NC-SI 1.2
PCIe Functions Representation (Work-in-Progress)

Version 0.3

September 19, 2017

Disclaimer

- The information in this presentation represents a snapshot of work in progress within the DMTF.
- This information is subject to change without notice. The standard specifications remain the normative reference for all information.
- For additional information, see the Distributed Management Task Force (DMTF) website.

PCIe Functions in NC-SI 1.2

- NC-SI 1.2 shall address the following associations of PCIe physical functions (PFs) with network ports
  - No network port partitioning within a host, single host
    - Single PF associated with each network port
    - Single PCIe root complex
  - No network port partitioning within a host, multiple hosts sharing a network port
    - 1 PF per host associated with a network port
    - Multiple PCIe root complexes, one per host
    - Network port partitioning within a host, single host: Multiple PFs associated with a network port
    - Network port partitioning within a host, multiple hosts sharing a network port: Multiple PFs per host associated with a network port
- A physical network port partitioned into multiple logical network ports relative to one or more hosts (PFs only)
- Multiple PFs within a single PCIe root complex, PFs are associated with the network port
- Multiple PCIe root complexes, 1 PF per PCIe root complex, each PF associated with the network port
- Multiple PCIe root complexes, multiple PFs per PCIe root complex, PFs associated with the network port
- A physical network port shared by multiple hosts shall support one or both BMC models below
  - Single BMC per host
  - Single shared BMC across all hosts
- Multiple PFs with multiple PCIe root complexes
- 1 PF/root complex as well as multiple PFs/root complex
- SR-IOV introduces another layer in hierarchy: VFs as children of PFs

PCIe Functions in NC-SI 1.2

- NC-SI 1.2 shall have the following representation of PFs and VFs
  - A PF associated with a specific NC-SI channel
  - A VF associated with a specific NC-SI channel that its parent PF is associated with
- NC-SI 1.2 may have the following representation of PFs and VFs
  - A PF associated with a package
  - A VF associated with a package
- NC-SI 1.2 shall support AENs to address change in associations of PCIe functions to network ports (e.g. hot plug events)

PCIe Function Identifier in NC-SI 1.2

- Not part of NC-SI control packet header
- Only used in commands targeted to PCIe functions
- First field in the NC-SI command packet payload
- What is the format for PCIe function id?
  - Options: PCIe BDF format or HRBDF format or PCIe ARI format
  - Recommendation: Use HRBDF format
- NC-SI 1.2 commands shall use PCIe function id to address a function
- A command targeted to a PCIe function can be either package command or channel command
  - Function identifier shall be validated before the command execution
  - Function to channel association shall be checked for channel cmd

Proposed NC-SI 1.2 Commands for PCIe Functions

- Query Device Capabilities (package level command)
  - Response provides partitioning, multi-host, and SR-IOV capabilities
  - This command should provide the maximum capabilities of the device
- Query Port Capabilities (channel level command)
  - Response provides partitioning, multi-host, and SR-IOV capabilities
  - This command should provide the maximum capabilities per port
- Get PF Information (channel command)
  - No function identifier in the command
  - Response provides identifiers of all PFs associated with the channel or package
- Get Associated VF Information (channel command)
  - PF identifier is specified in the command packet payload
  - Response provides identifiers of all VFs associated with a specific PF
  - Response provides all VF identifiers on the channel or package (special case)
  - This should include both maximum VFs available on the PF as well as number of VFs enabled by the host OS

Proposed NC-SI 1.2 Commands for PCIe Functions

- Get MAC Address (channel or package command)
  - Function identifier is specified in the command packet payload
  - Response provides permanent or configured MAC address on the function
- Set MAC Address (channel command)
  - Function identifier is specified in the command packet payload
  - MAC address/type (e.g. LAN) being configured is provided in cmd payload
  - Successful response means MAC address of a specific type is configured

Proposed NC-SI 1.2 Commands for PCIe Functions

- Get Transmit Bandwidth (channel or package command)
  - Function identifier is specified in the command packet payload
  - Response provides transmit bandwidth (min, max) of the function
- Set Transmit Bandwidth (channel command)
  - Function identifier is specified in the command packet payload
  - Bandwidth (min, max) being configured is provided in command payload
  - Successful response means transmit bandwidth is applied to that function
  - The sum of bandwidth across all enabled PFs on a channel should add up to port bandwidth.
  - Enabling and Disabling of PFs may change bandwidth allocation.
  - Maximum bandwidth can be oversubscribed.
  - For bandwidth allocation, there should be one common method specified.
  - Options for bandwidth allocation: 1) bandwidth allocation across all PFs using weighted, absolute, and percentage of bandwidth, 2) bandwidth setting per PF followed by a commit operation for all BWs to take effect

Proposed NC-SI 1.2 Commands for PCIe Functions

- Get Function Statistics (channel or package command)
  - Function identifier is specified in the command packet payload
  - Response provides packet statistics for the function
- Get Function Status (channel or package command)
  - Function identifier is specified in the command packet payload
  - Response provides status (enabled/disabled, driver status, etc.) of the function
- Get Boot Configuration (channel or package command)
  - Function identifier is specified in the command packet payload
  - Response provides boot config (e.g. iSCSI boot params) for the function
- Set Boot Configuration (channel or package command)
  - Function identifier is specified in the command packet payload
  - Boot configuration and parameters are provided in command pkt payload
  - Successful response means boot configuration is applied to that function

Teaming Support

- Teaming support for multiple ports per function should be a 2.0 spec item.
- In the 1.2 version, functions should be mapped to a single channel, while in 2.0 functions can be mapped to multiple channels.
- In the 1.2 version, the PCIe function representation should not prevent adding teaming support in 2.0 version.

Get Version ID Command

- This command reports ID of a single function on a channel.
- For multiple functions, we need a mechanism to address multiple functions IDs on a channel.
- Proposal TBD in NC-SI 1.2 spec.