ARTICLE

# Mirai: The IoT Bot That Took Down Krebs and Launched a Tbps DDoS Attack on OVH

Written by: Liron Segal

Date: October 7, 2016

The "Mirai" botnet has infected hundreds of thousands of Internet of Things (IoT) devices, specifically security cameras, by using vendor default passwords for Telnet access. This IoT botnet successfully landed a Terabyte attack on OVH[1], and took down KrebsOnSecurity[2] with an Akamai-confirmed 620+ Gpbs attack. Following Mirai's author post, dissecting the malware's source code and analyzing its techniques (including DDoS attack methods that are rarely seen like DNS Water Torture and GRE) we can definitely expect the IoT DDoSing trend to rise massively in the global threat landscape.

IoT devices are very attractive to the DDoS business as they don't require additional expenses, social engineering attacks, email infection campaigns, exploit kits or fresh zero-days. It is common for these devices to have poor security standards such that their remote administration ports are publically accessible and susceptible to brute force and dictionary attacks, the ports are "protected" with vendor default passwords, and they don't have an anti-virus solution in place to prevent malware infections. Combine these gaping security holes that make them "easy to exploit," with the device managers being people in their homes without security expertise, and these IoT

[1] OVH Terabyte DDoS Attack: https://twitter.com/olesovhcom/status/779297257199964160
[2] Krebs DDoS Attack: https://krebsonsecurity.com/2016/09/krebsonsecurity-hit-with-record-ddos/

devices being always online, ever-ready to serve the botmaster, makes this is a very suitable breeding ground for launching more massive DDoS attacks.

## Shifting DDoS Attack Varieties and Trends

As most typical volumetric attacks today rely on ICMP, SYN and a variety of UDP reflection and amplification attacks, the author of Mirai has interestingly introduced less common “DNS Water Torture” and “GRE flood” attacks. Though this DNS technique was already observed in the past, it’s not common to see nowadays.

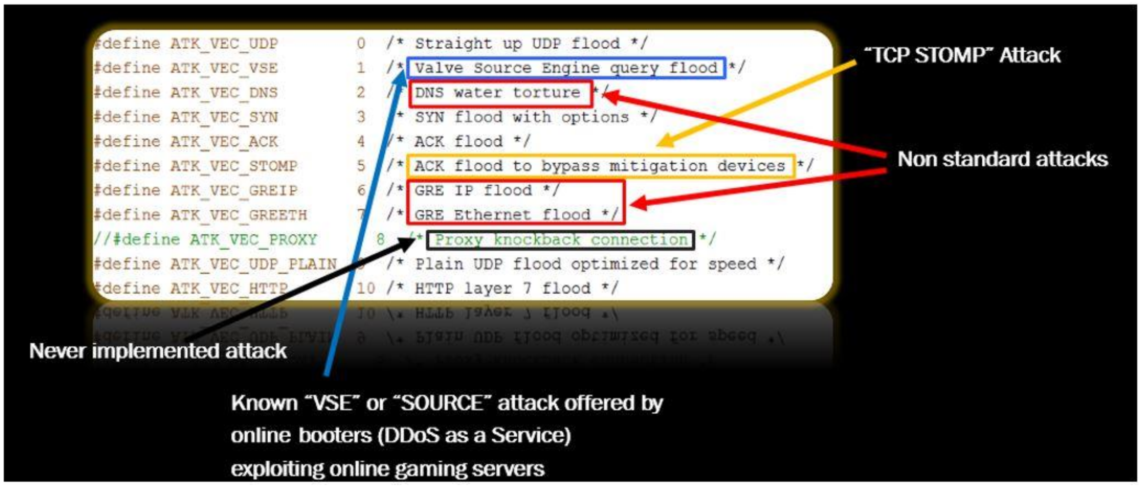

```
#define ATK_VEC_UDP          0  /* Straight up UDP flood */
#define ATK_VEC_VSE          1  /* Valve Source Engine query flood */
#define ATK_VEC_DNS          2  /* DNS water torture */
#define ATK_VEC_SYN          3  /* SYN flood with options */
#define ATK_VEC_ACK          4  /* ACK flood */
#define ATK_VEC_STOMP        5  /* ACK flood to bypass mitigation devices */
#define ATK_VEC_GREIP        6  /* GRE IP flood */
#define ATK_VEC_GREETH       7  /* GRE Ethernet flood */
//#define ATK_VEC_PROXY       8  /* Proxy knockback connection */
#define ATK_VEC_UDP_PLAIN    9  /* Plain UDP flood optimized for speed */
#define ATK_VEC_HTTP        10  /* HTTP layer 7 flood */
```

“TCP STOMP” Attack
Non standard attacks
Never implemented attack
Known “VSE” or “SOURCE” attack offered by online booters (DDoS as a Service) exploiting online gaming servers

Figure 1: DDoS Attack Methods

## “DNS Water Torture” Technique

This technique is different from the regular DNS reflection and amplification attack as it requires significantly less queries to be sent by the bot, letting the ISP’s recursive DNS server perform the attack on the target authoritative DNS server. In this attack, the bot sends a well formed DNS query containing the target domain name to resolve, while appending a randomly generated prefix to the name. The attack becomes effective when the target DNS server becomes overloaded and fails to respond. The ISP’s DNS servers then automatically retransmits the query to try another authoritative DNS server of the target organization, thus attacking those servers on behalf of the bot.

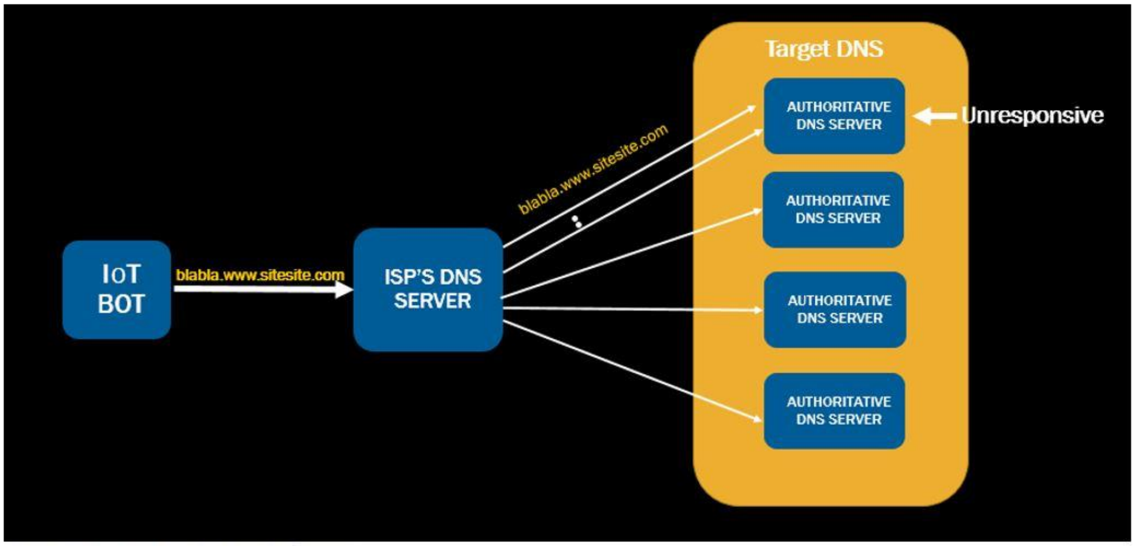

Figure 2: DNS Water Torture Technique

## GRE IP and Ethernet Floods

GRE (Generic Routing Encapsulation) is a tunneling protocol that can encapsulate a wide variety of network layer protocols inside virtual point-to-point links over an IP network, which is also used by DDoS scrubbing providers as part of the mitigation architecture.

Mirai holds two types of GRE attacks – with and without Ethernet encapsulation.

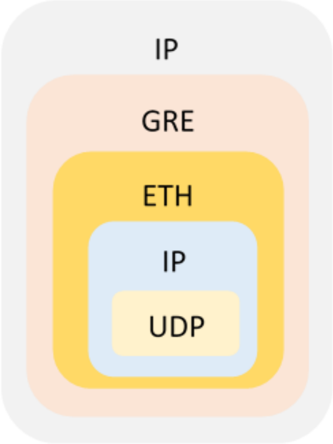

Figure 3: DNS Water Torture Technique

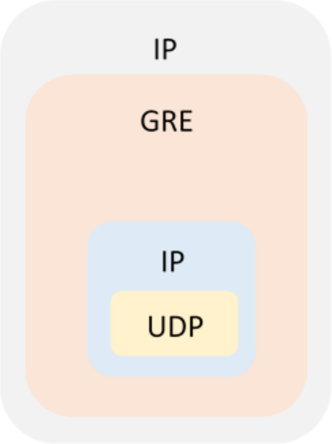

Figure 4: GRE IP

The encapsulated IP packet header uses the same parameters as the encapsulating IP header. The Transport Layer protocol for the encapsulated IP packet is UDP.

Most public routers will pass along the GRE packet because it's a widely used protocol for generating VPN connections. We speculate that GRE might be the protocol of choice due to its encapsulation nature, allowing huge payloads to be sent adding processing overhead of IP defragmentation to impact the target.

## The Hidden Attack

While there is no record of this attack and no supported command to invoke it, there is an implementation of a so called "cfnull" attack, which attacks the application layer. It is very similar to the GET/POST flood, but "cfnull" is designed to send a large POST payload of 80 MB of junk – a randomly composed alphabetic string - to the targeted server consuming webserver resources.

## The Hidden Attack

While analyzing Mirai's offered attacks, we took the perspective of how to mitigate it.

According to Mirai's creator, the so called "TCP STOMP" attack is a variation of the simple ACK flood intended to bypass mitigation devices. While analyzing the actual implementation of this attack it seems that the bot opens a full TCP connection and then continues flooding with ACK packets that have legitimate sequence numbers in order to hold the connection alive.

The Layer 7 "GET/POST" flood attacks support HTTP cookies and redirections that might handle simple bot challenges. While there is no actual support for bypassing more advanced challenges using JavaScript, several cloud DDoS scrubbing services are being fingerprinted by the Mirai bot.

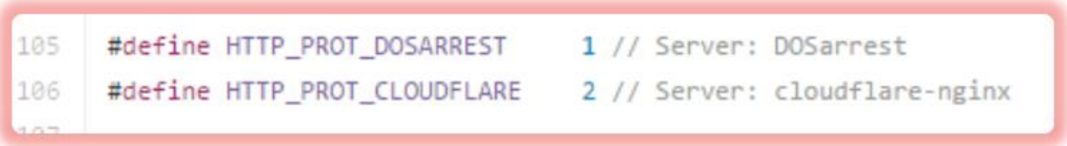

```
#define HTTP_PROT_DOSARREST     1 // Server: DOSarrest
#define HTTP_PROT_CLOUDFLARE    2 // Server: cloudflare-nginx
```

Figure 5: DDoS Scrubbing Services Fingerprinted

## IoT Nightmare

Judging by other leaked malware source code examples, such as the Zeus financial Trojan, we expect the underground market to adapt, combine, and improve the code, resulting in newer and enhanced variants.

Also, considering the low cost to maintain an IoT DDoS botnet, and referring to Gartner's forecast[3] stating connected things…will reach 20.8 billion by 2020, we can assume the IoT infection vector to grow. We should anticipate DDoS attacks over 1 Tbps to become more common in the near future, and see more "DNS Water Torture" and "GRE floods."

## Last Word

It seems that the bot creator named his creation after a Japanese series "Mirai Nikki (The Future Diary)" and uses the nickname of "Anna-senpai" referring to the "Shimoneta" series.

Figure 6: Mirai's Name Inspiration

[3] Gartner Forecast: http://www.gartner.com/newsroom/id/3165317

## About F5 Labs

F5 Labs combines the expertise of our security researchers with the threat intelligence data we collect to provide actionable, global intelligence on current cyber threats—and to identify future trends. We look at everything from threat actors, to the nature and source of attacks, to post-attack analysis of significant incidents to create a comprehensive view of the threat landscape. From the newest malware variants to zero-day exploits and attack trends, F5 Labs is where you'll find the latest insights from F5's threat intelligence team.

F5 Networks, Inc. | f5.com

US Headquarters: 401 Elliott Ave W, Seattle, WA 98119| 888-882-4447 // Americas: info@f5.com // Asia-Pacific: apacinfo@f5.com // Europe/Middle East/Africa: emeainfo@f5.com // Japan:f5j-info@f5.com