White Paper

Consolidation, without virtualising

Gaining all the benefits of server consolidation,
and more – using Oracle RAC

Gaining all the benefits of server consolidation, and more – using Oracle RAC

Modern IT departments are coming under more pressure to deliver higher returns on investment from IT infrastructure, as well as showing an effort to become more environmentally friendly. "Consolidation" has become one of the keywords when striving to achieve these investments, with IT managers attempting to "make the most" of IT infrastructure – achieving the dream of "doing more with less".

Less hardware of course means less capital investment into said hardware, alongside reduced operating costs – with software licences, support fees, physical data centre space, electricity bills, and administrative resources decreased, often with a cascading effect into supporting infrastructure such as backup, storage, and High-Availability (HA) / Disaster Recovery (DR) solutions. Total Cost of Ownership (TCO) decreases dramatically, Return On Investment (ROI) increases dramatically, and the IT department can demonstrably claim to have made genuine progress in becoming a 'green' organisation.

Oracle Real Application Clusters (RAC) technology provides an excellent method of hardware consolidation. Organisations with multiple Oracle databases can consolidate multiple database servers, perhaps with differing hardware and operating system versions, into a decreased hardware footprint – achieving the core goals of hardware consolidation, while also providing improved scalability, High Availability, and ease of administration – without any requirement for a new skill set.

This paper is intended to show the possibilities of Oracle RAC to achieve hardware consolidation – achieving all the benefits of more common methods such as virtualisation, while also providing the core benefits of Oracle RAC to provide High Availability – all without requiring any new software or skill sets, and potentially without requiring new hardware.

Layers of Consolidation

Consolidation is not a new concept within IT. For decades, IT managers have been attempting to reduce costs by consolidating resources – making more efficient use of a smaller amount of these resources, while also attempting to simplify and standardise the IT infrastructure, thereby providing ease of administration – reducing the support cost.

Several layers of consolidation exist – some of which have moved from being "best practice" to "normal practice", others which are driving the modern age of "Consolidation Projects"; and some which are still considered to be a "one step too far".

	Level of Consolidation →				
Consolidation:	Geographic	Storage	Cluster / Grid Computing	Server	Database
By sharing:	Data Centre	Storage Area Network	Cluster	Physical Server	Database Instance

Geographic location

The consolidation of servers within an organisation into one or more centralised data-centres has long ago moved from being "best practice" to "standard practice". This was forced by the need to accommodate an exponential growth in IT infrastructure, but also provided an impetus to investigate how to maximise efficiency of existing resources (data centres) by sharing such resources.

Storage

The move from each server with it's own attached storage to servers all connected to a consolidated Storage Area Network (SAN) has also long been "standard practice". This stemmed from the consolidation of the data centres, and the natural progression of the realisation that sharing resources allowed greater flexibility, and hence efficiency.

Cluster/Grid computing

A more recent development in consolidation options has resulted from clustering technology. Multiple servers can now utilise the availability of the shared storage capabilities of a SAN, to be linked together in a cluster. The most common use of a cluster today is to achieve a form of High Availability (HA) – whereby if one server was to 'crash' – the applications residing on this server could 'failover' to another server within the cluster.

Server

The current generation of "Server Consolidation" projects seen in IT departments is an attempt to reduce the number of physical servers – often achieved by converting each existing physical server to a "virtual machine" – and hosting multiple virtual machines on one larger physical server. There are other methods to achieve the goals of "Server Consolidation" however, and this paper seeks to address one of these methods, specific to servers hosting Oracle databases.

Database

A deeper layer of consolidation specific to databases is allowing multiple applications to share a single database. Although common practice with Microsoft SQL Server, this level of consolidation is rarely implemented in Oracle implementations. This may be due to organisations having a desire to keep the underlying data for disparate applications separated as much as possible – whether for reasons of security, differing availability/service-level requirements, identification of cost-of-ownership between departments. With modern versions of Oracle Database however, most if not all such concerns are easily remedied, and we may soon see the next generation of "Database Consolidation" projects.

Reasons for server consolidation

The general reasons for consolidation at any level – increase efficiency, reduction in total cost of ownership via less assets plus simplification of administration – are magnified at the level of Server Consolidation.

According to VMWare "Most servers operate at only 5-15% of their total load capacity." This shows an obvious and huge opportunity for consolidation. If one considers optimal capacity to be 80%, this suggests the capability for reducing the number of physical servers by a factor of anywhere from 5 to 16 with a consolidation project. For every physical server decommissioned, there are immediate savings in the areas of hardware support, software licence/support fees, electricity, and costs associated with the physical data centre rack space. In addition to this, consolidation at any level generally includes standardisation – therefore Server Consolidation ideally includes standardising hardware and operating systems, in turn greatly reducing the range of skills needed for technical support – and greatly reducing the cost of support.

Reduce hardware costs

The most obvious benefit of Server Consolidation is quite simply that the organisation needs to purchase less servers. The initial cost of purchasing an enterprise server is not inconsiderable, but this is not the end of the cost of owning an enterprise server…

Reduce operating costs

There are many costs associated with keeping a physical server running, and available for use. Hardware support agreements, operating system licenses, software licenses, electricity required by both the server & the data centre temperature control system – all add up to an ongoing cost.

Reduce data centre footprint

In addition to the cost of keeping a server in a data centre, in many cases there is a more intangible cost benefit, but more tangible physical benefit – in minimising the number of server racks populated. Many organisations have limited physical rack space in their data centre(s) – while others may wish to minimise use of rack space by internal systems, for example, to maximise use of rack space for revenue-gathering activities.

Reduce environmental impact

Many organisations today are striving to meet obligations, either regulatory or self-imposed, to reduce their environmental footprint. One of the largest environmental effects of any IT department is the data-centre. Reducing the number of physical servers maintained is a huge step in reducing CO2 emissions.

Centralised management

A smaller number of (ideally standardised) servers makes a centralised management structure much simpler to implement. A standardised consolidated strategy towards servers allows creation of a clean, simple view of the IT estate for IT management for monitoring both current state and future growth, while also enabling a 'Dashboard' type view of the estate for the executive level.

Consolidated administrative tasks

With fewer, and standardised, servers – standard administrative tasks such as backup, patching, monitoring, and reporting become simpler. Rather than multiple solutions for each task, for each subset of servers – a single backup strategy and/or monitoring solution can be implemented easily. The size of the patching database is reduced by a magnitude for every standard put into place, and the number of systems to patch reduces linearly with every server consolidated.

Cascading effects

With every consolidation the effects ripple down to supporting systems and efficiency can be raised in every area. For example, every server requires storage for it's operating system, common applications, and some space assigned - but not utilised - to allow for growth. Once servers are consolidated there is only the requirement for a single allocation of this space. This saves a certain amount of storage, but it also reduces the amount of data which needs to be backed up. The number of tapes required to backup the estate reduce, and possibly even the number of tape drives, and underlying infrastructure.

The same concept applies to High Availability & Disaster Recovery environments. Generally, consolidation of any primary environment will lead to reductions in hardware and storage which will be replicated in HA and/or DR infrastructure, cascading to overall operating costs.

Consolidating database servers with Oracle RAC

The concept of server consolidation using Oracle RAC is slightly different to other approaches, as it not only consolidates servers, but also at the cluster layer – thereby gaining not only the cost benefits of Server Consolidation, but also the technical capabilities of modern Grid computing.

But not only does consolidation using Oracle RAC add technical benefits over and above the traditional benefits of server consolidation – it is also a very simple process compared to other methods of consolidating servers. In order to demonstrate this, we will walk through two scenarios, utilising Oracle RAC to consolidate a number of existing Oracle database servers. First, however, an explanation of some of the features of Oracle RAC which make this strategy possible…

Features of Oracle RAC

Failover

Oracle RAC technology provides two different types of failover for database connections:

Connection failover – provides protection at the time of connection. When an attempt is made to establish a connection with the instance on any given node, if there is a failure then the application will failover, and instead establish a connection to an instance running on an available node.

Transparent Application Failover – provides protection for current connections. If there is failure of an instance, or the underlying node, or simply the communication with the instance – the connection fails over to another node. In some cases, the currently executing query will continue - using the new connection, from the point at which the failure occurred.

Both modes of failover allow for multiple automatic retries until a connection is successful.

RAC One Node

With 11gR2 – Oracle introduced the concept of a single instance RAC database. Providing some of the benefits of a RAC database, but for a database which may not require multiple instances, RAC One Node provides the ability for a single instance database to be migrated, with zero downtime, between nodes of a RAC cluster.

Using the Omotion feature, an Oracle RAC One Node instance – even under heavy workload – can be migrated between nodes, with no disruption to availability. Even in the middle of a critical month-end batch process, Omotion allows migration of an instance to another physical server – allowing for online migration to a more powerful server temporarily or permanently, or to allow hardware maintenance – all without end-users ever being aware.

Server Pools

A simplified method of administrating and partitioning an Oracle RAC cluster was introduced in 11gR2 – with Server Pools, and Policy-Managed databases. This feature was primarily designed in order to facilitate consolidation of databases (and applications) onto a single RAC cluster.

A new concept in Oracle's Grid Infrastructure – a Server Pool is a logical grouping of servers within the RAC cluster. Each Pool is defined with a minimum & maximum number of servers, along with an importance, or priority, rating.

A database is then allocated to a Pool, and the expected resource requirement defined. As many instances as required are then started within the server pool to meet this expected workload.

With any change to the number of physical servers available to the cluster – the number of servers actually made available to each Server Pool, and therefore the databases and services running in each pool, will dynamically change to best meet the defined workload requirements.

Instance caging

Another feature in Oracle 11gR2 is Instance Caging – providing greater control of resources per database, and enabling the RAC Cluster as a consolidation platform.

Instance Caging allows the administrator to define a maximum CPU limit used by each database instance – thereby allowing multiple instances to share physical servers, without risk that a poorly written ad-hoc query in the data warehouse creates a process which starves the financial database of CPU during a time-critical month-end reporting cycle. These limits can be dynamically altered, while all instances are online.

Grid plg and play (GPnP)

Oracle 11gR2 also introduced the concept of "Plug'n'Play" to the Oracle Real Application Cluster framework. Utilising the new Grid Naming Service (GNS), it has become possible to move the majority of networking name resolution into the control of the Oracle Grid Infrastructure, while the Single Client Access Name (SCAN) allows for a single virtual address to be used when configuring access to a RAC database.

At the time of creating the RAC cluster, initial configuration steps include creating a small number of DNS entries (one GNS virtual IP, and three SCAN virtual IPs) – along with providing a set of IP addresses within a subdomain assigned to the cluster. GNS then provides all name resolution for these IP addresses. This service allows dynamic allocation and removal of cluster nodes – including mapping hostnames and IP addresses - without any need to reconfigure DNS.

SCAN extends this transparency to the actual access to databases within a RAC cluster. A SCAN provides a single address – used to configure clients and applications – allowing connections to be made to the cluster. This means that the underlying hardware, hostnames, and IP addresses become invisible, and irrelevant, to client connections. Servers can be constantly added, and removed – allowing rolling hardware upgrades or expansion – without connection configuration ever needing to change.

Scenario 1 – Consolidate to new servers

The first scenario is that closest to the most common strategy of server consolidation: replacing a large number of existing small to medium sized servers with a small number of new larger servers.

Supposing an existing estate of several servers hosting Oracle databases – and of course, this is certain to be a mixture of versions, configurations, sizes, etc. There may exist several servers each hosting one single-instance Oracle database each, of differing versions and patch-sets. There may also be one or more existing RAC clusters. For example – suppose an existing Oracle estate looking something like the following – multiple servers, with multiple databases – and the most important database using the latest Oracle software in a RAC cluster:

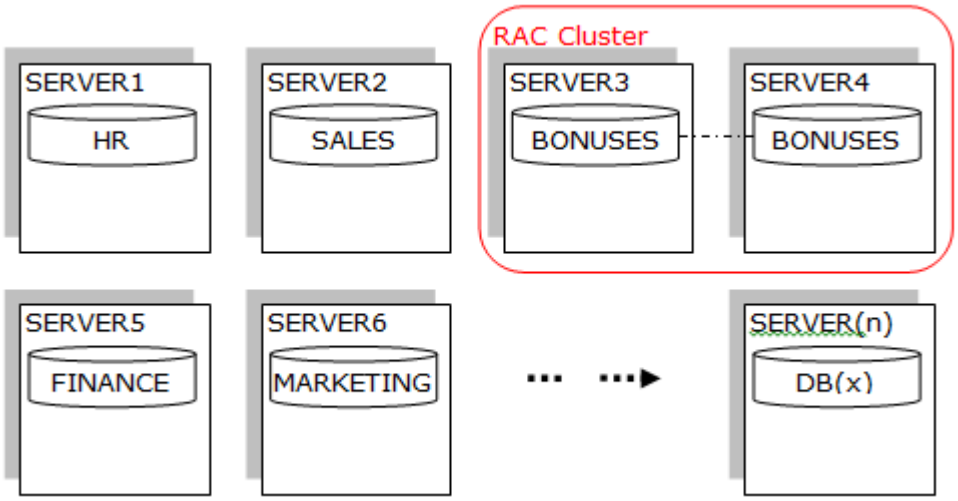

Server1 – HR database, Oracle 10.2.0.5

Server2 – SALES database, Oracle 11.2.0.1

Server3 – BONUSES database, instance 1 – Oracle 11.2.0.2

Server4 – BONUSES database, instance 2 – Oracle 11.2.0.2

Server5 – FINANCE database – Oracle 11.2.0.1

Server6 – MARKETING database – Oracle 10.2.0.4

…

Server(n) – Extensible to dozens, or even hundreds, of databases and/or servers.

The decision has been made to consolidate at the server level – purchasing a small number of new high-powered servers to host all databases which were previously spread across many smaller servers, while also adding the High Availability features of RAC to the SALES & FINANCE databases. The vision of this consolidation therefore would look like the following:

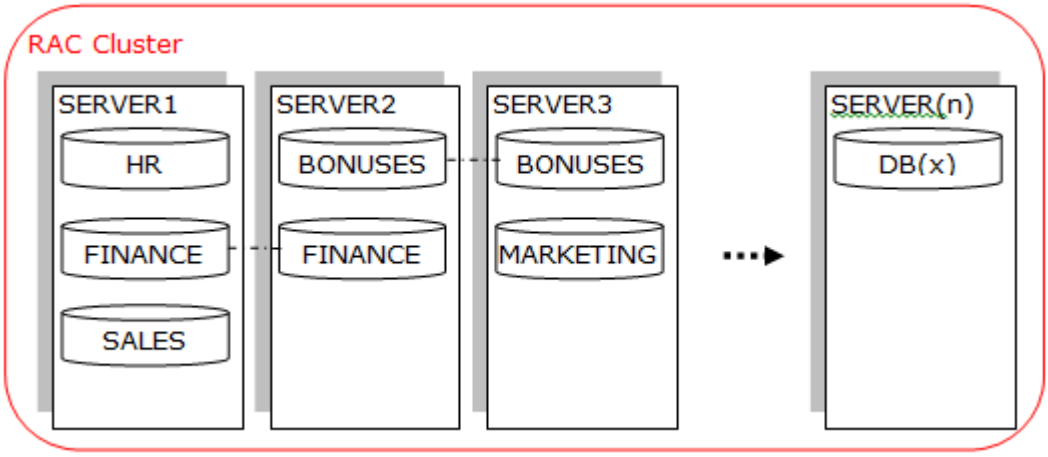

The process for this consolidation can be remarkably simple, and can be implemented with minimal disruption to availability. Depending on the nature of the application, the source operating system, and other factors – migrations can be performed with zero, or near-zero, downtime.

1. The first step, of course, is to provision the new servers – installing these into the data centre in line with standards, and installing the operating system of choice.
2. Install Oracle Grid Infrastructure. Installing only the latest version of Oracle Grid Infrastructure, this will form a cluster of all new servers, and allow as many versions of Oracle RAC Database to be installed in the cluster as necessary.
3. Install Oracle RAC Database. Ideally, this consolidation would be a perfect time to standardise versions of the database – upgrading databases as required. However, it is also possible, for expediency or other reasons, not to do this – and multiple versions of Oracle Database may be required. This is no issue – and multiple Oracle Database installations can be created, with varying versions and patch levels.
4. With the infrastructure in place – individual databases can be migrated to the consolidated platform at will, and upgraded to a standard version during the migration if desired. Various Oracle features can be used to aid the migration, depending on the circumstances for the migration. Oracle Recovery Manager (RMAN) allows migration using a simple backup/restore strategy, while Oracle DataGuard opens possibilities of a migration with minimal downtime.
5. Either during the migration, or at any time afterwards – single instance databases can be converted to One-Node RAC, or standard RAC databases – providing the full benefits of Oracle RAC technology, including High Availability with Transparent Application Failover, and increased flexibility of resource management.

The new estate may now look like the following:

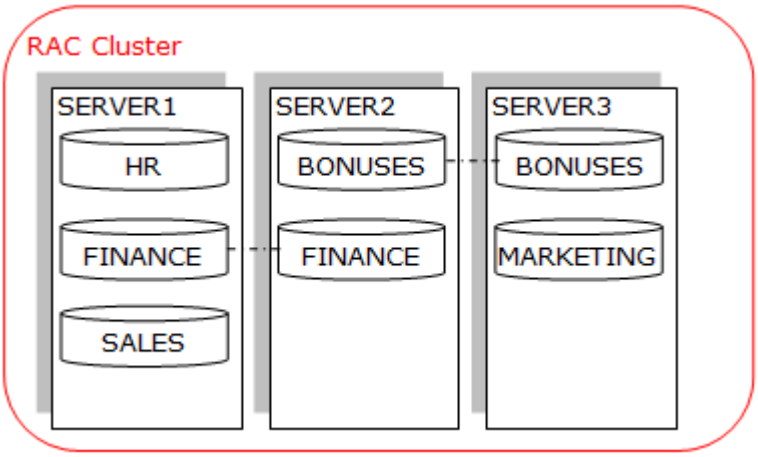

With minimal downtime, and without any new technologies – Server Consolidation has been achieved. However, rather than achieving the benefits of consolidation at a cost – the act of consolidation has actually improved the service provided by the database tier. These improvements are discussed below under Benefits of the RAC Consolidation.

Scenario 2 – Consolidate what you have

It is generally considered common knowledge that “you have to spend money to make money”, or in the case of consolidation – to save money. But, this needn’t be the case. It is possible to achieve server consolidation utilising existing hardware.

Considering the aforementioned statistic quoted by VMWare that on average 5-15% of any server's capability is utilised – there may not be any requirement to buy new higher-spec servers to replace existing servers. Instead – an organisation may opt to make better use of their existing hardware. The following scenario describes a process whereby all existing hardware is converted into a consolidated platform, allowing further growth without requiring additional hardware, and eventually making much more efficient use of existing hardware. During this process, of course, it may be decided that not ALL existing hardware should be reused, and that selected older hardware can be discarded during the process.

We start again with the estate described in scenario 1, however – the majority of the existing servers are still relatively young, and it has been deemed that these should not be discarded – when consolidation can easily be accomplished by reusing them. The old HR server, however, is nearing end of support, and the Finance server is running an operating system which is contrary to new standards – so both these servers will be discarded.

In this scenario, the consolidation is effectively taking place on in-use production servers. There is no migration involved – and changes are taking place in situ. However, the process can be performed in relatively small steps – with standard change control procedures governing the process throughout. The first step is to create a cluster consisting of all the servers which are intended to remain as the consolidated platform. In this particular scenario, we already have an existing RAC cluster – containing the BONUSES database. This RAC cluster is running the latest version of Oracle Grid Infrastructure – if it was running an older version – it should be updated to the latest version, to allow all versions of Oracle RAC Database to reside in the cluster. The first step, therefore, is to extend this cluster to include all other servers we intend to re-use.

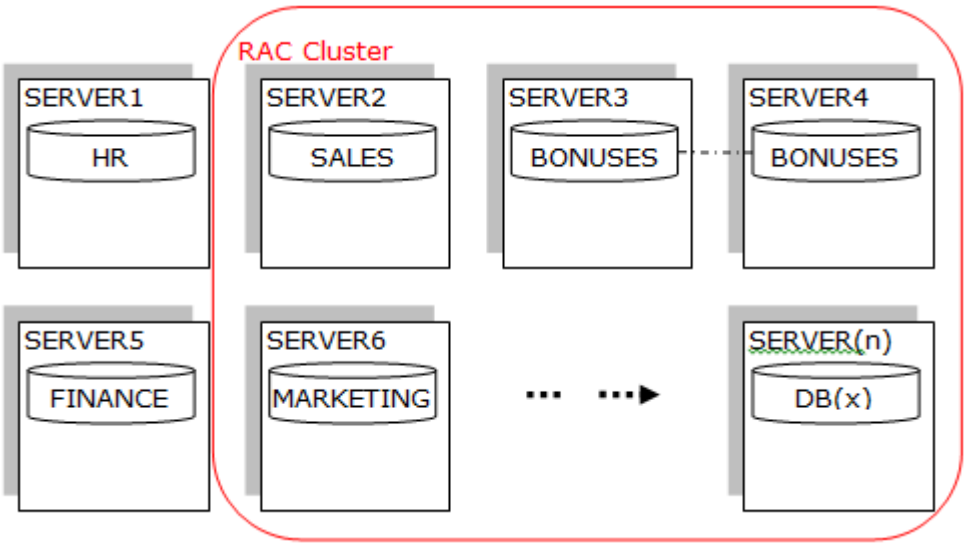

1. We now have a RAC cluster, with several single-instance databases residing on individual nodes. To allow for the use of RAC functionality, we now install Oracle RAC Database for all versions of the database we require.
2. The single instance databases can now be converted to multi-node or One-Node RAC databases – depending on High Availability requirements.
3. The databases residing on the servers to be decommissioned, HR & MARKETING – are simply migrated to the new cluster via RMAN or Oracle DataGuard, as per Scenario 1.

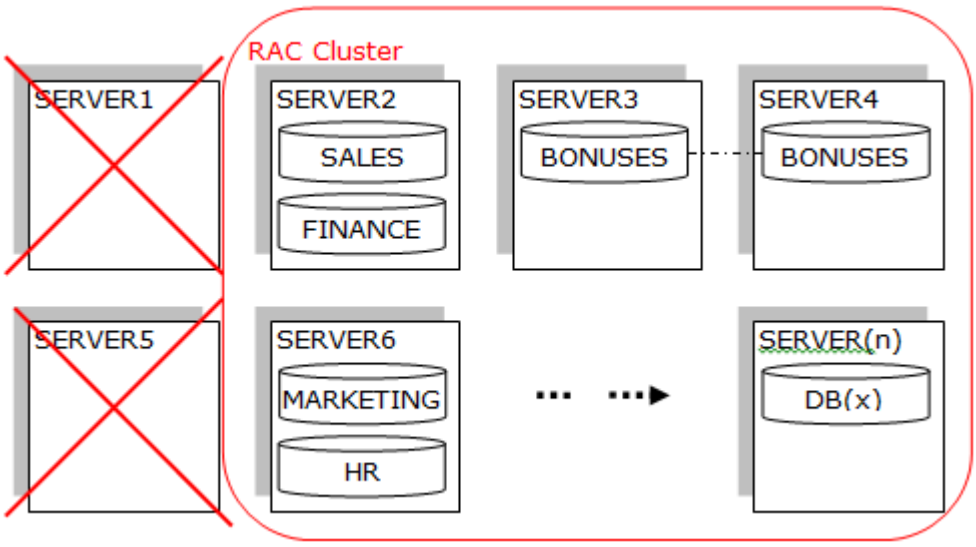

Benefits of the RAC consolidation

As mentioned previously, there are many benefits gained from Oracle RAC, on top of the results from server consolidation. With Oracle RAC, there are many situations where the end-result is far more positive than that in the pre-RAC environment.

Server crash

The most obvious case is also the most drastic. A hardware failure, or a kernel crash – can result in a server suddenly and definitively ceasing to function. Any single-instance database(s) residing on such a server will also become instantly unavailable. With Oracle RAC, in the case of any node in the cluster failing, other nodes will continue to function, providing services to clients. In many cases, a client will never be aware of the server becoming unavailable.

For example, with the estate described above, if Server5 was to crash, the database FINANCE becomes unavailable until the server is able to be restarted. Supposing the server crashed due to catastrophic hardware failure – this means that the FINANCE database is unavailable until either the faulty hardware is replaced, or the most recent backup can be restored to another server. In either case, this can conceivably be a period of several days during which the FINANCE database is unavailable. If this were to unfortunately occur during an important financial window, such as month-end, or even year-end – the impact could be enormous.

However, post-consolidation, with the FINANCE database sitting on a RAC cluster – the impact is drastically reduced, possibly eliminated. During consolidation, FINANCE database became a RAC database with 2 instances - therefore the loss of one server has no effect on the availability of the FINANCE service. The instances of FINANCE running on other server(s) continue to run, providing continuous availability. If, indeed, the FINANCE database had not been deemed worthy of multiple instances, but was instead converted to a One-Node RAC database – then at time of server failure, the database would also crash. However, the database could be configured to automatically restart on another node of the cluster – reducing the outage from days to minutes.

The difference between pre-consolidation and post-consolidation can be summarised as either being forced to explain to the Finance Director that a server crashed, and the Finance department will not have access to their data until next week; OR – having the option of mentioning in passing to the Finance Director that a server crashed, knowing that they would not have noticed.

Hardware maintenance

With Oracle RAC, trying to plan hardware maintenance becomes much simpler. The main issue with trying to schedule any downtime for a server is finding a suitable window for the outage. With Oracle RAC, however, there simply is no outage.

Supposing that a hardware fault has been identified within Server2 – the diagnosis is that the server will continue to function for some time, but a component needs to be replaced as soon as possible. However – the Sales team are adamant that their team are working 24x7 – and their work is far too important to cater for the required 2 hour outage of their SALES database.

With the old single-instance SALES database, this would mean that no outage would be possible, and eventually the hardware component would fail, resulting in an unplanned outage.

Post consolidation, however, the SALES database is a One-Node RAC. This allows the online migration of the SALES database from one node in the cluster to any other – without any interruption to the availability of the database – utilising Oracle One-Node RAC OMotion. Therefore, once the hardware maintenance requirement is identified – the SALES database can be moved from one server to another, transparently – and the server shutdown for maintenance – with zero downtime for end-users.

Scheduled patching

The regular requirement for software patching, whether operating system or Oracle – results in a similar situation to the above hardware maintenance.

With Oracle releasing Critical Patch Updates (CPUs) and Patch Set Updates (PSUs) quarterly – best practice would necessitate an outage of every single-instance database quarterly.

However, with a database infrastructure built on Oracle RAC, a “rolling patch” tactic becomes possible – whereby every node in the server is patched, one-by-one, with zero downtime.

Increased workload

Inevitably, workload for any given database may change over time – due to either increased amount of data, differing use of software, or implementation of extra modules. For example, the implementation of a new reporting tool may increase the load on the MARKETING database substantially.

In the case where the original server Server6 was underpowered to cope with this additional workload, the options are limited. It is necessary to either upgrade the server – involving an outage; or replace the server – with requisite migration and associated outage.

With the MARKETING database sitting on an Oracle RAC cluster, however – additional resources can be made available to the database dynamically. The database can either be moved to a more powerful server, using One-Node RAC’s OMotion feature (if the database was upgraded to 11gR2); or an additional instance can be added – scaling horizontally. Indeed, even if the existing hardware proves insufficient for the combined needs of all databases, additional hardware can be added to the cluster, and the resources made available to all databases in the cluster – all with zero downtime.

Oracle RAC vs Server virtualisation

One of the most common methods of consolidation is utilising virtualisation technologies. In fact, the terms 'consolidation' and 'virtualisation' are often used interchangeably. Many IT departments are in the midst of their "IT Consolidation Project" – of which a typical scope may actually be: "Take our 160 physical servers of varying sizes and specifications, and consolidate these into 160 corresponding virtual machines on 20 large physical servers".

One reason that virtualisation has become the primary method of server consolidation is that it allows a level of abstraction to the process. It is possible to reduce the number of physical servers, without making any other changes. Therefore, it is seen as a quick and easy method of achieving server consolidation.

However – although there are advantages with virtualisation, these do come at some cost. There will always be a negative impact on performance with virtualisation. Although virtualisation technology is continuously improving, there will always be a performance overhead – most notably due to the need for translating CPU instructions in the virtual server. Arguably the leader of the virtualisation industry – VMWare – currently demonstrates in their own performance benchmarking a 10% CPU overhead, compared to physical hardware.

Also – in the case of simply performing a "Physical-to-Virtual" (P2V) transformantion - at a logical level, nothing has changed – therefore the benefits of consolidation such as simplified management are not realised. In fact, after a pure "P2V" virtualisation/consolidation process – you have the same number of (now virtual) servers to administer – plus the physical hosts, and the underlying virtualisation software/operating system.

The following are the key areas which are seen as advantages of a virtualised infrastructure. In each area, using an Oracle RAC infrastructure can provide the same, if not better, benefits – as well as providing some additional benefits. This, of course, only applies to database servers – and virtualisation will likely have a part to play for consolidating other classes of servers. With database servers, however, all benefits can be gained using Oracle RAC, without any of the disadvantages of virtualisation.

Reduce hardware

The primary benefit of virtualisation is of course the actual physical server consolidation. This stems from the underutilisation of hardware. Virtualisation solutions aim to improve efficiency of hardware, thereby reducing the amount of hardware required – by hosting multiple virtual servers on one physical server.

However – one of the traditional downsides of virtualisation technology is the resource overhead that the actual virtualisation layer requires. As stated above, virtualisation software providers are continuously working to minimise this overhead, but this will always exist. The best-case scenario with current virtualisation technology can be expected to be at least 10% CPU, on top of any overhead pertaining to virtual networking and virtual memory. Indeed, VMWare's "Performance Team" has published results which show a nearly 20% reduction in throughput for an Oracle database residing on a virtual machine, reduced to 10% with a para-virtualised guest.

Oracle RAC, however, does not have any such overhead. Therefore, when consolidating a large number of database servers – an organisation can effectively achieve an additional 10% degree of consolidation, as compared to the use of a virtualisation strategy.

Simplify administration

Virtualisation commonly claims simplified and centralised administration/management as a benefit. This is true to an extent, in that a single view can be made via the virtualisation software of the current state of virtual servers. However – these virtual servers still consist of multiple installations of operating systems. Therefore – in effect, multiple 'servers' still need to be installed, configured, monitored, upgraded, patched, etc. And in the case of a pure physical-to-virtual consolidation, this will likely include different operating systems, different versions of operating systems, and different patch levels.

With an Oracle RAC consolidated architecture, however – there is a single operating system. With a single operating system installation per physical server to install, configure, patch, secure, upgrade, backup, and monitor - the workload of not only the physical server, but also the system administrator, is reduced.

Reduce time to provision servers

Another advertised benefit of virtualisation, is the greatly reduced time required to provision a new server. This, however, refers to the provisioning of a virtual server – on existing hardware. Therefore, there is not a direct correlation to the concept of an Oracle RAC infrastructure – other than comparing a non-consolidated environment to both a virtualised infrastructure and an Oracle RAC database infrastructure; relating to the process of requiring a new database.

In the non-consolidated environment – a new database would nearly always require new hardware. The time taken to provision a new database would therefore include the purchasing, delivery, and installation (racking, networking, etc) of a physical server; followed by the installation of an operating system, installation of Oracle software, and finally the creation of the database.

With a virtual infrastructure, this timeline greatly decreases. The virtual infrastructure administrator merely creates a new virtual machine – perhaps from a pre-existing template – and assigns it to a physical server, or server pool. The operating system is then installed, Oracle software is installed, and the database created.

With an Oracle RAC infrastructure, this timeline can be even shorter. In nearly all cases – the only step required is to create the database – using pre-existing hardware, pre-existing operating system, and pre-existing Oracle software. In very rare cases – a very specific version of Oracle software may be required – in which case one additional step of installing the Oracle software is added.

However, regardless of the platform used to create a consolidated infrastructure – there may always be a requirement to expand this infrastructure – either to allow for additional resource requirements which can not be met by the current underlying hardware, or simply to provide additional resources for performance and/or high availability flexibility.

With Oracle RAC 11gR2 – the Grid Plug and Play feature provides an extremely simply method for extending an existing cluster. A physical server can be added to the cluster, with GPnP negotiating appropriate network identities, acquiring configuration information from a profile, and making hostnames and addresses available on the wider network. Adding such a new server to an existing Server Pool immediately makes it available for existing services, which dynamically make use of the new resources provided.

High availability

Another stated claim of the leading virtualisation provider is supplying High Availability features – "decrease downtime and improving reliability with business continuity and built-in data disaster recovery".

Virtualisation can certainly accomplish this – with online migration of virtual machines between physical hosts to reduce maintenance outages, and automatic fail-over in case of physical hardware failure.

However, Oracle RAC – provides the same capabilities, only in a manner designed specifically for Oracle databases, therefore doing so with a little more elegance than the equivalent capability of virtualisation.

Virtualisation software typically provides online migration between physical hosts – by mirroring the complete memory state of the virtual machine, quiescing this, and eventually switching to the new, migrated virtual machine. With Oracle databases, however, there can be a huge amount of constantly changing data in memory (ie: the SGA). With even a medium workload on an Oracle database residing on a virtual machine, the process of mirroring memory between two physical hosts may take an inappropriately long time – or under a very heavy workload, may even be impossible to complete. With Oracle RAC, however, whenever an instance is migrated, either due to a Server Pool dynamically changing resource allocation, or a manual invocation of Omotion for RAC One Node – the workload is actually split between target and source nodes during the migration, without any need for quiescing the instance – and even the heaviest workload can be easily migrated.

Also – Oracle RAC allows online maintenance of the actual base software – the operating system and/or Oracle software. While virtualisation allows for migration away from a physical host, to allow maintenance of the host – there is still no way to patch the operating system or Oracle software without downtime. An Oracle RAC cluster allows this – as all instances are merely migrated to other nodes, to allow patching or maintenance of the operating system and/or Oracle software – usually performed in a rolling fashion across all nodes of the cluster.

In the case of automated failover after physical hardware failover, virtualisation software works at the virtual server level. Therefore, in the case of an Oracle database running on a virtual server – the failover is not transparent. The actual virtual machine does go down, and is brought back up on an alternative physical host once this is detected. The process of failover does still incorporate an unscheduled outage.

With Oracle RAC, the Transparent Application Failover feature allows the equivalent scenario to result in zero downtime, and end-users may never be aware there was an issue at all.

Summary

Oracle Real Application Clusters has always provided extremely high availability to mission critical Oracle databases. Functionality such as Transparent Application Failover allows applications, and even currently running sessions, to continue without disruption throughout the failure of a server. However - with features such as Server Pools, RAC One Node with Omotion, and Instance Caging – Oracle RAC 11g Release 2 now provides a viable alternative to virtualisation for server consolidation. The Server Pools feature allows for dynamic reassignment of physical hardware/resources – according to changes in hardware availability and workload priorities; while RAC One Node with Omotion allows for dynamic movement of database instances between physical nodes for purposes of online upscaling or maintenance; and Instance Caging ensures that any one Oracle instance cannot starve another of CPU resource.

Using Oracle RAC 11gR2 as a basis for a “Database As A Service” (DBAAS) type infrastructure tier provides multiple benefits as opposed to performing a straight physical-to-virtual migration of individual database servers. All the key goals of a server consolidation and/or virtualisation project can be met instead with an Oracle RAC infrastructure – plus several additional benefits.

A greater level of server consolidation can be achieved, as the full capability of physical hardware can be utilised, without the overheads of multiple operating systems, CPU translation, or routing of other virtual hardware devices to physical.

Administration of the environment is simplified greatly. Instead of multiple operating systems residing in virtual machines, a single operating system needs to be maintained by the system administrator – resulting in a fraction of the work effort required for maintenance, monitoring, backups, etc. Also – no new skills are required for system administration. The operating system will be a single operating system, chosen according to the organisation's standards - and the Oracle Grid Infrastructure – providing the RAC cluster - can be supported by the existing Database Administrator team.

Oracle RAC provides superior protection from failures – rather than a physical hardware/host failure causing a virtual database server to effectively also fail, before being restarted on an alternative host – an Oracle RAC Cluster allows currently connected database sessions to automatically failover from a failed node to another node, in some cases while being able to continue the currently running query. With One Node RAC – an instance will automatically restart on an available node – without waiting for an operating system to start.

Online maintenance becomes extremely simple with an Oracle RAC infrastructure. An Oracle RAC database can have any single instance shut down at any time, without impacting current or future client connections, while with the One Node RAC feature – a single instance database may also be migrated between nodes, with zero impact on connections, and without any need for mirroring memory structures, or quiescing of the database. Oracle RAC also allows even the patching of operating system and Oracle software, on top of hardware maintenance.

Workload management and efficiency is increased utilising an Oracle RAC consolidated platform. Server Pools allows for grouping of services together on a subset of physical nodes – while allowing for these physical nodes to be dynamically reallocated between these pools, depending on requirements. Allowing instances of multiple databases to share physical node, utilisation of an organisation's hardware can become much more efficient, ideally ~85%, while Instance Caging ensures that each instance will have access to resources when required.

The flexibility of a database tier based on Oracle RAC has become even more impressive with the features in 11gR2. The provisioning of a new database can now be achieved with a simple "CREATE DATABASE" command, or a few mouse clicks in Oracle Enterprise Manager or Database Configuration Assistant – as opposed to the lead time of ordering new hardware before installing operating system and Oracle software with a siloed "one-database-per-server" tactic, or even the time of provisioning a virtual machine – where the installation of operating system and Oracle software is still required.

Expansion of such an infrastructure is also simplified, with Oracle's Grid Plug and Play feature. The use of the Global Naming Service allows for dynamic addition, and removal, of physical servers from the cluster – without any need for reconfiguring DNS entries; and the use of a Single Client Access Name allows for the same dynamic changes to the underlying hardware - without any reconfiguration of application/client connection details.

Although this paper has focused on the use of Oracle RAC for consolidating existing Oracle databases into a flexible, robust, efficient and consolidated infrastructure – Oracle RAC technology can also be used as the platform for other applications. The Clusterware component of the Oracle Grid Infrastructure is easily configured to provide Oracle RAC's excellent features for any application which can be run on the underlying operating system. It is, in fact, possible to utilise Oracle Real Application Clusters to provide a single underlying infrastructure for all mission critical databases and applications.

Virtualisation, of course, does provide some complementary features. An organisation may have an extremely small amount of rack space - whereby a logical use of virtualisation would be to complement Oracle RAC in creating one or more development/test Oracle RAC clusters – replicating a production cluster. Utilising virtual machines as nodes of an Oracle RAC cluster allows all the benefits of online maintenance provided by Oracle RAC, allowing patching of the oracle software and guest operating system with zero downtime – ideal for time critical development and testing.