Experimental Research On The iGENI/International OpenFlow Testbed

Joe Mambretti, Director, (j-mambretti@northwestern.edu)
International Center for Advanced Internet Research (www.icair.org)
Northwestern University
Director, Metropolitan Research and Education Network (www.mren.org)
Co-Director, StarLight, PI-iGENI, PI-OMNINet (www.startap.net/starlight)

LHCOPN and LHCONE Joint Meeting
Paris, France
June 17-18, 2013

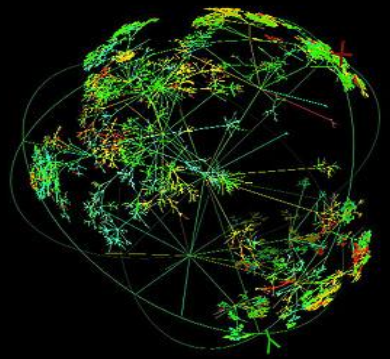

iCAIR

STARLIGHT SM

iGENI and GLIF

- iGENI Consortium Members Have Partnered with Many Other Participants of the Global Lambda Integrated Facility (GLIF) To Undertake Multiple Experimental Network Research Projects
- The iGENI Initiative Has Been Building On That Experience To Create and Exploring New Prototypes of Innovative Communication Services and Technologies.

StarLight: Founding Partner of the Global Lambda Integrated Facility
Available Advanced Network Resources

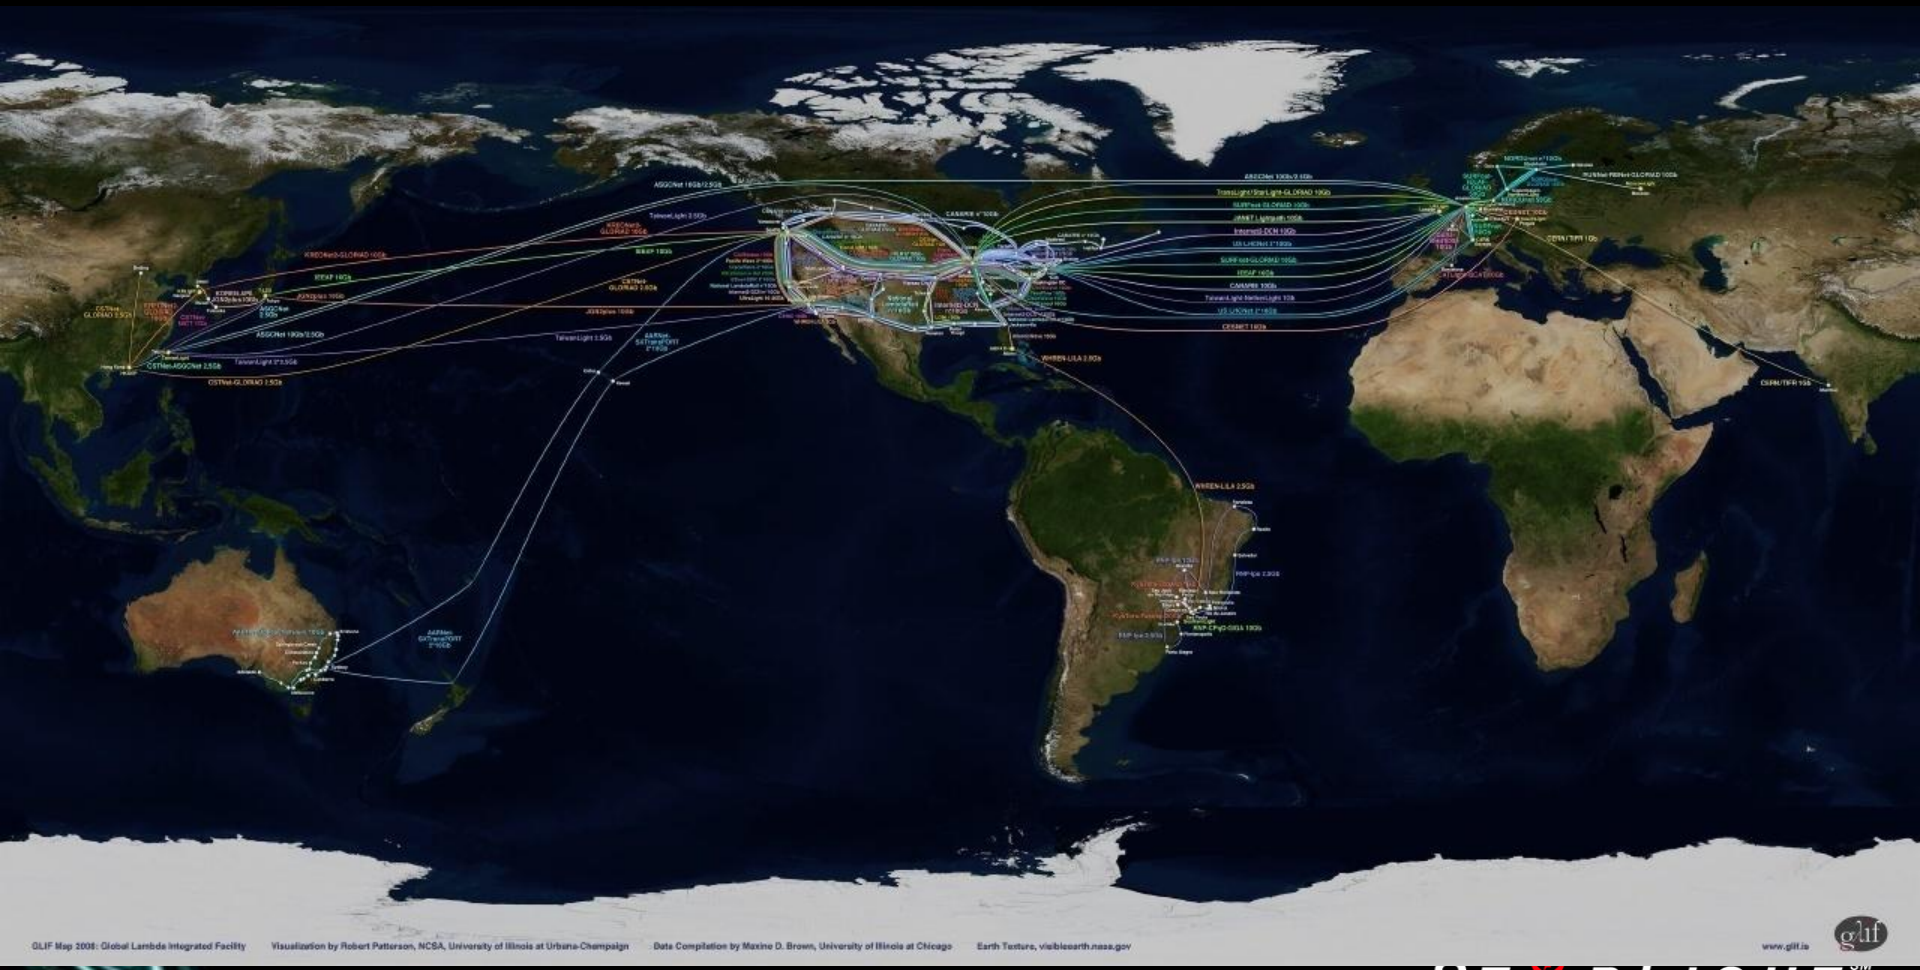


KUAS
NetFPGA
NOX
OpenFlow
NCKU
NetFPGA
NOX
OpenFlow
NOX
OpenFlow
NetFPGA
ETRI
NOX
OpenFlow
NetFPGA
KiSTi
University of Victoria
GENICloud
OpenFlow
InstaGENI
CRC
NOX
OpenFlow
NetFPGA
NLR
light the future
OpenFlow
JGN-X
東京大学
THE UNIVERSITY OF TOKYO
VNode
OpenFlow
cenic
PACIFIC
NORTHWEST
GIGAPOP
aarnet
NICTA
ExoGENI
OpenFlow
NCHC
NetFPGA
NOX
OpenFlow
KRLight
canarie
APPLIED RESEARCH CENTER FOR COMPUTER NETWORKS
OpenFlow
GLORIAD
SingLight
SingAREN
OpenFlow
Northern Wave
hp
GENICloud
renci
ExoGENI
OpenFlow
STARLIGHT
The Optical STAR TAP
MREN
OpenFlow
NLR
light the future
OpenFlow
BEN
Breakable Experimental Network
Duke
UNIVERSITY
ExoGENI
OpenFlow
CNC
OpenFlow
NORTHWESTERN
UNIVERSITY
NetFPGA
NOX
OpenFlow
GENICloud
InstaGENI
AMPATH
OpenFlow
INTERNET2
uen
UTAH EDUCATION NETWORK
OpenFlow
CPqD
OpenFlow
RNP
NLLight
NOX
OpenFlow
NetFPGA
THE UNIVERSITY OF UTAH
OpenFlow
InstaGENI
USP
InstaGENI
UNIVERSITY OF AMSTERDAM
ExoGENI
MAGPI
Princeton University
OpenFlow
InstaGENI
GENICloud
TECHNISCHE UNIVERSITÄT KAISERSLAUTERN
DFN
Deutsches Forschungsnetz
GEANT2
ja.net
University of BRISTOL
OpenFlow
University of Essex
OpenFlow
grnet
Belnet
dedicated connectivity
iMinds
CONNECT.INNOVATE.CREATE
InstaGENI

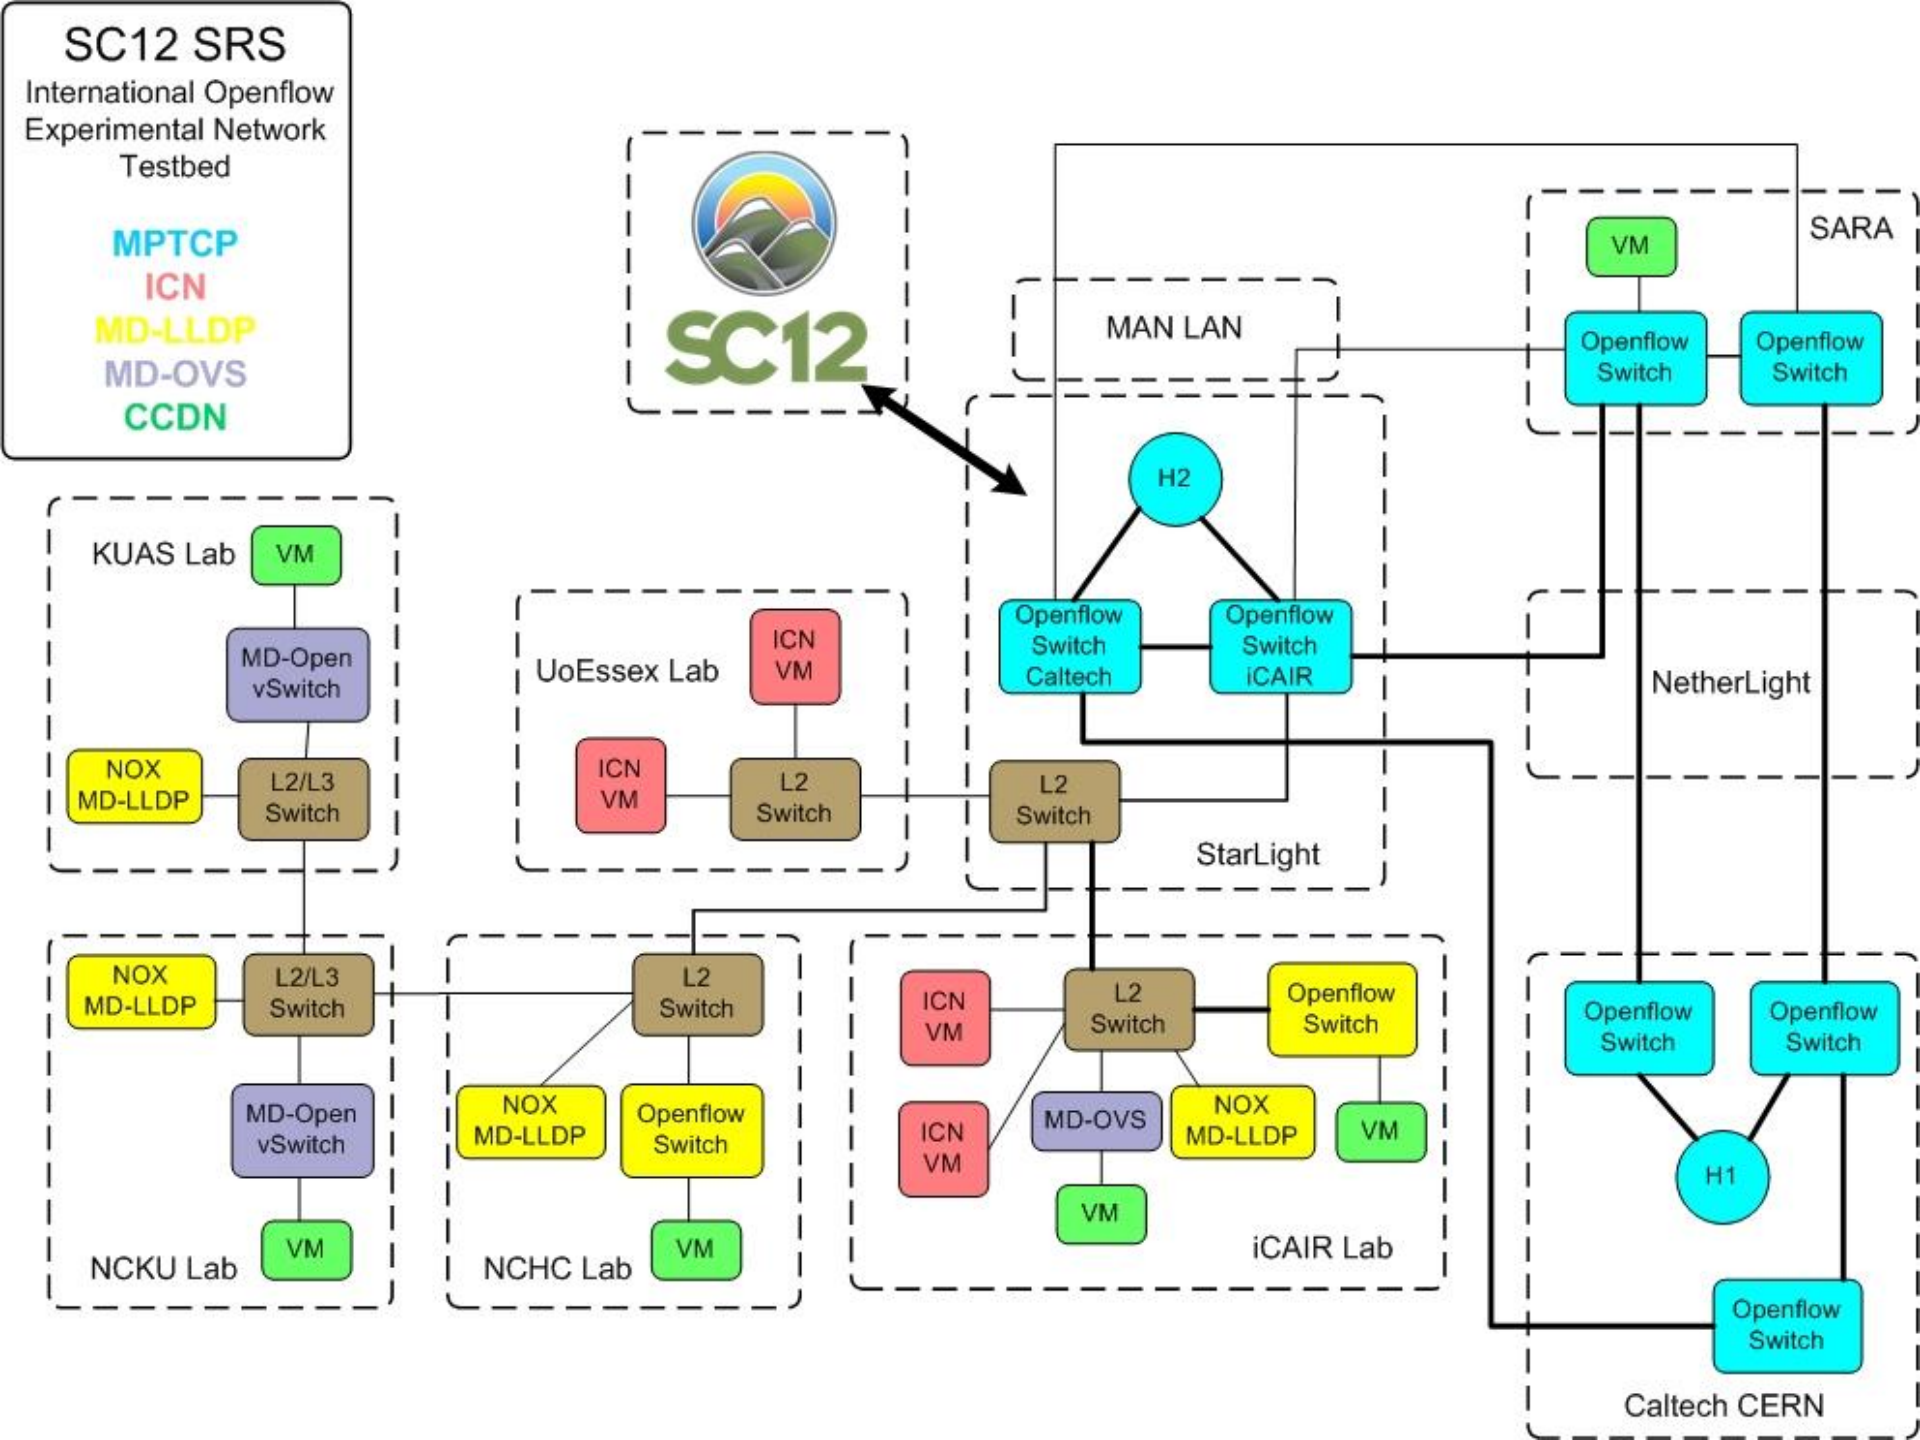
SC12 SRS
International Openflow Experimental Network Testbed
MPTCP
ICN
MD-LLDP
MD-OVS
CCDN
SC12
MAN LAN
SARA
VM
Openflow Switch
Openflow Switch
H2
Openflow Switch Caltech
Openflow Switch iCAIR
NetherLight
KUAS Lab
VM
MD-Open vSwitch
NOX MD-LLDP
L2/L3 Switch
UoEssex Lab
ICN VM
ICN VM
L2 Switch
L2 Switch
StarLight
NOX MD-LLDP
L2/L3 Switch
MD-Open vSwitch
VM
NCKU Lab
L2 Switch
NOX MD-LLDP
Openflow Switch
VM
NCHC Lab
ICN VM
ICN VM
L2 Switch
Openflow Switch
MD-OVS
NOX MD-LLDP
VM
VM
iCAIR Lab
Openflow Switch
Openflow Switch
H1
Openflow Switch
Caltech CERN

Multiple Experiments Share A Persistence International Openflow Testbed

– **MPTCP:** Multi-Path TCP

Lead: SARA, Surfnet, Caltech, CERN

– **ICN:** Information Centric Networking

Lead: University of Essex

– **MD-LLDP:** Multi-Domain Openflow Topology Discovery and Management with LLDP

Lead: NCHC, Taiwan

– **ML-OVS:** Multi-Layers Open vSwitch networking

Lead: KUAS, NCKU, Taiwan

– **CCDN:** Content Centric Distributed Network

Lead: iCAIR

Slice Using Forwarding Rules Based On Payload Content

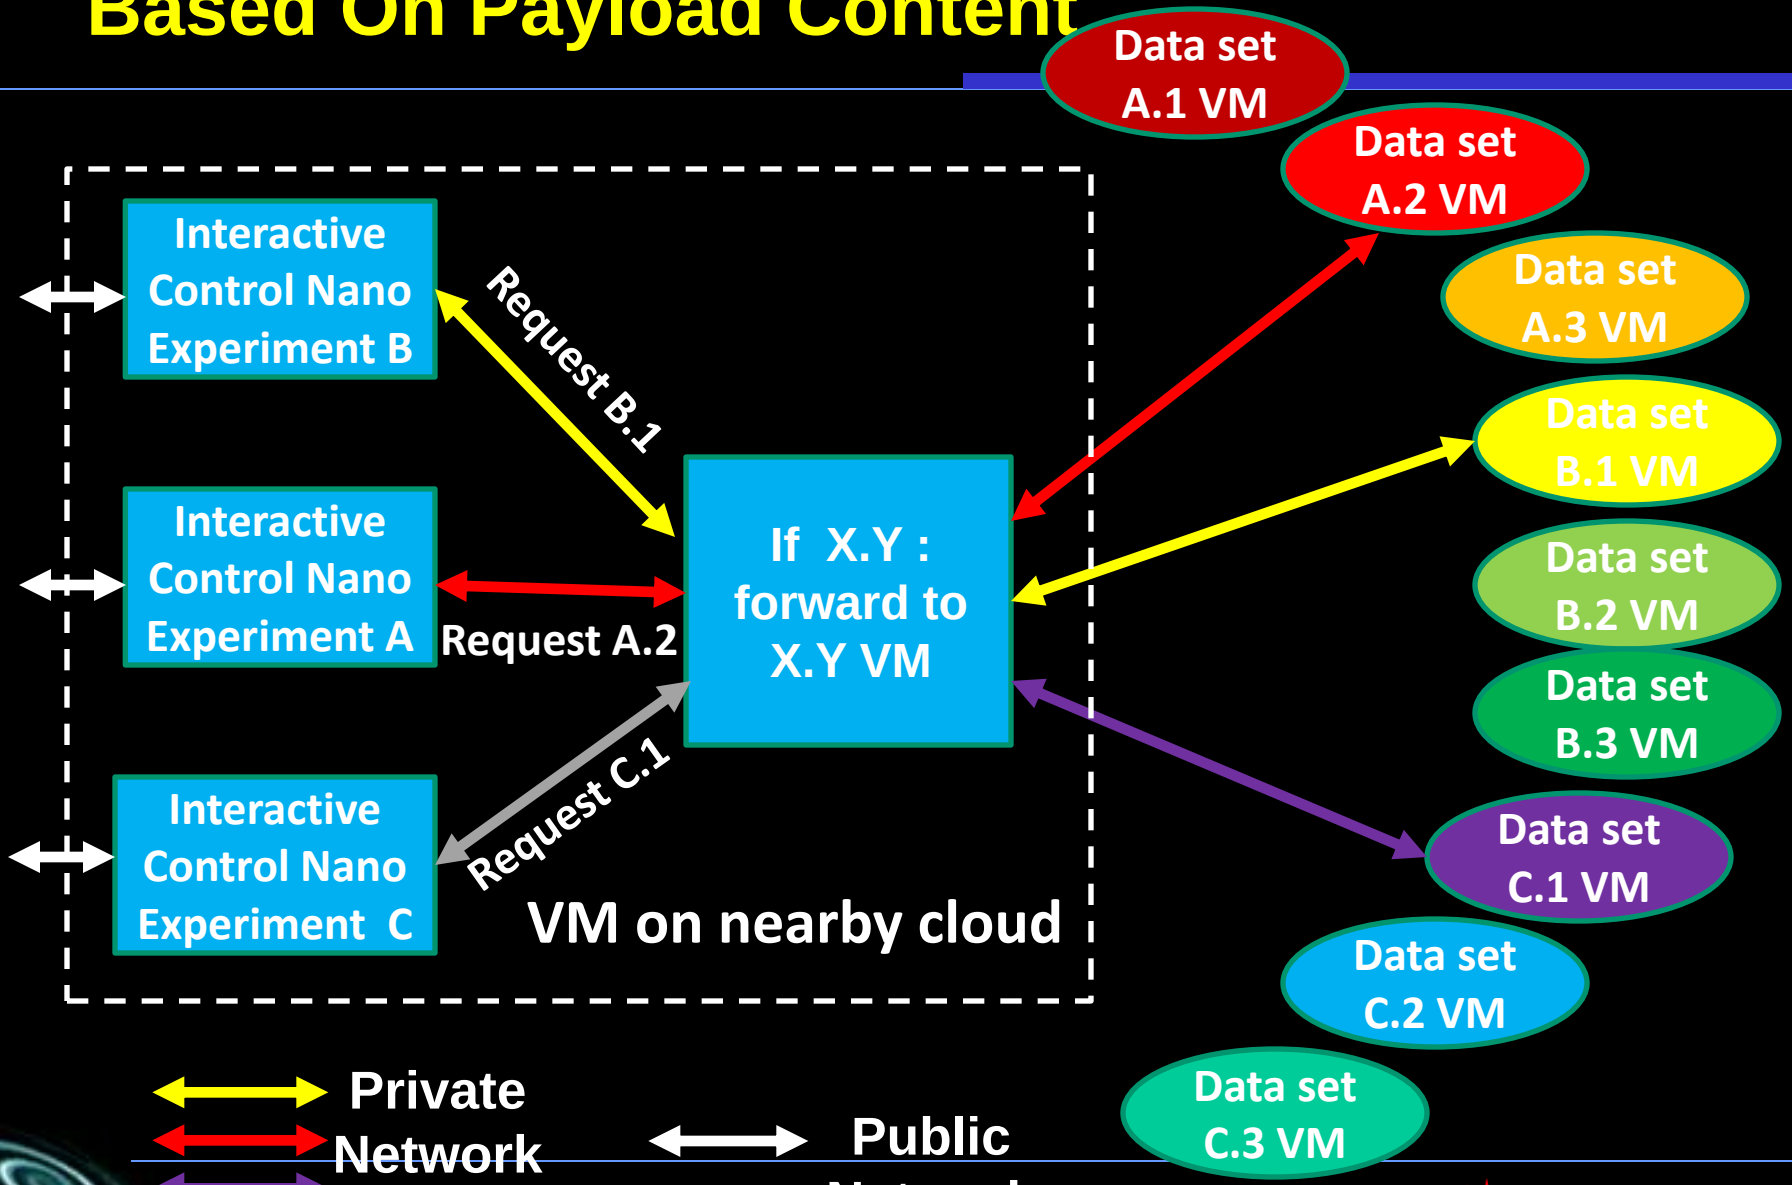

Photonic Band Gap

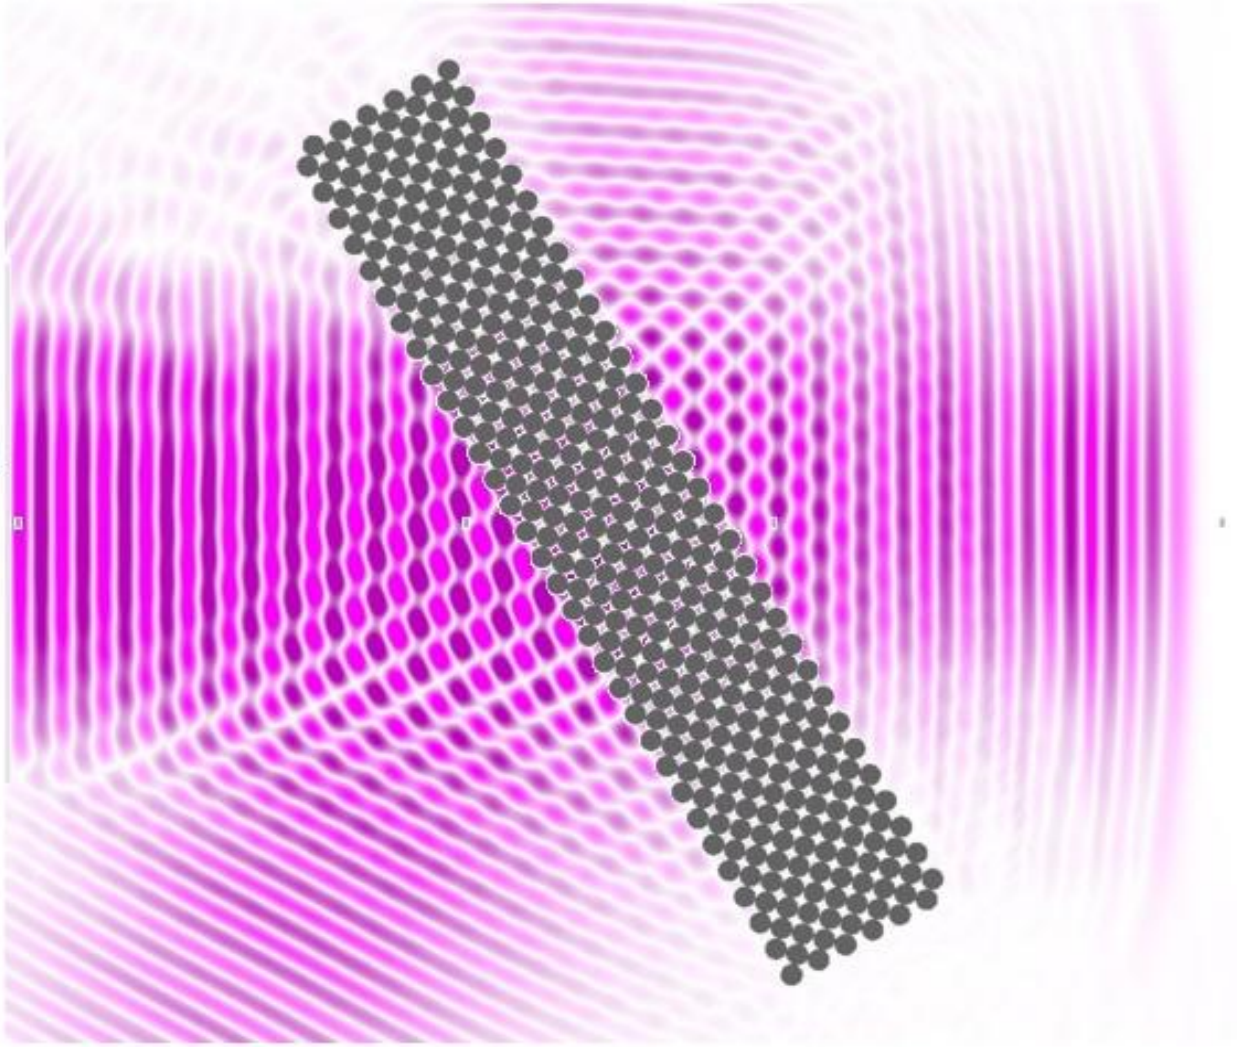

← 11140 nanometers →

Click the picture to zoom in (picture will appear in a new window)

1. Choose the color of light source.
Wavelength appears in nm

Violet (380 **nm)**

300 nm 980

Magnitude

\- 0 +

2. Choose particle size in nanometers

particle size is 200 **nm**

200 nm 380

3. Choose incident angle in degrees

Incident angle is 30 °

0 ° 45

Preview

1 **Frame # 13 , time is 39 fs** 15

4. Watch the animation

> ||

Information Centric Network

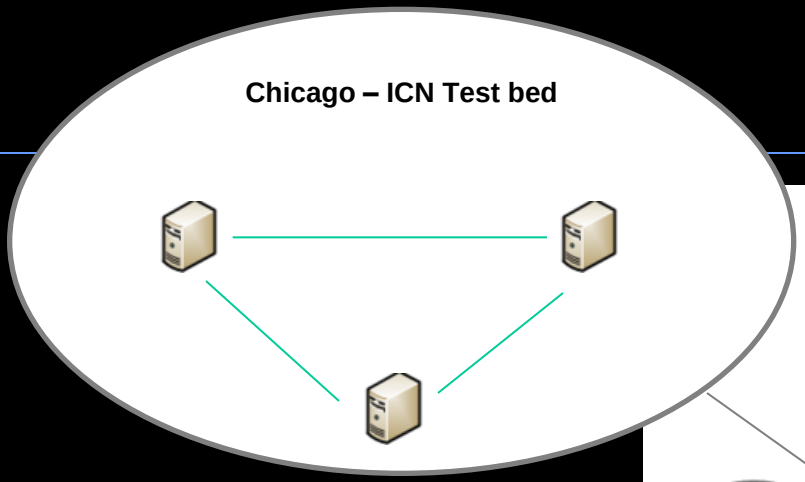

Beyond IP -- Content Routed Networking

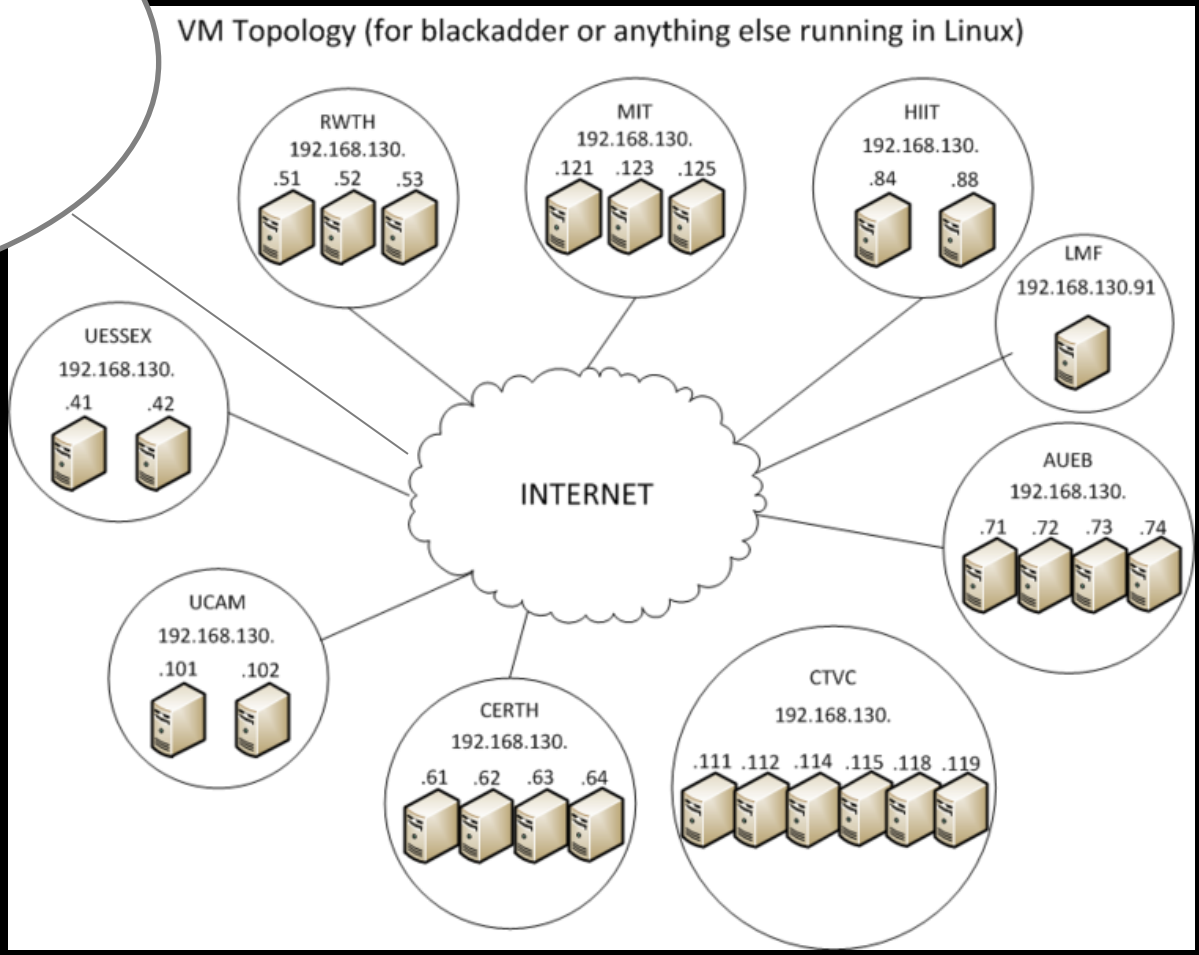

Slide Provide by University of Essex

Paper on This Experiment/Demonstrated Accepted By SC Proceedings

Multipath TCP streaming from Geneva to Chicago over OpenFlow controlled paths

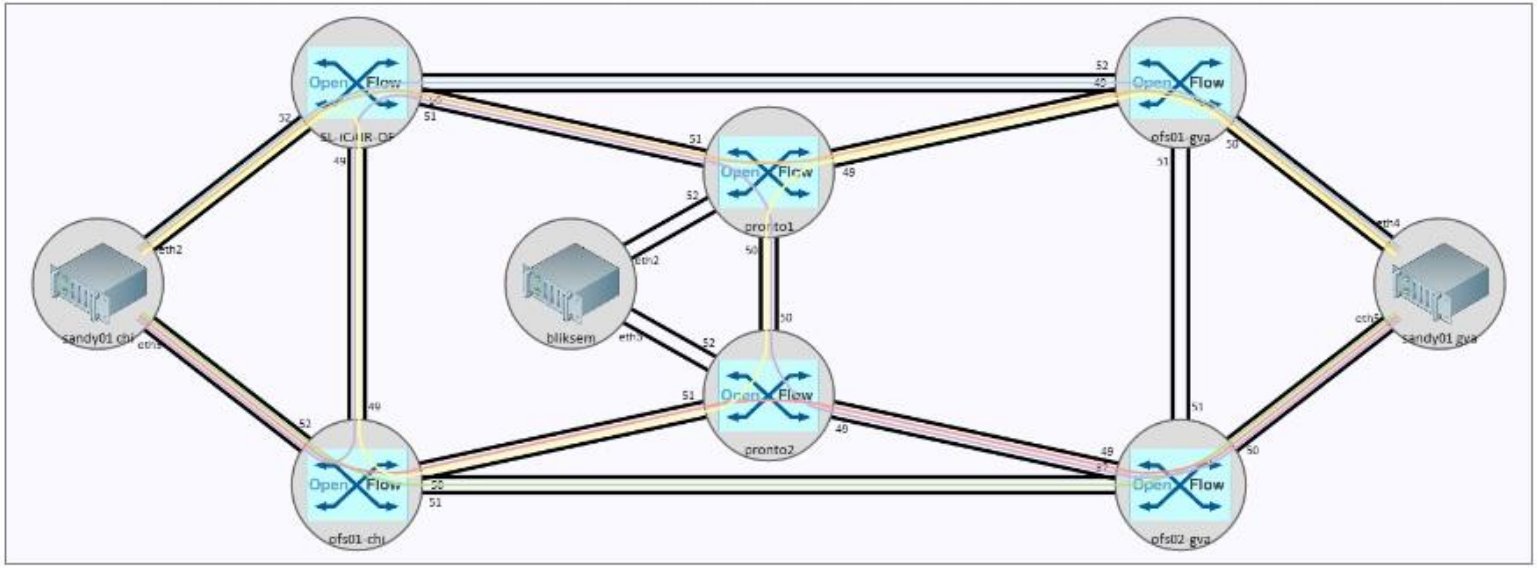

Demo partners:

iCAIR

sara

SURFNET

Current Traffic Rate

Name	Bandwidth
f1 Gva-Chi	0.00 b/s
f2 Gva-Chi	20.22 kb/s
f3 Gva-Chi	0.00 b/s
f4 Gva-Chi	0.00 b/s
f5 Gva-Chi	0.00 b/s
f6 Gva-Chi	0.00 b/s
Total	20.22 kb/s

Reset to defaults

Start Measurements

Stop Measurements

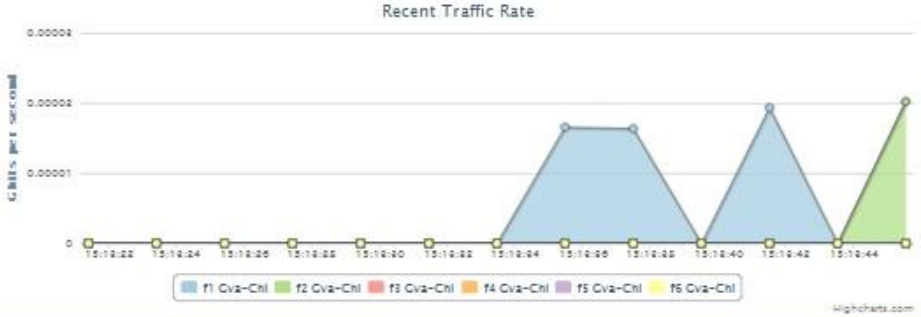

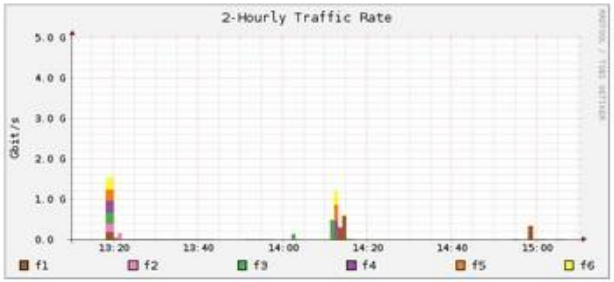

ow f1 Gva-Ams ends at [Port 50 of 00:00:e8:9a:8f:d2:80:57 (ofs01-gva)] (no flow entry for IP 10.0.200.2)
ow f2 Gva-Ams ends at [Port 50 of 00:00:60:eb:69:fe:49:14 (ofs02-gva)] (no flow entry for IP 10.0.201.2)
ow f3 Gva-Ams ends at [Port 50 of 00:00:60:eb:69:fe:49:14 (ofs02-gva)] (no flow entry for IP 10.0.202.2)

Slide Provide by Caltech,SARA, Surfnet

STARLIGHT SM

Inter-Domain Openflow
Topology Discovery & Monitoring
Slide Provide By NCHC (Paper on Technique Accepted for Publication)

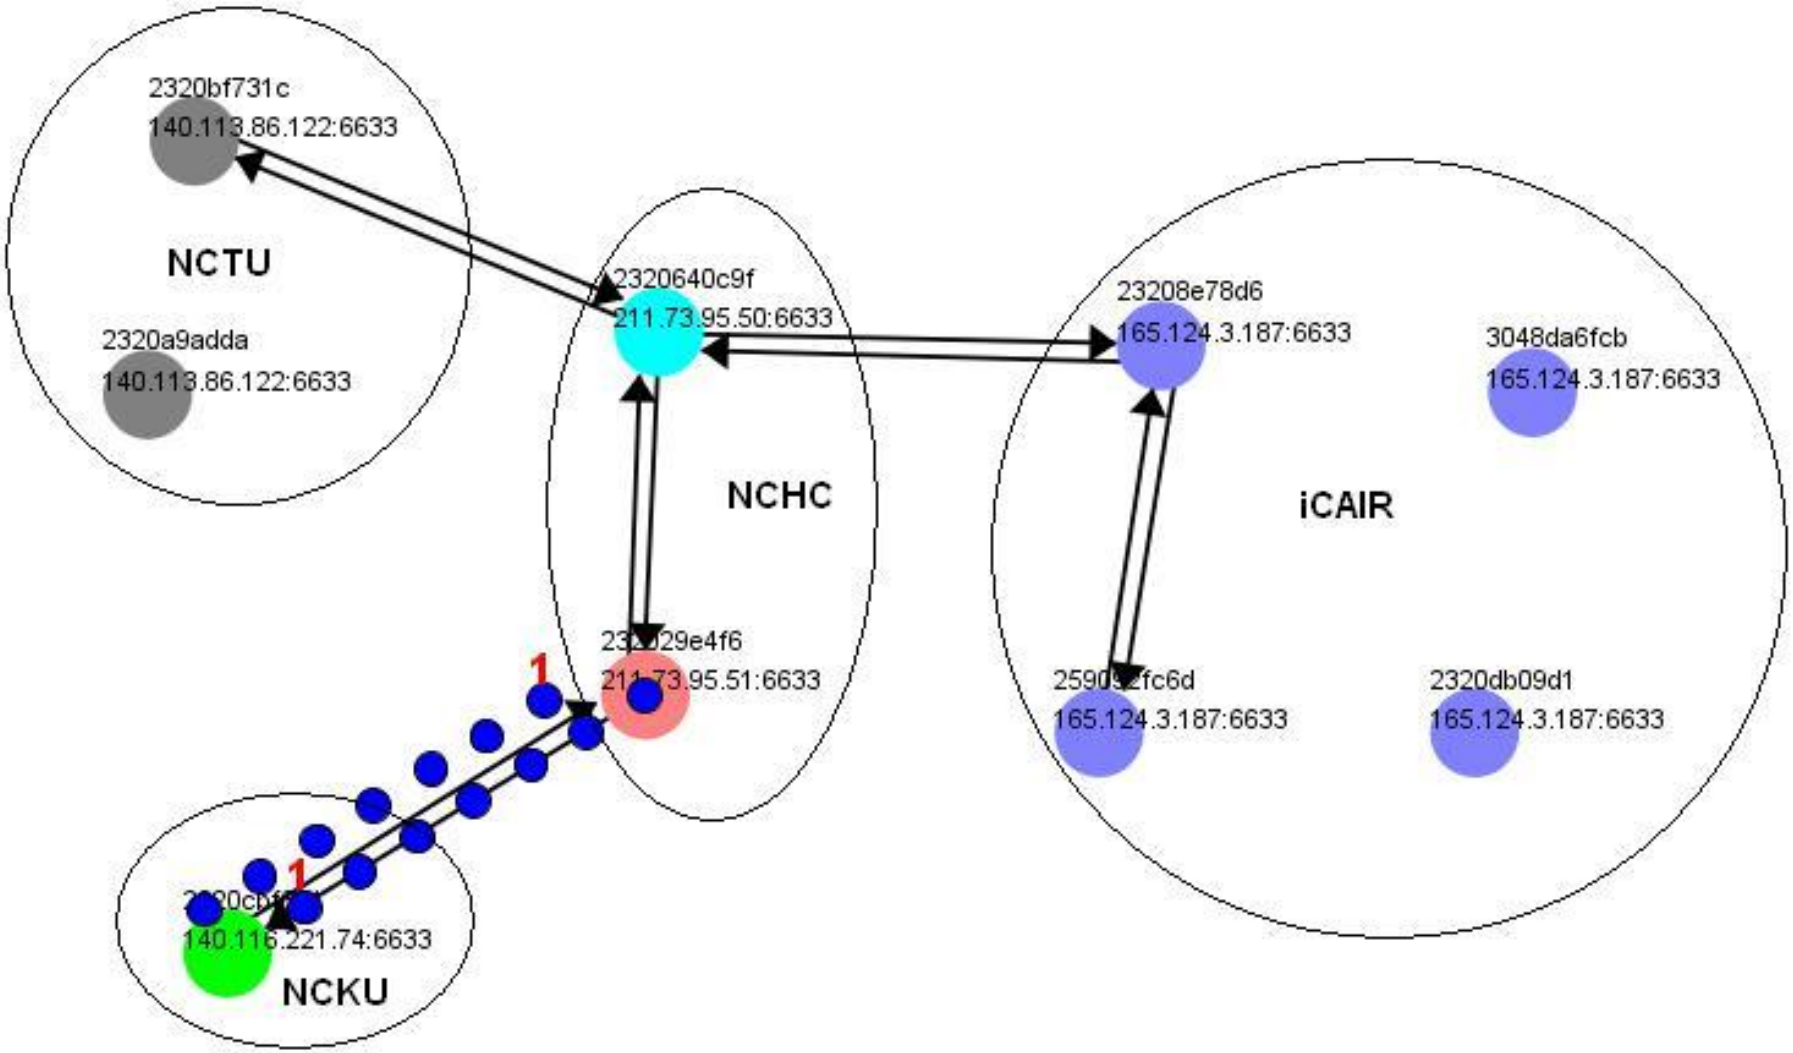

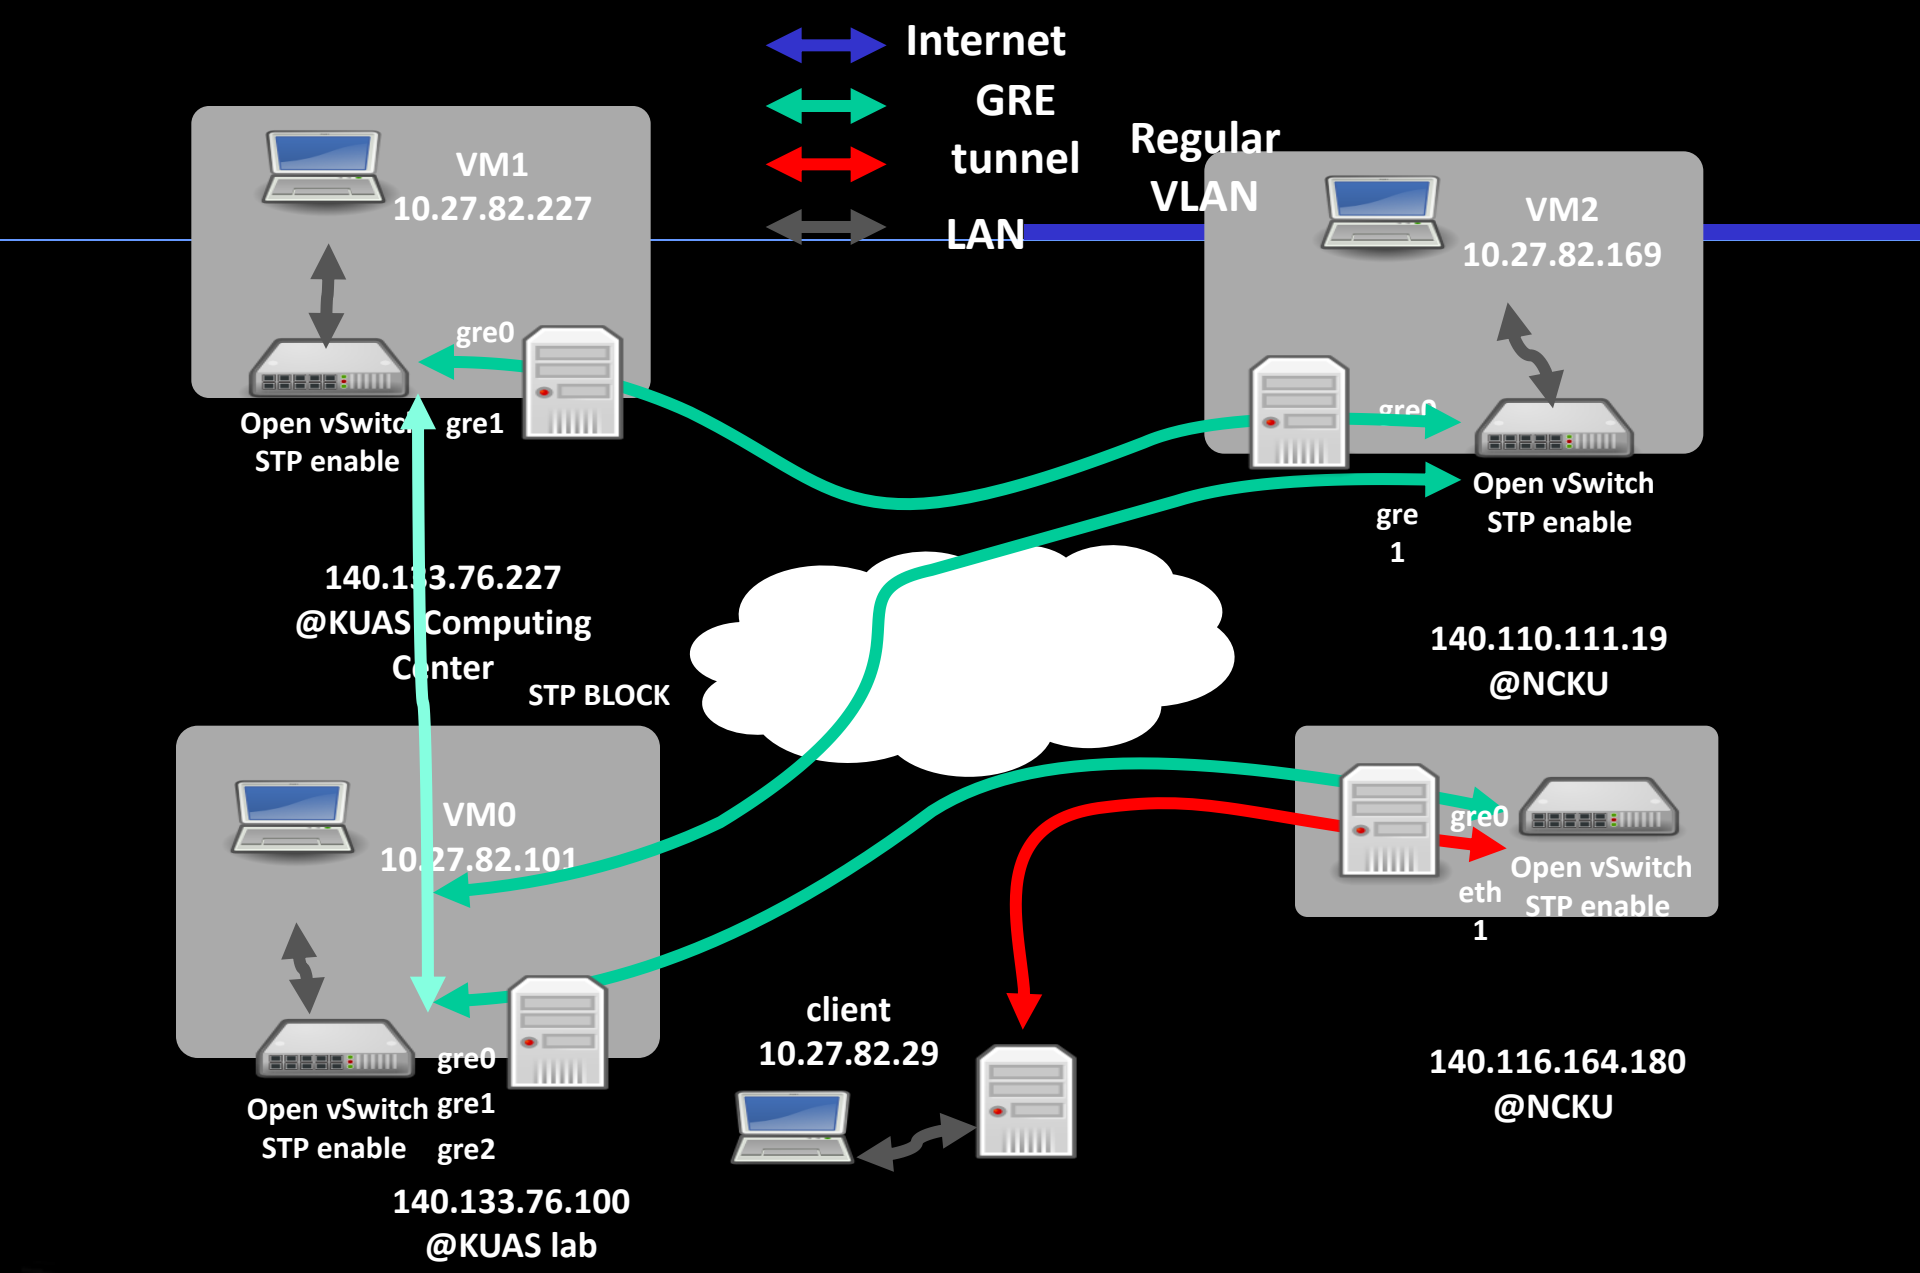

Multi-Layer Openflow OVS Network
Slide Provide by NCKU and KUAS

STARLIGHT℠

TransCloud Experiments and Demonstrations

Alvin AuYoung, Andy Bavier, Jessica Blaine, Jim Chen, Yvonne Coady, Paul Muller, Joe Mambretti, Chris Matthews, Rick McGeer, Chris Pearson, Alex Snoeren, Fei Yeh, Marco Yuen

TransCloud Today

TransCloud: Based on iGENI and GENICloud

- Transcontinental Federation of Cloud Systems
- Slice-Based Federation Architecture for sign on and trans-cluster slice management
- SFA cluster manager at each site
  - Currently, enhanced Eucalyptus
- Private 10 Gb/s transcontinental network linking sites
  - Thanks to GLIF, NLR, NetherLight, CAVEWave, StarLight, DFN

Roadmap

- Accept experimenters ***now***
- Federation expansion
  - TU Amsterdam immediately
  - Brazil, Asia by July
  - All interested parties at any time
- Full integration with PlanetLab Control Framework (July)
- High-level programming environment based on RePy and NaCl
- High-level distributed query environment

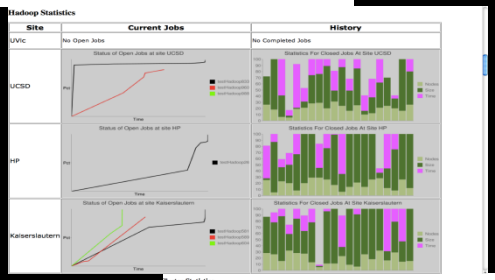

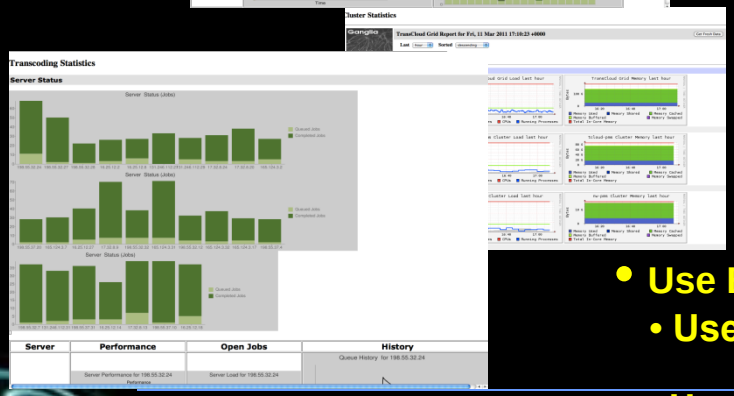

Example of working in the TransCloud

[1] Build trans-continental applications spanning clouds:

- **Distributed query application based on Hadoop/Pig**
- **Store archived Network trace data using HDFS**
- **Query data using Pig over Hadoop clusters**

[2] Perform distributed query on TransCloud, which currently spans the following sites:

- **HP OpenCirrus**
- **Northwestern OpenCloud**
- **UC San D**
- **Kaiser**

- **Use By Outside Researchers? Yes**
- **Use Involving Multiple Aggregates? Yes**
- **Use for Research Experiments? Yes**

Also Ref: Experiments in High Perf Transport at GEC 7

Demo: http://tcdemo.dyndns.org/

STARLIGHT℠

Digital Media TransCoding Demonstration

- TransCloud: Advanced Distributed Global Environment Enables Dynamic Creation of Communication Services, Including Those Based on Rapid Migration of Virtual Network and Cloud Resources
- TransCloud: Set of Protocols, Standards, Management Software Enables Interoperation of Distinct Cloud and Network Resources
- Example: Dynamic Cloud+Dynamic Network for Digital Media Transcoding Using Single Platform vs Multiple Infrastructures

Transcoding Cloud 1

Transcoding Cloud 2

Video Sources

OpenFlow Switches

Transcoding Cloud 3

STARLIGHT SM

GENI Engineering Conference 9 Washington DC, Nov 2010

Prototype

UCSD GENICloud

HP Labs GENICloud

iCAIR GENICloud

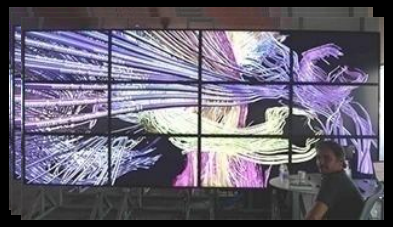

RLIGHT

VirtuLab Tile Display:
Directly Connected To National 10 Gbps Testbed
With Core at the StarLight Facility

Open Resource Control Architecture/ OSCARS (DOE)/ iCAIR Testbed

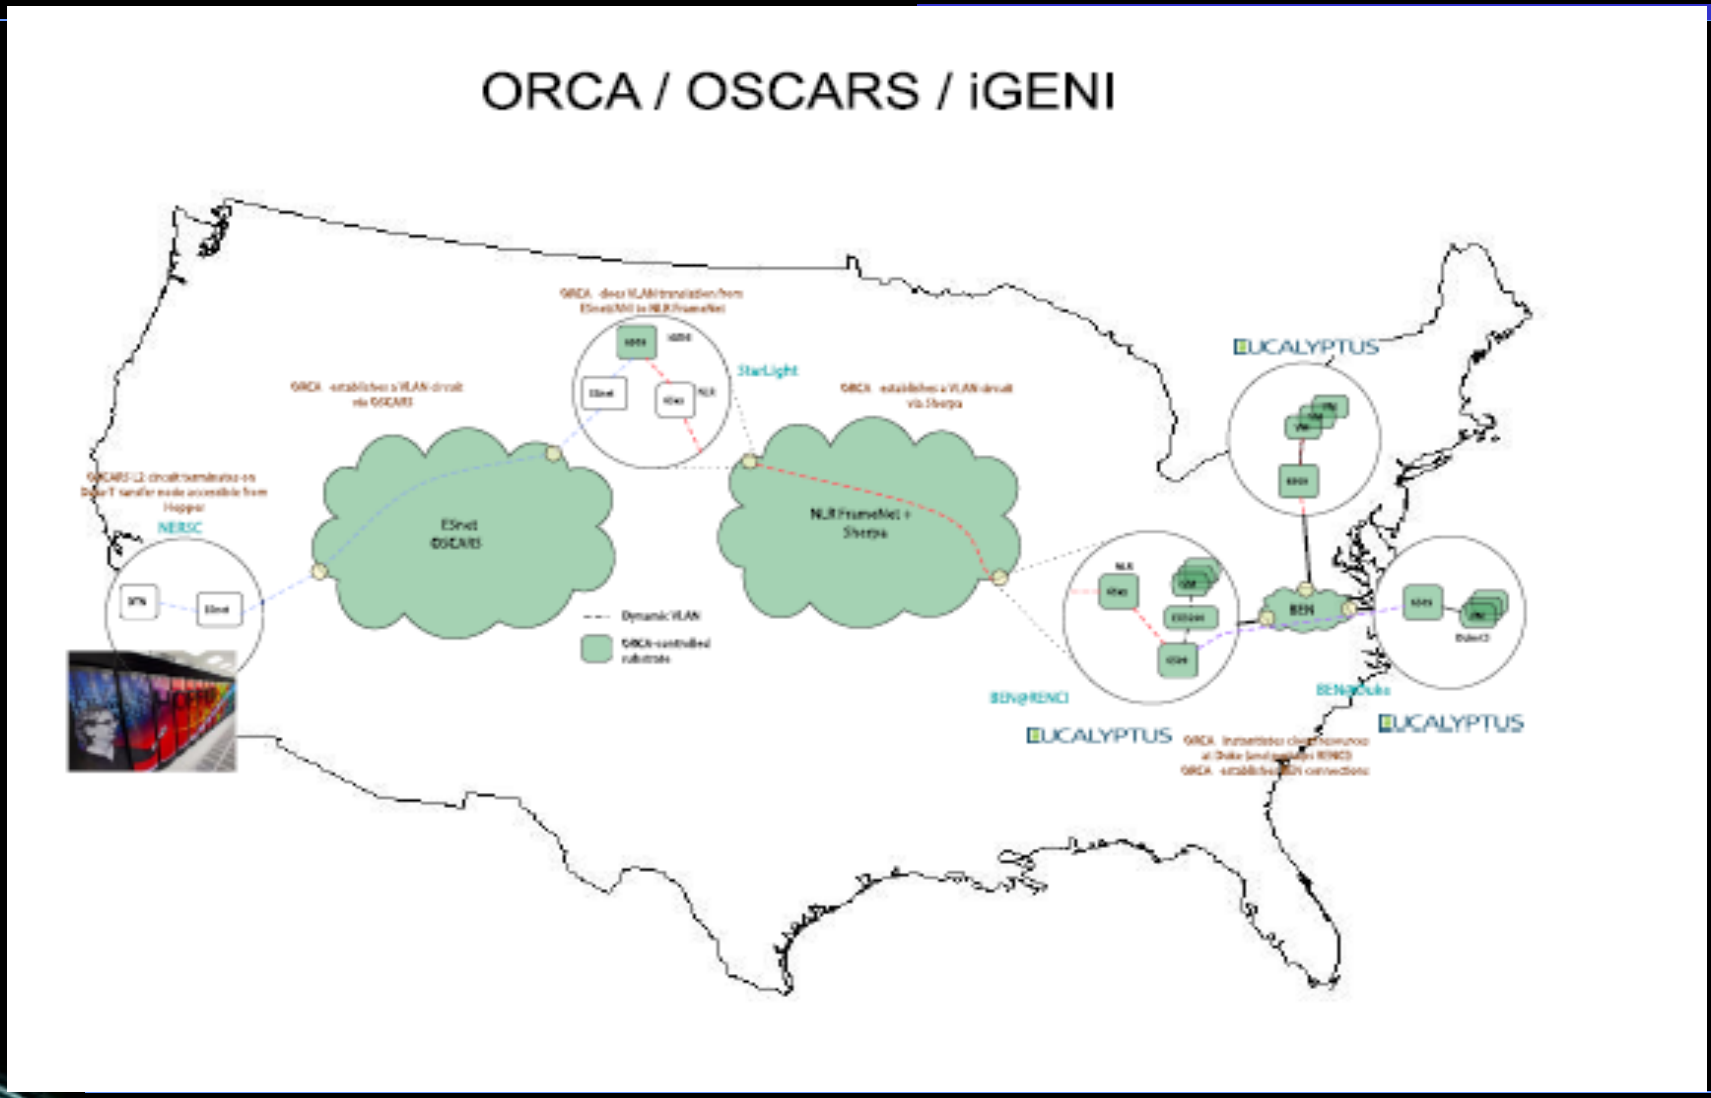

iGENI GIST-BEN-KREONET Testbed

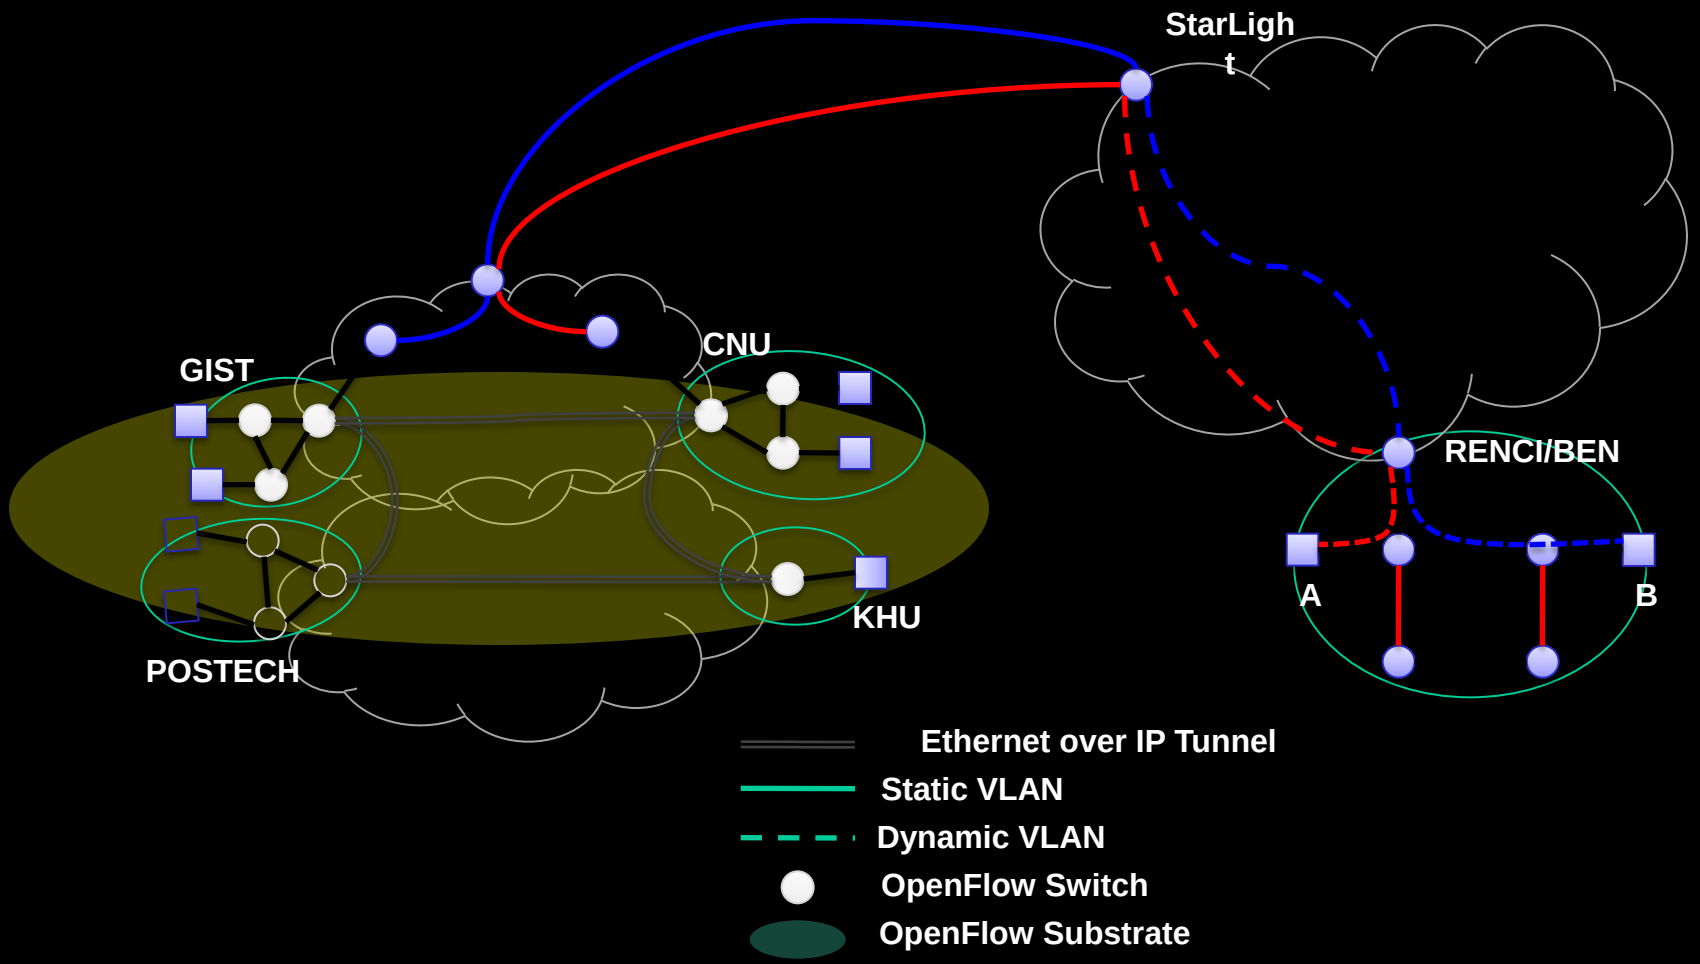

Nowcast Demo Data Flow

- Dynamic end-to-end Nowcasting
  - Mapping Nowcast Workflows onto GENI

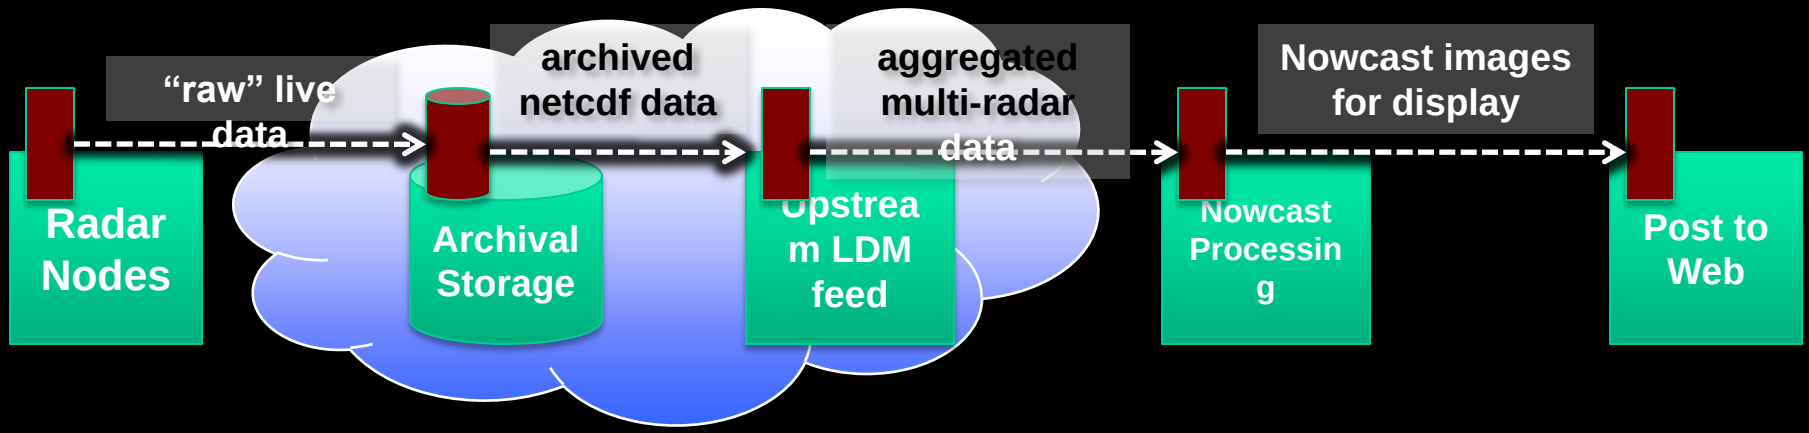

Source: UMass Amherst – Mike Zink

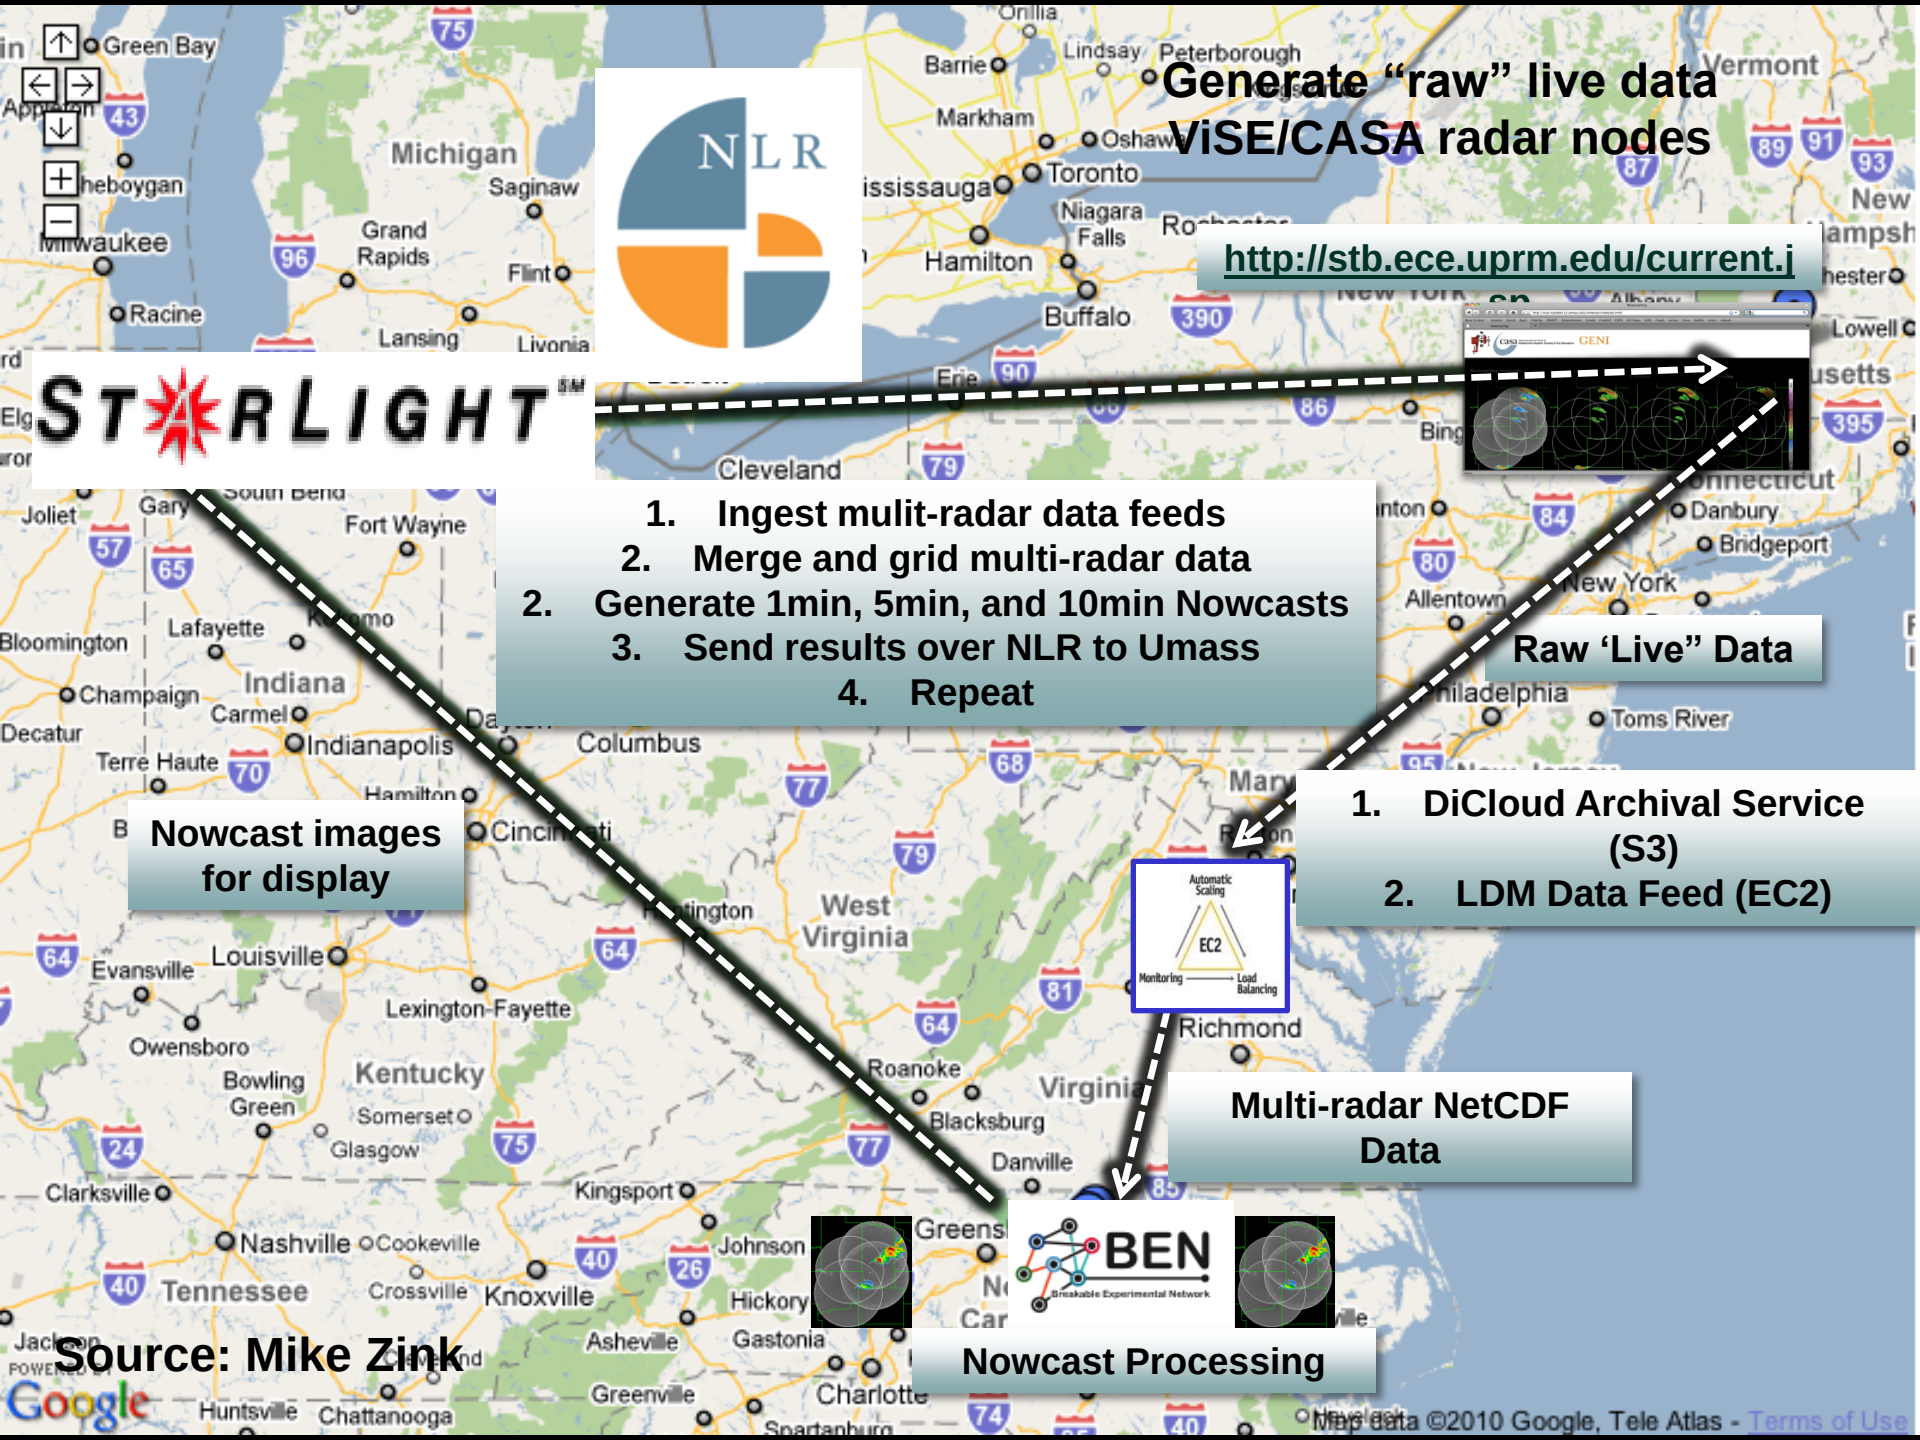
Generate “raw” live data
ViSE/CASA radar nodes
NLR
STARLIGHT
http://stb.ece.uprm.edu/current.j
1. Ingest mulit-radar data feeds
2. Merge and grid multi-radar data
2. Generate 1min, 5min, and 10min Nowcasts
3. Send results over NLR to Umass
4. Repeat
Raw ‘Live” Data
1. DiCloud Archival Service (S3)
2. LDM Data Feed (EC2)
Nowcast images for display
Multi-radar NetCDF Data
BEN
Nowcast Processing
Source: Mike Zink
Google
Map data ©2010 Google, Tele Atlas - Terms of Use

TransGeo: An Open, Distributed, Federated GIS/System Cloud –Rick McGeer Chris Matthews et al

- GIS Data Is Large, Collected By Many Sources, Needed All Over the World
- Use Today Is Mostly Desktop Fat Clients (Quantum GIS, ESRI)
- Many Want to Compute in the Cloud
- Open and Available To Everyone
- Distributed Swift as Federated Store
- Distributed Disco as MapReduce Computation Engine
- Open-Source Standard Tools For Point Computation (GRASS, GDAL)

green.cities.trans-cloud.net:8080/geoserver/www/maps/view.html

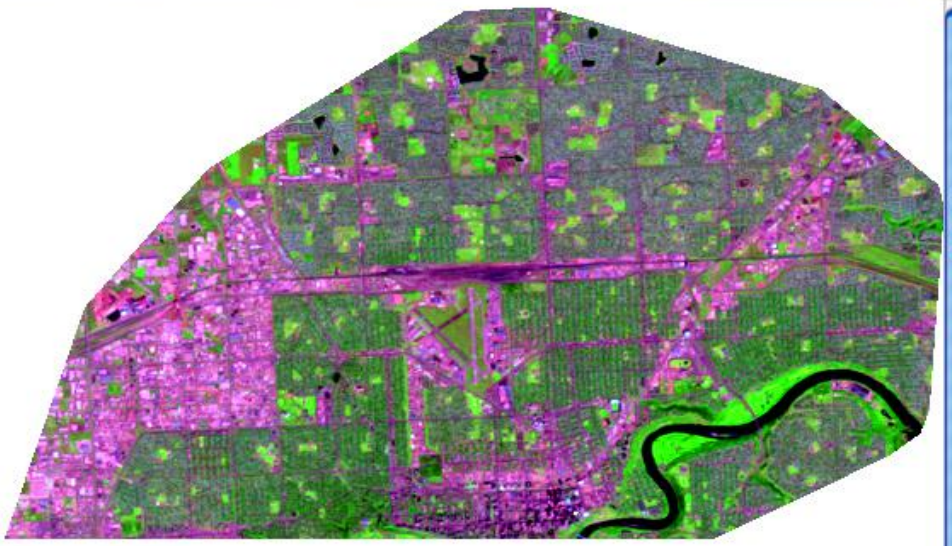

GreenCities Demo

A demonstration of TransGeo, an open-source multisite GIS Cloud developed at the University of Victoria

Big Data jobs are a key use case for federating clouds at multiple sites. In this project, we are developing a prototype open-source system to demonstrate Big Data processing at multiple sites. This system uses a single PostGIS database as a repository for meta-data, and landsat (or other) satellite image data in Swift repositories, and distributed Disco jobs to do the analysis

bing™

Image courtesy of NASA © 2012 Microsoft Corporation

-113.55200, 53.55856

DARPA-BAA-12-53_U....pdf

GRASS-6.4.2-5-Snow.dmg

Show All

STARLIGHT SM

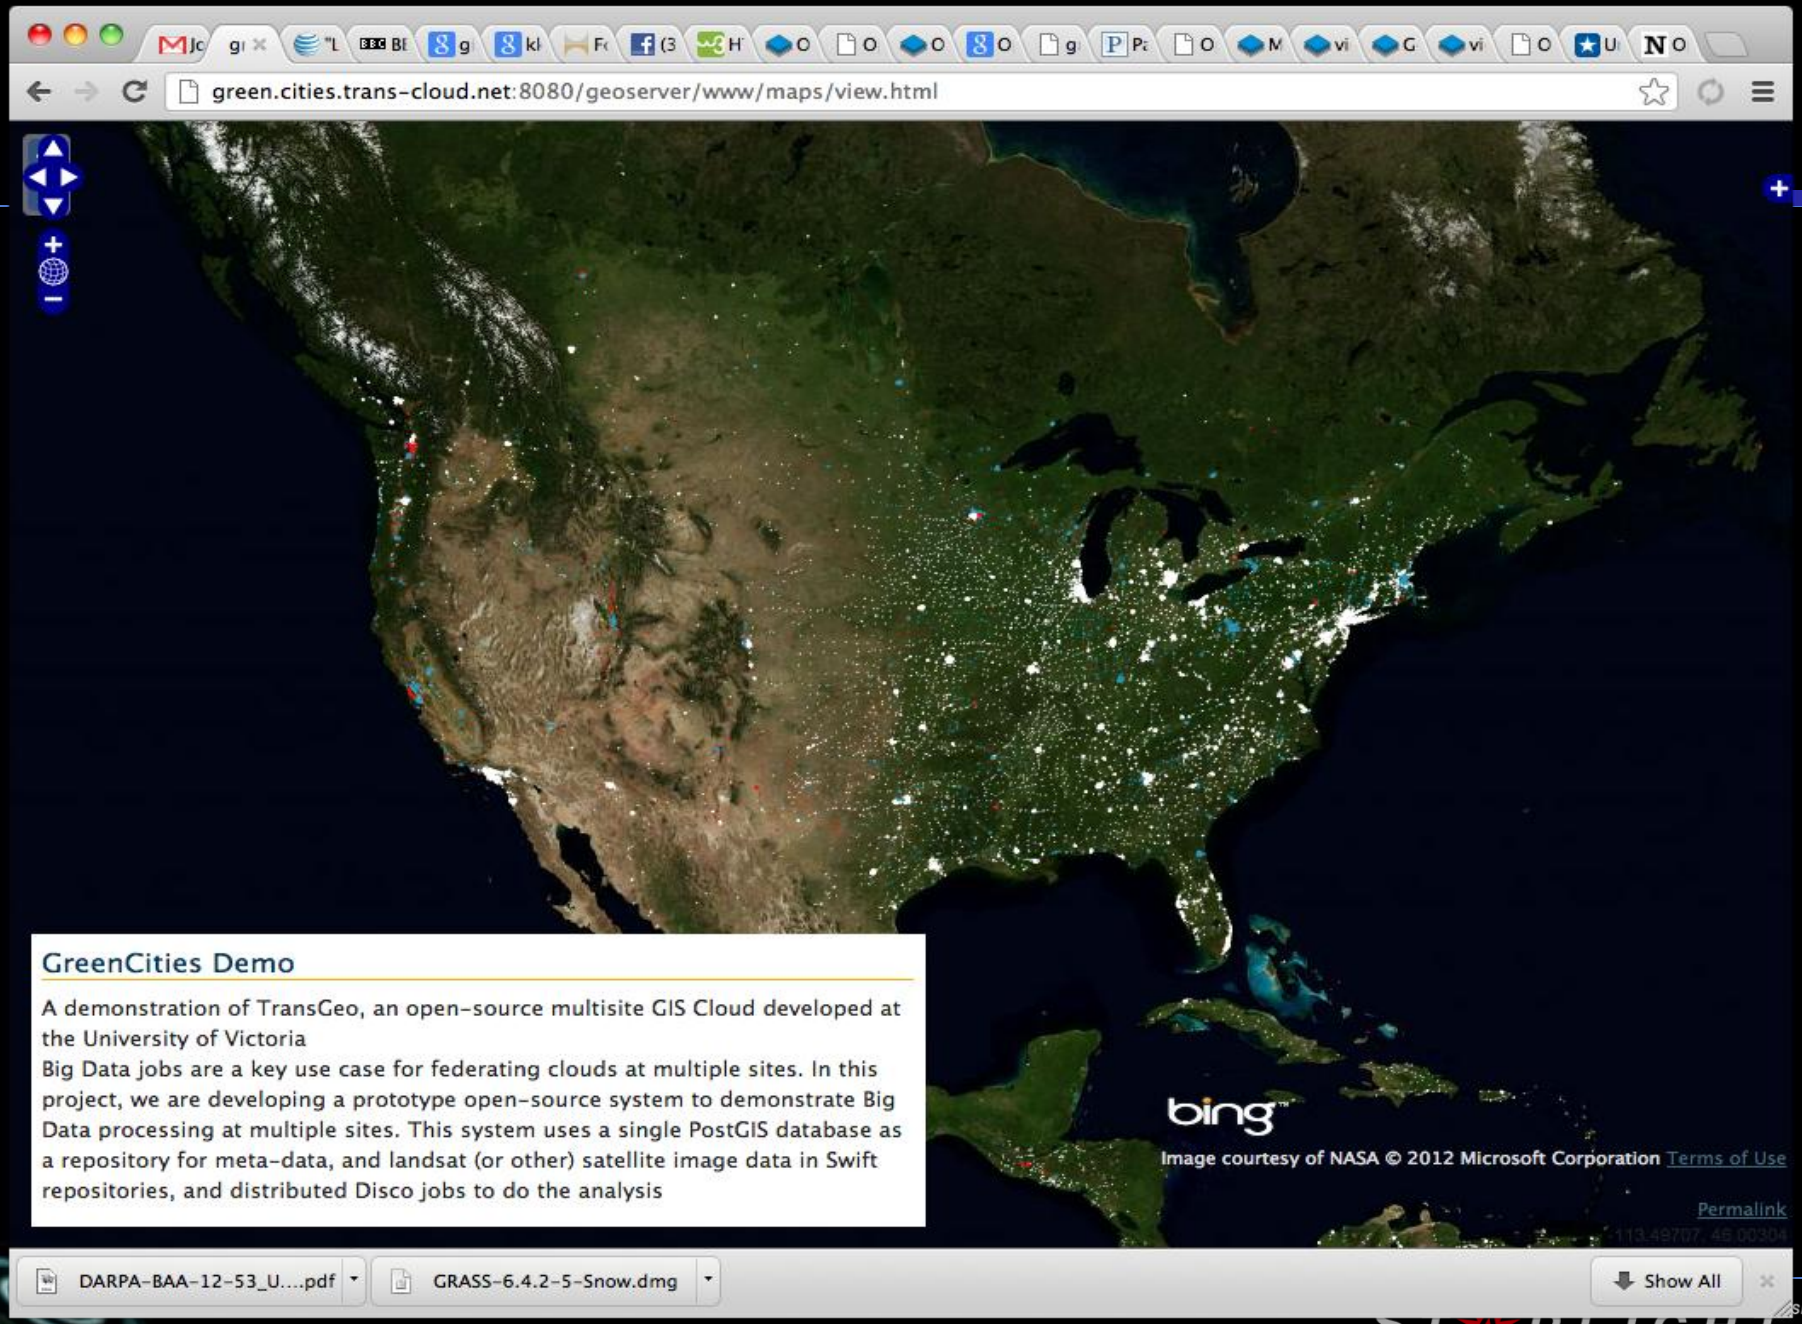
green.cities.trans-cloud.net:8080/geoserver/www/maps/view.html
GreenCities Demo
A demonstration of TransGeo, an open-source multisite GIS Cloud developed at the University of Victoria
Big Data jobs are a key use case for federating clouds at multiple sites. In this project, we are developing a prototype open-source system to demonstrate Big Data processing at multiple sites. This system uses a single PostGIS database as a repository for meta-data, and landsat (or other) satellite image data in Swift repositories, and distributed Disco jobs to do the analysis
bing
Image courtesy of NASA © 2012 Microsoft Corporation Terms of Use
Permalink
DARPA-BAA-12-53_U....pdf
GRASS-6.4.2-5-Snow.dmg
Show All

Aki's Packet Cache and XIA

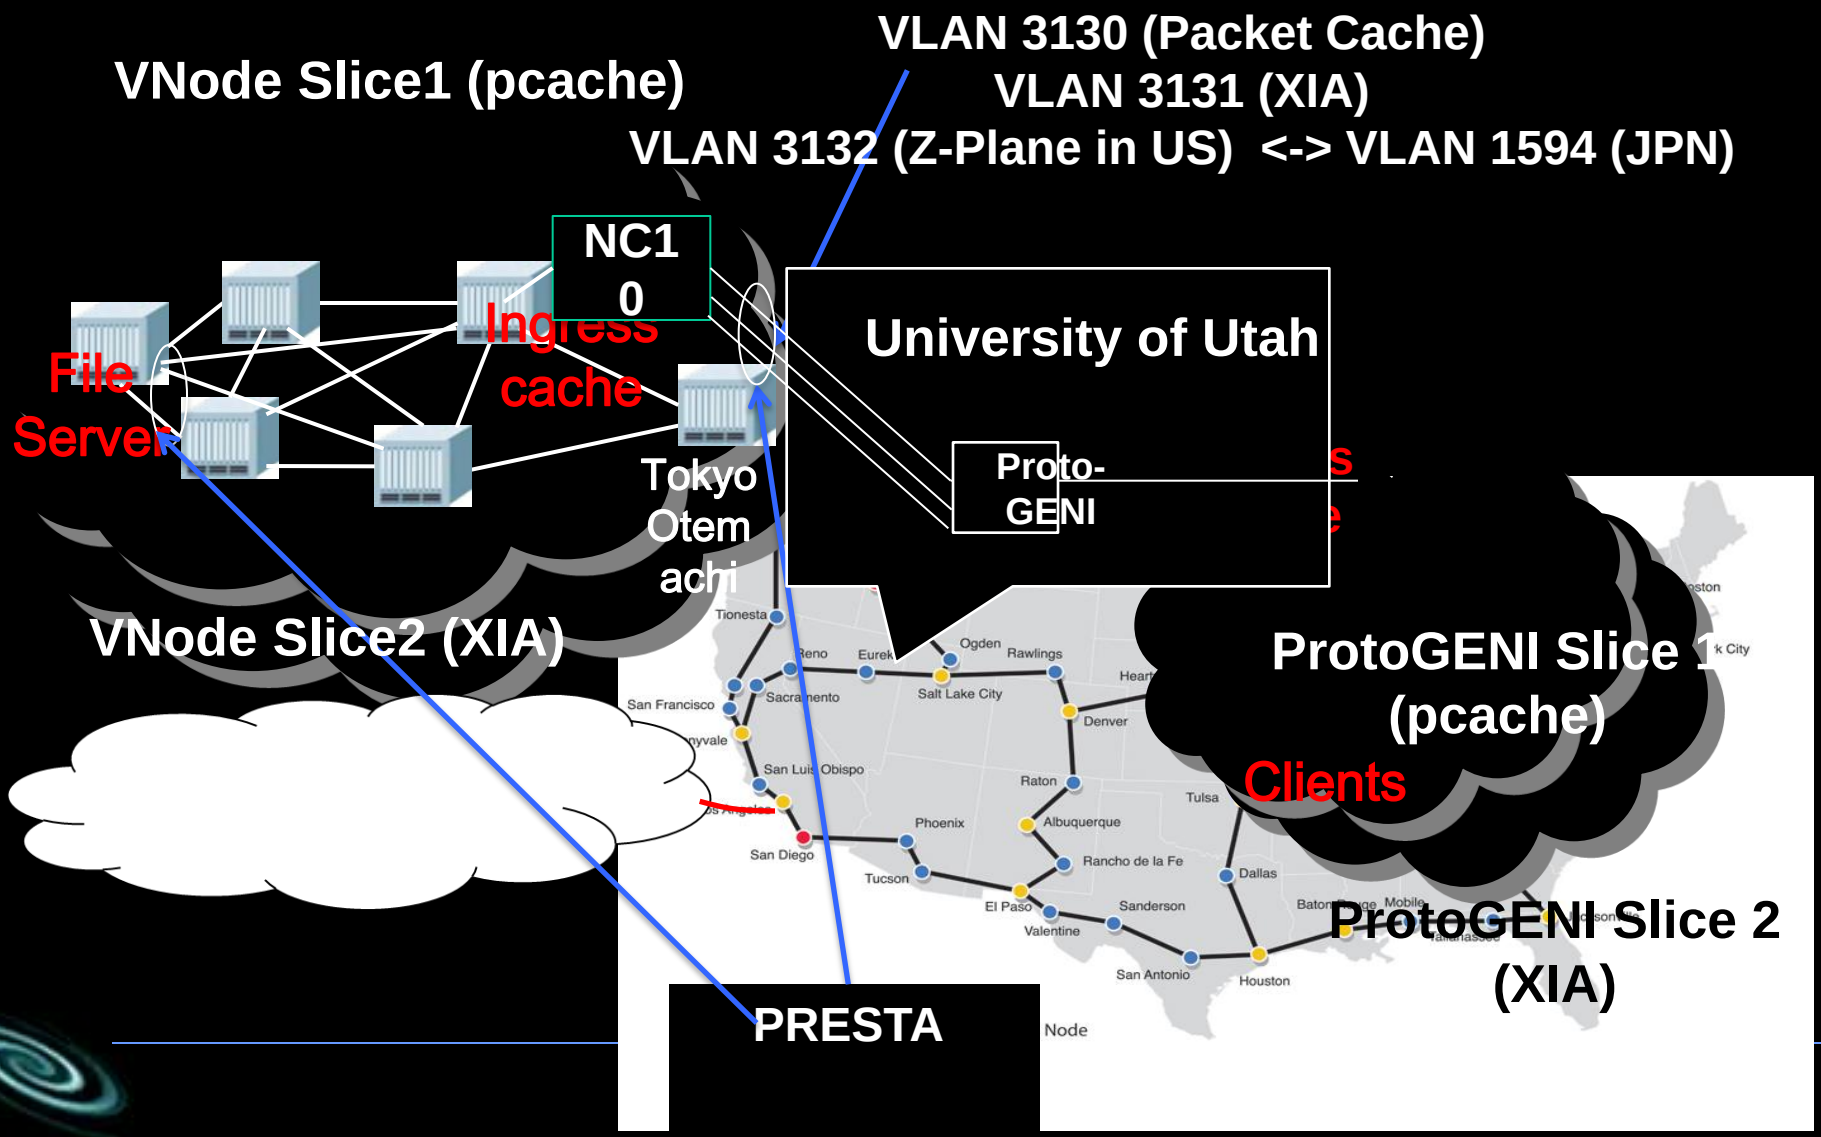

Kansei VLAN Connectivity Diagram

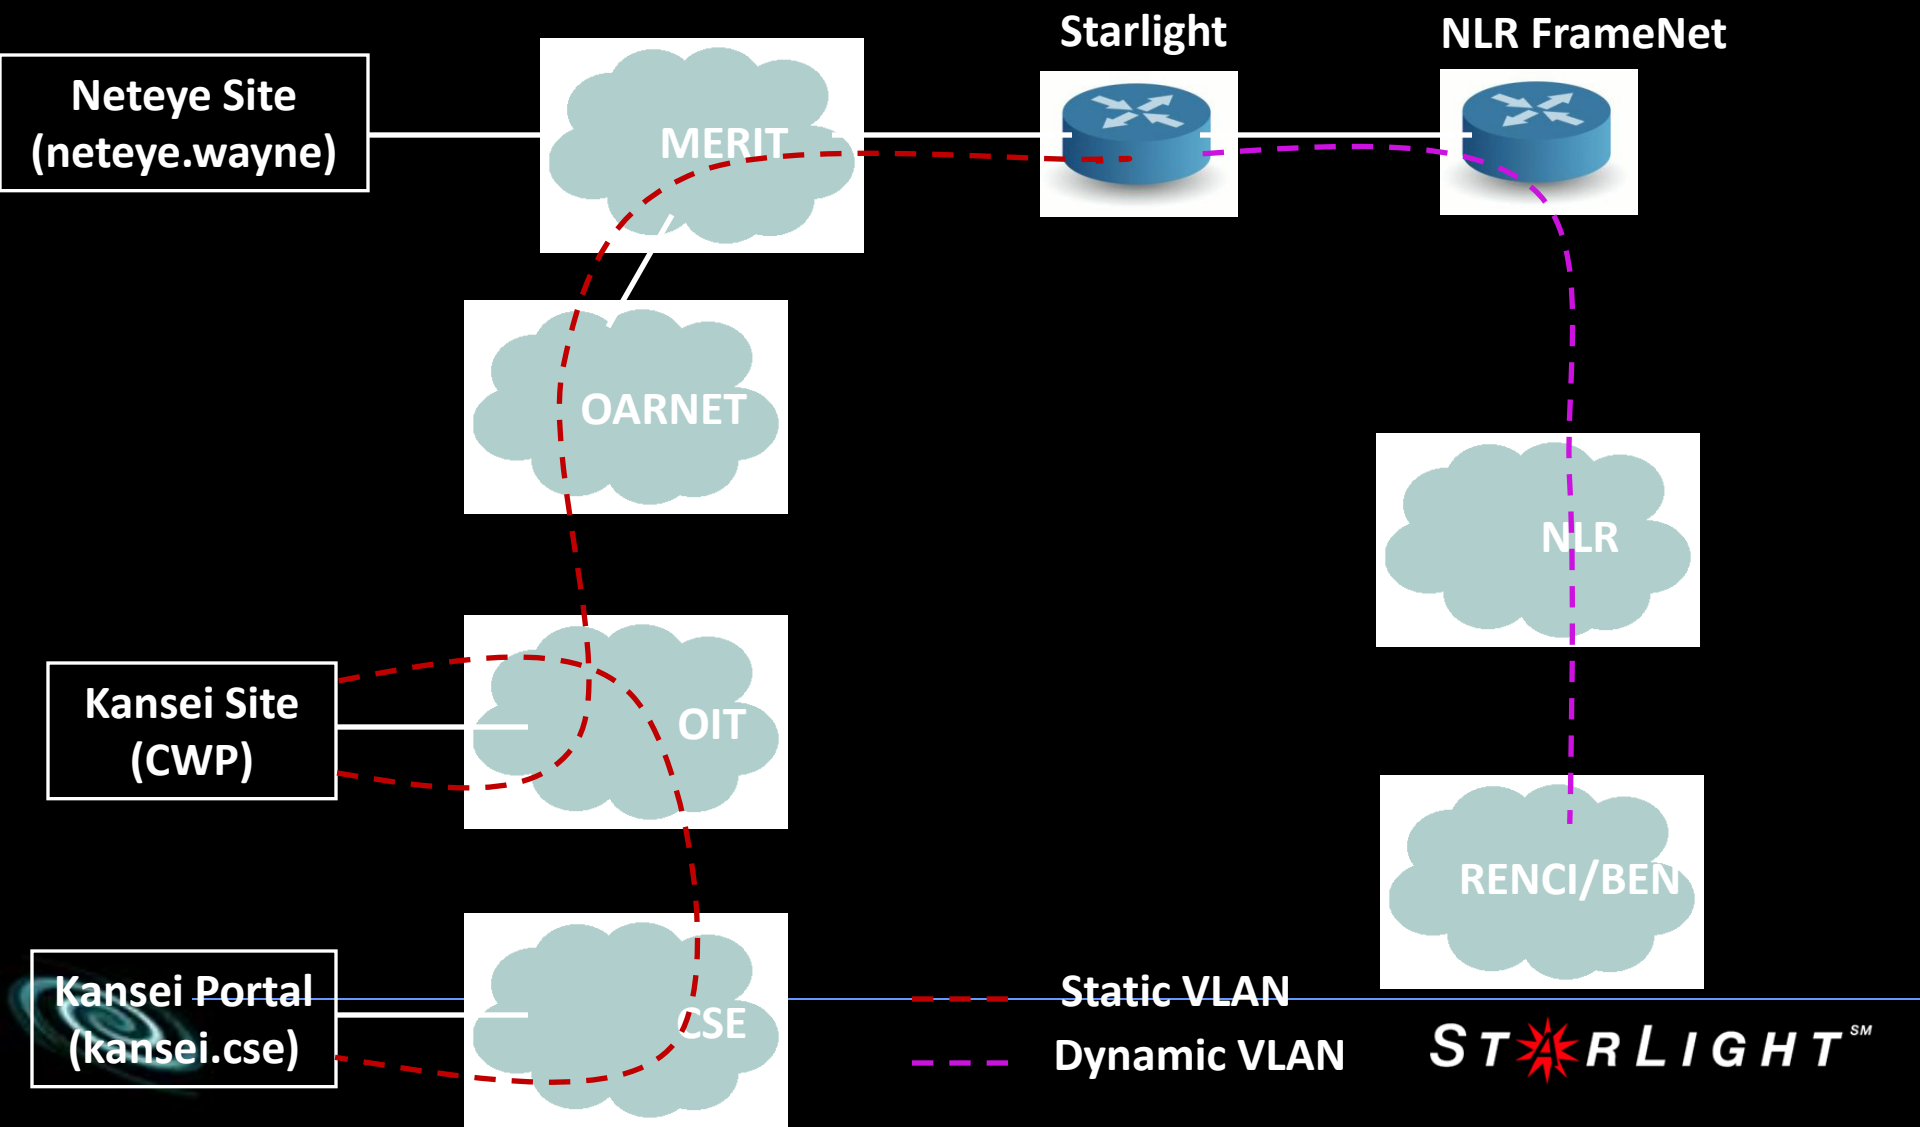

IGENI Project
StarLight to Elephant Cluster at UVic

Ian Gable
Version 0.3
2012-06-13
Note: Only a rough idea

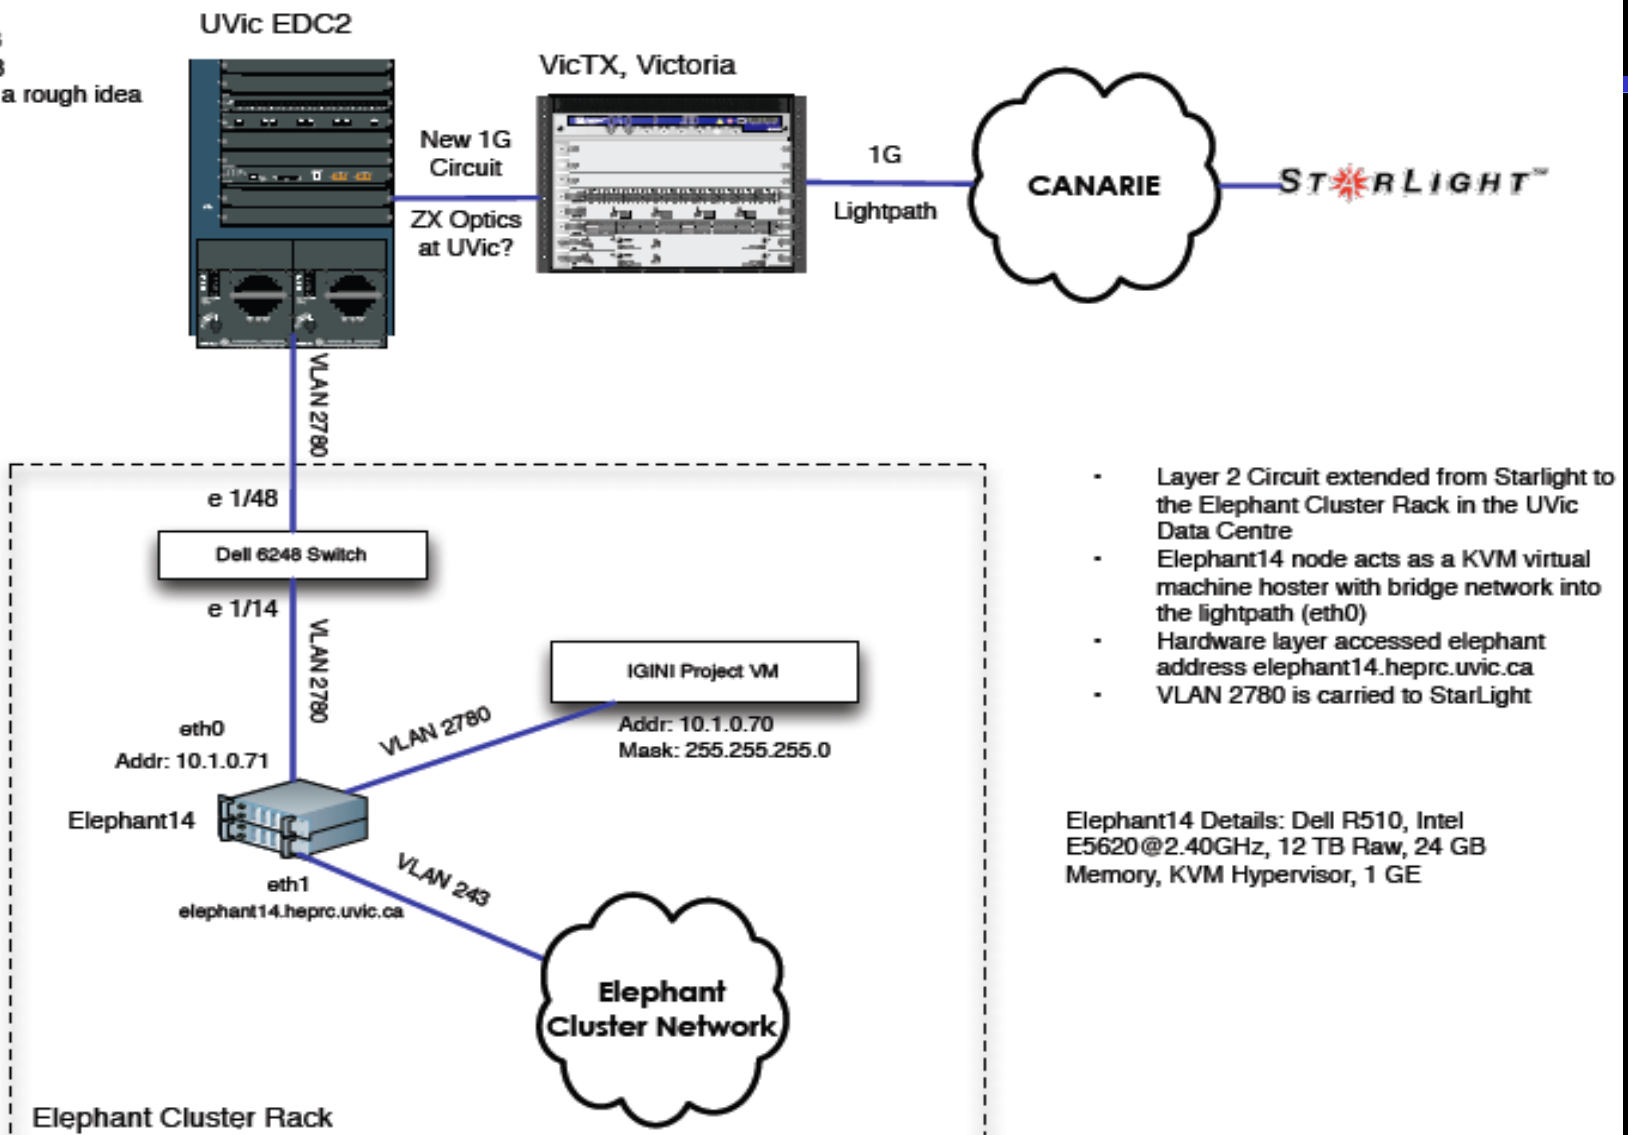

- Layer 2 Circuit extended from Starlight to the Elephant Cluster Rack in the UVic Data Centre
- Elephant14 node acts as a KVM virtual machine hoster with bridge network into the lightpath (eth0)
- Hardware layer accessed elephant address elephant14.heprc.uvic.ca
- VLAN 2780 is carried to StarLight

Elephant14 Details: Dell R510, Intel E5620@2.40GHz, 12 TB Raw, 24 GB Memory, KVM Hypervisor, 1 GE

StarWave: A Multi-100 Gbps Facility

- StarWave, A New Advanced Multi-100 Gbps Facility and Services Will Be Implemented Within the StarLight International/NationalCommunications Exchange Facility
- StarWave Is Being Funded To Provide Services To Support Large Scale Data Intensive Science Research Initiatives
- Facilities Components Will Include:
  - An ITU G. 709 v3 Standards Based Optical Switch for WAN Services, Supporting Multiple 100 G Connections
  - An IEEE 802.3ba Standards Based Client Side Switch, Supporting Multiple 100 G Connections, Multiple 10 G Connections
  - Multiple Other Components (e.g., Optical Fiber Interfaces, Measurement Servers, Test Servers
- ***GENI @ 100 Gbps***

GLIF/GENI/StarLight/StarWave/MREN Continually Progressing Forward!

www.startap.net/starlight