

ISIS Flooding Reduction in MSDC


draft-xu-isis-flooding-reduction-in-msdc-00

Xiaohu Xu (Huawei)

IETF98, Chicago

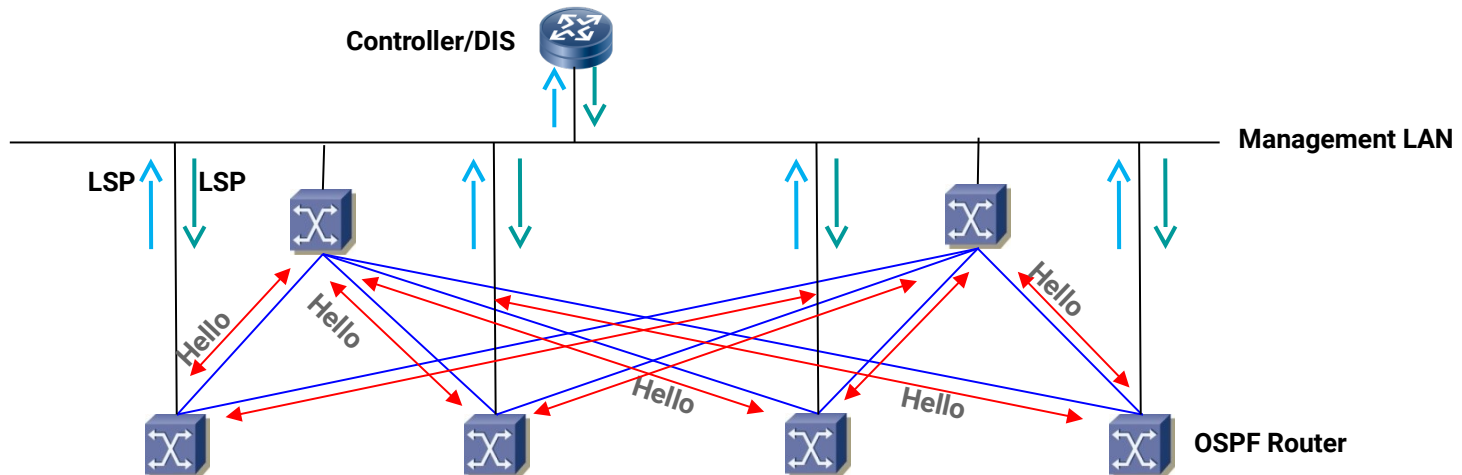
Problem Statement

- IS-IS is commonly used as a underlay routing protocol for Massively Scalable Data Center (MSDC) networks where CLOS is the most popular topology.
- Within the CLOS topology, a given IS-IS router would receive multiple copies of exactly the same LSP from multiple IS-IS neighbors. The unnecessary link-state information flooding wastes the precious process resource of IS-IS routers greatly and therefore IS-IS could not scale very well in MSDC networks.



Solution Overview

- **Mix of centralized link-state information distribution and distributed SPF calculation.**
 - All IS-IS routers within the CLOS network are connected with controllers via a management LAN.
 - IS-IS routers within the MSDC network just need to exchange IS-IS Hello packet among them so as to discover IS-IS neighbors.
 - The link-state information is only required to be exchanged between IS-IS routers and controller which is elected as IS-IS DIS for the management LAN.



Solution Details

- **After the exchange of IS-IS Hello packets among IS-IS routers, IS-IS routers would originate LSPs accordingly. However, these self-originated LSPs need not to be exchanged directly among them anymore. Instead, these LSPs just need to be sent solely to the controller being elected as IS-IS DIS for the management LAN.**
 - **IS-IS routers SHOULD NOT be elected as IS-IS DIS for the management LAN (This is done by setting the DIS Priority of those IS-IS routers to zero).**

Solution Details

(con't)

- **To further reduce the flood of multicast IS-IS PDUs over the management LAN, IS-IS routers SHOULD send IS-IS PDUs as unicasts.**
 - **IS-IS routers SHOULD send unicast IS-IS Hello packets periodically to the controller being elected as IS-IS DIS. In other words, IS-IS routers would not send any IS-IS Hello packet over the management LAN until they have found IS-IS DIS for the management LAN.**
 - **IS-IS routers SHOULD send other types of IS-IS PDUs to the controller being elected as IS-IS DIS as unicasts as well.**

Solution Details

(con't)

- **To avoid the data traffic from being forwarded across the management LAN, the cost of all IS-IS routers' interfaces to the management LAN SHOULD be set to the maximum value.**
- **When a given IS-IS router lost its connection to the management LAN, it SHOULD actively establish adjacency with at least one of its IS-IS neighbors within the CLOS network.**
 - **As such, it could obtain the full LSDB of the CLOS network while flooding its self-originated LSPs to the remaining part of the whole network through that IS-IS neighbor.**

Solution Details

(con't)

- **The controller being elected as IS-IS DIS would send ISIS PDUs as multicasts or unicasts as before.**
- **It SHOULD accept and process those unicast IS-IS PDUs originated from IS-IS routers.**
- **Upon receiving any new LSP from a given IS-IS router, the controller being elected as DIS MUST flood it immediately to the management LAN for two purposes: 1) implicitly acknowledging the receipt of that LSP; 2) synchronizing that LSP to all the other IS-IS routers.**
- **Furthermore, to decrease the frequency of advertising CSNPs on the controller being elected as DIS, it's RECOMMENDED that IS-IS routers SHOULD send an explicit acknowledgement with a PSNP upon receiving a new LSP from the controller being elected as DIS.**

Next Steps

- **Comments?**