

93rd IETF, July 2015, Prague, Czech Republic

Multiple Upstream Interface Support for IGMP/MLD Proxy

draft-asaeda-pim-multiif-igmpmldproxy-01

Hitoshi Asaeda (NICT)

Eiichi Muramoto (Panasonic)


Luis M. Contreras (Telefonica)

Document History

- Multiple Upstream Interfaces Support for IGMP/MLD Proxy
 - draft-asaeda-pim-mldproxy-multif-00, Oct. 15, 2012
 - draft-asaeda-pim-mldproxy-multif-01, Feb. 25, 2013
 - draft-asaeda-pim-multiif-igmpmldproxy-00, Mar. 23, 2015
 - draft-asaeda-pim-multiif-igmpmldproxy-01, Jul. 6, 2015

Background

- There are many situations an IGMP/MLD proxy multiply attached to same or different networks (e.g., Internet and Intranet) or different interfaces (e.g., ethernet and wireless link), yet RFC4605 does not support such multihoming situations.
- Enable an IGMP/MLD proxy device to use multiple upstream interfaces and receive multicast packets through these interfaces.



Objective

- Support multiple upstream interfaces for an IGMP/MLD proxy device
 - An IGMP/MLD proxy device enables to receive multicast sessions/channels through the different upstream interfaces
- Follow the requirement draft recently published

Benefits

- Load balancing
 - Subscriber-based upstream selection: One or more upstream interface(s) is selected per subscriber/receiver
 - Channel-based upstream selection: One or more upstream interface(s) is selected per channel/session
- Robust data reception
 - More than one upstream interface used per channel/session when more than one upstream interface is enabled for the channel/session
- Upstream interface takeover
 - Switch inactive upstream IF to other active (backup) IF
- Each function can be enabled/disabled by configuration

Candidate Upstream Interface Configuration

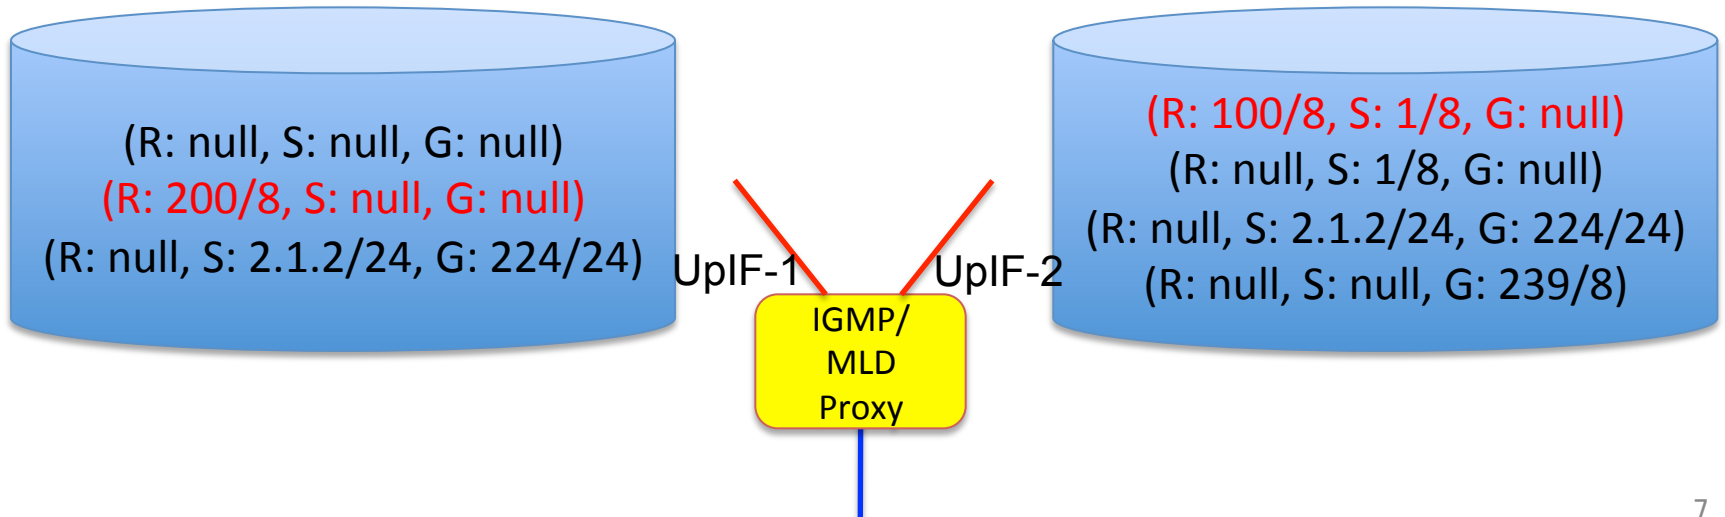
- Parameters for candidate upstream interface configuration
 - Subscriber address prefix
 - Channel/session ID
 - Source address prefix and multicast address prefix
 - Priority (fixed) -> Credit (variable)
 - Backup interface(s)
- Configuration syntax
 - (R: subscriber-addr-prefix, S: source-addr-prefix, G: multicast-addr-prefix) (C: value) (B: IF-name)
 - Default: (null, null, null) (0) (null)
- Decision order
 - Subscriber prefix > Channel ID > Credit > Lowest IP address



Called "address prefix record"

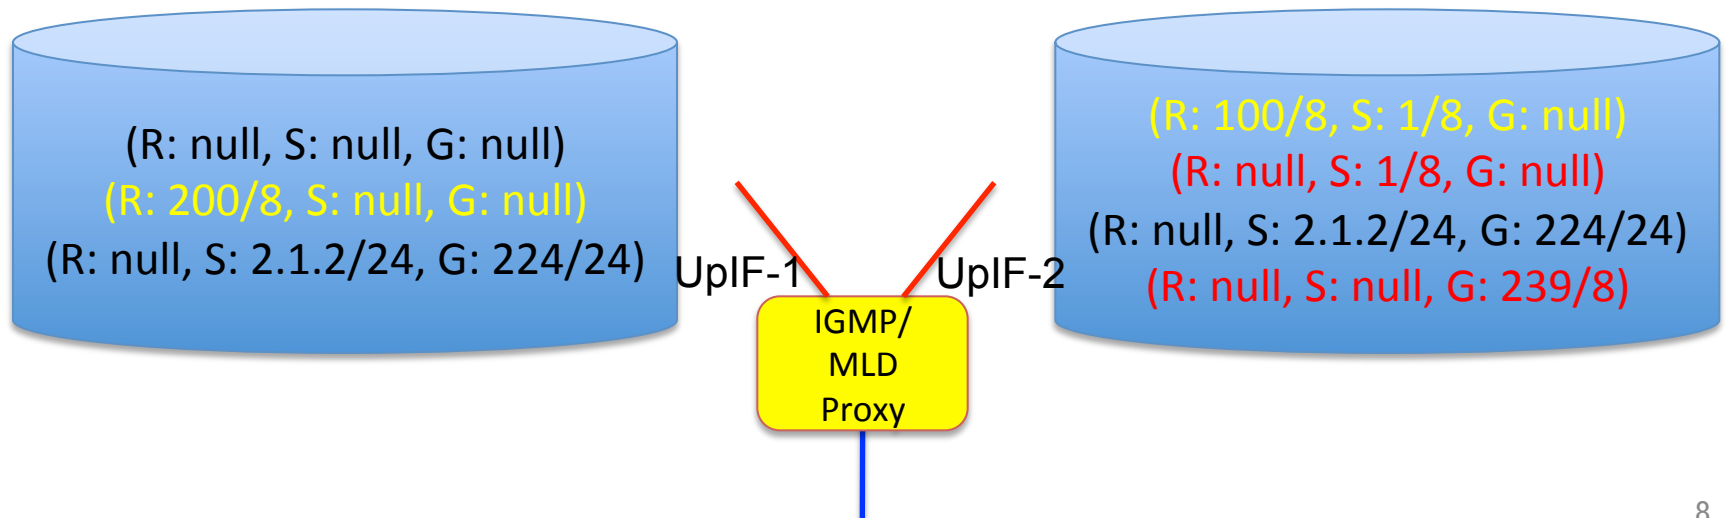
Subscriber-Based Upstream Selection

- Subscriber address prefix
 - Default values of subscriber address prefix is null value
 - Configuration with null subscriber addr. prefix is “channel-based upstream selection”



Channel/Session ID-Based Upstream Selection

- Channel/session ID: Source address prefix and multicast address prefix
 - Source address prefix takes priority over multicast address prefix
- Longest match
 - Longest prefix configuration is prioritized

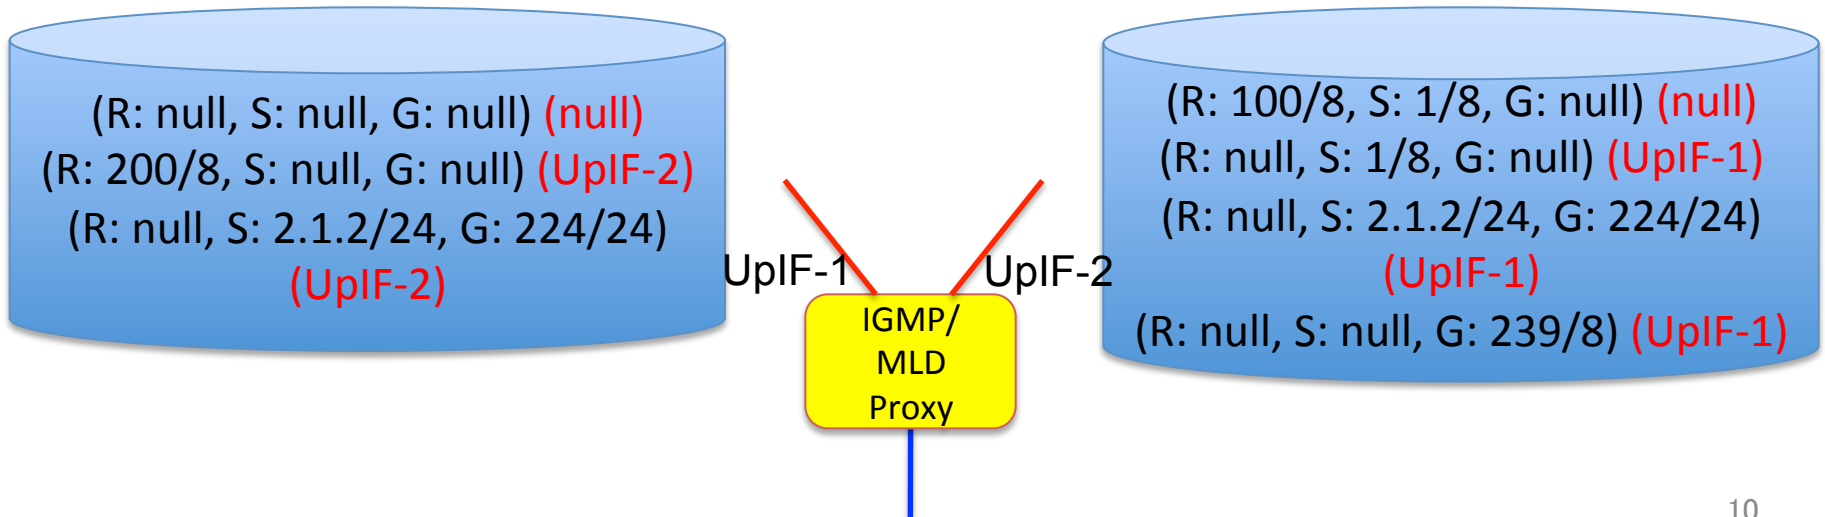


Default Interface

- The default of “address prefixes” is “(null, null, null)”
- The default of “credit” is (0)
- The default of “backup interface” is “null”
- When all values are default for all candidate upstream interfaces, the configured upstream interface having lowest IP address is selected as the upstream interface for all multicast channels

Backup Interface Configuration

- Backup interface can be configured per address prefix record
 - If no backup interface configuration for an address prefix record, no interface takeover happens



Upstream Interface Takeover


- Activate a “backup” upstream interface;
 - When a “selected” upstream interface or a network the selected upstream attaches goes down
- Open issue
 - How to quickly detect “inactive” interfaces?
 - Monitor IGMP/MLD Query and/or PIM Hello?
 - Not fast
 - Explicit tracking make the query period longer
 - Define a new IGMP/MLD message?

~~Priority~~ Credit

- Enable interface takeover due to continuous unreliability
 - When credit becomes 0, candidate interface having the highest credit next to the primary one becomes upstream IF
 - Specification of the action when (in what kind of condition) credit is subtracted, added (because of recovery), or reset is left to the implementation
- Each address prefix record can have own credit value
 - Default value is 0 (i.e. credit is not used)
- Credit and backup interface configurations are mutually exclusive
 - If credit is configured (i.e. not defaults), backup interface configuration is ignored for that interface


Multiple Upstream Interface Selection for Robust Data Reception

- If both an address prefix record (subscriber addr prefix, channel/session ID) and credit are identical for multiple candidate upstream IFs, these interfaces are selected for the specified addr prefix (and hence duplicate packets are transmitted via these IFs)
 - This does not guarantee that the packets come from disjoint paths.



How to Use Multiple IFs for All Channels?

- Multiple IFs having all default (i.e. (null,null,null)(0) (null)) do not enable redundancy
 - One of these IFs having lowest IP addr becomes upstream IF
- Each configured upstream interface has the default prefix but must have identical non-default credit value



Conclusion

- Load balancing and robust data reception
- Upstream interface takeover
- Configuration for each candidate upstream interface
 - Address prefix record, credit value, backup interface(s)
 - (R: subscriber-addr-prefix, S: source-addr-prefix, G: multicast-addr-prefix)(C: value)(B: IF-name)
- Open issues
 - How to quickly detect “inactive” interfaces?
 - Monitoring IGMP/MLD Query and/or PIM Hello?
 - New IGMP/MLD message?
 - Credit is useful configuration option?
 - Automatic upstream interface configuration is necessary?
- Comments always welcome