LHC Open Network Environment
LHC*ONE*

Artur Barczyk
California Institute of Technology
Internet2 Spring Member Meeting
Arlington, April 19th, 2011

LHC Computing Infrastructure

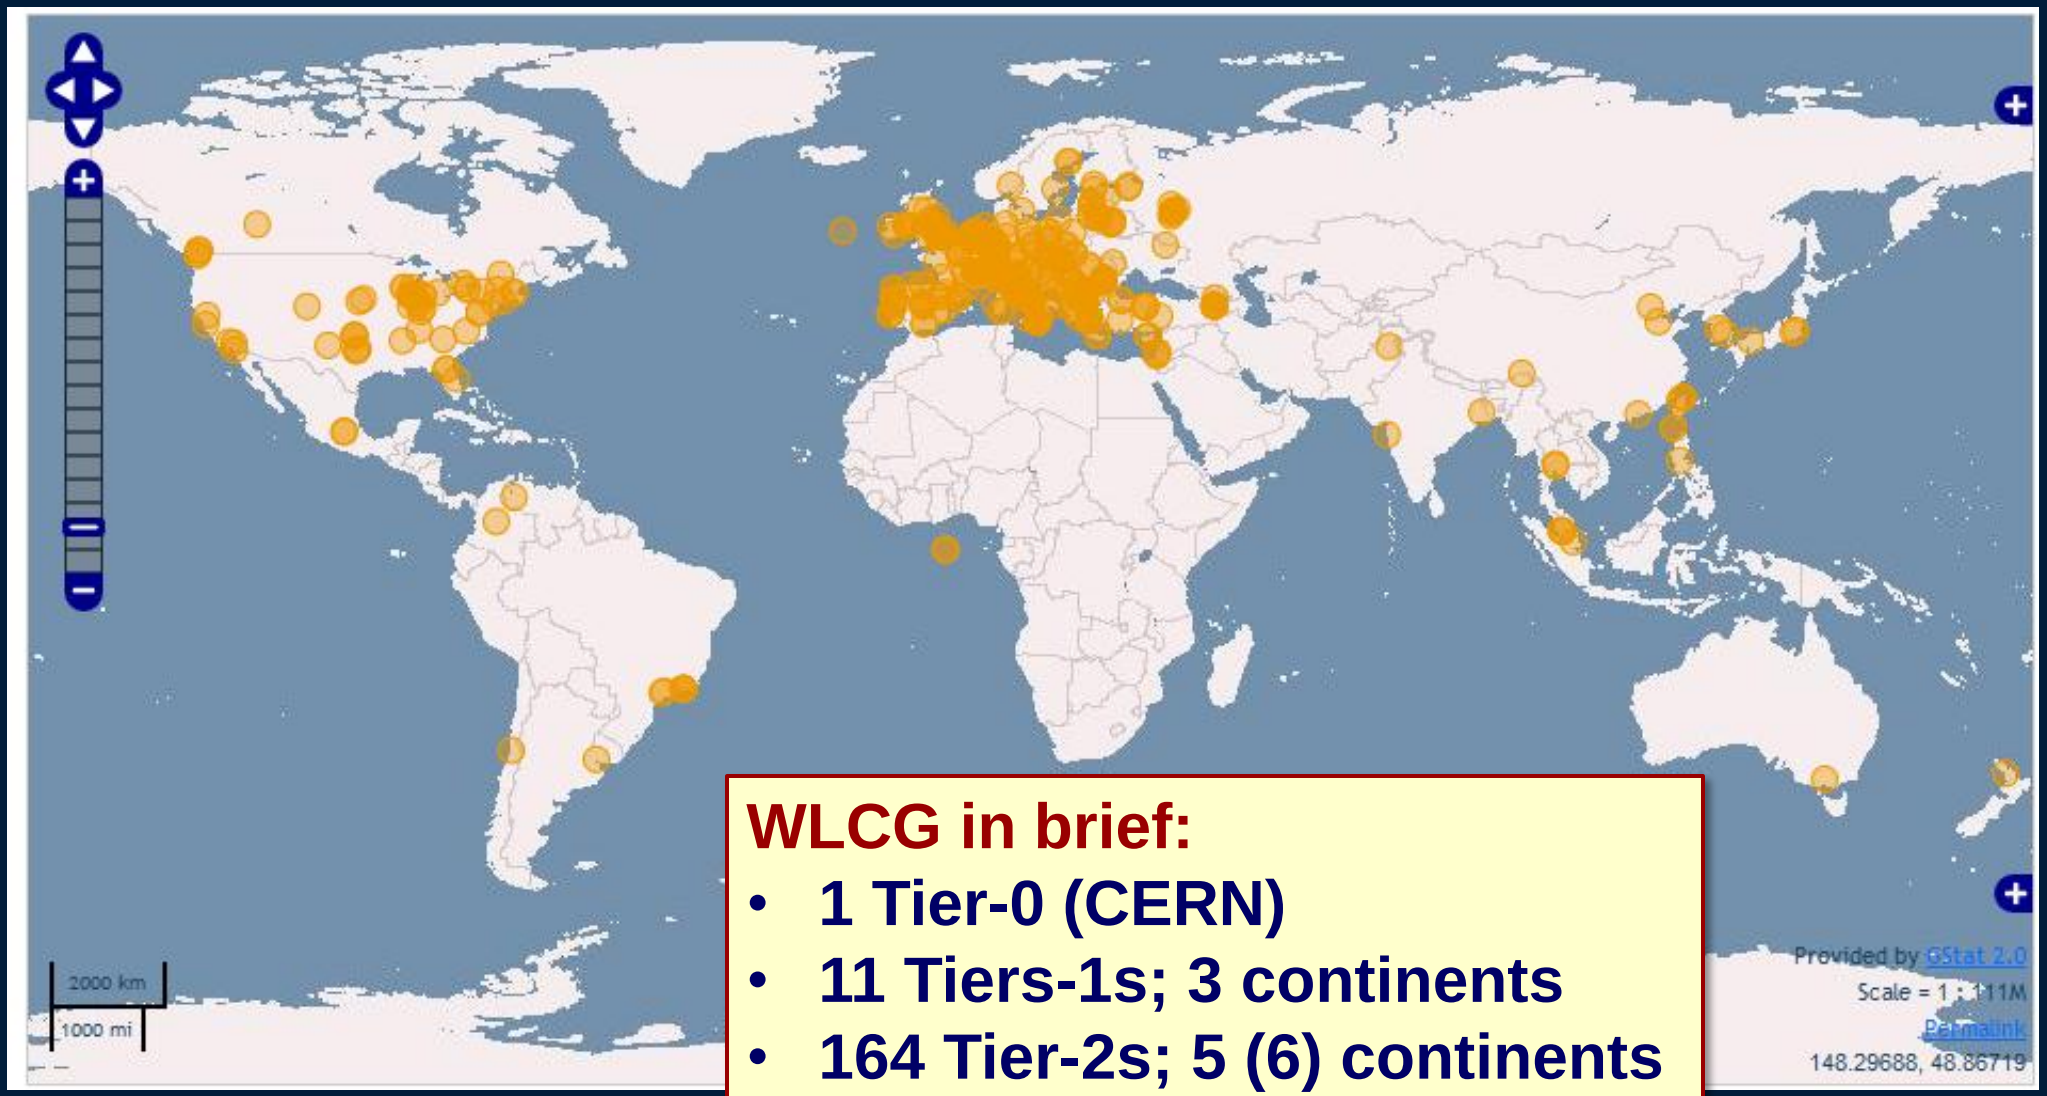

WLCG in brief:

- **1 Tier-0 (CERN)**
- **11 Tiers-1s; 3 continents**
- **164 Tier-2s; 5 (6) continents**

Plus O(300) Tier-3s worldwide

Status of the LHC

Compared to 2010
2.3 x better focusing
50 nsec bunch spacing
(vs 75ns in 2010)

Sa Late: 336 bunch operation

ching to 36 bunch trains

16:55h injection / ramp
and nice intensity distribution

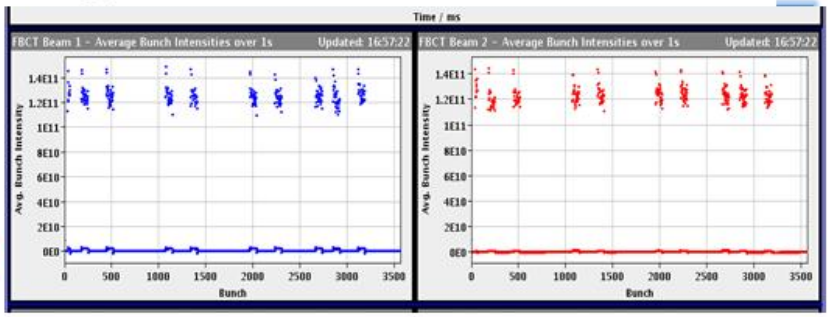

17:40h Luminosity

Last Saturday - Luminosity: 3.6x10^32
(1.7 x max of 2010)

Moving to New Data Models

- Moving away from the strict MONARC model
- 3 recurring themes:
  - Flat(ter) hierarchy: Any site can use any other site as source of data
  - Dynamic data caching: Analysis sites will pull datasets from other sites "on demand", including from Tier2s in other regions
    - Possibly in combination with strategic pre-placement of data sets
  - Remote data access: jobs executing locally, using data cached at a remote site in quasi-real time
    - Possibly in combination with local caching
- Expect variations by experiment

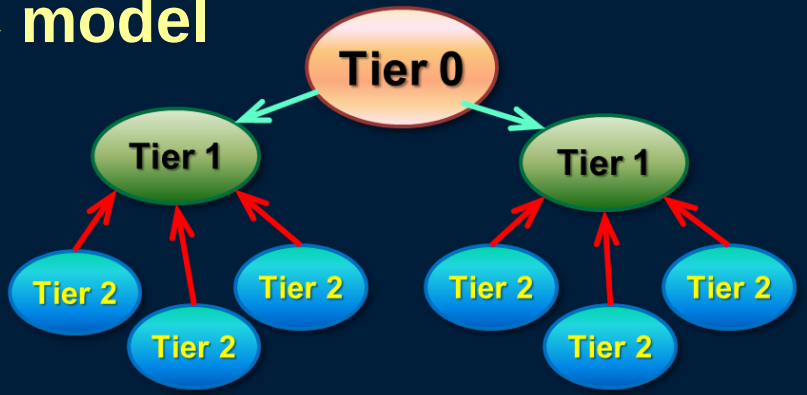

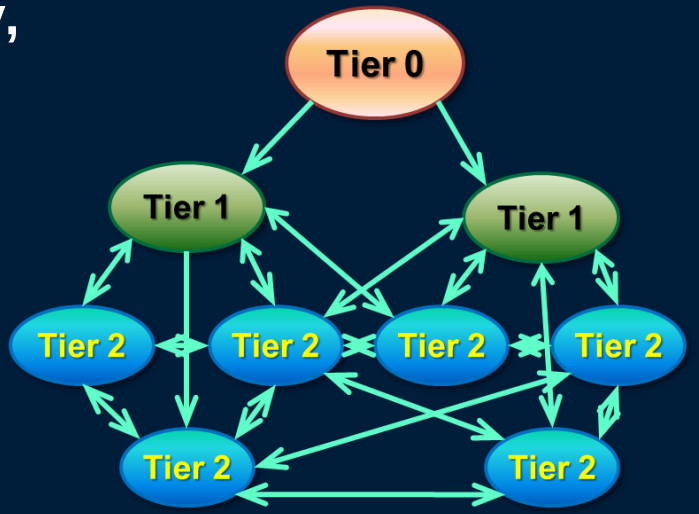

LHC*ONE*

HTTP://LHCONE.NET

The requirements, architecture, services

Requirements summary (from the LHC experiments)

- **Bandwidth:**
  - **Ranging from 1 Gbps (Minimal site) to 5-10Gbps (Nominal) to N x 10 Gbps (Leadership)**
  - **No need for full-mesh @ full-rate, but several full-rate connections between Leadership sites**
  - **Scalability is important,**
    - sites are expected to migrate **Minimal → Nominal → Leadership**
    - Bandwidth growth: Minimal = 2x/yr, Nominal&Leadership = 2x/2yr
- **Connectivity:**
  - **Facilitate good connectivity to so far (network-wise) under-served sites**
- **Flexibility:**
  - **Should be able to include or remove sites at any time**
- **Budget Considerations:**
  - **Costs have to be understood, solution needs to be affordable**

Some Design Considerations

- So far, T1-T2, T2-T2, and T3 data movements have been using General Purpose Network infrastructure
  - Shared resources (with other science fields)
  - Mostly best effort service
- Increased reliance on network performance → need more than best effort
  - Separate large LHC data flows from routed GPN
- Collaboration on global scale, diverse environment, many parties
  - Solution to be Open, Neutral and Diverse
  - Agility and Expandability
    - Scalable in bandwidth, extent and scope
- Allow to choose the most cost effective solution
- Organic activity, growing over time according to needs

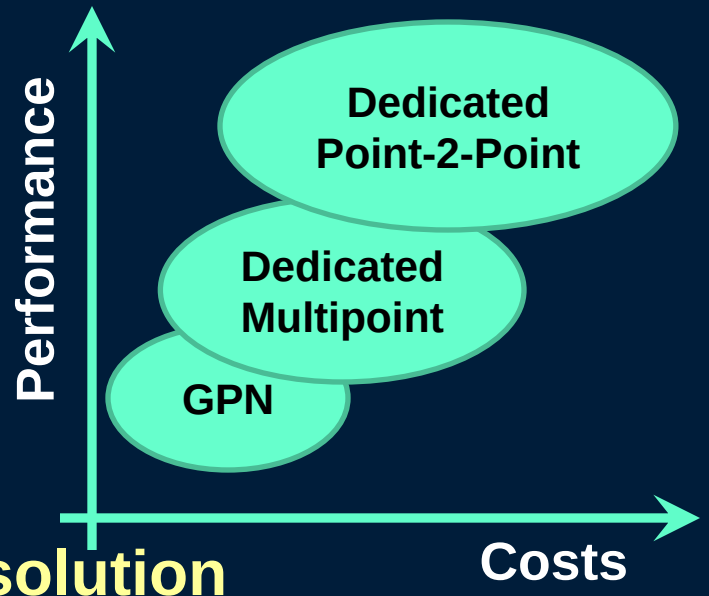

LHC*ONE* Architecture

- **Builds on the Hybrid network infrastructures and Open Exchanges**
  - **To build a global unified service platform for the LHC community**
- **LHC*ONE*'s architecture incorporates the following building blocks**
  - **Single node Exchange Points**
  - **Continental / regional Distributed Exchanges**
  - **Interconnect Circuits between exchange points**
    - Likely by allocated bandwidth on various (possibly shared) links to form LHCONE
- **Access method to LHCONE is chosen by the end-site, alternatives may include**
  - **Dynamic circuits**
  - **Fixed lightpaths**
  - **Connectivity at Layer 3, where/as appropriate**
- **We envisage that many of the Tier-1/2/3s may connect to LHCONE through aggregation networks**

High-level Architecture, Pictorial

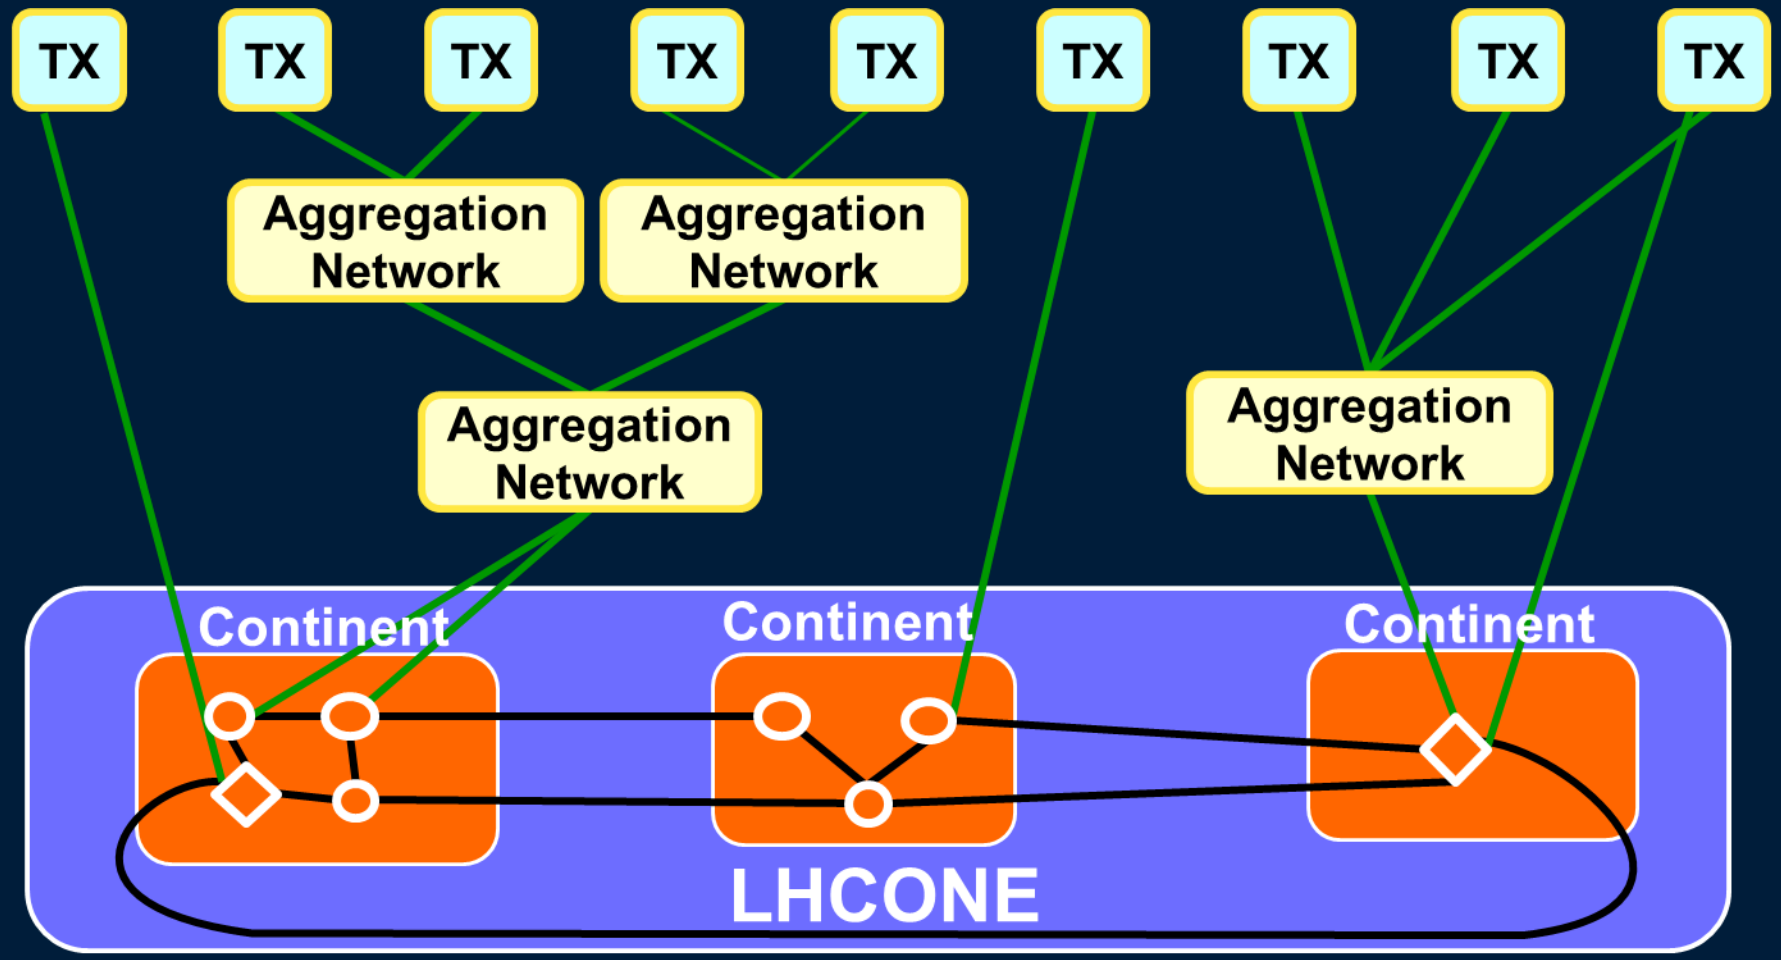

Single node Exchange Point Distributed Exchange Point

LHC*ONE* Network Services
Offered to Tier1s, Tier2s and Tier3s

- **Shared Layer 2 domains: separation from non-LHC traffic**
  - IPv4 and IPv6 router addresses on shared layer 2 domain(s)
  - Private shared layer 2 domains for groups of connectors
  - Layer 3 routing is between and up to the connectors
    - A set of Route Servers will be available
- **Point-to-point layer 2 connections: per-channel traffic separation**
  - VLANS without bandwidth guarantees between pairs of connectors
- **Lightpath / dynamic circuits with bandwidth guarantees**
  - Lightpaths can be set up between pairs of connectors
- **Monitoring: perfSONAR archive**
  - current and historical bandwidth utilization and availability statistics
- **This list of services is a starting point and not necessarily exclusive**
- **LHCONE** does not preclude continued use of the general R&E network infrastructure by the Tier1s, Tier2s and Tier3s - where appropriate

Dedicated/Shared Resources

- LHCONE concept builds on traffic separation between LHC high impact flows, and non-LHC traffic
  - Avoid negative impact on other research traffic
  - Enable high-performance LHC data movement
- Services to use resources allocated to LHCONE

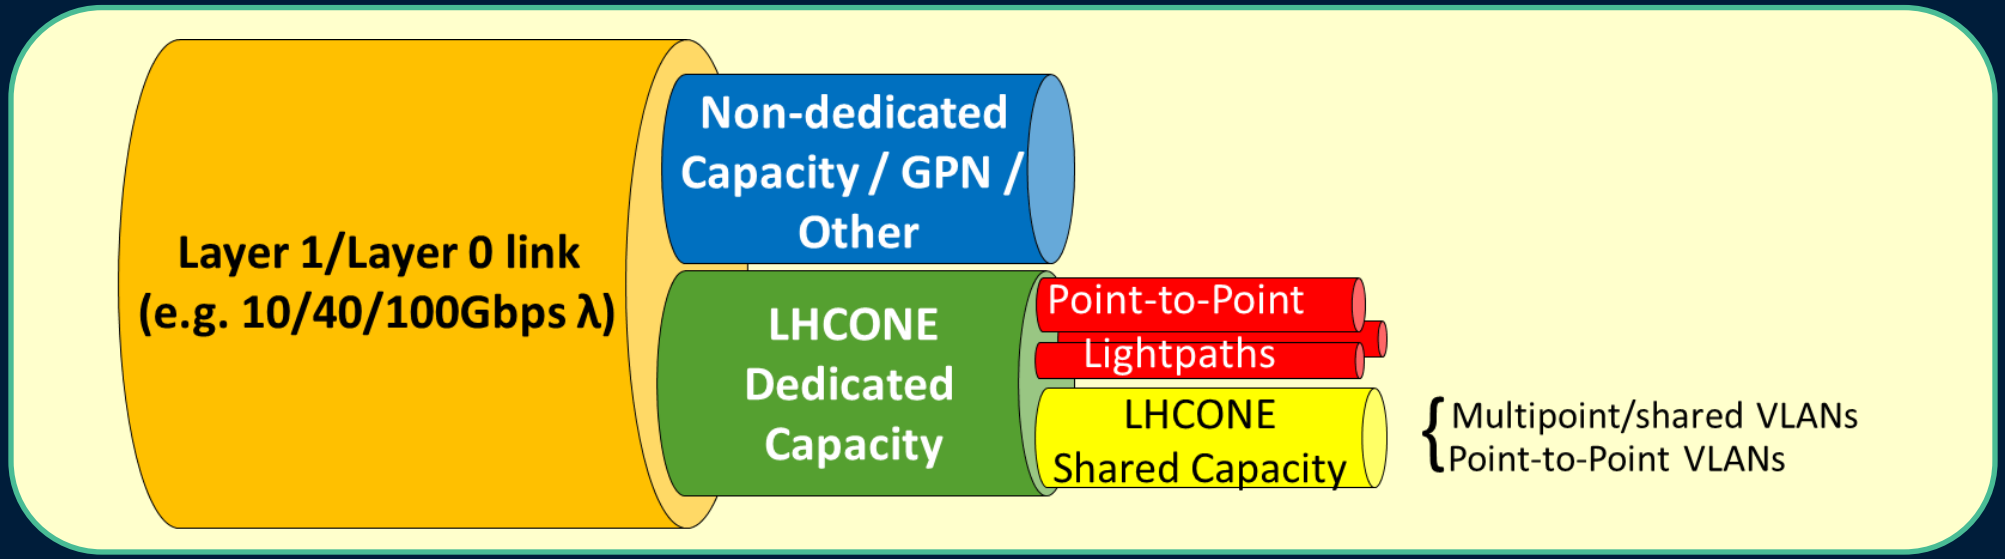

- Prototype might use non-dedicated resources, but need to be careful about evaluation metrics

LHC*ONE* Policy Summary

- **LHCONE policy will be defined and may evolve over time in accordance with the governance model**
- **Policy Recommended for LHCONE governance**
  - Any Tier1/2/3 can connect to LHCONE
  - Within LHCONE, transit is provided to anyone in the Tier1/2/3 community that is part of the LHCONE environment
  - Exchange points must carry all LHC traffic offered to them (and only LHC traffic), and be built in carrier-neutral facilities so that any connector can connect with its own fiber or using circuits provided by any telecom provider
  - Distributed exchange points: same as above + the interconnecting circuits must carry all the LHC traffic offered to them
  - **No additional restrictions can be imposed on LHCONE by the LHCONE component contributors**
- **The Policy applies to LHCONE components, which might be switches installed at the Open Exchange Points, or virtual switch instances, and/or (virtual) circuits interconnecting them**

LHC*ONE* Governance Summary

- Governance is proposed to be similar to the LHCOPN, since like the LHCOPN, LHCONE is a community effort
  - Where all the stakeholders meet regularly to review the operational status, propose new services and support models, tackle issues, and design, agree on, and implement improvements
- Includes connectors, exchange point operators, CERN, and the experiments, in a form to be determined.
- Defines the policies of LHCONE and requirements for participation
  - It does not govern the individual participants
- Is responsible for defining how costs are shared
- Is responsible for defining how resources of LHCONE are allocated

LHCONE Status

- **Start with a prototype, building on a Use Case document provided by CMS and Atlas**
  - **Includes a number of sites to participate in the initial stage**
- **Teams in North America and Europe are working on the prototype design**
- **In the meantime, a starting point has been deployed**
  - **"Shared vlan" service**
  - **4 Open Exchange Points**
    - CERNLight, Netherlight, MANLAN and Starlight
  - **Dedicated core capacity**
    - SURFnet, US LHCNet
  - **Route server operational (at CERNLight)**
  - **Validated with two LHC sites (CERN-T1, Caltech-T2)**

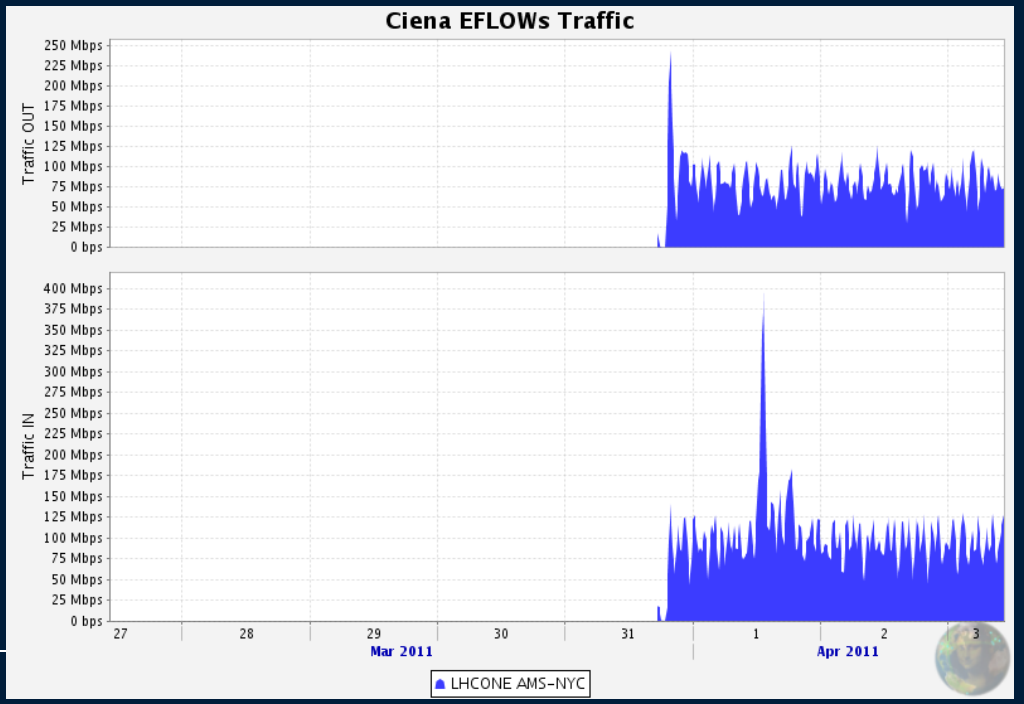

Next Steps

- **Prototype implementation (Seed)**
  - CMS & Atlas prepared a use case with a mix of Tier2 and Tier3 sites (including all Tier1s)
  - Small engineering groups are working out prototype designs in Europe, North America, Asia and other regions
  - Build initial implementation integrating all these efforts (target July 2011)
- **Follow-up roadmap**
  - In parallel with prototype implementation
    - Refine governance model
    - Refine service and policy definitions
    - Refine architecture
    - Grow LHCONE to full scale
- **LHCONE will grow "organically", as needs arise**
  - and as funding is or is made available

LHCONE - Open Environment

- LHCONE creates an Open Network Environment, with the intention that the LHC sites can connect
  - When the site decides
  - Where the site decides
  - How the site decides
- Driven by the experiments' and sites' requirements

- Connection policy is not in the core, it's between the sites
  - Sites can choose to peer with a subset or all other sites
  - Most probably in a way coordinated by the experiments' operations teams

THANK YOU!

http://lhcone.net

Artur.Barczyk@cern.ch