

# L3DL

## Layer 3 Discovery & Liveness

### draft-ietf-lsvr-l3dl-03.txt

LSVR WG

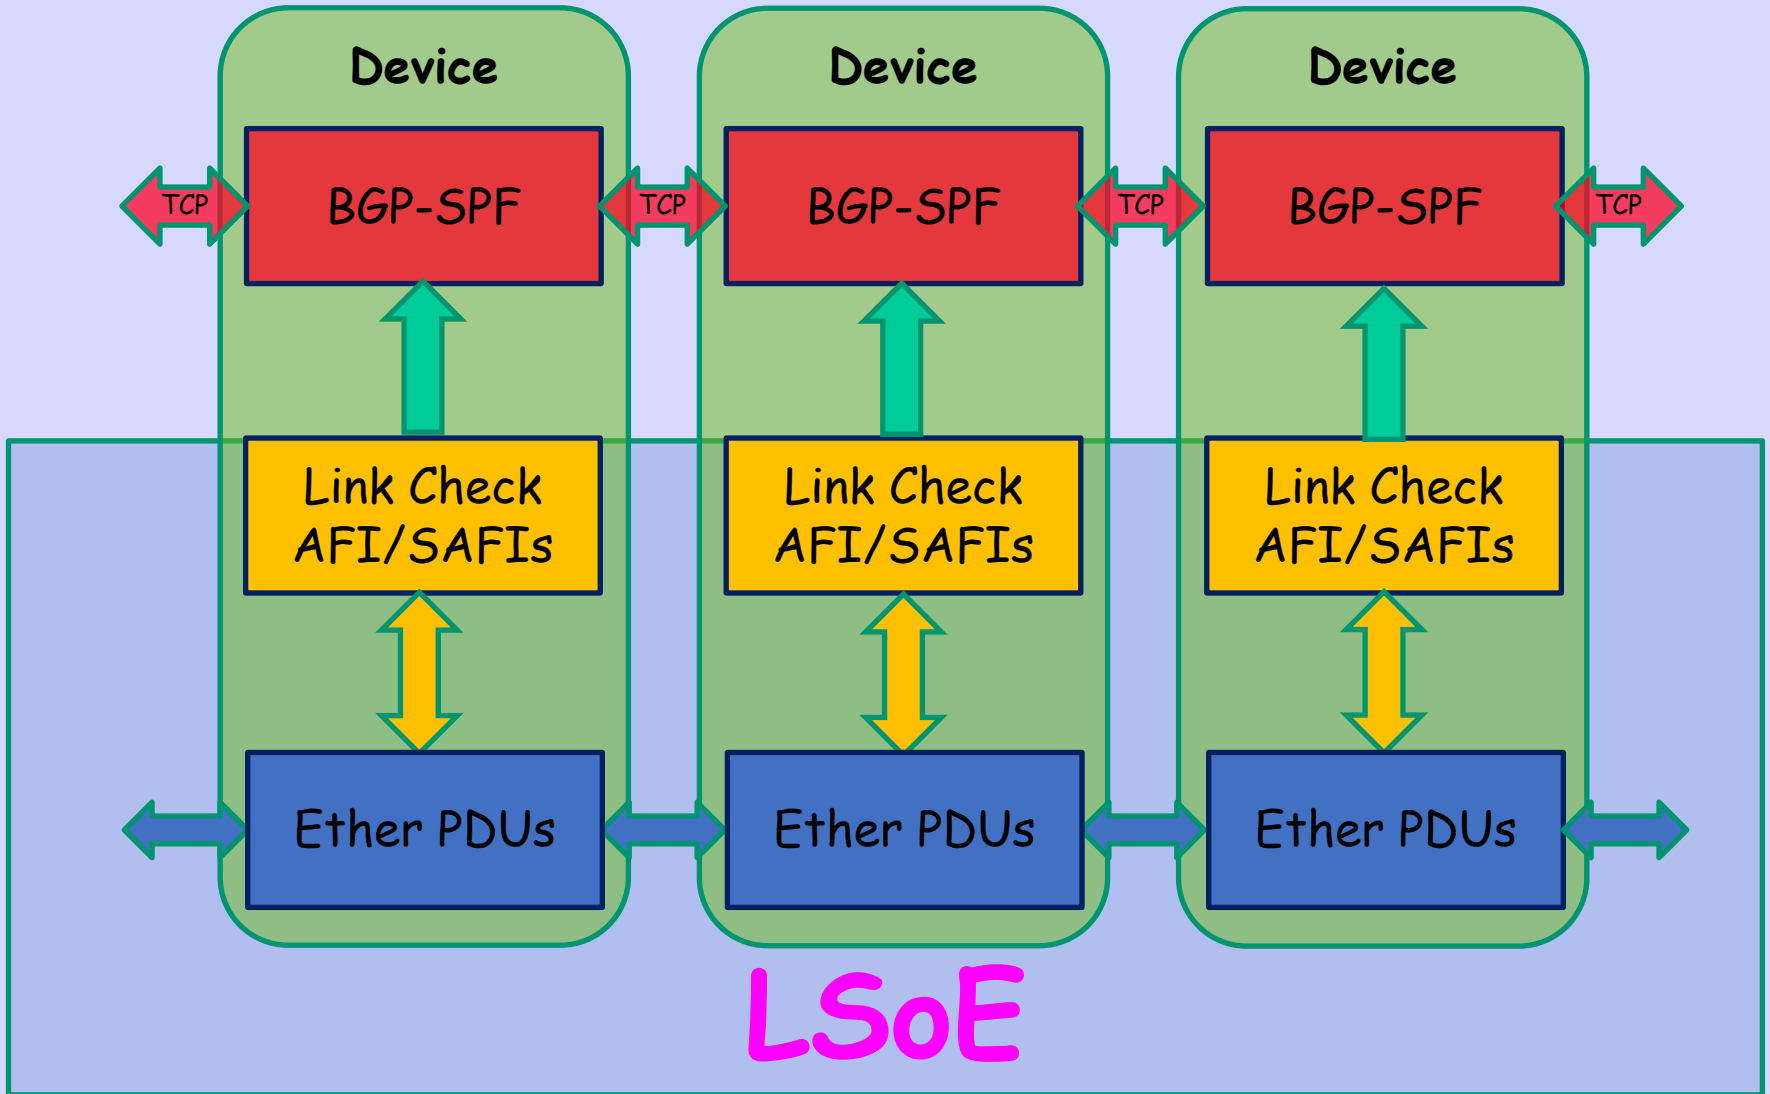
2019.07.25

randy@psg.com, sra@hactrn.net, keyur@arrcus.com

# Remember the Goal

Layer 3 Topology  
Discovery and Liveless  
for LSVR / BGP-SPF

# Just a Reminder



This is NOT a  
Routing Protocol

Discovers the  
Layer 3 Addresses  
on a PointToPoint Link

You Read the Draft



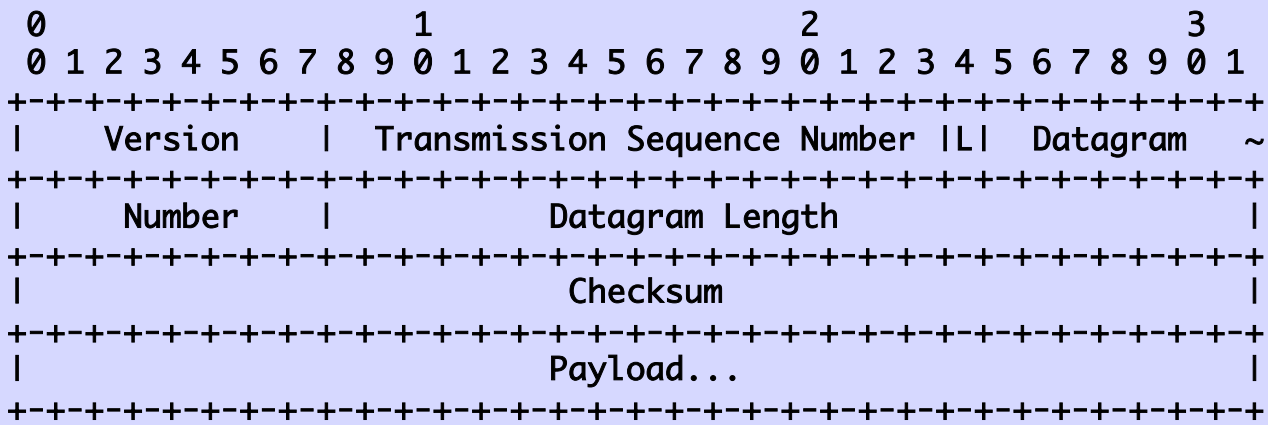
So What's New?

# Devil is in the Details



Transport Layer can  
handle  $2^{32}$  octet PDUs,  
and is now Clearly  
Delineated

$2^{24}$  Datagrams (in Ethernet Frames)  
 $2^{16}$  Octets/Datagram (except it is a Frame)





# Why not TCP?

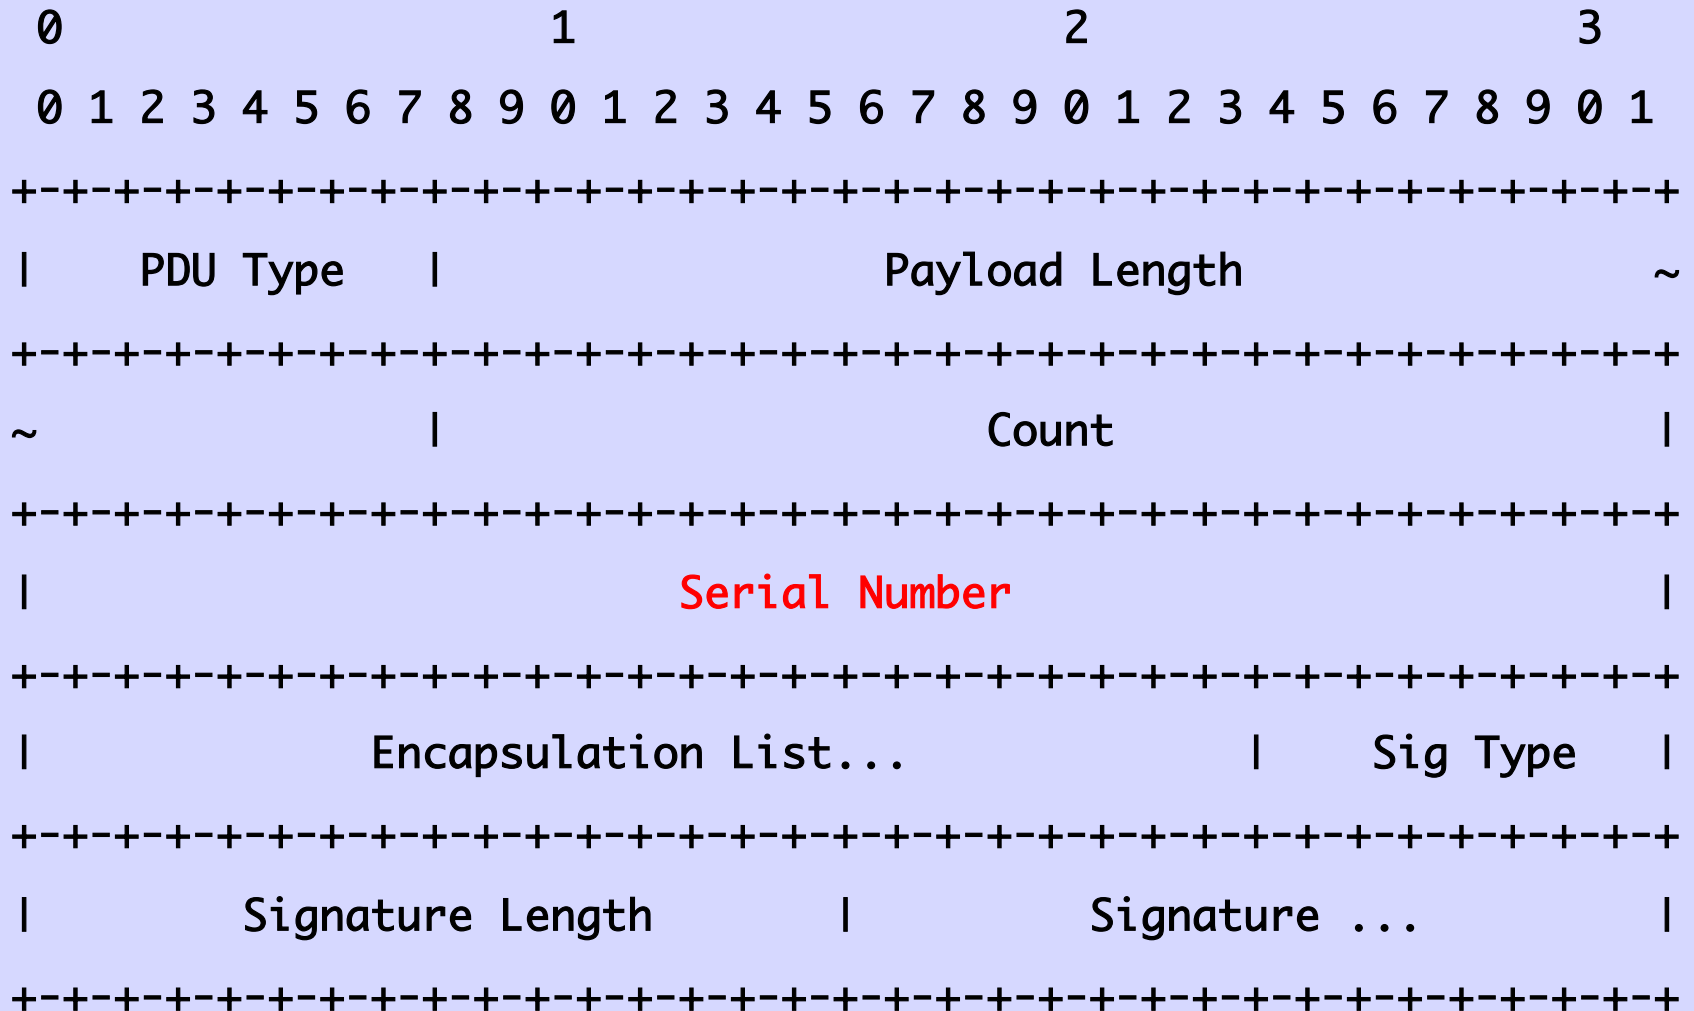
- When this runs, there are no IP Addresses
- This protocol is to Learn IP Addresses
- So it is a cheap TCP-like protocol
- Reassembly of out of order Datagrams
- Retransmission with Back-off
- PDUs are ACKnowledged
- ...

# Fully Stateful Session Per Peer

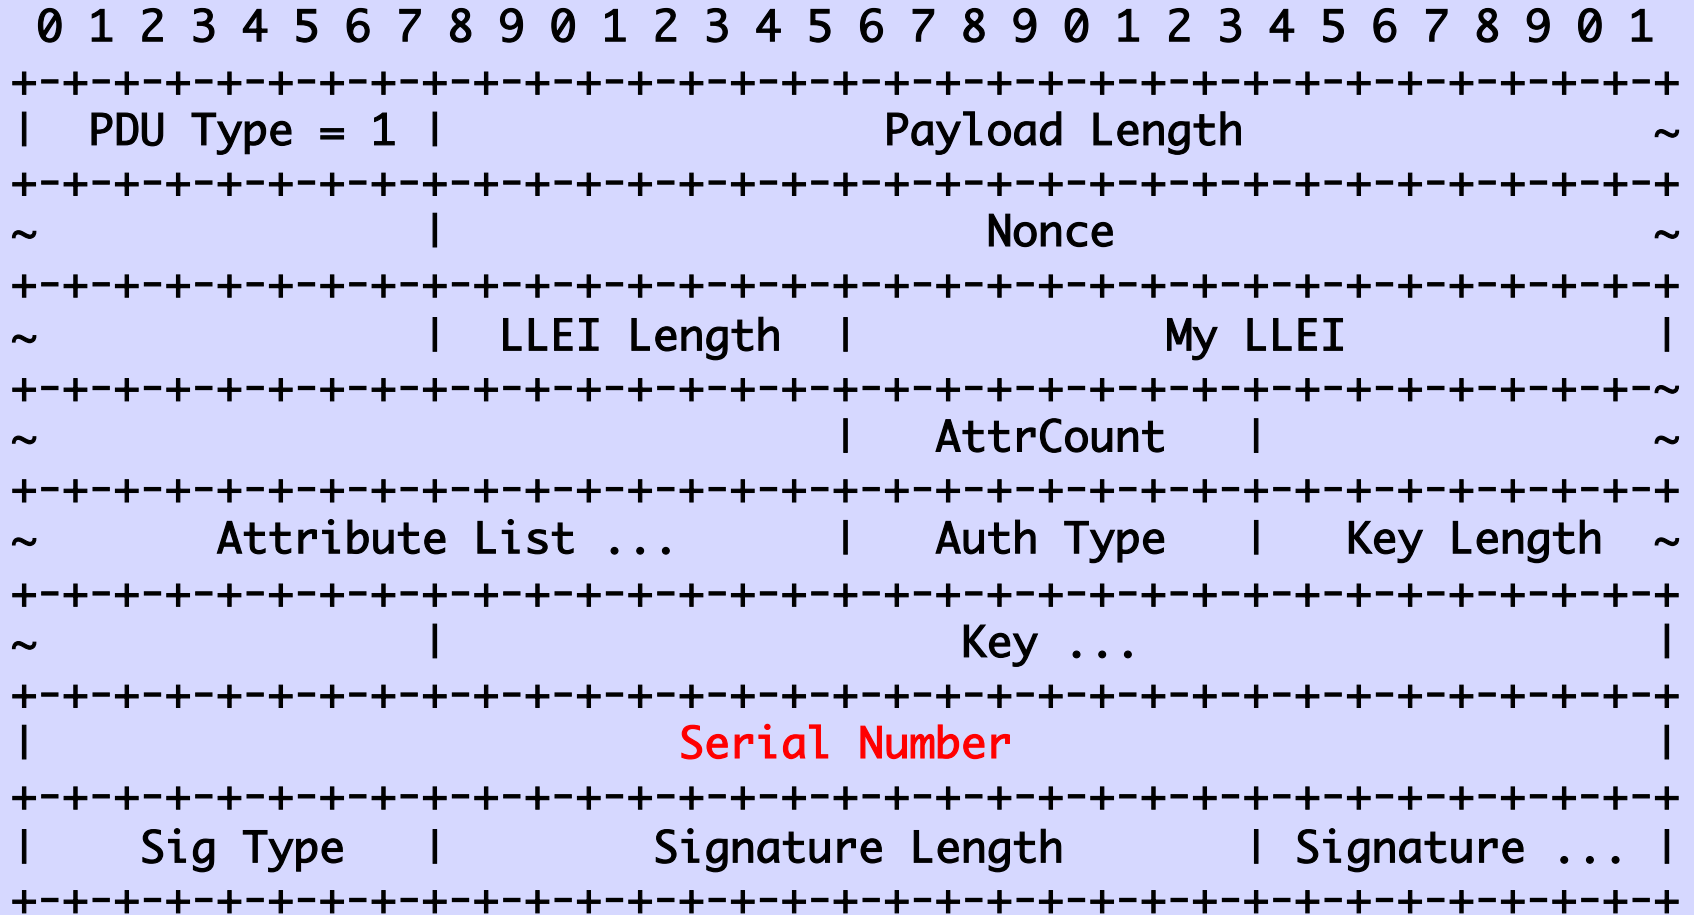
# Graceful Restart

State May Be Resumed  
à la BGP

# Encaps etc PDUs



# OPEN PDU



# Duplicate Datagram and PDU Detection

Retransmission  
on Time-Out for ACK

HELLO and KEEPALIVE  
are still separate

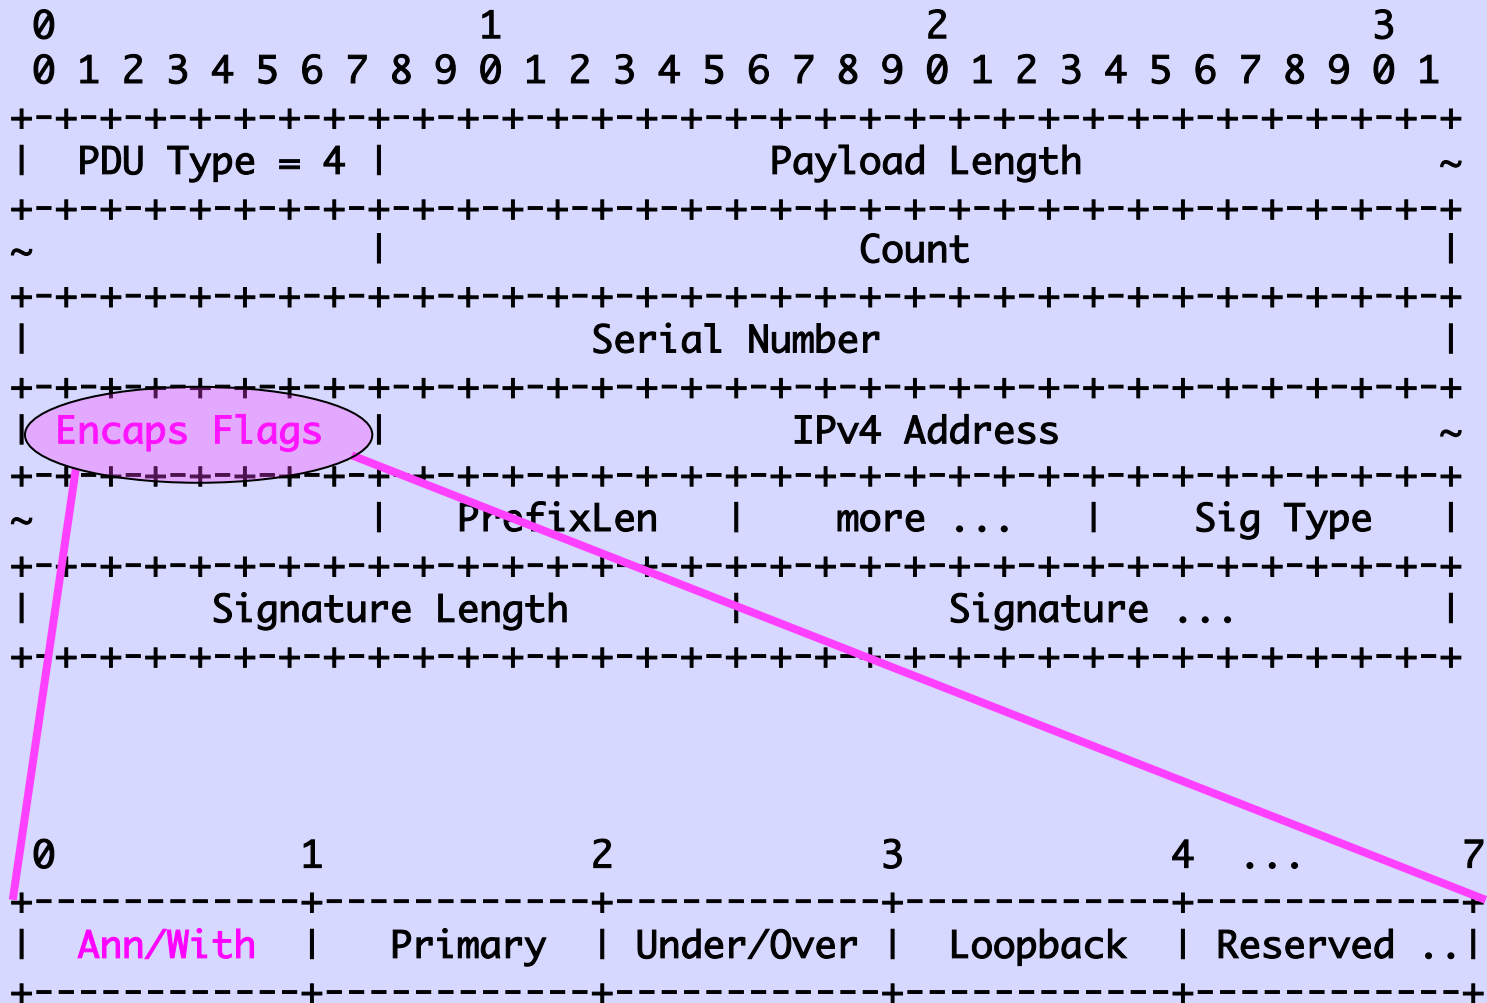
UniCast vs MultiCast  
& Separate Timers

HELLO does not repeat  
on a P2P Link

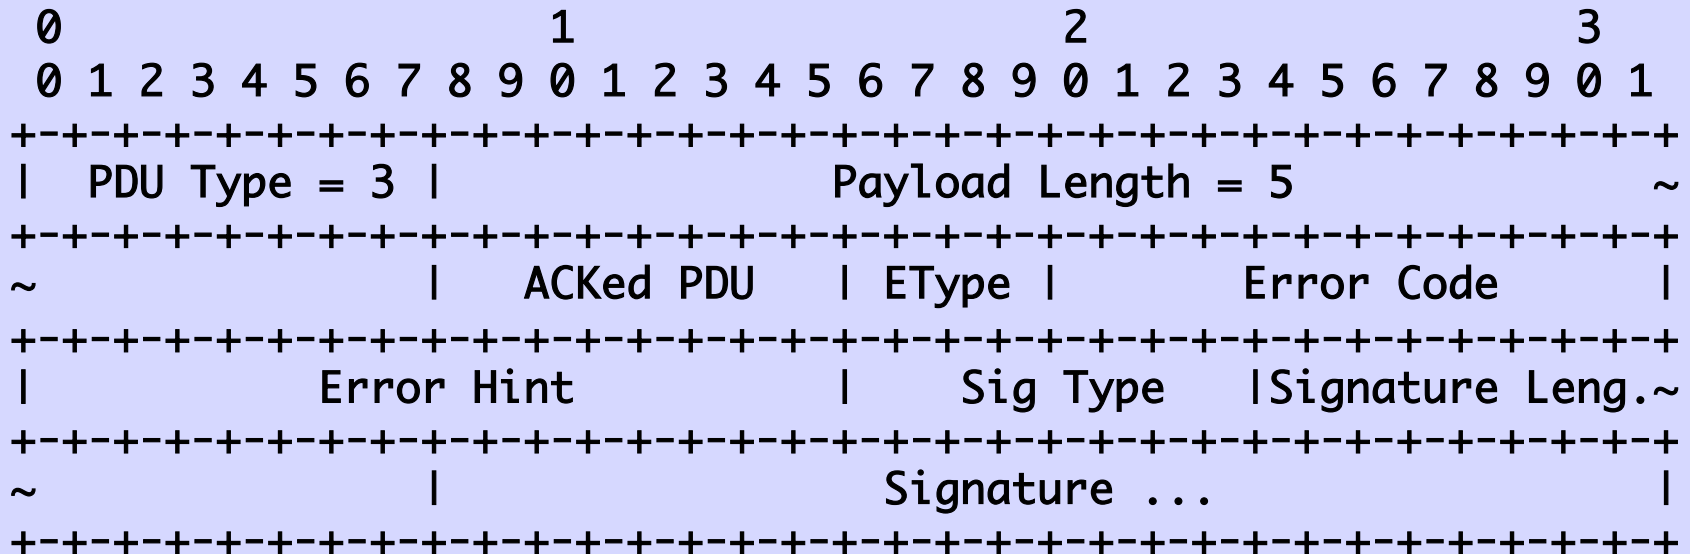
Only on Multi-Link



# Announce/Withdraw




# Explicit ACK/ERROR



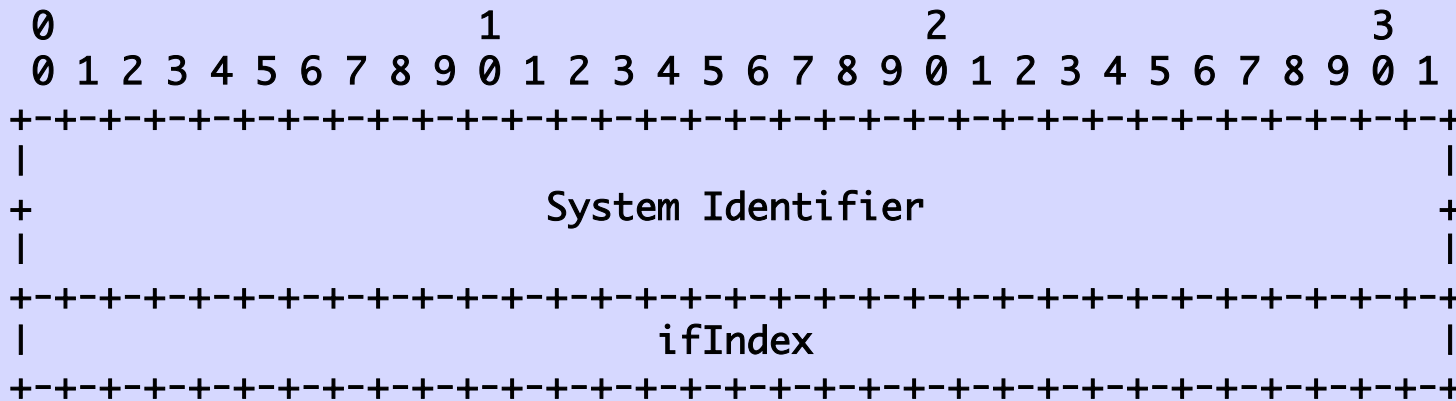

EType, Error Code, Error Hint = 0 /\* no error, just an ACK \*/

# Vendor Extension



# LLIE

## Logical Link Endpoint Identifier



For Interoperation,  
do we need a Mandatory  
to Implement  
LLEI?

We now have Two  
Implementations  
One Python3 (LSOE)  
One in Golang

# Going into ArcOS & Maybe a FOSS Interop in Singapore and WGLC

# Neeraj has Extended it



Always a Good Sign

draft-malhotra-bess-evpn-lsoe-00.txt



# DRAFT NEEDS REVIEW

Would prefer to get  
volunteer reviewers  
without having to  
WGLC it

# L3DL-ULPC

## Upper Layer Protocol Configuration

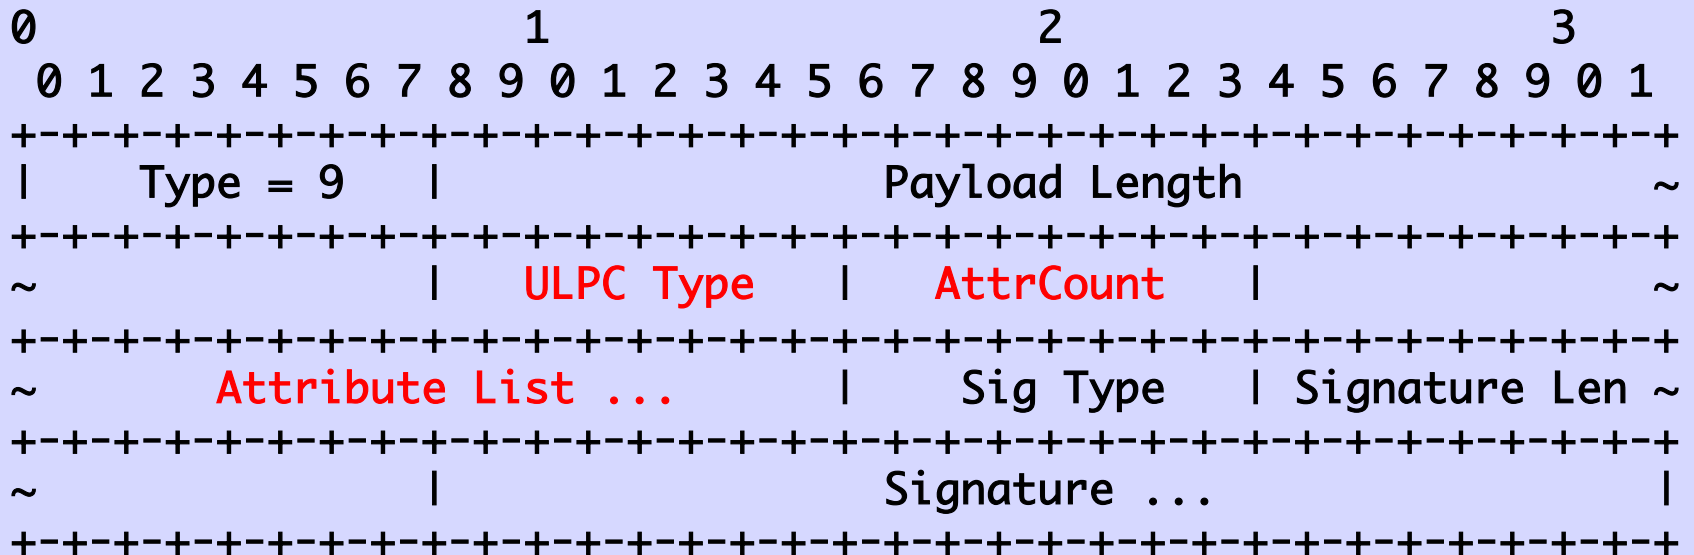
draft-ymbk-lsvr-l3dl-ulpc

LSVR WG

2019.07.25

randy@psg.com, sra@hactrn.net, keyur@arrcus.com

# L3DL PDU for ULPC



Provide the minimal set  
of configuration  
parameters for BGP  
OPEN to succeed

Not to replace or  
conflict with data  
exchanged by  
BGP OPEN

Multiple sources of  
truth are a recipe for  
complexity and pain

# ULPC for BGP

```

0                               1                               2                               3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Attr Type = 1 | Attr Len = 48 |                               My ASN                               ~
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
~
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

```

0                               1                               2                               3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Attr Type = 2 | Attr Len = 56 |           My IPv4 Peering Address           ~
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
~
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           Prefix Len           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

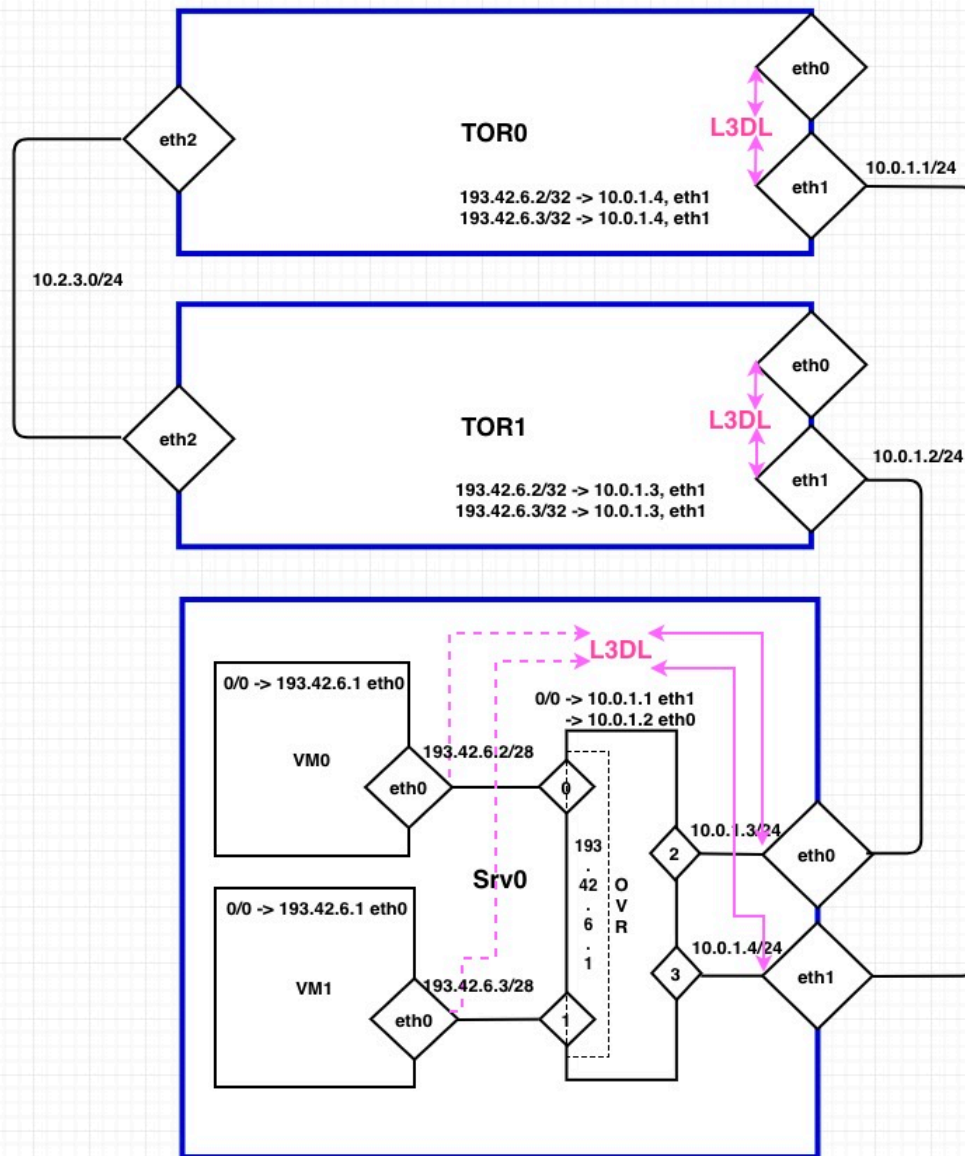
```

0                               1                               2                               3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Attr Type = 4 | Attr Len |                               ~
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
~
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           BGP Authentication Data ...           ~
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

# Backup Slides





**I3DL Config**  
**interface eth0**  
 redistribute vm0.eth0 as overlay  
 redistribute vm1.eth0 as overlay  
**interface eth1**  
 redistribute vm0.eth0 as overlay  
 redistribute vm1.eth0 as overlay

← active →  
 - - - passive - - -  
 passive means  
 - get the attributes of the interface  
 - but do not speak I3dl on it

# Overlay Multi-Tenancy

- Overlay SVI on OVR attached to VRF BLUE
- Overlay traffic tunneled with VRF VNI encapsulation between OVR/Server and TOR
- L3DL carries overlay IPs with VRF identifier attribute (VNI 10) from Server to TOR

