Blue Planet Framework In The Context Of Software Defined Infrastructure (SDI)

Joe Mambretti, Jim Chen
International Center for Advanced Internet Research (www.icair.org)
Northwestern University,
StarLight International/National Communications Exchange Facility
Rod Wilson, Marc Lyonnais, Gauravdeep Shami
Ciena Research Lab
Maxine Brown, Lance Long, Luc Renambot
Electronic Visualization Lab, University of Illinois Chicago
Linda Winkler
Argonne National Laboratory, StarLight International/National Communications Exchange Facility

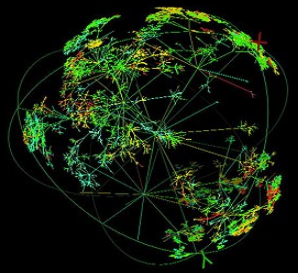

Global LambdaGrid Workshop
Miami, Florida
September 29-30, 2016

Revolution In New Communications Architecture

- Traditional Networking Architecture and Technology Are Designed For Supporting A Limited Number of Communication Services, Implemented With Multiple Restrictions
- Traditional Networking Services and Infrastructure Are Rigid/Static, e.g., Difficult To Expand, Customize And Enhance
- Consequently, ***A New Architectural Model Is Being Developed***
- ***Infrastructure As A Platform*** Is Replacing The Traditional Static/Rigid Network With a New Communication Services Foundation – a Highly Distributed Facility That Can Support Multiple Networks With Different Characteristics, Multiple Highly Differentiated Services, and Co, Dynamic Service Provisioning, Enhancements, Customization, Specialized Services, Real-Time Analytics And Adjustments

Paradigm Shift – Ubiquitous Services Based on Large Scale Distributed Facility vs Isolated Services Based on Separate Component Resources

Traditional Provider Services:
Invisible, Static Resources,
Centralized Management,
Highly Layered

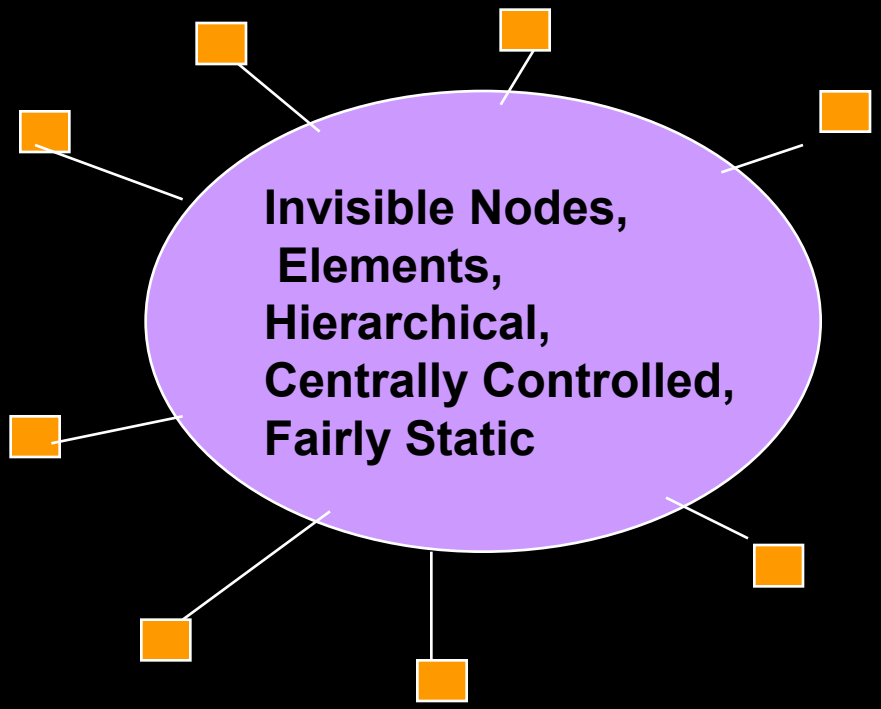

Limited Services, Functionality,
Flexibility, Expandability

Distributed Programmable Resources,
Dynamic Services,
Visible & Accessible Resources,
Integrated As Required, Non-Layered

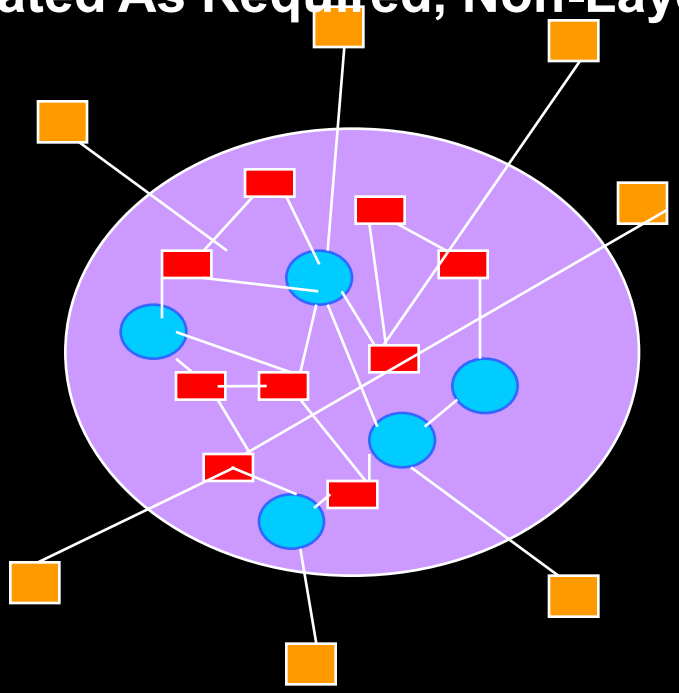

Unlimited Services, Functionality,
Flexibility, Expandability

A New Architectural Model: Software Defined Infrastructure (SDI)

- One Precursor: Programmable Grid Infrastructure
- Current: Software Defined Networking (SDN), Software Defined Computing (SDC), Software Defined Storage (SDS), Software Defined Infrastructure (SDI), Software Defined Everything (SDE)
- A Fundamentally New Architecture Is Required To Address The Issues That Arise From SDE
- Progress Is Being Made On Developing The Required Architecture
- A Special Focus For Our Community Is Meeting the Networking Requirements Of Data Intensive Science

GriPhyN
Data Intensive Science

NEPTUNE
A Fiber-Optic Telescope to Inner Space

NEES

GENOMES to LIFE
BIOLOGICAL SOLUTIONS FOR ENERGY CHALLENGES
U.S. DEPARTMENT OF ENERGY

NG Digital Sky Survey

NEON
National Ecological Observatory Network

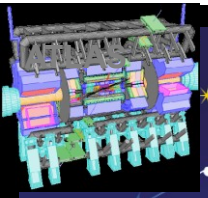

ALMA
ALMA: Atacama Large Millimeter Array

ANDRILL: Antarctic Geological Drilling
www.andrill.org

BIRN
BIRN: Biomedical Informatics Research Network
www.nbirn.net

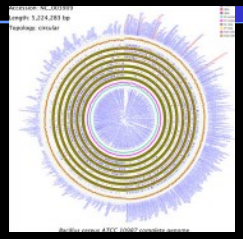

CAMERA metagenomics
camera.calit2.net

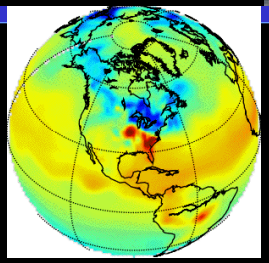

Carbon Tracker
www.esrl.noaa.gov/gmd/ccgg/carbontracker

CineGrid
CineGrid
www.cinegrid.org

LHCONE
www.lhcone.net

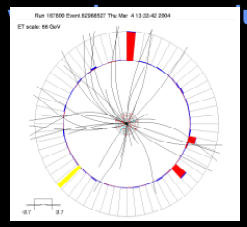

DØ (DZero)
www-d0.fnal.gov

GEON
GEON: Geosciences Network
www.geongrid.org

GLEON: Global Lake Ecological Observatory Network

OOI
OCEAN OBSERVATORIES INITIATIVE
CYBERINFRASTRUCTURE
Providing a link between ocean research and discovery
OOI-CI
ci.oceanobservatories.org

ISS: International Space Station
www.nasa.gov/station

CLASS
Comprehensive Large-Array Stewardship System
www.class.noaa.gov

LIGO
www.ligo.org

WLCG
Worldwide LHC Computing Grid
WLCG
lcg.web.cern.ch/LCG/public/

PRAGMA
Pacific Rim Applications and Grid Middleware Assembly
www.pragma-grid.net

TeraGrid
TeraGrid
www.teragrid.org

IVOA
IVOA: International Virtual Observatory
www.ivoa.net

Open Science Grid
OSG
www.opensciencegrid.org

the globus alliance
Globus Alliance
www.globus.org

SKA
SQUARE KILOMETRE ARRAY
SKA
www.skatelescope.org

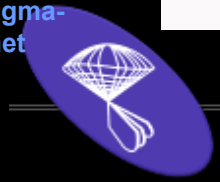

Sloan Digital Sky Survey
www.sdss.org

XSEDE
XSEDE
www.xsede.org

STARLIGHT

Compilation By Maxine Brown

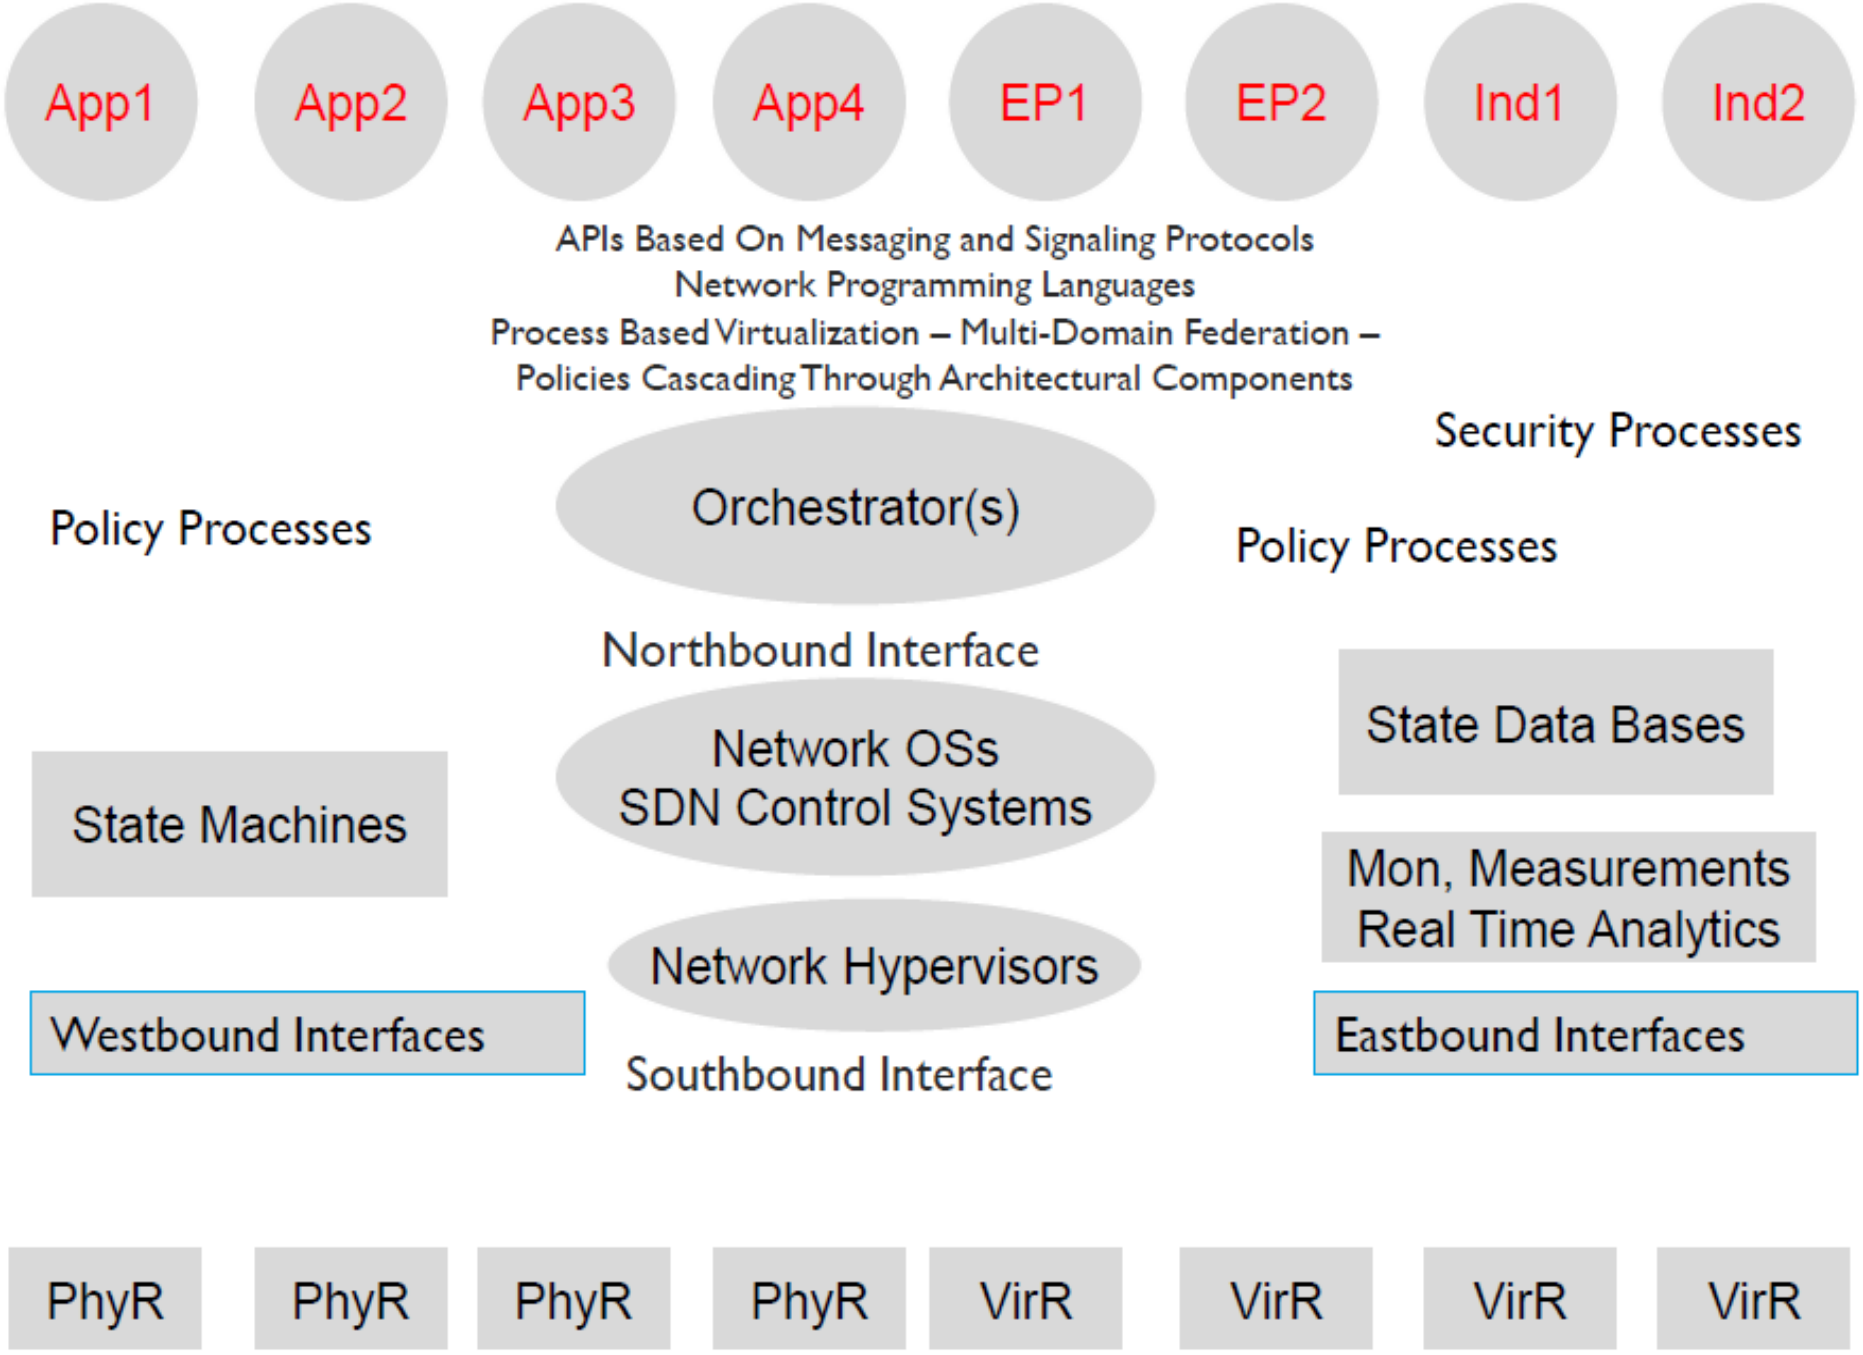
App1
App2
App3
App4
EP1
EP2
Ind1
Ind2
APIs Based On Messaging and Signaling Protocols
Network Programming Languages
Process Based Virtualization – Multi-Domain Federation –
Policies Cascading Through Architectural Components
Security Processes
Orchestrator(s)
Policy Processes
Policy Processes
Northbound Interface
State Data Bases
Network OSs
SDN Control Systems
State Machines
Mon, Measurements
Real Time Analytics
Network Hypervisors
Westbound Interfaces
Eastbound Interfaces
Southbound Interface
PhyR
PhyR
PhyR
PhyR
VirR
VirR
VirR
VirR

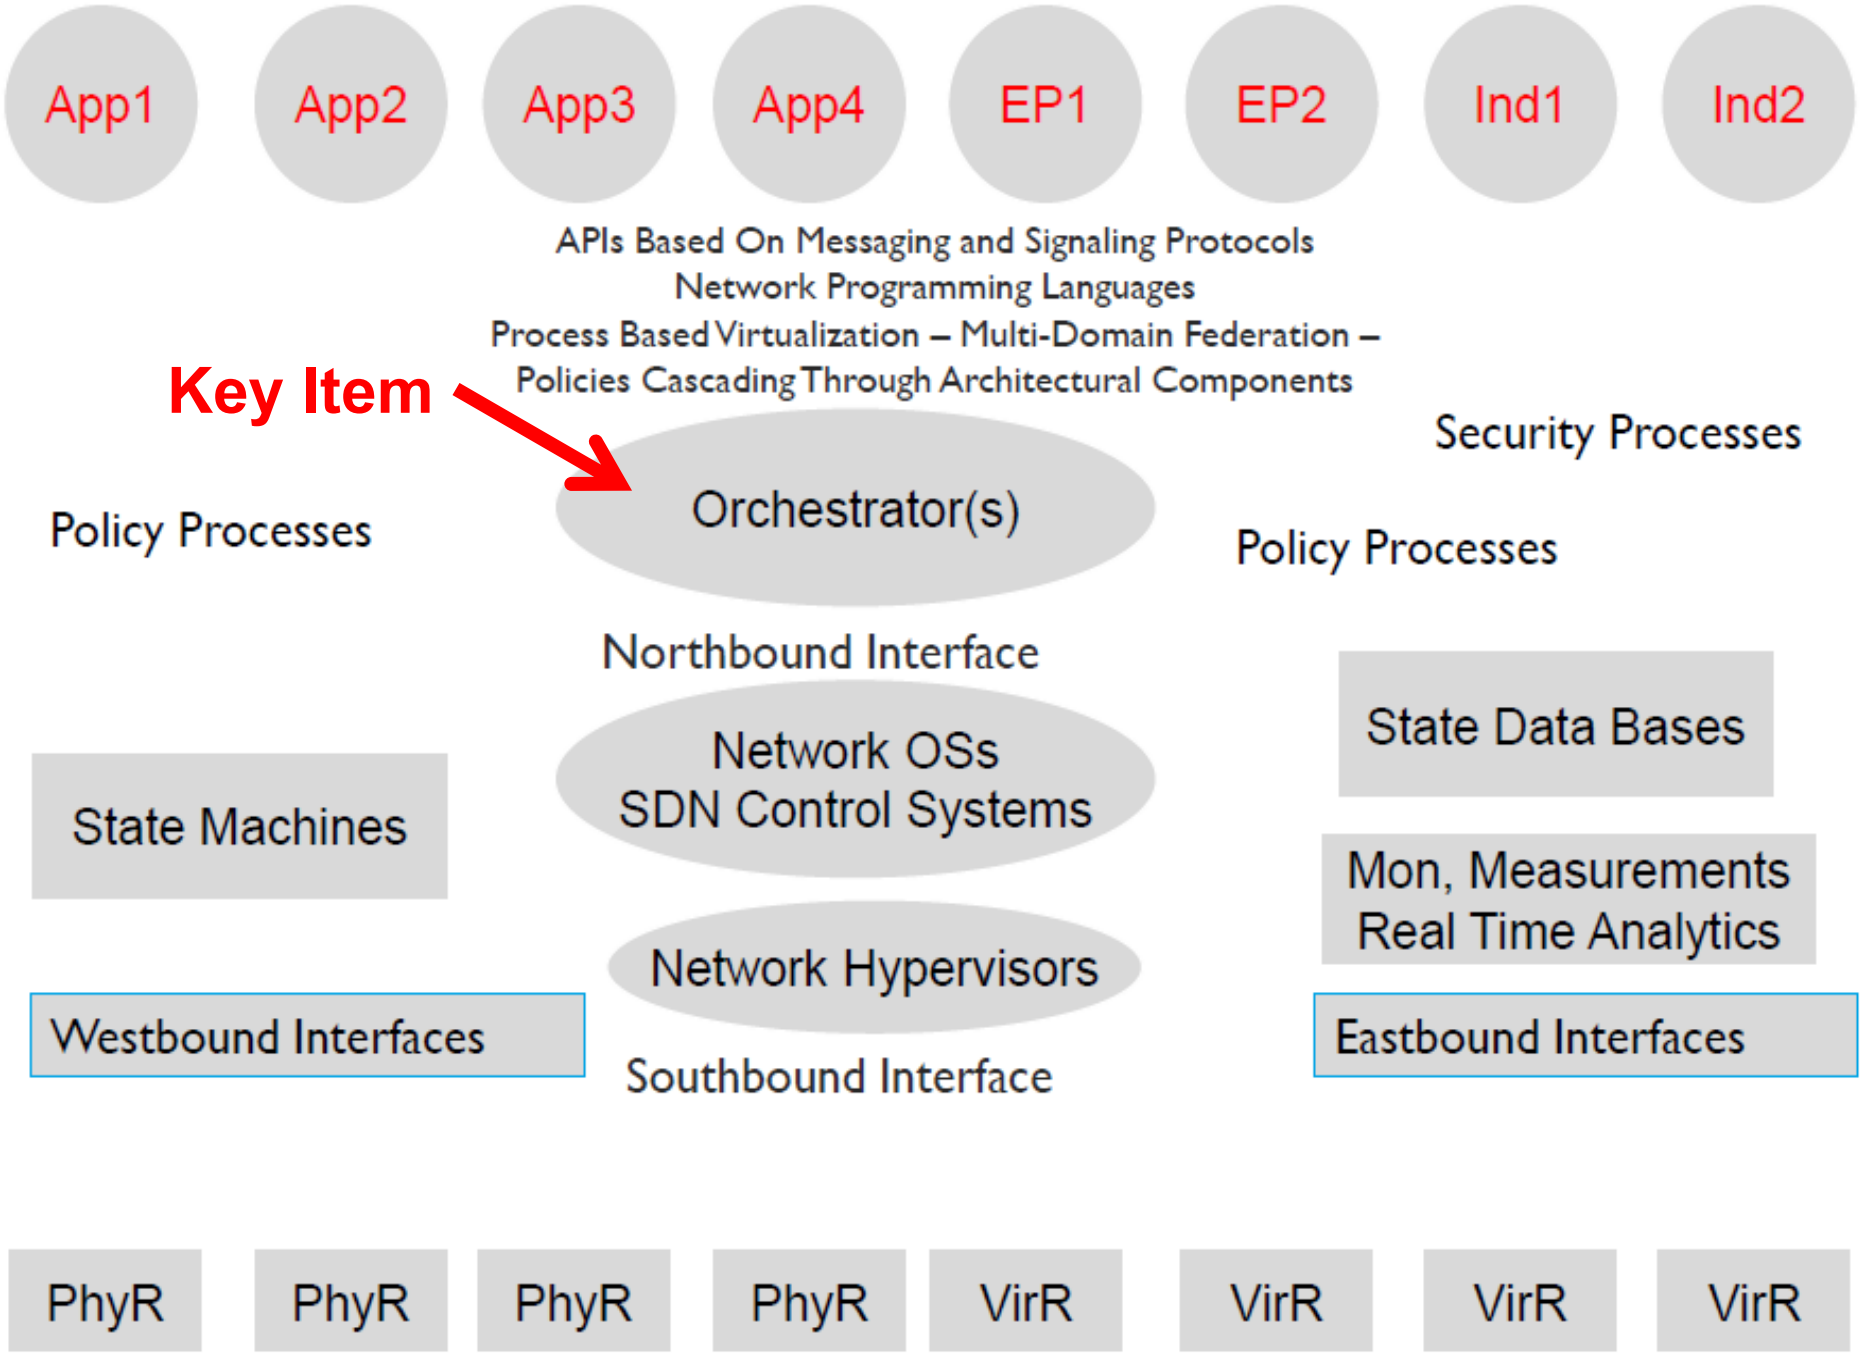
App1
App2
App3
App4
EP1
EP2
Ind1
Ind2
APIs Based On Messaging and Signaling Protocols
Network Programming Languages
Process Based Virtualization – Multi-Domain Federation –
Policies Cascading Through Architectural Components
Key Item
Security Processes
Policy Processes
Orchestrator(s)
Policy Processes
Northbound Interface
State Data Bases
Network OSs
SDN Control Systems
State Machines
Mon, Measurements
Real Time Analytics
Network Hypervisors
Westbound Interfaces
Eastbound Interfaces
Southbound Interface
PhyR
PhyR
PhyR
PhyR
VirR
VirR
VirR
VirR

Opendaylight 4th Release: Baryllium

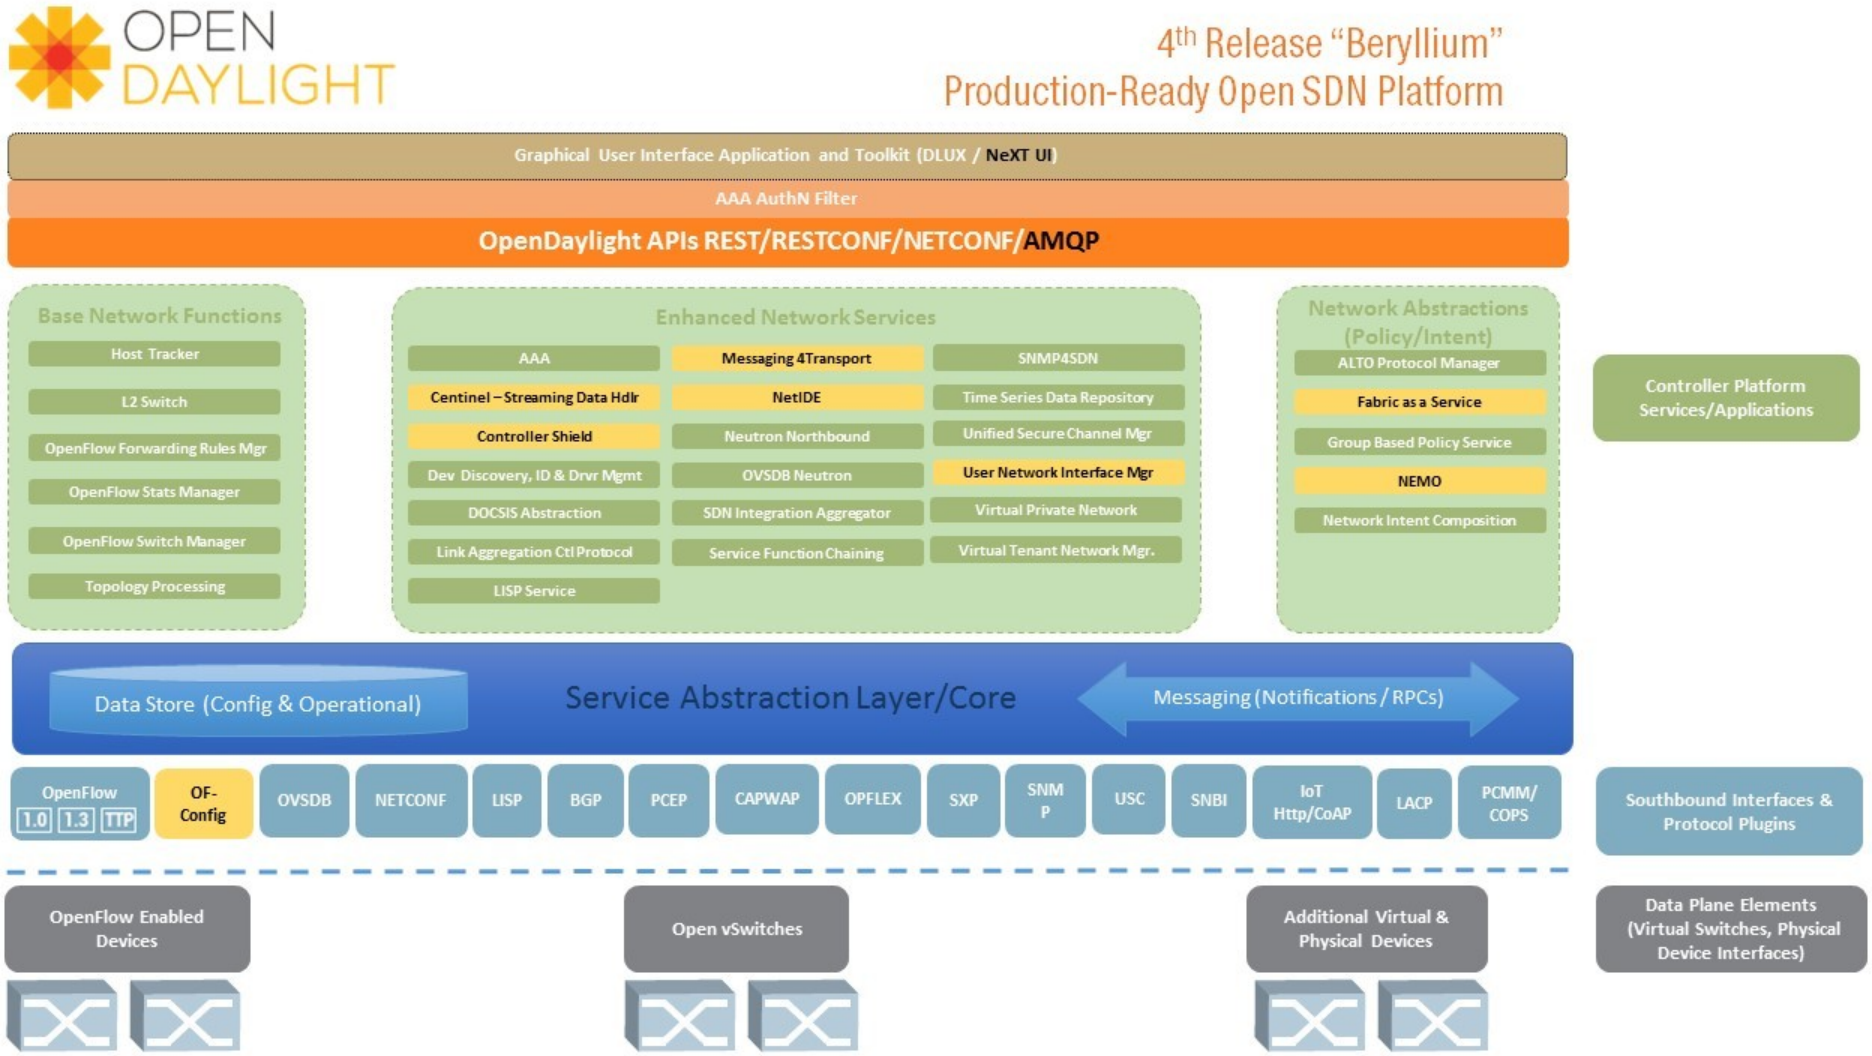

Next Step: Transition To Production Operations

- Production Operations Requires Sophisticated Orchestration, Topology, Real-Time Monitoring, Measurements, Analytics and Response
- Currently, Many Control Frameworks Are Being Investigated To Determine Its Potential For Achieving These Operational/Production Objectives
- One That Is Being Developed Is Ciena's Blue Planet
- Blue Planet Has A Large Number Of Components
- The StarLight Consortium Has Established a Research Project With Ciena To Experiment With, Investigate, and Demonstrate Several Of These Components

Ciena Blue Planet : SDN Management and Control

Web-Scale Technologies

SDN Applications

Multi-Vendor SDN Control

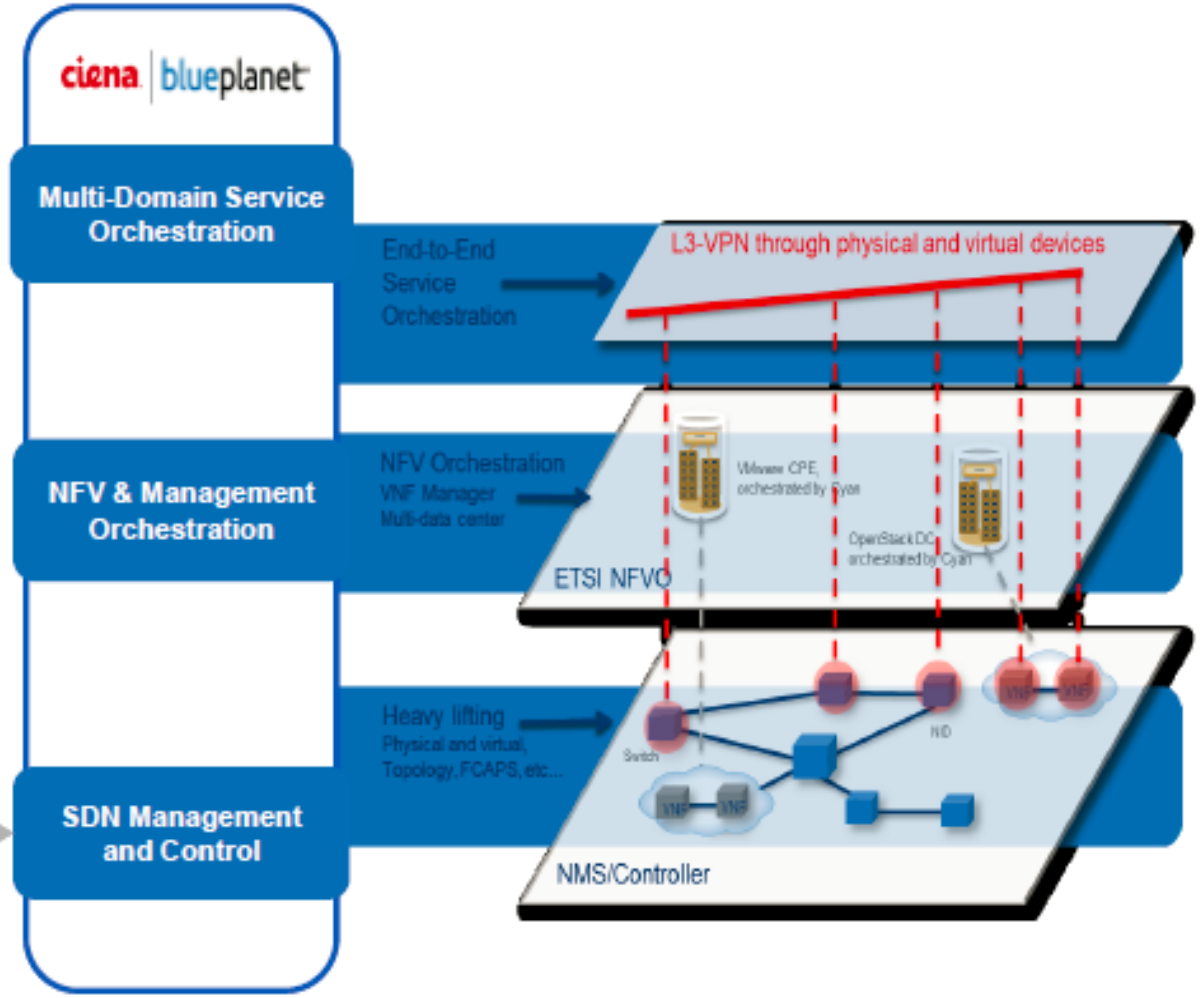

ciena | blueplanet

Global LambdaGrid Workshop Demonstration

- At the Global LambdaGrid Workshop, Capabilities For Using Blue Planet For Real Time Analytics Is Being Demonstrated By the StarLight Consortium and Ciena.
- This Is a Prelude/Pre-Staging Event To a Major Demonstration at SC16 In November In Salt Lake City Utah.
- This Is A Demonstration Of A Real Time “Blue Planet Analytic Probe”
- The Demonstration Is Being Supported By a Large Scale International 100 Gbps Testbed

Real Time Analytics Experiment/Demonstration

1) Goal: Develop A Mechanism For End-to-End Performance Monitoring Of A Specific Service

2) Basic Concept: Read Diagnostics And Performance Data From All Gear That Is Interconnected In The Network – And Analyze That Data In Real Time (!)

Approach:
Integrate The Data-Collecting Code Within the Blue Planet Analytics Platform Through A CLI Based Adapter That Logs Into The Network Equipment and Extracts Performance Monitoring Data.

Proof-of-Performance
Perform Real-Time Data Gathering Using Resource Adapter's (RA's) , To Enhance And Expand Previous Analytics Demonstrations With Real-Time Live CENI Network Operational Data.

Plan For Showcasing Emerging Capabilities For Real-Time Analytics

1) Select An Application
2) Implement Blue Planet On, Testbed Network Designed To Test and Experiment Proof Of Concepts and Stretch Objectives, As a Real, Live (Not Simulated) Network, In Collaboration With R&E Networks Operators And Analytics Designers (CENI)
3) Incorporate Blue Planet into CENI To Expand Statistics, Topology and Orchestration Techniques With New Specialized Resource Adapters
4) Augment Blue Planet Analytics Application with Performance Probes Created To Extract Much Link Data, Sampling On

   - perfSONAR Data
   - An 8700 Implemented On The CENI Testbed
   - Various Other Switches and Routers
   - Bare Metal and Virtual Machine's NIC Information
   - Processes Using CLI Prompts And Incorporate Relevant Statistics

The Application SAGE2 Based Scientific Visualization

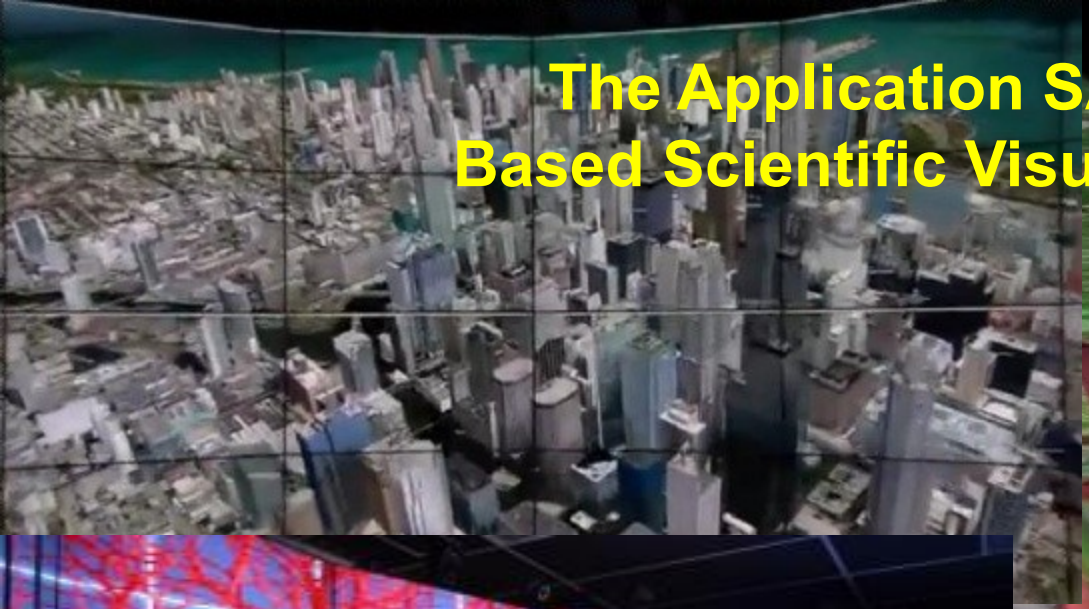

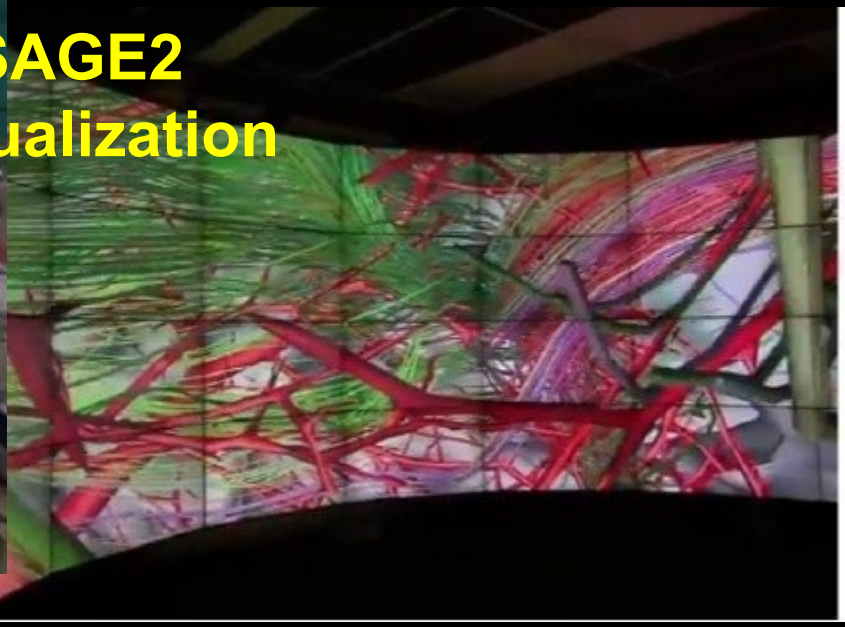

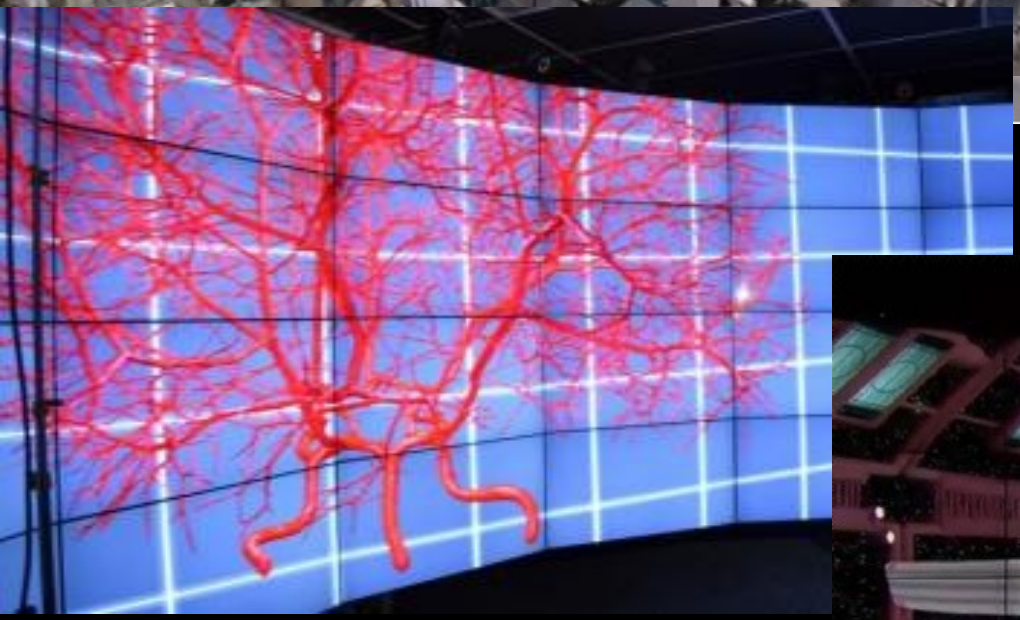

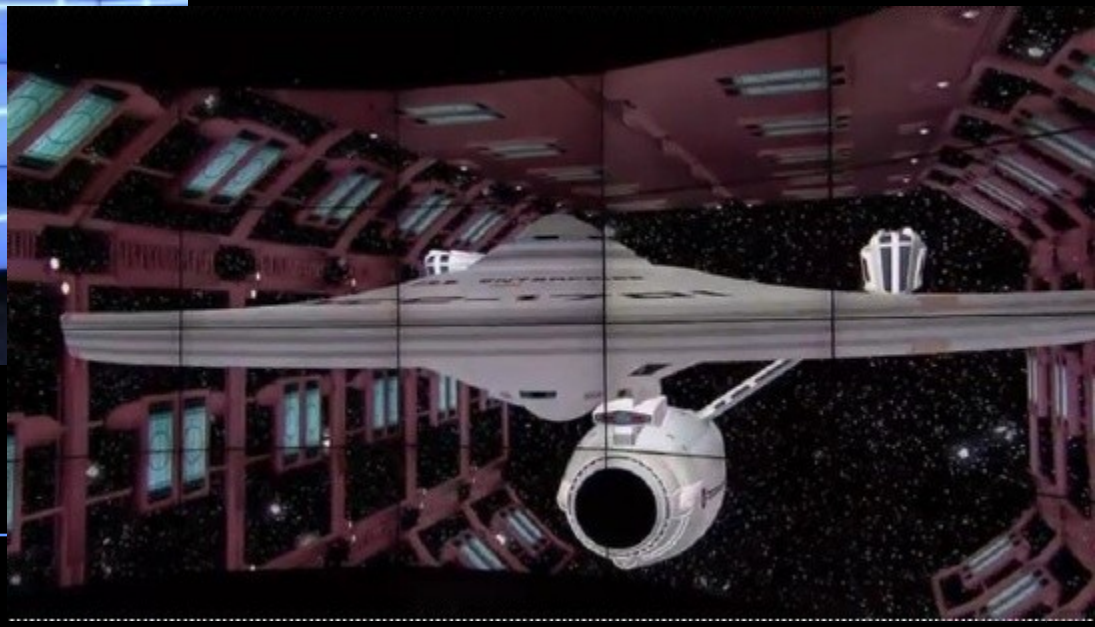

CAVE2: EVL UIC

What is SAGE2?

Scalable Amplified Group Environment

- Middleware: Access, Display, Share High-Resolution Digital Media On One or More Scalable Resolution Tiled-Display Walls
- Uses Web Technologies - Rewrite of SAGE: Scalable Adaptive Graphics Environment
- Multi-Touch Interaction (One or Many People)
- Can Push Laptop Screens Or Windows Onto Walls

Ciena Vector Summit 2015

Chicago, Canada, Amsterdam ExoGENI SAGE2/SDN Demo – May-June 2015

Using ExoGENI/SDN Technologies, a Single Slice Was Provisioned On the UvA ExoGENI Rack, (Slice = Integrated Resources – Here, A Virtual Machine (VM) In Amsterdam Interconnected Via Dedicated Private Network and Accessible From a Public Internet Address). The VM Supported the SAGE2 Server. Ciena VS Participants In Ottawa Could Open a Chrome Browser On Their Laptops, Connect To VM IP Address of the VM and Drag and Drop Content, Share Desktop To the tiled Display, Manipulate Content Already On the Display, All Through The Amsterdam SAGE2 server.

UIC/EVL, Northwestern/iCAIR, Ciena Research, University of Amsterdam (UvA)

Photo: Cees de Laat

The Testbed: CENI

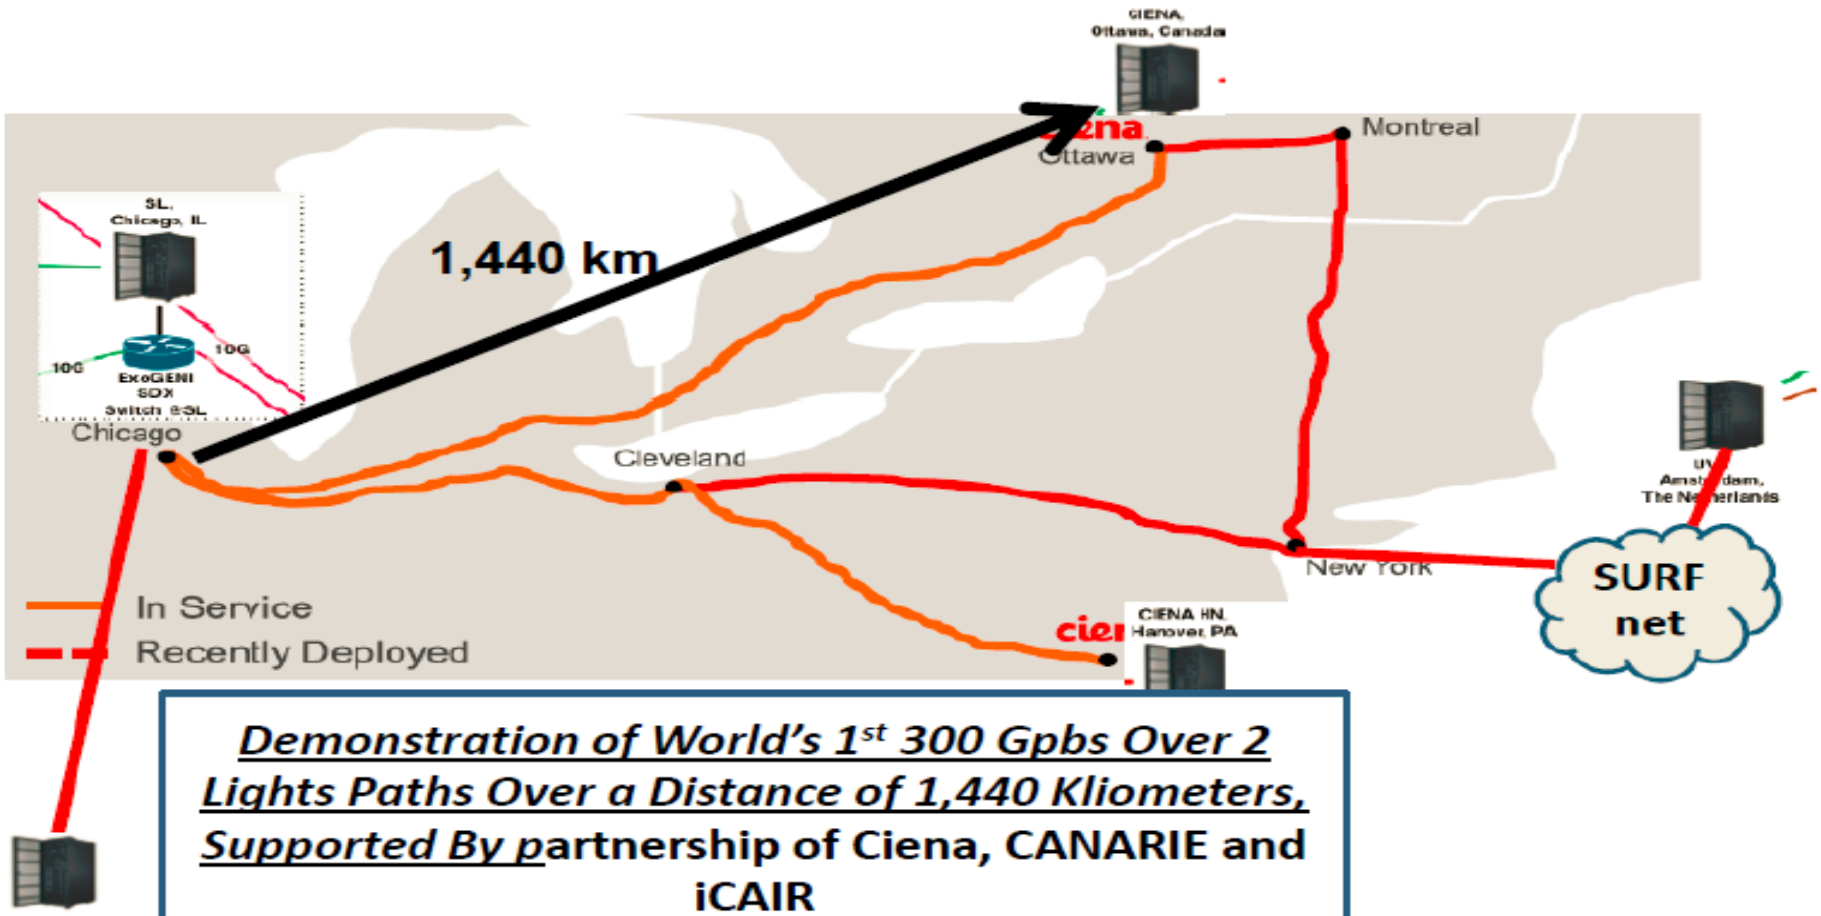

***Demonstration of World's 1st 300 Gpbs Over 2 Lights Paths Over a Distance of 1,440 Kliometers, Supported By p*artnership of Ciena, CANARIE and iCAIR**
May 16 2016

Network Diagram for Analytics Demonstration

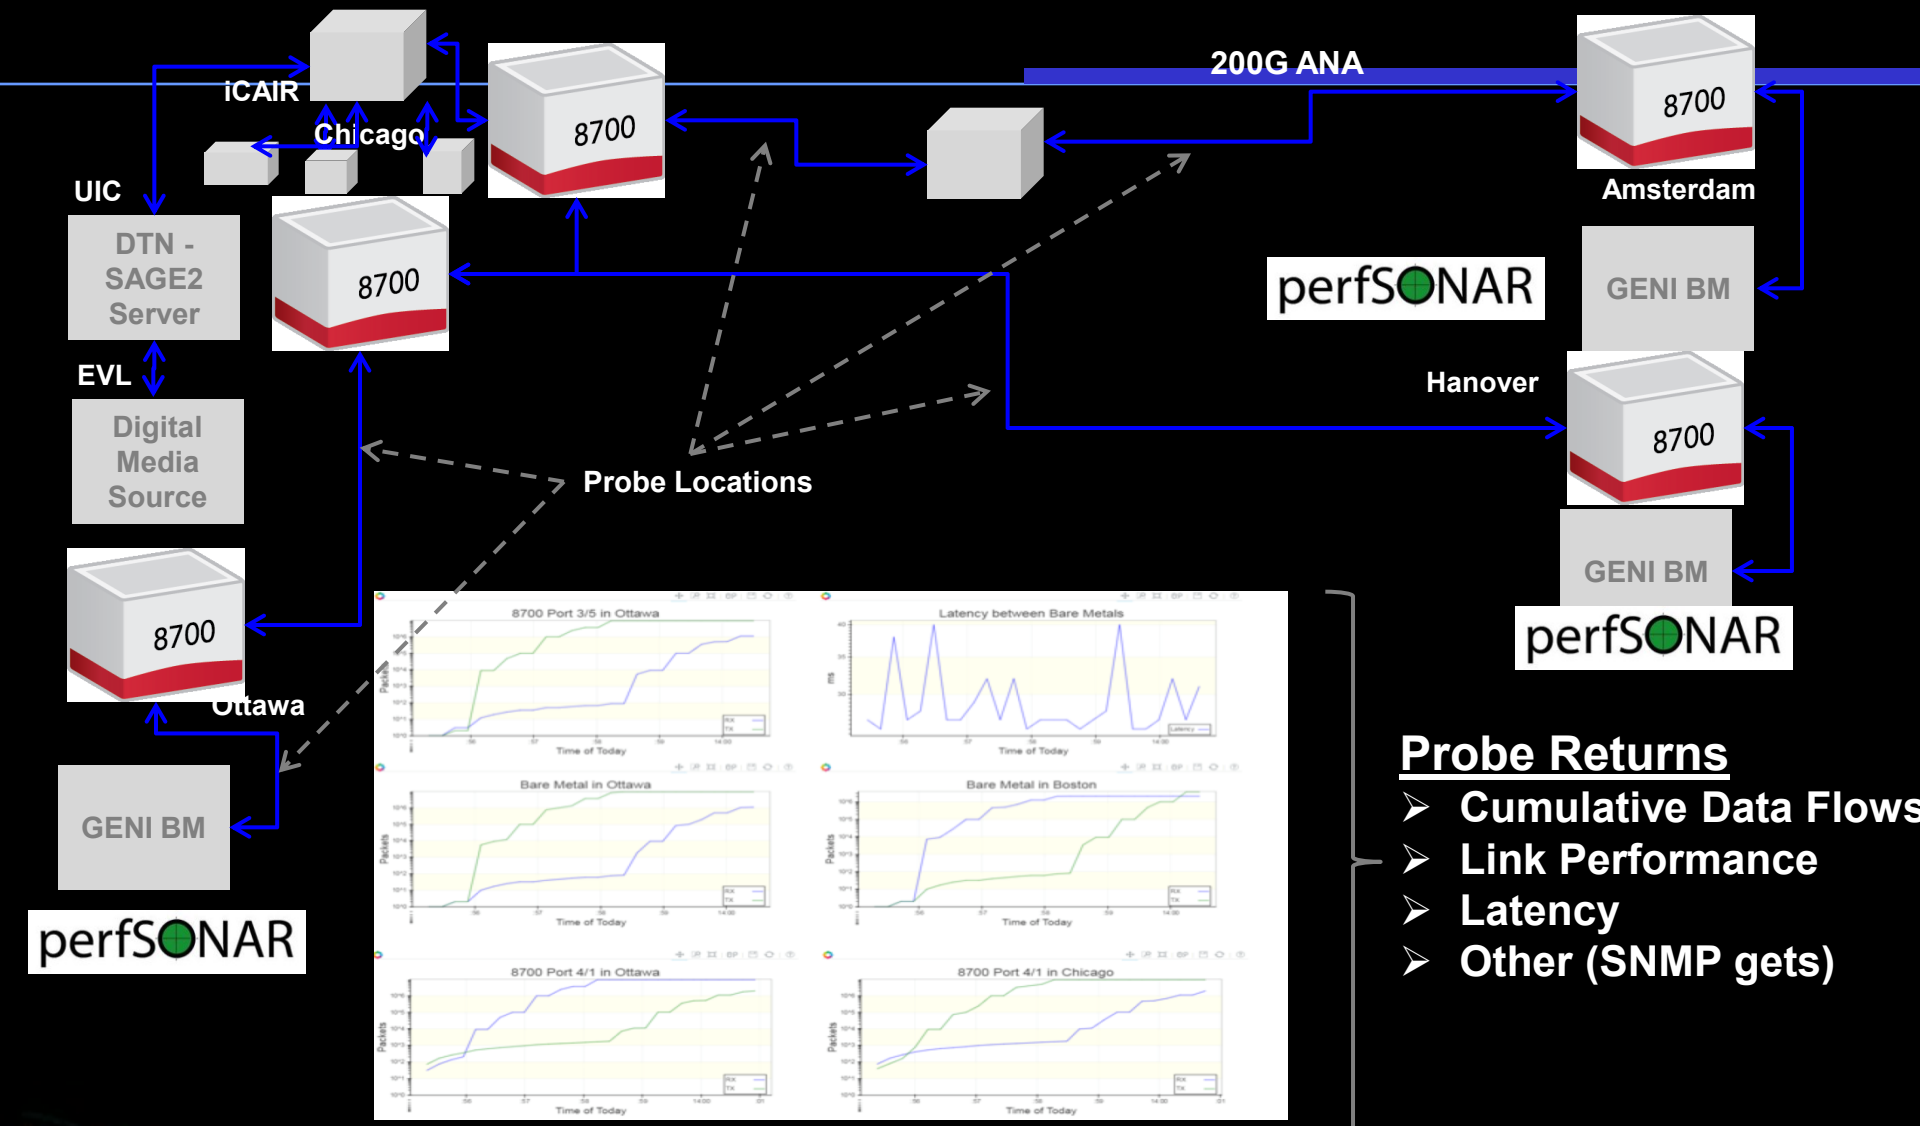

Probe Returns

- Cumulative Data Flows
- Link Performance
- Latency
- Other (SNMP gets)

E2E Real Time Service Analytics Over 100 G Paths Using the Blue Planet Framework On CENI

STARLIGHT SDX

The Processes

- Blue Planet Has Been Implemented On CENI Testbed Server In Ottawa and Chicago
- CENI Blue Planet Server Has Access To All The Devices Illustrated And Can Gather Data Using CLI-Based Resource Adaptors (RAs) Developed For This Project
- The Data Is Then Transformed And Reformatted for The BP Analytics Application
- The Analytics Application Creates Graphics Using This Data That Is Continuously Collected Live During The Demonstration.

RA Setup

Command Line (CLI)

Send Single Login-Logout Commands Through RA SDK.

- Uses JSON format for scheme, endpoint, and commands.

```
{
  "endpoint": "cli",
  "type": "bpprov.runners.simple.Sequence",
  "endpoint-parameters": {
      "command": “port show statistics"
  },

  "tests": [],
  "out-path": [

  ],
  "in-path": [
  ]
}
```

Need to know which commands are accepted by the device OS.

The Setup

- Clear All Counters, Bins and Statistics On All Devices On Network
- For the list of Port Statistics on 8700 and interface statistics on Bare Metals, Use:
  - Port Show Statistics and Ifconfig ethX Respectively
- Run Each Command In a For-Loop (Highly Inefficient)
  - Runs In Python
  - Create a Dictionary or JSON (key-value pairs) With the Output Data (Formatted Text)
  - Send This Organized Data To a Service (Kafka)

STARLIGHT SDX

Data Gathering Flowchart

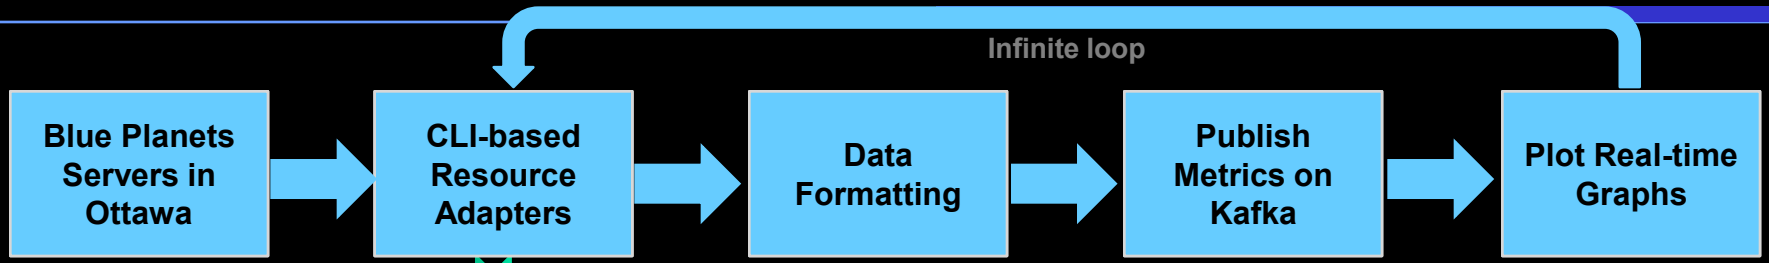


```
{"": ["Tx", "Rx"], "3/19": [50, 149], "3/18":
"HN_bond1": [73, 28], "3/20": [26, 13], "1/2"
 "3/6": [0, 0], "3/9": [0, 0], "2/1": [0, 0],
```


```
CNX08700#> port show statistics
------------------------ PORT STATISTICS SUMMARY (PACKETS) 20 SECOND SAMPLE ------------------------
|          Port          |          Pkt           |          Pkt           |
|                        |           Tx           |           Rx           |
|------------------------+------------------------+------------------------|
| 1/1                    |                      0 |                      0 |
| 1/2                    |                      0 |                      0 |
| 2/1                    |                      0 |                      0 |
| 2/2                    |                      4 |                      0 |
| 3/1                    |                    282 |                     82 |
| 3/2                    |                      4 |                      0 |
| 3/3                    |                      4 |                      0 |
| 3/4                    |                      4 |                      0 |
| 3/5                    |                      5 |                      0 |
| 3/6                    |                      0 |                      0 |
| 3/7                    |                      0 |                      0 |
| 3/8                    |                      0 |                      0 |
| 3/9                    |                      0 |                      0 |
| 3/10                   |                     76 |                     62 |
| 3/11                   |                      5 |                      0 |
| 3/12                   |                      5 |                      0 |
| 3/13                   |                     77 |                     61 |
| 3/14                   |                    247 |                     94 |
| 3/15                   |                     73 |                      0 |
| 3/16                   |                      0 |                      0 |
| 3/17                   |                      4 |                      0 |
| 3/18                   |                      0 |                      0 |
| 3/19                   |                    404 |                    530 |
| 3/20                   |                    158 |                     79 |
| 4/1                    |                    494 |                  1,158 |
| 4/2                    |                      5 |                      4 |
| HN_bond1               |                    440 |                    161 |
```

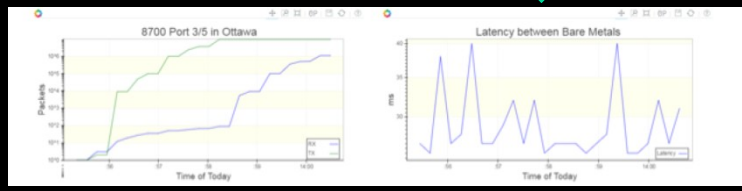

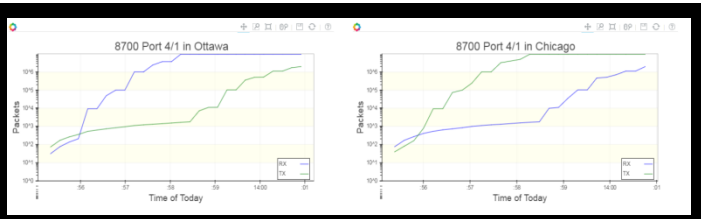

STARLIGHT SDX

Next Demonstrations: SC16 SDN/SDX/SDI 100 Gbps Demonstrations

- ***What's New=> Using Orchestrated SDN/SDX/SDI Services @100 Gbps Over WANs And 1 Tbps At SC16 Venue, In Salt Lake City, November 13-18, 2016***

ESnet
ENERGY SCIENCES NETWORK

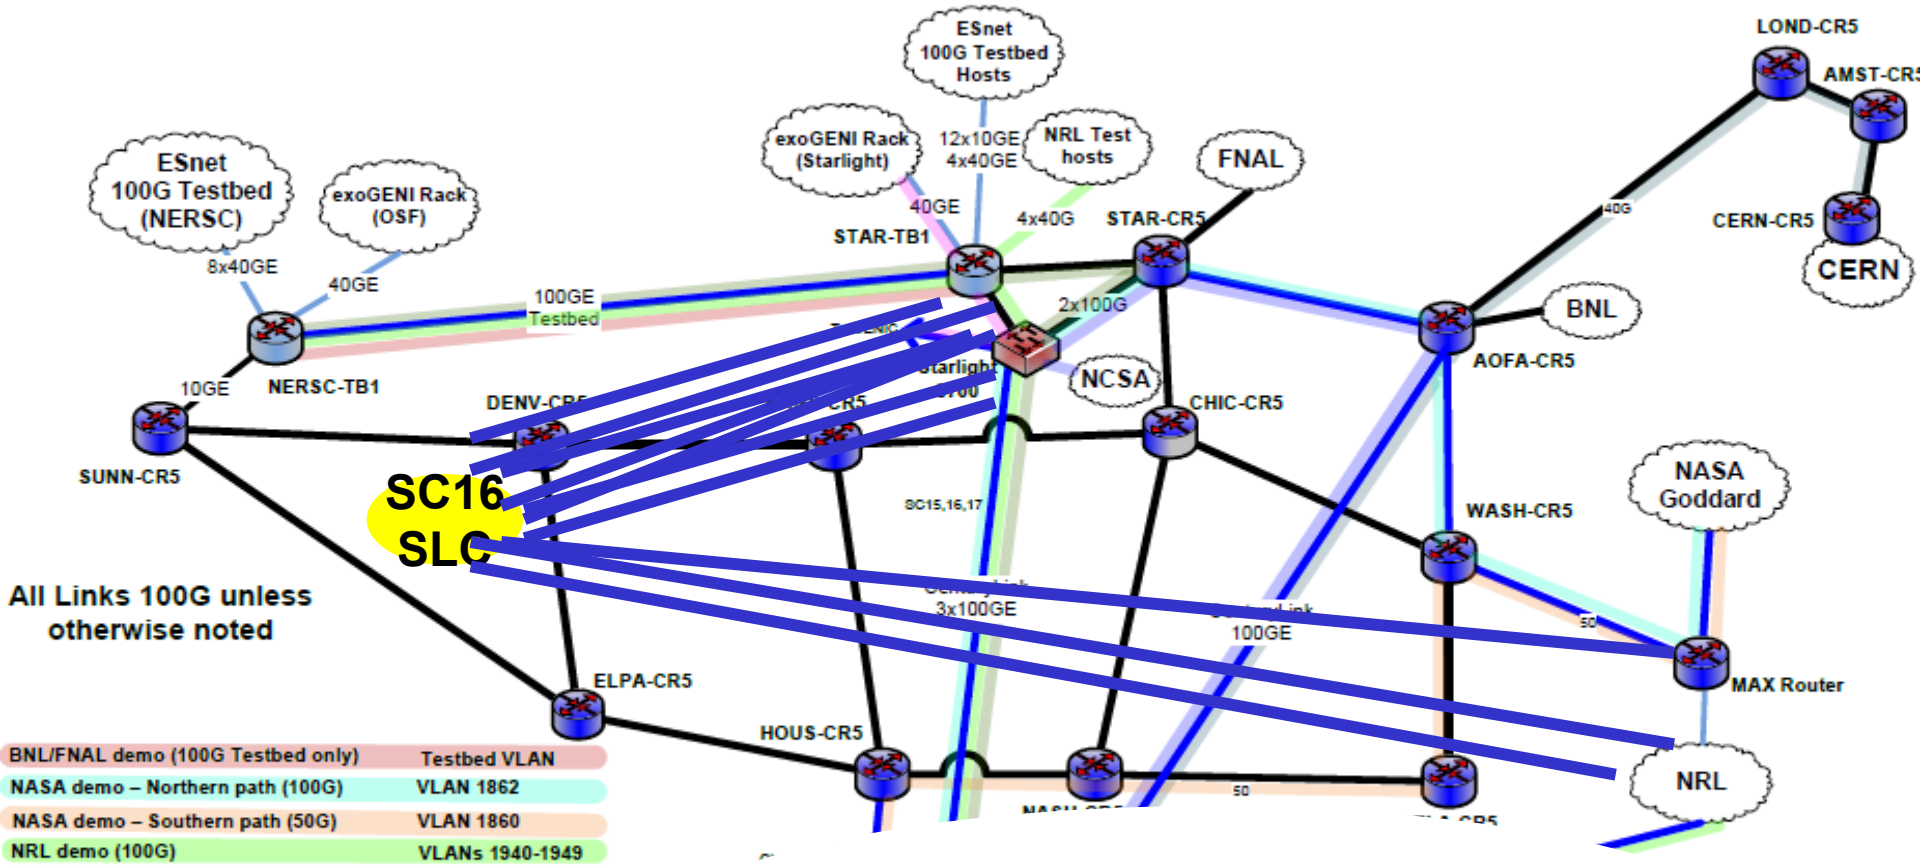

BNL/FNAL demo (100G Testbed only)	Testbed VLAN
NASA demo – Northern path (100G)	VLAN 1862
NASA demo – Southern path (50G)	VLAN 1860
NRL demo (100G)	VLANs 1940-1949
Caltech demo to NERSC (80G)	VLAN 2605
Caltech demo to CERN (40G)	VLAN 2606
StarLight/Ciena ExoGENI demo (40G)	VLAN 1779-1782
Aspera/NCSA demo (100G)	VLAN 2035

6*100 Gbps From StarLight
3* 100 Gbps from Wash DC

Other Demos 10G or Less:
- ESnet/RENCI demo: NERSC to ANL
- ESnet ENOS Demo: WASH, AMST, CERN
- ANL QoS Demo: DENV, ATLA

SC15 demos – ESnet
Brian Tierney, ESnet 11/6/2015
FILENAME SC15-DEMOS-V9.VSD

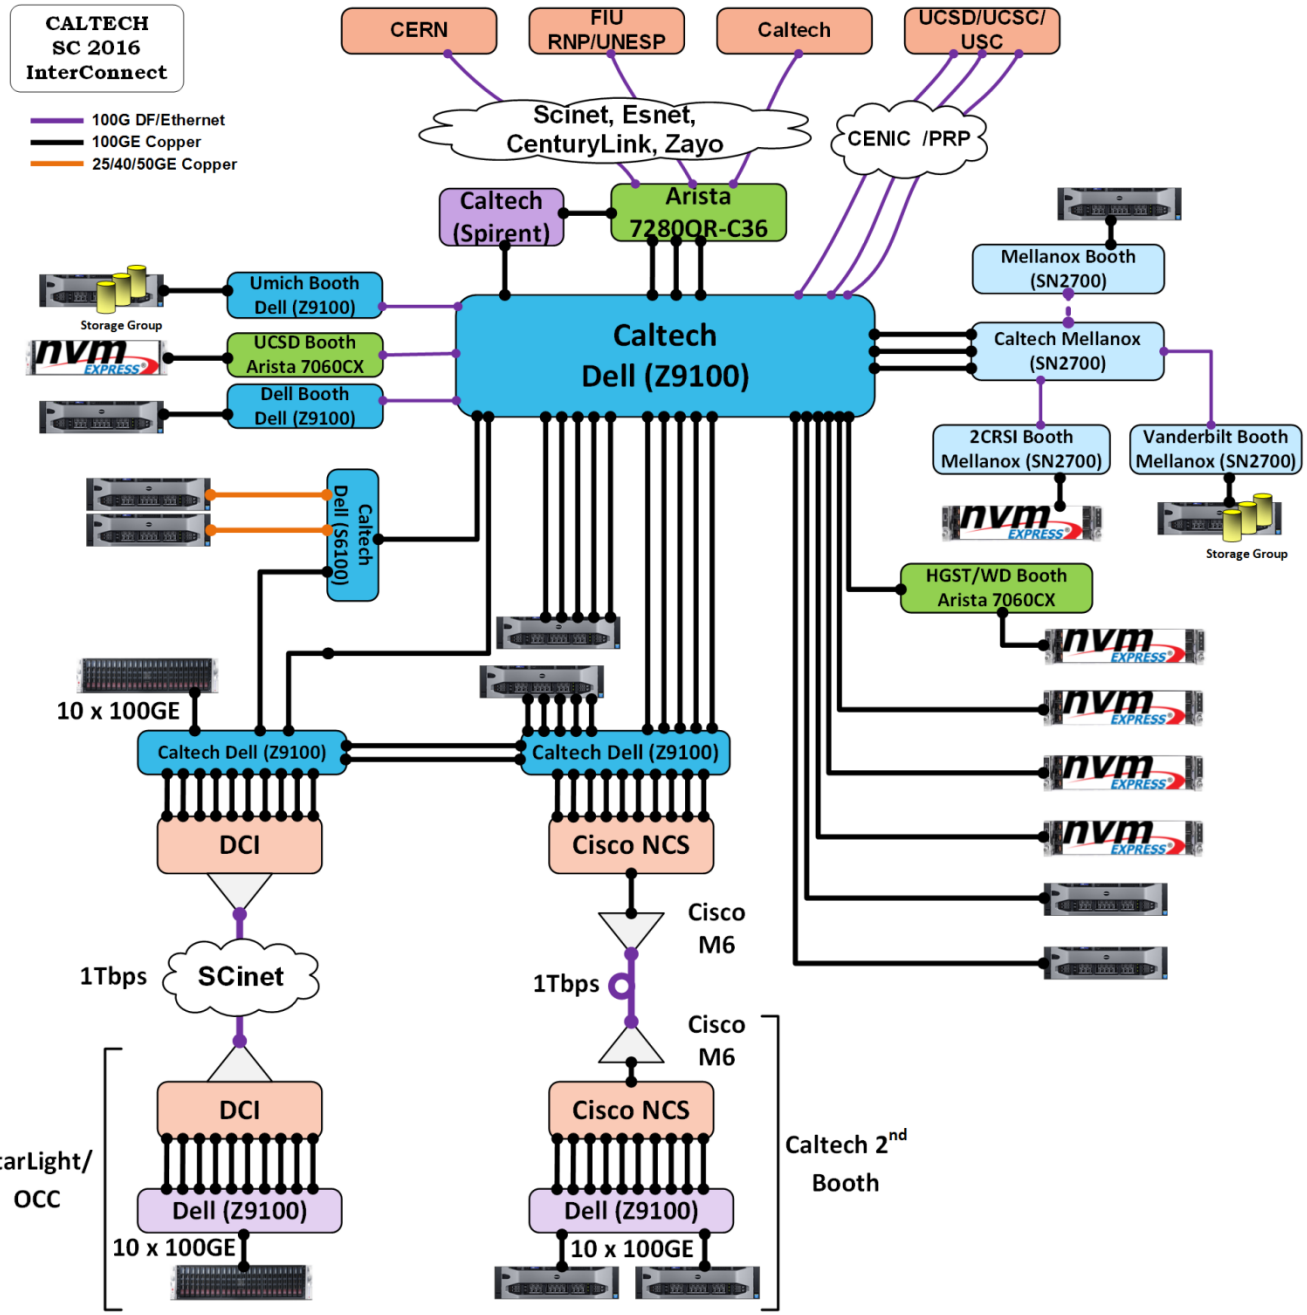
CALTECH
SC 2016
InterConnect
100G DF/Ethernet
100GE Copper
25/40/50GE Copper
CERN
FIU
RNP/UNESP
Caltech
UCSD/UCSC/
USC
Scinet, Esnet,
CenturyLink, Zayo
CENIC /PRP
Caltech
(Spirent)
Arista
7280QR-C36
Mellanox Booth
(SN2700)
Storage Group
Umich Booth
Dell (Z9100)
UCSD Booth
Arista 7060CX
Dell Booth
Dell (Z9100)
Caltech
Dell (Z9100)
Caltech Mellanox
(SN2700)
2CRSI Booth
Mellanox (SN2700)
Vanderbilt Booth
Mellanox (SN2700)
Storage Group
Caltech
Dell (S6100)
HGST/WD Booth
Arista 7060CX
10 x 100GE
Caltech Dell (Z9100)
Caltech Dell (Z9100)
DCI
Cisco NCS
Cisco
M6
1Tbps
SCinet
1Tbps
Cisco
M6
DCI
Cisco NCS
StarLight/
OCC
Dell (Z9100)
Dell (Z9100)
10 x 100GE
10 x 100GE
Caltech 2nd
Booth

Summary

- SDE Is Motivating A Network Revolution
- Hardware Defined => Software Defined
- Static=>Dynamic
- Reactive=>ProActive
- Delayed Analysis=> Real Time Analytics And Response
- Automated (vs Manual) Network Services and Processes
- Options For High Degrees Of Customization
- These Are important Trends For Our Global Community, Especially For Data Intensive Science

www.startap.net/starlight

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

iCAIR

STARLIGHT SDX