THE BROADCAST BRIDGE

Connecting IT to Broadcast

Reality of IP

Essential Guide

EG

ESSENTIAL GUIDES

Introduction from HPE

Hello, and thank you for downloading this Essential Guide on cutting-edge processing in the broadcast industry.

The rise and rise of content on-demand has taken the industry by storm, but it's not hit us without its difficulties. This heightened demand for constant streaming and the move to ever higher definition video formats are just two of the challenges facing not only content creators, but those managing content storage.

The infrastructure delivering this content needs to be scalable, dynamic and able to support the bandwidth and speed required for the highest-quality content, without overloading broadcast network data centres.

At Hewlett Packard Enterprise (HPE) Original Equipment Manufacturer (OEM) Solutions, we have 75 years of experience of building high-performance systems. And our experts are available 24/7 to help OEMs build next-generation broadcast solutions tailored to their individual needs.

That's why we partnered with The Broadcast Bridge to create three articles to share our learnings on the latest trends and solutions in broadcast video, audio and metadata processing:

1. **Reality of IP Part 1 - Real time** defines real-time processing and how serial data systems are processed in traditional hardware, especially serial data rates.

2. **Reality of IP Part 2 - Benefits of Virtualization** focuses on how ST2110 has separated video, audio and metadata into separate streams that are co-timed to allow independent processing without embedders and de-embedders.

3. **Reality of IP Part 3 - Data Acceleration** investigates solutions using accelerated intra-GPU on-premise resources and why this is superior to cloud systems.

Nokia, Imagine Communications and Starfish Technologies are the broadcast OEMs that have already benefited from a partnership with HPE OEM Solutions – their stories are included towards the end of the Guide.

We hope you find this Essential Guide useful. If you want to take advantage of HPE OEM Solutions' decades of expertise, please get in touch and we will help you take the next steps toward innovative, future-ready systems.

Sincerely,

Rod Anliker and Matt Quirk

Rod Anliker.

Matt Quirk.

Sponsored by Hewlett Packard Enterprise OEM Solutions

Reality of IP

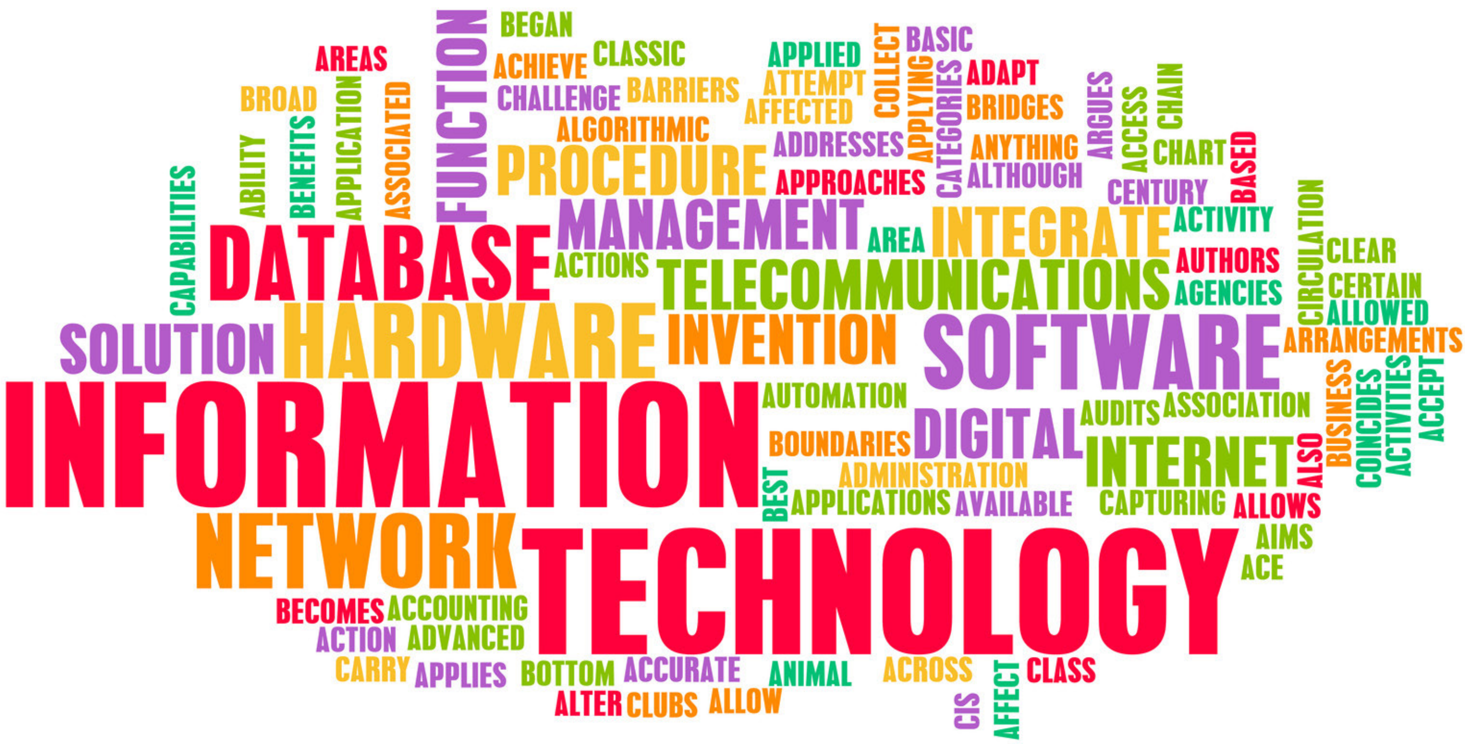

By Tony Orme, Technology Editor at The Broadcast Bridge

Part 1 - Real Time

Live broadcast television was once considered to be unique as every bit of data had to be delivered to the viewers television set in real-time. However, as IT continues to leverage its influence on television, we discover the uniqueness of broadcasting isn't as exclusive as we may have once thought.

Real-time is a term used across many industries and has many different meanings. There will always be some delay between an event occurring and the recipient witnessing it, whether it's due to processing delays in electronics and software, or transmission characteristics of electromagnetic propagation. The only question then is, how much delay is acceptable?

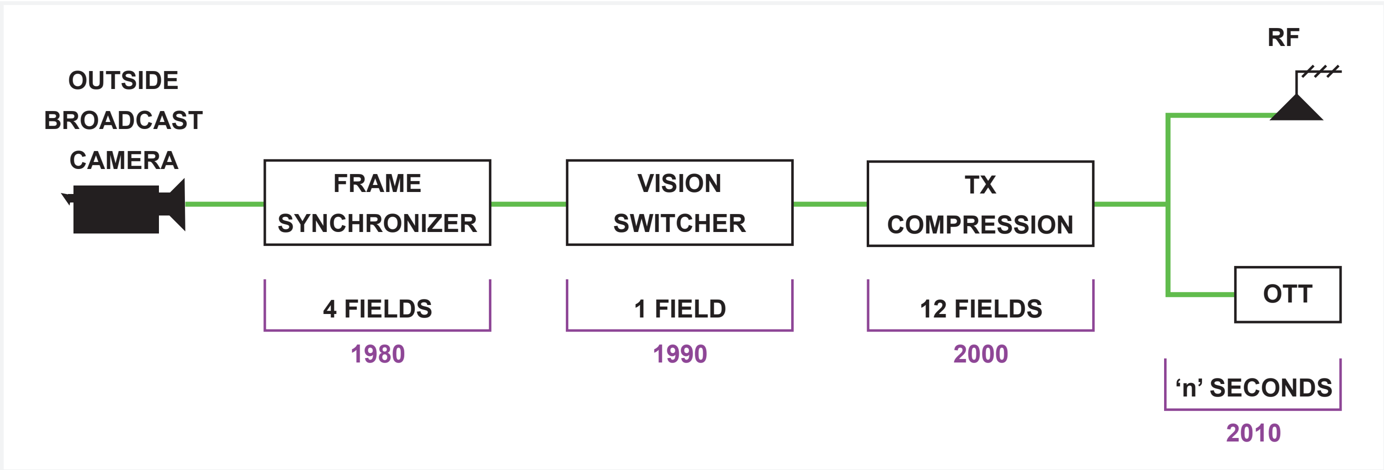

Diagram 1 – Every ten years a new technology has been introduced to improve the viewer experience, but has also increased latency and delay.

In the early days of broadcasting, the electronic beam of the viewers' cathode-ray-tube television set was synchronous with the scanning beam of the camera. Only a few microseconds of delay occurred due to propagation delays through amplifiers and the time taken for the signal to reach the viewer from the transmitter.

Progressive Latency

Frame synchronizers were built to facilitate outside broadcasts and live news feeds. They introduced multiple fields of delay resulting in fifty or sixty milliseconds of latency. When SDI was introduced, vision switchers provided up to six lines of delay to widen studio timing limits.

Inter-field video compression and motion compensation delivered anything from half a second to many seconds of delay dependent on the length of the GOP (Group of Pictures). And with the introduction of OTT (Over the Top) internet delivery, latency increased to thirty seconds or more due to technologies such as HTTP (Hypertext Transfer Protocol) and DASH (Dynamic Adaptive Streaming over HTTP).

In the space of twenty-five years, we've gone from transmission latencies of a few milliseconds to nearly a minute.

Sponsored by Hewlett Packard Enterprise OEM Solutions

Hewlett Packard Enterprise
OEM

But does latency matter? If viewers are watching live programs in isolation, then probably not. However, many viewers now watch live sport events whilst interacting on hand held devices. Watching a soccer match in the garden or yard can be spoiled if your neighbors have a less delayed service and suddenly cheer, thirty seconds before you knew a goal had been scored.

IT has been moving into television from three directions; software applications running on servers, storage, and IP distribution infrastructures. Although IP has been prevalent in off-line systems for many years, it has only started to make an impact in live transmission since the advent of SMPTE's ST2022-6. The recent introduction of ST2110 will only accelerate integration.

True Benefits of COTS

Many talk-up the benefits of COTS for broadcasting. The advantages are not restricted to hardware procurement, but instead, it's about being able to tap into the massive amount of innovation and research and development that other IT industries bring to the table.

Video and audio streaming now account for 75% of internet traffic and Cisco is reporting that it would take five million years for one person to watch the amount of video that will cross global IP networks each month by 2021. And with the number of worldwide smart phone users predicted to rise to 2.5 billion by 2020, there has clearly been a lot going on in the IT world.

New emerging technologies such as 5G are a good example of how broadcasters gain from telecommunications' R&D effort.

Low Latency Already Possible

In the UK, BT's mobile operator EE has demonstrated a 5G network with 2.8Gbit/s end-to-end speed and 2ms latency through the IP core and RAN (Radio Access Network). This is showing broadcast speeds and latencies are now achievable in IT, even over the airwaves.

More important is the business model this is built on. EE are the operator and have designed the virtualized 5G IP core and RAN in partnership with Huawei, who in turn have provided the software that is running on COTS servers provided by HPE.

Many broadcast engineers have grown up with the idea that to achieve real-time television, we must use dedicated video hardware systems, and up to recently, this has been a valid assertion. With 5G tests showing the kind of results EE is achieving just demonstrates the balance is now tipping to IT systems.

SLA's Deliver

Data speed and latency has been improving at unprecedented levels and the IT systems of ten years ago do not do justice to the speeds being achieved today.

The model demonstrated by EE shows another key advantage of COTS system - Service Level Agreements. SLA's are deeply entrenched within IT but relatively new to broadcasting. The IT industry has been working to improve SLA's for many years and the whole model for support has been refined and tweaked to make it highly efficient and fit for purpose.

Hardware used in high-frequency-trading and 5G is at the cutting edge of currently available technology and is directly applicable to real-time broadcasting. Consequently, there are few manufacturers that can design, build, and support hardware to the precision needed at this end of the market.

Single Point of Support

To keep trading floors, communication networks, and commercial websites running 24/7, IT manufacturers have built international support centers. They hold spares and provide follow-the-sun expert support to keep hardware infrastructures running at optimum capacity throughout the year.

By building alliances with IT manufacturers, software solutions providers are delivering the best of all worlds for their clients. In addition to providing cutting-edge solutions and level-one support, they will also have support contracts for the hardware with the IT manufacturers. So, if the hardware is at fault then parts can be shipped within hours and repaired on site.

The beauty of SLA's is that they are administered and serviced by the software solutions provider, so the end client only has one person to contact in the event of a failure.

Sponsored by Hewlett Packard Enterprise OEM Solutions

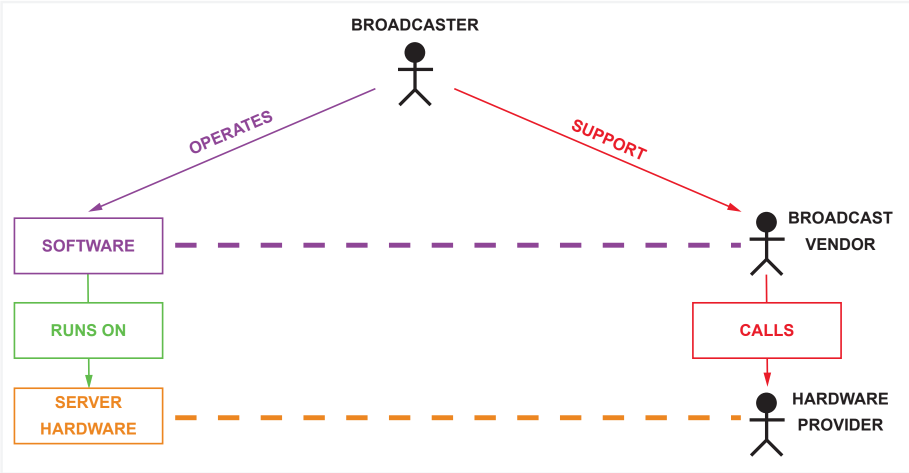

Diagram 2 - Hardware manufacturers are building alliances with broadcast vendors to provide world class SLA's. For the broadcaster, the principle point of contact is their vendor, if they can't resolve the issue then they will call their hardware partner and liaise with them to resolve the issue.

Due to the esoteric nature of broadcasting, it has been difficult for manufacturers to provide levels of support now common place in the IT industry. The volumes of scale are just not there. IT manufacturers can amortize the cost of their development and support over many industries and this is just not possible in broadcasting.

Virtualization, until recently, was also a limitation for broadcasters. But collaboration and developments in open-source software are making this technology viable for broadcasters.

NIC's (Network Interface Cards) regularly use buffers to hold data received from the Ethernet connection until the CPU has time to process the datagrams. But to keep latency low, ST2110 has imposed strict size limits on buffers and the amount of time datagrams can spend in them before being processed by the CPU.

IT is Already Achieving Low Latency

Low latency was once considered exclusive to broadcasting. But the same requirements are now demanded by video streaming in 5G and microsecond decision making of high-frequency-trading.

The challenge is further compounded in virtualized systems as many services may share the same NIC or group of NIC's, potentially adding further latency to the video streaming path and probably breaking the strict latency limits of ST2110, impacting heavily on real-time broadcasts.

To solve this problem, IT hardware manufacturers have invested heavily in R&D and found solutions, some of which will be discussed in a later article.

To bust a myth, COTS has nothing to do with being able to go down to your local computer store and buy replacements for your cutting-edge technology and high-speed infrastructure. But is instead about benefiting from the massive R&D investment IT manufacturers have already committed and will continue to build on. Many of the challenges we will face as we migrate to IP have already been solved by IT manufacturers.

As 5G is demonstrating, broadcasters can benefit greatly from innovation in IT. And the added advantage of SLA's makes deploying solutions based on IT hardware infrastructure compelling. In the Part 2, we will look at the benefits of virtualization.

Part 2 - Benefits of Virtualization

Riding on the back of IT innovation allows broadcasters to benefit from virtualization. In Part 2, we investigate those benefits and learn how they apply to television. Especially as we learn of the new trailblazers waiting in the wings.

Telephone network operators were suffering the same challenges found in many current day broadcast facilities. Proprietary hardware equipment was prevalent and even increasing as new technologies were developed. Space and power were at a premium and finding places to install hardware was proving more and more difficult.

Increasing energy costs, the need for more capital investment, and a skills shortage to design, install, and maintain complex hardware systems was proving unsustainable. And hardware systems quickly reach their end of life, especially when new technologies are constantly developed. This leads to very short design, install, and operate life cycles, resulting in poor return on investment for operators.

Virtualization to the Rescue

Virtualization addresses these problems by leveraging standard topologies found in IT systems. Building on the economies of scale gained by hardware manufacturers during their research and development cycles, broadcast service suppliers can achieve significantly reduced equipment costs and reduced power consumption.

Before virtualization was available, IT hardware was either under or over utilized. But virtualization enables more efficient use of IT infrastructure as it allows services to be moved seamlessly between resource and balanced with more efficiency to make better use of the available hardware.

Hardware design life-cycles can last many months and even years. But the economies of scale needed to design and manufacture hardware, such as computer servers and Ethernet switchers, is no longer applicable for software-based development. As data speeds increase, hardware manufacturers must purchase faster and faster test equipment, and continually invest in smaller skillset pools.

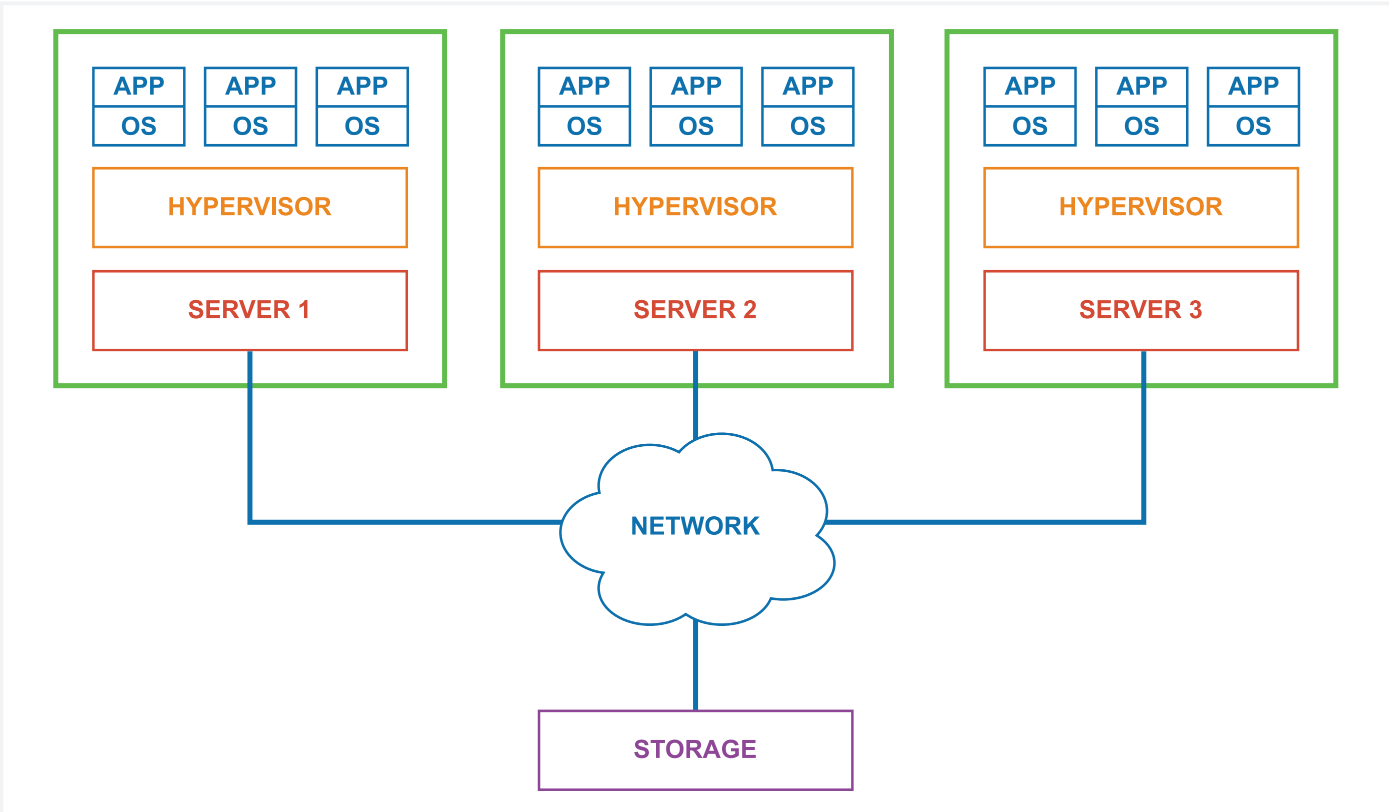

Diagram 1 – Three physical servers, each running multiple copies of different operating systems and applications via the hypervisor virtualized management software.

As device speeds increase, the cost of associated test equipment costs increases exponentially, and the salaries of the engineers designing the equipment increases as their skillset becomes more and more specialized. Thus, making the viability of high-speed hardware design more difficult for small and medium businesses.

In virtualized architectures, software vendors can design, build, and test their products on the same infrastructure. And with cloud systems, they can even employ follow-the-sun development strategies to program and release code 24 hours a day.

Innovators will Win

Existing broadcast service vendors may be looking at the developing virtualization platforms with some concern. But this doesn't have to be the case as many vendors already specialize in software development. In many cases, the hardware they designed was necessary to sell the service the software was providing. Instead, a standardized hardware platform could be an opportunity. As in many cases, their real value lies in the service their software is providing or problem it is solving.

Cloud infrastructures are deployed throughout the world with high-speed connectivity between data centers enabling localized processing to be brought closer to the viewer. This reduces latency considerably and improves the viewer experience.

Virtualization mechanisms are at the heart of cloud computing. Hardware virtualization is the process of creating a virtual machine that acts like a real computer with its own operating system. The host machine is the actual physical server running the virtualization system and the guest machine is the virtual instance. Many virtual instances with differing operating systems can run on a single host device.

Hypervisor and VMM

The software that provides the virtualization and creates the instances is called the hypervisor or Virtual Machine Manager (VMM). Both the VMM and hypervisor sit between the physical server hardware and the virtual instances.

Although virtualization was available in the earliest computers from 1967, it was not implemented in the first x86 architecture but did see a resurgence from around 2005. This was due to the improved hardware capabilities of servers, striving for greater cost efficiencies, better control of server render farms, and improved security and reliability through hypervisor architectures.

To help improve virtualization, Central Processing Units (CPU's), such as the Intel Xeon series, are built with support for virtualization directly into the silicon, thus offloading some of the hypervisor code from the CPU and providing a form of hardware acceleration.

Cloud infrastructures provided systems to facilitate automated initialization of instances and control of the devices. Monitoring software, also running in the cloud, establishes if more resource is needed and will spin-up new instances as required. When the work-load reduces, the instances will be switched off and deleted.

Few businesses have the privilege of building green-field systems and will instead develop organically. Cloud infrastructures exist on-prem and off-prem, they may be private or public. A broadcaster moving to the cloud may find they have multiple systems from different service suppliers and hardware manufacturers.

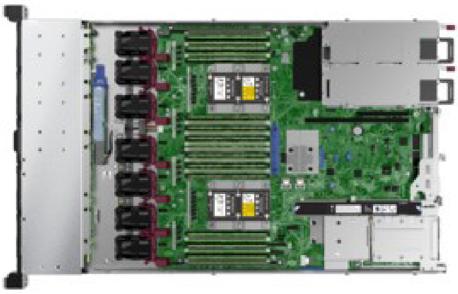

Diagram 2 – HPE ProLiant DL360 server supporting Intels Xeon processor with virtualization hardware support.

Integrated Control

Such systems can soon escalate in complexity and will become difficult to administer and control. Again, looking to the IT industry, there are many solutions. For example, HPE OneSphere brings together all the diverse cloud systems a broadcaster may be employing under one application, simplifying configuration and monitoring.

Understanding where servers and services are being deployed is critical when keeping control of the costs, especially with public cloud systems. As well as simplifying administration, a complete monitoring solution will provide increased granularity so that costs for projects, services and instances can all be gathered in real time. Preconfigured alarms help administrators keep control of costs and identify potential issues quickly.

Self-service is a relatively new concept and it relies on users being empowered to solve their own problems. By setting limits and providing preselected applications, IT administrators enable users to build and configure their own servers and launch applications. This is incredibly powerful for developers who want the flexibility to work on their own servers whilst knowing they are working in a safe environment where they cannot do damage, especially if a program under development develops a system critical bug.

Complete Resilience

Snapshots are copies of a virtualized instance and can be used as backups or as part of a disaster recovery system. The snapshot can be used later to create duplicate instances all the same. This is useful as the system administrator doesn't have to install an operating system from scratch when an instance is launched. Instead, they use the snapshot resulting in a faster launch time, and the administrators know the instance works to a certain level.

Rip-and-replace takes advantage of virtualization. If a service running on a server instance is problematic, for example a software upgrade may have failed, then the instance is simply deleted from the virtualization control console, and another instance created in its place. A snapshot instance can also be used to improve spin-up times.

Software Defined Infrastructures (SDI) describe the collective system virtualized servers, storage and networks. Using SDI, broadcasters can focus more on their operation and keep control of costs, and service providers can spend more time on delivering services to solve problems, allowing them to focus on their core skillset, instead of fighting the hardware.

Innovation has improved virtualized and cloud computing leaps and bounds over recent years to the point where it is now the de facto system in most IT installations. Public, private, on-prem, and off-prem systems further increase the number of options available to broadcasters and their vendors. Delivering systems that have unprecedented flexibility and speed.

Recent advances in input-output processing technology have further improved the prospect of virtualization for broadcasting. And in Part 3, we investigate the technologies needed to make super-high-speed, low-latency, ethernet data transfer and processing in SDI viable for broadcasters.

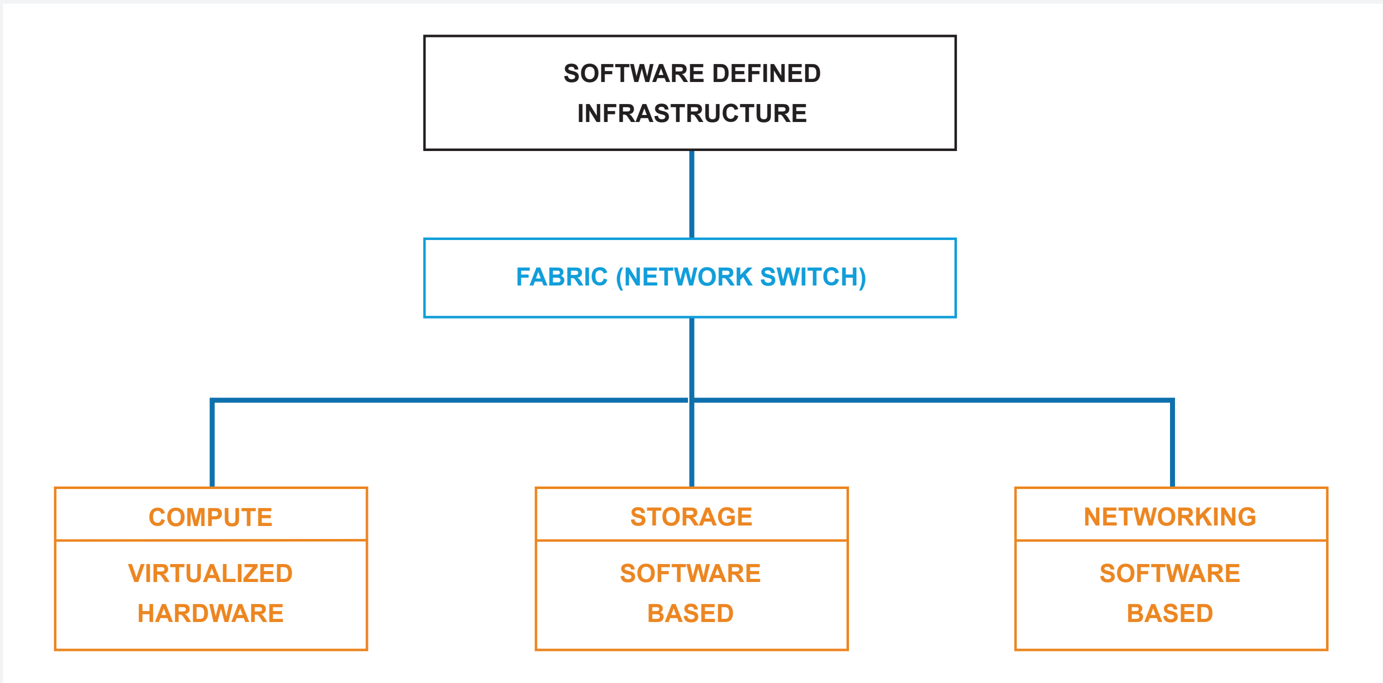

Diagram 3 – With a fully SDI data center, servers will be virtualized, storage will be software based and the network will be software defined, giving the most optimized and flexible system possible.

Part 3 - Data Acceleration

Network Interface Cards (NIC's) are often seen as the bottleneck of data processing for ST2110 and ST2022-6. IT manufacturers have witnessed similar challenges with high speed trading and 5G networks but have been able to provide real-time solutions to overcome latency and blocking. In Part 3, we investigate IT's achievements and how they are applicable to broadcast television.

A progressive HD video source creates just over 2.5Gb/s of data resulting in approximately 200,000 IP packets every second. SMPTE's ST2110 adds rate shaping to keep the variance of the packet rate within tight limits. The purpose of this is to keep latency low and avoid dropped packets.

ST2110 defines two parameters for rate shaping; Cmax and VRXfull. Cmax describes the variation of packets leaving the sender and VRXfull describes the variations of packets being read at the receiver.

In the ideal world, the rate at which a device writes to a buffer will exactly equal the read rate. However, in software systems, the short-term data rate being written to the buffer may spike briefly, even though the long-term average will be the same.

Packet Loss from Bursts

If a camera is sending packets with too much burst, then the short-term variance could be too high. Either a larger send-buffer will be needed in the camera or packets will be lost. A larger buffer will solve this issue but will result in increased latency as the packets are spending more time in memory.

Type	Cmax	VRXfull
N (narrow)	4	8
NL (narrow linear)	4	8
W (wide)	16	720

Table 1 – Size of relative packet variance for 1080p50 video stream.

Field Programmable Gate Arrays (FPGA's) can send and receive packets with very low variations. But software systems can easily create large variances due to the interaction of the application software, operating systems, and shared resource.

ST2110 has three parameters to define the sender rate shaping; narrow (N), narrow linear (NL) and wide (W). Table 1 above shows the relative variance sizes compared to the data rate of the video.

To put this into context, the 1080p50 HD stream requires approximately 3.7 IP packets per line of video. In the narrow model in Table 1, the maximum variance is 4 IP packets, or just one video line. But the wide variance is 16 packets, approximately 4 video lines at the sender, and 720 packets, approximately 194 lines at the receiver.

As we migrate to IP, manufacturers are writing new FPGA firmware to deliver ST2110 streams resulting in well gapped packets with very little variance that will easily match the Narrow model.

Wide and Narrow Incompatibilities

There are potential challenges for graphics processing applications running on servers. Its highly likely that developers of a pure software application would specify the wide model to provide the maximum IP packet variance allowable under ST2110. However, if it was streaming to a connected device using an FPGA implementation then it may well specify the narrow method. A wide output and narrow input are incompatible and will result in lost packets.

IT manufacturers have been working to resolve similar latency issues and have achieved better latency and data throughput in recent years. It is now common to aggregate computing on virtualized, cloud native clusters of COTS servers. 5G Telco clouds moved to this architecture and new IP-based broadcast systems will adopt similar architectures.

Aggregating compute and networking in a private cloud architecture yields significant bandwidth using commonly available IT servers. Today, a typical dual processor server node may have 16 cores per processor or 32 cores per node. A typical private cloud deployment may range from 4 to 32 nodes, or 128 to 1024 cores.

Offload Repetitive Tasks

To further improve CPU performance, packet processing tasks are offloaded to hardware controllers. A compromise between flexibility, off-the-shelf availability, and cost is constantly being balanced.

Using Linux as an example, a NIC will receive an IP datagram and copy it to memory. When a predefined memory limit has been reached, the NIC generates an interrupt causing the kernel to terminate the current process and then service the NIC's request to copy the IP packets from the NIC card to the server memory.

Even using the lightweight UDP protocol, kernel-based networking soon reaches the limits of its capabilities due to the speed of the processor's core performance ceiling, typically 8Gbits/s. As bandwidth demands increase, so does the required CPU performance resulting in increased and unpredictable latency.

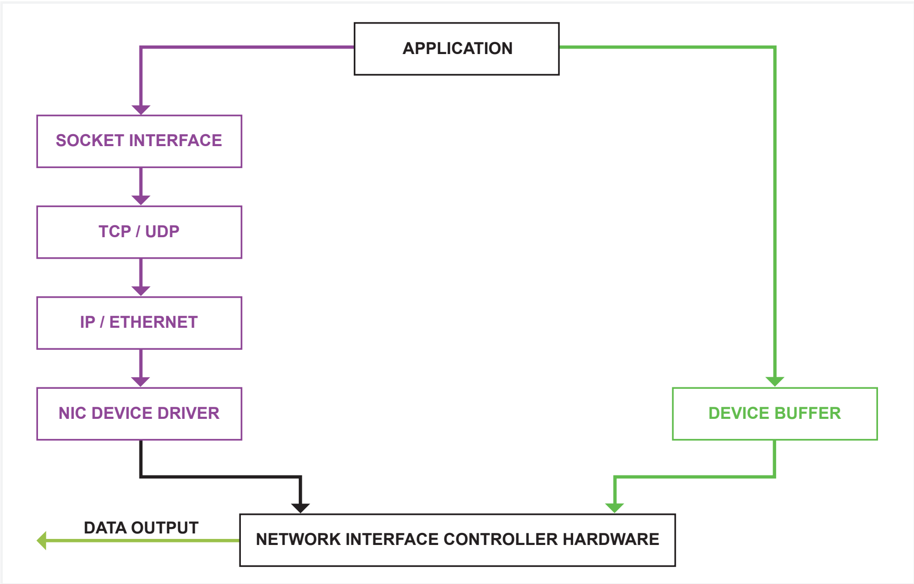

Diagram 1 – the left side shows data flow through the kernel to access the NIC, the right shows direct access to the NIC to reduce latency.

Kernel Bypass

Video processing is CPU intensive and we can't afford for the CPU to be tied up moving data backwards and forwards between its memory and the NIC's buffers. Kernel-bypass solves this issue by offloading input/output processing from the CPU to more intelligent NIC's.

Bypassing the kernel and IP stack reduces context switching, that is the overhead time taken to service the NIC's interrupt, and memory to buffer copies resulting in extremely low latency. Data received from the NIC is written directly into user-space facilitating faster data processing.

The Data Plane Development Kit (DPDK) is an open source project to deliver fast packet processing in networking applications using kernel-bypass technology. Essentially, DPDK is a programming framework that allows developers to load and configure libraries and drivers to meet the needs of their specific application and hardware.

The Environment Abstraction Layer (EAL) hides the hardware from the developer and provides a standard programming interface to libraries and available hardware accelerators such as performance network adaptors.

Polling Replaces Interrupts

To remove the overhead of interrupt processing, DPDK uses polling. Flags within the NIC will be read every few microseconds to determine if it has any data to transfer. To stop DPDK from blocking any other processes, a manager library is employed to implement lockless queues.

Using DPDK, Intel have reported the packet processing performance has been boosted by up to ten times on their Xeon E5 processor, achieving speeds of 233Gbits/s during IP forwarding.

Ethernet Just Keeps Getting Faster

Ethernet speeds continue to increase at an unprecedented rate. In just twenty years, Ethernet data rates increased ten thousand times from the humble 10BASE-T (IEEE-802.3i) 10Mbits/s in 1990, to 100Gbit/s (IEEE-802.3ba) in 2010. And in December 2017, IEEE ratified 802.3bs giving 200Gbit/s and 400Gbits/s.

IEEE-802.3bs specifies four distances to operate over; 100m, 500m, 2km (1.2 miles) and 10km (6.2 miles). Although 200GbE and 400GbE will initially be core network technology, as with other standards such as 25GbE and 100GbE, they will soon work their way closer to the edge of the network.

To put this into context, a 200GbE network connection can transport 66 simultaneous 1080P50 video services, or 132 over a 400GbE. Even an 8K service would be able to be distributed over 400GbE.

Better Economies of Scale

But the higher data rates deliver an unexpected bonus, they provide better economies of scale. 400GbE is not only four times as fast as 100GbE, but it allows a denser configuration. A 1U switch with 32 ports of 400GbE will have the same data throughput than the equivalent 100GbE switch with 128 ports, requiring at least 4U. Not only will this reduce the required rack space, but a single 1U unit will be cheaper to build than four 1U units delivering a similar capacity.

Technologies such as 5G are paving the way for broadcast television. With high speed, low latency networks, 5G operators have been able to prove bandwidths and latencies needed for broadcasters are achievable today. Adding the flexibility of COTS data center servers and network switches, the future is proving to be very optimistic for both television and media companies.

The time to market for traditional broadcast manufacturers has been slashed giving them untold opportunities to provide better services and business models. And new players will be able to hit the ground running, so they can share their new solutions quickly and efficiently.

HPE - Case Study

The Sponsors Perspective

Broadcast OEMs Innovate Faster with Intelligent Edge Technology

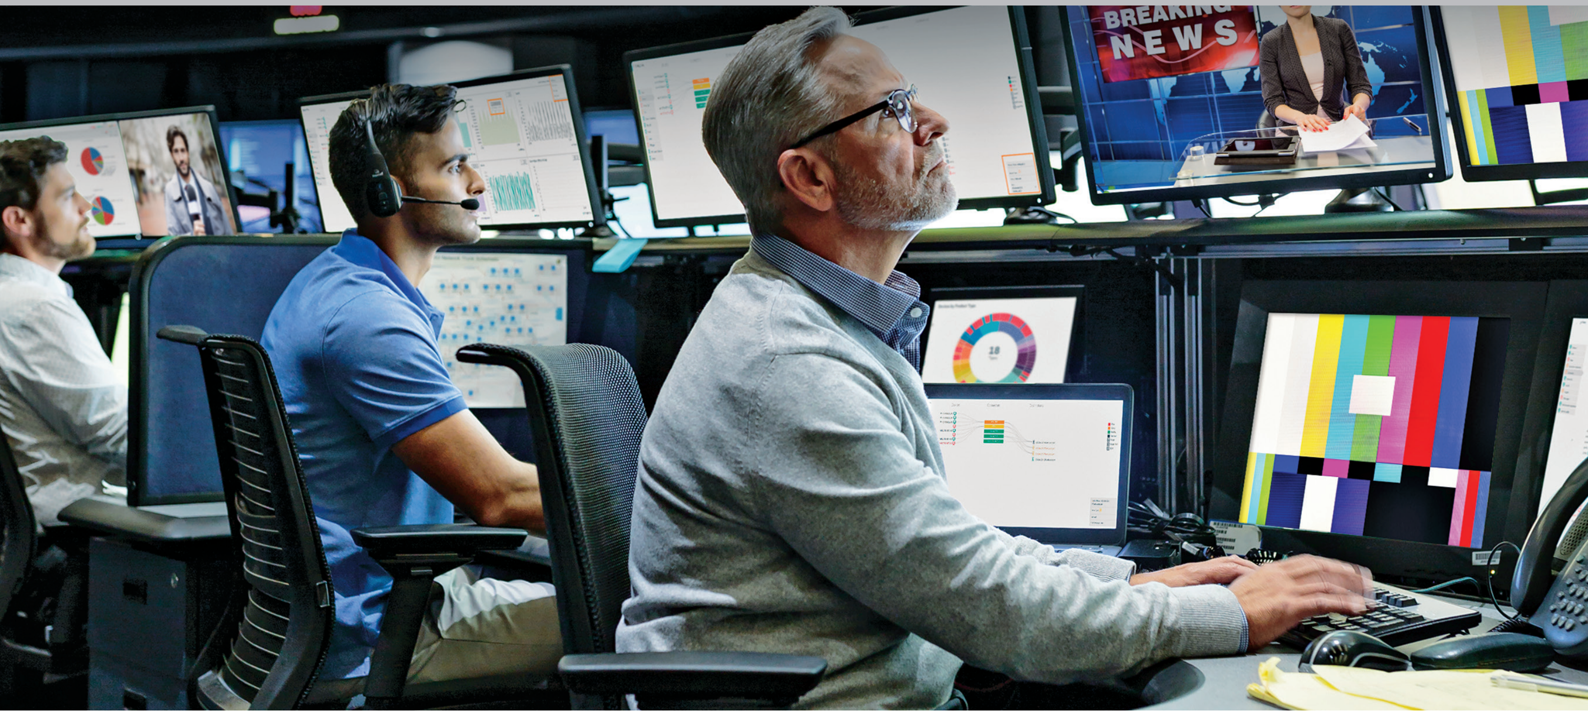

Original equipment manufacturers (OEMs) Nokia and Imagine Communications use Hewlett Packard Enterprise (HPE) OEM Solutions to process data faster and support customer desire for high definition, on-demand content.

The broadcast industry is undergoing radical change which is creating new opportunities for OEMs and their customers. Improvements in video compression technology, the migration from 4G to 5G, and the widespread adoption of streaming and viewing video on mobile technologies are just a few industry advancements broadcasters can monetize from today. But harnessing these new opportunities comes with technical challenges.

Sponsored by Hewlett Packard Enterprise OEM Solutions

Broadcasters also need to innovate to compete with over-the-top (OTT) providers such as Netflix, Apple and Amazon. OTT providers offer their customers streaming media services over the internet, reducing the need for traditional broadcasting companies. According to a study from Deloitte, the number of American households subscribing to a paid streaming service has increased 450 percent in less than a decade, highlighting the threat posed by OTTs[1].

[1] https://deloitte.wsj.com/cio/2018/05/02/viewers-migrate-to-streaming-tv-upending-media-models/

Broadcast OEMs need to provide next-generation products and services that will enable broadcasters to protect their existing customer base, secure new customers, and open up new revenue streams. To meet demand, OEMs need to build technologies that can speed up content delivery and bring these technologies to market, fast.

With its deep understanding of and experience in the broadcast industry, Hewlett Packard Enterprise (HPE) OEM Solutions helps broadcast OEMs design and introduce new, reliable services efficiently, and deploy Intelligent Edge solutions that capture data at the edge - where it is being created - and process it more rapidly than ever before.

What is the Intelligent Edge?

Connected 'things' are all around us, and they're generating more and more data, giving rise to new possibilities in our increasingly hyper-connected world.

The edge refers to where data is created and aggregated. Intelligent Edge solutions provide organizations with the ability to analyze this data in real-time, where it is captured.

HPE OEM Solutions' Edgeline portfolio consists of ruggedized enterprise systems specifically designed to exist in the harsh environments of the edge. And HPE OEM Solutions' full spectrum of edge-to-cloud products and solutions ensure data and analytics are transferred to the data center efficiently to drive real-time decisions and outcomes.

Intelligent Edge solutions help OEMs optimize operations, improve customer satisfaction and create differentiated business models. With data analysis at the edge, latency is reduced, so OEMs can offer their broadcast customers high definition, on-demand content almost instantly.

Imagine Communications and Nokia are just two OEMs working closely with HPE OEM Solutions to bring next-generation technology to market faster and improve customer services.

Imagine Communications: Helping broadcasters adapt to a changing media landscape

- Meeting the need for zero latency

For broadcasters to compete with OTTs and secure customer loyalty, they must deliver live and recorded broadcast content with zero latency. Broadcast customers also require on-demand content, so broadcasters need to offer a wide range of services that adapt to their end users' needs.

Imagine Communications, an OEM that provides end-to-end solutions to the media and entertainment industry, needed to develop solutions that would allow broadcasters to easily create and share content. This involves using technology that lets broadcasters store recorded content, monetize that content through advertising management solutions, and distribute content via multiple channels with zero latency.

As an HPE OEM Solutions partner, Imagine Communications is developing its applications on top of HPE servers and storage solutions, which include the HPE Moonshot System, HPE Edgeline System, HPE Apollo servers and HPE ProLiant servers. As a result, the company has been able to globally deploy its new technologies and applications much faster than ever before, increasing revenue opportunities in the process.

In fact, these technologies have worked so well that Imagine Communications has now achieved zero-latency and reliable performance for broadcasting content in multiple formats – including over the air, cable, satellite, web and mobile.

- Secure solutions for an uncertain world

Building and maintaining robust technological solutions is crucial – especially in the media industry, where going off air at any point is out of the question. HPE OEM Solutions is acutely aware of stringent security requirements and has added layers of protection across their software, hardware and firmware levels to keep OEM solutions more secure.

HPE OEM Solutions' new silicon root of trust technology ensures that HPE servers are protected for the full lifecycle –from the production process, through to supply chain and distribution, right to the customer's final location. HPE's servers can also recover essential firmware and restore operating systems in the event of a breach.

“For Imagine Communications, HPE OEM Solutions is not just a hardware provider. HPE OEM Solutions brings additional value because they have understood that beyond hardware, you need to have intelligent solutions that make hardware different and more secure,” said Dr. Glodina Connan-Lostanlen, Chief Marketing Officer at Imagine Communications.

Nokia: Enabling new revenue streams with quicker content distribution

Video streaming has opened up a wide variety of viewing options for consumers and gives them the power to select almost any form of content on any device of their choice. This flexibility creates a threat to traditional TV broadcasters.

However, IP-based networks are not always as reliable as broadcast networks, which can leave customers frustrated. This is where technology from OEMs like Nokia can help broadcasters separate themselves from OTT providers.

“With IP streaming, you have to maintain broadcast quality of experience. And you have to build the network for peak, which is driven by sporting events, new events and elections,” commented Paul Larbey, the head of the IP video division at Nokia.

If a customer tuned in to watch a World Cup soccer match, for example, only to find that the video feed wasn’t working, they may think twice about using this provider again.

In order to meet customer expectations for the highest quality of content delivered without interruption, Nokia partnered with HPE OEM Solutions to deploy HPE’s Edgeline Systems and move video processing closer the edge, optimizing the video experience through reduced latency.

With its Intelligent Edge technologies, Nokia allows broadcasters to capture data from customers in real-time, so they can offer personalized content to each customer. For example, the Nokia Velocix Media Delivery Platform, which runs on HPE ProLiant DL Servers and HPE Apollo Systems, lets broadcasters tailor program guides to each customer’s viewing interests and stream it on any IP video device including TVs, tablets, smartphones, game consoles and PCs.

WP
WHITE PAPERS

EG
ESSENTIAL GUIDES

MEDIA

CS
CASE STUDIES

Find Out More

For more information and access to white papers, case studies and essential guides please visit:

thebroadcastbridge.com

9/2018