AN ALTERNATIVE ACCESS CONTROL ARCHITECTURE FOR IP OVER ATM NETWORKS*

Olivier Paul, Maryline Laurent
ENST de Bretagne, RSM Department,
2 rue de la chataîgneraie, BP 78,
35512 Cesson-Sévigné,France.
Phone: (33) (0) 299127051.
Email: {Olivier.Paul|Maryline.Laurent}@enst-bretagne.fr

Abstract In this article, we describe a new architecture providing the access control service in both ATM and IP-over-ATM networks. This architecture is based on management agents distributed in network equipment. Several examples are given illustrating the benefits of this architecture. The comparison with other approaches shows that this architecture provides big improvements in ATM-level access control, scalability and QoS preservation.

Keywords Access Control, Management, Security, ATM, Agents, MIBs, IP-over-ATM.

1. INTRODUCTION

In the recent past, much attention has been paid to develop security services for ATM networks. This resulted in the creation of many working groups within (and outside) the standardization bodies. One of them is the security Working Group of the ATM Forum created in 1995, which is to the point to release its version 1.0 specifications. Confidentiality, authentication, integrity and some kind of access-control have been considered. Access control as defined by the ISO in [7498-2] is a security service used to protect resources against an unauthorized use.

* This work is funded by DRET.

B. Preneel (ed.), *Secure Information Networks*

The ATM technology has been specified to transport various kinds of flows and allows users to specify the QoS (Quality of Service) applying to these flows. Communications are connection oriented and a signaling protocol is used to set up, control and release connections. In this article we show that the classical approach supplying the access control service (commonly called firewall) is unable to preserve the QoS. We then describe a new access control architecture for ATM and IP-over-ATM networks which does not alter the negotiated QoS.
The next section analyses the current solutions providing the access control service in the ATM and IP over ATM networks. Section 3 describes the way to retrieve access control information from the MIBs (Management Information Bases) and to provide the access control service through our access control architecture. As a conclusion we make a comparison between our solution and other proposed approaches and we show that our architecture is a good alternative to current solutions.

2. PROPOSED SOLUTIONS

Several solutions have been proposed in order to provide some kind of access-control in ATM and IP over ATM networks. This section is divided into three parts. In the first part we consider the adaptation of the internet «classical» firewall architecture to ATM networks. In the second part we describe the solution proposed by the ATM Forum. In the third part we show various solutions proposed to improve the «classical» firewall solution.

2.1 Classical solution

The first solution [Ran92] is to use a classical firewall located between the internal and public networks in order to provide access-control at the packet, circuit and application levels. As such the ATM network is considered as a level 2 OSI layer offering point to point connections. As a result access-control at the ATM level is not possible and end to end QoS is no longer guaranteed.
At the IP and circuit levels, IP packets are reassembled from the ATM cells. Access-control is supplied using the information embedded in the TCP, UDP and IP headers. Packets are filtered by comparing the fields in the headers such as the source and destination addresses, the source and destination ports, the direction and the TCP flags with a pattern of prohibited and allowed packets. Prohibited packets are destroyed whereas allowed packets are forwarded from one interface to the other. When the same QoS is

negotiated on both sides of the firewall, the end to end QoS may be modified in the following ways:

- Reassembly, routing, filtering and deassembly operations increase the Cell Transit Delay.
- Internal operations done over IP packets may increase the Cell Loss Ratio.
- The time spent to reassemble and deassemble the packets is proportional to the packets sizes, which are variable. As a result, the Cell Transit Delay Variation may be different from the CTDV value negotiated on each side of the firewall.
- Routing and filtering actions operate at the software level. Thus the load of the system may cause variations in the Peak, Sustainable and Minimum Cell Rate.

Application procedures are then filtered at the application level by proxy applications in accordance with the security policy. Like with the IP or circuit level filters, the QoS is affected, but much more strongly, since the traffic has to reach the application level. Moreover since the filtering operations are provided in a multitasking environment, a desynchronization between the flows can occur.
This kind of solution is reported to have performance problems in high speed networks environment ([Data97], [JA98]). The latest tests ([KL98]) show that this access control solution is unsuccessful at the OC-3 speed (155Mb/s).

2.2 The access control service as considered by the ATM Forum

The access-control service as defined in the ATM Forum specifications ([SEC1.0]) is based on the access-control service provided in the A and B orange book classified systems. In this approach one sensitivity level per object and one authorization level per subject are defined. Those levels include a hierarchical level (e.g. public, secret, top secret...) and a set of domains modeling the domains associated with the information (e.g. management, research, education...). A subject may access an object if the level of the subject is greater than the level of the object and one of the domain associated with the subject includes one of the domain associated with the object.
In the ATM Forum specifications, the sensitivity and authorization levels are coded as a label, which is associated to the data being transmitted. This label may be sent embedded into the signaling, or as user data prior to any user data exchanges. The access-control is operated by the network equipment

which verifies that the sensitivity level of the data complies with the authorization level assigned to the links and interfaces over which the data are transmitted.
The main advantage of this solution is its scalability since the access control decision is made at the connection setup and doesn't interfere with the user data. However it suffers from the following drawbacks:

- The network equipment is assumed to manage sensitivity and authorization levels. This is not provided in current network equipment.
- A connection should be set up for each sensitivity level.
- The access-control service as considered in traditional firewalls (i.e. access-control to hosts, services) is voluntarily left outside the scope of the specification.

2.3 Specific solutions

The above limitations have been identified and many proposals have been made in order to supply the «traditional» access-control service in ATM networks. These solutions may be classified into two classes: industrial and academic solutions.

Industrial solutions

The first industrial solution (Cisco, Fore) uses a classical ATM switch that is modified to filter ATM connection set up requests based on the source and destination addresses. The problem with this approach is that the access-control is not powerful since the parameters are very limited.

The second one (Storagetek) is also based on an ATM switch, however this switch has been modified to supply access-control at the IP level. Instead of reassembling cells for packets headers examination like in traditional firewalls, this approach is expected to find IP and TCP/UDP information directly in the first ATM cell being transmitted over the connection. This approach prevents delays to be introduced during cell switching. Storagetek has also developed a specific memory called CAM (Content Addressable Memory) designed to speed up the research in the access-control policy. This approach is the first one taking into account the limitations introduced by the classical firewall approach. However some problems have not yet been solved:

- Access-control is limited to the network and transport levels. ATM and application levels are not considered.
- IP packets including options are not filtered since options may shift the UDP/TCP information in the second cell. This causes a serious security flaw.

- The device is not easy to manage especially when dynamic connections are required, since connection filters have to be configured manually.
- Performances of the device are not very scalable. An OC-12 (622 Mb/s) version of this product has been announced in 1996 but is not yet exhibited.

Academic solutions

Both academic solutions being proposed are based on the above Storagetek architecture, but they introduce some improvements to cope with Storagetek problems.

The first approach [Da98] uses a FPGA specialized circuit associated to a modified switch architecture. At the ATM level, the access control at connection establishment is improved by providing filtering capabilities based on the source and destination addresses. This approach also allows ATM level PNNI (Private Network to Network Interface) routing information to be filtered. At the IP and circuit levels the access-control service is similar to this provided by the Storagetek product.
This solution is interesting since it is the most complete solution being currently implemented. However it suffers from many limitations:

- Special IP packets (e.g. packets with optional fields in the header) are not processed.
- Only a small part of the information supplied by the signaling (i.e. source and destination addresses) is used.
- Access-control at the application level is not considered.

The second approach [XS97] is the most complete architecture being currently proposed. This solution provides many improvements over the Storagetek architecture. The most interesting idea is the classification of the traffic. The traffic is classified into four classes depending on the ATM connection QoS descriptors and on the processing allowed to be done over it. Class A provides a basic ATM access-control. ATM connections are filtered according to the information provided by the signaling (i.e. source and destination addresses). Class B provides traffic monitoring. The analysis of the traffic is made on a copy of the flow. When a packet is prohibited, the reply to this packet is blocked. Class C is associated with packet filtering. IP and transport packet headers are reassembled from the ATM cells and analysed. During this analysis the last cell belonging to the packet called LCH (Last Cell Hostage) is kept in memory by the switch. The analysis should be at least faster than the time spent by the whole packet to cross the switch. When the packet is allowed, the LCH is released, but when the packet is prohibited the LCH is modified so that a CRC error occurs and the

packet is rejected. For class D, the access control processing is similar to the firewall proxy's.
This classification expects the switch to separate traffics with QoS requirements from traffics without QoS requirements. As such the traffic with QoS requirements is allowed to cross the switch without being delayed. Table 1 gives the filtering operations depending on the level implementing the access control and the traffic QoS requirements.

Table 1: Use of the Access Control Classes

Level/application	With QoS Requirements	Without QoS Requirements
Application	No Access Control	Class D
Transport	Class B	Class C
ATM	Class A	Class A

This approach is interesting since it introduces many improvements (traffic classification, LCH) over all the other proposed solutions. However some problems remain:

- Few parameters are used to supply the access control service at the ATM level.
- Access control is not provided at the application level for applications requiring QoS.
- Traffic monitoring only applies to connection oriented communications, and UDP packets can not be filtered using this technique.
- This architecture is complex so that it is likely that scalability is not offered
- No implementation has been exhibited.

The problems most oftenly met are the lack of scalability and the impact on QoS introduced by the access control service. As a consequence, it appears interesting to develop a scalable architecture that could provide the access control service while maintaining the negotiated QoS.

3. AN AGENT BASED ACCESS CONTROL ARCHITECTURE

The goal of our architecture is to provide a scalable access-control service without altering the QoS negotiated for a connection. We selected a distributed architecture approach to have more scalability than in a

centralized approach. As stated in [Schu98] a distributed architecture induces many advantages. These advantages are as follows:

- Better fault tolerance. If a device providing the access control service fails, this device is the only one affected. Other devices are able to continue to communicate.
- Security level improvement. For an intruder to control the whole network, it is necessary to subvert all the access control devices one after the other.
- Protection against internal attacks. Internal attacks can be avoided and detected since all the devices are protected.
- Realistic information about the flows. [PN98] shows that firewalls and intrusion detection tools systems rely upon a mechanism of data collection which is fundamentally flawed. In this system, the system watches all the traffic on the network, and scrutinizes it for patterns of suspicious activity. However there isn't enough information on the wire on which to base conclusions about what is actually happening on networked machines Two classes of attacks (traffic insertion and evasion) which exploit this fundamental problem are exhibited thus showing that centralized traffic analysis systems cannot be fully trusted. A distributed architecture is not prone to these attacks since all the necessary information about the connections can be found on the end devices themselves.
- Performances improvement. For centralized devices to filter traffic, it is necessary to reassemble frames and packets in order to isolate flows that require filtering. As such, overhead is introduced by the controller. On the other hand, in a distributed architecture, the traffic is naturally reassembled. As a consequence, the access control processing introduces much less overhead than in the centralized approach.
- Scalability improvement. The access control processing can be distributed over several devices. As a result, very high rates can be supported, without needing a powerful centralized device.
- Efficiency improvement. As mentioned in section 1 many protocol stacks can be used above the ATM model. Providing access control mechanisms for all these protocols on a single device is not very efficient. In a distributed architecture, access control mechanisms and access control policy can be specific to the protocol stack being used. This results in a less complex and thus more efficient equipment.

A distributed framework has also some disadvantages. It is more difficult to manage. Detecting attacks against several devices requires each device to cooperate with one another, which is not an easy task. The main disadvantage is that every device on the network has to be modified in order

to supply the access control service. Another problem with a distributed architecture is the mean to exchange access control information.
In section 3.2 we show that Management Information Bases (MIBs) provide useful information from an access control point of view. Section 3.3 then shows how this information can be used by a modified management agent in order to supply the access control service. Finally section 3.4 gives some indication to solve the problem of managing the distributed architecture.

3.1 Management information Bases

Using the MIBs for security services provision is not new. The most common use is in the field of the intrusion detection. For example [TIB95] suggests to use information provided by the IETF MIBs in order to detect intrusions in a local area network. [AD97] proposed to introduce the time parameter in MIBs in order to improve intrusion detection techniques. MIBs have also been used to manage access control ([Kar98]) and to supply the access control service ([SKM97]).

In the field of local area network most of management information has been specified by the IETF and the ATM-Forum. The management framework relies on four simple concepts.
The management platform provides an interface between management data and the network manager. It also provides an interface allowing the network manager to get and set management information on remote agents. The management platform is also able to analyse the data received from the agents thanks to specific software. Management information is stored on a local database called Management Information Base (MIB).
The management agent is the other active element of this framework. The agent manages a set of physical or logical objects through their logical representation. This representation is coded using the Structure of Management Information (SMI) and stored in a MIB. Through this representation, the agent is able to configure and supervise physical devices. The agent is also able to send information to the management platform asynchronously when an unusual event occurs on a supervised device.
The management station and the management agent communicate through a protocol called SNMP (Simple Network Management Protocol).

Many MIBs have been specified in order to manage network protocols and applications. In this part we will only consider MIBs used to manage IP or ATM networks. In order to avoid a long list of management objects, only scenarios illustrating the information to be used and the way to use it are given. Interested readers can find a more complete analysis in [PL98]. Those

scenarios present typical rules from the real world. Each scenario applies to a single level in the protocol stack (ATM, Transport, Application).

Example 1; Restricting access to an ATM video server.

The ATM-MIB has been specified by the IETF in 1994 and gives general information about ATM connections. Since this MIB is quite old, a new MIB ([AToM98]) defining additional information is currently being discussed in the AToMib working group at the IETF.
In this scenario, the goal is to prevent external users to access an internal video server while allowing internal users to use it. Thus we have to identify connections between the video client and the video server. These connections are identified by three parameters: the client ATM address, the server ATM address and the VoD application identifier. These parameters can be retrieved from the ATM2-MIB objects ([AToM98]), namely:

- The *atmVclAddrTable* and *atmAddrVclTable* objects provide the source and destination addresses for each connection through the *atmVclAddrAddr* object.
- The VoD application can be uniquely identified by a set of parameters. This set called BHLI (Broadband Higher Layer Information) is transported through the signaling at the connection setup in order to correctly direct the connection on the destination host. The *atmSigDescrParamTable* provides the same set of information for established connections through the *atmSigDescrParamBhliType* and *atmSigDescrParamBhliInfo* objects. The values defining the VoD application are:
  - 0x04 for the *atmSigDescrParamBhliType* object.
  - 0x00A03E00000002 for the *atmSigDescrParamBhliInfo* object.

Example 2; Prohibiting telnet from the internet to internal hosts.

[RFC2012] and [RFC2013] define two MIBs in order to manage TCP and UDP communications. Like in the previous example, connections between the telnet client and the telnet server have to be identified. These connections are described by four parameters: the client and server IP addresses, the client source port and the server destination port. The client source port value may vary in time but the server destination port is fixed to value 23. Those parameters can be retrieved from the following TCP-MIB objects:

- The *tcpConnTable* gives the destination and source addresses and ports through the *tcpConnLocalAddress, tcpConnLocalPort, tcpConnRemAddress* and *tcpConnRemPort* objects.

Example 3; Prohibiting root remote connections.

Information provided by the MIBs at the application level is scarcer but we can still find some valuable information. For example [RFC1414] allows the connections' owner to be identified as follows:

- The *identTable* defined in the RFC1414-MIB gives the user ID through the *identUserid* object.
- The *sysApplElmtRunTable* defined in the sysApplMIB [RFC2287] provides the process and the user associated with an application through the *sysApplElmtRunName, sysApplElmtRunUser* and *sysApplElmtRunIndex* objects.

As such, it is possible to build a matching between a process and an application through the user name and the connection parameters. The connection parameters (source and destination ports) are used to make the matching with the processes. The other connection parameters (source and destination addresses) are used to check which connections are targeted to external hosts.

3.2 Access control enforcement

Our architecture is based on a modified management agent. This agent can be located on a terminal or on an intermediate device as described in figure 1.

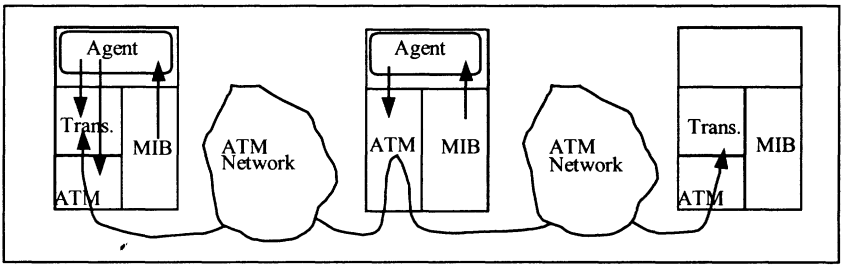

Figure 1. Access control enforced on terminal and intermediate devices

In order to supply the access control service our architecture requires our modified management agent and the MIBs described in section 3.2.

The agent has to be modified in order to introduce the access control operations. It periodically polls the objects described in section 3.2. It then compares the value of these objects with the allowed values. The allowed values describe part of the access control policy to be applied in that agent. Thus the allowed values may vary from one agent to another. In order to determine the prohibited and allowed values, the agent includes a table

containing a set of access control rules. When prohibited values are detected, the agent interacts with the protocol stack in order to stop the prohibited action. Stopping the prohibited action can take several forms:

- At the application level: Interrupting the process executing the prohibited operation. This can be used to provide application level access control.
- At the transport level: Blocking the prohibited communication by releasing the relevant connection. This method can only be used for connection oriented communications (TCP).
- At the ATM level: Blocking the prohibited communication by releasing the connection when dynamic connections are used, or by disconfiguring the relevant interface for permanent connections are employed. This method can be useful when ATM level access control or transport connectionless access control (UDP/ ICMP) is required.

Our architecture has the following advantages:

- The information used to provide the access control service is examined asynchronously by the agent at the application level. Thus no impact on the QoS can be induced.
- The modifications of the system providing the access control service are small. Only the management agent has to be changed.

However selecting the polling rate may not be easy. Indeed a too short interval of polling introduces unuseful overhead for the system whereas a too long interval of polling decreases the security level provided by the agent since some events described by the MIBs will be missed by the agent thus introducing possible security flaws.

3.3 Access control management

As explained in section 3.1 a distributed architecture is quite difficult to manage. To solve this management problem, the three elements depicted in figure 3 are defined.

The Access Control management application.

The Access Control management application is responsible for configuring each agent with the relevant access control rules. In order to perform that task the manager:

- Reads the Access Control Policy. This policy can be stored on a local file or completed through a graphical management interface.

- Selects the rules that should be applied in the agents to be configured. Agents located on terminal devices hold rules concerning their own security, whereas rules located on intermediate network devices can apply to various equipment.
- Codes these rules according to agent MIB syntax.
- Transfers this information to the Access Control Agent.

The access control management application should also retrieve access control results from the agents and should analyse them to detect distributed attacks.

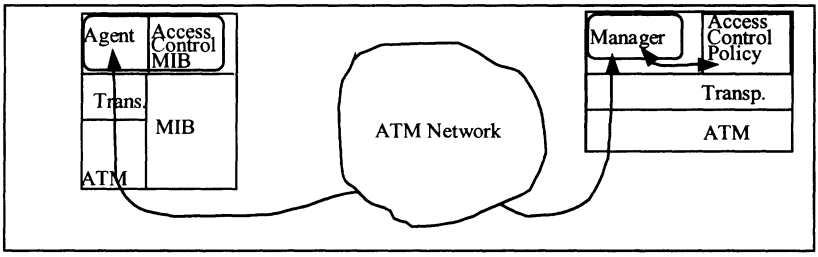

Figure 2. Access control management architecture

The Access Control MIB.

The access control MIB is located on the access control agent device and is remotely managed by the manager through the agent. This MIB includes both access control information, and results from the access control process. Access control information can be classified into three tables, a table at the ATM level, a table at the transport level and a table at the application level. Each table contains access control parameters relative the host holding the table (when this host is an end device) or to many hosts (when the host is an internal network device). Each table includes a set of rules and each rule is described using a generic format.

Access control results are stored into two distinct tables. The first one describes the communications that have been blocked. The second one describes various security alarms.

The MIB size remains relatively small because of the generic format of the rules and because it only contains information about the security of the host.

The Management Protocol

The management protocol used to carry information between the access control manager and the access control agent has to supply several services.

- Integrity. Access control management information should not be modified during their transfer between the manager and the agent.

- Authentication and access control. Only authorized users should be allowed to access the access control information stored on the hosts. The identity of the allowed users must be guaranteed using the authentication service.
- Confidentiality. Only authorized users should be able to read access control information during its transfer over the network.

The SNMPv2 [Sta93] and SNMPv3 [Ba98] protocols seem to be good candidates since they supply all these services.

4. CONCLUSION

As a conclusion, table 2 compares all the competing approaches designed to provide access control on both ATM and IP over ATM networks.

Table 2: Comparison of the different approaches

Property/ Approach	Fire-wall	ATM Forum	Filte-ring Switch	ATM Fire-wall	Dowd & al. Da98	Xu & al. Xs97	Paul & al.
ATM Level Access Control	No	No	Poor	No	Poor	Poor	**Good**
Transport Level Access Control	**Good**	No	No	Med.	Med.	**Good**	**Good**
Application Level Access Control	**Good**	No	No	No	No	Med.	Poor
Label Based Access Control	No	**Good**	No	No	No	No	No
Scalability	Poor	**Good**	**Good**	Med.	Med.	Med.	**Good**
Modification Level	**Poor**	Large	**Poor**	**Poor**	**Poor**	**Poor**	Med.
Impact on the QoS	Large	**No**	**No**	Poor	Poor	Poor	**No**
Security Level	**Good**	**Good**	Poor	Med.	Med.	**Good**	Med.
Management	**Good**	Poor	**Good**	Poor	Poor	**Good**	**Good**
Implement.	**Yes**	No	**Yes**	**Yes**	**Yes**	No	No

As we can see, our approach has the following advantages:

- Good access control at the ATM level.
- Very good scalability thanks to the distributed architecture.

- No impact on the QoS thanks to the asynchronous information retrieval process.
- Good manageability through a management and security integrated approach.

This work could be usefully continued in two directions. The first direction is its implementation since this might give us interesting feedback on the real performance and security level provided by the architecture. The second direction is the extension of our architecture to other types of networks because our architecture can easily be adapted to other kinds of network that are based on a layer 2 switching and that consider QoS as an important constraint.

5. REFERENCES

[7498-2] : ISO 7498-2:1989, Information processing systems -- Open Systems Interconnection -- Basic Reference Model -- Part 2: Security Architecture, ISO, 1989.

[Atom98] : Definitions of Supplemental Managed Objects for ATM Management, Faye Ly, Michael Noto, Andrew Smith, Kaj Tesink, Internet Draft. March 1998.

[Ba98] : Basking in Glory-SNMPv3, Dan Backman, Network Computing, August 1998.

[Da98] : An FPGA-Based Coprocessor for ATM Firewalls, J. McHenry, P. Dowd, F. Pellegrino, T. Carrozzi, W. Cocks, in proceedings of IEEE FCCM'97, April 1997.

[Data97] : Firewalls: Don't Get Burned, David Newman, Helen Holzbaur, and Kathleen Bishop, Data Communications, March 1997.

[JA98] : ATM Net Management: Missing Pieces, Joe Abusamra, Data Communications, May 1998.

[KL98] : Firewall Shootout Test Final Report, Keylabs, Networld+Interop'98, May 1998.

[Kar98] : Integrated Access Control Management, Günter Karjoth. In Lecture Notes in Computer Sciences, 1995.

[PL98] : Où trouver l'information de contrôle d'accès dans les réseaux ATM. Olivier Paul, Maryline Laurent, Technical report. ENST de Bretagne. September 1998.

[PLG98] : Manageable parameters to improve access control in ATM networks, Olivier Paul, Maryline Laurent, Sylvain Gombault, Proceedings of the 5th HP-OVUA Workshop, April 1998.

[PN98] : Insertion, evasion, and denial of service: eluding network intrusion detection, T. Ptacek, T. Newsham, Technical report, Secure Network, January 1998.

[Ran92] : A network firewall, M. Ranum, Proc. World Conference on System Administration and security, 1992.

[RFC2012] : SNMPv2 Management Information Base for the Transmission Control Protocol using SMIv2, RFC 2012, K. McCloghrie, November 1996.

[RFC2013] : SNMPv2 Management Information Base for the User Datagram Protocol using SMIv2, RFC 2013, K. McCloghrie, November 1996.

[RFC2233] : The Interfaces Group MIB using SMIv2, RFC 2233, K. McCloghrie, F. Kastenholz, November 1997.

[RFC2287] : Definitions of System-Level Managed Objects for Applications, RFC 2287, C. Krupczak, J. Saperia, February 1998.

[Schu98] : On the modeling, design and implementation of firewall technology, Christoph Schuba, Ph.D. Thesis, Purdue University, December 1997.

[SEC1.0] : ATM Security Specification Version 1.0, The ATM Forum Technical Committee. July 1998.

[SKM97] : System Security Management via SNMP, F. Stamatelopoulos, G. Koutepas, B. Maglaris, Proceedings of the 4th HPOVUA workshop. April 1997.

[Sta93] SNMP, SNMPv2 and CMIP, The pratical guide to network management Standards. William Stallings. Addison-Wesley. 1993.

[TIB95] : Détection d'intrusions dans les réseaux de communication, K. Tibourtine, Ph.D. Thesis. Université de Paris Sud. February 1995.

[XS97] : Design of a High-Performance ATM Firewall, J. Xu, M. Singhal, Technical report, The Ohio State University, 1997.