University of Missouri

VIMAN Lab
VIRTUALIZATION, MULTIMEDIA AND NETWORKING LAB

Deploying Network-aware Applications in GENI

Presenter: **Prasad Calyam**, Ph.D.
Department of Computer Science

Graduate Student: *Ronny Bazan Antequera*

GENI Webinar Series, Fall 2016

Tutorial Topics

- Future Internet Performance Architecture (FIPA)
  - Requirements and Design
- OnTimeMeasure-GENI: FIPA Reference Implementation
  - Hands-on Demo
    - Working with instrumentation and measurement services
    - Adding custom metrics of your Application
- FIPA Case Study: Dynamic Resource Allocation in Virtual Desktop Clouds

Tutorial Topics

- Future Internet Performance Architecture (FIPA)
  - Requirements and Design
- OnTimeMeasure-GENI: FIPA Reference Implementation
  - Hands-on Demo
    - Working with instrumentation and measurement services
    - Adding custom metrics of your Application
- FIPA Case Study: Dynamic Resource Allocation in Virtual Desktop Clouds

User Slice in Future Internet

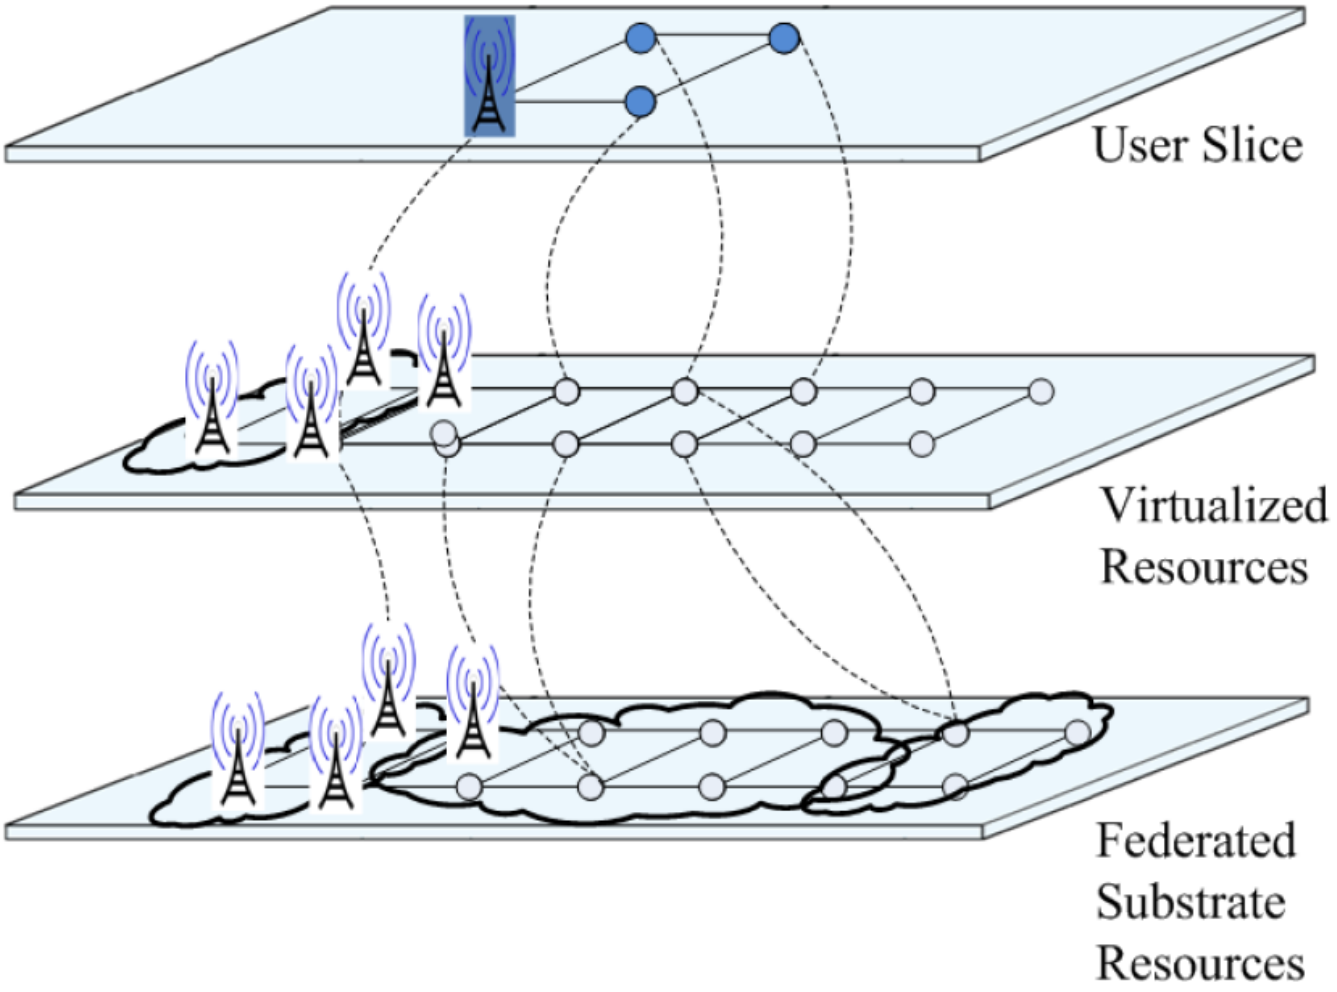

Future Internet Application Life-cycle

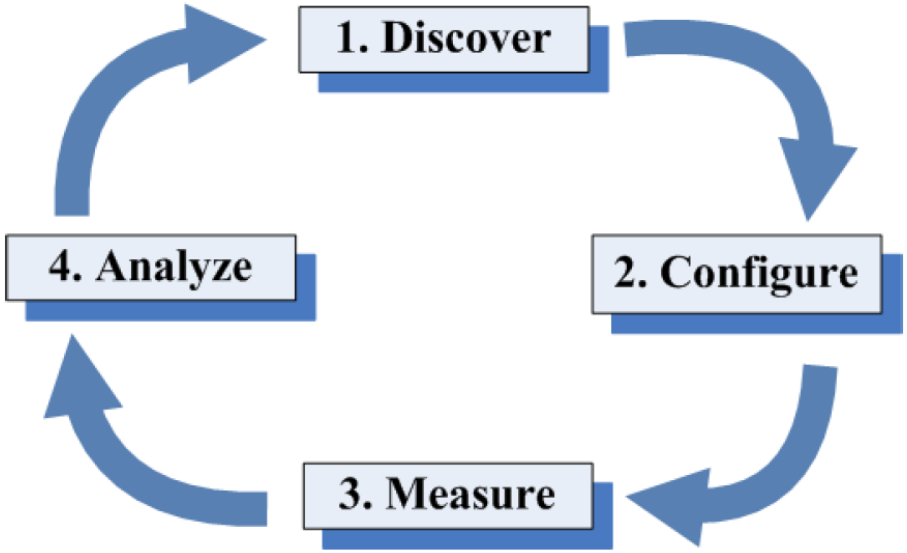

- Awareness in host and network components
- Future Internet application's *autonomicity* attributes
  - Self-configuring, Self-managing, Self-monitoring, Self-optimizing
- Enabling *performance intelligence* in Future Internet applications is vital

FI Performance Architecture (FIPA)

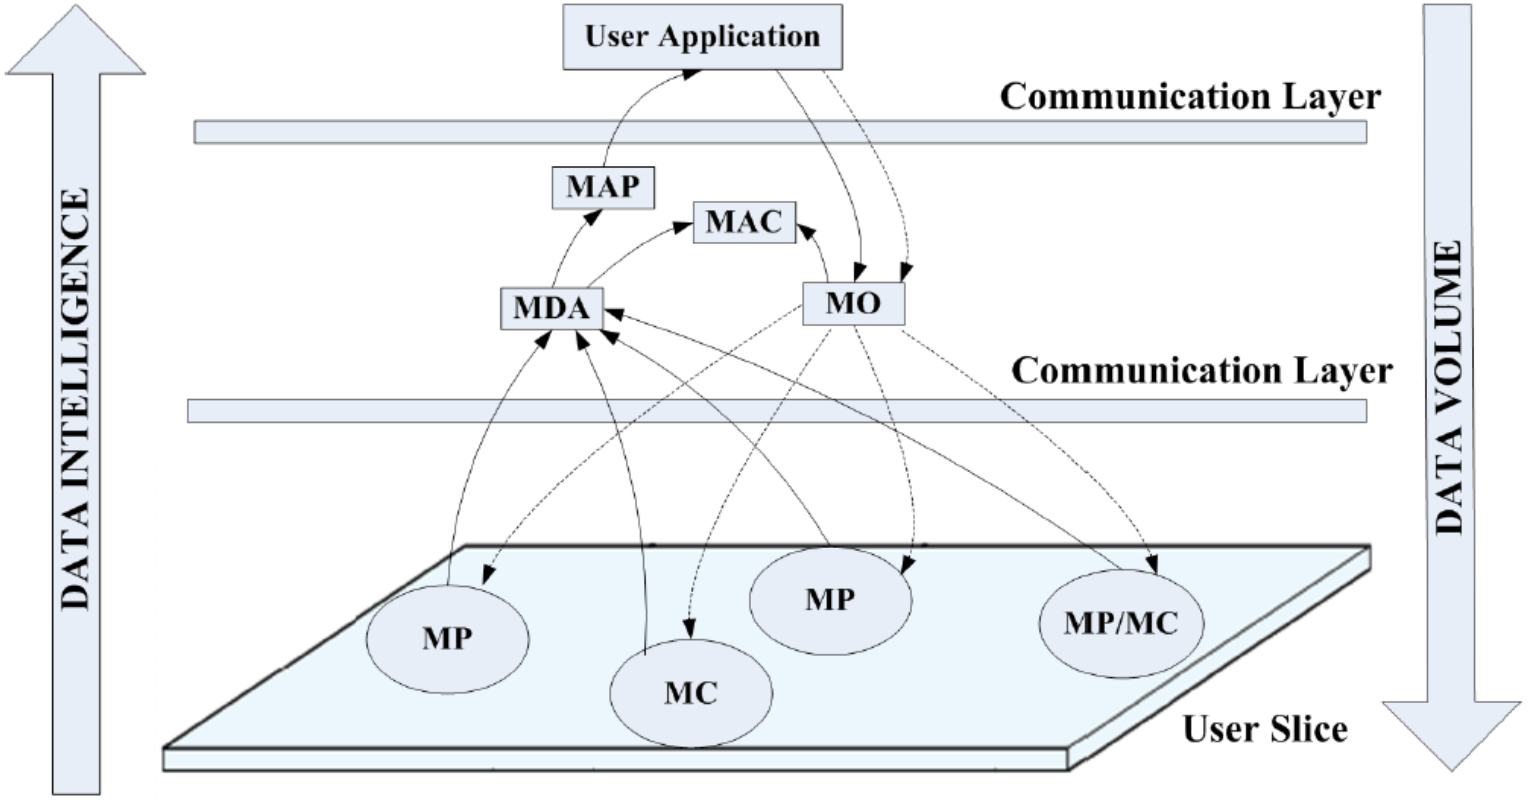

Legend:

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

DIY: I&M service for GENI Experimenters!

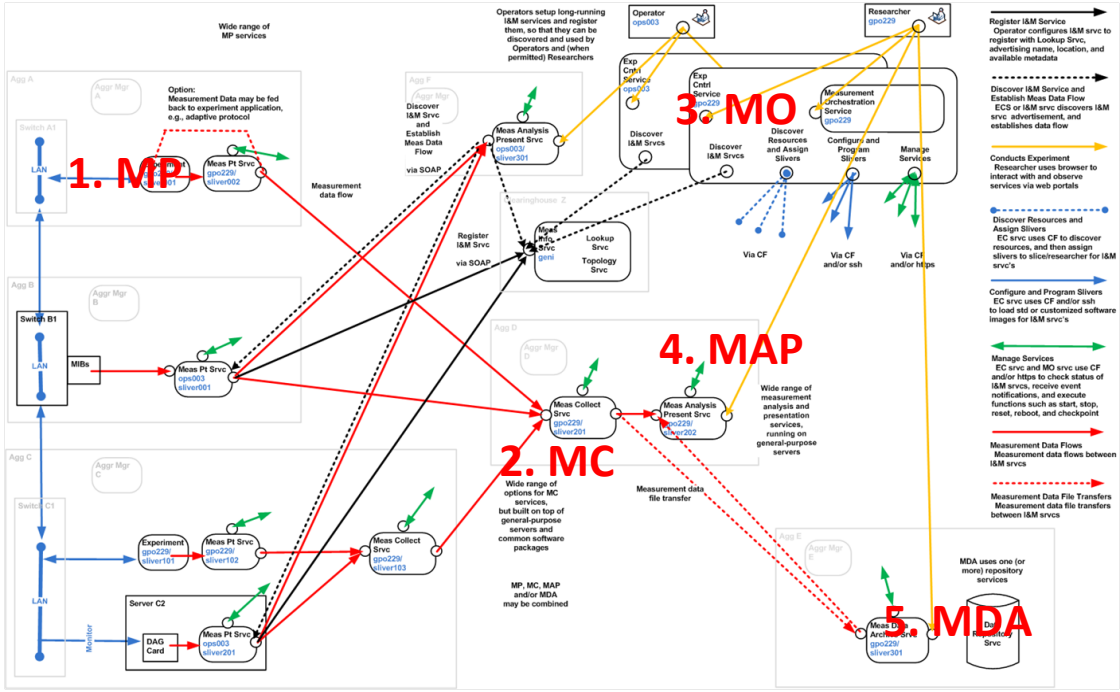

GENI I&M Architecture (Credit: Harry Mussman, GPO)

Every GENI experiment requires a custom I&M service; a GENI I&M service should remove most of that burden for an experimenter….

GENI Experimenter Use Cases

- Experiment Environment Monitoring
  - Have I got the system and network resources I asked in my slice?
  - Is the experiment environment functioning as expected over my slice lifetime?
  - Can I trace my non-intuitive results in my experiment to a problem in the slice environment?

- Experiment Instrumentation and Measurements (I&M) Management
  - I want to collect active and passive measurements at hop, link, path and slice levels for my experiments…
  - I need measurements feedback with anomaly detection and forecasts to control my experiment progress…
  - I would like to have a measurement framework in my experiment to be configured once and be re-usable for subsequent experiment runs…
  - I want an archive of the experiment measurements collected for offline analysis or to share with colleagues…

Tutorial Topics

- Future Internet Performance Architecture (FIPA)
  - Requirements and Design
- OnTimeMeasure-GENI: FIPA Reference Implementation
  - Hands-on Demo
    - Working with instrumentation and measurement services
    - Adding custom metrics of your Application
- FIPA Case Study: Dynamic Resource Allocation in Virtual Desktop Clouds

OnTimeMeasure-GENI Overview

GENI Project Wiki – http://groups.geni.net/geni/wiki/OnTimeMeasure

- Goal: Provide GENI community with capabilities for provisioning on-going and on-demand measurement requests
  - Used in forecasting, anomaly detection, and fault-location diagnosis in GENI experiments and GENI operations
- Outcomes:
  - *Software* to perform measurement orchestration and provisioning of measurements
    - For continuous monitoring, persistent measurements storage and processed network measurement feeds
  - *Measurement service* that enables users to utilize OnTimeMeasure software in GENI experiments
    - Registers users, slices, maintains meta-data, and allows user control of measurement service functions
      - Researcher Web-portal – http://ontime.rnet.missouri.edu for *interactive user control* of measurement service
      - Command-line tools for measurement service *control automation*

OnTimeMeasure Architecture

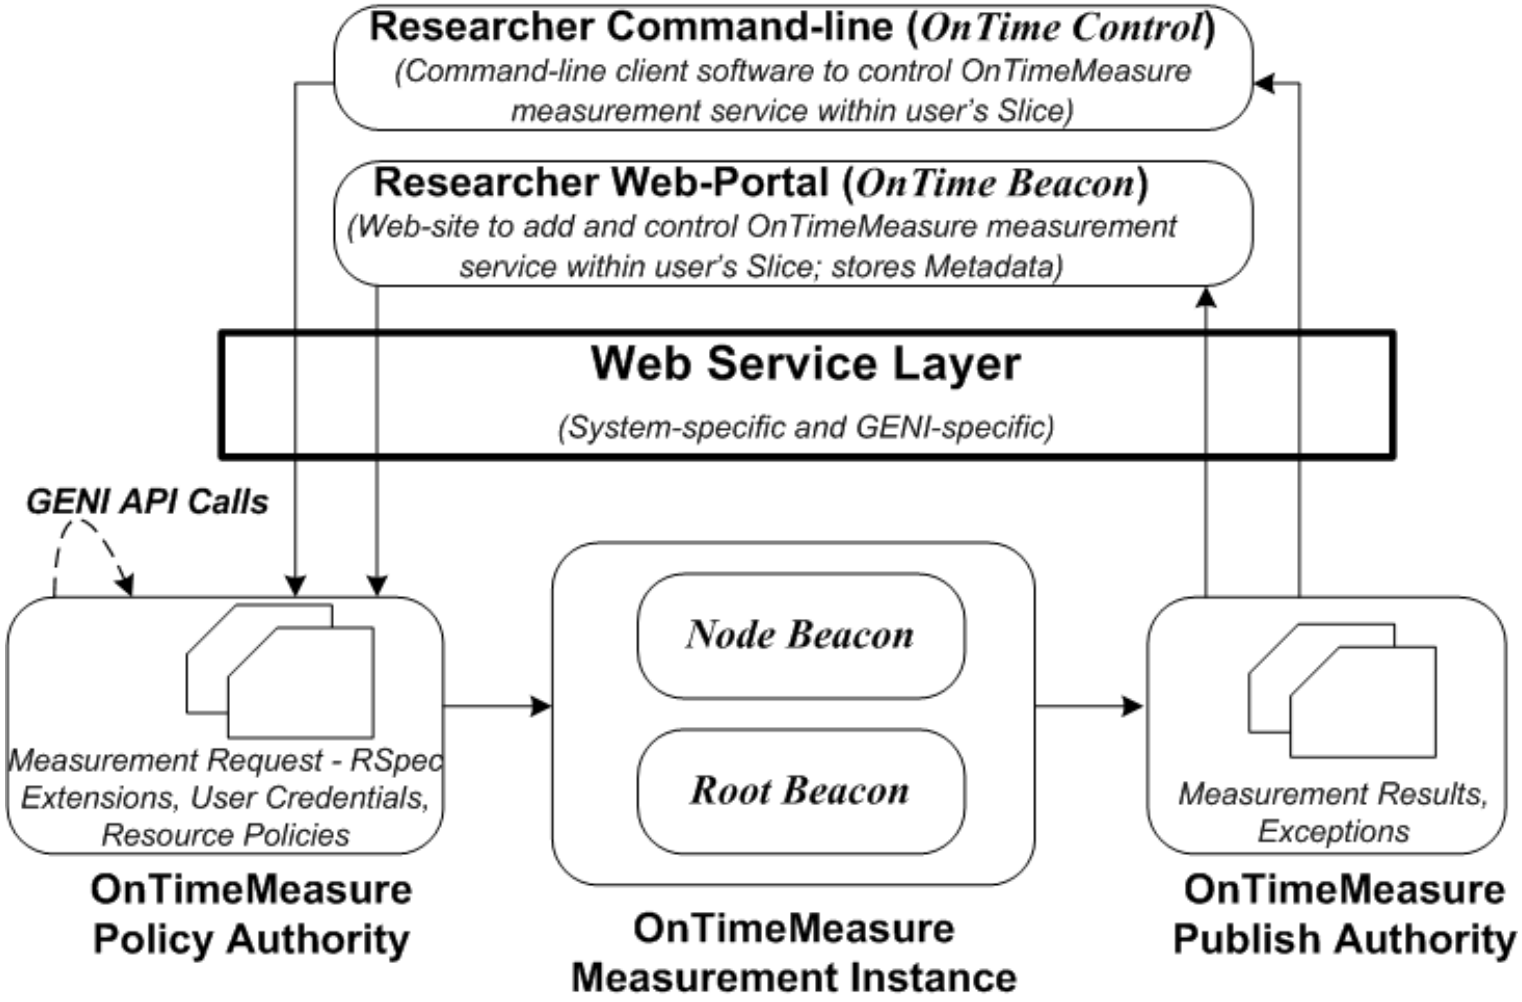

OnTimeMeasure Software Modules

- Customizable software developed by P. Calyam & Team at The Ohio State University
- Two main modules installed within a GENI experiment slice
  - **Node Beacon**
    - Installs tools that measure network health metrics such as: route changes, delay, jitter, loss, bandwidth
      - TCP/UDP Iperf, Traceroute, Ping, Pathload, OWAMP, etc.
    - Runs measurements based on a schedule and outputs results
  - **Root Beacon**
    - Installs Apache, MySQL and other packages
    - Creates database tables and configuration files
    - Generates measurement schedules for node beacons
    - Collects data and provides dashboard visualization, statistical analysis (i.e., anomaly detection and weather forecasting) with alarm generation

    NOTE: Typically, two or more Node Beacons and one Root Beacon need to be installed on separate dedicated servers in a slice; it is however possible to install Root Beacon on one of the Node Beacon servers

What OnTimeMeasure can provide Experimenters?

- Data Aggregation
  - Your distributed data sources can be controlled and accessed (start/stop/query) in a centralized manner via web-portal or command line
- Data Visualization
  - Measurement data graphs, dashboards
- Data Analysis
  - Time Series files with Anomalies/Time Series with Forecasts/others
  - Ability to use analysis of the measurements to reconfigure the measurement specifications without human intervention or use in research experiment schemes
- Data Archive
  - Slice owners can access and download data and metadata; transfer to GENI archive

OnTimeMeasure-GENI Hands-on Tutorial

Follow OnTimeMeasure Tutorial Steps at –
http://groups.geni.net/geni/wiki/GENIExperimenter/GENIWebinars/DeployingNetworkAwareApplicationsInGENI

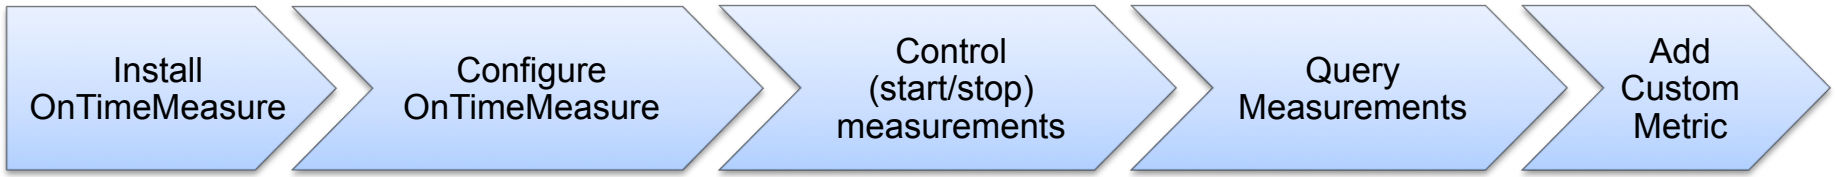

- **Purpose of the Tutorial**
  - **Perform Instrumentation and Measurement of a GENI Slice**
    - Install and configure a distributed measurement Web App viz., "OnTimeMeasure" within a slice. Configuration involves 'Node Beacons' at measurement points, and a 'Root Beacon' that acts as a central intelligence 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.
    - Add custom metrics on system health metrics such as CPU utilization

Screenshots

View OnTimeMeasure Demo Videos at – http://ontime.rnet.missouri.edu/demo

Centralized Service Control

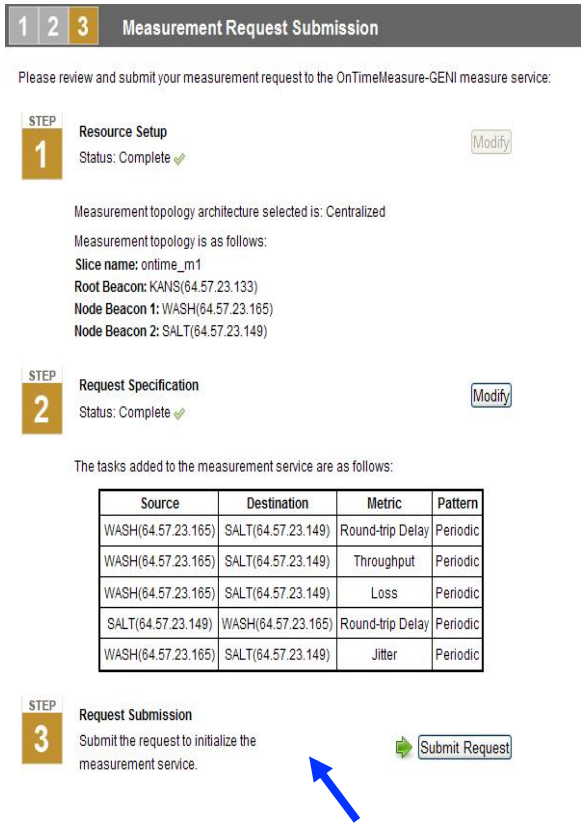

Centralized Measurement Request

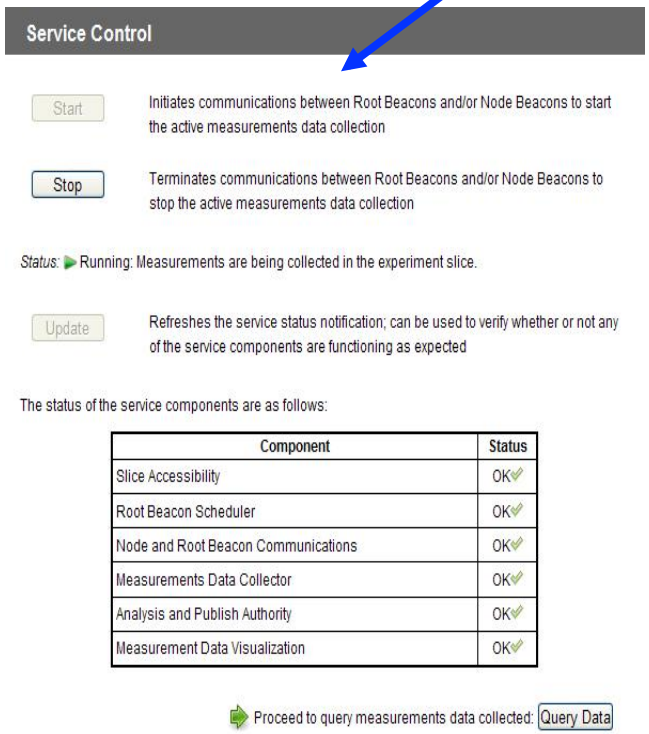

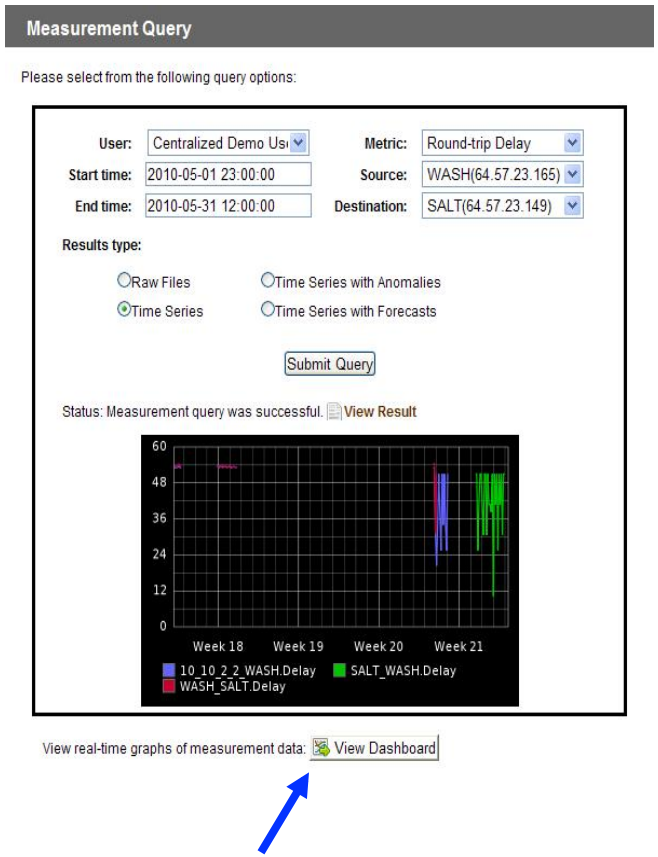

Centralized Measurement Query

Screenshots (2)

View OnTimeMeasure Demo Videos at – http://ontime.rnet.missouri.edu/demo

Researcher Web-Portal

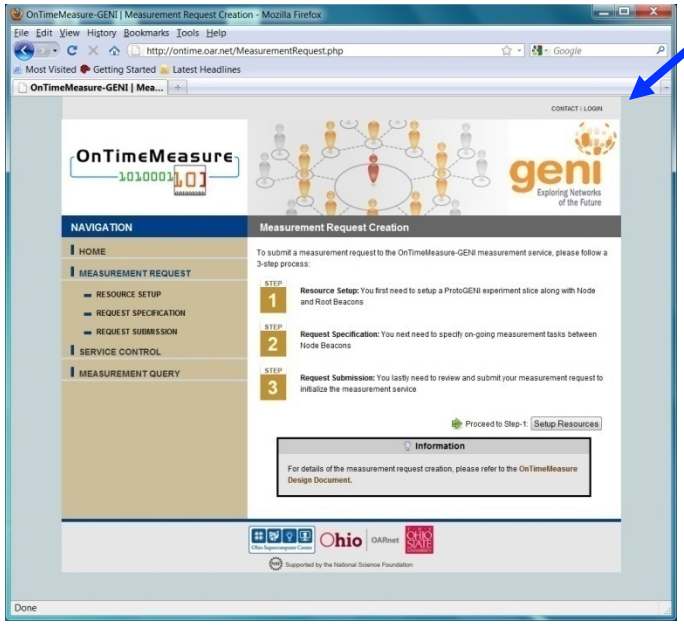

Measurement Result

Measurement completed, please see below. Download raw files: WASH

```
[ 15] local 64.57.23.149 port 5001 connected with 64.57.23.165 port 5001
[ 15]  0.0-10.0 sec    494 MBytes    414 Mbits/sec
[ 15] MSS size 1448 bytes (MTU 1500 bytes, ethernet)
bwctl: stop_exec: 3484246228.794817

RECEIVER END
```

```
[WASH]$ bwctl -c 64.57.23.149 -f m -u -b 768k -a 10

RECEIVER START
bwctl: exec_line: iperf -B 64.57.23.149 -s -f m -m -p 5001 -u -t 10
bwctl: start_tool: 3484246254.590082
------------------------------------------------------------
Server listening on UDP port 5001
Binding to local address 64.57.23.149
Receiving 1470 byte datagrams
UDP buffer size: 0.11 MByte (default)
------------------------------------------------------------
[  8] local 64.57.23.149 port 5001 connected with 64.57.23.165 port 5001
[  8]  0.0-10.0 sec  0.92 MBytes  0.77 Mbits/sec  0.011 ms    0/  655 (0%)
bwctl: stop_exec: 3484246280.789142

RECEIVER END
```

Distributed Measurement Result

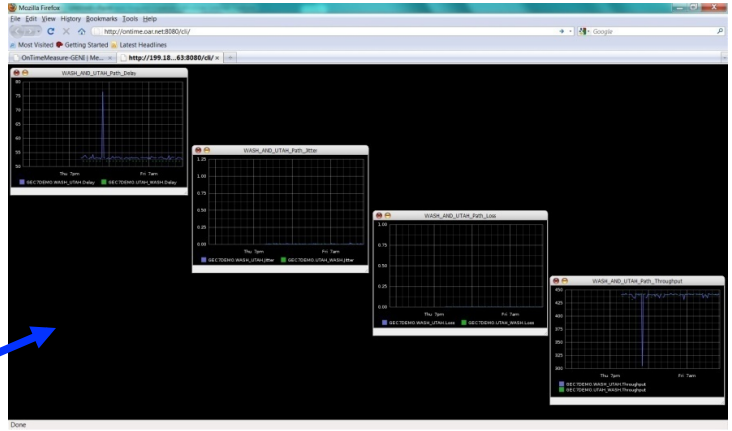

User Customizable Dashboard

Adding a Custom Metric into OnTimeMeasure

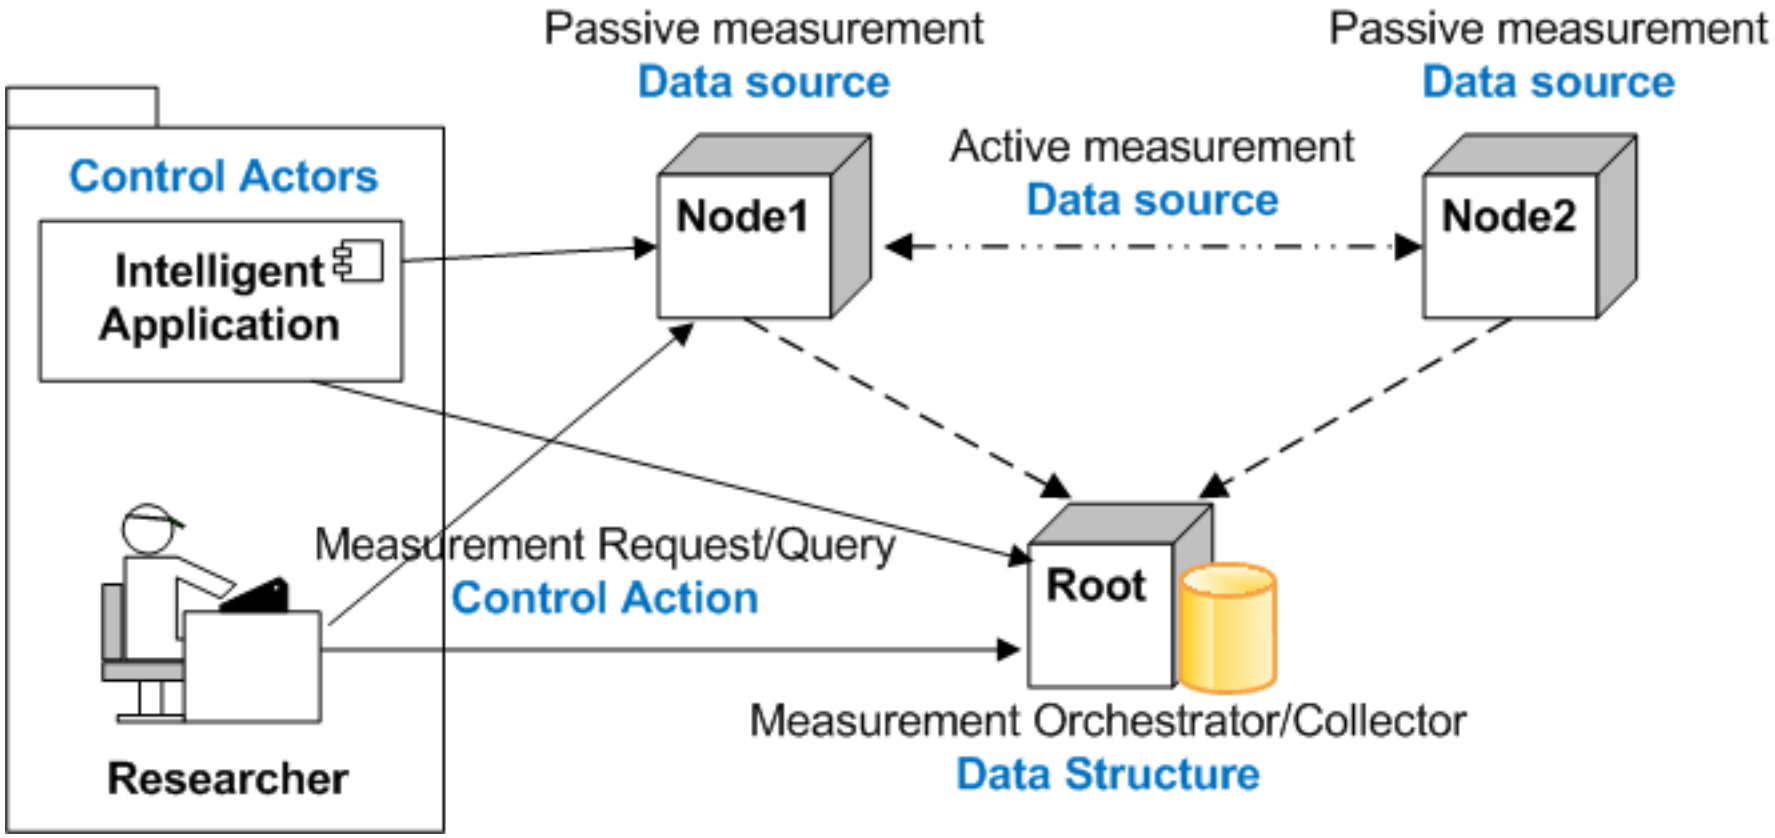

Definitions for Custom Metric Specification

- ***Control Actors***: Project members or your research experiment who can provide the "slice password" to the OnTimeMeasure instance in an experiment slice, and would access the I&M data or share the data with other actors
- ***Data Sources***: Data generation tools deployed in slice nodes; the tool would communicate with other nodes to perform active measurements or inherently collects passive measurements in on-going and on-demand manners
- ***Data Structures /Data Types***: Measurement data that needs to be stored in a database with appropriate data structure and data types
- ***Control Actions***: Control both start and stop of the data generation tool, control how to retrieve and utilize the data e.g., dashboard, plots, time series files with anomaly annotation

More Information at - http://groups.geni.net/geni/wiki/OTM-CustomMetric

Tutorial Topics

- Future Internet Performance Architecture (FIPA)
  - Requirements and Design
- OnTimeMeasure-GENI: FIPA Reference Implementation
  - Hands-on Demo
    - Working with instrumentation and measurement services
    - Adding custom metrics of your Application
- FIPA Case Study: Dynamic Resource Allocation in Virtual Desktop Clouds

Virtual Desktop Clouds (DaaS)

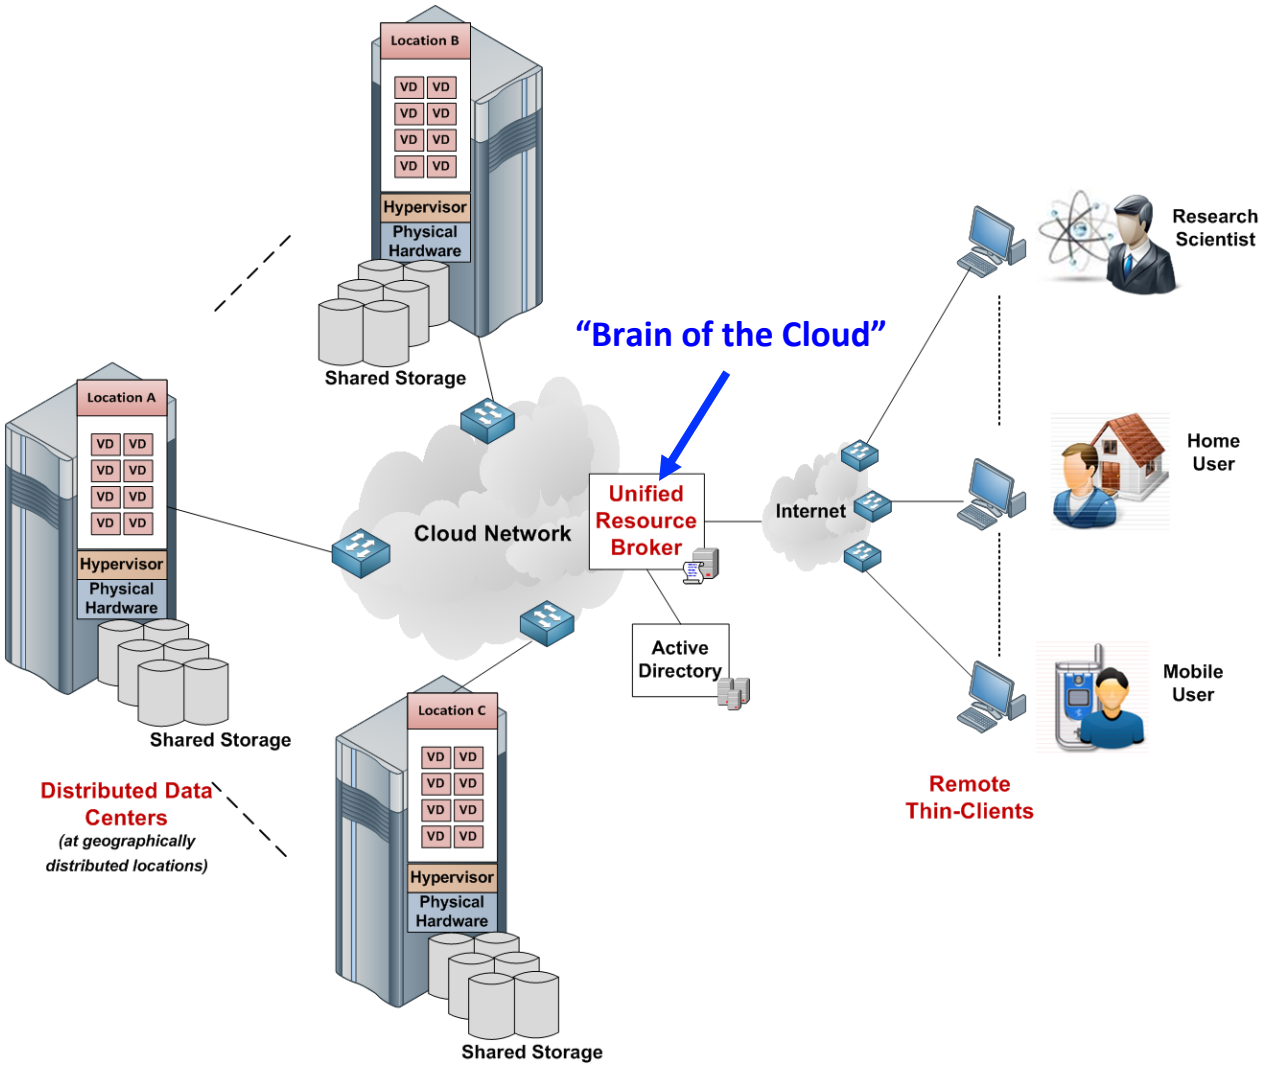

VDCs Today – Overprovisioning and Guesswork…

- High consistent CPU
- High consistent memory
- High bandwidth connectivity

Research Scientist

VDC Service Provider

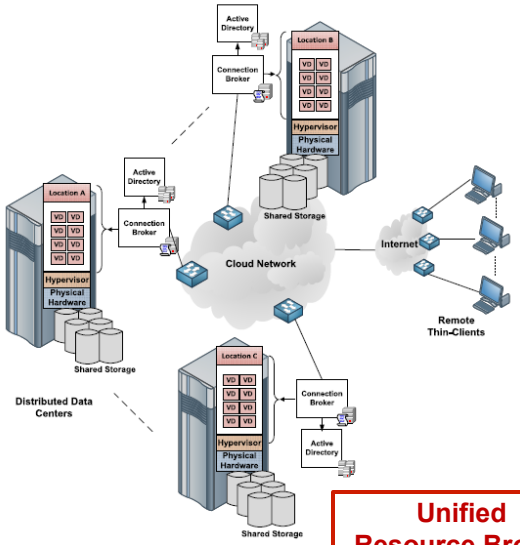

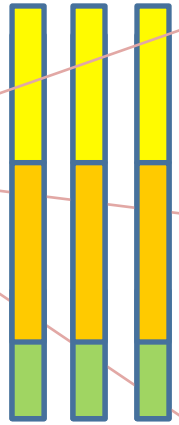

CPU Memory Bandwidth

Unified Resource Broker

Fixed Resource Allocation Model (F-RAM) = Available Resources / Number of Users

- Low bursty CPU
- Low bursty memory
- Medium bandwidth connectivity

Home User

- Low bursty CPU
- Low bursty memory
- Low bandwidth connectivity

Mobile User

Overprovisioning and Guesswork Fails!

- Calls from unhappy customers
- High operation $$

VDC Service Provider

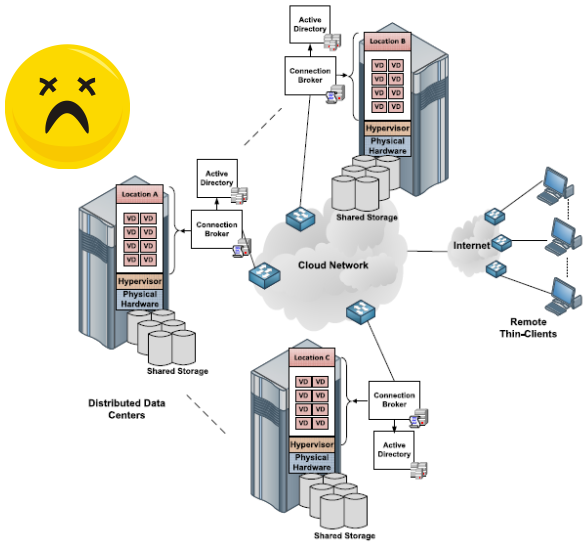

- Inadequate CPU, memory and bandwidth

(Impact e.g., Slow interaction response times)

Research Scientist

- Inadequate CPU, memory and bandwidth

(Impact e.g., IPTV with impairments and slow playback)

Home User

- Excess CPU, memory and bandwidth

(Impact e.g., Good interaction response times and smooth IPTV playback)

Mobile User

Problem: Resource allocation without awareness of system, network and user experience characteristics

VDCs in the Future – Smart set-top boxes at user sites

- Happy customers
- Low operation $$

- Utility-directed CPU, memory and bandwidth

(Impact e.g., Good interaction response times and smooth IPTV playback)

VDC Service Provider

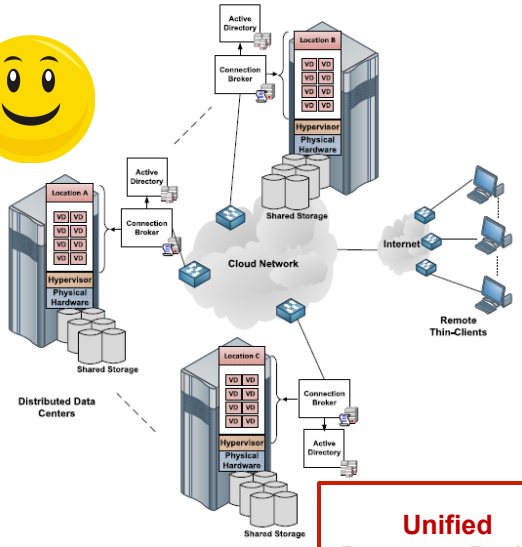

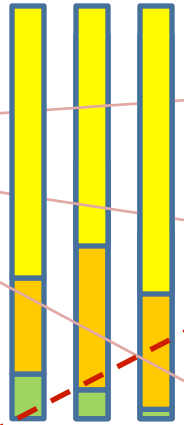

Unified Resource Broker

Utility-directed Dynamic Resource Allocation Model (U-RAM)

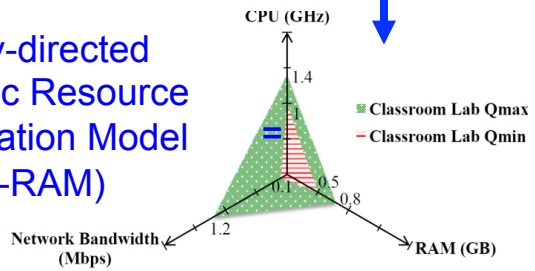

Research Scientist

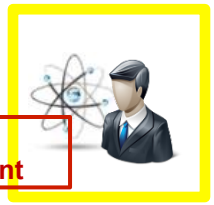

Smart Thin-Client

Home User

Smart Thin-Client

Mobile User

Virtual Desktop Cloud – FIPA Implementation

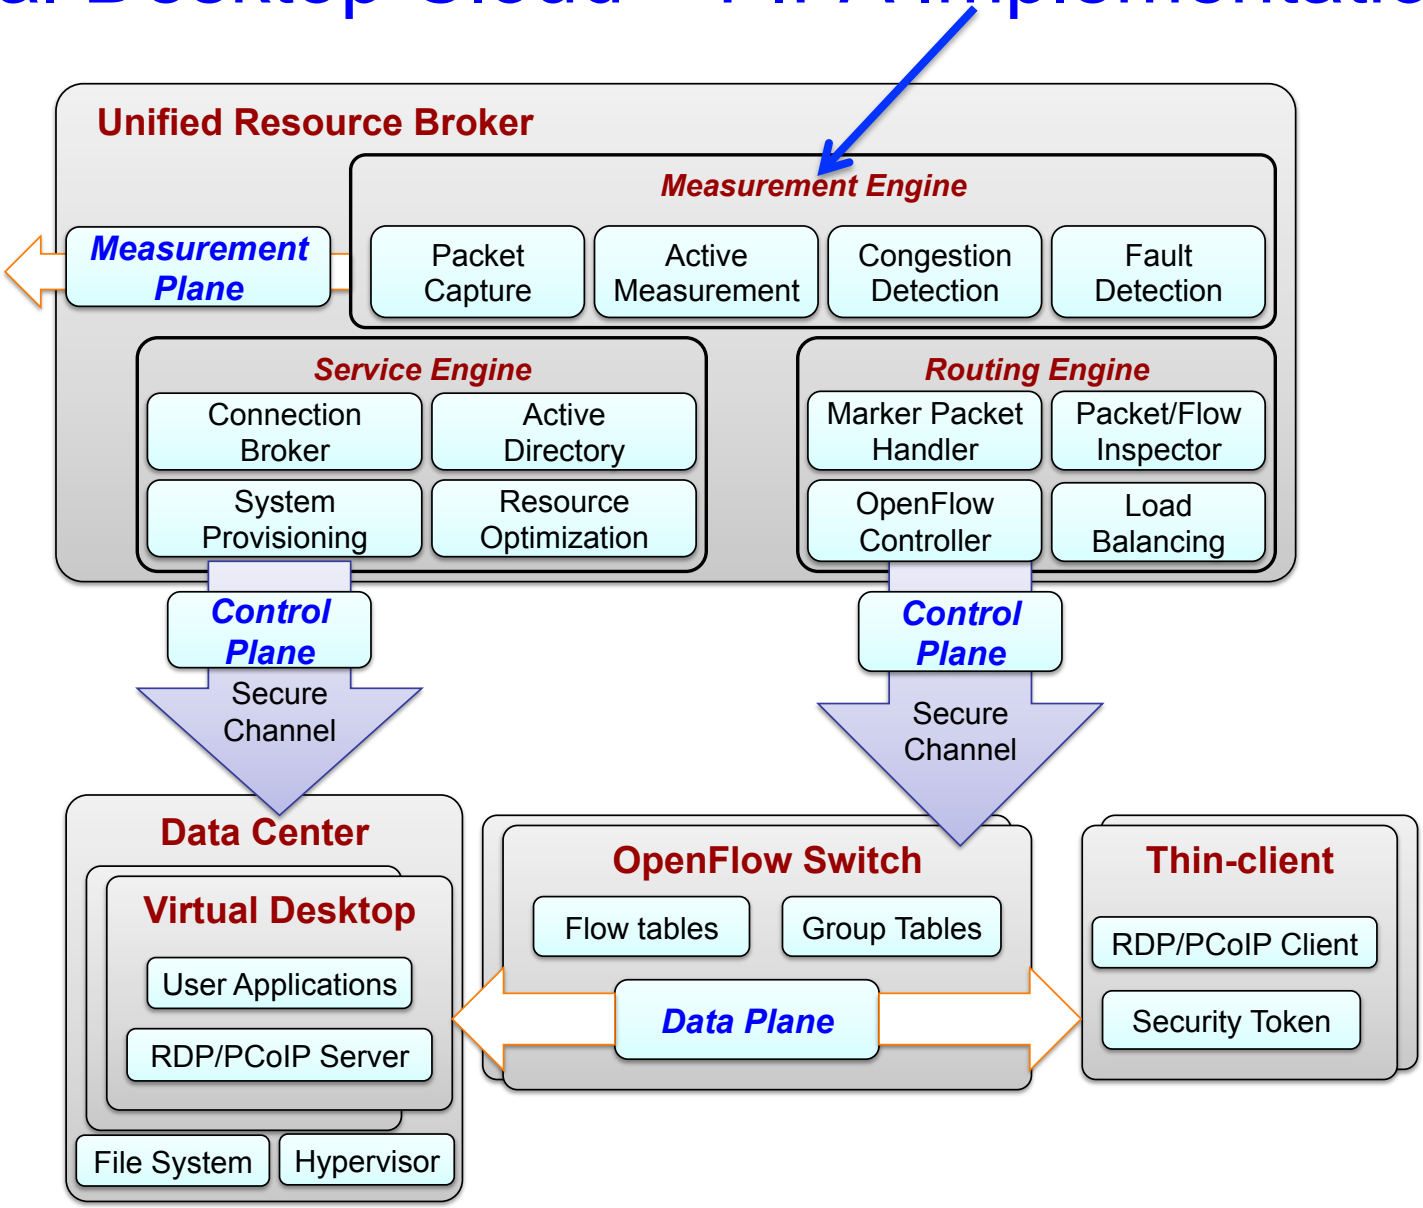

Virtual Desktop Provisioning Experiment Demonstration

- Compared U-RAM and F-RAM performance
- Created a datacenter in VMLab
  - One physical server each in VMLab for U-RAM and F-RAM
- VDC clients on ProtoGENI slice in GEC10
  - Developed a web-portal to launch VDC clients and control network emulation for demonstration
- Leveraged OnTimeMeasure's custom metric creation capability
  - Path-based measurements of network health such as delay, available bandwidth, loss
  - Host-based measurements from VMware VDI tools such as CPU, memory, number of VM connections

GEC10 Experiment Demonstration Setup

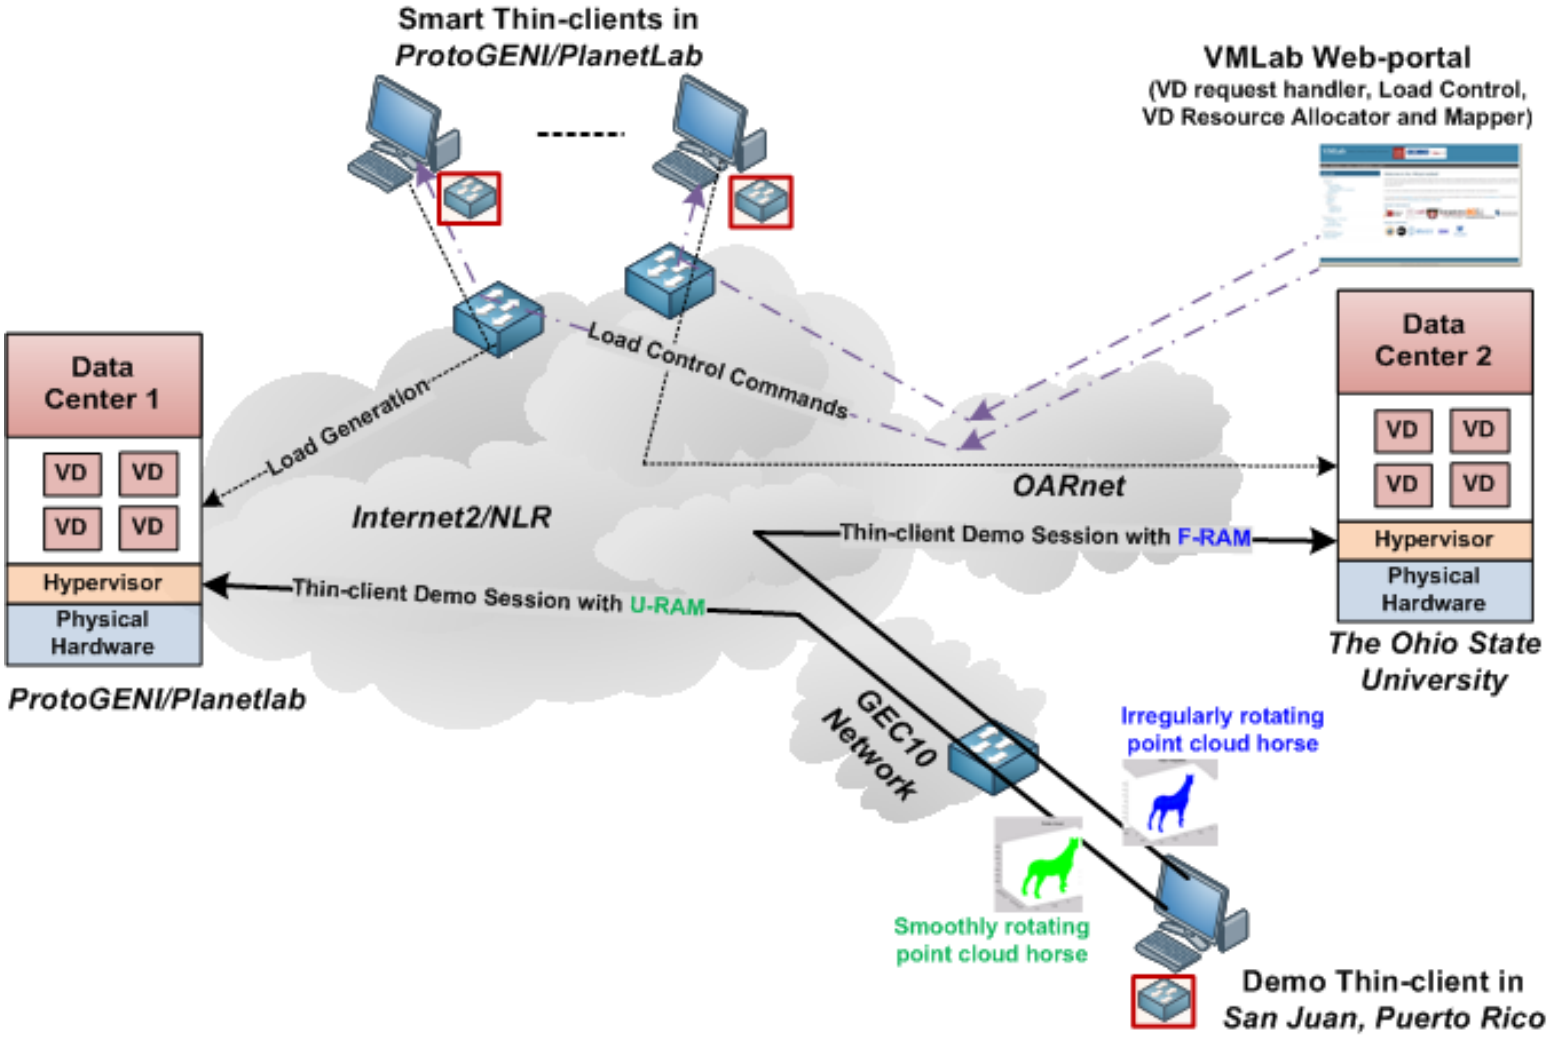

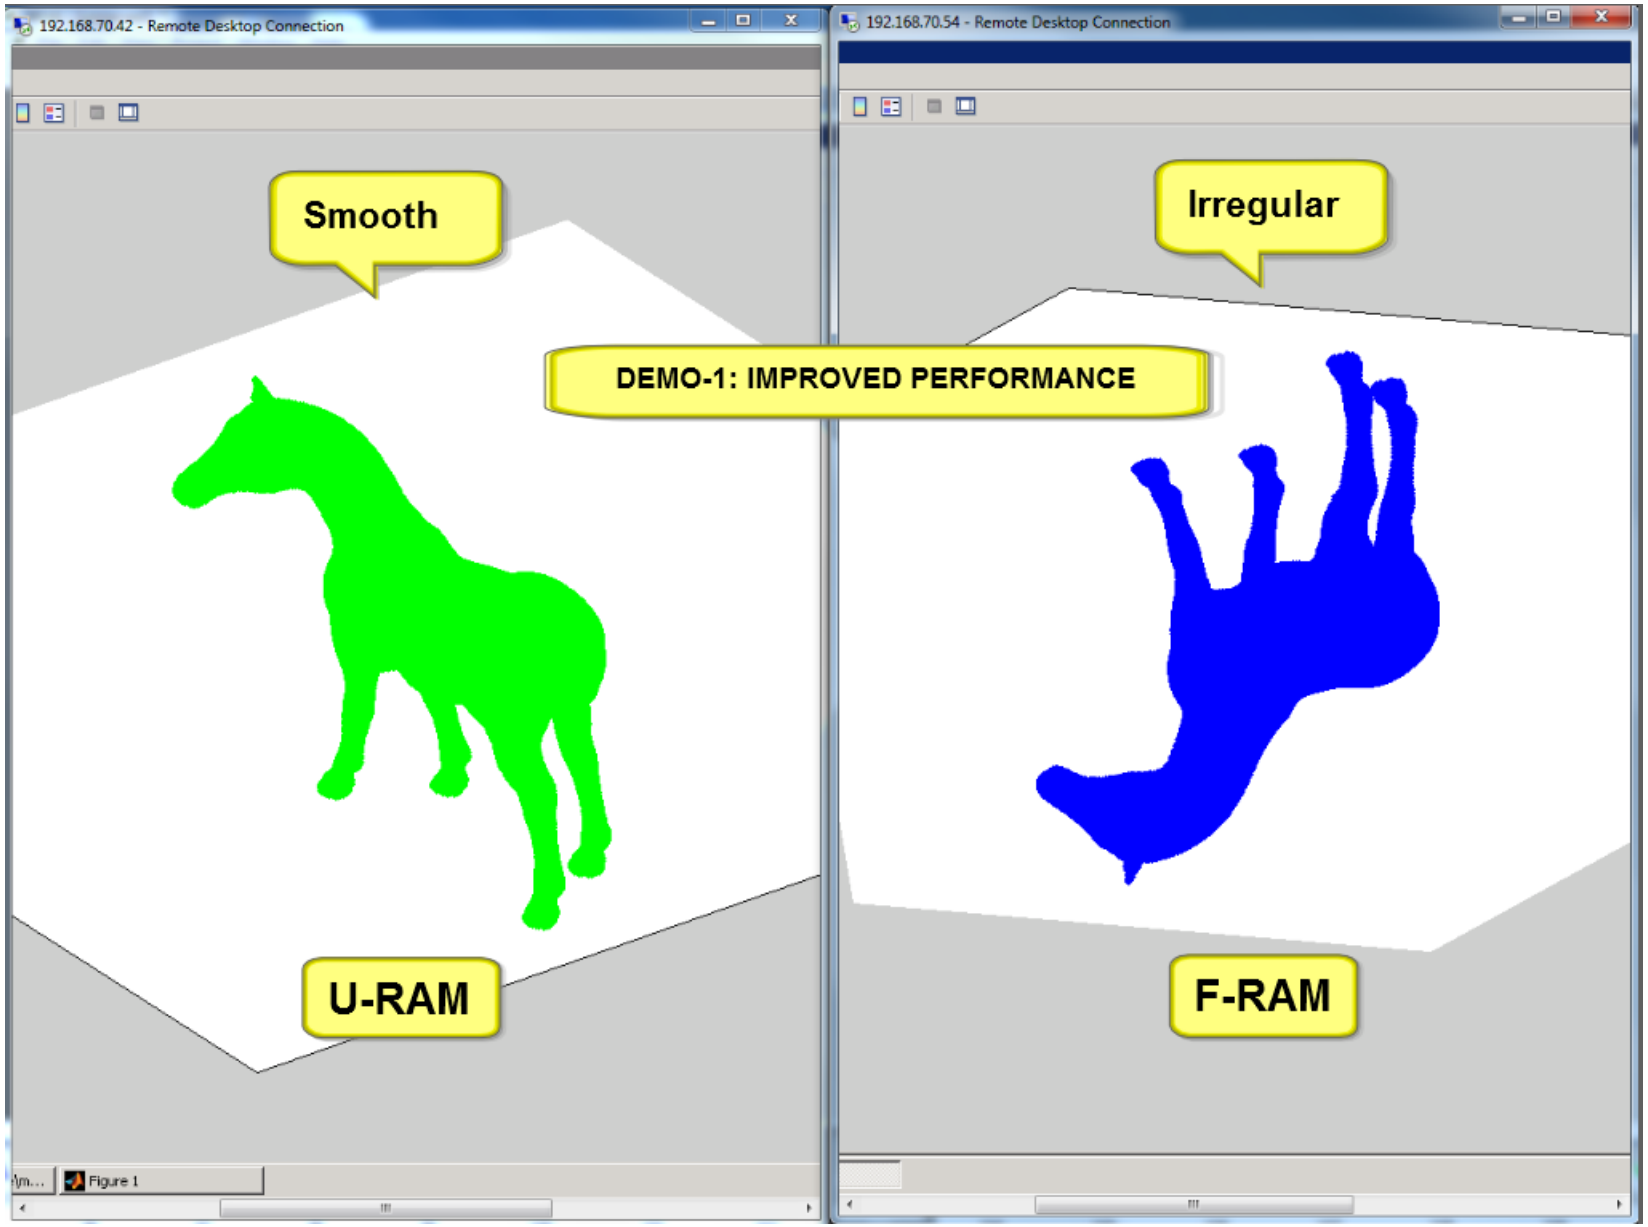
192.168.70.42 - Remote Desktop Connection
192.168.70.54 - Remote Desktop Connection
Smooth
Irregular
DEMO-1: IMPROVED PERFORMANCE
U-RAM
F-RAM
Figure 1

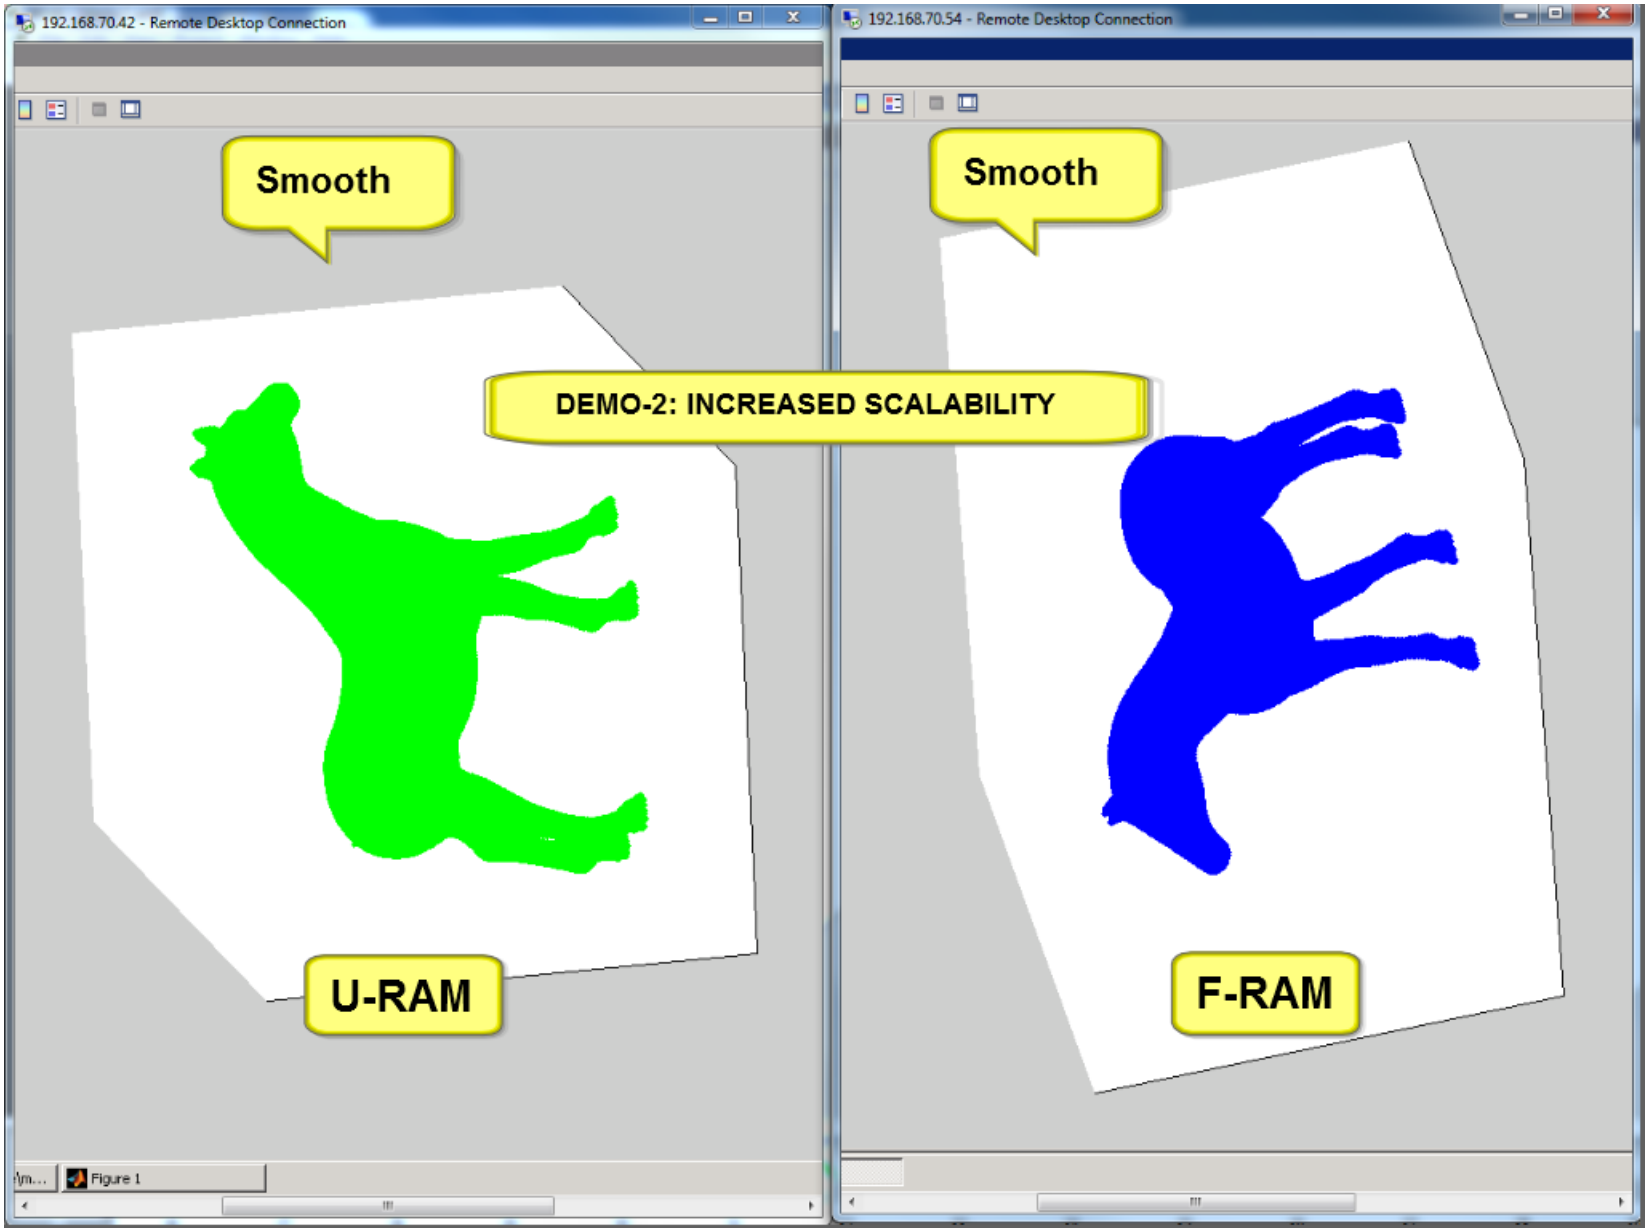
192.168.70.42 - Remote Desktop Connection
Smooth
U-RAM
Figure 1
192.168.70.54 - Remote Desktop Connection
Smooth
F-RAM
DEMO-2: INCREASED SCALABILITY

Virtual Desktop Placement after U-RAM Provisioning

- Placement decisions involving data centers are influenced by:
  - Session latency, Load balancing, Operation cost
- Compared network-adaptation performance when backup path switching occurred upon detecting a cross-traffic congestion scenario
- Multi-domain GENI testbed at GEC15 involved a slice setup
  - Using “OSU VMLab – Utah Emulab” VLAN
  - End-to-end Path: OSU – OARnet – Internet2/NLR Backbone – Utah
- Leveraged OnTimeMeasure’s custom metric creation capability
  - Path-based measurements of network health such as delay, available bandwidth, loss
  - Host-based measurements from VMware VDI tools such as CPU, memory, number of VM connections

Virtual Desktop Placement Experiment w/o Load-Balancing

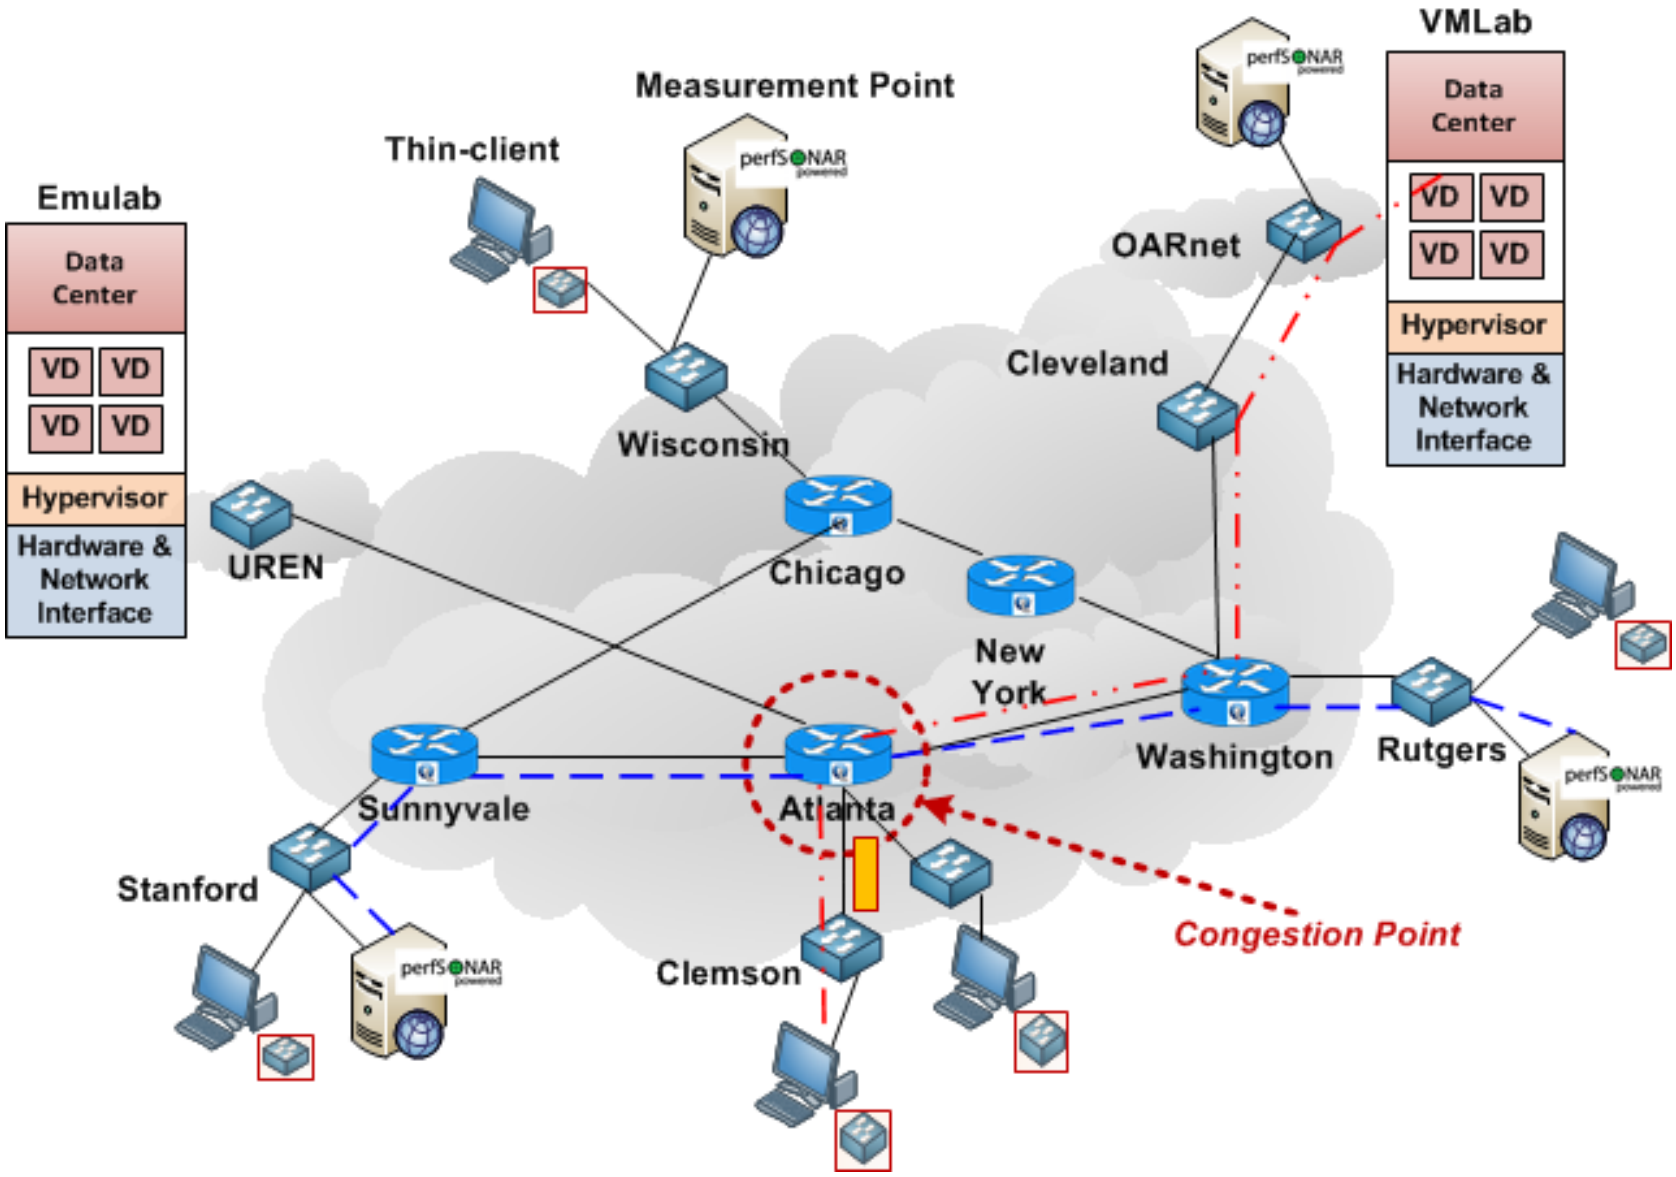

Virtual Desktop Placement Experiment w/ Load-Balancing

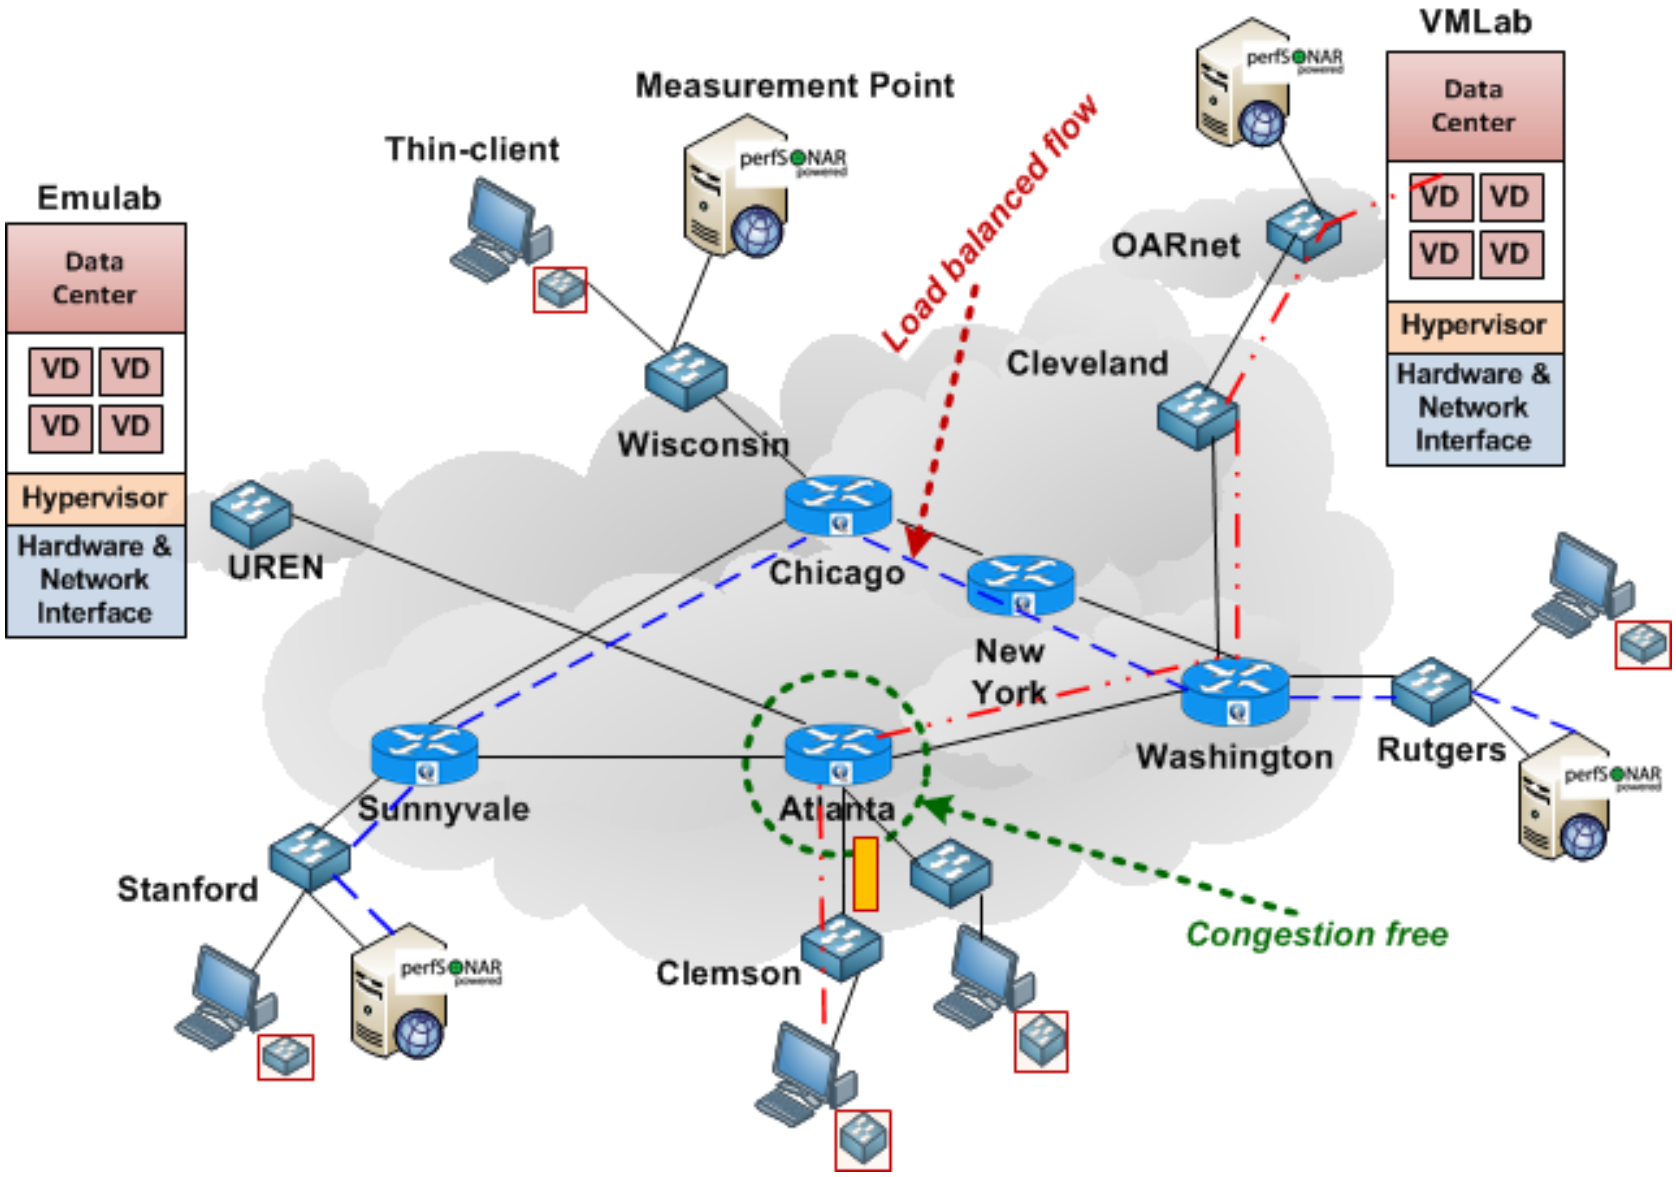

OpenFlow-based Network Adaptation Demonstration

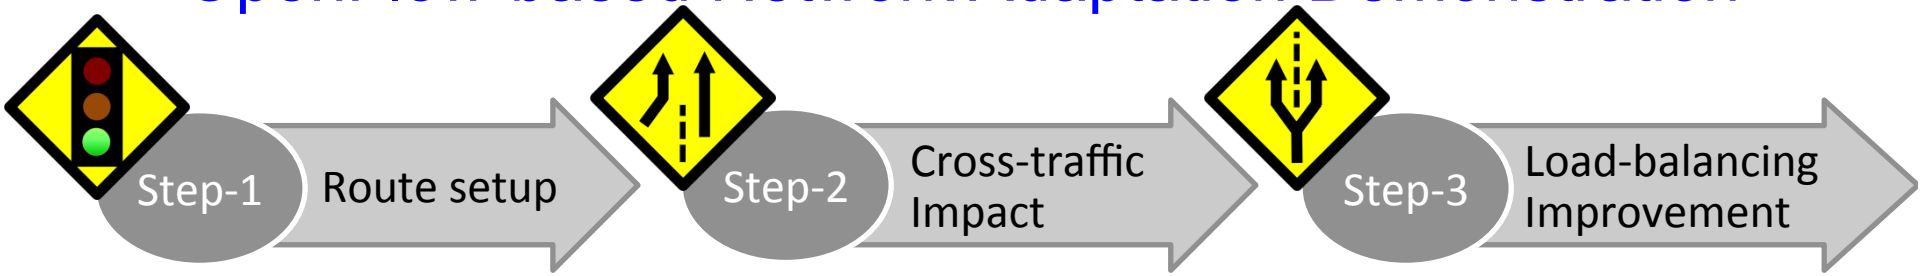

OpenFlow Switch	Client	In Port	Out Port
ATLA	PG46	20	52
ATLA	PG47	20	52

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

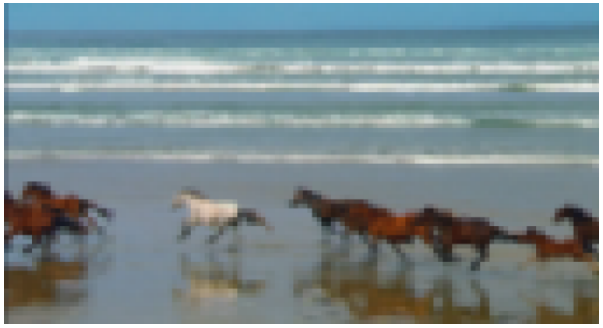

Video freezes, disconnects, GUI applications are not responsive

Video runs smooth, GUI applications are responsive

OpenFlow-based Network Adaptation Demonstration (2)

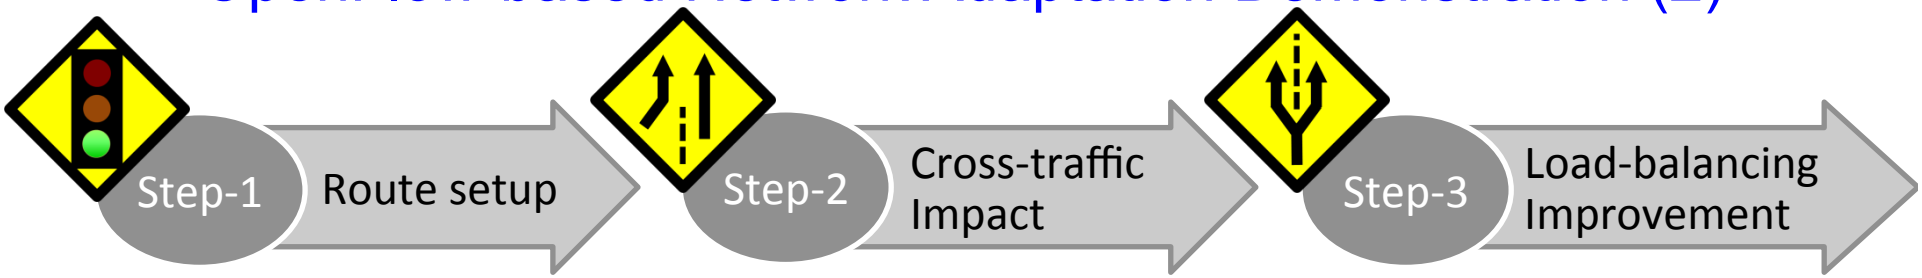

Bandwidth Consumed (Mbytes/s)

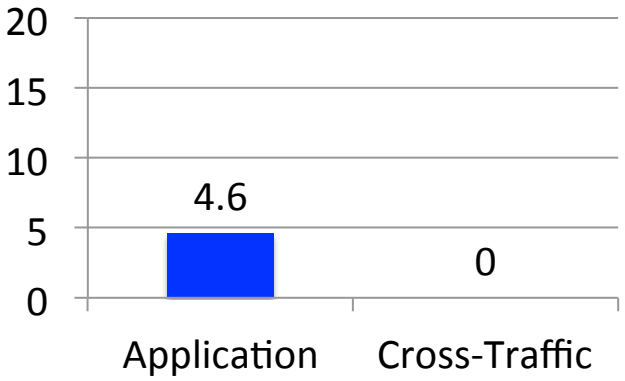

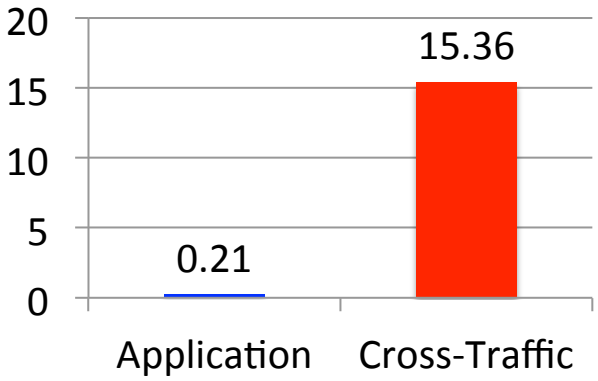

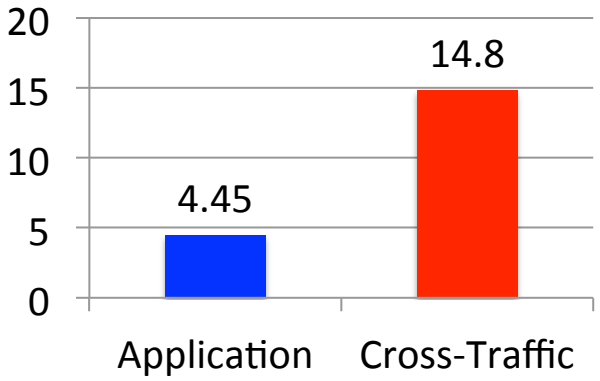

Video runs smooth, GUI applications are responsive

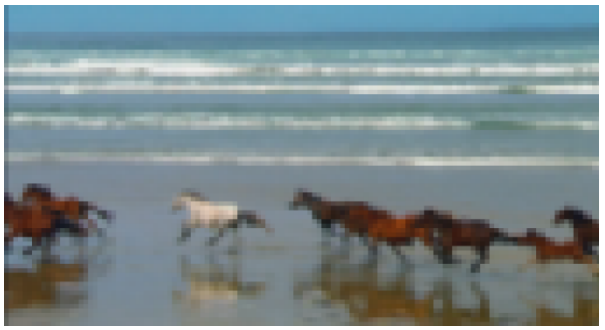

Video freezes, disconnects, GUI applications are not responsive

Video runs smooth, GUI applications are responsive

Thank you for your attention!

References:

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://faculty.missouri.edu/calyamp/publications/fi-ontimemeasure-vdcloud_jcn11.pdf

P. Calyam, S. Rajagopalan, S. Seetharam, A. Selvadhurai, K. Salah, R. Ramnath, "VDC-Analyst: Design and Verification of Virtual Desktop Cloud Resource Allocations", *Elsevier Computer Networks (COMNET)*, 2014.
http://faculty.missouri.edu/calyamp/publications/vdc-analyst-comnet14.pdf