LHCONE Site Provisioning Guidelines

V1.2 - 9/27/16

The LHCONE is a collaborative network implemented for the LHC research community as a controlled access “walled garden” routed IP internet.

In the global Internet, asymmetric routing is an undesirable but generally acceptable condition. However, at the customer edge, a state-full firewall will only permit connectivity if it is able to observe a flow on both the ingress and egress directions, preserving routing symmetry at the firewall. Asymmetric routing through a state-full firewall will fail in many common connection scenarios.

The majority of the LHCONE participating high performance compute centers have implemented some form of science DMZ perimeter architecture in order to attach to the LHCONE network. These sites rely on this type of high performance access controlled path as an alternative to their conventional perimeter access chokepoint. However, if these multiple paths through their secure site perimeters are not routed precisely, the potential for creating routing asymmetry increases dramatically.

Connection guidelines to ensure route symmetry at the DMZ

- Select the local LAN address ranges that are able to participate in LHCONE based on the requirements specified in the LHCONE Appropriate Use Policy[1]. Work with your NSP to advertise these address range prefixes into LHCONE.

- 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, failure to do so will result in asymmetric routing for unauthorized address ranges.

- 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.

[1] https://twiki.cern.ch/twiki/bin/view/LHCONE/LhcOneAup

Site Policy and Filtering Requirements

Site egress packets should be filtered based on the BGP address ranges advertised to LHCONE. Filter drop counters can be used locally to verify that unauthorized site sources are not egressing onto LHCONE. Site egress packet filtering is highly recommended for adding excellent operational perspective.

Inbound packet filtering should use a dynamic access control method such as Reverse Path Forwarding (RPF). This will allow only packets from routed LHCONE sources to enter an LHCONE connected interface. Dynamic methods are recommended over static filters, which will inevitably cause periods of instability as the LHCONE VRF routing table is updated. Inbound packet filtering is not a requirement and left to the sites discretion.

NSP Policy and Filtering Requirements

Objective: The LHCONE access control policy is based on the set of accepted site prefixes, the routing table. In order to maintain route symmetry and access control each NSP will implement policy to manage their connected customer address prefix ranges.

- Prefix Lists – negotiated between LHCONE connecting institutions and the NSP. NSP's will implement route prefix filter policy control all of the advertised ranges from their connecting sites. A connecting site must not have the ability to add a prefix to the LHCONE routing table without direct cooperation from their LHCONE NSP.

- Packet filtering – NSPs will implement packet filters on their customer edge that are based on customer address prefixes ranges. All accepted packets will be sourced from address ranges present in the LHCONE routing table. In that way an LHCONE site's security organization can be confident that only address ranges present in the LHCONE routing table will have the ability to transmit into their site through the LHCONE network.

- NSPs are strongly encouraged to implement maximum prefix limits and default route filters between LHCONE NSP peerings.

By monitoring filter counters NSPs will have a clear indication of routing problems at the customer edge.

A wide variety of routing implementations are available for sites to attach to LHCONE, while also ensuring adherence to the LHCONE AUP. Many of these approaches involve Policy Based Routing (PBR) applied to source addresses. PBR is generally considered non-standard since PBR features and capabilities vary widely between router manufacturers. As an alternative to PBR source routing, the following destination routing examples are offered for consideration.

Example 1: Destination Routing for LHCONE Site Connections

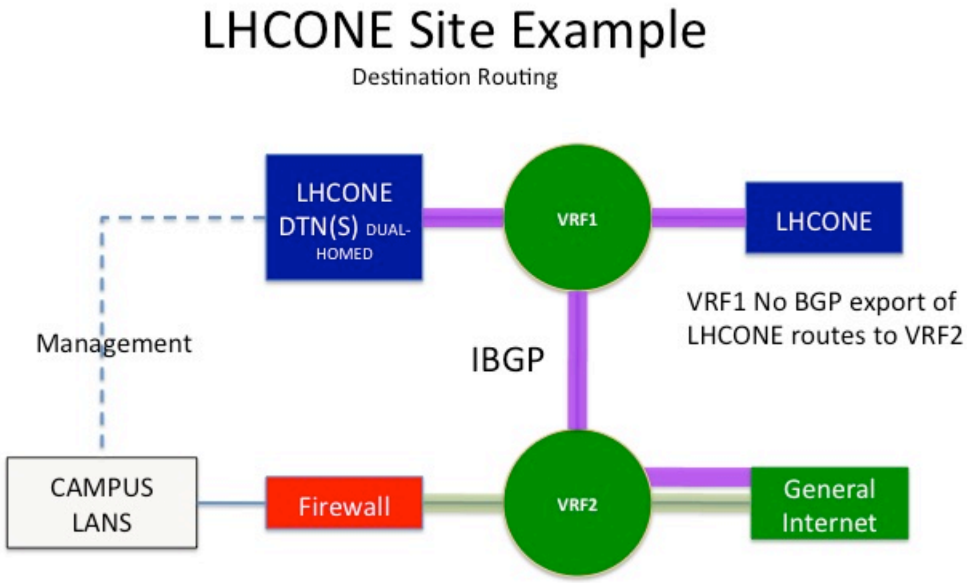

Example 2: Destination Routing for Multiple Overlay Connections

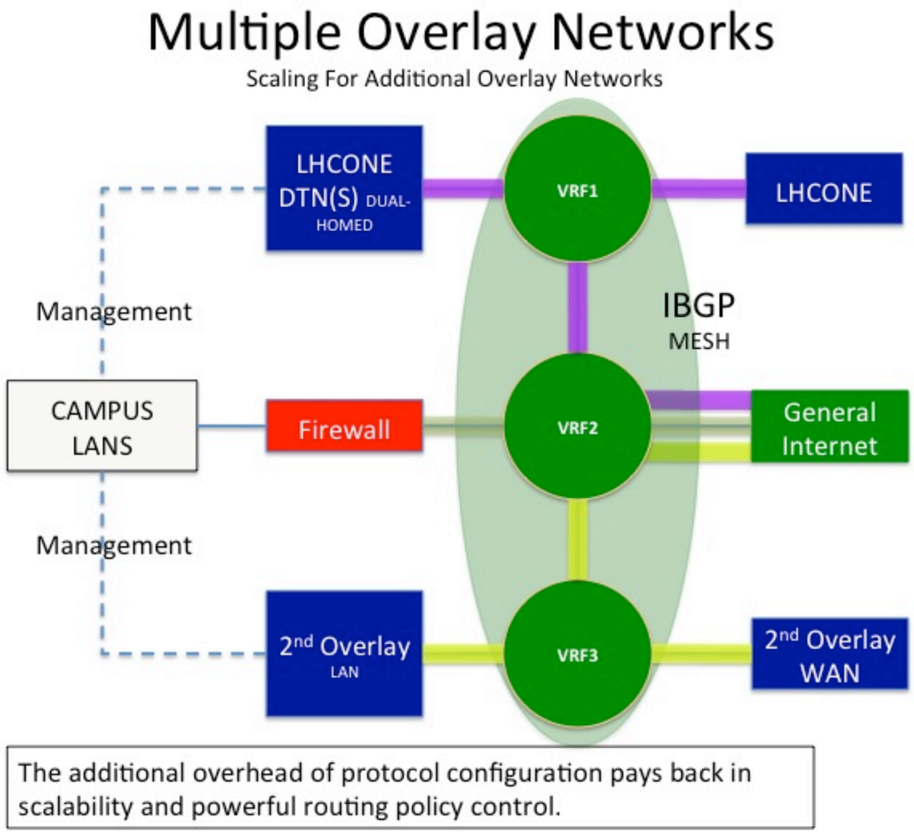

Multiple Overlay Networks

Scaling For Additional Overlay Networks

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