OSPFv3 over IPv4

I. Chen, F. Xie, and S. Abdallah

October 30, 2023

Evolution of OSPFv3

- OSPFv3 originally only supported IPv6 address family.
  - RFC5340 - OSPF over IPv6
- OSPFv3 standards enhanced OSPFv3 to carry both IPv4 and IPv6 address families over IPv6 transport.
  - RFC5838 - Support of Address Families in OSPFv3
- RFC7949 allows OSPFv3 to become the routing protocol of choice over both IPv4 and IPv6 transports.
  - RFC7949 - OSPFv3 over IPv4 for IPv6 Transition
  - An IP router can use OSPFv3 to replace OSPFv2 both in IPv4-only deployments and in mixed IPv4-IPv6 deployments

MITRE

Motivations

- Many existing networks
  - Have an existing IPv4 core
  - Difficult to migrate to IPv6 due to the sizable task of renumbering
  - Rely upon relatively low-speed radio links
- Replacing OSPFv2 with OSPFv3 over IPv4 transport
  - Is an intermediate step for eventual migration to IPv6
  - Allows the use of new features standardized in OSPFv3
  - Enables reduction in overhead bandwidth when compared to OSPFv3 over IPv6 transport

MITRE

Implementation

- Built a prototype implementation of RFC7949, based on FRR Release 7.3
- FRR provides an IETF standards-compliant implementation of OSPFv2/v3 with deployment experience both in commercial and tactical environments
- The modular design of OSPFv3 in FRR Release 7.3 makes the implementation of RFC7949 relatively straight forward.
- An indicator for the underlying IPv4 transport address was added to both the OSPFv3 interface and to the OSPFv3 neighbor structure.
- An AF_INET socket is added to handle OSPFv3 packets encapsulated in an IPv4 packet.
- When IPv4 is used as the transport mechanism, a different pseudo-header is used to calculate the OSPFv3 checksum and the corresponding AF_INET socket is used for transmitting and receiving OSPFv3 packets.
- The implementation improved FRR's sending of OSPFv3 packets by replacing the IPv6 link local address with an interface identifier. Replacing the address with the interface identifier increases modularity and reduces the potential for operational issues should an interface address change for any reason

Operational Considerations

- If the transport control is explicitly configured, either at the protocol/instance level or at the OSPFv3 interface level, the specified transport mechanism MUST be used.

- If an interface is not assigned an IP address required for the transport control configuration (except for unnumbered IPv4 interfaces), then the OSPFv3 interface becomes inactive.

- Removal of an IP address may result in de-activation of the corresponding OSPFv3 interface if the required transport configuration cannot be satisfied .

- Addition of an IPv4/IPv6 address may result in re-activation of the corresponding OSPFv3 interface if the required transport configuration can be satisfied.

- Changing the OSPFv3 transport at the process/instance level may result in activation/de-activation of the associated OSPFv3 interfaces depending on whether the new transport requirement on each interface can be satisfied.

- Changing the OSPFv3 transport at the OSPFv3 interface level may result in activation/de-activation of the OSPFv3 interface depending on whether the new transport requirement on the interface can be satisfied.

Conclusions

- Use IPv4 as the transport for OSPFv3
- Facilitates the eventual migration
- Enables reduction in overhead bandwidth when compared to OSPFv3 over IPv6 transport
  - Extremely important for many mission-critical tactical deployments

MITRE

Packet Analysis – PDU Byte Count Comparison

- n_1 is the number of LSAs
- n_2 is the number of LSAs requested
- n_3 is the number of LSAs acknowledged
- n_4 is the number of LSAs updated
- n_5 is the number of OSPF interfaces

Type	OSPFv2/IPv4	OSPFv3/IPv6	OSPFv3/IPv4
IP Header	20	40	20
OSPF Packet Header (OPH)	24	16	16
Authentication	Included in OPH	24	24
OSPF Hello Packet (OHP)	20	20	20
LSA Header (LH)	20	20	20
Database Description Packet (DBP)	OPH+8+n_1*LH	OPH+12+n_1*LH	OPH+12+n_1*LH
Link State Request Packet (LSRP)	OPH+ n_2*12	OPH+n_2*12	OPH+n_2*12
Link State acknowledgement Packet (LSAP)	OPH+n_3*LH	OPH+n_3*LH	OPH+n_3*LH
Link State Update Packet (LSUP)	OPH+n_4*LSA	OPH+n_4*LSA	OPH+n_4*LSA
Type-5 External LSAs	16	48	36
Type-4 Summary LSAs	8	12	12
Type-3 Summary LSAs	8	24	12
Type-2 Network LSAs	4	4	4
Type-1 Router LSAs	4 + n_5*8	4+n_5*16	4+n_5*16
OSPFv3 Link LSA	Not applicable	44	20
OSPFv3 Intra-Area-Prefix LSA	In Router LSA	52	12+8

MITRE

Six Node Topology Comparison

- Single area – No summary LSAs
- Packet capture from an OSPFv2 router of a network of six OSPFv2 routers
- Assumed the same sequence of events in both OSPFv2 and OSPFv3 over IPv4.
- Expect a 2% increase in protocol PDU total bytes for the duration of the run

Type	OSPFv2/IPv4	OSPFv3/IPv4	Change over OSPFv2
Hello Packet (count)	349	349	Not applicable
Hello Packet (byte)	30634	36218	+2%
Database Description Packet (count)	11	11	Not applicable
Database Description Packet (byte)	1170	1390	+2%
Link State Request Packet (count)	3	3	Not applicable
Link State Request Packet (byte)	246	294	+2%
Link State Acknowledgement Packet (count)	8	8	Not applicable
Link State Acknowledgement Packet (byte)	716	844	+2%
Link State Update (count)	39	39	Not applicable
Link State Update (byte)	7086	9258	+3%
Total (byte)	39852	48004	+2%

MITRE