

draft-dt-nvo3-encap-00

Encap Design team

# Design Team Goals, as communicated by WG Chairs

- Take one of the proposed encapsulations and enhance it to:
  - Address technical concerns
  - Allow future extensibility
  - Avoid a design that is burdensome on hardware implementations
  - Operate well with ICMP and in ECMP environments.
  - Extensibility should not require the consent of an entity outside of the IETF

# Current encapsulations

- Extensibility:
  - Geneve and GUE were defined with built-in extensibility, while VXLAN-GPE is not inherently extensible. Note that any of the three encapsulations can be extended using shim header (such as Network Service Header (NSH))
- Extension method:
  - Geneve is extensible using Type/Length/Value (TLV) that allows up to 64 TLVs, while GUE uses a bit-field based extensions with 16 bits for possible extensions.
- Length field:
  - Geneve and GUE include a Length field, indicating the length of the encapsulation header, while VXLAN-GPE does not include such a field.

# Common Encapsulation Considerations

- Useful Extensions Use cases
- Hardware Considerations
- Extension Size/ordering
- TLV vs Bit Fields
- Control plane considerations
- Larger VNI considerations

# Useful Extensions Use cases

- Telemetry extensions
  - In-band OAM
  - Tasks like debugging, troubleshooting, network planning/optimization, policy or SLA compliance checks
- Security/Integrity extensions
  - To authenticate the header and the outer IP addresses, thereby preventing attackers from injecting packets with spoofed VNIs
  - Payload security
- Group Base Policy
  - Carry source group information within a NVO3 header extension, such as access control and QoS to be applied between abstract groups rather than coupled to specific endpoint addresses

# Hardware Considerations

- NIC offloads
  - Support for efficient NIC offloads using length in the Encap header
- Switching ASIC
  - Hardware can support a subset of extensions as compared to software
  - Allows future hardware enhancements to support larger set
- Transit devices
  - Participates in the NVO3 data plane e.g. for OAM purposes
  - Able to skip over the Encap header

# Extension Size & Ordering

- Extension Size
  - Limit on maximum extension size
  - Does not restrict software flexibility
  - Control plane can further limit the size
- Extension Ordering
  - Control plane can restrict TLV ordering

# TLV vs Bit Fields

- TLVs are a lot more flexible
  - TLVs are self describing
  - Allows multiple vendor specific extensions
- Ordered TLVs are the same thing as bit fields in known order
  - Use TLV with restrictions on the size and alignment
  - Support control plane to enforce order of the TLVs

# Control plane considerations

- Control plane means both protocols such as EVPN and others, and also deployment specific configuration.
- Express care about particular extensions and the desired order of those extensions
- Different NVEs could desire different (sets of) extensions, which means that the source NVE needs to be able to place different sets of extensions in different NVO3 packets, and perhaps in different order

# Larger VNI considerations

- 24-bit VNI would avoid unnecessary changes to existing implementations
- Extensions are the right mechanism to support larger sized VNI

# Design team recommendations

- Geneve is most suitable as a starting point for proposed standard for network virtualization, for the following reasons:
  - TLVs provide necessary flexibility to meet diverse set of requirements
  - Total options length in header allow skipping over the options (for NIC offload operations and flow information in inner packet)
  - VNI is critical information for network virtualization and useful to be present in the header