Next gen cloud technologies that will accelerate business transformation

Raejeanne Skillern, Vice President Data Center Group
& General Manager Cloud Service Provider Platform Group

October 2018

@RaejeanneS

Cloud is Everywhere

ORACLE OPEN WORLD

intel

public cloud Growth continues

BY 2021

Digital Retail
$4.9T[1]

Digital Advertising
$400B[2]

Digital Video & Media
$120B[3]

Cloud Services
$300B[4]

1. Digital Retail – eMarketer Jan/March 2018
2. Digital Ads – eMarketer May 2018
3. Digital video/media – Juniper Research, Subscription Video on Demand, Dec 2017
4. Xaas (cloud services) – IDC Public Cloud Services Tracker Forecast 2017H2, May 2018

Cloud market Segment continues to accelerate

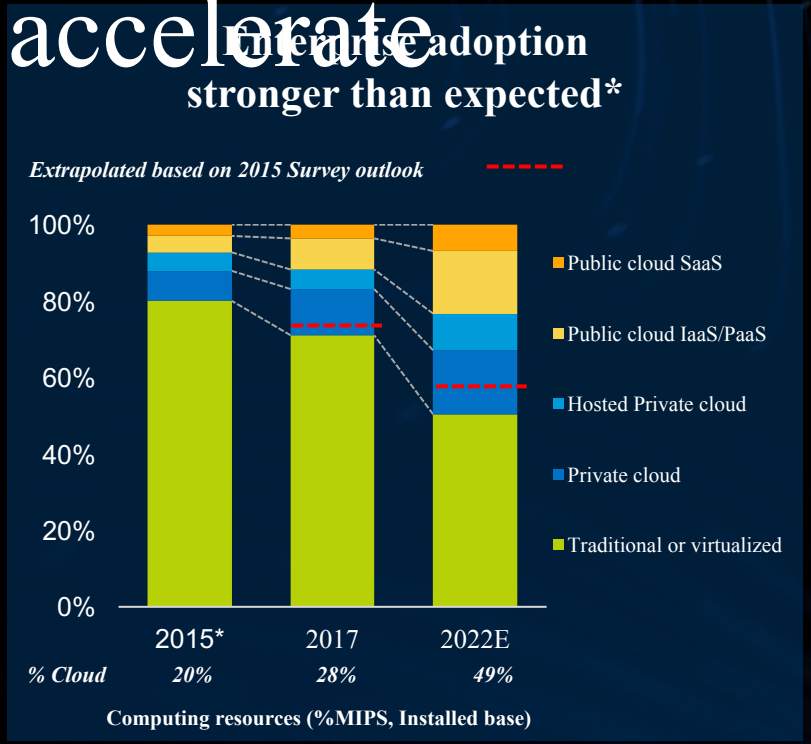

Complexity Rises By 2022
50% of enterprises will be intensively multicloud (10+ service providers), up from 10% in 2018 **

Workloads with highest-growth in Cloud expected in next 3-5 years: Business Communications, Web Infra/Apps, ERP, Decision Support, Database, and AI*

* Source: Intel Internal research
** IDC GRAC 2018 ICT Outlook

Cloud market Segment continues to accelerate

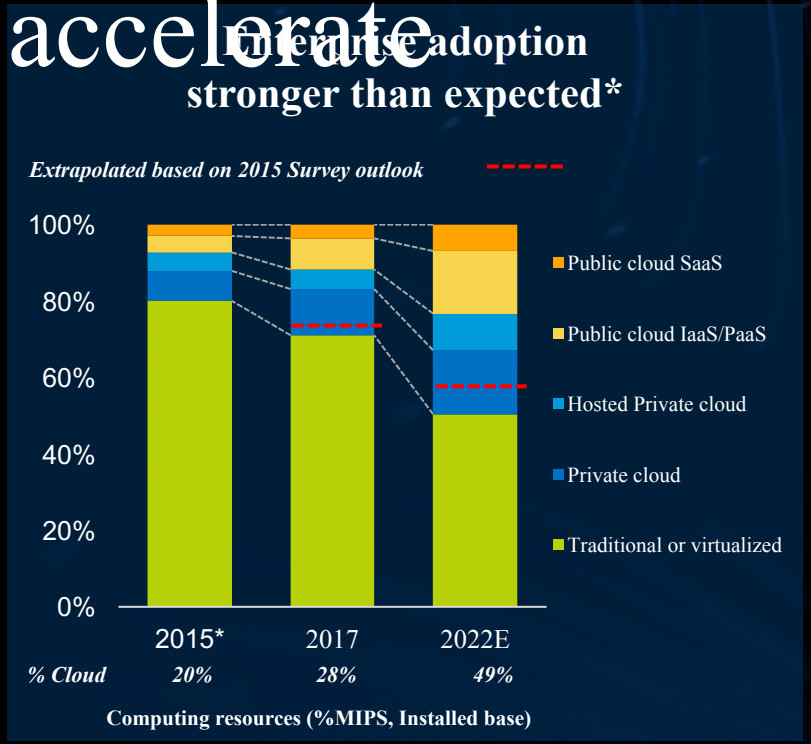

Complexity Rises By 2022
50% of enterprises will be intensively multicloud (10+ service providers), up from 10% in 2018 **

Workloads with highest-growth in Cloud expected in next 3-5 years: Business Communications, Web Infra/Apps, ERP, Decision Support, Database, and AI*

* Source: Intel Internal research
** IDC GRAC 2018 ICT Outlook

Our approach: Deep collaborations

» **Transformative platform technologies**

» **Solutions optimization & differentiation**

» **Accelerate the Journey to C**

ORACLE OPEN WORLD

intel

New Era Of Data Center technology

Data-Centric infrastructure

Process Everything

Move Faster

- intel SILICON PHOTONICS
- intel OMNI-PATH FABRIC
- intel ETHERNET

Store More

- intel OPTANE DC SOLID STATE DRIVE
- intel OPTANE DC PERSISTENT MEMORY

Process Everything

next Intel® Xeon® Scalable processor: Cascade Lake

Performance enhancements

Intel® Deep Learning Boost:
Vector Neural network Instruction (VNNI)

Intel® OPTANE™ DC PERSISTENT MEMORY

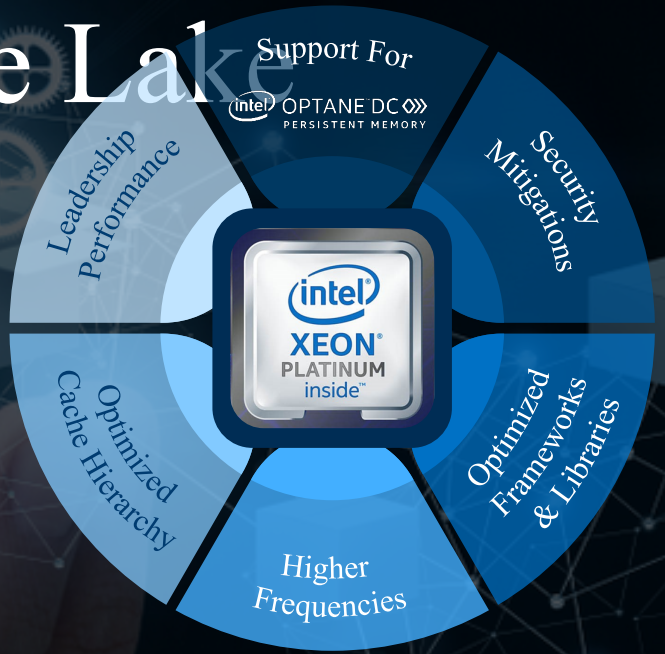

Reinventing Xeon for AI

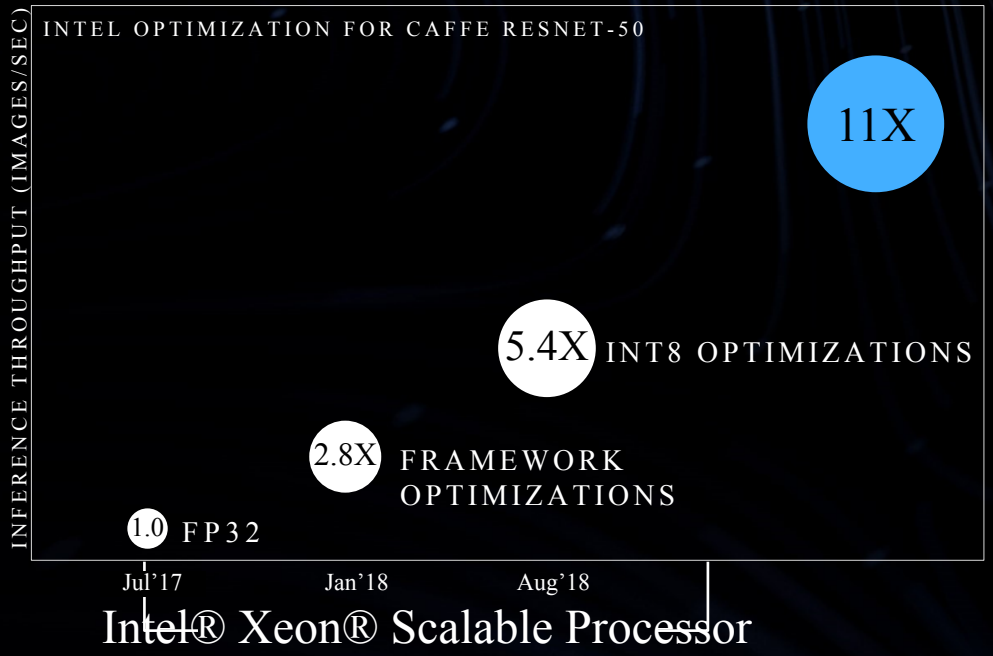

INTRODUCING

INTEL® DEEP LEARNING BOOST

VECTOR NEURAL NETWORK INSTRUCTION (VNNI)

For inference Acceleration

Framework & library support
Caffe mxnet TensorFlow
intel MKL-DNN

1 Intel® Optimization for Caffe Resnet-50 performance does not necessarily represent other Framework performance.
2 Based on Intel internal testing: 1X (7/11/2017), 2.8X (1/19/2018) and 5.4X (7/26/2018) performance improvement based on Intel® Optimization for Caffe Resnet-50 inference throughput performance on Intel® Xeon® Scalable Processor.
3 11X (7/25/2018) Results have been estimated using internal Intel analysis, and provided to you for informational purposes. Any differences in your system hardware, software or configuration may affect your actual performance.
Performance results are based on testing as of 7/11/2017(1x), 1/19/2018(2.8x) & 7/26/2018(5.4) and may not reflect all publically available security update. See configuration disclosure for details (config 1). No product can be absolutely secure.Optimization Notice: Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.
Other names and brands may be claimed as the property of others.

Reinventing Xeon for AI

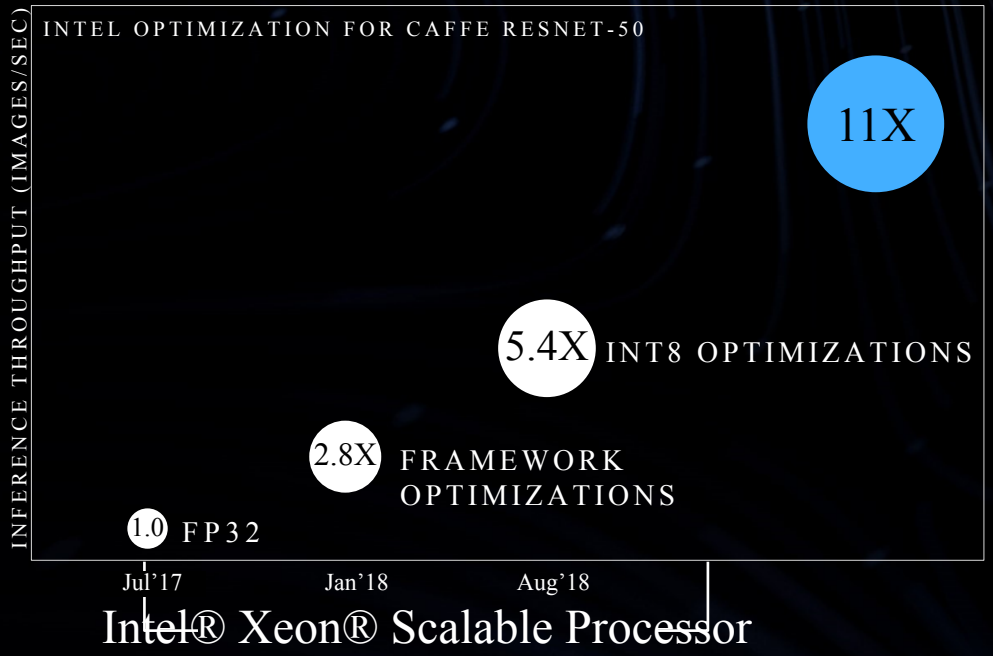

INTRODUCING

INTEL® DEEP LEARNING BOOST

VECTOR NEURAL NETWORK INSTRUCTION (VNNI)

For inference Acceleration

Framework & library support
Caffe mxnet TensorFlow
intel MKL-DNN

1 Intel® Optimization for Caffe Resnet-50 performance does not necessarily represent other Framework performance.
2 Based on Intel internal testing: 1X (7/11/2017), 2.8X (1/19/2018) and 5.4X (7/26/2018) performance improvement based on Intel® Optimization for Caffe Resnet-50 inference throughput performance on Intel® Xeon® Scalable Processor.
3 11X (7/25/2018) Results have been estimated using internal Intel analysis, and provided to you for informational purposes. Any differences in your system hardware, software or configuration may affect your actual performance.
Performance results are based on testing as of 7/11/2017(1x), 1/19/2018(2.8x) & 7/26/2018(5.4) and may not reflect all publically available security update. See configuration disclosure for details (config 1). No product can be absolutely secure.Optimization Notice: Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.
Other names and brands may be claimed as the property of others.

Move **Faster**

Intel Connectivity Portfolio

Intel® Omni-Path Fabric

LEADING HPC FABRICS

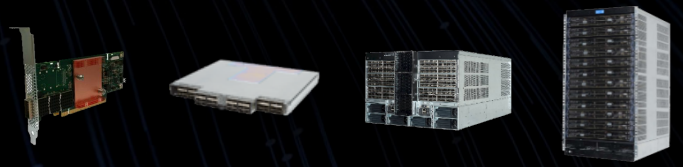

Intel® Ethernet

#1 MSS HIGH SPEED1 ETHERNET
COMING 2019
CASCADE GLACIER SMARTNIC

Intel® Silicon Photonics

SILICON **INTEGRATION**
SILICON **MANUFACTURING**
SILICON **SCALE**

1. High speed = 10GbE and above
Source: Connectivity TAM includes Ethernet, High Performance Fabrics, and Silicon Photonics and is based on amalgamation of analyst data and Intel analysis, based upon current expectations and available information and are subject to change without notice..

Move **Faster**

Intel Connectivity Portfolio

Intel® Omni-Path Fabric

LEADING HPC FABRICS

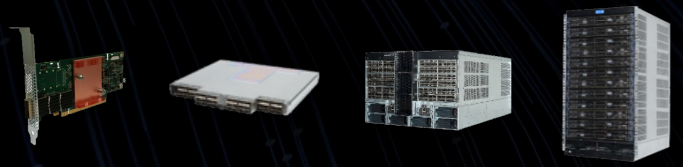

Intel® Ethernet

#1 MSS HIGH SPEED1 ETHERNET
COMING 2019
CASCADE GLACIER SMARTNIC

Intel® Silicon Photonics

SILICON **INTEGRATION**
SILICON **MANUFACTURING**
SILICON **SCALE**

1. High speed = 10GbE and above
Source: Connectivity TAM includes Ethernet, High Performance Fabrics, and Silicon Photonics and is based on amalgamation of analyst data and Intel analysis, based upon current expectations and available information and are subject to change without notice..

intel OPTANE™ DC PERSISTENT MEMORY

intel XEON PLATINUM inside™

Now shipping for revenue to select customers

- Big and Affordable Memory
- High Performance Storage
- Direct Load/Store Access
- Native Persistence

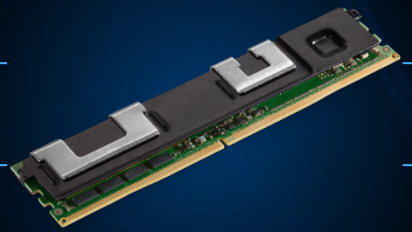

- 128, 256, 512GB
- DDR4 Pin Compatible
- Hardware Encryption
- High Reliability

intel OPTANE™ DC PERSISTENT MEMORY

intel XEON PLATINUM inside™

Now shipping for revenue to select customers

- Big and Affordable Memory
- High Performance Storage
- Direct Load/Store Access
- Native Persistence

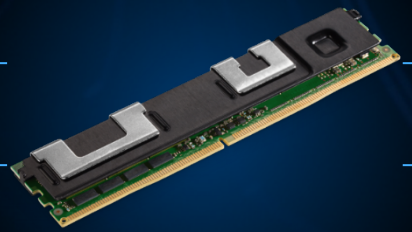

- 128, 256, 512GB
- DDR4 Pin Compatible
- Hardware Encryption
- High Reliability

Store More

Disrupting the storage/memory hierarchy

Oracle TimesTen In-Memory Database

Ananth Raghavan, VP Oracle TimesTen In-Memory Database

Store More

Disrupting the storage/memory hierarchy

Oracle TimesTen In-Memory Database

Ananth Raghavan, VP Oracle TimesTen In-Memory Database

TimesTen: World's Fastest OLTP Database

In-Memory Relational Database

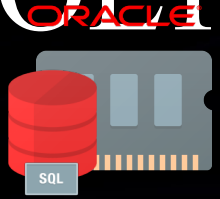

- Pure in-memory
- ACID compliant
- Standard SQL
- Entire database in RAM

Persistent and Recoverable

- Database and transaction logs persisted on local storage
- Automatic recovery after failure

Extremely Fast

- Microseconds response time
- 10s of millions of TPS
- Billions of queries per second

Highly Available, Extremely Scalable

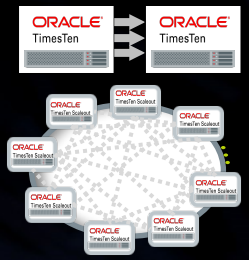

- High performance replication
- Elastic scalability with K-safety

*Based on YCSB benchmark and TPTBM benchmark testing

ORACLE

TimesTen: World's Fastest OLTP Database

In-Memory Relational Database

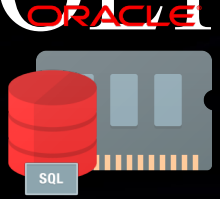

- Pure in-memory
- ACID compliant
- Standard SQL
- Entire database in RAM

Persistent and Recoverable

- Database and transaction logs persisted on local storage
- Automatic recovery after failure

Extremely Fast

- Microseconds response time
- 10s of millions of TPS
- Billions of queries per second

Highly Available, Extremely Scalable

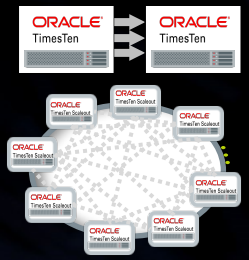

- High performance replication
- Elastic scalability with K-safety

*Based on YCSB benchmark and TPTBM benchmark testing

ORACLE

In-Memory Database Startup

Today

- Copy database image from persistent storage into volatile DRAM
- 1.35 TB database

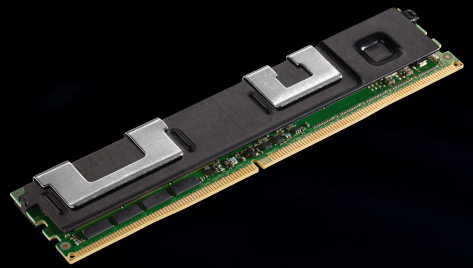

With Intel persistent memory

- Different implementation
- Database in Intel persistent memory **is** the persistent storage
- 'Startup' is instantaneous
- 2.7 TB database

In-Memory Database Startup

Today

- Copy database image from persistent storage into volatile DRAM
- 1.35 TB database

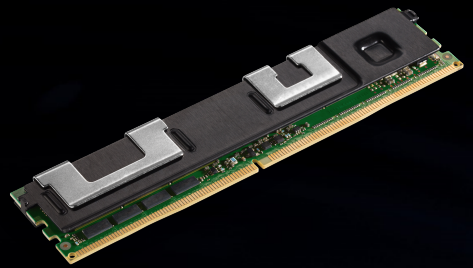

With Intel persistent memory

- Different implementation
- Database in Intel persistent memory **is** the persistent storage
- 'Startup' is instantaneous
- 2.7 TB database

In-Memory Database Startup

Today

- Copy database image from persistent storage into volatile DRAM
- 1.35 TB database

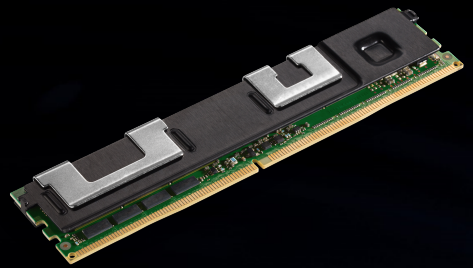

With Intel persistent memory

- Different implementation
- Database in Intel persistent memory **is** the persistent storage
- 'Startup' is instantaneous
- 2.7 TB database

In-Memory Database Startup

Today

- Copy database image from persistent storage into volatile DRAM
- 1.35 TB database

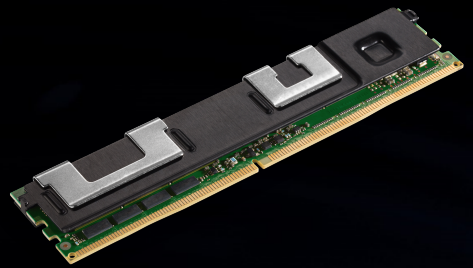

With Intel persistent memory

- Different implementation
- Database in Intel persistent memory **is** the persistent storage
- 'Startup' is instantaneous
- 2.7 TB database

intel OPTANE DC PERSISTENT MEMORY

In-Memory Database Startup

Today

- Copy database image from persistent storage into volatile DRAM
- 1.35 TB database

→ **> 10 min**

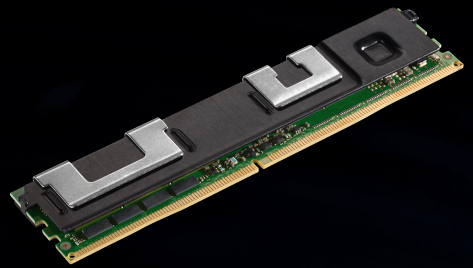

With Intel persistent memory

- Different implementation
- Database in Intel persistent memory **is** the persistent storage
- 'Startup' is instantaneous
- 2.7 TB database

→ **< 1 second**

Performance results are based on Oracle TimesTen IMDB, running TPTPM benchmark. Hardware configuration for Baseline and Persistent Memory systems is 2 Sockets Intel Skylake 8180M SKU, 28 cores, 2 threads per core. BIOS version1.0134. Storage is 2x6TB SSD DC P4608. Operating system is Redhat 7.5, Kernel 4.18.
Baseline system has 1,536Gb DDR4 quad rank @2666MHz.
Persistent Memory system has 192GB DDR4 dual rank @2666MHz and 6TB of Intel Optane DC Persistent Memory

ORACLE OPEN WORLD

intel

Transaction Durability

Today

- Log buffer in volatile DRAM
- Transactions commit to buffer…
- … then buffer written synchronously to storage

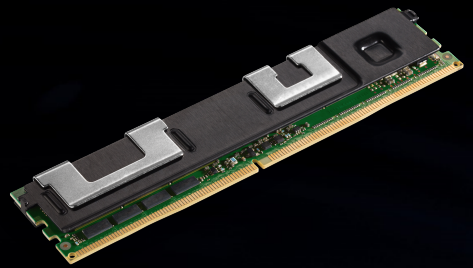

With Intel persistent memory

- Different implementation
- Log buffer is in Intel persistent memory
- Persistence is immediate on commit
  - No need to wait for write to storage

Transaction Durability

Today

- Log buffer in volatile DRAM
- Transactions commit to buffer…
- … then buffer written synchronously to storage

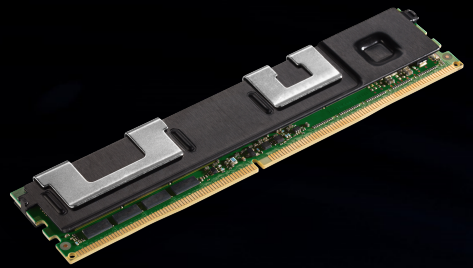

With Intel persistent memory

- Different implementation
- Log buffer is in Intel persistent memory
- Persistence is immediate on commit
  - No need to wait for write to storage

Transaction Durability

Today

- Log buffer in volatile DRAM
- Transactions commit to buffer…
- … then buffer written synchronously to storage

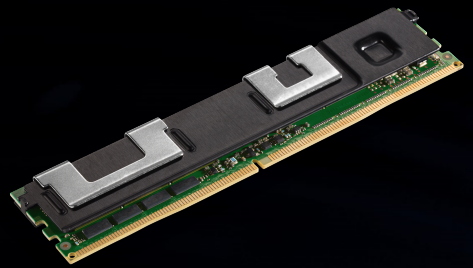

With Intel persistent memory

- Different implementation
- Log buffer is in Intel persistent memory
- Persistence is immediate on commit
  - No need to wait for write to storage

Transaction Durability

Today

- Log buffer in volatile DRAM
- Transactions commit to buffer…
- … then buffer written synchronously to storage

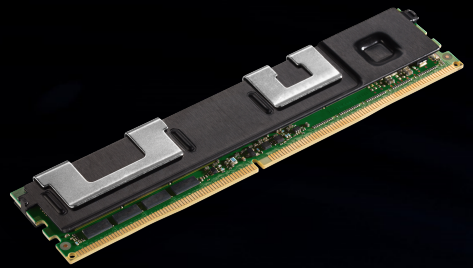

With Intel persistent memory

- Different implementation
- Log buffer is in Intel persistent memory
- Persistence is immediate on commit
  - No need to wait for write to storage

Transaction Durability

Today

- Log buffer in volatile DRAM
- Transactions commit to buffer…
- … then buffer written synchronously to storage
- Average throughput = **176K TPS**

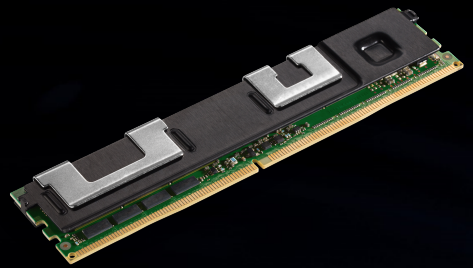

With Intel persistent memory

- Different implementation
- Log buffer is in Intel persistent memory
- Persistence is immediate on commit
  - No need to wait for write to storage
- Average throughput = **1.16M TPS**

6.49x improvement!

Performance results are based on Oracle TimesTen IMDB, running TPTPM benchmark. Hardware configuration for Baseline and Persistent Memory systems is 2 Sockets Intel Skylake 8180M SKU, 28 cores, 2 threads per core. BIOS version1.0134. Storage is 2x6TB SSD DC P4608. Operating system is Redhat 7.5, Kernel 4.18.
Baseline system has 1,536Gb DDR4 quad rank @2666MHz.
Persistent Memory system has 192GB DDR4 dual rank @2666MHz and 6TB of Intel Optane DC Persistent Memory

Summary

Intel persistent memory enables new capabilities and enhanced performance for In-Memory Databases

- From >10min to less than 1 second - faster database startup
- 6.49x faster durable transaction performance
- TimesTen Info:
  - Hands On Lab: 10/23 5:15pm, 10/24 12:45pm
  - TimesTen Session: 10/24 11:15 am
  - TimesTen Demo Booth

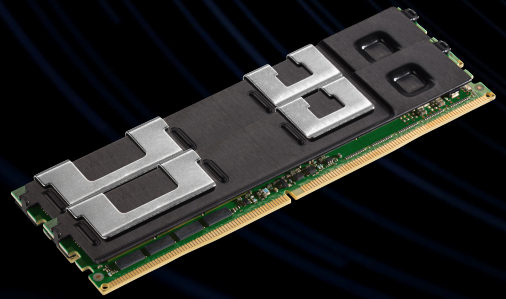

Summary

Intel persistent memory enables new capabilities and enhanced performance for In-Memory Databases

- From >10min to less than 1 second - faster database startup
- 6.49x faster durable transaction performance
- TimesTen Info:
  - Hands On Lab: 10/23 5:15pm, 10/24 12:45pm
  - TimesTen Session: 10/24 11:15 am
  - TimesTen Demo Booth

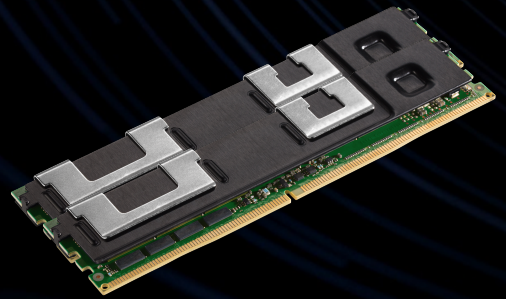

Our approach: Deep collaborations

- Transformative platform technologies
- Solutions optimization & differentiation
- Accelerate the Journey to C

Exposing hardware Features

Oracle Database In-Memory utilizes Intel® Advanced Vector Extensions for processing columnar data.

Oracle Database In-Memory and Intel AVX512 together speed up query performance by increasing the number of data values analyzed per second from Millions to Billions *

Source: *https://www.oracle.com/technetwork/database/options/database-in-memory-brief-2215626.pdf

Oracle-Intel Collaboration Model

ORACLE CLOUD

ORACLE OPEN WORLD

intel

Customer Success Story: Oracle Autonomous Data Warehouse Cloud

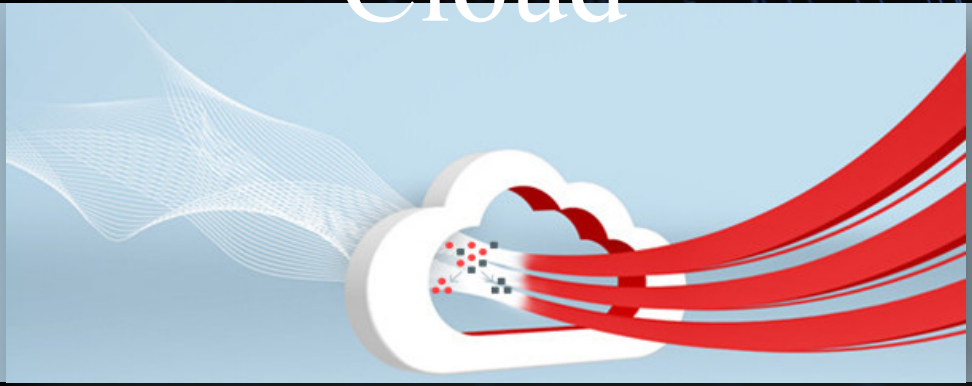

ORACLE OPEN WORLD

intel

Customer Success Story: Oracle Autonomous Data Warehouse Cloud

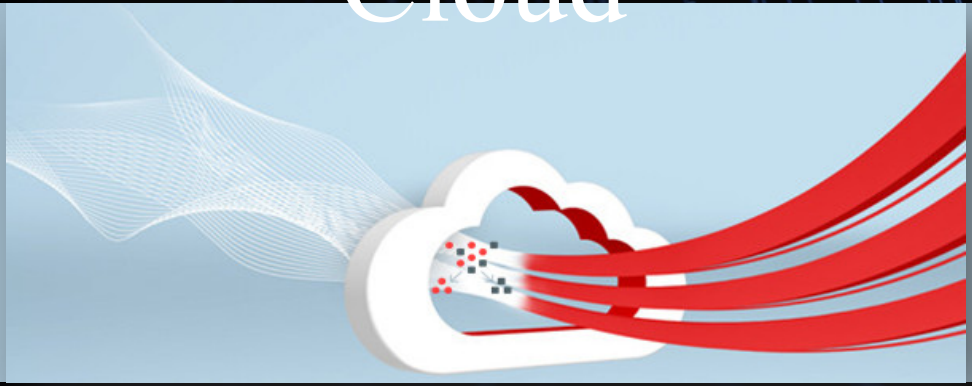

ORACLE OPEN WORLD

intel

Customer Success Story: Oracle Autonomous Data Warehouse Cloud

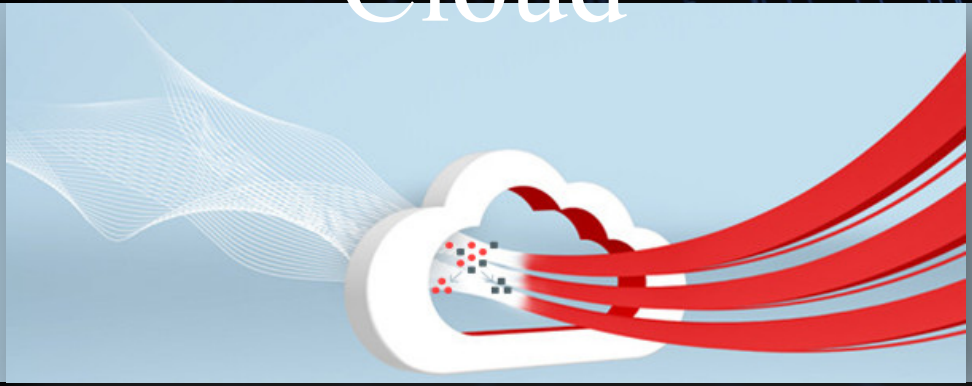

ORACLE OPEN WORLD

intel

Customer Success Story: Oracle Autonomous Data Warehouse Cloud

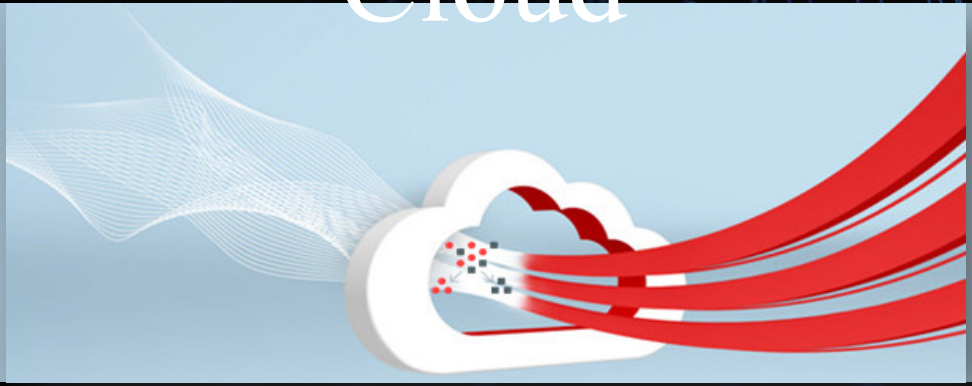

ORACLE OPEN WORLD

intel

Our approach: Deep collaborations

» Transformative platform technologies

» Solutions optimization & differentiation

» Accelerate the Journey to C

Oracle Cloud Infrastructure

James Stanbridge
ORACLE VP IAAS PRODUCT MANAGEMENT EMEA & JAPAC

Oracle Cloud Infrastructure

James Stanbridge
ORACLE VP IAAS PRODUCT MANAGEMENT EMEA & JAPAC

Oracle Cloud Infrastructure

James Stanbridge
ORACLE VP IAAS PRODUCT MANAGEMENT EMEA & JAPAC

Oracle Cloud Infrastructure

James Stanbridge
ORACLE VP IAAS PRODUCT MANAGEMENT EMEA & JAPAC

Oracle Cloud Infrastructure

James Stanbridge
ORACLE VP IAAS PRODUCT MANAGEMENT EMEA & JAPAC

Oracle Cloud Infrastructure

James Stanbridge
ORACLE VP IAAS PRODUCT MANAGEMENT EMEA & JAPAC

Oracle & Intel Cloud ignition p

- Accelerates Your Journey to Cloud
- Proof of Concepts
- Sample Cloud Migration Services

ORACLE OPEN WORLD

intel

Summary

Intel and Oracle are your #1 Technology Partner

Deep collaboration & customization drives innovative differentiation

Intel and Oracle are investing to accelerate your growth

Transformative platform technologies

Software and solutions co-engineering

large scale market acceleration programs

@RaejeanneS

ORACLE OPEN WORLD

intel

Oculus Rift Virtual Reality System + Acer Predator Intel Core i7 Gaming Laptop Bundle

THANK YOU

Notices & Disclaimers

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. No product can be absolutely secure.

Tests document performance of components on a particular test, in specific systems. Differences in hardware, software, or configuration will affect actual performance. For more complete information about performance and benchmark results, visit http://www.intel.com/benchmarks .

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 complete information visit http://www.intel.com/benchmarks .

Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

The benchmark results may need to be revised as additional testing is conducted. The results depend on the specific platform configurations and workloads utilized in the testing, and may not be applicable to any particular user's components, computer system or workloads. The results are not necessarily representative of other benchmarks and other benchmark results may show greater or lesser impact from mitigations.

Intel does not control or audit third-party benchmark data or the web sites referenced in this document. You should visit the referenced web site and confirm whether referenced data are accurate.