



Simple join failure notification for PIM-SM multicast routing

draft-hoerdt-pim-group-unreachable-00

Stig Venaas
venaas@uninett.no




Introduction




- | Multicast debugging is hard, especially interdomain
- | When someone is not able to receive, it is difficult to identify where and what the problem is
 - | For interdomain it's next to impossible to track down which domain is at fault
- | The idea in this draft, is to send a message downstream if join propagation fails
- | This mechanism is simpler than a multicast traceroute, but still very useful

Proposal




- | I will focus on SSM/SPT, but draft also considers ASM/RPT
- | If a router fails to send an (S,G)-join, it sends an ICMP error message downstream, which is propagated hop-by-hop
 - | Sent out on all interfaces in oif list to the all-nodes multicast address
 - | The next hops will accept it if it arrives on the incoming interface for S
 - | The next hops will resend it on all interface in oif list
- | Message contains information on why the join failed, and the address of the router that failed to send it

Example error codes




- | Interface pointing towards S is not multicast enabled
- | Next-hop is not a PIM neighbor
- | No route for source
- | Exceeding admin scope
- | Filtered
- | For ASM, the RP-address is wrong

How this can be used



- | Routers might log reception of such messages
- | Someone that fails to receive could look for logged errors on the router, or report to provider
 - | Provider can check logs and find out where the error is
 - | The entity responsible for the router at fault can be contacted
- | A common problem is that there is no state to diagnose problem when user complaint is received
 - | Having logs with timestamps really helps here
- | Also, user might run an application that can help diagnose multicast problems
 - | This application might report errors to users based on the received ICMP messages
 - | General multicast applications might also utilise this

Further details



- | A single message might contain reports for multiple join failures
 - | multiple (S,G)s etc.
- | These are split up as needed when being forwarded downstream
 - | For a given interface we only send message with reports for the (S,G)-entries where it is in the oif list
- | Messages should be rate-limited