A Tale of Two Botnets

netscout.com/blog/asert/tale-two-botnets

by Christopher Conrad on October 28th, 2021

Executive Summary

In June of 2021, a new botnet comprised of unpatched MikroTik routers emerged. Dubbed *Mēris* by security researchers who first reported it, this IoT botnet launched numerous application-layer HTTP and HTTP/S DDoS attacks against multiple targets worldwide, including Krebs On Security and Yandex. By some accounts, there are ~250,000 unpatched MikroTik routers worldwide which can potentially be compromised and incorporated into DDoS-capable botnets like *Mēris*.

Based on an analysis of identifiable botted MikroTik routers leveraged to launch DDoS attacks during the last four months, it appears that there are in fact at least two distinct MikroTik-based IoT botnets involved in these attacks. We have dubbed the second botnet *Dvinis*, which means 'twin' in Latvian (*Mēris* is Latvian for plague). The main difference between these is the use of HTTP Pipelining as a form of attack. *Mēris* leverages this attack method, while *Dvinis* does not.

Key Findings

- There are at least two DDoS-capable IoT botnets, *Mēris* and *Dvinis*, inhabiting the same population of unpatched, exploitable MikroTik routers.
- Since August of 2021, we observed multiple HTTP and HTTP/S application-layer DDoS attacks launched by *Mēris* and *Dvinis*, and assisted network operators in successfully mitigating these attacks.
- Both botnets are actively attempting propagation to expand and to date we are tracking approximately 4,800 *Mēris* and 3,500 *Dvinis* botted nodes.

NOTE: *NETSCOUT Arbor DDoS defense solutions can be used to detect, classify, traceback, and mitigate DDoS attacks launched by these botnets.*

Overview

In mid-2018, a vulnerability targeting the WinBox service of network vendor MikroTik's RouterOS was discovered, CVE-2018-14847. This vulnerability affected MikroTik RouterOS through version 6.42, permitting unauthenticated remote attackers to remotely read and write files via a directory traversal vulnerability in routers' WinBox interface. This allowed adversaries to extract admin passwords, create option packages that enabled developer backdoors, and then further allowed post-exploitation Telnet/SSH root user developmental access with administrative credentials.

While MikroTik patched these vulnerabilities soon after they were discovered and urged all their customers to upgrade to fixed versions of their software, a large population of these routers are still susceptible to intrusion because they have never been patched or were patched while still leaving remote administration mechanisms open to the Internet and without changing the credentials for administrative accounts.

A substantial number of these vulnerable MikroTik routers have been compromised and subsumed into the *Mēris* IoT botnet. The botnet has been used by attackers to launch multiple high-profile HTTP and HTTP/S DDoS attacks mainly targeting organizations in Central and Eastern Europe. *Mēris* has also been used in DDoS attacks against Russian networking conglomerate Yandex, as well as prominent security researcher Brian Krebs and his website Krebs on Security, who is well-known for his investigation of online criminal activities worldwide — including Russia and the Baltic states.

Mēris has been heavily publicized, and some security analysts initially assumed that all ~250,000 vulnerable devices were part of a single, massive botnet. However, after analyzing numerous DDoS attacks purportedly sourced from the *Mēris* botnet, assisting network operators to successfully mitigate many of these same DDoS attacks, and developing a methodology to identify and classify compromised MikroTik routers, we've been able to determine that the actual numbers of botted IoT devices is considerably fewer than 250,000, and that there are in fact at least two distinct MikroTik-based IoT botnets being used to launch HTTP and HTTP/S application-layer DDoS attacks.

Due to public reporting alleging that the *Mēris* botnet had been used to launch the largest DDoS attack on record, some organizations targeted by HTTP and HTTP/S DDoS attacks have initially assumed they were being targeted by *Mēris*, when in reality they were being attacked by *Dvinis*, or another botnet altogether. In fact, *Mēris* only accounts for a quarter of DDoS attacks we've observed being launched from compromised MikroTik devices. The remaining 75% of these attacks have actually been sourced from *Dvinis*.

The increased prevalence of DDoS-capable IoT botnets like *Mēris* and *Dvinis* is associated with a significant uptick in direct-path DDoS attacks ASERT has observed over the last several months. While this style of DDoS attacks has never completely vanished, the vast majority of DDoS attacks we've observed over the last several years have been various types of reflection/amplification attacks, in which the attacker spoofs

the IP address of the intended target and induces various abusable services to issue large, unsolicited responses directed towards the targeted network.

There are likely multiple reasons for the recent increase in direct-path DDoS attacks, but one salient factor may be a renewed emphasis on implementing source-address validation (SAV) by network operators in order to disallow spoofed packets from ingressing or egressing their networks. Without the ability to spoof the IP addresses of their intended targets, attackers can't launch reflection/amplification attacks, which can range in size from a few Gbps up to the very largest terabit-class DDoS attacks. Broader implementation of SAV is bad news for the criminal operators of spoofing-capable DDoS attack infrastructure, and we encourage all network operators to work on implementing SAV as broadly across their network edges as possible. More information on SAV and other network infrastructure best current practices (BCPs) can be found here (https://www.manrs.org).

Mēris Botnet Analysis

We've been tracking *Mēris* botnet activities since August of 2021. NETSCOUT's ATLAS Global DDoS attack telemetry revealed ~4800 *Mēris* nodes participating in numerous DDoS attacks. Our global honeypot network also observed multiple connections from 103 separate *Mēris* botnet nodes attempting to compromise and subsume additional vulnerable MikroTik devices into the botnet.

Mēris propagation scans for the following open ports:

- TCP/80
- TCP/8080
- TCP/3389
- UDP/123
- UDP/389
- UDP/1121 (the UDP ports are related to UDP reflection/amplification vectors).

The compromised MikroTik router infected with *Mēris* attempts to use the cited remote exploit on TCP/8291 as well as brute-force credentials over Telnet and SSH.
Mēris compromised devices exhibit the following characteristics:

- Utilize HTTP Pipelining
- TCP/2000 open - reports as Bandwidth test server and responds with a "SERVER HELLO" once a TCP 3-way handshake is complete.
- TCP/5678 open - reports as MikroTik Neighbor Discovery Protocol (MNDP)

To date, *Mēris* nodes have only been observed launching HTTP and HTTP/S application-layer DDos attacks, principally targeting Web servers with HTTP GET and POST request-floods. Many of these attacks — possibly all of them — are actually initiated by external attack harnesses under the control of the attackers, and are then relayed through *Mēris* bots via the SOCKS4/5 services present on the compromised MikroTik devices.

Mēris HTTP and HTTP/S DDoS attacks are somewhat unusual due to the incorporation of HTTP pipelining in the synthesized HTTP attack traffic. HTTP pipelining was intended to allow persistent TCP connections to be used to issue multiple successive HTTP/1.1 requests; however, it has been deprecated in almost all current Web browsers, and most Web servers do not support it.

An extensive set of some ~600 credentials are in use across the *Mēris* botnets. They are all easily brute-forced, and do not present a barrier to automated compromise techniques.

While there are *Mēris* nodes present in many geographies, we've observed significant concentrations in Central and Eastern Europe, South America, Southeast Asia, Northeast Asia, and the Indian subcontinent (**Figure 1**).

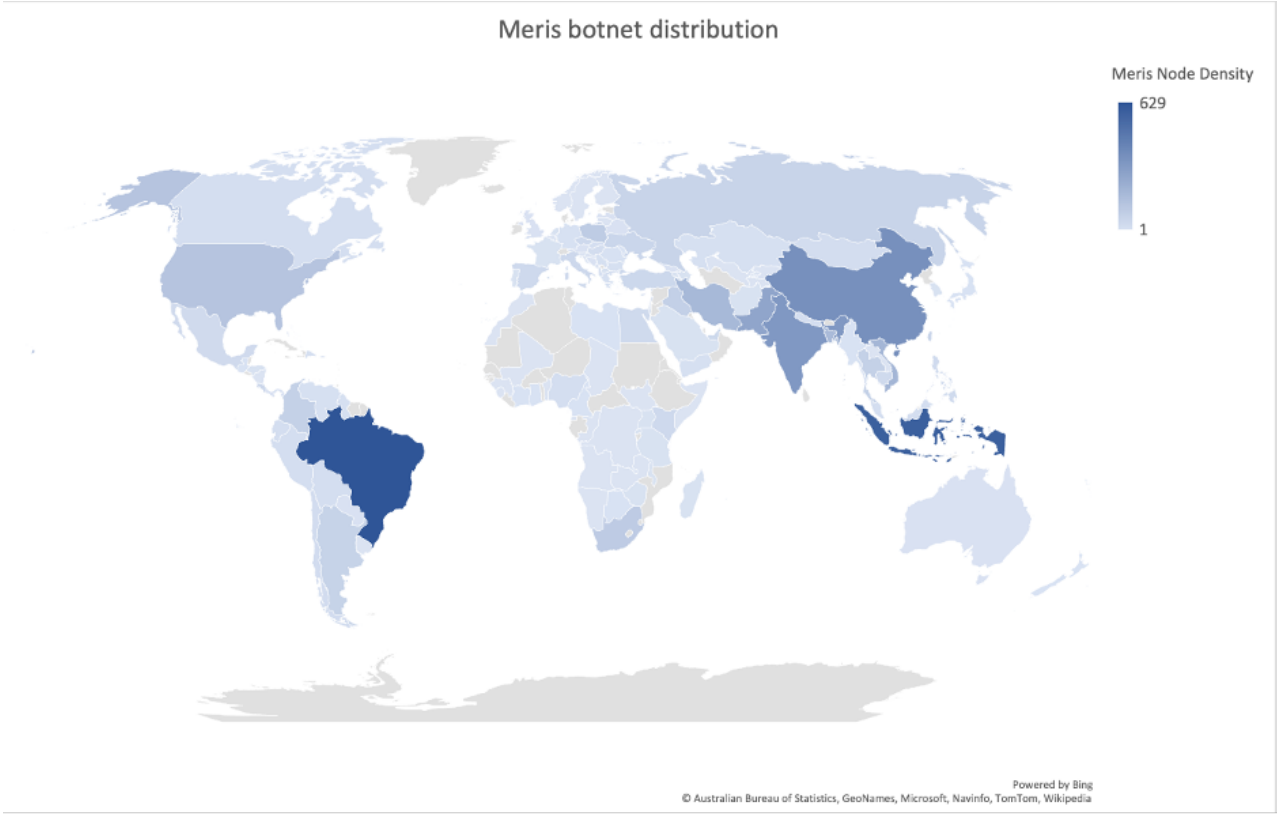

Figure 1: Mēris Botnet Distribution

The distribution of the botnet has higher concentration in some cities, though this isn't indicative of adversaries choosing a city to establish nodes, but rather a global popularity in MikroTik routers:

- Brazil (Recife, Sao Paulo, Rio de Janeiro, Belo Horizonte)
- Indonesia (Jakarta)
- China (Shanghai, Beijing, Zhuhai, Tonghua, Guangzhou)

Dvinis Botnet Details

The *Dvinis* botnet is very similar to the *Mēris* botnet described above. It also consists of

compromised MikroTik routers and is also used to launch HTTP and HTTP/S application-layer DDoS attacks. To date, we've observed ~3500 *Dvinis* nodes participating in DDoS attacks worldwide.

Unlike *Mēris*, *Dvinis* sourced HTTP and HTTP/S application-layer DDoS attacks don't appear to make use of HTTP pipelining; however, an apparent typo in the attack generators appends an extra '/' character to the end of the URIs targeted in HTTP POST and GET floods. It also appears that, as with *Mēris*, most, if not all of the observed HTTP and HTTP/S DDoS attacks sourced from Dvinis are also actually initiated by external attack harnesses and relayed via the SOCKS4/5 proxy subsystem built into the compromised MikroTik routers. The HTTP X-Forwarded-For field in captured attack packets includes the source IP addresses of the actual attack infrastructure being used to generate these attacks.

Dvinis compromised devices exhibit the following characteristics:

- TCP/2000 open - reports as Bandwidth test server and responds with a "SERVER HELLO" once a TCP 3-way handshake is complete.
- TCP/8291 open
- HTTP Post and GET Requests included a double slash (//) in the URI

We observed at least 50 of the ~3500 *Dvinis* bots attempt propagation to one of our honeypots. Propagation starts with a scan for the following open ports:

- TCP/80
- TCP/8080
- TCP/555
- UDP/123
- UDP/389
- UDP/3702
- UDP/5060
- UDP/11211 (as with *Mēris*, the UDP ports are associated with UDP reflection/amplification DDoS vectors, but includes some not scanned for by *Mēris*).

Dvinis brute-force propagation uses ~400 credential sets; some overlap with those used by *Mēris*, while others are unique to *Dvinis*.

We've found the largest concentration of *Dvinis* bots in Eastern Europe, Southeast Asia, and South America (**Figure 2**). As with Mēris, they're distributed on a number of networks in those regions and are primarily concentrated in large cities.

- Indonesia (Jakarta)
- Russia (Moscow)
- Brazil (Sao Paulo, Rio De Janeiro, Salvador)

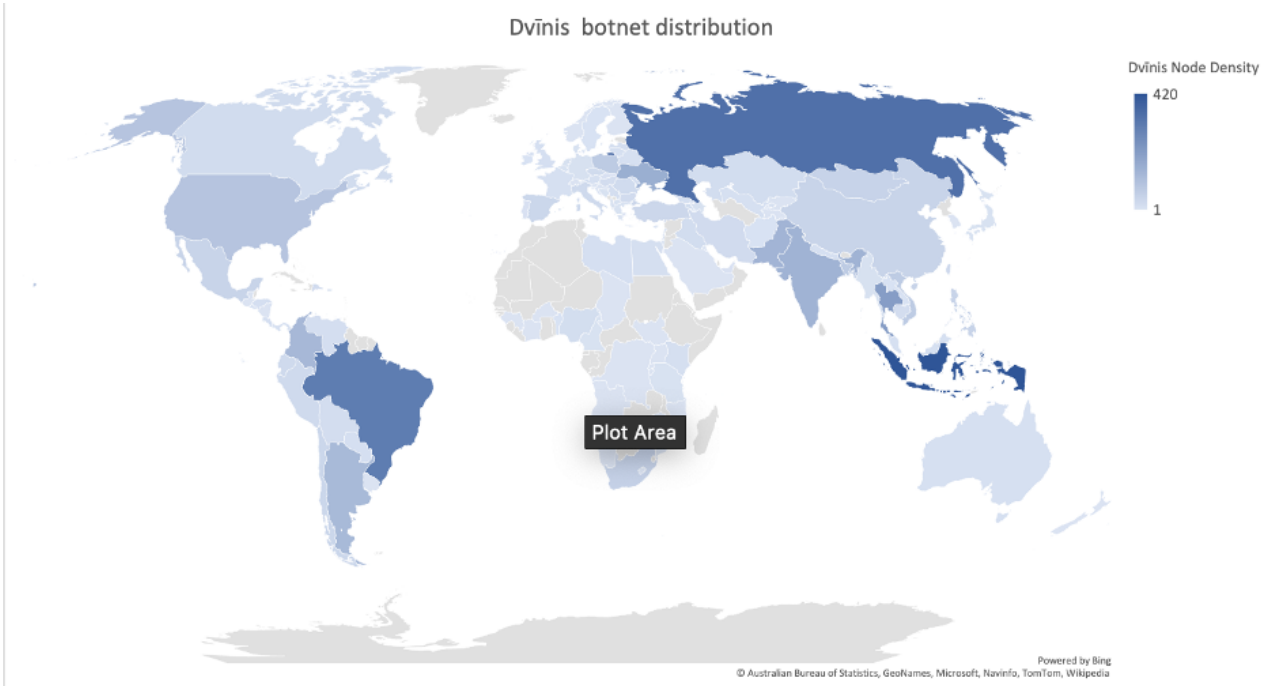

Figure 2: *Dvinis* Botnet Distribution

Conclusion

The rise of the *Mēris* and *Dvinis* IoT DDoS botnets almost certainly could've been prevented by operators of MikroTik routers ensuring their devices were kept up-to-date with security patches, as well as by implementing good password hygiene. These are issues we see time and time again in these cases.

These botnets essentially 'rose from the dead' thanks to an old exploit which enabled adversaries to build out not one, but at least two DDoS-capable IoT botnets. The good news is that the DDoS attacks launched by both botnets can be mitigated by implementing relevant BCPs; developing, updating, and rehearsing a comprehensive DDoS mitigation plan; and deploying an intelligent DDoS mitigation system (IDMS) such as NETSCOUT Arbor AED and NETSCOUT Arbor Sightline/TMS

References

https://blog.qrator.net/en/meris-botnet-climbing-to-the-record_142/
https://krebsonsecurity.com/2021/09/krebsonsecurity-hit-by-huge-new-iot-botnet-meris/
https://www.reuters.com/technology/russias-yandex-says-it-repelled-biggest-ddos-attack-history-2021-09-09/
https://blog.mikrotik.com/security/winbox-vulnerability.html
https://blog.mikrotik.com/security/meris-botnet.html

Contributors

Key Contributors: Roland Dobbins, Jon Belanger, John Kristoff, Steinthor Bjarnason