

MemStac™: A High Performance, Low Cost Memcached Solution



Up to 80x increase in transactions per second

**Up to 30x lower query load on backend database
leading to improved system responsiveness**

Up to 30x lower \$/GB/Hour

Executive Summary

OmniTier's MemStac™ is a high performance, Memcached-compatible key-value solution that provides a much higher cache capacity at significantly lower cost than Memcached. It achieves higher capacity by using tiered memory based on DRAM and NVMe SSD technologies. MemStac™ on AWS i3 instances can provide up to 30x higher capacity at 30x lower \$/GB/Hour compared to AWS ElastiCache instances. MemStac™ cache cluster delivers over 80x improvement in application performance through higher cache hit rates while greatly reducing the load on the backend database.

Why Memcached?

Business performance has been shown to correlate with application performance, as slow application performance

can cause customer churn [1]. Because relational databases are relatively slow, in-memory caches are frequently used to accelerate database performance. Memcached is a widely used in-memory (100% DRAM) caching solution for database caching. Its large-scale users include companies like Facebook, LinkedIn and Yahoo. Memcached libraries are available in most programming languages, thus enabling application developers to easily integrate the caching solution into their application stack.

MemStac™: High performance Memcached

MemStac™ is based on a tiered-memory architecture. By using advanced data classification and latency management algorithms, it optimally distributes data objects of varying sizes and frequency across a pool of DRAM and NVMe SSD resources. The result is a much greater cache capacity that scales at the price point of SSDs and mirrors the performance of DRAM-only solutions like Memcached.

Memcached running on a single high-performance server based on dual socket Intel Xeon® E5-2699v4 (44 vCPU) with 10GbE network can deliver over 4M requests per sec (RPS) with an average latency of less than 0.5ms. MemStac™ on the same server hardware with a tiered-memory architecture (DRAM and NVMe SSD) matches Memcached's throughput and full latency distributions. Detailed

performance analysis of MemStac™ under a variety of test workloads is available in [2]. This demonstrates that the DRAM-only solution can be replaced with MemStac™ on cheaper, higher-capacity SSDs with equivalent performance.

MemStac™ on AWS

MemStac™ can be launched on i3.4xlarge and i3.8xlarge instances to obtain Memcached-compatible capacities of 3800GB and 7600GB per instance, respectively. Any standard Memcached client can then be used to interact with the cache. No application level or client library level changes are needed to interface with MemStac™. MemStac™ on i3 instances can enable up to 30x higher cache capacity per instance compared to memory optimized AWS r4 instances (Figure 1).

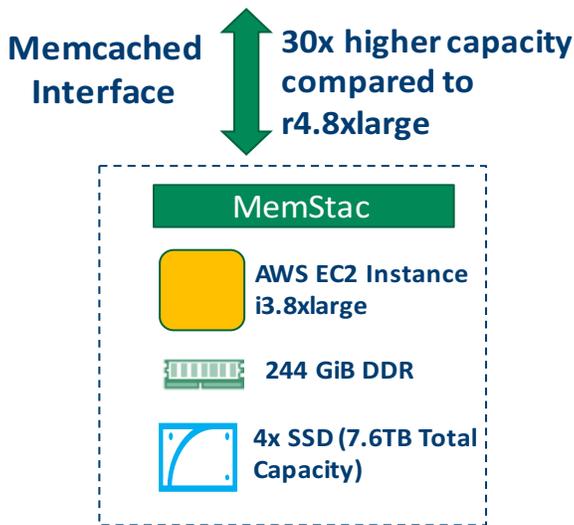


Figure 1: MemStac™ on AWS i3 instances

Figure 2 shows a performance comparison of MemStac™ and Memcached for 100% GET, 1KB key-value size, Zipfian key frequency distribution on AWS i3.8xlarge instances. MemStac™ delivers performance comparable to Memcached while providing 30x increase in cache capacity per instance. Increased cache hit rates due to higher cache capacity per instance can noticeably improve the application performance further. Apart from throughput

and average latency, MemStac’s intelligent data classification, tiering and advanced SSD management delivers full latency distributions similar to Memcached (100% DRAM) as shown in Figure 3.

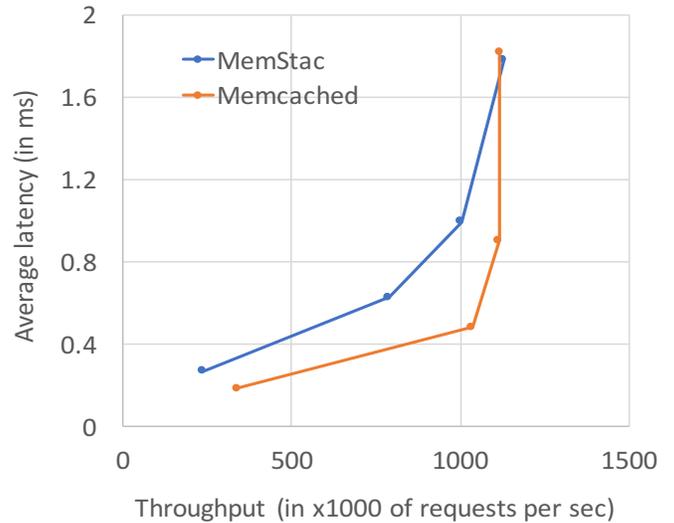


Figure 2: Throughput vs average latency for 100% GET, 1KB KV size, Zipfian key frequency distribution on AWS EC2 i3.8xlarge for 64 TCP connections; queue depth of 1,8,16 and 32.

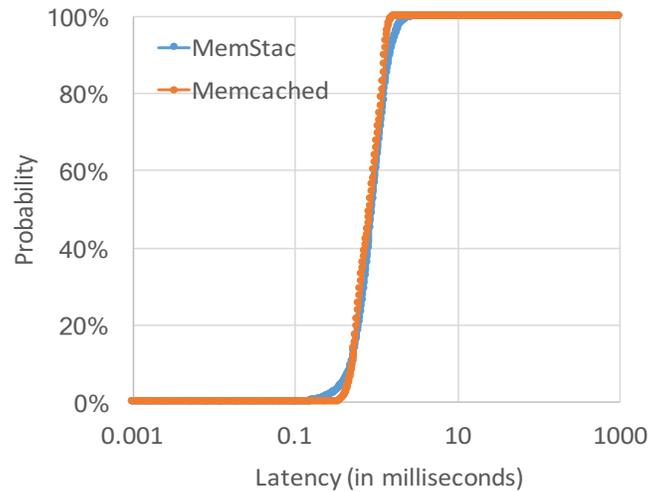


Figure 3: Latency distribution comparison for 100% GET, 1KB KV size, Zipfian key frequency distribution on AWS EC2 i3.8xlarge for 64 TCP connections and queue depth 16.

Application performance with MemStac™

Modern databases are several orders of magnitude slower than either MemStac™ or Memcached – tens of thousands of requests per second compared to millions. Since overall system performance is determined not only by the speed of the cache but the percentage of requests (hit rate) serviced by the cache, insufficient cache capacity can cause the database to be the primary performance bottleneck. Due to its inherently smaller cache size, Memcached may only be able to service a small fraction of the system workload. This limitation causes most of the requests to be directed to the database, leading to poor performance of overall system. MemStac™, on the other hand, with its significantly larger cache size, can economically cache a majority of the application’s requests and reduce the load on the backend database, as shown in the Figure 4, resulting in a system-level performance improvement up to 80x. Figures 5 and 6 summarize the performance improvements with MemStac™ relative to capacity constrained Memcached (with capacity at 10% of database size) for industry standard Yahoo Cloud Serving Benchmark workloads (YCSB) of 95% GET (workload-B) and 50% GET (workload-A), when AWS RDS/MySQL is used as the backend database.

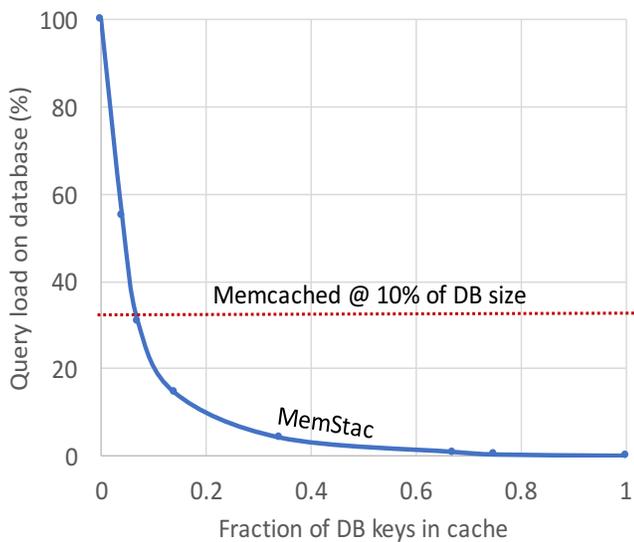


Figure 4: Database query load vs cache size relative to database

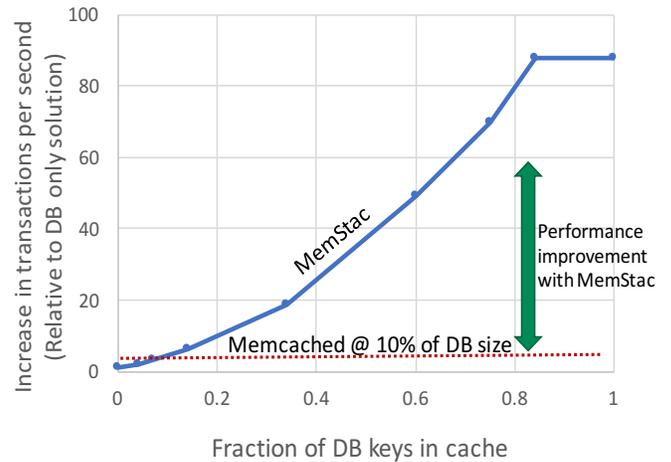


Figure 5: Increase in transactions per second relative to database only solution for YCSB workload-A

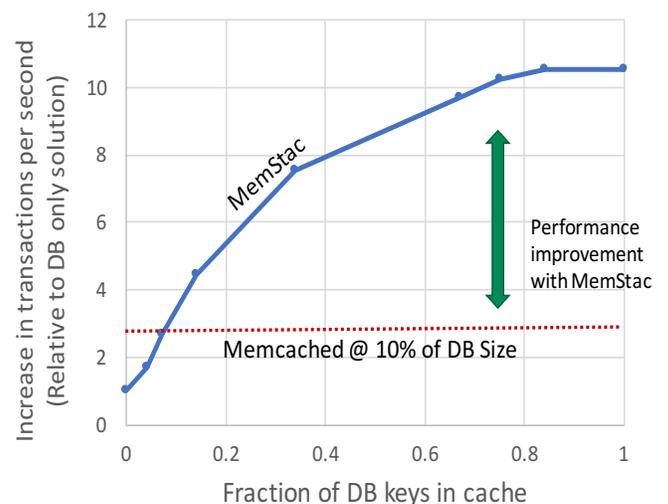


Figure 6: Increase in transactions per second relative to database only solution for YCSB workload-B

Cost savings with MemStac™

By using SSD technology, MemStac™ not only enables higher application performance through higher cache hit rates, but also brings a significant reduction in hourly expenses. Table 1 summarizes MemStac’s \$/100 GB/Hour cost. It can provide over 30x lower cost relative to other

AWS offerings. Lower \$/GB/Hour enables the use of caching tier for large datasets for which in-memory caching was previously impractical.

Memcached Solution	\$/100GB/Hour
ElastiCache (cache.r4.8xlarge)	\$1.79
Memcachier	\$4.40
Redislabs - Memcached Cloud	\$12.20
MemStac on i3.8xlarge	\$0.06

Table 1: MemStac™ economics on AWS

MemStac™ is a distributed caching solution compatible with Memcached. Multiple instances can be grouped as a cluster to achieve targeted capacity or transactions per second. Figure 7 summarizes the cost savings relative to Memcached (AWS ElastiCache) for a deployment that requires 1M requests per second and varying capacities. For very low cache cluster capacities, MemStac’s annual costs are similar to those of Memcached. But as cluster capacities grow, MemStac™ delivers 8x to 17x reduction in estimated costs. On the other hand, for a deployment that requires 10TB cache capacity with varying performance requirements, MemStac™ can deliver 3x to 14x reduction in estimated costs, as shown in Figure 8. Also, number of instances in a MemStac™ cluster can be scaled to handle time varying application needs, minimizing fixed upfront provisioning costs.

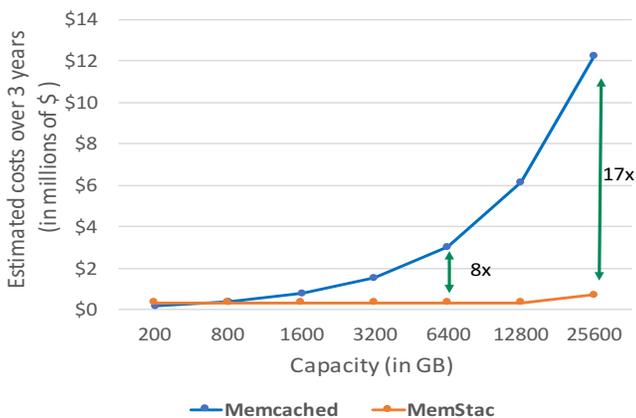


Figure 7: Estimated cluster level costs to achieve 1M transactions and varying cluster capacities from 200GB to 25.6GB for YCSB workloads

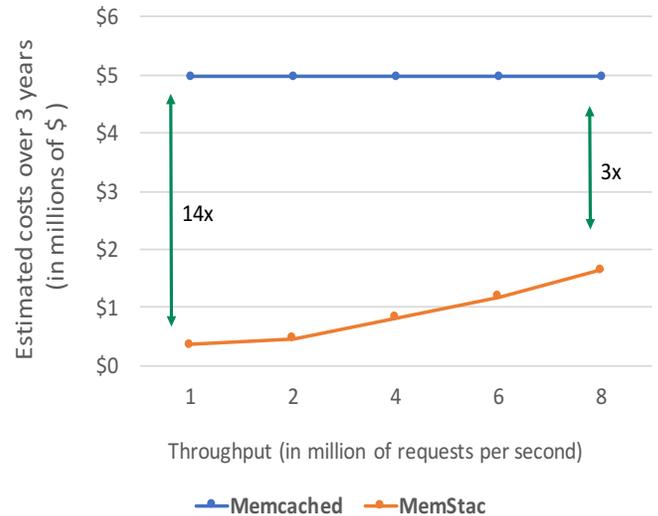


Figure 8: Estimated cluster level costs to achieve 10TB cluster capacity and transactions per second, varying from 1M to 8M for YCSB workloads.

MemStac™ application use cases

Database caching

When a webserver receives a request, it often generates multiple database queries to create the page for the visitor. Memcached is a widely used high-performance in-memory caching solution for database caching. Since MemStac™ is 100% compatible with Memcached, it can be used as a drop-in replacement for applications that are using Memcached today. By providing larger cache capacity and hit rates, MemStac™ can deliver improved response times when it is used as a caching tier for applications like eCommerce.

HTTP caching

In modern web architecture, content proxy servers act as a gateway between the user and origin/source servers to cache documents such as files, images, and html pages. This improves user-experience due to lowered latency, and reduces the resources needed at the source. MemStac™ can be easily integrated with web servers like Nginx and Apache HTTPD by using their respective Memcached modules.

Content management system caching

Content management systems like WordPress, Drupal, Joomla use in-memory caching layer to cache database queries and pre-rendered pages to improve user experience. Most of these CMS solutions already provide plugins to interface with Memcached, so users can readily take advantage of MemStac™.

Shared memory cluster for microservices

To handle cloud workload demands, developers are adopting microservice architectures to partition and modularize their applications. Complex applications can have hundreds of microservices that are stateless but linked through in-memory caches and persistent databases. Given its performance and capacity advantages, MemStac™ can be used as a distributed shared memory for microservices.

Scalable temporal cache

A MemStac™ cluster can be used as a large temporal cache that enables processing of very large datasets in gaming, medicine, automotive and other large-scale applications. The economics of MemStac™ enable application developers to cache datasets that were previously impractical, thereby enabling new features and capabilities.

Conclusion

MemStac™ enables dramatic performance improvements and cost savings when used as a replacement to Memcached. MemStac™ on AWS i3 instances can provide up to 30x higher capacity at 30x lower \$/GB/Hour when compared to AWS ElastiCache instances. MemStac™ cache cluster delivers over 80x improvement in application performance through higher cache hit rates while greatly reducing the load on the backend database. This enables companies to lower infrastructure cost while providing greater scalability for their deployments.

About OmniTier

OmniTier Inc., founded in 2015, develops and supports integrated software solutions for memory-centric infrastructure applications, including high performance object caching, scientific analysis for machine learning, AI, and genomics. Its leadership team has a track record of delivering many industry firsts in data storage and access across different media types. The company has offices in Milpitas, California, and Rochester, Minnesota.

References

¹ *How One Second Could Cost Amazon \$1.6 Billion In Sales*
<https://www.fastcompany.com/1825005/how-one-second-could-cost-amazon-16-billion-sales>

² *MemStac™ performance with Micron SSDs*
<https://omnitier.com/wp-content/uploads/2017/04/MemStac-Performance-on-Micron-SSDs.pdf>