
GRE-in-UDP Tunnels Discussion

draft-yong-tsvwg-gre-in-udp-4-encap-02

Edward Crabbe

Lucy Yong

Xiaohu Xu

Nov. 2013 Vancouver BC

GRE in Network Environments Today

- Increasingly common tunneling protocol
- Increasingly high traffic volumes
- Many devices unable to use GRE key field for input to LB hashing functions and
 - use GRE key field for input to LB hashing may cause a problem for some tunneled applications because of the key field usage
- High volume GRE source or GRE aggregator NEs problematic as a result

Goals

- GRE encapsulated traffic takes the ECMP capability provided in underlying network
- Make use of lowest common denominator field available as input across broad range of LB hash functions
- Provide a general purpose entropy-shim for use in a variety of environments
- minimize packet overhead
- conserve UDP ports
- Preserve GRE Key Field for other uses
 - NVGRE uses key field for VN ID
 - Other may use key field for payload encryption

GRE-in-UDP Encapsulation

IPv4 Header:

```
+-----+
|Version|  IHL  |Type of Service|          Total Length  |
+-----+-----+-----+-----+-----+-----+-----+
|          Identification          |Flags|    Fragment Offset  |
+-----+-----+-----+-----+-----+-----+
|  Time to Live | Protocol=17[UDP] |          Header Checksum  |
+-----+-----+-----+-----+-----+-----+
|          Source IPv4 Address          |
+-----+-----+-----+-----+-----+-----+
|          Destination IPv4 Address          |
+-----+-----+-----+-----+-----+-----+
```

UDP Header:

```
+-----+-----+-----+-----+-----+-----+-----+
|          Source Port = XXXX          |          Dest Port = TBD          |
+-----+-----+-----+-----+-----+-----+-----+
|          UDP Length          |          UDP Checksum          |
+-----+-----+-----+-----+-----+-----+-----+
```

GRE Header:

```
+-----+-----+-----+-----+-----+-----+-----+
|C| |K|S| Reserved0          | Ver |          Protocol Type          |
+-----+-----+-----+-----+-----+-----+-----+
|          Checksum (optional)          |          Reserved1 (Optional)          |
+-----+-----+-----+-----+-----+-----+-----+
|          Key (optional)          |
+-----+-----+-----+-----+-----+-----+-----+
|          Sequence Number (Optional)          |
+-----+-----+-----+-----+-----+-----+-----+
```

Advantages

- Works in most deployments
 - Almost all deployed NE's can use UDP src/dest as input to hash function
 - does not require support of additional excaps or signaling protocol in environments where they are otherwise not used
- Adds 16 bits of flow entropy
- Retains flexibility in use of GRE Key Field
- Minimal packet overhead
- Preserves UDP ports

Open Discussions

- Do we need two network virtualization overlay data encapsulation methods?
 - This draft and draft-yong-l3vpn-nvgre-vxlan-encap make enhanced NVGRE and VXLAN encapsulations very similar in the format, the difference between two:
 - Use different standard UDP port number
 - Use different bit (3 or 5) to indicate overlay header existence
- Should IETF standardizes one or both?
 - ✓ One: no need interworking or supporting both
 - ✓ Both: used in industry already, if two are very similar, hardware supports both at no cost