Measuring the Leakage of Onion at the Root

"A measurement of Tor's .onion pseudo-top-level domain in the global domain name system"

Matthew Thomas and Aziz Mohaisen

Verisign Labs, VA, USA,
mthomas,amohaisen [at] verisign.com

Abstract. The Tor project provides individuals with a mechanism of communicating anonymously on the Internet. Furthermore, Tor is capable of providing anonymity to servers, which are configured to receive inbound connections only through Tor—more commonly called hidden services. In order to route requests to these hidden services, a namespace is used to identify the resolution requests to such services. A namespace under a non-delegated (pseudo) top-level-domain (TLD) of *.onion* was elected. Although the Tor system was designed to prevent .onion requests from leaking into the global DNS resolution process, numerous requests are still observed in the global DNS. In this paper we will present the state of .onion requests received at the global public DNS A and J root nodes, potential explanations of the leakage, and highlights of trends associated with global censorship events. By sharing this preliminary work, we wish to trigger further discussions on the matter in the community.

Keywords: DNS, privacy, security, Tor

1 Introduction

The Domain Name System (DNS) has become a critical and reliable component of the Internet, allowing individuals to quickly match domain names with their corresponding IP-addresses. The DNS is a hierarchical system, in which at the top of the hierarchy is the root domain. Currently, the root consists of a combination of 13 groups of DNS servers located globally around the world. Each of those servers is named in the form [A-M].root-servers.net. These roots are responsible for the delegation of top-level-domains (TLDs) such as .com [1].

It is well known within the Internet research and engineering community that many installed systems on the Internet query the DNS root for a wide range of TLDs that are not delegated and will ultimately result in an error, or more commonly referred to as a NXDomain [2]. Many of these installed systems depend explicitly or implicitly on the indication from the global DNS that the domain name does not exist. For instance, many internal networks use a domain name suffix that is not currently delegated in the global DNS, such as .corp .home [3]. Due to the recent delegation of new gTLDs within the global DNS [4], several studies have measured the amount of internal name space leakage to

the DNS roots [5, 6]. These unintended leaked DNS queries have been shown to expose sensitive private information and present potential new security threat vectors [5–7]. During the analysis of potential colliding name spaces within the global DNS, queries suffixed in .onion appeared to be one of the more prevalent non-delegated TLDs at the global root DNS.

Tor is an example of a system that exploits the absence of a non-delegated namespace within the global DNS system for its internal use. Hidden services, a unique feature within Tor, provide additional anonymity for users to communicate with servers. To identify these services, Tor uses the .onion name space to identify such requests [8]. While the Tor system was designed to not route requests suffixed in .onion, there exists a clear conflict of interests between internal namespace routing and the global DNS namespace when .onion URLs are shared and or requested [9]. In fact, DNS leakage is a known and well-documented issue within the Tor community. Many tutorials on the Tor website have been published giving users instructions to mitigate the leakage through the use of proxies, disabling DNS pre-fetching within the browser or even installing a local DNS server which rejects .onion addresses [10]. However, non-technical Tor users likely do not practice these mitigation steps due to their complicated nature.

The leakage of .onion requests to the global DNS roots clearly presents some risk to Tor users and also has privacy implications that need to be explored. To this end, in this paper we present a first look at the .onion leakage at the DNS root. We use two root servers, A and J, that are operated by Verisign, and explore .onion resolutions seen at both of them over a period of time close to six months. Our findings highlight that a large amount of .onion traffic is observed at both servers and the requests originate from a diverse set of locations (at the recursive name server level). Furthermore, we illustrate .onion's heavy tailed distribution (with respect to the number of queries per .onion), and a very interesting weekly traffic pattern. We highlight various causes and scenarios of the leakage and call for further investigation into the leakage potential implication on users privacy.

The organization of the rest of this paper is as follows. In section 2 we introduce the DNS profile of the .onion data collected. In section 3, we examine longitudinal patterns of .onion traffic to the A and J root servers operated by Verisign from various network and second-level-domain (SLD) points-of-view, and highlight correlations between global events and increased .onion traffic volumes. In section 4, we explore potential reasons .onion traffic is being leaked to the roots and highlight considerations within the Internet engineering community to address the use of non-delegated TLDs. Finally, in section 5 we will present our conclusions and discuss future directions in which we will further explore the .onion leakage.

2 Data Set

Verisign operates the A and J root servers in the DNS root zone. NXDomain (NXD) responses for the non-delegated TLD .onion were captured over slightly more than *six months* from both root servers starting on September 10th, 2013

and ending March 31st, 2014. The data set consists of approximately 27.6 million NXD records spanning 81,409 unique SLDs. The DNS requests originated from a wide variety of sources: in total, they are sent from 172,170 IP addresses, 105,772 unique /24 net blocks, and 21,345 distinct Autonomous System Numbers (ASNs).

During the multi-month collection period, numerous NXD TLDs appeared at the roots. Based on the total query volume, we ranked the various TLDs and found that the .onion TLD ranked 461 out of 13.8 billion TLDs. The following section will further depict the traffic patterns and trends observed within the .onion TLD.

3 "Onion" DNS Trends and Characteristics

3.1 Traffic Volume and Diversity Measurements

To better understand the overall traffic pattern, a longitudinal study of query volumes including the total number of requests, number of distinct /24 net blocks and the number of distinct ASNs for a given day was conducted, and the results are represented in Figure 1. Overall, we observe that there is a clear upward trend in the total query traffic volume, increasing nearly 300% since the beginning of the collection period. Meanwhile, the diversity of the traffic sources also increased by approximately 20%. One common characteristic that many DNS researchers and network operators are familiar with is the weekly repeated pattern of the volume of requests, as shown in the ASN and /24 measures in Figure 1. These patterns and trends are clear in ".onion's" /24 and the ASN-level measurements; however, this weekly pattern is surprisingly absent when observing .onion total traffic volume. Many other NXD TLDs at the root have been shown to exhibit a regular weekly query volume pattern [11]. It is unclear to us at this time why .onion does not exhibit this common traffic pattern, and that warrants additional investigation to understand this phenomenon.

The data presented in Figure 1 only represents measurements taken from the A and J root nodes. In order to gauge the total global DNS leakage of ".onion" requests, we can segregate the unique SLDs received at each root node and compare their overlap. This measure will provide us with a SLD root affinity and a simple way of estimating total global DNS leakage if this trend was to be extrapolated over all roots.

Figure 2 depicts the number of unique SLDs observed at the A node, J node, and the combination of A and J nodes. In this figure, we can see that the combined A+J roots, on a daily basis, observe about 3300 unique SLDs; while each of the A and J nodes separately observe roughly 2500 unique SLDs—roughly 75% of the combined A+J root nodes. Prior work studying multi-root distinct SLD overlap [11] has shown that the combined traffic observed at A+J constitutes approximately 40% of all observed distinct SLDs for various TLDs spanning the global DNS roots. The .onion SLD root affinities and overlap between the A and J roots are comparable to the finding in the prior literature concerning other

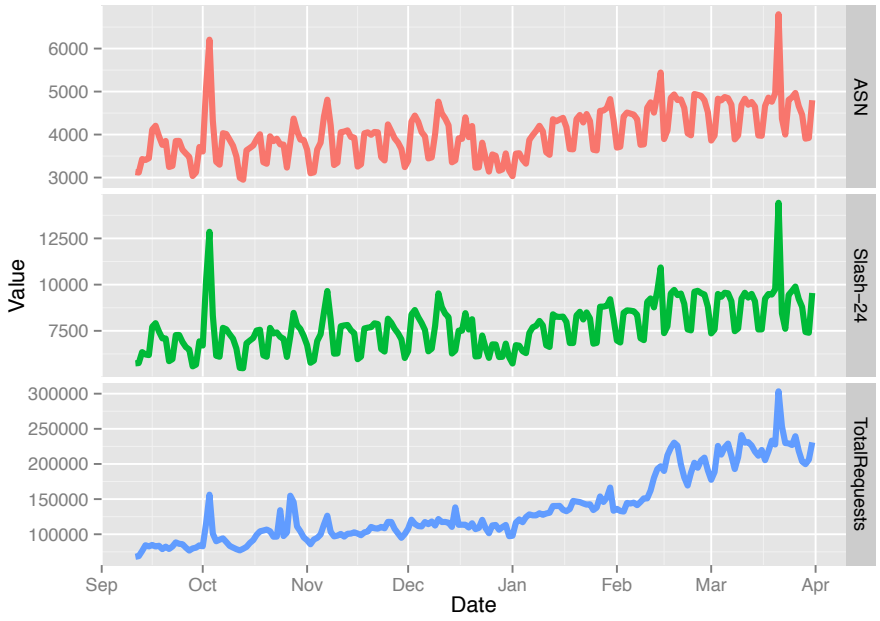

Fig. 1. "Onion" Traffic Measurements Observed at A and J Root DNS Nodes

TLDs [11]. Therefore, we postulate that the .onion traffic observed at A+J would continue such a trend and an appropriate sizing of total global .onion leakage could be roughly estimated. Based on the statistics in section 2, we estimate the total number of .onion NXD records at 69 million over the same period of time.

3.2 Hidden Service and Second Level Domain Measurements

Figure 2 shows a few days in which the absolute number of distinct SLDs dramatically increases from the average number of daily SLDs observed in the rest of the measurement period. We now turn our attention to the overall distribution of requests for a given SLD within the .onion TLD to better understand the DNS request dynamics of all .onion SLDs. Figure 3 provides three different plots of various traffic diversity measurements, namely the number of total requests, the count of distinct /24 net blocks, and ASNs a distinct SLD received during the collection period. The corresponding cumulative distributions of these measures are reflected in Figure 3.

Clearly, the vast majority of SLDs receive a minimal amount of DNS requests over the six months period covered in our data set, with 50% of the SLDs receiving only one request and nearly 90% of SLDs receiving less than 10 requests. A similar trend of minimal traffic source diversity for the majority of SLDs is displayed, where nearly 95% of the SLDs originate from fewer than 10 distinct ASNs; leaving very few SLDs with large amounts of traffic from a wide variety of

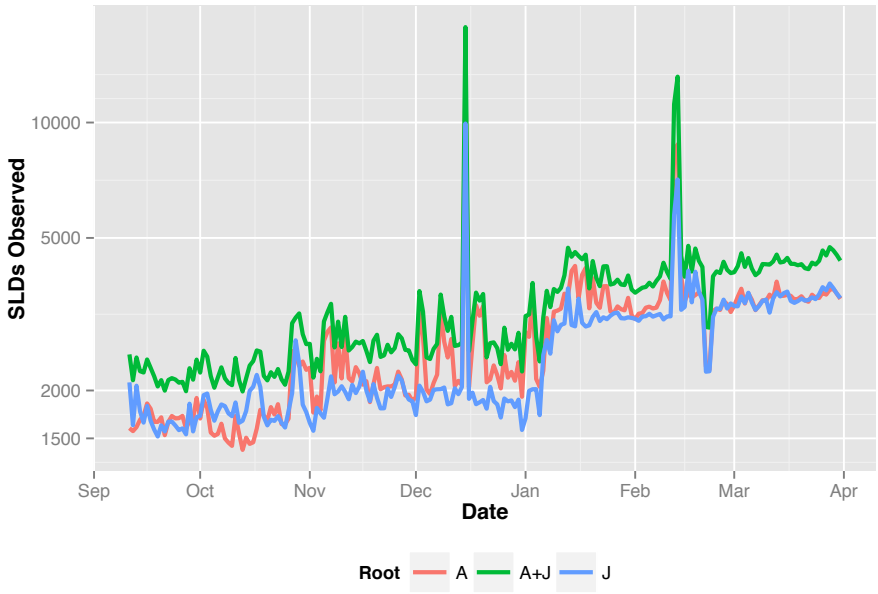

Fig. 2. Global DNS Estimation of Onion By Root

network locations. This pattern is in line with the general traffic characteristics and trend for other non-delegated TLDs.

Next we shift our focus to those few but very popular SLDs within the .onion TLD. Table 1 provides a list of the most requested hidden services along with their total percentage of .onion traffic and the type of service provided using them. The mapping of SLDs to their type of service was constructed manually by searching for references of the hidden service online. The SLDs listed in the table have been anonymized (masked) for privacy concerns, where the first and last two characters of each SLD are shown.

Rank	Anonymized SLD	Type of Service	Traffic (%)
1	Z6---------43	Hidden Tracker	26.5
2	DK---------II	Silk Road	2.1
3	DP---------PC	TorDir	1.7
4	SI---------FK	Silk Road	1.4
5	3G---------4M	Search Engine	1.3
6	JH---------JX	Tor Mail	1.2
7	XM---------SL	Search Engine	1.1
8	AG---------WW	Agora Marketplace	1.1
9	FO---------UI	Bitcoin	0.9
10	TO---------NS	TorLinks	0.9

Table 1: Most Popular SLD Hidden Services and Their Traffic Measurements

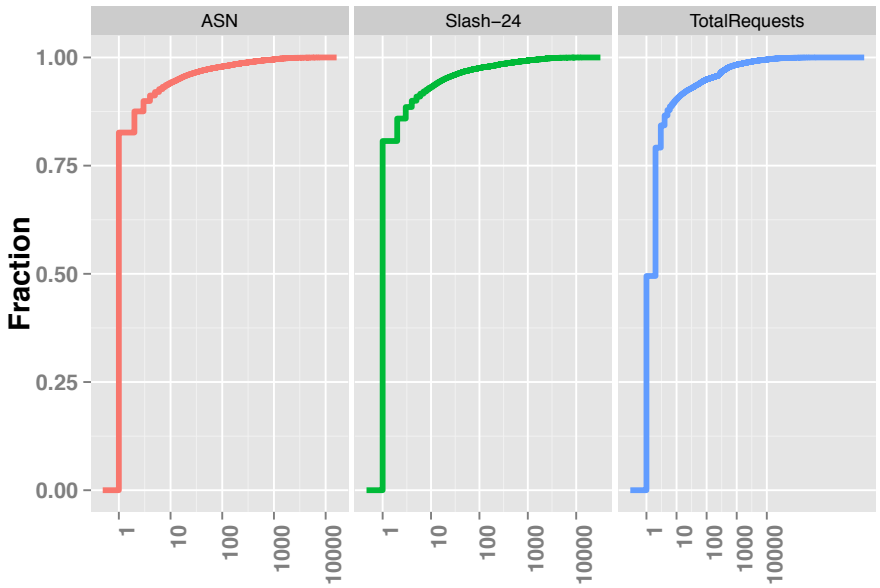

Fig. 3. Cumulative Distribution of SLD Traffic Measurements in "Onion"

From the statistics shown in Table 1, we observe that nearly 27% of all .onion traffic belongs to one hidden service whose focus is on Torrent tracking. The remaining traffic forms a long tailed distribution over the remaining hidden services with an emphasis on services surrounding search, commerce and currency exchange. The top 10 hidden services shown in Table 1 account for more than 38% of the traffic observed over the total period of time of our data set.

3.3 Traffic Source Measurements

In Tables 2, we examine the origination of the .onion DNS requests issued by recursive name servers to the A and J roots from an ASN and country perspective.

The geographical distribution of .onion requestors deviates from the Top-10 countries by directly connecting users as reported by the Tor project over the same period of time. At nearly 36%, the US is 3 times higher than reported from Tor. Other countries such as Germany, France, and Spain also differed significantly, with 7.7%, 7.23% 6.17% and 4.8% respectively [12]. While clearly leaked .onion queries to the global DNS roots and actual Tor connections are very different (e.g. measuring recursive name servers vs. direct connections), the variance in the distribution of the .onion requests may prove helpful in understanding the root cause of the leaked DNS queries.

Country Code	Requests	% Traffic
US	9878093	35.7
RU	2213691	8.0
DE	1482075	5.3
BR	1258468	4.5
CN	996130	3.6
GB	984059	3.5
KR	980656	3.5
PL	918948	3.3
CA	785184	2.8
FR	670103	2.4
AU	510745	1.8
NL	454441	1.6
ES	448171	1.6
IE	425469	1.5
IT	423550	1.5
AR	387594	1.4
MX	363389	1.3
IN	295122	1.0

Autonomous System	Requests	%Traffic
AS15169	2267250	8.2
AS7922	1222955	4.4
AS7018	654680	2.3
AS36692	571609	2.0
AS30607	561349	2.0
AS4766	560739	2.0
AS701	512989	1.8
AS7132	447528	1.6
AS22773	400657	1.4
AS6830	392233	1.4
AS20115	342716	1.2
AS3786	326885	1.1
AS28573	309751	1.1
AS5617	290577	1.0
AS3356	290160	1.0
AS7738	284726	1.0
AS22773	273845	0.9
AS4134	258832	0.9

Table 2: Top Geographical Countries and ASNs Requesting "Onion"

With such a large percentage of .onion requests originating in the United States, it is not surprising to observe the major Internet Service Providers (ISP) in Table 2. However, it is interesting to observe that nearly 8% of all .onion traffic originates from AS15169 (Google). We hypothesize that users/advocates of Tor would most likely not use their default ISP name servers and instead would choose to use public DNS providers such as Google Public DNS or OpenDNS (AS36692, which has a share of 2.06%).

3.4 Global Event Correlation

Global events, such as Internet censorship, political reform, and economic shifts, among others, spur the use of privacy enhancing technologies like Tor. The total traffic volume measured on a daily basis in Figure 1 exhibits several spikes in which .onion traffic significantly increases from its moving average. In order to better understand these events, we cross-correlated the spikes with news stories on global events. Table 3 lists the events and their impact on .onion traffic. These events typically manifest themselves in the form of increased traffic from a specific geographical region or the predominance of queries for a particular SLD. Figure 4 plots the events listed in Table 3 against the total daily ".onion" traffic volume, highlighting the spikes in relation with the rest of the traffic volume over the entire period of time observed in our data set.

Event	Date	Requests	Event
A	10/03/13	156312	Silk Road Shutdown [13]
B	10/24/13	134236	TorATM Traffic Spike [14]
C	10/27/13	154855	URL Posted on Reddit [15]
D	11/07/13	126398	New Silk Road URL [16]
E	12/15/13	138231	Pirate Bay URL Posted [17]
F	03/21/14	303347	Multiple URLs Posted on Reddit [18]

Table 3: Global Events and Elevated "Onion" Request Correlation

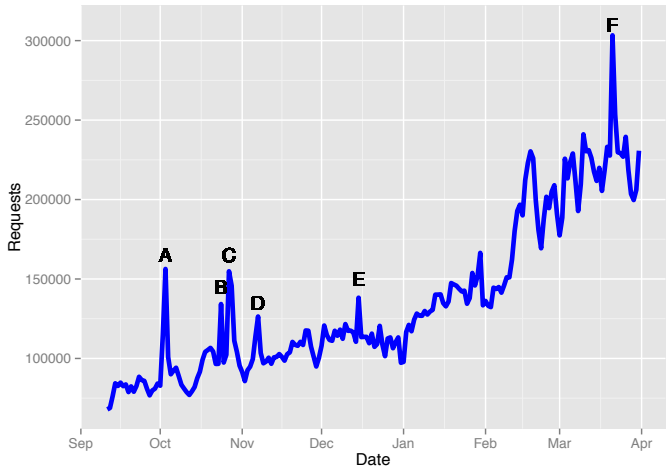

Fig. 4. Global Events and Elevated "Onion" Request Correlation

Certain global events such as the censorship of Internet domains in Turkey may span a longer period of time than a few days. Figure 5 depicts the number of requests for .onion domains originating from Turkey over the multi-month collection period. There is a clear upward trend and a sudden increase in the second half of March 2014 when many DNS-based censorship events took place. The requests originating from Turkey during the censorship spanned hundreds of unique SLDs and were spread over several ASNs.

4 Root Cause Exploration and Namespace Management

Applications electing to use non-delegated TLDs as a namespace in which they seed their routing and resolution processes face scenarios in which possible DNS leakage may occur. Tor has been specifically designed to prevent .onion requests from leaking within the application into the global DNS infrastructure. However,

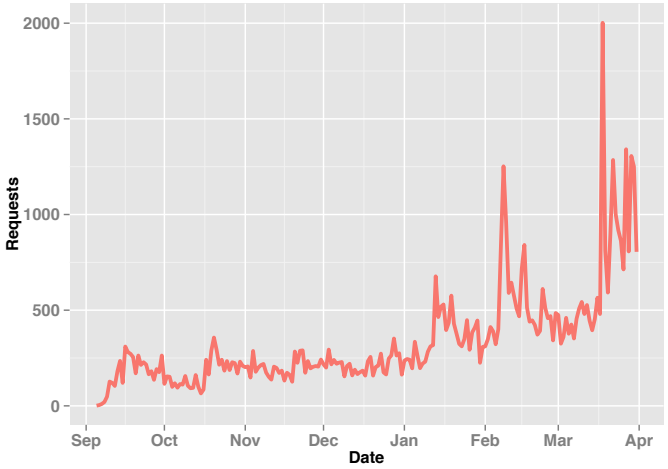

Fig. 5. "Onion" Traffic Measurements From Turkey

it is clear from the measurements we presented so far that a significant volume of requests are being issued to the global DNS root servers. Whether they are initiated by users by mistake or caused by a misconfiguration in the underlying application, such as Tor, or the web browsers, leaked DNS queries outside of the Tor network have a significant implication to individuals' privacy and safety. To that end, understanding the causes of the leakage may help reducing the risk at the user side.

There are many plausible reasons or mechanisms in which .onion queries could be generated and observed in the global public DNS; however, the root cause of how and why these queries are being requested within the global DNS remains unclear. We have seen in this paper numerous global events that spurred additional query volume. One potential explanation associated with surge in the volume of .onion domains in those times is users errors, in which users are not aware that the addresses of hidden services should be run on top of Tor (i.e., by first installing Tor plug-in associated with the browser). Other notable explanations may include browser prefetching, third party application or plug-ins, DNS suffix search lists, web crawlers, and malware.

Advanced families of malware are also now utilizing Tor within their Command and Control (C&C) infrastructure. Cyber-criminals may use Tor and its Hidden Services in order to avoid detection and prevent takedowns. Several cyber-criminals have now started actively using Tor to host malicious infrastructure via Hidden Services. Variants of Zeus [19], CyrptoLocker [20], Chew-Bacca [21], CryptorBit [22] and Torec [23] have all been found to use various aspects of the Tor network, including hidden services. Possible misconfigurations within these malware pieces could facilitate a percentage of the leaked DNS re-

quests. To that end, we observed numerous requests for .onion SLDs associated with these malware samples during our analysis.

Focus within the Internet Engineering community has recently increased on ways for applications to properly use non-delegated domains. A recent Internet draft describes several special-use domain names of peer-to-peer name systems and is seeking approval from the Internet Engineering Steering Group (IESG) [9]. Discussions about the proposal on the DNS operators mailing list have brought forth other generic solutions such as proposed .alt alternative TLD in which applications would safe anchor namespace under it [24]. Blurred lines of authority, privacy and security makes solving such a namespace problem difficult to solve and appease all parties.

5 Conclusion and Future Work

In this paper we looked at a sample of .onion DNS requests issued to the A and J root nodes of the global DNS infrastructure. We examined the unique characteristics of these requests longitudinally as well as the dynamics of requests received from a geographical and network location for unique SLDs. We found that increased traffic spikes within the global DNS for .onion requests corresponded with external global events, highlighting the potential human factor in those leakages (i.e., user error). While the root cause of these leaked DNS queries remains unknown, our preliminary investigation unveiled concerns to the severity of the leakage and to the possibility of more sensitive private information being unintentionally exposed. Our future work will continue the examination of leaked DNS queries to the root but will also extend to other non-delegated TLDs such as i2p and .exit. We will plan to further dissect the impact of global events and the role of malware in the leakage, and investigate the potential privacy consequences of the leakage under the various leakage causes. By sharing this preliminary work, we wish to trigger further discussion in the community.

References

1. Mockapetris, P., Dunlap, K.J.: Development of the domain name system. Volume 18. ACM (1988)
2. Security, I., (SSAC), S.A.C.: Invalid top level domain queries at the root level of the domain name system. http://bit.ly/1mDxRJO (2010)
3. Chapin, L., McFadden, M.: Reserved top level domain names. RFC 2606, http://bit.ly/1nIQ5cS (2011)
4. —: New Generic Top-Level Domains. ICANN, http://newgtlds.icann.org/en/ (2014)
5. —: New gTLD Security, Stability, Resiliency Update: Exploratory Consumer Impact Analysis. http://bit.ly/QB6ntp (2013)
6. Interisle Consulting Group, LLC: Name collision in the DNS. ICANN, http://bit.ly/1iQVj5F (2013)
7. Simpson, A.: Detecting search lists in authoritative DNS. In: Workshop and Prize on Root Causes and Mitigation of Name Collisions (WPNC'14). (2014)

8. The Tor Project: Tor: Overview. The Tor Project, http://bit.ly/1dZ2zvZ (2014)
9. Grothoff, C., Wachs, M., Wolf, H., Appelbaum, J.: Special-use domain names of peer-to-peer name systems. IETF Internet Draft (2013)
10. Garcia, R.: Preventing tor DNS leaks. The Tor Project, http://bit.ly/1royLtU (2014)
11. Thomas, M., Labrou, Y., Simpson, A.: The effectiveness of block lists in preventing collisions. In: Workshop and Prize on Root Causes and Mitigation of Name Collisions (WPNC'14). (2014)
12. The Tor Project: Tor metrics portal: Users. The Tor Project, http://bit.ly/1hrHqGp (2014)
13. —: FBI arrest 'silk road' owner Ross William Ulbricht, shut down tor's most notorious black market. Huffington Post UK, http://huff.to/1fu0tA7 (2013)
14. Bitcoin wiki: Toratm. https://en.bitcoin.it/wiki/TorATM (2013)
15. —: ELI5: What exactly is the "deep web". Reddit, http://bit.ly/1l7hLbz (2013)
16. Biggs, J.: Silk road 2.0 rises again. TechCrunch, http://tcrn.ch/QB5HnQ (2013)
17. Zournas, K.: Pirate bay relocates to thepiratebay.ac. http://bit.ly/1iQNEEz (2013)
18. —: People who have visited the 'deep web' what was it like and why did you do it? http://bit.ly/ROuupk (2014)
19. Tarakanov, D.: The inevitable move - 64-bit zeus has come enhanced with tor. SecureList http://bit.ly/1mIuAeR (2013)
20. Smtih, M.: Cryptolocker crooks charge 10 bitcoins for second-chance decryption service. Network World, http://bit.ly/ROxhPd (2013)
21. Schwartz, M.J.: Chewbacca malware taps tor network. Dark Reading, http://ubm.io/1nrFfKY (2013)
22. Abrams, L.: Cryptorbit and howdecrypt information guide and faq. Bleepingcomputer, http://bit.ly/1eoKEjh (2014)
23. Kovacs, E.: Backdoor.androidos.torec.a: First tor-based trojan for android. http://bit.ly/1pte18L (2014)
24. Wouters, P.: DNS operation mailing list. DNSOP, http://bit.ly/1roTIXw (2013)