

Enhanced Internet Security by a Distributed Traffic Control Service Based on Traffic Ownership

Matthias Bossardt, Thomas Dübendorfer, Bernhard Plattner

Computer Engineering and Networks Laboratory (TIK)

Swiss Federal Institute of Technology, ETH Zürich

E-mail: {bossardt,duebendorfer,plattner}@tik.ee.ethz.ch

19th June 2005

Abstract

Frequency and intensity of Internet attacks are rising at 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 the novel concept of traffic ownership and describe a 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, service level agreement validation, traceback, distributed network debugging, support for forensic analyses and many more. A use case illustrates how our system enables network users to prevent and react to DDoS attacks.

Keywords: traffic control, network management, network services, mitigation, distributed denial of service attack

1 Introduction

Recent massive Internet worm outbreaks have shown that a large number of hosts that goes into the millions [14] are patched lazily or are operated by security-unaware users. Such hosts can be compromised within a short time and misused 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 [7]. In recent years, such attacks were repeatedly used for blackmailing companies offering casino, sport bet or advertising distribution [10] services on the Internet. The attacks’ structures differ, but all aim at rendering 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 install malicious tools and launch new attack 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 even 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 enhance 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 such as for example new ways of collecting traffic statistics, distributed network debugging and support for forensic analyses.

The paper is organised as follows: In Section 2, we present DDoS attack scenarios. In Section 3, we analyse various mitigation mechanisms 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 traffic processing devices. The infrastructure we rely on for our service is explained in Section 5. Section 6 presents a use case describing in detail how our system can be applied to mitigate DDoS attacks. In Section 7, we draw our conclusions and give an outlook on future work.

2 Attack Scenarios

2.1 Distributed Denial of Service Attacks

In an Internet DDoS attack, compromised hosts of security unaware users are usually remotely controlled and organised by an attacker as a so called *amplifying network* of masters and agents. They are then misused to carry out attacks on a few or just a single host. 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. Technically, a partial or complete denial of service can be caused by *exploiting a system weakness* to make a specific host crash, by *exhausting a host's computational, storage, memory or other resources* or by triggering resource consuming operations. 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 or the very commonly used technique of *flooding* a target router, host or network link with a huge number of packets at fast rates such that many packet losses occur and legitimate traffic is hindered from reaching its destination. The many forms in which DDoS attacks occur in today's

Internet make it highly nontrivial to find a panacea for mitigating or stopping such attacks. Attackers can make use of Internet worms as it was done with e.g. MyDoom [25] for compromising hosts and installing a backdoor. 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 defend against as the victim is flooded with traffic from ordinary Internet servers that were not even compromised.




Figure 1: A generic DDoS reflector attack setup

A selection of DDoS reflector attacks is described in [19]. 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 TCP SYN requests and other TCP packets, or ICMP “time exceeded” or ICMP “host unreachable” messages as a reaction to certain IP packets.

Figure 1 shows that the agents send their packets with the spoofed source address set to the victim’s address (V) to “innocent” servers with IP addresses R_i . These servers act as reflectors. The source addresses of the actual attack packets (R_i) 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 that aims at locating the attack sources. In the third phase, countermeasures are deployed to mitigate the attacks.

A lot of prior work concentrated on tracing back packets with spoofed source addresses to their actual origin [20, 22, 23, 26]. While this is very valuable in forensics to find the origins 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. If DDoS attacks involve reflectors, traceback mechanisms will yield a wrong “attack source” – the reflectors – to be identified and possibly filtered. Hence, access to important services might be blocked because DNS or web servers are often abused as reflectors.

The authors of [13] propose that attacked hosts set filter rules to limit the incoming traffic at the last hop IP router. The network infrastructure is assumed to be able to deal with traffic bursts, while the attacked host is not able to process incoming traffic. An open question is, whether a host is still able to configure filter rules, if its resources are exhausted under a DDoS attack. Moreover, the authors of [13] propose a DDoS defence mechanism based on the *Internet Indirection Infrastructure (i3)* [24]. *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 mechanism, 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 [15] 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 routers on the path towards the source(s) of attack. 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. 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 [9] requires *all* routers on the attack paths to collaborate. 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 [8] rejects packets with a spoofed source address at the ingress of a network (e.g. to the Internet service provider's backbone network). As spoofed source addresses are used in several attacks, this approach when put into widespread 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 already proposed in 1998 to prevent attacks, it was only partially applied worldwide as current attacks show.

Secure overlay networks such as SOS [12] and Mayday [1] require each communicating user of a group to pre-establish a trust relationship with the other group members. Hence, 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 communication with popular web servers (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: attack detection, attack location (traceback) and attack mitigation (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.

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 [18] the authors show that ingress filtering combined with distributed packet filtering can already be 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 Internet service providers (ISP) and backbone service providers (BSP) 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 mechanisms that protect network users from attacks.

4 Distributed Traffic Control by IP Address Owners

Today's 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** [6]. 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 4.4, 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 traffic processing device (TPD)

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 traffic processing device (TPD) only if it carries an IP address as source or destination, which the TPD was setup for (see Section 5). Most traffic will use the direct "IP fast path" through the router.

When the TPD 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 hand-over is performed at each activated TPD on the network path of an IP packet.

4.2 Adaptive Traffic Processing Device Functionality

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 such as 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 traffic processing device is aware of whether it processes transit traffic of autonomous systems or only traffic from customers of a peripheral ISP ("stub network").

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 attack agents¹. Whereas ingress filtering itself is not new, the way we allow network users to remotely deploy such filtering for their IP range is novel.

Attacks based on protocol misuse, such as sending ICMP "host unreachable" or TCP RST messages to tear down TCP connections can also be filtered out. Without a distributed traffic control service as provided by our system, 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 Security Considerations

For the proposed distributed traffic control service to be accepted by network operators (namely 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 approach, the novel concept of traffic ownership: We restrict the traffic control for each network address owner to his/her own traffic, i.e. packets to and from owned IP addresses.

However, while this restriction is the most important one to ensure security and acceptance of our system, it alone does not yield a secure system. To prevent collateral damage caused by misconfigurations or malicious behaviour of users having access to such devices, we need to restrict traffic control even further. We do not allow the adaptive traffic processing device to modify the source and the destination IP address of a packet. Such rerouting could wreak havoc easily by causing routing loops or

¹Our ISP based ingress filtering strategy cannot prevent an IP spoofing attack originating from the same enterprise network domain as the attacked service resides in. However, enterprise firewalls can provide solutions for this case.

interference with other routing mechanisms, by allowing transparent source spoofing via changing the source address of an attack packet to a random one on a TPD close to the victim host, or by “forwarding” attack traffic by setting a new destination address in attack packets. Also the TTL (time to live) field of IP packets is a field that 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 less² compared to the amount of traffic entering it. I.e. packet size may only stay the same or become smaller. To ensure these restrictions, new service modules for an 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 the distributed traffic control service. In fact, the above mentioned restrictions assure that traffic owned by other parties is not negatively affected by dynamically deployed service functionality and that network operators do not lose control over their network albeit their customers’ extended traffic control realm.

4.5 Incentives for Deployment

As a reaction to a DDoS attack on its web servers in February 2005, a German publisher posted a 10’000 EUR reward for hints about the identity of the originator of the attack [11]. Although this publisher’s business model is based on selling hard copies of journals and not on the news they provide (for free) on the web, the reward shows how much companies value an uninterrupted web presence. For other companies whose business model depends on the web presence, our traffic control service is even more valuable. Hence, 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 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 such as logging data, collecting traffic statistics, or validating service level agreements.

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.

²For e.g. logging, statistics or trigger event services, we will allow a reasonable amount of additional traffic.

5.1 Network Model

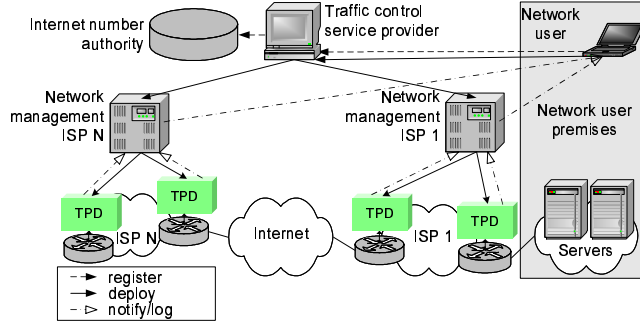


Figure 3: Network model

Our network model shown in Figure 3 distinguishes four different roles: *Internet number authority*, *traffic control service provider (TCSP)*, *Internet service provider (ISP)*, and *network user*. This section subsumes both types of organisations, ISPs and BSPs, under the role ISP. The TCSP manages the new traffic control (TC) service. It sets up contracts with many ISPs that subsequently attach adaptive traffic processing devices to some or all of their routers and enable their network management system to program and configure these devices.

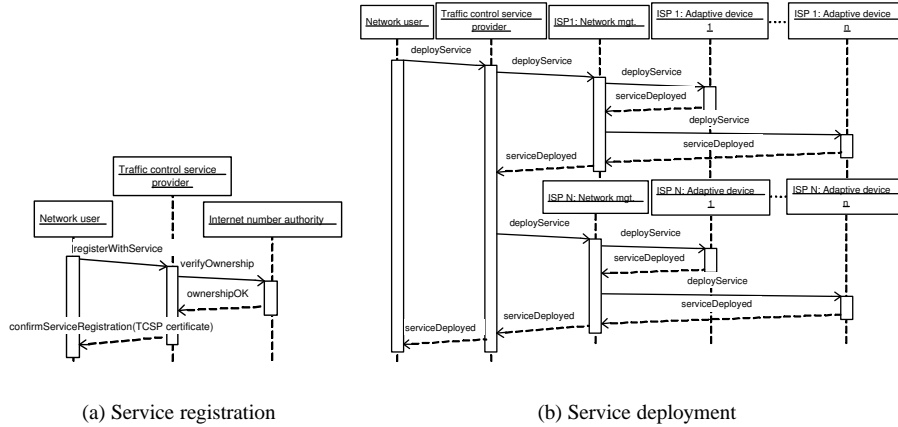


Figure 4: Message sequence diagrams

A network user must first register with the TCSP before using the traffic control service (Figure 4(a)). After checking the identity of the network user³, the TCSP verifies the claimed ownership of IP addresses, which the network user wants to control traffic for. Therefore, it checks with Internet number authorities⁴ whether the IP addresses are

³To 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.

⁴Ownership of IP addresses is maintained in databases of organisations such as ARIN, RIPE NCC, etc.

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 successful registration for 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 4(b)). The network user requests the TCSP to deploy the specific service in the network. The network user may scope the deployment according to different criteria (e.g. “only on incoming links of border routers placed in 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 *traffic processing devices* (TPD) (see Figure 3) 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 she sends corresponding requests to the TCSP, which relays them to the appropriate network management systems of the concerned ISPs.

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 she owns either as source or destination. For an efficient configuration of many traffic processing devices and routers, 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. For example, anti-spoofing protection and firewall-like services can filter closer to the source and therefore less network resources will be wasted.

5.2 Traffic Control Unit Architecture

A Traffic Control Unit (TCU) is the combination of one or more traffic processing devices (TPD) attached to a legacy router. More specifically, a TCU is defined by a router interface and the TPDs that traffic of the interface can be redirected to. Figure 5 shows a simple TCU consisting of one TPD connected to a router. The devices can be physically separate, even located at different sites, or integrated into future routers.




Figure 5: Traffic Control Unit architecture

Traffic of specific network users 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 [4]. Each component performs some well defined packet processing. The functionality of the components is restricted as described in Section 4.4 to prevent misuse of the system. A network user may define two different stages of packet processing. As discussed in Section 4.1, these processing stages determine the processing of packets having the network user's *source* and *destination* IP address, respectively.

5.3 Scalability Issues

The scaling factors that our distributed traffic control service depends on are 1) the number of service subscribers (i.e. network users), 2) the total number of ISPs deploying our service, 3) the number of rules installed per network user, and 4) the bandwidth of network links. This section discusses the parameters that are influenced by these scaling factors.

5.3.1 Service logic and state per TPD

Each user of our traffic control service will use some specific customisable service modules (e.g. ingress filtering, traceback support) that will result in service logic and per-user state being activated in our adaptive devices.

The traffic control service will be a charged premium service because it requires additional infrastructure to be installed and operated. 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 enterprises.

In July 2004, there were 285 mill. hosts [17] connected to the Internet, roughly 26 mill. thereof were active web servers (February 2005) [16] that hosted 59 mill. web sites. For active SMTP, VPN and other common servers no reliable numbers seem to be available. We estimate the number of potential traffic control service subscribers that meet the discussed criteria to about 1% of the active web servers, which would be about 260'000. Each subscriber is expected to request a few customised services to be executed on his behalf, which results in service logic to be deployed on a number of TPDs and in per-subscriber configuration information and potentially service state to be kept on the TPDs.

Although several subscribers may request the same service to be executed on a TPD, the service logic needs to be deployed only once because the logic can be shared. However, configuration and state information, if applicable, must be kept on a per-user basis. The number of TPDs involved in a service of a particular user is service-specific. That is, some services, (e.g. service level agreement validation) involve only a few TPDs, whereas others (e.g. ingress filtering) improve in quality the more TPDs are included.

If we assume that for each subscriber/user 10 services are run, that a service runs on 1% of the TPDs on average and each service includes per-user state, we find that each TPD must keep the configuration information and state for 2'600 users. Accounting for 1 kByte of configuration and state information per user and service, this results in 26 MByte of memory needed per TPD, which is a rather modest requirement.

5.3.2 Signalling Effort

The number of ISPs offering distributed traffic control influences the number of service deployment messages that need to be sent by the TCSP to the ISP management stations each time a network user requests the provisioning of a new service. According to the CIDR report [3] of February 2005 there were 18'918 autonomous systems, out of which 7'732 announced only one network prefix. Even if we assume that each AS corresponds to one ISP and that all ISPs offer our distributed traffic control service, signalling overhead due to the secure distribution of the small service deployment messages by the TCSP to a few thousand ISPs is not a bottleneck. Furthermore, proactive services like ingress filtering do usually not require instant deployment.

5.3.3 Traffic Processing Capacity

Due to high performance demands, a hardware based solution for our traffic processing devices is favorable. Research prototypes of FPGA based devices exist that can concurrently filter 8 mill. flows [21] on a 2.5 Gbps (OC-48) link. According to [2], faster FPGAs allow achieving advanced packet filtering at 10 Gbps (OC-192). Highly flexible solutions using commercial network processors are an alternative.

6 Use Case

In this Section, we give a sample scenario for a DDoS mitigation service, describe the deployment steps and explain the service mapping process using XML documents.

6.1 DDoS Mitigation Service

We assume that a (fictional) U.S. based company named Xorus operates the large online shop "Xorus perfumes". All its revenues stem from online sales and therefore Xorus is interested to have their shop available 24/7 without interruption. After some recent DDoS attacks on some other large online shops, Xorus registers for the traffic control service at the TCSP. It installs some proactive traffic control services that filter out UDP based attacks towards their website's IP address right at the source as there is no need for UDP packets to be received and processed by the Xorus online shop. In addition, Xorus requests a rate limiting service for TCP SYN requests, which will automatically become active for non-US IP addresses after a trigger condition on the rate of TCP SYN packets yields true on certain links connecting the rest of the world to the U.S. Internet backbone. This allows to provide a high service quality for U.S. customers even while under attack from abroad. However, early on a Monday, the Xorus shop is heavily flooded by TCP RST packets from a DDoS reflector attack that misuses tens of thousands of compromised computers as DDoS agents and U.S. based web servers as reflectors. Instantly, Xorus reacts and deploys an ingress filtering traffic control service on its shop's IP addresses that blocks all IP packets sent over an Internet uplink with the shop's IP addresses from a location different than Xorus' own uplink. This immediately stops the attack. In addition, Xorus deploys a logging service at the uplinks to trace attack traffic back to the real origin of the reflector attack in the hope to find the attacking agents and possibly also the attacker.

6.2 Service deployment

For our further explanations, we focus on the ingress filtering service, which is deployed to mitigate the attack based on reflected TCP RST messages described in the scenario above. As our network user (Xorus) has already registered with the TCSP, it can directly initiate the deployment of the ingress filtering service. Therefore, Xorus selects the service from the TCSP's web site together with service-specific parameters. Client authentication is used to make sure that only a legitimate network user can initiate service deployment. A service request is generated and used as input for the service mapping process in the TCSP layer. Figure 6 shows the details of the deployment process, which is subdivided into *TCSP*, *ISP*, *TCU* and *Device* layer. The complete deployment process is carried out at the management stations of TCSP and ISP. For each service a layer offers, a *service descriptor* specifies the following:

- The mapping of the service to sub-services offered by the layer below.
- The set of mandatory and optional parameters, their default values and their mapping to parameters for sub-services.
- Restrictions that direct the placement of service logic.

For each layer a database contains *context information* about the infrastructure relevant to that layer⁵. Information at the TCSP layer includes the identities of contracted ISPs and properties of their networks, e.g. whether they transport transit traffic or provide a stub network. At the ISP layer relevant information includes the location of the TCUs, e.g. whether it is located at the border of a network or in the core network. At the TCU layer details about the pairing of TPDs and routers are kept as context information. Finally, at the Device layer information about the make and version of TPDs and routers and their configuration interfaces must be kept. Additionally context information can contain dynamic state information about managed objects and deployed services.

Deployment logic on each layer *maps* the service request from the layer above to services provided by the layer below (right column of Figure 6) based on information provided by the service descriptors (left column of Figure 6). Taking into account restrictions specified in the service descriptor and context information from the databases, sub-services are placed on the managed objects of the corresponding lower layer (ISPs, TCUs, TPDs and routers, respectively). The middle column of Figure 6 describes the mapping process at the different layers as carried out for the ingress filtering service. The deployment process ends with the configuration of the devices that were previously selected to run part of the service logic.

7 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 based on the concept of traffic ownership that enables ISPs to deploy new applications within the network and to

⁵These logical databases may be merged into two physical databases located at the TCSP and ISP management station.

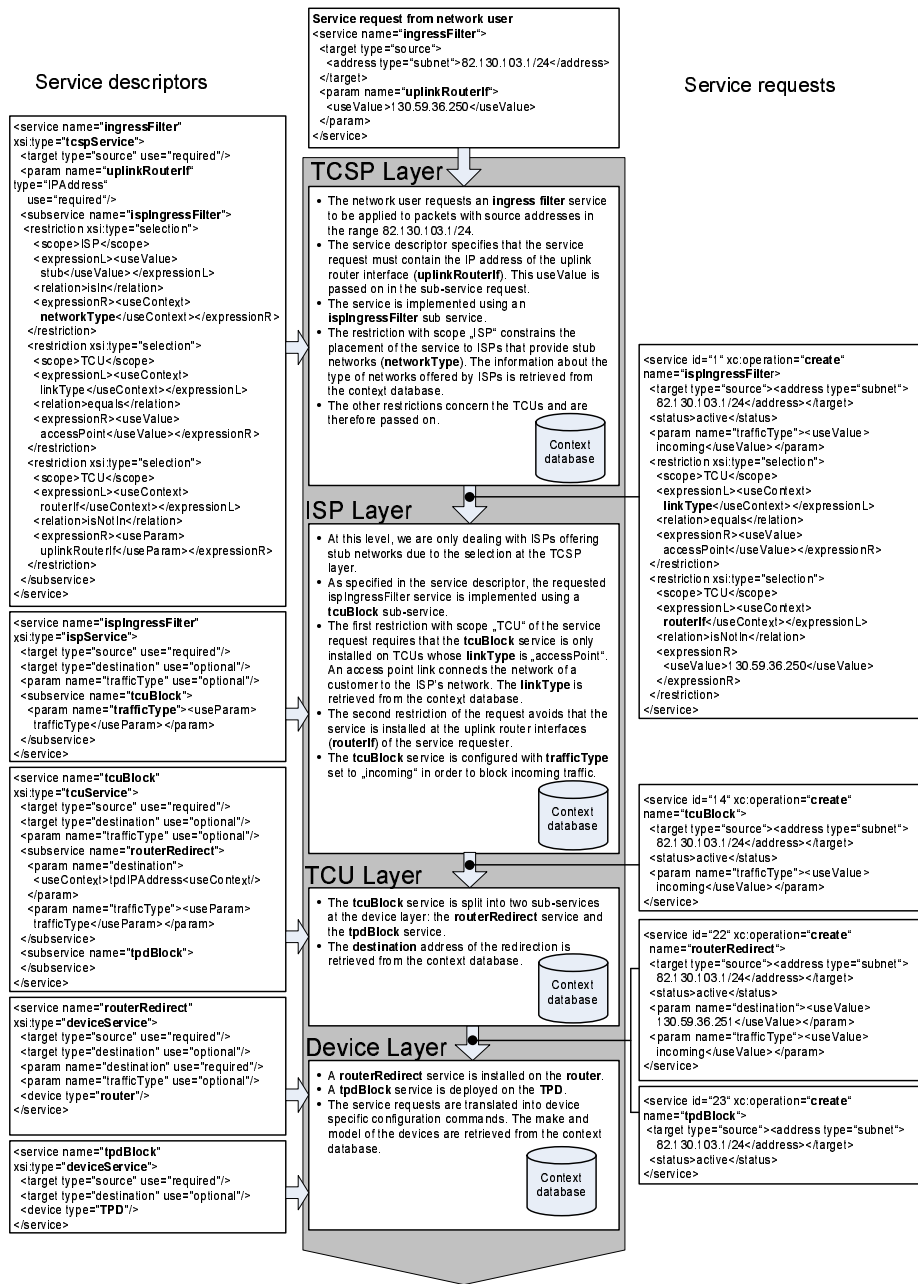


Figure 6: Service mapping

safely delegate partial network control to network users [5]. 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, are expected to 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 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. We are currently prototyping the distributed traffic control system including the “DDoS mitigation” use case. The prototype will allow us to get first experiences with such a system and to demonstrate its use and effectiveness.

References

- [1] D. G. Andersen. Mayday: Distributed Filtering for Internet Services. In *4th USENIX Symposium on Internet Technologies and Systems (USITS) 2003*, Seattle, USA, March 2003.
- [2] M. Attig and J. W. Lockwood. A Framework for Rule Processing in Reconfigurable Network Systems. In *Proceedings of IEEE Symposium on Field-Programmable Custom Computing Machines (FCCM)*, Napa, USA, April 2005.
- [3] T. Bates, P. Smith, and G. Huston. CIDR REPORT for 22 Feb 2005. <http://www.cidr-report.org/>, February 2005.
- [4] 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 2003. Springer Verlag.
- [5] T. Dübendorfer and M. Bossardt. Distributed Internet Traffic Control System, Patent PC-T/CH2004/000631, 2004.
- [6] T. Dübendorfer, M. Bossardt, and B. Plattner. Adaptive Distributed Traffic Control Service for DDoS Attack Mitigation. In *IEEE Proceedings of IPDPS 2005, International Workshop on Security in Systems and Networks SSN*, Denver, USA, 2005.
- [7] 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.
- [8] P. Ferguson and D. Senie. RFC 2267: Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing, January 1998.
- [9] S. Floyd, S. Bellovin, J. Ioannidis, K. Kompella, R. Mahajan, and V. Paxson. Pushback Messages for Controlling Aggregates in the Network. Internet Draft, work in progress.
- [10] Heise Online. Doubleclick knickt unter DDoS-Attacke ein. <http://www.heise.de/newsticker/meldung/49514,72004>.
- [11] Heise Online. Gezielte Denial-of-Service-Attacke gegen heise online [Update]. <http://www.heise.de/newsticker/meldung/55800,February2005>.
- [12] A. D. Keromytis, V. Misra, and D. Rubenstein. SOS: Secure Overlay Services. In *Proceedings of ACM Sigcomm 2002*, Pittsburgh, USA, August 2002.
- [13] 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.
- [14] R. Lemos. MSBlast epidemic far larger than believed. http://news.com.com/MSBlast+epidemic+far+larger+than+believed/2100-7349_3-5184439.html, 2004.
- [15] 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.
- [16] Netcraft. Web Server Survey. http://news.netcraft.com/archives/2005/02/01/february_2005_web_server_survey_finds_40_million_sites_on_apache.html, February 2005.
- [17] Network Wizards. Internet Domain Survey. <http://www.isc.org/>, July 2005.
- [18] 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.

- [19] V. Paxson. An Analysis of Using Reflectors for Distributed Denial-of-Service Attacks. *ACM Computer Communications Review (CCR)*, 31(3), July 2001.
- [20] 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.
- [21] D. V. Schuehler and J. W. Lockwood. A Modular System for FPGA-based TCP Flow Processing in High-Speed Networks. In *Proceedings of the 14th International Conference on Field Programmable Logic and Applications (FPL)*, Antwerp, Belgium. Springer LNCS 3203.
- [22] 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.
- [23] D. X. Song and A. Perrig. Advanced and Authenticated Marking Schemes for IP Traceback. In *Proceedings of IEEE Infocom 2001*, Anchorage, USA, April 2001.
- [24] I. Stoica, D. Adkins, S. Zhuang, S. Shenker, and S. Surana. Internet Indirection Infrastructure. In *Proceedings of ACM Sigcomm 2002*, Pittsburgh, USA, August 2002.
- [25] US-CERT. Technical Cyber Security Alert TA04-028A on W32/MyDoom, 2004.
- [26] 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.