

Stealth DoS

Marcel Waldvogel Michael Muncan Mahak Patidar*

Distributed Systems Laboratory, Department of Computer and Information Science

University of Konstanz, 78457 Konstanz, Germany

<firstname>.<lastname>@uni-konstanz.de

Abstract—Users and providers increasingly disagree on what Denial of Service (DoS) is. For example, an ISP might consider large multimedia downloads an attack to overload its infrastructure or have it pay high interconnection fees. On the other hand, a user will certainly consider selective bandwidth reduction that is used by ISPs as a countermeasure, as a DoS measure. Given the nature of their business relationship, neither side is likely to openly admit that they are fighting each other.

In this paper we attempt to formalise the concept of Stealth DoS, including listing mechanisms that may be used at high link speeds. We concentrate on mechanisms that might be used in one particular area, voice over IP (VoIP). We start evaluating them under the different aspects, including their cost, political suitability and the likelihood for countermeasures to succeed. We expect that this will give both sides better insight on their options and plea for peace, hopefully in an attempt to avoid and open war.

I. INTRODUCTION

Traditionally, the positions in DoS game were clear: Those sending the bulk data were the bad guys, those at the receiving end the good guys. The sides were set in the distributed DoS (DDoS) attacks from early 2000 [1], [2]. In the following years, the borders became increasingly blurred:

DDoS. There was inaction from the providers, some even charged the victims for the bandwidth consumed. The users did not see any improvements and some started blaming the ISPs for their inaction and lack of reliable countermeasures. This extended from traditional DoS/DDoS to E-Mail and Wiki/Blog/Bulletin board spam.

Multimedia. The advent of broadband allowed for multimedia files to be exchanged (first music and still pictures, later videos). Under the dominant policy of “flat-rate” fees toward the end users but per-megabyte pricing to the upstream provider, the ISPs came under financial pressure. The increasing number and size of software and update downloads put up even more pressure. ISPs started buying equipment that could detect and rate-limit many of the protocols used for multimedia sharing or cancelled contracts with bandwidth-hungry customers.

VoIP. The advent of broadband and VoIP started making telephone companies their own competitors: Every DSL line that was used for VoIP potentially reduced the number of actual telephone conversations. Some equipment vendors have started including VoIP detection mechanisms in their

products [3], claiming that they do not care what ISPs do with that information.

Users feel that their ISPs no longer care about their customers, but just about themselves. This makes users more likely to start implementing evasive measures against their ISP’s regulations, which in turn might be considered some form of service denial, potentially to other customers, by the ISP. As a result, skirmishes will start accumulating, possibly resulting in a big war or secession between ISPs and users. This is likely to result in the creation of an alternative network structure excluding ISPs.¹ Such a network split is unlikely in the interest of the ISPs and might neither be of advantage for society in general [4].

Despite these hidden tendencies, neither side is currently willing to openly start a fight. Instead, some members of either side have started to invest in appropriate tools that would prepare them and partly start using them. In this paper, we will look at the available arsenal of tools, both existing and hypothetical, and what the countermeasures are. We also provide a view into where we consider the Nash equilibrium to lie, i.e. whether investing in these tools actually pays. Our view is influenced by the fact that both sides will try to stay under cover for as long as possible, so the quality reduction needs to be stealthy. We picked VoIP as a good example, because already subtle changes to the network can cause noticeable quality degradation. This allows the degradation measures to be more easily labelled as maintenance effects or using other excuses. Nevertheless, our observations apply to a more general choice of traffic by adapting the traffic sensors or actors accordingly.

A. Cost and benefit for the ISP

The prices for broadband connection such as DSL or over cable television networks have been falling over the past years, yet the ominous “last mile” typically still remains in control of one or two former governmental monopolists. If at all possible, the infrastructure has to be rented by alternate service providers for comparatively high prices. Furthermore, customer charging typically is done on a flat-rate basis, while the ISP interconnection agreements frequently include

¹Such an alternate network is likely to use rooftop wireless networks in unlicensed spectrum together with routing technology developed by the mobile ad-hoc community. The secession reminds of the early ages of the Internet, which was (partially) built as an alternative to the system owned by the telephone operators, where the dataheads were unhappy with its regulations and tariffing.

*In the meantime, Mahak Patidar returned to IIT Bombay, Mumbai, India, <mahak@cse.iitb.ac.in>.

bandwidth-dependent charging components. As a result, the ISPs can hardly earn money on the basic Internet connection, especially not on users who use their bandwidth generously, e.g. using VoIP and other multimedia traffic.

Therefore, ISPs have to earn money on value-added services, such as offering VoIP themselves. That makes the ISPs willing to detect and degrade selected traffic flows, like VoIP traffic of competitors, even if they have to first invest money into the hardware infrastructure that enables the detection [3].

B. Paper organisation

The paper is structured as follows. We first look at what defines VoIP quality, then at how we can detect VoIP traffic, how it can be degraded and what the users' countermeasures are. We will conclude with an analysis of who might be able to survive this siege situation of mutual stealth DoS the longer and at lower cost.

II. VOIP QUALITY MEASURES

There are several mechanisms for the measurement of VoIP quality:

Delay is the sum of all delays that appear during the packet transfer. For effective interaction, the delay has to be low. ITU Recommendation G.114 addresses this issue and suggests that the round-trip delay should stay below 150..400ms to avoid disrupting the user experience, such as confusing human floor control (a summary can be found in [5]).

Packet Loss is originated by overload situations or transmission errors. In data networks, the loss of a packet does not necessarily lead to a loss of information, as the lost packet might be recovered from e.g. using retransmit. For a transmission in realtime, packet loss directly leads to quality reduction (in VoIP: lower intelligibility). As delay is critical, retransmission is rarely an option. Frequently, some amount of packet loss will be absorbed by the codec or the human recipient.

Jitter describes fluctuations in the inter-packet gaps. Jitter is typically compensated by a jitter buffer in the receiver, which evens out the spacing again. A large jitter buffer introduces unnecessary, disruptive delay, while too short a jitter buffer will result in packet loss (c.f. Figure 1). The choice of the optimal jitter buffer size is therefore somewhat of a black art.

Modification of the traffic means to either inject packets or to modify passing packets. Both modifications can heavily influence the VoIP traffic, resulting in packet loss (modifications which fail integrity/correctness checks) or codec misbehaviour/confusion of the human recipient (modifications which successfully pass the checks). The consequence may be that communication is no longer possible. Such events are highly improbable in today's Internet. Large amounts or specific instances of malicious modification can result in loud noise or even application crashes. Thus, they can hardly be considered stealthy and will not be discussed further.

Reordering changes the order in which packets arrive at the destination. This can be considered a relatively rare event in today's Internet and thus should not be used excessively if stealth is desired. Most VoIP user agents fix reordering in the jitter buffer; depending on the timing and the control logic of the jitter buffer, the effect can cause multiple packet drops at the receiver, when it made the decision to start playing a later packet.

Reasonable amounts of delay, packet loss, and jitter can be transparent and common in the wild, so that the user does not realise this as a quality mechanism. Heavy use of these mechanisms as well as traffic modification/reordering are very brute and obvious mechanisms.

III. VOIP TRAFFIC IDENTIFICATION

For simplicity, the following discussion focuses on SIP/RTP traffic, ignoring other protocols such as Skype, even though Biondi and Desclaux describe how to identify and block Skype traffic [6]. Nevertheless, there are little SIP/RTP specifics and all but one of the detection tools are generic.

Direct identification. The most reliable mechanism is to directly identify VoIP traffic. [7] describes how to identify RTP/RTCP traffic with high probability, [6] explains how to find Skype traffic. These methods can be expensive at high link speeds and may fail if not enough context is available or the traffic is simply tunnelled.

Mere tunnelling further raises the cost of packet processing for the detector by a significant factor as a significant part of the packet must be searched for weak signs which in turn must be correlated with other packets of the same flow before any statement about the packet carrying VoIP traffic can be made. Of course, encryption, maybe even with random padding, will make this form of detection all but impossible. Skype, which is known for its heavy reliance on encryption, can be detected and blocked, as the connection setup packet is not encrypted and highly predictable [6].

Transport identification. Most VoIP systems rely on UDP for transport. The presence of UDP alone is not a perfect indication, as UDP is also used for the Domain Name System (DNS), Network Time Protocol (NTP), security tunnels (e.g. OpenVPN), interactive online games, or other streaming applications, most of which are accepted or even encouraged and supported by many ISPs. Therefore, these protocols should not be affected. This can be achieved by white-listing these protocols.

Host/port identification. Many VoIP protocols use a default port or a port in a default range or contact a predefined host as part of their operation. This can be identified and then blacklisted, either on a per-packet, a per-session, or a per-host basis.

Packet size. VoIP packets can be distinguished from other multimedia streaming protocols by their short packet size. Again, short packets are a weak criteria if used exclusively, as it also matches many other UDP messages

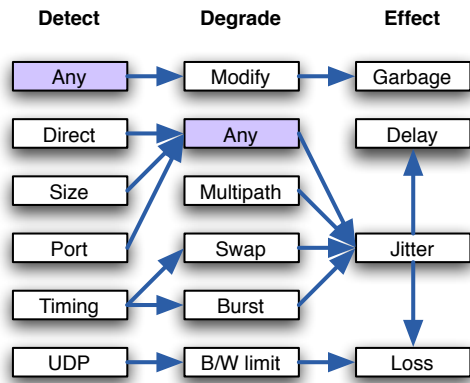


Fig. 1. Effect of detection and degradation mechanisms on the resulting VoIP quality

(NTP, DNS, games, control messages) as well as a large quantity of TCP control traffic.

Packet rate. The first of two timing-related aspects, packet rate, is the weaker of the two but easier to identify. Looking for a pre-defined minimum number of packets per unit time in a single stream is supported by many high-speed networking devices.

Packet frequency. The more complex mechanism is to watch for (relatively) evenly spaced packets. This does not seem to be supported by current network devices and seems to be very costly to implement at high link speeds.

IV. VOIP QUALITY DEGRADATION

When a potential VoIP stream has been identified, its quality can be reduced. The following mechanisms are readily available:

Drop. Drop specific or random packets in the stream at a higher rate than others. Raising the loss too high puts the ISPs plausible deniability at risk.

Delay. Delaying all packets by a specific or random period. Inserting constant delay only increases the perceived round-trip time while random delays might also cause the receiver to select too small a jitter buffer, causing packets to arrive too late and being thrown away by the receiver.

Burst creation. Queue packets from a given flow until a threshold count or maximum delay is reached, then release them. This introduces large amounts of jitter (and delay).

Packet swapping. Similar to burst creation, but re-inject the packets in reverse order. This might also convince certain user agents to start playing out the (original) later packet, making the (original) earlier packet(s) useless.²

Multipath routing. Send packets along different paths. This is harder to control than the latter two items, but may give more plausibility to the denials.

²If both the source and destination users' ISP use this technique, it will be reduced to burst creation.

Limit bandwidth. This approach apparently is in use at many ISPs for reducing file swapping traffic. It results in increased delay and/or drop, depending on the buffer size configured.

V. COUNTERMEASURES

There are several options available to the end user to escape the ISP's efforts, as listed below.

A. Evading detection

Direct identification. VoIP can only be identified directly, if it really looks like VoIP traffic, i.e. use encryption (like IPsec) or tunnelling (e.g. IP over DNS³) to hide the traffic in the data stream. Another way to avoid the direct identification is the insertion of confusing traffic [8], [9]. The traffic might be ignored by the end system or application, such as containing incorrect checksums, too low a TTL, illegal syntax, a (potentially slightly modified) replay of an earlier packet, or any other technique used to confuse intrusion detection systems.

Transport identification. Fake a different transport protocol, such as encapsulating in a TCP-look-alike protocol. It appears to be easy for the end system or a middlebox to create fake TCP headers which, at high link speeds, are hard to distinguish from legitimate TCP traffic.

Host/port identification. Use a different/multiple ports; avoid contacting well-known hosts (e.g. by using a peer-to-peer protocol which does not rely on central servers).

Packet size. Pad or encapsulate the data packets to make them longer. RTP supports adding up to 255 bytes of padding [7], while Skype seems to support arbitrary padding [6]. For flat-rate broadband users, this might turn into a backlash against the ISP: The users' bandwidth will be sufficient for additional traffic while the ISP might not like the interconnection fees associated with its higher traffic volume.

This can be especially attractive when combined with transport faking: About half of the TCP packets have traditionally been around 40...48 bytes in size (connection setup/teardown and acknowledgements [10]). Recent studies suggest that these packets now contribute close to 80% of the total packet count.⁴

Packet rate. Multiplex source/destination ports, use fake sources to avoid linking the packets into a single stream.⁵ Alternatively, the host may use multipathing, e.g. by participating in an overlay system such as RON [11] or an onion routing system such as TOR⁶. This works

³<http://savannah.nongnu.org/projects/nstx/>

⁴<http://ipmon.sprint.com/packstat/viewresult.php?NULL:pktsz:sj-00.0-050110>

⁵The ISP could avoid the host faking by installing source faking prevention, which would also help against some traditional DDoS attempts.

⁶<http://tor.eff.org/>; users would typically chose short anonymising chains for optimum voice quality.

especially well when the rate is limited on a “per flow”-basis, which is necessary if the measure is directed against VoIP traffic and not the user’s overall traffic generation.

Packet frequency. Spraying fake packets at random time intervals into the packet stream will require the ISP to use more complex (i.e., expensive) mechanisms to detect the inter-packet timing. Again, data might be injected in a fashion transparent to the end-system, as described above. In any case, frequency analysis seems to be very demanding on routers, even with the help of advanced flow grouping techniques [12].

Obviously, multiple mechanisms can be combined as necessary to avoid the detection or degradation. Also, advanced applications might adapt to provide minimum use of each of the mechanisms necessary to achieve the quality desired.

B. Avoiding degradation

Drop. To keep the dropping of packets stealth, we expect no more than 10-20% of the packets to be dropped. We distinguish between two types of drops: continuous drops and burst drops. For continuous drops send packet pairs or a checksum packet over the last 3 packets. If a packet gets dropped, the latency is still acceptable. The Real Time Streaming Protocol (RTSP) would be the appropriate mechanism. Burst drops are more problematic. A burst drop would be the loss of 5 packets following on each other, or the drop of all packets for 10 ms. Burst drops are no longer stealth. Use forward error correction (FEC) and a larger and/or variable jitter buffer.

Delay. With variable delay, use more jitter buffer at the receiving end. Higher delays will be less convenient during the discussion, but will not influence speech quality.

Burst creation, packet swapping, multipath routing.

Their impact can be effectively cancelled by sending multiple packets: duplicate, FEC, and dummy packets should all help equally well.

Limit bandwidth. This seems to be the only mechanism which is non-trivial to ignore. It needs to be overcome by evading detection or using multiple different streams, as described above.

VI. CONCLUSIONS AND FURTHER WORK

We provided a comprehensive list of subtle DoS measures to be used when selectively degrading the customers’ perceived quality without having to openly admit such measures. At the same time, we have shown that the end users can easily overcome almost all combinations of detection/degradation attempts with minimal effort and at almost no cost. Many of these countermeasures can be implemented easily using standard tools (e.g., tunnels) or without requiring co-operation from the receiving end (many of the fake/duplicate traffic insertions).

We therefore conclude that most investments by ISPs in this direction will provide a short market advantage at best. In the near term, users (and maybe even user agents) will start to introduce countermeasures, either by design or by

implementing privacy mechanisms such as encryption. This will make the investments in equipment and know-how moot, maybe even counter-productive, if interconnection traffic with other ISPs is charged on a volume basis.

To investigate the practicability of our thoughts, we are in process of implementing traffic modifiers for both sides of the game. Our early prototype is able to monitor and modify the traffic between SIP users by using a bridge. All the communication occurs through the bridge which in a Java framework based on Jpcap⁷ provides various functionality to control and modify the traffic. The bridge is used for selectively dropping packets, delaying RTP packets, modifying the the RTP headers, etc.

We are working on providing a generalised packet modification/delaying system built on top of our bridge which we plan to use for other forms of DoS research as well.

REFERENCES

- [1] Rich Pethia, Alan Paller, and Gene Spafford, “Consensus roadmap for defeating distributed denial of service attacks,” <http://www.sans.org/dosstep/roadmap.php>, 2000.
- [2] Jelena Mirkovic, Janice Martin, and Peter Reiher, “A taxonomy of ddos attacks and ddos defense mechanisms,” Tech. Rep. 020018, Computer Science Department, University of California, Los Angeles, 2002.
- [3] Steven Cherry, “The VoIP backlash,” *IEEE Spectrum*, vol. 42, no. 10, pp. 61–63, Oct. 2005.
- [4] David P. Reed, “That sneaky exponential—Beyond Metcalfe’s law to the power of community building,” <http://www.reed.com/Papers/GFN/reedslaw.html>, 1999.
- [5] Cisco Systems, “Understanding delay in packet voice networks,” White Paper 5125, Cisco, <http://www.cisco.com/warp/public/788/voip/delay-details.html#standarfordelaylimits>.
- [6] Philippe Biondi and Fabrice Desclaux, “Silver needle in the Skype,” Presentation at *BlackHat Europe*, <http://www.blackhat.com/presentations/bh-europe-06/bh-eu-06-biondi/bh-eu-06-biondi-up.pdf>, Mar. 2006.
- [7] Henning Schulzrinne, Stephen L. Casner, Ron Frederick, and Van Jacobson, “RTP: A transport protocol for real-time applications,” RFC 3550, Internet Engineering Task Force, July 2003.
- [8] Eric Cronin, Micah Sherr, and Matt Blaze, “The eavesdropper’s dilemma,” Tech. Rep. MS-CIS-05-24, University of Pennsylvania, 2005, <http://www.crypto.com/papers/internet-tap.pdf>.
- [9] Ronald L. Rivest, “Chaffing and winnowing: Confidentiality without encryption,” <http://theory.lcs.mit.edu/~rivest/chaffing.txt>, Mar. 1998.
- [10] National Laboratory for Applied Network Research, “WAN packet size distribution,” <http://www.nlanr.net/NA/Learn/packetsizes.html>, 1997.
- [11] David G. Andersen, Hari Balakrishnan, M. Frans Kaashoek, and Robert Morris, “Resilient overlay networks,” in *Proceedings of 18th ACM Symposium on Operating Systems Principles (SOSP)*, Banff, Canada, Oct. 2001.
- [12] Cristian Estan and George Varghese, “New directions in traffic measurement and accounting: Focusing on the elephants, ignoring the mice,” *ACM Transactions on Computer Systems*, vol. 21, pp. 270–313, Aug. 2003.

⁷<http://netresearch.ics.uci.edu/kfujii/jpcap/doc/>