DIETI. UNIVERSITA' DEGLI STUDI DI NAPOLI FEDERICO II

NM2 NetworkingMadeEasy

CloudSurf:

a platform for monitoring public-cloud networks

Valerio Persico*, Antonio Montieri◊, Antonio Pescapè*,◊

*University of Napoli "Federico II" (Italy)

◊NM2 srl (Italy)

IEEE RTSI 2016

Bologna, Italy, September 9th 2016

(Public) Cloud Paradigm

An increasing number of services and applications rely on Public Clouds

- *XaaS*: everything as a service
- Pay-as-you-go resources
- No upfront investments
- Real-time provisioning
- Ubiquitous access
- Autoscaling

Cloud infrastructures backed by huge investments from the providers

- Complex and continuously evolving infrastructures
- *Hot research topic*

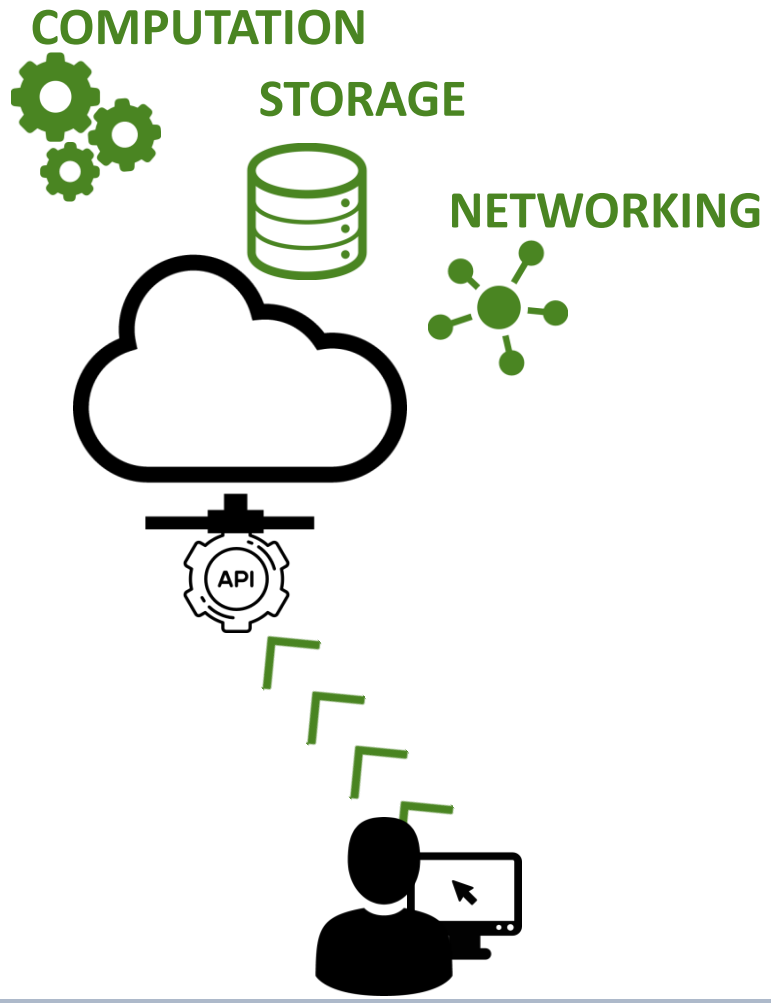

Cloud Networking, a hot research topic

Cloud network performance is critical

- Without high-performance networks, there would be no such thing as cloud computing*
- The network can be the bottleneck for computation (e.g. scientific computing, video processing, etc.)
- Poor network provisioning can severely compromise user experience

Cloud traffic is rapidly growing

- Since 2008, most of the IP traffic has originated or terminated in a datacenter
- More than 83% of datacenter traffic will be cloud traffic by 2019**
- Public cloud is growing faster than private cloud

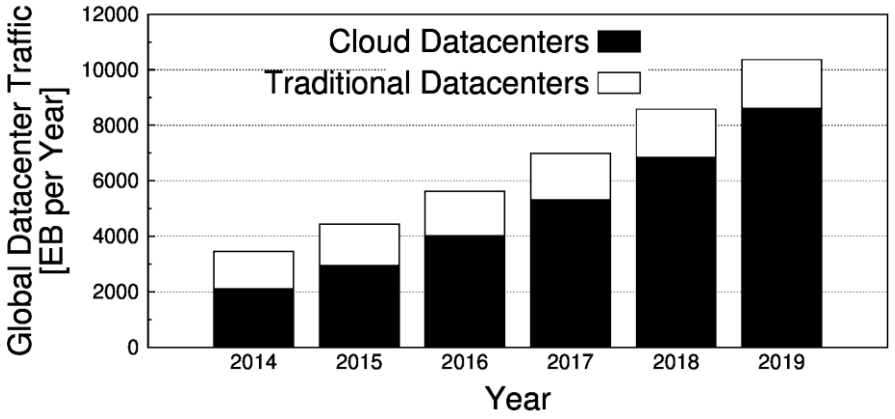

**Mogul and Popa, 2012*

***Cisco Global Cloud Index: Forecast and Methodology 2014-2019*

Public Cloud Networks

Investments aim at improving the performance of the cloud network infrastructure in all its composing areas

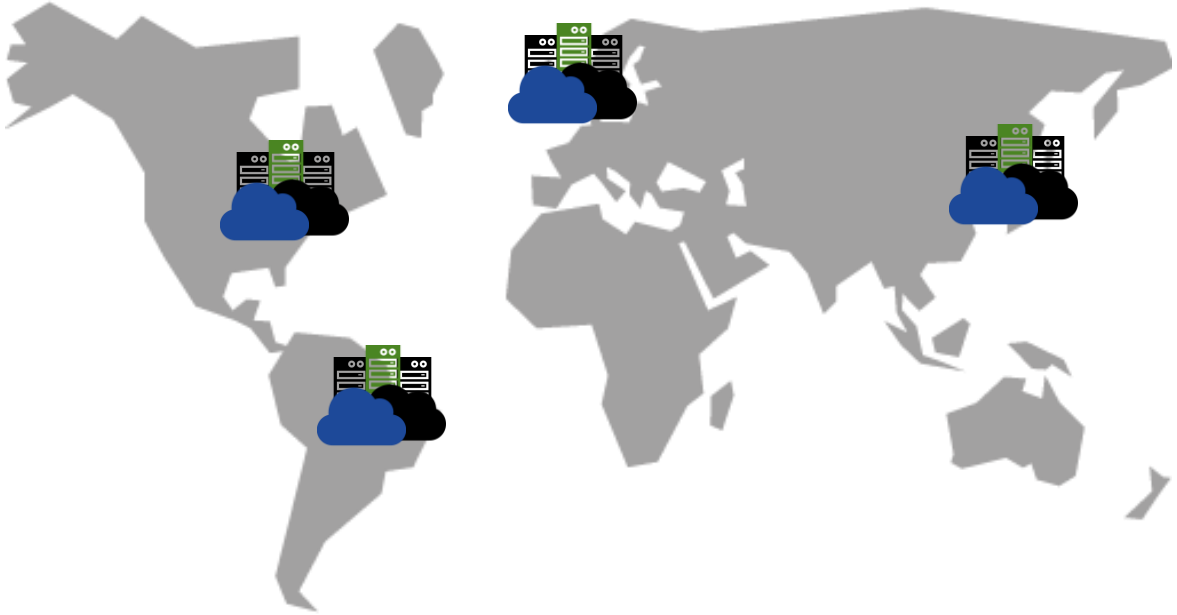

Public Cloud Networks

Investments aim at improving the performance of the cloud network infrastructure in all its composing areas

- Intra-datacenter

Public Cloud Networks

Investments aim at improving the performance of the cloud network infrastructure in all its composing areas

- Intra-datacenter
- Inter-datacenter

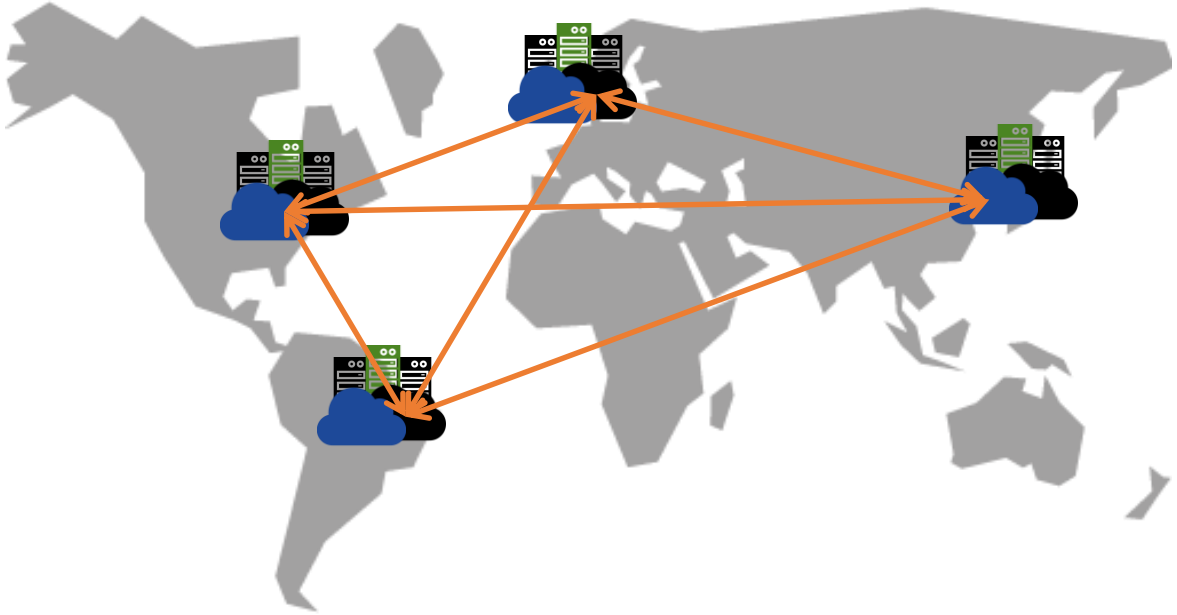

Public Cloud Networks

Investments aim at improving the performance of the cloud network infrastructure in all its composing areas

- Intra-datacenter
- Inter-datacenter
- Cloud-to-user

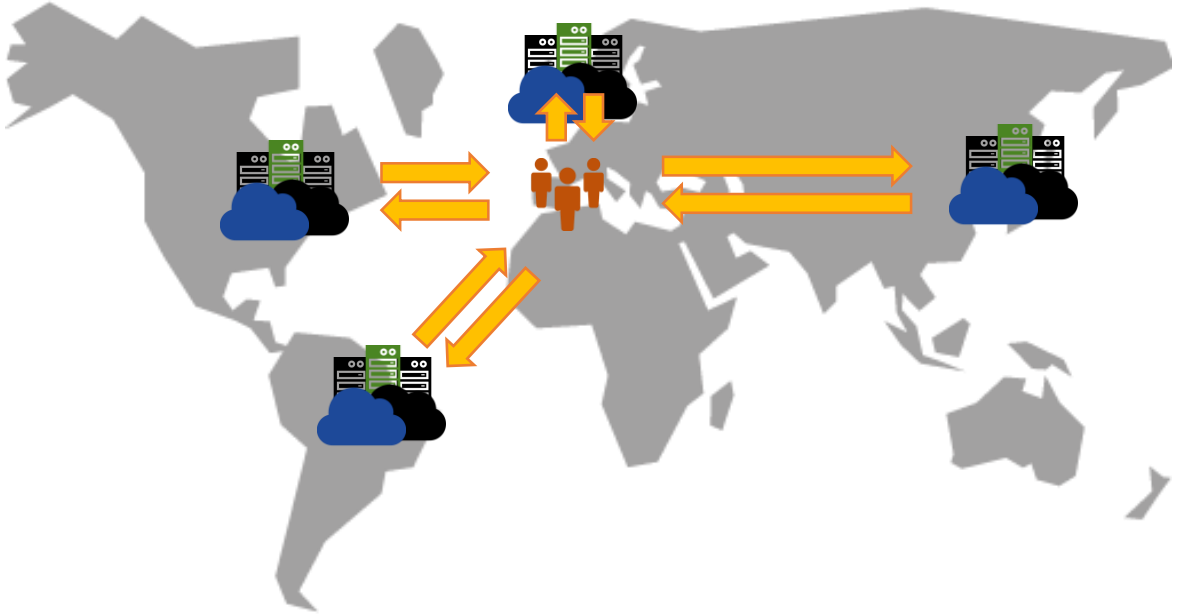

What kind of information about public-cloud networks customers are provided with?

Instance Type	vCPU	Memory (GiB)	Storage (GB)	Networking Performance
t2.micro	1	1	EBS Only	Low to Moderate
t2.small	1	2	EBS Only	Low to Moderate
t2.medium	2	4	EBS Only	Low to Moderate
m3.medium	1	3.75	1 x 4 SSD	Moderate
m3.large	2	7.5	1 x 32 SSD	Moderate
m3.xlarge	4	15	2 x 40 SSD	High
m3.2xlarge	8	30	2 x 80 SSD	High

amazon web services™

Specifications	Extra Small	Small	Medium	Large	Extra Large
CPU	1.0 GHz	1.6 GHz	2 X 1.6 GHz	4 X 1.6 GHz	8 X 1.6 GHz
Memory	768 MB	1.75 GB	3.5 GB	7 GB	14 GB
VM Local Storage	20 GB	225 GB	490 GB	1,000 GB	2,040 GB
Network I/O Performance	Low	Moderate	High	High	High
Allocated Bandwidth	5 Mbps	100 Mbps	200 Mbps	400 Mbps	800 Mbps
Cost per Hour	$0.05	$0.12	$0.24	$0.48	$0.96

Microsoft Azure

What kind of information about public-cloud networks customers are provided with?

Instance Type	vCPU	Memory (GiB)	Storage (GB)	Networking Performance
t2.micro	1	1	EBS Only	Low to Moderate
t2.small	1	2	EBS Only	Low to Moderate
t2.medium	2	4	EBS Only	Low to Moderate
m3.medium	1	3.75	1 x 4 SSD	Moderate
m3.large	2	7.5	1 x 32 SSD	Moderate
m3.xlarge	4	15	2 x 40 SSD	High
m3.2xlarge	8	30	2 x 80 SSD	High

amazon web services™

Specifications	Extra Small			Large	Extra Large
CPU				4 X 1.6 GHz	8 X 1.6 GHz
Memory				GB	14 GB
VM Local Storage				0	2,040 GB
Network I/O Performance				n	High
Allocated Bandwidth	5			400 Mbps	800 Mbps
Cost per Hour	$0.05			$0.48	$0.96

Microsoft Azure

Public cloud network monitoring

Effective, efficient, and accurate **monitoring** is constantly needed to manage increasing complexity

- *Providers* rarely make promises or expose details about the network
- *Customers* have to cope with the limited awareness about cloud networking environments
  - Performance unpredictability
  - Performance variability
  - Hidden datacenter topology and virtual machine (VM) location

Public cloud network monitoring

Effective, efficient, and accurate **monitoring** is constantly needed to manage increasing complexity

- *Providers* rarely make promises or expose details about the network
- *Customers* have to cope with the limited awareness about cloud networking environments
  - Performance unpredictability
  - Performance variability
  - Hidden datacenter topology and virtual machine (VM) location

- *Non-cooperative* monitoring approaches
  - do not require access to any restricted information
  - adopt the point of view of the general consumer
- *Customers* are able to
  - validate the (poor) information supplied by the provider
  - perform informed choices among different services or different providers

The blurred picture of the cloud network

- **Cooperative approaches**
  - privileged points of view (e.g., SNMP logs, datacenter traffic traces)
  - data privacy-concerns
  - datasets not publicly available

- **Non-cooperative approaches**
  - analyses rarely focused on the performance of the network
  - different methodologies adopted (measurement tool, observation period)
  - limited and non-univocally defined scenarios

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

CloudSurf
architecture

Main components

- Cloud probes
- Master
- Results repository

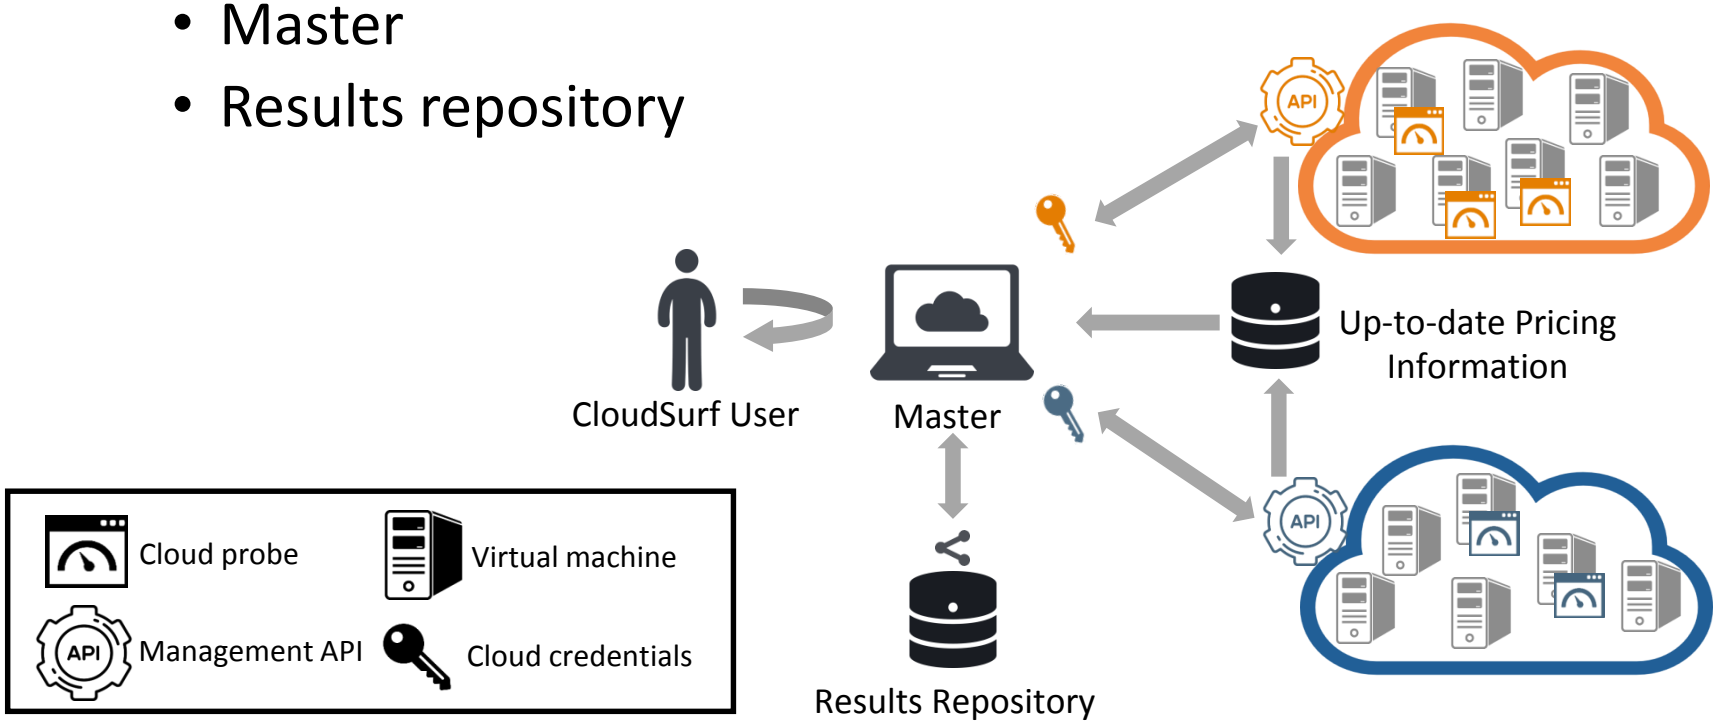

CloudSurf
architecture

Main components

- Cloud probes
- Master
- Results repository

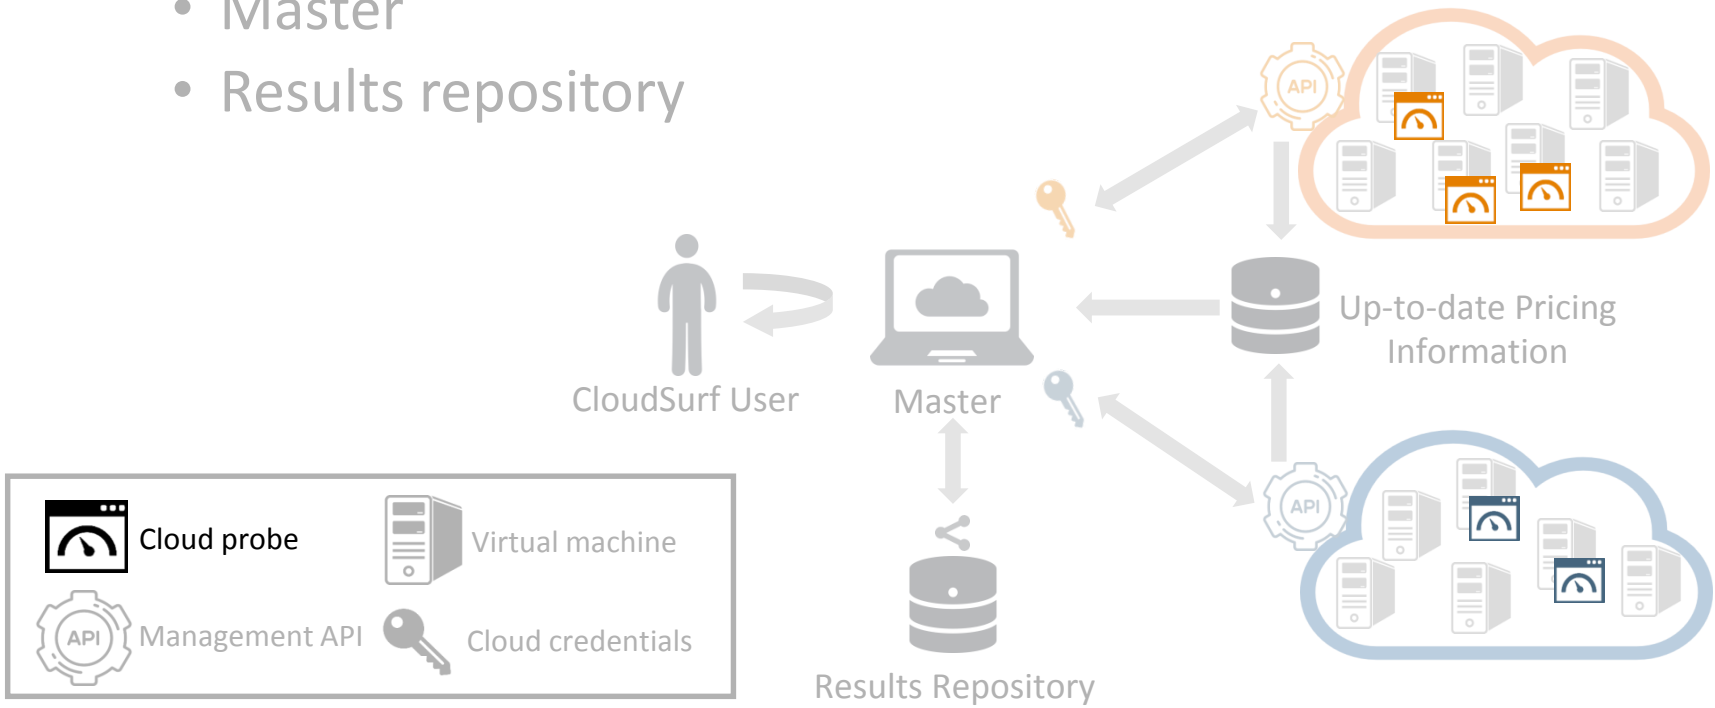

Cloud probes

Remote measurement servers

- deployed on demand by the master through the IaaS model
- integrate active monitoring tools
- expose services through the *XML-RPC* protocol

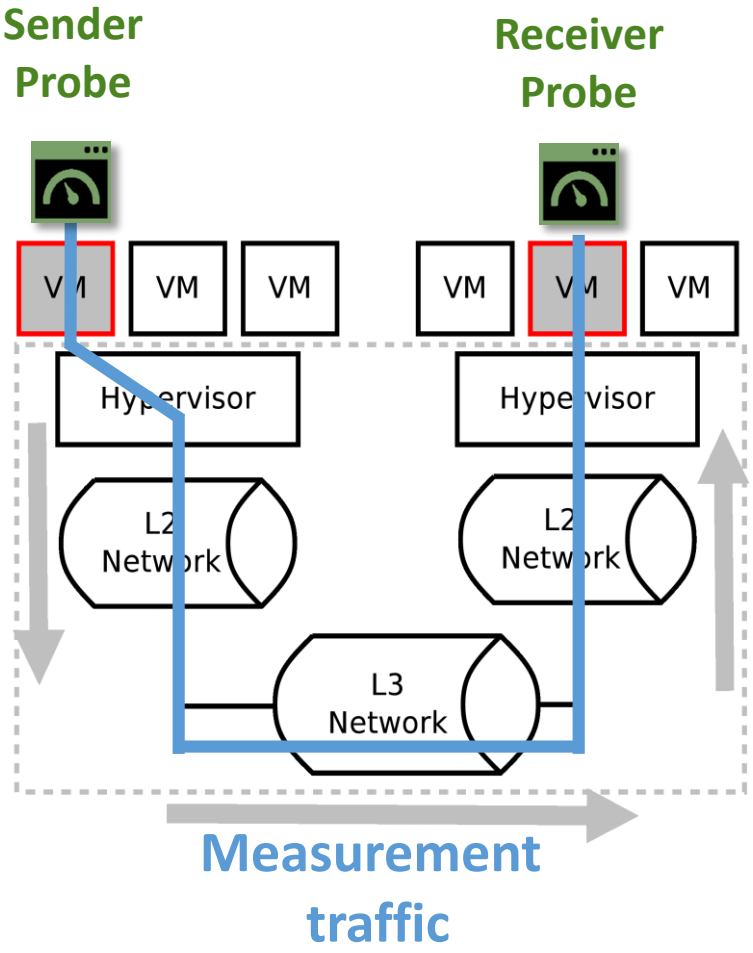

CloudSurf
architecture

Main components

- Cloud probes
- Master
- Results repository

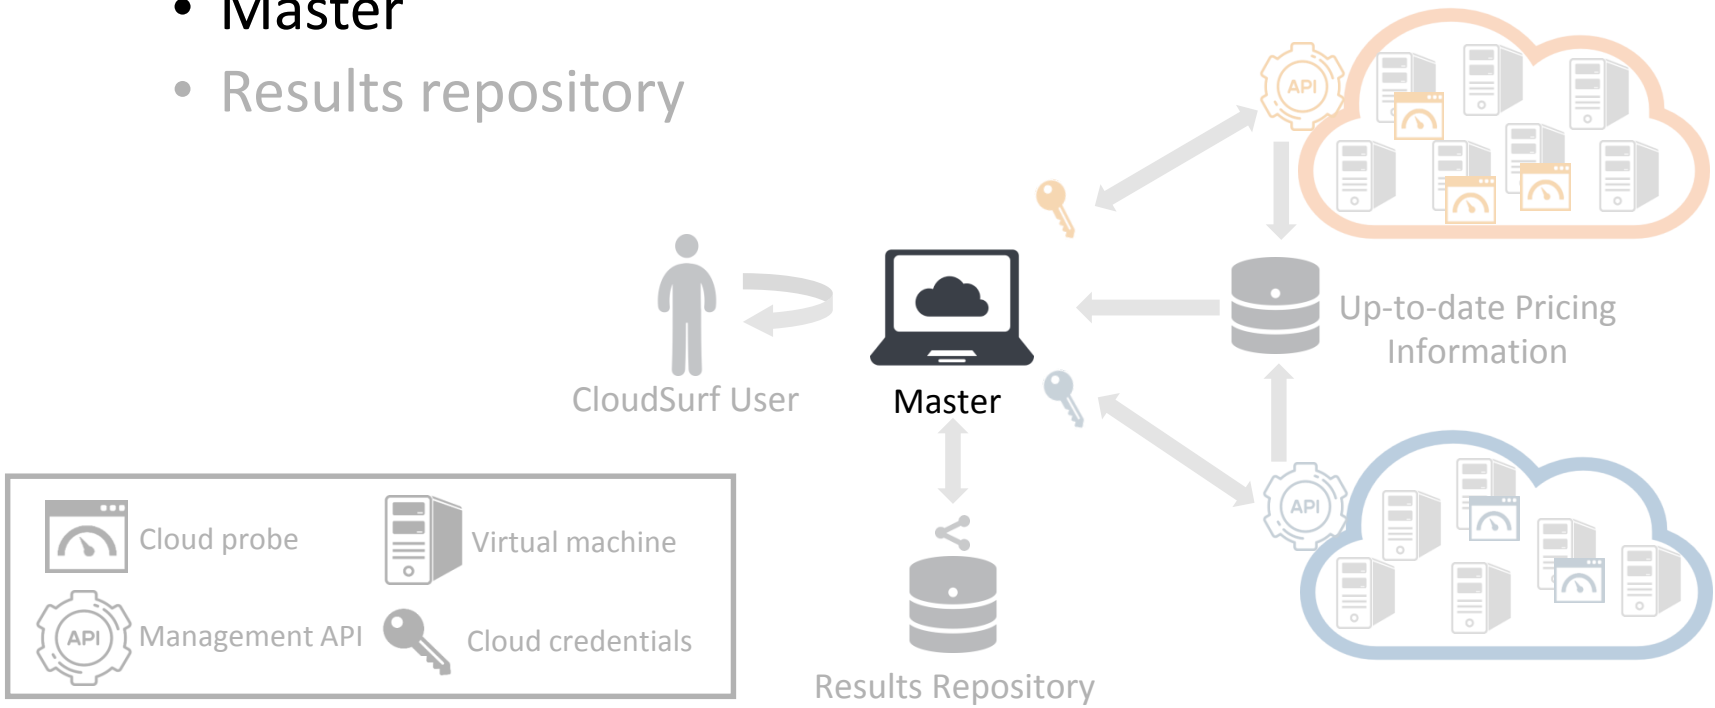

Master

- **Orchestrates the overall monitoring process**
  - Experiment-cost estimation
  - Cloud-environment setup
  - Probe deployment
  - Experiment life-cycle management

- **Is directly interfaced to the customer**
  - *Interactive* mode vs. *batch* mode
  - Usage is independent from heterogeneous provider's APIs
  - A set of *factors* of interest uniquely identifies scenarios

Experiment-cost estimation

$$Exp_{cost} = \left\lceil \frac{D}{3600} \right\rceil \times \left(C_{VM}^{sender} + C_{VM}^{receiver} \right) + R \times C_{traffic}$$

- $D[\mathrm{s}]$: experiment duration
- $C_{VM}^{sender}[^{€}/_{h}]$: hourly cost of the sender VM
- $C_{VM}^{receiver}[^{€}/_{h}]$: hourly cost of the receiver VM
- $R[^{GB}/_{s}]$: rate of measurement traffic
- $C_{traffic}[^{€}/_{GB}]$: data transfer charge

Master

- **Orchestrates the overall monitoring process**
  - Experiment-cost estimation
  - Cloud-environment setup
  - Probe deployment
  - Experiment life-cycle management

- **Is directly interfaced to the customer**
  - *Interactive* mode vs. *batch* mode
  - Usage is independent from heterogeneous provider's APIs
  - A set of *factors* of interest uniquely identifies scenarios

Cloud-environment setup

BEFORE THE MEASUREMENT ACTIVITIES

- managing user's credentials
- configuring firewalling rules
- launching VMs
- VM status check
- deploying the monitoring probes onto VMs

AFTER THE MEASUREMENT ACTIVITIES

- decommissioning the VMs
- resetting the state of the cloud-environment

Master

- **Orchestrates the overall monitoring process**
  - Experiment-cost estimation
  - Cloud-environment setup
  - Probe deployment
  - Experiment life-cycle management

- **Is directly interfaced to the customer**
  - *Interactive* mode vs. *batch* mode
  - Usage is independent from heterogeneous provider's APIs
  - A set of *factors* of interest uniquely identifies scenarios

Factors to identify scenarios

- Region e.g.,
  - North Virginia (US)
  - Ireland (EU)
  - Singapore (AP)
  - Sao Paulo (SA)

- VM type and size e.g.,
  - General purpose, compute optimized, storage optimized
  - Medium (M), Large(L), Extra Large (XL)

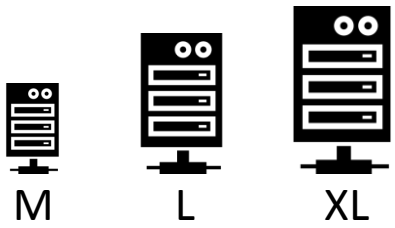

- Configuration
  - Availability zone
  - Affinity Group
  - Virtual Network

DEPLOYMENT FACTORS

- Transport protocol
- Packet size and rate
- Communication channel
- VM relocation and probe placement

EXPERIMENTAL FACTORS

Factors to identify scenarios

- Region e.g.,
  - North Virginia (US)
  - Ireland (EU)
  - Singapore (AP)
  - Sao Paulo (SA)

- VM type and size e.g.,
  - General purpose, compute optimized, storage optimized
  - Medium (M), Large(L), Extra Large (XL)

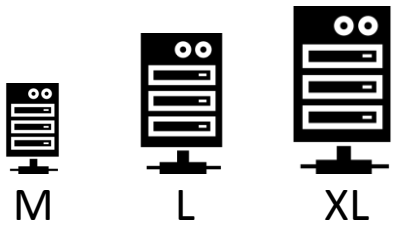

COMMON FACTORS

- Configuration
  - Availability zone
  - Affinity Group
  - Virtual Network

PROVIDER-SPECIFIC FACTORS

- Transport protocol
- Packet size and rate
- Communication channel
- VM relocation and probe placement

Factors to identify scenarios

- Region e.g.,
  - North Virginia (US)
  - Ireland (EU)
  - Singapore (AP)
  - Sao Paulo (SA)

- VM type and size e.g.,
  - General purpose, compute optimized, storage optimized
  - Medium (M), Large(L), Extra Large (XL)

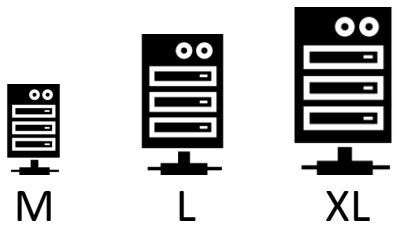

DIRECTLY IMPACT COSTS

- Configuration
  - Availability zone
  - Affinity Group
  - Virtual Network
- Transport protocol
- Packet size and rate
- Communication channel
- VM relocation and probe placement

Results repository

- Main components
  - Cloud probes
  - Master
  - Results repository

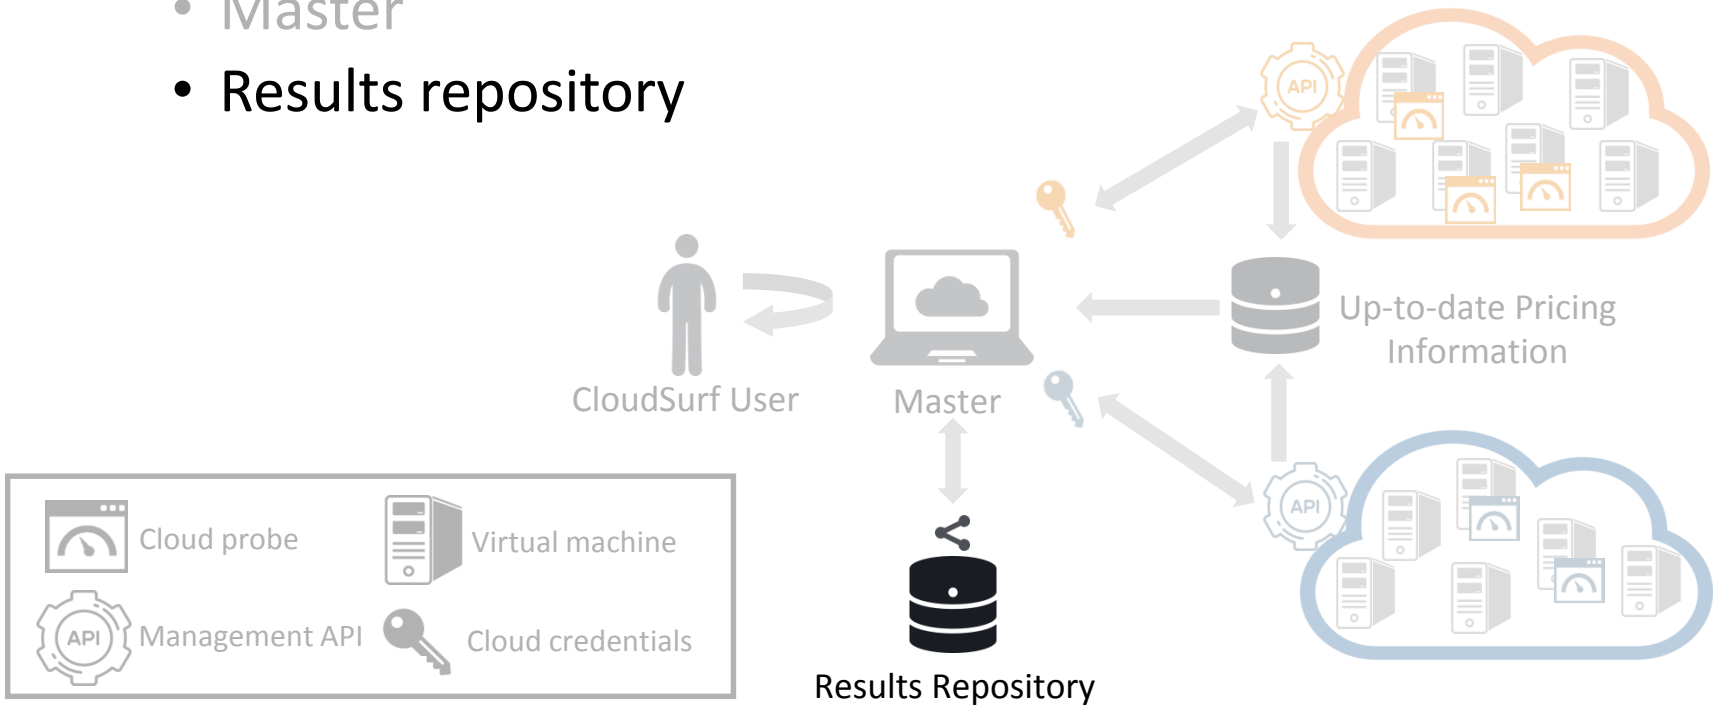

Results repository

Community archive where results are uploaded by the master

- JSON-encoded homogenous format
  to reduce the heterogeneity of the raw output of the tools
  - Metadata of the experimental campaign
  - Instantaneous values
  - Synthetic information

- Users can access information related to a wider set of configurations and scenarios reducing experimental cost
  - More than *1,000 hours* of experimental data already shared

CloudSurf allows to improve the knowledge about public-cloud networks

- *How network resources are allocated to customer VMs?*
- *Which provider performs better?*
- *How performance varies across different geographical regions?*
- *Which cloud region serves better a given customer?*
- *…*

Performance figures advertised by providers are coarse-grained

Amazon intra-datacenter throughput

- TCP throughput
- 24-hour-long campaign
- repeated 5-minute long experiments

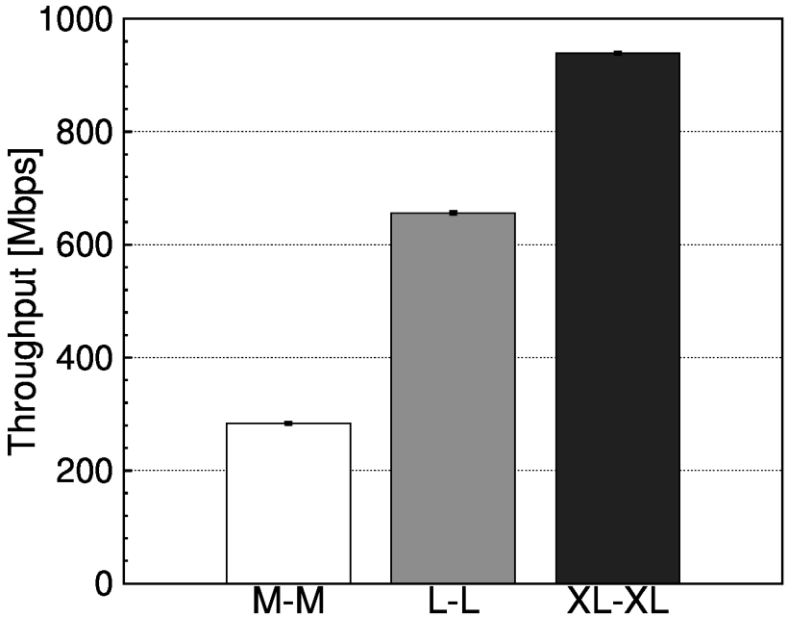

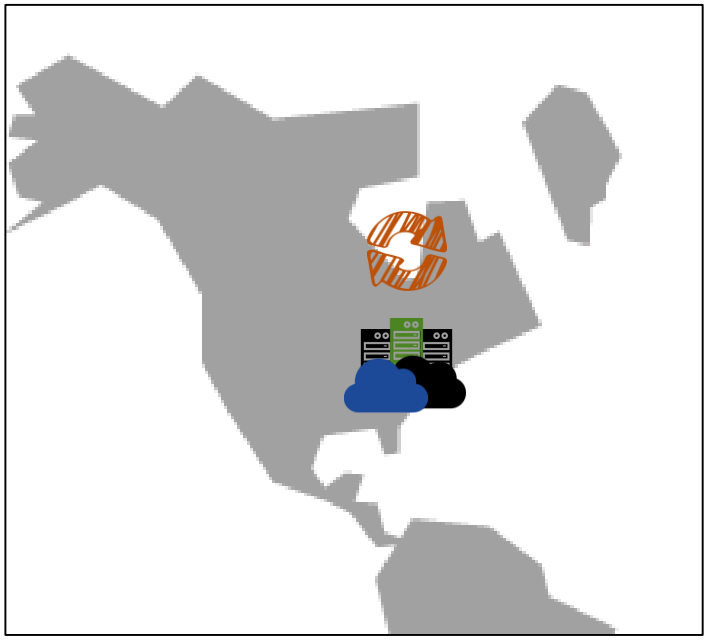

- Performance in terms of throughput varies with VM size (and cost)
- Negligible variability observed over time
- M-sized and L-sized VMs are both advertised to have *Moderate* network performance

Performance figures advertised by providers are coarse-grained

Amazon intra-datacenter throughput

- TCP throughput
- 24-hour-long campaign
- repeated 5-minute long experiments

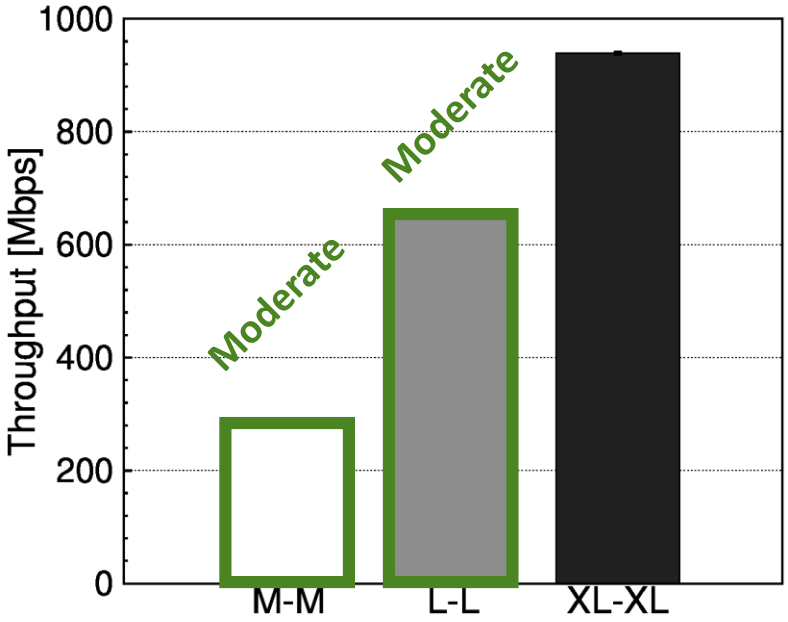

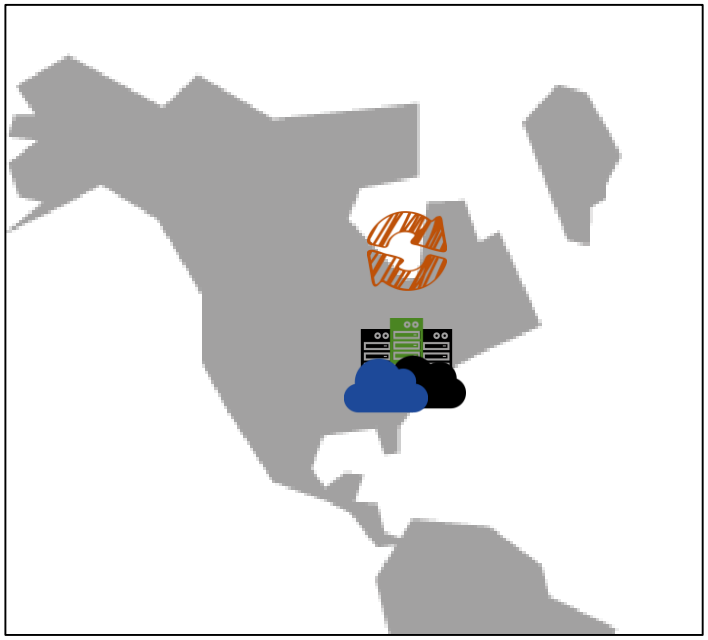

- Performance in terms of throughput varies with VM size (and cost)
- Negligible variability observed over time
- M-sized and L-sized VMs are both advertised to have *Moderate network performance*

Performance may heavily vary across different providers

Inter-datacenter TCP throughput

- Four different Amazon and Azure regions
  - Ireland
  - North Virginia
  - Singapore
  - Sao Paulo
- 300-hour-long campaign

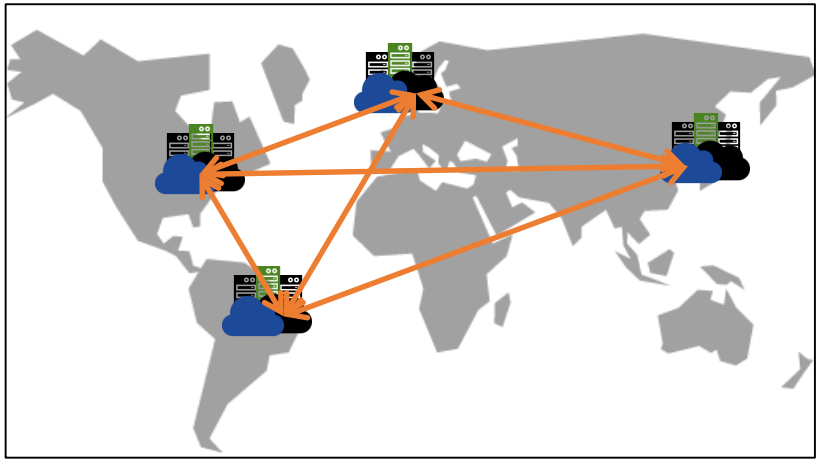

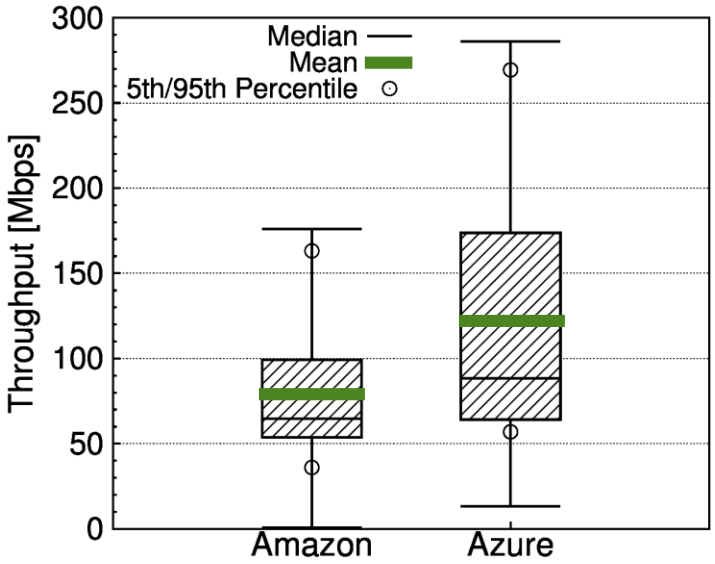

- VM size is not influential
- Azure inter-datacenter network performs better in terms of throughput
  - +56%, on average

Some cloud regions are connected to users better than others

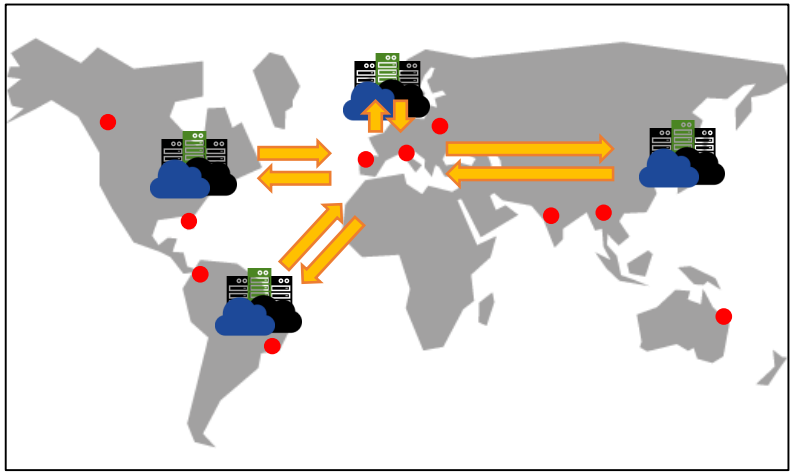

Cloud-to-user download performance

- Download 100MiB-content stored in 4 different cloud sites
- Vantage points spread worldwide (PlanetLab nodes)

- US and EU regions performs 45.5% better than SA and AP regions, on average.
- Counterintuitively, AP and SA are also associated to higher network-transfer costs with respect to EU and US.

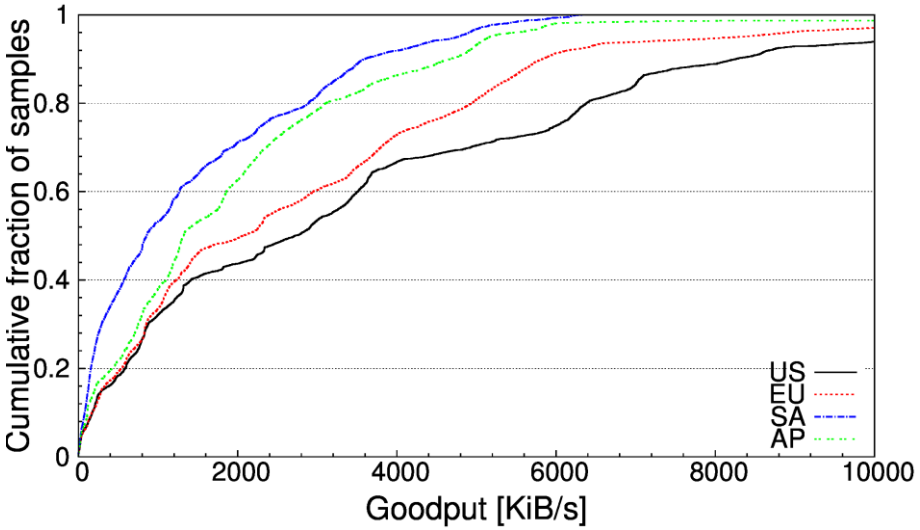

Conclusion

- *Cloudsurf* implements non-cooperative monitoring approaches to monitor public-cloud networks
  - Comprehensive
  - Easy to use
  - Freely available at http://traffic.comics.unina.it/cloudsurf/

- Main features
  - Customizable experiments
  - Two providers supported (as for now)
  - Intra-datacenter, Inter-datacenter, cloud-to-user network performance
  - Community repository

- *CloudSurf* helps to improve the knowledge about public-cloud network performance

Questions?

valerio.persico@unina.it

http://wpage.unina.it/valerio.persico

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

DIETI. UNIVERSITÀ DEGLI STUDI DI NAPOLI FEDERICO II

NM2 NetworkingMadeEasy