ECE 356/COMPSI 356
Computer Network Architecture

DNS and Content Distribution

Monday November 18th, 2019

Duke UNIVERSITY

Recap

- Previous lecture:
  - ➢Multimedia communications
  - ➢Internet Quality of Service (QoS)
    - DiffServ
- Readings for this lecture: **PD 9.3.1, 9.4.3**

2

Duke UNIVERSITY

Lecture Overview

- Domain Name System
  - ➢Introduction
  - ➢DNS servers
  - ➢DNS records and messages
  - ➢Inserting records into DNS
- Content Distribution Networks

Duke UNIVERSITY

Previously Saw DNS in Lecture on UDP

- Map an easy-to-remember name to an IP address
  - ➢ Without DNS, to send an IP packet, we'd have to remember
    - 66.102.7.99
    - 64.236.24.28
  - ➢ With DNS
    - **www.google.com**→ 66.102.7.99
    - **www.cnn.com**→ 64.236.24.28
    - Name does not have to change if the IP address changes

4

Duke UNIVERSITY

DNS Runs on UDP

- Response is needed quickly
  - First thing that happens after you enter a URL into the browser
  - Incurs a delay even before connection establishment starts
- Runs over UDP, uses port 53
- IP address cached at a *nearby server*
  - Helps to reduce DNS traffic on the network
  - Helps to reduce delay

5

Duke UNIVERSITY

Domain Name System: An Introduction

- Worldwide distributed directory service
  - "**Phonebook for the Internet**"
- *Distributed database* implemented in hierarchy of many *name servers*
- *Application-layer protocol:* hosts, name servers communicate to *resolve* names (address/name translation)
  - Note: core Internet function, implemented as application-layer protocol

6

Duke UNIVERSITY

Why Does DNS Need to be Distributed? (1/2)

- 1.01 **billion** Internet hosts in the DNS

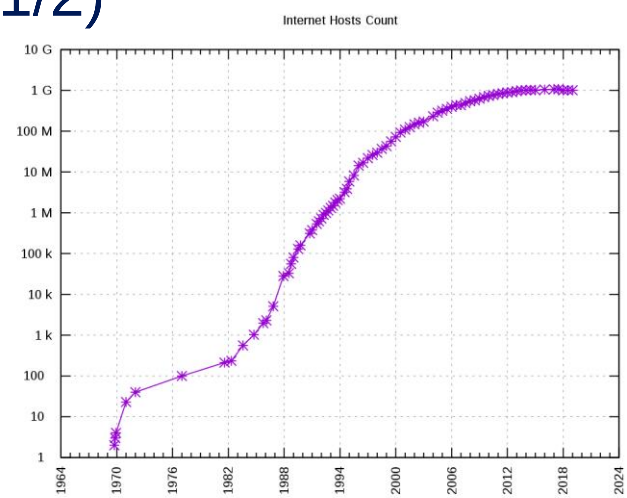

Duke UNIVERSITY

Why Does DNS Need to be Distributed? (2/2)

- What if we had one DNS server?
  - ➢Single point of failure
  - ➢Unmanageable traffic volume
  - ➢Unmanageable database maintenance needs
    - Frequent updates
  - ➢Far from many hosts

8

Duke UNIVERSITY

RTT Latency: Examples

Verizon Enterprise Latency Statistics (ms)												
	2019						2018					
	June	May	April	March	February	January	December	November	October	September	August	July
Trans Atlantic (90.000)	73.833	69.986	69.950	69.930	69.965	69.888	70.531	70.965	70.376	70.529	70.489	70.423
Europe (30.000)	10.978	11.706	11.234	10.592	11.099	11.478	10.954	10.070	11.215	11.257	11.239	11.237
North America (45.000)	30.927	31.352	31.531	33.523	33.782	36.083	36.084	39.243	38.468	37.999	37.618	35.244
Intra-Japan (30.000)	-	11.221	11.932	13.093	12.910	12.761	12.616	12.894	11.704	13.332	12.674	10.872
Trans Pacific (160.000)	134.714	99.336	99.320	99.238	99.237	99.242	99.240	99.250	103.168	102.561	101.381	101.369
Asia Pacific (125.000)	90.206	85.806	85.201	85.119	86.840	86.726	98.990	87.173	85.007	107.209	84.737	86.923
Latin America (140.000)	93.080	90.968	88.450	87.782	119.633	-	-	-	-	-	125.605	123.193
EMEA to Asia Pacific (250.000)	122.317	144.462	119.350	119.239	118.699	116.281	115.876	115.030	113.885	113.836	120.111	119.401

- 100 ms feels instantaneous to web users

9

Duke UNIVERSITY

Nslookup Program

- *gethostbyname()* in the first lab
- *nslookup* <website name>

```
C:\Users\Maria>nslookup durhamnc.gov
Server:  UnKnown
Address:  fe80::4a5b:39ff:fee8:4e5

Non-authoritative answer:
Name:    durhamnc.gov
Address:  208.90.188.133
```

```
C:\Users\Maria>nslookup nytimes.com
Server:  UnKnown
Address:  fe80::4a5b:39ff:fee8:4e5

Non-authoritative answer:
Name:    nytimes.com
Addresses:  151.101.65.164
          151.101.129.164
          151.101.1.164
          151.101.193.164
```

10

Duke UNIVERSITY

DNS: Additional Services

- Host *aliasing*
  - relay1.west-coast.enterprise.com, enterprise.com, www.enterprise.com
  - duke.edu, www.duke.edu
- Load distribution
  - Many web servers are replicated
  - IP addresses returned by DNS are rotated

11

Duke UNIVERSITY

Before There Was DNS ….

- Back in ARPANET days, already wanted memorable host names
  - Used *HOSTS.TXT* file maintained on a host at SRI Network Information Center (NIC)
  - Downloaded a HOSTS.TXT from a central server via FTP
  - Had to call SRI during office hours to get a new entry into HOSTS.TXT

Duke UNIVERSITY

DNS: Key Points to Remember

- Invoked every time URL is entered into a browser
  - Needs to be *fast*
- *Distributed database* implemented in hierarchy of many *name servers*

13

Duke UNIVERSITY

Lecture Overview

- Domain Name System
  - Introduction
  - **DNS servers**
  - DNS records and messages
  - Inserting records into DNS
- Content Distribution Networks

Duke UNIVERSITY

DNS Servers

- Large number of servers
  - Organized in a hierarchical fashion
  - Distributed around the world
  - **No single server has all the information**
- Three server classes:
  - Root servers
  - Top-level domain (TLD)
  - Authoritative servers

DNS: A Distributed, Hierarchical Database

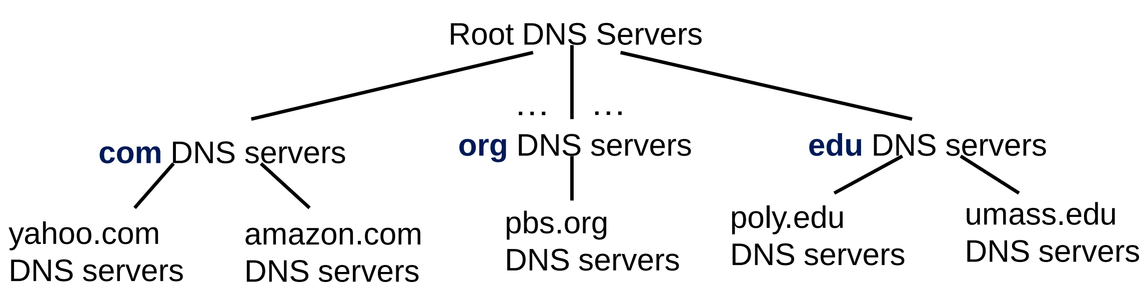

Client wants IP for www.amazon.com; 1st approximation:

- Client queries root server to find **.com** DNS server
- Client queries **.com** DNS server to get **amazon.com** DNS server
- Client queries **amazon.com** DNS server to get IP address for www.amazon.com

DNS Resolver: An Example

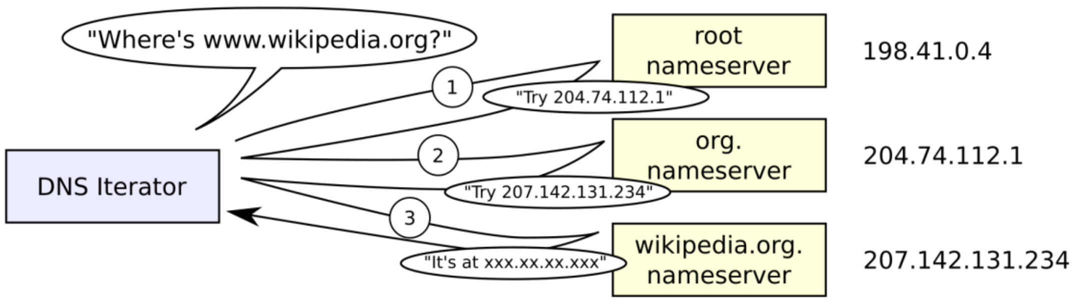

- In practice, caching is used to offload root servers
  - Only a small fraction of requests go to them

Root Name Servers

- Over 400, managed by 13 different organizations
- Know how to find the authoritative name servers for all top-level zones

Addresses of Root Servers

- Hard-coded into every DNS resolver

A.ROOT-SERVERS.NET.	(VeriSign, Dulles, VA)	198.41.0.4
B.ROOT-SERVERS.NET.	(ISI, Marina Del Rey CA)	192.228.79.201
C.ROOT-SERVERS.NET.	(Cogent Communications)	192.33.4.12
D.ROOT-SERVERS.NET.	(University of Maryland)	128.8.10.90
E.ROOT-SERVERS.NET.	(Nasa Ames Research Center)	192.203.230.10
F.ROOT-SERVERS.NET.	(Internet Systems Consortium)	192.5.5.241
G.ROOT-SERVERS.NET.	(US Department of Defense)	192.112.36.4
H.ROOT-SERVERS.NET.	(US Army Research Lab)	128.63.2.53
I.ROOT-SERVERS.NET.	(Stockholm, Sweden)	192.36.148.17
J.ROOT-SERVERS.NET.	(Herndon, VA)	192.58.128.30
K.ROOT-SERVERS.NET.	(London, United Kingdom)	193.0.14.129
L.ROOT-SERVERS.NET.	(IANA, Los Angeles)	198.32.64.12
M.ROOT-SERVERS.NET.	(WIDE, Tokyo)	202.12.27.33

Top-Level Domain (TLD) Servers

- Responsible for *com*, *org*, *net*, *edu*, *aero*, *jobs*, *museums*, and all top-level country domains, e.g.: *uk, fr, ca, jp*
- Network Solutions maintains servers for .com TLD
- Educause for .edu TLD
- Provide the IP addresses for authoritative DNS servers

Authoritative DNS Servers

- Each organization with publicly accessible hosts must provide publicly accessible DNS records
- Organization's own DNS server(s), providing authoritative hostname to IP mappings for organization's named hosts
- Can be maintained by organization or service provider (for a fee)
- Universities, large companies usually implement and maintain their own primary and secondary authoritative DNS servers

21

Duke UNIVERSITY

Local DNS Servers (1/2)

- Does not strictly belong to a hierarchy
- Each ISP (residential ISP, company, university) has one
  - Also called "*default name server*"
- When host makes DNS query, query is sent to its local DNS server
  - Has local cache of recent name-to-address translation pairs (but may be out of date!)
  - Acts as proxy, forwards query into hierarchy

22

Duke UNIVERSITY

Local DNS Servers (2/2)

- Close to the host
  - Same LAN, or only a few hops away
- Dramatically reduces latency, by both:
  - Proximity
  - Caching

Iterative Server Queries

- Host at cis.poly.edu wants IP address for gaia.cs.umass.edu
- Contacted server replies with name of server to contact
- “I don’t know this name, but ask this server”

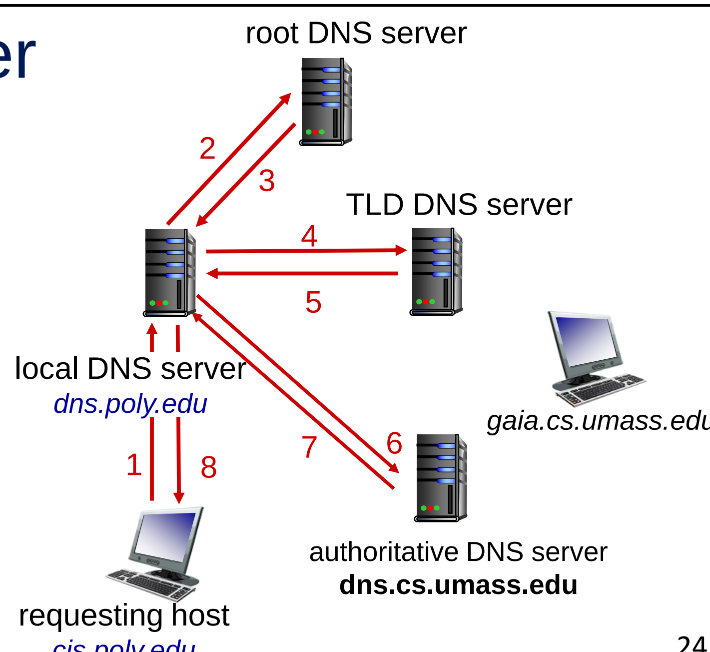

Recursive Server Queries

- Puts burden of name resolution on contacted name server
- Heavy load at upper levels of hierarchy?

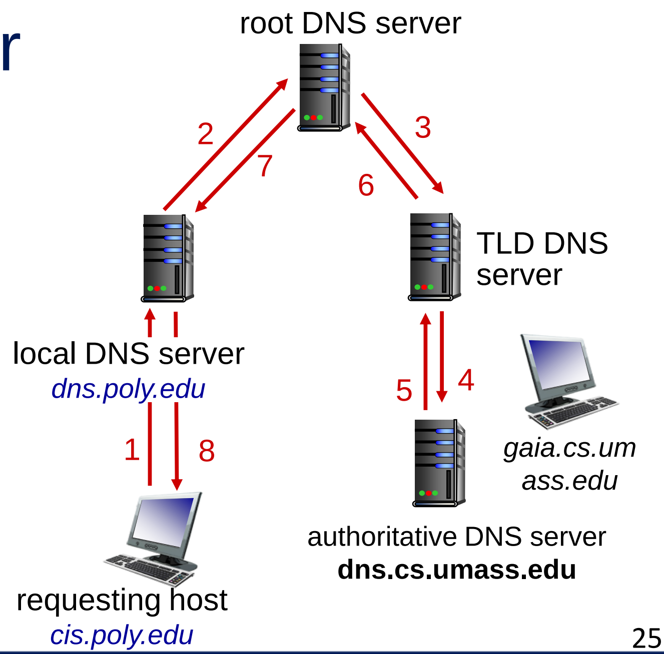

25

DNS Caching (1/2)

- Once (any) name server learns mapping, it *caches* mapping
  - ➢ Cache entries timeout (disappear) after some time (TTL)
  - ➢ TLD servers typically cached in local name servers
    - Thus root name servers not often visited
- Cached entries may be *out-of-date*
  - ➢ *Best effort* name-to-address translation!
  - ➢ if name host changes IP address, may not be known Internet-wide until all TTLs expire

DNS Caching (2/2)

- If an entry is sent from a cache, the reply from the server is marked as "*unauthoritative*"

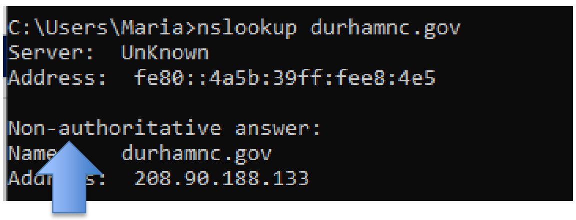

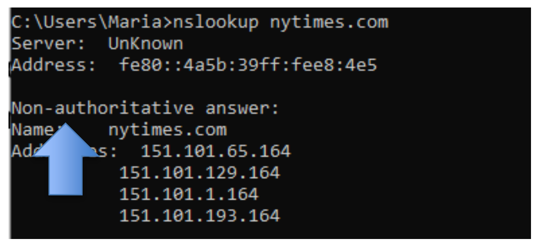

27

Duke UNIVERSITY

DNS Servers: Key Points to Remember

- Rely on a 3-level server hierarchy, plus local servers
  - Root servers
  - Top-level domain servers (TLD)
  - Authoritative servers
- Servers queries *iteratively* or *recursively*
- Servers cache entries to operate efficiently
  - "Best-effort" name to address translation

28

Duke UNIVERSITY

Lecture Overview

- Domain Name System
  - Introduction
  - DNS servers
  - **DNS records and messages**
  - Inserting records into DNS
- Content Distribution Networks

Duke UNIVERSITY

Resource Records

- Distributed database stores Resource Records (RRs)
- (**Name, Value, Type, TTL**)
  - TTL: how long the record should be cached

Duke UNIVERSITY

Resource Records: Name, Value, Type (1/2)

- Type A: Name: hostname, Value: IP address
  - Standard hostname-to-IP mapping
  - Example: (relay1.bar.foo.com, 145.37.93.126,A)
- Type NS: Name: domain, Value: hostname of an authoritative DNS server
  - Example: (foo.com, dns.foo.com, NS)

Duke UNIVERSITY

Resource Records: Name, Value, Type (2/2)

- Type CNAME: Name: alias hostname, Value: canonical hostname
  - Example: (foo.com, relay1.bar.foo.com, CNAME)
- Type MX: Name: alias mail server name, Value: canonical mailserver name
  - Example: (foo.com, mail.bar.foo.com, MX)

Duke UNIVERSITY

DNS Messages

- *Query* and *reply* messages, both with the same message format
- Header: first 12 bytes
- Data: the rest

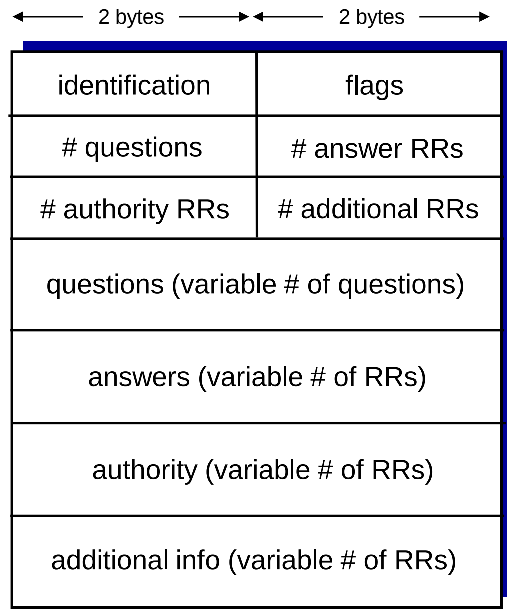

33

Duke UNIVERSITY

DNS Messages: Header

- ID of request/reply
- Flags:
  - ➢ Query/reply
  - ➢ Authoritative
  - ➢ Recursion desired
  - ➢ Recursion available
- # of the types of data to follow

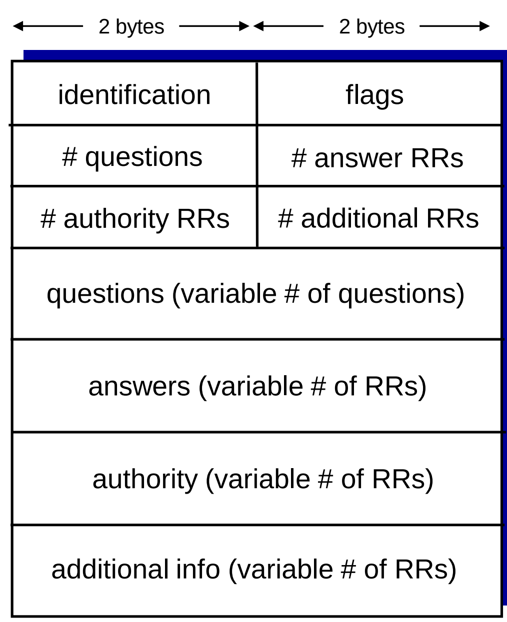

34

Duke UNIVERSITY

DNS Messages: Data

- Questions: name, type field of a query
- Answers: resource records
  - Can return multiple RRs in an answer
- Authority: records for authoritative servers

2 bytes	2 bytes
identification	flags
# questions	# answer RRs
# authority RRs	# additional RRs
questions (variable # of questions)	
answers (variable # of RRs)	
authority (variable # of RRs)	
additional info (variable # of RRs)	

35

Duke UNIVERSITY

DNS Records and Messages: Key Points to Remember

- Relatively simple *Resource Record* structure:
  - 4-tuple, (**Name, Value, Type, TTL**)
- One message format, for query and reply messages
  - Can return multiple RRs in an answer

```
C:\Users\Maria>nslookup nytimes.com
Server:  UnKnown
Address:  fe80::4a5b:39ff:fee8:4e5

Non-authoritative answer:
Name:    nytimes.com
Addresses:  151.101.65.164
          151.101.129.164
          151.101.1.164
          151.101.193.164
```

36

Duke UNIVERSITY

Lecture Overview

- Domain Name System
  - ➢Introduction
  - ➢DNS servers
  - ➢DNS records and messages
  - ➢**Inserting records into DNS**
- Content Distribution Networks

Duke UNIVERSITY

DNS Registrars

- Internet Corporation for Assigned Names and Numbers (ICANN) accredits ***registrars***
  - ➢ Registrar: commercial entity that verifies the uniqueness of a domain name and enters it into the DNS database, for a fee
- Only one company until 1999, now many
- Full registrar list: https://www.internic.net/
  - ➢ Examples: VeriSign, GoDaddy

38

Duke UNIVERSITY

Inserting Records into DNS

- Provide names, IP addresses of authoritative name server (primary and secondary)
  - Example: want to register networkutopia**.com**
  - Authoritative DNS servers: dns1.networkutopia.com, 212.2.212.1, dns2.networkutopia.com, 212.2,212.2
- Registrar inserts RRs into **.com** TLD server:
  (networkutopia.com, dns1.networkutopia.com, NS)
  (dns1.networkutopia.com, 212.212.212.1, A)
- Create authoritative server type A record for www.networkuptopia.com

39

Duke UNIVERSITY

DNS: Key Points to Remember

- Distributed hierarchical database
  - Hierarchy: 3 levels of servers
  - Additional *local* DNS servers
- Uses caching to reduce client latency and network loads

40

Duke UNIVERSITY

Lecture Overview

- Domain Name System
  - ➢Introduction
  - ➢DNS servers
  - ➢DNS records and messages
  - ➢Inserting records into DNS
- **Content Distribution Networks**

A Traditional Web Application

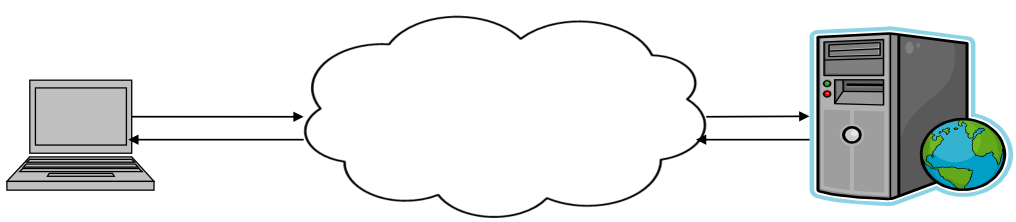

- HTTP request http://www.cs.duke.edu
- A DNS lookup on www.cs.duke.edu returns the IP address of the web server
- Requests are sent to the web site
- If DNS lookup returned a different IP address, request would be sent to a different web server, without user knowing

Some Domain Names Are Handled by Content Distribution Networks (CDNs)

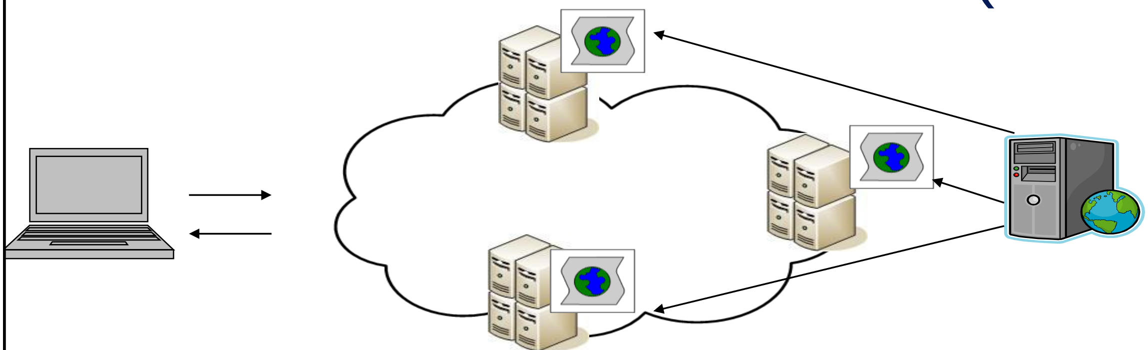

- A single provider that manages multiple replicas
- A client obtains content from a close replica

CDNs

- Akamai, AWS CloudFront, Fastly, …
- Use globally distributed "Points of Presence" (POPs), also called "edge servers":
  - ➢ Akamai: ~200,000, AWS CloudFront: 100 POPs

What Problems Does CDN Solve?

- Enhance web performance: reduce latency
- Reduce network load
- *Flash crowd* may overwhelm a server and the access network

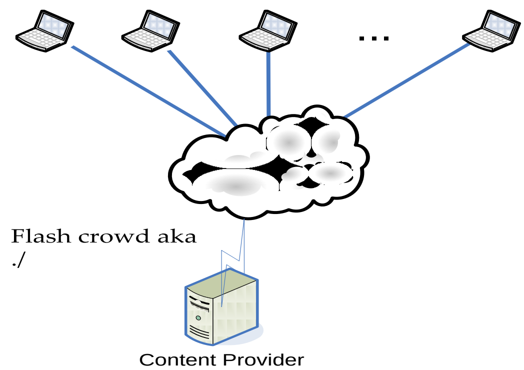

Duke UNIVERSITY

CDN and DNS

- *DNS redirection*
- Geographical lookup: based on DNS resolver's IP address, return an IP address for an edge server that is closest to the area
  - DNS request originating in Virginia and California will be given different IP addresses
- Can also load-balance

Duke UNIVERSITY

Accessing Content in a CDN

- CDNs are *caches*

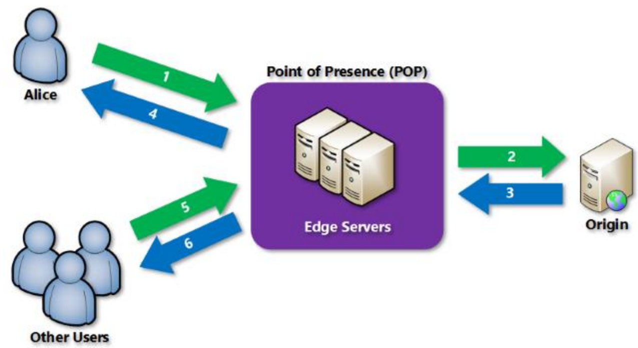

Content Delivery Networks: Key Points to Remember

- CDNs replicate content to multiple local servers
  - ➢ Mainly to reduce the total content loading latency
- Use *DNS redirection* to point clients to most appropriate edge server

Lecture Summary

- Domain Name System
  - ➢Introduction
  - ➢DNS servers
  - ➢DNS records and messages
  - ➢Inserting records into DNS
- Content Distribution Networks

Duke UNIVERSITY

Next Lecture

- Guest lecture:
  - ➢On-campus network infrastructures and operational challenges

50

Duke UNIVERSITY