

Intel® Optane™ DC Persistent Memory - Content Delivery Networks Use Case

Application Note

April 2019



You may not use or facilitate the use of this document in connection with any infringement or other legal analysis concerning Intel products described herein. You agree to grant Intel a non-exclusive, royalty-free license to any patent claim thereafter drafted which includes subject matter disclosed herein.

No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document.

All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest Intel product specifications and roadmaps.

The products described may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Copies of documents which have an order number and are referenced in this document may be obtained by calling 1-800-548-4725 or by visiting: <http://www.intel.com/design/literature.htm>

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Learn more at <http://www.intel.com/> or from the OEM or retailer.

No computer system can be absolutely secure.

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

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

Copyright © 2019, Intel Corporation. All rights reserved.



Contents

1.0	Introduction.....	5
1.1	Terminology	6
1.2	Reference documents.....	7
2.0	Using Intel® Optane™ DC persistent memory.....	8
2.1	CDN architecture	8
2.2	Configuring the server in memory mode.....	9
2.3	Test results.....	10
2.4	Bare metal and virtualized implementation.....	12
3.0	Intel® Select Solutions for Visual Cloud Delivery Networks.....	13
4.0	Summary	15

Figures

Figure 1.	CDN environment.....	8
Figure 2.	Testing environment.....	9
Figure 3.	Intel® Optane™ DC persistent memory topology.....	10
Figure 4.	Test results.....	11
Figure 5.	Intel® Select Solutions for Visual Cloud Delivery Network advantages.....	13
Figure 6.	Process for defining Intel® Select Solutions.....	14

Tables

Table 1.	Terminology	6
Table 2.	Reference documents.....	7



Revision History

Date	Revision	Description
April 2019	001	Initial release.



1.0 Introduction

A content delivery network (CDN) refers to a geographically distributed group of servers that work together to provide fast and efficient delivery of valuable content. Content that users need more frequently is typically moved physically closer to the user in regional or edge servers. Our society is moving from passive consumption of media to highly immersive and intelligent visual experiences. As CDNs continues to evolve to support emerging edge use cases, such as cloud gaming and virtual reality (VR), new approaches are needed to address more open, realtime, and cloud-based CDN solutions.

In CDNs, the ability to provision content closer to the user for a faster, better user experience is critical. Operators have turned to faster storage technologies (such as PCIe storage devices). Increasingly, larger memory subsystems are needed to host the most popular content directly in memory to make it even faster and more accessible to stream out to numerous customers. A rapidly growing percentage¹ of Internet content is now live linear content, which has no reason to be stored since it is immediately distributed to the end user. Large memory capacity is needed to buffer each of the many independent streams a server needs to handle.

Unfortunately large capacity memory DIMMs are expensive. Intel® Optane™ DC persistent memory can provide large memory capacity (up to 3TB per socket) at significantly lower costs (approx. 20-40% savings at 1.5 TB) than DRAM. It can also facilitate more streams per server when comparing against DRAM systems in equivalent price ranges. Intel® Optane™ DC persistent memory enables a rethinking of the memory/storage hierarchy to deliver new levels of cost-effectiveness and scalability for memory-dependent use cases such as CDNs.

Now there are two options to expand local memory in the server: 100% DRAM or a mix of small DRAM cache with Intel® Optane™ DC persistent memory as bulk memory. For live linear use cases and for the hottest VoD use cases, Intel® Optane™ DC persistent memory delivers similar performance as DRAM while meeting quality of service (QoS) requirements at significantly lower total memory system cost.

Conversely, for the same price, extra memory capacity can be deployed. This extra memory can be used to host additional video streams (e.g. users) on the same server because each stream requires a pre-defined chunk of memory dedicated to that stream. With scaling simple linear video play-through, busy CDN servers tend to run out of memory before they run out of CPU and/or IO resources.

¹ Source: Cisco Visual Networking Index: Forecast and Trends, 2017–2022, Nov. 28, 2018, <https://www.cisco.com/c/en/us/solutions/collateral/service-provider/visual-networking-index-vni/white-paper-c11-741490.html>



Benefits for Intel® Optane™ DC persistent memory include:

- Scale memory capacity at significantly lower costs than DRAM.
- Enable more streams per server compared to DRAM systems at the same price.
- Consolidate live linear and VoD nodes.
- Lower latency interaction than with SSD for most popular (hottest) VoD content
- Non-volatile memory behavior (where data persists when power is shut off)

Note: For live linear streaming CDNs use models, we recommend the use of Memory Mode (volatile).

Intel® Optane™ DC persistent memory addresses memory constraint challenges for popular and latency-sensitive use cases such as live linear video streaming in CDNs. This document explains how to apply Intel® Optane™ DC persistent memory in CDNs.

This document is part of the Network Transformation Experience Kit, which is available at: <https://networkbuilders.intel.com/>

1.1 Terminology

Table 1. Terminology

Term	Description
1LM	One-level memory
2LM	Two-level memory
App Direct mode	1 of 2 modes for Optane DCPMM – enables memory persistence and requires code enabling
CDN	Content Delivery Network
Hot VoD	Hottest (most popular) Video on Demand content
Live linear video	Video that is streamed either live or as linear (not stored) broadcast
Memory Mode	1 of 2 modes for Optane DCPMM – behaves as cached system memory (e.g. volatile) without requiring code changes but offering huge capacity boost
NGCO	Next Generation Central Office – an aggregation point of wireless and wireline data that resides at the network and utilizes industry standard servers plus virtualization to provide flexibility and scalability while delivering performant and efficient services
Persistent memory	Non-volatile memory – retains data throughout power loss
vCDN	Virtualized CDN – Content Delivery Network function implemented virtualized in a virtual machine or containers
VoD	Video on Demand



1.2 Reference documents

Table 2. Reference documents

Document	Location
Persistent Memory on Intel Developer Zone	https://software.intel.com/en-us/persistent-memory
Intro video: Intel Optane DC Persistent Memory Fills the Gap between DRAM and SSDs	https://youtu.be/f9plXw1ndRI
Intel® Optane™ DC persistent memory for CDNs	https://www.intel.com/content/www/us/en/architecture-and-technology/optane-dc-persistent-memory.html
Intel® Select Solutions	https://builders.intel.com/intelselectsolutions
Intel® Select Solutions for Visual Cloud Delivery Network	https://builders.intel.com/docs/networkbuilders/intel-select-solutions-for-visual-cloud-delivery-network.pdf
Intel® Optane™ DC solid state drive (SSD) Product Brief	https://www.intel.com/content/www/us/en/products/docs/memory-storage/solid-state-drives/data-center-ssds/optane-ssd-dc-p4800x-p4801x-brief.html
Creating the Next Generation Central Office White Paper	https://www.intel.com/content/www/us/en/communications/network-transformation/next-generation-central-office-paper.html

2.0 Using Intel® Optane™ DC persistent memory

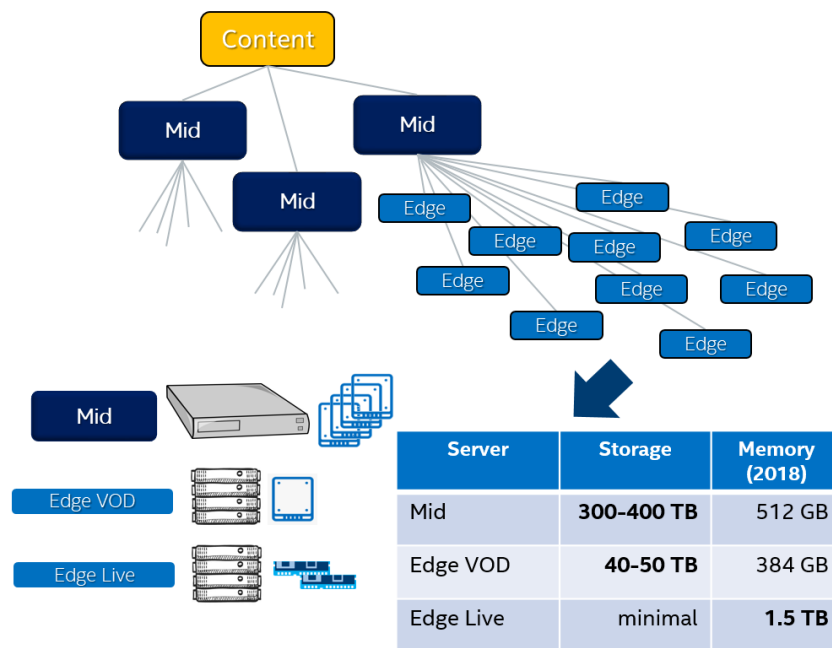
This section describes how to apply Intel® Optane™ DC persistent memory in your content delivery network (CDN).

2.1 CDN architecture

Industry practice² has shown that the hottest VoD content under extreme demand is best cached in DRAM, not storage, to maintain high quality of service (QoS). Classic storage-based CDN architectures are challenged to meet this demand because they need large amounts of memory instead of storage, but DRAM is costly. The solution is to increase memory capacity using Intel® Optane™ DC persistent memory to save total memory cost compared to classic DDR4 DRAM alone.


Figure 1 shows a simple diagram of a typical CDN high performance computing environment, which consists of local compute nodes, storage, and intermediate I/O nodes. The Server, Storage, and Memory data was obtained from Tier 1 CDN vendors serving as lead market definition partners.

Figure 1. CDN environment



The figure above shows a typical CDN architecture. Given the large DRAM needed in the Edge Live server, this becomes a good candidate for Intel® Optane™ DC persistent

² Tier 1 CDN vendor(s) serving as lead market definition partner(s).



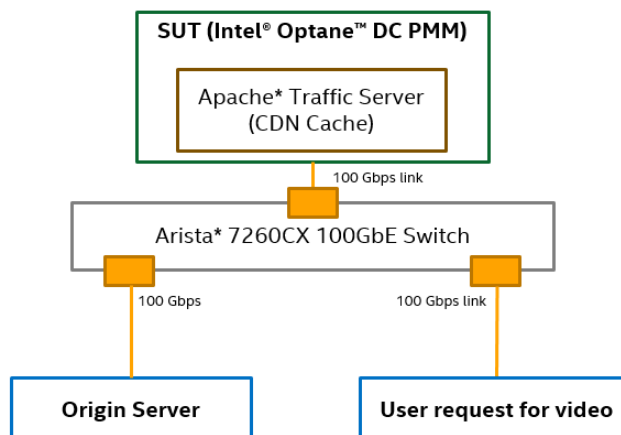
memory. As the market proceeds forward, the other memory footprints will also increase and more of the CDN solution becomes a candidate for this new memory technology. Some customers are exploring the combination of Edge Live and Edge VOD into a single server.

2.2 Configuring the server in memory mode

Figure 2 shows the setup we used for our benchmarking. The test performed the following operations:

- User request for video (Live Linear) was sent at 10 Gbps, 20 Gbps, 30 Gbps, 40 Gbps, up to 50 Gbps HTTP-GET request to the CDN server.
- If the content requested was missing from the CDN server, a request was made to the origin server.
- If the content requested was in the CDN server, it provided the content directly to User to satisfy the request. The CDN Server is preconditioned to address 80/20 hit ratio.

Figure 2. Testing environment



Hardware

The System Under Test (SUT) was Intel® Server Board S2600WF (formerly codenamed Wolf Pass) which includes:

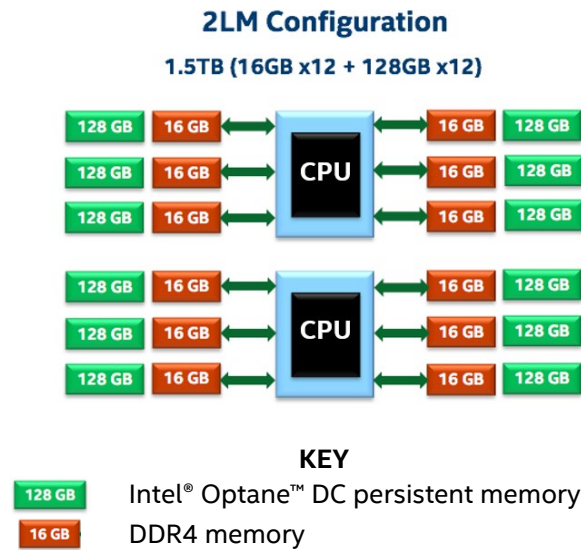
- 2nd generation Intel® Xeon® Scalable processor (formerly Cascade Lake)
25 Intel® Xeon® 6252 Gold 24C
- 12x 16GB DDR4 2667 MHz
- 12x 128GB Intel® Optane™ DC persistent memory module



- 2 Add in Cards Intel® Ethernet Network Adapter XXV710-DA2 populated NUMA aligned.
- 4x 25 Gbps Network ports are configured for 100 Gbps LAG (Link Aggregation Group)

The following figure shows the 2-2-2 Intel® Optane™ DC persistent memory topology for both sockets. For reference, the 1LM configuration was 1.5TB DDR4 (64GB x24).

Figure 3. Intel® Optane™ DC persistent memory topology



Software

Apache traffic server (ATS) and NGINX are Layer 7 applications that can be configured to serve HTTP-GET or HTTPS-GET or as reverse proxy web servers.

We used the following open source software in our benchmarking scenario:

- Apache traffic software: <https://trafficserver.apache.org/downloads>
- NGINX: <https://www.nginx.org>

2.3 Test results

Figure 4 shows the results of the benchmarking tests. The data measures Time to First Frame (TTFF), which is an indicator of the client experience, and Number of Requests, which indicates the number of requests the CDN server can support. The performance data shows that the compared servers provide the same capacity and similar performance, which leads to cost savings when using Intel® Optane™ DC persistent memory.

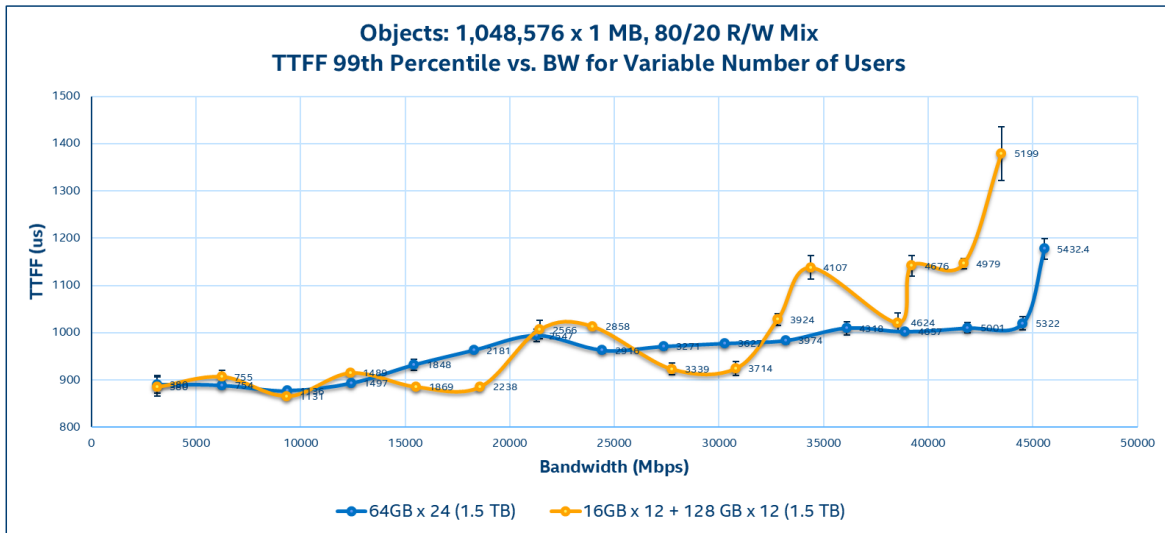


The results indicates that CDNs that leverage Intel® Optane™ DC persistent memory can perform comparably to existing DRAM-only based CDNs. In particular, CDN servers populated with Intel® Optane™ DC PMM can meet an SLA of 99th percentile HTTP GET response times of less than 1.5s for 4 Mbps HD content in 2s chunks.

Figure 4. Test results

Live Linear 80/20 – Time To First (video) Frame (TTFF)

Target: HTTP GETs 99pct < 1.5s for 4Mbps HD content in 2-second chunks



¹TTFF – Time To First Frame, derived from TTLB
Live Memory Cache Size: 1.1TB
Video Object Size: 1MB avg

480 concurrent connections at 50Gbps with 2M hugepages (64/node)
playlists (1.1TB memory +10%) distributed over 16 clients random/port

Note:

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 information go to www.intel.com/benchmarks.

Performance results are based on testing as of 2/6/2019 and may not reflect all publicly available security updates. See configuration disclosure for details. No product or component can be absolutely secure.

§ Configurations: CentOS Linux release 7.5.1804 (Core) 4.19.0-rc3+ (Host) ATS v7.1.4, BIOS: SE5C620.86B.0D.01.0250.112320180145 Test used: 80/20 Live Linear, tested by Intel, Date of Testing: 2/6/2019 (comprehensively)



2.4 Bare metal and virtualized implementation

The test results in the previous section were measured in a bare metal environment, however, Intel® Optane™ DC persistent memory also supports virtualized (vCDN) implementations. For data-intensive workloads, a bare metal server provides maximized processing power and network performance. Virtualized servers are scalable, customizable, and flexible, which is ideal for deployments such as NGCO.

Learn more about NGCO:

- White paper: <https://www.intel.com/content/www/us/en/communications/network-transformation/next-generation-central-office-paper.html>
- Video of vCDN value propositions in next generation central office (NGCO) installations: <https://youtu.be/ZWY6r3B9HMY>



3.0 Intel® Select Solutions for Visual Cloud Delivery Networks

The Intel® Select Solutions for Visual Cloud Delivery Networks consist of optimized hardware resources and an open-source software stack residing within a virtualized infrastructure. The solution stack leverages the most common and popular open source CDN caching frameworks such as NGINX* and Apache Traffic Server (ATS). It also leverages open source media libraries such as FFmpeg, Media Service Studio, and Scalable Video Technology for media transcoding. Acceleration is built into the system for key CDN workload functions such as cryptography, data compression and transcoding. Furthermore, the CDN configuration utilizes new memory and storage solution options to improved scalability, reduced latency and cost savings.

Intel® Select Solutions are pre-defined and pre-verified configurations designed to help address these challenges.

- Intel® Select Solutions are tightly specified in terms of hardware and software components to eliminate guesswork and speed up decision-making.
- Intel® Select Solutions are designed to increase efficiency in the testing process, speed up time to service delivery, and increase confidence in solution performance.
- Intel® Select Solution configurations are designed by Intel and our partners to deliver to a performance threshold for the workload, and are built on the latest Intel architecture foundation including the 2nd generation Intel® Xeon® Scalable platforms.

Figure 5. Intel® Select Solutions for Visual Cloud Delivery Network advantages

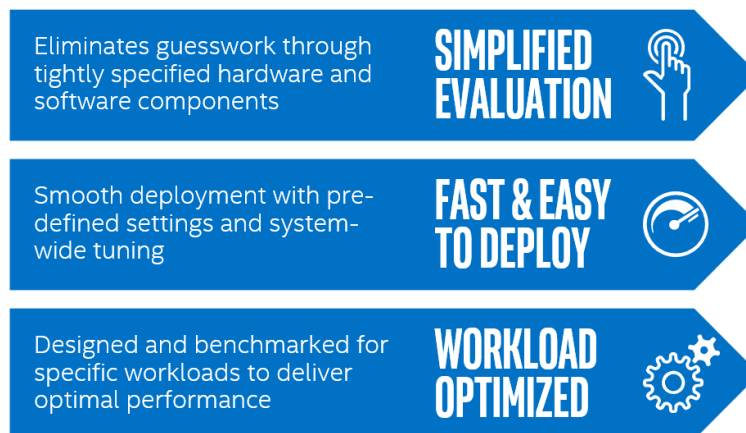
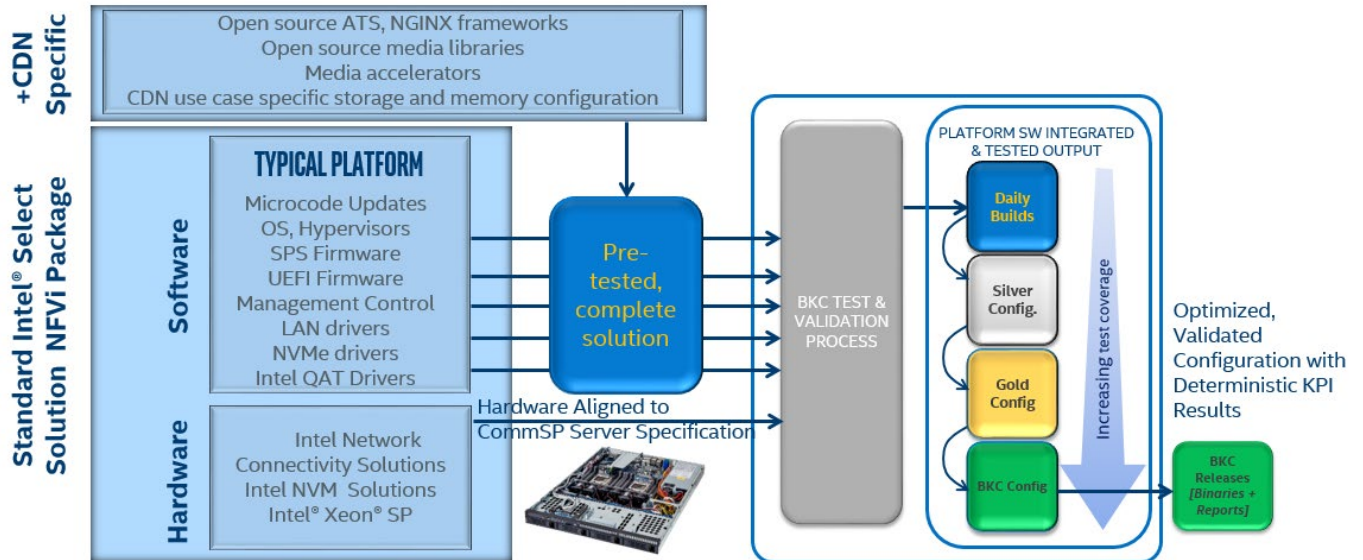


Figure 6 shows the process behind the optimized and well defined foundation, which starts by collaborating with service providers and ecosystem. The configuration is tailored for specific CDN workloads. It is tested with open source caching framework and leverages Intel internal testing throughout the lifecycle of product. The whole recipe is easy to reproduce and partners can follow the build recipe and certify their solution quickly.

Figure 6. Process for defining Intel® Select Solutions



The result of the process is the following:

- Tuned platform designed for the unique needs of Visual Cloud Delivery Network use cases with the following features:
 - Second generation Intel® Xeon® Scalable processor based servers with dense, form-factor optimized NVMe SSDs
 - Intel® Optane™ technology solutions for performance and cost benefits
 - Network throughput with support up to 100Gbps
 - NUMA Balanced I/O for consistent latency and maximum throughput
 - Acceleration for media transcode and crypto/compression to handle these increasingly common and demanding operations
- Open software stack optimized for highly efficient, high quality visual workloads
 - Ensures platform scalability with open source software and compliance to standard industry frameworks
 - CentOS* hardened, KVM hypervisor, optimized & validated solution stack including media transcode libraries
- Delivered by Intel and trusted vendors as tested, hardened platform solutions
 - Abstracts the HW/SW complexity typical of open platforms with vigorous stability and performance testing
 - Accelerates deployment and commercialization of innovative CDN services



4.0 Summary

Intel® Optane™ DC persistent memory helps to address memory constraint challenges for popular but latency-sensitive use cases such as live linear video streaming in content delivery networks (CDNs). Test results shown in this document prove this new memory type provides comparable performance, while saving significant cost when compared to classic DRAM at these capacity sizes. Intel® Optane™ DC persistent memory supports both bare metal and virtualized (vCDN) scenarios, enabling you to choose the right configuration for your application. For rapid time-to-market, Intel offers pre-validated reference solutions in the Intel Select Solution for Visual Cloud Delivery Networks.