

# Top 10 Network Issues in Apache CloudStack

Kirk Kosinski  
Escalation Engineer  
Citrix Systems

# Agenda

- Introduction
- Top 10 Network Issues
- Q&A



# VLAN Issues 1

- Switch misconfiguration
  - Symptoms
  - All VLANs trunked by default? Or denied?
- Other network problems




# VLAN Issues 2

- Hypervisor problems
  - NIC drivers
  - Bonding
    - xe-switch-network-backend?
  - Cabling
  - VLAN scaling
- Database hacking
  - Stop that!



# Open vSwitch

- Default on XenServer / XCP
  - VMware + Nicira?
- Weirdness
- `ovs-appctl`



# Security Groups

- KVM
  - Works out of the box (unless it doesn't)
- XenServer / XCP
  - Must enable Linux bridge back-end
  - Must install CSP (XS <6.1)
- vSphere...



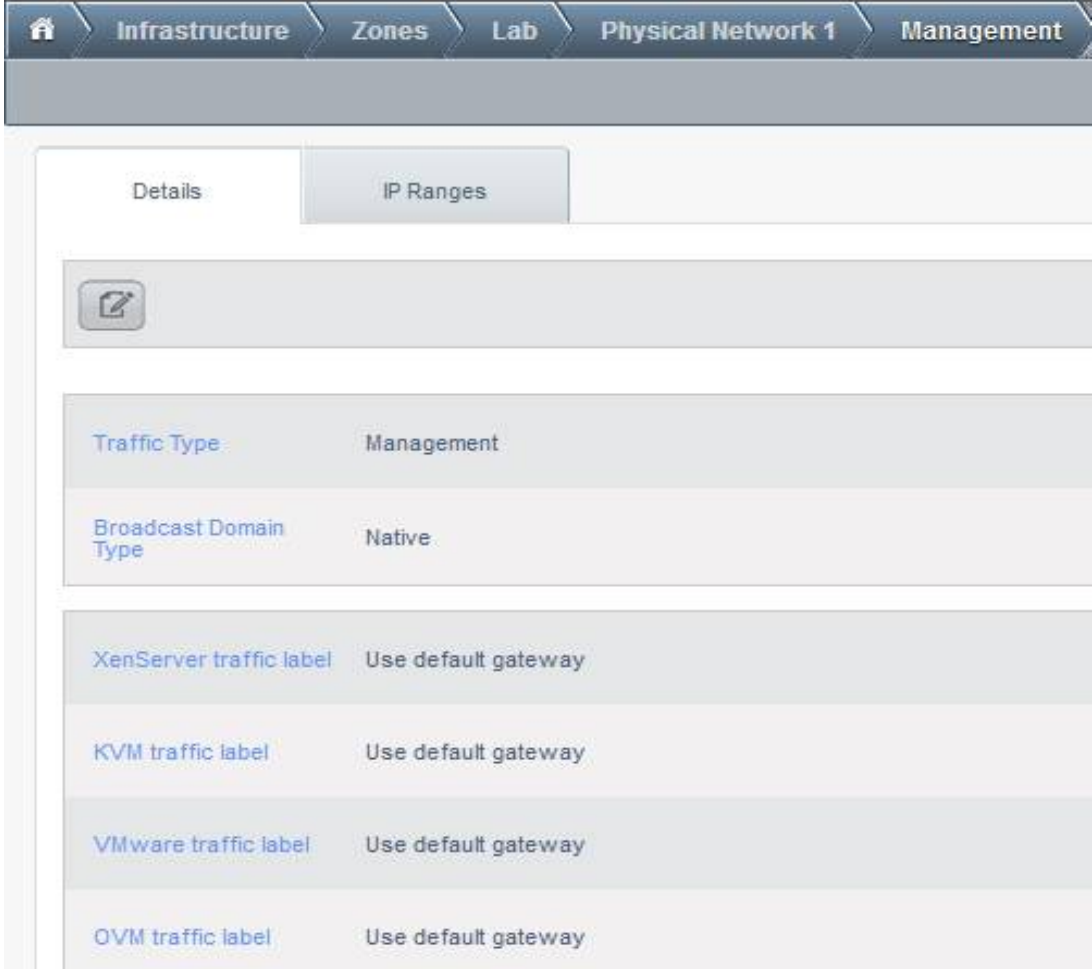
# “Host” Connectivity

- Hypervisors
- System VMs
- Secondary Storage
  - Alert status is normal



# CloudStack “Physical Networks”

- Not necessarily “physical”
- Traffic labels
  - Multiple NICs
  - Mgmt. traffic on a VLAN




The screenshot shows the CloudStack management interface for a Physical Network. The breadcrumb navigation at the top reads: Infrastructure > Zones > Lab > Physical Network 1 > Management. Below the navigation, there are two tabs: "Details" (selected) and "IP Ranges". A "Details" section contains a table of configuration parameters:

Traffic Type	Management
Broadcast Domain Type	Native
XenServer traffic label	Use default gateway
KVM traffic label	Use default gateway
VMware traffic label	Use default gateway
OVM traffic label	Use default gateway



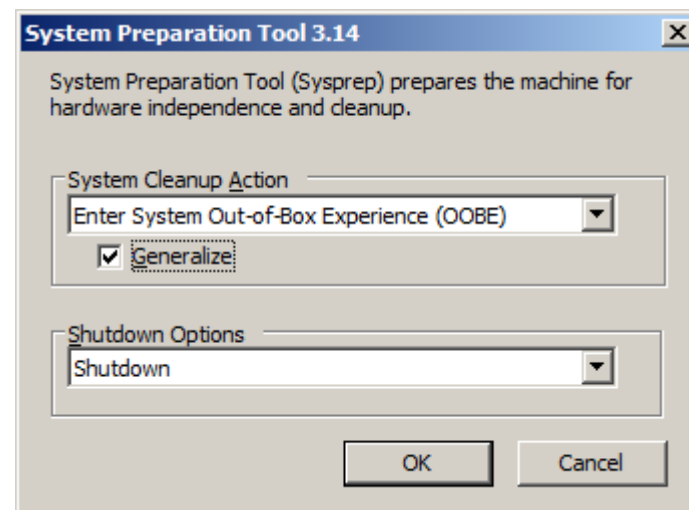
# Console Proxy VM

- Connectivity to CPVM
  - Management server
  - End-users (web browser)
- realhostip.com
- SSL certificate
  - Wildcard, or include all possible hostnames
  - a-b-c-d.yourdomain.tld
- Problem?
  - Destroy CPVM



# Templates

- eth0, or is it eth1? Or maybe p192p1?
- “sysprep” for Windows, your own solution for Linux
- Prepare in CloudStack environment?
- Can't “import” them?



# Password Reset Feature

- Reset script problems
  - DHCP client and version
- Daemon problems
  - 8080/tcp on virtual router
  - socat process, `serve_password.sh`

# User- and Meta-data

- Where is it?
- Problems?
  - Stop/start VM
  - Stop/start virtual router
  - Destroy/recreate virtual router
  - Check management-server.log

# Q&A

- Any questions?




# Top 10 Network Issues in Apache CloudStack

Kirk Kosinski  
Escalation Engineer  
Citrix Systems

# Agenda

- Introduction
- Top 10 Network Issues
- Q&A



# VLAN Issues 1

- Switch misconfiguration
  - Symptoms
  - All VLANs trunked by default? Or denied?
- Other network problems

Switch misconfiguration – common

Symptom – cannot ping across hosts, VMs cannot get DHCP sometimes (when DHCP server is on another host)

Detection – Where does the traffic stop?

XS (bridge) / KVM: tcpdump; ESXi, XS (OVS): dummy VMs? ARP, MAC address table

Switchport in trunk mode, not access

Solution – Fix the switch config, replace switches

Do VLANs need to be explicitly allowed?

Other network problems – firewall blocking ports; “weird” problems with application layer firewall (e.g. ping, nmap to 22 work, ssh fails); bad load balancer (can set “host” parameter in Global Settings to a LB, but the LB needs to actually work)



## VLAN Issues 2

- Hypervisor problems
  - NIC drivers
  - Bonding
    - xe-switch-network-backend?
  - Cabling
  - VLAN scaling
- Database hacking
  - Stop that!



Bad drivers – What, you actually want to use VLANs?

NIC bonding

Symptoms – similar to switch misconfiguration

Detection – NIC drivers / bonding – similar to switch misconfiguration, but traffic stopped elsewhere.

Bonding – check config (XS), check if traffic is dropped on the bond interface (ifconfig); disable one slave NIC, or force failover; confirm subinterfaces are on the right interface (the bond); change bond mode (active-passive vs. SLB).

Cabling – traffic “randomly” dropped

VLAN scaling – XS – high dom0 CPU, slowness


Check iptables/ebtables rules on host (KVM, XS).

DB hacking – “wrong” VLANs in use

Solution – Update drivers, replace NICs, unhack db

## Open vSwitch

- Default on XenServer / XCP
  - VMware + Nicira?
- Weirdness
- ovs-appctl



Great – as long as it works

VMware + Nicira – OVS more stable/robust on XS?

Detection – Usually something weird. Random “slowness” or dropped traffic for certain traffic – e.g. ping works, SSH or HTTP doesn't

ovs-\* commands – poor docs – use Google

Solution – Change to bridge mode (or not), patch XS, upgrade XS version, change hypervisor

## Security Groups

- KVM
  - Works out of the box (unless it doesn't)
- XenServer / XCP
  - Must enable Linux bridge back-end
  - Must install CSP (XS <6.1)
- vSphere...



Symptoms – VMs inaccessible or cannot reach anything (partial, complete)

KVM/XS – check iptables/ebtables

XS - Is the CSP installed? ARE YOU SURE?!

Some patches will blow it away.

Don't “optimize” your XS.

Host level – migrate VM to another host

## “Host” Connectivity

- Hypervisors
- System VMs
- Secondary Storage
  - Alert status is normal



Symptoms – connectivity, HA errors in management-server.log

Note – Secondary Storage “Alert” status is normal

Requirements – Mgmt to hypervisor, vice versa – varies by hypervisor

XS: SSH, HTTPS

KVM: SSH

vSphere: 443/tcp to vCenter

System VMs to mgmt – 8250/tcp (“host” param)

System VMs to Internet (ping gateway)

Mgmt to system VMs – ssh via hypervisor (KVM, XS) or direct (vSphere) – 3922/tcp

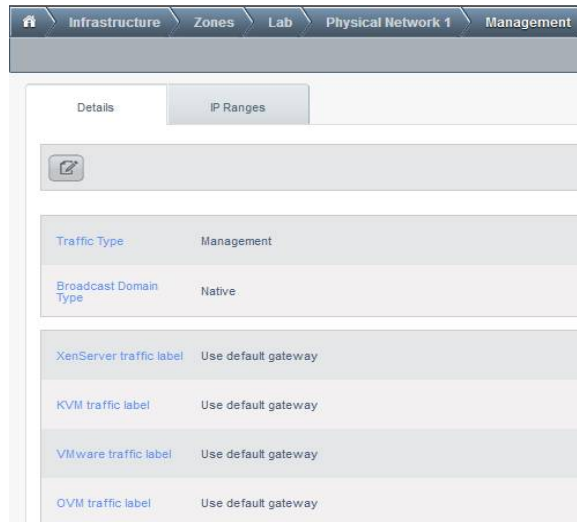
Mgmt to sec store, hypervisors to sec store – varies by hypervisor

SSVM to sec store

Mgmt to Mgmt (w/ multi-Mgmt) – 9090, 8250/tcp

## CloudStack “Physical Networks”

- Not necessarily “physical”
- Traffic labels
  - Multiple NICs
  - Mgmt. traffic on a VLAN



Typically map to NICs on the hypervisor. Can map to most hypervisor configs

KVM – bridge name

XS – network name

ESXi – vSwitch name, VLAN ID possible, to put management/private traffic on a VLAN – not supported on other hypervisors

## Console Proxy VM

- Connectivity to CPVM
  - Management server
  - End-users (web browser)
- realhostip.com
- SSL certificate
  - Wildcard, or include all possible hostnames
  - a-b-c-d.yourdomain.tld
- Problem?
  - Destroy CPVM



CPVM must reach mgmt server and hypervisors (management/private network) and end-user (public network)

CPVM proxies VNC from hypervisors to end-user

Public IPs must be accessible to end-users

CPVM uses realhostip.com domain by default – NOT a placeholder, it's a real domain

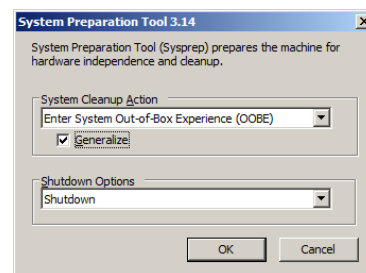
Traffic is over HTTPS using \*.realhostip.com wildcard cert by default

Change the domain and cert – must be valid!

Potential URLs are a-b-c-d.yourdomain.tld, where a-b-c-d are IPs with s/./-/g from public net

## Templates

- eth0, or is it eth1? Or maybe p192p1?
- “sysprep” for Windows, your own solution for Linux
- Prepare in CloudStack environment?
- Can't “import” them?



Templates preparation should follow best practices from OS vendor.

Use “sysprep” for Windows, custom scripts for Linux

Linux suggestions – clear udev persistent network device names, SSH keys, bash history, logs, temp files

Easier to prepare templates outside of CS – PV mode Ubuntu (XS) much easier – not always an option (e.g. slow connectivity to CS environment)

Setup password reset script

# Password Reset Feature

- Reset script problems
  - DHCP client and version
- Daemon problems
  - 8080/tcp on virtual router
  - socat process, serve\_password.sh





## User- and Meta-data

- Where is it?
- Problems?
  - Stop/start VM
  - Stop/start virtual router
  - Destroy/recreate virtual router
  - Check management-server.log

Virtual router - Apache on 80/tcp  
Standard location - /var/www/html  
Check docs for correct path to use

## Q&A

- Any questions?

