

LHCONE BGP Filtering Service Definition V1.0

M. O'Connor ESnet

[Overview](#)

[Supported BGP Communities](#)

[Algorithm](#)

[Example JUNOS:](#)

[Example Alcatel TIMOS](#)

[Compliance](#)

[Precision and Scope](#)

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.

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.

Supported BGP Communities

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.

Algorithm

Common community processing policy for LHCONE customer site BGP export.

#“Do not announce to ASxxxx”

if from community “65010:Site-ASN” where “Site-ASN” matches BGP “Peer-ASN”
Reject, do not export

“Prepending 1X to ASxxxx”

If from community “65001:Site-ASN” where “Site-ASN” matches BGP “Peer-ASN”
Prepend AS-PATH LOCAL-NSP-ASN 1X
Accept and export

“Prepending 2X to ASxxxx”

If from community “65002:Site-ASN” where “Site-ASN” matches BGP “Peer-ASN”
Prepend AS-PATH LOCAL-NSP-ASN 2X
Accept and export

“Prepending 3X to ASxxxx”

If from community “65003:Site-ASN” where “Site-ASN” matches BGP “Peer-ASN”
Prepend AS-PATH LOCAL-NSP-ASN 3X
Accept and export

Default Action: Next Policy

Example JUNOS:

(Site specific example using ESnet NSP ASN 293 and END-SITE ABCD ASN 100)

```
community lhcone-noexport-abcd members 65010:100;
community lhcone-prepend-1x-abcd members 65001:100;
community lhcone-prepend-2x-abcd members 65002:100;
community lhcone-prepend-3x-abcd members 65003:100;

policy-statement lhcone-abcd {
    term noexport {
        from community lhcone-noexport-abcd;
        then reject;
    }
    term prepend-1x {
        from community lhcone-prepend-1x-abcd;
        then as-path-prepend 293;
    }
    term prepend-2x {
        from community lhcone-prepend-2x-abcd;
        then as-path-prepend "293 293";
    }
    term prepend-3x {
        from community lhcone-prepend-3x-abcd;
        then as-path-prepend "293 293 293";
    }
    then next policy;
}
routing-instances {
    lhcone {
        description "LHCONE L3VPN";
        instance-type vrf;
        protocols {
            bgp {
                group LHCONE-Sites {
                    local-as 293;
                    neighbor 198.124.90.250 {
                        description "ABCD LHCONE";
                        local-address 198.124.90.249;
                        import [ lhcone-accept-abcd-routes reject-all ];
                        export [ lhcone-abcd accept-all ];
                        peer-as 100;
                    }
                }
            }
        }
    }
}
}
```

Example Alcatel TIMOS

```
community "lhcone-noexport-abcd" members "65010:100"
community "lhcone-prepend-1x-abcd" members "65001:100"
community "lhcone-prepend-2x-abcd" members "65002:100"
community "lhcone-prepend-3x-abcd" members "65003:100"

policy-options
begin
  policy-statement "lhcone-abcd"
  entry 20
    from
      community "lhcone-noexport-abcd"
    exit
    action reject
  exit
  entry 30
    from
      community "lhcone-prepend-1x-abcd"
    exit
    action next-policy
      as-path-prepend 293 1
    exit
  exit
  entry 40
    from
      community "lhcone-prepend-2x-abcd"
    exit
    action next-policy
      as-path-prepend 293 2
    exit
  exit
  entry 50
    from
      community "lhcone-prepend-3x-abcd"
    exit
    action next-policy
      as-path-prepend 293 3
    exit
  exit
  default-action next-policy
  exit
exit
exit
```

```
# BGP
neighbor 198.124.80.6
  description "ABCD"
  local-address 198.124.90.249
  import "lhcone-abcd-routes" "reject-all"
  export "lhcone-abcd" "accept-all"
  local-as 293
  peer-as 100
exit
```

Compliance

LHCONE “BGP Filtering” is required to be implemented by participating NSPs but not required to be used by “compute centers” or “sites”. Each site acts on it’s own behalf to affect redistribution of their BGP route prefixes by the remote end NSP as defined by the LHCONE collaboration.

In using the LHCONE BGP Filtering feature, a site also takes responsibility to reject in symmetric fashion the remote site BGP route prefixes provided to it by their LHCONE NSP. For example, if ASN X employs “BGP Filtering” to prevent the remote LHCONE 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.

Precision and Scope

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.

General Testing Procedure:

BGP community based filter testing between participating compute centers 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, in other words test from one site at a time.

Testing in this fashion is intended to only Impact the LHCONE path between the testing partners. If LHCONE is not the best path between testing sites then these tests can be run in a non-service affecting manner, assuming the primary path remains up throughout the tests. Non service affecting testing is not described in the following procedure.

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.

