CyberGreen

Internet Infrastructure Health Metrics Framework

Phase II: Methodology and Scoring

October 31, 2022

Introduction

This document is intended to help policymakers and scientists understand the scorecard work that CyberGreen has launched. It is part of an ongoing project to provide easily digestible information about the state of health of internet infrastructure.

Background on Cyber Public Health

CyberGreen is on a mission to establish a science of Cyber Public Health, dedicated to making the internet safer and more resilient for all. You can read more about the mission at cybergreen.net

This document is intended to be a self-contained description of how we create and update the scores in the scorecard, including the tradeoffs that we make to produce the scores. There is no single number that fully captures a person's health; the same is true of the numbers we present here. Many of them are aspects of Cyber Public Health, and together they may paint a picture. That picture is enhanced by time series information (the indicators going up or down with time) and by comparison between countries, and Internet entities within each country.

Many of the measurements we take are measures across, say, a nation's IP space. To the extent that the numbers differ by country, it is likely that they are impacted by policy choices. To the extent that they are changing at different rates between countries may reflect either extant policies having an effect, or new policies coming into effect. That is why we strive to make this understandable to policymakers.

Overview of the Scorecard Methodology

1. CyberGreen is actively collecting data on a variety of things which we (or others) can scan for at internet scale because those factors are visible to the public. We select these observable data points with the belief that they indicate something about the security state of the systems they represent. For example, Routing security refers to the technical security measures to ensure that internet traffic ends up where the sender means for it to go.
2. For each data *point*, we tie it to a country through a process called geolocation, sometimes relying on a commercial data source such as Maxmind. For each type of data, we define and measure a digital population for each country, so we can assess things not only country by country, but also by the prevalence of an indicator within a given country.
    a. These population measures include IP addresses, but also things like ASNs (autonomous systems) and domains (like CyberGreen.net)

   b. The count of each may differ substantially - a country might have a lot of IP addresses grouped into a few “Autonomous Systems” (AS), a few IP addresses in a few ASes. This is similar to how In the physical world, we might measure things by city or by state.
3. With the data points and populations, we look to assess the prevalence of each indicator (the fraction of the population with that status.) The differing ways to measure and group may lead to different orderings. This is similar to how a ranking that sorts by cities might give a different order than one that ranks by states. Rather than being a problem with the methodology, these lead to nuance in the results.
4. We produce one or more ranked lists per problem, including by count of problems, and by prevalence and incidence (the new cases in a given time period).
5. Over time, we measure churn in each list.
6. We normalize (some of) those lists to produce ranks from 1 to 100 to make comparisons between the ranked lists simpler. (We continue to explore methods for combining the raw scores and churn into rankings.)

Ranking for Routing Security

Routing Security: Why It Matters

Routing refers to the way the internet directs (routes) information from one computer to another. When routing is not secure, either or both the path that packets take or its endpoints can be manipulated. An attacker might alter the path to increase their ability to view or change the packets themselves. Thus the data are more vulnerable to their contents being disclosed or altered. And if the endpoints are changed, either resulting in a “monkey in the middle” (MITM) attack or a complete fake endpoint, then the user may submit information to the wrong place, not receive expected services, etc.

How Routing Security Works

According to APNIC:

> Resource Public Key Infrastructure (RPKI) is a public key infrastructure framework designed to secure the Internet's routing infrastructure, specifically the Border Gateway Protocol. RPKI provides a way to connect Internet number resource information (such as IP Addresses) to a trust anchor. Using RPKI, legitimate holders of number resources are

able to control the operation of Internet routing protocols to prevent route hijacking and other attacks.

Cloudflare also has a good writeup, RPKI - The required cryptographic upgrade to BGP routing, which includes “The simplest introduction to BGP possible.” This also explains ROA records, which we also refer to later in this section.

Cyber Public Health and Routing Security

No single entity can fix routing security. Routing security is an emergent property of many different ISPs choosing to improve how they tell other ISPs about the routes that they make available, and how they choose to listen to (or process) route information from others. This property is a bit like food safety. Even if many producers, individually, are producing food safely, there may be problems introduced either along the path from farm to grocery to restaurant or home, or during preparation or storage.

This coordination has largely been driven via internet governance processes, rather than by public policy, and perhaps the outcomes are appropriate. With this analysis, we provide policymakers more information which they can use to judge.

Another important property of routing security is that it’s externally measurable, and it is being measured by the Routinator project. At this stage, Cyber Public Health is best advanced by leveraging such data.

How We Measure and Rank Routing Security[1]:

1. We collect data using the Routinator APIs at RIPE/NCC. Routinator is a front end to the RPKI Repository Delta Protocol, and draws authoritative data from the five Regional Internet Registries. That data includes a list of which ASNs are implementing RPKI.
2. For each ASN, we check the country of registration, and add it to a list of ASNs by country.
3. For IP addresses, we geolocate it to a country level of granularity, and add them to lists of IP addresses.
4. We separate the ASNs into a set of lists, one list per country. For each country, we check each ASN in that country for a valid ROA record. (The ROA record securely ties assigned IP addresses to an ASN.)
5. We then calculate the percentage of ASNs in the country which have ROA records.
6. We can then sort the countries by the fraction of their ASNs that have ROA.

[1] For more detail, see Appendix 1

Time-Based Rankings

We calculate a small group of ranked changes over time, including:

1. For countries, the absolute improvement (going from a .5 to a .75, representing 25 points of improvement is better than going from .20 to .40, representing 20 points of improvement).
2. For countries, the relative improvement (going from a .2 to a .4 represent a 100% improvement, while going from a .5 to a .75 is a 50% improvement)
3. Within each country, the relative change within ASNs that contribute most to the ranking changes, positive or negative.
4. Within each country, the absolute improvement within each ASN

Time-Based Rankings, Explained

Each day, for each country, we compute the fraction of the ASNs that have a valid ROA record (relative to the total number of ASNs). Day to day, we compute the change in that fraction. That change can have several meanings.

To be precise and to show how the changes may happen, a little math. With n_t representing the total number of ASNs in a given country, and n_r representing those with ROA records, we're computing:

$$m = \frac{n_t - n_r}{n_t}$$

So, let's say on day 1, n_t is 100, and n_r is 50. That is, half the ASNs have ROA records. We compute (100-50)/100 = .5. If the next day, the number of ASNs assigned to the country doubles without any change to the ROAs, then we compute (200-50)/200 = .75. (Lower is better in this representation.) If the ROAs go up the next day to 90, then we get 200-90/200 = .55.

We compute the churn for a given day (c_t) as the day to day changes to m. A simplified view is here, and a more complete one is in Appendix A.

$$c_t = \frac{m_{t-1} - m_t}{m_{t-1}}$$

A positive churn (c_t) indicates improvement to routing security, while a negative churn indicates deterioration.

We can then average the day to day churn to discover the average churn across a time period.

Additionally, the combination of the churn over a time period and the rankings on a given day provide a method for ranking countries and ASNs by their security adoption and stability.

Ranking for Open Services

Open Services Security: Why It Matters

DDoS (Distributed Denial of Service) is a category of attacks which use many internet connected computers to connect to and overwhelm a target. Many DDoS attacks use “amplifiers,” which are systems that return a great deal of data compared to what’s transmitted to them. (The amplification can range from 2x up to 500x, and in the case of “memcache,” up to 51,000x). The attackers combine these amplified services with spoofed source addresses to have the attack impact a third party.

How Open Services Security Works

CyberGreen conducts five scans per week of IPv4 space, each of which focuses on the systematic probing of five different services on publicly accessible hosts:

- Domain Name System (DNS): The Internet’s equivalent of a phone book. One important function is that it maps human readable domain names to computer readable IP addresses;
- Network Time Protocol (NTP): Used for clock synchronization between varying computer systems and is widely used to disseminate accurate time to computers and network devices;
- Simple Network Management Protocol (SNMP): Used for exchanging management information between network devices and is widely used to monitor the health and welfare of these devices;
- Simple Service Discovery Protocol (SSDP): Used to determine what services are available on the network;
- Character Generator Protocol (CHARGEN): Used for testing and measurement purposes.

All of the services measured often run open or unmanaged which are the starting point for many successful DDoS attacks. Obtaining measurement data on the number of open services gives valuable information to ascertain where varying threats are more realizable and where more effective mitigation techniques may need to be deployed

Each of these services often have unauthenticated means of being utilized and can be abused to initiate amplification attacks. Amplification attacks are a type of DDoS where an initial small query turns into a much larger payload, targeted at a specific victim.

Cyber Public Health and Open Services Security

We use the term “Open Services” to refer to network services that can be used for DDoS amplification. There are attacks where a large quantity of traffic is created which causes disruption of service or renders a service unavailable. Those exploitable resources could be used as attack infrastructure to harm others, not just the asset owners.

DDoS is a serious issue which can disrupt critical Internet enabled services that citizens are dependent upon for their daily life and well-being. At the same time, the costs imposed by DDoS often fall on entities other than those managing the assets (open systems) that enable DDoS attacks. Active steps by governments are useful to overcome this market failure. The ultimate goal is to provide national stakeholders with the information they need to mitigate the vulnerabilities in their own ecosystems which pose a risk not only to their own country but to others as well.

How We Measure and Rank Open Services Security

For each ASN we count the number of distinct IPs responding for each open service, as seen by the Open Services scanners. We multiply the count by the amplification factor published by CISA, and then divide that new number by the number of advertised IPs for that ASN (data comes from RIPE at this API). Additionally, the final score number is multiplied by 100 for clarity.

$$ASN\ Score\ =\ \frac{\#\ Associated\ IPs\ for\ each\ ASN\ *\ Amplification\ Factor}{\#\ Advertised\ IPs\ by\ the\ ASN}\ *\ 100.0$$

For each country we count the number of IPs associated with it by adding the IP counts of their respective ASN (calculated above), multiply that total by the amplification factor listed above, and then divide that new number by the number of advertised IPs for that country (again, the data comes from RIPE). Additionally, the final score number is multiplied by 100 for clarity.

$$Associated\ IPs\ for\ Country\ =\ \Sigma\ \#Associated\ IPs\ for\ each\ ASN\ registered\ in\ the\ country$$

$$Advertised\ IPs\ for\ Country\ =\ \Sigma\ \#Advertised\ IPs\ for\ each\ ASN\ registered\ in\ the\ country$$

$$Country\ Score = \frac{Associated\ IPs\ for\ Country * Amplification\ Factor}{Advertised\ IPs\ for\ Country} * 100.0$$

This process is repeated across all tracked services (DNS, NTP, SNMP, SSDP, CHARGEN). We then use the same time based algorithms as for routing for each open service to see how much churn occurs for each open service. Each service's churn is calculated from one data collection to the next, and these data collection periods vary based on the effort to collect the data. (Open services data is collected over a week, while others can be collected as frequently as daily.)

A positive churn indicates improvement to Open Services security, while a negative churn indicates deterioration.

Ranking for DNS

DNS Security: Why It Matters

The domain name system is a globally distributed, loosely coherent dynamic database of information. It maps names to IP addresses and is also used for other types of information dissemination. It is a fundamental service that must be reliable, available and trusted.

How DNS Security Works

As Akamai explains:

> Domain Name System Security Extensions (DNSSEC) are cryptographic signatures that get added to DNS records to secure data transmitted over Internet Protocol (IP) networks. DNSSEC exists because the founding architects of DNS did not include any protocol security measures. This enabled attackers to discover opportunities to forge records and direct users to fraudulent websites. Therefore, the DNSSEC protocol was introduced to add a layer of authenticity and integrity to DNS responses.
>
> DNSSEC works by adding cryptographic signatures to existing DNS records to establish a secure DNS. The signatures get stored in DNS name servers alongside common record types, such as AAAA and MX. Then, by checking the signature that corresponds to a requested DNS record, you can verify that the record stems directly from its authoritative name server. This means that the record was never poisoned or otherwise tampered with during its digital transit — thereby preventing the introduction of fake records.

Cyber Public Health and DNS security

There are several key threats to the trustworthiness of DNS systems. The integrity of responses can be affected by forging DNS responses and causing traffic to be misrouted to malicious servers. DNS integrity can also be compromised using cache poisoning attacks where legitimate DNS queries receive falsified responses. This is also sometimes referred to as DNS spoofing since the DNS responses are "spoofed" or altered to redirect traffic to an attacker's chosen destination. Each of these can be addressed via DNSSec.

Lame delegations are issues which occur when a nameserver responsible for a specific domain is unable to authoritatively provide information about it which exposes the domains to performance, reliability, and security risks. Lame delegations are an operational issue, and their presence is an indicator that DNS records are being less carefully maintained. This likely correlates with other security issues, possibly more broadly than DNS.

We do not include DNS as a DDoS amplifier to keep measures independent.

How We Measure and Rank DNS Security

Ranking Country Domains by DNS Security

1. Worldwide consider the total number of domains being used in the measurement (N)
2. For each country who's zone files / authoritative country dataset we're using, identify the number of DNSSEC enabled domains (S) in the zone file. This results into a map as follows:
    a. <Country, [#Domains (N_d), #DNSSEC Enabled Domains (N_s)]>
    b. Compute the % of DNSSEC enabled domains $x_c = N_s/N_d$.
3. Additionally, we also obtain the number of lame delegations[2] per country for the domains (L):
    a. <Country, [#Domains (N_d), #Lame Delegations (N_l)]>
    b. Compute the % of Lame delegations $x_l = N_l/N_d$.
4. Identify the number of domains which are DNSSEC enabled but have a lame delegation, assign a penalty score of 1, while for non DNSSEC enabled zones with lame delegations assign a penalty score of 0.5.
    a. Compute $P=(N_{(S \cap L)} + 0.5\ N_{(L-S)})$
5. Normalize the number of domains per country by computing the exponent E=sum(*Xc-Xl+P*)for all the countries and dividing each *Xc-Xl+P* with E.
6. Rank the results accordingly in absolute terms.

[2] For more detail, see Appendix 2

We then use the same time based algorithms as for routing for DNS to see how much churn occurs for each domain. DNS churn is calculated from one data collection to the next, and these data collection periods vary based on the effort to collect the data. (DNS data is collected over a week, while others can be collected as frequently as daily.)

A positive churn indicates improvement to DNS security, while a negative churn indicates deterioration.

Evaluation Criteria: Understanding the Data and Scores

On looking at all the numbers, it's natural to ask "is that good?" This section enumerates current and initial ways for someone looking at a numeric score to contextualize that score. We expect this will change both as we gather more data, and as we refine our scoring algorithms.

This document provides graphs to give the reader a sense of the first snapshots, and also provides, median, mean and other characterizations of the data. The evaluation for particular countries is currently fully manual.

As examples, we'll look at China, Korea and Japan as consistently selected examples, and also to specific examples chosen to illustrate specific evaluations. The data shown in this report is generally from the file 2022-10-02-through-2022-10-27-created-2022-10-28-18-22.csv.

Routing scores

Routing scores are risk scores (higher is more risk, less security) currently 0-100. While the scores visually seem heavily weighted towards the higher risk (higher part of the graph), the mean score is 63, with a median of 61. 10 countries score at 0, while 31 score at 100.

In many of these charts, ISO country codes are sorted alphabetically, and represented by a number. Thus a 1 is .AD, a 2 is .AE, a 3 is .AF, and 250 is .ZW. (The country to number mapping may change if a country code TLD starts or stops routing.)

Figure 1: Routing Scores

Figure 1 shows risk scores on the vertical Y axis, and country code, represented by its number, on the X axis. The chart is intended to show the distribution of data, not to allow a specific data point to be selected.

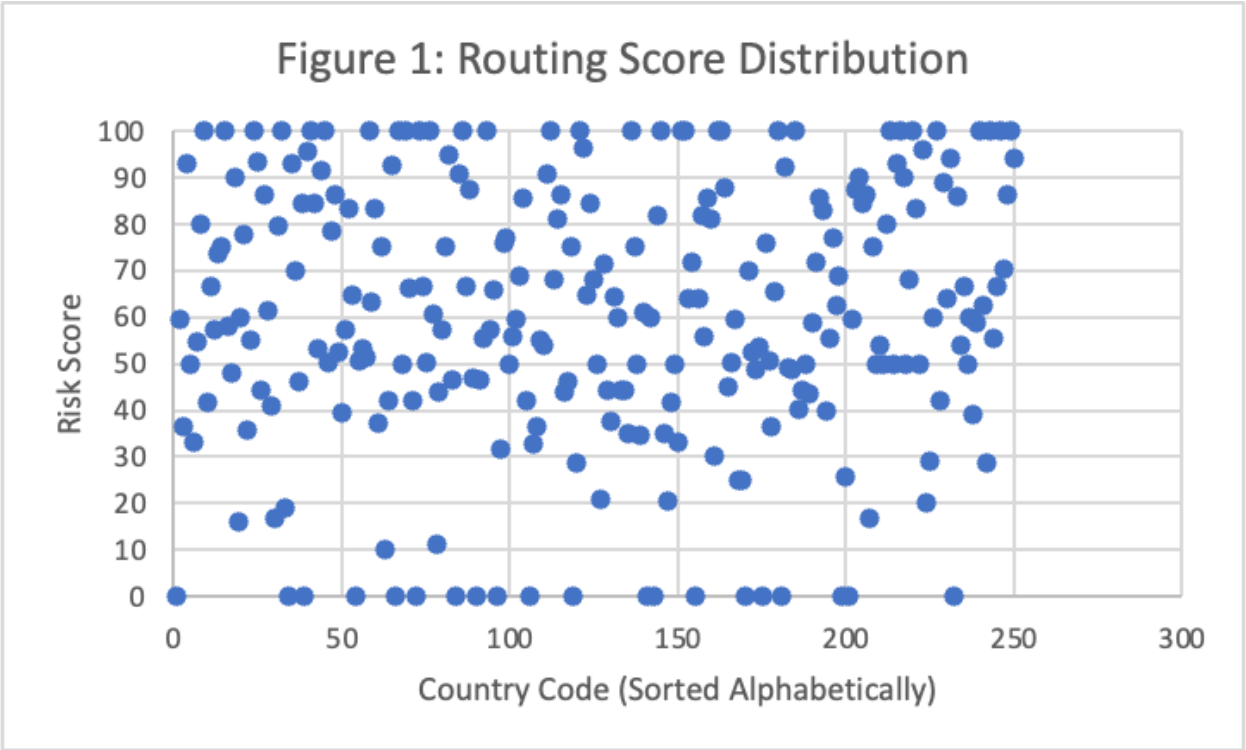

Figure 2: Risk scores, sorted

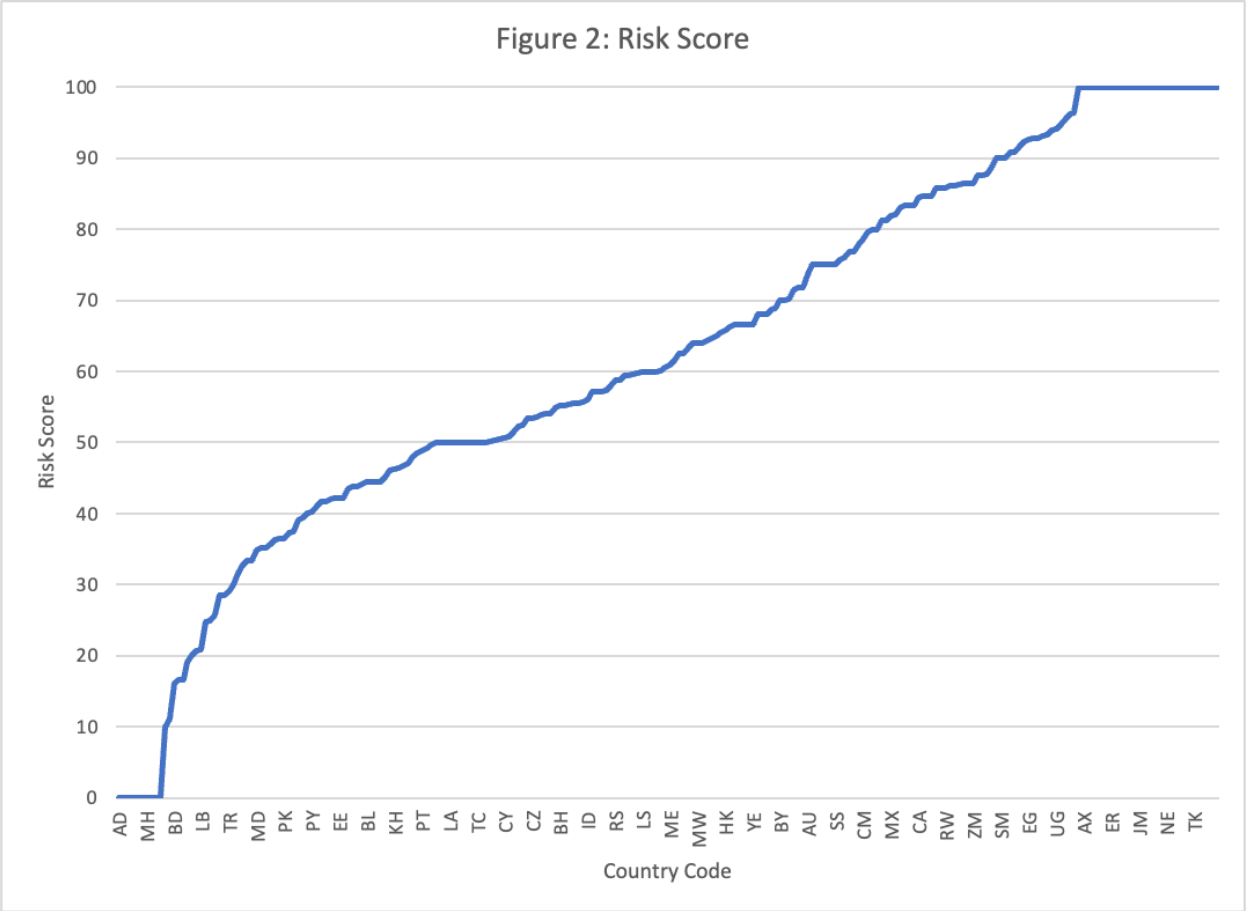

There are apparent discontinuities in the graph around 35 (the slope changes) and a gap close to 100. The left edge of the graph (at zero) is mostly small countries with minimal routing. The right edge of the graph at 100 is the 10 country codes which are not routing as of data gathering. At exactly 50 there is a set of 11 countries, and another 6 with a score under 51.5.

China (.cn) has a routing score of 86 (ranked #185), Japan scores 81 (ranking 168th), and Korea a 96 (ranking 210th).

Figure 3: Routing Rank

Rankings are an ordered list of scores. Unsurprisingly, the data are roughly evenly distributed. The longer horizontal segments in this graph where countries rank the same, except at the top and bottom, are somewhat surprising.

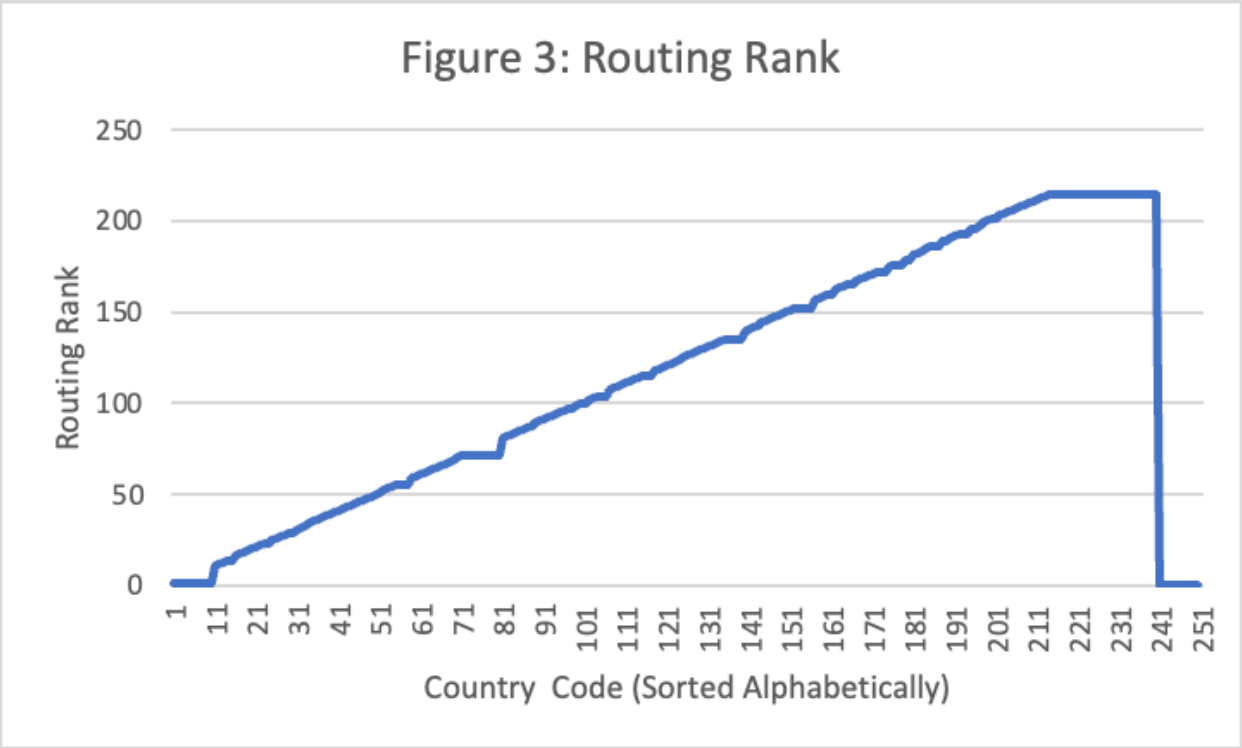

Figure 4: Aggregated Open Services score

Open service aggregated scores represent the denial-of-service capacity of open services in the country. Lower numbers are less capacity that can be used to attack. Scores are 0 to nearly infinity: the higher the score the more capacity, and, in this aggregate, are normalized against the number of IP addresses in the country.

Figure 4 shows aggregated open services. The median score is 185, while the average score is 2,034 and only 25 TLDs score above 5,000.

The vertical axis is dominated by the 5 scores above 10,000, which are:

CCTLD	Country/territory name	Score
TK	Tokelau	199870
SJ	Svalbard and Jan Mayen Islands	67990
AQ	Antarctica	63890
BV	Bouvet Island	24600
EH	Western Sahara	12300

Each of these very small territories may well have a high score because they have few IP addresses, resulting in a few open services being highly impactful on the calculated score.

China scores 154, Korea scores 116, and Japan 45 - each better than the median.

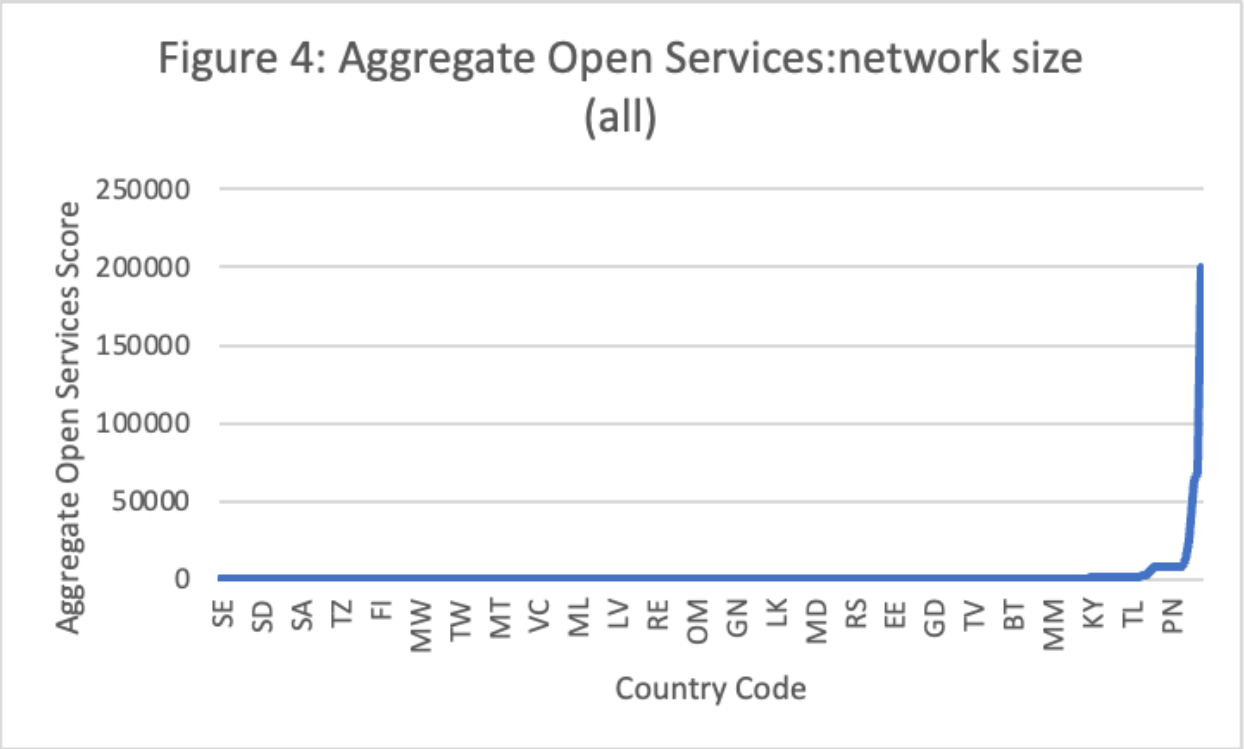

Figure 5: A subset of Figure 4's open services data.

Because Figure 4 is dominated by one very high score, Figure 5 shows the 236 countries in Figure 4 with a lower score, with the vertical axis scaled to make it easier to see the distribution.

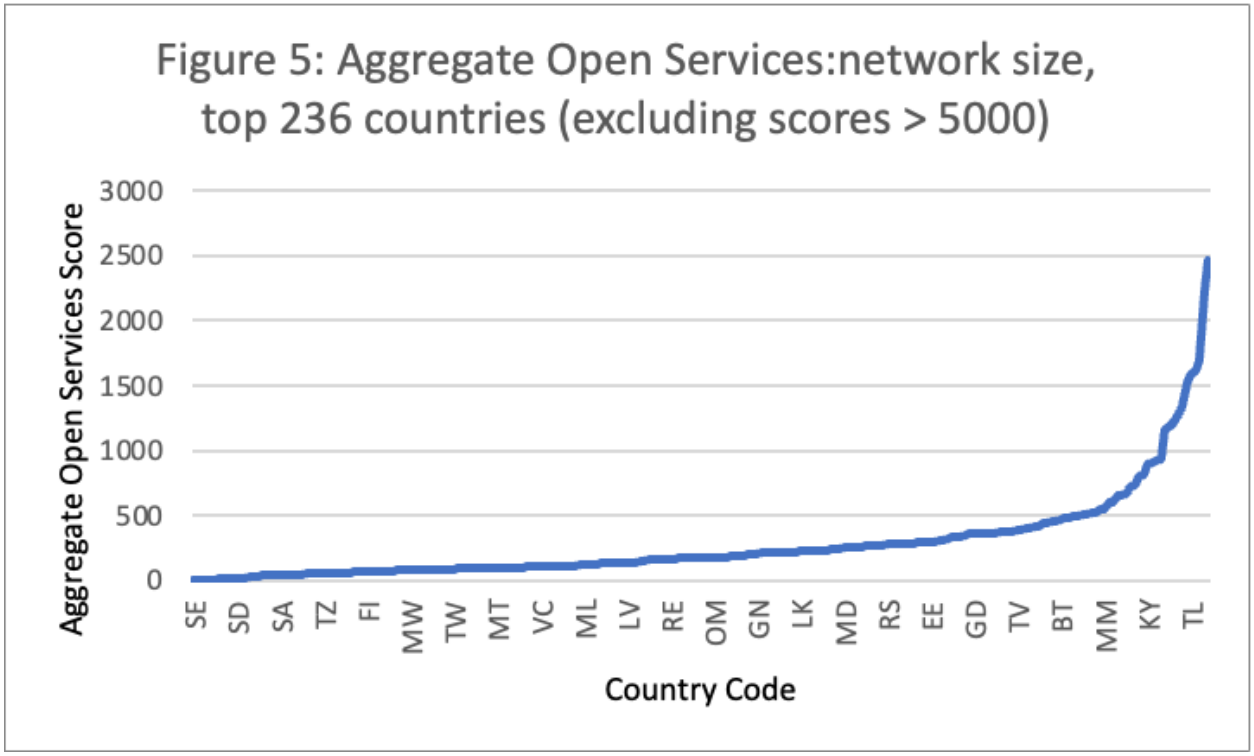

Figure 6: A different view of Figure 5 data

Figure 6 shows the same data as in Figure 5, but with the vertical axis being a log form to help see the distribution differently.

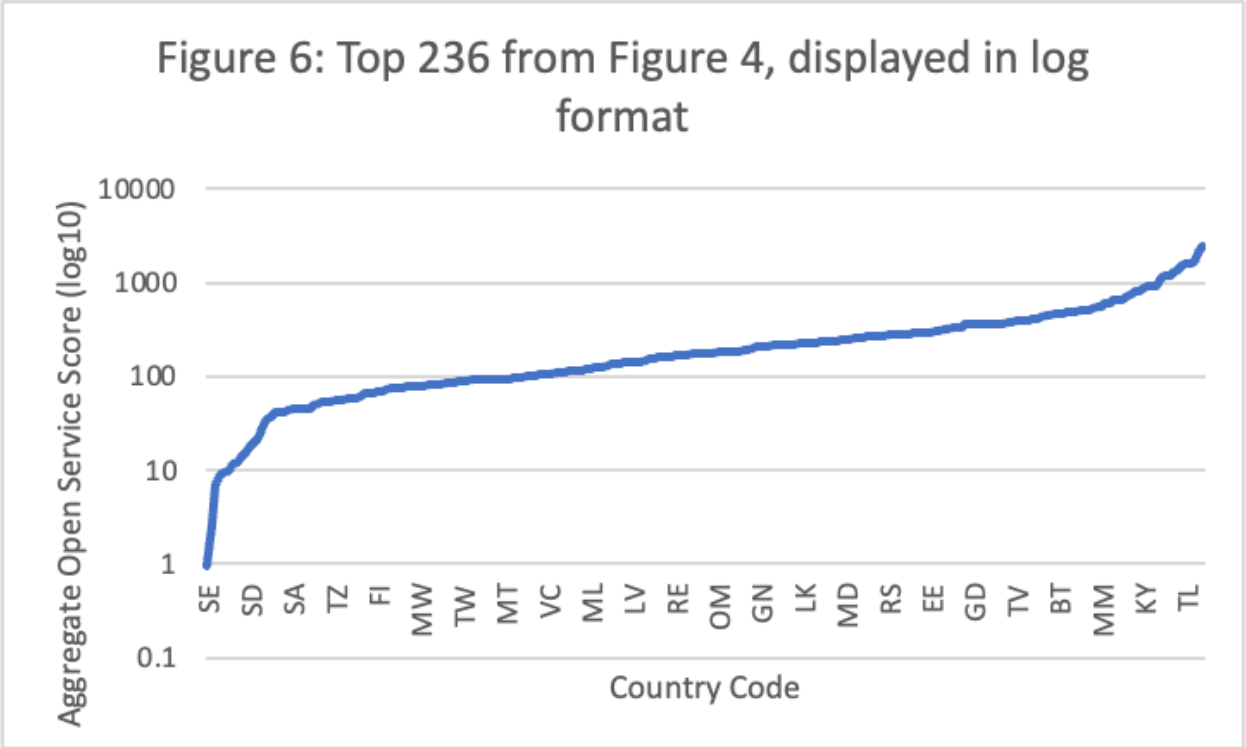

Figure 7: DNS Scores

DNS Scores incorporate data on DNSSec and on “lame delegations.” The scores are again a measure of insecurity. A domain with neither scores a zero, one with better security scores below a zero. For technical reasons, the data currently results in a range from roughly -10 to 10.

85 ccTLDs have neither DNSSec enabled nor lame delegations.

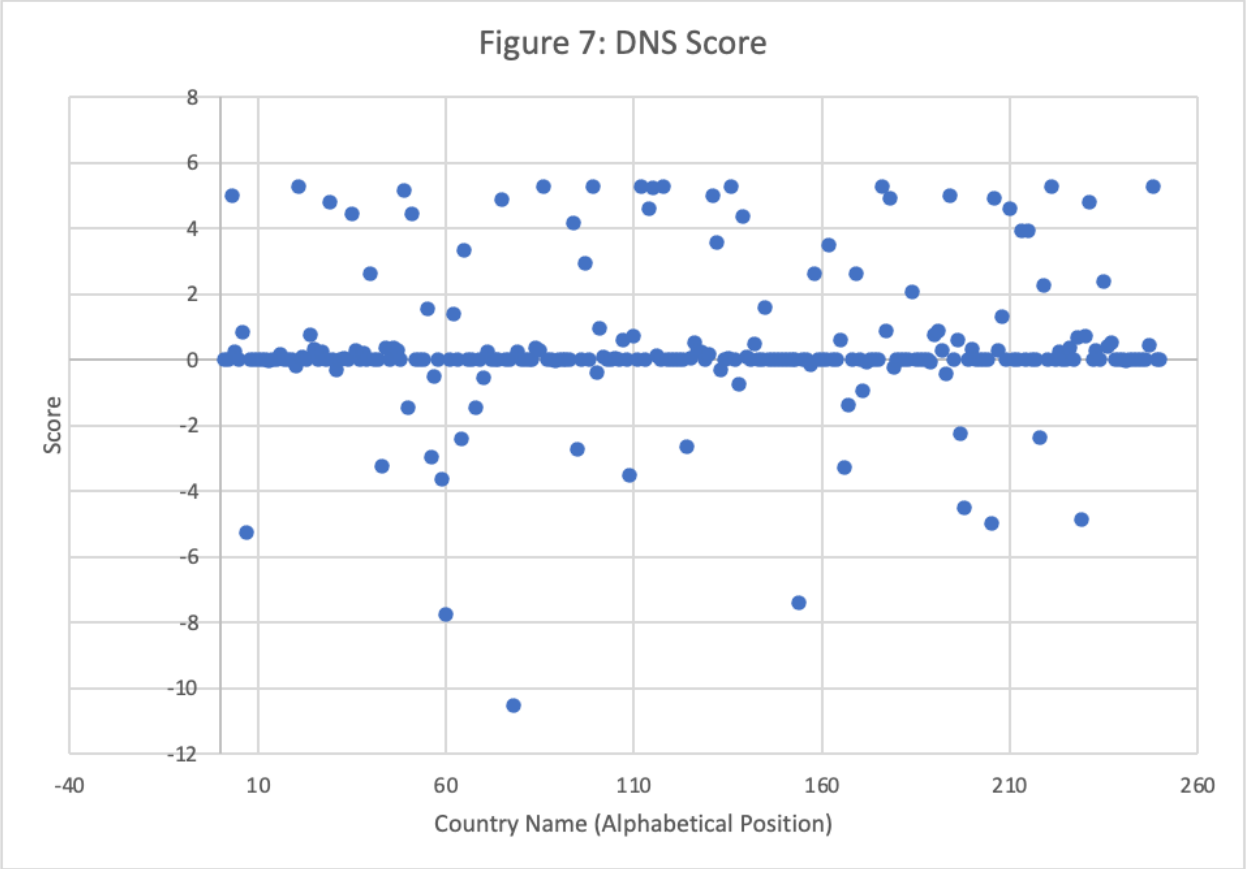

China has a score of .0017, with no DNSSec and 12 lame delegations.
Japan has a score of 4.6, with no DNSSec and 919 lame delegations. This is the fourth highest in our dataset.
Korea has a score of -.0043, 2 DNSSec enabled domains and 1 lame delegation.

Figure 8: DNSSEC Enabled Domains

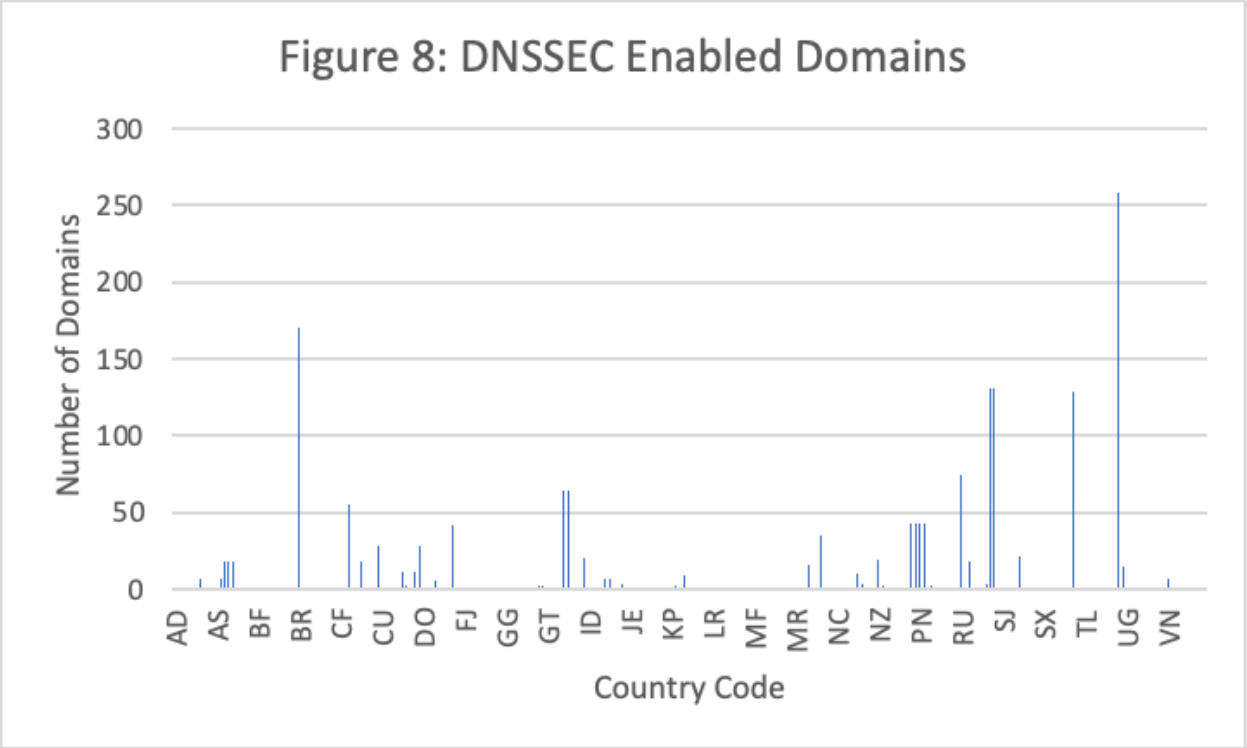

The average country has 6.7 DNS domains, with a median of 0. If we exclude the countries without DNSSec, the average is 29.

Figure 9: Lame Delegations

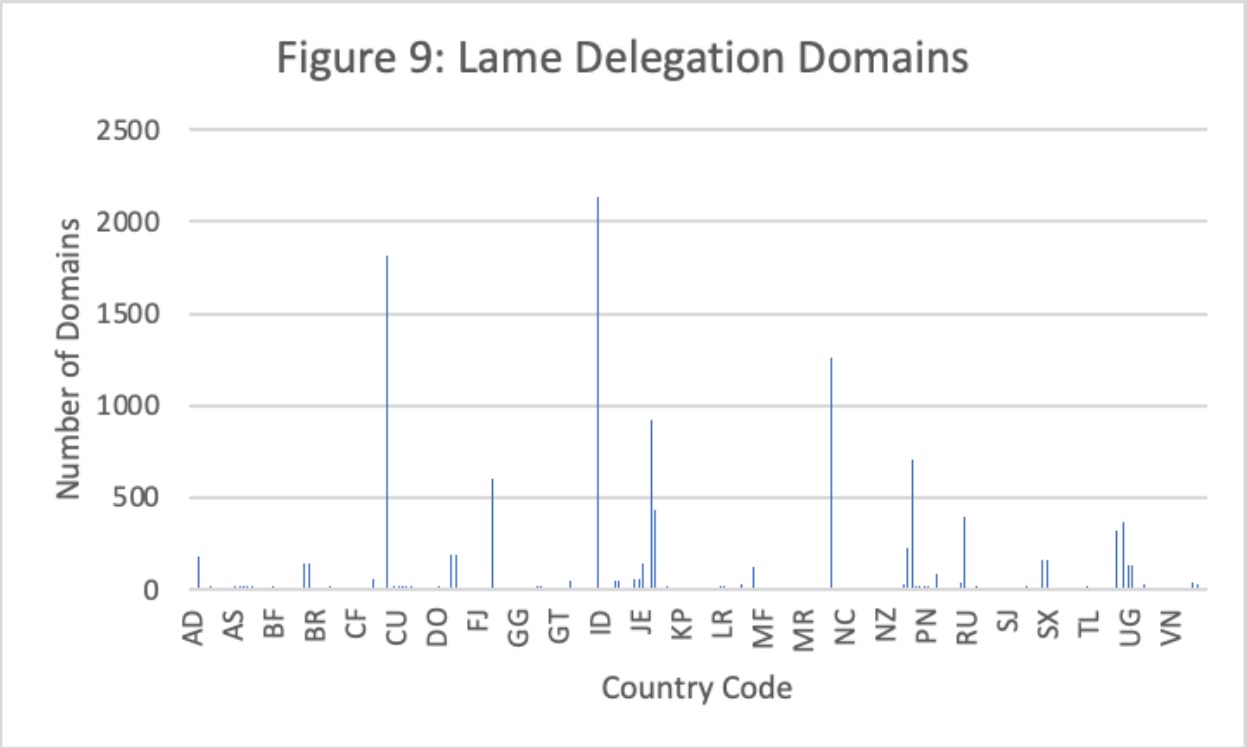

143 countries have Lame Delegations, with an average of 48, and a median of 1. Six countries (ID, CO, MY, JP, PK, and FR) have counts above 500:

Country code	Country name	Lame delegations
ID	Indonesia	2137
CO	Columbia	1812
MY	Malaysia	1263
JP	Japan	919
PK	Pakistan	707
FR	France	599

These countries seem meaningfully different from those shown in the table showing the outliers in Figure 4. It seems likely that management of these larger domains is more complex, giving more opportunities for error and improvement.

Figure 10: Lame Delegations Subset

Figure 10 shows the subset of domains with 1-499 lame delegations, sorted by count of lame delegations.

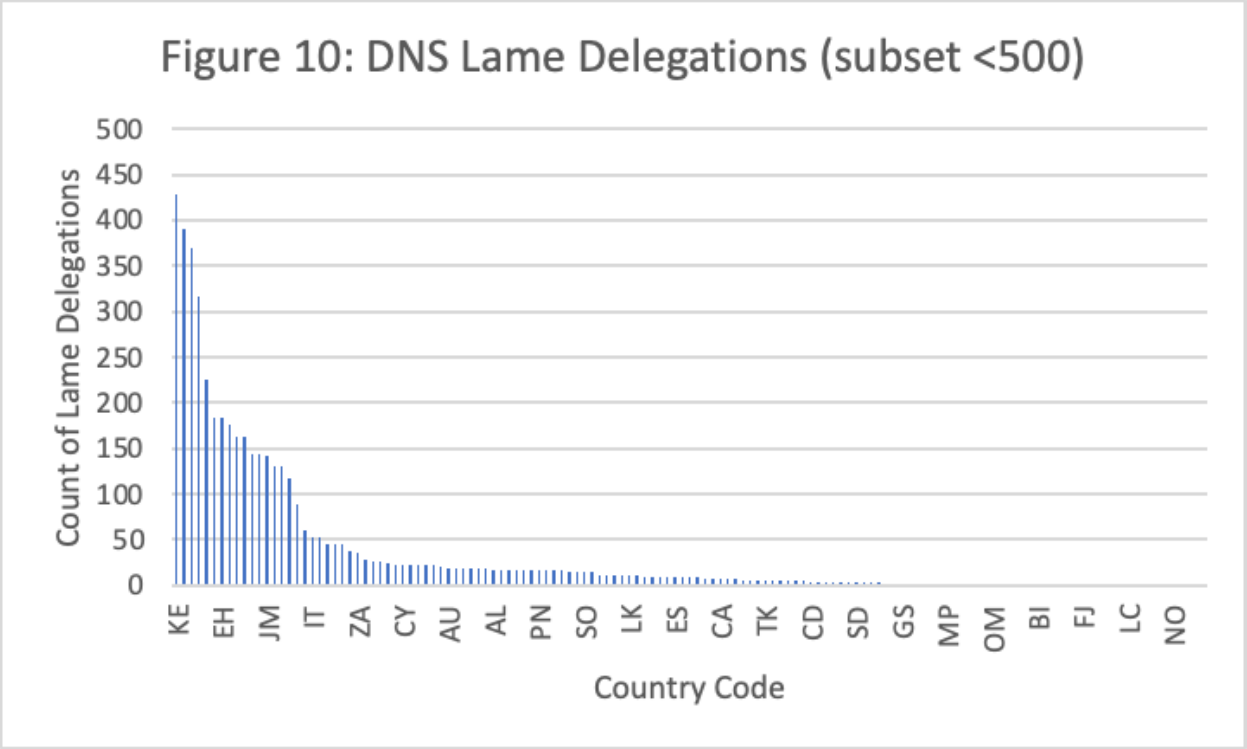

Known Challenges and Limits

This research is new. It points out important differences in the quality of internet services between countries, and represents important potential. Of course, it's preliminary, can be improved, and the improvements to make are determined by the uses to which the data and scorecard are put. That is to say, such improvements are best selected when we know that precision, accuracy or completeness would help us make better policy decisions. (Precision is how closely repeated measurements will match each other, accuracy is how close they are to the underlying truth. For example, if we take a several series of measurements of a coin flip and get a range of 49-51% heads, we would have high precision. We can also have high precision with a claim of 59-61% heads, which is either inaccurate, or we have a biased way of flipping the coin.)

Some of these challenges will be similar across our projects, including issues stemming from population counting and data quality issues. Others are specific to a project, such as routing spanning countries, while yet a third category are related to the software performing the analysis.

Populations.

1. Various population counts may not always line up perfectly. For example, an IP address within an ASN might be geolocated to one country while the containing ASN is allocated to another. So a count by IP may differ from count by ASN.
2. IPs and ASNs are categorized via geolocation to be connected to a particular country, but can either span multiple countries or be categorized as being in a country other than their actual location.
3. We say “country” and mean “two letter country code,” or ccTLD as defined by IANA. There are 193 countries recognized by the UN, and there are 308 ccTLDs, of which we generally collect data for roughly 250. We also have an artificial ccTLD of “zz” which we use when we have trouble allocating data to a real ccTLD.
4. IP addresses can be used in various ways, including routers, servers, and clients. Routers may be running Network Address Translation (NAT) and so may be the internet gateways for an arbitrarily large number of systems. The NAT pattern is used heavily by enterprise networks, home networks and mobile networks. What is “behind” each varies.
5. We don’t have a population of routes. There’s some information in databases like routeview, but those are advertised routes, but don’t include transit routes.

Time

In measuring churn day to day, we may be obscuring bigger patterns; we hope to investigate this in the future.

Data quality.

Our data is subject to measurement errors because of code quality, availability of scanned resources and possibly other factors. It is also dependent on what is available to us at reasonable cost and by choices made by companies. For example, if Google registers all of their data centers as being associated with an American company, while Amazon has local operating companies, that would influence our data. Google’s decisions would influence US numbers, while Amazon’s would influence those of each country where they operate.

Our use of geolocation may introduce issues because of challenges geolocating entities. We address this first by using data at a country granularity, which is generally believed to have fewer inaccuracies than more precise data. Second, we use commercial data sources, who we believe are accurate and up-to-date.

Routing-specific issues

Routes and ASes will often span countries - we allocate to one or another for tractability. We have multiple sets of routing data for historical reasons, and there are places at the margin where they don’t line up.

Software

Rate of change calculations: For simplicity of calculation, we are currently using an averaging function that results in different results than going across longer time periods. For example,

dropping 10% per day for 5 days results in a day 5 value of .59, where dropping 10% of the initial value each day would result in a day 5 value of .5.

That is, for $c_t = \frac{m_{t-1} - m_t}{m_{t-1}}$ gives us different results than computing the churn across longer time periods. Because we are applying this consistently we believe the results are reasonable.

Challenges with "Improvement" measures

Measures of improvement are both important and risky. For many scales (1-100, % improvement) those who are already performing well cannot improve as much as those who were performing poorly, either on an absolute basis or on a relative basis. It is important to consider both absolute performance and improvement.

Glossary

ASN: An Autonomous System (AS) is a group of one or more IP prefixes (lists of IP addresses accessible on a network) run by one or more network operators that maintain a single, clearly-defined routing policy. Network operators need Autonomous System Numbers (ASNs) to control routing within their networks and to exchange routing information with other Internet Service Providers (ISPs). (https://www.arin.net/resources/guide/asn/)

Hosts: Internet connected computers that do not provide routing services including desktops,

ROA: Route Origin Authorization is an attestation of a BGP route announcement. It attests that the origin AS number is authorized to announce the prefix(es). The attestation can be verified cryptographically using RPKI.

Appendix 1: Routing Scoring

1 High Level Factors to Rank By

- Routing
  1. Security
     (a) Valid ROA attested by RPKI
     (b) BGPSEC
  2. Resilience
     (a) Number of Direct Peers
     (b) Alternate Path reach ability
     (c) Identifying Potential unavailability

2 Proposed Ranking Algorithm

2.1 Snapshot Ranking - Daily

Algorithm 1 Snapshot of Daily Routing Ranking

Require:
$A_c \leftarrow \{R : A \mapsto C\}$ List of ASNs per Country
$R_c \leftarrow \{R : A \mapsto C\}$ List of ASNs with valid ROA Records grouped by country
$I_a \leftarrow \{R : \text{IP} \mapsto A\}$ List of IP ranges allocated to ASN

Ensure: the computation of the following metrics:
$M_1 \leftarrow$ Percentage of ROA routes per ISP in Country (Eq. 2)
$M_2 \leftarrow$ Percentage of ASNs per country with ROA valid records (Eq. 1)
$M_1 \wedge M_2$

$$M_1 = \frac{\mathbb{N}_{A_c} - \mathbb{N}_{\alpha_c}}{\mathbb{N}_{A_c}} * 100.0 \; \forall \; \alpha_c \in R_c \wedge \{c \in A_c = \; c \in R_c\} \tag{1}$$

$$M_2 = \langle c, \langle A_c, \frac{|\alpha_c|}{|I_a|} * 100.0) \rangle \rangle \; \forall \; c \in A_c \wedge a \in A_c \ni a = c \tag{2}$$

2.1.1 Daily Snapshot Ranking Explained

1. Group the list of ASNs by country of registration.
2. Group the list of ASNs with valid ROA Records by country of registration.
3. Identify the list of IP ranges allocated to each ASN.
4. For each ASN in the grouped list by country:
5. N1 = Identify the number of ASNs with Valid ROA records
6. N2 = Identify the number of ASNs registered
7. Compute $\frac{(N2-N1)}{N2} * 100$ to obtain the result M_1 which provides the country level rankings.
8. To provide ranking of ASNs within each country, perform:
9. For each country:
10. For each registered ASN:
11. Nr = Identify the number of valid ROA prefixes
12. Nt = Identify the number of total prefixes advertised by the ASN
13. Compute $\frac{(Nt-Nr)}{Nt} * 100.0$ to obtain M_2 which is ther ranking of individual ASNs within the country.
14. Present the results M_1 and M_2 in a ranked order by sorting the percentages computed in descending order and ranking the indices in ascending order (indexed at 1).

2.2 Time Based Ranking

Notes

1. u is the percentage of IP space in use. For example, an ASN after registration with a RIR obtains the ownership for 128 IP prefixes, it decides to advertise and use only 64 of them, the u is 50%.
2. Similarly, the p value represents the number of IP prefixes secured through ROA. Continuing the example, the ASN signs the advertisement routes for 48 IP prefixes of the advertised 64, this makes the p value 75%.

3. In the implementations so far and the data we have, we only compute p and do not compute u since we do not authoritatively know the ownership space. The current assumption is p = u where advertised prefixes = ownership i.e. 100

4. the `Zip` function is a combination function, rather than a compression. The output is a list of $< Country, List(ASN_1..ASN_n). >$

5. The `ASC` produces a numerically ascending sort.

Algorithm 2 Ranking Countries by Routing Security

Require:

$U \leftarrow$ Universe of IP Prefixes
$C \leftarrow$ List of Countries
$A \leftarrow$ ASN Identity
$\mathcal{I} \leftarrow \text{Zip}(C, A)$
$\mathcal{O} \leftarrow$ Ownership space of IP Addresses; $\mathcal{O}_i \subseteq U \ \forall \ i \ \epsilon \ \mathcal{I} \mid \sum_{i}^{|\mathcal{I}|} \mathcal{O}_i \ = U|$
$\mathcal{S} \leftarrow \{A, A \mapsto \{C, \mathcal{O}_i\}\} \ \forall \ i \ \epsilon \ \{\mathcal{I}\}$
$\alpha_A \leftarrow$ Observed Advertised IP Prefixes per ASN
$\beta_A \leftarrow$ Valid ROA Prefixes with an attested RPKI Trust Anchor
$T \leftarrow$ Total number of recordings made over time
With a valid precondition $\beta_A \subseteq \alpha_A \subseteq \mathcal{O}_A \subset U$

Ensure:

For each country $c \ \epsilon \ C$
$\exists$ a set of ASN A_c and satisfying the precondition
Compute ASN Churn as presented in Equation 3
Churn for the time period is $\frac{(\S 3)}{|T|}$
Compute Per ASN IP Advertising Churn as presented in Eq. 4 with M_1

$\rho \leftarrow$ Proportion of churn compared to IP space $(N_{\text{churn}}/\alpha_A)$
$\rho_t \leftarrow \frac{\rho}{|T|}$ This can be ranked relatively using a CDF for ASNs within the country
For the Country level information
`GROUP BY` C and `COMPUTE` $k = \int_0^{|A_c|} \frac{1}{|T|} \int_0^{|T|} \rho_t \cdot dt \ \forall \ c \in A_c$
`ORDER BY` k `ASC`
`COMPUTE` CDF, Perform `PERCENTILE` bucket ranking or `FRACTIONAL` rank tie break

$$\sum_{t=0}^{|T|} \frac{\frac{(M_2)_t - (M_2)_{t-1}}{(M_2)_t}}{|T|} \tag{3}$$

$$
\begin{aligned}
u &= \frac{\alpha_A}{\mathcal{O}_A} * 100.0 \\
a &= 100 - u \\
p &= \frac{\beta_A}{\alpha_A} * 100.0 \\
\frac{\sum_{t=0}^{T} u_t}{|T|} &\times \sum_{t=0}^{T} Pr(u_t \leq A_c) + ...
\end{aligned}
\tag{4}
$$

u is the utilized IP address space, a is the available address space, p is the fraction of used address space which is routing secure with an ROA record.

Design Question: Do we want to incentivize higher IP space availability? How do we treat that as a vector for an even larger scoring algorithm?

2.2.1 Time Based Ranking Generation

On a day-to-day basis, Given the universe of IP prefixes U, the pre-requisities for the algorithm need the mapping of ⟨Country, ASN⟩, and subsequent ⟨ASN, Prefixes⟩. The combination of these pre-requisites helps establish the relationship between the ASN, Country, and the Ownership of Prefixes advertised. The ⟨ASN, Prefixes⟩ are of three different types

- $\mathcal{O}$: Corresponding to the ownership of the IP address space (Allocated)
- α: Corresponding to the advertised IP prefixes (In Use)
- β: Corresponding to the RPKI Validated Prefixes (In Use with RPKI)

For each day, For each country in the list of countries, Identify the list of ASNs from the ⟨Country, ASN⟩. To rank at the country level over time T, for each day in the interval, compute the percentage of the number of ASNs in the country to the count of the number of ASNs with *any* valid RPKI ROA route. This result is the same as the computation of M_1 from the snapshot based ranking computation. The result of each of these computations are stored in an equivalent data structure $\langle t_i, \langle \text{Country}, M_1 \rangle \rangle$.

Compute the daily churn by considering $N-1$ entries of the interval T, and identify the daily percent change in the snapshot scores resulting in $\frac{M_{1\,t-1} - M_{1\,t}}{M_{1\,t-1}}$ indicating the relative percent change per day. A positive change here indicates improvement while a negative change indicates deterioration and 0 indicates stability. For the $N-1$ intervals ranging from $(t_{N-1}, t]$ compute the average of of the daily percent changes by summing the daily changes and dividing it by the duration of the interval $|T|$. The current result here generates two possible ranking outputs to consider for policy decisions:

1. Ranking the Countries with the most significant positive/negative changes in the given time period.

2. Coupled with the daily snapshot in Algorithm 1, allows for ranking based on absolute percentage value ranked in descending order combined with the ascending order ranking of the relative change in given interval, the equivalent of `SORT BY [COL1 DESC, COL2 ASC]`

Similarly, the computations happen a level deeper within each country to rank the ASNs per country over time. Similar to the description above, this results into the ranking of ASNs within a country which has the most significant positive/negative changes in the given time period and needs to be coupled with the daily snapshot to provide a ranking over the time $|T|$. To do this, consider the same time interval and the datastructure $\langle t_i, \langle \text{Country}, \langle \text{ASN}, M_2 \rangle \rangle \rangle$. Within each country, compute the average relative change in number of RPKI ROA prefixes to the advertised prefixes by the ASN. This value is associated with the ASN within the country.

Therefore the current sets of computations result into 4 individual rankings which are generated as the baselines.

1. Country wise ranking over Time indicating most improvement
2. Country wise ranking including snapshot and relative improvement in routing security.
3. Within each country:
   (a) ASN level ranking over time indicating most improvement.
   (b) ASN level ranking taking snapshot and relative improvement into account.

In addition to the above results, a good indicator for the ranking security is the number of IP addresses which are protected by the routing, this can be computed by identifying the total number of IP addresses which are advertised by the ASN (N_α), and the number of IP addresses among those advertised which are covered by a valid secure RPKI signed route (N_β). The resulting churn is computed for each of time interval for $N-1$ intervals as N_{churn} and the proportion ρ is computed by $\frac{N_{\text{churn}}}{\alpha}$. The daily distribution as a CDF allows us to understand the total IP addresses in use and those which are secured per country.

Visualization Note: X axis = Number of ASNs in ranked order (most IPs to least IPs), Y axis = CDF of IPs advertised, CDF of IPs RPKI secured.

2.3 Identifying `intersection`

ASNs have the ability to advertise larger IP prefixes but sign the underlying set of IP addresses in smaller CIDR sizes. For example, it is possible to advertise a `/16` IP range, but sign two `/17` IP ranges or four `/18` IP ranges, or eight `/19`s etc.., In the naïve computation of N_1, N_2 in the snapshot algorithms, this results into negative numbers. However, this is definitely valid. The goal to identify

N_2 is to identify how many prefixes the RPKI prefixes cover in the advertised prefixes. This is done by construction a `trie` or a radix-tree. (Refer to easy to use libraries like `pytricia`)

1. $\mathcal{T} \leftarrow$ Construct an empty tree.
2. For each prefix p in advertised prefixes by ASN P_a:
3. `T.Insert(p, [ASN, p])`
4. Construct a map $M \leftarrow \langle \text{P}, [\,] \rangle$ for each prefix in ASN's advertised prefixes
5. For each RPKI prefix r signed by the ASN R_a:
6. $a, p =$ `T.lookup(r)`
7. `M.Insert(p, r)` if $a ==$ ASN
8. N_2 is the number of keys in the map M for which the `len(value) > 0`

Note the above computation of N_2 is slightly relaxed, a pedantic check would be to see for each prefix (key) if the entire IP range is covered by the values ($list[r]$) inserted as a part of the `intersection` procedure.

Appendix 2: Identifying Lame DNS Delegations

To identify a lame delegation, we can query a domain's authoritative nameservers. This provides the NS records and corresponding *glue* records for the domain. By querying each name server listed for a specific query, if each of the name servers respond with the data authoritatively (`AA` bit set), the DNS records are configured correctly, else the records indicate a lame delegation.

Let us take an example of the following DNS configuration to understand.

D = example.com.
NS = [ns1.example.com, ns2.example.com]

Algorithm:

1. Start with the DNS root zone and retrieve the corresponding NS records for the root servers. (Example result here)
2. Use one of the name servers to query for the next section of the zone chain i.e. TLD (.com.) in this case. This returns the NS for the .com TLD and the corresponding glue records indicating the IP addresses. (here)
3. Query the domain's Nameservers from the TLD NS records by connecting to a NS IP among the TLD name servers. Here in this example, we'll query for example.com's NS records from the .com TLD nameservers. This returns a result presented here. These results could be presented with/without glue records depending on how they're configured. Resolve accordingly for the IP addresses of the nameservers listed here.
4. For each name server in the list of nameservers for the domain D
    a. Query the DNS Query type (A / AAAA / MX etc..,) and capture the response
    b. Check if the response has DNS AA bit set.
    c. An example of resolving the DNS A record for example.com from the nameservers in step 3 (a.iana-servers.net.) and (b.iana-servers.net) are here (A response, B response).
    d. Identify name servers which perform the following actions:
        i. Timeout / Fail to respond
        ii. Return SERVFAIL/NXDOMAIN
        iii. Don't respond with the AA bit set in the response
        iv. If any of the conditions above i, ii, iii meet, classify the domain D as a domain having a lame delegation.