NoSQL databases on Red Hat Enterprise Linux

Highlights

Deploy Intel processor and memory technologies running Red Hat Enterprise Linux for predictable, security-hardened, and tested infrastructure.

Employ memory tiering to achieve database scalability and cost-efficient infrastructure.

Reduce database restart times and eliminate time-consuming database rescans by hosting the primary database index on persistent tiered memory.

Memory tiering supports cost-effective real-time computing at petabyte scale.

Scaling and protecting real-time applications

Non-relational (NoSQL) databases offer reliable high performance against large transaction volumes, creating exciting new application opportunities. Digital payment systems, fraud detection and prevention systems, risk analysis applications, recommendation engines, data management platforms, user profile stores, and more now demand near instantaneous response at tremendous scale. At the same time, designing and optimizing infrastructure can be challenging, with a low tolerance for latency or downtime. Cost-effective infrastructure scaling is also vital for sustainable operations.

In the face of rapidly evolving technology, Intel and Red Hat are collaborating to help organizations meet service level agreements (SLAs), reduce downtime, and improve total cost of ownership (TCO).

- **Intel.** Processor, memory, and storage technologies are constantly evolving. Given the importance of memory and storage to real-time applications, Intel has long focused on the growing gap between exponential business data growth and slower dynamic random access memory (DRAM) scaling. As a result, Intel processors, systems, and memory technologies accommodate memory tiering approaches that promote cost-effective scalability for large NoSQL databases while reducing system acquisition costs.
- **Red Hat.** Organizations need a security-focused and stable foundation from the public cloud to the edge that offers scalability and consistent performance. Certified on hundreds of clouds and with thousands of hardware and software vendors, Red Hat® Enterprise Linux® provides a consistent administrative and management experience, automated security and compliance, and comprehensive performance monitoring, tracing, and analysis.

Memory tiering with NoSQL databases

Some NoSQL databases can hold the database index in high-speed DRAM memory and the associated data in high-performance solid-state drives (SSDs). Until recently, the amount of DRAM in the system constrained how much data a given server node could host, limiting how much data the overall cluster could support. For example, a database could use Intel® Optane™ Persistent Memory (PMem) in App Direct mode to place the index in persistent memory for dramatic scalability and cost savings (Figure 1).

The path to CXL

Intel® Optane™ PMem provides a compelling memory tiering solution today for applications like NoSQL databases, along with a path toward open standards like Compute Express Link (CXL). As an industry-supported, cache-coherent interconnect for processors, memory expansion, and accelerators, CXL promises to facilitate breakthrough performance for emerging usage models while supporting an open ecosystem for datacenter accelerators and other high-speed enhancements. As a member of the CXL Board of Directors, Intel is taking an active role in these exciting new developments.

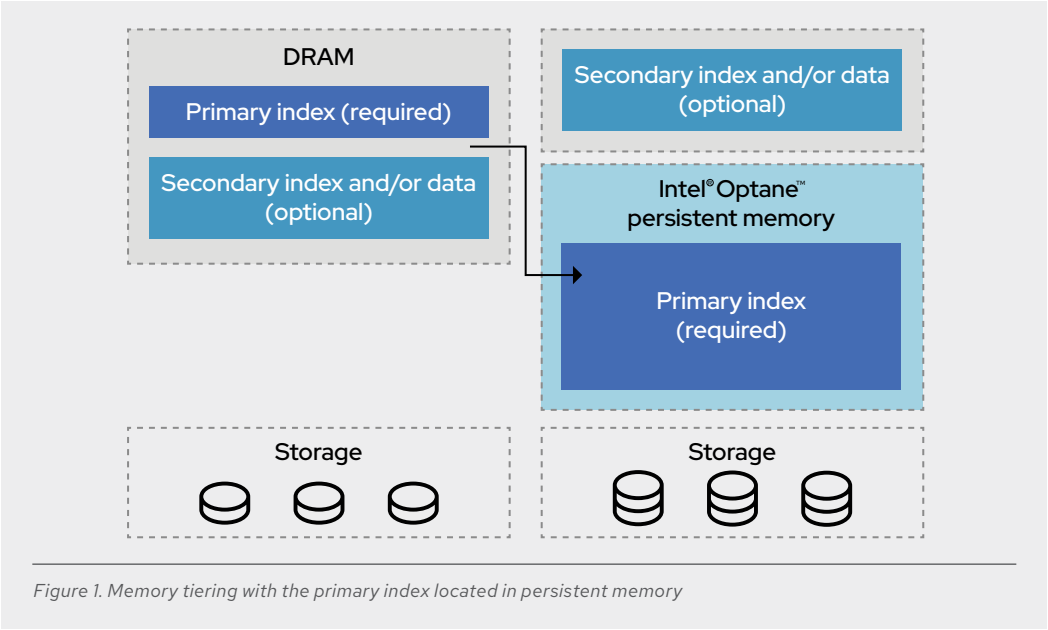

Figure 1. Memory tiering with the primary index located in persistent memory

Using memory tiering to host the primary index of a NoSQL database provides several distinct benefits:

- **Increased scale.** When using only 64GB DRAM dual in-line memory modules (DIMMs) for system memory, a single 2-socket server node is limited to about 1.5 terabytes (TB). In contrast, using memory tiering combined with a small amount of DRAM, a 2-socket server node can readily be configured with up to 6TB of memory. With memory tiering, each node of a database cluster can potentially support up to 4 times as much data as a DRAM-only solution—scaling the cluster's overall capacity.
- **Comparable performance.** Most NoSQL databases are tuned to use the memory resources on each server node efficiently. For example, data placement is typically optimized to take advantage of slightly faster DRAM performance and the higher capacity obtainable with memory tiering. The result is performance that can approach that of an all-DRAM solution at a much lower cost.
- **High uptime.** With a traditional DRAM-only database solution, the database index is lost whenever a node crashes or powers down for maintenance or upgrades. Data must then be scanned during startup to rebuild the indexes—a process that can require several hours. With the primary index placed on a persistent memory tier, the entire database index is retained across crashes or reboots, and the database can be restarted in seconds.

Performance testing on Red Hat Enterprise Linux

Intel conducted extensive testing to evaluate the effectiveness of placing a NoSQL database primary index in Intel Optane PMem 200 Series compared to locating it in DRAM. Tests exercised a 50TB database of 50 billion records (25 billion unique records with 2 times replication). A bare-metal cluster with 3 database servers was accessed by 4 client nodes, as shown in Figure 2.

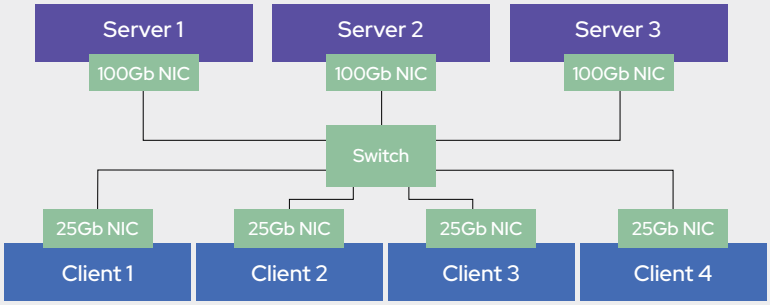

Figure 2. Intel test configuration for NoSQL database testing

Transaction throughput

Testing confirmed that the memory-tiered approach offered nearly the performance of a DRAM-only configuration.[1] Expressly, the memory-tiered configuration provided up to 96.83% of the performance of servers configured to store the primary index in DRAM (Figure 3).

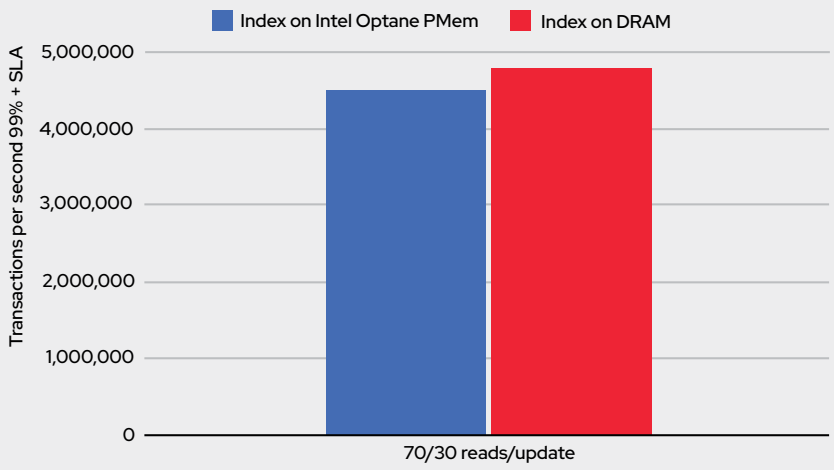

Figure 3. 70/30 read/update test results placing the primary index in Intel Optane PMem vs. DRAM

1 *BASELINE: DDR4 DRAM: Test by Intel as of 02/07/23. 3-nodes, 2x Intel(R) Xeon(R) Platinum 8358 CPU @ 2.60GHz, 32 cores, HT On, Turbo On, Total Memory 2048GB (32x64GB DDR4 2933 MT/s [2933 MT/s]), BIOS 1.2, microcode 0xd000375, 4x Ethernet Controller E810-C for QSFP, 2x Ethernet Controller 10-Gigabit X540-AT2, 1x 447.1G INTEL SSDSC2BB48, 6x 7T INTEL SSDPF2KX076TZ, 1x 745.2G INTEL SSDPF21Q800GB, Red Hat Enterprise Linux 8.7 (Ootpa), 4.18.0-425.3.1.el8.x86_64, Aerospike Enterprise Edition 6.1.0.1, Aerospike Benchmark Utility Version 1.5.3, C Client Version 6.0.0, score=4,839,900 transactions per second on 70/30 read/update workload with 99%+ transactions completing within 1 millisecond.*
MEMORY TIERING: Intel Optane persistent memory 200 series: Test by Intel as of 02/01/23. 3-nodes, 2x Intel(R) Xeon(R) Platinum 8358 CPU @ 2.60GHz, 32 cores, HT On, Turbo On, Total Memory 2048GB (16x128GB Logical non-volatile device 3200 MT/s [2933 MT/s]); 1024GB (16x64GB DDR4 2933 MT/s [2933 MT/s]), BIOS 1.2, microcode 0xd000375, 2x Ethernet Controller 10-Gigabit X540-AT2, 4x Ethernet Controller E810-C for QSFP, 1x 447.1G INTEL SSDSC2BB48, 6x 7T INTEL SSDPF2KX076TZ, 1x 745.2G INTEL SSDPF21Q800GB, Red Hat Enterprise Linux 8.7 (Ootpa), 4.18.0-425.3.1.el8.x86_64, Aerospike Enterprise Edition 6.1.0.1, Aerospike Benchmark Utility Version 1.5.3, C Client Version 6.0.0, score=4,686,633 transactions per second on 70/30 read/update workload with 99%+ transactions completing within 1 millisecond.

Restart time

Intel testing found that placing the primary index in Intel Optane PMem also provided significant advantages for database restart time. A system with the index in Intel Optane PMem restarted up to 67 times faster than a server with the primary index hosted in DRAM. The memory-tiered server avoided having to scan the database and rebuild the index—a process requiring hours for a large database.

System cost of acquisition

Table 1 shows the significant system cost reduction available with Intel memory tiering. Cost savings can approach 50% even across a small 3-node cluster.[2]

Table 1. Cost of acquisition for a 3-node cluster

	CPU cost	Memory subsystem cost	Per-system cost	3-node cluster cost
DRAM-based system	$9,214	$71,440	$80,654	$241,962
Memory-tiered system	$9,214	$31,853	$41,067	$123,201

Learn more

By deploying NoSQL databases on Intel hardware and Red Hat Enterprise Linux, organizations can benefit from a strong partnership and technologies tested to work together. With memory tiering, databases can scale by tens of billions of records with a minimal performance tradeoff over a DRAM-only system while lowering system acquisition cost and database restart times. For more information on this solution, please see vendor documents concerning Intel memory tiering and Red Hat Enterprise Linux.

2 *System comparison based on two-socket servers powered by 32-core Intel Xeon Platinum 8358 CPUs. DRAM-based system configured with 1024GB DDR4 DRAM memory per socket (16x64GB) for a total of 2048GB. Memory-tiered system configured with a total of 1152GB per socket (total 2304) consisting of Intel Optane PMem 200 Series (8x128GB) and DRAM (8x16GB) as calculated by Intel internal cost estimation tool. Performance results are based on testing as of February 2023, and may not reflect all publicly available security updates. For more complete information about performance and benchmark results, visit www.intel.com/benchmarks.*

About Red Hat

Red Hat helps customers standardize across environments, develop cloud-native applications, and integrate, automate, secure, and manage complex environments with award-winning support, training, and consulting services.

facebook.com/redhatinc
@RedHat
linkedin.com/company/red-hat

North America
1 888 REDHAT1
www.redhat.com

Europe, Middle East, and Africa
00800 7334 2835
europe@redhat.com

Asia Pacific
+65 6490 4200
apac@redhat.com

Latin America
+54 11 4329 7300
info-latam@redhat.com

redhat.com
#339190_0523