

Virtualized Outside Broadcast Van at the Olympic Winter Games 2022

Intel, as the Official Processor Partner of the Olympic and Paralympic Games, and the Olympic Broadcasting Services (OBS) conducted a Proof of Concept (PoC) of an open reference architecture for a software-defined outside broadcast van using COTS hardware. This model has the potential to reduce costs and complexity for live event video production while also improving flexibility and scalability.

Table of Contents

Reference Architecture Proof of Concept (PoC) at the Olympics ...	1
PoC Objectives	2
Venue	2
Production Technical Details	2
SMPTE ST 2110 Suite of Standards	2
Production Format and Standards	3
Video	3
Audio	3
Sources	3
Processing Application Pillars	3
Architecture	4
Control and Orchestration	4
Software Architecture	4
Rack and Network Architecture	4
Intel® Architecture Building Blocks for Video Production	5
Intel Server System M50CYP Family	5
Intel® FPGAs	6
Intel GPU	6
Intel® Network Adapters and the Intel® ST 2110 Streaming Library	7
Intel® Virtual RAID on CPU (Intel® VROC)	7
Intel® Silicon Photonics	7
Intel® Tofino™ Series Programmable Ethernet Switch ASIC	7
Conclusion	7

Live video production is transitioning from its historic dependence on inflexible, costly proprietary standards and equipment to a more open model based on commercial off-the-shelf (COTS) hardware. Adopting an IP software-defined deployment model for video production enables the industry to benefit from innovations from the larger information and communications technology (ICT) arena. In doing so, solution architects for video processing realize advantages in flexibility, scalability, and cost by replacing proprietary infrastructure with general-purpose IT technologies. Joining OBS's expertise in live video production and video production technology with Intel's expertise in deploying compute and networking platforms common across the ICT landscape, we've been able to build a fully functional "virtual OB."

The Society of Motion Picture and Television Engineers (SMPTE) ST 2110 standards for IP-based video transmission catalyze the transition of video production from broadcast-specific serial digital interface (SDI) to open architectures. This transformation also enables remote, centralized, or distributed cloud-based video production workflows that help reduce the cost of creating video assets. The updated approaches enable a single team to produce multiple events within a time-constrained window, eliminating the need to send dedicated production crews and proprietary equipment to live event venues or to outsource production tasks to third-party providers.

Reference Architecture Proof of Concept (PoC) at the Olympics

The ongoing collaboration between Intel and the Olympic Broadcasting Services defines a reference architecture for a software-defined outside broadcast van under the working name of the Virtualized Outside Broadcast Van (vOB). The goal is a fully virtualized architecture based on a common platform using COTS hardware managed under broadcast and software-defined network (SDN) orchestrators.

The defined architecture uses open-standard application programming interfaces (APIs) and retains the user experience familiar to broadcast engineers and operators using traditional broadcast appliances. The standards-based platform enables multiple best-in-class software applications from one or more vendors of choice to be deployed on the same physical platform. This approach facilitates simple scalability of physical hardware resources to match the complexity and compute requirements for various broadcast events.

Olympic Broadcasting Services and Intel jointly defined an architecture addressing networking, processing, storage, software, control, and orchestration requirements for a live video production based on COTS hardware.

Piloted at the curling event during the Olympic Winter Games 2022, the reference architecture mirrored the standard broadcast production made available to international broadcasters. Multiple vendors contributed to the project, with their software applications deployed on a common platform for processing audio and video streams. The first stage of this project prioritizes functionality and interoperability, ingesting and processing 1080p50 standard dynamic range (SDR) video and audio channels. The next phase of the project will scale to 4K video resolution and include support for High Dynamic Range (HDR) video. The transformative nature of the reference architecture is expected to deliver the following benefits to broadcasters:

- **Reduced costs** include both equipment capital expenses (CapEx) and logistical operating expenses (OpEx) associated with producing events. The use of drop-in, general-purpose COTS hardware in place of shipping specialized, dedicated equipment streamlines logistics and enables a smaller physical footprint. COTS resources can also be repurposed during non-event times, helping amortize costs across other workloads, and CapEx can be shifted to OpEx by paying only for the hardware and software resources needed, on an on-demand basis. New codecs and features can be added with simple software updates, rather than requiring forklift upgrades.
- **Improved flexibility** is enabled by open workload placement and choice of best-in-class software applications among solution vendors, as well as support for multiple software deployments on a single common platform. Separate workflows can be efficiently distributed across different locations including on-prem, off-prem, and cloud, and the architecture can easily scale as needed by adding additional compute resources to the existing infrastructure.
- **Unified platforms** based on COTS architecture provide a homogeneous architecture across video production pipelines that emulates general-purpose IT infrastructure used elsewhere in the organization, simplifying deployment and maintenance.
- **Scalability** is enhanced by ongoing improvements in availability and price points of bandwidth and IP switching capacity, so infrastructure capacity can be gradually increased at a constant level of financial investment.
- **Redundancy** improvements provided by the unified platform include the benefits of N+1 redundancy versus 1+1 redundancy to reduce costs and physical footprint at events. Redundancy can also be provided by public cloud infrastructure that is utilized only if required, further reducing costs.
- **Carbon footprint** is reduced by avoiding the need to send broadcast vans, equipment, and personnel to event venues, particularly for international events.

PoC Objectives

By validating the reference architecture at the Olympic Winter Games 2022, Intel and the Olympic Broadcasting Services achieve the following objectives:

- **Reduce logistical and operational complexity** compared to traditional broadcast infrastructure, including planning, transportation, and setup.
- **Provide an open architecture** that replaces proprietary monolithic broadcasting equipment with virtualized and containerized services designed to run on COTS hardware and cloud infrastructure.
- **Increase flexibility** by enabling a single system to fulfill multiple roles across a variety of live sports and other events, allowing testing and commissioning outside of restrictive timelines.
- **Reduce the overall broadcast footprint** at the venues and the IBC (International Broadcasting Center).

Venue

The curling event at the winter games took place at the National Aquatic Centre (NAC) venue in the Olympic Park in Beijing, which has been transformed from a venue with swimming pools used during the Olympic Games Beijing 2008 into a venue with several curling sheets (fields of play), as illustrated in Figure 1. The production was performed at the venue with all equipment also located on site.



Figure 1. Curling sheets.

Production Technical Details

SMPTE ST 2110 Suite of Standards

Traditional broadcast infrastructure built with SDI baseband routers, coaxial cables, and BNC connectors is transitioning to Ethernet using IP network switches and SDN control. The increasing speeds supported by Ethernet switches and the corresponding improvement in aggregate non-blocking switch and network throughput have paved the way for the use of IP and Ethernet to transport critical broadcast applications in a cost-effective manner. At the same time, this new approach delivers parity in robustness and stability compared to legacy SDI operations, with increased agility, flexibility, and scalability to meet shifting requirements from ongoing media-format evolution.

To meet this challenge, a collaboration between industry partners resulted in several new standards, specifications, protocols, and recommendations from the media, internet, and IT sectors. As a result, several new standards have been approved and ratified for live production, notably the [SMPTE ST 2110 Professional Media Over Managed IP Networks suite of standards](#) and the [Advanced Media Workflow Association \(AMWA\) Networked Media Open Specifications \(NMOS\)](#).

SMPTE ST 2110 is quickly becoming the standard for transmission of low latency uncompressed or visually lossless compressed video and audio essence streams over IP networking for live video production. The standards provide for video, audio, and ancillary data to be sent as individual essence streams that are synchronized using Precision Time Protocol (PTP), with up to nanosecond precision. ST 2110 was used within the PoC to transport audio and video streams between endpoints with SDI signals converted to ST 2110 using SDI-to-ST 2110 IP gateways.

Production Format and Standards

Video

Table 1. Video production format and standards.

Video Format	HD 1080p50SDR
System Nomenclature	1920 x 1080/50/P SMPTE 274 System 3
Resolution	1920 x 1080
Picture Aspect Ratio	16:9
Frame Rate	50 frames per second
Scan Mode	Progressive
Color Gamut	Rec. ITU-R BT.709
Color Sampling	Y'C'rC'b 4:2:2
Color Depth	10-bit

Audio

Audio specifications used in the reference architecture included the following:

- Pulse code modulation (PCM) uncompressed at 48 kHz sample rate/24-bit with -18dBFS reference level.
- Production audio formats of stereo 2.0 and immersive sound 5.1.4.

Sources

Table 2 shows the feeds that were implemented as inputs to production.

Table 2. Production input sources.

Input Types	Total # of Inputs
Camera Splits from the Host Coverage	14
Dedicated Cameras for the PoC	4
Replay PGM/PVW Outputs	4
GFX Signals	2
Replay Audio Channels	48
Dedicated Microphones for the PoC	12
Mono Audio Splits from the Host Coverage	64

Processing Application Pillars

The key control and processing applications within the project were known as the “project pillars,” as illustrated in Figure 2. Multiple vendors collaborated with the Olympic Broadcasting Services and Intel in defining and implementing the live video production workflow on the defined vOB architecture. The main priority of the PoC was to test functionality and interoperability among the vendors’ pillar applications. AMWA NMOS and SMPTE ST 2110 compliance were mandatory for each pillar. The compute and processing applications were all provisioned on a common base platform with support for hardware acceleration.

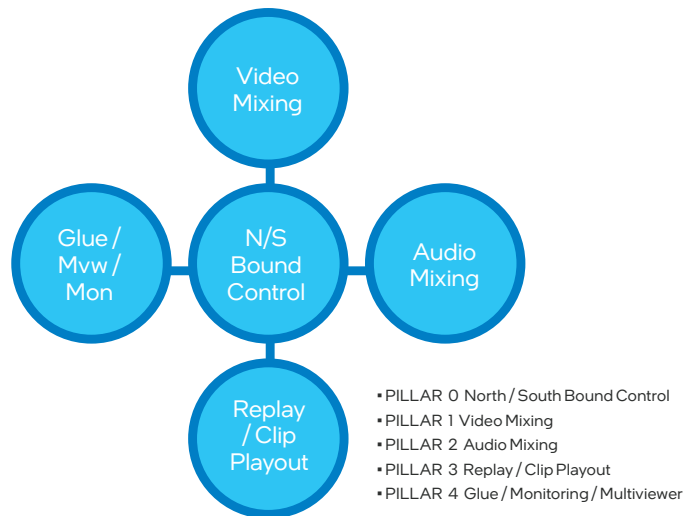


Figure 2. Processing application pillars.

The processing applications implementing the core video mixing, audio mixing, replay, and multiviewer pillars within the project were provided by Evertz, EVS, and Grass Valley. The control pillars were based on solutions from EVS and Nevin.

Architecture

Control and Orchestration

- **Northbound (broadcast) controller** fulfills the role of a traditional broadcast controller, including discovering endpoints and configuring connections, tally, and under monitor display control.
- **SDN controllers** manage both the real-time IP networks and the media endpoints. The SDN controller has the same role as traditional SDI routers, connecting inputs with outputs and managing flows within the network.
- **NMOS** was key to the control and provisioning of the network. Each ST 2110 Rx and Tx endpoint announced its presence to the broadcast controller using IS-04, with IS-05 and IS-06 managing connections between devices.

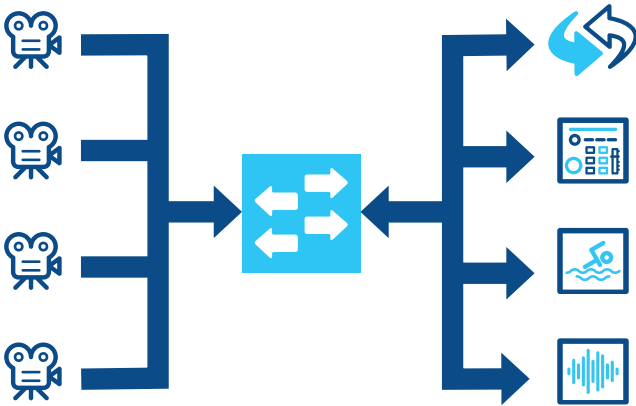


Figure 3. Solution data flow.

Software Architecture

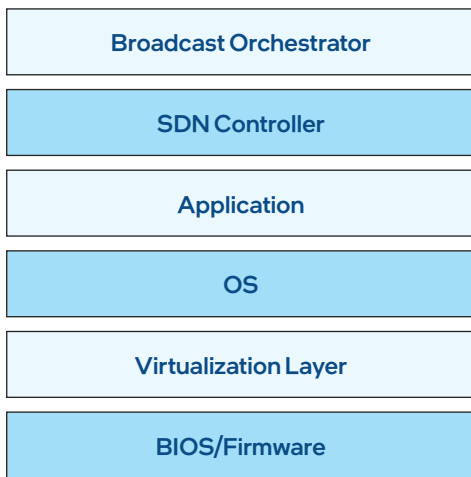


Figure 4. Solution software architecture.

Rack and Network Architecture

The vOB rack architecture is based on data-center-style hyper-converged infrastructure (HCI) that combines compute, network, and storage resources in a single node for flexible, simple deployment. Orchestration and compute nodes are based on the Intel® Server System M50CYP Family, with a one-to-one mapping between project “pillars” and compute/controller nodes.

The orchestration node hosts all control and management services running on Intel® Xeon® Platinum 8352M processors, with the application workloads distributed between Intel Xeon processors and accelerators based on PCIe add-in cards. The compute nodes are populated with both graphics processing units (GPUs) and field programmable gate array (FPGA)-based accelerator cards to meet diverse performance, latency, and user experience requirements.

The reference architecture ensures the independent function of the application control, streaming, and server management networks. The Intel Server System M50CYP Family provides a dedicated 1 Gbps RJ45 remote management port for remote access and control using intelligent platform management interface (IPMI) and a baseboard management controller (BMC).

Application control and data streaming connectivity are provided using separate interfaces to provide isolation for data and network traffic controls, following best practices for IT network topology. This consideration is particularly important for network topologies that support ST 2110 data traffic, to reduce the possibility of data loss.

The Arista 7170 network switch based on the Intel® Tofino™ ASIC was used for the high bandwidth 100 GbE media network with a low-bandwidth switch used for control and platform management.

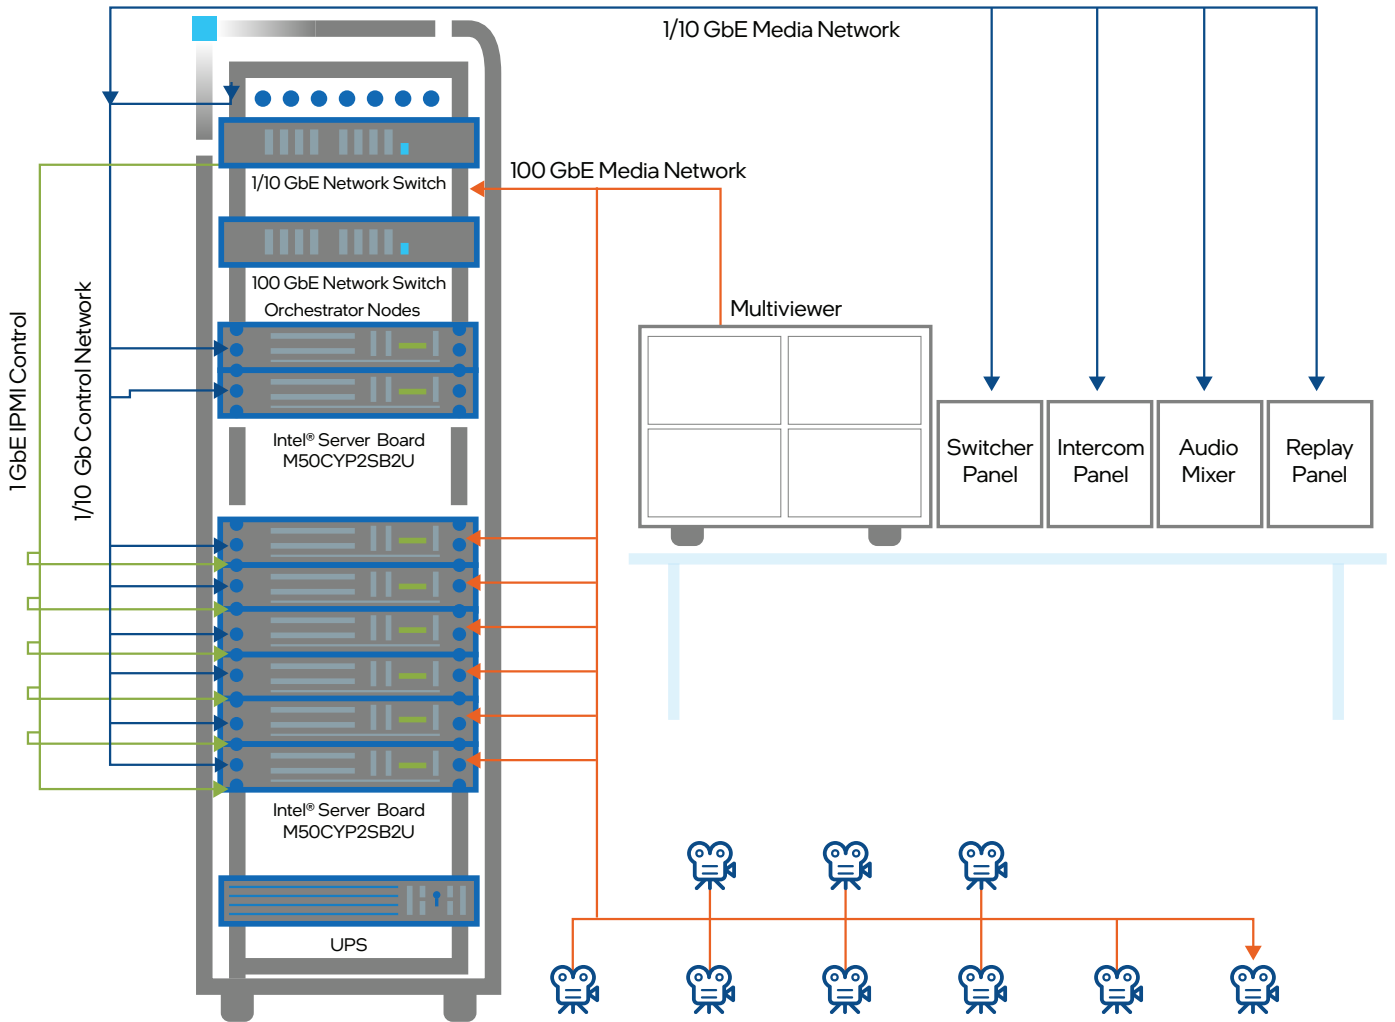


Figure 5. Solution network architecture.

Intel® Architecture Building Blocks for Video Production

Intel Server System M50CYP Family

The compute architecture is based on the Intel Server System M50CYP family, illustrated in Figure 5, a dual-socket system for mainstream data center workloads. The server is powered by 3rd Generation Intel® Xeon® Scalable processors, which provide up to 40 high-speed cores per socket (80 cores/160 software threads per server) and accelerated I/O with up to three Intel® Ultra Path Interconnects between processors and up to 64 PCIe 4.0 lanes per socket.

The Intel Server System M50CYP Family is validated for a wide range of add-in cards and accelerators. The systems used in the PoC are configured with 100 Gbps network interface cards for ST 2110 streaming and control networks, as well as discrete GPUs and FPGA cards for video and audio processing where required. The architecture is based on the Intel Xeon Platinum 8352M processor, a 32-core 185-watt CPU that is optimized for media processing workloads with optimizations to deliver better performance per watt for workloads including media and AI. The Intel Server System M50CYP family of products is developed and supported by the Intel® Data Solutions Group, which provided on-site break-fix support during the Games.

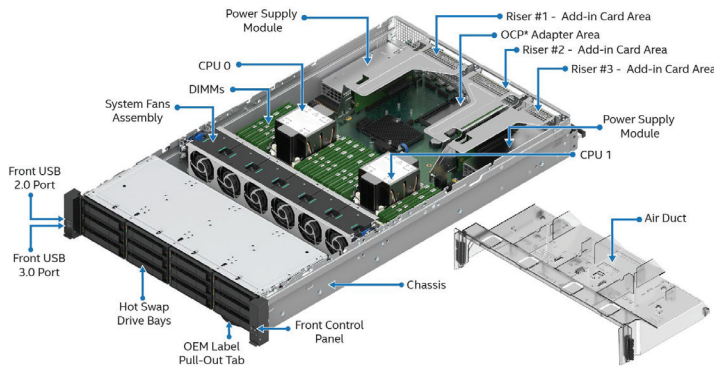
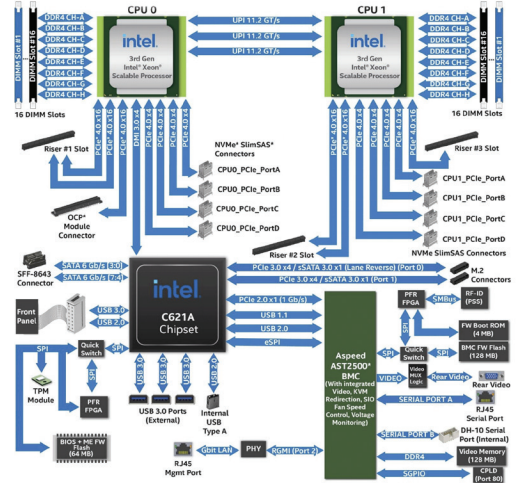
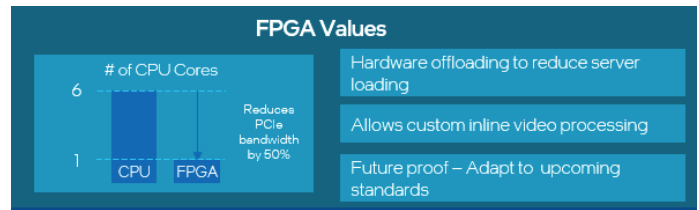


Figure 6. Intel® Server System M50CYP family.



Intel® FPGAs

Intel® FPGA SmartNIC N6000-PL Platform, illustrated in Figure 7, enables vendors to implement their FPGA-based IP on the defined COTS platform. Also available with an integrated dual port 100GbE Intel® Ethernet Controller E810 (the Intel FPGA SmartNIC N6001-PL Platform), this high-performance PCIe-based FPGA acceleration card is ideal for workloads that require both inline and lookaside acceleration. The Intel FPGA SmartNIC N6000-PL Platform can be used to offload video and audio streams, providing high density for multiple and simultaneous audio processing, including ST 2110 network connectivity on chip. The device offers a compact and reliable solution to offload control and monitoring functions from the host CPU. It provides the performance and versatility of FPGA acceleration and is one of several platforms supported by the acceleration stack for Intel® Xeon® CPUs with Intel FPGAs. FPGAs can be reconfigured on the fly to address a wide range of video and audio processing requirements. This programmability makes FPGAs a great choice to implement very diverse high performance and low latency media processing pipelines, as required in modern video production.



Results based on Intel estimates using Intel® Xeon® Platinum 8358 processor, Intel® Agilex™ F-Series 014 FPGA, and Intel Partner Solution IP (Macnica ST2110-21A and ST2022-7) as of March 2022.

Intel GPU

Intel launched the Intel® Server GPU, its first discrete data center graphics card, in 2021. The next generation graphics card is coming to market in 2022. Intel data center graphics cards are an excellent fit for video production workflows with the right mix of media processing capabilities and excellent graphics and general compute performance.

Intel GPUs have dedicated media codec IP blocks with support for AVC 8b and HEVC 8b/10b 420/422/444 decoding and encoding, HDR and SDR processing, and macroblock/coding unit granularity of controls. Intel GPUs support OpenGL, OpenCL, and Data Parallel C++ for high-performance general compute and graphics processing.

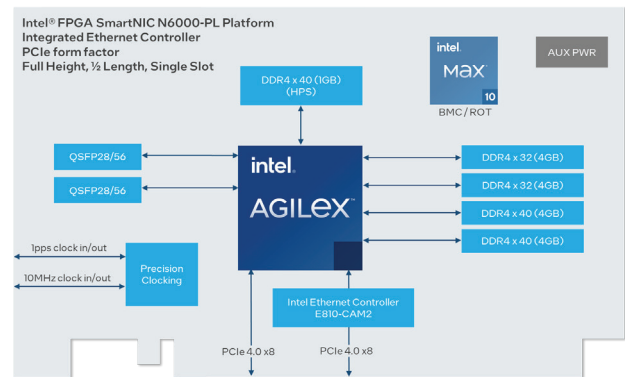


Figure 7. Intel® N6000 Acceleration Development Platform

Intel® Network Adapters and the Intel® Media Transport Library

Intel® Ethernet 800 Series Network Adapters deliver speeds up to 100 GbE and are available in both PCIe and Open Compute Project (OCP) form factors. The Intel Server System M50CYUP family has a single OCP slot, which supports the Intel Ethernet 800 series OCP card. The Intel Media Transport Library enables SMPTE ST 2100 compatibility with Intel Ethernet 800 Series and 700 Series Network Adapters.

The library takes advantage of the data plane development kit (DPDK) for kernel bypass and high bandwidth network performance. It supports ST 2110-10, ST 2110-20, ST 2110-22, ST 2110-30, and ST 2110-40, including using hardware-assisted capabilities of the Intel Ethernet 800 Series Network Adapter to achieve 'N' type pacing up to 4K60.

Intel® Virtual RAID on CPU (Intel® VROC)

Intel VROC is an enterprise hybrid RAID solution, specifically designed for NVMe SSDs connected directly to the CPU, that takes advantage of the Intel® Volume Management Device (Intel® VMD) on Intel Xeon Scalable Processors. Intel VMD enhances the dependability of NVMe connections. Intel VROC uses Intel VMD to enable a simple RAID solution that requires no additional hardware, unleashing the full potential of NVMe drives.

Intel® Silicon Photonics

Intel® Silicon Photonics encompasses a wide product range, including transceivers and active optical cables (AOCs) that address the 100 Gbps to 400 Gbps market. In the Beijing Olympics PoC, the Intel Silicon Photonics 100G CWDM4 dual fiber transceiver, shown in Figure 8, was used for communication between the cameras and rack switch. Compute nodes were connected to the 100 Gbps network switch using 100 Gbps AOC cables for inter-rack high-bandwidth connectivity.




Figure 8. Intel® Silicon Photonics optical transceiver.



Performance varies by use, configuration and other factors. Learn more at www.Intel.com/PerformanceIndex

Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See backup for configuration details.

Intel technologies may require enabled hardware, software, or service activation.

No product or component can be absolutely secure.

Your costs and results may vary.

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.

0322/DL/MESH/PDF 346441-001US

Intel® Tofino™ Series Programmable Ethernet Switch ASIC

The Intel® Tofino fixed-function application-specific integrated circuit (ASIC) delivers full P4 programmability for enabling custom feature and function implementation within the network switch. This is ideal for broadcast workflows, allowing the implementation of custom features that address the unique requirements of the video broadcast market.

Conclusion

Virtualization will redefine and simplify broadcast production requirements and workflows. This approach allows the broadcast production environment to be scaled to provide for changes in demand and workload, with production workflows that can be spun up and down as required in a matter of seconds. It optimizes setup time by enabling systems to be configured remotely before moving to the host city. In the future, it may be possible for production crews to operate from their own home-country premises, eliminating the need for dedicated full production crews on-site and avoiding the long-term rental of worldwide broadcast equipment for live production at the Games.

With mature, cloud-based infrastructure now prevalent across the broadcast industry, helping to process and distribute content faster and more accurately than ever before, virtualization opens new opportunities that could provide the groundwork for producing video coverage of the Olympic Games in a whole new way in the near future. Virtualization relies on software to simulate hardware functionality and create a virtual computing system in the cloud. The full adoption of an IP-enabled infrastructure, moving the functions of the in-venue production units away from the hardware traditionally on premise, will provide greater flexibility and scalability, while reducing the overall broadcast footprint.