L3VPN Service Specification

This file contains a high-level overview of the L3VPN service components and mechanics. It also contains the set of technical parameters for the service.

1 L3VPN service overview

The L3VPN service follows the peer-to-peer model. The HEAnet network is presented to the client sites as a large, dedicated, and distributed virtual router. Those are the main attributes of the L3PVN service: scalability, segregation and IP routing between the client and provider.

There is no upper end on the number of sites that can take part of a L3VPN. Because the service is based on IP routing between the client and the provider, and the VPN prefixes are distributed by means of BGP, growth is unlimited.

When part of a VPN, traffic between sites is segregated from the rest of the traffic in the HEAnet network using tags. Traffic is not encrypted, it is separated (a possible analogy is the VLAN ID in a Layer 2 network, where traffic in each VLAN is separated from the rest by means of their VLAN ID).

Finally, there is IP routing interaction between the client and the provider (HEAnet in this case), except that the routing instance in the provider side is a dedicated Virtual Routing and Forwarding (VRF) table, which is again segregated from the main and public routing instance (continuing with the similarities, same concept of chroot in UNIX systems).

The glue, the transport, between all the sites are Multi-Protocol Label Switching (MPLS) tunnels across the HEAnet network.

The way all the above components are put together is:

- Each peer ("client") network connects to a Layer 3 VPN using a dedicated interface.
- These interfaces are associated on the HEAnet ("provider") routers with a specific Virtual Routing and Forwarding (VRF) table.
- The client router and the HEAnet router do exchange IP prefixes by means of IP routing. These prefixes are restricted to the VRF.
- Multi-Protocol BGP (MP-BGP) is used between the HEAnet routers to exchange prefixes and to ensure that only this VRF is populated with routes advertised via the associated client interfaces.
- MPLS tunnels are used across the HEAnet network to keep the traffic contained to that specific Layer 3 VPN.

For detailed technical information on the service mechanics, the L3PVN is no other than the BGP/MPLS VPN described in rfc4364.

2 L3VPN service specification

2.1 Interfaces, UNI

- Service is delivered over the same physical interface than the IP service.
- If the L3VPN requires a new port, standard (one-off) price will be charged.
- Service is delivered over a sub-interface, 802.1q.
- VLAN ID for the service are in the range 100 to 200.
- There is no other feature on the UNI except VLAN ID and IP addressing.
- There is no rate limit in the UNI.
- There is no firewall filters/ACL in the PE end of the UNI.

2.2 UNI, CE-PE IP addressing

- Allocated by HEAnet. Not an option to use any other addressing.
- Managed by client.
- IP addressing is from the HEAnet ranges 87.44.68.0/22 and 2001:0770:0100::/48. IPv4 addresses will be /30, with the lowest IP on the HEAnet side, and the highest on the client side. IPv6 addresses will be /64 with xxx::1/64 on the HEAnet side and xxx::2/64 on the client side.

2.3 Address families supported

- Supported
    - IPv4 unicast
    - IPv6 unicast
- Not supported
    - IPv4 nor IPv6 multicast.
    - Any other address family not explicitly supported.

It is necessary, not optional, that if IPv6 (same applies to IPv4) is enabled, it is enabled in all the UNI CE endpoints. It is not possible to have IPv6 in a section of the L3VPN and not in another. Monitoring and SLA measure and report any to any connectivity, hence they will fail if some parts of the L3VPN are not enabled for a particular address family.

2.4 Routing in the CE-PE segment

- Supported
    - Connected. Secondary IP addresses are not supported. The change of IP addressing requires a service rebuild (decommission and commission).
    - BGP.
- Not supported
    - Static
    - OSPF
    - RIP
    - Any other routing protocol not explicitly supported

2.5 Routing policy within the VPN

- HEAnet does not add, modify nor delete any BGP attributes received from the client.
- Client is free to use the BGP attributes to determine the routing policy within the VPN (traffic sinks, generation and treatment of default route(s) and resilience within the VPN, etc.).

2.6 CE AS

- The BGP AS for the CE is allocated by the client. If the site is expected to have IP transit from HEAnet at any time in the future, please consult AS allocation with HEAnet first.
- Each site must use unique AS number in each site.

Notes:

- The change of BGP AS on a given CE requires a rebuild (decommission and commission) of the service in that site.

2.7 Number of prefixes

The number of dynamic routes that a customer can send into the HEAnet network per location, per VPN, is limited to 30 (24). The maximum number of routes for any single Virtual Routing and Forwarding (VRF) table is 1000 (600).

2.8 Topologies

Supported

- Any to any. All the client sites have visibility of all the other client sites.

Not supported

- Hub and spoke. Can have more than one hub, but only one route-target.

2.9 Demarcation point

In each site, the sub-interface of where the service is presented. See section 2.1 Interfaces, UNI.

2.10 Monitoring

The L3VPN service is monitored for faults by HEAnet's fault monitoring systems (Argus/Icinga2). It monitors service components and end to end traffic/data plane. All checks are done on the HEAnet router end of the UNI, on or from the PE.

Components:

- UNI administrative status must be up
- UNI operations status must be up
- BGP session PE-CE must be Established

Data plane:

- From two different PE, each of those PE will ping all the CE UNI addresses. If any of the pings fails (from any of the two PE), the service is declared down.

SLA compliance/not is calculated on the success/failure of the data plane checks.

It is necessary, not optional, that the CE does allow the reception of ICMP echo request and send of ICMP echo reply from/to the UNI IP ranges (see Section 2.2). The client routers need to have, if they have them, their control plane access lists allowing ICMP with the UNI IP ranges.

2.11 SLA

Please refer to the SLA document for the details. This section only addresses the technical implications of what is being measured.

The service level agreement for the service is measured on the availability per month for connectivity to all the demarcation points. This means that all the UNI CE IP addresses must respond to ICMP from the PE issuing the ICMP echo requests. Again, this illustrates the necessity of having ACL allowing such in CE. Equally, the measure of service uptime must be understood to be the maximum common success of all the PE-CE ICMP. The service uptime is the composite intersection of all the individual successes in the PE-CE ICMP.

3 Conventions (HEAnet internal)

- Interface naming: follow convention in https://wiki.heanet.ie/RMAN#Interface_Naming_and_SHIBA
- Service name: follow convention in https://wiki.heanet.ie/RMAN#CSD_service_.22Name_Field.22_and_Interface_Description_Format
- Route-target: auto allocated. Unless the project requires and allocates a specific one.
- Route distinguisher: auto allocated. Unless the project requires and allocates a specific one.