

Taming the 800 Pound Gorilla: The Rise and Decline of NTP DDoS Attacks

Jakub Czyz
University of Michigan

Christos Papadopoulos
Colorado State University

Michael Kallitsis
Merit Network

Michael Bailey
University of Michigan
University of Illinois

Manaf Gharaibeh
Colorado State University

Manish Karir
Merit Network

ABSTRACT

Distributed Denial of Service (DDoS) attacks based on Network Time Protocol (NTP) amplification, which became prominent in December 2013, have received significant global attention. We chronicle how this attack rapidly rose from obscurity to become the dominant large DDoS vector. Via the lens of five distinct datasets, we characterize the advent and evolution of these attacks. Through a dataset that measures a large fraction of global Internet traffic, we show a three order of magnitude rise in NTP. Using a large darknet, we observe a similar rise in global scanning activity, both malicious and research. We then dissect an active probing dataset, which reveals that the pool of amplifiers totaled 2.2M unique IPs and includes a small number of “mega amplifiers,” servers that replied to a single tiny probe packet with gigabytes of data. This dataset also allows us, for the first time, to analyze global DDoS attack victims (including ports attacked) and incidents, where we show 437K unique IPs targeted with at least 3 trillion packets, totaling more than a petabyte. Finally, ISP datasets shed light on the local impact of these attacks. In aggregate, we show the magnitude of this major Internet threat, the community’s response, and the effect of that response.

Categories and Subject Descriptors

C.2.3 [Computer-Communication Networks]: Network Operations—*Network management, Network monitoring*

General Terms

Measurement, Networking, Security

Keywords

DDoS; NTP; Darknet

1. INTRODUCTION

Though Distributed Denial of Service (DDoS) attacks have become increasingly commonplace, the scope and magnitude of the firepower that can now be unleashed on a target is unprecedented. A

comprehensive recent study described various DDoS amplification attack vectors and their attack potential, predicting the rise of Network Time Protocol (NTP) [23] as one of the most potent [32]. In the weeks since, data published by efforts such as the OpenNTPProject.org [2] painted a bleak picture of millions of devices that were readily available to facilitate these attacks. Indeed, in the first quarter of 2014, 85% of DDoS attacks larger than 100 Gbps were using NTP reflection [24]. Used for everything from a gaming edge [18] to extortion [28], DDoS attacks have been a scourge for years, but this new attack vector has made them even more powerful.

We study the dramatic rise of NTP-based DDoS attacks, in a span of just a few months, from a mere trickle to rapidly become the primary DDoS vector for large attacks. We show who the amplifiers are, what they are capable of, who they attack, and what attacks looked like. We present findings from the vantage points of five unique datasets: a global traffic profile and attack dataset from a large analytics and DDoS mitigation vendor; a global probing dataset that sheds light on potential amplifiers, their victims, and attacks; a large darknet that observes Internet-scale scanning activities; and, finally, two regional Internet service providers that give us a local perspective on attack mechanics and evolution.

NTP DDoS Mechanics: Much has been written on the general ideas of reflection and amplification (e.g., [27], [32]), so, we discuss attack mechanics only briefly. For these attacks, the first step is the identification of vulnerable *amplifiers*, which attackers can accomplish via large-scale scanning. An amplifier is simply a host running a protocol (e.g., NTP, DNS) which, when sent a query packet, responds with one or more packets whose aggregate size is larger than the query it received. Once suitable amplifiers have been identified, an attacker, directly or via intermediate hosts he controls, sends small UDP packets with the spoofed *source* address set to the IP of the intended attack victim and the destination address a vulnerable amplifier. Such spoofing is possible because many networks do not follow best security practices (e.g., BCP 38/84) [36]. In turn, as is the goal of these *volumetric* DDoS attacks, large amounts of traffic from amplifiers may saturate bandwidth at the victim. We say these attacks are *reflected* as they are executed via an intermediate host (the amplifier); attacks are *amplified* in that more bandwidth is used at the victim than the attacker (or its bots) need to expend. In the case of NTP, the protocol feature that has been used in large attacks is the *monlist* command. A command intended only for diagnostics, it returns the last 600 clients of the amplifier, producing a typically very large, multi-packet reply to a single small query packet—an ideal amplification attack vector.

The rest of our paper and highlighted findings are as follows. In § 2 we describe a rapid rise of NTP traffic to become the most potent DDoS attack vector. Only constituting 0.001% of traffic four months earlier, by mid-February NTP had climbed to use up 1% of global

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.
IMC'14, November 05–07, 2014, Vancouver, BC, Canada
Copyright 2014 ACM 978-1-4503-3213-2/14/11\$15.00
<http://dx.doi.org/10.1145/2663716.2663717>.

Internet traffic, surpassing even the ubiquitous DNS. Simultaneously, the majority of large DDoS attacks observed globally were now using the NTP vector, one not even on the radar in November. In § 3 we explore the amplifier pool that enabled this rapid rise in traffic and attacks—a pool of 2.2M vulnerable NTP servers. In this pool we find several surprises: first, while a typical server provides just 4x amplification, we find a subset we term “mega amplifiers,” who respond to a single packet with gigabytes of data; second, we find that a large fraction of amplifiers are not infrastructure servers but end hosts; comically, we also discover that nearly a fifth of NTP servers do not serve the correct synchronized time.

As for target victims, in § 4 we seek to answer the questions of who is being DDoSed, when, and how badly. We find tens of thousands of victim IPs across thousands of autonomous systems, and discover that many appear to be gamers. The 437K unique victims we saw were hit over fifteen weeks with at least 3 trillion packets, which is, at median amplifier response, over 1.2 petabytes. Luckily, things are improving, which we begin to outline in § 5 and § 6 where we first show a dramatic increase in scanning for NTP over eight months as attackers scramble to build lists of vulnerable IPs. Then, we discuss how the pool of NTP monlist amplifiers being scanned for has itself decreased rapidly, especially relative to other amplifier pools.

Zooming in from the global perspective to the local, in § 7 we highlight the effect on and mitigation at two regional networks that actively worked to respond to attacks, including local confirmation of some of our global dataset findings and several new insights. Finally, in § 8 we describe some related work, and in § 9 we summarize key conclusions and future work.

2. GLOBAL NTP TRAFFIC AND ATTACKS

2.1 Global Internet NTP Traffic

In late 2013 and early 2014, public attack disclosures (e.g., [12], [18], and [30]) suggested that NTP-based DDoS attacks were very large and increasingly common. To gauge the global prevalence of these types of attacks, we obtained traffic statistics from Arbor Networks, a provider of network analytics and attack mitigation services [9]. Arbor Networks collects traffic data, via appliances that export network flow statistics, from a global set of over 300 generally large and typically global network operators, including tier-1 ISPs, tier-2/regional ISPs, mobile service providers, as well as two dozen large enterprise customers. Arbor Networks estimates that their netflow traffic dataset represents between a third and a half of all Internet traffic.

In Figure 1 we show the relative fraction of all measured Internet traffic that NTP and DNS represent, observed by Arbor over six months, between November 1st, 2013 and May 1st, 2014. The data lines are the ratio of each protocol’s daily bits-per-second averages to all Internet traffic seen. The overall daily average of Internet traffic represented in this dataset is 71.5 Tbps. As the figure shows, NTP starts this period off at a level that constitutes only about 0.001% of daily bits per second transferred. By March 1st, NTP had grown to use up about 1% of of global Internet traffic, surpassing even DNS, which hovers at around 0.15% of traffic throughout this period. The dramatic three order of magnitude rise of NTP traffic in both absolute and relative terms, translates to noticeable financial impact, even at service providers who host not only victims but the NTP amplifiers used in attacks. It is possible that this impact is part of the reason for rapid remediation, which we discuss in § 6. After peaking on February 11th, NTP traffic appears to have begun a decline and is at around 0.1% of Internet traffic at the beginning

of May, still two orders of magnitude higher than at the start of November 2013, but a tenth of what it was at peak.




Figure 1: Fraction of Internet traffic that is NTP and DNS for November 2013 – April 2014. There is a nearly three order of magnitude growth in the fraction of global traffic that is NTP, peaking at 1% of all traffic in mid-February, but then dropping, to around 0.1%.

2.2 Global NTP DDoS Attacks

In addition to the traffic statistics described in § 2.1, the Arbor Networks devices also collect labeled attack counts, which detail the types of attacks that are being seen by Arbor Networks’ customers, including the prevalent protocols used in the attacks (e.g., ICMP flooding, TCP SYN flooding, bandwidth exhaustion using DNS, etc.). The exact mechanism for labeling a traffic spike as an attack is proprietary, and any method is likely to miss some attacks—especially small ones—but our aim here is merely to show the relative trend in attack categories.

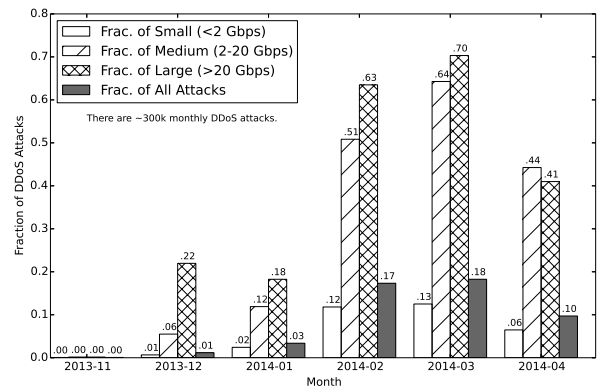


Figure 2: Fraction of all monthly global DDoS attacks in three size ranges and overall that are NTP-based. Note that NTP quickly rises from obscurity to dominate medium and large attacks. Mitigation appears to have had an impact with NTP-based attacks declining substantially in April.

In Figure 2 we show the attacks seen by Arbor Networks’ customers in the six months starting in November 2013. For each month, we bin attacks into three size categories: *Small*, consisting of aggregate bandwidth less than 2 Gbps; *Medium*, which are attacks

that are between 2 and 20 Gbps; and *Large*, which are any attacks larger than 20 Gbps. In the figure, a different bar represents the fraction of all DDoS attacks in each category as well as of all attacks that were NTP-based. We do not show the frequency of attacks in each category, but there are approximately 300K monthly attacks seen globally by Arbor Networks, and approximately 90% of them are Small, 10% Medium, and 1% Large. While the raw number of attacks was never dominated by NTP, the majority of Medium and Large DDoS attacks in February and March were, a dramatic change from just three months earlier, when only 0.07% of attacks involved NTP. The first quarter 2014 attack fractions were confirmed by Prolexic, who reported NTP used in 17% of attacks [31] (also see Goodin, [18], who quotes a third DDoS protection vendor that reported a majority of attacks in January were via NTP). In the first quarter of 2014, 85% of attacks exceeding 100 Gbps were using NTP reflection [24]. Fortunately, as evidence that the community’s mitigation efforts are starting to pay off, the fraction of attacks using NTP has started to decline in April and is now below February levels.

These traffic and attack statistics point to NTP-based DDoS as a major new vector and as another reminder of the power of reflected amplification DDoS attacks. This was a category of attacks that have been seen for some time but gained notoriety in recent years when unsecured open DNS resolvers started being used to amplify attackers’ ability to exhaust targets’ bandwidth. To better understand the amplifiers that miscreants are using to generate the NTP traffic, we next delve into the vulnerable NTP server population.

3. GLOBAL NTP AMPLIFIERS

The linchpin in NTP-based reflected amplification attacks is a large pool of globally-reachable NTP servers that are configured in a way that allows them to serve as reflectors and amplifiers.

3.1 The Amplifier Population

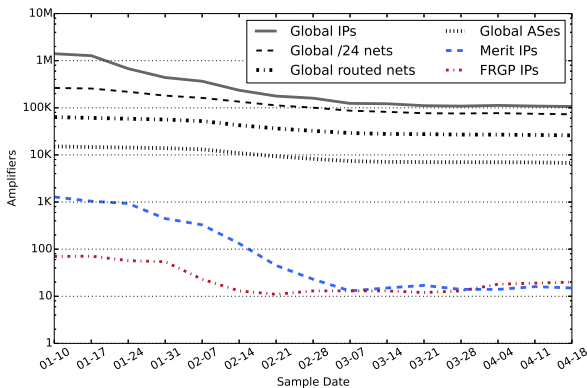


Figure 3: Count of NTP monlist amplifiers, including aggregations at the /24, routed block, and AS level, and subsets of IPs under Merit and CSU/FRGP address space. Remediation at IP levels has been swift, though many networks and nearly half of ASes still have some amplifiers.

Since January 10th, 2014, the OpenNTPProject.org [2] has been conducting weekly Internet-wide scans of the entire IPv4 address space in an effort to estimate and identify the global population of NTP servers that are vulnerable to being used as reflectors and amplifiers in DDoS attacks. These measurements entail simply sending a single NTP packet that requests a target server return the results of a *monlist* query, as described in § 1, and capturing all response packets. Since the size of the single query packet is known,

when servers reply with their packets of monlist table data, we are able to determine if each server can be used for attacks, and we can measure its amplification factor. We analyzed the response data that the OpenNTPProject.org shared with us (*ONP data*) in order to learn about the amplifier population as well as other facets of the NTP DDoS phenomenon. We report on fifteen weekly measurements through April 18th 2014, except where noted.

Before continuing, we discuss several limitations of this data. First, we note that one important source of error in the ONP dataset has to do with the type of monlist query packet that was issued. There are several implementations of the NTP service, and they do not all respond to the same packet format. The Linux `ntpd` tool for example, when used to query a server with the monlist command, tries each of two implementation types, one at a time, before failing. The ONP monlist scans in the dataset shared with us only used one of the two implementation values in the one NTP packet they send each IP. However, the implementation used appears to be more commonly present in the attacker scripts and darknet scanning data we have observed. Further, the raw numbers we report match those published by an independent scanning effort, reported in Kühner *et al.* [20]. Although the same error might affect both datasets. We also note that Internet-wide scans conducted from a single source IP, as the ONP scans were, are prone to be affected by some networks eventually blocking or filtering the scanning IP over time. The ONP data was not measured for this effect, but in their study, Kühner found an additional 9% amplifiers when scanning from a secondary network versus repeated scans from their main network. For these reasons, our numbers likely under-represent the full population of amplifiers. Finally, we note that the data only speaks to the amplifiers that respond to *monlist* and *version* (discussed later) queries, and not other NTP queries. Although these two are most commonly discussed in literature, there are other commands. Those others have typically lower amplification than monlist, however.

In Figure 3 we show the number of global NTP monlist amplifiers, and include aggregation at network levels, which underlines the global dispersion and varying remediation rates. The figure also shows amplifier count lines labeled Merit and CSU(FRGP), which we will elaborate on later. As can be seen, the global Amplifier population starts at approximately 1.4M in the first half of January and, as the Internet community begins to patch servers and block certain NTP requests, the population begins a mostly steady decline through mid-March, when it levels off at around 110K amplifiers. We discuss operator mitigation and amplifier remediation in § 6.

We were next interested in measuring the churn in the amplifier population. Over the course of the 15-week measurement period, we learn 2,166,097 unique amplifier IPs. We found that the first weekly sample only sees about 60% of these unique IPs seen, and that some new amplifiers are discovered on every scan conducted. About half of the amplifiers are only seen during one of the fifteen weekly scans, which is partly due to the rapid remediation of the amplifier population. A second factor is that a non-trivial fraction (13% – 35%) of these NTP servers are running on end-user (i.e., residential) machines and, thus, may be subject to DHCP churn. The left half of Table 1 shows the percentage of each sample’s IPs that are “end hosts” according to the Spamhaus Policy Block List (PBL) [34], taken on April 18th, 2014. The PBL identifies end-user IP addresses, and there is independent evidence that PBL-labeled IPs are, indeed, residential [7]. Finally, we calculate the average number of IPs per routed block, which starts at a peak of 22 and declines toward 4. This tells us that, initially, the pool of vulnerable servers included many large groups of closely-addressed (and, thus, likely managed together) server machines, whereas the population left in

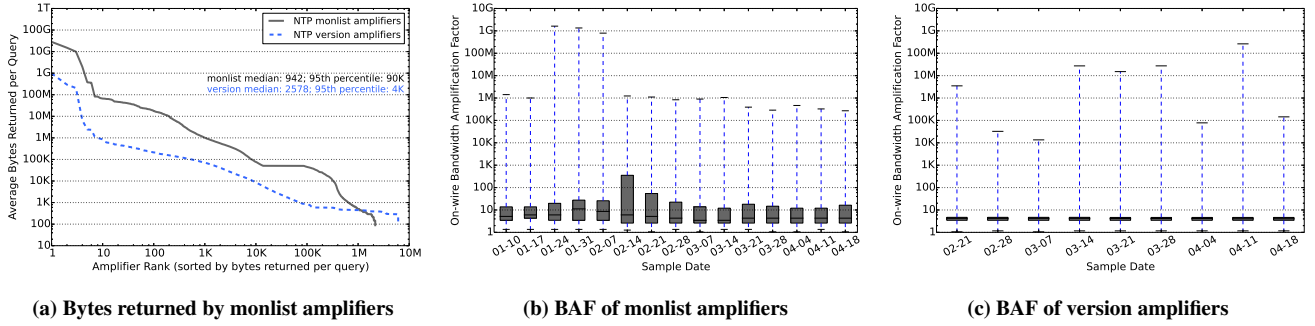


Figure 4: (a) Average on-wire bytes returned by monlist and version amplifiers. Boxplots of the resultant bandwidth amplification factors (BAF) for monlist (b) and version (c). Boxplots show the minimum, first quartile, median, third quartile, and maximum BAF. There are a small number of very large outliers in both types.

Table 1: For both amplifiers and victims seen in the fifteen weeks of ONP data, the IP counts, unique routed blocks, unique origin AS numbers, percentage of IPs that are end hosts, and number of IPs per routed block.

Date	Global Amplifiers						Global Victims					
	IPs	Blocks	ASNs	End Hosts	End Host %	IPs per Block	IPs	Blocks	ASNs	End Hosts	End Host %	IPs per Block
2014-01-10	1405186	63499	15131	260252	18.5	22.13	49979	16233	4797	15571	31.2	3.08
2014-01-17	1276639	61070	14671	207647	16.3	20.90	59937	18722	5373	19321	32.2	3.20
2014-01-24	677112	58519	14339	90889	13.4	11.57	66373	19690	5334	25504	38.4	3.37
2014-01-31	438722	56376	13903	74781	17.1	7.78	68319	20561	5351	28614	41.9	3.32
2014-02-07	365724	52229	13095	70053	19.1	7.00	81284	23062	5624	36765	45.2	3.52
2014-02-14	235370	42719	10961	63164	26.8	5.51	94125	25302	6154	42070	44.7	3.72
2014-02-21	176931	36411	9335	54578	30.9	4.86	121362	28235	6261	60866	50.1	4.30
2014-02-28	159629	32376	8241	51551	32.3	4.93	156643	31802	6702	83178	53.1	4.93
2014-03-07	123673	29159	7403	43531	35.2	4.24	153541	31111	6435	81684	53.2	4.94
2014-03-14	121507	27849	7115	40934	33.7	4.36	169573	32533	6585	88840	52.4	5.21
2014-03-21	110565	27590	7036	38870	35.2	4.01	167578	32748	6700	87550	52.2	5.12
2014-03-28	108385	27003	6997	37808	34.9	4.01	160191	31485	6512	82881	51.7	5.09
2014-04-04	112131	26947	7000	37880	33.8	4.16	143422	28656	5975	69340	48.4	5.00
2014-04-11	108636	26514	6925	36493	33.6	4.10	108756	24425	5272	52371	48.1	4.45
2014-04-18	106445	25976	6751	35683	33.5	4.10	107459	23264	5009	53233	49.5	4.62

April tends to be sparse, helping explain remediation slow-down, perhaps.

3.2 NTP monlist Amplifier Power

To begin to characterize the threat that this amplifier pool posed, we first aggregated the on-wire bytes of all monlist responses from each queried amplifier over the course of the ONP data collection weeks. As shown in Figure 4a, which plots the average per-sample on-wire bytes (i.e., packet bytes plus all Ethernet framing and overhead), there is a large range of data returned for the single monlist query packet. We find a median of 942 bytes for monlist responses, and maximum sizes for a given sample were typically in the tens to a hundred megabytes. Surprisingly, however, a small fraction of amplifiers responded with much more data than the monlist command should ever return; in one case, this was as high as 136 Gigabytes. We discuss these “mega amplifiers” in § 3.4. The figure also shows the response sizes to the *version* command, discussed in § 3.3.

A key feature of a good candidate service for use in amplification attacks is that it has a high asymmetry of responses to queries, (i.e., a high *packet- or bandwidth amplification factor (BAF)*). Thus, the servers that return the most packets or bytes for every packet or byte sent are the most powerful weapons for attacks. For simplicity, we focus just on bandwidth amplification in our analyses. We also caution that attackers may “prime” their amplifiers by first making connections from various IPs in order to make sure that the monlist table returns the maximum number of entries (600) when later sending traffic to victims. Thus, actual effects on victims may be larger when attackers make this effort.

To measure this relative power of the global population of vulnerable amplifiers over time, we calculated the aggregate on-wire

bytes from each amplifier in the ONP data and divided that by the on-wire bytes of a minimum monlist query packet. We used the 64 byte minimum Ethernet frame plus preamble and inter-packet gap, which total 84 bytes, to obtain the “on-wire” bandwidth amplification factor (BAF). Note that with respect to using all UDP, IP, and Ethernet frame overhead (including all bits that take time on the wire), our BAF calculations are lower than [32] but more accurately represent real bandwidth exhaustion effects via the most common (Ethernet) data links, as the actual load on the wire in both directions is considered. Figure 4b shows boxplots for the BAFs seen in each of the fifteen ONP monlist query response samples. As we can see, there is a wide range of BAFs in any sample, but the median is fairly steady across samples at around 4 (4.31 in the last five samples), and the maximum is generally around 1 million, except for the three samples starting on January 24th, when the maximum is around 1 billion. The third quartile BAF is typically around 15, except for the middle two samples in February, when it spikes to between 50 and 500. This suggests that, while the typical monlist-responding NTP server can provide an on-wire amplification of just 4x, a quarter of the amplifiers still seen in the wild can provide at least a 15x amplification. Using just one or a handful of such amplifiers, an attacker with a 100 Mbps Internet connection can easily overwhelm a service with a 1000 Mbps connection.

3.3 Threat of the Version Command

Our main focus in this paper is on the threat posed by the NTP *monlist* command, as it is known to have a high BAF, is of low utility for normal NTP operation, and has been used in high-profile DDoS attacks. However, NTP supports other commands that return more data than is sent (e.g., *version*, *showpeers*) though these have not

been as widely reported in attacks. As of February 21st 2014, the ONP data also includes separate Internet-wide NTP mode 6 *version* command probes. These are conducted in the same fashion as the monlist scans, in that every IP in the IPv4 address space is sent a single packet with the NTP version command and all response packets are stored. As of the April 18th ONP sample, the global pool of version responders is around 4M unique IPs. In Figure 4c we show the version BAFs observed. The measurements reveal several noteworthy differences between the version command and the monlist command threat. First, the pool of NTP server IPs that respond to the version query is much larger (4M vs 110K). Second, the version pool has not reduced substantially over the nine weeks that it has been measured. Third, there is much less variance in the BAF, (the 25th, 50th, and 75th percentiles are almost exactly the same at around 3.5, 4.6, and 6.9 throughout the nine samples). Fourth, there are still some outliers, as with monlist, with the maximum BAF as high as 263M, possibly due to the same routing loop-like behavior seen for the largest monlist amplifiers discussed in § 3.4. This all means that, while the threat from NTP DDoS using monlist may be starting to wane, the amplifier pool available for amplification and reflection using the version command is much larger and the median BAFs are comparable (though the higher end, 75 or 95 percentile, are much lower for version).

We were curious to know what fractions of scanning host or victim attack packets involved the version versus monlist command. To measure this, we tabulated the mode flag for likely victim or scanning NTP clients listed in the monlist tables that ONP-probed amplifiers return (detailed in § 4.1). We found that interest in the version command (mode 6) by both scanners and attackers relative to monlist (mode 7) appears to have grown somewhat since mid-February, with both showing the highest fraction of scanner or victim IPs contacting the sampled amplifiers in the final, April 18th sample (19% of scanners and 0.3% of victims). These values, especially for victims, should be interpreted with caution, since the global pool of NTP servers responding to the version command is nearly 40 times the size of the current monlist pool (4M vs 110K), and shrewd attackers may simply be using the former for version-based attacks and the latter for monlist. However, as monlist remediation reduces the pool of those amplifiers, this ratio may change.

Global NTP Versions and Systems: We parsed the responses to version command probes included in the ONP data between February 21st and March 28th. We aggregated version information for the samples, which include the OS, system, and version strings, as well as the NTP *stratum* of each server [23]. Table 2 shows the strings most commonly found in the OS field. No other systematic patterns in the data were prominent. We did make one surprising finding, however; of the 5.8M unique IPs returning data, nearly a fifth, 19%, reported stratum 16, which indicates that the NTP server is unsynchronized to the correct time [23]. This suggests poor management, as providing the correct time is the reason for an NTP server's existence. We also extracted the compile time year from all version strings, which was present in 1.1M of the server samples. Only 21% had compile dates in 2013 or 2014; We found that 59% were compiled before 2012, 48% before 2011, and 23% before 2010. Surprisingly, 13% were compiled before 2004, over ten years ago. Such poor state of updates and management is perhaps one reason vulnerabilities can have impact long after they are discovered.

3.4 The Case of the Mega Amplifiers

The maximum number of table entries that the monlist command returns (which we've confirmed empirically) is 600, and each entry includes just a handful of small fields, which we discuss in § 4.1. The expected maximum amount of data returned for a query is

under 50K. Indeed, as Figure 4a shows (note the log scale), the vast majority of amplifiers (99%) return less than 50K in aggregate.

However, as previously shown in Figures 4a and 4b, there is a small set of amplifiers that, at least in one or more of the fifteen weekly samples, behaves in an unusual and most devastating way. These "mega amplifiers," when sent a single packet of size less than 100 bytes, reply with megabytes or gigabytes of aggregate response packets. We found six amplifiers that responded with more than a gigabyte of aggregate packets, and the largest amplifier returned over 136 Gigabytes in a single day sample. In total, about 10 thousand amplifiers responded with more than 100KB of data, double or more than the command should ever return.

Since April, we have also been conducting twice-daily probes of a set of 250K IPs that were monlist amplifiers in any of the March 2014 ONP data samples. Between 60K and 15K of these IPs (decreasing over time) have been responding with monlist tables. Between April 2nd and June 13th, a set of nine IPs from seven ASNs had, on at least one occasion replied with more than 10,000 packets (at least 5MB). In parallel to monlist probes, we have been running packet captures to identify the amplifiers that exhibit this mega amplifier behavior. Several of these amplifiers did so on multiple samples, including the largest of the nine, which replied with more than 20M packets on each of at least a dozen samples during this period. This indicates that the behavior was not briefly transient but likely a systematic misconfiguration or bug. On May 31st, a single of these amplifiers sent 23M packets totalling over 100 gigabytes in just the first hour after our probe. Traffic data shows that this IP continued replying for hours afterwards. These mega amps often caused a steady stream of ≈ 50 Mbps of traffic, and spikes above 150Mbps were common, with the largest peak around 500Mbps, likely when more than one such mega amplifier was triggered. Strangely, all nine of these amplifiers were located in Japan, according to GeoIP information. We contacted JPCERT about these IPs and the operators were notified. After several weeks, these IPs no longer responded with excessive volume. However, we never received confirmation as to the root cause of the phenomenon.

If an attacker was lucky enough to either identify an amplifier or, by chance, happen to trigger an amplifier into behaving this way, he would hit the DDoS jackpot. Even a single host sending a small number of packets to a handful of such amplifiers could theoretically use up gigabits of victim bandwidth, and a single such amplifier would be enough to effectively knock a home Internet user, such as a gamer, offline, possibly for hours.

To understand what could cause this unusual behavior, returning to the 15-week ONP data, we first examined the *monlist* and *version* responses from both normal and unusually-large amplifiers, and found that these mega amplifiers did not differ systematically from the overall pool of 6M NTP servers or other amplifiers, with perhaps the exception of system strings, which we show in Table 2. We observe that in the overall pool, nearly half of the systems responding to the version command list their system as "Cisco," followed by 31% that list "Unix" (some Cisco devices running the IOS-XR OS apparently also report system as "UNIX"). In the mega amplifier (top 10k) pool (of which about half responded to the version command, allowing labeling) the reported system is most likely to be Linux (44%) or Junos (36%). In spite of this systematic difference, there is a large variety of systems represented in this mega amplifier pool, suggesting that a single bug or common configuration is likely not the cause of the behavior. For a second clue, we turned to the contents of the replies, whose parsing we detail in § 4.2.

Briefly, we looked at the monlist tables that the probed servers replied with, reconstructing the table from the packet payloads just as the NTP tools would do. We found that the mega amplifiers were

Table 2: Mega amplifier operating system strings, versus all amplifiers and all NTP servers reporting version information.

Rank	Mega (10k)		All Amplifiers		All NTP	
	OS	%	OS	%	OS	%
1	linux	44.18	linux	80.22	cisco	48.39
2	junos	35.85	bsd	11.08	unix	30.64
3	bsd	9.18	junos	3.43	linux	18.97
4	cygwin	4.82	vmkernel	1.42	bsd	0.97
5	vmkernel	2.41	darwin	0.92	junos	0.33
6	unix	2.01	windows	0.84	sun	0.21
7	windows	0.42	unix	0.56	darwin	0.13
8	sun	0.37	secureos	0.49	OTHER	0.14
9	secureos	0.25	sun	0.25	vmkernel	0.10
10	isilon	0.23	qnx	0.22	windows	0.07
11	OTHER	0.21	cisco	0.17	secureos	0.03
12	cisco	0.06	OTHER	0.41	qnx	0.02

incrementing the count for the ONP scanning server and resending an updated monlist table, continuously, up to thousands of times. Since we know that the ONP probe server only sent one packet, this behavior is consistent with a routing or switching loop or a similar networking stack flaw, which resulted in the ONP query packet being re-transmitted along the path or re-processed by these mega amplifiers’ ntpd processes. Other entries in these repeating (re-sent) tables showed that different clients of these servers had previously also seen the same multiple response behavior. These repeated responses typically occurred on a single sample week, but, in several cases, the same behavior was observed more than one week in a row for the same amplifier, suggesting something other than a brief transient failure. We are unable to definitively confirm that a network loop was to blame, but evidence points in that direction.

4. VICTIMOLOGY

NTP DDoS attacks involve at least three categories of systems: victims, amplifiers, and the attacker-controller nodes that spoof source addresses and trigger amplifiers to generate attack traffic. Our perspectives do not shed much light on the third category, but, in addition to the amplifiers described above, thanks to monlist for the first time we have insight into the population of NTP *victims*. This is because the monlist command used to conduct attacks can itself be used to gain an understanding of who the attacked targets are. Recall that the command returns the most recent set of clients that communicated with the NTP service, up to 600 (the median we saw was 6, the mean 70). Since NTP attacks use the NTP service by spoofing victim IPs, this list of “clients” will also include the IPs and other characteristics of actual DDoS victims and attacks.

4.1 Understanding the monlist Table

Table 3 shows two truncated examples of tables returned by an NTP server in response to a monlist query with IPs obfuscated. We have also eliminated three unrelated columns for clarity (the local address, client version, and restricted field). Thus, the columns of interest to us are: (i.) remote address (client or victim IP); (ii.) client source port; (iii.) count of packets; (iv.) client request mode; (v.) average inter-arrival time; and (vi.) last seen time (seconds ago).

In monlist table 3a, we first observe that the ONP scanning IP appears as the topmost address in the table. This is typically but not always the case, perhaps because some implementations may only update the table after replying to the probe. The *last seen* for the ONP IP is thus 0, and the *inter-arrival time* is typically around 600 thousand, since ONP probes IPs once a week. Second, we note that the count field for our probe is in the single digits, and that the mode column shows 7, which, along with mode 6, are the special NTP modes that include commands such as monlist and version. In

Table 3: Partial monlist table examples, showing the ONP probe, normal clients, as well as apparent NTP DDoS victims and the fields used to identify each.

(a) monlist Table A

Address	Src. Port	Count	Mode	Inter-arrival	Last Seen
ONP-IP	57915	6	7	526929	0
client.a1	10151	19	6	154503	310
client.a2	123	3281	4	1024	345
client.a3	54660	2	7	823	20795
client.a4	36008	1	3	0	104063

(b) monlist Table B

Address	Src. Port	Count	Mode	Inter-arrival	Last Seen
ONP-IP	47188	1	7	0	0
client-b1	59436	3358227026	7	0	0
client-b2	43395	25361312	7	0	0
client-b3	50231	158163232	7	0	0
client.b4	80	2189	7	0	2

addition to the ONP IP, we see four other clients, with request modes 3,4,6 and 7, that have also communicated with this server between 29 hours and five minutes ago (*last seen*). Client 1a sent 19 packets up to this point, about once every 43 hours, using mode 6; so, it is likely another probe. Checking the hostname of that IP confirms that it is indeed a research project. Client a3 is also a benign Internet survey host. Clients a2 and a4 appear to be normal NTP clients, contacting the server over mode 3 and 4 (normal operational modes) and using expected source ports and *inter-arrival times* (the time for a4 is 0, because it only sent one packet).

Moving to the second example table, 3b (assembled from entries in two actual tables for illustration), we again see the ONP probe packet. We also see four client IPs, each using mode seven, some with average inter-arrival times of 0 seconds and last seen within the last second. Further, the packet count sent by some clients is very high, with the largest of the four clients showing a count of over 3 billion packets. These all appear to be victims of the example amplifier. What is more, one of the clients is using source port 80, which, for UDP, is not a well-known service port but, as our victim port analysis (section 4.3.2) reveals, is commonly targeted.

4.2 monlist Table Analysis

For each of the weeks of ONP data, we parsed the responses to ONP monlist probes from each amplifier, 5M amplifier-week pairs. We applied the protocol logic found in the NTP ntpdc tool to rebuild the monlist tables present in the packets captured for each probed amplifier. If an amplifier sent repeated copies of the table (i.e., was a “mega-amplifier” discussed in § 3.4) we used the final table received that sample day. For each client IP in each table parsed, we use a filter we describe next to bin each client into one of three categories: a *non-victim*, a *scanner/low-volume victim*, or an actual apparent *victim* of that amplifier.

Identifying Victims: While the normal NTP modes could conceivably be used for reflection, there is little reason for attackers to do so, as no amplification is provided. Thus, if the mode value was less than 6, we classified the client as a *non-victim*. For clients that were using modes 6 or 7, we applied the following thresholds: if a client sent less than three packets to this amplifier (via the *count* field) or the average *inter-arrival* time was greater than 3600, indicating that it had received no more than one packet per hour from the amplifier, on average, we categorized the client as a *scanner/low-volume victim*, otherwise it was labeled a *victim*. We only report results for *victims*. This limits our reporting of victims to those above a low threshold of packets. Again, when a client received more than

three packets from an amplifier *and* averaged more than one packet per hour, we labeled it a victim of that amplifier. While this may seem like a low threshold, note that no legitimate NTP client would send these mode packets in the course of normal operation, and that individual researchers scanning the IP address space are not likely to re-scan a given IP more than once per hour.

For each victim seen by an amplifier, we extracted the *count* of packets it received, the average *inter-arrival* time, and the *last seen* time that the amplifier saw the victim, which is the attack end time for that amplifier/victim pair. We also estimated the duration of attacks, using simply the packet *count* multiplied by the average *inter-arrival* time. As the start time is not part of the monlist table, we estimate it using the duration and *last seen* time values, via similarly simple arithmetic including this calculated duration and time the table was received by the ONP scanner.

View Provided By Tables: Across all ONP weekly samples, the median largest *last seen* time in the monlist tables is about 44 hours. Thus, the median window within which we see attacks is approximately the previous two days. Arguably, this suggests that our samples will underestimate the number of victims, packets, or attacks by roughly a factor of 3.8 (i.e. there are 168 hours per week and we see 44). We don't show a plot, but, as expected, the largest *last seen* time in tables shrinks between the first sample in January and the peak of attacks in mid-February as the volume of attacks increases (thus shrinking the view of the 600-entry-capped monlist tables), and then increases again; following the same pattern seen in our other measurements (e.g., Internet NTP traffic). For this reason, our February peaks are likely underestimating attacks more than in the lower-volume months before or since.

4.3 Victims and Attacks

We next seek to understand who is the target of NTP DDoS attacks, when attacks occurred, how long the attacks lasted, and whether we see evidence for any of the public attacks.

4.3.1 Victim Locations and Networks

Our victim list includes victims from 184 countries in six continents, and, as Table 1 shows, the victim population in our samples spans up to 33K routed blocks and 6700 ASes at peak. We find that about half of victims are end host IPs, though this has grown from 31% on January 10th. The average number of IPs attacked in routed blocks is between 3–5, suggesting most attacks target a very small number of selected hosts in targeted organizations. Together with the fraction of target hosts that are end-hosts, this adds evidence to the idea that many DDoS attacks are launched against individuals.

We plot a CDF of the contribution of victim packets by autonomous system in Figure 5 for both amplifier and victim ASes. *Just 100 amplifier ASes are responsible for 60% of the victim packets measured.* Victim ASes are even more concentrated, with the *top 100 ASes receiving three quarters of all attack packets.* In examining top ASes with victims, we note that, out of the top ten ASes, eight are hosting providers and two are telecom companies. The top attacked AS is French hosting firm OVH, the purported target of large NTP DDoS attacks in mid-February (see § 4.4).

4.3.2 Attacked Ports

We were interested in seeing what services or ports were being attacked at these victim networks. For each victim seen at each amplifier, we tallied the source port used. Table 4 shows the top twenty attacked ports, along with the count of amplifier/victim pairs and fraction of all. We also include a column describing common use of each port, where known.

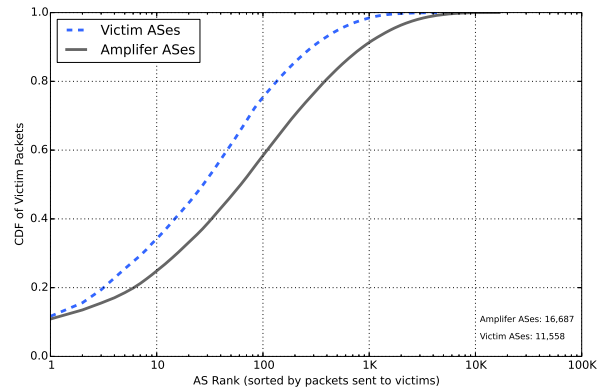


Figure 5: CDF of aggregate victim packets sent or received by autonomous systems. Just 100 amplifier ASes are responsible for 60% of the victim packets measured. Victims are even more concentrated, with 75% in the top 100 ASes.

Table 4: Top 20 ports seen in victims at amplifiers. Note prevalence of ports known to be associated with games, marked with (g), which add up to at least 15% of the victim ports attacked in the top 20 (we did not manually label ports above rank 20). Port 80 is also used by game systems, though over TCP.

Rank	Attacked Port	Fraction	Common UDP Use
1	80	0.362	None. via TCP:HTTP (g)
2	123	0.238	NTP server port
3	3074	0.079	XBox Live (g)
4	50557	0.062	Unknown
5	53	0.025	DNS; Xbox Live (g)
6	25565	0.021	Minecraft (g)
7	19	0.012	chargen protocol
8	22	0.011	None. via TCP:SSH
9	5223	0.007	Playstation (g); other
10	27015	0.006	Steam/e.g. Half-Life (g)
11	43594	0.004	Runescape (g)
12	9987	0.004	TeamSpeak3 (g)
13	8080	0.004	None. via TCP:HTTP alt.
14	6005	0.003	Unknown
15	7777	0.003	Several games (g); other
16	2052	0.003	Star Wars (g)
17	1025	0.002	Win RPC; other
18	1026	0.002	Win RPC; other
19	88	0.002	XBox Live (g)
20	90	0.002	DNSIX (military)

The top port attacked is port 80. Since port 80 in UDP is not a well-known application port (unlike in TCP, where it is the most used port, supporting HTTP), we speculate that perhaps attackers hoped that port 80 packets might be less likely to be filtered or blocked, enabling a more effective attack; this attack pattern has been reported by others [33]. Other ports prominently seen in attack traffic, include port 123 itself, which is the NTP *server* port (we also find this in our local impact datasets, described in § 7).

Game Wars: In addition to 80 and 123, other notable ports that show up in the top 20 list include at least 10 associated with gaming (e.g., ports for Xbox Live, Playstation, and specific games, such as Minecraft, Half-Life, and Counter-Strike). If we include port 80, which in TCP is used by many games (and non-games, of course) and may be mistakenly targeted, the total fraction possibly related to games may be even higher; likewise many of the ports above the top-20 list are also used by games. Along with the end-host fraction mentioned earlier, this targeting of game ports adds evidence to our conclusion that *a large fraction of NTP DDoS attacks are perpetrated against gamers*, as previously reported in press and industry reports (e.g., Goodin [18], Prolexic [31], and

Arbor Networks [8]). Minecraft (6th highest), in particular, was a confirmed large attack target, according to Goodin’s sources.

4.3.3 Attack Volume

Having discussed the target networks and ports, we turn to the volume of attacks. The aggregate attack packet count across all amplifiers over fifteen weeks is 2.92 trillion packets. The median bytes on the wire seen in the ONP response data is 420 bytes, thus, assuming each packet had that size, 1.2 petabytes may be a reasonable estimate of the aggregate data these particular attack numbers represent. However we believe that these samples underestimate the size of the victim population; this is because we do not see version victims from non-monlist-returning NTP servers, we only sample once per week, we only see the last 600 victims, and some amplifiers have been seen to attack thousands of victims at a time (e.g., see Table 5). As discussed in § 4.2, we under-sample by a median factor of 3.8. Thus, the likely actual volume of traffic to victims is probably closer to 3.8 times 1.2 petabytes. Of course, we caution that our data is lossy and sourced from often mis-manged devices, so these estimates may be considerably off. In general, we believe our numbers underestimate the amplifier population, victim population, attack counts, and attack sizes, because of the limitations discussed here and earlier.

Figure 6 shows the total number of packets that observed victims receive. We see that median attacks are quite small, at around 300-1000 packets, but the average is high at 1-10M, driven by a relatively small number of heavily-attacked victims. The 95th-percentile had been in the range of 400K to 6M until mid February, but since then has been in the 110k-200k range, suggesting the effect of remediation.

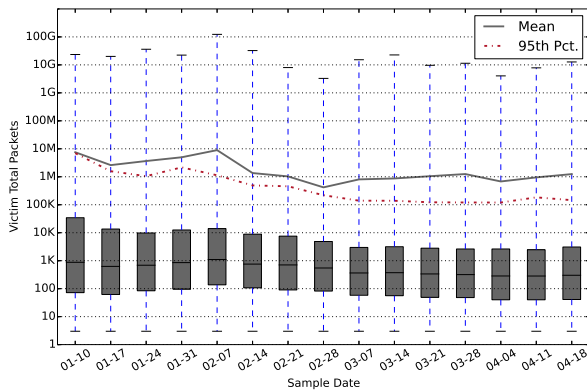


Figure 6: Total packets victims received from amplifiers. The mean packets are skewed by mega-amplifiers, but the 95th percentile has gone down by two orders of magnitude and median by a factor of 3.

4.3.4 Attack Counts and Durations

Because of the noise and lossy nature of our data, we make some simplifying assumptions when discussing attack counts and durations. First, we count each unique IP targeted in a given weekly sample as a single victim and attack. Of course, this is a simplification because: (i.) a given attack campaign may involve several IPs in a network block or several autonomous systems (ASes); (ii.) a single IP may be a host to multiple websites or users that are targets of independent attacks; and, (iii.) multiple attacks launched by one or more attackers may be targeting the same IP in small time frames that are too fine-grained for our data to disambiguate. Second, to determine start time of an attack, we use the median start time cal-

culated across all amplifiers seen attacking the victim. In § 4.2 we explained how we derive start times from the monlist tables.




Figure 7: Time series of attack counts seen in ONP monlist table data, including some time before first table. Attack counts are a lower bound. Mean: 514/hr; median: 280/hr.

In Figure 7 we use these derived median attack start times to show a time series of attack counts per hour seen in the ONP data. The attack counts are a lower bound, because, as discussed previously, the tables only show a median of 44 hours of amplifier activity every week. Samples start on January 10th, but some tables include evidence of older activity by virtue of not yet flushing older victims. Our calculations thus allow us to partially identify attacks with start and end times prior to our first sample, indicated with a vertical dashed line. We plot the attacks seen binned by hour as well as a daily average line. Note that the daily average peaks on February 12th, which is the time of the largest publicly-disclosed CloudFlare and OVH attacks (discussed in § 4.4), that started on February 10th. This peak corresponds to the same daily peak in NTP traffic observed in the Arbor Networks Internet traffic graph in Figure 1; likewise, the general trend up through mid-February and down afterwards matches the trajectory of global NTP traffic, suggesting that the attacks indeed drove global NTP traffic. Median attacks over our sample period only lasted approximately 40 seconds in the samples since mid February, and about half or a quarter of that in previous samples. On the other hand, the 95th percentile duration attacks lasted about six and a half hours in the January 10th sample and have been declining since, with 50 minutes being the 95th percentile duration in the later April samples.

4.4 Validation

One of the early massive NTP-based DDoS attacks that was observed in the wild occurred in early February and was disclosed by CloudFlare [30]. The attack started on February 10th, and was reportedly near 400Gbps in size, a record for NTP DDoS at the time. Purportedly, the attack targeted servers at OVH, a French-based hosting firm that was protected by CloudFlare [13]. The firm is one of the largest hosting providers in the world, and includes services targeting game servers [6]. In our rankings of networks that have been the targets of attacks, out of 11,558 ASes, OVH (AS number 16267) is the top AS (Cloudflare itself ranks 18th). Our data shows the OVH AS getting hit with over 170Bn aggregate packets from the amplifiers we studied, nearly 6% of all attack packets. OVH also features prominently in the Colorado State University (CSU) top-10 victims list (see § 7), accounting for five of the top 10 most attacked IPs. The attack campaign against it appears to be long-lasting; for example, OVH is the top AS in at least one weekly sample during each of the four months in our dataset, and shows up as one of the

top ten ASes that are attacked in 13 of the 15 weeks. Attackers target multiple unique IPs at OVH, peaking at nearly 4K in March.

CloudFlare also published the list of 1,297 ASes that hosted vulnerable NTP servers used as amplifiers in this attack. Of these 1,297, 1,291 were also seen in the ONP data, which totals 16,687 total amplifier-hosting ASes. The 1,291 that overlapped were, in aggregate, responsible for 60% of all victim packets. In addition to cross-dataset confirmation discussed throughout, this example supports the validity of our approach and gives us greater confidence.

5. ATTACKERS AND DARKNET SCANNING

5.1 View from a Darknet

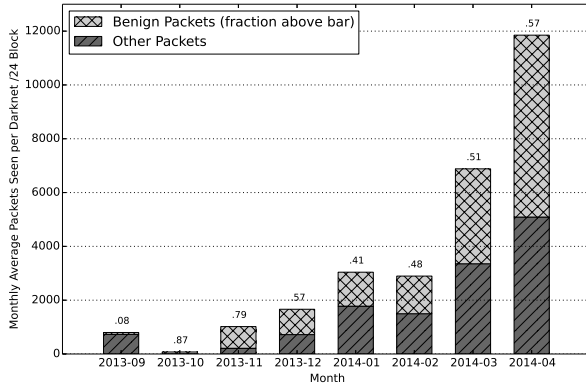


Figure 8: NTP scanning packet volume (in avg. packets per destination /24) detected by an ≈ 9 darknet over eight months, broken down by known benign (e.g., academic research projects) and other, which are from suspected malicious scanners.




Figure 9: Darknet scanning activity increased a week before the NTP attacks became significant.

Darknet or network telescope based observations of unused portions of the Internet IP address space (e.g., [35]) have often been used to detect the impact of large-scale phenomenon. The core idea behind such observations is that any large-scale significant physical or network event will ultimately have some spillover into unused address space (e.g., by not seeing expected noise packets from some remote network). Such analyses have, for example, been used to observe the impact of earthquakes on the Internet infrastructure, worm propagation, DDoS backscatter, misconfigurations, and censorship, among other events (e.g., [10, 11, 15]).

Based on the scale of the NTP events, we expected to find significant evidence of attacker activity (i.e., scanning) in darknets. Our

darknet dataset consists of full packet captures for roughly 75% of an IPv4 /8 address block¹ operated by Merit Network.

Figure 8 shows the volume of NTP traffic observed at the darknet by an average /24 network block equivalent. We notice a 10-fold increase in NTP related scanning activity starting from December 2013 to April 2014. It is interesting to note that we observe not just an increase in malicious scanning, but also scanning from various research efforts that were attempting to track the vulnerable NTP population (we identified these by their hostnames). Roughly half of the increase in scanning can be attributed to research efforts.

Figure 9 shows a time series of NTP scanning activity in terms of the number of unique IP addresses observed in our darknet. We are clearly able to pinpoint the onset of large-scale NTP scanning in mid December 2013. The figure also shows aggregate NTP traffic volume on the operational (non-dark) portions of Merit’s network (details in § 7). We observe that the rise in scanning activity precedes actual NTP attack traffic increases by roughly a week. This highlights the importance of using darknets to build active early warning systems of new and emerging threats [10]. It should also be noted that scanning traffic volumes continue to be high even as the global vulnerable NTP population has seen a dramatic decline. This indicates a continuing interest in finding vulnerable NTP servers.

Since at least one study reported seeing some UDP amplifiers (it is unclear if these were NTP) in the IPv6 space, we were curious to see if there was scanning activity observable in IPv6 [32]. We examined collected packets to dark (unused) address space in a large IPv6 darknet we operate, which includes covering prefixes for four of the five Regional Internet Registrars, including the RIRs for North America, South America, Asia, and Africa [14]. We searched for evidence of NTP scanning in the IPv6 darknet data in Nov. 2013, Dec. 2013, and Feb. 2014, but saw mostly errant point-to-point NTP connections, and no evidence of broad scanning. Likewise, the Arbor Networks netflow data for IPv6 does not list NTP within the top 200 UDP ports (it was 12th in IPv4), and, thus, did not show a noticeable level of NTP traffic.

5.2 Attacker Ecosystem and Motivations

The concept of “attacker” in this type of DDoS activity is both nebulous to define and difficult to measure. Unfortunately, most of the available datasets shed very little light on who the actual attackers are that perpetrate these NTP-based DDoS attacks or what their motivations may be. However, there are a few small clues that public reports have revealed as well as a few tidbits in our data.

Dissecting the DDoS ecosystem and understanding attackers is complicated by the fact that several types of actors are involved in launching attacks. We’ve reserved the term “amplifier” for the vulnerable or misconfigured boxes that are leveraged in attacks to flood victims with traffic. These are part of the problem, but they are more enablers than aggressors. The attack ecosystem has several other entities that better fall under the umbrella of “attacker.” First, there are the nodes that send spoofed-source packets to amplifiers in order to elicit the large responses that flood victims. These may or may not be the machines that are owned by the humans launching attacks. In many cases, they are actually compromised Internet user machines (“zombies” or “bots”) that can be remotely commanded to perform such actions on behalf of a “botmaster.” Thus, the second entity we might label an attacker is the botmaster himself. Certainly, this person or group and the system they use can be

¹We typically have around 75% coverage for the darknet /8 in terms of effective unique /24 that are advertised and can receive traffic. However, the size of the darknet varies over time due to routing changes and suballocations. To account for this, we normalize the data to average packets per effective dark /24s that month.

labeled culpable. But, blame might not stop there. These botmasters could be individuals acting on behalf of a black market DDoS (“booter”) service, many of which are advertised on underground forums (e.g. [5]) [19]. They or the service may have been hired by the party that is actually motivated to cause damage to the victim.

So, while a botmaster or booter service is likely to be motivated by money, the person that actually wants the attack to happen could be motivated by anything that normally motivates people to attack others. This, of course, includes money (e.g., via extortion [28]), revenge, political reasons, competitive advantage (see Poulsen for an early example [29]), etc. According to a large 2014 survey of global network operators, political/ideological reasons top the list of perceived motivations for being attacked [8].

We discuss several clues about attackers and their motives that appear in our own data where that data is presented. For instance, in § 4.3.2 we mentioned that a significant amount of victims were targeted on game-related ports. This is congruent with the idea that many DDoS attacks are perpetrated against gamers, by rivals or for financial gain, as reported previously (e.g., [4, 18, 19, 26, 31]). Another clue is discussed briefly in § 7.2, where we found that packet TTLs of scanning packets indicate they are mostly Linux, while packets sent by nodes generating traffic to amplifiers indicates they run Windows. As botnet nodes are often Windows machines, while individual miscreants with enough sophistication to conduct broad Internet scans may be Linux users, this clue fits the story that attackers are using botnet hosts to indirectly launch attack traffic.

6. REMEDIATION

One of the most encouraging observations regarding the NTP DDoS attacks has been the community response. Community response to the NTP amplifier threat has been swift [17], with the number of vulnerable monlist NTP servers dropping dramatically from a high of 1.4M when first measured on January 10th, down to less than half (678K) just two weeks later and continuing to fall to around 110K, where it has held steady since March 21st. Some interesting facets of how remediation occurred are presented next.

6.1 Subgroup Remediation Rates

One interesting aspect of how monlist amplifier remediation is proceeding is its varying nature across several axes.

Network Levels: First, we look at network granularity. We already noted that the overall set of amplifier IPs has been reduced in cardinality from approximately 1.4M in early January to 110K by April 18th; a reduction of 92%. However, when we aggregate amplifier IPs by /24 subnets, we find that the reduction is from 264K to 73K, or 72%. There are at least two possible reasons for this discrepancy; given IPs in a /24 might be managed by different operators (e.g., because they are home PCs in a residential ISP), or the individual hosts may be more or less difficult to patch, for example, due to their role. When we aggregate up to the routed block level, the percentage reduction falls again, from 64K routed blocks to 26K, or 59% reduction. At the autonomous system level, it is only 55%, from 15k origin ASes to 6.8K. These trends highlight the difficulty in completely eliminating a vulnerability from all of a single network, let alone such a large number of independently-managed networks.

Regional Levels: The second axis along which remediation differs is the regional one. When we aggregate the amplifiers according to their continent, we find that North America has remediated 97% of its amplifiers, Oceania 93%, Europe 89%, Asia 84%, Africa 77%, and South America 63%. These differences in the speed to remediate amplifiers across region may be caused by various socio-economic factors, by the relative regional impact that these amplifiers have caused, and by network management practices or norms.

End Host Composition: A third axis is the relative composition of amplifiers that are end hosts. We again used the PBL [34] to label each IP seen in the weekly samples as either an end host or not. As Table 1 showed, as the pool of amplifiers was remediated, the fraction of amplifiers that are end hosts approximately doubled from 17% in the first two weeks to 34% in the last, suggesting that perhaps remediation was more likely to happen at servers that are professionally managed versus at workstations.

6.2 Comparison to Open DNS Resolvers

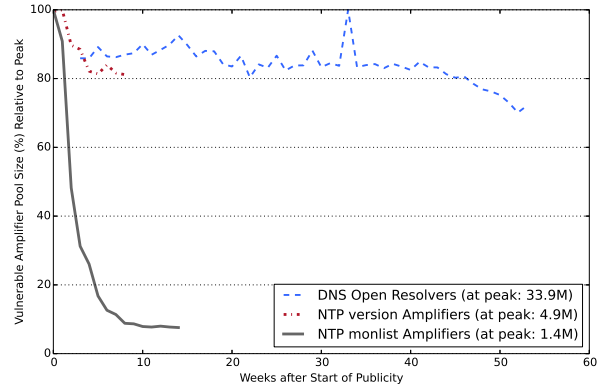


Figure 10: Size of vulnerable NTP (and, for comparison) DNS amplifier pools, relative to peak, versus weeks since the OpenNTPProject.org and OpenResolverProject.org, respectively, began publicizing vulnerable server counts. (We ablate the first three DNS samples and a single DNS outlier, artificially low due to collection and methodology issues.) Monlist amplifiers have been remediated dramatically faster than the other two pools.

The initial rapid remediation of over 90% of vulnerable servers is remarkable in how it compares to another, related, Internet threat, DNS-based DDoS amplifiers. In Figure 10 we show the fraction of amplifiers seen in the wild versus the number of weeks since the OpenNTPProject began collecting data and raising awareness of the threat. For comparison, we show the counts of open DNS resolvers, which are susceptible to use in amplification/reflection attacks for the same reasons that NTP servers are attractive. The OpenResolverProject.org [3], which is run by the same people as the NTP project, has been conducting identical IPv4-wide surveys of the size of the vulnerable DNS server population for about a year. As the figure shows, that pool has not decreased much in relative terms. A possible key difference is that open DNS resolvers are often found on customer premises equipment, which is much more difficult to update or replace than are infrastructure servers, like those that typically run NTP. We measured the intersection between the NTP monlist amplifier IP pool and the corresponding open DNS recursive resolver pool. It is about 7K out of 107K monlist amplifiers in the latest sample. The aggregate unique IPs seen over 15 ONP samples and the DNS IPs over the same period show an overlap of 199K or 9.2%. These badly mis-manged IPs may remain vulnerable for some time to come given the length of time since the open resolver threat has been known and that a non-trivial fraction of networks are mis-managed for reasonable metrics of management [36].

6.3 The Effect of Remediation

The drastic amplifier reduction we saw was also evident in the number of amplifiers seen per each victim IP, which also decreased by an order of magnitude across the ONP samples. On the other hand, we don’t show this in a figure, but the number of packets

that the average amplifier sends all victims has actually gone up by about an order of magnitude, to somewhat compensate for the reduced number of amplifiers. In other words, remaining amplifiers are being put to more use. Likewise, as Table 1 showed, more victims are being attacked over time, though this stopped increasing in April. As shown in Figure 6, the average number of packets that a victim sees has decreased by about an order of magnitude (from 10M in January to 1M starting in mid-February), and the median has decreased to about a third.

It is possible, and, as discussed in § 3.3, likely, that other NTP commands may be used by attackers as the pool of monlist amplifiers is reduced. Figure 10 also shows the size of the version command amplifier pool, which only decreased 19% since peak.

Because our victim and attack counts come from a parsing of the monlist tables, as our view of the global NTP population is reduced when the command is blocked or patched, we will see fewer victims, attacks, and packets. Thus, our estimates of these populations and the attack traffic are a lower bound. Finally, those small number of mega amplifiers, if they are being exploited by more attackers, might increase actual overall attack traffic seen by victims.

6.4 Operator Motivation and Notification

We learned that part of the reason that monlist amplifiers were remediated more quickly than other amplifier populations may be due to an aggressive notification effort that was conducted via CERTs and direct operator contact, as recently reported by Kühner *et al.* [20] just as we were going to press. While it is plausible and likely that communication had an impact on remediation, causality, unfortunately, can not be determined. Our discussions with experts in DDoS at Arbor Networks and with one global ISP operator also suggested that operator motivation to remediate amplifiers or otherwise mitigate NTP traffic in other ways was likely increased by the large impact that this traffic was having on the operators’ own networks. Understanding the reasons and mechanisms for the dramatic reduction in the NTP monlist population was beyond our scope, but we mention these two possible causes as it is important for the community to work to understand what encourages beneficial operator response to such global Internet threats. Clearly, remediation and mitigation of amplifiers have positive externalities—they are behaviors that not only help the operator’s own network but help the Internet at large be more secure.

7. A VIEW FROM REGIONAL NETWORKS

In the previous sections we have focused on the global view of the NTP DDoS phenomenon. We now turn our attention to how these events appeared at individual networks and the commonalities between them.

7.1 Local Impacts

The two local views include Merit Network and Front Range Gigapop (FRGP). Merit [22] is a large regional ISP serving multiple educational and non-profit institutions via a state-wide fiber backbone in Michigan while FRGP [1] is the major ISP serving Colorado State University (CSU).

Overall aggregate Merit traffic ranges from 15-25 Gbps. On a normal day, NTP constitutes a negligible fraction of this traffic. The NTP attacks first became visible at Merit on the third week of December 2013. Figure 11 shows an almost instantaneous increase in both inbound and outbound aggregate NTP volume with peaks exceeding 200MB/sec. At FRGP/CSU, however, we see the first sign of NTP attacks almost a month later. Figure 12 shows the aggregate NTP volume at CSU and FRGP. Due to network size differences and the limited FRGP vantage points, the attack volumes

observed were an order magnitude less than at Merit. The CSU NTP servers were secured on January 24th, 2014. This is the point at which the NTP traffic volumes returned to pre-attacks levels. We do however, note that other networks within FRGP were not nearly as proactive. A number of vulnerable NTP servers within FRGP were used to launch attacks throughout the month of February. NTP traffic volume continues to grow throughout our observation window of 3 months. The distinctive spikes in FRGP ingress traffic in Figure 12 were NTP reflection attacks directed at specific hosts within FRGP prefixes. The largest one on the 10th of February lasted for just under 23 minutes with an attack rate of close to 3GBps and a total of 514 GB of attack traffic.

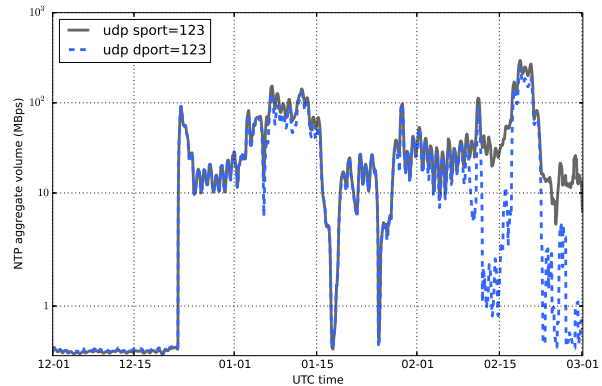


Figure 11: Merit NTP traffic (3 months).

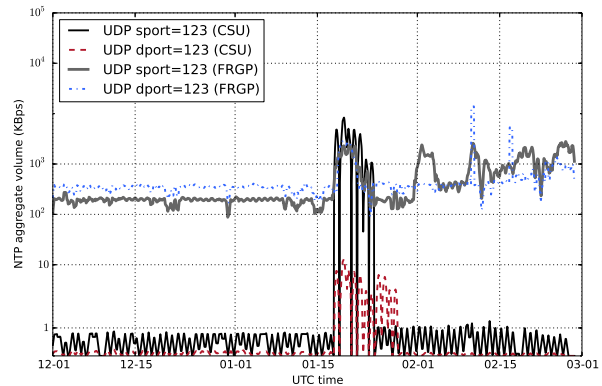


Figure 12: CSU/FRGP NTP traffic (3 months).

For the purposes of conducting detailed forensics we focused our attention on a 12-day period for the Merit dataset, and a 19-day period for CSU and FRGP datasets starting January 25th and January 18th, 2014 respectively. During the corresponding periods, we identified 50 NTP amplifiers inside Merit with an average amplification factor of 187, nine amplifiers at CSU with an average amplification factor of 436, and 48 amplifiers at FRGP². Note that, as both Merit and FRGP are ISPs that each encompass multiple independent networks, these numbers of NTP servers are larger than a single enterprise would typically run. Further, depending on configuration, potentially any server or high-end router can act as an NTP server.

Table 5 shows the five worst amplifiers at Merit and CSU and their BAF, unique victims contacted over the periods studied, and the total volume sent in gigabytes³. Our analysis shows the extent to which

²BAF was not computed for FRGP due to incomplete picture of egress and ingress volume of the dataset.

³In this section, amplifiers, victims and the BAF are defined similar to [32]. Here, a *victim* is a client receiving at least 100KB from




Figure 13: Time series of top-5 affected victims (Table 6).

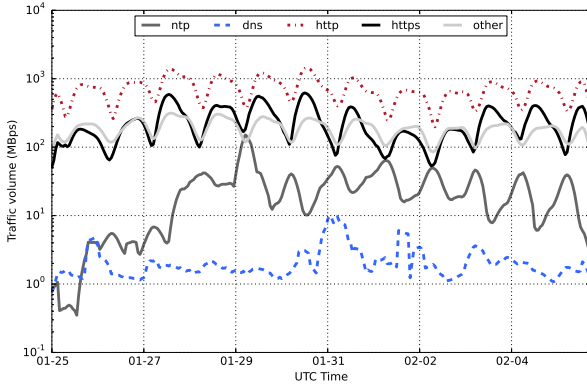


Figure 14: Time series of all traffic at Merit.

some of the NTP amplifiers in these networks participated in this global phenomenon. Notice, for example, that the five amplifiers alone identified in Merit were abused to target thousands of victims generating terabytes of data over a short time span.

Overall, we identified 13,386 unique victims at Merit and 5,659 at FRGP and CSU. There were 291 victims common between the two sites. We identified target networks and countries to which these systems belong. Table 6 shows a characterization of the top 5 victims at Merit and at CSU. Our data also clearly shows signs of coordination in these attacks as we frequently see several amplifiers being used to target the same victim at each site (Figure 15). All of the 9 CSU amplifiers were observed to attack many victims in a coordinated fashion. Further, in several cases at Merit, more than 35 of the identified amplifiers were used in a coordinated manner in attacks that lasted multiple days. Figure 13 provides a stacked-graph visualization of these victims. Interestingly, we observe a diurnal pattern of traffic destined to the victims perhaps suggesting a manual element in the attacks. Note that the larger attacks that used the most amplifiers also lasted longer (e.g., see top half of Table 6).

Since the attack volume constitutes a significant amount of ISP traffic, we investigate whether these attacks had monetary impact. Figure 14 shows the aggregate NTP traffic volume at Merit, along with other traffic, illustrating NTP’s steep rise. We estimate that attacks resulted in over 2% additional traffic at Merit, overall, which would incur extra transit costs. Whether it did or would at another ISP depends on the billing model used (e.g., a 95th percentile model, which Merit uses with its upstream), the aggregation window size, and when the extra traffic was transited [16]. For example, if the

an amplifier with a ratio of amplifier’s bytes received to bytes sent of at least 100. An amplifier sent at least 10MB and had a ratio of sent/received traffic greater than 5. BAF is UDP payload ratio.

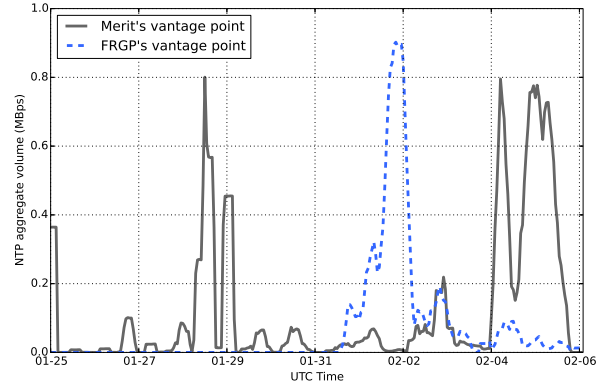


Figure 15: Common Merit/FRGP victims volume.

Table 5: Top-5 amplifiers at Merit and CSU

Amplifier	BAF	Unique victims	GB sent
Merit-A	1297	1966	375
Merit-B	1148	1626	4697
Merit-C	1004	3072	5808
Merit-D	993	1801	316
Merit-E	948	2740	4369
CSU-F	805	38	162
CSU-G	797	33	163
CSU-H	796	33	163
CSU-I	469	238	223
CSU-J	465	236	222

extra traffic resulted in a higher billed traffic level after removing the 5% of peak traffic in a 95th percentile model, it might have produced more transit cost at the ISP, which was the case at Merit.

Remediation: Figure 3 showed the rate at which remediation efforts at both these networks progressed. At Merit, trouble tickets were used to track the status of each identified amplifier, and customer notifications were used to encourage prompt patching of systems. During the early stages of the attacks, Merit also put in place traffic rate limits on NTP traffic to minimize the impact of these attacks to its customers. At CSU, due to the small set of servers, patching happened rapidly, within a single day, though remediation in the rest of FRGP is ongoing. Likewise, some holdouts remain under Merit.

7.2 Individual Attacker/Victim Activity

One of the unique aspects of our datasets from these two sites is that they offer us a level of detail that is not present in some of our higher level global datasets. In particular, we are able to examine flow and packet level details that offer additional insight about the activity of individual attackers and victims as well as validate our global datasets.

The attack volume to common targets is shown in Figure 15. We plot the traffic volume to common targets as recorded from our two vantage points, Merit and FRGP. We found 291 common targets attacked by amplifiers at both sites. However, the attack

Table 6: Top-5 victims at Merit and CSU

Victim	ASN	Country	BAF	Amplifiers	Dur. Hours	GB
Merit-A	AS4713	Japan	105	42	114	5887
Merit-B	AS4837	China	1380	4	143	4542
Merit-C	AS30083	USA	202	7	166	4017
Merit-D	AS8972	Germany	165	7	166	1703
Merit-E	AS8972	Germany	147	7	166	1595
CSU-F	AS16276	France	730	9	31	17
CSU-G	AS39743	Romania	658	9	143	14
CSU-H	AS28666	Brazil	670	9	30	12
CSU-I	AS12390	UK	670	9	51	10
CSU-J	AS16276	France	669	9	74	10

volumes were fairly low. We also find that the overlap in target ports between attack traffic at Merit and the global data shown in Table 4 is remarkable. The ports shown in that table constitute 98.91% of the total traffic destined to victims identified with Merit’s flow data.

The information we obtained from the local views confirms some of our observations from the global ONP data in terms of attack magnitude and attack occurrences. For example, the BAFs reported via traffic analysis of the local datasets Tables 5 and 6 confirms those depicted in Figure 4b after allowing for differences due to packet headers. Additionally we also note that FRGP data shows servers in that network actively participated in the attacks on OVH servers described in section 4.4. Furthermore, the remediation efforts at each site were clearly visible in the ONP datasets giving us increased confidence in the completeness of the ONP data.

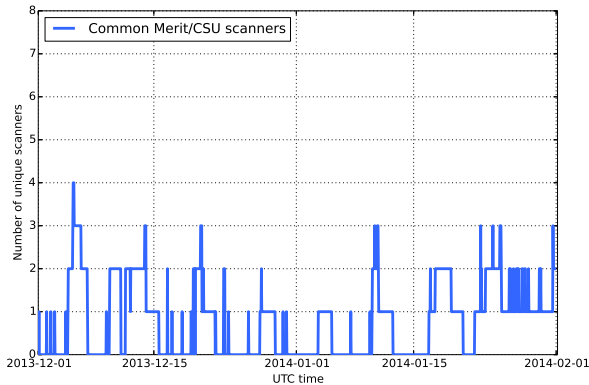


Figure 16: Common scanners Merit/CSU.

We also attempted to identify common scanning activity that might have been visible from both sites. However, we find little evidence of this in our datasets from Merit and CSU. Figure 16 shows a timeline of the trickle of activity from these common scanners. Though we identify 42 IP addresses in common, most of these were determined to be a result of research scanning. We speculate that true malicious scanning activity may be well distributed in time, and, therefore, the likelihood of two distinct sites observing synchronized scans is relatively low—whereas research scanning is being conducted in the open and at much more aggressive rates.

Another attacker behavior of interest is whether it is possible to estimate if scanning activity (to find amplifiers) and attack traffic (for amplifiers to reflect) is sourced from the same systems. To study this, we analyzed TTL values from the CSU dataset corresponding to both scanning activity as well as spoofed attacks. Surprisingly, we find that while the scanning activity appears to be largely sourced from Linux-based systems (mode TTL:54), the attack traffic appears to be originating from Windows-based systems (mode TTL: 109), perhaps botnet nodes.

8. RELATED WORK

The threat of DDoS attacks leveraging reflection has been well-known for many years (e.g., [27]), and the 2014 study by Rossow [32] is the latest and most comprehensive example of work examining the potential vectors for amplification in these reflected attacks. The Rossow paper characterized how 14 UDP protocols, including SNMP, chargen, and NTP with the monlist command, may be used for reflection/amplification DDoS attacks. For each of these protocols, broad and multifaceted passive and active measurements of a limited scale (e.g., scanning a random sample of 1M IPv4 addresses, partial crawling of P2P networks, and measurements of scanning in a medium-sized (/17) and small (/27) darknet over four weeks) were

conducted. In contrast to that comprehensive examination of amplification in general, we instead focus just on the threat and actual harm caused by attacks leveraging NTP. Further, our measurements of this particular amplification vector are much broader, over a longer timespan, and include a deep analysis of the actual attack and victim characteristics seen in the wild, and at Internet-scale. At the time Rossow’s paper was written, no major attacks using NTP had been reported. Since then, NTP had become the dominant attack vector for large DDoS attacks, and our work aims to chronicle that rise, explore who was attacked and how, and show how the mitigation of the vector has impacted the threat.

As our manuscript was undergoing final preparation for publication, new work by Kühler *et al.* [20] was published exploring several DDoS attack vectors and aspects. Pertinent to our work here, the study examined the size of the NTP amplifier populations via global scanning, conducted a notification campaign that may have (though causality can not be shown) speeded remediation of monlist amplifiers, and reported on several aspects of the amplifier population over 13 weeks, starting in November 2013. The numbers reported in Internet scans for NTP amplifiers match our numbers closely, as expected. In addition, characterizations of the amplifier pool (e.g., that nearly half of the version command amplifiers are Cisco), likewise match our analysis. The study also examined two interesting but unrelated to our work facets of DDoS in general, one being a TCP attack vector and the other a technique for remotely identifying networks that allow IP spoofing. Unlike the Kühler *et al.* study, our work digs deeper into who NTP attacks target and with what force, as well as explores other features of the attacks (e.g., attacked ports). We also bring to bear a dataset on NTP scanning from a large (≈ 8) darknet, large global traffic and attack data, and data from local ISPs impacted by these attacks.

While DDoS attacks in general and mitigation strategies in particular have seen much work, reports on large-scale measurements of DDoS attacks are few. We are not aware of any studies characterizing DDoS attacks at-scale since 2006, aside from proprietary industry reports based on commercial deployments (e.g., [31], [30]), whereas the DDoS threat landscape has evolved significantly in the ensuing years. The 2006 measurement studies examining DDoS attack activity focused on backscatter (e.g., [25]) and flow analysis (e.g., [21]). Backscatter is evident in random spoofed-source flood attacks (esp. SYN-flood), which makes it inapplicable to the specific type of attack we focus on, NTP-based reflection/amplification attacks. The relative fraction of attacks based on NTP is similar to that using SYN floods [31], though it is unclear what fraction of modern SYN floods utilize random spoofed source addresses, which are required for detection in darknets, as done by [25], and [21]. To address some of the limitations of backscatter, Mao *et al.* [21], argued for direct measurement of DDoS attacks (e.g., flow data), which is one of the approaches we take here.

To our knowledge, ours is also the first study to measure amplification-type DDoS attack activity via a direct global survey of records on the amplification hosts themselves. In addition, we bring to bear passive datasets with both global and local perspective.

9. CONCLUSION

Using data from a variety of vantage points, we chronicle the rapid rise and steady decline of the NTP DDoS attack phenomenon. Our analyses serve to characterize the global scale of attacks, both in terms of amplifiers as well as victims. We confirm the value of actively monitoring darknet address space, as it can help detect attack precursors, such as large-scale scanning and probing, observed prior to the onset of the first large-scale attacks. We demonstrate that, in addition to countless hours spent by engineers worldwide

to install patches and filters, these events had direct measurable costs in terms of increased bandwidth loads as measured at example edge networks. Though this paper documents the lethal power of the largest DDoS attacks observed to date, our conclusions include a positive one. The network research and operations community worked to actively mitigate the effects of these attacks and these efforts have had a visible impact in diminishing the vulnerable amplifier population and reducing attack traffic. There are, however, limits to the effectiveness of such remediation efforts, as the tapering of mitigation shows. Since rapid remediation is how such attack vectors are thwarted, we are interested in future work examining why some networks remediate faster than others.

Acknowledgments

This work was supported in part by the Department of Homeland Security Science and Technology Directorate under contract numbers D08PC75388, FA8750-12-2-0314, and FA8750-12-2-0235; the National Science Foundation under contract numbers CNS 1111699, CNS 091639, CNS 08311174, CNS 0751116, CNS 1330142, and CNS 1255153; and the Department of the Navy under contract N000.14-09-1-1042. We would like to thank Jared Mauch for sharing the OpenNTPProject.org dataset as well as Kirk Soluk and team at Arbor Networks for sharing traffic and attack statistics. Finally, we are grateful to Roland Dobbins, Christian Rossow, Denis Foo Kune, anonymous reviewers, and our shepherd, Sharon Goldberg, for valuable feedback on earlier drafts.

10. REFERENCES

- [1] Front Range GigaPop. <http://www.frgp.net/frgp-overview-2014-03-27.pdf>.
- [2] Open NTP Project. <http://openntpproject.org/>.
- [3] Open Resolver Project. <http://openresolverproject.org/>.
- [4] Arbor Networks Solution Brief: DDoS Attacks in the Gaming Industry, 2013. www.arbornetworks.com/docman-component/doc_download/687-gaming-company-defends-against-ddos-attacks.
- [5] Hack Forums “Server Stress Testing” marketplace forum, Aug. 2014. <http://www.hackforums.net/forumdisplay.php?fid=232>.
- [6] The OVH offering expands with new lines of dedicated servers, Feb 2014. https://www.ovh.com/us/newsroom/cpl355.the_ovh_offering_expands_with_new_lines_of_dedicated_servers.
- [7] M. Allman. Comments on Bufferbloat. *ACM Computer Communication Review*, 43(1), Jan. 2013.
- [8] D. Anstee, A. Cockburn, G. Sockrider, and C. Morales. Arbor Networks Worldwide Infrastructure Security Report, 2014. <http://pages.arbornetworks.com/rs/arbor/images/WISR2014.pdf>.
- [9] Arbor Networks. www.arbornetworks.com.
- [10] M. Bailey, E. Cooke, F. Jahanian, J. Nazario, and D. Watson. The internet motion sensor: A distributed blackhole monitoring system. In *Proceedings of Network and Distributed System Security Symposium (NDSS '05)*, pages 167–179, 2005.
- [11] K. Benson, A. Dainotti, k. Claffy, and E. Aben. Gaining Insight into AS-level Outages Through Analysis of Internet Background Radiation. In *Proceedings of the 2012 ACM Conference on CoNEXT Student Workshop*, CoNEXT Student '12, 2012.
- [12] S. O. Blog. Hackers Spend Christmas Break Launching Large Scale NTP-Reflection Attacks, Dec 2013. <http://www.symantec.com/connect/blogs/hackers-spend-christmas-break-launching-large-scale-ntp-reflection-attacks>.
- [13] L. Constantin. OVH's Own NTP Servers Used in Attack, Feb 2014. <http://news.techworld.com/security/3501549/attackers-use-ntp-reflection-in-huge-ddos-attack/>.
- [14] J. Cxyz, K. Lady, S. G. Miller, M. Bailey, M. Kallitsis, and M. Karir. Understanding IPv6 Internet Background Radiation. In *Proceedings of the ACM SIGCOMM Internet Measurement Conference (IMC'13)*, Barcelona, Spain, 2013.
- [15] A. Dainotti, C. Squarcella, E. Aben, K. C. Claffy, M. Chiesa, M. Russo, and A. Pescapé. Analysis of Country-wide Internet Outages Caused by Censorship. In *Proceedings of the 2011 ACM SIGCOMM Internet Measurement Conference (IMC'11)*, pages 1–18. ACM, 2011.
- [16] X. Dimitropoulos, P. Hurley, A. Kind, and M. P. Stoecklin. On the 95-percentile billing method. In *Proceedings of the Passive and Active Network Measurement Conference (PAM'09)*, 2009.
- [17] J. Fleury. Good News: Vulnerable NTP Servers Closing Down, Feb 2014. <http://blog.cloudflare.com/good-news-vulnerable-ntp-servers-closing-down>.
- [18] D. Goodin. New DoS attacks taking down game sites deliver crippling 100Gbps floods, Jan 2014. <http://arstechnica.com/security/2014/01/new-dos-attacks-taking-down-game-sites-deliver-crippling-100-gbps-floods/>.
- [19] M. Karami and D. McCoy. Understanding the Emerging Threat of DDoS-as-a-Service. In *Presented as part of the 6th USENIX Workshop on Large-Scale Exploits and Emergent Threats*. USENIX, 2013.
- [20] M. Kührer, T. Hupperich, C. Rossow, and T. Holz. Exit from Hell? Reducing the Impact of Amplification DDoS Attacks. In *Proceedings of the 23rd USENIX Security Symposium*, August 2014.
- [21] Z. M. Mao, V. Sekar, O. Spatscheck, J. Van Der Merwe, and R. Vasudevan. Analyzing Large DDoS Attacks Using Multiple Data Sources. In *Proceedings of the 2006 SIGCOMM workshop on Large-scale attack defense*, pages 161–168. ACM, 2006.
- [22] Merit Network, Inc. www.merit.edu.
- [23] D. Mills, J. Martin, J. Burbank, and W. Kasch. Network Time Protocol Version 4: Protocol and Algorithms Specification. RFC 5905, 2010.
- [24] M. Mimoso. Volume of NTP Amplification Attacks Getting Louder, Apr 2014. <http://threatpost.com/volume-of-ntp-amplification-attacks-getting-louder/105763>.
- [25] D. Moore, C. Shannon, D. J. Brown, G. M. Voelker, and S. Savage. Inferring internet denial-of-service activity. *ACM Transactions on Computer Systems (TOCS)*, 24(2):115–139, 2006.
- [26] K. Orland. Multiple gaming platforms hit with apparent DDoS attacks, Jan 2014. <http://arstechnica.com/gaming/2014/01/multiple-gaming-platforms-hit-with-apparent-ddos-attacks/>.
- [27] V. Paxson. An analysis of using reflectors for distributed denial-of-service attacks. *ACM SIGCOMM Computer Communication Review*, 31(3):38–47, 2001.
- [28] N. Perlroth. Tally of Cyber Extortion Attacks on Tech Companies Grows, Jun 2014. <http://bits.blogs.nytimes.com/2014/06/19/tally-of-cyber-extortion-attacks-on-tech-companies-grows/>.
- [29] K. Poulsen. FBI busts alleged DDoS Mafia, Aug. 2004. <http://www.securityfocus.com/news/9411>.
- [30] M. Prince. Technical Details Behind a 400Gbps NTP Amplification DDoS Attack, Feb 2014. <http://blog.cloudflare.com/technical-details-behind-a-400gbps-ntp-amplification-ddos-attack>.
- [31] Prolexic. Prolexic Quarterly Global DDoS Attack Report: Q1 2014, Apr. 2014. <http://www.prolexic.com/knowledge-center-ddos-attack-report-2014-q1.html>.
- [32] C. Rossow. Amplification Hell: Revisiting Network Protocols for DDoS Abuse. In *Proceedings of the 2014 Network and Distributed System Security Symposium, NDSS*, San Diego, CA, 2014.
- [33] C. Systems. Cisco Event Response: Network Time Protocol Amplification Distributed Denial of Service Attacks, Feb. 2014. <http://www.cisco.com/web/about/security/intelligence/ERP-NTP-DDoS.html>.
- [34] The Spamhaus Project - PBL. <http://www.spamhaus.org/pbl/>.
- [35] E. Wustrow, M. Karir, M. Bailey, F. Jahanian, and G. Houston. Internet Background Radiation Revisited. In *Proceedings of the 10th ACM SIGCOMM Conference on Internet Measurement (IMC '10)*, Melbourne, Australia, November 2010.
- [36] J. Zhang, Z. Durumeric, M. Bailey, M. Karir, and M. Liu. On the Mismanagement and Maliciousness of Networks. In *Proceedings of the Network and Distributed System Security Symposium (NDSS '14)*, San Diego, CA, February 2014.