LHCONE Site Connection Guidelines

Mike O’Connor, Network Engineer

ESnet Engineering Group

LHCOPN and LHCONE Joint Meeting

Paris (FR)

June 17, 2012

LHCONE Site Provisioning Guidelines

The Version 1.0 draft of the LHCONE Site Provisioning Guidelines are available on the LHCONE Web at twiki.cern.ch.

These guidelines were drafted following the LHCONE meeting in Washington DC, Jan. 2013.

1. Connection guidelines to ensure symmetry

2. NSP Policy and Filtering

3. Site connection information

Site Connection Guidelines

Connection guidelines to ensure route symmetry at the DMZ

- Select the local LAN address ranges that are able to participate in LHCONE. Work with your NSP to advertise these address range prefixes into LHCONE.
- Be prepared to accept all BGP route prefixes advertised by the LHCONE community through the LHCONE connection.
- Ensure that only hosts in your locally defined LHCONE ranges have the ability to forward packets into the LHCONE network.
- Routing preference: Ensure that the LHCONE paths are preferred over general R&E IP paths or any other path that has employs a state-full firewall.

Reverse Path Forwarding (RPF) filtering is suggested as a dynamic access control method for sites that wish to filter packets. This will allow only packets from LHCONE sources to enter your LHCONE connected interface. A dynamic method is recommended since static filters will inevitably cause periods of instability as the LHCONE routing table is updated.

Site Information Questionnaire

Site Information:

Site Name

AS Number:

Network representative

Email

Operations Organization

Email

Site Information Questionnaire (Cont.)

LHCONE Connection Information:

- From which NREN(s) does your organization receive LHCONE connectivity?
- Has your site implemented a separate WAN connection for LHCONE? As part of a science DMZ for instance.
- Please describe your sites connection to LHCONE in terms of: Bandwidth, shared/dedicated, location.
- What proportion (by host address) of the prefixes in your ASN do you announce to LHCONE and by what mechanism are they separated from the general purpose LAN?
- As per LHCONE Guideline do you only announce the prefixes assigned to LHC resources?
- What architectural model or protocol policy techniques are in place to ensure routing symmetry to the LHCONE in accordance with the LHCONE Site Provisioning Guidelines. If your NREN performs this service, please respond “NREN provided service”.

Site Information Questionnaire (Cont.)

LHCOPN and General Purpose R&E connection information:

- Does your site also connect to LHCOPN?
- Do you prefer LHCOPN Prefixes over LHCONE?
- Do you prefer LHCONE Prefixes over General Purpose IP?

Policy Blowback

The response and commentary following the distribution of the LHCONE Site Provisioning Guidelines were limited, typically implying general acceptance of the policies within the community. However I did receive this very interesting implementation case study.

Commentary from the University of Chicago Network Engineer group based on UOCs' diligent attempt to implement the policies stated in our Site Provisioning Guidelines.

"I just read the mail regarding the LHCONE VRF site connection guide and I'm a bit confused. In working with our LHCONE connection over the last year or so we've run into nothing but trouble with the policies as they stand. I've heard similar grumbles from other Network Engineers."

-Ryan Harden

UOC Experience in Summary

(See notes for original text)

1. They immediately ran into problems upon turn-up of LHCONE.
2. The Policy Based Routing at several end-sites broke the UOC connections to those institutions when UOC started preferring LHCONE.
3. UOC shutdown the LHCONE peering and requested that all of the unreachable end-sites fix their policies.
4. UOC LHC/ATLAS contacts were unable to easily determine which subnets would connect to LHCONE and which would not.
5. After a month of endless coordination and instability UOC announced their entire campus address space into LHCONE eliminating these issues.

Upon connecting to LHCONE, detailed and timely knowledge of collaborating network state became very important, creating great instability.

LHCONE Routing Table Prefix Size

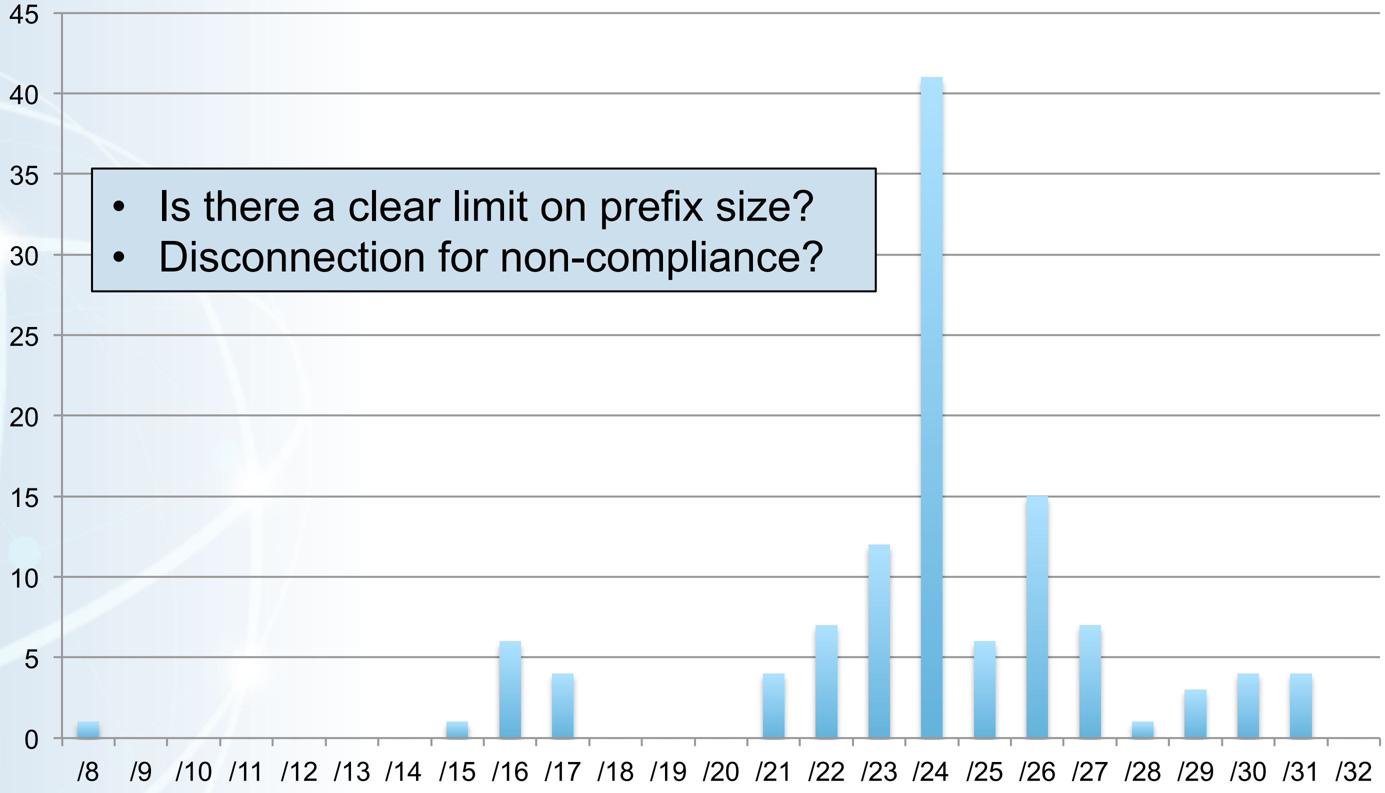

- Is there a clear limit on prefix size?
- Disconnection for non-compliance?

Since the routing table is very small (116 routes), prefixes with netmasks greater than /24 are perfectly reasonable, in contrast to general Internet practice.

A=A’ Method

Advertise ALL LAN prefixes to LHCONE

Effectively uses a single CE router without the need for PBR. Remember to prefer LHCONE paths over general R&E IP.

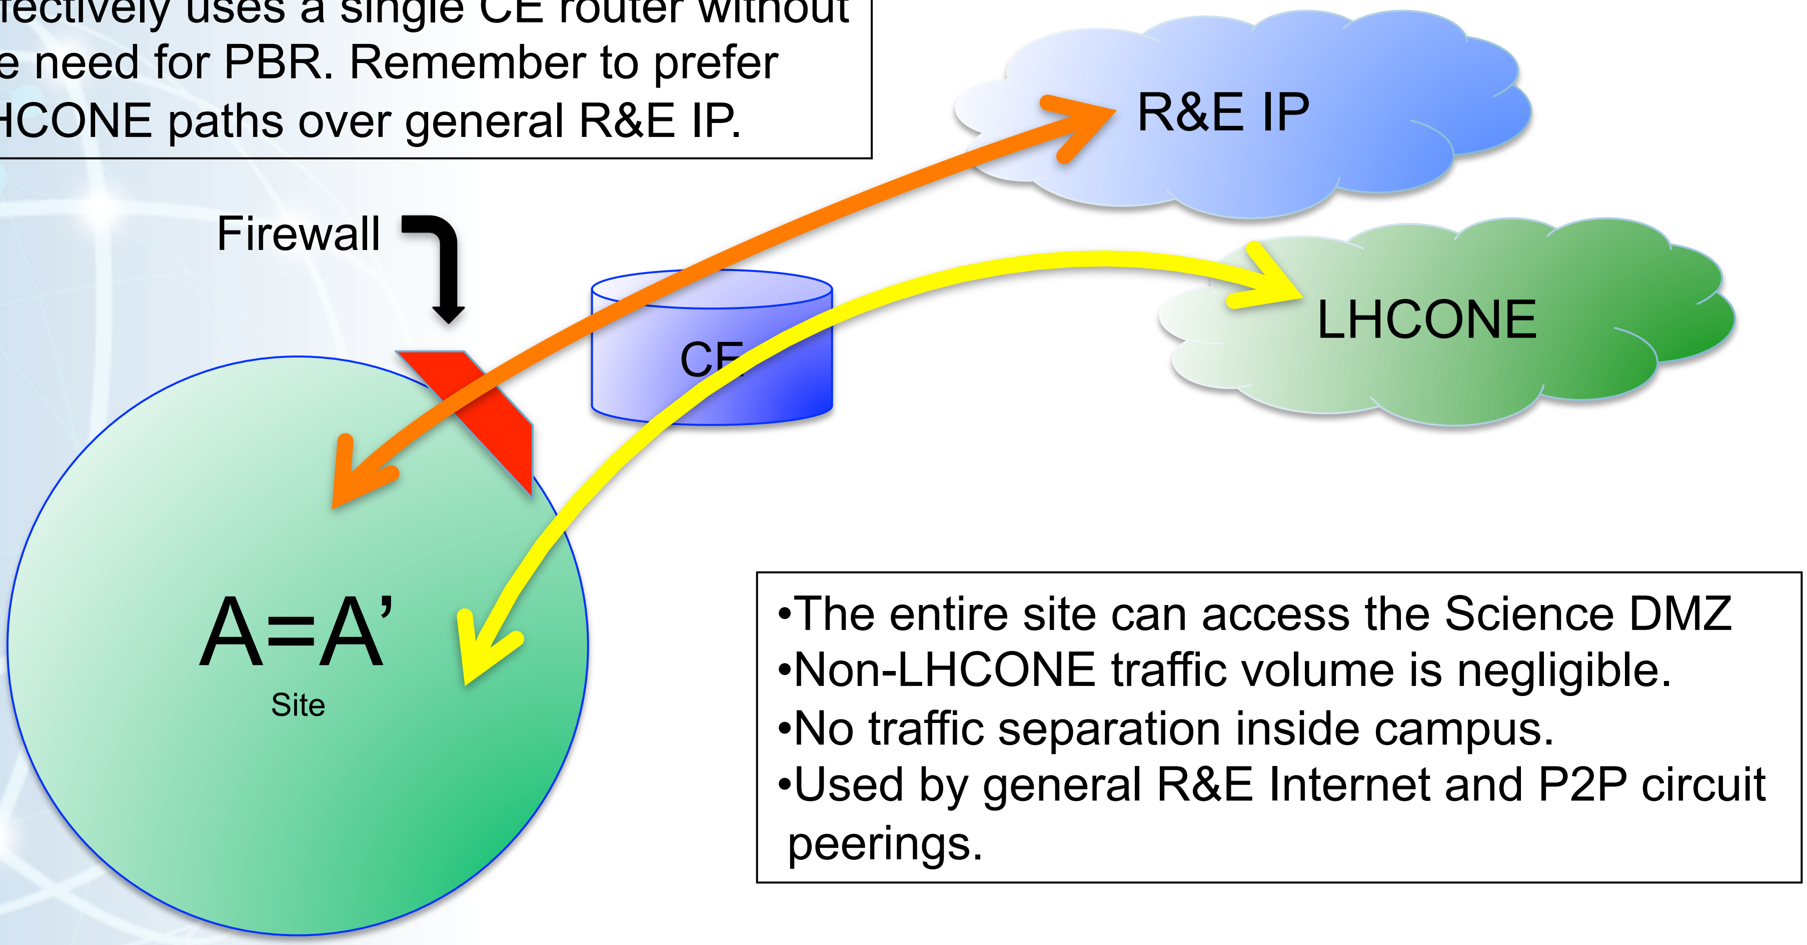

- The entire site can access the Science DMZ
- Non-LHCONE traffic volume is negligible.
- No traffic separation inside campus.
- Used by general R&E Internet and P2P circuit peerings.

R&E IP, P2P Circuits, LHCONE

Internet routing operates on the premise that if you want to have connectivity to your CIDR block you must advertise it's prefix.

P2P circuit services implemented within the LHC community generally use a BGP policy that follows the A=A' routing model.

LHCONE routing exists within the context of the general Internet and P2P services that follow different routing models.

UOC is an example of a sites that after experiencing service affecting routing problems that have no quick and easy solution, solves it by implementing a routing scheme that matches the Internet and P2P services.

The Site Connection Guidelines are not specific with regard to appropriate use and do not explicitly prohibit the A=A' routing model, should we leave it as is or more strictly enforce elimination of this routing model in LHCONE?

VRF Connection Architecture

The science DMZ requires planning and dedicated HW that many LHC affiliated enterprises have yet to either implement or even plan and budget for. UOC is in the planning and budgeting phase of their science DMZ.

A well developed science DMZ enables fine grain routing policy control necessary to enable site access control to the LHCONE.

Architectural recommendations for connecting to the LHCONE could be used as input to the planning and budgeting process of connecting sites that intend to implement tighter access controls on special purposes network connections.

Should the LHCONE site connection guidelines contain architectural recommendations?

Dual CE Routers

Maintains Traffic Separation without PBR
An Example Method

Home all A' LANs to CE2
Advertise only A' prefixes to LHCONE

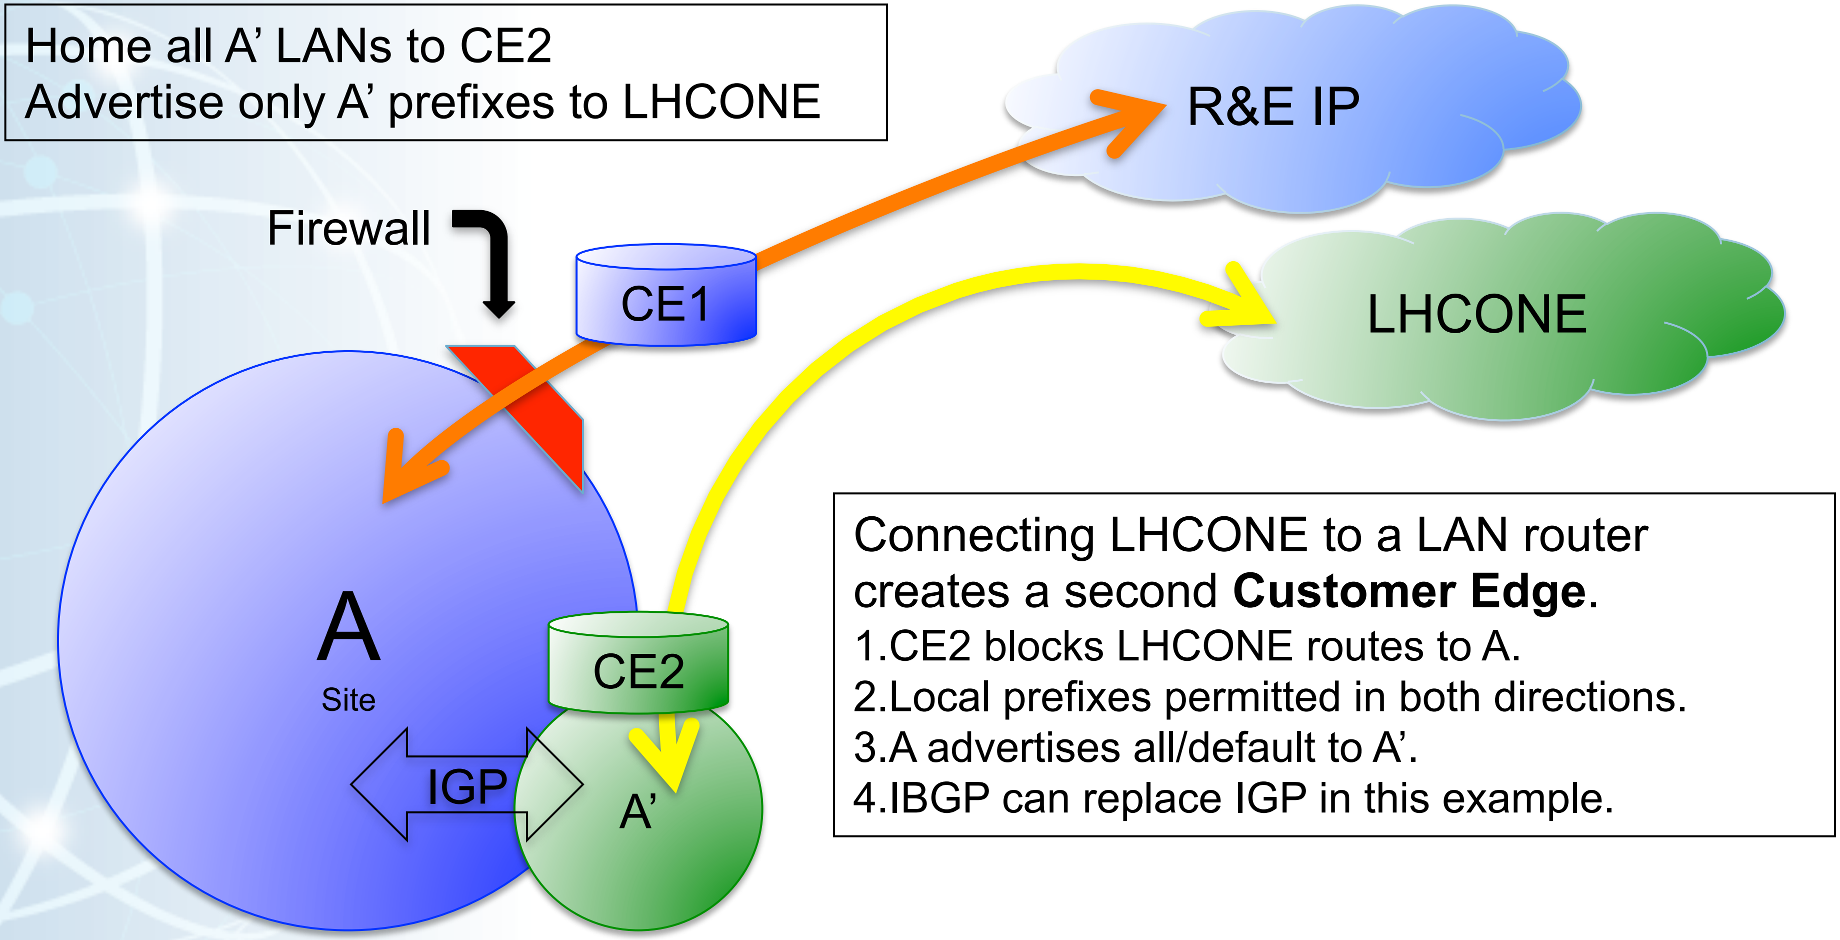

Connecting LHCONE to a LAN router creates a second **Customer Edge**.

1.CE2 blocks LHCONE routes to A.
2.Local prefixes permitted in both directions.
3.A advertises all/default to A'.
4.IBGP can replace IGP in this example.

Questions?

Michael O’Connor

ESnet Network Engineer

moc@es.net

631 344-7410