|galois|

Three Mechanisms for Protecting the Digital Ecosystem from Botnet Threats

Galois, Inc. | July 2017

The open, unstructured nature of the global Internet has provided a number of significant advantages to the United States, including dramatically increased commercial activity. However, adversaries have used this same decentralized structure to create effective, hidden attack systems. These systems, incorporated into botnets, provide adversaries with a platform to launch a variety of attacks against the ecosystem, including both nuisance-level attacks (e.g., spam) and attacks against critical infrastructure (e.g., distributed denial of service, or DDoS, attacks against telecommunications hubs).

While there are many ways to address the problem of botnets, many of them fail in light of the vast existing infrastructure already deployed. Many proposed hardware and software solutions are infeasible, as even the threat posed by botnets is considered insignificant compared to the cost of a wholesale restructuring or upgrade attempt. In addition, some mitigations potentially violate the contracts between core Internet backbone providers.

We thus suggest three avenues for the National Telecommunications and Information Administration (NTIA) and the Department of Homeland Security (DHS) to pursue, each of which can be incrementally incorporated into existing systems.

Avenue #1: Automated, Collaborated Analysis & Defense

Over the last two years with funding from the DHS Science & Technology Directorate, Galois has been pursuing the development of *DDoS Defense for a Community of Peers*, or 3DCoP[1]. A key insight of 3DCoP is that one of the problems with addressing botnet attacks, and DDoS attacks in particular, is that the work required to understand their root cause happens too slowly. In particular, all the work required to understand the true origin of a DDoS attack happens as a result of human research, and is limited by when network administrators finally notice an attack is happening, review logs, and can contact their peers. During this time the attack continues unfettered, and the attacker is rewarded and encouraged by their success.

With 3DCoP, we automate much of the reasoning performed by human operators using a collaborative, peer-to-peer network. As a result, investigation can begin as soon as a network monitor senses something unusual, and requests for further information happen rapidly and automatically. In addition, because the system is peer-to-peer rather than centralized, this information can flow even when the network is under attack; there is no central node to bring down, and communication among endpoints can be carried by low-bandwidth emergency backhauls if their main lines to the Internet are down.

At its current state of development, 3DCoP is capable of recognizing early indications of a DDoS attack and performing investigations regarding the origins of the attack. Its ability to perform these investigations is limited by the number of nodes in the 3DCoP network, but it has performed interesting analyses even in our pilot deployments.

We believe that the collaborative investigation capabilities of 3DCoP represent a dramatic leap forward in network defense. Historically, the detection, dissemination, and mitigation of attacks has happened at human speed, and at the mercy of human intuition. With 3DCoP, these processes can operate at computer and network speeds, with the full reasoning power of modern technology at their disposal.

[1] The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the Department of Homeland Security or the U.S. Government.

3DCoP is currently available for pilot partners. Please contact Adam Wick (awick@galois.com) for more information on 3DCoP and open pilot opportunities.

Avenue #2: Improved Protocols for Attack Mitigation

One interesting problem we have discovered in building 3DCoP is that understanding how to respond to an attack is often as difficult as understanding the attack itself. Consider, for example, a typical DDoS attack against a medium-sized organization like a hospital: the attack must be noticed, and a system administrator notified. If this administrator is not already at the office, they must get on site; remote access will be blocked by the DDoS. At this point, they can attempt to discover what kind of traffic is involved in the attack.

At the current level of DDoS attack, though, this is not enough information. They may be able to update their firewall to block attack traffic, but it is likely that the attack involves so much information that blocking the traffic at the target endpoint is ineffective.[2] At this point, the scramble begins. The system administrator must dig through their notes and find the contact and account information for their upstream ISP(s). They must call each of them and convince them that (a) there really is a problem, (b) they are who they say they are, and (c) the course of action they recommend is acceptable. Comprehensively, this process can take hours.

To improve this response time, we suggest the creation of a (human) protocol for requesting mitigation of an attack. In particular, we recommend the creation of a standardized process that allows the victim of an attack to request action of their service provider. This process must present to the service provider the following information:

The identity and account information required to prove that the request comes from authorized personnel;

1. Evidence of the attack, in a form that is easily consumable by the ISP;
2. An action, or series of actions, that the ISP can take to mitigate the attack; and
3. Evidence that this action will not cause problems for the ISP;
4. Expectations of when the ISP will cease and/or remove mitigation elements (with either or both of human-initiated removal and automatic time-based removal)[3]

We believe that it may be possible to meet #4 by limiting the actions that can be taken in #3, but list them separately as they address different concerns.

Finally, we suggest that this particular effort focus on person-to-person protocols, rather than computer-to-computer protocols. However, we note that the DDoS Open Threat Signaling Protocol (DOTS), currently in IETF draft[4], may be an interesting foundation for automating some of this effort. The group working towards this protocol may also be compelling seed members for developing the working group we suggest.

---

[2] Consider: current attacks involve tens or hundreds of gigabits of attack traffic, while the vast majority of medium-sized organizations have less than one gigabit of bandwidth. In such cases, firewalls are irrelevant because the Internet link is saturated regardless.

[3] In many cases, the attack mitigation itself may be harmful to the target, it is simply less harmful than the attack. For example, one mitigation against an attack might be to block incoming traffic from a given geographical region. Doing so stops the attack, but also inhibits sales to that region. Establishing this limit for mitigation also helps the ISP manage the additional personnel or physical equipment resources needed to implement the mitigation. Thus, developing a time limit for a mitigation is a key part of the process.

[4] https://datatracker.ietf.org/doc/draft-teague-dots-protocol/

The development of such a protocol will require input from ISPs, their customers, and security experts. We recommend that NTIA convene such a group to draft this standard, with the goal of rolling it out to interested parties.

Avenue #3: Protection from The Endpoints

Finally, any solution to the problem of botnets must address the millions, perhaps billions, of existing devices deployed around the world. Each of these devices is a potential target of attack, and thus a potential member of a future botnet. This threat is particular dangerous in light of the increasing use of IoT devices in the home. These devices are typically manufactured with minimal security requirements, installed and managed by security laypeople, and inconsistently updated.

To address this threat, we suggest a focus on detecting IoT devices in the home and limiting their capabilities *at the router*. In particular, given the increasing use of carrier-provided networking equipment, including wireless access points and other home networking solutions, we believe that *a significant improvement in the fight against botnets can occur through automated detection and sandboxing of IoT components*.

As an example, consider an Internet-connected thermostat. As part of its operation, the thermostat will contact a small set of Internet servers in order to provide service to the homeowner: the manufacturer's website for updates and management, the local weather service to determine heating or cooling trends, and possibly the power company to try to limit costs.

We imagine an update to current router firewall software that can use this information to limit the consequences should the thermostat be hacked. In particular, we imagine that the firewall can let expected traffic function as normal, but abnormal traffic be rate-limited. Thus, even if the thermostat is hacked, its ability to participate in DDoS and other attacks will be limited.

The key to this mechanism will be the ability to determine "expected traffic." However, we note that tools such as `nmap` provide an interesting mechanism. These tools can, via a set of carefully crafted probes, make very intelligent guesses regarding what is running on a network. By expanding this database—and, in particular, with the aid of device manufacturers—ISPs could generate comprehensive databases that match `nmap` fingerprints with detailed information about devices, including the list of sites each device can be expected to visit.

Galois has performed an exploratory analysis of this idea and would be interested in building out this capability with support from partners, including Commerce or the Department of Homeland Security.

Conclusion

Securing every device connected to the Internet is an unsolvable problem, but we *can* take significant steps to reduce the impact of botnets. In this document we have presented three: automated sharing of attack information, protocols for more rapid response to attacks, and tackling the IoT problem at the home Internet gateway. Each of these items addresses a different aspect of the botnet problem: understanding what is happening, taking effective action in a timely manner, and stopping attacks at their sources. In addition, each of them is practical, with potential deployments in as little as 6-12 months. Afterwards, while botnets will still exist, we will have significantly hampered their effectiveness and provided a clear improvement in overall Internet security.