

Clusters of Re-used Keys

Stephen Farrell,
Trinity College Dublin
stephen.farrell@cs.tcd.ie

Article version 0.3, Compiled on 2018/04/07 at 14:39
work-in-progress

Abstract—This article describes a survey of long-term cryptographic public keys observed in deployments of secure-shell, e-mail and web protocols in four similarly-sized countries – Ireland, Estonia, Finland and Portugal. We find that keys are very widely re-used across multiple IP addresses, and even autonomous systems. From one run scanning 18,268 hosts in Ireland that run at least one TLS or SSH service, approximately 53% of the hosts involved are using keys that are also seen on some other IP address. If two IP addresses share a key, then those two IP addresses are considered members of the same cluster. In the same scan we find a maximum cluster size of 1,991 hosts and a total of 1,437 clusters, mostly with relatively few hosts per cluster (median cluster size was 26.5, most common cluster size is two). In that scan, of the 54,447 host/port combinations running cryptographic protocols, we only see 20,053 unique keys (36%), indicating significant key re-use across hosts and ports. Scans in the other countries demonstrate the same issue. We describe the methodology followed and the published source code and public data sources that enable researchers to replicate, validate and extend these results. Clearly, such key sharing can create undesirable security and privacy dependencies between cluster members. The author contacted some local (Irish) asset-owners to try establish the reasons for this key sharing and to possibly assist with improving network posture, and is incorporating those findings on an ongoing basis.

Index Terms—Internet measurement, security, privacy, cryptographic key management

I. INTRODUCTION

This article describes scans of six populations of hosts on the Internet, two each in Ireland (IE) and Estonia (EE) and one each in Finland (FI) and Portugal (PT). The hosts in question offer some mail service. That is, the hosts in question are contactable at IPv4 addresses that listen on TCP port 25. Our scans record information relating to the long-term cryptographic keys used by a number of services on those hosts. We see unexpectedly large-scale re-use of cryptographic keys across clusters of hosts and Autonomous Systems (ASes) in these scans, Figure 1 is one of the more structured clusters seen in these scans (the graph format is described in Section V).

At the time of writing the author is in the process of contacting asset-owners for some of the hosts in the Irish (IE) population scanned in order to try establish why and how key re-use has occurred, whether or not that was deliberate, and whether the asset-owners can or would prefer to move away

from re-using long-term cryptographic keys on multiple hosts. This version is therefore a work-in-progress to assist in those discussions and will be updated as those proceed.

II. BACKGROUND

Over the last five years, the proportion of Internet traffic that is encrypted, particularly using the Transport Layer Security (TLS [1]) protocol has been increasing consistently. [2], [3]¹ The Secure-Shell (SSH [4]) protocol has also been nearly ubiquitously used for systems administration for many years. [5]

These increases in the use of encryption are consistent with the consideration that “Pervasive Monitoring is an attack” [6] [7] which reflects some of the general drivers behind recent increased deployment of cryptographic protocols. However, in addition to making use of these protocols, we also want endpoints to move beyond opportunistic security [8] and properly manage keys, especially long-term asymmetric keys, as otherwise we won’t gain all the intended benefits of increased deployment of security protocols.

For TLS and SSH, the long term key pairs in which we are interested are essentially used for host authentication, though historically they may also have been used directly for RSA key transport. Re-using the same long-term asymmetric key pair values for many different instances of services can create vulnerabilities or increase the probability or impact of some attacks. This can be relatively easily avoided if different key pairs are used for each instance, and if keys are regularly rotated.

In our scans we do see many keys re-used in clusters on different hosts. A **cluster** is defined here to be the largest set of IP addresses in our population such that, each host in the cluster shares at least one public key (for some service) with another host in the cluster. In other words, two hosts are in the same cluster, if they share a private key, regardless of the port with which that private key is used.

A. Safe Key Re-Use

Not all re-use of keys is bad. There are certainly situations where it is reasonable to re-use a key pair for different services on a single host, for example if a public key is certified for multiple related DNS names (perhaps smtp.example.com and mail.example.com) that resolve in the public DNS to the same IP address, and where the same server instance listens on

Manuscript received MMM dd 2018; Stephen Farrell is with Trinity College, Dublin 2, Ireland (email: stephen.farrell@cs.tcd.ie, <https://www.cs.tcd.ie/Stephen.Farrell/>)

¹<https://letsencrypt.org/stats/#percent-pageloads> has recent graphs.




Fig. 1: Cluster 835 from run FI-20180326 is an 80 host cluster with many key re-uses and interesting structure. Nodes represent hosts. Node colour represents AS – this cluster has hosts in only one AS. Edges represent (possibly different) ports where the same key is used by both hosts for that port. Edge colour maps to the combination of ports as per the legend (Figure 4). See Section V for more detail of graphs.

both ports. In such cases it seems fairly clearly reasonable to use one key pair to protect services on say port 587 (email submission) and port 25 (email submission and mail transfer).

Where services on the same host are less related, for example, web and mail service, it seems a little less reasonable to re-use the same key pair, however, it is certainly arguable that if an attacker manages to gain access to a host with software-only key storage, a successful attack that gains access to one key pair on the host also results in access to all the key pairs on that host.

There are also cases where re-using a key across multiple hosts can be defended. For example, if the set of hosts for which the key is re-used are all virtual machines that are, and always will be, running on the same physical hardware, and all under the control of one entity. Another reasonable case might be if the set of hosts for which the key is re-used are all “behind” a middlebox or hypervisor that is not externally visible at the IP layer and where that middlebox offers a cryptographic front-end with the full knowledge of the hosts visible on the public Internet.

Sometimes people might mirror two hosts for redundancy, including mirroring the keys used. If the private keys on both mirrors are considered to only be at risk from the same entities (e.g. both are hosted by one entity), then this scenario also seems defensible. If two different hosters were used for better redundancy, then having the same private key at risk to two entities seems less defensible.

There could be cases where a Hardware Security Module (HSM) that holds keys in secure storage is used by different hosts (say in one rack) and where the HSM for some reason cannot operate with different keys for each host. In such a case, having hosts re-use keys that never leave a HSM might be considered defensible.

There may also be cases where our scans show the same key being used, but where the “real” application will use a unique key, if for example, a TLS Server Name Indication [9] for the intended name/application is used. Given our scans

start off based on IP addresses, we would miss such cases. Nonetheless, even having the same default key pair for port 443 on multiple hosts can be problematic.

Of course, simply using different keys for a contactable IP address does not in itself imply that a service is more secure, better managed, nor that it is independent of other hosts, given current virtualisation trends. Note though that from Internet vantage points, observers cannot easily distinguish between these and less desirable cases.

B. Dangers of Key Re-Use

While the previous sub-section describes some situations where key re-use can be defended, there are also clear dangers potentially associated with key re-use. Note that any cluster member holding a relevant private key is in a position to attack other members sharing that key, when those are different entities, and our results do show cases where this is possible.

Risk of Leaks: Where many copies of private keys are stored as files on disk, which is common, a leak from one host affects all. If we notionally say the cost of a leak of any of the host keys is the same and is C and there are n hosts, then the cost of one cluster key leaking is $n * C$. If the probability of a key leaking from any host is p , then we can say that the probability of some key leaking from some host is $n * p$. If those considerations do apply then the risk of the cluster scenario could be $n^2 * p * c$ instead of $n * p * c$ if each host has a different key. In other words, the cluster scenario is riskier, and the bigger the cluster grows, the faster the risk get worse. In the argument here we’re making reasonable assumptions about independent events and equal costs and probabilities. While those assumptions will not always apply, they often will, and the external observer cannot know when they do not apply, and so would be wise to assume they do, when considering what services to use.

Masquerade: If used for authentication, then re-use of key pairs across hosts means that a breach of any host in a cluster

enables a successful attacker to also masquerade as any host from the cluster in question.

Misdirected Mail: If an attacker can manipulate mail routing (via MX resource records) or routing information (via BGP), and the attacker has a copy of relevant private keys, then the attacker can masquerade as the mail recipient’s domain and intercept email, even if the recipient’s domain enforces strong authentication of mail transport, e.g., via MTA-STS. [10]

Credentials: Being able to masquerade as any of the hosts in a cluster is likely to allow an active attacker to capture user credentials, for example passwords sent in IMAP or SMTP transactions. Even worse, if SSH password logins are still enabled, such an attacker can make use of those privileged user credentials. Given user credentials (especially passwords) are frequently re-used in many places, this kind of attack may extend beyond the cluster to other hosts at which those credentials can be used.

Web Origins: If the hosts in a cluster represent different web origins then the ability to masquerade as any host in the cluster would allow theft of web cookies, breaking the same origin policy on which the web depends. Given new web technologies such as alt-svc [11] and the ORIGIN frame [12] in HTTP version 2, breach of one web server in a cluster should allow masquerading as any other, all from the initial host.

Key Transport: If a long term RSA key pair is or has ever been used for key transport, and if an attacker has a record of those sessions, then breach of any host using that key pair allows the attacker to directly decrypt all recorded sessions.

Million Messages: Any unpatched vulnerability related to use of a private key that requires multiple messages to be sent by the attacker (e.g. the ROBOT² attack [13]), can likely be exploited more efficiently and more stealthily, if there are more service endpoints to which the attacker can send messages.

Cross Protocol: Key re-use increases the potential for cross-protocol attacks, and the likelihood that some service software supports older versions of SSH or TLS/SSL, that may be vulnerable to such attacks.

Revocation: If a key is known to be compromised, then it ought to be revoked if the public component is part of some Public Key Infrastructure (PKI[14]). With key re-use it may be that the same private key continues to be used on some hosts whilst being replaced on others. If the reason for revocation related to potential private key compromise, then those less-well managed hosts will remain at risk even after revocation of some certificates and replacement of some key pairs.

Theory: It is now feasible to do security proofs for realistic protocols, and indeed TLS1.3 has been the subject of such formal studies. It is not clear that widespread key re-use was considered in such studies. The impact, if any, on such proofs is uncertain.

Laziness: Re-use of keys can clearly be a result of careless management. Advertising attributes that can reasonably be taken to indicate carelessness to the public Internet seems like a bad plan for a service operator.

C. Research Questions

The overarching research question behind this work is to investigate whether or not local measurement of Internet security posture is more useful (compared to Internet-scale measurement) in helping asset holders to improve security posture. This study tackles one aspect of that work.

As part of this work, we define a metric related to key re-use: the percentage of hosts in a population doing some cryptography using keys that are known to be used for multiple IP addresses; versus hosts where we only see keys being used for one IP address. (Of course, the scope of our measurement means that hosts could be mis-identified as being in the latter set, if in fact some key is also used elsewhere.) Call this the “hosts are re-using keys” (HARK) percentage.

One hypothesis of this work is that reducing HARK could correlate with improvements in security, and be a reasonable indicator of whether a population are managing security more carefully than in the past. We do not necessarily aim to reduce HARK to zero, nor do we currently claim to know what value might be optimal or a good target. But of course leaving room for future work is also a fine feature for a proposed new metric.

III. METHODOLOGY

During October and November 2017, we extracted scans of Irish (IE) and Estonian (EE) hosts that listen on port 25 from the censys infrastructure. [15]³ Subsequently, starting in March 2018, we ran scans locally using ZMap/ZGrab [16] ourselves, for Ireland, Estonia, Finland and Portugal, but still limited to hosts who listen on port 25 (according to ZMap).

Our expectation was that hosts that run the SMTP protocol (mail servers that listen on port 25) could be a useful population to examine as they hopefully represent hosts with a better than average probability of being well managed. Initial examination of that data showed that keys were being re-used more frequently than expected and that there were clusters of hosts re-using private keys in various ways.

For each host we record SSH or TLS details for each of the ports listed in Table I. (Censys scans did not include port 587 in November 2017.) We developed analysis code to identify and analyse the clusters of hosts re-using keys as described below.

All source code required to replicate this study or do a similar-scale scan (of perhaps up to 30k hosts) has been published⁴ under the MIT license. Each scan of that scale requires a few days to run on our very modest infrastructure, due to limits of that infrastructure and built-in (100ms) delays between scan stages to consume less bandwidth. About 3GB of storage is required per run. Our tools take as input a country-code (e.g. IE) and should work for other countries with similar numbers of port 25 listeners.

A speedier or much larger-scale scan could likely be done with minor code changes given better bandwidth, processing power and storage. Note though that one of the research questions we want to explore is whether such “local” scans can be (more) effective in helping asset owners mitigate risks,

² <https://robotattack.org/>

³<https://censys.io>

⁴<https://github.com/sftcd/surveys/>

TABLE I: Ports Scanned

Port	Protocol	Key Fingerprint Used
22	secure-shell	SSH host key
25	SMTP	TLS Server Public
110	Pop3	TLS Server Public
143	IMAP	TLS Server Public
443	HTTPS	TLS Server Public
587	SMTP Submission	TLS Server Public
993	IMAP	TLS Server Public

so we’re sanguine about the speed and scale limitations, and happy to be able to demonstrate that quite limited resources are sufficient.

Data sets from the scans are not being published as doing so could assist an attacker with “lateral” movement to other hosts in a cluster, or attract attackers to members of clusters. However, it is important to note that, as shown by this article, it is relatively simple to detect these clusters.

A. Scanning Process

The following are the steps in the process run by our code, (we don’t go into significant detail, as consulting the code is the better option for such):

Select addresses. The first step is to establish the set of IP addresses to scan. That can be based on a previous scan or a set of prefixes from a geo-location database. Our co-ordinating script (“skey-all.sh”) either copies addresses from a previous scan or (via “IPsFromMM.py”) uses a geo-location database and ZMap to decide which ranges to scan. Running ZMap can take from a few hours to a day depending on the number of addresses in the scan, and the available bandwidth. Typically, about 1% or so of the hosts scanned will have a port 25 listener and will therefore be retained for further scanning. From our scanning hosts, having ZMap send about 150 probes per second seems to result in few or no lost answers.

Grab. The next step is to use ZGrab to attempt to connect with the host and port in question and to record details seen, including keys and other cryptographic parameters. This is done via the “FreshGrab.py” script and takes a number of hours. Before calling ZGrab we check if the IP address being considered is correctly geo-located, as in some cases addresses may be in the “wrong” country. (This could be due to changes in routing, or due to ambiguities in how MaxMind and ZMap work.)

Analyse. Having accumulated data, we analyse the records and detect the key re-use. The “SameKeys.py” script does this, and additionally compares some of the names found (e.g. in SMTP banners or X.509 certificates subject and SubjectAltName (SAN) fields) against the forward and reverse DNS.⁵

Report We then produce graphs for each cluster using graphviz “dot” format [17] (optionally rendered as SVG files)

⁵We only query up to 20 SANs per certificate - in runs we have seen some gigantic certificates with more than 1500 SANs - querying each would be too time consuming for the benefit gained - 20 SANs should be enough to identify any asset-owner, which is why we’re interested in SANs in this study.

and create a single JSON file per cluster with relevant details. This is done by the “ReportReuse.py” script.

All the steps above are run by the “skey-all.sh” shell script. Consult that script (or the “README.md” file in the code repository) for more detail. There are also some additional scripts for installing required components (“install-deps.sh”) and for validating clusters (“check-keys.sh”) as described further below.

B. Fingerprints and Clustering

We base clusters on the SHA256 [18] fingerprint of the keys, as reported by ZGrab. We also record whether e.g. TLS server certificates are “browser-trusted” or not and some other meta-data. For TLS services, the hash input is the encoded SubjectPublicKeyInfo field of the X.509 certificate presented by the server. For SSH, we use the SSH key hash, as produced by “ssh-keygen.”

We include checks for cross-protocol key re-use and do see quite a lot of that both on individual hosts and between hosts. It is not uncommon to see the same key apparently being used for port 25 on one IP address and port 443 on another. The only cross-protocol re-use that we have not yet seen is between port 22 (SSH) and other ports, though we do see many cases where SSH host keys are being re-used across multiple hosts.

Once two hosts have the same key fingerprint for any pair of ports, then we assign those hosts to the same cluster. This involves iterating over the set of records more than once, e.g., when we “join” two existing clusters having found that a host shares keys with both.

C. Scoping and Geo-location

Both censys and our local re-scans make use of MaxMind⁶ for geo-location. We use the free GeoLite2 databases (for ASNs, City and Country). The “mm_update.sh” script downloads the databases needed for our scans. When using censys, we select the set of IPv4 addresses that have listeners on port 25 and that have the relevant Country Code (IE or EE) for the sample concerned. When re-scanning locally, we start with the list of country-specific prefixes from MaxMind and later discard any specific IP addresses that no longer appear to have the correct country code. Censys’ geo-location appears to be more accurate than our local scans, which is unsurprising. However, some inaccuracy in geo-location doesn’t affect our main conclusions with respect to key re-use.

D. Other Tooling

There are some additional tools in the code repo in the “clustertools” directory – the “overlaps.sh” script compares two runs and says which clusters overlap with which, see Section VI-A for discussion of the output of this tool. The “fpsfromcluster.sh” script shows host many occurrences of each fingerprint are seen in a cluster. The “check-no-ssh-cross-protocol.sh” script checks for any cases where an SSH host-key is used for a TLS port - no such case has been seen in

⁶ <https://www.MaxMind.com/>

```

$ check-keys.sh -i cluster835.json
Running check-keys.sh at 20180402-141156
Starting at 20180402-141156, log in
  validation-results-20180402-141156.out
Doing cluster835.json
infile,ipcount,22count,matches,mismatches
835,80,66,1042,0
infile,ipcount,tlscount,matches,mismatches
835,80,140,544,41

```

Fig. 2: Output from check-keys.sh for cluster 835 (Figure 1). This indicates no discrepancies for the re-used SSH host-keys but 41 mismatches out of 585 pairs of ports involving re-used TLS keys. Most of those are due to a lack of response from some host/port combinations, others are due to four real key changes.

the five runs reported here. The “ClusterGetCerts.py” script makes a fresh connection to the TLS ports from a cluster file and extracts the X.509 certificates seen in text form. There is also a script for producing the LaTeX source (“make-tex.sh”) used in the results section below. Additional tools may be added as the analysis proceeds.

E. Validation

In order to increase our confidence that these clusters are real, we have a validation script (“check-keys.sh”) that reads a list of cluster files and uses different tooling to check if the cluster is as before. This is to guard against e.g., bugs in ZGrab or in our clustering code. For SSH, we use the “ssh-keyscan” binary to connect to the hosts in question and re-check the key hashes. For TLS services (all the others), we use the “openssl” binary (in “s_client” mode) to re-check the fingerprints.

During validation, we often see additional SSH host keys, as it appears that ZGrab (at least as we use it) finds fewer keys than ssh-keyscan. If you run the validation script from a network that e.g. doesn’t allow outbound port 25 connections (which is not uncommon), then you’ll get some false negatives, as the validation script won’t be able to connect to port 25 on the hosts in the cluster. Similarly, if a host is not accessible at all during validation, that will show up as discrepancies that may disappear in a later run. We do also see some real discrepancies for some clusters, but that is to be expected, e.g., due to key rotation for browser-trusted certificates that have expired.

We do not currently re-validate all clusters as part of runs, but do that selectively when looking at individual clusters of interest. For example, a validation run of Cluster 835 (shown in Figure 1), produced the output shown in Figure 2.

F. (Lack of) Infrastructure

As stated previously, being able to do local scanning using very modest “infrastructure” seems like a benefit. In this case, we did make initial use of the presumably well-engineered censys infrastructure, but, thanks to open-source technology and open databases, we are able to run our scans from an extremely modest virtual server or normal laptop.

TABLE II: A very modest scanner

Parameter	Value
Processor	AMD Opteron 62xx class CPU
CPU	25% of 1 Core
RAM	0.75 GB
Disk	7.5 GB
Bandwidth	Unlimited @10 Mbps

Specifically, we currently run scans from a modest Virtual Private Server (VPS) with the parameters listed in Table II. Only the “Grab” phase of the process needs to be run on this host - at the scale of scan we’re doing, the rest can be done just fine on a typical laptop.

G. Ethical Considerations

As we’re doing active scans it is appropriate to consider whether there are ethical implications of this work. Given that our current scans are of hosts that listen on port 25 (i.e., email servers) we feel that these scans have fewer ethical implications than those that might involve hosts that are operated by individual users.

Given our preference for modest scanning infrastructure, our scanning rate is low enough that we are not likely to affect any running services. We use the default ZMap block list and have published a web page and DNS TXT record that can be found from the source address from which we scan. So far, nobody has asked us to not scan them, if someone did, we would add them to the ZMap block list. As stated previously, we do not intend to publish scan data, as that could assist attackers in some cases.

IV. RELATED WORK

Early work surveying the use of cryptographic keys on the Internet included Heninger et al’s seminal work [19] identifying re-used keys and keys with common factors. Since ZMap/ZGrab and censys have become available many people have studied the properties of populations of cryptographic keys for example [2], [20], [21], [3].

The properties of email security deployments have also been studied, for example, by Durumeric et al. [22] and Holz et al. [23]. Albrecht et al, carried out Internet-wide scans of SSH usage in 2015 [5] finding “about 2^{24} ” servers in their scans.

To our knowledge, those and the many other studies of the TLS and SSH ecosystems have focused more on the protocol or cryptographic properties seen, and did not consider the clustering aspect studied here.

V. RESULTS

Table III provides the overviews of each of the five runs done for this article.

Figure 3 shows the cluster size distributions for these runs.

For each specific run below we report the following:

Country: IE, EE, FI, or PT in this article.

Scan start/end: dates on which the scan was started and ended, scans can run over multiple days. For the 2017 scans

TABLE III: Overview of runs

Country (year)	IE(2017)	IE(2018)	EE(2017)	EE(2018)	FI(2018)	PT(2018)
Scan start	2017-11-30	2018-03-16	2017-11-30	2018-03-24	2018-03-26	2018-04-03
Scan end	2018-03-27	2018-03-25	2018-03-28	2018-03-29	2018-04-01	2018-04-05
Number of IPs from ZMap	23616	24774	12775	17827	37012	19782
Judged “out of country”	0	1233	0	1334	506	63
“In country” IPs	23616	23541	12775	16493	36506	19719
No crypto seen	13277	5273	997	1519	26106	4169
Some Crypto	10339	18268	11778	14974	10400	15550
Percent with some crypto	43%	77%	92%	90%	28%	78%
Total crypto host/ports	31832	54447	47555	80019	34263	63907
Total unique keys	9234	20053	11867	20014	11686	12202
Percent keys vs. max	29%	36%	24%	25%	34%	19%
Hosts with only local keys	5406	8570	4053	3303	4675	4143
Hosts in clusters	4933	9698	7725	11671	5725	11407
HARK	47%	53%	65%	77%	55%	73%
Number of clusters	815	1437	310	639	1029	1512
Biggest cluster size	671	1991	2838	2402	373	2016
Median cluster size	21	26.5	29.5	42	24	30
Average cluster size	63.23	87.78	134.31	98.04	50.65	117.51

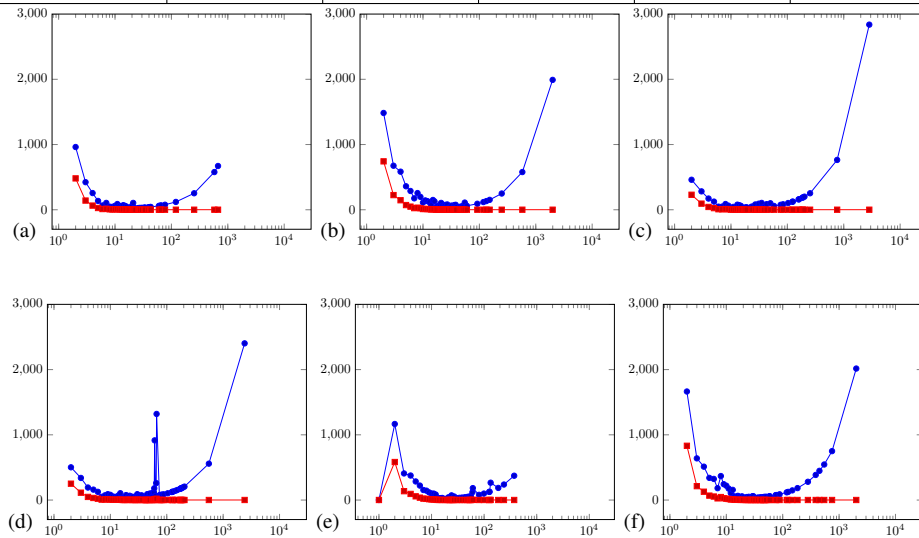


Fig. 3: Clustersize distributions for runs (a) IE-20171130, (b) IE-20180316, (c) EE-20171130, (d) EE-20180324, (e) FI-20180326, (f) PT-20180403. Blue circles show the number of hosts in clusters of given size, red squares reflect the number of clusters of given size. The x-axis is logarithmic.

here the end is not close to the start due to the time taken to develop the analysis code.

IPs from ZMap: is the number we got from censys for 2017 scans, or the (overestimated prefixes) we got from MaxMind for 2018 scans.

Judged “out of country”: is a count of specific IP addresses that MaxMind doesn’t consider to be in the right country. Note there is the usual ambiguity here with respect to Ireland/Northern Ireland/UK here. In our scans, we don’t further consider a host if MaxMind considers it to be in the “UK” though we could presumably extend to the island of Ireland via lat/long co-ords if those were available. For Estonia, there

were quite a few addresses considered to be in Sweden, but the author is not sufficiently locally knowledgeable to know if there is any rationale behind that other than some inaccuracy in the prefixes in the geo-location database.

No crypto seen: these are the in-country hosts that, when contacted, ran no SSH or TLS services that our scanner could detect.

Some crypto: these hosts have at least one port (not necessarily port 25) where a server key can be detected. There is no quality judgement as to whether the specific keys are good or bad, certified or not, certificate expired or not, etc. (We do store that but it’s not reflected in this number.)




Fig. 4: Legend for edge colours in graphs. Nodes are port numbers, edge colours are as used between host-nodes in other graphs when the pair of ports re-use the same key. The loopback colours (e.g. black from p22 to itself) are used when a key is re-used on the same port on two hosts.

Percent with some crypto: = $100 * \text{some-crypto} / \text{in-country}$

Total crypto hosts/ports: is a count of all of the ports seen on all of the hosts that run TLS or SSH

Total unique keys: is a count of all of the key fingerprints seen, across all hosts/ports - keys are only counted once, regardless of the number of hosts on which a key is seen.

Percent keys vs. max: if a different key were seen on every possible host/port combination this would be 100%

Hosts with only local keys: this is the number of hosts for whom none of their keys are seen on any other host in the run

Hosts in clusters: this is the total number of hosts that are in some cluster

HARK: the “Hosts Are Re-using Keys” percentage - which is $100 * \text{hosts-in-clusters} / \text{some-crypto}$

Biggest cluster size: the number of hosts in the biggest cluster

Median and average cluster size: as you’d expect

A. Graphing Clusters

Graphs provide a sometimes useful way to visualise clusters. These were very useful during debugging when e.g., impossible asymmetries in graphs showed up problems that needed to be addressed. Simply flicking through the graph images was a useful way to spot such anomalies.

For the medium to large clusters in these runs, the more complex graphs aren’t that useful, but do at least give an impression of scale. Some few of these graphs do show some structure that could prove useful when investigating causes, e.g. Figure 1.

For the largest graphs, the graphviz package fails to render the graph as an image. The “try-render-problematic.sh” script attempts to use graphviz in various ways and does succeed in generating images for all but the largest graphs in our runs to date. Graphviz has a number of tools for rendering graphs that perform differently, some slower, making nicer graphs (e.g. “sfdp”), but failing for more graph instances, others (e.g. “neato”) quicker and more robust, but producing less readable output in general for our graphs. All of the graphs below were produced using the “sfdp” tool, except in Figure 11 which was




Fig. 5: Cluster 112 from run EE-20180324 is a 14 host cluster with mail and web key re-uses.




Fig. 6: Cluster 60 from run FI-20180326 is a 30 host cluster involving mostly SSH host-key re-uses, but with some re-use of key on mail ports. 13 of the hosts are in one AS, and 17 in another.

produced by “neato” as sfdp times out in our build with that input as we impose a two-minute timeout for building each specific graph.

Figures 5, 6, 7 and 8 are sample cluster graphs from scans, to give a sense of the range of clusters we see.

Node numbers in the graphs are local indexes of the IP addresses in our data set. These are essentially determined by ZMap which hashes the set of input ranges we give it, resulting in node and cluster numbers changing (sufficiently) unpredictably for each run. The “ReportReuse.py” script that produces the graphviz dot files used to generate these graphs has a command line argument specifying whether to anonymise the IP addresses like this. Node colour is set based on the ASN of the host. Edge colours are specific to




Fig. 7: Cluster 144 from run PT-20180403 is a 48 host cluster with two clear “lobes”




Fig. 8: Cluster 240 from run FI-20180326 is a 25 host, 2 AS, cluster with mail and /web key re-uses.




Fig. 9: Cluster 355 from run FI-20180326 is a 21 host cluster with various mail/web key re-uses and some apparent structure.

the combination of ports on which the same key is seen - the colours used are as shown in Figure 4.

Once there are more than 10 hosts in a cluster, we no longer distinguish re-use for the same protocol via different edge colours, but simply add one edge when the same mail related ports use the same keys on a pair of hosts. So if two hosts in a larger cluster use the same key for both ports 25 and 143, then only one edge will be created on the graph for that cluster and that edge will be coloured for “mail.” Even in such larger graphs, if two hosts share a key on different ports, e.g. port25 on one host has the same key as port 143 on the second host, then, as these situations are less common, we continue to add edges for those, and the graph will have a “p25-p143” entry in the legend.

Relatively few of these graphs display interesting structure – one that does is cluster 9 from the IE-20171130 run, shown in Figure 18. The set of hosts in run IE-20171130/Cluster-9 turns into run IE-20171130/Cluster-333 shown in Figure 17. One additional host is added to the cluster and some of the linkages have undergone changes. Cluster 355 from the FI-20180326 run, shown in Figure 9, also shows some interesting structure as does cluster 835 (Figure 1) from the same run.

B. Cross-border Overlaps

The “fpoverlaps.sh” script checks for fingerprint overlaps between the clusters found in different runs. We find overlaps between the runs reported on in this article as shown in Figure 10. Note that this only checks hosts for which key re-use has already been detected within a run – it is entirely possible that there are additional re-uses not detected by this script.

Overall we see 89 cases of cross-border cluster overlap. One of those is between IE-20180316/227 (a 15-host, 10 AS cluster, Figure 19), EE-20180324/165 (a 2-host, 2 AS, cluster) and FI-20180326/688 (a 2-host cluster) and PT-20180403/1184 (a 5-host, 4 AS cluster). This may be related to some common software or hardware as one vendor’s name is mentioned in banners for all four clusters.

C. Word Clouds for Clusters

While cluster graphs help to understand the scale and possible structure of a cluster, if we aim to understand more we need to delve into the details as to which hosts, belonging to which asset-holders, are sharing keys. The cluster details stored include names from AS names, banners, reverse DNS, and X.509 certificates, but it can be tedious to extract that data from whatever bulk format (in our case JSON files) in which it is stored. To help with this, we have experimented with generating a word-list for each cluster, based on the names we find, repeated as often as they are in the underlying cluster data. From those word-lists, we build a word cloud image⁷ Doing that for all clusters and flicking through the images has been useful for identifying some asset-holders for local (Irish) clusters, but hasn’t been useful for clusters from other locales. Local knowledge therefore does seem to help in identifying some asset-holders – whether that results in much change is still to be determined. The word-list and word-cloud images are built using the “wordle.sh” script. If we get permission from some asset-holder to include an example, we will, but for now, we don’t as that would be identifying information.

VI. IRISH RESULTS

In this section we report in more detail on the Irish runs, with additional detail of selected clusters.

A. IP Address Overlaps Between Runs

We see 819 clusters from the later run (IE-20180316) that have no overlap with clusters from the earlier (IE-20171130) run. These are new clusters.

There are 612 clusters from the later run that overlap with just one cluster in the earlier run. These clusters may have expanded or contracted.

There were 6 clusters in the later run that represent a merge between two or more clusters from the earlier run, as shown in Table IV.

Looking at things the other way about, 199 clusters from the earlier run have no overlap with clusters from IE-20180316. Presumably that means those clusters were broken up into singletons, or those hosts weren’t part of the later scan, were not responsive or are no longer geo-located in Ireland. (That latter could happen due to routing/hosting changes.)

There are 607 clusters in the earlier run that have some overlap with clusters from the later run, and so presumably evolved into the cluster in the later run.

There were 9 clusters in the earlier run that appear to break up into two or more new clusters in the later run as shown in Table V.

Cluster 1107 from IE-20180316 has a key re-use overlap with IE-20171130/467 but also has a link to IE-20171130/308. A certificate in the latter cluster contains a SAN that resolves to an IP address that is a member of cluster 1107, even though there’s no key in common between these two clusters. (This link was discovered due to a bug in the script used to calculate the overlaps above.) We ignore this kind of linkage for now.

⁷https://en.wikipedia.org/wiki/Tag_cloud

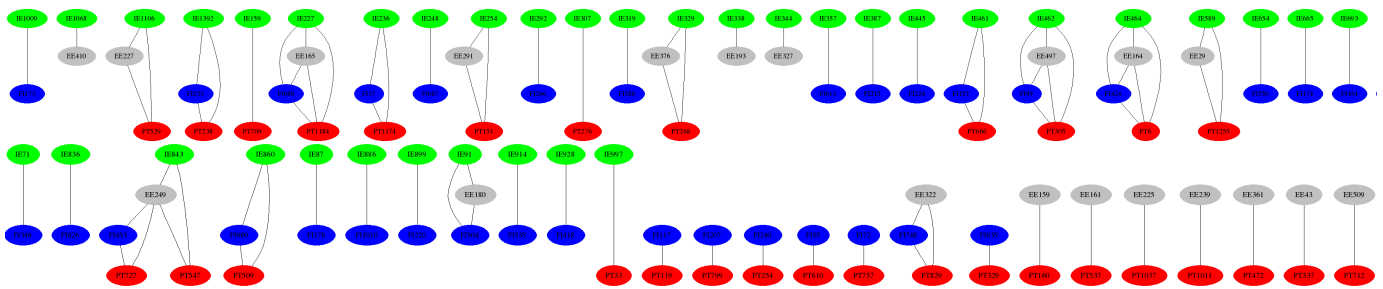


Fig. 10: Cross-border Cluster Links. Nodes in the graph represent clusters, edges mean that there is at least one key fingerprint in common between the clusters in different runs. (green=IE-20180316;blue=FI-20180326;gray=EE-20180324,red=PT-20180403)

TABLE IV: Merging Clusters - from IE-20171130 to IE-20180316

Merged into	Merged from
80	58 612
254	635 392 782
220	549 672
419	219 651
768	128 184
1003	600 714

TABLE V: Clusters Breaking Up - from IE-20171130 to IE-20180316

Original	Broken into
104	640 1380
19	100 431
253	1291 821
320	523 665
504	728 154
506	540 1434 3
515	59 6
63	780 1245
80	847 1260

We could in principle extend the concept of clustering to names and certificates as well as keys. However, as of this version, we do not extend the clustering concept, as using names would inevitably introduce ambiguity, (similar to the public suffix list, briefly described in [24]). Links based on certificates would inherit those issues with names but probably also involve too many degrees of freedom, and lead to spurious linkages.

B. Cluster-Specifics

We now describe a selected subset of the clusters seen in more detail. We do not yet attempt to present a fully consistent level of description of each cluster – we may evolve towards that, or it may be that the aspects of interest for each cluster can only be established via inspection, in which case different descriptions for each are a reasonable outcome. We hope to learn more as we discuss these clusters with asset-owners.

1) IE-20180316-181141/Cluster-3:

- This cluster has 1991 hosts and is the largest in the run.
- All the IP addresses are in the same ASN, which is one of the world’s hyper-scalers who happen to host machines in Ireland, so this cluster doesn’t seem to be particularly associated with Ireland.
- There are a total of 1995 host/port combinations. 1991 of those use the same key. The other 4 use one other key.
- 4 of the hosts have the same key for port 25. None of those are marked as browser-trusted.
- A run of “check-keys.sh” shows that the p25 keys appear to have changed.
- All 1991 hosts have the same key for port 443. All of those are marked as browser-trusted. A CT search via crt.sh for that SPKI hash seems to indicate that key has been used since late 2014.
- The hosts in the cluster do appear to be listening on port 25 - 3 randomly selected hosts from the cluster all accepted connections from telnet on port 25.
- There are no other cryptographic ports in use.
- 1972 of the hosts have the same public key certificate, which is a wild-card certificate. It may be that these hosts are used for marketing related services.
- There are 1980 different SMTP banners. For each one, it appears that the hostname in the banner is used as the hostname for a DNS name that matches the wildcard certificate. So if the SMTP banner is “foo” and the certificate is for “*.example.com” then we see a DNS entry for “foo.example.com” as one might expect.
- Many of the banner hostnames are of the form “[company]-mkt-dev1-1” or “[company]-mkt-prod2” There are 293 different company names in the set, some of which are internationally known brand names, and only a few of which are related to Ireland. An HTTP GET to port 443 on one of those hosts returned a 404. We didn’t further investigate the web site content.
- All TLS host/port combinations report use of the TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256 ciphersuite,⁸ so there may not be a current concern with RSA key transport.

2) IE-20180316-181141/Cluster-665:

- This cluster has 578 hosts and is the second largest in the run.

⁸See <https://www.iana.org/assignments/tls-parameters/tls-parameters.xhtml#tls-parameters-4>




Fig. 11: Cluster 52 has one hundred and eighteen hosts sharing Web server keys and is the largest cluster that renders with graphviz for this run. There are six larger clusters.

- The hosts in this cluster are in five ASes. The number of hosts in each of these ASes is 433, 111, 17, 15 and 2. The ASN with 433 hosts is a well-known local hoster/registrar. The ASN with the 111 hosts is a local ISP/hoster. The ASNs with 17 and 15 hosts are international with a local presence. The ASN with 2 hosts is a local Internet/computing consultancy.
- 517 of the TLS server certificates seen are browser-trusted whereas 2400 are not.
- There are a total of 3405 host/port combinations with crypto. 1298 of those use the same key. The next most commonly used key is used for 570 ports, then 375, 219, 86, and so on down to 2 - there are 65 keys used twice. There are 480 keys used only on one port.
- There are 457 reverse DNS names associated with this cluster and 20 unique names from TLS certificates - two for tool names, three as wildcards, i.e. of the form “*.example.com” and 15 for what appear to be regular hostnames one might find in DNS.
- From looking at the names from banners and certificates, it appears that some default keys for a particular hosting control panel tool are being re-used. Forum posts for that tool do appear to indicate that controlling these settings is challenging.

3) IE-20180316-181141/Cluster-9:

- This cluster has 249 hosts and is the third largest in the run.
- All 249 hosts in this cluster are in the same AS, which is a local hosting company.
- 96 of the TLS server certificates seen are browser-trusted whereas 1052 are not.
- There are a total of 1396 host/port combinations with crypto. 911 of those use the same key. The next most commonly used key is used for 90 ports, then 51, 18, 18, and so on down to 2 - there are 9 keys used twice. There are 145 keys used only on one port.

4) IE-20180316-181141/Cluster-52:

- This cluster has 118 hosts and is the seventh largest in the run. See Figure 11.
- All 118 hosts in this cluster are in the same AS, which is a hyper-scaler with a local presence.
- All 118 of the TLS server certificates seen are browser-




Fig. 12: Cluster 199 has 48 hosts




Fig. 13: Cluster 103, has twenty-five hosts sharing SSH host keys and is the largest “pure” SSH cluster in this run.

trusted and use the same certificate, issued in mid-2016.

- There are a total of 118 host/port combinations with crypto.
 - The name associated with these keys sounds like it’s related to mail delivery. But who knows?
 - The certificate for these hosts includes a name like “example.com” with a SAN for “www.example.com” but there is not A record for the relevant www.example.com in the DNS today
- ### 5) IE-20180316-181141/Cluster-199:
- Cluster 199 has 48 hosts who make lots of use of one key. See Figure 12.
 - All 48 hosts are part of a small local telco ASN. (With which the author was unfamiliar.)
 - There are 285 host/port combinations, 231 of which use the same key. 48 of those use another key, and then 5 ports use unique keys. 232 ports use browser-trusted certificates, whereas 5 do not.
 - check-key.sh produced no discrepancies.
 - All but one of the SMTP banners use the same hostname. The one exception appears to be a certificate with a SAN for what could be a .ie second-level domain (2LD) except that that value is not registered. The rest have names of the form “cpanel.[foo].com”.
- ### 6) IE-20180316-181141/Cluster-103:
- This cluster has 25 hosts and is the largest “pure-” SSH cluster in the run. See Figure 13.




Fig. 14: Cluster 144, has twenty-six hosts.




Fig. 15: Cluster 177, has twenty-six hosts.

- All 25 hosts in this cluster are in the same AS, which is a hyper-scaler with a local presence.
- There are a total of 25 host/port combinations with crypto.
- “check-keys.sh” shows some discrepancies. The discrepancies all seem to relate to hosts not responding, except for one host that does appear to now be using different host-keys.

7) *IE-20180316-181141/Cluster-144:*

- This cluster has 26 hosts and is one of two matching the median cluster size See Figure 14.
- All 26 hosts in this cluster are in the same AS, which belongs to a local ISP.
- There are a total of 125 host/port combinations with crypto. 109 of those use the same key. Each of the others has a unique key. The most-used key is used mail ports (110,143,587,993) on all 26 hosts. That same key is used for port 443 on 5 hosts. 16 other hosts use a unique key for port 443.
- “check-keys.sh” detected no mismatches - all 125 host/ports remain the same

8) *IE-20180316-181141/Cluster-177:*

- This cluster has 26 hosts and is one of two matching the median cluster size See Figure 15.
- All 26 hosts in this cluster are in the same AS, which belongs to a local ISP.
- There are a total of 128 host/port combinations with crypto. Three keys are used for all of these, one 78 times, one 26 times and one 24v times.




Fig. 16: Cluster 194 has twenty-two hosts that re-use lots of keys.



Fig. 17: Cluster 333 - Fig. 18 grown a little few months later




Fig. 18: Cluster 9 from IE-20171130-000000.

- “check-keys.sh” shows some change on port 443 since the test run.

9) *IE-20180316-181141/Cluster-194:*

- This cluster has 22 hosts. See Figure 16.
- All 22 hosts in this cluster are in the same AS, which belongs to a local ISP.
- There are a total of 121 host/port combinations with crypto. One key is used for 88 of those, one key for 21, and 12 keys are used for one port each. For 96 ports, a browser-trusted certificate is used, 4 ports use a certificate that is not browser-trusted.
- “check-keys.sh” shows a number discrepancies for SSH, but none for mail or web.
- An interesting variety of names are used for SMTP banners. Some relate to bitcoin, others to finance, and more to gaming.
- A certificate associated with the most-used key was issued in late 2017.

10) *IE-20180316-181141/Cluster-333:*

- Cluster 333 has 21 hosts. See Figures 17 and 18.
- All hosts are within a hyper-scaler’s local AS.
- There are 135 host/port combinations, with one key used 47, 26, 15, 10, 5 and 2 times each, and with 30 other keys used for one port each. 7 ports use certificates that are browser-trusted, 109 ports do not. One of the browser-trusted certificates uses a key that has been in use since 2014, across multiple certificates/CAs, another uses a key certified only once so far in early 2018, a third was issued in March 2018, and for a Dutch-sounding .eu DNS name not reflected in the SMTP banners.
- check-key.sh shows up some mail and web port discrepancies, with cases of hosts not answering and of changed




Fig. 19: Cluster 227 has 15 hosts sharing variously over ten different ASes.




Fig. 21: Cluster 462 has 14 hosts sharing variously on ports 25 and 443 web over 6 different ASes.




Fig. 20: Cluster 138 from run IE-20171130-000000 overlaps cluster 227.

fingerprints. All of the SSH ports were the same, about 20% of the mail and web ports showed discrepancies.

- SMTP banners are mixed. There are 21 different values, 19 with a common 2LD, which seems to be that of a Dutch Internet consultancy, but with two more outside that namespace. Of those, one has a .nl 2LD that matches the hostname used in some of the 19 banners, the last one seems to have no relation to the others.
- A search at crt.sh shows certificates for this public key dating back to late 2014.
- We saw this cluster as cluster 9 in our IE-20171130-000000 scan.

11) IE-20180316-181141/Cluster-227:

- Cluster 227 has 15 hosts. See Figures 19 and 20.
- 10 different ASes are involved in this cluster, of many kinds, from hyper-scalers, through local ISPs to small local hosters and even the local NREN.
- There are 20 host/port combinations, 19 of which use the same key. The odd one out is a port 443 listener. The others are for ports 25 or 443. One of those certificates is browser-trusted and is for a medium-sized food processing company in Ireland.
- SMTP banners have names that vary widely.
- check-keys.sh reports about 50% mismatches, mostly it seems due to inability to connect to some hosts.
- Since run IE-20171130-000000, this cluster has grown in size by one host, but two others have changed, i.e., since then, two hosts have dropped out, and three have joined. (Assuming our scan is perfect and hosts weren't just unavailable at the wrong moment.)
- This cluster overlaps with clusters in other runs - see Section V-B for details.




Fig. 22: Cluster 1227 has five hosts that re-use the same SSH host-keys, and two of whom share a key for port 25.

12) IE-20180316-181141/Cluster-462:

- Cluster 462 has 14 hosts. See Figure 21.
- 6 different ASes are involved in this cluster, mainly local ISPs.
- There are 16 host/port combinations, 15 of which use the same key. The most used key is a 1024 bit RSA key.
- The SMTP banners vary but generally name Irish organisations.
- check-keys.sh reports mismatches on port 443, but none on port 25.
- The certificate used on port 25 in most cases has a SAN of the form “[vendor-appliance] Demo Certificate.” It is not known if these hosts are offering a live service or are just used for testing. In at least one case the IP address of the HOST matches the MX address for the domain as per the SMTP banner seen, which implies these keys are in live use.

13) IE-20180316-181141/Cluster-1227:

- Cluster 1227 has 5 hosts. See Figure 22.
- All hosts are within a single AS operated by a hyper-scaler
- There are 13 host/port combinations, 5 are the shared SSH host-keys. Two are the shared port 25 keys. 6 other unique keys are used for port 25 and port 443 on the hosts that don't share a port 25 key. The three port 443 certificates are browser-trusted, no other certificates are browser-trusted.
- check-key.sh reports no discrepancies.




Fig. 23: Cluster 111 has five hosts that re-use the same keys for almost everything. As can be seen the individual edges become less useful at this point.




Fig. 24: Cluster 1003 has five hosts that re-use keys for almost everything. Cluster 1003 overlaps with clusters 600 and 714 from run IE-20181130-000000.



Fig. 25: Cluster 600 from 2017.




Fig. 26: Cluster 714 from 2017.

- There are four different SMTP banners, that appear unrelated. The two identical banners are on the hosts that share a key for port 25.

14) IE-20180316-181141/Cluster-111:

- Cluster 111 has 5 hosts who make lots of use of one key. See Figure 23.
- Four hosts are within a local hoster's AS. One is within a local telco/ISP's AS.
- There are 31 host/port combinations, all but one use the same key which is used on ports 110, 143, 443, 587 and 993. The odd one out is the only SSH host-key in use.
- check-key.sh says all keys remain the same.
- Each host has an SMTP banner of the form "[foo].example.com" where foo is either "server1" (one occurrence) or "hosting2" (four occurrences), whilst the (same) public key certificate for each host is for "*.example.com".
- A search at crt.sh shows certificates for this public key dating back to late 2014.

15) IE-20180316-181141/Cluster-1003:




Fig. 27: Cluster 8 - a representative four host cluster where the same TLS key is used on port 443.




Fig. 28: Cluster 76 consists of three hosts, in two different ASes and shows SSH host-key re-use.

- Cluster 1003 has 5 hosts who make lots of use of one key. See Figures 24 and Figures 25 and 26.
- All 5 hosts in this cluster are in the same AS, which is a hyper-scaler with a local presence.
- There are 35 host/port combinations. 16 use the same key for various mail ports, (but not web), 11 use the same key for the same, 5 use the same key for SSH and three unique keys are used for the web.
- browser-trusted certificates are used for various ports, the most-used key in this cluster (used for 16 ports over 6 hosts), has been certified since early 2013 (5 years, according to crt.sh).
- check-key.sh reports no discrepancies.
- SMTP banners and X.509 certificate names in this cluster refer to a .de 2LD.

16) IE-20180316-181141/Cluster-8:

- Cluster 8 has 4 hosts who share one key for port 443. See Figure 27
- All hosts are within one AS operated by a hyper-scaler.
- There are 9 host/port combinations, 4 use the same key, all for port 443. Each host has a unique key for port 22. One of the hosts has a unique key for port 25. The port 443 keys are browser-trusted, the port 25 key is not.
- check-key.sh says there are 3 discrepancies - one of the hosts seems to be unresponsive.
- Each host has an SMTP banner of the form "foo.example.com" whilst the (same) public key certificate for each host is for "*.example.com". A search at crt.sh shows only one matching certificate issued in mid 2017.

17) *IE-20180316-181141/Cluster-76:*

- Cluster 76 has 3 hosts who share SSH host-keys. See Figure 28.
- Two of the hosts are within one local AS operated by a hosting company. Another is within an AS operated by a local ISP/hoster.
- Each host only does crypto on port 22.
- check-key.sh says all 3 keys remain the same.
- There are two SMTP banners with names that don't appear to obviously match. One is a DNS bogon.

18) *IE-20180316-181141/Cluster-648:*

- This cluster has 2 hosts who share an SSH host-key and is the 3rd smallest (file) in the run.
- The SSH host-keys are 1024-bit RSA, so presumably quite old.
- Both hosts are in an ASN operated by the local NREN. Reverse DNS indicates they belong to a 3rd level educational institute.
- Neither host has any other cryptographic ports.
- The SMTP banner indicates a name that matches one of the hosts. Reverse DNS for that host also matches the banner. The other host has
- "check-keys.sh" reports both keys are still active.

19) *IE-20180316-181141/Cluster-804:*

- This cluster has 2 hosts who share an SSH host-key and is the 2nd smallest (file) in the run.
- Both hosts are in an ASN operated by the local NREN, in a range that appears to be for non-universities.
- Neither host has any other cryptographic ports.
- "check-keys.sh" reports only one of those keys are still the same at the time of writing. That could be some Firewall/Intrusion Detection System (IDS) behaviour, ssh-keyscan results for one of the hosts varies depending on whether one uses names or IP addresses, and IPv4 or IPv6. Manual inspection shows that the host keys remain the same.
- Neither IP is in our IE-20171130-000000 scan.

20) *IE-20180316-181141/Cluster-1389:*

- This cluster has 2 hosts who share an SSH host-key and is the smallest (file) in the run.
- Both hosts are in an ASN operated by a major multinational, and hoster. Reverse DNS suggests these hosts are in the ranges used for the multinational's hosted customers.
- Neither host has any other cryptographic ports.
- "check-keys.sh" reports those keys are still the same at the time of writing.
- Neither IP is in our IE-20171130-000000 scan.

21) *Summary of Clusters of Size 2:* There are 742 clusters of size 2, of those:

- 150 only involve common port 25 keys.
- 143 only involve common port 443 keys.
- 41 only involve common keys for both ports 25 and 443 (one unique key per host).
- 27 only involve common SSH host-keys.
- 24 only involve independently common keys for ports 25 and 443 (two unique keys per host).

In total, there are 170 different combinations of ports seen in this set of clusters.

VII. DISCUSSION

A. Confirmed Reasons for Key Re-Use

In this section we describe the causes for key re-use that have been confirmed with asset-owners.

SSH Host Key Reuse: Some SSH host-key clusters were caused by starting the "sshd" daemon (thus generating the SSH host key pair) on a host prior to creating a virtual machine image as a clone of that host, so that subsequent loads of that image onto another of host caused the re-use. This is not obvious to clients, as the host-key hash presented is perhaps unlikely to be recorded. This was confirmed for cluster IE-20180316/Cluster 35. The asset-owner in that case now has new scripting so that cluster shouldn't grow further and may disappear.

B. Possible Reasons for Key Re-Use

As of now, this section is speculative. We hope that changes as we engage further with asset owners, so this section is expected to change significantly.

Clusters could have been created deliberately or accidentally. The former case could be as a result of an attack or, more likely, to make someone's life easier or due to limitations in tooling that don't make key re-use apparent to administrators. Asset owners may or may not know or care about the existence of the clusters.

Mirrored Hosts: One asset holder contacted the author describing a mirroring arrangement as outlined in Section II-A. Presumably this kind of mirroring is the cause for a number of the clusters of size two and perhaps even three.

Copied configs/VMs: We expect that some of the key re-use detected here is due to config files being copied between hosts, either as part of some orchestration process, or within virtual machine images.

Old/New Hosts: It could be that some of the size-two clusters represent hosts where one is a "new" version of the other, but where the administrators forgot to turn off or re-configure the old machine after setting up the new.

Bad Random Number Generators: It is possible that some of the key re-uses detected here are the result of bad random number generators causing the same keys to be generated at different times and places.

Software/Hardware with Hard-coded Keys: If some software package or piece of equipment ships with hard-coded or default keys, then that would clearly show up in our results.

Test Equipment left Running: If test equipment is installed with default or test keys, and never updated to use real keys then key re-use is highly likely. Some of the SMTP banners we see in results seem to imply that this may have occurred.

C. De-Clustering

If one considers these clusters undesirable, then the question arises as to how one might migrate away from large clusters.

In this section we suggest ways in which administrators might move away from key re-use.

Measurement: If the adage “you can’t manage what you don’t measure” is considered to have some validity, then presumably a first step in de-clustering could be to monitor for the existence of key re-use. Systems administration and monitoring tools could relatively easily integrate the kind of key re-use detection described here, and thereby enable administrators to decide to take action (or not) when they see unexpected re-use.

Regular Key Rotation: The most basic way to avoid being a part of one of these clusters is likely to be to simply rotate all cryptographic keys at some frequency that is acceptable for the application context. This can of course cause problems if e.g. hashes of keys have been stored in applications. However, not rotating keys brings with it the significant risk that e.g. former employees may continue to have access to systems in contravention of local policy, so the pain is likely worthwhile.

Use a PKI that encourages key rotation: As previously pointed out, the combination of certbot and LetsEncrypt results in keys being changed every few months, so setting up cron jobs on each host to renew certificates in that manner with LetsEncrypt or some other Certification Authority (CA) will remove re-uses fairly quickly. Note that this can as easily be done for certificates used for mail as for the web and this could go a long way to de-clustering generally. There may be a perceived downside to doing this as the certificates will likely end up in Certificate Transparency logs, however, the fact that scans such as ours can in any case see those keys and detect key re-use seems to imply that that is not a very convincing argument.

SSH Client Notification: Whether or not having SSH clients warn about the re-use of host keys would be an effective improvement is something that could be tested with systems administrators. In principle, SSH clients could warn a user that the same host key has been seen for multiple entries in the “known_hosts” file.

VIII. FUTURE WORK

Talk to Asset Owners: The author has begun discussing these results with local (Irish) asset holders and will be updating this as events warrant. But there are clearly plenty of clusters to go around, and also different locales so replication of this work would be interesting, both to validate (or falsify!) these results, but also (in the former case) to investigate whether some common local approaches to de-clustering emerge in different locales.

Efficiency: Whilst it isn’t a problem that these scans take days, as the underlying data is unlikely to change quickly, it would nonetheless be better to improve the efficiency of the tools so they could be run on even more modest machines. At present they are fairly memory intensive and slow in parts - use of a database would likely be a significant improvement.

Infrastructure: Building a relatively modest server, with modest storage (a few terabytes) and excellent bandwidth would also speed up the scans. The author plans to investigate such infrastructure locally in Ireland and would be happy to

help anyone else who wanted to do that instead, or as well, in some other locale.

Other Populations: The populations scanned to date are geographically bounded and all run SMTP listeners. It would also be interesting to scan sets of related hosts, e.g. belonging to the same sector or making use of the same technologies. The scanning tools used here could also be used within enterprise networks to check internal and externally visible hosts.

Check the rest of the Internet: While the local clusters are interesting, it seems like an obvious extension to check if keys being re-used locally are also used elsewhere. Starting from locally detected clusters may be a useful way to approach that at Internet-scale.

Check in CT: We have not, but could, check for additional information based on searching Certificate Transparency (CT) logs [25], for example the crt.sh⁹ web interface offered by Commodo does allow searching based on the TLS fingerprints we find. We have not (yet) done the work to develop an application to use a CT API for the populations we have scanned, but have verified that at least some of our fingerprints do occur in CT logs.

Slowly doing more runs: Another obvious thing to do is to run this for other locales, and to extend to other geographic scopes, e.g. city or regional scale.

IPv6: Investigating the IPv6 addresses associated with the names detected here could also be of interest, if it extends any of the clusters. (Generic IPv6 scanning is of course of interest but not as a specific extension of this work.)

Mitigations and Incentives: Whilst discussing with asset holders, it will be interesting to investigate better mitigations and the incentives that might motivate administrators to do better in this space. For the former, one could speculate that administrative tools may be making accidental key re-use too likely and run-time/monitoring tools are presumably not looking out for key re-uses.

Key Rotation: Even if a host starts out as part of a cluster, it ought to be a normal part of applications using cryptography to periodically rotate to new key pairs. That should result in clusters being broken up relatively quickly. TLS server certificates using the LetsEncrypt (tm) CA and certbot do seem to default to changing keys during certificate renewal which is also relatively frequent. If CAs had policies that called for key rotation, or notified key holders that key re-use has been detected, that could be a significant help in breaking up clusters.

Longitudinal Studies Running cron jobs with these scans over time may produce interesting results in terms of how clusters form, live and (hopefully) dissipate.

Better Metrics: While the HARK metric is simple and easy to understand, it does not capture the changes in risk possibly associated with cluster sizes and density. It could be interesting to investigate whether metrics used for other clusters of risk (e.g. health related metrics) could provide more accurate metrics. And of course, determining whether or not any metric related to these clusters is useful in practice could also be interesting.

⁹<https://crt.sh/>

IX. CONCLUSION

The HARK numbers were a surprise to the author. One conclusion is that doing measurement is a good, perhaps especially when researchers first dip a toe in these waters with a background that others haven't previously brought to the space. The clusters seen here do seem to indicate some failings in key management, possibly due to a mixture of technology limitations and operators still being less familiar with managing keys at scale.

In the end though - rotate the keys!

ACKNOWLEDGEMENTS

Initial data on which this survey was built was made available at no cost to the author as a researcher by Censys.io. Thanks for that.

Thanks to all those who attended the September 2017, "responsible" workshop¹⁰, the discussion at which provided the inspiration for this work.

REFERENCES

- [1] T. Dierks and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2," RFC 5246 (Proposed Standard), RFC Editor, Fremont, CA, USA, pp. 1–104, Aug. 2008, updated by RFCs 5746, 5878, 6176, 7465, 7507, 7568, 7627, 7685, 7905, 7919. [Online]. Available: <https://www.rfc-editor.org/rfc/rfc5246.txt>
- [2] A. P. Felt, R. Barnes, A. King, C. Palmer, C. Bentzel, and P. Tabriz, "Measuring https adoption on the web," in *26th USENIX Security Symposium*, 2017, pp. 1323–1338.
- [3] J. Amann, O. Gasser, Q. Scheitle, L. Brent, G. Carle, and R. Holz, "Mission accomplished?: Htpps security after digninotar," in *Proceedings of the 2017 Internet Measurement Conference*. ACM, 2017, pp. 325–340.
- [4] T. Ylonen and C. Lonvick (Ed.), "The Secure Shell (SSH) Transport Layer Protocol," RFC 4253 (Proposed Standard), RFC Editor, Fremont, CA, USA, pp. 1–32, Jan. 2006, updated by RFCs 6668, 8268, 8308, 8332. [Online]. Available: <https://www.rfc-editor.org/rfc/rfc4253.txt>
- [5] M. R. Albrecht, J. P. Degabriele, T. B. Hansen, and K. G. Paterson, "A surfeit of ssh cipher suites," in *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 2016, pp. 1480–1491.
- [6] S. Farrell and H. Tschofenig, "Pervasive Monitoring Is an Attack," RFC 7258 (Best Current Practice), RFC Editor, Fremont, CA, USA, pp. 1–6, May 2014. [Online]. Available: <https://www.rfc-editor.org/rfc/rfc7258.txt>
- [7] S. Farrell, "Why pervasive monitoring is bad," *IEEE Internet Computing*, vol. 18, no. 4, pp. 4–7, 2014.
- [8] V. Dukhovni, "Opportunistic Security: Some Protection Most of the Time," RFC 7435 (Informational), RFC Editor, Fremont, CA, USA, pp. 1–11, Dec. 2014. [Online]. Available: <https://www.rfc-editor.org/rfc/rfc7435.txt>
- [9] D. Eastlake 3rd, "Transport Layer Security (TLS) Extensions: Extension Definitions," RFC 6066 (Proposed Standard), RFC Editor, Fremont, CA, USA, pp. 1–25, Jan. 2011. [Online]. Available: <https://www.rfc-editor.org/rfc/rfc6066.txt>
- [10] D. Margolis, M. Risher, B. Ramakrishnan, A. Brotman, and J. Jones, "SMTP MTA Strict Transport Security (MTA-STX)," Internet Engineering Task Force, Internet-Draft draft-ietf-uta-mta-sts-15, Apr. 2018, work in progress. [Online]. Available: <http://tools.ietf.org/html/draft-ietf-uta-mta-sts-15>
- [11] M. Nottingham, P. McManus, and J. Reschke, "HTTP Alternative Services," RFC 7838 (Proposed Standard), RFC Editor, Fremont, CA, USA, pp. 1–20, Apr. 2016. [Online]. Available: <https://www.rfc-editor.org/rfc/rfc7838.txt>
- [12] M. Nottingham and E. Nygren, "The ORIGIN HTTP/2 Frame," RFC 8336 (Proposed Standard), RFC Editor, Fremont, CA, USA, pp. 1–11, Mar. 2018. [Online]. Available: <https://www.rfc-editor.org/rfc/rfc8336.txt>
- [13] H. Bck, J. Somorovsky, and C. Young, "Return of bleichenbacher's oracle threat (robot)," Cryptology ePrint Archive, Report 2017/1189, 2017, <https://eprint.iacr.org/2017/1189>.
- [14] D. Cooper, S. Santesson, S. Farrell, S. Boeyen, R. Housley, and W. Polk, "Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile," RFC 5280 (Proposed Standard), RFC Editor, Fremont, CA, USA, pp. 1–151, May 2008, updated by RFC 6818. [Online]. Available: <https://www.rfc-editor.org/rfc/rfc5280.txt>
- [15] Z. Durumeric, D. Adrian, A. Mirian, M. Bailey, and J. A. Halderman, "A search engine backed by internet-wide scanning," in *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*. ACM, 2015, pp. 542–553.
- [16] Z. Durumeric, E. Wustrow, and J. A. Halderman, "Zmap: Fast internet-wide scanning and its security applications," in *Usenix Security*, 2013, pp. 605–620.
- [17] E. Gansner, E. Koutsofios, and S. North, "Drawing graphs with dot," 2006.
- [18] D. Eastlake 3rd and T. Hansen, "US Secure Hash Algorithms (SHA and SHA-based HMAC and HKDF)," RFC 6234 (Informational), RFC Editor, Fremont, CA, USA, pp. 1–127, May 2011. [Online]. Available: <https://www.rfc-editor.org/rfc/rfc6234.txt>
- [19] N. Heninger, Z. Durumeric, E. Wustrow, and J. A. Halderman, "Mining your ps and qs: Detection of widespread weak keys in network devices," in *USENIX Security Symposium*, vol. 8, 2012, p. 1.
- [20] L. Valenta, D. Adrian, A. Sanso, S. Cohnney, J. Fried, M. Hastings, J. A. Halderman, and N. Heninger, "Measuring small subgroup attacks against diffie-hellman," *IACR Cryptology ePrint Archive*, vol. 2016, p. 995, 2016.
- [21] M. Hastings, J. Fried, and N. Heninger, "Weak keys remain widespread in network devices," in *Proceedings of the 2016 Internet Measurement Conference*. ACM, 2016, pp. 49–63.
- [22] Z. Durumeric, D. Adrian, A. Mirian, J. Kasten, E. Bursztein, N. Lidzorski, K. Thomas, V. Eranti, M. Bailey, and J. A. Halderman, "Neither snow nor rain nor mitm...: An empirical analysis of email delivery security," in *Proceedings of the 2015 Internet Measurement Conference*. ACM, 2015, pp. 27–39.
- [23] R. Holz, J. Amann, O. Mehani, M. Wachs, and M. A. Kaafar, "Tls in the wild: An internet-wide analysis of tls-based protocols for electronic communication," *arXiv preprint arXiv:1511.00341*, 2015.
- [24] A. Barth, "HTTP State Management Mechanism," RFC 6265 (Proposed Standard), RFC Editor, Fremont, CA, USA, pp. 1–37, Apr. 2011. [Online]. Available: <https://www.rfc-editor.org/rfc/rfc6265.txt>
- [25] B. Laurie, A. Langley, and E. Kasper, "Certificate Transparency," RFC 6962 (Experimental), RFC Editor, Fremont, CA, USA, pp. 1–27, Jun. 2013. [Online]. Available: <https://www.rfc-editor.org/rfc/rfc6962.txt>



Stephen Farrell is a research fellow at Trinity College Dublin, from which he received his PhD in 2008. His research interests include security and privacy and communication in unusual and stressed environments. Between 2011 and 2017, Stephen was Internet Engineering Task Force (IETF) security area director. Stephen is a senior technical advisor for M3AAWG and a co-founder of Tolerant Networks Limited.

¹⁰ <https://responsible.ie/>