

MPLS

MPLS SHIM Header

LABEL	EXP	S	TTL
20b	3b	1b	8b

- 0- IPv4 explicit null (QoS)
- 1- router alert
- 2- IPv6 explicit null
- 3- implicit null (PHP)
- >15 - reserved

ImplicitNull adv. for DirectConn./Summaries

- labels: fixed length, local significance
- IOS: normally per-platform label space
- distribution:
 - **unsolicited downstream (IOS)** / downstream on demand
- retention:
 - **liberal (IOS)** / conservative
- control
 - **independent (IOS)** / ordered

MPLS VPN RRs

RRs differ from other vpnv4 speakers in that they DON'T reject routes with RT unknown, but forward them unmodified.

RR Groups

We can have RR groups that only permit some RT (for example Even/Odd)

Hub&Spoke: usually 1 RD / spoke, 2 RTs

- **no-decrement-TTL** (Juniper, spec. LSP)
- **no-propagate-TTL** (interoperable)

Default: propagate TTL

TTL Expiration: ICMP Time Exceeded goes EgressLSR and is sent IP from there.

MPLS VPNs PE-CE protocols:

- **connected** (redistribution)
- **static** (redistribution)
- **RIPv2** (not v1!) - set default-metric!
- **OSPF**
 - redistribution normally transforms routes to External, so instead they are sent through MP-BGP as **Type 3**
 - OSPF metric is saved to MED
 - Ext. community defined: **OSPF RouteType Attribute:** Area, OSPF RID, Domain ID, Metric Type 1/2
 - if **Domain ID != OSPF RID** (Juniper compares with local DomainID) advertised as Type 5 (E2)
 - **Sham Link:** if backdoor link used, all routes through it (type1, 2, 3) will be preferred against Type3 through MPLSVPN. Sham Link is a fake link between PEs treated as demand circuit and all LSAs are sent through it normally. Its endpoints must be adv. in iBGP vpnv4
 - **Down Bit** (in Lsa Type 3): if a PE adv. a route in OSPF area, sets Down Bit and another PE will not readvertise it back in MPLS VPN iBGP. Avoids loops.
 - **Domain Tag/Route Tag:** same thing but for Type 5.
- **EIGRP**
 - routes arrive as External (like OSPF) so new Extended Communities created for each needed element
- **ISIS**
- **eBGP** (not iBGP!)
 - **AS Override** (change VPN AS with ISP AS so that the prefix is allowed on the other side)
 - **Allows-in** (disable check of AS-path loop)

SOO can be used for loop prevention (extended community) - set with route-map when BGP used PE-CE

L3 VPN.

New NLRI:

- mask
- Label
- RD (64bit) (uniq per VRF/VPN)
- prefix

Ext. Community: Route Target RT

Routes discarded if non-local RT

L2 VPN Martini (LDP)

Uses **targeted LDP** between PE

Packet:

- FEC type
- C-Bit & VC Type (encapsulation)
- Group ID (not used)
- VC ID
- IF Params (MTU,...)

L2 VPN Kompella (BGP)

NLRI:

- RD
- Label Block Offset & Label Base (more labels can be announced)
- SUBTLVs

L2 Community:

- Ext. community type
- Encap Type (4 ETHVlan, 5 ETH)
- L2 MTU

VCID in Junos: **remote-site-id**

In **ATM** cells are reassembled by PE, sent as AAL5, segmented back at end Vlans must match! DLCI not. **IP Interworking:** different encaps at ends.