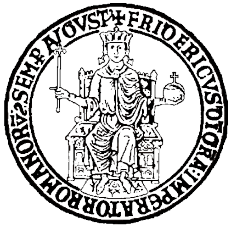

DIETI. UNIVERSITA' DEGLI STUDI DI NAPOLI FEDERICO II

NM2 NETWORKING MADE EASY

A first look at public-cloud inter-datacenter network performance

Valerio Persico, Alessio Botta, Antonio Montieri, Antonio Pescapè
University of Napoli "Federico II" (Italy) and NM2 srl (Italy)

IEEE Globecom 2016
Washington, DC, USA, December 5 2016

Cloud networks

- Companies more and more leverage cloud solutions to supply services across the Internet
- Top players have made *huge investments in networks of datacenters* to cope with this increasing demands
- These expensive investments are justified by cloud-traffic trends
- *Without high-performance networks, there would be no such thing as cloud computing**

*Mogul2012

Cloud-traffic growth

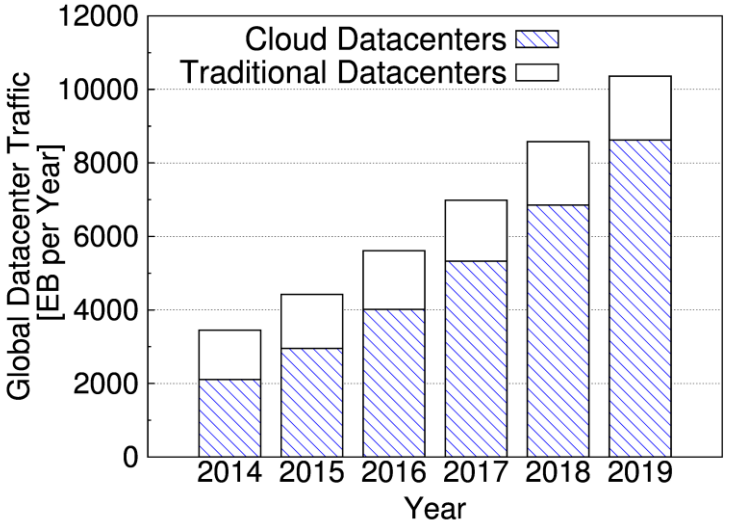

- Since 2008, most of the Internet traffic has originated or terminated in a datacenter
- Datacenter traffic is expected to continue to dominate Internet traffic
- More than three-quarters of datacenter traffic will be **cloud traffic** by 2018

*Cisco Global Cloud index

Cloud inter-datacenter networks

- Traffic between datacenters is growing faster than both **traffic to end users** and **traffic within the datacenter**

	Expected grow* (CAGR 2014 – 2019)
Intra-datacenter traffic	24%
Inter-datacenter traffic	31%
Cloud-to-user traffic	25%

*Cisco Global Cloud index

- The performance of wide-area networks interconnecting geographically distributed cloud nodes is gaining more and more interest
  - Data shuttling among clouds
  - Data replication across datacenters
  - Novel solutions leveraging the inter-datacenter WANs to support high performance applications in spreading multimedia contents world-wide

Issues for cloud customers

The cloud-to-user interface provides high-level abstraction

PROS

- No need for hardware maintenance
- Real-time resource provisioning
- Ubiquitous access to resources
- Resource automatic scaling
- No upfront investments

CONS

- Cloud providers rarely expose detailed information about performance
- Poor network provisioning impacts user experience
- Customers have to cope with limited awareness about cloud networking environment
  - Application performance unpredictability
  - Application performance variability

Contribution

- Experimental evaluation of the performance of the inter-datacenter network for the two leading public-cloud providers

- Our work depicts a clear picture of the performance of the inter-datacenter networks
  - (mainly) in terms of **throughput** and **latency**
  - considering the impact of several **configuration factors** under customer control
  - providing **insights into the infrastructure** leveraged by cloud providers
  - analyzing phenomena generated by **traffic management strategies**

Resource setup

- Two global **providers** dominating the cloud market
  - Amazon + Azure = 40% market share

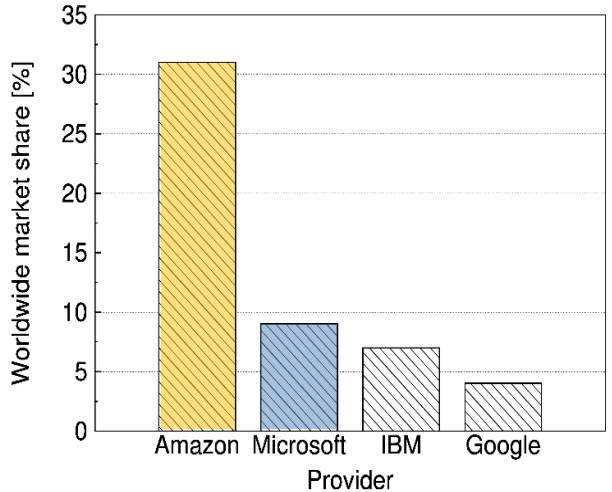

- Four **regions** hosting datacenters for both providers
  - Ireland (EU)
  - North Virginia (US)
  - Sao Paulo (SA)
  - Singapore (AP)

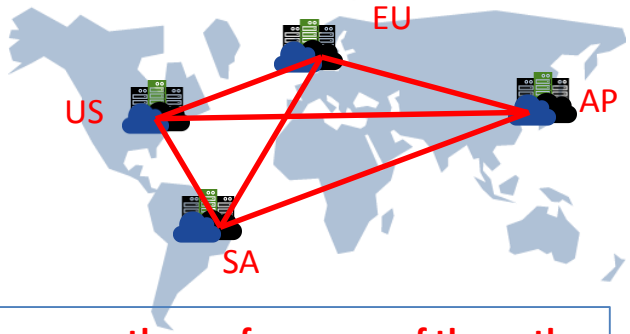

We assess the performance of the paths interconnecting these geographically distributed datacenters

Reference architecture and tools

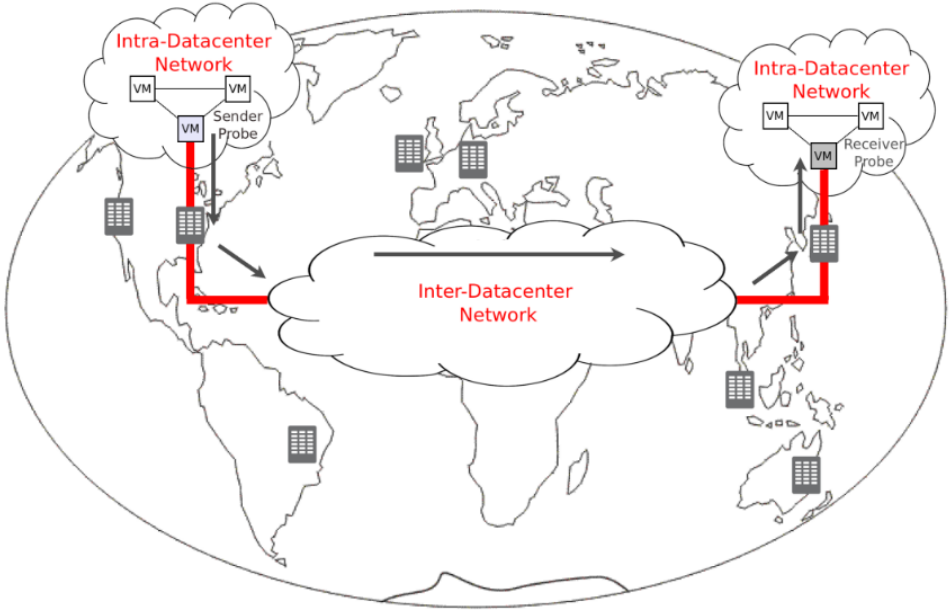

- The inter-datacenter WAN is assumed to be the bottleneck of the communication
  due to technological and physical considerations
- *Nuttcp*
  to inject **synthetic traffic** into the network
  for measuring the raw TCP and UDP throughput and latency

Dataset

Measurements performed acting as general customers

Experimentation subjected to provider fees

Details about leveraged resources

- 4 cloud regions
- 12 combinations of the 4 regions
- General purpose eXtra-Large Virtual Machines (VMs)

General information

- Data gathered from Mar. to Nov. 2015
- 300-hour-long traffic generation results
- Repeated 5-minute-long measurement experiments

Dataset freely available at:

http://traffic.comics.unina.it/cloud

Overall view of the performance

Azure infrastructure performs better than Amazon's in terms of (TCP) throughput

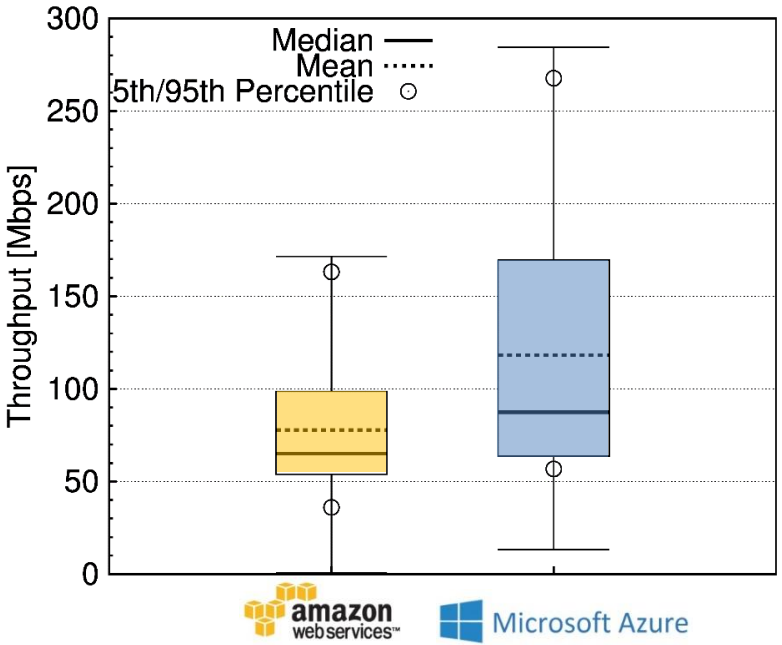

- **Average**:
  - +52%
  - 77.8 Mbps vs. 118.2 Mbps
- **Max:**
  - +65%
  - 284.5 Mbps vs. 171.6 Mbps
- **Min:**
  - 1 Mbps for Amazon
  - 13 Mbps for Azure

Region breakdown

A significant difference of performance can be observed across different regions

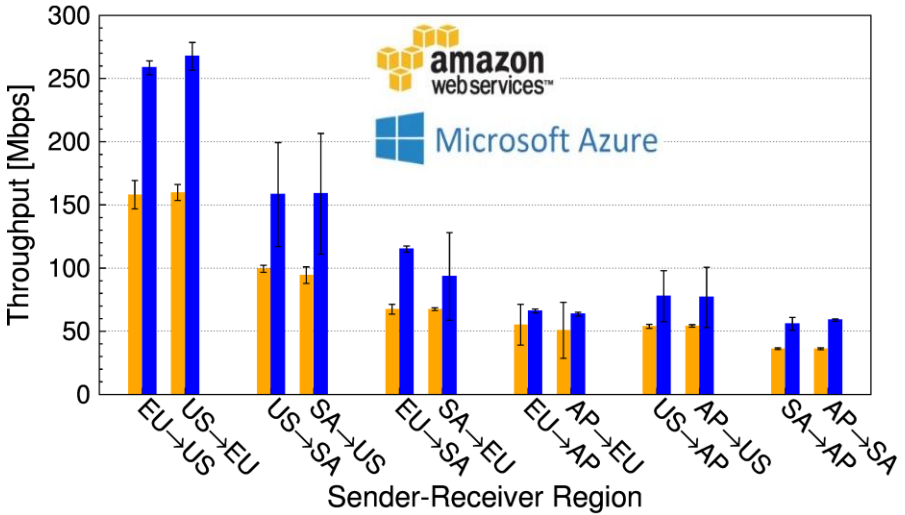

- Up to about 80% variation
- Very low variability within a fixed region pair
  - some Azure region pairs show a larger standard deviation
- Same region-pair ranking for the two providers
  - only one exception: i.e., US↔AP vs EU↔AP pair for Azure

UDP Throughput / path capacity

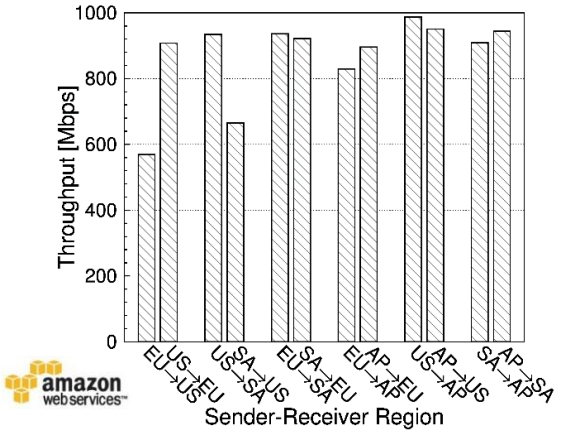

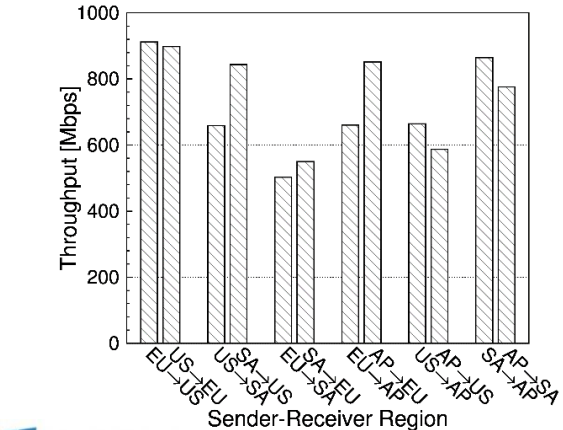

- UDP throughput is significantly higher than TCP
- In some cases UDP inter-datacenter throughput durably reaches the performance figure of attainable within the datacenter
- UDP throughput appears mainly limited by the bottlenecks imposed by providers at source

Lower TCP performance
(shown in previous analyses)
may depend upon the different number of users

Traffic-management policies

Providers enforce traffic-management policies along the path

- Throughput is not stable over time
  - performance variation within each of the 5-minute-long experiments
- Characterization of the high-to-low throughput transition
  - always happens around 100 seconds
  - in the 80% of the cases 10 GB have been transferred

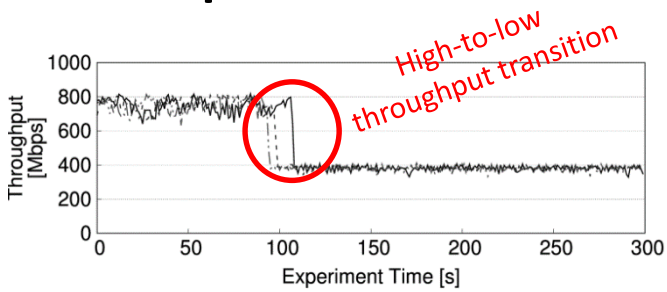

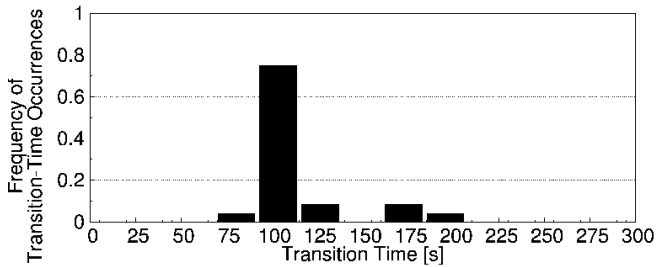

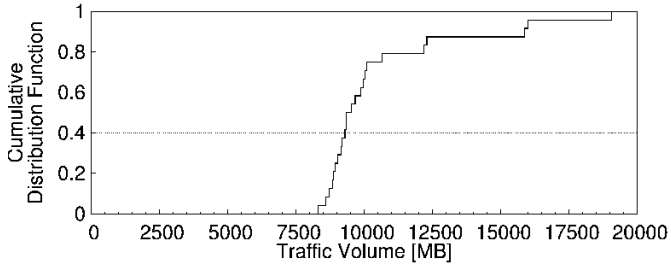

Microsoft Azure

US→EU region

Performance vs. Cost

Higher costs for the customers **do not imply** Higher network performance

The size of the VM may have no effect on inter-datacenter network performance

- Smaller VMs reported the same performance in terms of throughput
- VM advertised to have *Moderate* and *High* network performance expose the same performance in terms of throughput

Higher networking cost for the customer is related to worse network performance

- Worst performance typically related AP and SA regions
- Data transfer from AP and SA is subjected to higher costs with respect to EU and US regions
  - Up to 8× for Amazon,
  - and up to 3.2× for Azure

Network latency

Latency is very stable over time

- **CoV values**
  - smaller than 0.1 in all but two of around 70 experiments for Amazon
- **Latency values**
  - are symmetric
  - are smaller than those reported in previous works

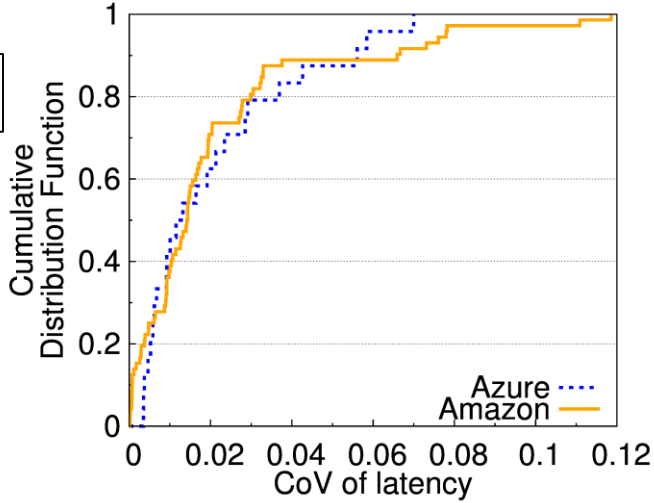

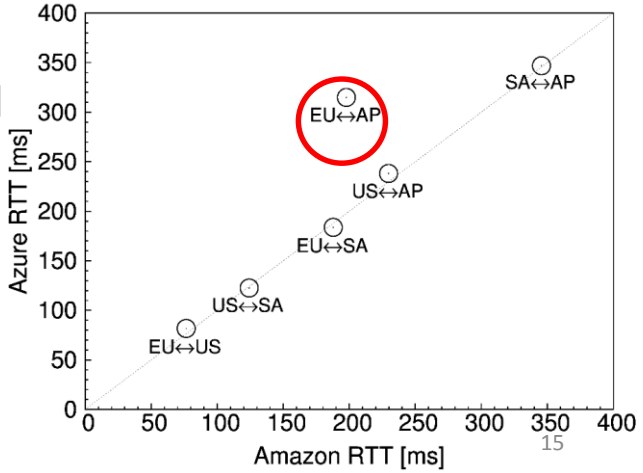

Markedly higher latency for Azure for EU←→AP

CloudSurf
a platform to monitor public-cloud networks

- Non-cooperative approaches
  - adoption of the point of view of the general customer
  - no need for access to information restricted to the provider
- Comprehensiveness
  - different providers
  - different use scenarios
- Predictability of experimental costs
  - computation cost
  - network-usage cost
- Results sharing
  - community results repository
- Ease of use
  - on demand measurements
  - no specific monitoring skills needed
- Public availability
  http://traffic.comics.unina.it/cloudsurf

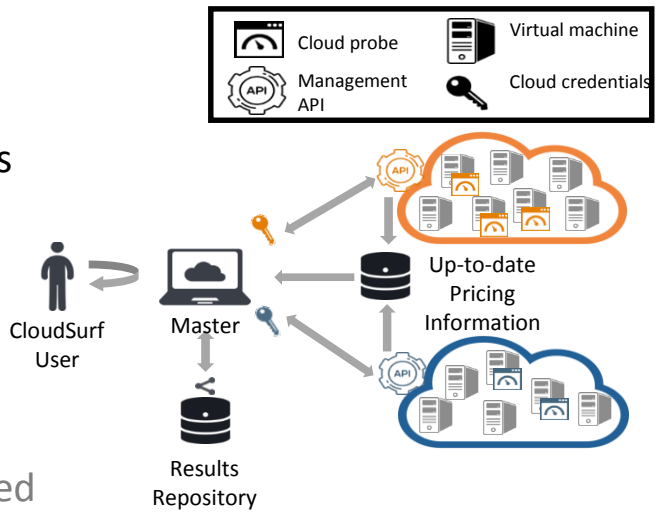

Conclusion

- Cloud inter-datacenter networks are gaining more and more interest
- In this work we have provided an experimental assessment of these network for Amazon and Azure
- Main findings
  - Azure inter-datacenter network performs better than Amazon (+52%, on average) probably because of the smaller number of customers
  - Counterintuitively, higher costs do not imply better network performance
  - Network latency is comparable for the two providers, with the remarkable exception of the path interconnecting AP to EU
  - Traffic engineering policies enforced have been identified and characterized

Questions?

valerio.persico@unina.it

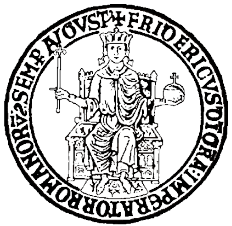

DIETI. UNIVERSITA' DEGLI STUDI DI NAPOLI FEDERICO II

NM2 NetworkingMadeEasy

This slide left intentionally blank

Performance (a)symmetry

Performance between regions is (typically) symmetric

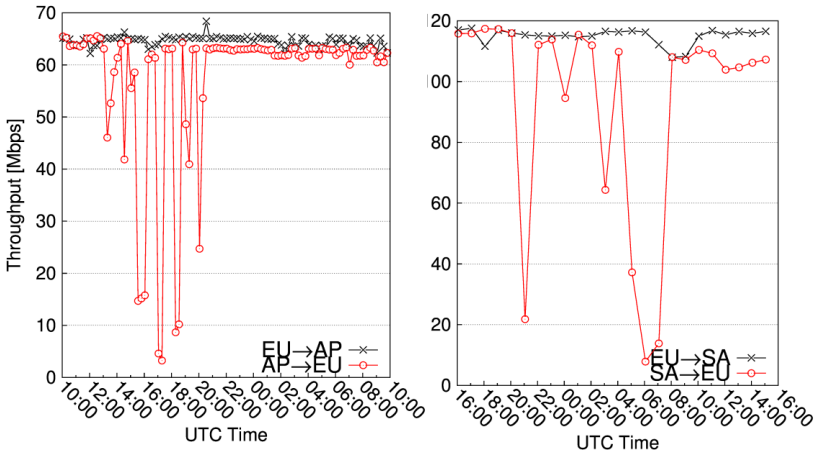

Some notable exceptions found

- Severe performance degradations lasting for several hours
- Throughput dropping down to values smaller than 5 Mbps for only one direction of the communication

Impact of placement inside a region

- The AZ* does not clearly impact the achievable throughput.
- In a limited number of cases severe performance degradations involving only one AZ
  - throughput dropped down to values smaller than 5 Mbps
  - lasting for several hours

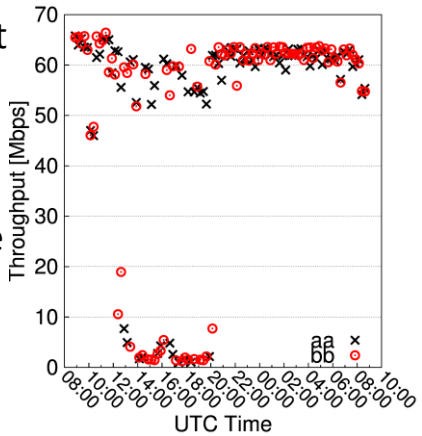

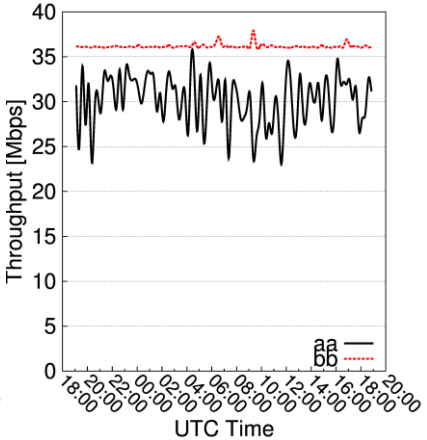

***Availability zone (AZ)**: isolated locations made available inside a region by Amazon