

Wireshark Developer and User Conference

A-4: Wireshark vs. “The Cloud”

Tuesday June 14, 2011. 3:30pm – 4:45pm

Jasper Bongertz

Senior Consultant | Fast Lane Institute for Knowledge Transfer

SHARKFEST '11

Stanford University

June 13-16, 2011

Agenda

- Uh, Cloud?
- Physical vs. Virtual
- Cluster Basics
- VMs on the Move
- Capture Methods



„Cloud“ - Terminology


- Using Cloud resources means
 - Servers and locations are more or less irrelevant
 - You don't know where your „stuff“ actually is
- Resources may be
 - shared between multiple users/teams/companies
 - paid for what is actually used, when it's used
- Offering OS platforms, installation platforms, software catalogs
- Allows self service access to shared computing resources

Cloud types

- Private cloud
 - Hosted on private virtualization servers
 - Offering services to internal users/teams
- Public cloud
 - Offering services to anyone („with a credit card“)
- Hybrid cloud
 - Mix of private and public cloud
 - Usually used to outsource computing power in times of high usage, or as fallback

Virtual Environments

- Virtual Environments consolidate multiple servers on one or more virtualization hosts
- Physical hardware runs an virtualization layer with virtualized servers on top
- Virtual Servers share
 - CPU cycles and memory
 - Storage
 - Network adapters




Enterprise Virtualization

- Common virtualization solutions found in datacenters today are:
 - Citrix XenServer
 - Microsoft Hyper-V
 - Red Hat Enterprise Virtualization
 - VMware vSphere
- Basically all enterprise virtualization solutions have the same basic features
 - or will have them sooner or later

Host Virtualization Example #1

- Virtualization host runs multiple Virtual Machines on a single NIC
- The host may use the NIC for its own data communication, too
- Potentially dozens of virtual servers showing up with their own virtual MAC address on the physical NIC




Capturing virtual servers

- In virtual/cloud environments, virtual servers, applications, services may run everywhere
- Multiple virtual servers on physical hosts share network cards
- If you have access to the virtualization host you can SPAN/TAP its connections
- Challenges:
 - Find and capture the correct NIC
 - Isolate traffic for the virtual server/application
 - Servers with 10GBit or even faster links


Host Virtualization Example #2

- There may also be „internal only“ switches making things complicated
- Data on internal switches never leaves the virtualization host
- No physical pickup possible



NIC Teaming Strategies

- Port ID based
- Source MAC hash
- Source/Destination IP hash
- Based on physical NIC load




Virtualization Cluster Basics

Trouble Brewing

Virtualization Cluster Example

- Group of virtualization hosts combined into a cluster



Cluster Basics

- Server clusters are always difficult to capture
 - Even without virtualization you usually don't know where the connection will end up
- Possible solutions include
 - Forcing specific connections to certain cluster members that can be captured
 - Capturing a common cluster uplink if available
 - Las Vegas style: capture somewhere and hope that you'll catch the relevant frames 😊

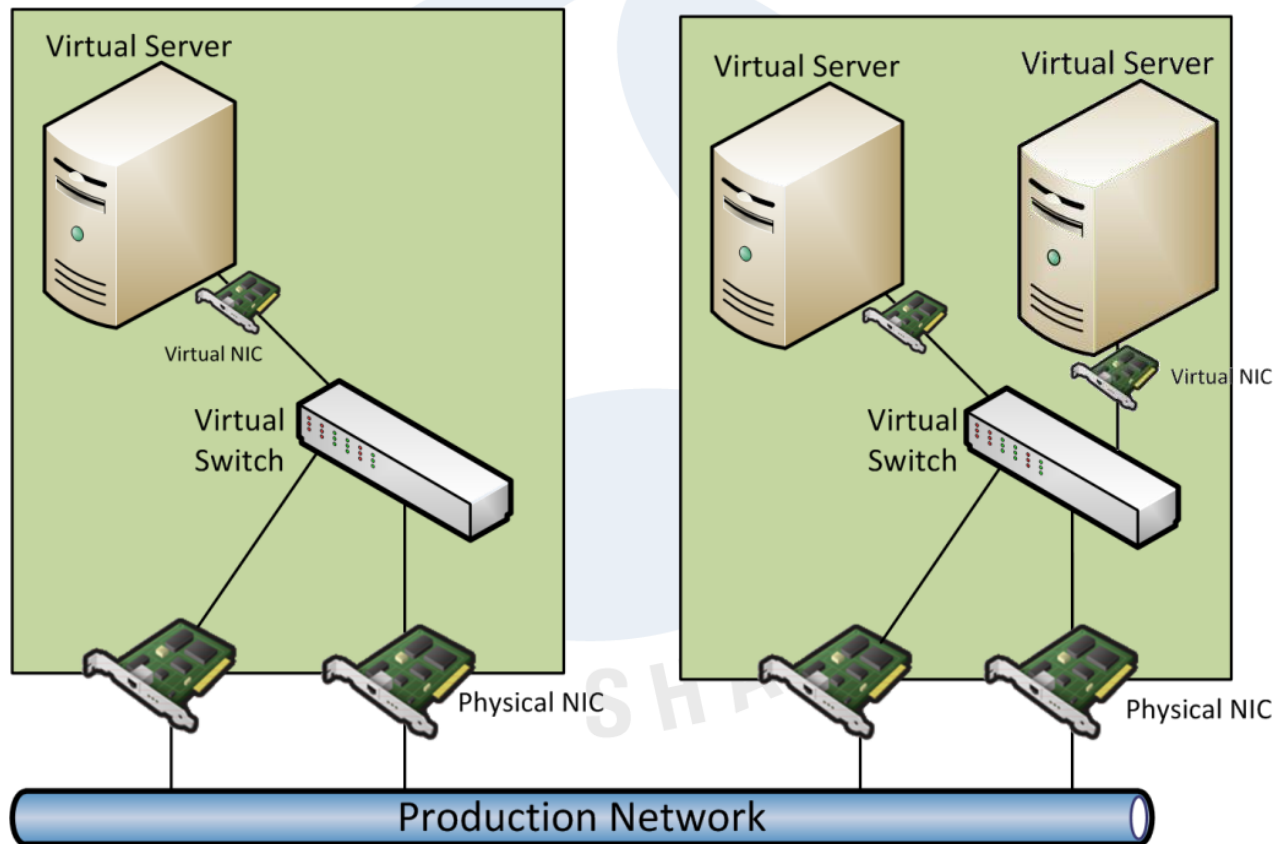
Virtualization Clusters

- Virtualization clusters are even more complex than clusters of physical servers
 - Load Balancing of virtual machines
 - High Availability / Failover
- Virtual machines may move from host to host without warning, at any given time!
- Requires shared storage
 - Fibre Channel, iSCSI, NFS
 - Lets better hope you never have to capture those... unless you like megatons of data

VMs on the move

Live Moving of Virtual Machines

- Virtual Machines may move from host to host while running




Cluster Movement Features

- High Availability (sort of)
 - Restart virtual machines on other hosts if there is a host crash
- Real High Availability
 - Running an “invisible” hot standby VM on a secondary host that is kept in sync
- Fully automatic live VM moving
 - Load Balancing virtual machines across virtualization hosts

Capture Strategies


Capture Strategy #1

- Install Wireshark on the virtual system of interest
- Advantages:
 - Can capture, even on VMs with internal only NICs
 - Sometimes your only option
- Disadvantage:
 - Changes the environment
 - Gets funny results (way too often)
 - May crash the VM




Capture Strategy #2

- Capture at virtualization host uplink (TAP/SPAN)
- Maybe your only option when you have no better access to the virtual infrastructure
- Advantages:
 - Easy to do in simple setups
 - Usually gets good data
 - Most familiar way to get data since its similar to physical captures




Capture Strategy #2

- Disadvantages:
 - May get you tons and tons of data to sort
 - Server uplink may be too fast for your capture device or the SPAN port
 - VM may be live-moved off the server, interrupting the capture
 - Worst case: you don't even know *where* to capture!



Real World Example



„Too much data“

- Ways to handle „too much data“ (a.k.a „dropped frames“) on physical captures:
 - use frame slicing if possible
 - SPAN only as few affected ports or VLANs as possible
 - use a filtering TAP
 - Capture Filters on the Wireshark itself may help, too
 - Use dumpcap on command line

New Capture Strategies

Virtual captures for a virtual environment

New Capture Strategies

- Virtualization technologies may or may not offer additional capture strategies
- The big question usually is „what can you do with that virtual switch?“
- Worst case: the vSwitch behaves like a dumb switch (a.k.a. Desktop Switch) – out of luck

- Promiscuous Mode on virtual switch or port group
 - Puts switch/port group into „hub“ mode
 - Security risk
 - Increased traffic
- Virtual SPAN sessions
 - Only on some devices, like Cisco Nexus 1000v
- Virtual TAPs
 - Probably: additional license required
 - Installation/configuration ain't exactly child's play

Demo





Questions?