

A Full Bandwidth ATM Firewall

Olivier Paul **, Maryline Laurent, and Sylvain Gombault

ENST de Bretagne, 2 rue de la chataigneraie, 35510 Cesson-Sévigné, France.
{paul|mlaurent|sylvain}@rennes.enst-bretagne.fr

Abstract. In this paper we describe an architecture providing an high speed access control service for ATM networks. This architecture is based on two main components. The first one is a signalling analyser which takes the signalling information as an input and produces dynamically the configuration for our second module. This second module called IFT (Internet Fast Translator) is used to analyse the information located in the ATM cells and currently operates at 622 Mb/s. The complete architecture provides the access control at the ATM, IP and transport levels without packet reassembling.

1 Introduction

In the recent past, much attention has been paid developing security services for ATM networks. This resulted in the creation of many working groups within (and outside) the standardisation bodies. One of them is the ATM Forum security Working Group created in 1995, which released its version 1.0 specifications in February 1999. Confidentiality, authentication, integrity and some kind of access control have been considered. Access control as defined by the ISO in [1] is a security service used to protect resources against unauthorised use. 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 signalling 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. This architecture called CARAT does not alter the negotiated QoS. The next section analyses current solutions providing the access control service in the ATM and IP over ATM networks. Section 3 describes the CARAT architecture. This one is based on two main components. The first one is a signalling analyser which takes the signalling information as an input and produce dynamically the configuration of our second module. This second module called IFT (Internet Fast Translator) is used to analyse the information located in the ATM cells at 622 Mb/s. As a conclusion we perform a comparison between our solution and other proposed approaches and we show that CARAT is a good alternative to current solutions.

** This work is funded by DRET and done in collaboration France Telecom - RD

2 Related Work

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 four 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 describe various solutions proposed to improve the "classical" firewall solution. Finally, part four compares existing solutions and highlights the main problems.

2.1 Classical Solution

The first solution [9] 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 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 disassemble the packets is proportional to the packet sizes, which are variable. As a result, the Cell Transit Delay Variation may be different from the CDVT 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 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, desynchronisation between the flows can occur. This kind of solution is reported to have performance problems in a high speed network environment ([3], [5]). The latest tests ([6]) show that this access control solution is unsuccessful at the OC-3 (155 Mb/s) speed. Finally, [17] has shown that load balancing techniques between several firewalls could partially solve this performance problem but at a very high cost since the speedup/cost ratio generated by this technique is far from being linear.

2.2 The access control service as considered by the ATM Forum

The access-control service as defined in the ATM Forum security specifications ([10]) 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 authorisation level per subject are defined. These levels include a hierarchical level (e.g. public, secret, top secret, etc.) and a set of domains modelling the domains associated with the information (e.g. management, research, education, etc.). A subject may access an object if the level of the subject is greater than the level of the object and one of the domains associated with the subject includes one of the domains associated with the object. In the ATM Forum specifications, the sensitivity and authorisation levels are coded according to the NIST [4] specification as a label, which is associated with the data being transmitted. This label may be sent embedded into the signalling, 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 authorisation 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 set-up and does not interfere with the user data. However it suffers from the following drawbacks:

- The network equipment is assumed to manage sensitivity and authorisation 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 [14], Celotek, GTE) 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 [13]) 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 packet 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 being

introduced during cell switching. Storagetek also uses 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 was announced in 1996 but has not been 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 [2] uses an FPGA specialised circuit associated with a modified switch architecture. At the ATM level, the access control at connection establishment time 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 the one 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 signalling (i.e. source and destination addresses) is used.
- Access-control at the application level is not considered.
- Only a small part of the proposal has been implemented.

The second approach [11] is the most complete architecture being currently proposed. This solution provides many improvements in comparison with 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 signalling (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 crossing 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 that of the firewall proxy.

Table 1. Access Control Classes

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

This classification expects the switch to separate traffic with QoS requirements from traffic 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.

This approach is very interesting since it introduces many improvements (traffic classification, LCH) over all the other proposals. 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 cannot be filtered using this technique.
- The LCH technique is useless against information leakage since an internal user can decide to bypass the integrity checks on two end systems on both sides of the firewall.
- This architecture is complex. No implementation has been exhibited.

2.4 Conclusion

As a conclusion, table 2 compares all the competing approaches designed to provide access control on both ATM and IP over ATM networks. The ATM forum proposal has not been included in this comparison since this solution requires deep changes in the existing equipment.

A comparison between the remaining proposals shows that several problems remain unsolved. For example, existing solutions only allow the security officer

Table 2. Comparison of different approaches

Property/Approach	Classical Firewall [9]	Filtering Switch [14]	ATM Firewall [13]	McHenry Xu & al. [2]	Xu & al. [11]	CARAT
ATM level A.C.	No	Poor	No	Poor	Poor	Good
TCP/IP A.C.	Good	No	Average	Average	Average	Average
Application A.C.	Good	No	No	No	Good	No
Impact on the QoS	Large	Low	Low	Low	Low	Low
Manageability	Good	Good	Poor	Poor	Good	Good
Implementation	Yes	Yes	Yes	Partly	No	Yes
Bandwidth (Mb/s)	150	622	155	155	/	622

to filter ATM connections on addresses. One of the goals of CARAT is solve this problem by using an improved signalling analyser which allows the security officer to control almost all the parameters that can be used to describe an ATM connection.

Another point is the balance between performance, impact on the QoS and quality of the access control. Most of the current solutions either provide a good security level while offering poor performance and large QoS modifications or provide a lower security level while offering good performance and small impact on the QoS. The goal of CARAT in this field is to reach a security level similar to the one provided by a traditional stateless packet filter while offering better performance than the existing cell based proposals. Similarly to [11], CARAT can be easily extended to provide application level access control for applications without QoS requirements. As a consequence the security level provided by CARAT is at least as good as the Xu and al. proposal which has not been implemented.

The last point is that current cell based proposals rely on the caching of access control decisions in order to speed-up the cell classification process. This cache is filled through a slow cell classification process that analyses cells that can not be classified through the information located in the cache. However, as demonstrated in [18], these kind of architecture is subject to denial of service attacks because hackers can produce a traffic that will always generate cache misses thus forcing the cell classification process to work at the speed of the slow software classification scheme. These attacks can reduce dramatically the performance of cache based proposals. On the other hand, CARAT succeeds to store the complete access control policy ([15]) by using a policy compression technique and a patented storage method. Consequently CARAT is not subject to similar DoS attacks.

3 Proposed solution

As depicted in figure 1, CARAT is based on two main parts. The first part is dedicated to the ATM signalling analysis. The result of this analysis is then used to build a dynamic configuration which is used by our second part to control

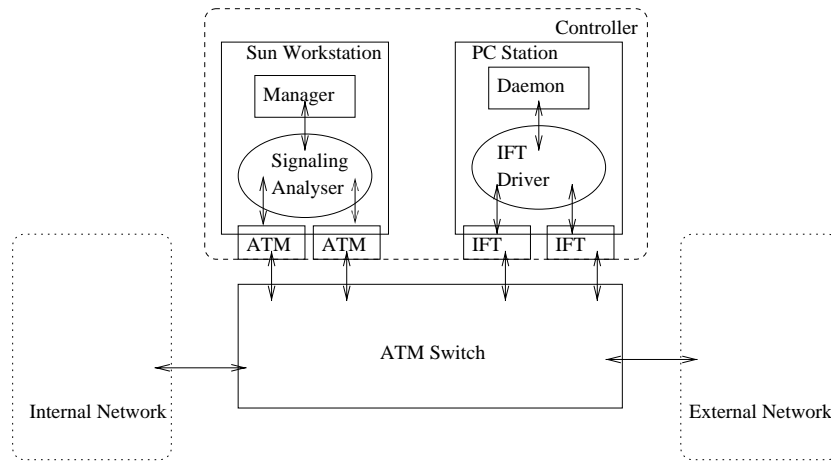


Fig. 1. Controller prototype architecture

the ATM cells. This second part is able to retrieve the ATM, IP and transport level information in order to decide whether a communication has to be allowed or denied. The configuration of the whole controller is made through a single language.

3.1 An Access Control Policy Definition Language

In order to express the access control policy we define an Access Control Policy Description Language (ACPDL). This language is based on draft proposal for a policy description language [7] which had been defined in 1998 by the policy working group at the IETF. In this language an access control policy is described by a set of rules. Each rule consists of a set of conditions and one action which has to be executed when the conditions are met. The following BNF (Backus-Naur Formalism) expression describes the rule syntax. Rule ::= IF $\{$ Conditions $\}$ THEN $\{$ Action $\}$. All the conditions have the same generic structure (BNF notation):

Condition ::= <ACCESS CONTROL PARAMETER> <RELATIONAL OPERATOR>
<VALUE>

Depending on the level in the protocol stack, various access control parameters may be used:

- At the ATM level useful access control parameters have been described in [8], which include the traffic type, connection identifiers, addressing information, QoS descriptors and service identifiers.
- At the transport level most of the included parameters are commonly used to provide access control in firewalls (e.g. addressing information, ports, TCP flags, ICMP codes, etc.).

Actions also have a generic structure (BNF notation).

```
Action ::= <ACTION> <ACTION LEVEL>
```

The action can be to permit or to deny the communication. The level describes the layer (i.e. ATM, Transport) where the action has to be executed.

```
IF ( SRC ADDRESS = 47.0073000000000000000000002402.08002074E457.00 ) AND  
( DST ADDRESS = 47.0073000000000000000000002404.0800200D6AD3.00 ) AND  
( BHLI TYPE = 04 ) AND ( BHLI ID = 00A03E00000002 ) THEN DENY.
```

Fig. 2. Access Control Rule Example

Figure 2 provides an example of how a rule prohibiting connections between two ATM devices for the Video On Demand service can be expressed using the ACPDL. In this example both devices are identified by their ATM addresses and the video service is identified by the Broadband Higher Layer Identifier (BHLI).

3.2 The Manager

The policy defined by the security officer using this language is used to configure the two parts of our access controller. However the policy cannot be used directly by our access control tools. As a result the manager has to translate the access control policy in relevant access control configurations for both our components. The whole translation process is described in figure 3.

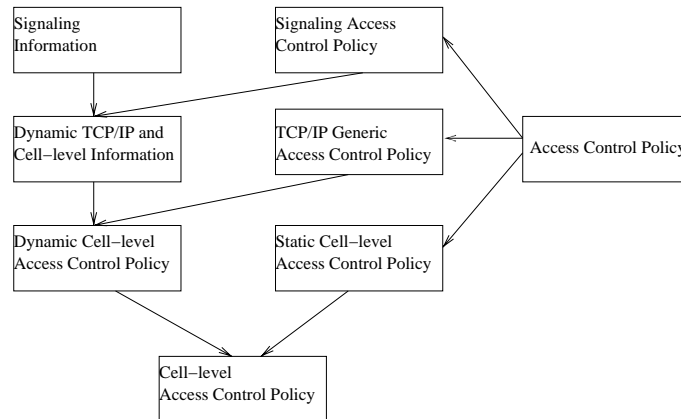


Fig. 3. Access Control Policy Generation Process

This translation process can be divided into two main parts. The first one translates the policy into three static sets of configuration data.

- At the ATM signalling level, this configuration includes a description of the communications that have to be controlled. Each communication is described by a set of Information Elements (IEs) and an action (DENY or ALLOW). This configuration is sent to the signalling analyser.
- At the TCP/IP level, the configuration includes a description of packets that have to be controlled. This part of the policy is generic which means that this configuration is not dedicated to a specific ATM connection.
- At the cell level, the configuration includes a description of the cells that have to be controlled. These cells are divided into a set of fields. The set of values that each field can take is described through a tree. This configuration is directly sent to the IFT modules.

The second part of the configuration process occurs when a connection request is received by the signalling analyser. Once the access control process has been completed, the signalling analyser sends to the manager the pieces of information needed to complete the dynamic configuration of the IFTs. This dynamic configuration process is important since it allows the size of the configuration stored in the IFTs to be reduced in comparison to a static configuration. The size of the configuration is an important issue because the delay introduced by the IFT during the access control process depends on it. The information provided by the signalling analyser includes:

- The Vci and Vpi connection identifiers.
- The source and destination ATM addresses.
- The service descriptor (Classical IP over ATM (CLIP), Native ATM Application). When an additional layer is used above the ATM model, the signalling analyser also provides the encapsulation (with or without SNAP/LLC headers).
- The direction of the communication.

In a CLIP environment, the manager uses the ATM source and destination addresses to find the corresponding IP addresses. This translation is done directly by using a local file describing the matching between ATM and IP addresses. However this process could be improved by taking advantage from an ATM Address Resolution Protocol Servers. The manager then uses the TCP/IP generic access control policy to find a match between the IP addresses and the TCP/IP level access control rules. The subset of matching rules is used along with the other pieces of information (addresses, encapsulation, connection identifiers, direction) to complete the configuration of the IFT cards. They are kept during the connection life. At the connection release, the manager receives a message from the signalling analyser to reconfigure IFTs cards and clean their configuration. The manager then destroys the information associated to the connection.

3.3 The Signalling Analyser

The signalling analyser process relies on two capabilities. The first one is to redirect all the signalling messages coming from our external and internal networks

to a message filter located on a SUN workstation. The second one is the ability to decompose these messages according to the UNI 3.1 specification [12] and to forward or drop them according to the access control policy description provided by the manager.




Fig. 4. ATM Switch Configuration

In order to redirect the signalling, the ATM switch has to be reconfigured to redirect signalling messages to the SUN workstation as described in figure 4. This configuration can be achieved by disabling the signalling protocol on interfaces 1, 2, 3 and 6. A Virtual Path has then to be defined between each pair of interfaces for each possible signalling virtual channel. These virtual channels are identified by the vci connection identifier 5.

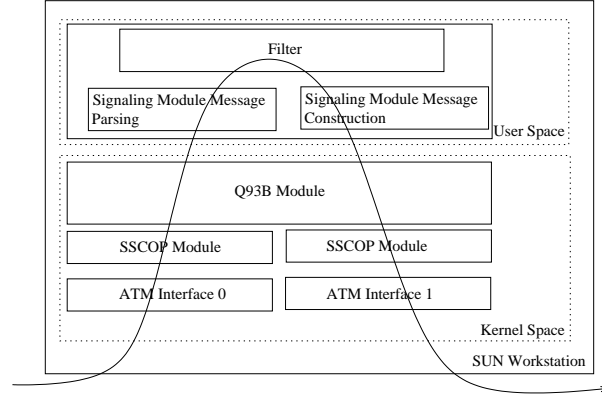


Fig. 5. Signalling Filtering Process

With the previous configuration, signalling messages coming from the external network reach the ATM interface 1 on the SUN workstation, whereas

messages coming from the internal network reach the ATM interface 0. As described by figure 5, all signalling messages are usually multiplexed by the Q93B module whose goal is to establish, manage and release ATM connections. In order to prevent signalling messages from being rejected by the Q93B module, this module has to be modified to forward all the signalling messages to the filter application located in the user space without further analysis. In order to differentiate the filtering process for incoming and outgoing messages, signalling messages are associated with the originating ATM interface. This information is provided to the signalling filter by the Q93B module.

When signalling messages are received by the signalling analyser, these messages are parsed by the message parsing module into Information Elements (IEs) according to the UNI 3.1 specification. IEs are then parsed into basic connection descriptors such as addresses, connection identifiers, call reference, QoS descriptors and service identifiers. The analyser then checks whether the message can be associated with an existing connection through the type of the message and the call reference information. If the connection is new, a connection description structure is constructed. When the connection already exists, the structure is updated according to the new connection description parameters. The resulting set of parameters is associated with the connection state, the originating interface and identified by a connection identifier. The whole structure is then sent to the filter for analysis.

When the filter receives a new connection descriptor, it compares the connection parameters with the set of communications described by the access control policy. If a match is found, the filter applies the action described by the access control policy. When the action is to deny the communication, the filter destroys the corresponding connection description structure. Otherwise the connection identifier is sent to the message construction module. For CONNECT signalling messages, a subset of the connection parameters is sent to the manager as described in the previous section so that the dynamic part of the cell-level access control policy can be generated:

- Vci and Vpi are retrieved from the Connection Identifier IE.
- Source and destination addresses are retrieved from the Called and Calling Party Identifier IEs.
- The service descriptors can be retrieved from the Broadband Higher Layer Identifier (BHLI) and Broadband Lower Layer Identifier (BLLI) IEs.
- The direction is provided by the interface name associated with the connection identifier. For RELEASE COMPLETE messages, the connection identifier is sent to the manager. The communications between the filter and the manager are realised by using a shared memory segment which allows the manager to send the access control policy to the signalling analyser and the filter to send the results from the filtering process to the manager.

When the message construction module receives a connection identifiers from the filter, a new signalling message is constructed according to the information included in the connection description structure. The message is then associated

with the outgoing interface and sent to the Q93B module. When the connection state associated to the connection identifier indicates that a RELEASE COMPLETE message has been sent to release the connection, the construction module frees the resources associated with the corresponding connection.

Another functionality provided by the message composition module is the ability to modify the ATM source address when the communication comes from the internal network to hide the structure of the network. This functionality is provided by changing the source address into the ATM address of the workstation external ATM interface.

The delay introduced by the whole signalling analysis process has few impact on the communication since the standardised signalling timeout values have been greatly oversized (for example a 14 seconds delay is allowed between SETUP and CONNECT messages).

3.4 IFT cards

The Internet Fast Translator (IFT) card ([15], [16]) is a product designed and manufactured by the research branch of our industrial partner. These card have been originally designed to implement an high speed IP packet routing engine. However this card integrates several interesting features that can be used by our ATM firewall:

- It allows the first cell of the AAL5 frame to be analysed and the connection identifiers to be modified according to the analysis.
- The current prototype works at 622Mb/s thanks to a patented cell analysis scheme.
- The delay introduced by the cell analysis process can be bounded and depends on the cell analysis configuration.
- It can be configured dynamically during the cell analysis without interrupting ongoing operations.
- It can be integrated into an off the shelf personal computer using the Solaris x86 operating system.

Table 3 describes the information available for analysis in the first ATM cell with the CLIP and CLIP (without LLC-SNAP encapsulation) protocols. The UD and TD fields indicate the beginning of UDP and TCP data segments. However optional fields such as IP options have not been indicated and may shift TCP or UDP related information in a second ATM cell. Our policy for these kind of cells is currently to drop all the packets including IP options. One may consider this policy to be a severe limitation, however similar access control policies are usually implemented on traditional firewalls since these options are used most of the time by hackers to generate DoS attacks or bypass existing routing mechanisms.

The first part of the cell-level access control process is to direct all the traffic coming from the external and internal networks to the IFT cards. However the configuration of the switch needs to preserve the configuration used by the

Table 3. First ATM Cell Analysis

Byte	1	2	3	4	5	6	7	8	9	10	11	12
CLIP1	ATM Header					AA	AA	03	00	00	00	08
CLIP2	ATM Header					45		Length				
Byte	13	14	15	16	17	18	19	20	21	22	23	24
CLIP1	XX	45		Length							P	
CLIP2			P			Src IP Addr			Dst IP			
Byte	25	26	27	28	29	30	31	32	33	34	35	36
CLIP1	Src IP Addr					Dst IP Addr			Src Port	Dst		
CLIP2	Addr	Src Port	Dst Port							UD		
Byte	37	38	39	40	41	42	43	44	45	46	47	48
CLIP1	Port					UD					D	
CLIP2			D							TD		
Byte	49	50	51	52	53							
CLIP1												
CLIP2												

signalling analysis process. As a result the switch has to be configured to create a virtual channel for each value of vci different from 5 and 31 between each pair of interfaces (1,4 and 5,6). Virtual channels identified by a 31 vci value and later called trash VCs are voluntarily left unconfigured to allow the switch to drop cells belonging to a communication that has to be denied.

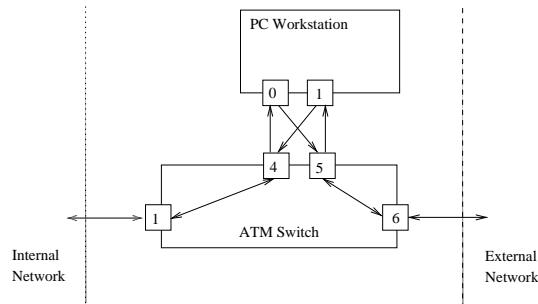


Fig. 6. ATM Switch Configuration

IFT cards allow only unidirectional flows to be controlled. This means that incoming and outgoing flows have to be separated. This operation is particularly simple when dealing with a Mono Mode Fiber physical support since emission and reception fibers are physically separated. Figure 6 shows how emission and reception fibers have to be connected between IFTs and switch ports on both sides of the switch.

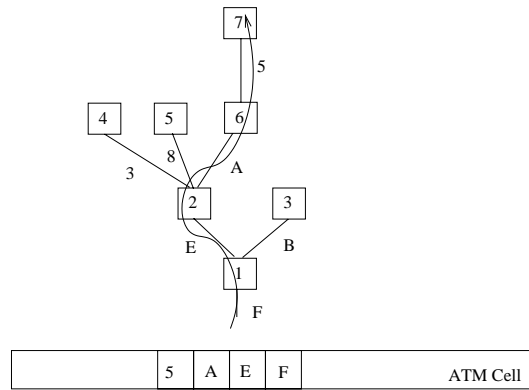


Fig. 7. Analysis Example

The second part is the configuration of the IFTs to provide the access control service. This configuration is done by the manager. IFTs have been originally designed to be managed remotely by concurrent managers. As a result an RPC daemon has been developed to serialise configuration requests to the IFT driver. On the manager side, a library gives access to configuration functions. This library translates the local calls to remote calls on the Solaris PC. The communication between the workstation and the PC are done through a dedicated external Ethernet network.

The IFT analysis process is based on the trie memory ([19]). Trie memory is especially interesting in order to take advantage of the redundancy that can be found in access control policies. The configuration of this process relies on a description of the communications as a set of trees. Each branch within a tree describes a 4 bit value that can be matched during the analysis process. The root of each tree describes a gate by which to begin the tree analysis. An example of analysis is provided in figure 7. Additional information can be provided in a node to allow the analysis to jump from one tree to another or to terminate the analysis and return the connection identifiers values that have to be modified. Configuration functions allow the manager to build, update and remove this set of trees and entries within a given tree while the IFTs are operating. The translation between the information provided by the dynamic cell-level policy generation process and trees can be done as follows:

- Each possible field is coded into a tree. The values described by the security policy are then sliced in 4 bit words and attributed to the nodes of the tree. Range described by multiple conditions on the same field can be described by generating the nodes for each possible word inside the range.
- An AND operation between two conditions on two different fields is coded as a jump from one tree to another.
- The action (DENY or ALLOW) is coded through a special node ending the analysis process and returning the connection identifier that will be at-

tributed to the cells belonging to the corresponding AAL5 frame. A DENY action is coded by directing the frame to a trash virtual channel on the switch. An ALLOW action is coded by leaving the connection identifiers fields unmodified.

Preliminary experimental results ([15]) show that the use of trie memory makes it possible to store the complete TCP/IP level access control policy from a well know French ISP in 2.8M 4 bits words. This is far behind the capacity of the current IFT prototype which is 4M 4 bits words. Additional results ([16]) show that the worst case delay introduced by the access control process during the test was around $1.7\mu s$.

4 Conclusion

In this article, we describe in detail how an ATM firewall can be constructed by using existing components. This ATM firewall has the ability to provide the access control service at the ATM, IP and transport levels at 622Mb/s while maintaining the QoS that has been negotiated. As we can see, our approach has the following advantages:

- Good access control at the ATM level.
- Small impact on the QoS thanks to a bounded delay cell level access control process.
- Improved access control speed at the cell level.
- Is not subject to performance DoS attacks like existing cache based architectures.
- Can easily be adapted to provide the access control service for other kinds of ATM usage (LANE, MPOA, MPLS) with reduced software developments.

Our proposal could be improved in the following directions:

- Application level access control could be easily provided for applications with no quality of service requirements by using the IFTs to direct the flows generated by these applications to a classical firewall where these flows could be analysed in depth. The filter would have to be modified to send a QoS flag to the manager along with the set of parameters currently used to describe the connection. This solution would also provide an answer to the IP options problem.
- The manager and the message filter could be modified to provide filtering capabilities for other kinds of ATM usage such as LAN Emulation, MPOA or MPLS.
- The access control speed could be improved. Our industrial partner is currently working on a new version of the IFT that would be able to handle a several gigabits bandwidth.

5 Acknowledgements

The authors would like to thank all the people that are involved in the CARAT project, namely Yves le Pape Anthony Lucas, Benoit Martin and Jean-Jacques Maret at DGA, Pierre Rolin, Christian Duret, Joel Lattmann and Jean-Louis Simon at CNET for their support and useful comments.

References

1. ISO, ISO 7498-2:1989, Information processing systems – Open Systems Interconnection – Basic Reference Model – Part 2: Security Architecture, 1989.
2. J. McHenry, P. Dowd, F. Pellegrino, T. Carrozzi, W. Cocks, An FPGA-Based Coprocessor for ATM Firewalls, in proceedings of IEEE FCCM'97, April 1997.
3. D. Newman, H. Holzbaaur, and K. Bishop, Firewalls: Don't Get Burned, Data Communications, March 1997.
4. National Institute of Standards and Technology, Standard Security Label for Information Transfer, Federal Information Processing Standards Publication 188, September 1994.
5. J. Abusamra, ATM Net Management: Missing Pieces, Data Communications, May 1998.
6. Keylabs inc., Firewall Shootout Test Final Report, Networld+Interop'98, May 1998.
7. J. Strassner, S. Schleimer, Policy Framework Definition Language, draft-ietf-policy-framework-pfdl-00.txt, Internet Engineering Task Force, November 1998.
8. O. Paul, M. Laurent, S. Gombault, Manageable Parameters to improve Access Control in ATM Networks, in proc. of the 5th HPOVUA Workshop, April 1998.
9. M. Ranum, A network firewall, in proc. of the World Conference on System Administration and Security, 1992.
10. The ATM Forum Technical Committee, ATM Security Specification Version 1.0, February 1999.
11. J. Xu, M. Singhal, Design of a high-performance ATM Firewall, in proc. of the 5th ACM Conference on Computer & Communications Security, 1998.
12. The ATM Forum Technical Committee, ATM User-Network Interface Specification, Version 3.1 (UNI3.1), July 1994.
13. B. Kowalski, Atlas Policy Cache Architecture, White paper, Storagetek Corp., 1997.
14. Cisco Corp., LightStream 1010 Multiservice ATM Switch Overview, 1999.
15. M. Accarion, C. Boscher, C. Duret, J. Lattmann, Extensive packet header lookup at Gb/s speed for an application to IP/ATM multimedia switching router, in proc. of the WTC/ISS2000 Conference, May 2000.
16. Centre National d'Etude des Tlcommunications - France Telecom, IP Fast Translator, FT.BD/CNET/DSE/SDL/226/CD, December 1999.
17. C. Benecke, A parallel Packet Screen for High Speed Networks, in proc. of the 15th Annual Computer Security Applications Conference, December 1999.
18. T.V. Laksham, D. Stiliadis, High-Speed Policy-based Packet Forwarding Using Efficient Multi-Dimensional Range Matching, in proc. of ACM SIGCOMM'98, September 1998.
19. E. Fredkin, Trie Memory, Communications of the ACM, Vol 3, September 1960, pp 490-499.