

Forwarding Actions

draft-kompella-mpls-mspl4fa

Kireeti Kompella (Juniper) [presenter]

Vishnu Beeram, Tarek Saad (Juniper)

Israel Meilik (Broadcom)

Genesis of draft

- to initiate discussion into the use of the TC/TTL fields to increase the expressiveness of a bSPL (and potentially, eSPL)
- to capture multiple forwarding actions with a single bSPL
- to expand the data fields associated with an SPL from 20 to 31 bits
- to allow for dynamic/policy-based definition of the data fields
- to **do this efficiently** in today's forwarding pipelines

Indicators

- A bSPL “indicates” forwarding actions and possible associated data
 - Forwarding actions: router alert, NFFRR
 - Associated data, either in the label stack: ELI, Flow ID, GISS; or after the label stack: EHI, GAL, GDF, OAM
 - New requests for bSPLs are similar (NFFRR, GDF, GISS, Flow ID, EHI, ...)
- The FAI proposes to use “indicator flags” instead of an entire label for this purpose
 - to solve the above problems, including ...

bSPL scarcity

- Among the topics discussed in the DT
- 8 of 16 bSPLs have been allocated
 - RFC 7274 introduces extended SPLs, but each needs two labels
- Currently, there are requests for 6 bSPLs of the 8 remaining
- The FAI proposal allows a single bSPL to capture 10* actions; with each extension, another 30* actions can be captured for four octets
 - The bSPL scarcity problem can be put to rest
 - We may never need extended SPLs (besides the two already allocated)

* using one bit for extension

Two Tracks for the DT

Extend the MPLS architecture to take into account forwarding requirements and forwarding hardware for the 21st century

- within the label stack (LS)
- after the EoS (payload or PL)

This draft focuses on LS data

- Indicators for PL data are also needed
- One proposal is that PL data be self-describing
 - If so, detailed indicators are not needed in the FAI

-00 version

- Had an “Opaque Data” field, where the semantics of the data was determined by policy (not elaborated in the draft)
 - (some issues to solve; see slide 9)
- Had a “FA Header” for expansion of the flags (not elaborated)

- Both were raised in the DT discussions

-01 version

- Elaborated on the use of the FA header for expansion
- Started rationale for when data should be in the LS vs in the PL
- Moved flags around, made space for “edist”
 - Distance in words to the end of stack, to optimized PL access
- Updated processing of the LS FAD and examples

Discussion Points (1)

- The FAI *must* be processed efficiently in the data path
 - A TLV structure is flexible, but not the most efficient; hence the use of flags
 - Other approaches?
- Should future bSPLs use a similar approach?
 - What this draft proposes may be the first: “FAI-1”
 - Even eSPLs can use this approach
- Rationale of when to put data in LS versus in PL
 - Centers on the use of data and the efficiency of the data plane
- Multiple FAIs in a label stack (e.g., FAI per segment)
- The value of having “edist” vs 4 more indicator flags
 - Label stacks are getting very large, so finding the EoS efficiently is useful

Discussion Points (2)

- The idea of flexible definition of forwarding actions and associated data (using policy or other means) should be debated
 - This allows flags/fields to be programmable and associated actions to be on-demand
- Potential issues are:
 - Efficiency
 - Consistency across the network
 - Multivendor consistency
 - Incremental change of definition
- Standard actions and flexible FAI are not mutually exclusive
 - The Q flag indicates that such data exists