Crossbow Virtual Wire: Network In a Box

Sunay Tripathi, Nicolas Droux,
Kais Belgaied, Shrikrishna Khare

November 5th, 2009
USENIX LISA 09, Baltimore, MD

Nicolas Droux, Senior Staff Engineer
Solaris Kernel Networking, Sun Microsystems Inc.
nicolas.droux@sun.com

Key Issues in Network Virtualization

- Fair or Policy based resource sharing in virtualized environments
  - Bandwidth
  - NIC Hardware resources including Rx/Tx descriptors
  - Processing CPUs
- Overheads due to Virtualization
  - Latency, Throughput
- Management
  - Isolation between distributed applications
  - Network fabric configuration
- Security
  - New threats to L2 network
- Where to solve the problem?
  - Switches
  - L3/L4 devices
  - Hosts

Crossbow: Solaris Networking Stack

- 8 years of development work to achieve
  - Scalability across multi-core CPUs and multi-10gigE bandwidth
  - Virtualization, QoS, High-availability designed in
  - Exploit advanced NIC features
- Key Enabler for
  - Server and Network Consolidation
  - Open Networking
  - Cloud computing

Crossbow “Hardware Lanes”

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

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

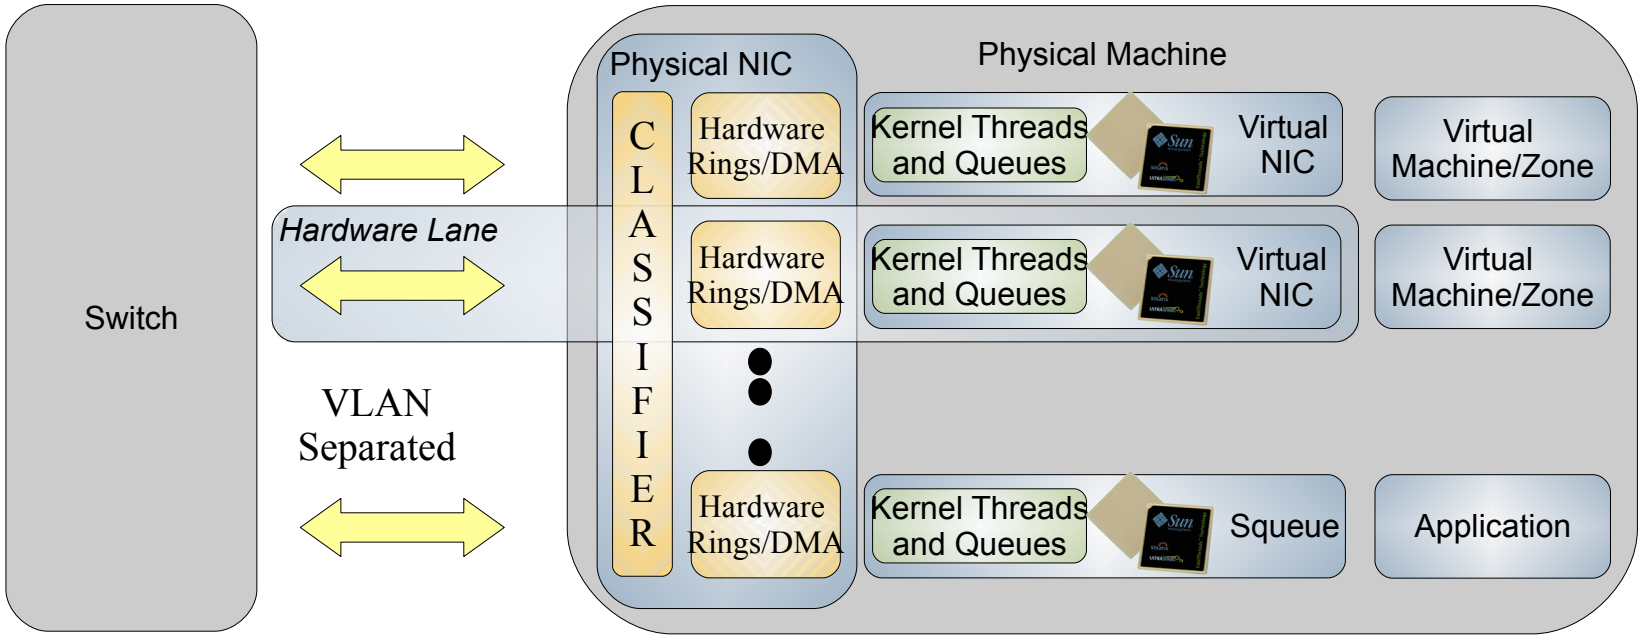

Hardware Lanes and Dynamic Polling

- Partition the NIC Hardware (Rx/Tx rings, DMA), kernel queues/threads, and CPU to allow creation of “Hardware Lane” which can be assigned to VNICs & Flows
- Use Dynamic Polling on Rx/Tx rings to schedule rate of packet arrival and transmission on a per lane basis
- Effect of dynamic polling

Mpstat (older driver)

```
intr   ithr  csw   icsw  migr smtx  srw  syscl  usr sys wt idl
10818  8607  4558  1547  161  1797  289  19112  17  69  0  12
```

Mpstat (GLDv3 based driver)

```
intr   ithr  csw   icsw  migr smtx  srw  syscl  usr sys wt idl
2823   1489  875   151   93   261   1    19825  15  57  0  27
```

~75% Fewer Interrupts

~85% Fewer Ctx Switches

~85% Fewer Mutexes

~15% More CPU Free

Crossbow Virtual NICs (VNICs)

- Pseudo MAC instances
  - Can be managed as if they were physical NICs
  - Per VNICs stats, reuse existing management tools
  - Link speed derived from configured bandwidth limit
  - High-Availability by creating VNICs on link aggregations or combining VNICs in IPMP groups
- Dedicated per-VNIC hardware and kernel resources
- Data path pass-through, no bump in the stack
- Standards based End-to-End Network Virtualization
  - VLAN tags and Priority Flow Control (PFC) assigned to VNIC extend Hardware Lanes to Switch

Crossbow Virtual Switching

- A virtual switch is created implicitly each time >2 VNICs are created on a data link
- The MAC layer provides packet switching semantics equivalent to an ethernet switch
  - Data path between VNICs created on top of the same data link
  - Connectivity between VNICs and physical network
  - Per VLAN broadcast domain, isolation between VLANs
- VNICs can be created on etherstub to create virtual switches independent from hardware

Crossbow Virtual Switching Example

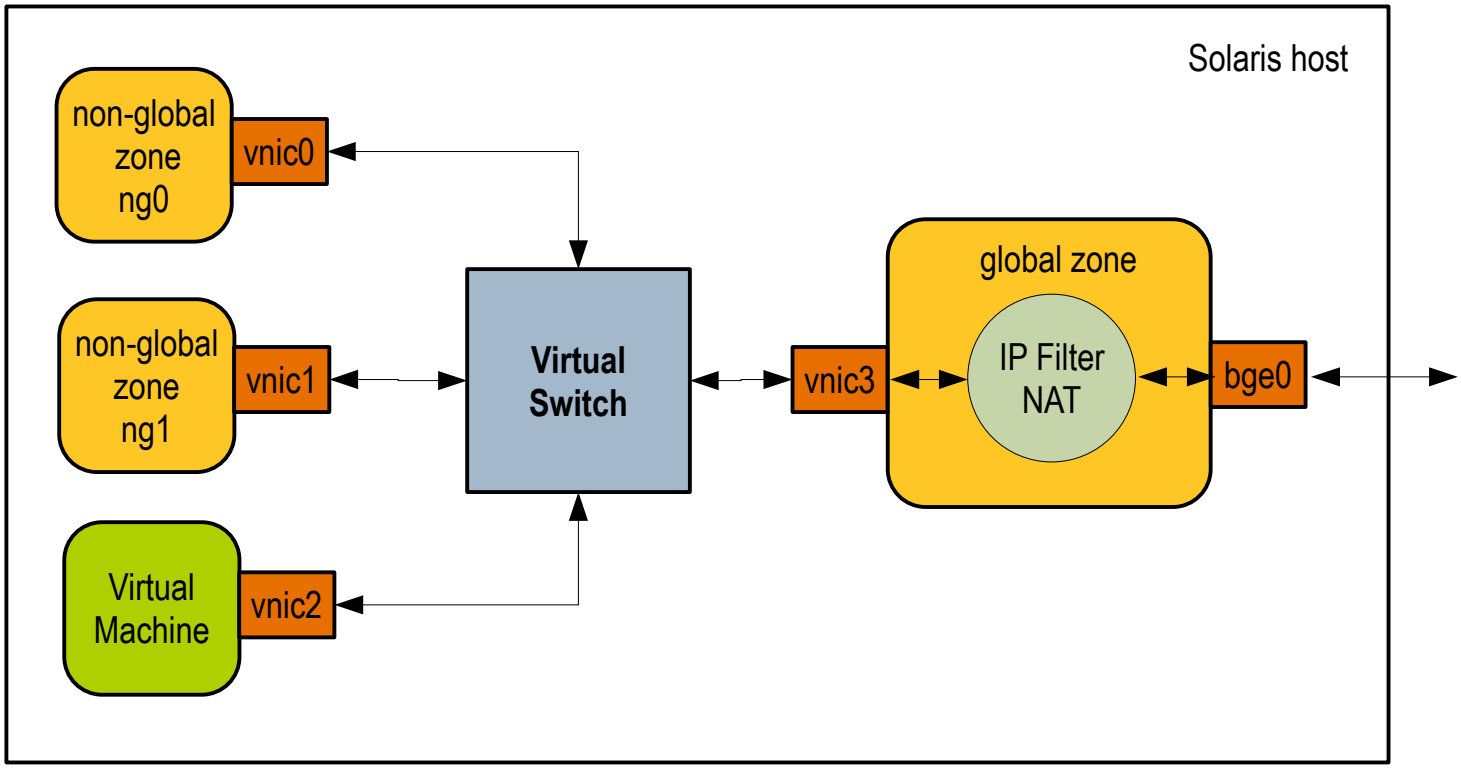

Virtual NIC & Virtual Switch Usage

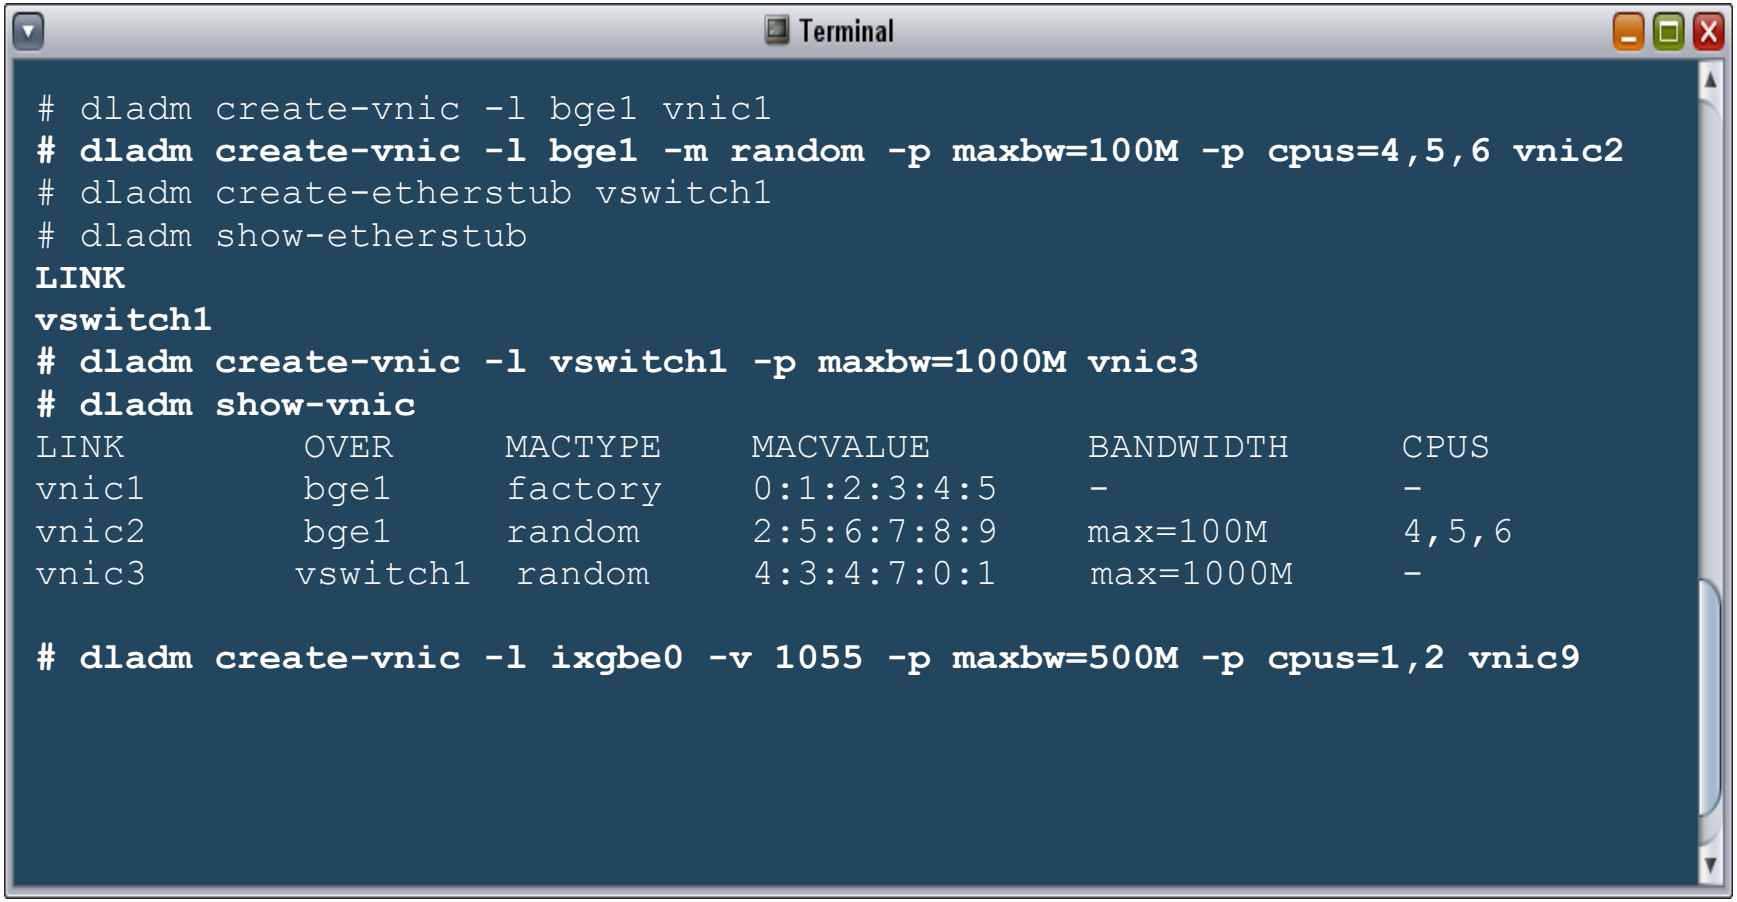

Physical Wire w/Physical Machines

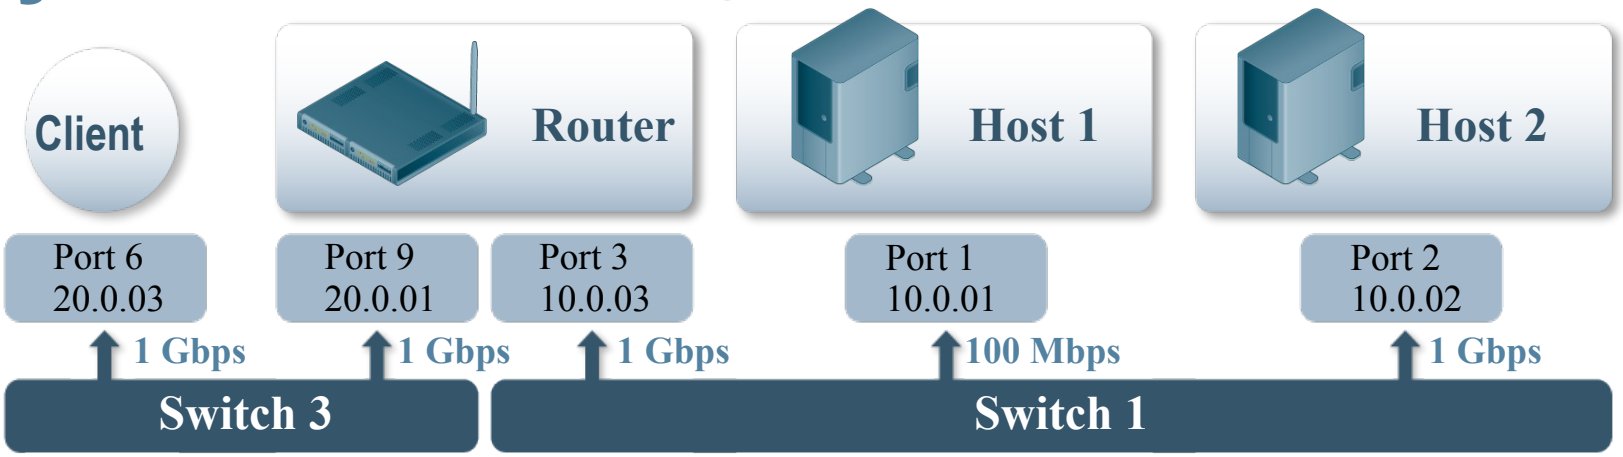

Virtual Wire w/Virtual Network Machines

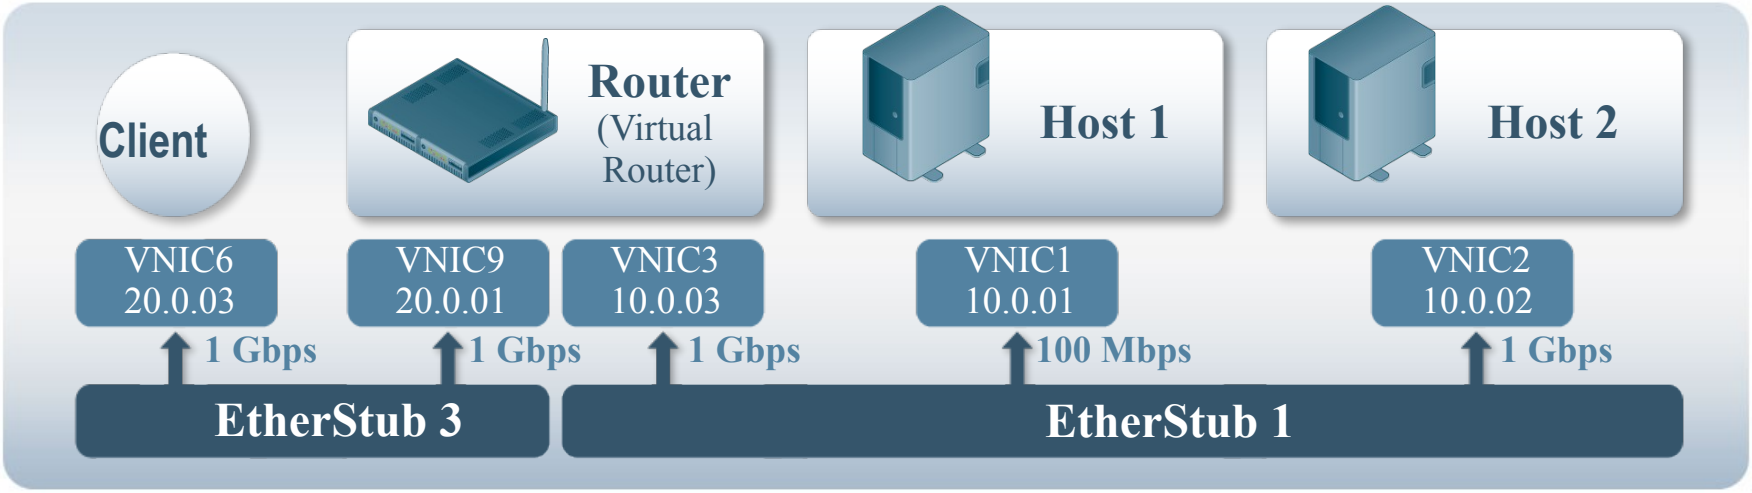

Virtual Network Machines

- A Virtual Network Machine (VNM) is a Zone or Virtual Machine associated with a set of network functions (routing, firewall, load balancing, etc)
- A VNM has dedicated VNIC(s) with configured link speed, CPUs
- Multiple VNMs can run on a single host, connected through virtual private networks (etherstubs) or to the physical network
- Use for simulation, consolidation, testing, etc

Crossbow Flows

- Crossbow **flows** based on the following attributes
  - > Services (protocol + remote/local ports)
  - > Transport (TCP, UDP, SCTP, iSCSI, etc)
  - > IP addresses and IP subnets
  - > DSCP labels

- The following properties can be set on each flow
  - > Bandwidth limits
  - > Priorities
  - > CPUs

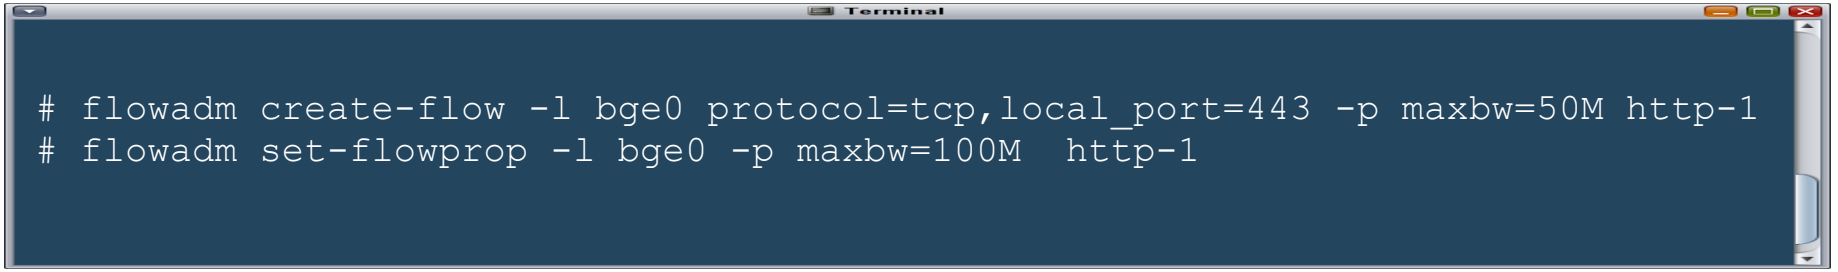

Join Us...

- Beer @ Crossbow and Solaris Networking BoF
  - Tonight 10:30-11:30pm (Dover A&B)
  - Presentation by Ben Rockwood (Joyent)
  - vWire demo and deep-dive discussions
- OpenSolaris project and community
  - http://www.opensolaris.org/os/project/crossbow
  - crossbow-discuss@opensolaris.org
  - networking-discuss@opensolaris.org

Crossbow Virtual Wire: Network In a Box

Nicolas Droux
nicolas.droux@sun.com
Solaris Kernel Networking