

XenGT: 一款工作在英特尔® 处理器显卡上的高性能图形虚拟化解决方案

Kevin Tian, Software Architect, Intel Corporation
David Cowperthwaite, Software Architect, Intel Corporation

SFTS008

议程

- 为什么需要GPU虚拟化?
- 完全GPU虚拟化之路
- 架构综述
- 关键技术
- 总结

议程

- 为什么需要GPU虚拟化?
- 完全GPU虚拟化之路
- 架构综述
- 关键技术
- 总结

GPU应用案例




3D图形

媒体

计算



虚拟化应用案例



应用案例

虚拟数据中心

云计算

虚拟远程桌面

富客户端虚拟化

自带设备

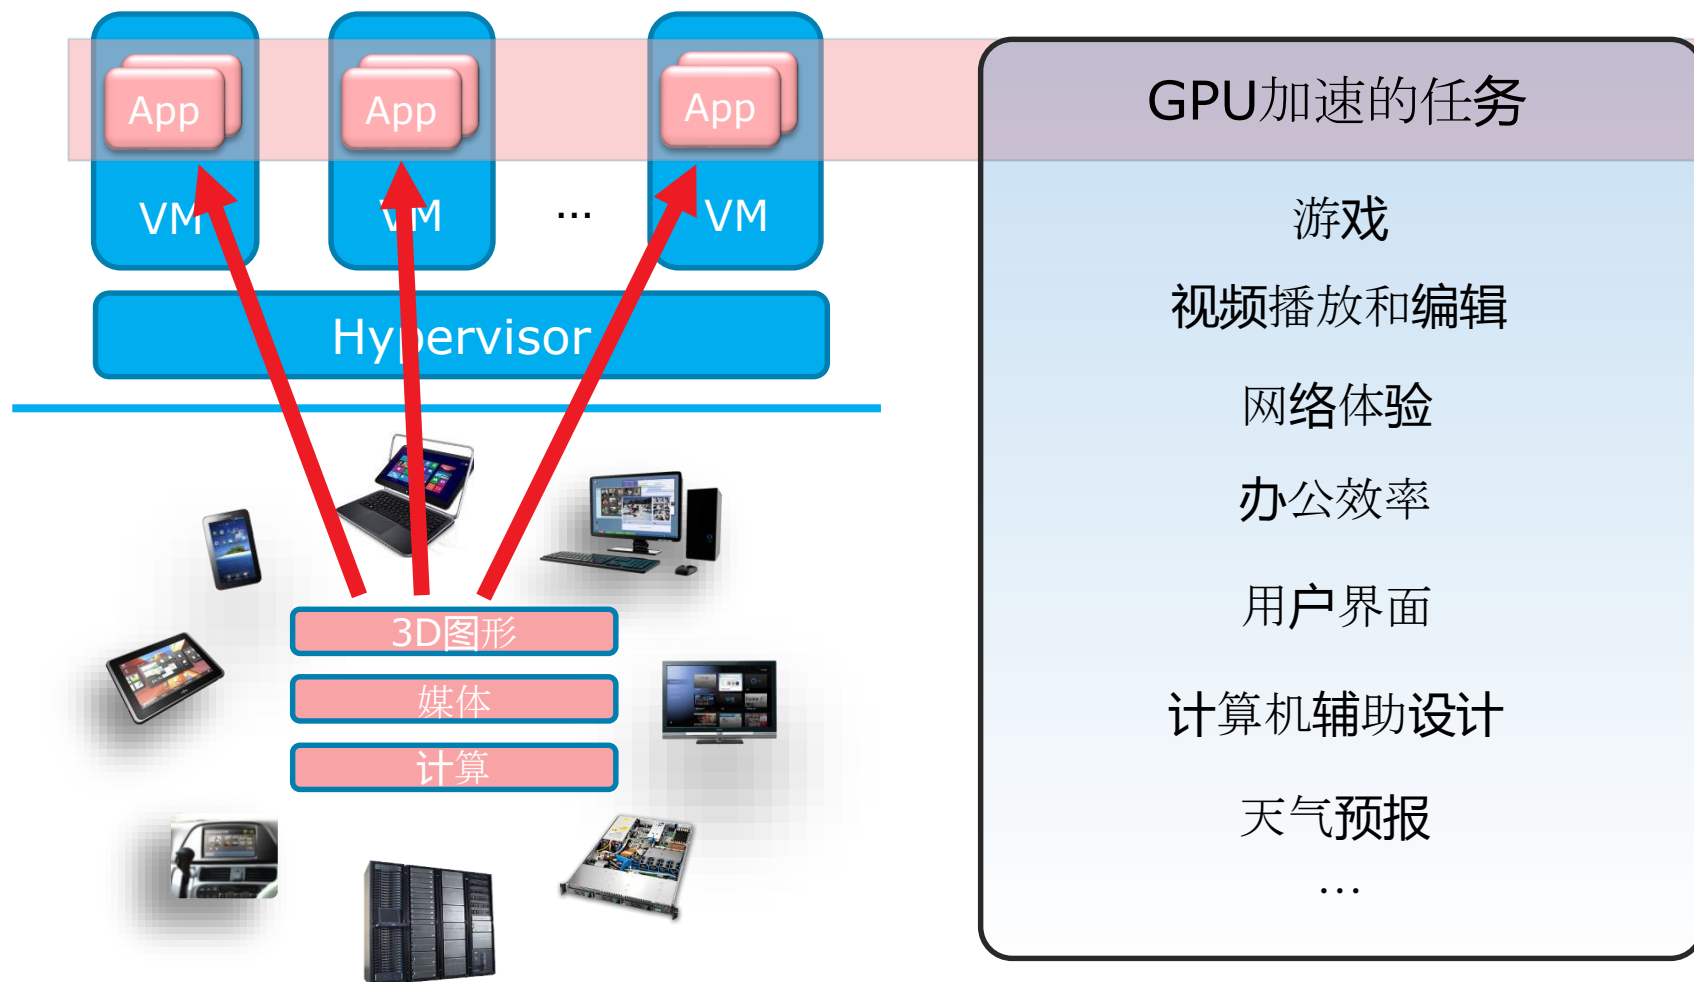
智能电视

多屏信息娱乐

安全电子商务


...

所以……GPU虚拟化



GPU虚拟化成为了一种基本需求

Intel® 处理器显卡



22nm

集成在第四代Intel® Core™ 处理器



GT3

新的GT3带来2倍的shader处理能力
提升- **Intel® Iris™ Graphics**



EDRAM

GT3e使用128MB快速缓存节省带宽-
Intel® Iris™ Pro Graphics



Intel® Quick Sync Video

H.264/MPEG-4 AVC, VC-1 高速视
频编解码

IDF14

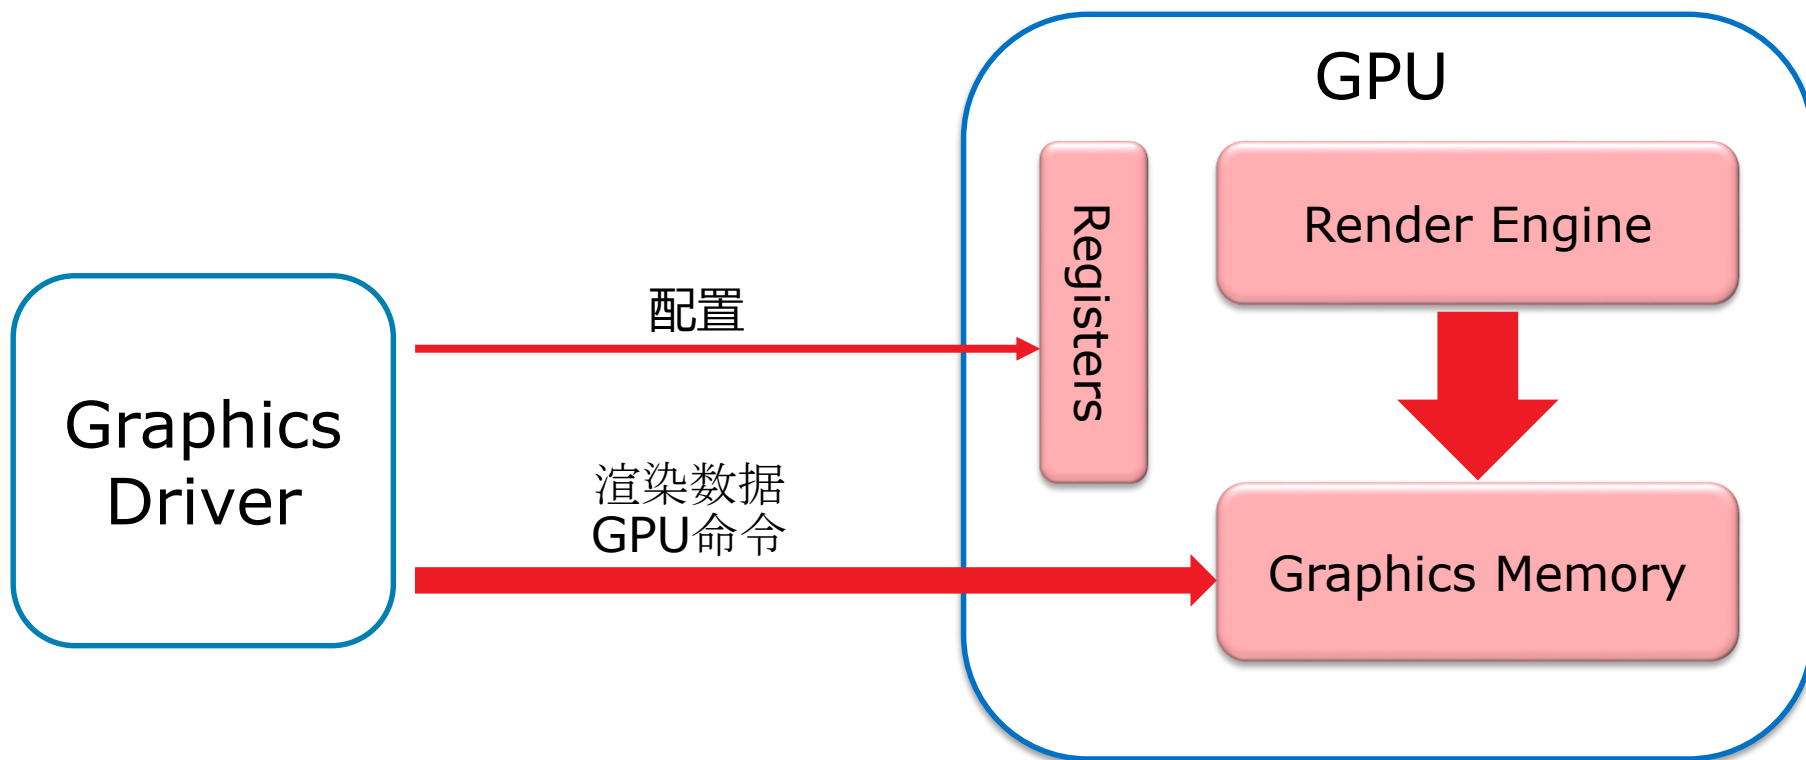
议程

- 为什么需要GPU虚拟化?
- 完全GPU虚拟化之路
- 架构综述
- 关键技术
- 总结

传统显卡

- 显示是唯一目的
- 只提供简单功能
 - 少数寄存器
 - 显存容量小
- 传统显卡的模拟很简单
 - 大多数设备模拟器提供的基本功能

现代GPU



模拟一个渲染引擎是不现实的！

GPU虚拟化的要求



性能



GPU直接加速



功能



一致的视觉体验

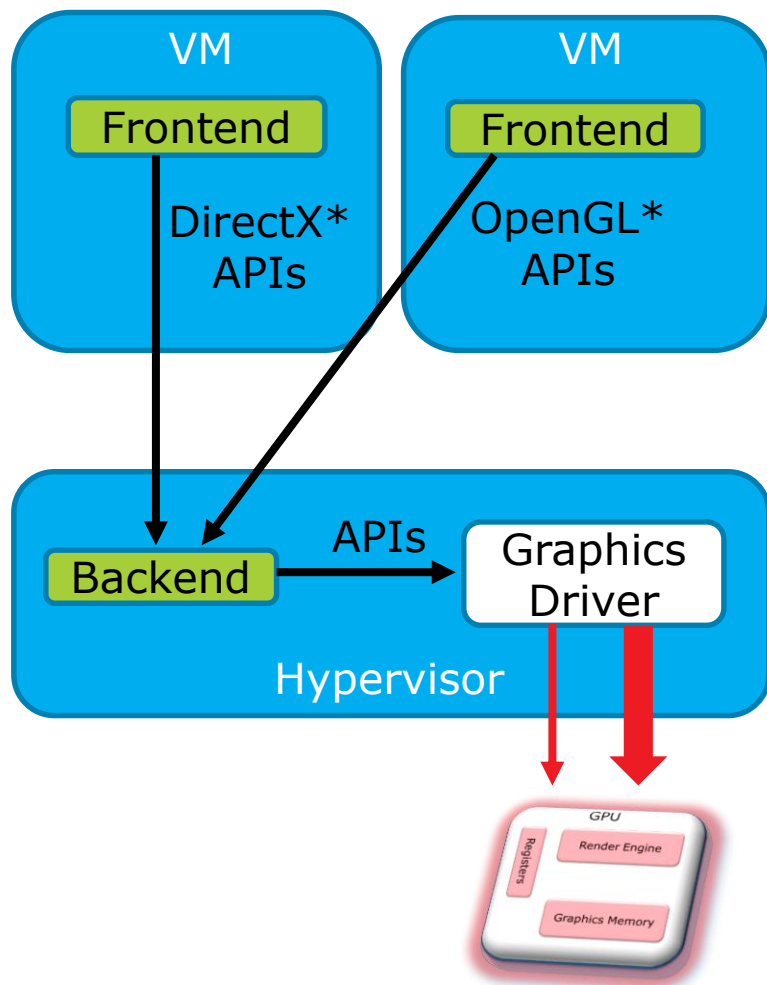


共享



同时加速多个虚拟机

API转发



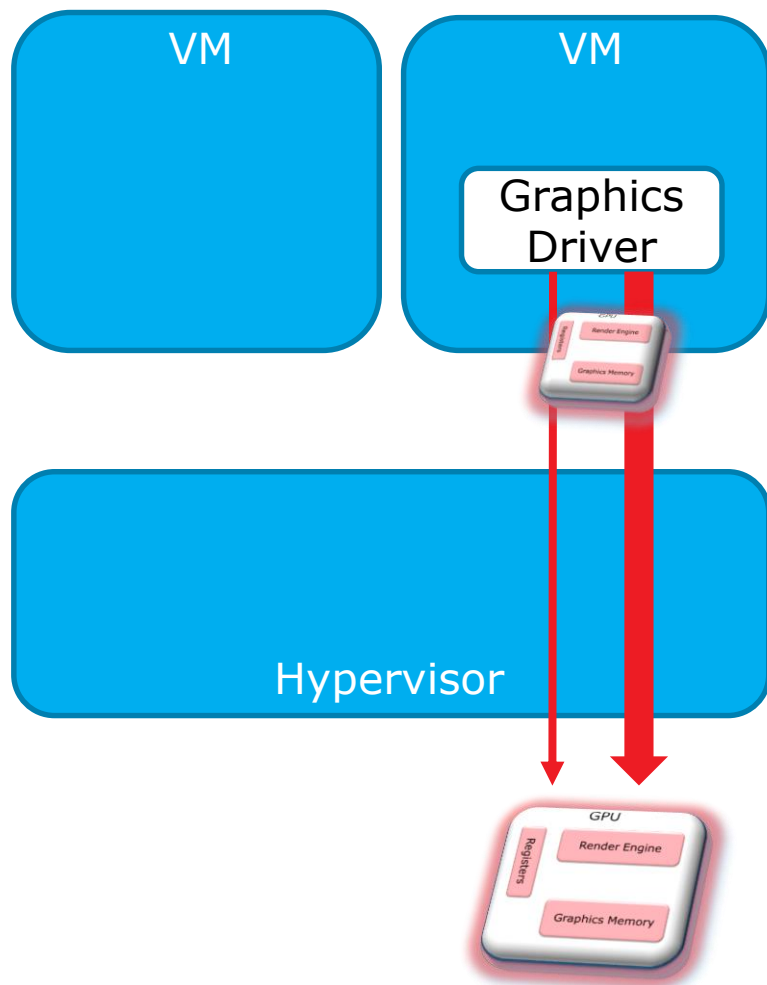
优点

- 性能佳
- 可以共享

缺点

- 功能滞后

直通设备



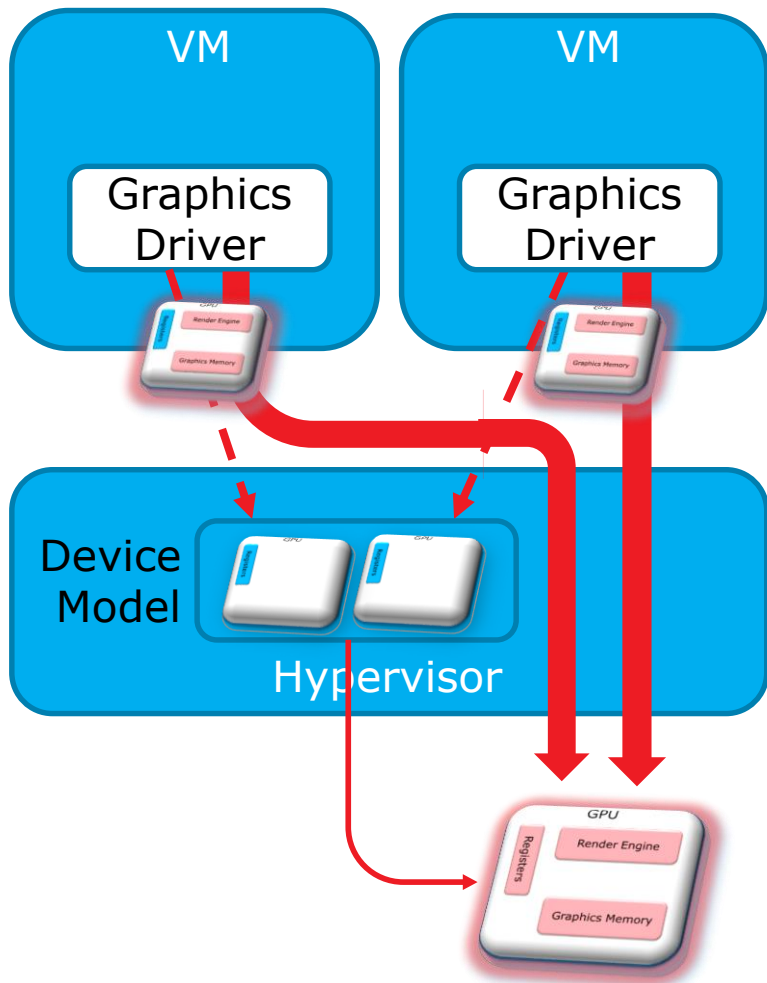
优点

- 性能佳
- 功能完备

缺点

- 不能共享

完全GPU虚拟化



优点

- 性能佳
- 功能完备
- 可以共享

可以在虚拟机里直接运行native图形栈！

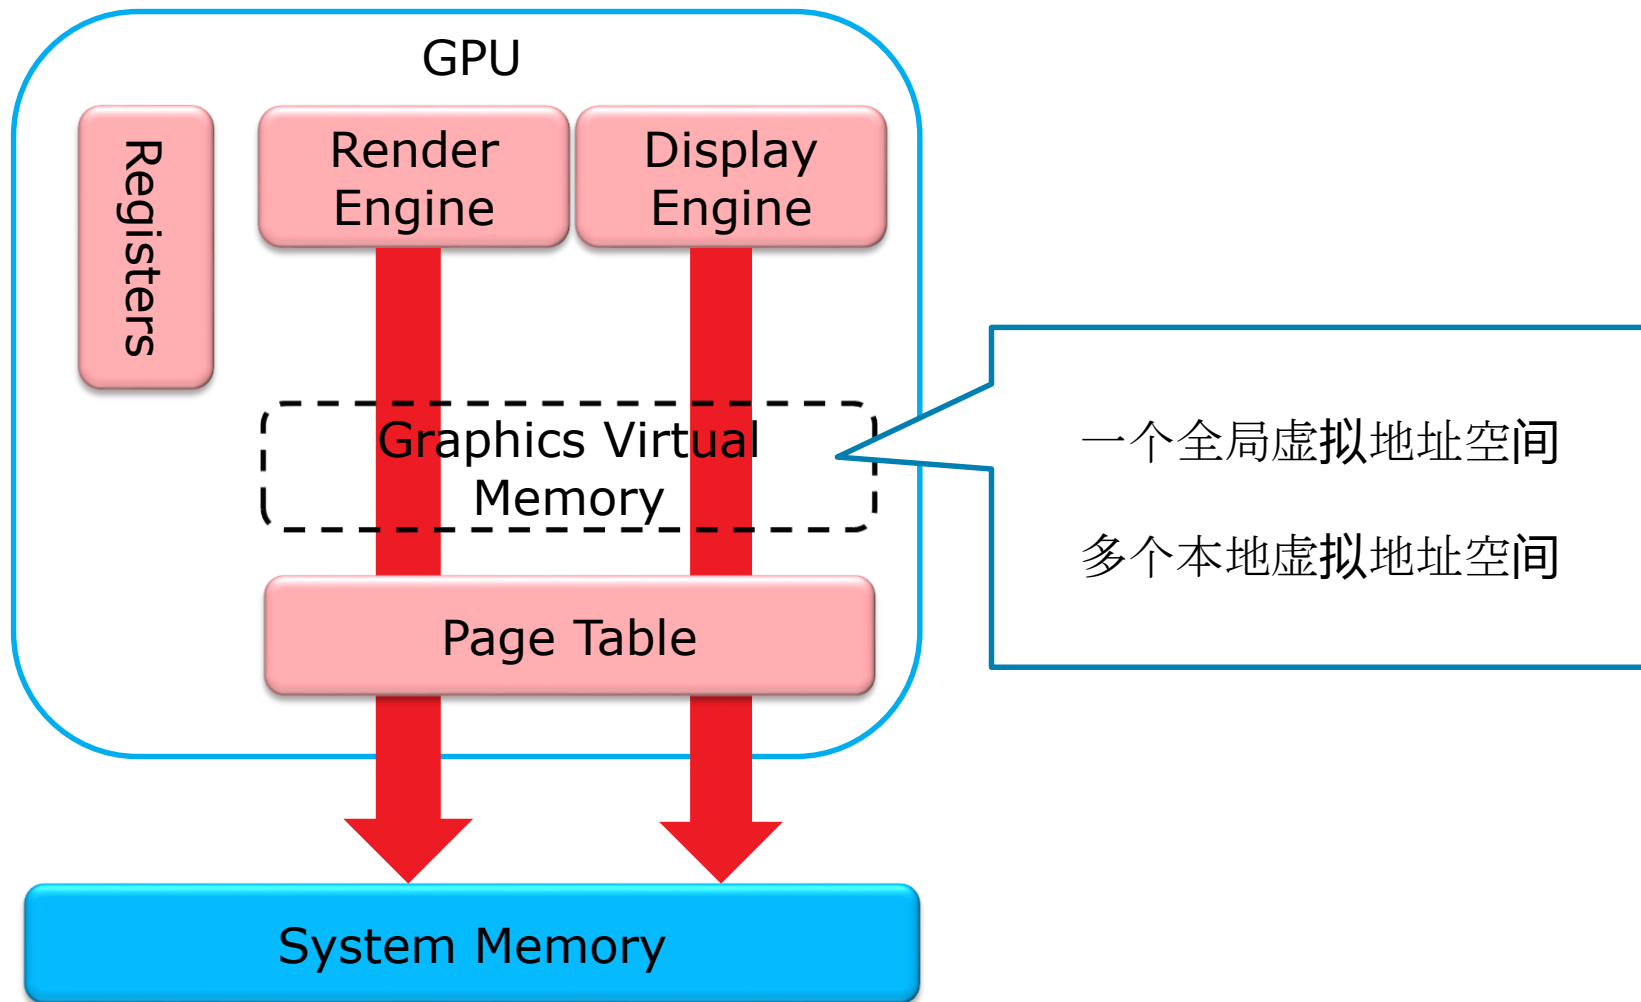
议程

- 为什么需要GPU虚拟化?
- 完全GPU虚拟化之路
- 架构综述
- 关键技术
- 总结

XenGT: 一款完全GPU虚拟化解决方案


- 基于mediated pass-through框架
 - 陷入和模拟特权I/O操作
 - 直通性能关键操作
- 虚拟GPU (vGPU)模拟器
 - 具备和物理Intel®处理器显卡一致的功能
- 在虚拟机中直接运行native图形驱动
 - 沿用驱动中已有的优化和稳定性补丁
- 第一版实现基于Xen Hypervisor
 - 核心设备模拟器可以重用在其他hypervisor中

处理器显卡：组成部件




虚拟化策略

直通



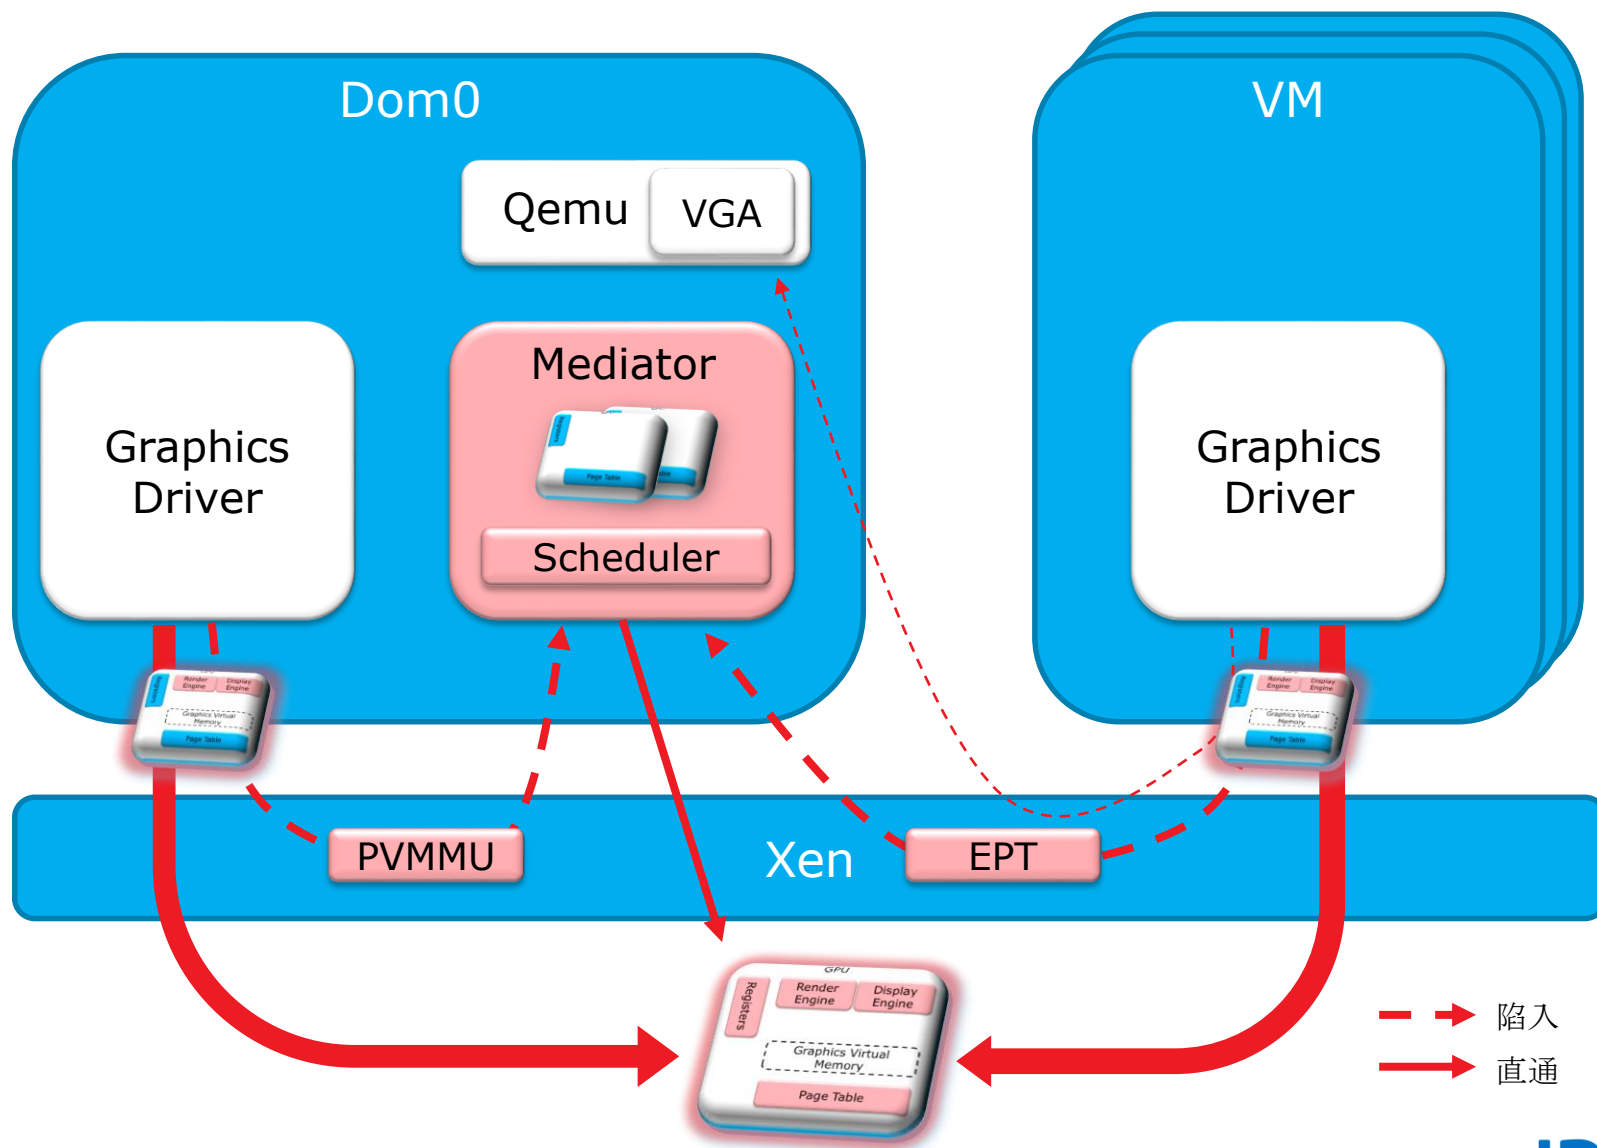
- 图像缓存
 - 命令缓存
- (都位于图形内存中)

陷入和模拟



- 内存映射I/O寄存器
- 端口I/O寄存器
- PCI配置空间寄存器
- GPU页表项


总体架构



能力

性能


通过GPU直接加速来达到接近native的性能



Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark* and MobileMark*, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more information go to <http://www.intel.com/performance>.


功能

运行native图形栈来保证一致的用户视觉体验



共享

同时加速3个虚拟机+Dom0

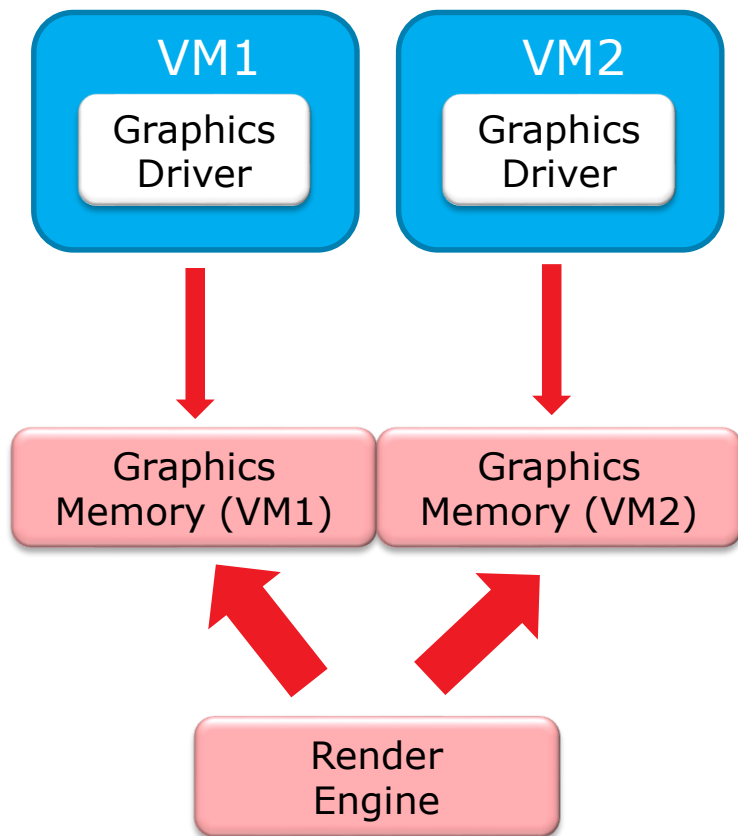


IDF14

议程

- 为什么需要GPU虚拟化?
- 完全GPU虚拟化之路
- 架构综述
- 关键技术
- 总结

直通：图形内存



- 来自于CPU和GPU的并行访问
 - CPU/GPU独立调度
- 图形内存分区
 - 每个VM分配专有的资源
- 本地图形内存完全直通

直通：图形地址空间

图形内存分区引起的不一致观点

虚拟机1的观点 

虚拟机2的观点 



宿主机的观点  

(虚拟机的观点)

- 资源变少
- 从地址0开始

地址空间ballooning带来一致观点

虚拟机1的观点  

虚拟机2的观点  


宿主机的观点  

(虚拟机的观点)

- 资源大小不变
- 从任意地址开始
- 其他虚拟机的资源被保留

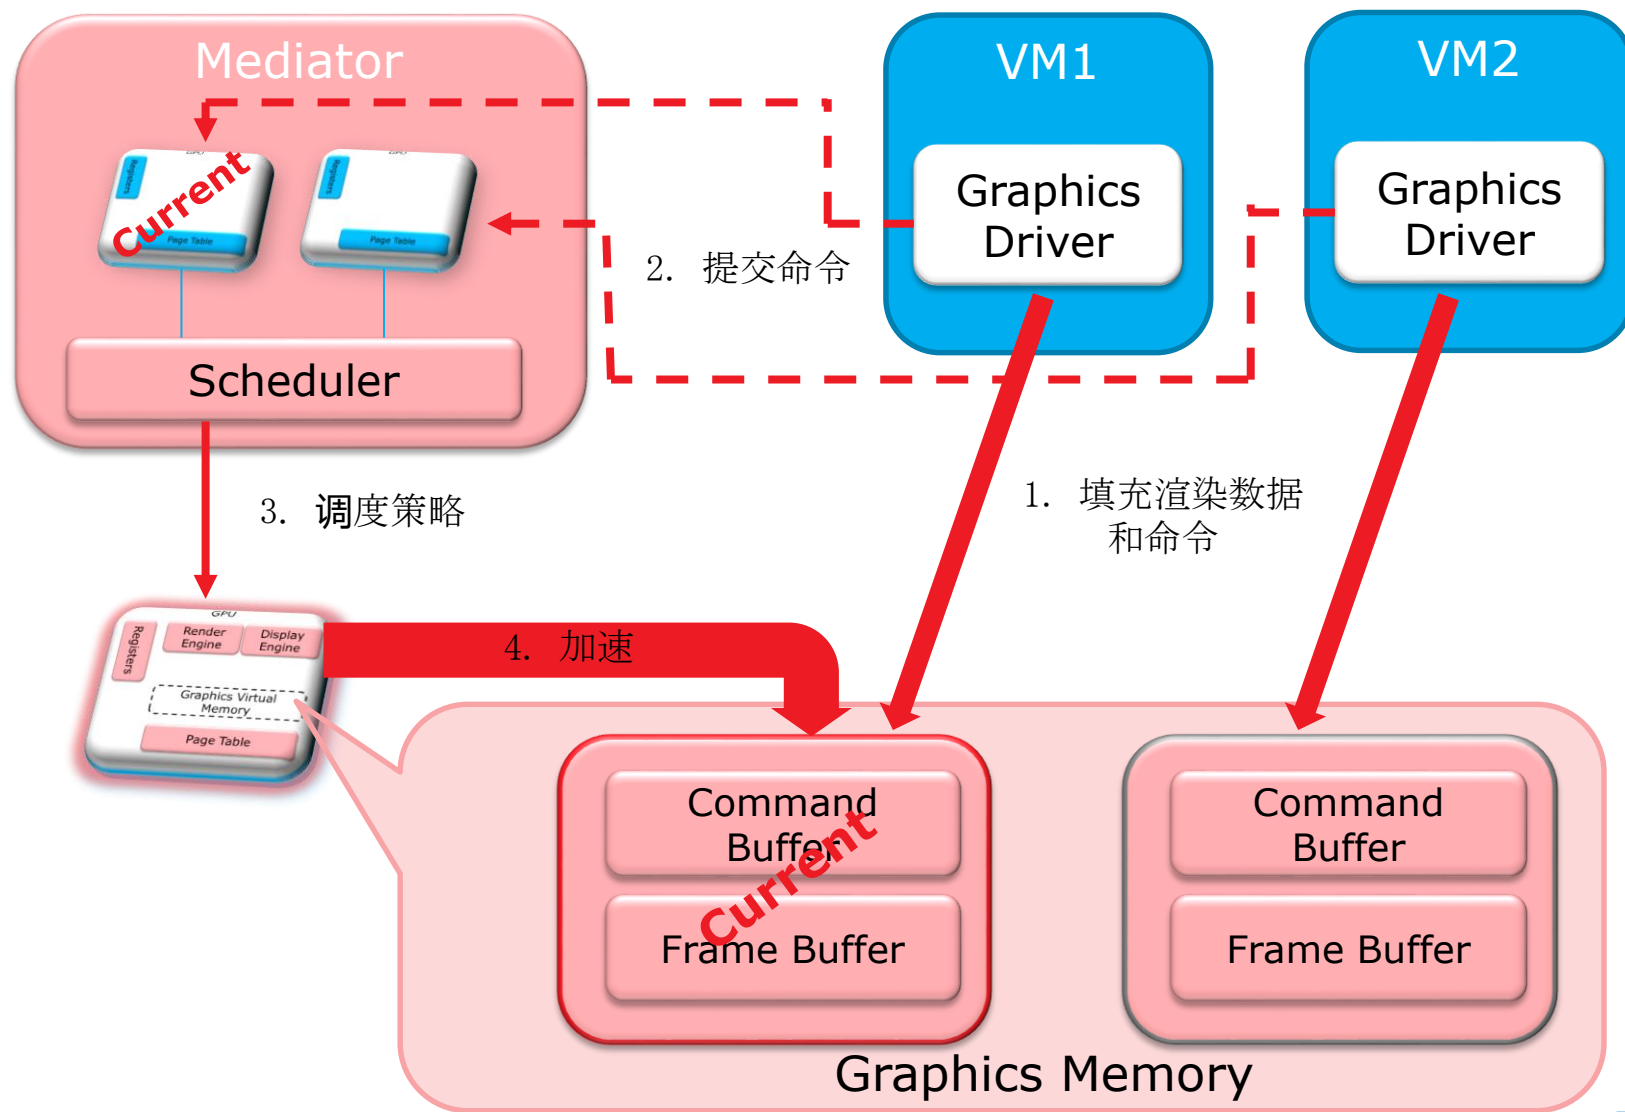
避免了
地址转换！

GPU页表虚拟化





- GPU页表
 - 将系统内存映射进虚拟地址空间
- 影子页表
 - 客户页框号 \leftrightarrow 主机页框号

GPU调度程序



显示共享



议程

- 为什么需要GPU虚拟化?
- 完全GPU虚拟化之路
- 架构综述
- 关键技术
- 总结

总结

- 在许多虚拟化应用中，GPU虚拟化正在成为一个基本需求
- 完全GPU虚拟化在性能、功能和共享中达到一个很好的平衡
- XenGT是一款工作在英特尔® 处理器显卡上的完全GPU虚拟化解决方案，可以支持在虚拟机里运行native图形驱动

下一步

- 开源项目：尝试并反馈您的意见
 - <https://github.com/01org/XenGT-Preview-kernel>
 - <https://github.com/01org/XenGT-Preview-xen>
 - <https://github.com/01org/XenGT-Preview-qemu>

本课程演示文稿（PDF）发布在技术课程目录网站：
www.intel.com/idfsessionsSZ。该网址同时列于会议指南手册中专题讲
座日程页的上方。

请填写在线课程反馈表

即刻登录，惊喜大奖等你拿！

在本课程结束前，您将会收到一封电子邮件，内含在线课程反馈表的网址。请在明天上午10点前提交在线反馈，以便获得抽奖资格。

获奖者将会通过电子邮件发布

抽奖规则 可在四层服务台索取
所有在线课程反馈必须在 4 月 4 日星期五下午 5 点前提交

问答

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.

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 the important factors that could cause actual results to differ materially from the company’s expectations. Demand could be different from Intel’s expectations due to factors including changes in business and economic conditions; customer acceptance of Intel’s and competitors’ products; supply constraints and other disruptions affecting customers; changes in customer order patterns including order cancellations; and changes in the level of inventory at customers. Uncertainty in global economic and financial conditions poses a risk that consumers and businesses may defer purchases in response to negative financial events, which could negatively affect product demand and other related matters. Intel operates in intensely competitive industries that are characterized by a high percentage of costs that are fixed or difficult to reduce in the short term and product demand that is highly variable and difficult to forecast. Revenue and the gross margin percentage are affected by the timing of Intel product introductions and the demand for and market acceptance of Intel’s products; actions taken by Intel’s competitors, including product offerings and introductions, marketing programs and pricing pressures and Intel’s response to such actions; and Intel’s ability to respond quickly to technological developments and to incorporate new features into its products. The 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; start-up costs; excess or obsolete inventory; changes in unit costs; defects or disruptions in the supply of materials or resources; product manufacturing quality/yields; and impairments of long-lived assets, including manufacturing, assembly/test and intangible assets. 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. Expenses, particularly certain marketing and compensation expenses, as well as restructuring and asset impairment charges, vary depending on the level of demand for Intel’s products and the level of revenue and profits. Intel’s results could be affected by the timing of closing of acquisitions and divestitures. Intel’s results could be affected by adverse effects associated with product defects and errata (deviations from published specifications), and by litigation or regulatory matters involving intellectual property, stockholder, consumer, antitrust, disclosure and other issues, such as the litigation and regulatory matters described in Intel’s SEC reports. 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. 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.

Rev. 1/16/14