

# Sorting out the Confusion of Multicore Processor Virtualization Features

Jan 2012

Rajan Goyal (Cavium)


Bryan Chin (Cavium)

Surender Kumar (NSN)

# Agenda

- Virtualization Introduction
- Working Group Charter
  - Software Virtualization Strategy (SVS)
  - Multicore Virtualization Profiles (MVP)
- Current Progress
  - Proposal for MVP
- Contact Info

# Agenda

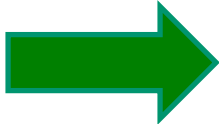


- Virtualization Introduction
- Working Group Charter
  - Software Virtualization Strategy (SVS)
  - Multicore Virtualization Profiles (MVP)
- Proposal for MVP
- Discussion
  - Next Steps

# Virtualization Introduction

- Mainstream feature in server based computing (cloud, enterprise)
- Migrating to embedded computing
  - Support for legacy operating systems and applications
  - Isolation and security
  - Many more requirements for embedded (lots of variety in I/O)

# Agenda

- Virtualization Introduction
-  Working Group Charter
  - Software Virtualization Strategy (SVS)
  - Multicore Virtualization Profiles (MVP)
- Proposal for MVP
- Discussion
  - Next Steps

# Work Items Proposal Background

- Current software architecture approach/trend for multicore applications:
  - Linux or equivalent – Control Plane
  - RTOS/Bare Metal OS – Data Plane
- Hence, hypervisor and/or virtualization required to:
  - Mediate
  - Isolate
  - Inter domain communication/sharing
- Issues:
  - Higher Complexity
  - Higher Development cost

# Charter - SVS

- Newer Trend or Goal
  - Single environment (Linux/FreeBSD) with support for both data plane and control plane.
  - Low overhead at run time.
  - Flexibility / Capability of single OS environment (e.g. Linux)
- Hence, need for virtualization is different/evolving.
- Charter # 1
  - Enable transition in current software paradigm to newer architecture by providing standard virtualization techniques/guidelines.

# Charter - MVP (First Milestone)

- Come up with different grading of virtualization hooks in multicore SoCs based on different market/application space. For example for:
  - Networking
  - Mobile Infrastructure
  - Data Center
  - Cloud
  - Others?
- Different grading/profile could be:
  - Level 1 (Basic)
  - Level II (Intermediate)
  - Level III (Advanced)
- Working Group will generate detailed specification for different grading and mapping multicore SOC features to various grades.

# Charter - MVP (contd...)

- MultiCore SoC features will be in following categories:
  - Core virtualization
  - I/O virtualization
  - H/W accelerator virtualization
  - Others
- Finally, standardize various profiles and get it accepted by system companies viz. Cisco / Juniper / NSN / Huawei / etc. and processor companies (MIPS/ARM/PowerPC).

# Charter - MVP

- Processor vendors can make claims as to what grade of virtualization their chips support for the different features based on the specification.
- RFQ:
  - system vendors make it as part of RFQ and chip vendors reply with certified grade for their individual chips.
- What it achieves:
  - Today - vendors claim some support for virtualization in SoCs and it is confusing.
  - help system companies chose right part for their needs and it will become standard.

# Charter - MVP (Next Step)

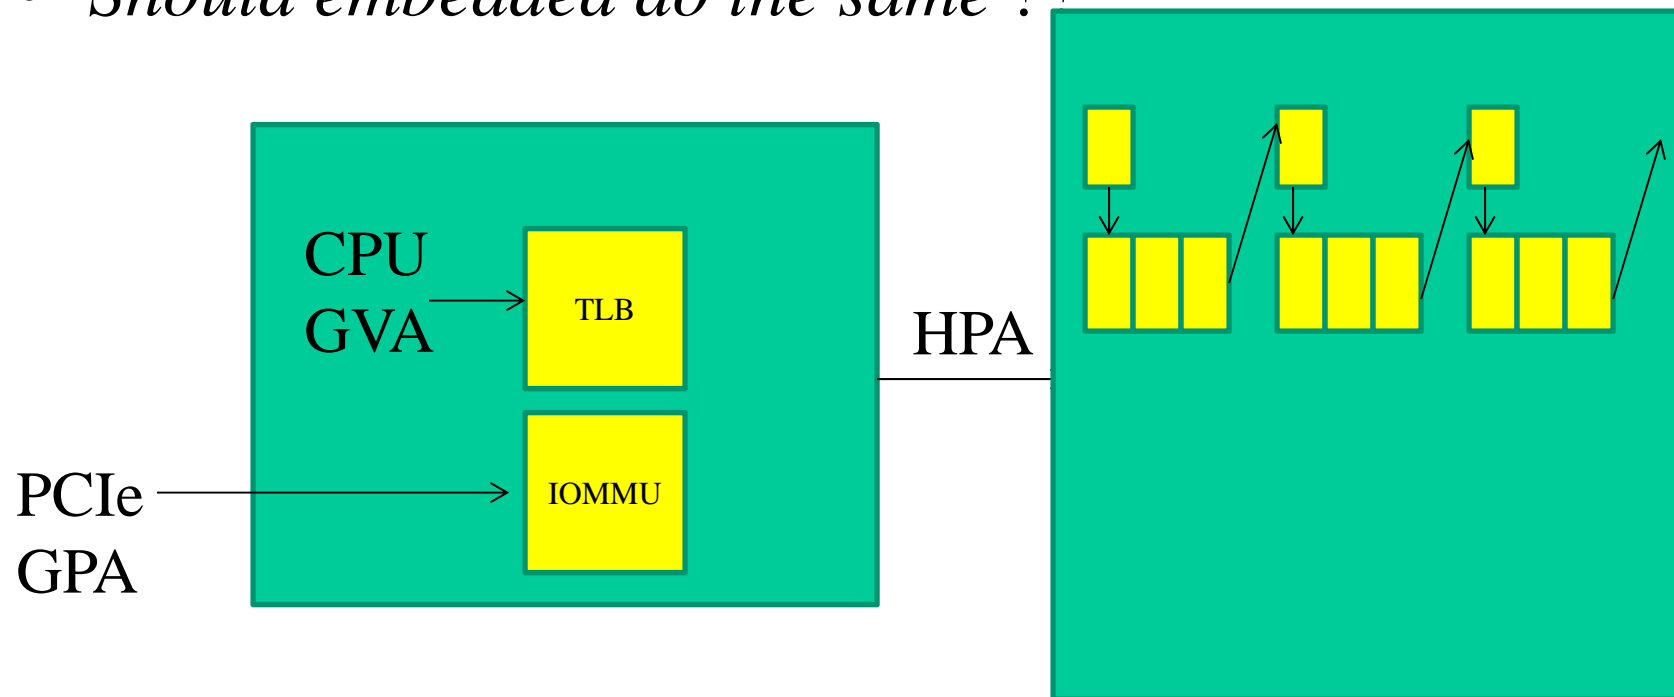
- WG works with another independent association (e.g EEMBC) to facilitate certification of SoCs to particular virtualization grade, based on features a particular processor is implementing.
  - i.e. either chip vendors submit their HRMs/data sheet to technical committee of WG AND/OR WG (along with EEMBC) develops benchmark suites to qualify particular features.
- Benchmark suites will be to validate feature capabilities and performance. (TBD) Based on results, committee will certify a chip to a particular grade.

# Agenda

- Virtualization Introduction
- Working Group Charter
  - Software Virtualization Strategy (SVS)
  - Multicore Virtualization Profiles (MVP)
-  Proposal for MVP
- Contact Info

# Example: Desktop and Server

- Addressing scheme defined by AMD and Intel
- Nested Page Tables
- IOMMU/IOTLB
- Hardware Page table walker
- *Should embedded do the same ??*



# Application Classes

- General Purpose Computing
  - Enterprise Servers
  - Desktop Computers
  - Cloud Computers
- High End Embedded
  - Networking/Communications
    - Wireline Networking SoCs
    - Wireless SoCs
  - Storage
  - Imaging
- Consumer
  - Portable devices
  - Non-portable devices

# H/W Virtualization Features

- Memory Management
  - Memory partitioning schemes with para-virtualized addressing
  - Block translation schemes
  - Nested Page Table
  - IO Address Translation
- I/O
  - PCIe SR-IOV
  - “physicalization”
  - Endpoint ATC (Address Translation Cache)
  - Interrupts direct to Guest
- H/W assist (co-processor) Virtualization
  - Packet Processing Block
  - Compression Block
  - Regex Block
  - Crypto Block
  - etc.
- Performance Features
  - Virtualization contexts – scale UP focused
  - Virtualization contexts – scale out focused

# Profiles Requirements/Features

- Multiple profiles to address different markets

	Consumer	Hi End Embedded	General Purpose Computing
Simple Memory Protection	Ra		
Nested Page Tables	Ra	R	R
IO Address Translation	O	R	R
PCIe SR-IOV	O	R	R
Interrupts to Guest	O	R	O
Scale UP Virtualization	Rb	O	R
Scale Out Virtualization	Rb	R	O

Diagram illustrating requirements (R) and optional features (O) for three profiles across different market segments:

- Profile I** (Consumer): Includes Simple Memory Protection (Ra), Nested Page Tables (Ra), Scale UP Virtualization (Rb), and Scale Out Virtualization (Rb).
- Profile II** (Hi End Embedded): Includes Nested Page Tables (R), IO Address Translation (R), PCIe SR-IOV (R), Interrupts to Guest (R), and Scale Out Virtualization (R).
- Profile III** (General Purpose Computing): Includes Nested Page Tables (R), IO Address Translation (R), PCIe SR-IOV (R), Scale UP Virtualization (R), and Scale Out Virtualization (O).

O – optional; R required; Rx – x identifies alternates

# Agenda

- Virtualization Introduction
- Working Group Charter
  - Software Virtualization Strategy (SVS)
  - Multicore Virtualization Profiles (MVP)
- Proposal for MVP
-  Contact Info

# Contact Information

- MCA Virtualization Working Group  
URL:  
<http://www.multicore-association.org/home.php>
- Working Group Mailing Alias:  
[virtualization@multicore-association.org](mailto:virtualization@multicore-association.org)
- Chairpersons:  
Rajan Goyal ([rgoyal@cavium.com](mailto:rgoyal@cavium.com))  
Surender Kumar ([surender.reddy@nsn.com](mailto:surender.reddy@nsn.com))