IMPROVING IP ACCOUNTING FOR SECURE BORDER ROUTERS

Kai Simon and Thomas Schwenkler
IT Security Department
Fraunhofer Institute Experimental Software Engineering
Sauerwiesen 6
67661 Kaiserslautern, Germany
email: {simon, schwenkl}@iese.fraunhofer.de

Stephan Groß
Institute for System Architecture
Technische Universität Dresden
01062 Dresden, Germany
email: st.gross@inf.tu-dresden.de

ABSTRACT

Today, security is a major issue in design and operation of computer networks. To reduce a network's vulnerability, the effort should not be restricted to a firewall as a single point of network traffic control. Only multi-layered security models can effectively protect a network. Thus, a border router with proper access control embodies the outermost security layer. Unfortunately, the rejection of potentially harmful packages can have a negative impact on traffic accounting mechanisms applied on a border router that has been secured this way. In this paper we discuss the state of the art for both access control and traffic accounting techniques. We show that one cannot solely trust current accounting mechanisms because they often suffer from inadequate accuracy and that this problem becomes even worse in secure environments with consequently applied access control. We confirm this statement with an experiment using Cisco's accounting technologies IP Accounting and NetFlow. Going on, we demand a better traffic measurement to meet the security requirements in future network operations and make a first proposal to enhance NetFlow in this direction.

KEY WORDS
Network Management, IP Networks, Network Operations, Applications and Case Studies

1 Introduction

By the end of December 1999 the US National Infrastructure Protection Center announced a security alert comprising intruders installing Distributed Denial of Service (DDoS) tools on various computer systems [1]. From their point of view, "this has been done to create large networks of hosts capable of launching coordinated packet flooding denial of service attacks". Six weeks later, massive DDoS attacks shut down the service of several major Internet sites including Yahoo!, Amazon, and CNN. As a consequence of these early 2000 attacks, network administrators should be aware of the impact DDoS attacks can have. Part of the attack's damage could have been prevented if the network filtering devices in the security perimeter had been configured to correctly deny any spoofed incoming IP packets.

This so-called ingress filtering on the border routers does not only prevent the internal hosts from being misused as a client in a DDoS attack but also blocks malicious IP packets targeted to the router itself. Unfortunately, implementing ingress filtering access control lists (ACLs) also has its disadvantages. Adversities arise if customers want to keep tab on the network traffic they are charged by their Internet Service Provider (ISP). This is in fact of interest as the traffic volume they are charged for by their ISPs is usually an estimated value. Unfortunately, current traffic measurement solutions like Cisco's NetFlow [2]—the traffic measurement solution most widely used by ISPs—suffer from a lack of accuracy. Due to the traffic filtering, this problem becomes even worse. For example, incoming traffic being counted by ISP and then failing the ingress ACL on the customer's border router is not measured on the customer's side. Therefore, we need better and more precise traffic measurement mechanisms that especially take account of security requirements.

The remainder of this paper is structured as follows: After starting with a rough overview on related work in section 2 we discuss in detail ingress filtering as an effective access control technique, its objectives, and which attacks it can prevent (section 3). In section 4 we first of all describe the testbed of our experiment. Thereafter we present several state of the art techniques for IP accounting and discuss their (dis)advantages acquired from the experiment. In section 5 we compare the three accounting techniques and summarize their capabilities. Finally, we discuss some enhancements of Cisco's NetFlow to overcome current drawbacks.

2 Related Work

As the incidents in February 2000 have clearly revealed, the importance of proper ingress filtering on border routers should not be a new thread to network administration. Nevertheless, IP packets with spoofed source addresses are ubiquitous. Even ISPs do not consequently filter obviously spoofed IP packets from their customers. Thus, invalid network traffic is routed to other ISPs and customers although many current threats like DOS attacks could be prevented by consequently filtering this traffic. In the remainder of this section we give an overview on publications concerning best practices for secure router configurations and actual trends in traffic measurement.

On the one hand, valuable information on secure router configuration can be received directly from the router manufacturers. Besides the respective manuals and command references to their various systems, Cisco, for example, provides continuative documentation on specific topics like network security [3, 4]. On the other hand, increasing commercial relevance of the Internet has led to national organizations building up regulations for secure operation of data networks all over the world. Examples are the U.S. National Security Agency (NSA)[1] or the German Federal Office for Security in Information Technology (BSI)[2]. In [5] the NSA gives detailed recommendations on secure configuration of Cisco routers whereas the BSI IT Baseline Protection Manual [6] discusses the deployment and operation of routers as part of a secure IT system on a higher abstraction level. Additionally, several civil organizations concentrate on topics of IT security. Here, the CERT Coordination Center [7] and the SANS Institute [8] should be mentioned. To give just one example for recent books on this topic, Stephan Northcutt et al. have published a comprehensive overview in their book "Inside Network Perimeter Security: The Definitive Guide to Firewalls, Virtual Private Networks (VPNs), Routers, and Intrusion Detection Systems" [9].

Traffic measurement is also a topic in current research activities as operators need to determine the composition of the traffic mix to ensure frictionless operation of their networks. In addition, it is used for provisioning and security reasons. Consequently, there are several publications about the shortcomings of current methods and possible approvements. For example the authors of [10] propose an enhancement of Cisco's NetFlow by adapting the sampling rate to achieve robustness without sacrificing accuracy and by introducing an optional Flow Counting Extension to enable counting of non-TCP flows. [11] investigates NetFlow's capabilities and describes FLOW-REDUCE, a tool to overcome the limitations of Cisco's unspecific and resource dependent definition of NetFlow. The authors come to the conclusion that NetFlow is a valuable tool for network measurement although one looses some accuracy.

3 Securing Routers with Access Control

Access control on a border router generally means deciding whether to deny and drop or to accept and forward an incoming IP packet. Distinguishing characteristics are typically source and destination addresses (for IP-based filtering) or protocol and port (for service-based filtering). Here, ICMP (Internet Control Message Protocol [12]) packets have a special status because they assure a smooth functioning of the Internet. However, they can be misused for DDoS attacks and, thus, should also be filtered.

As already mentioned in the introduction, there are plenty of different kinds of attacks that can be prevented by correctly implemented access control lists. Some of the most famous are the Land Attack, Smurf/Fraggle Attack, and TCP-SYN-Flooding. The first two are based on a forged source IP address and lead to a denial of service of the host specified by this address. In the third one the attacker tries to initiate as many TCP connections as possible without actually completing the TCP three-way-handshake. Thus, the victim's connection queue fills up with unanswered TCP requests preventing it from further processing any valid connection requests. For further details on these and many more common attacks we refer to [13]. Today, several freely available tools like Stacheldraht, Trinoo or TFN2K make rising such attacks as easy as a mouse click so that a protection is not longer a nice-to-have option but a mandatory feature. For a detailed overview on current DDoS tools we refer to [14].

3.1 IP-Based Filtering

IP-based ingress filtering mainly blocks network packets with invalid source addresses. Invalid destination addresses rarely are a major problem, because routing mechanisms in the ISP's backbone typically prevent incorrectly targeted IP packets. For better understanding, the invalid source addresses are classified into three areas: reserved addresses, unallocated addresses, and internal network addresses.

- **Reserved addresses:** About 14% of all theoretically available IP addresses are reserved for special purposes. These include private IP address ranges [15], address ranges reserved for benchmark communication of network devices [16], IP multicast address assignments [17], and address ranges reserved for various other purposes [18], e. g. network addresses denoting *this network* or *this host*.

- **Unallocated address range:** Though the IP address space is strictly limited, a noteworthy amount of subnets (approximately one third) is still unallocated. In the first place, the Internet Assigned Numbers Authority (IANA) is responsible for the allocation of IPv4 addresses.

  IP packets pretending to originate from one of the unallocated address ranges have to be treated as invalid packets. Admittedly, the allocation of IP addresses is a dynamic process so that respective denying ACL clauses have to be revised continuously (see [19]).

- **Address ranges used in the internal network:** Besides the unallocated addresses and address ranges reserved for special purposes there is another group of incoming invalid IP packets. For a border router which is seperating an internal network from the Internet, any incoming packet (from the Internet) pretending to originate from the inside network behind the router is an invalid packet, too. As RFC 2827

[1] http://www.nsa.gov

[2] http://www.bsi.de

[20] clearly describes, this traffic should already be blocked at the time it is injected into the Internet (at the "attackers" ISP), but invalid IP packets of this type are observed repeatedly.

By the way, RFC 2827 should not only be implemented at the ISP's site, but also at the end customer's site. This conspicuously reduces the probability of being misused in a DoS attack.

Usually, IP-based filtering is realized either by assigning appropriate ACLs to the respective interfaces of the router or by applying so-called "null routes"[3] to the respective subnets. Though routing to the virtual Null interface increases performance, it has its drawbacks, too. The routers internal logging mechanism typically requires an ACL whose clauses block a specific part of all incoming IP packets. This can be controlled by the so called hit counter (on Cisco devices) that gives an overview of all filtered packets and the according ACL clause which led to the packet's rejection. Unfortunately, implementing a null route does not allow the logging in this detail.

3.2 Service-Based Filtering

There is no unanimous opinion about the accomplishment of service-based filtering. However, filtering IP traffic on a per port basis can significantly reduce the amount of malicious network traffic. Even though service-based filtering may help to protect against DDoS attacks only slightly, it anyhow increases the routers overall impact on security. Therefore, the NSA [5] gives the advice to block several privileged and non-privileged ports (approximately 40).

- **Filtering privileged ports ($<$ 1024):** Blocking IP packets targeted to specific privileged ports on hosts in the internal network is not a replacement for a firewall at all. Port-based ingress filtering will rather reduce the amount of malicious packets entering the still essential firewall behind the border router.

- **Filtering non-privileged ports ($>$ 1023):** IP packets targeted to non-privileged ports normally describe packets in reply to communication originating from the internal network. However, well-known malicious software oftentimes uses non-privileged ports to connect to hidden services running on internal hosts.

Because the border router itself ideally has to be protected from part of the presented network traffic, service-based filtering is accomplished on an ingress basis at best. To avoid being misused as a client in a DDoS attack, service-based filtering can also be activated for egress traffic to the Internet.

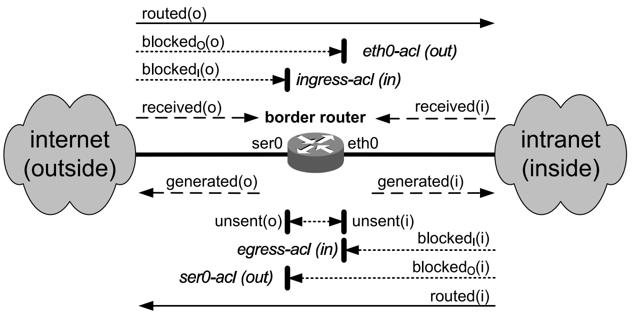

Figure 1. Inbound, outbound, routed, and blocked network flows on a routing device with two interfaces. Letters in brackets indicate the referring network (**i**nside or **o**utside).

3.3 Dealing with ICMP

On the one hand, ICMP is needed to assure smooth functioning of the Internet. Therefore, blocking ICMP packets in general is not advisable. On the other hand, ICMP forms a possible threat to the Internet and adjacent private networks as it is used by several common attacks. A possible way out of this dilemma is restricting ICMP traffic to the inevitable message types [5]. The `ping`-command, for example, uses message types 8 (Echo) and 0 (Echo Reply) [12]. While it is generally helpful to allow ping replies, it is sometimes a good idea to block inbound echo requests to prevent flooding attacks. ICMP messages of type 5 (Redirect) affect the routing table of network interconnecting devices and, thus, should be blocked completely or at least be restricted to the minimum to prevent man-in-the-middle attacks.

4 Traffic Accounting on Cisco Routers

In the previous section we have pointed out the necessity for proper access control on a border router. Beside these security aspects, there are also economical facets to consider. To keep track of the total amount of incoming and outgoing traffic, accounting concentrates on the single point of transit, i.e. the border router.

In the following, we will describe three current accounting techniques for routers. Particularly, we will point out the problems which arise from the combination of security (proper egress and ingress filtering) and economy (correct accounting of network traffic).

In our experimental setup a router with two interfaces was confronted with several types of network traffic, while the device's behaviour was recorded. Figure 1 shows the border router which connects an internal network (on interface eth0) to the Internet (on interface ser0). Letters in brackets indicate the referring network (**i**nside or **o**utside). This is typically the originating network except from traffic generated by the router. In this experiment, there are basically four different types of network traffic:

[3] Null routes can be compared with the /dev/null device in the Unix operating system. All traffic forwarded to this virtual Null interface is dropped.

1. **inbound traffic:** $received(i)$ and $received(o)$
2. **outbound traffic:** $generated(i)$ and $generated(o)$
3. **routed traffic** from and to the internal network: $routed(i)$ and $routed(o)$
4. **blocked traffic**
   (a) packets blocked by an inbound access control list: $blocked_i(i)$ and $blocked_i(o)$
   (b) routed packets blocked by an outbound access control list: $blocked_o(i)$ and $blocked_o(o)$
   (c) generated traffic blocked by an outbound access control list: $unsent(i)$ and $unsent(o)$

Consequentially, each of the interfaces in our experiment is configured with an inbound and an outbound access control list. In contrast, traffic filtering in practice is conducted only in inbound direction as described in section 3.

4.1 IP Accounting

On Cisco routers, the corresponding mechanism is easily activated on a per interface basis [21]. The `ip accounting output-packets` command logs different information about the routed traffic, for example the number of bytes or packets transmitted through the router, in outbound direction. Consequentially, the transit traffic $routed(i)$ on interface ser0 and $routed(o)$ on interface eth0 is included to the accounting data.

Unfortunately, traffic targeted to and received by the router never passes an interface in outbound direction. Thus, the inbound packets $received(o)$ and $received(i)$ are unconsidered regarding the accounting database. Furthermore, traffic generated by the router is not accounted either. Therefore, the outbound network traffic $generated(i)$ and $generated(o)$ also do not add to the accounting database.

Besides the valid packets described so far, IP accounting also provides information that identifies IP traffic failing access control lists. On Cisco routers, this additional feature is activated on a per interface basis by enabling IP accounting of access control list violations with the command `ip accounting access-violations`. This is an important fact because the remaining traffic consists of packets which are blocked by an access control list—either an ingress or an egress ACL.

Our experiment has shown that Cisco has implemented an erroneous treatment of access control lists: Traffic blocked by traditional Numbered ACLs is correctly accounted with the access violation mechanism. In IOS version 12.0 the feature of Named ACLs was introduced, but traffic blocked by these access control lists is not accounted with the access control mechanism. Particularly, with regard to the recommended ingress filtering [20], outbound IP Accounting lacks essential correctness. However, from the total amount of 172.754.545 packets monitored with the aid of Cisco's ACL hit counter [4] 2.651.249 packets have been blocked by an ingress ACL. At the same time, these packets (approximately 1.53% did not appear in the IP Accounting database. Additionally, traffic generated by the router is not subject to any given outbound access control list and will not be blocked at all. Thus, for Cisco routers $unsent(i)$ and $unsent(o)$ do not exist by definition.

4.2 NetFlow

A potential solution to solve the outbound IP Accounting problem without decreasing security is to use Cisco's NetFlow technology [2]. Like IP Accounting, NetFlow can be configured on individual interfaces. Additionally, NetFlow allows for accounting network traffic in inbound direction. All information collected by IP Accounting is also picked up by NetFlow. Furthermore, packets are logged with information about their input and output interface, respectively. Blocked packets are easily recognized by a nonexistent output interface[5]. Unlike IP Accounting, the collected information is not held locally on the routing device but exported to a collector server somewhere in the network. Typically, the collector server does not only log accounting information but also optimizes data storage, for example by aggregation of similar information. To give just one example, CFLOWD [22] is a shareware tool which collects the user datagram protocol (UDP) packets sent by NetFlow. To send NetFlow data to the collector server, NetFlow has to be activated on the appropriate router interfaces with the `ip route-cache flow` command.

In contrast to IP Accounting, NetFlow includes blocked IP packets in the accounting database regardless of the type of ACL. Furthermore, traffic targeted to the router is also acccounted because packets are counted as soon as they enter the device. Nevertheless, NetFlow has its drawbacks, too: Firstly, NetFlow is not fully accurate with regard to the number of connections—or network flows—in that it fails to differentiate between packets with identical IP addresses and ports within a given time slot [11]. Nevertheless, Aggregating IP traffic like this does not affect the correctness of the accounted network traffic. Secondly, packets generated on the router are not accounted either. That is to say, the outbound traffic $generated(i)$ and $generated(o)$ remains unaccounted. Here, only a proper and restrictive configuration can reduce the ignored IP traffic to a minimum. Additionally, $unsent(i)$ and $unsent(o)$ do not exist as described earlier.

4.3 sFlow

Another way of keeping track of network flows is by means of the sFlow Monitoring Technology which is described in

[4] For each packet matching any rule of an ACL the corresponding hit counter of this rule is incremented by one. A concluding rule at the end of an ACL permitting or denying any packet garantees that all packets match exactly one rule.

[5] In fact, the output interface number is set to `Null`

RFC 3176 [23]. Multiple hardware developers and vendors have joined together to form the industry standard sFlow [24] which was designed to monitor traffic high performance broadband networks. Because sFlow captures packets down on the second OSI Layer (Data-Link Layer) [25], there are no restrictions regarding the protocol. Apart from the TCP/IP protocol family, sFlow is able to monitor packets from other protocol suits, like IPX/SPX or AppleTalk. Based on this information, all IP network flows in figure 1 will be monitored.

Unfortunately, sFlow, too, has its disadvantages: The traffic monitoring technique is based on packet sampling [26]. This means, not all packets entering or leaving the device are captured but only a sample set. Depending on the ratio of the set of all packets to the set of sample packets, more or less precise accounting values are collected. Performance side effects can be decreased significantly while accounting only 1 out of N packets on average, but all measurements will be based on statistical estimations. In current implementations of sFlow N has to be set to a minimum value of 2—that is to say, at least half of the traffic is not monitored. Even though accuracy in this case is at its best, sFlow lacks of the neccessary accounting correctness, particularly with respect to low overall network usage.

5 The Need for a Better Accounting

As shown in the previous section there is currently no accurate accounting mechanism for routers. This is especially true for border routers with activated ingress or egress ACL filtering. Table 1 gives a recapitulating comparison of the accounting capabilities for IP Accounting, NetFlow, and sFlow.

Traffic	**IP Accounting**	**NetFlow**	**sFlow**
$received(x)$	no	yes	only statistically
$generated(x)$	no	no	
$routed(x)$	yes	yes	
$blocked_i(x)$	erroneously	yes	
$blocked_o(x)$	erroneously	yes	
$unsent(x)$	no	no	

Table 1. Comparison of accounting capabilities for IP Accounting, NetFlow, and sFlow.

Even though sFlow monitors all data streams in our experimental setup depicted in figure 1, it nevertheless lacks the neccessary correctness because of its statistical packet monitoring technique. For high bandwidth networks sFlow may be the best solution because it has only slight effects on router performance and it produces best results with high load. For low bandwidth network connections as in most small and medium businesses sFlow is far too imprecise.

One of the advantages of IP Accounting is its availability on most of Cisco's routers even in the consumer market. Table 1 on the other hand side makes its main disadvantage quite evident: IP Accounting only keeps track of routed traffic error-free. However, on small routing devices IP Accounting may be the only way to account the network traffic.

NetFlow is probably the most promising accounting technique on Cisco routers. Although NetFlow is only available on midrange to enterprise devices, it shows the highest accuracy from all accounting techniques described above. To further advance NetFlow and to maximize its accuracy, it is only necessary to include traffic generated by the router—regardless of any blocking ACL. Unfortunately, for high bandwidth networks performance still leaves room for improvement in general. For state of the practice high load network connections a software based solution can only be accomplished with a dynamic sampling rate [10] like sFlow uses. To circumvent the problem of accuracy loss, a hardware acceleration or support is indispensable. With the Network Analysis Module (NAM) Cisco presents a dedicated hardware for network traffic analysis including NetFlow technology [27]. This hardware support addresses known performance problems but does not solve the deficiencies in the implementation. There are also some performance problems under high load conditions [28].

6 Conclusion

In this article we have given an overview of the state of the art in both access control and accounting techniques for routers. We identified the lack of accuracy as the major drawback of current traffic measurement techniques. In an experimental setup we have proven that this drawback becomes even worse when considering security measures like access control. In our opinion improved accounting techniques are absolutely necessary to keep up with future requirements of network operations. We strongly believe that Cisco NetFlow has the potential to meet these demands even if it still has some drawbacks and glitches now.

References

[1] TRINOO/Tribal Flood Net/tfn2k. http://www.nipc.gov/warnings/alerts/1999/trinoo.htm, December 1999.

[2] Cisco IOS Software: NetFlow. http://www.cisco.com/warp/public/732/Tech/nmp/netflow/index.shtml.

[3] Michael Wenstrom. *Managing Cisco Network Security.* Cisco Press, 2001.

[4] Richard A. Deal. *Cisco Router Firewall Security. Harden perimeter routers with Cisco firewall functionality and features to ensure network security.* Cisco Press, August 2004.

[5] V. Antoine, R. Bongiorni, A. Borza, P. Bosmajian, D. Duesterhaus, M. Dransfield, B. Eppinger, K. Gallicchio, J. Houser, A. Kim, P. Lee, T. Miller, D. Opitz, F. Richburg, M. Wiacek, M. Wilson, and N. Ziring. Router Security Configuration Guide, Version 1.1b. Technical Report C4-040R-02, System and Network Attack Center (SNAC), National Security Agency (NSA), December 2003.

[6] Bundesamt für Sicherheit in der Informationstechnik (BSI). IT Baseline Protection Manual. http://www.bsi.bund.de/gshb/english/etc/menue.html, October 2003.

[7] CERT Coordination Center. http://www.cert.org.

[8] The SANS (SysAdmin, Audit, Network, Security) Institute. http://www.sans.org.

[9] Stephan Northcutt, Lenny Zeltser, and Scott Winters. *Inside Network Perimeter Security: The Definitive Guide to Firewalls, Virtual Private Networks (VPNs), Routers, and Intrusion Detection Systems*. New Riders, 2002.

[10] Cristian Estan, Ken Keys, David Moore, and George Varghese. Building a Better NetFlow. In *Proceedings of the annual conference of the ACM Special Interest Group on Data Communication (SIGCOMM)*, Portland, Oregon, USA, August 2004. ACM.

[11] Robin Sommer and Anja Feldmann. NetFlow: Information loss or win? In *Proceedings of the second ACM SIGCOMM Workshop on Internet measurment*, pages 173–174, 2002.

[12] J. Postel. RFC792: Internet Control Message Protocol. Technical report, Internet Engineering Task Force (IETF), September 1981.

[13] Adrian Brindley. Denial of service attacks and the emergence of "intrusion prevention systems". SANSGSEC Practical Assignment 1.4b, The SANS (SysAdmin, Audit, Network, Security) Institute, November 2002.

[14] Dave Dittrich. Distributed Denial of Service (DDoS) Attacks/tools. http://staff.washington.edu/dittrich/misc/ddos/. Last visited on August 26, 2004.

[15] Y. Rekhter, B. Moskowitz, D. Karrenberg, G. J. de Groot, and E. Lear. RFC1918: Address Allocation for Private Internets. Technical report, Internet Engineering Task Force (IETF), February 1996.

[16] S. Bradner and J. McQuaid. RFC2544: Benchmarking Methodology for Network Interconnect Devices. Technical report, Internet Engineering Task Force (IETF), March 1999.

[17] Z. Albanna, K. Almeroth, D. Meyer, and M. Schipper. RFC3171: IANA Guidelines for IPv4 Multicast Address Assignments. Technical report, Internet Engineering Task Force (IETF), August 2001.

[18] IANA. RFC3330: Special-Use IPv4 Addresses. Technical report, Internet Engineering Task Force (IETF), September 2002.

[19] The Internet Assigned Numbers Authority. Internet protocol v4 address space. http://www.iana.org/assignments/ipv4-address-space, January 2004.

[20] P. Ferguson and D. Senie. RFC2827: Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing. Technical report, Internet Engineering Task Force (IETF), May 2000.

[21] Configuring IP Services: Configuring IP Accounting. http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr%/fipr_c/ipcprt1/1cfip.htm#1086216.

[22] cflowd: Traffic Flow Analysis Tool. http://www.caida.org/tools/measurement/cflowd/.

[23] P. Phaal, S. Panchen, and N. McKee. InMon Corporation's sFlow: A Method for Monitoring Traffic in Switched and Routed Networks (RFC3176). Technical report, Internet Engineering Task Force (IETF), September 2001.

[24] sFlow Monitoring Technology. http://www.sflow.org.

[25] J. D. Day and H. Zimmermann. The OSI Reference Model. *Proceedings of the IEEE*, 71, 12:1334–1340, 1983.

[26] Jonathan Jedwab. Traffic estimation for the largest sources on a network, using packet sampling with limited storage. HP Labs Technical Reports HPL-92-35, Hewlett-Packard Laboratories, March 1992.

[27] Cisco Network Analysis Module. http://www.cisco.com/go/nam, 2004.

[28] Cisco catalyst 6500/7600 network analysis module deploy guide. http://www.cisco.com/warp/public/cc/pd/ifaa/6000nam/prodlit/nam31_wp.pd%f, 2003.