DNS Names

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

Structure of DNS Names

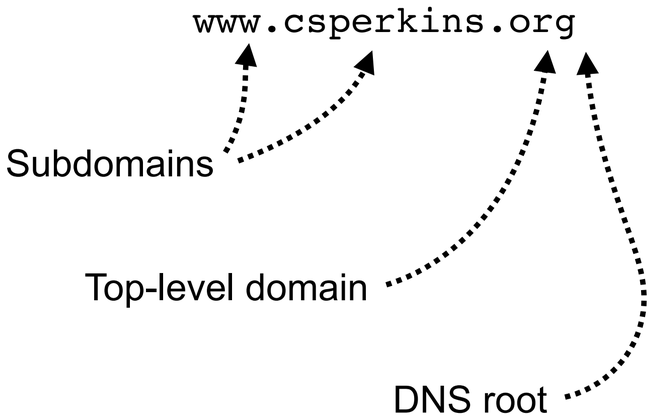

- DNS names are assigned in a hierarchy – subdomains delegated from a top-level domain delegated from the DNS root
  - What top-level domains exist?
  - What policy does each top-level domain have for allocating sub-domains?
  - Who decides when to add new top-levels domains?
  - Who controls the DNS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

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

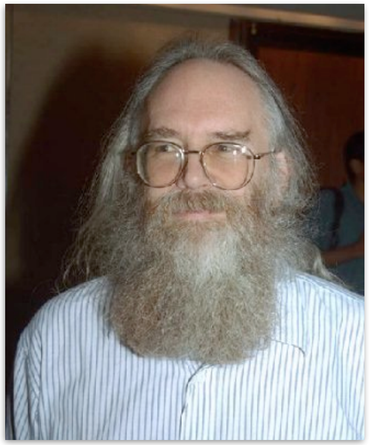

Jon Postel

From IANA to ICANN (2/2)

- ICANN has a complex governance model:
  - Board of governors; generic names supporting organisation – *gTLDs*; country code names supporting organisation – *ccTLDs*; address supporting organisation – *regional Internet registries*; governmental advisory committee – *representatives from each of the 112 UN-recognised countries*; at-large advisory committee; root server advisory committee; security and stability advisory committee; and technical liaison group
  - 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

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

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, UN special agencies (ITU, IMF, UNESCO, WHO, …), 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
      - `.cs` – JANET Name Resolution Systems and the "Czechoslovakia problem" (historical)
    - 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`

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
    - ~1,500 gTLDs registered
    - e.g., `.scot` is a gTLD

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
  - It has one current use: reverse DNS
- Forward DNS lookup:

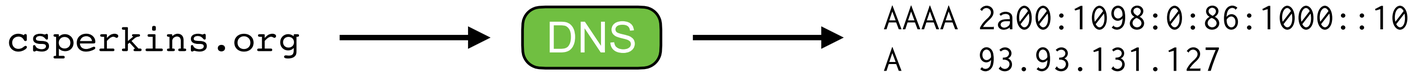

- Reverse DNS lookup:

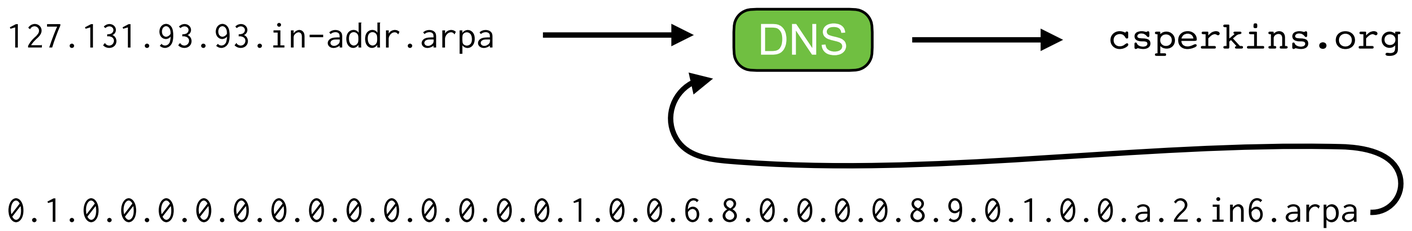

Top-Level Domains (5/5)

- Six special-use domains exist:
  - **`.example`** – examples and documentation (**`example.com`**, **`example.org`**, etc., also exist)
  - **`.invalid`** – guaranteed never to exist
  - **`.local`** – represents the local network
  - **`.localhost`** – represents the local machine
  - **`.onion`** – gateway to Tor hidden services (see https://datatracker.ietf.org/doc/rfc7686/)
  - **`.test`** – for testing

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, **in principle**, 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`
    - `Bahnhof München-Ost` → `Bahnhof Mnchen-Ost-u6b`
    - `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)

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 relies → 13 root servers is all that will fit into a single UDP packet

The DNS Root (2/3)

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

- Heavily US-based, for historical reasons – discussed later
- The IP addresses of root servers cannot be changed – they're too widely known – who operates the root servers could change

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 IP addresses used by the root servers, but many more physical servers

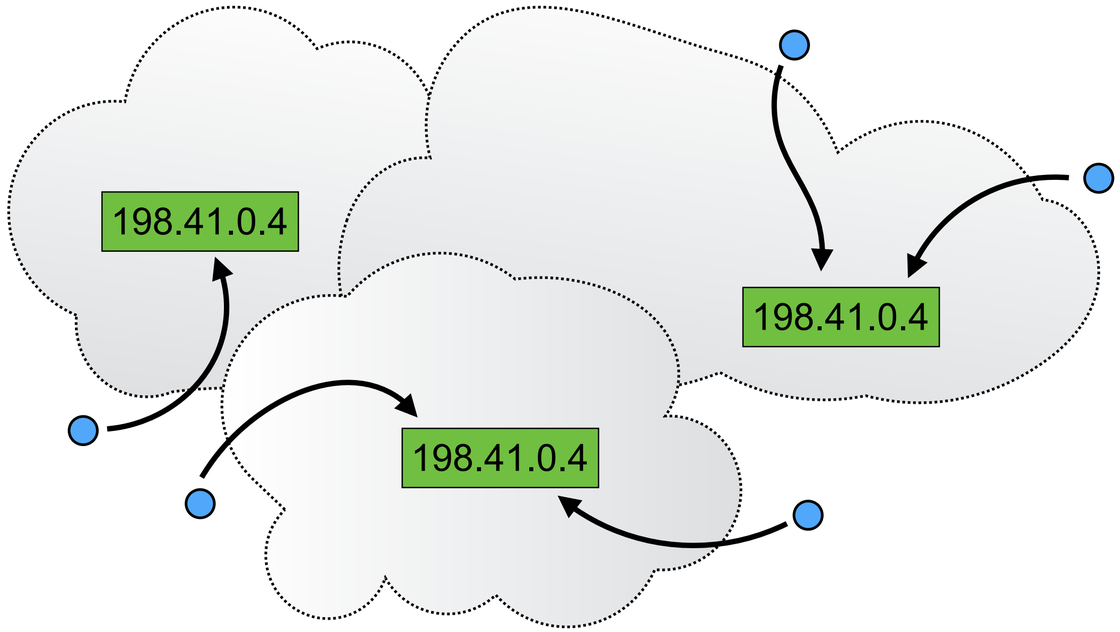

DNS Names

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