

基于英特尔® 图形虚拟化技术（英特尔® GVT-g） 和 OpenStack* 的媒体云



Xiao Zheng – 英特尔公司软件工程师

SFTS002

芯动，行动，共创未来！ 



议程

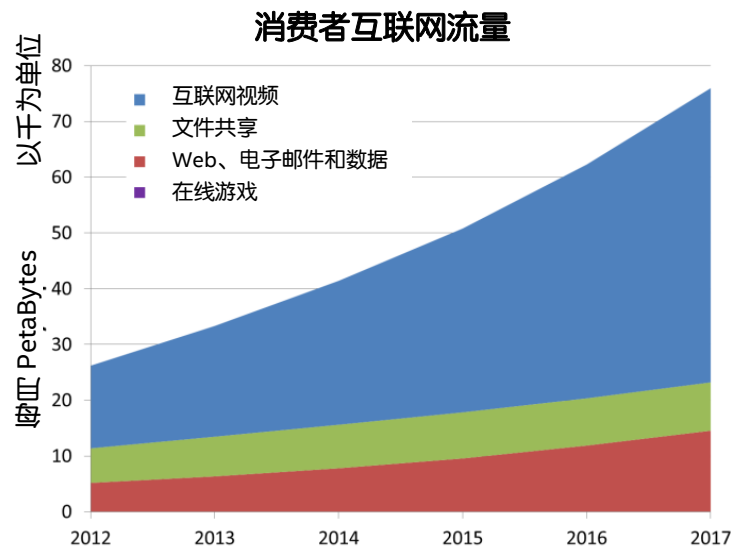
- 媒体云
- 媒体云基础架构
- 案例研究：虚拟化媒体服务器
- 针对虚拟化媒体服务器而优化



媒体云

IDF15

媒体处理机遇



资料来源: Cisco® Systems Inc, 视觉网络指数 (VNI), 2013, 2015



资料来源: 2015.1 iResearch Inc., www.iresearch.com.cn

互联网视频流量的年复合增长率预计为 29%，到 2017 年将占据消费者流量的 69%

行业前景



视频投递

存储/视频流, 转码
离线/实时处理、
OTT、VOD、IPTV



视频分析

搜索、监控



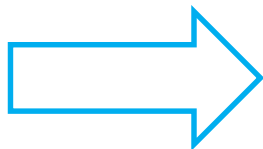
媒体云



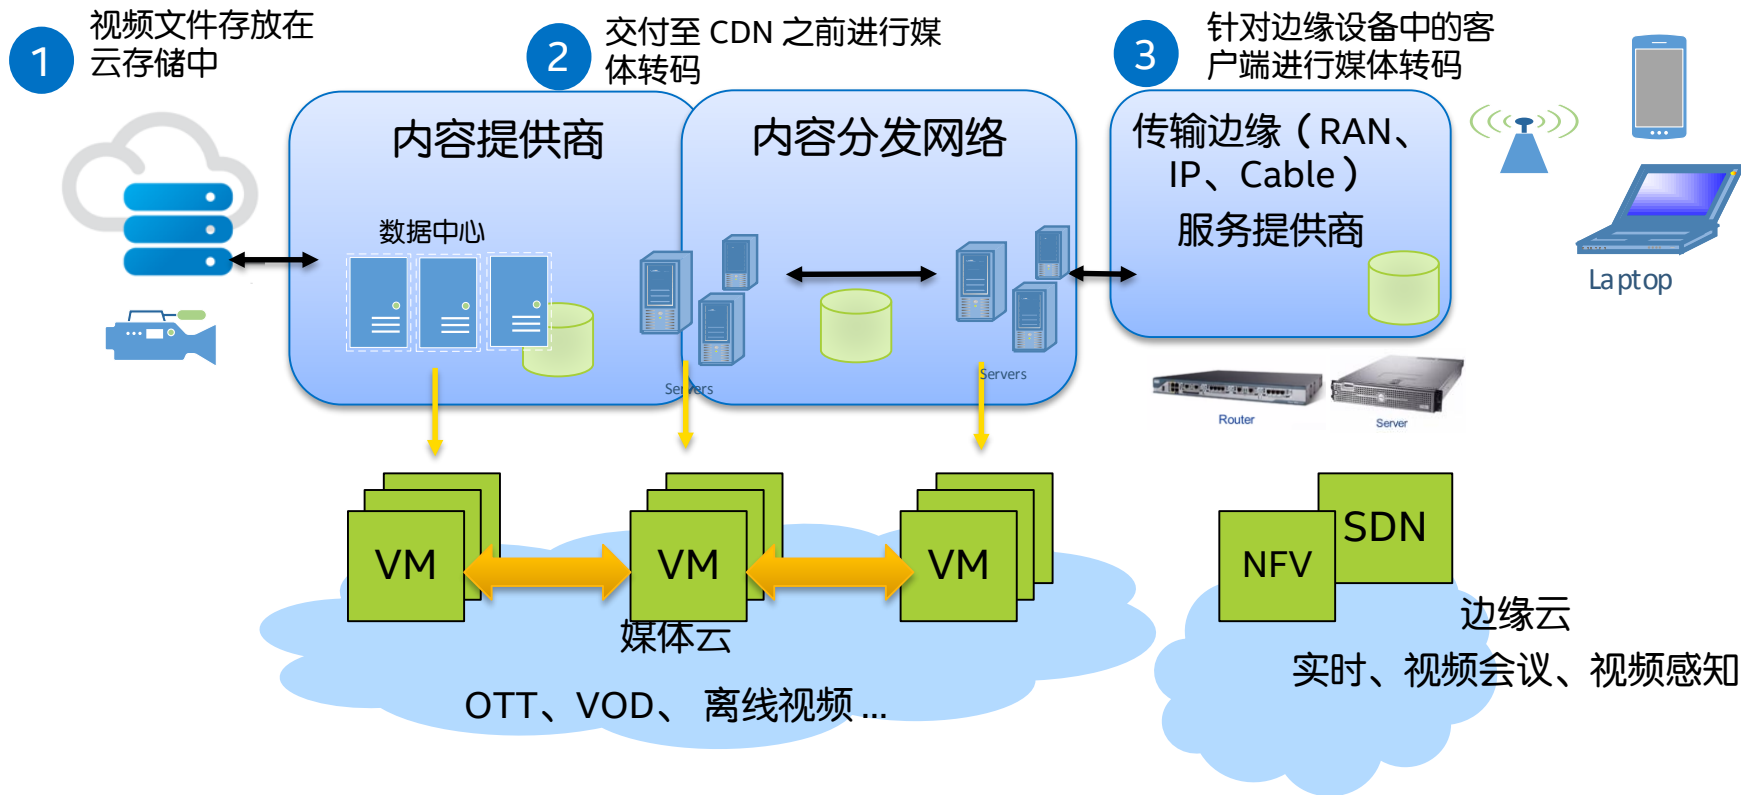
视频交付
存储/流, 转码
离线、实时



视频分析
搜索、监控

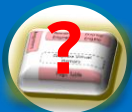


媒体交付示例



技术障碍

无 GPU 虚拟化



在没有 GPU 加速功能的情况下，低成本 CPU 转码吞吐量非常低

DSP 解决方案难以在云环境中集成

无云协调



无 GPU 实例感知能力
无 vGPU 调度
无 vGPU 资源监控功能

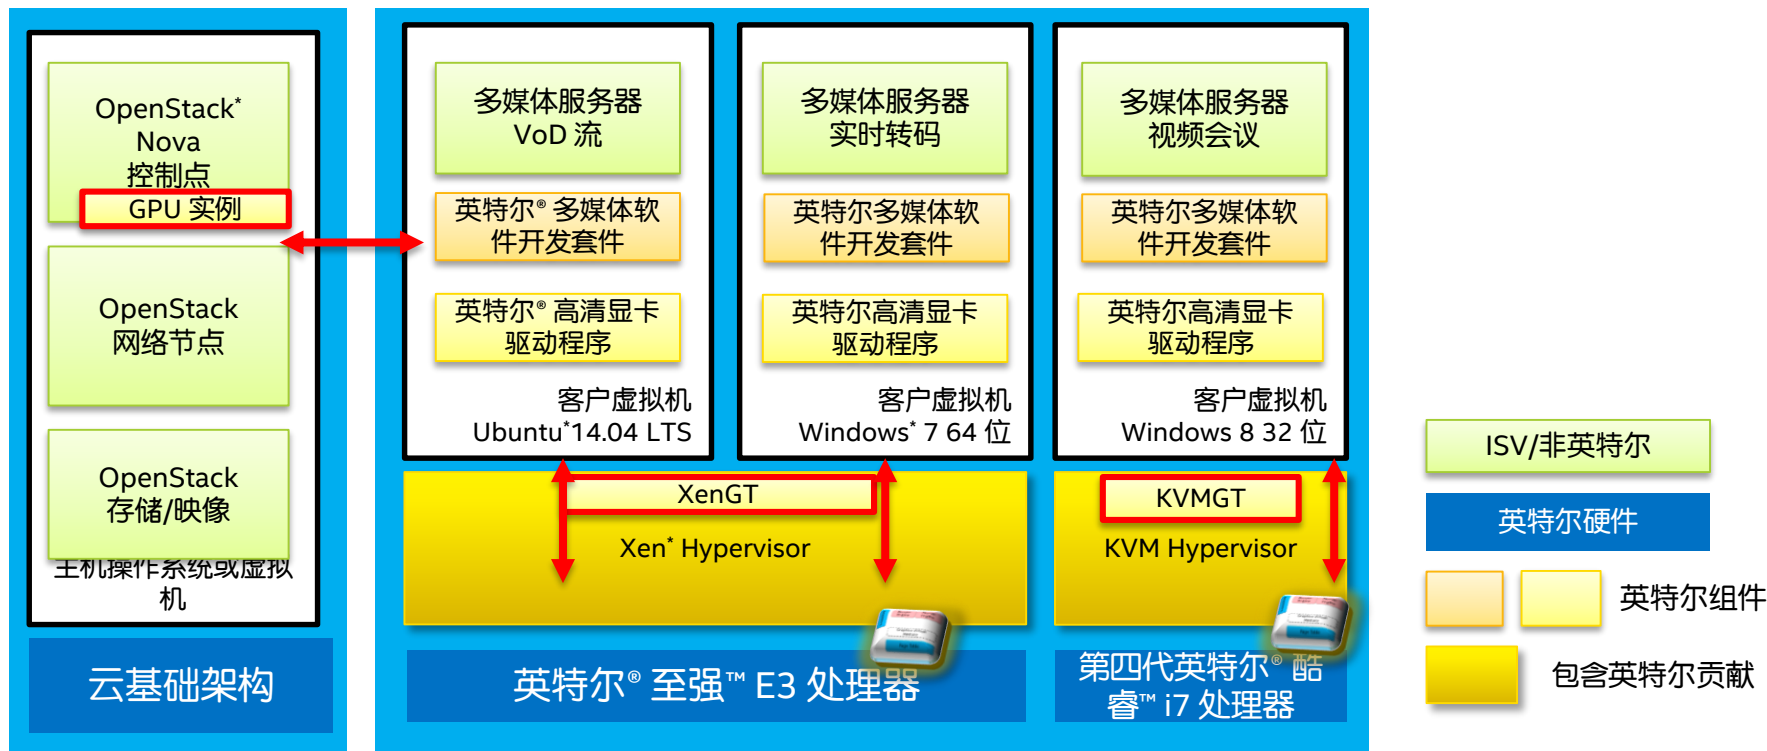
基于英特尔® 图形虚拟化
解决方案的媒体云





媒体云基础架构

基本架构图



GPU 虚拟化要求



性能



直接 GPU 加速。



功能



多标准的兼容性

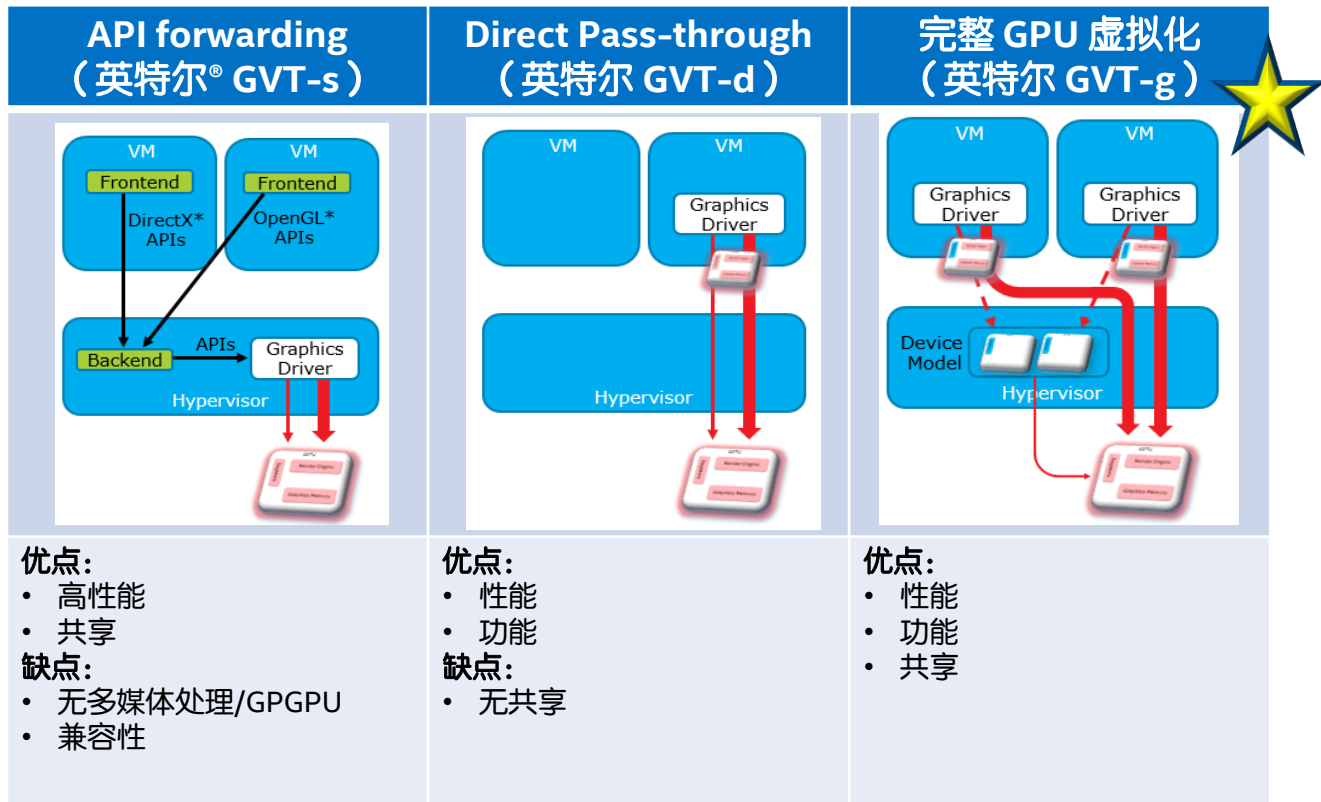


共享



多虚拟机共享

GPU 虚拟化方法



英特尔® 图形虚拟化技术（英特尔® GVT-g）

英特尔® 图形虚拟化技术（英特尔® GVT-g），支持基于 vGPU 的共享功能

- 面向 Xen* 的英特尔® GVT-g (XenGT)
- 面向 KVM 的英特尔® GVT-g (KVMGT)

性能

3DMark:80%
H.264 转码: 90%（本机性能）

特性

运行本机驱动程序
DirectX* 11.1
OpenGL* 4.2
OpenCL* 1.2
MediaSDK 16.2

共享

多个虚拟机
支持 Ubuntu* 客机
支持 Windows* 7 x32/x64
支持 Windows 8 x32/x64

面向媒体云的协调

- 中间件扩展
 - 支持英特尔® 图形虚拟化技术 API (libvirt)
- OpenStack* 扩展
 - GPU 实例 (Flavor)
 - GPU 感知调度
 - 查找匹配的 vGPU 功能
 - QoS
 - GPU 资源监控和分配



案例研究：虚拟化的媒体服务器

视频会议使用案例



媒体平台 NFV 演示上的 4K 视频会议（2015 年世界移动通信大会）

IDF15

媒体面应用转至英特尔® 架构



英特尔与华为* 在 MWC15 上联合演示

- H264/MPEG2/VC1/JPEG/MJEG 解码
- H264/MPEG2 编码/转码
- 后期视频处理

资源指示板监控

intel

huanan

Media Plane NFV with Intel® Media Server Studio

OVERVIEW

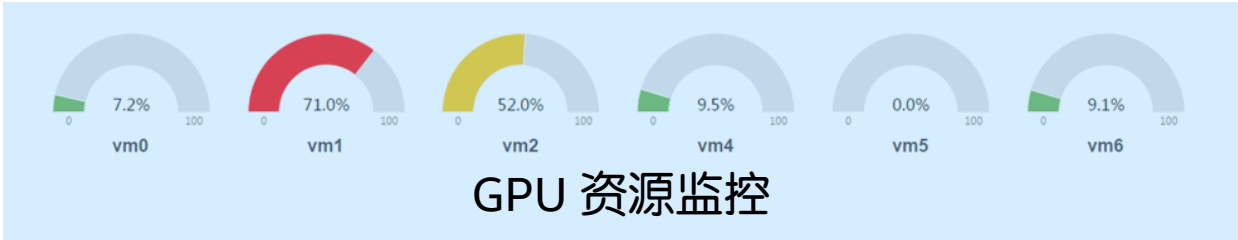
SOCKET 1

SOCKET 2

SOCKET 3

SOCKET 4

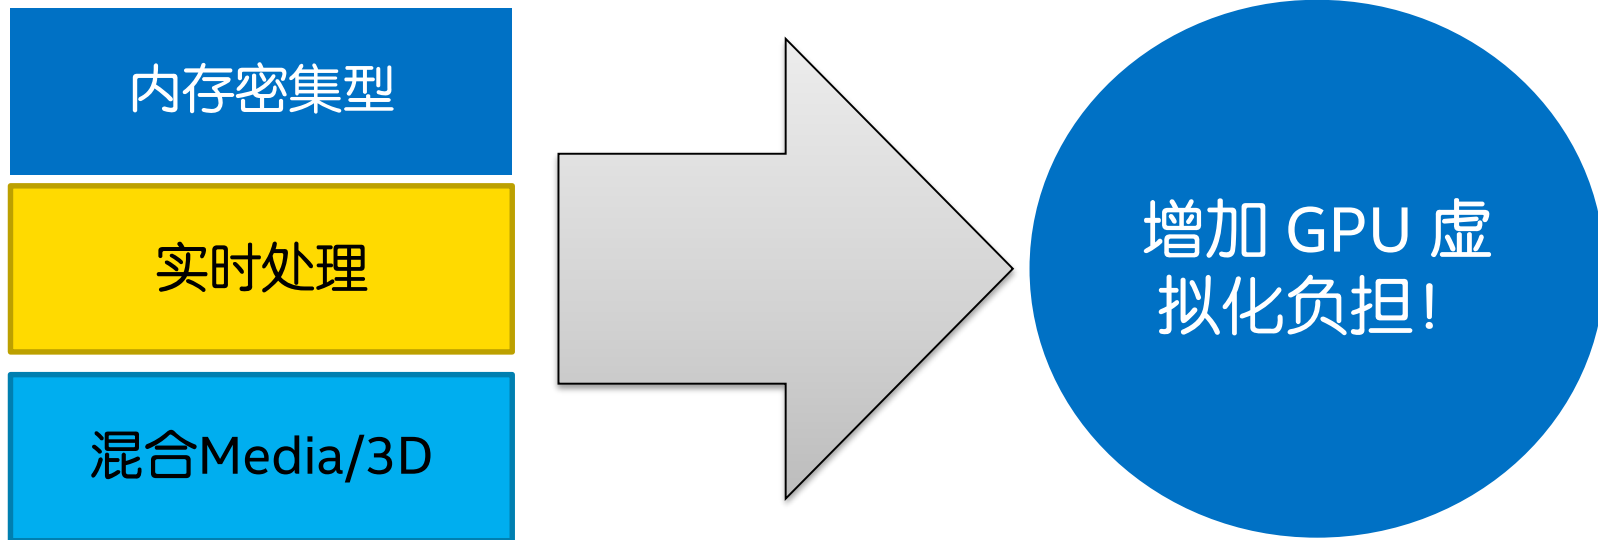
Blade	Usage	VM	Status	IP Address	Workload
Socket 1	Video On Demand	VM0	●	192.168.1.33	video on demand server
		VM1	●	192.168.1.52	media transcoding server
		VM2	●	192.168.1.53	media transcoding server
		VM3	●	N/A	
Socket 2	Video Conferencing	VM4	●	192.168.1.34	
		VM5	●	192.168.1.56	4k video conferencing server
		VM6	●	192.168.1.57	4k video conferencing server
		VM7	●	N/A	
Socket 3	OpenStack Control Node		●	192.168.1.11	OpenStack Nova, Glance, Keystone, Horizon
Socket 4	OpenStack Network Node		●	192.168.1.21	OpenStack Neutron





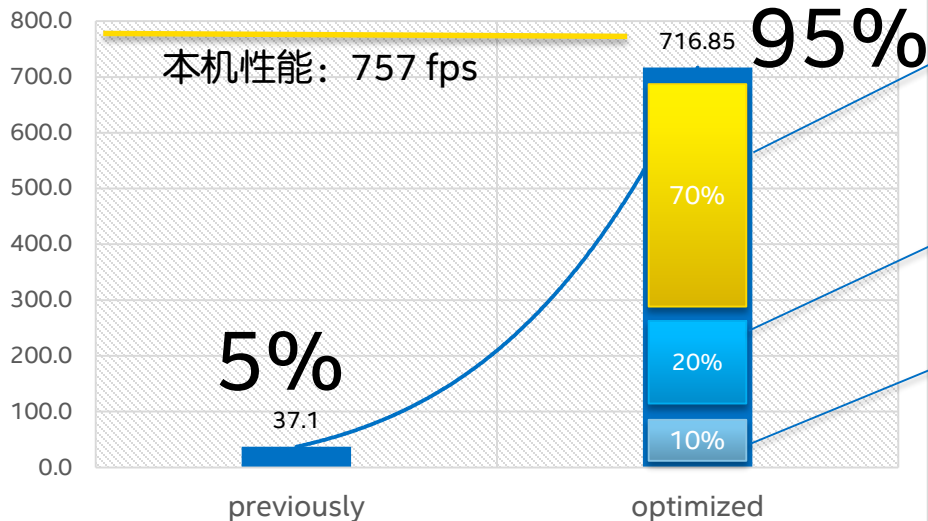
虚拟化多媒体服务器的优化

新的挑战



优化

15ch-720p_2Mbps_transcoding



智能影子 (Smart shadow) GPU 页表

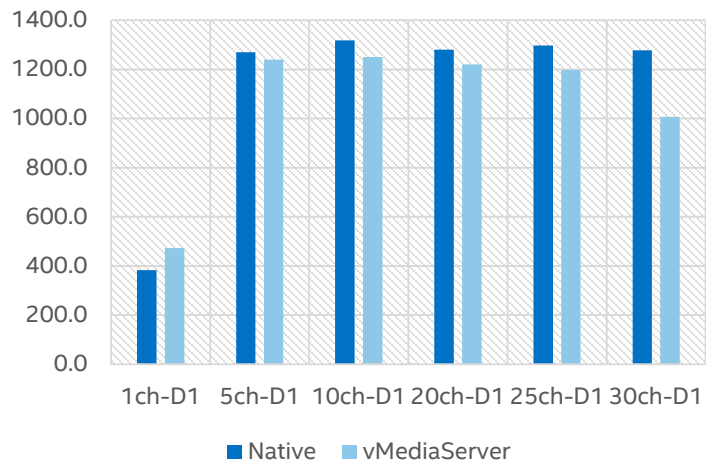
支持跨引擎同步

提高系统内存

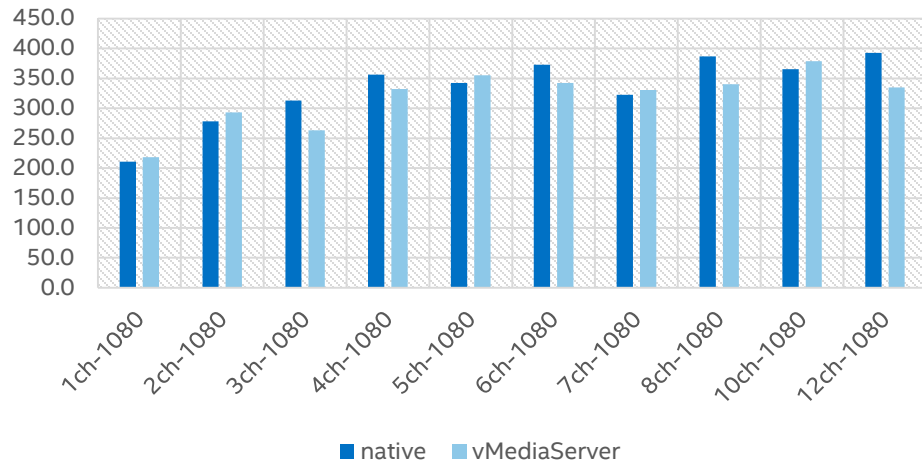
配置: i7 4770, 客机 Ubuntu* 14.04LTS、4GB 内存、1.5G GraphicMem、MediaSDK

性能摘要

高密度工作负载：480p 转码



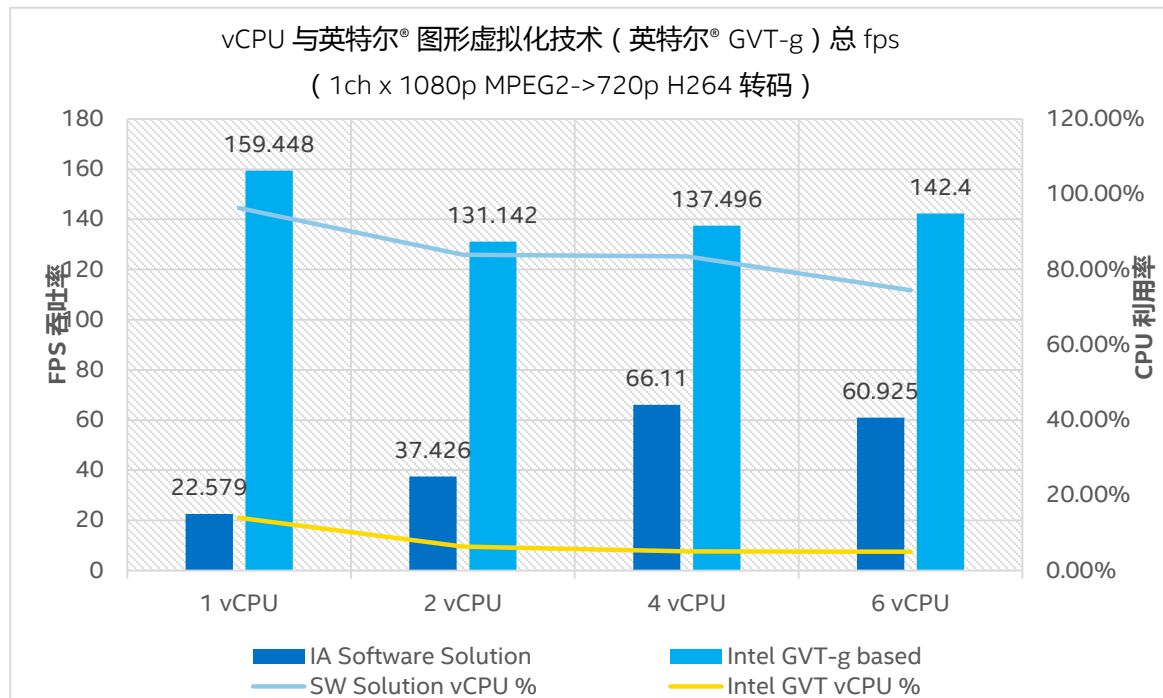
高密度工作负载：1080p 转码



配置：I7 4770，客机 Ubuntu* 14.04LTS、4GB 内存、1.5G GraphicMem、MediaSDK

接近本机性能！

CPU 转码与 GPU 转码



英特尔 GVT-g 助力
性能显著提升!

实验室数据。配置: I7 4770、客机 Windows* 7_x64、4GB 内存、1.5G GraphicMem、MediaSDK

总结

- 行业媒体处理领域每年以较快的速度增长，市场机遇巨大
- 采用英特尔® 图形虚拟化技术（英特尔® GVT-g）的媒体云提供了接近本机的性能、灵活性、可扩展性以及相对更低的成本。

查看更多相关信息

- 您可通过IDF网站的技术课程目录下载此次课程的PDF演讲文稿
www.intel.com/idfsessionsSZ 也可以在技术课程目录的首页链接下载。
- 您可以在4月13日以后通过该目录收看所有课程的视频录像
- 更多网络信息：
 - 项目: <https://01.org/xen/blogs>
 - 演示视频: <https://www.youtube.com/watch?v=V2i8HCcAnY8>
 - 演示视频 2: http://v.youku.com/v_show/id_XNzQ5MDg1MTg4.html

其他技术课程

课程编号	标题	日期	时间	教室
DATS004	采用英特尔® 快速视频同步技术的高密度媒体解决方案	星期四	14:30 – 15:30	JING 景

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.

© 2015 Intel Corporation.

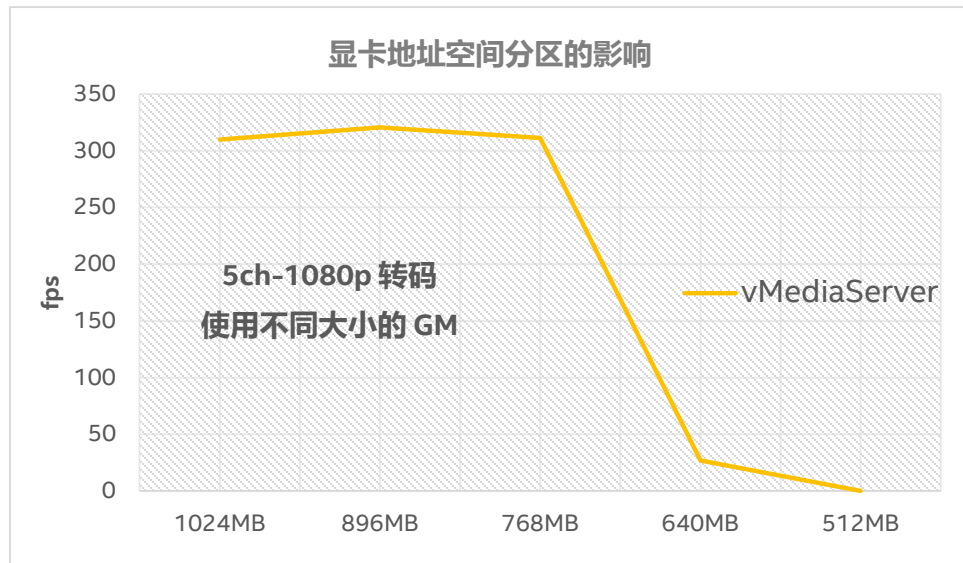
Risk Factors

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.

媒体云总结

媒体云要求	专用硬件（DSP、FPGA、ASIC）	纯软件	英特尔 GVT 上的 MSS
转码吞吐率	良好	良好，但是过于昂贵	良好
实时与低延迟	良好	困难，但是过于昂贵	良好
虚拟化	困难	轻松	轻松
云集成	困难	轻松	轻松
开发、运营与升级成本	高	低	低

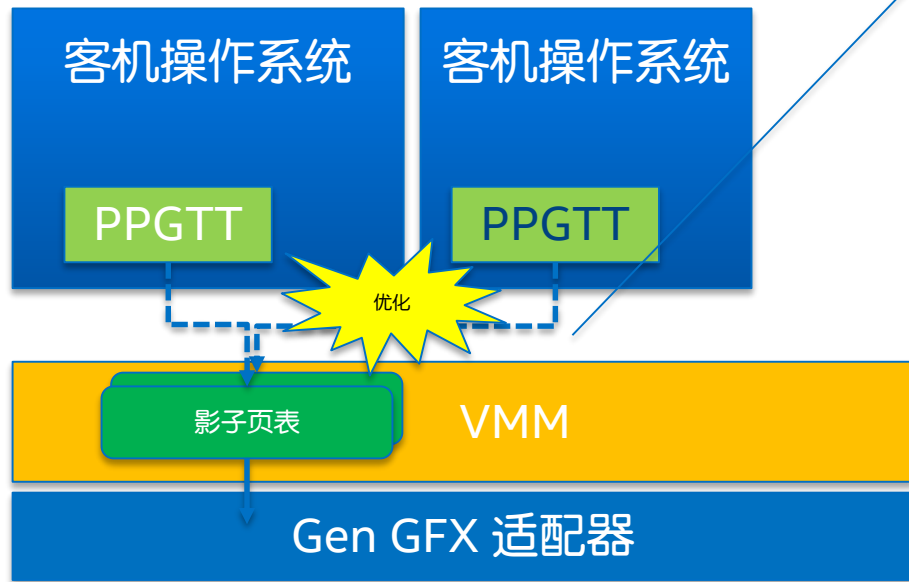
高密度媒体工作负载：GPU 内存影响



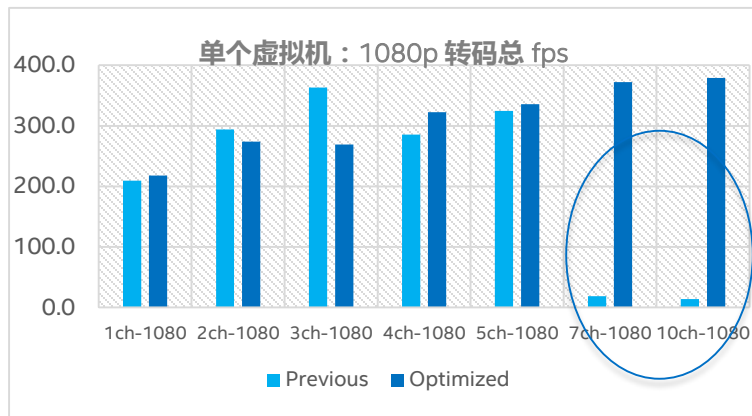
配置：I7 4770、客机 Ubuntu^{*} 14.04LTS、4GB 内存、显存 512MB-1024MB

- 每个虚拟机的默认 GM 分区为 512MB
- 5ch-1080p 转码的 fps 降低 90%

智能影子 (Smart shadow) 页表



- 内存密集型工作负载会造成大量 PPGTT 访问
- 优化：缓存客机中的访问，并只在 GPU 硬件要访问的 Shadow PPGT 的时候一次性同步到影子页表



配置：I7 4770、客机 Ubuntu* 14.04LTS、4GB 内存、显存 768MB

跨引擎同步

