A proactive and collaborative DDoS mitigation strategy for the Dutch critical infrastructure

Cristian Hesselman[1], Jeroen van der Ham[2], Roland van Rijswijk[3], Jair Santanna[2], Aiko Pras[2]
[1]SIDN, the Netherlands
[2]University of Twente, the Netherlands
[3]SURFnet, the Netherlands
Corresponding author: cristian.hesselman@sidn.nl

April 4, 2018

Banks and government agencies in the Netherlands repeatedly suffered from outages in the past few months as a result of relatively small DDoS attacks. This is worrisome, because DDoS attacks will only get bigger and more complex, for instance as a result of the emerging "Internet of Things". We therefore argue for a proactive and collaborative DDoS mitigation strategy for the Dutch critical infrastructure, which revolves around providers of critical services continually collecting information on potential and active DDoS sources and automatically sharing this information with other providers as part of a to-be-developed system called the "national DDoS-radar".

The Internet in 2018: essential, but vulnerable

The Internet has hugely eased and enriched our daily lives and has become an essential part of it. The services we use increasingly depend on the Internet, which includes services that the Dutch government has flagged as critical because their "failure or disruption … would result in severe social disruption and poses a threat to national security". Examples are electronic payments, Internet access, energy supply, and emergency communications.

Unfortunately, the Internet also makes critical services vulnerable to disruptions as a result of DDoS attacks on the providers of these services, such as ISPs, banks, government agencies, and DNS operators. They have to deal with increasing attack sizes, which currently peak at terabits per second. For example, Github and Dyn both had to handle attacks of more than 1 terabit per second, two orders of magnitude larger than the attacks on Dutch banks and government agencies earlier this year (estimated at 40 gigabit per second) and four orders of magnitude larger than the attacks on Estonia in 2007 (estimated at 90 megabit per second). DDoS attacks are also getting more complex because they are being launched from an increasing range of DDoS sources that can be located anywhere in the world and that can hit multiple targets simultaneously. Examples of DDoS sources include insecure devices in the "Internet of Things" (IoT), open "memcached" servers, and "booter" sites that offer DDoS attacks for a few tens of Euros.

Acting reactively and individually is insufficient

Our position is that the Dutch critical infrastructure's capabilities to sustainably handle the increasing size and complexity of DDoS attacks are currently insufficient. This is because providers of critical services predominantly fight DDoS attacks reactively, usually by sending the DDoS traffic they receive to a commercial DDoS scrubbing company that separates legitimate traffic from attack traffic. In addition, critical service providers often operate individually and do not have a current view on DDoS attacks on other providers. This reactive and individual mitigation strategy increases the probability that a service provider is insufficiently prepared for an attack and that this will result in service disruptions. In addition, commercial scrubbing companies have a financial interest to stick to this model, which means that any change will need to come from the critical service providers themselves.

Towards a proactive and collaborative strategy

We therefore propose that the Dutch critical infrastructure introduces a *proactive and collaborative* DDoS mitigation strategy based on the "national DDoS radar", a to-be-developed system that that continually creates "fingerprints" of potential and active DDoS sources and automatically shares them with connected critical service providers. The fingerprints for instance represent the Internet addresses of DDoS sources, the identifiers of the networks where they reside, and detailed descriptions of the DDoS traffic they generate. The DDoS radar creates fingerprints based on sensors of critical service providers, such as DDoS logs, crawlers that scan the Internet for booter sites, and honeypots that attract and map out compromised IoT devices.

The DDoS radar enables the proactive mitigation of DDoS attacks because its fingerprints provide critical service providers with more insight into potential DDoS attacks. This allows them to detect and filter DDoS sources more quickly, thus reducing the probability of service outages. In addition, the collaborative nature of the DDoS radar enables critical service providers to widen their view from just their own DDoS sensors and network traffic to the entire Dutch critical infrastructure. For example, a bank could use the DDoS radar to share the fingerprints of a set of Mirai DDoS sources it is dealing with in real-time, after which ISPs that host the Mirai-infected IoT devices contact their customers to help them cleaning their devices. Other service providers use the information to configure their detection and filtering systems to prepare for a potential attack. Another example is when a critical service provider uses the DDoS radar to share information on a booter site it detected in the .nl domain, after which SIDN (operator of the .nl top-level domain) and relevant hosting providers use the information to take the site offline.

The DDoS radar also helps reaching a collective decision on when to incrementally and temporarily disconnect "Dutch" networks from the global Internet in case of extremely large DDoS attacks and subsequently enforce this decision in collaboration with the former Trusted Networks Initiative and the Dutch Continuity Board.

In addition, the DDoS radar contributes to collecting more data about the perpetrators of DDoS attacks (attribution), allowing the government to track them down and prosecute them better and more quickly. This also discourages potential new attacks on the Dutch critical infrastructure.

Anti-DDoS manifesto

We propose to capture the DDoS mitigation strategy in an anti-DDoS manifesto in which participating critical service providers confirm they continually collect information on DDoS sources and share it with other providers as part of the DDoS radar. Network operators use a the construct of a manifesto to route Internet traffic more securely.

The anti-DDoS manifesto also includes operational agreements, such as which types of information service providers share through the DDoS radar, how receivers should use it, how to collaborate with criminal investigation services, and which message formats the DDoS radar supports (such as IODEF). This includes use of best practice anti-DDoS measures, such as BCP38.

Embedding

We propose to embed the manifesto in a national anti-DDoS organization that drafts, manages, and enforces the manifesto. This organization will need the support from a wide range of critical service providers, such as ISPs, mobile operators, internet exchanges, hosting providers, banks, DNS operators, and government agencies.

The anti-DDoS organization could consist of a small board of independent experts supported by members of the operational security community. They also work together to track new DDoS developments to handle larger and different types of attacks in the future. The scientific community will also need to actively contribute here.

An important prerequisite is that the anti-DDoS organization operates as transparently as possible by publishing the manifesto, its membership, as well as attack details and mitigating actions. This will help increasing the trust of citizens, companies, and politicians that critical service providers will be able to handle the problem of DDoS attacks.

What's next?

We call upon the former partners of the Trusted Networks Initiative, the members of the Dutch Continuity Board, the NAWAS, banks, universities, and governmental agencies to discuss the feasibility of the anti-DDoS manifesto, thus allowing for a proactive and collaborative level of DDoS resilience for the Dutch critical infrastructure and our online society.

Bios

Cristian Hesselman is with SIDN, the operator of the Dutch national top-level domain, .nl. He directs SIDN Labs, SIDN's research team.

Jeroen van der Ham is a guest researcher at the University of Twente. He is an assistant professor in Internet security and ethics at the Design and Analysis of Communication Systems (DACS) group.

Roland van Rijswijk is with SURFnet, the Dutch National Research and Education Network (NREN) where he fulfills the role of Innovator Internet Security.

Jair Santanna is with the University of Twente. He is an assistant professor in Internet security at the Design and Analysis of Communication Systems (DACS) group.

Aiko Pras is with the University of Twente. He is a full professor in Internet security at the Design and Analysis of Communication Systems (DACS) group.

== END