

LHCONE BGP Filtering - Site to Site Test Procedure V1.0

M. O'Connor ESnet 2015

Participants:

BNL - J. Bigrow bigrow@bnl.gov

CERN - E. Martelli edoardo.martelli@cern.ch

ESnet - M. O'Connor moc@es.net

KIT - B. Hoeft bruno.hoeft@kit.edu

Objective:

To verify the LHCONE BGP community based filtering functionality for participating Network Service Providers through the definition of tests procedures that can be followed by any pair of LHCONE compute centers to verify this service feature.

Overview:

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**” BGP community filters specified in the following table. The BGP community based prefix filtering capability in LHCONE will be referred to simply as “BGP Filtering”, within the LHCONE networking context. This more generic term permits some expansion to other specific types of BGP filtering if necessary and presently there are no exclusive forms of LHCONE BGP filtering that require any particular distinction.

LHCONE “BGP Filtering” is required to be implemented by participating NSPs but not required to be employed by “compute centers” or “sites”. Each site acts on it's own behalf to affect redistribution of the sites own BGP route prefixes by the remote end NSP as defined by the LHCONE collaboration. The site also takes responsibility to reject in symmetric fashion the target BGP route prefixes provided to it. For example, if ASN X employs “BGP Filtering” to prevent the remote side NSP from advertising their prefixes to ASN Y, then ASN X must also take care to reject all ASN Y prefixes it receives from it's LHCONE NSP.

BGP Filtering is only precise to the ASN and all prefixes advertised by a site to LHCONE must be marked consistently. In order to preserve symmetry, if a site intends to have the LHCONE NSPs filter their prefixes to specific customers, the chosen filtering communities will be applied to “ALL” of their LHCONE advertised prefixes uniformly, this filtering capability is not intended to omit specific prefixes from specific sites as this would require source based routing policy outside the scope of this service to implement symmetrically.

Table of supported BGP communities:

https://twiki.cern.ch/twiki/bin/view/LHCONE/LhcOneVRF#BGP_communities

Community	Type	Meaning	Notes
65001:XXXX	Operational	prepend 1x to ASxxxx	Mandatory
65002:XXXX	Operational	prepend 2x to ASxxxx	Mandatory
65003:XXXX	Operational	prepend 3x to ASxxxx	Mandatory
65010:XXXX	Operational	do not announce to ASxxxx	Mandatory

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:

- Compute centers not using BGP to connect to their NSP can not participate in a scheme relying on BGP.
- Compute centers using private ASNs in their BGP configuration.

General Procedure:

BGP community based filtering will be tested between participating compute centers. This is an evolving document and will be refined over time, it is intended to be exercised by a pair of end sites to verify the correct service implementation by each of their respective first hop LHCONE NSPs. This is the testing procedure designed for a pair of sites to test the use of “LHCONE BGP Filtering” by the designated “local site” and will test control at a single end. However both sites should be tested by designating each as the “local site” individually in a series of two procedures.

Establish pre-test routing baseline

Designate:

1. Local site and a remote site and note their ASNs.
2. Local and remote LHCONE hosts for testing.
3. Run continuous pings between hosts to indicate RTT and general reachability between sites.

Record:

1. (edge router) LHCONE BGP peer IP address(es).
2. (edge router) LHCONE BGP peering(s) import/export route prefixes
3. (edge router) LHCONE BGP route table
4. (host) Ping from local to remote test host, note round trip time (RTT).
5. (host) Traceroute from local to remote test host, verify LHCONE path.

Check:

1. Did traceroute output list one of the LHCONE BGP peer IP addresses from above?
 - a. if yes proceed.
 - b. if no, examine BGP route preferences for all paths to the remote ASN to determine why LHCONE is not the preferred path to the remote ASN.

Implement Changes

Modify LHCONE BGP export policy:

1. Tag all route prefixes exported into LHCONE with the BGP community being tested.
2. Note time of change.

Modify LHCONE BGP import policy:

1. Reject all route prefixed imported from LHCONE that have an as-path containing the target site ASN.

Verify Results

Local Verification:

1. (edge router) show BGP routes advertised to LHCONE verify BGP community tags on all routes advertised to all LHCONE peers. (possibly on multiple routers)
2. (edge router) show BGP routes received from all LHCONE peers, verify reject policy changes.
3. (host) ping from local host to remote host.
4. (host) traceroute from local host to remote host.

Remote Verification:

1. (edge router) show BGP route received from all LHCONE peers. Repeat until change is observed or a reasonable BGP propagation period has elapsed. Make a note of the time.
2. (host) ping from remote host to local host.
3. (host) traceroute from remote host back to local host.

Check:

1. Is the remote host still reachable with ping? (if not traffic between sites may be affected and you may need to back out changes.)
2. Is traceroute able to reach the remote site/host and are your edge LHCONE peer addresses still present in the output? (indicating an LHCONE path)

Evaluation**Success Criteria:**

1. The remote target site LHCONE BGP peering no longer advertises any route prefixes containing the designated local ASN to the remote ASN.
2. When after some reasonable propagation period, reachability between sites is preserved and the traceroute output in both directions no longer lists LHCONE peer addresses in the path output.