Attacks from Within:
Windows Spreads Mirai to Enterprise IoT - Draft

Steinthor Bjarnason
Arbor Networks, ASERT
sbjarnason@arbor.net

Jason Jones
Arbor Networks, ASERT
jasonjones@arbor.net

Abstract

When the Mirai IoT Bot surfaced in September 2016, it received a lot of publicity, not only because of the large-scale attacks it launched against highly visible targets, but also due to the large scale compromise of IoT devices. This allowed the attackers to subsume 100,000's of vulnerable, poorly secured IoT devices into DDoS bots, gaining access to resources that could launch powerful DDoS attacks.

However, as the original Mirai bot code scanned public Internet addresses to find new devices to infect, in most cases it was unable to detect and compromise IoT devices provisioned behind firewalls or NAT devices. As most firewalls stop these kind of scanning attacks, the (potential millions of) IoT devices behind firewalls were safe against detection and compromise. Or so most people thought…

1 Enter the Mirai Windows Seeder

In early February of 2017, a multi-stage Windows Trojan containing code to scan for vulnerable IoT devices and inject them with the Mirai bot code was detected in the wild.

This weaponization of a Windows Trojan to deliver IoT bot code reveals an evolution in the threat landscape that most organizations are completely unprepared to deal with: *DDoS attacks from within*. Windows machines infected by the Seeder will now actively scan for IoT devices whenever they establish a network connection. For example, if a laptop gets compromised by the Windows Mirai Seeder on a public wireless network, it will start scanning for vulnerable IoT devices as soon as it makes a network connection. This includes connecting to internal corporate networks via VPN, connecting to Wireless networks, or by using a physical network connection.

This is somewhat related to the old paradigm of attacking medieval castles. The castle walls (analogy: modern firewalls) were usually very effective at keeping the enemy outside the walls and stopping common attacks. However, they were useless if you could convince someone on the inside into becoming a traitor or by planting a spy inside the walls.

If there were no defenses inside the castle, the traitor/spy could now open the castle gates (disable the firewall), attack critical resources from the inside or simply burn down the entire castle. In medieval times, treachery was one of the most common cause of castle defenses being breached.

Any IoT device which gets compromised (scanners, printers, vending machines, light bulbs) will now be under the control of the threat actor, allowing him to launch DDoS attacks from inside the Enterprise against external and internal targets.

2 The Internals of a Traitor: The Mirai Windows Seeder

The Windows Mirai Seeder appears to be a refurbished version of a Windows Trojan which was discovered in the wild in early 2016. This Trojan was designed to attack CPE routers by brute forcing administrative passwords and then modifying DNS settings such that any devices on the inside would receive DNS replies from DNS servers under the attackers control.

Both the new Seeder and the older Trojan use brute force login attacks against Microsoft SQL servers, My SQL server and RDP with the goal of gaining administrative privileges on the target computer. It then proceeds to inject the malicious binary into the target computer, gaining full administrative control of the computer and launching the scanning process.

Post compromise, the Seeder will connect to its hardcoded Command & Control server (C&C) and download various files. This includes the Mirai bot code, scanning parameters, and information on the Mirai C&C servers.

The scanning process of the Windows Mirai Seeder has been modified from the original Trojan scanning process such that it now uses the same

scanning algorithm that the Mirai bot code uses. The Seeder will scan the IP ranges which were downloaded from the C&C and will attempt to detect vulnerable IoT devices on TCP ports 22 (SSH), 23 (Telnet), 5555 and 7547 (TR-069 SOAP management). If a vulnerable device is detected, it will try to brute force the Telnet and SSH usernames and passwords using a dictionary downloaded from the C&C. If the brute force login is successful, the Seeder will proceed to upload the Mirai bot code to the device, turning it into a Mirai bot which will then act in the same way as traditional Mirai bots[1].

3 The Nefarious Traitor – Turning Innocent IoT Devices into Zombies

Almost all networks, from the small SoHo to the largest Enterprise have a (large) number of IoT devices deployed on their internal networks. This can be anything from the smart TV in your living room to intelligent network enabled thermostats in a large Enterprise. These devices are, in most cases, protected by network firewalls making them unreachable by scans from malicious devices on the open Internet.

The Mirai Windows Seeder is a game changer because compromised Windows computers can now scan for vulnerable IoT devices whenever they connect to the internal network via VPN, Wireless or physical connections.

Unless proper care is taken to segment the internal network, this will make any device with an IP stack a potential target for compromise. Currently the Mirai bot infects devices like Web cameras and DVR recorders but it can easily be modified to attack other devices like printers, scanners, HVAC controllers and numerous other devices. Any device subsumed will start scanning for other vulnerable IoT devices and will proceed to infect those if detected.

There have already been reports of infected soda vending machines and light bulbs being used to launch DDoS attacks, confirming that the attackers are constantly finding new vulnerable devices to infect.

Coming back to the castle scenario, a single traitor can now rapidly subsume the innocent population of the castle into zombies, commanding them to attack the castle defenses or other internal or external assets.

[1] https://www.arbornetworks.com/blog/asert/mirai-iot-botnet-description-ddos-attack-mitigation/.

4 The DDoS Extortion Attack

A clever attacker could use the multi-stage Trojan to get inside the network, subsuming vulnerable IoT devices and computers on the internal network into his botnet and then scan the internal network to identify vulnerable network devices and critical services.

The attacker could then use this information to direct the bots on the inside to launch a devastating short-lived attack against the network itself and against critical services from the inside of the network, potentially disrupting the entire network. This would provide a proof-of-concept attack which proves to the victim that the attacker is now in control and continued availability of the service is based on the victim paying the attacker an extortion fee.

If the network hasn't been designed to withstand these kind of internal attacks, it will be a very time consuming and complex task to redesign and secure the network. Basically, the entire network security posture would have to be redone from scratch, beginning by shutting down all communication on all links, including any Internet connections.
If a network which hasn't been designed to withstand these kinds of attacks comes under attack, it will be very complex and time consuming to resume normal operations. Re-architecting the network is not something you want to do while under attack.

5 The Impact of Infected IoT Devices on Your Network

If a device infected by the Mirai Windows Seeder is active on an internal network, the following will be observed:

- There will be high volumes of scanning activity on internal networks as the Seeder searches for vulnerable Windows and IoT devices. As more devices get infected, the scanning activity will increase, potentially causing serious issues and outages with network devices like firewalls, switches and other stateful devices. These kinds of outages have repeatedly happened in the wild, both during the NIMDA, Code Red and Slammer outbreaks in 2001 and also recently during large scale Mirai infections at large European Internet Service Providers.
- Infected devices will contact their C&C server and will be subsequently used to launch DDoS

attacks. These attacks will result in high volumes of DDoS attack traffic which can potentially fill Internet and WAN links, resulting in loss of network connectivity. In addition, network based services like IP based voice services will be impacted, potentially resulting in IP phone service outages.
- Stateful devices like Firewalls and load balancers will also be at risk as they use state tracking to control traffic flows. These state tables will rapidly be exhausted due to the sheer traffic volume generated by the DDoS attacks, resulting in these devices no longer being able to pass network traffic. Firewalls and load-balancers are also often deployed in series and in front of each other. If one goes down, all network traffic will stop.
- When a device gets compromised, it will be under full control of the threat actor. It can now be used to perform reconnaissance on internal networks, launch DDoS attacks against internal targets, attack database servers and do whatever nefarious activity the threat actor is interested in performing.

This has the potential to turn your network into a virtual battleground where your (previously innocent) IoT devices actively attack external and internal targets, consuming valuable network resources including outgoing network bandwidth and capacity. Additionally, collateral damage in the form of network devices failing due to the sheer scanning and attack volume can occur.

6 Why Most Network Architectures Fail at Stopping this kind of Threat

Most network security architectures are designed for defending against external threats, it is **very** uncommon to see network security designs that treat both insiders and outsiders as potential threats.

This allows a well-equipped spy to enter the network using multi-stage Trojans which, after infecting the victim's computers, launch a second stage attack when the infected computers are connected to the often-unsecured internal network.

7 Network Impact of Bot Scanning

The Windows Trojan, has two main purposes. It scans for vulnerable Windows computers to propagate a copy of itself and it will also scan for vulnerable IoT devices to convert into bots. In addition, infected IoT devices will also launch their own scanning process to find additional IoT devices to attack.
Potentially the attacker could instruct the Trojan to scan for specific services or subnets, mapping out the internal network to find critical services. This kind of scanning hasn't been seen in the wild yet, but several other Trojans already have this capability.

All this scanning will result in:
- Large volumes of ARP (IPv4) / Neighbor discovery (IPv6) requests
- A flood of small scanning packets on network segments with infected devices.

Whenever a Layer 2 network switch receives an ARP packet for a specific IP, it will broadcast it out on all ports associated with the same network segment (physical/VLAN) as the one which the packet was received on. If there is a device with that IP address on the network segment, it will reply to the originating device, thereby providing it with its L2 MAC address. If there are multiple devices all scanning at the same time, the network switch might get overloaded by the flood of ARP packets, prohibiting it from performing its normal duties. Basically, it stops forwarding packets and the users won't be able to reach their services. This happened late 2016 at a large Internet Services Provider during a large scale Mirai infection.

In addition, this high scanning activity can also impact other devices on the same network segment, also resulting in high CPU loads and loss of functionality.

8 Network Impact of Internally Launched DDoS Attacks

When vulnerable IoT devices have been subsumed into the attacker's botnet, they will connect to their Command and Control (C&C) server and await instructions.

The botmaster can now instruct the bots to launch various types DDoS attacks. For example, the Miari bot is capable of launching the following attacks:

- UDP/ICMP/TCP packet flooding
- Reflection attacks using UDP packets with spoofed source IP addresses
- Application level attacks (HTTP/SIP attacks).
- Pseudo random DNS label prefix attacks against DNS servers.

The pseudo random DNS label prefix attack is designed to cause resource starvation of DNS servers. If this attack would be launched against an internal recursive DNS server, it would quickly result in the DNS server using up all its resources. This would then impact all network services which depend on DNS resolution, including web traffic,

network based services and potentially IP telephony services as those often use DNS for translating numbers to Uniform Resource Identifiers (URI).

The attack traffic for the flooding and reflection attacks will be generated as quickly as possible, potentially reaching high packet-per-second rates very quickly. A typical low end IoT device using a CPU similar to what is used in the Raspberry Pi computers can generate up to 8,000 packets per second which is enough to fill a 100Mbit link with large packets. A more powerful IoT device, for example an Internet connected HD network camera, can easily saturate a Gigabit Ethernet link with traffic.

A DDoS attack launched using internally based IoT devices could therefore potentially result in a flood of packets reaching Gigabit throughput. This malicious traffic will have to traverse the internal network on its way to its target on the Internet, sometimes traversing internal WAN links and traversing devices which are in many cases not capable of forwarding such high volumes of traffic. This could then lead to network outages, both on internal WAN/LAN links but also on external links due to the high traffic volume.

In addition, if the attack would use the infected IoT devices to launch DDoS attacks against internal targets, the impact could potentially be very high as most Enterprises do not protect internal resources against high-volume DDoS attacks originating from the inside.

9 How to Mitigate this New Threat

Defending against DDoS attacks from the internet is not trivial, especially if the network defenses are not secured properly to withstand such attacks. A well architected multi-layer design using Intelligent DDoS Mitigation Systems (IDMS) is capable of withstanding almost any kind of DDoS attack. However, such defenses are, in almost all cases, focused on defending against external attacks, not from attacks originating from the inside.

This new threat vector means that the network security designer will have to design the network to be resistant against attacks from both the inside and the outside. Also, care has to be taken to harden the network against collateral damage from scanning activities and the sheer volume of potential attack traffic traversing the network.

Interestingly, most Internet Service Providers have been doing this successfully for more than 20 years and there is considerable amount of Security Best Current Practices available which can help the network security administrator to properly secure his network.

Among those are:

- Cisco Systems (equivalent functionality is provided in network infrastructure devices from other vendors):
  - Service Provider Security Best practices[2]
  - Router Security Strategies[3]
- Arbor Networks:
  - Collection of security BCPs[4]
- NANOG:
  - An Architecture for Automatically Detecting, Isolating, and Cleaning Infected Hosts[5]

The information available is very comprehensive so a summary of the main phases for dealing with attacks are listed below:

1. Preparation: Prepare and harden the network against attack
2. Identification: Identify that an attack is taking place
3. Classification: Classify the attack
4. Traceback: Where is the attack coming from
5. Reaction: Use the best tool based on the information gathered from the Identification, Classification and Traceback phases to mitigate the attack
6. Post-mortem: Learn from what happened, improve defenses against future attacks.

One of the most important aspects of successful network defense are visibility and understanding what is going on. Without enough information, any kind of reaction has the potential to cause more harm than good. A well-known quote from Sun Tzu explains this very well:

> *"If you know the enemy and know yourself, you need not fear the result of a hundred battles. If you know yourself but not the enemy, for every victory gained you will also suffer a defeat. If you know neither the enemy nor yourself, you will succumb in every battle."*

The most important priority during attack is to **keep the network up and running**. If the network is down, no traffic will be able to traverse the network.

---

[2] http://bit.ly/2kUnZ1Y

[3] http://bit.ly/2mhJP0m

[4] https://app.box.com/s/4h2l6f4m8is6jnwk28cg

[5] https://www.nanog.org/meetings/abstract?id=662

A brief overview of the most relevant security tasks is provided in the following sections.

10 Mitigating Collateral Damage from Scanning Activity

As explained earlier, a network of compromised IoT devices and Trojans will see high levels of scanning activity. The scanning itself is not deliberately malicious but due to the high scanning volume, it can result in collateral damage on network devices like switches, routers and firewalls.

To mitigating the impact of scanning activity, the following tasks should be implemented:

- Segment the network such that devices with similar services/control are kept in their own segments.
- Implement IP source guard and DHCP snooping to block devices from masquerading as other hosts using spoofed source IP addresses.
- Only allow host devices and servers to communicate with the default gateway using Private VLANs thereby blocking the ARP packets from being seen by other devices on the same network segment.
- Implement "storm control" on the network devices to stop floods of packets.
- Implement the appropriate Control Plane Policing (CoPP) policies on network devices. If done properly, scanning activity with not impact the network devices.
- Use infrastructure Access Control Lists (iACLs) to control the flow of traffic between devices on the same network segment and between networks. Care has to be taken not to use stateful devices for this purpose as they have a tendency to collapse under heavy load, especially if a lot of small packets are being transmitted or if a DDoS attack is being launched from inside the network.

11 Blocking Trojan and Bot Infection Vectors

Both the Trojan and the Mirai IoT bot use network scanning to detect devices to attack. The Trojan uses brute force login attacks against Microsoft SQL servers, MySQL server and RDP with the goal of gaining administrative privileges on the target computer. Both the Trojan and the Mirai IoT bots scan for devices on TCP ports 22 (SSH), 23 (Telnet), 5555 and 7547 (TR-069 SOAP management) and will use brute for login attacks against SSH and Telnet and exploiting a known vulnerability against TR-069 configuration protocol.

To mitigate these activities:

- Implement network segmentation to separate IoT devices and client computers into separate network segments; additionally, each group of IoT devices should be grouped into their own segments.
- Implement strict control of network traffic to and from the individual network segments. These controls should be implemented using non-stateful controls like iACLs.
- Only allow client devices and IoT devices to communicate with their default gateway, no inter communication should be allowed. One example of such controls is Private VLAN.
- Wherever possible, separate Management traffic from data traffic and only allow management traffic originating from a specific set of IP ranges.

Coming back to our castle scenario, a well-designed castle had multiple layers of castle walls, with guards monitoring external and internal activities.

12 Mitigating the Impact of DDoS Attacks Launched from the Inside

A DDoS attack launched using IoT devices located on the inside of an enterprise network will cause very high traffic volumes, measured in both Bandwidth and packets-per-second. Even if the attack is destined towards external targets, the attack traffic will first have to traverse the internal network. This can result in network link congestion on WAN and LAN segments and high CPU load on network devices, all potentially leading to network outages.

To mitigate the impact of such attacks, the following should be implemented:

- Implement flow telemetry (i.e., NetFlow, IPFIX, et. al.) export, collection, and analysis, along with collection and analysis of recursive DNS queries and responses. This will provide comprehensive visibility into network traffic and will quickly detect any abnormalities and internally launched DDoS attacks.
- Implement Control Plane policing on all network devices. This will allow the network devices to withstand both direct attacks against the network elements and from having attack traffic traversing impacting the network device.
- Secure Routing protocols against attacks and overload. Without routing, no traffic can traverse the network.

- Implement Management Plane Protection to secure and protect management traffic. Also, reserve bandwidth and capacity on WAN and LAN links for management plane traffic. If you are not able to communicate with the network elements, the attack cannot be mitigated.
- Implement Unicast Reverse Path Forwarding (uRPF) policing to drop spoofed packets, this will stop all DDoS reflection attacks.
- Implement Data plane protection to filter and control what traffic should be allowed through the network. Examples:
    - A DNS server farm should only receive DNS traffic.
    - Client computers should only communicate with specific services on specific ports, not each other.

  Data plane protection should be implemented using non-stateful controls like iACLs, stateful controls have to tendency to crash and burn during heavy attacks.
- Do not trust any Quality-of-service tags made by clients, downgrade those such that management plane traffic has highest priority.
- Implement Remote Triggered Blackhole (RTBH) and Source-based RTBH (sRTBH) mitigation on network devices to allow for mitigation of attacks based on destination and source address. Properly implemented, RTBH/sRTBH are capable of stopping DDoS attacks with minimal impact to network devices.
- Implement Flowspec on network devices to allow for granular mitigation of attack traffic.
- Implement a quarantine system to isolate compromised devices. By utilizing flow telemetry collection/analysis, recursive DNS collection/analysis, and other forms of detection and classification, make use of recursive DNS poisoning to implement a universal 'soft' quarantine, and both VLAN- and WiFi channel-based 'hard' quarantine mechanisms to isolate botted devices.

13 Summary

The Windows Mirai Seeder is a simple delivery vehicle for the more dangerous Mirai IoT bot. However, as it will infect computers inside the Internet firewall, the attack surface has expanded tremendously, allowing for the creation of even larger Mirai botnets that will consequently have the capability to cause inadvertent collateral damage and to launch DDoS attacks against internal devices. A situation which most enterprise networks are not prepared to defend against.

A new threat scenario has emerged which has the potential to cause a myriad of issues in the future for networks with weak or non-existent defenses inside the corporate firewall.

A network designed and secured using the security BCP's described herein will be highly resistant to such compromise and the ramifications thereof. If one of your Windows systems becomes a traitor, it will not be able to subsume your innocent IoT population into an army of raving zombies…