

# INTEL<sup>®</sup> OPTANE<sup>™</sup> DC PERSISTENT MEMORY

Harald Servat  
MNHack19

# NOTICES AND 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 or component 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® Advanced Vector Extensions (Intel® AVX)\* provides higher throughput to certain processor operations. Due to varying processor power characteristics, utilizing AVX instructions may cause a) some parts to operate at less than the rated frequency and b) some parts with Intel® Turbo Boost Technology 2.0 to not achieve any or maximum turbo frequencies. Performance varies depending on hardware, software, and system configuration and you can learn more at <http://www.intel.com/go/turbo>.

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.

Cost reduction scenarios described are intended as examples of how a given Intel-based product, in the specified circumstances and configurations, may affect future costs and provide cost savings. Circumstances will vary. Intel does not guarantee any costs or cost reduction.


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.

Intel, the Intel logo, Intel Optane, Intel Xeon and VTune are trademarks of Intel Corporation in the U.S. and/or other countries.


\*Other names and brands may be claimed as property of others.

© 2019 Intel Corporation.

# INTRODUCING INTEL® OPTANE™ DC PERSISTENT MEMORY



**FAST MEMORY**



**ENHANCE DATA INSIGHTS BY  
REDEFINING THE MEMORY AND  
STORAGE HIERARCHY**

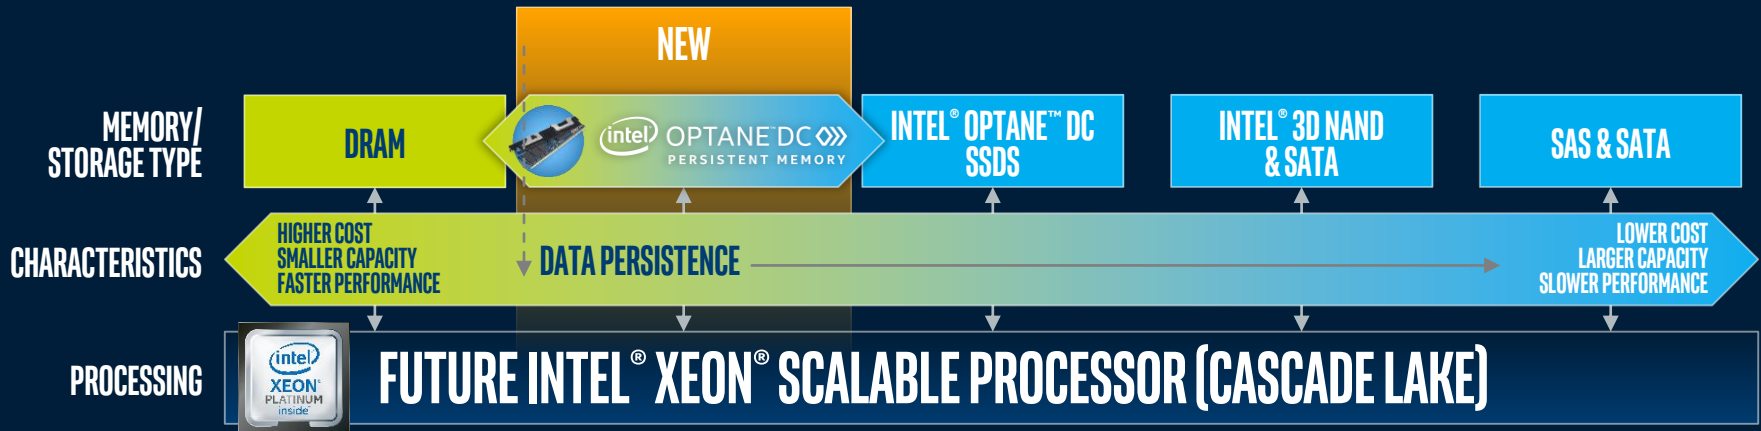


**SIZE AND DATA PERSISTENCE  
OF STORAGE**

Supported on 2<sup>nd</sup> Generation  
Intel® Xeon® Scalable Processors  
Platinum and Gold SKUs



# REDEFINING THE MEMORY AND STORAGE HIERARCHY



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

# INTEL® OPTANE™ DC PERSISTENT MEMORY


CPU's support Large Memory  
 "M" = 2.0TB, "L" = 4.5TB

28C 2.7G	205W	8280
28C 2.2G	165W	8276
24C 2.4G	165W	8260
18C 2.6G	150W	6240
22C 2.1G	140W	6238
10C 2.5G	85W	5215

Intel 82xx processors



Intel 62xx  
Intel 52xx processors



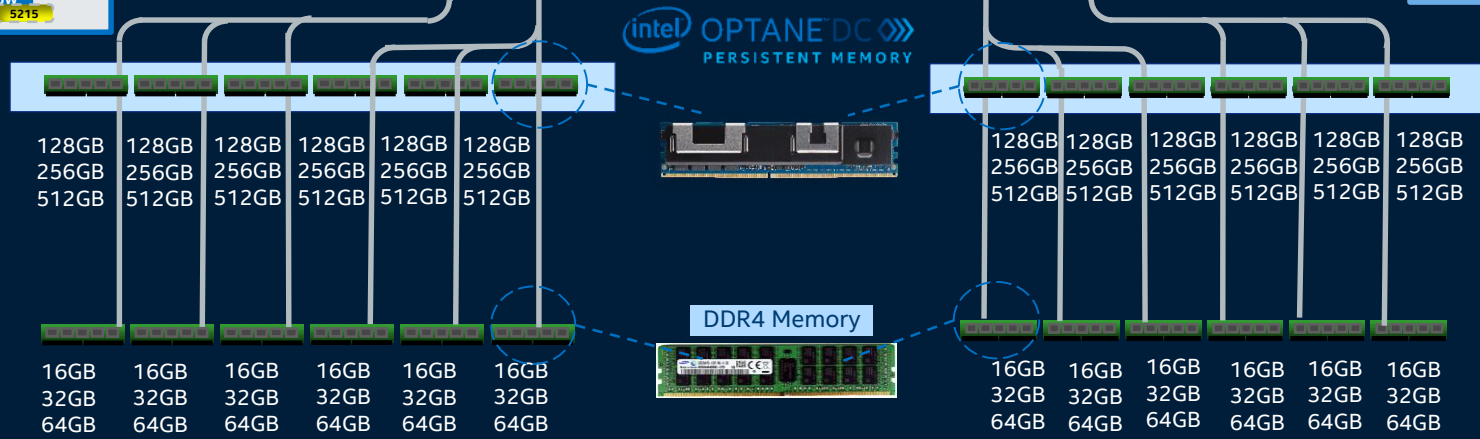
Intel 82xx processors



Intel 62xx  
Intel 52xx processors

## INTEL OPTANE DC PERSISTENT MEMORY

- DIMM Capacity**
  - 128, 256, 512GB
- Speed**
  - 2666 MT/sec
- Capacity per CPU**
  - Up to 3TB (not including DRAM)




Intel Server (Purley) Platforms support up to (12) DDR4 Slots per CPU  
 Intel DCPMM can be configured up to (6) DIMMs per CPU  
 General ratio (DRAM:DCPMM) is between 1:4 to 1:16, with minimum 1:2

# USAGE OPTION 1: MEMORY MODE

- PMEM used as Large Volatile Memory
  - DRAM acting as L4 cache
- Looks to SW stack as large memory
- up to 6TB in a 2S server at 90% the performance of DRAM
- Good for Memory-constrained workloads
  - Redis/Memcached cache tier
  - Analytics
  - Higher VM & Container density
- Infrastructure cost savings by doing more with each compute node (scale-in)

## Memory Mode



DRAM as L4 Cache  
Hidden from OS



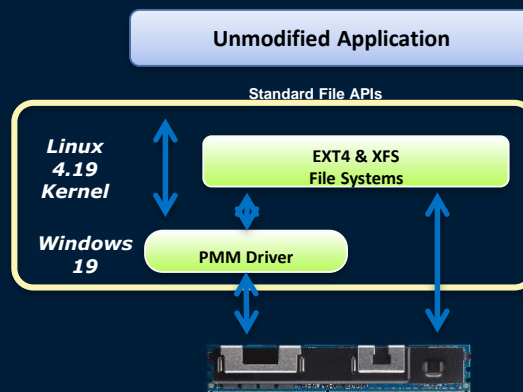
Persistent  
Memory



Volatile Memory Capacity  
visible to the OS

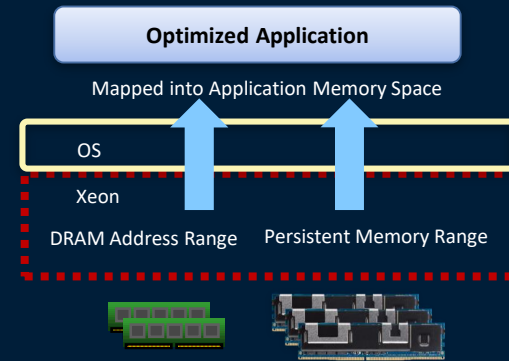
# USAGE OPTION 2: AS FAST STORAGE

- PMEM presented as Fast Storage
- ANY Application
  - On Linux 4.19 or later kernel: RHEL 7.6, SLES 12, CentOS 7.6, Ubuntu 18.04.x
  - Windows Server '19
- Value Prop: Lowest latency and highest BW storage
- Good for IOPS-limited Apps
  - Technical Compute
  - Databases: MySQL, MongoDB, Cassandra, etc.



# USAGE OPTION 3: APPLICATION-DIRECT (PERSISTENT) MEMORY WITH DAX

- Application directly uses persistent memory as byte addressable, load/store memory
- Highest Performance Lowest Latency, Highest Capacity, Highest Availability
- Persistent or Volatile through memkind
- 4.19+ Linux kernel
- Windows Server 2019



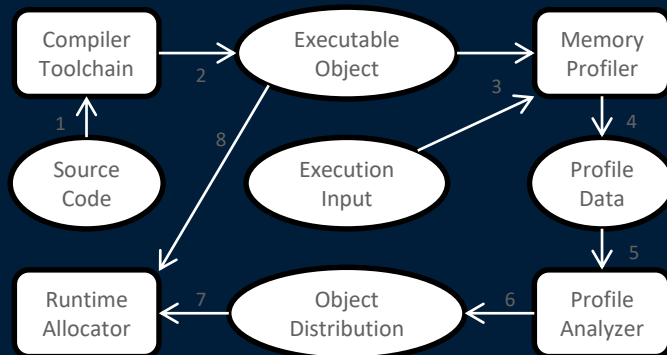


# (PART OF) BSC-INTEL COLLABORATION

Object-differentiated data-oriented profiling + distribution algorithm (analysis)

- 1) Profile to determine per-object last-level cache misses / avg. access time
- 2) Assess the optimal distribution of the different objects among the memory subsystem
  - Minimize processor stall cycles

Similar approach used earlier in  
Intel® Xeon Phi™ with a number of  
positive results



# DEVELOPER RESOURCES FOR INTEL® OPTANE™ DC PERSISTENT MEMORY

Find the PMDK  
(Persistent Memory Development Kit)  
at <http://pmem.io/pmdk/>

## Getting Started

- Intel IDZ persistent memory - <https://software.intel.com/en-us/persistent-memory>
- Entry into overall architecture - <http://pmem.io/2014/08/27/crawl-walk-run.html>
- Emulate persistent memory - <http://pmem.io/2016/02/22/pm-emulation.html>

Persistent Memory  
Programming Video Series -  
<https://software.intel.com/en-us/persistent-memory/get-started/series>

## Linux Resources


- Linux\* Community Pmem Wiki - <https://nvdimm.wiki.kernel.org/>
- Pmem enabling in SUSE Linux Enterprise 12 SP2 - <https://www.suse.com/communities/blog/nvdimm-enabling-suse-linux-enterprise-12-service-pack-2/>

## Windows\* Resources

- Using Byte-Addressable Storage in Windows Server 2016 - <https://channel9.msdn.com/Events/Build/2016/P470>
- Accelerating SQL Server 2016 using Pmem - <https://channel9.msdn.com/Shows/Data-Exposed/SQL-Server-2016-and-Windows-Server-2016-SCM--FAST>

## Other Resources

- SNIA Persistent Memory Summit 2018 - <https://www.snia.org/pm-summit>
- Intel manageability tools for Pmem - <https://01.org/ixpdimm-sw/>




# INTEL® PARALLEL STUDIO XE

CREATE FASTER CODE...FASTER



COMPOSER EDITION	PROFESSIONAL EDITION	CLUSTER EDITION
<b>BUILD</b> Compilers & Libraries	<b>ANALYZE</b> Analysis Tools	<b>SCALE</b> Cluster Tools
<b>C / C++ Compiler</b> Optimizing Compiler	<b>Intel® VTune™ Amplifier</b> Performance Profiler	<b>Intel® MPI Library</b> Message Passing Interface Library
<b>Fortran Compiler</b> Optimizing Compiler	<b>Intel® Inspector</b> Memory & Thread Debugger	<b>Intel® Trace Analyzer &amp; Collector</b> MPI Tuning & Analysis
<b>Intel® TBB</b> C++ Threading Library	<b>Intel® Advisor</b> Vectorization Optimization & Thread Prototyping	<b>Intel® Cluster Checker</b> Cluster Diagnostic Expert System
<b>Intel® MKL</b> Fast Math Kernel Library		
<b>Intel® IPP</b> Image, Signal & Data Processing		
<b>Intel® DAAL</b> Data Analytics Library		
<b>Intel® Distribution for Python*</b> High Performance Scripting		
Intel® Architecture Platforms		
Operating System: Windows*, Linux*, MacOS <sup>1</sup> *		



**This suite is already installed in the target machine.**  
**Check /apps/INTEL and module avail in MN4 for further information.**

**More Power for Your Code - [software.intel.com/intel-parallel-studio-xe](https://software.intel.com/intel-parallel-studio-xe)**

