
PCI
SIG®

PCIe™ Trusted Config Space ECN
&
PCIe Link Speed Controls

Joe Cowan

Computer Systems Architect

Hewlett-Packard Company

Trusted Config Space ECN Session Outline

- TCS ECN Background & Key Terms
- Trust Issues with PCIe Platforms
- TCS ECN Details
  - ✓ Trusted Config Space and TCS Transactions
  - ✓ Trusted Config Access Mech (TCAM)
  - ✓ Standard vs Trusted Config Access
  - ✓ New Capability Structures
  - ✓ TCS Support in Root Ports, Switches, & Bridges
- TCS “Does not…” List
- Example Trusted Computing Platform
- Revisiting the Trust Issues
- Key Takeaways
- Recommendations

TCS ECN Background & Key Terms

- Trusted Configuration Space (TCS) is an ECN for the PCIe 1.1 Base spec
  - ✓ Brought to the Protocol Working Group in May 2004
  - ✓ Refined in a Sub-Workgroup through December 2004
  - ✓ Currently out for 30-day PCI-SIG Member Review
  - ✓ Final approval targeted for Summer 2005
- Key Terms
  - ✓ Trusted Device: An Endpoint, Root Port, or Switch that implements Trusted Config Space
  - ✓ Trusted Software: software that has a reliably established notion of identity
  - ✓ Trusted Software Environment: a protected environment for running Trusted Software

High-Level Trust Issues

- An increasing number of systems have critical private data that needs protecting
  - ✓ Financial Services and Banking
  - ✓ Healthcare
  - ✓ Personal Computers
- Security technologies typically use unique identifiers and keys to provide protection
  - ✓ Care must be used to avoid the use of unique identifiers to be linked with Personally Identifiable Information (PII)
  - ✓ TCS provides a means to isolate unique identifiers so that they cannot be linked to PII

Trust Issues with Existing PCIe Platforms

- Trusted Device perspective
  - ✓ Did the Request I received come from Trusted Software?
  - ✓ Did the Request I received come from some (untrusted) device instead of the Root Complex?
- Trusted Software perspective
  - ✓ Am I operating in a Trusted Software Environment?
  - ✓ Am I accessing a Trusted Device?
  - ✓ Is the device virtualized in an authorized manner?
  - ✓ Has the device been repurposed?

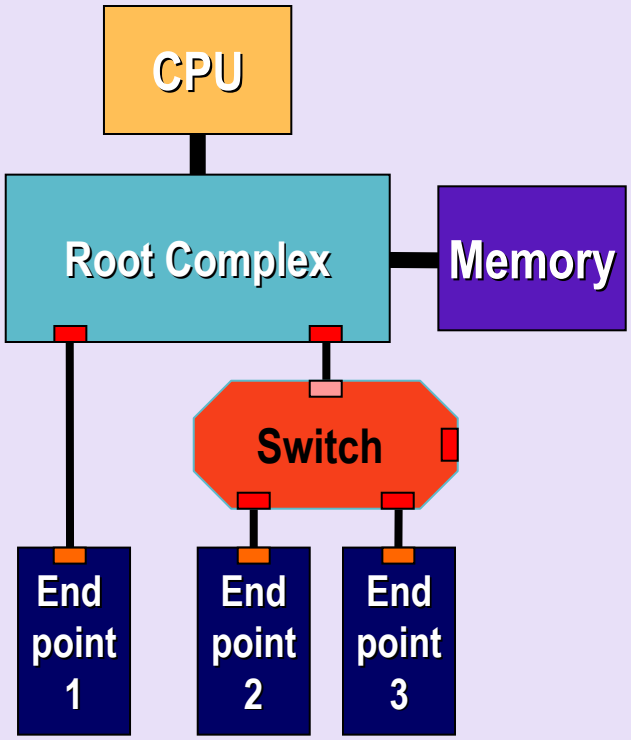

Will revisit these issues near the end to determine which are addressed by TCS and to what extent.

TCS ECN Details: TCS and TCS Transactions

- TCS is a new PCIe address space, distinct from I/O, Memory, Message, and (standard) Config
- 2 new PCIe Requests – Trusted Config Read & Trusted Config Write
- Parallels with (standard) Configuration Space:
  - ✓ Single byte-enabled DWORD payload per transaction
  - ✓ ID-based (Bus/Device/Function) routing
  - ✓ 4KB space per Function
- Contrasts with (standard) Configuration Space:
  - ✓ TCS is optional for Endpoints, Switches, & Root Ports
  - ✓ A distinct software access mechanism for TCS (TCAM)
  - ✓ TCS transactions are usable only after standard Config transactions have assigned hierarchy bus numbers

Trusted Config Access Mech (TCAM)

- The TCAM is patterned after the PCIe Enhanced Config Access Mechanism (ECAM), but is distinct.
- Parallels with the ECAM:
  - ✓ Memory-mapped regions, 1 MB per bus number
  - ✓ Base Addrs & Bus number ranges reported by firmware
  - ✓ Single or multiple TCAMs per platform or partition
  - ✓ Each TCAM maps n bus number bits, with $1 \leq n \leq 8$
  - ✓ Supports single- and multi-segment models
  - ✓ Platforms can provide firmware access in lieu of TCAMs
- Contrasts with the ECAM:
  - ✓ Usable only by the Trusted Software Environment (specifics are outside the scope of the specification)
  - ✓ On some platforms, usable only when enabled by a hardware signal from a Trusted Platform Module (TPM)

Standard vs Trusted Config Access

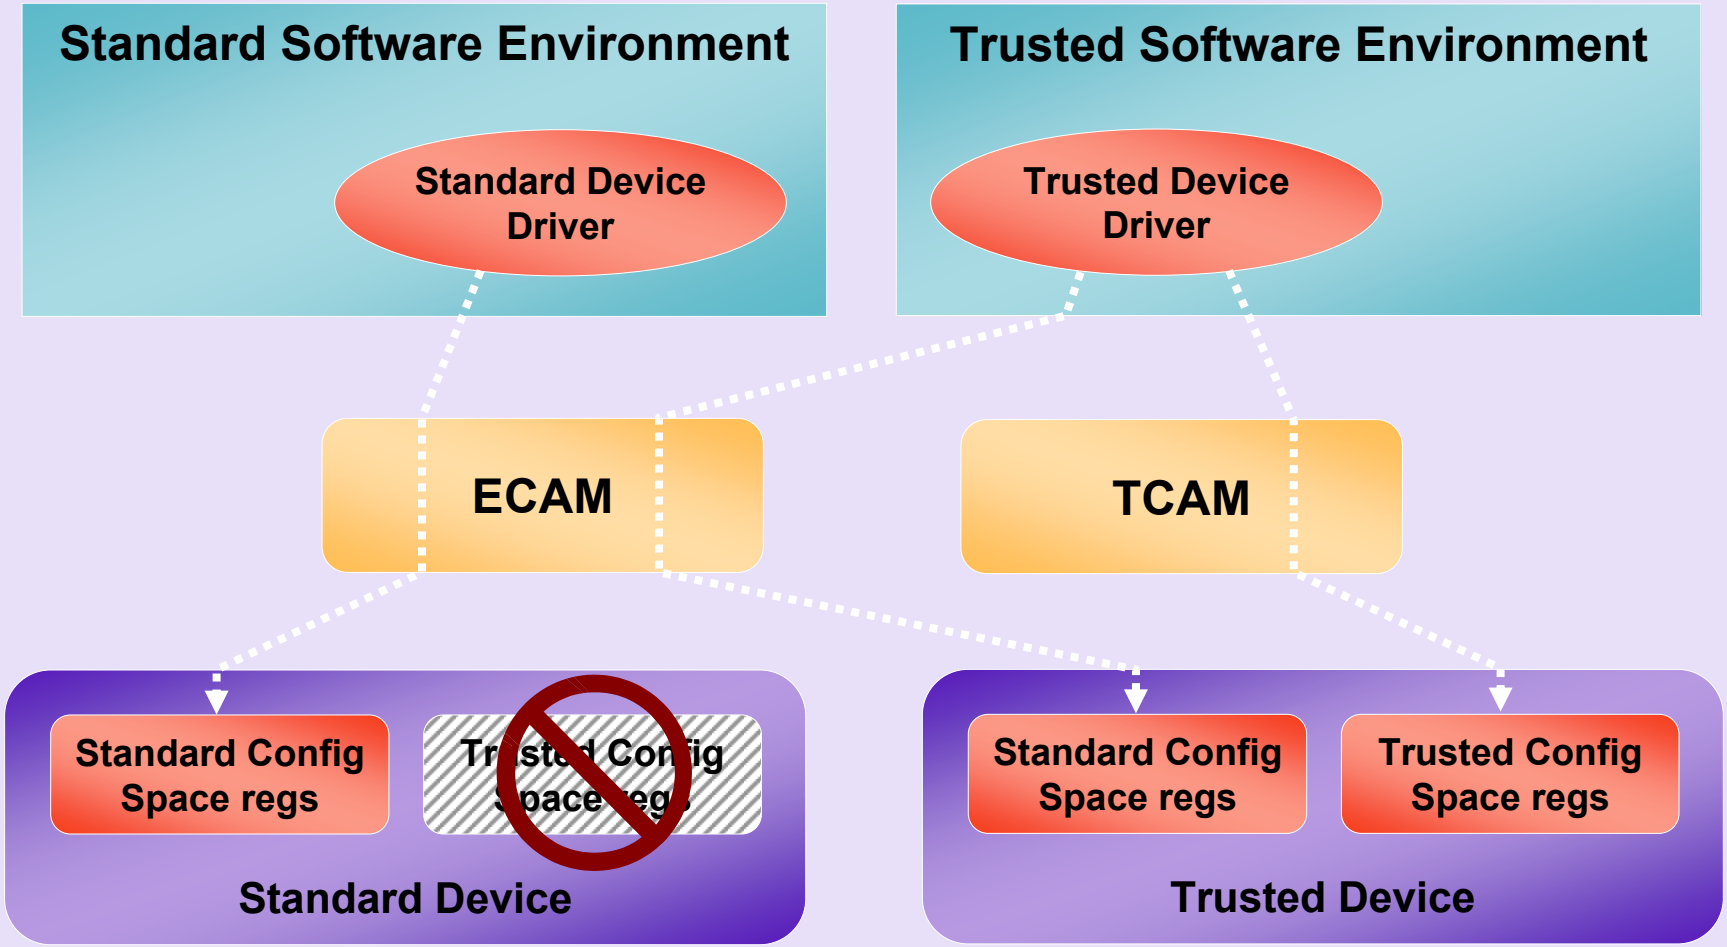

Issue: how can Trusted Software verify that its standard Config & Trusted Config accesses are indeed reaching the same Device?

New Capability Structures

- Standard Config Space Capability Structures
  - ✓ Config Access Correlation (CAC) Capability Struct: used to verify correlation between Trusted & Standard Config Spaces; existence also indicates the existence of a Trusted Config Space in this Function
- Trusted Config Space Capability Structures
  - ✓ Config Access Correlation (CAC) Trusted Capability Struct: used to verify correlation between Trusted & Standard Config Spaces; required by all Trusted Devices
  - ✓ Vendor-Specific Trusted Capability (VSTC): similar to a Vendor-Specific Extended Capability in standard Config Space, however vendors can define “public” VSTCs, which other vendors are permitted to use. “Private” VSTCs can also be defined.

Indicating TCS-Related Capabilities

- Firmware indicates the presence of TCAMs.
- Root Ports and Switches indicate their ability to route TCS Requests via a new bit in the PCIe Capabilities Register in standard Config Space.
- Functions indicate their implementation of a Trusted Config Space by the existence of a new CAC Capability Struct in standard Config Space.

TCS Support in Root Ports, Switches, & Bridges

- Adding Root Port and Switch routing support for TCS Requests is relatively easy:
  - ✓ Recognize two new Transaction Layer Packet types
  - ✓ Handle TCS Request routing identically to standard Type 1 Config Request routing
  - ✓ Indicate this support by setting one bit in the PCIe Capabilities register
- PCIe to PCI/PCI-X Bridge routing support for TCS Requests is not defined. There are no plans to add TCS to PCI-X or Conventional PCI.
- Root Ports, Switches, & Bridges are permitted to implement Trusted Config Space as Completers, but the value of doing so is debatable
  - ✓ Root Ports and Switches are not required to implement Trusted Config Space in order to support TCS routing

TCS Itself Does Not …

- Provide a complete solution for building Trusted Computing Environments
  - ✓ Substantial additional support required from hardware, OSs, drivers, applications, & firmware (some platforms)
- Handle Authentication (confirmation of identity)
  - ✓ One envisioned approach is to include certificates or digital signatures in VSTCs.
- Prohibit virtualization of Trusted Devices
  - ✓ Virtualization is primarily a platform policy.
- Provide high-bandwidth access to/from a device
  - ✓ TCS itself is a low-bandwidth channel on most platforms, but can be used as a secure channel to set up high-bandwidth channels, e.g. that use encryption.

Example Trusted Computing Platform

- The Trusted Computing Group™ has specified a Trusted Platform Module (TPM), which as part of a trusted subsystem can provide:
  - ✓ Protected storage
  - ✓ Protected capabilities
  - ✓ Authentication of the platform
  - ✓ Measurement of platform integrity
  - ✓ Attestation of platform integrity
- The TPM can be used to assert a hardware signal that enables the TCAM for use only if/when platform integrity has been attested.

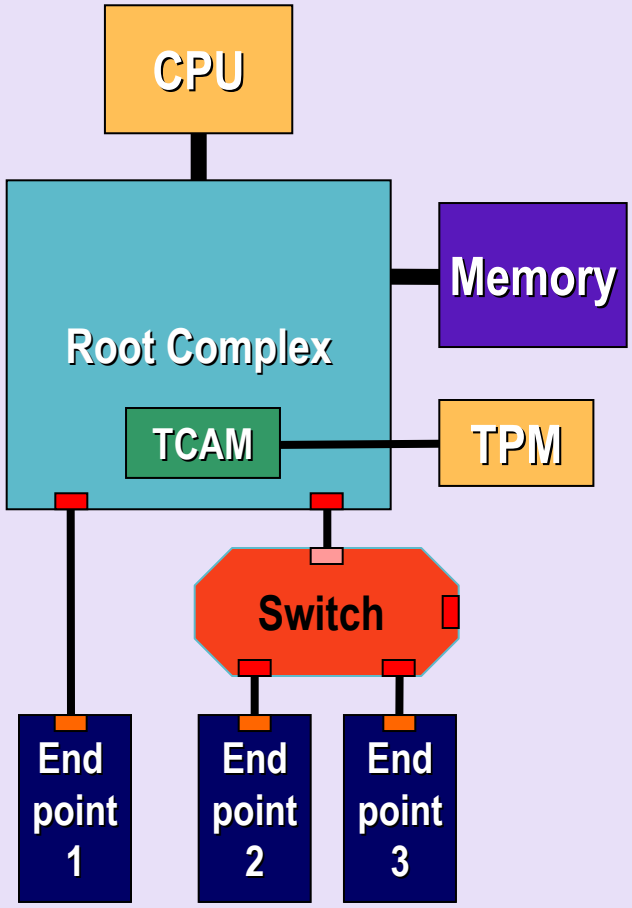

Revisiting the Trust Issues: Device Perspective

- Did the Request I received come from Trusted Software?
  - ✓ Only Trusted Software operating within the Trusted Software Environment is permitted to generate TCS Requests.
  - ✓ Other technologies are needed to assess platform integrity.
- Did the Request I received come from some (untrusted) device instead of the Root Complex?
  - ✓ Endpoints are prohibited from generating TCS Requests.
  - ✓ Root Ports and Switches are prohibited from routing TCS Requests upstream or peer-to-peer.

Revisiting the Trust Issues: Software Perspective

- Am I operating in a Trusted Software Environment?
  - ✓ Other technologies are needed to assess platform integrity.
- Am I accessing a Trusted Device?
  - ✓ Trusted Software can determine if a device implements TCS.
  - ✓ Other technologies are needed to assess device integrity.
- Is the device virtualized in an authorized manner?
  - ✓ Device virtualization is a policy belonging primarily to the Platform.
  - ✓ High-end servers are likely to support virtualization environments.
  - ✓ Personal computers are envisioned less likely to support them.
- Has the device been repurposed?
  - ✓ Untrusted devices with FPGAs might be reprogrammed “in the field” to masquerade as Trusted Devices.
  - ✓ Wide-spread use of FPGAs throughout the industry makes it unreasonable for the ECN simply to prohibit them.
  - ✓ Other technologies are needed to assess device integrity.

Key Takeaways

- TCS provides part of the foundation for Trusted Computing Environments, and is intended to work in synergy with other independent industry efforts such as those within the Trusted Computing Group™.
- TCS requires special hardware support in the Root Complex, Switches, and Trusted Devices.
- TCS is envisioned to become an important feature on a wide range of PCIe platforms, from personal devices to personal computers to all classes of servers.

Recommendations

- Chipset Designers: add support for TCS in your chipsets
- System OEMs: add support for TCS in your platforms
- Switch Vendors: add support for TCS routing in all Switches
- Adapter Vendors: add support for TCS if appropriate for your particular device

Additional Resources

http://www.pcisig.com

TCS Draft ECN for Member Review (PCI SIG Members only)

http://www.pcisig.com/specifications/pciexpress/review_zone

https://www.trustedcomputinggroup.org *

* Note: the Trusted Computing Group™ has no formal involvement with the development of TCS.

PCIe Link Speed Controls

PCIe Link Speed Controls Session Outline

- Link Speed Controls for PCIe 2.0 Signaling
- Link Bandwidth Change Notification Mech ECR

Note:
The following is still under active development and subject to change!

Link Speed Controls for PCIe 2.0 Signaling

- Being developed in the Protocol and Software Workgroups to complement the work being done in the Electrical Workgroup
- External Link Speed Management Model
  - ✓ By default, hardware automatically trains to the greatest common speed
  - ✓ For whatever reason, software is permitted to place an upper bound on the speed
  - ✓ For Link reliability, hardware is permitted to place an upper bound on the speed
  - ✓ Still investigating if hardware should be permitted to change the speed autonomously for power management purposes
  - ✓ New mechanism supporting software control for entering/exiting Compliance Mode
- New/Modified Registers for External Links
  - ✓ *Maximum Link Speed* renamed to *Supported Link Speeds* (Link Cap reg)
  - ✓ *Target Link Speed* (Link Control Reg)
  - ✓ *Hardware Autonomous Speed Control* (Link Control Reg)
  - ✓ *Enter Compliance* bit (Link Control Reg)
  - ✓ *Current Link Speed* (Link Status Reg)
  - ✓ *Link Speed Change Attempt Failed* bit (Link Status Reg)
- Root Complex Internal Links
  - ✓ Can report their supported and current speeds via similar changes to their Regs
  - ✓ Will not be controllable via these mechanisms

Link Bandwidth Change Notification Mech ECR

- New ECR under development within the Protocol Workgroup
- Motivation
  - ✓ Need mech for PCIe-aware software to be notified when Link bandwidth (speed or width) changes, due to hardware-autonomous Link retraining
  - ✓ Can help reduce vendor support costs by having software notify users if marginal Links retrain to a lower bandwidth, impacting system performance
- Mechanism
  - ✓ *Link Bandwidth Changed Notification Capability* bit (Link Cap reg)
  - ✓ *Link Bandwidth Changed Interrupt Enable* bit (Link Control reg)
  - ✓ *Link Bandwidth Changed Status* bit (Link Status reg)
- Why packaged separately from Link Speed Controls
  - ✓ Want it available ASAP for all new PCIe components, not just PCIe 2.0
  - ✓ Though separate, still logically coupled with Link Speed Controls
  - ✓ ECR timing will be somewhat tied to Link Speed Controls stabilizing
- Intend to make this a mandatory feature for PCIe Base 2.0

Thank you for attending the
PCI-SIG Developers Conference 2005.

For more information please go to
www.pcisig.com


PCI
SIG®