Position Paper: Multicast in NGI

Sponsor: C Wilkinson / Architecture / 2021-MAR-02
Contributors: M Brochu, M Mullins, C Wilkinson

Summary:

In the initial deployment of the NGI network, Layer 3 PIM-based multicast will no longer be supported; multicast using Layer 2 multipoint VPLS/EVPN will continue to be supported. Assuming there is a clear community driver, future L3 support can be added as additional software features are added to IOS-XR - this is likely in 2022.

Background:

At the current time, Layer 3 PIM-based multicast is supported on the Internet2 core through the base-instance. This feature is lightly used by a few legacy programs (EUMETSET and NOAA) and has been a matter of long-term community debate due to the complexity it adds to the network configurations, limitations to future network virtualization (lack of support in L3VPNs), and a general misunderstanding / lack-of-training prevalent in our community. While these discussions are on-going, clear consensus has proved to be elusive. Similar to what we observed historically with Openflow support, vendor support for generalized multicast features lags behind the development of other platform features; this limits options for implementation and can overly constrain the deployment of other services on the platform.

Options Considered:

IOS-XR as implemented in the Cisco 8200s has limited options for the support of PIM-based multicast.

The first option would be to continue to support multicast using protocol independent multicast (PIM), in the core backbone. Many operators have expressed a desire to remove protocols such as PIM in the backbone in order to run a “BGP Free Core”. This means all service traffic on the backbone be label switched. Running PIM/Legacy Multicast in the global table runs counter to this objective.

In line with running a “hollow core”, software reconfigurable network, Internet2 desires to modernize and streamline it’s routing configuration by moving the R&E network from the global table into a VRF. This philosophy is mutually exclusive with the PIM/Global option above.

There are mechanisms to deliver multicast in a VPN while allowing the core to remain PIM-Free. One mechanism that has been recently supported on the Cisco 8200 platform is BGP Multicast VPNs. However, this feature cannot take advantage of novel multicast solutions that Segment Routing offers, such as BIER. As such, Internet2 will have to enable legacy protocols on the

backbone to facilitate transport of the multicast sources and receivers, e.g multipoint LDP (mLDP).

BGP multicast VPNs present a potential solution to the layer3 multicast problem, but one needs to consider if the complexity and risk associated with the solution outweigh the benefits. Consider the following:

1.) It is unclear how many real customers are implementing BGP MVPNs on the 8200 platform. Internet2 may be venturing into lightly-treaded territory.
2.) BGP MVPN involves the addition of several new protocols which introduce additional state into the control plane of the network. For example, PIM joins will now be propagated across the BGP backbone, and mLDP will need to be initialized between all interfaces. The addition of this state increases the overall risk of NGI, especially due to the amount of new protocols and software being deployed using IOS-XR.
3.) To properly support BGP MVPN, additional engineering will be required. The current timelines for deployment are stretching engineering resources. There is a high likelihood that if a solution is engineered, there will be minimal staff trained to troubleshoot issues. Adding to this, since L3 multicast is such a lightly/mostly unused service, there are not many opportunities for real experience with the protocols.

Plan of Record:

Q2 2021 initial deployment:

We will message to the community that Layer 3 PIM-based multicast will no longer be supported when the network is migrated to the NGI platform. With the current project timelines, these migrations are planned to take place in May-June.

EUMETsat has been identified as being potentially impacted and solution development will also be pursued. ASM is still required by EUMETSAT unless an NREN is unable to support it. In these cases, GÉANT provides GRE tunnels to the end site specified by EUMETSAT. GÉANT is currently exploring AMT with EUMETSAT as an evolution to the existing multicast setup. However, it has not yet become a production service.

Q1 2022 planned improvement:

The NTAC is planning to form a multicast working group, which we can leverage to inform decisions on future multicast implementations after the core network migration is complete.