

# Reducing Oracle RAC Wait Events by Using Instance-Specific Block Allocation for Production Applications

Murali Natti

Lead Database Engineer | DevOps Lead | Database Architect | Cloud Infrastructure Solutions Expert | DB Security Lead

---

## Article Info

### Article history:

Received Aug, 2023

Revised Aug, 2023

Accepted Aug, 2023

---

### Keywords:

Application Performance

Cache Fusion

Data Partitioning

Database Scalability

High Availability

Instance-Specific Block

Allocation

Load Balancing

Oracle RAC

Resource Contention

Wait Events

---

## ABSTRACT

Oracle Real Application Clusters (RAC) is a robust and high-availability solution designed to enable multiple database instances to share the same physical database, offering benefits such as scalability and fault tolerance. However, while Oracle RAC can support critical production environments, it introduces significant complexities, particularly with regard to wait events. One of the primary performance bottlenecks in Oracle RAC is inter-instance communication, commonly referred to as cache fusion, where instances must exchange and synchronize shared data blocks across the cluster. This overhead becomes particularly problematic for production applications that access commonly used tables or objects, leading to increased wait times, slower response rates, and reduced throughput. This paper outlines a novel approach that aims to alleviate Oracle RAC wait events by binding specific application instances to individual Oracle RAC nodes. By allocating frequently accessed tables or objects to specific nodes, this method reduces contention, optimizes database access, and enhances overall application performance.

*This is an open access article under the [CC BY-SA](https://creativecommons.org/licenses/by-sa/4.0/) license.*



---

## Corresponding Author:

Name: Murali Natti

Institution: Lead Database Engineer | DevOps Lead | Database Architect | Cloud Infrastructure Solutions Expert | DB Security Lead

Email: [murali.natti@gmail.com](mailto:murali.natti@gmail.com)

---

## 1. INTRODUCTION

The Challenge of Oracle RAC Wait Events Oracle Real Application Clusters (RAC) [1] is a high-availability solution designed for large-scale, critical production environments. It allows multiple instances to access a single database, thereby providing horizontal scalability, fault tolerance, and continuous availability. However, the distributed nature of Oracle RAC also introduces inherent complexities. One of the most significant challenges faced by organizations running Oracle RAC in production environments is the occurrence of

high wait events, particularly related to inter-instance communication. In an Oracle RAC environment, multiple instances share data blocks stored in a centralized storage location. When an instance requires a data block that another instance currently holds in memory, a communication process known as cache fusion [2] is triggered. This process allows one instance to transfer a copy of the data block to another, enabling it to continue processing. While cache fusion is essential for ensuring data consistency across the cluster, it comes with a substantial performance overhead. As instances must frequently synchronize shared data blocks, wait events such as gc cr request,

gc buffer busy, and gc current block busy can cause significant delays in application performance. This behavior is particularly detrimental to applications that frequently access tables or data objects with high contention. The Impact on Applications For production-critical applications, high wait times can lead to noticeable performance degradation. When applications spend considerable time waiting for data blocks to be transferred between instances, increased application latency, resource contention, and performance inconsistency become common. The most direct consequence of cache fusion-related wait events is increased latency in application transactions. As database blocks are moved between nodes, the response times for queries and updates are significantly delayed, resulting in slower transaction processing. The excessive communication required between instances to transfer data blocks consumes CPU, I/O, and network resources inefficiently. This not only hinders the performance of the database but also increases the overall resource load on the system, contributing to potential bottlenecks. The performance of applications in Oracle RAC environments can be highly inconsistent, especially during peak periods. High contention for shared blocks may cause fluctuating response times, leading to unpredictable performance behavior under heavy workloads. This paper proposes a solution that minimizes inter-instance communication, thereby addressing these issues. By binding specific application instances to individual Oracle RAC nodes and allocating commonly accessed tables or objects to those nodes, we aim to reduce contention, optimize data access, and enhance overall application performance.

## **2. THE PROPOSED SOLUTION: INSTANCE-SPECIFIC BLOCK ALLOCATION FOR CRITICAL APPLICATIONS**

In traditional Oracle RAC configurations, all instances within the cluster share access to the same set of data blocks. This shared access increases the likelihood of cache contention, particularly when multiple

instances access the same blocks simultaneously. As multiple instances attempt to read or modify the same data blocks, cache fusion [2] operations must synchronize and transfer these blocks between instances, leading to increased latency and resource overhead. This contention results in high wait events, degraded system performance, and significant slowdowns in transaction processing, particularly for high-traffic applications that require frequent access to shared data objects. The proposed solution seeks to mitigate these issues by binding application workloads to specific Oracle RAC nodes and configuring instance-specific services for critical applications. The primary goal of this approach is to ensure that frequently accessed tables, partitions, or objects are predominantly accessed by a designated instance, thereby reducing the need for inter-instance communication. This can be achieved by implementing a targeted allocation strategy where database resources are distributed based on access patterns and workload requirements.

By limiting the number of instances that need to access a particular set of data blocks, the solution effectively reduces inter-instance communication required for cache fusion. This minimizes wait events associated with global cache management, leading to improved performance, lower transaction latency, and a more balanced utilization of database resources. Moreover, this instance-specific allocation approach allows database administrators to implement fine-tuned performance optimizations [3], such as prioritizing access to specific data partitions, reducing network traffic between nodes, and improving response times for mission-critical applications. Through careful analysis of workload distribution, database administrators can design an allocation strategy that ensures optimal data locality. This involves evaluating historical access patterns, transaction frequency, and contention hotspots to determine which data objects should be assigned to specific instances. Additionally, integrating load-balancing mechanisms and failover

capabilities ensures that the system remains highly available and resilient, even in the event of node failures. By adopting this structured and instance-specific block allocation approach, organizations running Oracle RAC can significantly enhance database performance [4], reduce contention-related delays, and ensure a more efficient and scalable production environment.

### 3. HOW THE SOLUTION WORKS

**Step-by-Step Breakdown** The first step in the implementation of this solution is to analyze the application's database workload and identify the tables or objects that are frequently accessed. High-traffic data blocks that are commonly accessed by multiple instances tend to cause the most contention when distributed across the entire cluster. To achieve this, tools such as Automatic Workload Repository (AWR) reports, V\$Session, and V\$SQL views can be utilized to pinpoint high-volume tables or objects that contribute to inter-instance contention. Once the critical data objects have been identified, the next step is configuring node-specific services in Oracle RAC. Each node within the RAC cluster is configured with a dedicated service that binds specific application workloads to the corresponding node. This service configuration ensures that applications connect only to the nodes that store their most relevant data. The Oracle RAC service configuration is designed to improve locality by ensuring that an application instance accesses data from the same node, thus reducing the need for cross-node cache fusion [5]. The most critical step in reducing wait events is the partitioning of high-access tables and data objects. In Oracle RAC, it is possible to partition tables so that specific partitions are stored on specific nodes. This partitioning reduces the need for inter-node communication by ensuring that frequently accessed data is localized to a particular node. After configuring services and partitioning data, it is essential to validate the solution's effectiveness. This can be achieved by benchmarking the application performance before and after implementing the solution.

### 4. RESULTS

The implementation of this solution led to notable improvements in application performance. Specifically, there was a significant reduction in wait events related to cache fusion [5], including gc cr request, gc buffer busy, and gc current block busy. These wait events, which are typically caused by excessive inter-instance communication, were substantially minimized due to the instance-specific block allocation strategy. By partitioning data and binding application workloads to specific nodes, the data that the application required was largely available locally, reducing the need for inter-instance block transfers. This resulted in fewer global cache synchronization requests, leading to a more streamlined and efficient database operation. The solution resulted in a drastic decrease in application response times, both for read and write operations. As inter-instance communication was minimized, transaction processing became faster, leading to enhanced user satisfaction and better application performance. Business-critical operations that previously suffered from high wait times and unpredictable performance [6] behavior saw measurable improvements in speed and reliability. The application performance became more predictable, allowing for better capacity planning and improved user experience.

The distribution of the application workload across the nodes was more balanced, leading to better CPU and I/O utilization. Since each instance predominantly accessed its local data blocks, there was a significant reduction in contention for shared resources. This balance improved overall system throughput and ensured that computing resources were efficiently allocated, preventing CPU and memory bottlenecks. In addition, unnecessary network traffic between nodes was reduced, improving overall system efficiency and lowering the operational costs associated with excessive data movement. Comprehensive performance monitoring and benchmarking confirmed these improvements. Metrics such as transaction latency, query execution times, and system throughput indicated a consistent

reduction in wait times and enhanced overall database efficiency. By adopting this approach, organizations were able to significantly improve their Oracle RAC environments, leading to a more scalable and high-performance database infrastructure.

## 5. CONCLUSION

By binding critical application workloads to specific Oracle RAC nodes and optimizing the partitioning of high-traffic data objects, this solution provides a scalable and efficient method for reducing Oracle RAC wait events. The targeted allocation of frequently accessed data minimizes inter-instance communication, particularly cache fusion, which is one of the most significant contributors to performance bottlenecks in Oracle RAC environments. By implementing this strategy, database administrators can significantly enhance application response times, reduce contention for shared resources, and improve overall database performance. This approach is especially beneficial for large-scale applications that experience high-volume, transactional workloads with significant cache contention. By ensuring that applications primarily access data stored on

their designated nodes, organizations can mitigate the inefficiencies associated with excessive cache fusion operations. The localized access strategy not only improves query and transaction speeds but also enhances database stability and consistency. Furthermore, the scalability of this solution allows businesses to accommodate growing workloads without experiencing significant degradation in performance. As data volumes and user demands increase, organizations can extend this strategy by further refining data partitioning [7] and workload distribution to maintain optimal efficiency. This ensures that Oracle RAC continues to deliver high availability [8] and fault tolerance while maintaining superior performance. By reducing wait events and optimizing data access, businesses can achieve better scalability, improved resource utilization, and enhanced end-user satisfaction, all while maintaining the high availability [8] and fault tolerance that Oracle RAC is known for. The instance-specific block allocation approach presents a sustainable and effective method for optimizing Oracle RAC environments, ensuring that mission-critical applications run smoothly with minimal latency and maximum efficiency.

## REFERENCES

- [1] Oracle Corporation, "Oracle Real Application Clusters Documentation," 2023. [https://docs.oracle.com/cd/E11882\\_01/rac](https://docs.oracle.com/cd/E11882_01/rac).
- [2] Oracle Support, *Reducing Cache Fusion Wait Events in Oracle RAC*. Oracle White Papers., 2020.
- [3] J. Smith, *High-Performance Oracle RAC: Strategies for Optimization*. Pearson, 2021.
- [4] I. M. Review, "Case Studies in Oracle RAC Performance Optimization," 2021. <https://www.itmanagementreview.com>
- [5] Database Journal, "Understanding Cache Fusion Performance in Oracle RAC," 2023. <https://www.databasejournal.com>.
- [6] P. Kumar, *Advanced Techniques in Oracle RAC Performance Tuning*. Springer., 2020.
- [7] S. P. Blog, "Improving Oracle RAC Performance Through Data Partitioning and Instance Services," 2022. <https://www.sqlperformance.com>.
- [8] Oracle Press, *Oracle RAC Best Practices for High Availability*. McGraw-Hill., 2022.