APPLICATION NOTE

Intel Corporation

Intel® Optane™ DC Persistent Memory – Telecom Use Case Workloads

Author

David Kennedy

1 Introduction

This document provides an overview of Intel® Optane™ DC persistent memory for 2nd generation Intel® Xeon® Scalable processors (formerly Cascade Lake). The document contains a description of the technology and discusses some use cases for telecom environments.

The requirement for large memory footprints for artificial intelligence and media/streaming servers, as well as low latency access to large subscriber databases and routing tables in telecom networks provides a challenge for customers in deploying manageable nodes in the network. Often, instances of these databases are stored across multiple nodes, which increases the total cost of ownership and manageability costs. Also, due to the different latency requirements, there are often multiple copies of database instances across the network, plus data duplication for redundancy.

Intel Optane DC persistent memory is highly beneficial for applications that require high capacity, low latency, and resiliency, which are applicable across a number of telecom specific use cases, including:

- In-memory databases
- Analytics
- Virtualization
- HPC clusters

Intel Optane DC persistent memory provides the opportunity to:

- Consolidate databases into smaller, easier to manage instances.
- Position large datasets, such as network telemetry, right next to the CPU, directly in-memory, for faster big data analytics for network automation.
- Enable a whole new class of shared, distributed data layer to enable cloud native infrastructure with ultra low latency, similar to DRAM, but without the cost of DRAM.
- Cache rapidly emerging live linear video content or highest performance video-on-demand in a CDN with capacity-like storage, but performance-like memory, at a lower cost.
- Provide faster redundancy for failover.
- Take advantage of being able to reuse pre-existing databases/routing tables through a reboot cycle due to persistent memory mode, which further increases high availability of the network.

This document is part of the Network Transformation Experience Kit, which is available at: https://networkbuilders.intel.com/network-technologies/network-transformation-exp-kits

Table of Contents

Figures

Tables

1.1 Intended Audience

This document is intended for Telecom Equipment Manufacturers (TEMs) and Communication Service Providers (CSPs) who are planning and deploying virtualized mobile core infrastructure running on the latest Intel® Xeon® Scalable processors and want to understand the benefits of Intel Optane DC Persistent Memory.

1.2 Terminology

Table 1. Terminology

ABBREVIATION	DESCRIPTION
BIOS	Basic Input / Output System
CDN	Content Delivery Network
CSP	Communication Service Provider
DBaaS	Database as a Service
DIMM	Dual In-Line Memory Module
DRAM	Dynamic Random Access Memory
GPRS	Global Packet Radio Service
GSM	Global System for Mobile Communication
HLR	Home Location Register
HPC	High Performance Compute
HSS	Host Subscriber Service
IMC	Integrated Memory Controller
IMDB	In-Memory Database
MME	Mobility Manager Entity
NTB	Non Transparent Bridge
PCRF	Policy and Charging Rule Function
RDMA	Remote Direct Memory Access
TEMs	Telecom Equipment Manufacturers
UDR	Unified Data Repository
VM	Virtual Machine
VNF	Virtual Network Function

1.3 Reference Documents

Table 2. Reference Documents

REFERENCE	SOURCE
3D XPoint™ Technology Revolutionizes Storage Memory Informational Video	https://www.youtube.com/watch?v=Wgk4U4qVpNY
Assembling VoLTE CDRs based on Network Monitoring – Challenges with Fragmented Information	https://www.researchgate.net/publication/317003510_Assembling_VoLTE_CDRs_based_on_Network_Monitoring_-_Challenges_with_Fragmented_Information
Intel® Optane™ DC Persistent Memory - Content Delivery Networks Use Case Application Note	https://builders.intel.com/docs/networkbuilders/intel-optane-dc-persistent-memory-content-delivery-networks-use-case.pdf

2 Intel Optane DC Persistent Memory Overview

Intel Optane DC persistent memory combines DRAM-like low latency performance with SSD-class capacities and persistency characteristics, allowing users to take advantage of very large memory footprints close to the processor. Intel Optane DC persistent memory utilizes Intel® 3D XPoint™ technology, a memory and storage breakthrough developed jointly with Micron Technology* (informational video). Intel 3D XPoint technology uses a stackable cross-gridded data access, with bit storage based on a change of bulk resistance.

Intel Optane DC persistent memory is physically and electrically compliant with the DDR4 specification and can be plugged directly into the existing memory connectors. Table 3 and Figure 1 highlight the features of Intel Optane DC persistent memory.

Table 3. Intel Optane DC Persistent Memory Capabilities

Capability	Value
Device Raw Capacity (GB)	128, 256, 512
Usable Capacity (GB)	98% of Device Capacity
12V On Support	Yes
S5 Support	Yes
S3 Support with 12V On	No
S3 Support with 12V Off	Yes
DIMM Connector Power Supplies	12V, 2.5V, 1.2V
12V Icc Inrush + Max. (A)	1.5 <= 10ms; 2.5 <= 10us
Tj min-max	0 – 85C
Operating Temperature (Ambient)	0 – 70C
DDRT Speed (MT/s)	1866, 2133, 2400, 2666
SPD Speed Max. (kHz)	400
Average Power Budget	128GB: 10W – 15W 256GB: 12W – 18W 512GB: 12W – 18W
Security	Encryption, Signed FW, key/passphrase

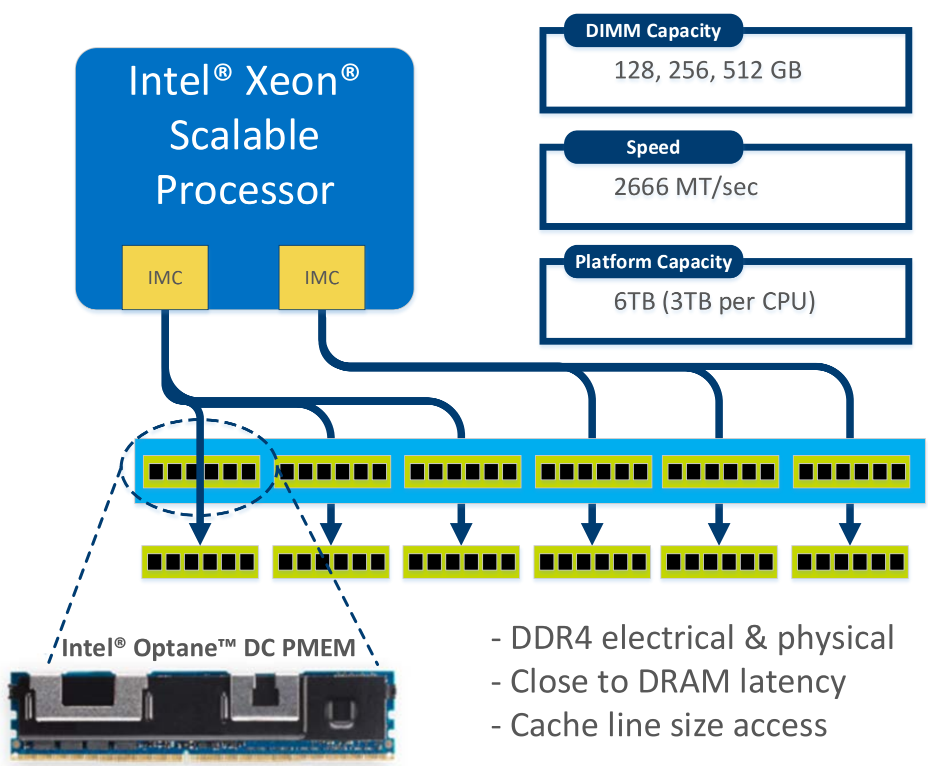

Figure 1. Intel Optane DC Persistent Memory

2.1 Modes

As shown in Figure 2, Intel Optane DC persistent memory has two available modes. It can run in either App Direct Mode, Memory Mode, or a hybrid of both.

- App Direct Mode
  - Persistent memory storage
  - DIMMs act as an independent memory resource under direct load/store control of the application
  - Byte addressable memory mapped to system physical address space
    - Directly accessible from applications
  - Applications can also access memory through block storage for fast direct-attach storage
  - Requires software enablement
- Memory Mode
  - Appears as high capacity DRAM type memory
  - Memory is volatile
  - No software enablement required

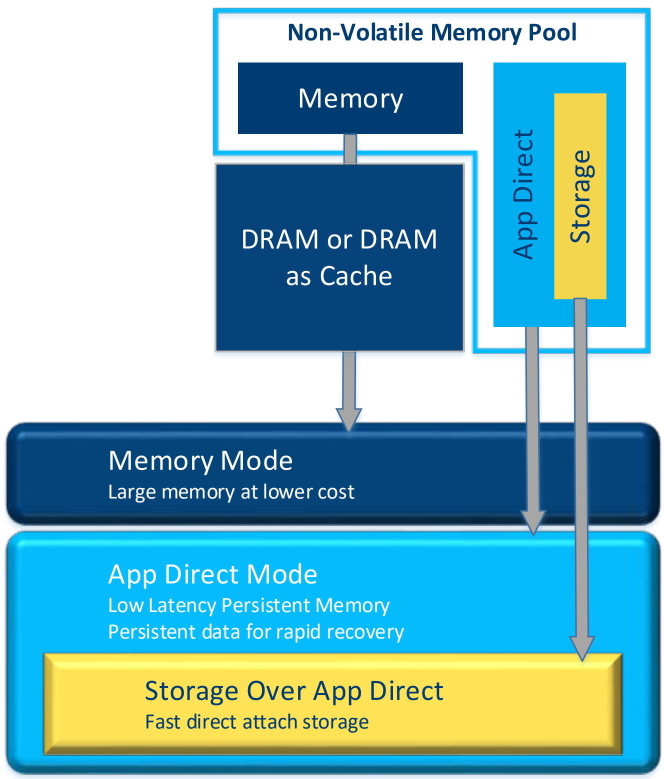

Figure 2. Intel Optane DC Persistent Memory Modes

2.2 Data Resiliency

One of the biggest contributors to downtime in a telecom network is the ability to resume context upon reboot, especially in cases for unplanned downtime. Applications can spend up to 30 minutes rebuilding context from previous boot, loading the subscriber database from a backup location, and rebuilding routing tables for ongoing sessions. This affects the overall availability indicator for network operators. Intel Optane DC persistent memory in Application Direct Mode provides the capability to maintain persistency through reboot and reduces initial boot time from minutes to seconds.

The traditional model, on the left side of Figure 3, shows how routing tables are stored in DDR memory and backed up in SSD memory. On initial boot or during a recovery scenario, the routing tables are restored from SSD memory into DDR memory.

The right side of Figure 3 shows the advantage of using Intel Optane DC persistent memory in Application Direct Mode to store the routing tables. After a reboot, the routing tables are kept persistent and are instantly available to the application when reboot is completed.

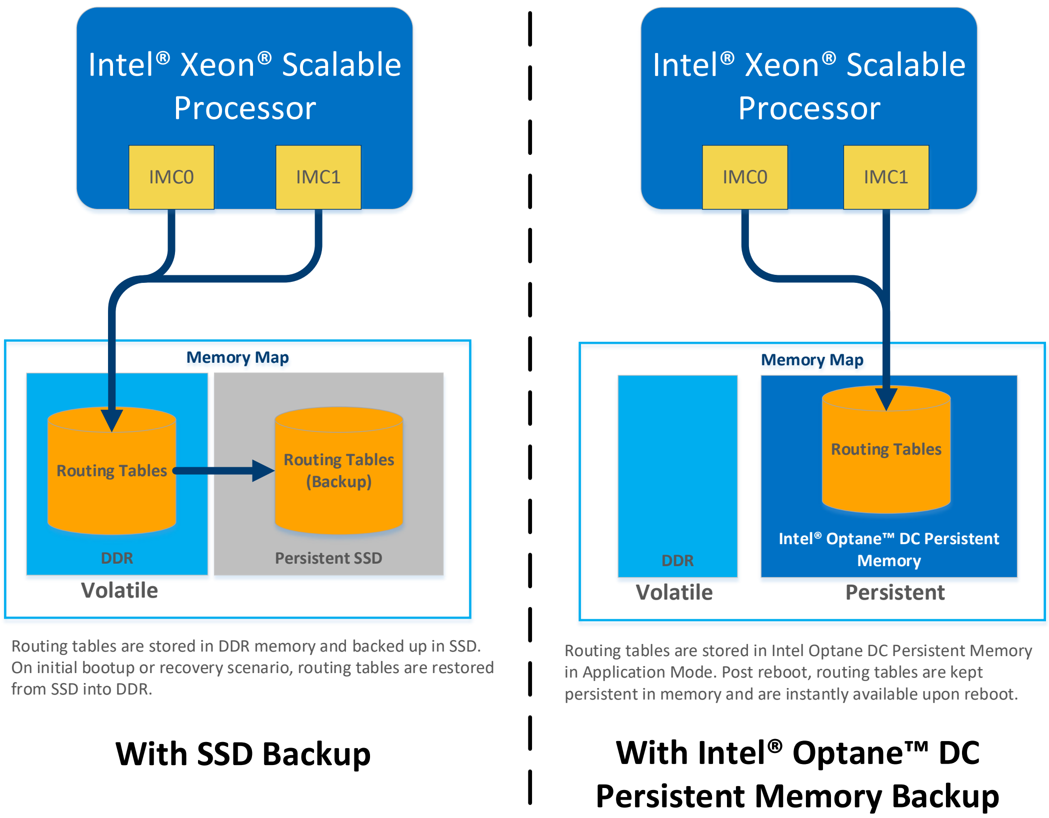

Figure 3. Intel Optane DC Persistent Memory Resiliency Through Reboot

2.3 Data Replication

Another expectation of telecom networks is the requirement for redundancy. This is especially true for databases, because a potential loss of this type of data would be detrimental to the operation of the network. Live databases have a redundant backup as well as a disaster recovery built into the function. Therefore, there is a requirement to maintain a coherent copy of all database instances across the network.

To ensure coherency across live databases, there is a requirement to do read and write transactions between live and backup databases. This has a cost from a CPU perspective, because a read/write has to go through the CPU and results in compute cycles being used. RDMA over Ethernet or PCIe provides the option to directly read and write to remote databases from Intel Optane DC persistent memory without having to directly access CPU resources, by doing direct memory to memory writes, thus reducing the latency of accesses between memory devices.

Figure 4 shows Intel Optane DC persistent memory used in a data replication scenario with two servers.

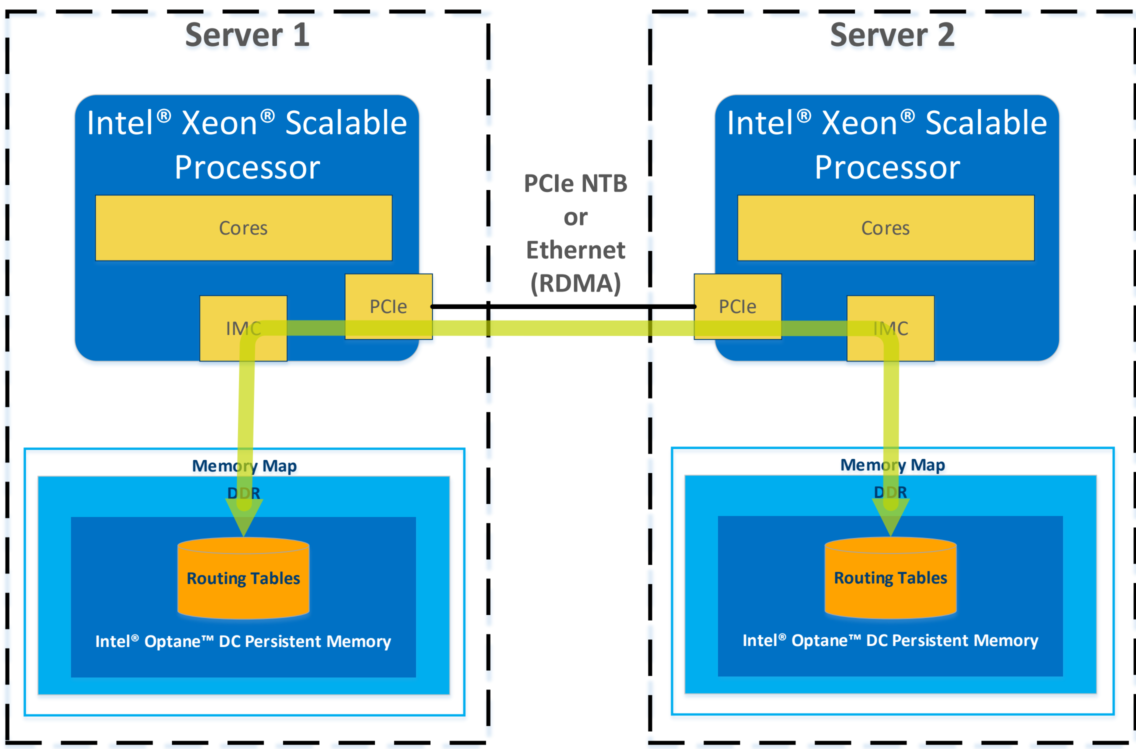

Figure 4. Intel Optane DC Persistent Memory Data Replication

This mechanism can optionally be done either through RDMA over PCIe using a Non Transparent Bridge (NTB) or over GbE devices that support RDMA, including Intel® Ethernet Converged Network Adapters X710 (formerly Fortville) and Intel® Ethernet Controller 810 Series (formerly Columbiaville). If a low latency requirement is critical, then we recommend a physical connection between the servers, instead of using a cloud connection.

3 Intel Optane DC Persistent Memory in Telecom Networks

Telecom networks contain numerous instances of subscriber information of all of their users (as well as roaming users from other networks), to allow for validation of users, billing applications, device location, and others. This section describes some commonly used databases.

3.1 Home Subscriber Service (HSS)

Home Subscriber Service (HSS) is a master database of subscriber information for the network. Generally it consists of the following:

- Subscriber profiles
- Authentication and authorization of the user
- Device location and IP information

3.2 Home Location Register (HLR)

Home Location Register (HLR) is a central database of each mobile device. Typically, a network contains multiple instances of HLRs, which can contain data such as:

- SIM card details (IMSI)
- Unique identifiers relating to mobile phone numbers (MSISDN)
- GSM services
- Call divert services
- Device physical location
- GPRS settings for packet services

3.3 Mobility Management Engine (MME)

The Mobility Management Engine (MME) is the control node for the access network and is responsible for bearer action, etc.

3.4 Content Delivery Network (CDN)

A Content Delivery Network (CDN) is a geographically distributed network of proxy servers and their data centers, whose goal is to provide high availability and high performance by distributing the service spatially relative to end-users. Passive consumption of video-on-demand (VoD) is increasingly giving way to live linear content and to immersive / intelligent visual experiences. CDNs are evolving to support emerging edge use cases, such as cloud gaming and virtual reality (VR). New approaches are needed to address more open, real-time, and cloud-based CDN solutions. CDNs serve a large portion of the Internet content today, including web objects (text, graphics, and scripts), downloadable objects (media files, software, documents), applications (e-commerce, portals), live streaming media, on-demand streaming media, and social media sites.

3.5 Operational Support System (OSS)

The Operational Support System (OSS) manages networks supporting management functions such as network inventory, service provisioning, network configuration and fault management.

In general, an OSS covers the following five functions:

- Network management systems
- Service delivery
- Service fulfillment, including the network inventory, activation and provisioning
- Service assurance
- Customer care

OSS/BSS and network management is becoming increasingly complex, especially as the market moves toward 5G and network slicing. There is a strong market need to collect and analyze all types of data, such as network telemetry, customer call reports, billing, and roaming in order to detect fraud, network intrusion, sentiment analysis, network traffic dynamics, and thereby keep end users safer, optimize network performance, minimize energy consumption, automatically predict fault/maintenance and intelligently or autonomously operate the network. Intel Optane DC persistent memory allows capturing huge amounts of data for real-time or near real-time analytics, machine learning and making sense of all the data in order to deliver the aforementioned services.

3.6 Policy Control Rules Function (PCRF)

This entity is responsible for aspects of charging of subscriber and billing information.

3.7 Other Components

Some of the network elements such as the SGSN require routing of packets to a static device, and thus build huge routing tables of data relating to packet steering. There is a potentially large benefit of having these tables stored on a single low latency memory map with persistency to save context through reboot.

4 Current Topology

4.1 Managing Subscriber Databases in 5G

In existing deployments of infrastructure such as an HSS node, TEMs provide a large cluster of storage nodes to address the full footprint for in-memory database applications. These databases are spread across a number of servers to meet the low latency access requirements using DDR memory, and to stripe the data across a number of nodes to remove the IO bandwidth constraint in accessing a single node.

In the current infrastructure, there are two database classes:

- Type 1
  - Hyperscale (e.g. HSS, UDR)
  - Subs – 10-100M

  - Clients – 10s
  - TR Rate – V High
  - TR Size – Medium (>10kB)
  - IMDB/Server > 1TB
  - Latency – High ms
- Type 2
  - Large/Distributed (e.g. MME, HLR, PCRF)
  - Subs – 1-10M
  - Clients – 1
  - TR Rate – V High
  - TR Size – Small (>10kB)
  - IMDB/Server > 500GB
  - Latency – Low us

A front end application is responsible for managing access to memory locations across the servers. Also, these servers will need to provide a redundant copy for availability.

Intel Optane DC persistent memory can help to consolidate the database instances into a single footprint on a server, and simplify management and total cost of ownership for network operators by either referencing a single location for database access or by distributing databases across a number of servers.

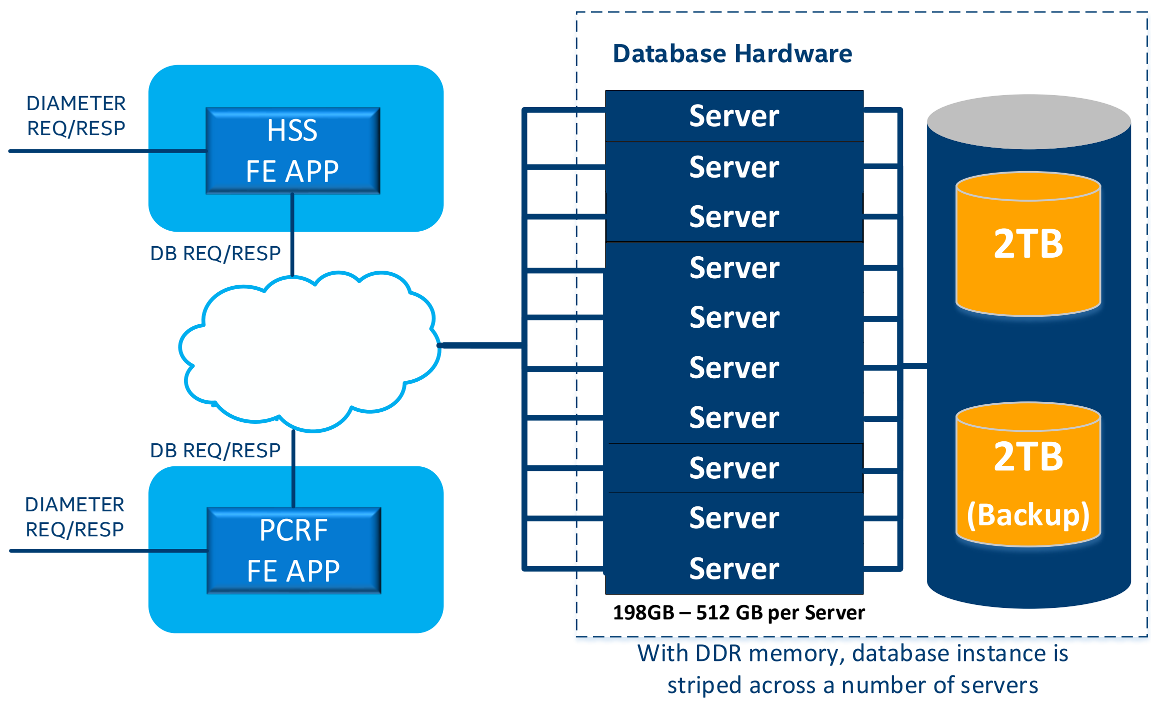

Figure 5. HSS Node Topology with DDR

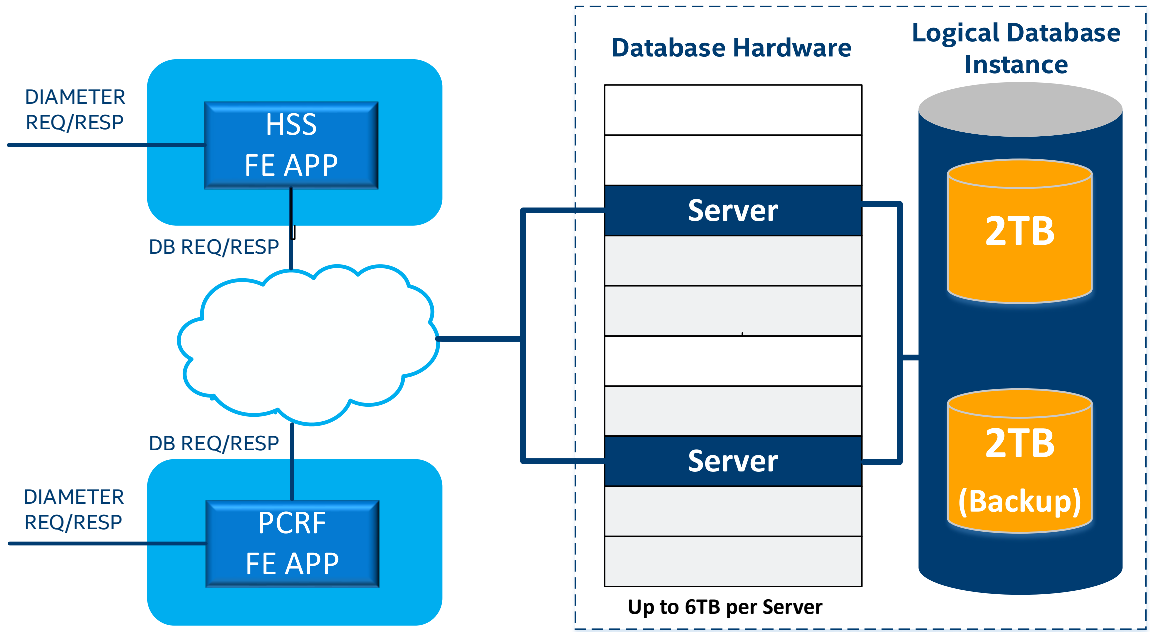

Figure 6. HSS Node Topology with Intel Optane DC Persistent Memory

Figure 6 describes a typical instance of a HSS node using Intel Optane DC persistent memory. Typically, subscriber data would be need to be striped across a number of servers as typical deployment only have between 200-500GB of DRAM memory per server, as shown in Figure 5. Using Intel Optane DC persistent memory, we can support up to 3TB per socket of DRAM-like memory (a second server is also required for resiliency). This provides a cost benefit by reducing the overall number of servers being used to support HSS function, as well as reducing the complexity of managing a database across a number of servers.

There are also a number of additional benefits to this topology:

- Operators can take advantage of persistent memory to help reduce downtime during failover.
- Utilize RDMA over PCIe to provide a coherent mapping of redundant databases and offload traffic from IO.
- Build a single low latency access point for multiple applications.
- Scale for greater database instances for 5G.

4.2 Content Delivery Network

In CDNs, the ability to provision content closer to the user for a faster, better user experience is critical. Intel Optane DC persistent memory can provide large memory capacity (up to 3TB per socket) at significantly lower costs than DRAM. It can also facilitate more streams per server when comparing against DRAM systems in equivalent price ranges.

For further details on this use case, please refer to the white paper:
https://builders.intel.com/docs/networkbuilders/intel-optane-dc-persistent-memory-content-delivery-networks-use-case.pdf

4.3 Artificial Intelligence

One of the hallmarks of Artificial Intelligence and Machine Learning is that it requires large datasets for training workloads. Many of the telecom use cases also require low latency access to datasets:

- Network Provisioning
  - Orchestration of an entire network which relies on huge volumes of telemetry
  - Requires fast decision making on where to deploy VNFs/VMs/Containers
  - Would benefit from large memory footprint to help in decision making
- Network Health Monitoring
  - OSS type infrastructure to monitor the overall system health
  - Identification of failures within the overall network
  - Troubleshooting and root causing of issue
- Data Management
  - Optimization of network traffic
  - Management of Quality of Service for customers
  - Power saving for radio interfaces

All of the above use cases require large datasets which could benefit from Intel Optane DC persistent memory at various different locations in the network, such as OSS, NGCO, and Orchestrator.

5 Subscriber Database Types

Subscriber databases account for around $2.5B of total addressable market in telecom networks per annum. Due to the size and complexity of in-memory database (IMDB) in 5G, operators are looking at simplifying how subscriber data is accessed by the application layer. The below models that can be used to deploy subscriber database servers. Intel Optane DC persistent memory can provide a cost and performance benefit for each of these models.

5.1 Single Database Instance

As shown in Figure 7, this model replaces the current infrastructure by replacing all of the existing database nodes into a single master database, which will manage all data. This will help to reduce multiple instances of data being replicated in multiple locations across the network, but may impact some of the applications that require lower latency access to data.

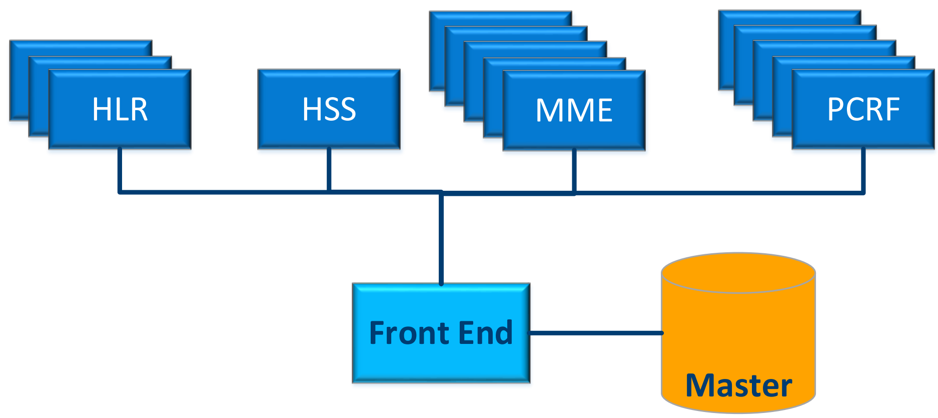

Figure 7. Single Database Instance

5.2 Distributed Database Instance

In this scenario, shown in Figure 8, databases are distributed across nodes in the network and are managed by a front end application, which will be responsible for accessing shared data across the network (i.e. subscriber data can be directly accessed directly from the HSS from the front end). This will help to meet requirements of low latency access, but may increase complexity of management across the network.

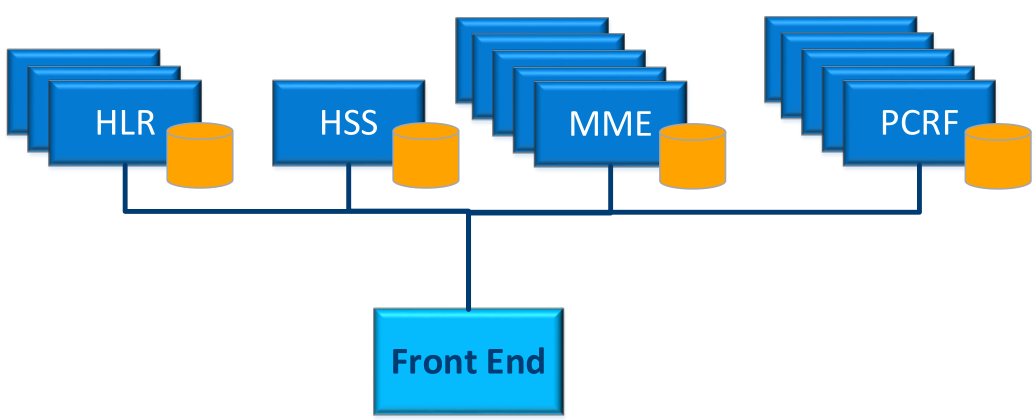

Figure 8. Distributed Database Instance

5.3 Hybrid Model

The Hybrid model shown in Figure 9 is similar to the one currently being deployed, where back end such as the HSS provide the backbone of the network, while other applications, such as MME, HLR etc. will supplement this data with data for more localized users.

Intel Optane DC persistent memory can help to consolidate instances onto single blades and help to reduce the overall management complexity of delivering this infrastructure in DDR.

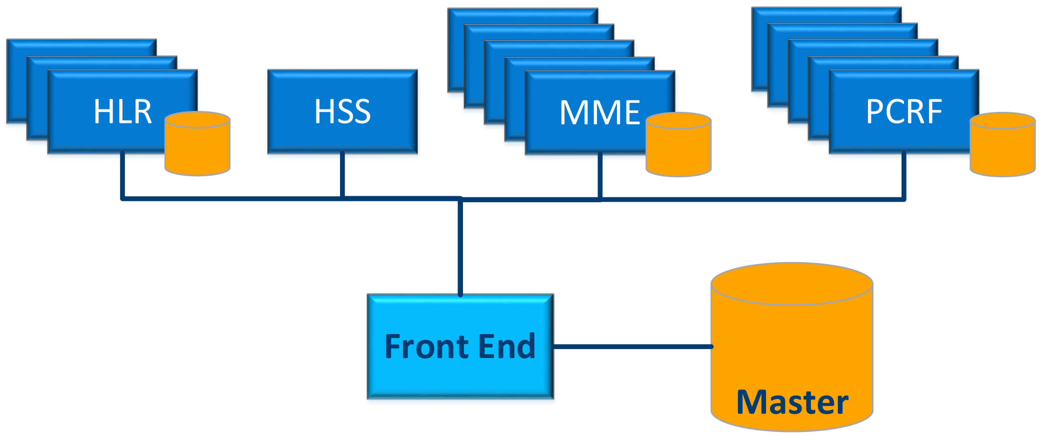

Figure 9. Hybrid Model

6 Summary

There are a wide variety of use cases within the telecom network that can take advantage of the high capacity and low latency performance of Intel Optane DC persistent memory. It provides a cost-optimal way to provide a large memory footprint for multiple applications, as well as providing a low latency access.

Intel Optane DC persistent memory also can help to reduce the overall complexity of having large clusters of compute blades with the sole purpose of providing striping of memory to a front end application. For AI use cases, which often require large memory footprints for telemetry etc., this provides a useful way to meet the memory applications.

The first generation of Intel Optane DC persistent memory is currently being deployed with the 2nd generation Intel® Xeon® Scalable processors (formerly Cascade Lake), and performance and capacity is expected to increase in the next generation.