A Technical Overview of Content Blocking Methods

May, 2014

by
Pier Carlo Chiodi

Preface

I wanted to write this document in the wake of a discussion born on the RIPE Cooperation Working Group mailing-list; the main goal is to obtain a draft to work on with the group in order to reach a final guide for policy makers and legislators who are called upon to solve social issues by using web blocking mechanisms.

Because I wanted to keep a very low technical level I made some approximations; please drop me a comment if they were too significant.

Summary

Goals and intended audience

Goals of this document are quite ambitious:

- explaining how the Internet works in a manner as simple as possible, so that even non-expert people could understand the mechanisms which let it to go on;
- having an overview of web blocking measures application contexts;
- explaining advantages and disadvantages of various blocking measures;
- focusing on cross-border / human rights issues.

Only technical topics related to web blocking measures are covered, or those which are needed to be known to fully understand implications of web blocking.

This document is particularly addressed to legislators, agencies, stakeholders, courts and to whoever may be involved in Internet governance and engaged in law enforcements on the network, and also to who is interested in this topic and would like a basic understanding of mechanisms and approaches used by whoever to block or prevent access to contents.

How Internet works

Internet is a network of networks, each one connected with one or more of the other. Every network is administered by a company which autonomously applies its local policies and rules on the traffic traversing its devices. Of course, every company is subject to local laws.

Let's assume that all these networks could be summarized in 3 macroscopic groups:

- Internet Service Providers (ISPs), which allow users to connect their devices to Internet;
- Content Providers and Web Sites Operators, which host web sites, mail boxes, services and media contents on their servers;
- Transit Providers, which "just" allow other networks to exchange traffic with each other.

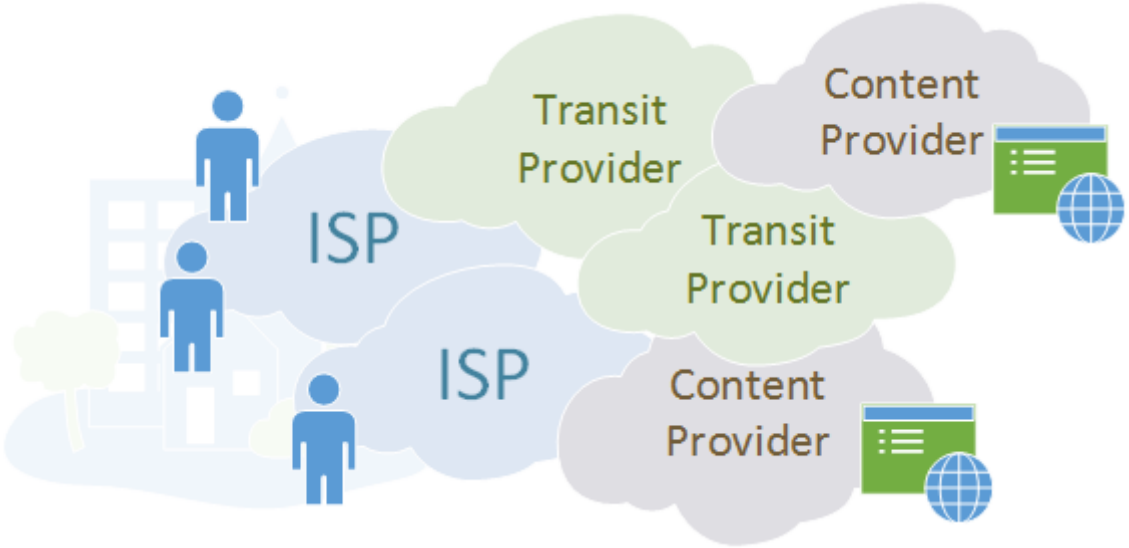

Users are connected to their ISPs; in turn, ISPs may be directly connected to as many Content Providers as possible or may pay one or more Transit Providers to reach those Content Providers that they don't have a straight link to. Content Providers host resources on their servers on various forms and protocols: HTTP (web sites, personal pages, blogs and so on…), SMTP, POP3, IMAP (for mail boxes), FTP (file sharing) and many others.

All connected devices talk each other by exchanging "packets"; each packet traverses links and networks from the source to the destination and carries commands and instructions to setup a communication between parties and to exchange data, images, videos and audio streams. Everything happens in a client/server model, where a requesting device (the client) asks information to another one (the server) using a specific protocol that both understand.

Packets are addressed using a numerical identifier, an IP address (Internet Protocol); every device has one or more IP address assigned to it and use them to set the source and the destination of packets. Since it would be difficult for people to remember all those numbers, a facility has been introduced: DNS, Domain Name System. DNS allows to associate a numerical IP address to a name used to represent and locate an Internet resource or service. Using DNS, for example, we can reach a website by typing www.example.com instead of 93.184.216.119. This system works in a hierarchical and distributed manner, where each component of the hierarchy delegates the one below, so that a chain is made:

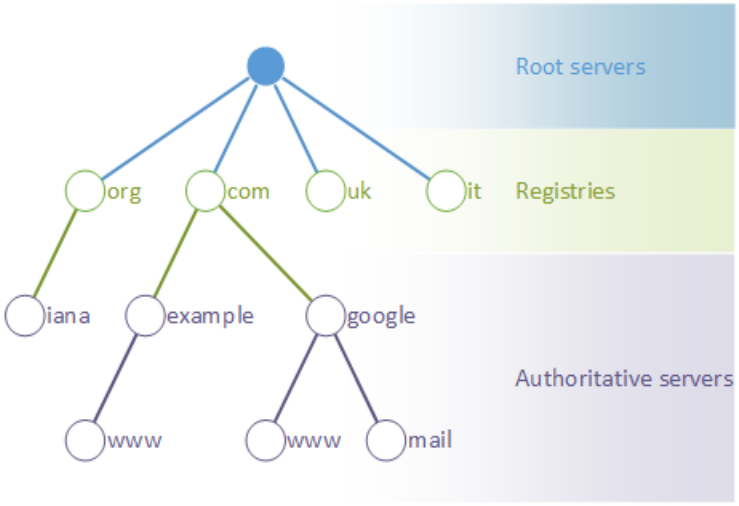

- the "root servers", which are spread across the world, contain a list of other servers responsible for the "Top Level Domains" (TLDs), such as .com, .org, .net or geographic areas like .uk, .it, .de, .fr and so on;
- "Registries" are the entities which manage those servers in charge of TLD zones (.com in the case of www.example.com); they hold the list of all the domain names (like example.com) which belong to the TLD zone they handle;
- "Authoritative" servers, the last leaves of the tree, are in charge of the "zone" of each domain name and contain the list of the service that each domain offers (www in the example).

Please, keep in mind that all these servers are often operated by different companies, on different countries and under different jurisdiction.

While the above servers hold the hierarchical structure of Internet names, another kind of servers allow users to use it: they are the "recursive resolvers", or simply "resolvers". Devices connected to Internet are (manually or automatically) configured with the IP address of one or more of this kind of servers; often ISPs provide their own resolvers to their customers, but many other companies offer the same service on Internet, even for free.

Once an user types an Internet resource address (URL) on his/her browser, or when he/she follows a link present on a web page that he/she is already visiting, or when an image is loaded, the following happens (assume that the user wants to open the web page at http://www.example.com/page1.html):

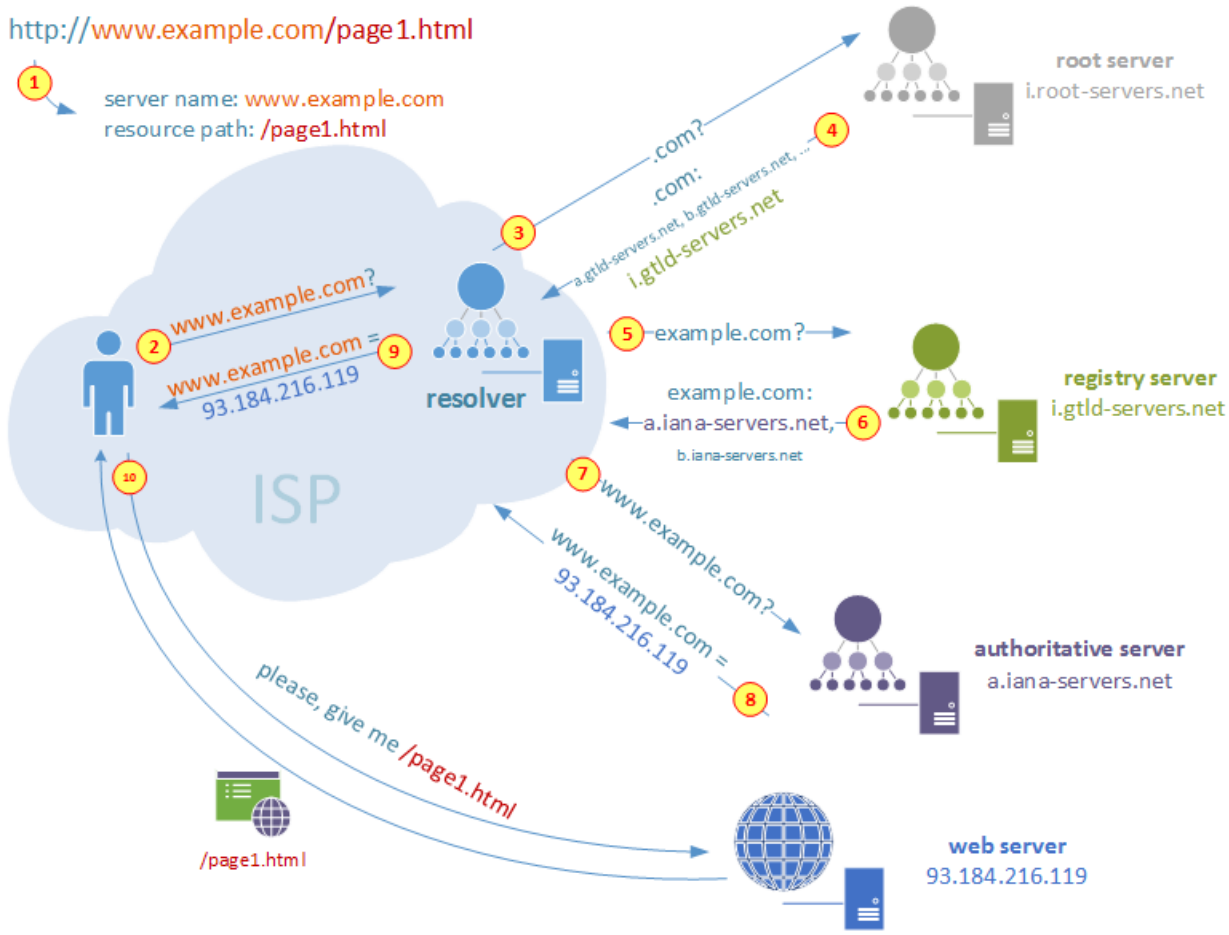

1. the user's application splits the URL in two parts and extract the server's name (www.example.com) and the requested resource's path (/page1.html);
2. the device then sends a DNS query to the configured resolver asking for the IP address of the web server which hosts the requested resource (www.example.com);
3. the resolver starts descending the DNS hierarchy by first asking who is the server responsible for the .com TLD to one of the root servers it knows (i.root-servers.net in the example);
4. the requested root server answers and refers a list of servers maintained by the Registry in

charge of the .com TLD;

5. the resolver picks one server from the list (i.gtld-servers.net) and sends a query to it asking which is the Authoritative server holding the domain name zone (example.com);
6. the Registry's server answers with a list of two Authoritative servers (a.iana-servers.net and b.iana-servers.net) and …
7. … the resolver sends the final query to one of them (a.iana-servers.net), asking the IP address of the server which handles the requested resource (www).
8. Finally the resolver receives from the Authoritative server the wanted response…
9. … and forwards it to the user's device.
10. The user's web browser now knows the IP address of the web server and sends it a request for the wanted resource (/page1.html), receiving it back.

All these request and response packets are often expected to cross many networks and countries, depending on the location of the ISP and the Content Provider and also on their local policies and interconnection agreements. It's pretty obvious that DNS is a crucial point regarding the security of Internet. The (not so) "simple" hijacking of a response given by a server may lead to a traffic diversion and bring users on servers managed by criminal entities, which can then steal data, passwords and money.

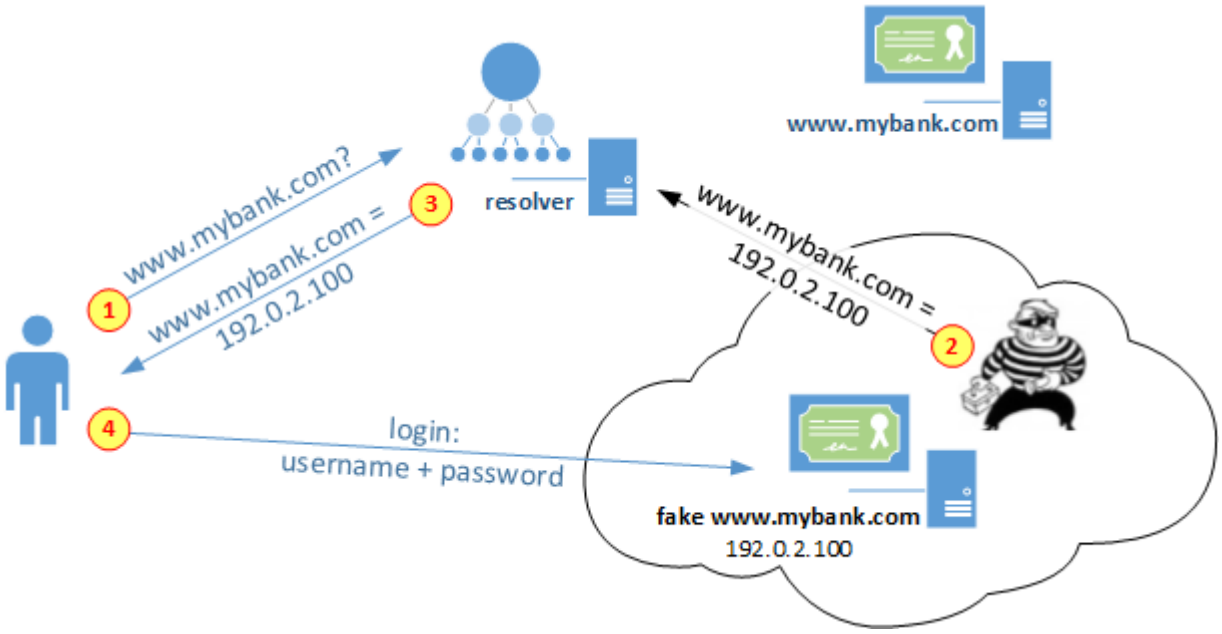

In order to solve this kind of problem a DNS extension has been developed, DNSSEC, whose adoption rate is growing day after day within the whole Internet community.
DNSSEC adds security to the Domain Name System; using cryptographic mechanisms each response provided by DNSSEC-aware servers is digitally signed so that any forgery would be detected by resolvers and users' devices. Every DNS response is authenticated using a chain of trust having its anchor on the cryptographic key at the root zone, managed by trusted root servers, so that resolvers could verify both Registries' and Authoritative Servers' responses.

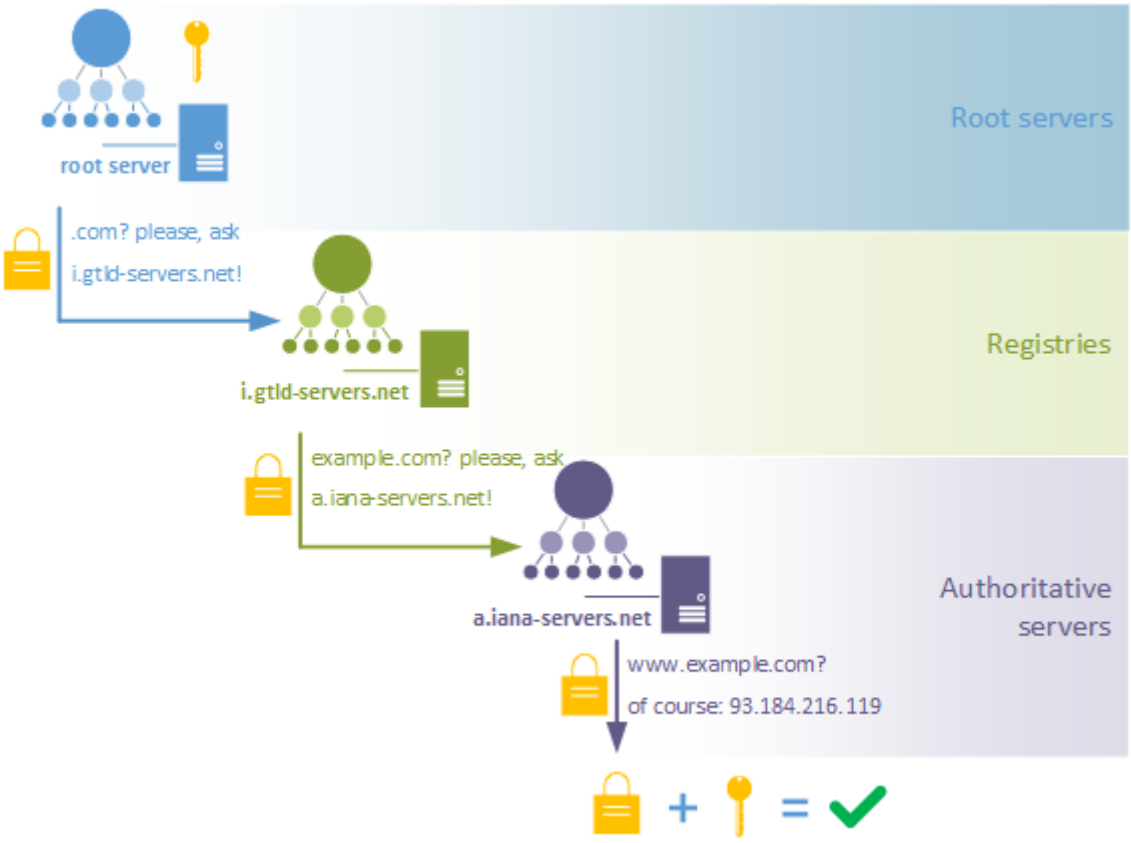

While at this time most of the implementations of DNSSEC are limited to resolvers, the main goal of DNSSEC is to extend this model to end users' applications, so moving the validation function from resolvers to users' operating systems.

Web filtering needs

Web blocking and filtering are measures usually requested by governments or law enforcement agencies (LEAs), addressed to prevent access to illicit contents such as pedo-pornography, unhautorized gaming and gambling, piracy, or even to constrain access to opposing political or religious contents or to quiet debates that threaten the parties in power.

These measures are particularly used when the undesired content is hosted on servers that are out of the jurisdiction of the requesting party, so when it's not possible or very difficult to order the website operator to remove the unwanted material from its servers. In such cases ISPs operating under the jurisdiction of the requestor are imposed to prevent their customers to access the identified resources.

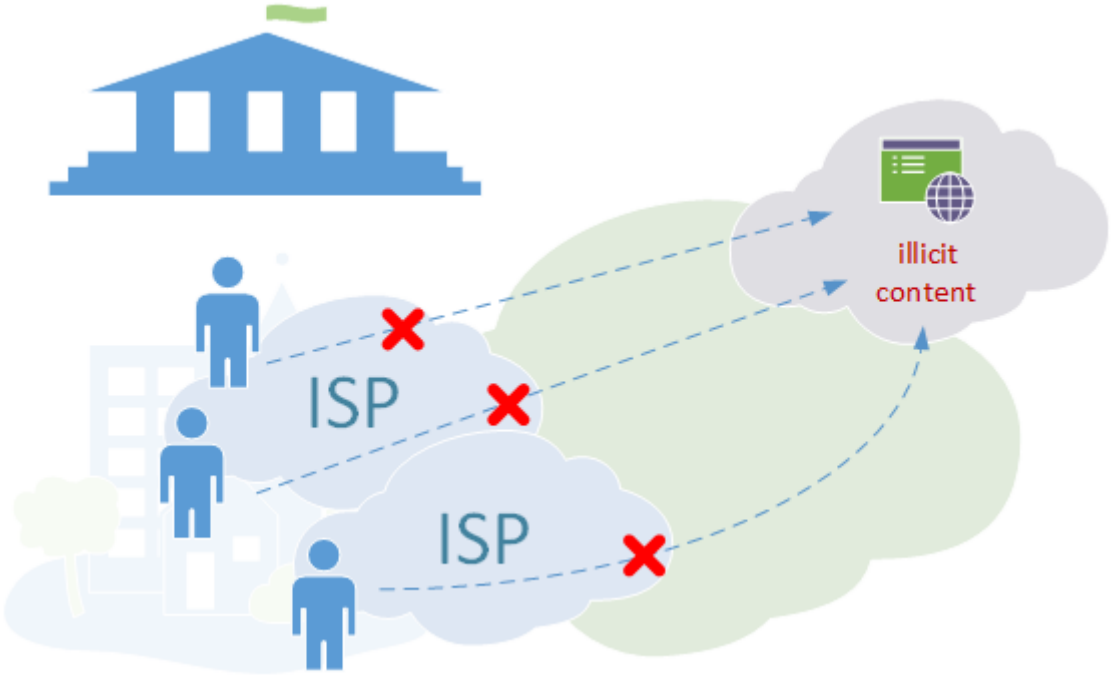

Optionally, they may be asked to redirect customers who try to access the forbidden content to a web page reporting additional information, such as the legal notice about the blocking measure ("stop-page").

Many methods can be used to comply with this purpose but, as we'll see later, all of them have drawbacks that may compromise Internet operations, security and reliability and also human rights, privacy and freedom of expression.

Control points and methods

In order to prevent access to web resources, a block or a filter must be placed along the path between users and contents. End users' devices make up a such wide set of publicly available hardware and software which is almost impossible to control or to constrain, so they can't be considered as a valid point where to implement filtering methods. Content Providers' networks (and also Transit Providers' ones) can't be considered too because, on the basis of our assumption, they are out of the jurisdictional control of the entity which requests the block (if they were, there would be no need to request a web filtering measure but they could be imposed to remove unwanted content from their server).

Following is a list of control points and methods that can be used to block access to resources; please refer to the following diagram to have a better view of enforcement points:

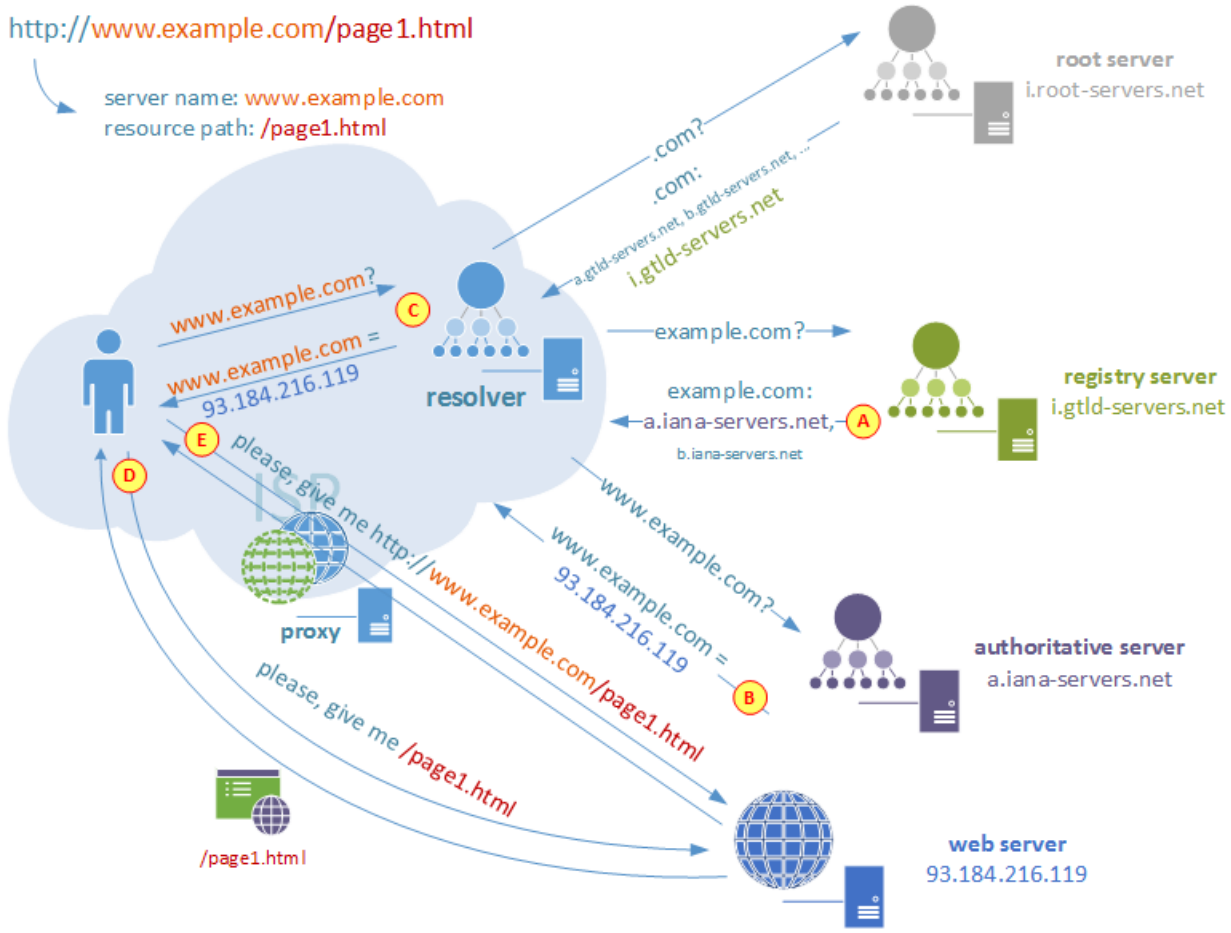

DNS Registries (A)

Requestors order a Registry to stop answering DNS queries for a target domain, or to answer with an Authoritative name-server which redirects users to a stop-page.

DNS queries made by users for that domain name have no answers or are resolved with the IP address of a server which only publishes a stop-page.

DNS Authoritative Servers (B)

This case is similar to the above, but here requestors send the take-down order to the Authoritative Servers manager.

ISP DNS Recursive Resolvers (C)

This is a control point closer to the end users than the previous two; requestors order ISPs to block a target domain on their recursive resolvers, optionally redirecting customers to a web stop-page.

ISP IP address block (D)

In this case requestors order ISPs to block the IP addresses associated to the server which publishes the undesired content. With no regards to DNS translation, every request made toward the target IP address is dropped on the ISP's network.

ISP Web Proxies (E)

The networks operated by ISPs may be instructed to divert all web traffic toward specific devices,

called “web proxies”, which perform an analysis of each web request and eventually drop those toward blocked domain names or URLs.

ISP Deep Packet Inspection

Like the previous scenario, ISPs’ networks may be configured to perform a deep packet inspection (DPI) of all the data traversing them (not only web) and eventually drop traffic patterns which match those reported by requestors.

Collateral damage

Before moving on to a deep analysis of various filtering methods, here is a brief overview of some collateral damage they may produce; it will be helpful to fully understand the analysis that follows. Moreover, an additional analysis of side effects will be provided in the rest of the document too.

Overblocking

With the exclusion of some particular cases (such as country-wide Internet shutdown imposed by regimes) the enforcement of blocking measures may be requested against aggregate resources (domain seizure) or specific resources (content filtering).

While a domain seizure is aimed to prevent access to all the resources published under a specific domain name, a content filter is focused only on a subset of them.

Measures which lead to domain name seizure, when used to prevent access to only a subset of resources, produce the heavy collateral damage of a complete black-out for the target domain name. For example, the domain seizure for example.com would prevent access to both www.example.com/GoodContent.html and www.example.com/BadContent.html.

This behaviour is particularly dangerous for those platforms which host blogs, discussion forums and personal pages where many people and companies publish their thoughts and projects: to block a single offending web page the whole platform is cut off.

For example, on 20 November 2012 the European Court of Human Rights ordered[1] Turkey to refund an Internet user whose personal website (http://sites.google.com/a/ahmetyildirim.com.tr/academic/), hosted on “Google Sites” platform and used to publish academic works, was shut down as a result of a take-down order for another website, hosted on the same platform (http://sites.google.com/site/kemalizminkarinagrisi/benimhikayem/atatuerk-koessi/at) and aimed to insult the memory of Atatürk, founder of the Turkish Republic. The blockade prevented the whole platform (sites.google.com) to be accessed by Turkish users and the European Court deliberated that this fact amounted to a violation of Article 10 of the European Convention on Human Rights and Fundamental Freedoms, guaranteeing the freedom of expression “without interference by public authority and regardless of frontiers”.

DNSSEC breakage

As described before, DNSSEC is a technology developed to ensure data authenticity along the whole DNS hierarchy. It prevents forged responses to be sent in place of real ones, a technique used by

[1] “Case of Yıldırım v. Turkey, Application no. 3111/10”, European Court of Human Rights: http://hudoc.echr.coe.int/sites/fra/pages/search.aspx?i=001-115705

criminal entities to hijack web traffic and divert it on unofficial servers to gather confidential information and passwords from users.
Many DNS blocking methods break the chain of trust on which DNSSEC is based because they exactly inject false responses, indistinguishable from those that DNSSEC fights. The DNSSEC goal is to provide a full chain of trust extended to users' applications and devices, but if it's broken by requestors orders within ISPs networks every effort is vanished. Many other security-oriented protocols put their basis on DNSSEC, so the whole suite would be compromised if DNSSEC can't be trusted anymore.

> *The world's economy can either have secure Internet naming and therefore secure Internet applications, or have effective content blocking via Internet DNS – but not both.* [2]

Analysis of blocking/filtering methods

On the basis of the IETF draft "Technical Considerations for Internet Service Blocking and Filtering"[3], five criteria will be used for analysis of blocking methods (one more than IETF work):

- scope: to evaluate which users are blocked;
- granularity: to evaluate how specific is the blockage/filter and how it impacts on other services and contents;
- efficacy: to evaluate how difficult it is for users to avoid the blocking measure and keep accessing the forbidden resource;
- security: to evaluate impacts of the blocking measure on the security of Internet, meant in terms of availability of service, authenticity, confidentiality and integrity of information;
- feasibility: to evaluate difficulties and costs related to the implementation of the method.

DNS Registries

This method implies the removal of a domain name from a Registry, or its configuration with a name-server which only redirects users to a stop-page.

Scope

Since the Registry is the top most level of the DNS hierarchy which holds data regarding a specific domain name, blocking at this level has impacts on the whole Internet, with disregard of jurisdictional borders of the entities requesting the block.

Granularity

The removal of a domain name from a Registry cause every DNS query made toward that domain, for any service related to it (web, email, VoIP, …) to fail like if that domain had never been created. Whether the blocking request was made for a specific content or for the whole domain name, every service provided by that domain and any sub-domain is interrupted. For example, in case of Registry removal of example.com it would no longer be possible to access www.example.com/BadContent.html but also www.example.com/GoodContent.html or

[2] "SAC 056 - SSAC Advisory on Impacts of Content Blocking", ICANN: http://www.icann.org/en/groups/ssac/documents/sac-056-en.pdf
[3] "Evaluation of Blocking Design Patterns, Criteria for evaluation", IETF: http://tools.ietf.org/html/draft-iab-filtering-considerations-04#section-4.1

blog.example.com/MyPersonalPage.html nor send an email to info@example.com.

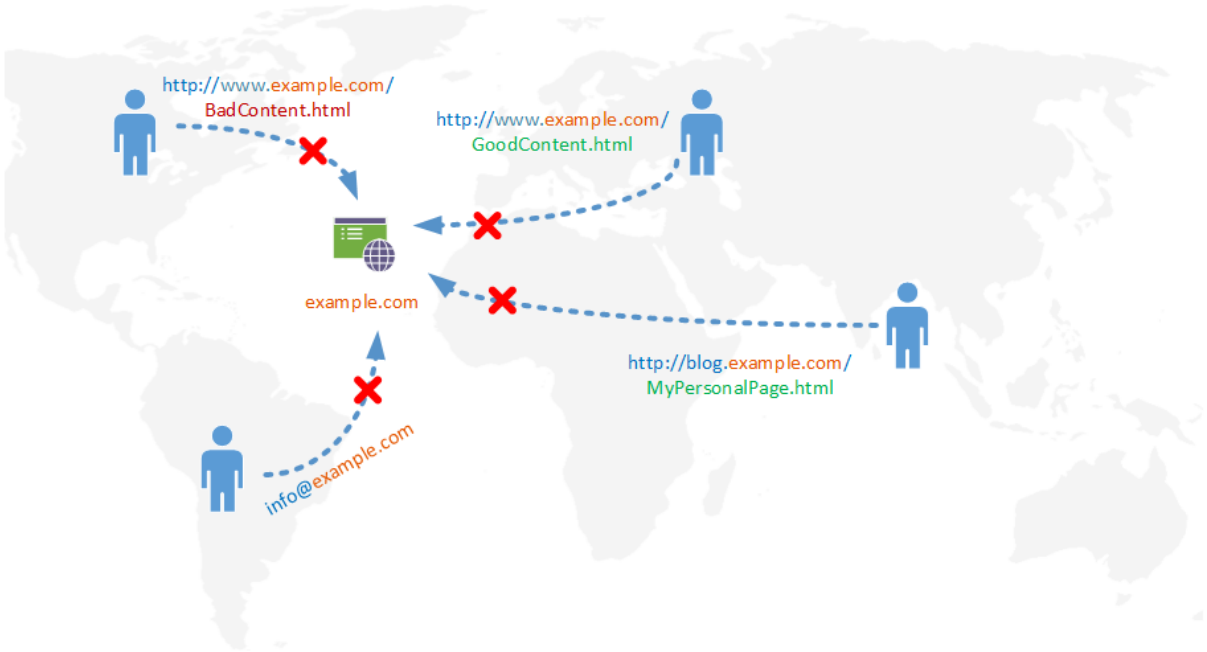

If only a specific content filtering has been requested (/BadContent.html) heavy overblocking collateral damage may be caused.

Efficacy

In this scenario, even if the target domain name has been seized at a global level, unwanted content is still online and can be accessed using other domain names or, in some cases, URLs containing the IP address of the server. For example, if iana.org was seized its contents would be accessible using the URL http://192.0.32.8/. Also every user can force his/her devices to resolve the seized domain name even if it has been blocked in the DNS chain: on most operating systems this can easily be accomplished by simply editing a configuration text file, adding a static entry line like “192.0.32.8 www.iana.org”.

Registering new domain names implies additional costs and time; furthermore, even if a new domain name was created, links and URLs that have already been distributed keep to refers the old one, vanishing any form of advertising or phishing technique made to attract users.

Security

This scenario is compatible with DNSSEC deployment provided that the Registry, when asked to remove or redirect a domain name, would also remove DNSSEC data. If that data is left unchanged, then any DNSSEC-aware application would fail, preventing any stop-page or seizure notice to be displayed too.

Feasibility

Due to the distributed nature of DNS hierarchy, name-servers operated by Registries may be out of the jurisdiction of the government which requests the block.

Costs and technical requirements to operate a domain removal from a Registry are minimal.

DNS Authoritative Servers

In this scenario requestors send a take-down notice to Authoritative Server manager in charge of the target domain.

Scope

Like in the case of Registries-level take-down, the domain name removal may have global impact on the whole Internet but only if it is implemented on all the Authoritative servers that handle the target domain name. In fact, for the sake of redundancy, a domain name may have many authoritative servers, spread around the world and also operated by different companies.
Moreover there is no way to predict which authoritative name-server will be used by a resolver to obtain resource IP addresses; some resolvers may use a server subject to their same jurisdiction, so they would block the undesired content, while others may use an Authoritative server out of the borders, which doesn't apply the blockade.

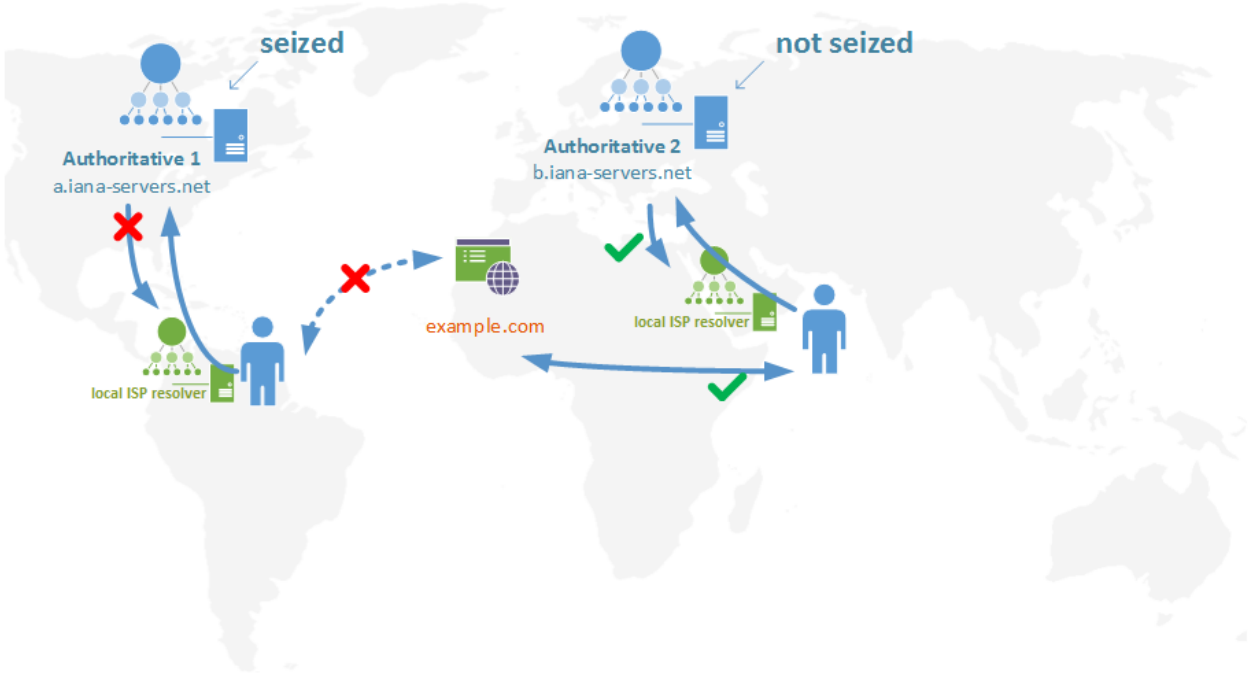

The scope of this blocking measure is therefore dependant on how many authoritative servers can be involved in its implementation: if all the authoritative servers are blocked it has a global scope, otherwise it leads to an Internet Balkanization with an unpredictable behaviour from the users perspective.

Granularity

As the filter is added to the servers which manage the domain zone it may be focused on just those services impacted by the blocking request, leaving the others fully operational. If a blocking request is made for www.example.com, only the "www" resource may be filtered out, leaving other services unchanged (info@example.com and blog.example.com/MyPersonalPage.html would continue to work).

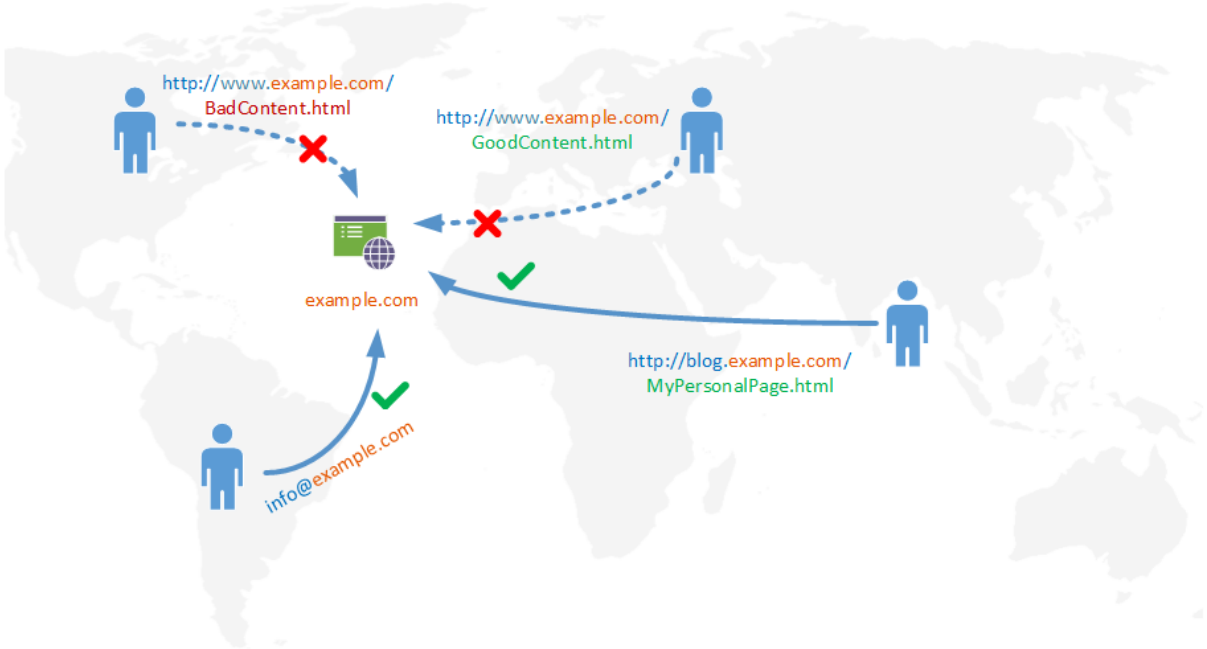

In this case overblocking collateral damage are limited to contents hosted on the same server: it would not be possible to block www.example.com/BadContent.html and keep www.example.com/GoodContent.html reachable.

Efficacy

The same considerations already seen for DNS Registries seizure efficacy apply to this scenario.

Security

Unlike the previous case, Authoritative name-server blockades may break the chain of trust of DNSSEC. Changes introduced by this blocking measure are the same that DNSSEC has been developed to identify and fight; DNSSEC capable devices or resolvers would consider any response from the Authoritative server as not valid and would drop it. Only in particular circumstances the digital signature can be kept valid, that is when the Authoritative server's provider also manage the cryptographic keys on behalf of the domain name owner.

Feasibility

As mentioned in the Scope section, a domain name may have many Authoritative name-servers spread across the world and also operated by different companies. A take-down order must be considered and applied by every company which provides this service in order to comply its objective, otherwise only Internet Balkanization and a patchy block would be achieved.
Also in this case costs and technical requirements are negligible.

ISP DNS Recursive Resolvers

Requestors order ISPs to block a target domain name on their recursive resolvers, optionally redirecting users to a "stop-page".

Scope

This method, that wants the take-down order to be sent to ISPs operating within the borders of the requesting government, is the most focused of those based on DNS. Its implementation takes place on the recursive resolvers that ISPs offer to their customers along with Internet access contracts.

Since the order is only sent to companies registered under the jurisdiction of the requesting government, it is expected to be implemented on those access networks that are only used by users also subject to the same laws, with no impacts on other parts of Internet.

In order to achieve a country-wide effect orders must be sent to every ISP which operates an access network on the territory, otherwise a non-homogeneous and discriminatory treatment would be reserved to people.

Granularity

Like the other DNS-based blocking methods, this one imposes heavy-impacting overblocking collateral damage too. Even though it may be applied against a specific sub-domain or resource only (www, blog, …), every content published under that name space will be filtered out: both www.example.com/BadContent.html and www.example.com/GoodContent.html will be cutted off from the Internet for the users of the ordered ISPs.

Efficacy

With regards of this blocking method, in addition to considerations already made for the previous two DNS-based solutions, that are still valid for this scenario, another issue must be considered. As already mentioned, recursive resolvers are offered by ISPs along with Internet access contracts, but also by many other providers and entities on Internet (Google, OpenDNS, …), even for free. Users wanting to keep access to forbidden contents can simply change the configuration of their devices and set the parameters of a resolver out of the jurisdiction of the requesting government.

Security

One kind of implementation of this measure requires the client device to be redirected to a remediation server, which displays a message about the take-down order notice. Since unwanted contents may be on any form and may be published using every protocol (HTTP for web pages, but also FTP for file storage, Telnet or SSH for access to servers management software, SMTP for open relay servers, …) the remediation server also must implement every Internet protocol and it must be ready to provide the message in any form, otherwise a connection error would be detected by the user's client application.

Another kind of implementation is based on response codes provided by resolvers to client devices; many of these response codes may be interpreted by clients as a server malfunctioning and may lead operating systems to remove the resolver from the list of those to be used. The iteration of this behaviour for every configured resolver could lead to all the resolvers to be marked as not working, rendering the user's device unable to resolve Internet names and addresses.

DNSSEC also is highly impacted by resolver level blockade. As already stated, all the changes introduced along the DNS hierarchy which have not a corresponding digital signature are considered invalid and untrustable by DNSSEC. Since resolvers don't have access to the cryptographic key used to sign the domain zone, they can't rebuild the digital signature and are forced to break the chain of trust. Moreover, when the target domain is protected by DNSSEC end users can't be redirected to the remediation server which shows them the stop-page containing the take-down order notice;

they only receive an error message and may be urged to call ISP help-desk, pushing costs to 3rd parties and impairing service experience.

Feasibility

The take-down notice has to be notified only to ISPs operating within the borders of the requesting government, so no governments cooperation is required to reach all the involved companies.
Costs and technical requirements also are negligible for this implementation.

ISP IP address block

Requestors order ISPs to block a specific IP address on their network; packets sent to and received from the target IP address are blocked and never forwarded to destination, preventing any kind of communication.

Scope

Such as any filter implemented at ISP level this method allows to block only those users subject to the same jurisdiction of the requesting government, with no impacts on other parts of Internet. Like Recursive Resolvers blocking method, it requires every ISP operating on the territory to be ordered to block the IP address to avoid a blockade that only apply to a portion of users.

Granularity

With regards of overblocking collateral damage, this is the method with the highest possible impact on contents and services. Since an IP address may be used to host more than one web site or resources with different names[4], whenever a take-down order is composed by a "domain name + mapped IP address" bundle it causes the shut down not only for the target domain, but also for every domain name hosted on that same IP address. This is the case, for example, of Content Providers which offer hosting solutions to their customers by providing them with domain name registration and web space storage where to publish their personal pages, blogs or projects. In such case the whole hosting platform will be cutted off and any web site will be prevented to be visited by customers of the notified ISP.

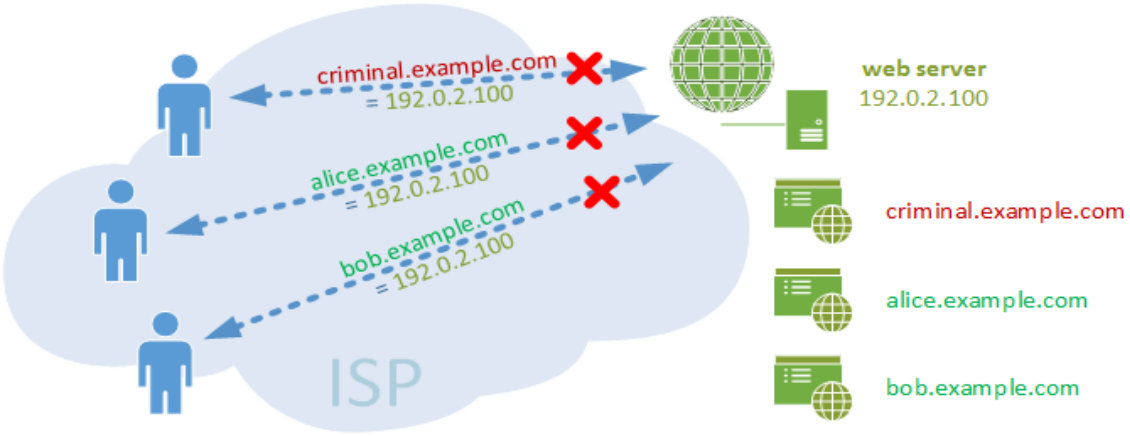

[4] "For the COM, NET and ORG top-level domains [...] the average number of website instances sharing a single IP address is 7.5" from "Study into Websites Sharing Internet Protocol Addresses", Ofcom: http://stakeholders.ofcom.org.uk/binaries/internet/websites-sharing.pdf

An opposite effect exists too: sometimes, the same content is distributed on many servers with different IP addresses for the sake of load-balancing or to benefit of caching mechanisms. In such case, to obtain a global coverage of the blocking measure every server involved in the content distribution must be identified and added to the list of those IP addresses to block.

Efficacy

Even if contents that are kept online on a server can't be reached anymore, nor if users try to change their DNS resolvers nor if they try to setup static entries in configuration file, countermeasures exist for this kind of method too. Users can circumvent the blockade by routing their traffic away from the blocking enforcement, by bouncing on networks which are out from the jurisdictional scope of the requestor.

Many companies and entities spread over the Internet offer free anonymous web proxy servers hosted on "offshore" platforms, which may be easily configured on devices; instead of traversing the ISP's network on the route toward the blocked IP address, packets are sent to these proxies which act as intermediaries and keep communications with the blocked server on behalf of end users' devices. The large number of anonymous web proxies present on Internet, their ease of setup and shutdown and the volatility of their IP addresses make impossible to block this kind of servers and prevent them to be used.

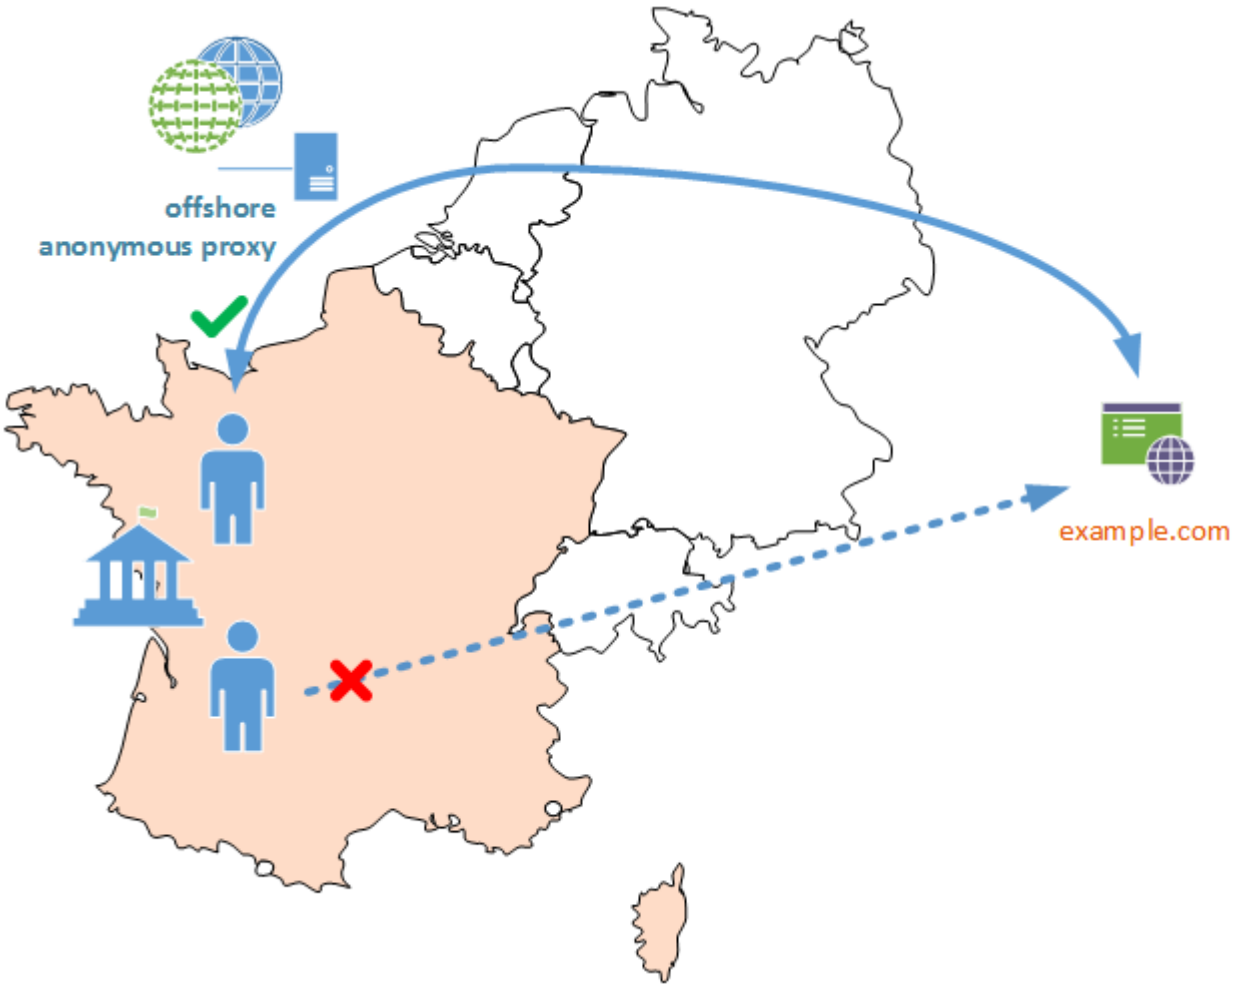

As well as anonymous web proxies, also VPNs (Virtual Private Networks) may be used to bypass IP blocks. Like web proxies, VPNs allow devices to route any traffic toward a focal point which, in turn, forward it to the original destination. They offer more stability, security and privacy than web proxies and they can handle any protocol (while proxies are usually only for web); while normally they are not free and require a paying account to be used, many commercial self-interested websites provide them to their customers to let them to bypass blocks and access their contents (online casino, gambling, …).
Furthermore, contents publishers can change the IP address of their server and update the domain zone which refers to it, bringing the content back online in a few hours.

Security

While DNSSEC is not involved in this blocking method and no direct impacts exist on the security of Internet, many other side effects exist, which will be covered in the rest of the document.

Feasibility

Even if the blockade is spread along the whole ISP network and it's not focused on a rendezvous point like DNS Recursive Resolvers, on most cases it may be not so difficult to be implemented. Most ISPs already have core routing points where all the users network traffic is concentrated, many have mechanisms used to prevent network attacks which may be used to block IP addresses too, so in many cases costs and deployment difficulties may be considered to be low.
Because only companies operating within the borders must be involved, no governments cooperation is required.

ISP Web Proxies

Requestors order ISPs to filter any web request made from their customers and to block those involved in take-down orders.

Scope

The same considerations already seen for ISP IP address blockade apply to this scenario.

Granularity

Proxies bring one of the most specific filtering method that can be implemented on a network, even if they are quite totally focused on HTTP/HTTPS traffic. Any web request is diverted to a proxy which analyzes the requested URL and then decides to handle it or to drop it, on the basis of lists provided by requesting parties in take-down orders. Discrimination may be specific and may reach the resource content level, distinguishing specific web pages from others on the same domain name too (www.example.com/BadContent.html may be blocked while www.example.com/GoodContent.html may be allowed).

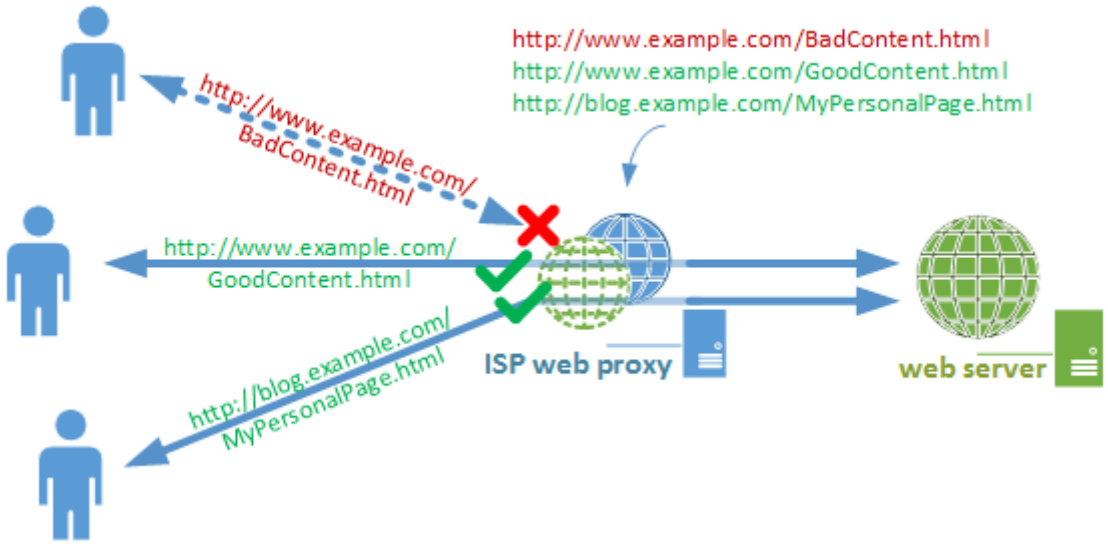

Efficacy

Users' devices must be configured to use web proxies provided by their ISPs. Blocking enforcements must be placed by ISPs along their network in order to prevent users from bypassing the usage of these web proxies and eventually redirect web traffic toward them.

Even if all the web traffic may be constrained through web proxies, other protocols can not and have to be kept open, unless a complete Internet principles distortion would be provided: anonymous web proxies and VPNs (already seen in the previous method considerations) may be still used.

VPNs traffic is not subject to web rules and it's not forced to go through the ISP proxies so VPNs may be deployed in order to bypass filters. Furthermore many anonymous web proxies implement camouflage mechanisms which allow them to be used even if web blocking enforcements exist on the ISP network.

Security

Most of the issues regarding web proxies security are about encrypted communications over HTTPS. In order to establish which policy to use regarding a web request a proxy must have access to some information; encryption, often used to add security and confidentiality to sensitive communications such as e-commerce, home banking, e-government, alters these data in a way that only endpoints can access them. To obtain the required information web proxies must impersonate the endpoint of the communication and decrypt the content, by breaking the authentication of the protocol and introducing a security weakness along the path that users consider fully trusted.

Also, modern browsers and applications implement new protocols developed just to detect third-party entities which impersonate secure communications endpoints; these protocols are focused on detecting so called "man in the middle" attacks, used by criminals to break secure path and gather sensitive data and passwords. When such an event occurs, the browser blocks any communication and acts as if the service is blocked, preventing any website which implements these new protocols from being accessed, even if not involved in take-down orders.

Feasibility

Even if it does not require cross-borders cooperation (like other ISPs based solutions) this method requires big efforts to be implemented. Many networking equipments are conceived to switch packets only on the basis of their source and destination IP addresses and they suffer huge performance deteriorations when used to analyze other information, even if only deployed in blocking enforcement points to distinguish web traffic from other protocols and route it toward web proxies. Additional hardware and software must be deployed and maintained to run web proxies too; depending on the ISP's network size and topology, more than one group of proxies may be needed in order to assure load-balancing, traffic optimization, scalability and fault-tolerance. Furthermore, every ISP existing network has to be redesigned and engineered on the basis of these new functional elements.

With regards to the depth of the analysis needed to discriminate the content of communications, issues may rise about privacy and human rights too.

ISP Deep Packet Inspection

Like the previous method, requestors order ISPs to monitor content of every packet on their network and block those packets that matches a specific protocol data pattern.

Scope

Same considerations already seen for ISP IP address blockade and web proxies apply to this method.

Granularity

Like web proxies, Deep Packet Inspection (DPI) software can analyze the content of packets while they traverse the ISP network and decide which may keep going and which have to be dropped, on the basis of pieces of data (called "signatures") that are particular to various protocols. While web proxies only analyze web requests (HTTP, HTTPS), DPI appliances can use a wider range of policies to classify (and eventually block) communications on various protocols (email, VoIP, file sharing, …). This method of analysis allows DPI to be a very specific solution and to block only those contents which are the real targets of take-down orders.

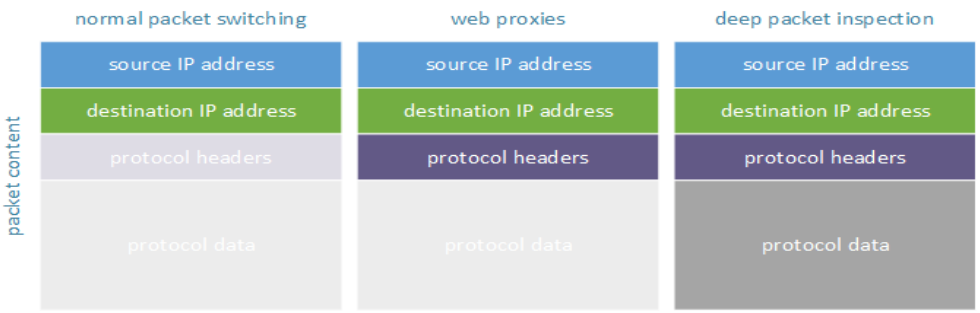

Efficacy

Many DPI software can handle obfuscated and encrypted protocols, reducing the chances of circumvention by using offshore proxies or VPNs. Efficacy of this method is closely bound up with obfuscation and encryption techniques which are over and over improved.

Security

No direct impacts are introduced by this method on the security of Internet, but many side effects exist, which will be covered in the rest of the document

Feasibility

Like other ISP-based solution, this method does not require governments cooperation because it involves only companies operating within the same jurisdiction borders of the requesting one. Costs and difficulties introduced by this filtering technique are very high. DPI solutions are expensive and, usually, they can't handle high loads of traffic, so many inspection points have to be spread along the ISP network; moreover, existing networks have to be deeply changed in order to house them.
In order to identify the traffic to block DPI signatures must be very focused and targeted to match a specific protocol data pattern; a heavy effort is needed to code them and to keep them updated and aligned with the ever increasing number of applications and protocols developed on Internet. LEAs may be in need of hiring security experts to develop signatures and to build a strong coordination protocol with ISPs in order to maintain them up to date.
Legal questions must be evaluated about impact on users' privacy, data protection and human rights.

Side effects

Blocking and filtering measures may have some side effects on the everyday usage of Internet. Side effects are not directly due by protocol breakages or data corruption but are caused by abnormal and unusual behaviours of users and by misuse of some tools. If users know that the contents they want to access are still online but to access them something has to be changed on their devices, it is likely that these changes develop quickly on large scale. Users may be urged to adopt out of borders open resolvers or proxy servers to access forbidden contents or to reach resources which have been over-blocked by an improper blocking measure. These behaviours, which all lead to exposure of users to threats, are expected to grow with the growth of the number of over-blocked resources.

Extended trust on automatic configuration script

Even if changing the device configuration is not a difficult operation, some users may not be able to accomplish it successfully on their own. These users may be encouraged by sites operators to use external automatic scripts to achieve that goal, by granting to these programs maximum privileges on their operating systems and allowing them to change the configuration on their behalf. The content providers which want the users to change their configuration may be reliable entities or unreliable ones, with many ulterior (fraudulent) motives; once executed, these programs allow them to fully control a user's device.
Extending the trust of inexpert users on external scripts and helper programs may break basic security principles.

Use of untrusted resolvers and proxies

When users are persuaded to change their DNS resolvers or to use an anonymous web proxy by a blocked website, changes impact every Internet activity they perform, not only that related to the filtered website. Take the case of an online casino which suggests users to change their resolver and use the one it provides; the new server will certainly handle queries for the casino website in the right way, but it may also redirects queries for an home-banking website toward a fake server and let criminal entities to gather accounts credentials. The same would happens for anonymous proxies.

Defeat of anti-cybercrime activities

DNS blocking techniques may be used to defeat cybercrime too, by blocking those domain names which are dedicated to frauds, phishing or malware distribution (viruses, trojans, …). If users decide to change their device configuration and use public open resolvers to access (over-) blocked content any local anti-cybercrime activity is vanished.

Impacts on Content delivery networks

Some Content Providers distribute their contents on more servers around the world, so that the same resource may be accessed through servers that are closer to the user than other; this is done to improve performances and user experience. Some content delivery networks base their decisions about which server to redirect user to on the DNS resolver used by that user to resolve domain name. By using open resolvers outside the own country may lead this technique to fail and to bring to an increased latency and worst performances than usual.

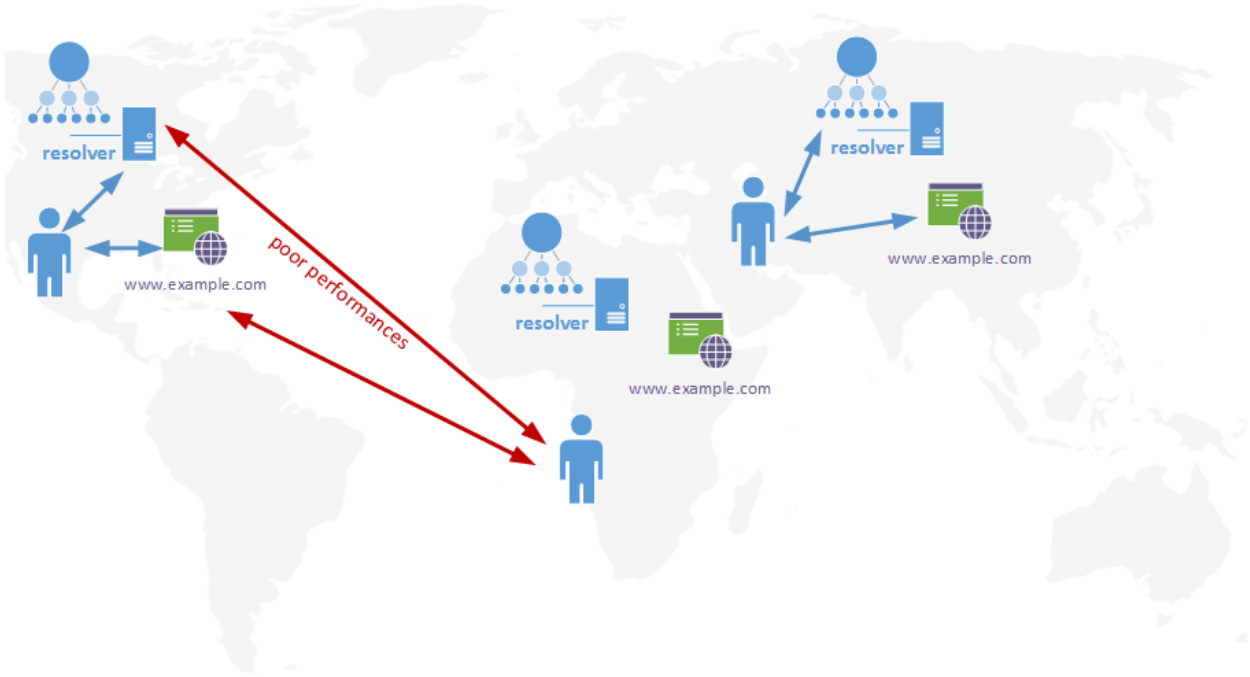

Conclusions

Blocking methods may lead to different results whenever used for global domain seizure or for content filtering purposes.

ISP-focused enforcements have always a better scope than other, because they are closer to the users which must be prevented to access the undesired contents.
For what concerns granularity, DNS based blockades are good when a global domain seizure is requested but represent a very poor solution when used for content filtering purposes.
Efficacy is a critical point of every blocking measure; traffic redirection and encryption allow every enforcement to be circumnavigated, even with little effort.
Internet security is compromised by most of the methods, with major impacts when considering side effects and human rights violation too.
Feasibility is worst for more effective methods and better for less ones, mostly because of government cooperation needs and additional costs for hardware and software.

Method	Domain seizure				
	Scope	Granularity	Efficacy	Security	Feasibility
DNS Registries	√ √	√ √ √	√ √	√ √	√
DNS Authoritative Servers	√	√ √ √	√ √	√ √	√ √
ISP DNS Recursive Resolvers	√ √ √	√ √ √	√	√	√ √ √
ISP IP address block	√ √ √	√	√ √	√ √ √	√ √ √
ISP Web Proxies	√ √ √	√ √ √	√ √	√ √	√ √
ISP Deep Packet Inspection	√ √ √	√ √ √	√ √	√ √ √	√

Method	Content filtering				
	Scope	Granularity	Efficacy	Security	Feasibility
DNS Registries	√ √	√	√ √	√ √	√
DNS Authoritative Servers	√	√ √	√ √	√	√ √
ISP DNS Recursive Resolvers	√ √ √	√ √	√	√	√ √ √
ISP IP address block	√ √ √	√	√ √	√ √ √	√ √ √
ISP Web Proxies	√ √ √	√ √ √	√ √	√ √	√ √
ISP Deep Packet Inspection	√ √ √	√ √ √	√ √	√ √ √	√

(√ √ √ = best; √ = worst)

Efforts spent for years by ISPs and service providers to educate users about good practices and safe behaviors may be vanished by risky operations spread to bypass (improper) enforcements. Moreover, the growth of new protocols developed to strengthen Internet security risks to be impaired or delayed.

Further Reading

- "SAC 056 - SSAC Advisory on Impacts of Content Blocking via the Domain Name System", ICANN: http://www.icann.org/en/groups/ssac/documents/sac-056-en.pdf
- "Consequences of DNS-based Internet filtering", AFNIC: http://www.afnic.fr/medias/documents/conseilscientifique/SC-consequences-of-DNS-based-Internet-filtering.pdf
- "Technical Considerations for Internet Service Blocking and Filtering", IETF: http://tools.ietf.org/html/draft-iab-filtering-considerations-04
- "Site Blocking to reduce online copyright infringement", Ofcom: http://stakeholders.ofcom.org.uk/binaries/foi/2011/october/1-186872101-attachment1.pdf
- "Internet Society Perspectives on Domain Name System (DNS) Filtering", Internet Society: http://www.internetsociety.org/sites/default/files/pdf/dns-filtering_20110915.pdf

Author's address

Pier Carlo Chiodi
pierky@pierky.com / pc.chiodi@gmail.com
http://pierky.com/aboutme