



Big Memory for Financial Services



Learn about Big Memory



Time	Topic	Presenter
12:00 – 12:05	Introductions	Andrew Degnan
12:05 – 12:25	Optane Persistent Memory	Mark DeMarseilles
12:25 – 12:45	Big Memory	Charles Fan
12:45 – 12:55	Q&A	Andrew, Mark & Charles
12:55 – 01:00	Wrap-up & Raffle	Andrew Degnan

Tech Talk Hosts



Andrew Degnan
VP Sales
MemVerge



Mark DeMarseilles
Technical Sales
Intel




Charles Fan
CEO
MemVerge



Have Some Fun on Draft Day



Attend the entire Tech Talk and you get a MemVerge "What happens in memory stays in memory" t-shirt.



Fill out your Top 10 picks in our mock draft. The person or people that best match the actual Top 10 picks win an NFL jersey of their choice.



<https://bit.ly/memverge>



MemVerge Mock Draft


Mock Draft

For each selection 1-10, pick from the choices provided, or submit a name in the "other" box.

7. #1 Pick - Cincinnati Bengals

8. #2 Pick - Washington Redskins

9. #3 Pick - Detroit Lions



INTEL® OPTANE™ DC PERSISTENT MEMORY REVOLUTIONIZING MEMORY

Mark DeMarseilles Intel Optane Persistent Memory Technical Specialist

CLOUDIFICATION
EXPANSION
AND SCALE TO THE
POINT OF USE

CLO
UD

Deliver new services faster to handle the deluge of data

COM
MS

deliver operational efficiencies

Deliver secure infrastructure to protect data privacy

ENTERPRISE

DATA
CENTRIC
ERA

90%


OF WORLD'S DATA
WAS CREATED
IN THE PAST TWO
AND ONLY

2%


IS BEING USED

INTRODUCING


INTEL® OPTANE™ DC PERSISTENT MEMORY



FAST MEMORY



Enhance data
insights by
Redefining the
Memory and
Storage Hierarchy



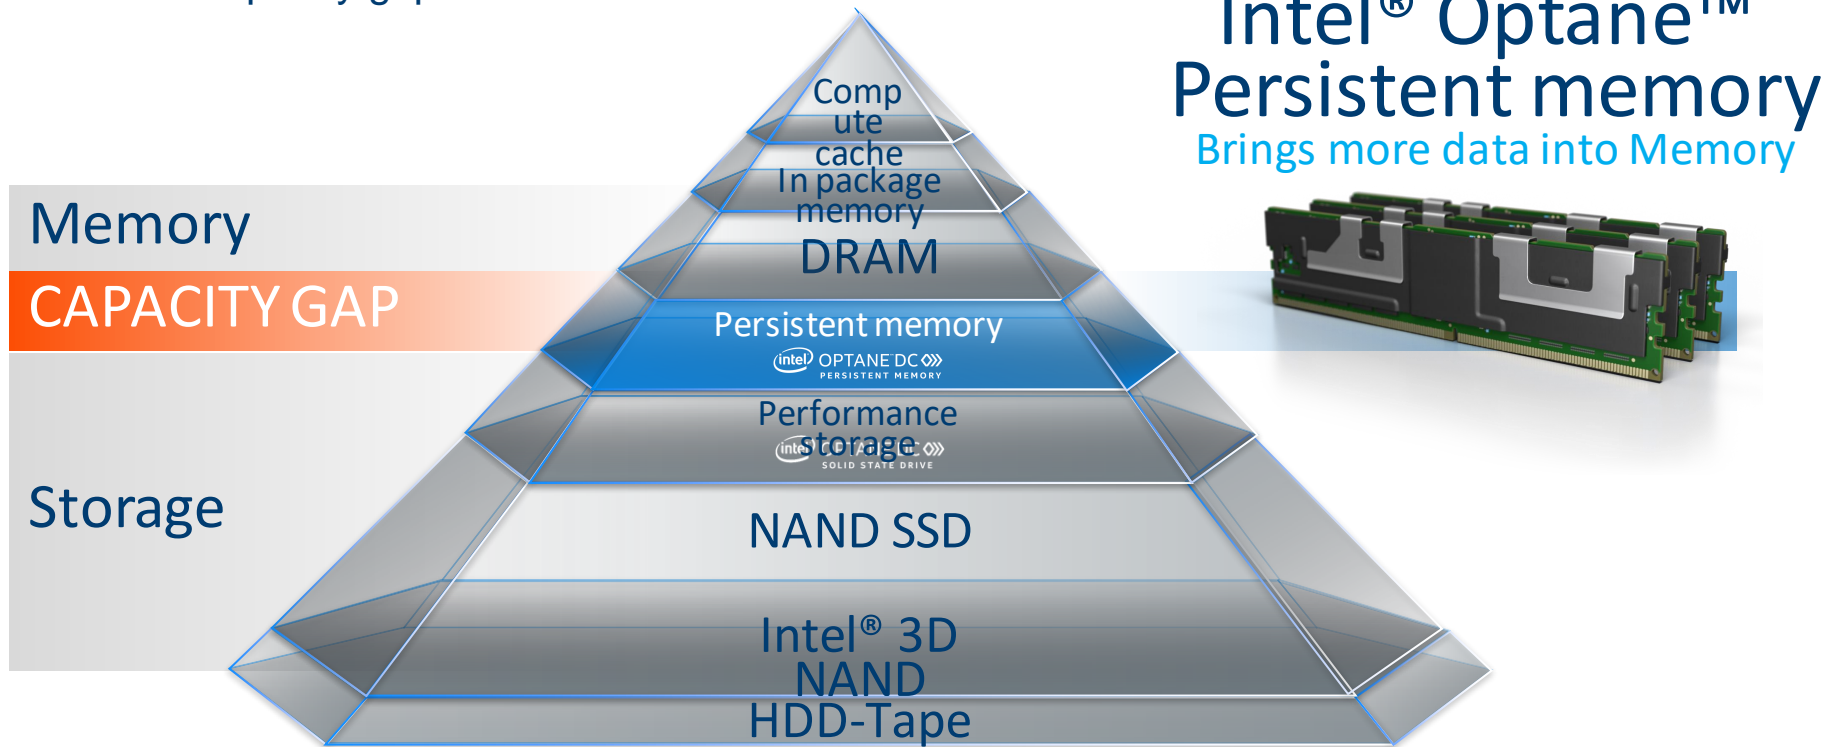
SIZE AND DATA
PERSISTENCE
OF STORAGE

Supported on 2nd Generation
Intel® Xeon® Scalable Processors
Platinum and Gold SKUs



Delivering to the Needs of Growing Data

Close the capacity gap





✓ Big and Affordable Memory

128, 256, 512GB Modules
DDR4 pin compatible

✓ Byte Addressable


Direct load/store access

✓ High Performance Storage

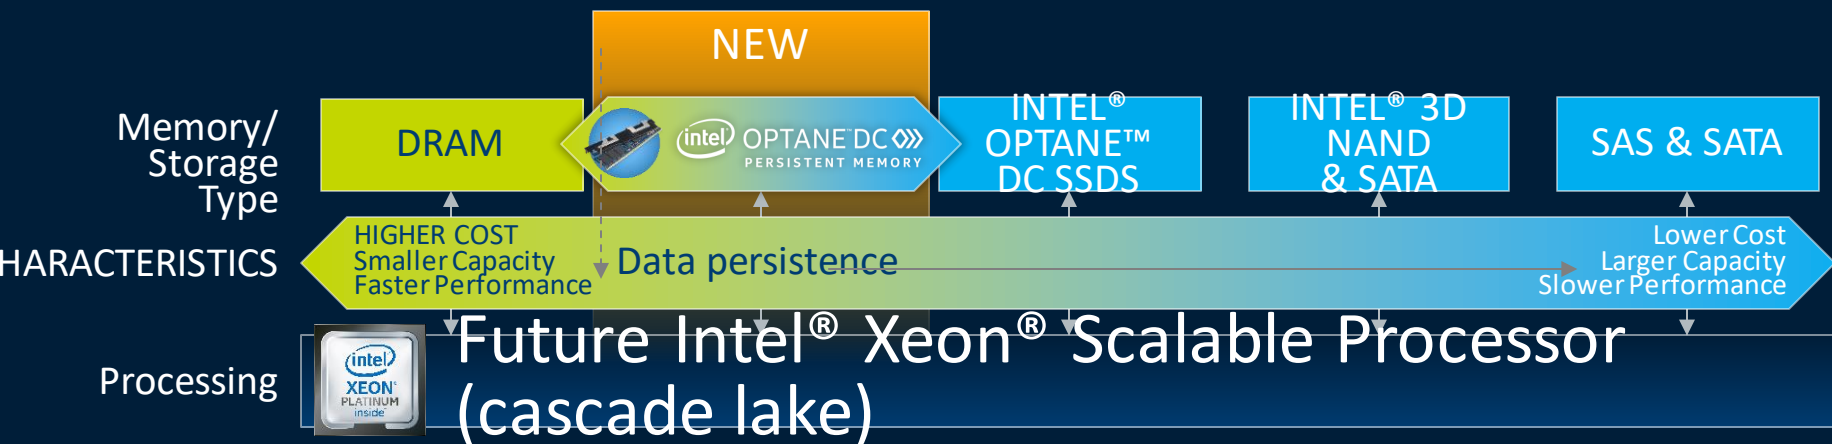
Native persistence

✓ High Reliability and Security

✓ Two Operational Modes



REDEFINING THE MEMORY AND STORAGE HIERARCHY



FAST MEMORY, PERSISTENCE OF STORAGE.
FLEXIBLE AND SCALABLE TO ACCELERATE YOUR DATA INSIGHTS.

Providing value on multiple vectors

→ Customer Pain Points ←

DDR4 is too expensive

Scale up is expensive

Not enough capacity

Operational inefficiencies

Workload **performance**

Storage is too slow

USE INTEL® OPTANE™ DC PERSISTENT MEMORY TO...

SAVE MORE

DRAM-ONLY
ALTERNATIVE

HELP
IMPROVE TCO

128 GB DDR4 compared to
128 GB DCPMM **45% lower**

256GB DDR4 compared to
256GB DCPMM **61% lower**

DO MORE

INCREASE
MEMORY SIZE

CONSOLIDATE
WORKLOADS

152 VMs 190 VMS **25% more**

768GB DDR4 192GB DDR4
1TB DCPMM **33% More**
18% Lower
Cost

GO FASTER

BREAK THE IO
BOTTLENECKS

ADD HIGH SPEED
STORAGE

DCPMM 180-340 nano sec **8GB/S**

NAND 100 micro sec **2GB/S**

What Problems Does Persistent Memory Address?

Current Problems

DRAM too costly	Scale up too expensive	Not enough capacity	Operational inefficiencies	Poor workload performance	Storage too slow
-----------------	------------------------	---------------------	----------------------------	---------------------------	------------------

Use INTEL® OPTANE™ DC persistent memory for...

💰 Cost savings		📊 productivity		⚙️ performance	
DRAM	Improve TCO	Increase Memory Size	Consolidate Workloads	Break I/O Bottlenecks	Add High-Speed Storage
Servers greater than 512GB	Workloads that need large or persistent memory	Large memory or software license fees per core	High VMs, with low CPU utilization	High disk I/O traffic	Tiered storage subsystem

OPERATING MODES

"APP DIRECT" MODE

APP/WORKLOAD DIRECT ACCESS TO




High SPEED, HIGH CAPACITY STORAGE
Persistent and much larger capacity



High availability/
less downtime



Significantly faster
storage



intel OPTANE DC
PERSISTENT MEMORY

MEMORY MODE

PLATFORM/OS/APP ACCESS TO
High SPEED, HIGH CAPACITY MEMORY



High capacity

Targeting >1.2X More VMs¹



Affordable capacity
128GB, 256GB and 512GB
Modules


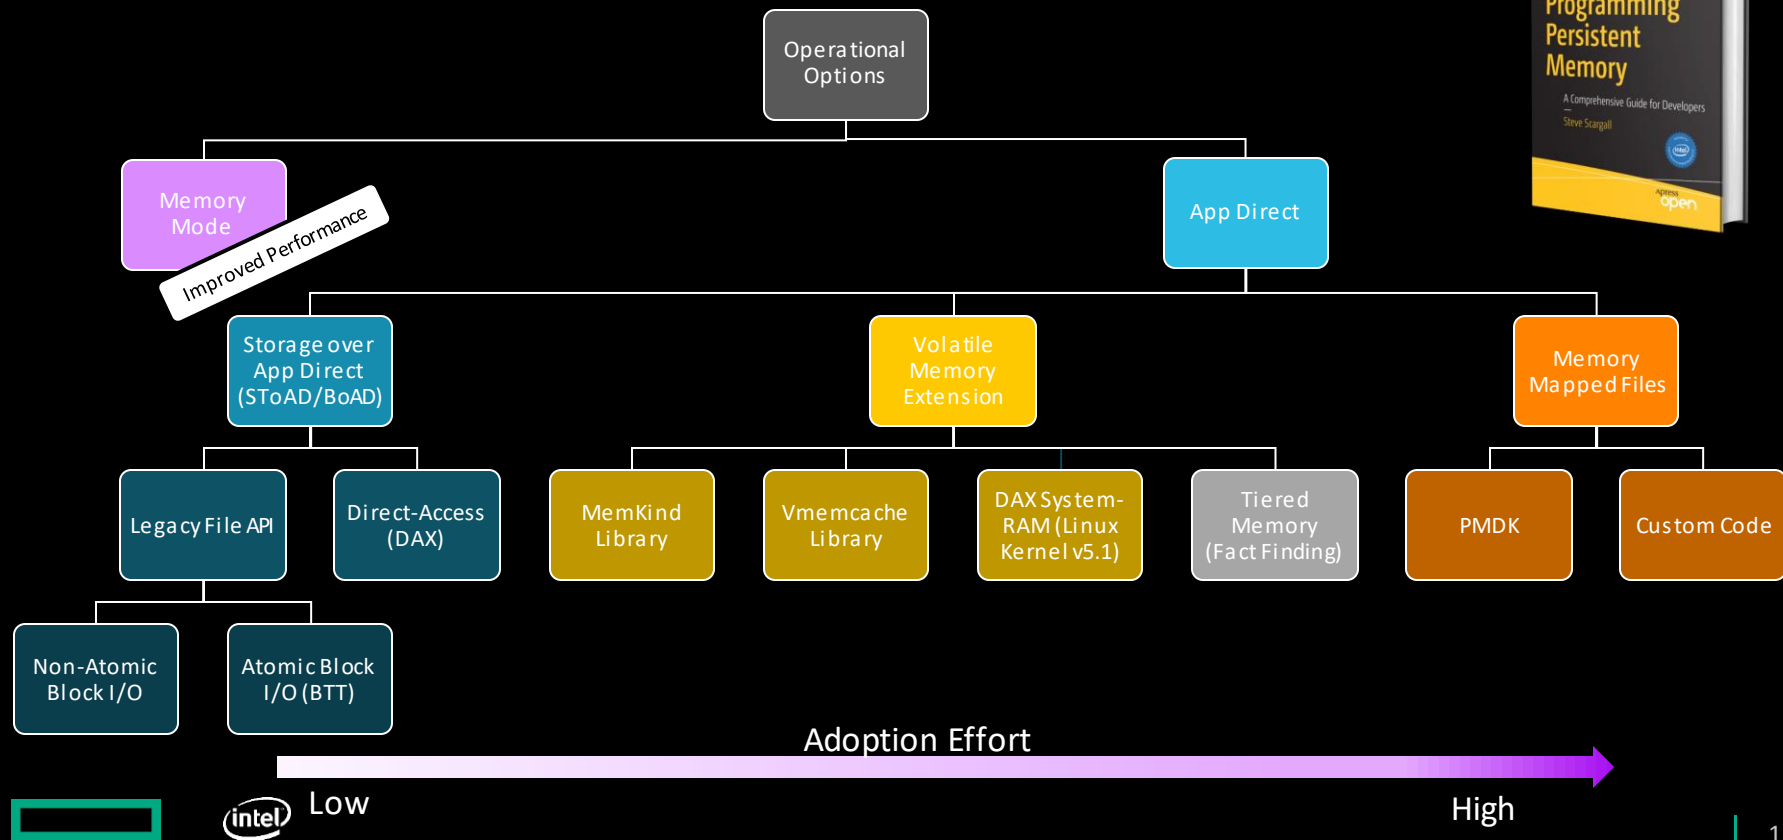


Ease of adoption
No code changes required

BUILT-IN FLEXIBILITY TO USE BOTH MODES SIMULTANEOUSLY

¹Performance results are based on testing as of dates shown in configuration and may not reflect all publicly available security updates. See configuration disclosure for details. No product or component can be absolutely secure. For more complete information about performance and benchmark results, visit www.intel.com/benchmarks.

ARCHITECTURAL OPTIONS FOR LINUX





Intel® Optane™ DC Persistent Memory

MEMORY INNOVATION 10 YEARS IN THE MAKING

ECOSYSTEM
SUPPORT

SOLUTION
OPTIMIZATION

TECHNOLOGY
INNOVATIONS

up to
36TB
to 8 SOCKET SYSTEM

9.
1B
NEW WORLD RECORDS
SAP
BW on HANA

up to
8
X
more VM instances
MEETING SLA


Performance results are based on testing: 8X (2/19/2019), and may not reflect all publicly available security updates. No product can be absolutely secure. See configuration disclosure for details. Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more information go to www.intel.com/benchmarks.

#data-centric



LEARN MORE AT


intel.com/optanedcpersistentmemory



intel OPTANE™ DC 
PERSISTENT MEMORY

Supported on future
Intel® Xeon® Scalable Processors
Platinum and Gold SKUs







Big Memory for Financial Services

Charles Fan
CEO
MemVerge



Introducing MemVerge



Founded by:

Shuki Bruck

XtremIO co-founder and
Caltech professor



Charles Fan

VMware storage BU leader
and creator of VSAN



Yue Li

Caltech post-doc and top researcher
on non-volatile memory



Introducing MemVerge



World-class team in Silicon Valley assembled from:



Our BIG MEMORY vision

All applications live in memory


Our mission

Open the door to Big Memory

A world of abundance,
persistence and high availability



MemVerge



INTEL® OPTANE™ DC PERSISTENT MEMORY REVOLUTIONIZING MEMORY


Why Not All Apps Can Run in Memory yet...



Not plug-and-play
App rewrite needed



Can't share memory
Siloed in servers



No Data Services
Crash recovery is slow




MemVerge Memory Machine™




Software
Subscription



Virtualizes
DRAM & PMEM




Low Latency
PMEM over RDMA



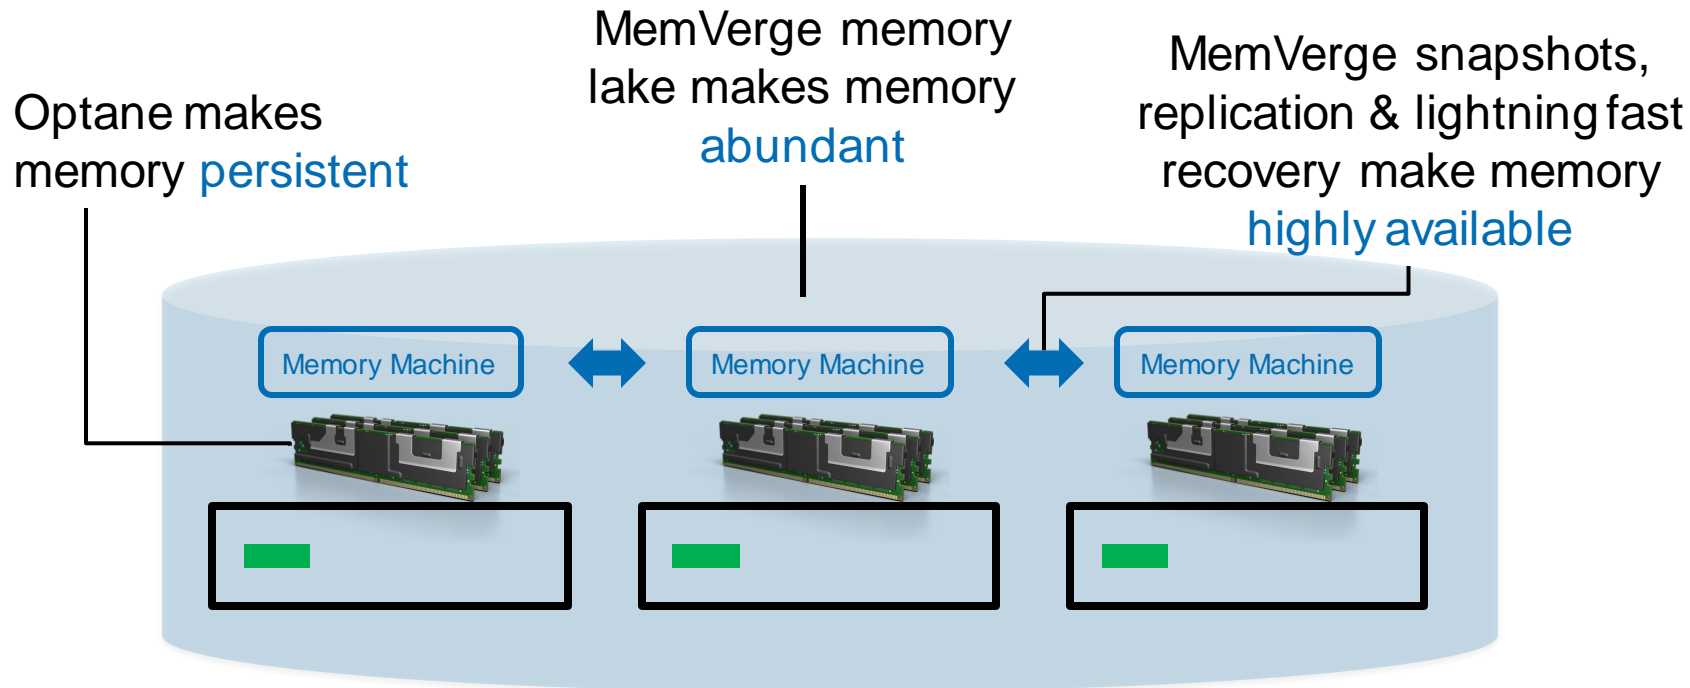
Plug Compatible
No re-writes



Memory Data Services
Snapshot, Replication, Tiering




Big Memory = Optane + Memory Machine



Financial Services are Killer Apps for Big Memory




Microseconds matter
In-Memory databases used



Data > memory
More capacity needed

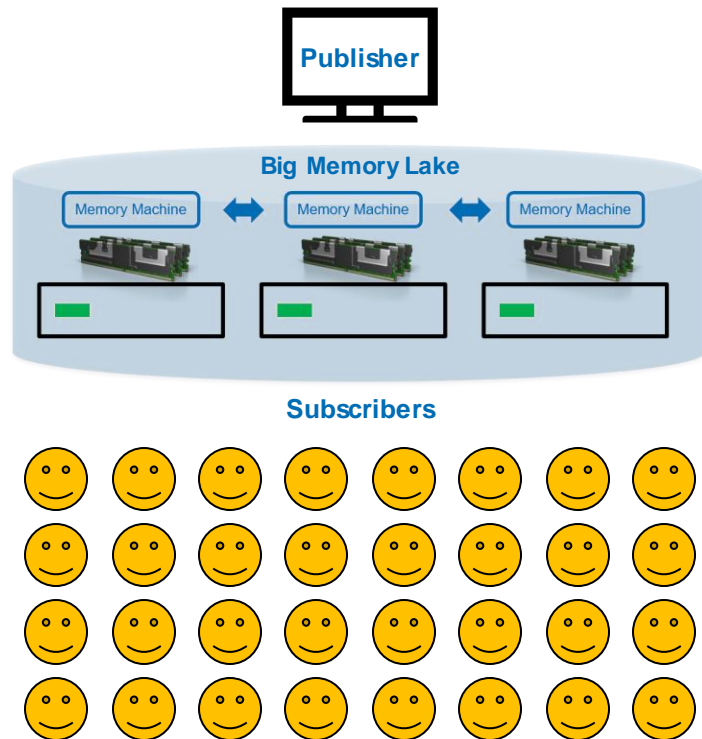


Growing blast zone
Memory data services needed




Objectives

1. Market data event stream published to multiple subscribers with the lowest latency
2. Achieve fairness between the subscriber processes
3. Persist the event stream without incurring significant performance penalties




Solution

1. Memory Machine™ software writes market data event stream to an in-memory bus
2. Background process commits the event stream to Persistent Memory synchronously or asynchronously
3. The event stream is replicated over RDMA to memory of other servers
4. Subscriber processes across all servers read the event stream with low latency



Results with 210 Subscribers





Results

0.5 uS

Avg. latency local host

3 uS

Avg. latency remote host


<2x avg.

99.9% tail latency



Problem

1. Need to run analytics, reporting or dev/test but concerned about taking performance hit on Primary instance
2. Application takes a long time to restart after crash or planned shutdown





In-Memory Database Cloning & Crash Recovery



Solution

1. Memory Machine takes instant snapshot, as frequently as every 1 minute
2. In-Memory Cloning easily creates a read replica of the primary instance using snapshot plus log replay
3. Fast restart from the database crash using snapshot plus log replay



+



In-Memory Database Cloning & Crash Recovery



Results

Every Minute

Fine grained Snapshots

Fast

Clone and Crash Recovery

Zero

Performance hit


Problem

1. When data is greater than the size of DRAM, AI/ML performance slows down dramatically
2. Memory-intensive inference jobs take a long time to load and restart




Solution

1. Create big memory lakes consisting of DRAM and PMEM to provide capacity needed for all data including models and embeddings
2. Fast data recovery and restart by using in-memory data snapshot




Test results: 1000 libs (1 million records) case

TPS



Data access latency





Results

Up to 4x

Transactions per second

Up to 100x

Lower latency

50%

Cost savings vs DRAM

MemVerge Vision for Big Memory Industry

By 2025...

1. Persistent Memory will be mainstream and data Infrastructure will be memory-centric.
2. Big Memory, consisting of PMEM and DRAM, will achieve petabyte-scale over clusters of servers interconnected by next-gen memory fabrics
3. Big Memory software will be needed to offer data services in memory, and every application will be run in-memory.

MemVerge Vision for Big Memory Industry





**What happens in memory
stays in memory...**





Q&A

Thank-you



Try it

Contract Andrew to sign-up for a PoC

Learn more about


Optane Persistent Memory at www.intel/optane

MemVerge at www.memverge.com

Contact:

mark.demarseilles@intel.com

andrew.degnan@memverge.com



<https://www.apress.com/us/book/9781484249314>



Raffle Winner

