



# OpenStack and OVS: From Love-Hate Relationship to Match Made in Heaven

Chloe Jian Ma

Senior Director, Cloud  
Market Development

Erez Cohen

Senior Director,  
CloudX Program



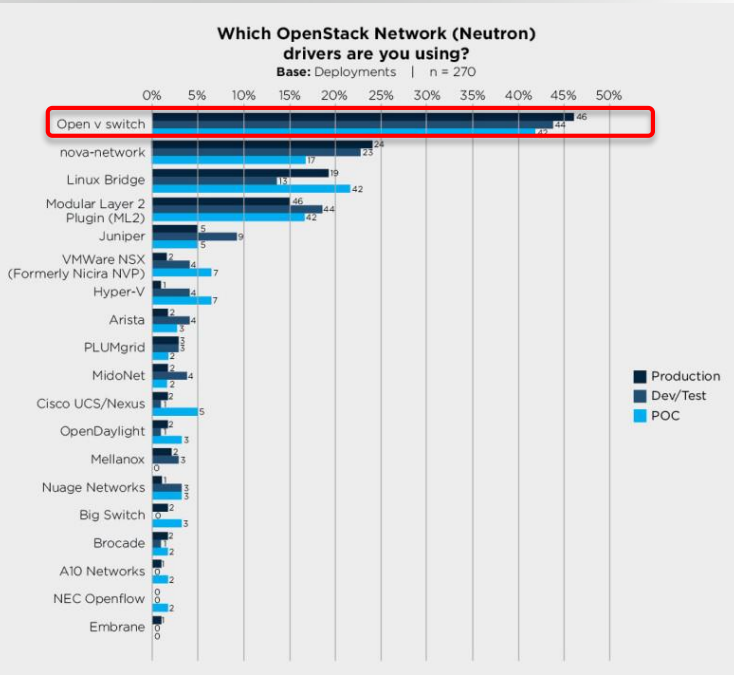
# What Are We Covering in this Session

- Challenges with Using OVS for OpenStack Networking
- Mellanox OVS Offload Overview
- Demo!

# What Are We Covering in this Session

- Challenges with Using OVS for OpenStack Networking
- Mellanox OVS Offload Overview
- Demo!

# OpenStack and OVS: A Love-Hate Relationship

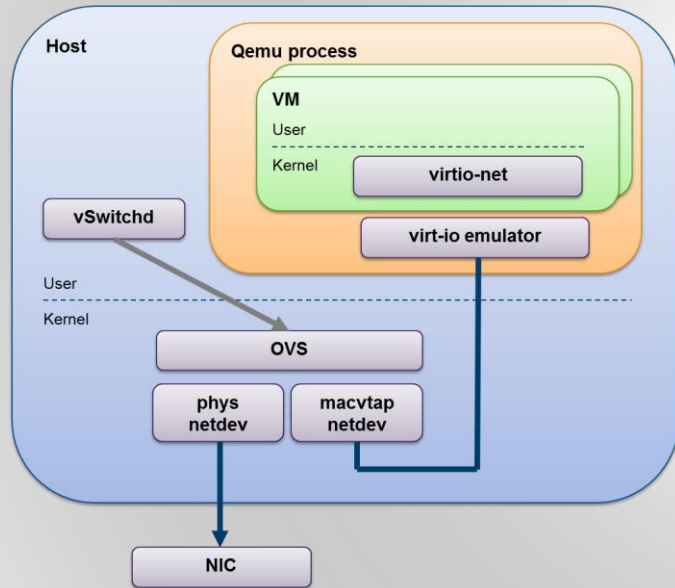


- Man, It is **SLOW!**
- What do you mean it **drops my packets?**
- It **burns CPU** like there is no tomorrow!

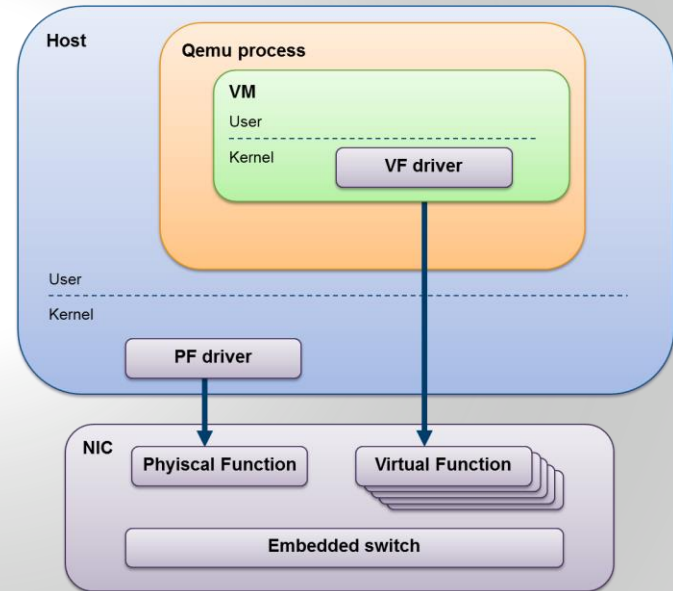


# Comparison of Existing I/O Virtualization Solutions

- Paravirt - Control



- SRIOV - Performance



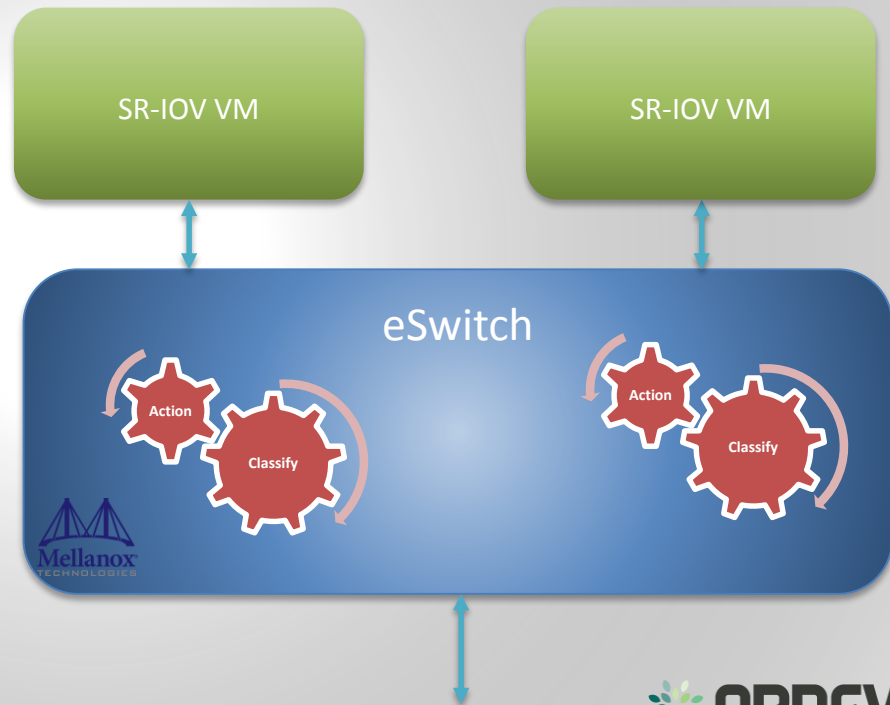
# What If We Could Enjoy the Best of Both Worlds?

# What Are We Covering in this Session

- Challenges with Using OVS for OpenStack Networking
- Mellanox OVS Offload Overview
- Demo!

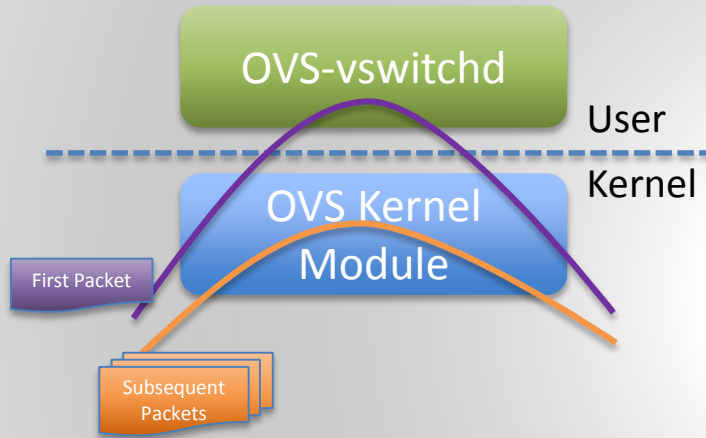
# Mellanox Embedded Switch (eSwitch)

- Advanced flow-based switch
- Sophisticated classification engines
- Multiple actions supported including:
  - Steering and Forwarding
  - Drop / Allow
  - Encap/Decap





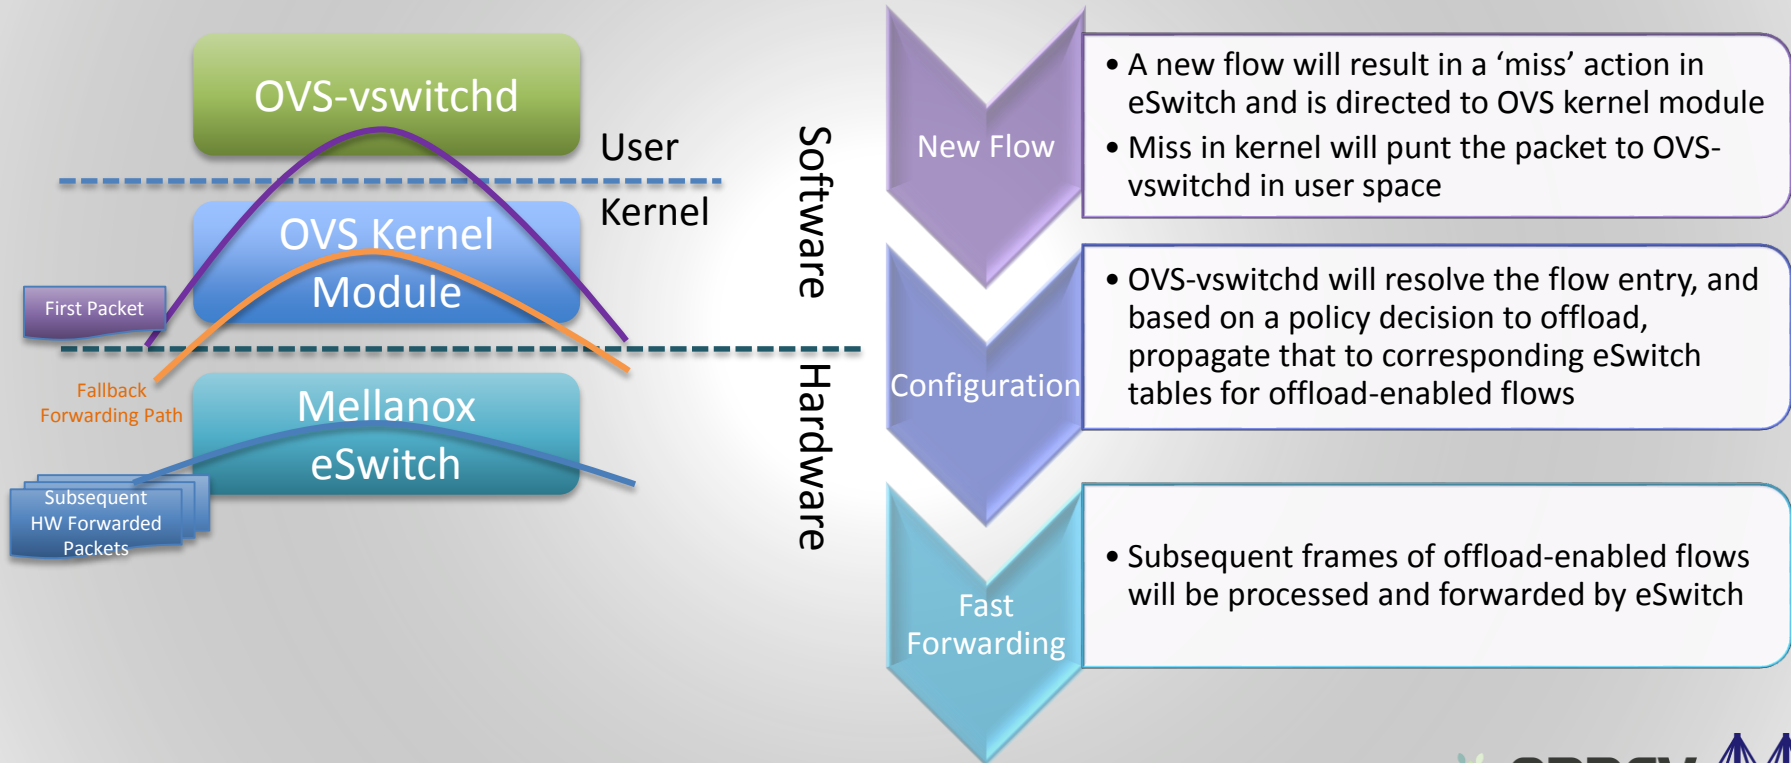
# OVS Architecture and Operations



- Forwarding

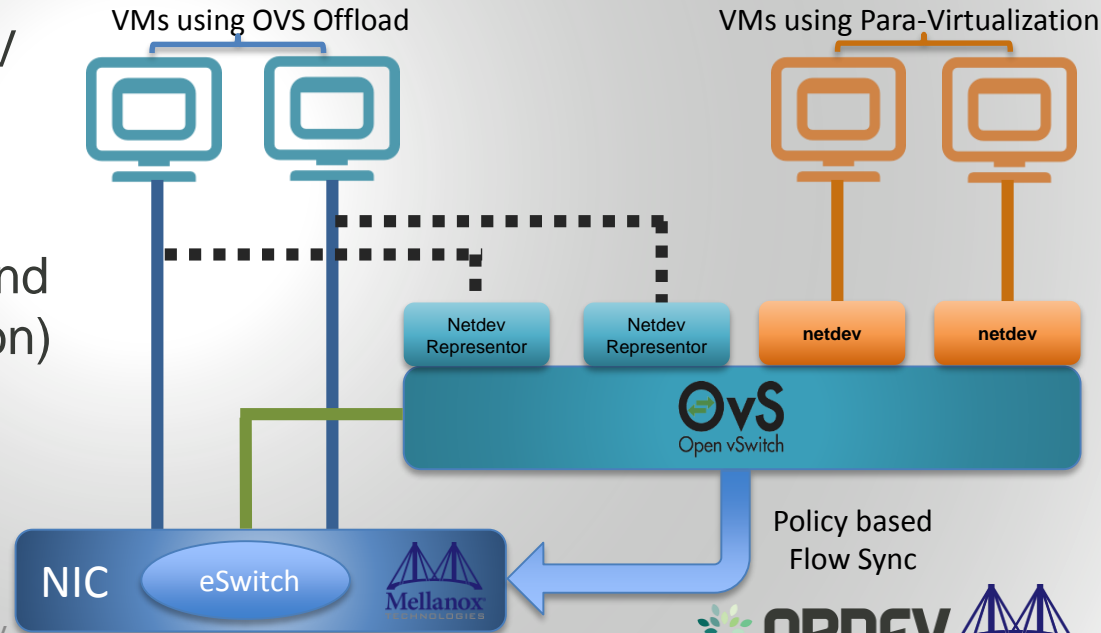
- Flow-based forwarding
- First packet of a new flow (match miss) is directed to user space (ovs-vswitchd)
- ovs-vswitchd determines flow handling and programs kernel (fast path)
- Following packets hit kernel flow entries and are executed in fast path

# OVS Offload – Let the Hardware Do the Heavy-lifting



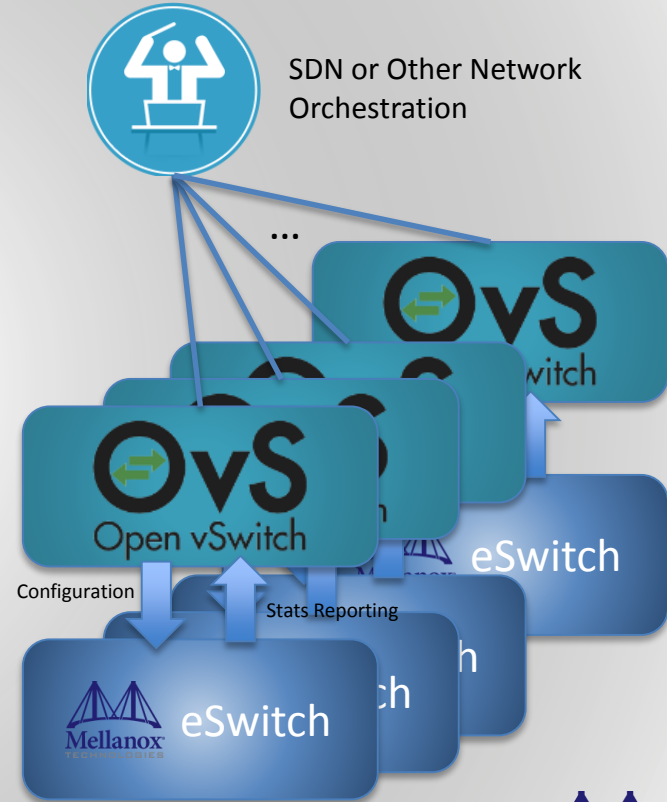
# OVS and SRIOV, Isn't it Oil and Water?

- Representer ports enable OVS to “know” and service those VMs that uses SR-IOV
- Representer ports are used for eSwitch / OVS communication (miss flow and PV to SR-IOV communication)



# Software Defined Networking, at Full Speed

- Leverage Open vSwitch control-plane and Software Defined Networks (SDN) capabilities to control eSwitch forwarding-plane
- Enhance forwarding performance while maintaining network programmability
- Benefits:
  - Open vSwitch interfaces to the user remain untouched
    - The hardware offloads are transparent to the user
  - User does not need changes in his environment



# Key OVS Offload Capabilities

Enable/Disable OVS Offload on a per flow basis.



OVS Control Path with SRIOV performance

Classification and steering offload



Support “Mega-Flow” (Wild cards).

# Conclusion

- Key OVS Offload Benefits
  - Highest performance (Offload is increasingly important as server I/O speed goes up)
  - Low CPU overhead, higher infrastructure efficiency
  - Software defined
  - Everything In-Box (All changes will be up-streamed, no proprietary OVS or kernel patches)

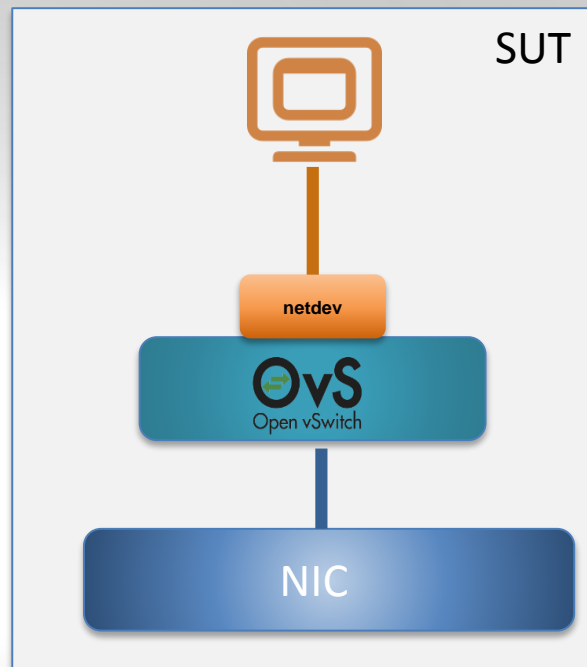
# What Are We Covering in this Session

- Challenges with Using OVS for OpenStack Networking
- Mellanox OVS Offload Overview
- Demo!



## Demo – OVS PV

- Measure 64B packet rate
- Measure Network CPU load

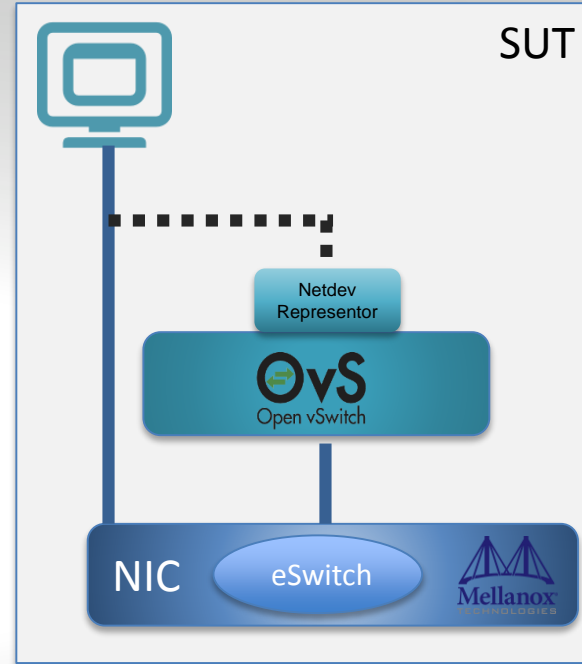


**ixia**



## Demo – VS. OVS SRIOV

- Measure 64B packet rate
- Measure Network CPU load





THANK YOU

Q&A