

GENI Laboratory Exercises for a Cloud Computing course

Prasad Calyam, Ph.D.

Assistant Professor, Department of Computer Science

NSF Workshop on GENI in Education, October 26th 2013

What is Cloud Computing?



- **Classic Computing**

- ‘Buy and Own’
- Approx. 18-month cycle for hardware/software procurement
- Plan for peak needs of resources, could be expensive!
- Client software installation (e.g., MS Word), dedicated

- **Cloud Computing**

- ‘Subscribe and Use’ like any other “utility” such as water, gas, etc.
- Instantly use and pay \$ for only what you use
- Elastic resource reservations
- Web Apps (e.g., Google Docs), Mobility, Multi-tenancy

VIMAN Lab Cloud Computing Research

<http://people.cs.missouri.edu/~calyamp>



- Cloud Resource Allocation
 - Computer and network virtualization models, algorithms, tools
- Cloud Performance Monitoring
 - Software-defined measurements and performance diagnosis
- Cloud Testbeds for Apps, Marketplaces
 - e.g., Manufacturing/Healthcare/Education
- Cloud Security
 - Cyber attacks defense, Federated Policy Management

Cloud Computing Course Information

- **Pre-requisites:**
 - Object Oriented Programming
 - (Operating Systems-I), (Computer Networks-I), (Cyber Security-I)
- **Students:**
 - 23 Graduate, 7 Undergraduate, 5 visitors! ☺
- **Textbook:**
 - Distributed and Cloud Computing - by K. Hwang, J. Dongarra, and G. Fox; First Edition [ISBN: 9780123858801] (required)
- **Reference books:**
 - Programming Amazon EC2 - by J. Vliet and F. Paganelli [Online copy available] (optional)
 - Cloud Computing: Principles and Paradigms - by R. Buyya, J. Broberg, A. Goscinski (optional)
- **Teaching Assistant:**
 - Sripriya Seetharam – attended GENI Summer Camp, GECs, Thesis Research that uses GENI infrastructure
 - Works closely with GPO (Thanks! Niky, Sarah and Vic)
 - Monitors and answers student issues in Blackboard Discussion Forum

Blackboard Discussion Forum

Discussion Board

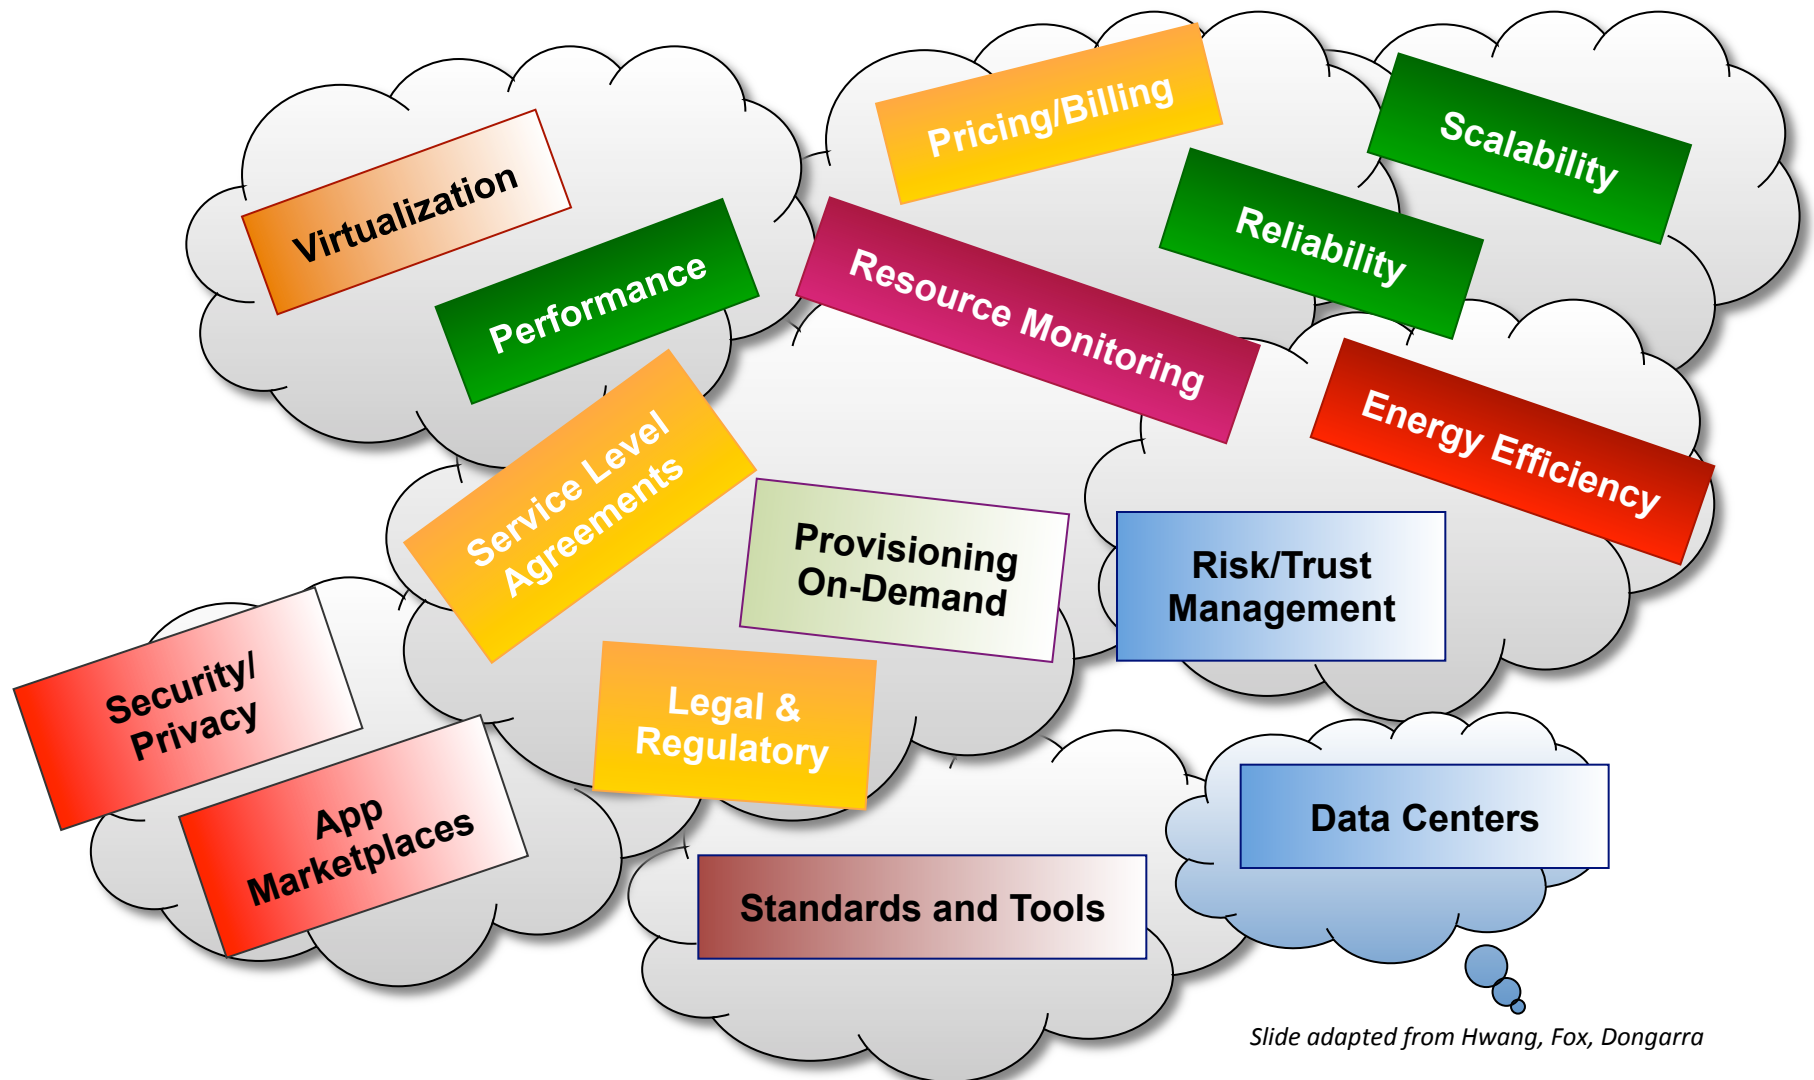
Forums are made up of individual discussion threads that can be organized around a particular subject. Create Forums to organize discussions. [More Help](#)

Create Forum

[Delete](#)

<input type="checkbox"/> Forum	Description	Total Posts
<input type="checkbox"/> AWS Lab-1 Discussion		7
<input type="checkbox"/> GENI Lab-1 Discussion		15
<input type="checkbox"/> Share Cool Tools for Labs/Projects	Please share any cool tools for AWS or GENI that you feel other students can use to more effectively do their labs and the final project.	2
<input type="checkbox"/> AWS Lab-2 Discussion		5
<input type="checkbox"/> GENI Lab-2 Discussion		8
<input type="checkbox"/> Register for MU Data Center Tour	<p>Hi All,</p> <p>As mentioned in class, I am arranging a MU Data Center Tour for you to see a "real data center"! The tour will be on November 6th at 1:30p. I will send data center address and meet-up details later.</p> <p>Please put a "Yes, I will attend the tour" note into this forum if you can come to the tour. This will help me in getting a count of students, and who will come to the tour.</p> <p>Thanks</p> <p>-Prasad</p>	15

Cloud Computing Challenges?



Slide adapted from Hwang, Fox, Dongarra

Cloud Computing Course Objectives

- Be competent with principles and architectures that integrate computing theories and information technologies
- Be competent with distributed system models, computer virtualization concepts, network virtualization concepts
- Be competent with the design, programming and application of *distributed* and *cloud computing* systems

and...

- Be familiar with cyber-security and service-oriented concepts relating to use/design of cloud computing systems
- Be familiar with App customization of cloud computing infrastructures, and standards through hands-on experience

Course Lab Assignments

- 6 labs to develop technical background and skills for working with distributed system and related software environments
- 3 Amazon Web Services (AWS) labs - <http://aws.amazon.com>
 - Received \$100 usage credit per student through their academic program
 - Focus areas:
 - Instance setup with Cloud Watch alarm for billing alerts
 - Web App setup with EC2, S3 and Cloud Watch
 - Autoscaling to handle user load bursts in a Web App
- **3 GENI labs** - <http://groups.geni.net/geni/wiki>
 - “Calyam UMissouri Fall2013” Project in GENI portal (<http://portal.geni.net>)
 - **Focus areas:**
 - **Slice setup for “Hello GENI” experiment**
 - **Instrumentation & Measurement Web App setup with GENI Rack VMs and OnTimeMeasure software**
 - **Software-defined networking and user QoE performance adaptation under degrading network health conditions**

Final Course Project

- **Objective:**
 - To develop programming and performance evaluation skills
 - To reinforce the understanding of major concepts in the course
- **Approach:**
 - Solve a “real” problem in a GENI/AWS testbed
 - **Teams:** *Core Logic Team, Testbed Team, User Interface Team*
 - *Graduate Student team projects* require use of either C/C++ or Java, as well as cloud-platform specific tools and applications
 - *Undergraduate student participation is optional (extra credit)*

Course Topics Schedule

- **Week 1 – 4: Introduction**

- Distributed system models and enabling technologies
 - Scalable, utility computing evolution and concepts
 - Networked systems related technologies
 - Performance, security, and energy efficiency issues
- AWS overview discussion
- **GENI overview discussion**
- AWS Lab 1
- **GENI Lab 1**

- **Week 5 – 6: Virtualization**

- Virtual machines and virtualization of clusters, networks and data centers
 - Levels of virtualization implementation
 - Structures/tools and mechanisms for resource management
- AWS topics discussion
- AWS Lab 2

Course Topics Schedule (2)

- **Week 7 – 9: Architectures**

- Cloud platform architecture with security over virtualized data centers
 - Public, private and hybrid clouds ecosystem
 - Public clouds and service offerings such as IaaS, PaaS, SaaS, DaaS
 - Inter-cloud resource management
 - Cloud security and trust management
- **GENI topics discussion**
- **GENI Lab 2**

- **Week 10 – 12: Services**

- Service-oriented architectures for distributed computing
 - Web services and message-oriented middleware
 - Discovery, registries, metadata, publish-subscribe
- AWS topics discussion
- **GENI topics discussion**
- AWS Lab 3
- **GENI Lab 3**

Course Topics Schedule (3)

- **Week 13 – 16: Applications**
 - Cloud programming and software environments
 - Cloud application requirements and constraints
 - MapReduce, Hadoop library, Pig Latin
 - Programming support, approaches on real cloud platforms
 - Teams formation & 2-Page Project Plan
 - In-Class “Circle Time” for Project Teams
 - Project Team Presentations/Demos

Final Course Projects

- **Platform-as-a-Service for Hybrid Clouds**
 - Based on requirements from “ElderCare-as-a-Service” testbed
- **Cloud Performance Engineering**
 - Based on requirements from “Virtual Classroom Lab” testbed
- **Software-defined Networking for Multi-tenancy**
 - Based on requirements from a MU-OSU “Science DMZ” testbed

GENI Lab-1 Exercise

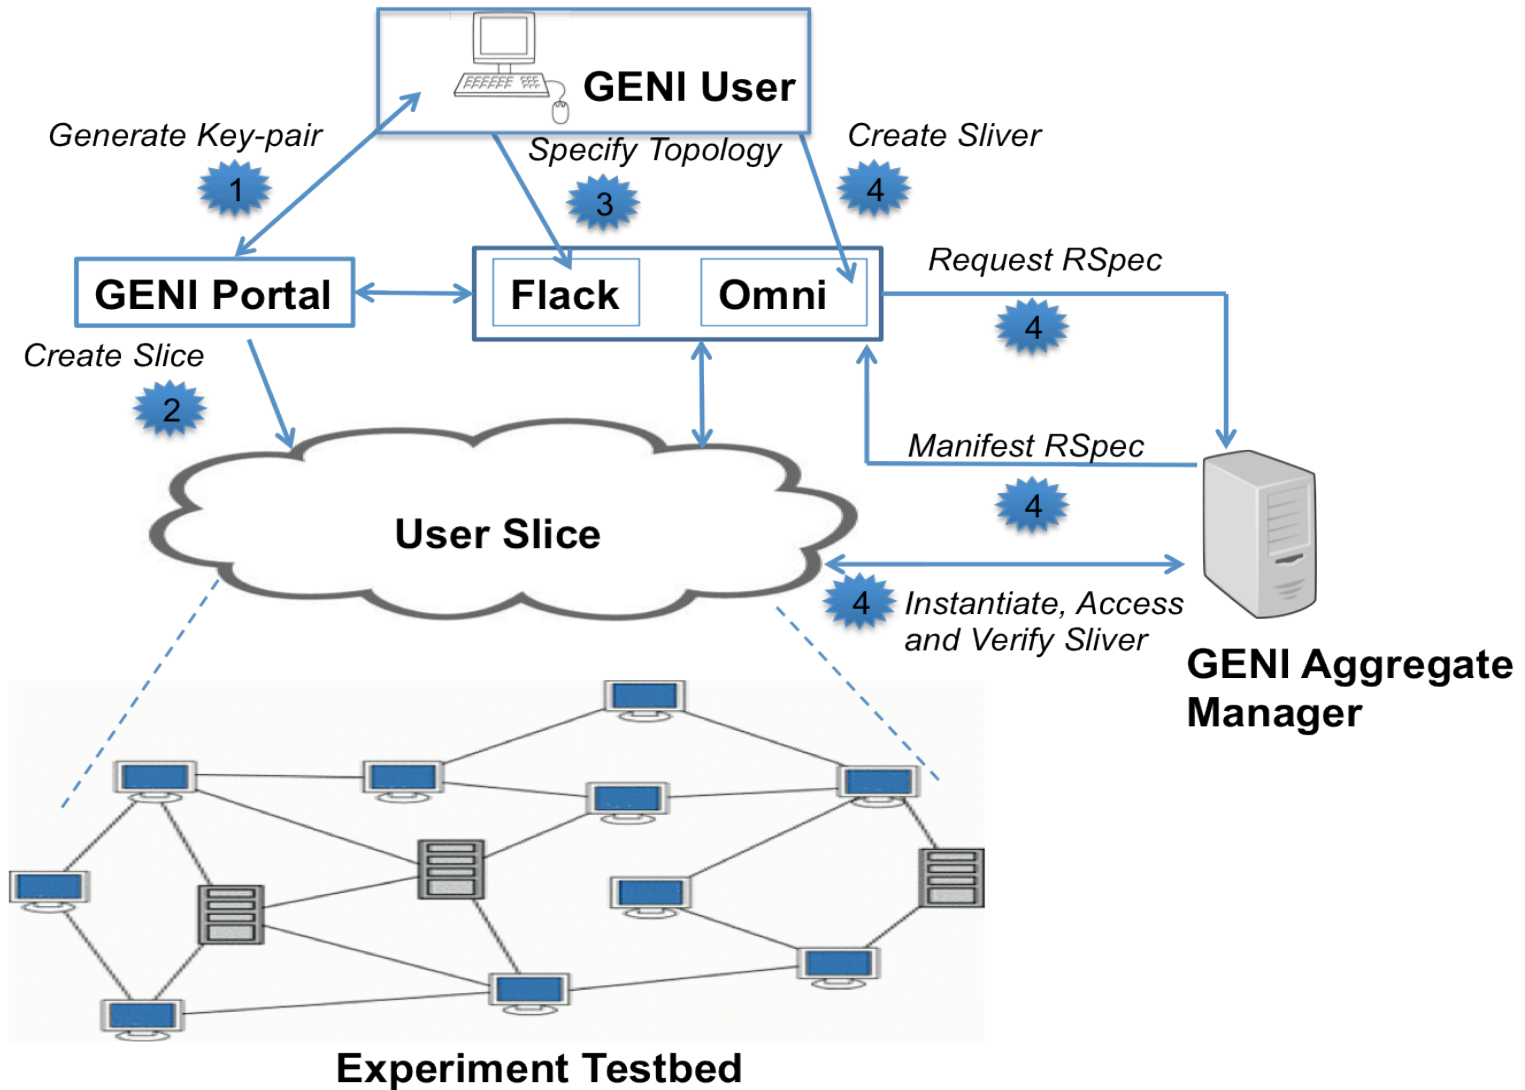


- **Lab # 1 – GENI Account Setup and Services Overview**
- **Purpose of the Lab**
 - Setup your GENI account, and use it to reserve an experiment environment a.k.a. “slice” that features 2 compute nodes communicating over a Layer 2 connection
 - Run a “Hello GENI” experiment with GUI (Flack) and Command-line (Omni) tools in your reserved slice by automatic installation and execution of software to perform basic bandwidth measurement tests between the nodes

References for GENI Lab-1

- Chapter 1, Distributed and Cloud Computing, Hwang, Fox & Dongarra
- GENI Terminology/Architecture:
<http://groups.geni.net/geni/wiki/GeniNewcomersWelcome>
- GENI Tutorial Paper:
<http://www.sigcomm.org/sites/default/files/ccr/papers/2012/January/2096149-2096161.pdf>
- GENI Omni Tool: <http://trac.gpolab.bbn.com/gcf/wiki/Omni>
- GENI Flack Tool: <http://www.protogeni.net/ProtoGeni/wiki/FlackManual>
- GENI Understanding Aggregate Manager exercise:
<http://groups.geni.net/geni/wiki/GENIEducation/SampleAssignments/UnderstandAMAPI/Procedure>


GENI Experiment Workflow



GENI Lab-1 Q&A

1. Provide a screenshot of the reserved slice resources in your GENI portal account
2. What would be the added capabilities and benefits in performing an experiment on GENI Future Internet infrastructure versus the commercial Internet?
3. Define: (i) slice, (ii) sliver, (iii) aggregate manager and (iv) Rspec
4. How does Federated Identity and Access Management in the GENI portal allow you to use your Pawprint and password to login? What are the benefits of this approach?
5. Explain the role of “Experiment Control Tools” such as Omni, Gush and Flack in GENI. What are the advantages of using the Omni command line tool in comparison to using the Flack GUI tool?
6. What is the significance of generating the SSH and SSL keys in your .ssh folder?

GENI Lab-2 Exercise

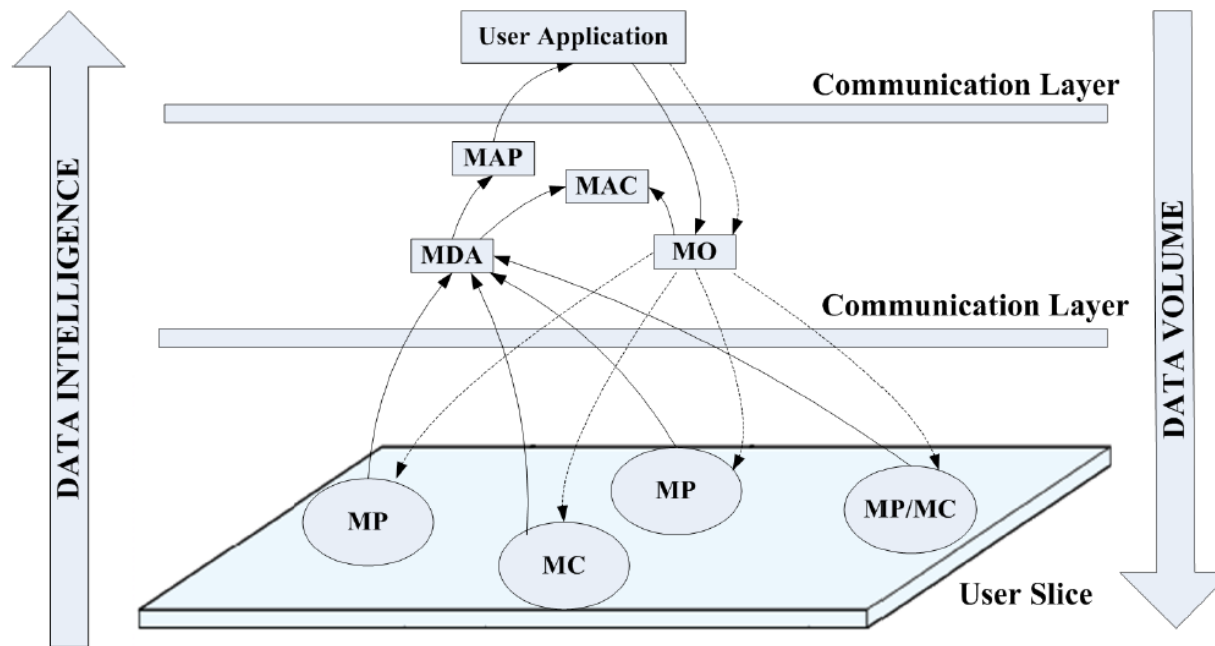


- **Lab # 2 – Instrumentation and Measurement of GENI Slice**
- **Purpose of the Lab**
 - Install and configure a distributed measurement Web App viz., “OnTimeMeasure” within a slice
 - Configuration involves ‘Node Beacons’ at measurement points, and a ‘Root Beacon’ that acts as a central intelligence module
 - Schedule measurements from the Root Beacon and query/visualize performance between two compute nodes that host Node Beacons
 - Performance data collection comprises of network health metrics such as round trip delay, loss, jitter and throughput

References for GENI Lab-2

- Chapter 3 Distributed and Cloud Computing, Hwang, Fox & Dongarra
- OnTimeMeasure Tutorial: <http://groups.geni.net/geni/wiki/OnTimeMeasure>
- OnTimeControl: <http://groups.geni.net/geni/wiki/OTM-Control>
- OnTimeMeasure Custom Metric Integration:
<http://groups.geni.net/geni/wiki/OTM-CustomMetric>
- OnTimeMeasure demo video on Graphite:
<http://ontime.oar.net/demo/dashboard.htm>
- Miscellaneous OnTimeMeasure demo videos: <http://ontime.oar.net/demo>

GENI I&M Overview



Legend:

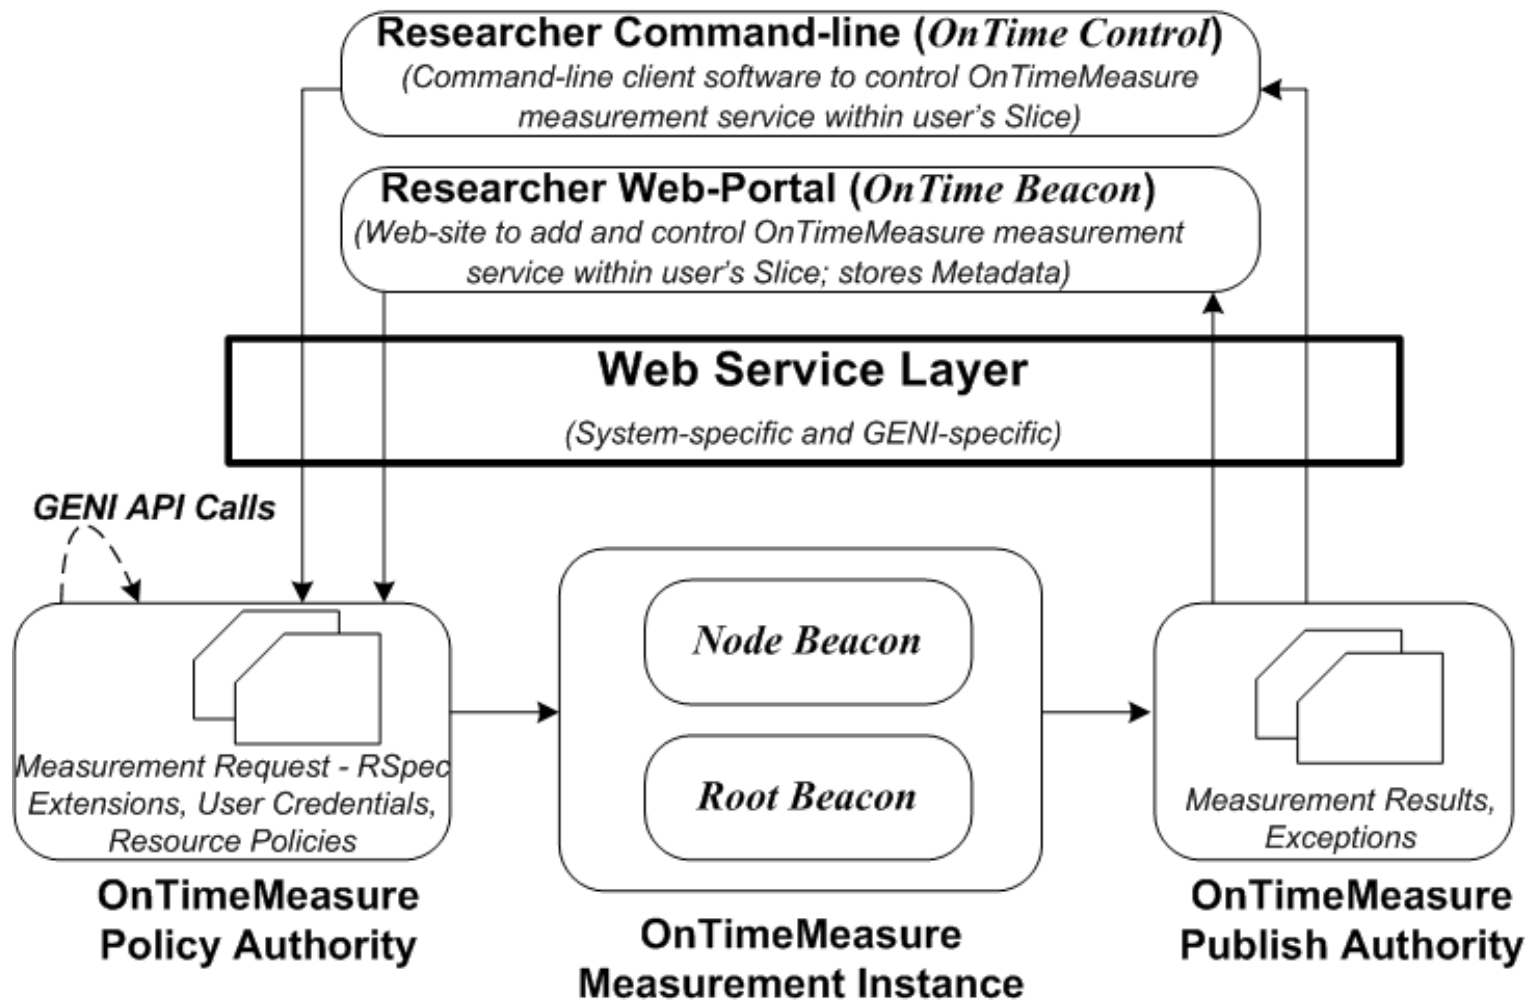
—▶	Data Flow	MDA –	Measurement Data Archive
- - -▶	Control Flow	MO –	Measurement Orchestration
MP –	Measurement Point	MAC –	Measurement Access Control
MC –	Measurement Collector	MAP –	Measurement Analysis and Presentation

Reference:

P. Calyam, M. Sridharan, Y. Xiao, K. Zhu, A. Berryman, R. Patali, "Enabling Performance Intelligence for Application Adaptation in the Future Internet", *Journal of Communications and Networks (JCN)*, 2011.

http://groups.geni.net/geni/attachment/wiki/FirstGenCalyam/fi-ontimemeasure-vdcloud_jcn11.pdf


OnTimeMeasure I&M Overview



GENI Lab-2 Q&A

1. Provide screenshots of your OnTimeMeasure instance' Graphite page for 'RoundTripDelay' of node2 -> node1, node1-> node2 and 'MemFree' of node1 and node2
2. Explain the role and functions of "Instrumentation and Measurement Tools" such as OnTimeMeasure, and GEMINI Tool Set in GENI infrastructure
3. Briefly explain in your own words the architecture of the 'OnTimeControl' framework (refer OnTimeMeasure wiki)
4. Describe the workflow that was involved when you added the custom metric feature to your OnTimeMeasure instance in the last step of your experiment


GENI Lab-3 Exercise (Planned)



- **Lab # 3 – SDN and QoE Measurement in GENI Slice**
- **Purpose of the Lab**
 - Install and configure a virtual desktop App (e.g., video player), and a network emulator within a slice
 - GENI Delay Node for network emulation - <http://geni-app-developer-documentation.readthedocs.org/en/latest/resources/delay.html>
 - Understand topology setup and OpenFlow controller configuration
 - Explore client-side encoding, and network-side transit selection adaptation and measure QoE improvement
 - Performance data collection comprises of metrics such as:
 - Perceptible Impairment MOS (Subjective QoE)
 - Bandwidth consumption (Objective QoE)


SDN and QoE Measurement Overview

(Experiment w/o Load-Balancing)

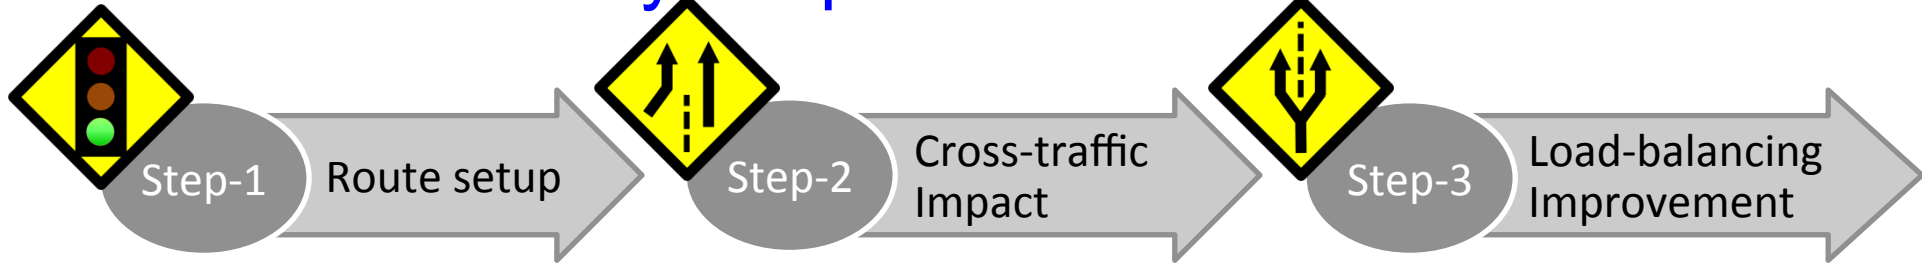


SDN and QoE Measurement Overview

(Experiment w/ Load-Balancing)



VDC-Analyst OpenFlow Demonstration




OpenFlow Switch	Client	In Port	Out Port
ATLA	PG46	20	52
ATLA	PG47	20	52

OpenFlow Switch	Client	In Port	Out Port
SUNNW	PG48	50	51
SUNNW	PG49	50	51
ATLANTA	PG46	52	52
ATLANTA	PG47	52	52
ATLANTA	PG46	20	52
ATLANTA	PG47	20	52

OpenFlow Switch	Client	In Port	Out Port
ATLANTA	PG46	20	52
ATLANTA	PG47	20	52
SUNNW	PG48	50	52
SUNNW	PG49	50	52



Video runs smooth, GUI applications are responsive

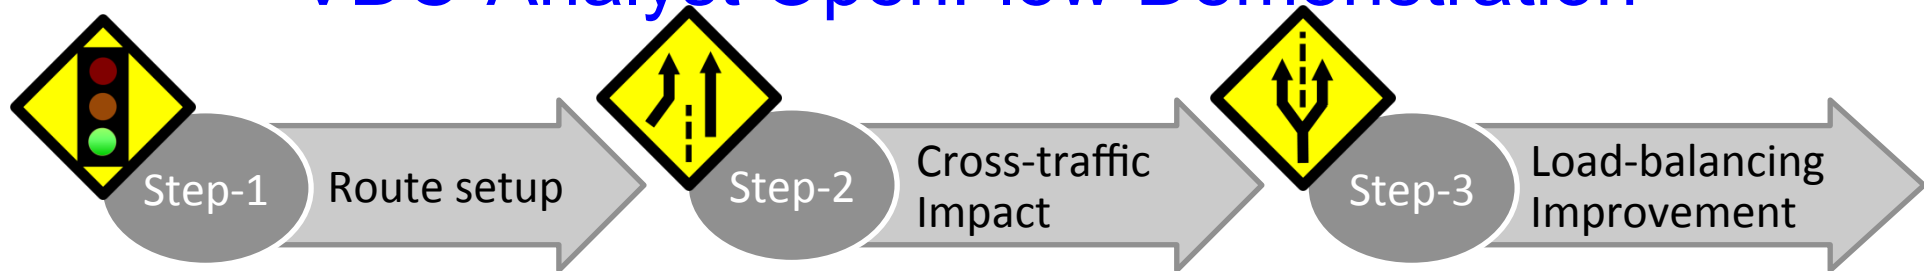


Video freezes, disconnects, GUI applications are not responsive

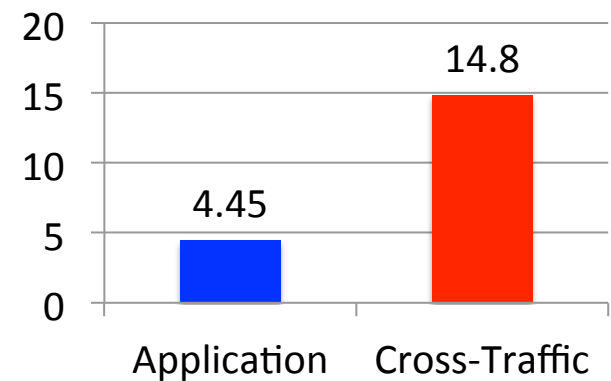
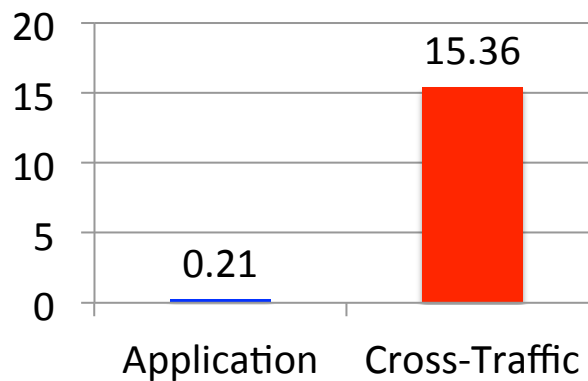
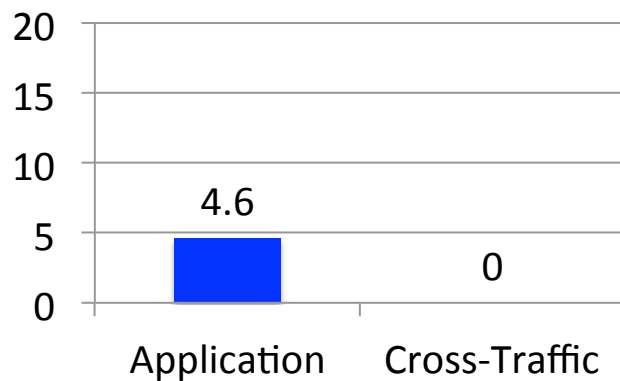


Video runs smooth, GUI applications are responsive


VDC-Analyst OpenFlow Demonstration



Bandwidth Consumed (Mbytes/s)



Video runs smooth, GUI applications are responsive




Video freezes, disconnects, GUI applications are not responsive



Video runs smooth, GUI applications are responsive


OpenFlow Topology (Planned)



Thank you for your attention!

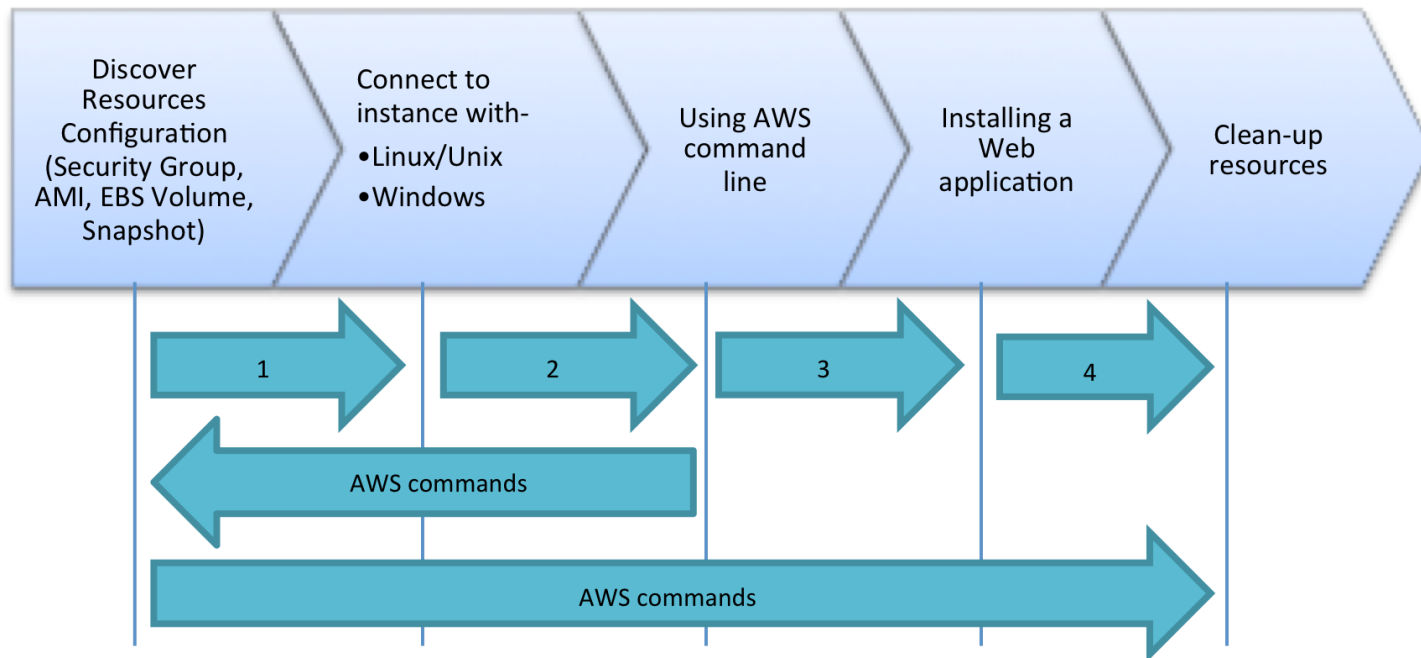


AWS Lab-1 Steps




- **Lab # 1 – AWS Account Setup and Services Overview**
- **Purpose of the Lab**
 - Understand definitions of various Amazon Web Services (AWS) and their use in cloud computing based web applications that are accessible over the Internet through an AWS account

AWS Lab-2 Steps



- **Lab # 2 – AWS Resource Discovery and Instance Setup**
- **Purpose of the Lab**
 - Use the AWS account for the discovery, reservation and access of virtual compute/storage infrastructure instances; setup a platform along with web-application and related basic firewall configurations within your reserved infrastructure resources

AWS Lab-3 Steps



- **Lab # 3 – Platform/Application Provisioning and Auto Scaling Adaptation**
- **Purpose of the Lab**
 - Launch new instances through Images taken from your current instance state, include them to a load balancer configuration, and apply CloudWatch alarms for automatically scaling up and scaling down using AutoScaling service adaptations based on usage load demands