13th ANNUAL WORKSHOP 2017

INTEL® OMNI-PATH FABRIC VIRTUAL NETWORK INTERFACE CONTROLLER (OPA VNIC)

Niranjana Vishwanathapura, Software Development Engineer

Intel Corporation

[March 29, 2017]

LEGAL DISCLAIMERS

INTRODUCTION

Supporting Ethernet over Omni-Path fabric allows us to make full use of standard Ethernet support provided by the Operating System (including VLAN etc.) over the fabric without having verbs layering in the stack.

Intel Omni-Path (OPA) Virtual Network Interface Controller (VNIC) feature supports Ethernet functionality over Omni-Path fabric by encapsulating an Ethernet packet within an Omni-Path packet.

Agenda:

- OPA VNIC Architecture
- OPA VNIC Driver Design

OMNI-PATH VNIC ARCHITECTURE

ARCHITECTURE

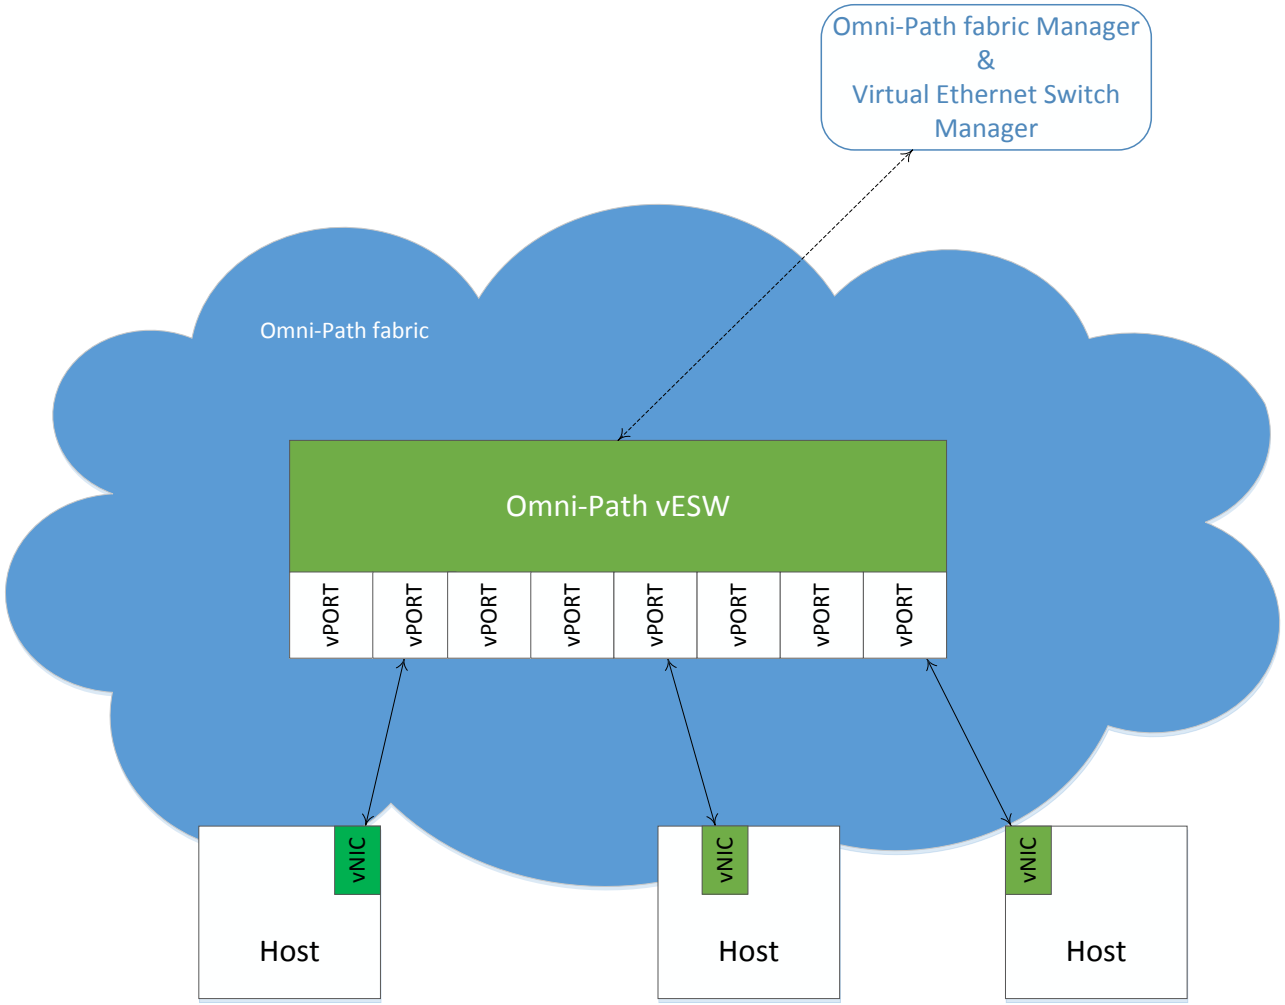

- An Omni-Path virtual Ethernet switch (vESW) is a **logical abstraction** achieved by configuring the hosts on the fabric for header generation and processing

- The configuration is performed by an Ethernet Manager (EM) which is part of the trusted Fabric Manager (FM) application

ARCHITECTURE

Multiple Omni-Path vESW example

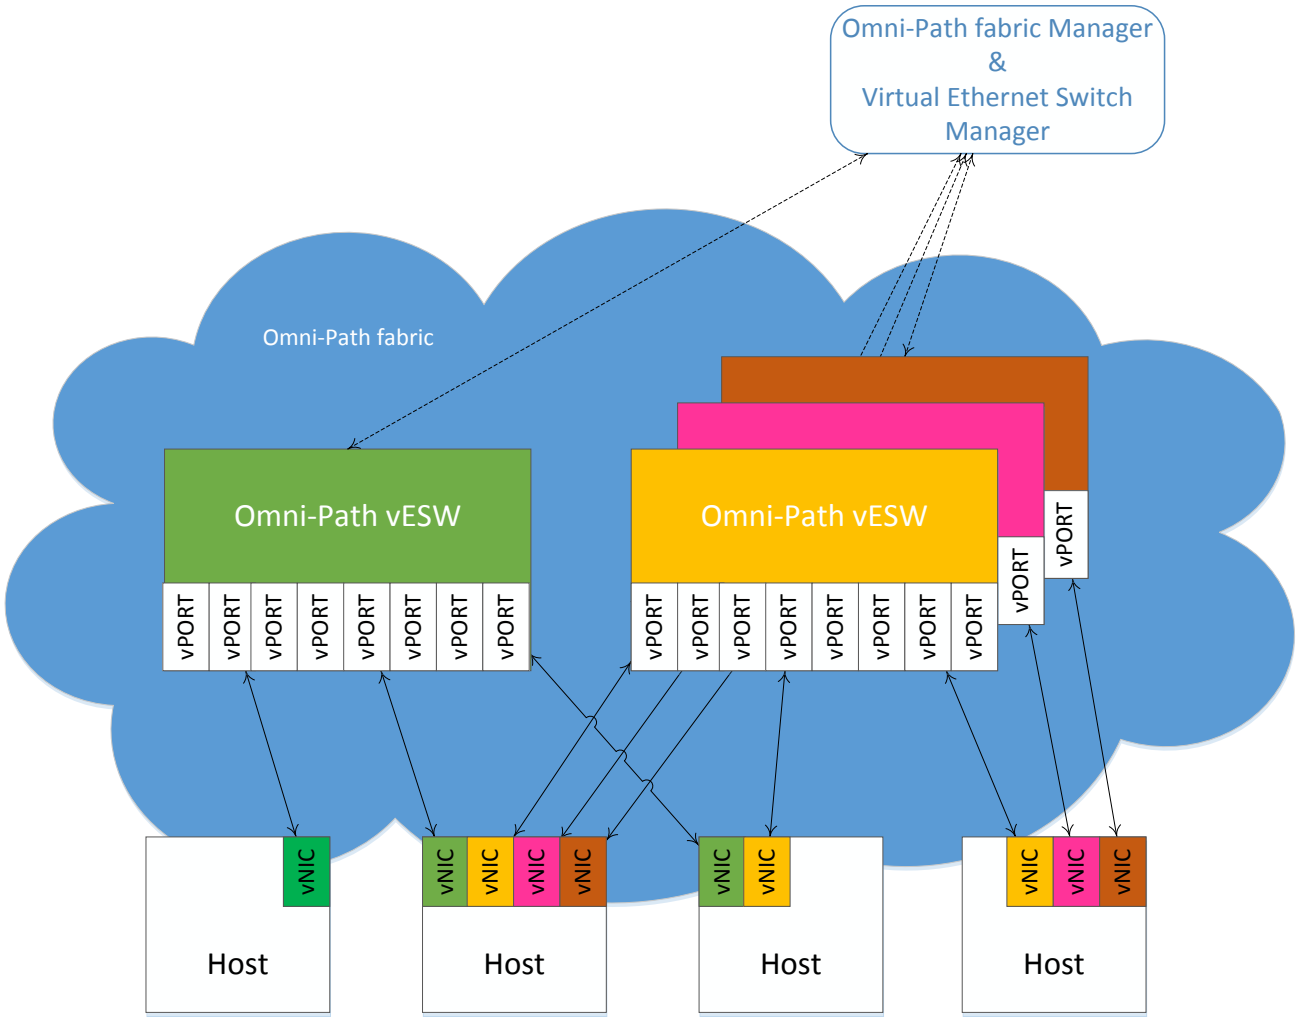

- There can be multiple Omni-Path vESWs in the fabric
- Hosts can have multiple vNICs each connected to a different Omni-Path vESW

ARCHITECTURE

Packet format

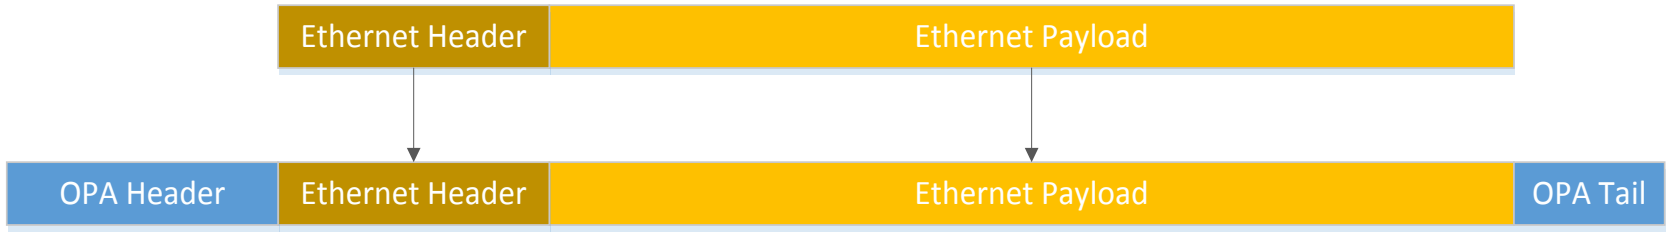

Omni-Path encapsulation of Ethernet Packet

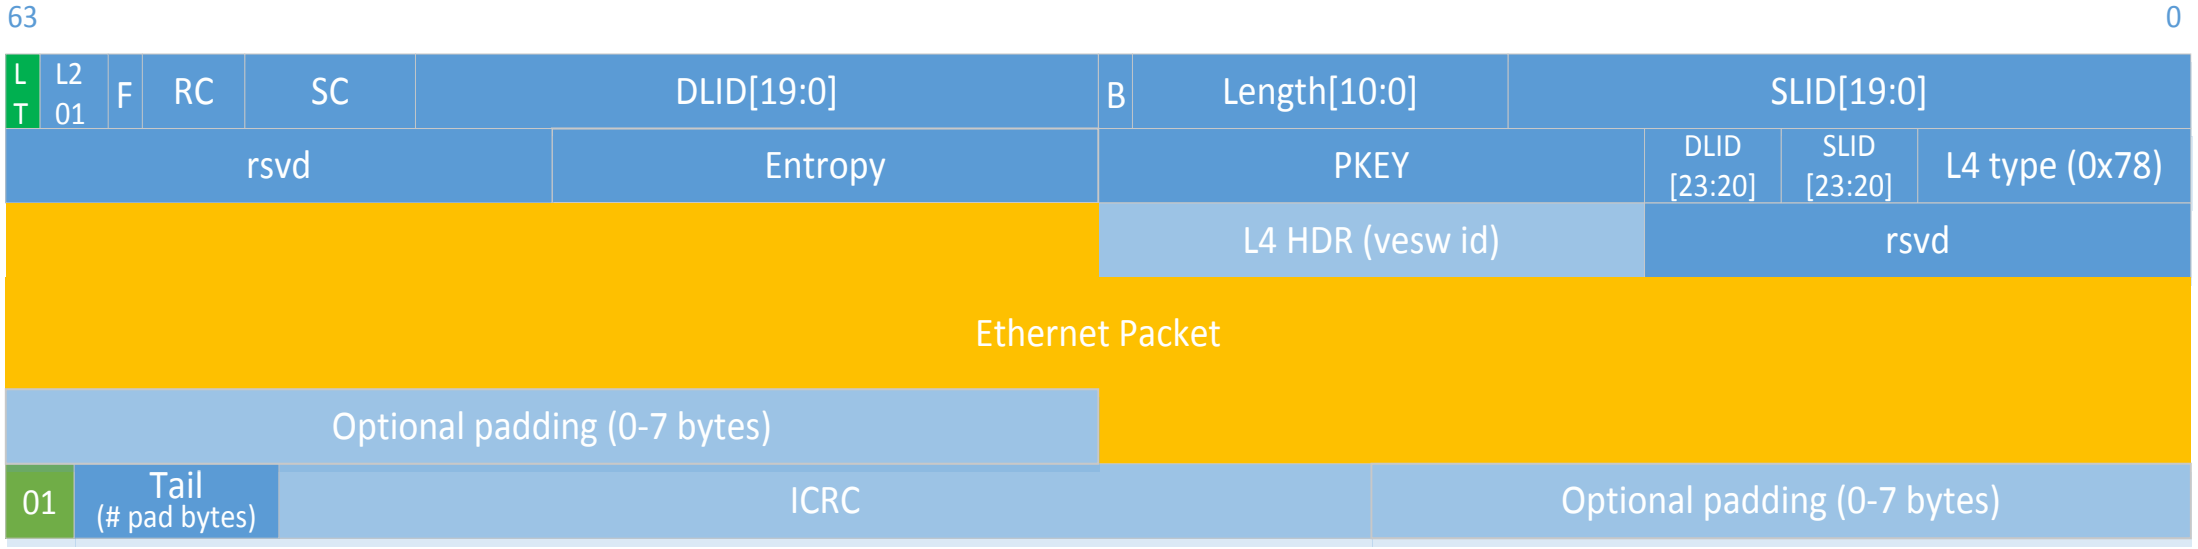

Omni-Path encapsulated Ethernet Packet format

OMNI-PATH VNIC DRIVER DESIGN

DRIVER DESIGN

Omni-Path VNIC SW stack

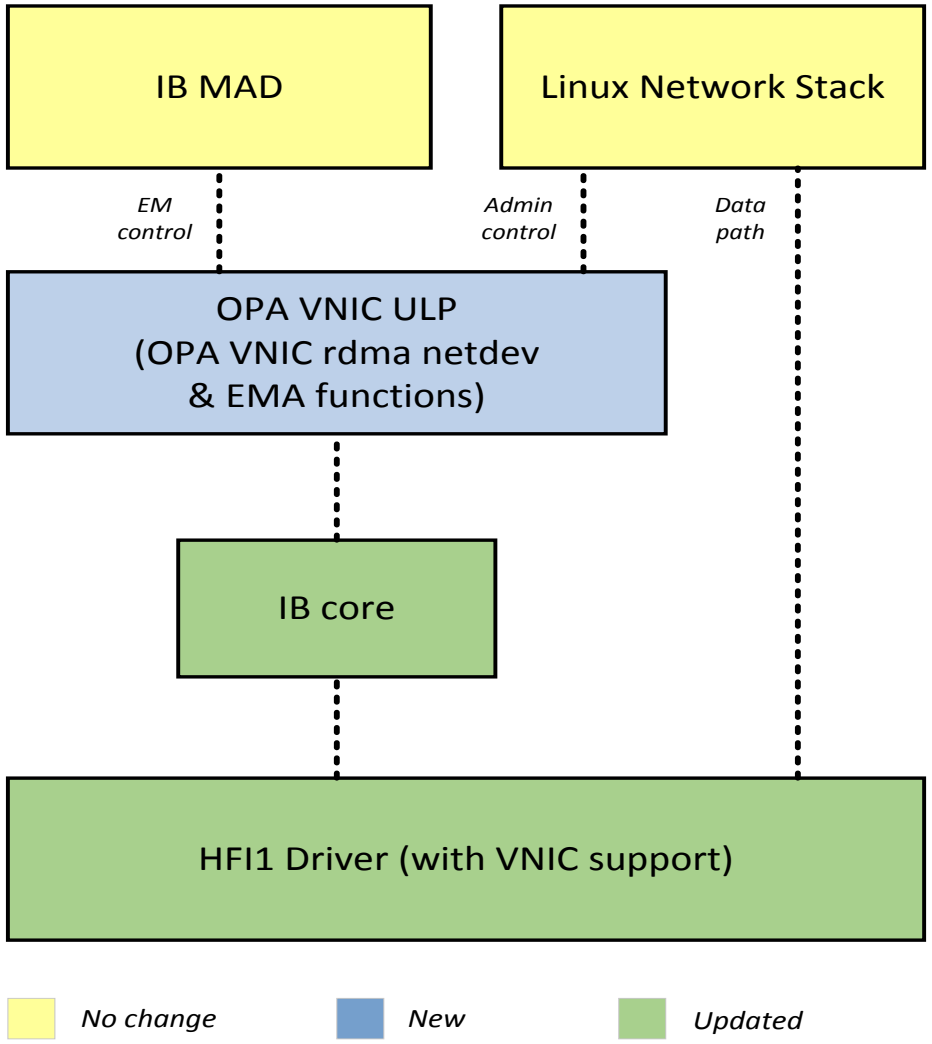

- OPA VNIC ULP is an '*ib_client*'
- EMA is an '*ib_mad_agent*'
- Linux network stack's SKB interface is used and no translation to verbs API required
- HW Driver (HFI1) defines '*net_device_ops*' and can interact directly with the network stack on data path
- OPA VNIC module can override the '*net_device_ops*' defined by HW driver to implement control plane operations and encapsulation

DRIVER DESIGN

rdma netdev

- Requirements
  - Allow OFA device drivers to interface directly with Linux network stack
  - No translation to Verbs Interface required thus providing optimization
- '**rdma netdev**' - A generic netdev interface to OFA device drivers where Linux network stack interfacing is required
- Ability to support different kind of rdma netdev devices
- Address OPA_VNIC and IPoIB use case requirements
- Not adding any overhead on the data path

DRIVER DESIGN

rdma netdev (Omni-Path VNIC example)

```
/**
 * struct rdma_netdev - rdma netdev
 * For cases where netstack interfacing is required.
 */
struct rdma_netdev {
        void                    *clnt_priv;
        struct ib_device  *ibdev;
        u8                       port_num;

        /* control functions */
        void (*set_id)(struct net_device *netdev, int id);
};

/* rdma netdev type - specifies protocol type */
enum rdma_netdev_t {
        RDMA_NETDEV_OPA_VNIC
};

struct ib_device {
        …
        /* rdma netdev operations */
        struct net_device *(*alloc_rdma_netdev)(struct ib_device *device, u8 port_num,
                                                enum rdma_netdev_t type, const char *name,
                                                unsigned char name_assign_type, void (*setup)(struct net_device *));
        void (*free_rdma_netdev)(struct net_device *netdev);
        …
}
```

```
/* opa vnic rdma netdev's private data structure */
struct opa_vnic_rdma_netdev {
        struct rdma_netdev rn;  /* keep this first */

        /* followed by device private data */
        char *dev_priv[0];
};

/* Get ULP's (OPA_VNIC) private data */
static inline void *opa_vnic_priv(const struct net_device *dev)
{
        struct rdma_netdev *rn = netdev_priv(dev);

        return rn->clnt_priv;
}

/* Get driver's (HFI1's VNIC) private data */
static inline void *opa_vnic_dev_priv(const struct net_device *dev)
{
        struct opa_vnic_rdma_netdev *opa_rn = netdev_priv(dev);

        return opa_rn->dev_priv;
}
```

DRIVER DESIGN

Omni-Path VNIC ULP

- Implements required netdev control operations. Allocates rdma netdev and registers netdev with network stack.
- Does OPA encapsulation of Ethernet packets
- Implements EMA IB MAD agent to interact with EM
- Implements Ethtool interface

- **EM Interface**
  - Attributes:
    - CLASS_PORT_INFO
    - VESWPORT_INFO
    - VESWPORT_MAC_ENTRIES
    - IFACE_UCAST_MACS
    - IFACE_MCAST_MACS
    - DELETE_VESW
    - VESWPORT_SUMMARY_COUNTERS
    - VESWPORT_ERROR_COUNTERS
  - Traps:
    - IFACE_UCAST_MAC_CHANGE
    - IFACE_MCAST_MAC_CHANGE
    - ETH_LINK_STATUS_CHANGE

DRIVER DESIGN

HFI1 VNIC support

- HW resource management for VNIC traffic
  - Allocates and frees receive contexts
  - Implements RSS using HFI1 RSM engine
- Implements TX path
  - uses hfi1 SDMA engines
  - supports multiple TX queues (VL based)
  - supports TX queue halt and wakeup
- Implements the Rx path
  - Implements multiple Rx queues (RSM)
  - Implements NAPI interface
- Implements VNIC statistics support
  - Supports standard netdev and rmon counters
  - Supports EM defined counters

DRIVER DESIGN

Queue Mapping

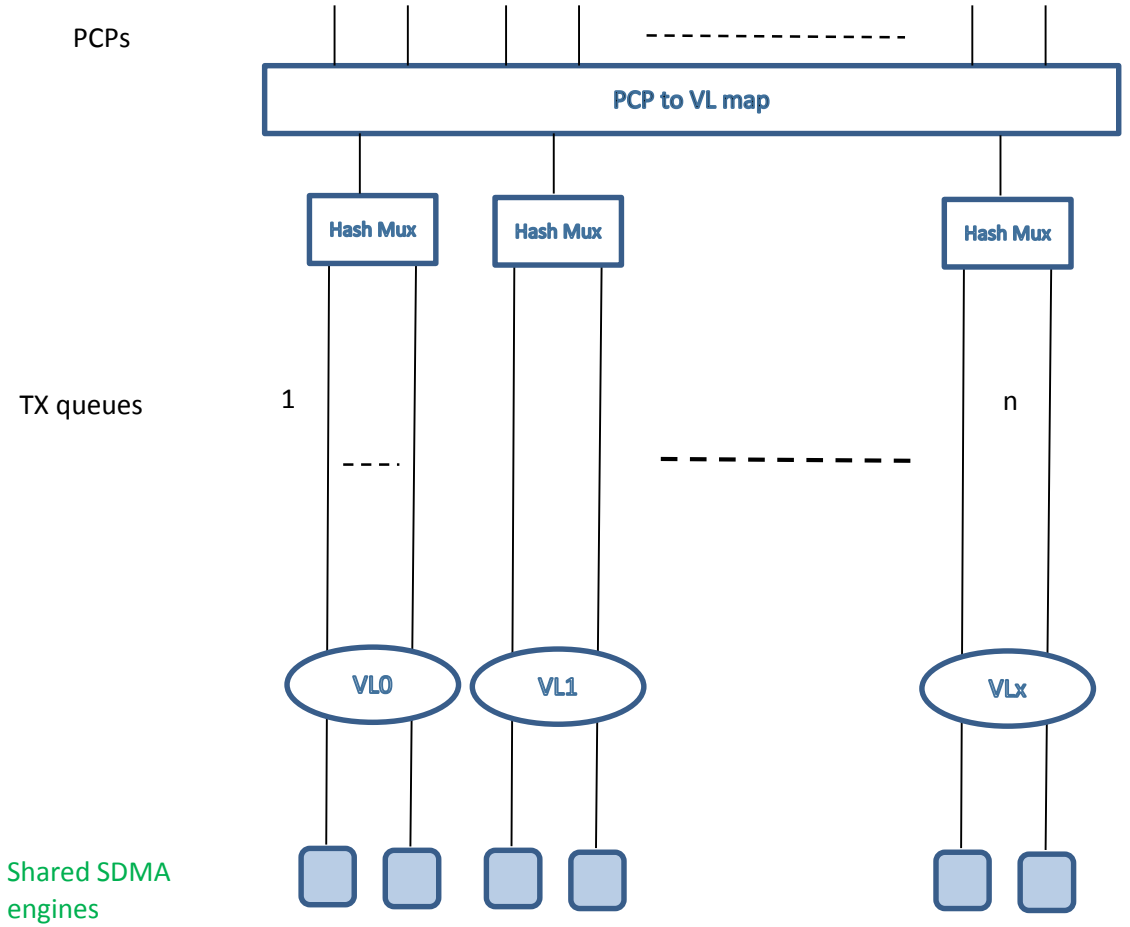

TX queue mapping

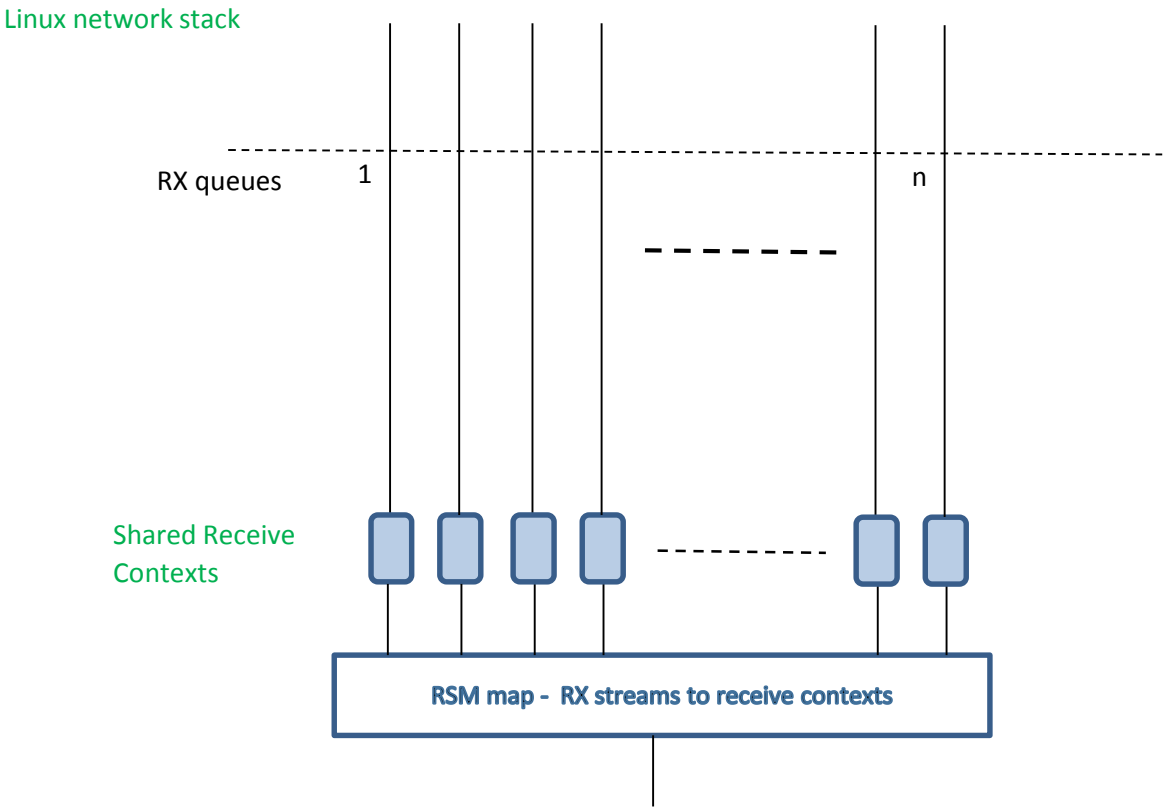

RX queue mapping

STATUS & NEXT STEPS

Status:

- **Currently the OPA_VNIC patch series is posted on LKML**
  - https://www.spinics.net/lists/linux-rdma/msg46604.html

Next Steps:

- **RDMACM address resolution using VNIC interface**
  - Currently exploring options for RDMACM to use VNIC interface (instead of IPoIB) to translate destination node's IP address to LID

13th ANNUAL WORKSHOP 2017

THANK YOU

Niranjana Vishwanathapura, Software Development Engineer

Intel Corporation