

Control Frameworks Based On OpenFlow Integrated Within the NSF's Large Scale Global Environment for Network Innovations (GENI)

Joe Mambretti, Director, (j-mambretti@northwestern.edu)

International Center for Advanced Internet Research (www.icaair.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



Global Environment for Network Innovations (GENI)

- **GENI Is Funded By The National Science Foundation's Directorate for Computer and Information Science and Engineering (CISE)**
- **GENI Is a Virtual Laboratory For Exploring Future Internets At Scale.**
- **GENI Is Similar To Instruments Used By Other Science Disciplines, e.g., Astronomers – Telescopes, HEP - Synchrotrons**
- **GENI Creates Major Opportunities To Understand, Innovate and Transform Global Networks and Their Interactions with Society.**
- **GENI Is Dynamic and Adaptive.**
- **GENI Opens Up New Areas of Research at the Frontiers of Network Science and Engineering, and Increases the Opportunity for Significant Socio-Economic Impact.**



Global Environment for Network Innovations (GENI)

- **GENI**
 - **Supports At-Scale Experimentation on Shared, Heterogeneous, Highly Instrumented Infrastructure**
 - **Enables Deep Programmability Throughout the Network,**
 - **Promotes innovations in Network Science, Security, Technologies, Services and Applications**
 - **Provides Collaborative and Exploratory Environments for Academia, Industry and the Public to Catalyze Groundbreaking Discoveries and Innovation.**
- **Current Round of GENI Funding Accelerates Prototyping of Suite of Infrastructure for the GENI Project With Federation and Early Experiments to Guide Future GENI System Design.**
- **GENI To a Large Degree Is Infrastructure-As-A-Service (IaaS)**
- **Major Potential For Supporting Wide Range of Science Research**



iGENI: The International GENI

- **The iGENI Initiative Has Designed, Developed, Implemented, and Now Operates a Major New National and International Distributed Infrastructure.**
- **iGENI Placed the “G” in GENI Making GENI Truly Global.**
- **iGENI Is a Unique Distributed Infrastructure Supporting Research and Development for Next-Generation Network Communication Services and Technologies.**
- **This Infrastructure Has Been Integrated With Current and Planned GENI Resources, and Is Operated for Use by GENI Researchers Conducting Experiments that Involve Multiple Aggregates At Multiple Sites.**
- **iGENI Infrastructure Has Connected Its Resources With The GENI National Backbone Network, With GENI Regional Related Networks, and With International Research Networks and Projects.**



Selected GENI/iGENI Components

- **GENI/iGENI Core Infrastructure**
- **L1/L2 Core Paths**
- **Control Plane**
- **Management Plane**
- **Resource Management Processes**
- **Resource Data Paths**
- **Core Resources (Including Edge Clouds)**
- **Experimental Resources**
- **Overall Architecture – “3 Concentric Spheres”**



External Aggregate Interfaces Are L1/L2 Paths (e.g., vLANs, Tunnels, Et AI)

- **MPLS (L 2.5)**
- **vLANs**
 - Mapped at ingress/egress or
 - Tunneled (Q-in-Q)
- **Tunnels**
 - Layer 2, such as OpenVPN
- **SONET/SDH, GFP, vCAT Channels,**
- **Layer 1**
 - Lightpaths
- **Provisioning**
 - Statically Provisioned Based on Requirements
 - Partially Dynamically Provisioned Based on Requirements
 - Dynamically Provisioned Based on Requirements

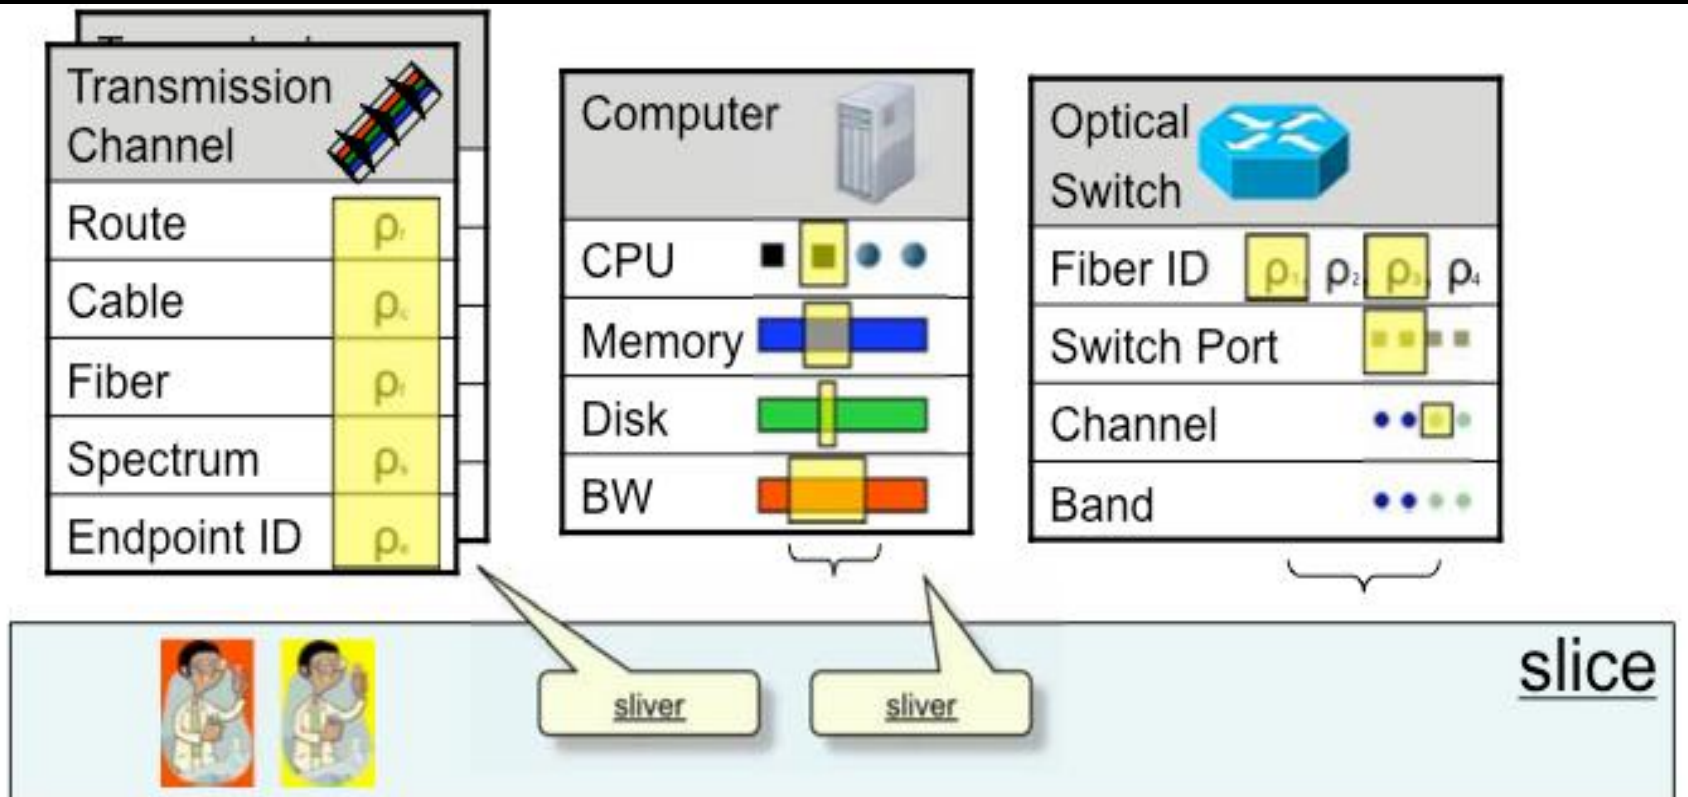


Building Experimental Networks

- **GENI Network Abstractions (Provided by Tools) For Experimental Resources**
 - Automatic Path Discovery and Allocation
 - Control Over Path/Network Selection for Experimenter
 - Various Parameters, For Example, Bandwidth.
 - Bare Bones Control Over Everything Available to the ER Tools
 - Resource Description Semantics (Rspecs)



GENI Slicers and Slivers



Source: BBN GENI Program Office

GENI Control Frameworks

- **Four Primary Control Frameworks**
 - Open Resource Control Architecture (ORCA)
 - ProtoGENI
 - PlanetLab
 - ORBIT
- **Integrated Through The GENI Aggregate Manager API, Which**
 - Specifies A Set Of Functions For Reserving Resources
 - Describes Common Format for Certificates And Credentials to Enable Compatibility Across All Aggregates

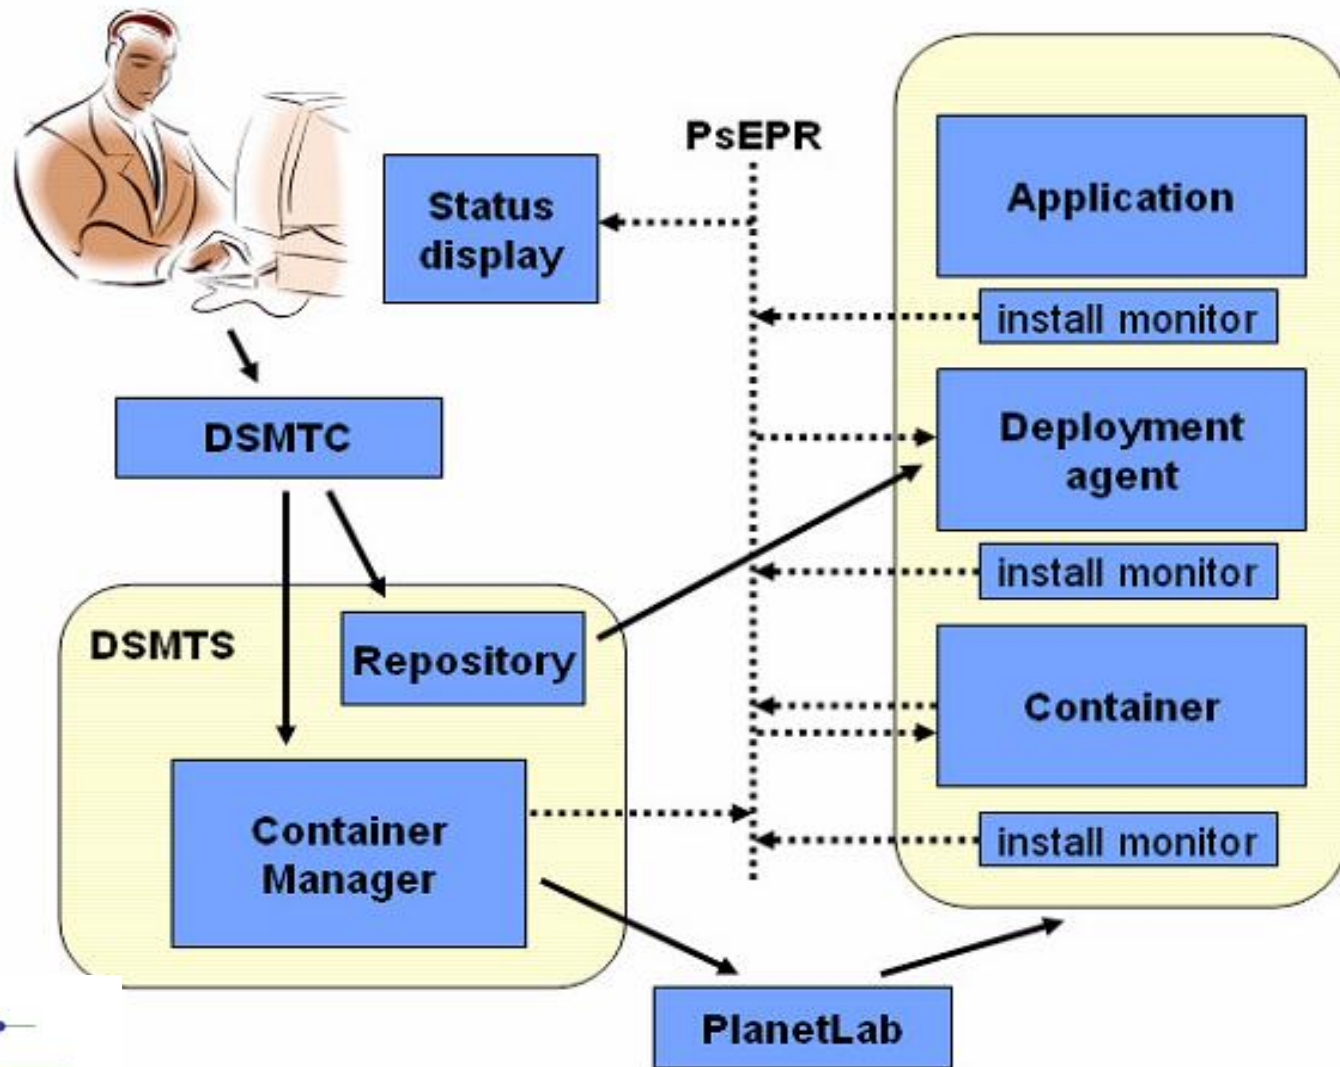


Invocation Mechanism

- **A Slice Federation Architecture (SFA) Defines Global Identifiers (GID)**
- **XML-RPC over HTTPS [RFC 2818] Used As Invocation Mechanism for GENI API operations.**
- **XML-RPC: Remote Procedure Call Mechanism Using HTTP For Transport And XML As The Encoding.**
- **XML-RPC: Readily Secured By SSL, Can Be Used From a Variety of Programming Languages.**
- **Within GENI Used By Specific Control Frameworks.**
- **However, Can Also Be Used By Alternative Control Frameworks**
- **Can Also Be Used Independently**

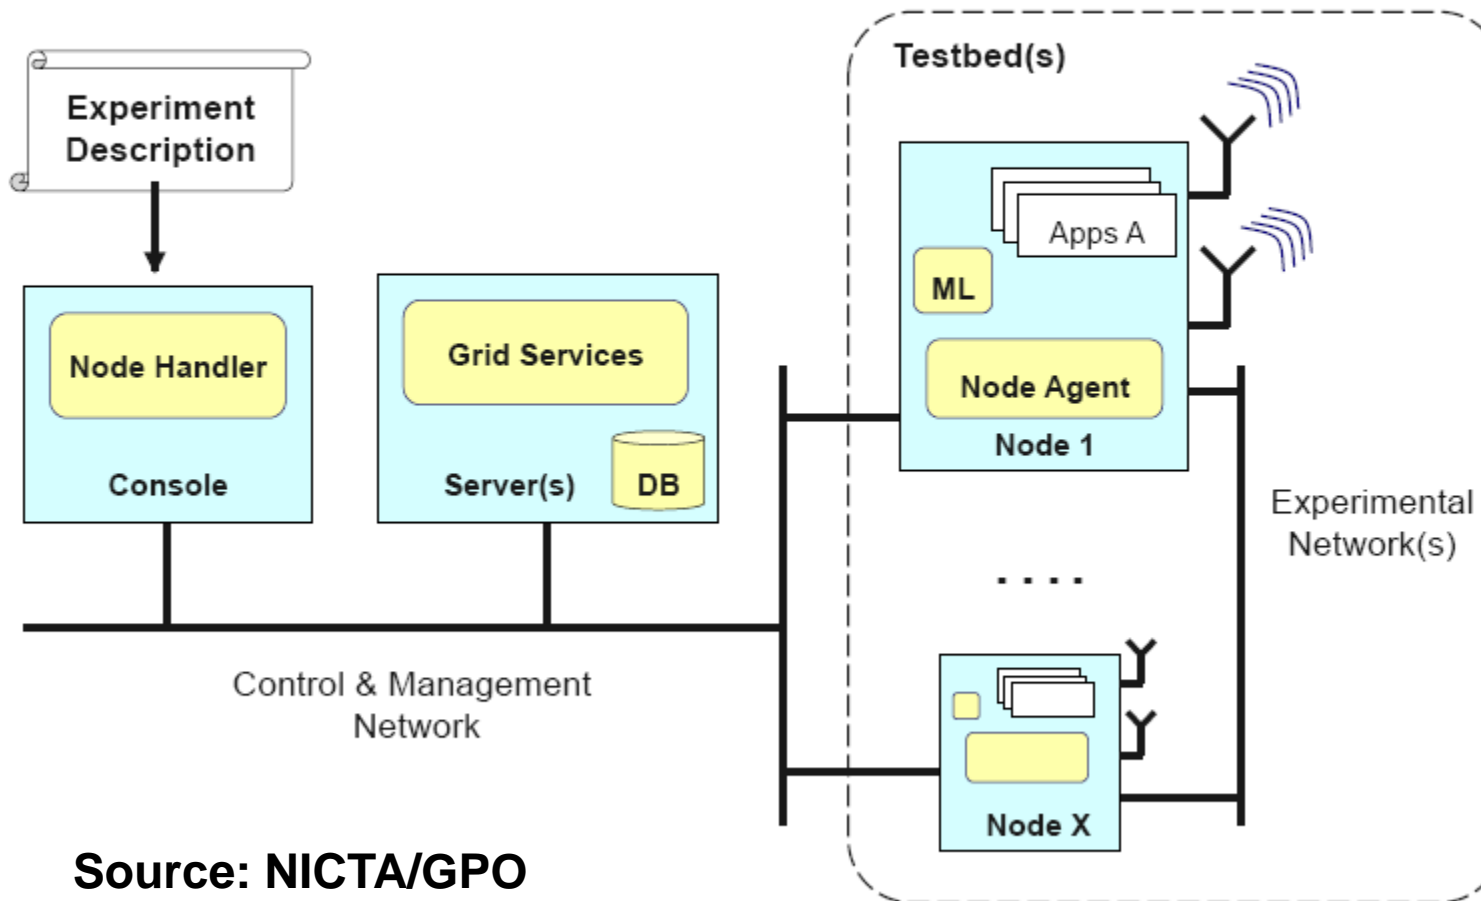


PlanetLab



ORBIT Management Framework

OMF - Operator View



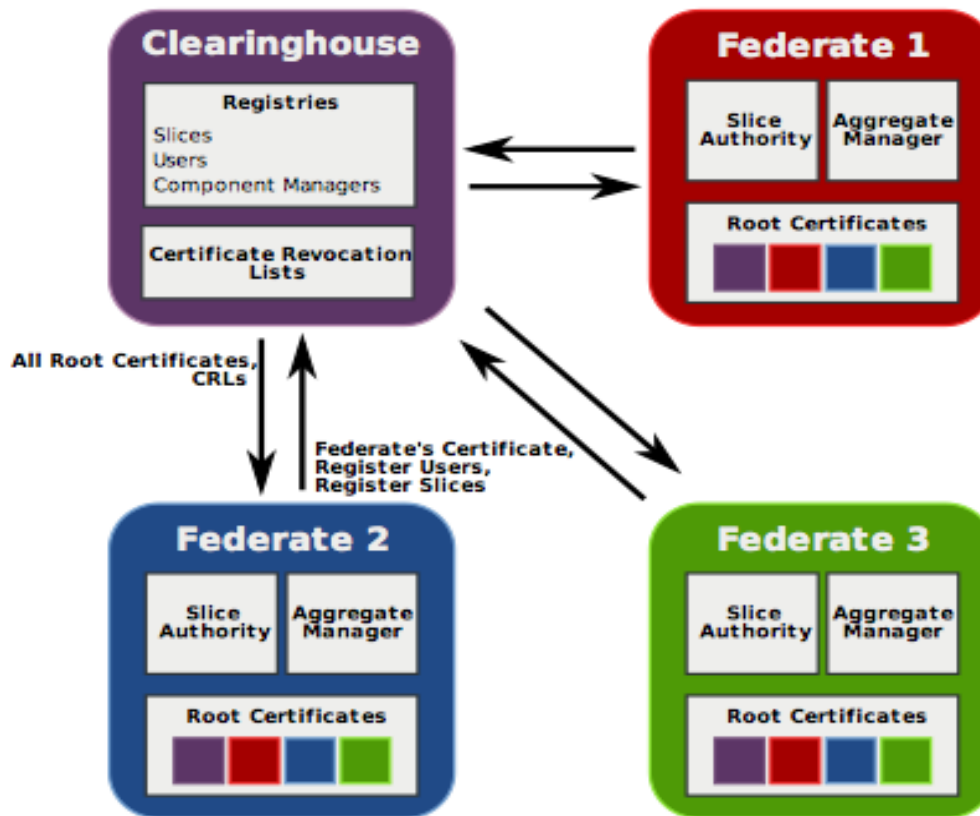
Source: NICTA/GPO

ProtoGENI

- **ProtoGENI Is Based On Emulab (Linux Based Experimental Environment)**
- **It Supports Hardware as a Service (HaaS)**
- **It Is Essentially An Instantiation Of Platform as a Service (PaaS)**
- **It Enables The Installation Of Custom OS Images on Bare Metal**
- **It Has Been Integrated With FlowVisor OpenFlow Agg Mgr (FOAM)**
- **It Provides OpenFlow Programmability For Experiment Dataplanes**
- **It Uses OpenVZ images for Shared Nodes**
- **It Supports PlanetLab Node Images**

ProtoGENI Federation Design

Federation Design



Flack: Experimenter Interface

The screenshot displays the Flack web interface in a browser window. The browser's address bar shows the URL `www.protogeni.net/trac/protogeni/wiki/Flack`. The interface includes a top navigation bar with menu items like File, Edit, View, History, Bookmarks, Tools, and Help. Below the browser window, the Flack application is visible, featuring a sidebar with a list of managers and a main area showing a network diagram. The network diagram consists of several nodes representing PCs and their interconnections via LANs. Each node is labeled with its IP address, hostname, and other details. The interface also includes a search bar at the bottom with the text "Find: buildroutes" and navigation controls like "Previous", "Next", "Highlight all", and "Match case".

Renew Floodlight: Renewed. All expire at the same time in 3 days, 17 hours, 12 minutes. v1.4.46

Flack

Floodlight View Import Output GENI...

psingha

Slices New

Show All Floodlight

Managers Add

Show/Hide

- ansp.br.cm
- bbn-pgeni.cm
- beelab.cm
- cis.fiu.edu.cm
- cmulab.cm
- cron.loni.org.cm
- ETRI-CM1.cm
- genicloud.hplabs...
- jonlab.cm
- mygeni.cm
- plc.sa
- shadowgeni.cm
- ukoeni.cm

Add Resources Get Status

utahemulab.cm Renew Sliver: Renewed

PC-0 pc508.emulab.net:22
excl. raw-pc
UBUNTU12-64-STD
=> pc508

lan6 lan4 lan2 lan3 lan5 lan7

PC-1 pc482.emulab.net:22
excl. raw-pc
UBUNTU12-64-STD
=> pc482

lan1 lan0

PC-2 pc501.emulab.net:22
excl. raw-pc
UBUNTU12-64-STD
=> pc501

PC-3 pc502.emulab.net:22
excl. raw-pc
UBUNTU12-64-STD
=> pc502

PC-4 pc490.emulab.net:22
excl. raw-pc
UBUNTU12-64-STD
=> pc490

PC pc505.emulab.net:22
excl. raw-pc
UBUNTU12-64-STD
=> pc505

Delete 3d 14h 13m Extend 3 Days

Find: buildroutes Previous Next Highlight all Match case

InstaGENI




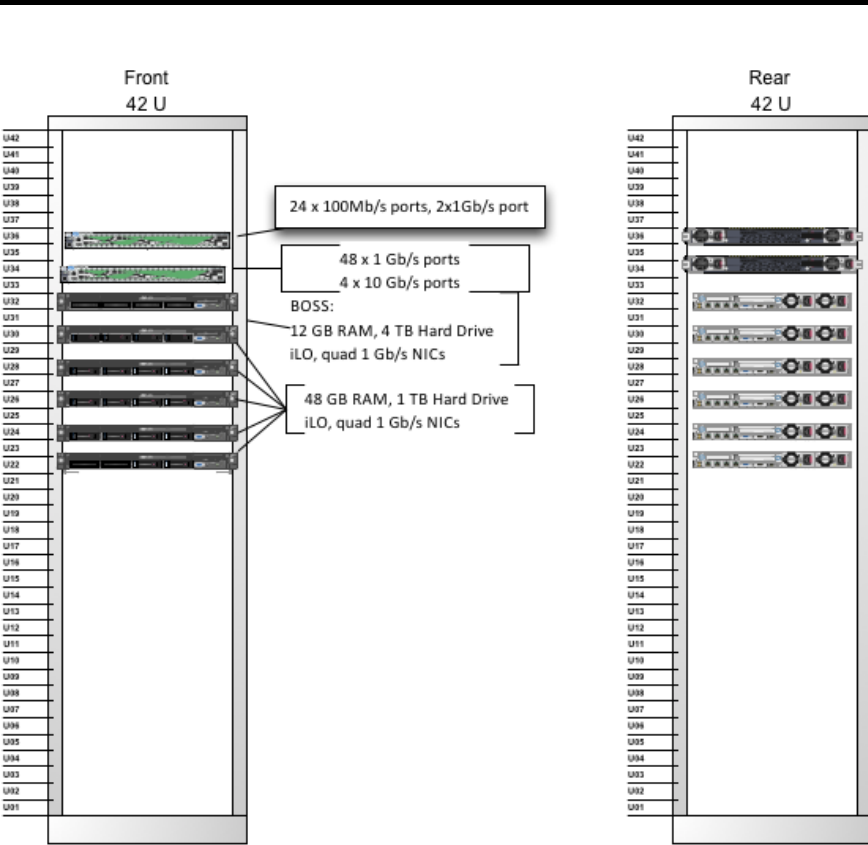
geni

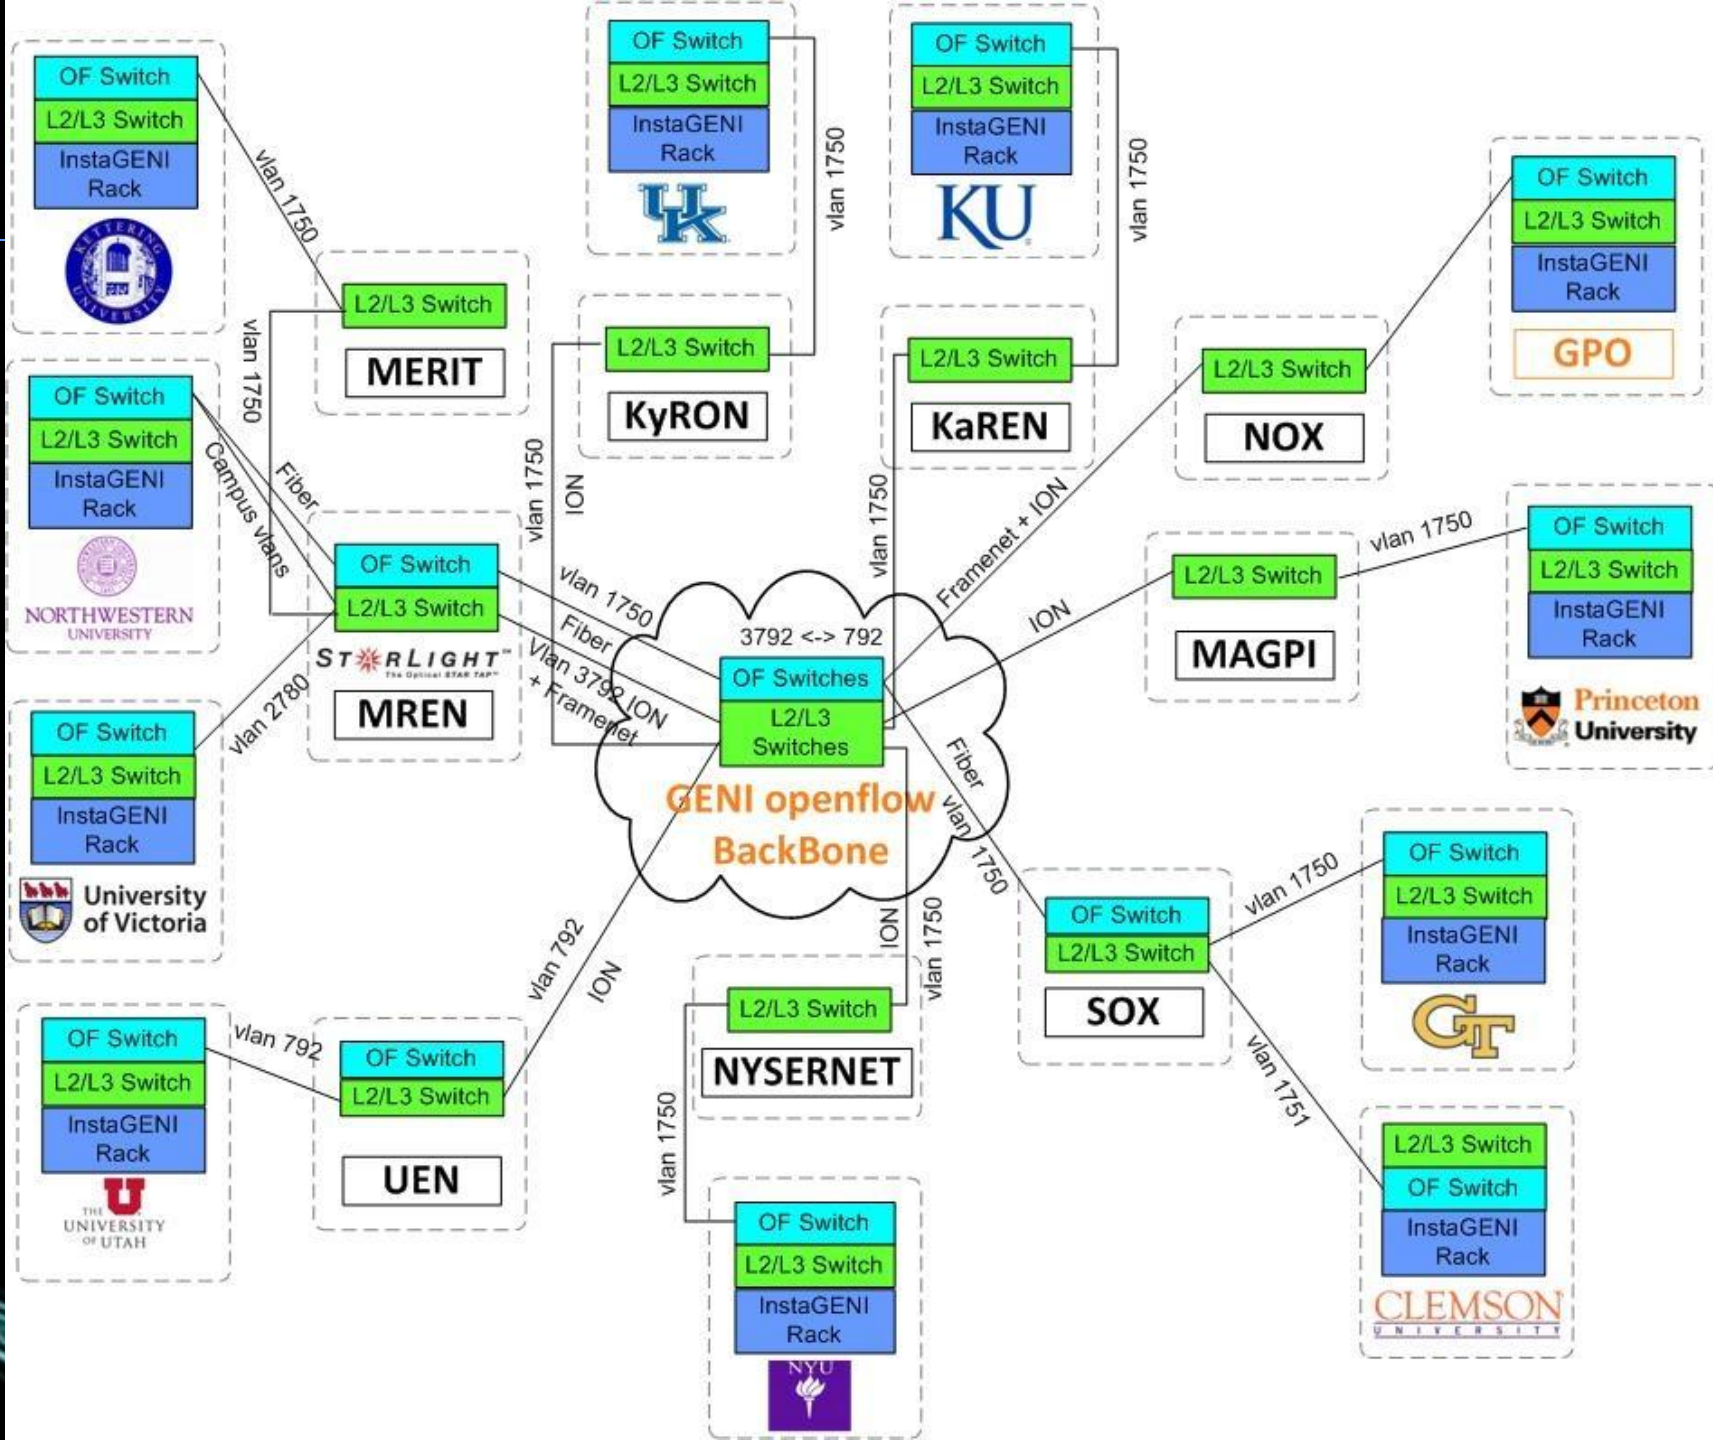
Exploring Networks
of the Future

www.geni.net



InstaGENI Racks and Topology





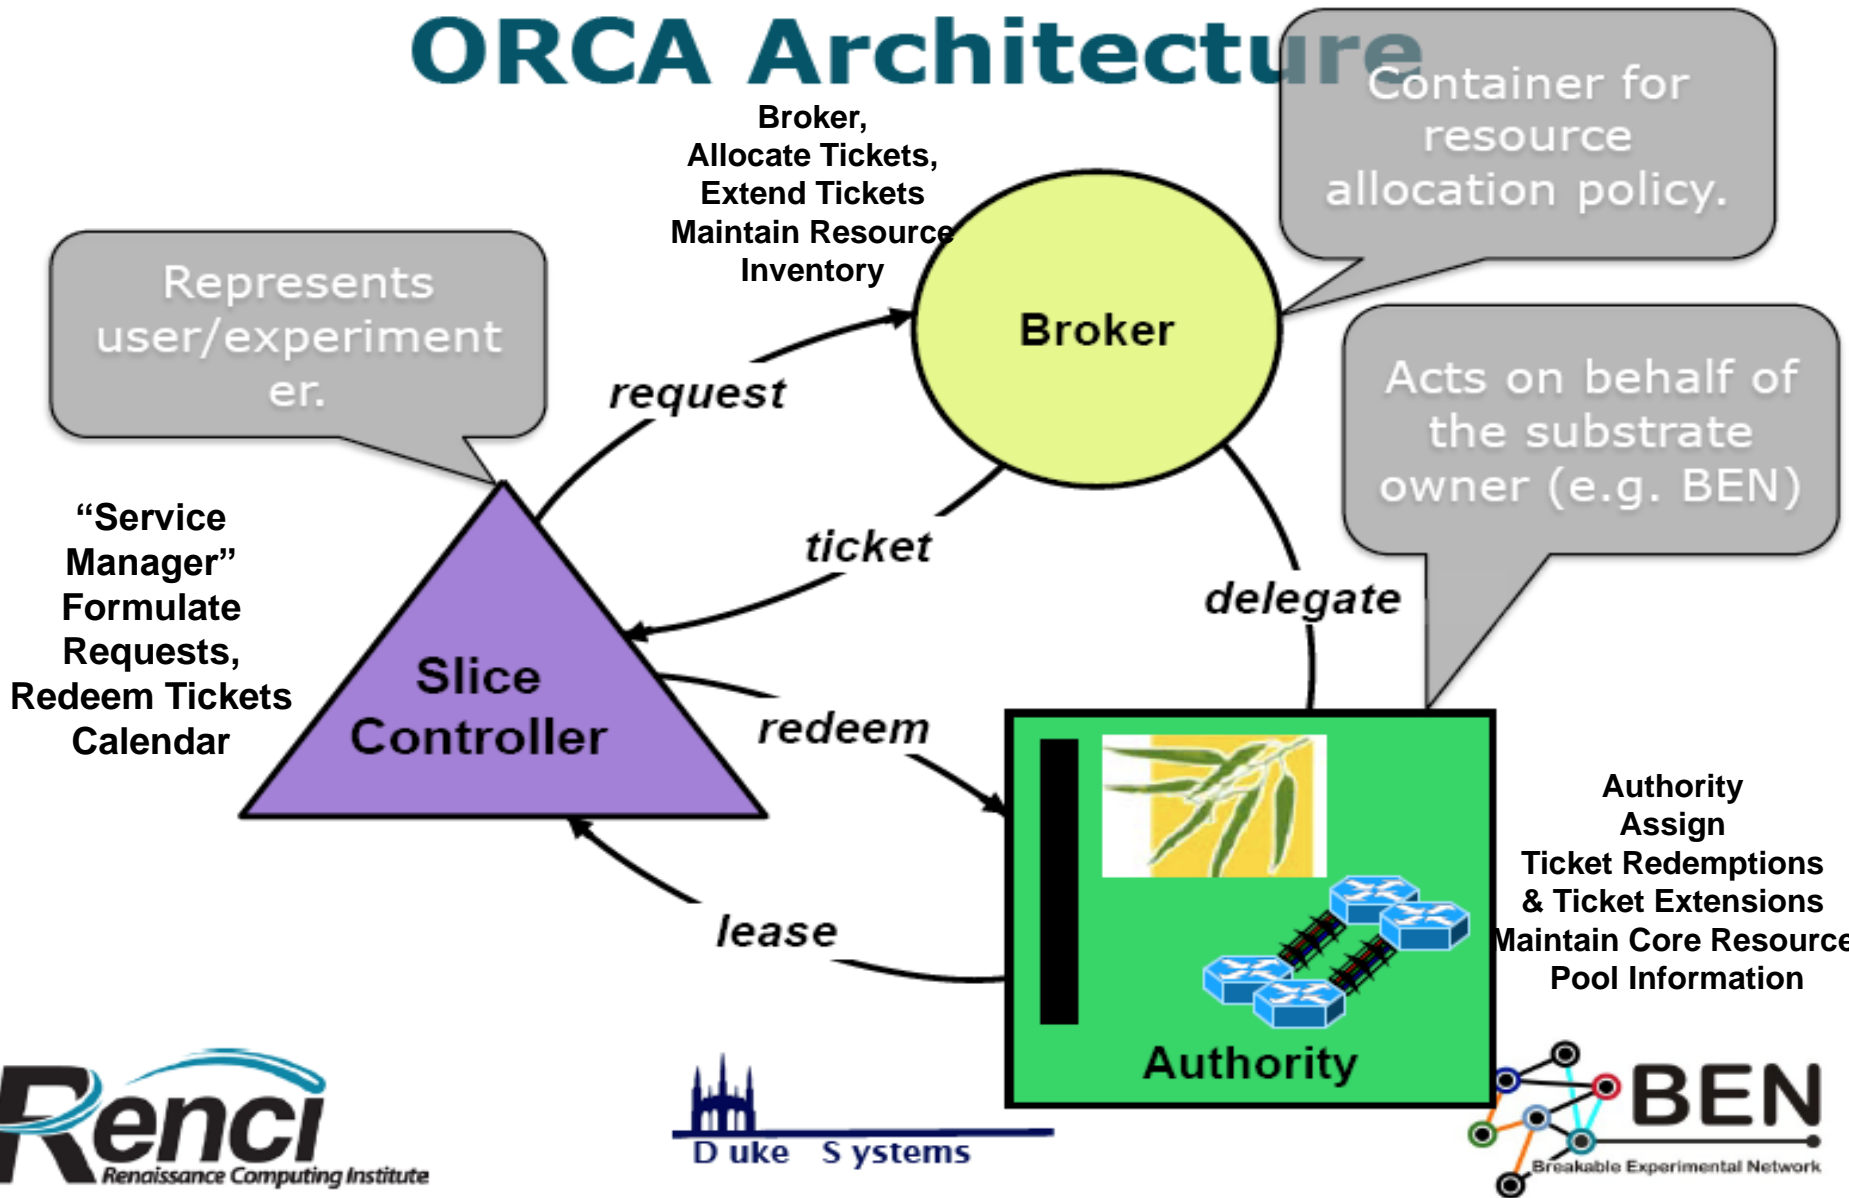
Open Resource Control Architecture (ORCA)

- **Cross Cloud Provisioning: Connections, Clusters, and Topologies**
- **Provisioning Among Multi-Layered Network Domains**
- **Leveraging Semantic Web Technologies**
- **API and Resource Description Translation**
- **ORCA Enables The Creating And Instantiation Of Application Specific Topologies Among Diverse Independent Clouds – Interconnected By Multi-Layer Networks**

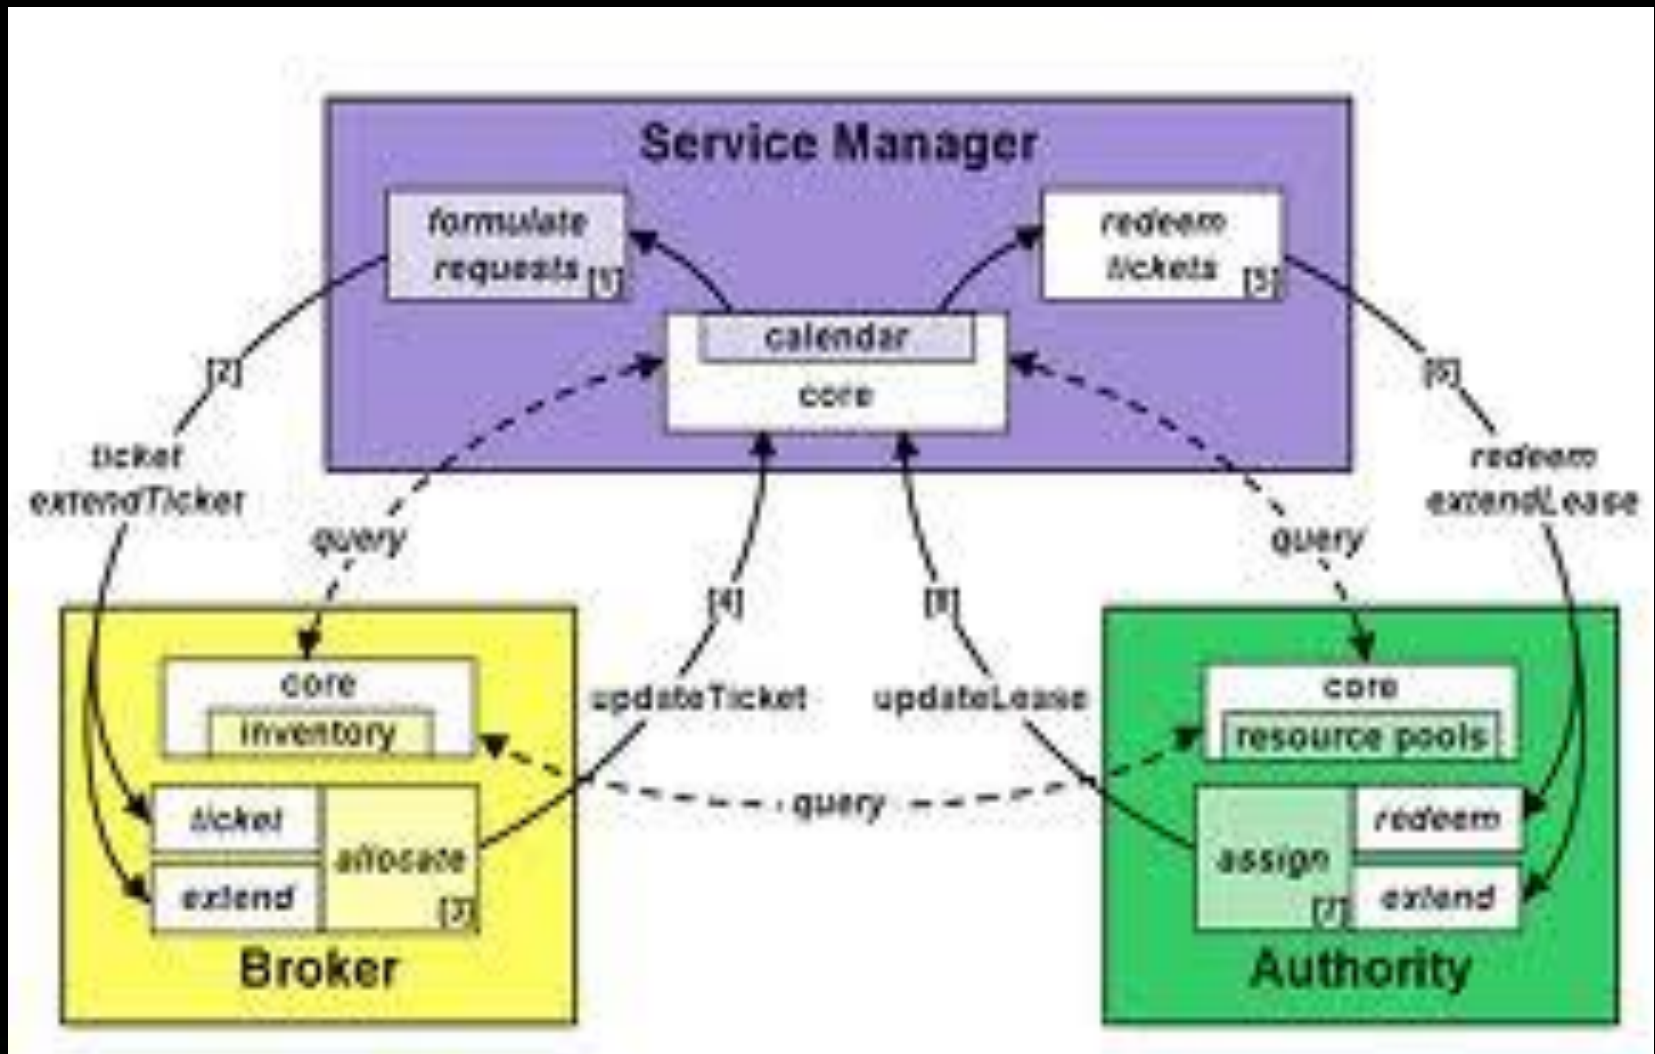


Being Developed @: RENCIBEN-Duke-GENI

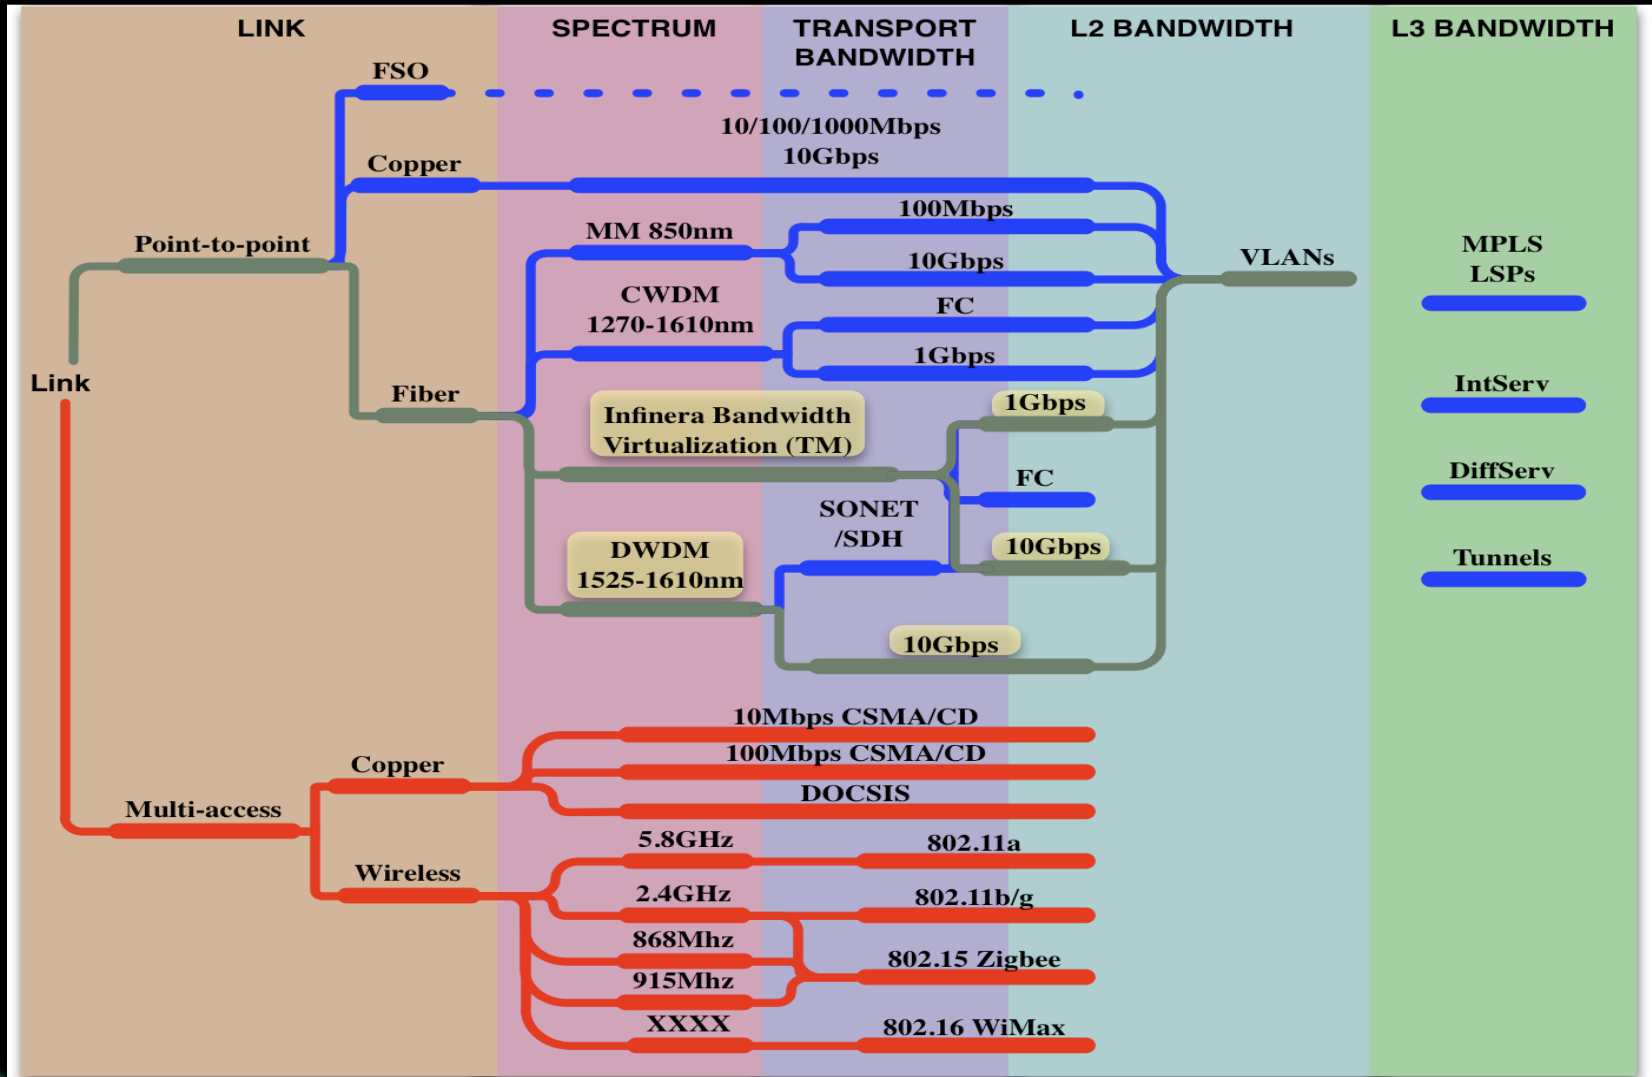
ORCA Architecture



ORCA Architecture

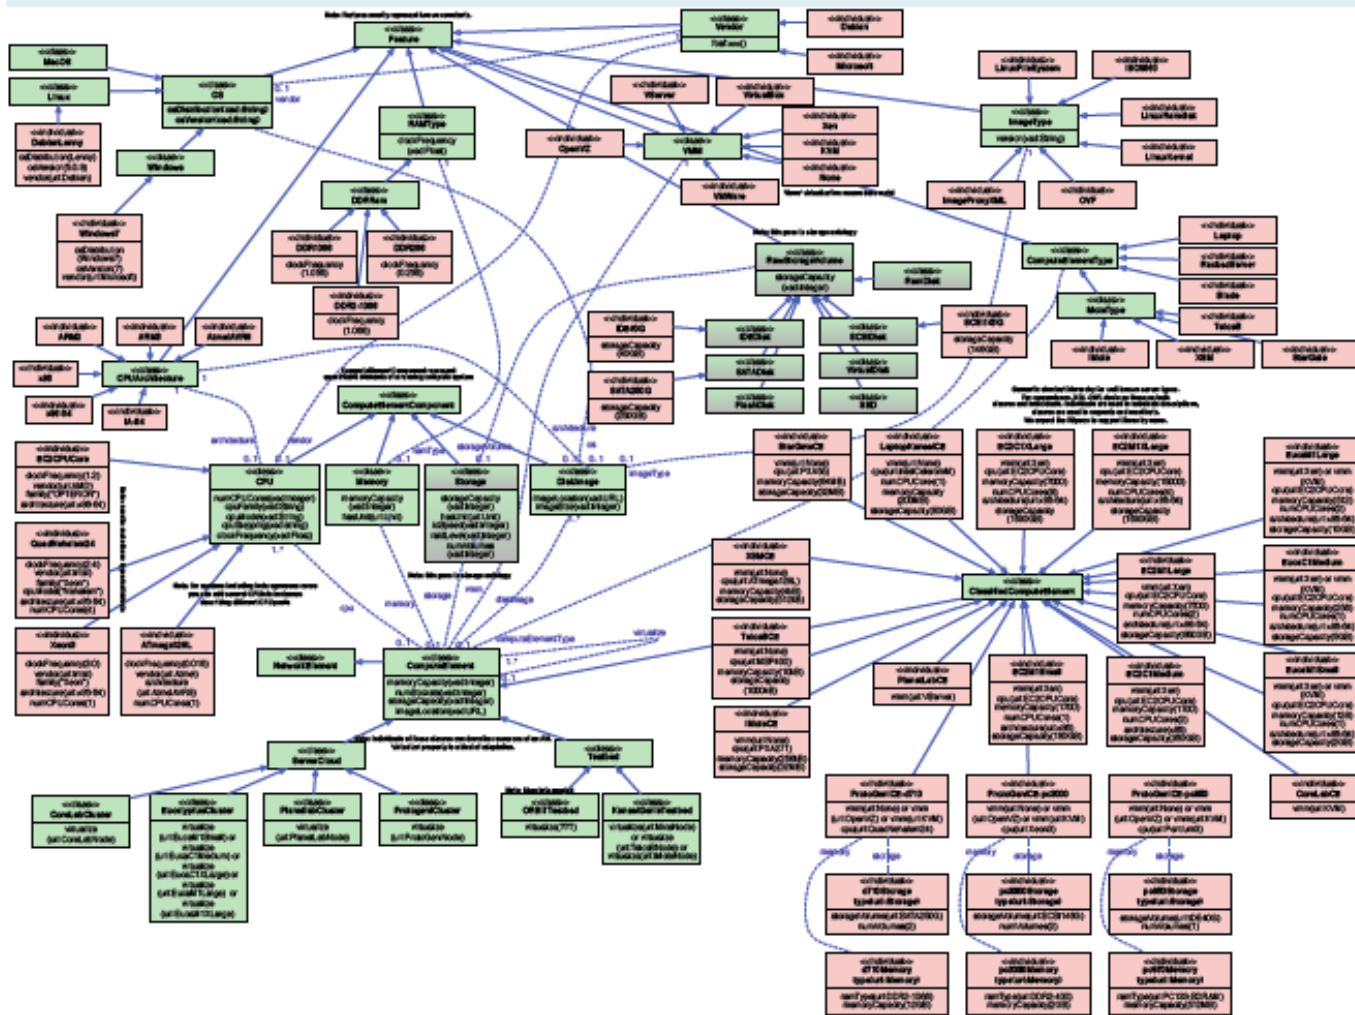


ORCA "Link" Slivering

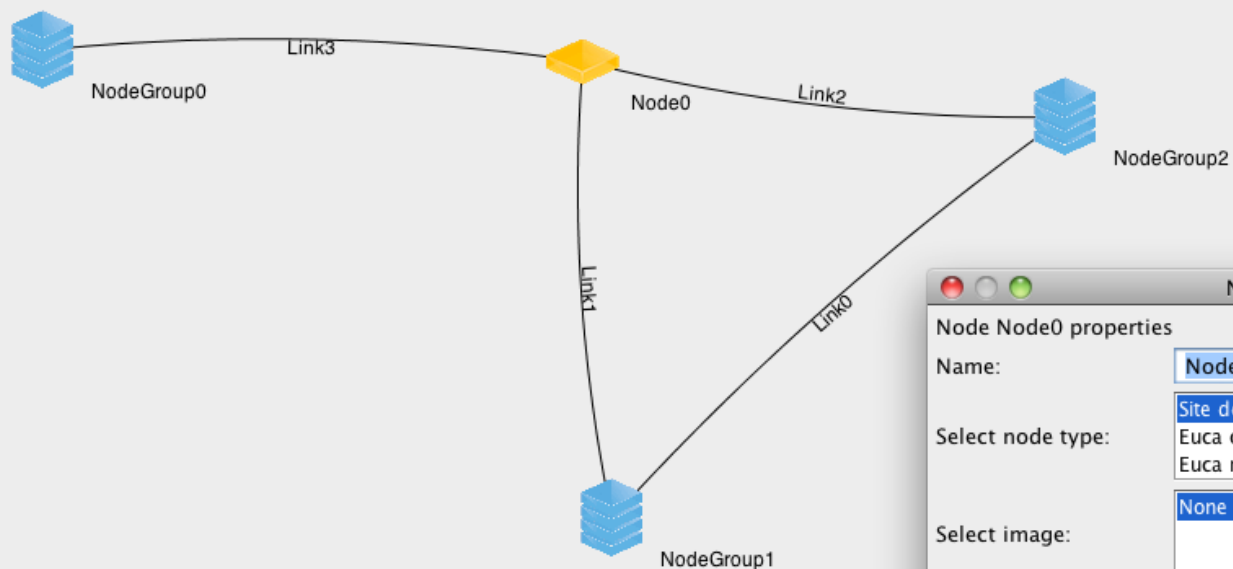


ORCA Semantics

Compute.owl



Add Nodes Add Node Groups Client Images Edit Reservation



Node Properties

Node Node0 properties

Name:

Select node type:
 Euca c1.medium
 Euca m1.small

Select image:

Select domain:
 Duke CS
 NERSC

Select dependencies:
 NodeGroup2
 NodeGroup1

Link1 IP Address:
 Link2 IP Address:
 Link3 IP Address:

Additional properties:



- **14 GPO-funded racks**
 - Partnership between RENCI, Duke and IBM
 - IBM x3650 M3/M4 servers
 - 1x146GB 10K SAS hard drive +1x500GB secondary drive
 - 48G RAM 1333Mhz
 - Dual-socket 6-core Intel X5650 2.66Ghz CPU
 - Dual 1Gbps adapter
 - 10G dual-port Chelseo adapter
 - BNT 8264 10G/40G OpenFlow switch
 - DS3512 6TB sliverable storage
 - iSCSI interface for head node image storage as well as experimenter slivering
- **Each rack is a small networked cloud**
 - OpenStack-based (some older racks run Eucalyptus)
 - EC2 nomenclature for node sizes (m1.small, m1.large etc)
 - Interconnected by combination of dynamic and static L2 circuits through regionals and national backbones
- <http://www.exogeni.net>




StarWave: A Multi-100 Gbps Facility

- **StarWave, A New Advanced Multi-100 Gbps Facility and Services Will Be Implemented Within the StarLight International/National Communications Exchange Facility**
- **StarWave Is Being Funded To Provide Services To Support Large Scale Data Intensive Science Research Initiatives**
- **Facilities Components 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)**
- **Also, GENI @ 100 Gbps**



ExoGENI at a glance



iGENI and StarLight

- **iGENI Integrates Multiple Network Resources, Including:**
 - **Resources at the StarLight International Communications Exchange in Chicago - StarLight Current Supports Over 20 Major Network Research Testbeds, Including National and International Fabrics**
 - **Segments of National Research and Education Network Infrastructures**
 - **A National Wide-Area Private Testbed Networks**
 - **Components of the International Optical-Networking -- GLIF Fabric.**



StarLight – “By Researchers For Researchers”

StarLight is an advanced national and international communication exchange facility optimized for high-performance data intensive applications


World’s “Largest”
10G and 100 Gbps
Exchange
Over 100 10 Gbps
Channels
Soon Over 25
100 Gbps Paths
Interoperability
Services
At All Layers



View from StarLight



Abbott Hall, Northwestern University's Chicago downtown campus




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



GLIF Map 2009: Global Lambda Integrated Facility Visualization by Robert Patterson, NCSA, University of Illinois at Urbana-Champaign Data Compilation by Maximo D. Brown, University of Illinois at Chicago Earth Texture, visbleearth.nasa.gov

www.glif.is glif

www.glif.is

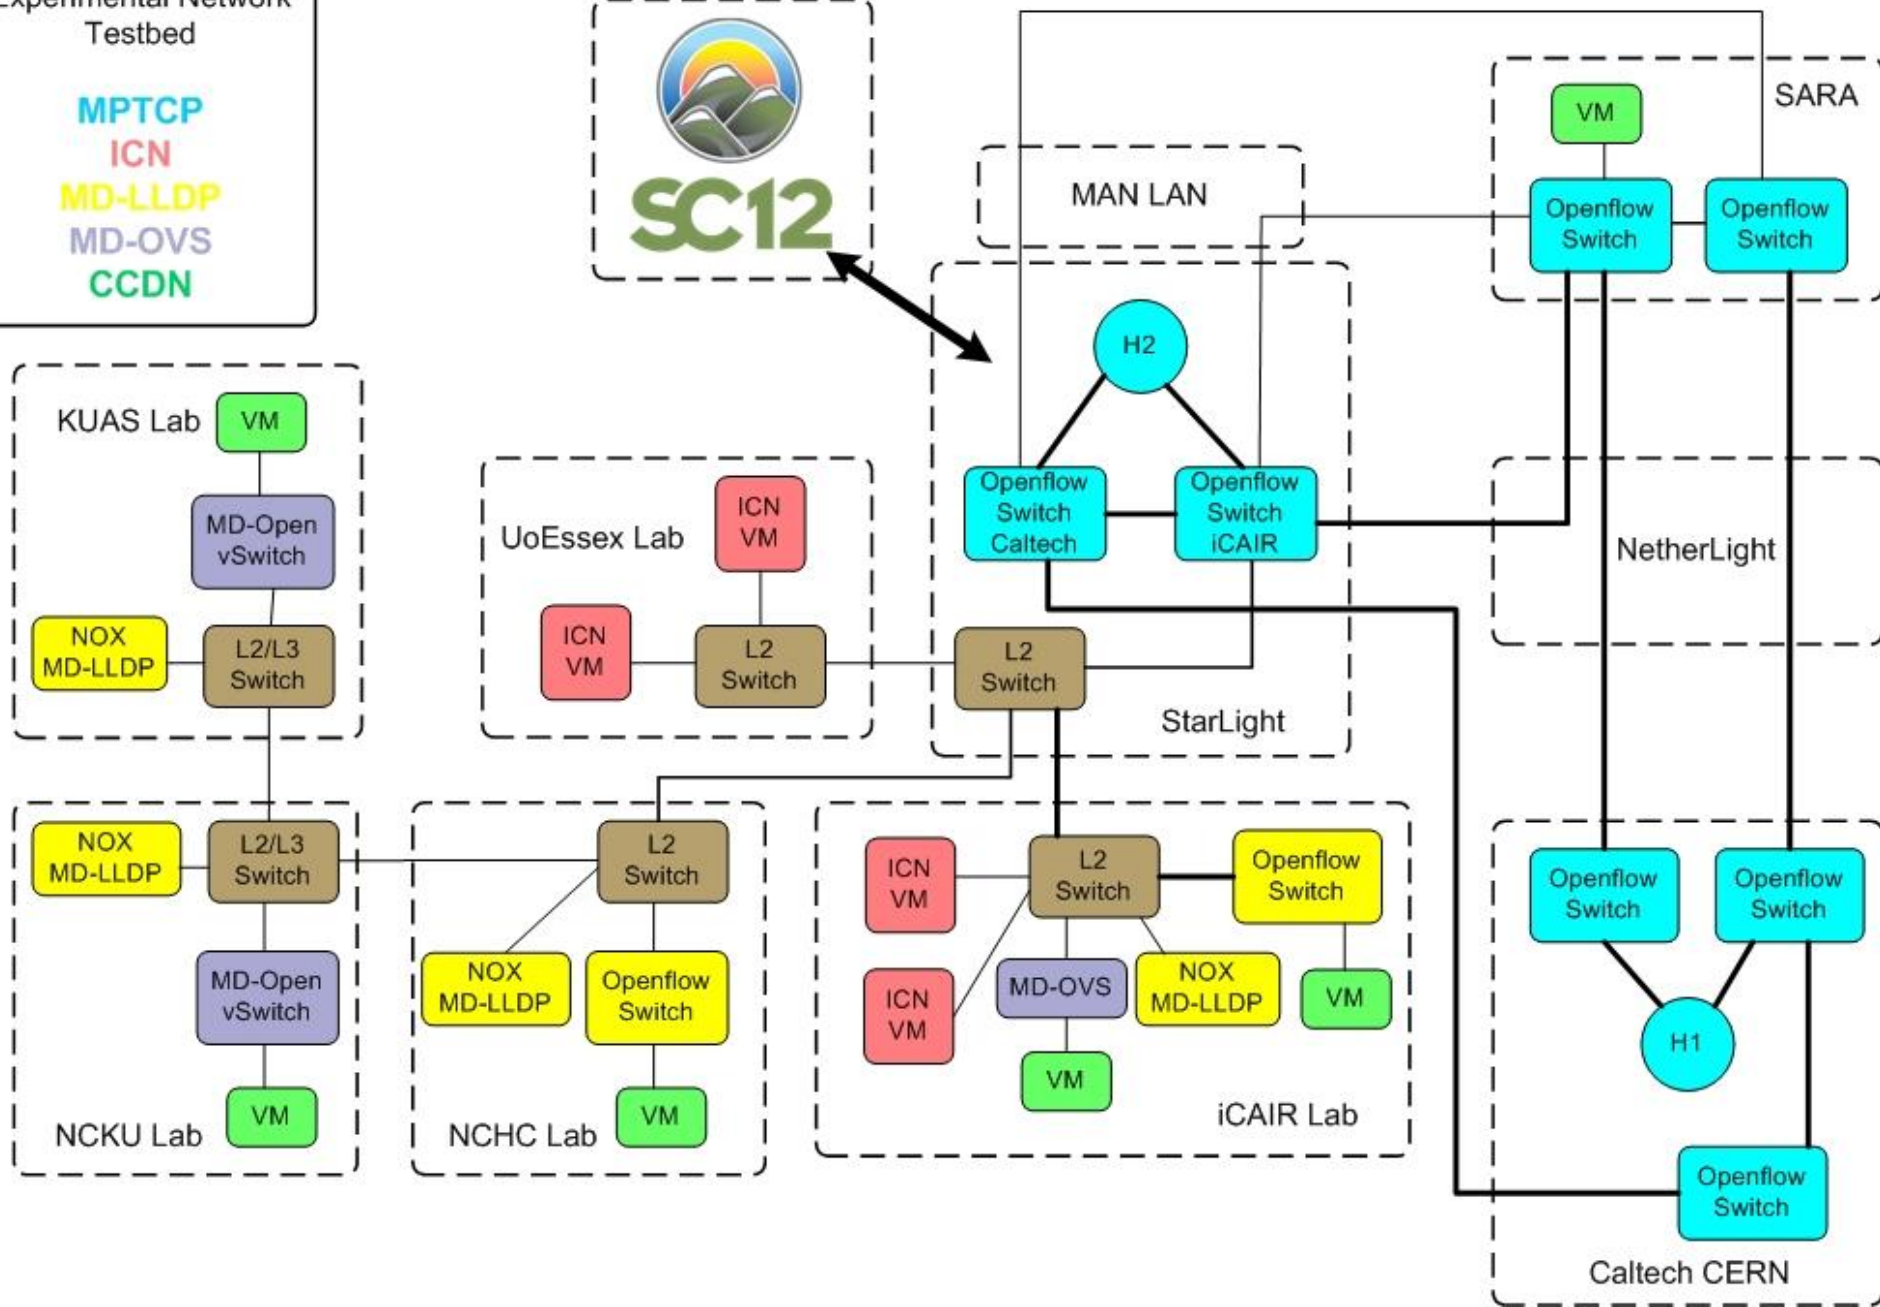
STARLIGHT™



SC12 SRS

International Openflow
Experimental Network
Testbed

MPTCP
ICN
MD-LLDP
MD-OVS
CCDN



www.startup.net/starlight

Thanks to the NSF, DOE, NASA, DARPA,
Universities, National Labs,
International Partners,
and Other Supporters

