Leveraging New Database Persistent Memory Features in 21c & 19cR12 to Eliminate I/O Latency & Boost Performance

Kai Yu
ORACLE ACE Director
Distinguished Engineer
Dell Technologies, Inc

About Me: Kai Yu

DELL Technologies

- **Distinguished Engineer, Dell Technical Leadership Community**
- **Dell Database Solutions Engineering**
- **28+ years working in Tech Industry**
- **Specializing in Oracle Database, Virtualization/Cloud and Machine Learning**
- **Author and Frequent Speaker at IEEE and Oracle Conferences**
- **IOUG Cloud Computing SIG Co-founder and VP**
- **Oracle ACE Director since 2010**
- **Co-recipient of the OAUG Innovator of Year Award**
- **Oracle Technologist of the Year: Cloud Architect by Oracle Magazine**
- **My Blog: http://kyuoracleblog.wordpress.com/**

Agenda

- Introduction to Persistent Memory (PMem)
- PMem Memory Mode for Oracle Database
- PMem App-Direct Mode for Oracle Database
  - Oracle PMem Filestore in Oracle 21c
  - Oracle Memory Speed File System in Oracle 19c
- POCs: PMem App-Direct for 19c and 21c Oracle databases
- Persistent Memory in Oracle Exadata X8M
- Intel's Latest Decision on PMem

Introduction to Persistent Memory

Persistent Memory (PMem):

- A solid-state high-performance byte-addressable memory device
- Co-developed by Intel and Micro based on 3D XPoint memory
- Resides on the memory bus. A new tier sits between memory and storage
- Deliver the best of both through the convergence of memory and storage

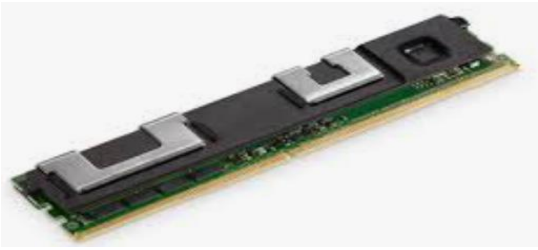

HDDs	1-10ms	1x
NAND SSD	10-100us	100x
Persistent Memory	< 1us	1000x
DRAM	80-100ns	10000x

- Provide large capacity and affordable memory: 128GiB, 256GiB, 512GiB
- Supports volatile or non-volatile, also serves as high performance storage
- Compatible with DDR4 socket and can co-exist with conventional DDR4 DRAM.
- Modules fit into standard DDR4 DIMMS slots

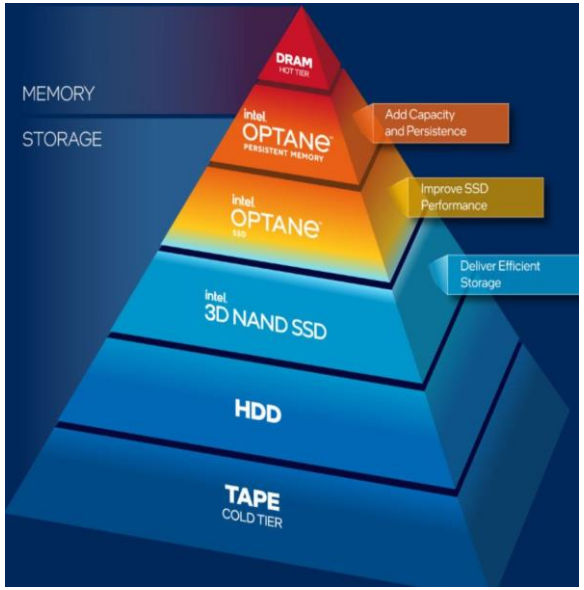

Introduction to Persistent Memory

Generations of Persistent Memory

First generation of PMem:

Product Name: Intel® Optane™ Persistent Memory 100 Series
Module Capacity: 128GB, 256GB, 512GB
Product code name: Apach Pass
CPU: 2nd Generation Intel® Xeon® Scalable processors,

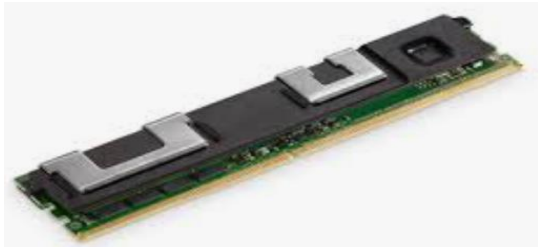

Second generation of PMem:

Product Name: Intel® Optane™ Persistent Memory 200 Series
Module Capacity: 128GB, 256GB, 512GB
Product code name: Barlow Pass
CPU: 3rd Gen Intel® Xeon® Scalable processors

Third generation of PMem (when?)

Product Name: Intel® Optane™ Persistent Memory 300 Series
Product code name: Crow Pass
CPU: 4th Gen Intel Xeon Scalable processors (codenamed “Sapphire Rapids”).

Introduction to Persistent Memory

Example: Dell PowerEdge R750 CPU-to-memory subsystem:

- Processor: dual 3rd General Intel® Xeon® Scalable processors ,up to 2x40 cores
- 8 Memory channels per processor, two memory slots per channel
- PMem: Intel® Optane™ Persistent Memory 200 Series

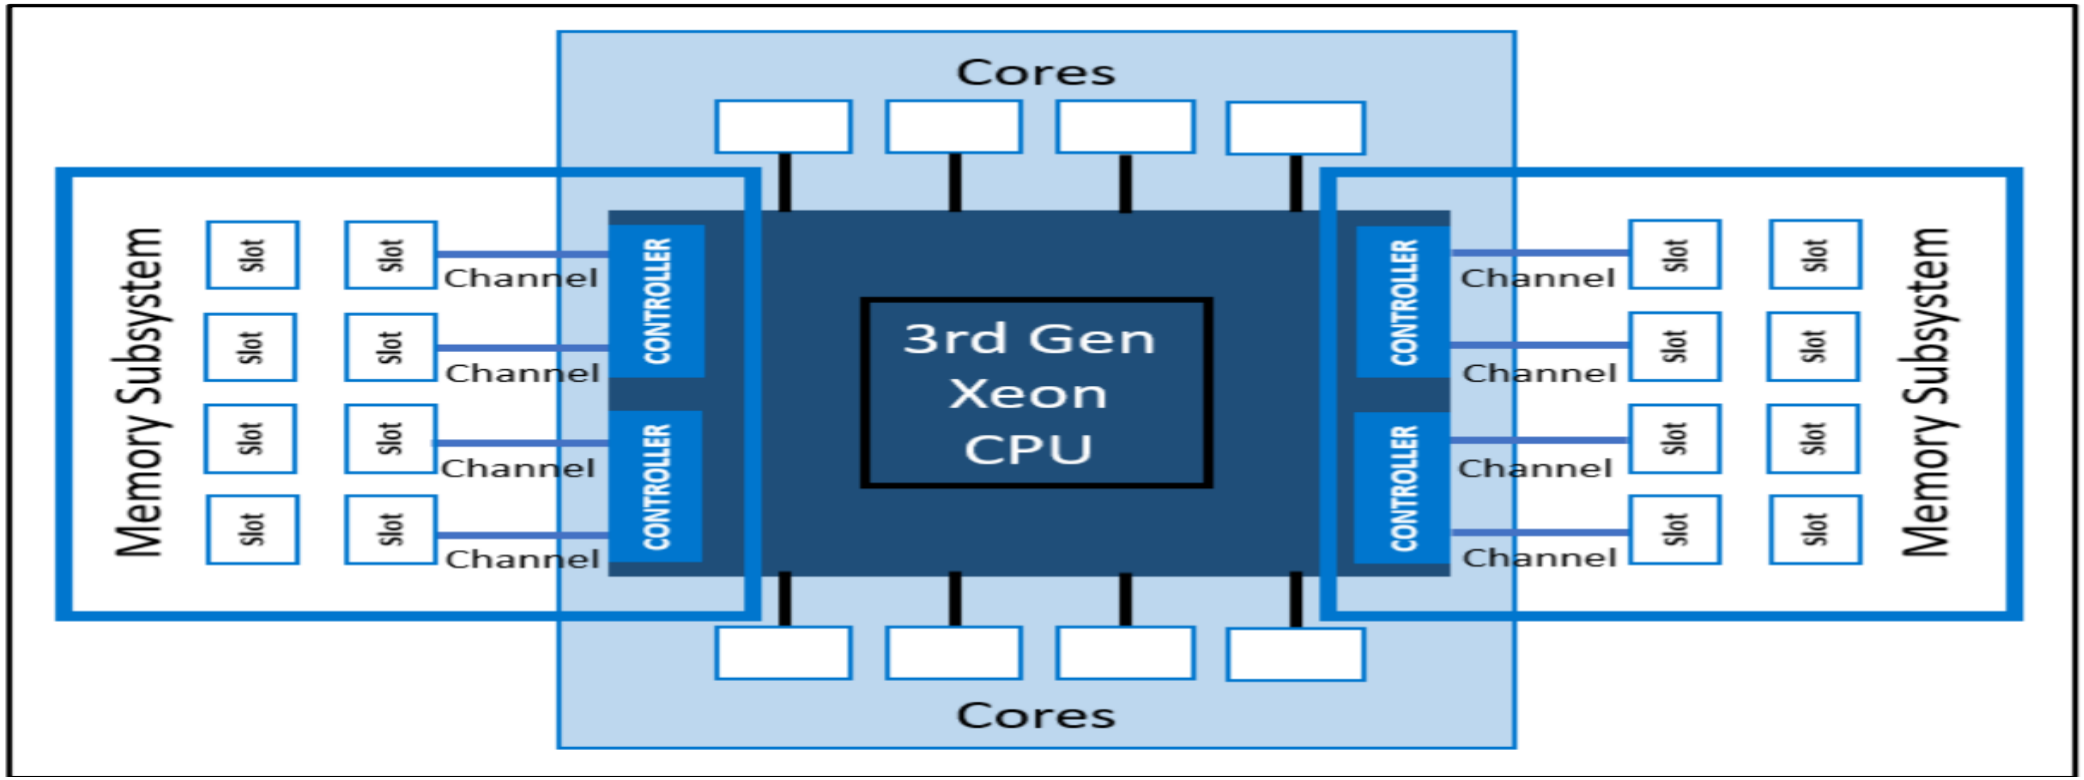

Introduction to Persistent Memory

- R750 with two sockets: Total 32 DDR4 DIMM slots A1-A16, B1-B16
- Support 6 different configurations of mixing DRAMs and PMems
- For example, 16(A1-A8, B1-B8) for DRAMs, 16(A9-A16, B9-B16) for PMems,
  Total 16 Intel Persistent Memory 200 series (BPS) slots, 16x512GB=8 TB max

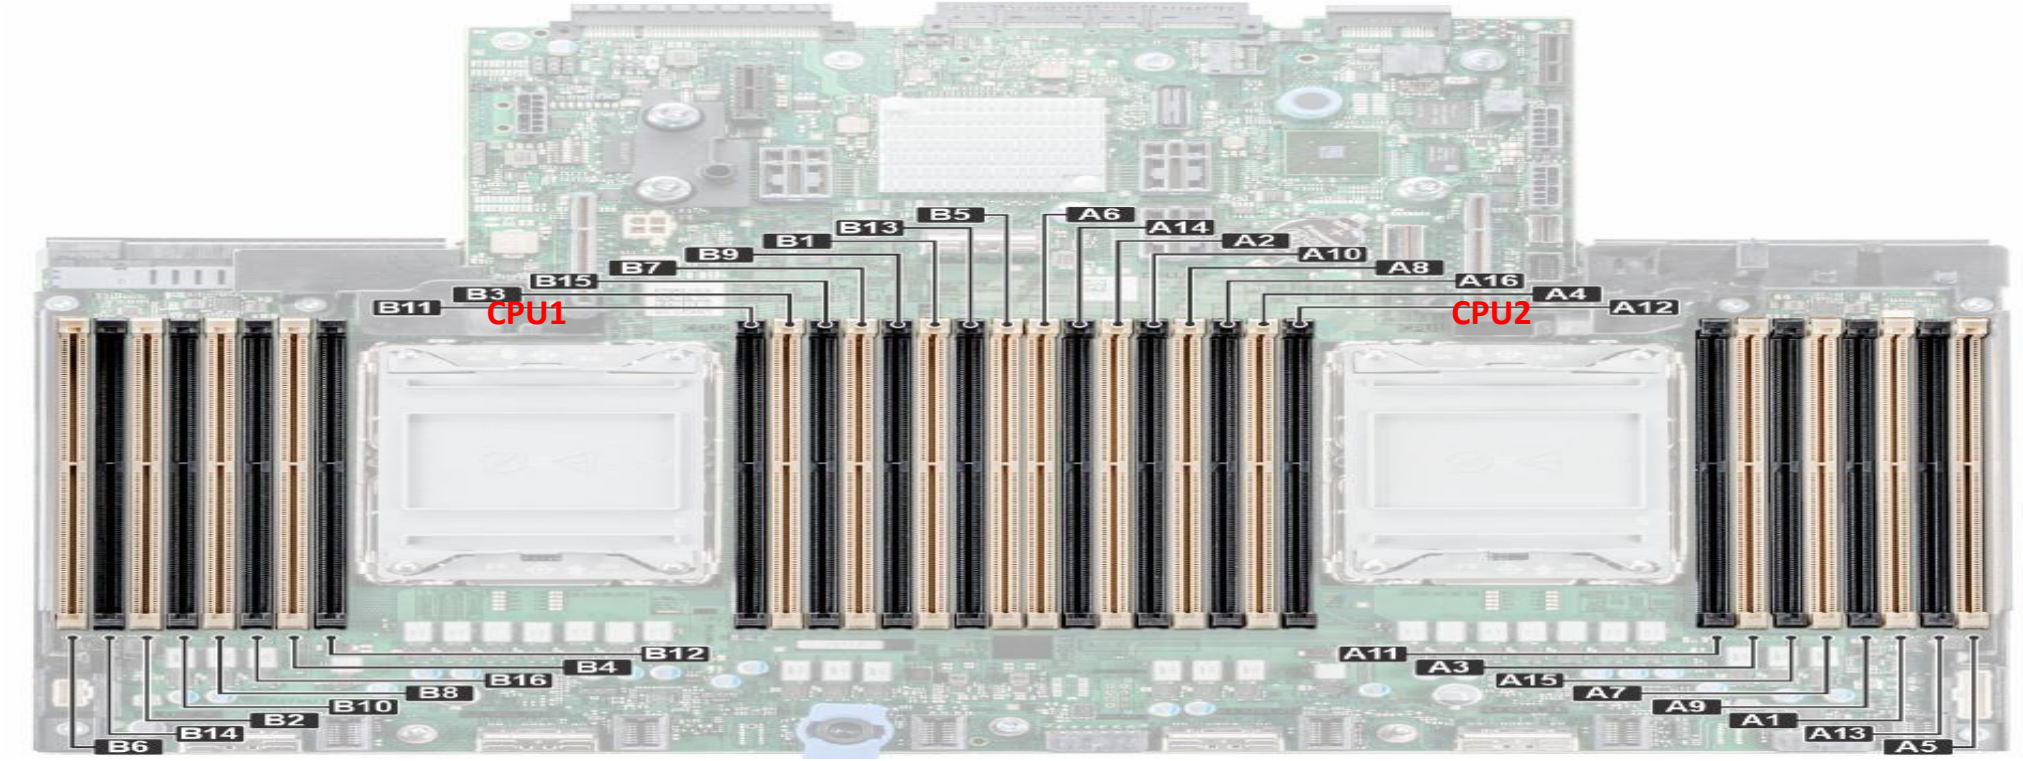

Introduction to Persistent Memory

Persistent Memory operation modes:

- Memory Mode
  - Data stored in PMems is volatile
  - Extends the amount of available volatile memory to OS
  - DRAM used as Cache for PMems
  - PMems used as addressable main memory
- App-Direct Mode
  - PMem in this mode is persistent as storage, act like an SSD
  - Bytes addressable like memory
- Dual Mode (not officially supported)
  - A subset of PMems is configured as Memory mode
  - The remaining is in App-Direct Mode.

Introduction to Persistent Memory

Multi-tiered memory and storage hierarchies for optimal workload performance

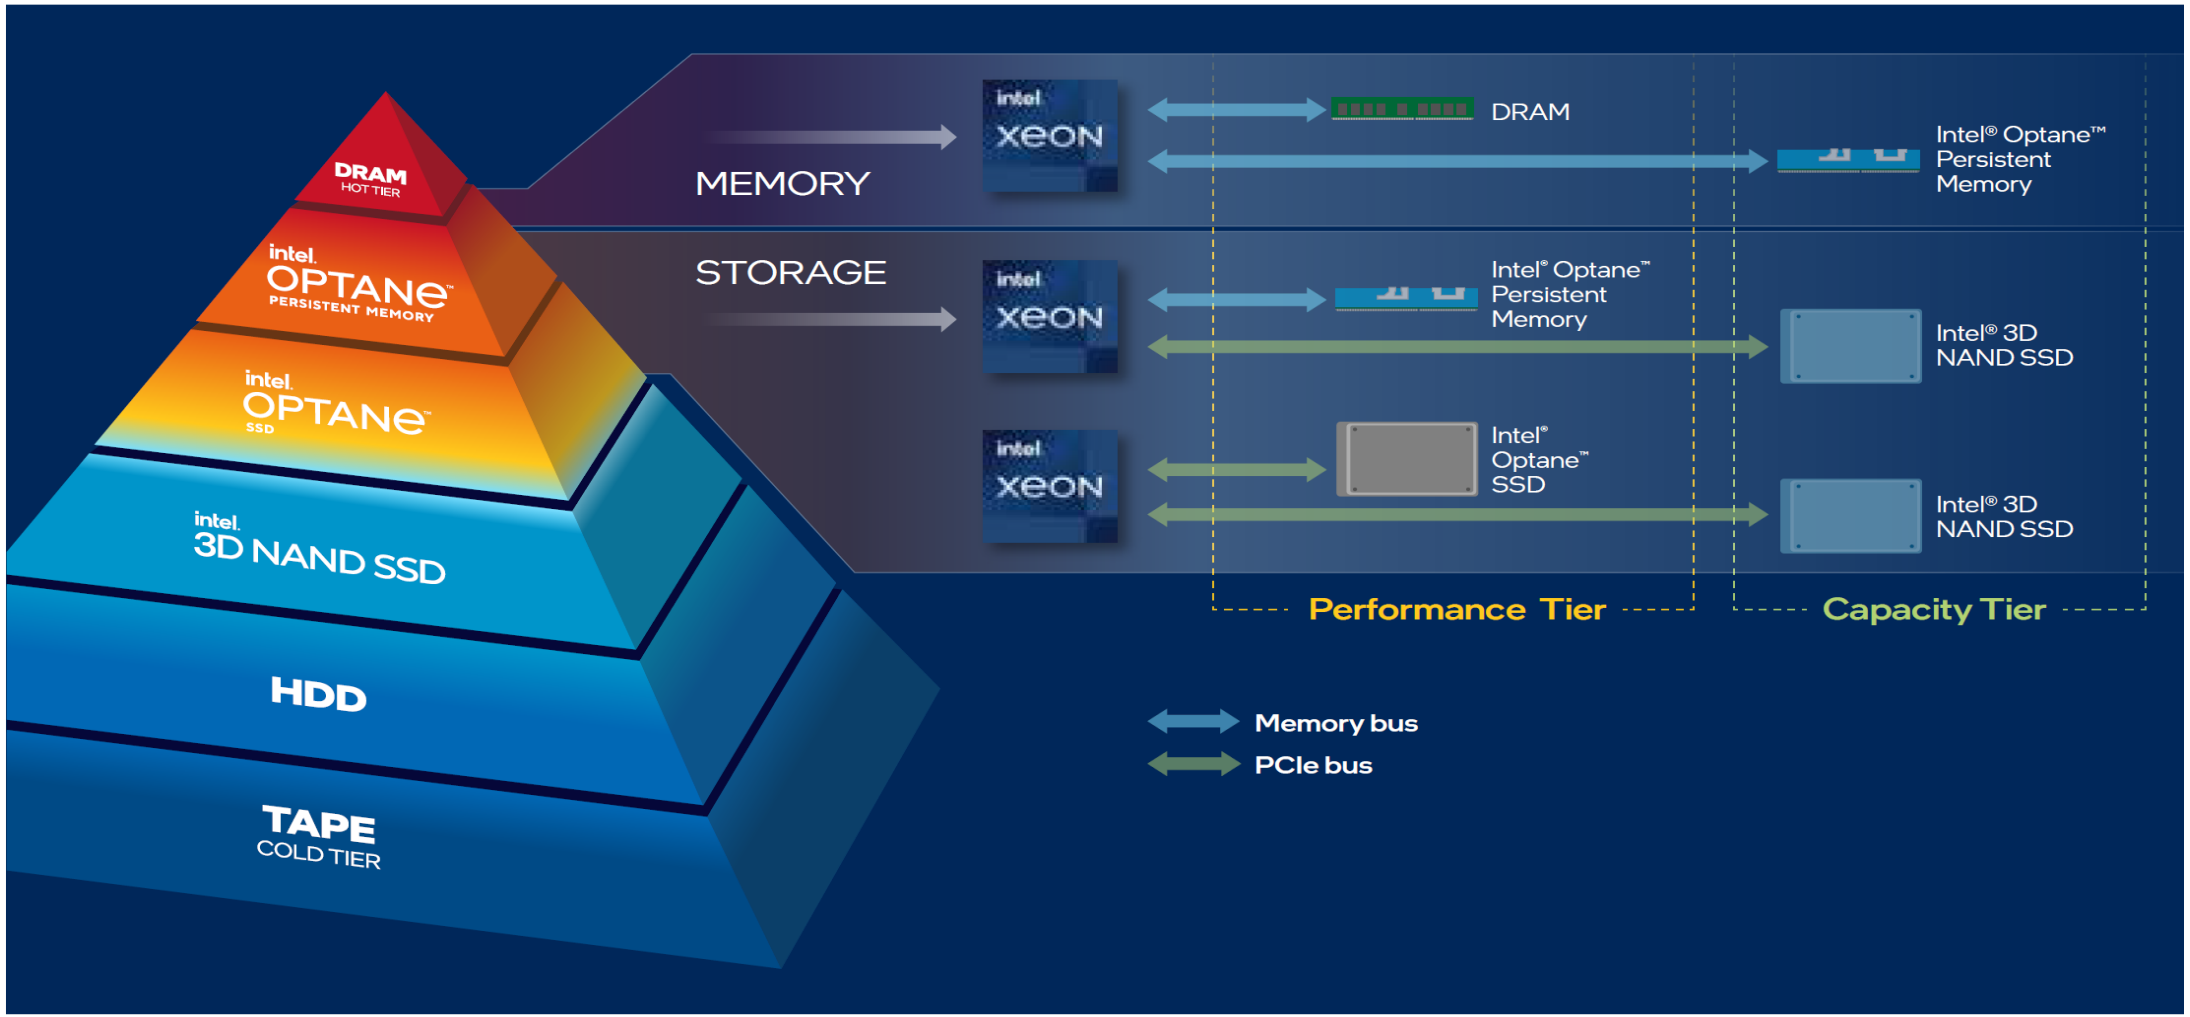

Introduction to Persistent Memory

Some Performance comparison study: (https://code.kx.com/q/architecture/optane-tests/
Query response times: similar to DRAM
3.8x to 12x faster than NVMe

	DRAM RDB 2 query processes	PMem IDB 8 query processes	NVMe HDB 8 query processes	PMem vs. DRAM[5]	PMem vs. NVMe
			1 query at a time		
Mean response time (ms)	23	26	319	1.17	12.10
Mean Payload Size (KB)	778	778	668	1	1
			100 queries at a time		
Mean response Time (ms)	100	82	310	0.82	3.77
Mean Payload Size (KB)	440	440	525	1	1

Introduction to Persistent Memory

- **Configure Persistent Memory in Server: Configured on server's BIOS and OS**

Intel Persistent Memory

	Current Value
Raw capacity	4.0 TB
App Direct capacity	4.0 TB
Memory capacity	0 B
Unconfigured capacity	0 B

```
[root@r75019c ~]# ipmctl show -memoryresources
MemoryType   | DDR         | PMemModule   | Total
==========================================================
Volatile     | 1024.000 GiB | 0.000 GiB    | 1024.000 GiB
AppDirect    | -           | 4048.000 GiB | 4048.000 GiB
Cache        | 0.000 GiB    | -            | 0.000 GiB
Inaccessible | 0.000 GiB    | 11.748 GiB   | 11.748 GiB
Physical     | 1024.000 GiB | 4059.748 GiB | 5083.748 G
```

region(s): interleaved set(s) of PMEMs

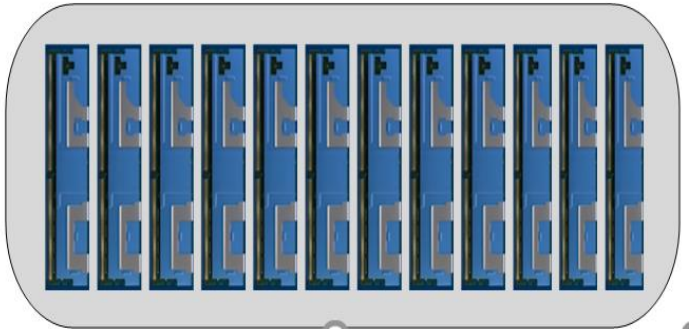

```
[root@r75019c ~]# ndctl list --regions --human
[
 {
  "dev":"region1",
  "size":"2024.00 GiB (2173.25 GB)",
  "available_size":0,
  "max_available_extent":0,
  "type":"pmem",
  "iset_id":"0x9c1b2120c77e1110",
  "persistence_domain":"memory_controller"
 },
```

```
{
  "dev":"region0",
  "size":"2024.00 GiB (2173.25 GB)",
  "available_size":0,
  "max_available_extent":0,
  "type":"pmem",
  "iset_id":"0xf8932120258e1110",
  "persistence_domain":"memory_controller"
 }
```

```
[root@r75019c ~]# more /proc/partitions | grep pmem
major minor #blocks  name

 259        4 2089154560 pmem0
 259        5 2089154560 pmem1
```

PMem Memory Mode for Oracle Database

- How PMem Memory Mode works
  If there is a cache-hit, data or instructions will be fetched from DRAM ➔ DRAM-like performance
  Otherwise, data will be fetched from PMems slightly longer latency

- Recommended 1:4 ratio of DRAM:PMem

- Benefits of the PMem memory mode:
  Deliver DRAM-like performance
  Larger memory capacity
  Affordable cost (PMem less expensive than DRAM)
  Application-agnostic, no application change need

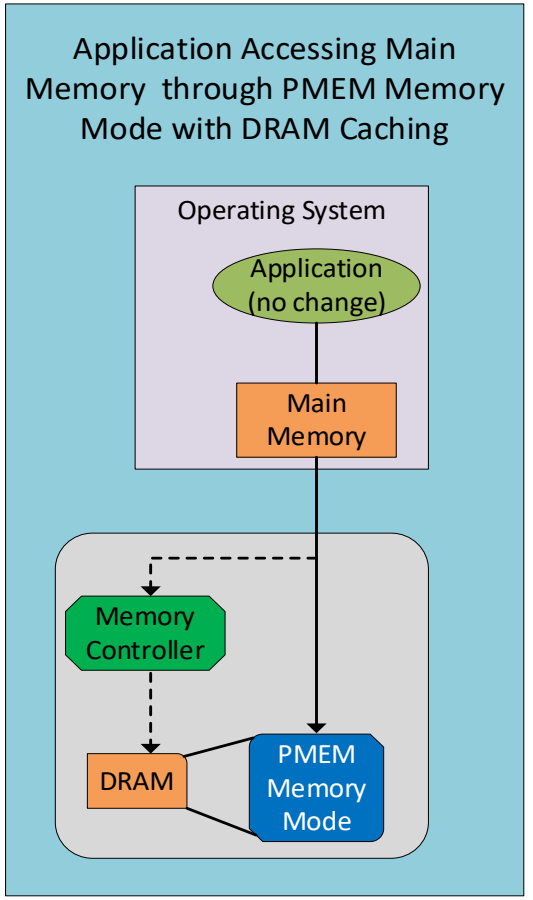

PMem Memory Mode for Oracle Database

- **Expand Oracle database memory beyond DRAM to reduce physical IOs**

Expand SGA or specifically the buffer cache size
. Improve the cache hit ratio with
more logical reads from buffer cache
. Caching tables and indexes :

```
sql> alter system set db_keep_cache_size=100G;
sql> create table sales tablespace data
       storage (buffer_pool keep) ;
sql> create index sales on sales(id) tablespace
     data storage (buffer_pool keep);
alter table products storage (buffer_pool keep);
```

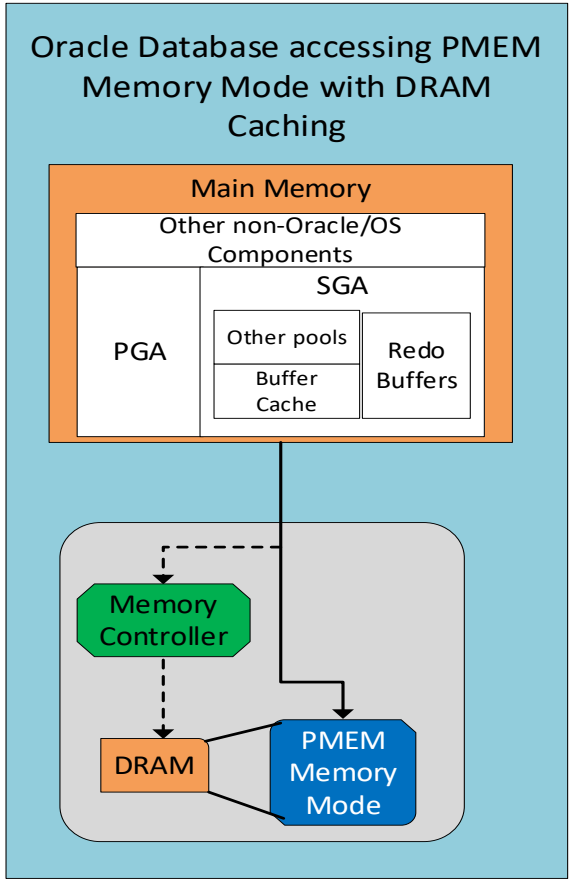

PMem Memory mode for Oracle Database

Full database In-memory caching: caching the entire database in memory

sql> alter database force full database caching;

- PMem in Memory Mode expands the database memory cache
  - Increasing the size of database memory cache with a lower cost
  - Larger database memory cache → reduce # of db file reads
  - less db file read → improve the read latency

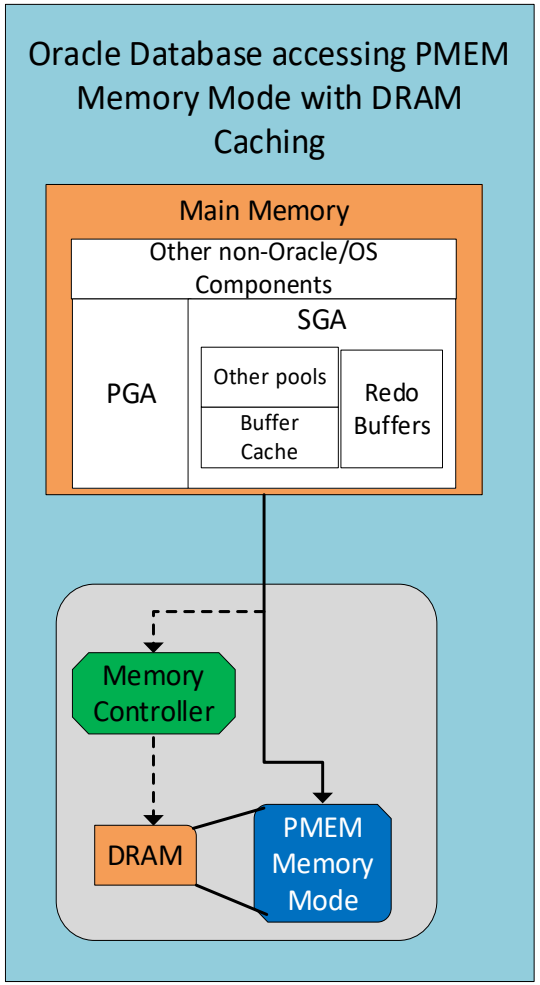

Top 10 Foreground Events by Total Wait Time

Event	Waits	Total Wait Time (sec)	Avg Wait	% DB time	Wait Class
db file sequential read	33,746,785	167.9K	4.98ms	71.7	User I/O
log file sync	4,890,364	55.4K	11.33ms	23.7	Commit
DB CPU		9239.3		3.9	

PMem Memory Mode for Oracle Database

Leverage PMem Memory mode for Oracle In-Memory Columnar stores

- PMems for Oracle In-Memory Database Dual Memory architecture:
  A new component called In-Memory Area in SGA
  SQL> alter system set in_memory_size = 1000G scope=spfile;

- Overcome the memory size limitation posted by DRAM size and DRAM cost

- Help Analytical processing through reading data from the In-memory column store

- Help OLTP by allowing you drop indexes that were created for reporting

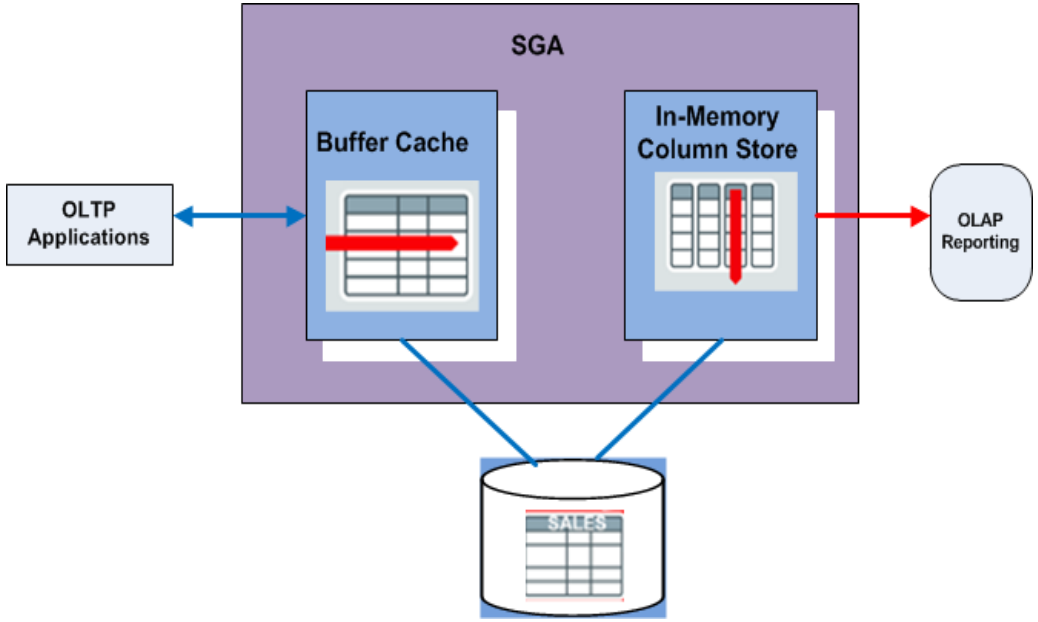

Oracle In Memory Database for OLAP workloads.
Refer to Oracle solution team paper:
Refer to Dell Oracle solution team paper:
Design Guide—Oracle Database In-Memory Use Cases with Intel Optane Persistent Memory (PMem)

PMem Memory Mode for Oracle Database

- Performance studies with PMems for ICMS done in Dell Oracle solution team's Whitepaper
  - HammerDB TPC-H like test:
    a single user , 22 sequential OLAP queries, parallel degree of 56

. 1 TB Database with 768 GB DRAM vs 1.5 TB PMem

Test1: Capacity High In-Memory compression : Compare performances between DRAM vs PMem used as IMCS.

Test Config	Memory config	Compression	IMCS size	Query Run time(relative)
IMCS on DRAM	768 GB DRAM	Capacity High	300GB	100: baseline
IMCS on PMem	1.5 TB PMems	Capacity High	300GB	126

Test2: Query Low In-Memory compression: DRAM is not big enough to fit IMCS

Test Config	Memory config	Compression	IMCS size	Query Run time
IMCS on DRAM	768 GB DRAM	Query Low		578, Data spill to disk
IMCS on PMem	1.5 TB PMems	Query Low	960 GB	132

PMem Memory Mode for Oracle Database

- 3 TB database with 1.5 TB DRAM vs 3 TB PMem + 768 GB DRAM

Test1: Use Capacity High In-Memory compression

Test DB size	Memory config	Compression	IMCS size needed	Query Run time
3 TB	1.5 TB DRAM	Capacity High	920GB	100
3 TB	3 TB PMems	Capacity High	920GB	103

Test2: Use Capacity Low In-Memory compression

Test DB size	1. Memory config	Compression	IMCS size needed	Query Run time
3 TB	1.5 TB DRAM	Capacity Low	1200GB	196, 80% data in IMCS
3 TB	3 TB PMems	Capacity Low	1200 GB	102

Test3: Use Query High In-Memory compression

Test DB size	Memory config	Compression	IMCS size needed	Query Run time
3 TB	1.5 TB DRAM	Query High	1780GB	203, 69% data in IMCS
3 TB	3 TB PMems	Query High	1780 GB	88

PMem Memory Mode for Oracle Database

How to leverage PMem to improve Oracle database performance:

- Memory Mode for database memory cache
  - Scope: Reduce physical reads by reading data from cache
  - Benefits: Increasing the size of Oracle memory with a lower cost
    Greatly improving read performance: more for OLAP workloads
  - The performance issues not being addressed
    How to reducing the high IO latency on Data writes and Redo Log writes
    =>Improving the IO bottlenecks for OLTP workloads
  - Is it possible to leverage PMem App-Direct mode to improve database I/O performance ?

Top 10 Foreground Events by Total Wait Time

Event	Waits	Total Wait Time (sec)	Avg Wait	% DB time	Wait Class
db file sequential read	33,746,785	167.9K	4.98ms	71.7	User I/O
log file sync	4,890,364	55.4K	11.33ms	23.7	Commit
DB CPU		9239.3		3.9	

PMem App-Direct Mode for Oracle Database

- **PMem in App-Direct Mode as storage devices to improve database I/O performances**
  - Use Case: Use PMems as the persistent storage to store data, redo logs, temp files.
  - Benefits:
    - Improve R/W performance of data and redo logs: I/O latency, IOPs
    - these I/Os have been the major database performance bottlenecks
  - Scope: Improve I/O performance for **all types of databases** including OLTP database
  - Database systems needs to add the support for storing database files in the byte-addressable PMem device

Persistent Memory App-Direct Mode

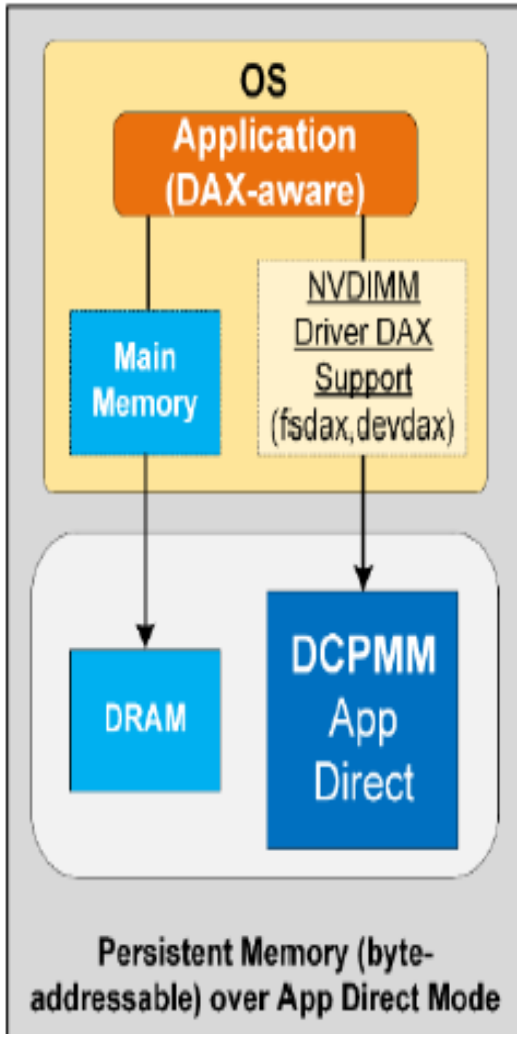

How PMem App-Direct Mode Works:

- PMems acts like persistent storage media, no data loss across power cycles.
- PMems are byte-addressable media, vs block device as SSDs/HDDs
- Direct Access (or DAX)
  - Directly access Persistent Memory contents without copying data into buffers in DRAM (bypass the DRAM)
  - By-passes the operating system's I/O stack
  - Important for performance
  - The Applications must be DAX-aware and leverage and support the direct-access feature,
  - The application must tolerate the level of atomicity provided by the Persistent Memory device

Persistent Memory App-Direct mode

- **Persistent Memory over App-Direct (block-addressable)**
  - Applications don't need to be modified to directly access PEMEM as a traditional persistent block storage using either 'raw' or 'sector' namespace.
  - **PMem configured as 'raw' :**
    It can guarantee a power-fail atomicity1 of only eight bytes
    PMem at the block level (512 bytes) to the risk of 'torn writes' which could result in potential data loss or inconsistency in the event of a power failure.
  - **PMem configured as in 'sector' :**
    It uses Block Translation Table (BTT) driver, which implements the copy-on-write optimization on top of the PMem driver. This helps to provide power-fail atomicity and lower the risk of data loss/inconsistency.
  - **Both of these are not suitable to store Oracle database files directly.**
- **Persistent Memory over App-Direct (byte-addressable)**
  - The PMem driver in operating system supports the direct-access (DAX) feature
  - The application code is modified to support the direct-access feature
  - Oracle Database pre-19c* older versions do not support the direct-access (DAX) feature

Persistent Memory for Oracle Database

- **PMem App-Direct mode support by Oracle database**:
  - Currently Oracle support for PMem in App-Direct mode for Oracle database
    - Oracle's Exadata X8M or later versions of Exadata
    - Oracle PMem Database/Filestore in Oracle 21c released on 8/13/2021
      - Available in Oracle 21c release
      - Support a single node Oracle database,
      - Include Directly Mapped Buffer Cache feature

    - Oracle Memory Speed File System in Oracle 19c database (release on July 28, 2021)

      Use Persistent Memory in App Direct mode to store Oracle database files and logs files
      - Available with Oracle database 19c Patch set 12 (19.12)
      - Support a single node Oracle database

  - An POC project was started right after these two features were available with collaboration with Oracle 19c development team and 21c development team

Oracle PMem Filestore in Oracle database 21c

Oracle 21c Pmem File store features

- Oracle Persistent Memory Filestore (PMem Filestore) can store data files, online redo log files, and control files.
- Directly Mapped Buffer Cache:
  - Traditional database buffer cache resides in DRAM. Oracle reads data from the storage to buffer cache first.
  - With PMem, Queries can read bytes *directly* from PMem without first migrating blocks into the database buffer cache.
  - This high-performance mechanism avoids data redundancy and unnecessary I/O.
- PMem Filestore is supported through Oracle 21c Persistent Memory feature
- Only works in a single instance Oracle database.
- PMem Filestore is the underlying file store used for a Persistent Memory database

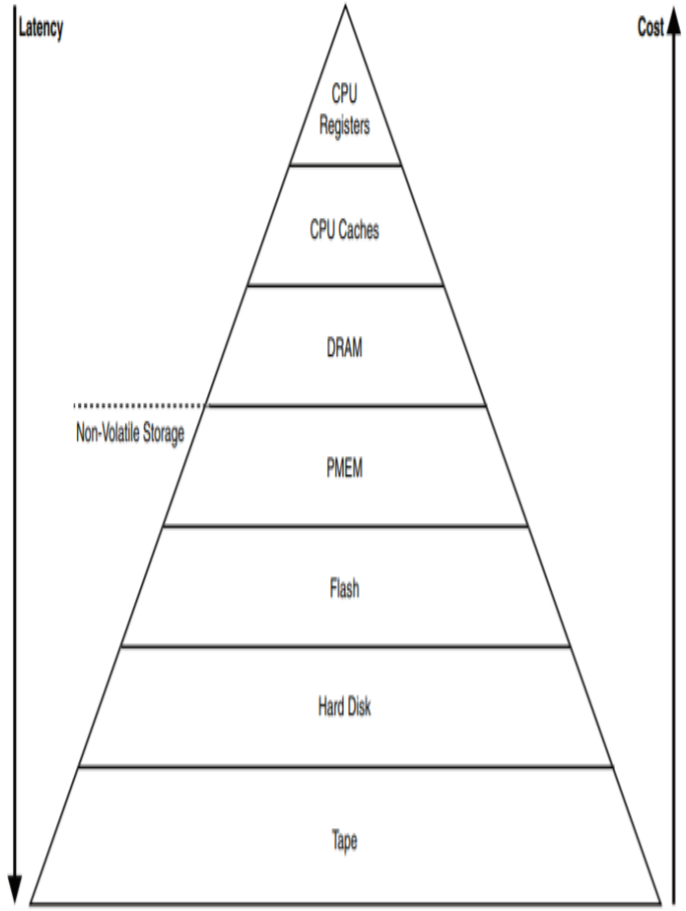

Oracle PMem Filestore in Oracle Database 21c

PMem Filestore Storage Architecture and Configuration

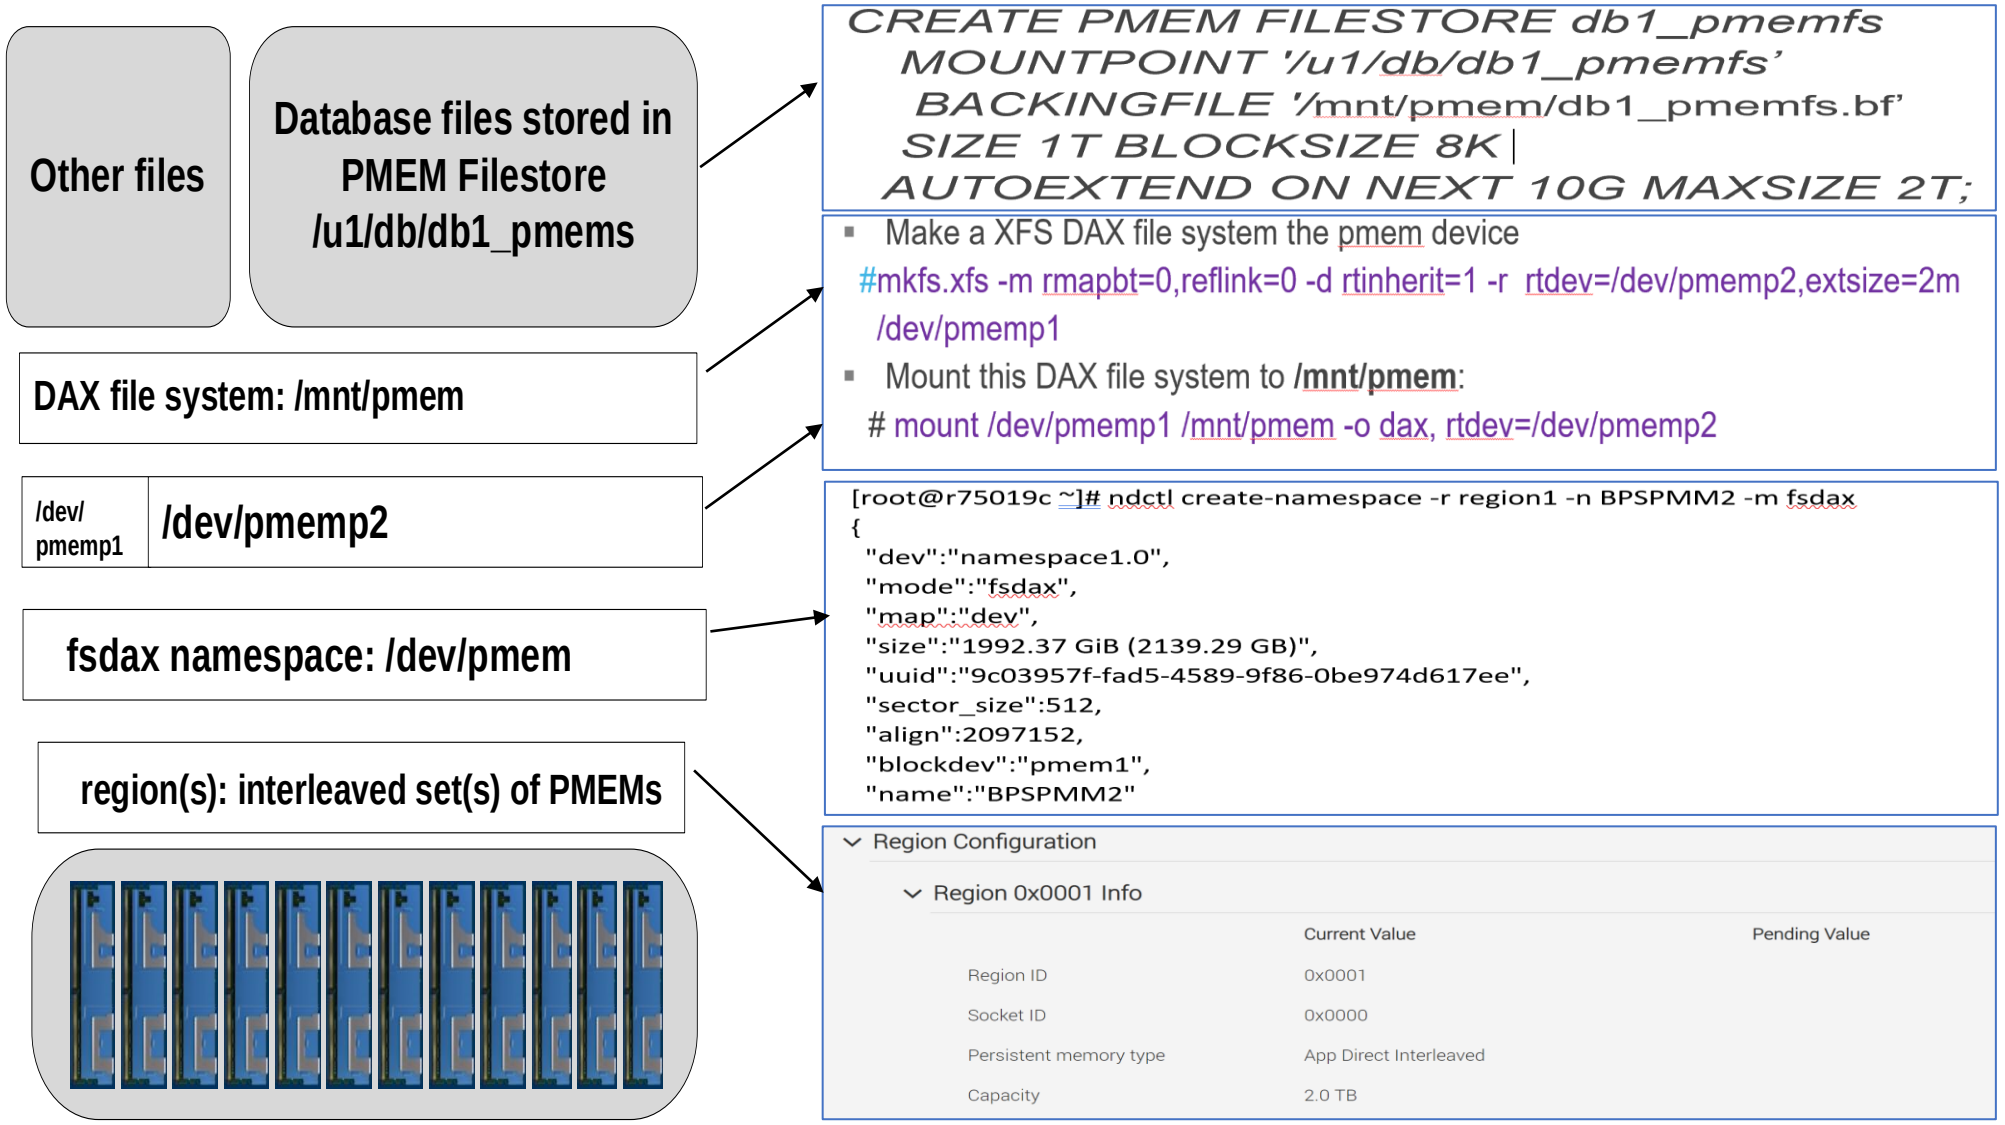

Oracle PMem Filestore in Oracle Database 21c

How PMem filestore stores database files:

Storing all the datafiles and redo logs files to the */u1/db/db1_PMemfs* :

```
SQL> select FILE_NAME from dba_data_files;

FILE_NAME
--------------------------------------------------------------------------------
/u1/db/db1_pmemfs/datafile/o1_mf_system_jlyljn1k_.dbf
/u1/db/db1_pmemfs/datafile/o1_mf_sysaux_jlyljrfn_.dbf
/u1/db/db1_pmemfs/datafile/o1_mf_undotbs1_jlyljsbt_.dbf
/u1/db/db1_pmemfs/datafile/o1_mf_users_jlylk1m5_.dbf
```

```
SQL> create bigfile tablespace pmem_data datafile
  2  '/u1/db/db1_pmemfs/datafile/tpccdata.dbf' size 1100G;

Tablespace created.
```

```
SQL> select MEMBER from v$logfile;

MEMBER
--------------------------------------------------------------------------------
/u1/db/db1_pmemfs/onlinelog/group8.log
/u1/db/db1_pmemfs/onlinelog/group9.log
/u1/db/db1_pmemfs/onlinelog/group10.log
/u1/db/db1_pmemfs/onlinelog/group11.log
```

```
> select name from v$controlfile;

NAME
--------------------------------------------------------------------------------
/u1/db/db1_pmemfs/controlfile/o1_mf_jlyljmbw_1.ctl
/u1/db/db1_pmemfs/controlfile/o1_mf_jlyljmbw_2.ctl
/u1/db/db1_pmemfs/controlfile/o1_mf_jlyljmbw_3.ctl
```

Oracle Memory Speed File System (OMS) in Oracle 19c

Oracle Memory Speed File System (OMS) Architecture and configuration

Database files stored in
OMS file System mounted on
/oracle/omsfs

Create OMS file system on
/mnt/pmem/omsuberfile.$ORACLE_SID

Uber file:
/mnt/pmem/omsuberfile.$ORACLE_SID

DAX file system: /mnt/pmem

fsdax namespace: /dev/pmem

region(s): interleaved set(s) of PMEMs

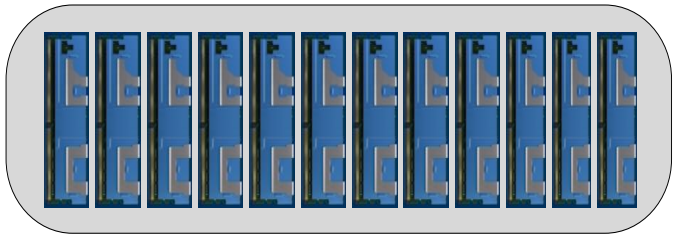

```
$ORACLE_HOME/bin/omsfscmds
OMS> mkfs /mnt/pmem0/omsuberfile.test
OMS> mount /mnt/pmem0/omsuberfile.test  /oracle/omsfs
```

```
fallocate -l 1400G /mnt/pmem0/omsuberfile.$ORACLE_SID
xfs_bmap /mnt/pmem0/omsuberfile.$ORACLE_SID
```

```
# mkfs.xfs -f -d su=2m,sw=1 /dev/pmem0
# mount -o dax /dev/pmem0 /mnt/pmem0
```

```
[root@r75019c ~]# ndctl create-namespace -r region1 -n BPSPMM2 -m fsdax
{
  "dev":"namespace1.0",
  "mode":"fsdax",
  "map":"dev",
  "size":"1992.37 GiB (2139.29 GB)",
  "uuid":"9c03957f-fad5-4589-9f86-0be974d617ee",
  "sector_size":512,
  "align":2097152,
  "blockdev":"pmem1",
  "name":"BPSPMM2"
```

Region Configuration

Region 0x0001 Info

	Current Value
Region ID	0x0001
Socket ID	0x0000
Persistent memory type	App Direct Interleaved
Capacity	2.0 TB

Oracle Memory Speed File System (OMS) in Oracle 19c

Deploying Oracle Database with Oracle Memory Speed file system

- Enabling the OMS ODM Library OMS file system /oracle/omsfs

```
[root@kaipmmdb oracle]# cd $ORACLE_HOME/rdbms/lib
[root@kaipmmdb lib]# make -f ins_rdbms.mk oms_on
rm -f /u01/app/oracle/product/19.3.0/rdbms/lib/odm/libomsodm19.so; \
   cp /u01/app/oracle/product/19.3.0/lib/libomsodm19.so
/u01/app/oracle/product/19.3.0/rdbms/lib/odm/libomsodm19.so
```

- Migrate Oracle control files, data files, redo log files, temporary files to the OMS file system: /oracle/omsfs

POC1: PMem App-Direct for 19c Oracle database

Showed a low I/O latency using PMems App Direct mode as storage

With standard HammerDB workload config, it showed 30 us I/O latency for db file I/Os.

(30 us = 0.03 ms = 0.000,030 second)

Top 10 Foreground Events by Total Wait Time

Event	Waits	Total Wait Time (sec)	Avg Wait	% DB time	Wait Class
DB CPU		8889		97.3	
db file sequential read	54,025,552	1627.8	30.13us	17.8	User I/O
log file sync	5,900,209	481	81.52us	5.3	Commit

The Performance monitor shows No I/O waits (Blue), all sessions are using CPUs (Green)

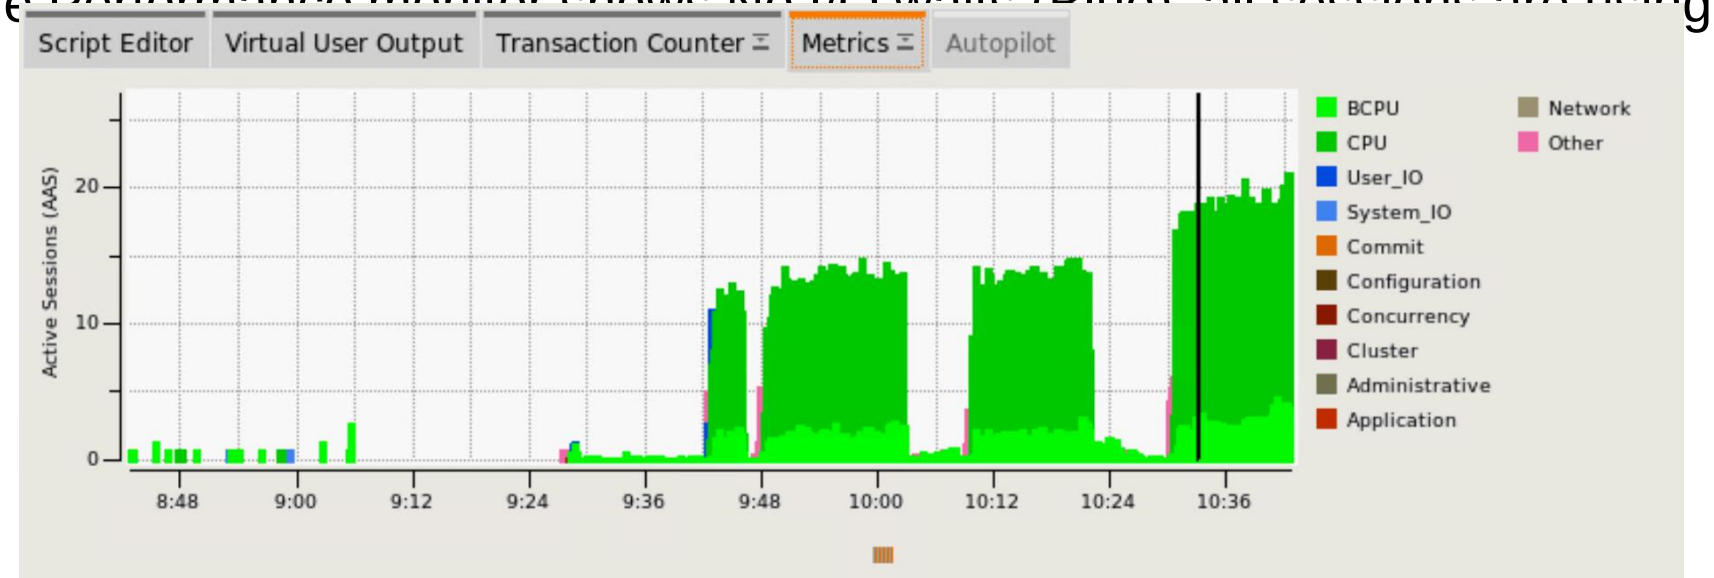

POC1: PMem App-Direct for 19c Oracle database

An extremely high performance database with PMems as storage:

- Even higher transaction throughputs (TPMs):
  Over 3 million tpms (Average 2,769,198 tpms)

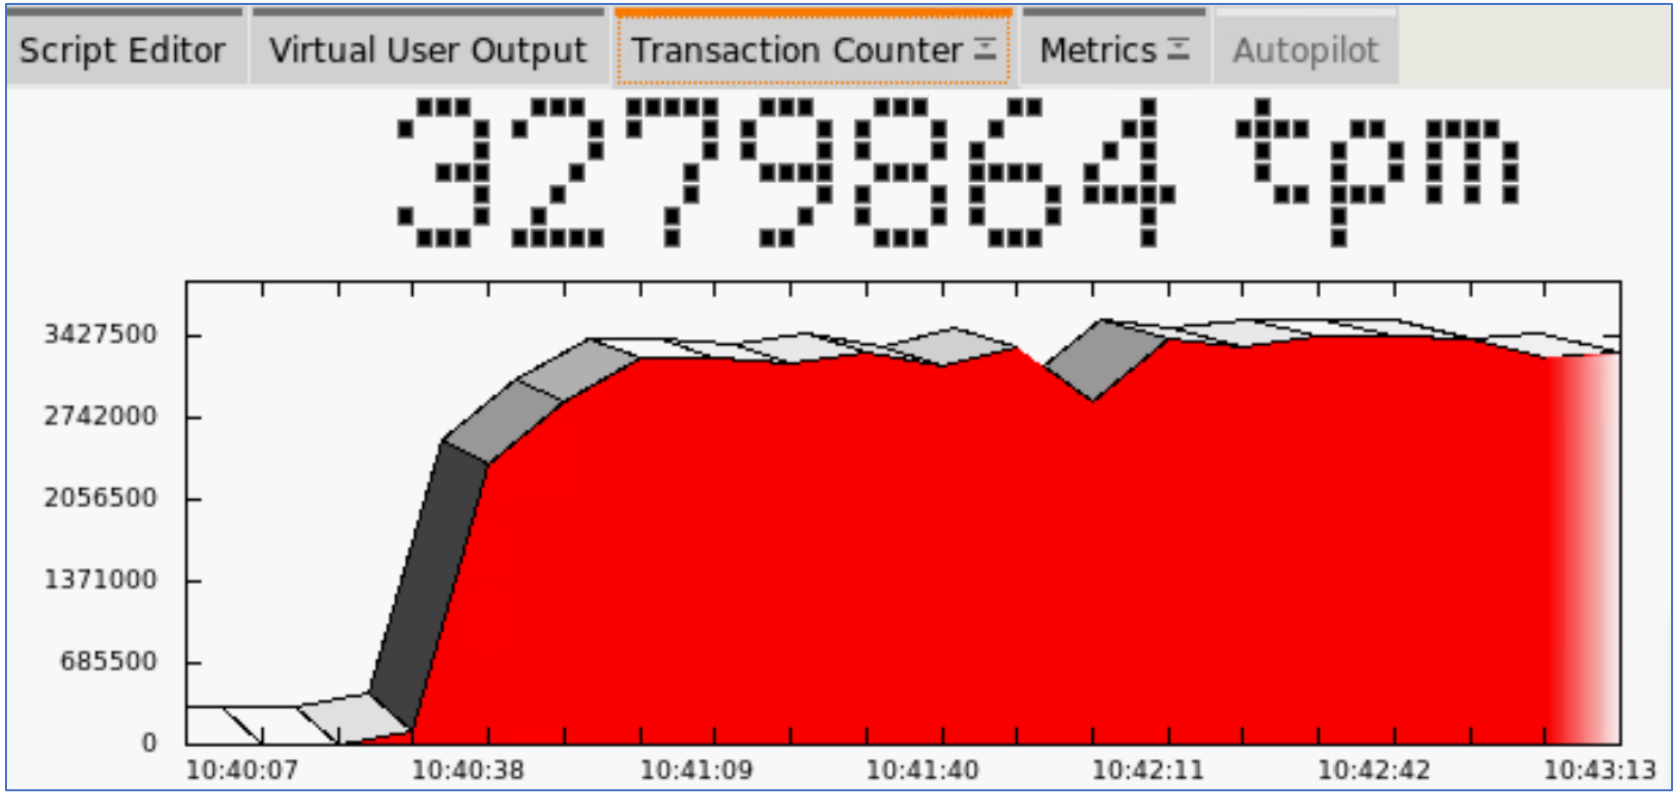

POC1: PMem App-Direct for 19c Oracle database

Performance Studies over more workloads by increasing the concurrent users

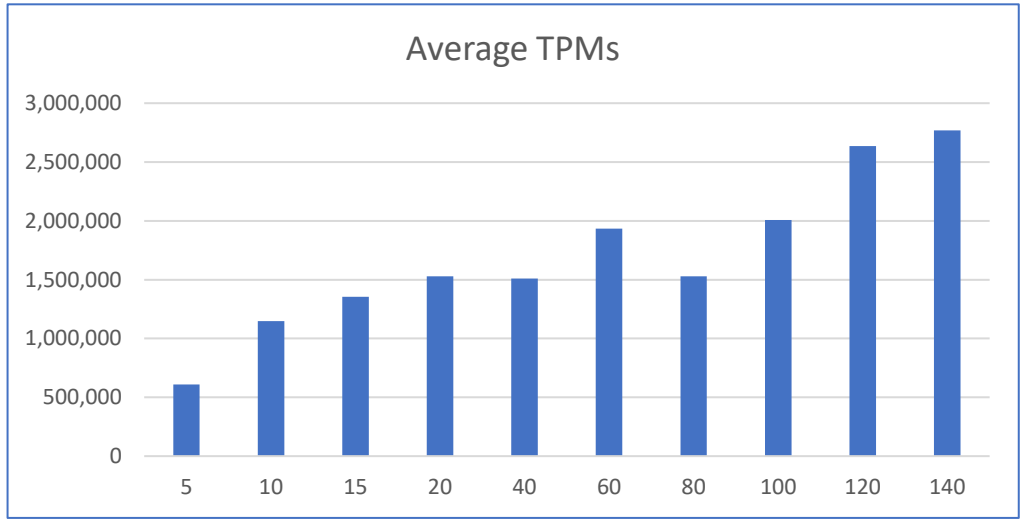

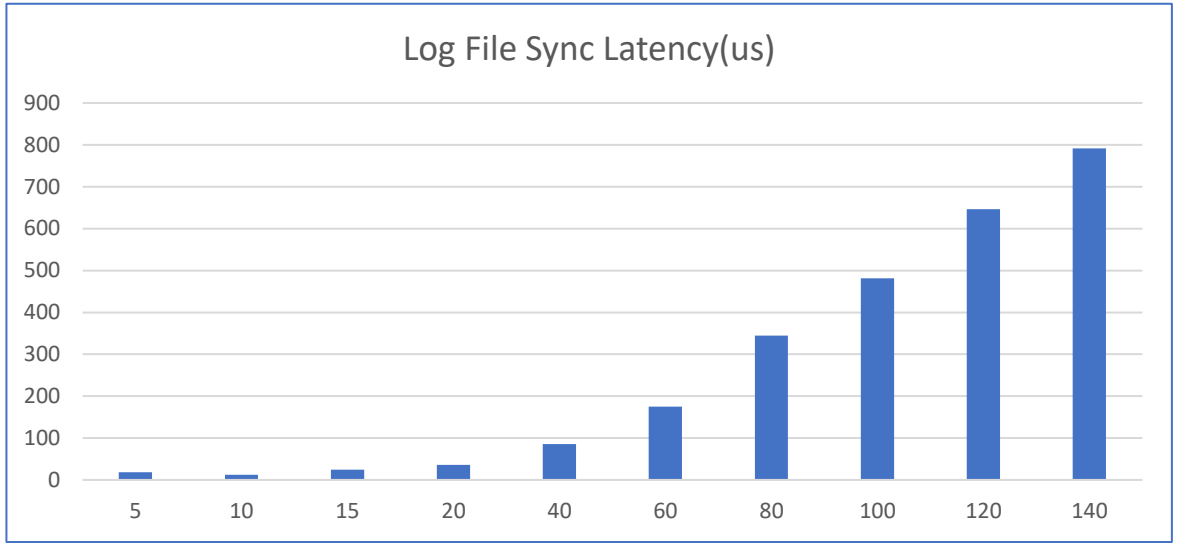

POC1: PMem App-Direct for 19c Oracle database

Performance Studies over more workloads by increasing the concurrent users

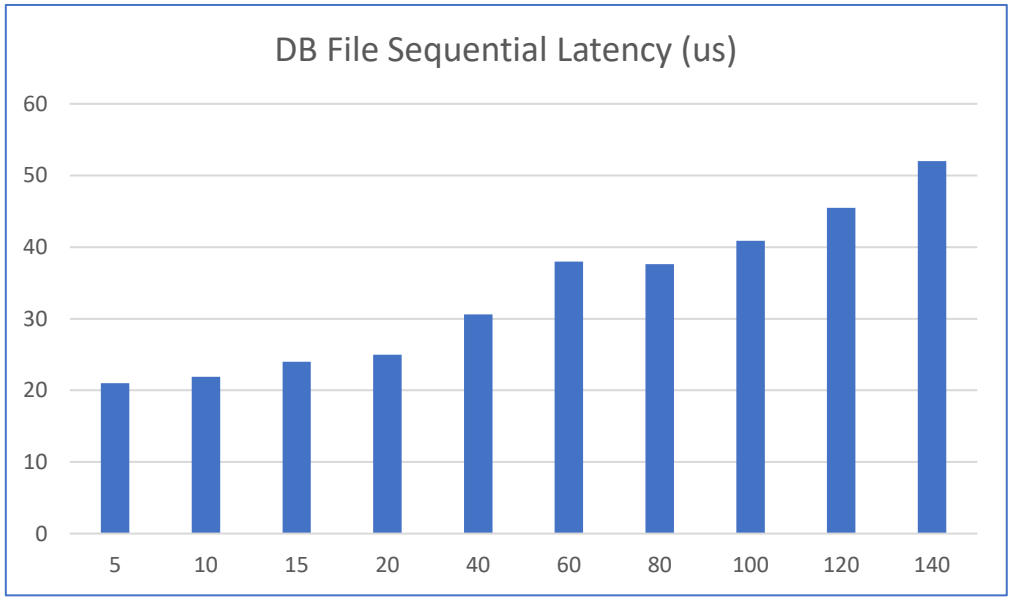

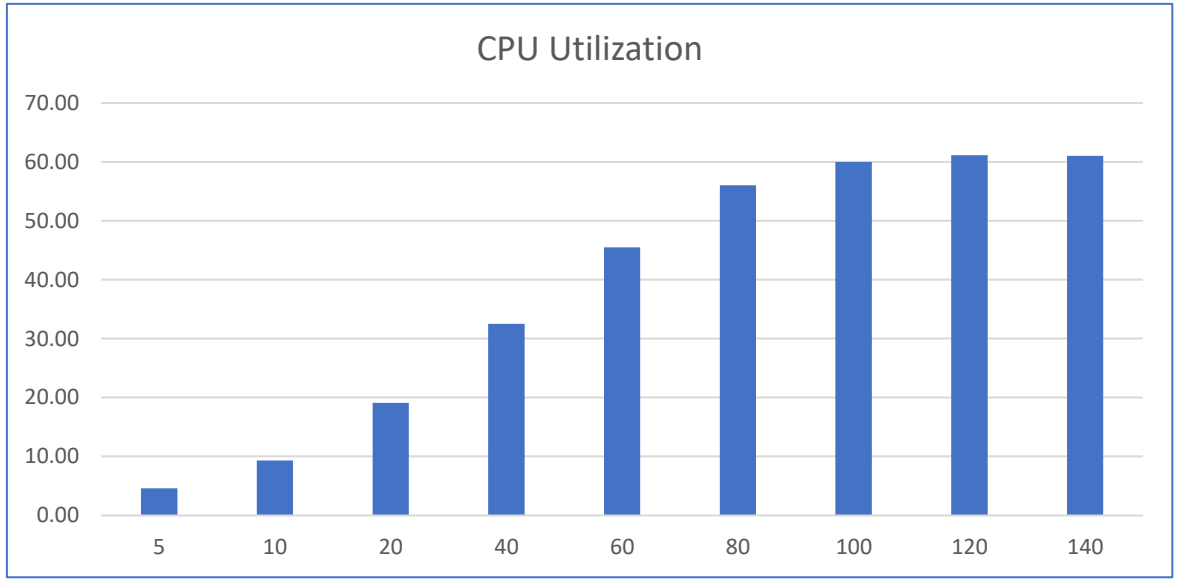

POC2: PMem App-Direct for 21c Oracle database

An extreme low I/O latency using PMem App Direct mode as storage

It showed 24 us I/O latency for db file I/O while it is common to see over 1ms=1000us I/O latency

(24 us = 0.024 ms = 0.000,024 second)

Top 10 Foreground Events by Total Wait Time

Event	Waits	Total Wait Time (sec)	Avg Wait	% DB time	Wait Class
DB CPU		14.8K		88.3	
log file sync	11,517,762	2255.5	195.83us	13.4	Commit
db file sequential read	38,619,811	905.6	23.45us	5.4	User I/O

The Performance monitor shows No I/O waits (Blue), all sessions are using CPUs (Green)

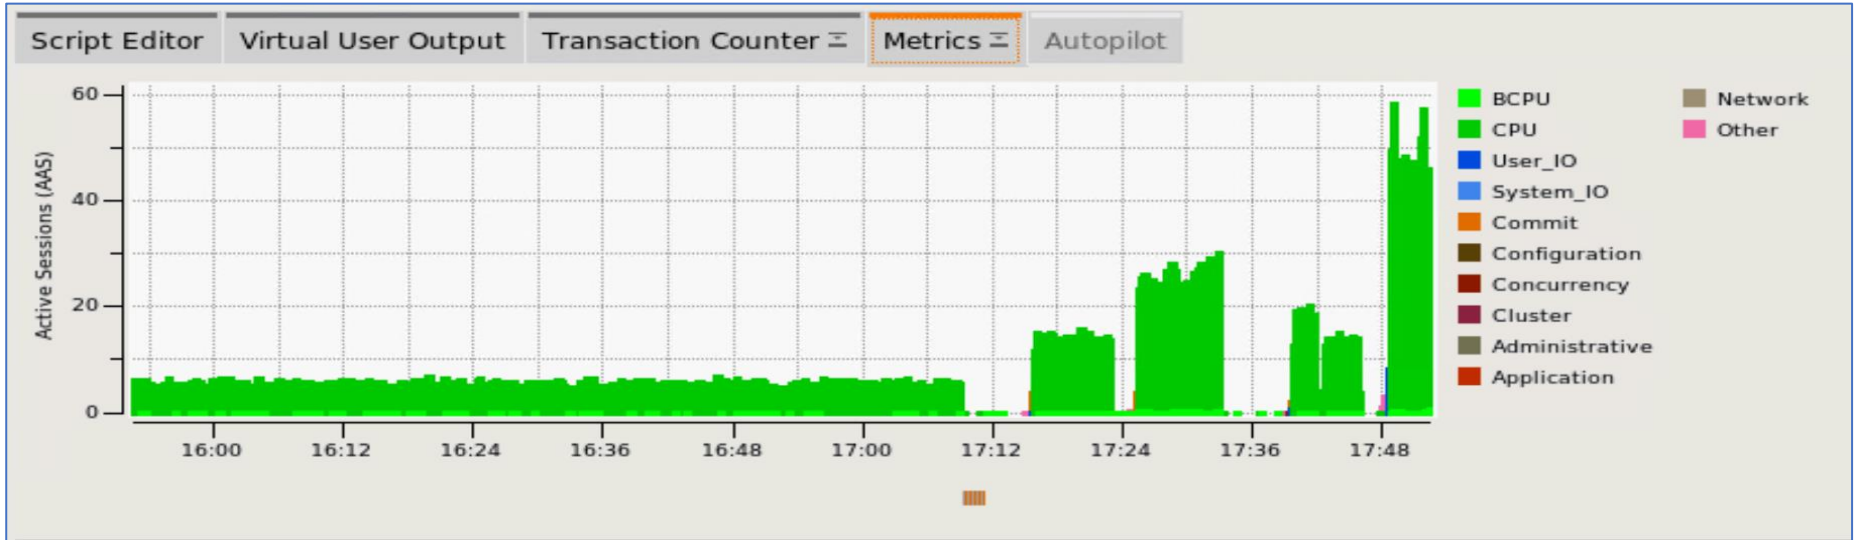

POC2: PMem App-Direct for 21c Oracle database

- An extremely high performance database with PMems as storage:
  - Even higher transaction throughputs (TPMs):

    Over 3 million tpms (Average 3,020,442 tpms)

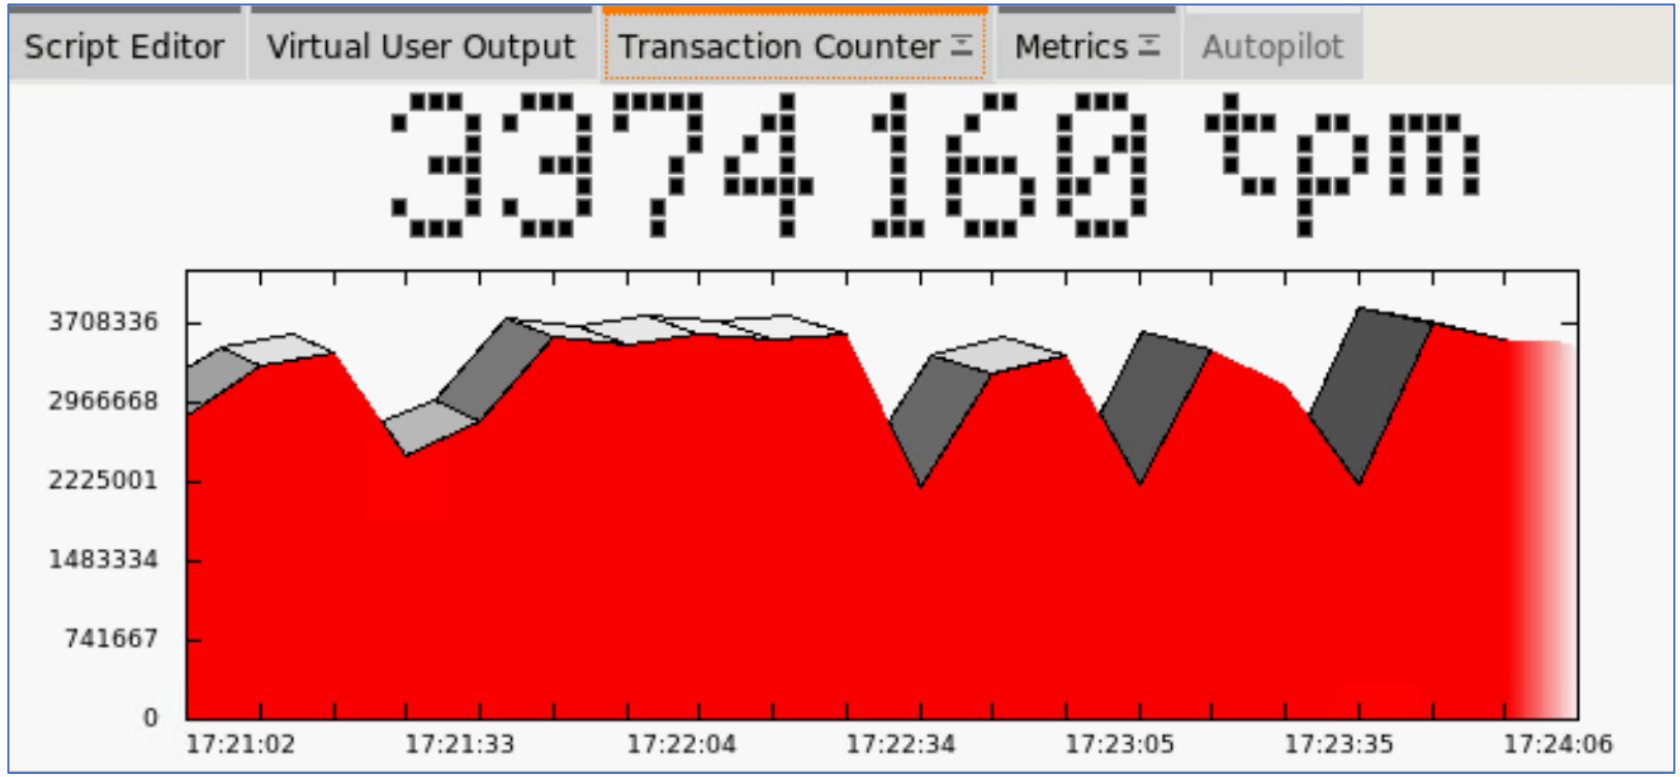

POC2: PMem App-Direct for 21c Oracle database

- Performance Studies over workloads

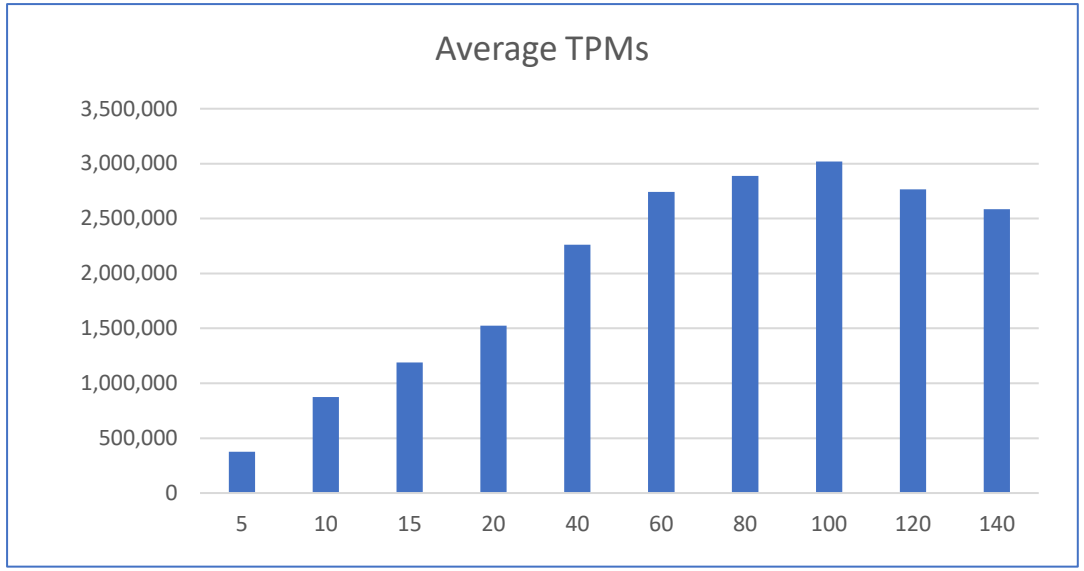

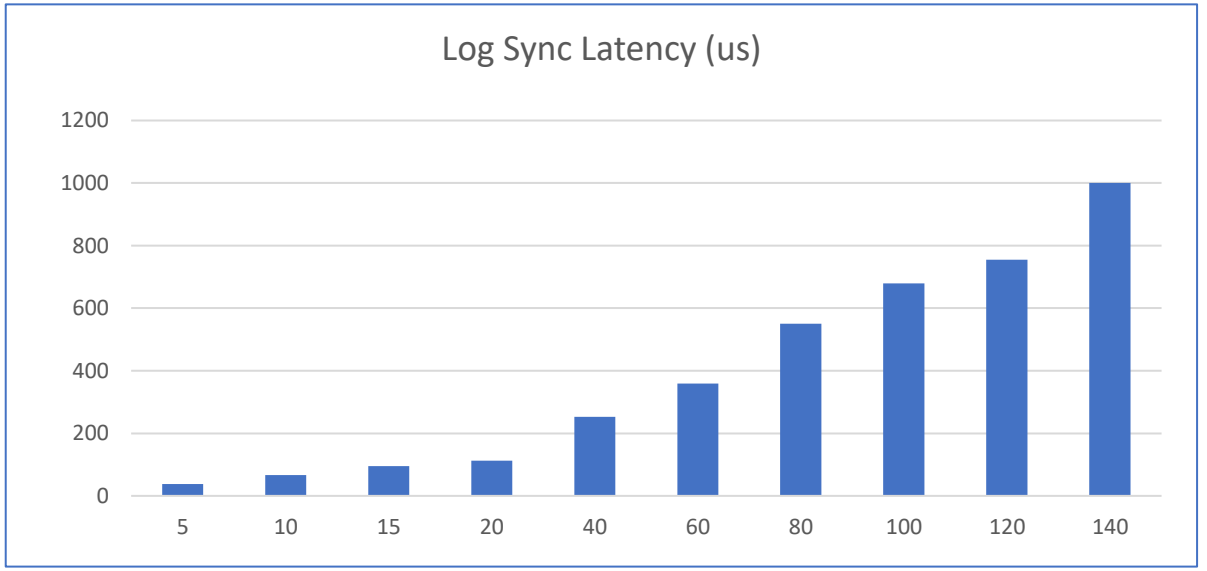

POC2: PMem App-Direct for 21c Oracle database

- Performance Studies over workloads

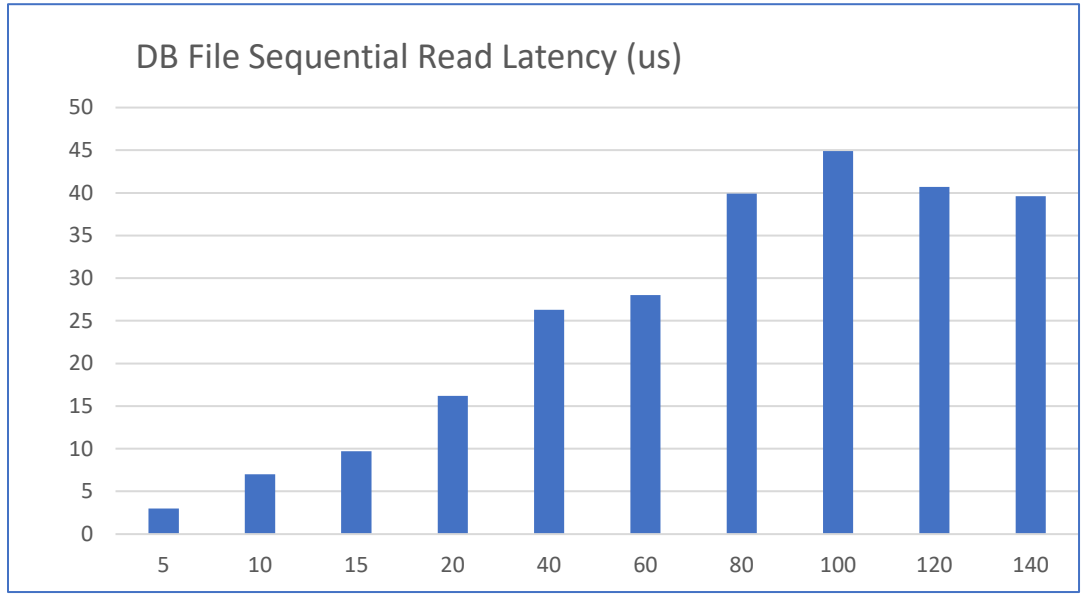

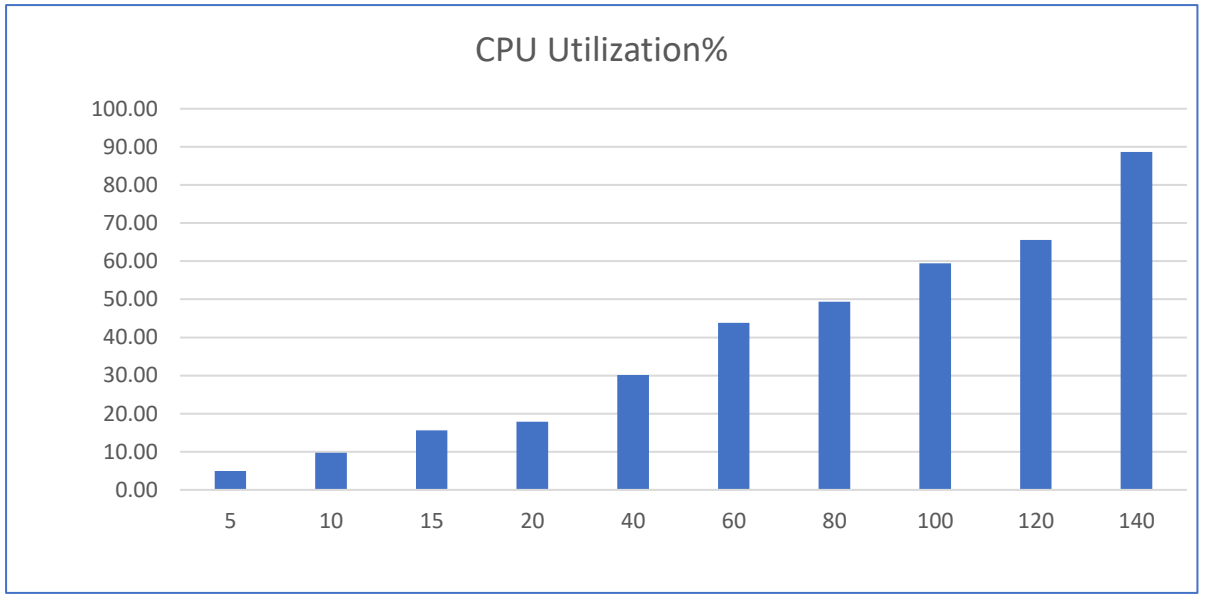

PMems used in Oracle Exadata Systems

Oracle Exadata X8M and X9M: PMems in storage serve

- Use Persistent memory in each storage server in front of Flash memory
- Uses RDMA instead of I/O to read remote PMems.
- PMem sharable across DBs : Used as a cache for hottest data
- PMem mirrored across storage serves for fault-tolerance
- <19 microseconds IO latency and 16 Millions IOPs for 8k I/Os

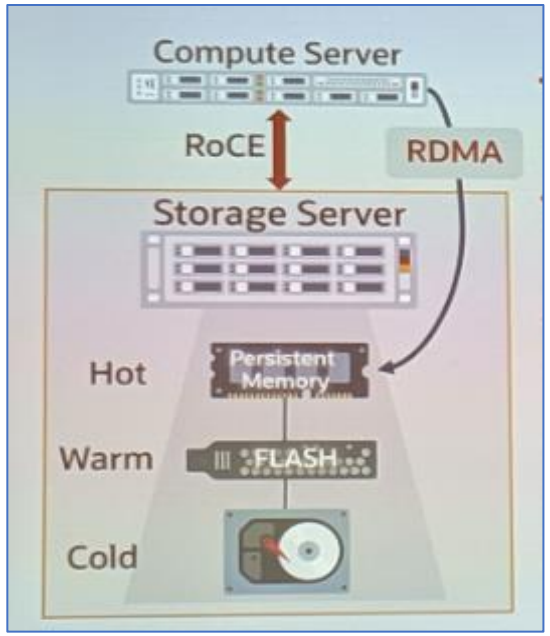

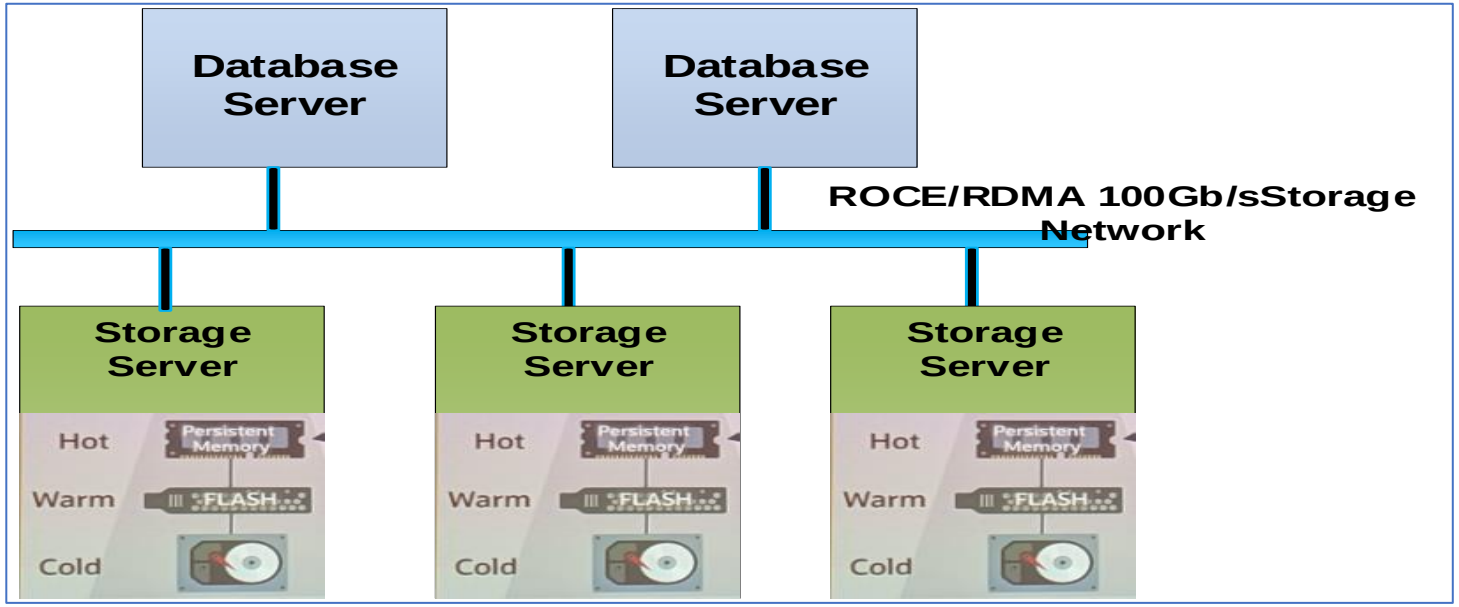

Oracle Sharding Database with Persistent Memory

- **Oracle Sharding Database Archietcure :**

  Horizontal partitioning of data across independent database (called shards)
  - Each shard holds a subset of the data in sharded tables
  - Sharded table: same columns, subset of rows.
  - Can be single node or RAC

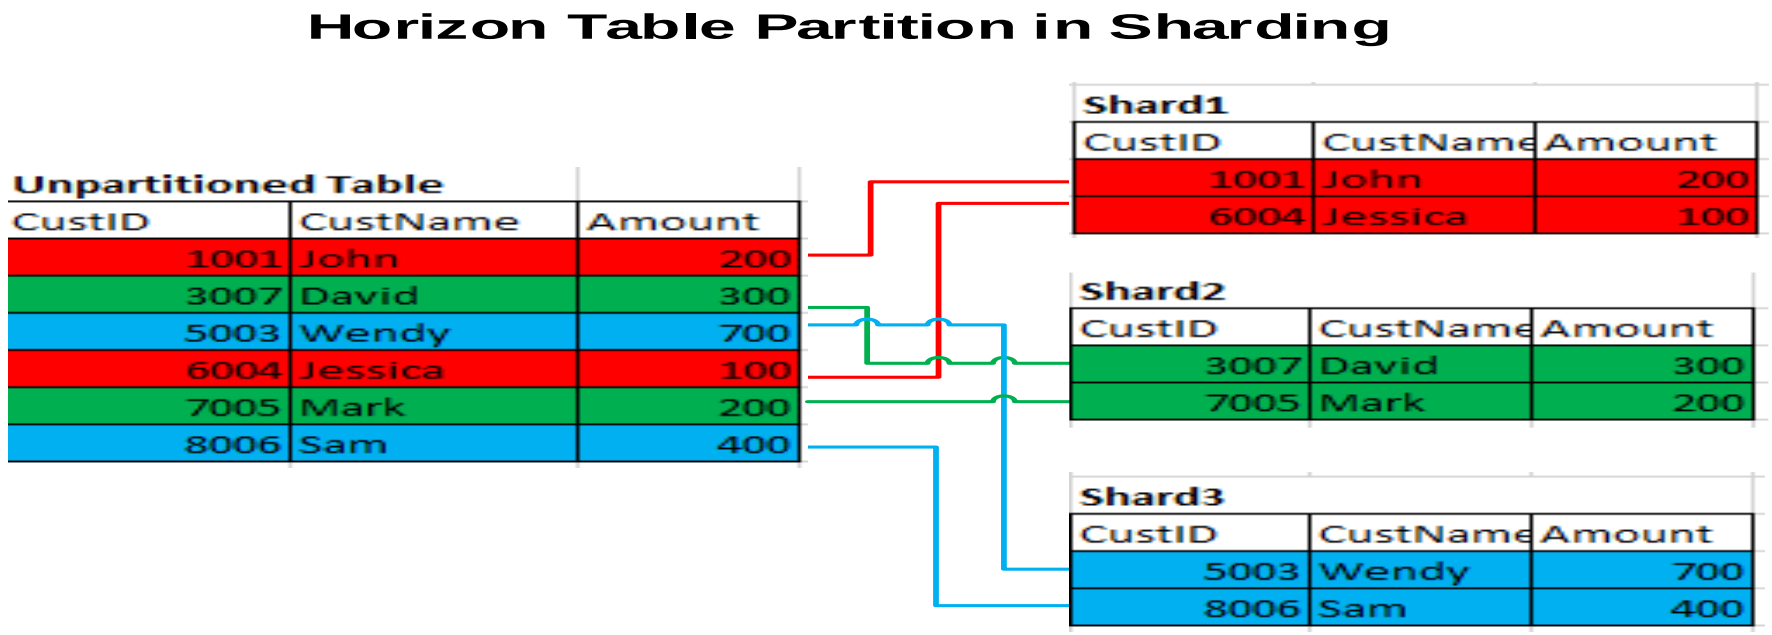

Oracle Sharding Database with Persistent Memory

- **Shared-nothing between sharded databases**
  - Each in dedicated server
  - Shards don’t share any hardware
  - Different architecture from RAC
- Application point of view:
  - Logical group as a single database
  - Data partition transparent to application

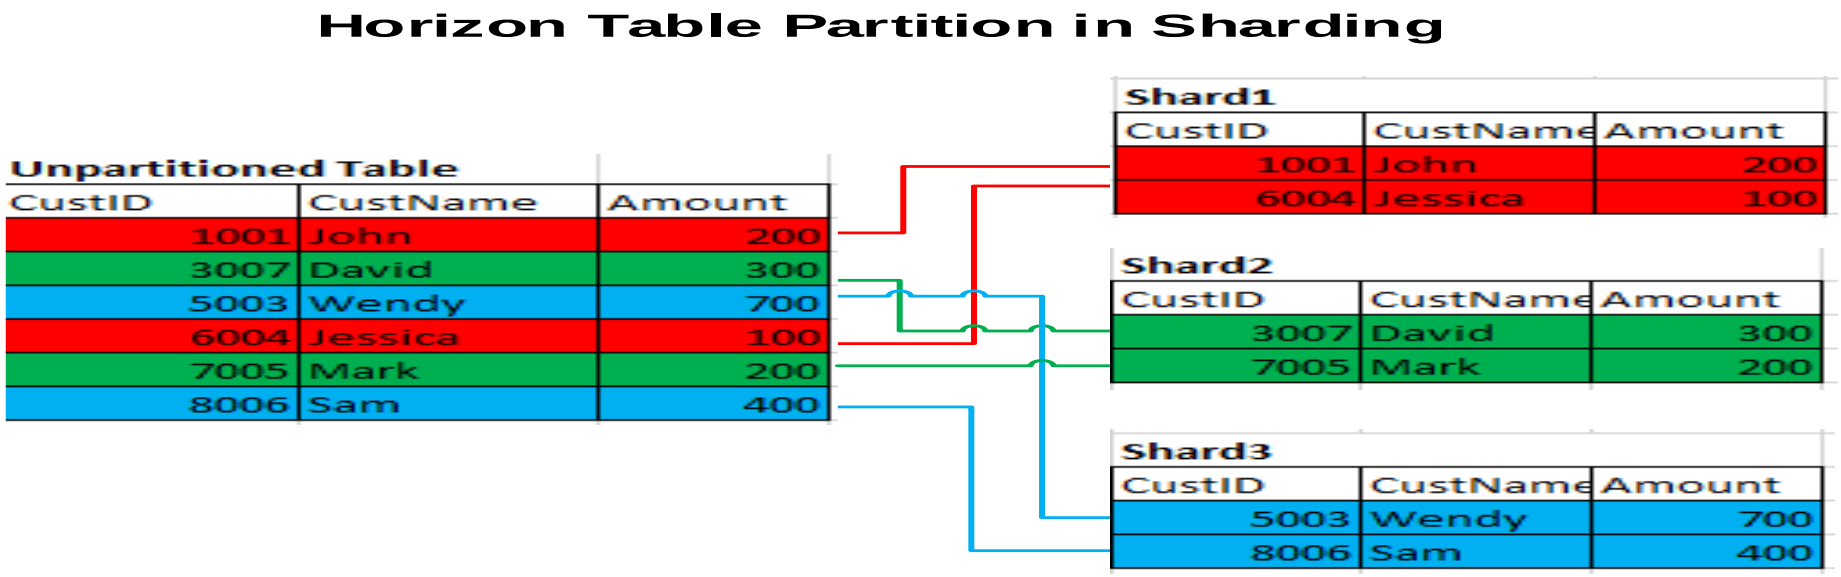

Oracle Sharding Database with Persistent Memory

Oracle 21c sharding database support Persistent memory as storage devices for shard databases

. Leverage the low latency and high performance provided by persistent memory

. Leverage the mass scalability of the sharding database architecture

. The database size is not limited by the maximum capacity of the persistent memory of each database node,

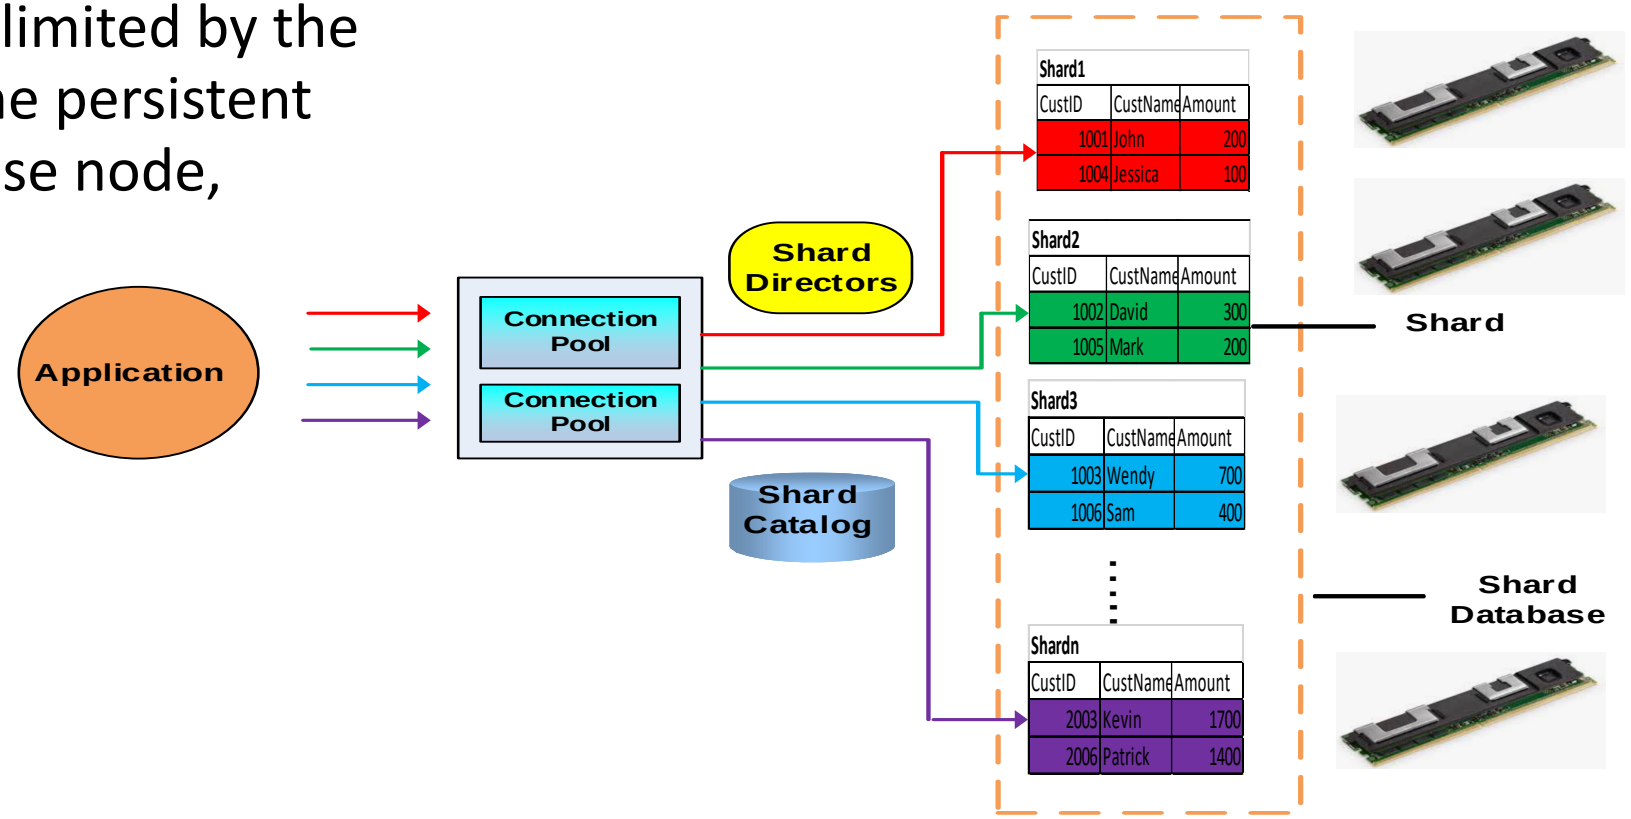

Intel's Latest Decision on Optane PMem

- Intel has announced that it is winding down its Optane PMem business entirely Intel's Q2 quarterly results
  - Intel developed a line of Optane SSDs and persistent memory DIMMs in 2020 based on the 3D XPoint technology it developed with Micron in 2015.
  - Micron sold off the 3D XPoint fab it operated in 2021, and left the market
  - Intel initiated the winding down of its Optane Persistent memory business in Q2
  - Intel has stopped manufacturing Optane products since Micron's departure
  - Intel will incur a $559 million inventory impairment/write-off as it exits the market
  - Intel has two years' worth of supplies on hand
  - Intel commits to supporting existing Optane Pmem and Optane SSD product lines through end-of-life, including product warranty terms for PMem 100 series, PMem 200 series, etc.
- Intel will also continue development of the Intel Optane Pmem 300 series on 4th Gen Intel Xeon Scalable processors (codenamed "Sapphire Rapids"). **

** https://blogs.vmware.com/vsphere/2022/08/intel-optane-memory-optimization-and-vsphere.html (8/28/2022)

** Gifts From Intel's Optane Memory (forbes.com)

Intel’s Latest Decision on Optane PMem

- What this means to customers:
  - Intel has two years' worth of supplies for server vendors such as Dell and Oracle etc.
  - Three major PMem major use cases in Oracle Databases with the continuous supplies of PMems
    - PMems in memory mode to expand DRAMs
    - PMems in App-Direct mode for the database storage : Oracle 21c PMem store and 19c OMS
    - PMems used as hot cache for hottest data in Oracle Exadata systems
    - What this means to customers who uses PMems :
  - VMware also announced its continuous support Intel Optane PMem with vSphere 7.x and 8.x releases ***
- What is in future: alternative to PMem:
  - Intel has invested a technology called Compute Express Link(CXL) introduced in 2019
  - CXL defines a cache-coherent interface for connecting CPUs, memory, accelerators, and other peripherals.
  - CXL 1.1 enables memory to be attached directly to the CPU over the PCIe 5.0 link, shipped with Sapphire Rapids Xeon Scalable and AMD's fourth-gen Eypc Genoa and Bergamo processors later this year
  - Emulating some aspects of Intel's Optane persistent memory may have to wait until the first CXL 2.0-compatible CPUs

*** https://blogs.vmware.com/vsphere/2022/08/intel-optane-memory-optimization-and-vsphere.html

Summary and Q&A

- Persistent Memory Technology
- PMem Memory Mode for Oracle Database Performance
- Challenges of using PMem App-Direct mode for Oracle Database
  - Oracle PMem Filestore in Oracle 21c
  - Oracle Memory Speed File System in Oracle 19c
- Performance studies POCs with PMems in Oracle 19c and 21c databases
- Persistent Memory in Oracle Exadata X8M and x9M
- Intel’s latest decision on PMem

Thank You!

kai_yu@dell.com
http://kyuoracleblog.wordpress.com/

Acknowledgement

I would like to thanks Dell Oracle Solutions Engineering team's work on Oracle Database In-Memory with Intel Optane Persistent Memory and the technical whitepaper
https://www.delltechnologies.com/resources/en-us/asset/technical-guides-support-information/solutions/h18295-inteloptanePMem-oracleinmem-dg.pdf

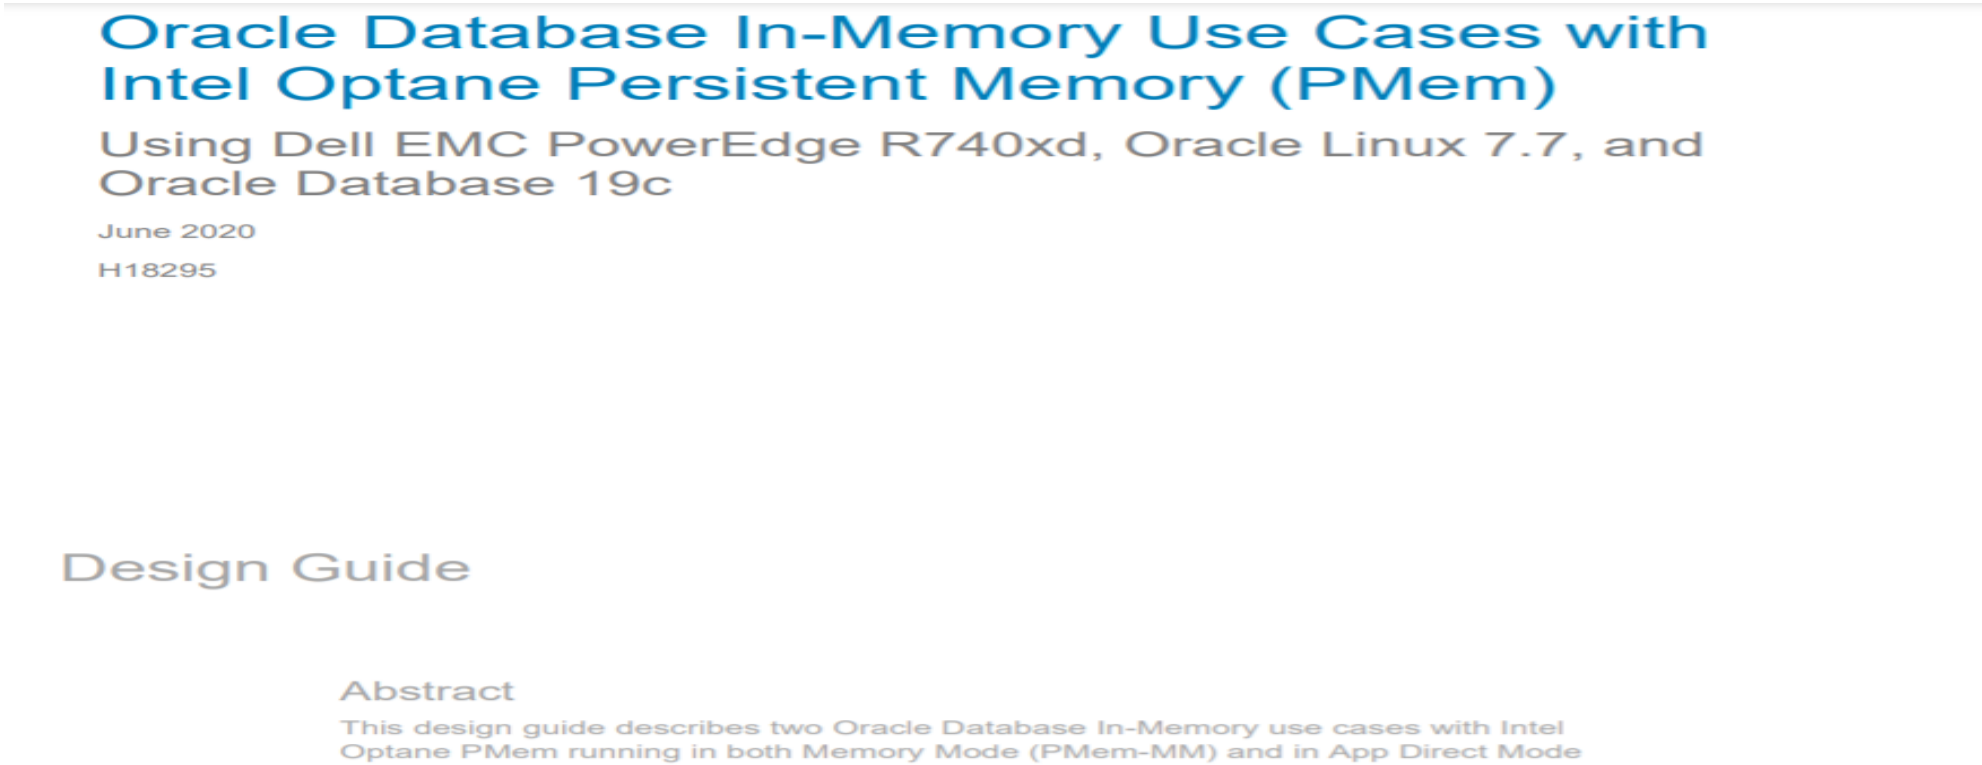

Oracle Database In-Memory Use Cases with Intel Optane Persistent Memory (PMem)

Using Dell EMC PowerEdge R740xd, Oracle Linux 7.7, and Oracle Database 19c

June 2020

H18295

Design Guide

Abstract

This design guide describes two Oracle Database In-Memory use cases with Intel Optane PMem running in both Memory Mode (PMem-MM) and in App Direct Mode (PMem-AD) inside a Dell EMC PowerEdge server. Our goal is to determine if PMem technology can provide an alternative option to traditional DRAM and to other fast-tiered storage disks in a Database In-Memory environment. It also provides the steps to configure Oracle Database 19c with PMem as a persistent block-addressable storage device.