White Paper

Video Service Providers, Broadcasters, Communications Service Providers

Broadpeak Improves Caching, Cloud DVR Performance with SPDK

Broadpeak uses Storage Performance Development Kit to leverage fast, low latency PCIe Gen5 solid-state drives; tests[i] two of its content delivery solutions to highlight performance achievements

Video service providers are modernizing their infrastructure to serve more users with fewer servers. This involves creating denser systems that offer higher network and storage capacity while minimizing rack-space, energy usage, and cooling requirements. Intel® Xeon® Scalable CPUs and PCIe Gen5 SSDs, which deliver significantly improved read and write performance compared to PCIe Gen4 SSDs, are essential components for building these high-density systems.

I/O-oriented applications such as read-heavy content delivery networks (CDNs) and write-intensive recording servers in CloudDVRs are ideal candidates to benefit from high-density hardware configurations. Nevertheless, the significant performance gains offered by modern SSDs can be hindered by inefficiencies in traditional software storage stacks.

These software stacks, written with conventional synchronous programming models, fail to leverage the full potential of high-throughput SSDs due to lock contention, unnecessary memory copying, synchronous / blocking IO operations, etc. To address these limitations, Broadpeak incorporated Intel-developed open source Storage Performance Development Kit (SPDK) into its High Performance Cache (formerly known as BkS450) and cDVR solutions.

In this paper we'll first introduce Broadpeak solutions. Subsequently, we'll examine the specific challenges faced by I/O-intensive applications and how the SPDK programming model can be a potential remedy. Finally, we'll present the results and benefits of integrating SPDK into Broadpeak's software stack.

Table of Contents

Broadpeak High Performance Cache and Cloud DVR Solutions .. 1

High Performance Cache 1

Cloud DVR 2

CDNs Powered by 5th Gen Intel® Xeon® Scalable Processor 3

IO Challenges in Content Delivery 3

SPDK to the Rescue 3

Performance Achievements 3

Conclusion 5

Learn More 5

Broadpeak High Performance Cache and Cloud DVR Solutions

High Performance Cache

Broadpeak High Performance Cache is a high performance cache management and streaming engine that is deployed at edge compute sites, reducing contention for the origin server and offering high streaming capacity[ii] to respond to the increasing demands of end-users (see Figure 1).

The High Performance Cache supports HTTP adaptive streaming, HLS, Smooth Streaming, HDS, and MPEG-DASH protocols. Unlike standard HTTP cache servers that have no knowledge of the content they are delivering nor the recipients they are delivering content to, the High Performance Cache is fully aware of video characteristics and viewer profiles. With this information, the High Performance Cache can optimize the video delivery and build a consistent video session report for use by the video service provider.

The High Performance Cache supports HTTPS for secure video delivery and uses token-based access and an anti-deep-linking system to protect against unauthorized access.

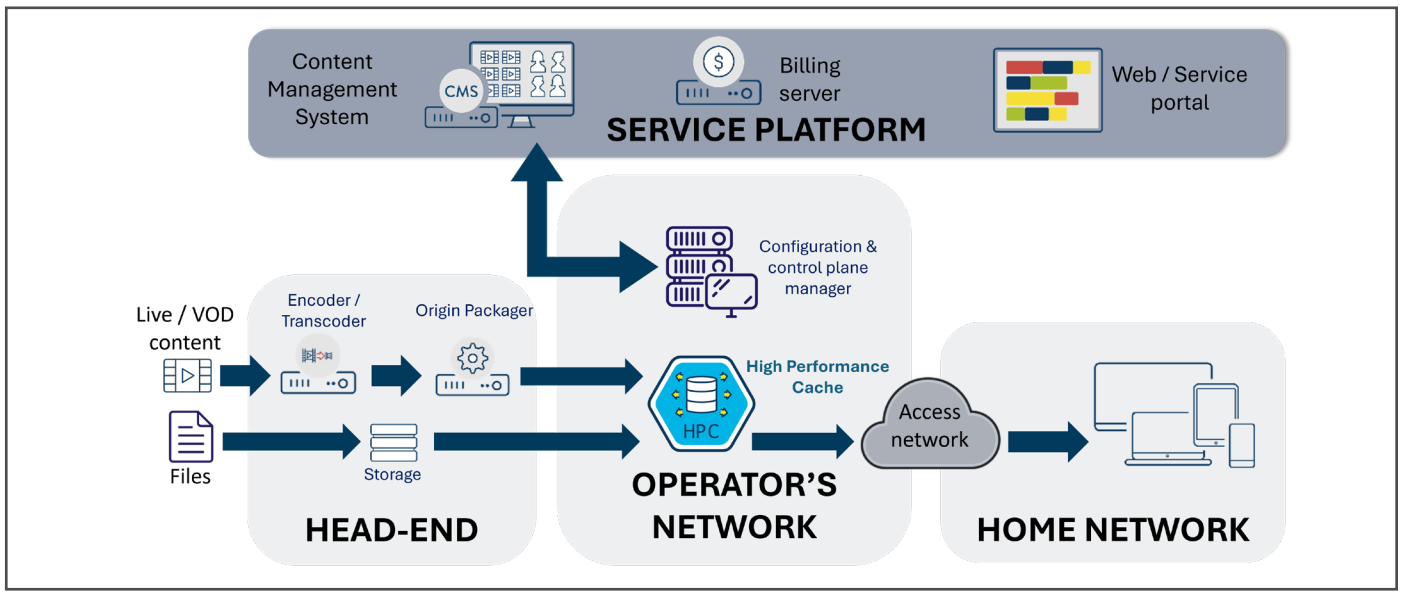

Figure 1. Block diagram shows how the Broadpeak High Performance Cache works in operators' networks to deliver content with low latency and minimal impact on the origin server.

Cloud DVR

The Broadpeak Cloud DVR solution includes a live origin server, a packager and a recording manager as key components to support cloud DVR features (see Figure 2). The solution is scalable thanks to its microservice architecture, enabling service providers to separately and automatically scale services modules according to customer demand.

The solution supports major streaming formats including HLS, DASH, and CMAF Low Latency with all of the leading audio, video, and subtitle codecs with advanced advertisement management.

The Broadpeak Cloud DVR solution has built-in cache that offers scalable storage from a few terabytes to hundreds of petabytes, with support for scale-out NAS and software-defined storage. Broadpeak Analytics is a complete cloud DVR statistics suite, and the entire system is compliant with all multiscreen devices, including the latest Android TV and Apple devices, and is integrated with major CMS (Content Management Systems) solutions, DASH, and CMAF.

The Recording Manager component oversees the copy policy and the scheduling of the recording. With rolling buffer and automatic archiving functions, service providers can mutualize short-term catch-up TV and long-term cloud DVR.

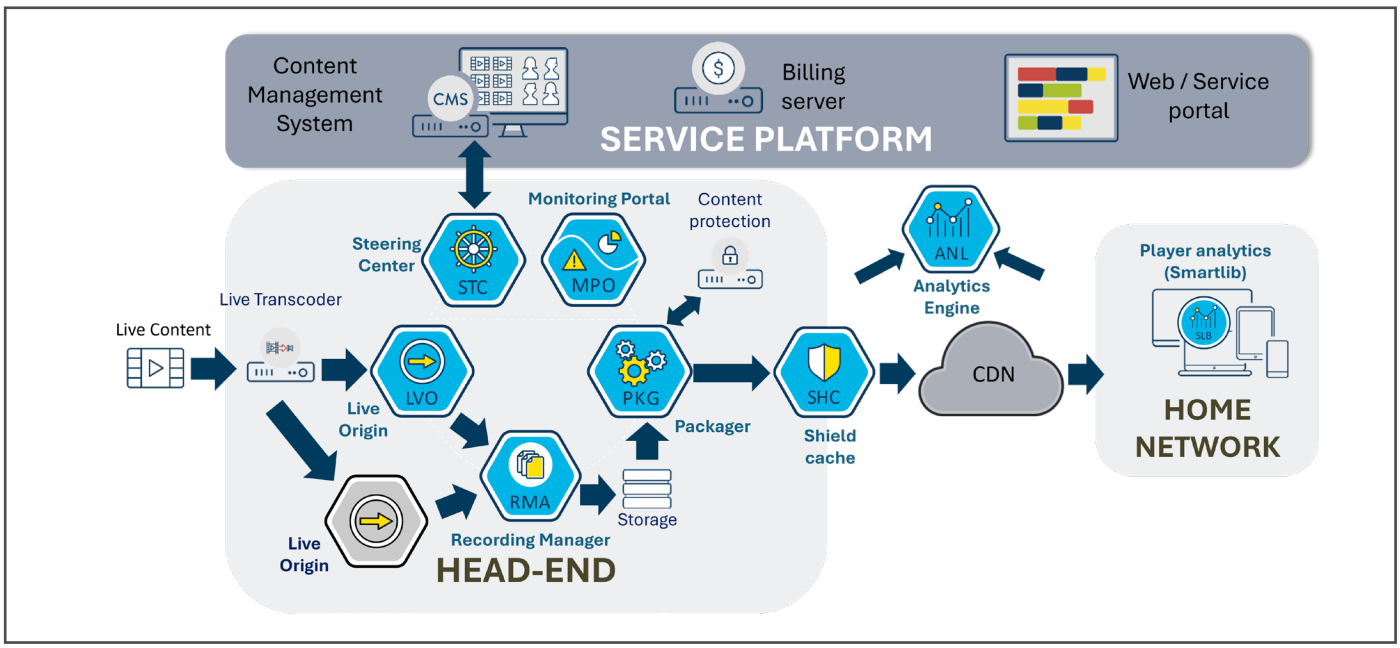

Figure 2. Broadpeak Cloud DVR solution and its components.

CDNs Powered by 5th Gen Intel® Xeon® Scalable Processor

5th Gen Intel® Xeon® Scalable processors deliver exceptional performance, efficiency, and scalability, making them ideal for handling demanding CDN workloads. This processor family has models with up to 80 lanes of PCIe Gen5 connectivity for faster data transfer rates between the processor and connected devices, such as storage devices.

Support for DDR5 memory improves compute performance by overcoming data bottlenecks with higher memory bandwidth improving performance, capacity, power efficiency, and cost. The processors feature up to 64 cores and base frequency up to 3.6 GHz enabling parallel processing of multiple tasks and improving overall performance.

IO Challenges in Content Delivery

Video access patterns in CDN and Cloud DVR systems differ from those assumed for traditional filesystems. Also, Linux filesystems provide storage capabilities that enable simultaneous read and write access to files and allow partial read/partial write, which introduces unnecessary processing overhead for content delivery applications. This overhead adds significant processing latency for I/O intensive workloads.

Some caching systems have shifted away from the traditional approach of treating each cached object as a separate file. Similarly, recording systems often utilize large, fixed-size files with an application-specific internal layout. This approach introduces several challenges, including ensuring data durability in the event of a disk / hardware failure, minimizing overhead imposed by the filesystem due to synchronous I/O operations, and enabling software stacks to fully exploit the performance capabilities of PCIe Gen5 SSDs, which require optimizing for higher throughput and lower latency to avoid bottlenecks.

SPDK to the Rescue

To alleviate these issues, Broadpeak turned to SPDK for the storage backend for the Broadpeak High Performance Cache and Cloud DVR solutions.

SPDK makes it easier for ISVs to develop fast storage by providing a set of tools and libraries for writing high performance, scalable, user-mode storage software. SPDK is designed to help application developers reduce the transaction time needed by the storage stack in order to match the high throughput and low latency offered by SSDs.

SPDK drivers bypass the OS kernel and run in Linux user space. By running the device driver code in user space, the driver avoids kernel context switching and interrupts that require a lot of processing overhead. These CPU cycles can now be used for storing the data, which means better performance and latency.

The other performance feature of SPDK is the use of poll mode drivers (PMDs). Using PMDs addresses the OS interrupt challenge by changing the way that I/O is handled. With PMDs, an application submits the request for a read or write, and then checks back at a defined interval once the I/O has been completed. This avoids the latency and overhead of using interrupts and allows the application to improve I/O efficiency.

SPDK can handle millions of IOPS, saturating high-performance storage devices, with far fewer CPU cores utilized and without the use of additional offload hardware (see Learn More section for a link to the report).

To determine the impact of SPDK on two content delivery solutions, Broadpeak, an Intel® Industry Solution Builders member, worked with Intel to test[i] its software performance with SPDK enabled for two services.

Performance Achievements

The ability of Broadpeak software to leverage the high throughput of PCIe Gen5 SSDs is demonstrated in the two tests conducted by Broadpeak on servers powered by 5th Gen Intel Xeon Scalable processors.

Test 1: SPDK in Broadpeak High Performance Cache

This test showcases the resilience of the Broadpeak High Performance Cache CDN stack starting from cold cache. The objective is to achieve high-throughput content delivery while simultaneously populating the cache from the origin (storing data on disk or in memory), all without compromising response times.

Traffic Pattern / Test Conditions

The testbed is made up of 10,000 HTTPS clients that randomly request 50 million objects, 1MB each, following a uniform distribution. The test is run until the cache is filled to 25TB of cached data, starting from an empty state (0 bytes).

Test Setup

Broadpeak, in collaboration with Intel, has developed the Broadpeak Benchmarking Tool (details available in the "Learn More" section). This tool is deployed on five load-generator servers to simulate requests to the Broadpeak High Performance Cache via a 32-port 400 Gbps switch.

In order to avoid being bottlenecked on the origin server side, the test is set up for the Broadpeak High Performance Cache to generate and modify content internally, placing it directly into its cache. This enables testing without being limited by throughput constraints on the origin side.

Test Results / Analysis

The test systems successfully cached 25 million objects, 1MB each, in 6 minutes and 15 seconds. Figure 3 illustrates the NVMe I/O performance and cache TCP egress throughput during the process.

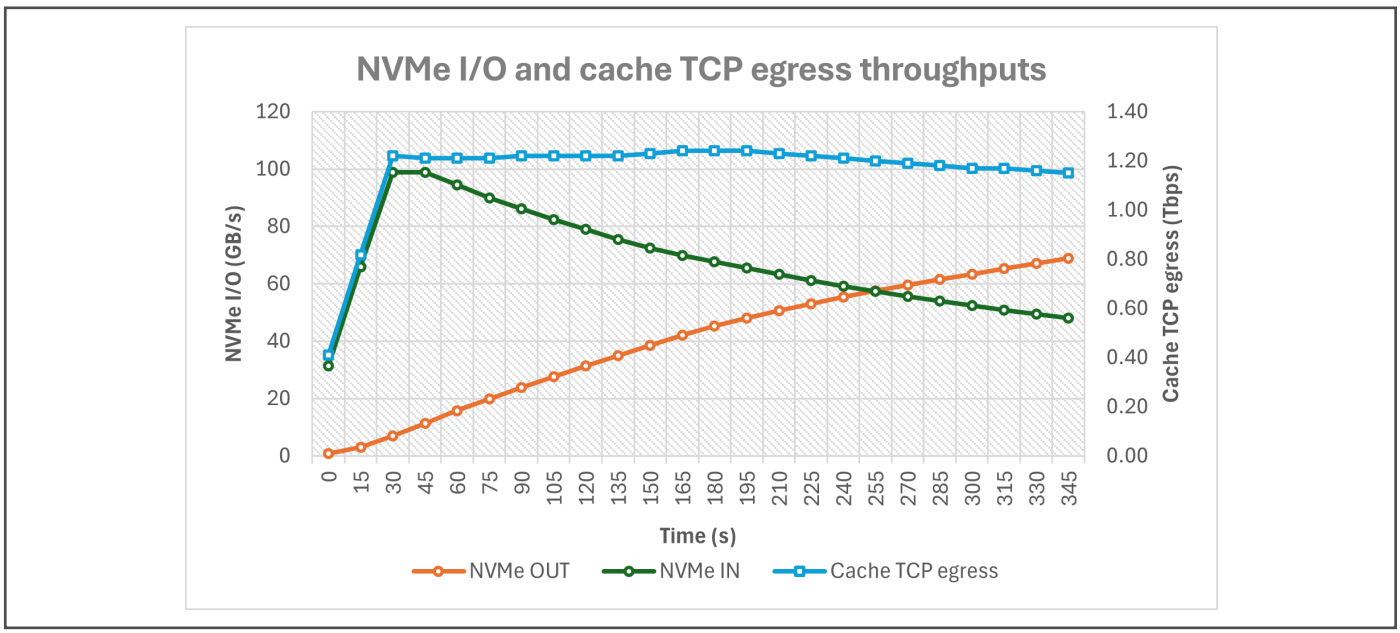

Figure 3. NVMe I/O throughput and cache TCP egress over time.

Figure 4 depicts the percentile P99 and P95 serving times.

The Broadpeak High Performance Cache achieved a peak content ingress of 100GB/s, while simultaneously delivering an impressive 1.2Tbps of egress data over the network. This demonstrates the ability of the High Performance Cache to efficiently serve content at high speeds while simultaneously populating its storage from origin, without compromising user experience as shown in Figure 4 where the High Performance Cache maintained P99 service times under 20 ms. Leveraging SPDK, the High Performance Cache storage/IO operations are optimized to deliver the highest possible CDN throughput within available CPU resources.

Test 2: SPDK used in Broadpeak CloudDVR Service

The goal of this test is to demonstrate that the use of the SPDK in the Broadpeak CloudDVR stack enables maximum user density per DVR server as shown by the efficiency in handling recording requests from a large number of concurrent live streaming users.

Traffic Pattern / Test Conditions

The test setup has one Live Origin instance that is co-located with the Recording Manager software on the same SUT. The tests involve recording a separate copy of a live stream for each user ("private copy" model).

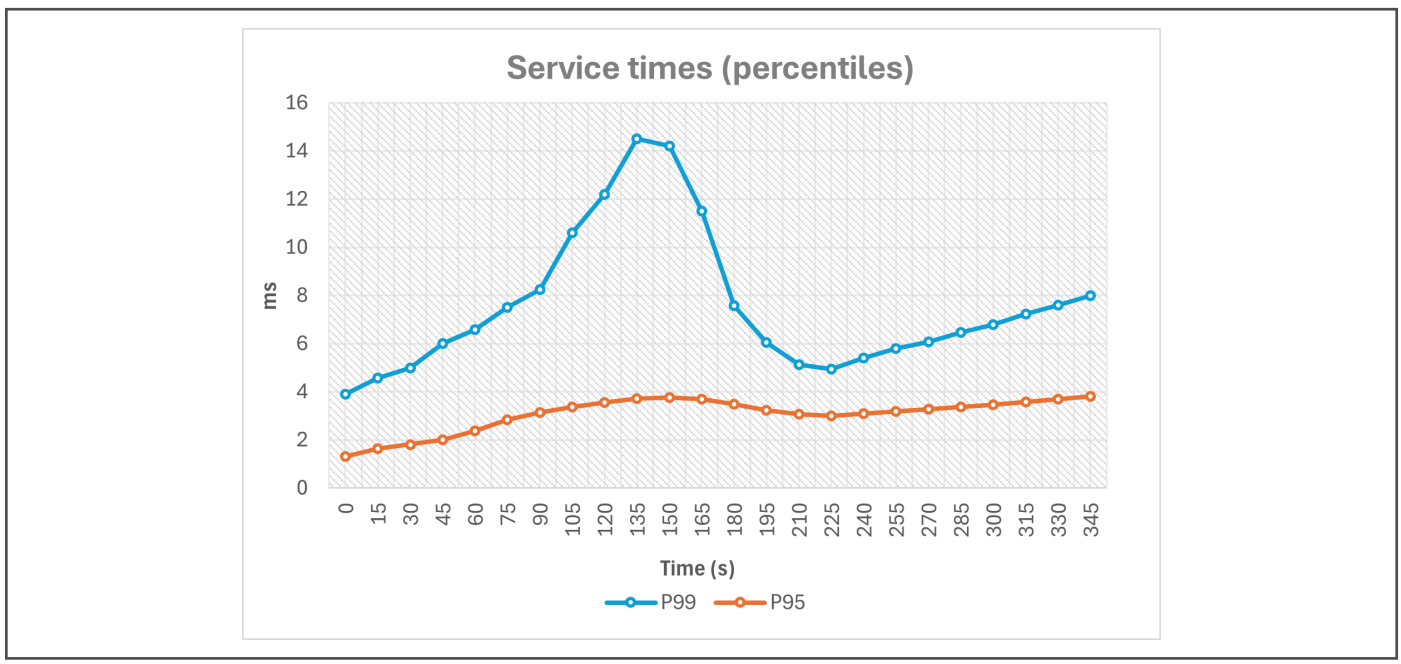

Figure 4. P99 and P95 service times over time.

The tests use two live stream formats/patterns:

- Low-resolution 48Kbps stream and 11 million concurrent private copies
- High-resolution 20 Mbps stream and 600,000 concurrent copies

The IOPS and GB/s limits for these formats are as in Table 1.

	IOPS	Throughput (GB/s)
48 Kbps live	11.3 million	46
20 Mbps live	600,000	83

Table 1. Write performance of Broadpeak Cloud DVR solution with SPDK.

Test Results / Analysis

Efficient storage processing is crucial to recording as many concurrent copies as possible per server. With the use of SPDK, Broadpeak was able to get close to the underlying hardware processing / IO limits, and record millions of private copies with low-resolution streams (small IO transactions, CPU intensive), and up to hundreds of thousands of items with high-resolution streams (larger IO transactions, storage intensive).

Conclusion

The availability of PCIe Gen5-based SSDs is providing content delivery systems with extremely fast I/O rates just at a time when video service providers need to improve the efficiency of their networks.

To fully harness the performance of high performance storage and network hardware, software needs to be redesigned. Intel's SPDK framework is a key enabler toward this goal as evidenced by the test results presented in this paper.

For the Broadpeak High-Performance Cache, the tests demonstrate consistent and reliable I/O performance, even under scenarios with unexpectedly low cache hit ratios requiring frequent origin pulls and a high write transaction rate. This determinism allows the High Performance Cache to deliver a superior user experience. For the Broadpeak CloudDVR solution, tests reveal significant improvements in user density per server. This translates to a reduction in the number of servers required to process large number of concurrent recording requests, and thus reduced operational costs and a smaller environmental footprint.

Learn More

Broadpeak Homepage

Broadpeak Builds No-Compromise CDN Benchmarking Tool

SPDK.io

SPDK NVMe-oF TCP (Target & Initiator) Performance Report

5th Gen Intel® Xeon® Processors

Intel® Industry Solution Builders

[i] SUT: Lenovo SR650v3 powered by dual Intel® Xeon® 6538N CPUs each with 32 cores and 64 threads. Total DDR5 memory was 512GB GB (16 slots/ 32GB/ 4800 MHz); microcode 0x21000230; Intel® Hyper-Threading Technology enabled; Intel® Turbo Boost Technology enabled. BIOS version: ESE124B-3.11. Application storage: 16x Kioxia CM7-R 3.84T; Network controller: Three third-party 400GbE Ethernet adapters and two 100GbE adapters. Software: OS was Archlinux; kernel was 6.9.10-arch1-1. Benchmark software: Broadpeak Benchmarking Tool: ref. 36baf1719f2588505d56d8b3a2fb3fa884b91653. Workload software: Broadpeak High Performance Cache (version 01.00.01) and Broadpeak Cloud DVR (version 01.05.09). Compiler was gcc 14.2; Libraries were SPDK (24.05). Test conducted by Broadpeak on 2024-07-15 for cDVR, 2024-07-25 for High Performance Cache.

[ii] https://broadpeak.tv/newsroom/bks450-streaming-software-passes-1tbps/. The private DVR is a write-mostly use case. This is why only disk ingress is measured. Both IOPS and GB/s throughput metrics are measured at the SPDK level.

Notices & Disclaimers

0225/CG/DJA/PDF Please Recycle 363890-001US