


BROCADE



Bridge Requirements for Server Virtualization

new-dcb-ghanwani-server-virtualization-0908-v1

Anoop Ghanwani (anoop@brocade.com)
Roy D'Cruz (rdcruz@brocade.com)
Satish Sathe (ssathe@brocade.com)
Suresh Vobbilisetty (svobbili@brocade.com)

September 2008

Motivation

- Many data center customers are in various stages of evaluating or deploying server virtualization; VMware, Microsoft Virtual Server
 - Increases server CPU utilization especially as multi-core CPUs are introduced
 - Changes admin ratios – from 20-30 servers/admin for physical machines to 60-100 servers/admin for virtual machines
- The market for virtual machine software is growing rapidly [IDC]
 - \$1.78B in 2007 from \$1.05B in 2006
 - \$5.5B by 2012
- This introduces some new problems in the area of networking for virtualized server entities



Definitions

- VM: Virtual Machine
 - An instantiation of a server within a physical machine
 - Does everything that regular servers do
 - Communication channel to the external world is shared with other VMs
 - May communicate with other VMs within the same machine
 - May or may not be under the administrative control of the entity managing the physical machine in which it is instantiated
 - The OS running in the VM is sometimes referred to as a “guest OS”
- Hypervisor
 - Performs administrative functions for managing the virtualized environment
 - e.g. VMs register with the hypervisor when instantiated
 - Implements inter-VM communication (when that communication is implemented in software), sometimes called a “soft switch”
 - Under administrative control of the entity managing the physical machine



Problems that need to be addressed (1)

- Performance/scalability
 - Implementing the communication path between VMs in software uses precious CPU cycles
 - Software-based solution is not cost effective as we get to 10G and beyond
- Security
 - Authentication of the VMs and the hypervisor
 - Guard against VM (guest OS) spoofing a MAC address
 - What assumptions can be made about the integrity of frame contents coming from a VM?
 - Inside the physical machine/on the NIC
 - Outside the physical machine in an external bridge

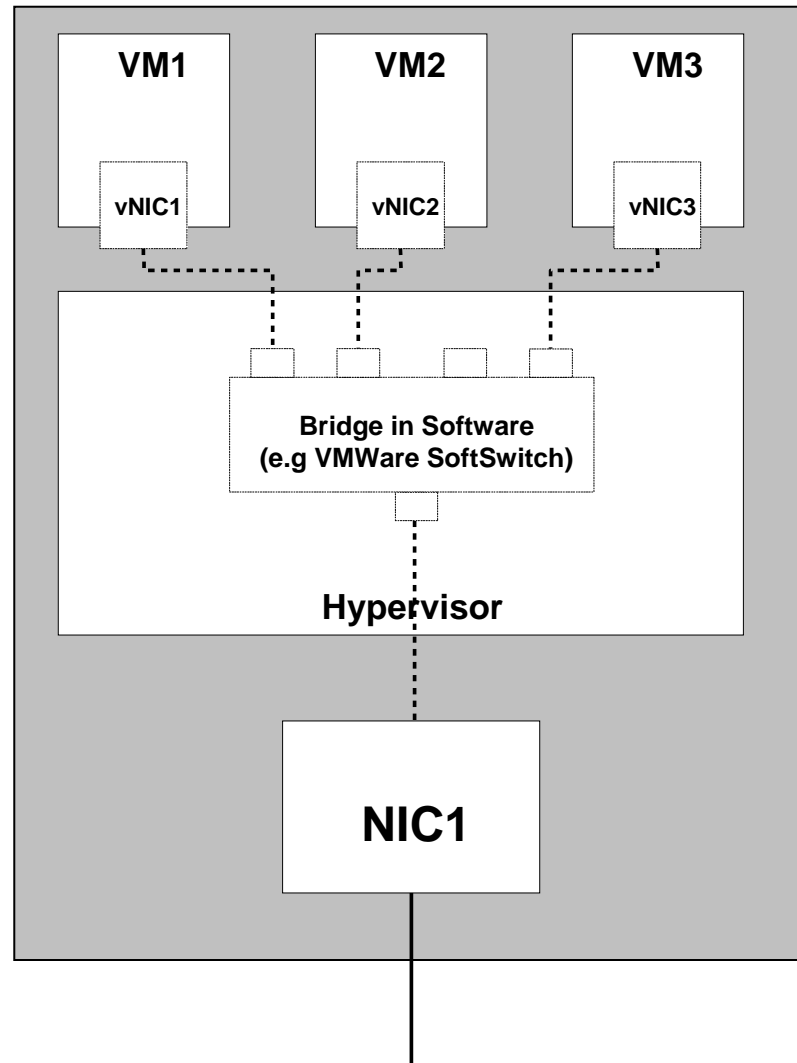


Problems that need to be addressed (2)

- Manageability
 - Does the communication entity function as a piece of the server or of the network?
 - Can security features be implemented inside the server or outside?
 - Dealing with resource mobility in a seamless fashion – A VM moving from one physical machine to another
 - How to control the proliferation of bridges that need to be managed?
 - As many bridges as there are servers!
- Traffic management
 - Multiple VMs share the same physical channel to the external bridge
 - Dedicated CN/PFC/ETS per VM?

Implementation options

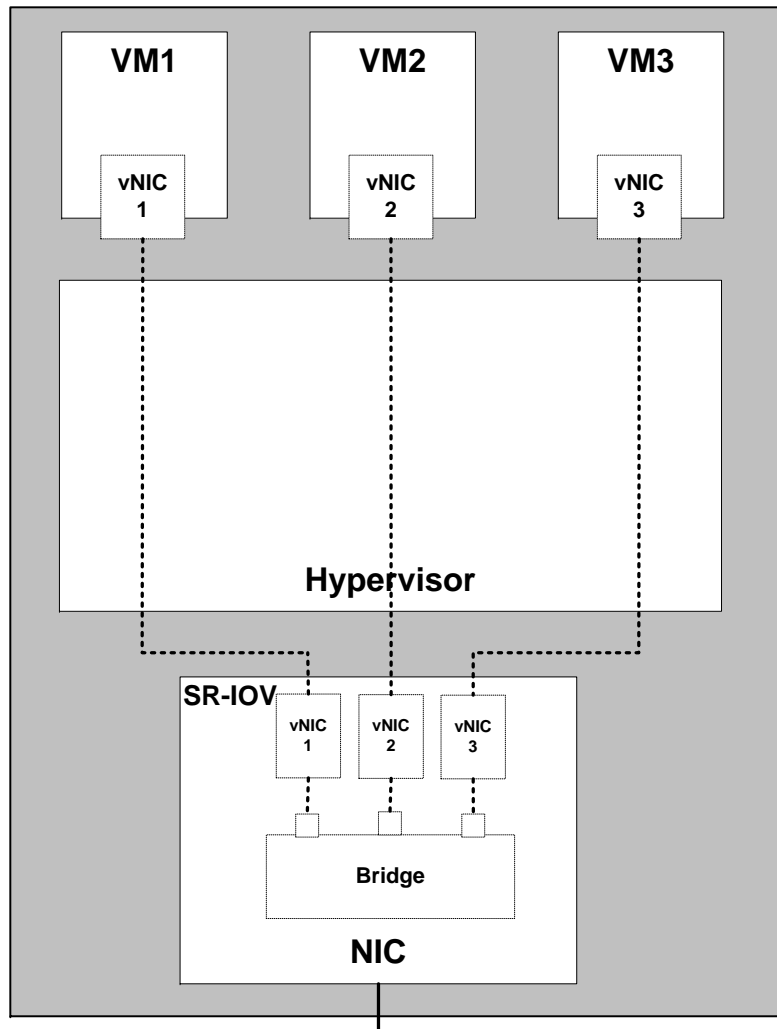
Software-based implementation in the hypervisor



- This is how it's done today
- Implementation of regular expression ACLs would be too much of a performance hit

Implementation options


Inside the physical machine/server



- + No changes to on-the-wire frame formats
- + Easier to implement things like per-VM traffic management
- Requires the NIC implementer to build a switch, and possibly all the security components that go with it – increases cost and complexity
- Potential scaling issues with respect to number of VMs that can be supported
- Management could be a problem – server guys and switch guys are usually independent

Implementation options

In an external bridge




- + Most of the functionality already exists
- + Allows centralized management
- ? May or may not need changes to on-the-wire frame format depending on what fields in the existing frame can be trusted
- Potential scaling issues with respect to bandwidth
 - Frames must go to an external device only to be bridged back out the same port
 - How to handle frames that need to be multicast to VMs within the same physical machine?
- Implementing per-VM traffic management could be tricky

Implementation options

Combined server and external switch

- Best of both?
- Traffic mgmt in the adapter
- Administrative controls in the external bridge

BROCADE



THANK YOU

