Secure Multicast Routing Infrastructure: The Network Operator's Viewpoint

Zainab KHALLOUF

Thesis prepared at
France Télécom R&D and INRIA Rhône Alpes, Planete project-team

Advisors:

Prof.	Andrzej DUDA	Thesis Director (LSR-IMAG)
Dr.	Vincent ROCA	Supervisor (INRIA Rhône Alpes)
Eng.	Sébastien LOYE	Supervisor (France Télécom R&D)

The Multicast Model (RFC-1112)

Multicast : the sender send the same packet to a set of hosts identified by a single address.

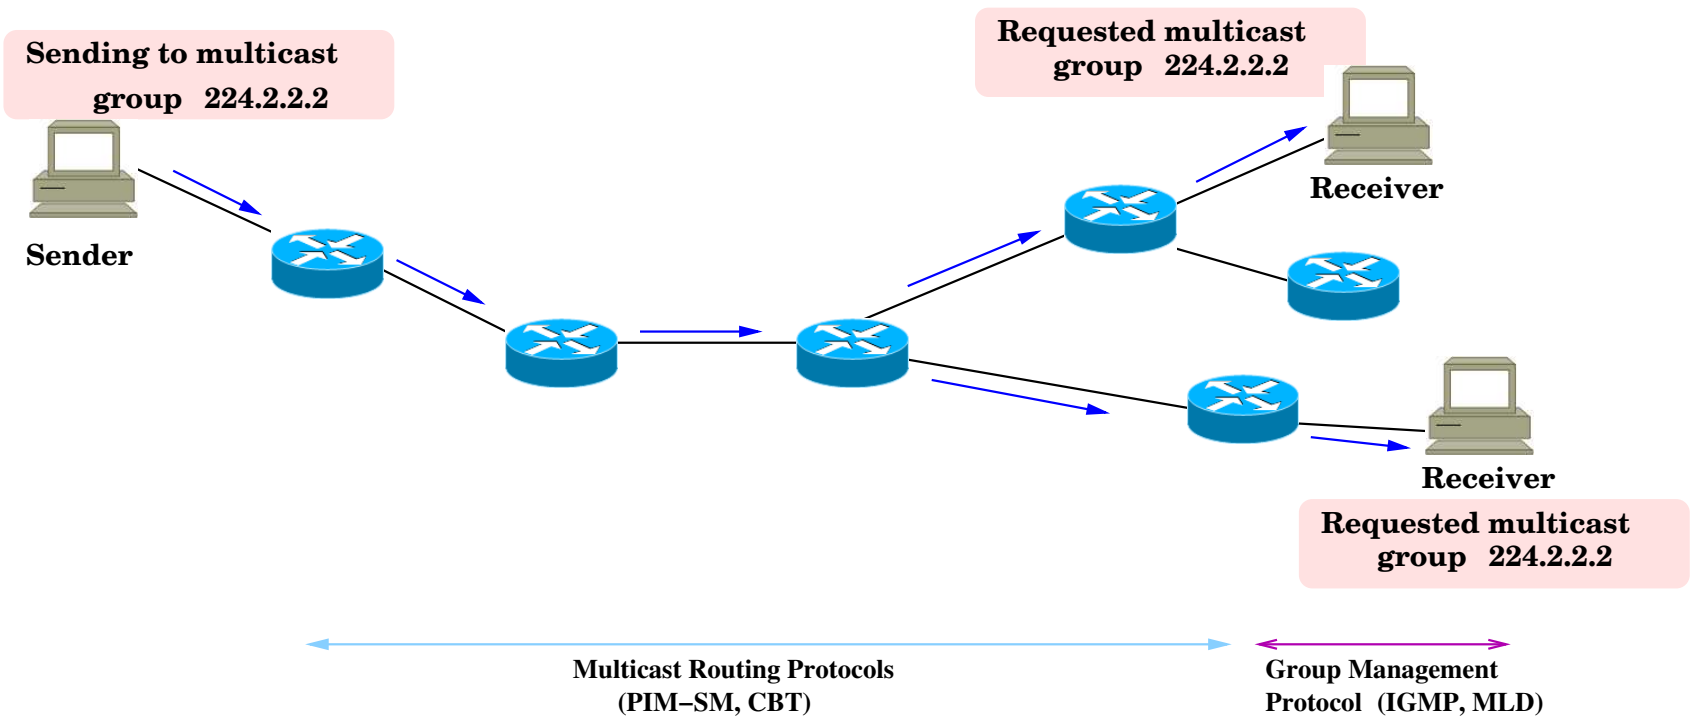

Why This Work?

The security issues preclude the deployment of multicast in the operators networks.

The multicast routing infrastructure is highly vulnerable to attacks launched intentionally (by bad guys) or not.

Many theoretically ideal proposals but rarely accepted by the operators community.

$\Rightarrow$ It's time to listen to the operator....

Why This Work?

This work aims to:

Understand the operator's requirements in terms of security (largely different from that of end-users or content providers).

Analyze the multicast routing infrastructure security threats from the network operator's standpoint.

Analyze the existing security propositions and understand why they have not been implemented in the operational network.

Propose realistic security mechanisms that can response to the operator requirements.

Outline

1. Part I: Who is the Network Operator?
2. Part II: Multicast Routing Infrastructure Security from the Network Operator Viewpoint: State of the Art
3. Part III: Our Filtering Proposal
4. Part IV: Experimental Evaluation
5. Part V: Discussion
6. Part VI: Conclusion

Part I

Who is the Network Operator?

The Various Actors

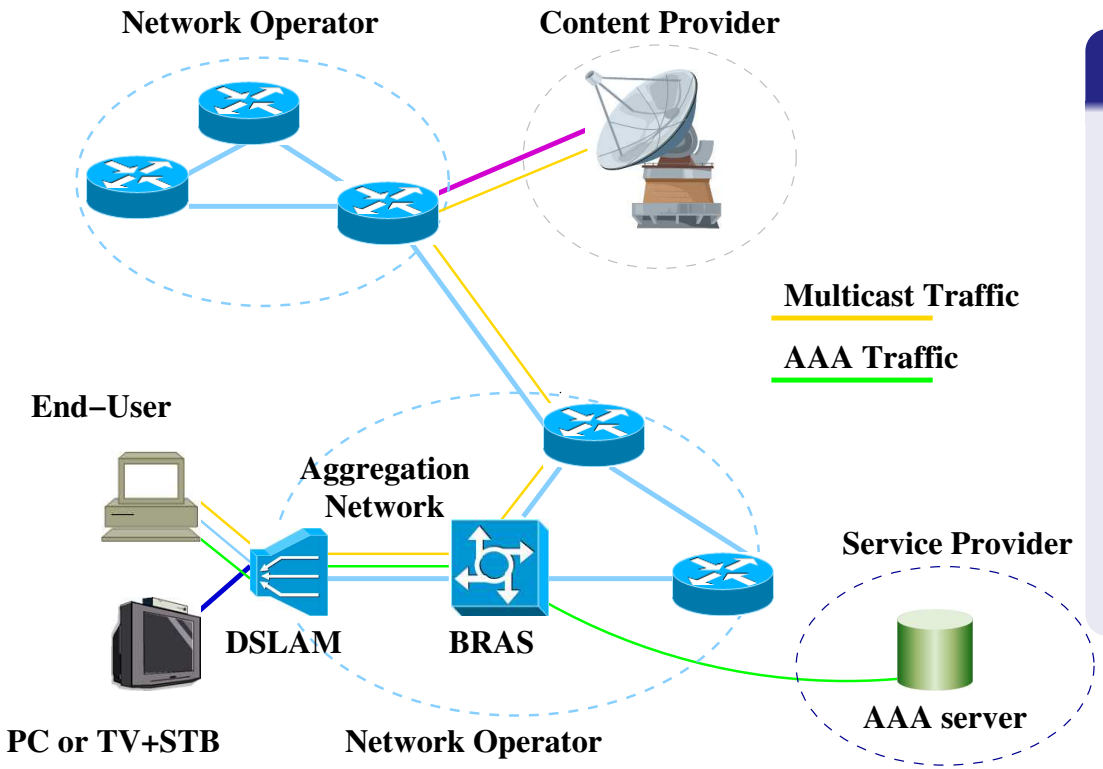

The Actors

1. The Network Operator (Carrier)
2. The Internet Service Provider (ISP)
3. The Content Provider
4. The End User

$\Rightarrow$ The network operator manages the routing infrastructure

The Multicast Deployment in the Operator Network

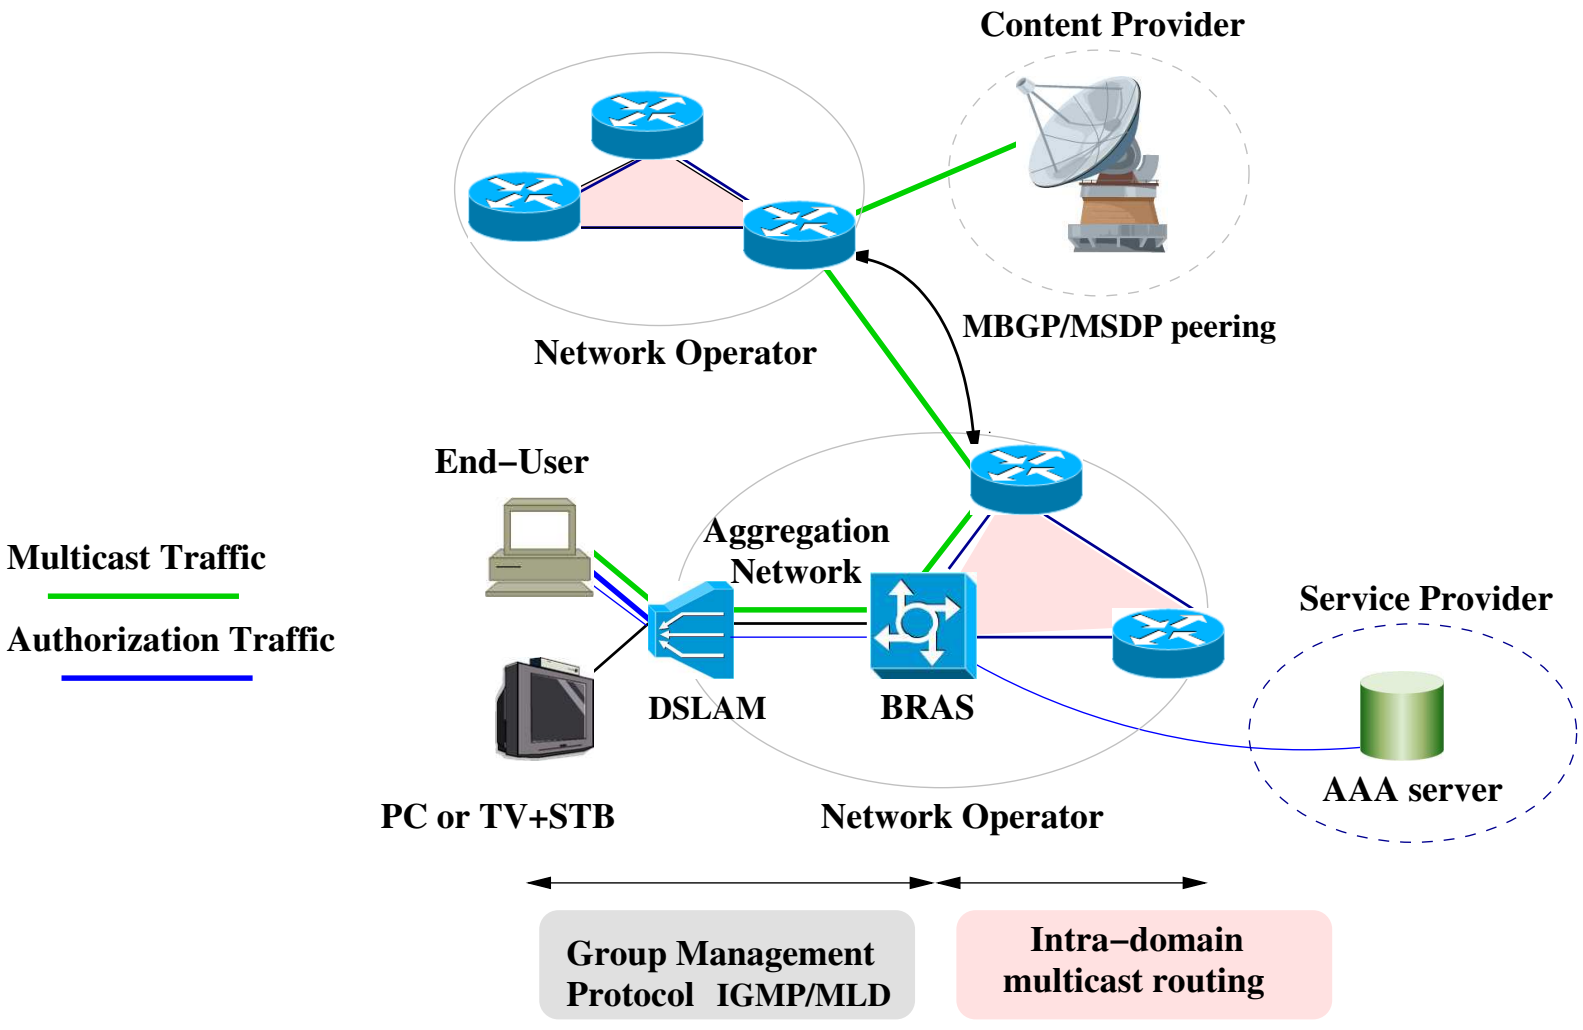

Part II

Multicast Routing Infrastructure Security: the Network Operator Viewpoint

A Taxonomy of Possible Attacks of Interest to the Network Operator

Their origin:

Internal attacks

Mounted from within the multicast distribution tree of a network operator.

Edge attacks

Mounted from the edge of the operator network.

- Senders attacks.
- Receivers attacks.

A Taxonomy of Possible Attacks of Interest to the Network Operator (cont.)

Their type:

Data plane attacks

Disturb the data forwarding functions in the routers.

- Flooding the router with a large amount of multicast traffic.

Control plane attacks

Disturb the signaling functions in the router.

- Sending a forged routing messages to change the forwarding table.
- Increasing the amount of multicast states information in routers above a manageable level.

Real Attack Example

Ramen Worm (January 2001): triggered an MSDP attack

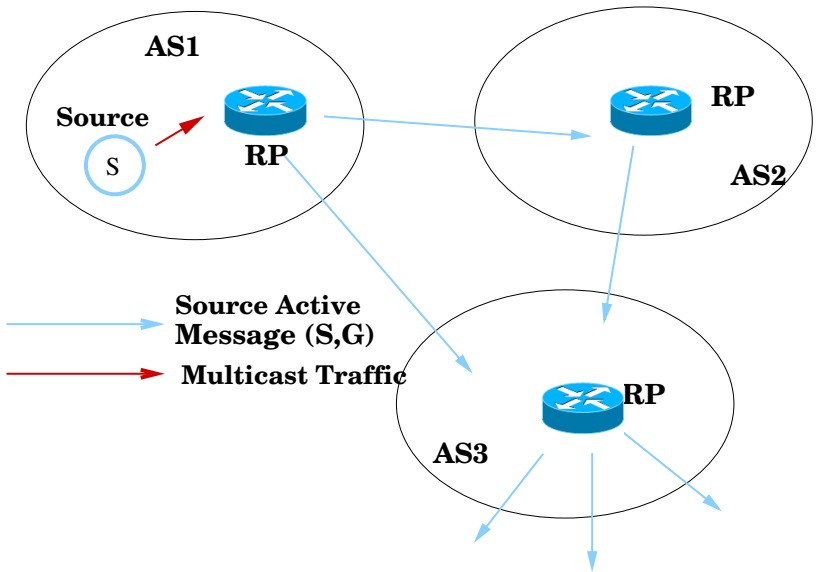

Real Attack Example: Ramen Worm (January 2001) (cont.)

- The story begun by accident: port-scan a random set of IP addresses by sending ICMP messages.
- A portion of these addresses were multicast addresses.

Real Attack Example: Ramen Worm (January 2001) (cont.)

- The story begun by accident: port-scan a random set of IP addresses by sending ICMP messages.
- A portion of these addresses were multicast addresses.

Real Attack Example: Ramen Worm (January 2001) (cont.)

- The story begun by accident: port-scan a random set of IP addresses by sending ICMP messages.
- A portion of these addresses were multicast addresses.

Real Attack Example: Ramen Attack (January 2001) (cont.)

- The number of generated SAs (Source Active) ranges from 10000 to 45000 /minute
- Melt down the routers.

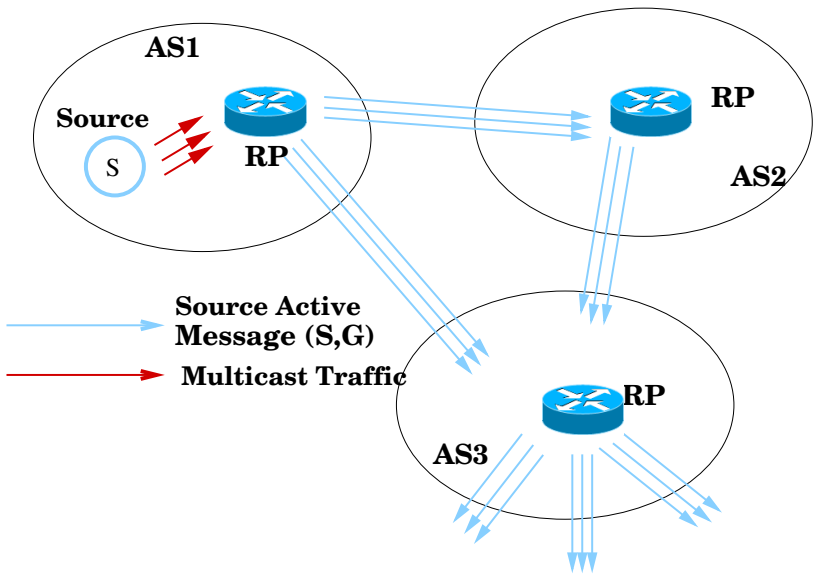

Discussion of Possible Attacks

- Edge attacks in particular DoS attacks
  Easily launched and they are not simple to counter.
  $\Rightarrow$ High probability.
- Congestion control attacks
  Most multicast-enabled applications can trigger this attack, either intentionally or not.
- Core attacks
  Require generally a strategically placed intruder.
  $\Rightarrow$ Small probability.

Discussion of Possible Attacks

- Edge attacks in particular DoS attacks
  Easily launched and they are not simple to counter.
  $\Rightarrow$ High probability.
- Congestion control attacks
  Most multicast-enabled applications can trigger this attack, either intentionally or not.
- Core attacks
  Require generally a strategically placed intruder.
  $\Rightarrow$ Small probability.

Discussion of Possible Attacks

- Edge attacks in particular DoS attacks
  Easily launched and they are not simple to counter.
  $\Rightarrow$ High probability.
- Congestion control attacks
  Most multicast-enabled applications can trigger this attack, either intentionally or not.
- Core attacks
  Require generally a strategically placed intruder.
  $\Rightarrow$ Small probability.

Discussion of Possible Attacks

- Edge attacks in particular DoS attacks
  Easily launched and they are not simple to counter.
  ⇒ High probability.
- Congestion control attacks
  Most multicast-enabled applications can trigger this attack, either intentionally or not.
- Core attacks
  Require generally a strategically placed intruder.
  ⇒ Small probability.

A Taxonomy of Existing Mechanisms to Secure the Multicast Routing Infrastructure

Three categories:

Attack Avoidance Approaches (Preventive)

Control the ability of entities (routers/receivers/senders) to take part in the multicast routing tree for a given group.

Attack Resiliency Approaches (Reactive)

Detect and mitigate the effects of attacks.

Hybrid Approaches

Prevention + Reaction.

The Limitations of Existing Mechanisms

- Participant authentication and authorization is not always feasible.
  - It is not feasible in case of free content delivery services.
  - It is only feasible if a network operator/content provider agreement exists.
  - Nothing guarantees that an authenticated/authorized client will behave correctly.
- Modifying existing protocols
  - Operators are reluctant to change their infrastructure.

The Limitations of Existing Mechanisms

- Participant authentication and authorization is not always feasible.
  - It is not feasible in case of free content delivery services.
  - It is only feasible if a network operator/content provider agreement exists.
  - Nothing guarantees that an authenticated/authorized client will behave correctly.
- Modifying existing protocols
  - Operators are reluctant to change their infrastructure.

The Limitations of Existing Mechanisms

- Participant authentication and authorization is not always feasible.
  - It is not feasible in case of free content delivery services.
  - It is only feasible if a network operator/content provider agreement exists.
  - Nothing guarantees that an authenticated/authorized client will behave correctly.
- Modifying existing protocols
  - Operators are reluctant to change their infrastructure.

The Limitations of Existing Mechanisms (cont.)

- Inter-domain confidence
  - Solutions that require collaborations between different operators are almost impossible.
    $\Rightarrow$ Limited (No) trust in other operators.
- Simple filtering does not provide differentiation between good or ill-behaved clients.

The Limitations of Existing Mechanisms (cont.)

- Inter-domain confidence
  - Solutions that require collaborations between different operators are almost impossible.
    $\Rightarrow$ Limited (No) trust in other operators.
- Simple filtering does not provide differentiation between good or ill-behaved clients.

The Limitations of Existing Mechanisms (cont.)

- Inter-domain confidence
  - Solutions that require collaborations between different operators are almost impossible.
    $\Rightarrow$ Limited (No) trust in other operators.
- Simple filtering does not provide differentiation between good or ill-behaved clients.

The Security from the Network Operator Viewpoint

Survivability

The group communication service provided to its clients (i.e. end users or other network operators with whom he has peering relationships) must **be operational** at any time.

The security is not the goal, but a mean to achieve the network operator's "continuation of service no matter what happens" goal.

The Security from the Network Operator Viewpoint

Survivability

The group communication service provided to its clients (i.e. end users or other network operators with whom he has peering relationships) must **be operational** at any time.

The security is not the goal, but a mean to achieve the network operator's "continuation of service no matter what happens" goal.

The Security from the Network Operator Viewpoint

The priority is for the edge attacks.

The operator wants realistic and simple solutions.

In This Work We aim to:

Mitigate edge attacks in particular group management flooding attacks.

Avoid the limitations of the existing proposals.

- A cost effective, scalable and transparent mechanism.
- Do not rely on authorization and the authentication of clients.
- Intelligent solution: protect well-behaved clients against ill-behaved ones.

Our main goal is "continuation of service no matter what happens".

In This Work We aim to:

Mitigate edge attacks in particular group management flooding attacks.

Avoid the limitations of the existing proposals.

- A cost effective, scalable and transparent mechanism.
- Do not rely on authorization and the authentication of clients.
- Intelligent solution: protect well-behaved clients against ill-behaved ones.

Our main goal is "continuation of service no matter what happens".

In This Work We aim to:

Mitigate edge attacks in particular group management flooding attacks.

Avoid the limitations of the existing proposals.

- A cost effective, scalable and transparent mechanism.
- Do not rely on authorization and the authentication of clients.
- Intelligent solution: protect well-behaved clients against ill-behaved ones.

Our main goal is "continuation of service no matter what happens".

In This Work We aim to:

Mitigate edge attacks in particular group management flooding attacks.

Avoid the limitations of the existing proposals.

- A cost effective, scalable and transparent mechanism.
- Do not rely on authorization and the authentication of clients.
- Intelligent solution: protect well-behaved clients against ill-behaved ones.

Our main goal is "continuation of service no matter what happens".

Part III

Our Filtering Proposal

Our Filtering Proposal

Intelligent filtering approach to thwart some DoS attacks that are based on IGMP or MLD flooding.

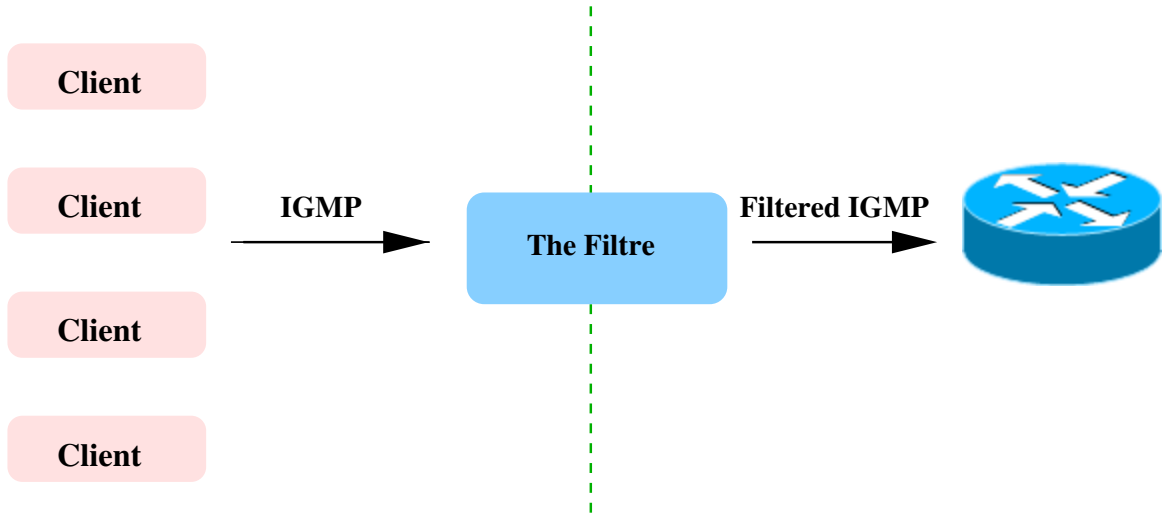

Our Filtering Proposal: Basic Architecture

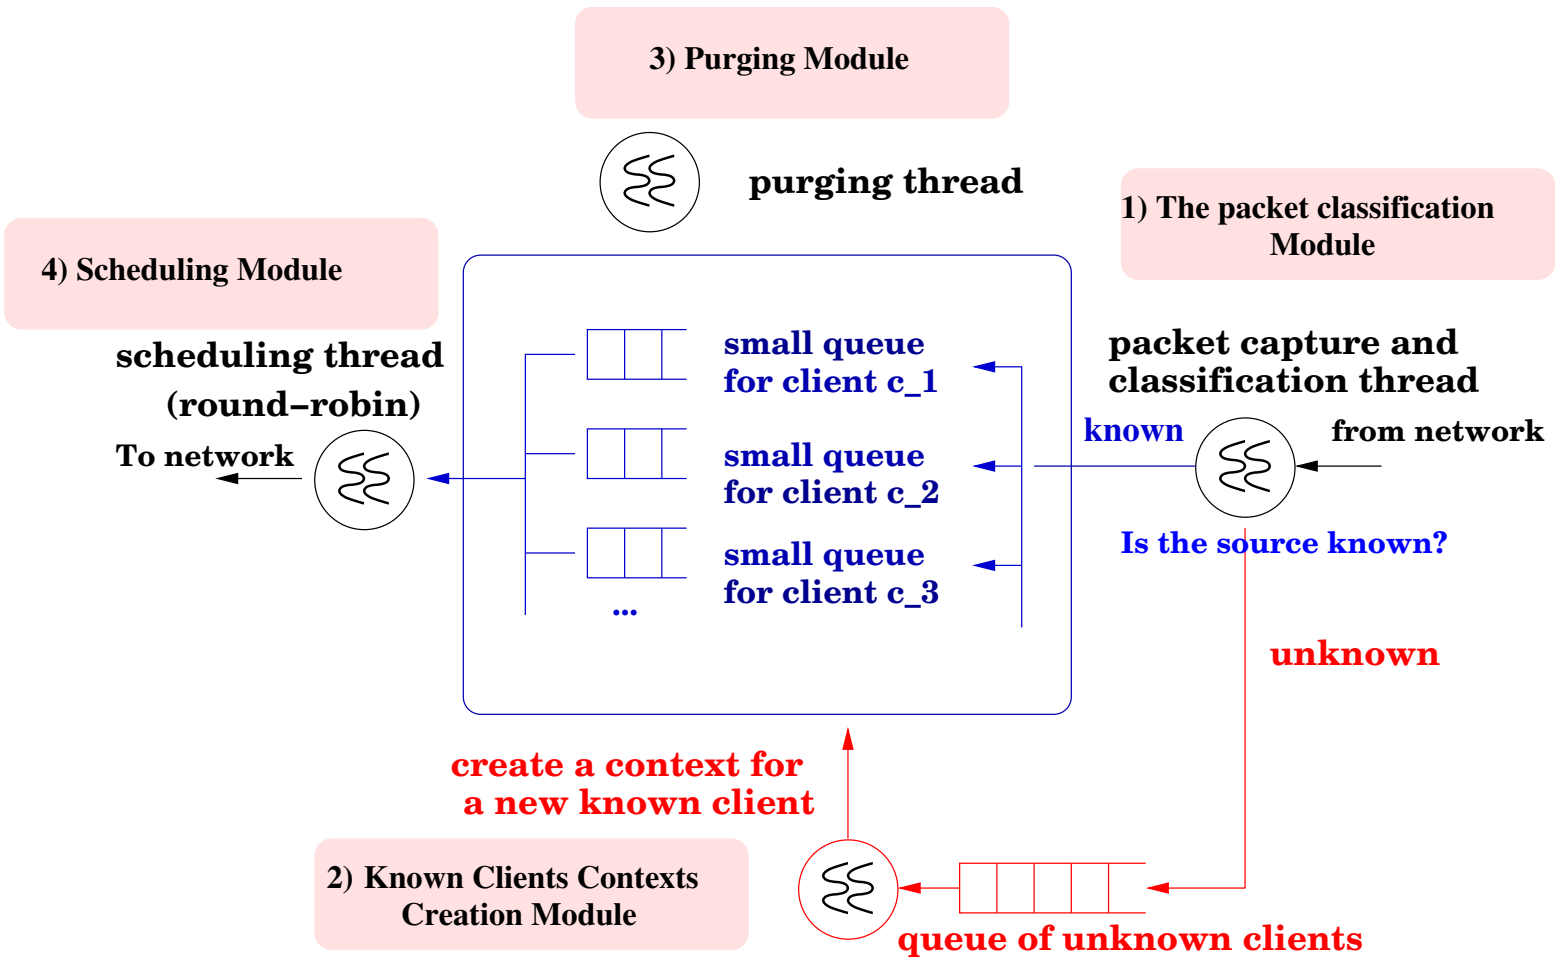

Our Filtering Proposal: Other Extensions

Other Extensions are Possible:

- IP/MAC addresses spoofing resiliency component.
- Use of learning mechanisms.
- Group management policy.
- Congestion control policy.

$\Rightarrow$ We will discuss these extensions later.

Two Possibilities: Internal versus External Deployment

External

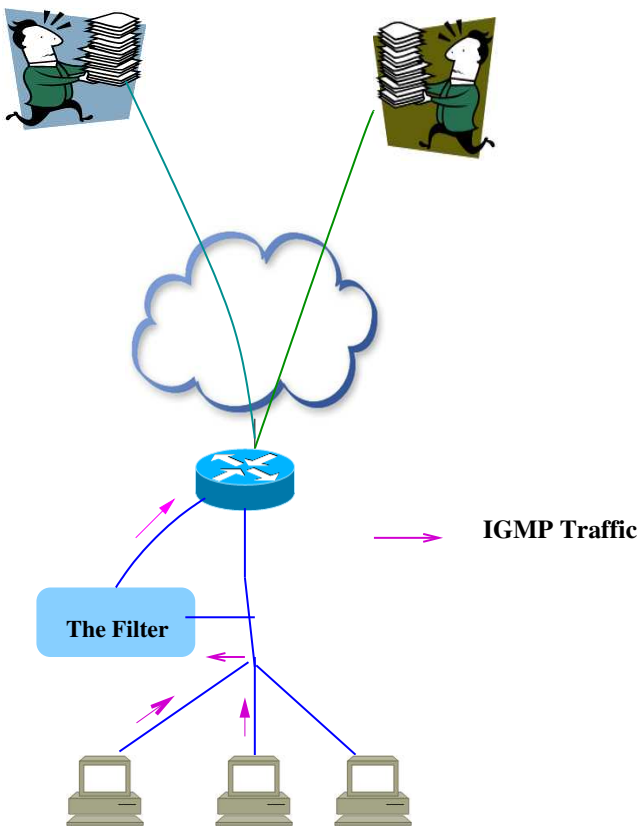

Internal

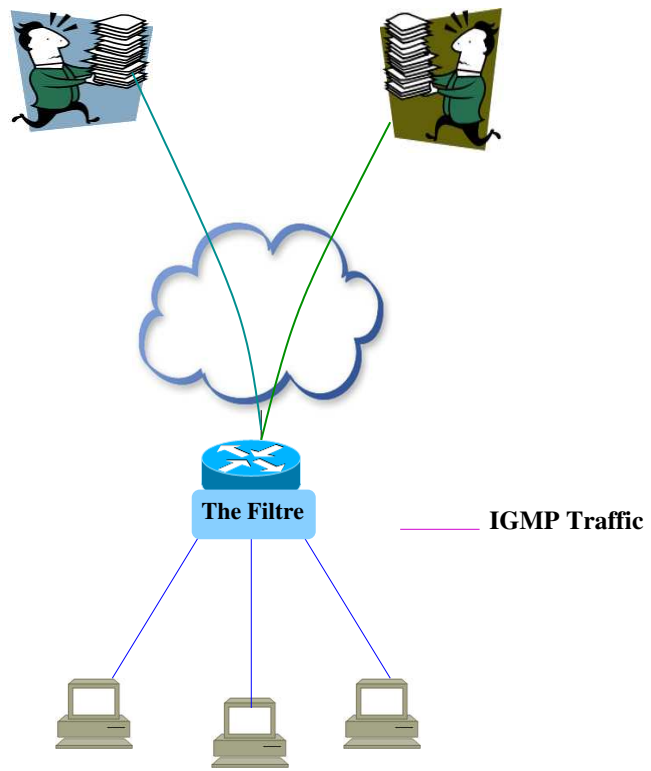

Deployment in the Operator Network

Example: DSLAM is an IGMP proxy:

- IGMP proxy acts in a dual mode as an IGMP router and IGMP client.
- IGMP processing is done in the DSLAM and the BRAS.

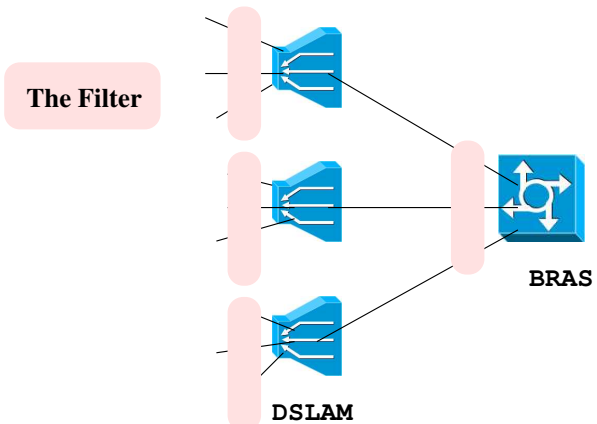

- The identification can be done per VLAN ID, MAC/IP addr, interface...

Deployment in the Operator Network (cont.)

Example: DSLAM is an IGMP snooper:

- The snooping DSLAM listens promiscuously to transitions between clients and the BRAS.
- IGMP processing is done in the BRAS.

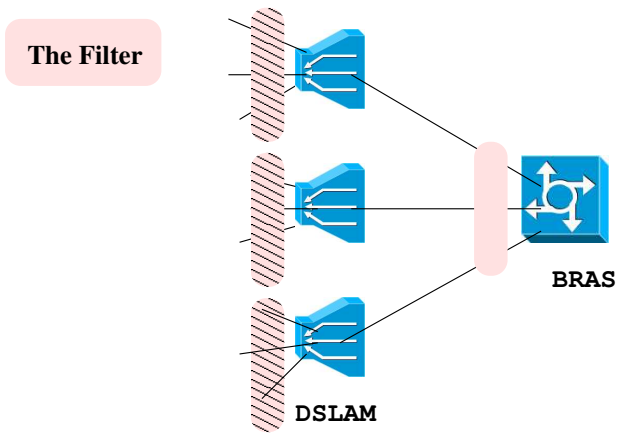

- Adding the filter before/in the DSLAMs in this case is optional.
- The identification can be done per MAC/IP addr, interface.

Filter Dimensioning
How to Initialize the Filter?

We identified a set of key parameters:

- Some parameters are obtained theoretically via a mathematical model.
- Other parameters are initialized by the operator.
  - Scheduling rate for known clients' packets.
  - Maximum number of groups per client.

They depend on:

- The underlying infrastructure (LAN, Point to Point lines).
- The timer settings of the IGMP/MLD protocols.
- The users behavior: generally modelized by Poisson process.

Filter Dimensioning:

Dimensioning the UCQ (Unknown Clients Queue): Two traffic models:

- Constant traffic: $D/D/1/K$ queue model.
- Poisson traffic: $M/D/1/K$ queue model.
  $\Rightarrow$ The parameters of UNQ are chosen according to a wanted acceptance probability $P = 1 - P_K$
  P_K is the probability of having K packets in the system.

Dimensioning the known clients queues

- Each client is modelizd by a: $M/M/1/G_{max}$ queue model.
  M depends on the scheduling rate.

Filter Dimensioning (cont.):

We determine the purging period Δ as a function of the known clients number and their activities.

Part IV

Experimental Evaluation

Experimental Environment

Platform:

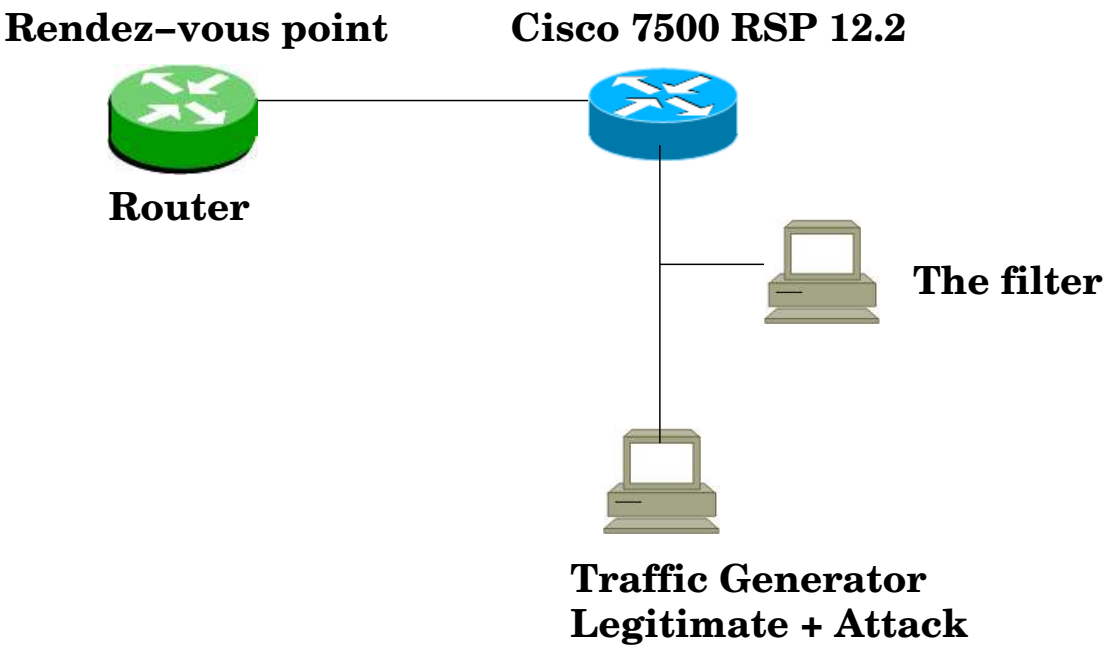

- The DR is a Cisco 7500 RSP 12.2, running PIM-SM.
- The PCs are equipped with Intel Pentium/Mandrake Linux v10.0.
- The IGMP timers: default values.

We have implemented the filter and the traffic generator with different traffic patterns on Linux.

Test Scenarios

Three tests:

First test

An attacker spoofs 1 source address. It reports over 500 group addresses other than those requested by the legitimate clients.

Second test

The attacker spoofs 10 source addresses.

Third test

The attacker spoofs 1000 source addresses.

Filter Parameters

- Size of the unknown client queue (K) = 20 packets
- Unknown client queue service rate (μ) = 15 pps
- Purging period Δ = 80 s
  $\Rightarrow$ Obtained via the mathematical model.
- Scheduling rate for known clients' packets (s) = 20 pps
- Maximum number of groups per client (G_{max})= 6 packets
  $\Rightarrow$ chosen empirically.

Results Test 1: Single Forged IP Addresses by the Attacker and IGMP Requests Arrive With Deterministic Rate

Incoming traffic

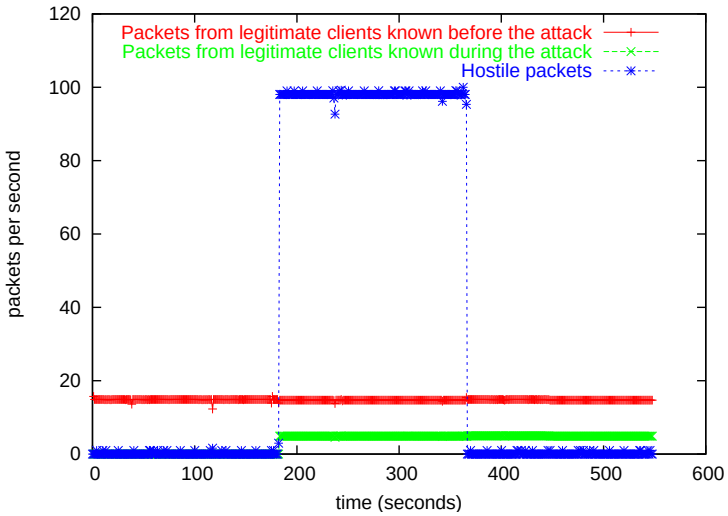

outgoing traffic

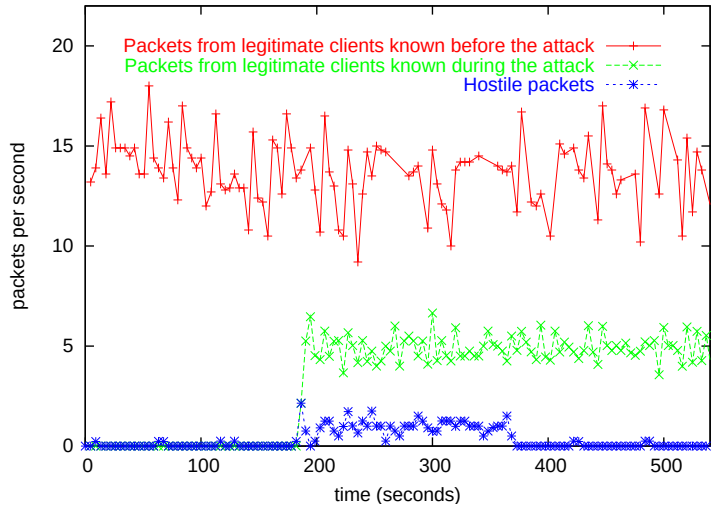

⇒ The attacker receives a significantly lower share than its incoming traffic rate and new clients are still accepted during the attack.

Results Test 1: Single Forged IP Addresses by the Attacker and IGMP Requests Arrive With Deterministic Rate (cont.)

Without filter

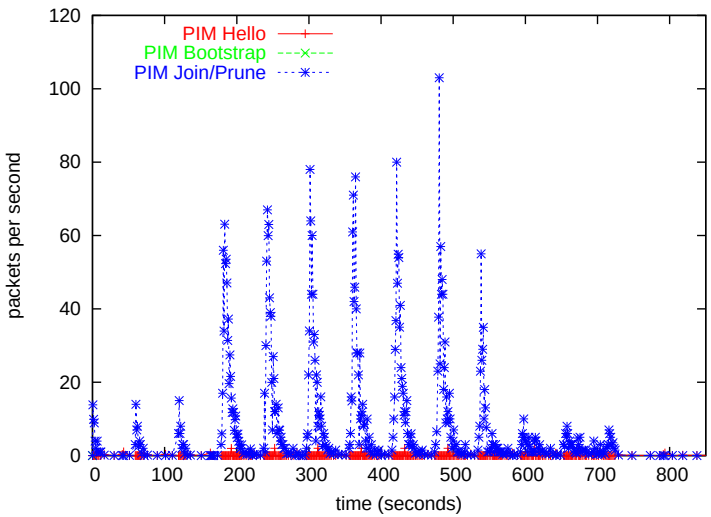

With filter

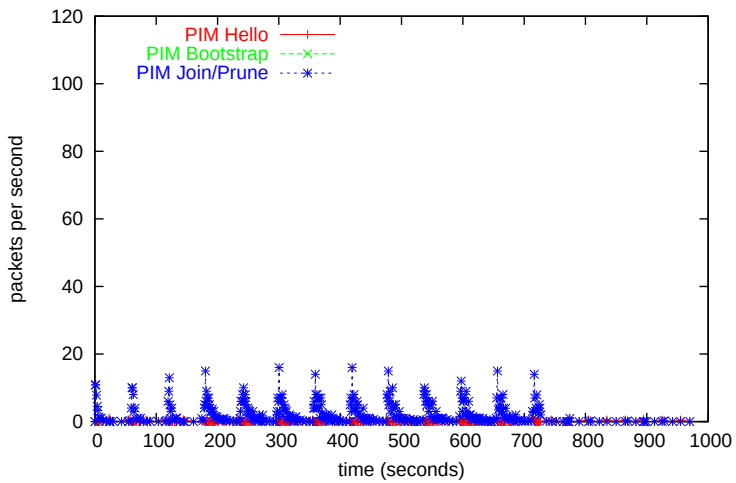

⇒ The number of PIM Join/Prune messages is significantly reduced (Maximum of 16 pps versus 105 pps).

Results Test 1: Single Forged IP Addresses by the Attacker and IGMP Requests Arrive With Deterministic Rate (cont.)

Memory Consumption of the Cisco Router Without/With Filter

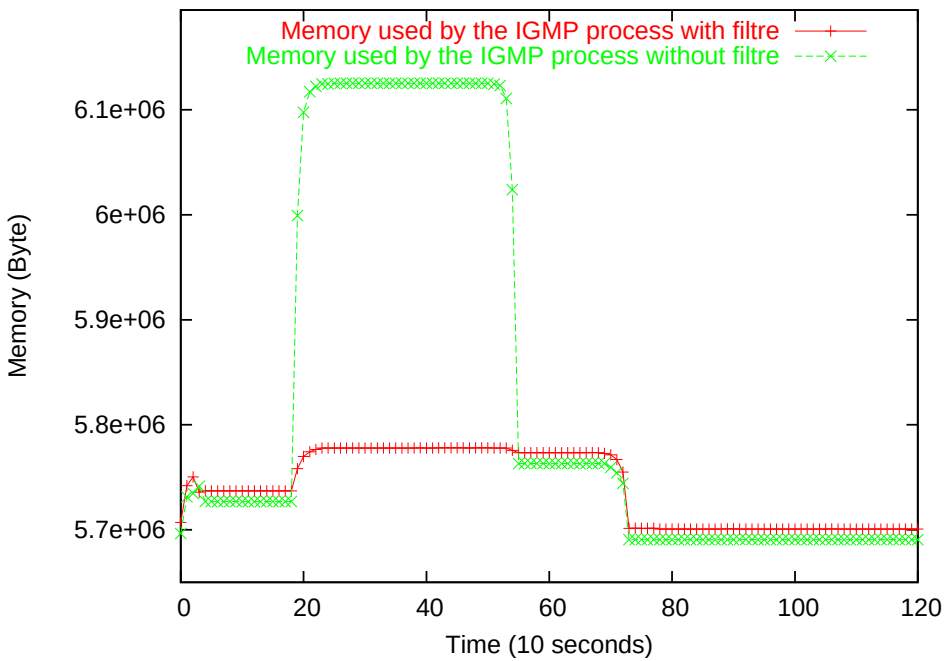

⇒ Router memory consumption is largely improved

Results Test 3: 1000 Forged IP Addresses by the Attacker and IGMP Requests Arrive With Deterministic Rate

Incoming traffic

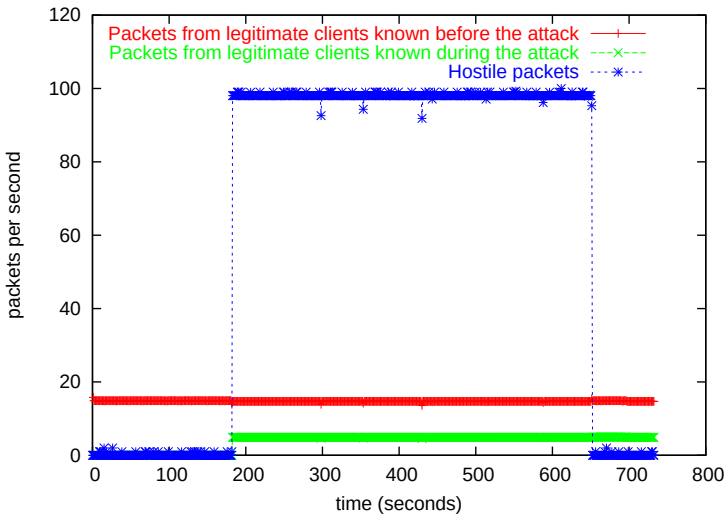

Outgoing traffic

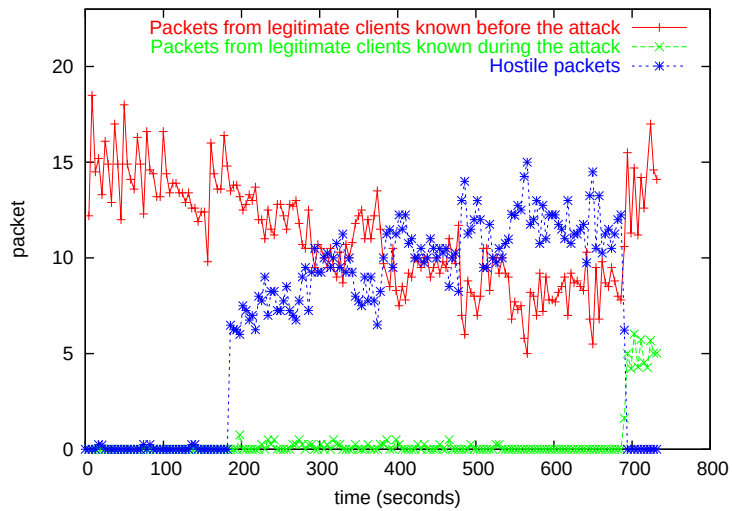

⇒ hostile clients entering the system increases but the DoS attack does not impact the whole multicast routing infrastructure.

The Good News...

For the Operator: Service continuity goal is achieved:

- Service continuity for clients accepted before the attack.
- Total outgoing traffic, leaving the filter and entering the router, is limited even in the presence of an attack.
- Router memory consumption is largely reduced.
- The core of the network is protected from being flooded by routing messages.
- Lightweight mechanism.

The Bad News...

Vulnerable to IP addresses spoofing:

- When the attacker uses random IP address spoofing:
  The legitimate new clients arriving with the attack traffic can be largely affected when the number of the spoofed addresses is large.
- When the attacker uses targeted IP address spoofing:
  The known clients can be largely penalized.

The Bad News...

⇒ No surprise!
⇒ Some extensions are possible thanks to keeping a state per client in the system.

Part V

Discussion of Possible Extensions

IP/MAC Addresses Spoofing Resiliency: Source Addresses Inspection

- Source addresses inspection.
- Interaction with the DHCP server to protect against unused subnet addresses spoofing.
- Using DHCP relay option 82 to protect against IP and MAC addresses spoofing.

IP/MAC Addresses Spoofing Resiliency: Source Addresses Inspection (cont.)

- DHCP relay option 82 associates each IP/MAC with DSL-ID, VC/PVC,...

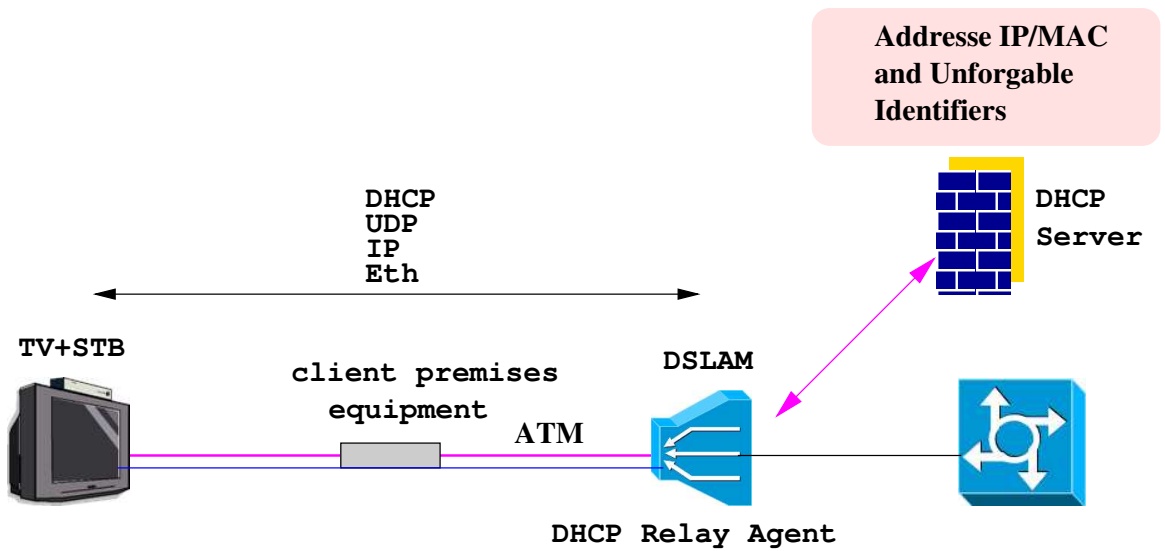

IP/MAC Addresses Spoofing Resiliency: Using Unforgeable Identifiers

- In the broadband networks some identifiers can be used to uniquely identify the end users: VPI/VCI, DSL-ID, VLAN tag...

IP/MAC Addresses Spoofing Resiliency: Using Unforgeable Identifiers (cont.)

We distinguish between two cases:

- The filter is implemented in the DSLAM:
  The DSLAM knows these identifiers.

- The filter is implemented in the B-RAS:
  - The DSLAM can know the valid IP address associated to unforgeable identifier thanks to the DHCP server (option 82).
  - The DSLAM drops packets with invalid (source IP address/unforgeable identifier).

IP/MAC Addresses Spoofing Resiliency: Using Unforgeable Identifiers (cont.)

We distinguish between two cases:

- The filter is implemented in the DSLAM:
  The DSLAM knows these identifiers.

- The filter is implemented in the B-RAS:
  - The DSLAM can know the valid IP address associated to unforgeable identifier thanks to the DHCP server (option 82).
  - The DSLAM drops packets with invalid (source IP address/unforgeable identifier).

IP/MAC Addresses Spoofing Resiliency: Using Unforgeable Identifiers (cont.)

We distinguish between two cases:

- The filter is implemented in the DSLAM:
  The DSLAM knows these identifiers.

- The filter is implemented in the B-RAS:
  - The DSLAM can know the valid IP address associated to unforgeable identifier thanks to the DHCP server (option 82).
  - The DSLAM drops packets with invalid (source IP address/unforgeable identifier).

IP/MAC Addresses Spoofing Resiliency: Adding Learning Mechanism

- In environments where the use of unforgeable identifiers is not possible.
- Aim to take decisions before accepting an IP address (client) and to determine the lifetime and the priority of known clients in the system.

IP/MAC Addresses Spoofing Resiliency: Adding Learning Mechanism

- In environments where the use of unforgeable identifiers is not possible.
- Aim to take decisions before accepting an IP address (client) and to determine the lifetime and the priority of known clients in the system.

IP/MAC Addresses Spoofing Resiliency: Adding Learning Mechanism (cont.)

Three parameters:

- The burst factor of the arriving packets.
- The IGMP finite state machine conformity. (Example: a LEAVE message is received without a previous JOIN...)
- The proximity of the IP addresses.

IP/MAC Addresses Spoofing Resiliency: Adding Learning Mechanism (cont.)

These factors are used to take many decisions:

- Dropping the arriving packets immediately.
- Accepting the new clients.
- Classifying the accepted clients into clusters of priorities. Use the K-means algorithm.
- Degrading or upgrading the priorities of known clients.

Discussion of Possible Extensions: Group Management Policy

- The maximum number of groups for each clients.
- Who can subscribe to what group.

$\Rightarrow$ Adding such a policy can improve the filtering efficiency.

Discussion of Possible Extensions: Congestion Control Policy

Aim to avoid data plane attacks.

- The congestion state of a client determines if he is authorized to subscribe to a particular multicast group.

To obtain the congestion state of each line DSL or of each VLAN the filter can interact with access network topology discovery mechanism.

Discussion of Possible Extensions: Congestion Control Policy

Aim to avoid data plane attacks.

- The congestion state of a client determines if he is authorized to subscribe to a particular multicast group.

To obtain the congestion state of each line DSL or of each VLAN the filter can interact with access network topology discovery mechanism.

Part VI

Conclusion

Conclusion

Our proposal:

- Allows the operator to survive IGMP/MLD flooding attacks.
- Protects the well-behaved clients against the ill-behaved ones.
- Keeps the changes to the current used multicast routing infrastructure as minimum as possible.
- Avoids non realistic assumptions and techniques.
- Is easily extensible.

Future Works

- Implementing the proposed extensions
  - Using IP addresse spoofing resiliency component.
  - Adding learning component.
  - Adding group management policy.
  - Adding congestion control policy.

  $\Rightarrow$ These extensions can help to thwart other IGMP/MLD attacks.
- Using the filter with other protocols: MSDP, PIM/SM,...

A Few More Words...

- The security of the multicast routing infrastructure is a challenging task and a tough problem.
- Many open points.
- Our proposal solves pragmatically a subset of them with reasonable costs.

Many thanks for your attention!
Questions?

Our Achievements:

- Zainab Khallouf, Vincent Roca, Renaud Moignard, and Sébastien Loye, Infrastructure Sécurisée de Routage Multipoint: le Point de Vue de l'Opérateur de Réseau, 3ème Conférence sur la Sécurité et Architectures Réseaux (SAR'04), La Londe, Cote d'Azur, France, June, 2004.
- Zainab Khallouf, Vincent Roca, Renaud Moignard, and Sébastien Loye, A Filtering Approach for an IGMP Flooding Resilient Infrastructure, 4ème Conférence sur la Sécurité et Architectures Réseaux (SAR'04), Batz sur Mer, France, June, 2005.
- Contribution to the writing of the deliverable D2.5: Requirements on Security in 6QM Project. QM Consortium, 20/6/2003. http://www.ist-mome.org/documents/6qm_pp_d2_5_v3_1.pdf.
- Software (C/AWK/Linux), Implementation of a filtering mechanism to mitigate IGMP flooding attacks.