

IS-IS Dst/Src Routing

draft-baker-ipv6-isis-dst-src-routing-03

Fred Baker · fred@cisco.com


David Lamparter · david@opencourcerouting.org

IETF 93, Prague, July 2015

Changes from -02 to -03

Completely removed "critical sub-TLV"
now using RFC5120 multitopology routing

Context refresher




⇒ ensuring correct exit taken based on packet source address

```
+-----+
| TLV 236/237, dst ::/0, etc.
| +-----+
| | TLV X, src 2001:db8:1::/48
| +-----+
+-----+
```

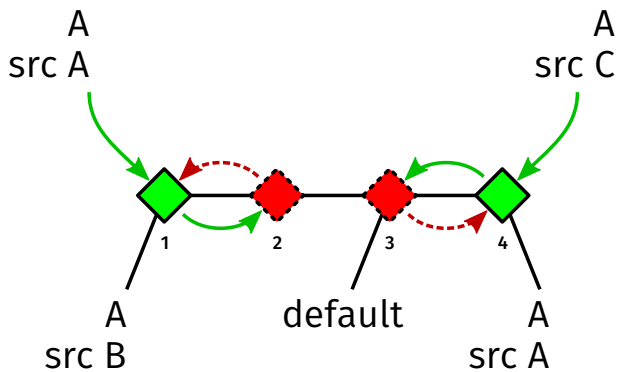
- ▶ sticking the source prefix in is straightforward
- ▶ implemented, tested, demo'd @ IETF 90
- ▶ open source: <https://git-us.netdef.org/projects/OSR/repos/openwrt-isis-hnet>

Interop / Compatibility

- ▶ naively mixing dst-src routers and non-dst-src routers leads to persistent loops
- ▶ might be out of scope for homenet
- ▶ certainly in scope for SMB & Campus applications



Loop cases



Using MT as separation mechanism


To fix these scenarios, we need:

- ▶ to fix loop #1: capability indication
- ▶ to fix loop #2: hiding reachabilities from “old” systems


MT provides both:

- ▶ participation in separate topology is capability indication
- ▶ TLVs in separate topology are invisible to non-participants

Loop cases

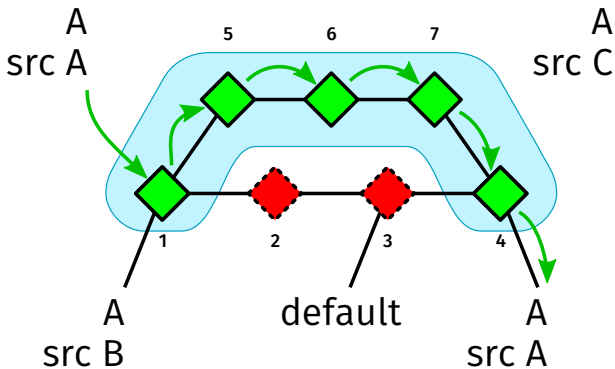


Loop cases



- ▶ A/A hidden from router #3 \Rightarrow default route works
- ▶ router #1 detects non-reachability of A/A

Loop cases



- ▶ connecting MT islands allows router #1 to reach A/A

Route installation

- ▶ MT SPF calculations as usual: separate for IPv6 default and D/S topologies
- ▶ routes are installed in the same table
- ▶ no overlap between the two sources
- ▶ D/S extended longest-match fulfills its purpose

Route installation

MTID 2

(IPv6 default topology)

```
TLV 237 2001:db8:1234::/48
TLV 237 2001:db8:5678::/48
TLV 237 ::/0
```



SPF
topology #2

```
2001:db8:1234::/48 nh A
2001:db8:5678::/48 nh B
::/0 nh C
```

MTID n

(IPv6 D/S topology)

```
TLV 237 2001:db8:abcd::/48
  subTLV x src 2001:db8:1111::/48
TLV 237 ::/0
  subTLV x src 2001:db8:2222::/48
TLV 237 ::/0
  subTLV x src 2001:db8:3333::/48
```



SPF
topology #n

```
...:abcd::/48 src ...:1111:/48 nh D
::/0 src ...:2222:/48 nh E
::/0 src ...:3333:/48 nh F
```



common routing table

Other relevant details

- ▶ separate D/S MTID per "parent" topology
- ▶ draft contains appendix with some (shallow) correctness considerations

Next steps

- ▶ hoping to get (positive) feedback on MT approach here
- ▶ expecting to do -04 to address comments
- ▶ updating implementations to actually use MT
- ▶ asking for WG adoption at between now and IETF 94