IDF15

Media Cloud Based on Intel® Graphics Virtualization Technology (Intel® GVT-g) and OpenStack*

Xiao Zheng – Software Engineer, Intel Corporation

SFTS002

Make the Future with China!

intel

Agenda

- Media Cloud
- Media Cloud Infrastructure
- Case Study: Virtualized Media Server
- Optimize for Virtualized Media Server

Media Cloud

Media Processing Opportunity

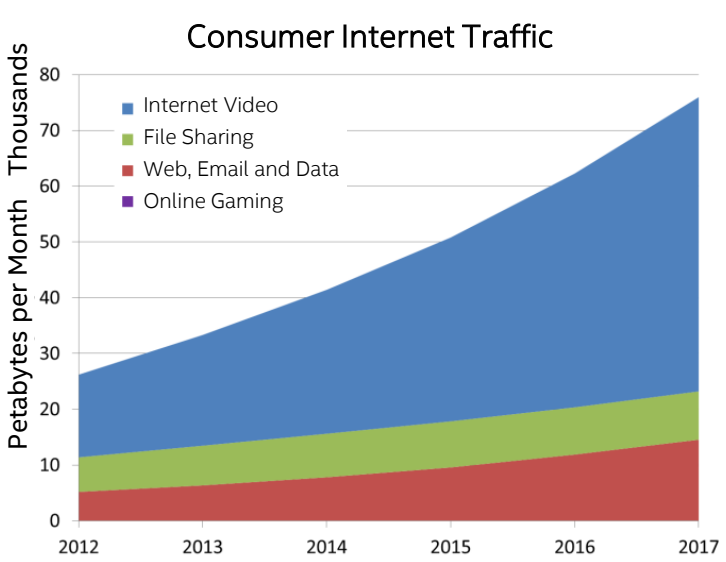

Source: Cisco* Systems Inc., Visual Network Index (VNI), 2013, 2015

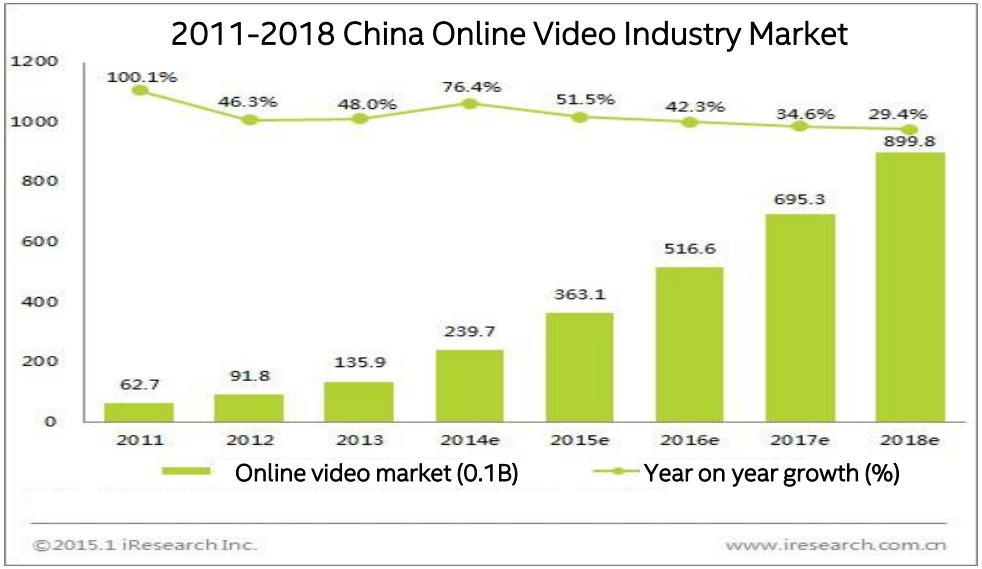

Source: 2015.1 iResearch Inc., www.iresearch.com.cn

Internet video traffic is forecasted to grow at 29% CAGR and will represent 69% of consumer traffic by 2017

Industry Landscape

Video Delivery

Store/Stream, Transcode
Offline, real-time,
OTT, VOD, IPTV

Visual Understanding

Search, Surveillance

Media Cloud

Video Delivery

Store/Stream, Transcode
Offline, real-time

Visual Understanding

Search, Surveillance

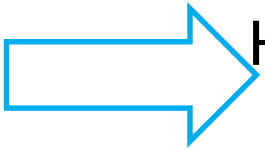

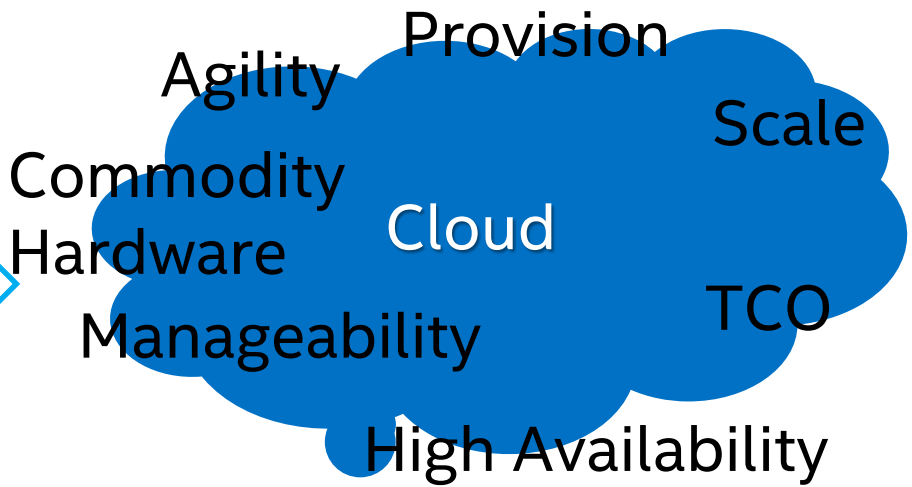

Media Delivery Example

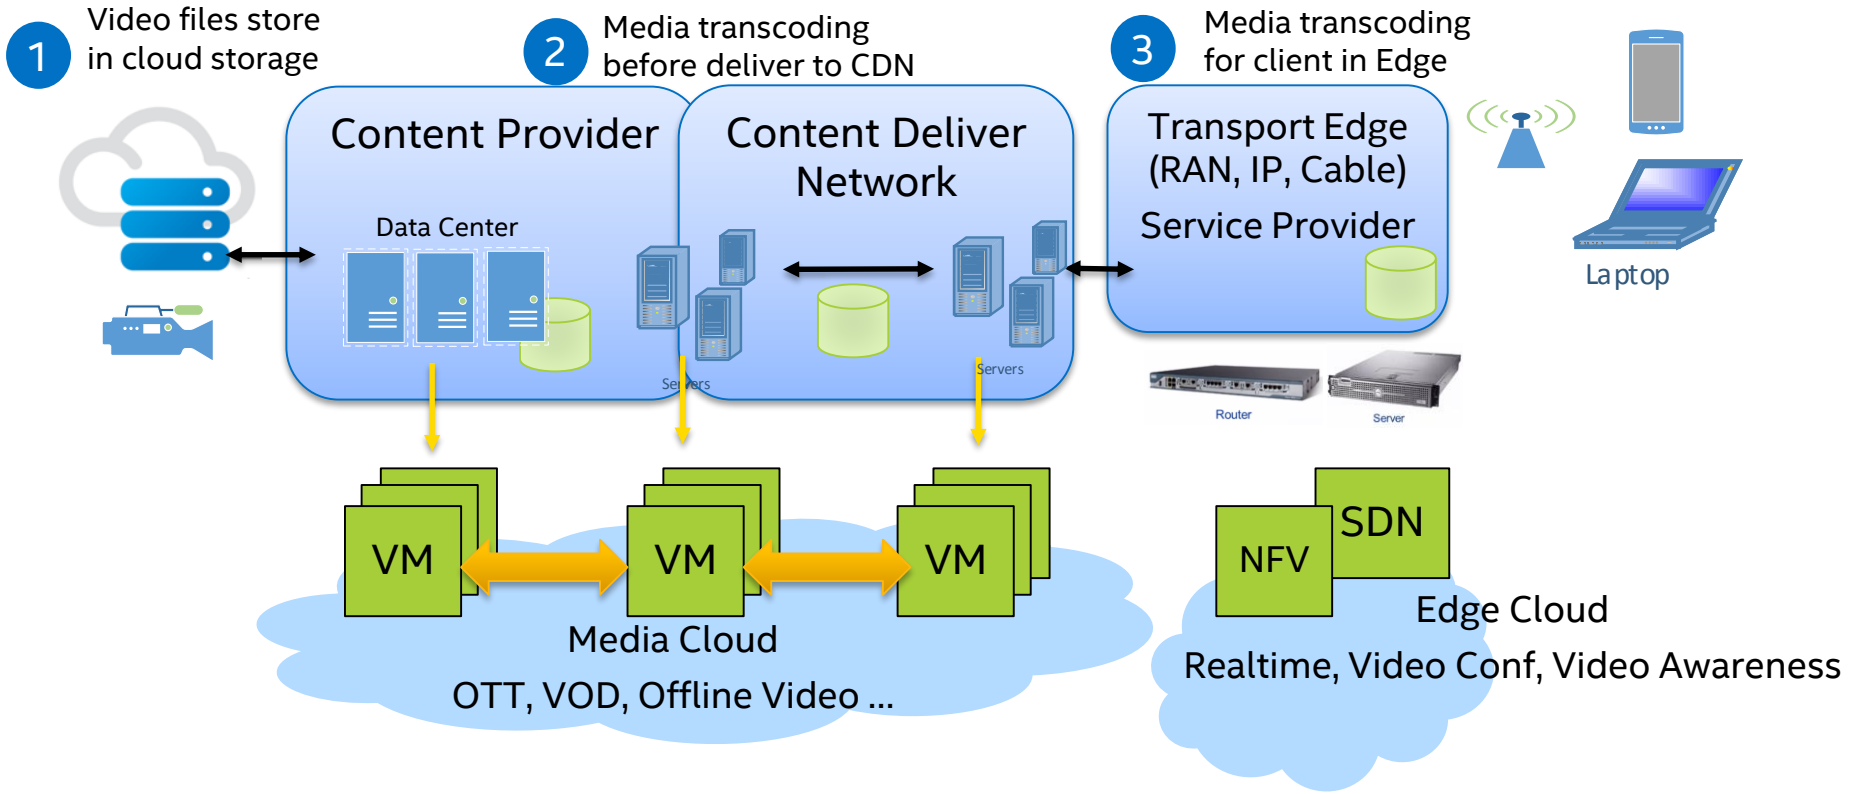

Technology Gaps

No GPU Virtualization

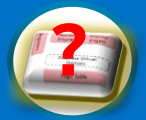

Low cost CPU transcoding throughput is much lower without GPU acceleration

DSP solution difficult to integrate in Cloud

No Cloud Orchestration

No GPU instance awareness
No vGPU capability scheduling
No vGPU resource monitoring

Media Cloud based on Intel® Graphic Virtualization Solution

Media Cloud Infrastructure

IDF15

Building Blocks

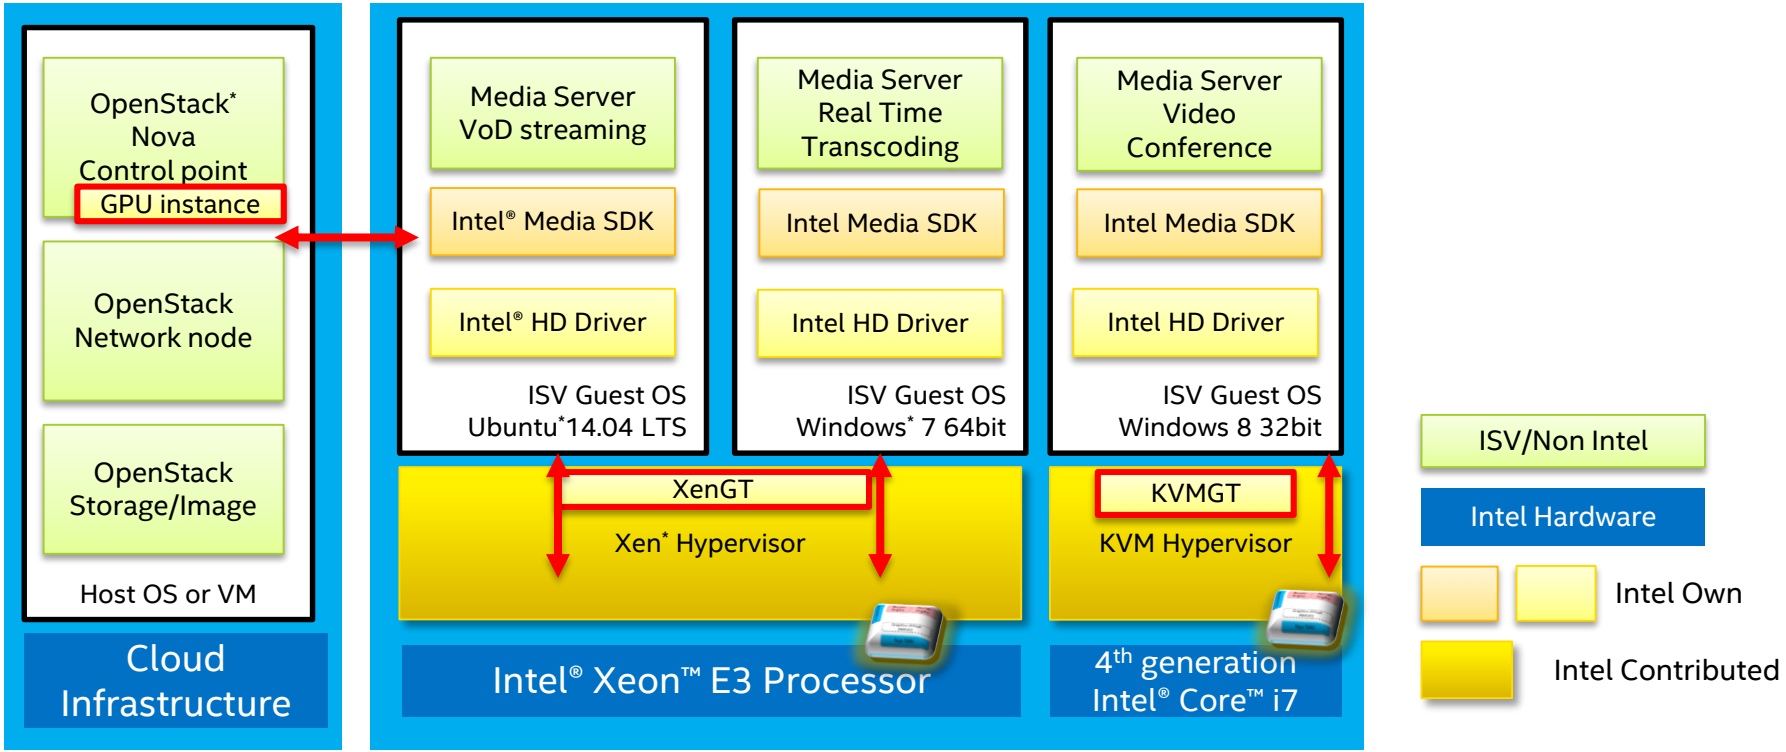

Requirements of GPU Virtualization

Performance → Direct GPU acceleration

Feature → Consistent visual experience

Sharing → Multiple virtual machines

GPU Virtualization Approaches

API Forwarding (Intel® GVT-s)	Direct Pass-through (Intel GVT-d)	Full GPU Virtualized (Intel GVT-g)
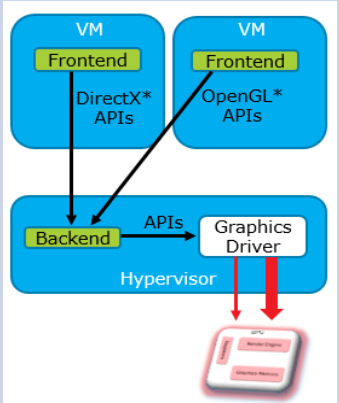	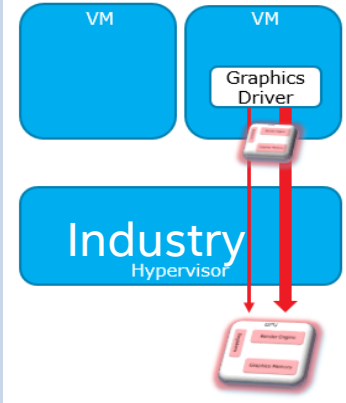	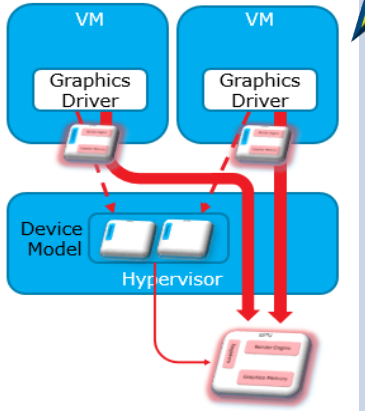
Pros: • Performance • Sharing **Cons:** • No media/GPGPU • Compatibility	**Pros:** • Performance • Capability **Cons:** • No sharing	**Pros:** • Performance • Capability • Sharing

Intel® Graphics Virtualization Technology (Intel® GVT-g)

Intel® Graphics Virtualization Technology (Intel® GVT-g) for vGPU based sharing

- Intel® GVT-g for Xen* (XenGT)
- Intel® GVT-g for KVM (KVMGT)

Performance	Features	Sharing
3DMark: 80% H.264 transcoding: 90% (of native performance)	Running Native Driver DirectX* 11.1 OpenGL* 4.2 OpenCL* 1.2 MediaSDK 16.2	Multiple VMs Support Ubuntu* Guest Support Windows* 7 x32/x64 Support Windows 8 x32/x64

Orchestration for Media Cloud

- Middleware extension
  - Support Intel® Graphics Virtualization Technology APIs in libvirt

- OpenStack* extension
  - GPU instance flavor
  - GPU aware scheduling
    - Find matching vGPU capability
    - QoS
  - GPU resource monitoring and allocation

Case Study: Virtualized Media Server

IDF15

Video Conference Usage Case

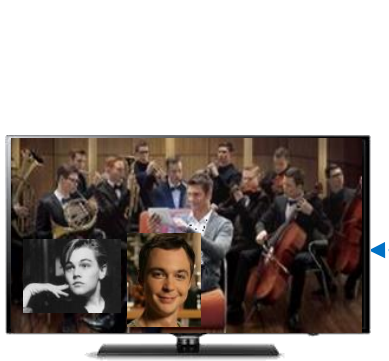

4K Display

4K Camera

1080P Camera

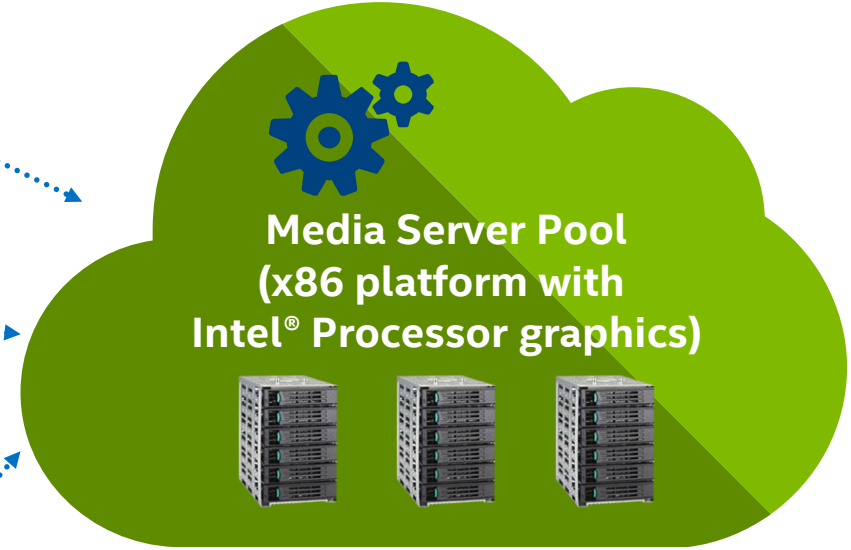

Intel Processor graphics-based Encoding, Decoding, and Composing.
OpenStack* Media Cloud

4K Video Conferencing on Media Plane NFV Demo @ MWC'15

Conversion of Telco Media to Intel® Architecture

Intel and Huawei* Joint Demo on MWC15

- H264/MPEG2/VC1/JPEG/MJEG decode
- H264/MPEG2 encode/transcode
- Video Post Processing

Resource Dashboard Monitor

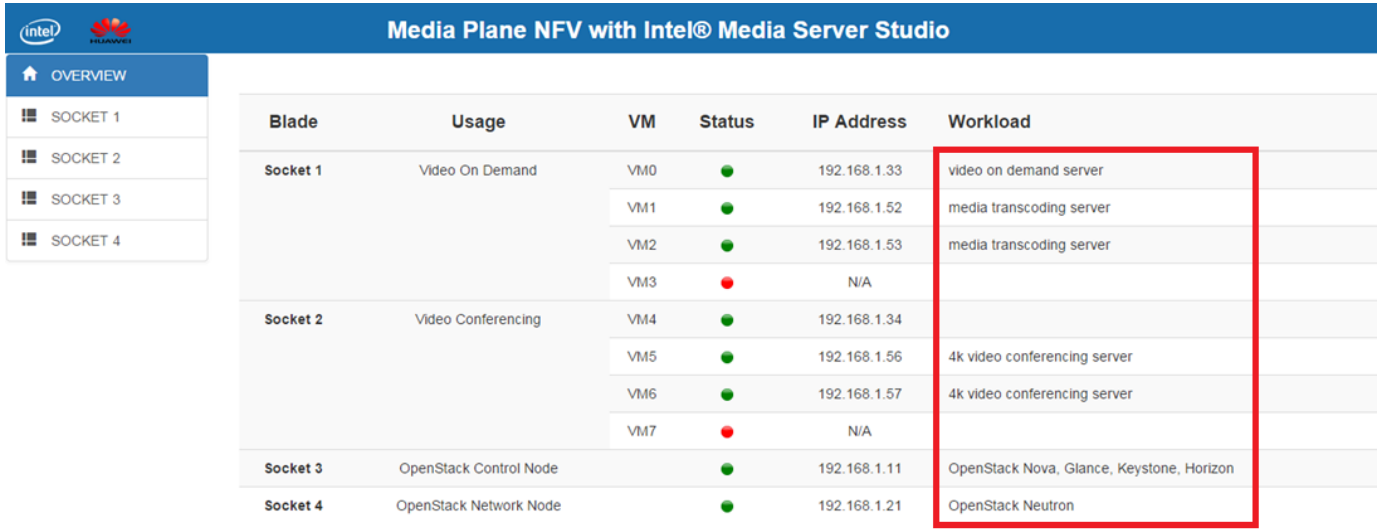

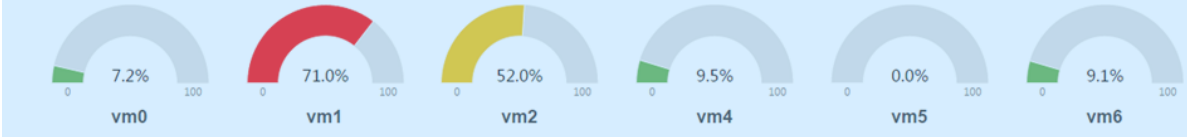

GPU Resource Monitoring

Optimize for Virtualized Media Server

IDF15

New Challenges

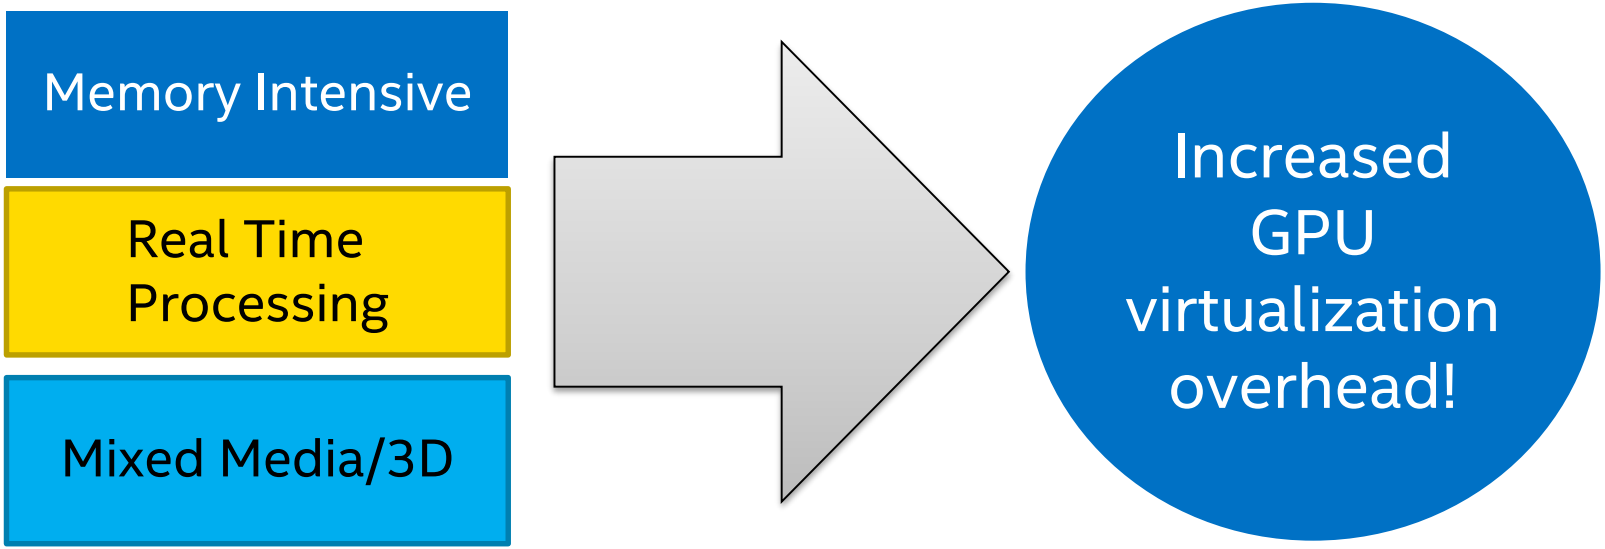

Optimizations

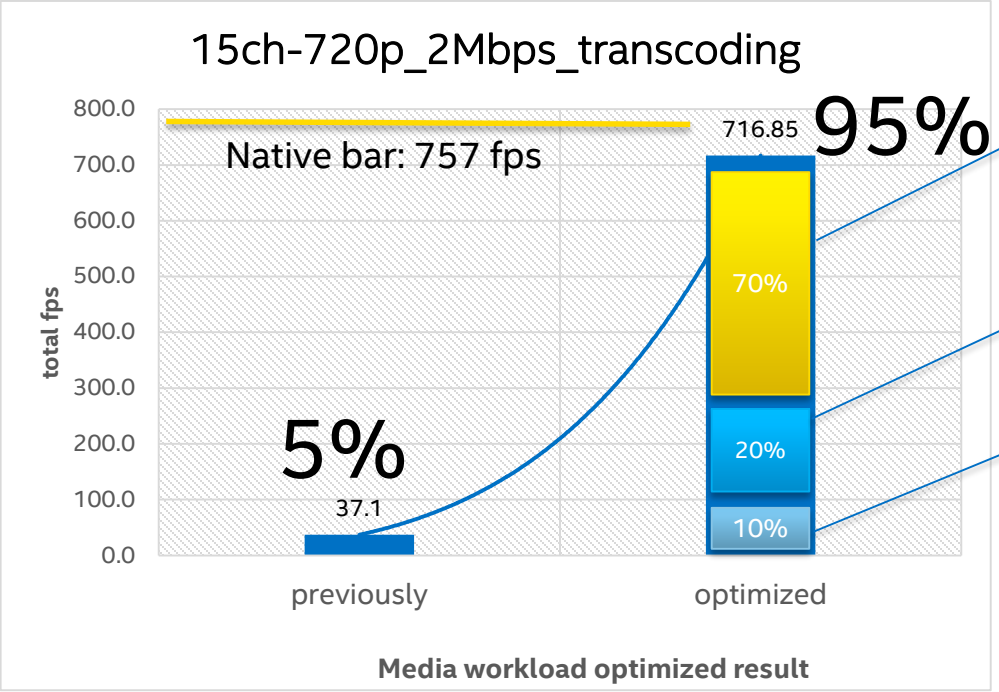

Smart shadow GPU page table

Enable cross-engine synchronization

Increase system memory

Config: I7 4770, Guest Ubuntu* 14.04LTS, 4GB mem, 1.5G GraphicMem, MediaSDK

Performance Summary

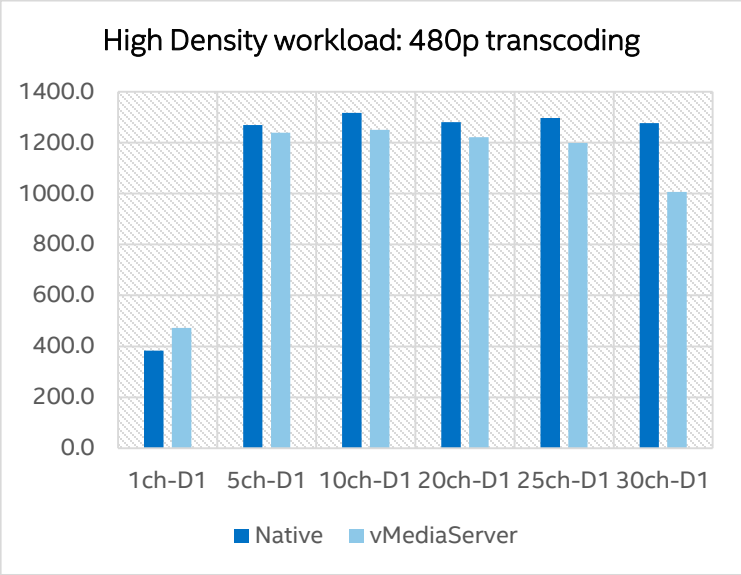

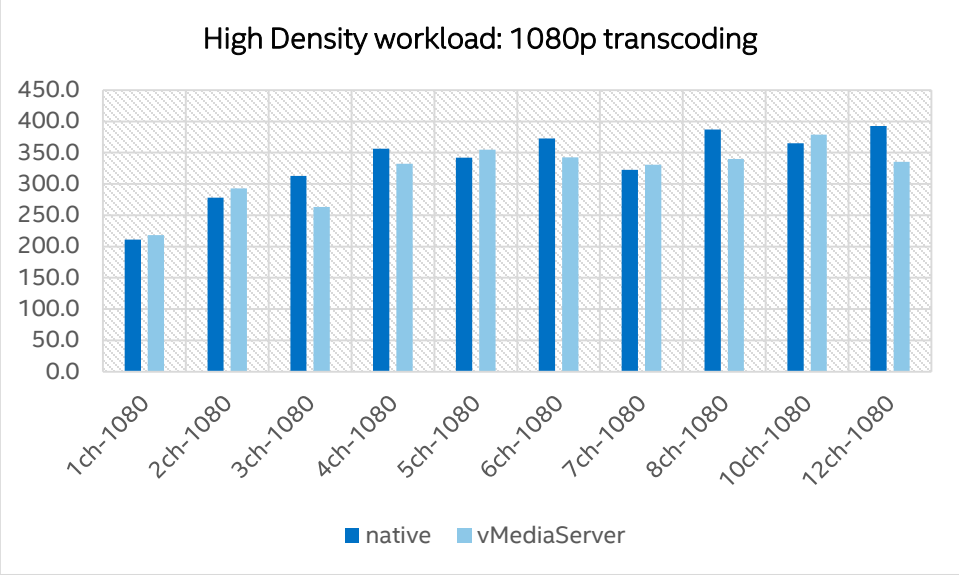

Config: I7 4770, Guest Ubuntu* 14.04LTS, 4GB mem, 1.5G GraphicMem, MediaSDK

Close to native performance!

CPU Transcoding vs. GPU Transcoding

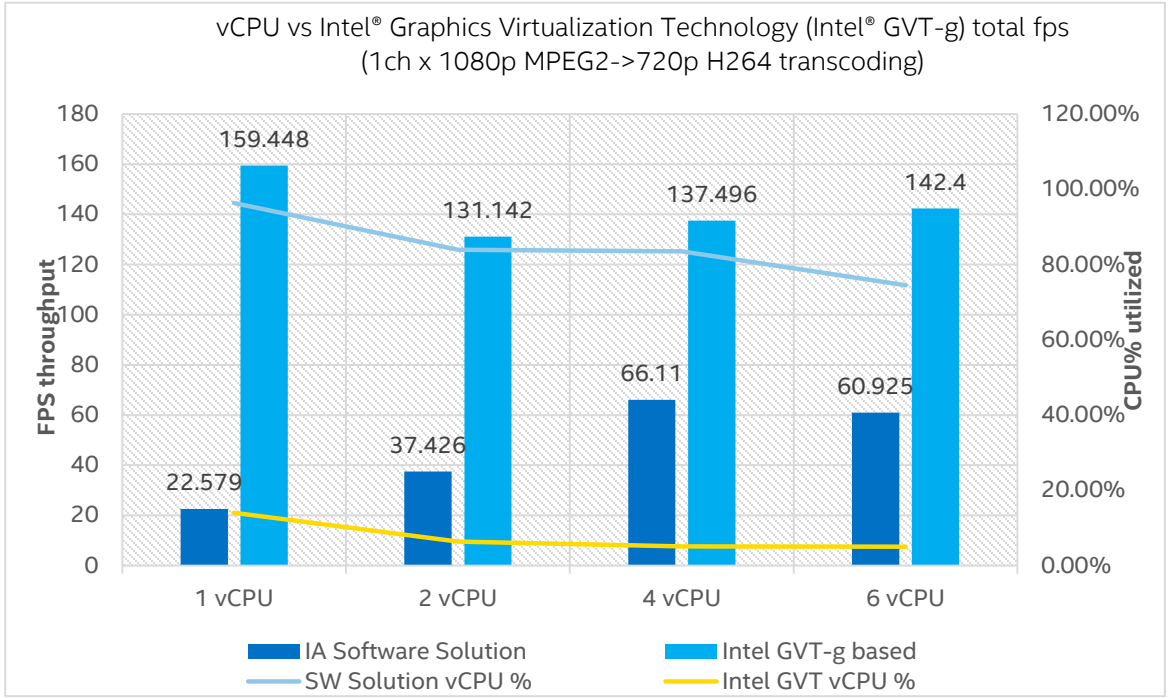

Lab data. Config: I7 4770, Guest Windows* 7_x64, 4GB mem, 1.5G GraphicMem, MediaSDK

Performance boost with Intel GVT-g!

Summary

- Industry media processing increases every year with big opportunity

- Media Cloud with Intel® Graphics Virtualization Technology (Intel® GVT-g) technology provides close to native performance, flexibility, scalable and relative low cost solution

Additional Sources of Information

- A PDF of this presentation is available from our Technical Session Catalog: www.intel.com/idfsessionsSZ. This URL is also printed on the top of Session Agenda Pages in the Pocket Guide.
- Audio recordings of sessions will be added to this catalog by April 13

- More web based info:
  - Project: https://01.org/xen/blogs
  - Demo video: https://www.youtube.com/watch?v=V2i8HCcAnY8
  - Demo video2: http://v.youku.com/v_show/id_XNzQ5MDg1MTg4.html

Other Technical Sessions

Session ID	Title	Day	Time	Room
DATS004	High-Density Media Solutions with Intel® Quick Sync Video	Thurs	14:30 – 15:30	Jing

Legal Notices and Disclaimers

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Learn more at intel.com, or from the OEM or retailer.

No computer system can be absolutely secure.

Tests document performance of components on a particular test, in specific systems. Differences in hardware, software, or configuration will affect actual performance. Consult other sources of information to evaluate performance as you consider your purchase. For more complete information about performance and benchmark results, visit http://www.intel.com/performance.

Cost reduction scenarios described are intended as examples of how a given Intel-based product, in the specified circumstances and configurations, may affect future costs and provide cost savings. Circumstances will vary. Intel does not guarantee any costs or cost reduction.

This document contains information on products, services and/or processes in development. All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest forecast, schedule, specifications and roadmaps.

Statements in this document that refer to Intel's plans and expectations for the quarter, the year, and the future, are forward-looking statements that involve a number of risks and uncertainties. A detailed discussion of the factors that could affect Intel's results and plans is included in Intel's SEC filings, including the annual report on Form 10-K.

The products described may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document.

Intel does not control or audit third-party benchmark data or the web sites referenced in this document. You should visit the referenced web site and confirm whether referenced data are accurate.

Intel, Core and the Intel logo are trademarks of Intel Corporation in the United States and other countries.

OpenCL and the OpenCL logo are trademarks of Apple Inc. used by permission by Khronos.

*Other names and brands may be claimed as the property of others.

Risk Factors

The above statements and any others in this document that refer to plans and expectations for the first quarter, the year and the future are forward-looking statements that involve a number of risks and uncertainties. Words such as "anticipates," "expects," "intends," "plans," "believes," "seeks," "estimates," "may," "will," "should" and their variations identify forward-looking statements. Statements that refer to or are based on projections, uncertain events or assumptions also identify forward-looking statements. Many factors could affect Intel's actual results, and variances from Intel's current expectations regarding such factors could cause actual results to differ materially from those expressed in these forward-looking statements. Intel presently considers the following to be important factors that could cause actual results to differ materially from the company's expectations. Demand for Intel's products is highly variable and could differ from expectations due to factors including changes in the business and economic conditions; consumer confidence or income levels; customer acceptance of Intel's and competitors' products; competitive and pricing pressures, including actions taken by competitors; supply constraints and other disruptions affecting customers; changes in customer order patterns including order cancellations; and changes in the level of inventory at customers. Intel's gross margin percentage could vary significantly from expectations based on capacity utilization; variations in inventory valuation, including variations related to the timing of qualifying products for sale; changes in revenue levels; segment product mix; the timing and execution of the manufacturing ramp and associated costs; excess or obsolete inventory; changes in unit costs; defects or disruptions in the supply of materials or resources; and product manufacturing quality/yields. Variations in gross margin may also be caused by the timing of Intel product introductions and related expenses, including marketing expenses, and Intel's ability to respond quickly to technological developments and to introduce new features into existing products, which may result in restructuring and asset impairment charges. Intel's results could be affected by adverse economic, social, political and physical/infrastructure conditions in countries where Intel, its customers or its suppliers operate, including military conflict and other security risks, natural disasters, infrastructure disruptions, health concerns and fluctuations in currency exchange rates. Results may also be affected by the formal or informal imposition by countries of new or revised export and/or import and doing-business regulations, which could be changed without prior notice. Intel operates in highly competitive industries and its operations have high costs that are either fixed or difficult to reduce in the short term. The amount, timing and execution of Intel's stock repurchase program and dividend program could be affected by changes in Intel's priorities for the use of cash, such as operational spending, capital spending, acquisitions, and as a result of changes to Intel's cash flows and changes in tax laws. Product defects or errata (deviations from published specifications) may adversely impact our expenses, revenues and reputation. Intel's results could be affected by litigation or regulatory matters involving intellectual property, stockholder, consumer, antitrust, disclosure and other issues. An unfavorable ruling could include monetary damages or an injunction prohibiting Intel from manufacturing or selling one or more products, precluding particular business practices, impacting Intel's ability to design its products, or requiring other remedies such as compulsory licensing of intellectual property. Intel's results may be affected by the timing of closing of acquisitions, divestitures and other significant transactions. A detailed discussion of these and other factors that could affect Intel's results is included in Intel's SEC filings, including the company's most recent reports on Form 10-Q, Form 10-K and earnings release.

Backup

Summary for Media Cloud

Media Cloud Requirements	Specialized Hardware (DSP, FPGA, ASIC)	Pure Software	MSS on Intel GVT
Transcoding Throughput	Good	Good, but expensive	Good
Real-time & low-latency	Good	Hard, but expensive	Good
Virtualization	Hard	Easy	Easy
Cloud Integration	Hard	Easy	Easy
Cost on Dev, Ops & Upgrade	High	Low	Low

High Density Media workload: GPU Memory Impact

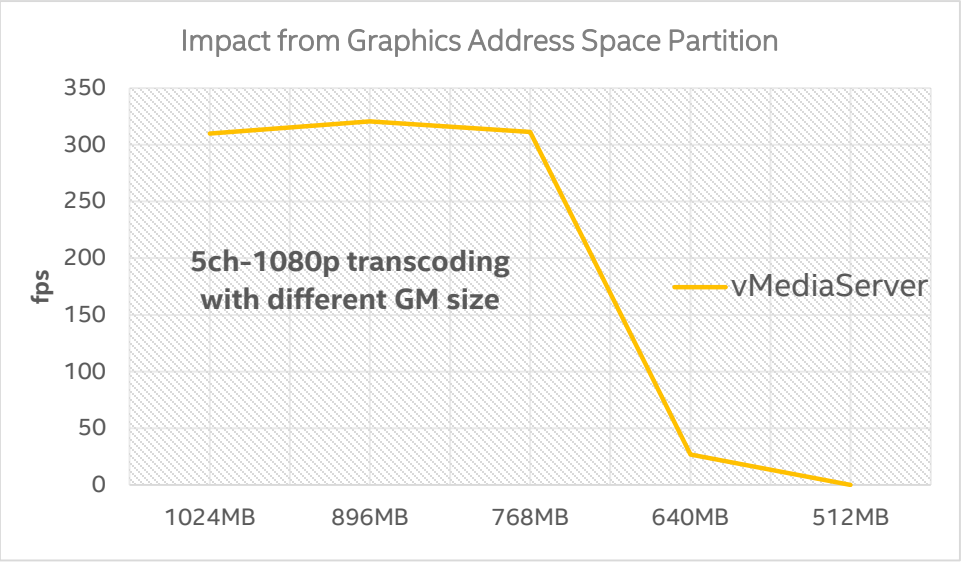

Config: I7 4770, Guest Ubuntu* 14.04LTS, 4GB mem, Graphic Mem 512MB-1024MB

- Default GM Partition 512MB per VM
- fps 90% drop for 5ch-1080p transcoding

Smart Shadow Page Table

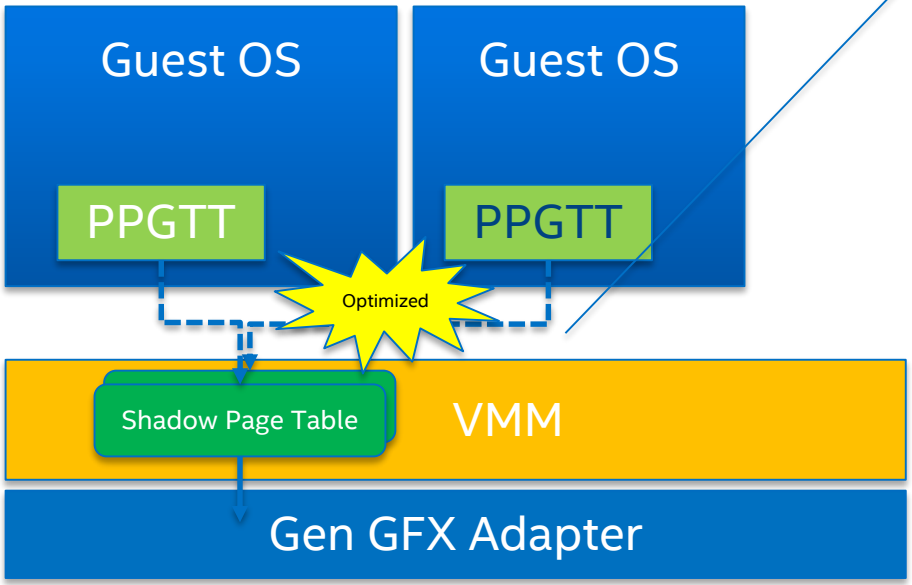

- Lots of PPGTT access due to memory intensive workload
- Optimization: cache the access in Guest and flush to Shadow PPGTT only GPU HW is about to access

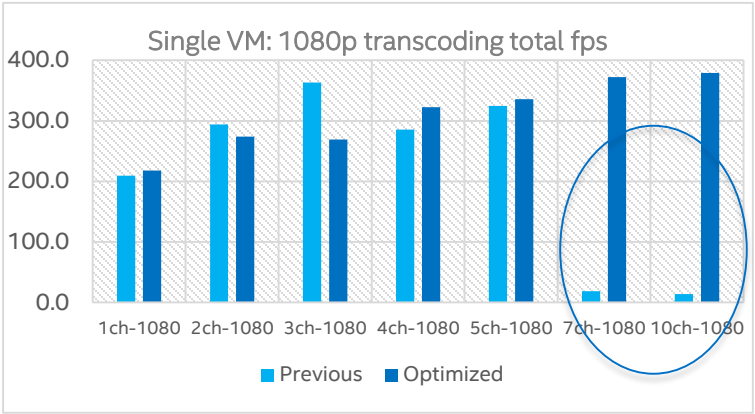

Config: I7 4770, Guest Ubuntu* 14.04LTS, 4GB mem, Graphic Mem 768MB

Cross Engine Synchronization