CAN YOUR APPLICATION BENEFIT FROM INTEL® OPTANE™ PERSISTENT MEMORY?

Vineet Singh

Software Engineer
Intel Corporation

Storage Performance Development Kit (SPDK)
Persistent Memory Development Kit (PMDK)
Intel® VTune™ Profiler

Virtual Forum

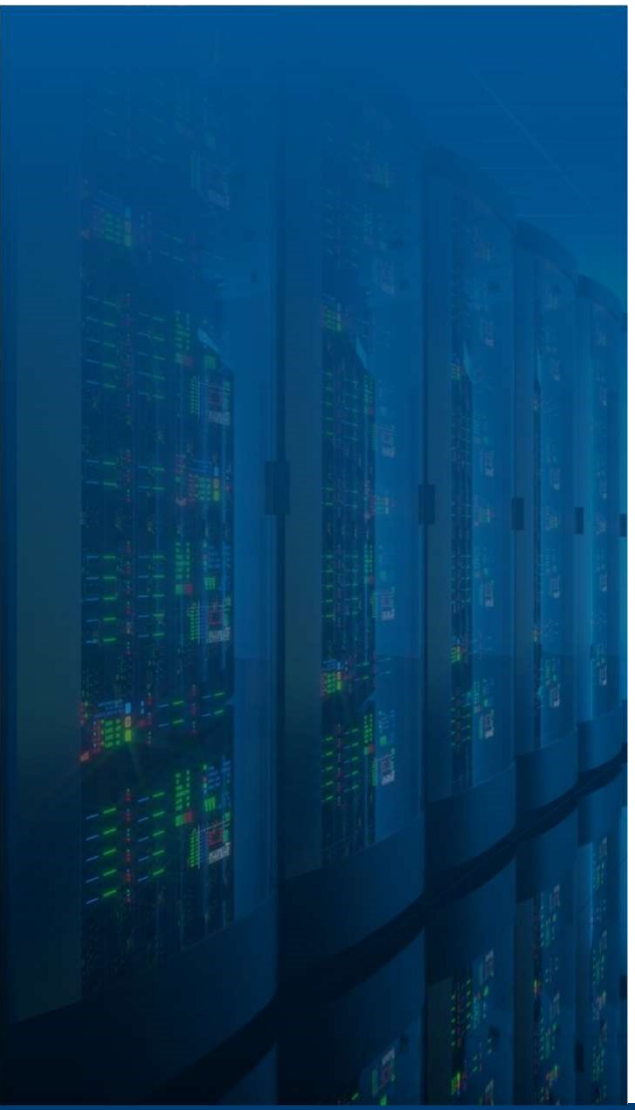

AGENDA

1. Intel® Optane™ persistent memory
   Introduction, benefits and tradeoffs.

2. What do you need to know about your application?
   Metrics and Tools

3. Where should you use Memory Mode?
   Analyze your application on DRAM-only system and identify if Memory Mode is useful for you

4. Where should you use App Direct mode?
   Analyze your application on DRAM-only system and identify if App Direct mode is useful for you

Storage Performance Development Kit (SPDK)
Persistent Memory Development Kit (PMDK)
Intel® VTune™ Profiler

Virtual Forum

Intel® Optane™ Persistent Memory

Introduction, modes, benefits and tradeoffs

INTEL® OPTANE™ PERSISTENT MEMORY

- Sits on the same slot as DRAM memory modules (DIMMs)
- Up to 6 persistent memory modules per socket
- Modes: Memory Mode and App Direct Mode

Why should you use PMem

- Cheaper than DDR
- Higher capacity
  - 128GB, 256GB and 512GB DIMMs
- Persistent or volatile
- Orders of magnitude lower latency than SSD

Things you need to understand before using PMem

- Performance varies based on traffic pattern
  - Sequential vs random
  - Read vs write
- Performance slightly lower than DDR (2x-5x latency)
- DRAM DIMMs share the same bus with persistent memory modules

Storage Performance Development Kit (SPDK)
Persistent Memory Development Kit (PMDK)
Intel® VTune™ Profiler

Virtual Forum

What do you need to know about your application?

Metrics and tools

UNDERSTANDING YOUR APPLICATION: RELEVANT METRICS

Application
- Data structures and allocation
- File I/O patterns
- Working Set size/Active memory/Hot data size

System Configuration
- CPU
- Memory
- Storage device
- Network device

Memory
- Utilization
- Bandwidth
- Read/Write ratio

CPU
- Utilization
- Kernel Utilization
- I/O wait
- Intel Top-down Microarchitecture Analysis (TMA) metrics

Storage
- IOPS
- Throughput
- Queue depth

Network
- Throughput

UNDERSTANDING YOUR APPLICATION: TOOLS

Intel® VTune™ Profiler - Platform Profiler

- Provides an overview of the system
- Covers main subsystems:
  CPU, Memory, Storage, Network
- Analyzes PMem performance

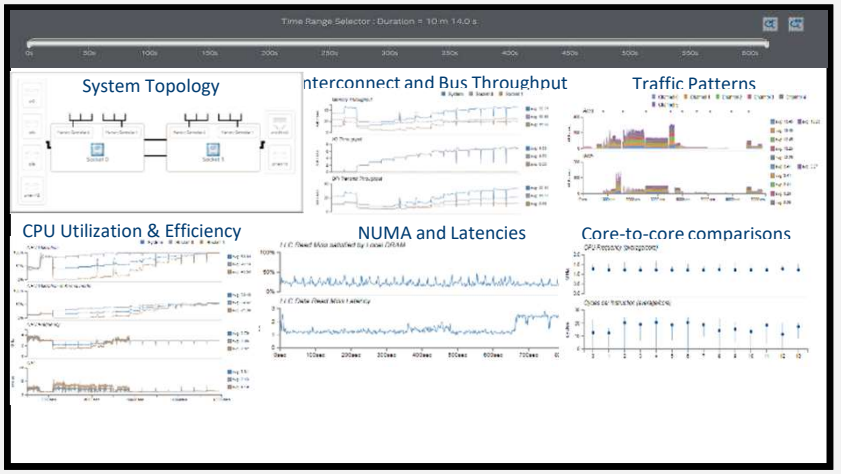

Intel® VTune™ Profiler - Hotspots Analysis

- Identifies where your application spends time

Intel® VTune™ Profiler - Microarchitecture Exploration Analysis

Intel® VTune™ Profiler - Memory Analysis

- Analyzes how efficiently your code is using the underlying hardware

Storage Performance Development Kit (SPDK)
Persistent Memory Development Kit (PMDK)
Intel® VTune™ Profiler

Virtual Forum

Where should you use Memory Mode?

- Understanding Memory Mode
- Understanding where Memory Mode can help

PMEM MEMORY MODE

- Does not require code change

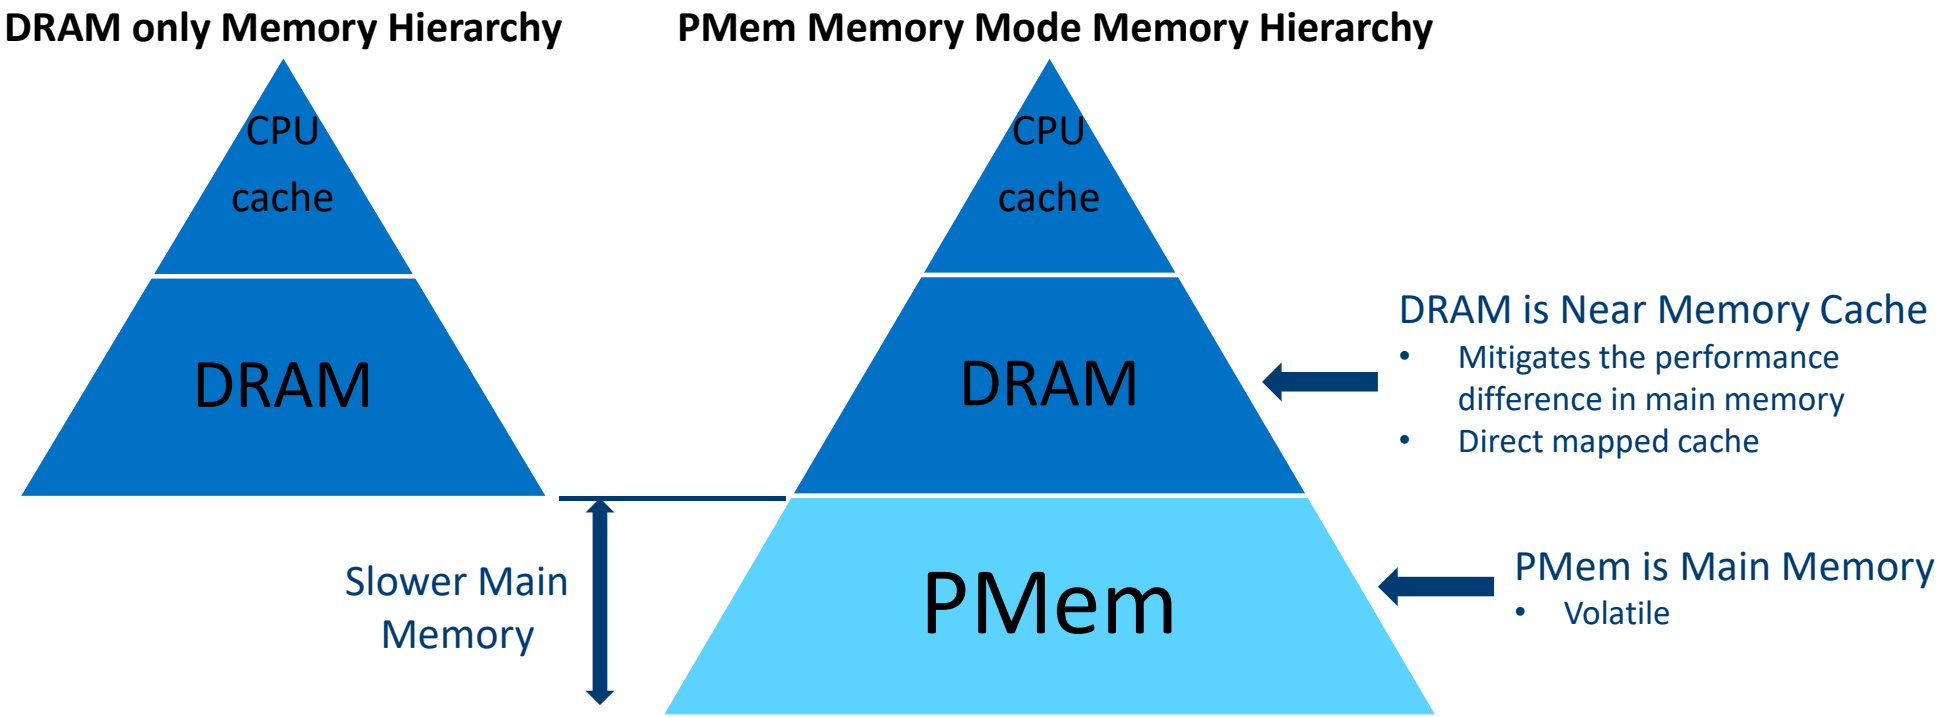

SCENARIO BLUEPRINT

Overview
- High level overview of scenario where PMem will help you

Potential Value Proposition
- What is the value for using PMem in this scenario

Required Metrics
- What information do you need about your application behavior on a DRAM-only system to identify this scenario

Evaluation
- Steps you can use to identify the scenario using performance information from a DRAM-only system

MEMORY CAPACITY BOUND - SINGLE NODE

Overview

- Your workload performance is taking a hit due to limited size of addressable memory

Potential Value Proposition

- Increase performance of a single node

Required Metrics

- System Configuration: Total Memory Size
- Memory utilization
- CPU utilization
- CPU utilization in Kernel Mode
- Major page faults
- Storage throughput/IOPS
- Maximum achievable DRAM bandwidth
- Actual DRAM throughput (read and write)
- CPU I/O Wait

Evaluation

- Total memory utilization(capacity) close to 100%
- Major page faults.
- High CPU utilization in Kernel Mode (>10%) or disk I/O activity matching 100% memory utilization.
- Low CPU utilization matching 100% memory utilization
- DRAM bandwidth not saturated

MEMORY CAPACITY BOUND

Memory Utilization ~ 100%

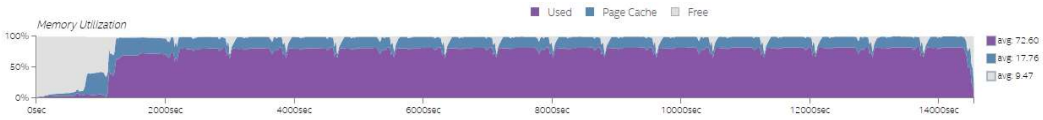

High Disk Activity matching memory utilization

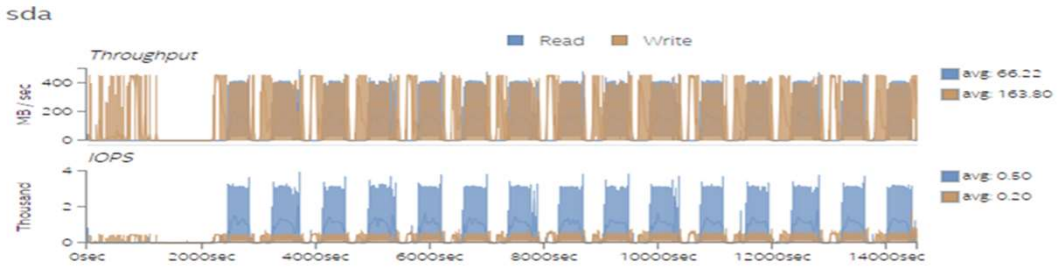

Low CPU Utilization matching 100% memory utilization

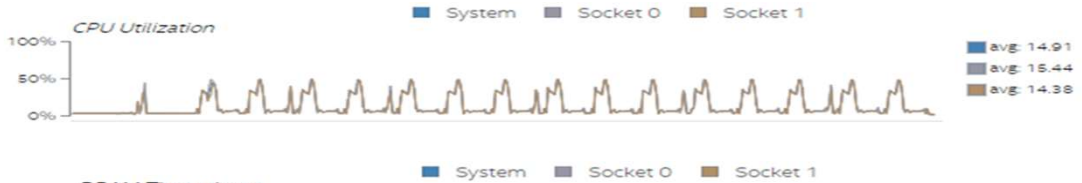

Memory Bandwidth not saturated

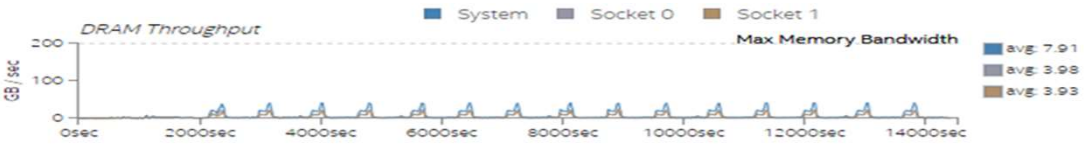

DRAM DISPLACEMENT

Overview

- Your workload performance is bound by system components other than the memory bandwidth and latency

Potential Value Proposition

- Decreasing cost of a single node

Required Metrics

- System Configuration: Total size of available memory
- System Configuration: Memory topology and dim placement
- System Configuration: Size and type of DRAM memory modules (DIMMs)
- Working set size of the workload (size of hot/frequently accessed data)
- TMA metrics

Evaluation

Check “TMA Memory_Bound(%)”

AND check “write bandwidth ratio (DRAM write BW/DRAM total BW)”

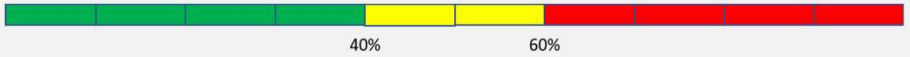

If previous “Memory_Bound” AND “write BW ratio” check fails
Check “working set size”
Working set size < planned near memory cache size (size of DRAM on the planned PMem memory mode configuration)

DRAM DISPLACEMENT

Low TMA Memory Bound % ~20%

Read Intensive workload

NODE CONSOLIDATION

Overview

- Your workload is distrusted across multiple worker nodes, number of nodes required for the workload depends on the total data and the size of memory available per node

Potential Value Proposition

- Reduce total cost of ownership

Required Metrics

- System Configuration: CPU, Memory, Storage and Network
- Total in-memory data size across all nodes
- Memory utilization
- Intel Top Down metrics
- CPU utilization
- Storage throughput, IOPS and queue depth
- Actual DRAM throughput (read and write)
- Network utilization

Evaluation

- Available memory per node is limiting the size/chunk of work assigned to the node
- Check for memory BW/latency sensitivity
- Check aggregate CPU requirement
- Check aggregate Storage requirement
- Check aggregate Network requirement
- Check SLA requirements.

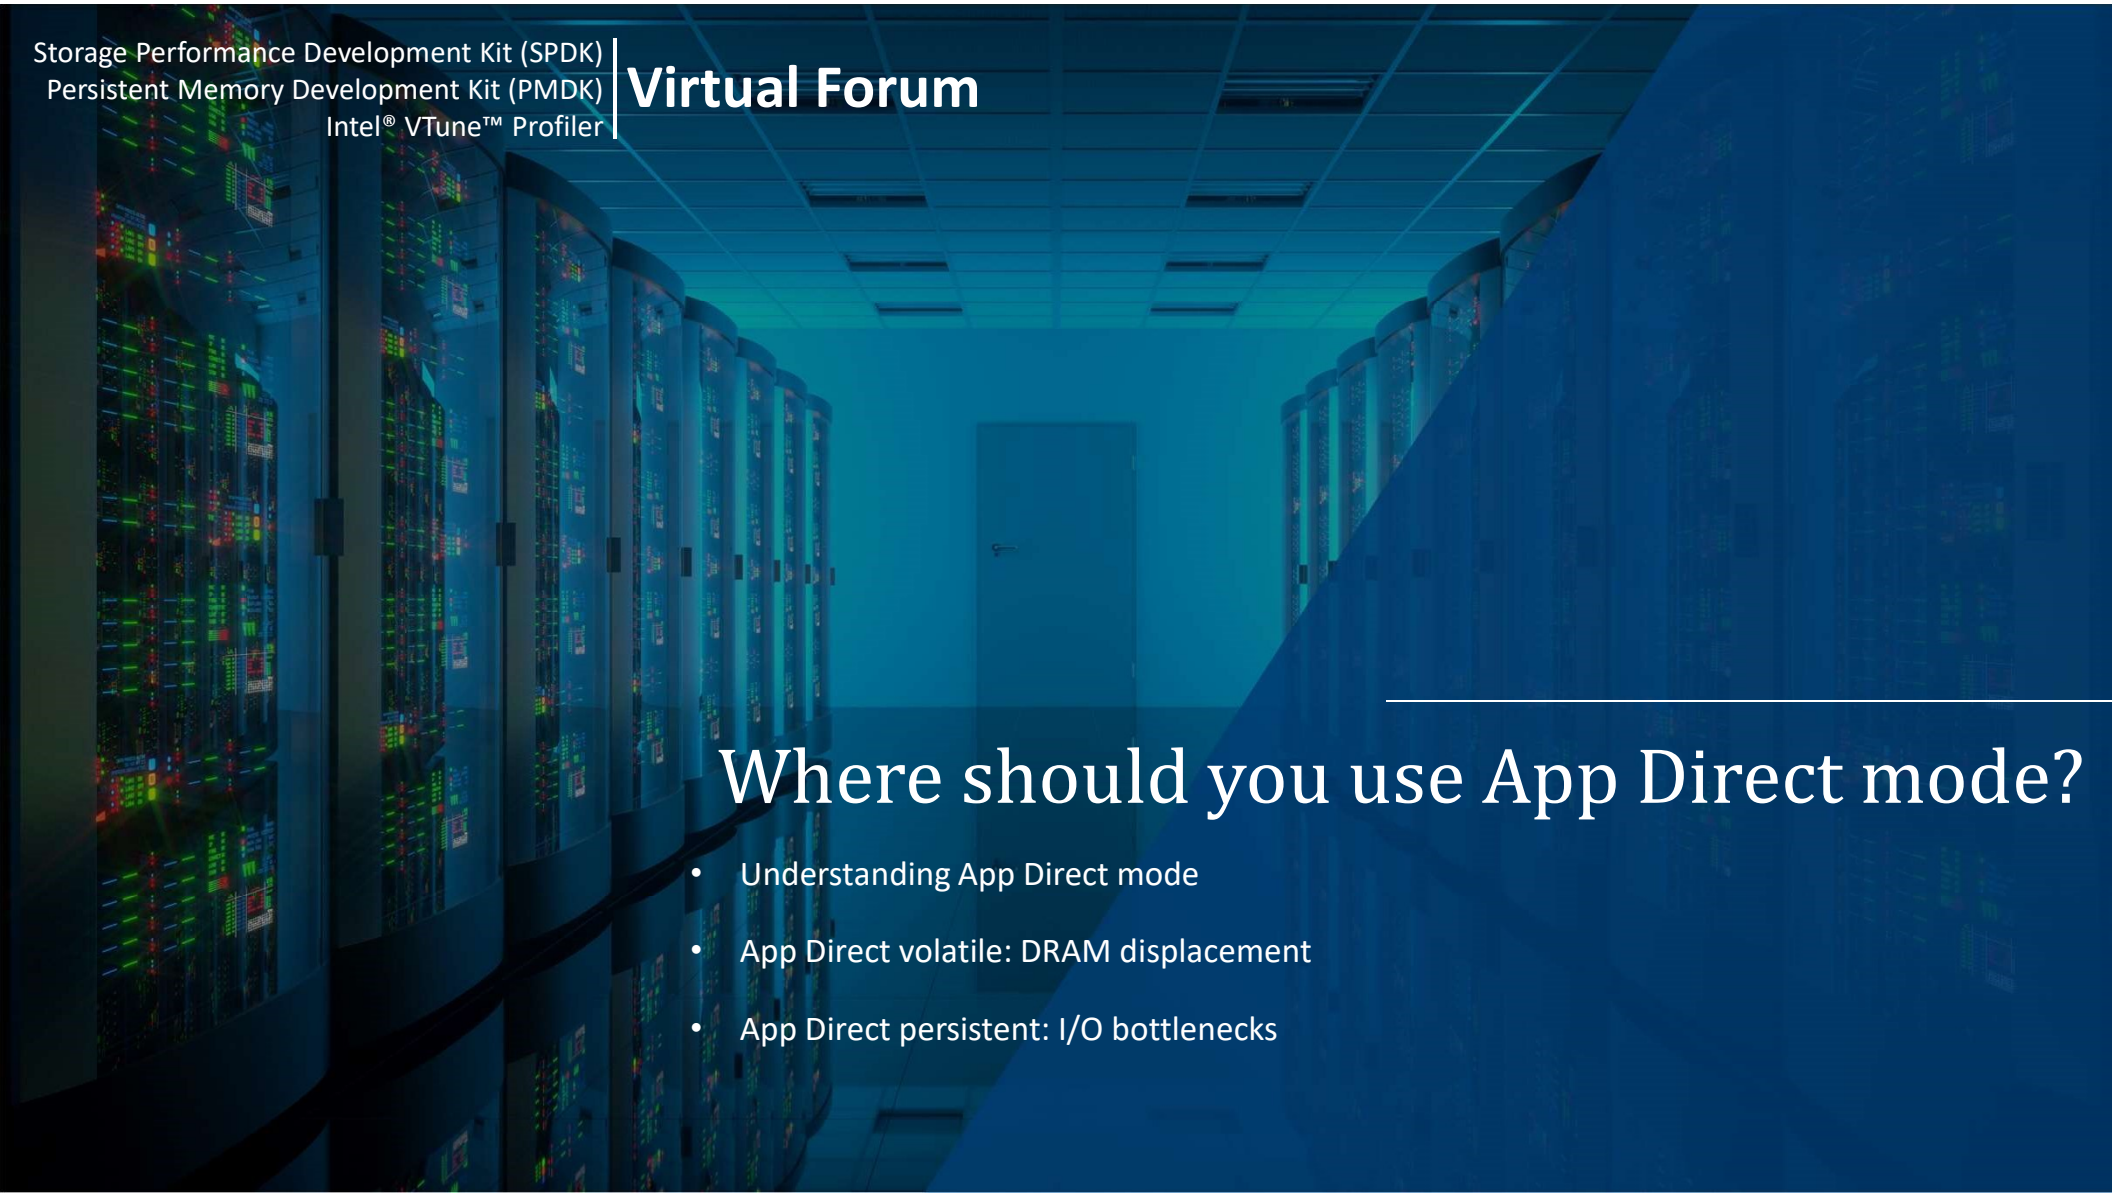

Storage Performance Development Kit (SPDK)
Persistent Memory Development Kit (PMDK)
Intel® VTune™ Profiler

Virtual Forum

Where should you use App Direct mode?

- Understanding App Direct mode
- App Direct volatile: DRAM displacement
- App Direct persistent: I/O bottlenecks

PMEM APP DIRECT MODE

- Requires code change
- Programmer chooses whether to allocate data in DRAM or PMem
- Use PMDK

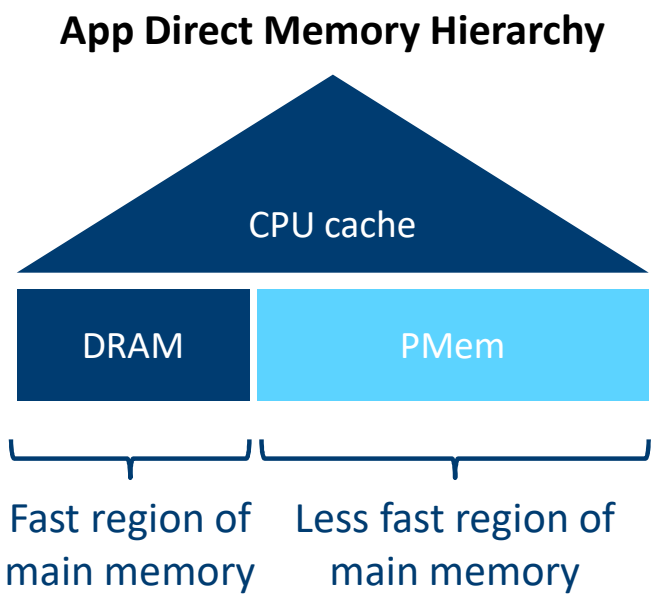

- PMem is an additional resource that the software has to manage
- Can be volatile or persistent

VOLATILE USE CASE: DISPLACE DRAM

Overview

- Your in-memory data structures can be divided into hot data structures and warm data structures. Moving the warm data structures to PMem enables you to save on DRAM cost

Potential Value Proposition

- Decreasing cost of a single node

Required Metrics

- System configuration: memory size
- Memory utilization
- Per memory object access characteristics

Evaluation

- Memory utilization > 50%
- Two categories of memory objects, hot and warm
- Size of hot memory objects < size of DRAM on the planned PMem App Direct configuration

PERSISTENT USE CASE: I/O BOTTLENECKS

Overview

- Your workload is I/O intensive. Moving some of the frequently use data to PMem will improve overall performance

Potential Value Proposition

- Better performance

Required Metrics

- System configuration: CPU, memory, storage
- CPU utilization
- CPU I/O Wait
- Storage throughput, IOPS and queue depth
- Per file I/O traffic

Evaluation

- CPU utilization is low
- CPU I/O Wait is significant
- Frequently accessed files fit on the proposed PMem capacity

AND MORE…

This is not an exhaustive list of use cases

- **Time sharing applications**
- **New applications that can take advantage of large memory**
- **I/O Caching**
- **New Persistent data structures**

How people are using Intel® Optane™ persistent memory in real world

- **Intel® Optane™ persistent memory**
- **DAOS**
- **HPC Applications with PMem**
- **VMWare –Memory Mode**
- **In-memory Graph Analytics**

Resources

- Software tools for Intel® Optane™ persistent memory. Free downloads and technical articles software.intel.com/persistent-memory/tools
- Intel® VTune™ Profiler software.intel.com/vtune


Storage Performance Development Kit (SPDK)
Persistent Memory Development Kit (PMDK)
Intel® VTune™ Profiler
Virtual Forum