




# IP Anycast

Prof. Anja Feldmann, Ph.D.

Dr. Oliver Gasser




# Introduction



## Terminology

- Unicast: Communication between a single source and a single destination




# Introduction



## Terminology

- **Unicast:** Communication between a single source and a single destination
- **Multicast:** Communication between a single source and multiple destinations
  - Useful for streaming and conferencing applications
  - Heavily used in IPv6 for signaling




# Introduction



## Terminology

- **Unicast:** Communication between a single source and a single destination
- **Multicast:** Communication between a single source and multiple destinations
  - Useful for streaming and conferencing applications
  - Heavily used in IPv6 for signaling
- **Anycast:** Communication between a single source and **any** destination
  - Useful for communication to the closest server of a group of servers



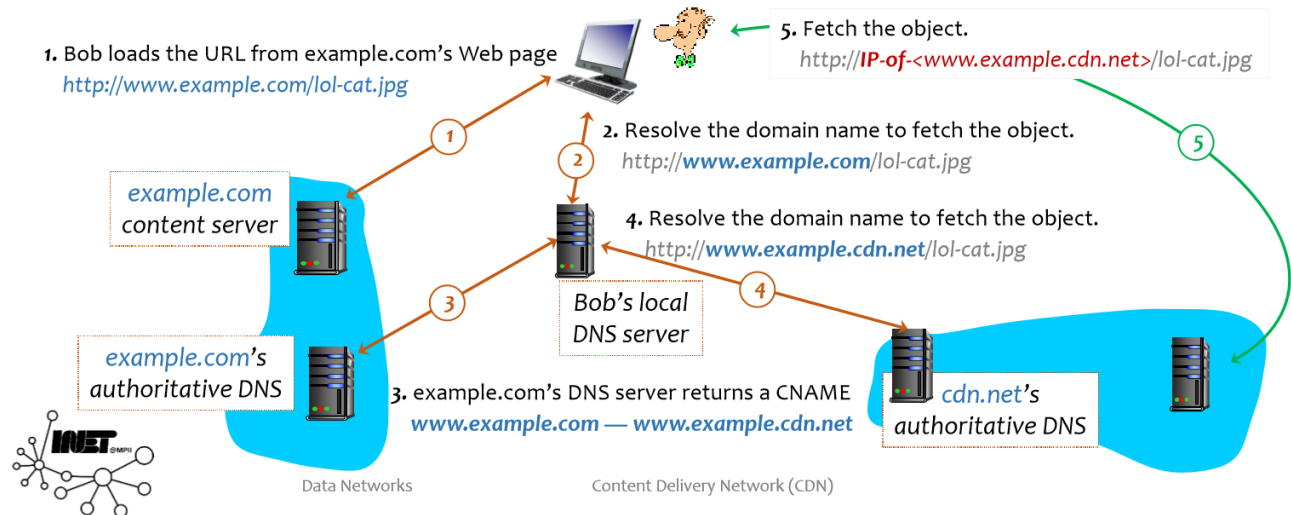
# Recall: CDN Lecture




## Serving content via CDN



- User requests image at URL <http://www.example.com/lol-cat.jpg>
- Image stored in CDN at <http://www.example.cdn.net/12lol34cat56>



17



Data Networks

Anycast

# Recall: CDN Lecture


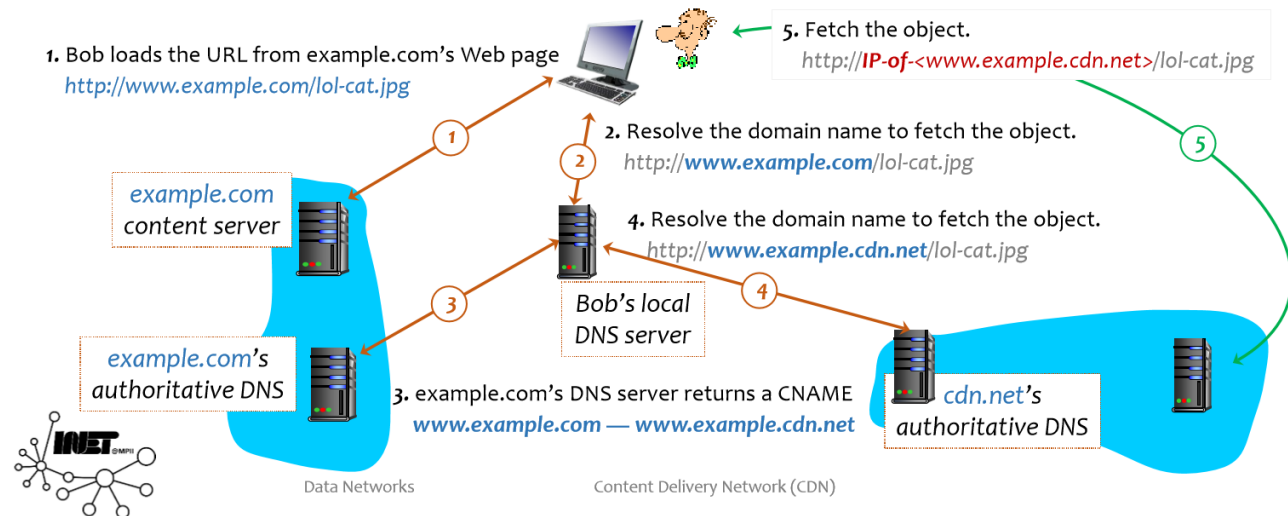


## Serving content via CDN



- User requests image at URL <http://www.example.com/lol-cat.jpg>
  - Image stored in CDN at <http://www.example.cdn.net/12lol34cat56>


- DNS-based load-balancing
  - DNS CNAME records
- IP Anycast
  - BGP



# IP Anycast




- IP packets are sent to a specific destination IP address



# IP Anycast



- IP packets are sent to a specific destination IP address
- How can we send a packet to an anycast destination IP address?
  - No difference to regular unicast IP address





# IP Anycast




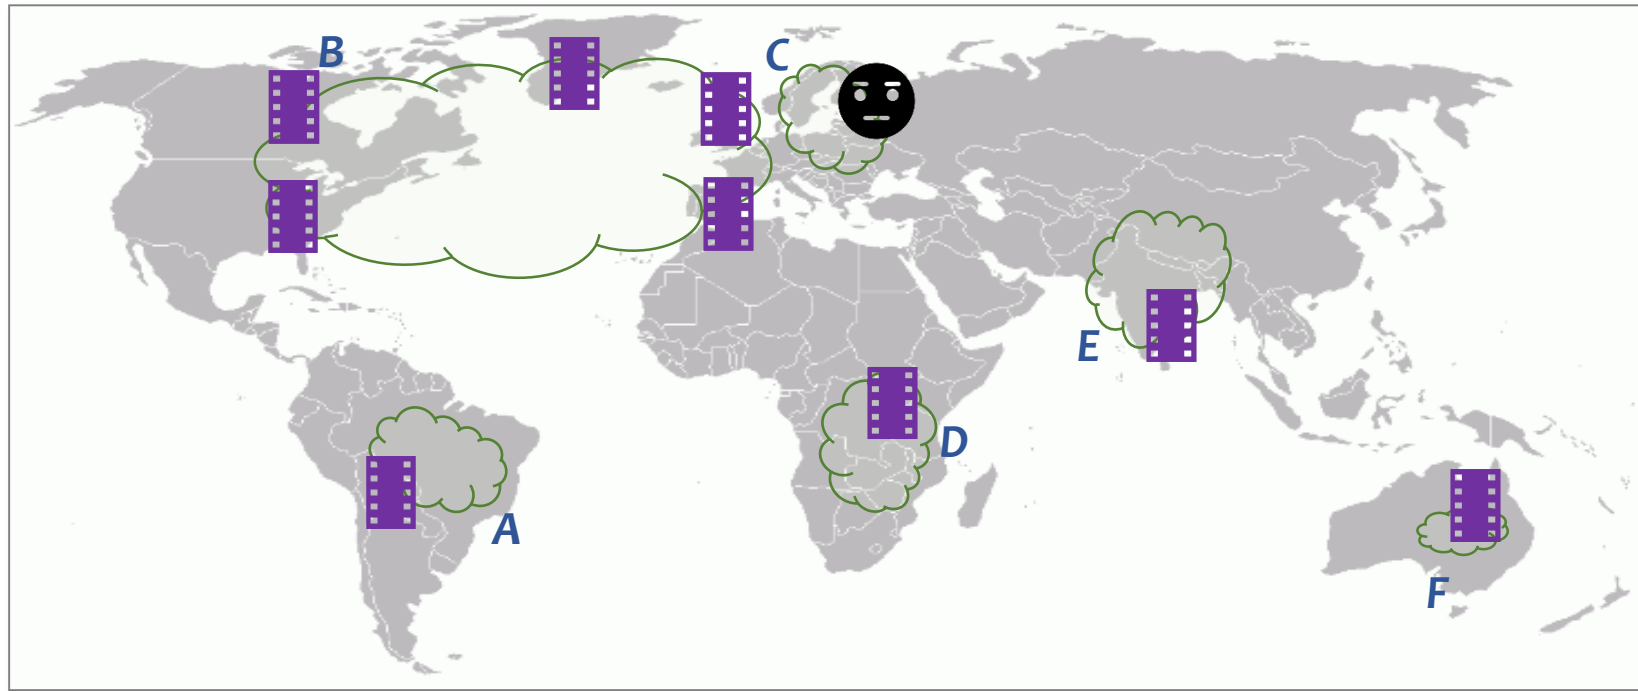
- IP packets are sent to a specific destination IP address
- How can we send a packet to an anycast destination IP address?
  - **No difference** to regular unicast IP address
- How can we serve content from an anycast IP address?
  - Announce the same prefix **from multiple locations** in BGP



# IP Anycast: BGP




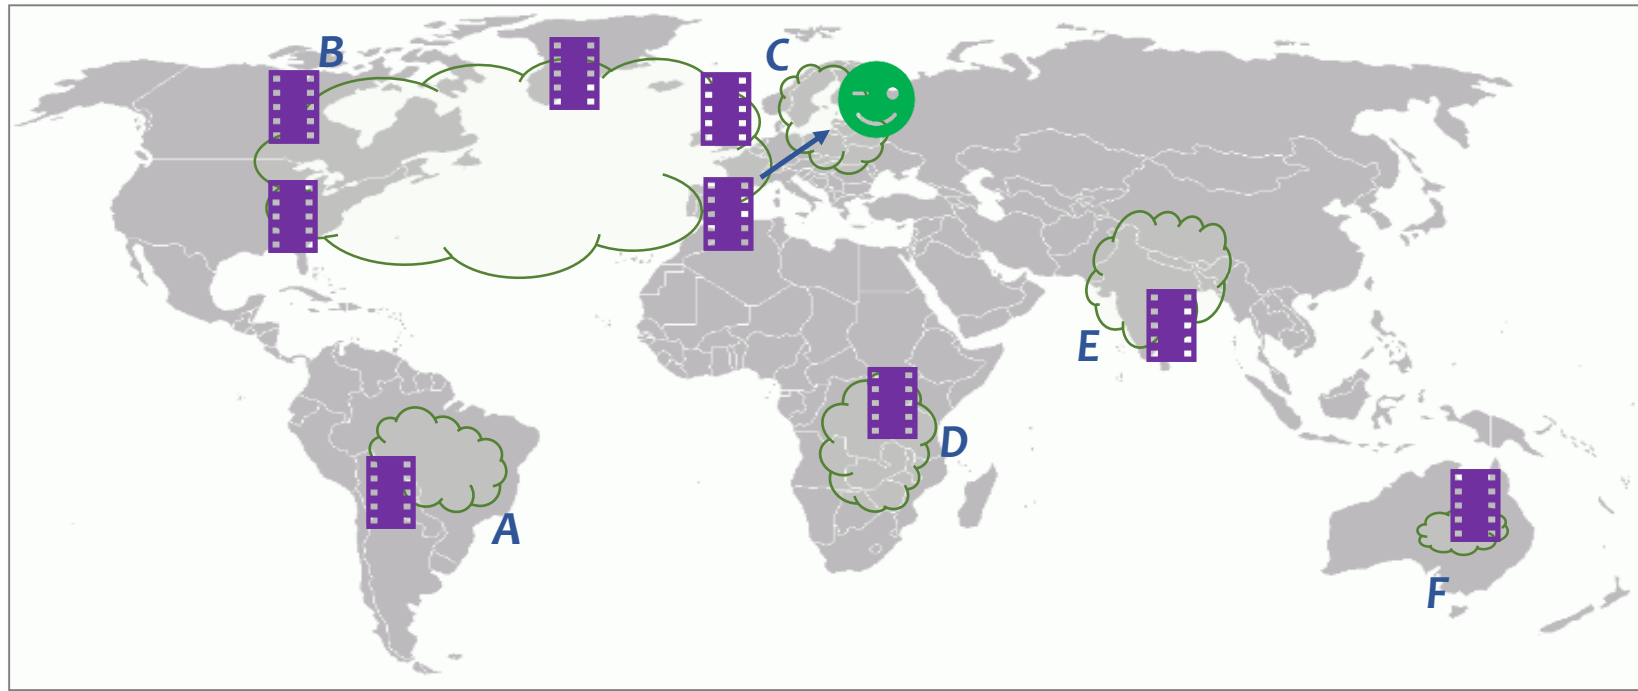
- Announce the same IP prefix from multiple locations



# IP Anycast: BGP




- Announce the same IP prefix from multiple locations




# IP Anycast: BGP



- Announce the same IP prefix from multiple locations



- BGP automatically “chooses” the nearest location



# Challenges with IP Anycast




- Many factors influence BGP routing decision: AS path, local preference, BGP communities, MED,...
- Chosen anycast replica can change repeatedly
  - cf. “Broad and Load-Aware Anycast Mapping with Verfploeter” by de Vries et al.
- Connectionless services are best suited for IP anycast



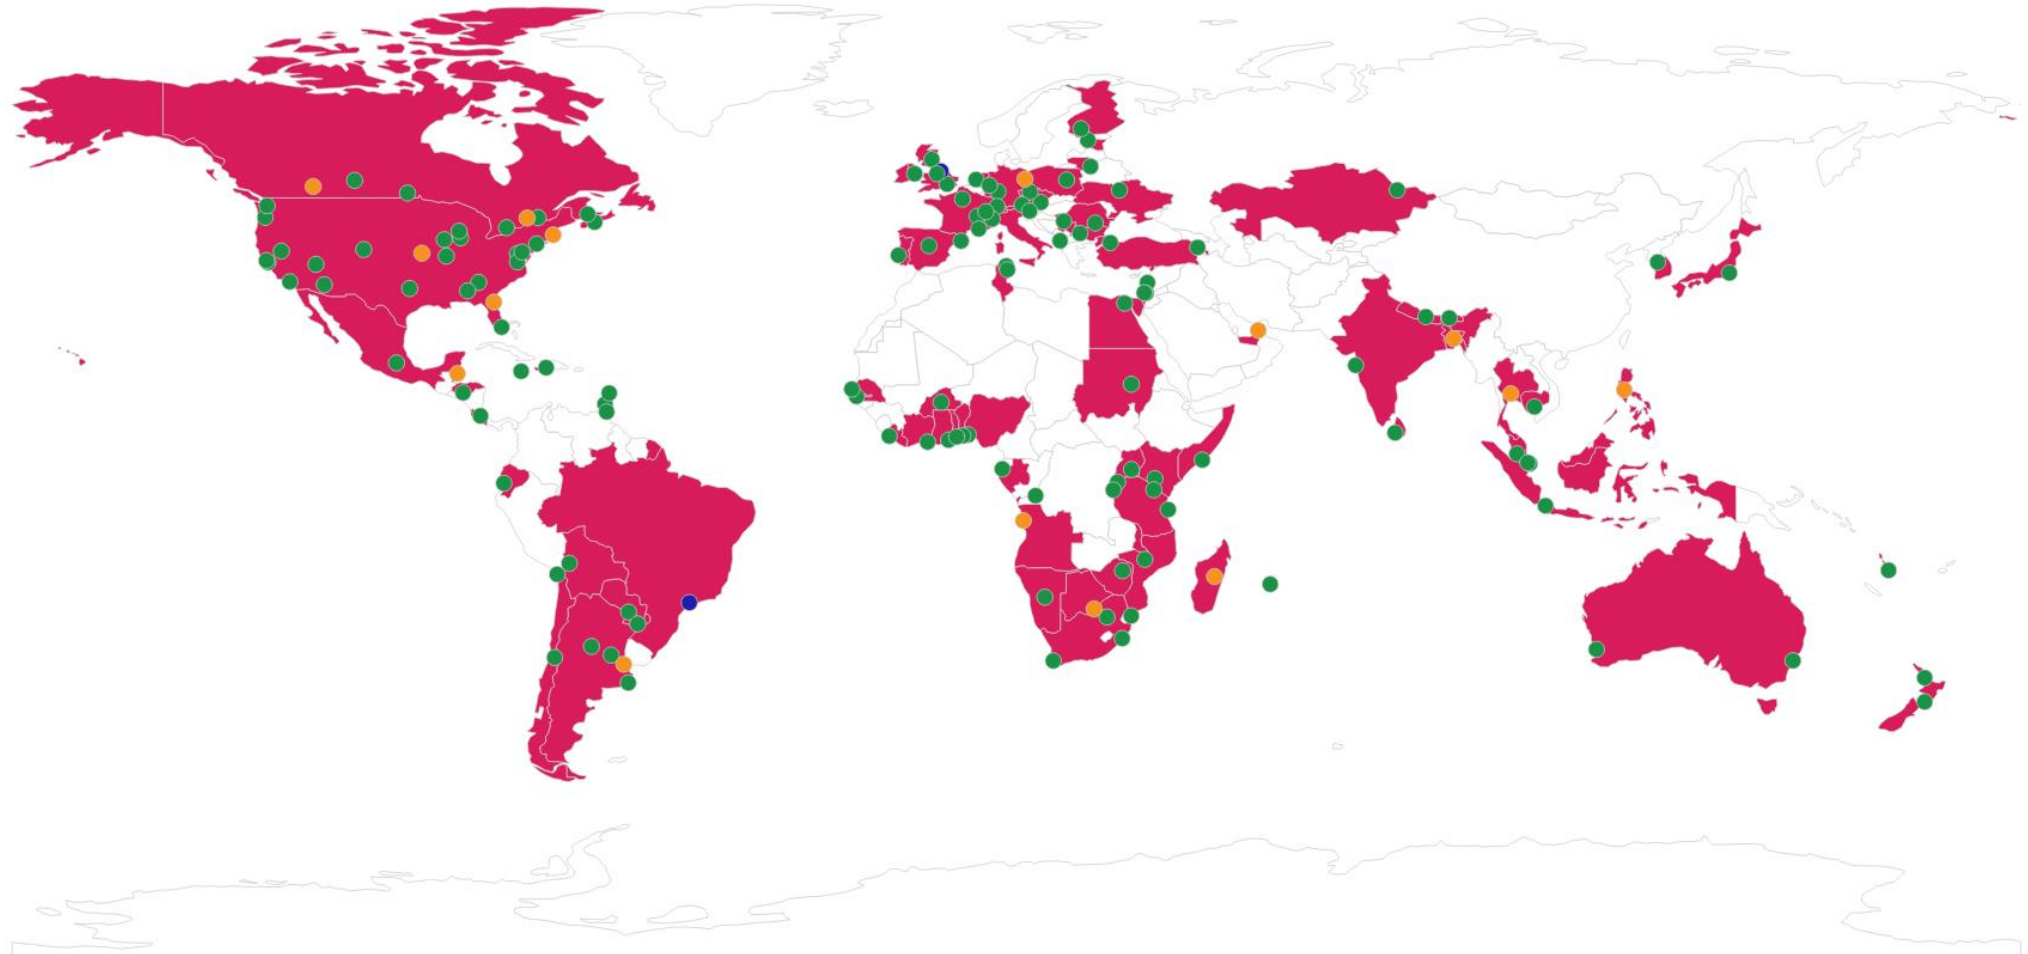
# Use Case: DNS Open Resolvers



- Google Public DNS
  - 8.8.8.8, 8.8.4.4, 2001:4860:4860::8888, 2001:4860:4860::8844
- 1.1.1.1: operated by Cloudflare
  - 1.1.1.1, 1.0.0.1, 2606:4700:4700::1111, 2606:4700:4700::1001
- Quad9: operated by PCH
  - 9.9.9.9, 149.112.112.112, 2620:fe::fe, 2620:fe::9



# Quad9: 131 Anycast Replicas




POP Status: 131 online, 17 in bypass, 0 in maintenance, 0 degraded, 2 coming soon

Data Networks

Anycast

Source: [quad9.net](https://quad9.net)



# EDNS Client Subnet



- Problem: the authoritative nameserver does not know the original client's IP address
  - DNS-based load-balancing using CNAMEs not really possible





# EDNS Client Subnet



- Problem: the authoritative nameserver does not know the original client's IP address
  - DNS-based load-balancing using CNAMEs not really possible
- Solution: EDNS Client Subnet DNS extension
  - Recursive resolver signals the original client's prefix
  - Authoritative nameserver can answer based on the client prefix



# Recap



- IP anycast
- BGP enables anycast
- Use case: DNS open resolvers

