WHITE PAPER

Media and Communications
Content Creation and Distribution

# Rethinking Visual Cloud Services for Evolving Media

## Creating a New Model

With visual computing workloads growing at an accelerating pace, cloud service providers (CSPs), communications service providers (CoSPs), and enterprises are rethinking the physical and virtual distribution of compute resources to more effectively balance cost and deployment efficiency while achieving exceptional performance.

Visual cloud deployments accommodate a diverse range of streaming service workloads, encompassing media processing and delivery, media analytics, immersive media, cloud graphics, and cloud gaming.

Contending with the onslaught of new visual workloads will require more nimble, scalable, virtualized infrastructures; the capability of shifting workloads to the network edge when appropriate; and a collection of tools, software, and hardware components to support individual use cases fluidly. Advanced network technologies and cloud architectures are essential for agile distribution of visual cloud workloads.

A 2017 report, *Cisco Visual Networking Index: Forecast and Methodology, 2016–2021*, projected strong growth in all internet and managed IP video-related sectors. Compound annual growth rate (CAGR) figures during this time span, calculated in petabytes per month, included these predictions:

- Content delivery network (CDN) traffic: 44 percent increase globally
- Consumer-managed IP video traffic: 19,619 petabytes per month (14 percent increase) by 2021
- Consumer internet video: 27 percent increase for fixed, 55 percent increase for mobile

The impact of this media growth on cloud-based data centers will produce a burden on those CSPs, CoSPs, and enterprises that are not equipped to deal with large-scale media workloads dynamically. Solutions to this challenge include:

- **Increasing flexibility and optimizing processing**: Virtualization and software-defined infrastructure (SDI) make it easier to balance workloads on available resources. Open platforms and widely available acceleration technologies give CSPs and CoSPs a flexible way to optimize processing of diverse workloads. Learn more in the section titled *Balancing Dynamic Workloads in a Cloud-Based Data Center.*
- **Scaling compute, storage, and network resources**: Dynamic elasticity is a major advantage when contending with visual cloud workloads. Adopting a modern cloud infrastructure powered by a new generation of scalable processors increases resource availability substantially. Pervasive high-bandwidth and low-latency networks will also play a vital role for delivering rich user experiences and ensuring business model scalability. Learn more in the section titled *TCO Considerations*.

Table of Contents

- **Enhancing development processes**: High performance media, AI and graphics software components—along with sample reference pipelines that demonstrate how to construct new visual cloud services quickly and easily—support rapid application development and help reduce time-to-market for new revenue. Learn more in the section titled *The Open Visual Cloud*.
- **Deploying specialized solutions for select use cases and edge computing**: Specialized hardware—including discrete graphics processing units (GPUs), integrated GPUs, field programmable gate arrays (FPGAs), video processing units (VPUs), and image-processing units (IPUs)—can boost performance for select applications in which a single, targeted workload must be handled on a large-scale basis. A cohesive approach to solutions that are software only through hardware accelerated is an important consideration given the diversity of services in the visual cloud. Reference examples showing where to start provide significant flexibility across the solution continuum. Learn more in the section titled *The Open Visual Cloud*.
- **Implementing modern cloud architectures**: For performing operations on large, complex data files and delivering elastic processing power to efficiently handle ebbs and flows in intensive operations, modern cloud architectures hold the key to effective workload distribution.

## What is the Visual Cloud?

The visual cloud consists of a set of capabilities for remotely consuming content and services that center around efficient delivery of visual experiences from the cloud—both live and file-based. As shown in the figure below, the visual cloud supports five major services, each providing a set of related visual cloud use cases.

**Visual Cloud Services**

**All require high performance, high scalability, and full hardware virtualization**

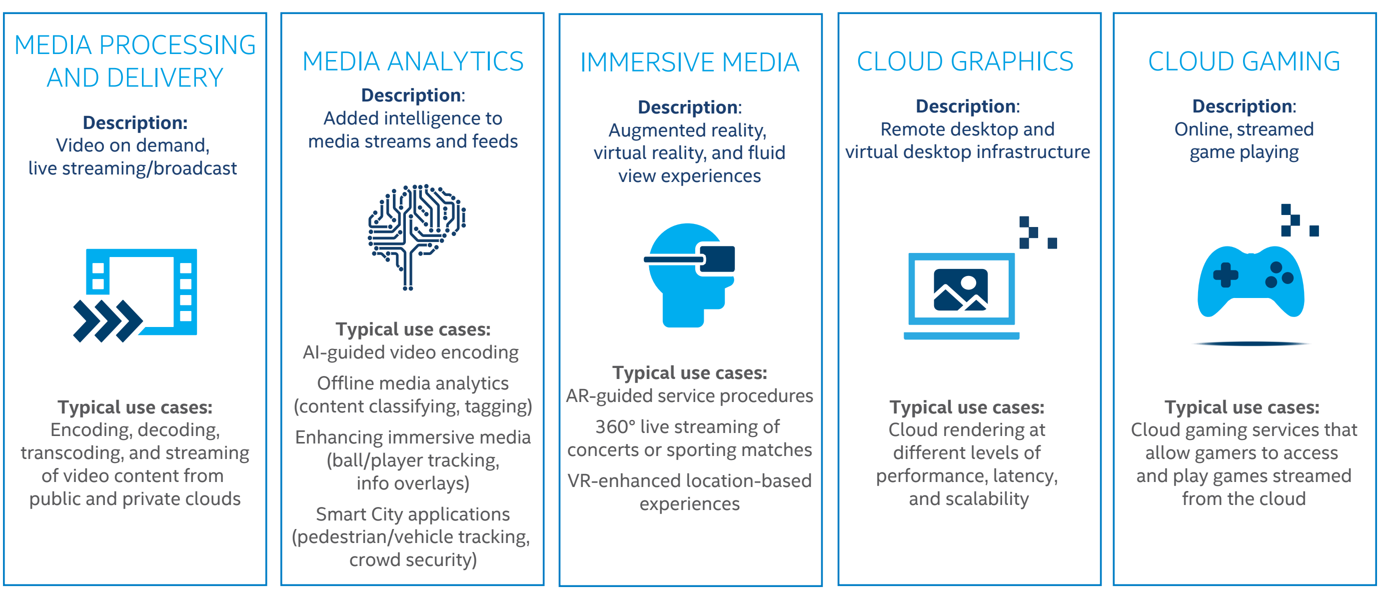


## Adapting to the Evolving Media Landscape

Media is rapidly evolving and has become a foundational building block for many new services. Visual cloud extends beyond traditional media, encompassing richer, more immersive use cases such as 360-degree virtual reality (VR) experiences. Visual cloud includes delivering more intelligence with video that is being distributed, incorporating analytics as a part of the media workflow. Analytics, performed on the contents of video sent to the end user, make it possible to deliver additional services or value. Media content as a part of Internet of Things (IoT) solutions is growing in importance. Sometimes visual workloads will need to be handled at the network edge, to reduce latency and provide near real-time responsiveness. Other times dense visual workloads will be best processed in the centralized data center using the most advanced processors and other scalable technologies.

Certain kinds of services—especially those subject to sudden surges in traffic—create challenges in building a supporting infrastructure that can adapt to quickly changing conditions. For example, NFL game views can skyrocket from hundreds to thousands to millions in a matter of several minutes—fueled by social media—and drop back down again just as quickly.

A virtualized, software-based infrastructure provides many benefits for service providers delivering content in many use cases growing in popularity, such as:

- User-generated live streaming (for example, Facebook Live*).
- Over-the-top linear streaming (Yahoo or Twitter* streaming of NFL games).
- Cloud-gaming applications (such as Shadow* or LiquidSky*), which can create substantial compute demands when large numbers of users are participating simultaneously.
- Rising video content generated through social media applications (for example, Facebook and Snapchat* typically handle eight billion average daily video views).

All of this is part of the visual cloud.

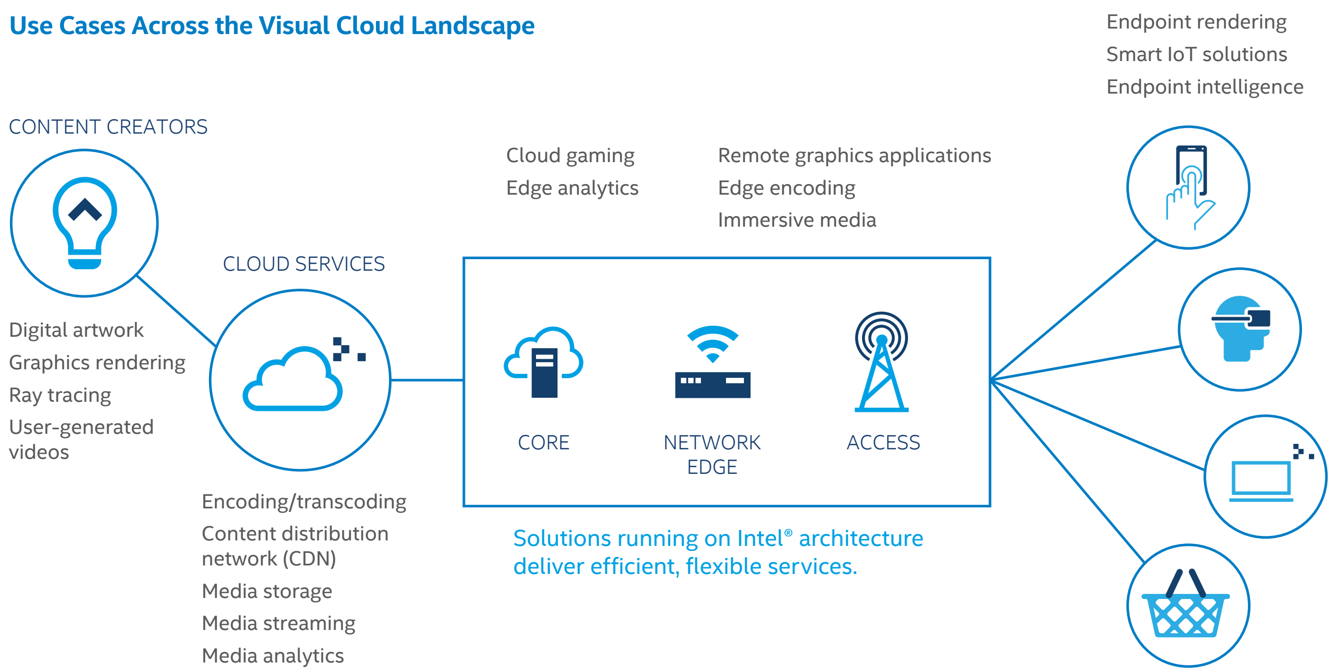


## Media Processing and Delivery

Cloud-based media processing and delivery demands in response to escalating video traffic across wired and wireless channels have become a vitally important service capability for CSPs and CoSPs. The predominant use cases involving media delivery and processing, encompassing both content creation and streaming, fit into two basic categories:

- **Broadcast media**—Communications media providers, such as Comcast and DirecTV (AT&T), produce and deliver both live and streaming video content to audiences—at least partially using their own CoSP network infrastructure.
- **Over-the-top (OTT) media**—OTT content is captured, produced, and delivered over the internet, either through mobile services or wired connections typically provided by a CoSP. This category includes video on demand content as well as live streaming video from companies such as Apple TV, Amazon, and Netflix, and web entities such as Facebook Live and YouTube*.

### Examples of Media Processing and Delivery Use Cases

Among the most noteworthy use cases involving media processing and delivery are:

- **Transcoding**—A large part of handling media workloads involves converting video formats, adjusting bit rates, and controlling compression intelligently, both in real time and offline. Intel® Advanced Vector Extensions 512 (Intel® AVX-512), part of the Intel® Xeon® Scalable processor platform, makes it possible to accelerate media workloads, supporting ultra-wide 512-bit vector operations, and handling 32 double-precision and 64 single-precision floating point operations per second, per clock cycle.
- **Distribution of workloads**—The nature of CPU-centric infrastructures makes it easy to perform cloud load balancing of media workloads. The additional flexibility of Intel®-based infrastructures employing virtualization technology, based on open standards and representing the clear majority of cloud service installations, gives CSPs and CoSPs more precise management control and agility in processing workloads dynamically.

## Media Analytics

Media analytics performed on live media or distributed video streams can help service providers, content aggregators, and content delivery networks better understand the nature of the visual content and derive useful intelligence from it.

> By 2022, MarketsandMarkets estimates that the video analytics market will reach USD 11.17 billion.[1] According to Tractica, deep learning revenues will increase from USD 655 million in 2016 to USD 35 billion by 2025.[2]

Sophisticated media analytics applications can do everything from detecting suspicious intruders in a video surveillance feed to surveying the traffic patterns in a smart city to better control flow. With billions of pieces of visual content exchanged daily, the market opportunities are expansive for creating new, useful services and adding features and capabilities to existing services.

### Media Analytics Use Cases

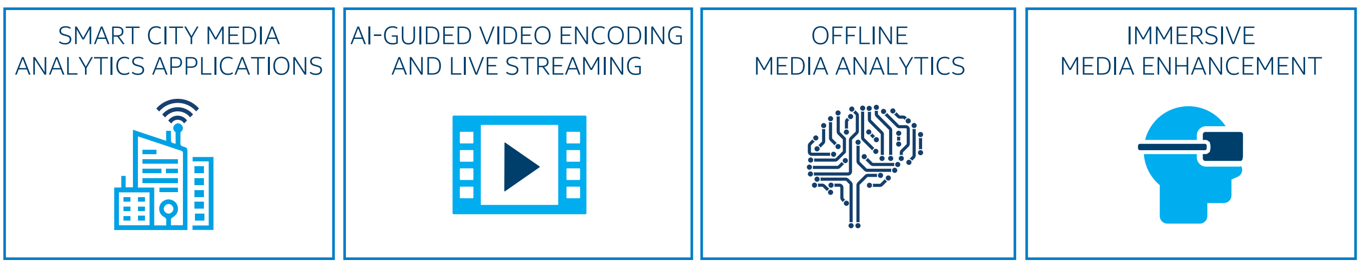


### Examples of Media Analytics Use Cases

- **Smart City**—Media analytics applications within the smart city range from identifying vehicles, including their license plates, for traffic monitoring and toll collection to performing facial recognition of individuals near buildings or public spaces, for both security reasons and for crowd analysis. Often, visual recognition tasks need to be performed in real time, working most efficiently in edge computing implementations.
- **AI-guided video encoding and live streaming**—By adding AI capabilities to video encoding algorithms, broadcasters and OTT producers can control quality and bandwidth parameters to meet specified goals. An AI algorithm at the heart of this process can be trained for optimal quality or bandwidth from a collection of video streams, and then automatically adjust the encoder to ensure goals are achieved.
- **Real-time and offline media analytics**—By using object and performer detection and classification, AI applications can provide a wide variety of services to help CSPs and CoSPs create new revenue channels. For example, AI techniques can identify and tag potentially violent movies, movies that include certain performers, or provide metadata to recommendation services based on movie content. These kinds of services can be applied at different points in the delivery system, in real time or offline, to add intelligence to the packaging and distribution of content.
- **Immersive media enhancement**—Applying analytics to immersive media applications opens opportunities in both 360-degree content display and AR (augmented reality) applications. For example, sports presentations can benefit from tracking players or the ball, using AI to control the encoding for providing a selective view of the action. Interactive applications include being able to automatically overlay video content with relevant information (such as a player's batting average or basketball foul shot percentage), or an international football team's ball possession or passing accuracy statistics.

## Immersive Media

Innovation in AR and VR solutions is changing the way that human beings interact with the world. These technologies are also inspiring innovators to introduce new products, services, and business models into the visual computing market, creating opportunities for CSPs, CoSPs, and broadcast companies.

> Potential revenues, as projected by International Data Corporation, show strong growth opportunities, with USD 45 billion in sales of headset hardware forecast by 2021. Software sales in this market also look to be substantial; expected to exceed USD 35 billion by 2025, according to Statistica.[3]

Advanced immersive media applications are just beginning to gain traction in the industry with the more practical uses leaning toward AR, such as guiding worker tasks through smart glasses. Consumer-driven interests and entertainment tend to be aligned toward VR. On the broadcast side, 3D renderings of stadiums and other types of facilities can be integrated into media presentations, whether live or on demand. Streaming of live concert experiences or sports using 360-degree video content is another application that is already being implemented by companies, including Tiledmedia and Netzyn. Recent advances in moving media workloads to the network edge are providing for more quality experiences by addressing the bandwidth concerns and latency issues that have been a challenge for VR scenarios.

### Examples of Immersive Media Use Cases

The following immersive media use cases are among the more promising in this sector:

- **360-degree live streaming**: Companies are now able to live-stream events, such as concerts or sporting matches, to a variety of consumption platforms with high visual fidelity.
- **AR-guided service procedures**: AR offers a strong potential for guiding factory workers, warehouse employees, and technicians in daily procedures, equipment operation, and maintenance tasks, using instructions overlaid on AR head-mounted displays or smart glasses. Guidance can be delivered by video, audio, or as text lists of steps in a procedure. More sophisticated applications can provide immediate feedback by observing and commenting on worker actions.
- **VR-enhanced location-based experiences**: On-premises immersive experiences are being used broadly in a number of consumer-focused ways. The number of VR arcades has been growing globally, as have the number of businesses augmenting tourist destination experiences.

To learn about optimizing content delivery networks, read this case study: Orange* Builds and Tests Virtual Content Delivery Network.

## Cloud Graphics

By being able to tap into powerful cloud-based servers to perform graphics-intensive operations, organizations can selectively choose the level of performance, latency, and scalability that best fits their business needs. Using new technologies, such as cloud-hosted workstations powered by the Intel® Xeon® processor E3-1500 processor v5 with Iris® Pro graphics, enterprises can equip technical staff members to work collaboratively and access centralized graphics applications remotely from wherever they want to work. With data and applications accessible securely from the data center, complex renderings and visualizations can be handled by high-performance servers in a virtualized environment in which workload processing can be adjusted dynamically to meet requirements. Remote rendering of 3D animations or 3D models represents another promising use of cloud graphics.

**Cloud Graphics Use Cases**

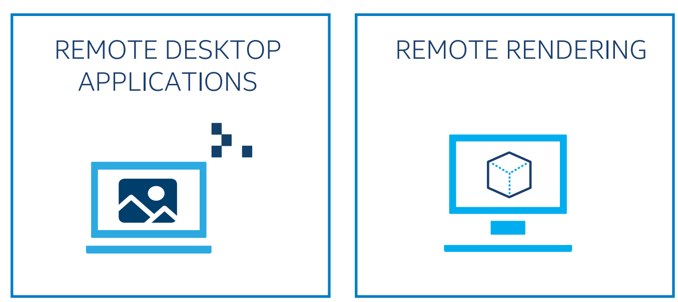


### Examples of Cloud Graphics Use Cases

- **Remote desktop applications**—Desktop and workstation remoting solutions, powered by the Intel Xeon E3-1500 processor v5 family with Iris Pro graphics, deliver professional graphics capabilities to geographically diverse technical teams, adding to productivity, improving collaboration, and minimizing the risk of lost or stolen intellectual property.
- **Remote rendering**—The intensive compute cycles required to process complex animations, scientific data visualizations, ray traced content creation, and 3D models lend themselves well to remote rendering in the cloud where virtualization can be used to dynamically provision compute, storage, and network resources to handle demanding media creations. The memory footprint required to render massively complex data sets often exceeds that of a GPU, making the Intel® Rendering Framework's CPU-based ray tracing and rasterization the choice of many in Hollywood and the scientific visualization community.

## Cloud Gaming

In much the same way as movie distribution evolved from DVDs to online streaming as internet connectivity speeds increased, online game services have begun to take advantage of faster connection rates and improved compression techniques. Video streaming lets users stream games in real time from remote servers without requiring that the game be downloaded first. Some benefits of cloud gaming are instant play, no downloads required, and playing games on multiple screens or devices.

According to Knowledge Sourcing Intelligence, the cloud gaming market is projected to grow at a CAGR of 26.12 percent between 2018 and 2023, reaching a total market size of USD 4.284B by 2023 (up from USD 1.064B in 2017).[4]

### Examples of Cloud Gaming Benefits

- **Online gaming services revenues**—The promising growth of online games, served from the cloud, represents additional revenue for CSPs and a complement to current use models for gaming. Media technology can be used to stream the game content to the end user, applying media encoding or the media pipeline processing part of the stack to deliver the game content efficiently. Visual cloud software stack components can be used for both creation and delivery of the media.
- **Any time, any place access to games**—a unique characteristic of online gaming is that the delivery model from the cloud makes it possible for gamers to access and play games from their choice of devices—from smartphones to tablets to desktop machines—wherever they happen to be. This is sometimes called the any pane of glass approach, signifying that because the compute operations are mainly taking place in the cloud, users can enjoy good performance on any device across the board.

## Balancing Dynamic Workloads in a Cloud-Based Data Center

Significant increases in complexity for visual cloud processing and delivery workloads in the data center call for rethinking the approaches and technologies for handling current and future processing requirements.

The trends and projections are bullish as new media use cases are developed and deployed. The Cisco Visual Networking Index* (VNI) report for 2017 forecasts that, by 2021, Ultra HD will represent 20.7 percent of IP video traffic. The same report projects a seven times increase in internet video surveillance traffic between 2016 and 2021. According to the World Economic Forum, immersive media, including AR and VR, will grow to a USD 95 billion market by 2025, spurred by live events, gaming, and video entertainment.

The capabilities of virtualization allow diverse workloads to access and share a common collection of compute, storage, and network resources. In environments that are demand-driven and that frequently experience processor-intensive workloads, virtualized hardware can scale freely and relocate workloads as required across infrastructures. This has spawned the massive increases of efficiency that customers have found in the cloud.

Cloud-based data centers built using modern cloud infrastructures can capitalize on network function virtualization (NFV), software-defined networking (SDN) and SDI to support functions historically delivered by dedicated network appliances. These new technologies adapt well to the demands of shifting media workloads. As visual computing use cases evolve and mature, enterprises, CSPs, and CoSPs—relying on Intel® architecture-based infrastructures—will have a wide range of choices as to where and how to process workloads. This may involve using the network edge, the cloud, or the core data center, depending on the scenario, data sizes involved, security issues, and workload distribution. The broad choice of tools and technologies available from Intel to efficiently handle the workloads—from pure software solutions to those based on hardware acceleration—also give solution architects the freedom to design, develop, and deploy high-performance visual solutions in the most cost-efficient and scalable way.

## TCO Considerations

Throughout the ecosystem, running visual compute workloads on pure software architectures gives customers the flexibility of making improvements to media processing and delivery as well as customizations to the visual quality. To maximize the business value of infrastructure investments, CoSPs, CSPs, and enterprises need servers flexible enough to accommodate a diverse range of both compute and visual workloads. Several key factors affect the total cost of ownership (TCO) of a media-processing system. Factors to consider include:

- Scalability of server infrastructure to address all workloads for maximizing IT investments.
- Optimized software features that enhance the value of data center and cloud resources.
- Hardware platform characteristics with respect to cost, cooling, overall utility, deployment complexity, and maintenance.

- GPU license fees by other vendors have an annual cost of licensing.
- Server density considerations and overall resource utilization in the data center, factoring in virtualization capabilities and workload balancing tools.
- The degree of flexibility and agility for managing the balance between video quality and the streaming bit rate, as well as availability of customizations and ongoing performance enhancement.
- Tradeoffs between investing in an open standards hardware platform with widely available open source software components and a proprietary hardware platform subject to vendor lock-in and a greatly reduced range of supported applications.
- Depth and breadth of the ecosystem supporting the hardware/software platform and whether it is adequate to provide reasonable choices and fulfill business objectives.

There is tremendous potential to take advantage of existing infrastructures, and the majority of the world's leading data centers are populated with Intel® Xeon® processor family servers. Infrastructure managers can use low-demand use periods to handle media-processing applications that have been optimized and tuned for these servers. This—in addition to leveraging open source and off-the-shelf packaged open source-based projects—helps minimize the need to acquire additional hardware.

### Meeting Service-Level Expectations

Optimizing the performance of the infrastructure to meet service-level expectations, a key consideration for CSPs, can be accomplished by increasing the workload density within the data center. This, in turn, can help lower the TCO for the CSP. By upgrading to Intel Xeon Scalable processors, a CSP can substantially increase the density of virtual machines in the data center, letting them support more customers per rack.[5]

Intel Xeon Scalable processors deliver exceptional performance and provide fluid scalability for dense media workloads. To improve media operations, Intel Advanced Vector Extensions 512 (Intel AVX-512) feature double the number of floating point operations per second (FLOPS) per clock cycle, a significant leap forward compared to the past-generation capability. Demanding tasks that are vital to a visual cloud environment gain performance boosts, including media analytics, video data encoding and decoding, digital content creation, 3D modeling and simulation, and visualization. Additional microarchitecture and instruction sets to drastically enhance visual cloud capabilities like VNNI (Vector Neural Network instructions), that improve throughput of multiply-add operations with int8 and int16 data types to achieve performance gains in low precision convolution, and matrix-matrix multiplication operations used in deep neural networks.

For cost-effective processing, products such as the Intel Xeon D processor series benefit from software optimizations that bring new performance capabilities for handling media streams. For example, the Intel Xeon D-2191 processor now supports multiresolution transcoding to an adaptive bitrate (ABR) profile of a 4Kp60 stream in a single socket, thanks to optimized Scalable Video Technology for High Efficiency Video Coding (HEVC) (SVT-HEVC) software. Wide selections available in the Intel Xeon processor family can help resolve the challenge of balancing performance and cost for any given deployment.

### A Software-Based Approach to Visual Cloud Processing

A software-based approach to visual cloud processing offers flexibility and the opportunity for new revenue channels. Without having to rework the physical architecture or update the hardware, a software-based model—running code on a CPU—provides these advantages:

- Supporting new video codecs that emerge, through simple software updates, rather than expensive hardware changes.
- Adding intelligence to video segments through analytics software that can be rapidly iterated on with new algorithms.
- Adding efficiency to encoding operations through artificial intelligence, enabled by software.
- Advancing video compression to new levels through software enhancements.
- Balancing video quality with bandwidth, mediated by software.

Ninety-five percent of the server processors powering the cloud are based on Intel® technology.[6] The vast supporting ecosystem and deep expertise surrounding open platforms based on Intel architecture ensures a reliable, long-term foundation for launching media applications and services. It also provides a framework in which software updates are delivered frequently, through open source contributions as well as media toolkits that are regularly updated, adding to the long-term business value.

## The Open Visual Cloud

### Software Building Blocks

CSPs, CoSPs, and enterprises increasingly rely on open source solutions and components to drive innovation, reduce time-to-market, lower costs, and minimize obstacles when commercializing products. Optimized open source building blocks are needed to accelerate services innovation for the visual cloud. There are essentially four core building blocks for most visual cloud workloads: decode, inference, render, and encode. Services are defined by the order in which these core building blocks are placed in a pipeline. For example, a simple transcode service is realized with decode + encode core building blocks. Insertion of an inference building block (decode + inference + encode) would result in a media analytics service relevant for digital security and surveillance or user-generated content ad-insertion use cases where intelligent content analysis is required.

**Pipelining Visual Cloud Building Blocks**

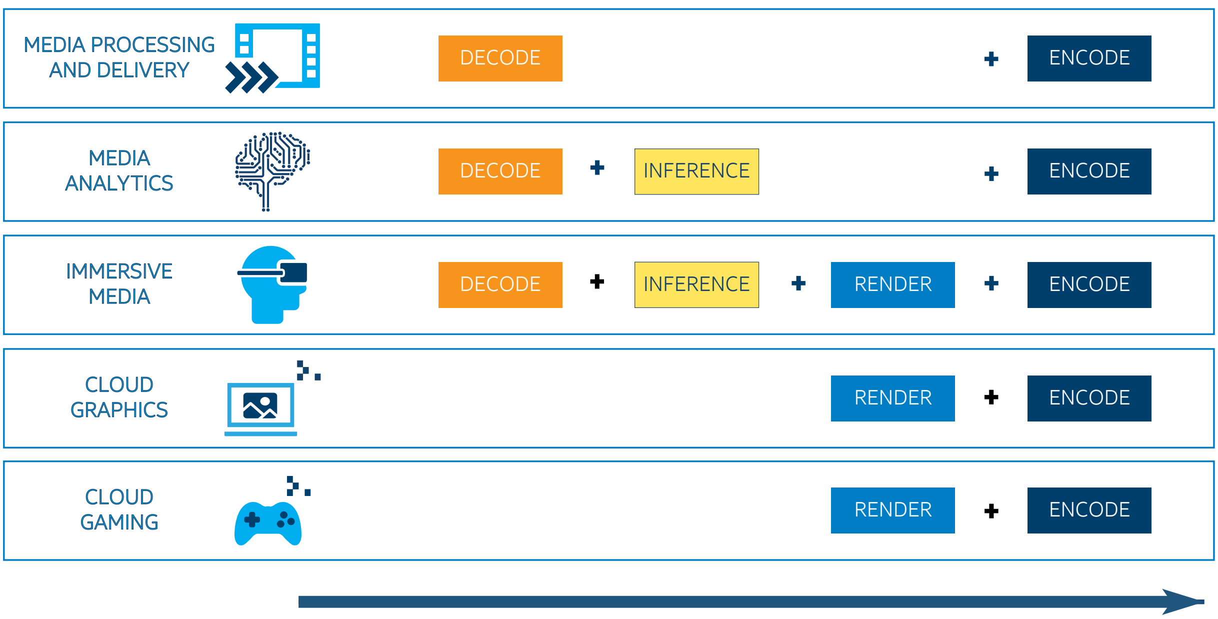


To help strengthen the ecosystem and provide ready access to the core building blocks and pipelines for cost-effective visual cloud innovations, Intel is providing reference pipeline recipes for visual cloud services using existing open source functions from Intel in an open source project called the Open Visual Cloud. The Open Visual Cloud is an open source project consisting of highly optimized cloud native media, AI, and graphics components and sample reference pipelines to easily construct visual cloud services. The focus is on accelerating the availability of high performance, high quality, open source, validated building blocks—across encode, decode, inference, and rendering core building blocks—that support visual cloud service workloads. The goal is to minimize barriers to services innovation and provide interoperable building blocks and sample reference pipelines for quickly and easily creating, storing, distributing, and monetizing visual cloud solutions.

The Open Visual Cloud provides support for familiar industry standard frameworks in media (FFMPEG and gstreamer), AI (TensorFlow*, Caffe*, Apache MXNet*, Open Neural Network Exchange (ONNX*), Kaldi*) and graphics (OpenGL*, DirectX*) to encompass the larger open source community. The entire ecosystem will benefit from Intel's continued upstreamed performance optimizations to these important frameworks as well as the interoperability of building blocks to pipeline services.

### Open Visual Cloud: Media Building Blocks

Intel is contributing to each core building block with new projects and enhanced performance. With the observation that encode is a required building block across all the visual cloud services, Intel has released several Scalable Video Technology (SVT) core encoders to support the ecosystems needs. First released in September 2018, SVT-HEVC is an H.265-compliant encoder core with significant speed gains at equivalent video quality levels over other open source encoders, and has been highly optimized for Intel® Xeon® Scalable processors. Additional experimental versions of SVT-AV1 (an Alliance for Open Media Video 1 compliant encoder core) and SVT-VP9 were released in January and February 2019 respectively. FFMPEG and gstreamer plugins for all SVT encoder cores are provided to simplify developer integration, evaluation, and deployment.

SVT is a CODEC-agnostic architecture, highly optimized for Intel Xeon Scalable processors taking advantage of Intel® Advanced Vector Extensions 512 (Intel® AVX-512) and providing outstanding scalability in the cloud utilizing as many CPU cores as allocated. Key features include 3 levels of parallelism (process, segment, and picture), a human visual system optimization classifier, and 13 presets of speed/quality tradeoffs ranging from video on demand (VOD) to live linear broadcast. All SVT-based encoder cores are provided under the highly permissive Open Source Initiative* (OSI) approved BSD + patent license.

- **SVT-HEVC**: An extremely scalable encoder core with great visual quality/density tradeoffs and released to open source, this advanced, ultra-high-definition encoder delivers performance speedups of two orders of magnitude compared to X265 at equivalent video quality levels.
- **SVT-AV1**: This encoder core is an alternative implementation of the open source AV1 standard, providing interoperability with the latest Intel processors, and based on the same architecture as SVT-HEVC, with similar scalability and attention to visual quality/density tradeoffs. The royalty-free AV1 codec delivers 4K ultra-high-definition video with provisions for supporting even higher definition video content as new formats emerge. An SVT-AV1 decoder will also be part of this project, contributed by SVT ecosystem partners, allowing for shared tools between the encoder and decoder, and to provide comprehensive AV1 transcode support.
- **SVT-VP9**: The latest addition to the SVT portfolio is another open and royalty-free CODEC standard, VP9. In contrast to HEVC, VP9 support is common among web browsers.

## Open Visual Cloud: AI

In addition to the encode and decode building blocks, optimized inference is also needed to support emerging media analytics services. This includes personalized, localized content delivery tuned to user preferences, along with inference intelligence such as image/text/speech detection and recognitions, along with many more usages. Whether the provider is trying to recommend content based on viewing habits or inserting better targeted ads into live over-the-top broadcasts, it requires intelligence and analytics. This drives the need for inference as another core building block.

The Open Visual Cloud building block for inference is the OpenVINO™ toolkit with the Deep Learning Deployment Toolkit for Intel CPU and GPU (Intel® Processor Graphics) and heterogeneous plugins and Open Model Zoo pre-trained models and samples. Those open source OpenVINO toolkit components provide deep learning and traditional computer vision analytics capabilities and enhance the development of media analytics applications, using a write once, deploy everywhere model aligned with Intel architecture-based platforms, including pure software and several hardware acceleration options. These capabilities include a model optimizer that supports industry standard deep learning frameworks, so you can work with models you may have already developed. It supports pre-trained models from the Open Model Zoo along with more than 100 open source and public models in popular formats such as Caffe, TensorFlow, MXNet and ONNX. The Intel® Distribution of OpenVINO™ toolkit is also available with additional, proprietary plugin support for Intel® FPGAs, Intel® Movidius™ Neural Compute Stick (NCS), and Intel® Gaussian Mixture Model - Neural Network Accelerator (Intel® GMM-GNA), and provides optimized traditional computer vision libraries (OpenCV*, OpenVX*).

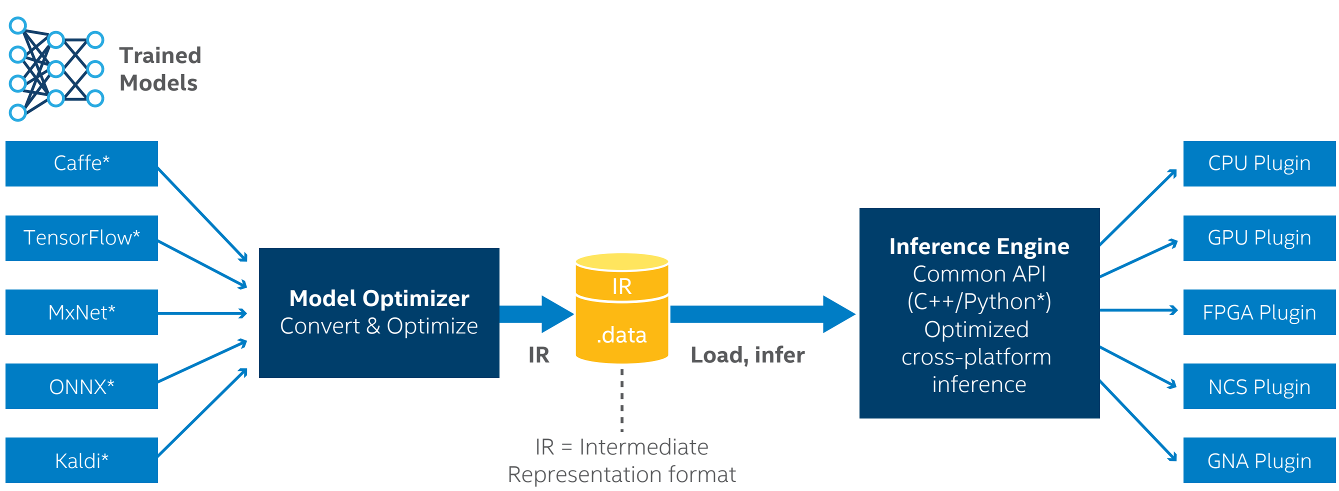

## Open Visual Cloud: Graphics

The Open Visual Cloud building block for rendering is the Intel Rendering Framework. This open source framework consists of four ingredients to deliver optimized CPU rendered raytracing and rasterization. It has been used extensively by professional animation and movie studios for photo-realistic scene creations, as well as by researchers for scientific data visualization due to the massively large complex data sets involved.

### Intel® Embree

- CPU optimized ray tracing algorithms that take full advantage of all available CPU cores for interactive rendering to address large models using full system memory
- Tool kit for building ray tracing apps
- Broadly adopted by third-party ISVs into more than 80 commercial applications. Autodesk* currently uses Intel Embree as a rendering engine in several of their top graphics applications.
- For more information, visit http://embree.github.io

### Intel® OSPRay

- Scalable rendering engine based on Intel Embree
- API designed to ease creation of visualization software
- For more information, visit http://ospray.org

### OpenSWR

- High-performance CPU rasterization
- Fully integrated into MESA* v12.0+
- Supports ParaView*, Visit*, VTK*, EnSight*, VL3
- For more information, visit http://mesa3d.org ; www.openswr.org

### Intel® Open ImageDenoise

- High-performance denoising library for ray tracing
- Efficient use of deep learning filter trained to handle a wide range of samples per pixel
- Suitable for both preview and final-frame rendering
- For more information, visit https://openimagedenoise.github.io

Innovate Quickly with Open Visual Cloud Pipelines

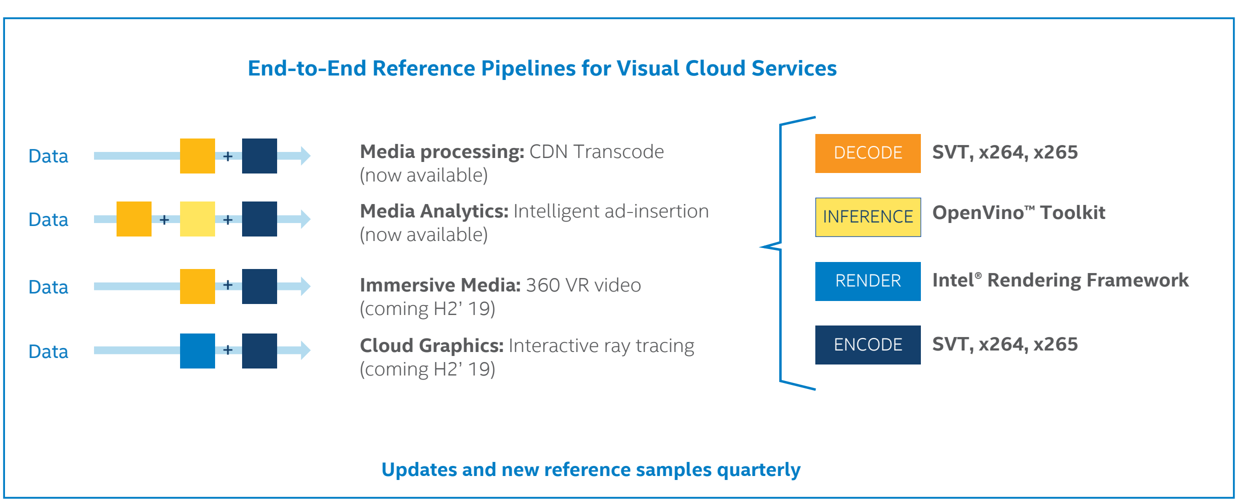


Open architecture, Intel Xeon processor optimized for cloud infrastructures, industry standard framework hardware and software plugins

### Open Visual Cloud: Reference Pipelines

The Open Visual Cloud goes beyond simply providing optimized software building blocks and additionally offers full, end-to-end sample reference pipelines to demonstrate how to use the building blocks most efficiently and to serve as a starting point for enhanced or new services. Highly flexible and optimized CPU pipelines will be delivered first, followed by hardware-accelerated versions. The reference pipelines are provided as Docker* containers, reducing setup, pipeline construction and configuration from what may take days to minutes. Developers have the ability to experience and evaluate a live visual cloud service enabled by a reference pipeline on public or private clouds.

**Media processing and delivery**

- Reference Pipeline for CDN streamed 1:N transcoding service
  - Apache traffic server + NGINX + SVT (FFMPEG and gstreamer plugins)
  - Full software-optimized pipeline—get started without any special hardware
  - Hardware-accelerated transcoding (FFMPEG and gstreamer plugins) for use with Intel® Visual Cloud Accelerator 2 add-in card
    - Increase stream density when needed without changing the software architecture

**Media analytics**

- Reference Pipeline for live over-the-top ad insertion service
  - OpenVINO toolkit + SVT
  - Full software-optimized pipeline—get started without any special hardware
  - Hardware-accelerated inference (OpenVINO toolkit plugins) for use with Intel® inference accelerators (FPGA, GPU, VPU)
  - Increase inference performance when needed without changing the software architecture

Additional sample reference pipelines will be provided on a regular basis, as well as additional platform support via the plugin architecture.

With Open Visual Cloud software running on an x86-based infrastructure, the development community has a platform from which to create comprehensive end-to-end visual cloud services, with powerfully optimized building blocks, sample reference pipelines to simplify adoption and deployment, and the flexibility, enhancements and support of the open source community.

## Get Ready for Visual Cloud Opportunities

Evolving media applications are transforming entire industries, creating new cloud-based services, including media processing and delivery, media analytics, immersive media, cloud graphics, and cloud gaming. All of these services are part of an next-generation infrastructure model known as the visual cloud. As this transformation accelerates, CSPs, CoSPs, and enterprises have unique opportunities to attract new business customers and meet changing user expectations, as well as make maximum use of available compute, storage, and network resources across a cohesive platform architecture.

The building blocks and components that support the visual cloud include a robust hardware portfolio from Intel for meeting diverse power and performance needs, optimized open source software to streamline development of solutions, and enabling ingredients provided through the Open Visual Cloud to improve design results and reduce time-to-market.

Intel offers comprehensive, interoperable visual cloud solutions powered by the Intel Xeon Scalable processor family, giving service providers a way to focus technology investments on rapid innovation and agile deployment of services built for the visual cloud. This platform is richly supported by a vibrant ecosystem and contributions from an active open source community. To keep pace with rapidly growing visual cloud opportunities, Intel is collaborating with industry leaders to enable transformative solutions through the latest cloud architectures, advanced networks, and flexible, scalable compute platforms.

For more information visit **http://www.intel.com/visualcloud** and **https://01.org/OpenVisualCloud**.

[1] https://www.marketsandmarkets.com/PressReleases/iva.asp
[2] Tractica, 2Q, 2017
[3] https://www.statista.com/chart/4602/virtual-and-augmented-reality-software-revenue/
[4] *Cloud Gaming Market Forecasts from 2018 to 2023*. Knowledge Sourcing Intelligence, 2016
[5] *Workload Optimization: A Guide for Cloud Service Providers.* Intel, 2018. https://plan.seek.intel.com/WorkloadOpteGuide_REG
[6] http://www.itprotoday.com/cloud-data-center/intel-owns-95-percent-cloud-and-you-didnt-know-it-until-now