

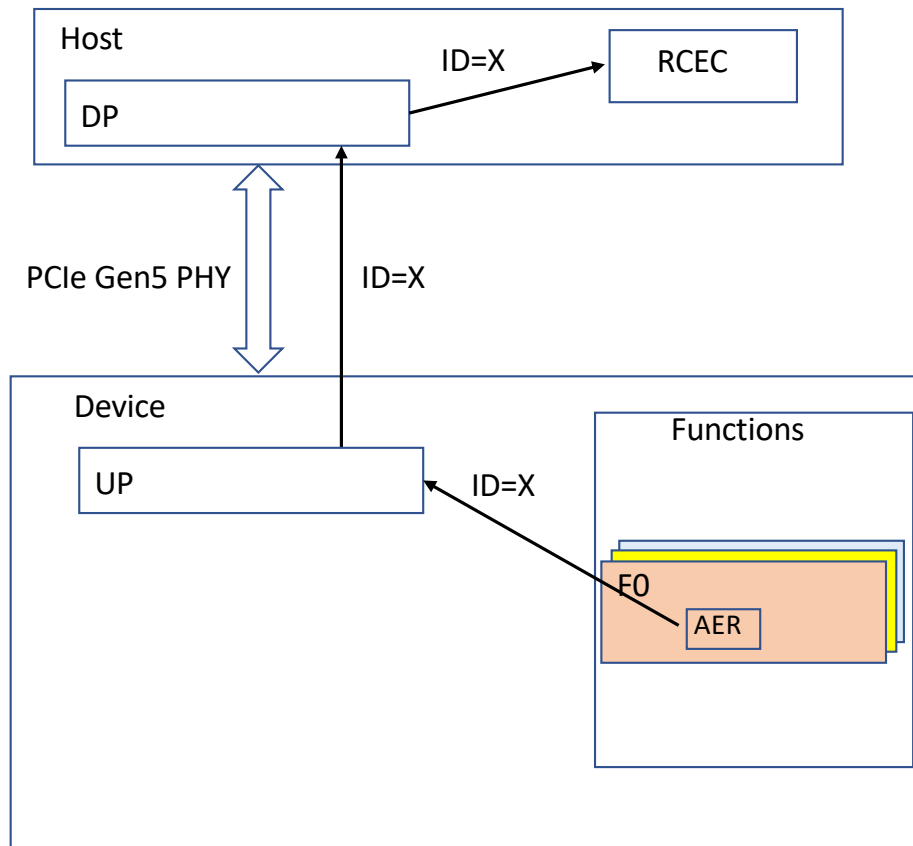
# AER handling for RCECs

Sean V Kelley

# What is RCEC?

- Root Complex Event Collectors
  - Provides support for terminating error and PME messages from Root Complex Integrated Endpoints (RCiEPs).
  - Resides on a Bus in the Root Complex.
  - An RCEC will explicitly declare supported RCiEPs through the Root Complex Endpoint Association Extended Capability.  
(See PCIe 5.0-1, sections 1.3.2.3 (RCiEP), and 7.9.10 (RCEC Ext. Cap.))
- To be discussed in this session
  - Native vs Non-native handling: The firmware takes action in place of an RCEC.
    - AER native: Root Port or RCEC
    - AER via APEI: Root Port, RCEC, or other PCIe device / non-existent RCEC (non-native)

# Native Case



- RCIEP Detected errors
  - Errors are logged in local AER extended capability per function.
  - Error messages sent to RCEC.
  - RCEC logs errors and generates MSI if enabled.

When an RCEC device signals error(s) to a CPU core, the CPU core needs to walk all the RCiEPs associated with that RCEC to check errors.

OS error handler may begin by inspecting the RCEC AER Extended Capability and follow PCI Express rules to discover the source of the error.

# Non-Native Case

- ACPI Platform Error Interface (APEI)
  - In this case, an RCEC may not even need to exist.
  - Platforms could have a hardware entity performing the role of an RCEC, but not visible to the OS.
  - An RCiEP would thus have no visible association with an RCEC.
  - This is true where APEI [1] is used to report the AER errors via a GHES[v2] HEST entry [2] and relevant AER CPER record [3] and non-native handling is in use.

[1] ACPI Specification 6.3 Chapter 18 ACPI Platform Error Interface (APEI)

[2] ACPI Specification 6.3 18.2.3.7 Generic Hardware Error Source

[3] UEFI Specification 2.8, N.2.7 PCI Express Error Section

# Discussion

- Error Recovery
  - When attempting error recovery for an RCiEP associated with an a native RCEC device, there needs to be a way to update error status and severity.
  - Depending on the Native versus Non-Native Case, it may be that the expected error registers do not even exist.
  - Thus a “bridge” could point to any of the following:
    - bridge points to a device (RP, DP, RCEC)
    - bridge points to an associated RCiEP’s own RCEC
    - bridge does not even exist (AER/APEI case)
  - What additional considerations should be considered or addressed in handling the call to `pcie_do_recovery()`, including the subsequent reset if needed?

Backup

# History of RCEC

- Current patch series:
  - <https://lore.kernel.org/linux-pci/20200812164659.1118946-1-sean.v.kelley@intel.com/>
- Native versus Non-Native Handling discussion:
  - <https://lore.kernel.org/linux-pci/20200818100150.00007f29@Huawei.com/>