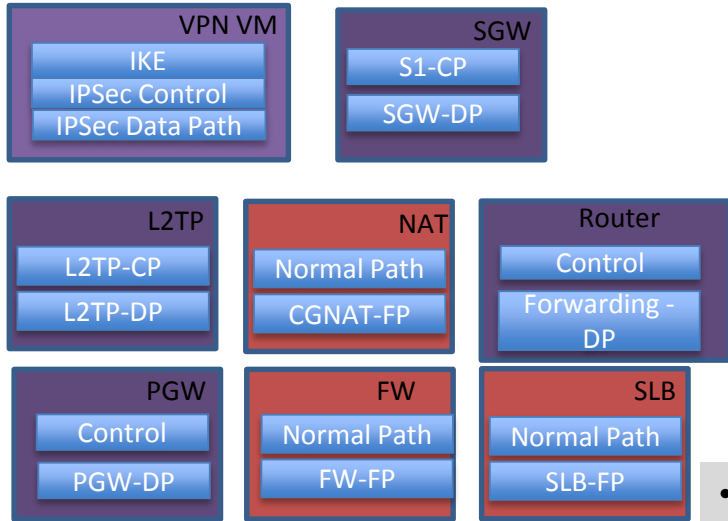




Flow Processing for Fast Path & Inline Acceleration

Srinivasa Addepalli – Intel
Denis Crasta - Freescale

Background : Performance requirements & Solutions



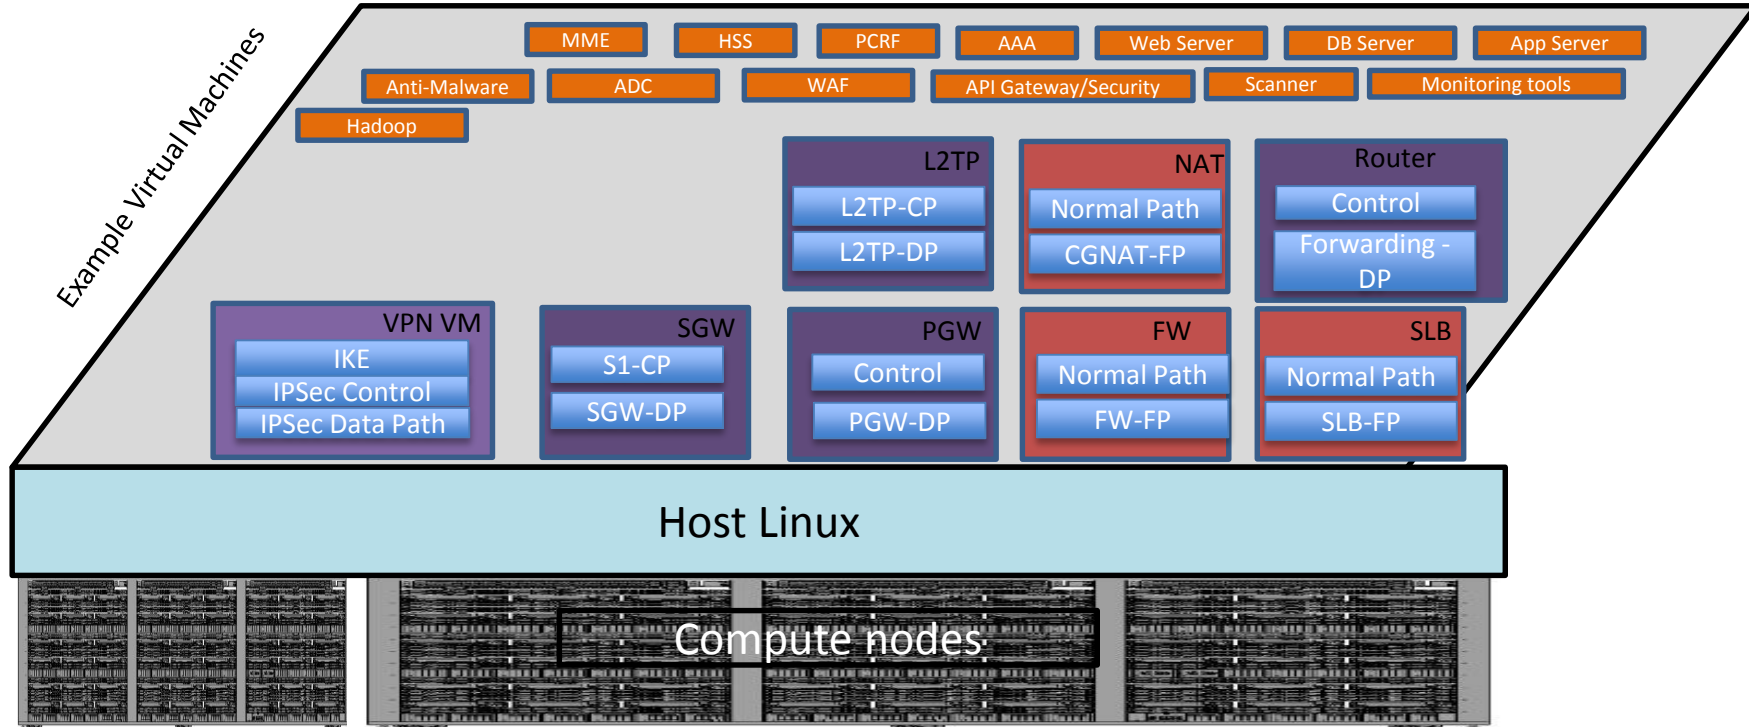
Few examples of network functions

- High PPS (Packets/sec)
- Low Latency
- Low Jitter
- High single flow performance
- Packet Order maintenance in a flow



- Data Path for many network services is simple and can be separated from Control-path/Normal-Path
- Data Path in physical appliances is normally implemented using specialized packet processors Specialized Packet Processors. Some examples include
 - Network processors (eg. PCIe based, Network based)
 - Inbuilt (to SoC) specialized network processing
 - FPGA

Background : Clouds - Virtual Machines with both Control and Data Planes together



Full inspection Applications

No data packet inspection Apps

Partial inspection apps

CP-DP separation is possible for partial and no-data inspection services/applications



Virtualization Challenges (Over Physical appliances)

Performance Challenges

- vSwitch overhead

Unable to utilize the specialized packet processors

Virtualization Considerations

Any vNF should run on compute nodes.

Supply Chain – Multiple suppliers in creating cloud environment

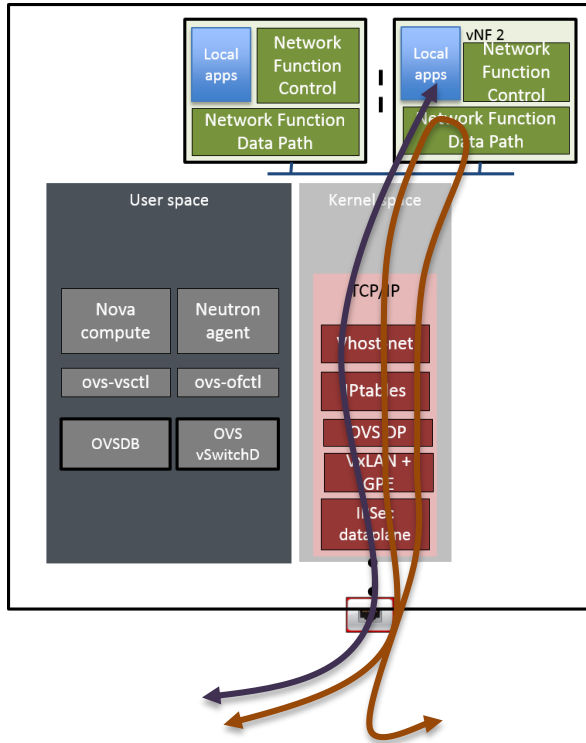
- Compute node hardware vendors
- VMM vendors
- Multiple vNFs & Multiple virtual appliance vendors.
- Multiple accelerator vendors

No/Minimal dependencies among the suppliers

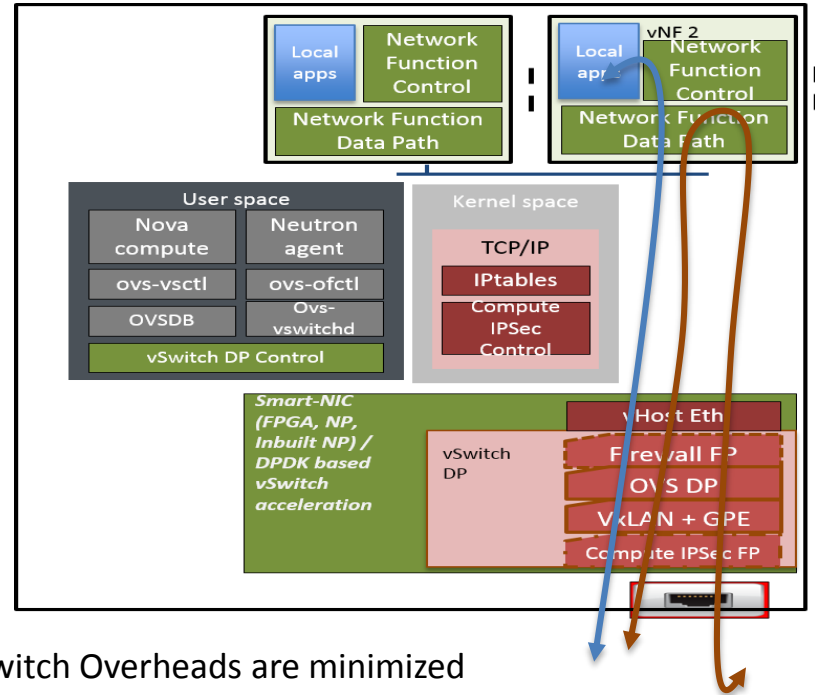
- Virtual appliances must continue to work with newer hardware (Accelerator and compute node) vendors without operator having to go back to virtual appliance vendors for image upgrades
- New virtual appliances should be able to take advantage of performance features.
- Standardization and programmability is the key.

Background : vSwitch Overheads – Mitigations

Traditional



Accelerated vSwitch compute nodes

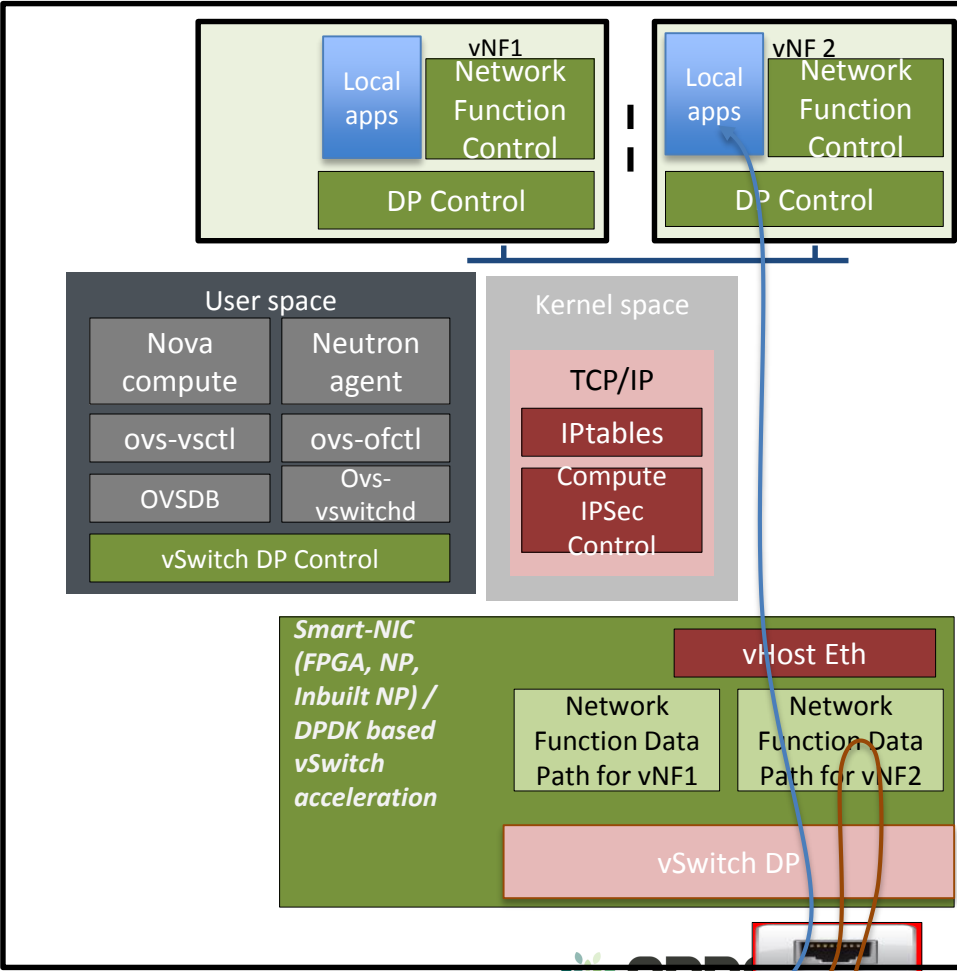


- vSwitch Overheads are minimized
- Direct connectivity with no changes to vNF is achieved using virtio interface

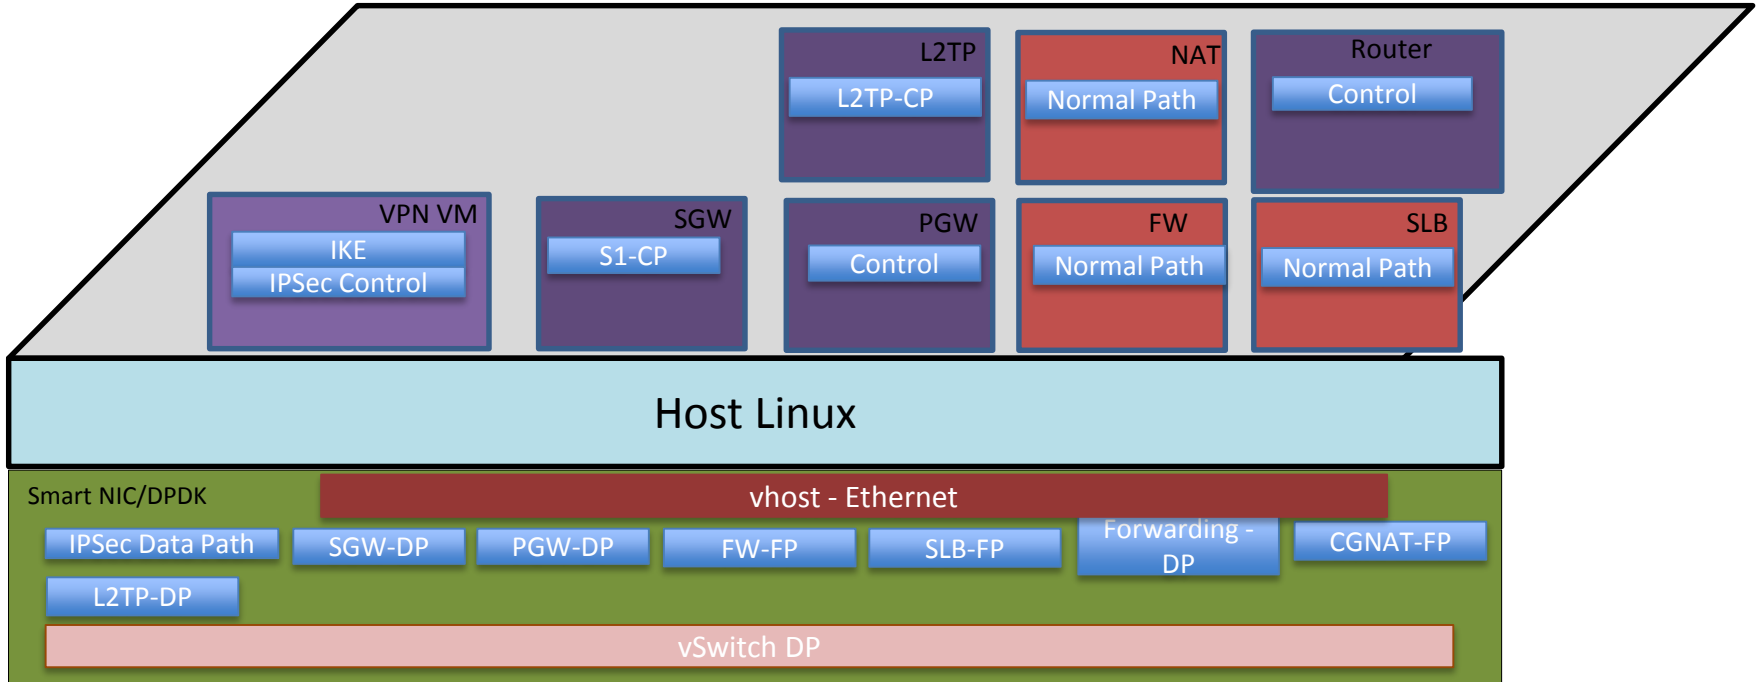
Requirement : vNF Data Paths in Packet processors

- Mitigate performance challenges
- Bring performance closer to the physical appliances

But... Cloud challenges....



Challenge – Proliferation of Data Path functions



- Many Data Path Functions now and in future → Could be a challenge for smart-NIC vendors → Challenge for operators in coordinating between virtual appliance vendors and smart-NIC vendors.
- Virtual appliance vendors might be satisfied with DP functions provided by smart-NIC vendors. vNF vendors might want to add their own extensions
- Choices :
 1. Provide a way for uploading of new DP modules in smart-NIC
 2. Flow programmable DPs (Flow processor)

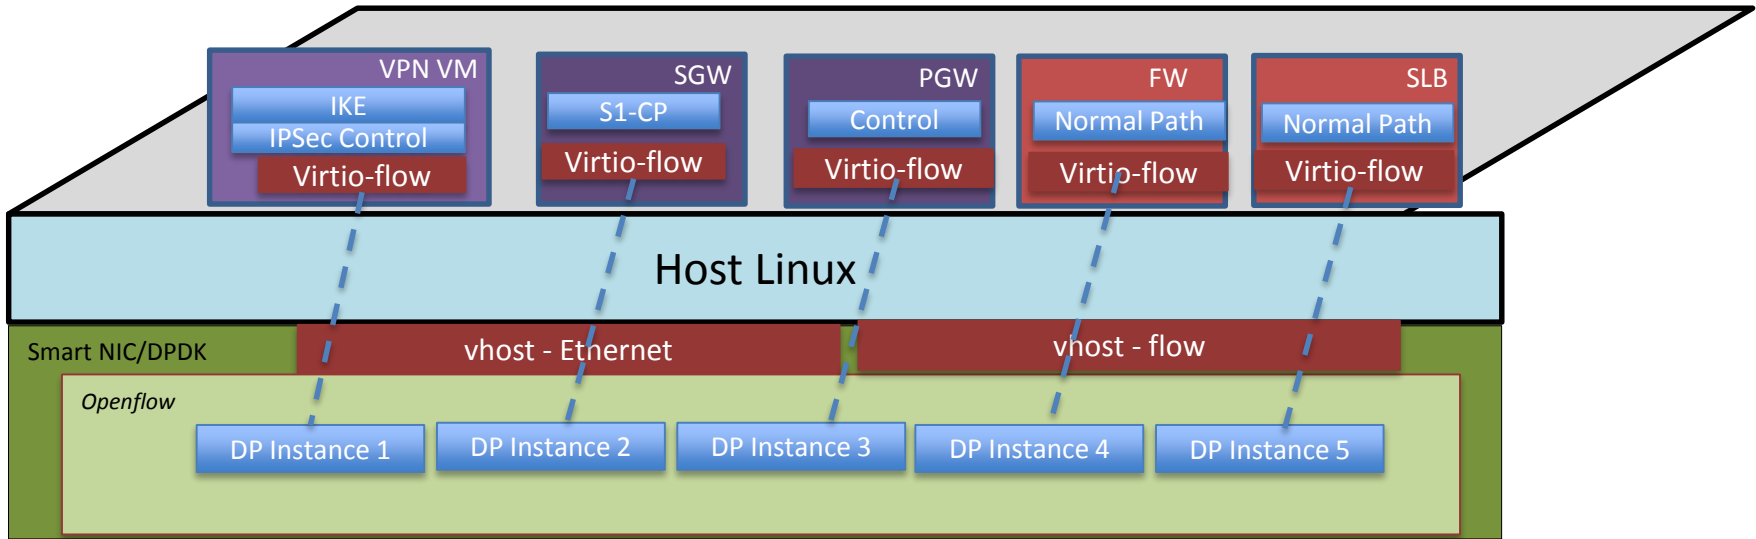
Challenges with Dynamic DP Modules

- **Many smart-NIC vendors and Smart-NIC technologies**
 - Virtual appliance vendors supporting many of them could be a challenge
- **Reliability Concerns**
 - Smart-NIC may not provide good isolation among DPs -> Malicious and poorly coded DP could make entire smart-NIC unstable.
- **Security Concerns**
 - Performance Isolation concerns
- **Coordination & Code Bloating Concerns**
 - Multiple versions of VMs of same type and multiple DP versions -> Code bloating.
 - Version mismatch
 - Dependencies among suppliers etc...

Flow Processing

- One Logic (One code piece)
- Run time programmable using messages from the control logic (No code upgrade).
- Multiple Instances of Data Paths.
 - Each instance can be programmable independently with tables, flows in tables and action in flows.
 - Pipeline of flows across the tables.
- Known best flow processing today
 - Openflow
 - Standardized by ONF.
 - Ability to dynamically program the flows and actions.
 - Ability to create multiple instances.
 - Ability to assign the ownership of control to instances.
 - Ability to control memory usage
 - Possible to implement processing isolation across instances
 - Proven/WIP to create DPs for Forwarding, IPsec, Firewall/NAT, SGW, PGW etc..
 - Possible to extend or create complementary specifications to meet our requirements.

Mitigation of proliferation of DP logic modules using Openflow



- As many DP instances as number of vNFs in the compute node.
- Smart-NICs to have only one logic module (Openflow).
- Each vNF controls its own DP instance by programming tables, flows and actions at run time.
- May require upgrade to smart-NIC software only for bug fixes & to implement OF extensions.

DPACC Proposal

- Identify the components to enable Openflow based Data Paths
 - Openstack Controller
 - Openstack agents
 - QEMU/KVM side
- Define the API and messages to communicate from Control to Data Planes
 - For standardization.
 - Binary compatibility and source compatibility.
- Identify the gaps & Document them.
 - Security and reliability aspects
 - Functional limitations
 - Any other...
- Create a PoC
 - Define the 1st functional PoC
 - Identify the existing Open Source
 - Identify the gaps.

Backup

