

IGP Extensions for Intra-Domain SAV

draft-chen-savnet-lsr-intra-00

Huaimo Chen (Futurewei)

Weiqiang Cheng (China Mobile)

Aijun Wang (China Telecom)

Gyan S. Mishra (Verizon)

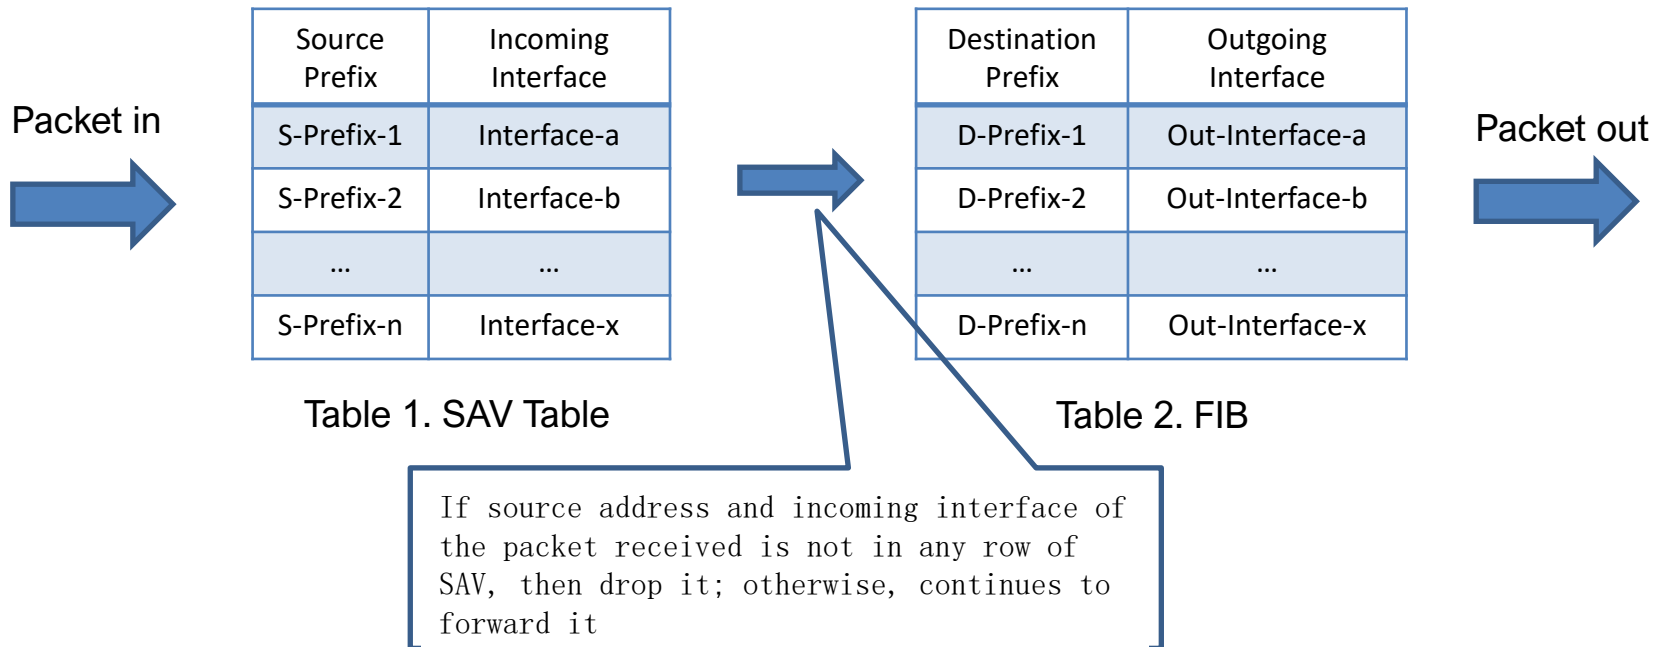
Yanhe Fan (Casa Systems)

Xufeng Liu (Alef Edge)

IETF 118

Overview

- IGP on every node builds its SAV Table based on its LSDB or RIB.
- SAV table is used to validate source address of packet received.



IGP extensions for

1. **Fast Convergence** of SAV table on Routing Changes (0 to 1 SPF)
2. **Accurate** Validation (Illegal Packet dropped)
3. **Backward Compatible** (Work in Incremental/Partial Deployment)

Intra-Area SAV table (AS has 1 Area)

- When every link in AS/Area is symmetric,
IGP on each node **builds SAV table using RIB (0 SPF for SAV)**
by having a row in SAV table for each prefix with a NH interface in RIB

- When a link in AS/Area is not symmetric,
IGP on each node X builds SAV table in 3 steps (1 SPF for SAV):
 1. **Builds reserve shortest path tree (RSPT)**.
builds a SPT from X to other nodes using cost of each link in reverse direction. **(1 SPF)**
 2. **Builds reserve routing table (RRT/RRIB) using RSPT.**
For a shortest path from X to Y with a next hop interface in its RSPT, adds an entry for each prefix attached to Y into its RRT. The entry has the prefix as the destination and the next hop interface as the next hop.
 3. **Builds SAV table using RRT/RRIB.**
has a row in SAV table for each prefix with a next hop interface in RRT/RRIB.

Options for scope of prefixes to be validated:

- a) Prefixes attached to every node (the above)
- b) Prefixes attached to each ASBR and ABR (consider only Y=ASBR, ABR in step 2)
- c) Prefixes indicated/configured on any node (consider only prefixes attached to node Y and indicated/configured by Y in step 2).

Inter-Area SAV table (AS has 1+ Areas)

IGP on each node X in area A builds SAV table in 4 steps (0-1 SPF for SAV):

1. Gets area shortest path tree (ASPT)

The ASPT is a tree from X as root to the other nodes in area A. If every link in A is symmetric, ASPT is SPT to the other nodes for RIB (reused); otherwise (i.e., asymmetric link in A), ASPT is RSPT from X to the other nodes in A.

2. Extends ASPT

Adds non-ABR's prefixes to ASPT; If ASPT is SPT and every path in A' between an ABR and a summary prefix is symmetric (no reverse cost to prefix Sub-TLV from A'), extended ASPT reuses portion of SPT with summary prefixes attached to ABR for RIB; otherwise, attaches summary prefixes to ABR using reverse cost to prefix.

3. Builds reserve routing table (RRT/RRIB) using extended ASPT


For a shortest path from X to Y with a next hop interface in its extended ASPT, adds an entry for each prefix attached to Y into its RRT/RRIB.

4. Builds SAV table using RRT/RRIB

has a row in SAV table for each prefix with a next hop interface in RRT/RRIB.

The last two steps here are similar to the last two steps for one area case.

```
1. IF A is symmetric {ASPT = SPT}
   ELSE {ASPT = RSPT};
2. adds non-ABR's prefixes to ASPT;
   IF ASPT==SPT and A' is symmetric {
     reuses SPT portion from
     ABR1/2 to Prefix for RIB
   } ELSE {
     attaches Prefix to ABR1/2
     using reverse cost
   }
}
```



Extensions to IGP

- New indication (e.g., **S-Flag**) indicating a prefix to be validated **when option 3** is selected
- A new Sub-TLV, called **Reverse Cost to Prefix Sub-TLV**, for ABR to **advertise** the cost of the shortest path from prefix to ABR **when the path** between ABR and prefix is **not symmetric**).

Next Step

- Comments


Comparisons with another

Differences:

1. Number of SPF's for SAV table
0-1 SPF vs Multiple SPF's
2. Options for Scope of Prefixes to be Validated
3 options vs 1 option (i.e., option 3)
3. IGP Extensions Depend on Options
Yes vs No
4. Extensions
S-flag in existing Sub-TLVs vs New Sub-TLVs
5. Reverse cost from ABR to summary prefix is distributed
if asymmetric vs in any case (i.e., if symmetric or asymmetric)

Extensions to IGP: S-Flag

- New **S-Flag** (SAV Prefix Flag) indicating a prefix to be validated **when option 3** is selected
 - OSPFv2:
 - 0x20 - S-Flag (SAV Prefix Flag) in existing OSPFv2 Extended Prefix TLV:
Set when the prefix is configured for SAV (i.e., to be validated).
 - IS-IS:
 - Bit 5 - S-Flag (SAV Prefix Flag) in existing Prefix Attribute Flags Sub-TLV:
Set when the prefix is configured for SAV (i.e., to be validated).
 - OSPFv3:
 - 0x01 - S-Flag (SAV Prefix Flag) in new OSPFv3 Prefix Attribute Flags Sub-TLV:
Set when the prefix is configured for SAV (i.e., to be validated).




Extensions to IGP: Reverse Cost to Prefix Sub-TLV

- A new Sub-TLV, called **Reverse Cost to Prefix Sub-TLV**, for ABR **to advertise** the cost of the shortest path from prefix to ABR **when the path** between ABR and prefix is **not symmetric**).


- **OSPFv2:**

New OSPFv2 Sub-TLV in Existing OSPFv2 Extended Prefix TLV



- **OSPFv3:**

New OSPFv3 Sub-TLV in Existing OSPFv3 Intra-Area-Prefix TLV



- **IS-IS:**

New IS-IS Sub-TLV in Existing TLV 135, 235, 236 and 237 for prefix

