Oracle's Clustered Architectures
Which is Best for You?

Ian Cookson
Oracle

Keywords:

Clusterware RAC Real Application Clusters ASM ACFS Database

Introduction

Oracle's Clustering solutions were introduced to provide high availability protection and scalability across the cluster for databases and applications. And, while much of the early emphasis was on providing these functions for database software, in the form of Oracle RAC and RAC One, there has been a steadily growing recognition of the same requirements for applications too.

Now, as demands for public and private cloud initiatives intensify, with greater demands for ease of provisioning, deployment, management and operations, there are new approaches to architecting Oracle clustered solutions. And, while many may wish to adopt new approaches, there are now alternatives that will help users meet these changing demands in a sustainable and prompt fashion.

Clusters in 11.2 (and earlier) – One Cluster for All

Oracle clusters in version 11.2 and earlier all conformed to a single architectural model, in supporting databases and applications, or a combination of both. They consisted of one or more cluster nodes with any amount of shared disk storage attached.

Typically, the cluster nodes ran the Oracle database, while applications ran on other servers (application servers), on users' desktops, or on the cluster nodes themselves. Thus, the only real variation in the deployment architectures was where and how the applications would be run.

While this was a simple architecture to understand and support, with each cluster node 'normally' referred to as a database node or database server, there was not a lot of opportunity to optimize deployments for any particular mix of applications and databases.

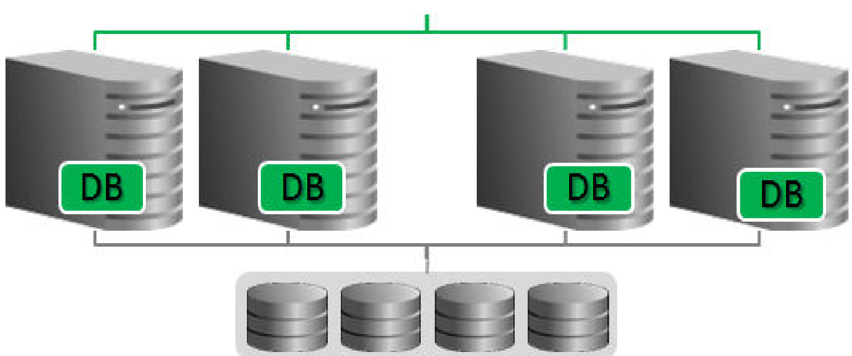

Illustration 1: Oracle 11.2 Cluster Architecture

Clusters in 12.1 – to 'Flex' or not to 'Flex'

With the advent of Oracle 12.1, there were now deployment options that added flexibility in the cluster architectures. These were introduced to try to address the shortcomings in the 'One Cluster for All' solutions previously available.

While the original, pre-12c cluster architecture commonly known as the 'Standard Cluster' was still available (see *Illustration 2*), Oracle introduced the 'Flex Cluster'. This new architecture enabled formal lighter-weight support for applications within an otherwise database-centric cluster.

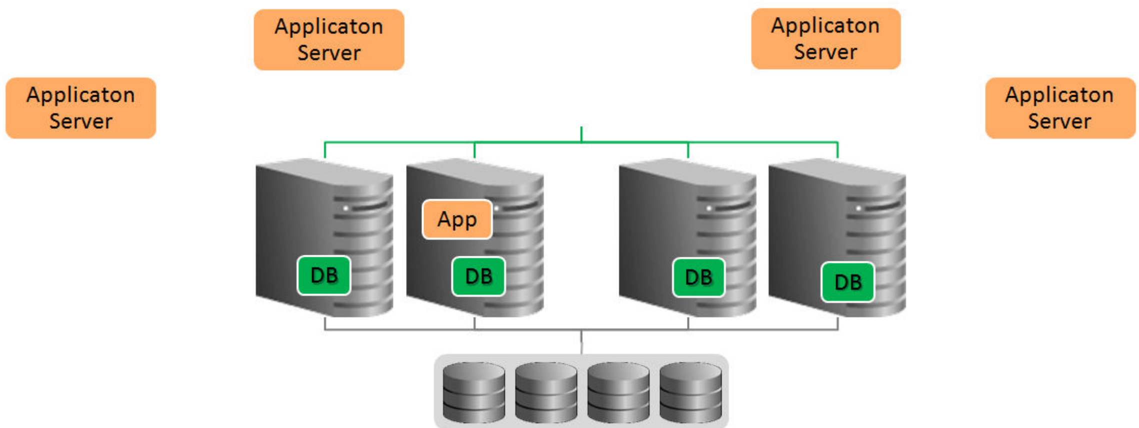

Illustration 2: Oracle 12.1 'Standard Cluster' Architecture

The Flex Cluster architecture allowed for Hub Nodes (on which database instances were intended to run) that had direct access to the shared disk storage in a traditional tightly-coupled arrangement, and Leaf Nodes (on which applications were intended to run) that did not have direct access to shared disk storage in a loosely-coupled arrangement (see *Illustration 3*).

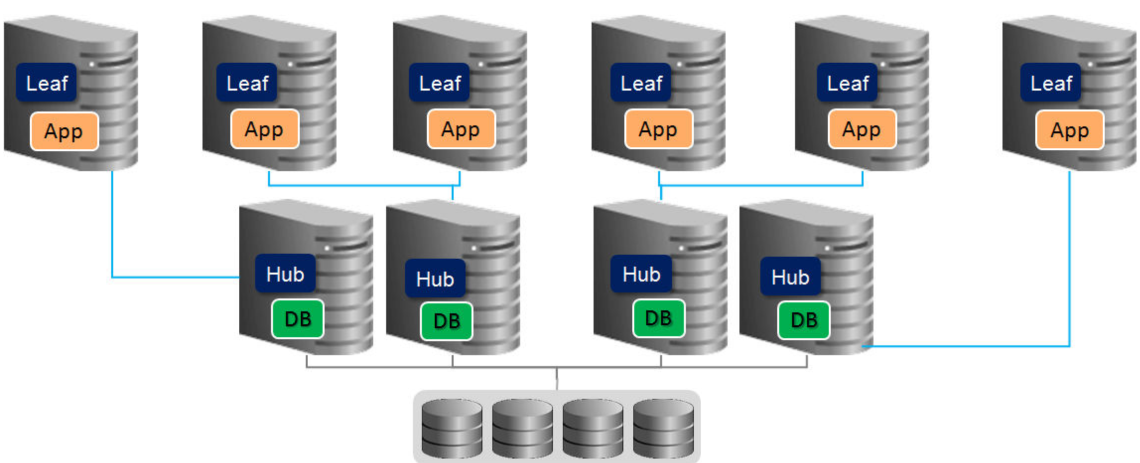

Illustration 3: Oracle 12.1 'Flex Cluster' Architecture

Now, the applications and related database instances could be managed as a single cluster, tying the application and database resources closely together. In addition, the Leaf Nodes could be sized much smaller than the Hub Nodes, since they would not host database instances and access shared storage. And, as the move to a virtualized environment was gaining momentum, a mix of physical and virtual server nodes was now available.

While there were many other features added to the Oracle Clustered stack in 12.1, two of particular note were Flex ASM and the Grid Infrastructure Management Repository. Flex ASM had been introduced to remove the dependency of database instances on the need for a local ASM instance. With this new ASM functionality, there was no longer any need to have an ASM instance running on every Hub Node. And while this obviously reduced the resource load on the nodes without ASM instances, the primary benefit was the removal of a single point of failure for local database instances (as they no longer depended upon having a viable local ASM instance).

The second feature of note was the introduction of the Grid Infrastructure Management Repository, or GIMR. It was designed to host cluster-wide information and statistics that were used by the Cluster Health Monitor to help diagnose cluster-wide performance or viability events.

The GIMR was required, but stored in its own wholly self-managed database within ASM. And, while shared storage was necessarily assigned from within established ASM disk groups, the actual storage volume was small, as was the IO activity. The Management Database (in which the GIMR was stored) was actually configured as a single-instance PDB within a Multi-tenant CDB.

Illustration 4: The Grid Infrastructure Management Repository ->

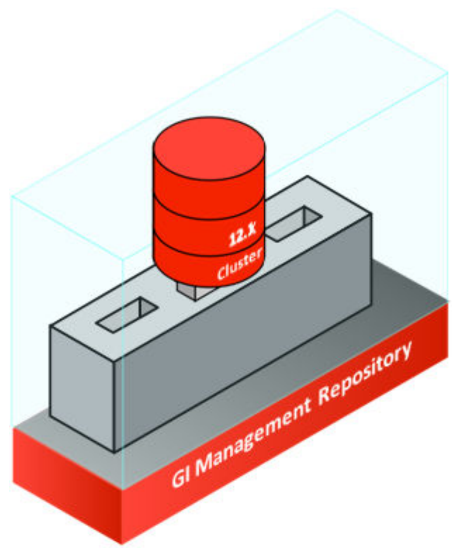

The Next Era of Clustering Architectures – Which is Best for You?

The future of Oracle's clustering architectures is expected to introduce yet more alternatives, to meet growing demands for ease and assurance in provisioning, deployment, management and operations. Based upon a shared services model, these new architectures are designed to help meet the diverse requirements of the private and public clouds, and on-premise solutions.

Using this shared services model, any number of clusters would be configured to use, or consume, the shared services. These services would provide the following:

- consolidated storage management and accesses,
- centralized location for offloading each cluster's GIMR,
- a unified, web-enabled trace file analyzer store and interface, and
- provisioning support (for patches, deployments and upgrades) via RHP

The clusters using these shared services could then be configured specifically to meet the technical and usage needs of the applications and databases that they would host, without having to be limited by the infrastructure requirements.

Once the shared infrastructure was deployed and configured, the provisioning of new clusters would be greatly simplified. It would fall into one of the following four categories:

1. A cluster to host databases and possibly applications in relatively complete physical isolation (limited resource sharing) – like a cluster from 12.1 or 11.2, with local shared disk storage, or
2. A cluster to host applications (with no capacity to host database instances) with no direct access to shared disk storage (a good option for virtualized deployments), or
3. A database cluster with direct access to shared disk storage (connected to a SAN), but with ASM disk management offloaded, or
4. A database cluster with no direct access to shared disk storage (for instance, a virtualized deployment with no provisioning for shared disk access).

The increased physical isolation of the deployment that hosts its own ASM instance (and shared disk storage), yet offloads its GIMR, would benefit those systems that need to strictly control possible resource sharing contention and security exposures. It would also suit customers that have a relatively small requirement for clustered deployments and might not realize the gains that the fully shared infrastructure could provide others.

A cluster deployment specifically designed for applications would continue and improve upon the features of the 12.1 Flex Cluster. In this case, the resource requirements for the cluster nodes would be greatly reduced, as the configuration would exclude direct support for databases. Yet, in every other way the cluster nodes and applications could be protected from failures, just as the database instance would be on other clusters.

The database member clusters could be configured to use directly attached shared disk (i.e. the cluster nodes are directly connected to the shared storage, as in a SAN) or the use remote IO services provided directly by ASM. Cluster deployments requiring support for higher IO demands, or those already configured with locally attached shared disk storage would fall in the former category, while clusters not having the same IO demands, or without being provisioned with shared disk storage would fall into the latter category.

Provisioning for each of these components, including the shared infrastructure, would be provided by the Oracle Universal Installer (OUI), using specific, focused dialogues that would streamline the process. Following a step-by-step process, the customer would be directed to provide the required directions for deploying best-practice configurations of the cluster architectures that would best suit their needs.

Summary

The evolution of Oracle’s Clustering solutions continues with the introduction of a shared service model, providing alternative deployment architectures that are designed to better meet customer requirements for ease of provisioning, deployment, management and operations.

While the existing cluster architectures will continue to be provided and supported, the new options are expected to assist customers better meet the shifting demands of the current cloud-focussed world. Whether the need is for application availability, high performance and highly available databases, in a virtual or physical scenario, there are options available for easy deployment.

No longer will customers be tied to the ‘One Cluster for All’ architecture. They will now be enabled and empowered to deploy multiple architectures to meet multiple different demands.

Contact address:

Ian Cookson
Oracle America, Inc
500 Oracle Parkway
Redwood Shores, CA
USA 94065

Phone: +1 650-633-9992
Mobile: +1 408-616-0399
Email : ian.cookson@oracle.com