# Q35 - QEMU

Marcel Apfelbaum
August, 2016

# Agenda

- **Q35 chipset - Overview**
- Q35 emulation
- Q35-only features
- Q35 limitations
- Q35 use cases

# Q35 chipset - Overview

- Data sheet

| Essentials | |
|---|---|
| Status | Launched |
| Launch Date | Q3'07 |
| Supported FSBs | 1333MHz / 1066MHz / 800MHz |
| FSB Parity | No |
| Embedded Options Available | Yes |
| TDP | 15 W |
| Recommended Customer Price | N/A |
| Datasheet | Link |

| Memory Specifications | |
|---|---|
| Max Memory Size (dependent on memory type) | 8 GB |
| Memory Types | DDR2 667/800 |
| Max # of Memory Channels | 2 |
| Max Memory Bandwidth | 6.4 GB/s |
| Physical Address Extensions | 36-bit |

| Expansion Options | |
|---|---|
| PCI Express Revision | 1.1 |
| PCI Express Configurations ‡ | 1x16 |

| Advanced Technologies | |
|---|---|
| Intel® Virtualization Technology for Directed I/O (VT-d) ‡ | Yes |
| Intel® Fast Memory Access | No |
| Intel® Flex Memory Access | Yes |

| Package Specifications | |
|---|---|
| Max CPU Configuration | 1 |
| $T_{CASE}$ | 106°C |
| Package Size | 34mm x 34mm |
| Low Halogen Options Available | See MDDS |

# Q35 chipset - Overview

- Topology
  - North Bridge: MCH
  - South Bridge: ICH9

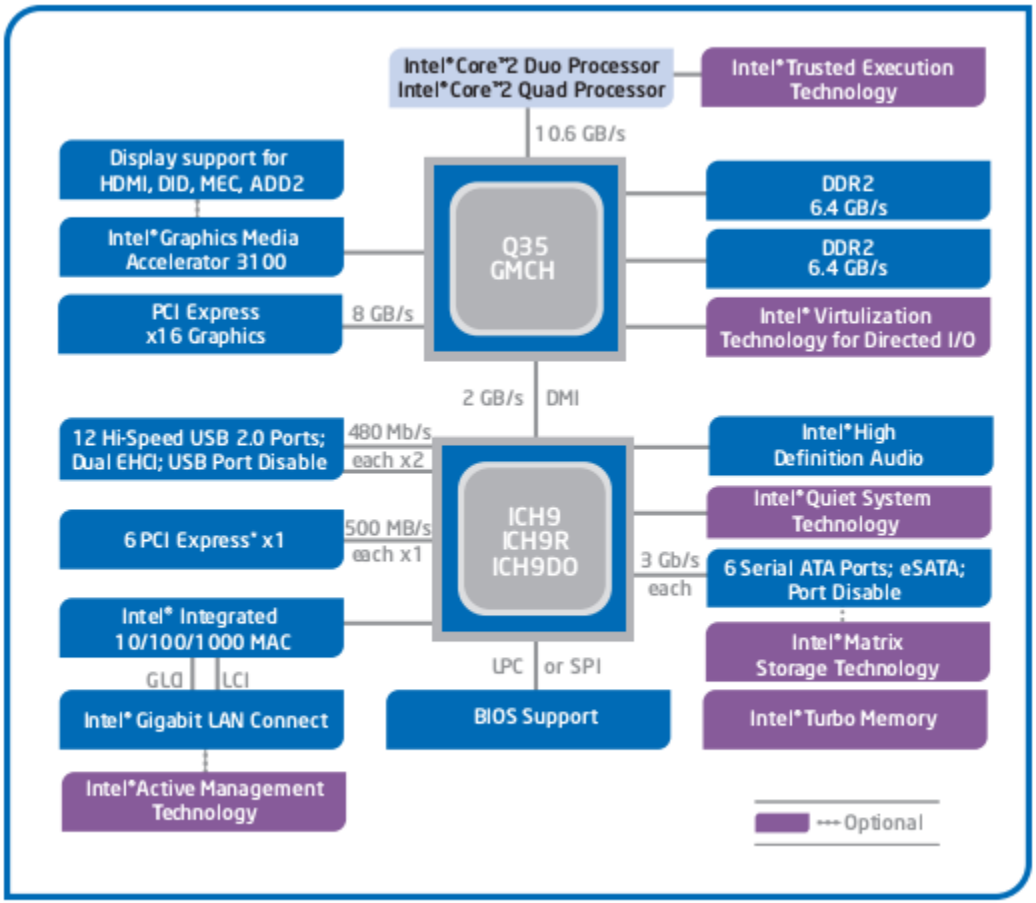


Intel® Q35 Express Chipset Block Diagram

# Agenda

- Q35 chipset - Overview
- **Q35 emulation**
- Q35-only features
- Q35 limitations
- Q35 use cases

# Q35 emulation

- Info qtree

```
Q35:
  dev: hpet, id ""
  dev: kvm-ioapic, id ""
  dev: q35-pcihost, id ""
    MCFG = 0xb0000000
    pci-hole64-size = 0 B
    below-4g-mem-size = 2 GiB
    above-4g-mem-size = 1 GiB
    bus: pcie.0
      dev: e1000, id ""
      dev: VGA, id ""
      dev: ICH9 SMB, id ""


      dev: ich9-ahci, id ""
        bus: ide.2
          dev: ide-cd, id ""
        bus: ide.0
          dev: ide-hd, id ""
      dev: ICH9-LPC, id ""
        class ISA bridge, addr 00:1f.0
        bus: isa.0
          dev: i8257, id ""
          dev: i8257, id ""
          dev: port92, id ""
          dev: vmmouse, id ""
          dev: vmport, id ""
          dev: i8042, id ""
          dev: isa-parallel, id ""
          dev: isa-serial, id ""
          dev: isa-pcspk, id ""
          dev: kvm-pit, id ""
          dev: mc146818rtc, id ""
          dev: kvm-i8259, id ""
          dev: kvm-i8259, id ""
      dev: mch, id ""
        addr = 00.0
        class Host bridge, addr 00:00.0
  dev: fw_cfg_io, id ""
  dev: kvmclock, id ""
  dev: kvmvapic, id ""
```

```
I440FX:
  dev: hpet, id ""
  dev: kvm-ioapic, id ""
  dev: i440FX-pcihost, id ""
    pci-hole64-size = 16 EiB



    bus: pci.0
      dev: e1000, id ""
      dev: VGA, id ""
      dev: PIIX4_PM, id ""
        bus: i2c
          type i2c-bus
          dev: smbus-eeprom, id ""
      dev: piix3-ide, id ""
        bus: ide.1
          dev: ide-cd, id ""
        bus: ide.0
          dev: ide-hd, id ""
      dev: PIIX3, id ""
        class ISA bridge, addr 00:01.0
        bus: isa.0
          dev: isa-fdc, id ""
          dev: i8257, id ""
          dev: i8257, id ""
          dev: port92, id ""
          dev: vmmouse, id ""
          dev: vmport, id ""
          dev: i8042, id ""
          dev: isa-parallel, id ""
          dev: isa-serial, id ""
          dev: isa-pcspk, id ""
          dev: kvm-pit, id ""
          dev: mc146818rtc, id ""
          dev: kvm-i8259, id ""
          dev: kvm-i8259, id ""
      dev: i440FX, id ""
        class Host bridge, addr 00:00.0
  dev: fw_cfg_io, id ""
  dev: kvmclock, id ""
  dev: kvmvapic, id ""
```

# Q35 emulation

- lspci

```
Q35
-------
00:00.0 Host bridge: Intel Corporation 82G33/G31/P35/P31 Express DRAM Controller
00:01.0 VGA compatible controller: Device 1234:1111 (rev 02)
00:02.0 Ethernet controller: Intel Corporation 82540EM Gigabit Ethernet Controller (rev 03)
00:1f.0 ISA bridge: Intel Corporation 82801IB (ICH9) LPC Interface Controller (rev 02)
00:1f.2 SATA controller: Intel Corporation 82801IR/IO/IH (ICH9R/DO/DH) 6 port SATA Controller [AHCI mode]
00:1f.3 SMBus: Intel Corporation 82801I (ICH9 Family) SMBus Controller (rev 02)

I440FX
------
00:00.0 Host bridge: Intel Corporation 440FX - 82441FX PMC [Natoma] (rev 02)
00:01.0 ISA bridge: Intel Corporation 82371SB PIIX3 ISA [Natoma/Triton II]
00:01.1 IDE interface: Intel Corporation 82371SB PIIX3 IDE [Natoma/Triton II]
00:01.3 Bridge: Intel Corporation 82371AB/EB/MB PIIX4 ACPI (rev 03)
00:02.0 VGA compatible controller: Device 1234:1111 (rev 02)
00:03.0 Ethernet controller: Intel Corporation 82540EM Gigabit Ethernet Controller (rev 03)
```

# Q35 emulation

- PCIe/PCI Topology

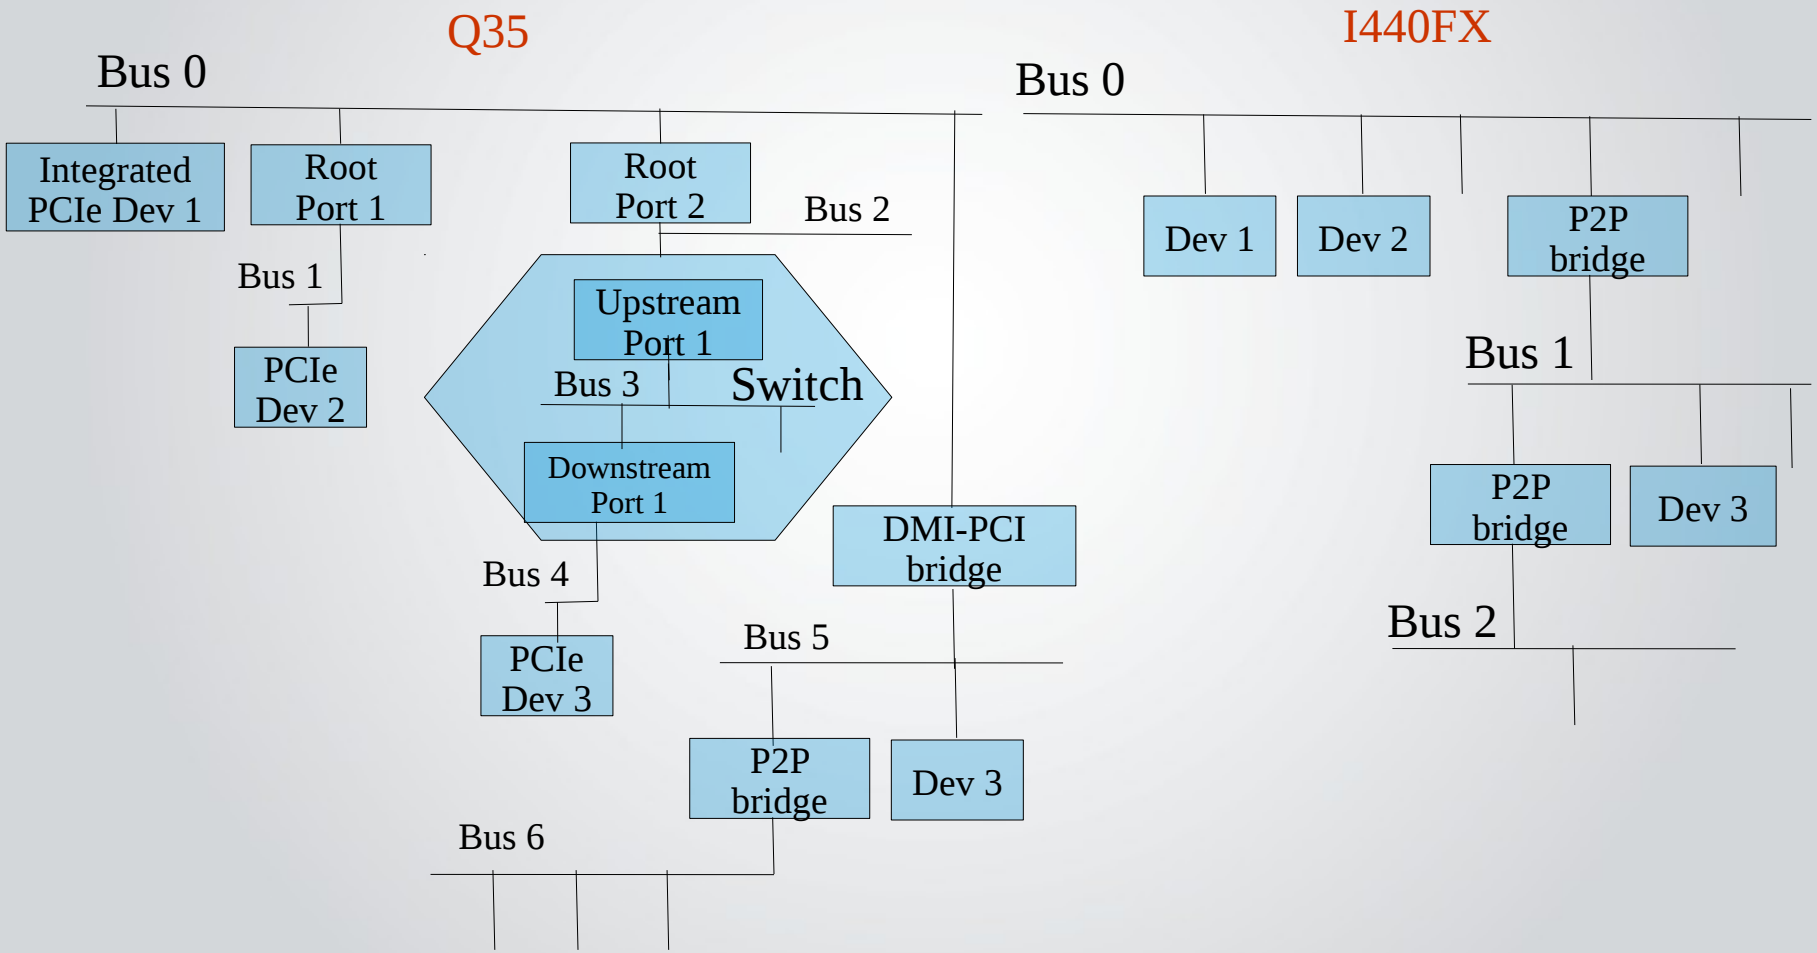

# Q35 emulation

- PCI/PCIe devices placement

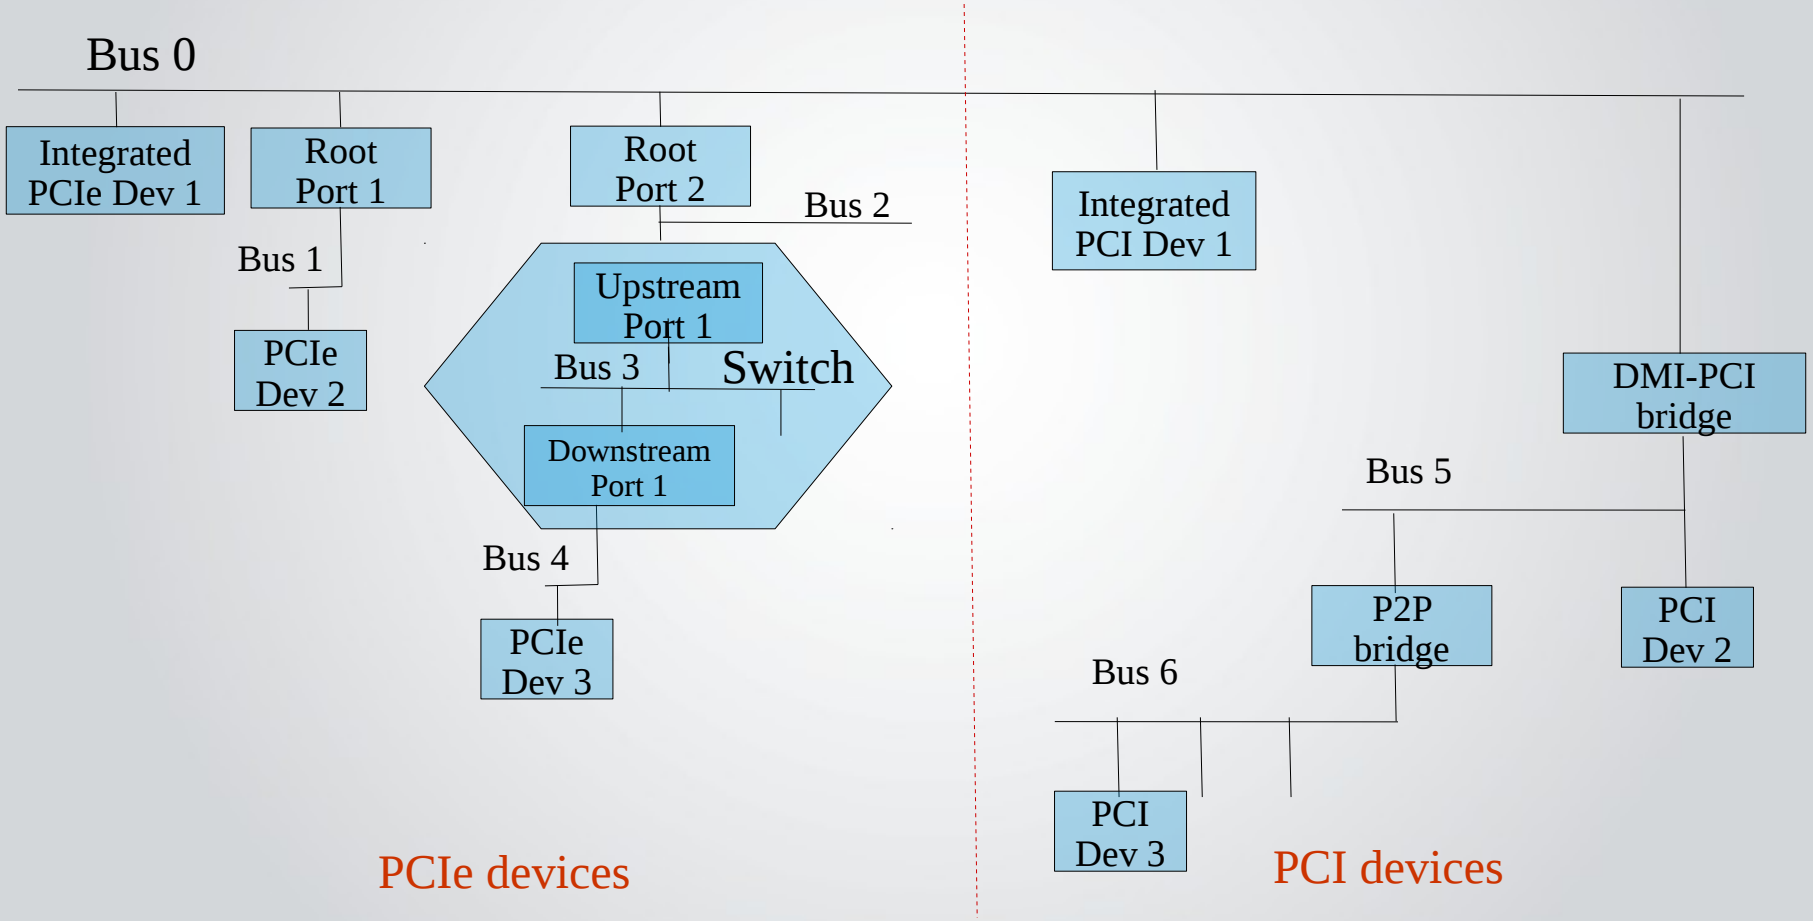

# Q35 emulation

- PCIe/PCI ACPI Hotplug

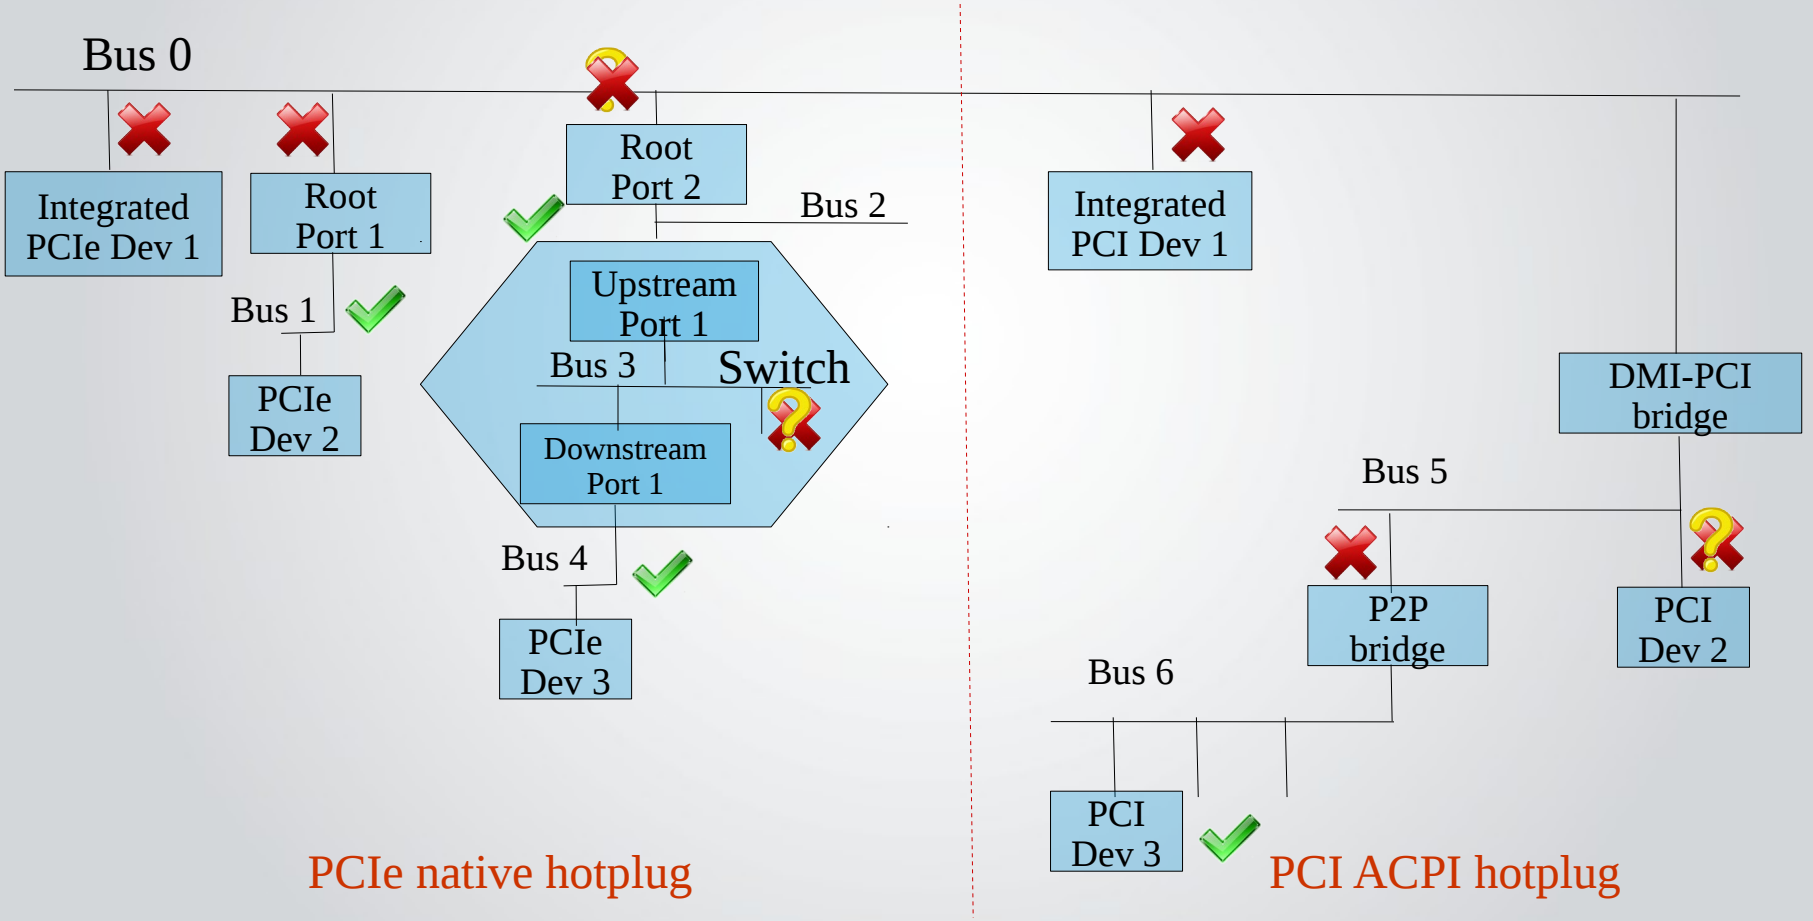

# Q35 emulation

- Virtio

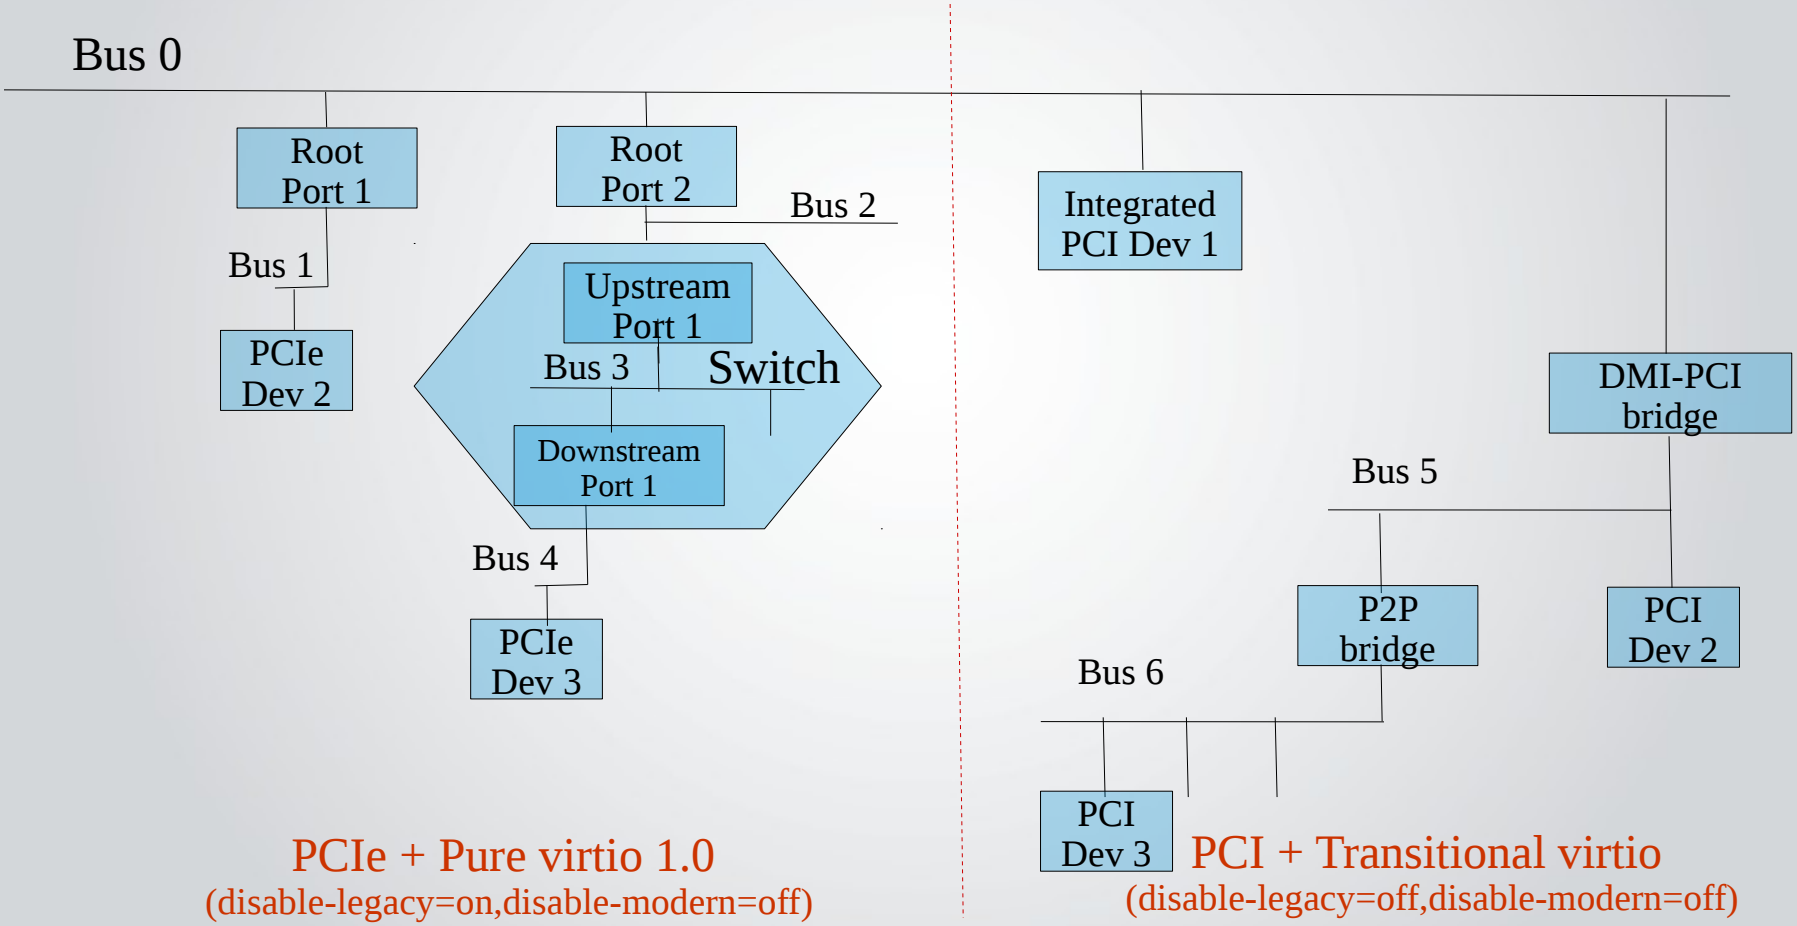

# Agenda

- Q35 chipset - Overview
- Q35 emulation
- **Q35-only features**
- Q35 limitations
- Q35 downstream

# Q35-only features

- **PCIe** “goodies”
  - Extended configuration space (MMCFG)
  - PCIe native hotplug
  - Advanced Error Reporting (AER)
  - Alternative Routing-ID Interpretation (ARI)
  - Native Power Management
  - Function Level Reset (FLR)
  - Address Translation Services (ATS)
- **AHCI** storage controller
- **vIOMMU** emulation
- “Secure” **Secure Boot**

# Agenda

- Q35 chipset - Overview
- Q35 emulation
- Q35-only features
- **Q35 limitations**
- Q35 downstream

# Q35 limitations

- **No support** for legacy guests (Windows XP/2000).
- **Questionable support** for legacy QEMU devices.
- **Limited IO space** can affect the number of devices used by a single Q35 machine:
  - Each device behind a separate PCI bridge.
  - Each bridge requires 4K IO range.
  - Several solutions available:
    - Plug only PCIe devices into PCIe ports.
    - Use smaller/non-standard IO windows for bridges.
    - ...

# Agenda

- Q35 chipset - Overview
- Q35 emulation
- Q35-only features
- Q35 limitations
- **Q35 use cases**

# Q35 use cases

- Current use-cases:
  - **P2V**: the only “out of tech preview” case
  - **Secure boot** (+OVMF)
  - **vIOMMU** (NFV)


# Thank you!