

# Next Generation Clouds, The Chameleon Cloud Testbed, and Software Defined Networking (SDN)

Joe Mambretti, Director, ([j-mambretti@northwestern.edu](mailto:j-mambretti@northwestern.edu))

International Center for Advanced Internet Research ([www.icaair.org](http://www.icaair.org))  
Northwestern University

Director, Metropolitan Research and Education Network ([www.mren.org](http://www.mren.org))

Co-Director, StarLight, PI- StarLight SDX, PI-iGENI, PI-OMNINet, Co-PI  
Chameleon, ([www.startap.net/starlight](http://www.startap.net/starlight))

International Conference on Cloud Computing  
Research and Innovation  
Singapore  
October 28-29, 2015




# Introduction to iCAIR:




Accelerating Leading Edge Innovation and Enhanced Global Communications through Advanced Internet Technologies, in Partnership with the Global Community

- **Creation and Early Implementation of Advanced Networking Technologies - The Next Generation Internet All Optical Networks, Terascale Networks, Networks for Petascale Science**
- **Advanced Applications, Middleware, Large-Scale Infrastructure, NG Optical Networks and Testbeds, Public Policy Studies and Forums Related to NG Networks**
- **Three Major Areas of Activity: a) Basic Research b) Design and Implementation of Prototypes c) Operations of Specialized Communication Facilities (e.g., StarLight)**





ALMA: Atacama Large Millimeter Array




DØ (DZero) [www-d0.fnal.gov](http://www-d0.fnal.gov)



IVOA: International Virtual Observatory [www.ivoa.net](http://www.ivoa.net)



OSG [www.opensciencegrid.org](http://www.opensciencegrid.org)



ANDRILL: Antarctic Geological Drilling [www.andrill.org](http://www.andrill.org)



BIRN: Biomedical Informatics Research Network [www.nbirn.net](http://www.nbirn.net)




GLEON: Global Lake Ecological Observatory Network




WLCG [lcg.web.cern.ch/LCG/public/](http://lcg.web.cern.ch/LCG/public/)



Globus Alliance [www.globus.org](http://www.globus.org)



CAMERA metagenomics [camera.calit2.net](http://camera.calit2.net)



Carbon Tracker [www.esrl.noaa.gov/gmd/ccgg/carbontrack](http://www.esrl.noaa.gov/gmd/ccgg/carbontrack)



OOI-CI [ci.oceanobservatories.org](http://ci.oceanobservatories.org)



PRAGMA Pacific Rim Applications and Grid Middleware Assembly [www.pragma-grid.net](http://www.pragma-grid.net)




SKA [www.skatelescope.org](http://www.skatelescope.org)



Sloan Digital Sky Survey [www.sdss.org](http://www.sdss.org)



CineGrid [www.cinegrid.org](http://www.cinegrid.org)



ISS: International Space Station [www.nasa.gov/station](http://www.nasa.gov/station)



TeraGrid [www.teragrid.org](http://www.teragrid.org)



XSEDE [www.xsede.org](http://www.xsede.org)



LHCONE [www.lhccone.net](http://www.lhccone.net)



CLASS Comprehensive Large-Array Stewardship System [www.class.noaa.gov](http://www.class.noaa.gov)



Compilation By Maxine Brown

STARLIGHT<sup>SM</sup>

# Petascale Computational Science



For Decades, Computational Science  
Has Driven Network Innovation  
Today –  
Petascale Computational Science



National Center for Supercomputing Applications, UIUC



STARLIGHT<sup>SM</sup>

# TeraGrid=> XSEDE



Concept of Private Network As Backplane To Distributed Computational Environment Continues In Next Iteration

# Open Science Grid: Selected Investigations



DNA Modeling



Gravity Wave Modeling



Nutrino Studies

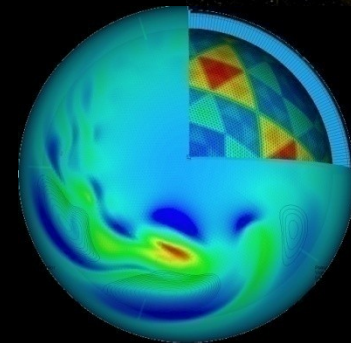





Usage



This Distributed Facility  
Supports Many Sciences

# HPC Cloud Computing



**DOE Magellan Initiative: Testbed  
To Explore Cloud Computing  
For Science**

# Multiple HPC Cloud Computing Testbeds Specifically Designed for Science Research



At Scale Experimentation  
Integrated With High Performance Networks





# Global Environment for Network Innovations (GENI)

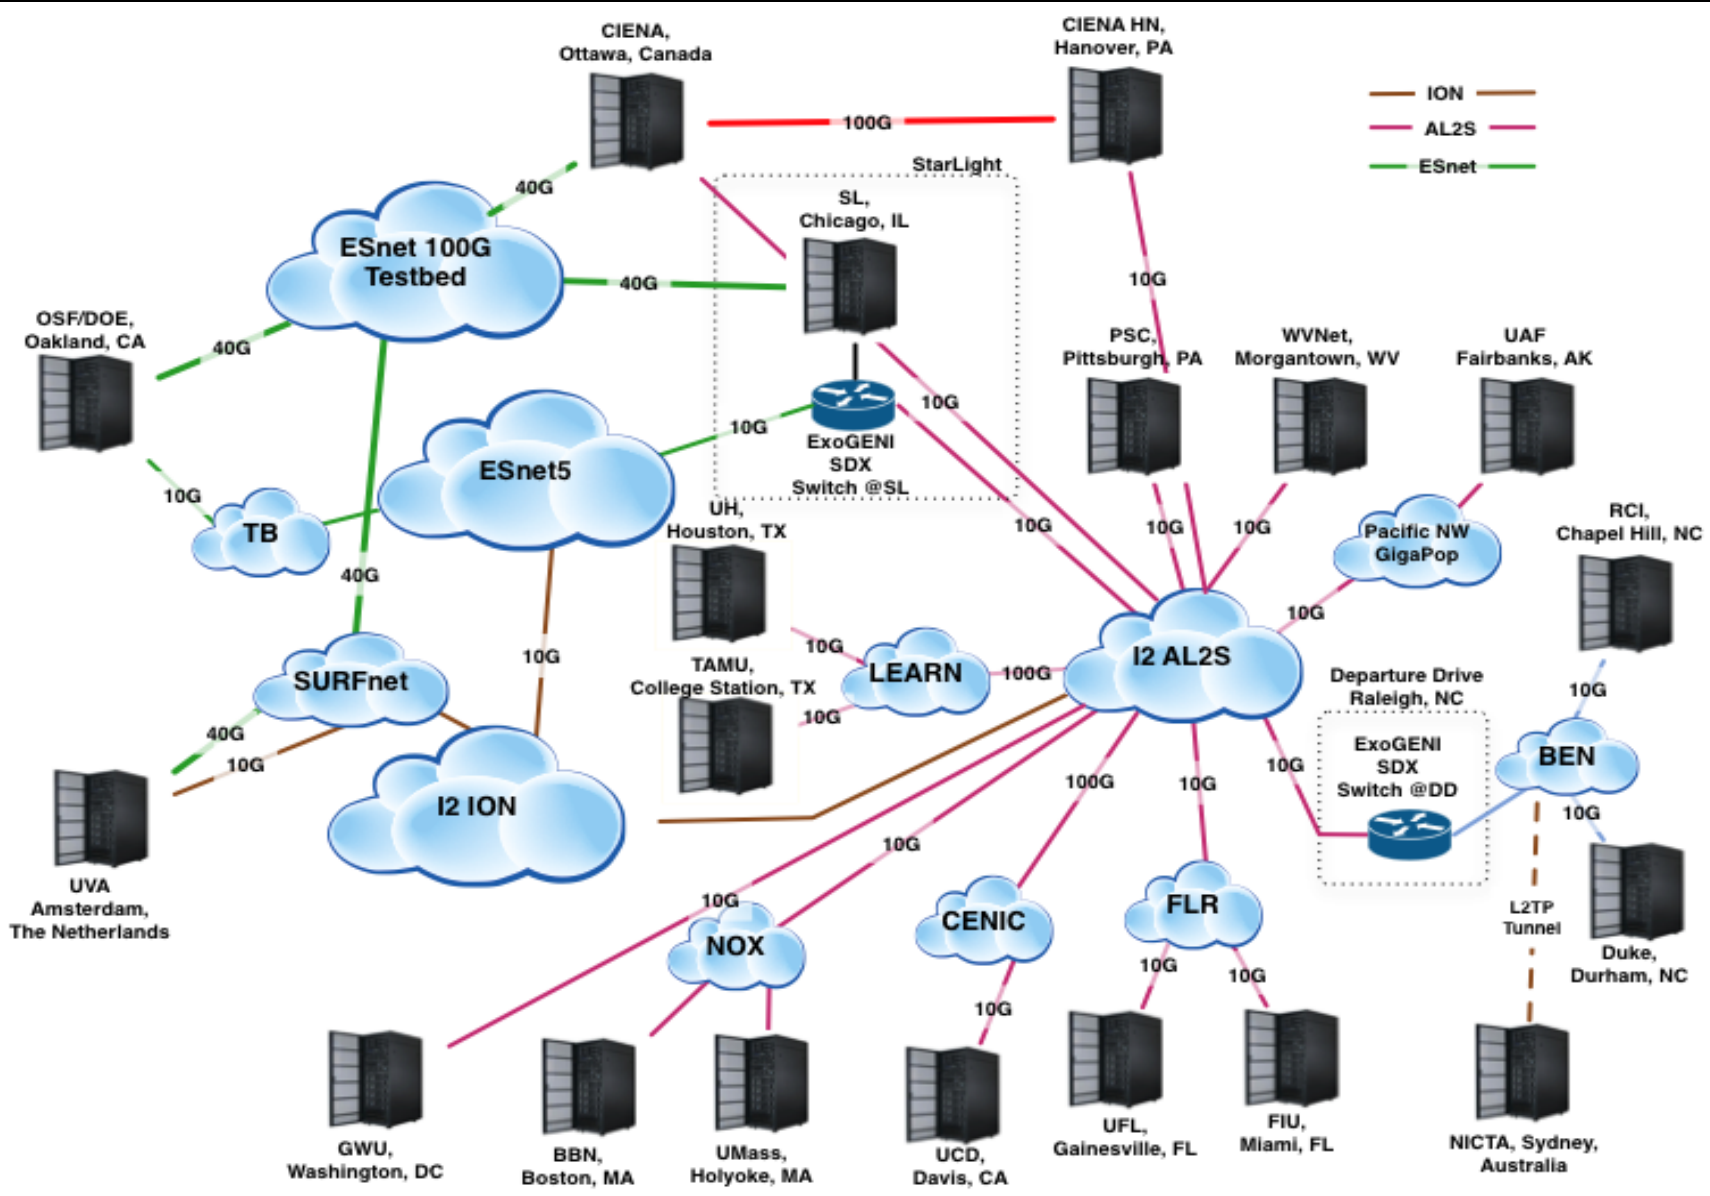
- **GENI**

- 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.





# International 40G and 100 G ExoGENI Testbed



App1

App2

App3

App4

EP1

EP2

Ind1

Ind2

APIs Based On Messaging and Signaling Protocols  
Network Programming Languages

Process Based Virtualization – Multi-Domain Federation –  
Policies Cascading Through Architectural Components

Security Processes

Policy Processes

Orchestrator(s)

Policy Processes

Northbound Interface

Network OSs  
SDN Control Systems

State Data Bases

State Machines

Mon, Measurements  
Real Time Analytics

Network Hypervisors

Westbound Interfaces

Eastbound Interfaces

Southbound Interface

PhyR

PhyR

PhyR

PhyR

VirR

VirR

VirR

VirR



# An Experimental Testbed For Computer Science Research

[www.chameleoncloud.org](http://www.chameleoncloud.org)

CHAMELEON:  
A LARGE-SCALE, RECONFIGURABLE EXPERIMENTAL  
ENVIRONMENT FOR CLOUD RESEARCH

Principal Investigator: Kate Keahey

Co-PIs: J. Mambretti, D.K. Panda, P. Rad, W. Smith, D. Stanzione

AUGUST 29, 2014



STARLIGHT<sup>SM</sup>

# TESTBED TO SUPPORT THE COMMUNITY'S RESEARCH CHALLENGES

*The community builds the testbed,  
and afterwards the testbed will shape the  
community*

## Big Data

Data volume,  
velocity and  
variety

## Big Compute

A wide range of  
data analytics

Programmable networks  
cheap, ubiquitous sensors  
and other emergent trends

## Big Instruments

Cyber-Physical  
Systems,  
Observatories

- **Build the right testbed**
- **Make the environment**

- **Reach the right community**
- **Have the right team**

# CHAMELEON: A POWERFUL AND FLEXIBLE EXPERIMENTAL INSTRUMENT

- ▶ Large-scale
  - ▶ Targeting Big Data, Big Compute, Big Instrument research
  - ▶ Over 650 nodes, 5 PB disk, 100G network
- ▶ Reconfigurable
  - ▶ Bare metal reconfiguration, single instrument, graduated approach for ease-of-use
- ▶ Connected
  - ▶ Workload and Trace Archive, partners with production clouds
- ▶ Complementary
  - ▶ Complementing GENI, Comet, Wrangler, XSEDE
  - ▶ Partnering with GENI, Grid'5000, OCC, et al (FIRE?)
- ▶ Sustainable
  - ▶ Strong industry connections

# RESEARCH COMMUNITIES AND CAPABILITIES

## Users

New models, algorithms, platforms, auto-scaling HA, etc.,  
Application and educational uses

*Persistent, reliable, shared cloud*

## Core Researchers and Users

Repeatable experiments in new models, algorithms,  
platforms, auto-scaling, HA, etc.

*Isolated partition, pre-configured images reconfiguration*

## Core Researchers

Virtualization technology (SR-IOV, accelerators, etc.)  
Infrastructure-level resource management

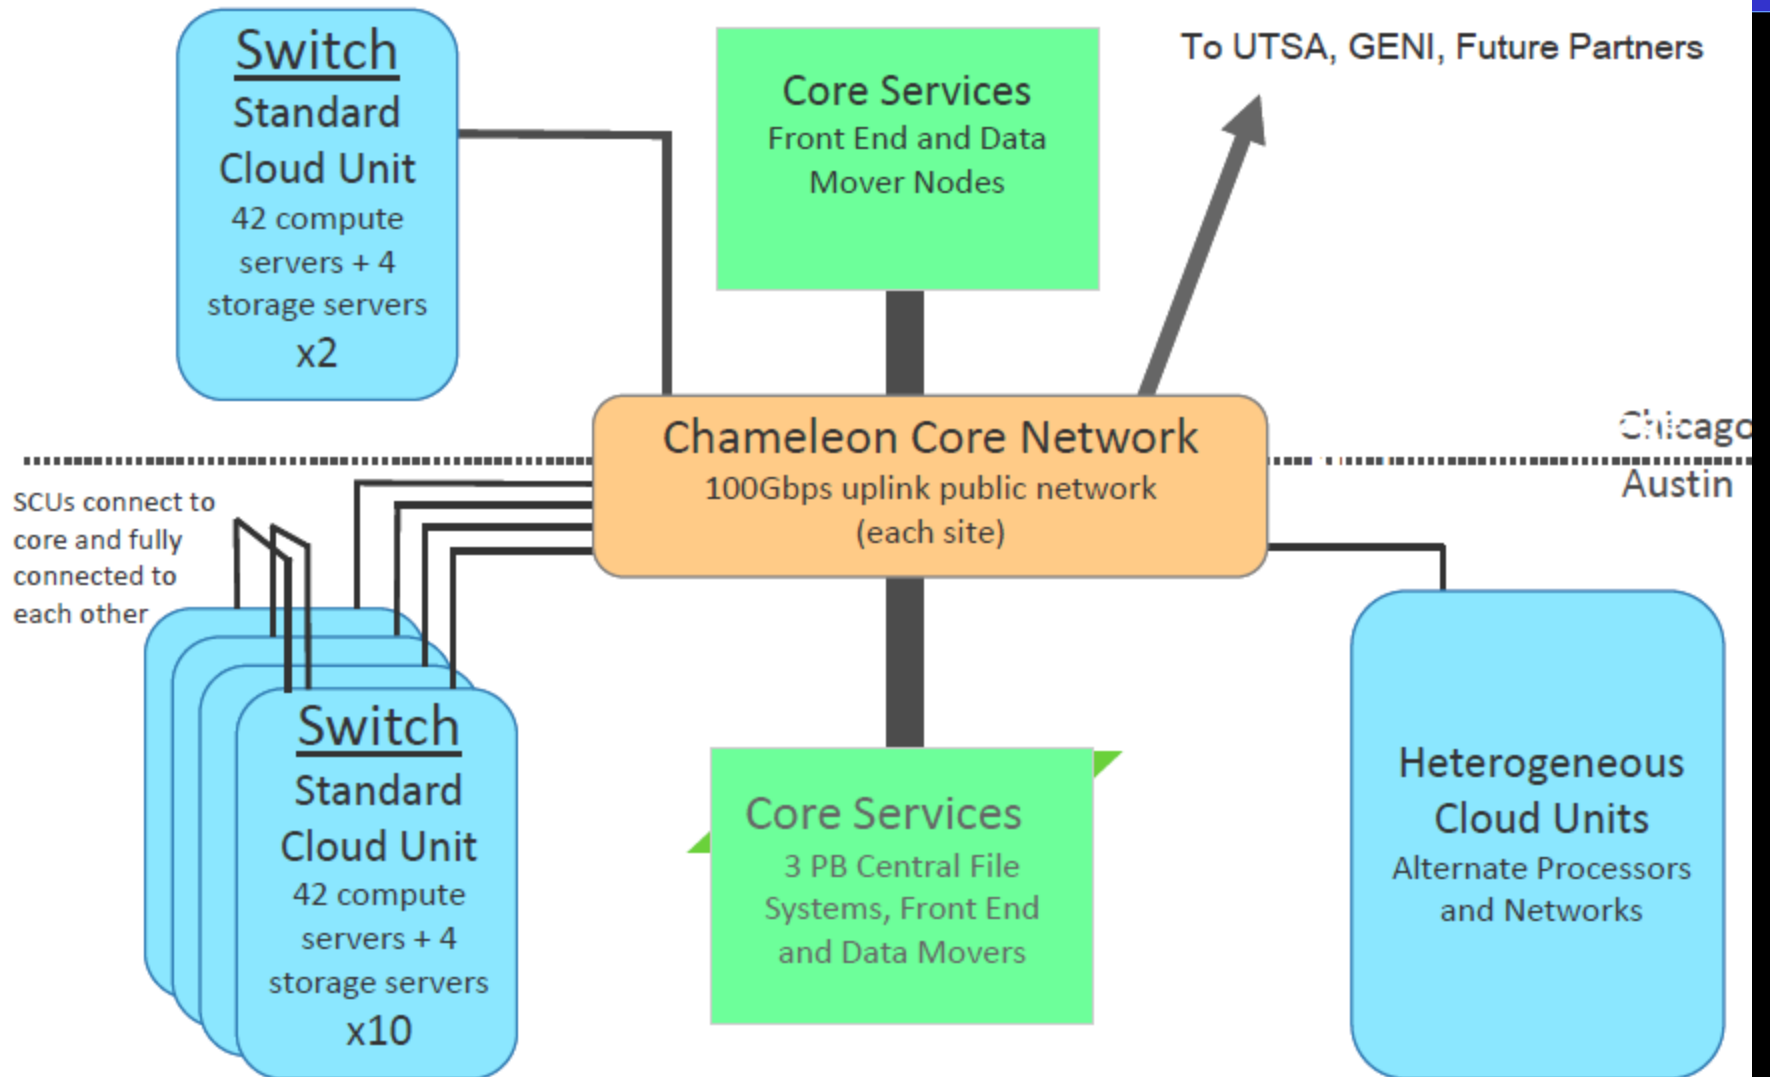
*Isolated partition, full bare metal reconfiguration*



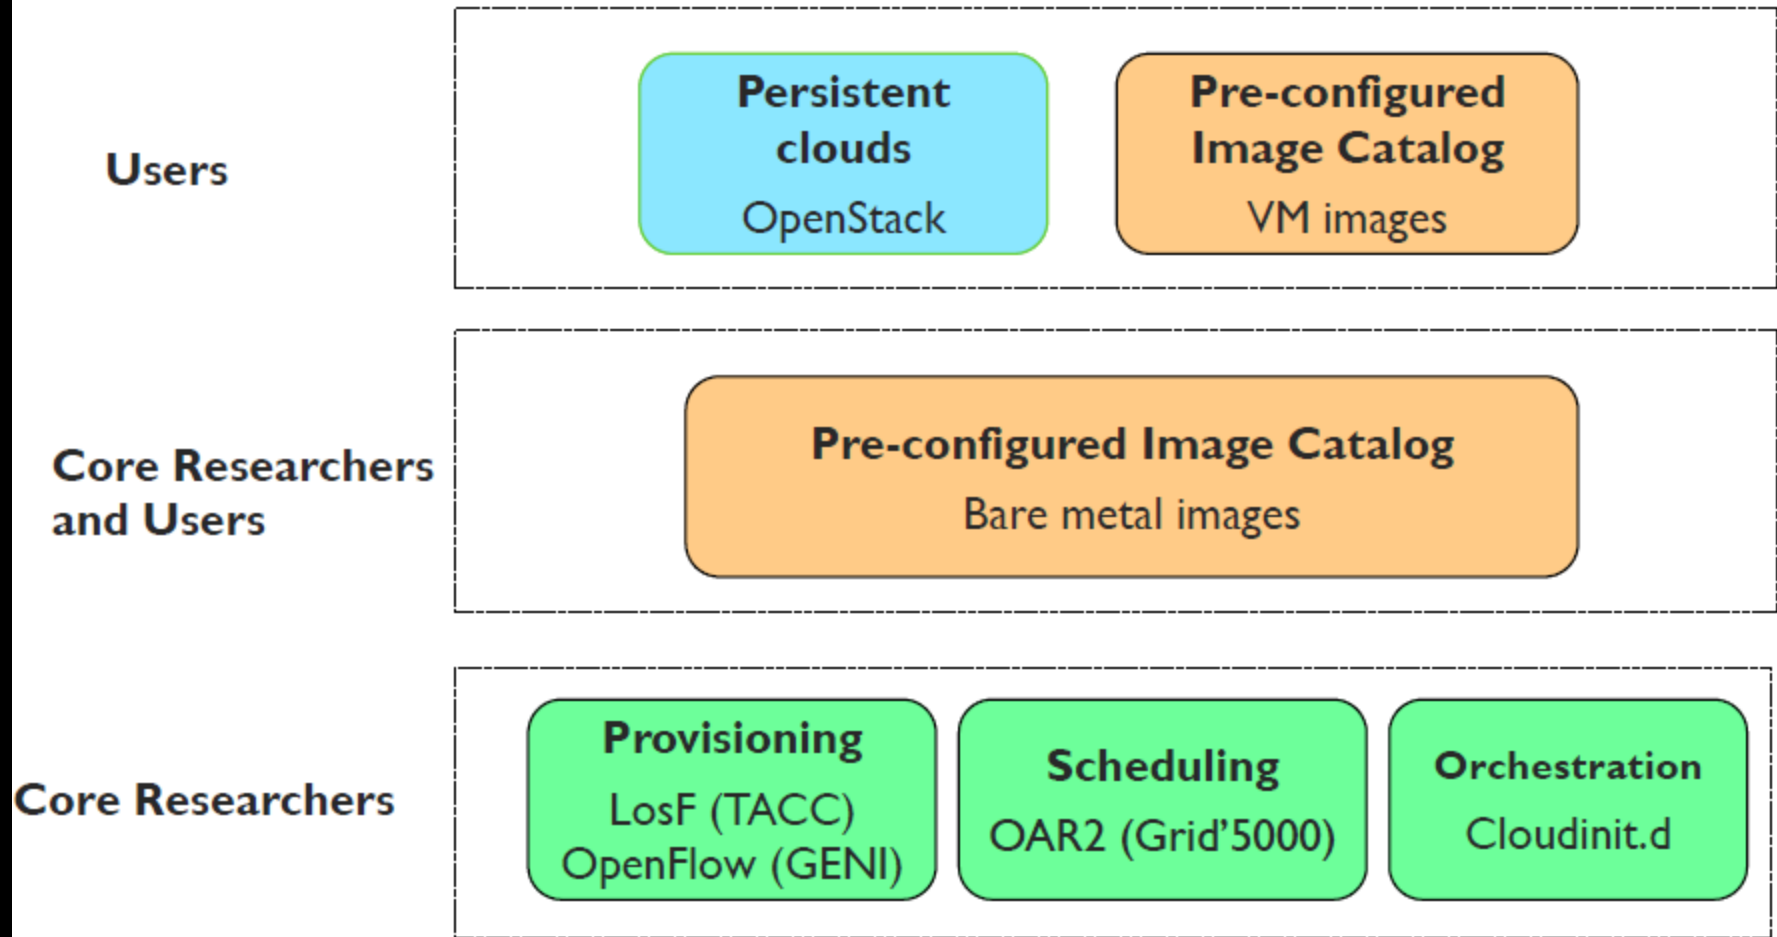
# SUPPORTED APPLICATIONS

- ▶ CPS
  - ▶ Offloading, multi-criteria trade-off analysis (response time vs cost), auto-scaling, high availability, etc.
- ▶ Machine learning, data mining
  - ▶ Mix of Big Compute and Big Data simulations and models, design of novel data processing frameworks
- ▶ System Software/Virtualization
  - ▶ Hypervisors optimizing a range of qualities, SR-IOV, virtualizing accelerators, etc.
- ▶ Networking
  - ▶ Programmable networks & QoS, refinement and effects of SR-IOV, large dataflows, end-to-end QoS

# ARCHITECTURE



# SYSTEM SOFTWARE: CORE CAPABILITIES



# CHAMELEON SERVICES AND FEATURES

- ▶ User Services
  - ▶ Allocation management through reservations, automatic image deployment
  - ▶ Dedicated Web portal for reservations, docs, stats, etc.
- ▶ Experiment Support
  - ▶ Trace and Workload Archive
  - ▶ Experiment enhancement (e.g., load generators)
- ▶ Additional Features
  - ▶ Reconfigurable, connected instrument
  - ▶ Development-focused approach
  - ▶ Ease-of- use: one stop shopping for experimental needs
  - ▶ Distinct from off-the-shelf cloud services
  - ▶ Code ownership and collaboration

# NETWORKING CAPABILITIES

- ▶ Expose SDN, OpenFlow, etc. to users
  - ▶ Isolation
  - ▶ Hybrid Network Capabilities
  - ▶ Programmable Topologies
  - ▶ Integration With Other Resources Within and External to the Testbed
- ▶ Pushing 100G Networks To Their Limit
  - ▶ Using 100G + SDN Optimally
  - ▶ Chameleon appliances and services allow experimenters a highly granulated view into -- and control -- over traffic flows
- ▶ Integration/Federation with GENI
- ▶ Common Policy Context

# Federation Among Multiple International Testbeds (Federation-as-a-Service)

- **Pair-Wise Approaches Among Multiple Is A Challenge**
- **Centralized Approaches Do Not Scale**
- **Decentralized Processes Will Scale**
- **SDXs Are Key Resources For Federation**
- **Requirements**
  - **Multi-Domain**
  - **Multi-Resource**
  - **Discovery**
  - **Integration**
  - **Policy Based Resource Utilization**
  - **Flexibility**
  - **Selectivity (e.g. A ⇔ N ⇔ T ⇔ NOT X ⇔ Conditionally Z)**



# Future Internet Research and Experimentation

## What is FIRE?

### FIRE Research

Research

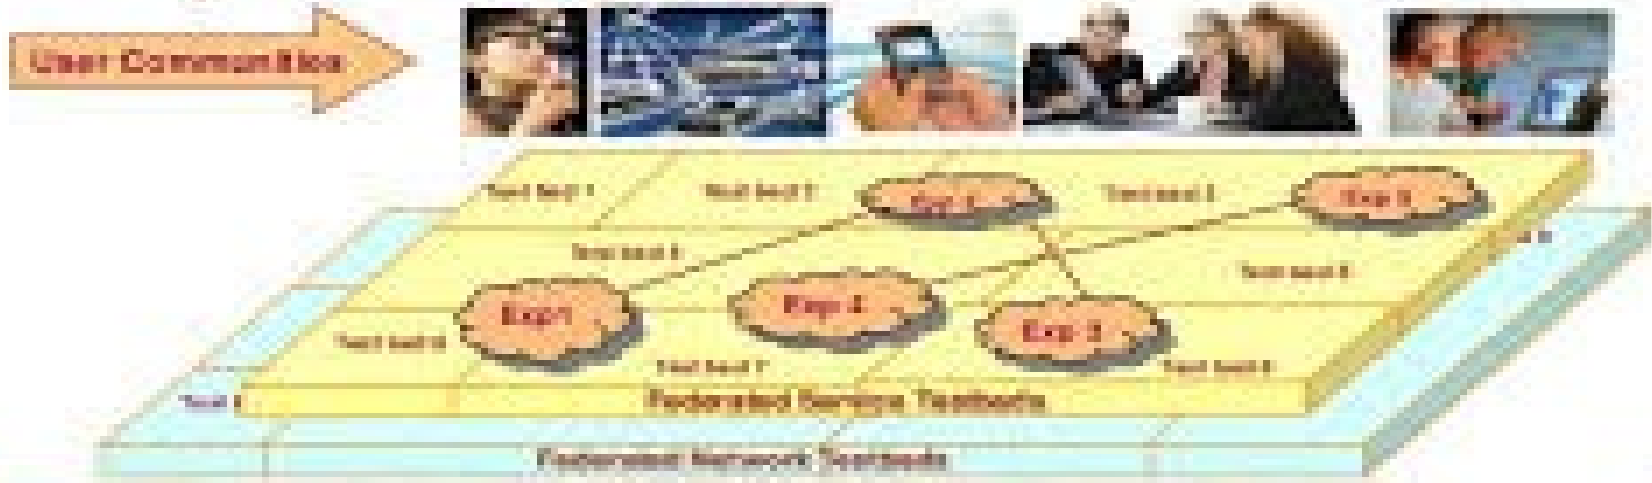


Validation

Large Scale Experiments

Requirements

### FIRE Experimental Facility



**FIRE ↔ FIRE Federation Project ↔ GENI ↔ Chameleon**



**FED4FIRE**

**STARLIGHT<sup>SM</sup>**





# iGENI: The International GENI

- **The iGENI Initiative Designed, Developed, Implemented, and Operated a Major 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 Operated for Use by GENI Researchers Conducting Experiments that Involve Multiple Aggregates At Multiple Sites.**
- **iGENI Infrastructure Has Connected Its Resources With Current GENI National Backbone Transport Resources, With Current and Planned GENI Regional Transport Resources, and With International Research Networks and Projects**





# StarLight International/National Communications Exchange Facility— “By Researchers For Researchers”

StarLight Is an Innovation Platform For Advanced Communications Services Architecture and Technologies, Including Experimental Testbeds Optimized For High-Performance Data Intensive Applications

Multiple  
10GE+100 Gbps  
Over Optics –  
World’s “Largest”  
10G/100G Exchange  
First of a Kind  
Enabling Interoperability  
At L1, L2, L3  
Also, StarWave  
Multi-100 Gbps Exchange



View from StarLight

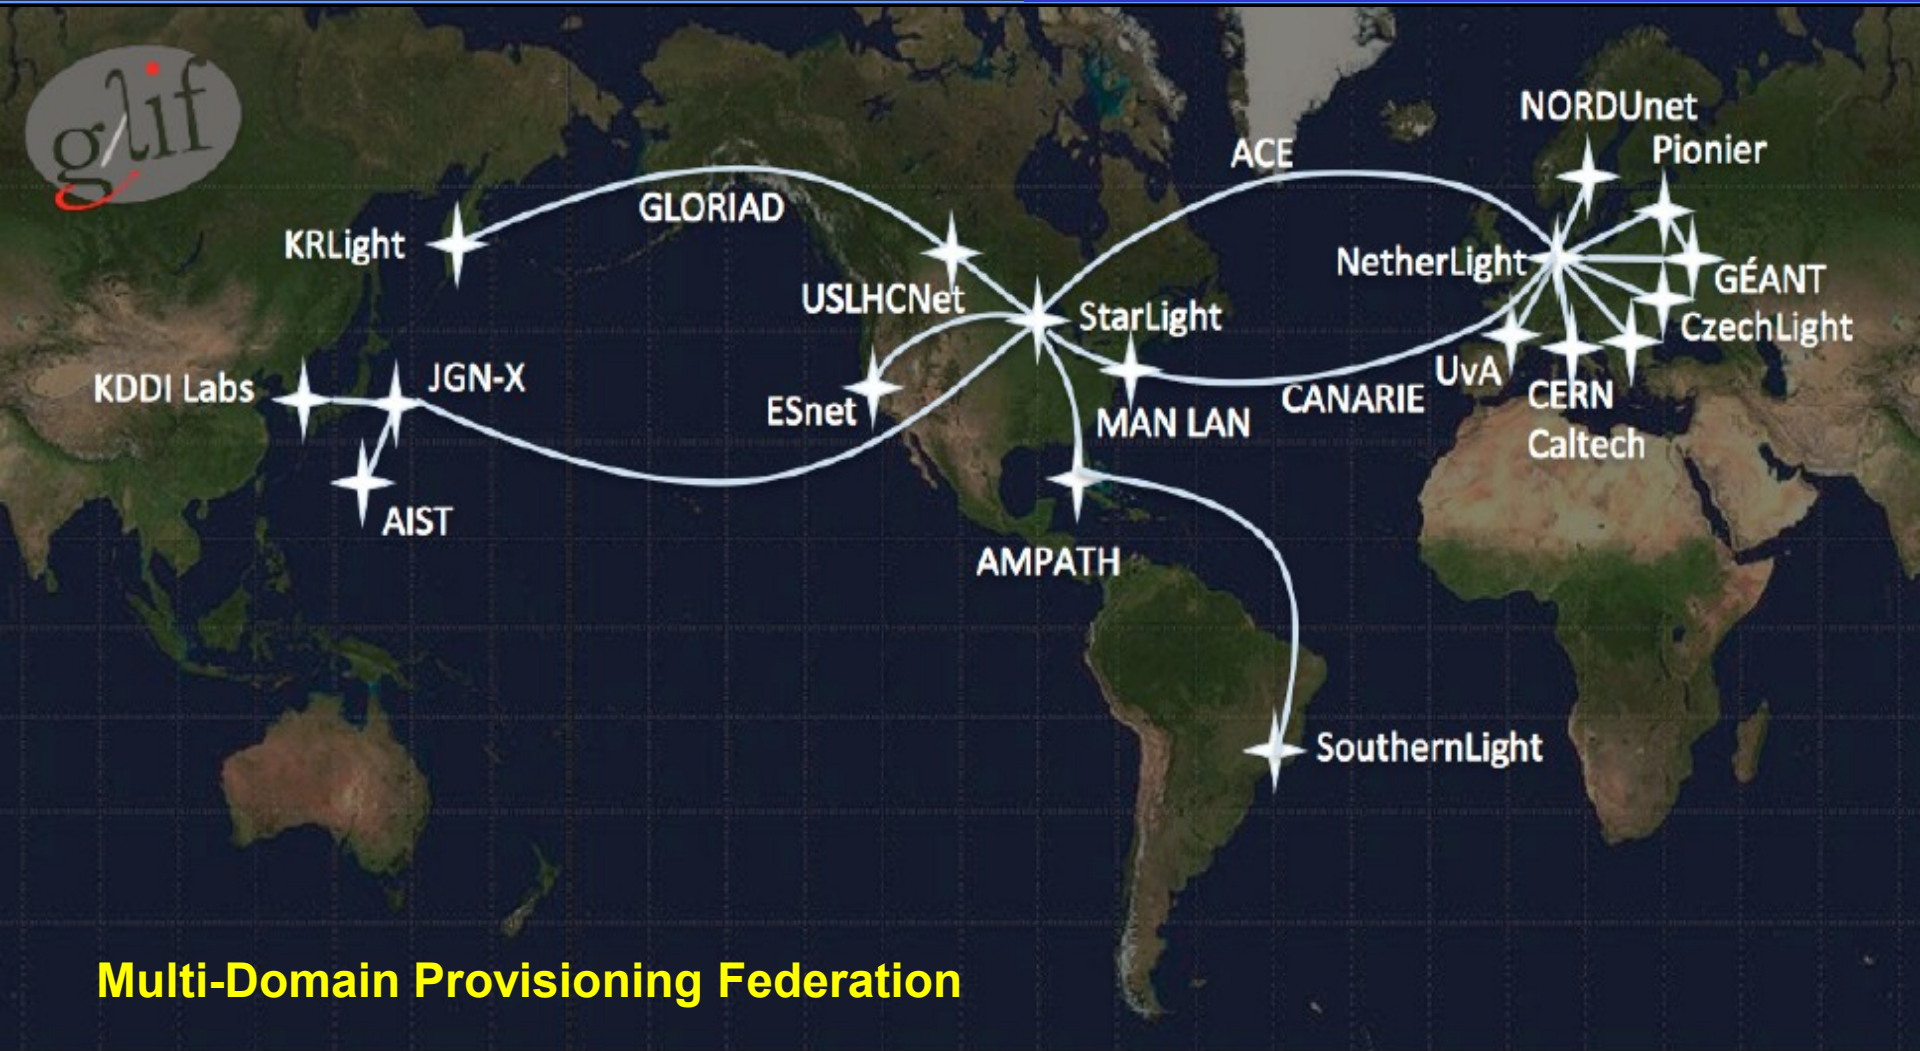


Abbott Hall, Northwestern University's Chicago Campus






# AutoGOLE Initiative



**Multi-Domain Provisioning Federation**

# AutoGOLE Dashboard (In Development)

Control Plane




Legend

peersWith



Data Plane



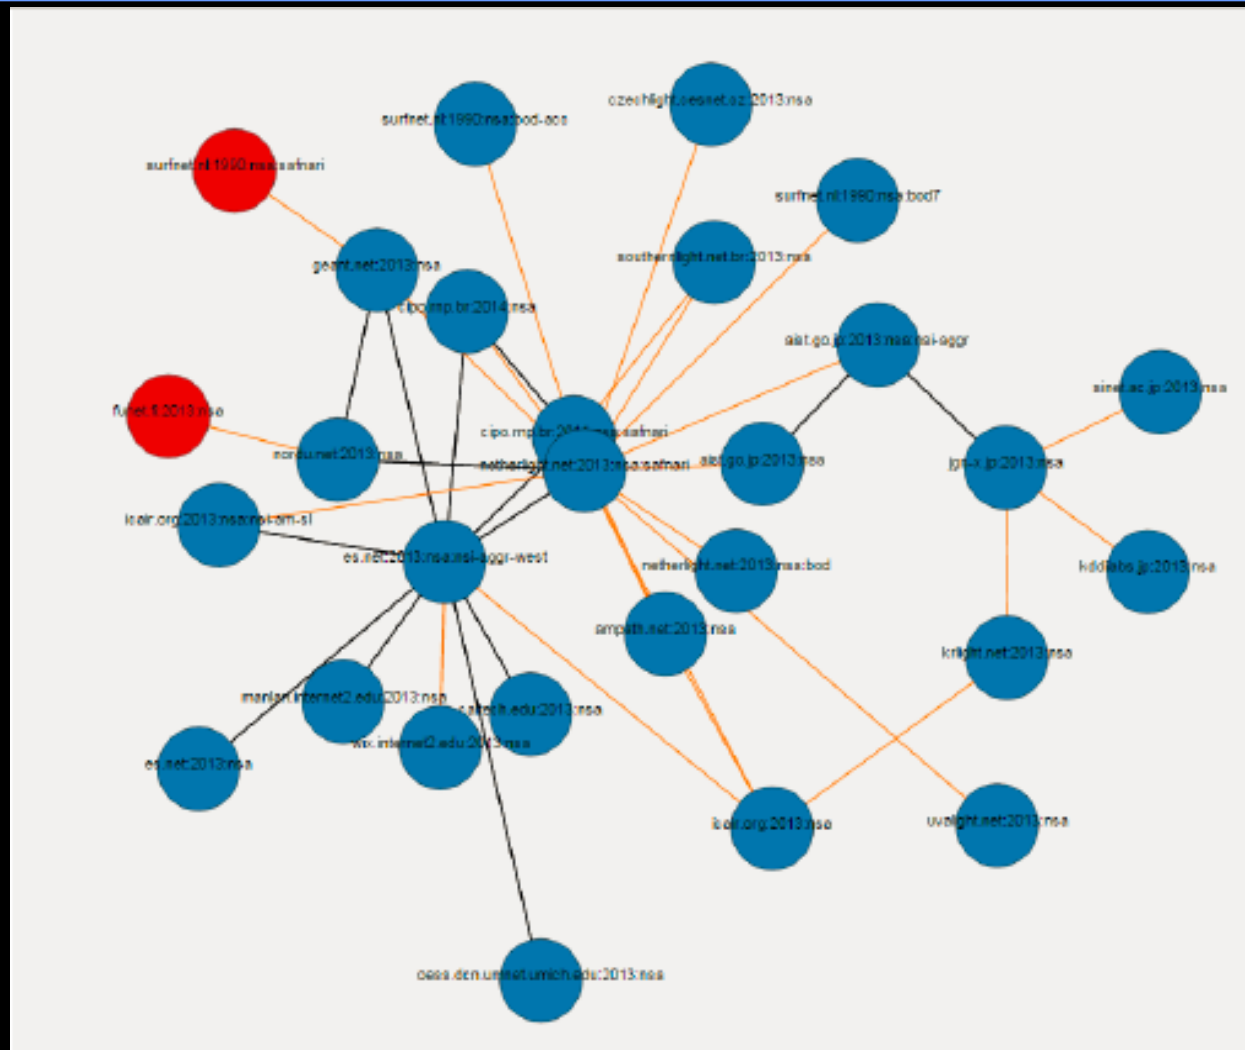
**isAlias mismatches**

'isAlias' in domain mismatches 'isAlias' in domain  
'isAlias' in domain mismatches 'isAlias' in domain  
'isAlias' in domain mismatches 'isAlias' in domain  
'isAlias' in domain mismatches 'isAlias' in domain

**NSA ID mismatches**

NSA ID X mismatches NSA ID Y  
NSA ID X mismatches NSA ID Y  
NSA ID X mismatches NSA ID Y

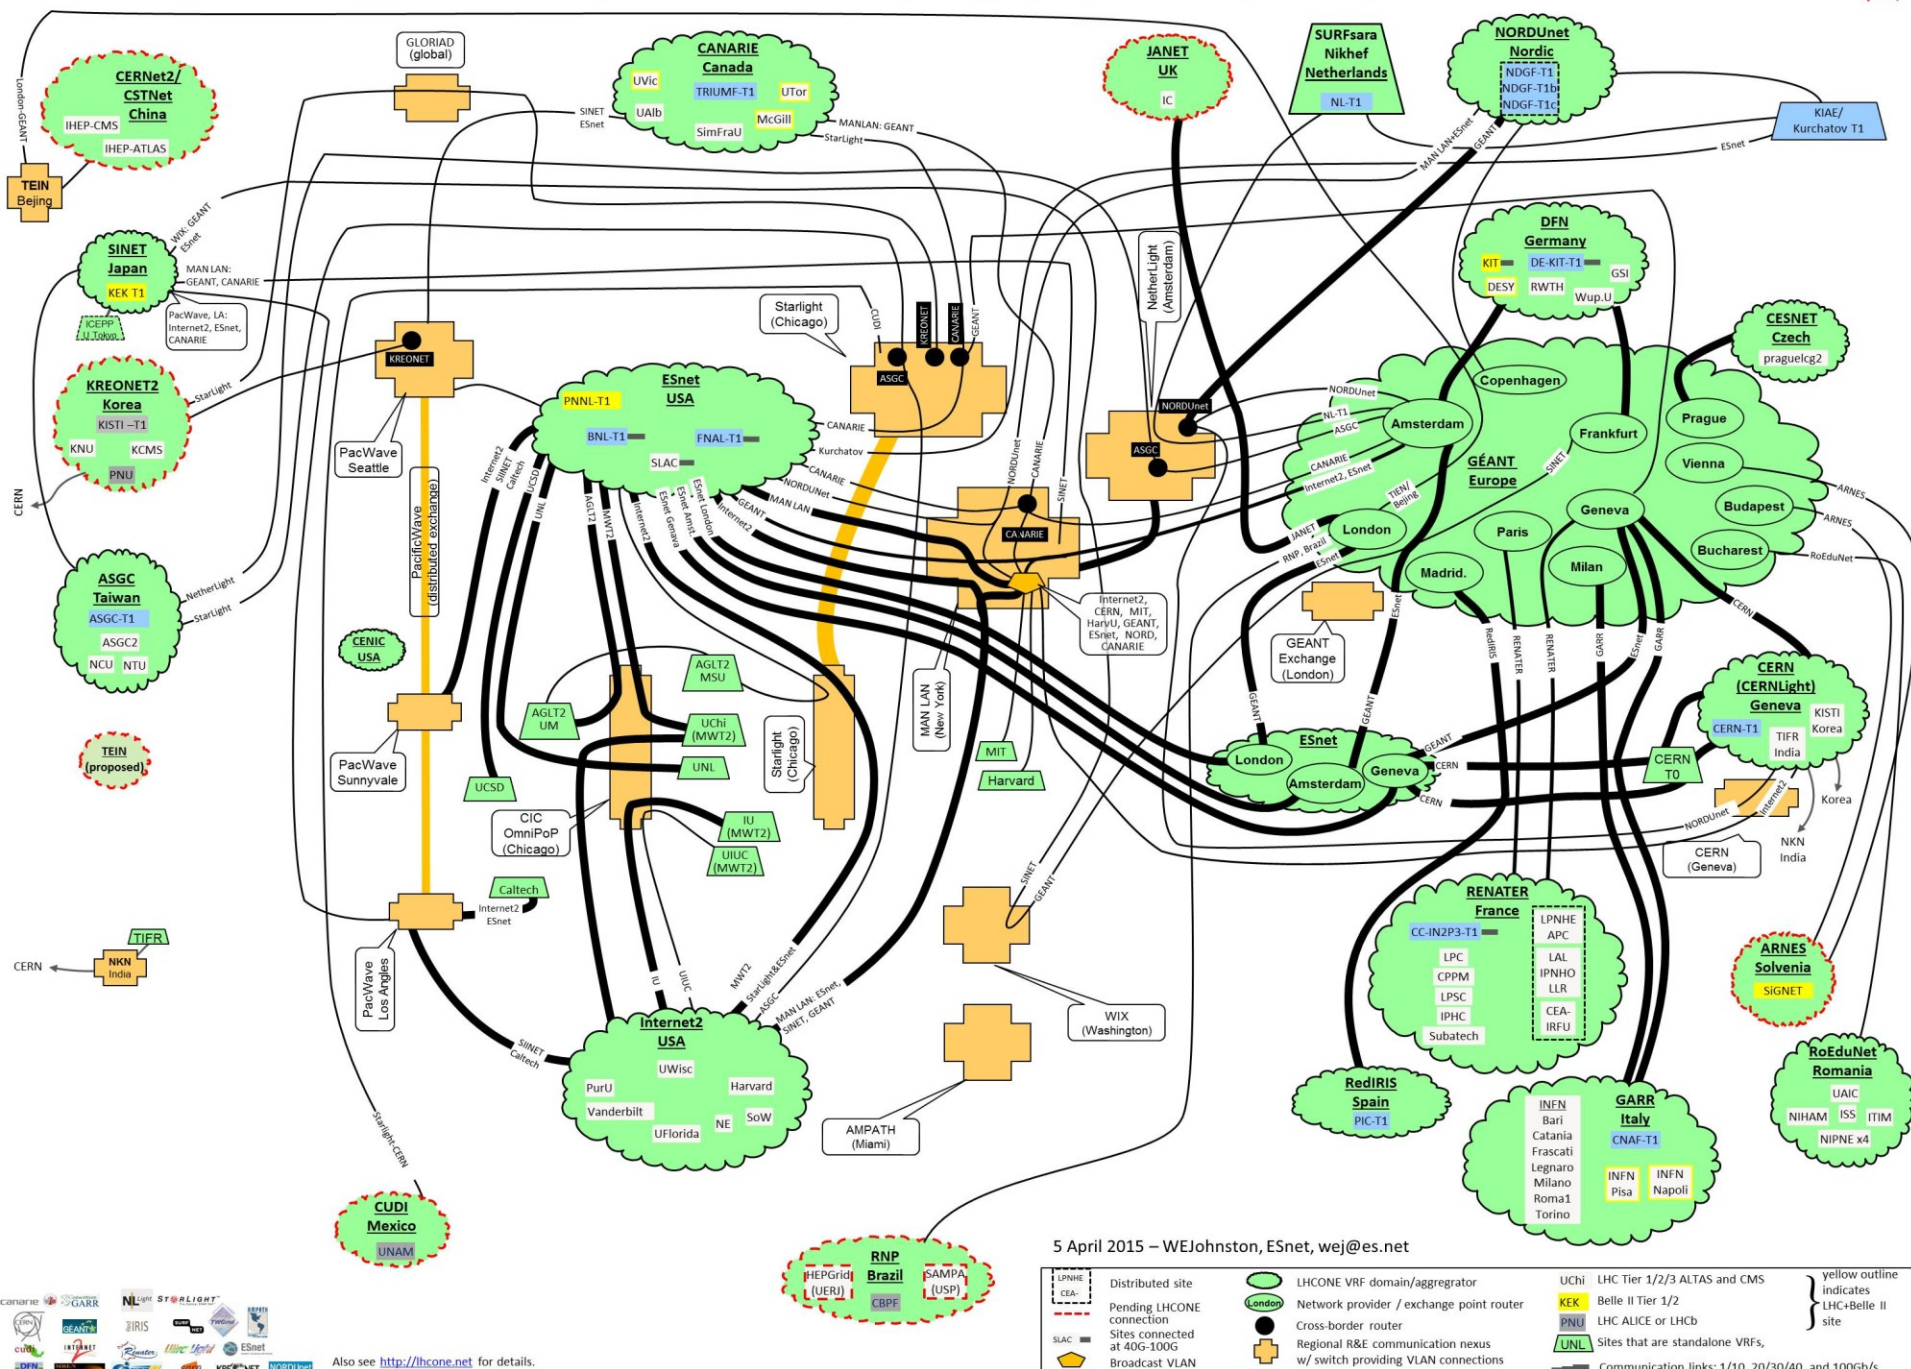
# Current Version of Dashboard







# LHCONE: A global infrastructure for the High Energy Physics (LHC and Belle II) data management



5 April 2015 – WEJohnston, ESnet, wej@es.net

LPNHE CEA	Distributed site	LHCONE VRF domain/aggregator	Uchi	LHC Tier 1/2/3 ATLAS and CMS
Pending LHCONE connection		London	KEK	Belle II Tier 1/2
SLAC	Sites connected at 40G-100G	Cross-border router	PNU	LHC ALICE or LHCb
Broadcast VLAN		Regional R&E communication nexus w/ switch providing VLAN connections	UNL	Sites that are standalone VRFs, yellow outline indicates LHC-Belle II site
			RENATER	Communication links: 1/10, 20/30/40, and 100Gb/s




Also see <http://lhcone.net> for details.


# Software Defined Networking Exchanges (SDXs)

- **With the Increasing Deployment of SDN In Production Networks, the Need for an SDN Exchange (SDX) Has Been Recognized.**
- **Many Motivations Exist for SDXs**
  - **Bridging SDNs (Which Are Single Domain & Centralized Controller Oriented)**
  - **Granulated Engineering Over Flows**
  - **High Degrees Of Exchange Customization**
- **Required: Capabilities for Multi-Domain Distributed SDN Resource Discovery, Signaling, Provisioning, Federation, Operational Functions, Fault Detection and Recovery**
- **These Are Fairly Challenging Issues**

# Selected SDX Architectural Attributes

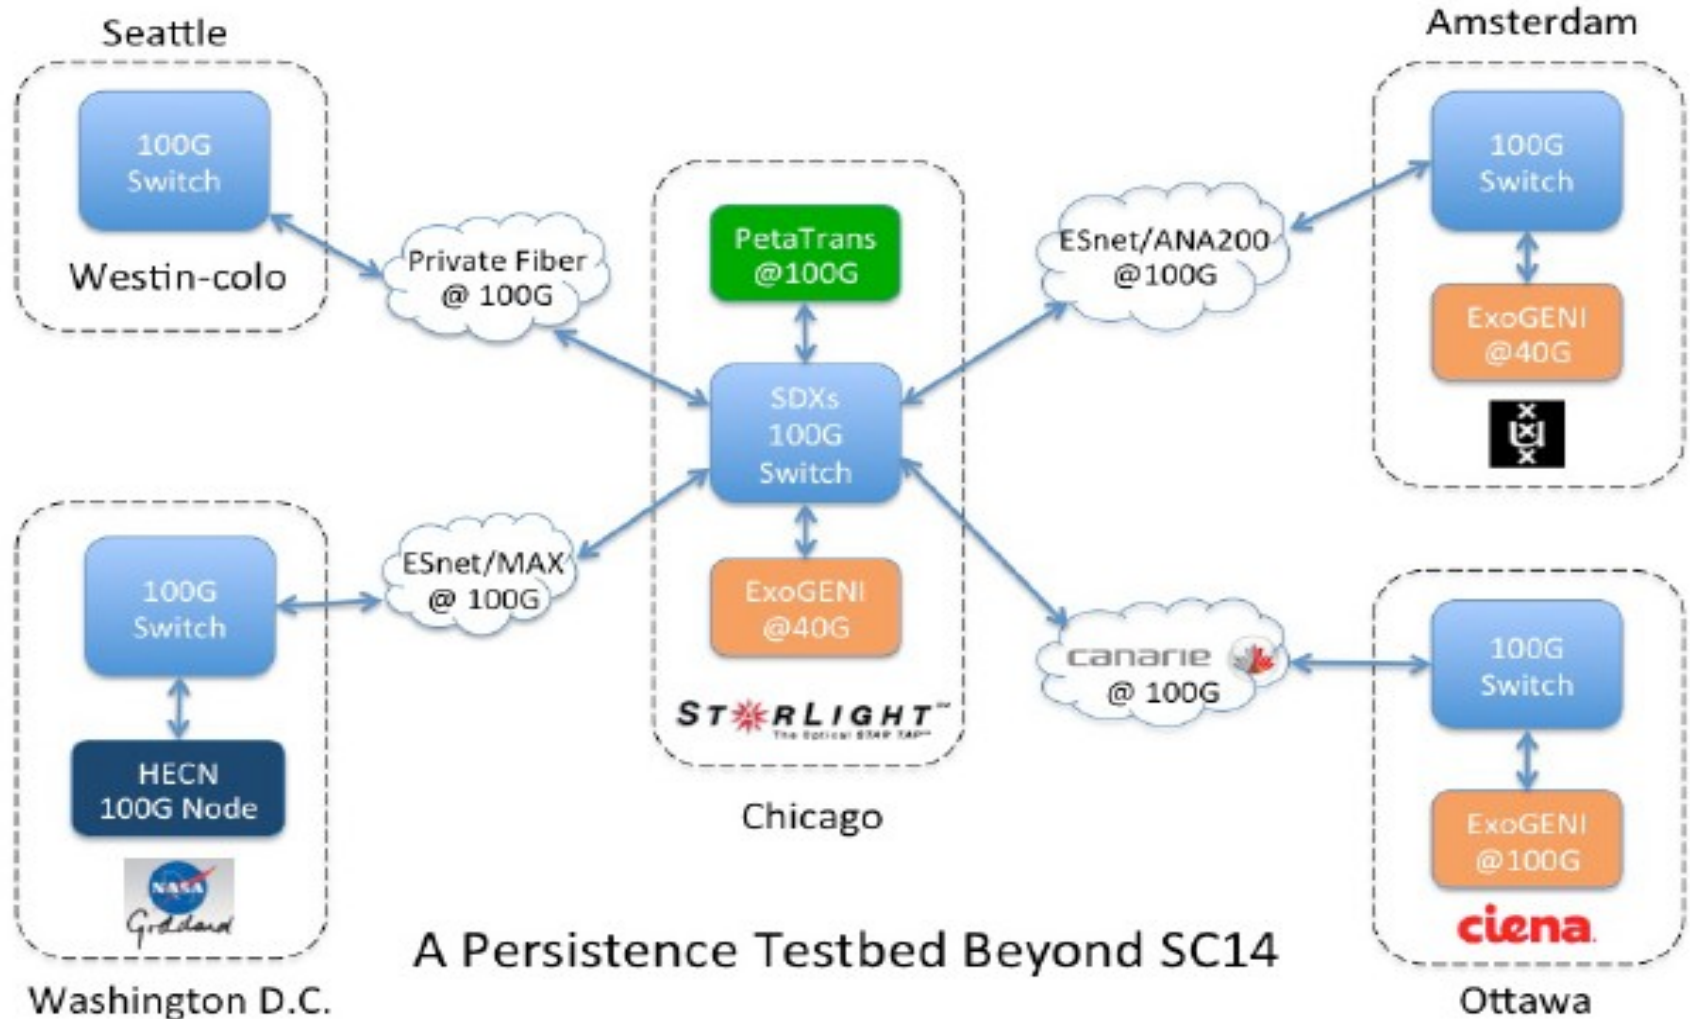
- **Control and Network Resource APIs**
- **Multi Domain Integrated Path Controllers (With Federation)**
- **Controller Signaling, Including Edge Signaling**
- **SDN/OF Multi Layer Traffic Exchange Services**
- **Multi Domain Resource Advertisement/Discovery**
- **Topology Exchange Services**
- **Multiple Highly Customized Services At All Layers**
- **Granulated Resource Access (Policy Based), Including Through Edge Processes, Including To individual Streams**
- **Foundation Resource Programmability**
- **Various Types of Gateways To Other Network Environments**
- **Integration of OF and Non-OF Paths, Including 3<sup>rd</sup> Party Integration**
- **Programmability for Large Scale Large Capacity Streams**





SL-SDX-0.3 Prototype Logical Diagram 05/2015

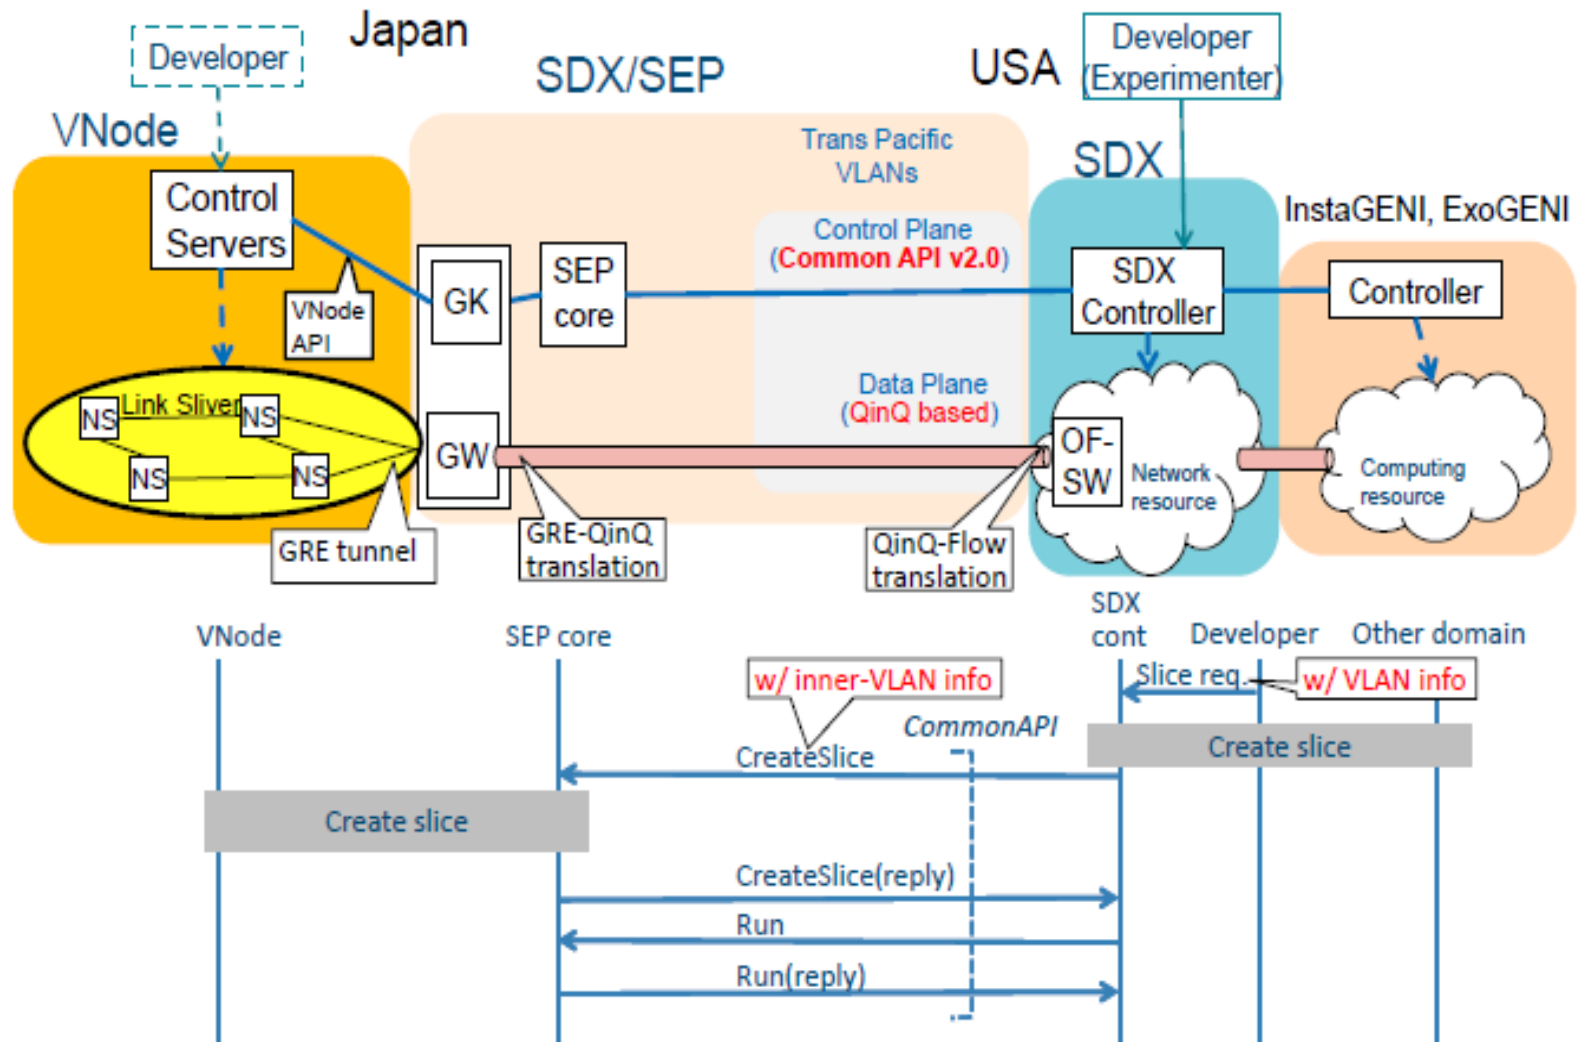
# PetaTrans: Petascale Science Data Transfer



Global Software-Defined Dynamic Circuits for Data Intensive Science  
(PhEDEx - ANSE - PANDA - OpenDayLight)

# vNode/SEP & StarLight Inter-SDX federation 2014

## Multi-architecture Federation



© KDDI R&D Laboratories Inc, HITACHI LTD, The University of Tokyo

# StarLight Software Defined Networking Exchange (SDX)



- **The StarLight SDX Will Provide The Services, Architecture, and Technologies Designed To Provide Scientists, Engineers, and Educators With Highly Advanced, Diverse, Reliable, Persistent, and Secure Networking Services, Enabling Them to Optimally Access Resources in North America, South America, Asia, South Asia (including India), Australia, New Zealand, Europe, the Middle East, North Africa, And Other Sites Around the World.**
- **The StarLight SDX Initiative Undertakes Continued innovation and Development of Advanced Networking Services and Technologies.**
- **Potential For Providing Federation-as-a-Service**

[www.startap.net/starlight](http://www.startap.net/starlight)



Thanks to the NSF, DOE, NIH, USGS, DARPA  
NOAA, Universities, National Labs,  
International Partners,  
and Other Supporters

