

More Resilient GMP

draft-anggawijaya-pim-resilient-gmp-00

IETF-94 Yokohama

Hermin Anggawijaya

Allied Telesis Labs, NZ

GMP – Quick Recap

- Unsolicited membership reports sent from listeners wanting to receive multicast data
- GMP router forwards multicast data, based on group membership
- GMP router (querier) – sends periodic query to see if listeners still want the multicast data
- Solicited reports are sent by all listeners in response to the query

Lost GMP Packets

- Less reliable link layer
 - WiFi, issues identified in:
 - [draft-mcbride-mboned-wifi-mcast-problem-statement]
 - [draft-vyncke-6man-mcast-not-efficient]
 - Other wireless link technologies
- More likely on large L2 broadcast domains spanning multiple devices:
 - Temporary congestions on certain links
 - Transient states on devices – device restarting, temporary CPU congestions (during L2 topology change)

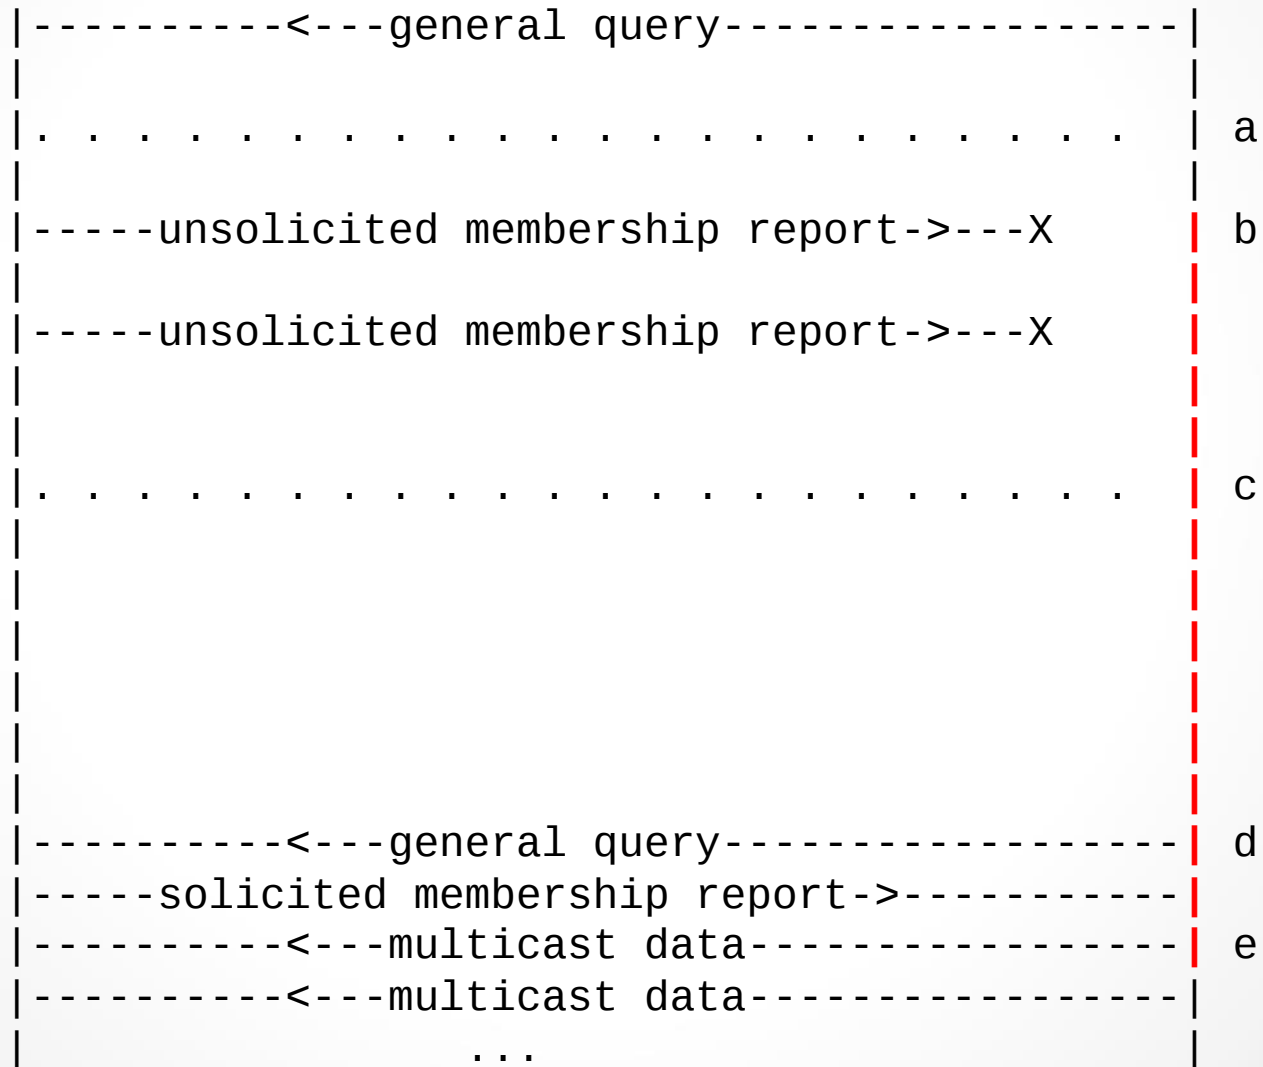
Problem Statement

- GMPs membership reports rely on retransmissions of packets
 - (Un)solicited membership reports are retransmitted:
 - robustness variable (RV) – default 2
 - random interval $(0, [\text{Unsolicited Report Interval}])$ – default $(0, 1)$
 - All listeners do retransmission
- **If the ‘RV’ number of reports are lost, ‘dead period’ may result**
- Consider the following scenario....

'Dead Period' in GMP

Multicast listener

GMP Querier



T
i
m
e
|
|
V

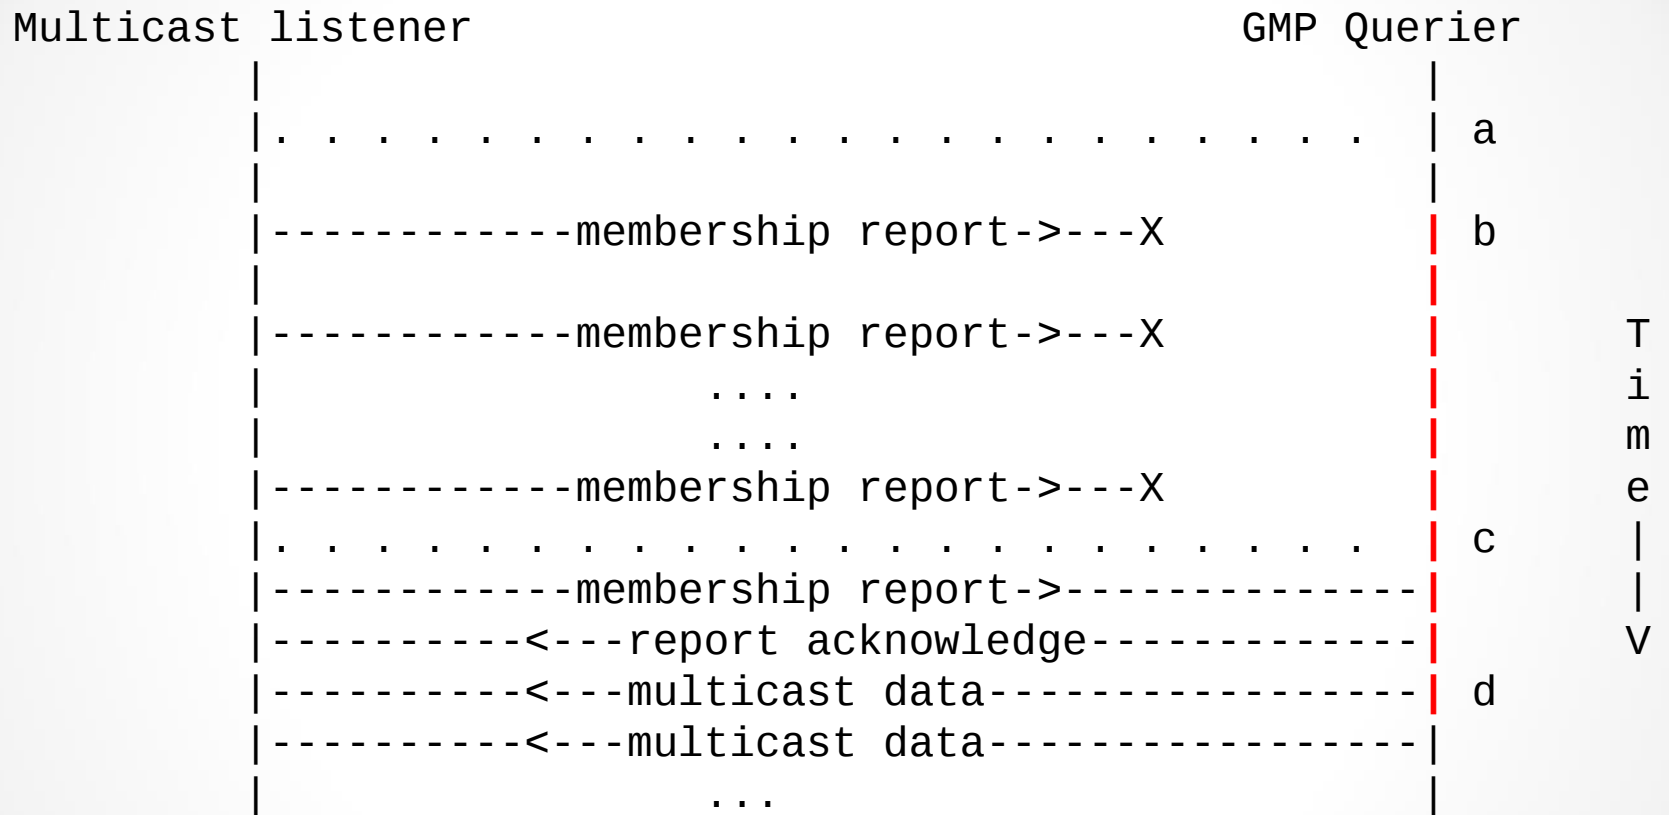
Mitigation Strategies

- Leave the responsibility to app layer to reset the socket, but better to do it in network layer
- Increase the robustness variable – increase chattiness for all listeners
- Decrease the query interval – also increase chattiness for all listeners
- Longer group membership timeout – applicable only to existing/known listeners only. This requires incrementing rv or query interval or QRI


Proposed Solution

- Listeners retransmit membership reports beyond the robustness variable – *maybe disabled administratively*
- Only stops when router acknowledge the reports or when no router is detected (2.5 x Query Interval)
- Acknowledgement is sent unicast to listener - reflected membership report message
- Router announce capability to acknowledge reports in query packets

Proposed Solution – cont.



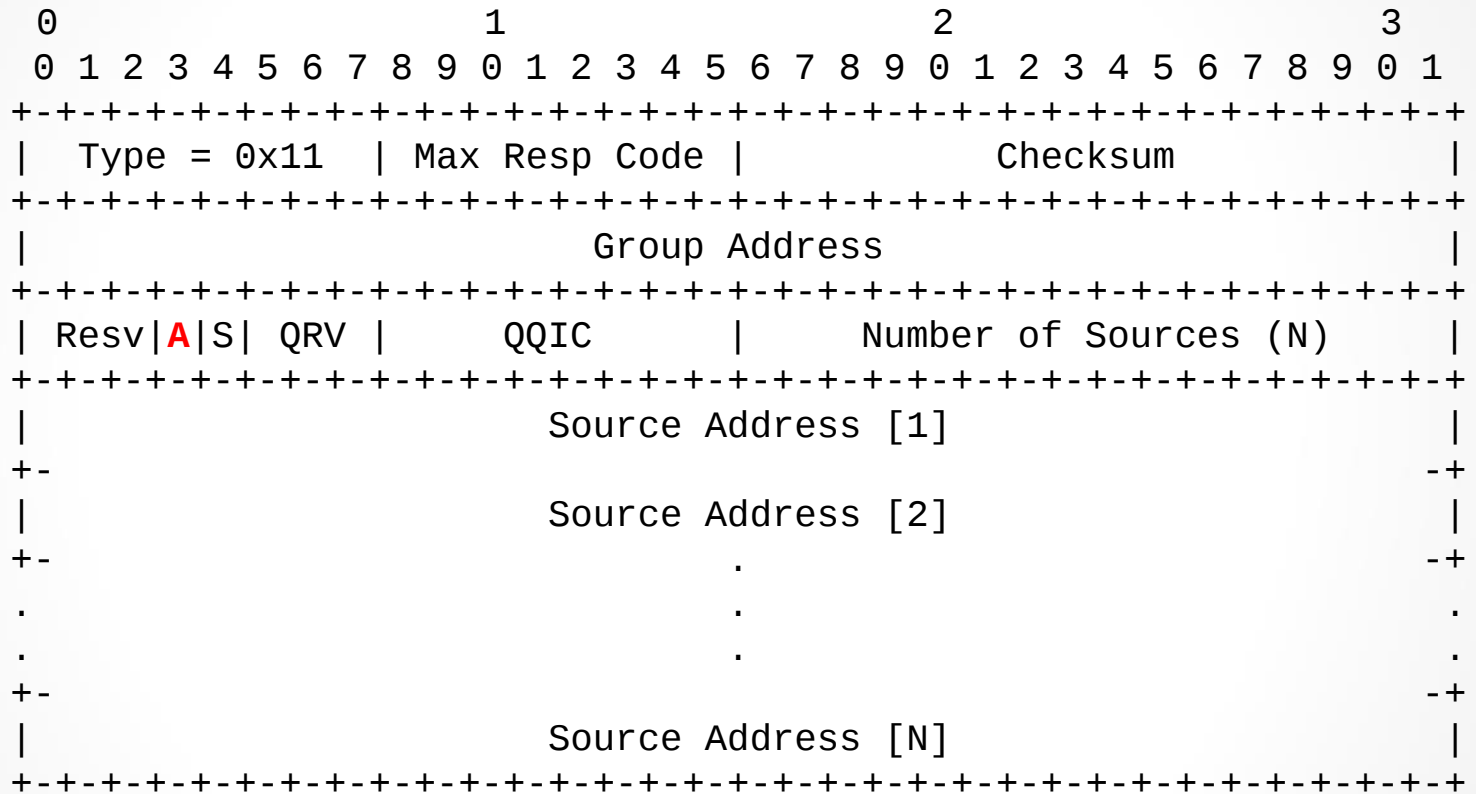
Resilient IGMP - Report



Resilient IGMP Report

- R: 'Request for Acknowledgement' flag – requires an acknowledgement in response
 - Only set if querier indicate it can ack reports
- A: 'Membership report acknowledgement' flag – set on reflected report by router as acknowledgement to the original report

Resilient IGMP Query



Resilient IGMP Query

- A : 'Acknowledgement Capable Querier' flag - indicates that the querier is capable of sending acknowledgements to multicast listeners.
- Querier only sends query with A-flag set *If enabled administratively*
- Listeners may ignore the A-flag from querier

Resilient MLD Report

MLDv2 Membership Report

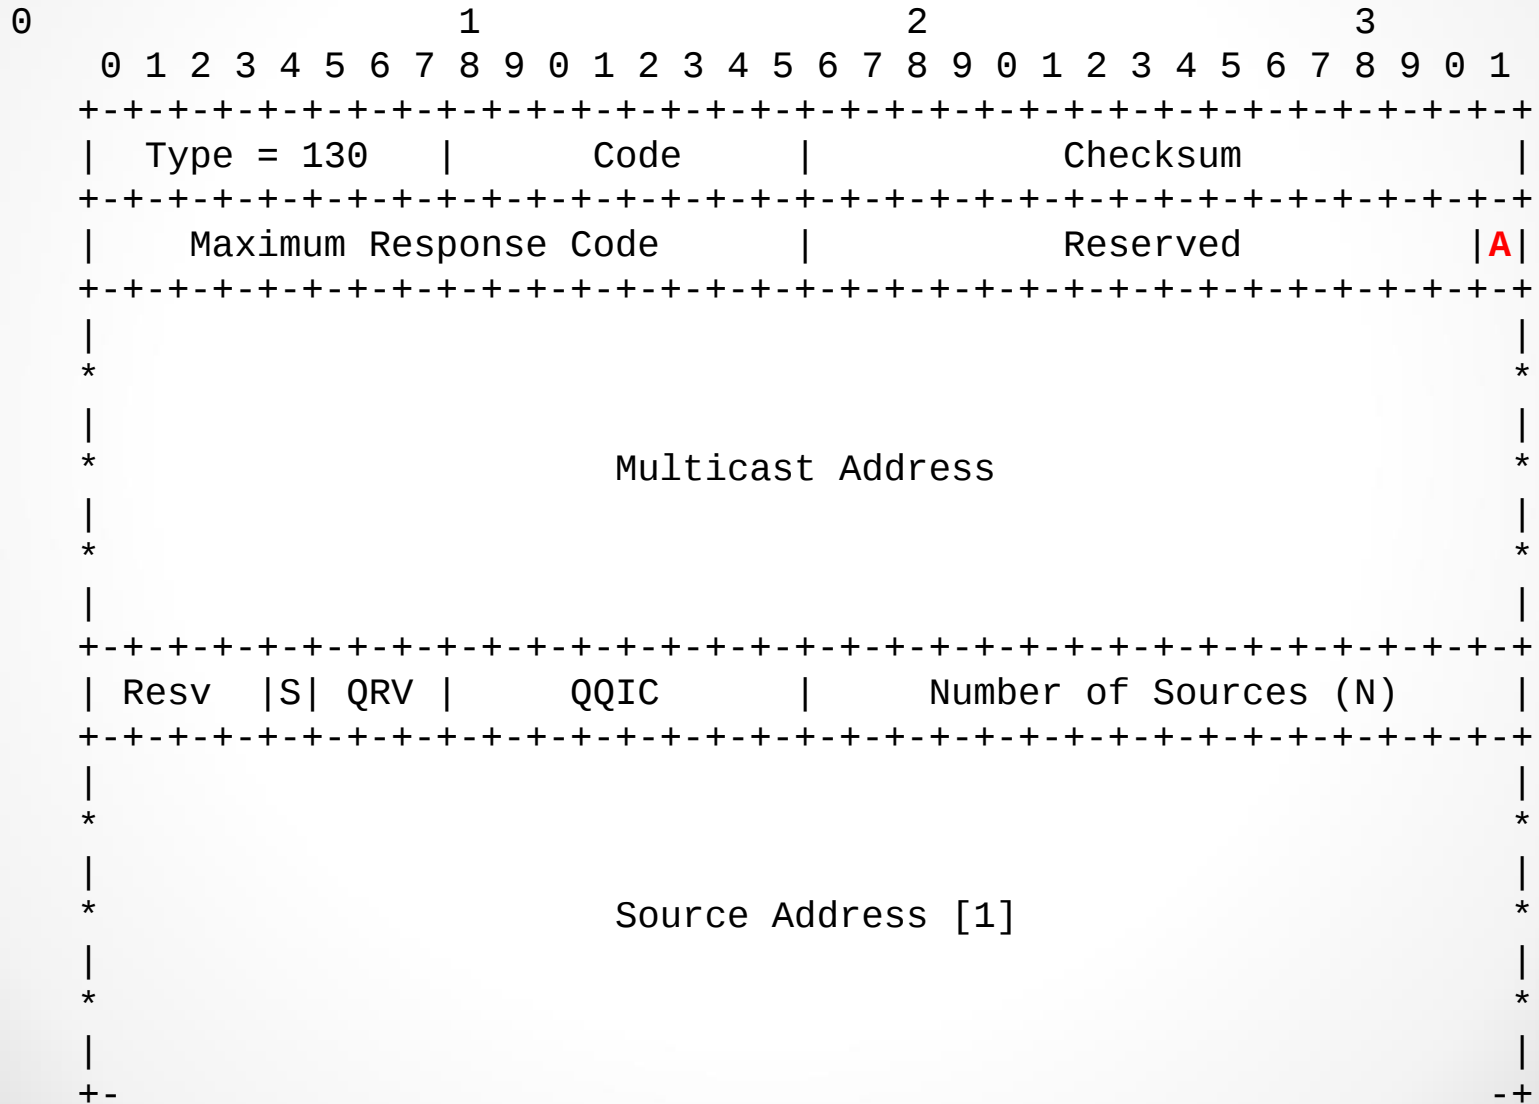
```

      0                               1                               2                               3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|  Type = 143   |R|A|  Reserved |                               Checksum |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                               Reserved |Nr of Mcast Address Records (M)|
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|
:
.                               Multicast Address Record [1]                               .
:
|
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|
:
.                               Multicast Address Record [2]                               .
:
|
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|
:
.
|
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|
:
.
|
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|
:
.
```

Resilient MLD Report

- R: 'Request for Acknowledgement' flag – requires an acknowledgement in response
- A: 'Membership report acknowledgement' flag – set on reflected report packet by router as acknowledgement to the original report

Resilient MLD Query



Resilient MLD Query

- A : 'Acknowledgement Capable Querier' flag - indicates that the querier is capable of sending acknowledgements to multicast listeners
- Querier only sends query with A-flag set *If enabled administratively*
- Listeners may ignore the A-flag from querier

Backwards Compatibility

- Listeners fallback to current behaviour when receiving query w/o A-flag set
- Router fallback to current behaviour when receiving query w/o A-flag set from other querier
- Also imply that when IGMPv2, or IGMPv1 are present, all fallback to current behaviour
- Similar behaviour for MLDv2

Observations

- Shorter 'dead period' – until router receives the next report (vs. solicited report received after the next general query)
- Not chatty unnecessarily – retransmissions only done by yet to be acknowledged listeners (vs. all listeners)
- *Some implementations processes unicast MLDv2 and IGMPv3 report – may need to change type for acknowledgement message?*

Operational Considerations

- The proposed mechanism is more effective when the querier is also the DR
 - Most of simple deployments have the querier on the same device as DR
- Does not work w/ addressless snooper
- Suitable for sparse L2 network, or more capable router needed for dense L2 network
- Security
 - Fake reports w/ R-flag set
 - Fake queries w/ A-flag set

Next Steps

- Questions
- Suggestions and Comments

Thank You

hermin.anggawijaya@alliedtelesis.co.nz

Allied Telesis Labs, NZ