ESnet

ENERGY SCIENCES NETWORK

LHCONE Update

Michael O'Connor moc@es.net

ESnet

Network Engineering

LHCOPN/LHCONE workshop

CERN Geneva, CH

January 14, 2020

U.S. DEPARTMENT OF ENERGY

Office of Science

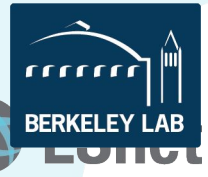

BERKELEY LAB

LHCONE Architecture

In general each NSP/NREN implements LHCONE as a virtual overlay network or VRF.

L3 VPN Advantages

- NREN **connections** similar to traditional Internet, for sites and peers.
- **Scales** well, using a network of networks model.
- **Robust routing** around failures.
- **Cost effective** infrastructure sharing, upgrades, maintenance, operations and monitoring.

Unique LHCONE qualities

- Limited access **science only** network.
- **Distributed** management model.

ESnet

NSP Requirements and Responsibilities Overview
NSP BGP Import Policy

Prefix Lists will be negotiated between connecting institutions and their NSP within the constraints imposed by the LHCONE AUP.

LHCONE NSPs have agreed to to configure:

1. BGP import filters
2. Source address packet filters

End sites are encouraged to implement source address filters at their edge in order to count their own unroutable LHCONE packets. NSPs will generally discard these packets without informing the site.

Connecting institutions/sites will not add prefixes to the LHCONE routing table without direct cooperation with their NSP.

ESnet

NSP Requirements and Responsibilities Overview Packet Filtering

All LHCONE Traffic is subject to the following conditions:

- Traffic injected into the LHCONE must only be originated from addresses within an LHCONE routable prefix
- Only address ranges present in the LHCONE routing table should be transported on the network
-

Objective: In order to maintain route symmetry and access control, each NSP will implement policy and packet filters to manage their connected customer address prefix ranges.

- Ensures that a return route exists in the LHCONE network
- Blocks spoofed packets (Similar to BCP 38)

https://twiki.cern.ch/twiki/pub/LHCONE/LhcOneVRF/LHCONEconnectionguide-1.2.pdf

ESnet

Edge Filtering Special Case

L3 Network Exchange Fabrics

An exchange is like an NSP:

- BGP import filtering
- Packet filtering
- Community based BGP filtering

An exchange is like a site:

- Require the full LHCONE table via a transit NSP
- Packet filters are configured and require maintenance

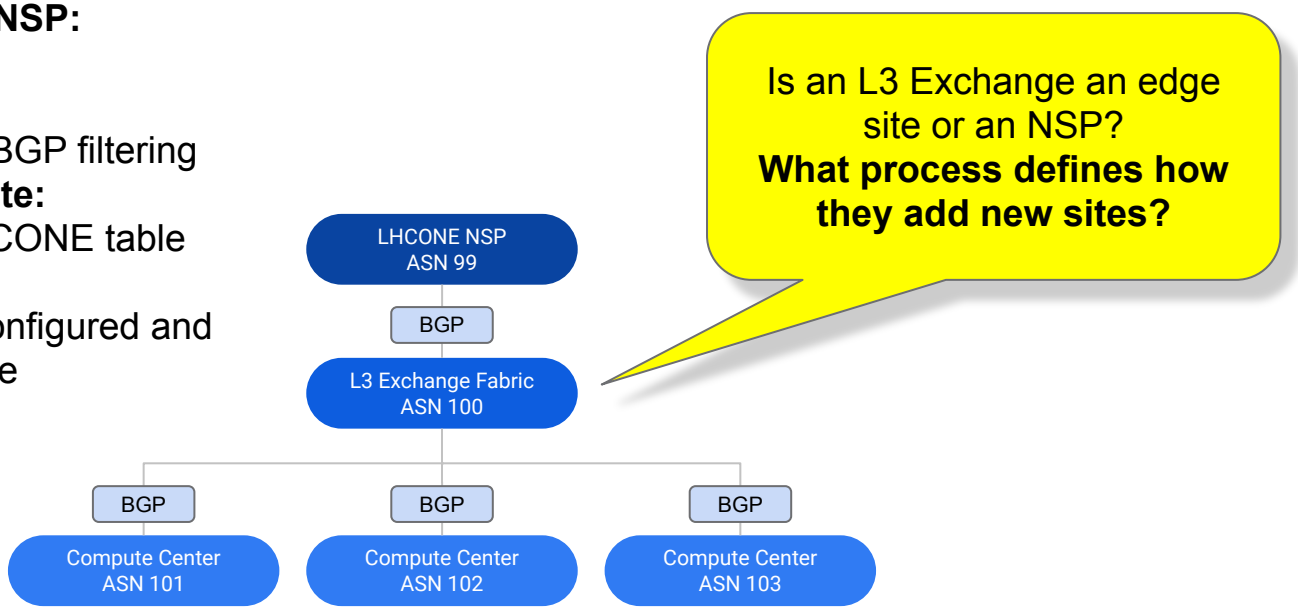

Indiana GigaPOP and SOX are ESnet examples.

- Will L3 Exchange Fabrics implement and maintain LHCONE specific services?
- Should there be an LHCONE defined role for these network organizations?
- Are they permitted to attach new sites?

ESnet

Unroutable Packets
Distributed Policy Failures

Unroutable traffic originates at sources that have no return path in the LHCONE routing table.

Detection: potential approaches

- Regularly scheduled monitoring?
- Periodic NSP self run audits?

Prevention

- Edge Site filter configuration
  - RPF → too strict?
  - Templated route policies & packet filters

Information

- Regular AUP updates to address special cases
- Improve “onboarding” process for new networks to provide comprehensive management and operations guidance
- Sharing configuration best practices

ESnet

LHCONE Interdomain Trust
BGP Community Based Filtering

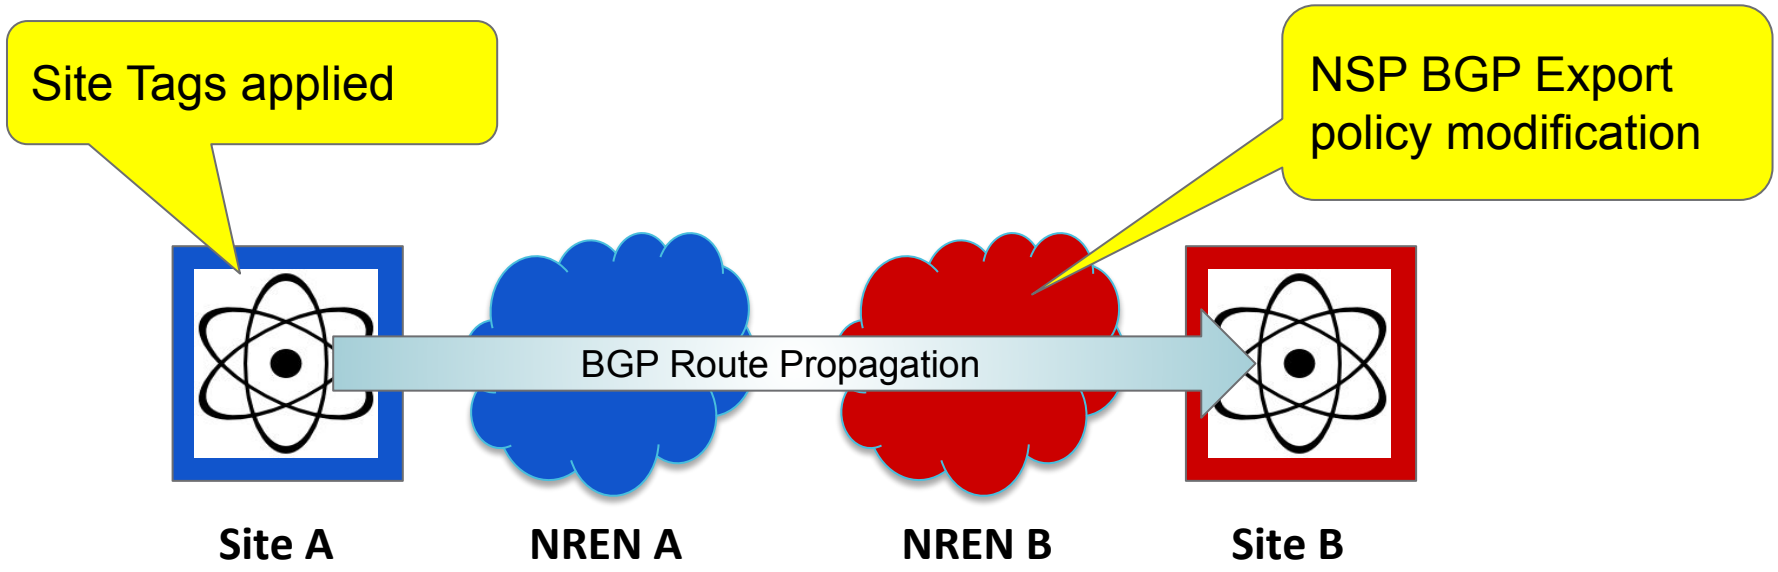

The LHCONE community based filtering system is a successful example of an interdomain trust agreement.

LHCONE NSP's dynamically adjust their export policies based on BGP communities set by remote collaborating compute centers.
Why does it work:

- A site only has the ability to add communities to routes that it originates.
- The NSPs trust the tags applied by remote collaborating sites.

ESnet

LHCONE BGP Community Tagged Routes 1/13/2020

```
ITEP(2148)
        194.85.68.0/23     DO NOT ANNOUNCE(65010) to SLAC(3671) VANDERBILT(7212) UOK(25776)
UTNET(2501)
        133.11.127.244/30       PREPEND 1X(65001) to CERN(513)
        133.11.255.181/32       PREPEND 1X(65001) to CERN(513)
        133.11.59.48/28         PREPEND 1X(65001) to CERN(513)
        150.99.198.220/30       PREPEND 1X(65001) to CERN(513)
        157.82.112.0/21         PREPEND 1X(65001) to CERN(513)
HEPNET-J(2505)
        202.13.197.192/26       PREPEND 1X(65001) to CERN(513)
        202.13.203.128/25       PREPEND 1X(65001) to CERN(513)
SINET(2907)
        117.103.111.128/30      PREPEND 1X(65001) to CERN(513)
        133.11.254.16/30        PREPEND 1X(65001) to CERN(513)
        138.44.226.12/31        PREPEND 1X(65001) to CERN(513)
        144.206.255.144/30      PREPEND 1X(65001) to CERN(513)
        202.13.223.192/29       PREPEND 1X(65001) to CERN(513)
        202.13.223.52/30        PREPEND 1X(65001) to CERN(513)
        202.180.40.0/30         PREPEND 1X(65001) to CERN(513)
        202.180.40.4/30         PREPEND 1X(65001) to CERN(513)
        62.40.126.176/31        PREPEND 1X(65001) to CERN(513)
        62.40.126.22/31         PREPEND 1X(65001) to CERN(513)
FNAL(3152)
        131.225.13.128/25  DO NOT ANNOUNCE(65010) to ASGARR(137) UIUC(38)
        131.225.67.0/24    DO NOT ANNOUNCE(65010) to ASGARR(137) UIUC(38)
        131.225.69.0/24    DO NOT ANNOUNCE(65010) to ASGARR(137) UIUC(38)
ERX-HEPCAS-AS(3460)
        202.122.32.160/27       PREPEND 1X(65001) to CERN(513)
        202.122.32.45/32        PREPEND 1X(65001) to CERN(513)
        202.122.33.0/24         PREPEND 1X(65001) to CERN(513)
        202.122.35.0/24         PREPEND 1X(65001) to CERN(513)
        202.122.36.0/24         PREPEND 1X(65001) to CERN(513)
        202.38.128.0/24         PREPEND 1X(65001) to CERN(513)
        202.38.129.0/24         PREPEND 1X(65001) to CERN(513)
NSCKIPT(35296)
        193.239.180.128/27 DO NOT ANNOUNCE(65010) to ERX-ERNET(2697) KIAE-TRANSIT(57484)
        193.239.180.208/29 DO NOT ANNOUNCE(65010) to ERX-ERNET(2697) KIAE-TRANSIT(57484)
CERN(513)
        128.142.0.0/16          PREPEND 2X(65002) to KIAE(59624)
        188.184.128.0/17        PREPEND 2X(65002) to KIAE(59624)
        188.185.128.0/17        PREPEND 2X(65002) to KIAE(59624)
        188.185.48.0/20         PREPEND 2X(65002) to KIAE(59624)
AARNET(7575)
        138.44.13.124/30        PREPEND 1X(65001) to CERN(513)
```

ESnet

Tier1 connectivity

LHCOPN and LHCONE Networks

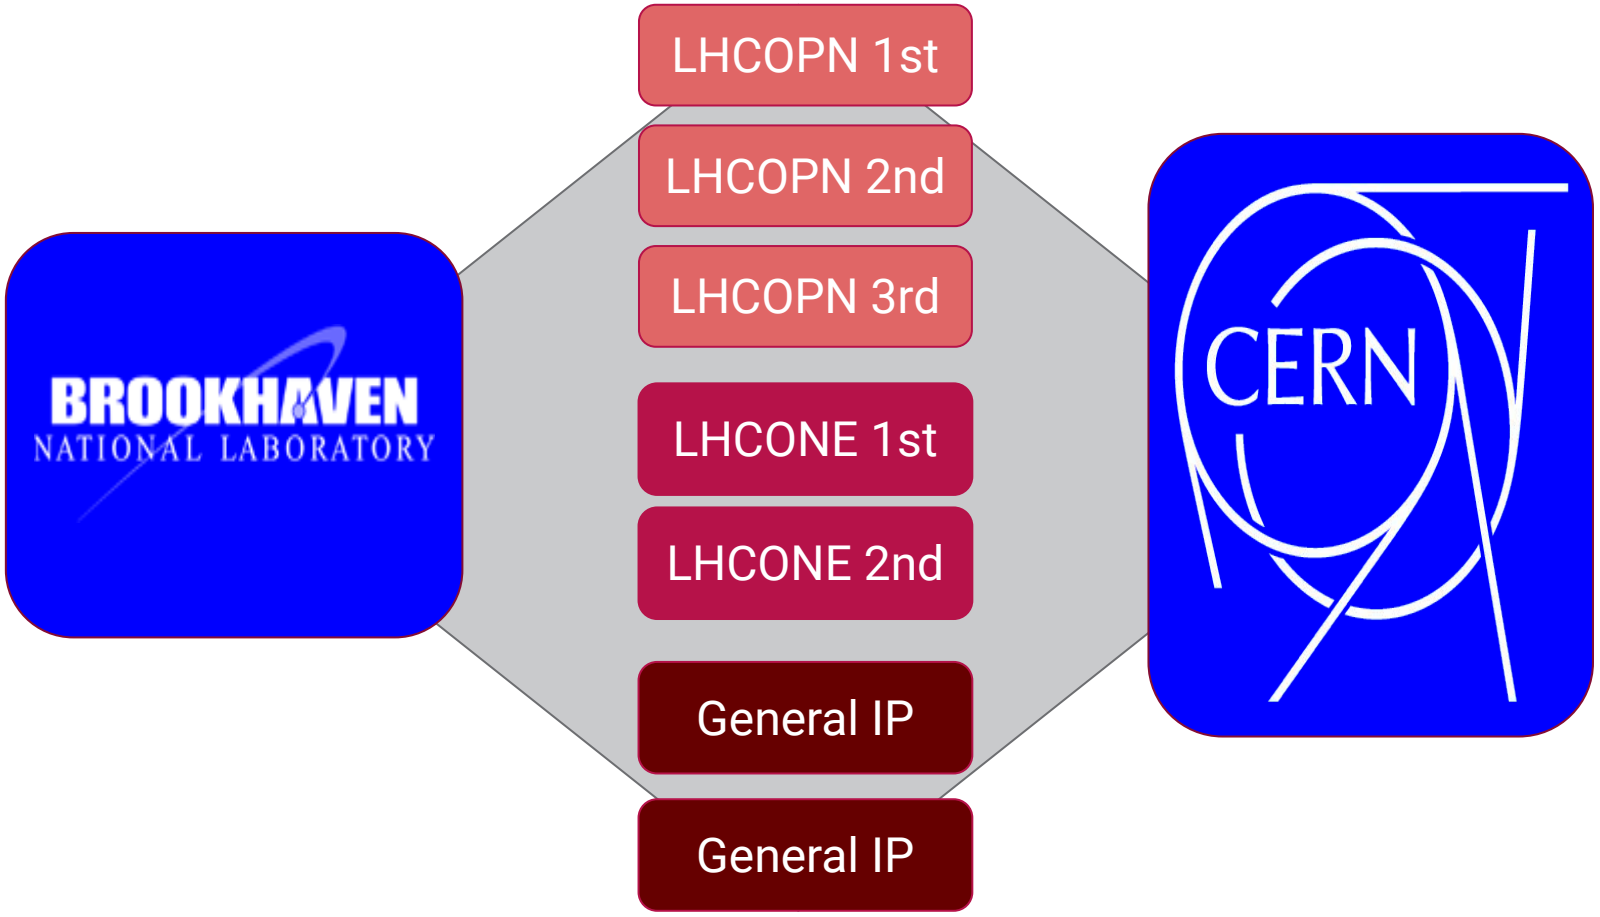

Paths between BNL and Cern

ESnet

Brookhaven Lab Policy Routed Perimeter With 100G Diversity

Background information

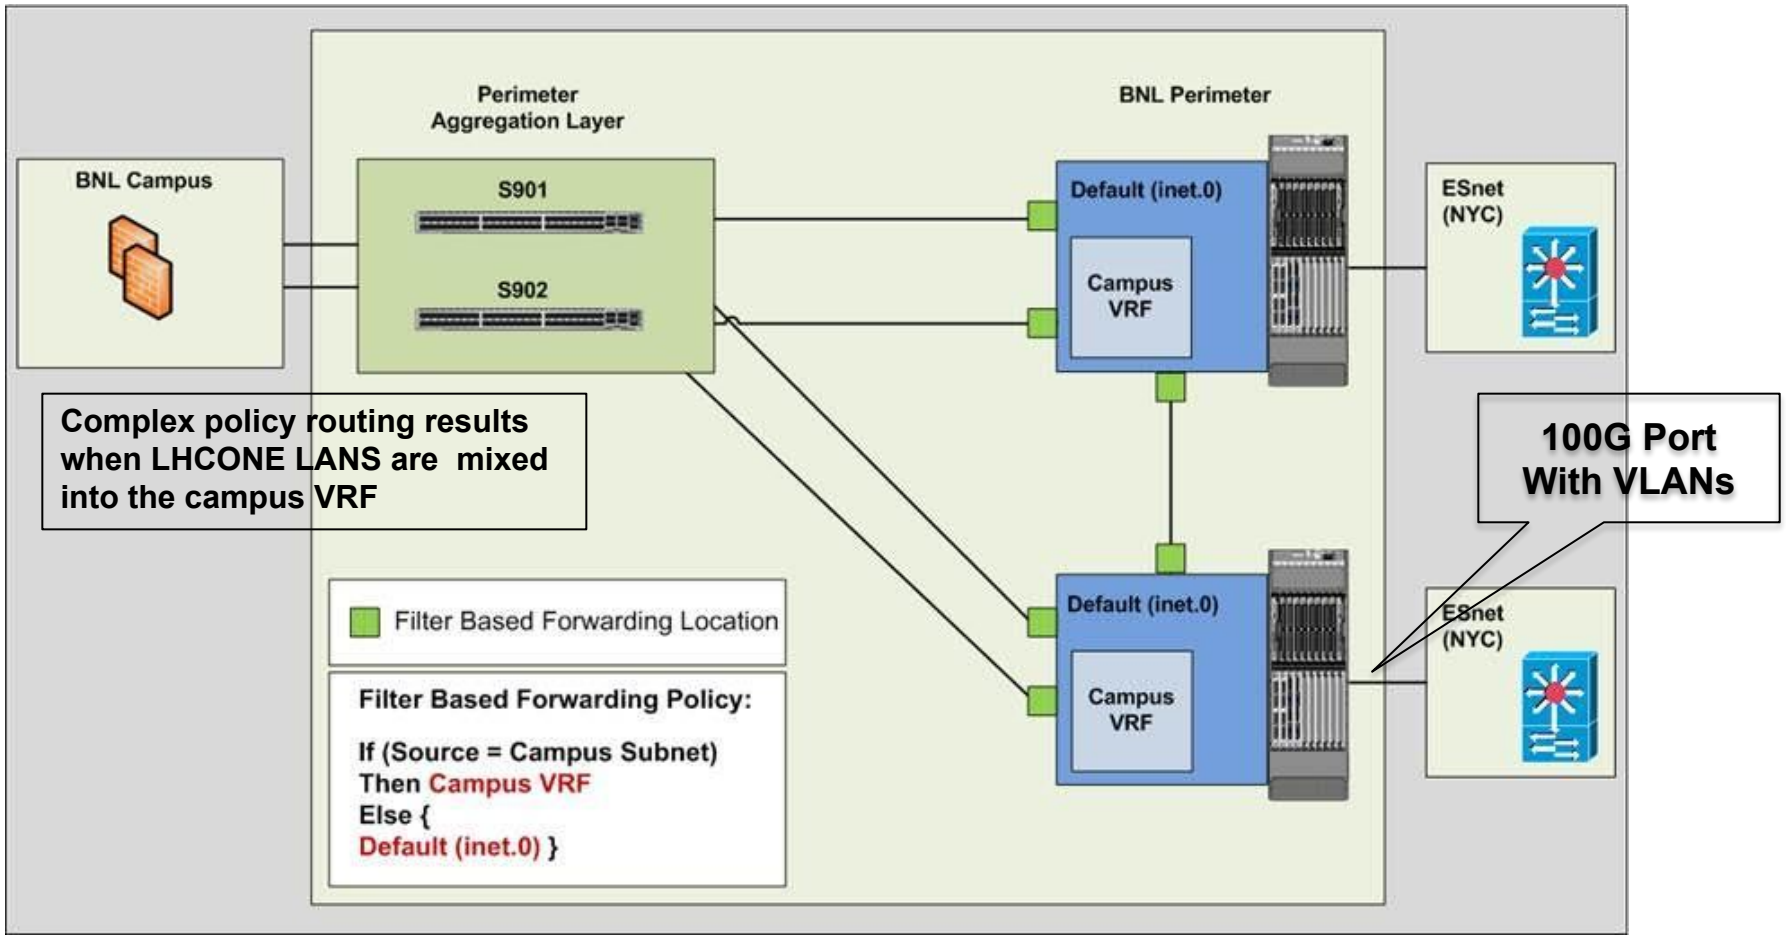

BNL checks **all** egress traffic against the LHCONE source route policy, if the source is from a Campus LAN (not LHCONE) it takes a conditional default to ESnet general IP. LHCONE sources egress using destination routing toward the best path.

ESnet

Policy Routed General IP Egress

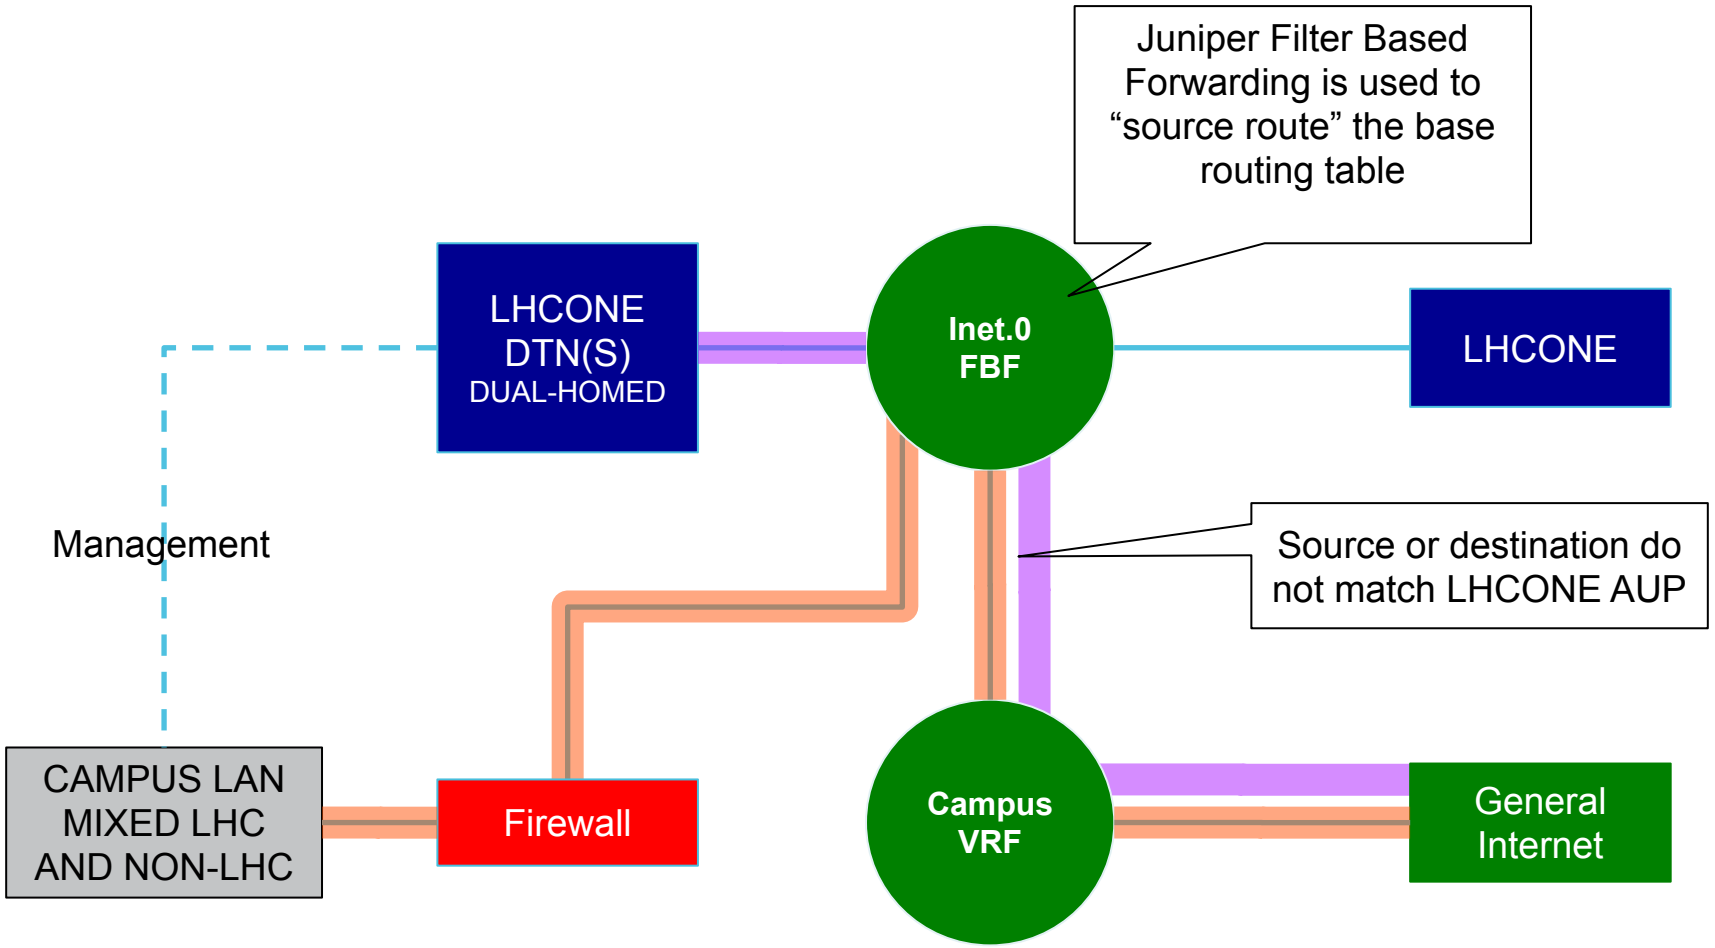

Mixing the Campus and LHCONE LANs required multi-stage routing for generic IP traffic.

ESnet

LHCONE Site Example

Destination Routing

- The architecture recommended by the LHCONE community for new sites.
- PBR is not required in this architecture.

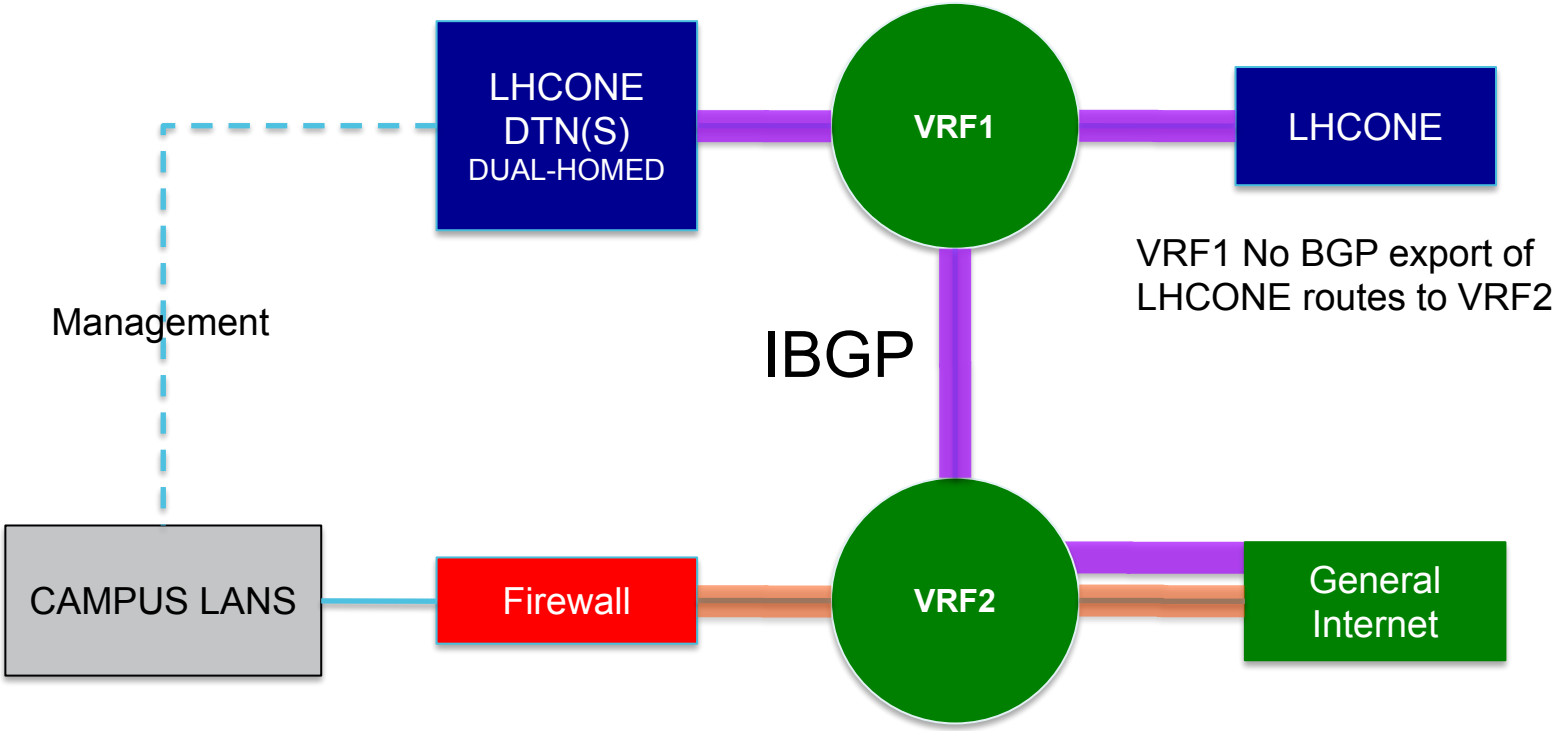

If the LHCONE can be separated into a VR or VRF then standard destination based routing can be used

ESnet

Multiple Overlay Networks

Scaling For Additional Overlay Networks

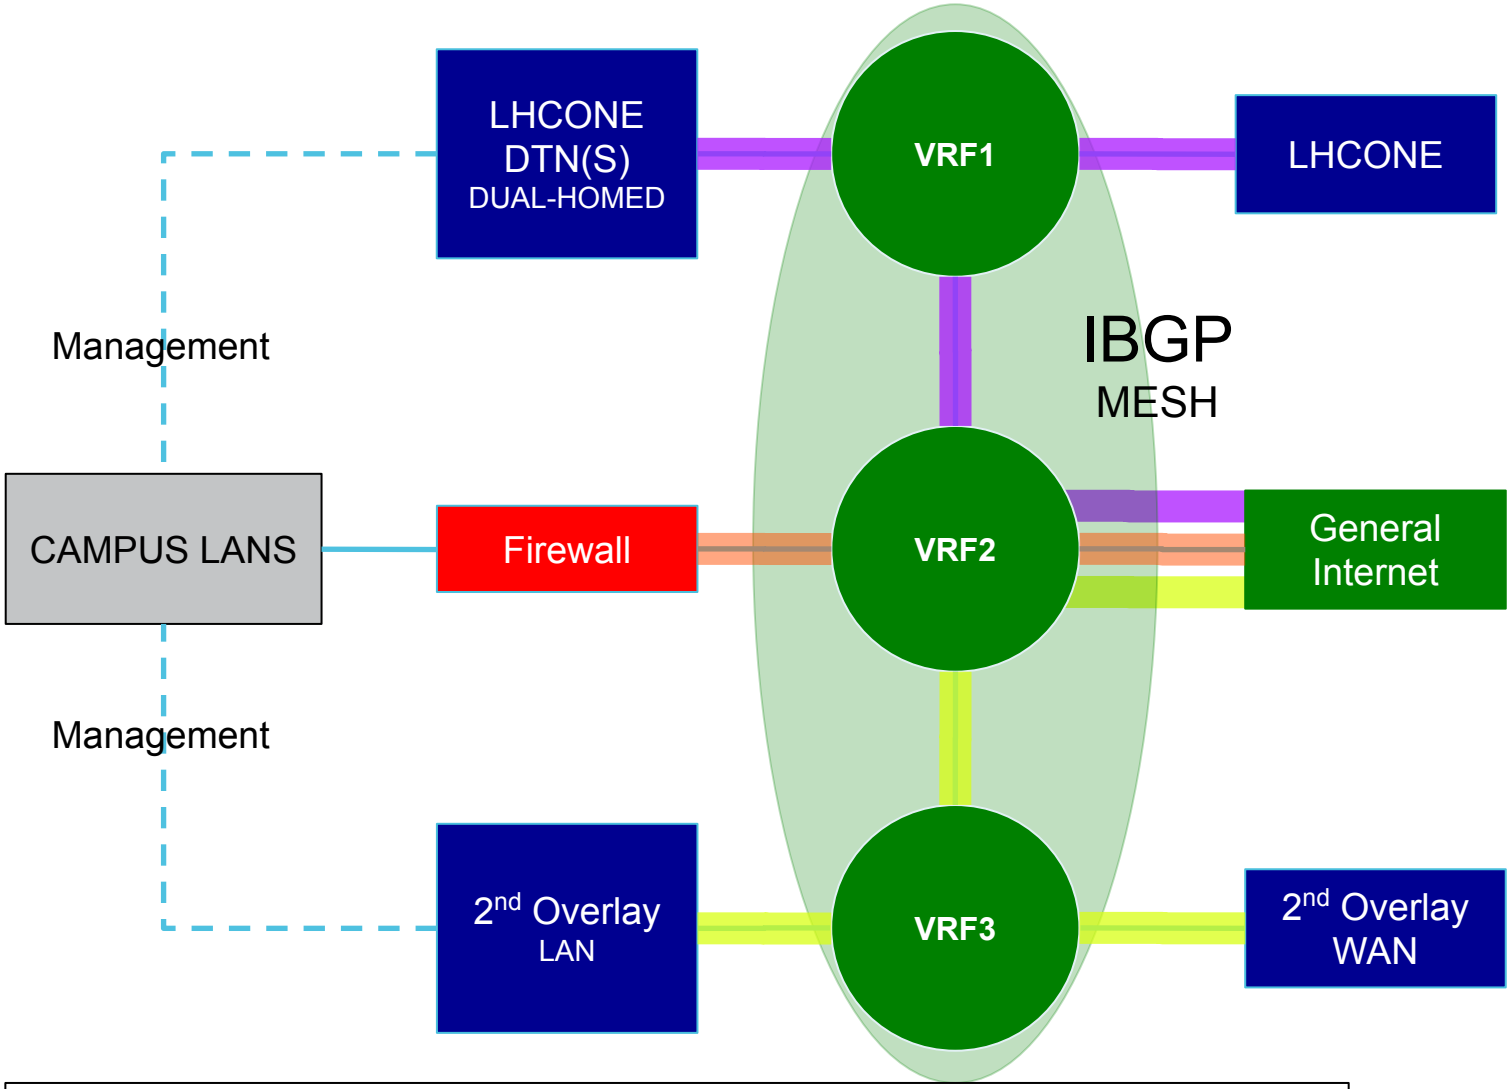

The additional overhead of protocol configuration pays back in scalability and powerful routing policy control.

ESnet

ESnet

ENERGY SCIENCES NETWORK

Questions?

Michael O'Connor moc@es.net

ESnet

Network Engineering

Conference

Location

Date

U.S. DEPARTMENT OF ENERGY

Office of Science

BERKELEY LAB