Network performance in virtual infrastructures

A closer look at Amazon EC2

Alexandru-Dorin GIURGIU

University of Amsterdam
System and Network Engineering Master

03 February 2010

Coordinators: Paola Grosso & Rudolf Strijkers

Research question

What is the network performance level offered by commercial cloud services and are there any limitations imposed by their network infrastructure?

Cloud computing

What is cloud computing?

- many definitions
- my focus is on IaaS(Infrastructure as a Service) clouds
- resources are offered as an abstracted service
- billing is done on the amount of usage
- usually the hardware is virtualized
- resource allocation is dynamic and transparent

Why cloud computing?

- easy scalability
- agility
- automation through API's
- can be cost-effective

Amazon EC2

Amazon Elastic Compute Cloud(EC2)

- since August 2006
- based on the Xen hypervisor
- easy to use API
- 3 standard instance types: small, large and extra large
  - small: 1.7GB memory, 1 EC2 compute units
  - large: 7.5GB memory, 4 EC2 compute units
  - extra large: 15GB memory, 8 EC2 compute units
- all tests were done on a small instance in the Northern Virginia region

Experimental tests

I recorded the data from the following tests in order to find the characteristics of the Amazon EC2 network:

- response time
- link throughput
- hop count
- packetloss & jitter
- local network scan on a VM instance

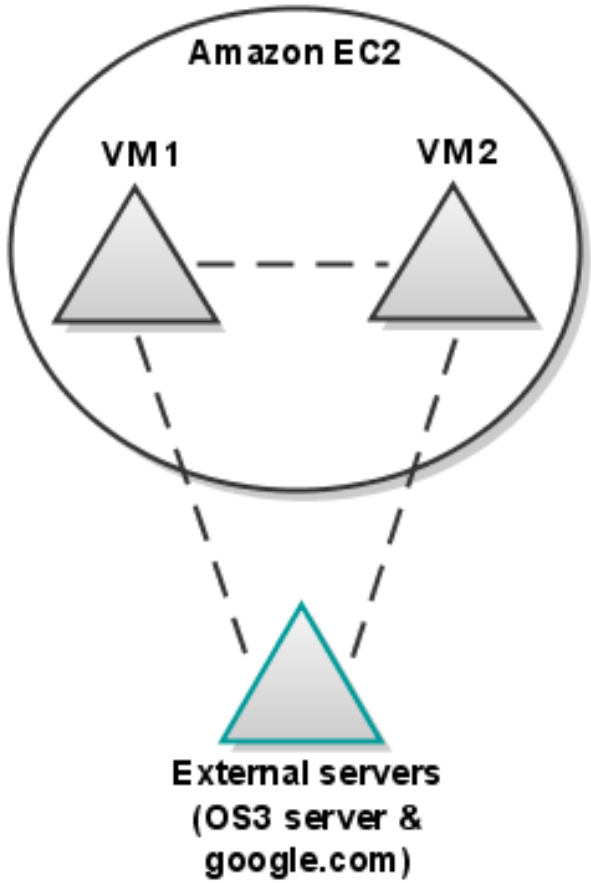

Test suite

The data gathering was automated using a test suite that ran the tests sequentially, on each pair of virtual machines.

- written in Ruby
- uses the ping, traceroute and iperf tools
- reads test configuration from a YAML structured file
- instantiates new virtual machines using the Amazon EC2 API
- saves results in a CSV file

Network topology

- structured in /23 network
- the first IP of the network is always the gateway
- networks are interconnected through variable number of routers
- each VM gets allocated in a different network

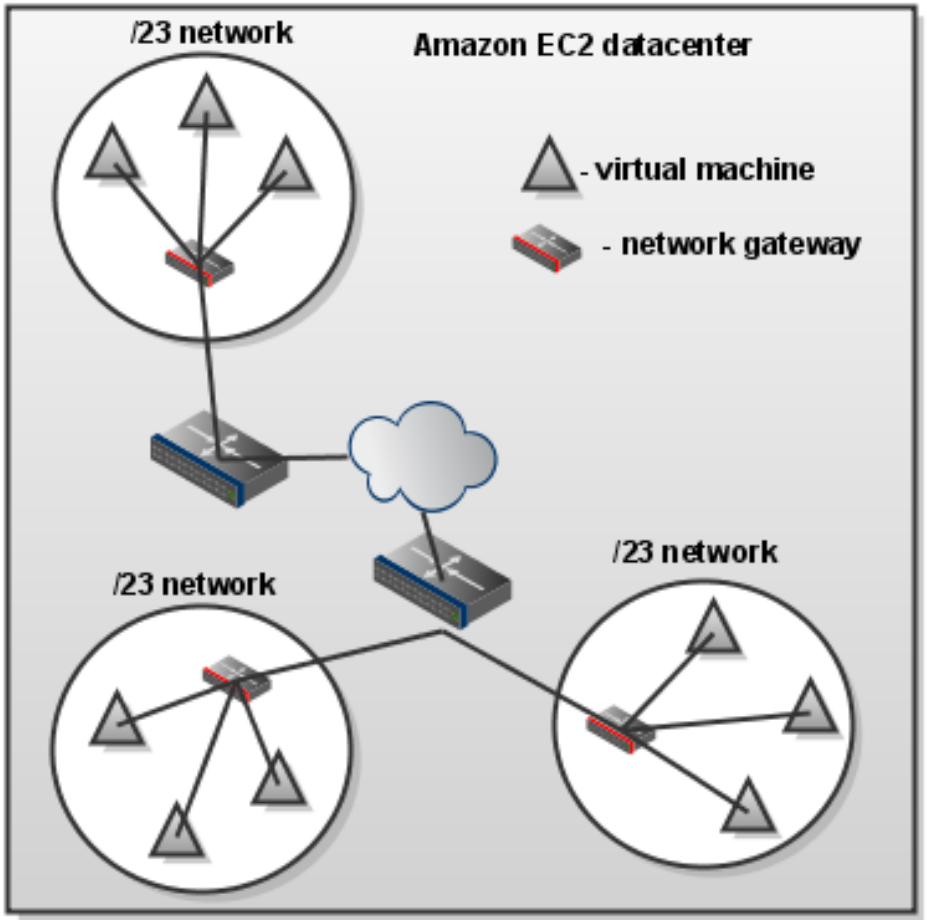

Internal latency

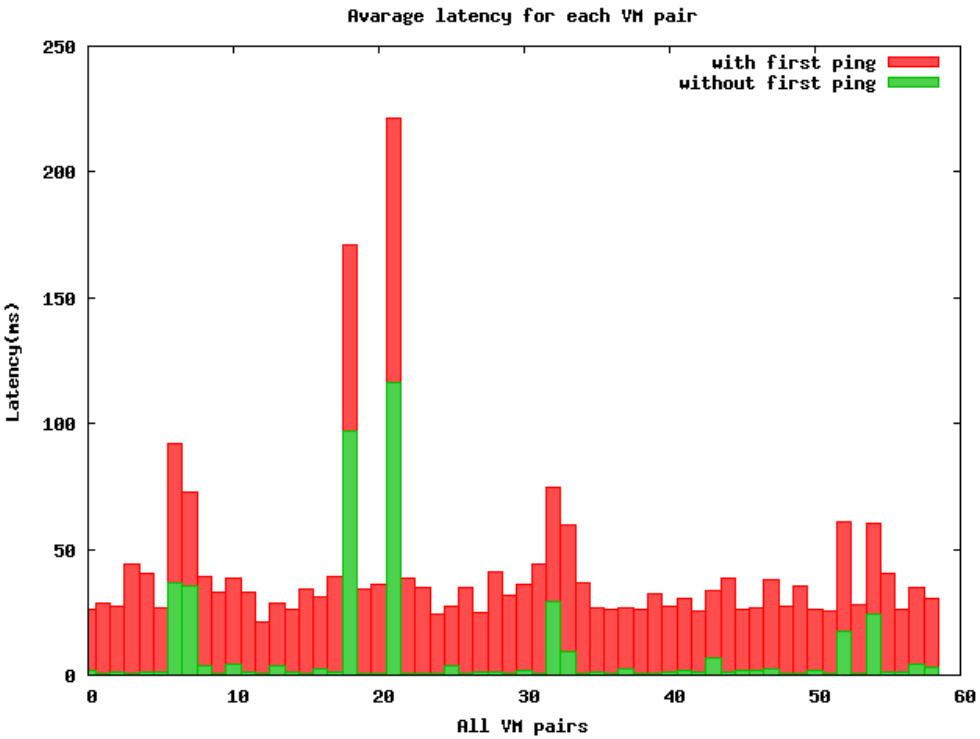

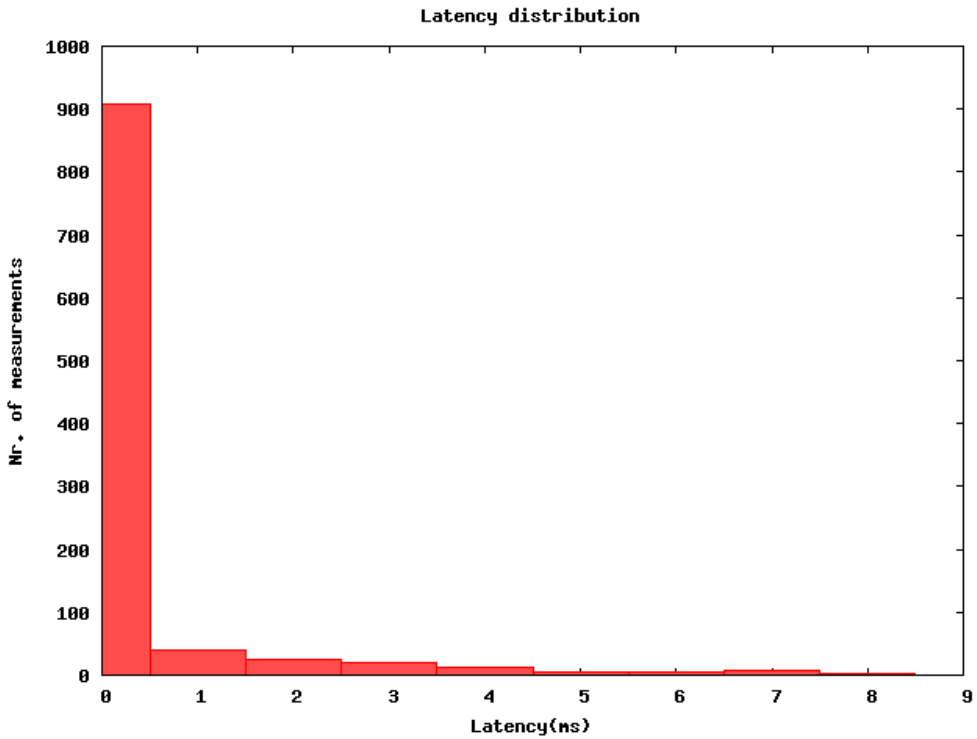

Very high averages for some of the VM pairs.

External latency

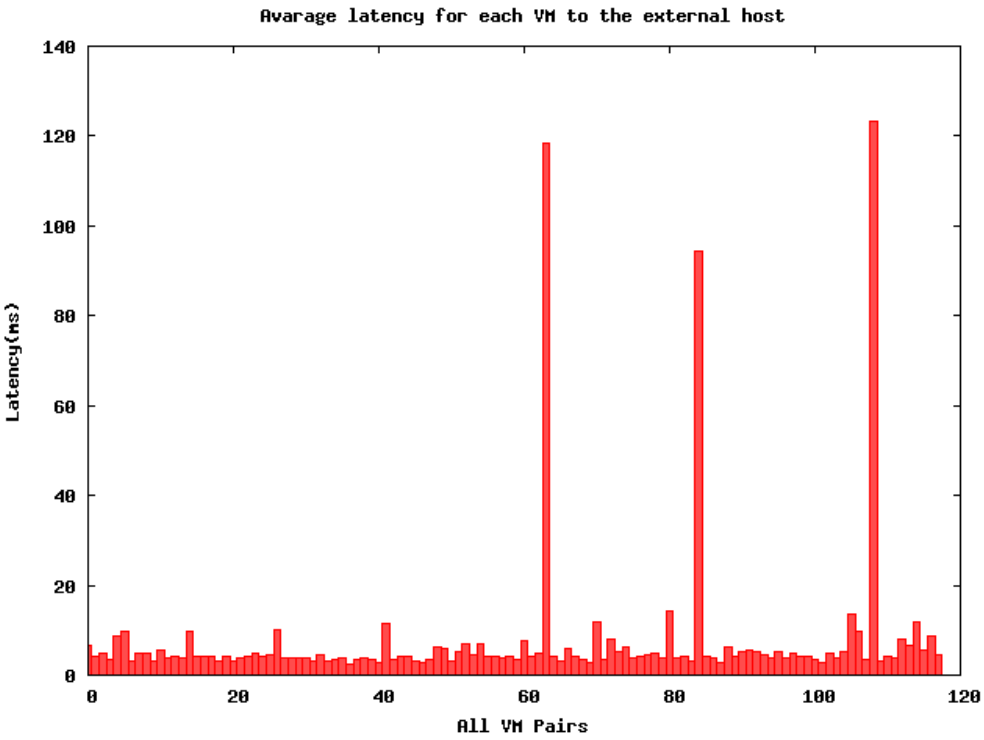

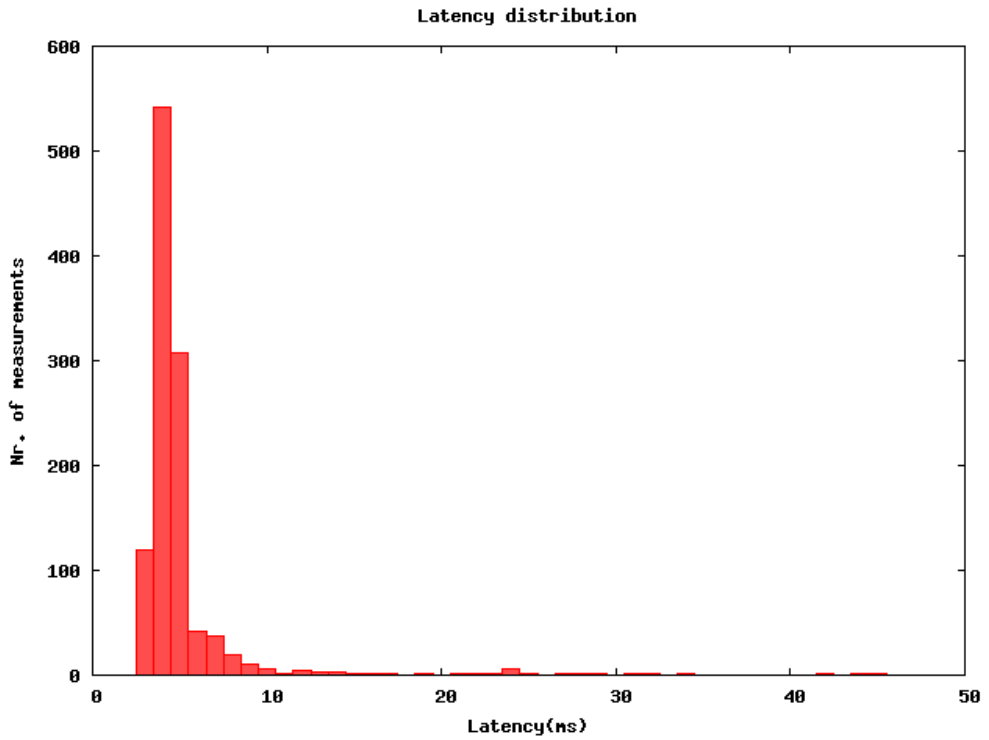

Majority of measurements have expected latency. 3 VM pairs have very high averages, probably caused by high resource usage on the physical host.

Internal bandwidth

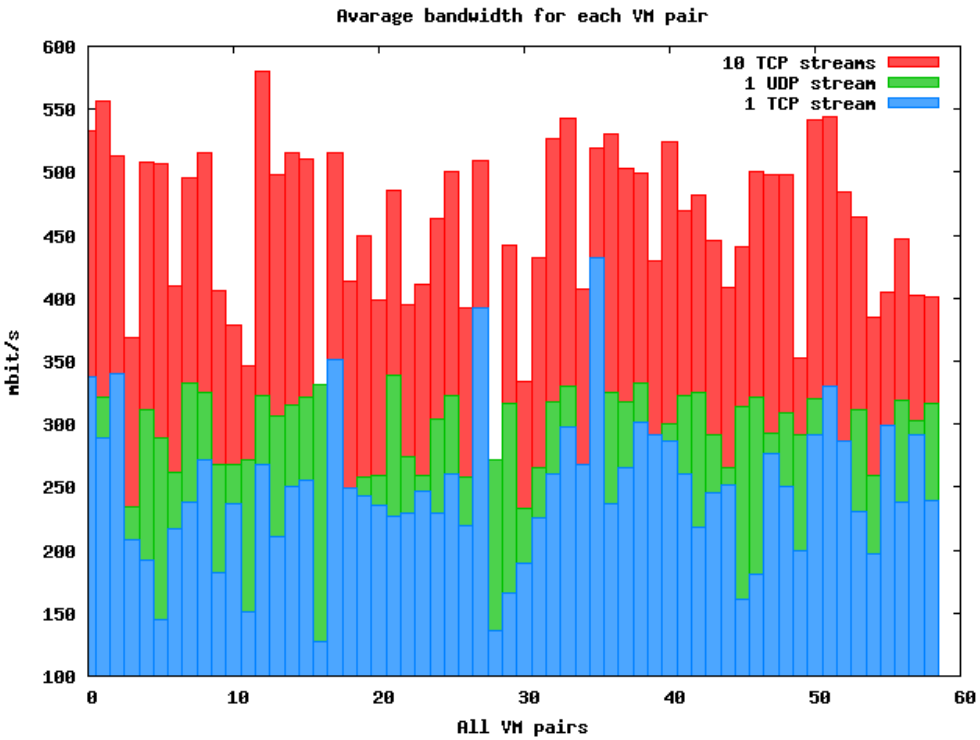

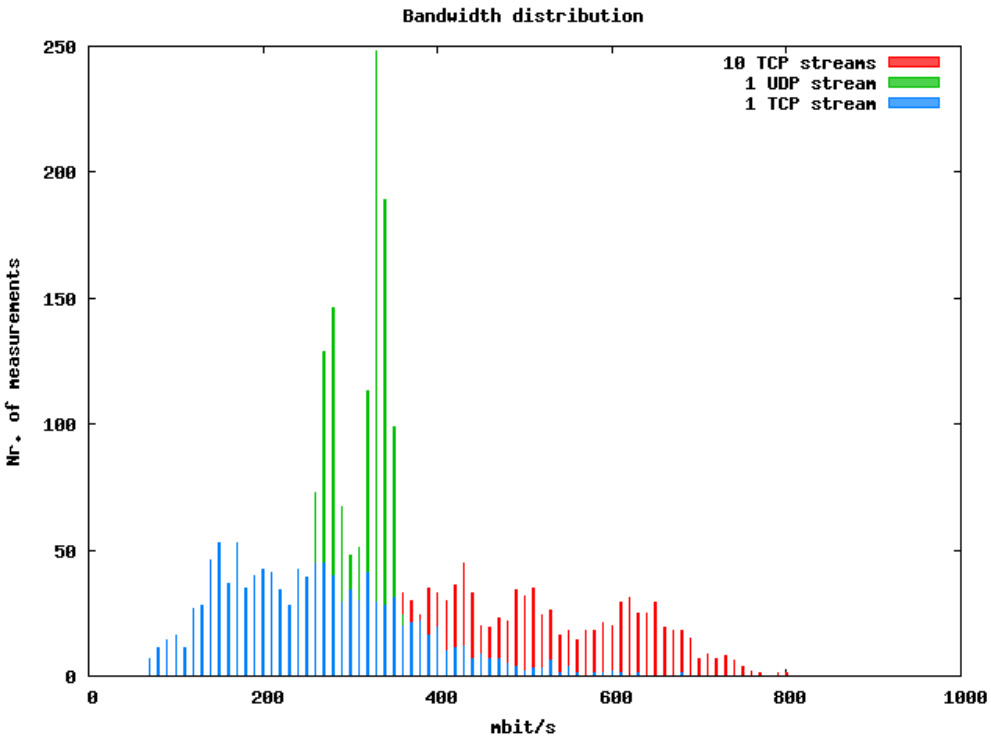

High variability, bandwidth between VM's provided on a best effort basis.

External bandwidth

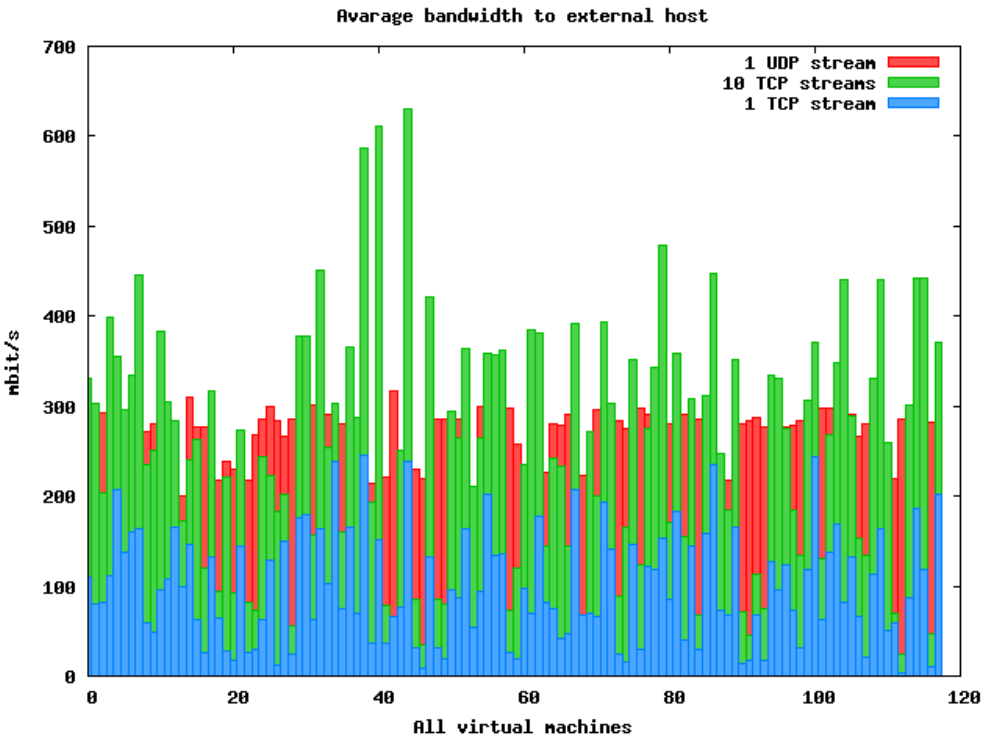

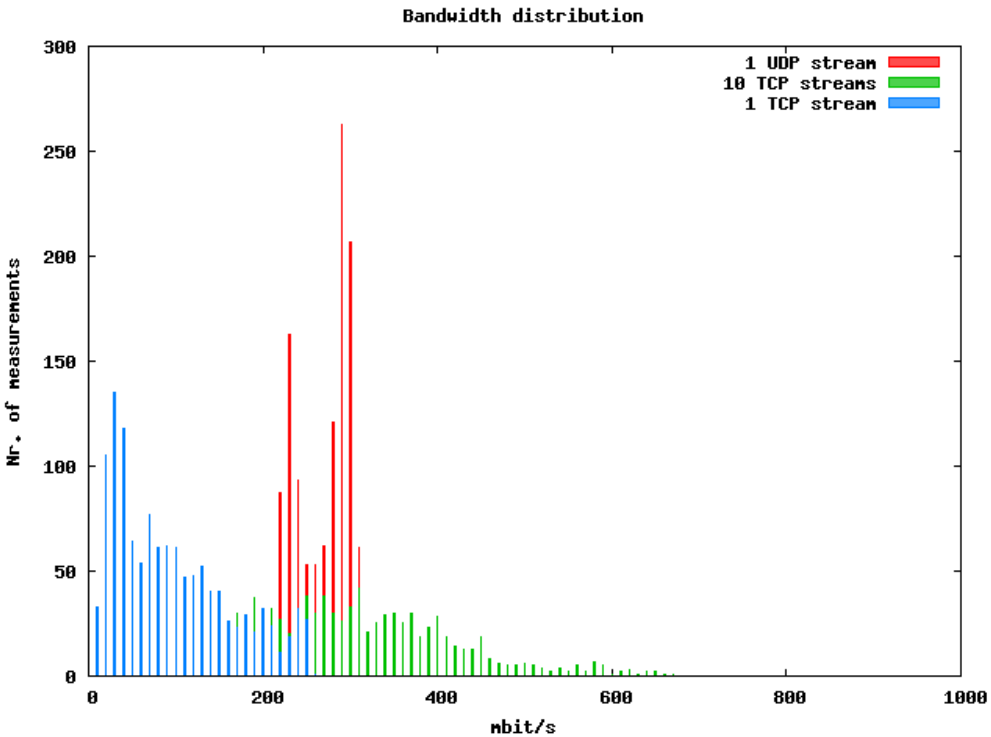

Packet loss and jitter

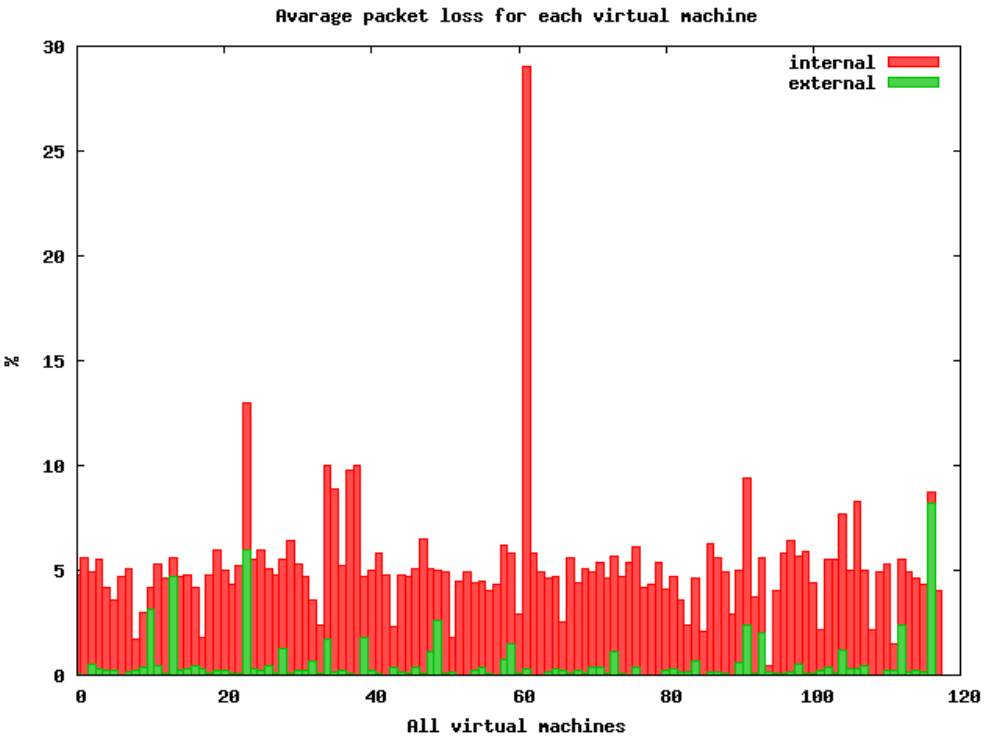

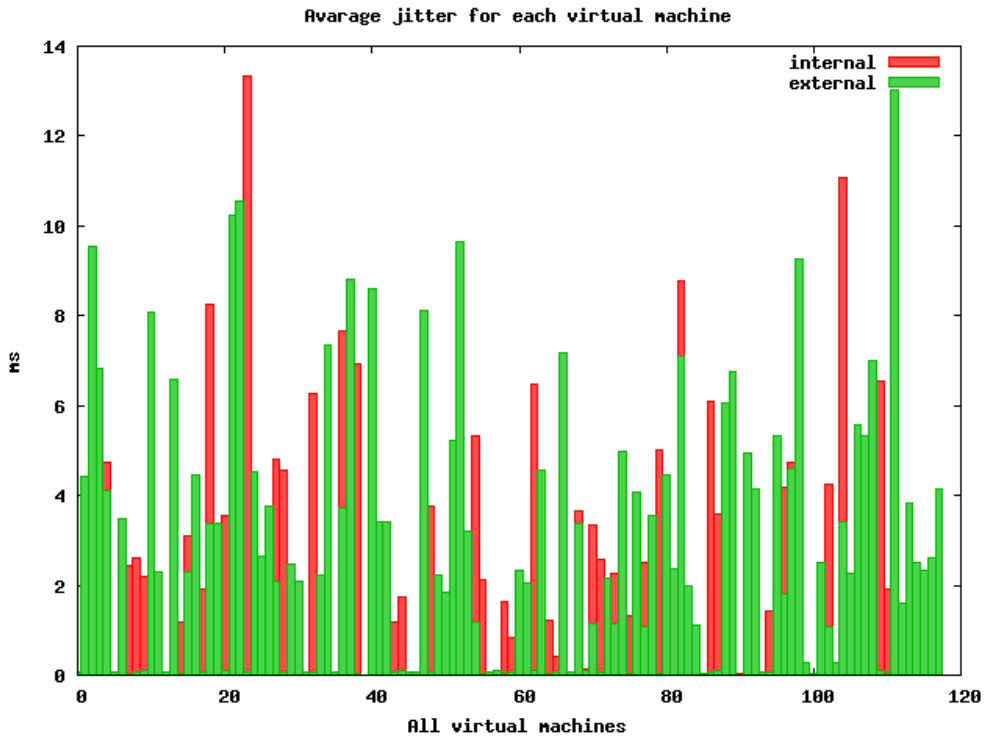

While jitter is within acceptable limits, packet loss is very high for applications like VoIP and video streaming.

Network performance

Summary

Latency
- spikes as big as 400ms between virtual machines, majority under 1ms
- variations from 2 to 170ms, majority is in the 2 to 6 ms range.
- concerning high averages for 5 of the VM pairs

Throughput
- high variability
- VM to VM - 10 TCP streams: 200 to 800 mbits, 1 TCP stream: 60 to 680 mbits, 1 UDP stream: 220 to 380 mbits
- VM to OS3 server - 10 TCP streams: 20 to 670 mbits, 1 TCP stream: 2 to 250 mbits, 1 UDP stream: 124 to 327 mbits

Jitter from 0 to 13 ms, in both internal and external tests

Packet loss average of 5% loss between VM's, as high as 27%

Network connectivity and isolation

Two kind of tests:

- layer 2 discovery using the *arp-scan* tool
- open port scan on layer 3 using the *nmap* tool

State of the network:

- no layer 2 connectivity
- only IPv4 on layer 3, IPv6 not possible
- no multicasting

Network security

- Amazon firewall implemented in the XEN hypervisor
- Virtual machines are organized in security groups
- Unrestricted communication between VM's in the same security group
- Communication between VM's from different security groups is firewalled
- Good security but at the cost of flexibility

Conclusions

- Variations in the network performance are very big. Bandwidth is offered on a best effort basis.
- Packet loss can be quite high between the VM instances which would prevent sensitive applications like VoIP and video streaming from working properly.
- Hardware usage levels already impact the network performance. With higher usage the performance will degrade even more and this is a big concern as more and more users start using cloud services.
- Network has limited flexibility with only a strict set of protocols allowed, preventing many applications from migrating to the cloud.

Possible future improvements

Network performance

- Quality of Service for critical traffic like VoIP and VoD
- allocate virtual machines on servers that are physically close to each oher
- reserve CPU time for proccessing network packets

Network flexibility

- Short term - VPN's between the virtual machines, emulates layer 2 or layer 3, big overhead
- Long term - Vlans or network virtualization technologies like OpenFlow for sharing the same physical network infrastructure

Test suite: http://nidoran.studlab.os3.nl/testsuite.tar.gz
OpenFlow project: http://www.openflowswitch.org

Thank you!

Questions?