



ESnet

ENERGY SCIENCES NETWORK

LHCONE BGP Filtering

Michael O'Connor moc@es.net

ESnet

Network Engineering

LHCOPN-LHCONE meeting

Cambridge, (UK)

February 9-10, 2015



U.S. DEPARTMENT OF
ENERGY
Office of Science



BGP Communities

Overview

- The BGP communities path attribute is an optional transitive attribute of variable length.
- Prior to the creation of BGP communities, the distribution of routing information was solely based on either IP address prefixes or on the value of the AS_PATH attribute.
- BGP communities facilitate and simplify the control of routing information by grouping destination prefixes so that the routing decision can also be based on the identity of a group, significantly simplifying a BGP speaker's configuration that controls distribution of routing information. (RFC-1997)
- **External policy control** - BGP communities were devised as a way to control routing policy in an upstream provider network. This is the primary reason for considering the use of communities in the LHCONE context.
- Communities are assigned and policy is implemented in the BGP import and export policy chains on a peer by peer basis.

The Decision to Implement LHCONE BGP Filtering

Ann Arbor meeting ...

The LHCONE collaboration based on those present in Ann Arbor at the University of Michigan meeting Sept. 15-16, 2014 agreed to implement the complete list of “Mandatory” LHCONE BGP communities. The BGP community based prefix filtering capability in LHCONE will be referred to simply as “BGP Filtering”, within the LHCONE networking context.

Operations calls ...

Subsequent LHCONE operations calls were used to refine the requirements, coordinate the implementation and initiate testing.

LHCONE BGP communities are listed at:

<https://twiki.cern.ch/twiki/bin/view/LHCONE/LhcOneVRF>



LHCONE

Mandatory Operational BGP Communities

Community	Meaning
65001:ASN	Prepend Local ASN 1X to Export to Peer ASN
65002:ASN	Prepend Local ASN 2X to Export to Peer ASN
65003:ASN	Prepend Local ASN 3X to Export to Peer ASN
65010:ASN	Do not Export to Peer ASN

Due to a lack of consistency in using the term “Tier1” as a classifier, which may vary over time and between experiments. **Support was dropped for the “Do not advertise to any Tier1” community** and it’s status changed from mandatory to optional. (operations call Jan. 5, 2015) This action is generally accepted to reduce the risk of routing asymmetry that would arise when filtering by this method.

2014 Operations call. M. O’Connor would bring this to the attention of the wider collaboration at the Cambridge meeting.



LHCONE BGP Filtering Guidelines

Service Definition document V1.0

The BGP filtering service is intended to be used by an LHCONE end site to prevent the distribution of their BGP route prefixes to another LHCONE end-site.

1. An individual BGP community tag will be used for each and every remote end site that is filtered.
2. A site will tag ALL of the route prefixes it exports into LHCONE uniformly.
3. NSP ASNs are NOT valid for use in LHCONE BGP Filtering communities.
4. NSPs will only provision this service at their customer edge and will NOT provision it on internal LHCONE NSP/NSP BGP peerings.
5. NSPs only filter prefixes for their directly attached customers on export to those customers. Otherwise they pass LHCONE BGP Filtering communities along without modification.

<https://twiki.cern.ch/twiki/pub/LHCONE/LhcOneVRF/LHCONEBGPFilteringServiceDefinition.pdf>



LHCONE BGP Filtering Limitations

The LHCONE BGP filtering approach is enabled through policies provisioned by NSPs and are designed to be exercised by LHCONE collaborating compute centers.

Limitations to the LHCONE BGP filtering approach:

1. Compute centers not using BGP to connect to their NSP can not participate in a scheme relying on BGP. ie: Static routing
2. Compute center routes originated from their NSP ASN.
3. Compute centers using private ASNs in their BGP configuration.

LHCONE Networks That Are Responsible For Implementing Customer Edge BGP Filtering

North America	Europe	
<p>CANET(6509)</p> <ul style="list-style-type: none"> • BCNET(271) • UTORONTO(239) • UVIC(16462) • MCGILL(15318) • TRIUMF(36391) • UALBERTA(3359) <p>ESNET(293)</p> <ul style="list-style-type: none"> • FNAL(3152) • BNL(43) • SLAC(3671) • PNNL(3428) <p>I2(11537)</p> <ul style="list-style-type: none"> • UIUC(38) • UNL(7896) • MIT(3) • AGLT2(229) • MICH-Z(230) • UOC(160) • CSUNET(2153) • ULTRALIGHT(32361) • VANDERBILT(39590) • OKLAHOMA(25776) <p>INDIAN(19782)</p> <ul style="list-style-type: none"> • IUPUI(10680) 	<p>CERN-LHC1(20641)</p> <ul style="list-style-type: none"> • CERN-WIGNER(6133) • CERN(513) <p>DFN(680)</p> <ul style="list-style-type: none"> • KIT(34878) • DESY(1754) <p>GEANT(20965)</p> <ul style="list-style-type: none"> • ROEDUNET(2614) • ASGARR(137) • ARNES-NET(2107) • CZECH-ACAD-SCI(281) <p>LHCONE-RENATER(2091)</p> <ul style="list-style-type: none"> • IN2P3(789) • CEA-SACLAY(777) <p>NORDUNET(2603)</p> <ul style="list-style-type: none"> • NDGF(39590) <p>REDIRIS(766)</p> <ul style="list-style-type: none"> • PIC(43115) <p>SURFSARA(1162)</p> <ul style="list-style-type: none"> • NIKHEF(1104) 	<ol style="list-style-type: none"> 1. CANET(6509) 2. CERN(513) 3. CERN-LHCONE(20641) 4. DFN(680) 5. ESNET(293) 6. GEANT(20965) 7. I2(11537) 8. INDIAN(19782) 9. LHCONE-RENATER(2091) 10. NORDUNET(2603) 11. REDIRIS(766) 12. SURFSARA(1162)



BGP Filtering End Site Configuration Steps

In order to take advantage of LHCONE BGP Filtering a site needs to:

1. Tag **ALL** BGP routes **exported** to LHCONE with a consistent set of supported “BGP Filtering” community tags.
2. For every use of the “Do Not Advertise to ASN” community, a corresponding “Reject From ASN” **import** policy is required for symmetry.



The LHCONE BGP Filtering feature enables a compute center to control only the **export** of their routes. While essential to using this feature, **import** policy remains a local responsibility.


Initial Test Pair: BNL, KIT

Example: KIT blocks BNL

KIT tags their own routes with “Do not export to BNL”.

DFN and GEANT pass KIT routes with communities in tact.

ESnet rejects KIT routes on export to BNL, but exports KIT routes to other ESnet sites, FNAL, SLAC etc.



With ESnet support, BNL and KIT tested the LHCONE BGP Filtering capability using a two phase bidirectional process.



Testing Results

The general concept is well understood among participants, however the implementations across NSPs and end sites are based on differing assumptions.

- End sites might assume that their communities will be respected and preserved across network domains. However, NSP policies vary with regard to BGP community transit, some may ignore or clear them.*
- Regional NSPs may not realize that they are responsible for implementing LHCONE BGP filtering on their LHCONE customer peerings. Better documentation including configuration examples became useful.*
- NSPs are not always clear on which LHCONE BGP peerings require this configuration.*
- Both BNL and KIT are ATLAS Tier1 sites, preferring LHCOPN over LHCONE. Since these tests did not affect the “best path”, access to edge routers at both ends was required to verify correctness of the BGP neighbor routes rather than simply using traceroute as specified in the test procedure document.*

Even though the orchestra knows the music, rehearsal is critical to a successful performance.



Summary & Conclusions

- LHCONE BGP Filtering can become a reliable feature of the network service if implemented uniformly by all required NSPs.
- Testing has only begun and needs to be performed by a much broader set of participants.
- The participation and cooperation of LHCONE NSPs is extremely valuable in the testing process.
- Once uniformly implemented, this feature has potential for addressing issues in other areas.
- Service Definition V1.0:
<https://twiki.cern.ch/twiki/pub/LHCONE/LhcOneVRF/LHCONEBGPFilteringServiceDefinition.pdf>
- Site to Site Test Procedure V1.0:
<https://twiki.cern.ch/twiki/pub/LHCONE/LhcOneVRF/CommunityFilterTestPlan.pdf>

Contact: Michael O'Connor (moc@es.net)

