Data Plane Optimization in Open Virtual Routers

ROYAL INSTITUTE
OF TECHNOLOGY

M.Siraj Rathore
siraj@kth.se

Markus Hidell
mahidell@kth.se

Peter Sjödin
psj@kth.se

ROYAL INSTITUTE OF TECHNOLOGY

Outline

- PC based Virtual Routers
- Challenges
- Virtual Router Design
- Data plane Optimization
- Performance Evaluation
- Conclusion

ROYAL INSTITUTE OF TECHNOLOGY

Open Virtual Routers

- Commodity hardware, Open source softwares
- Run multiple independent virtual instances in parallel on the same hardware
- A virtualization technology enforces resource limiting among virtual routers
- Each virtual router maintains its own set of virtual network interfaces, protocols, routing tables, packet filtering rules

ROYAL INSTITUTE OF TECHNOLOGY

Challenge

- Router virtualization is associate with performance penalties
- Virtualization overhead is introduced in terms of how packets are processed in the router
- How to combine software modules to form an open virtual router with minimum virtualization penalty

ROYAL INSTITUTE OF TECHNOLOGY

Linux Virtual Routers

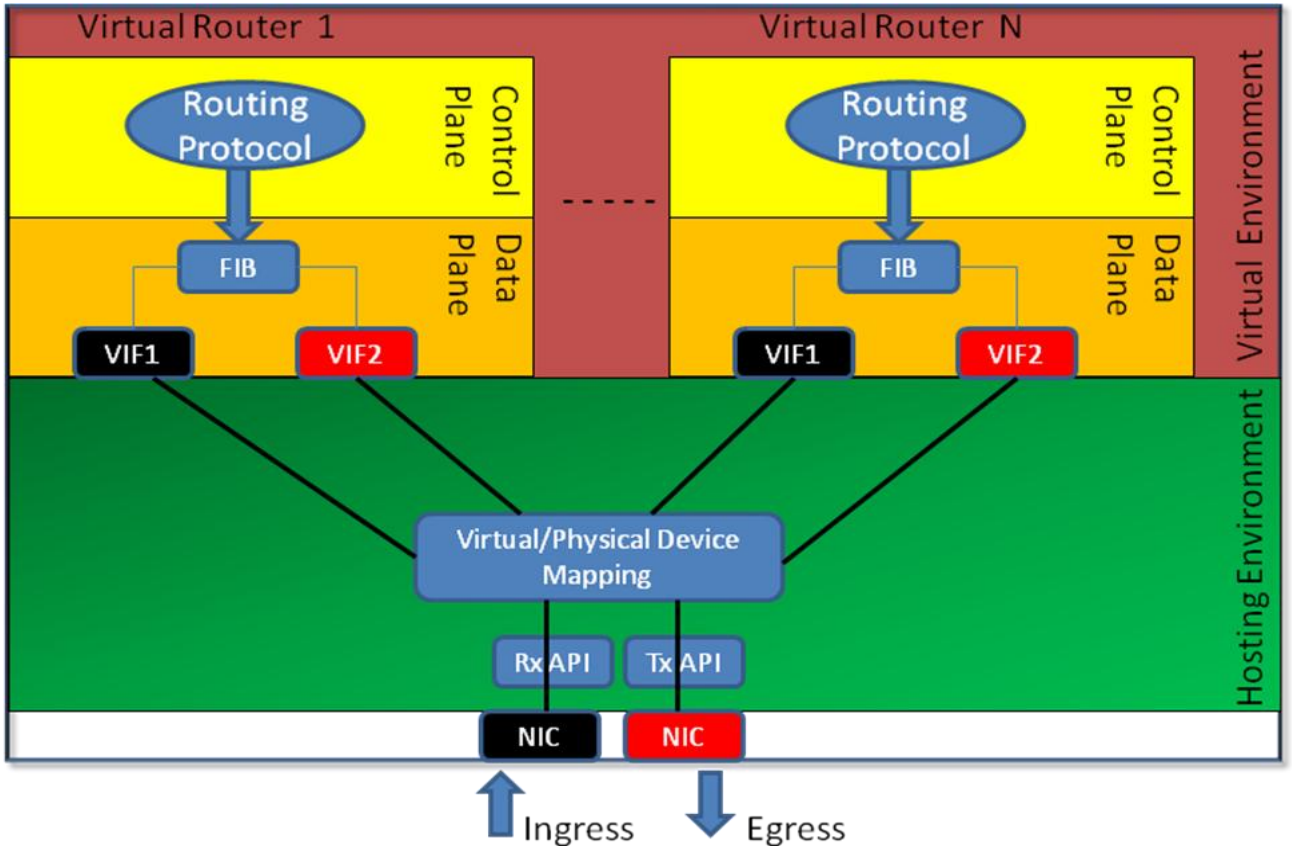

ROYAL INSTITUTE OF TECHNOLOGY

Virtualization Technologies

- Hypervisor: It runs on top of the physical hardware and it virtualizes hardware resources to be shard among multiple guest operating systems

  E.g. VMware, Xen

- Container: The operating system resources are virtualized (e.g. files, system libraries) to create multiple isolated execution environment on top of a single operating system.

  E.g. OpenVZ, Linux Namespaces

ROYAL INSTITUTE OF TECHNOLOGY

Virtualization Technologies

- Macvlan, a virtual network device operating at layer 2 (with its own MAC address). It also provides a mechanism of physical/virtual device mapping

ROYAL INSTITUTE OF TECHNOLOGY

Data Plane of macvlan based Virtual Router

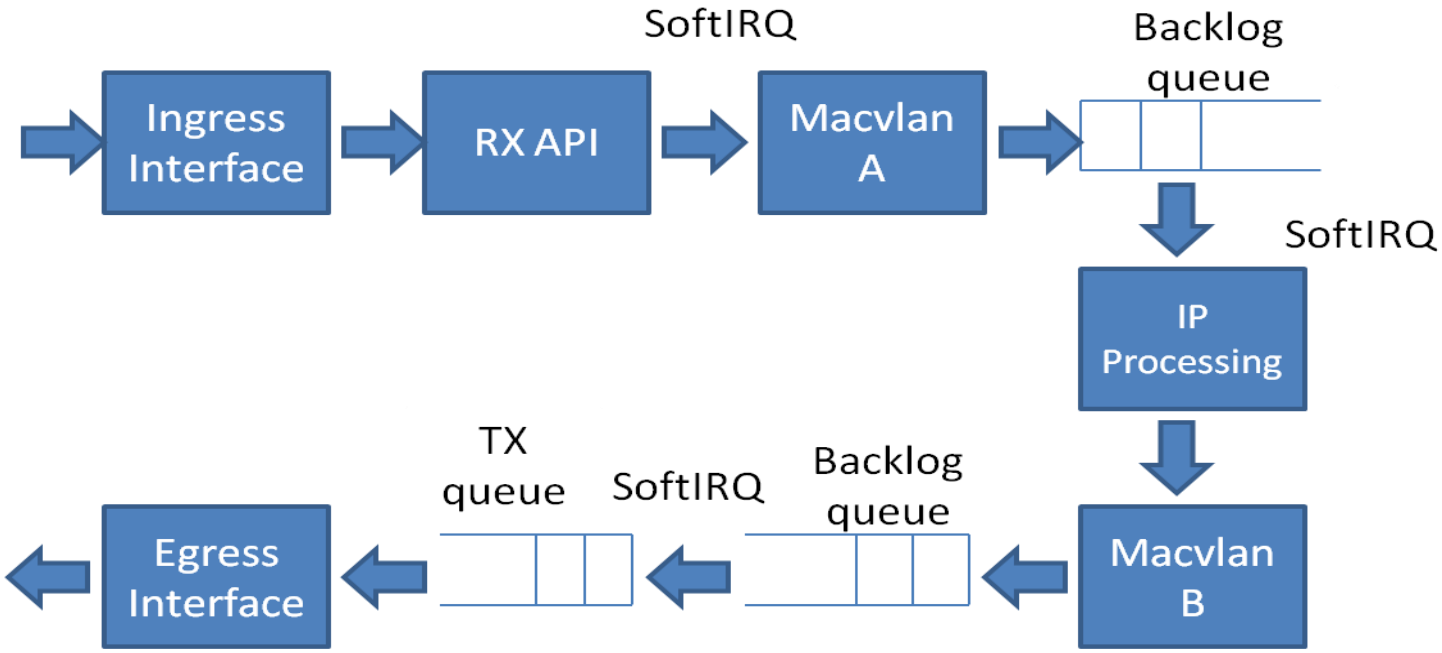

ROYAL INSTITUTE OF TECHNOLOGY

Data Plane of macvlan based Virtual Router

- Backlog queue congestion is observed under high loads
- The backlog queue is maintained on a per CPU basis, each virtual router running on the same CPU will share the queue. This may result in poor traffic isolation
- Multiple queuing points along the data path may cause unnecessary delays in packet processing, and lead to inefficient usage of CPU

ROYAL INSTITUTE OF TECHNOLOGY

Data Plane of NAPI-macvlan based Virtual Router

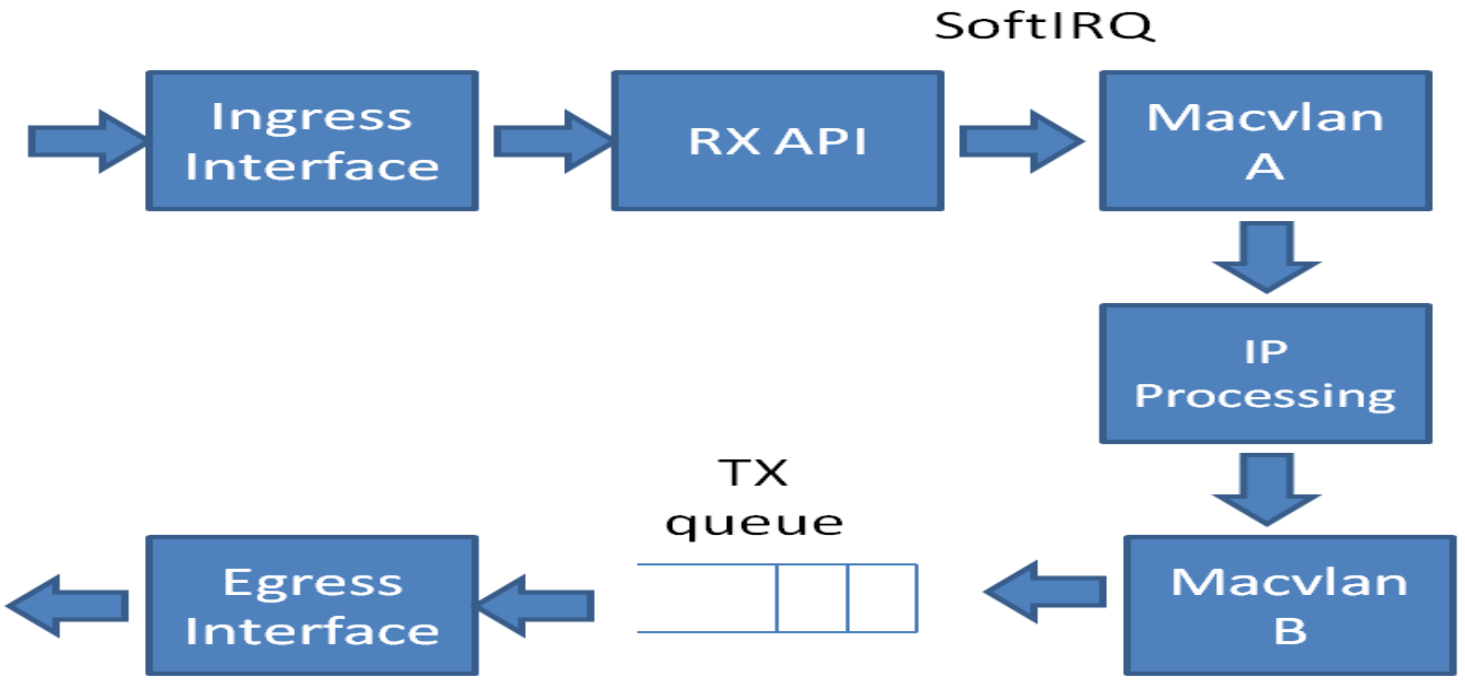

ROYAL INSTITUTE OF TECHNOLOGY

Throughput result with two interfaces (Case I)

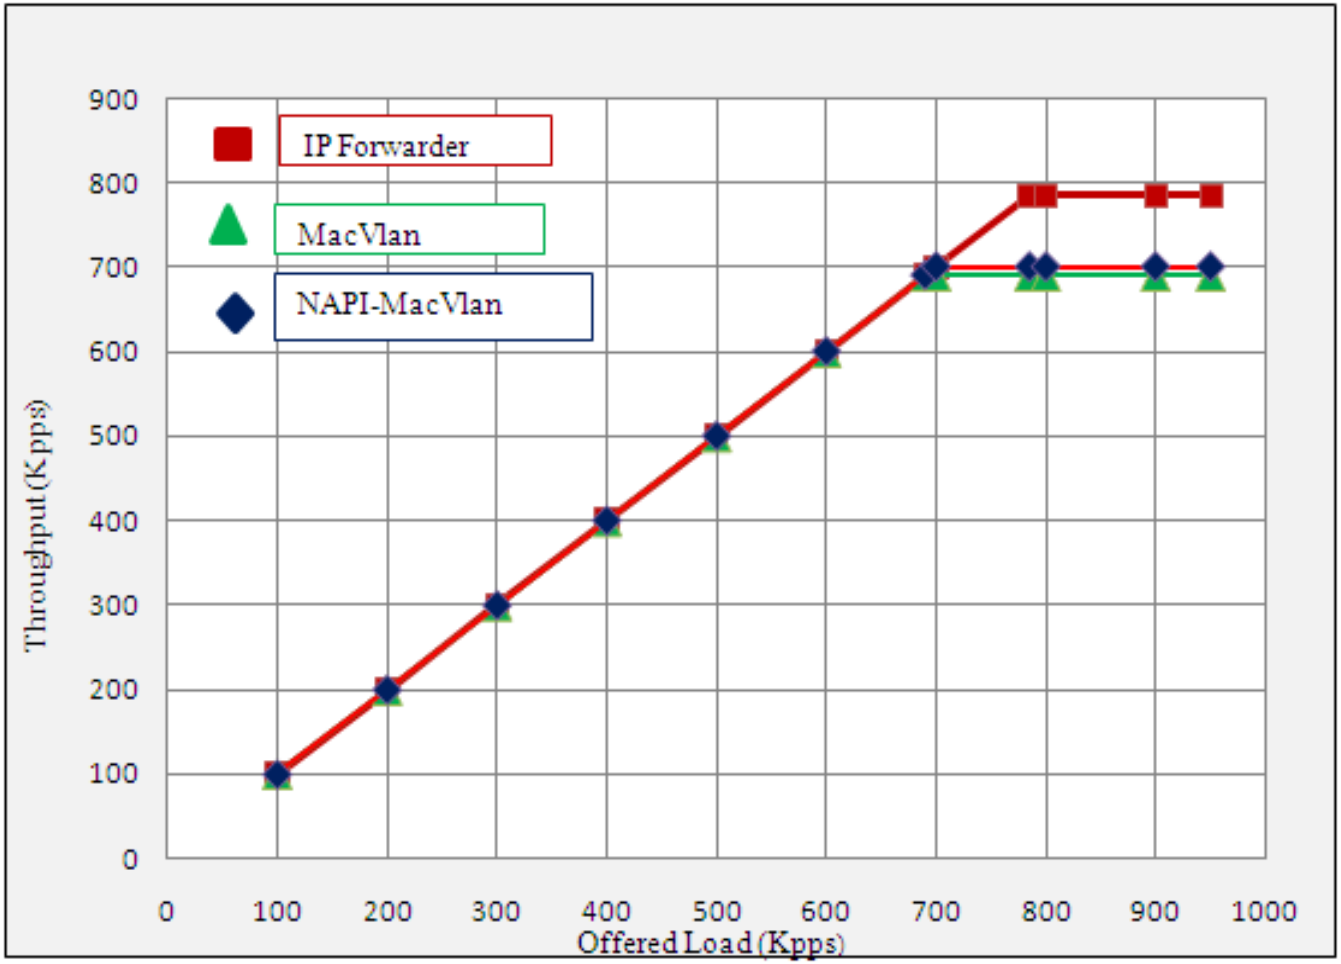

ROYAL INSTITUTE OF TECHNOLOGY

Throughput results: three (Case II) and four interfaces (Case III)

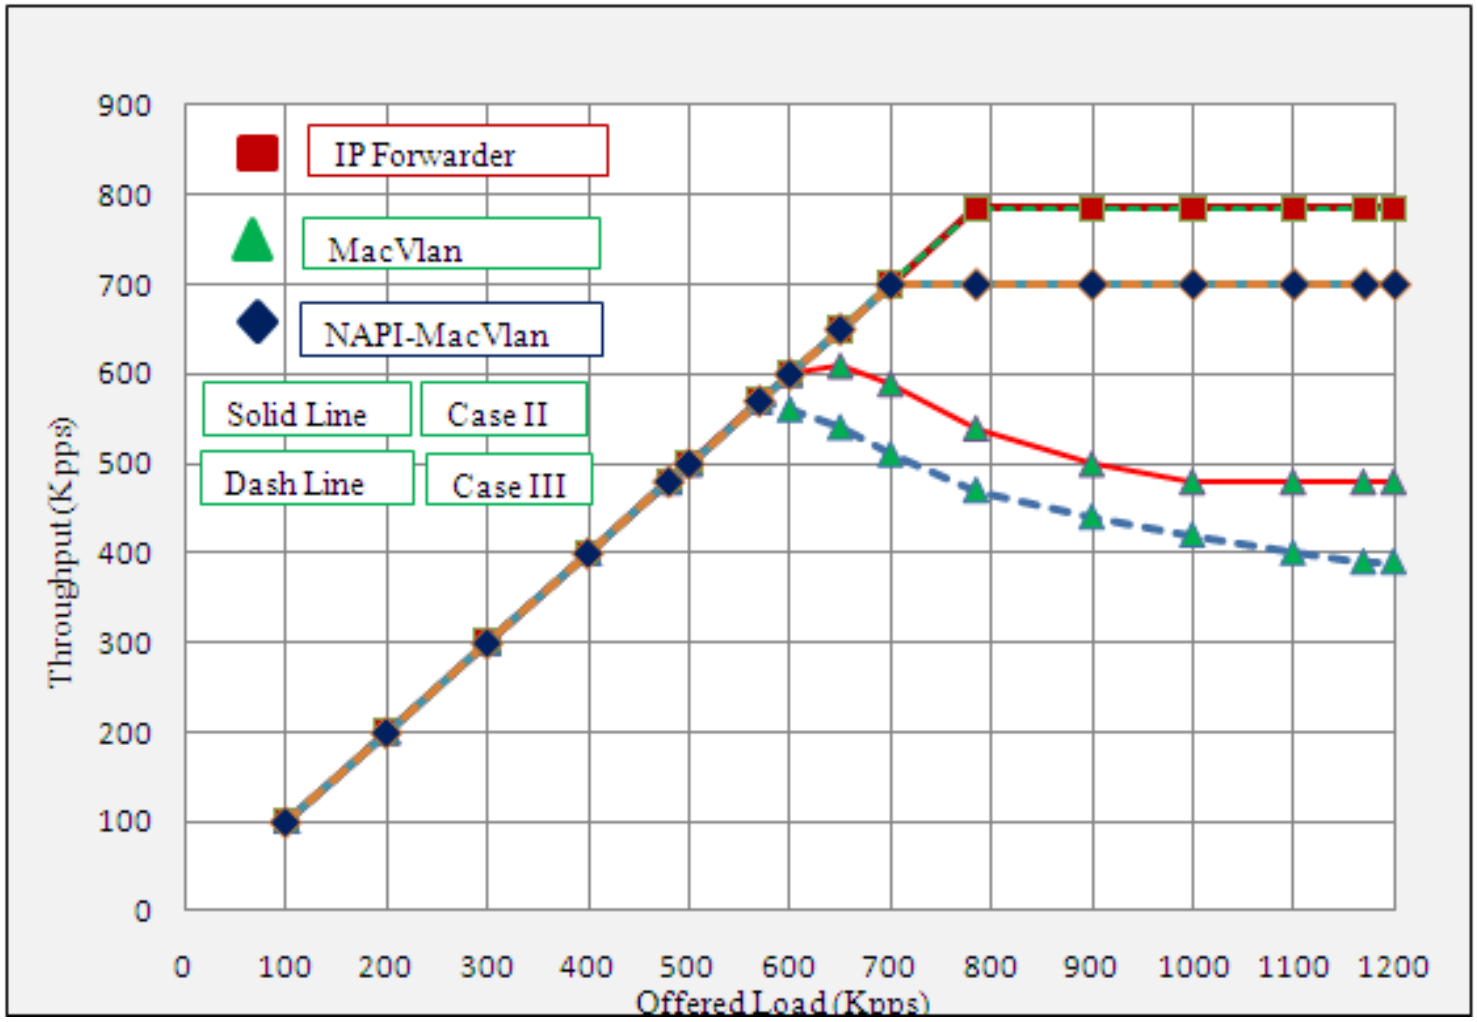

ROYAL INSTITUTE OF TECHNOLOGY

Total number of interrupts under high load (900kpps)

Setup	Case	Interrupts (HW)
Macvlan	I	104322
Macvlan	II	869406
Macvlan	III	1633493
NAPI-Macvlan	I	10221
NAPI-Macvlan	II	16336
NAPI Macvlan	III	82626

ROYAL INSTITUTE OF TECHNOLOGY

Latency result

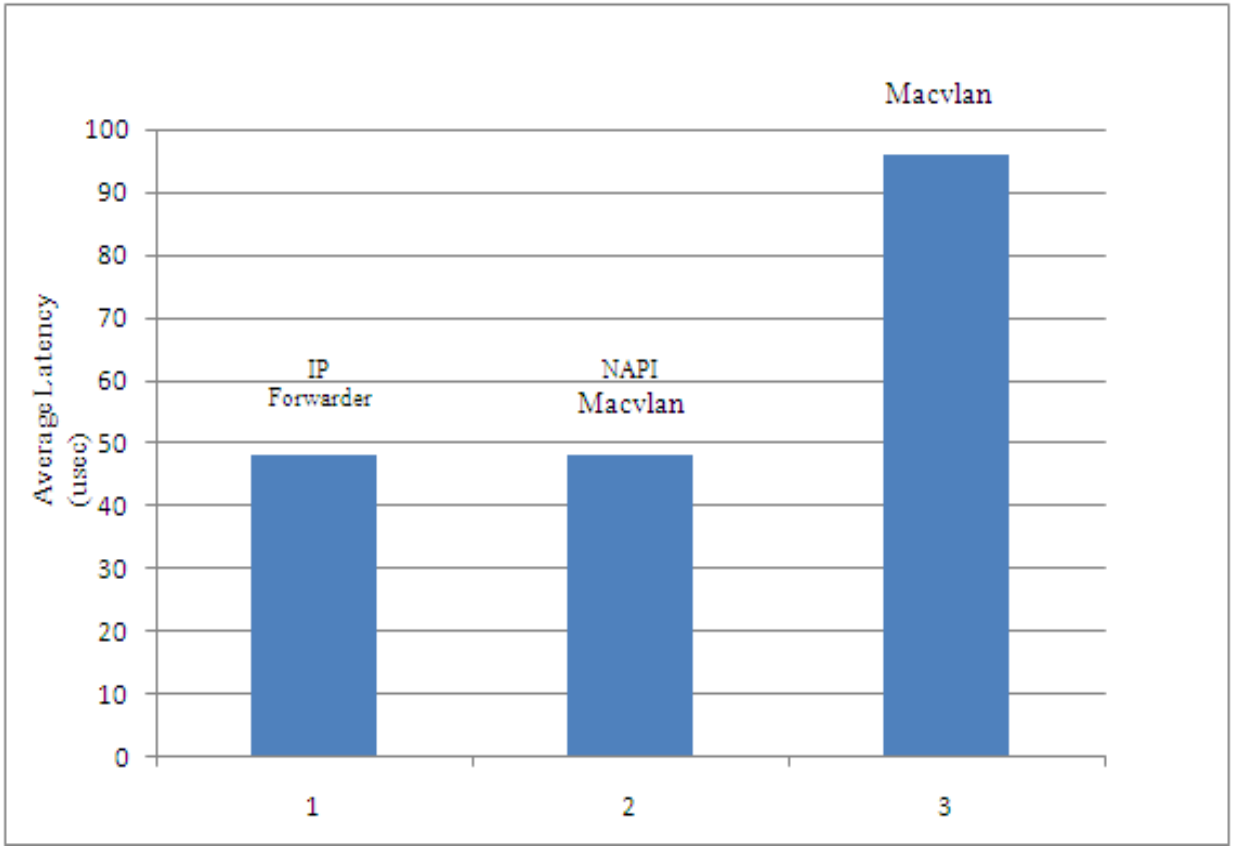

ROYAL INSTITUTE OF TECHNOLOGY

Traffic isolation between virtual routers

Setup	Packet rate (kpps)					
	Offered load			Throughput		
	VR1	VR2	Total	VR1	VR2	Total
Macvlan	1000	300	1300	480	55	535
NAPI-macvlan	1000	300	1300	350	300	650

ROYAL INSTITUTE OF TECHNOLOGY

Conclusion and future work

- It is important to know how virtual devices communicate with kernel
- Backlog queue incurs significant performance penalty
- NAPI-macvlan is an attractive alternative
- NAPI-macvlan scalability is required to be tested on a multi-core architecture

ROYAL INSTITUTE
OF TECHNOLOGY

- Thanks for listening
- Questions ?