

SWEEPING THE IP SPACE: THE HUNT FOR EVIL ON THE INTERNET

Dhia Mahjoub
OpenDNS, USA

Email dhia@opendns.com

ABSTRACT

The total IPv4 space consists of 4 billion addresses, the public ASN visible space consists of 46,000+ AS numbers, and the BGP prefix space consists of 520,000+ prefixes. Together, they form the foundation of addressing, routing and hosting on the Internet. Most of the current reputation systems used for network-level threat detection derive scores for IPs, BGP prefixes or ASNs based on hosted content.

In this paper, we take a novel approach by exploring the AS graph which models the interconnections between ASNs. We uncover hotspots of maliciousness by analysing AS graph topology, hosted content and IP space reservation, and shed some light on suspicious relationships between ASNs and abusive IP sub-allocations.

This exploration methodology enriches classical scoring mechanisms that are based on the counting of malicious domains/IPs hosted on ASNs. This method also provides actionable intelligence and can be used pre-emptively to detect and block malicious IP infrastructures before or immediately after they are set up for waging malware campaigns. We will go over multiple relevant use cases of attack domains detected by this system, such as trojan C&Cs, exploit kit domains and malware domains.

OVERVIEW

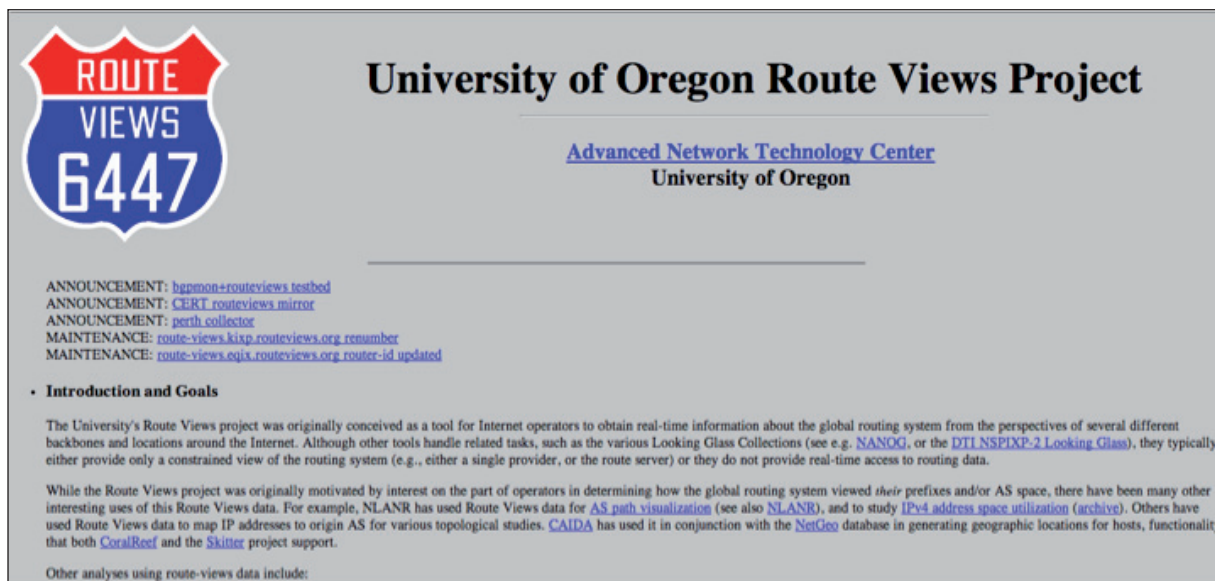
Classical reputation systems used for network-level threat detection assign scores to IPs, BGP prefixes and ASNs based on counting the number of hosted malicious domains or IPs. In this study, our goal is to assess malicious IP ranges in certain ASNs from a new perspective. We look beyond the simple counting of the number of bad domains and IPs hosted on prefixes of an ASN, by exploring the topology of the AS graph, and looking at a finer granularity than the BGP prefix (sub-allocated ranges within BGP prefixes).

Previous research has been conducted on malicious ASNs. For example, in [1], Stone-Gross *et al.* assign scores to rogue ASNs based on the number of events involving hosts engaged in phishing or spamming, hosting drive-by download malware, or generating botnet traffic. In [2] and [3], the authors use visualization to track security incidents and malware events drawn from blacklist databases, and [4] explores ASNs providing transit for malicious ASNs.

ASN GRAPH

Every globally routable network on the Internet is identified by an Autonomous System Number (ASN). An Autonomous System (AS) represents a collection of IPv4 and IPv6 network prefixes or CIDRs administered by the same entity and sharing a common routing policy. In practice, an AS announces the prefixes under its authority to its peering ASNs, and these announcements propagate across routers where they are stored to help with routing decisions.

The BGP table is the accumulation of all announced prefixes with their reachability information (AS paths). An AS path is a sequence of ASNs through which an announced prefix can be reached [5]. The BGP table is not only important for packet forwarding and loop detection on every Internet router, it is also very useful to study the evolution of the Internet from a topology



ROUTE VIEWS 6447

University of Oregon Route Views Project

Advanced Network Technology Center
University of Oregon

ANNOUNCEMENT: [bgpmon+routeviews.testbed](#)
ANNOUNCEMENT: [CERT.routeviews.mirror](#)
ANNOUNCEMENT: [perth.collector](#)
MAINTENANCE: [route-views.kixp.routeviews.org.renumber](#)
MAINTENANCE: [route-views.eqix.routeviews.org.router-id.updated](#)

- **Introduction and Goals**

The University's Route Views project was originally conceived as a tool for Internet operators to obtain real-time information about the global routing system from the perspectives of several different backbones and locations around the Internet. Although other tools handle related tasks, such as the various Looking Glass Collections (see e.g. NANOG, or the DTLNSPIXP-2 Looking Glass), they typically either provide only a constrained view of the routing system (e.g., either a single provider, or the route server) or they do not provide real-time access to routing data.

While the Route Views project was originally motivated by interest on the part of operators in determining how the global routing system viewed *their* prefixes and/or AS space, there have been many other interesting uses of this Route Views data. For example, NLANR has used Route Views data for [AS path visualization](#) (see also [NLANR](#)), and to study [IPv4 address space utilization](#) ([archive](#)). Others have used Route Views data to map IP addresses to origin AS for various topological studies. CAIDA has used it in conjunction with the [NetGeo](#) database in generating geographic locations for hosts, functionality that both [CoralReef](#) and the [Skitter](#) project support.

Other analyses using route-views data include:

Figure 1: Route Views website.

and security threat perspective. For that, we need to build an AS graph which represents the interconnections between peering ASNs.

In this study, we built an AS graph using publicly available data sources. Our primary source is the Route Views data from the University of Oregon [6], which provides a global BGP table by collecting BGP data from hundreds of Autonomous Systems worldwide. We can also use the BGP tables from all the routers we operate within our *OpenDNS* global network of 23 data centres. The current global BGP table counts 500,000+ IPv4 BGP prefixes and 46,000+ ASNs.

Other valuable data sources that are useful for studying the IP, BGP prefix and ASN landscapes are the CIDR report [7] and Hurricane Electric Internet Services website [8].

BUILDING THE ASN GRAPH

BGP data is collected in MRT format. When we dump it in text format, an entry is as follows:

```
TABLE_DUMP2|1392422403|B|96.4.0.55|11686|67.215.94.0/24
411686 4436 2914 36692||GPI96.4.0.55|O|NAG||
```

We mark the fields that are of interest to us in bold. In this entry, 67.215.94.0/24 is an example network prefix, and 11686 4436 2914 36692 is one associated AS path. The ASN that appears at the end of the AS path is the origin ASN of the prefix, which is 36692 in this case. Typically, it is the ASN that announces (advertises) that prefix.

We represent the AS graph as a directed graph, where an ASN is denoted by a node and there is a directed edge between an ASN and every one of its upstream ASNs. For example, in the BGP table entry above, 36692 is the origin ASN for 67.215.94.0/24, and 2914 is an upstream ASN of 36692 (the last ASN before reaching the origin ASN when packets are travelling towards an IP in the origin ASN), therefore that entry can be graphically represented as shown in Figure 2.



Figure 2: Graphic representation of an entry of the BGP table.

An alternative method to build the AS graph is to use the entire AS path on every prefix entry of the BGP table. In this case, from the example above, we can build the following edges in the graph: **36692->2914**, **2914->4436**, **4436->11686**. The AS graph is built by parsing the BGP table line by line.

In the directed AS graph, an AS node can have incoming and/or outgoing edges. The outgoing edges point to upstream ASNs and incoming edges originate from downstream ASNs. Below, we define a few terms describing the AS graph nodes from a directed graph topological perspective [9].

A ‘source’ ASN is an ASN that only has outgoing edges and no incoming edges, i.e. the ASN only has upstream ASNs that it relies upon for connectivity and for propagating its prefix announcements. A ‘leaf’ ASN is a special case where an ASN has a single outgoing edge and no incoming edge. This is described as a ‘stub’ ASN in the BGP routing terminology.

We define ASNs that are ‘source’ ASNs (or ‘leaves’) that share the same parents (upstream ASNs) as ‘sibling’ ASNs. For clarity, we will use the more intuitive term ‘peripheral’ ASNs to denote source ASNs for the remainder of this paper.

The BGP table/ASN graph is constantly changing as new prefixes (with their AS paths) are announced, old prefixes are dropped, new ASNs are registered and start advertising prefixes, and others cease to exist and withdraw all their prefixes. Most common changes are probably caused by new AS relations, new peers or previously unseen relations.

This dynamic state can be the result of multiple factors: intentional technical and business decisions, human errors, hardware faults, route hijacking, etc. By parsing the entries of the BGP table, we can extract two types of useful data: the upstream and downstream ASNs of every ASN, and IP to ASN maps (via prefix to ASN mapping). For this, we can load the prefix and the origin ASN data into a radix tree. With the radix tree (given an IP as input), we can quickly find the best matching prefix, and consequently, matching ASN.

Alternatives are to use services like BGPMON.net (e.g. `whois -h whois.bgpmon.net 8.8.8.8`), *Team Cymru* IP to ASN mapping [10], GeoIPASNum.dat from maxmind [11], or `http://ipinfo.io/` (e.g. `curl ipinfo.io/8.8.8.8/org` returns the AS number and AS name of *Google Inc.*). In this study, we discuss interesting patterns in the AS graph topology – typically, suspicious peripheral ASNs that are siblings, i.e. they share common parents (upstream ASNs) in the AS graph. By clustering peripheral nodes in the AS graph by country, we found that certain peripheral sibling ASNs in a few countries have been delivering similar suspicious campaigns.

USE CASE 1: SUSPICIOUS SIBLING PERIPHERAL ASNS

During manual investigation of suspicious domains and IPs that we detected in our traffic, we observed several cases of sibling peripheral ASNs that are hosting similar malware payloads. In this section, we will describe one such use case.

In Figure 3, we show a snapshot of a suspicious ASN subgraph taken on 8 January 2014, consisting of 10 sibling peripheral ASNs (57604, 8287, 50896, 49236, 29004, 45020, 44093, 48949, 49720 and 50818) sharing two upstream ASNs (48361 and 31500). We colour the ASNs that were hosting malicious payloads in red. The malicious payload is identified by some anti-virus vendors as Trojan-Downloader.Win32.Ldmon.A [12, 13] and described as a Trickler [14]. Notice that most of these peripheral ASNs are small-scale with a single prefix, as shown in Table 1.

Figure 4 shows the same suspicious ASN subgraph as shown in Figure 3, but with the snapshot taken six weeks later (on 21 February). Notice the change in subgraph topology: more leaves




Figure 3: Malicious ASN subgraph.



Figure 4. Malicious ASN subgraph six weeks later.

ASN	No of prefixes	Prefixes
57604	1	91.233.89.0/24
8287	3	91.213.72.0/24 91.213.93.0/24 91.217.162.0/24
50896	5	195.78.108.0/23 91.198.127.0/24 91.200.164.0/22 91.201.124.0/22 91.216.3.0/24
49236	1	62.122.72.0/23
29004	1	195.39.252.0/23
45020	1	194.29.185.0/24
44093	1	193.243.166.0/24
48949	1	95.215.140.0/22
49720	1	194.242.2.0/23
50818	1	194.126.251.0/24

Table 1: Sibling peripheral ASN prefixes.

started hosting the same suspicious payloads (via new resolving domains or directly on the IPs). Additionally, AS31500 detached itself from the leaves by ceasing to forward their prefix announcements.

We observed that a large pool of contiguous IPs in the /23 or /24 prefixes of these ASNs were hosting the same aforementioned type of payload. In most cases, the payload URLs were live on the entire range of IPs before any domains were hosted on them. Furthermore, the IPs were set up with the same server infrastructure. For instance, we took a random sample of 160 live IPs from this subgraph.

In this sample, 50 IPs had a similar nmap fingerprint:

```
22/tcp open  ssh          OpenSSH 6.2_hpn13v11
(FreeBSD 20130515; protocol 2.0)
8080/tcp open  http-proxy  3Proxy http proxy
Service Info: OS: FreeBSD
```

and 108 IPs shared the following fingerprint:

```
22/tcp open  ssh          OpenSSH 5.3 (protocol 1.99)
80/tcp open  http?
```

In total, this subgraph featured 3,100+ malware domains on 1,020+ malware hosting IPs, and it is clear this IP infrastructure across multiple ASNs was set up in bulk and in advance to deliver the same rogue campaign [13].

USE CASE 2: ROGUE ASN DE-PEERED OR HIDDEN

In this section, we discuss one case among many we observed of rogue peripheral ASNs that serve various malware content. In this example, it is AS48031, XSERVER-IP-NETWORK-AS PE Ivanov Vitaliy Sergeevich 86400 that had a single upstream provider, AS15626. AS48031 has been hosting browser-based ransomware, porn sites, spam, and radical forums.

Browser-based ransomware, or ‘Browlock’, is a rudimentary piece of ransomware that consists of an HTML page that loads when the user visits the browlock domain. It locks the browser screen (through HTML or JavaScript code) and demands payment, supposedly either for possession of illegal material or for usage of illegal software [15]. This is more of a scam than real ransomware (which corrupts or encrypts the user’s data), because the browlock alert can be neutralized simply by killing the browser task. Despite its simplicity, Browlock has been around for a couple years, is targeting users in a large number of countries, and seems to be generating profit for the criminals. Browlock has been delivered by dedicated domains (domains specifically registered for malicious intent) as well as compromised ones.



Figure 5: Browlock web page.

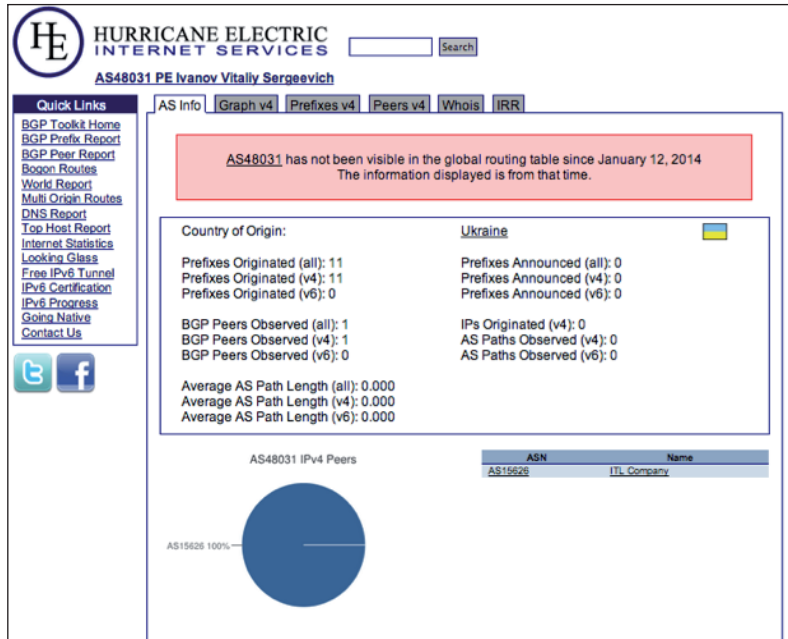


Figure 6: AS48031 disappears off the global BGP routing table.

In Table 2, we show the prefixes announced by AS48031 in the autumn of 2013. A few months later, in January 2014, AS48031 stopped advertising prefixes and disappeared from the global routing table, as shown in Figure 6.

However, those prefixes did not actually disappear, and AS48031’s only parent in the AS graph, its upstream peer AS15626, took over announcing them, as shown in Figure 7.

The rogue IPs in those prefixes continued to host malware content.

The question remains as to whether AS15626 had been abused by its downstream client AS48031 to host malware, and it acted responsibly by ceasing to announce those prefixes when it took notice of the malicious content on AS48031 prefixes, or whether both AS48031 and AS15626 are complicit in hosting malware,

and AS15626 is simply being evasive by hiding AS48031 from the global routing table and yet retaining connectivity to the rogue IPs by announcing their prefixes. There are several such suspicious cases that occur on the BGP routing space.

Prefixes
176.103.48.0/20
193.169.86.0/23
193.203.48.0/22
193.30.244.0/22
194.15.112.0/22
196.47.100.0/24
91.207.60.0/23
91.213.8.0/24
91.217.90.0/23
91.226.212.0/23
91.228.68.0/22
93.170.48.0/22
94.154.112.0/20

Table 2: Prefixes announced by AS48031 in autumn 2013.

Prefix	Description
5.34.176.0/21	ITL Company
31.44.188.0/24	VERATON PROJECTS LTD
46.28.64.0/21	ITL Company
82.118.16.0/21	ITL Company
91.207.60.0/23	PE Ivanov Vitaliy Sergeevich
91.213.8.0/24	Hosting Solutions Ltd.
91.217.90.0/23	PE Ivanov Vitaliy Sergeevich
91.226.212.0/23	PE Ivanov Vitaliy Sergeevich
91.228.68.0/22	Hosting Solutions Ltd.
91.235.128.0/23	PE Dobrogivskiy Muroslav Petrovich
91.238.102.0/24	Puchkov Yuriy Volodimirovich PE
130.0.232.0/21	3nt solutions LLP
146.185.235.0/24	VPS & DEDICATED SERVERS SOLUTIONS
146.185.240.0/24	pool for VPS & dedicated servers
176.103.48.0/20	PE Ivanov Vitaliy Sergeevich
193.30.244.0/22	Maxim Odintsev
193.169.86.0/23	PE Ivanov Vitaliy Sergeevich
193.203.48.0/22	PE Ivanov Vitaliy Sergeevich
193.238.152.0/23	PF "Volodymyr Lyakh"
194.15.112.0/22	Hosting Solutions Ltd.
195.54.162.0/23	PE Dobrogivskiy Muroslav Petrovich
217.12.192.0/19	ITL Company

Figure 7: Former prefixes of AS48031 now announced by the upstream AS15626.

USE CASE 3: MALICIOUS SUB-ALLOCATED RANGES

In this section, we summarize a study we conducted for five months between October 2013 and February 2014 that consisted of monitoring rogue sub-allocated ranges on OVH IP space [16], where these ranges are reserved by recurring suspicious customers and used to serve Nuclear Exploit Kit domains. In this type of infection, visitors are lead to the exploit landing sites through malvertising campaigns, then malware is dropped

on victims' machines (e.g. Zbot). The results of the study were published in [17].

For several months, OVH IP ranges had been abused. Notably, the IPs were used exclusively for hosting Nuclear Exploit subdomains, with no other sites sharing the IPs. These IPs were reserved in small ranges from OVH Canada and set up with identical services (nmap fingerprint). Consulting ARIN's referral whois database showed the reserved ranges and customer IDs. As an evasive measure, on 7 February, the bad actors moved their activities to besthosting.ua, a Ukrainian hosting provider. RIPE's whois service, which covers European IP space, does not always give details of reserved ranges and customers, but in this case the Ukrainian IPs were still set up with identical services. Therefore, we flagged them as prone to serve the same Nuclear campaign. On 14 February, the bad actors moved to a Russian provider, pinspb.ru, with a similar bulk IP range set-up. On 22 February, they moved back to OVH, notably changing their MO: the IPs being used have been allocated and used in the past for other content. This could be an evasion technique or a case of resource recycling.

However, although the bad actors have migrated between hosting providers to host the Nuclear Exploit serving domains, they still kept the name servers infrastructure (authoritative for the Nuclear domains) on ranges reserved on OVH by the same customers, which allows us still to track them. Thanks to a great collaboration with the non-profit security research group MalwareMustDie, a large number of Nuclear Exploit domains that were active at the time have since been taken down [18].

Subsequently, bad actors have been circulating between OVH and other hosting providers. Lately, compromised domains, especially GoDaddy domains, have been used to host Nuclear and Angler Exploit kit domains (as we will cover later).

USE CASE 4: PREDICTING THE IP INFRASTRUCTURE OF MALICIOUS DOMAINS

As part of the study described in the previous section, we have been monitoring IP ranges reserved on OVH Canada by the suspicious customer(s) who reserved the ranges hosting Nuclear Exploit Kit domains. Table 3 shows the number of reserved ranges, the total number of IPs they represent, and the number of IPs effectively used for malicious purposes during the months of December 2013, January 2014, February 2014 and early March 2014. These IPs were used to host Nuclear Exploit Kit domains, Nuclear domains' name servers, and Browlock domains.

Reservation dates	No. ranges	No. IPs	No. IPs used
1 to 31 Dec 2013	28	136	86
1 to 31 Jan 2014	11	80	33
1 to 28 Feb 2014	4	28	26
1 to 20 Mar 2014	43	364	215
7 Mar 2014	40	352	208
10 Mar 2014	3	12	7

Table 3: IP ranges reserved by suspicious customers.

Looking at the prefixes to which these malicious reserved sub-ranges belong, we notice that all 86 ranges described in Table 3 are concentrated in four large OVH prefixes, as shown in Table 4.

No. IPs	BGP prefix
388	198.50.128.0/17
128	192.95.0.0/18
80	198.27.64.0/18
12	142.4.192.0/19

Table 4: BGP prefixes of the rogue reserved ranges.

We used two investigative techniques to track rogue IP ranges: the first is to monitor sub-allocated ranges reserved by suspicious customers. The second technique is to monitor the IPs' service fingerprints. Below, we review a few examples of the IP ranges used to host Nuclear Exploit domains [17]:

1. For the IPs hosted on besthosting.ua, the live IPs in the range 31.41.221.131 to 31.41.221.143 all have the same server set-up (nmap fingerprint):

```
22/tcp open  ssh      OpenSSH 5.5p1 Debian
6+squeeze4 (protocol 2.0)
80/tcp open  http     nginx web server 0.7.67
111/tcp open rpcbind
```

2. For the IPs hosted on pinspb.ru, the IPs in the range 5.101.173.1 to 5.101.173.10 have the following fingerprint:

```
22/tcp open  ssh      OpenSSH 6.0p1 Debian 4
(protocol 2.0)
80/tcp open  http     nginx web server 1.2.1
111/tcp open rpcbind
```

3. For the IPs hosted on OVH, the IPs in the range 198.50.143.64 to 198.50.143.79 have the following fingerprint:

```
22/tcp open  ssh      OpenSSH 5.5p1 Debian
6+squeeze4 (protocol 2.0)
80/tcp open  http     nginx web server 0.7.67
445/tcp filtered microsoft-ds
```

The IPs used to host the name servers also had the same fingerprints [17]. Notice that initially the malware IPs in a given range used to become active in bulk and sequential order, but later, as an evasion method, the bad actors started bringing them up at random, one by one or a few at a time, right when they are about to deliver the exploit kit attack.

The combination of the two investigative techniques made it possible to predict the next attack IPs with practically no false positives. As hosting providers have become more aggressive in suspending rogue customers' accounts and swifter in taking down malware IPs, and as bad actors choose hosting providers on IP space where the RIRs' whois service does not always provide full information about reserved ranges and customers (e.g. RIPE), the first technique might not always work. The second technique of fingerprint tracking, however, still provides accurate results when combined with other intelligence.

USE CASE 5: DETECTING MALICIOUS SUBDOMAINS UNDER COMPROMISED DOMAINS

In this section, we discuss the results of a five-month study we conducted between February and June 2014 that followed the study of use case 3. For this project, we designed a system to pre-emptively detect malicious subdomains injected under compromised domains (particularly GoDaddy domains) and track their IP infrastructure. The phenomenon of compromised GoDaddy domains serving malware has been around for at least two years [19]. The compromise can happen through at least two methods: hacking GoDaddy accounts or injecting malicious redirection scripts into vulnerable GoDaddy




Figure 8: Top hosting malicious ASNs.

websites. When the compromise is successful, subdomains (third-level domains) are injected under the *GoDaddy* domains (second-level domains), and these subdomains resolve to malicious sites.

Most abused ASNs

By monitoring this threat from February 2014 to the present day, we observed that the subdomains resolve to IPs serving exploit kit attacks (typically Nuclear [20, 21] and Angler [22, 23]), and also browser-based ransomware. We recorded several hundred IPs hosting these malicious subdomains over the period of the study. In Figure 8, we show the top ASNs hosting these malware subdomains over that period. The x-axis represents the AS numbers, and the y-axis represents the percentage of IPs from the total dataset that map to a particular ASN.

We see that the top five abused ASNs are:

- 16276 OVH SAS
- 24961 myLoc managed IT AG
- 15003 Nobis Technology Group, LLC
- 41853 LLC NTCOM
- 20473 Choopa, LLC

AS16276, which is *OVH*, hosted 18% of the total malicious IPs. In this specific case, since the abuse of *OVH* has been exposed through February 2014 (particularly for hosting Nuclear Exploit domains [17]), bad actors have changed their MO: they switched temporarily to other hosting providers, and started using recycled IPs (not reserved exclusively for exploit domains). Additionally, *OVH* took action by suspending rogue accounts. However, by monitoring the compromised domains' campaigns, we observed that *OVH* was still being abused by bad

actors to host malicious content. These were the general changes in the bad actors' MO that we observed:

- From a domain perspective, for a while, bad actors have been abusing various ccTLDs (e.g. .pw, .in.net, .ru, etc.) facilitated by rogue or victim registrars and resellers. Then, they supplemented that approach with using compromised domains, particularly *GoDaddy* domains, under which they inject subdomains to host exploit kit landing URLs and Browlock. (Notice that using compromised domains for attacks goes further back in the past for other different campaigns.)
- From an IP perspective, bad actors used to bring the attack-hosting IPs online in contiguous chunks, then they started bringing them up in randomized sets or one IP at a time.
- The other notable fact is that bad actors used to abuse *OVH Canada* (attached to ARIN) where rogue customers were reserving re-assigned small ranges (/27, /28, /29, etc.). By consulting the ARIN RWhois database, it was possible to correlate the rogue customers with the IP ranges they reserved and therefore predict and block the IP infrastructures they set up for exploit kit attacks. As the adversaries changed MO, this method of tracking became less effective.
- The shift became clear when they started to use ranges on *OVH's* European IP space (which is attached to RIPE) more frequently, as well as other European providers. Typically, we saw small gaming hosting providers being abused among other platforms.

Additionally, although the standard geolocation of *OVH* European IP space maps to France (FR), the attack IP ranges were reserved from *OVH's* server pools in various European countries (France, Belgium, Italy, UK, Ireland, Spain, Portugal,

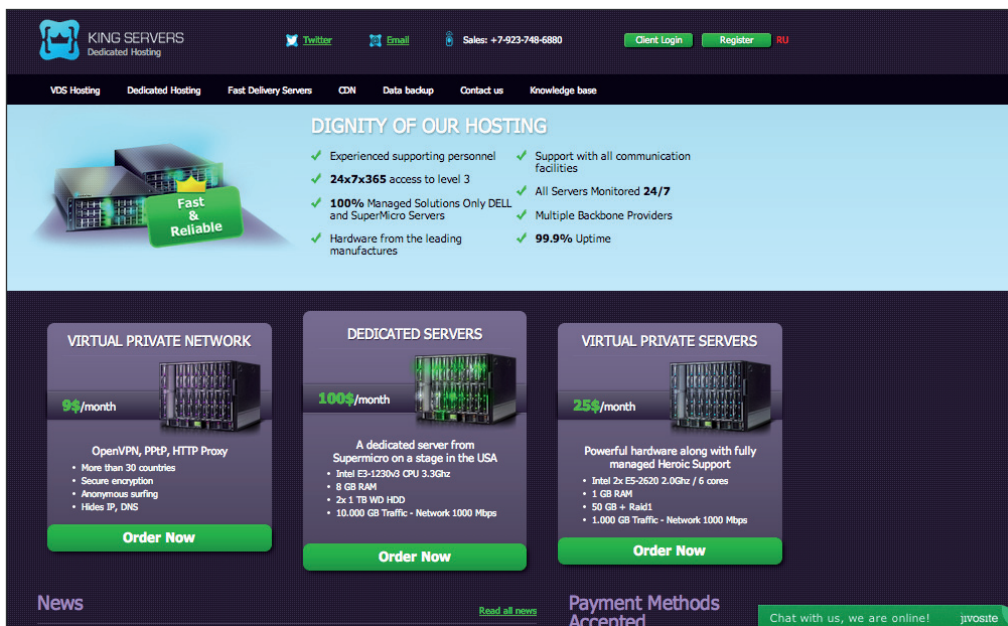


Figure 9: King Servers main website.

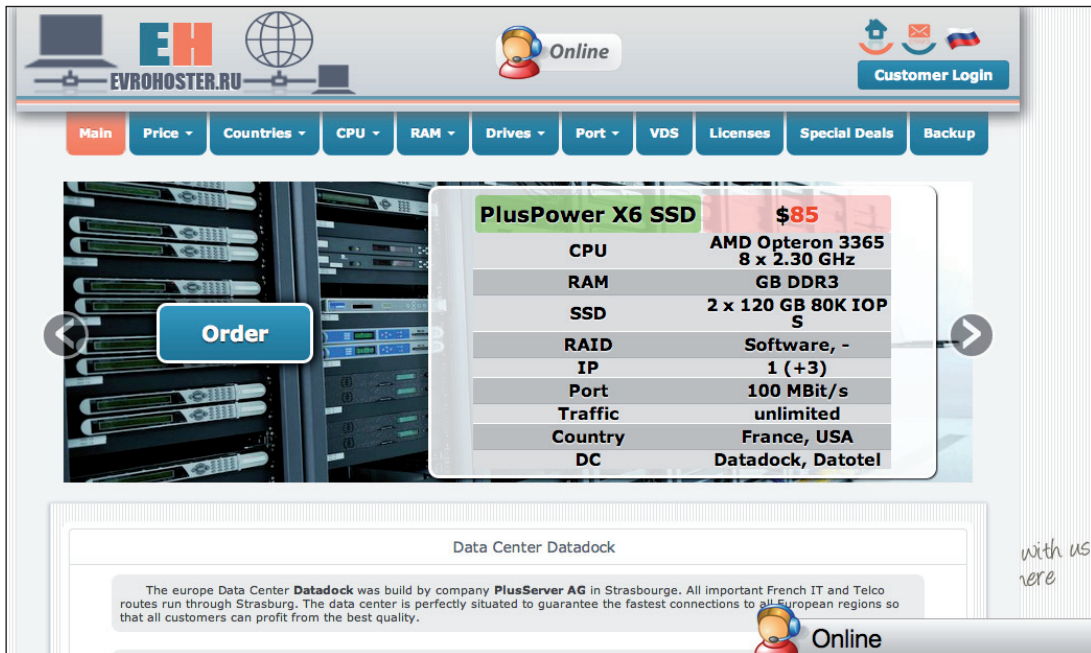


Figure 10: Evrohoster.ru main website.

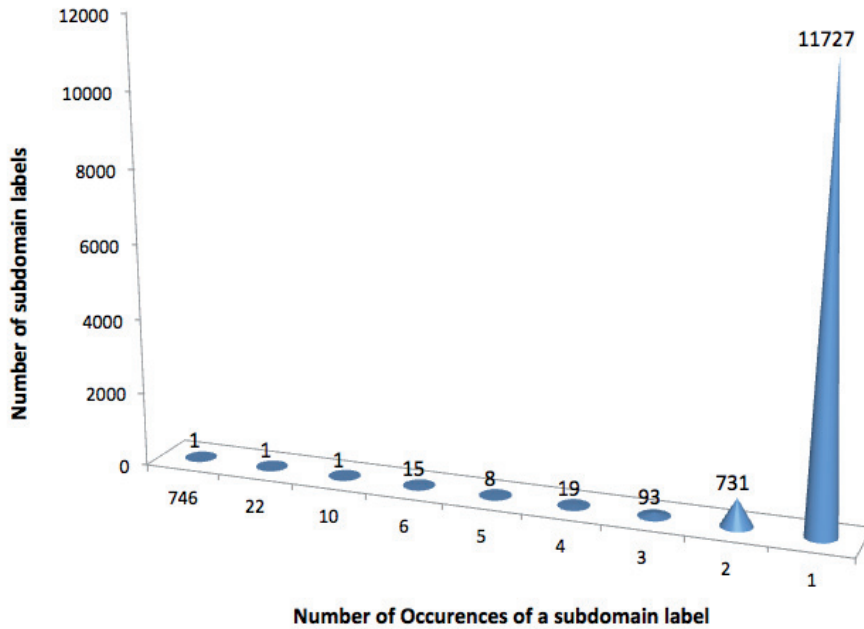


Figure 11: Frequency of occurrence of subdomain labels.

Germany, Netherlands, Finland, Czech Republic and Russia). This clearly shows that the adversaries are diversifying their hosting assets, which affords them redundancy and evasive capabilities. Notice also that RIPE has stricter data protection laws so it would be more difficult to obtain information about customers, and that could explain the shift in hosting infrastructures by the bad actors.

More generally, we list a few of the small-scale hosting providers involved in hosting the attack subdomains. These hosting providers could either be abused, complicit with the bad actors, or simply lax about the maliciousness of the content they host. Note that the rogue providers among these will often switch prefixes by dropping dirty ones and reserving new ones from the backbone providers to which they are attached.

- <http://king-servers.com/en/> has been observed to host exploit kit domains (Angler, Styx), porn, dating sites and pharma sites [24, 25]. It was also described by a comment on Web Of Trust as ‘Offers bulletproof hosting for Russian-Ukrainian criminals (malware distributors, etc.)’ [26]
- <http://evrohoster.ru/en/> hosted Browlock through redirections from porn sites [27].
- <http://www.xlhost.com/> hosted Angler Exploit Kit domains [28]
- <https://www.ubiquityhosting.com/> hosted Browlock
- <http://www.qhoster.bg/> hosted Nuclear Exploit Kit domains
- <http://www.codero.com/>
- <http://www.electrickitten.com/web-hosting/>
- <http://hostink.ru/>

String analysis of domain names

During this study, we recorded 19,000+ malicious subdomains injected under 4,200+ compromised *GoDaddy* 2LDs. By analysing the strings used for the subdomains, we recorded 12,000+ different labels. We show the list of top five labels used; ‘police’, ‘alertpolice’, ‘css’, ‘windowsmoviemaker’ and ‘solidfileszsr’. ‘Police’ and ‘alertpolice’ were the most common labels for hostnames serving Browlock. The remaining labels were used for hostnames mainly serving exploit kit attacks. In Figure 11, we show the frequency of occurrence of all the labels used.

One label occurred 746 times (‘police’), one label occurred 22 times (‘alertpolice’), one label occurred 10 times (‘css’), 15 labels occurred six times (‘windowsmoviemaker’ and ‘solidfileszsr’ among them), and 11,727 distinct labels occurred just a single time.

CONCLUSION

In this paper, we have covered methodologies for exploring malicious IP space from new angles: in addition to known techniques of assigning maliciousness scores to IPs, prefixes and ASNs based on counting volume of hosted content, we considered the topology of the AS graph, and looked at a granularity smaller than the BGP prefix. In the first case, we showed cases of rogue sibling peripheral ASNs that are delivering common suspicious payloads. In the second case, we studied sub-allocated IP ranges and shed light on the MO of bad actors to abuse these allocations from providers and avoid detection. Our system provides actionable intelligence and helps pre-emptively detect, quarantine, and monitor or block specific rogue IP space.

REFERENCES

- [1] Stone-Gross, B.; Kruegel, C.; Almeroth, K.; Moser, A.; Kirda, E. Finding rogue networks. Annual Comp. Security Applications Conference, ACSAC ‘09.
- [2] Roveta, F.; Di Mario, L.; Maggi, F.; Caviglia, G.; Zanero, S.; Ciuccarelli, P. BURN: Baring Unknown Rogue Networks. 8th Intl. Symposium on Visualization for Cyber Security, VizSec ‘11.
- [3] Yu, T.; Lippmann, R.; Riordan, J.; Boyer, S. Ember: a global perspective on extreme malicious behavior. 7th Intl. Symposium on Visualization for Cyber Security, VizSec ‘10.
- [4] Wagner, C.; Francois, J.; State, R.; Dulaunoy, A.; Engel, T.; Massen, G. ASMATRA: Ranking ASs Providing Transit Service to Malware Hosters. IEEE International Symposium on Integrated Network Management (IM 2013), 2013.
- [5] Broido, A.; Claffy, K. Analysis of RouteViews BGP data: policy atoms. Network Resource Data Management Workshop, May 2001.
- [6] <http://archive.routeviews.org/bgpdata/>.
- [7] <http://www.cidr-report.org/as2.0>.
- [8] <http://bgp.he.net/>.
- [9] [http://en.wikipedia.org/wiki/Vertex_\(graph_theory\)](http://en.wikipedia.org/wiki/Vertex_(graph_theory)).
- [10] <http://www.team-cymru.org/Services/ip-to-asn.html>.
- [11] <http://dev.maxmind.com/geoip/legacy/geolite/>.
- [12] <https://www.virustotal.com/en/ip-address/5.254.120.124/information/>.
- [13] <http://pastebin.com/X83gkPY4>.
- [14] <http://telussecuritylabs.com/threats/show/TSL20130715-08>.
- [15] http://www.f-secure.com/v-descs/trojan_html_browlock.shtml.
- [16] <http://www.ovh.com/>.
- [17] Mahjoub, D. When IPs go Nuclear. <http://labs.opendns.com/2014/02/14/when-ips-go-nuclear/>.
- [18] <http://blog.malwaremustdie.org/2014/02/tango-down-of-nuclear-packs-174.html>.
- [19] <http://nakedsecurity.sophos.com/2012/11/23/hacked-go-daddy-ransomware/>.
- [20] <http://www.malware-traffic-analysis.net/2014/05/08/index.html>.
- [21] <http://www.malware-traffic-analysis.net/2014/05/13/index.html>.
- [22] <http://www.malware-traffic-analysis.net/2014/05/25/index.html>.
- [23] <http://www.malware-traffic-analysis.net/2014/06/03/index.html>.
- [24] <http://urlquery.net/report.php?id=1397035856786>.
- [25] <https://www.virustotal.com/en/ip-address/184.105.139.31/information/>.
- [26] <https://www.mywot.com/en/scorecard/king-servers.com/comment-15984778#comment-15984778>.
- [27] <https://www.virustotal.com/en/ip-address/62.75.195.244/information/>.
- [28] <http://urlquery.net/report.php?id=1399060473120>.