



a Western Digital brand

TABLE OF CONTENTS

2 WHAT IS IN-MEMORY COMPUTING (IMC)

- Benefits of IMC
- Concerns with In-Memory Processing
- Advanced In-Memory Computing using Supermicro MemX¹

3 MEMX ARCHITECTURE

- MemX Functionality and Operation
- At Boot Time
- At Runtime
- MemX Benefits
- System Memory (GB)
- Internal Storage (GB)

5 MEMX APPLIANCE SPECIFICATION

- MemX PCIe NVMe Drives

5 PERFORMANCE TESTING: MEMCACHED WEB-CACHING

- Test Configuration

6 PERFORMANCE TEST RESULTS

- Bandwidth (Transaction per Second)
- Latency

7 PERFORMANCE TESTING: MYSQL

8 CONCLUSION

WHITE PAPER

ADVANCED IN-MEMORY COMPUTING USING SUPERMICRO MEMX SOLUTION

EXECUTIVE SUMMARY

With the rise of real time and on-demand applications, massive data is generated which needs better and faster processing that derives meaningful patterns from the data generated. In Big Data generation, it is said that bigger is always better but that holds true only if data can be processed quickly to be utilized by key stake holders. Failure to achieve desired performance leads to revenue losses and highlights the need of in-memory computing.

July 2017



WHAT IS IN-MEMORY COMPUTING (IMC)

With the explosion of data gathering we are seeing today, organizations need to find new ways of collecting, sorting, and analyzing this data in order to make real-time business decisions. Having a large amount of data at your disposal is only one part of the equation. Being able to use the data is just as important. Enterprise storage systems can no longer keep up with the demand of real-time analytics, so Big Data System Administrators have turned to large scale-up systems populated with massive amounts of expensive DRAM, or set up distributed in-memory grids consisting of many servers.

Benefits of IMC

By turning to in-memory computing, businesses can quickly discover patterns, analyze massive data volumes, and perform business intelligence operations in real time. In-Memory Computing is used in diverse industries such as financial services, telecommunications, healthcare and life sciences, government, energy, transportation, among many others.

Some of the benefits of in-memory computing include

- **Mitigating Errors and Avoiding Future Failures**
Real-time data analytics allows businesses to quickly predict, react and/or avoid future operating problems. Doing so protects business continuity and keeps customers satisfied.
- **Advanced Business Intelligence**
Staying ahead of the competition is key. Using in-memory computing for BI applications allows businesses to become dynamic, and transform their strategies on the fly.
- **Fraud Detection**
Security is paramount. Cyber-attacks lead to financial losses and business down-time. Real-time security systems expose these threats and allow them to be mitigated instantly.

Concerns with In-Memory Processing

There are some physical limitations to the size of data to be placed into memory:

1. Capacity limitations
 - a. Memory module capacity
 - b. Server total memory capacity
2. Cost limitations
 - a. At the time of writing, commodity 64GB DIMMs are less than \$1000 per unit while 128GB DIMMs are uncommon and costly, e.g. more than \$4,000 per unit.
 - b. For distributed applications, it is possible to increase total memory by adding server nodes, but the additional expenses of systems, processors, and networking become cost prohibitive.

Advanced In-Memory Computing using Supermicro MemX¹

By replacing large monolithic servers with Supermicro's high-performance MemX solution (using NVMe Solid State Drives from Western Digital Corporation and Software-defined Memory), system administrators can now extend DRAM to much larger system memory. Further, MemX uses direct attached NVMe storage to replace slow external SAN or NAS storage. This greatly reduces initial acquisition costs by up to 90%, while reducing operational costs and providing added scale-out flexibility over competing DRAM In-Memory Computing solutions.

¹ MemX Software-defined Memory technology licensed from ScaleMP

MEMX ARCHITECTURE

Transition from DRAM to Hybrid DRAM-NVM System Memory *Breaking the memory-storage barrier*

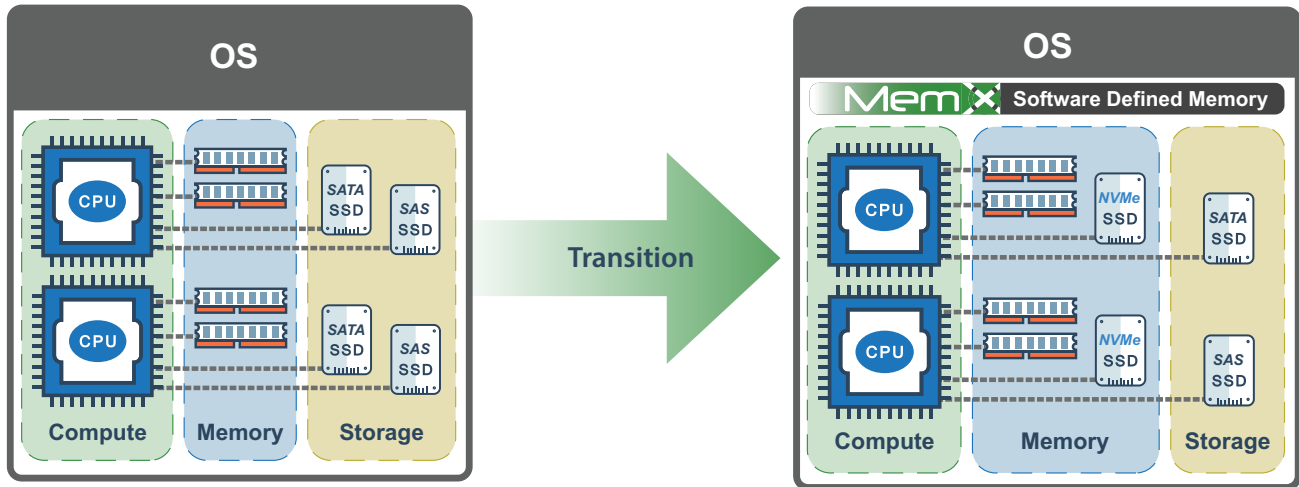
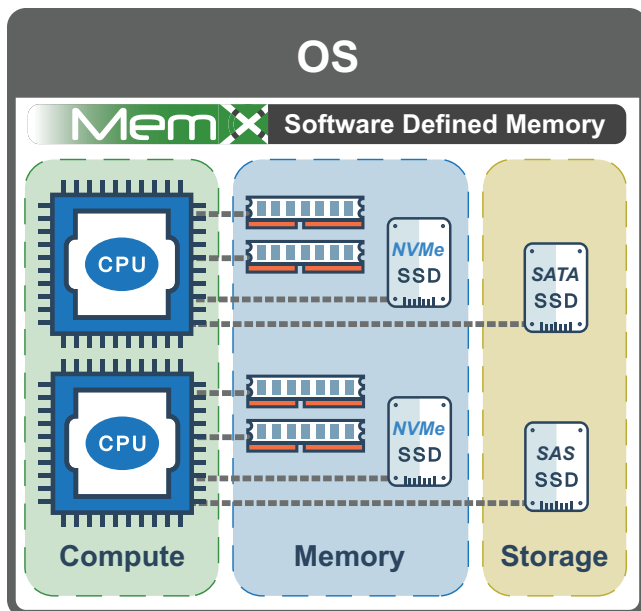


Figure 1: before and after use of Software-defined Memory (SDM) over Storage

MemX Functionality and Operation



At Boot Time

- BIOS table list the increased memory map by the added capacity of the NVMe SSDs
- The NVMe SSDs are hidden from PCI Bus Scanning

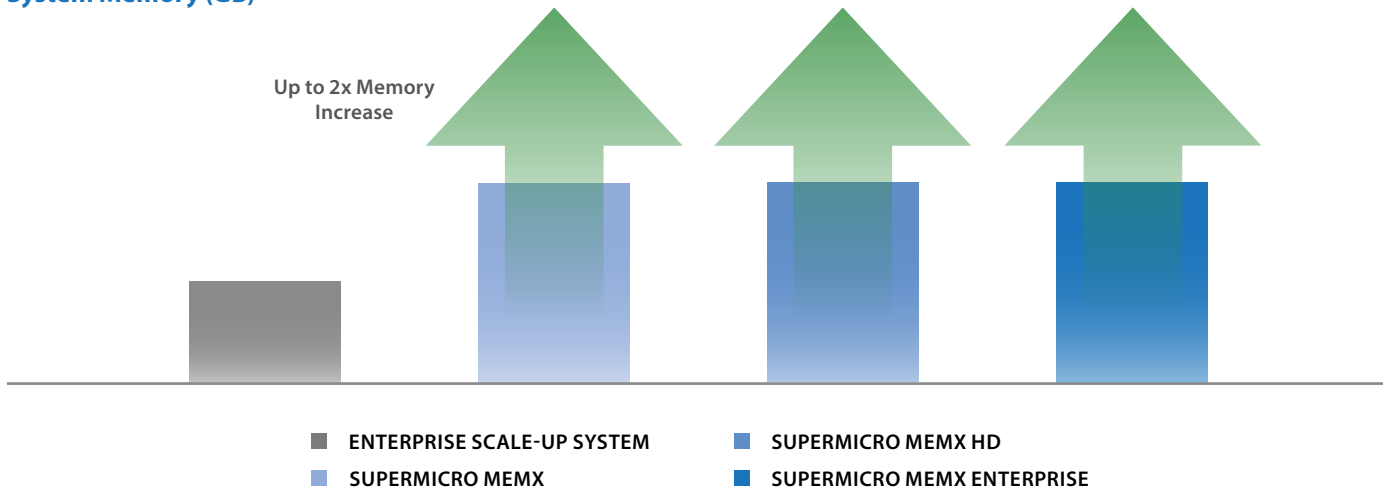
At Runtime

- When OS or application access anything that is present in DRAM, the CPU handles that directly. The DRAM is used as a cache.
- When access occurs to data that is not present in DRAM (cache-miss), the MemX software-defined memory kicks in and brings the required data into DRAM.
- The MemX software-defined memory uses machine learning to continuously prefetch from the NVMe SSDs to significantly reduce misses.

MemX Benefits

With Supermicro's MemX Solution, end users can increase their system memory up to 2X and their internal storage up to 10X, compared to an Enterprise 8-socket system. This allows for greater server consolidation, making Supermicro MemX the most flexible, scalable IMC solution on the market.

System Memory (GB)



Internal Storage (GB)

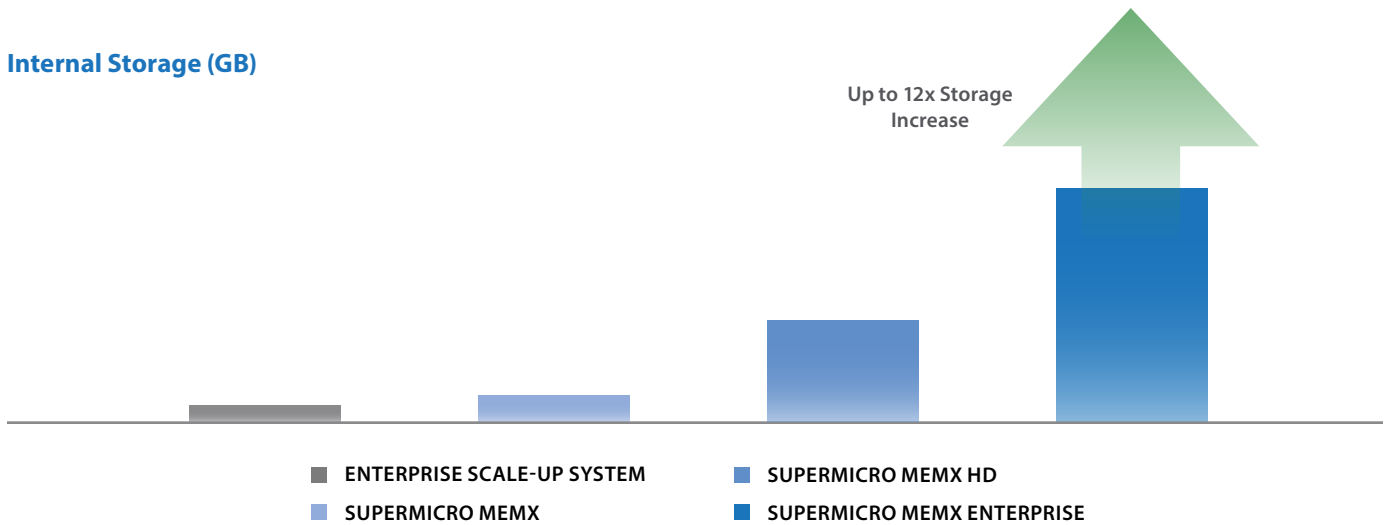


Figure 2: Scale advantages for Supermicro MemX over Enterprise scale-up system

MEMX APPLIANCE SPECIFICATION



	MEMX CLOUD (PER NODE) SYS-2028BT-HNR+	MEMX HD SYS-1028U-TN10RT+	MEMX SYS-2028U-TN24R4T+	MEMX ENTERPRISE SSG-2028R-NR48N
Processors	E5-2683 v4	E5-2683 v4	E5-2683 v4	E5-2699 v4
Total Cores	32 cores @ 2.1GHz	32 cores @ 2.1GHz	32 cores @ 2.1GHz	44 cores @ 2.2GHz
Total System Memory (GB)*	3,328 / 11,776	3,328 / 11,776	3,328 / 11,776	6,656 / 11,776
Total SSD Storage Size (GB)	1,920	9,600 / 38,400	60,800 / 145,920	137,600 / 330,240
SSD Storage Model	2 x 10GBase-T ports (SIOM)	2 x 10GBase-T ports	2 x 10GBase-T ports	2 x 10GBase-T ports

MemX PCI-E NVMe Drives

The Supermicro MemX utilizes two types of PCI-E NVMe flash drives for memory extension and for storage capacity.

- For memory extension, HGST Ultrastar SN200® NVMe drives of capacities 800GB, 1.6TB, and 3.2TB are utilized with memory sizes of 512GB, 1024GB, and 2048GB respectively.
- For storage in MemX Cloud offering, 1.92TB SanDisk™ Skyhawk™ NVMe SSDs are used. For the other MemX offerings, HGST Ultrastar SN200® drives are used, with up to 7.68TB per SSD.

PERFORMANCE TESTING: MEMCACHED WEB-CACHING

Test Configuration

Memcached is a general-purpose distributed memory caching system. It is often used to speed up dynamic database-driven websites by caching data and objects in RAM to reduce the number of times the persistent or back-end data source (such as a database) must be read. Memcached allows for a cluster to be used as cache, meaning if one wants a larger cache, they can add servers to the memcached cluster. In this test, we measure the number of 1K transactions that a single physical memcached server can process per second, by loading it from a much more powerful load-generation system to drive it to its limits. We then compare the results obtained by a regular server with 768GB of DRAM serving as the memcached server, to the results obtained by a MemX server reconfigured for only 768GB capacity, of which only 1/8 is DRAM (96GB) and 7/8 come from two NVMe SSDs.

TEST SYSTEM	
(MEMX HD, SCALED DOWN IN ORDER TO BE COMPARABLE IN SIZE TO STANDARD DUAL-SOCKET SYSTEM)	
CPU	2 x E5-2699v4
MEMORY	96GB
MEMX MEMORY EXTENSION	768GB
INTERNAL STORAGE:	3.2TB
MEMCACHED SYSTEM	2x E5-2699 v4 @ 2.20GHz - 88 CPUs
MEMASLAP SYSTEM (LOAD-GENERATOR)	4x E7-8890 v3 @ 2.50GHz - 144 CPUs
NETWORK TOPOLOGY	10GbE dedicated back-to-back link
LOAD CHARACTERISTICS	key:value=128B:1024B, get:set=90%:10%

PERFORMANCE TEST RESULTS

Bandwidth (Transaction per Second)

Initial performance tests for an in-memory caching software – memcached. MemX HD, reconfigured to only have 96GB of DRAM, expanded to a total of 768GB of memory using MemX architecture (NVMe SSDs and Software-Defined Memory), was able to deliver 98% of the performance (Transaction per Second) of a standard server configuration with 768GB of DRAM.




Figure 3: MemX Software-defined Memory delivers near-DRAM performance in rate of transactions

Latency

The charts below illustrate how the MemX solution delivers low, predictable latency – close to that of the DRAM-only server.

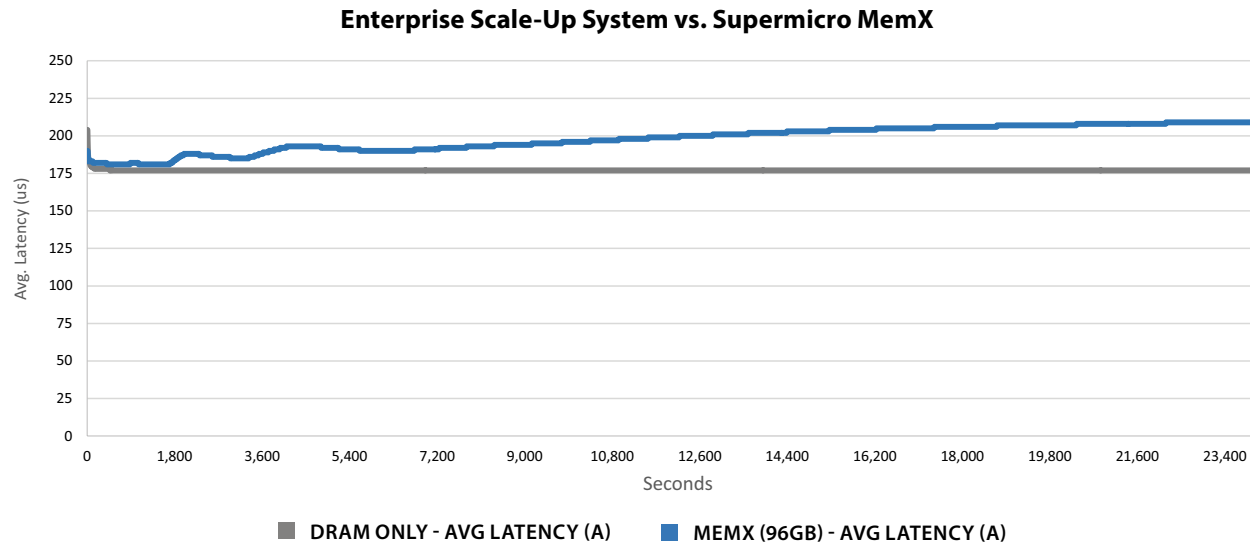


Figure 4: MemX Software-defined Memory delivers consistent low-latency performance

PERFORMANCE TESTING: MYSQL

In a separate performance characterization test using an Online Transaction Processing (OLTP) workload (TPC-C) with MySQL for a multi-tenant use case, the MemX HD solution using 96GB of DRAM to deliver 768GB of system memory, was able to deliver 80% and 74% of the performance (TPC-C) for two industry standard DRAM servers with an Intel® Xeon® processor E5-2637 v4 3.5GHz with 16 CPU threads and an E5-2699 v4 2.2GHz processor with 88 CPU threads respectively.




Figure 5: MemX delivers 74%-80% of DRAM performance for in-memory TPC-C using MySQL

CONCLUSION

These performance characterizations and budgetary estimates point to a very compelling proposition for customers looking into large memory configurations for in-memory processing requirements.

Three benefits of this solution are as follow:

First, with the memory extension solution, customers will be able to derive the full benefits of in-memory processing by increasing the system memory available.

Second, as the in-memory processing requirements grow, customers can extend the system memory using high performance NVMe flash when the memory requirement grows above one or two terabytes. No need to add additional servers. Customers now have more flexibility to scale based on CPU-utilization rather than maxing out on the amount of memory they can install in a server.

Third, by reducing the number of servers needed to support large memory configurations, customers are able to reduce the TCO and operational costs for in-memory deployments. In addition, customers have the added capability to scale their storage capacity.

In-memory computing is vital to all of today's organizations, but using the right IMC solution is paramount. Supermicro's advanced in-memory computing solution – MemX – delivers the right insight faster, with a much lower CAPEX and OPEX, allowing any business to take advantage of the data they've gathered.

NOTES

A series of horizontal dashed lines for taking notes, spanning the width of the page.

About Super Micro Computer, Inc.

Supermicro® (NASDAQ: SMCI), the leading innovator in high-performance, high-efficiency server technology is a premier provider of advanced server Building Block Solutions® for Data Center, Cloud Computing, Enterprise IT, Hadoop/Big Data, HPC and Embedded Systems worldwide. Supermicro is committed to protecting the environment through its “We Keep IT Green™” initiative and provides customers with the most energy-efficient, environmentally-friendly solutions available on the market.

www.supermicro.com

No part of this document covered by copyright may be reproduced in any form or by any means — graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system — without prior written permission of the copyright owner.

Supermicro, the Supermicro logo, Building Block Solutions, We Keep IT Green, SuperServer, Twin, BigTwin, TwinPro, TwinPro², SuperDoctor are trademarks and/or registered trademarks of Super Micro Computer, Inc.

Ultrabook, Celeron, Celeron Inside, Core Inside, Intel, Intel Logo, Intel Atom, Intel Atom Inside, Intel Core, Intel Inside, Intel Inside Logo, Intel vPro, Itanium, Itanium Inside, Pentium, Pentium Inside, vPro Inside, Xeon, Xeon Phi, and Xeon Inside are trademarks of Intel Corporation in the U.S. and/or other countries.

All other brands names and trademarks are the property of their respective owners.

© Copyright 2017 Super Micro Computer, Inc. All rights reserved.

