CrossBow and VirtualWire

Nick Street

October, 2009

CrossBow Motivation

- Datacentre Resource Consolidation
- Dedicated Network Interface Cards too costly
- Need to share physical NICs across Containers / VM
  - Increase bandwidth utilisation
  - Flexible network resource allocation
  - Reduce per-port cost: host, switch, wiring, power

Use Cases

- ISP offering web and e-mail services
  - Consolidate multiple hosts on a single machine
  - All users expect good performance on their virtual host
- Financial services
  - Consolidate multiple services on a single machine
  - Some services have minimum performance requirements, or higher priority
- Mixing UDP & TCP traffic

Segregated Stacks

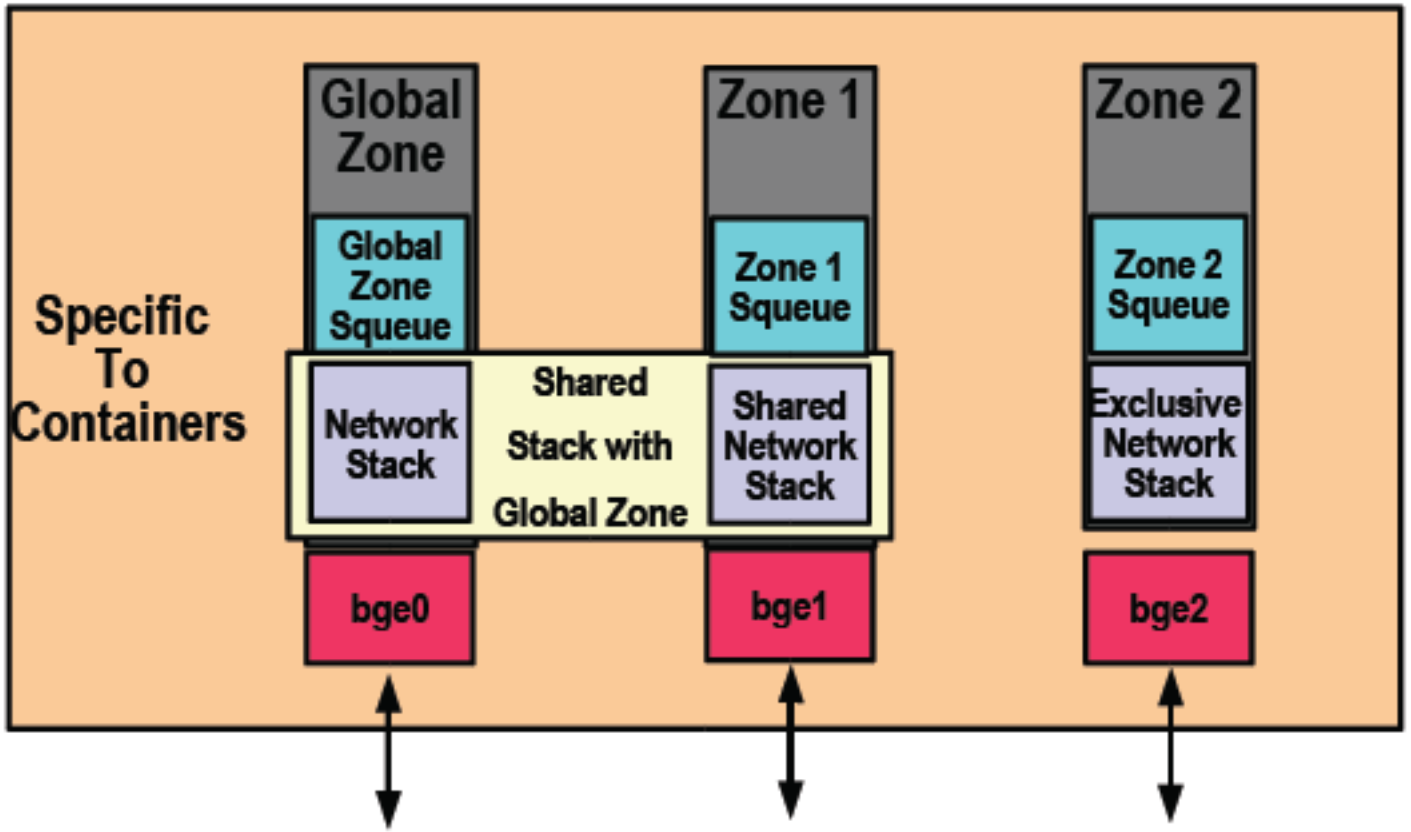

Crossbow “Hardware Lanes”

Ground-Up Design for multi-core and multi-10GigE

- Linear Scalability and Assured Throughput
- '*Hardware Lanes*' with dedicated resources
- Network Virtualization and QoS designed in the stack
- More Efficiency due to Dynamic Polling and Packet Chaining

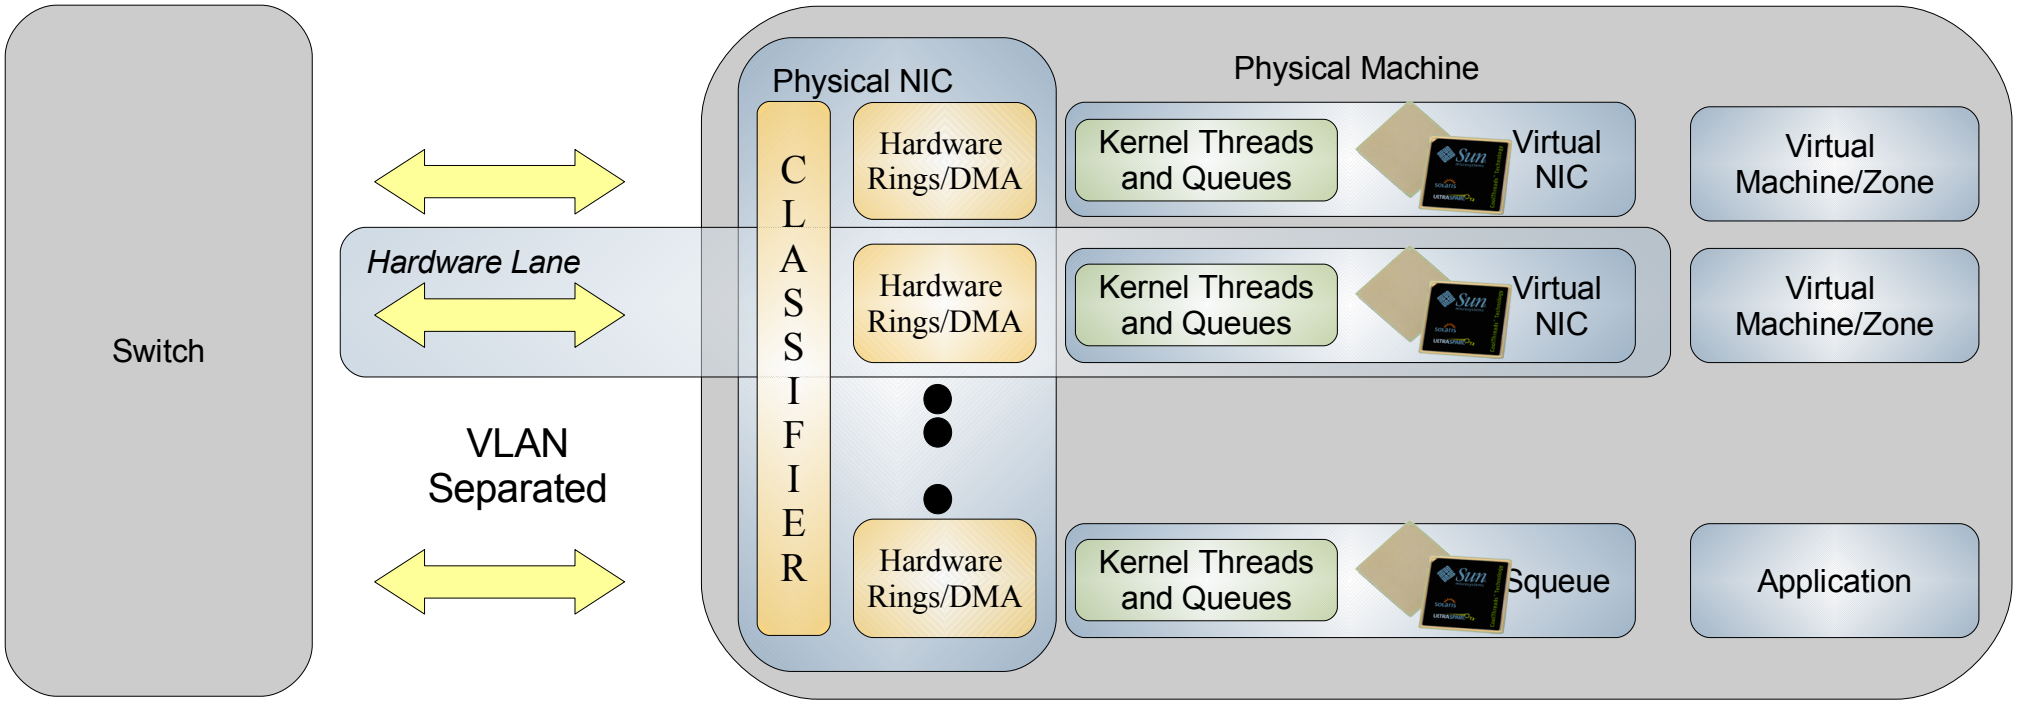

Crossbow Network Stack

- *Fairness:* Ability to allow various traffic types to share the bandwidth and associated compute resources in a fair Manner
- *Priority:* Ability to prioritize service, protocols or containers
- *Resource control:* Ability to manage network bandwidth and associated compute resources for each service, protocol, or Container
- No performance penalty
- Features must be integrated as part of the architecture

Virtual Network Stack

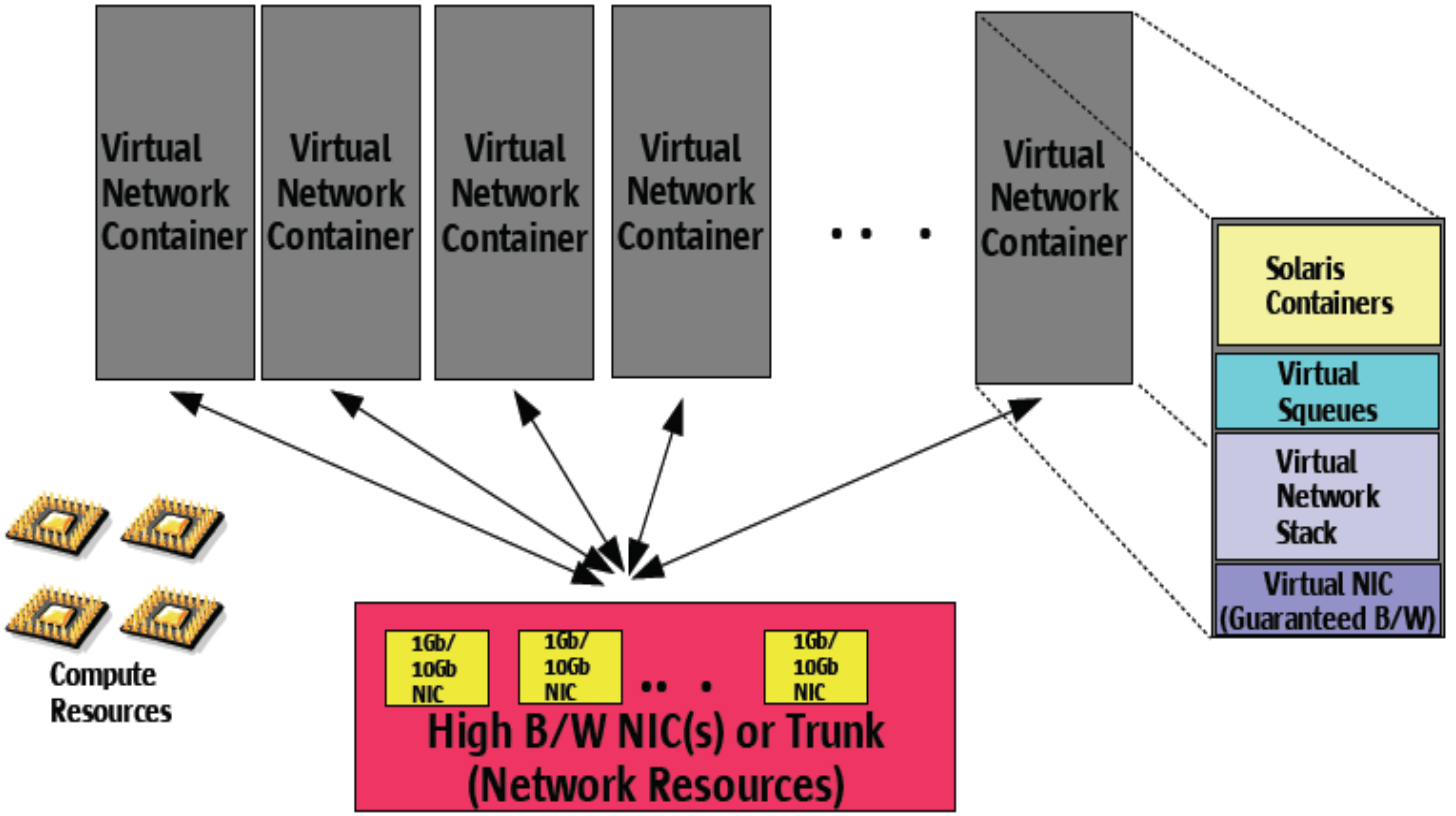

Virtual Stacks: Services & Protocols

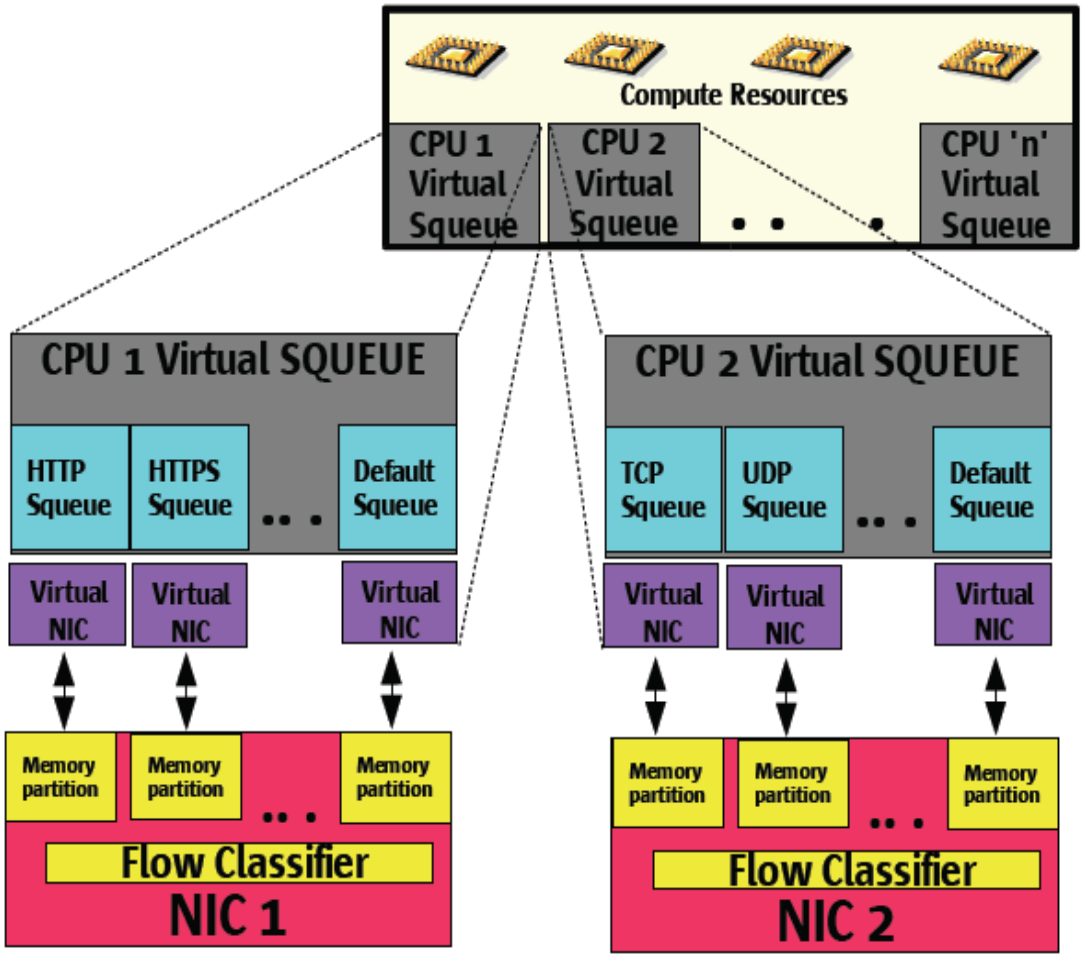

- Classifying Traffic
- Bandwidth Control
- DOS avoidance

Crossbow Smart NIC & Xen

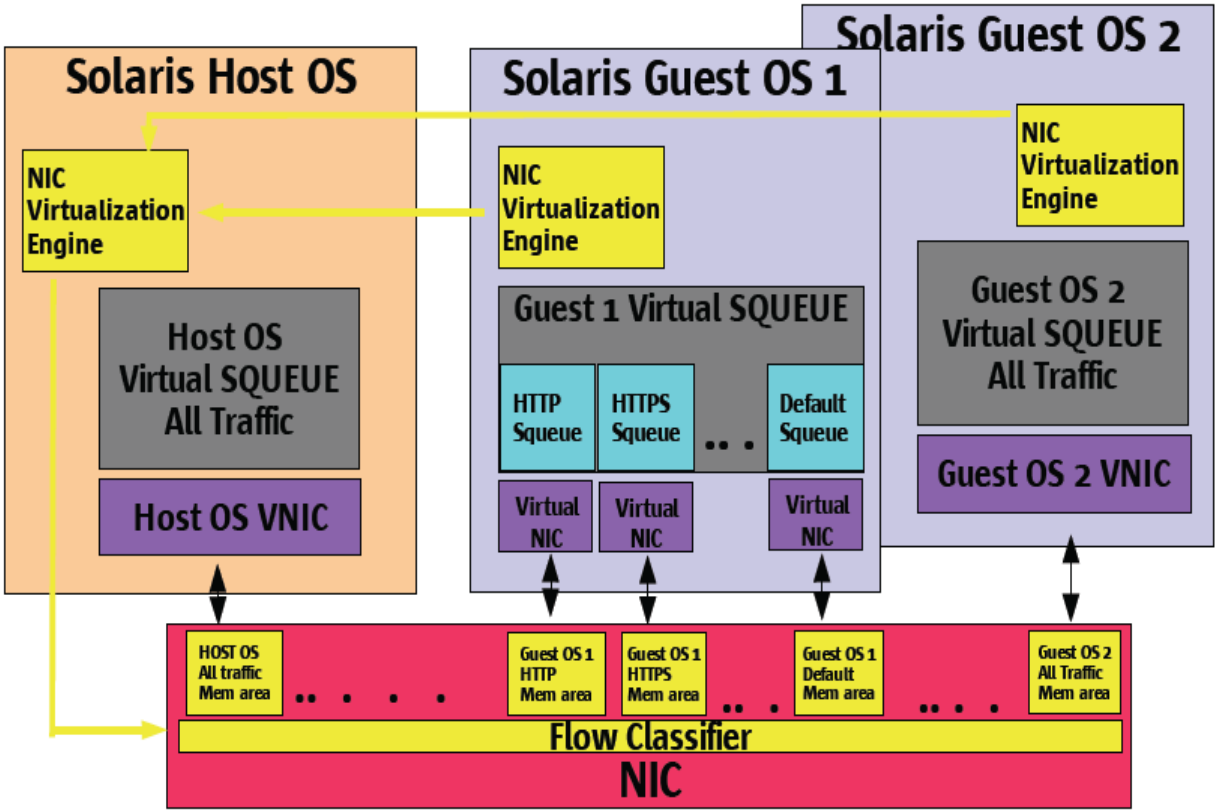

Virtual Stacks with Dumb NICs

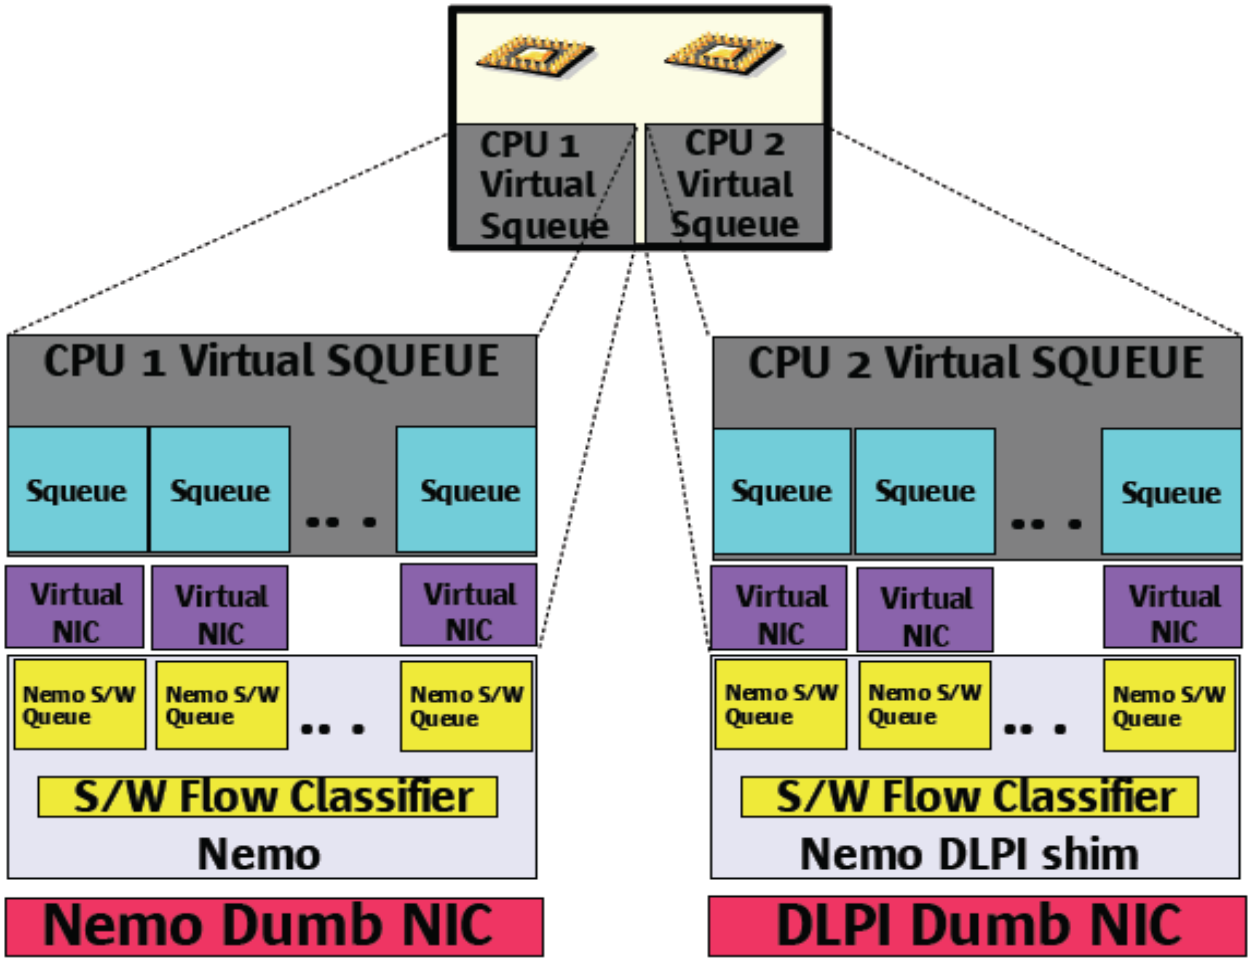

- *'Highest Common Denominator'*

Virtual Network Machine for the Cloud

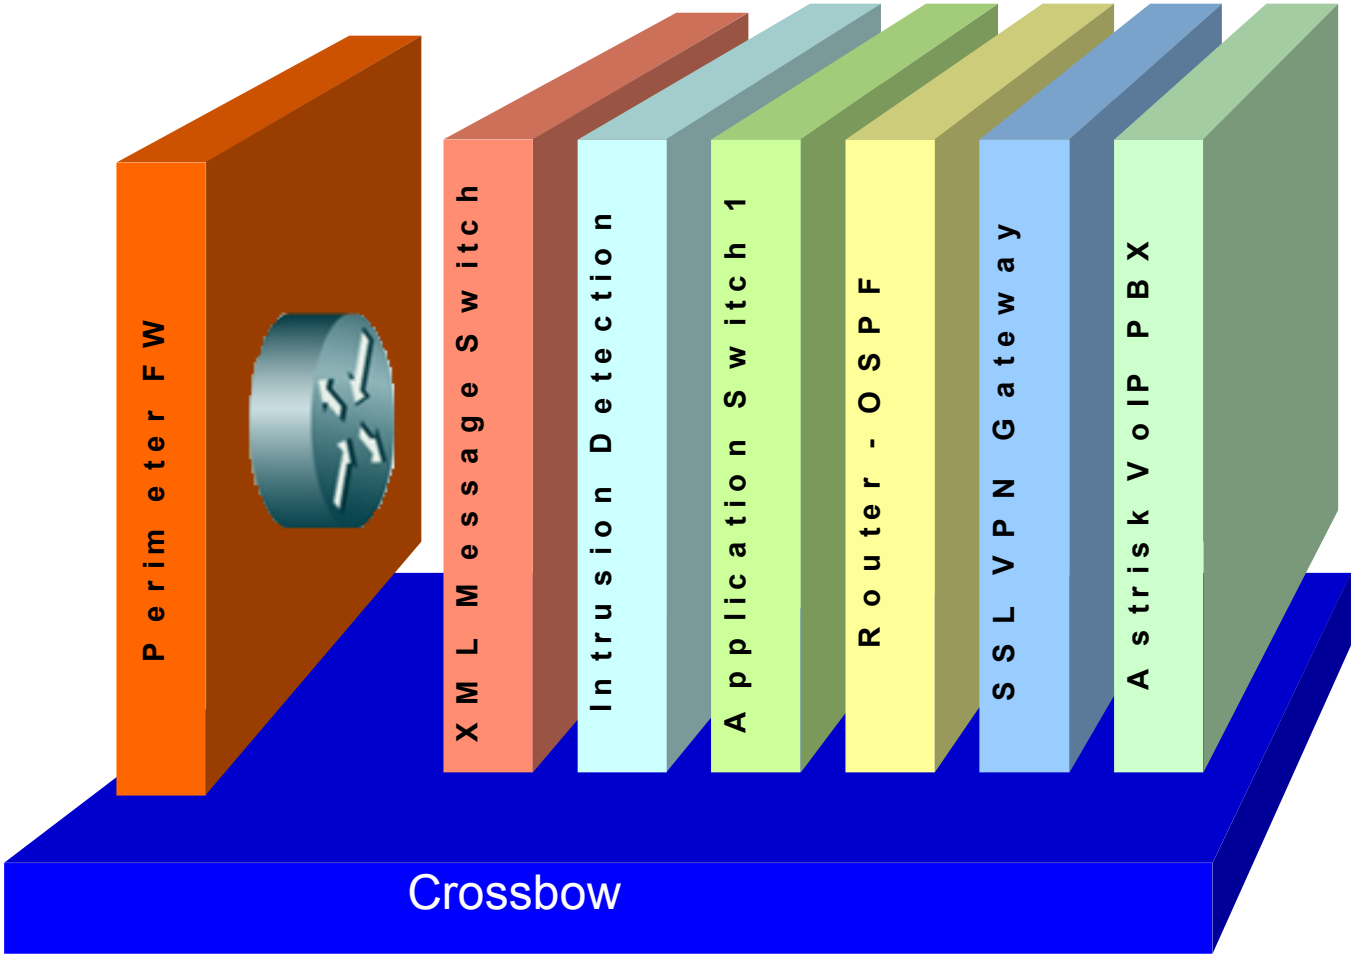

The network is the computer

- *Router, Firewall, DNS Server, DHCP Server, Quagga*

Physical Wire with Physical Machines

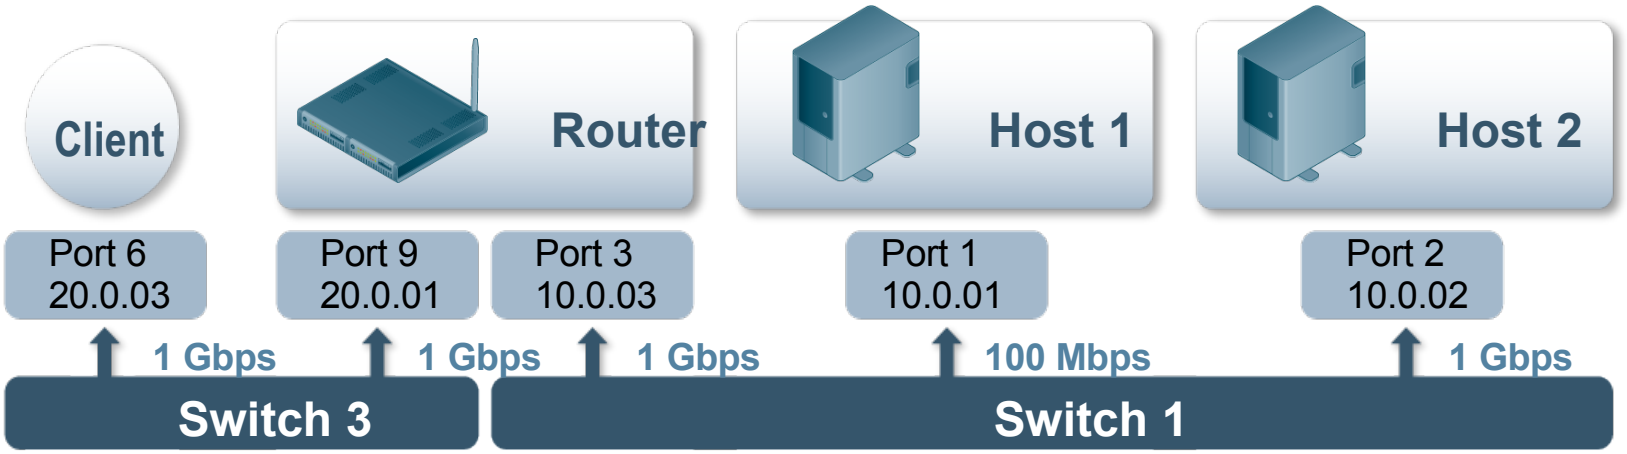

Virtual Wire with Virtual Network Machines

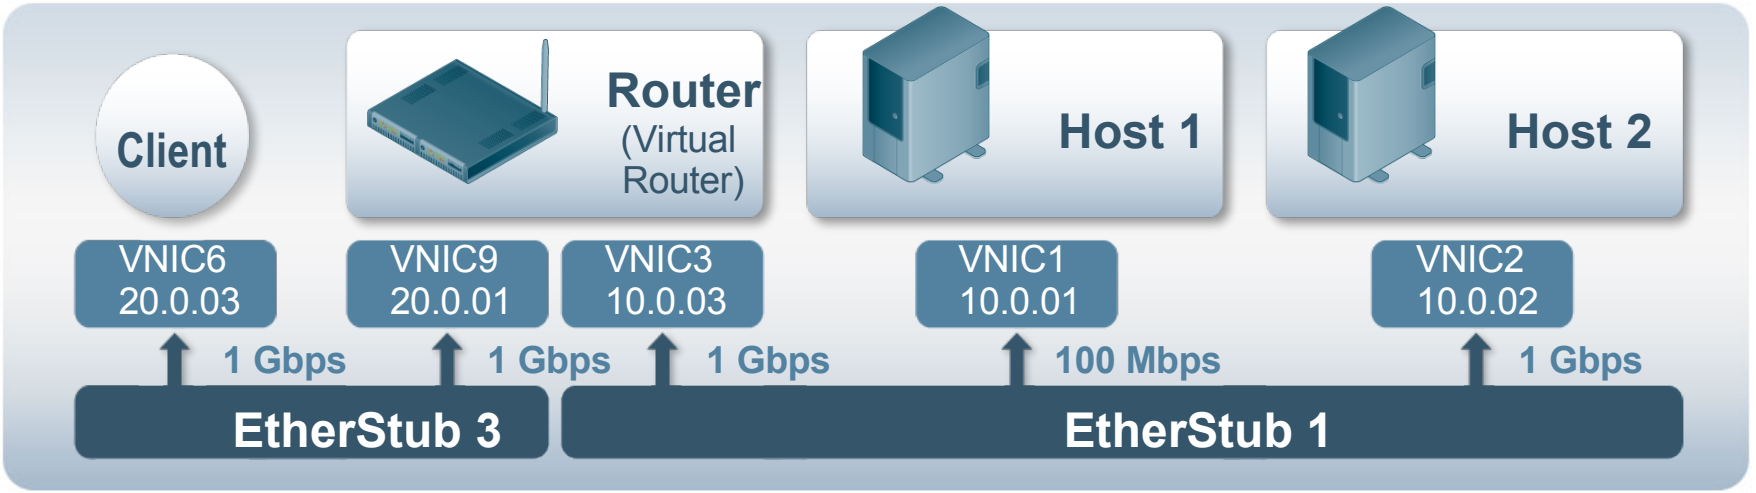

VirtualWire: Network in a Box

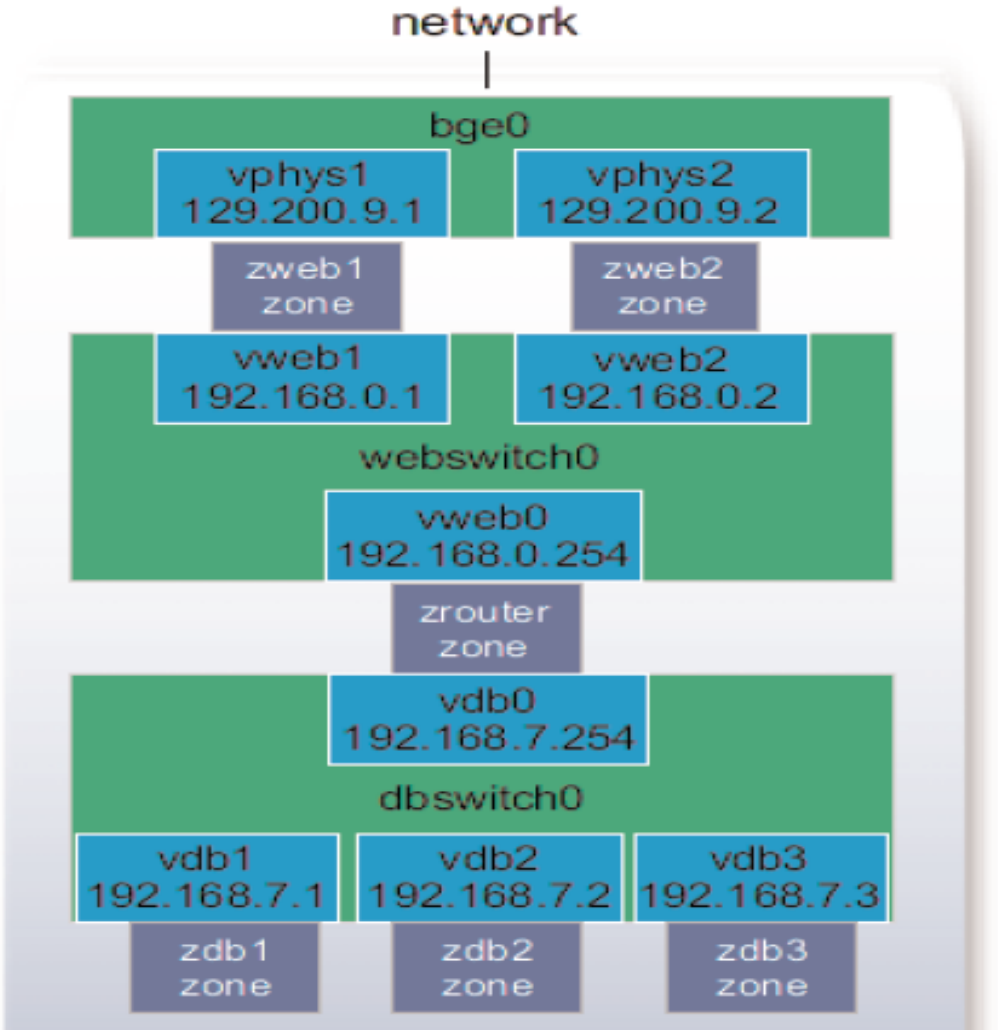


```
# dladm create-etherstub webswitch0

# dladm create-etherstub dbswitch0
```

Admin Interfaces

```
iprcm [-b B/W] [-p pri] [-i phys_iface] [-v virt_iface] ip_addrs and services

zonezonecfg -z new_zone
zonecfg:new_zone> create
zonecfg:new_zone> net phys=bge1
zonecfg:new_zone> net virt=eth0
zonecfg:new_zone> net bw=30Mbps
zonecfg:new_zone> net pri=hi
zonecfg:new_zone> net virt_stacks=3
zonecfg:new_zone> net ip_addr=a.b.c.dcfg:new_zone> net virt=eth0

dladm -i phys_intface -v virt_intface -b b/w -p pri
ipvstack -v virt_intface ipaddr and services
zonecfg:new_zone> net virt=virt_intface


flowadm create-flow -l bge0 protocol=tcp,local_port=443 -p maxbw=50M http-1
flowadm set-flowprop -l bge0 -p maxbw=100M  http-1
```

http://www.opensolaris.org/project/crossbow/demo/Crossbow_HowToGuide.pdf

Crossbow VWire Demo Tool

- CAD style GUI
- Configure (Compiles CLI commands)
  - dladm
  - zonecfg
  - Flowadm
- Observe Traffic
- Generate Traffic
- Allocate Bandwidth

http://www.opensolaris.org/project/crossbow/demo

Acknowledgements

- Crossbow Team *esp.*
  - Sunay Tripathi
  - Kais Belgaied
  - Nicolas Droux
- For Presentation Materials

(and a Cool Product )