

INSTITUT NATIONAL POLYTECHNIQUE DE GRENOBLE

N° attribué par la bibliothèque

--	--	--	--	--	--	--	--	--	--	--	--

THÈSE

pour obtenir le grade de

DOCTEUR DE L'INPG

Spécialité: «Systèmes et Logiciels»

préparée à France Télécom R&D et l'INRIA Rhône Alpes, projet Planète

sous d'une convention CIFRE

dans le cadre de l'**École Doctorale « MATHÉMATIQUE, SCIENCES ET
TECHNOLOGIE DE L'INFORMATION »**

présentée et soutenue publiquement

par

Zainab Khallouf


le 15 Mars, 2006

Titre:

**Infrastructure Sécurisée de Routage Multipoint:
Le Point de Vue de l'Opérateur Réseau**

Composition de Jury

Prof.	Guy Mazaré, ENSIMAG	Président
Dr.	Isabelle Chrisment, (LORIA, Université Henri Poincaré (Nancy 1))	Rapporteur
Prof.	Bernard Cousin, (IRISA / Université de Rennes 1)	Rapporteur
Dr.	Olivier Paul, INT, Evry	Examineur
Prof.	Andrzej Duda, (LSR-IMAG, ENSIMAG)	Directeur de thèse
Dr.	Vincent Roca, INRIA Rhône Alpes	Encadrement
Eng.	Sébastien Loye, France Télécom R&D	Encadrement



To My Parents with Love and Gratitude.

Acknowledgements

*Gratitude is when memory is stored
in the heart and not in the mind.*

- Lionel Hampton

Sincere appreciation from my heart can not be conveyed in words to all the people who helped me carry out the research reported in this dissertation, by their support and encouragement, and to all those who have made my stay in France, the land of liberty, equality, and fraternity, an exciting and rewarding experience by their friendship and kindness.

My sincere thanks to my advisors, Dr. Vincent Roca and Eng. Sébastien Loyer. I am grateful for their advice, guidance, and patience. I am deeply indebted to Dr. Vincent Roca for giving me the chance to work with him, for his support, patience, and for all what I've learned from him. Many thanks also to Prof. Andrzej Duda the director of my thesis.

All my gratitude to prof. Bernard Cousin and Dr. Isabelle Chrisment for the time they spent reading and commenting my thesis, and for the reports they wrote. And I also would like to thank Dr. Olivier Paul and Prof. Guy Mazaré for taking part in the committee of my thesis.

I would like to express my immense appreciation to Dr. Renaud Moignard the head of the MTM (MPLS Traffic Engineering and Multicast) team at France Télécom R&D for his support during all this thesis period. I also wish to express appreciations to all my colleagues in the MTM team. A special thank to Mr. Jean-Luc Lehagré for his assistance in implementing the test-bed used in this study.

I am grateful to Dr. Christian Guillemot, director of the CPN (Core Packet Network for NGN) laboratory for welcoming me at France Télécom R&D, and I would like to thank all the members of the CORE/CPN laboratory at Lannion they all helped me in a way or in another.

I'm grateful to all the past and present members of the Planète team, for their support and helps. I was honored to work with them. I thank them sincerely for all the memorable moments, for their helps, and for their kindness.

I would like to thank the many people who helped, advised, and gave fruitful discussions during this work. Many thanks to Eng. Jacqueline Boyer from France Télécom R&D at Lannion. I'm grateful for her help in the theoretical study presented in this thesis. I also would like to thank Dr. Mohammed Alchemlal from France Télécom R&D at Caen.

Special thanks to Mr. Jean Louis Le Roux, Mrs. Sana Abada, Mrs. Dina Sol, Mr. Patrick Fleming, Mr. Yacine Mami, Miss. Imène Chaieb, Mrs. Thi Hue Nguyen, Mr. Nawras Aldibbiat and all the other friends for their kindness and friendship. Thank you for everything!

Last but not least, I would like to thank my family, I'm grateful to my sisters **Khouzam** and **Hadeel**, my brother in law **Morhaf** and the small angle my nephew **Karam** for their love and support. And, it is beyond all expression my gratitude to **my parents** for their endless love, support and encouragement. I appreciate, with all my heart, all what they have done for me, for learning me that knowledge is liberator, for respecting my choices, and for encouraging me to be what I want to be. To them I would like to dedicate this dissertation.

Our Achievements: Published Papers and Implementations

Conferences Papers

- Z. Khallouf, V. Roca, R. Moignard, S. Loye, *A Filtering Approach for an IGMP Flooding Resilient Infrastructure*. 4th Conference on Security and Network Architectures (SAR'05), Batz sur Mer (France), June 06-10, 2005.
- Z. Khallouf, V. Roca, R. Moignard, S. Loye, *Infrastructure Sécurisée de Routage Multipoint: le Point de Vue de l'Opérateur Réseau*. 3rd Conference on Security and Network Architectures (SAR'04), La Londe, Cote d'Azur, France, June 2004.

Software (C/AWK/Linux)

- Implementation of a filtering mechanism to mitigate IGMP flooding attacks.

Abstract

Multicast is a promising technology for the distribution of streaming media, bulk data and many other added-value applications, yet the deployment of multicast still in its infancy.

Providing security is still one of the main challenge that hinders the introduction of multicast in the existing infrastructure. It is therefore critical to provide sound security mechanisms that can protect the ISPs and the carriers infrastructure against multicast threats and allow them in the same time to get all the benefits of introducing multicast in their networks.

This work considers one of the most challenging features of multicast deployment: the security of the multicast routing infrastructure.

In this thesis we adopt the security from the multicast Network Operator's viewpoint. The kind of security required by a network operator, who manages and operates the multicast routing infrastructure, largely differs from that of end-to-end security. More specifically, the operator is concerned by service continuity no matter what happens. In other words, the operator wants the group communication service provided to its clients (i.e. end users or other network operators with whom has peering relationships) must be operational at any time, in spite of anomalies in the multicast flows, no matter whether they are intentional (i.e. are the result of deliberate attacks) or not (e.g. are caused by a misbehaving component).

While many theoretically ideal proposals have been done to secure the routing protocols, they have rarely been accepted by the operators community. For instance, because they require to modify existing and widely deployed protocols, or they introduce heavy authentication mechanisms, which is in practice almost impossible to deploy in legacy networks, and even infeasible, since a corrupted host may be the source of a DoS attack, even if it has been authenticated.

In this thesis analyze in depth the threats to the multicast infrastructure. We show that the vulnerability of the multicast model comes largely from the edge. More specifically, several attacks arise from the use of group management protocols, IGMP for IPv4 and MLD for IPv6. In the light of this analysis, we introduce and evaluate a simple, yet efficient filtering approach to thwart some DoS attacks that are based on IGMP or MLD flooding, and that threaten the whole operator's infrastructure.

A key feature of our proposal is that it follows a realistic and pragmatic approach, and in particular it does not require any modification to the existing, widely deployed protocols.

Keywords: Multicast, Multicast Routing Protocols, Network Operator, IGMP/MLD, Infrastructure Security, DoS attacks.

Résumé

Le multicast est un mécanisme efficace qui permet à un grand nombre de récepteurs de recevoir le même contenu puisqu'un paquet traverse une fois et une seule un lien donné. Le multicast a été étudié de longue date et pourtant à ce jour aucun déploiement à grande échelle n'a eu lieu. Aujourd'hui l'une des raisons du non déploiement du multicast auprès des opérateurs de réseaux est la problématique de la sécurité. Or, deux niveaux complémentaires de sécurité doivent être considérés: (1) la sécurité applicative qui est essentiellement la préoccupation des clients finaux et les fournisseurs de contenu, et (2) la sécurité de l'infrastructure de routage multicast, qui est la préoccupation de l'opérateur de réseau.

Ce travail considère la sécurité de l'infrastructure multicast du point de vue de l'opérateur de réseau. Or, l'opérateur est essentiellement concerné par un problème de "continuité de service en toutes circonstances", même dans le cas où son réseau est victime d'une attaque.

Dans cette thèse nous identifions les attaques possibles, nous les classons selon leurs dangers pour l'opérateur, et identifions divers mécanismes de sécurité pour y faire face. Cette étude révèle que l'infrastructure est fortement vulnérable aux attaques DoS consommant les ressources de réseau, ce dernier devenant alors lent voire indisponible. Ces attaques sont faciles à lancer de la périphérie de réseau (intentionnellement ou non) en utilisant les protocoles de gestion de groupe IGMP/MLD. Notre étude révèle également les limites des approches proposées pour répondre à ces attaques.

A la lumière de l'analyse détaillée de la problématique, des vulnérabilités, et des solutions actuelles, nous proposons une nouvelle approche pour aider le réseau de l'opérateur à se défendre contre les attaques basées sur IGMP ou MLD. Notre proposition suit une approche pragmatique et flexible, qui garantit qu'elle sera facilement déployée dans les infrastructures existantes, et vise également à protéger les clients légitimes en cas d'attaque.

Mots clés: Multicast, Les Protocoles de Routage Multicast, l'Opérateur Réseau, IGMP/MLD, La sécurité de l'infrastructure, Les attaques DoS.

Contents

1	An Overview	1
1.1	Introduction and Motivations	1
1.2	Goals of the Thesis	3
1.3	Dissertation Roadmap	4
2	Multicast Deployment by the Network Operator: a State of the Art	7
2.1	Multicast Delivery: the Operational Model	7
2.1.1	The Various Actors and the Services Provided	7
2.2	Deployment Examples	9
2.2.1	First Example: Commercial Delivery of Video Over xDSL	10
2.2.2	Second Example: Multicast Delivery of Free Content	11
2.3	Conclusion	11
3	Multicast Routing Protocols: Overview	13
3.1	Multicast Service Models	13
3.2	Group Management Protocols	14
3.2.1	Internet Group Management Protocol (IGMP)	14
3.2.2	Multicast Listener Protocol (MLD)	22
3.2.3	Group Management Protocols Proxying	22
3.2.4	Group Management Protocols in Layer2 Switching environment	22
3.3	Multicast Routing Protocols	23
3.3.1	Intra-Domain Routing Protocols	23
3.3.2	Inter-domain multicast routing architectures and protocols	26
3.4	Consluion	28
4	Taxonomy of Possible Attacks of Interest to the Network Operator	29
4.1	Introduction	29
4.2	Internal Attacks	30
4.2.1	Internal data attacks	30
4.2.2	Internal control attacks	30
4.3	Edge attacks	31
4.3.1	Edge Data Attacks	31
4.3.2	Edge control attacks	33

CONTENTS

4.4	Conclusion	33
4.4.1	Classification of Attacks According to their Importance and Likelihood	34
4.5	Conclusion	35
5	A Focus on Group Management Protocols Specific Attacks	37
5.1	Introduction	37
5.2	Attacks on the IGMP/MLD routers	38
5.2.1	Querier Impersonation	38
5.2.2	Querier Paralysis	39
5.2.3	IGMP/MLD Querier Degradation	40
5.3	Attacks on the Hosts	40
5.4	Classifying the Group Management Protocols Threats According to their impacts on the Operator	41
5.5	Conclusion	42
6	A Survey of Proposals to Secure the Multicast Routing Infrastructure	43
6.1	Attack Avoidance Approaches:	44
6.1.1	Securing the Edges of the Multicast Tree	44
6.1.2	Protecting the Core of the Multicast Tree	47
6.2	Attack Resiliency Approaches:	48
6.3	Hybrid Techniques	49
6.4	Conclusion	50
6.4.1	Participant Authentication and Authorization	50
6.4.2	Modifying Existing Protocols	51
6.4.3	Inter-Domain Confidence	51
6.4.4	Simple filtering	51
7	Our Proposal: A Filtering Approach to Mitigate IGMP/MLD Flooding Attacks	53
7.1	Architectural Overview	54
7.1.1	Packets Capturing	54
7.1.2	Packets Classification	54
7.1.3	Clients Purging	55
7.2	A Closer View of the Various Building Blocks	55
7.2.1	The packet capture and classification thread:	55
7.2.2	The known clients queues creation thread:	56
7.2.3	The main scheduling thread:	56
7.2.4	The purging thread:	56
7.3	Deployment	57
7.4	Benefits in Front of a DoS Flooding Attack	57
7.4.1	Naive Flooding DoS Attack Without IP Address Spoofing	57
7.4.2	Flooding DoS Attack With Random IP Address Spoofing	58
7.4.3	Flooding DoS Attack With Targeted IP Address Spoofing	58

7.4.4	What About other Group Management Specific Attacks?	59
7.5	Conclusion	59
8	Filter Dimensioning: Theoretical Study	61
8.1	Key Parameters	61
8.2	Dimensioning the Unknown Clients Queue: a Theoretical Study	64
8.2.1	Deterministic Arrival Model	64
8.2.2	Poisson Arrival Model	65
8.3	Dimensioning the Known Client Queues	66
8.3.1	Estimating the Number of Known Client Queues	67
8.3.2	Estimating the Number of the Waiting Packets for the Known Client Queues	68
8.3.3	Evaluating the Waiting Time	70
8.4	Conclusion	70
9	Evaluation	71
9.1	Introduction	71
9.2	Experimental Environment	72
9.3	Tests Scenarios	73
9.4	IGMP Traffic Sent at a Constant Rate	74
9.4.1	Filter Parameters	74
9.4.2	Results Test 1: Single Forged IP Addresses	75
9.4.3	Results for Test 2: 10 Forged IP Addresses	77
9.4.4	Results for Test 3: 1000 Forged IP Addresses	80
9.5	IGMP Traffic Sent with a Poisson Distribution	80
9.5.1	Introduction to Poisson Distribution	81
9.5.2	Generating Traffic with Poisson Distribution	81
9.5.3	Filter Parameters	82
9.5.4	Results Test 1: Single Forged IP Addresses	82
9.5.5	Results Test 2: Ten Forged IP Addresses	84
9.5.6	Results for Test 3: 1000 Forged IP Addresses	87
9.6	Conclusion	88
10	Discussion	89
10.1	Filter Deployment: Practical Examples	89
10.1.1	Filter Deployment in the Operator Network	89
10.1.2	Filter Deployment in a Campus Network	92
10.2	Discussion of Some Extensions to Improve the Filtering Efficiency	93
10.2.1	IP/MAC Addresses Spoofing Resiliency	93
10.3	Conclusion	94

CONTENTS

11 Conclusion and Future Works	95
11.1 Reminder of the Objective of the Thesis	95
11.2 Our Proposal: A Pragmatic and Deployable Filter-Based Solution	96
11.3 Future Works	97
11.3.1 Membership Policy	97
11.3.2 Signature Based Anomaly Detection	97
11.3.3 Using Learning Mechanisms	98
11.4 Final Words	98
12 Infrastructure Sécurisée de Routage Multipoint:	
Le Point de Vue de l'Opérateur Réseau	101
12.1 Introduction	101
12.2 Les Différents Acteurs dans le Cadre d'un Service de Diffusion Multicast	102
12.2.1 Vue générale	102
12.2.2 Premier exemple : service commercial de TV sur ADSL	103
12.2.3 Deuxième exemple: diffusion libre d'un contenu gratuit	104
12.3 Sécurité de l'Infrastructure du Point de Vue de l'Opérateur Réseau	104
12.4 Taxonomie des Attaques Intéressant l'Opérateur Réseau	105
12.4.1 Attaques internes	105
12.4.2 Attaques venant de la périphérie	106
12.5 Taxonomie des Mécanismes de Protection de l'Infrastructure de l'Opérateur Réseau	107
12.5.1 Mécanismes de Défense Préventifs	107
12.5.2 Mécanismes Réactifs	107
12.5.3 Mécanismes Hybrides	108
12.6 Discussion sur les Attaques et les Techniques de Défense	108
12.6.1 Classification des attaques selon leur impact sur le réseau de l'opérateur	108
12.6.2 Discussion sur les Techniques de Défense	109
12.7 Notre Proposition	110
12.8 Le Déploiement de Filtre	111
12.9 Le Déploiement de Filtre dans le Réseau de l'Opérateur:	112
12.10 Comment Initialiser le Filtre ?	112
12.10.1 Les paramètres de filtre	112
12.10.2 Le dimensionnement de Filtre	112
12.11 Les Résultats Expérimentaux	113
12.11.1 La plateforme de tests	113
12.11.2 Les scénarios de tests	113
12.12 Discussion des Extensions Possibles	114
12.12.1 l'Utilisation des identifiants non forgeables	116
12.12.2 l'Utilisation des mécanismes d'apprentissage	116
12.12.3 L'utilisation d'une politique de gestion de groupes	117
12.12.4 l'utilisation d'une politique de control de congestion	117

CONTENTS

12.13 Conclusions	117
References	119

CONTENTS

List of Figures

2.1	The different actors during a multicast content delivery session.	10
3.1	The Reports Suppression Mechanism.	16
3.2	IGMPv1 packet format.	16
3.3	IGMPv2 packet format.	17
3.4	IGMPv3 packet format.	19
3.5	IGMPv3 group record.	20
3.6	IGMPv3 query.	21
3.7	PIM-SM: Building a new branch of the forwarding tree.	25
3.8	MSDP source active messages propagation.	27
4.1	A General Classification of Possible Attacks of Interest to the Network Operator.	30
4.2	A General Classification of Possible Attacks of Interest to the Network Operator.	34
7.1	Architecture of the filter.	56
7.2	External versus internal deployment of the filter.	57
8.1	Successive waiting times.	65
8.2	Acceptance Probability (P_K) for the Unknown Client Queue (UCQ).	65
8.3	Mean Number of Packets in the UKQ.	66
8.4	Number of known clients	68
8.5	The number of waiting packets as a function of traffic intensity	69
8.6	The number of waiting packets as a function of number of attackers	69
8.7	Evaluating the Average Waiting Time	70
9.1	The Testbed.	73
9.2	Traffic Entering and Leaving the Filter (Test 1).	75
9.3	Memory Consumption of the Cisco Router Without/With Filter (Test 1)	76
9.4	PIM Control Traffic Without/With Filter (Test 1).	76
9.5	Known Clients (Test 1)	77
9.6	Waiting Packets (Test 1)	77
9.7	Average Waiting Time (Test 1)	78

LIST OF FIGURES

9.8	Traffic Leaving the filter (Test 2).	78
9.9	Memory Consumption of the Cisco Router (Test 2).	78
9.10	Known Clients (Test 2)	79
9.11	Waiting Packets (Test 2)	79
9.12	Average Waiting Time (Test 2)	79
9.13	Traffic Entering and Leaving the filter (Test 3).	80
9.14	Memory Consumption of the Cisco Router (Test 3).	81
9.15	Traffic Entering and Leaving the filter (Test 1-Poisson).	83
9.16	Known Clients Number (Test 1-Poisson)	83
9.17	Simulation of the Number of Known clients (SIMSCRIPT)	84
9.18	Waiting Packets (Test 1-Poisson)	84
9.19	Average Waiting Time (Test 1-Poisson)	85
9.20	Incoming traffic (Test 2-Poisson)	85
9.21	scheduled traffic (Test 2-Poisson)	85
9.22	Known clients (Test 2-Poisson)	86
9.23	Simulation of the Known Clients Number (Test 2-Poisson)	86
9.24	Waiting Packets (Test 2-Poisson)	86
9.25	Average Waiting Time (Test 2-Poisson)	87
9.26	Traffic entering and leaving the filter (Test 3).	87
10.1	Deploying the Filter in the DSALM.	90
10.2	Deploying the Filter in the BRAS.	91
10.3	Example of a Campus Network Architecture.	92
12.1	Les Différents Acteurs dans le Cadre d'un Service de Diffusion Multicast.	104
12.2	L'Architecture de Filter.	111
12.3	Le Déploiement de Filtre.	111
12.4	The Testbed.	113
12.5	Le Trafic Entrant et Sortant (Test 1).	114
12.6	PIM Messages Sans/Avec Filtre (Test 1).	115
12.7	La Consommation de Mémoire sur le Routeur Cisco Avec/Sans filtre (Test 1)	115
12.8	Les trafic Entrant et Sortant de Filtre (Test 3).	116

List of Tables

3.1	IGMPv1 message types	17
3.2	IGMPv2 messages types	18
3.3	IGMPv3 messages types	19

List of Acronyms

AAA	Authentication, Authorization, and Accounting
ADSL	Asymmetric Digital Subscriber Line
AAL	ATM adaptation layer
ACL	Access Control List
AH	Authentication Header
ASM	Any Source Multicast
ATM	Asynchronous Transfer Mode
BGP	Border Gateway Protocol
BGMP	Border Gateway Multicast Protocol
B-RAS	Broadband Remote Access Server
CGMP	Cisco Group Multicast Protocol
CHAP	Challenge Handshake Authentication Protocol
CIDR	Classless Interdomain Routing
CPU	Central Processing Unit
DHCP	Dynamic Host Configuration Protocol
DR	Designated Router
DSL	Digital Subscriber Line
DSLAM	Digital Subscriber Line Access Multiplexer
DoS	Denial of Service
DVMRP	Distance Vector Multicast Routing Protocol
EAP	Extensible Authentication Protocol
EARTH	EAsy IP multicast Routing THrough
FIB	Forwarding Information Base
FIFO	First in First Out
HMAC	Hashed Message Authentication Code
IANA	Internet Assigned Numbers Authority
IETF	Internet Engineering Task Force
IDS	Intrusion Detection System Module
IGMP	Internet Group Management Protocol
IGMPv2	IGMP version 2
IGMPv3	IGMP version 3
ISP	Internet Service Provider
IP	IP Address

LIST OF TABLES

LAN	Local Area Network
MFIB	Multicast Forwarding Information Base
MIB	Management Information Base
MLDv2	Multicast Listener Discovery
MPLS	Multi-protocol Label Switching
MSDP	Multicast Source Discovery Protocol
MAC	Media Access Control; Message Authentication Code
MARS	Multicast Address Resolution Service
MBGP	Multicast Border Gateway Protocol
MBone	Multicast Backbone
MD5	Message Digest 5
MSEC	Multicast Security Group
NAS	Network Access Server
NAT	Network Address Translation
NO	Network Operator
PEP	Policy Enforcement Point
PPP	Point-to-Point Protocol
PIM	Protocol Independent Multicast
PIM DM	Protocol Independent Multicast Dense Mode
PIM SM	Protocol Independent Multicast Sparse Mode
PIM SSM	Protocol Independent Multicast Source-Specific Multicast
PKI	public key infrastructure
PPPoE	Point-to-Point Protocol over Ethernet
PVC	Permanent Virtual Circuit (or connection)
RADIUS	Remote Access Dial-In User Service
RGMP	Router-Ports Group Management Protocol
RIF	Routing Information Field
RLM	Receiver-Driven Layered Multicast
RP	Rendezvous Point (router)
RPF	Reverse-Path Forwarding
STP	Spanning Tree Protocol
STB	Set-Top Box
SVC	Switched Virtual Circuit
SVI	Switched Virtual Interface
SFM	Source-Filtered Multicast
SSM	Source Specific Multicast
TCP/IP	Transmission Control Protocol/Internet Protocol
TTL	Time To Live
UDP	User Datagram Protocol
UCQ	Unknown Client Queue
VLAN	Virtual LAN
VENUS	Very Extensive Non-Unicast Service
xDSL	Combined term used to refer to ADSL, HDSL, SDSL, and VDSL

Chapter 1

An Overview

Contents

1.1	Introduction and Motivations	1
1.2	Goals of the Thesis	3
1.3	Dissertation Roadmap	4

Multicast is a promising technology for the distribution of streaming media, bulk data and many other added-value application, yet the deployment of multicast still in its infancy. In this chapter, we glance over IP multicast technology. We summarize the advantages of multicast and the challenges facing its deployment by the ISPs (Internet Service Provider) and the carriers community. Finally, we introduce the goals and the organization of this thesis.

1.1 Introduction and Motivations

IP multicasting is the transmission of an IP datagram to a "host group", a set of zero or more hosts identified by a single IP destination address. A multicast datagram is delivered to all members of its destination host group with the same "best-efforts" reliability as regular unicast IP datagrams. The membership of a host group is dynamic; that is, hosts may join and leave groups at any time. There is no restriction on the location or number of members in a host group. A host may be a member of more than one group at a time. A host need not be a member of a group to send datagrams to it. [24].

Currently, three different service models for multicast exist. The initial multicast model, which corresponds to the definition above, and referred to as *Any Source Multicast (ASM)* [24]; a new service model which was lately introduced called *Source-Specific Multicast (SSM)*; and the *Source-Filtered Multicast (SFM)*

With ASM, hosts use the group management protocols (Internet Group Management Protocol (IGMP)v1, v2 [24] with IPv4 or the Multicast Listener Discovery Protocol

1. AN OVERVIEW

(MLD)v1 [25] with IPv6 to report their group membership interests to the Designated Router.

The SSM model provides host applications with a "channel" abstraction, in which each channel has exactly one source and any number of receivers. An IP datagram is transmitted by a source S to SSM destination address G, and receivers can receive this datagram by subscribing to channel (S,G).

In order to support the new features of the SSM and SFM models, a new version of the group management protocol was introduced: IGMP v3 [16] with IPv4 and MLD v2 [72] with IPv6.

The multicast distribution tree is constructed by exchanging routing messages between multicast capable routers according to an intra-domain routing protocol, the most widely used being currently the PIM-SM [25] protocol. Providing the reachability and path information between the multicast domains is achieved through inter-domain routing protocols: the Multicast Source Discovery Protocol (MSDP) [28] or the Border Gateway Multicast Protocol (BGMP) [70].

Using multicast delivery to send the same content to multiple receivers reduces the aggregate bandwidth required from the network compared to a unicast solution. However, multicast is not yet widely deployed in the Internet and it continues to be a challenging task for the carriers community [50]. Furthermore, as has been cited in [26]: *Multicast is included with the standard set of protocols shipped with most commercial routers, but most IP carriers have not yet enabled the service in their networks. A number of issues have stalled the widespread use of multicast.*

- *Group management, including authorization for group creation, receiver authorization, and sender authorization.*
- *Distributed multicast address allocation.*
- *Security, including protection against attacks on multicast routes and sessions, as well as support for data integrity mechanisms.*
- *Support for network management.*

Providing security is still one of the main challenges that hinder the introduction of multicast in the existing infrastructure. It is therefore critical to provide sound security mechanisms that can protect the ISPs and the carriers against multicast threats and allow them at the same time to get all the benefits of introducing multicast in their infrastructure. More specifically, two complementary levels of security must be considered:

- *The application level security:*
which is essentially the concern of the end users and the content providers. Securing the multicast data content is commonly made possible by encryption methods. A group key is shared by the members that belong to the group and this key needs to change every time a member joins (leaves) the group for backward access

control (forward access control). Application level security has been intensively studied [60] [55][77] [17], in particular, within the MSEC IETF working group [2].

- *The multicast routing infrastructure security:*

Securing the routing infrastructure have been the goal of many researches. We can divide these approaches into three broad categories: we classify these approaches into three broad categories:

- *Attack avoidance approaches:* attacks avoidance is performed by controlling the ability of entities (i.e. routers, receivers and senders) to take part in the multicast routing tree for a given group.
- *Attack resiliency approaches:*the attacks resiliency techniques aim to survive those attacks that could not be completely repelled.
- *hybrid approaches:* in this category we find the works whereby the access control and the filtering techniques are combined.

The Network Operator (Carrier), is the main actor in this thesis. This entity owns, manages and deploys the IP infrastructure (i.e. the routers, the Network Access Server/Broadband Access Server (NAS/BAS), the Digital Subscriber Line Multiplexer (DSLAMs), etc.). The Network Operator runs the IP multicast routing protocols. It also carries out the multicast peering relationships with other IP multicast Network Operators (e.g. to exchange the active source/group pairs in Multicast Domains with MSDP). The Network Operator provides IP multicast access to the end user. More specifically, the multicast operator infrastructure can be divided into two parts, the core and the edge. The core contains all the multicast routers and servers, and multicast routing protocols are used to construct the multicast distribution trees. The edge contains multicast-enabled edge nodes (edge routers and/or NAS/BAS) and hosts, and the IGMP protocol is used to communicate the group membership subscriptions. In this thesis we focus on the security of multicast from the network operator viewpoint.

1.2 Goals of the Thesis

This work considers the security from the multicast Network Operator's viewpoint. The kind of security required by a network operator, who manages and operates the multicast routing infrastructure, largely differs from that of end-to-end security.

While many theoretically ideal proposals have been done to secure the routing protocols, they have rarely been accepted by the operators community. For instance, because they require to modify existing and widely deployed protocols, or they introduce authentication mechanisms, which is in practice almost impossible to deploy in legacy networks, and even infeasible, since a corrupted host may be the source of a DoS attack, even if it has been authenticated.

The Network Operator's security point of view for its group communication service can be summarized as follows:

1. AN OVERVIEW

Secondly, the group communication service provided to its clients (i.e. end users or other network operators with whom has peering relationships) must be operational at any time, in spite of anomalies in the multicast flows, no matter whether they are intentional (i.e. are the result of deliberate attacks) or not (e.g. are caused by a misbehaving component). Security is not the goal, but a mean to achieve the network operator's "continuation of service no matter what happens" goal. Security should not impact too much the unicast and multicast forwarding performances on the operator's network.

The Network Operator can have additional requirements. For instance, it may want to guarantee that the traffic exchanged is not altered while traveling on its own routing infrastructure. It may also want to ensure some confidentiality of the traffic, in case an attacker eavesdrops on a link or a subverted router. These additional security considerations are more commonly addressed by end-to-end security.

Similarly, the physical failures affecting the links, the routers or the servers are not considered in this thesis either, even if they will affect the "continuation of service" requirement. Note that our definition shares some similarities with the notion of "network survivability" [79], but this latter is broader since it also considers physical failures.

The goal of this thesis is to analyze in depth the threats to the multicast infrastructure, to understand the requirements of the network operator and to propose a realistic solution that can respond to these requirements and relieve the networks during a possible attacks. More specifically, this thesis has the following objectives:

1. *Mitigating multicast specific attacks through a cost effective, scalable and transparent mechanism. This mechanism aims, in particular, to protect the well-behaved clients against ill-behaved ones*
2. *Keeping the changes to the currently used multicast routing as minimum as possible*
3. *Avoiding non realistic assumptions and techniques. For instance, a cryptographic authentication of each control packet (e.g. IGMP/MLD) sent by clients is not realistic (section 4.4.1, and as we will see, does not prevent attacks (section 6). This objective is close related to the point "2" which states that changes must be kept to a minimum.*
4. *Keeping track of routers and network resources at the network layer and allows a fair resource regulation.*

1.3 Dissertation Roadmap

The rest of the document is organized as follows:

- In chapter 2, we put in light the deployment aspects of multicast by the network operator.

1.3 Dissertation Roadmap

- In chapter 3, we present the different components and protocols of the group communication service.
- In chapter 4, we introduce and discuss the vulnerabilities of the multicast routing infrastructure from the point of view of the network operator. We classify the multicast specific threats into several categories according to their harmfulness to the operator network.
- In chapter 5, we present the group management protocols specific threats. We classify these attacks according to their potential impact on the network operator infrastructure.
- In chapter 6, we present the state of the art in the field of multicast routing security. We present a taxonomy of the different propositions, and we discuss the pros and cons of the different categories of solutions.
- In chapter 7, we introduce our filtering proposal, its architecture and the different building blocks involved in it. We also discuss some parameterizing and performance aspects.
- In chapter 8, we present a theoretical study that allows to easily dimension the filter and to find the different initializing parameters as a function of the traffic arrival models and the underlying infrastructure.
- In chapter 9, we evaluate experimentally the efficiency of our proposal. We compare the experimental and the theoretical results presented in the chapter 5. Then, we conclude with a detailed discussion of several key issues that emerged from this study.
- In chapter 10, we explore the deployment of the filter in the operational networks then we discuss some extensions to improve the filtering efficiency.
- In chapter 11, we conclude the dissertation by revisiting the lessons learned from this work and presenting direction for future researches in this area.

1. AN OVERVIEW

Chapter 2

Multicast Deployment by the Network Operator: a State of the Art

Contents

2.1 Multicast Delivery: the Operational Model	7
2.1.1 The Various Actors and the Services Provided	7
2.2 Deployment Examples	9
2.2.1 First Example: Commercial Delivery of Video Over xDSL . . .	10
2.2.2 Second Example: Multicast Delivery of Free Content	11
2.3 Conclusion	11

In this chapter we present an operational overview of the multicast deployment in the network operator. Understanding the different components in multicast content delivery is essential to understand the security issues related to the multicast service. Moreover, Different deployment architecture have different repercussions on the security of the operator infrastructure or the subscribers.

2.1 Multicast Delivery: the Operational Model

The operational model describes the different actors participating in the multicast delivery service and the different transactions between them. Our goal of introducing the operational model is to understand the role of the multicast operator and its relationships with the other actors.

2.1.1 The Various Actors and the Services Provided

Providing a multicast access service to receivers scattered over wide area networks requires the collaboration of different actors, which have relevant relationships and

2. MULTICAST DEPLOYMENT BY THE NETWORK OPERATOR: A STATE OF THE ART

their own set of roles.

In this thesis we consider the following entities (Figure 2.1):

- *The Network Operator (Carrier)*: This entity owns, manages and deploys the routing infrastructure (i.e. the routers, the Network Access Server/Broadband Access Server (NAS/BAS), the Digital Subscriber Line Multiplexers (DSLAMs), etc.). The operator network is composed from the following networks:
 - *Access Network*: The access network encompasses the xDSL lines and the DSLAMs. The xDSL connect the xDSL modems in the CPE to the DSLAMs. The xDSL modem can operate as a *bridge* or a *router*. In the bridge mode a Point-to-Point Protocol (PPP) [35] connection is established between the CPE and the BRAS, whereas in the router mode the xDSL modem establishes a direct IP connection. PPP allows to transport IP traffic over point-to-point links. PPP also established a standard for the assignment and management of IP addresses. Another way to assign IP address is the use of DHCP and DHCP relay [73]. Two variants of the PPP protocol are deployed: PPP over Ethernet (PPPoE) [35] and PPP over ATM (PPPoA) [53]. Additionally, there is a combination of these protocols called PPP over Ethernet over ATM or PPPoEoA.

The *DSLAM* performs the following functions:

- * Line Termination of the ADSL subscriber lines.
- * Concentration/multiplexing of the ADSL subscriber lines toward the aggregation network.
- * Termination of customer ATM signaling Channels (in the case of ATM based network)

The access network, as the gatekeeper for the multicast group, has a larger set of responsibilities during the IGMP processing [67].

- *Aggregation Network*: The aggregation network encompasses the Broadband Remote Access Server (BRAS). The BRAS is the aggregation point for the subscriber traffic. It provides aggregation capabilities (e.g. IP, PPP, ATM) between the Regional/Access Network and the ISP. Beyond aggregation, it is also responsible for policy management and IP QoS in the Regional/Access Networks.
- *The Core Network*: The core networks follows a continuous evolving process. Today's core network has four layers: IP and other content-bearing traffic, ATM for traffic engineering, a Synchronous Optical Network/Synchronous Digital Hierarchy (SONET/SDH) transport network, and dense wave division multiplexing (DWDM) for fiber capacity.

The Network Operator runs the IP multicast routing protocols. It also carries out the multicast peering relationships with other IP multicast Network Operators (e.g. to exchange the active source/group pairs in Multicast Domains with MSDP.

The Network Operator provides IP multicast access to the end user but it does not have commercial relationships with them. It neither knows the user identity nor manages the users database.

- *The Internet Service Provider (ISP) and the Service Provider:* The ISP is in charge of commercial relationships to the final user. It provides Internet access to the end user as well as other basic services (e.g. E.mail, web, chat, portal).

The Service Provider mainly focuses on the provisioning and marketing of service offers (digital TV broadcasting, VoD, ...) to the end user, and gives them a real-time access to valuable contents.

Both providers enable a user to subscribe or unsubscribe to a service. These providers own the user database: user login and password, access right, and charging. They work in close relationship with the network operator but do not manage the infrastructure themselves. In the remaining of the thesis, in order to simplify the discussion and also because we expect it will usually be the case, *we assume that the same company plays the two roles*, and we use the generic term "Service Provider".

- *The Content Provider:* This entity provides the information included in the service provided to users or via copyright to service providers. These added-value information can consist of text, images, video (digital TV or VoD), audio (radio, music), etc.
- *The Content Aggregator:* This entity builds bundles of channels from the content provided by one or more Content Providers. It collects the contents, turns them into the appropriate digital coding and broadcasts them on its IP network. The content aggregator has peering relationships and contracts with one or more Network Operators so as to forward its traffic up to the end users.
- *The End User:* The user owns an IP multimedia terminal, which can be either a PC or a TV/STB (Set Top Box).

One can notice that a same actor could play the role of several entities. For instance, a service provider could also be the network operator if it owns the network infrastructure. On the contrary, within the same enterprise, different entities could play different roles with internal pseudo commercial relationships. Additionally, a multicast operator could play only a transit role domain for the multicast traffic of other operators of ISPs.

2.2 Deployment Examples

The previous discussion is illustrated with two major case studies:

- A commercial delivery of video over xDSL, and
- A multicast delivery of free content.

2. MULTICAST DEPLOYMENT BY THE NETWORK OPERATOR: A STATE OF THE ART




Figure 2.1: The different actors during a multicast content delivery session.

These two case studies are largely different in their assumptions. We believe that they perfectly illustrate the two classes of situations that can be found: An access to a commercial versus a free service.

2.2.1 First Example: Commercial Delivery of Video Over xDSL

In this section we describe a commercial multicast video delivery service over xDSL (i.e. ADSL, VDSL, etc.). xDSL provides an attractive platform for the delivery of services requiring ever-greater bandwidths like voice and video delivery.

Practically, in this type of service the service provider offers Internet broadband access service to the client. In this context, he performs several functions such as IP address configuration and user authentication, authorization and accounting. This is the *first level of client authentication/authorization and it enables a user to get IP access and connectivity*.

Some content, like basic TV programs already available on hertzian channels, will probably continue to be provided free of charge over xDSL. But commercial services like pay-per-view or premium TV programs must be controlled. Service usage should be monitored and logged in an Authentication Server like a RADIUS (Remote Access Dial In User Service) server. This is the *second level of client authentication/authorization and it enables a user to get access to the service*.

Accessing a video channel requires that the user joins the associated multicast channel. The IGMP with IPv4 or MLD with IPv6 is used between the User and the Network Operator.

Management flows, such as AAA flows to authenticate and charge a user, are exchanged (or relayed) between the Network Operator and the Service Provider. In order to deliver the TV channels or videos to the end user, the Service Provider relies on the IP multicast transport service provided by the Network Operator. The Network Operator owns the IP network infrastructure and runs the IP multicast routing protocols so as to build multicast distribution trees from the content sources located at the Content

Aggregator's network to the end users. Therefore, a technical and physical relationship exists between the Network Operator and the Content Aggregator: Their IP multicast enabled networks are interconnected and Network Operator and Content Aggregator have multicast peering relationships so as to exchange the active source/group pairs in each Multicast Domains (through MSDP sessions for example). Therefore, content flows are transmitted from the Content Aggregator network, then transit in the Network Operator's network to get to the end user without going through the Service Provider premises. The only flows transiting between the User and the Service Provider (through the Network Operator's network) are the management flows such as AAA flows.

2.2.2 Second Example: Multicast Delivery of Free Content

The situation is completely different here, even if the access network can also be xDSL. This difference lies in the fact that clients that join the session to download the content (e.g. a software update, a video clip, etc.) have not previously subscribed to any service. They cannot be authorized since there is a free access to the content (said differently, the content provider needs no AAA server). *Client authentication/authorization is only performed at the Service Provider level* in this scenario, when the client gets connected to the Internet.

2.3 Conclusion

The purpose of this chapter is to only make an introduction to deployment of the multicast in the operational networks. As has been shown, providing a multicast access service to receivers scattered over wide area networks requires the collaboration of different actors, which have relevant relationships and their own set of roles. This thesis focuses on the network operator (Carrier) who owns, manages and deploys the routing infrastructure. The operator has a large set of responsibilities during the multicast groups management processing.

2. MULTICAST DEPLOYMENT BY THE NETWORK OPERATOR: A STATE OF THE ART

Chapter 3

Multicast Routing Protocols: Overview

Contents

3.1	Multicast Service Models	13
3.2	Group Management Protocols	14
3.2.1	Internet Group Management Protocol (IGMP)	14
3.2.2	Multicast Listener Protocol (MLD)	22
3.2.3	Group Management Protocols Proxying	22
3.2.4	Group Management Protocols in Layer2 Switching environment	22
3.3	Multicast Routing Protocols	23
3.3.1	Intra-Domain Routing Protocols	23
3.3.2	Inter-domain multicast routing architectures and protocols .	26
3.4	Consluion	28

In this chapter we present the different components and protocols of the group communication service.

3.1 Multicast Service Models

Currently, three different service models for multicast exist. The initial multicast as has been introduced by Deering and referred to as *Any Source Multicast (ASM)* [24], the *Source-Specific Multicast (SSM)* a new service model and the *Source-Filtered Multicast (SFM)*

- *Any-Source Multicast (ASM)*: With this model, an IP datagram is transmitted to a "host group", a set of zero or more end-hosts (or routers) identified by a single IP destination address taken from the class D space (224.0.0.0-239.255.255.255). ASM is an open model: any host can join any multicast group as a listener, and similarly, any host can start to transmit multicast traffic to a multicast group.

3. MULTICAST ROUTING PROTOCOLS: OVERVIEW

- *Source-Specific Multicast (SSM)*: This is the multicast service model defined in [41]. SSM provides host applications with a "channel" abstraction, in which each channel has exactly one source and any number of receivers. SSM is derived from earlier work in EXPRESS [42]. An IP datagram is transmitted by a source S to a SSM destination address G, and receivers can receive this datagram by subscribing to channel (S,G). The address range 232/8 has been assigned by IANA for SSM service in IPv4. For IPv6, the range FF3x::/96 is defined for SSM services.
- *Source-Filtered Multicast (SFM)*: This is a variant of the ASM service model, and uses the same address range as ASM. It extends the ASM service model by introducing sources filtering. That is, each "upper layer protocol module" can now request data sent to a host group G by only a specific set of sources, or can request data sent to host group G from all BUT a specific set of sources. Support for source filtering is provided by version 3 of the Internet Group Management Protocol (or IGMPv3) for IPv4, and version 2 of the Multicast Listener Discovery (or MLDv2) protocol for IPv6. Note that while SFM is a different model than ASM from a receiver standpoint, there is no distinction between the two for a sender.

3.2 Group Management Protocols

3.2.1 Internet Group Management Protocol (IGMP)

The Internet Group Management Protocol (IGMP) is used by IPv4 end hosts to report their group membership interests to their immediately neighboring multicast routers. Routers periodically query their attached subnetworks to determine if known group members are still active. IGMP is only locally significant; IGMP packets are not forwarded by routers (note that in the case of IGMP proxy the IGMP packets are generated by the router on behalf of the receivers and not forwarded). Based on the group membership information learned from IGMP, a router is able to determine which (if any) multicast traffic needs to be forwarded to each of its "leaf" subnetworks. Multicast routers use this information, in conjunction with a multicast routing protocol, to support IP multicasting across the Internet. An IP multicast router may itself be a member of one or more multicast groups, in which case it performs both the "multicast router part" of the protocol (to collect the membership information needed by its multicast routing protocol) and the "group member part" of the protocol (to inform itself and other, neighboring multicast routers of its memberships).

Three versions of IGMP has been defined. Version 1, specified in RFC-1112 [24], was the first widely-deployed version and the first version to become an Internet Standard. Version 2 [29] adds support for "low leave latency", that is, a reduction in the time it takes for a multicast router to learn that there are no longer any members of a particular group present on an attached network. Version 3 [16] adds support for "source filtering", that is, the ability for a system to report interest in receiving packets only from specific

source addresses, or from all but specific source addresses, sent to a particular multicast address. Version 3 is designed to be interoperable with Versions 1 and 2, Version 2 to be interoperable with Version 1.

IGMP Version 1

IGMP Version 1 was specified in RFC-1112 [24]. According to the specification, a host sends a *Group Membership Report* message to receive the flows destined for a multicast group. The report messages are sent with the address of the multicast group with a TTL of 1.

Multicast routers periodically send *Membership Query* messages to determine which groups have members on their directly attached networks. *Query* messages are addressed to the all-hosts group (224.0.0.1) and have an IP TTL of 1.

When a host receives a *Query* message, it responds with a *Host Membership Report* for each host group to which it belongs. In order to avoid a storm of *Reports*, each host starts a randomly chosen delay timer for each of its group memberships. If, during the delay period, another *Report* is heard for the same group, the local host resets its timer to a new random value. Otherwise, the host transmits a *Report* to the reported group address, causing all other members of the group to reset their *Report* message timers. This procedure guarantees that Reports are spread out over a period of time and that Report traffic is minimized for each group with at least one member on the subnetwork.

When a host first joins a group, it immediately transmits a *Report* for the group rather than waiting for a router *Query*. This guarantees that the host will receive traffic addressed to the group if it is the first member of that group on the subnetwork.

It should be noted that a router does not need to maintain a detailed list of which hosts belong to each multicast group; the router only needs to know that at least one group member is present on a network interface. Multicast routers periodically transmit Queries to update their knowledge of the group members present on each network interface.

If the router does not receive a *Report* for a particular group after a number of *Queries* (by default this number equals 2), the router assumes that group members are no longer present on the interface and the group is removed from the list of group memberships for the directly attached subnetwork. This report suppression mechanism is depicted in the figure 3.1:

IGMP Version 1 Packet Format:

IGMPv1 packet is encapsulated in an IP packet with a protocol identifier of 2. Figure 3.2 depicts the packet format for IGMPv1.

IGMP Version 2

IGMP Version 2 [29] enhances and extends IGMP Version 1 while also providing backward compatibility with Version 1 hosts. IGMP Version 2 defines a procedure for the

3. MULTICAST ROUTING PROTOCOLS: OVERVIEW




Figure 3.1: The Reports Suppression Mechanism.




Figure 3.2: IGMPv1 packet format.

election of the multicast *Querier* for each LAN. In IGMP Version 2, the router with the lowest IP address on the LAN is elected the multicast querier. In IGMP Version 1, the *Querier* election was determined by the multicast routing protocol. This could lead to potential problems because each multicast routing protocol used different methods for determining the multicast querier.

IGMP Version 2 defines a new type of *Query* message: the *Group-Specific Query* message. *Group-Specific Query* messages allow a router to transmit a *Query* to a specific multicast group rather than all groups residing on a directly attached subnetwork.

Finally, IGMP Version 2 defines a *Leave Group* message to reduce IGMP's "leave latency". When the last host to respond to a *Query* with a *Report* wishes to leave that specific group, the host transmits a *Leave Group* message to the all-routers group (224.0.0.2) with the group field set to the group to be left. In response to a *Leave* group message, the router begins the transmission of *Group-Specific Query* messages on the interface that received the *Leave* message. If there are no *Reports* in response to the *Group-Specific Query* messages, the group is removed from the list of group memberships for the directly attached subnetwork.

IGMP Version 2 Packet Format:

IGMPv1 packet is encapsulated in an IP packet with a protocol identifier of 2. Figure 3.3 depicts the packet format for IGMPv2.

0x1	Membership Query
0x2	Membership Report

Table 3.1: IGMPv1 message types




Figure 3.3: IGMPv2 packet format.

IGMP Version 2 Timers:

IGMPv2 introduces many timers, these timers are configurable and allow to tune the burstiness of IGMP traffic in the network.

- *Robustness Variable*: The Robustness Variable allows tuning for the expected packet loss on a subnet. If a subnet is expected to be lossy, the Robustness Variable may be increased. IGMP is robust to (Robustness Variable-1) packet losses.
- *Query Interval*: The Query Interval is the interval between General Queries sent by the Querier. Default: 125 seconds.
- *Query Response Interval*: The Max Response Time inserted into the periodic General Queries. Default: 10 seconds.
- *Group Membership Interval*: The Group Membership Interval is the amount of time that must pass before a multicast router decides there are no more members of a group on a network. Group Membership Interval = ((the Robustness Variable) times (the Query Interval)) plus (one Query Response Interval).
- *Other Querier Present Interval*: The Other Querier Present Interval is the length of time that must pass before a multicast router decides that there is no longer another multicast router which should be the querier. This value equals to ((the Robustness Variable) times (the Query Interval)) plus (one half of one Query Response Interval).
- *Startup Query Interval*: The Startup Query Interval is the interval between General Queries sent by a Querier on startup. Default: 1/4 the Query Interval.
- *Last Member Query Interval*: The Last Member Query Interval is the Max Response Time inserted into Group-Specific Queries sent in response to Leave Group

3. MULTICAST ROUTING PROTOCOLS: OVERVIEW

0x11	Membership Query
0x16	Version 2 Membership Report
0x12	Version 1 Membership Report
0x17	Version 2 Leave Group

Table 3.2: IGMPv2 messages types

messages, and is also the amount of time between Group-Specific Query messages. Default: 1 second.

- *Unsolicited Report Interval*: The Unsolicited Report Interval is the time between repetitions of a host's initial report of membership in a group. Default: 10 seconds.
- *Version 1 Router Present Timeout*: The Version 1 Router Present Timeout is how long a host must wait after hearing a Version 1 Query before it may send any IGMPv2 messages. Value: 400 seconds.

IGMP Version 3

IGMP Version 3 [16] introduces support for Group-Source Report messages so that a host can request to receive traffic from only specific source addresses or from all but specific source addresses. IGMP Version 3 will help conserve bandwidth by allowing a host to select the specific sources from which it wants to receive traffic. Also, multicast routing protocols will be able to make use of this information to conserve bandwidth when constructing the branches of their multicast delivery trees. In order to enable this source access filtering mechanism, a system's IP service interface must support the following operation: `IPMulticastListen` (socket, interface, multicast-address, filter-mode, source-list) where:

- *Interface*: the network interface on which reception of the specified multicast address is to be enabled/disabled.
- *Multicast address*: the IP multicast address to which the request pertains.
- *Filter-mode*: this parameter can be either `INCLUDE` or `EXCLUDE`. In `INCLUDE` mode, reception of packets sent to the specified multicast address is requested only from those IP source addresses listed in the source-list parameter. In `EXCLUDE` mode, reception of packets sent to the given multicast address is requested from all IP source addresses except those listed in the source-list parameter.
- *Source-list*: An unordered list of zero or more IP unicast addresses from which multicast reception is desired or not desired, depending on the filter mode.

IGMP Version 3 provides support for *Group-Source Leave* messages. This feature allows a host to leave an entire group or to specify the specific IP address(es) of the (source, group) pair(s) that it wishes to leave.

Version 3 Membership Report Message:

IGMP Version 3 Membership Reports are sent by IP systems to report (to neighboring routers) the current multicast reception state, or changes in the multicast reception state, of their interfaces. Reports have the following format 3.4:

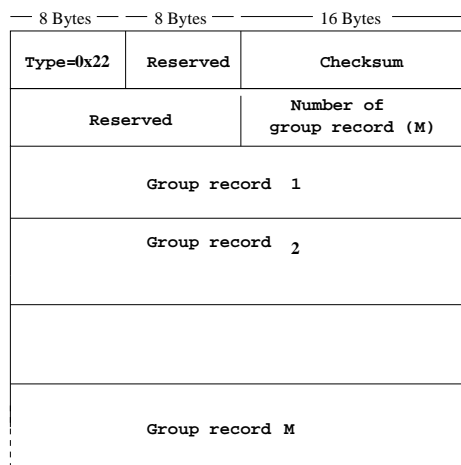


Figure 3.4: IGMPv3 packet format.

0x11	Membership Query
0x22	Version 3 Membership Report
0x12	Version 1 Membership Report
0x16	Version 2 Membership Report
0x17	Version 2 Leave Group

Table 3.3: IGMPv3 messages types

where each Group Record has the following format 3.5:

The most interesting fields are:

- **Group Record Types** There are a number of different types of Group Records that may be included in a *Report* message:
 - **Current-State Record:** This indicates the current reception state with respect to one multicast group at a given interface. It contains the filter mode (include or exclude) and the set of related sources. A Current-State record type is either *MODE_IS_INCLUDE* or *MODE_IS_EXCLUDE*.

3. MULTICAST ROUTING PROTOCOLS: OVERVIEW

8 Bytes	8 Bytes	16 Bytes
Record Type	Aux Data Length	Number of Sources(N)
Multicast Address		
Source Address 1		
Source Address 2		
.....		
Source Address N		
Auxiliary Data		

Figure 3.5: IGMPv3 group record.

- Filter-Mode-Change Record: This indicates that the filter-mode of the reception state has changed. It contains the new filter-mode and the set of related sources. A Filter-Mode-Change record type is either *CHANGE_TO_INCLUDE_MODE* or *CHANGE_TO_EXCLUDE_MODE*.
- Source-List-Change Record: This indicates that the group’s associated sources have changed. A Source-List-Change record type is *ALLOW_NEW_SOURCES*, when data from a new set of sources are to be received. It is *BLOCK_OLD_SOURCES*, when data from an existing set of sources are not required.

- The Number of Group Records (M) field which specifies how many Group Records are present in this Report. Each Group Record is a block of fields containing information pertaining to the sender’s membership in a single multicast group on the interface from which the Report is sent.
- The Number of Sources (N) field specifies how many source addresses are present in this Group Record.
- The Multicast Address field contains the IP multicast address to which this Group Record pertains.

IGMP Version 3 Reports are sent with an IP destination address of 224.0.0.22, to which all IGMPv3-capable multicast routers listen. Since in IGMPv3 these packets are not addressed to the group joined (as in the previous versions) the router will effectively learn all hosts desiring to receive a given multicast group. This enables new enhancements in terms of expedited leave and simplified IGMP snooping architectures. Another feature added in IGMP v3 is that when a router sends, a query all hosts must respond with their group of interest to the address 224.0.0.22.

Version 3 Membership Query Message:

Membership Queries 3.6 are sent by IP multicast routers to query the multicast reception state of neighboring interfaces.

The functionality of an IGMPv3 querier is complex compared to the earlier versions (v1 or v2). An IGMPv3 querier has to handle the six types of group record information (discussed above) and maintain the state information for each reachable multicast group associated with each network interface. The state information consists of the following: multicast address, group timer, filter mode, and source records. Each source record contains source address and a source timer.

The filter mode for a multicast group in a querier is known as router filter mode. It is determined after processing all the received state record information from the neighboring IGMPv3 hosts for the group. When the filter mode is INCLUDE, the source record list contains the list of sources whose data are to be forwarded on the attached network. When the filter is EXCLUDE, the source record list contains two types of sources.

The first set (type) contains sources whose data needs to be forwarded by some routers while the second set contains sources whose data are not to be forwarded. The group timer is used when the filter mode is exclude.

When a Current-State Record with record type *MODE_IS_EXCLUDE* is received, the router filter for the group becomes exclude. When the host that had reported a Current-State Record with *MODE_IS_EXCLUDE* stops reporting and the group timer expires, the router filter changes to include.

Queries have the following format: The most interesting fields (IGMPv3 specific) are:

Record Type	Reserved	Checksum
Reserved		Number of Sources (N)
Multicast Address		
Source Address 1		
Source Address 2		
.....		
Source Address (N)		

Figure 3.6: IGMPv3 query.

- **Group Address:** this field is set to zero when sending a General Query, and set to the IP multicast address being queried when sending a Group-Specific Query or Group-and-Source-Specific Query.
- **Number of Sources (N):** this field specifies how many source addresses are present in the Query. This number is zero in a General Query or a Group-Specific Query, and non-zero in a Group-and-Source-Specific Query.
- **Source Address [i]:** these fields are a vector of n IP unicast addresses, where n is the value in the Number of Sources (N) field. In IGMPv3, General Queries

3. MULTICAST ROUTING PROTOCOLS: OVERVIEW

are sent with an IP destination address of 224.0.0.1, the all-systems multicast address. Group and Group-and-Source Queries are sent with an IP destination address equal to the multicast address of interest.

3.2.2 Multicast Listener Protocol (MLD)

MLD works with IPv6 in the same way that IGMP manages multicast groups for IPv4. MLD version 1 [25] implements the functionality of IGMP version 2, whereas MLD version 2 [25] implements the functionality of IGMP version 3. MLD is a sub-protocol of ICMPv6, that is, MLD message types are a subset of the set of ICMPv6 messages, and MLD messages are identified in IPv6 packets by a preceding Next Header value of 58. All MLD messages are sent with a link-local IPv6 Source Address, an IPv6 Hop Limit of 1.

3.2.3 Group Management Protocols Proxying

In some deployment scenarios, such in the case of an edge box with only one connection to the core network side and many connections to the customer side, it is not necessary to run a multicast routing protocol. It is sufficient to learn and proxy group membership information and simply forward multicast packets based upon that information. One typical example of such tree topology can be found on an edge aggregation box such as Digital Subscriber Line Access Multiplexer (DSLAM) [25].

3.2.4 Group Management Protocols in Layer2 Switching environment

Layer 2 switches cannot handle multicast MAC address. By default, when seeing a multicast MAC address the switch will forward the packet out of ALL its ports. This loads the switch up and results in multicast traffic being unnecessarily sent to hosts that have not requested it. To resolve this issue, the following approaches have been proposed. These approaches allow the switch to determine which interfaces should be used for egress processing. Therefore, reduce drastically the amount of unwanted multicast traffic.

- IGMP/MLD Snooping: A layer-2 switch supporting "IGMP/MLD snooping" can snoop on IGMP *Query, Report and Leave* (IGMP version 2) packets transferred between IP Multicast Routers/Switches and IP Multicast hosts to learn the IP Multicast group membership. It checks IGMP/MLD packets passing through it, picks out the group registration information, and configures multicast forwarding tables accordingly. However, without some type of hardware (ASIC) assist, this additional decode process can be quite taxing on a central CPU-based switch. In fact, IGMP/MLD snooping may cause such switches to arbitrarily discard a large number of packets during times of multicast peaks. However, with the proper hardware assist, IGMP/MLD snooping is a viable solution. IGMP/MLD snooping is available on a variety of mid-to-high end switches.

- Group Address Resolution Protocol (GARP) [71]: This is a IEEE approach. The primary purpose of GARP is to maintain VLAN group information. GARP can be extended to also provide (S,G) lists that allow the switch to map multicast groups to egress ports in much the same way that a VLAN is a list of MAC addresses that belong to a specific broadcast domain.
- Cisco's Group Management Protocol (CGMP) [1]: CGMP is proprietary to Cisco and involves a router-to-switch multicast-group information exchange protocol. A mid-to-high end Cisco switch can receive multicast-group *Join/Leave* messages from a multicast-enabled Cisco router. These Join/Leave updates are then used by the switching logic to provide better multicast filtering through the switch fabric.
- The Router-Port Group Management Protocol (RGMP) [78]: This is also a Cisco proprietary layer 2 protocol. RGMP enables a router to communicate to a switch (or a networking device that is functioning as a Layer 2 switch) the multicast group for which the router would like to receive or forward traffic. RGMP restricts multicast traffic at the ports of RGMP-enabled switches that lead to interfaces of RGMP-enabled routers.

3.3 Multicast Routing Protocols

3.3.1 Intra-Domain Routing Protocols

Although a number of protocols such as PIM-DM [8], CBT [12], DVMRP [74], etc. have been proposed to construct multicast tree inside the same administrative domain, PIM-SM is the most widely used protocol. Version 1 of PIM-SM was created in 1995, but was never standardized by the IETF. It is now considered obsolete, though it is still supported by Cisco and Juniper routers. Version 2 of PIM-SM was standardized in RFC 2117 (in 1997) and updated by RFC 2362 (in 1998). Version 2 is significantly different from and incompatible with version 1. However, there were a number of problems with RFC 2362, and a new specification of PIM-SM version 2 is currently being produced by the IETF. In the following we give an overview of the PIM-SM protocol [27].

Overview of the Protocol Independent Multicast (PIM-SM)

The Protocol-Independent Multicast (PIM) routing protocol is being developed to provide a multicast routing protocol between members of sparsely distributed groups. PIM-SM like other multicast routing protocols needs the activation of IGMP/MLD within LANs. Thanks to IGMP/MLD, hosts directly connected to LANs may indicate their current group memberships to DR routers, that is, routers designated to act as local centralized registration points. DR routers, in turn, send *Join/Prune* messages towards Rendez-vous Point (RP) routers in order to establish and/or possibly prune branches of delivery trees.

3. MULTICAST ROUTING PROTOCOLS: OVERVIEW

PIM-SM *by default uses shared trees*, which are multicast distribution trees rooted at some selected node (in PIM, this router is called the Rendezvous Point, or RP) and used by all sources sending to the multicast group.

PIM-SM also *supports the use of source-based trees*, in which a separate multicast distribution tree is built for each source sending data to a multicast group. Each tree is rooted at a router adjacent to the source, and sources send data directly to the root of the tree. Source-based trees enable the use of source addresses filtering.

PIM-SM may use source-based trees in the following circumstances.

- To avoid multicast data sent to an RP having to be encapsulated, the RP may join a source-based tree.
- To optimize the data path, a last-hop router may choose to switch from the shared tree to a source-based tree.
- For SSM, a last-hop router will join a source-based tree from the outset.

PIM-SM is a *soft-state protocol*. That is, all state is timed-out a while after receiving the control message that instantiated it. To keep the state alive, all PIM *Join* messages are periodically (by default each 60s) retransmitted.

PIM is not dependent on a specific underlying unicast routing protocol, hence it uses the traditional IP multicast model of receive-initiated membership.

The reverse path check (RPF) check is a mandatory condition, upon which multicast flows are forwarded one hop further through all interfaces specified in the Outgoing Interfaces (OIF) list of the corresponding route entry. The RPF control consists in verifying that the multicast flow has been received on its RPF interface. That is, its next hop router neighbor interface, along the path towards the RP.

The PIM-SM protocol architecture is based on the following parts:

- *Hello messaging*: PIM routers periodically send *Hello* messages to discover neighboring PIM routers. *Hello* messages are multicast using the address 224.0.0.13 (ALL-PIM-ROUTERS group).
- *Joining the shared tree*: When a DR gets a membership indication from IGMP for a new group G, it looks up the associated RP. Then it creates a new state (*,G) in its multicast FIB and sends a *Join/Prune* message towards the RP router. Hop by hop, this *Join/Prune* message is processed by every router along the path to the RP router, resulting in a new branch (due to the Join part of the *Join/Prune* message) for the corresponding multicast delivery tree.

The RP address is included in a special field in the route entry and is included in periodic upstream *Join/Prune* messages. The outgoing interface is set to that included in the IGMP membership indication for the new member. The incoming interface is set to the interface used to send unicast packets to the RP.

When there are no longer directly connected members for the group, IGMP notifies the DR. If the DR has neither local members nor downstream receivers, the (*,G) state is deleted.




Figure 3.7: PIM-SM: Building a new branch of the forwarding tree.

- *RP discovery*: Within a PIM domain, there is a particular router, called the BSR router, which is in charge of collecting and dispatching updated information about the set of active RP routers. Thanks to *Candidate-RP-Advertisement* messages, the BSR router may determine the set of active RP routers which to send advertisements for. That is carried out by means of broadcasting BSR messages all over the domain. Therefore, DR routers hold knowledge of the mapping of RP routers to multicast address scopes.
- *Registering with the RP*: When an active source for group G starts sending multicast datagrams, its DR begins a process for Registering this source with the corresponding RP and requesting the RP to build a tree back to that router. The source router encapsulates the multicast data from the source in a special PIM-SM message called the *Register* message and unicasts the data to the RP. When the RP receives the *Register* message, it decapsulates the multicast data packet inside of the *Register* message and forwards it down the Shared Tree. The RP also sends series of *Register-Stop* messages. The reception of these PIM control messages by the source router triggers the ending of sending multicast datagrams encapsulated within *Register* messages. Hence, multicast datagrams flow natively from the source to members going through the RP router.
- *Shortest-Path Tree (SPT) switching*: A router R with directly-connected members first joins the shared RP- tree. The router can switch to a source's shortest path tree (SP- tree) after receiving packets from that source over the shared RP-tree. The recommended policy is to initiate the switch to the SP-tree after receiving a significant number of data packets during a specified time interval from a particular source. To realize this policy the router can monitor data packets from sources for which it has no source-specific multicast route entry and initiate such an entry when the data rate exceeds the configured threshold. In general, This value is set to zero. Where the default behavior for the PIM-SM leaf routers attached to active receivers is to immediate join the SPT to the source as soon

3. MULTICAST ROUTING PROTOCOLS: OVERVIEW

as the first packet arrives via the (*,G) Shared Tree.

To join the SPT the router R sends a (S,G) *Join* message. This (S,G) *Join* message travels hop-by-hop to the first-hop router (the router connected directly to the source) thereby creating another branch of the SPT. This also creates (S,G) state in all the routers along this branch of the SPT. Finally, special (S,G) RP-bit *Prune* messages are sent up the Shared Tree to prune off this (S,G) traffic from the Shared Tree. If this were not done, (S,G) traffic would continue flowing down the Shared Tree resulting in duplicate (S,G) packets arriving at the receiver. At this point (S,G) traffic is now flowing directly from the first-hop router to the last-hop router and from there to the receiver.

The RP will normally send (S,G) Prunes back towards the source to shutoff the flow of now unnecessary (S,G) traffic to the RP if it has received an (S,G) RP-bit Prune on all interfaces on the Shared Tree. (This is not shown above).

At this point, the RP no longer needs the flow of (S,G) traffic since all branches of the Shared Tree have pruned off the flow of (S,G) traffic. As a result, the RP will send (S,G) Prunes back towards the source to stop the flow of the now unnecessary (S,G) traffic to the RP. This will if the RP has received an (S,G) RP-bit Prune on all interfaces on the Shared Tree.

3.3.2 Inter-domain multicast routing architectures and protocols

The Multicast Source Discovery Protocol (MSDP)

MSDP protocol addresses the *issue of interconnecting multiple shared trees* (such as those which are based upon the PIM-SM) whether these trees belong to the same administrative entity (a single AS) or they are deployed among different autonomous systems.

A RP within a domain will have an MSDP peering relationship with an RP in another domain, which will be established over a TCP connection between the two peers. Thanks to this peering relationship, the multicast sources of a given SM domain which may present an interest for receivers which belong to another domain will be advertised to the corresponding MSDP peer, thus allowing the establishment of an *inter-domain source-tree*. The TCP connections between MSDP peers may well be established thanks to the activation of a BGP routing process, but the MSDP protocol can be activated within a single autonomous system, where multiple SM domains have been deployed.

When a PIM sparse-mode RP that is running MSDP becomes aware of a new local source, it sends *source-active* message to its MSDP peers. When a source-active message is received, a peer-reverse-path-forwarding (peer-RPF) check (not the same as a multicast RPF check) is done to make sure this peer is toward the originating RP. If not, the source-active message is dropped. This message is counted as a *"rejected"* source-active message.

In a given SM domain, when a source generates traffic towards a multicast group,




Figure 3.8: MSDP source active messages propagation.

the DR which is directly connected to the source will send a PIM *Register* message to the RP which happens to be a MSDP peer. From that perspective, the MSDP peer knows if there are active sources within the domain. On the other hand, any RP within an SM domain knows if there are receivers, thanks to the reception of *Join* messages. Thus, any RP of a given SM domain, can tell an RP in a peer domain what sources are active, therefore indicating the peering RP that it can actually join a source tree rooted at the source in the peer domain. The corresponding information is sent over the TCP connection which has been established between a pair of MSDP peers, and it consists in generating a *Source Active (SA)* message which contains the following fields:

- source address of the multicast data source;
- group address this data source sends to;
- IP address of the RP which sends the SA message.

Each MSDP peer which receives such SA messages will forward them according to what the MSDP specification calls a peer-RPF flooding kind of forwarding. This peer-RPF flooding mechanism is based upon the following procedures: the BGP routing table is examined to determine which peer is the next hop towards the originating RP of the SA message. The above-mentioned peer is called an RPF peer; if the MSDP peer receives the SA from a non-RPF peer towards the originating RP, it will drop the message. Otherwise, the message will be forwarded to all its MSDP peers. When each MSDP peer of a given SM domain receives an SA message, they determine (as classical RPs (or core routers in a CBT flavored environment)) if there are receivers/group members which may be interested in the group which is described in the SA message. If the (*, G) state of an RP exists while the oif list is not empty, then the domain is interested in the group and the RP will trigger an PIM(S, G) join towards the source. This sets up a branch of the source-tree, and the corresponding multicast datagrams which arrive at the RP (of the domain where the source does not reside) will be forwarded down the shared tree of the corresponding domain.

3. MULTICAST ROUTING PROTOCOLS: OVERVIEW

The MBGP Protocol

The Multiprotocol BGP (MBGP) [14] adds capabilities to BGP to enable multicast routing policy throughout the Internet and to connect multicast topologies within and between BGP Autonomous Systems. In other words, MBGP is an enhanced BGP that carries IP multicast routes. BGP carries two sets of routes, one set for unicast routing and one set for multicast routing. The routes associated with multicast routing are used by the Protocol Independent Multicast (PIM) to build data distribution trees. MBGP introduces two additional attributes - the *MP_REACH_NLRI* (which is used to indicate what are the destination prefixes that can be reached, whatever the address format), and the *MP_UNREACH_NLRI*, which is designed to carry a set of unreachable destinations.

The Border Gateway Multicast Protocol (BGMP)

BGMP [70] is a protocol for inter-domain multicast routing. Like CBT and PIM Sparse Mode, BGMP chooses a global root for a delivery tree. However, the root is a domain, not a single router, so if there is any path available to the domain connectivity can be maintained. BGMP builds a bidirectional, shared tree of domains. Similarly to the unicast EGP/IGP split, BGMP is used as the inter-domain or external protocol, while domains can run any multicast IGP internally (such as CBT or PIM Sparse Mode), and can build source-specific shortest-path distribution branches to supplant the shared tree where needed.

3.4 Conclusion

This chapter describes concepts and mechanisms behind several multicast routing protocols. This chapter is intended to be only the first step in understanding the different characteristic of the multicast routing protocols. Therefore, analyzing their potential vulnerabilities. This will be the context of the next chapter.

Chapter 4

Taxonomy of Possible Attacks of Interest to the Network Operator

Contents

4.1	Introduction	29
4.2	Internal Attacks	30
4.2.1	Internal data attacks	30
4.2.2	Internal control attacks	30
4.3	Edge attacks	31
4.3.1	Edge Data Attacks	31
4.3.2	Edge control attacks	33
4.4	Conclusion	33
4.4.1	Classification of Attacks According to their Importance and Likelihood	34
4.5	Conclusion	35

4.1 Introduction

In this chapter we present the threats of the multicast operator routing infrastructure. We give a taxonomy of these threats and we classify them according to their harmfulness to the network operator.

The attacks on the operator multicast infrastructure can be classified according to the following orthogonal criteria [64] [23] [40]:

- *Their origin:* We distinguish edge attacks, mounted from the edge of the network, and internal attacks, mounted from the core of the network.
- *Their target:* We distinguish attacks located in the transfer plane (or "data attacks" for short), and attacks located in the signaling plane (or "control attacks" for short) (Figure 4.1).

4. TAXONOMY OF POSSIBLE ATTACKS OF INTEREST TO THE NETWORK OPERATOR

The term "attack" must be understood in its broadest meaning, and sometimes refers to non-intentional anomalies (e.g. a defective software can cause DoS attacks) [48].




Figure 4.1: A General Classification of Possible Attacks of Interest to the Network Operator.

4.2 Internal Attacks

An internal attack by definition originates from within the multicast distribution tree of a Network Operator, either from a compromised router or a tapped-line.

4.2.1 Internal data attacks

These attacks target the data traffic. For instance, the internal attacker can *alter the content* of data packets or *inject spurious additional traffic* into the data stream. The attacker can also *copy the content of a group to another group* at the content owner expense, or *intercept information and replay it at a later time*. The attacker can issue bogus data packets that will be received by all members of a group.

If these attacks essentially concern the clients and the content providers, the multicast Network Operator is concerned too. Indeed, these attacks can cause large amounts of additional data to be forwarded in its network, wasting bandwidth resources, which may also affect other existing services.

These internal data attacks could also be internal passive attacks (i.e observation, or eavesdropping) which results mainly in the disclosure of information such as traffic type, content, frequency and presence/absence. If purely passive, this attack does not compromise the continuation of service requirement of the Network Operator, yet it affects its credibility in front of its clients. End to end encryption techniques remain the best counter measures, but this is out of the scope of the Network Operator role.

4.2.2 Internal control attacks

The multicast routing protocols have many vulnerabilities when facing a strategically placed intruder. An intruder can *create, reply to, monitor, or delete any control packet* exchanged by the routing protocols. This attack can easily compromise the multicast delivery tree creation and management process. Issuing a large number of forged control

messages can also consume router processing resources. For example, the PIM-SM protocol, is vulnerable to several attacks based on forged PIM control messages [27] [36]. An attacker who has compromised a topologically-strategic router may also be able to introduce modifications to the routing table in such a way that a considerable amount of traffic is pulled towards the router which may result in congestion problems. The situation is worse if the compromised router happens to be a strategic and sensitive router (e.g. the Rendez-vous Point (RP) in PIM-SM).

4.3 Edge attacks

These attacks exploit the vulnerabilities of the open model followed by the current IP multicast service, which does not include any access control mechanism. Two major categories of attacks exist:

- Edge Data attacks, and
- Edge Control attacks.

4.3.1 Edge Data Attacks

These attacks affect the data plane. More precisely, two sub-categories can be distinguished:

- Sender attacks, and
- Receiver attacks.

Sender attacks

The traditional ASM (Any Source Multicast) model does not define any mechanism to prevent any host (group member or not) from sending multicast data to a group. This is an easy way of creating denial of service attacks since an attacker can *inject a large number of bogus data packets to an existing group* consuming a large part of the available bandwidth. Amplified variants of this attack, that involve several hosts are also possible. Practically, a sender attack which starts in the data plane can affect rapidly the control plane: *generating multicast traffic* (a few packets is sufficient) *to a large number of different multicast groups*, no matter whether these groups exist or not. will generate many register-encapsulated packets, loading the Designated Router (DR), the RP, and each router between them. If inter-domain multicast is enabled, this attack is even worse (i.e. "Ramen Worm" attacks [61]). With the SSM (Source Specific Multicast) model, the flooding attack on an existing (S, G) channel is more difficult because only source of IP address S can join this channel. The attack therefore requires that the attacker manages to spoof the IP address of the legitimate sender. Besides a multicast routing protocol based on the Reverse-Path Forwarding (RPF) algorithm forms an implicit safeguard against masquerading, since multicast packets

4. TAXONOMY OF POSSIBLE ATTACKS OF INTEREST TO THE NETWORK OPERATOR

that arrive on an interface other than that which would be used to reach the source are automatically discarded by the multicast routers. Therefore, (1) an attacker must be on the shortest reverse-path for the authentic source from the perspective of a multicast router, and (2) must use source address spoofing for the attack to succeed.

Receiver attacks

These attacks happen when an attacker simply *join one or several groups*, causing the tree to expand and multicast traffic to be forwarded to him. Even if the content is protected thanks to enciphering mechanisms, the encrypted packets would still be forwarded to the attacker, thereby consuming bandwidth along the distribution tree. In that scheme, the attacker simply discards the encrypted packets he receives. In such an attack, the main victim is mainly the "attacking" host's network, which presumably would have prior been hacked by the real intruder. Such an attack can have heavy consequences for the Network Operator since it wastes network bandwidth resources. Its effects can here also be magnified by a distributed attack, with several slave attackers joining several groups.

A particular type of edge data attacks are *congestion control attacks*. Congestion Control Attacks are located in the transfer plane of the IP multicast model and are twofold.

- *Inflate subscription congestion control attacks*

Since a multicast traffic is of UDP/IP type, it can easily lead to router congestion. Multicast-enabled applications must therefore include a congestion control protocol, and several of these protocols are currently being standardized by the IETF: WEBRC [52] for multi-rate transmission reliable multicast protocols (e.g. (Asynchronous Layered Coding) ALC), PGMCC [63] and TFMCC for single-rate reliable multicast protocols (e.g. (Pragmatic General Multicast) PGM [68] or (NACK-Oriented Reliable Multicast Protocol) NORM [9]). Other congestion control schemes may also be used, in particular by multimedia applications (e.g. DSG [20] defines a multi-rate solution whereby the number of groups and the individual transmission rate on each of them is automatically adjusted so as to provide a maximum satisfaction to the clients). But there is a high risk that some client or source applications do not use them (e.g. a coarse application that does not implement any congestion control), or misbehave, whether deliberately or not (e.g. does not join an appropriate number of groups in multi-rate schemes, or the right group in DSG). Then the resulting unresponsive flow can easily congest routers, thereby creating another kind of DoS attack.

- *Deflate subscription congestion control attacks*

Here the attacker sends feedback messages to the source with the intention to reduce as much as possible the transmission rate of a session. This typically occurs with single rate transmission schemes (e.g. PGM or NORM reliable multicast protocols) that take into account feedback messages from a set of receivers (as in

PGMCC and TFMCC). Multi-rate transmission schemes are immune to this kind of attack since they are receiver oriented (e.g. no feedback message is possible in an ALC session using WEBRC). Yet hybrid schemes, like DSG, can be affected. While this attack does not endanger the Network Operator's infrastructure, it leads to a DoS to end users.

4.3.2 Edge control attacks

These attacks aim to increase the amount of multicast states information in the routers by exploiting the control aspects of IGMP or MLD. An attacker can therefore issue one or more IGMP *Report* messages indicating its interest in joining one or more groups. These IGMP/MLD messages trigger the multicast routing protocol and eventually cause the multicast distribution tree to be extended towards the host's network. The lack of host identification information in the IGMP/MLD subscription process makes it even harder for an edge router to identify an illegal host. As explained before this attack wastes network bandwidth resources, but also memory resources used by the routers to maintain the multicast state information and processing power used to process the control messages.

An attacker can start this DoS attack even if the group is idle (i.e. no multicast traffic is flowing) or does not exist. Indeed, current multicast routing protocols create multicast data forwarding states in all routers along the reverse path from the receiver to the sender (SSM case) or from the receiver to the Rendez-vous Point (ASM case) before the multicast packets flow into the network.

An attacker that *subscribes to thousands of multicast groups* can cause routers to create huge amounts of forwarding states that may exceed the router capabilities. Eventually a router would deny other legal service requests. This is particularly true to routers that are roots or near the roots of the multicast distribution trees since they maintain most of the multicast forwarding states. Such routers are RPs in PIM-SM, and Designated Routers (DRs) in PIM-SSM. The effects can here also be magnified by a distributed attack, with several slave attackers, driven by a master attacker, simultaneously sending thousands of IGMP/MLD *Report* messages for thousand of multicast groups.

4.4 Conclusion

We have so far identified the possible attacks and identified the various security mechanisms (Figure 4.2). In this section we discuss some complementary aspects: we first classify attacks in decreasing importance order, and then we discuss assumptions that cannot be made when considering the network operator standpoint.

4. TAXONOMY OF POSSIBLE ATTACKS OF INTEREST TO THE NETWORK OPERATOR




Figure 4.2: A General Classification of Possible Attacks of Interest to the Network Operator.

4.4.1 Classification of Attacks According to their Importance and Likeliness

The attacks need to be classified according to the threats they create on the Network Operator and the probability they occur.

1. Edge Attacks (excluding congestion control attacks): These attacks are easily launched. They are not so simple to avoid since IGMP/MLD is directly linked to the clients (e.g. generating a high number of IGMP/MLD requests with a wrong authentication can create a DoS even if IGMP is secured, because of the extra load generated by authentication). Sender attacks directed to MSDP have also shown how simple a DoS can be setup with some of the existing multicast routing protocols. Besides, the consequences of these attacks are serious and often compromise the whole group communication service.
2. Congestion control attacks of inflate subscription type: This is a pretty easy kind of attack. Most multicast-enabled application can trigger this attack, either intentionally or not. Risks are very high for the operator to see a subset of its multicast network become congested and unusable.
3. Congestion control attacks of deflate subscription type: This attack requires a modification of the client application or an application sending forged packets

in order to mislead the sending host congestion control protocol. Since it will only affect a limited number of clients (those of the service which is currently attacked), without any influence on other sessions nor on the Network Operator infrastructure, the importance is lower than with previous kinds of attacks.

4. Attacks on routers and physical infrastructure: This is one of the most unlikely kind of attack. Routers traditionally accept configuration commands only from the console or from a remote host belonging to the same subnet.
5. Internal attacks: Since these attacks require that the intruder be strategically placed, the risk is extremely low.

4.5 Conclusion

The multicast routing infrastructure has many security vulnerabilities and security problems exist in every aspect of the topology, from end-hosts to core routers. In this chapter, we list some of these security problems, discuss how various vulnerabilities can be exploited, and describe the effects of some of the attacks on the operation of multicast.

We clearly see that most edge attacks are pretty easy to launch. On the contrary protecting against the inside attacks and the attacks targeted to the routers and physical infrastructure are not a priority. This classification should be considered when designing the security framework that will enable the "continuation of service" goal (section 1.2).

4. TAXONOMY OF POSSIBLE ATTACKS OF INTEREST TO THE NETWORK OPERATOR

Chapter 5

A Focus on Group Management Protocols Specific Attacks

Contents

5.1	Introduction	37
5.2	Attacks on the IGMP/MLD routers	38
5.2.1	Querier Impersonation	38
5.2.2	Querier Paralysis	39
5.2.3	IGMP/MLD Querier Degradation	40
5.3	Attacks on the Hosts	40
5.4	Classifying the Group Management Protocols Threats According to their impacts on the Operator	41
5.5	Conclusion	42

Many works have explored the vulnerabilities of the group management protocols. However, these works consider mainly the LAN environment. In this chapter, we detail the threats of the group management protocols in considering other underlying environments. In particular, that of the network operator. We classify the group management specific attacks according to their potential impact on the network operator infrastructure.

5.1 Introduction

The group management protocols represent a mutual relationship between the router and their on-link hosts. Each part in this relationship do not need to know the other. Moreover, no part can verify the authenticity of the message sent to him. This makes the group management protocols highly vulnerable to many attacks. More specifically, IGMP/MLD have several characteristics that can lead to attacks [23].

5. A FOCUS ON GROUP MANAGEMENT PROTOCOLS SPECIFIC ATTACKS

- *Mandatory Query response*: While IGMP/MLD doesn't have any particular message authentication, any device which is acting as a router may force mandatory signaling responses from other hosts on-link.
- *Queries Elect the Querier Router*: Sending of IGMP/MLD Queries can be used to elect or de-select a Querying multicast router. This may be used to modify parameters on a network and conceivably support
- *Previous version of IGMP or MLD can bid down the recent versions*: Backward compatibility mechanisms for interworking between the current IGMPv2,v3/MLD (MLDv2) and IGMPv1/MLDv1 (MLDv1) allow hosts to change the MLD compatibility state on a router by sending Reports. This may be used to force changes in the source model used for off-link multicast routing.

- *Reports cause off-link changes*: The reports which are sent for joining arbitrary multicast groups cause changes to off link routing state when new groups are joined, or when routing halts after a group or source is excluded.

For instance, PIM-SM *Join* messages are initiated when a PIM-SM when an entities join a specific group or a specific source sending to the group. If this is due to a IGMPv1,2/MLDv1 (or IGMPv3 or MLDv2 Report with a zero-length EXCLUDE list), then a PIM-SM Join is sent as a (*,G) Join towards the RP.

If the join is triggered by an IGMPv3 or MLDv2 state change that affects source information, the PIM-SM join is sent as a (S,G) join towards the specific source.

PIM-SM *Prune* messages are initiated when a PIM-SM router determines that there are no entities interested in a specific group, or a specific source sending to the group. If this is triggered by either receiving a Report with an EXCLUDE or if a specific group with IGMPv1,2 or (Source/Group with IGMPv3) times out, then an (S,G) *Prune* is sent towards the upstream router.

- *Reporting can cause Querying*: Host transmitted *Report* messages can be used to instigate Queries from a router when the last group member leaves a group.
- *Unprivileged multicast API*: Access to arbitrary multicast groups is typically available through the host API. This allows generic tasks on a host computer to join or abuse multicast groups.

5.2 Attacks on the IGMP/MLD routers

5.2.1 Querier Impersonation

An entity can become a *Querier* by configuring its IP address to be lower than that of the legitimate *Querier*, then sending forged *Query* message. This will cause *Querier* duties to be assigned to the false *Querier*. If the later then sends no more *Query* messages, other routers' *Other Querier Present* timer will time out and one will resume

the role of *Querier*. During this time, if the attacker ignores *Leave* messages, traffic might flow to groups with no members for up to *Group Membership Interval*. During its role as *Querier* the attacker can stage a variety of attacks on the hosts as ignoring *Report* messages and flooding the network with *Query* messages.

These attacks could be specifically harmful to the multicast hosts, however, they do not create any multicast state in the multicast operator infrastructure. Moreover, the effect of *Query* message forging could be avoided by giving the *Querier* router (for instance, the DSLAM) the smallest IP address.

5.2.2 Querier Paralysis

When the router becomes unable to perform its routing functionalities, this includes data forwarding functions (i.e traffic conditioning, traffic classification, scheduling and buffer management) or the signaling functions in the routers, we say that the router is paralyzed.

Such an attack happens when the router's memory is flooded by multicast traffic and routing states or when the router enters in an endless computing operation which consumes its CPU.

While some events such as flash crowds can flood the router, in the following we consider only anomalous behaviors. That is, when an host or multiple hosts intentionally attacks the *Querier* router using the vulnerabilities of IGMP/MLD protocol. To achieve this goal an attacker needs to forge the control messages of IGMP/MLD protocol. This is an easy target, as IGMP/MLD does not provide any mechanism to authenticate the source of the message as been mentioned previously.

Increasing Multicast Routing States

When receiving an IGMP/MLD message for a new group, the router creates a multicast state for this group: (*,G) in PIM-SM or (S,G) in PIM-SSM. In order to saturate the router with these states, a host forges *Report* messages for a large number of groups.

Such an attack could have a direct impact on the operator network, as it creates multicast states in the first IGMP enabled equipment (DSLAM, edge router...) and the whole infrastructure if PIM-SM is active.

CPU Consumption Attacks

A forged *Leave* message will cause the *Querier* to send out *Group-Specific Queries* for the group in question. This causes extra processing on each router and on each member of the group. An amplified version of this attack could be as follows: For each *Report* heard in the network the attacker sends a *Leave*, causing the *Querier* to enter in an endless computing operation.

A forged *State-Change Report* message will cause the *Querier* to send out *Group-Specific* or *Source-and-Group-Specific Queries* for the group in question. This causes extra processing on each router and on each member of the group, but can not cause loss of desired traffic.

5. A FOCUS ON GROUP MANAGEMENT PROTOCOLS SPECIFIC ATTACKS

Router Saturating

A malicious host can use IGMP/MLD to elicit unfair share of the available bandwidth. Since a multicast traffic is of UDP/IP type, it can easily lead to router congestion. Thus, to flood the data plane in the router a host can for example subscribe to a large number of high bandwidth groups or he can simple misuse the congestion control protocols.

The multicast congestion control protocols showed their great diversity with respect to the definition of the congested state, session structure, and mechanisms for congestion notification and reliability. Despite the differences, the multicast protocols share a common feature their congestion control assumes that each party always adheres to guidelines for fair sharing of the network bandwidth [33].

So, with these protocols there is a high risk that some clients or sources applications do not use them (e.g. a coarse application that does not implement any congestion control), or misbehave, whether deliberately or not (e.g. does not join an appropriate number of groups in multi-rate schemes, or the right group in DSG). Then the resulting unresponsive flow can easily congest routers, thereby creating another kind of DoS attack.

In the case of feed-back driven protocols. A host can distort the congestion summary, to trick the sender into unfairly high transmission.

5.2.3 IGMP/MLD Querier Degradation

A forged IGMP Version 1 *Report* message may put a router into "version 1 members present" state for a particular group, meaning that the router will ignore *Leave* messages. This can cause traffic to flow to groups with no members for up to *Group Membership Interval*. This can be solved by providing routers with a configuration switch to ignore Version 1 messages completely. This breaks automatic compatibility with Version 1 hosts, so should only be used in situations where "fast leave" is critical.

With IGMPv3, a forged *Version 2 Report* message may put a router into "version 2 members present" state for a particular group, meaning that the router will ignore IGMPv3 *source-specific state* messages. This can cause traffic to flow from unwanted sources for up to *Group Membership Interval*. This can be solved by providing routers with a configuration switch to ignore Version 2 messages completely. Again, this breaks automatic compatibility with Version 2 hosts, so should only be used in situations where source INCLUDE and EXCLUDE are critical.

5.3 Attacks on the Hosts

Hosts have no idea which is a valid *Querier*. In consequence, when the *Querier* misbehaves or when a another entity impersonates the legitimate *Querier*, this intruder can stage a variety of attacks on the end hosts.

For example, when host suppression is not in use, a router specifying a very small *Maximum Response Time* in its *Query* messages may cause multicast report bombing at

5.4 Classifying the Group Management Protocols Threats According to their impacts on the Operator

fine granularity with a single message. In some cases, this may have severe consequences in terms of packet loss or delay on other data sources or signaling.

A DoS attack on a host could be staged through forged *Group-and-Source-Specific Queries*. The attacker can find out about membership of a specific host with a general *Query*. After that it could send a large number of *Group-and-Source-Specific Queries*, each with a large sources list and the *Maximum Response Time* set to a large value. The host will have to store and maintain the sources specified in all of those queries for as long as it takes to send the delayed response. This would consume both memory and CPU cycles in order to augment the recorded sources with the source lists included in the successive queries.

Some hosts can attack other hosts by misusing the trust prerequisite in the congestion control protocols. Thus, a host can misuse the congestion control protocols to elicit self-beneficial share of the bandwidth at the expense of cross traffic. We distinguish between this self-beneficial attacks and the DoS attack mentioned in the previous section where an attacker aims to totally paralyze the router by requesting high bandwidth groups.

5.4 Classifying the Group Management Protocols Threats According to their impacts on the Operator

In the light of our definition of the security of the multicast routing infrastructure from the operator point of view, we classify now the group management protocols threats according to their impact on the operator infrastructure as follows:

- *Querier Impersonating* IGMP/MLD are naturally vulnerable to *Querier* impersonating. In the access network the Querier could be the DSLAM or the B-RAS. Thus, impersonating such devices could have severe consequences of the operator network and on its clients. As has been cited in [18]: *we believe that the problem related to Querier selection is solvable. For example, all routers of a link could share a secret key. It would then be enough for routers to verify the authenticity of the Query messages. The problem related to fake routers causing extraneous traffic by sending fake Query messages is much harder to solve. This problem, known as the fake bank teller problem, is not specific to group communication. This issue of validating routers or, in general, nodes that offer services remains an unsolved authorization problem in the Internet at large.*

- *Querier Paralysis*

Making the *Querier* router totally saturated by multicast states, and thus disturbing the routing and forwarding functions is quite easy using the characteristics of IGMP/MLD. As been shown in the previous section, incapacitating the router could be done by flooding the control plane, data plane or both. In both cases these attacks have serious consequences on the whole infrastructure. Thus, it's a priority to find a pragmatic solution to protect the routing functionality against these attacks.

5. A FOCUS ON GROUP MANAGEMENT PROTOCOLS SPECIFIC ATTACKS

- *Attacks on the Hosts*

As has been shown the hosts could undergo attacks coming from the router side or from the other hosts. Thus, securing the *Querier* router itself plays an important role in protecting the hosts. While it's not possible to secure the hosts against attacks coming from other hosts, especially in environment where the end-host can reach each other directly and without passing by the operator networks. However, in some attacks such as self-beneficial attacks caused by misusing the congestion control protocols, the operator can protect the hosts by using secure congestion control protocols. That is, by enforcing some policies that allow to control the subscription level of each user. Thus, only an uncontested receiver can report for its current or a higher subscription level.

5.5 Conclusion

The group management protocols represent a mutual relationship between the router and their on-link hosts. Each part in this relationship do not need to know the other. Moreover, no part can verify the authenticity of the message sent to him. This makes the group management protocols highly vulnerable to many attacks in particular DoS attacks which could impact the whole routing infrastructure.

All the group management specific attacks are harmful for the operator. However, making the *Querier* router totally saturated by multicast states, and thus disturbing the routing and forwarding functions is quite easy using the characteristics of IGMP/MLD and could paralyze all the multicast delivery service.

In the following chapters we will focus on this attack. Moreover we will present a pragmatic and intelligent solution that improve the resiliency of the access infrastructure towards this attack.

Chapter 6

A Survey of Proposals to Secure the Multicast Routing Infrastructure

Contents

6.1	Attack Avoidance Approaches:	44
6.1.1	Securing the Edges of the Multicast Tree	44
6.1.2	Protecting the Core of the Multicast Tree	47
6.2	Attack Resiliency Approaches:	48
6.3	Hybrid Techniques	49
6.4	Conclusion	50
6.4.1	Participant Authentication and Authorization	50
6.4.2	Modifying Existing Protocols	51
6.4.3	Inter-Domain Confidence	51
6.4.4	Simple filtering	51

The multicast routing infrastructure contains the multicast routers which exchange control messages to construct the multicast routing trees and it is connected directly to the clients premises and the content provider networks. Trying to secure the routing infrastructure have been the axe of many researches. In this section, we classify theses approaches into three broad categories:

- *Attack avoidance approaches:*

Attacks avoidance is performed by controlling the ability of entities (i.e. routers, receivers nad senders) to take part in the multicast routing tree for a given group. Such an access control could be enforced via: 1) *cryptographic methods* which rely on key management and distribution protocols[43], or 2) *non-cryptographic methods* which aim to control the access of a given entity by simply enforcing a

6. A SURVEY OF PROPOSALS TO SECURE THE MULTICAST ROUTING INFRASTRUCTURE

lightweight access list which contains, for example, the authorized IP addresses. These approaches focus on two issues: *authentication*, which requires that entities prove their identities before being allowed to join the multicast session and *authorization*, which implies that only these entities with specific permission are allowed to participate in the multicast session for particular group.

The efforts in this category focus on two issues: (1) the security of the *core of the multicast tree*, and (2) the security of the *edges of the multicast tree*. The core of the multicast tree contains the multicast routers which exchange control messages to construct the multicast routing trees and the edges of the multicast tree include the (receivers/senders) and the access routers.

- *Attack resiliency approaches:*

These approaches suppose that no system is breaches-proof. Thus, in this category we find techniques that aim to detect attacks and to mitigate their affects on the network. The attacks resiliency techniques aim to survive those attacks that could not be completely repelled.

- *hybrid approaches:*

In this category we find the works whereby the access control and the filtering techniques are combined. They generally introduce a management policy that controls multicast parameters like the maximum number of groups individual hosts can join or send traffic to, or the rate of the incoming traffic.

For each category we present some approaches and we discuss these propositions from the multicast operator viewpoint.

6.1 Attack Avoidance Approaches:

The efforts in this category focus on two issues: (1) the security of the *core of the multicast tree*, and (2) the security of the *edges of the multicast tree*. The core of the multicast tree contains the multicast routers which exchange control messages to construct the multicast routing trees and the edges of the multicast tree include the (receivers/senders) and the access routers.

6.1.1 Securing the Edges of the Multicast Tree

1. *Securing the Group Management Protocols*

These mechanisms aim to secure the communication between the access routers and the receivers, by: (1) Secure IGMP/MLD: which aim to control the ability of end users to join the multicast groups. The goal is to enable the routers to determine if a given host is authorized to join a given group. This problem is sometimes referred to as the *Proof-of-Membership problem* [37]. Such mechanisms appear under different names: *Multicast Group Access Control* or *Secure IGMP/MLD*;

and (2) Securing the congestion control protocols. Since the identity based access control alone cannot protect the leaves of the multicast tree against self-beneficial congestion attacks.

(a) *Secure IGMP/MLD:*

[13], was first to present a new version of IGMP that allows receivers to be authorized before joining the group. The architecture includes authorization servers that possess ACLs distributed by an initiator. This approach did not meet any reasonable scalability requirements, besides it is vulnerable to IP address spoofing.

In [37], Hardjono and Cain proposes a solution to the IGMP proof-of-membership problem that uses a Key server. The proof consists of a symmetric-key, IGMP-key that is used by the receiver and the Multicast router to protect the IGMP message.

This solution seems to work for multicast group within a domain only. It is not clear how a group that has members in different domains could use this proposal. Furthermore, the scalability property of this solution is questionable. A domain key server must have tokens for each potential member of each multicast group of the domain. Additionally each router must have tokens for each potential member of each multicast group of its links [18].

Other similar works are [47], [31], and [19].

The previous proposals do not support mobile hosts that visit a foreign domain. In [18], Castelluccia and Montenegro propose a solution to secure MLD in the context of mobile IPv6 environment where the Proof-of-Membership problem is exacerbated, as routers do not necessarily know the multicast listeners. In this work the authors propose to extend the Cryptographically Based Addresses(CBA) concept to group address to solve the proof-of-membership problem. This approach was first to handle the security of MLD. The authors discuss the possibility to extend this approach to cover other group management specific attacks. The use of CBA addresses allow a distributed solution and it does not need any pre-established security association with the listeners.

Short Discussion:

The previous approaches need new version of IGMP/MLD which make their acceptance by the operators community a disputable issue. Moreover, these proposals do not protect the operator infrastructure in the situations where authorization and authentication are not needed, or when an authorized user launch a flooding attacks.

(b) *Secure Congestion Control Protocols*

[33] is the only in this category. This work considers receivers driven congestion control mechanisms but not sender based solutions (i.e. the sender references the feedback from the group members to control the overall rate of this group).

6. A SURVEY OF PROPOSALS TO SECURE THE MULTICAST ROUTING INFRASTRUCTURE

To protect the network against inflated subscription, this work introduces a group access control mechanism based on the congestion status of receivers. Thus, the congestion status of a receiver determines its eligibility to access a multicast group. To that goal, each group is guarded by three keys: top key, decrease key and increase key. The conditions for keys distribution are as follows:

- i. An uncontested receiver should obtain updated keys for its current groups.
- ii. An congested receiver of g groups should obtain updated keys for its lower $g - 1$ groups. It can obtain an updated key for group g only if the protocol authorizes an upgrade to group g , and groups 1 through $g - 1$ do not lose packets.
- iii. When authorized, an uncontested receiver of g groups should obtain an updated key for group $g + 1$.

Short Discussion:

This proposition is an original work in the field of congestion control robustness. However, this work does not protect against more arsenal bandwidth attacks. And in the case of the operator network, if the DSLAM plays the role of IGMP proxy, this work does not protect the shared line between the DSLAM and the BRAS from being congested. This motivates introducing other mechanisms adaptable for the Network Operator.

2. *Sender Access Control*

Sender access control is needed to prevent group members from receiving undesired data, but also to protect the multicast infrastructure from being flooded by irrelevant multicast traffic.

The source specific and uni-directional shared trees, where information sources can be authorized or authenticated at the single root or Rendezvous Point (RP), provide an intrinsic protection against such an attack. However, the problem is exacerbated in the bi-directional trees.

Several source authentication schemes have been suggested, but none of these schemes is satisfactorily efficient in all prominent parameters [60].

In [76], the authors propose a policy for sender access control for bi-directional multicast routing so that irrelevant data is policed and discarded once it arrives at an on-tree router. Short Discussion:

the operators community privilege largely the proposal that alleviates the burden of cryptographic methods. This is still true in the context of sender access control. The proposition presented in [76] is a lightweight solution that meet the need of the operator. However, as all the access-list oriented solutions, it does not protect the participants against IP address spoofing. Moreover, this solution relies largely on the cooperation between different independent domains. This could be a main hindrance to its adoption by the operators.

6.1.2 Protecting the Core of the Multicast Tree

In the context of unicast, many cryptographic mechanisms have been introduced to secure the routing protocols messages [43]. They aim mainly to prevent the control messages from being illegally modified in transit, and preventing bogus routing information from being injected into the network [38]. These authentication mechanisms can be used to protect the multicast routing messages. However, they cannot defend the routing infrastructure in its integrity; the open-model nature of the multicast that allows group membership to be transparent to the sender and the host to join and leave the group in any moment make the unicast oriented mechanisms ineffective against the multicast specific threats.

[65], is the first work pointing out the importance of the *authorization* in the context of multicast routing security. *Authorization*, implies that only these entities with specific permission are allowed to participate in the multicast session for particular group.

This work introduces KHIP, a secure multicast routing protocol. KHIP defines a new protocol, the Ordered Core Based Tree protocol (OCBT), a variant of the bi-directional core-based multicast routing (CBT) [12].

KHIP focuses on the authentication and authorization issues. It uses an authentication service (AS) [32] to authenticate and authorize the routers. The authentication service maintains the list of who is allowed what access to specific multicast groups. When a router wishes to become part of the secure multicast tree, it must first request a certificate for the group from the authentication service using its public key. If authenticated and approved, it receives a certificate consisting of its IP address, its public key, the multicast group or range of authorized groups. The router can use this certificate then to join the tree using a special join message.

Other works [56] [36] [66] which target the core of the multicast tree rely on a key distribution method based on the use of a combination of symmetric and asymmetric keys to authenticate the control messages for the PIM-SM protocol. So, only authenticated router can participate in the routing functions.

Short Discussion:

Such mechanisms are undoubtedly form theoretically ideal solutions to secure the core multicast routing. However, their acceptance by operator community could be a tenuous issue. First, these proposals consider only the communication between the routers constructing the multicast tree; they do not describe the communication between hosts and routers or the group management infrastructure. Without secure version of IGMP/MLD protocol the leaves of the multicast tree are highly vulnerable to DoS attacks, as acknowledged by the authors of KHIP themselves. Second, they rely on authorization and authentication mechanisms which need to modify a widely used infrastructure this will preclude their adoption, by the operators who are very conservative on this point. Third, they prevent partially flooding attacks or other multicast DoS attacks. So, other mechanisms still needed to guarantee the service continuity requirement.

6.2 Attack Resiliency Approaches:

These approaches suppose that no system is breaches-proof. Thus, in this category we find techniques that aim to detect attacks [21] [75] and to mitigate their affects on the network. The attacks resiliency techniques aim to survive those attacks that could not be completely repelled.

1. *Rate Limiting Techniques*

These techniques impose a rate limit on the control traffic. To that goal, they require to define thresholds (statically or dynamically). Another example is the IGMP State Limit (ISL) feature of CISCO IOS [5] which provides a mechanism to limit the number of IGMP states that can be joined on per interface, per sub-interface, or at the global router level. [30] presents a new solution to achieve rate limiting par source address. Thus this solution can protect up to a certain point the legitimate sources in the case of attack.

2. *Offending traffic filtering*

These techniques use the signals raised by an Intrusion detection system (IDS) to filter out the malicious streams. Intrusion detection has been studied for a long time. Its primary goal is the detection of abuse of computer systems. IDS systems can be generally classified as either signature-based or statistical-based. *Signature-based* detection requires to understand the attack in order to identify them by their state transition sequences or patterns. *Statistical-based* detection relies on the comparison of previously observed normal statistical profiles and under/after attack statistical profiles.

Many intrusion detection works have been proposed to secure the routing infrastructure. However, most of the previous proposals consider unicast. It has been clearly shown that the use multicast monitoring for detecting attacks introduces many challenges [7]. One of the few efforts done in this area can be found in [62]. This work deals with the MSDP SA flooding and presents a reactive approach to detect and deflect this attack. This approach employs a signature-based intrusion detection scheme that uses a dynamic prediction of normal traffic. This prediction is used to find dynamic threshold to filter the traffic per sources basis. If a source advertises more than the SA threshold set for it, this source is considered rogue and all it's SAs are dropped.

3. *Attack Isolation Mechanisms*

Isolating the attack could be achieved by modifying the network topology by either adding more resources or isolating the compromised areas [22] [69]. In [79] we find a framework for survivable connection oriented group communications. Here providing survivability is expressed as a multidimensional optimization problem where the overall goal is to make failures imperceptible to the users.

We think that these techniques can largely benefits from the recovery techniques introduced by the Multiprotocol Label Switching (MPLS) [57] infrastructure.

Short Discussion:

The filtering solutions has long been the only way to defend against flooding attacks. However, they include a risk of accidentally denying service to legitimate users. Even worse, an attacker can use filtering mechanisms as a tool to launch denial of service attacks [54]. They face also the problem related to the defining of a static threshold, as well as, with the static source addresses (no new addresses will be accepted by this mechanism).

6.3 Hybrid Techniques

In this category we find the works whereby the access control and the filtering techniques are combined [15]. They generally introduce a management policy that controls multicast parameters like the maximum number of groups individual hosts can join or send traffic to, or the rate of the incoming traffic. [51], presents *MCOP* an intra-domain multicast management solution. It provides multicast management functions with a centralized information database located at a Multicast Control Server (*MCS*). These information include the networks or hosts which can or can't send traffic to a group, the networks can receive traffic from a group, the maximum number of groups that an individual host can receive traffic from, and the total rate of traffic to multicast groups that an individual source can send.

The *MCOP* protocol is used between the *MCS* and the router with directly connected multicast sources or receivers.

As acknowledged by the authors themselves, this approach however has some limitations. The *MCS* can face scalability problems if the number of entities is large. A *DoS* attack on the *MCOP* is also possible when an attacker sends a large number of bogus *IGMP/MLD* to the *MCOP* router.

MAFIA [18], is an another multicast management solution for access control and traffic filtering. This approach aims to 1) control multicast group membership, that is, limit the use of multicast to only trusted hosts and groups, as well as, to control which host can be a member of a certain multicast group, 2) to filter multicast traffic flowing in and out of the enterprise using state gathered from the multicast routing protocol, and 3) prevent multicast denial of service attacks using multicast access control as a prevention technique.

[49], introduces a new filtering architecture to thwart some *DoS* attacks that are based on *IGMP/MLD* flooding. This work shares some similarities with [30] and [18]. However, its main advantage in compared with the previous proposals is that it includes a learning mechanism. More specifically, it creates two classes of clients: known clients, who have been already accepted by the system, and new clients. Known clients are then served equally by the filter that forwards their *IGMP/MLD* messages to the first hop multicast router. Thanks to this feature, well-behaved clients, already known by the system, will be served with a higher priority during an attack. As cited by the authors, several peripheral components can be added to the filtering component, taking into account the specificities of the target environment.

6. A SURVEY OF PROPOSALS TO SECURE THE MULTICAST ROUTING INFRASTRUCTURE

Short Discussion:

The management proposals undoubtedly form an attractive family of solutions from the operators community viewpoint. A major advantage is that these proposals do not require any special support from network routers and can therefore be deployed universally.

6.4 Conclusion

6.4.1 Participant Authentication and Authorization

Several proposals for a secure multicast routing infrastructure require that participants (source or clients) be authenticated and authorized to use the group communication service [65]. We argue here that this requirement is not always *feasible* or *efficient*:

- *It is not feasible in case of free content delivery services:*
Authentication/authorization assume that the client is registered somewhere, for instance in a RADIUS server. It is always the case with a commercial service like a video over ADSL (section 2.2.1). But there is no such RADIUS server with services requiring no preliminary registration, for instance in case of a free download service if a group of friends sets up a video-conference between them (section 2.2.2). Such an user-initiated private group communication is by definition uncontrollable.
- *It is only feasible if a network operator/content provider agreement exists:*
Indeed, doing authorization requires that the network operator has access to the client database of the content provider (we assume a commercial delivery service here). This is usually done by accessing the RADIUS server, which is often hosted by the network operator. But what about the situations where a client connected to a first operator wants to access the services provided by a content provider who has no agreement with the network operator (e.g. the content provider may have agreements with another network operator)?
No authorization is possible, and the first network operator either accepts the Join request, considering that the authorization task will be performed by the second network operator, or refuses.
- *It is not always efficient:*
Indeed, nothing guarantees that an authenticated/authorized client will behave correctly. This client can use an ill-behaving application, thereby leading to congestion problems. This host can also be subverted by an attacker (thanks to a virus, a Trojan horse, or a root-kit installed on the client's PC). This attacker can then take advantage of the legitimate client to create for instance DoS attacks.
- *It is not feasible when NAT is used:*

Let's consider a user using a home gateway¹ which also performs NAT (Network Address Translation). In that case, once the user has joined the group communication service, all the traffic to/from his home gateway is considered legitimate by the network operator. If the user has set up a Wireless LAN, then any host in the vicinity can easily join the WLAN (802.11b is known to offer a very limited security level) and create DoS attacks, taking advantage of the legitimate user's authentication/authorization. The fundamental problem here is that NAT prevents IP-level authentication/authorization till the end-host.

6.4.2 Modifying Existing Protocols

Many proposals for securing the multicast infrastructure rely on the modification of existing, largely deployed protocols, in order to add security features. Examples are the addition of authentication into IGMP/MLD, the robustification of various routing protocols, like PIM-SM. Even more extreme solutions lead to design and secure new routing protocols, like KHIP [65] which assumes the presence of a variant of CBT [12] (whereas this latter has never been implemented and deployed within operators). Even if the resulting protocols are extremely robust in front of certain types of attacks (the ones for which they are supposed to be robust), they do not form a realistic class of solutions for our goal.

The Network Operator has many reasons not to move to such new solutions. Even if such solutions become standardized by the relevant IETF working group (usually a long process), the path can be long before it becomes widely deployed. And even in that case, there is a risk that some client (or other peering networks) keeps on using the old non-secure version.

6.4.3 Inter-Domain Confidence

A Network Operator is responsible of the communication services provided to its clients. Therefore security solutions relying on a collaboration between all the Network Operators implied in a multicast tree cannot be deployed easily. The security requirements of a Network Operator, as defined in section 1.2 are essentially selfish, and should not be limited in any way by dependencies on other Network Operators.

6.4.4 Simple filtering

Simple filtering does not help as it penalizes the legitimate and the attack packets in the same time.

¹ An equipment providing WAN access and several convenient services for the internal home network, like NAT, an Ethernet hub/switch, an IEEE802.11b access point, routing/forwarding between the various networks, firewall.

6. A SURVEY OF PROPOSALS TO SECURE THE MULTICAST ROUTING INFRASTRUCTURE

Chapter 7

Our Proposal: A Filtering Approach to Mitigate IGMP/MLD Flooding Attacks

Contents

7.1 Architectural Overview	54
7.1.1 Packets Capturing	54
7.1.2 Packets Classification	54
7.1.3 Clients Purging	55
7.2 A Closer View of the Various Building Blocks	55
7.2.1 The packet capture and classification thread:	55
7.2.2 The known clients queues creation thread:	56
7.2.3 The main scheduling thread:	56
7.2.4 The purging thread:	56
7.3 Deployment	57
7.4 Benefits in Front of a DoS Flooding Attack	57
7.4.1 Naive Flooding DoS Attack Without IP Address Spoofing . .	57
7.4.2 Flooding DoS Attack With Random IP Address Spoofing . .	58
7.4.3 Flooding DoS Attack With Targeted IP Address Spoofing . .	58
7.4.4 What About other Group Management Specific Attacks? . .	59
7.5 Conclusion	59

In this chapter, we introduce a simple yet efficient filtering approach to thwart some DoS attacks that are based on IGMP or MLD flooding, and that threaten the whole operator's infrastructure. A key feature of our proposal is that it follows a realistic and pragmatic approach, and in particular it does not require any modification to the existing, widely deployed protocols [49]. More specifically, our approach has the following objectives:

7. OUR PROPOSAL: A FILTERING APPROACH TO MITIGATE IGMP/MLD FLOODING ATTACKS

1. *Mitigating multicast specific attacks through a cost effective, scalable and transparent mechanism. This mechanism aims, in particular, to protect the well-behaved clients against ill-behaved ones.*
2. *Keeping the changes to the currently used multicast routing as minimum as possible.*
3. *Avoiding non realistic assumptions in the existing approaches.*

7.1 Architectural Overview

The proposed solution relies on a filter, managed by the network operator, and located between the clients and the first hop multicast router. This filter captures the IGMP (or MLD) packets generated by clients, filters them traffic according to specific rules that will be detailed later, and sends them back to the network. The filter does not take part in any way to the packet forwarding functionality (still managed by the router), hence, it is transparent for the multicast packet forwarding.

7.1.1 Packets Capturing

Packets capturing module sniffs all IGMP *REPORT* and *LEAVE* packets. The *QUERY* packets are not considered, since clients are not expected to issue them.

7.1.2 Packets Classification

After capturing the IGMP packets the filter classifies them according to their source IP address. Two categories of packets are defined:

- those coming from a client that is already known by the system, and
- those coming from a new, unknown client.

To that goal the filter keeps a context for each known client. This context contains the source IP address, the date of the last IGMP packet received, and a queue containing a maximum number of G_{max} packets for G_{max} different group addresses.

The Case of Known Clients

An IGMP *REPORT* or *LEAVE* packet issued by a *known client*, S_i , and related to group G_1 , is enqueued to the associated queue of S_i . It may erase a previously enqueued packet if the queue already contains a packet related to group G_1 . If the filter is correctly initialized, then IGMP packets issued by legitimate clients should not accumulate in the filter. This point is further discussed in next chapter.

Periodically, a certain number of the IGMP packets enqueued in the known client lists are selected and sent back to the network. Those packets are the ones that will be accepted by the first hop multicast router. In order to guaranty fairness within the set of known clients, this scheduling follows a round-robin policy.

The Case of Unknown Clients

On the opposite, an IGMP *REPORT* or *LEAVE* packet issued by an *unknown client*, S_j , is systematically enqueued in a dedicated FIFO whose maximum size is strictly enforced.

Periodically, a certain number of IGMP packets enqueued in the unknown-client FIFO queue are elected, and the associated clients are accepted by the filter. A context is created for each client, and they are now "known" by the system. Therefore future incoming IGMP packets arriving from these clients will be directly accepted by the system and enqueued in their associated list.

7.1.3 Clients Purging

Because clients will disappear, a purging system is set up in the filter. Periodically a thread monitors all known clients and checks if each of them has been active during the period, i.e. has sent at least one IGMP packet. If the client has been silent, then it is dropped from the list of known clients and its context is removed.

Because the IGMP version 1 and version 2 *REPORT* suppression mechanism does not oblige each client to reply to a *QUERY* request, the purging period must be an order of magnitude larger than the IGMP *QUERY* polling period (often equal to 125 seconds, but this can be changed).

In case of IGMP version 3, there is no *REPORT* suppression mechanism and the purging period can be set to a value a little bit higher than the IGMPv3 *QUERY* polling period¹. This is the optimal situation, and the list of active clients closely matches the reality, which warrants optimal classification performances and minimum memory requirements. Besides, if an attack occurs, creating contexts for ghost clients, then these context will quickly be removed from the system.

7.2 A Closer View of the Various Building Blocks

7.2.1 The packet capture and classification thread:

Packet capturing module relies on the `libpcap` library [45] to capture all IGMP *REPORT* and *LEAVE* packets. The IGMP *QUERY* packets are not considered, since clients are not expected to issue them. The *PCAP* library interfaces directly with the kernel of the operating system in order to avoid the copy of packets data from the kernel to the user level space before filtering. *PCAP* is based on the BPF (Berkeley Packet Filter). The BPF listens directly on the link layer interface and applies its filtering rules as required by the user before passing data to the user level program. Definitions of packets to be filtered can be written in a simple human readable format using boolean

¹ Note that the explicit tracking of clients associated to a multicast group functionality can be enabled on Cisco routers when using IGMPv3. Network operators are typically interested by this functionality, even if it is not made mandatory by the IETF documents, since it will help to improve the network behavior.

7. OUR PROPOSAL: A FILTERING APPROACH TO MITIGATE IGMP/MLD FLOODING ATTACKS

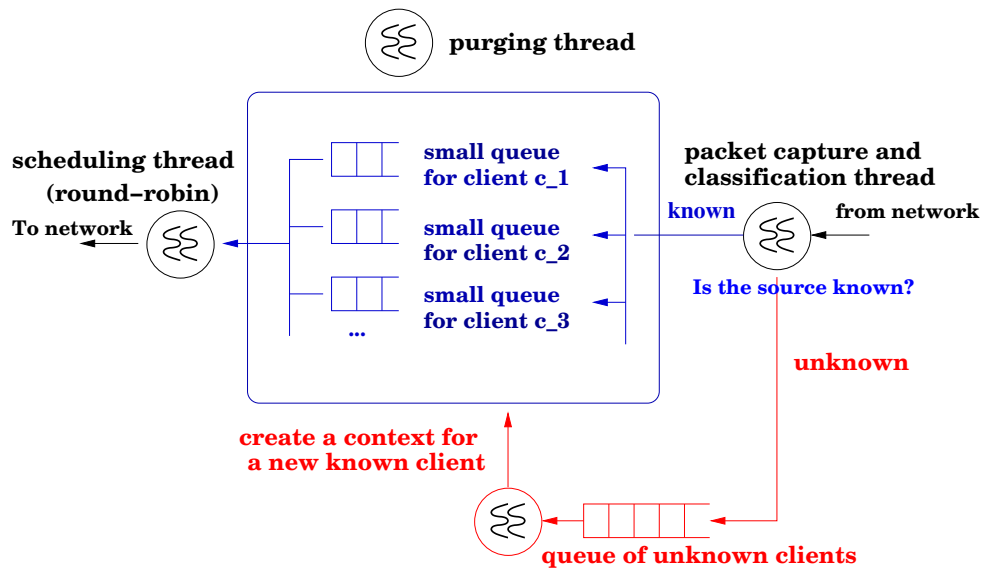


Figure 7.1: Architecture of the filter.

operators and can be compiled in a pseudo-code to be passed to the BPF device driver by a system call.

This thread also performs a classification of each packet based on the source IP address, as explained above. For performance purposes, a hash-based search algorithm is used by the classification function. The hash table is made of a direct address table (a static table) which is addressed by an index obtained through an hash function applied to the source IP address of the packet.

7.2.2 The known clients queues creation thread:

This thread elects, periodically, a certain number of IGMP packets enqueued in the unknown-client FIFO queue and creates queues for the associated clients.

7.2.3 The main scheduling thread:

This thread implements the round-robin scheduling of the waiting packets for the known clients. Each packet is then sent back to the dedicated link that links directly the filter and the first hop multicast router. The original source address of the IGMP packet must be kept in order to enable a correct operation of IGMPv3 (which unlike IGMPv2 is a stateful protocol).

7.2.4 The purging thread:

This thread periodically checks whether a client has been active or not. As explained above, this thread is only activated with a very low frequency in order to prevent "false

positive” problems in IGMPv1 and v2 mode. This is different in IGMPv3 mode since the frequency will be much higher.

7.3 Deployment

Two different deployments are possible (figure 7.2):

- an external deployment (figure 7.2(a)) where the filter is implemented in an independent host, located at the operator’s premises, beside the first hop multicast router, and it is connected both to the public LAN and to the multicast router through a dedicated link. Therefore this solution is universal, the only feature required from the router being the possibility to use an Access Control List (ACL) to only consider IGMP packets coming from the filter.
- an internal deployment (figure 7.2(b)) where the filter is integrated to the first hop multicast router.




Figure 7.2: External versus internal deployment of the filter.

7.4 Benefits in Front of a DoS Flooding Attack

The filtering mechanism has been designed to improve the resiliency of the multicast routing infrastructure in front of IGMP (or MLD) flooding DoS attacks. More specifically, two kinds of situations must be considered:

- the naive DoS attack, where the IGMP packets contain the attacker’s IP address,
- the DoS attack where the IGMP packets contain a spoofed source IP address.

7.4.1 Naive Flooding DoS Attack Without IP Address Spoofing

If the source address of packets used by the attacker is its real address, this latter is most probably already known by the system, or if it is not the case, he will quickly be known. So packets are systematically enqueued in the associated list, whose size is by

7. OUR PROPOSAL: A FILTERING APPROACH TO MITIGATE IGMP/MLD FLOODING ATTACKS

definition limited to at most G_{max} packets (all with a different group address). Since the attacker's IGMP packet arrival rate will most probably exceed its fair share of the filter outgoing rate, the associated queue will always overflow and the attack will easily be defeated.

7.4.2 Flooding DoS Attack With Random IP Address Spoofing

In the second case, the attacker uses systematically forged source IP addresses. If these addresses are chosen randomly, or if the subnet addressing space is much in excess of the possible legitimate clients, then most of the attacker's packets will not be known by the system, and therefore will enter the "unknown clients" FIFO queue. Since the service rate of this queue is low in front of the attacker's sending rate (by definition of a flooding attack), the FIFO will overflow. Of course some of the forged addresses will be accepted by the system and a context will be created for them. Yet if the attacker continues to randomly use source IP addresses, the accepted addresses will probably represent only a small fraction of the total number of addresses. Besides the maximum number of packets that will go through the filter for these accepted forged addresses will anyway be limited to the fair share of the filter outgoing rate. Both mechanisms will automatically defeat the attack, especially if this one has a limited duration.

7.4.3 Flooding DoS Attack With Targeted IP Address Spoofing

A more intelligent variant of this attack consists in using the range of possible addresses of the subnet, which in case of IPv4 networks, will most probably be limited to a few tens or hundreds of hosts. This kind of attack will more easily limit the benefits of our filter than the previous two kinds of attacks: (1) either these hosts are already known by the system (i.e. the legitimate client using this address has recently issued an IGMP *REPORT* or *LEAVE* packets) and the attacker's packets will immediately be accepted, or (2) since there is a limited number of possibilities, if the attack duration is long enough, most addresses will finally be known by the system, after which further packets will immediately be accepted. In both cases, even if the impacts on legitimate clients will be serious, for instance preventing their legitimate IGMP packets from being sent to the first hop multicast router, the outgoing rate of packets sent to this router will not exceed the nominal outgoing rate of the filter. Since this rate takes into account the capabilities of the multicast routing system (chapter 8), this attack will not have any other impact than preventing legitimate clients of this subnet from using multicast services. The DoS attack is, in that case, confined to a few clients, but does not impact the whole multicast routing infrastructure of the operator.

In another variant, the attacker deliberately uses the filter's source IP address, in order to make its IGMP packets be accepted by the first hop multicast router. This attack is in fact easily defeated by having a direct link between the multicast router and the filter (the "additional link" of figure 7.2(a)). The first hop multicast router only accepts IGMP packets arriving from this link, rather than from the shared subnet. Similarly, the filter is configured to ignore IGMP packets having one of its IP addresses

as a source address. Note also that in the integrated implementation, where the filter is integrated to the first hop multicast router itself rather than being implemented within an independent host attached to the multicast router, the attack is also easily defeated.

7.4.4 What About other Group Management Specific Attacks?

In this section, we discuss the robustness of the filter in front of the other, more subtle, attacks that take advantage of group management specificities (see chapter 5).

Forged IGMP/MLD Report Messages

An attacker can use forged *Report* messages to subscribe to a large number of groups, thus, saturating the router and the network by multicast states. The filter can limit the harmfulness of this attack since it limits the number of groups for each client. However, the filter alone cannot totally stop these attacks. In particular, when the attacker spoofs a large number of IP address and group addresses.

Forged IGMPv3/MLDv2 State-Change Report

A forged *State-Change Report* message will cause the *Querier* to send out *Group-Specific* or *Source-and-Group-Specific Queries* for the group in question. An amplified version of this attack causes extra processing on the router and on each member of the group. Again the filter can limit the harmfulness of this attack if the attacker does not use spoofing.

Report and Leave Messages Storm Attack

A forged *Leave* message will cause the *Querier* to send *Group-Specific Queries* for the group in question and will cause the receivers to send *Reports*. This causes extra processing on the first hop multicast router and on each member of the group.

An amplified version of the attack is also possible: for each *Report* heard in the network, the attacker sends a *Leave* message causing the resources of the *Querier* to be exhausted. Of course, this attack requires that the attacker and the other clients share the same medium (e.g. a LAN).

The filter can limit the impact of this attack as the forged *Leave* message will be only handled in its turn, that is, the legitimate clients still have their messages serviced thank to the scheduling process. However, the filter cannot totally counter such an attack and other extensions are need.

7.5 Conclusion

This chapter presents our proposal, a simple yet efficient filtering approach to thwart some DoS attacks that are based on IGMP or MLD flooding, and that threaten the whole operator's infrastructure. A key feature of our proposal is that it follows a realistic

7. OUR PROPOSAL: A FILTERING APPROACH TO MITIGATE IGMP/MLD FLOODING ATTACKS

and pragmatic approach, and in particular it does not require any modification to the existing, widely deployed protocols.

Chapter 8

Filter Dimensioning: Theoretical Study

Contents

8.1	Key Parameters	61
8.2	Dimensioning the Unknown Clients Queue: a Theoretical Study	64
8.2.1	Deterministic Arrival Model	64
8.2.2	Poisson Arrival Model	65
8.3	Dimensioning the Known Client Queues	66
8.3.1	Estimating the Number of Known Client Queues	67
8.3.2	Estimating the Number of the Waiting Packets for the Known Client Queues	68
8.3.3	Evaluating the Waiting Time	70
8.4	Conclusion	70

Dimensioning the system correctly is essential for an optimal behavior of the filter. An optimal behavior means a good protection against DoS attacks, but at the same time minimized size effects: small false-positive error rate (when legitimate clients requests will be blocked), and little impacts (ideally no impact) on the IGMP/MLD mechanism (e.g. legitimate IGMP *Report* and *Leave* requests should not be delayed).

In this chapter we present a theoretical study that allows to easily dimension the filter and to find the different initializing parameters as a function of the traffic arrival models and the underlying infrastructure.

8.1 Key Parameters

There is a small set of key parameters that must be initialized suitably for a given environment:

8. FILTER DIMENSIONING: THEORETICAL STUDY

- the size of the Unknown Client FIFO Queue (UCQ): this queue cannot be too small, since many new but legitimate clients may be interested by joining a session at a given time, for instance because they have been informed of its existence synchronously;
- the unknown client queue service rate: this is the number of packets coming from new clients that can be accepted by the system per time unit. Its value is a balance between the reactivity of the system, in front of new legitimate clients, and the protection in front of an attacker that would spoof the source IP address of the IGMP packets it generates;
- the scheduling rate for known client packets: this is the number of IGMP/MLD packets from known clients that can be accepted by the round-robin scheduling thread per time unit. Its value is directly related to the processing capabilities of the first hop multicast router (CPU and memory), and that of the whole multicast routing infrastructure since IGMP/MLD packets may trigger the creation or pruning of multicast branches in the operator's infrastructure;
- the maximum number of waiting IGMP/MLD packets for different groups per known client: this parameter protects the filter from an attacker known by the system that would generate a large number of IGMP/MLD packets. Since it is expected, in a correctly dimensioned filter, that IGMP/MLD packets for known clients do not stay too long in the filter, this maximum number of waiting packets per client should be low;
- the maximum number of known clients managed by the system: this parameter should be adjusted according to the simultaneous number of potential clients, which is usually known by the operator. An upper bound exists (especially in IPv4 environments), namely the number of IP addresses made possible by the IP subnet. Of course the maximum number of clients should not exceed the processing capabilities of the filter (CPU and memory);
- the purging period: the purging period must be adapted to the specificities of the target environment, and in particular whether only IGMP version 3 hosts are deployed or not (section 7.1.3);

There are clearly trade-offs to find when initializing the filter. These trade-offs must take into account various external parameters that are specific to the target environment: the simultaneous number of potential clients, the number of groups they may be interested in, the frequency with which they may join or leave these groups (e.g. because of zapping in a TV/ADSL environment).

Obtaining the optimal parameters depends largely on the arrival rate of the IGMP/MLD traffic, which in turns is affected by the following factors:

- The underlying infrastructure:
The underlying infrastructure affects the number of IGMP/MLD reports to be

sent by the multicast participants in response to the periodic *Queries*. On shared medium networks (such as Ethernet) where all hosts share a common transmission channel only one host sends a *Report* in response to a specific *Query*, in contrast, on a point-to-point oriented packet-switched infrastructure such as ATM VCs or PPP connections, the IGMP Report suppression mechanism is not applicable and every host should answer when receiving a *Query*.

- The timer settings of the IGMP/MLD protocols:
Many timers have been defined in the IGMP/MLD protocols and most of the IGMP/MLD timers are configurable 3.2.1. These timers affect directly the burstiness of IGMP traffic on a subnet.
- The users behavior:
The IGMP/MLD arrival rate depends on how users respond to the presence of multicast flows. The membership in a multicast group on a given interface is dynamic and changes over time as hosts join and leave the group. Some multicast services such as TV/ADSL or on-line games trigger a dynamic behavior more than other services such as VoD or distribution of a new software. Moreover, this behavior follows other factors such as the session duration and the application popularity. The few works done to model the multicast participants behavior [10] [39] [11] show that there is a steady increase in user arrivals as sessions start, and a steady decrease in users as sessions conclude. Further, there is a correlation between the number of users and the number of requested groups. This behavior can be modeled by a Poisson process.

In the following evaluation we consider two cases: in the first case the IGMP packets arrive at a constant rate and in the second the packets arrive with a Poisson distribution. The first case could be considered as a special case of the second. We are in particular interested by studying how the arrival distribution affects the different parameters of the filter.

We define the following parameters:

- K : size of the UCQ;
- λ : average packet arrival rate coming from unknown clients (we assume here that each packet has a different client source address);
- μ : service rate for the UCQ;
- ρ : traffic intensity, equal to λ/μ . Traffic intensity is a measure of the congestion of the system. If it is near zero, there is very little queuing. As the traffic intensity increases (near 1 or even greater than 1), queuing increases;
- G_{max} : maximum number of groups authorized per clients.
- s : scheduling rate.

8. FILTER DIMENSIONING: THEORETICAL STUDY

- *purge_period*: the interval between two purges (where inactive receiver states are removed from the system).

In this evaluation we consider the system composed of two components: the first is the Unknown Client Queue (UCQ) and the second is the set of known client queues.

8.2 Dimensioning the Unknown Clients Queue: a Theoretical Study

We now explain, through a simple model, how the Unknown Client Queue (UCQ) can be initialized, taking into account some statistical information on the target environment that the network operator can easily define.

8.2.1 Deterministic Arrival Model

If the IGMP *Reports* arrive with a constant rate λ , then we model the UCQ using a $D/D/1/K$ queue [34]. $D/D/1/K$ is a queue of the single server type, with a deterministic arrival rate λ , a deterministic service rate μ and total of K buffers. The arriving packets are serviced according to a FCFS (First Come First Served) rule. Theoretically, the number of packets in the queue at time t , $n(t)$ is given by the equation:

$$\begin{aligned}
 n(t) &= \text{number of arrival in the interval } (0, t] \\
 &\quad - \text{number of served packets in the interval } (0, t] \\
 &= \text{int}\left(\frac{t}{1/\lambda}\right) - \text{int}\left(\frac{t - 1/\lambda}{1/\mu}\right) \\
 &= \text{int}(t\lambda) - \text{int}\left(\mu t - \frac{\mu}{\lambda}\right)
 \end{aligned} \tag{8.1}$$

where $\text{int}(x)$ is the greatest integer $\leq x$. This equation is valid only up until the first balk occurs. So, if t_i is the time until the first balk, this equation 8.1 is written as follows:

$$n(t) = \begin{cases} 0 & \text{if } t < 1/\lambda \\ \text{int}(t\lambda) - \text{int}\left(\mu t - \frac{\mu}{\lambda}\right) & \text{if } 1/\lambda \leq t < t_i \\ K & \text{if } t \geq t_i \end{cases} \tag{8.2}$$

AWT the waiting time in the queue for the n arrived packet is given by the recurrence equation (Figure 10.2):

$$AWT^{n+1} = AWT^n - 1/\mu - 1/\lambda \tag{8.3}$$

$$AWT^n = \begin{cases} (1/\mu - 1/\lambda)(n - 1) & \text{if } n < \lambda t_i \\ (K - 1)1/\mu & \text{if } n \geq \lambda t_i \end{cases} \tag{8.4}$$

The equations 8.1 and 8.3 show that the $n(t)$ and AWT^n go to zero when $\lambda \leq \mu$ and they increase in the opposite case. This increase however is limited because of the finite size of the buffer.

8.2 Dimensioning the Unknown Clients Queue: a Theoretical Study




Figure 8.1: Successive waiting times.

8.2.2 Poisson Arrival Model

As the service rate of the new clients queue is deterministic, the $M/D/1/K$ presents a suitable model. The $M/D/1/K$ is a queue of the single server type, with a Poisson arrival, a deterministic service time and a total of K buffers. For simplicity purpose, instead of modeling the new client queue as a $M/D/1/K$ queue we model it as a $M/M/1/K$ queue. That is, as a queue of the single server type, with Poisson arrival, exponential service time and a total of K buffers. For $M/M/1/K$ the probability of having n packets in the queue is [34]:

$$P_n = \begin{cases} \frac{(1-\rho)\rho^n}{1-\rho^{K+1}} & \text{if } \rho \neq 1 \\ \frac{1}{1+K} & \text{if } \rho = 1 \end{cases} \quad (8.5)$$

Where $\rho = \lambda/\mu$.

The probability of not overflowing the queue, also known as the acceptance probability, is:

$$P = 1 - P_K \quad (8.6)$$

Figure 8.2 shows that for traffic intensity $\rho = 0.5$ we need at least $K = 5$ buffers to have an




Figure 8.2: Acceptance Probability (P_K) for the Unknown Client Queue (UCQ).

acceptance probability approaching 1. On the opposite, when $\rho > 1$, whatever the UCQ size is, the acceptance probability is small, and there will always be a significant number of lost

8. FILTER DIMENSIONING: THEORETICAL STUDY

packets. The packet loss rate is given by:

$$\lambda_{lost} = \lambda(1 - P) \quad (8.7)$$

The mean number of packets in the queue, N , is:

$$N = \begin{cases} \frac{\rho}{1-\rho} - \frac{(K+1)\rho^{K+1}}{1-\rho^{K+1}} & \text{if } \rho \neq 1 \\ \frac{K}{2} & \text{if } \rho = 1 \end{cases} \quad (8.8)$$

The mean number of packets in the UCQ depends only on the traffic intensity and the capacity of the queue. Figure 8.3 shows the mean number of packets for varying ρ and for $k = 10$.




Figure 8.3: Mean Number of Packets in the UKQ.

The average waiting time, AWT , in the UCQ is (Little law):

$$AWT = \frac{N}{\lambda P} \quad (8.9)$$

This study demonstrates how to dimension the UCQ. The equation 8.6 shows how to dimension the UCQ using the parameters: K , λ and ρ . More specifically, the K and μ parameters can be deduced from the equation 8.5 after setting a target acceptance probability value.

Example

Let's imagine that new clients send IGMP packets with an average arrival rate $\lambda_l = 15$ pps, following a Poisson distribution. In this case, for $\mu = 15$ and $K = 20$, the acceptance probability for the legitimate clients equals to 95 % (eq. 8.6).

8.3 Dimensioning the Known Client Queues

The previous study allows to dimension the UCQ. We now look at the second component of the filter: the known clients queues.

8.3.1 Estimating the Number of Known Client Queues

To find the number of known clients as a function of time, we first introduce a new variable: $ACT_client(i)$ which represents the activity of a client i , that is, the rate at which a known client i sends requests. If a client i becomes known (active) at an instant t where $nPurge_period \leq t \leq (n+1)Purge_period$, this client i keeps its place in the system as long as: $ACT_client(i) > Purge_frequency$. Let's that the clients' requests arrive with a Poisson distribution with rate λ and the total number of clients is N . Then $ACT_client(i) = \lambda/N$ (the rate at which each individual clients send requests). In this section, we are interested in calculating the number of clients between each purge. The number of clients increases exponentially between two successive purges and it decreases when the purging thread awakes. Let consider two successive purges: $(K-1)\Delta$ and $K\Delta$, where Δ is the purging period. Let $A_{k\Delta+}$ be the number of clients just after the purge and $A_{k\Delta-}$ the number of clients just before the purge.

Calculating the Number of Clients After the Purge $A_{k\Delta+}$

$A_{k\Delta+}$ is the number of clients kept in the system after the purge. Under the assumption that the arrival requests are modeled as a Poisson process:

$$P(\text{one client inactive}) = P(X = 0) = e^{-\frac{\lambda\Delta}{N}} \quad (8.10)$$

Let $P(A_{k\Delta+} = l)$ be the probability that the number of clients equals to l after the purge. Then $P(A_{k\Delta+} = l)$ is the probability that l clients have send at least one request before the purge and $N-l$ haven't sent any requests. These $N-l$ will be deleted. Then:

$$P(A_{k\Delta+} = l) = \binom{N}{l} (1 - e^{-\frac{\lambda\Delta}{N}})^l e^{-\frac{\lambda\Delta}{N}(N-l)} \quad (8.11)$$

The equation follows the binomial distribution, the mean is given by the equation:

$$E[A_{k\Delta+}] = N(1 - e^{-\frac{\lambda\Delta}{N}}) \quad (8.12)$$

Calculating the Number of Clients before the Purge $A_{k\Delta-}$

$$P(A_{k\Delta-} = m) = \sum P(m)P(m|A_{(k-1)\Delta+} = l) \quad (8.13)$$

After the reduction we find:

$$P(A_{k\Delta-} = m) = \binom{N}{m} (1 - e^{-\frac{2\lambda\Delta}{N}})^m e^{-\frac{2\lambda\Delta}{N}(N-M)} \quad (8.14)$$

Which gives in turn:

$$E[A_{k\Delta-}] = N(1 - e^{-\frac{2\lambda\Delta}{N}}) \quad (8.15)$$

The equations 8.15 and 8.12 give the number of clients before and after the purge. The following example illustrates the use of these equations.

8. FILTER DIMENSIONING: THEORETICAL STUDY

Example

Let's consider a scenario where 1000 legitimate sources send 15 IGMP *REPORT* per second and the *purge_period* = 80s. If we suppose that the requests are sent randomly over the time, then in average: $E(A_{k\Delta+}) = 696$ and $E(A_{k\Delta-}) = 907$

The figure 8.4 shows this process.



Figure 8.4: Number of known clients

8.3.2 Estimating the Number of the Waiting Packets for the Known Client Queues

In this section we study how to estimate the number of waiting packets in the known client queues WP_{known}

The system must be tuned to make the number of waiting packets in the known clients queues very small, even negligible. This number increases steadily if the system is under attack.

In the following, we study how to quantify the number of waiting packets WP_{known} as a function of the parameters of the filter. To that goal, we consider again that the legitimate clients' requests arrive with a Poisson distribution with a total rate λ , the total number of legitimate clients is N_{leg} and N_{att} the number of intruders in the system.

To that goal, we model all the legitimate clients queues as $M/M/1/N_{leg}G_{max}$. Then the average number of packets for this queue is given by the equation:

$$L_{legitimate} = \begin{cases} \frac{\rho_{leg}}{1-\rho_{leg}} - \frac{(N_{leg}G_{max}+1)\rho_{leg}^{N_{leg}G_{max}+1}}{1-\rho_{leg}^{N_{leg}G_{max}+1}} & \text{if } \rho \neq 1 \\ \frac{N_{leg}G_{max}}{2} & \text{if } \rho = 1 \end{cases} \quad (8.16)$$

Where

$$\rho_{leg} = \lambda_{leg} / \left(N_{leg} \frac{s}{N_{leg} + (N_{att})} \right) \quad (8.17)$$

Now, we model all the intruder queues as $M/M/1/N_{leg}G_{max}$, then the average number of

8.3 Dimensioning the Known Client Queues

packets for this queue is given by the equation:

$$L_{attack} = \frac{\rho_{att}}{1 - \rho_{att}} - \frac{(N_{att}G_{max} + 1)\rho_{att}^{N_{att}G_{max}+1}}{1 - \rho_{att}^{N_{att}G_{max}+1}} \quad (8.18)$$

Where

$$\rho_{att} = \lambda_{att} / \left(N_{att} \frac{s}{N_{leg} + N_{att}} \right) \quad (8.19)$$

This makes:

$$WP_{known} = L_{legitimate} + L_{attack} \quad (8.20)$$

Figure 8.5 shows the increase of waiting packets with the increase of traffic intensity. This increase is limited by the number of total available buffers. The figure 8.6 depicts the increase

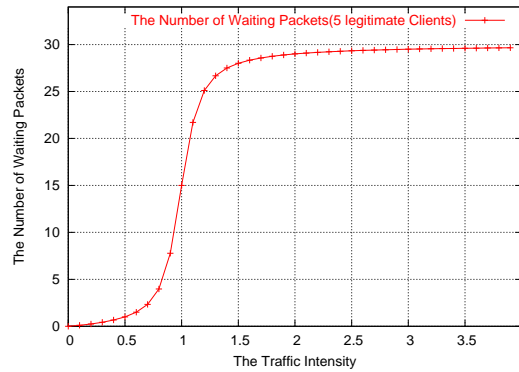


Figure 8.5: The number of waiting packets as a function of traffic intensity

of the waiting packets with the increase of the accepted attackers in the system.

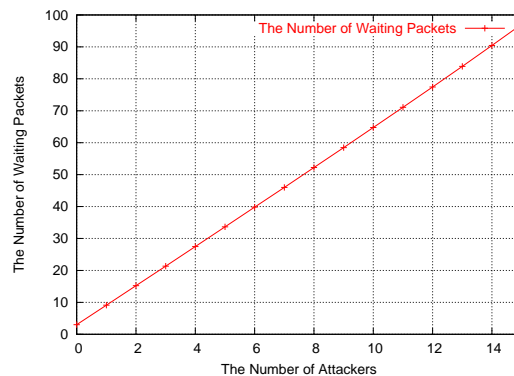


Figure 8.6: The number of waiting packets as a function of number of attackers

8. FILTER DIMENSIONING: THEORETICAL STUDY

8.3.3 Evaluating the Waiting Time

We obtain the waiting time in the system using again the Little's law:

$$AWT' = WP_{known}/\lambda_{eff} \quad (8.21)$$

where λ_{eff} is the effective arrival rate and can be obtained from the equations 8.6 and 8.7.

$$\lambda_{eff} = \lambda_{att}^{eff} + \lambda_{leg}^{eff} \quad (8.22)$$

Figure 8.7 depicts the evolution of the waiting time in the system with the increase of the number of attackers queues. This simulation was done for the values: $N_{leg} = 100$, $\lambda_{leg} = 15$, $\lambda_{att} = 100$, and $s = 20pps$.




Figure 8.7: Evaluating the Average Waiting Time

8.4 Conclusion

In this chapter we present a theoretical study that allows to dimension the filter and to find the different initializing parameters as a function of the traffic arrival models and the underlying infrastructure. Two types of IGMP/MLD traffic arrival models are studied, deterministic and Poisson models. The first case could be considered as a special case of the second. We are in particular interested in studying how the arrival distribution affects the different parameters of the filter.

In this evaluation we consider that the system is composed of two components: the first is the Unknown Client Queue (UCQ) and the second is the set of known client queues. We depicted the equations that allows to dimension and parameterize each component of the system.

Chapter 9

Evaluation

Contents

9.1	Introduction	71
9.2	Experimental Environment	72
9.3	Tests Scenarios	73
9.4	IGMP Traffic Sent at a Constant Rate	74
9.4.1	Filter Parameters	74
9.4.2	Results Test 1: Single Forged IP Addresses	75
9.4.3	Results for Test 2: 10 Forged IP Addresses	77
9.4.4	Results for Test 3: 1000 Forged IP Addresses	80
9.5	IGMP Traffic Sent with a Poisson Distribution	80
9.5.1	Introduction to Poisson Distribution	81
9.5.2	Generating Traffic with Poisson Distribution	81
9.5.3	Filter Parameters	82
9.5.4	Results Test 1: Single Forged IP Addresses	82
9.5.5	Results Test 2: Ten Forged IP Addresses	84
9.5.6	Results for Test 3: 1000 Forged IP Addresses	87
9.6	Conclusion	88

In this chapter, we evaluate experimentally the efficiency of our filter. We compare the experimental and the theoretical results presented in the previous chapter. Finally, we conclude with a detailed discussion of several key issues that emerged from this study.

9.1 Introduction

In chapter 5, we have analyzed the group management protocols specific threats. In the light of this analysis, we classified and discussed these threats according to their harmfulness to the operator infrastructure. Our study shows that there is an emerging need to cope with attacks (intentional or not) that aim to paralyze the querier router. Namely, by flooding the router with IGMP/MLD *Reports* for a large number of groups (See section:5.4).

9. EVALUATION

Our filter proposal was dedicated response to this requirement. A simple yet efficient filtering approach to thwart DoS attacks that are based on IGMP or MLD flooding, and that threaten the whole operator's infrastructure. The main objective of our filter, as previously illustrated, is to design a realistic and pragmatic approach, which does not require any modification to the largely deployed infrastructure. This proposal aims as well to protect the well-behaved clients against ill-behaved ones.

To evaluate the efficiency of the filter and its capacity to respond to its design goals, we have conducted the experimental study presented in this chapter. This experimental study has the following main objectives:

1. *Study the possibility to integrate the filter with the existing Infrastructure. To that goal, we use in our evaluation a commercial cisco router.*
2. *Study the benefits of by the filter in front of IGMP flooding attacks, to this goal, we measure the memory consumption on the router in the case of attack with and without filter.*
3. *Evaluating how the filter can protect well behaved clients against ill-behaved ones.*

We conduct our experimental evaluation in the light of the theoretical study presented in the precedent chapter. The theoretical study allows to parameterize the filter in order to achieve an optimal filtering efficiency.

9.2 Experimental Environment

Our experimental environment consists of a small testbed. This testbed consists of one end-host, running the IGMP traffic generator for both legitimate clients and the attacker, the filter, a PIM DR and a PIM RP.

The filter is attached to the same Ethernet LAN, as well as the first hop multicast router, a Cisco 7500 RSP 12.2, running PIM-SM and acting as the IGMP *Querier*. Connected to this router is the RP, a Linux router running PIM-SM. The host running the filter is equipped with Intel Xeon-2.00 GHz and 1 GB RAM, while the other PCs are equipped with Intel Pentium IV-2.5 GHz processors. The operating system is Mandrake Linux v10.0 [12.4](#).

The IGMP timers are tuned as has been described in the section [3.2.1](#):

- *Robustness Variable = 2.*
- *Query Interval:* the interval between General Queries sent by the Querier equals to 125 seconds.
- *Query Response Interval:* the Max Response Time inserted into the periodic General Queries equals to 10 seconds.
- *Group Membership Interval:* the Group Membership Interval is the amount of time that must pass before a multicast router decides there are no more members of a group on a network. $\text{Group Membership Interval} = ((\text{the Robustness Variable}) \times (\text{the Query Interval})) + (\text{one Query Response Interval})$. Hence, in our test the Group Membership Interval = 260s.
- *Last Member Query Interval:* the Last Member Query Interval is the Max Response Time inserted into Group-Specific Queries sent in response to Leave Group messages, and is also the amount of time between Group-Specific Query messages. Default: 1 second.

- *Unsolicited Report Interval*: The Unsolicited Report Interval is the time between repetitions of a host's initial report of membership in a group. Default: 10 seconds.
- *Version 1 Router Present Timeout*: The Version 1 Router Present Timeout is how long a host must wait after hearing a Version 1 Query before it may send any IGMPv2 messages. Value: 400 seconds.




Figure 9.1: The Testbed.

9.3 Tests Scenarios

In this evaluation we carried out three kinds of tests. For the first test, the attacker uses a single forged source IP address. This test corresponds to the naive flooding DoS attack (section 7.4.1). For the second test, the attacker uses a set of 10 addresses, and for the third one he uses a set of 1000 addresses. The last two tests are therefore representative of severe attacks, where distinguishing legitimate traffic from attacker's traffic becomes complex or impossible. These tests are also representative of distributed DoS attacks (DDoS) with respectively 10 and 1000 compromised hosts, each of them launching a naive flooding DoS attack.

Each test consists from three periods: (1) a first period where there are only legitimate clients, (2) a second period where a flooding DoS attack is launched, and in parallel new clients arrive, and (3) a third period where there are only legitimate clients.

The first and second tests last 540 seconds and use the following scenario:

- [0; 540s] (whole test): 500 legitimate sources send 15 IGMP *REPORT* per second over 50 groups during the whole test duration;
- [180; 540s]: in parallel, another 500 legitimate sources send 5 IGMP *REPORT* per second over 50 groups different from the groups reported by the first legitimate clients. The goal of introducing new clients here is to evaluate the impact of the attack and the filter on new unknown but legitimate clients;
- [180; 360s]: in parallel, an attacker spoofs 1 source address in the first test and 10 in the second test. The attacker sends 100 IGMP *REPORT* packets per second, with group addresses chosen randomly within a set of 500 addresses other than those requested by the legitimate clients.

9. EVALUATION

The third test lasts a total of 718 seconds, attack begins after 179 seconds and finishes at time 638 second. The number of packets sent is the same, the only difference compared to the previous two tests being the number of spoofed IP addresses used by the attacker.

In this evaluation we consider two cases: in the first case the IGMP traffic is sent with a constant rate and in the second case the IGMP traffic follows a Poisson distribution. The first choice has been done in the beginning to simplify the evaluation of the filter while the second choice seems to be more realistic since it models more accurately the IGMP traffic [10] [39] [11].

9.4 IGMP Traffic Sent at a Constant Rate

9.4.1 Filter Parameters

To tune the filter, we consider the previous theoretical study. As has been previously mentioned, finding the good parameters is essential for optimal traffic filtering.

- Initializing the UCQ

To find the suitable parameters for the UCQ, we use the study presented in section 8.2.1. The equations 8.1 and 8.2 show that the size of the buffer does not impact the mean number of waiting packets, but it plays a role in the waiting time of the packets. Hence, we choose the size of the unknown client queue to be 20 packets. Hence, to make the number of waiting packets $n(t)$ approaches zero when $\lambda = 15$ pps needs that $\mu = 15$ pps (from the equation 8.1).

- Initializing the known client queues

- scheduling rate for known clients' packets (s)

From the equation 8.18, for $\rho > 1$ we guarantee that the waiting packets in the known client queues be very small. To that goal we choose $s = 15$.

- maximum number of authorized groups per client (G_{max})

This number should be specified by the operator in function of the number of proposed groups and the capacity of the router. In the following evaluation we choose G_{max} to be 6.

- purging period

Equations 8.15, 8.11 give the number of clients as a function of the arrival rate λ and total number of clients before and after the purge. As has been illustrated in the section 8.3 the clients activity $ACT_{client}(i)$ plays an important role in finding the purging period. In our test, the 500 legitimate clients send IGMP traffic with 15pps thus each client sends traffic each 33s approximately. By choosing the purging period to be 80s we guarantee that the legitimate clients will not be deleted prematurely.

In summary, the filter is initialized as follows:

- size of the unknown client queue (K) = 20 packets
- unknown client queue service rate (μ) = 15 pps
- scheduling rate for known clients' packets (s) = 20 pps
- maximum number of waiting packets per client at a given time (G_{max}) = 6 packets
- purging period = 80 s

9.4.2 Results Test 1: Single Forged IP Addresses

Figure 12.5(a) shows the traffic entering the filter and resulting from both the legitimate clients and the attacker, and figure 12.5(b) the traffic leaving the filter and entering the first hop multicast router. The attacker receives a share of the outgoing flow that is significantly lower (≈ 100 times) than its incoming traffic rate. Indeed, once the attacker has joined the "known clients" list, a maximum of $max_nb_waiting_pcks_per_client = 6$ packets will be kept in its context. The attacker's packets, that are much in excess (100 pps are sent), will make the attacker's list of 6 packets overflow, and the attack will naturally be filtered. The attacker's flow receives at most a $\frac{max_nb_waiting_pcks_per_client}{total_number_active_pcks}$ fraction of the outgoing flow. Since the legitimate clients send a total of 20 packets per second (once all new clients of the second group are known by the system), $total_number_active_packets = 26$, and the scheduling rate equals to 20 pps. Tgen, the attacker's share of the outgoing traffic is therefore $\leq 1/20 = 0.05$.

Besides, figure 12.5(b) also shows that new clients can still be accepted during the attack. This proves the effectiveness of the filter in limiting the IGMP flooding traffic while limiting the side effects on legitimate users. Let's now look at the router. Figure 12.7 displays the

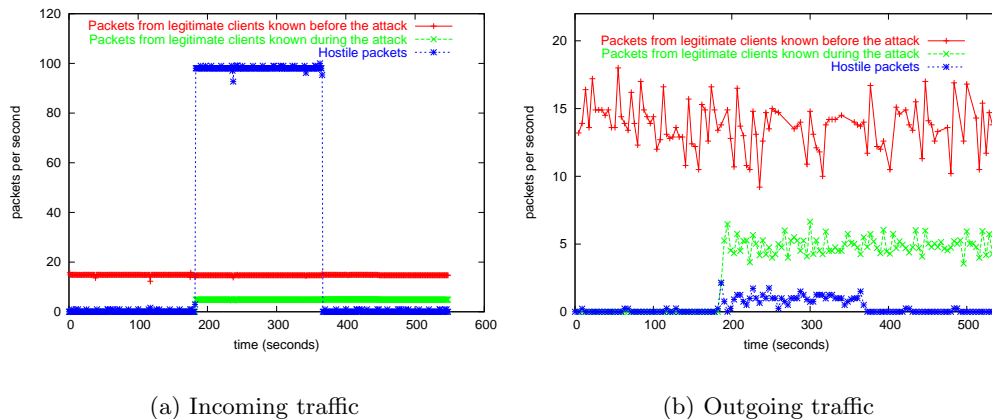


Figure 9.2: Traffic Entering and Leaving the Filter (Test 1).

instantaneous memory consumption on the Cisco router with or without the filter¹. This figure highlights the high benefits of the filter since the impact of the flooding attack is rather limited. Note that the total time when the router exhibits a higher memory consumption is larger (up to time ≈ 740 seconds) than the total test duration (540 seconds). This is caused by the various IGMP/PIM-SM protocol timers that oblige the router to keep some state information for a given group a few minutes after having received the associated IGMP *REPORT* message. Figures 12.6(a) and 12.6(b) show the PIM-SM control traffic leaving the Cisco router towards the root of the multicast tree without and with the filter. These curves show once again the benefits of the filter since the control traffic (PIM Join/Prune messages) is significantly reduced (peaks amount to 16 pps with the filter instead of 105 pps without).

Let's now look at the filter. The figure 9.5 shows the number of known client queues. The

¹This memory consumption has been obtained through the following command:
`show ip mroute count | include routes`

9. EVALUATION


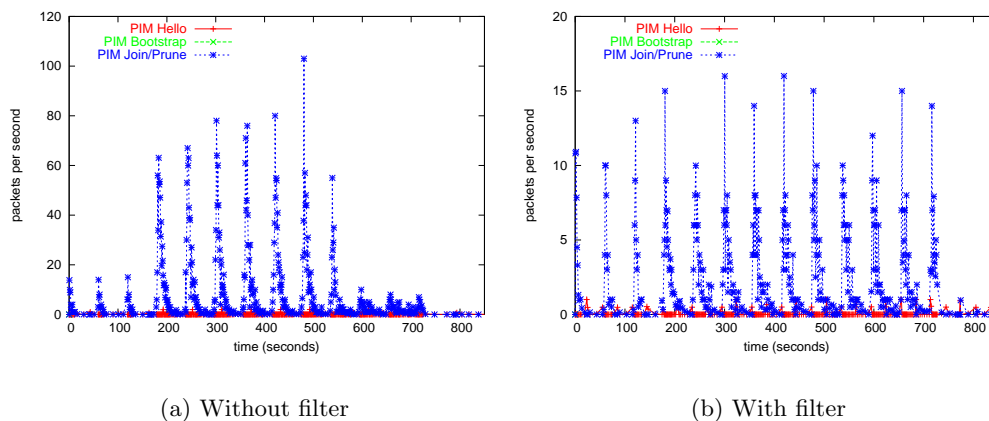


Figure 9.3: Memory Consumption of the Cisco Router Without/With Filter (Test 1)



(a) Without filter

(b) With filter

Figure 9.4: PIM Control Traffic Without/With Filter (Test 1).

successive decreases in the number of queues correspond to the *purging thread frequency* = 80s, these values are so close from the theoretical values of $A_{k\Delta+} = 696$ and $A_{k\Delta-} = 907$ that were obtained by modelization in the example1 in the section 8.3. Figure 9.6 shows the number of waiting packets in the known clients queues. As been illustrated in section 8.3.2, by suitably parameterizing the filter, the mean number of waiting packets in the system should be close to zero when the system does not suffer any attack. However, figure 9.6 shows that the mean number of waiting packets could increase suddenly even if all the packets in the system are legitimate. This increase is caused by a scheduling problem between the different threads in the program. Executing the program on other operating systems such as Solaris which supports explicitly threads scheduling features could largely improve this result.

Figure 9.7 depicts the average waiting time in the system. The average waiting time in this case has very little impact on the IGMP leave and join in comparison to benefits gained from the filter. The increase in the waiting time when the attack takes place comes especially from the implementation aspects as in the case of an attack the packets capturing threads works

9.4 IGMP Traffic Sent at a Constant Rate

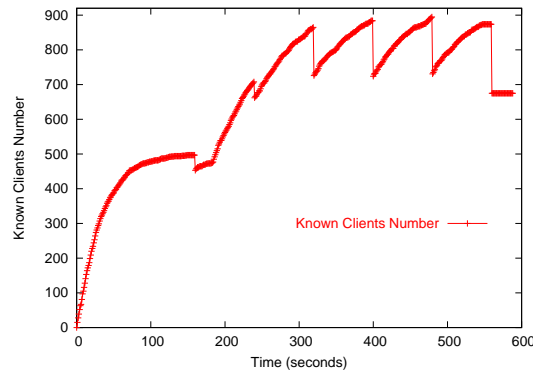


Figure 9.5: Known Clients (Test 1)




Figure 9.6: Waiting Packets (Test 1)

more than the others threads.

9.4.3 Results for Test 2: 10 Forged IP Addresses

With test 2, Figure 9.8 shows that the number of scheduled hostile packets increases only slightly when the attacker uses 10 forged source IP addresses. The benefits when considering the memory consumption on the Cisco router (Figure 9.9) are here also very significant and close to those observed with test 1. The filter is therefore very effective when the attacker spoofs a limited number of addresses, or said differently, the filter is very effective in case of a Distributed DoS attack involving a small number of hosts (10 in this test). Figure 9.10 shows the number of known client queues. Figure 9.11 shows the number of waiting packets. The figure 9.12 shows the average waiting time. The impact on the IGMP leave and join latency is still acceptable in this test. However we observe that number of waiting packets and the average waiting time in the system increase significantly compared to the first test. This increase is explained by the increase of hostile queues accepted in the system. Where for these clients $ACT_client(i)$ is much higher than for legitimate clients. This result could be inferred directly

9. EVALUATION




Figure 9.7: Average Waiting Time (Test 1)

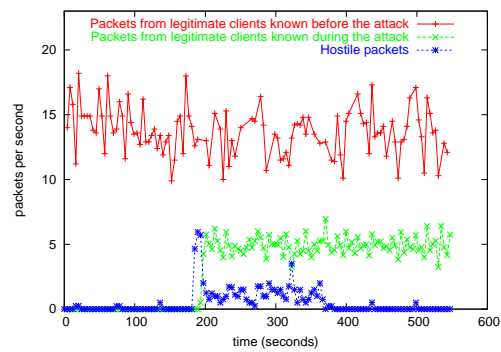


Figure 9.8: Traffic Leaving the filter (Test 2).




Figure 9.9: Memory Consumption of the Cisco Router (Test 2).

from the equation 8.19.

9.4 IGMP Traffic Sent at a Constant Rate

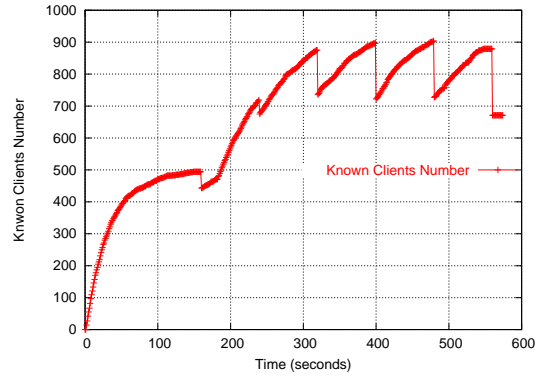


Figure 9.10: Known Clients (Test 2)




Figure 9.11: Waiting Packets (Test 2)




Figure 9.12: Average Waiting Time (Test 2)

9. EVALUATION

9.4.4 Results for Test 3: 1000 Forged IP Addresses

Let's now consider test 3, where the attacker uses a high number (1000) of forged source addresses. Figure 12.8(b) shows that the attacker's traffic is progressively accepted by the filter, once the forged source address become known by the system. The attacker succeeds in obtaining a little bit more than half of the outgoing traffic. It also prevents almost all new legitimate clients from being accepted during the attack. These latter must wait the end of the attack before reaching the known client status, and having their traffic go through the filter. Figure 9.14 shows that the filter does not significantly change the memory consumption for the

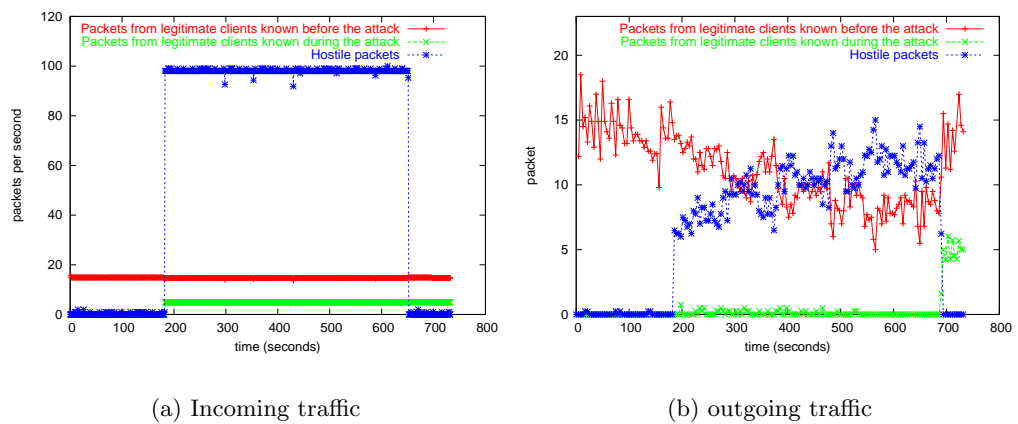


Figure 9.13: Traffic Entering and Leaving the filter (Test 3).

Cisco router. This is caused by the high number of different multicast groups for which the attacker sends reports and that will progressively create multicast states in the router since the filter does not remove them as efficiently as before, with tests 1 and 2.

When the number of hostile clients entering the system increases over the scheduling rate and especially if the duration of the attack is large enough, the legitimate clients could be largely penalized. This is visible in Figure 12.8.

However, the DoS attack does not impact the whole multicast routing infrastructure of the operator. In summary, the experimental and theoretical evaluations both show that the filter can up to certain degree sustain the availability of the multicast routing infrastructure and achieve a fair service guarantee for the legitimate clients. However, other extensions are needed to improve the filtering efficiency. We will discuss some perspectives to improve the filtering efficiency in the future works.

9.5 IGMP Traffic Sent with a Poisson Distribution

In this section we consider that the IGMP traffic arrives with a Poisson distribution. This assumption seems to be more realistic to modelize the IGMP traffic [10] [39] [11].

9.5 IGMP Traffic Sent with a Poisson Distribution




Figure 9.14: Memory Consumption of the Cisco Router (Test 3).

9.5.1 Introduction to Poisson Distribution

The Poisson distribution is a mathematical rule that assigns probabilities to the number of occurrences. The probability density function (PDF) of a Poisson variable is given by:

$$P_{X=x} = \frac{e^{-\lambda} \lambda^x}{x!} \quad (9.1)$$

λ is the shape parameter which indicates the average number of events in the given time interval. The random variable of the intervals $t_i = x_i - x_{i-1}$ between these events is exponentially distributed with PDF:

$$f(t) = \lambda e^{-\lambda t} \quad (9.2)$$

Therefore, to generate these intervals we use the equation:

$$f(t) = -\frac{1}{\lambda} \ln(x), \quad 0 \leq x < 1 \quad (9.3)$$

9.5.2 Generating Traffic with Poisson Distribution

Generating traffic with a Poisson distribution means that the delay between the generated packets follows a Poisson distribution. In ANSI c, the function `rand()` satisfies: $0 \leq \text{aninteger} \leq \text{RAND_MAX}$. Thus, to obtain x we use: $x = \text{rand}() / (\text{double})\text{RAND_MAX} + 1.0, 0 \leq x < 1$.

The following code allows to specify the intervals between packets to obtain traffic with a Poisson distribution.

9. EVALUATION

```
/* generates an exponential random variable
given the mean in micro seconds*/
double Rnd(double mean)
{
    double rnd,exp;
    rnd=(double)rand()/((double) RAND_MAX + 1.0);
    exp=-mean*log(1 - rnd);
    return(exp);
}
```

9.5.3 Filter Parameters

- Initializing the UCQ

To find the suitable parameters for the UCQ, we use the study presented in section 8.2.2. The equation 8.8 shows that the mean number of packets in the UCQ depends only on the traffic intensity and the capacity of the queue.

Hence, if the size of the unknown client queue to be 20 packets, then for $\lambda = 15pps$, and to make $\rho \leq 1$ (so all legitimate packets will be accepted, μ should be greater or equals to $15pps$. By choosing $\mu = 15$ then the mean number of waiting packets in the UCQ will be $10p$

- Initializing the known client queues

- scheduling rate for known clients' packets (s)

From the equation 8.18 for $\rho > 1$ we guarantee that the waiting packets in the known client queues be very small. To that goal we choose $s = 15$.

- maximum number of authorized groups per client (G_{max}): This number should be specified by the operator in function of the number of proposed groups and the capacity of the router. In the following evaluation we choose G_{max} to be 6.

- purging period

The equations 8.15 and 8.11 give the number of clients before and after the purge as a function of the arrival rate λ and total number of clients. Tuning this parameters allows the operator to determine the total number of queues in the system. We choose the purging period to be $80s$.

In summary, the filter is initialized as follows

- size of the unknown client queue (K) = 20 packets.
- unknown client queue service rate (μ) = 15 pps.
- scheduling rate for known clients' packets (s) = 20 pps.
- maximum number of waiting packets per client at a given time (G_{max}) = 6 packets.
- purging period = 80s.

9.5.4 Results Test 1: Single Forged IP Addresses

Figure 9.15(a) shows the traffic entering the filter and resulting from both the legitimate clients and the attacker, and figure 9.15(b) the traffic leaving the filter and entering the first hop multicast router. In the case of a Poisson arrival, the attacker receives a share of the outgoing flow

9.5 IGMP Traffic Sent with a Poisson Distribution

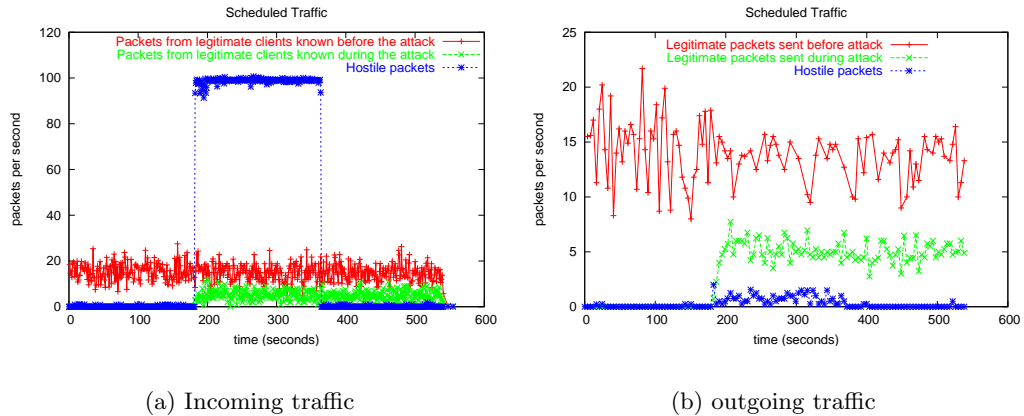


Figure 9.15: Traffic Entering and Leaving the filter (Test 1-Poisson).

that is significantly lower (≈ 100 times) than its incoming traffic rate. Besides, figure 9.15(b) also shows that new clients can still be accepted during the attack. This proves the effectiveness of the filter in limiting the IGMP flooding traffic while limiting the side effects on legitimate users. Figure 9.16 shows the number of known client queues. Again, the successive decreases in




Figure 9.16: Known Clients Number (Test 1-Poisson)

the number of queues correspond to the frequency of the purging thread which equals to 80s. These experimental values are so close from the values of $A_{k\Delta+} = 696$ and $A_{k\Delta-} = 907$ that were obtained by modelization in the example1 in the section 8.3. In order to verify these numbers we have simulated the number of known clients using SIMSCRIPT II.

SIMSCRIPT II.5 is a simulation language which is in active use for major simulations, particularly in engineering applications. The result illustrated in 9.17, shows clearly that the theoretical and experimental results are extremely closes. Figures 9.18 shows the number of waiting packets number. As in the previous tests, this figure shows that the number of waiting packets increases suddenly when the attack takes place.

9. EVALUATION




Figure 9.17: Simulation of the Number of Known clients (SIMSCRIPT)




Figure 9.18: Waiting Packets (Test 1-Poisson)

Figure 9.19 shows the average waiting time. The maximum waiting time shows a sudden increase also in the case of attack.

9.5.5 Results Test 2: Ten Forged IP Addresses

Figure 9.20 shows the traffic entering the filter and resulting from both the legitimate clients and the attacker, and figure 9.21 the traffic leaving the filter and entering the first hop multicast router. Figure 9.22 shows the number of known clients. In this test we have evaluated the number of known clients theoretically, using SIMSCRIPT II. The results are illustrated in 9.23

Figures 9.24 and 9.25) show the number of waiting packets and the maximum waiting time. Compared to the previous test, we note the impact of the increase of accepted attackers on the these values. The waiting packets number and the maximum waiting time increase significantly with the increase of the number of intruders in the system. However as shown in the figure the system returns to its stable state almost immediately after the attack.

9.5 IGMP Traffic Sent with a Poisson Distribution




Figure 9.19: Average Waiting Time (Test 1-Poisson)




Figure 9.20: Incoming traffic (Test 2-Poisson)




Figure 9.21: scheduled traffic (Test 2-Poisson)

9. EVALUATION

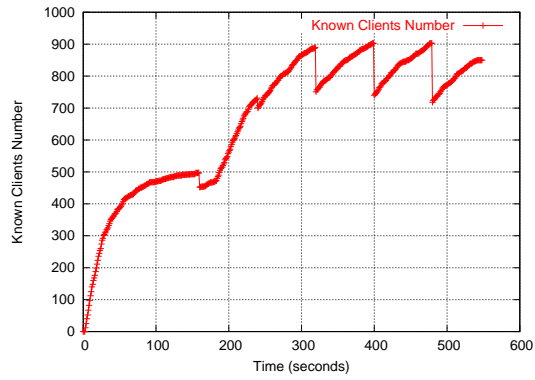


Figure 9.22: Known clients (Test 2-Poisson)




Figure 9.23: Simulation of the Known Clients Number (Test 2-Poisson)




Figure 9.24: Waiting Packets (Test 2-Poisson)

9.5 IGMP Traffic Sent with a Poisson Distribution




Figure 9.25: Average Waiting Time (Test 2-Poisson)

9.5.6 Results for Test 3: 1000 Forged IP Addresses

In this test the attacker uses a high number (1000) of forged source addresses. The purpose of this test is to study how the increase in the spoofed addresses affects the system and how the system can mitigate the results of the attack.

Figure 9.26(b) shows that the attacker's traffic is progressively accepted by the filter, once the forged source address becomes known by the system. The attacker prevents almost all new legitimate clients from being accepted during the attack. These latter must wait the end of the attack before reaching the known client status, and having their traffic go through the filter.

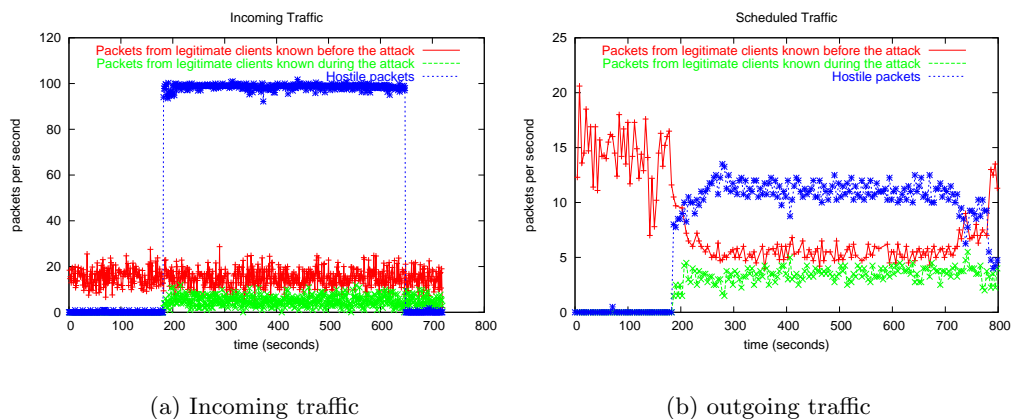


Figure 9.26: Traffic entering and leaving the filter (Test 3).

As with the third test in the case of traffic coming at a constant rate, when the number of hostile clients entering the system increases over the scheduling rate and especially if the duration of the attack is large enough, the legitimate clients could be largely penalized. This is visible in Figure 9.26. However, the DoS attack does not impact the whole multicast routing

9. EVALUATION

infrastructure of the operator. In summary, the experimental and theoretical evaluations both show that the filter can up to certain degree sustain the availability of the multicast routing infrastructure and achieve a fair service guarantee for the legitimate clients. However, other extensions are needed to improve the filtering efficiency. We will discuss some perspectives to improve the filtering efficiency in the future works.

9.6 Conclusion

In this chapter we have presented some experimental results to evaluate our filtering proposal. While more extensive tests are needed, the tests shows that the filter exhibits excellent performances when the attacker uses a small number of forged source addresses. In this case, using the filter makes the attack almost unnoticeable to the first hop multicast router. Besides new clients, coming during the attack, will still be accepted by the filter and their traffic will go through the filter normally. These results remain true during high packet rate attacks, since the small list associated to the attacker(s) quickly overflow.

While the results achieved with test 3 seem to be very negative, one would keep in mind:

- that the filter, as described up to now, has no way to distinguish the legitimate traffic from the attacker's traffic. Several extensions will be introduced in the future works on in order to improve its behavior in cases similar to test 3.
- the total outgoing traffic, leaving the filter and entering the router, is limited to a total of 20 pps, even in presence of an attack where the incoming traffic amounts to a total of 120 pps. The same result would have been achieved, no matter the attacker's *REPORT* message sending rate. It means that the filter makes attacks that try to exhaust CPU resources more complex.
- some underlying infrastructures have intrinsic features, such as using identifiers that can be used to identify the end-clients: VC/PVC, DSL-ID, Relay DHCP ID, VLAN tag. These identifiers can largely reduce or eliminate the problem of addresses spoofing as will be explained in next section.

Chapter 10

Discussion

Contents

10.1 Filter Deployment: Practical Examples	89
10.1.1 Filter Deployment in the Operator Network	89
10.1.2 Filter Deployment in a Campus Network	92
10.2 Discussion of Some Extensions to Improve the Filtering Efficiency	93
10.2.1 IP/MAC Addresses Spoofing Resiliency	93
10.3 Conclusion	94

In this chapter we first discuss the deployment of the filter in the operational networks. We present then some extensions to improve the resiliency of the filter in front of addresses spoofing.

10.1 Filter Deployment: Practical Examples

In this section, we study the deployment of the filter in the operational networks. More specifically, we consider two cases: 1) the operator network; and 2) a campus network.

10.1.1 Filter Deployment in the Operator Network

Although IGMP itself is a well known standard, the implementation and use of IGMP in a broadband access network is not governed by any best practices and thus not implemented in a single consistent manner [67].

Moreover, the operator network follows continuous evolution: ATM is widely deployed by network operators as it offers a proven and reliable way to support very large multi-service networks. However, Ethernet has emerged as viable high speed access technology and has a clear advantage over ATM.

In the following we present two broadband architectures examples, the first is based on ATM and the other on the Ethernet. We explore how to place the filter in these architectures.

10. DISCUSSION

ATM broadband access and aggregation network

In the "ATM over ADSL" model there are two possibilities for ATM connection between end-user and the operator network: ATM PVCs (Permanent Virtual Channel) and ATM SVCs (Switched Virtual Circuit). A PVC is established statically by the network administrator following a service order and cannot be altered by the user, while a SVC connection is established in real-time in response to signaling messages from the customer. SVCs would greatly reduce the effort to provision service to a new customer, and would also permit customer to freely roam among Information Service Providers. However, the SVC is not largely deployed in the operational networks. So, in the following we consider only the PVC technology. All the ATM connections will be aggregated in the DSLAMs then in the B-RAS which is a focal point for all DSLAM PVCs. There may be more than one PVC from each customer to the B-RAS. For example, separate voice and data PVCs may be used to differentiate the two types of traffic. In the case of IGMP a special PVC per use could be used for signaling purposes.

The IGMP packets are handled typically by the DSLAM which plays the role of an IGMP/MLD proxy 3.2.3. And membership in an IP multicast group is equivalent to adding the ATM PVC to a leaf for a multicast tree [6].

As a result, the filter should be integrated in the DSLMA (Figure 10.1).

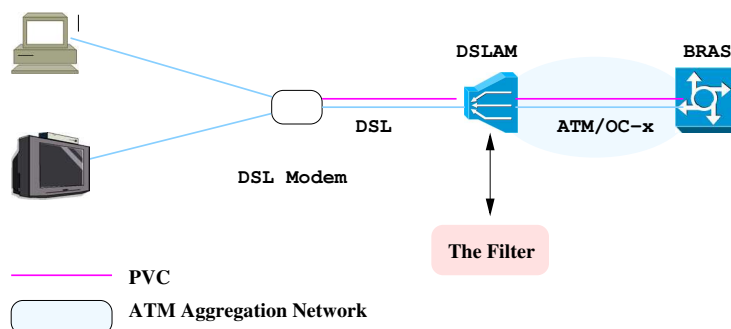


Figure 10.1: Deploying the Filter in the DSALM.

The PVCs are used in this case to identify the clients.

Note that this deployment protects efficiently the DSLAMs but not the BRAS which connects multiple DSLAMs, especially, in the case where a particular DSLAM is corrupted (software anomaly for example). Integrating the filter in the BRAS can guarantee the continuity of service in such a scenario (Figure ??).

Ethernet broadband access and aggregation network

In the Ethernet networks the VLANs (virtual LAN) are used to separate user traffic from each other. The frames belonging to each VLAN is marked with a VLAN ID number that is a 12-bit field of IEEE 802.1Q tag. VLAN-IDs are used to identify a group of users. Thus in this case the scheduling process in the filter should be done per VLAN basis rather than per end client basis.

When it is possible to use C-VLAN which is a point to point VLAN providing VC-like connectivity the scheduling can be done per client basis.

Over pure Ethernet that does not support VLAN, the end user have be identified by theirs IP/MAC addresses and when it is necessary the clients could be authenticated thank to the use




Figure 10.2: Deploying the Filter in the BRAS.

the IEEE 802.1x-2001 [4].

Over Ethernet, there is essentially only one mode of operation for multicast in a bridging DSLAM: participants are added to or removed from multicast group lists based on the detection of IGMP messages in the downstream traffic [46]. The only nuance is that these IGMP messages can either be passed on into the network with their original source address (snooping), or the IGMP messages can be proxied by the DSLAM (section 3.2.4). In the following we explore these two cases and we study the repercussions on the filter deployment.

- *IGMP Proxy*

When acting as the proxy, the DSLAM performs the host portion of the IGMP task on the upstream interface as follows:

- When queried, sends group membership reports to the group.
- When one of its hosts joins a multicast address group to which none of its other hosts belong, sends unsolicited group membership reports to that group.
- When the last of its hosts in a particular multicast group leaves the group, sends an unsolicited leave group membership report to the all-routers group (244.0.0.2).

In this case, the placement of the filter is similar to that in the ATM network. However, over Ethernet the external and the internal deployment are possible and the clients can be identified by their addresses IP/MAC, the VLAN-ID and the interface over which the IGMP message has arrived.

- *IGMP Transparent Snooping*

The IGMP snooping allows a network device (switch) to monitor IGMP queries and reports to determine if a downstream host should receive multicast frames or stop receiving multicast frames (section 3.2.4). The snooping agent will add an interface to its OIF table when a join is seen coming from an interface. In a similar manner, the snooping agent can stop sending multicast traffic when leave is discovered from an IGMP client [67].

The snooping device does not participate in the IGMP host messaging and promiscuously listens to transactions between clients and routers.

The snooping device are highly vulnerable to L2 attacks, however here we consider only the IGMP specific attacks.

10. DISCUSSION

With snooping, the filter could be placed before/in the switch and the filtering should be done per port and MAC addresses basis. In addition, the filter must be integrated in the upstream router (i.e. the B-RAS) and the filtering could be done per IP/MAC addresses and interface basis.

Some mechanisms such as the layer2 control access topology discovery allow the B-RAS to gain knowledge about the topology of the access network, the various link being used and their respective rates. Therefore, the B-RAS could identify the clients by unforgeable identifiers such as the DSL-ID even if these clients are not directly connected to the B-RAS.

10.1.2 Filter Deployment in a Campus Network

The common description of a campus network is a group of LAN segments within a building or group of buildings that connect to form one network. Typically, one company owns the entire network, including the wiring between buildings. This local area network (LAN) typically uses Ethernet, Token Ring, Fiber Distributed Data Interface (FDDI), or Asynchronous Transfer Mode (ATM) technologies.

Many campus network models exist. In this section we consider only one architecture [3] (Figure 10.3).




Figure 10.3: Example of a Campus Network Architecture.

- Core
The core layer is literally the core of the network. At the top of the hierarchy, the core layer is responsible for transporting large amounts of traffic both reliably and quickly.
- Distribution
The primary function of the distribution layer is to provide routing, filtering, and WAN access and to determine how packets can access the core. The distribution layer determines the best path, then it forwards the request to the core layer. The core layer is then responsible for quickly transporting the request to the correct service. The distribution layer is the place to implement policies for the network.
- Access

The access layer controls user and workgroup access to internetwork resources. This layers compromise the L2 switches.

10.2 Discussion of Some Extensions to Improve the Filtering Efficiency

In such an architecture the switches in the access layer snoop IGMP messages and the router in the distribution layer handle these messages and replicate multicast traffic. In this scenario the filter could be integrated in each L2 switch and the scheduling process have to be done per port and MAC addresses basis. In addition the filter have to be placed before/in the router in the distribution layer and the scheduling in this case could be done per IP/MAC addresses and interface basis.

10.2 Discussion of Some Extensions to Improve the Filtering Efficiency

10.2.1 IP/MAC Addresses Spoofing Resiliency

The experimental results presented in the previous chapter show clearly that the filter exhibits excellent performances when the attacker uses a small number of forged source addresses. In this case, using the filter makes the attack almost unnoticeable to the first hop multicast router. Though, the filter shows some limitations when the attacker uses a large number of spoofed addresses. Preventing address spoofing, if possible, is therefore highly beneficial.

Using unforgeable identifiers

When applicable, using unforgeable client identifiers solves totally spoofing problems and makes the filtering component highly efficient in front of attacks. In the operator networks these identifier could be the: VC/PVC, DSL-ID, and VLAN tag. Using these identifiers depends on the deployed access and core infrastructures (section 2.1.1), as well as, the location of the filter as has been discussed previous section.

The use of these identifiers could be achieved practically via the DHCP relay agent information option [59]. Actually, this option enables a Dynamic Host Configuration Protocol (DHCP) relay agent (located generally in the DSLAM) to include information about itself when forwarding client-originated DHCP packets to a DHCP server. The DHCP server can use this information to implement IP address or other parameter-assignment policies [59]. This feature communicates information to the DHCP server using a sub-option of the DHCP relay agent information option called agent remote ID. The information sent in the agent remote ID includes an IP address identifying the relay agent and information about the ATM interface and the PVC over which the DHCP request came in, and the DSL ID. In general, the DHCP server may be extended to maintain a database with the "triplet" of (client IP address, client MAC address, client remote ID).

To sum up, the use of DHCP relay agent information has many improvements on the security:

- IP address spoofing The DSALM (DHCP relay agent) may associate the IP address assigned by a DHCP server in a forwarded DHCP Ack packet with the circuit to which it was forwarded. Then the DSALM prevent forwarding of IP packets with source IP addresses -other than- those it has associated with the receiving circuit. This prevents simple IP spoofing attacks on the Central LAN, and IP spoofing of other hosts.
- MAC address spoofing

By associating a MAC address with an Agent Remote ID, the DHCP server can prevent offering an IP address to an attacker spoofing the same MAC address on a different remote ID.

10. DISCUSSION

Thus, we can imagine that the filter can interact with the DHCP server to obtain these identifiers.

Other Situations

When unforgeable identifiers cannot be used, or when techniques such as DHCP relay agent, cannot be used the situation becomes more complex. A possible sanity check consists in verifying that the source address contained in the IGMP/MLD messages are valid, i.e. correspond to existing or possible clients. This sanity check has of course limitations, but it can at least protect against attacks where the attackers chooses random addresses rather than targeted source addresses.

Going further would require to use packet authentication mechanisms (e.g. by adding a digital signature to each IGMP/MLD packet sent). Since it contradicts one of our main requirements, the need to keep the existing multicast routing infrastructure and protocols unchanged, we do not further elaborate this solution.

In such situations the use of a learning technology seems to be an attractive solution. We will present in the future work a learning proposition that allows to avoid spoofing and to make the filter robust against other IGMP attacks in the same time.

10.3 Conclusion

In this chapter we have discussed the deployment of the filter in the operational networks and we have presented some extensions to improve its resiliency towards addresses spoofing. As has been shown the operator network has many intrinsic properties such as the use of unforgeable identifiers that allows to improve largely the robustness of the filter in front of address spoofing. When the use of unforgeable identifiers is not possible the use of a learning technology seems to be an attractive solution. We will present in the future work such a learning mechanism.

Chapter 11

Conclusion and Future Works

Contents

11.1	Reminder of the Objective of the Thesis	95
11.2	Our Proposal: A Pragmatic and Deployable Filter-Based Solution	96
11.3	Future Works	97
11.3.1	Membership Policy	97
11.3.2	Signature Based Anomaly Detection	97
11.3.3	Using Learning Mechanisms	98
11.4	Final Words	98

In this chapter, we conclude the dissertation by revisiting the lessons learned from this work and presenting directions for future researches in this area.

11.1 Reminder of the Objective of the Thesis

Multicast is a promising technology for the distribution of streaming media, bulk data and many other added-value applications. Yet the deployment of multicast still in its infancy. This work considers one of the most challenging features of multicast: the security. More specifically this thesis focuses on *the security of the multicast routing infrastructure Security from the Network Operator Point of View*.

The kind of security required by a network operator, who manages and operates the multicast routing infrastructure, largely differs from that of end-to-end security. The Network Operator's security point of view for its group communication service can be summarized as follows:

The group communication service provided to its clients (i.e. end users or other network operators with whom has peering relationships) must be operational at any time, in spite of anomalies in the multicast flows, no matter whether these anomalies are intentional (i.e. are the result of deliberate attacks) or not (e.g. are caused by a misbehaving component). *Security is not the goal, but a mean to achieve the network operator's "continuation of service no matter what happens" goal.* Security should not impact too much the unicast and multicast forwarding performances on the operator's network.

11. CONCLUSION AND FUTURE WORKS

The Network Operator can have additional requirements. For instance, he may want to guarantee that the traffic exchanged is not altered while traveling on its own routing infrastructure. He may also want to ensure some confidentiality of the traffic, in case an attacker eavesdrops on a link or a subverted router. These additional security considerations are more commonly addressed by end-to-end security, and they are not considered in the present work.

11.2 Our Proposal: A Pragmatic and Deployable Filter-Based Solution

While many theoretically ideal proposals have been introduced to secure the routing protocols, they have rarely been accepted by the operators community. Most of them require to modify some existing and widely deployed protocols, or introduce authentication mechanisms, which is in practice almost impossible to deploy in legacy networks (and even useless since a corrupted host may be the source of a DoS attacks, even if it has been authenticated).

In our work, *we have deliberately avoided any kind of solution that would upset the existing infrastructure already deployed at the network operator or in the user premises.* Of course, this constraint has closed many doors, but we are convinced that this choice largely increases the acceptance of our proposal within the operator community.

In this thesis we have analyzed in depth the threats to the multicast infrastructure and identified the requirements for the network operator. This analysis has shown that the vulnerability of the multicast model comes largely from the edge. More specifically, several attacks arise from the use of group management protocols, IGMP for IPv4 and MLD for IPv6.

In order to cope with these attacks, directed to the group management protocol, we proposed a *realistic and pragmatic solution*, based on a filtering component, located between the first hop multicast router and the set of clients. This filter creates two classes of clients: known clients, who have been already accepted by the system, and new clients. Known clients are then served equally by the filter that forwards their IGMP/MLD messages to the first hop multicast router.

Our solution has the following features:

1. *It mitigates IGMP/MLD flooding attacks through a cost effective, scalable and transparent mechanism.*
2. *It includes an automatic learning mechanism, by which the filter learns who are the regular clients.* Because of this auto learning feature, the solution is *easily manageable*. Once a few key parameters have been initialized, taking into account the target environment, the filter will self-initialize.
3. *It protects the well-behaved clients against the ill-behaved ones.* This is a result of both the learning mechanism that keeps a state for each known client, and of the various small waiting queue that will easily be overflowed in case of an attack. Thanks to these two features, well-behaved clients, already known by the system, will be served with a higher priority during an attack, at least up to a certain point.
4. *It Keeps the changes to the currently used multicast routing infrastructure as minimum as possible, and it avoids non realistic assumptions and techniques.* This makes this solution easily deployable.
5. *It is easily extensible.* Several peripheral components can be added to the filtering component, taking into account the specificity of the target environment.

11.3 Future Works

11.3.1 Membership Policy

- Group management policy

Adding a policy that associates each client with the groups for which he is allowed subscribe if this information is available, or the maximum rate at which he can subscribe to new groups.

- Congestion control policy Some works, like [33], have studied the attacks on the congestion control protocols. However, few works consider the topology of the network operator. To the best of our knowledge, only [58] proposes a solution to handle this problem. In [58], the authors propose an IP multicasting rate control system using layered content, called Multi-Rate Filtering (MRF), to optimize the streaming rate by controlling the number of content layers for each VLAN when a shared link of an edge node is congested. MRF can allocate the bandwidth of the shared line to VLANs according to their utilization and then transmit a multicast content to clients with different stream rates reflecting the users' intentions. This proposal takes into consideration the topology of the access network, however it needs to modify the existing infrastructure which contradicts our goals.

We think that adding a congestion control policy to the filter could be helpful to protect against receivers based congestion control attacks. This policy determines if a receiver can access a particular multicast group upon its congestion state. To obtain the congestion state of each line DSL or of each VLAN the filter some mechanisms such as the layer2 control access topology discovery witch allow the B-RAS to gain knowledge about the topology of the access network, the various link being used and their respective rates could be used.

11.3.2 Signature Based Anomaly Detection

Signature-based anomaly detection requires to understand the attack in order to identify them by their state transition sequences or patterns. Testing the conformity of the arriving IGMP messages with the IGMP state transition sequences can help to counter many attacks:

Forged IGMPv3/MLDv2 State-Change Report

A forged *State-Change Report* message will cause the *Querier* to send out *Group-Specific* or *Source-and-Group-Specific Queries* for the group in question. This causes extra processing on each router and on each member of the group. To make the filter somehow robust toward this attack, the filter can store the history of each known client and drop the IGMPv3/MLDv2 *State-Change Report* messages if the client does not have a state for this group.

Report and Leave Messages Storm Attack

A forged *Leave* message will cause the *Querier* to send *Group-Specific Queries* for the group in question and will cause the receivers to send reports. This causes extra processing on the first hop multicast router and on each member of the group.

An amplified version of the attack is also possible: for each *Report* heard in the network, the attacker sends a *Leave* message causing the resources of the *Querier* to be exhausted. Of

11. CONCLUSION AND FUTURE WORKS

course, this attack requires that the attacker and the other clients share the same medium (e.g. a LAN). Using signature based detection can counter these attacks. Some signature are:

- *If any receiver sends a Leave for a group for which he has never been a member*
- *If any receiver sends a Leave for a group for which there are other members*

11.3.3 Using Learning Mechanisms

In this section, we propose to extend the filter by a learning mechanism. This mechanism aims to improve the resiliency of the filter in front of addresses spoofing when the use of unforgeable identifiers is not possible.

The proposed learning mechanism takes four parameters:

1. The interval statistics of the arriving packets: Accepting the packets in the system or dropping them depends on the intensity of the arrival traffic. When the packets arrive with intensity larger than a predefined threshold they will be dropped directly, in contrast, when the intensity of arriving traffic is less than this threshold they will be accepted in the system.
2. The IGMP finite state machine conformity, as has been explained in the previous section.
3. The IP/MAC addresses familiarity. When the use unforgeable identifiers is not possible we propose to use a learning mechanisms that takes decisions on the arrived packets upon the familiarity of their source addresses. That is, we define first a set of acceptable IP/MAC address then the arrived packet will be accepted or dropped in function of the distance between its source address and the predefined set of addresses.

These factors are used the to take many decisions on the arriving packets:

- Dropping the arriving packets immediately.
- Accepting the new client.
- Classifying the accepted clients into clusters of priorities. The clustering could be performed thanks to well known classifying algorithms such as the K-mean algorithm [44].
- Degrading or upgrading of known clients between the different clusters.

We think that such a mechanism can improve the filtering efficiency and protect the well behaved client against ill behaved ones by giving them higher priorities.

11.4 Final Words

The security of the multicast routing infrastructure is a large problem with many open and challenging issues. Our filtering module presents a pragmatic and realistic solution for most of these issues with reasonable implementation costs.

Our solution, in its simplest form, introduces an efficient architecture that complies with the thesis goals and can, in particular, be integrated into any network operator infrastructure. If it cannot prevent all kinds of edge attacks, it is easily extensible by adding more intelligence features.

Résumé Français

Chapter 12

Infrastructure Sécurisée de Routage Multipoint: Le Point de Vue de l'Opérateur Réseau

12.1 Introduction

Le multicast (ou les communication point a multipoints ou multipoints a multipoints) est un modèle de diffusion efficace des données permettant à une source d'émettre une seule copie de trafic à destination de plusieurs récepteurs dispersés sur l'Internet. Le modèle de service initial, tel qu'il a été introduit par Deering [24], est le modèle ASM (Any-Source Multicast). Ce modèle permet à n'importe quel terminal de s'abonner à un groupe, d'envoyer et de recevoir des flux multicast. Récemment, un modèle simplifié a été proposé au sein de l'IETF (Internet Engineering Task Force), le modèle SSM (Source Specific Multicast) spécialement conçu pour supporter la diffusion de type $\langle 1 \text{ vers } n \rangle$. Ce modèle introduit la notion de canal qui est l'association d'une adresse de groupe G et de l'adresse S de la source du trafic multicast.

Afin de s'abonner de déclarer leur appartenance à un groupe donnée les récepteurs utilise les protocoles de gestion de group protocole IGMPv1,v2,v3 (Internet Group Management Protocol) [24] [29] [16] en IPv4 et MLDv1,v2 (Multicast Listener Discovery) [25] [72] en IPv6.

Pour acheminer un paquet multicast, un arbre multicast doit être construit entre la source et les destinataires (membres du même groupe). Une seule copie du paquet est acheminée sur chaque branche de l'arbre multicast. Les protocoles chargés de construire l'arbre multicast et d'acheminer les paquets le long de cet arbre sont appelés les protocoles de routage multicast. Ces protocoles de routage multicast peuvent être classés en deux catégories, selon qu'ils sont propres à un "domaine multicast" (c'est le cas de MOSPF, CBT, PIM-SM [27] et PIM-DM) ou qu'ils sont inter-domaines (comme MBGP/PIM-SM/MSDP [28] et BGMP [70]).

Bien que le modèle de diffusion multicast ait été défini depuis plusieurs années et les implantations de la famille de protocole multicast soient disponibles dans presque tous les équipements réseau, à ce jour cette technologie n'est pourtant pas mise en oeuvre et déployée à grande échelle dans les réseaux des opérateurs. Plusieurs raisons permettent d'expliquer cet état de fait, tels que le problème de l'allocation dynamique des adresses multicast, la fiabilité de transmission

12. INFRASTRUCTURE SÉCURISÉE DE ROUTAGE MULTIPOINT: LE POINT DE VUE DE L'OPÉRATEUR RÉSEAU

ou la sécurité. Cet article n'aborde que la problématique de sécurité. En pratique, lorsque l'on parle de la sécurité appliquée au domaine du multicast, deux niveaux complémentaires doivent être considérés: la sécurité applicative d'une part, qui vise à la protection des données, et la sécurité de l'infrastructure de routage IP multicast de l'opérateur d'autre part. La première problématique a été largement étudiée [55] [60] [77] [17] et plusieurs travaux ont été réalisés au sein du groupe de travail MSEC [2] à l'IETF pour garantir une sécurité des données multicast transportées. Malgré l'efficacité de ces travaux, il est important de souligner que ces flux sécurisés sont transportés par l'infrastructure réseau multicast, qui elle n'est pas sécurisée et donc vulnérable aux attaques, en particulier aux attaques par déni de service (ou DoS). Le principe de ces attaques est de submerger le réseau de trafic pollueur ou de requêtes, afin de saturer la bande passante des liens ou les capacités de traitement des équipements (routeurs, serveurs), et rendre le réseau indisponible. Tous les services, y compris les services "non multicast", sont alors impactés. Les attaques DoS sont nombreuses et variées, certaines exploitent la vulnérabilité du protocole IGMP de souscription à un groupe de diffusion, d'autres ciblent les protocoles de routage multicast intra-domaine comme PIM-SM ou inter-domaine tels MSDP et BGMP. Enfin, d'autres types d'attaques, délibérées ou non, peuvent provenir d'un comportement non conforme du contrôle de congestion de l'application. On voit que le niveau de sécurité exigé de la part de l'opérateur de réseau, qui gère et exploite l'infrastructure de routage multicast, diffère largement de la sécurité de bout en bout de niveau application. C'est l'objet de ce papier que de se focaliser sur la sécurité de l'infrastructure réseau et la problématique de l'opérateur réseau. Ce résumé est organisé comme suit : la section 2 identifie et analyse les différents acteurs impliqués dans un service de diffusion multicast d'un contenu. Cette analyse est illustrée par deux études de cas complémentaires: un service de diffusion vidéo sur ADSL, et un service de diffusion libre. À la lumière de cette discussion, la section 3 analyse les spécificités et exigences de sécurité de l'opérateur réseau. La section 4 classe les attaques possibles sur l'infrastructure, tandis que la section 5 présente les grandes familles de solutions possibles pour sécuriser cette infrastructure. La section 6 offre une discussion critique des problèmes et des solutions proposées, qui reposent souvent sur des hypothèses peu réalistes dans notre contexte. La section 7 présente notre proposition. Les sections 8 et 9 discutent les scénarios de déploiement dans les réseaux opérationnels. On étudie dans la section 10 comment paramétrer le filtre en pratique. La section 11 présente les résultats expérimentaux. Tandis que la section 12 présente des extensions qui peuvent permettre d'améliorer le filtrage. Finalement, la section 12 conclut ce résumé.

12.2 Les Différents Acteurs dans le Cadre d'un Service de Diffusion Multicast

12.2.1 Vue générale

Comprendre la problématique de sécurité de l'infrastructure réseau nécessite de comprendre la façon dont l'architecture de service elle-même est organisée, d'identifier les différents acteurs qui collaborent à la fourniture du service de diffusion multicast, et de comprendre leurs relations, contractuelles ou non. Dans cet article nous distinguons (Figure 12.1) :

- *L'opérateur réseau:*

C'est celui qui possède, déploie et gère l'infrastructure physique (c'est à dire les routeurs multicast, les NAS/BRAS (Network Access Server/Broadband Access Server), les DSLAMs (Digital Subscriber Line Multiplexer), etc.). L'opérateur met en oeuvre le

12.2 Les Différents Acteurs dans le Cadre d'un Service de Diffusion Multicast

routage multicast et met en oeuvre des relations de peering avec les autres opérateurs (pour échanger les paires de source/groupe avec MSDP). L'opérateur fournit physiquement l'accès multicast IP à l'utilisateur final mais il est important de souligner qu'il n'a pas de relations commerciales avec celui-ci.

- *Le fournisseur d'accès Internet (ou ISP) et le fournisseur de service:*

L'ISP a la charge des relations commerciales, relatives à l'accès réseau, avec l'utilisateur final. Il fournit l'accès, la connectivité Internet et les services associés (courrier, portail web, messagerie instantanée). Le fournisseur de service quant à lui est chargé de la définition et de la commercialisation de l'offre de service (diffusion de TV numérique, ou VoD) à l'utilisateur final, et possède aussi à ce titre une relation avec l'utilisateur final. Ces deux rôles peuvent être joués par un seul et même acteur. Dans ce papier nous utiliserons le terme générique de "Fournisseur de Service" pour désigner l'un ou l'autre de ces fournisseurs.

- *Le fournisseur de contenu:*

Cette entité fournit les informations et le contenu qui composent le service proposé aux utilisateurs finaux par le fournisseur de service. Ce contenu peut être du texte, des images, de la vidéo (TV ou VoD numérique), ou de l'audio (radio, musique).

- *L'agrégateur de contenu:*

Cette entité construit les bouquets de chaînes à partir du contenu fourni par un ou plusieurs fournisseurs de contenu. Il collecte le contenu, applique le codage approprié sur ce contenu et le diffuse à partir de son réseau. L'agrégateur de contenu a des relations de peering et des contrats avec un ou plusieurs opérateurs de réseau afin d'atteindre les utilisateurs finaux.

- *L'utilisateur final:*

L'utilisateur possède un terminal multimédia qui peut être un PC ou un ensemble TV/Set Top Box. En premier lieu, l'utilisateur obtient un accès Internet de la part de son ISP, ensuite il peut utiliser le service de diffusion multicast proposé par l'opérateur réseau afin de s'abonner aux chaînes/groupes qu'il désire, et ainsi recevoir le contenu associé. Notons, qu'un même acteur peut très bien jouer le rôle de plusieurs entités.

12.2.2 Premier exemple : service commercial de TV sur ADSL

Cette section décrit un service commercial de multicast de TV sur ADSL (Asymmetric Digital Subscriber Line). En pratique, dans ce type de service le Fournisseur de Service offre l'accès Internet haut débit au client et gère les aspects d'identification, authentification, autorisation, allocation d'adresse IP et facturation. Ce premier niveau d'authentification/autorisation permet au client d'obtenir un accès Internet à haut débit. Il est probable que l'accès à certains contenus (telles les chaînes TV hertziennes) restera gratuit avec l'ADSL. En revanche l'accès à des chaînes à péage ou à des services comme la vidéo à la demande (VoD) sont des services payants soumis à un contrôle d'accès auprès d'un serveur d'authentification comme le serveur RADIUS (Remote Access Dial In User Service). Afin de recevoir les flux vidéo, le client utilise IGMP (avec IPv4) ou MLD (avec IPv6). Les messages "IGMP reports" sont traités par les équipements multicast du réseau de l'opérateur réseau, à savoir le DSLAM ou le BRAS (en fonction de la maturité des fonctions multicast de ces équipements ou des choix d'ingénierie opérationnelle). Les flux

12. INFRASTRUCTURE SÉCURISÉE DE ROUTAGE MULTIPOINT: LE POINT DE VUE DE L'OPÉRATEUR RÉSEAU




Figure 12.1: Les Différents Acteurs dans le Cadre d'un Service de Diffusion Multicast.

multicast sont donc transmis à partir du réseau de l'agrégateur de contenu, puis transitent sur le réseau de l'opérateur réseau avant d'atteindre le client final. Les flux de gestion, tels que les flux AAA (Authorization/ Authentication/ Accounting) qui permettent d'authentifier et de facturer l'utilisateur, sont échangés et relayés entre l'opérateur de réseau et le fournisseur de service. Ce sont les seuls flux qui existent entre le client et le fournisseur de service.

12.2.3 Deuxième exemple: diffusion libre d'un contenu gratuit

Dans le cas de la diffusion libre d'un contenu (logiciels gratuits, vidéo clips, visio-conférence entre amis, etc.), la chaîne de bout en bout (Fournisseur de contenu, utilisateurs finaux) est toujours valide. En revanche le client n'a pas à s'authentifier pour recevoir le contenu. Dans ce schéma de service, il n'y a qu'un seul niveau d'authentification, lors de l'obtention de la connectivité Internet.

12.3 Sécurité de l'Infrastructure du Point de Vue de l'Opérateur Réseau

L'opérateur réseau a un point de vue très spécifique sur la sécurité de son réseau : Le service de communication de groupe fourni à ses clients (utilisateurs finaux, opérateurs réseaux avec lesquels il a des relations de peering, agrégateur de contenu) doit être opérationnel en permanence, en dépit des anomalies pouvant survenir, que celles-ci soient intentionnelles (attaques) ou non (dysfonctionnement d'un composant quelconque, chez l'utilisateur ou l'opérateur). La sécurité n'est pas un but en soi, mais un moyen d'atteindre l'objectif de « continuation de service envers et contre tout ». D'autre part, les mesures de sécurité mises en oeuvre ne doivent pas influencer négativement les performances des services unicast et multicast fournis par l'opérateur. L'opérateur réseau peut avoir des exigences supplémentaires en matière de sécurité. Il peut vouloir garantir que le trafic échangé ne soit pas modifié lorsqu'il transite sur son propre réseau. Il peut aussi vouloir garantir un certain niveau de confidentialité au trafic, au cas où un attaquant puisse écouter le trafic à partir d'un lien ou d'un routeur corrompu. Ces considérations de sécurité supplémentaires ne sont cependant pas l'objectif premier de l'opérateur réseau, et sont mieux adressées par la sécurité de bout en bout. Par conséquent nous ne les considérons pas dans cet article. De même, les pannes physiques affectant des liens,

12.4 Taxonomie des Attaques Intéressant l'Opérateur Réseau

des équipements, routeurs ou serveurs, ne sont pas considérées dans cet article, bien qu'elles aient un impact fort par rapport à l'objectif de «continuité de service» de l'opérateur. Notons que notre définition partage quelques similitudes avec celle de « survivabilité du réseau » [79], même si cette dernière est un peu plus large puisqu'elle inclut les pannes matérielles.

12.4 Taxonomie des Attaques Intéressant l'Opérateur Réseau

Le modèle de diffusion multicast IP est par définition un modèle ouvert, où n'existe aucun mécanisme de contrôle à l'accès du réseau. Or, plusieurs attaques récentes, tel le ver Ramen [61], ont clairement montré la vulnérabilité de l'infrastructure multicast aux attaques, en particulier les attaques DoS. Dans cette section nous présentons une taxonomie des attaques dirigées contre l'infrastructure de l'opérateur :

- *Selon leur origine:*

Ces attaques peuvent provenir soit du coeur de réseau soit de l'accès. Le réseau de coeur contient l'ensemble des routeurs multicast qui exécutent les protocoles de routage multicast. En revanche, le réseau d'accès est schématiquement constitué des équipements de concentration de trafic qui opèrent IGMP/MLD, et des utilisateurs finaux.

- *Selon leur type:*

On peut distinguer les attaques dirigées dans le plan de transfert (échange de données) des attaques dirigées dans le plan de contrôle (visant les protocoles). A chaque fois le terme "attaque" est à prendre au sens large et dénote soit une agression intentionnelle, soit non intentionnelle.

12.4.1 Attaques internes

Une attaque interne provient d'un élément de l'arbre de distribution multicast de l'opérateur réseau : routeur corrompu ou lien réseau piraté. Ces attaques exploitent les vulnérabilités des équipements physiques au sein du réseau, comme les routeurs et les serveurs. En prenant le contrôle d'un équipement du réseau, un intrus peut alors déclencher une variété d'attaques.

Attaques internes dirigées dans le plan de transfert

Les attaques dans le plan de transfert sont nombreuses : un intrus peut modifier le contenu des paquets, ou submerger le réseau par du trafic parasite qui va être reçu par tous les récepteurs; il peut aussi copier le contenu d'un groupe, interpréter les informations et les rejouer plus tard. Ces attaques, qui pénalisent essentiellement les clients et fournisseurs de contenu, ont aussi des conséquences directes pour l'opérateur puisqu'elles génèrent des flux parasites qui gaspillent de la bande passante et rendent le réseau lent ou indisponible. Des attaques passives, où l'attaquant profite de son contrôle sur un routeur corrompu pour espionner le trafic, ne mettent pas en péril le fonctionnement du réseau de l'opérateur mais affectent la crédibilité de l'opérateur auprès de ses clients fournisseurs de service.

12. INFRASTRUCTURE SÉCURISÉE DE ROUTAGE MULTIPOINT: LE POINT DE VUE DE L'OPÉRATEUR RÉSEAU

Attaques internes dans le plan de contrôle

Les protocoles de routage en général, et ceux de routage multicast en particulier, sont très vulnérables aux attaques face à un intrus stratégiquement placé, car en pratique les messages de contrôle ne sont pas authentifiés dans les réseaux opérationnels. Cet intrus peut alors créer, rejouer, espionner, ou supprimer des messages de routage, ce qui entraîne rapidement des dénis de service pour les utilisateurs finaux. Le protocole PIM-SM par exemple, est vulnérable aux attaques basées sur des messages de contrôle falsifiés [27] [36]. Ainsi un faux message (Join/Prune) peut diriger les flux multicast vers des récepteurs illégitimes. Un intrus ayant pris le contrôle d'un routeur peut, en injectant de faux paquets de contrôle, modifier les tables de routage multicast afin de submerger de trafic un routeur particulier. L'attaque est d'autant plus grave que le routeur victime est stratégique (ce peut être le RP).

12.4.2 Attaques venant de la périphérie

Aucun mécanisme de contrôle d'accès n'étant assuré par le modèle traditionnel ASM (et guère plus avec le modèle SSM), un intrus peut facilement exploiter cette faiblesse pour déclencher des attaques très graves pour le réseau de l'opérateur. C'est ce que nous abordons dans cette section.

Attaques périphériques dans le plan de transfert

On doit distinguer les attaques venant des sources de celles venant des récepteurs.

- *Attaques issues des sources*

Le modèle ASM est très vulnérable aux attaques DoS (et DDoS). Avec ce modèle ouvert, un intrus peut inonder le réseau de l'opérateur par des flux multicast. Ces attaques qui pénalisent essentiellement les clients finaux et fournisseurs de contenu ont aussi des conséquences directes pour l'opérateur réseau puisqu'elles génèrent des flux parasites qui gaspillent de la bande passante et rendent le réseau lent ou indisponible. Et une attaque de ce type, qui commence dans le plan de transfert, peut facilement affecter le plan de contrôle. Ainsi générer du trafic destiné à un grand nombre de groupes distincts, existants ou non, peut épuiser les ressources du RP, et si le routage inter-domaines est actif, l'infrastructure est alors submergée de messages MSDP de type « Source Active ». C'est ce qui s'est passé avec le vers "Ramen". Le modèle SSM est plus robuste face à ces attaques. Pourtant il reste vulnérable face un intrus qui se fait passer pour une source légitime en lui piratant son adresse IP. Comme le routage multicast basé sur l'algorithme RPF (Reverse Path Forwarding) offre une certaine protection contre des sources qui modifient leur adresse IP source (les paquets arrivant sur une interface différente de celle qui serait utilisée pour atteindre la source sont rejetés), pour réussir son attaque l'intrus doit être sur le plus court chemin vers la source authentique. Cela rend l'attaque plus compliquée à conduire mais ne l'empêche pas totalement.

- *Attaques issues des récepteurs*

Les modèles ASM et SSM n'ont aucun mécanisme de contrôle d'accès coté récepteurs. Un intrus peut facilement utiliser IGMP/MLD pour joindre un grand nombre de flux multicast existant, ce qui consomme de la bande passante de réseau de l'opérateur .

Une autre possibilité vient des protocoles de contrôle de congestion des applications multicast. Actuellement plusieurs protocoles sont en cours de standardisation à l'IETF: WEBRC [52] pour les protocoles de multicast fiable multi-débit (tel ALC), PGMCC [63]

12.5 Taxonomie des Mécanismes de Protection de l'Infrastructure de l'Opérateur Réseau

et TFMCC pour les protocoles de multicast fiables du type PGM ou NORM. En cas de défaillance (volontaire ou non), un récepteur peut demander plus de trafic que raisonnable, en évitant de réagir aux indications de congestion, perturbant ainsi les flux réactifs tels TCP. Ceci est d'autant plus inquiétant qu'un récepteur est incité à se conduire ainsi puisque c'est une façon d'obtenir plus que sa part équitable de la bande passante.

Attaques périphériques dans le plan de contrôle

Ces attaques exploitent les vulnérabilités des protocoles IGMP et MLD. Elle peuvent consister à générer un grand nombre de messages IGMP de type REPORT; ou bien générer des messages IGMPv1 (ancienne version) et ainsi forcer les routeurs à basculer dans le mode de compatibilité, bien moins efficace; ou encore générer une succession de messages REPORT suivis de messages LEAVE, ce qui a pour effet de forcer les équipements à passer par une nouvelle phase de scrutation afin de déterminer s'il reste ou non des récepteurs pour ce groupe (forcer le basculement en mode IGMPv1 amplifie encore cette attaque). Bien entendu ces attaques peuvent se faire de façon distribuée, et l'adresse source de ces messages sera bien souvent usurpée.

12.5 Taxonomie des Mécanismes de Protection de l'Infrastructure de l'Opérateur Réseau

Nous nous intéressons maintenant aux techniques de défense. Nous pouvons classer ces mécanismes en trois catégories : les mécanismes préventifs qui visent à renforcer la sécurité en amont, les mécanismes réactifs qui prennent des mesures correctives afin d'assurer une continuité de service, les des mécanismes hybrides.

12.5.1 Mécanismes de Défense Préventifs

Les mécanismes de défense préventifs visent à ne permettre qu'aux seules entités (i.e. récepteurs, sources, et routeurs) autorisées de construire les branches des arbres de diffusion multicast et ainsi de participer à la diffusion multicast.

Les propositions qui ont été faites pour contrôler l'accès des récepteurs vérifient généralement l'identité d'un récepteur avant de lui permettre de revoir des flux multicast [13], [37], [47], [31], et [19]. Or, il y a des autres critères qui peuvent être utilisés pour autoriser un récepteur à participer à une session multicast (i.e. l'état de congestion de récepteur) [33].

Dans cette catégorie on trouve aussi des mécanismes qui visent à contrôler l'accès de sources [76] et les routeurs [65].

12.5.2 Mécanismes Réactifs

L'objectif de ces mécanismes est de tolérer les attaques qui n'auront pu être parées par les mécanismes préventifs, et assurer ainsi une continuité de service. Ces mécanismes réactifs sont en général une combinaison de techniques de reconnaissance et de détection de trafic malicieux (IDS) et de techniques de protection/restauration qui permettent au réseau de «survivre» à l'attaque.

- *Systèmes de détection d'intrusion (IDS)*

La détection d'intrusion a pour objectif de détecter toute violation de la politique de sécurité d'un système informatique. Elle permet ainsi de détecter les attaques (en temps

12. INFRASTRUCTURE SÉCURISÉE DE ROUTAGE MULTIPOINT: LE POINT DE VUE DE L'OPÉRATEUR RÉSEAU

réel ou en différé) portant atteinte à la sécurité de ce système. Plusieurs solutions ont été proposées afin de sécuriser l'infrastructure de routage unicast [21] [75]. Mais aucune d'elle n'a été faite dans le contexte du multicast qui soulève beaucoup de défis, en particulier pour collecter de façon automatisée les données représentant l'activité des systèmes de routage.

- *Mécanismes de continuité de service*

Plusieurs techniques permettent d'assurer la continuité de service et préserver le réseau dans le cas où l'intrus réussit à violer toutes les mesures précédentes :

- *Limitation de débit:*

Le principe consiste à limiter le débit des flux considérés suspects par les mécanismes IDS, ou à limiter les débits pour toutes les sources d'une façon équitable [5][30]. Bien que largement utilisée, cette solution a des limites. Ainsi elle permet des attaques distribuées, même si chacun des flux "malicieux" est individuellement limité en débit. De plus le paramétrage de la limitation en débit n'est pas chose aisée et nécessite de définir des seuils appropriés, statiques ou dynamiques.

- *Filtrage:*

On se sert des signaux reçus de la part d'un système IDS, pour bloquer les flux considérés non conformes [62]. Ces mécanismes sont largement utilisés, mais ont à leur tour des limitations puisque d'une part il y a un risque de filtrer des flux légitimes, et d'autre part un attaquant peut utiliser ces mécanismes de filtrage comme un outil pour déclencher une attaque DoS. Mécanismes d'isolement des attaques : Ces techniques réagissent aux attaques en modifiant la topologie réseau, soit en ajoutant des ressources supplémentaires, soit en isolant les parties du réseau victimes de l'attaque et en basculant le service vers les parties «saines».

12.5.3 Mécanismes Hybrides

Ces mécanismes utilisent une combinaison de control d'accès d'une coté et le filtrage d'autre coté. Plus précisément on se sert des politiques de control d'accès pour bloquer les flux considérés non-conformes [51] [18].

12.6 Discussion sur les Attaques et les Techniques de Défense

Nous avons jusqu'ici identifié les attaques possibles et les divers mécanismes de sécurité. Cette section apporte une discussion critique complémentaire, tout d'abord en classant les attaques selon leur dangerosité et probabilité, ensuite en soulignant les limites de certaines approches de sécurité qui reposent sur des hypothèses qui s'appliquent difficilement au contexte de l'opérateur réseau.

12.6.1 Classification des attaques selon leur impact sur le réseau de l'opérateur

Les attaques discutées en section 4 peuvent être classées en fonction de leur dangerosité relative pour le réseau de l'opérateur (ordre décroissant) :

12.6 Discussion sur les Attaques et les Techniques de Défense

- *Attaques venant de la périphérie: (en excluant les attaques de type contrôle de congestion)*
Ces attaques, qu'elles visent le plan de contrôle ou de transfert, sont faciles à lancer mais très difficiles à éviter (ainsi générer un grand nombre des requêtes IGMP peut créer une attaque DoS même si IGMP est sécurisé). Les attaques survenues sur MSDP ont également montré la fragilité de l'infrastructure multicast vis-à-vis des attaques DoS venues de la périphérie.
- *Attaques sur le protocole de contrôle de congestion:* Mettre en oeuvre des mécanismes de contrôle de congestion est indispensable, mais ces mécanismes, implantés dans les applications, sont également facilement modifiables. De plus des utilisateurs seront incités à rendre leur flux non (ou moins) réactifs aux indications de congestion puis qu'ils bénéficieront d'un meilleur service au détriment des autres flux type TCP. Ceci a un impact direct sur le réseau de l'opérateur.
- *Attaques internes:* Lancer ces attaques nécessite en général de disposer d'un emplacement stratégique au sein du réseau de l'opérateur, ce qui est peu probable.

En pratique, un opérateur réseau doit se focaliser essentiellement sur les attaques périphériques, très faciles à lancer et qui peuvent avoir des conséquences sérieuses sur son réseau. Se prémunir contre les attaques sur les routeurs internes et l'infrastructure physique n'est pas une priorité et ne devrait être entrepris que dans une deuxième étape.

12.6.2 Discussion sur les Techniques de Défense

Nous examinons maintenant d'un oeil critique certaines des hypothèses faites par ceux qui proposent des mécanismes de sécurité.

- *Authentication et autorisation des participants à une session multicast:*

De nombreux travaux liés à la sécurité de l'infrastructure multicast partent du principe que les participants à une session multicast sont authentifiés et autorisés. En fait, ces hypothèses sont rarement valides ou effectives : L'authentification/autorisation suppose que le client soit enregistré dans un serveur d'authentification (par exemple de type RADIUS). C'est vrai dans le cadre d'une offre de service commerciale (TV sur ADSL), mais cette hypothèse est irréaliste si l'on veut permettre des services de diffusion libres (ainsi organiser une visio-conférence entre amis, section 2.3) où aucune inscription préliminaire n'est nécessaire.

- *L'authentification/autorisation est seulement possible si une coopération opérateur/fournisseur de service existe:*

puisque l'opérateur de réseau doit interroger ou accéder à la base de données «client» du fournisseur de service. Ceci sera impossible dans certaines situations (ainsi si l'opérateur joue seulement un rôle de transit). Rien ne garantit qu'un client authentifié/autorisé se comportera correctement. Un client peut exploiter les vulnérabilités de certaines applications de façon à lancer des attaques sur le protocole de contrôle de congestion, ou bien simplement être victime d'un virus (un cheval de Troie, ou un root-kit installé sur le PC du client) qui lui-même profitera de l'authentification/autorisation du client pour lancer une attaque DoS. En présence de mécanismes de translation d'adresses NAT/PAT (Network Address Translation, Port Address Translation), l'authentification/autorisation du client sur la base de l'adresse IP est inefficace car l'adresse IP de la machine terminale

12. INFRASTRUCTURE SÉCURISÉE DE ROUTAGE MULTIPOINT: LE POINT DE VUE DE L'OPÉRATEUR RÉSEAU

est masquée au réseau de l'opérateur. Prenons le cas d'un utilisateur qui se connecte via sa passerelle domestique faisant du NAT et qui a construit un réseau local sans fil. Dès que cet utilisateur a rejoint le service de communication de groupe, tous les flux issus de sa passerelle sont considérés légitimes. Cependant un utilisateur dans le voisinage peut rejoindre ce réseau sans fil (les réseaux 802.11b sont connus pour avoir des failles de sécurité) et en profiter pour lancer une attaque DoS.

- *Modification des protocoles existants:*

Certaines propositions de sécurisation modifient les protocoles multicast existants, en ajoutant par exemple des mécanismes d'authentification à IGMPv3, ou définissent même de nouveaux protocoles de routage sécurisés, comme KHIP [65] qui s'appuie sur une variante du protocole CBT [12] (protocole qui n'a pourtant jamais été ni implanté ni déployé dans les réseaux opérationnels). Même si ces nouveaux protocoles sont extrêmement robustes à certains types d'attaques, ces solutions ne sont pas réalistes et n'ont pratiquement aucune chance d'être un jour déployées par un opérateur de réseau, très attentif à modifier le moins possible son infrastructure et ses protocoles réseau existants. En outre, même si une proposition est standardisée à l'IETF, il lui reste un long chemin à parcourir avant d'être déployée largement dans les réseaux opérationnels, et se posera alors le problème de l'interopérabilité avec des réseaux tiers qui n'ont pas migré vers la version sécurisée du protocole en question.

- *Confiance/dépendance entre domaines multicast:*

Un opérateur de réseau est responsable du niveau et de la qualité de service offert à ses clients. Aussi, des solutions de sécurité nécessitant la collaboration de plusieurs opérateurs réseau pour la création d'arbres multicast inter domaines sont difficiles à mettre en place et à déployer. Par conséquent, les exigences de sécurité d'un opérateur réseau, telles que définies dans la section 3, doivent être essentiellement égocentriques et ne pas être dépendantes d'autres opérateurs réseau.

- *Le Filtrage Classique:*

Le filtrage classique peut aider l'infrastructure mais il pénalise largement le trafic légitime dans le cas d'attaque.

12.7 Notre Proposition

Notre approche est basée sur un filtre, contrôlé par l'opérateur et située entre les clients et le routeur d'accès multicast. Le filtre capture les paquets IGMP (ou MLD) générés par les clients, les filtre selon des règles spécifiques, ensuite les renvoie au réseau. Le routeur multicast est configuré pour accepter les paquets IGMP provenant exclusivement du filtre, les autres paquets IGMP sont automatiquement rejetés. Lorsque un paquet IGMP est envoyé par un client inconnu, il est systématiquement mis dans une file d'attente FIFO dédiée dont la taille maximale est strictement imposée. Périodiquement un certain nombre de paquets IGMP mis dans la file d'attente FIFO de clients inconnus sont élus, et un contexte est créé pour chaque client qui est maintenant considéré comme "connu". Les futurs paquets IGMP arrivant des clients connus seront donc directement acceptés par le système et mis dans leur file d'attente

associée. Périodiquement, un certain nombre de paquets IGMP des files d'attente des clients connus sont choisis et renvoyés au réseau. Comme les clients vont disparaître, un système de purge est installé dans le filtre (Figure 12.2).

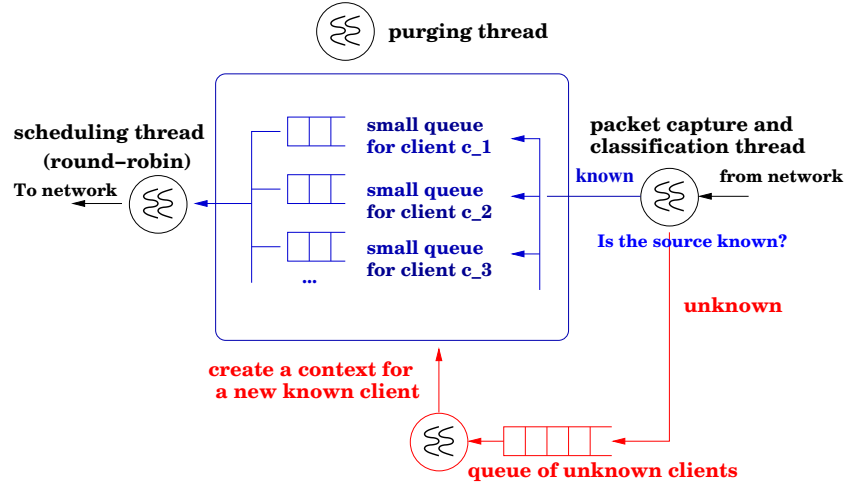



Figure 12.2: L'Architecture de Filtre.

12.8 Le Déploiement de Filtre

En général, deux déploiements sont possibles (Figure 12.3):

- *Déploiement Externe:* Dans ce cas le filtre est implémenté dans une machine indépendante, située entre les clients finaux et le routeur d'accès. Ce déploiement est universel vue qu'il n'est pas besoin de faire des modifications sur le routeur d'accès.
- *Déploiement Interne:* Dans ce cas le filtre est intégré dans le routeur d'accès. Ce déploiement nécessite de faire des modifications sur le routeur d'accès. Ainsi ce déploiement nécessite une collaboration technique avec le constructeur.



(a) Déploiement Externe

(b) Déploiement Interne

Figure 12.3: Le Déploiement de Filtre.

12. INFRASTRUCTURE SÉCURISÉE DE ROUTAGE MULTIPOINT: LE POINT DE VUE DE L'OPÉRATEUR RÉSEAU

12.9 Le Déploiement de Filtre dans le Réseau de l'Opérateur:

Cette section se focaliser sur deux déploiements exemples:

- *Le DSLAM est un IGMP proxy:*

Dans ce cas le DSLAM apparaîtra comme un routeur IGMP pour les clients finaux et comme un client pour le BRAS. Ainsi le DSLAM gère les abonnements aux groupes de ses clients. Le DSLAM envoie des requêtes IGMP vers le BRAS lorsque il reçoit une demande pour un nouveau group aussi il envoie un message LEAVE pour libérer un group lorsque le dernier membre de quitte ce groupe.

Le traitement de requêtes IGMP se fait principalement dans le DSLAM. Le filtre doit être implanté dans le DSALM. Or, afin de protéger le BRAS contre une attaque de la part d'un DSLAM corrompu, le filtre doit être placé avant/dans le BRAS.

- *Le DSLAM fait de IGMP snooping*

Dans ce cas la le DSLAM ne traite pas les requêtes IGMP. Plus précisément, les requêtes IGMP passent normalement au BRAS. Le DSLAM enregistre l'adresse MAC de group multicast et l'interface par laquelle les requêtes ont été reçues.

Comme peu de traitement se fait dans le DSLAM, le déploiement de filtre devant le DSLAM est optionnel. En revanche, il est nécessaire de déployer le filtre devant le BRAS.

12.10 Comment Initialiser le Filtre ?

12.10.1 Les paramètres de filtre

L'initialisation de filtre se fait par fixer les valeurs d'un ensemble de paramètres:

- K : la taille de la file d'attente de nouveaux clients UCQ;
- μ : le taux le taux de service de UCQ.
- *purge_period*: le temps entre deux purges.
- G_{max} : le nombre maximal de groupes par clients.
- s : le taux de sortie de filtre.

Certain paramètres peuvent être obtenues en appliquant un modèle mathématique. En revanche, les autres paramètres sont à fixer directement par l'opérateur en fonction des capacités des routeurs. Plus précisément, K , μ , et *purge_period*, sont obtenus théoriquement. Ainsi, G_{max} , et s sont à fixer par l'opérateur.

En pratique, l'ensemble de paramètres varie en fonction de l'infrastructure utilisée, les temporisations de IGMP/MLD et le comportement de client multicast.

12.10.2 Le dimensionnement de Filtre

Afin de dimensionner le filtre on considère qu'il est composé de deux parties, la première partie c'est la file d'attente des nouveaux clients et la deuxième partie c'est l'ensemble des files d'attente des clients connues.

Pour dimensionner la file d'attente de nouveaux clients on considère deux modèles pour l'arrivée de trafic IGMP.

- *Le trafic IGMP arrive avec un débit constant*: dans ce cas on modélise le filtre par une queue D/D/1/K, en revanche, dans le deuxième modèle le trafic suit la loi de poisson. Et dans ce cas on modélise le file d'attente de nouveau client par une file M/D/1/K. à partir de ces deux modèles on peut fixer les différents paramètres du filtre d'attente de nouveau client on est intéressé en particulier par le nombre de paquets qui sont en attente à un instant t ainsi que le taux de service.
- *Le trafic IGMP arrive selon un processus de Poisson* Maintenant pour dimensionner les files d'attente des clients connus on modélise chaque file d'attente par une file M/D/1/Gmax.

On peut aussi déterminer la bonne valeur pour la période de purge en modélisant celle-ci comme une fonction de nombre de clients dans le système ainsi que leur activité dans le système.

12.11 Les Résultats Expérimentaux

12.11.1 La plateforme de tests

Notre plateforme de tests est composée de deux routeurs et deux PC. Le filtre a été implémenté dans Linux. On a implémenté aussi le générateur de trafic qui génère le trafic IGMP selon deux modèles soit un modèle constant soit un modèle de poisson. Le filtre se situe entre le générateur de trafic et le routeur IGMP qui est un routeur Cisco 7500 (Figure 12.4).




Figure 12.4: The Testbed.

12.11.2 Les scénarios de tests

Plusieurs tests ont été faits afin d'évaluer les bénéfices de filtre en face d'une attaque IGMP par inondation. Plus précisément, ces tests visent à étudier comment l'utilisation

12. INFRASTRUCTURE SÉCURISÉE DE ROUTAGE MULTIPOINT: LE POINT DE VUE DE L'OPÉRATEUR RÉSEAU

de filtre peut améliorer la robustesse de l'infrastructure aux attaques. Ainsi comment le filtre peut protéger les clients légitimes qui arrivent avant et pendant une attaque.

On a fait trois types de tests: Dans le premier test l'attaquant spoof une seule adresse et il envoie de requêtes IGMP sur 500 groupes multicast différent de ceux des clients légitimes. Dans le deuxième test L'attaquant spoof 10 adresses et dans le dernière test l'attaquant spoof 1000 adresses.

Les figures 12.5(a) 12.5(b) montrent le trafic entrant et sortant de filtre. On voit d'emblée les bénéfices de filtre dans ce cas. Le nombre de requêtes appartenant à l'attaquer et qui sont servi par le filtre est beaucoup moins de ceux qui arrivent au filtre. Ainsi, les clients légitimes qui arrivent avec l'attaque sont servis normalement.

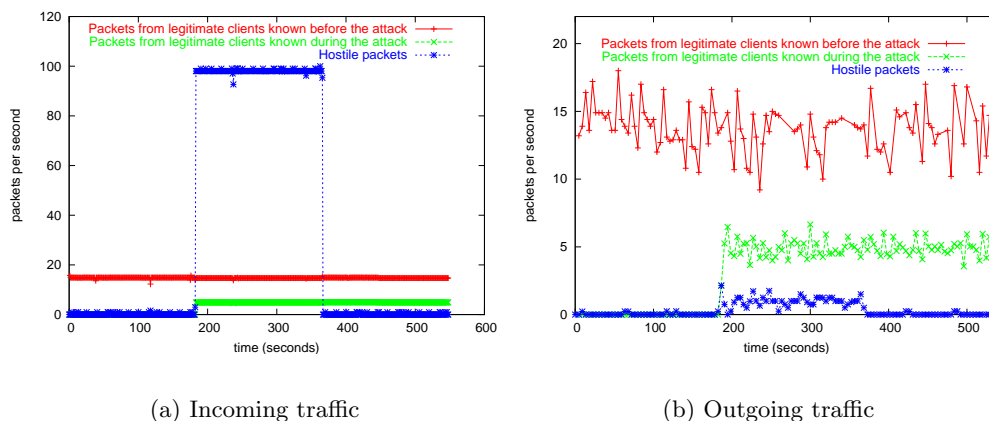


Figure 12.5: Le Trafic Entrant et Sortant (Test 1).

Les figures 12.6(a) et 12.6(b) montrent le nombre de messages PIM générés sans et avec filtre. Cette figure montre l'amélioration que le filtre peut rapporter sur toute l'infrastructure.

La figure 12.7 montre la consommation de mémoire de routeur sans et avec filtre. Cette figure montre qu'il y a une épargne considérable dans la consommation de mémoire avec le filtre. Or, la consommation de mémoire lorsque le filtre est utilisé est 1/6 par rapport au cas sans filtre.

La figure 12.8(b) montre le cas extrême ou l'attaquant spoof 1000 adresses. Dans ce cas les clients légitimes peuvent être largement pénalisés mais le routeur d'accès reste protégé.

12.12 Discussion des Extensions Possibles

Les résultats expérimentaux montrent que notre approche concertiez l'objective ultime de l'opérateur qui est la continuation de service en débit de tout anomalies. Plus précisément:

- Les clients arrivant avant les attaques reste servi normalement.

12.12 Discussion des Extensions Possibles

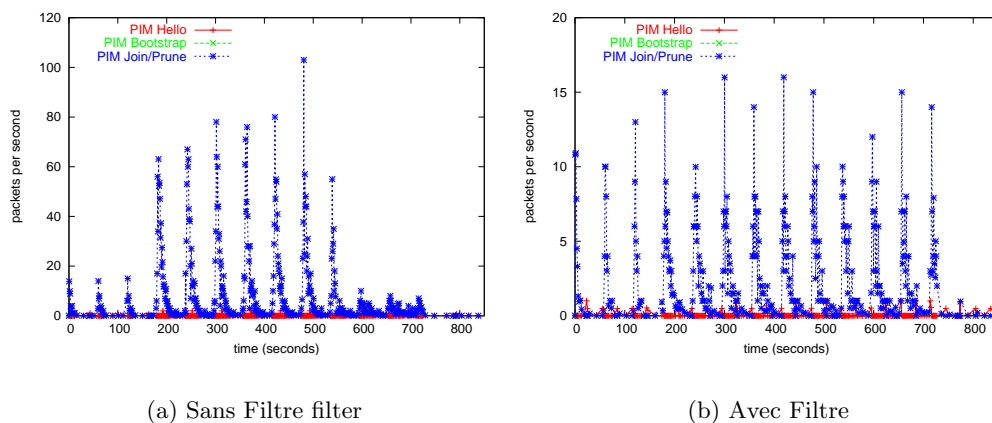


Figure 12.6: PIM Messages Sans/Avec Filtre (Test 1).




Figure 12.7: La Consommation de Mémoire sur le Routeur Cisco Avec/Sans filtre (Test 1)

- Le taux de trafic qui quitte le filtre est toujours contrôlé.
- il protège le mémoire de routeurs.
- il protège le réseau de coeur.
- il n'a pas besoin de modifier l'infrastructure existant.

En revanche, le filtre est vulnérable aux adresses spoofing. Dans le suivant on va discuter des extensions possibles pour améliorer la robustesse de filtre au spoofing.

En général, il y a des politiques de sécurité qui peuvent aider pour améliorer la robustesse de réseau au spoofing. Quelques politiques possibles :

- Vérifier l'adresse source de paquets et rejeter n'importe quel paquet avec une adresse suspect.

12. INFRASTRUCTURE SÉCURISÉE DE ROUTAGE MULTIPOINT: LE POINT DE VUE DE L'OPÉRATEUR RÉSEAU

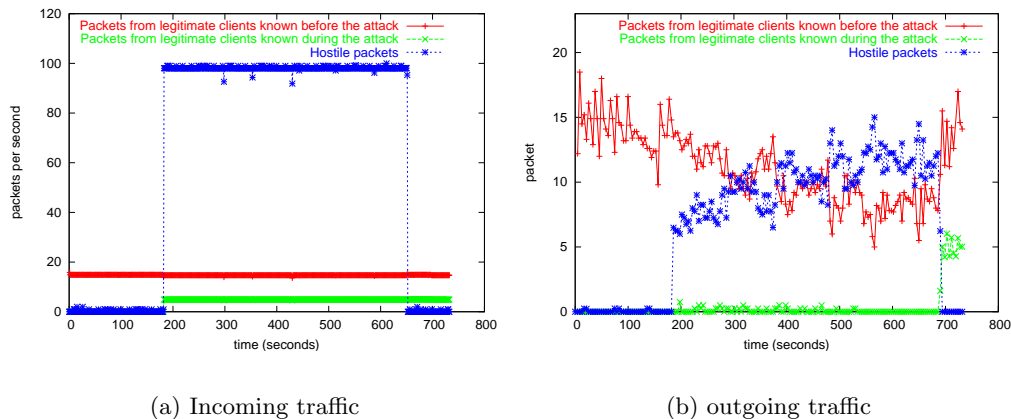


Figure 12.8: Les trafic Entrant et Sortant de Filtre (Test 3).

- Interagir avec le serveur DHCP afin d'obtenir les couples IP/MAC. Ainsi utiliser un nouveau option de protocole DHCP qui est l'option 82. cet option permet d'associer a chaque adresse IP/Mac un identifiant non forgeable.

12.12.1 l'Utilisation des identifiants non forgeables

Le réseau de l'opérateur fourni un ensemble des identifiants non forgeable (i.e. les adresses ATM, le VLAN-tag ou le DSL ID) qui peuvent être utilisés par afin d'éviter le spoofing des adresses.

Afin d'utiliser ces identifiants on peut distinguer entre deux cas:

- Le filtre est implémenté dans le DSLAM:
Le DSLAM peut alors obtenir ces identifiants directement vue que c'est lui qui termine les liens DSL.
- Le filtre est implémente dans le BRAS :
Dans ce cas l'obtention des identifiants non forgeable pour chaque client peut s'avérer difficile parfois. Mais l'utilisation de quelques mécanismes comme le DHCP relay avec l'option 82 permet de diminuer le spoofing et ainsi le besoin de ces identifiants.

12.12.2 l'Utilisation des mécanismes d'apprentissage

Le but de ces mécanismes est de prendre des décisions sur les paquets avant de les accepter dans le filtre. Ainsi ces mécanismes de classifier les clients en clusters de priorités. Or, l'ordonnancement des paquets dans les fils d'attente de clients acceptés se fait en considérant ces priorités. Donc, des clients avec des mauvais préretraités auront moins de chance à être servis.

L'algorithme d'apprentissage prendre trois paramètres.

- * L'intensité de trafic.
- * La conformité avec les automates IGMP.
- * La familiarité des adresses IP/MAC.

En fonction de ces paramètres plusieurs décisions seront prises, comme par exemple rejeter les paquets directement, ou accepter le nouveau client et lui donner une priorité de service. L'utilisation des mécanismes d'apprentissage améliorer le filtrage d'une coté et permet de lutter contre des autres attaques basées sur IGMP/MLD d'autre côté.

12.12.3 L'utilisation d'une politique de gestion de groups

L'objectif de cette politique est d'associer chaque client par les groups auquel Il est autorisé a participer.

12.12.4 l'utilisation d'une politique de control de congestion

L'objectif de la politique de control de congestion est que les clients seront autorisés à s'abonner à un group en fonction de leur état de congestion.

12.13 Concluions

Ce travail cible la sécurité de l'infrastructure multicast du point de vue de l'opérateur réseau. Ses exigences de sécurité sont largement différentes de celles des utilisateurs finaux ou des fournisseurs de service. Ainsi est-il essentiellement préoccupé par une exigence de « continuité de service », notamment lorsque son réseau est victime d'une attaque. L'infrastructure de l'opérateur est particulièrement vulnérable aux attaques DoS qui consomment des ressources de bande passante et compromettent les capacités de traitement des routeurs. Ces attaques DoS, intentionnelles ou non, sont faciles à lancer par des utilisateurs localisés à la périphérie du réseau de l'opérateur. L'opérateur doit donc porter tous ses efforts afin de se prémunir de ces attaques. En revanche, vouloir sécuriser les protocoles de routage multicast dans le réseau coeur de l'opérateur afin d'éviter qu'un routeur corrompu n'altère les arbres de distribution multicast a une importance très secondaire puisque cette attaque est hautement improbable. Plusieurs hypothèses couramment faites doivent aussi être évitées. En premier lieu, sécuriser IGMP par une authentification/autorisation systématique du client n'est ni toujours possible, ni efficace (un utilisateur légitime et authentifié peut conduire des attaques s'il est par exemple infecté par un virus). En second lieu, proposer des modifications à des protocoles existants largement déployés n'est pas efficace car les opérateurs, très conservateurs, recherchent avant tout des solutions pratiques et réalistes plutôt que des solutions intellectuellement idéales.

A la lumière de l'analyse détaillée de la problématique, des vulnérabilités et des solutions actuelles, nous proposons une nouvelle approche pour aider le réseau de l'opérateur à se défendre contre les attaques basées sur IGMP ou MLD. Notre proposition suit une approche pragmatique et flexible, qui garantit qu'elle sera facilement déployable dans les infrastructures existantes, et vise également à protéger les clients légitimes en cas d'attaque.

12. INFRASTRUCTURE SÉCURISÉE DE ROUTAGE MULTIPOINT: LE POINT DE VUE DE L'OPÉRATEUR RÉSEAU

References

- [1] Cisco Group Management Protocol. <http://www.javvin.com/protocolCGMP.html>. 23
- [2] Multicast Security (MSEC) Working Group. <http://www.ietf.org/html.charters/msec-charter.html>. 3, 102
- [3] Gigabit Campus Network Design– Principles and Architecture, 1999. Cisco White Paper, <http://www.cisco.com/>. 92
- [4] Ieee std 802.1x-2001, port-based network access control, 2001. 91
- [5] IGMP State Limit, Cisco IOS Documentation, 2003. <http://www.cisco.com/>. 48, 108
- [6] Allied Telesyn, White Paper: Architectures for IP Video Providing IP Video Services via DSL and FTTx Networks, July 2004. <http://www.alliedtelesyn.com/>. 90
- [7] Andrew Adams, Tian Bu, Ramón Cáceres, Nick Duffield, Timur Friedman, Joseph Horowitz, Francesco Lo Presti, Sue Moon, Vern Paxson, and Don Towsley. The Use of End-to-End Multicast Measurements for Characterizing Internal Network Behavior. *IEEE Communications Magazine*, May 2000. 48
- [8] Andrew Adams, Jonathan Nicholas, and William Siadak. Protocol Independent Multicast-Dense Mode (PIM-DM(Revised)): Protocol Specification, January 2005. IETF (Request for Comments) RFC 3973. 23
- [9] B. Adamson, C. Bormann, M. Handley, and J. Macker. Negative-Acknowledgment (NACK)-Oriented Reliable Multicast (NORM) Building Blocks, November 2004. IETF (Request for Comments) RFC3941. 32
- [10] Kevin C. Almeroth and Mostafa H. Ammar. Multicast Group Behavior in the Internet’s Multicast Backbone (MBone). *IEEE Communications*, 35(6):224–229, June 1997. 63, 74, 80
- [11] Sara Alouf, Eitan Altman, Chadi Barakat, and Philippe Nain. Estimating Membership in a Multicast Session. *ACM SIGMETRICS Performance Evaluation Review*, 31(1):250–260, June 2003. 63, 74, 80
- [12] Tony Ballardie. Core Based Trees (CBT) Multicast Routing Architecture, September 1997. IETF request for comments RFC2201. 23, 47, 51, 110

REFERENCES

- [13] Tony Ballardie and Jon Crowcroft. Multicast-Specific Security Threats and Counter-Measures. In *Symposium on Network and Distributed System Security (SNDSS'95)*, San Diego, California, February 1995. 45, 107
- [14] Tony Bates, Ravi Chandra, Dave Katz, and Yakov Rekhter. Multiprotocol extensions for bgp-4, June 2000. IETF request for comments RFC2858. 28
- [15] João B. D. Cabrera, Lundy Lewis, Xinzhou Qin, Wenke Lee, Ravi K. Prasanth, B. Ravichandran, and Raman K. Mehra. Proactive detection of distributed denial of service attacks using mib traffic variables-a feasibility study. In *The Seventh IFIP/IEEE International Symposium on Integrated Network Management (IM 2001)*, Seattle, WA, May 2001. 49
- [16] Brad Cain, Steve Deering, Bill Fenner, Isidor Kouvelas, and Ajit Thyagarajan. Internet Group Management Protocol, Version 3, October 2002. IETF request for comments RFC3376. 2, 14, 18, 101
- [17] Ran Canetti and Benny Pinkas. Taxonomy of Multicast Security Issues, April 1999. IETF Internet draft, draft-irtf-smug-taxonomy-01.txt, work in progress. 3, 102
- [18] Claude Castelluccia and Gabriel Montenegro. Securing Group Management in IPv6 with Cryptographically Generated Addresses. In *Proceedings of the Eighth IEEE International Symposium on Computers and Communications*, page 588, Washington, DC, USA, 2003. 41, 45, 49, 108
- [19] Ghassan Chaddoud. *Sécurisation de Communication de Groupes Dynamiques*. PhD thesis, université Henri Poincaré - Nancy 1, 2002. 45, 107
- [20] Shun Yan Cheung, Mostafa Ammar, and Xui Li. On the Use of Destination Set Grouping to Improve Fairness in Multicast Video Distribution. In *Fifteenth Annual Joint Conference of the IEEE Computer and Communications Societies*, San Francisco, USA, March 1996. 32
- [21] Steven Cheung. *An Intrusion Tolerance Approach for Protecting Network Infrastructures*. PhD thesis, University of California, Davis, September 1999. 48, 108
- [22] Abdur Chowdhury, Ophir Frieder, and Peng-Jun Wan. On the design, development, deployment, and network survivability analysis of the dynamic routing system protocol. *The Journal of Supercomputing*, March 2002. 48
- [23] Greg Daley and Gopi Kurup. Trust Models and Security in Multicast Listener Discovery, July 2004. IETF/MAGMA Internet draft, draft-daley-magma-sml-d-prob-00.txt, work in progress. 29, 37
- [24] Steve Deering. Host Extensions for IP Multicasting, August 1989. IETF request for comments RFC1112. 1, 13, 14, 15, 101
- [25] Steve Deering, Bill Fenner, and Brian Haberman. Multicast Listener Discovery (MLD) for IPv6, October 1999. IETF request for comments RFC2710. 2, 22, 101

REFERENCES

- [26] Christophe Diot, Brian Neil Levine, Bryan Lyles, Hassan Kassem, and Doug Balensiefen. Deployment issues for the ip multicast service and architecture. *IEEE Network magazine special issue on Multicasting*, January/February 2000. 2
- [27] Deborah Estrin, Dino Farinacci, Ahmed Helmy, Dave Thaler, Steve Deering, Van Jacobson, Mark Handley, Charley Liu, Puneet Sharma, and Liming Wei. Protocol Independent Multicast-Sparse Mode (PIM-SM): Protocol Specification, June 1998. IETF/IDMR (Request for Comments) RFC2362. 23, 31, 101, 106
- [28] Bill Fenner and David Meyer. The Multicast Source Discovery Protocol (MSDP), October 2003. IETF (Request for Comments) RFC3618. 2, 101
- [29] William C. Fenner. Internet Group Management Protocol, Version 2, November 1997. IETF request for comments RFC2236. 14, 15, 101
- [30] Lenny Giuliano. Juniper Networks, Inc, Application Note, Calm During the Storm: Best Practices in Multicast Security, 2005. <http://www.juniper.net/solutions/>. 48, 49, 108
- [31] A.F. Gomez-Skarmeta, A.L. Mateo-Martinez, and P.M. Ruiz-Martinez. Icmpv3-based method for avoiding dos attacks in multicast-enabled networks. In *25th Annual IEEE Conference on Local Computer Networks (LCN'00)*, Tampa, Florida, November 2000. 45, 107
- [32] L. Gong and N. Shacham. Trade-offs in routing private multicast traffic. In *In Proceedings of GLOBECOM '95*, Singapore, November 1995. 47
- [33] Sergey Gorinsky, Sugat Jain, Harrick Vin, and Yongguang Zhang. Robustness of multicast congestion control to inflated subscription. In *ACM SIGCOMM 2003*, Karlsruhe, Germany, August 2003. 40, 45, 97, 107
- [34] Donald Gross and Carl M. Harris. *Fundamentals of Queueing Theory (2nd ed.)*. John Wiley & Sons, Inc. New York, NY, USA, June 1985. 64, 65
- [35] George Gross, Manu Kaycee, Arthur Lin, Andrew Malis, and John Stephens. PPP Over AAL5, July 1998. Network working group request for comments RFC2364. 8
- [36] Thomas Hardjono and Brad Cain. PIM-SM Security: Interdomain Issues and Solutions. In *Communications and Multimedia Security CMS*, Leuven, BELGIUM, September 1999. 31, 47, 106
- [37] Thomas Hardjono and Brad Cain. Key Establishment for IGMP Authentication in IP Multicast. In *IEEE European Conference on Universal Multiservice Networks (ECUMN)*, Colmar, France, October 2000. 44, 45, 107
- [38] Thomas Hardjono and Lakshminath R. Dondeti. *Multicast and Group Security*. Artech House, Incorporated, June 2003. 47
- [39] Tristan Henderson and Saleem Bhatti. Modelling User Behaviour in Networked Games. In *Proceedings of the ninth ACM international conference on Multimedia*, pages 212–220, Ottawa, Canada, 2001. 63, 74, 80

REFERENCES

- [40] Mickaël Hoerdt, Damien Magoni, and Jean-Jacques Pansiot. Evaluation de l'Impact d'Attaques Distribuées par Déni de Service Utilisant un Protocole Multipoint à Source Unique. In *JDIR'04: 6èmes Journées Doctorales Informatique et Réseau*, Lannion, France, Novembre 2004. 29
- [41] Hugh Holbrook and Brad Cain. Source-specific multicast for ip, September 2004. IETF Internet Draft:draft-ietf-ssm-arch-06.txt, work in progress. 14
- [42] Hugh W. Holbrook and David R. Cheriton. IP Multicast Channels: EXPRESS Support for Large-Scale Single-Source Applications. In *SIGCOMM '99: Proceedings of the conference on Applications, technologies, architectures, and protocols for computer communication*, pages 65–78, New York, NY, USA, August 1999. 14
- [43] Yih-Chun Hu, Adrian Perrig, and David B. Johnson. Efficient Security Mechanisms for Routing Protocols. In *The 10th Annual Network and Distributed System Security Symposium*, San Diego, California, February 2003. 43, 47
- [44] Volker Hösel and Sebastian Walcher. Clustering techniques: A brief survey. 98
- [45] Van Jacobson, Craig Leres, and Steven McCanne. PCAP - Packet Capture Library. <http://www.tcpdump.org/>. 55
- [46] Jeanne De Jaegher, Piet Vandaele, Dominique Chantrain, and Davy Damien. Multi-Service Ethernet Broadband Access Solutions, 2004. Alcatel White Paper, <http://www.alcatel.com/>. 91
- [47] Paul Judge and Mostafa Ammar. Gothic: A Group Access Control Architecture for Secure Multicast and Anycast. In *IEEE INFOCOM 2002*, NY,USA, June 2002. 45, 107
- [48] Zainab Khallouf, Vincent Roca, Renaud Moignard, and Sébastien Loyer. Infrastructure Sécurisée de Routage Multipoint : le Point de Vue de l'Opérateur de Réseau. In *SAR2003: 3ème Conférence sur la Sécurité et Architectures Réseaux (SAR'04)*, La Londe, Cote d'Azur, France, June 2004. 30
- [49] Zainab Khallouf, Vincent Roca, Renaud Moignard, and Sébastien Loyer. A Filtering Approach for an Icmp Flooding Resilient Infrastructure. In *SAR2004: 4ème Conférence sur la Sécurité et Architectures Réseaux (SAR'04)*, Batz sur Mer, France, June 2005. 49, 53
- [50] Arnaud Legout, Jorg Nonnenmacher, and Ernst W. Biersack. Bandwidth-allocation Policies for Unicast and Multicast Flows. *IEEE/ACM Transactions on Networking (TON)*., 9(4):464–478, 2001. 2
- [51] Rami Lehtonen and Jarmo Harju. Controlled Multicast Framework. In *The 27th Annual IEEE Conference on Local Computer Networks (LCN)*, Tampa, Florida, USA, November 2002. 49, 108
- [52] Michael Luby and Vivek Goyal. Wave and Equation Based Rate Control (WEBRC), April 2004. IETF Internet draft, draft-ietf-rmt-bb-webrc-04.txt, work in progress. 32, 106

- [53] Louis Mamakos, Kurt Lidl, Jeff Evarts, David Carrel, Dan Simone, and Ross Wheeler. A Method for Transmitting PPP Over Ethernet (PPPoE), February 1999. Network working group request for comments RFC2516. 8
- [54] Jelena Mirkovic, Janice Martin, and Peter Reiher. A Taxonomy of DDoS Attacks and DDoS Defense Mechanisms. Technical Report CSD-TR-0200180, University of California, Los Angeles (UCLA), 2002. 49
- [55] Suvo Mittra. Iolus: A Framework for Scalable Secure Multicasting. In *ACM SIGCOMM '97*, Cannes, France, October 1997. 3, 102
- [56] Annelies Van Moffaert and Olivier Paridaens. Security Issues in Protocol Independent Multicast -Sparse Mode (PIM-SM), February 2002. IETF Internet draft, draft-irtf-gsec-smrac-00.txt, work in progress. 47
- [57] D. Ooms, B. Sales, W. Livens, A. Acharya, F. Griffoul, and F. Ansari. Overview of IP Multicast in a Multi-Protocol Label Switching (MPLS) Environment, August 2002. IETF (Request for Comments) RFC3353. 48
- [58] Hiroyuki Oouchi, Ken Takahashi, Hiromichi Nagata, and Kouichi Kamasawa. Multi-Rate Control Method Using Layered Content. In *SAINT '05: Proceedings of the The 2005 Symposium on Applications and the Internet (SAINT'05)*, pages 311–317, Washington, DC, USA, 2005. IEEE Computer Society. 97
- [59] Michael Patrick. DHCP Relay Agent Information Option, November 1998. IETF Internet Draft , work in progress. 93
- [60] Adrian Perrig, Ran Canetti, J.Đ. Tygar, and Dawn Song. The TESLA Broadcast Authentication Protocol. *RSA CryptoBytes*, 5(2), Summer 2002. 3, 46, 102
- [61] Prashant Rajvaidya, Krishna Ramachandran, and Kevin C. Almeroth. Detection and Deflection of DoS Attacks Against the Multicast Source Discovery Protocol. Technical report, Department of Computer Science, University of California, Santa Barbara, July 2002. 31, 105
- [62] Prashant Rajvaidya, Krishna N. Ramachandran, and Kevin C. Almeroth. Managing and Securing the Global Multicast Infrastructure. *Journal of Network and Systems Management*, 12(3):297–326, September 2004. 48, 108
- [63] Luigi Rizzo, Gianluca Iannaccone, Lorenzo Vicisano, and Mark Handley. PGMCC Single Rate Multicast Congestion Control: Protocol Specification, July 2004. IETF Internet Draft: draft-ietf-rmt-bb-pgmcc-03.txt, work in progress. 32, 106
- [64] Pekka Savola and James Lingard. Last-hop Threats to Protocol Independent Multicast (PIM), January 2005. IETF Internet draft, draft-savola-pim-lasthop-threats-01.txt, work in progress. 29
- [65] Clay Shields and J. J. Garcia-Luna-Aceves. KHIP-A Scalable Protocol for Secure Multicast Routing. In *ACM SIGCOMM Computer Communication Review*, Harvard University Science Center, Cambridge, Massachusetts, USA, October 1999. 47, 50, 51, 107, 110

REFERENCES

- [66] Young-Chul Shim. A New Approach for Secure Multicast Routing in a Large Scale Network. In *ICICS '01: Proceedings of the Third International Conference on Information and Communications Security*, pages 95–106, London, UK, 2001. Springer-Verlag. 47
- [67] Scott Shoaf and Jerome Moisand. Juniper Networks, Inc, Application Note, IGMP Capabilities in Broadband Network Architectures, Mars 2005. <http://www.juniper.net/>. 8, 89, 91
- [68] T. Speakman, J. Crowcroft, J. Gemmell, D. Farinacci, S. Lin, D. Leshchiner, M. Luby, T. Montgomery, L. Rizzo, A. Tweedly, N. Bhaskar, R. Edmonstone, R. Sumanasekera, and L. Vicisano. PGM Reliable Transport Protocol Specification, December 2001. IETF (Request for Comments) RFC3208. 32
- [69] Subramaniam R. Sthanu and Steven R. Lerman. Survivability through dynamic reconfiguration. In *The Advanced Telecommunications and Information Distribution Research Program Conference*, College Park, Maryland, February 1998. 48
- [70] Dave Thaler. Border Gateway Multicast Protocol (BGMP): Protocol Specification, January 2004. IETF Internet Draft: draft-ietf-bgmp-spec-06.txt, work in progress. 2, 28, 101
- [71] Fouad A. Tobagi, Pablo Molinero-Fernández, and Mansour J. Karam. Study of ieee 802.1p garp/gmrp timer values, September 1997. Computer Systems Laboratory, Stanford University. 23
- [72] Rolland Vida and Luis Henrique Maciel Kosmowski Costa. Multicast Listener Discovery Version 2 (MLDv2) for IPv6, June 2004. IETF request for comments) RFC3810. 2, 101
- [73] B. Volz. Reclassifying Dynamic Host Configuration Protocol version 4 (DHCPv4) Options, November 2004. IETF (Request for Comments) RFC3942. 8
- [74] David Waitzman, Craig Partridge, and Stephen E. Deering. Distance Vector Multicast Routing Protocol, November 1988. IETF request for comments RFC1075. 23
- [75] Feiyi Wang. *Vulnerability Analysis, Intrusion Prevention and Detection for Link State Routing Protocols*. PhD thesis, Carolina State University, Dec. 2000. 48, 108
- [76] Ning Wang and George Pavlou. Scalable Sender Access Control for Bidirectional Multicast Routing. *Computer Networks: The International Journal of Computer and Telecommunications Networking*, 43, 2003. 46, 107
- [77] Chung Kei Wong, Mohamed Gouda, and Simon S. Lam. Secure group communications using key graphs. In *IEEE/ACM Transactions*, New York, USA, November 2000. 3, 102
- [78] Ishan Wu and Toerless Eckert. Router-port Group Management Protocol (RGMP), February 2003. IETF (Request for Comments) RFC3488. 23

REFERENCES

- [79] William Yurcik and David Tipper. A Survivability Framework for Connection-Oriented Group Communications. In *IEEE Pacific Rim International Symposium on Dependable Computing 2000 (PRDC 2000)*, University of California at Los Angeles (UCLA), USA, December 2000. [4](#), [48](#), [105](#)