



# Redfish Model for SmartNICs

DPU, IPU, FPGA-based SmartNIC, etc.

Work-in-Progress

©2021DMTF

# 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 DMTF website: [www.dmtf.org](http://www.dmtf.org)

# Overview

- SmartNICs have reached the point of prevalence in the industry where consistent manageability via Redfish is necessary
- SmartNICs support a variety of networking specific features:
  - Metering/limits per network function
  - Firewalling capabilities per network function
  - Offload capabilities, such as RDMA and iSCSI
  - Address ranges or other switching related properties
- Other modeling aspects of SmartNICs
  - Physical and virtual functions
  - Devices, such as drives, GPUs, and FPGAs, that can be assigned to process workloads

# Overview

- Some types of SmartNICs contain entire systems and require this to be represented in the Redfish model
- All SmartNICs are represented with their own Chassis resource
  - *ChassisType* is usually *Card*, but can be others such as *Module*, *Zone*, or *Chip*
  - Chassis is also used to model network adapters, power, and thermal characteristics
- Redfish contains many components today for SmartNICs
  - NetworkAdapter, NetworkInterface, NetworkDeviceFunction, and Port for the network functions, physical containment, and traffic association
  - AddressPool and Endpoint for switching related properties
  - ComputerSystem for any integrated systems on the SmartNIC
  - Processor for any integrated FPGAs, GPUs, or other processing elements
  - Manager for SmartNICs with BMCs

# ACD Model for SmartNICs

- Why the ACD (Advanced Communication Device) model?
  - NetworkAdapter, NetworkInterface, NetworkDeviceFunction, and Port resources model the mapping of interfaces on a network function to different systems
  - Endpoint resources can be mapped to different network functions for representing placement in a fabric
  - Metrics for these are already defined for representing traffic and other telemetry

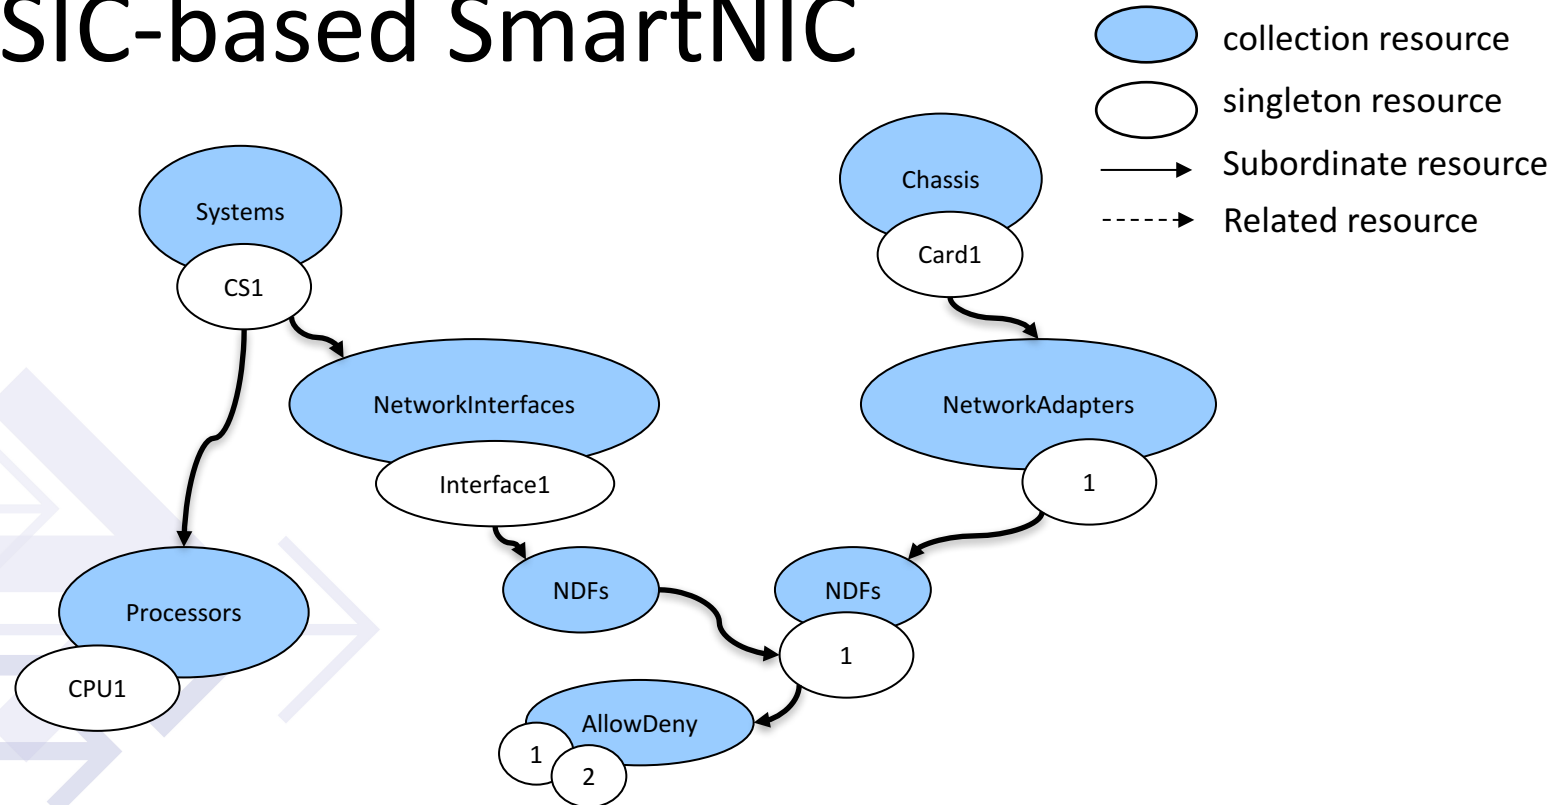
# Changes to the ACD Model

- Add byte and packet limits to `NetworkDeviceFunction`
- Add *AllowDeny* to `NetworkDeviceFunction`
  - Firewall configuration of a network function
- Add *DPU* to *SystemType* in `ComputerSystem`
  - Allows clients to distinguish a standard system in the system collection from the system contained in a SmartNIC
- Add *OffloadProcessors* to `NetworkDeviceFunction` and *Processors* to `NetworkAdapter`
  - For FPGA-based SmartNICs where dedicated programmable FPGAs are used for workloads
- Add *OffloadSystem* to `NetworkDeviceFunction`
  - For SoC-based SmartNICs where a DPU is used for workloads

# Representing Various SmartNICs

- Start with the ACD model and use *Links* to show relationships between the different components
- Model each SmartNIC as its own Chassis
- ASIC-based: Use the ACD model as-is (no changes needed)
- FPGA-based
  - Represent FPGAs as Processor resources
  - Add *Processors* in the NetworkAdapter resource to show offload processors
- SOC-based
  - Chassis resource references a ComputerSystem that representing the SoC
  - The ComputerSystem resource for the SoC contains *DPU* for its *SystemType*

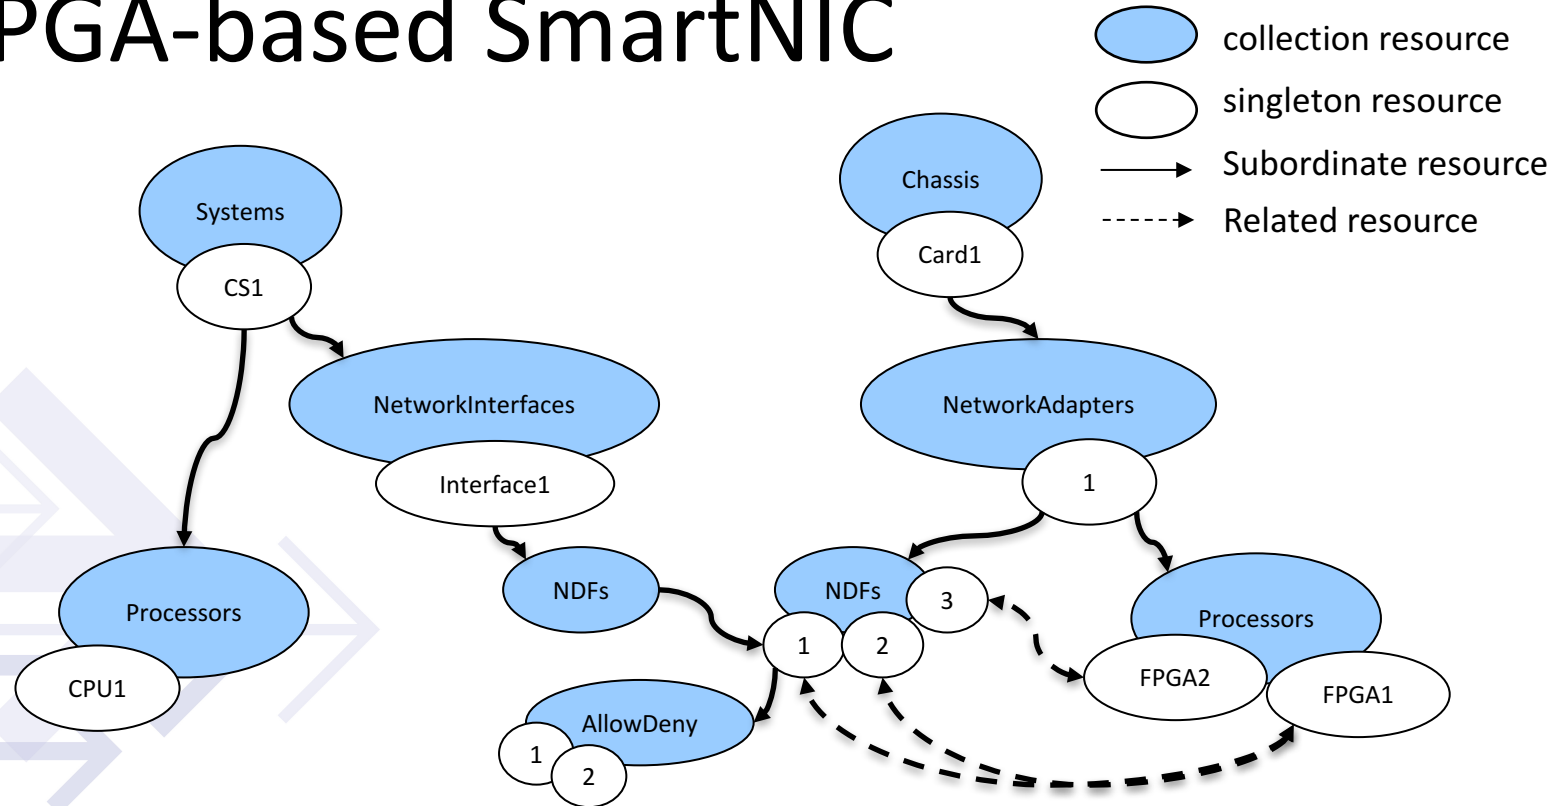
# ASIC-based SmartNIC



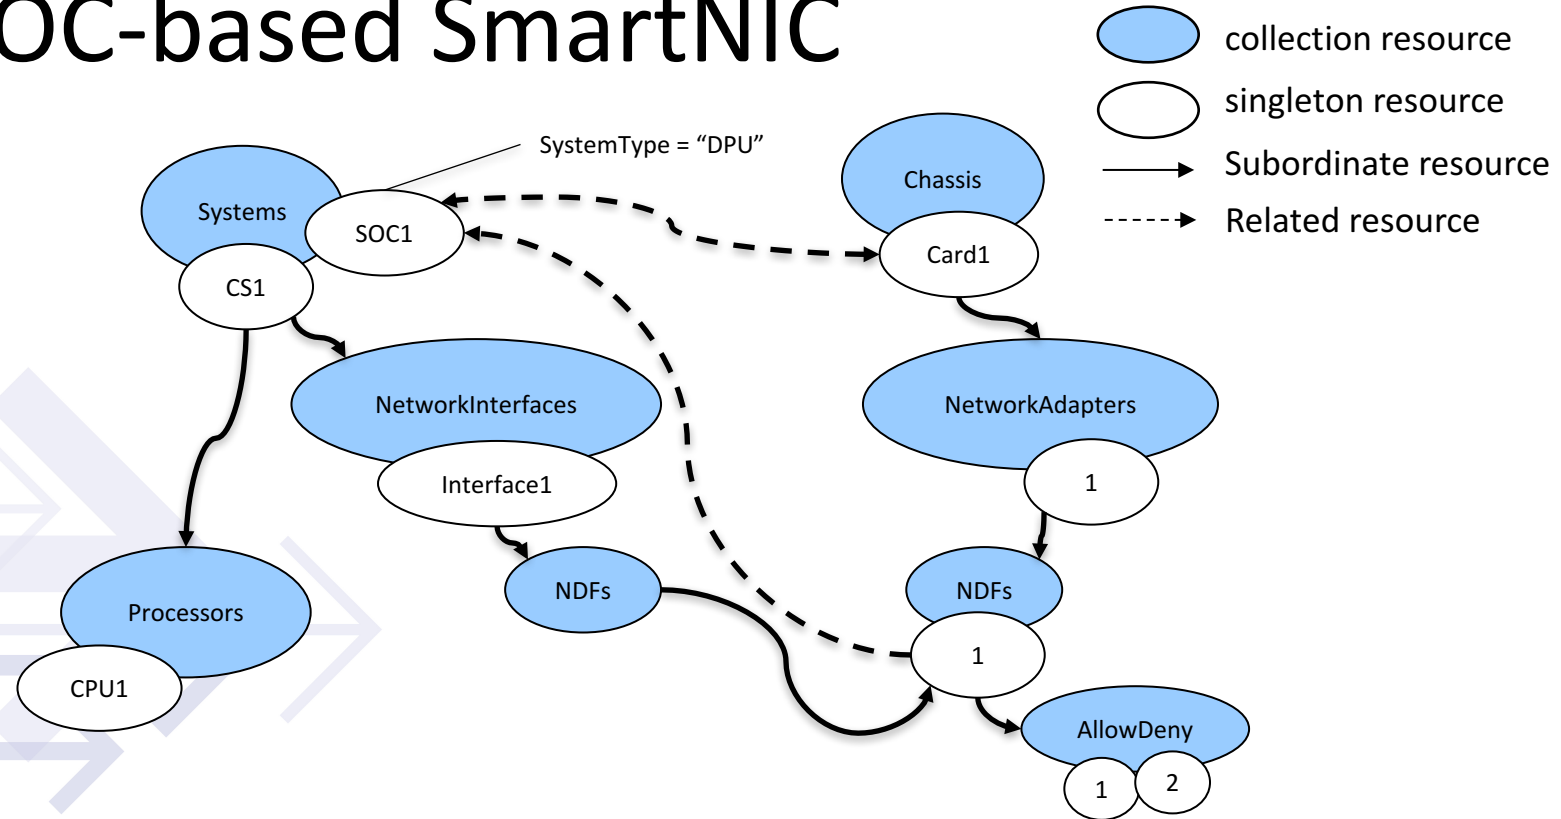




# FPGA-based SmartNIC



# SOC-based SmartNIC



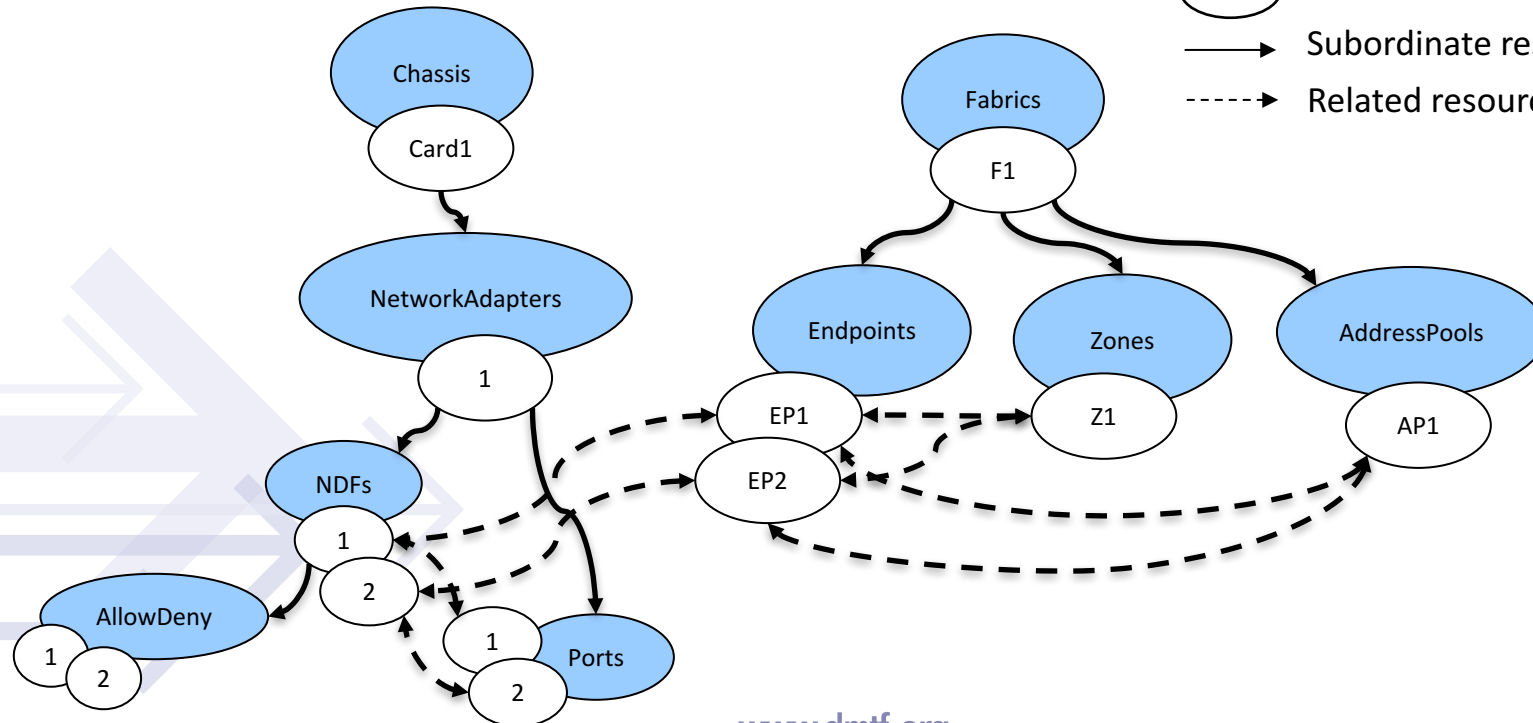

# Fabric Model for SmartNICs


- Why the fabric model?
  - Many SmartNICs have configurable switching properties, but modeling switches themselves isn't necessary
  - Zones can be used to show subnets and endpoint associations in an AddressPool resource
  - The AddressPool resource has eBGP settings as well as the IP addresses that can be assigned to the endpoints associated with the network functions
  - Endpoints also represent the addressability of the network function in the fabric

# Changes to the Fabric Model

- Existing AddressPool resource already contains the majority of the settings needed for SmartNICs
  - Gateway, DNS, NTP, DHCP, address ranges, and others
- Add *CIDR* and *Enabled* to BGP neighbor object in AddressPool

# SmartNIC Fabric Model





# Redfish