
NVGRE and VXLAN Encapsulation for L3VPN Extension

draft-yong-l3vpn-nvgre-vxlan-encap-00

Lucy Yong Xiaohu Xu

March 2013 Orlando FL

Problem Statement

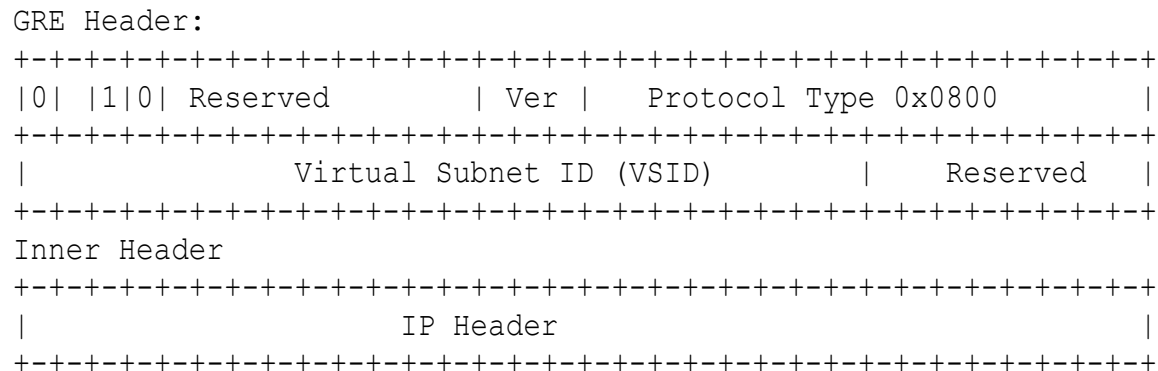
- Both NVGRE (draft-sridharan-virtualization-nvgre-01) and VXLAN (draft-mahalingam-dutt-dcops-vxlan-02) are originally specified for layer 2 vitalization overlay data encapsulation
- Several drafts propose to extend L3VPN into Data Center as layer 3 virtualization overlay
 - draft-fang-l3vpn-end-system-requirements-01
 - draft-fang-l3vpn-virtual-pe-framework-01
 - draft-drao-bgp-l3vpn-virtual-network-overlays-00
 - draft-fang-l3vpn-virtual-pe-01
 - draft-rfernando-l3vpn-service-chaining-01
- The above drafts also suggest using existing NVGRE or VXLAN data formats for layer 3 virtualization overlay beside MPLS based solution
 - carry an unnecessary inner Ethernet header per packet
 - an L3 overlay edge may interwork with an L2 overlay edge directly, where it prunes an error easily

About this draft

- Propose NVGRE enhancement for L3 virtualization overlay data encapsulation
- Propose VXLAN enhancement for L3 virtualization overlay data encapsulation

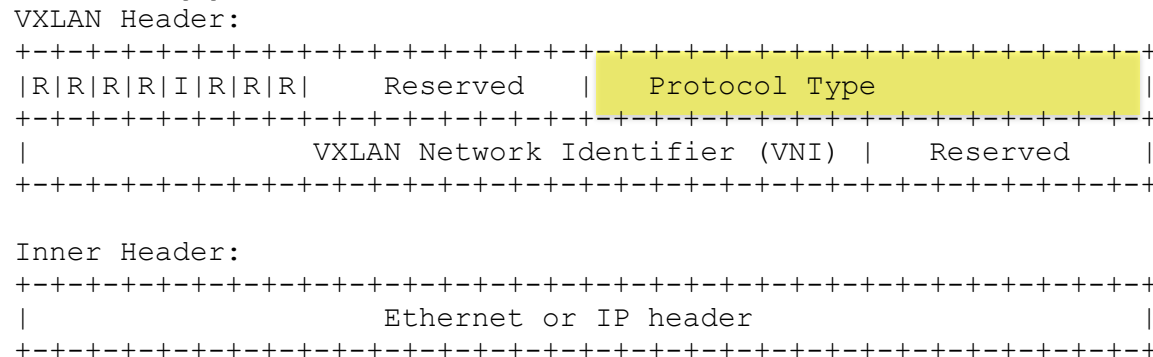
NVGRE Enhancement

- Proposed 0x0800 as the protocol type for IP payload
 - 0x6558 is the protocol type for Ethernet payload [NVGRE]
 - the version field in inner IP header can further differentiate IPv4 or IPv6 payload
- No change to other fields in NVGRE header
 - The usage of other field remains the same too
- No change to outer header
- Only IP header in the inner header if the protocol type value is 0x0800



VXLAN Enhancement

- 16 reserved bits in VXLAN header as protocol type field
 - Specify 0x0800 for IP payload and 0x6558 as Ethernet payload
 - To support backward compatibility, value 0x0000 is also for Ethernet payload
- No change to other fields in VXLAN header
 - The usage of other fields remains the same too
- No change to outer header
- Inner header may be Ethernet or IP depending on the value in protocol type



Benefit of the Enhancements

- Maintain L3VPN implementation natively and decouple it completely from L2 overlay implementation
- Interwork with existing L3VPN customers
- Seamless support L2 and L3 overlay interworking
- BGP control plane works consistently with the data plane in term of multi-protocols support
 - the inner header on a data packet matches the address family being advertised by BGP
- Save 14 bytes in each packet for a native L3 overlay
 - lower the probability of the packet fragmentation
 - eliminate the unnecessary packet process

Next Step

- Welcome comment and feedback on this
- Update the draft