

# Update on BIER Architecture and BIER MPLS Encapsulation

- A few changes and additions since early revisions of the drafts
- Some issues to which the WG should attend

# MPLS Encapsulation Changes

- Entropy field length: 8 → 20 bits
  - Now same length as MPLS entropy label
  - Note: when tunneling through non-BFRs
    - BIER encaps entropy should → tunnel encaps entropy
    - Thus tunnel “rewrite string” depends on tunnel payload
- BFIR-id field: optional 32-bit → mandatory 16-bit
- Bit flags: 8 → 16 bits
  - Will probably change further
  - Please, if you think you need a flag, do not resort to self-help

# MPLS Encapsulation Issues

- Newly Raised Issue re *Version Field*
  - First nibble declared to be version field
    - 4 and 6 “reserved” to avoid spurious IPv4/IPv6 ECMP treatment at transit nodes
  - Given use of first nibble in ECMP and other data plane heuristics, maybe not a good place for version field
  - Fixed value probably better:
    - 5 could be a poor man’s protocol type ☺
    - 0 aligned with PW usage, maybe would then use 1 for OAM
- (Let’s try to avoid the MPLS payload protocol type rathole)

# Additions to Architecture

- Sub-domains
- Automatic tunneling through non-BFRs
- Deterministic ECMP (presently optional, but ...)
- BitStringLength Issues

# Domains and Sub-Domains


- *Domain* is scope of control plane (e.g., IGP signaling scope)
  - BFR-prefix: per-domain (control plane) scope
- A domain contains one or more *sub-domains*:
  - Each domain contains at least the default sub-domain, sub-domain 0
  - BFR-id: per-sub-domain scope
- What is a sub-domain?
  - An assignment of BFR-ids to BFIRs and BFERs
  - A routing underlay (e.g., an IGP topology)

# Sub-Domains

- Per architecture doc, each BFR is *provisioned* to know:
  - the sub-domains to which it belongs
  - The routing underlay for each of those sub-domains
  - Its BFR-id (if any) for each of those sub-domains
    - BFR-id zero not legal, reserved for use by control plane to mean “no BFR-id”
- **No such thing as dynamically joining a sub-domain**
- BIER-MPLS label (BoS label in BIER packet) bound to <Sub-domain, BitStringLength, SetIndex>
  - This is what is needed to properly interpret the BitString
  - BitStringLength not part of sub-domain


# Automatic Tunneling

## From BFR to BFR through non-BFRs




Whoops, need to get packet from A to {D,G}, but E hasn't advertised MPLS label for packet's sub-domain

- Act as if topology is one of:



or



- Tunneling can be as simple as pushing on MPLS unicast label for F or G (with suitable entropy, TTL)
- Same procedure useful if E goes down

# BitStringLength

- Each BFR provisioned with:
  - BitStringLength to use as BFIR
  - BitStringLengths supported as BFR/BFER
  - Not specific to sub-domain
  - All BFRs/BFERs in domain **SHOULD** support any BitStringLength that may be originated in that domain
- But what if BFR next hop doesn't support the BitStringLength in a packet that you have to forward?
  - **MAY** drop the packet
  - **MAY** modify the BitStringLength (might require making a second copy of packet)
  - **MAY** tunnel through (as per slide 5)



# Deterministic ECMP

- Deterministic ECMP: (S,G) path from BFIR-A to BFER-B depends only on routing underlay plus entropy
- Without: path may change as other BFERs join/prune
- OAM problem when ECMP is non-deterministic:
  - BFER-B complains about intermittent problems
  - Problems occur only when BFER-C joins the stream (resulting in change of path), but no one knows this
  - Difficult to troubleshoot
  - Deterministic ECMP eliminates this problem (cost is possible extra packet-hops)
- Troubleshooting does require knowledge of entropy, but this is just function of (S,G)