



# Intel<sup>®</sup> Graphics Virtualization Technology

Kevin Tian  
Graphics Virtualization Architect



# Legal Disclaimer

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT.

A "Mission Critical Application" is any application in which failure of the Intel Product could result, directly or indirectly, in personal injury or death. SHOULD YOU PURCHASE OR USE INTEL'S PRODUCTS FOR ANY SUCH MISSION CRITICAL APPLICATION, YOU SHALL INDEMNIFY AND HOLD INTEL AND ITS SUBSIDIARIES, SUBCONTRACTORS AND AFFILIATES, AND THE DIRECTORS, OFFICERS, AND EMPLOYEES OF EACH, HARMLESS AGAINST ALL CLAIMS COSTS, DAMAGES, AND EXPENSES AND REASONABLE ATTORNEYS' FEES ARISING OUT OF, DIRECTLY OR INDIRECTLY, ANY CLAIM OF PRODUCT LIABILITY, PERSONAL INJURY, OR DEATH ARISING IN ANY WAY OUT OF SUCH MISSION CRITICAL APPLICATION, WHETHER OR NOT INTEL OR ITS SUBCONTRACTOR WAS NEGLIGENT IN THE DESIGN, MANUFACTURE, OR WARNING OF THE INTEL PRODUCT OR ANY OF ITS PARTS.

Intel may make changes to specifications and product descriptions at any time, without notice. Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined". Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them. The information here is subject to change without notice. Do not finalize a design with this information. The products described in this document 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.

Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.

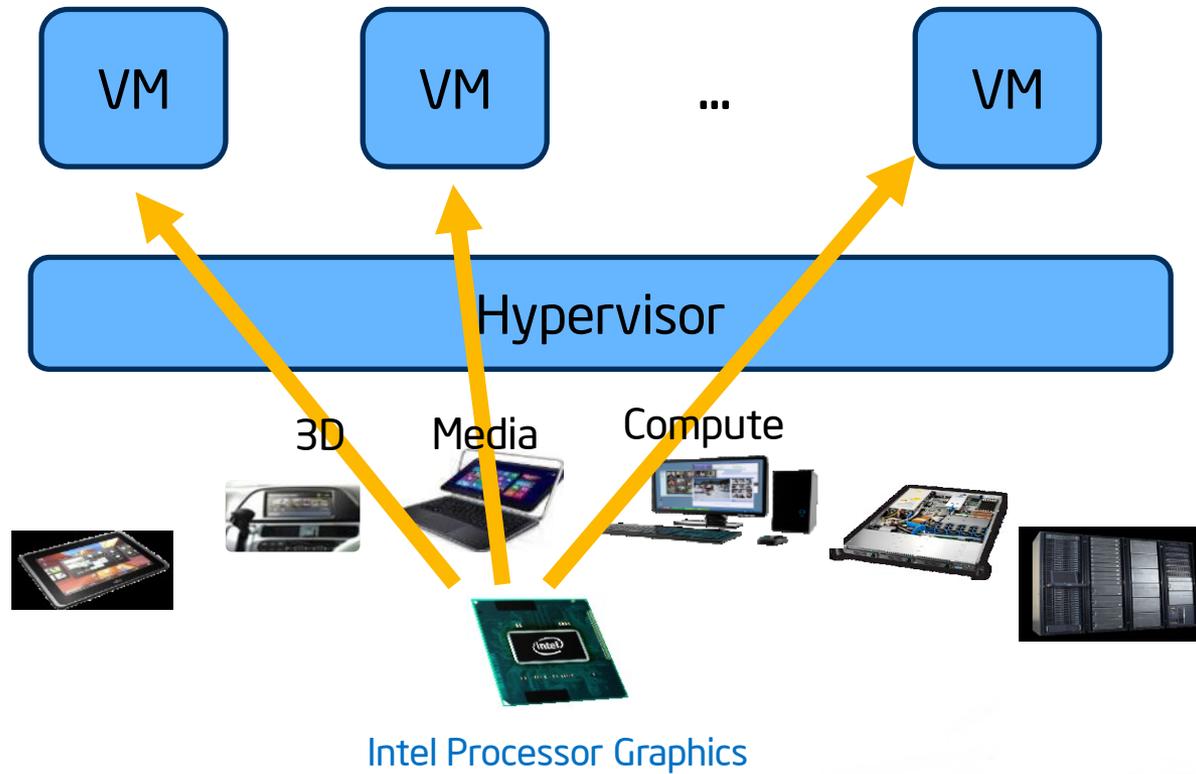
Copies of documents which have an order number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725, or go to: <http://www.intel.com/design/literature.htm>

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

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

Copyright ©2014 Intel Corporation.

# GPU Virtualization

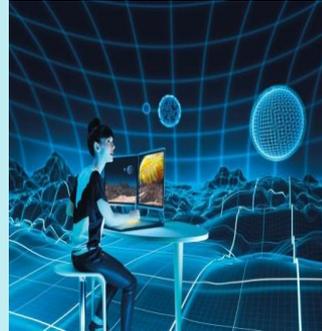


# GPU Cloud



## Video Delivery

Store/Stream, Transcode



## Cloud Graphics

Gaming, Remote Apps,  
Rendering



## Visual Understanding

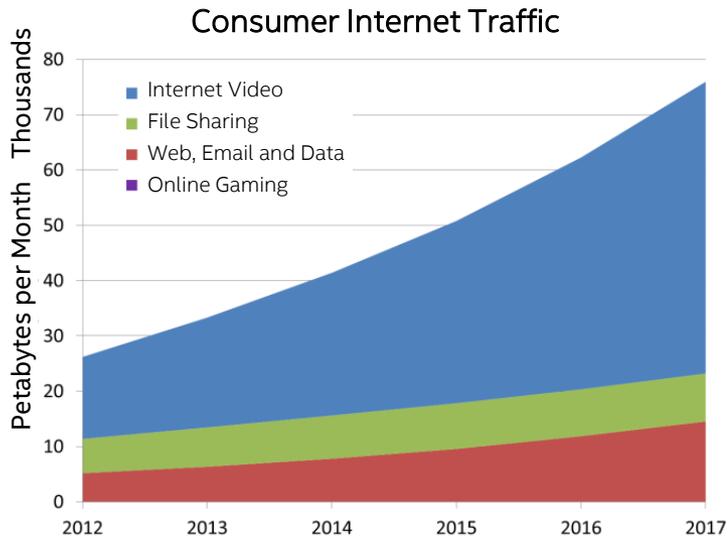
Search, Surveillance

GPU-as-a-service

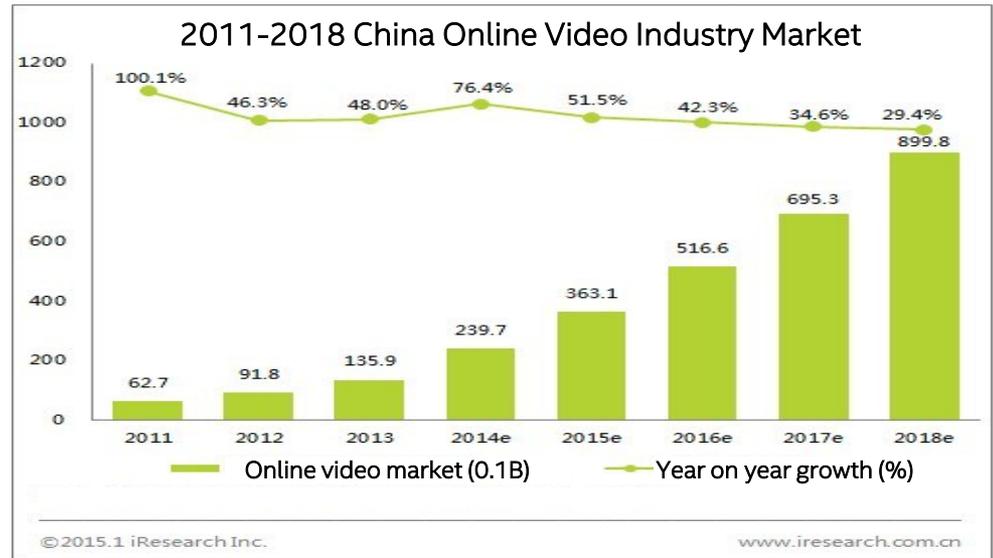
GPU Cloud



# Media Processing Opportunity



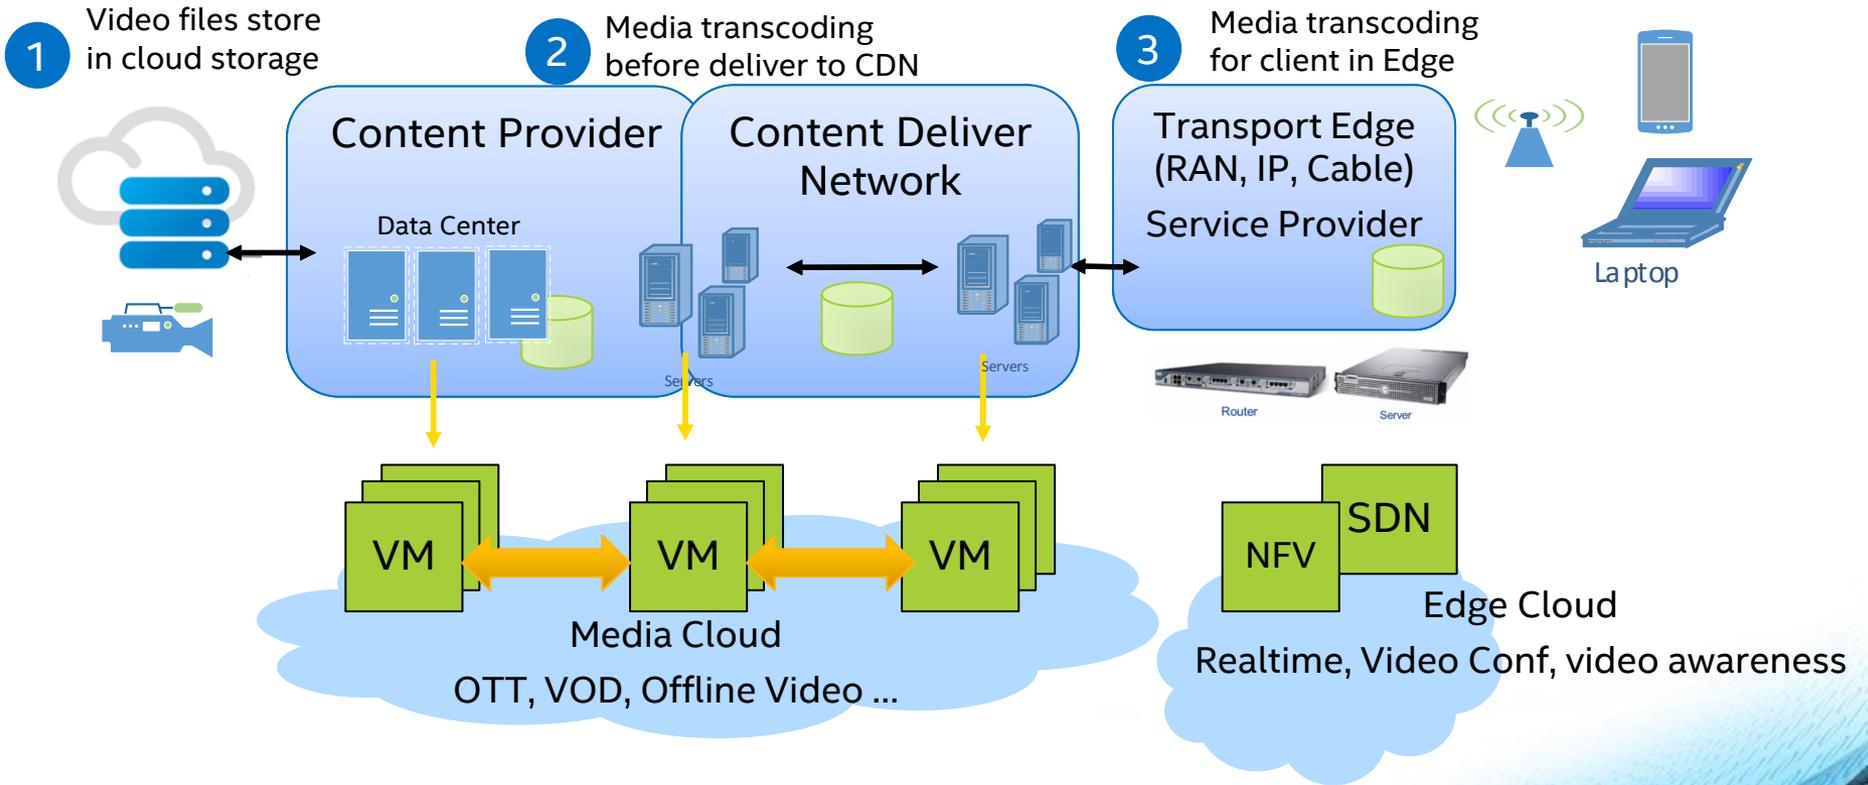
Source: Cisco Systems Inc., Visual Network Index (VNI), 2013, 2015



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

# Media Delivery Example



# Requirements of GPU Virtualization



Performance



Direct GPU acceleration



Capability



Consistent visual experience



Sharing



Multiple Virtual Machines

# GPU Virtualization Approaches



API Forwarding	Direct Pass-thru	Full GPU Virtualization
<p>The diagram shows two VMs at the top. The left VM contains 'DirectX* APIs' and the right VM contains 'OpenGL* APIs'. Both have a 'Frontend' box. Arrows from both Frontends point to a 'Backend' box within a 'Hypervisor' container. From the Backend, an arrow labeled 'APIs' points to a 'Graphics Driver' box, also within the Hypervisor. From the Graphics Driver, two red arrows point down to a 'GPU' component.</p>	<p>The diagram shows two VMs at the top. The right VM contains a 'Graphics Driver' box. An arrow from this Graphics Driver points down to a 'Hypervisor' container. From the Hypervisor, two red arrows point down to a 'GPU' component.</p>	<p>The diagram shows two VMs at the top, each containing a 'Graphics Driver' box. Arrows from both Graphics Drivers point down to a 'Device Model' box within a 'Hypervisor' container. From the Device Model, two red arrows point down to a 'GPU' component.</p>
<p><b>Pros:</b></p> <ul style="list-style-type: none"> <li>• Performance</li> <li>• Sharing</li> </ul> <p><b>Cons:</b></p> <ul style="list-style-type: none"> <li>• No media/GPGPU</li> <li>• Compatibility</li> </ul>	<p><b>Pros:</b></p> <ul style="list-style-type: none"> <li>• Performance</li> <li>• Capability</li> </ul> <p><b>Cons:</b></p> <ul style="list-style-type: none"> <li>• No sharing</li> </ul>	<p><b>Pros:</b></p> <ul style="list-style-type: none"> <li>• Performance</li> <li>• Capability</li> <li>• Sharing</li> </ul>

# Intel® Graphics Virtualization Technology

- Intel® GVT-s for API level GPU sharing
  - Make existing API forwarding protocols running best on Intel platform
- Intel® GVT-d for direct GPU passthrough
  - Xen GPU passthrough upstreaming in progress
  - KVM PoC patch in community
- Intel® GVT-g for vGPU based sharing
  - Achieve a good balance of performance, feature and sharing
  - Xen implementations (a.k.a XenGT) in production quality on HSW
    - BDW support is in alpha quality in 2015/Q1 release
  - KVM support (a.k.a KVMGT) in prototype quality on HSW

# Intel GVT-g: Full GPU Virtualization

## Performance

3DMark: **80%**

H.264 transcoding: **90%**  
(of native performance)

## Feature

**Native driver**

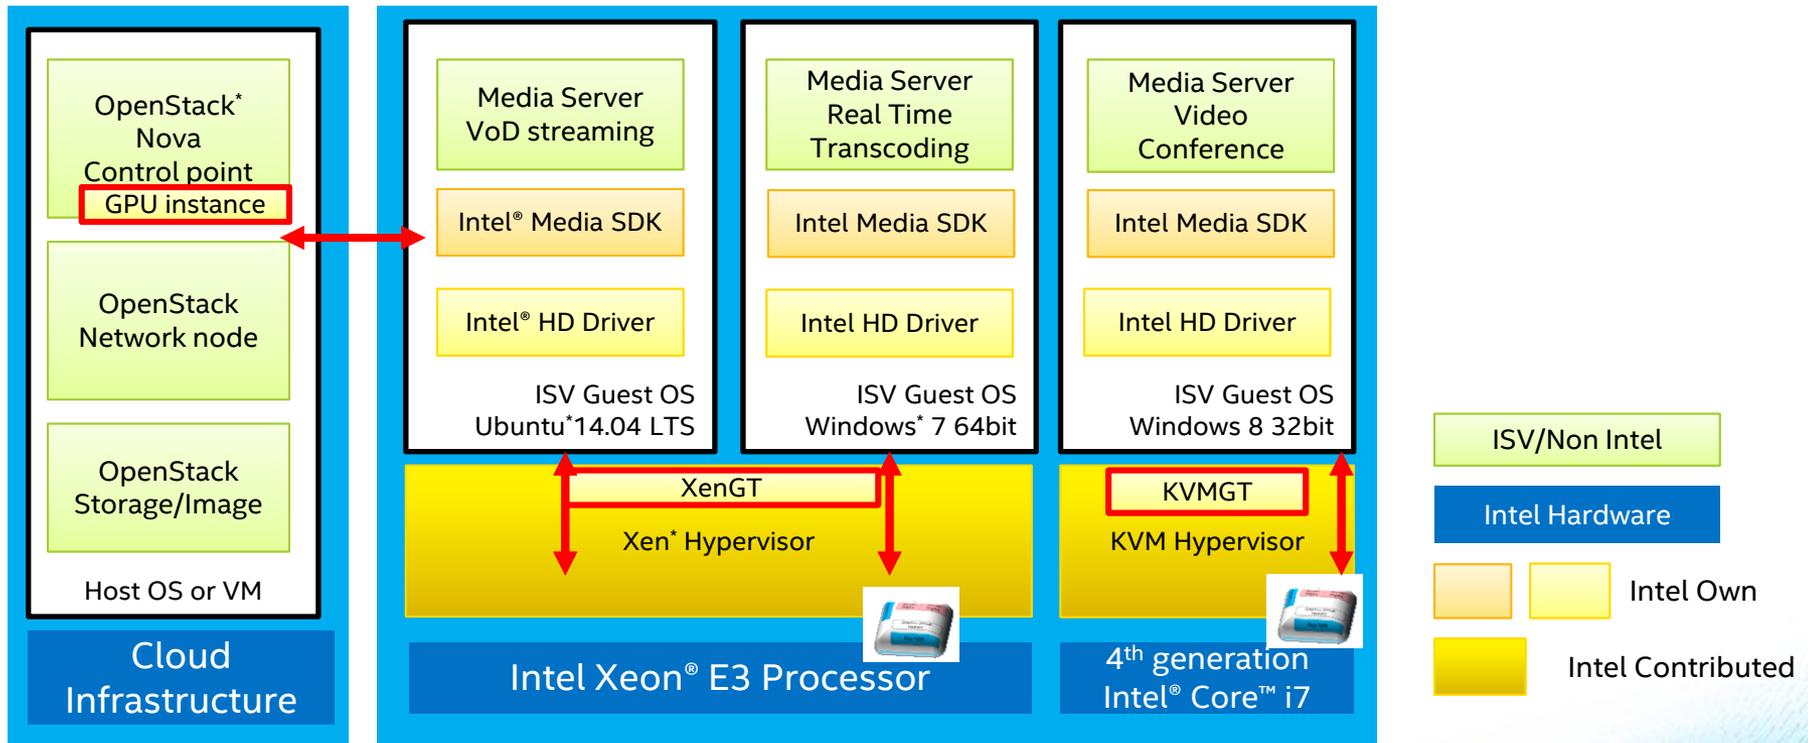
DirectX 11.1, OpenGL 4.2, OpenCL  
1.2, MediaSDK x.x

## Sharing

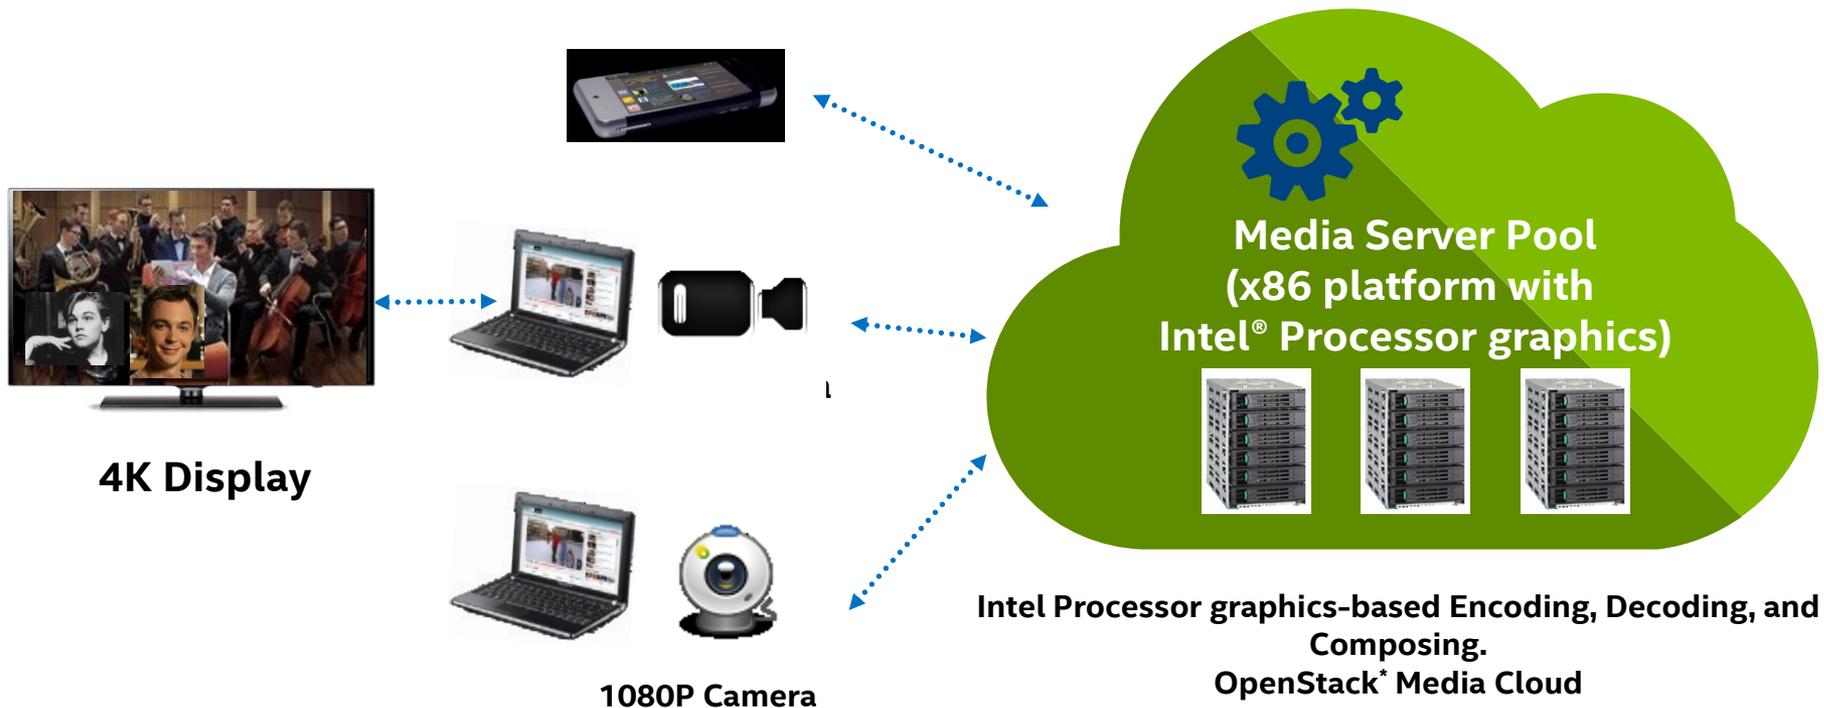
**Simultaneously**  
accelerate multiple  
VMs



# Building Blocks

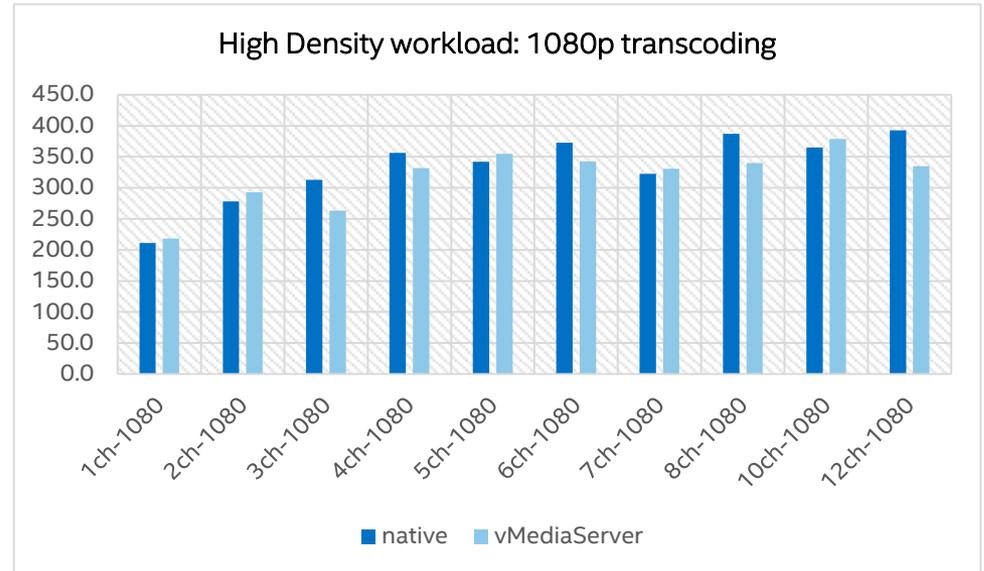
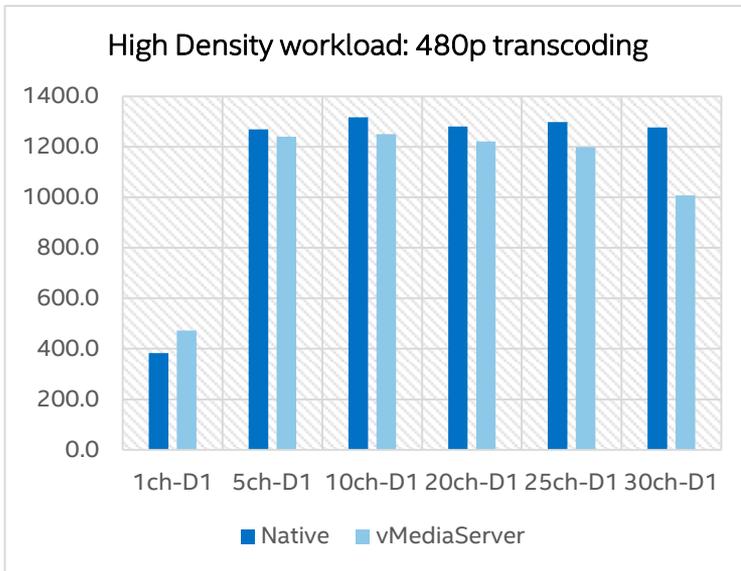


# Video Conference Usage Case



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

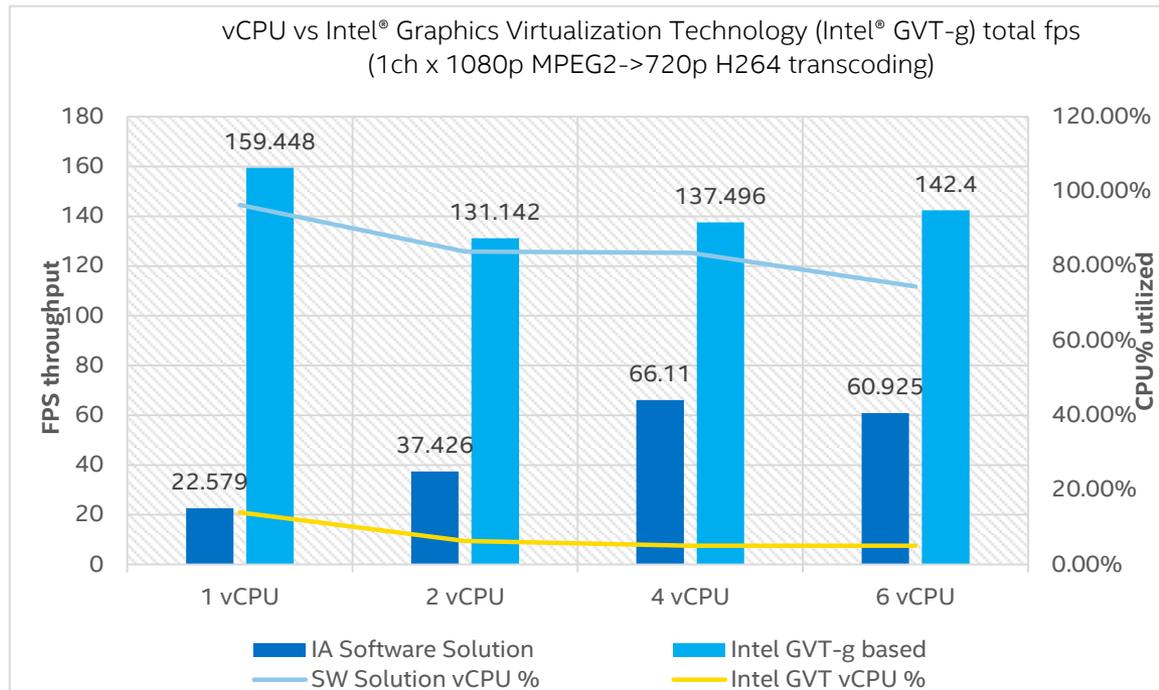
# Performance Summary



Config: I7 4770, Guest Ubuntu<sup>®</sup> 14.04LTS, 4GB mem, 1.5G GraphicMem, MediaSDK

Avg S90% of native H.264 transcoding performance

# CPU Transcoding vs. GPU Transcoding



Performance boost  
with Intel GVT-g!

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

# Q&A

