TLP GREEN

7.1.15

Threat Advisory: RIPv1 Reflection DDoS

Risk Factor - Medium

1.1 / Overview / PLXsert has been monitoring an uptick in a form of DDoS reflection previously thought of as mostly abandoned. This attack vector, which involves the use of an outdated routing protocol in RIPv1, began showing up in active campaigns again on May 16th after being dormant for more than a year. The latest attacks observed, as described later, are apparently making use of only a small number of available RIPv1 source devices.

RIPv1 was first introduced in 1988 under RFC1058, which is now listed as a historic document in RFC1923. The historic designation means the original RFC is actively deprecated. One main reason for this is that RIPv1 only supports classful networks. So if the network advertised by RIPv1 happens to be a class A network, such as 10.1.2.0/24, this will be sent in an advertisement as 10.0.0.0/8. This, among other things, further limits the usefulness for RIPv1 as a viable option for internal networks, much less the Internet.

Highlighted Campaign Attributes /

- Peak bandwidth: 12.8 Gigabits per second
- Peak packets per second: 3.2 Million Packets Per Second
- Attack Vector: RIPv1 reflection and amplification
- Source port : 520
- Destination port: random

1.2 / RIPv1 Overview / Routing Information Protocol version 1 has been available for many years now. RIPv1 is considered to be a quick and easy way to dynamically share route information in a small multi-router network.

A typical router communication would appear as shown in the table below. A request is sent by a router running RIP when it is first configured or powered on. Any other device listening for requests will respond to this request with a list of routes. Updates are also sent periodically as broadcasts.

To leverage the behavior of RIPv1 for DDoS reflection, a malicious actor can craft the same request query type as above, which is normally broadcast, and spoof the IP address source to match the intended attack target. The destination would match an IP from a list of

```
Router initial request for routes(sent as broadcast):
15:53:50.015995 IP 192.168.5.2.520 > 255.255.255.255.520: RIPv1, Request, length: 24

Listening router response for routes(sent as a unicast reply to request IP):
15:53:50.036024 IP 192.168.5.1.520 > 192.168.5.2.520: RIPv1, Response, length: 24

Regular periodic update sent every 30 seconds by default(broadcast):
15:54:26.448383 IP 192.168.5.1.520 > 255.255.255.255.520: RIPv1, Response, length: 24
```

Figure 1 - Normal communication for RIPv1

known RIPv1 routers on the Internet. Based on recent attacks, attackers prefer routers that seem to have a suspiciously large amount of routes in their RIPv1 routing table. This query results in multiple 504-byte payloads sent to a target IP per a single request. The multiple responses are also a result of the 25-route max that can be contained in an RIP packet. The payloads in the next figure are from an actual attack against an Akamai customer. At the target site, the only traffic visible are the unsolicited responses to RIPv1 queries as shown. The replies all source from udp port 520 used by RIP.

Payload Samples

A typical RIPv1 request contains a 24-byte payload. The responses above contain 504 bytes. This particular reflector responds with ten 504-byte payloads and one payload of 164 bytes.

```
Attack Traffic:
20:30:23.606170 IP X.X.X.X.520 > Y.Y.Y.Y.55735: RIPv1, Response, length: 504
20:30:23.606173 IP X.X.X.X.520 > Y.Y.Y.Y.16235: RIPv1, Response, length: 504
20:30:23.606174 IP X.X.X.X.520 > Y.Y.Y.Y.16235: RIPv1, Response, length: 504

Attack Payload detail(varies depending on router network):
RIPv1, Response, length: 504, routes: 25
         172.99.0.0, metric: 1
         172.99.1.0, metric: 1
         172.99.2.0, metric: 1
         172.99.3.0, metric: 1
          <snip>
RIPv1, Response, length: 504, routes: 25
RIPv1, Response, length: 504, routes: 25
RIPv1, Response, length: 504, routes: 25
RIPv1, Response, length: 504, routes: 25
RIPv1, Response, length: 504, routes: 25
RIPv1, Response, length: 504, routes: 25
RIPv1, Response, length: 504, routes: 25
RIPv1, Response, length: 504, routes: 25

RIPv1, Response, length: 164, routes: 8
         172.99.250.0, metric: 1
         172.99.251.0, metric: 1
         172.99.252.0, metric: 1
         172.99.253.0, metric: 1
         172.99.254.0, metric: 1
         172.99.255.0, metric: 1
         X.0.0.0, metric: 1
         0.0.0.0, metric: 2
```

Figure 2 - Attack traffic contains RIPv1 routes from queried router producing a total of 5,204 byte payload

When calculating the amplification factor, including IPv4 headers [IP(IP (10)UDP10) UDP(8)], the resulting amplification for a single RIPv1 request is 131.24 (over 13,000%) based on the above response. This varies depending on the number of routes in the router's table. During the first attack, a good majority of the amplifiers contained multiple 504-byte replies as shown in the next chart.

Non-intrusive scanning of attack sources reveals that the victims being leveraged for RIP reflection are likely not using enterprise-grade routing hardware.

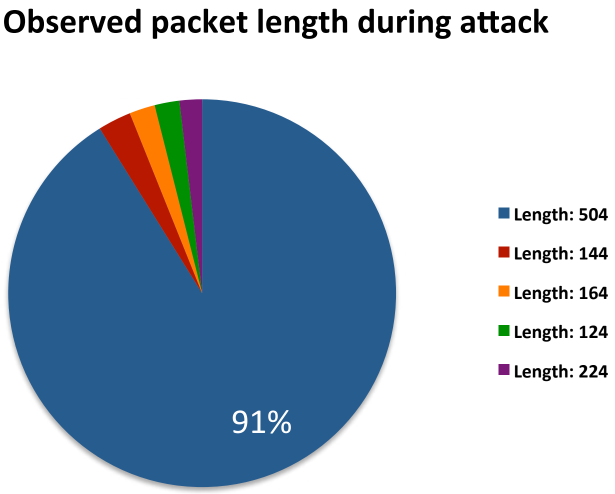
Observed packet length during attack

Figure 3 - Payload length analysis based on captured attack traffic sample

Most were detected as running custom SOHO router firmware (DD-WRT) or NAS devices like BlueArc Titan. Later scanning on an Internet scale resulted in similar findings with regards to the use of SOHO routers. Either way, the devices are not completely at blame here. RIPv1 is working as designed and malicious actors continue to exploit this method for reflection and amplification.

1.3 / RIPv1 Poisoning / Would a malicious actor be able to increase the amplification factor by forcing the router to learn extra routes? Although the concept is feasible, three main factors impede this scenario from being effectively used in a DDoS attack

The first factor is split horizon, a default on some devices running RIPv1. Simply put, the router receiving the crafted route update would not send that same update back through the interface from which it was received. This means an Internet connected device would not respond back to a target with the newly injected networks. It will, however, provide an update of those routes once they become stale with a metric of 16.

This leads to the next mitigating factor, route tables are often purged and will not keep stale routes in their table for an extended period of time. Once they expire, another update is required in order for the router to keep these routes. In Cisco devices for example an injected route wouldn't show up until it became unusable (metric 16) and would only persist for 1 minute after being flagged as unusable at which point it will be purged from the routing table. This means continuous and timed updates would be needed to keep tables poisoned. These repeated updates would affect the amplification, since instead of the one 24-byte payload, multiple 504-byte or smaller payloads (depending on number of routes injected) are being sent to maintain the routes as well as reflect the responses.

The third issue is RIPv1 responses originating from networks not directly connected to the router will be ignored. This means updates would need to be spoofed to match a network directly connected to the router for the update to even be accepted.

With all of these possible hurdles in the way, it is unlikely that poisoning a RIPv1 route table would be worth the effort and overhead for a malicious actor.

If local access to the device were possible, then the routes could be manipulated depending on the device's capabilities. This may be a possible scenario if the router is left with default or no credentials. This is a simple configuration oversight that unfortunately occurs more often than it should.

The example below is a single RIPv1 response from a lab router that has been populated with 255 route entries via an unsecured administrative interface.

```
Lab router response:
14:29:37.010393 IP 10.0.20.1.520 > 10.0.20.16.520: RIPv1, Response, length: 504
14:29:37.010559 IP 10.0.20.1.520 > 10.0.20.16.520: RIPv1, Response, length: 504
14:29:37.010630 IP 10.0.20.1.520 > 10.0.20.16.520: RIPv1, Response, length: 504
14:29:37.010985 IP 10.0.20.1.520 > 10.0.20.16.520: RIPv1, Response, length: 504
14:29:37.011117 IP 10.0.20.1.520 > 10.0.20.16.520: RIPv1, Response, length: 504
14:29:37.011227 IP 10.0.20.1.520 > 10.0.20.16.520: RIPv1, Response, length: 504
14:29:37.011336 IP 10.0.20.1.520 > 10.0.20.16.520: RIPv1, Response, length: 504
14:29:37.011446 IP 10.0.20.1.520 > 10.0.20.16.520: RIPv1, Response, length: 504
14:29:37.011560 IP 10.0.20.1.520 > 10.0.20.16.520: RIPv1, Response, length: 504
14:29:37.011683 IP 10.0.20.1.520 > 10.0.20.16.520: RIPv1, Response, length: 504
14:29:37.011822 IP 10.0.20.1.520 > 10.0.20.16.520: RIPv1, Response, length: 84
Verbose output:
14:29:37.010393 IP (tos 0xc0, ttl 255, id 0, offset 0, flags [none], proto UDP (17),
length 532)
    10.0.20.1.520 > 10.0.20.16.520:
       RIPv1, Response, length: 504, routes: 25
         192.168.1.0, metric: 1
         <snip>
14:29:37.010559 IP (tos 0xc0, ttl 255, id 0, offset 0, flags [none], proto UDP (17),
length 532)
    10.0.20.1.520 > 10.0.20.16.520:
       RIPv1, Response, length: 504, routes: 25
         192.168.26.0, metric: 1
         <snip>
14:29:37.010630 IP (tos 0xc0, ttl 255, id 0, offset 0, flags [none], proto UDP (17),
length 532)
    10.0.20.1.520 > 10.0.20.16.520:
       RIPv1, Response, length: 504, routes: 25
       <snip>
14:29:37.010985 IP (tos 0xc0, ttl 255, id 0, offset 0, flags [none], proto UDP (17),
length 532)
    10.0.20.1.520 > 10.0.20.16.520: RIPv1, Response, length: 504, routes:
25192.168.76.0, metric: 1
         <snip>
```

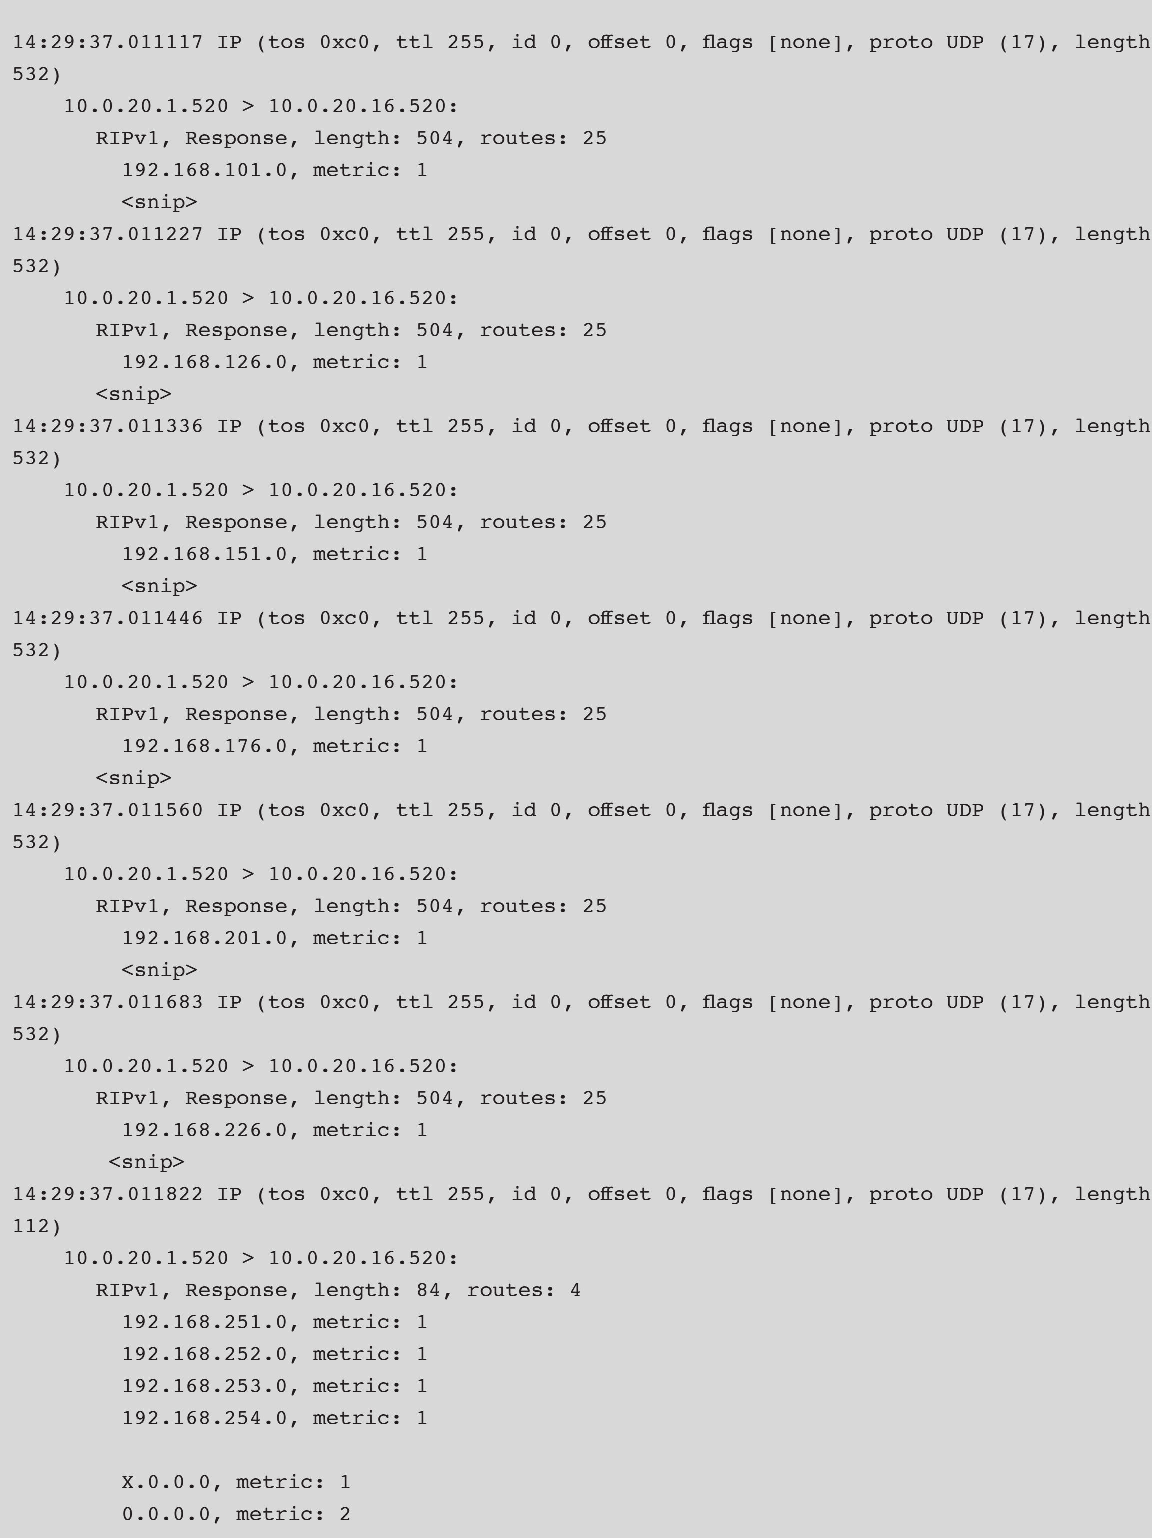

```
14:29:37.011117 IP (tos 0xc0, ttl 255, id 0, offset 0, flags [none], proto UDP (17), length
532)
    10.0.20.1.520 > 10.0.20.16.520:
      RIPv1, Response, length: 504, routes: 25
        192.168.101.0, metric: 1
        <snip>
14:29:37.011227 IP (tos 0xc0, ttl 255, id 0, offset 0, flags [none], proto UDP (17), length
532)
    10.0.20.1.520 > 10.0.20.16.520:
      RIPv1, Response, length: 504, routes: 25
        192.168.126.0, metric: 1
      <snip>
14:29:37.011336 IP (tos 0xc0, ttl 255, id 0, offset 0, flags [none], proto UDP (17), length
532)
    10.0.20.1.520 > 10.0.20.16.520:
      RIPv1, Response, length: 504, routes: 25
        192.168.151.0, metric: 1
        <snip>
14:29:37.011446 IP (tos 0xc0, ttl 255, id 0, offset 0, flags [none], proto UDP (17), length
532)
    10.0.20.1.520 > 10.0.20.16.520:
      RIPv1, Response, length: 504, routes: 25
        192.168.176.0, metric: 1
      <snip>
14:29:37.011560 IP (tos 0xc0, ttl 255, id 0, offset 0, flags [none], proto UDP (17), length
532)
    10.0.20.1.520 > 10.0.20.16.520:
      RIPv1, Response, length: 504, routes: 25
        192.168.201.0, metric: 1
        <snip>
14:29:37.011683 IP (tos 0xc0, ttl 255, id 0, offset 0, flags [none], proto UDP (17), length
532)
    10.0.20.1.520 > 10.0.20.16.520:
      RIPv1, Response, length: 504, routes: 25
        192.168.226.0, metric: 1
       <snip>
14:29:37.011822 IP (tos 0xc0, ttl 255, id 0, offset 0, flags [none], proto UDP (17), length
112)
    10.0.20.1.520 > 10.0.20.16.520:
      RIPv1, Response, length: 84, routes: 4
        192.168.251.0, metric: 1
        192.168.252.0, metric: 1
        192.168.253.0, metric: 1
        192.168.254.0, metric: 1

        X.0.0.0, metric: 1
        0.0.0.0, metric: 2
```

Figure 4 - Lab router RIPv1 response simulating response used in live attack

1.4 / Highlighted Campaign / The attack on May 16th, 2015 peaked at 12.81 Gbps and 3.2 Mpps. Looking at the breakdown of bandwidth by location in the next figure, the larger amounts of traffic were observed from European sources. Combined, London and Frankfurt saw 4.75 Gbps peak traffic.

Attack distribution by scrubbing center RIPv1 Reflection attack

Data Center	Peak Bits Per Second	Peak Packets Per Second
Hong Kong	228.00 Mbps	57.00 Kpps
DCA2	3.38 Gbps	839.00 Kpps
SJC2	1.82 Gbps	457.00 Kpps
FRA2	1.80 Gbps	459.00 Kpps
LON2	2.95 Gbps	743.00 Kpps
TYO2	2.63 Gbps	651.00 Kpps

Figure 5 - Attack distribution for largest RIPv1 reflection attack in May 2015

1.5 / RIPv1 Internet scan results /A total of 53,693 devices on the Internet responded to RIPv1 queries. Although many of these sources would not be suitable as DDoS amplification sources, they are still otherwise vulnerable to reflection and other attacks due to the inherently weak security offered by RIPv1. In the active attacks observed against Akamai customers, it appears that only a handful of the 53,693 possible sources are being leveraged, as only about 500 unique sources were identified in attack-traffic samples. However, since a majority of these sources sent packets predominantly of the 504-byte size, it's pretty clear as to why they were leveraged for attack purposes. As attackers discover more sources, it is possible that this vector has the potential to create much larger attacks than what we've observed thus far.

If poisoning of the routes was possible as described in the previous section, these mostly unused devices could be leveraged in larger and more distributed attacks. Right now, most of the 53,693 possible sources respond with one unique route — making them regular DDoS reflection sources without additional amplification.

Based on the sizes of RIPv1 responses we received from the Internet, there are plenty of devices that offer at least some factor of amplification. We identified 24,212 devices on the Internet that offered at least an 83% amplification rate. Below are the top 5 associated packet lengths of replies received during testing.

Response Packet Length (Bytes)	Number Of Packets Received
24	38,298
44	27,401
504	3,224
64	2,699
84	583

Another interesting finding is that RIPv1 implementations found in the wild can be sent an improperly structured RIPv1 packet which will cause them to respond with a reply that doesn't contain their real routing table information. While this kills amplification and downgrades the attack to a simple reflection, it could be leveraged to diversify attack traffic from a single source.

```
// VALID REQUEST & RESPONSE
17:23:36.743092 IP (tos 0x0, ttl 64, id 1, offset 0, flags [none], proto UDP (17), length
52)
    XX.XX.XX.XX.31337 > YY.YY.YY.YY.520:
        RIPv1, Request, length: 24, routes: 1
          AFI 0, 0.0.0.0, metric: 16
17:23:36.835398 IP (tos 0x0, ttl 51, id 0, offset 0, flags [DF], proto UDP (17), length 92)
    YY.YY.YY.YY.520 > XX.XX.XX.XX.31337:
        RIPv1, Response, length: 64, routes: 3
          192.168.4.0, metric: 1
          YY.YY.YY.0, metric: 1
          0.0.0.0, metric: 1

// MALFORMED REQUEST & RESPONSE
17:23:36.803073 IP (tos 0x0, ttl 64, id 1, offset 0, flags [none], proto UDP (17), length
52)
    XX.XX.XX.XX.1337 > YY.YY.YY.YY.520:
        RIPv1, Request, length: 24, routes: 1
          AFI 0, 0.0.0.0, metric: 4096
17:23:36.895476 IP (tos 0x0, ttl 51, id 0, offset 0, flags [DF], proto UDP (17), length 52)
    YY.YY.YY.YY.520 > XX.XX.XX.XX.1337:
        RIPv1, Response, length: 24, routes: 1
          AFI 0, 0.0.0.0, metric: 16
```

Example of using malformed requests to trigger different responses

The IPs discovered were broken down further in order to determine models for devices or if particular ISPs are handing out devices with these configurations in place. From the list of 53,693 responding IPs, over 20,000 were also listening on TCP port 80 (typically for a web admin interface). We attempted to fetch and log these requests looking for information leaks (WWW-Auth realm, model numbers in landing pages, HTTP headers, etc.) that would help us identify some of the devices affected. The following figure displays the top 3 models discovered running RIPv1. The Netopia devices are likely older hardware deployed during the initial boom in ADSL broadband Internet service.

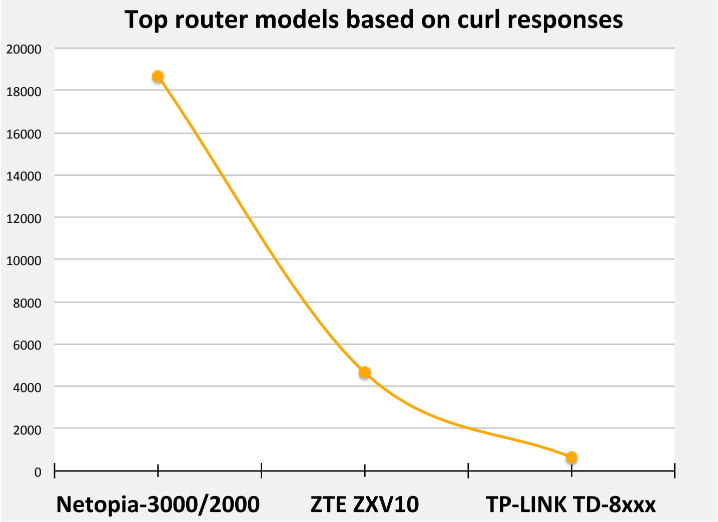

Figure 6 - Top 3 routers confirmed to be running RIPv1 and running a web interface

The next figure shows the breakdown of whois information for the Netopia source IPs. These devices are mostly provided by or used by AT&T customers.

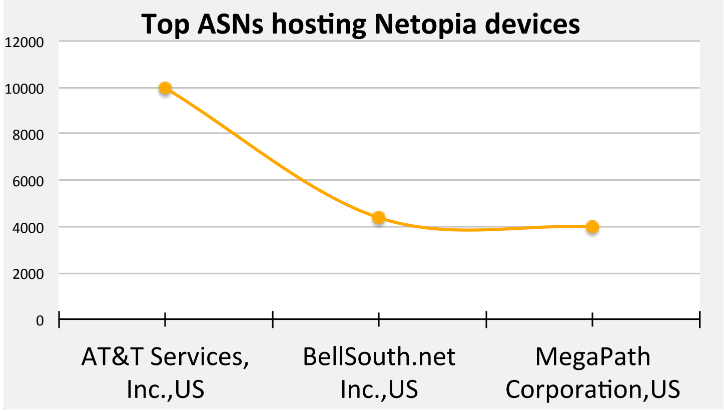

Figure 7 - Top 3 Internet Service Providers identified

The complete list of 53,693 available sources was mostly in the USA, as depicted in the next figure. As mentioned earlier, the sources confirmed to be used in recent attack campaigns were mostly based out of Europe. This leaves a lot of potentially untapped resources that could fall victim to abuse in amplification and reflection attacks.

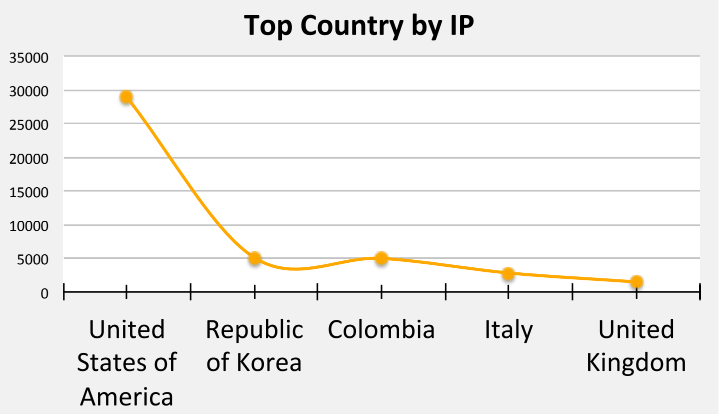

Figure 8 - Top 5 source countries for all devices responding to RIPv1

The breakdown of sources by country as leveraged in DDoS attacks is provided in the next figure. There is a common factor among the selected attack sources — they each have an unusually large amount of routes returned to each RIPv1 query. The payloads provided in Figure 2 are an example of this.

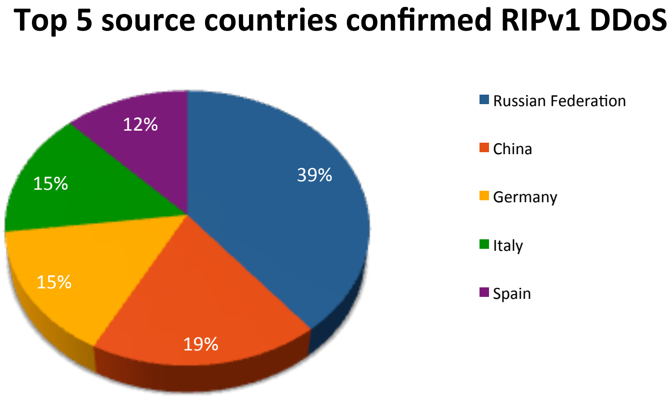

Figure 9 - Top 5 countries observed during largest RIPv1 DDoS reflection attack

1.6 / Recommended Mitigation / For victims of this attack, i.e. reflector sources, there are several ways to avoid this attack method:

- Switch to RIPv2 or later and enable authentication
- If RIPv1 is required, assess the need to expose RIP on your WAN interface. If it's not needed the WAN side interface should be marked as a passive interface (where supported).
- Access to RIP can also be restricted via ACL to only allow known neighbor routers.

For targets of an RIPv1 reflected DDoS attack, an ACL can be used to restrict UDP source port 520 from the Internet. If the attack is too large, it may require a DDoS mitigation provider such as Akamai Technologies.

1.7 / Conclusion / The list of available reflection vectors is by no means small, and some vectors have proven more difficult to keep under control than others due to their pervasive nature (i.e. DNS, SSDP). That being said, there is little reason for RIPv1 to continue as an available resource for DDoS attacks. Most of these sources appear to be from outdated hardware that has been running in home or small-office networks for years.

The large amount of Netopia devices on the list is a good indicator of this. ISPs will tend to leave the original hardware used to deploy service in place for years without pushing for end users to upgrade their devices/infrastructure. If there is no issue with the devices, why replace them? Over time this leaves devices on the Internet that expose outdated protocols such as RIPv1 and possibly leaves devices with exposing vulnerabilities in their software. In this case, the ISPs would likely have the biggest impact on cleanup efforts. For starters, UDP port 520 should not be exposed on the Internet. This would greatly mitigate the risk of abuse in DDoS reflection attacks. Another option, although more costly, would be to provide consumers with upgrade options on old modems/routers that were provided. This is especially true for devices that are no longer supported by the manufacturer.

PLXsert is currently monitoring ongoing campaigns. Future advisories and updates will be provided if warranted.

The Prolexic Security Engineering and Research Team (PLXsert) monitors malicious cyber threats globally and analyzes these attacks using proprietary techniques and equipment. Through research, digital forensics and post-event analysis, PLXsert is able to build a global view of security threats, vulnerabilities and trends, which is shared with customers and the security community. By identifying the sources and associated attributes of individual attacks, along with best practices to identify and mitigate security threats and vulnerabilities, PLXsert helps organizations make more informed, proactive decisions.

Akamai® As the global leader in Content Delivery Network (CDN) services, Akamai makes the Internet fast, reliable and secure for its customers. The company's advanced web performance, mobile performance, cloud security and media delivery solutions are revolutionizing how businesses optimize consumer, enterprise and entertainment experiences for any device, anywhere. To learn how Akamai solutions and its team of Internet experts are helping businesses move faster forward, please visit www.akamai.com or blogs.akamai.com, and follow @Akamai on Twitter.

Akamai is headquartered in Cambridge, Massachusetts in the United States with operations in more than 57 offices around the world. Our services and renowned customer care are designed to enable businesses to provide an unparalleled Internet experience for their customers worldwide. Addresses, phone numbers and contact information for all locations are listed on www.akamai.com/locations.

©2015 Akamai Technologies, Inc. All Rights Reserved. Reproduction in whole or in part in any form or medium without express written permission is prohibited. Akamai and the Akamai wave logo are registered trademarks. Other trademarks contained herein are the property of their respective owners. Akamai believes that the information in this publication is accurate as of its publication date; such information is subject to change without notice. Published 06/15.