Protocol found in webcams and DVRs is fueling a new round of big DDoSes

reader comments

51 with 40 posters participating

Share this story

Hackers have found a new way to amplify the crippling effects of denial-of-service technique by abusing an improperly implemented tool found in almost 1 million network-connected cameras, DVRs, and other Internet-of-things devices.

The technique abuses WS-Discovery, a protocol that a wide array of network devices use to automatically connect to one another. Often abbreviated as WSD, the protocol lets devices send user datagram protocol packets over port 3702 that describe the device capabilities and requirements. Devices that receive the probes can respond with replies that can be tens to hundreds of times bigger. WSD has shipped with Windows since Vista and is one of the ways the operating system automatically finds network-based printers.

IoT strikes again

The WSD specification calls for probes and responses to be restricted to local networks, but over the past few months, researchers and attackers have started to realize that many Internet-of-things devices allow devices to send probes and responses over the Internet at large. The result: these

improperly designed devices have become a vehicle capable of converting modest amounts of malicious bandwidth into crippling torrents that take down websites. Depending on the device, responses can be anywhere from seven to 153 times bigger, an amplification that puts WSD among the most powerful techniques for amplifying distributed denial of service attacks.

Researchers with content delivery network Akamai were recently in the process of investigating WSD amplification when a customer in the gaming industry was hit with just such an attack. At its peak, it generated 35Gb per second of junk traffic. That's nowhere close to record-setting attacks of 620 Gbps in 2016 and 1.7Tb per second last year. Still, the new amplification method is concerning, in part because the pool of available devices—which Akamai estimates is more than 802,000—is so large.

"It's going to be pretty bad, especially once the bad guys figure it out," Akamai researcher Chad Seaman told Ars. "It's bad enough that most people should be concerned about being hit with it."

A researcher with Netscout, meanwhile, told Ars the DDoS mitigation service has seen 1,000 WSD-based attacks in the past three months, 473 of them in the past 30 days. The biggest attack delivered about 150Gbps and about 35 million packets per second. In a recent report, Netscout said it first saw attacks in May. The technique can amplify bandwidth by about 300 fold.

A brief history of amplification

In late February 2018, researchers reported that DDoS vandals had begun using a then-previously obscure method that delivered responses that were 10,000 times bigger than their initial query and, in theory, could be as many as 51,000 times larger. The amplification source was poorly configured servers

running memcached, a database caching system for speeding up websites and networks. In theory, the attacks meant a single home computer with a 100 megabit-per-second upload capacity was capable of bombarding a target with a once-unimaginable 5Tb per second of traffic. In practice, the attacks started with volumes of about 500Gbps.

A week later, Akamai reported memcached servers were being abused to deliver a record-setting 1.3Tbps attack against Github and, a few days later, the previously mentioned 1.7Tbps DDoS against an unnamed target located in the United States. In 2014, DDoS vandals took down *League of Legends*, EA, and other online services by abusing what was then a previously unknown amplification vector in the Network Time Protocol used to sync time on the Internet.

One of the earlier forms of amplification attacks include those that abuse poorly configured domain name system servers. The technique was used to knock out anti-spam organization Spamhaus in 2013 and disrupt DNS provider NS1 in 2016.

The common thread among WSD, memcached, NTP, and other widely abused amplification vectors is the user datagram protocol. UDP traffic is often described as "stateless" and "connectionless," because all parameters are contained in each packet at the time it is sent. That makes UDP traffic susceptible to forgeries that misidentify the party sending the data. Amplification attacks seize on this weakness. An attacker sends a server or device a large number of queries that replace their true origination location with the IP address of the DDoS target. The device or server then sends the target an equal number of replies that are much, much larger than the spoofed request it just received.

Tweaking malicious probes

The pool of 802,115 improperly implemented WSD servers found by Akamai deliver average and median amplifications of about

two times and 1.9 times the size of the initial request. But with certain tweaks, Akamai researchers were able to devise probes that generated responses that were much bigger. By sending probes that exploited a buffer overflow found in about 2,000 devices from an unnamed manufacturer, for instance, the researchers amplified responses by a factor of 153. Other tweaks amplified responses by seven- and 20-fold.

The modified probes worked in part by adding "padding" payloads to them that contained a MessageID. These IDs worked as a session tracking mechanism that's similar to browser cookies.

"In short, this means that, when a client sends a payload with a MessageID value, WSD at the application layer associates this value with future responses directed back at the respective client/session," Akamai researcher Jonathan Respeto wrote in a post scheduled to go live Wednesday morning. "In some cases, this translates to a persistent field that attackers can leverage to increase the overall responses returned from the WSD service."

It's tempting to think that repelling a flood of junk WSD responses is a simple matter of blocking port 3702. While the move will prevent the malicious traffic from entering the target's network, it does nothing to prevent the torrent from deluging the Internet provider that's upstream from the target. Internet providers often terminate customers when that happens. As a result, targets will likely need the help of a DDoS mitigation service. The more sustainable solution is for device makers to implement WSD in a way that restricts traffic to local networks. But that could take decades.

Not the first time

This isn't the first time that IoT devices have been found to implement a network-discovery protocol in a way that threatens its users. In 2013, researchers identified 81 million unique

IPv4 addresses that responded to universal plug-and-play discovery requests, even though the UPnP standard isn’t supposed to communicate with devices that are outside a local network. The incorrectly implemented standard, which was the result of flaws in home and business routers, exposed networks to hacks.

The more intermediate step for mitigating the threat is to find networks that are leaking large amounts of WSD traffic and get them to stop.

“Typically when you see the reflection vectors, there’s a concerted organizational effort across the community to find these [networks], figure out who owns them, and notify them of the potential they cause,” Seaman told Ars. “You can never get to 100%, but in the past, we have seen it put big dents in some of these networks.”

https://arstechnica.com/?p=1570245