Load-Skew–Driven Cache Fusion Traffic Characteristics in Oracle RAC Clusters

Sophia Renwick, Marcus Ellery

Abstract

Cache Fusion efficiency is critical to the scalability and performance of Oracle Real Application Clusters (RAC), where global cache coherency depends on rapid block transfers and synchronized ownership arbitration across nodes. This study evaluates Cache Fusion traffic patterns under controlled heterogeneous node load conditions, introducing selective CPU throttling, memory pressure, and resource contention to isolate the effect of node performance asymmetry on cluster coherency. Results show that a single degraded node significantly increases global cache transfer latency, hot block ownership churn, and GCS/GES queue depth, even when it processes a smaller share of total workload. Performance metrics confirm that RAC throughput decreases by up to 37% under heavy throttling, whereas removal of the impaired node restores stability and improves throughput despite reducing cluster size. These findings establish that RAC scalability depends more on latency uniformity across nodes than on node count or compute volume. The study concludes that real-world RAC optimization must prioritize equalized node responsiveness, interconnect determinism, and proactive latency anomaly detection to prevent cascading coherency stalls in mixed workload deployments.

Keywords: Cache Fusion; Oracle RAC; Global Cache Synchronization; Cluster Latency Uniformity

1. Introduction

Oracle Real Application Clusters (RAC) enable multiple database instances to operate concurrently against a shared set of datafiles, coordinating cache state through Cache Fusion, a mechanism that transfers data blocks between buffer caches of different instances to maintain global cache coherency. The efficiency of Cache Fusion communication strongly influences cluster throughput, commit latency, and overall workload scalability. Foundational studies on shared-disk coordination and synchronized state consistency emphasize that reliable operation depends on disciplined propagation of shared state and controlled contention under concurrent access [1,2]. In cloud-deployed RAC systems, where resource pools and I/O characteristics vary dynamically, workload behavior often reflects sensitivity to how data blocks migrate across instances under mixed OLTP and analytical traffic [3]. Empirical observations from enterprise operational environments show that when nodes differ in CPU, memory, or I/O bandwidth capacity, global coordination mechanisms may become unbalanced, leading to non-uniform access latencies [4,5]. Similar sensitivity to synchronized state propagation has been observed in distributed application-tier deployments that rely on consistent execution semantics across multiple runtime endpoints [6,7].

When RAC nodes process heterogeneous workload intensities, the distribution of hot blocks may become skewed toward specific instances, increasing the rate of global cache transfers. Performance modeling research demonstrates that under mixed transactional workloads, transaction locality, update frequency, and buffer reuse patterns drive synchronization overhead and interconnect pressure [8,9]. In environments with irregular or burst-type access patterns, anomaly-driven cache hotspots may emerge, requiring careful monitoring of ownership transitions to prevent coherency stalls [10,11]. Systems that front-end RAC databases with orchestration or analytics layers can further amplify these effects, as multi-session stateful workflows implicitly concentrate access patterns on selected instances [12,13].

At the protocol level, Cache Fusion minimizes disk I/O by transferring current block images directly between instances, trading physical I/O latency for interconnect dependency. Shared-disk concurrency research shows that as contention rises, synchronization messages, lock remastering, and coordination cycles dominate latency behavior rather than disk operations [14,15]. Maintaining predictable behavior under these conditions requires consistent metadata propagation and disciplined control of authorization domains that govern instance ownership decisions, particularly during runtime load shifts [16,17].

Cluster interconnect topology plays a decisive role in Cache Fusion efficiency. Oracle RAC operational guidance emphasizes that private network tuning, buffer sizing, and retransmission thresholds significantly affect read-intensive versus write-intensive workload performance [18,19]. In heterogeneous compute environments, differences in NUMA layout, processor architecture, and memory latency can cause diverging block residency lifetimes across instances, even under identical configuration, increasing effective global cache access cost [20].

Workload scheduling and load-balancing behavior further shape Cache Fusion traffic. Studies on RAC service placement show that uneven CPU or I/O saturation can unintentionally bias workload affinity

toward specific instances, increasing synchronization overhead [21]. Latency amplification research indicates that under high-contention commit phases, localized congestion can propagate across the cluster, degrading global performance even when only a subset of nodes are stressed [22]. Runtime integration patterns observed in hybrid application–database deployments confirm that scaling behavior is highly sensitive to compute provisioning and session multiplexing strategies [23]. In low-code or orchestration-driven access frameworks, workload characteristics may shift rapidly as user-driven or automated workflows alter query locality, further modifying cache block residency patterns over time.

Empirical observations in large-scale enterprise deployments reinforce that Cache Fusion performance is governed not merely by cluster size, but by the interaction between workload variability, resource heterogeneity, and access locality dynamics. Ensuring long-term stability therefore requires aligning RAC configuration with disciplined data engineering practices that regulate workload shape and synchronization pressure. Automation strategies for workload orchestration and execution governance further help mitigate cache instability by smoothing access patterns and reducing sudden contention spikes across instances.

2. Methodology

The methodology for analyzing Cache Fusion traffic under heterogeneous node loads in Oracle RAC was designed to isolate the workload, interconnect, and buffer state variables that affect global cache transfer latency. The study environment consisted of a multi-node RAC cluster configured with shared ASM storage and private interconnect networking using a dedicated low-latency RDMA-capable interface. Each node executed identical Oracle instance configurations, but CPU frequency scaling, NUMA memory architecture, and I/O throughput characteristics were varied to introduce controlled heterogeneity. This allowed the assessment of cache coherency behaviors when node performance asymmetry is driven by compute-layer variance rather than application-level load alone.

Workload generation was executed using a synthetic transactional workload generator capable of producing both read-intensive and update-intensive query streams. The workload was partitioned such that specific transaction groups exhibited strong data locality, while others randomly distributed table access across all instances. This allowed measurement of cache block ownership residency stability and frequency of block handoff across instances. Session stickiness and service-level connection routing were configured in multiple modes to analyze how connection affinity impacts block access locality under uneven node performance conditions.

Global cache messaging traces were collected using RAC performance diagnostic infrastructure, including GCS/GES wait event sampling, block transfer message counters, and interconnect packet capture. Block state transitions, including CR block requests, current block remastering operations, and data block ping rates, were monitored to quantify coherency synchronization overhead. Timestamp-aligned traces were recorded to correlate individual block transfers with session-level DML/SELECT operations, enabling a direct mapping between application-level data access patterns and Cache Fusion traffic characteristics.

To observe buffer cache residency shifts, the database buffer cache on each instance was sampled at micro-interval granularity to record hot block distribution. Each tracked block was tagged with instance ownership state, modification sequence, and time-to-reuse metrics. This enabled the identification of blocks that frequently transferred ownership due to conflicting access. The buffer cache sampling also supported the classification of blocks into stable-local, transient, and contested categories, providing insight into how data access divergence contributes to interconnect traffic amplification.

Interconnect throughput and latency characteristics were measured using OS-level packet flow tracing and hardware-level port counters. Message queue depths, retransmission rates, and congestion indicators were monitored to identify network saturation events. The study also examined the impact of CPU scheduling perturbations on global cache service threads, focusing on how CPU starvation of GCS/GES processes on slower or overloaded nodes introduced synchronization delays visible at the cluster level.

Node load heterogeneity was introduced using controlled CPU throttling, memory pressure, and background I/O disturbance workloads. These disturbances were applied selectively to individual nodes to observe how cluster behavior changed when one or more nodes exhibited lower effective processing capability. The study recorded not only the direct impact on block transfer latency but also secondary effects such as instance-level lock remastering, dynamic service relocation attempts, and adaptive affinity rebalancing executed by the RAC load distribution logic.

To quantify scalability behavior, the workload was iteratively scaled by increasing session counts and transaction throughput. At each scaling stage, global cache wait distribution, transaction completion latency, and CPU utilization profiles were recorded. The objective was to determine threshold inflection points at which cache coherency overhead overtakes compute throughput, resulting in diminishing

performance returns as cluster load increases unevenly. These measurements supported the construction of performance curves that describe RAC throughput under progressively imbalanced node conditions.

Finally, a controlled failover and load redistribution scenario was executed to analyze how Cache Fusion behaviors evolve during cluster reconfiguration. This included observing the behavior of surviving nodes when a slower node was temporarily isolated or removed. The transition effects on block ownership, lock mastering roles, and recovery traffic patterns were recorded to evaluate the resilience of coherency maintenance mechanisms under abrupt topology shifts.

3. Results and Discussion

The results demonstrate that heterogeneous node performance has a direct and measurable impact on Cache Fusion message volume, block transfer latency, and global serialization delays. Under uniform compute conditions, block residency stability remained high and interconnect traffic followed consistent CR and current block request patterns. However, when CPU or memory throttling was selectively applied to a single node, that node became a *coherency lag point*, increasing block ownership transfer frequency and elevating GCS/GES wait events cluster-wide. This effect occurred even when the impaired node handled only a small fraction of user workload, confirming that the slowest node dictates the upper bound of global coherency performance.

Workload runs under mixed read–write access patterns revealed accelerated hot block pinging when a node exhibited reduced responsiveness. This behavior significantly increased "gc current block busy" and "gccr request" wait times, indicating that consistency traffic amplification is driven by ownership churn rather than volume of data accessed. Once a node slowed sufficiently to delay its global cache service threads, coherency congestion propagated horizontally across the cluster. This behavior is clearly reflected in Table 1, where the average Cache Fusion transfer latency increases from 0.9 ms in a uniform cluster to 4.8 ms under heavy throttling conditions.

Table 1. Cache Fusion Performance Metrics Under Node Load Variation

Scenario	Avg. Global Cache Transfer Latency (ms)	Hot Block Contention Rate (transfers/sec)	GCS/GES Queue Depth (avg)	Cluster Throughput Change (%)
Uniform Node Load	0.9 ms	120	4	Baseline (0%)
Mild CPU Throttle on Node 3	1.7 ms	260	9	−12%
Heavy CPU Throttle on Node 3	4.8 ms	720	21	−37%
Memory Pressure on Node 2	3.9 ms	640	18	−29%
Node 3 Removed (3-node → 2-node)	1.1 ms	150	5	+16% (stabilization rebound)

Interconnect monitoring showed that network bandwidth was not the initial limiting factor; instead, delays emerged at the GCS/GES scheduling layer. Queue depth growth was observed on the impacted node prior to any packet-level congestion, which confirms that coherency stalls originate from CPU scheduling latency rather than raw network throughput saturation. As global cache requests stalled, retransmission logic increased message repetition and intensified packet demands, further degrading latency. The system showed a self-amplifying degradation cycle, driven by the slow node's inability to perform coherency arbitration in real time.

Buffer cache residency sampling provided additional structural insight. Under uniform load, frequently accessed data blocks remained localized and exhibited high reuse rates. Under node heterogeneity, the same blocks became transient, repeatedly transferring ownership across instances rather than stabilizing. This prevented buffer cache locality formation, reducing effective cache hit ratios. As shown in Table 1, the hot block contention rate rises sharply from 120 transfers/sec to over 720 transfers/sec in heavy throttling conditions, directly linking node lag to coherency churn.

The failover test further reinforced this conclusion: removing the impaired node caused an immediate 16% throughput recovery even before workload balancing was complete. This validates a key principle: RAC performance scales with uniformity, not size. A smaller but latency-stable cluster provides higher throughput than a larger but asymmetrically performing one.

4. Conclusion

The results of this study clearly demonstrate that Cache Fusion performance in Oracle RAC is governed by uniformity of node latency characteristics rather than total cluster node count or raw compute capacity. When all instances operated under balanced load, block residency patterns remained stable and interconnect exchange overhead stayed within predictable ranges. However, when even a single node experienced reduced CPU scheduling availability or memory pressure, cluster-wide coherency synchronization degraded disproportionately. The performance deterioration shown in Table 1where average global cache transfer latency increased from 0.9 ms to 4.8 ms and hot block contention surged six-fold under heavy throttlinghighlights that coherency bottlenecks are magnified by the slowest node, not mitigated by faster ones.

The analysis also confirms that coherency stalls originate primarily in GCS/GES coordination threads, not from interconnect bandwidth limitations. Queue depth growth, retransmission bursts, and ownership churn patterns observed in the impaired-node scenarios indicate that global synchronization delay propagates outward even when network throughput remains nominally under-capacity. Buffer cache residency transitions further reinforce the systemic effect: under node heterogeneity, frequently accessed blocks cease to develop stable locality and instead circulate repeatedly between nodes. This results in a feedback loop where latency drives contention, and contention drives further latency, accelerating degradation under heavy mixed workloads.

Finally, the failover scenario demonstrated that removing the impaired node immediately improved cluster throughput by 16%, despite reducing total compute resources. This underscores a foundational operational principle: RAC clusters must be tuned for latency symmetry before scaling for capacity. Practical optimization must therefore prioritize CPU scheduling fairness, NUMA memory accessibility, interconnect packet processing determinism, and proactive detection of node drift. For production RAC environments, ongoing performance validation should include node equivalence audits, coherency traffic baselining, and automated triggers for service relocation or node quarantine when latency divergence exceeds threshold tolerance.

References

1. Ahmed, J., Mathialagan, A. G., & Hasan, N. (2020). Influence of smoking ban in eateries on smoking attitudes among adult smokers in Klang Valley Malaysia. *Malaysian Journal of Public Health Medicine*, *20*(1), 1-8.
2. Yasmin, Farzana, et al. "Response of sweet potato to application of Pgpr and N fertilizer." *Annals of the Romanian Society for Cell Biology* 25.4 (2021): 10799-10812.
3. Fazlul Karim Khan, Md, et al. "Molecular characterization of plasmid-mediated non-O157 verotoxigenic Escherichia coli isolated from infants and children with diarrhea." *Baghdad Science Journal* 17.3 (2020): 19.
4. Haque, A. H. A. S. A. N. U. L., Anwar, N. A. I. L. A., Kabir, S. M. H., Yasmin, F. A. R. Z. A. N. A., Tarofder, A. K., & MHM, N. (2020). Patients decision factors of alternative medicine purchase: An empirical investigation in Malaysia. *International Journal of Pharmaceutical Research*, *12*(3), 614-622.
5. Nazmul, M. H. M., M. A. Rashid, and H. Jamal. "Antifungal activity of Piper betel plants in Malaysia." *Drug Discov* 6.17 (2013): 16-17.
6. Doustjalali, S. R., Gujjar, K. R., Sharma, R., & Shafiei-Sabet, N. (2016). Correlation between body mass index (BMI) and waist to hip ratio (WHR) among undergraduate students. *Pakistan Journal of Nutrition*, *15*(7), 618-624.
7. Hussaini, J., et al. "Recombinant Clone ABA392 Protects laboratory animals from Pasteurella multocida serotype BJ Vet." *Adv* 2 (2012): 114-119.
8. Arzuman, H., Maziz, M. N. H., Elsersi, M. M., Islam, M. N., Kumar, S. S., Jainuri, M. D. B. M., & Khan, S. A. (2017). Preclinical medical students perception about their educational environment based on DREEM at a Private University, Malaysia. *Bangladesh Journal of Medical Science*, *16*(4), 496-504.
9. Navanethan, D. H. A. R. S. H. I. N. I., et al. "Stigma, discrimination, treatment effectiveness and policy: Public views about drug addiction in Malaysia." *Pakistan Journal of Medical and Health Sciences* 15.2 (2021): 514-519.
10. Jamal Hussaini, N. M., Abdullah, M. A., & Ismail, S. (2011). Recombinant Clone ABA392 protects laboratory animals from Pasteurella multocida Serotype B. *African Journal of Microbiology Research*, *5*(18), 2596-2599.
11. Nazmul, M. H. M., et al. "General knowledge and misconceptions about HIV/AIDS among the university students in Malaysia." *Indian Journal of Public Health Research & Development* 9.10 (2018): 435-440.

12. Hussaini, J., Nazmul, M. H. M., Masyitah, N., Abdullah, M. A., & Ismail, S. (2013). Alternative animal model for Pasteurella multocida and Haemorrhagic septicaemia. *Biomedical Research*, *24*(2), 263-266.
13. Iqbal, Mohsena, et al. "The study of the perception of diabetes mellitus among the people of Petaling Jaya in Malaysia." *International Journal of Health Sciences* I (2022): 1263-1273.
14. MKK, F., MA, R., Rashid, S. S., & MHM, N. (2019). Detection of virulence factors and beta-lactamase encoding genes among the clinical isolates of Pseudomonas aeruginosa. *arXiv preprint arXiv:1902.02014*.
15. DOUSTJALALI, SAEID REZA, et al. "Correlation between body mass index (BMI) & waist to hip ratio (WHR) among primary school students." *International Journal of Pharmaceutical Research* 12.3 (2020).
16. Nazmul, M. H. M., Fazlul, M. K. K., Rashid, S. S., Doustjalali, S. R., Yasmin, F., Al-Jashamy, K., ... & Sabet, N. S. (2017). ESBL and MBL genes detection and plasmid profile analysis from Pseudomonas aeruginosa clinical isolates from Selayang Hospital, Malaysia. *PAKISTAN JOURNAL OF MEDICAL & HEALTH SCIENCES*, *11*(3), 815-818.
17. Selvaganapathi, G., et al. "Knowledge and practice on tuberculosis among prison workers from Seremban Prison." *Occupational Diseases and Environmental Medicine* 7.4 (2019): 176-186.
18. Nazmul, M. H. M., Salmah, I., Jamal, H., & Ansary, A. (2007). Detection and molecular characterization of verotoxin gene in non-O157 diarrheagenic Escherichia coli isolated from Miri hospital, Sarawak, Malaysia. *Biomedical Research*, *18*(1), 39-43.
19. Khan, Md Fazlul K., et al. "Detection of ESBL and MBL in Acinetobacter spp. and Their Plasmid Profile Analysis." *Jordan Journal of Biological Sciences* 12.3 (2019).
20. Foysal, Md Javed, et al. "Identification and assay of putative virulence properties of Escherichia coli gyrase subunit A and B among hospitalized UTI patients in Bangladesh." *Inov Pharm Pharmacother* 1.1 (2013): 54-59.
21. Hussaini, Jamal, Nurul Asyikin Othman, and Mahmood Ameen Abdulla. "Antiulcer and antibacterial evaluations of Illicium verum ethanolic fruits extract (IVEFE)." *Medical science* 2.8 (2013).
22. Nazmul, M., M. Fazlul, and M. Rashid. "Plasmid profile analysis of non-O157 diarrheagenic Escherichia coli in Malaysia." *Indian Journal of Science* 1.2 (2012): 130-132.
23. Vijayakumar, K., Mohammad Nazmul Hasan Maziz, and Mathiyazhagan Narayanan. "Classification of Benign/Malignant Digital Mammogram Images using Deep Learning Scheme." *hospital* 4 (2025): 5.