ASM instances on the DSC. The I/O path between the database instances and the locally attached shared storage would still be direct.

Configuring the Database Member Cluster to use an indirect I/O path to storage is simpler still, requiring no locally configured shared storage, thus dramatically improving the ease of deploying new clusters, and changing the shared storage for those clusters (adding disks to the storage is done at the DSC – an invisible operation to the Database Member Cluster). Instead, all database I/O operations are channeled through the IOServer processes on the DSC. From the database instances on the Member Cluster, the database's data files are fully accessible and seen as individual files, exactly as they would be with locally attached shared storage. The real difference is that the actual I/O operation is handed off to the IOServers on the DSC instead of being processed locally on the nodes of the Member Cluster. The major benefit of this approach is that new Database Member Clusters don't need to be configured with locally attached shared storage, making deployment simpler and easier.

Trace File Analyzer Service

The Trace File Analyzer (TFA) Service provides centralized storage for TFA collections, for on-site analysis of diagnostic information via the TFA Receiver and for uploads to Oracle Support to assist with Support Request submissions. New functionality offered by the TFA Service further enhance previous releases of TFA by using Applied Machine Learning to apply the latest diagnostic models to the data in order to better diagnose issues that have occurred.

Rapid Home Provisioning Service

The Domain Services Cluster may also be configured to host a Rapid Home Provisioning (RHP) Server. RHP is used to manage the provisioning, patching and upgrading of the Oracle Database and GI software stacks and any other critical software across the Member Clusters in the Cluster Domain. Through this service, the RHP server is used to maintain the currency of the installations on the Member Clusters as RHP clients, thus simplifying and standardizing the deployments across the Cluster Domain.

Member Clusters

There are four possible configurations currently supported for Member Clusters, each configured during installation.

Application Member Clusters

The distinguishing features of an Application Member Cluster are a much smaller memory footprint and fewer required IP addresses and has no pre-configured shared storage (in ASM or otherwise). This enables the use of smaller servers for supporting application-only clusters.

This means that the Application Member Cluster is dedicated and optimized to support only applications or programs that may or may not access remote databases, but without the overhead required to support locally configured database instances. By utilizing Oracle Clusterware functionality for defining and managing resources, deploying mission critical applications on the Application Member Cluster is a simple way to provide an initial layer of high availability and scalability.

In addition, the availability of the applications or programs can be further enhanced with the deployment of XAG Agents, enabling the customization of high availability capabilities for those applications running on the Application Member Clusters.

New with Oracle Clusterware 18c, new Application Member Cluster deployments can be configured to use the Public Network for Cluster Interconnect Traffic. This removes the need for a dedicated private interconnect network over which the cluster nodes would pass cluster membership messages.

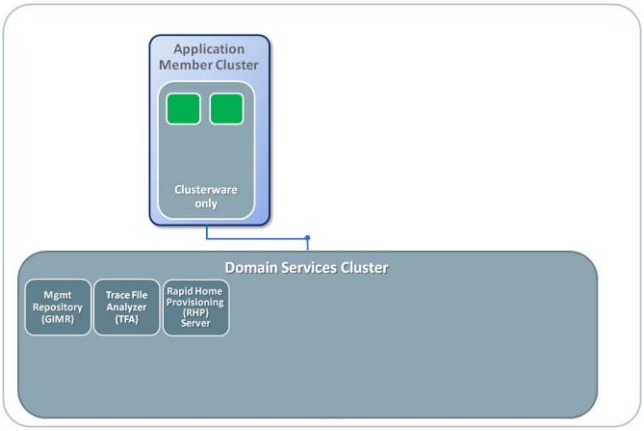

Figure 4: Application Member Cluster

Database Member Cluster – configured with local ASM and storage

This is the a Database Member Cluster on which ASM is configured to run locally, thus the database accesses only locally mounted shared storage. In this scenario, only the Management Database is offloaded to the Domain Services Cluster and stored in the centralized GIMR. Otherwise, this is identical to the Oracle Standalone Cluster.

Customer taking their initial steps in the world of the Cluster Domain may favour this option, since it is a relatively simple step to move an established Oracle Standalone Cluster into the Cluster Domain and only offload the Management Database. It provides the same high degree of isolation for the workload, resource and storage, and allows for the offloading of the Management Service's overhead to the DSC.

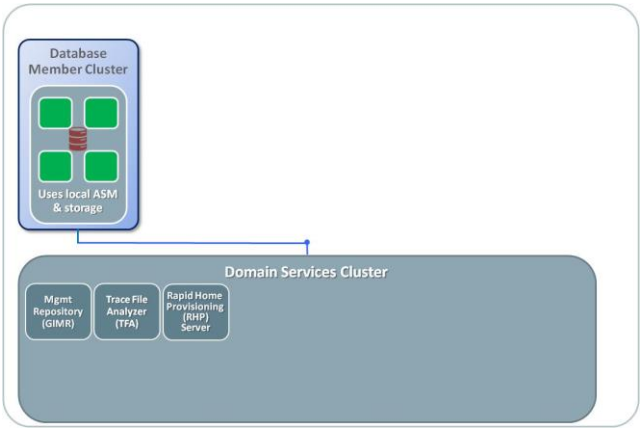

Figure 5: Database Member Cluster with local ASM

Database Member Cluster – configured with remote ASM

There are two alternative architectures for Database Member clusters that subscribe to the Storage Service on the DSC. One has the shared storage connected directly to the cluster, while the other has no locally configured shared storage and uses the IOServer running on the DSC for its I/O.

Database Member Cluster – configured with directly attached shared storage

This configuration takes the concept of the Oracle Standalone Cluster that one step further, by not only offloading the Management Database, but also offloading the ASM storage management to the DSC. Storage connectivity is still seen as locally mounted, but now the ASM instances are actually running on the DSC.

The database instances running on this Database Member Cluster register with the remote ASM instances that provides disk access information as required for the database files associated only with that particular database. This preserves the I/O paths that would have been available with the Oracle Standalone Cluster, but removes the dependency upon a locally available ASM instance. The ASM instances running on the DSC are still highly available, but now provide consolidated storage management for all the Member Clusters that are registered with the DSC.

This deployment emphasizes the sharing of storage for multiple Member Clusters, providing I/O performance stability and speed with the increased disk storage pool, and yet preserves the isolation of the I/O path itself. It is best suited to those applications and databases that would benefit from overall I/O performance stability and not be impacted by any issues with shared I/O paths. Examples of databases that would fit this criteria would be high-transaction rate OLTP systems or those systems typified with high random I/O's.

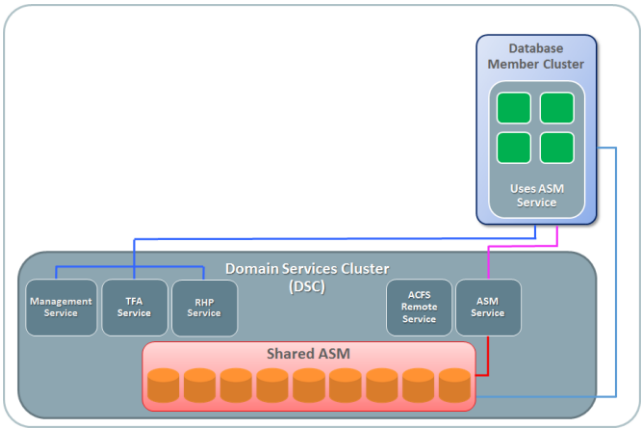

Figure 6: Database Member Cluster with Direct Storage Access

Database Member Cluster – configured with indirect access to shared storage

Under this configuration, the storage managed by ASM on the DSC is not mounted on the Database Member Cluster. Instead, the database instances on the Database Member Cluster submit their I/O requests over the network to the IOServer processes on the DSC which then pass on those requests to the I/O subsystem on the DSC nodes. These IOServer processes essentially act as pass-thru processes for handing off the I/O requests, without having to do any real work. Thus, they will not add overhead to the I/O processing.

From the perspective of the database instances on the Database Member Cluster, the DBA still sees their database files exactly as they are used to seeing them. They can manage and access those files exactly as they have always done, through their normal database management toolset.

This architecture for a Database Member Cluster that requires no locally configured shared storage would be particularly attractive as it simplifies the deployment of RAC and RAC One Node clusters, plus inherently benefits from the consolidation of the storage itself. The simplicity and ease of deployment (due to not having to allocate shared local storage) would be greatly beneficial for temporarily deployed clusters, such as for test and development, or for virtualized environments in which the speed of deployment is highly regarded. Highly consolidated systems and those that change in character (due to changes in workload or constituent databases) would be good fits for this type of deployment.

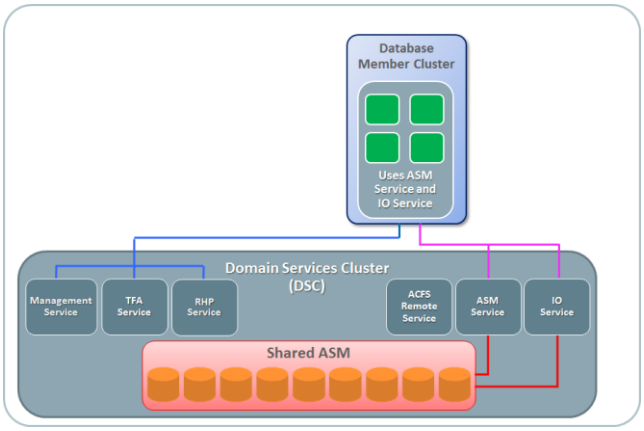

Figure 7: Database Member Cluster with Indirect Access to Storage

Cluster Domain Adoption

The Cluster Domain architecture was introduced initially for new deployments only. With this cluster architecture newly available it was believed that customer would most likely want to test it with new cluster deployments, rather than with established clusters.

Starting with Oracle Database 18c, administrators have the ability to convert Standalone Clusters to Database Member Clusters. This conversion will effectively bring the Standalone Cluster into the Cluster Domain and enroll that cluster to use the Management Service on the Domain Services Cluster. The data from the local Grid Infrastructure Management Repository (GIMR) would be transferred to the centralized GIMR, and the new Database Member Cluster would operate exactly as it had as a Standalone Cluster, but now would take advantage of the benefits of the centralized Management Service.

Benefits of the Cluster Domain Architecture

The Cluster Domain architecture is primarily a solution for managing an ever-expanding group of Oracle Clusters, thereby reducing management overhead in deploying new clusters, while consolidating storage management and offloading non-critical infrastructure. In doing so, there is also the potential for I/O performance gains as storage is consolidated, and also reducing the wastage in allocating storage on a per-cluster basis.

The Cluster Domain is a service-oriented architecture, enabling the services to be configured once and reused many times over. It provides the security and isolation required for consolidated environments, and has the potential to provide significant benefits without unacceptable costs.

In essence, by adopting the Cluster Domain architecture, there are benefits to be realized in risk reduction, storage optimization, ease of management and deployment, and better optimization of resources. In this environment, it is simple to envision DBA's focusing exclusively on their databases, while storage administrators managed the shared storage on the DSC, and cluster administrators deployed new clusters in a timely fashion to readily meet the demands of the business for more processing, more databases, more resources and more platforms.

References

For further details and reading on the features of Oracle Grid Infrastructure 18c, please refer to the following links:

- Oracle Webpages
  - http://www.oracle.com/goto/clusterware
  - http://www.oracle.com/goto/asm
  - http://www.oracle.com/goto/ahf
  - http://www.oracle.com/goto/rac
  - http://www.oracle.com/goto/rhp
- Oracle Documentation
  - Oracle Clusterware Administration & Deployment Guide (covers RHP in addition to Clusterware)
  - Oracle Automatic Storage Management Administrator's Guide
  - Oracle Autonomous Health Framework User's Guide
  - Oracle Real Application Clusters Administration and Deployment Guide

ORACLE

Oracle Corporation, World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065, USA

Worldwide Inquiries
Phone: +1.650.506.7000
Fax: +1.650.506.7200

CONNECT WITH US

blogs.oracle.com/oracle
facebook.com/oracle
twitter.com/oracle
oracle.com

Integrated Cloud Applications & Platform Services

Oracle Grid Infrastructure – Cluster Domains
February 2018
Author: Ian Cookson
Contributing Authors: Markus Michalewicz, Mark Scardina, Anil Nair, Jim Williams, Ricardo Gonzalez

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