

Network Functions

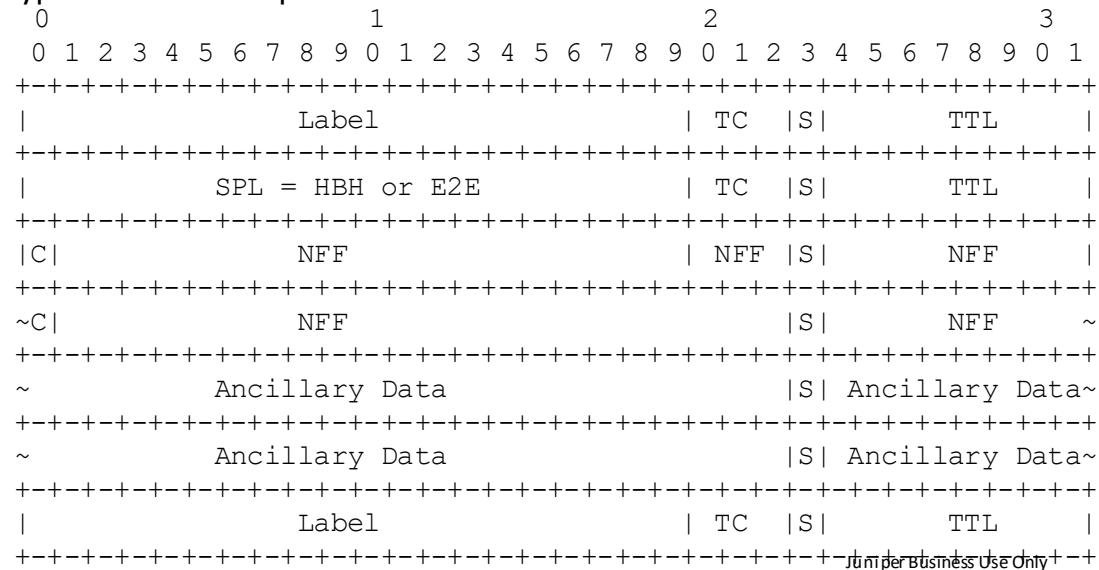
John Drake and Adrian Farrel

Network Functions (NF)

- Derived from <https://datatracker.ietf.org/doc/html/draft-kompella-mpls-mspl4fa-01>
 - Identify the forwarding actions to execute using information in the label stack
 - An SPL indicates that a bit-map follows that identifies forwarding actions
 - Additional data may be included
- Attempt to generalize and flesh out Kireeti's proposal
 - Forwarding Actions -> Network Functions (NFs)
- Two SPLs
 - Hop by Hop (HBH) and End to End (E2E)
 - P-nodes only process HBH SPLs
 - Suggested by Adrian and Stewart
 - Allows both types of NFs to evolve independently w/o interference
 - Both are completely extensible
 - Non-label fields are unchanged
- When defining (via an RFC) a new NF
 - Specify whether it is HBH or E2E
 - Specify its bit in the Network Functions Flags (NFFs) (see below)
 - Specify whether it has ancillary data
 - If so, what is it and is it "in-stack" or "after-stack"

Network Function Label Stack Block (NFLSB)

- Common format for both types of NFs (HBH and E2E)
- SPL, NFFs, and in-stack ancillary data
- SPLs
 - Two different SPLs make it easy for transit nodes to find which functions to process
- NFFs
 - Contiguous set of bits, one per each currently defined NF understood by the ingress node
 - Set of one or more label stack entries (LSEs)
 - Each LSE has a Continue bit, 19 NFFs, 3 NFFs, the BoS bit, and 8 NFFs
 - Continue bit indicates another LSE with NFFs
- In-stack ancillary data
 - 31 bits of ancillary data for each NF whose NFF is set
 - In the same order as the NFFs
- Delimited by another SPL, or an LSE with BoS set, or ancillary data with BoS set
 - Care must be taken that ancillary data does not match an SPL
- If both types of NFLSB are present HBH MUST be first



Ancillary Data

- A node understands the NFFs up to a given bit position N
 - Does not need to support all defined NFFs less than N
 - Skips over ancillary data for NFFs it doesn't support
- Stops processing the NFLSB for NFFs > N
- Given packet will have a set M whose NFFs are set and understood by the node
 - Size of (M) <= N
 - Within M
 - Set X (NFFs w/o ancillary data)
 - Set Y (NFFs w/ in-stack ancillary data)
 - Set Z (NFFs w/ after-stack data)
 - Size of (X + Y + Z) = Size of (M)
 - Ancillary data for Y in NFLSB after NFFs in the same order as elements in Y
 - Ancillary data for Z after-stack in the same order as elements in Z
 - If a set is empty there is no ancillary data of that type

