

# 9/8/21 - Domain Name System

Max Layer

CSC-6730 Advanced Networking

# Research Question

- "Is Internet routing undermining DNS anycast benefits?"
  - DNS
    - Design goals
    - Architecture
  - Anycast
    - Internet routing
      - BGP
  - Tie it together
    - DNS Anycast
      - Performance
      - Problems
    - Impact of internet routing
      - Undermining?
        - Why? Why not?



# Ancient Times: A Pre-DNS World

- ARPANET
- HOSTS.TXT
  - 1.2.3.4 name
- A *lot* more hosts...
- Aside: Still a backup!



# Design of DNS

- Design Assumptions: [1]
  - At least same information as HOSTS.TXT
  - Ability to maintain distributed database
  - No obvious data size limits
  - Connect as many environments as possible
  - *Tolerable* performance



# Design of DNS

- Derivative Constraints: [1]
  - Only worth it if extensible
    - Independent from network topology
    - Encapsulate other name spaces
  - Avoid forcing architectures on users
    - Allow users to customize implementation



# DNS Architecture

- 2 Main Components
  - Name Servers
  - Resolvers
- Name space
  - Tree
    - Nodes
      - Label
      - Data

Example: csc.tntech.edu



# Design of DNS - Success?

- Design Assumptions:
  - ~~At least same information as HOSTS.TXT~~
  - ~~Ability to maintain distributed database~~
  - ~~No obvious data size limits~~
  - ~~Connect as many environments as possible~~
  - **Tolerable performance**
- Derivative Constraints:
  - ~~Only worth it if extensible~~
    - ~~Independent from network topology~~
    - ~~Encapsulate other name spaces~~
  - ~~Avoid forcing architectures on users~~
    - ~~Allow users to customize implementation~~



# DNS Performance


- Room for improvement
- What if you want:
  - Load Balancing
  - DDoS Protection
  - Low Latency
- More name servers, more localized
  - Expensive
  - Time Consuming





# What is Anycast?

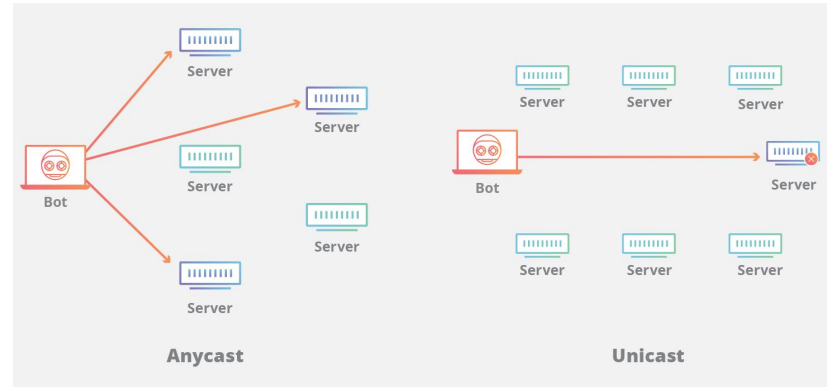
- One **IP**, many **hosts**
  - How is this different from DNS?
    - DNS: One **name**, many **IPs**
    - Who decides?
      - DNS Resolver VS ISP/BGP
- BGP: Border Gateway Protocol
  - Route between Autonomous Systems
    - Typically, shortest routes
  - TL;DR: ISP level routing
  - More next week!



Anycast Visualization [6]

# Why Use Anycast?


- Duplicate service, different geographic location
  - Redundancy
  - Latency
  - 'Free' load balancing/DDoS resiliency
    - Cheaper!
- What does this sound perfect for?



Anycast DDoS Protection [5]

# Who Uses Anycast?

- Content Delivery Networks
- Root DNS Servers
  - The *original* CDN
  - CDN for HOSTS.TXT
  - DNS Anycast
    - Same IP address
    - Different locations
    - Ex: 1.1.1.1



Anycast Visualization [6]

# Using Anycast For Your CDN

- All you need to do:
  - Have the same endpoints you would have anyways
  - Endpoints share the same IP
    - Instead of same domain name
  - Hope the user's request goes to the best one
    - Wait, *what?*



# Anycast CDN Performance [2]

- Anycast-based CDN
- Measure Performance
  - Passively (Logs)
  - Actively
    - Automatically query for random frontends
      - One via anycast
      - Three via unicast
    - Record the latency of requests



# Anycast CDN Performance [2]

- Results:
  - ~20% of clients get suboptimal frontend
  - Persistently poor performance (over time)
  - Front end affinity




Figure 7: The cumulative fraction of clients that have changed front-ends at least once by different points in a week

[2]




Figure 5: Daily poor-path prevalence during April 2015 showing what fraction of client /24s see different levels of latency improvement over anycast when directed to their best performing unicast front-end.

[2]

# Anycast Problems

- Lack of control
  - Free load balancing, but to where?
  - Reliance on internet routing
- BGP doesn't guarantee "best" path
  - Review: End to end argument
    - Application layer VS network layer error correction
    - Lower layer had enough information to solve the problem
      - **Most** of the time
  - BGP advertisement path *length*
    - Least path changes  $\neq$  fastest route
      - Highways
    - Big geographical leaps happen often
    - What if the shortest path is to an overloaded endpoint?



# Improving Anycast Performance [2]

- BGP lacks information to make best decisions
- More Information inside of BGP
  - Advertise more data on each AS
    - Geography?
    - Load?
  - More information allows more accurate decisions to be made
- Hybrid anycast and DNS-based redirection
  - Use DNS for clients with poorer anycast performance
  - Predict better frontends for clients using DNS information
    - EDNS - Prediction based on client's prefix
    - LDNS - Local DNS maps latency, chooses its best frontend

[2]




Figure 9: Improvement over anycast from making LDNS or ECS-based decisions with prediction using 25th percentile prediction metric. Negative x-axis values show where anycast was better than our prediction. Values at 0 show when we predicted anycast was the best performing. Positive x-axis values show our improvement.



# Research Question

- "Is Internet routing undermining DNS anycast benefits?"
  - **YES...** *to a degree*
    - Internet routing (BGP) typically prioritizes least amount of network hops
      - Not necessarily fastest route (Undermines latency)
      - Not necessarily least busy route (Undermines load balancing)
- Is Internet routing undermining *all* DNS anycast benefits?
  - **NO**
    - Still cheaper than managing load balancing DNS servers
    - *Most* of the time it works as intended
      - Reasonable trade off?
      - Used because it's **good enough**



# Questions?

## References:

- [1] P. Mockapetris, K. Dunlap, “Development of Domain Name System”, ACM SIGCOMM Computer Communication Review, 1988
  - [2] Matt Calder, et. al. 2015. Analyzing the Performance of an Anycast CDN. IMC 2015
  - [3] Zhihao Li, Dave Levin, Neil Spring, and Bobby Bhattacharjee. 2018. Internet anycast: performance, problems, & potential. In Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication (SIGCOMM '18)
  - [4] L. Colitti, E. Romijn, H. Uijterwaal, and A. Robachevsky. Evaluating the effects of anycast on DNS root name servers. In RIPE document RIPE-393, 2006
  - [5] <https://www.cloudflare.com/learning/dns/what-is-anycast-dns/>
  - [6] <https://en.wikipedia.org/wiki/File:Anycast-BM.svg>
- 