University of Glasgow

Internet Protocols Laboratory

Naming and the Tussle for Control

Networked Systems (H)

Lecture 8

Colin Perkins, Elizabeth Boswell | https://csperkins.org/ | Copyright © 2026 University of Glasgow | This work is licensed under the Creative Commons Attribution-NoDerivatives 4.0 International License. To view a copy of this license, visit http://creativecommons.org/licenses/by-nd/4.0/ or send a letter to Creative Commons, PO Box 1866, Mountain View, CA 94042,

Lecture Outline

- What is the DNS and DNS resolution
- DNS naming structure and organisation
- Methods for DNS resolution
- The politics of names

- Associated reading: “Tussle in cyberspace: Defining tomorrow’s Internet” discusses the issues of control over the network, and how protocol design influences this – the DNS is one area where we see this tussle

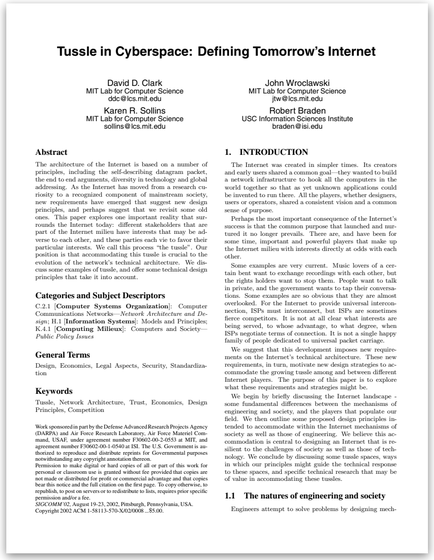

Tussle in Cyberspace: Defining Tomorrow’s Internet

David D. Clark
MIT Lab for Computer Science
ddc@lcs.mit.edu

John Wroclawski
MIT Lab for Computer Science
jtw@lcs.mit.edu

Karen R. Sollins
MIT Lab for Computer Science
sollins@lcs.mit.edu

Robert Braden
USC Information Sciences Institute
braden@isi.edu

Abstract

The architecture of the Internet is based on a number of principles, including the self-describing datagram packet, the end to end arguments, diversity in technology and global addressing. As the Internet has moved from a research curiosity to a recognized component of mainstream society, new requirements have emerged that suggest new design principles, and perhaps suggest that we revisit some old ones. This paper explores one important reality that surrounds the Internet today: different stakeholders that are part of the Internet milieu have interests that may be adverse to each other, and these parties each vie to favor their particular interests. We call this process “the tussle”. Our position is that accommodating this tussle is crucial to the evolution of the network’s technical architecture. We discuss some examples of tussle, and offer some technical design principles that take it into account.

Categories and Subject Descriptors

C.2.1 [**Computer Systems Organization**]: Computer Communications Networks—*Network Architecture and Design*; H.1 [**Information Systems**]: Models and Principles; K.4.1 [**Computing Milieux**]: Computers and Society—*Public Policy Issues*

General Terms

Design, Economics, Legal Aspects, Security, Standardization

Keywords

Tussle, Network Architecture, Trust, Economics, Design Principles, Competition

Work sponsored in part by the Defense Advanced Research Projects Agency (DARPA) and Air Force Research Laboratory, Air Force Materiel Command, USAF, under agreement number F30602-00-2-0551 at MIT, and agreement number F30602-00-1-0540 at ISI. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon.
Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
SIGCOMM'02, August 19-23, 2002, Pittsburgh, Pennsylvania, USA.
Copyright 2002 ACM 1-58113-570-X/02/0008 ...$5.00.

1. INTRODUCTION

The Internet was created in simpler times. Its creators and early users shared a common goal—they wanted to build a network infrastructure to hook all the computers in the world together so that as yet unknown applications could be invented to run there. All the players, whether designers, users or operators, shared a consistent vision and a common sense of purpose.

Perhaps the most important consequence of the Internet’s success is that the common purpose that launched and nurtured it no longer prevails. There are, and have been for some time, important and powerful players that make up the Internet milieu with interests directly at odds with each other.

Some examples are very current. Music lovers of a certain bent want to exchange recordings with each other, but the rights holders want to stop them. People want to talk in private, and the government wants to tap their conversations. Some examples are so obvious that they are almost overlooked. For the Internet to provide universal interconnection, ISPs must interconnect, but ISPs are sometimes fierce competitors. It is not at all clear what interests are being served, to whose advantage, to what degree, when ISPs negotiate terms of connection. It is not a single happy family of people dedicated to universal packet carriage.

We suggest that this development imposes new requirements on the Internet’s technical architecture. These new requirements, in turn, motivate new design strategies to accommodate the growing tussle among and between different Internet players. The purpose of this paper is to explore what these requirements and strategies might be.

We begin by briefly discussing the Internet landscape - some fundamental differences between the mechanisms of engineering and society, and the players that populate our field. We then outline some proposed design principles intended to accommodate within the Internet mechanisms of society as well as those of engineering. We believe this accommodation is central to designing an Internet that is resilient to the challenges of society as well as those of technology. We conclude by discussing some tussle spaces, ways in which our principles might guide the technical response to these spaces, and specific technical research that may be of value in accommodating these tussles.

1.1 The natures of engineering and society

Engineers attempt to solve problems by designing mech-

D. D. Clark, J. Wroclawski, K. R. Sollins, and R. Braden, “Tussle in cyberspace: Defining tomorrow’s Internet”, ACM SIGCOMM Conference, August 2002. https://dx.doi.org/10.1145/633025.633059

Colin Perkins, Elizabeth Boswell | https://csperkins.org/ | Copyright © 2026 University of Glasgow

University of Glasgow

DNS Name Resolution

- What is the DNS?
- Structure of names
- Name resolution

CC BY ND

What is the DNS? (1/2)

- IP packets contain addresses rather than names
  - Designed for efficient processing by routers determining where to forward the packet
  - Not human readable – people prefer names, not addresses
- Domain name system (DNS) is a distributed database; maps names to IP addresses

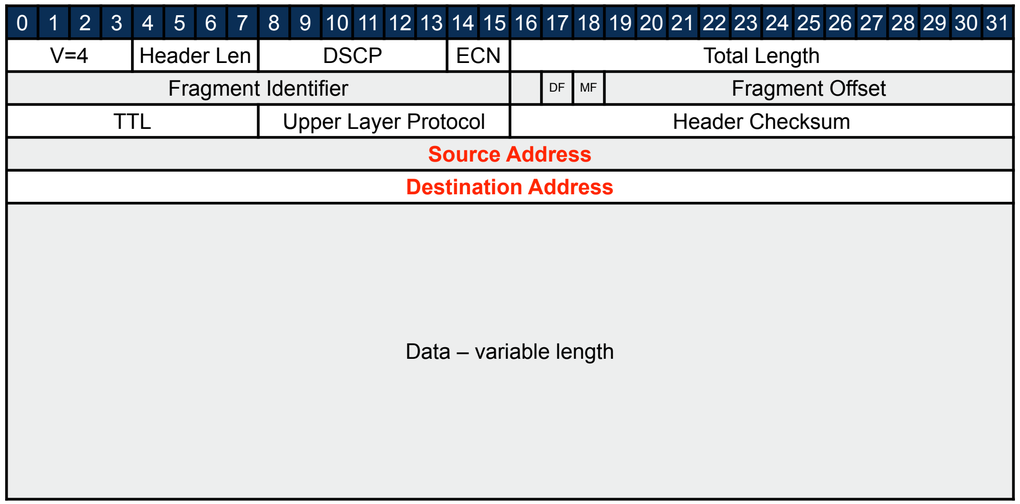

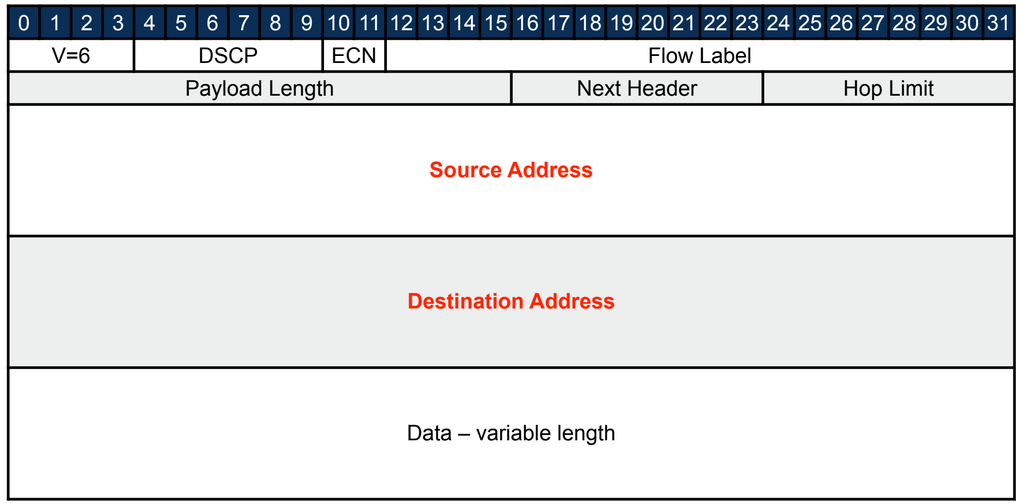

Colin Perkins, Elizabeth Boswell | https://csperkins.org/ | Copyright © 2026 University of Glasgow

What is the DNS? (2/2)

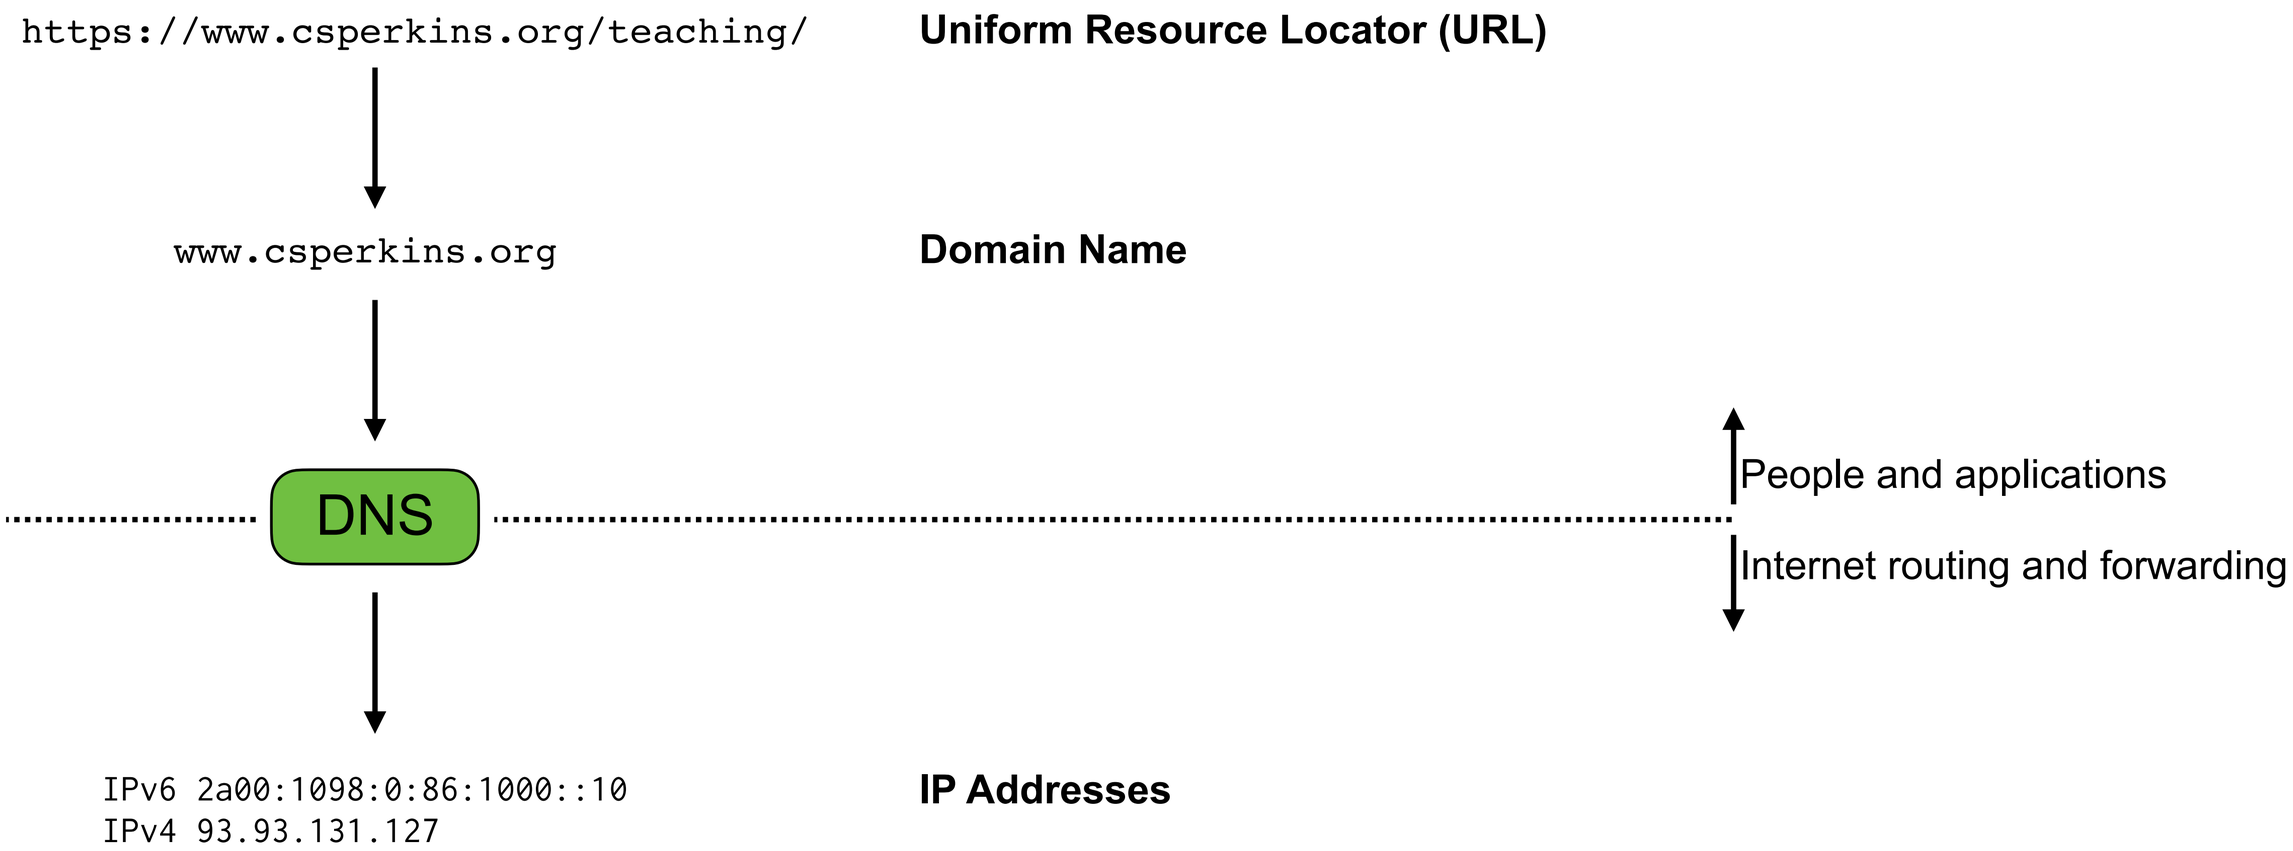

Colin Perkins, Elizabeth Boswell | https://csperkins.org/ | Copyright © 2026 University of Glasgow

Structure of DNS Names (1/2)

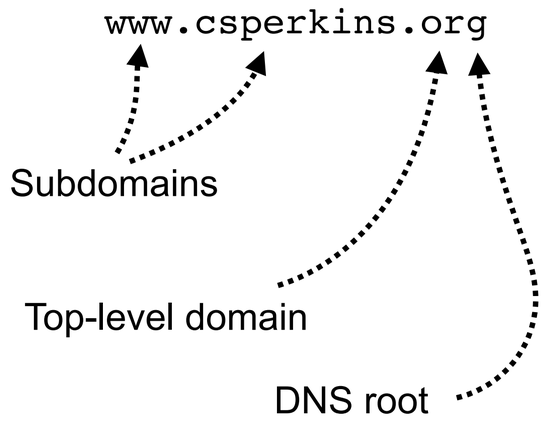

- Domain names consist of **labels** separated by dots
- Naming is hierarchical
  - Rightmost (top) label is **top-level domain** (TLD)
    - Country-code top-level domains (ccTLDs)
      - `.uk`, `.de`, `.cn`, `.io`, `.ly`, …
    - Generic top-level domains (gTLDs)
      - `.com`, `.org`, `.net`, …
  - Preceded by any number of sub-domains

Colin Perkins, Elizabeth Boswell | https://csperkins.org/ | Copyright © 2026 University of Glasgow

Structure of DNS Names (2/2)

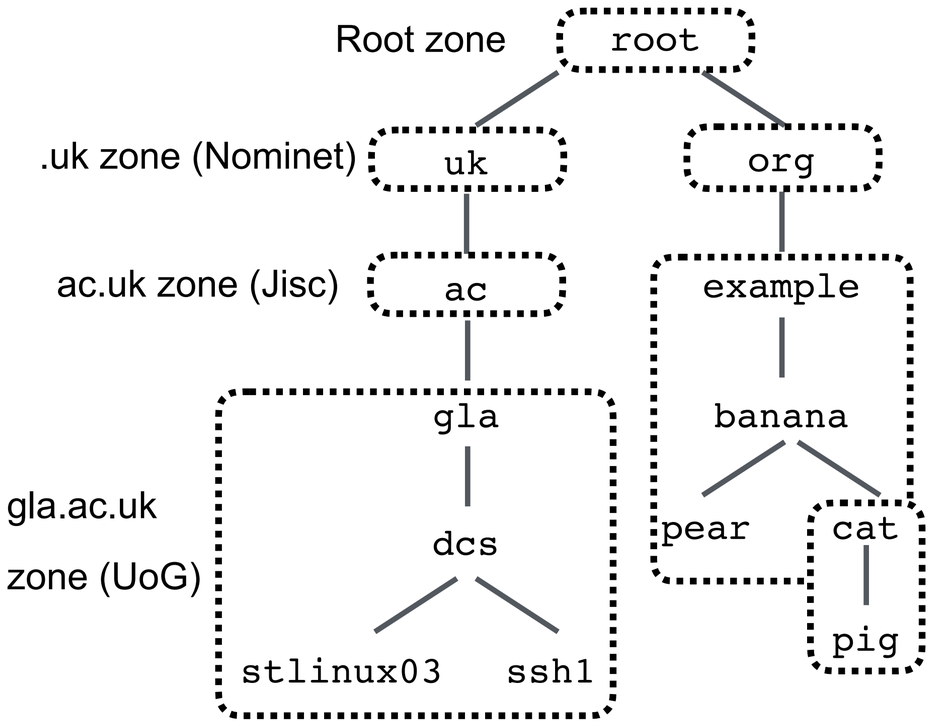

- DNS is a distributed database - in authority and implementation
  - Domain names form a tree structure
  - Split into subtrees (called **zones**), which are are administered independently
- Information associated with each zone are served by the zone's **nameservers**
  - IP addresses etc of domain names in zone
  - names of nameservers of zone below
- Top-level domains live within the **DNS root zone**
  - The **root servers** advertise the top-level domains
  - They have well-known, fixed, IP addresses – new DNS resolvers need to reach them to find the TLDs before they can answer DNS queries

Colin Perkins, Elizabeth Boswell | https://csperkins.org/ | Copyright © 2026 University of Glasgow

DNS Name Resolution (1/3)

- The DNS is used for **name resolution**
  - Given a name, lookup a particular type of **record** relating to that name
  - Many different types of record: `A`, `AAAA`, `CNAME`, `MX`, `NS`, `SRV`, …
  - Most common are `A` and `AAAA` records, that map hostnames to IPv4 and IPv6 addresses and `NS` records that identify the **nameserver** for a domain
- A DNS **client** asks its **resolver** to perform the lookup by calling `getaddrinfo()`

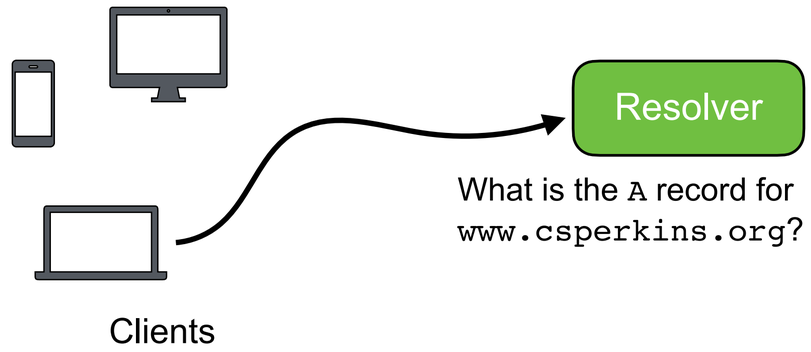

Colin Perkins, Elizabeth Boswell | https://csperkins.org/ | Copyright © 2026 University of Glasgow

DNS Name Resolution (2/3)

- If the resolver has no information, it makes a recursive query
- Starts by querying the root
  - Doesn't have the answer, instead refers the resolver to the TLD nameservers
  - Query the TLD nameservers, get referred to lower level nameservers, etc
  - Finally receive the answer

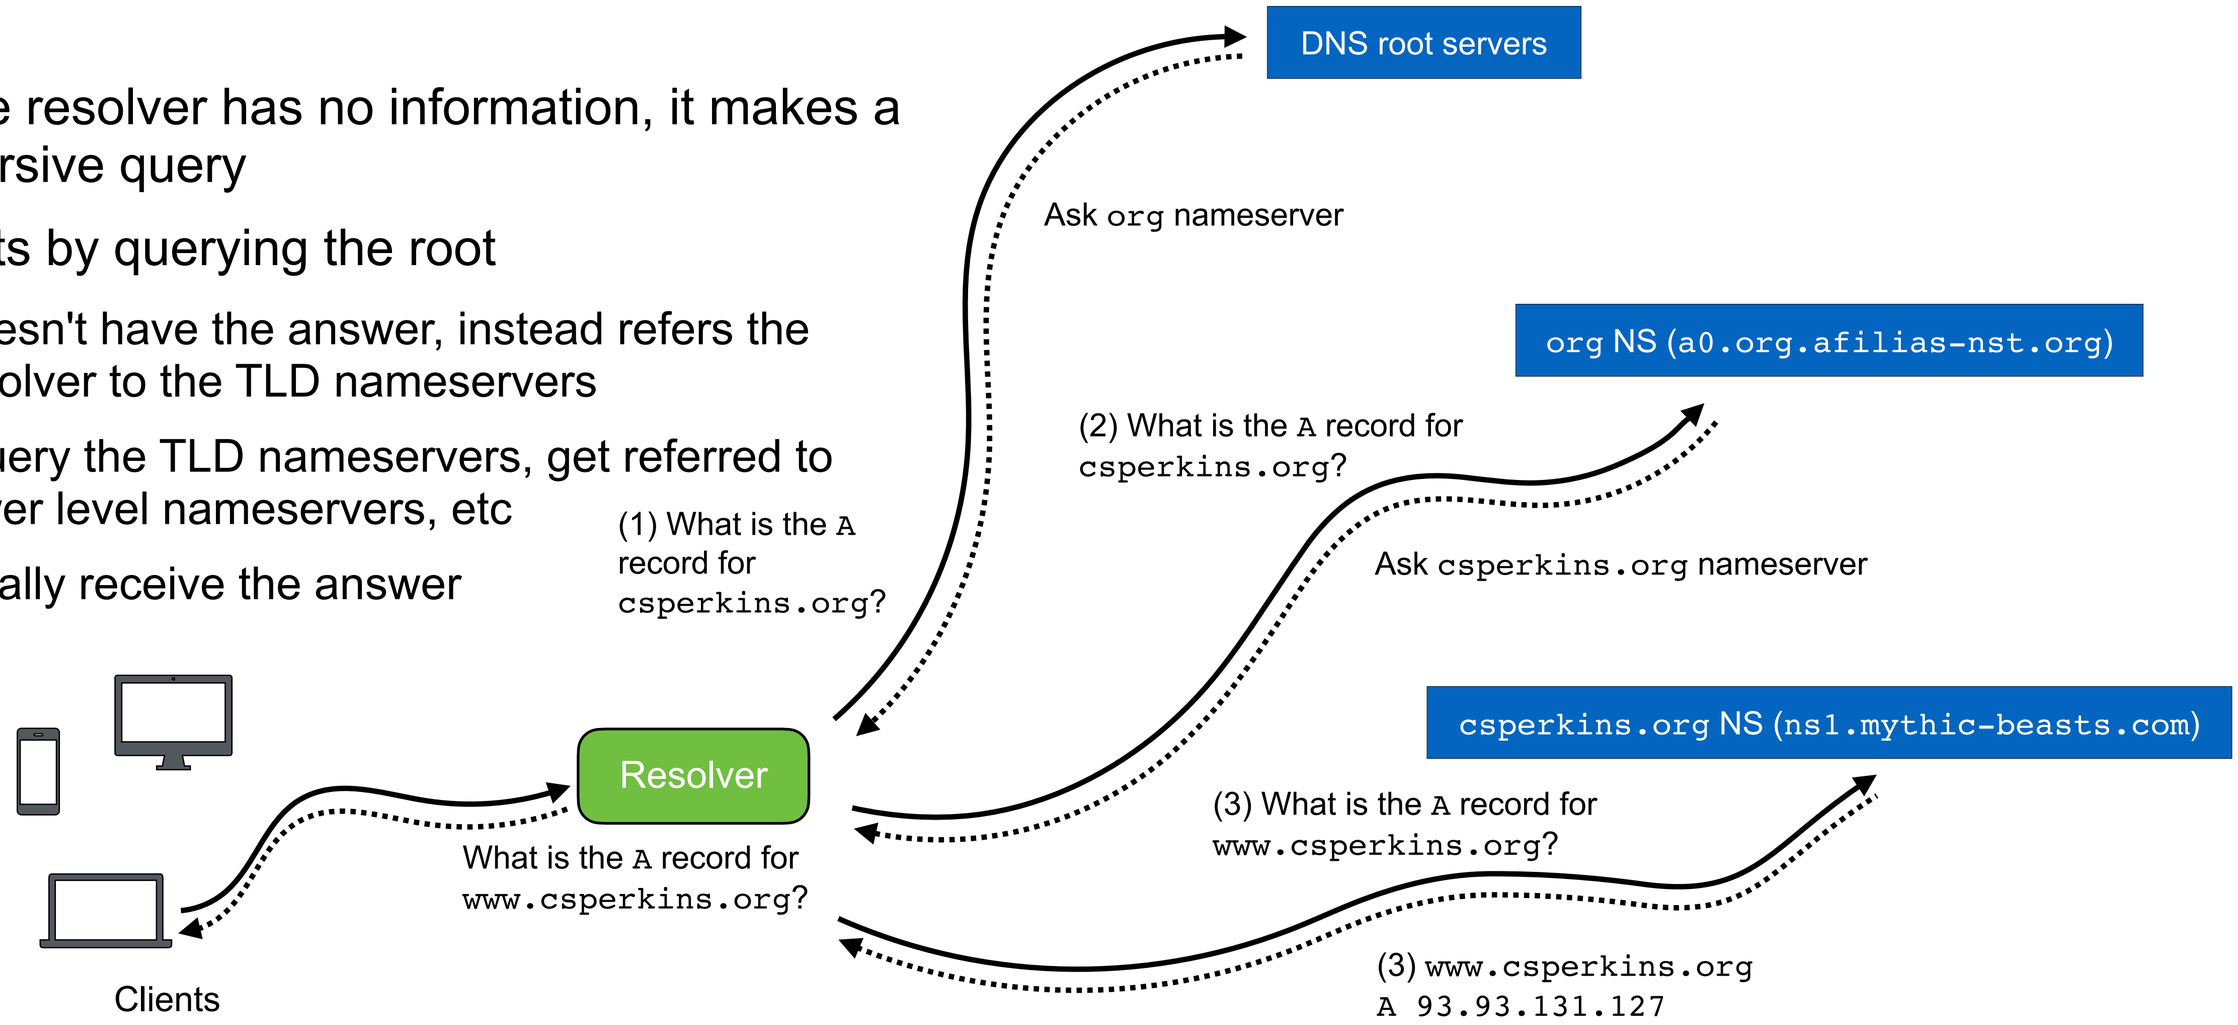

Colin Perkins, Elizabeth Boswell | https://csperkins.org/ | Copyright © 2026 University of Glasgow

DNS Name Resolution (3/3)

- Responses include a **time to live** that lets a resolver cache the value for a certain time
  - Subdomains can set a short TTL and give different answers each time a particular name is requested for load balancing; CDNs directing queries to local caches
- Subsequent queries are answered from the cache, where possible
  - If the cached entry times out, it's refreshed from the next level up in the hierarchy
  - Eventually reaching the DNS root

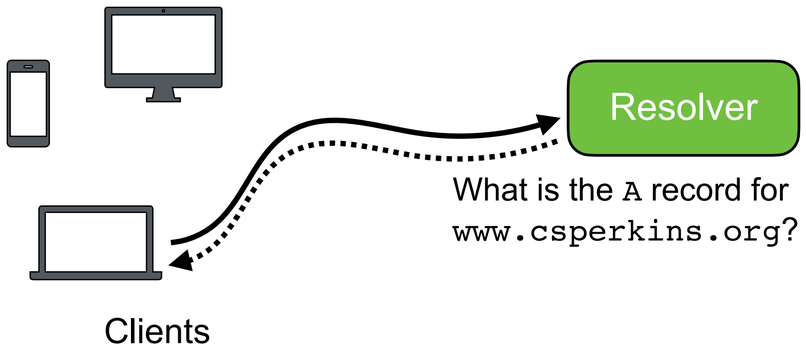

Colin Perkins, Elizabeth Boswell | https://csperkins.org/ | Copyright © 2026 University of Glasgow

University of Glasgow

DNS Names

- Who controls the DNS?
- Top-level domains
- Internationalised DNS
- The DNS root

CC BY ND

Structure of DNS Names

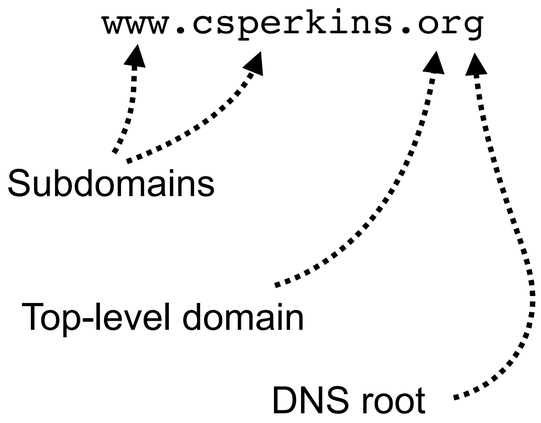

- DNS names are assigned in a hierarchy – subdomains delegated from a top-level domain delegated from the DNS root
  - What top-level domains exist?
  - What policy does a top-level domain have for allocating sub-domains?
  - Who decides when to add new top-levels domains?
  - Who controls the DNS root?

Colin Perkins, Elizabeth Boswell | https://csperkins.org/ | Copyright © 2026 University of Glasgow

From IANA to ICANN (1/2)

- The set of top-level domains is controlled by the Internet Corporation for Assigned Names and Numbers (ICANN)
- ICANN has a complex history
  - ARPANET project – US government funded research, 1966-1990
    - Developed initial versions of Internet protocols, TCP/IP, etc.
    - ARPANET needed protocol specifications and a parameter registry – Jon Postel volunteered as RFC Editor and IANA, originally as a graduate student at UCLA, later at USC/ISI
    - Postel handled domain name allocation, as IANA, funded by US Government
    - Informal, part-time, role – gradually formalised in the late 1990s, leading to formation of ICANN
    - http://www.ietf.org/rfc/rfc2468.txt – “I remember IANA”
  - ICANN formed, incorporating IANA, September 1998
    - Dedicated organisation to manage domain names in the public interest, as a global multi-stakeholder forum, as the Internet started to become widely deployed and commercialised
    - A US not-for-profit corporation based in Los Angeles
    - US Government contractual control of ICANN ended on October 2016

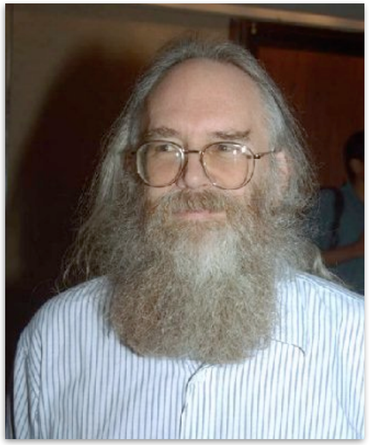

Jon Postel

Colin Perkins, Elizabeth Boswell | https://csperkins.org/ | Copyright © 2026 University of Glasgow

From IANA to ICANN (2/2)

- ICANN has a complex governance model
  - stakeholders include operators of generic and country-code TLDs, governmental representatives, regional Internet registries and root server operators
  - Regular public meetings, 3x per year; annual budget ~$140M
- ICANN is **political** – many countries and organisations want to influence how domain names are managed and allocated

Colin Perkins, Elizabeth Boswell | https://csperkins.org/ | Copyright © 2026 University of Glasgow

Top-Level Domains (1/5)

- There are four types of top-level domain:
  - Country code top-level domains (ccTLDs)
  - Generic top-level domains (gTLDs)
  - Infrastructure top-level domains
  - Special-use top-level domains

Colin Perkins, Elizabeth Boswell | https://csperkins.org/ | Copyright © 2026 University of Glasgow

Top-Level Domains (2/5)

- DNS includes Country Code Top-Level Domains (ccTLDs)
  - ISO standard 3166-1 defines two-letter country names
    - Member states of the United Nations or UN special agencies, parties to the International Court of Justice
  - Every code included in ISO 3166-1 is added to the DNS root zone
    - `.uk`, `.fr`, `.de`, `.cn`, `.us`, `.io`, `.ly`, …
      - Each country has it's own policy for sub-domains of the ccTLD
      - Includes some internationally used TLDs, e.g. `.io` belongs to Chagos Islands - formerly BIOT, (likely) soon to be removed from the ISO standard
    - Exceptions:
      - `.gb` – The UK should use `.gb` to match ISO 3166-1 (`.gov.uk` used to be `.hmg.gb`, but `.gb` never widely used)
      - `.su` – The domain for the Soviet Union still exists and accepts new registrations
      - `.eu` – The EU is not an ISO 3166-1 country, but has a ccTLD
      - `.oz` – Australia, sadly, changed from `.oz` to `.au`

Colin Perkins, Elizabeth Boswell | https://csperkins.org/ | Copyright © 2026 University of Glasgow

Top-Level Domains (3/5)

- DNS also includes Generic Top-level Domains (gTLDs)
  - Core set of gTLDs:
    - `.com`, `.org`, `.net` – unrestricted use
    - `.edu` – higher educational organisations (restricted use; primarily US-based)
    - `.mil` – US military
    - `.gov` – US government
    - `.int` – International Treaty Organisations (United National, Interpol, NATO, Red Cross, …)
  - ICANN has since massively expanded the set of gTLD registrations
    - ~1,200 gTLDs registered
    - e.g., `.scot`, `.box`, `.hsbc`

Colin Perkins, Elizabeth Boswell | https://csperkins.org/ | Copyright © 2026 University of Glasgow

Top-Level Domains (4/5)

- The infrastructure top-level domain, **`.arpa`**, is a historical relic
  - Used in the transition from the ARPANET – the precursor to the Internet
  - ARPA now stands for Address and Routing Parameter Area
- Main current use: reverse DNS
  - Forward DNS lookup:

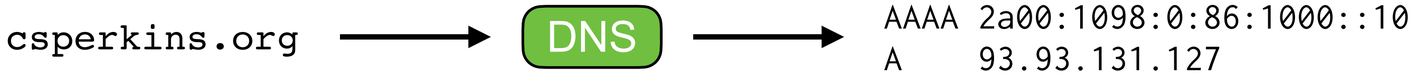

  - Reverse DNS lookup:

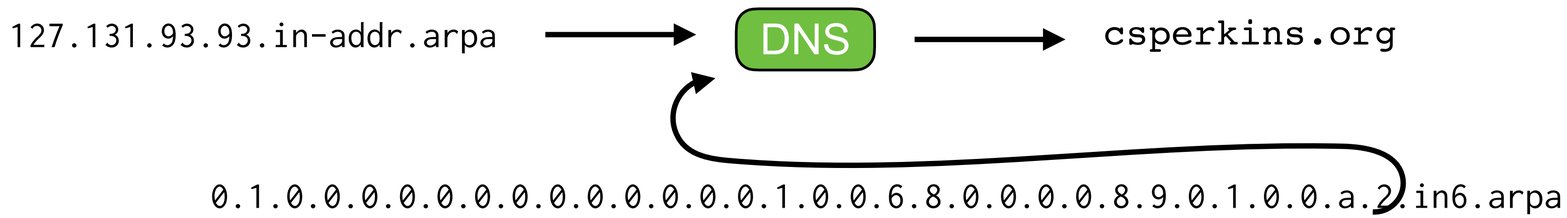

- Other technical uses, such as `home.arpa` (for home networking) or `ipv4only.arpa`

Colin Perkins, Elizabeth Boswell | https://csperkins.org/ | Copyright © 2026 University of Glasgow

Top-Level Domains (5/5)

- Six special-use TLDs exist:
  - **`.example`** – examples and documentation (**`example.com`**, **`example.org`**, etc., also exist)
  - **`.invalid`** – guaranteed never to exist
  - **`.local`** – used in networks with multicast DNS (protocol for local service discovery)
  - **`.localhost`** – represents the local machine
  - **`.onion`** – gateway to Tor hidden services (see https://datatracker.ietf.org/doc/rfc7686/)
  - **`.test`** – for testing

Colin Perkins, Elizabeth Boswell | https://csperkins.org/ | Copyright © 2026 University of Glasgow

Internationalised DNS

- DNS names should be available in any language
  - Initial TLDs and sub-domains were in ASCII, but UTF-8 ought to be allowed
  - DNS should work with UTF-8 name – in practice it doesn't, due to protocol ossification
- Internationalised DNS works around this by translating non-ASCII names into ASCII:
  - Punycode (see https://datatracker.ietf.org/doc/rfc3492/)
    - Encodes **any** unicode text as a sequence of ASCII letters, digits, and hyphens
    - `München` → `Mnchen-3ya`
    - Part after final hyphen is a base-36 (a-z0-9) encoded representation of a sequence of Unicode character and the locations where they should be inserted
  - Internationalised DNS names use Punycode, prefixed with `xn--`
    - e.g., `http://Яндекс.рф` translates to `http://xn--70akdum1a.xn--p1ai` (Yandex, a Russian search engine)

Colin Perkins, Elizabeth Boswell | https://csperkins.org/ | Copyright © 2026 University of Glasgow

The DNS Root (1/3)

- ICANN decides the set of legal top-level domains – the root servers then advertise the name servers for these domains
- What are the root servers?
  - The set of 13 servers that advertise the name servers for the top level domains
    - a.root-servers.net → m.root-servers.net
    - Also have well-known IPv4 and IPv6 addresses
  - Why 13 servers?
    - Want to be able to ask a DNS resolver to return a list of the root servers
    - DNS over UDP has a size limit on replies → 13 root servers is all that will fit into a single UDP packet

Colin Perkins, Elizabeth Boswell | https://csperkins.org/ | Copyright © 2026 University of Glasgow

The DNS Root (2/3)

- Heavily US-based, for historical reasons – discussed later
- The IP addresses of root servers rarely change (8 changes 2000-2023) – but who operates the root servers could change

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
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, Elizabeth Boswell | https://csperkins.org/ | Copyright © 2026 University of Glasgow

The DNS Root (3/3)

- There are not really 13 servers → **anycast routing** (lecture 9)
- Same IP address advertised from multiple places in the network – which replica you reach depends where you're located
- There are 13 sets of IP addresses used by the root servers, but almost 2000 physical servers
  - check https://root-servers.org/ for interactive map

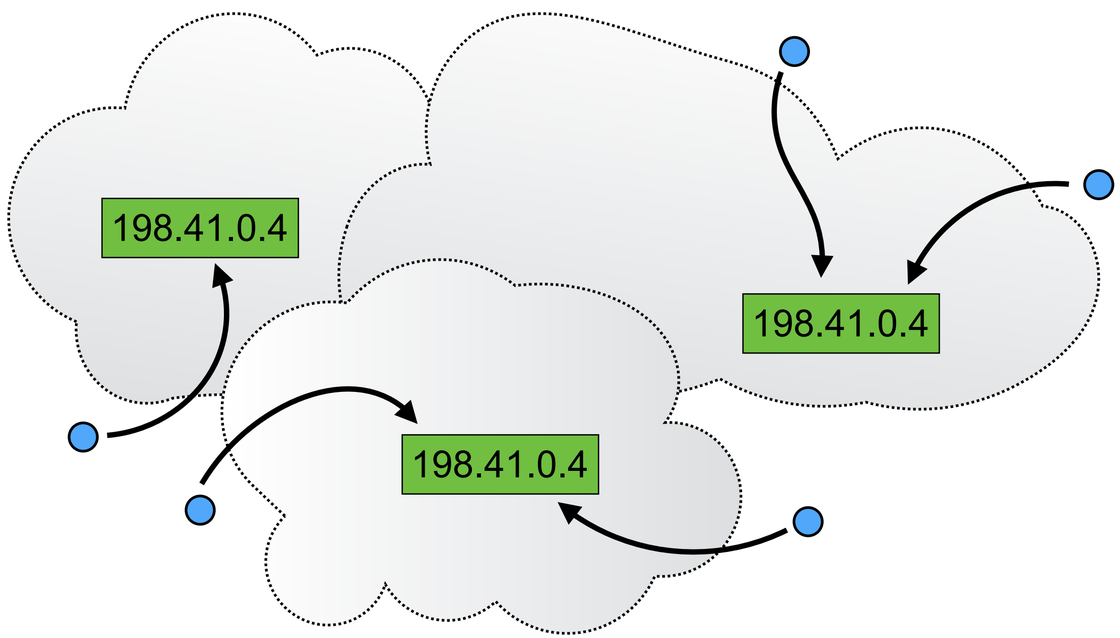

Colin Perkins, Elizabeth Boswell | https://csperkins.org/ | Copyright © 2026 University of Glasgow

University of Glasgow

Methods for DNS Resolution

- DNS security
- DNS resolution over UDP, TLS, HTTPS, and QUIC

CC BY ND

DNS Security (1/3)

- **DNS has historically been completely insecure**
  - Queries and responses delivered via unencrypted and unauthenticated protocol
  - Responses do not include a digital signature to verify authenticity of data
- Trivial to eavesdrop on who is looking up what name
- Trivial for on-path attackers, or malicious resolvers, to forge replies

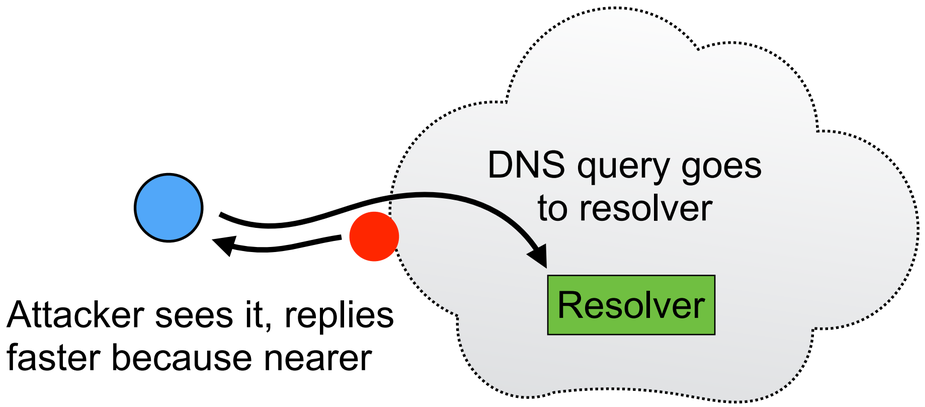

e.g., possible if victim and attacker are in the same cafe, using insecure Wi-Fi

Colin Perkins, Elizabeth Boswell | https://csperkins.org/ | Copyright © 2026 University of Glasgow

DNS Security (2/3)

- Two components of securing DNS:
  - Transport security
    - Make DNS queries, and receive replies, over TLS (or some other secure channel)
    - Queries and responses are encrypted, so can't be understood or modified by attacker
    - This secures the path between **client and resolver** - but work is being done to extend transport security to communication between resolver and nameservers (https://datatracker.ietf.org/doc/rfc9539/)
    - If you trust the resolver, this protects against attack

Colin Perkins, Elizabeth Boswell | https://csperkins.org/ | Copyright © 2026 University of Glasgow

DNS Security (3/3)

- Two components of securing DNS:
  - Transport security
  - Record security – DNSSEC
    - Add a digital signature to DNS responses that client can verify to check the data is valid
      - ICANN signs the root zone
      - Root servers sign information they provide about TLDs
      - TLDs sign information they provide about sub-domains
      - …
    - Allows a client to verify signatures back to the root, providing a chain of trust to demonstrate ownership of a domain – protects against malicious resolvers
    - Makes extensive use of public key cryptographic techniques – details are complex
    - Implemented, but not widely used
- **Need both transport and record security for fully secure DNS**

Colin Perkins, Elizabeth Boswell | https://csperkins.org/ | Copyright © 2026 University of Glasgow

DNS Over UDP (1/4)

- DNS queries generally made over UDP port 53
  - Queries and responses are generally small enough to fit into a single packet
  - TCP reliability isn't needed – if no answer, retransmit the query
  - Congestion control isn't needed – can't adjust the rate you send a single packet

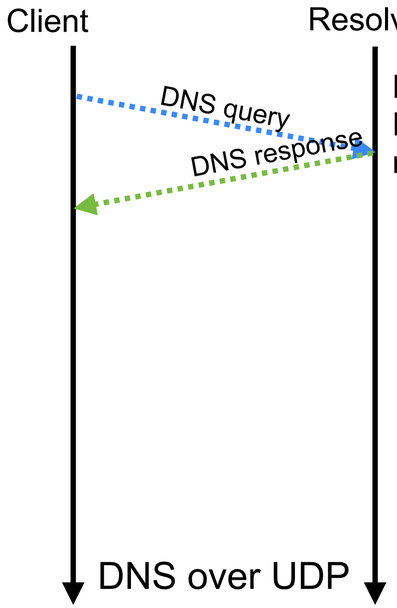

DNS query over UDP completes in one RTT – one packet for query, one for response

DNS over TCP exists, but is slower; initial and final TCP handshakes take several RTT and send extra packets.

Mainly used as fallback if message is too big for UDP.

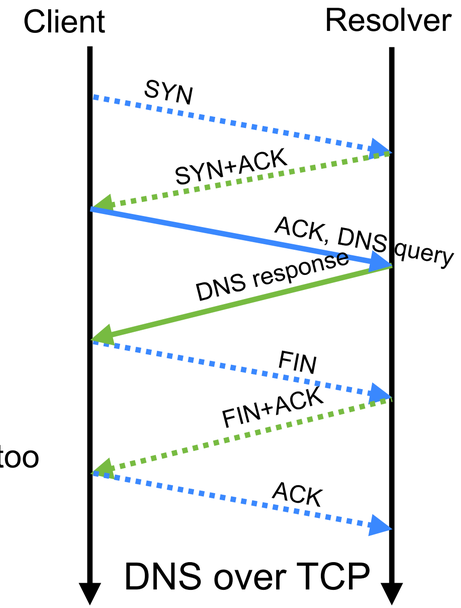

Colin Perkins, Elizabeth Boswell | https://csperkins.org/ | Copyright © 2026 University of Glasgow

DNS Over UDP (2/4)

- Question section:
  - List of domain names and requested record type
    - e.g., what is the `AAAA` record for domain `csperkins.org`

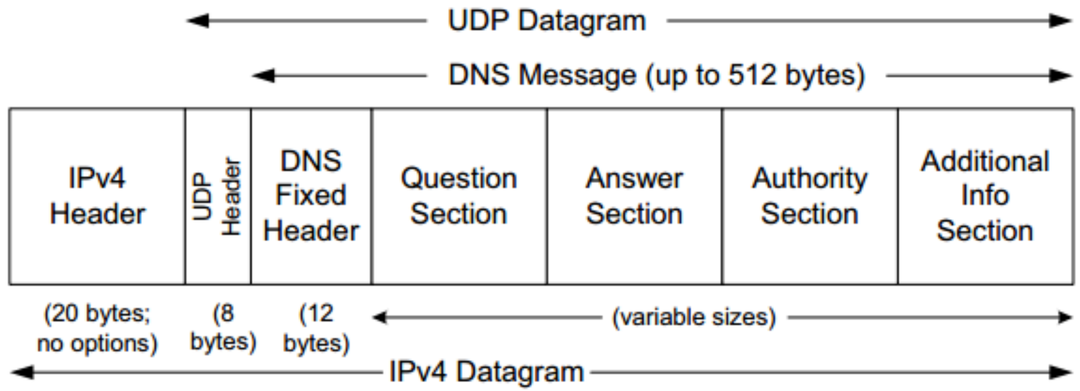

Colin Perkins, Elizabeth Boswell | https://csperkins.org/ | Copyright © 2026 University of Glasgow

DNS Over UDP (3/4)

- Answer, authority, and additional information sections:
  - List of domain names and record type, record data, and time-to-live
  - Answer section answers a question made in a previous query
    - e.g., the `AAAA` record for domain csperkins.org is 2a00:1098:0:86:1000::10 and it's valid for 1 hour
  - Authority describes where the answer came from
  - Additional section contains auxiliary information

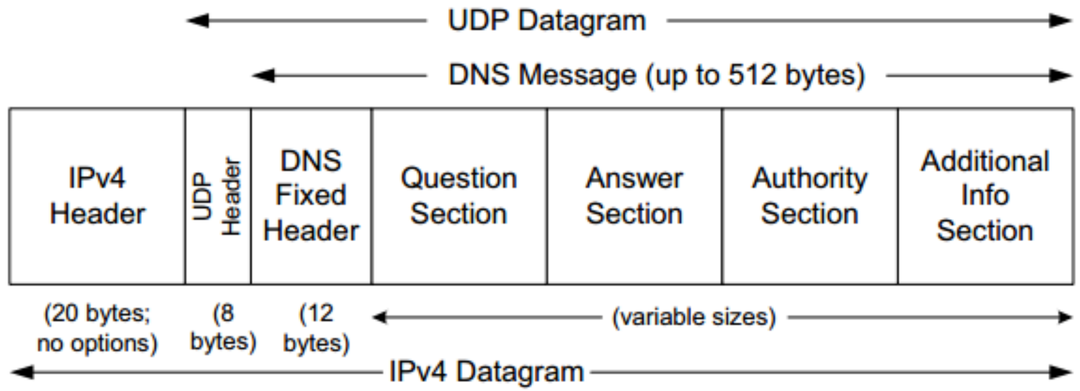

Colin Perkins, Elizabeth Boswell | https://csperkins.org/ | Copyright © 2026 University of Glasgow

DNS Over UDP (4/4)

```
[stlinux02] > dig csperkins.org

; <<>> DiG 9.11.4-P2-RedHat-9.11.4-9.P2.el7 <<>> csperkins.org
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 51409
;; flags: qr rd ra; QUERY: 1, ANSWER: 1, AUTHORITY: 2, ADDITIONAL: 5

;; OPT PSEUDOSECTION:
; EDNS: version: 0, flags:; udp: 4096
;; QUESTION SECTION:
;csperkins.org.         IN A

;; ANSWER SECTION:
csperkins.org.    2681  IN A  93.93.131.127

;; AUTHORITY SECTION:
csperkins.org.    78278 IN NS ns2.mythic-beasts.com.
csperkins.org.    78278 IN NS ns1.mythic-beasts.com.

;; ADDITIONAL SECTION:
ns1.mythic-beasts.com.  70870  IN   A    45.33.127.156
ns2.mythic-beasts.com.  157301 IN   A    93.93.128.67
ns1.mythic-beasts.com.  70870  IN   AAAA 2600:3c00:e000:19::1
ns2.mythic-beasts.com.  157301 IN   AAAA 2a00:1098:0:80:1000::10

;; Query time: 0 msec
;; SERVER: 130.209.244.1#53(130.209.244.1)
;; WHEN: Wed Mar 04 18:26:53 GMT 2020
;; MSG SIZE  rcvd: 199
```

Use **`dig`** on Linux or macOS to make DNS queries

(Tip: use `dig +trace` to see the queries made by a resolver)

Colin Perkins, Elizabeth Boswell | https://csperkins.org/ | Copyright © 2026 University of Glasgow

DNS Over TLS (DoT)

- DNS over UDP is insecure
  - The packets are not encrypted or authenticated
  - Devices on the path between client and resolver can see DNS queries and responses – and can forge responses
- DNS over TLS solves this problem
  - DNS client opens a TCP connection to the resolver (port 853)
  - DNS client and resolver negotiate a TLS 1.3 session on the TCP connection
  - DNS client sends query, and receives response, over the TLS connection
    - DNS over TLS messages are formatted exactly the same as DNS over UDP, and contain exactly the same information – only difference is that they're sent over TLS not UDP
    - Slower and higher overhead than DNS over UDP – due to need to negotiate TCP and TLS – but more secure. Can reduce overhead by keeping the connection open

Colin Perkins, Elizabeth Boswell | https://csperkins.org/ | Copyright © 2026 University of Glasgow

DNS over HTTPS (DoH)

- DoH allows a client to send queries to a resolver using HTTPS
  - Can use with either GET or POST methods in HTTPS

```
GET /dns-query?dns=AAABAAABAAAAAAAAA3d3dwdleGFtcGxlA2NvbQAAAQAB HTTP/1.1
Accept: application/dns-message
```

Base-64 encoded version of the data that would be sent in a DNS-over-UDP query

Uses /dns-query as the URL path

```
POST /dns-query HTTP/1.1
Accept: application/dns-message
Content-type = application/dns-message
Content-length = 33

<33 bytes of UDP query, exactly as if sent in a UDP packet>
```

  - HTTP response has `Content-Type: application/dns-message` and contains the exact same data that would be sent in a UDP-based DNS response

Colin Perkins, Elizabeth Boswell | https://csperkins.org/ | Copyright © 2026 University of Glasgow

DNS over QUIC (DoQ)

- DNS over QUIC was standardised in 2024
  - https://datatracker.ietf.org/doc/rfc9250/
  - Same principle as DNS over TLS:
    - Client opens a QUIC connection to the resolver
      - Negotiates TLS security as part of the connection setup
    - Sends the query and receives the response over that connection
      - Queries and response contain exactly the same data as DNS over UDP – just sent via QUIC

Colin Perkins, Elizabeth Boswell | https://csperkins.org/ | Copyright © 2026 University of Glasgow

Methods for DNS Resolution

- Increasingly many ways of making DNS queries:
  - DNS over UDP
  - DNS over TCP
  - DNS over TLS
  - DNS over HTTPS
  - DNS over QUIC
- The contents of the query and the response are **identical** in all cases
  - They change how the query is delivered to the resolver and how the response is returned, but not the contents of the messages
  - They change the security guarantees provided
  - They potentially gives clients more flexibility to query different resolvers

Colin Perkins, Elizabeth Boswell | https://csperkins.org/ | Copyright © 2026 University of Glasgow

University of Glasgow

The Politics of Names

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

CC BY ND

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 5G 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, Elizabeth Boswell | https://csperkins.org/ | Copyright © 2026 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, Elizabeth Boswell | https://csperkins.org/ | Copyright © 2026 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, Elizabeth Boswell | https://csperkins.org/ | Copyright © 2026 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 and 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, Elizabeth Boswell | https://csperkins.org/ | Copyright © 2026 University of Glasgow

Intellectual Property and the DNS (1/2)

- 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, Elizabeth Boswell | https://csperkins.org/ | Copyright © 2026 University of Glasgow

Intellectual Property and the DNS (2/2)

- ICANN allows interested parties to apply for new gTLDs
  - provided they pay ~$200,000
- What if two parties are interested in the same gTLD?
  - especially relevant to TLDs belonging to specific communities (e.g. religious/political/cultural/social groups)
  - ICANN has a process for deciding which party gets control of the TLD
- Some new gTLDs are in widespread **internal** use
  - e.g. Fritz!Box home router allows users to access configuration page on `http://fritz.box`
  - `.box` TLD was added to the DNS in 2023, leading to `fritz.box` becoming available on the public DNS - this is a security risk

Colin Perkins, Elizabeth Boswell | https://csperkins.org/ | Copyright © 2026 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
  - What if someone registers <your name>`.xxx` to inflict reputational damage?
- Who controls what TLDs ICANN permits? Who should control it?



XXXtortion? Inferring Registration Intent in the .XXX TLD

Tristan Halvorson
trhalvorson@cs.ucsd.edu
Kirill Levchenko
klevchen@cs.ucsd.edu
Stefan Savage
savage@cs.ucsd.edu
Geoffrey M. Voelker
voelker@cs.ucsd.edu
Department of Computer Science and Engineering
University of California, San Diego

ABSTRACT

After a decade-long approval process, multiple rejections, and an independent review, ICANN approved the .xxx TLD for inclusion in the Domain Name System, to begin general availability on December 6, 2011. Its sponsoring registry proposed it as an expansion of the name space, as well as a way to separate adult from child-appropriate content. Many independent groups, including trademark holders, political groups, and the adult entertainment industry itself, were concerned that it would primarily generate value through defensive and speculative registrations, without actually serving a real need. This paper measures the validity of these concerns using data gathered from ICANN, whois, and Web requests. We use this information to characterize each .xxx domain and infer the registrant's most likely intent. We find that at most 3.8% of .xxx domains host or redirect to potentially legitimate Web content, with the rest generally serving either defensive or speculative purposes. Indeed, registrants spent roughly $13M up-front to defend existing brands and trademarks within the .xxx TLD, and an additional $11M over the course of the first year. Additional evidence suggests that over 80% of annual domain registrations are for purely defensive purposes and do not even resolve.

Categories and Subject Descriptors

C.2.m [**Computer Communication Networks**]: Miscellaneous; K.4.1 [**Computers and Society**]: Public Policy Issues—*regulation, use/abuse of power*

Keywords

DNS; TLDs; Internet Economics; Abuse; Registration Intent

1. INTRODUCTION

Domain names are a unique and exclusive resource. Once a name is registered it is no longer available for others. Thus, names that correspond to popular trademarks, brands, proper names, common words, acronyms and so on, can accrue tremendous commercial value for those who hold them. On the theory that expanding the name space would therefore increase competition and create more naming opportunities, ICANN has created a number of new “top-level domains” (TLDs) in recent years (e.g., biz, aero, xxx, etc.) and is today considering a range of bids to dramatically increase the domain name space further. However, there has been considerable debate about the extent to which new TLDs will actually serve a real need, or whether they will instead exploit existing brand holders who will feel compelled to protect their marks on a defensive basis.

In this paper, we examine this question through the unique lens of one particular TLD: xxx. The xxx TLD was created to serve the needs of the adult entertainment industry and its approval by ICANN was a matter of great controversy (both from within the adult industry and without). Critics argued that the TLD did not serve any real need and that its creation would largely be monetized via speculators and defensive registrations. With the benefit of hindsight (xxx opened for general availability in December of 2011) we examine this question empirically. While it is impossible to measure the intent of a registrant directly, we argue that such intent can be inferred through how each of their domains are used (i.e., does it resolve?, serve content?, what kind of content?, etc.)

By combining public ICANN records, analysis of the xxx zone files, active whois requests and crawling the content of live sites, we have constructed what we believe is the most comprehensive picture to date of how xxx is used in practice. We find that almost half of all xxx domains do not appear in the zone file at all, and of the approximately 107k domains that do resolve, three-fourths are objectively defensive registrations. A significant fraction of the remainder point to parked sites, suggesting speculative use. Finally, our calculations suggest that the clear majority of registration revenues, both during the pre-registration period and for ongoing operation, is driven by defensive concerns rather than entrepreneurship. While the xxx TLD situation is clearly unique, we believe it is a telling data point as the Internet community considers a range of far broader expansions in the name space.

2. BACKGROUND

The Domain Name System (DNS) is widely used to provide memorable identifiers for Internet hosts and services. At a technical level, DNS simply translates human-readable strings to machine addresses. In the Internet's modern incarnation, the DNS is overseen by the Internet Corporation for Assigned Names and Numbers (ICANN) who in turn licenses individual *registries* to operate distinct name spaces. Thus, each registry maintains exclusive management rights over its names, and is responsible for licensing *registrars* (who sell registration services on a retail basis) and for ensuring authoritative name resolution for such names.

Copyright is held by the International World Wide Web Conference Committee (IW3C2). IW3C2 reserves the right to provide a hyperlink to the author's site if the Material is used in electronic media.
WWW'14, April 7–11, 2014, Seoul, Korea.
ACM 978-1-4503-2744-2/14/04.
http://dx.doi.org/10.1145/2566486.2567995

Tristan Halvorson, Kirill Levchenko, Stefan Savage, and Geoffrey M. Voelker. 2014. XXXtortion? Inferring registration intent in the .XXX TLD. International conference on World Wide Web (WWW '14). https://doi.org/10.1145/2566486.2567995

Colin Perkins, Elizabeth Boswell | https://csperkins.org/ | Copyright © 2026 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, Elizabeth Boswell | https://csperkins.org/ | Copyright © 2026 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 benefit in controlling a DNS root server?
  - Is there 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, Elizabeth Boswell | https://csperkins.org/ | Copyright © 2026 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, Elizabeth Boswell | https://csperkins.org/ | Copyright © 2026 University of Glasgow

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