Workgroup: intarea
Internet-Draft:
draft-giuliano-blocking-considerations-00
Published: 7 March 2022
Intended Status: Best Current Practice
Expires: 8 September 2022
Authors: L. Giuliano M. Aelmans T. Li

Regional Internet Blocking Considerations

Abstract

Geopolitical conflicts can cause policy makers to question whether or not blocking the Internet connectivity for an opposing region is a constructive tactic. This document provides an overview of the various technologies that can be used to implement regional blocking of Internet connectivity and discusses the implications of these options. This document does not advocate any policy or given blocking mechanism, but does attempt to articulate the implications of these blocking technologies for policy makers. The document also intends to help inform policy makers from countries who could be exposed to such blocking techniques on the implications of these methods.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at https://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 8 September 2022.

Copyright Notice

Table of Contents

1. Introduction

Geopolitical conflicts can cause policy makers to question whether or not blocking the Internet connectivity for an opposing region is a constructive tactic. This document provides an overview of the various technologies that can be used to implement regional blocking of Internet connectivity and discusses the implications of these options. This document does not advocate any policy or given blocking mechanism, but does attempt to articulate the implications of these blocking technologies for policy makers. The document also intends to help inform policy makers from countries who could be

exposed to such blocking techniques on the implications of these methods.

The content expressed in this document solely reflects the views of the authors and do not necessarily reflect the views or positions of any of our organizations, affiliates, friends, or enemies.

1.1. Requirements Language

The key words "**MUST**", "**MUST NOT**", "**REQUIRED**", "**SHALL**", "**SHALL NOT**", "**SHOULD**", "**SHOULD NOT**", "**RECOMMENDED**", "**NOT RECOMMENDED**", "**MAY**", and "**OPTIONAL**" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2. Scope

The scope of this document it limited to a description of well-known methods for disrupting core Internet services including physical cabling, Internet routing, filtering, and the Domain Name System.

The document does not intend to give any political directions or advocate for implementing the described methods, nor does it intend to be a guide for malicious attackers hence it will purely describe concepts and does not provide actual configuration or implementation methods.

3. Disconnection Methods

There are many ways of blocking a region's Internet connectivity. In this section, we discuss some of them, their implications, capabilities, advantages, and disadvantages.

3.1. Physical layer

Disconnection at the physical layer is the most definitive method. Cutting cables and severing fibers will most definitely stop bits from passing. Unfortunately, this approach is also the most expensive to repair. In the optimistic view that disconnection is only intended to be temporary, creating downstream costs of physical repair is distinctly suboptimal. This approach may also be selected by the unscrupulous who have physical access to the media, but do not have further physical or managerial access.

A less destructive physical layer disconnection is simply disconnecting the fiber or cable, either at the terminating device, optical module, patch panel, amplifier, repeater, or transponder. This is easily repaired, but still requires physical access. Unscrupulous parties that wish to prevent easy repairs would be unlikely to select this option.

The simplest physical layer disconnection is administrative shutdown. Managerially disabling a physical circuit is a trivial configuration option that will sever communications. It is trivial to revert.

3.2. Routing layer

The Internet is a collection of many enterprises, web and cloud hosting, access-providers, etc. networks connected to each other using the Border Gateway Protocol (BGP) [RFC4271] to exchange routing information between those networks.

The simplest explanation of how BGP works is to compare it to a group of networks using the earlier described physical connections to gossip with each other on their knowledge about their own and neighboring networks. In other words, they exchange routing information describing how to reach destinations on the Internet.

Connected entities play different roles in this ecosystem, some will know how to reach all destinations on the Internet. These networks are so-called Tier 1 providers and provide connectivity to Tier 2 and Tier 3 networks. They offer services on a global scale and connect thousands of networks.

Tier 2 providers are large regional or national operators (for example the national service providers) offering services in country or regional. They have connections to many other networks, provide services to Tier 3 networks, but also purchase services from Tier 1 networks to reach destinations on the Internet they cannot reach themselves.

Tier 3 providers don't provide routing information knowledge to other networks and are dependent from Tier 1 and Tier 2 networks to reach destinations on the Internet. In this category are small service providers or webhosting providers.

The BGP protocols offers several options to manipulate the routing information that is exchanged between these networks.

3.2.1. Autonomous System Number Filtering

Networks participating in the exchange of routing information with other networks use unique Autonomous System Numbers (ASNs) to identify themselves. These numbers are assigned to them by the Regional Internet Registries.

The ASN is used to construct a path to a destination prefix on the Internet. For example, a Tier 1 advertises its prefix to a Tier 2 originating from its own ASN. Next the Tier 2 advertises the prefix

to a Tier 1 and adds its own ASN to the path. The route to reach this prefix now looks as follows: ASN2 ASN1.

As networks are highly connected there are many ASN paths through the Internet to reach destinations. The 'further away' the destination is the longer the ASN path will be. On average most destinations on the Internet are reachable in a maximum of 5 hops. In other words, most destinations on the Internet can be reached via a maximum of 5 networks.

Networks that have a need to filter out another network with which they don't have a direct peering session completely have, next to filtering prefixes, the option to filter on ASN. When applying an ASN filter, it will filter out all prefixes originating from that specific ASN.

Mitigating ASN filtering requires similar measures as mitigating prefix filters; networks with many upstream connections to Tier 1 and 2 networks will have a much lower chance of being completely filtered as, if one out of many upstream peers filters the ASN (and so its originating prefixes), others might still propagate them. This could still result in prefixes not being globally reachable anymore, but the chances are much lower.

3.2.2. De-peering

BGP uses a TCP session between two networks to exchange routing information. Such a session is called a peering session. Disabling such a session is referred to as de-peering.

3.2.2.1. De-peering Tier 3 networks

In many cases Tier 3 networks are using a single Tier 1 or 2 network for their connectivity to the Internet. In that case it's relatively easy to disconnect such networks from the Internet by disabling their peering sessions on the Tier 1 or 2 side.

3.2.2.2. De-peering Tier 2 networks

As described earlier these networks have multiple connections to other Tier 2 providers and typically between 2-8 Tier 1 providers to provide connectivity to the Internet. Subsequently, they could also receive routing information via Internet Exchange Points giving them even more options to reach destinations on the Internet.

De-peering such a network is much harder as one would need to disable peering sessions in many networks and at multiple (probably international) locations. Tier 2 networks will likely have international connections as well. Pursuing networks to disable

these peering sessions in another jurisdiction could be very complicated.

3.2.2.3. De-peering Tier 1 networks

By their nature, Tier 1 networks have global span and have thousands of connections with other Tier 1, 2 and 3 networks. Fully disconnecting such networks is considered almost impossible without having physical and administrative access to the network itself. Pursuing other networks to de-peer a Tier 1 network is impossible because of the many countries they are present in and their jurisdictions.

3.2.3. Countering De-peering

Entities that want to protect themselves against de-peering would have a diversified connectivity strategy including multiple Tier 1 and 2 peers, actively peering on Internet Exchange Points, and preferably possessing its own physical infrastructure to connect to other networks in different countries or regions.

Tier 3 networks are most vulnerable to de-peering.

3.2.4. Prefix Filtering

Each network that is part of the Internet uses unique IPv4 and IPv6 address prefixes ranges to expose services to its directly connected (local) customers but also those connected via the Internet. These prefixes are advertised over a BGP peering session to the neighboring network so they will learn which prefixes originate from their neighbor and know how to reach them. Subsequently, they will advertise any routing knowledge they have about their neighboring networks and the neighboring network of their neighbors, etc. This way every network builds it own view of the Internet and map of how to reach destinations.

For example, Tier 1 and 2 networks will both have 'downstream' (customer) peering sessions with networks of which they have knowledge about; the prefixes they are advertising. If one of these networks wants to filter a neighbor, they could de-peer them as discussed earlier but that would basically filter all prefixes. In many cases, for example when intending to filter out social media, a subset of the prefixes is enough to accomplish this goal.

With this method a Tier 1 could also filter out prefixes from a Tier 3 that it learns via a Tier 2. De-peering the Tier 2 would result in all Tier 2 and all its customer prefixes becoming unreachable via this Tier 1. If only prefixes advertised by the single Tier 3 need to be filtered, the Tier 1 applies a prefix filter to the peering session from which it receives the advertisements.

Contrary, networks with many upstream connections to Tier 1 and 2 networks will have a much lower chance of being completely filtered as, if one out of many upstream peers filters the prefixes, others might still propagate them. This could still result in the prefix not being globally reachable anymore, but the chances are much lower.

3.3. Packet Filtering

Most network layer devices have the ability to filter traffic. The mechanism for doing this is commonly called an "Access Control List" (ACL). This is a possible mechanism for implementing a disconnection. Typically, an ACL allows filtering on a combination of source address, destination address, protocol number, and TCP/UDP source or destination port number.

3.3.1. GeoIP ACLs

The question then becomes one of ACL construction. However, this is not simple. IP address space is delegated in large sets, commonly known as 'prefixes.' Each prefix is assigned to an organization. Some organizations, such as Internet Service Providers (ISPs) will in turn delegate a portion of their address space to a customer. Customers and service providers do not necessarily fall along clean regional lines. Large multi-national corporations can receive a prefix from an ISP in one region and may use it in an entirely different region or even globally. They may also receive a prefix directly from a Regional Internet Registry (RIR). Service providers can obtain a prefix from an RIR and delegate parts of that prefix to customers from another region. This can create both false positives and false negatives when trying to map between a prefix and a region.

There are services which attempt to provide mappings from an IP address or prefix to a region, commonly called 'GeoIP' services. However, due to the above issues, these services cannot guarantee their accuracy. Constructing an ACL based on GeoIP services is likely to have unintended consequences, both filtering unintended addresses and not filtering intended addresses. Some commercial applications (notably streaming video) are willing to accept these inaccuracies, but this may not be acceptable in all circumstances.

Virtual Private Networks (VPNs) and other tunneling mechanisms can be used to create virtual topologies. If a single VPN server within a target region is not blocked, then it can provide access to innumerable other systems within the region, effectively bypassing GeoIP filtering services. When these are discovered, they are typically added to GeoIP databases, but this creates an ongoing battle between VPN service providers and GeoIP providers. As a

result, this is an imperfect solution that may or may not be sufficiently accurate.

3.4. DNS

Blocking DNS capabilities can be an effective method for inhibiting end users from easily accessing Internet resources in a given region. For example, removing the delegation entries in the root servers for a given country code can prevent users from resolving the names of all domains for that country code. This approach can be circumvented to an extent with the creation of stub zones on resolving nameservers, which can provide a shortcut delegation to the country code top-level domain servers (ccTLDs) that are authoritative for that country code. But these stub zone entries would have to be manually created on any resolving nameserver that serves the resolution requests of users seeking resolution of domains for that particular country code.

In the opposite direction, blocking resolution requests can inhibit users coming from a region from easily accessing Internet resources. Specifically, filters can be used to block resolving nameservers from a given region, or can block resolution requests from end users within a given region from making resolution requests to resolving nameservers that reside outside that region.

4. Gaps

The mechanisms discussed above cover the salient technical points for blocking a region. In this section, we discuss the various other considerations that are relevant to regional blocking.

4.1. Information Dissemination

At the very lowest level, the Internet copies bits from one location to another. Bits that are injected at one point are packetized, forwarded, and hopefully show up at their intended destination. The technology of the Internet does not care what is encoded in those bits. Whether it is state secrets or yesterday's grocery list, the Internet will happily ship it all the way around the world in milliseconds. The intrinsic value, properties, and attributes of the information conveyed in those bits is immaterial at the technological level.

4.1.1. Information Value

Policies considering blocking the transfer of information must also consider the value of the information that is being blocked. Filtering mechanisms can be extremely coarse and block all information, and this may not match the purposes of the policy.

Thus, a blocking policy may need to be extremely specific about its goals and purposes.

A policy may want some information to be able to enter into a region. Sending certain messages into the region may be beneficial to the policy maker. Similarly, being able to get information out of a region may be beneficial. Further, parties within a region may be depending on global Internet connectivity to coordinate activities. A policy that blocks too much information may be counterproductive to the aims of the policy maker. A more selective policy would want some information to be communicated and not other information. Further, a selective policy is likely to be highly directional. Information that should flow into the region may not be permitted back out, and vice versa.

4.2. Information Concealment

If a policy allows any information to transit a boundary, then there is the technological possibility that other information may also transit that boundary. Information can be disguised or concealed through the use of cryptography, steganography, or other techniques. Policy makers should assume that any mechanism that allows any information to transit a boundary would eventually be used to transfer information against the purposes of the policy.

4.3. Misinformation

If a policy blocks information from flowing into a region, that may allow parties within that region to generate misinformation that is not disputed by outside information. This may be highly advantageous to the parties within the region. In the past, there have been many occurrences when parties within a region disconnected from the Internet precisely so that internal information could not be disputed.

4.4. Target Inaccuracy

The Internet infrastructure does not assign address space or ASNs according to strict regional, national, or continental boundaries. While there is some rough correlation, that is the result of administrative convenience. Thus, a prefix that is allocated from the general pool of European address space may end up covering part of Europe and Greenland. An ASN that was allocated for Singapore may be used in Australia.

This is further complicated by the fact that the parties that receive an ASN or prefix are not obligated or constrained to use it in a given region. If an organization acquires an ASN and subsequently grows outside of its original region, it may still use that ASN. If a company is assigned a prefix and the company is

acquired by another firm, then that prefix could be used in a completely different hemisphere.

Consequently, if a policy elects to block traffic based on ASNs and prefixes, it may have unintended consequences, potentially blocking unintended traffic and not blocking proscribed traffic.

4.4.1. Accuracy of Registry Information

Many public resources are available to query Internet routing related information including, IPv4, IPv6 and ASN resource holders, routing intentions and actual reachability data. Unfortunately, the data doesn't always represent the actual situation, can be incomplete and in quite a few occasions outdated.

4.4.1.1. Internet Routing Registries

Internet Routing Registry (IRR) databases hold information about network operators routing intentions. For example, ASN holders can specify with whom they have peering relationships. This could give an indication which networks a specific ASN is connected to, however the data is entered (manually or automated) by network operators and isn't per se verified.

In practice IRR databases are between 40-70% accurate. However, some show an accuracy of around 95%.

4.4.1.2. RPKI

Resource Public Key Infrastructure (RPKI) is a public key infrastructure (PKI) framework to support improved security for the Internet's BGP routing infrastructure. The most important property is that in RPKI only legitimate resource holders can make statements about the IPv4, IPv6 and ASN resources they hold. This means that any information, right or wrong, found in RPKI databases represents the intention of, or at least is entered into RPKI by, the rightful holder.

RPKI is therefore considered to be 100% accurate. The downside of RPKI is that there aren't records for every resource and a large portion of the IPv4, IPv6 and ASN resources don't have records in RPKI.

4.5. Spoofing ASNs and Hijacking Prefixes

If a policy attempts to filter routing advertisements based on an ASN, then the opposition may attempt to counter that filtering attempt by using an alternate ASN. The alternate ASN may be an unused one, an ASN that has been assigned but is not actively in use elsewhere, or could be one that is actively assigned to another

party. Using this ASN, the opposition could advertise its prefixes into BGP, bypass the ASN filter, and regain connectivity.

Similarly, if a policy attempts to filter routing advertisements or implement forwarding plane filters based on assigned prefixes, then the opposition may attempt to circumvent these policies by obtaining, advertising, and deploying alternate prefixes. As with ASNs, these prefixes could come from unassigned address space, address space that has been assigned but is not actively advertised, or even address space that is actively advertised by other parties.

There are security mechanisms that have been developed to help counter these possible attacks (IRR filtering [RFC7682], RPKI [RFC6811], and BGPsec [RFC8205]), but they are not ubiquitously deployed and may or may not be effective, depending on the operational procedures of ISPs that provide connectivity to the region.

4.6. Porous Borders

The Internet is, by design, a decentralized system of interconnections. Thus, it is nearly impossible to completely block Internet access for a region. Simply put, there will always be ways to circumvent any blocking attempts by sufficiently motivated parties. However, there are certain chokepoints and various methods, as described above, that can significantly inhibit connectivity and throughput for users going to/coming from a given region.

4.7. Acknowledgments

This document was inspired by the thoughtful comments of many friends and colleagues.

5. IANA Considerations

This document makes no requests of IANA.

6. Security Considerations

This document discusses technical and policy considerations of blocking Internet access for regions, and their potential impact on global security.

This document does not present new attack or defense strategies and merely discusses the implications of a variety of technical approaches. This document does not advocate or dissuade any policy about blocking Internet connectivity, it discusses various considerations that policy makers should understand prior to setting policy.

7. References

7.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

7.2. Informative References

[RFC4271] Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, DOI 10.17487/RFC4271, January 2006, <https://www.rfc-editor.org/info/rfc4271>.

[RFC6811] Mohapatra, P., Scudder, J., Ward, D., Bush, R., and R. Austein, "BGP Prefix Origin Validation", RFC 6811, DOI 10.17487/RFC6811, January 2013, <https://www.rfc-editor.org/info/rfc6811>.

[RFC7682] McPherson, D., Amante, S., Osterweil, E., Blunk, L., and D. Mitchell, "Considerations for Internet Routing Registries (IRRs) and Routing Policy Configuration", RFC 7682, DOI 10.17487/RFC7682, December 2015, <https://www.rfc-editor.org/info/rfc7682>.

[RFC8205] Lepinski, M., Ed. and K. Sriram, Ed., "BGPsec Protocol Specification", RFC 8205, DOI 10.17487/RFC8205, September 2017, <https://www.rfc-editor.org/info/rfc8205>.

Authors' Addresses

Lenny Giuliano

Email: lenny@lenny.net

Melchior Aelmans

Email: melchior@aelmans.eu

Tony Li

Email: tony.li@tony.li