

# Anycasting the DNS

---

John Crain

Internet Corporation for Assigned Names and Numbers



# What is Anycast?

- ▶ “Anycast is a network addressing and routing scheme whereby data is routed to the "nearest" or "best" destination as viewed by the routing topology.” – Wikipedia
  - ▶ unicast – IP address is a unique interface
  - ▶ multicast – Multiple locations at same IP address all receiving packets.
  - ▶ anycast – IP address points to multiple locations but only one visible.

# Limitations

- ▶ Because routing information can change!
  - ▶ UDP protocols
  - ▶ Short lived TCP sessions
    - ▶ This sounds a lot like DNS packets :)
- ▶ Not good for longer TCP sessions
  - ▶ Think of a web session that can last seconds or videostreaming that can last even longer!

# Why do it?

- ▶ Adds Capacity
  - ▶ 2 x 100Mb is more than 1 x 100Mb
- ▶ Localizes traffic
  - ▶ Means closer, hence faster replies
  - ▶ During attacks it also localizes problems
- ▶ It's FUN!!!!

# Is it hard?

- ▶ Not really but there are some considerations
- ▶ Managing the anycast instances takes some thought
- ▶ Need multiple IP address ranges...
- ▶ More on that.....

# Implementing Anycast

# Separate DNS server into it's own network

- ▶ Get a dedicated routable /24 of IPv4 space and a range of IPv6 space
- ▶ Also get a dedicated Autonomous System Number (ASN)
  - ▶ This is necessary for creating a unique routing entity
    - ▶ ICANN L-ROOT (NET6-2001-500-3-1)  
2001:0500:0003:0000:0000:0000:0000:0000 –  
2001:0500:0003:FFFF:FFFF:FFFF:FFFF:FFFF
    - ▶ ICANN L-ROOT (NET-199-7-83-0-1) 199.7.83.0 – 199.7.83.255
    - ▶ ASNumber: 20144

# A name server

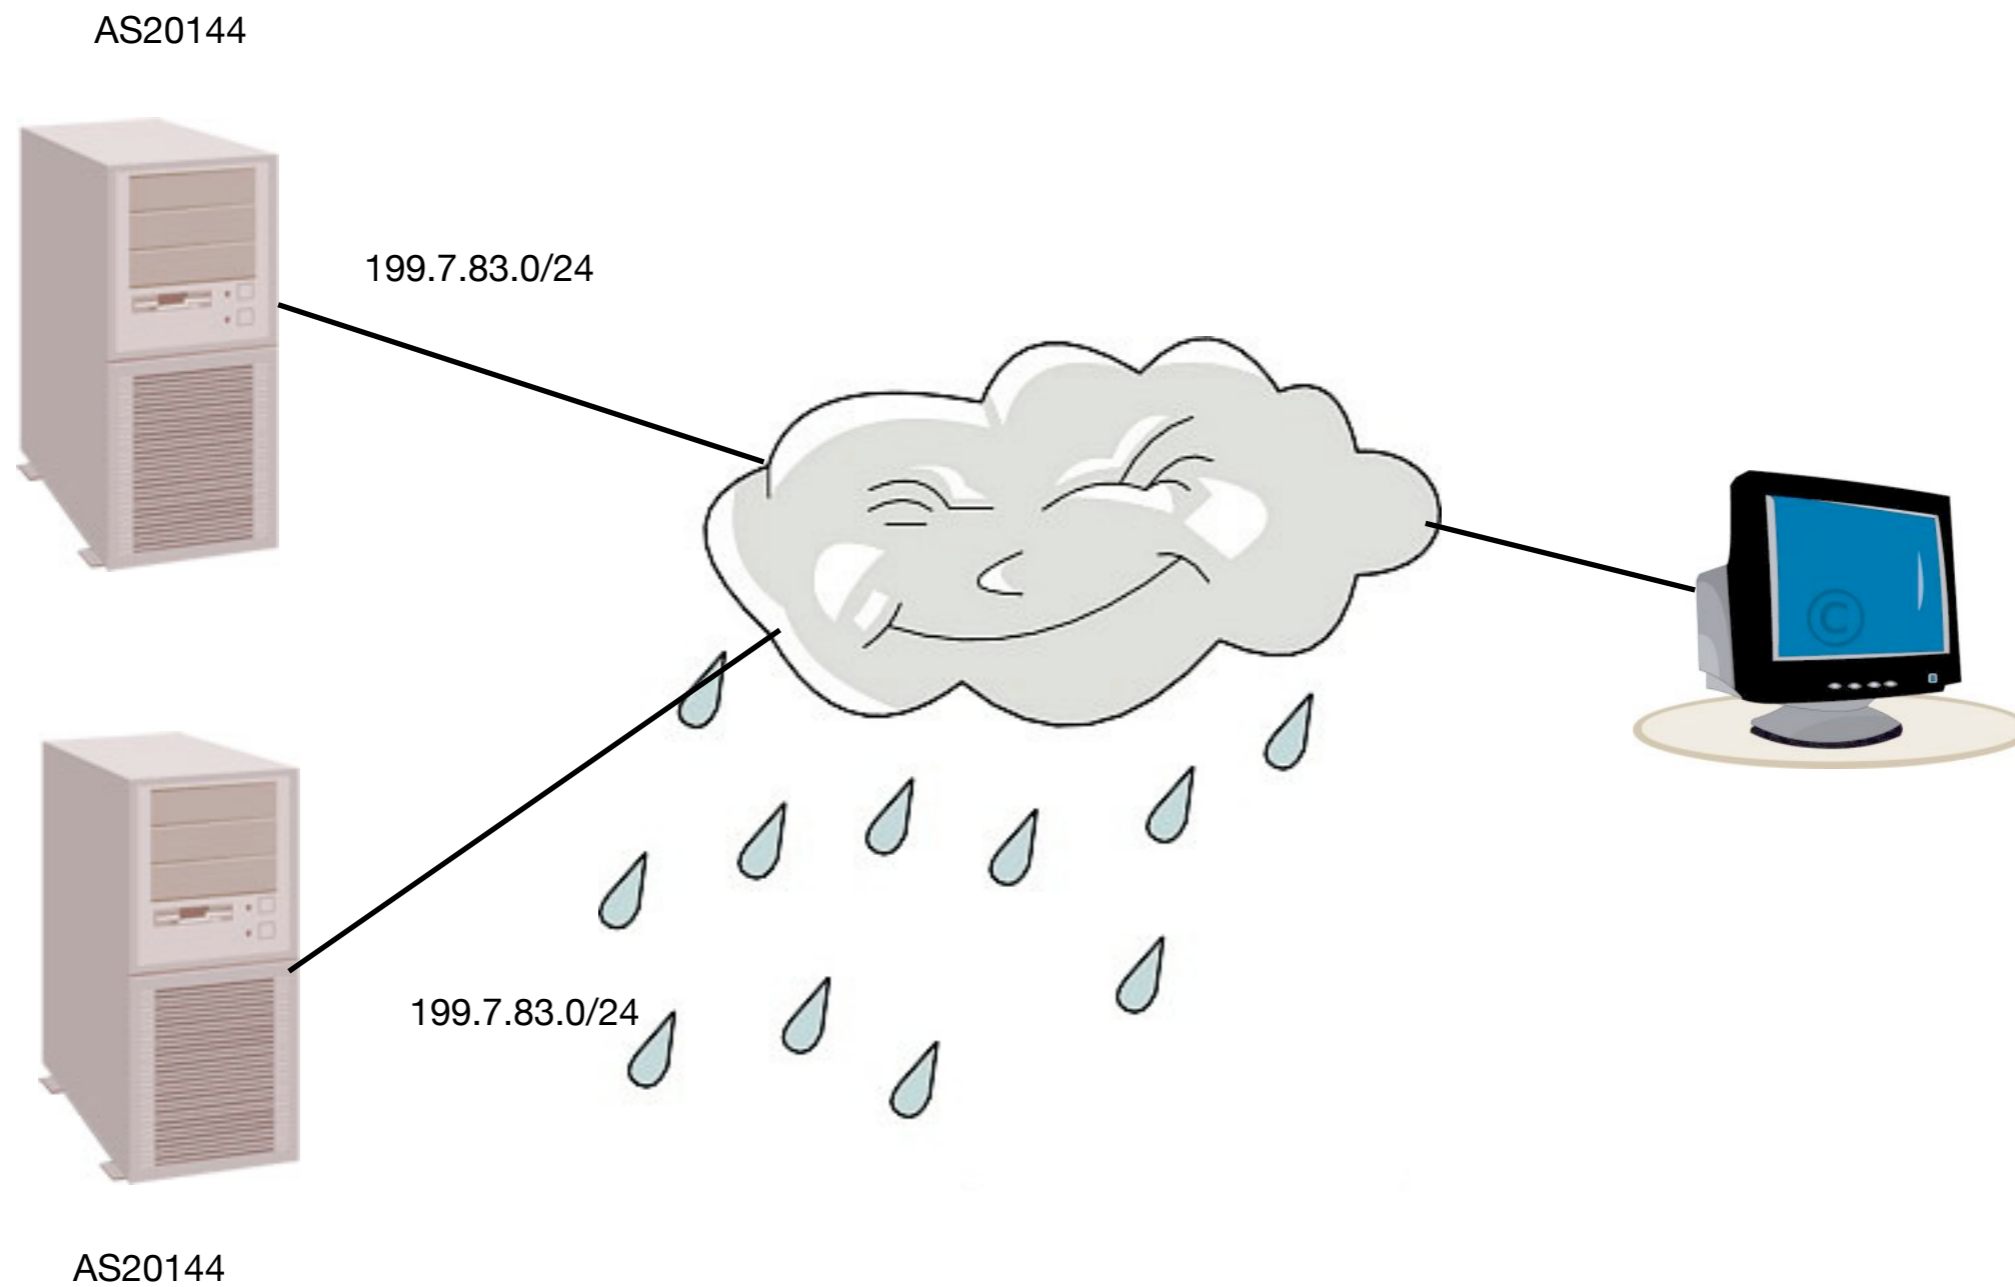
- ▶ Has an IP address on which it answers DNS queries. `l.root-servers.net > 199.7.83.42`
  - ▶ (Also `2001:0500:0003::42`)



- ▶ Announce `199.7.83.0/24`
  - ▶ Normal = unicast




# Anycast



# What are the issues?

- ▶ From any point on the Network I only see one route to the system at anytime
- ▶ How do I as administrator get to a specific instance?
- ▶ Administrative/Management IP addresses!
  - ▶ These are used to get to a specific host

# Administation addresses



# Important factors to think about

- ▶ If one of the systems stop answering (NSD/BIND stops?)
  - ▶ Make sure that the route gets withdrawn
- ▶ Have a “Backdoor” incase your system is not reachable.
- ▶ Have someone local to the machines who can pull the plug in an emergency

# Who does this?

- ▶ Various TLD operators
- ▶ The root-servers
  - ▶ Technical Limitation of 13 root-server names
    - ▶ a.root-servers.net – m.root-servers.net
    - ▶ In a unicast world this also meant 13 locations.
    - ▶ In an anycast world this looks different
- ▶ <http://www.root-servers.org/~wnagele/map/version2.html>
- ▶ <http://www.icann.org/en/maps/root-servers.htm>

# Thankyou for your attention!

---

John Crain  
[john.crain@icann.org](mailto:john.crain@icann.org)

