The Politics of Names

- Choice of DNS resolver
- Intellectual property and the DNS
- What domains should exist?
- Who controls the DNS root?

Colin Perkins | https://csperkins.org/ | Copyright © 2020 University of Glasgow

Implications of Choice of DNS Resolver (1/4)

- **How is the DNS resolver chosen?**
- When connecting to network, hosts use DHCP (dynamic host configuration protocol) to discover network settings and configuration
  - DHCP tells the host what DNS resolver to use for the network
  - If a host has multiple network interfaces, it may use a different DNS resolver for each
    - The `getaddrinfo()` call takes a `hints` parameter, that can include the local IP address
    - e.g., consider a device connected to 4G cellular network and a private company Ethernet – the Ethernet might make available names of internal services that aren't accessible to the public
- Possible to configure the DNS resolver manually
  - e.g., to talk to Google public DNS resolver (IPv4 address 8.8.8.8)

Colin Perkins | https://csperkins.org/ | Copyright © 2020 University of Glasgow

Implications of Choice of DNS Resolver (2/4)

- DNS resolution has typically been a system-wide service
  - Operating system implements a DNS resolution service; all DNS queries use that service
  - A consistent mapping of names to addresses
- **DoH is changing this**
  - JavaScript web applications can now easily perform DNS queries via any HTTP website
  - Each application may get different answers for the same query, depending on the server; it's no longer easily possible to enforce policy via the DNS

Colin Perkins | https://csperkins.org/ | Copyright © 2020 University of Glasgow

Implications of Choice of DNS Resolver (3/4)

- Giving applications ability to securely access arbitrary DNS servers allows them to avoid local observation and/or filtering of DNS traffic
- Is this flexibility for each application to perform DNS queries differently a concern?

	No	**Yes**
	Applications should have the ability to use a secure DNS server they trust to avoid phishing attacks, malware, monitoring, etc.	Network operators filter DNS responses to block access to malicious sites and prevent malware spreading – allowing applications to bypass this is a security risk
	Why should network operators be able to see DNS queries and modify responses? This is a privacy and security risk	Network operators filter DNS responses to enforce legal or societal constraints – e.g., Internet Watch Foundation DNS block list to stop UK-based access to sites hosting child sexual abuse material

Colin Perkins | https://csperkins.org/ | Copyright © 2020 University of Glasgow

Implications of Choice of DNS Resolver (4/4)

- **Can a network restrict the choice of DNS resolver?**
- Firewalls can block access to DNS-over-UDP and DNS-over-TLS resolvers
  - Block access to UDP port 53
  - Block access to TCP port 853
  - For all destination IP addresses except those of allowed DNS resolvers
- Difficult to block DNS-over-HTTPS
  - Network cannot distinguish DNS-over-HTTPS from any other traffic over HTTPS
  - May be able to tell from the destination IP address
    - e.g., Google use IPv4 address 8.8.8.8 for public DoH services, but not other traffic
  - But if a web server handles a mix of web and DNS traffic over HTTPS, cannot block one without blocking the other
  - Many ISPs and governments concerned that DoH prevents use of DNS as a control point

Colin Perkins | https://csperkins.org/ | Copyright © 2020 University of Glasgow

Intellectual Property and the DNS

- Intellectual property is managed on a national basis
  - e.g., a company might own a trademark in the UK while a different company owns the same trademark in the Republic of Ireland
  - Which of those companies owns *`trademark`*`.ie` and which owns *`trademark`*`.co.uk` is a straightforward legal question
  - Which of those companies owns *`trademark`*`.com` is likely harder to decide
    - Especially since `.com` is operated by a US-based organisation, and a third company might own the trademark in the US
- A ccTLD clearly operates under legal regime of a particular country
- Use of a gTLD might lead to legal complications

Colin Perkins | https://csperkins.org/ | Copyright © 2020 University of Glasgow

What Domains Should Exist?

- **Should a particular gTLD be allowed to exist?**

- For example, should `.xxx` exist to host “adult” content?
  - If so, who gets to decide what content should (must?) sit within that gTLD?
  - Different countries have very different norms and standards in this area
- Who controls what TLDs ICANN permits? Who should control it?

Colin Perkins | https://csperkins.org/ | Copyright © 2020 University of Glasgow

What Domains Should Exist?

- **Should a particular subdomain be allowed to exist?**

- Significant differences around freedom of speech and permissible topics in different parts of the world
- A ccTLD can enforce local conventions and rules
- What rules apply to a gTLD?
  - If a particular country/group finds a site objectionable, should it be taken down?
  - If country X decides particular content is illegal, but it is legal in country Y, should a gTLD operated out of country Y but accessible in country X permit such content?
  - e.g., Holocaust denial is illegal in Germany but not in the US – should `.com`, operating from the US, permit sites hosting such content?

Colin Perkins | https://csperkins.org/ | Copyright © 2020 University of Glasgow

Who Controls the Root Servers?

- DNS root servers are mostly controlled by US-based organisations
- Is this a risk for other countries?
- Should the root servers be controlled by a broader mix of countries?
  - If so, who gets to decide – ICANN? The UN?
  - Is there a benefit in controlling a DNS root server?
  - Is there a benefit in controlling a gTLD server? Who gets to host `.com`, for example?

Server	IPv4 Address	IPv6 Address	Operator
A	198.41.0.4	2001:503:ba3e::2:30	Verisign
B	199.9.14.201	2001:500:200::b	USC-ISI
C	192.33.4.12	2001:500:2::c	Cogent Communications
D	199.7.91.13	2001:500:2d::d	University of Maryland
E	192.203.230.10	2001:500:a8::e	NASA Ames Research Center
F	192.5.5.241	2001:500:2f::f	Internet Systems Consortium
G	192.112.36.4	2001:500:12::d0d	US Defense Information Systems Agency
H	198.97.190.53	2001:500:1::53	US Army Research Lab
I	192.36.148.17	2001:7fe::53	Netnod
J	192.58.128.30	2001:503:c27::2:30	Verisign
K	193.0.14.129	2001:7fd::1	RIPE NCC
L	199.7.83.42	2001:500:9f::42	ICANN
M	202.12.27.33	2001:dc3::35	WIDE Project

Colin Perkins | https://csperkins.org/ | Copyright © 2020 University of Glasgow

Should There Be a Single DNS Root?

- Should all TLDs be accessible from everywhere?
  - Should there be a single global DNS?
  - Should the same name always resolve to the same site?
    - With global content distribution networks, how can you tell?
  - Should different countries be allowed to filter DNS?
    - If so, how should such restrictions be implemented?
    - It is difficult to distinguish modifications to DNS responses made to conform to government-mandated filtering requirements from those made by malware, phishing attacks, etc. – is this a feature or a bug?

Colin Perkins | https://csperkins.org/ | Copyright © 2020 University of Glasgow

Naming and the Tussle for Control

- What is the DNS?
- How are DNS queries made?
- Who controls the names?

CC BY ND