

Performance Evaluation of Open Virtual Routers



ROYAL INSTITUTE
OF TECHNOLOGY

M.Siraj Rathore
siraj@kth.se

Outline

- Network Virtualization
 - PC based Virtual Routers
 - Challenges
 - Virtual Router Design
 - Performance Evaluation
 - Conclusion
-

Network Virtualization

- A solution to provide network researchers to run experiments on a shared substrate network
 - Network virtualization means to virtualize all network components (Hosts, Links and Routers)
 - A major challenge is to virtualize the actual network elements, Switches and Routers
-

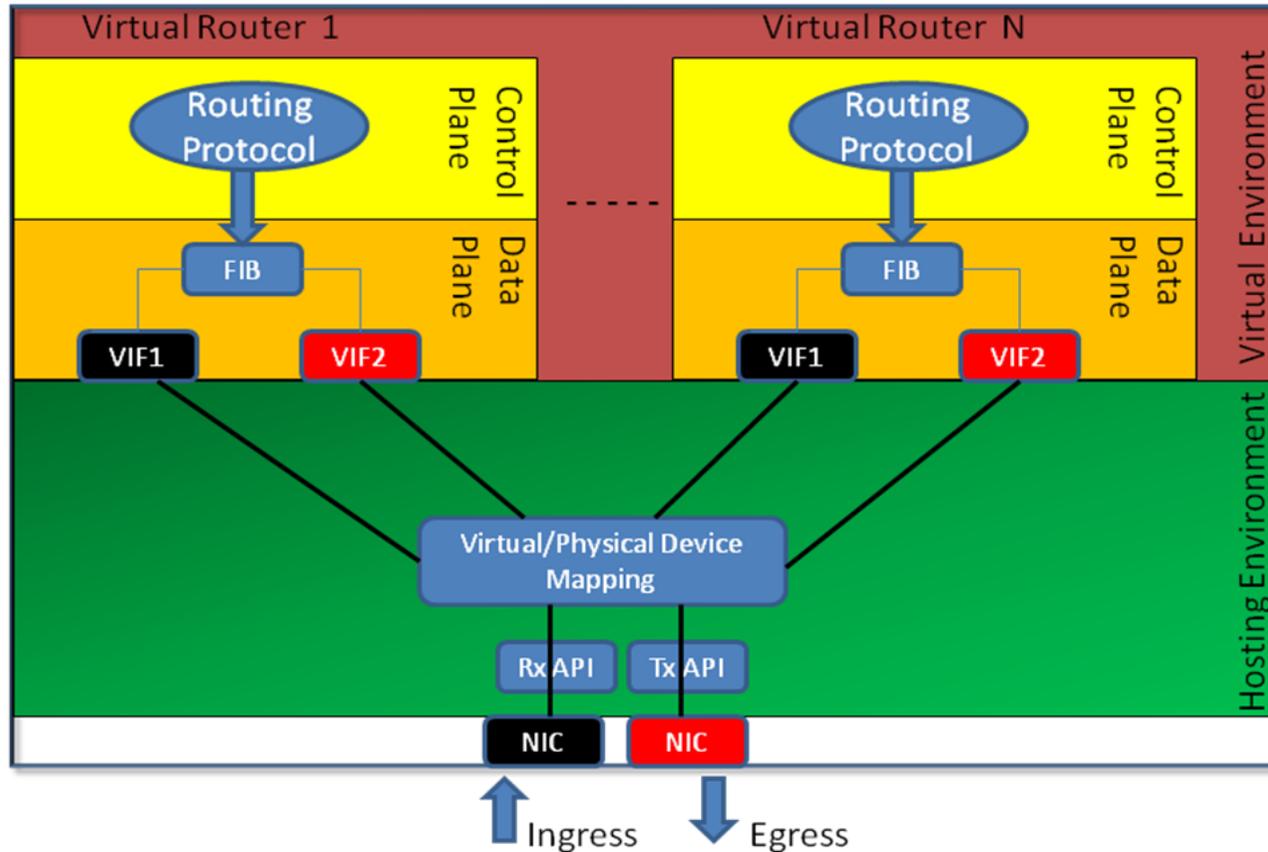
Open Virtual Routers

- 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 (i.e. separate data and control planes)
-

Challenge

- Router virtualization is associated 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
-

Linux Virtual Routers



Virtualization Technologies

- **Hypervisor:** It runs on top of the physical hardware and it virtualizes hardware resources to be shared 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 environments on top of a single operating system.
E.g. OpenVZ, Linux Namespaces
-

OpenVZ based Virtual Routers

- **Virtual devices**

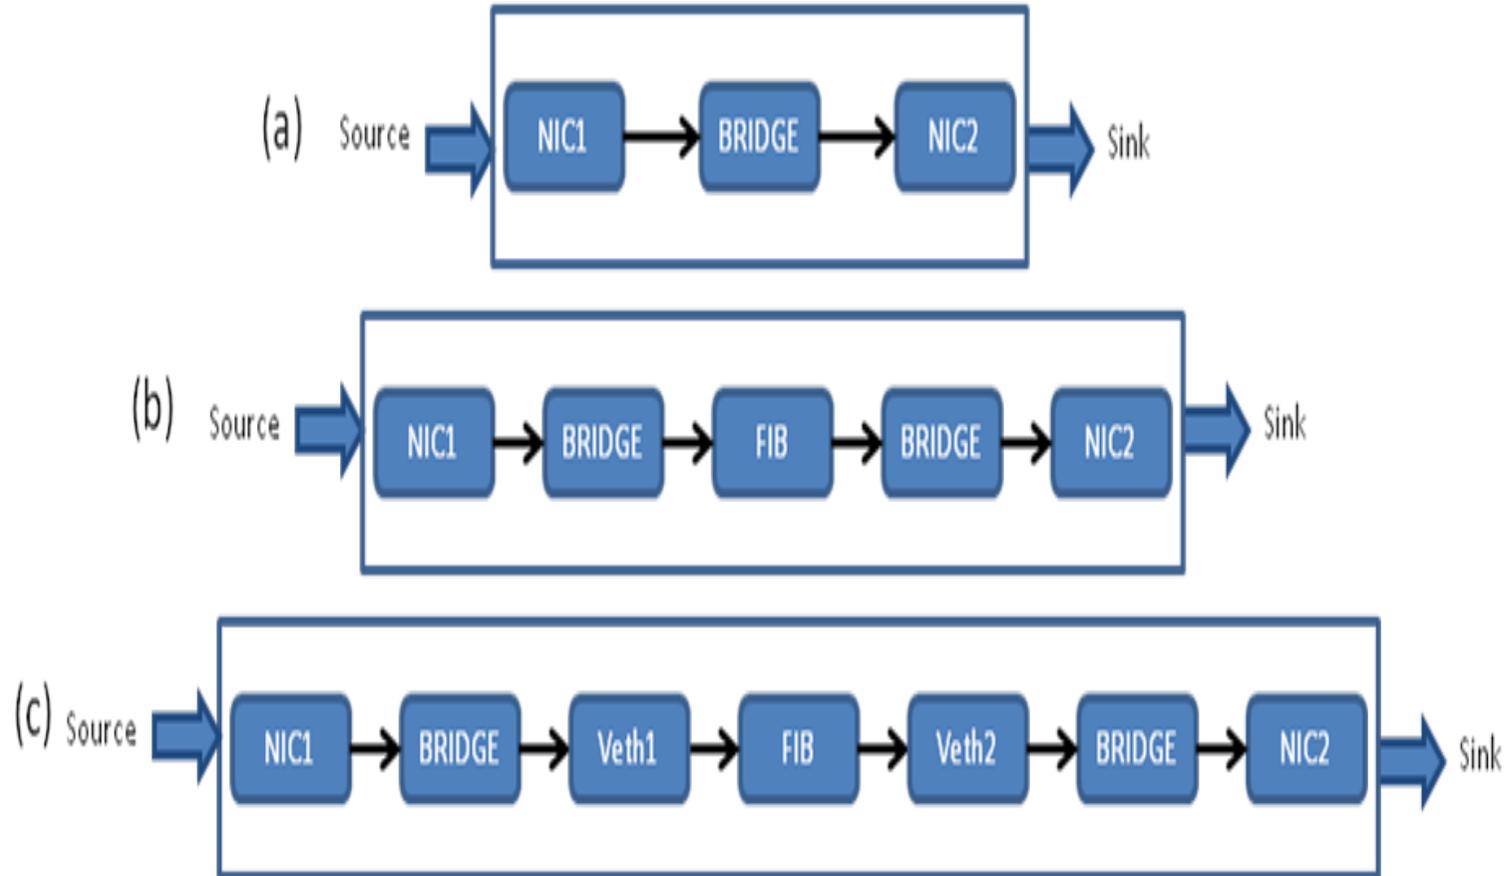
Virtual Network Device (venet): Operates at layer 3.
An IP address is local and unknown from external networks

Virtual Ethernet Device(veth): Ethernet-like device operating at layer 2 with its own MAC address

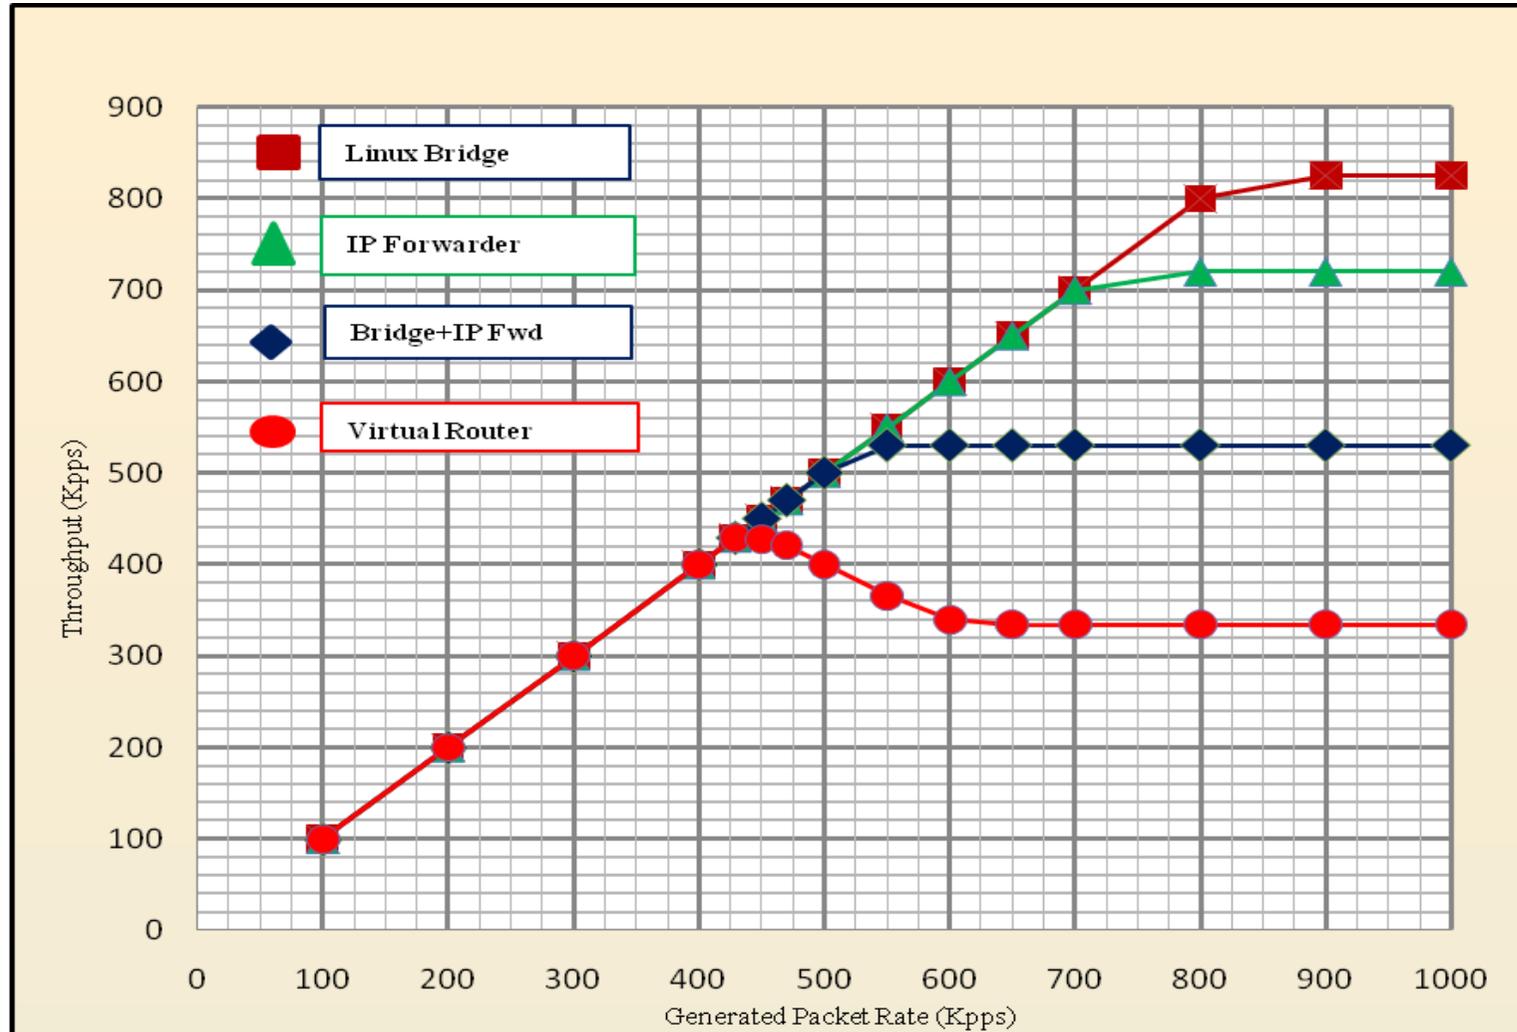
- **Physical/virtual device mapping**

Linux software bridge, IP forwarding, Virtual switch etc.

Building a Virtual Router: 3 step process



Impact of adding virtual components



IP Forwarder vs. Virtual Router

- **IP Forwarder**

Throughput: 720kpps

Packet drop: Ingress physical interface, CPU saturation observed at the offered load of 720kpps

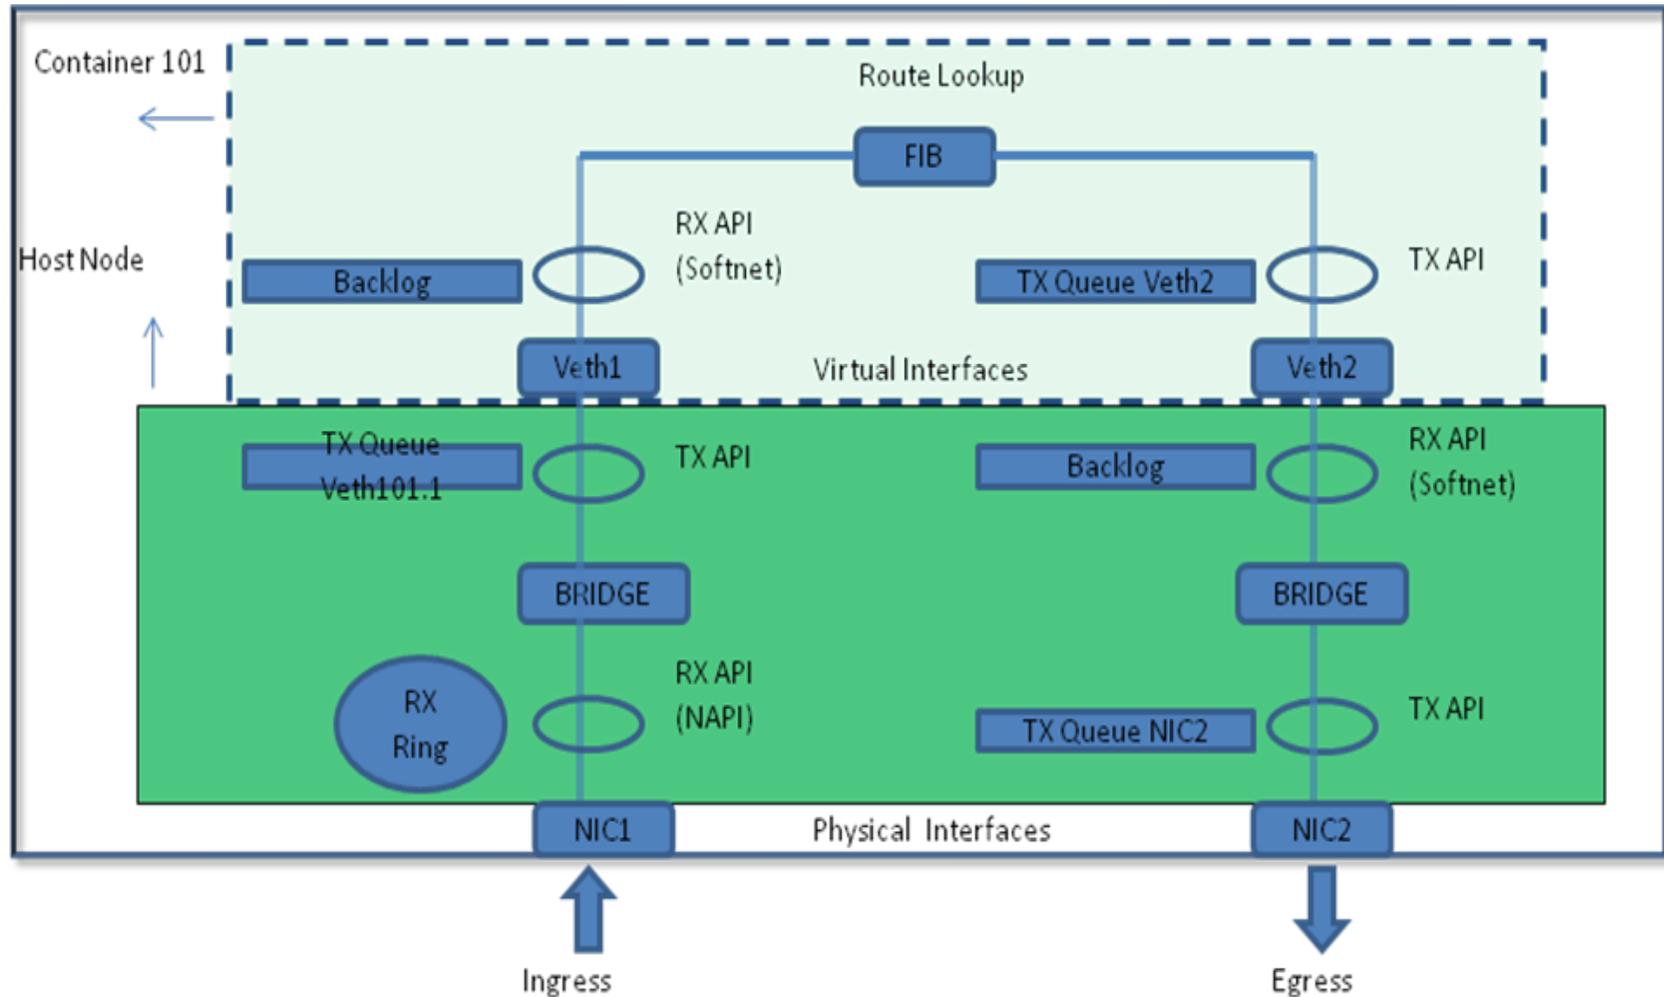
- **Virtual Router**

Throughput: 334kpps

Packet drop: Backlog queue congestion occurred at the offered load of 429kpps

Ingress physical interface, CPU saturation observed at the offered load of 650kpps

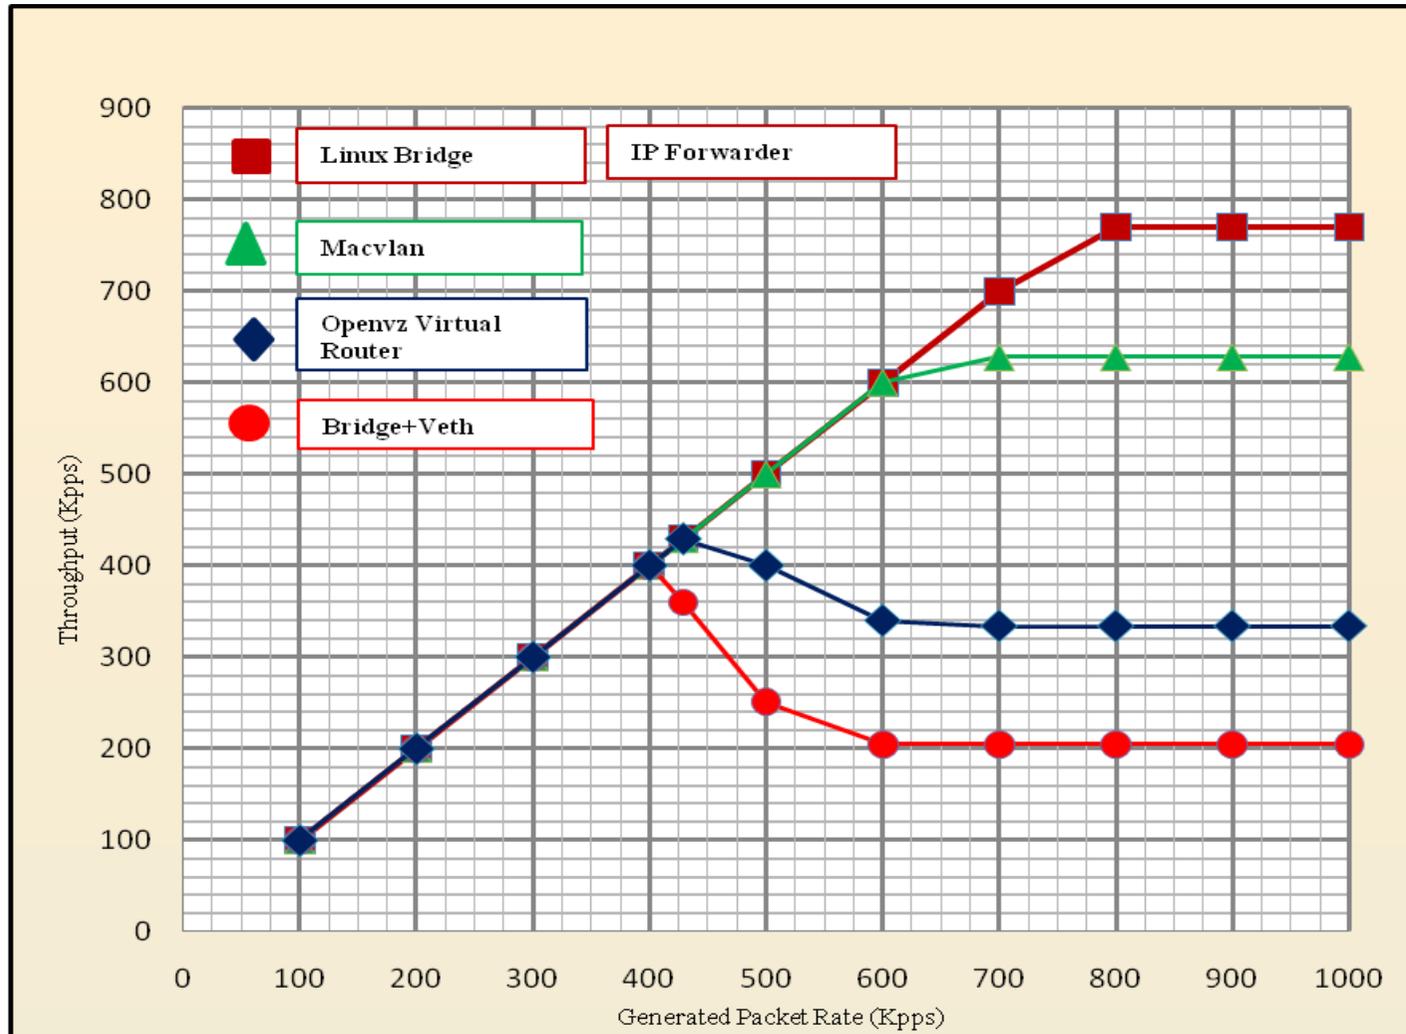
Virtual Router Design Internals



Virtual Router Design: An alternative approach

- Linux Namespaces, an emerging container based virtualization
 - Macvlan, a virtual device provides a built in mechanism of physical/virtual device mapping
 - Both bridge and veth are replaced with macvlan device
-

OpenVZ vs. Namespace Virtual Router



Virtual devices CPU usage

Packet Rate (kpps)	CPU %age Usage			
	<i>Kernel 2.6.27-openvz chistyakov</i>			<i>Kernel 2.6.34 Net-Next</i>
	<i>Linux Bridge</i>	<i>Veth</i>	<i>Total</i>	<i>Macvlan</i>
200	9	1.5	10.5	2.3
429	11	1.9	12.9	3.5
450	16	1.9	17.9	3.6
600	17	2.2	19.2	4.6
650	18	2.3	20.3	5
800	18	2.3	20.3	5

Conclusion and future work

- Apart from any virtualization technology, the way in which devices are mapped is important
 - Linux bridge is a CPU intensive device (MAC learning, forwarding database updates etc)
 - Macvlan is an attractive alternate
 - It is important to know how virtual devices communicate with kernel
 - Backlog is still there which may become performance bottleneck
-



ROYAL INSTITUTE
OF TECHNOLOGY

- Thanks for listening
 - Questions ?
-