Adaptive Distributed Traffic Control Service for DDoS Attack Mitigation

Thomas Dübendorfer, Matthias Bossardt, Bernhard Plattner
Computer Engineering and Networks Laboratory (TIK)
Swiss Federal Institute of Technology, ETH Zurich
{*duebendorfer, bossardt, plattner*}*@tik.ee.ethz.ch*

Abstract

Frequency and intensity of Internet attacks are rising with an alarming pace. Several technologies and concepts were proposed for fighting distributed denial of service (DDoS) attacks: traceback, pushback, i3, SOS and Mayday. This paper shows that in the case of DDoS reflector attacks they are either ineffective or even counterproductive. We then propose a novel concept and system that extends the control over network traffic by network users to the Internet using adaptive traffic processing devices. We safely delegate partial network management capabilities from network operators to network users. All network packets with a source or destination address owned by a network user can now also be controlled within the Internet instead of only at the network user's Internet uplink. By limiting the traffic control features and by restricting the realm of control to the "owner" of the traffic, we can rule out misuse of this system. Applications of our system are manifold: prevention of source address spoofing, DDoS attack mitigation, distributed firewall-like filtering, new ways of collecting traffic statistics, traceback, distributed network debugging, support for forensic analyses and many more.

1. Introduction

Recent massive Internet worm outbreaks such as Slammer [14], Blaster [24] or Sasser [25] have shown that a large number of hosts that goes into the millions [12] are patched lazily or are operated by security-unaware users. Such hosts can be compromised within a short time to run arbitrary and potentially malicious attack code transported in a worm or virus or injected through installed backdoors. Distributed denial of service attacks (DDoS) use such poorly secured hosts as attack platform and cause degradation and interruption of Internet services, which result in major financial losses, especially if commercial servers are affected [6]. In recent years, such attacks were repeatedly used for blackmailing companies offering casino, sport bet or advertising distribution [1] services on the Internet. The attacks' structures differ, but all aim at rending a service unavailable for legitimate clients. A large number of malicious hosts sends unsolicited network traffic and hereby exhausts network or host resources.

Keeping a commercial server up and running 24/7 is an asymmetric struggle: while attackers are able to exploit the processing and bandwidth resources and the flexibility of a huge number of compromised hosts to construct new attack tools and variants, operators of Internet servers are left without appropriate means to counteract attacks. Widespread availability of attack tools makes it easy for non-experts (i.e. script kiddies) to carry out large-scale attacks. As a consequence, new attacks appear frequently, while defence strategies lag far behind. We believe that current security technologies and concepts that focus on end system and access networks soon cannot cope anymore with the growing number and the increasing intensity of Internet attacks. We are convinced that large-scale attacks can only be efficiently handled by providing increased security within the network.

In this paper, we present a novel distributed traffic control service, which can help to improve Internet security significantly. At its core is a safe delegation of network management capabilities. It is based on adaptive network traffic processing devices that can be deployed incrementally in the Internet close to routers. As one specific application domain, we show how such a service can fight DDoS reflector attacks, which are tracked down unsatisfactorily and in some cases are handled even counterproductively by existing security mechanisms. Our service can help to stop attack traffic within the network as close to the Internet uplink of an attacker as possible. Our adaptive traffic control service is in no way limited to security related applications. It also enables many other new applications.

The paper is organised as follows: In Section 2, we present DDoS attack scenarios. In Section 3, we analyse various mitigation strategies and show the ineffectiveness of several proposed techniques and systems against DDoS attacks. In Section 4, we propose our new traffic control service based on adaptive devices. The infrastructure we rely on for our service is explained in Section 5. In Section 6, we draw our conclusions and give an outlook on future work.

2. Attack Scenario

2.1. Distributed Denial of Service Attacks

In an Internet DDoS attack, compromised hosts of security unaware users are usually remotely controlled and or-

ganised by an attacker into a so called *amplifying network* of masters and agents. They are then misused to carry out attacks on few or just a single host. Such attacks can also be targeted at core Internet infrastructure components such as routers, central services (e.g. domain name system) or low to medium bandwidth links.

The common aim of DDoS attacks is to deny certain services or resources to prospective users. A large diversity of attack forms exists in the wild. In [20] a taxonomy of denial of service attacks in the context of networks is presented. Technically, a partial or complete denial of service can be caused by *exploiting a system weakness* to make a specific host crash, hang with a blue screen or similar or cause it to reboot, by *exhausting a host's computational, storage, memory or other resources* with the initiation of expensive calculations (e.g. public key de-/encryption) or with triggering resource consuming operations (e.g. complex database queries) or using up all disk space (e.g. by making a host write huge log files). Other ways to cause denial of service are the *misuse of protocols* that make the victim host seem to be temporarily unavailable due to faked protocol signalling (e.g. sending ICMP unreachable messages or TCP reset packets) or the very commonly used technique of *flooding* a target router, host or network link with huge amounts of packets at fast rates such that many packet losses occur and stop legitimate traffic from reaching its destination. The many forms in which DDoS attacks occur in todays Internet make it highly nontrivial to find a panacea for mitigating or stopping such attacks.

DDoS attacks nowadays typically no longer require laborious manual hacking into poorly secured machines over the Internet. Attackers can make use of Internet worms as it was done with MyDoom [27] to do this dirty job. This allows to build up a huge amplifying network of several ten thousand hosts in a short time.

2.2. DDoS Reflector Attacks

A rather new variant of DDoS attacks became known as DDoS "reflector" attack. This attack form is especially difficult to defence against as the victim is flooded with traffic from ordinary Internet servers that were not even compromised.

A selection of DDoS reflector attacks is described in [16]. Any server that supports a protocol which replies with a packet after it has received a request packet can be misused as a reflector without the need for a server compromise. Some prominent examples are web servers, Gnutella servers that even initiate new connections on behalf of other hosts, FTP servers, DNS servers and routers. They return SYN ACKs or RSTs in response to the TCP SYN requests and other TCP packets or ICMP time exceeded or ICMP host unreachable messages upon certain IP packets.

Figure 1 shows that the agents send their packets with the

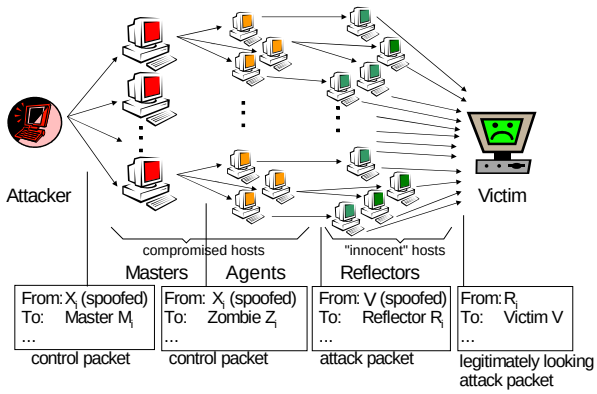

Figure 1: A generic DDoS reflector attack setup

spoofed source address set to the victim's address to "innocent" servers, which act as reflectors. The source addresses of the actual attack packets received by the victim are not spoofed. They belong to legitimate uncompromised servers. Stopping traffic from these sources will also terminate access to Internet services that the victim might rely on.

DDoS attacks organise master and agent hosts in the way of an *amplifying network* as shown in Figure 1. Such a network amplifies the **rate** of packets (a few control packets of the attacker to the masters cause many attack packets to be sent by the agents to the victim), the **size** of packets (if request packet size $<$ reply packet size) and the **difficulty** to trace back an attack to the initiating attacker. We will come back to these core properties when discussing security aspects of our traffic control service.

3. Analysis of Mitigation Mechanisms

This section presents related work that addresses mitigation strategies against DDoS attacks. We distinguish two basic mitigation schemes, *reactive* and *proactive*, which are analysed in more detail and discussed with regard to their mitigation effectiveness and implementation complexity. We show that earlier proposed mitigation schemes fall short of counteracting certain classes of DDoS attacks. In some cases mitigation schemes even amplify the effects of an attack as legitimate servers or complete networks are cut off from the network.

3.1. Reactive Mitigation Strategies

Reactive schemes often proceed in three phases. In the first phase, distributed monitoring components try to detect on-going DDoS attacks. Once an attack is detected, the detector triggers the second phase resulting in the deployment of countermeasures. In the third phase, when the DDoS attack subsides or stops, countermeasures are relieved or removed.

A lot of prior work concentrated on tracing back packets with spoofed source addresses to their actual origin [19, 21, 22, 28]. While this is very valuable in forensics to find the origins and maybe the originator of the attack, it deals with neither detecting attacks nor deploying any dispositions against ongoing attacks. Traceback mechanisms play an important role in other reactive mitigation schemes to determine where countermeasures should be deployed and which filtering rules should be applied. Reactive strategies involving traceback mechanisms, will yield a wrong "attack source" – the reflectors – to be identified and possibly filtered, if DDoS attacks involve reflectors. Relying on traceback mechanisms and subsequently filter outbound traffic of reflectors might block access to important services, because reflectors often host DNS or web servers.

The authors of [11] propose that attacked hosts set filter rules limiting the traffic to specific ports at the last hop IP router. The idea is that the network infrastructure is able to deal with traffic bursts occurring during a DDoS attack, while the attacked host is not able to process incoming traffic. An interesting open question is, whether a host is still able to configure filter rules, if its computing or memory resources are exhausted under a DDoS attack.

In addition to the approach above, the authors of [11] propose a DDoS defence mechanism based on the *Internet Indirection Infrastructure (i3)* [23]. i3 aims to generalise TCP/IP's point-to-point communication to support communication schemes such as multicast, anycast and mobility. i3 is implemented as an overlay that is used to route a client's packets to a *trigger* and from there to the server. Due to performance concerns, i3 would only be used if a server were under attack. Otherwise, communication would be established directly between client and server. To use i3 as a defence mechanisms, IP addresses of the attacked servers are assumed to be hidden from the attackers. It remains unclear how server IP addresses can be hidden under attack, when they are known under normal operation.

Pushback [13] performs monitoring by observing packet drop statistics in individual routers. Once a link becomes overloaded to a certain degree, the pushback logic, which is co-located with routers, classifies dropped packets according to source addresses. The class of source addresses with the highest dropped packet count is then considered to originate from the attacker. Filter rules to rate limit packets from the identified source address(es) are automatically installed on the concerned router. Routers on the path towards the source(s) of attack are informed about the detected attacks and install the same rules. In this way, the attack is *pushed back* and confined.

Pushback assumes that DDoS attacks result in overloaded links. In many cases, however, an attacked server's resources are exhausted before its uplink is overloaded. In particular, this is the case for servers that are hosted in farms, where the communication link is provisioned to feed a large number of servers. Moreover, rate limiting flows based on source addresses is not adequate, if addresses are spoofed. In this case, legitimate sources may experience severe service degradation. The pushback protocol [8] requires all routers to collaborate. If a router on a path between attacker(s) and victim does not speak the protocol, the pushback of filter rules stops to extend further on that particular path.

An inherent problem of reactive mechanisms is that it is very difficult to detect DDoS attacks. None of the discussed systems with the exception of Pushback addresses this issue.

3.2. Proactive Mitigation Strategies

Proactive strategies intend to reduce the possibility of successful DDoS attacks by taking appropriate provisions prior to attacks.

Ingress filtering [7] rejects packets with a spoofed source address at the ingress of a network (e.g. ISP's backbone). As spoofed source addresses are used in several attacks, this approach when put widespread into operation renders many attacks inefficient. Attacks involving reflectors with legitimate source addresses, however, are only affected if ingress routing is applied on paths between agents and reflectors (see Figure 1). Performing ingress filtering puts a management burden on ISPs, because they must keep all filtering rules up to date and defective rules will disgruntle their customers. Even though ingress filtering was proposed in 1998 to prevent attacks, it was only partially applied worldwide as current attacks show.

Secure overlay networks such as SOS [9] and Mayday [4] reduce the risk that a DDoS attack severely affects the communication among members of the overlay network to a minimum. Secure overlay networks require each user of a group wanting to communicate to pre-establish a trust relationship with the other group members. In addition, a user may be required to participate in many groups. As management of many trust relationships is costly and potentially large amounts of traffic is routed among overlay nodes, overlay-based proactive solutions are not adequate for generic communication scenarios (e.g. Yahoo, Google, ebay, etc.), which include millions of communicating hosts. Furthermore, keeping malicious users out of an overlay will be a challenge for a large user base.

3.3. Discussion of Mitigation Effectiveness

We have seen that the described reactive mitigation schemes fail to be effective against DDoS attacks in all three phases: detection, traceback and filtering. What makes DDoS attacks so hard to come by is the fact that attack traffic generally contains spoofed source addresses. In *DDoS reflector* attacks this is even more complex, because the

victim does not receive traffic from the DDoS agents directly, but from legitimate sources without spoofed source addresses. If source spoofing were impossible, reflector attacks could be prevented. Furthermore, complex traceback mechanisms would not be needed, because the originator of malicious packets could be identified by the source address in those packets.

Making source address spoofing impossible requires proactive mechanisms, since measures have to be taken *before* an attack. Proactive approaches may be implemented directly in the IP network or as an overlay network. An advantage of overlay-based solutions is that they can be deployed incrementally, without requiring the cooperation of ISPs. Users only participate in a secure overlay, if the risk of DDoS attacks against them and resulting costs exceed their effort to participate in the overlay.

More effective defence strategies are possible within the IP network. Performing ingress filtering, a single router is capable of blocking traffic from a big number of malicious nodes. In [15] the authors show that ingress filtering is already highly effective against source address spoofing even if only approximately 20% of the autonomous systems have it in place.

As a consequence, the network itself should offer appropriate means for defence. Defence mechanisms must be implemented by the ISPs and BSPs, because they control the traffic entering their network and have access to technology that allows them to deal with large volumes of traffic. However, ISPs currently lack any incentive to implement proactive mechanisms.

4. Distributed Traffic Control by IP Address Owners

Todays' Internet is controlled by network operators, namely Internet and backbone service providers. Network users are restricted to control traffic at their Internet uplink and cannot manage or control network traffic within the Internet.

4.1. Network Traffic Control Service

We propose a novel service that enables network operators to safely delegate specific traffic control to network users. That for, we introduce the new concept of **traffic ownership**. We declare a network packet to be owned by these network users, who are officially registered to hold either the destination or the source IP address or both of that packet. The delegation of certain network management capabilities from network operators to network users is safe in the way that our system assures that a network user can only get control over the IP packets he or she owns. By adding even further restrictions on the traffic control capabilities, as discussed in Section 4.5., we can prevent misuse and malicious interference with other traffic. If the source and destination address of a network packet belong to different parties, a packet can be controlled subsequently by two different parties. Traffic control can be executed by a designated party on behalf of a network address owner.

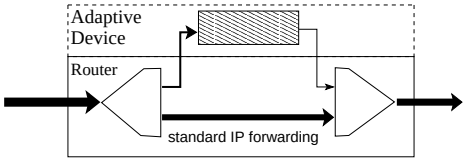

Figure 2: Router extension with adaptive device

Our system consists of remotely programmable network traffic processing devices as shown in Figure 2. The owner of a network address or range gets access to the management of some or all of these devices after having registered for the distributed traffic control service. Traffic entering a router is redirected to a nearby adaptive device only if it carries an IP address as source or destination, which the adaptive device was setup for (see Section 5.). Most traffic will use the direct path through the router.

When the adaptive device processes a network packet, it first executes traffic control on behalf of the owner of the source IP address. Subsequently, it executes traffic control on behalf of the owner of the IP destination address. This is analogous to the high-level communication process of first sending an Internet packet by the source (and hence under its control) and then receiving it by the destination (and consequently under the recipient's control). This control handover is performed at each activated adaptive device on the network path of an IP packet.

4.2. Adaptive Device Functionalities

As the name "adaptive" implies, the functionality of the device can be extended and modified by installing new software (or hardware) modules when new demands arise. Furthermore, upon routing updates, the configuration of modules that depend on the topology can be either automatically adapted or they can be temporarily disabled.

In the context of DDoS attack mitigation, we think of firewall-like services like anti-spoofing filtering, packet dropping, payload deletion, source IP blacklisting or traffic rate limiting. Rules that match traffic by header fields, payload (or payload hashes), or timing characteristics etc. can be installed, configured and activated instantly. During attacks, triggers can automatically activate predefined additional configurations.

To make such a distributed firewall even more powerful, each such device must provide contextual information depending on where it is attached to the network. Additionally, if made available by the network operator, the

router's state and configuration (e.g. static routing information, packet drop rates, congestion parameters, traffic mix, router load etc.) can also be provided. We can e.g. only prevent source spoofing effectively, if the adaptive device is aware of whether it processes transit traffic of autonomous systems or only traffic from customers of a peripheral ISP.

4.3. Attack Prevention and Defence

For stopping a DDoS reflector attack to a specific web site, the owner of that web site's IP address can, by using our proposed traffic control system, almost instantly deploy worldwide ingress filtering rules. These rules will block all traffic that enters the Internet from customers of a peripheral ISP and that carries this web site's spoofed IP address in the packets. Of course, transit traffic, the traffic of the peripheral ISP, where this web site is attached to, and traffic to clients located at peripheral ISPs must not be blocked, as we want the web site's reply packets to reach the legitimate hosts requesting service from it. The more ISPs offer such a distributed traffic control service, the more effective such a defence will be. Our service allows for filtering traffic close to the source of the attack. Hence, we can heavily reduce collateral damage caused by compromised hosts acting as DDoS reflector attack agents. Whereas ingress filtering itself is not new, the way how we allow network users to remotely deploy such filtering for their IP range is novel.

Attacks based on protocol misuse like e.g. sending ICMP unreachable or TCP reset messages to tear down TCP connections can also be filtered out. Without such a distributed traffic control service, worldwide filtering of illegitimate packets is almost impossible due to the many network operators involved that have to be contacted individually for setting up filter rules all over the globe.

4.4. Emerging Applications

Our adaptive devices are in no way limited to firewall-like functionality as new software and hardware modules can be installed when needed. Other services can be based on logging data, collecting traffic statistics, or triggering events. To illustrate the multipurpose nature of our infrastructure, we briefly describe additional use cases of our system.

Traceback: Our system could be used to implement a worldwide packet traceback service such as SPIE [21] by storing a backlog of packet hashes. This would enable support for network forensics by sampling traces of suspicious network activity. Such a service would allow the network user to investigate the origin of spoofed network traffic.

Automated reaction to network anomalies: Our system allows placing triggers in the network in a distributed manner. Triggers generate events if a specific condition is met and thus can be used to signal the activation of a traffic filter function. Automated reaction to network anomalies could be implemented by placing triggers that fire an event if the traffic statistics (e.g. rate of connection attempts from/to a particular server) indicate values exceeding expected boundaries. As a consequence, a rule that rate limits the anomalous traffic could be activated.

Network debugging and optimisation: Our system provides means to collect traffic statistics within the network. Link delays or packet loss on intermediate links could be measured for network debugging purposes. As an example, such information could help providers of content distribution services to optimize their (overlay) network.

4.5. Security Considerations

For the proposed distributed traffic control service to be accepted by ISPs and BSPs, it is vital, that such a device will keep the network manageable by the network operators and that it cannot be misused for an attack itself. This is addressed by the core of our novel approach: We restrict the traffic control for each network address owner to his/her own traffic, i.e. packets to/from owned IP addresses. This allows our service to assure that traffic owned by other parties is not affected. Hence, collateral damage caused by misconfigurations or malicious behaviour of users having access to such devices can be prevented. In addition, ISPs and BSPs do not loose control over their network.

As any misuse of such a novel service must be prevented from the very beginning for gaining acceptance by network operators, we restrict it even further. We do not allow the adaptive device to modify the source and the destination IP address of a packet. Such rerouting could wreak havoc easily (causing routing loops, interference with other routing mechanisms, transparent source spoofing, or "forwarding" of attack traffic). Also the TTL (time to live) field of IP packets is a field we cannot allow to be modified as it aims to set an upper bound of network resources a packet is able to use. Furthermore, we need to prevent that the service can cause amplifying network-like effects as discussed in Section 2. The traffic control must not allow the *packet rate* to increase. In addition, the amount of the network traffic leaving the adaptive device must be equal or less[1] compared to the amount of traffic entering it. I.e. packet size may only stay the same or become smaller. New service modules for the adaptive device must be checked for security compliance before deployment.

Consequently, the danger of delegating partial control of the network from the network operator to the customers is very limited as countermeasures against effects of misconfigurations and misuse were taken into consideration when designing this new service.

The *end-to-end principle* in system design [18] favours a

[1] For e.g. logging, statistics or trigger event services, we will allow a reasonable amount of additional traffic.

network with only a simple but powerful packet forwarding service. However, in [17] the authors argue this principle should be interpreted on a case by case basis. We are convinced that the increase in network security and functionality outweighs the disadvantages of the complexity added to the network by deploying our traffic control service.

4.6. Incentives for Deployment

We see many incentives for ISPs and BSPs to deploy such a distributed traffic control system. It can be offered as a new premium service to customers that e.g. need to protect their commercial Internet servers from attacks, or that want to gather distributed traffic statistics for their sites. Besides using it for new security services, there are many other possible applications as stated in Section 4.4.

Malicious or illegitimate traffic can now be filtered closer to the source. This frees valuable bandwidth resources and makes them available for transporting legitimate traffic. Collateral damage is limited mostly to poorly managed access networks where infected or compromised machines are hooked up to the Internet. This is because attack traffic can be filtered by the new traffic control service at the uplink of such an ISP to a more security-aware ISP or BSP. Other advantages are that ISPs can offer new services and generate additional revenue, whereas customers of an ISP get better service and, e.g. can rapidly reconfigure the adaptive devices in the network to their needs.

5. Infrastructure

This section describes the deployment of our traffic control service and its underlying network infrastructure.

5.1. Network Model

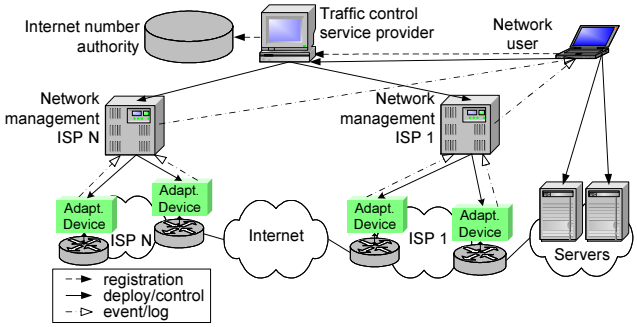

Figure 3: Network model

Our network model shown in Figure 3 distinguishes four different roles: *Internet number authority, Traffic control service provider (TCSP), Internet service provider (ISP)*[2], and *Network user*.

The TCSP manages the new traffic control (TC) service. It sets up contracts with many ISPs that subsequently attach adaptive devices to some or all of their routers and enable their network management system to program and configure these adaptive devices.

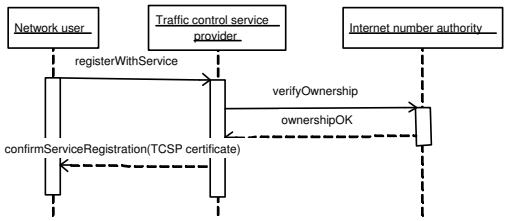

Figure 4: Service registration

A network user must first register with the TCSP before using the traffic control service (Figure 4). The TCSP checks the identity of the network user[3] and verifies the claimed ownership of IP addresses, which she wants to control traffic for. Therefore, the TCSP checks with Internet number authorities[4] if the IP addresses are indeed owned by the service requester. If everything is ok, access to the traffic control service is granted. The binding of a network user to the set of IP addresses owned and the subsequent verification when using the traffic control service (TC service) could be implemented with digital certificates signed by the TCSP. After successfully registering to the basic TC service, a network user may initiate the deployment of a specific service (e.g. ingress filtering), which is implemented on top of the TC service (Figure 5). The network user requests the TCSP to deploy the specific service in the net-

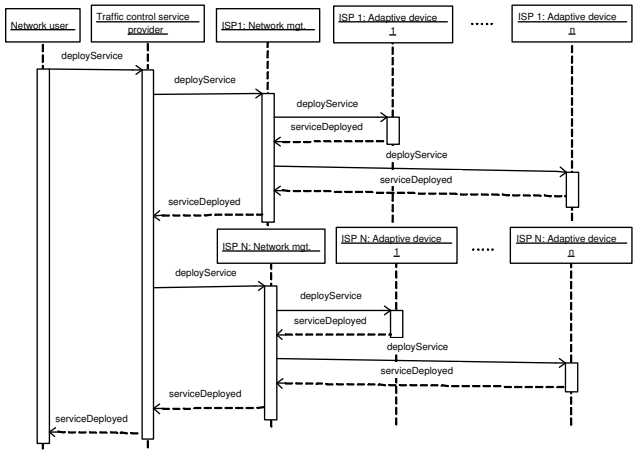

Figure 5: Service deployment

[2]This section subsumes both type of organisations, ISPs and BSPs, under the role ISP.

[3]To check the network user's identity the TCSP performs similar actions as a digital certification authority (CA), e.g. offline verification of an official identity card or online verification of a digital certificate issued by a trusted CA.

[4]Ownership of (ranges of) IP addresses is maintained in databases of organisations such as ARIN, RIPE NCC, etc.

work. The network user may scope the deployment according to different criteria (e.g. “only on border routers of stub networks”). The TCSP maps the request to service components and instructs network management systems of appropriate ISP’s to deploy and configure the service components. ISPs in turn deploy and configure the components on adequate adaptive devices and configure their routers accordingly. Once the service is deployed, a network user may activate, modify specific parameters or read logs of the service. Therefore it sends corresponding requests to the TCSP, which relays them to the appropriate ISP’s network management systems.

Our infrastructure offers an alternative way to activate, modify specific parameters or read logs of the service. A network user may directly interact with the ISPs’ network management systems to *control* the processing of packets that contain an IP address he owns either as source or destination. For an efficient configuration of many adaptive devices, an ISP’s network management system can forward requested configurations to other ISPs’ network management systems upon request of the network user. This approach is particularly useful if the network conditions are such that the TCSP can no longer be reached, e.g. because of an ongoing DDoS attack on the TCSP.

In principle, each ISP could establish a mini-TCSP and offer the traffic control service limited to his network. However, this would make worldwide deployment of traffic control based services cumbersome. The introduction of a TCSP helps to scale the management of our service. Only a single service registration is needed instead of a separate one with each ISP.

The infrastructure can be deployed incrementally. Most traffic control based services will be useful even if not all ISPs offer it. They become more effective when more ISPs join. E.g. anti-spoofing protection and firewall-like services can filter closer to the source and therefore less network resources will be wasted.

5.2. Node Architecture

Our node is based on a legacy Internet router with basic filtering and redirection mechanisms. The router is extended with a programmable traffic processing device (Figure 6). The device can be separate or integrated into future routers.

Network user traffic can be redirected permanently to the traffic processing device. The traffic is processed according to the service requested by the network user. Services are composed of components that are arranged as directed graphs [10, 5]. Each component performs some well defined packet processing. The functionality of components is restricted as described in Section 4.5.

A network user may define two different stages of packet processing. As discussed in Section 4.1., these processing

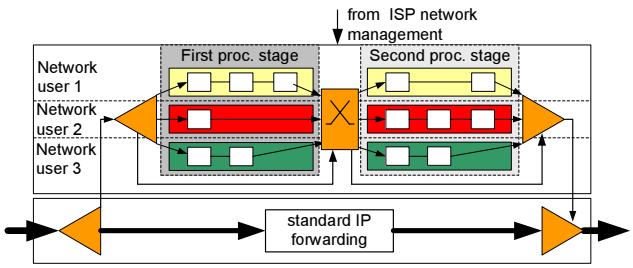

Figure 6: Node architecture

stages determine the processing of packets having the network user’s *source* and *destination* IP address, respectively.

5.3. Scalability

According to the Computer Industry Almanac [3] the Internet had 666 million users in 2002. In August 2004, there were 21.7 mill. hosts [2] connected to the Internet and roughly 10.000 autonomous systems. We target our distributed traffic control service at large organisations that are strongly dependent on Internet communication for their revenue, for vital information exchange or their reputation: Large online shops, large companies, organisations that make heavy use of VLANs to connect their subsidiaries, business to business portals, governmental organisations and others. We do not target our service at home users or small entreprises.

Such a service will be a paid premium service as it requires a new infrastructure to be built and operated. For large organisations that have a clear interest in keeping their Internet services available even when under attack, such a service covers a vital business interest. Therefore, we think that the total number of users of this service will be in the tens of thousands rather than in the millions.

Each user of our traffic control service will use some specific customizable services (e.g. ingress filtering, traceback support) that will result in new rules and modules being activated in our adaptive devices. If the user base grows larger than an adapative device can afford, the ISPs can simply install additional adaptive devices and connect them the to their routers. Redirecting traffic to one of several adaptive devices connected to a router is straight forward. Such a stepwise extension of our infrastructure is affordable as it is only needed when the customer base will increase.

It is important to notice that no additional rules must be installed in our adaptive devices when more users join the Internet or when additional computers are attached. Only if the bandwidth of an ISP implementing our service increases, a faster adaptive device or several devices might be needed. The scaling factors that our service depends on is the total number of autonomous systems (i.e. large ISPs) deploying our service, the resulting number of rules

installed (derived from the tens of thousands of subscribers to our service and their specific traffic control needs) and the bandwidth at which traffic must be filtered and processed in the adaptive devices.

6. Conclusions and Future Work

Our analysis of earlier proposed DDoS attack mitigation systems revealed several inherent weaknesses, which impede those systems to cope with certain classes of DDoS attacks. In particular, such systems may completely cut off legitimate servers or networks under a DDoS reflector attack, thus amplifying the effects of the attack.

We proposed a new distributed traffic control system that enables ISPs to deploy new applications within the network and to safely delegate partial network control to network users. We described how such a system can be used to *prevent* DDoS reflector attacks, which earlier proposed DDoS attack mitigation systems failed to counteract as our analysis showed. Ultimately, our system effectively stops attack traffic close to the source. Herewith, it frees network resources that are nowadays wasted for transporting attack traffic around the globe and that harm not only the target system but also cause collateral damage like network congestion. Many new applications, also not security related ones, will emerge once such a system is available.

Leveraging acceptance by ISPs for such a system will be vital. We think that our traffic control system [26] offers many incentives for ISPs and at the same time a high level of security against misuse, which was a major concern with other approaches in the field of active and programmable networks. In a next step, we build a prototype to get first experiences with such a system.

References

[1] Doubleclick knickt unter DDoS-Attacke ein. `http://www.heise.de/newsticker/meldung/49514`, 7 2004.

[2] RIPE NCC Hostcount. `http://www.ripe.net/info/stats/hostcount/2004/08/index.html`, 8 2004.

[3] Computer Industry Almanac. `http://www.c-i-a.com/pr1202.htm`, December 2002.

[4] D. G. Andersen. Mayday: Distributed Filtering for Internet Services. In *4th USENIX Symposium on Internet Technologies and Systems (USITS 2003)*, Seattle, USA, March 2003.

[5] M. Bossardt, R. Hoog Antink, A. Moser, and B. Plattner. Chameleon: Realizing Automatic Service Composition for Extensible Active Routers. In *Proceedings of the Fifth Annual International Working Conference on Active Networks, IWAN 2003*, LNCS, Kyoto, Japan, December 2004. Springer Verlag.

[6] T. Dübendorfer, A. Wagner, and B. Plattner. An Economic Damage Model for Large-Scale Internet Attacks. In *13th IEEE International Workshops on Enabling Technologies (WET ICE 2004); Enterprise Security*. IEEE, June 2004.

[7] P. Ferguson and D. Senie. RFC 2267: Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing, January 1998.

[8] S. Floyd, S. Bellovin, J. Ioannidis, K. Kompella, R. Mahajan, and V. Paxson. Pushback messages for controlling aggregates in the network.

[9] A. D. Keromytis, V. Misra, and D. Rubenstein. SOS: Secure Overlay Services. Pittsburgh, USA, August 2002.

[10] E. Kohler, R. Morris, B. Chen, J. Jannotti, and M. F. Kaashoek. The Click modular router. *ACM Transactions on Computer Systems*, 18(3):263–297, August 2000.

[11] K. Lakshminarayanan, D. Adkins, D. Perrig, and I. Stoica. Taming IP Packet Flooding Attacks. In *Proceedings of ACM Hot Topics in Networking Workshop (HotNets-II)*, Cambridge, USA, November 2003.

[12] R. Lemos. MSBlast epidemic far larger than believed. `http://news.com.com/MSBlast+epidemic+far+larger+than+believed/2100-7349%_3-5184439.html`, 2004.

[13] R. Mahajan, S. M. Bellovin, S. Floyd, J. Ioannidis, V. Paxson, and S. Shenker. Controlling High Bandwidth Aggregates in the Network. *Computer Communications Review*, 32(3), July 2002.

[14] D. Moore, V. Paxson, S. Savage, C. Shannon, S. Staniford, and N. Weaver. Inside the Slammer Worm. *IEEE Security and Privacy*, 4(1):33–39, July 2003.

[15] K. Park and H. Lee. On the Effectiveness of Route-Based Packet Filtering for Distributed DoS Attack Prevention in Power-Law Internets. In *Proceedings of ACM Sigcomm 2001*, San Diego, USA, August 2001.

[16] V. Paxson. An Analysis of Using Reflectors for Distributed Denial-of-Service Attacks. *ACM Computer Communications Review (CCR)*, 31(3), July 2001.

[17] D. P. Reed, J. H. Saltzer, and D. D. Clark. Comment on Active Networking and End-to-End Arguments. *IEEE Network*, 12(3), May/June 1998.

[18] J. H. Saltzer, D. P. Reed, and D. D. Clark. End-to-End Arguments in System Design. *ACM Transactions in Computer Systems*, 2(4), November 1984.

[19] S. Savage, D. Wetherall, A. Karlin, and T. Anderson. Practical Network Support for IP Traceback. In *Proceedings of ACM Sigcomm 2000*, Stockholm, Sweden, August 2000.

[20] C. Shields. What do we mean by Network Denial of Service? In *Proceedings of the 2002 IEEE Workshop on Information Assurance and Security*, West Point, USA, June 2002.

[21] A. C. Snoeren, C. Partridge, L. A. Sanchez, and C. E. Jones. Hash-Based IP Traceback. In *Proceedings of ACM Sigcomm 2001*, San Diego, USA, August 2001.

[22] D. X. Song and A. Perrig. Advanced and Authenticated Marking Schemes for IP Traceback. In *Proceedings of IEEE Infocom 2001*, Anchorage, USA, April 2001.

[23] I. Stoica, D. Adkins, S. Zhuang, S. Shenker, and S. Surana. Internet Indirection Infrastructure. Pittsburgh, USA, August 2002.

[24] Symantec Security Response. W32.Blaster.Worm. `http://securityresponse.symantec.com/avcenter/venc/data/w32.blaster.worm.html`, 2003.

[25] Symantec Security Response. W32.Sasser.Worm. `http://securityresponse.symantec.com/avcenter/venc/data/w32.sasser.worm.html`, 2004.

[26] Thomas Dübendorfer, Matthias Bossardt. Distributed Internet Traffic Control System, Patent PCT/CH2004/000631 filed, Oct. 2004.

[27] US-CERT. Technical Cyber Security Alert TA04-028A on W32/MyDoom, 2004.

[28] A. Yaar, A. Perrig, and D. Song. Pi: A Path Identification Mechanism to Defend against DDoS Attacks. In *Proceedings of IEEE Symposium on Security and Privacy*, Berkeley, USA, May 2003.