Stacked Tunnels, OSPF/ISIS as Label Distribution Protocols, and MPLS Architecture

RFC 3031 ("Multiprotocol Label Switching Architecture") provides all the concepts and terminology needed to describe (a) source routed tunnels, including the stacked-label variation, and (b) OSPF/ISIS as a label distribution protocol.

This presentation places stacked-label source routed tunneling, and OSPF/ISIS as a label distribution protocol in the context of the MPLS Architecture, using the terminology defined there.

MPLS Architecture and Explicit Routing

- *Explicitly Routed Tunnels*
  - "If a Tunneled Packet travels from Ru to Rd over a path other than the Hop-by-hop path, we say that it is in an "Explicitly Routed Tunnel"" (section 3.27.2 of RFC3031)
- Explicitly routed LSPs as a way to realize explicitly routed tunnels
  - "An "Explicitly Routed LSP Tunnel" is a LSP Tunnel that is also an Explicitly Routed LSP" (section 3.27.3 of RFC3031)
    - RSVP-TE (or CR-LDP) as the label distribution protocol for explicitly routed LSPs

MPLS Architecture and LSP Hierarchy

- *LSP Hierarchy*: LSP Tunnels within LSPs (see section 3.27.4 of RFC3031)
- "The label stack mechanism allows LSP tunneling to nest to any depth" (section 3.27.4 of RFC3031)

Stacked LSPs for Explicitly Routed Tunnels

- Instead of an explicitly routed LSP, one can use LSP hierarchy (stack of LSPs) to realize explicitly routed tunnels
- All LSPs in the stack have the same "LSP ingress"
- "LSP egress" of a given LSP in the stack is an intermediate point of the next LSP in the stack
- LSP egress of a given LSP in the stack could be either single or multi-hop away from the LSP egress of the next LSP in the stack
- Such stack of LSPs provides the functionality to forward a packet through a sequence of "LSP egress" of the LSPs on the stack
  - the sequence of "LSP egress" represents the explicit route
- The label stack mechanism allows stack of LSPs to nest to any depth

Incoming and Outgoing label in Label Swap

- From RFC3031:
  - label swap: the basic forwarding operation consisting of looking up an incoming label to determine the outgoing label, encapsulation, port, and other data handling information
- It is legal in the MPLS Architecture to preserve the label as a packet transits an LSR
- This means the incoming and outgoing label values just happen to be the same, and so the swap produces the same output as its input
- If for a given LSP all the LSRs traversed by the LSP have the same incoming and outgoing label, then the value of the label stays the same from the LSP ingress all the way to the LSP egress

MPLS Architecture and Label Distribution Protocol (1)

- "*THE ARCHITECTURE DOES NOT ASSUME THAT THERE IS ONLY A SINGLE LABEL DISTRIBUTION PROTOCOL*. In fact, a number of different label distribution protocols are being standardized." (section 3.6 of RFC3031)

MPLS Architecture and Label Distribution Protocols (2)

- *Local/remote label distribution peers*
  - When two LSRs are IGP neighbors, we will refer to them as “local label distribution peers”. When two LSRs may be label distribution peers, but are not IGP neighbors, we will refer to them as “remote label distribution peers” (section 3.27.5 of RFC3031)

OSPF/ISIS as a Label Distribution Protocol (1)

- OSPF/ISIS advertisements by a router carry label bindings for LSPs that transit through the router
- The router can be either single or multi-hop away from routers that receive these bindings
  - Local label distribution peers are the IGP neighbors of the router
  - Remote label distribution peers are other routers in the same IGP domain
- The router could be either single or multi-hop away from the egress of these LSPs
  - Existing label distribution protocols (LDP, RSVP-TE, etc.) can be used to establish multi-hop LSP fragments
  - See also slide 5

OSPF/ISIS as a Label Distribution Protocol (2)

- Each OSPF/ISIS router passes Link State Advertisements originated by other routers unmodified
  - Including the label binding information
  - Similar to how BGP Route Reflectors handle routes/labels
- Provides scalable support for remote label distribution peering
  - Label distribution protocol messages (Link State Advertisements) are exchanged only between IGP neighbors
    - Control plane peering only between IGP neighbors
  - No control plane peering between a router and each of its remote label distribution peers

Example

Assume that R1 wants to forward certain set of packets along the explicitly routed tunnel (R1, R2, R3, R4, R5)

- (R1, R2), (R2, R3), (R3, R4), and (R4, R5) are local label distribution peers (as they are IGP neighbors); all other pairs are remote label distribution peers

Explicitly routed tunnel (R1, R2, R3, R4, R5) is realized via the following LSP Stack:

- LSP1: (R1, R4, R5)
- LSP2: (R1, R3, R4)
- LSP3: (R1, R2, R3) - top of the stack

To get from the first hop of LSP1 to the second hop of LSP1, the packet has to go through an LSP tunnel, LSP2.To get from the first hop of LSP2 to the second hop of LSP2, a packet has to go through an LSP tunnel, LSP3.

R1 receives (via OSPF/ISIS) from its local label distribution peer R2 label bindings originated by R4, R3, and R2.

Label Stack construction at R1:

Step 1: R1 uses label binding L1 originated by R4 for LSP1 (R1, R4, R5). So, R1 starts building the label stack by pushing L1 onto the label stack.

Step 2: R1 uses label binding L2 originated by R3 for LSP2 (R1, R3, R4). So, R1 pushes L2 into the stack. At this point the stack contains (L2, L1).

Step 3: R1 uses label binding L3 originated by R2 for LSP3 (R1, R2, R3). So, R1 pushes L3 into the stack. At this point the stack contains (L3, L2, L1).

Step 4: Since R1 and R2 are local label distribution peers, label stack construction is completed

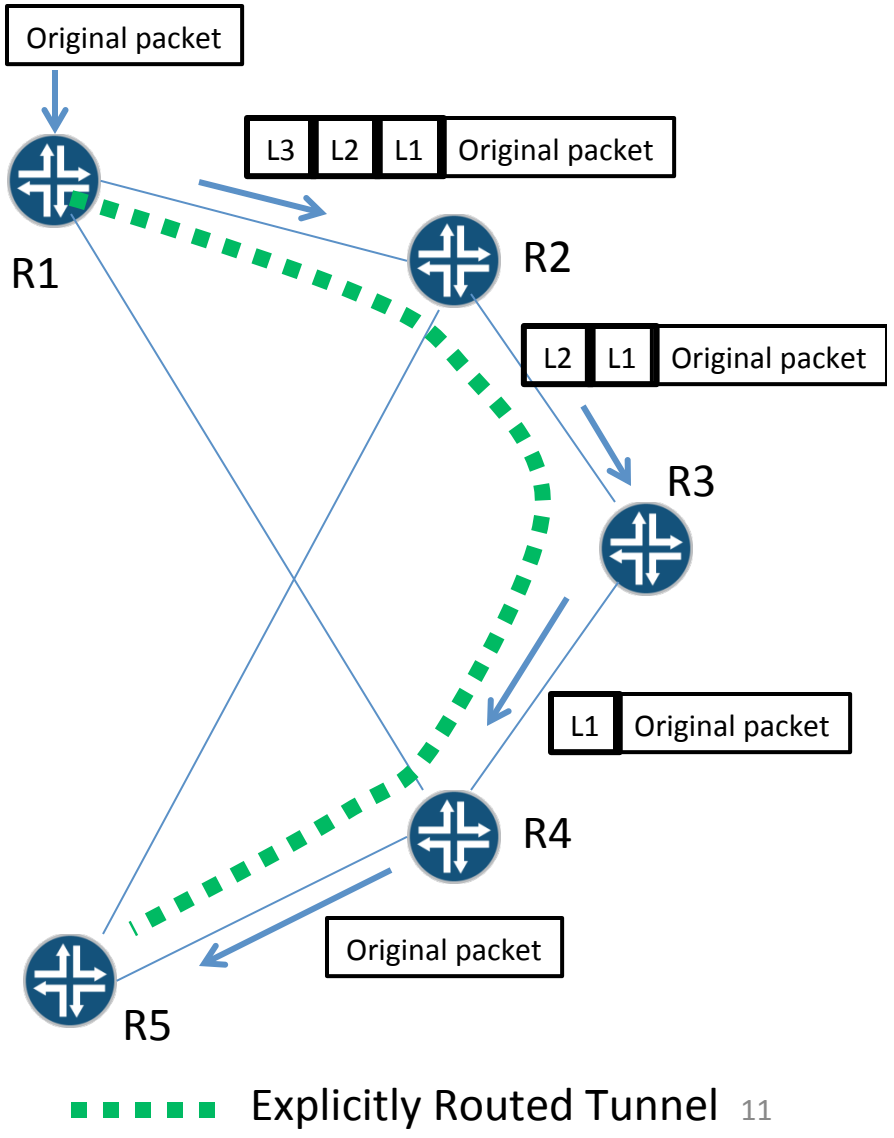

Conclusion

- Stacked LSPs for explicit/source routed tunnels and use of ISIS/OSPF as a label distribution protocol, both fit comfortably within the MPLS architecture.
- As the MPLS architecture defines the necessary concepts and terminology, no new architecture, concepts or terms are required.