Eric Boyd
Rob Vietzke
December 1-2, 2011 – LHCONE Architecture

LHCONE – Meeting the needs of the LHC Stakeholders

Internet2's Seven Focus Areas

Advanced network and network services leadership
Services at scale: Services "above the network"
U.S. UCAN – Community Anchor Network Program
Industry partnership development and engagement
Global reach and engagement
Research community development and engagement
National/Regional Partnership

INTERNET2

Translating Focus Areas into Action

- Create a campus and networking environment conducive to data intensive science
  - Campus
  - Regional networks and backbone network
- Create a networking environment conducive to scientific collaboration on a global scale
  - Campus
  - Regional networks and backbone network
- Explore advanced networking technologies that will be needed to support future needs of data intensive science
- LHC = Harbinger of data intensive science virtual organizations of the future
  - Solution for LHC should be a template for other VOs

LHCONE Goals

- Ease of Connection
  - Provide a collection of access locations that are effectively entry points into a network that is private to the LHC T1/2/3 sites.
  - As soon as a T1/2/3 site is connected to LHCONE, it ought to be able to easily exchange data with any other T1/2/3 site over an infrastructure that is sized to accommodate that traffic.
  - LHCONE must accommodate both IP connections and several variations of circuit-based connections.
  - T1/2/3 sites may connect directly or via their network provider (e.g. National Research and Education Network (NREN), US Regional Optical Network (RON), ESnet, etc.).
- Exploitation of Infrastructure
  - Build on the familiar idea of exchange points
  - Provide a mechanism to better utilize available transoceanic capacity.
- Ease of Operation
  - Provide an infrastructure with appropriate operations and monitoring systems to provide the high reliability (in the sense of low error rates) that is essential for the high bandwidth, high volume data transfers of the LHC community
  - Provide a test and monitor infrastructure that can assist in ensuring that the paths from the T1/2/3 sites to LHCONE are also debugged and maintained in the low error rate state needed for LHC traffic

LHCONE High-level Architecture

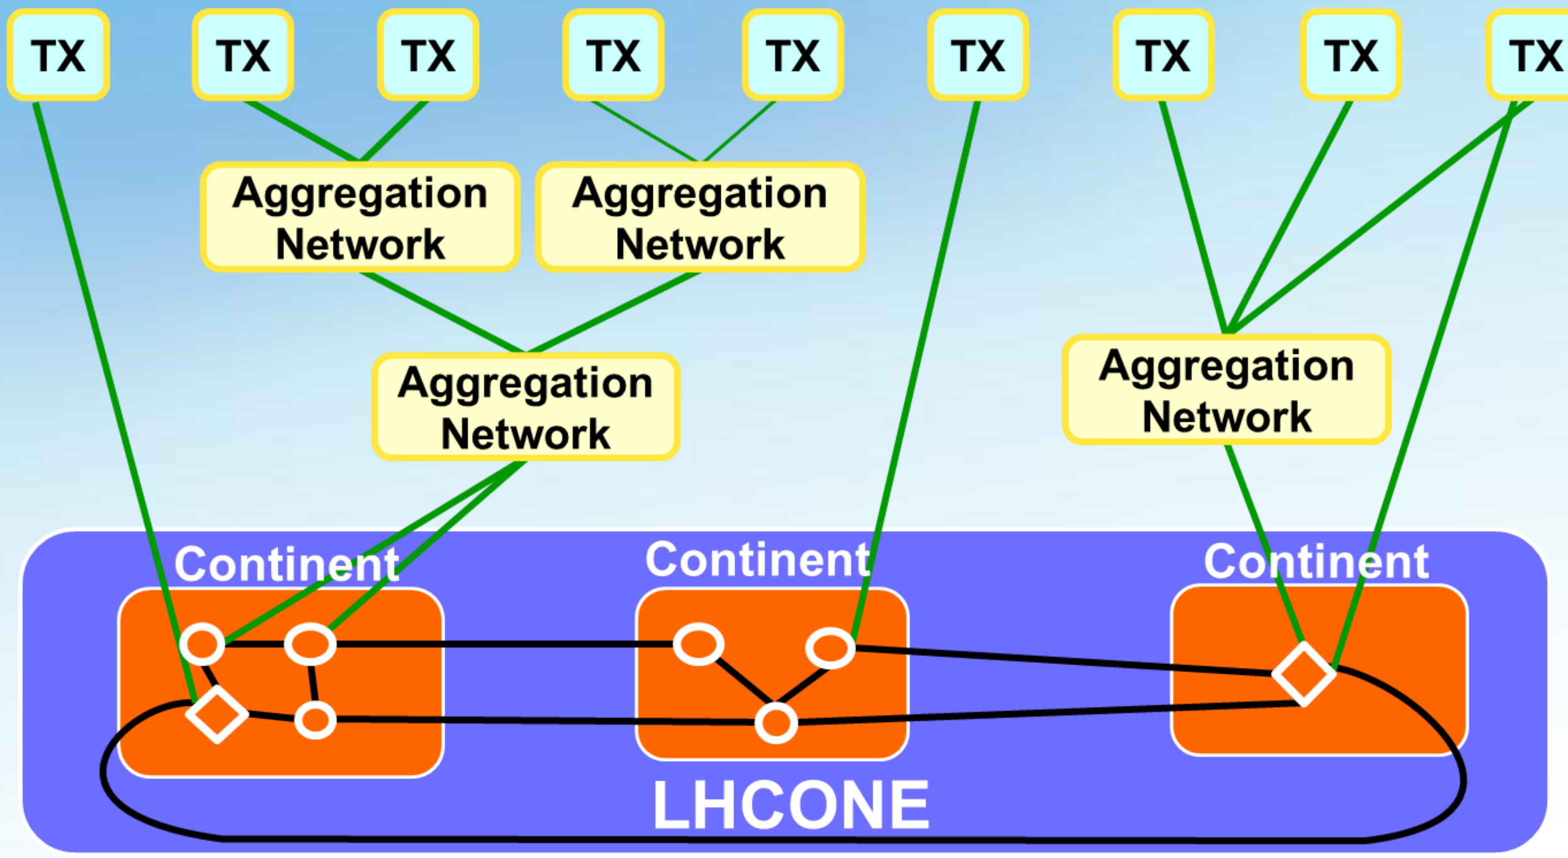

Single node Exchange Point Distributed Exchange Point

LHCONE Design Considerations

- LHCONE complements the LHCOPN by addressing a different set of data flows.
- LHCONE enables high-volume data transport between T1s, T2s, and T3s.
- LHCONE separates LHC-related large flows from the general purpose routed infrastructures of R&E networks.
- LHCONE incorporates all viable national, regional and intercontinental ways of interconnecting Tier1s, Tier2s, and Tier 3s.
- LHCONE uses an open and resilient architecture that works on a global scale.
- LHCONE provides a secure environment for T1-T2, T2-T2, and T2-T3 data transport.
- LHCONE provides connectivity directly to T1s, T2s, and T3s, and to various aggregation networks, such as the European NRENs, GÉANT, and North American RONs, Internet2, ESnet, CANARIE, etc., that may provide the direct connections to the T1s, T2s, and T3s.
- LHCONE is designed for agility and expandability.
- LHCONE allows for coordinating and optimizing transoceanic data flows, ensuring the optimal use of transoceanic links using multiple providers by the LHC community.

LHCONE Services

- Multipoint Service
- Point-to-Point Service
  - With bandwidth guarantees
  - Without bandwidth guarantees
- Diagnostic Service

Stakeholders

- LHC Stakeholders (invested in success of LHC)
  - LHC Experiments in Aggregate
  - Individual LHC researchers
  - LHC Software Stack Developers
  - LHC Network Operators
    - e.g. USLHCnet, owners of dedicated LHC circuits
- Network Stakeholders (invested in support of data intensive science, with LHC as an exemplar)
  - Campus Infrastructure Providers
    - CIOs and Vice-Presidents for Research
  - Federal Lab Infrastructure Providers
  - Regional Networks
  - Exchange Points
  - Backbone Network Providers

What Internet2 is hearing

- There's no need to rush to a solution.
- LHCONE is a pilot, not production infrastructure
  - ATLAS has identified a small number of pilot sites
  - CMS?
- We should be working towards a solution that is 3 years out (just prior to the resumption of the experiment after shutdown)
- Intercontinental layer-2 multipoint pilot has encountered significant challenges

What Internet2 is doing

- DYNES – Building a nationwide distributed virtual instrument interconnecting campuses
- NDDI – Experimenting with advanced network services based on software-defined networking
- OS3E – Providing a persistent Layer 2 VLAN service over NDDI with dedicated connections; providing regional and distributed open exchanges
- ION – Providing a persistent Layer 2 VLAN service over MPLS (bundled with IP connections)
- Working with ESnet, GÉANT, and USLHCnet to create
  - Interoperable Layer 2 VLAN service
  - Interoperable Diagnostic service

INTERNET2

DYNES Projected Topology (November 2011)

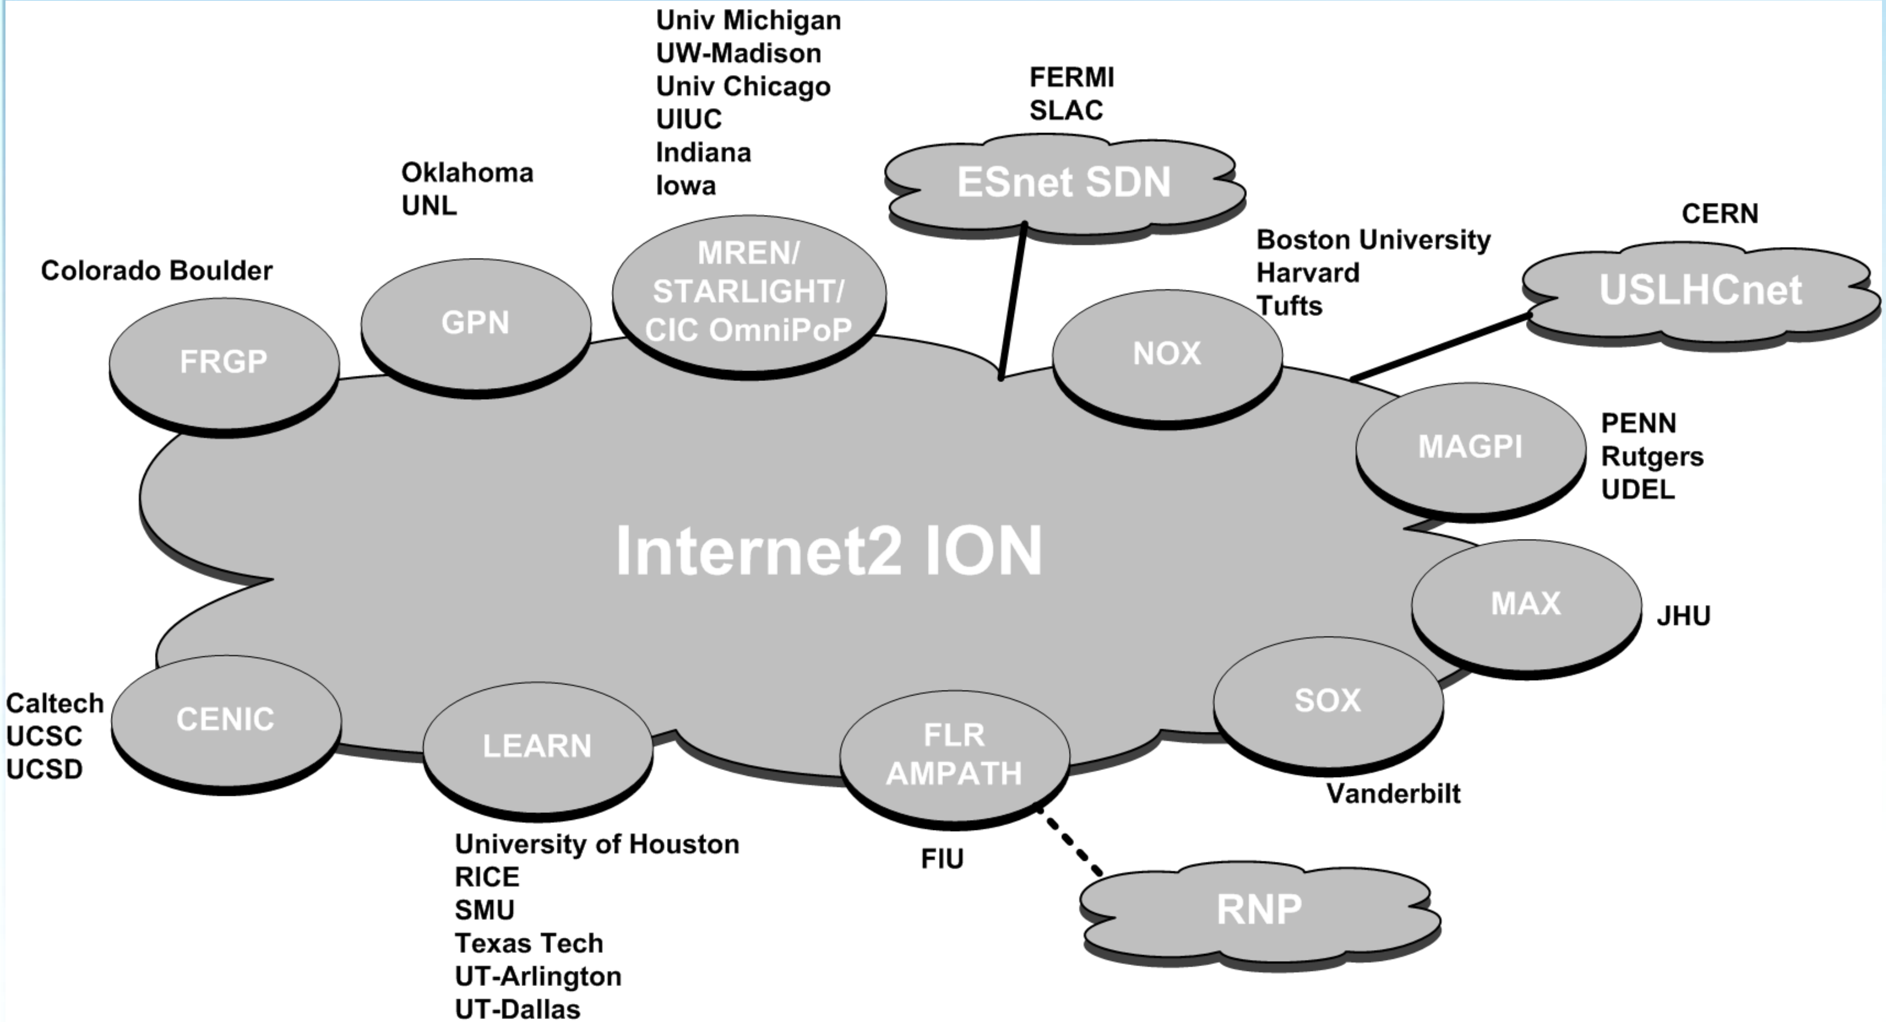

OS³E

The Open Science, Scholarship & Services Exchange

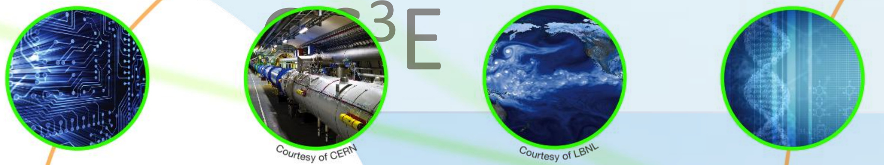

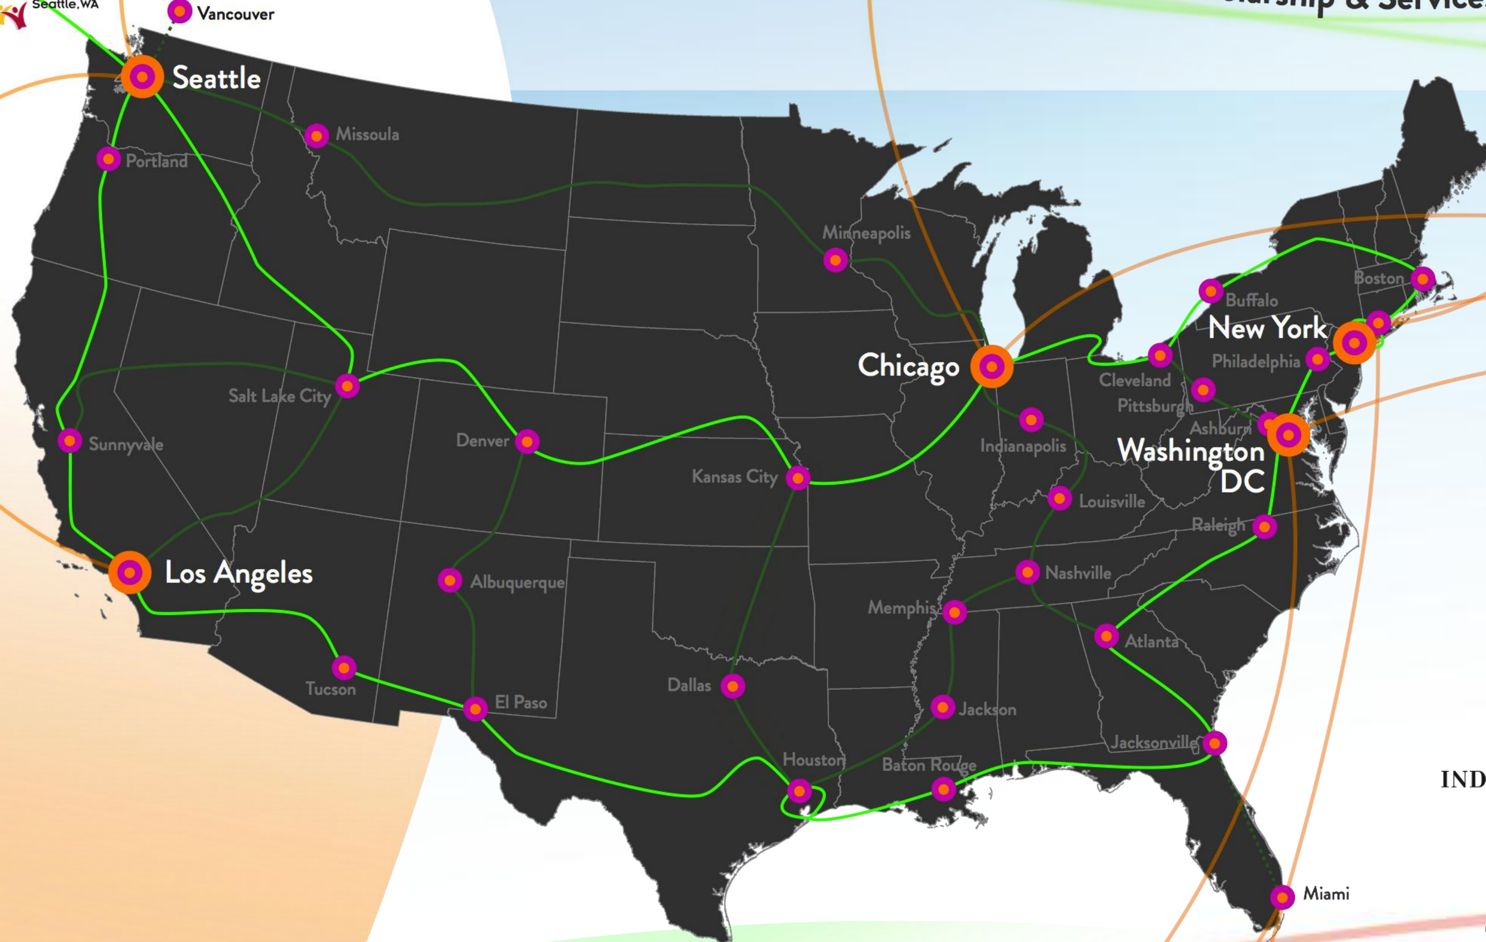

INTERNET2

INDIANA UNIVERSITY

STANFORD UNIVERSITY

INTERNET2

The Way Forward

- What do we envision the networking environment should look like?
- What questions do we need to answer to create that environment?
- What experiments (pilots) do we need to try to answer those questions

Networking Environment in 2014

- Compute, storage, and networking fully integrated into science software stacks
- 100G networks commonplace
  - 40G transoceanic network links commonplace
- Software-defined networking is mainstream
  - Data intensive science begins to adapt to take advantage of new networking technologies
  - Layer 2 networking at the core / Layer 3 networking at the edge
- In many parts of the world, researchers will have multiple options

Open Questions in 2011

- What are the appropriate APIs to delineate the boundary between compute, storage, and networking elements?
- How can science afford 100G capital costs?
  - Port costs on routers are high; port costs on switches are low
- What are the business models for sustainable research networks?

What experiments are needed?

- Operating a transatlantic Layer 2 multi-point VLAN service?
- Operating a multidomain, transoceanic point-to-point Layer 2 VLAN service?
- Operating a multidomain, transoceanic diagnostic service?
- Operating multi-domain software defined networks?
- Integrating network APIs into LHC software stack?

INTERNET2

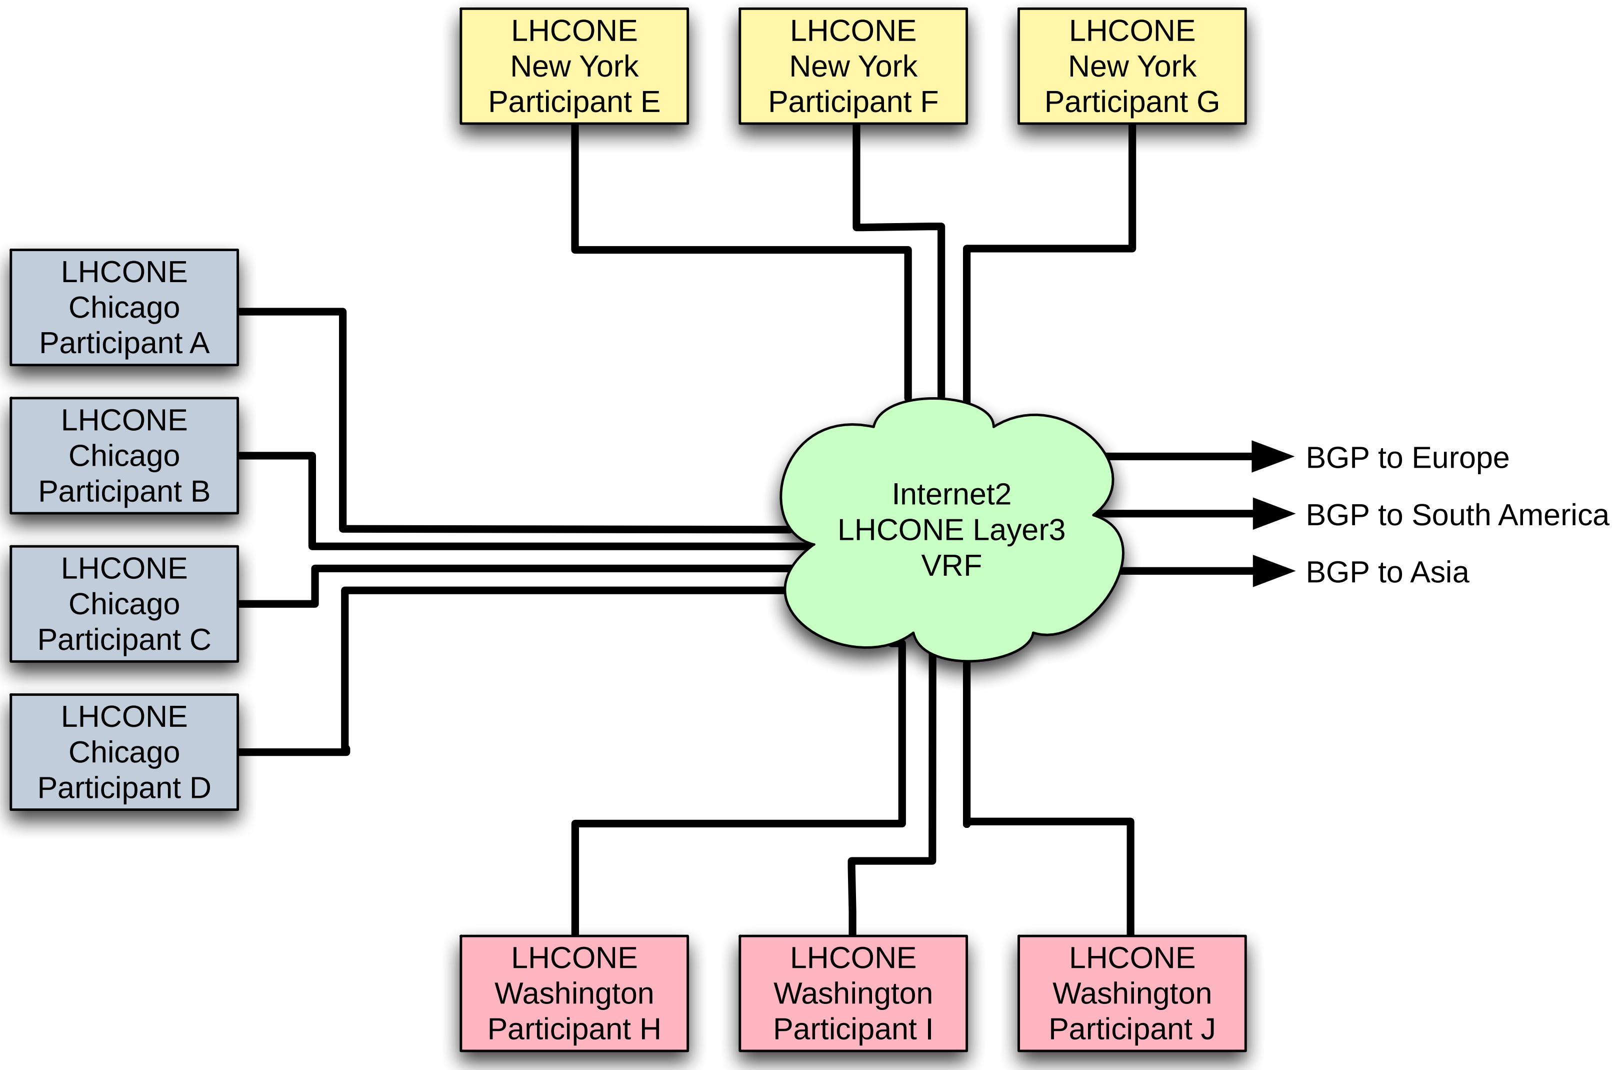
LHCONE
New York
Participant E
LHCONE
New York
Participant F
LHCONE
New York
Participant G
LHCONE
Chicago
Participant A
LHCONE
Chicago
Participant B
LHCONE
Chicago
Participant C
LHCONE
Chicago
Participant D
Internet2
LHCONE Layer3
VRF
BGP to Europe
BGP to South America
BGP to Asia
LHCONE
Washington
Participant H
LHCONE
Washington
Participant I
LHCONE
Washington
Participant J

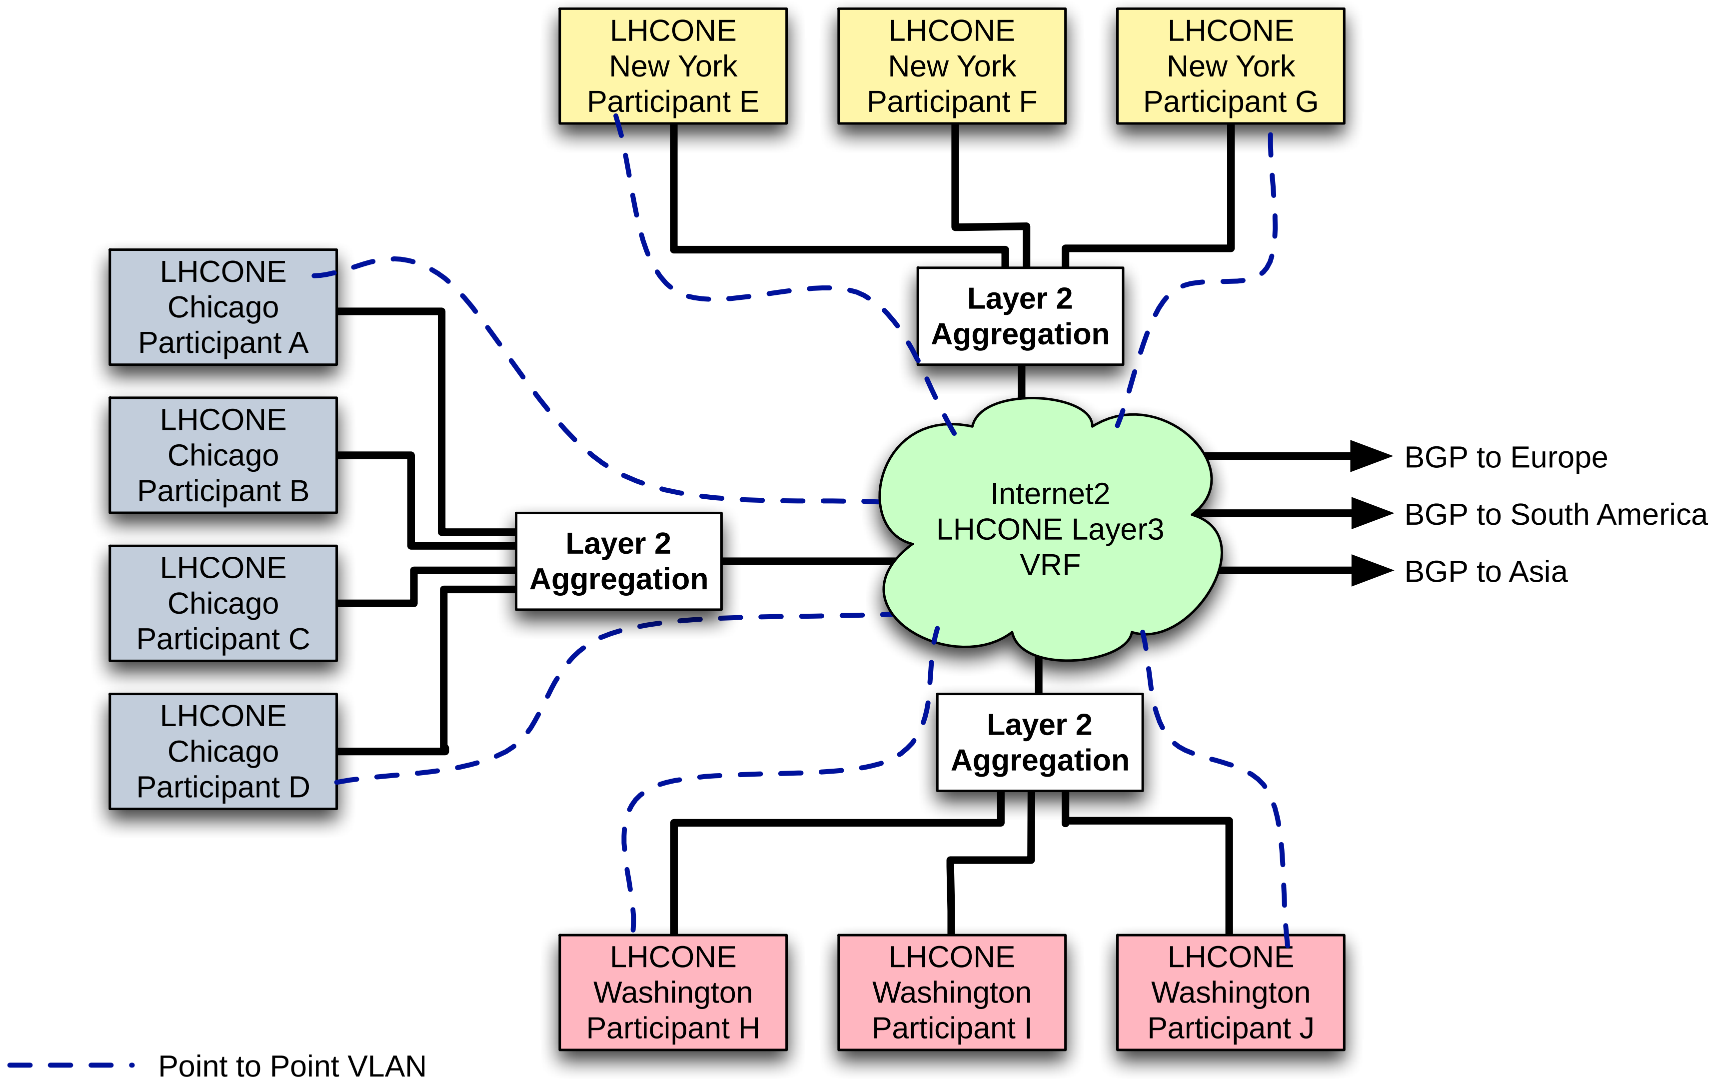
LHCONE
New York
Participant E
LHCONE
New York
Participant F
LHCONE
New York
Participant G
LHCONE
Chicago
Participant A
LHCONE
Chicago
Participant B
LHCONE
Chicago
Participant C
LHCONE
Chicago
Participant D
Layer 2
Aggregation
Layer 2
Aggregation
Internet2
LHCONE Layer3
VRF
BGP to Europe
BGP to South America
BGP to Asia
Layer 2
Aggregation
LHCONE
Washington
Participant H
LHCONE
Washington
Participant I
LHCONE
Washington
Participant J
Point to Point VLAN

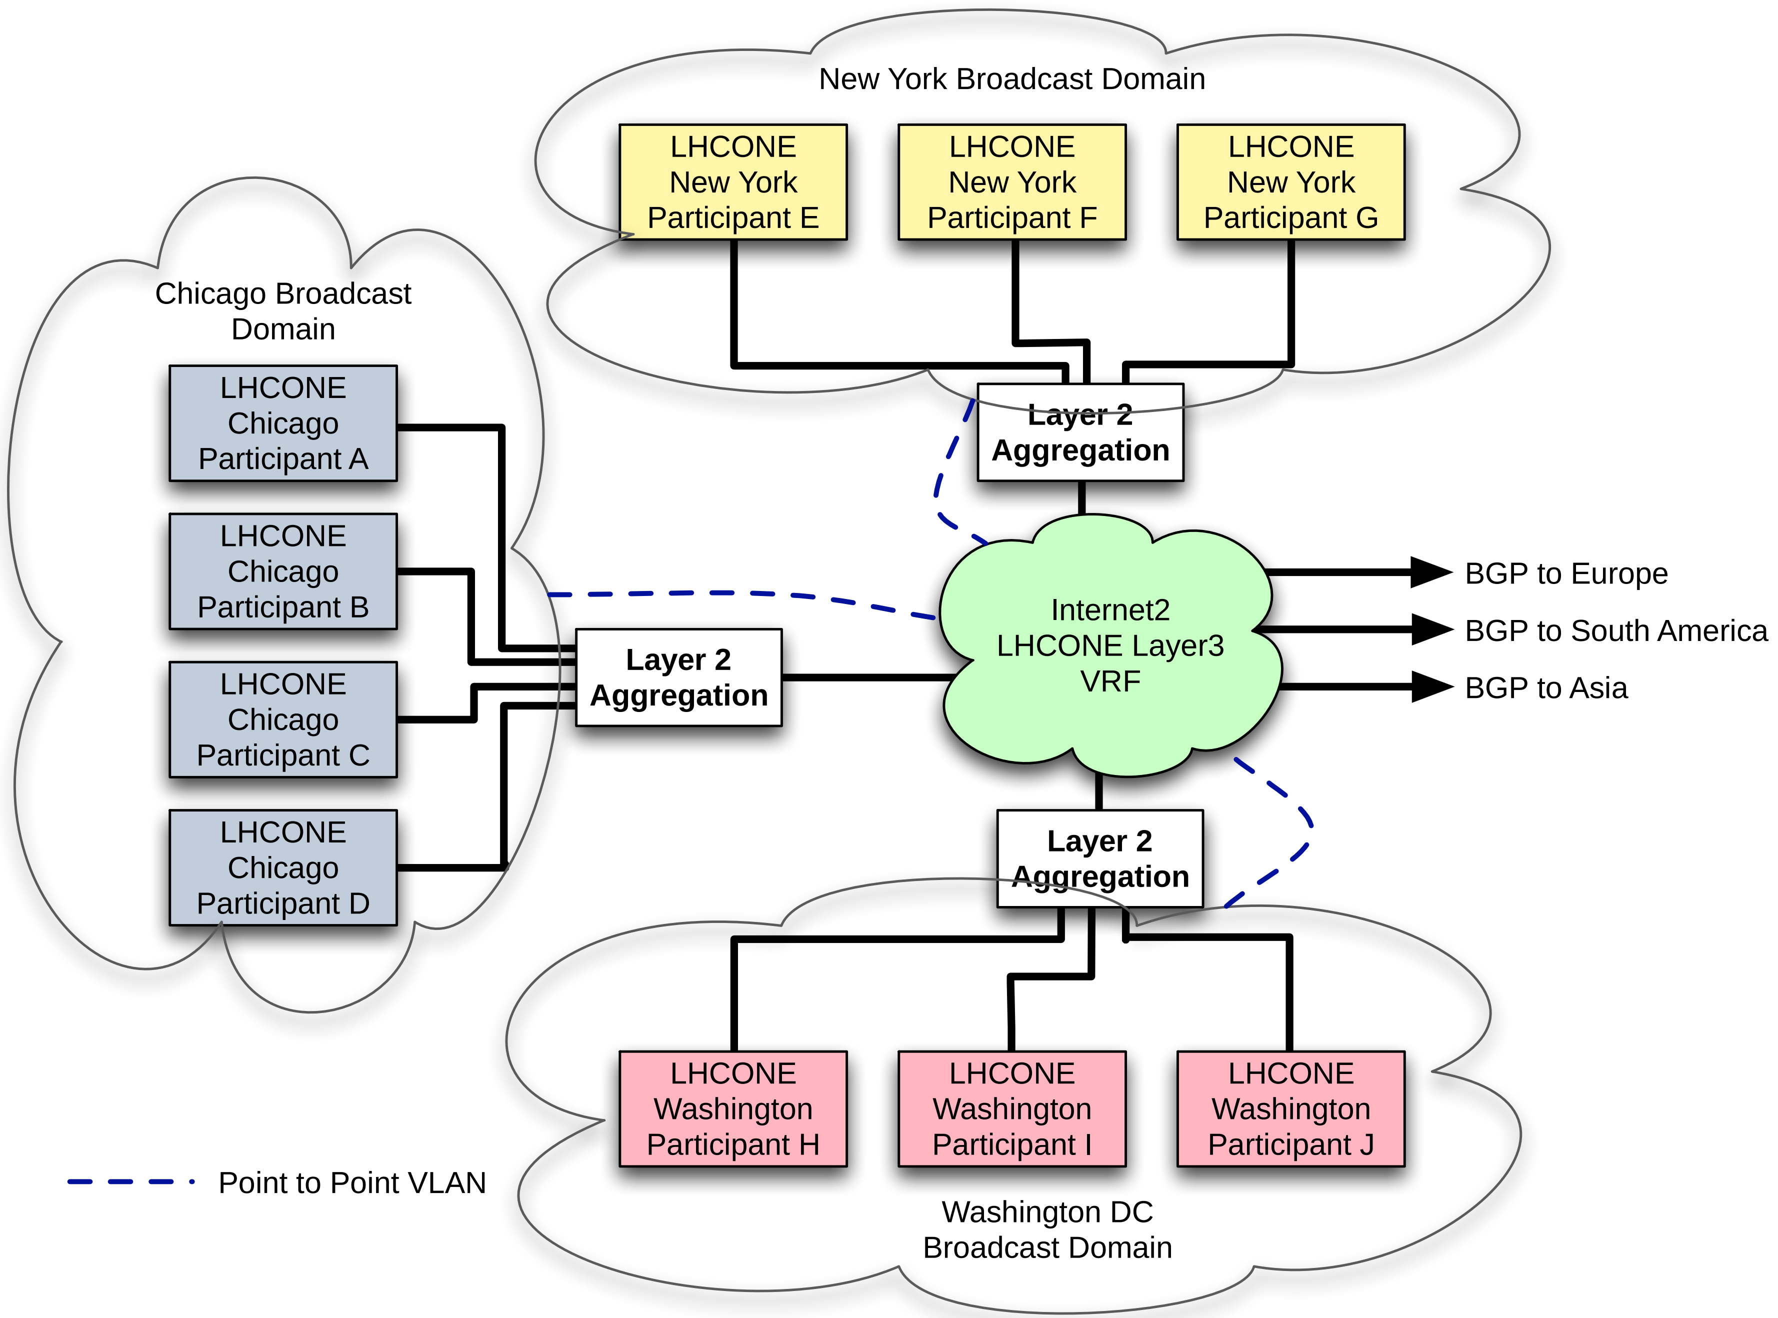
New York Broadcast Domain
LHCONE New York Participant E
LHCONE New York Participant F
LHCONE New York Participant G
Layer 2 Aggregation
Chicago Broadcast Domain
LHCONE Chicago Participant A
LHCONE Chicago Participant B
LHCONE Chicago Participant C
LHCONE Chicago Participant D
Layer 2 Aggregation
Internet2 LHCONE Layer3 VRF
BGP to Europe
BGP to South America
BGP to Asia
Layer 2 Aggregation
LHCONE Washington Participant H
LHCONE Washington Participant I
LHCONE Washington Participant J
Washington DC Broadcast Domain
Point to Point VLAN

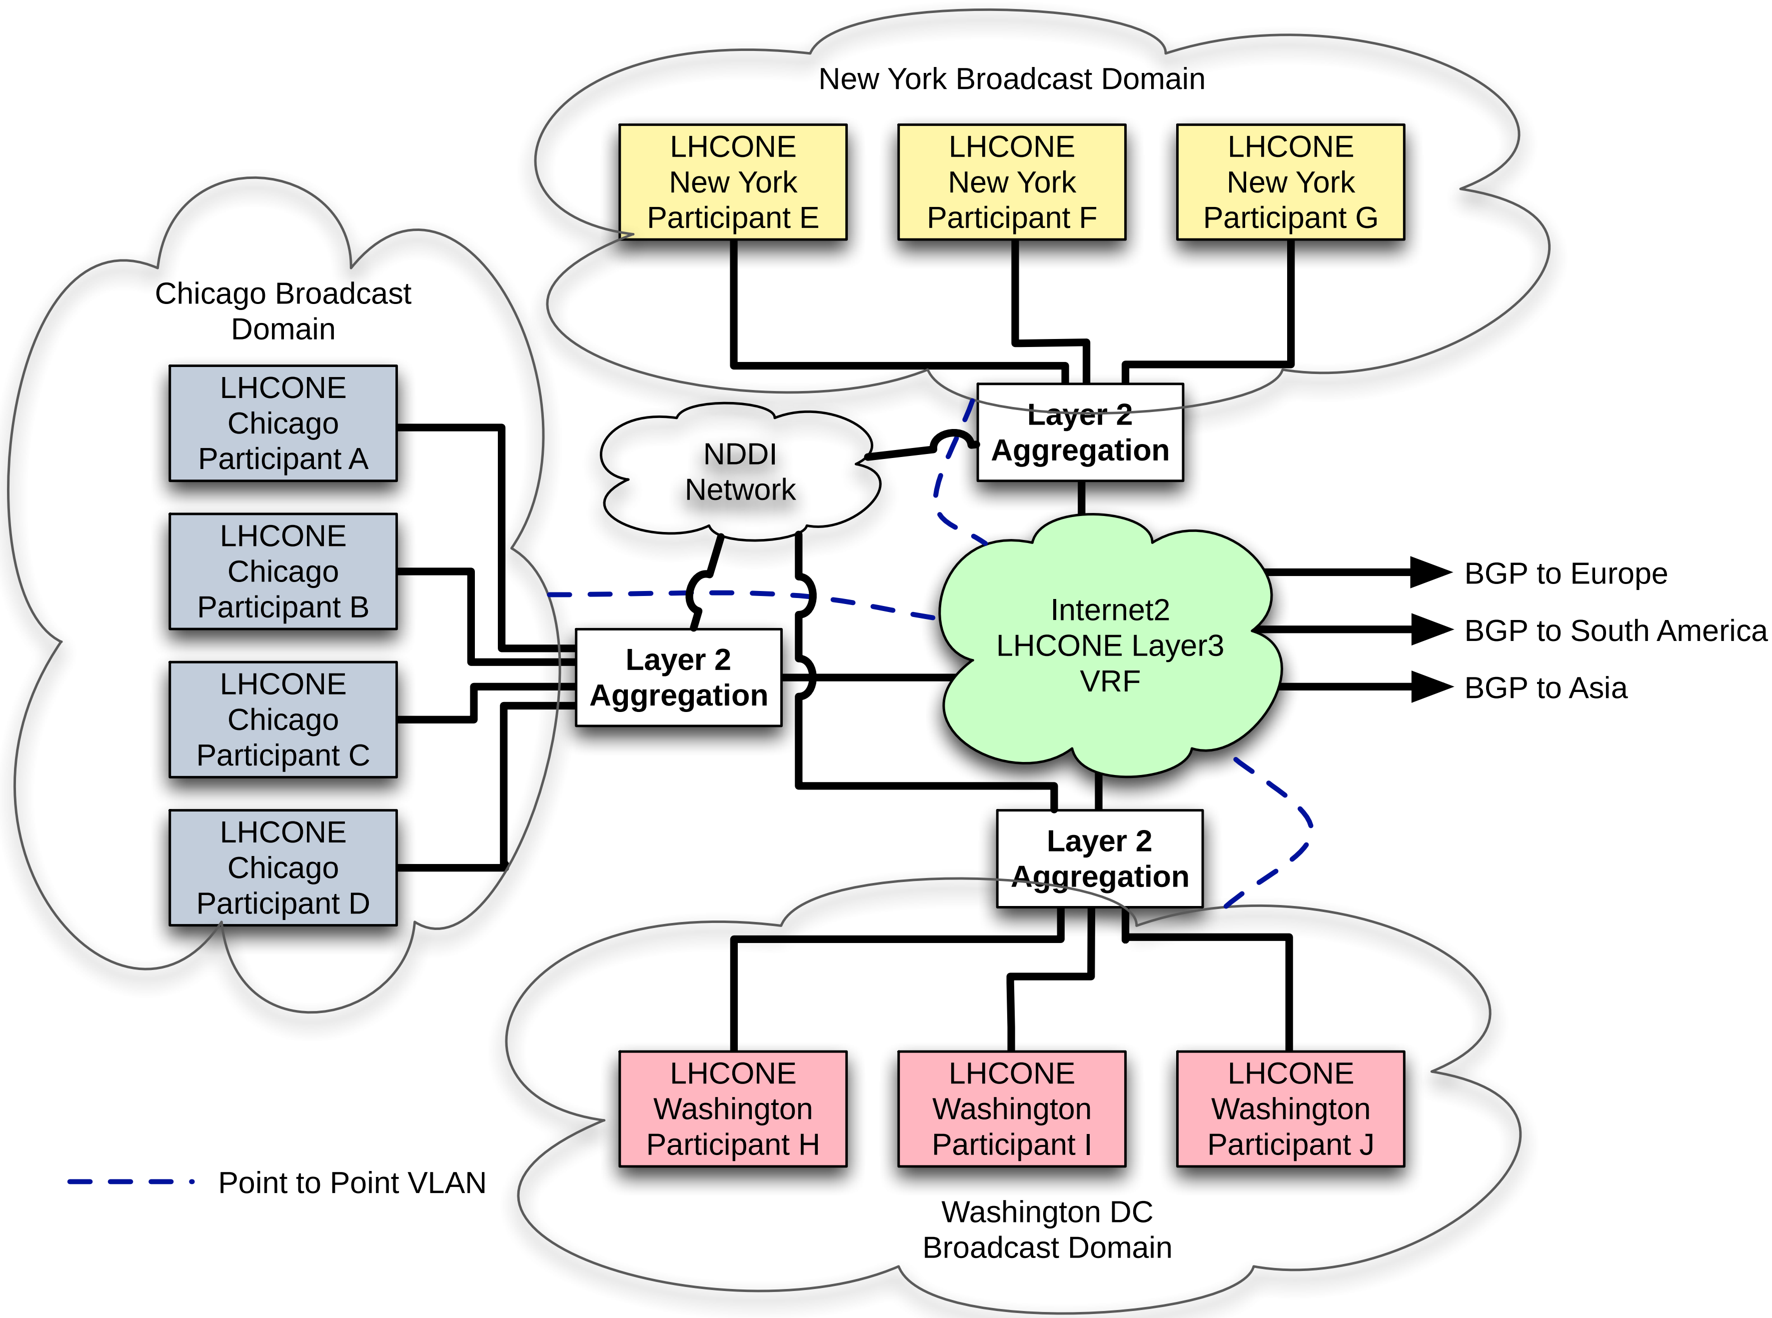
New York Broadcast Domain
LHCONE New York Participant E
LHCONE New York Participant F
LHCONE New York Participant G
Chicago Broadcast Domain
LHCONE Chicago Participant A
LHCONE Chicago Participant B
LHCONE Chicago Participant C
LHCONE Chicago Participant D
NDDI Network
Layer 2 Aggregation
Layer 2 Aggregation
Internet2 LHCONE Layer3 VRF
BGP to Europe
BGP to South America
BGP to Asia
Layer 2 Aggregation
LHCONE Washington Participant H
LHCONE Washington Participant I
LHCONE Washington Participant J
Washington DC Broadcast Domain
Point to Point VLAN

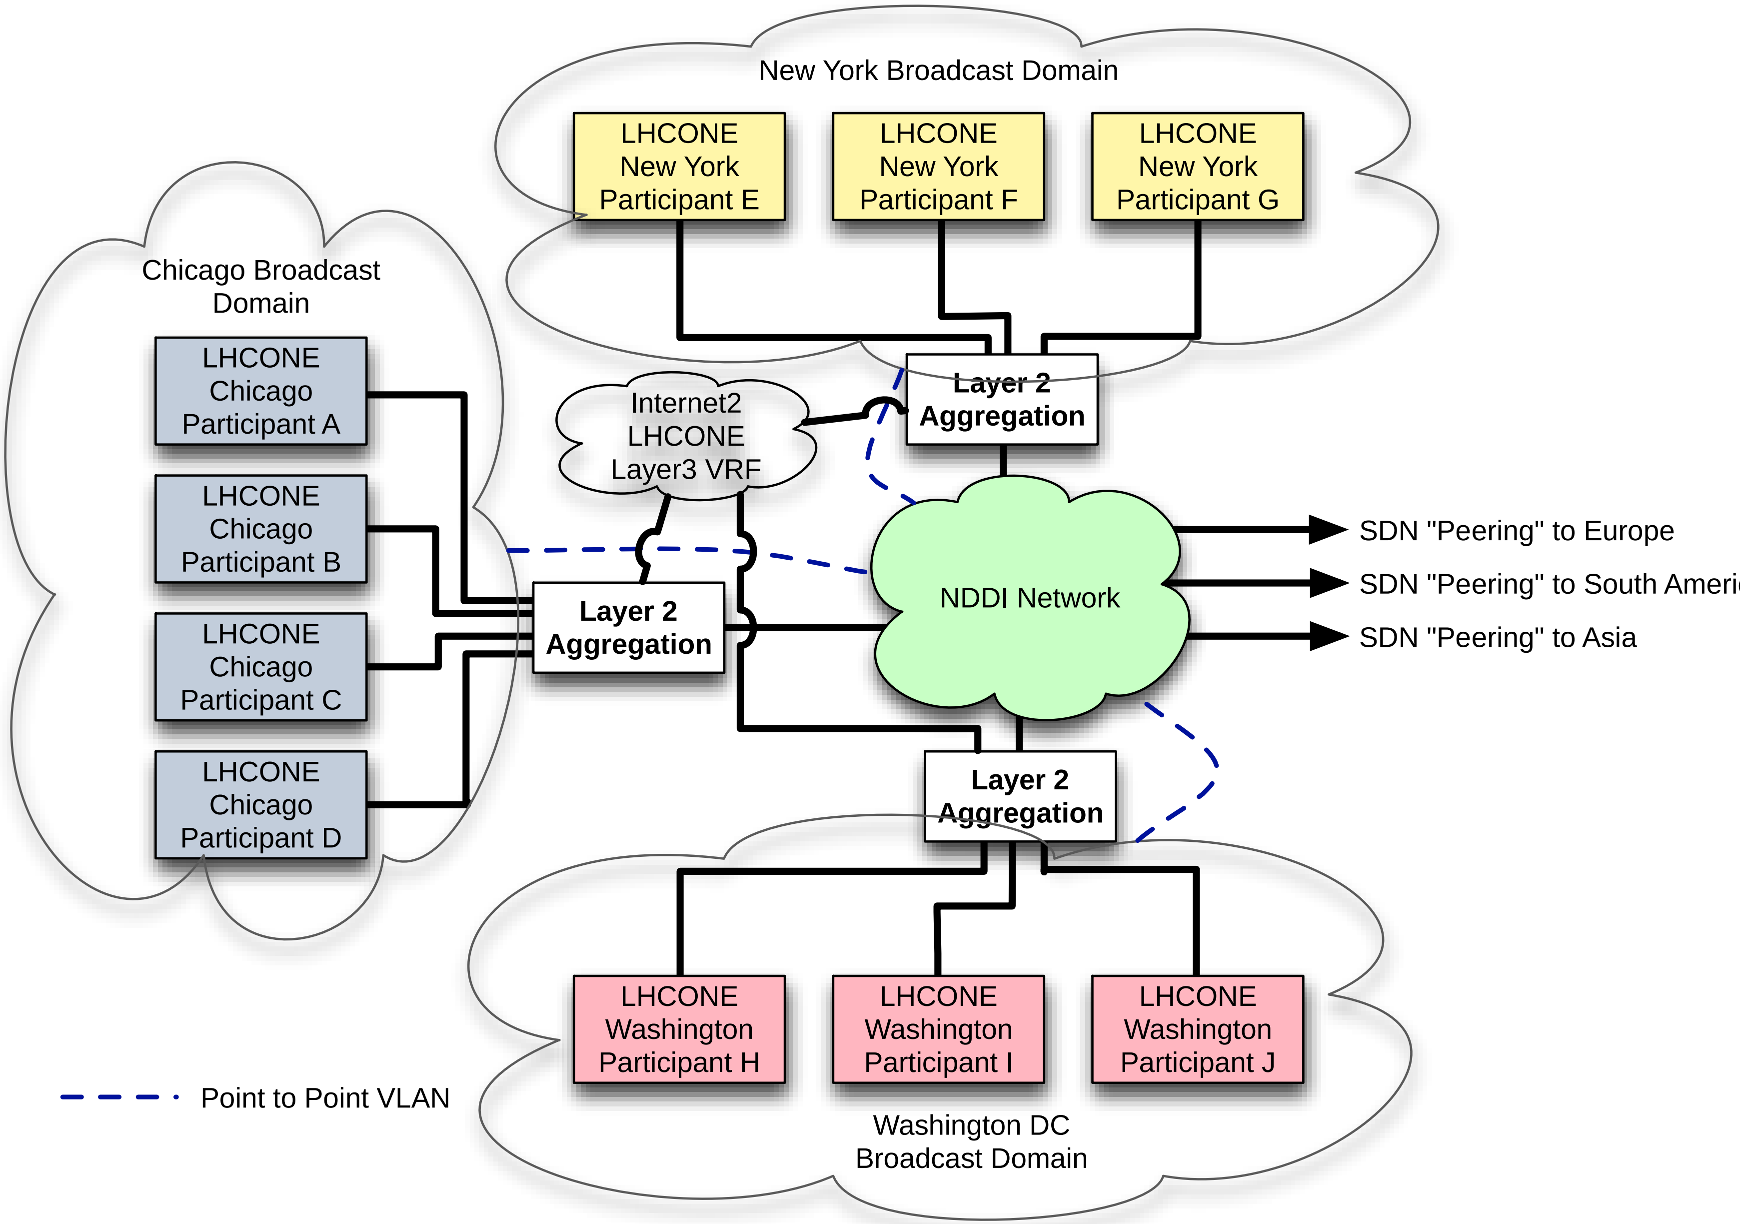
New York Broadcast Domain
LHCONE New York Participant E
LHCONE New York Participant F
LHCONE New York Participant G
Chicago Broadcast Domain
LHCONE Chicago Participant A
LHCONE Chicago Participant B
LHCONE Chicago Participant C
LHCONE Chicago Participant D
Internet2 LHCONE Layer3 VRF
Layer 2 Aggregation
Layer 2 Aggregation
Layer 2 Aggregation
NDDI Network
SDN "Peering" to Europe
SDN "Peering" to South Americ
SDN "Peering" to Asia
LHCONE Washington Participant H
LHCONE Washington Participant I
LHCONE Washington Participant J
Washington DC Broadcast Domain
Point to Point VLAN

Integrated Software Stack

- Physicists should not need to be network engineers
- LHC software stack should include network APIs under the hood
- Goal is software for the LHC community that:
  - Maximizes performance
  - Optimizes use of compute, storage, and network elements
  - Adapts to changing network conditions

Integrated Software Stack

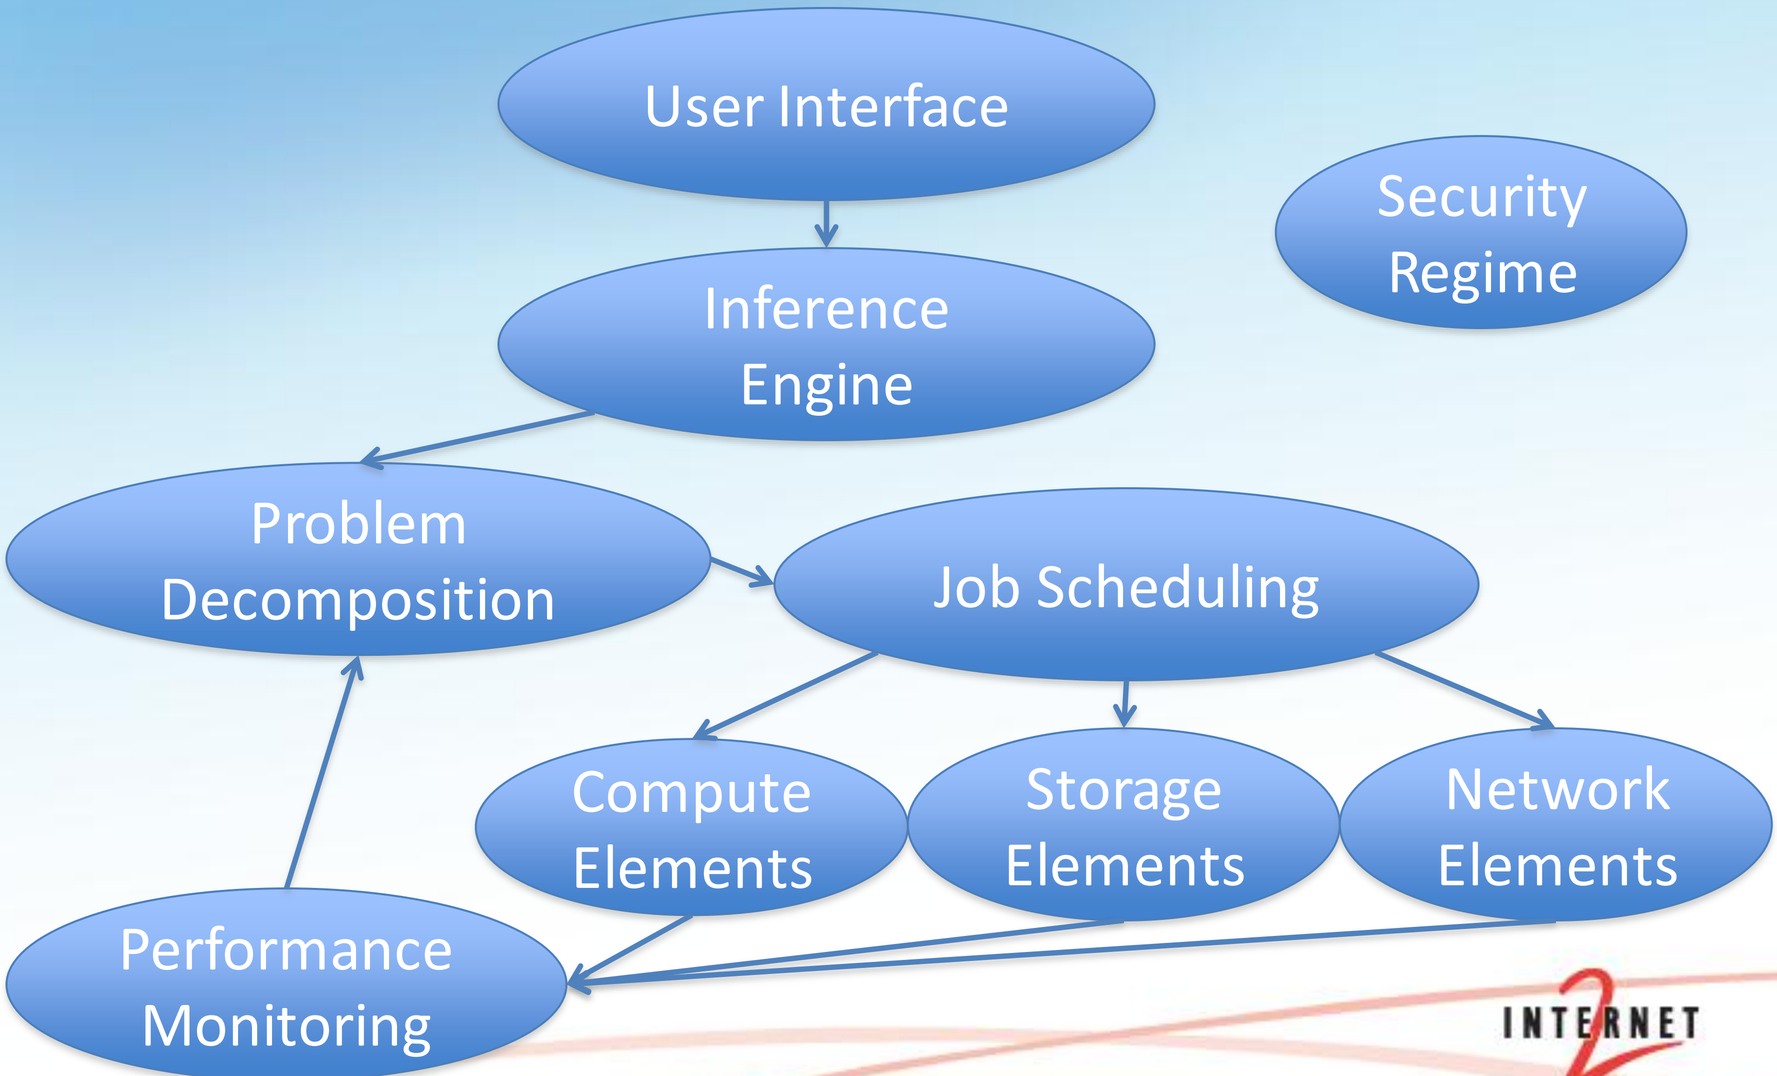

LHCONE – Meeting the needs of the LHC Stakeholders

Eric Boyd
Rob Vietzke

December 1-2, 2011 – LHCONE Architecture

