WS-DISCOVERY AMPLIFICATION ATTACK

WS-Discovery amplification is one of an ever-growing-array of weapons that attackers use to create very large DDoS attacks. This paper talks about it from a DDoS defender's perspective and provides some recommendations to mitigate the threat.

EXECUTIVE SUMMARY

The Web Services Discovery or WS-Discovery is a UDP-based protocol that allows an attacker to spoof request with the victims IP that causes a reflected and amplified response toward the victim's infrastructure or service. In this advisory, we will provide a brief overview from the perspective of a DDoS defender and offer advice for protection.

Summary of WS-Discovery research findings:

- There are >800k potential WS-Discovery amplification sources
- Nearly 50 percent of the hosts respond with random high UDP ports
- Observed amplification factor of up to 95 times
- IoT devices account for a significant portion of the exploitable hosts

PART #: A10-MS-23239-EN-01 SEP 2019

OVERVIEW

The WS-Discovery protocol uses TCP and UDP ports 3702 to respond to a multicast address that gives information on services provided in the local network. In some cases, the response provides the requester quite a bit of data about these services (see below).

```
<?xml version="1.0" encoding="UTF-8" ?>
<s:Envelope
    xmlns:s="http://www.w3.org/2003/05/soap-envelope"
    xmlns:sc="http://www.w3.org/2003/05/soap-encoding"
    xmlns:d="http://schemas.xmlsoap.org/ws/2005/04/discovery"
    xmlns:a="http://schemas.xmlsoap.org/ws/2004/08/addressing"
    xmlns:dn="http://www.onvif.org/ver10/network/wsdl"
    xmlns:tds="http://www.onvif.org/ver10/device/wsdl">
    <s:Header>
        <a:MessageID>uuid:471c3ba3-b770-4e31-9744-89c3ef252502</a:MessageID>
        <a:To>urn:schemas-xmlsoap-org:ws:2005:04:discovery</a:To>
        <a:Action>http://schemas.xmlsoap.org/ws/2005/04/discovery/ProbeMatches</a:Action>
        <a:RelatesTo>urn:uuid:ce04dad0-5d2c-4026-9146-1aabfc1e4111</a:RelatesTo>
    </s:Header>
    <s:Body>
        <d:ProbeMatches>
            <d:ProbeMatch>
                <a:EndpointReference>
                    <a:Address>uuid:85f381e3-b806-40e4-aaca-8e0db335d0dd</a:Address>
                </a:EndpointReference>
                <d:Types>dn:NetworkVideoTransmitter tds:Device</d:Types>
                <d:Scopes>onvif://www.onvif.org/location/country/China onvif://www.onvif.org/name/General onvif://www.onvif.org/hardware/XVR
                onvif://www.onvif.org/Profile/Streaming onvif://www.onvif.org/type/ onvif://www.onvif.org/extension/unique_identifier</d:Scopes>
                <d:XAddrs>http://[redacted]/onvif/device_service</d:XAddrs>
                <d:MetadataVersion>1</d:MetadataVersion>
            </d:ProbeMatch>
        </d:ProbeMatches>
    </s:Body>
</s:Envelope>
```

Protocol	UDP Source Port	UDP Dest Port	IP Length	IPID
UDP	3702	45943	1500	0x5a24 (23076)
IPv4			1500	0x5a24 (23076)
IPv4			1500	0x5a24 (23076)
IPv4			326	0x5a24 (23076)
UDP	3702	45943	1500	0x5a25 (23077)
IPv4			1500	0x5a25 (23077)
IPv4			1500	0x5a25 (23077)
IPv4			326	0x5a25 (23077)
UDP	3702	45943	1500	0x5a26 (23078)
IPv4			1500	0x5a26 (23078)
IPv4			1500	0x5a26 (23078)
IPv4			326	0x5a26 (23078)
UDP	3702	45943	1500	0x5a27 (23079)
IPv4			1500	0x5a27 (23079)

```
▶ Frame 13110: 1514 bytes on wire (12112 bits), 1514 bytes captured (12112 bits
0000  f2 3c 91 1d c9 9f 50 87  89 40 a1 c1 08 00 45 00   ·<····P· ·@····E·
0010  05 dc 5a 24 20 00 3a 11  e4 8e b0 c0 6e f6 ac 69   ··Z$ ·:· ····n··i
0020  50 3e 0e 76 b3 77 12 8a  0e 99 3c 3f 78 6d 6c 20   P>·v·w·· ··<?xml
0030  76 65 72 73 69 6f 6e 3d  22 31 2e 30 22 20 65 6e   version= "1.0" en
0040  63 6f 64 69 6e 67 3d 22  55 54 46 2d 38 22 3f 3e   coding=" UTF-8"?>
0050  0a 3c 53 4f 41 50 2d 45  4e 56 3a 45 6e 76 65 6c   ·<SOAP-E NV:Envel
0060  6f 70 65 20 78 6d 6c 6e  73 3a 53 4f 41 50 2d 45   ope xmln s:SOAP-E
0070  4e 56 3d 22 68 74 74 70  3a 2f 2f 77 77 77 2e 77   NV="http ://www.w
0080  33 2e 6f 72 67 2f 32 30  30 33 2f 30 35 2f 73 6f   3.org/20 03/05/so
0090  61 70 2d 65 6e 76 65 6c  6f 70 65 22 20 78 6d 6c   ap-envel ope" xml
00a0  6e 73 3a 53 4f 41 50 2d  45 4e 43 3d 22 68 74 74   ns:SOAP- ENC="htt
00b0  70 3a 2f 2f 77 77 77 2e  77 33 2e 6f 72 67 2f 32   p://www. w3.org/2
00c0  30 30 33 2f 30 35 2f 73  6f 61 70 2d 65 6e 63 6f   003/05/s oap-enco
00d0  64 69 6e 67 22 20 78 6d  6c 6e 73 3a 78 73 69 3d   ding" xm lns:xsi=
00e0  22 68 74 74 70 3a 2f 2f  77 77 77 2e 77 33 2e 6f   "http:// www.w3.o
00f0  72 67 2f 32 30 30 31 2f  58 4d 4c 53 63 68 65 6d   rg/2001/ XMLSchem
0100  61 2d 69 6e 73 74 61 6e  63 65 22 20 78 6d 6c 6e   a-instan ce" xmln
0110  73 3a 78 73 64 3d 22 68  74 74 70 3a 2f 2f 77 77   s:xsd="h ttp://ww
0120  77 2e 77 33 2e 6f 72 67  2f 32 30 30 31 2f 58 4d   w.w3.org /2001/XM
0130  4c 53 63 68 65 6d 61 22  20 78 6d 6c 6e 73 3a 63   LSchema"  xmlns:c
0140  31 34 6e 3d 22 68 74 74  70 3a 2f 2f 77 77 77 2e   14n="htt p://www.
0150  77 33 2e 6f 72 67 2f 32  30 30 31 2f 31 30 2f 78   w3.org/2 001/10/x
0160  6d 6c 2d 65 78 63 2d 63  31 34 6e 23 22 20 78 6d   ml-exc-c 14n#" xm
0170  6c 6e 73 3a 77 73 75 3d  22 68 74 74 70 3a 2f 2f   lns:wsu= "http://
0180  64 6f 63 73 2e 6f 61 73  69 73 2d 6f 70 65 6e 2e   docs.oas is-open.
0190  6f 72 67 2f 77 73 73 2f  32 30 30 34 2f 30 31 2f   org/wss/ 2004/01/
```

If the number of services is large enough, the responses come in many full-sized packets with trailing UDP fragments.

To the left is a packet capture of a WS-Discovery amplification attack filtered to the output of one node.

Notice a one-packet request was replied to with nearly 40 full-sized and fragmented packets. Each blue line shows a complete UDP packet with each trailing fragment shown in white.

Topic / Item	Count	Average	Min val	Max val	Rate (ms)	Percent
▼ Packet Lengths	20000	979.33	60	1514	0.0607	100%
0-19	0	-	-	-	0.0000	0.00%
20-39	0	-	-	-	0.0000	0.00%
40-79	1467	68.16	60	79	0.0045	7.33%
80-159	3891	94.10	80	159	0.0118	19.45%
160-319	859	204.65	160	317	0.0026	4.29%
320-639	274	467.30	322	621	0.0008	1.37%
640-1279	2707	1166.02	640	1279	0.0082	13.54%
1280-2559	10802	1449.75	1280	1514	0.0328	54.01%
2560-5119	0	-	-	-	0.0000	0.00%
5120 and greater	0	-	-	-	0.0000	0.00%

Over half of all packets in this attack are between 1,280 and 1,514 bytes.

The rest of the attack is made up of various sizes of UDP fragments

Protocol	UDP Source Port	UDP Dest Port	IP Length	IPID
IPv4			54	0xc77c (51068)
IPv4			54	0xa8fa (43258)
UDP	64118	42704	1450	0x0000 (0)
UDP	50896	42704	1448	0x0000 (0)
UDP	59265	42704	1447	0x0000 (0)
UDP	59266	42704	1448	0x0000 (0)
UDP	59267	42704	1447	0x0000 (0)
UDP	56930	42704	1500	0xc77c (51068)
UDP	50896	42704	1448	0x0000 (0)
UDP	56930	42704	1500	0xa8fa (43258)
UDP	64118	42704	1450	0x0000 (0)
UDP	59266	42704	1448	0x0000 (0)
UDP	59267	42704	1447	0x0000 (0)
UDP	59265	42704	1447	0x0000 (0)
UDP	64118	42704	1450	0x0000 (0)
UDP	59267	42704	1447	0x0000 (0)
UDP	59266	42704	1448	0x0000 (0)
UDP	59265	42704	1447	0x0000 (0)
UDP	64118	42704	1450	0x0000 (0)
UDP	59267	42704	1447	0x0000 (0)
UDP	59266	42704	1448	0x0000 (0)
UDP	59265	42704	1447	0x0000 (0)

```
▶ Frame 6586: 1514 bytes on wire (12112 bits), 1514 bytes captured (12112 bits)
▶ Ethernet II, Src: Cisco_40:a1:c1 (50:87:89:40:a1:c1), Dst: f2:3c:91:1d:c9:9f
0000  f2 3c 91 1d c9 9f 50 87  89 40 a1 c1 08 00 45 00   ·<····P· ·@····E·
0010  05 dc a8 fa 20 00 33 11  69 a6 52 c5 00 04 ac 69   ···· ·3· i·R····i
0020  50 3e de 62 a6 d0 05 ea  03 d4 3c 3f 78 6d 6c 20   P>·b···· ··<?xml
0030  76 65 72 73 69 6f 6e 3d  22 31 2e 30 22 20 65 6e   version= "1.0" en
0040  63 6f 64 69 6e 67 3d 22  55 54 46 2d 38 22 20 3f   coding=" UTF-8" ?
0050  3e 0d 0a 3c 73 6f 61 70  3a 45 6e 76 65 6c 6f 70   >··<soap :Envelop
0060  65 20 78 6d 6c 6e 73 3a  78 73 69 3d 22 68 74 74   e xmlns: xsi="htt
0070  70 3a 2f 2f 77 77 77 2e  77 33 2e 6f 72 67 2f 32   p://www. w3.org/2
0080  30 30 31 2f 58 4d 4c 53  63 68 65 6d 61 2d 69 6e   001/XMLS chema-in
```

As with any protocol, not all implementations are alike. As such, there is a difference in behavior between the above node and the single node to the left.

Even though the scan packet was destined to UDP 3702, the node replies back from an ephemeral UDP port.

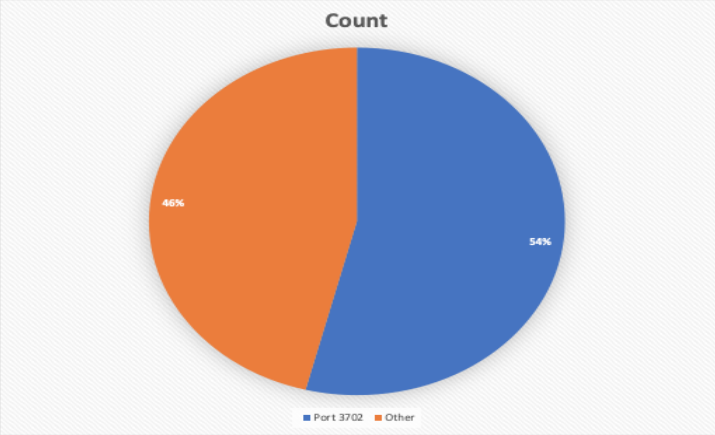

Recent scans performed by A10 Networks Research have yielded as many as 850k sources that respond to WS-Directory requests with a properly formatted Web Services response.

Some targets responded with UDP source port 3702 but slightly less than half of targets responded from the ephemeral range.

Since the source port is not deterministic, this makes blocking based on UDP port less effective.

OVERALL SCOPE

We decided to get a full understanding of the size and overall scope of this weapon by compiling its geo-location and the organization ID of the owner of the source IP as gleaned through BGP. Then, we dug a bit deeper to understand the stated equipment and location by the devices themselves through the WS-Discovery reply.

GEOGRAPHIC BREAKDOWN

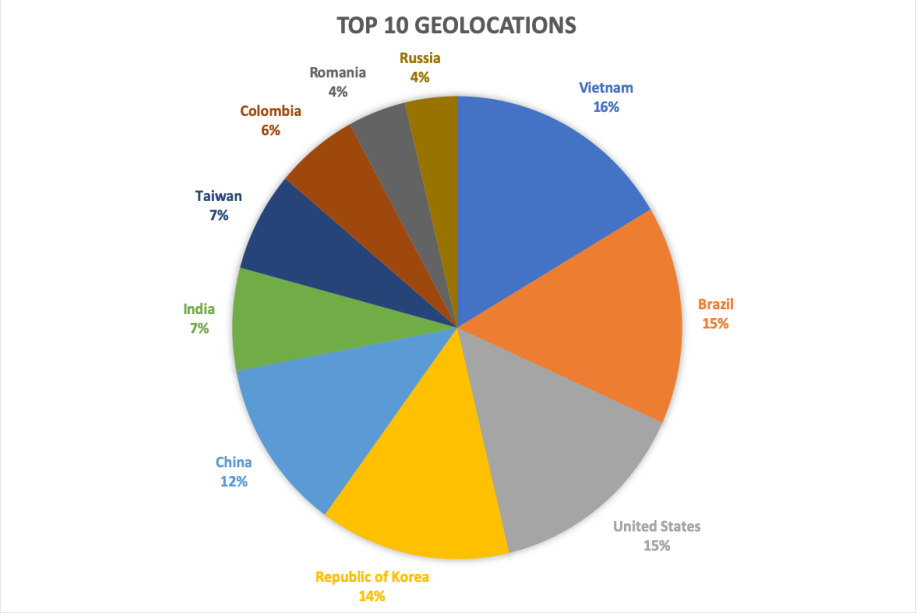

The top-five locations are a bit too close to call as Vietnam, Brazil, U.S., Korea, and China all come within two percentage points of each other. Interestingly, the location string stated within the Web Services response itself is overwhelmingly China, which shows the location that the offending device was manufactured.

BREAKDOWN BY BGP ORGANIZATION ID

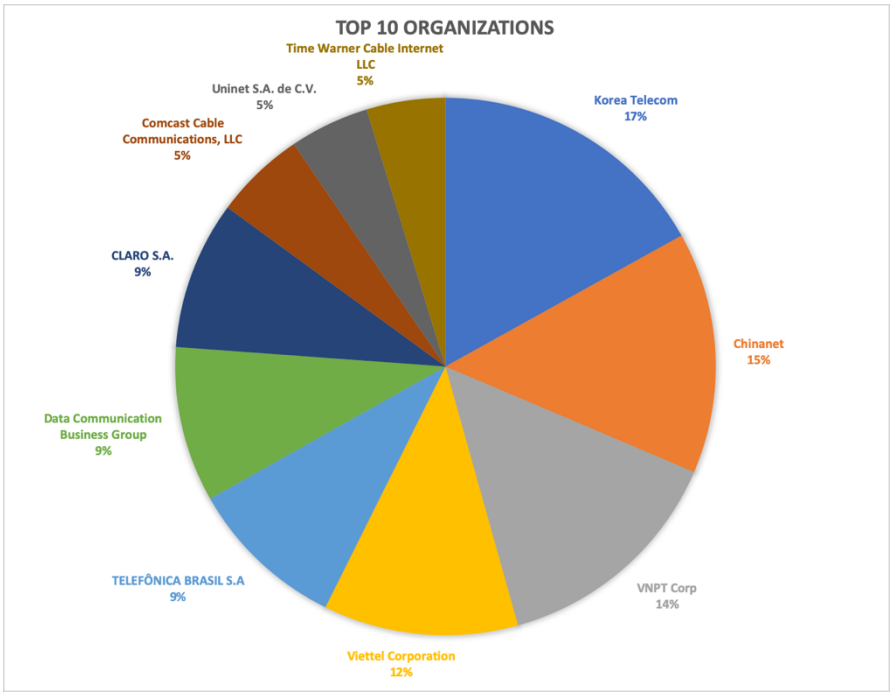

The top-five organizations, however, have a similar, but not the same, breakdown as their geographies. Here, Korea Telecom, Chinanet, VNPT Corp, Viettel, and Telefonica Brazil take the top honors, while U.S. carriers are diffused across many smaller ASNs.

IMPORTANT FACTS ABOUT WSD

1. WS-Discovery functions like SSDP are service discovery protocols. The WS-Discovery specification came much later in 2009 and the responses are in SOAP unlike the case of SSDP, which would give a HTTP like response.
2. 15 percent of the hosts that are potential WS-Discovery amplifiers respond to SSDP, as well. This means that the devices in the network use multiple service discovery protocols.
3. 160,000 devices, 60 percent of which are IP cameras and DVRs, are exposed to the internet because of the WS-Discovery protocol. This makes them susceptible to being hijacked to become botnets.
4. Here is a list of the top-three manufacturers of the IP cameras and DVRs along with their numbers:

MANUFACTURER	NUMBER OF DEVICES
Dahua	112K
IntelBras	42K
Hikvision	31K

5. The fact that Dahua and Hikvision are on top of the list is supported by the fact that 99 percent of these devices have a ONVIF geo-location of China.
6. The IP geo-location distribution graph shows that Brazil houses 15 percent of these WS-Discovery sources, which is backed by the fact that IntelBras is number-two on the list of top manufacturers of the hardware.

PROTECTING YOURSELF FROM THE THREAT

Here we will outline various ways that organizations can to protect themselves from these attacks, along their strengths and weaknesses.

BLOCKING AT LAYER 4 ONLY (THE ACCESS LIST OR FLOWSPEC METHOD)

Unlike many amplification protocols such as DNS and NTP, which send attack packets on deterministic source ports (UDP 53 and 123 respectively), WS-Discovery attack packets happen over a combination of UDP 3702 and the ephemeral range of UDP ports. Using the ephemeral range is non-deterministic in nature and blocking it will create too much collateral damage for any service that is being protected, those that are real UDP requests. As such, blocking this particular attack at Layer 4 only cannot be done.

BLOCKING UDP FRAGMENTS AT ATTACK TIME

Because UDP fragments (any packet with the MF bit set or a fragment offset not equal to 0) make up a large portion of this type of attack, one strategy that could be undertaken is blocking all UDP fragments at attack time. This would be an effective method to reduce the attack by a large amount but becomes problematic with all of the valid UDP fragmentation that occurs in the network or streaming in from the internet. Depending on the service being offered or the way the network uses DNS, this approach could be a problem and should be avoided due to unpredictable bad behavior. If low collateral damage is the goal, then this method should be avoided.

BLACKLISTING DURING THE ATTACK

Although the attack starts with spoofing the victim's IP address, each of these amplification sources are real IP addresses and not spoofed. As such, there is a finite number of these systems that could be attacking the service at any point in time. As discussed above, A10 Networks Research maintains an up-to-date list of WS-Discovery-enabled systems (over 800k at this writing) and can provide this as a means to block these systems. As with any scan-based data set, however, the number received depends highly on external factors including the time of day, provider changes in DHCP leases and so on. To make up for this, we also recommend that the next method should be used in conjunction.

A10 NETWORKS ZERO-DAY AUTOMATED PATTERN RECOGNITION (ZAPR)

A10 Networks created this functionality for this exact problem. As attack patterns tend to be ephemeral, any amount of up-front preparation will still have cracks and some of the attack will leak through. To ensure that this is not the case, we created a machine-learning-based system to look at each packet in an attack, understand the important features of the packet and analyze and cluster it together to find the appropriate pattern to block. As this happens without the need for a peacetime data set, it is perfect for these sorts of situations.

CONCLUSION

It is our feeling that understanding the functionality, size, and overall scope of a DDoS weapon is important information to have before attempting mitigation. This advisory provides a good foundation of the steps to take to protect the network from the WS-Discovery threat, as well as the upsides and downsides of the various mitigation options.

ABOUT A10 NETWORKS

A10 Networks (NYSE: ATEN) provides Reliable Security Always™, with a range of high-performance application networking solutions that help organizations ensure that their data center applications and networks remain highly available, accelerated and secure. Founded in 2004, A10 Networks is based in San Jose, Calif., and serves customers globally with offices in more than 80 countries worldwide. For more information, visit: a10networks.com and @A10Networks.