

Domain Name Service



To do ...

- ❑ Names
- ❑ Naming hierarchy and DNS
- ❑ Challenges and new trends

Names, identifiers and addresses

- **Names** are used to denote **entities** in a system
 - Hosts, printers, files, processes, users
- A **namespace** is the set of all possible names
 - Could be flat or hierarchical
- To operate on an entity, e.g. print a file, you need to access it at an **access point**
- An **address** is the name of an access point
 - An entity can offer one or more access points (think phone #s)
 - Address of an access point of an entity = address of an entity (telephone # - telephone - person)

*742 Evergreen Terrace
Springfield*



Names and resolution

- A **naming system** maintains a collection of bindings of names to values (e.g., addresses)
- A **resolution mechanism** is a procedure that, when invoked with a name, returns the corresponding value
- A **name server** is a specific implementation of a resolution mechanism

Domain Name System

- Naming for the Internet before 1983
 - Each computer retrieved HOST.TXT, maintained by the Network Information Center (NIC), from a computer at SRI
 - To add a name, email hostname/address pair to NIC
 - A legacy – a fossil host file still exist in most modern OS
- DNS (Paul Mockapetris, then at UC Irvine) [RFC 1034, 1035]
 - A distributed database implemented in a hierarchy of servers
 - An application-layer protocol for hosts to query the database

Domain Name System

- Main service
 - Mapping between two name spaces – hostnames and IP addresses
- Some other services provided by DNS
 - Host aliasing – mapping between alias names and the canonical name of a host
 - Mail server aliasing
 - Good to have a mnemonic address for email
 - Load distribution
 - Rather than a single IP, a list of IP associated with one canonical name
 - Replay with all IPs, but in rotating order

```
$ dig CNAME www.northwestern.edu
...
;; ANSWER SECTION:
www.northwestern.edu. 299 IN CNAME common.wideip.northwestern.edu.
```

```
$ dig cs.northwestern.edu MX
...
;; ANSWER SECTION:
cs.northwestern.edu. 21599 IN MX 0 cuda.eecs.northwestern.edu.
cs.northwestern.edu. 21599 IN MX 0 barra.eecs.northwestern.edu.
```

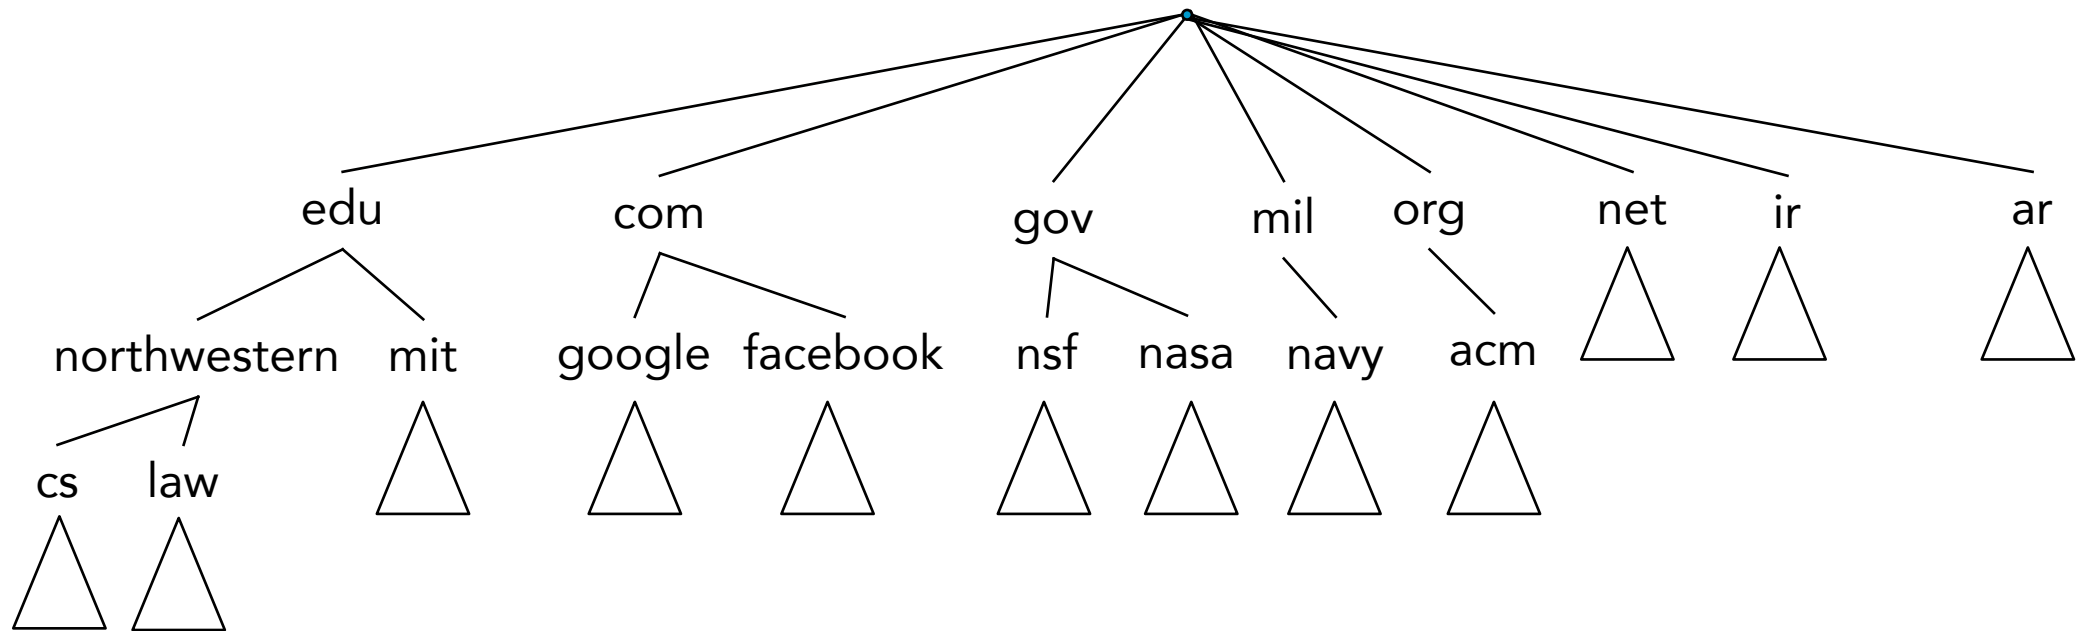
A simple solution – A centralized directory

What could be wrong with it?

- A single point of failure – If it crashes, the entire Internet does
- Huge traffic volume to handle all queries from million of hosts
- A single DNS server would be far a way from most nodes, so requests would have to travel long, maybe over congested links
- A huge DB that would have to update all the time

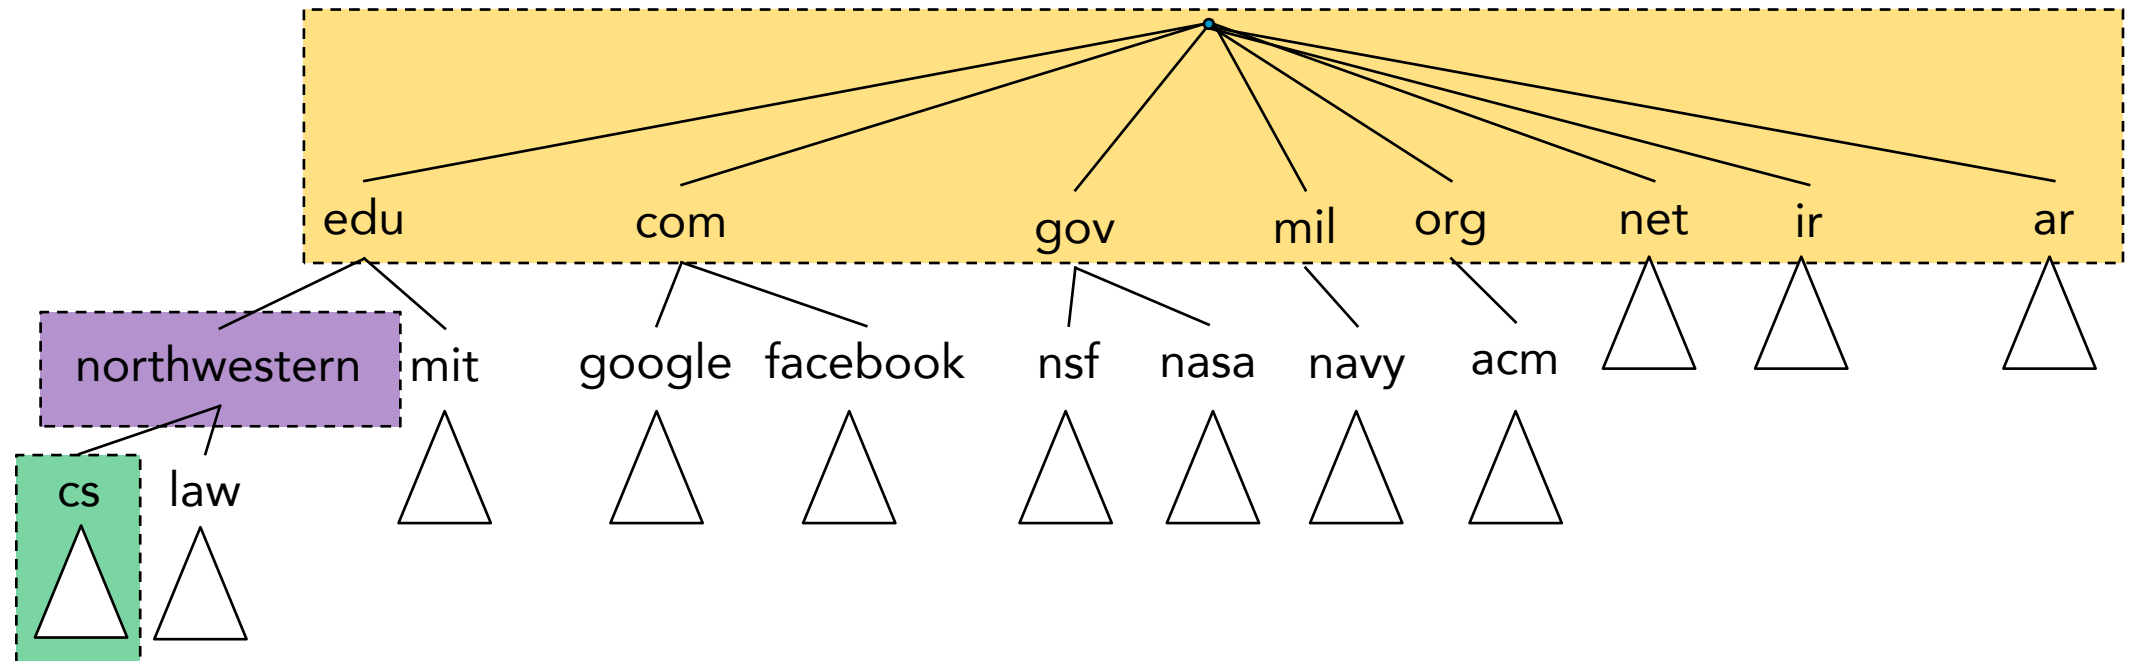
Domain hierarchy

- DNS implements a hierarchical name space for Internet objects
 - Names are interpreted right to left, using "." as separators
zappa.cs.northwestern.edu
 - Recent expansion at the top level which has been sort-of narrow



Partitioning into zones

- For scaling, partition it into zones
 - Each corresponding to some administrative authority
 - And a fundamental unit of implementation – the name server
 - Top: Internet Corporation for Assigned Names and Numbers (ICANN)



DNS – A distributed, hierarchical database

- A first approximation

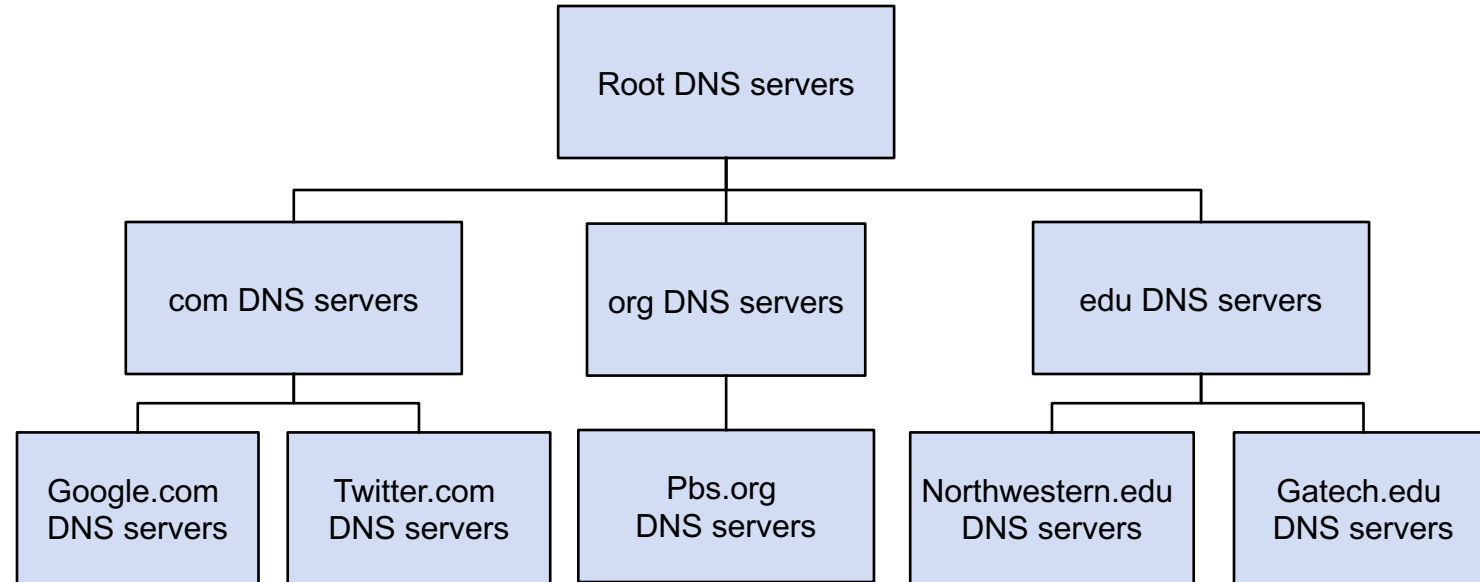
Root DNS servers managed by different organizations

Top-level domain servers:

Generic TLDs (gTLDs) - .com, .edu, .net

Country-code TLDs - .ar, .uk, .in, .cn, ...

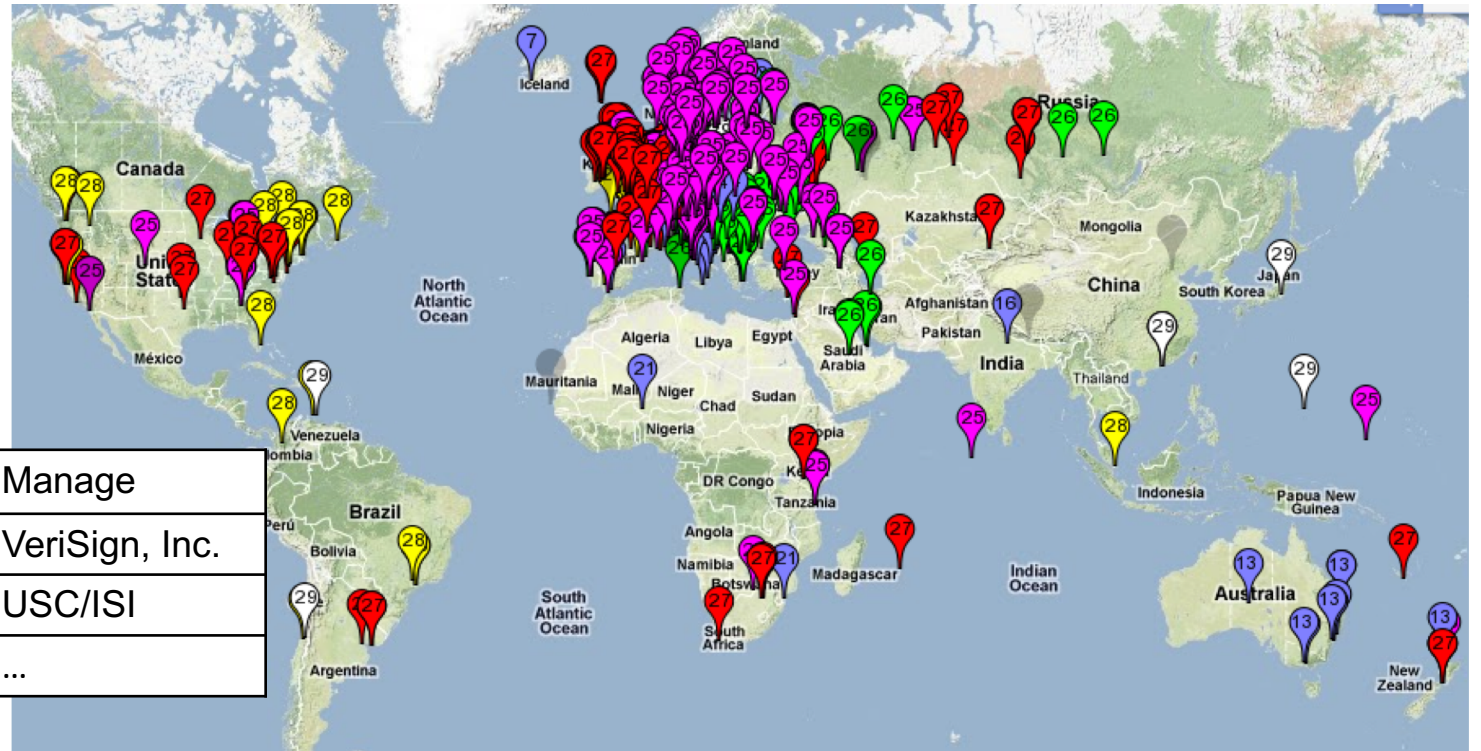
Authoritative DNS servers where the DNS records of an organization are hosted (could be done by a third-party)



Resolving zappa.cs.northwestern.edu?

Root DNS servers

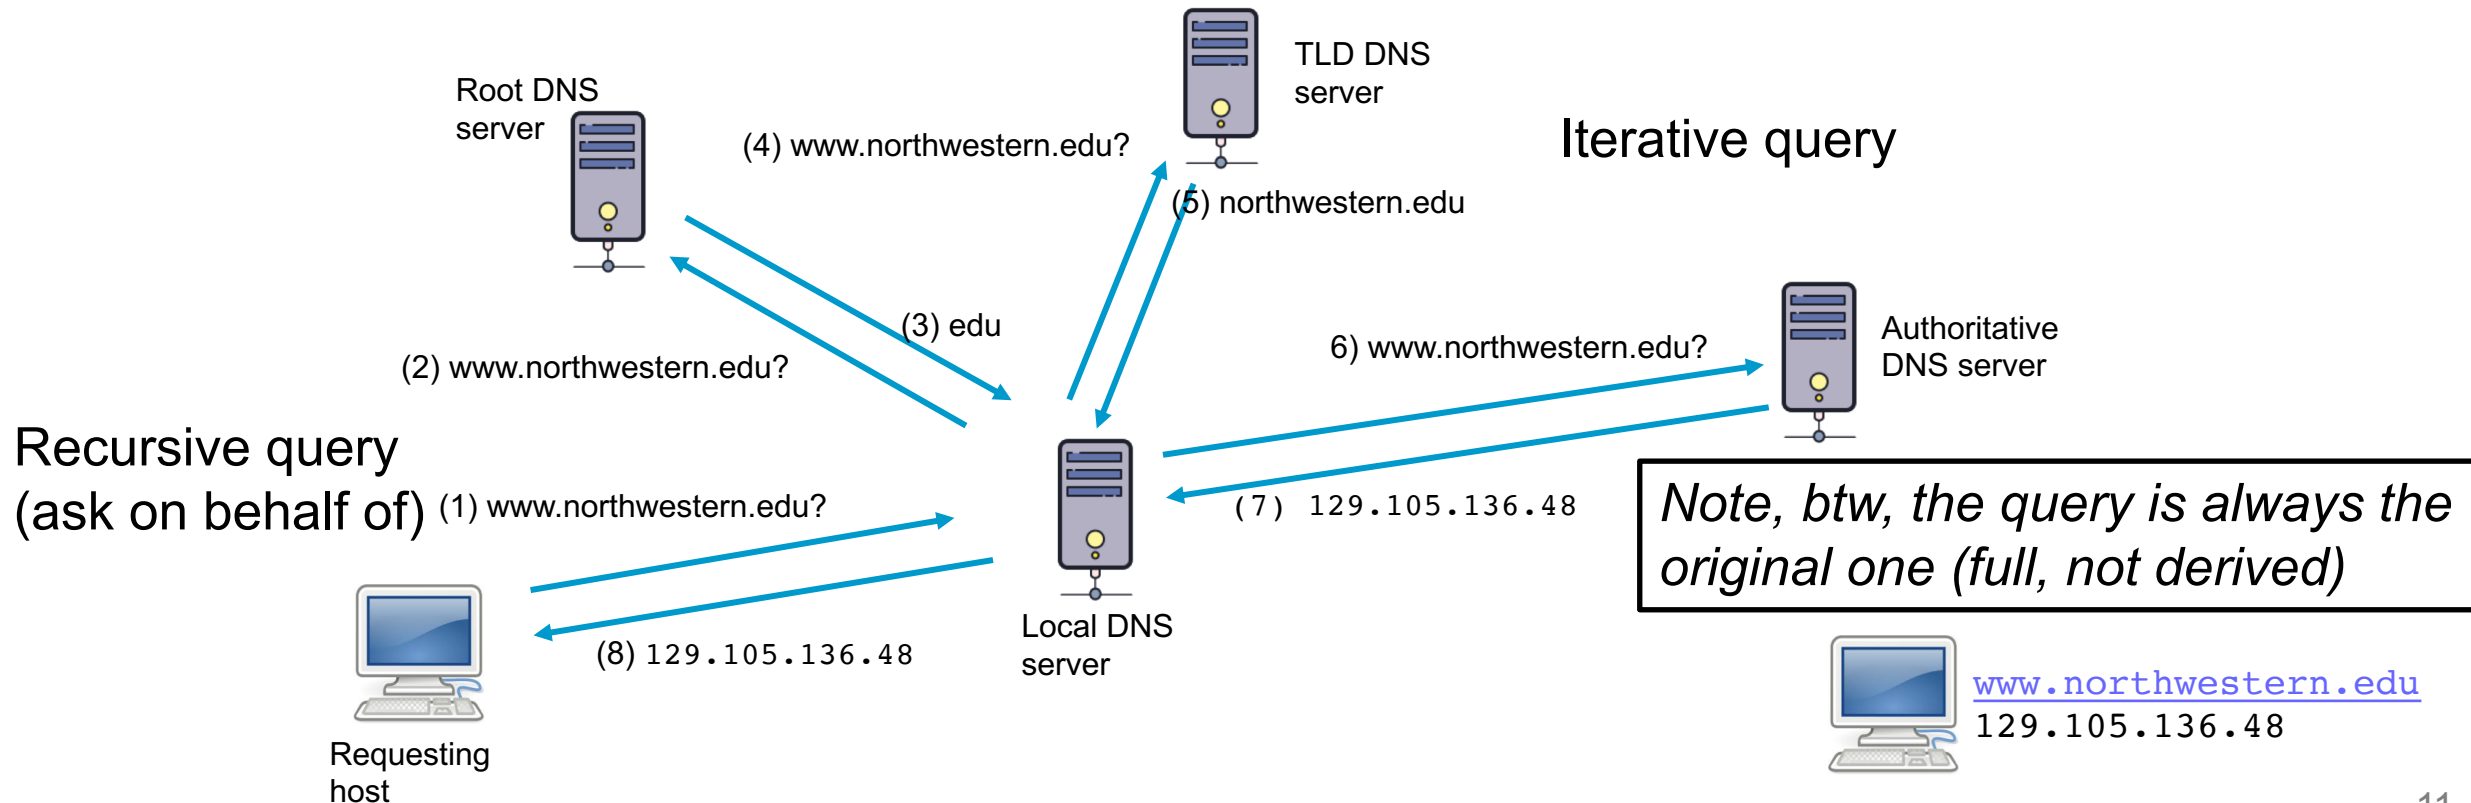
- Keep track of authoritative nameservers for TLDs
- Nameservers are statically configured with 13 IP addresses for the root servers (named authorities, over 400 servers)



Hostname	IP address	Manage
a.root-servers.net	198.41.0.4, 2001:503:ba3e::2:3	VeriSign, Inc.
b.root-servers.net	199.9.14.201, 2001:500:200::b	USC/ISI
...

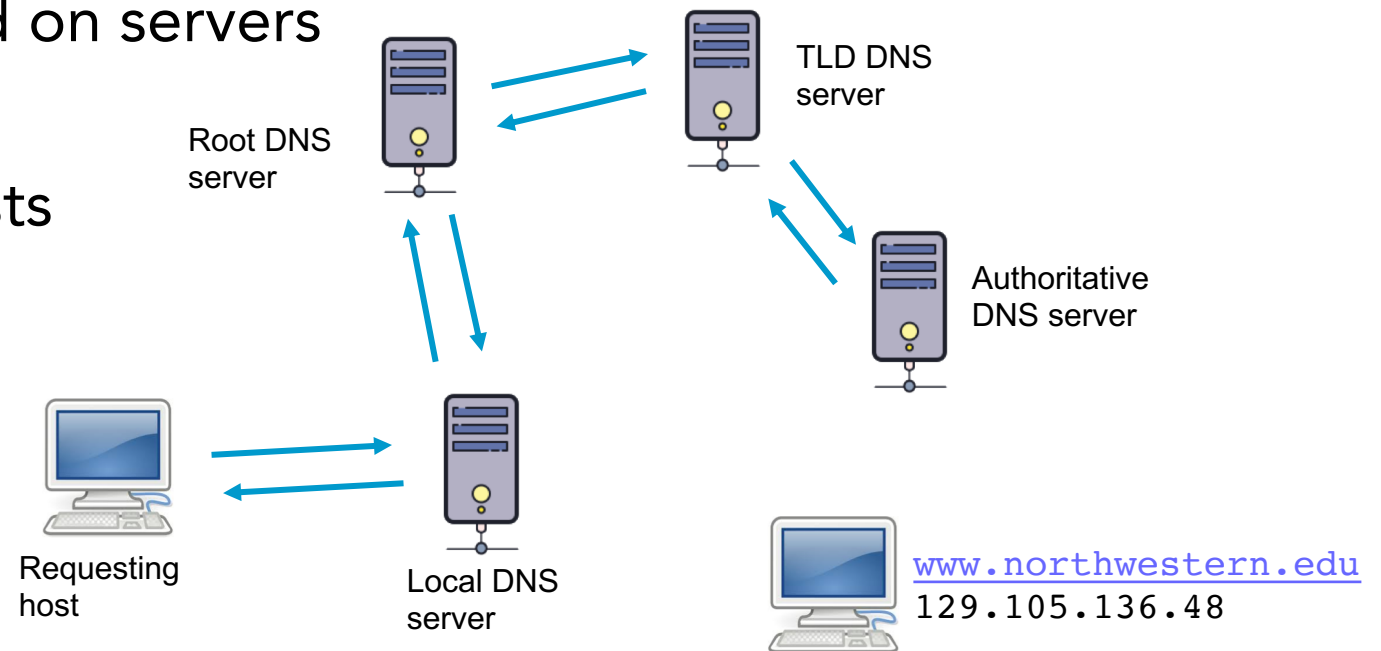
DNS hierarchy in action

- Another important server, not part of the hierarchy: local DNS server (aka, default name server, recursive resolver)
 - Provided by the ISP when a host connects, typically through DHCP
 - Client, the stub resolver, issues queries to the recursive resolver



Iterative / recursive name resolution

- Interactive – client drives the resolution
 - Caching by clients only (a second client's resolution of the same name has to go through the same sequence ...)
 - Potentially costly communication
- Recursive – a name server passes result to next server
 - Higher performance demand on servers
 - More effective caching
 - Reduced communication costs



DNS records

- Each name server implements a zone as set of Resource Records
- Each DNS reply carries one or more RR
- A RR is a tuple containing (name, value, type, class, TTL)
 - TTL is time-to-live of the record in a cache
 - Meaning of name and value depend on the type
 - A – IP address of the host this node represents; the “usual” mapping
 - MX – Canonical name of mail server to handle mail address to this node
 - NS – Name server that implement the represented zone
 - CNAME – Canonical name of the host (alias implemented by a node storing a CNAME record)
 - ...
 - Class allows entities other than the NIC to define useful record types; today the only widely used class is the one used by the Internet - IN

Some other DNS records

- SOA records store information about who created the DNS records and how they should be cached.
- SRV records list server for a given service (generalization of MX):
 - `_sip._udp.columbia.edu` → `laurel.cc.columbia.edu:5060`
 - Tells us which server handles VoIP phone calls to user@columbia.edu.
- TXT records are a generic key-value store
 - DKIM records (stored in a TXT record) store email signature public key:
 - `key1._domainkey.example.com` → `k=rsa;p=J8eTBu224i086iK`
- SPF records list valid outbound mail servers for the domain
- PTR records store reverse DNS records, IP address → hostname

DNS messages

- Query and reply messages, both have the same format
 - First field is a query ID (copied into the reply)
 - Various flags, query/reply, authoritative, recursion desired, ...
 - Number of following fields
 - Questions – name and type of question
 - Response, maybe >1 RRs
 - Records of authoritative servers
 - Additional info like type A record of a MX canonical name server

12 bytes header

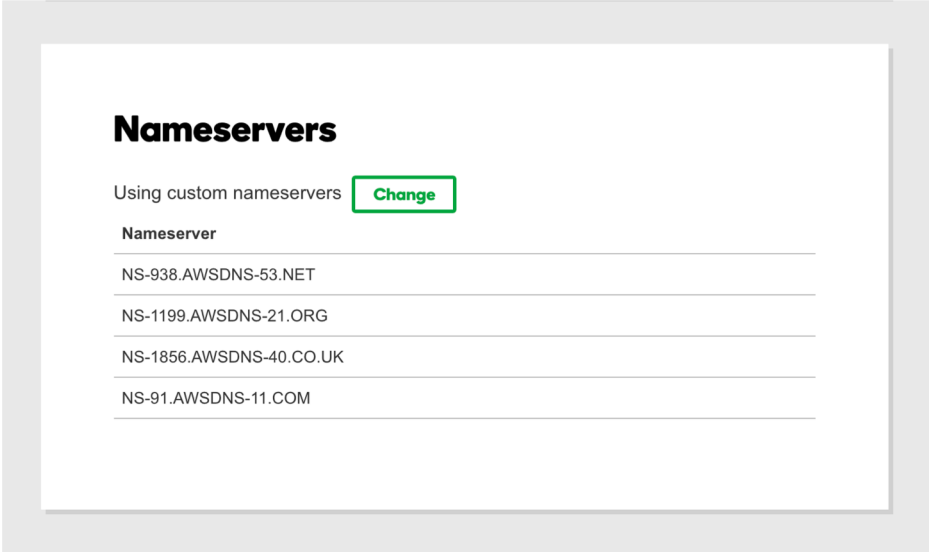
identification	flags
Number of questions	Number of answers RRs
Number of authority RRs	Number of additional RRs
Questions (variable number of questions)	
Answers (variable number of RRs)	
Authority (variable number of RRs)	
Additional information (variable number of RRs)	

DNS scalability

- Scalability through partitioning, replication and caching
 - Caching frequently used records, IP addresses of TLD servers, ...
- Tree sub-divides into zones beginning at the root
 - Each zone could be 1+ domains and sub-domains
- Zone files – the txt file that describes a zone
 - Includes name and address for 2+ authoritative servers and for delegated subdomains
 - Management parameters (e.g. caching) and RR
 - Information in a zone is kept in 2+ name servers (redundancy)
- Any server can cache data from other servers
 - If a non-authoritative server caches data, it notes the TTL

Domain Name Registrars sell domain names

- Eg., GoDaddy, Namecheap, AWS
- Must be accredited (approved) by the TLD registry
 - ICANN appoints an organization to manage each TLD
 - Eg., .com TLD is managed by Verisign, Inc.
 - >330 million .com domains (in 2018)
- Registrar collects your money and nameserver list, then passes the list on to the TLD.



Nameservers

Using custom nameservers [Change](#)

Nameserver

NS-938.AWSDNS-53.NET
NS-1199.AWSDNS-21.ORG
NS-1856.AWSDNS-40.CO.UK
NS-91.AWSDNS-11.COM

The image shows a screenshot of a domain registrar's interface for configuring nameservers. It features a heading 'Nameservers', a status indicator 'Using custom nameservers' with a 'Change' button, and a list of four nameserver addresses: NS-938.AWSDNS-53.NET, NS-1199.AWSDNS-21.ORG, NS-1856.AWSDNS-40.CO.UK, and NS-91.AWSDNS-11.COM.

Great design, now under stress

- Increases in malicious behavior
 - Delegation bottleneck – number of nameservers that need to be compromised to control the domain
 - Worst at the network level – fewer gateways ...
 - 33% at a single gateway (Microsoft DDOS attack 2001)
 - Buggy implementations with known vulnerabilities
 - 2% with a known buffer overflow bug
- Explosion in client population,
 - Zipf-like query distribution – low performance
 - In 2000, 29% of queries took >2"
- Hierarchy implies higher load at the higher levels
 - 2002 DDOS left 9/13 root servers unresponsive

Bottlenecks	All Domains	Top 500
1	0.82%	0.8%
2	78.44%	62.80%
3	9.96%	13.20%
4	4.64%	13%

(old data but the problems remain)

Dynamic DNS (DDNS)

- Static mapping and dynamic assignments via DHCP?
- Dynamic DNS – two different uses and user populations
 - Standard based DNS updates [RFC 2136] typically used with DHCP – authorized DHCP server updates
 - Provider's specific, as a DDNS service registering hosts within its own domain name or a client specific one
 - Custom or DDNS services for security appliances such as cameras
 - Using a simple HTTP-based update API



FreeDNS

 SECUREPOINT
Dynamic DNS Service

Securepoint Dynamic DNS Service:
Free and secure dyndns service



Public DNS Services

- DNS is critical for most Internet applications
- Failure of your local nameserver is very often the source of “Internet outages” you may experience
- Google provides free public DNS servers at 8.8.8.8 and 8.8.4.4
 - It’s a useful backup option if your local DNS resolver is not working
- There’s no free lunch
 - Why does Google and others provide this service for *free*?
 - DNS requests tell them even more about your web surfing habits, and this helps their advertising business

Provider	Primary	Secondary
Google	8.8.8.8	8.8.4.4
Quad9	9.9.9.9	149.112.112.112
OpenDNS Home	208.67.222.222	208.67.220.220
Cloudflare	1.1.1.1	1.0.0.1

DNS over HTTPS (DoH)

- Privacy concerns with DNS
 - The question is always the original
 - Nearly all request is sent in clear (unencrypted) and over UDP
- DoH
 - A protocol for performing DNS resolution via HTTPS
 - Encrypt data between the DoH client and the DoH-based DNS resolver
- Status
 - An RFC 8484, but still a work in progress
 - Support on the resolver side: Google, Cloudflare, Quad9 (IBM), CleanBrowsing, Adguard
 - No native OS support: get it through a browser, on a proxy in your local resolver or machine



Recap

- Nearly all network interactions involve a DNS request
- DNS services beyond name/address mapping (CDN localization)
- Key to performance
 - Visiting most websites require multiple DNS requests – 50% of the top 100k Alexa sites → >20 DNS queries
- ... and privacy
 - Most everything you do in the Internet starts with a DNS lookup
- In constant evolution