



September 15, 2005

Oracle RAC Gains Momentum

Reliability And Cluster Of Low-Cost Servers Are The Key Reasons For Growth

by Noel Yuhanna

EXECUTIVE SUMMARY

In the past year, Forrester has interviewed 35 Oracle customers using Oracle Real Application Clusters (RAC) to find out about their implementation. Most customers stated that RAC met or exceeded their high-availability requirements to support mission-critical database applications. Oracle RAC has come a long way in establishing itself as a key technology in delivering unparalleled high availability and scalability solutions for Oracle databases. Forrester estimates that there are more than 1,200 customers who are currently using RAC in production, and this is likely to double in the next 12 to 18 months as enterprises look to deploy low-cost commodity servers and extend their applications to support higher availability.

RAC HAS BECOME A KEY TECHNOLOGY DIFFERENTIATOR FOR ORACLE

Oracle RAC addresses two key enterprise requirements for mission-critical database applications: high availability and scalability. Although most enterprise database management systems (DBMS) products support high-availability and scalability features, Oracle RAC delivers the most innovative and superior solution in the industry. It supports all kinds of applications, such as data warehouses, OLTP, or mixed environments. Oracle continues to make considerable investment in R&D for RAC, and this is visible in the 10g release, where RAC is one of the key components in Oracle's Grid solution. Oracle claims to have already tested RAC with more than 100 node clusters and continues to work on improving performance and scalability.

RAC Adoption Has Increased Because Of Several Factors

Although Oracle RAC was released in June 2001, it existed for several years under another product name, Oracle Parallel Server (OPS), in the early '90s. But it has only been during the past two years that RAC has shown significant adoption in the industry. This increase in adoption has been mainly due to:

- **The reliability of the RAC solution.** RAC has matured over time, delivering high reliability and stability in a high-availability solution. Although Oracle9i RAC has had initial issues with memory and recovery management, Oracle9i Release 2 RAC and Oracle RAC Database 10g have proven to be reliable in most customer deployments. Customers who have migrated to RAC 10g in the past year have found that it delivers the best reliability for their Oracle environment.
- **The simplicity to manage a large number of nodes.** Oracle9i RAC required more administrative resources to manage RAC, mainly because of the complexity in setup, install, and management.



Headquarters

Forrester Research, Inc., 400 Technology Square, Cambridge, MA 02139 USA
Tel: +1 617/613-6000 • Fax: +1 617/613-5000 • www.forrester.com

However, with Oracle Database 10g, RAC offers a higher degree of automation in managing the cluster and simplified administration using Enterprise Manager Grid Control 10g. We are likely to see even more automation in RAC management in the coming years.

- **The clustering of low-cost commodity servers.** Today, more customers are leaning toward using low-cost commodity servers that run Linux and Windows to support their mission-critical database applications. Using RAC that runs on low-cost servers, customers claim to have saved 40% or more over large expensive symmetric multiprocessing servers. In 2003, the maximum number of RAC nodes used in any production environment was eight, and this number has doubled in 2005. We are likely to see a RAC cluster with as many as 32 nodes in the next two years.

RAC Still Has Limitations That Oracle Needs To Address

Although RAC offers a good high-availability and scalability solution, it still has limitations. It still:

- **Does not support heterogeneous infrastructure.** RAC technology depends highly on identical hardware servers; even the operating system and patch levels need to be identical for RAC to work. This often means acquiring new identical servers for RAC deployment. Although Oracle has touted a heterogeneous RAC, at present, there are no indications that this will be available any time soon.
- **Does not support availability outside of a data center.** RAC technology can only be used with shared disk storage, such as storage area networks or network-attached storage. Therefore, it imposes a limit on how far nodes can be located, and this is why it is generally confined within a data center. RAC delivers limited disaster-recovery capability and therefore needs to be supplemented with Oracle Data Guard (log shipping) or remote disk mirroring technology.¹

Other DBMS Vendors Still Lag Behind, For Now

Although other DBMS vendors offer high-availability features, Oracle RAC delivers the most innovative and reliable solution in the industry. RAC is based on shared-disk architecture, while IBM DB2 Universal Database (UDB) uses a shared-nothing architecture and Microsoft SQL Server uses a federated database architecture. Compared with other DBMS vendor products, Oracle RAC offers several advantages, including the capability to recover in-flight transactions on a failed node using Transparent Application Failover (TAF) APIs, recovering transactions quickly in a non-TAF environment and offering both scalability and high-availability functionality together. However, both IBM and Microsoft continue to extend their high-availability solutions.

- **IBM's DB2 UDB offers several high-availability features.** DB2 UDB offers many high-availability features, including a data partitioning feature (DPF), high-availability disaster recovery (HADR), log shipping, and failover clustering. The Informix HADR feature was

recently made available to DB2 UDB V. 8.2, which supports automatic client redirection and database failover capability. DB2 UDB uses DPF to deliver scalability for data warehousing applications using scale-out, shared-nothing architecture.

- **Microsoft's SQL Server is gearing up with new high-availability features.** SQL Server offers log shipping and an operating-system-level failover clustering solution for high availability. Starting with SQL Server 2005, a new feature called database mirroring will become available to support high-availability requirements, as well. The database mirroring feature uses a separate server to monitor two or more database servers that replicate transactions, and in the event of a failure, clients are connected to another server.

Case Study: Overstock.com

Overstock.com is an online retailer that offers discount, brand-name merchandise for sale over the Internet. The company offers its customers an opportunity to shop for bargains conveniently, while offering its suppliers an alternative inventory liquidation distribution channel. Overstock.com, headquartered in Salt Lake City, is a publicly traded company with an annual revenue of more than \$600 million.

Forrester spoke with Shawn Schwegman, the senior vice president of technology at Overstock.com. He claimed that the company's mission-critical shopping application ran on Oracle RAC, in addition to several other critical applications, including customer services and distribution application. "All of our critical applications in our company run on Oracle RAC — basically, everything we do today is against Oracle RAC databases," he told us.

Overstock.com has several RAC clusters to support these application, the largest being a nine-node RAC cluster that runs the shopping Web site application. Although the company currently uses Dell PowerEdge 2650s running SUSE LINUX with EMC Symmetrix DMX3000 storage array, it is in the process of migrating to IBM eServer pSeries running AIX with EMC storage. Why the migration? According to Schwegman: "Linux's version of Oracle 10g did not come out on the IBM pSeries until recently, so we wanted to give it a year before we migrated. We are going to AIX for one or two years, before Linux becomes mature on the pSeries box. But eventually, we will be using pSeries with Linux."

Schwegman says that the IBM pSeries servers will deliver more performance and scalability for mission-critical applications, and AIX with RAC has been very stable. He claims that RAC:

- **Supports large mission-critical applications that need high availability.** Overstock.com claims that its Web site supports more than 1.5 million unique visitors per day. Its largest RAC database is around 700 GB, with around 2 TB of disk storage allocated. Schwegman stated that if one node goes down in the RAC cluster, it still provided the company with availability for its

mission-critical applications. “We found that RAC performs extremely well and did not find any performance degradation,” he said.

- **Accommodates growth very easily by adding additional nodes.** One of the key benefits of using RAC was that Overstock.com could easily add more nodes in a cluster at any time to support additional workload. With nine nodes in a cluster, the company is already looking to add more in the near term.
- **Delivers stability to the environment, despite some initial hiccups.** Overstock.com did have initial issues with Oracle9i RAC. As Schwegman puts it: “Oracle9i RAC was a little buggy initially, but we found Oracle 10g RAC to be very stable. And it’s not an Oracle9i RAC issue but memory management issues with a Red Hat operating system that caused problems. However, we moved away from Red Hat to SUSE, because they have a better memory module, and that is essential for our high-end performance requirement. Stability and performance in 10g RAC is by far the best I have ever seen from Oracle.”

Case Study: Starwood Hotels & Resorts Worldwide

Starwood Hotels & Resorts Worldwide is one of the leading hotel and leisure companies in the world, with 733 properties in more than 80 countries. Forrester spoke with Arup Nanda, the director of database engineering and operations at Starwood, to understand the company’s RAC implementation.

Nanda claimed that Starwood’s business depends on Oracle RAC, which supports several mission-critical applications. Starwood’s key hotel and resort reservation application runs on RAC and is used by customers, travel agents, and internal staff. Starwood currently uses several RAC clusters for its applications, and the main database comprises a two-node RAC cluster running HP servers, each with 16 processors and 64 GB of RAM on HP-UX. Nanda mentioned that Starwood soon plans to go with three nodes to support increased workload. Starwood’s current workload executes more than 10,000 concurrent transactions at any given point in time. He claims that RAC:

- **Delivers very high availability for their critical application.** Starwood was of the opinion that it could take a single node down if any work needed to be done on it, such as hardware maintenance or changes, and these could be done without affecting the overall availability of the application — overall, it has been transparent. Nanda did mention that the company’s availability for mission-critical applications has been at its highest level ever, which has improved its service-level agreements consistently.
- **Maintains its simplicity but requires the right resources.** Nanda stated, “Anything that is complex would be difficult to manage, but the whole issue with RAC is getting the right resources to manage your RAC environment, because it’s a brand-new thing.” He is of the opinion that RAC expertise is difficult to find at the moment but is likely to change as more customers use this technology.

- **Provides the capacity on demand to meet growing business demands.** Starwood uses RAC for OLTP applications, and this allows it to increase and decrease the computing power by simply adding or removing nodes in real time. Nanda gave an example, saying that during peak season, Starwood was able to add more nodes to support increased usage and was even able to easily move some of its data warehousing nodes to OLTP nodes.

RECOMMENDATIONS

RAC SHOULD BE PART OF YOUR ORACLE DBMS STRATEGY

To be successful, Oracle RAC deployment requires careful planning, testing, and a good understanding of its architecture. Enterprises looking to deploy RAC should:

- **Start small and grow — don't rush.** RAC deployments take time and effort, and it is best to start with a few nodes initially to adapt to your environment and then add more nodes to deliver increased availability and scalability.
- **Get administrators trained on RAC.** While Oracle RAC 10g has improved dramatically over Oracle9i RAC, it is still not your typical Oracle database. It requires a certain level of expertise and knowledge; thus, you should train your database administrators.
- **Have a similar or scaled-down RAC test configuration.** When deploying patches or upgrades on RAC environments, you should apply them first in a test environment that is configured similarly to ensure that they do not cause any problems.
- **Understand that not all applications need RAC.** Not all databases need high availability and scalability; therefore, your DBMS strategy should be to focus on only those key applications that require such functionality.
- **Look at both availability and scalability requirements together.** RAC offers both availability and scalability. Look at both of these options together when choosing your hardware platform.
- **Understand RAC's limitations in disaster recovery, and plan accordingly.** Consider using RAC along with Oracle's Data Guard or other disk mirroring features for better protection against disaster situations.
- **Consider RAC 10g over Oracle9i RAC because of increased reliability.** Oracle RAC 10g has proven to be more stable and reliable in customer installation, in addition to requiring less effort in cluster manageability.

ENDNOTES

- ¹ Although the majority of Oracle customers are using RAC within a data center, there are a few enterprises that support RAC nodes that are up to 25 km apart.