SOLUTION BRIEF

Telecommunications
Visual Cloud

Unleash Innovation in the Visual Cloud

Collaboration between ATEME* and Intel is creating a new era of data center software and hardware technology that can improve Visual Cloud cost efficiency and enable the world to create, connect and immerse at new levels

Solution Benefits

- **Low total cost of ownership (TCO).** Modern Visual Cloud software can decrease content distribution costs without damaging video quality; combined with the latest data center technology, the result is significant savings.
- **High-quality, low-latency video and happy customers.** Consumers using a wide range of devices appreciate the premier viewing experience provided by powerful Intel® processors and intelligent transcoding applications such as those developed by ATEME*.
- **Great scalability and flexibility.** ATEME's software and Intel's data center technology provide a scalable solution that can be upgraded or adapted quickly to meet demand.

Authors

Nagesh Puppala
Media Segment Director
Network Platforms Group, Intel

Remi Beaudouin
Chief Strategy Officer
ATEME

Executive Summary

With an apparently insatiable appetite for video, consumers are putting pressure on communications service providers (CoSPs) to provide increasing amounts of live streaming video, video on demand (VoD) and immersive experiences such as augmented reality (AR) and virtual reality (VR). Cloud gaming, cloud graphics and media analytics are all contributing to the growth of the Visual Cloud—a set of capabilities for remotely consuming content and services that center around the efficient delivery of intelligent visual experiences. And yet, as network traffic grows, CoSPs struggle to cost-efficiently distribute, store and analyze these massive amounts of content.

Intel, ATEME* and other Visual Cloud ecosystem members are working together to remove the roadblocks to help CoSPs focus on services innovation and differentiation. For example, ATEME has developed Content Adaptive Encoding in its TITAN* transcoding solution to lower distribution bit rate and storage requirements, helping CoSPs lower total cost of ownership (TCO). Intel has also contributed greatly to Visual Cloud software efforts, such as the Scalable Video Technology-AV1 (SVT-AV1), Open Visual Cloud, Intel® Select Solutions for Visual Cloud and optimized plug-ins for popular frameworks and rendering applications.

Of course, software alone is not enough to power the Visual Cloud. New hardware technology innovations such as workload-optimized 2nd generation Intel® Xeon® Scalable processors, Intel® Optane™ DC persistent memory and Intel® Deep Learning Boost (Intel® DL Boost) provide the compute, memory and artificial intelligence acceleration that emerging Visual Cloud use cases demand.
The Visual Cloud will inevitably expand with the advent of 5G. The combination of innovative software and hardware together with teamwork throughout the ecosystem will drive cost efficiency and power new levels of content creation, connectivity and interactivity.

Figure 1. Forward-thinking communications service providers (CoSPs) are using modern software and hardware from ATEME* and Intel to meet the challenges of the Visual Cloud.

Business Challenge: Efficient Delivery of Intelligent Visual Experiences

CoSPs are expected to deliver visual experiences that are richer, highly immersive, interactive, personalized and available from the cloud anywhere, anytime and on any device. 82 percent of internet traffic will be video traffic by 2022[1], and mobile video traffic is expected to grow 10X between 2018 to 2023[2]. Consumers also expect low latency—video watchers hate jitter and buffering delays. Another challenge for CoSPs is handling increasingly complex codecs that provide excellent compression but also require extensive computing power.

All of the above raises pressing concerns about costs. Increasing processing, storage and bandwidth costs can drive up total cost of ownership (TCO), jeopardizing a CoSP's financial viability. This is especially true as the number of simultaneous viewers spread around the globe continues to rise. Sending thousands of MBs to millions of viewers increases distribution costs for streaming video services; storing a large number of pieces of content for a video-on-demand (VoD) service can escalate storage costs considerably.

In the face of ever-growing visual content that threatens to overwhelm legacy network infrastructures, CoSPs seek a cost-effective solution that enables them to meet these challenges. They need:

- Advanced networks and smart edge capabilities
- Scalable platforms and modernized software that enable the distribution of content close to consumers to minimize latency
- Powerful processing farms that can handle new complex codecs and improve the user experience
- The ability to reduce the size of the content to reduce TCO

Five Popular Visual Cloud Use Cases

The Visual Cloud relies on four basic building blocks (see Figure 2): decode (decompress the video stream); infer (analyze the video stream); render (draw or augment the video stream) and encode (compress the video stream).

These building blocks can be used to build a limitless number of services; here are five common examples:

- **Media processing and delivery.** The backbone of the Visual Cloud, this service streams content by decoding it and then encoding into different resolutions (480p, 720p, 1080p, 4K etc.) at various bit rates (7Mbps, 4Mbps, 2 Mbps etc.), then packaging it into various formats (HTTP Live Streaming (HLS), Dynamic Adaptive Streaming over HTTP (DASH)) that are supported by the different devices that consume the content. This service uses only two of the building blocks (decode and encode).
- **Media analytics.** This service uses inference (machine learning and deep learning) to extract information from a video stream. For example, looking for specific images or patterns. This service uses three of the building blocks (decode, infer and encode).
- **Immersive media.** AR and VR are rapidly growing use cases for the Visual Cloud. This service decodes the content, infers something from it, renders something—potentially adding something on top of the content—and then encodes it to send it to the individual users. This service uses all four building blocks.
- **Cloud graphics.** Cloud-based virtual desktop infrastructure (VDI) is an example of a cloud graphics use case. With cloud-based VDI, enterprises can easily provision virtual, cloud-based desktops for users, providing them access to the documents, applications and resources they need, anywhere, anytime, from any supported device. This service uses the render and encode building blocks.
- **Cloud gaming.** The global cloud gaming market size is expected to exceed USD 2.6 billion by 2022[3]. Gamers like cloud gaming because with cloud-streamed video, they can play games on a low-end laptop like it's an expensive gaming PC, as well as on-the-go with mobile devices. Like cloud graphics, this service uses the render and encode building blocks.

Solution Value: Bandwidth-Efficient Processing with Software

ATEME, a company that designs software video compression solutions to enable innovative live or on-demand video delivery services while preserving picture quality, has a long-standing history of collaborating with Intel on video compression. Both companies believe that software-based video solutions that take advantage of the power of new Intel® technologies can unleash a new level of innovation in the Visual Cloud. Working together, the two companies design, optimize, and deploy Intel® architecture-based bandwidth-efficient compression solutions for content providers and service providers.

Although ATEME's Visual Cloud solutions do sometimes incorporate specialized hardware such as Intel® FPGAs, the company primarily focuses on using software to decrease the size of the content, leading to cost-efficient video processing. In particular, ATEME has developed Content Adaptive Encoding for over-the-top (OTT) delivery, based on the following properties:

- Dynamic choice of profiles is based on an analysis of the content complexity, leading to optimized video quality and compression.

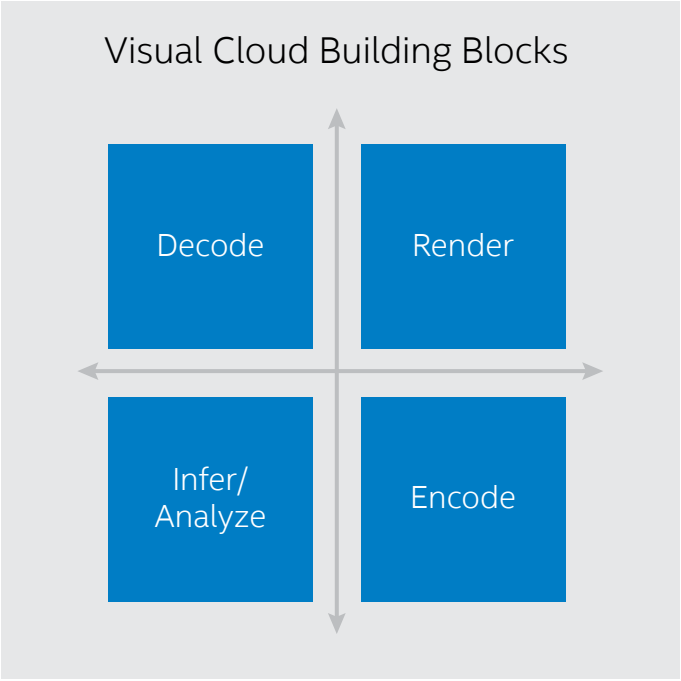

Figure 2. The Visual Cloud relies on four basic building blocks that can be combined in various ways to empower a wide range of Visual Cloud use cases.

- Capped variable bit rate (VBR) encoding enables the allocation of more bits to complex scenes, and fewer bits to simple ones, which results in constant video quality.
- Dynamic chunking of content, aligned to scene cuts (with minimum and maximum duration), reduces the occurrence of large I-frames (also called keyframes).
- In ATEME testing of public cloud-based live streaming, ATEME's Content Adaptive Encoding enabled a 12 percent reduction in bit rate while maintaining video quality. This lowered streaming costs by a similar amount . In another test, ATEME evaluated monthly VoD distribution costs, and found that Content Adaptive Encoding running on Intel architecture lowered the bit rate by up to 33 percent without decreasing video quality.
- By combining ATEME's Content Adaptive Encoding techniques with Intel's latest innovation in data center technology, CoSPs can achieve the following benefits:
- Low TCO (both capital and operational expense) through low bit rates as well as high-density racks that result in energy efficiency and small data center footprints
- Excellent viewing experience (resulting in customer preservation and revenue generation) on a wide number of devices by taking advantage of powerful CPUs to deliver high-quality video at low latency, either on-premises or in the hybrid cloud
- Future-ready, scalable solutions that allow in-field upgrades and fast adaptation to changing consumer demands
- Better Together: Intel® technology and Visual Cloud software
- Intel is focusing its Visual Cloud investments on system and software innovation to abstract software and hardware complexity to improve time to market for new Visual Cloud services. Here are just a few examples:
- To unleash innovation, simplify development, and accelerate time to market for Visual Cloud services, Intel is collaborating with the open source community on the Open Visual Cloud project (see sidebar on page 5). The Open Visual Cloud is an open source project that offers a set of pre-defined, optimized reference pipelines for various target visual cloud use cases.
- Through collaboration with Netflix*, Intel developed Scalable Video Technology for AV1 (SVT-AV1). This open source codec provides real-time AV1 encoding of 4K content using only CPUs.
  SVT-AV1 offers a production-grade AV1 encoder core that can be used to create a full transcoding solution.

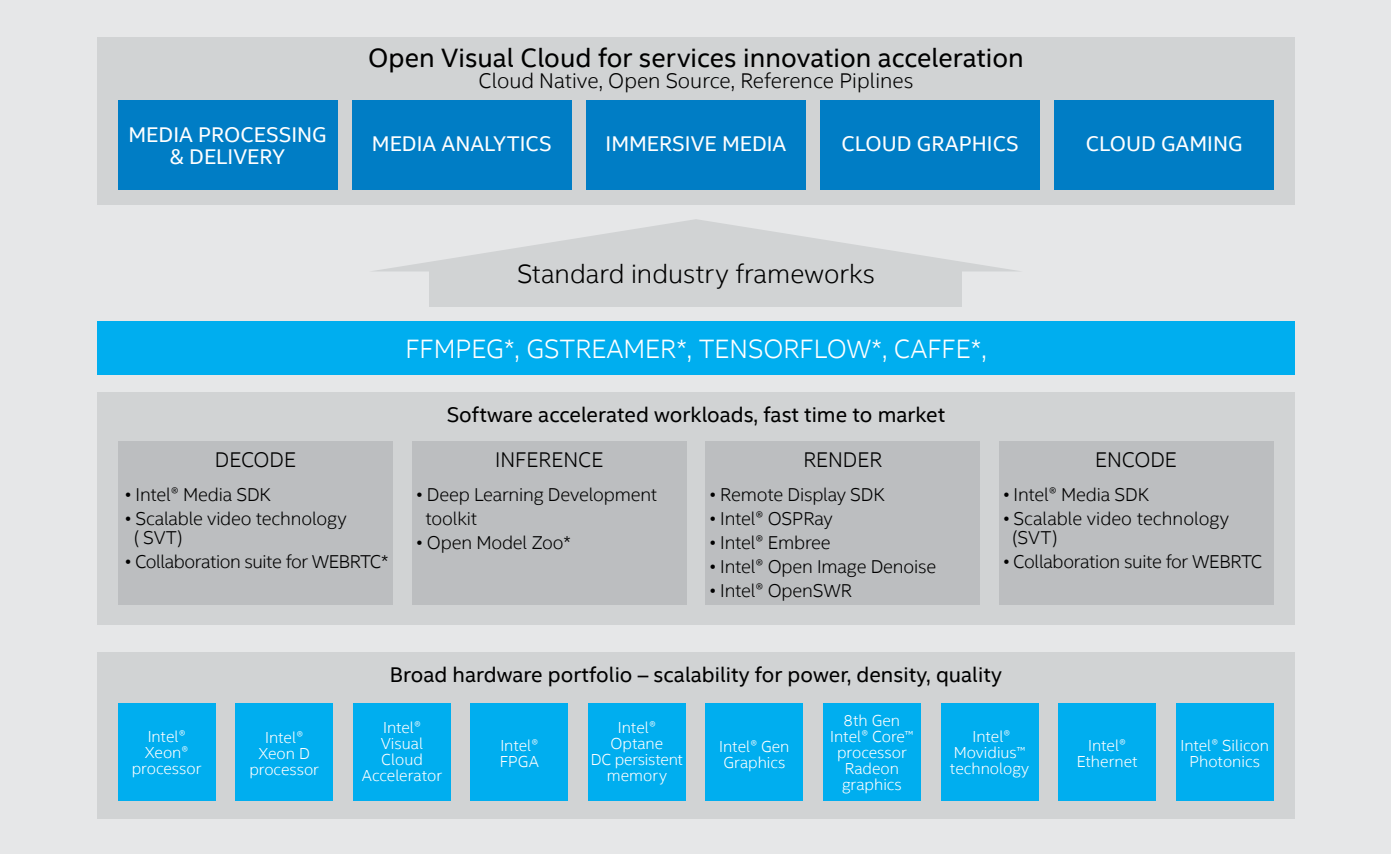

Figure 3. The combination of open source, optimized software and innovative hardware opens new performance vistas for the Visual Cloud, especially as 5G gains momentum.

SVT-AV1 uses parallelization at several stages of the encoding process, which allows it to adapt to the number of available cores, including the newest 2nd generation Intel® Xeon® Scalable processors with significant core count.

- Intel® Select Solutions for the Visual Cloud include a solution for simulation and visualization and a solution for Visual Cloud delivery networks. Intel Select Solutions eliminate guesswork with rigorously benchmark tested and verified solutions optimized for real-world performance. These solutions accelerate infrastructure deployment on 2nd generation Intel Xeon Scalable processors for today's critical workloads.
- By optimizing rendering engines for Intel architecture, Intel accelerates workloads such as cloud graphics and cloud gaming.
- By founding Visual Cloud services on software rather than building customized hardware-based solutions, CoSPs can maintain flexibility. If a service needs an accelerator or a different processor later, the application doesn't need to be rewritten to accommodate the new hardware.
- As shown in Figure 3, Intel's software-focused efforts are complemented by a broad hardware portfolio that provides the compute, memory, networking, and accelerator technologies that the Visual Cloud demands. For example, 2nd generation Intel Xeon Scalable processors have up to 56 cores per socket, and can scale from one to eight sockets—the perfect match for compute-hungry transcoding and rendering workloads. Intel® Optane™ DC persistent memory enables CoSPs to expand the system memory pool without having to invest in expensive DRAM, and allows in-memory software to deliver new levels of insight. The Intel® Visual Compute Accelerator 2 (Intel® VCA 2) and Intel FPGAs can speed rendering and transcoding. Intel® Deep Learning Boost (Intel® DL Boost) powers the artificial intelligence (AI) that is critical to delivering interactive personalized visual experiences and to monetizing video assets. Intel® SSDs can provide fast and affordable storage for content catalogs. Networking technologies from Intel, such as Intel® Ethernet and Intel® Silicon Photonics can help solve Visual Cloud content delivery challenges.

Conclusion

Media is rapidly transforming, becoming ever-more rich, interactive, intelligent and ubiquitous. The Visual Cloud consists of a set of capabilities for efficient processing and delivery of intelligent visual experiences globally. Because distribution accounts for the majority of cost in delivering visual experiences, innovation in new video encoders and encoding efficiency can significantly reduce CoSPs' distribution costs. Intel® technology, such as 2nd generation Intel Xeon Scalable processors, is enabling innovations in software-based video encoding to optimize for bit-rate efficiency while still maintaining the video quality consumers expect.

Find the solution that is right for your organization. Contact your Intel representative or visit **intel.com/networktransformation.**

Learn More

You may find the following resources helpful:

- ATEME* home page
- Intel's Visual Cloud home page
- IT Peer Network Visual Cloud blog posts
- 2nd generation Intel® Xeon® Scalable processors
- Intel® Optane™ DC persistent memory
- Free foundational 5G training from Intel® Network Academy
- Intel® Select Solutions
- Open Visual Cloud

The Open Visual Cloud

To help strengthen the ecosystem and provide ready access to the four Visual Cloud building blocks and pipelines for cost-effective Visual Cloud innovations, Intel is providing reference pipeline recipes for Visual Cloud services using existing open source functions from Intel in an open source project called the Open Visual Cloud. The Open Visual Cloud provides high-performance, high-quality, open source and validated building blocks—across encode, decode, inference, and rendering—as well as reference pipelines that support Visual Cloud workloads. The goal is to minimize barriers to innovation for quickly and easily creating and monetizing Visual Cloud services. Support for familiar industry-standard frameworks from the larger open source community include media (FFMPEG* and gstreamer*), artificial intelligence (TensorFlow*, Caffe*, MXNet*, ONNX*, and Kaldi*), and graphics (OpenGL* and DirectX*). The reference pipelines are provided as Dockerfiles* to simplify container image construction and deployment in cloud environments.

To find out more about the Open Visual Cloud, download code and unleash innovation, visit www.01.org/OpenVisualCloud.

Solution Provided By:

[1] Cisco Visual Networking Index: Forecast and Trends, 2017–2022

[2] ABI Research: Next-Generation Video and Content Delivery

[3] Global Cloud Gaming Market 2018-2022

[4] According to costing data provided to ATEME by Amazon Web Services* (AWS*) as of [date].
For live streaming the following equation was used:
Total Cost of Ownership (TCO) = Preparation Cost + $/GB x (Number of Simultaneous Viewers) x (GB per Viewer)
Without ATEME Content Adaptive Encoding: $761.84 = $14.84 + $747.07
bit rate = 4
With ATEME Content Adaptive Encoding: $667.42 = $13.74 + $653.69
bit rate = 3.5

[5] According to costing data provided to ATEME by Amazon Web Services* (AWS*) as of [date].
For video on demand (VoD), the following equation was used:
Total Cost of Ownership (TCO) = Preparation Cost + (Cost per GB per Viewer) + (Storage Cost per GB per Catalog Content)
Storage costs based on usual street price based on customer survey by ATEME.
Without ATEME Content Adaptive Encoding: $1,036,800 = $43,200 + $993,600
bit rate = 2.7
With ATEME Content Adaptive Encoding: $640,000 = $28,800 + $642,600
bit rate = 1.8

0719/FP/CAT/PDF xxxxx-001-001US