Unmatched Global Market Coverage

Telecom COTS World

Broadband Broadcast IoT AI Convergence

Telecom IT - HPC - A.I. - IoT
Infrastructures - Networks - Edge - RAN
Cloud - Data Centers - Storage
Video Networks - Broadcast - Digital TV
From Chips to Rack Scale Systems

Telecom COTS World is a Global Publication of e2mos

Sep-Oct 2019

intel

Global Video Demand Requires New Hardware and Software Solutions

Headlines

- Microsoft and ServiceNow Announce Strategic Partnership
- ServiceNow Named a Leader in Gartner Magic Quadrant for IRM
- Evolution of Mobile Video: Two-Screen World becomes many with 5G
- NETCONF & YANG - Network Automation "The TL;DR"
- Huawei's Network Cloud Engine Completes NETCONF/YANG Interoperability Tests

Products & Technology from Silicon to Application-ready Platforms

In this Edition:

Cover Story:

Global Video Demand Requires New Hardware and Software Solution - Intel
Millions of data centers and telecom towers strategically placed across the planet, over 885,000 kilometers of undersea cable, billions of devices, and nearly all of it in service of one thing: online video.

Headlines:

• Microsoft and ServiceNow Announce Strategic Partnership
- Together the two companies will accelerate digital transformation for enterprise and government customers.
- ServiceNow moves workloads to Microsoft Azure for highly regulated industries.

• ServiceNow Named a Leader in Gartner Magic Quadrant for Integrated Risk Management (IRM) for the Second Year in a Row
ServiceNow recognized based on its completeness of vision and ability to execute.

• Evolution of Mobile Video
TWO-SCREEN WORLD BECOMES MANY WITH 5G.

• Google Cloud Adds Compute, Memory-Intensive VMs
Google added virtual machine (VM) types on Google Compute Engine including second-generation Intel Xeon scalable processor machines and new VMs for compute- and memory-heavy applications.

• NETCONF & YANG - Network Automation "What you need to know first"
- DISADVANTAGES of CLI and SNMP
- NETCONF is an IETF standard with broad market adoption incl. Cisco, Juniper, Huawei and many others
- With NETCONF & YANG you can add and remove devices without pain.

• Huawei's Network Cloud Engine Completes EANTC's PCEP and NETCONF/YANG Interoperability Tests

Daniel Dierickx CEO & co-Founder and Acting Chief Editor at e2mos

40 Years Chips & Embedded Computing Business & Market Expertise
Together with my Partners we have built a PREMIER Global Customer Database

Dear Reader,

Here is your free copy of
Telecom COTS World
one of our six High Focus e-magazines

Your benefits:

- Selected subjects: HIGH FOCUS
- Easy reading
- TIME SAVER: includes many direct links
- UNIQUE: Dedicated Website for each e-magazine, remain visible for years
- FREE Distribution Worldwide
- FREE Subscription

Discover our 6 e-magazines
Click on the logos below

Semi Update World | Embedded Systems World
aiworld | Telecom COTS World (Broadband Broadcast IoT AI Convergence)
IoT World | ATCA World

Editor/Publisher: e2mos
WEB: www.e2mos.com
Contact: mgt@e2mos.com

Global Video Demand Requires New Hardware and Software Solutions

Written by Lynn Comp, Intel, VP Data Center Group, General Manager Visual Infrastructure Division & NPG Strategy
September 13, 2019

Millions of data centers and telecom towers strategically placed across the planet, over 885,000 kilometers of undersea cable, billions of devices, and nearly all of it in service of one thing: online video.

According to Cisco, IP video traffic will be 82 percent of all IP traffic in the world by 2022, increasing by four-fold since 2017. Video on demand alone is doubling in the same period, an amount equal to 10 billion DVDs per month. Service providers need platforms that can scale in order to accelerate services and meet what has become an insatiable demand. The vast majority of visual cloud services are delivered today on Intel® Xeon® Scalable processors, with powerful and flexible open source software delivering scalable transcode, analytics, and more. In some circumstances, hardware accelerators may be useful to address density and power requirements; for example, at the edge where latency and bandwidth are advantaged and power and real estate are at a premium.

Efficiently scaling video processing and distribution requires hardware and software optimized for specific workloads, and Intel is at the forefront of both, investing in scalable architectures to accelerate innovation.

Hardware Accelerators

At IBC 2019, Intel is proud to announce two new hardware products to help meet global video demand. Joining Intel's diverse, world-class processors are two PCIe card accelerators specifically designed for visual cloud workloads.

First, the Intel® Visual Cloud Accelerator Card – Analytics (VCAC-A) is optimized for the highest density offload acceleration for use in network edge servers. The new acceleration solution combines both on-card decode and video inference acceleration, along with high efficiency, FP16 precision that can manage 24 streams at 1080p 30FPS, up to 144 streams per server with six cards.

... to next page

Global Video Demand Requires New Hardware and Software Solutions

... from previous page

Second, the Intel® Visual Cloud Accelerator Card – Render (VCAC-R) is optimized for cloud gaming, graphics, VR, and virtual desktops. With high performance rendering and encoding, low latency, competitive 3D performance, and multiple Windows games per card, VCAC-R can handle four Windows game streams at 1080p 30FPS, and at 720p 30FPS can handle 20 Android game streams.

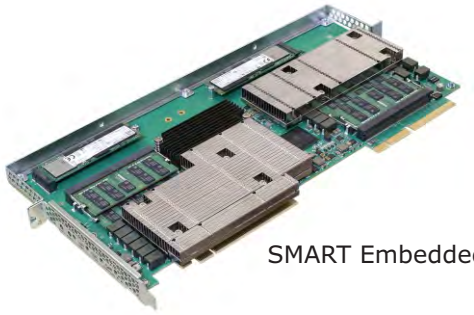

PCIE-7217
SMART Embedded Computing

Both accelerators are available today from our ODM partner Celestica and, like the previously released PCIE-7217 from SMART Embedded Computing (formerly Artesyn), are the result of Intel's close collaboration with ecosystem partners to produce targeted hardware that can handle the most challenging demands of video delivery today.

In partnership with SMART Embedded Computing and Gamestream, Intel will host a cloud gaming demonstration at IBC this year to show how AAA games can achieve low latency (25ms including 10ms network latency) with the SMART Embedded Computing MaxCore™ rendering system, which can host up to 15 SMART EC PCIE-7217 cards, supporting up to 58 concurrent gamers.

Open Software

Hardware is only one half of the video delivery equation; software is just as important. I wrote a recent blog about how focusing on bit rate reduction through software can result in drastic TCO benefits. Intel continues to invest heavily in several software projects that target key steps along the data journey.

Open Visual Cloud leverages a modular interoperable approach to software supporting core decode, inference, render, and encode functionality with optimized open source ingredients such as Scalable Video Technology (SVT), OpenVINO™ Toolkit, and the oneAPI Rendering Framework. In addition, several different sample reference pipelines for visual cloud services are provided as Docker files and images so that service providers can get started quickly to deploy customized services without further platform investments.

Intel is updating two new Open Visual Cloud reference pipelines (provided as Docker containers) for release at IBC 2019, including:
• Smart City: The smart city reference pipeline shows how to integrate various media building blocks, with analytics powered by the OpenVINO™ Toolkit, to accelerate converged media pipelines. This is a smart city planning sample that can be referenced to ease application development challenges.
• Ad-insertion: The intelligent ad-insertion reference pipeline shows how to integrate various media building blocks, with analytics powered by the OpenVINO™ Toolkit, to accelerate converged media pipelines. This is a server-side ad insertion sample that can be referenced by developers, now supported on the Intel VCAC-A accelerator card.

Intel® Select Solutions

Of course it's when hardware and software elements are strategically combined that the full potential of optimization can be realized. Intel® Select Solutions consist of optimized bundled hardware and software recipes to deploy workloads easily and meet performance thresholds for a premium customer experience. The first of these solutions for the visual cloud market was designed to meet the rapidly emerging needs for content delivery networks (CDNs). Intel® Select Solutions for Visual Cloud Delivery Network are now available through Quanta Cloud Technology.

At IBC 2019, we are continuing our commitment to these solutions by announcing that the new Intel® Select Solutions for Media Analytics powered by Intel® Xeon processors and Intel® QuickAssist Technology for software-enabled security, authentication, and compression will be available soon. Available in Base and Plus configurations, the Plus configuration will take advantage of the newly announced VCAC-A. Like all of the other Intel Select Solutions, the latest offering is a rigorously benchmark tested and verified solution optimized for real-world performance.

If you're not able to attend the conference in person, you can download the "IBC 360 Live" app on iOS App Store, Google Play, Oculus Go, or Oculus Gear VR Store to see presentations live or download after the event in 8k 360° VR thanks to a collaboration between several partners. For more information about Intel visual cloud solutions, visit intel.com/visualcloud.

Microsoft and ServiceNow Announce Strategic Partnership

Together the two companies will accelerate digital transformation for enterprise and government customers

ServiceNow moves workloads to Microsoft Azure for highly regulated industries

Redmond, Wash., and Santa Clara, Calif. — July 9, 2019 — Microsoft Corp. (NASDAQ: MSFT) and ServiceNow (NYSE: NOW) today announced a broader strategic partnership intended to significantly enhance the integration and optimization of the companies' products, platform and cloud capabilities. Through this expanded partnership, the companies will enable enterprise customers in certain highly regulated industries, as well as government customers, to accelerate their digital transformation and drive new levels of insights and innovation. And, for the first time, ServiceNow will house its full SaaS experience on Azure in addition to its own private cloud. The expanded partnership will elevate ServiceNow to one of Microsoft's strategic partners in its Global ISV Strategic Alliance Portfolio.

"There is an enormous opportunity for customers — including in the public sector — to apply the power of the cloud to become more efficient and responsive," said Satya Nadella, CEO of Microsoft. "Our partnership combines ServiceNow's expertise in digital workflows with Azure, our trusted cloud, so that customers can accelerate their digital transformation, while meeting their security and compliance needs."

"Expanding our strategic global relationship with Microsoft enables ServiceNow to more fully leverage and integrate our platform and products with Microsoft's leading enterprise technology and capabilities," said John Donahoe, president and CEO of ServiceNow. "Together, ServiceNow and Microsoft will help our enterprise and government customers accelerate their digital transformation, creating great experiences and unlocking productivity."

The expanded agreement builds on a partnership announced last fall by Microsoft and ServiceNow. As leading enterprise technology platforms, Microsoft and ServiceNow make it easier for customers to integrate and optimize across the two companies' products and platforms. By collaborating on next-generation experiences, Microsoft and ServiceNow will leverage technology to bring further cognitive services and intelligence to products across the Now Platform® with Microsoft 365 and Azure.

ServiceNow Selects Microsoft Azure for Certain Highly Regulated Industries

ServiceNow will use Azure Cloud as part of its preferred cloud platform for certain highly regulated industries, benefiting from Microsoft's deep expertise in data protection, security, and privacy, including the most comprehensive set of compliance offerings of any cloud service provider. ServiceNow will first be available through Azure Regions in Australia and Azure Government in the United States, followed by additional markets in the future.

With ServiceNow available through Azure Government, U.S. government agencies will be able to leverage the compliance coverage across regulatory standards available through Azure. Microsoft is committed to supporting the full spectrum of government data to help agencies quickly and easily achieve their necessary requirements. Azure Government was built specifically to address the capabilities, performance and compliance needs of U.S. government customers and their partners. Azure Government enables innovation with deeply integrated cloud services, data and advanced analytics, and an open application platform that provides the building blocks to rapidly develop, deploy and manage intelligent solutions.

The U.S. federal government continues to look to ServiceNow as a strategic partner as it modernizes its IT infrastructure and accelerates its use of modern technology to digitally transform how it operates.

Microsoft Selects ServiceNow to Digitize Its Workflows

As part of a separate transaction, Microsoft will implement ServiceNow's IT & Employee Experience workflow products across its own business to improve operations, enhance employee experiences, and deliver stronger business outcomes. With ServiceNow, Microsoft will bring even more digital workflows into its organization, so employees can spend less time on manual tasks.

About Microsoft

Microsoft (Nasdaq "MSFT" @microsoft) enables digital transformation for the era of an intelligent cloud and an intelligent edge. Its mission is to empower every person and every organization on the planet to achieve more.

About ServiceNow

ServiceNow (NYSE: NOW) is making the world of work, work better for people. Our cloud-based platform and solutions deliver digital workflows that create great experiences and unlock productivity for employees and the enterprise. For more information, visit: www.servicenow.com.

ServiceNow Named a Leader in Gartner Magic Quadrant for Integrated Risk Management (IRM) for the Second Year in a Row

ServiceNow recognized based on its completeness of vision and ability to execute

SANTA CLARA, Calif., July 18, 2019 — ServiceNow (NYSE: NOW), the leading digital workflow company making work, work better for people, today announced that it has again been positioned by Gartner as a Leader in the 2019 Gartner Magic Quadrant for Integrated Risk Management.[1]

ServiceNow Governance, Risk, and Compliance (GRC) transforms inefficient processes by combining security, IT and risk capabilities into a unified risk program built on the Now Platform. Continuous monitoring, prioritization, and automation ensure rapid identification and response to risks.

"Organizations are betting on digital initiatives to drive competitive advantage and they can't generate value if they don't take risks," said Sean Convery, vice president and general manager, ServiceNow Security and Risk. "The key to unlocking the strategic upside of this disruption is properly managing digital risk. ServiceNow GRC brings IT, security and risk teams together with a unified approach to managing risk so they can enable the business while handling risk with confidence."

A Foundation for Managing Digital Risk

According to Gartner, "Forward-looking organizations increasingly look for ways to improve the relationship between risk management investment and business outcomes. Striking a balance between taking risks and imposing controls requires risk management principles and insights into strategic decision making. It also means risk teams must develop digital capabilities to harness risk intelligence across the enterprise. Such a vision — supported by the right IRM program, processes and technology — is adaptive and well-suited to address new risks from cyber incident disruption, new regulatory obligation and the imperative to build customer trust."

More information: download the report 2019 Gartner Magic Quadrant for Integrated Risk Management

ServiceNow connects the business, security and IT within an integrated risk framework and on the same platform. With ServiceNow GRC, customers can:

- **Control risk exposure.** Use continuous monitoring and real-time dashboards to get actionable information about high-risk areas, noncompliance, vendor status, and significant audit findings.
- **Improve risk-based decision making.** Plan and make decisions more strategically with a single integrated risk management program. Simplify communication and use context to assess business impact and prioritize activities.
- **Increase performance.** Boost performance and productivity with consistent cross-functional automation. Reduce errors and give your team more time to focus on higher-value tasks.

Built on the Now Platform, GRC includes powerful capabilities that drive cross-functional communication and processes, including a single data model to eliminate silos, automated workflows to reduce bottlenecks, and knowledge management to manage policies in one location.

Gartner does not endorse any vendor, product or service depicted in its research publications, and does not advise technology users to select only those vendors with the highest ratings or other designation. Gartner research publications consist of the opinions of Gartner's research organization and should not be construed as statements of fact. Gartner disclaims all warranties, express or implied, with respect to this research, including any warranties of merchantability or fitness for a particular purpose.

About ServiceNow

ServiceNow (NYSE: NOW) is making the world of work, work better for people. Our cloud-based platform and solutions deliver digital workflows that create great experiences and unlock productivity for employees and the enterprise. For more information, visit: www.servicenow.com.

Evolution of Mobile Video

White Paper / Jul 2019 / ABI Research
Analyst: Michael Inouye / Content Manager: Eric Abbruzzese

1. TWO-SCREEN WORLD BECOMES MANY WITH 5G

The arrival of YouTube in 2005, followed by Apple's first iPhone in 2007, and Google's Android platform in 2008, serves as the preamble to what has become a two-screen video market. A decade and a half might seem like a long period of time, but put into context of television's history, it is a remarkably short timeframe. Color TVs, for example, arrived in 1953, and it wasn't until the late 1990s before HDTVs hit the market. Viewing behavior itself had remained largely unchanged for decades, where consumption of video largely occurred in designated areas of the house, quite often as a shared viewing experience among household members. Video sources were comparatively limited and curated, resulting in a generally homogenous viewing experience.

Streaming media began to "democratize" the viewing experience, shifting choice and controls from the content owners and service providers to the viewer. The introduction of other technologies like DVR and VOD services created further diversions from the previous viewing status quo. While content began to flow to more devices, including outside of the home (for example, workplaces), the smartphone and later the connected TV introduced what have become arguably the most significant changes to viewer behavior and the video market at large.

These shifts in the video market have yielded a more personal viewing experience that spans the entire spectrum of what, when, and where content is viewed. In mature markets like the United States, consumers have started to shift entertainment budgets from traditional pay TV services to OTT, and the MVPDs have responded by pushing TV Everywhere (TVE) services, and more recently OTT options, which address the changing viewing climate, but also extend the operator's customer reach beyond its network.

This interruption to the typical linear broadcast approach to support on-demand streaming has also impacted the hardware deployed by operators into households. STBs, for example, are no longer sold with just a broadcast stack, but now include platforms that help bridge the gap between traditional and OTT services (e.g., HBBTV, Linux, Android TV, etc.).

As the video market undergoes its current transformations, the arrival of new technologies—5G in particular— will usher in further changes that will see the market move from a two-screen paradigm (mobile and TV/CTV) to many displays. Within the home, the screen could extend to multiple devices including appliances, walls, windows, mirrors, and tables. Display technologies like flexible screens, transparent displays, and modular designs will fuel this expansion of screens within the home. In the public domain, digital signage along with new touchpoints like autonomous and public transportation will extend the screen beyond the home and workplace. Wearables, like smart glasses, will likely have the greatest impact by creating displays anywhere, making the screen virtually ubiquitous. The proliferation of physical displays will also occur over a long period of time due to cost and development/rollout of supporting technologies.

1.1. CONTINUITY OF CONTENT AND NEW OPPORTUNITIES

The road to this idea of the pervasive screen will require the development and integration of a range of technologies and markets to ensure continuity of content across displays to produce a seamless viewing experience. The seeds for this transition were sown with the arrival of the iOS and Android ecosystems, giving rise to the two-screen paradigm, and via application/content stores, served as a key role in the monetization of OTT content. Apple further had a significant role in establishing HLS as the predominant choice for ABR in today's streaming marketplace.

As a result, many aspects of the video market have already moved away from the linear viewing experience on predefined channels; in the future, however, the viewing experience will further transcend today's conception of TVE or portable/mobile viewing on the go, where content continuity more directly refers to content and service portability. The pervasive screen represents the blending of new elements being introduced into the video market today with technologies that are on the horizon.

... to next page

Chart 1: ***Ingredients of the Pervasive Screen***
World Market, Forecast: 2018 to 2024 *(Source: ABI Research)*

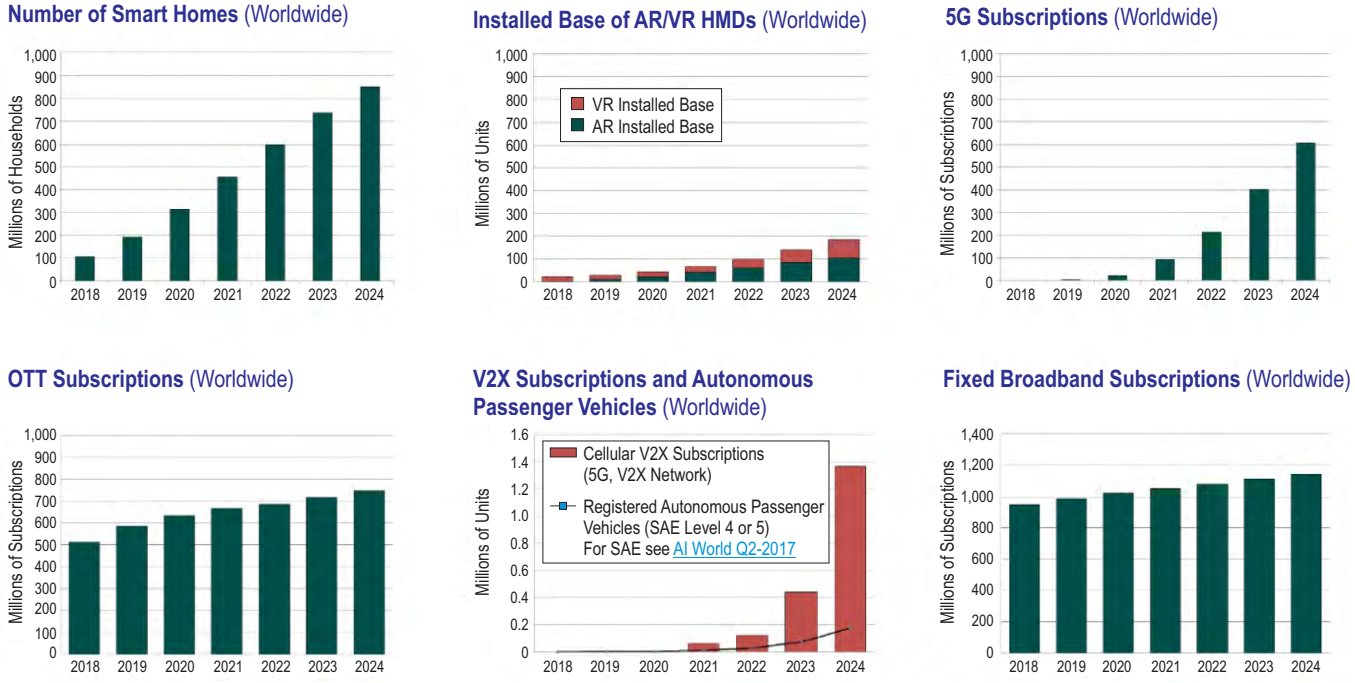

1.1.1.1. SPECTRUM OF THE PERVASIVE SCREEN

It is important to distinguish the two halves of the pervasive screen spectrum. The first half speaks to the pervasiveness of content, which started with the rise of mobile screens and will continue as additional displays are added to the mix, and the latter half points to the ubiquity of screens (and applications) through the addition of other technologies like Augmented Reality (AR) and computer vision. Some technologies like 5G are critical pieces of the fabric connecting and making these experiences possible and span (or will span) the entire spectrum of the pervasive screen.

There are a number of ingredients that need to come together to yield the pervasive screen, and many of the markets are at different developmental stages. Fixed and mobile broadband in some market segments are quite mature, while other market segments like AR and Virtual Reality (VR) are still in the more nascent stages of development. In addition, key networking technologies like 5G are only just starting to come to the market. The chart above shows a snapshot of some of these ingredients and their expected growth vectors over the next several years.

The development and formation of the pervasive screen does not have a specific completion or arrival date. Instead, it should be viewed as a rolling trend that is expanding and developing as new technologies and markets come onboard and integrate into the larger whole. The timeline on the following page illustrates when some of the markets above will start to have an impact on the pervasive screen. Note that this does not mean these aspects of the market are necessarily at mainstream adoption levels, but rather when they will start to have a larger role. Over time, the market will move from the two- or few-screen paradigm into a transitional stage where screens are added but not necessarily cohesively interconnected. The latter stages of the pervasive screen will see a stronger level of continuity and seamlessness between screens and content.

Figure 1: ***Timeline for the Pervasive Screen*** *(Source: ABI Research)*

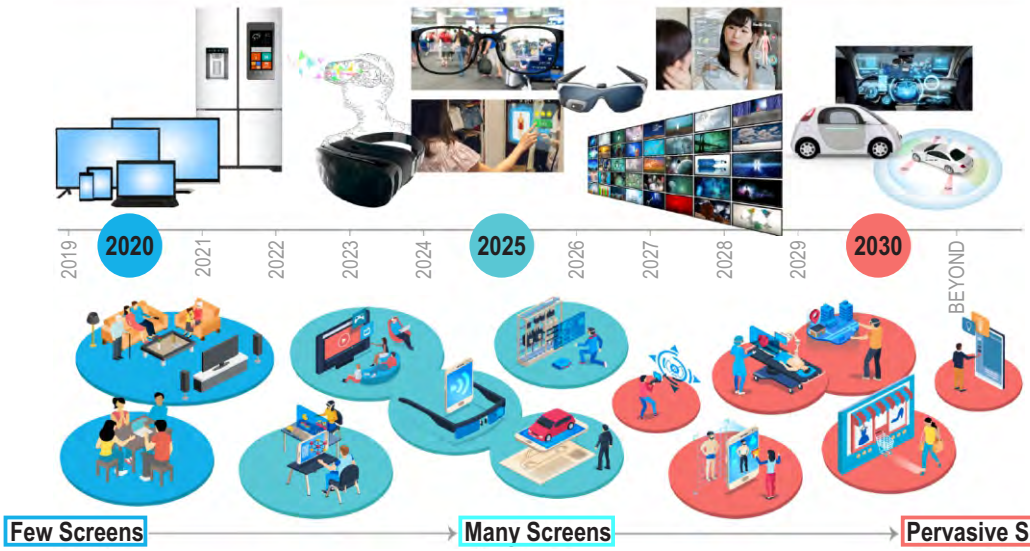

... to next page

... from previous page

1.1.1.2. IN THE HOME

Within the home, voice control has become a common user interface employed by connected CE manufacturers and pay TV operators alike. Silicon providers are also embedding AI into their solutions to address the interest from operators and hardware manufacturers that are making or are evaluating the move to virtual assistants (and other features) in order to integrate other platforms like smart home/security, information gathering, content recommendations, and advanced advertising. While privacy concerns remain a critical consideration, the value-added benefits will help with user opt-in. Tasks such as programming a DVR (including coordinating with personal calendars), reminders about favorite TV shows and sporting events, content recommendations, and integrations with smart home features (e.g., automatically starting a VOD movie and setting the lights) will help offset the loss of some privacy to enable these features.

In time, computer vision could play a larger role to augment some of these services and features. While this aspect of the home environment is more tenuous—attempts and trials by operators and CE manufacturers (e.g., game console manufacturers) in the past have not yielded strong results or support from users—it would greatly enhance the user experience while offering operators and services a wealth of user information critical for targeted advertising and affiliate commerce/marketing opportunities. For example, users in future applications could be identified through inference (e.g., facial recognition) to activate user profiles (for content and preferences) and proactively manage the users' digital entertainment to match with their personal preferences and schedules. A user could have a news feed playing on the mirror display in the morning, and during a commercial break ask for additional information about a watch advertisement, and then see it superimposed on their arm, either via the mirror display or the use of smart glasses. Unlike most targeted ads employed by operators and CTV platforms today, these advertisements could more frequently include a call to action like purchases (or additions to shopping carts) or directing users to brands' websites, allowing these types of ads to yield similar KPIs more commonly found on mobile and PC platforms, which would also help with cross-platform attribution.

The requisite camera/sensor for computer vision does present some potential hurdles and will not be universally accepted by all users or for all applications, particularly in today's climate where privacy is often placed front and center. While differences exist at a country level, where countries like China employ significantly more monitoring technologies to track their citizens (e.g., facial recognition, Internet monitoring, etc.), platforms more generally will have to work across different forms of recognition to identify users to make the experience transferable between devices or displays and acceptable to the widest number of users. This could entail a beaconing system or voice recognition, assuming that the virtual assistant is used for the majority of the user experience. This is an area where AR could help mitigate some of these privacy concerns since the user can better define when and where to use these features. If a user wanted to ensure that the camera (or passive listening device in the case of virtual assistants) was not active in a certain room or for a particular use case, the smart glasses could simply be removed and pocketed or turned off. While these options exist for fixed displays and devices as well, the mere presence of a camera (even if it is off) could still create discomfort among some users. Despite these hurdles, the value-added features and benefits will find acceptance among users as a tradeoff for personal data as has historically occurred with affinity cards for retailers, tracking across the Internet, and social media.

The image on the following page illustrates some examples of in-home applications for the pervasive screen, both in the context of video and the integration of other features and services like the smart home. Physical displays beyond the TV and mobile devices will enter the home over time and are anticipated to develop at a slower rate than AR due to higher prices and need for technologies (e.g., transparent displays) to come to the consumer market.

Figure 2: ***Examples of the Pervasive Screen in the Home*** *(Source: ABI Research)*

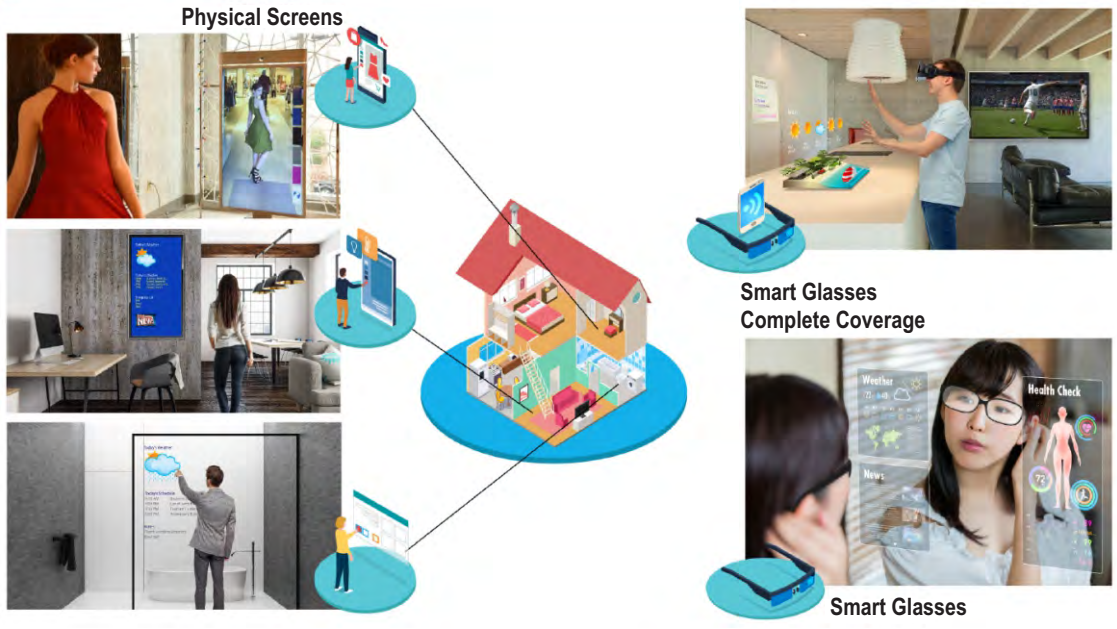

... to next page

... from previous page

Within Connectivity, particularly as more displays are added to the household bandwidth, will also come to the fore-front within the home. To accommodate the growing number of connected devices, operators and vendors are already marketing more advanced Wi-Fi networking within the home—for example, pushing Wi-Fi mesh and network management. As video quality increases, coupled with new displays and touchpoints, the need for greater bandwidth could push some households to consider fixed wireless alternatives like 5G as a way to create a cohesive data plan across devices both in and out of the home. While an option, most households will continue with traditional fixed broadband services, although fixed wireless will play a stronger role in rural areas where fixed broadband infrastructure is too costly to deploy.

1.1.1.3. OUT IN PUBLIC AND ON THE GO

At the most basic, users will access services and content on displays at multiple touchpoints in the public environment. This could occur via a mobile device or further into the future through displays in public transportation (including autonomous vehicles). In these instances, many engagements will stem from preexisting SVOD subscriptions or purchased content, and later from mass transit operators selling access to video services/ content and using these displays for informational and/or advertising purposes. As these opportunities grow, 5G will play a critical role to meet capacity needs and to provide a consistent service throughout the user's journey by maintaining network uniformity and through simultaneous multi-connectivity (multiple 5G access points, 4G, and Wi-Fi).

Figure 3: ***Examples of the Pervasive Screen in Public*** *(Source: ABI Research)*

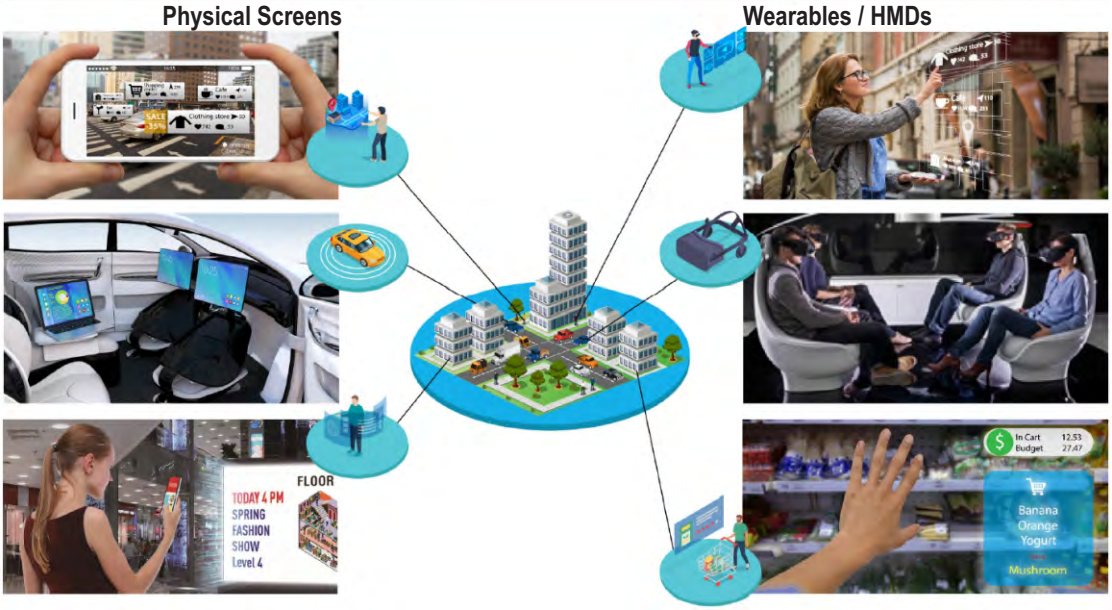

In highly dense urban environments and locations (e.g., sporting venues), mmWave 5G will provide the necessary capacity (both up and down) to support the demands of the public. For example, users at a sporting event will be able to record and share content with family and friends, which requires network capacity and low latency. The potential also exists to deploy interactive elements to the stadium experience—for example, using AR to show stats, highlights, and introducing some gamification like trivia. As evidence of 5G's efficacy to meet the demands of sporting events (e.g., high-quality video and low latency), field trials have already been conducted using 5G to move significant parts of a broadcaster's production operations to a remote offsite location, saving cost and consolidating workflows. Huawei also announced plans for a 5G 8K TV, perhaps more as a prelude of what is to come (or possible) with 5G and the distribution/reception of data-intensive content like 8K and 360 video/VR. The TV could also serve as a router for LAN access, making it an entry point for fixed wireless broadband.

5G will also engender opportunities for localization. While these applications exist today (e.g., via Bluetooth beacons, GNSS, Wi-Fi, cellular), 5G has the potential to offer better localization and richer, more cohesive experiences as part of a hybrid (existing technologies) positioning system. These applications reach both the consumer and enterprise markets, where the latter could extend to a range of applications including automotive, industrial and UAV operations.

On the consumer front, receiving consumer opt-in remains a priority, and many of the core features (e.g., tracking location, collecting visual/auditory data, and receiving requests by companies for further opt-in) will need to be activated at the platform (e.g., Android, iOS) or browser level where consumers will later grant individual opt-in to companies based on proximity, in a similar fashion as websites do when informing users about or requesting opt-in for cookies. As an example, users could be reminded about goods they may have added to a virtual cart but did not complete the transaction. Users could also be guided to products of interest within the store as well. Regarding products within a store, users could receive prompts to get additional information, see ads/promotions, or view trailers for items related to movies or animated features playing in the theaters. AR could also render virtual billboards and signage, with customized advertisements and displays. This might require visual tags and computer vision or indoor navigation for interior environments, but AR and mixed reality is a critical ingredient to these more advanced and interactive features because it enables seamless transitions between environments and interactions without requiring the user to purposefully take out and use a smartphone.

... to next page

1.1.1.4. COMPUTER VISION IN OTHER APPLICATIONS

The focus of this paper is on the video and entertainment markets, but 5G will present new opportunities for video within the commercial and enterprise verticals as well. Examples include:

- Communications between field workers and remote staff
- Remote viewing, monitoring, and/or control of devices (e.g., UAVs)
- Uploading content and coordination between devices (e.g., surveillance systems)

While privacy concerns and regulations may limit computer vision in some applications like surveillance (outside of countries like China), 5G connectivity could still connect networked surveillance cameras and support the need for higher data throughputs and low latencies. High-resolution surveillance cameras, for example, could cover a wider field of view, and therefore require fewer installed cameras, albeit requiring a more robust data connection. In applications such as this, the generation and flow of video does not always go to a screen; in many cases, the video is automatically analyzed by a computer vision algorithm and then stored in a NAS or SAN.

2. MEETING THE QUALITY OF EXPERIENCE NEEDS OF THE PERVASIVE SCREEN

2.1. NEAR-TERM QoE CONSIDERATIONS

As more households shift viewing time to streaming media, the video industry is pushing to bring these experiences closer to parity with current premium pay TV services. Efforts here have brought elements like low latency to the forefront. Currently, common streaming latencies are in the range of 30 to 60 seconds using ABR (e.g., HLS), with most broadcasts at sub-10 seconds (ignoring intentional delays from live to allow for censorship/editing). Matching broadcast levels is particularly important for sports programming where key moments could get ruined if a stream is lagging significantly behind broadcasts, in the case of others cheering nearby or through social media notifications. To reach these latency levels, companies are employing different techniques; most, however, are using chunked transfer encoding such as using shorter HLS chunks, which extends to standards like CMAF (standardized packaging), which helps to reduce packaging costs/time.

For a large majority of video viewed by consumers, hitting these levels of low latency is all that is needed. There are forms of video today, however, that benefit from lower, ultra-low, latencies. Video related to auctions and gambling, for example, while niche, often needs as close to real-time performance as possible; real-time communication like video conferencing by nature benefits from as low latency as possible. The low latencies touted by 5G, however, go well above and beyond what is necessary for most of today's streaming media.

5G will help with capacity issues and in time support higher bitrate videos like UHD when mobile devices make the transition to higher resolutions. It is worth noting, however, that the increased availability of UHD content, while certainly a driver for higher data consumption, will not necessarily push data consumption up at an equitable level as it becomes available. Companies today are already considering and offering solutions to cut costs while delivering an optimal viewing experience. For example, if a user won't notice the difference in quality between a 4K or 1080p stream, the user will receive the latter in order to save on transit and egress fees. These cost-saving measures are necessary since video will account for the majority of data (fixed and mobile). It is also important to note that the mix of video data is not uniform across regions, particularly in the splits between mobile and fixed viewing; in some countries within Asia-Pacific, for example, it is more common to see mobile-first services and viewers who primarily consume video over these devices. In these cases where mobile viewing is particularly high, 5G could have a more immediate need to address these volumes.

While the highest data bandwidths promised by 5G may remain the domain of densely populated urban centers (where mmWave 5G and its massive MIMO antennae arrays are more financially viable), it will help rural and less populated areas to receive higher data rates where fiber and related fixed infrastructure is too costly. Similarly, in homes with enough displays and connected devices to strain their current broadband service, these households may also consider upgrading to wireless 5G instead of moving up to more premium fixed broadband tiers. Positioning 5G as a premium service with very high, if not "unlimited," data caps could also lure in households that consume a particularly large amount of data. For telcos that are invested in content and streaming services, the wider reach of fixed wireless broadband services will also create more opportunities to sell triple-play packages.

2.1.1.1. CLOUD GAMING

Other markets, like gaming, will soon place additional demands on data traffic as the industry collectively makes a migration to cloud gaming. While not occurring in the next generation of consoles, more cloud gaming services will launch in the near term, which will in time usher in a more comprehensive shift to the cloud. While the cloud gaming experience from the user's perspective is ideally indistinguishable from a console or PC, the content workflows (once it has been rendered, encoded, and packaged) are based on video streaming technologies. This allows the games to reach a wide breadth of screens, much like typical streaming video (hence, a good fit for the pervasive screen), but due to gaming's interactivity/controls, it further pushes current streaming technologies towards ultra-low latency. Gaming is also starting to embrace more continuity between platforms, not unlike streaming video services. There are increased cases of cross-platform game availability and even multiplayer support that spans mobile, PC, and/or console. This allows a game and the users' profile and experiences to remain connected in the home on a PC or game console (using IP network) and then on the go via a smartphone or tablet (e.g., via 5G). Usage habits and expectations, however, also differ between gamers and typical video streaming.

... to next page

The In the home, video streamers often binge-watch during the weekends and have more typical viewing rates during the week. Gamers, on the other hand, can spend considerably more hours throughout the week on their hobby. Gamers will also demand higher bitrates in order to support higher resolutions (new cloud gaming services will support up to 4K gaming) and higher framerates. The video game industry is progressing faster towards higher visual quality than streaming video, and cloud gaming will reflect these trends. Cloud gaming will also drive more traffic outside of the home as consumers will use mobile displays while on the go. Latency is a critical area for cloud gaming and was a factor that hindered the adoption of services in years past.

Despite the recent attention that cloud gaming is receiving, its roots trace further back to as early as E3 2000 where Finnish company G-Cluster presented its cloud gaming technology. Previous attempts at cloud gaming suffered from poor experiences, and in many regards, it was a solution ahead of its time. Improvements to existing streaming protocols, coupled with 5G, will allow services and operators to offer consumers the best cloud gaming experience, creating a point of differentiation and a monetization opportunity for 5G. This extends beyond just latency to include the continuity of the gaming experience, being able to follow a user from in the home and into the public, be it in a coffee shop, public transportation, or in a future autonomous vehicle. While the demands for mobile data will certainly grow, there are some factors that could accelerate this demand if adoption of technologies like AR and VR move faster than anticipated.

2.1.1.2. AR/VR AND ULTRA-LOW LATENCY IN THE ENTERPRISE

AR and VR have not resonated with consumers as many thought would happen when this current generation of devices launched. There are many reasons for the slower uptake of these devices; high pricing and limited content in the early stages, for example, tempered some excitement. Like 3D, there are also individuals who remain reluctant to use a wearable like glasses or a Head-Mounted Display (HMD) (or found the form factors disagreeable), and for others, some have cited poor resolution and lack of reactivity as a cause for dizziness/discomfort. The slower adoption of these technologies has curtailed investment and interest by many companies throughout the value chain in the media and entertainment markets. Many of the issues most cited by users today are getting addressed via updated hardware and the introduction of new technologies/standardizations. Looking at a more extended time horizon, the possibility also exists for devices that offer a mixed or combined experience of AR and VR. These devices would allow the user to select between an AR experience, which combines the real world with virtual elements, and VR, which isolates the user.

Within the nearer term, focus has shifted to the enterprise markets where these devices have gained better traction for applications like training, remote assistance, and design. While many of these use cases are currently deployed and used on premise with limited needs for high data rates, some applications such as field workers using remote assistance or monitoring/ inspections will make use of advancements made in wide-area connectivity l ike 5G where low latency and reliable connectivity are more critical factors.

2.2. LONGER-TERM QoE CONSIDERATIONS

As the more advanced elements of the pervasive screen come online, the need to establish an optimal Quality of Experience (QoE) will shift towards edge computing and ultra-low latencies. In this segment, there are two key markets to consider and their impacts on the consumption of video: autonomous transportation and AR/VR.

2.2.1.1. AUTONOMOUS TRANSPORTATION AND VR

While many within the automotive industry have pushed targets out for more mainstream presence of autonomous vehicles (SAE Level 4 to 5 for high and full automation respectively), this mode of transportation will in time have an impact on the video market. Displays in these vehicles will provide opportunities to access video (and gaming) content, serve as a marketing platform (particularly for local businesses), and provide a new avenue for VR use outside of the home or office. In many regards, the screens used in automobiles will mirror the developments within the home as screen time is split between physical displays (in-car and mobile devices) and wearables (i.e., smart glasses and VR HMDs). The primary difference is movement, which in some cases means adapting the entertainment content to the driving (as a passenger) experience—for example, movements from the vehicle could be reflected within a gaming environment to help minimize motion sickness.

... to next page

The In all cases, network uniformity is necessary to ensure a consistent and uninterrupted experience, particularly critical for gaming and VR applications. Cloud gaming and VR also require very low latencies; for VR applications that use an HMD, the Motion-to-Photon (MTP, or time between user input and display update) latency needs to fall under 20 ms, if not 15 ms (lower is better, of course). Cloud gaming latency requirements vary on the type of gameplay, but at a minimum, a service catering to more casual players would need to support sub-150 ms latencies, and ideally target below 100 ms. If the game service is delivering game types that require even lower latencies, like First-Person Shooter (FPS) or racing games, then latencies need to target sub 50 ms, and ideally match the targets set for VR at 20 ms or less. 5G and edge computing can deliver these lower latencies and meet the bandwidth and network uniformity requirements. Hitting the lowest levels of latency may require the edge, but the telco network will support a wide breadth of use cases and applications to satisfy most users.

Bandwidth requirements for these services could also reach high levels, particularly for cloud VR applications, which could range from 50 Mbps to 1 Gbps and higher depending on the complexity of the user experience and type of content. At the higher data requirements, 5G is again the best candidate to bring these types of services.

Localization will also play a role where displays in autonomous vehicles will provide both marketing and upsell opportunities for mass transit operators and local businesses. For example, a passenger travelling from work to home may receive ads or promotions for food places along the route, and the user could accept a promotion and the vehicle could drive to that location to pick up food before resuming the journey to their destination.

2.2.1.2. AR AND SMART GLASSES

Moving the experience down to the pedestrian level creates the need for further refinements to localization (particularly indoor) and ultra-low latencies for vision-related applications. Positioning accuracy (e.g., relative versus absolute) will depend on the applications and use cases, but 5G coupled with IoT implementations will address these needs, particularly in indoor environments. Computer vision applications will also benefit from 5G for uninterrupted connectivity and hybrid computing for applications like deep/ machine learning for computer vision and large databases.

It is worth reiterating here the need for platform support in order to make opt-in and use of these features as efficient as possible. Ascribing to an app model, where the user downloads individual brand/ retailer apps to enable these features, will not generate enough traction among the user base. Rather, global opt-in for the tracking and computer vision must occur at the platform level and then individual brand/retailer opt-in can occur in two parts—upfront checklists, and then organically as the user navigates public spaces and receives opt-in requests.

3. MARKET OPPORTUNITIES

The pervasive screen will create opportunities throughout the video value chain and incentivize additional investment from external markets. Brands and retailers are already increasingly allocating resources to video, be it for advertising/brand awareness or internal communications, and the pervasive screen will accelerate these investments, particularly if advertising and promotional opportunities extend to public areas and personalization yields higher CPMs and other ad rates.
Companies operating within the video industry will help their customers deliver video at proper latencies and quality depending on target screens and applications. This will become considerably more challenging as displays come to have less standardized sizes/ratios and resolutions, particularly among modular and virtual displays. For services and operators, the pervasive screen will generate additional data to enhance content recommendations & personalization, present new marketing & sales opportunities & tighten integration into other areas like smart home.
Some aspects of the pervasive screen, like the spread of physical displays within the home, will take time to develop due to cost and the need for advancements in some technologies, but more near-term developments in AR and 5G will introduce both the concepts and value propositions necessary to springboard the market forward. In the meantime, the two-screen market will continue to serve as the status quo, and market updates like 5G, while critical for many of the more advanced elements of the pervasive screen, will afford services and operators the necessary headroom to meet the needs of their customer base as consumption continues to grow. Cloud gaming also represents a key nearer-term market opportunity for telcos, allowing these operators to push more advanced services, reduce churn, and grow ARPU. Cloud gaming (due to its cloud nature) also fits well in the evolving landscape where content will flow more seamlessly between devices, platforms, and locations, be it in the home or on the go.
The pervasive screen will also offer consumers tangible benefits beyond just faster downloads and higher data limits from 5G; although, for some markets, like cloud gaming, higher data caps (or unlimited) would certainly make fixed wireless an alluring alternative. Positioning, network uniformity, and ultra-low latencies will help usher in new applications to further 5G's (and related markets) value. While many of these aspects of the market will take time to develop, this affords the operators adequate time to fully deploy 5G services to ensure the technology is ready to deliver the experiences envisioned by the pervasive screen.

Google Cloud Adds Compute, Memory-Intensive VMs

Jessica Lyons Hardcastle | Managing Editor at SDxCentral
August 14, 2019 6:48 PM

Google added virtual machine (VM) types on Google Compute Engine including second-generation Intel Xeon scalable processor machines and new VMs for compute- and memory-heavy applications.

The former, available in beta, are general-purpose VMs. They provide greater than 20% price-performance improvement for many workloads and support up to 25% more memory per virtual CPU compared with first generation machines, according to Google. These N2 VMs offer a balance of compute, memory, storage, and network resources for general-purpose workloads such as web and application servers, enterprise applications, gaming servers, content and collaboration systems, and most databases.

They are available in Google's U.S.-Central, Europe-West, and Asia-Southeast region now and will be available in most Google Cloud Platform (GCP) regions in the next few months.
The second type of new VM is optimized for compute-intensive applications such as high-performance computing (HPC), video encoding, and massive multiplayer games. These C2 VMs provide the highest performance per core available on Google Cloud, according to the vendor.

IT consulting firm Burwood Group is an early customer, and said that after moving one of its customers to C2 VMs it saw a 40% improvement in performance, while using less hardware. "One of the critical jobs/processes took 82% less time to complete at 42% less cost," said Dan Speck, vice president of technology research and development at Burwood Group, in a blog post.

These VMs are available in four GCP regions: U.S. Central, Europe-West, Asia-Northeast, and Asia-East, with expansion planned for U.S.-West, U.S.-East, and Europe-West in the coming weeks.

The third, new memory-optimized VMs, are available in beta and are SAP Certified. Google added 6 terabyte (TB) and 12 TB VMs to its memory-optimized machine types (M2), which it says makes them the largest SAP Certified VMs available from a public cloud provider.

They also include cloud features such as live migration and committed use discounts to migrate growing SAP databases from a 4 TB instance to the new 6 TB VM while keeping the current price.

Jessica Lyons Hardcastle | Managing Editor

Jessica is Managing Editor at SDxCentral, covering next-generation data center technologies and security. She has worked as an editor and reporter for more than 15 years at a number of B2B publications including Silicon Valley Business Journal, Environmental Leader, and Solar Novus Today.

Network Automation with NETCONF & YANG - What you need to know first

- CLI DISADVANTAGES: CLI INPUTS(commands) & OUTPUTS differ from vendor to vendor, and more
- SNMP DISADVANTAGES: Unreliable as it inherently uses UDP as its transport protocol, NO rollbacks etc
- NETCONF is an IETF standard with broad market adoption incl. Cisco, Juniper, Huawei and many others
- Discover more articles www.fir3net.com - www.ixiacom.com and video's youtubeV55Y7ibjXcw cooE3wZ7O4I

NETCONF and YANG why you need it

By Tomas Hedqvist, Enea Software

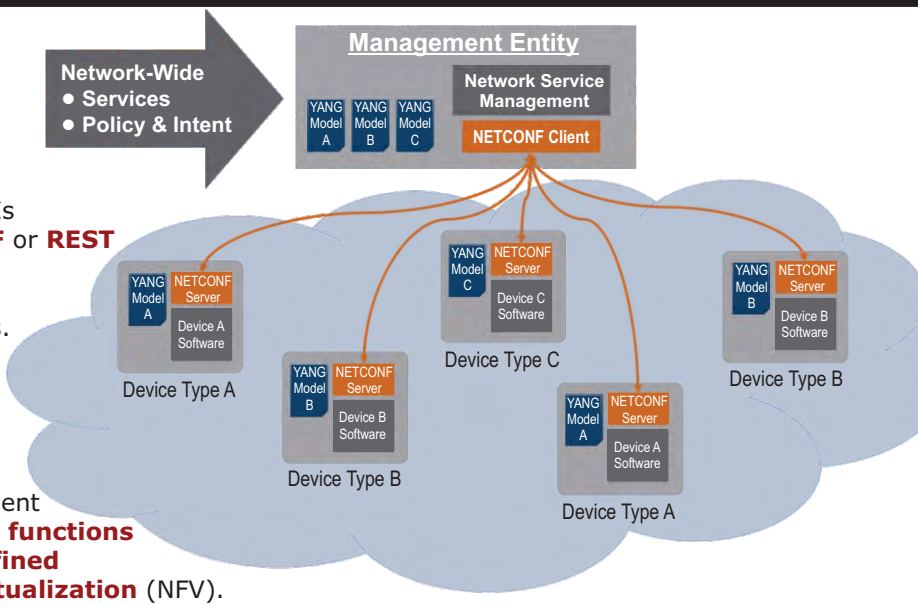

How do you manage your devices? Still using CLIs and scripting? Or have you considered **NETCONF** or **REST**

Nowadays you have a lot options when choosing management interfaces, some better than others. You have a lot to gain from choosing the right management protocols and data models if you want to automate network management, or ensure your devices can be managed in multi-vendor networks. You have even more to gain from choosing a modern, standardized management protocol if you want to manage **virtual network functions** (VNF) and leverage the benefits of **software defined networking** (SDN) and **network functions virtualization** (NFV).

I have listed four questions that could help you choose network interface and decide if NETCONF and **YANG** is something you need.

1 - Do you need to configure a network or just a device?

In very small networks with up to a few dozen devices, managing them one by one may not be a very big concern. But in larger networks, with tens of thousands of devices, it sure is. Scripting has been the traditional way to solve this but is unreliable and definitely not standardized.

Configuration changes are best performed using transactions. Using transactions, you make sure that either the complete configuration, consisting of a sequence of changes, is applied or rolled back entirely. No half-done configuration changes in other words.

REST APIs works well for machine communication, especially when connecting to cloud services, which makes it useful for integration with orchestration. But it does not provide transactions, so it is a less optimal choice for device management.

NETCONF on the other hand is designed for device configuration and provides network-wide transactions, making it possible to change configuration for all devices in the network at once. A change will either succeed on all devices or it will not be activated at all. This way you make sure you have the same configuration in all devices.

2 - Do you need a standardized management interface?

Most networks, at least larger ones, are built with devices from multiple vendors. No need to say, standardization makes it much easier to manage multi-vendor networks.

Standardization starts with the data model. YANG, for example, is a standardized modelling language possible to use with both NETCONF and REST interfaces. Both configuration data and state data is modeled with YANG, providing a standardized way to describe devices. Northbound interfaces can be automatically rendered from the model ensuring consistency between the device and the management entity.

REST however is not a single, standardized protocol, but an architectural style. So while all REST APIs follow the same principles and guidelines, they are essentially proprietary.

NETCONF is an **IETF standard** with broad market adoption. Many service providers and network operators even require devices to be manageable using NETCONF and YANG if they are to be installed in their networks.

3 - Do you need robust management that doesn't fail?

As mentioned NETCONF uses transactions for configuration changes, making network management more robust.

Another property if NETCONF and YANG is that you as network manager does not have to keep track of the order in which changes are done to the device. This minimizes the risk of changes being committed in the wrong order, causing a faulty state in the device.

Devices modeled with YANG can also include constraints and data validation rules, making it possible to automate validation. This is important because you can both validate the model itself, and the data you put into it before you push it to a device.

... to next page

NETCONF and YANG why you need it ... from previous page

NETCONF is designed around a configuration datastore concept, making rollbacks possible and even automatically applied if the device loses connection with the management entity because of a faulty configuration change.

4 - Do you want to automate?

With more and more network functions being virtualized, the number of devices and their locations are not static. If you would have to manually configure virtualized devices you will lose much of the agility you were looking for in the first place when virtualizing. But automation is important also with physical devices.

When you automate management you bring down opex and add agility. But remember that CLI:s are for humans, not automation. Automation needs an API making the device programmable, and that API can be provided by a YANG model and accessed over NETCONF or a REST API. YANG, because it is a standardized modelling language, makes it possible to manage multi-vendor networks with a high degree of automation.

REST, coming from the cloud end of things, integrates superbly with cloud services, making it a suitable option for integration with orchestration, but not directly from a managed device. Southbound from an EMS/NMS, NETCONF is the better alternative, but northbound, REST is a perfect option.

So what now?

If you want to know more about software for network device management interfaces, you should definitely read this report: **Selecting NETCONF/YANG management interface software**.

Huawei's Network Cloud Engine Completes EANTC's PCEP and NETCONF/YANG Interoperability Tests

[Paris, France, April 10, 2019] Huawei's **Network Cloud Engine (NCE)** has achieved a strong performance and completed multi-vendor interoperability tests in both the **Path Computation Element Communication Protocol (PCEP)** and **NETCONF/YANG** fields, according to the 2019 Interoperability Test White Paper released by the **European Advanced Networking Test Center (EANTC)** at **MPLS+SDN+NFV World Congress 2019**.

The interoperability test brought together 20 vendors and focused on PCEP, NETCONF/YANG, Segment Routing (SR), Ethernet VPN (EVPN), clock, and microwave. In the PCEP field, Huawei NCE functions as a path computation element (PCE) to compute intra- and inter-domain paths and optimize paths in interoperability tests with multi-vendor devices. In the NETCONF/YANG field, Huawei NCE functions as a NETCONF/YANG client to create VPN services and complete NETCONF/YANG configuration in interoperability tests with devices. Huawei NCE is the software product that has completed both the PCEP and NETCONF/YANG tests, achieving exceptional test results that show NCE's capability in terms of network openness and function integrity.

"The Huawei NCE implementation has shown great multi-vendor interworking," said Carsten Rossenhoevel, Managing Director, EANTC. "NCE has impressed us with regards to the interoperability of its PCEP and NETCONF/YANG implementation. This success reflects the consistent commitment of the NCE team to multi-vendor evaluation at EANTC over multiple years."

Guo Dazheng, President of Huawei NCE-Data Communication Domain, said, "Huawei actively participates in interoperability and openness tests initiated by industry organizations and has participated in EANTC interoperability tests for more than 10 years. At this test event, Huawei NCE has successfully interconnected with multiple vendor devices, which fully reflects Huawei's strength and determination to promote and develop cutting-edge technologies. "Huawei will continue to work with EANTC and industry partners to promote relevant technical standards and commercial processes, build a more intelligent and open network, and allow SDN-enabled networks to enter the autonomous driving network era, thereby creating a better connected world with optimal experience."

NCE is the industry's first automation network platform that integrates management, control, analysis, and AI functions. It additionally provides Design Studio, which is an open and programmable integrated development environment, and a developer community; and also connects to third-party network controllers or network devices in the southbound direction.

To date, Huawei has completed integration certification or interoperability testing with more than 40 industry partners and players. Based on the unified open NCE platform, Huawei works together with partners in upstream and downstream industries to create an open industry ecosystem.