



Evaluating the Network Performance of ExoGENI Cloud Computing System



System and Networking Engineering

Andreas Karakannas

Anastasios Poulidis

Cloud Computing

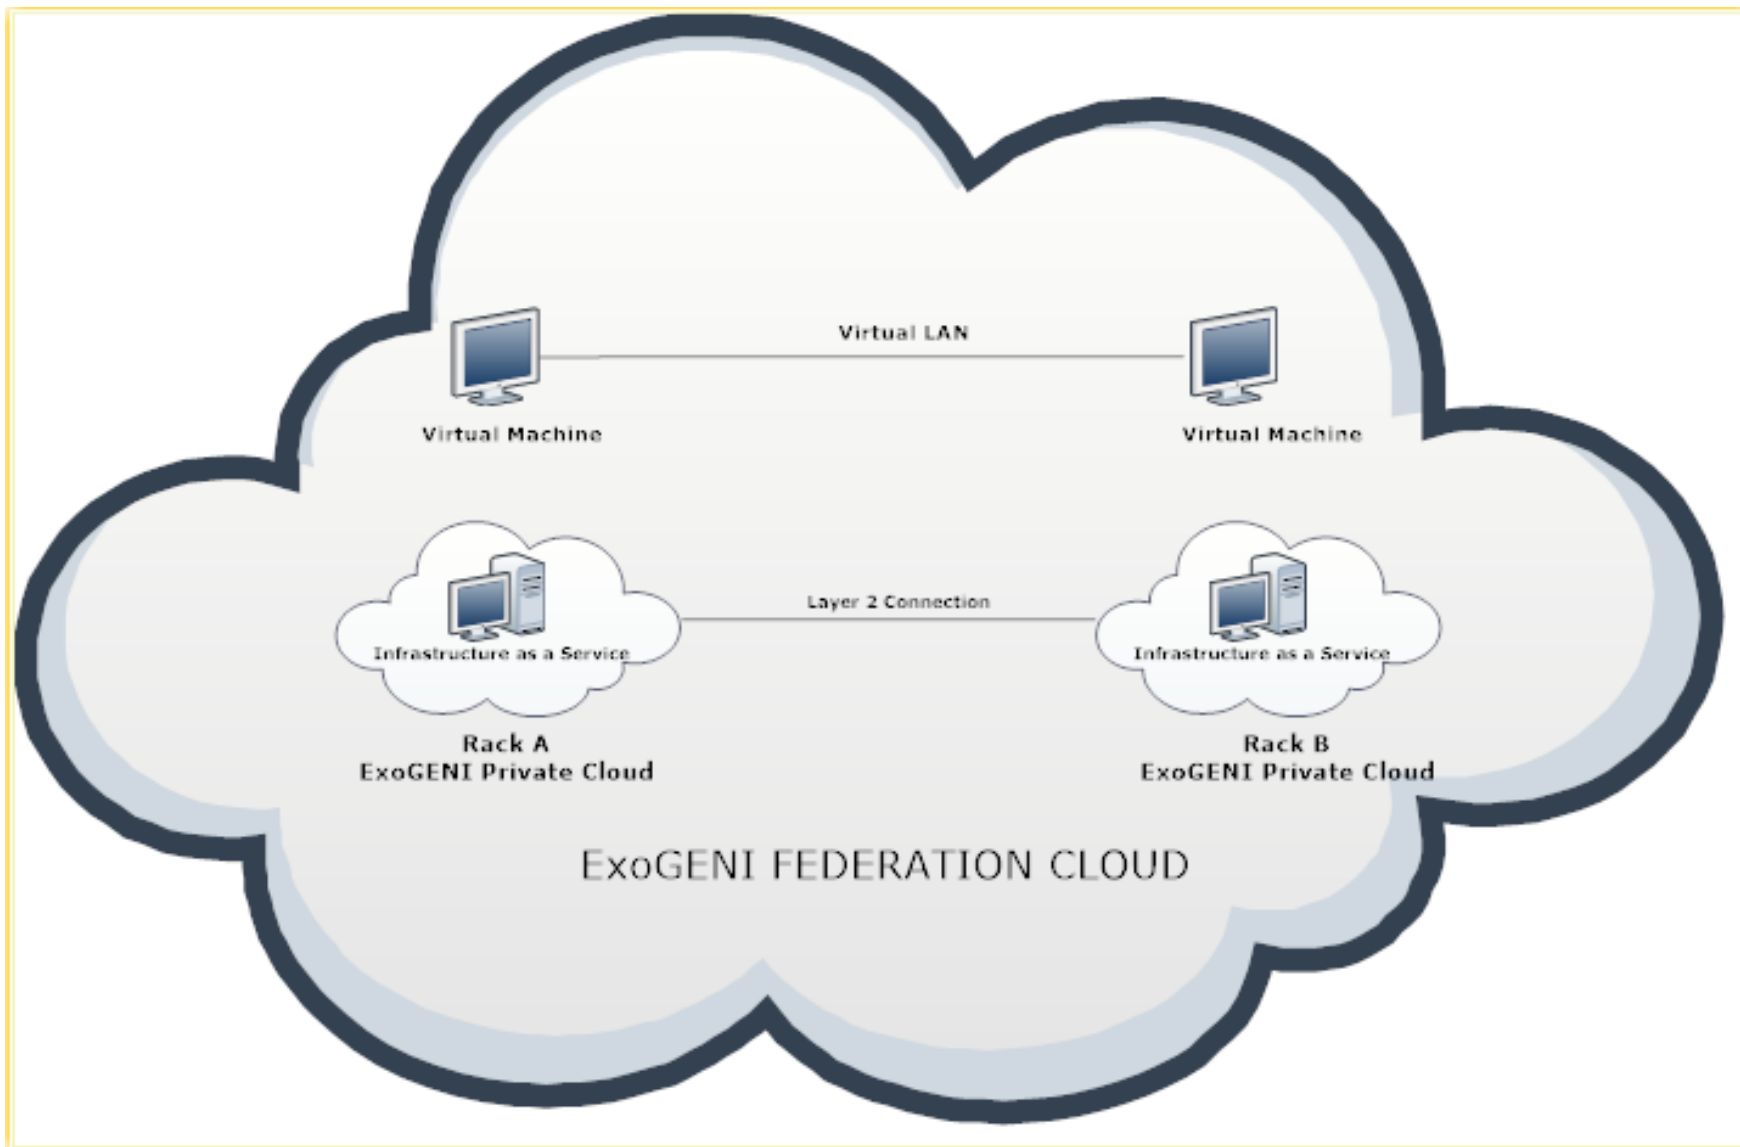
- ▶ Fundamental Technology
 - Virtualization
- ▶ Infrastructure as a Service
 - The user can create his own virtual network by combining virtual computers, storage, network devices and other computing resources from the Cloud.
- ▶ The User Problem
 - The user has no knowledge about the physical Infrastructure of his virtual network

ExoGENI Cloud Computing System

- ▶ Federated Cloud Computing System
 - Offers IaaS
 - Designed to support Research and Innovation in Networking
- ▶ Mostly used for Data-Intensive Applications
 - Network Performance Critical

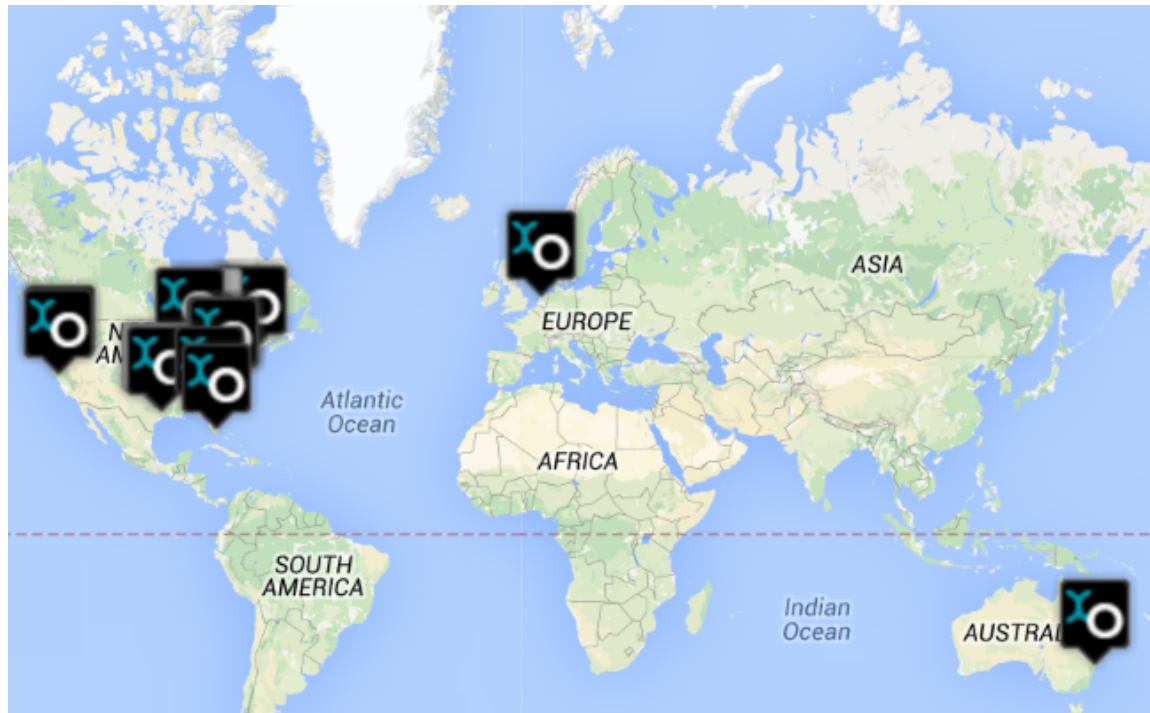
Research Questions

- What is the network performance on ExoGENI and how suitable is for data-intensive applications?
- Is the network performance on ExoGENI reproducible when the virtual network topologies are reconstructed from scratch with the same attributes?



ExoGENI Cloud System Virtualization


Geographical Location of Private ExoGENI Clouds



➤ <http://www.exogeni.net/locations/>

Private Cloud	Location
RENCI	North Carolina, USA
UFL	Boston, USA
NICTA	Sydney, Australia
UH	Huston, USA
FUI	Florida, USA
UFL	Florida, USA
DU	North Carolina, USA
SL	Illinois, USA
UVA	Amsterdam, Netherlands
UDC	California, USA
OSF	California, USA

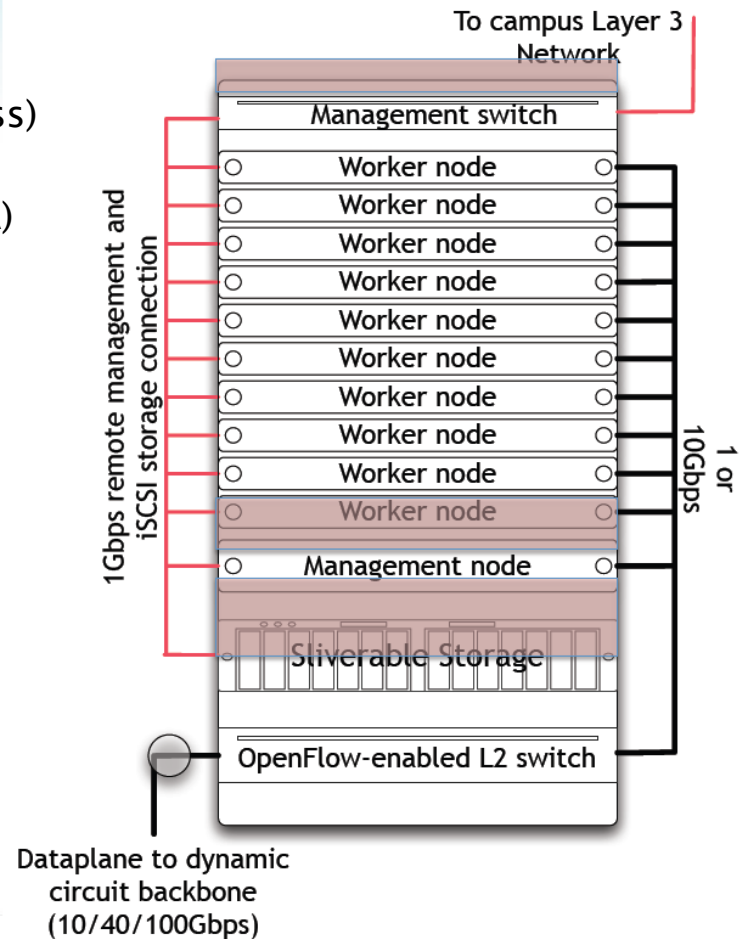
Physical Connection of ExoGENI Private Clouds through Circuit Providers



➤ <https://wiki.exogeni.net/doku.php?id=public:experimenters:topology>

ExoGENI Rack Architecture

- 11 X3650m4 Servers
 - 10 Worker Nodes (User Access)
 - 1 Management Node (Management Access)
- 1 iSCSI Storage (OS images, Measurement DATA)
- 1 /10G Ethernet Infrastructure (Machines Interconnection)
- 1 8052 1/10G management switch (Provisioning and Managing the Rack)
- 1 8264 10/40/100G OpenFlow-enabled Dataplane Switch (Interconnection with a circuit provider)



➤ <http://groups.geni.net/geni/attachment/wiki/GEC12GENIDeploymentUpdates/GEC12-ExoGENI-Racks-campuses.pdf?format=raw>

Control framework

▶ ORCA

- Provision resources by using leases
- Uses OpenStack


▶ Provisioning Resources Problems

- Not available resources
- Failing nodes
- Technical problems
- 5 maintenances

Creating Network Topologies on ExoGENI

- ▶ FLUKES: User tool for creating network topologies on ExoGENI through a GUI.
 - NDL-OWL
 - Functionalities
 - Create
 - Modify
 - Inspect

- Node
- Node Group
- Broadcast Link
- Storage
- StitchPort



Client Images

Auto IP

Reservation Details

Submit Request

test_topology

Node Properties

Node Node0 properties

Name:

Select node type:

- Site default
- XO Large
- Euca c1.medium

Select image:

- None
- OpenDaylight
- Debian 6 (squeeze) v1.0.9

Select domain:


- System select
- BBN/GPO (Boston, MA USA) XO Rack
- FIU (Miami, FL USA) XO Rack

Select dependencies:

- No dependencies
- Node1

Link0 IP Address: . . . /

Additional properties:



Client Images

Auto IP

Reservation Details

Submit Request

test_topology

Node Properties

Node Node0 properties

Name:

Select node type: →
XO Large
Euca c1.medium


Select image:
OpenDaylight
Debian 6 (squeeze) v1.0.9

Select domain:
BBN/GPO (Boston, MA USA) XO Rack
FIU (Miami, FL USA) XO Rack

Select dependencies:
Node1

Link0 IP Address: /

Additional properties:



Client Images

Auto IP

Reservation Details

Submit Request

test_topology

Node Properties

Node Node0 properties

Name:

Select node type:


Select image:

Select domain:

Select dependencies:

Link0 IP Address:

Additional properties:



Client Images

Auto IP

Reservation Details

Submit Request

test_topology

Node Properties

Node Node0 properties


Name:

Select node type:

- Site default
- XO Large
- Euca c1.medium

Select image:

- None
- OpenDaylight
- Debian 6 (squeeze) v1.0.9

Select domain: 


- System select
- BBN/GPO (Boston, MA USA) XO Rack
- FIU (Miami, FL USA) XO Rack

Select dependencies:

- No dependencies
- Node1

Link0 IP Address: . . . /

Additional properties:



Client Images

Auto IP

Reservation Details

Submit Request

test_topology

Node Properties

Node Node0 properties

Name:

Select node type:

- Site default
- XO Large
- Euca c1.medium

Select image:

- None
- OpenDaylight
- Debian 6 (squeeze) v1.0.9

Select domain:


- System select
- BBN/GPO (Boston, MA USA) XO Rack
- FIU (Miami, FL USA) XO Rack

Select dependencies:

- No dependencies
- Node1

Link0 IP Address: . . . /

Additional properties:



Client Images

Auto IP

Reservation Details

Submit Request

test_topology

Node Properties

Node Node0 properties

Name:

Select node type:

- Site default
- XO Large
- Euca c1.medium

Select image:

- None
- OpenDaylight
- Debian 6 (squeeze) v1.0.9

Select domain:


- System select
- BBN/GPO (Boston, MA USA) XO Rack
- FIU (Miami, FL USA) XO Rack

Select dependencies:

- No dependencies
- Node1

Link0 IP Address: . . . /

Additional properties:



Client Images

Auto IP

Reservation Details

Submit Request

test_topology

Node Properties

Node Node0 properties

Name:

Select node type:

- Site default
- XO Large
- Euca c1.medium

Select image:

- None
- OpenDaylight
- Debian 6 (squeeze) v1.0.9


Select domain:


- System select
- BBN/GPO (Boston, MA USA) XO Rack
- FIU (Miami, FL USA) XO Rack

Select dependencies:

- No dependencies
- Node1

Link0 IP Address: . . . /

Additional properties: 



Client Images

Auto IP

Reservation Details

Submit Request

test_topology

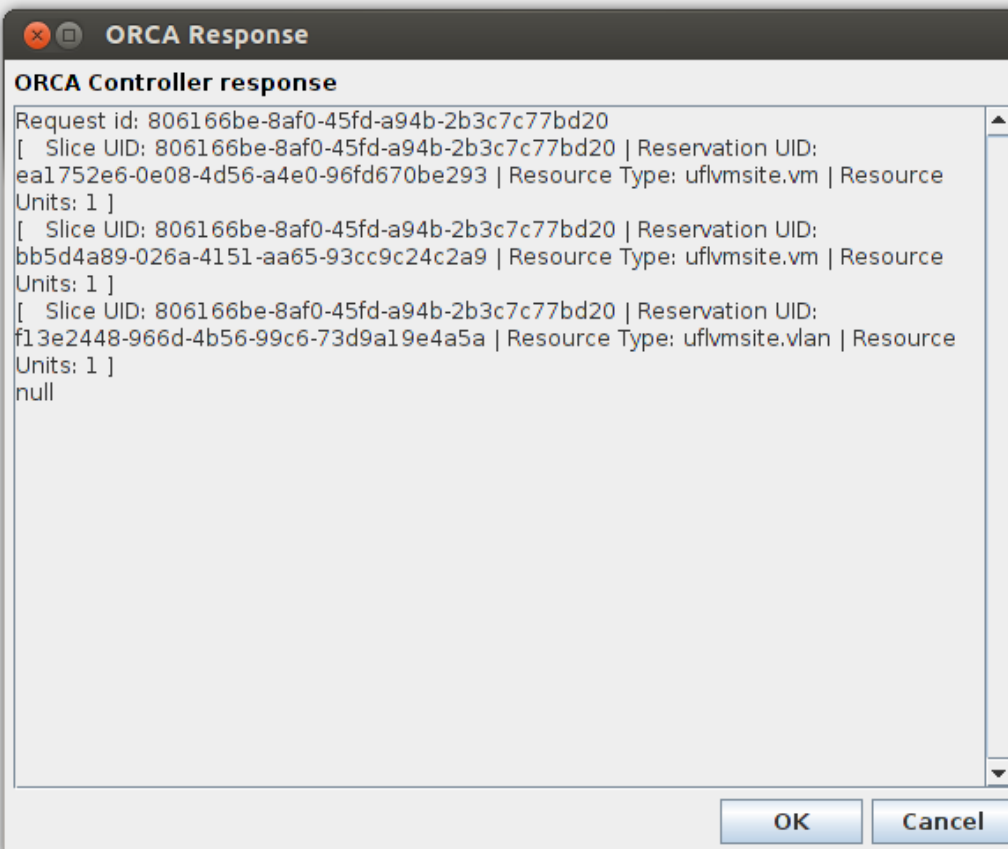


Node0

Link0



Node1






ORCA FLUKES - The ORCA Network Editor

FileOrca ControllerOutput FormatGraph LayoutHelp

Resource ViewRequest ViewManifest View

My SlicesQuery for Manifesttest_topologyView Raw ResponseExtend ReservationCommit Modify ActionsClear Modify A





View current resource states.

Resource states (start: Sun Feb 02 16:36:39 PST 2014 end: Mon Feb 03 16:36:39 PST 2014):

Resource Name	Resource State
Node0/0	Ticketed
Node1/0	Ticketed
Link0	Active

OK Cancel



View current resource states.

Resource states (start: Sun Feb 02 16:27:17 PST 2014 end: Mon Feb 03 16:27:17 PST 2014):

Resource Name	Resource State
Node0/0	Active
Node1/0	Active
Link0	Active

OK Cancel

Resource View

Request View

Manifest View

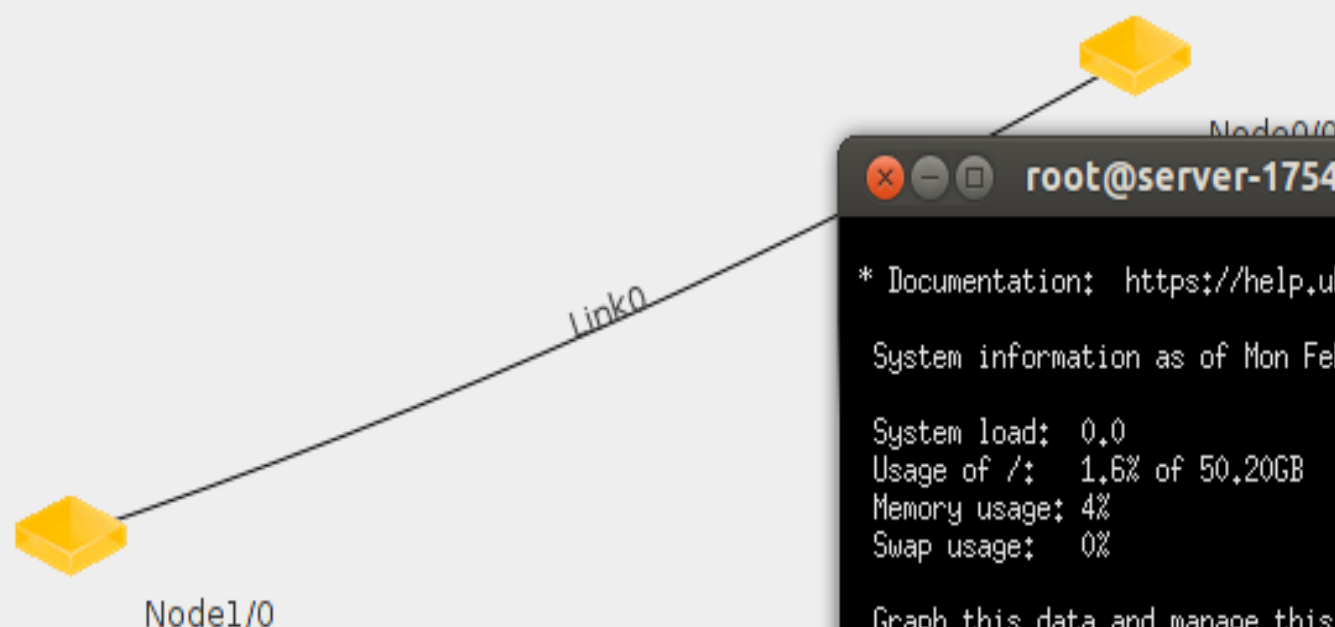
My Slices

Query for Manifest

test_topology

View Raw Response

Extend Reservation



root@server-1754: ~

* Documentation: <https://help.ubuntu.com/>

System information as of Mon Feb 3 12:09:42 UTC 2014

System load:	0.0	Processes:	67
Usage of /:	1.6% of 50.20GB	Users logged in:	0
Memory usage:	4%	IP address for eth0:	10.103
Swap usage:	0%		

Graph this data and manage this system at <https://landscape>

88 packages can be updated.

42 updates are security updates.

Get cloud support with Ubuntu Advantage Cloud Guest

<http://www.ubuntu.com/business/services/cloud>

root@server-1754:~#

Experimental Scenarios


Cloud	Scenarios	Communication	Distance	Virtual Links Bandwidth
Inter-Racks	Experiment 1	Point to Point	Short – Long	10Mbps
	Experiment 2	Point to Multiple Points	–	10Mbps
Intra-Racks	Experiment 3	Point to Point	Same Server – Different Server	100Mbps
	Experiment 4	Point to Multiple Points	–	100Mbps
Both	Reproducibility	All	All	Both

Network Performance Metrics

Experimental Scenario 1 – 4

Metric	Measurements	Measurement Interval (Minutes)	Measurement Time(Second)
TCP Throughput	100	10	60
UDP Throughput	100	10	60
Packet Loss	100	10	60
RTT	100	10	60


Experimental Scenario 1



Experiment 1

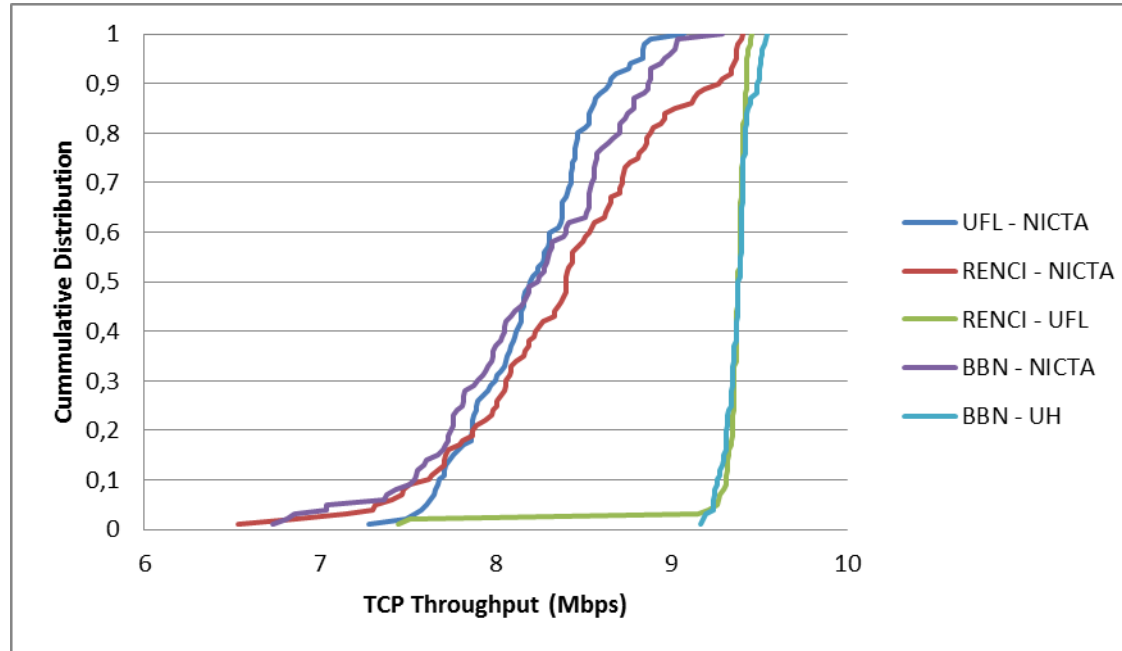
A/A	Rack A	Rack B	Distance
1	RENCI, USA	NICTA, AUSTRALIA	Long
2	UFL, USA	NICTA, AUSTRALIA	Long
3	BBN, USA	NICTA, AUSTRALIA	Long
4	RENCI, USA	UFL, USA	Short
5	BBN, USA	UH, USA	Short

Results: Experiment 1



- 5-times bigger RTT on long distances
- Minor abnormalities

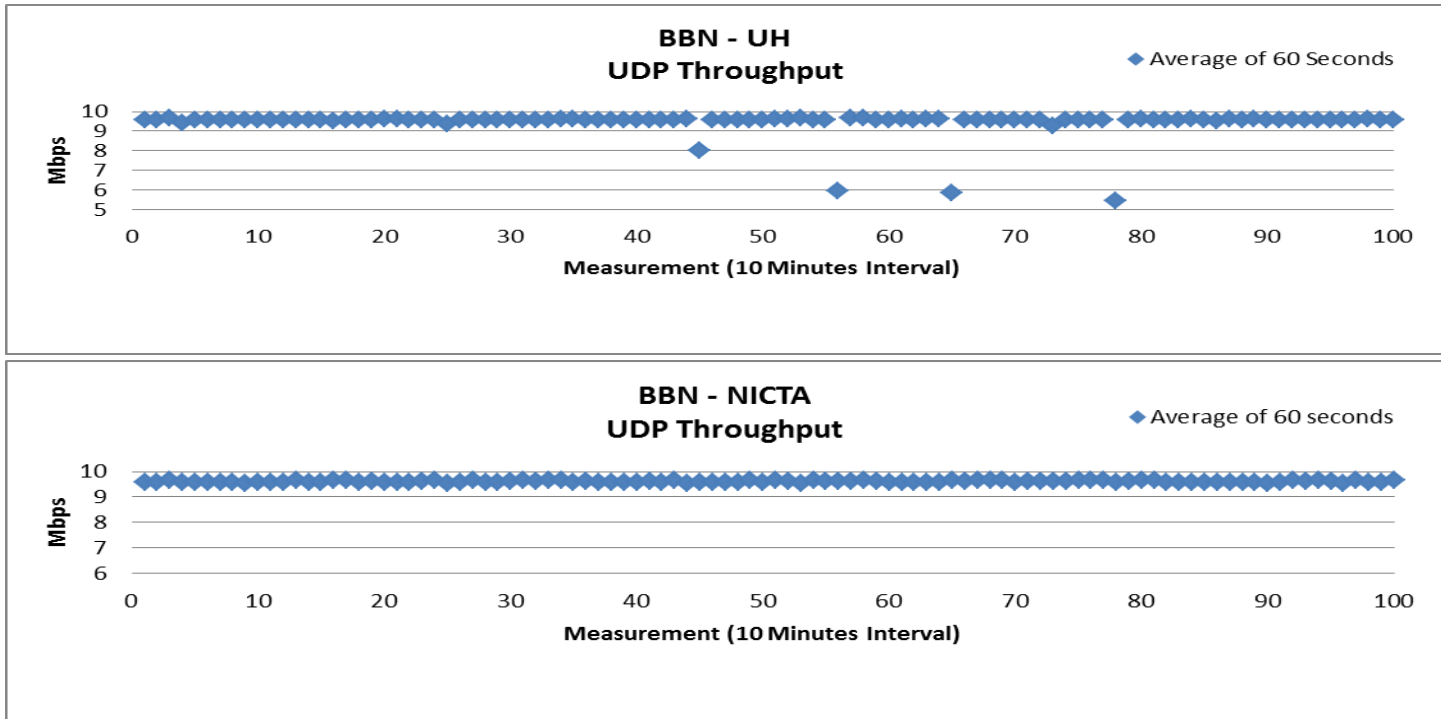
Results: Experiment 1



Long Distance Connections have lower average TCP Throughput because of higher RTT


Results: Experiment 1

UDP Throughput




UDP Throughput for short and long distance connections is approximately the same. (No ACK needed on UDP packets => No RTT for ACK)
4 cases of high packet loss rate(40%) for UDP short distance connection = >
Long Distance more Stable for UDP Connections.
Packet Loss : BBN - UH(5.4%), BBN - NICTA(4%)

Experimental Scenario 2




Results: Experiment 2



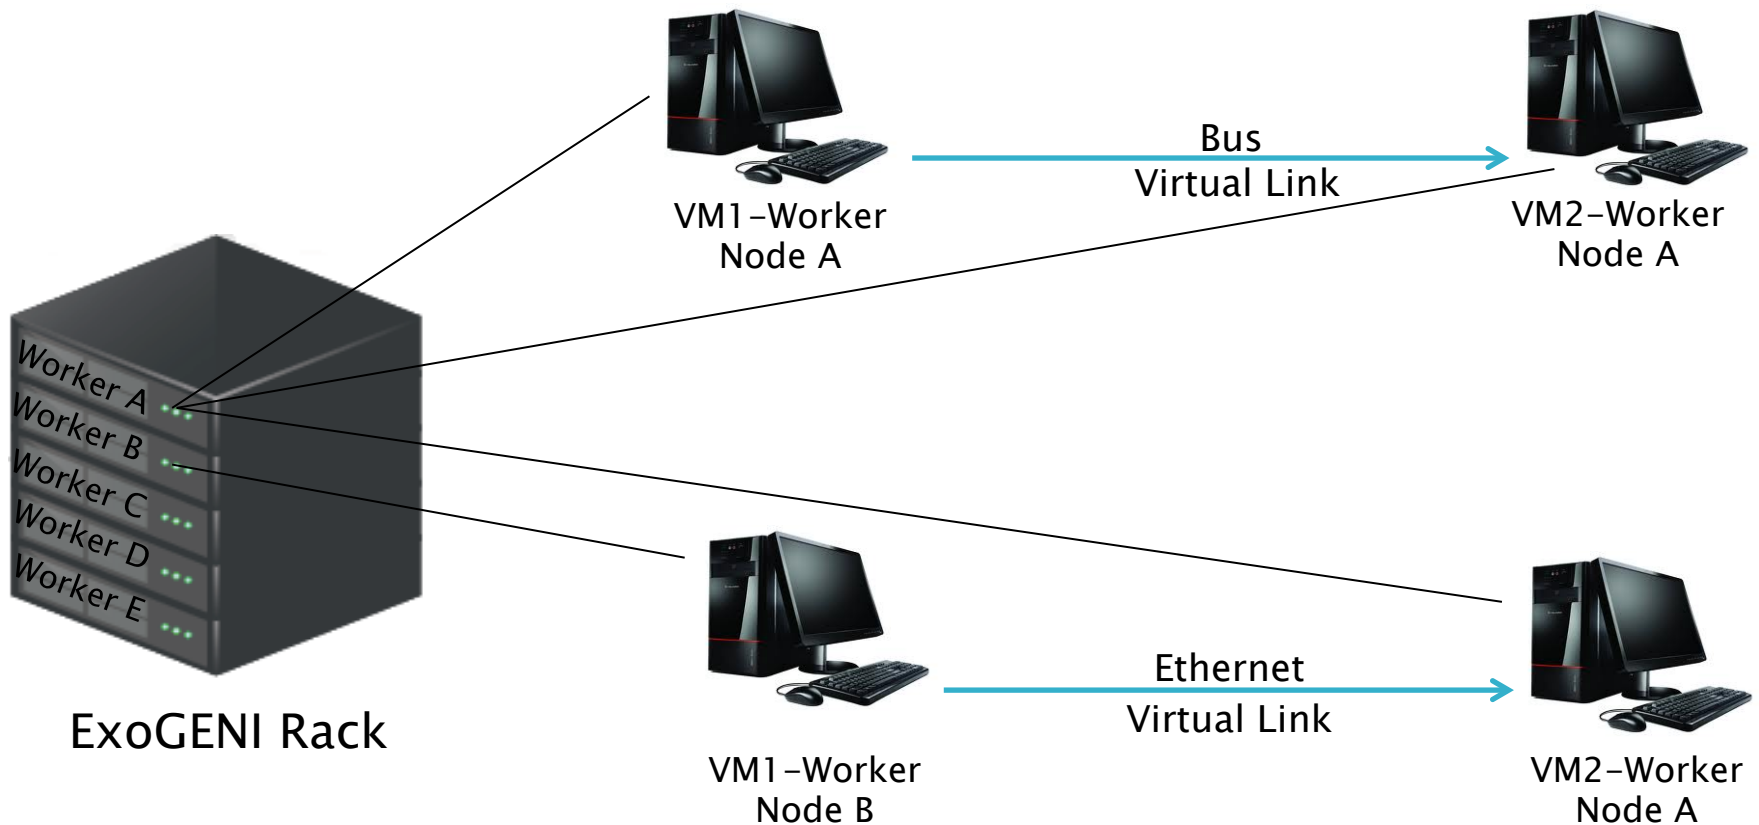
Abnormal Behavior of TCP Connections

Results: Experiment 2



- UDP Throughput implies no competition upon the physical Infrastructure
- Packet Loss(4%)

Experimental Scenario 3



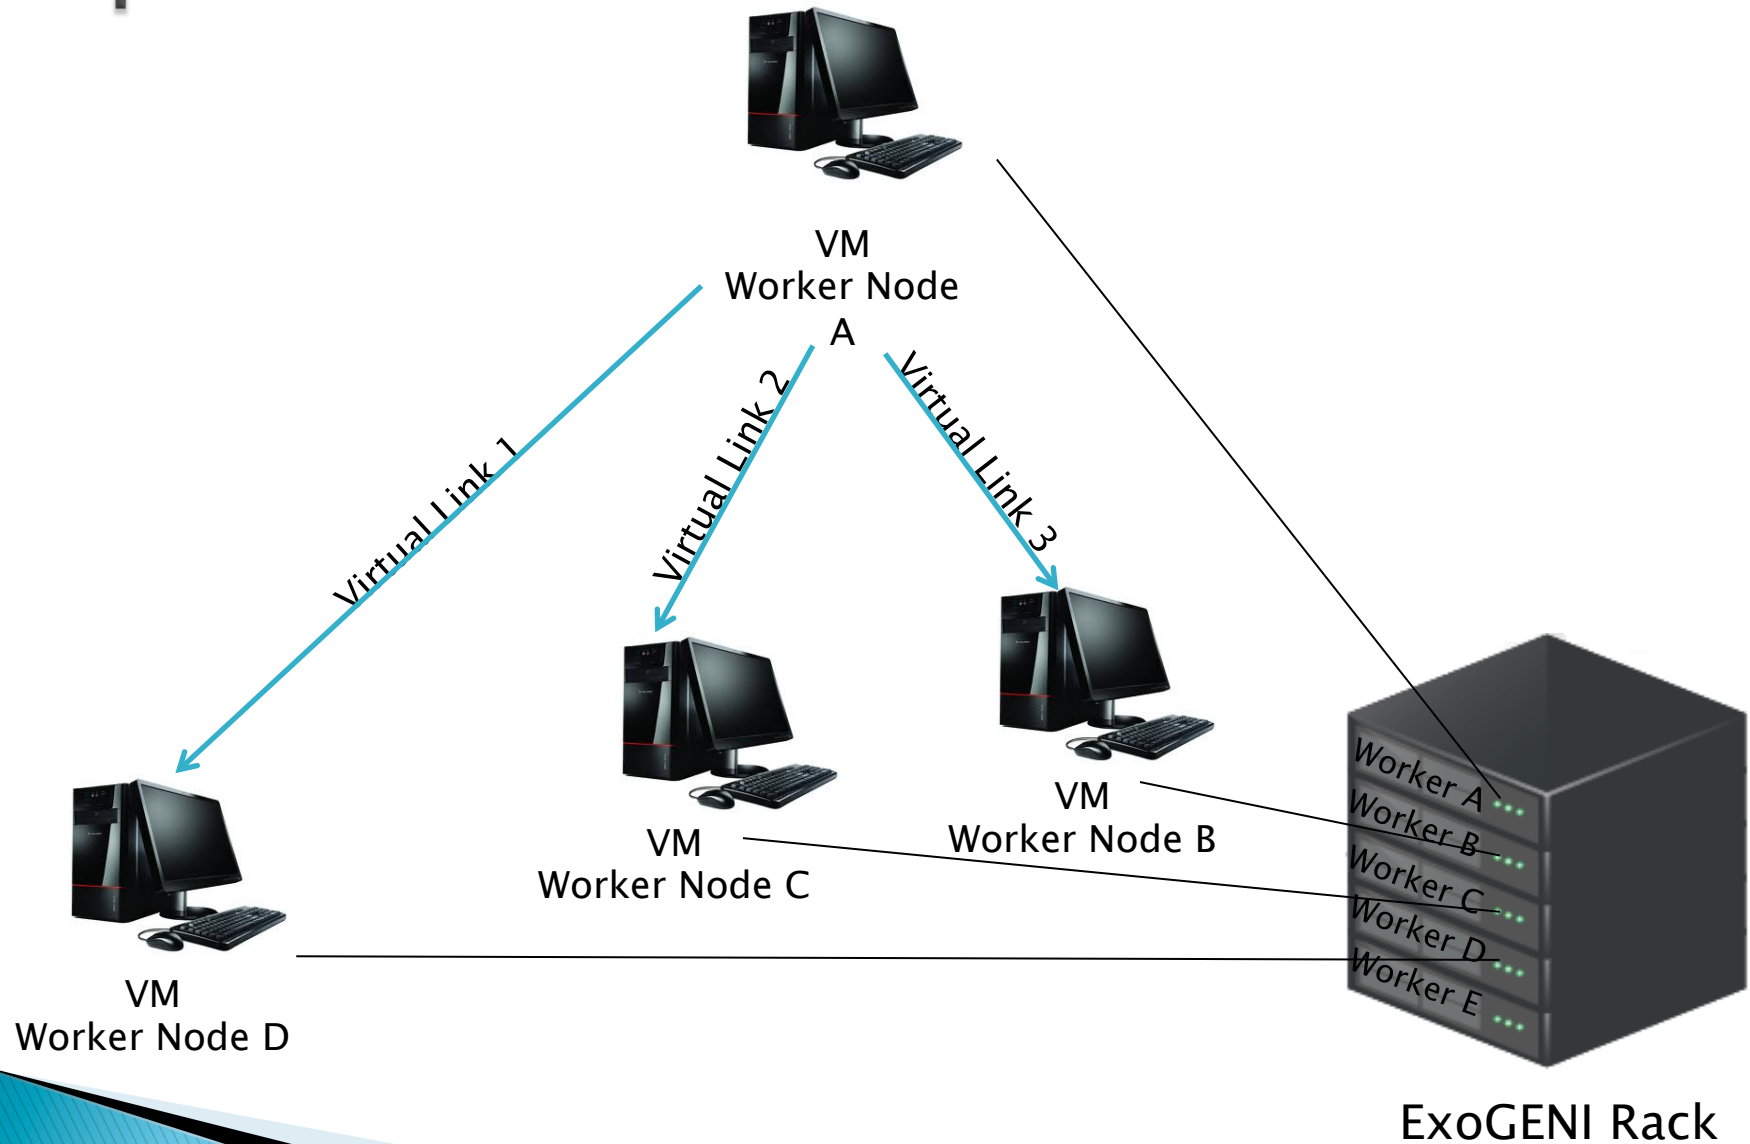
Results Experiment 3

Exo GENI Rack	TCP Throughput (Mbps)		UDP Throughput (Mbps)		Packet Loss (%)		RTT Milliseconds	
	VMs on Same Worker Node	VMs on Different Worker Node	VMs on Same Worker Node	VMs on Different Worker Node	VMs on Same Worker Node	VMs on Different Worker Node	VMs on Same Worker Node	VMs on Different Worker Node
UFL	100	99,7	95,9	95,9	4	4	0,337	0,832
UH	100	99,6	95,8	95,7	4	4	0,341	0,811

Network performance is the same independently on which Worker Node the Virtual Machines are located.

RTT for VMs on different worker node is ~2 times bigger than VMs on the same worker node

Experimental Scenario 4



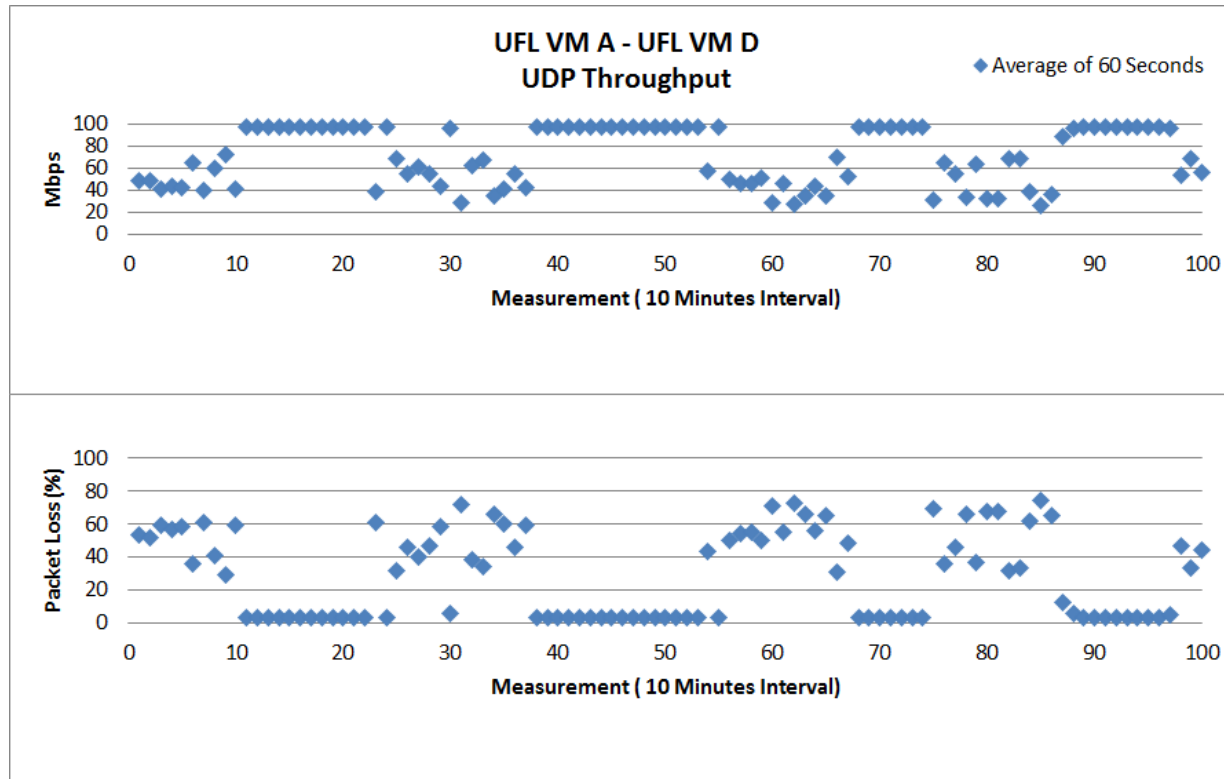
Results: Experiment 4

ExoGENI Rack	TCP Throughput (Mbps)			UDP Throughput (Mbps)		
	VM Worker Node B	VM Worker Node C	VM Worker Node D	VM Worker Node B	VM Worker Node C	VM Worker Node D
UFL	99,9	99,9	99,9	96,6	95	72
UH	99,9	99,9	99,9	95	94	77

TCP Throughput is the same for all connections for both RACKS
UDP Throughput has an abnormal behavior for the VM on Worker Node D on both Racks.

Results: Experiment 4

UDP low average UDP throughput is caused by a lot of packet loss in specific time intervals.



Reproducibility of Experiments

The topology was deleted and recreated 100 times with 5 minutes interval between each same experiment repetition.

The above measurements were taken for each repetition.

Metric	Measurements	Measurement Interval (Minutes)	Measurement Time(Second)
TCP Throughput	10	5	20
UDP Throughput	10	5	20
Packet Loss	10	5	20
RTT	10	5	20

Results: Reproducibility

	Reproducibility	Results
Experiment 1	Not Available Resources	–
Experiment 2	Not Available Resources	–
Experiment 3	Possible	Same as Initial Experiment
Experiment 4	Possible	Same as Initial Experiment

Conclusion

▶ Federated Cloud

- Short Distance end to end point communication – TCP and UDP stable.
- Long Distance end to end point communication – UDP stable, TCP unstable.
- One to Multiple communication – UDP stable, TCP unstable.
- Reproducibility of experiments – No Results.

▶ Private Cloud

- End to end point communication – TCP and UDP stable.
- One to Multiple communication – TCP stable, UDP unstable.
- Reproducibility of network performance – 100%.

Thank you!

Questions ?