

# BGP and Traffic Engineering with Akamai

Kams Yeung  
Akamai Technologies  
VNIX-NOG 2023  
6<sup>th</sup> Oct 2023



# Agenda

## Akamai Introduction

- Who's Akamai?
- Akamai Connected Cloud
- Akamai network deployment in Vietnam

## How Akamai map traffic (and why most of the BGP Traffic Engineering doesn't work with Akamai)

- AS Path Prepending
- MED
- More Specific Route

## Best practices and recommendations

- Setup own DNS resolvers
- Maintain complete and consistent route announcements
- Do not filter traffic
- Avoid CGNAT and enable IPv6

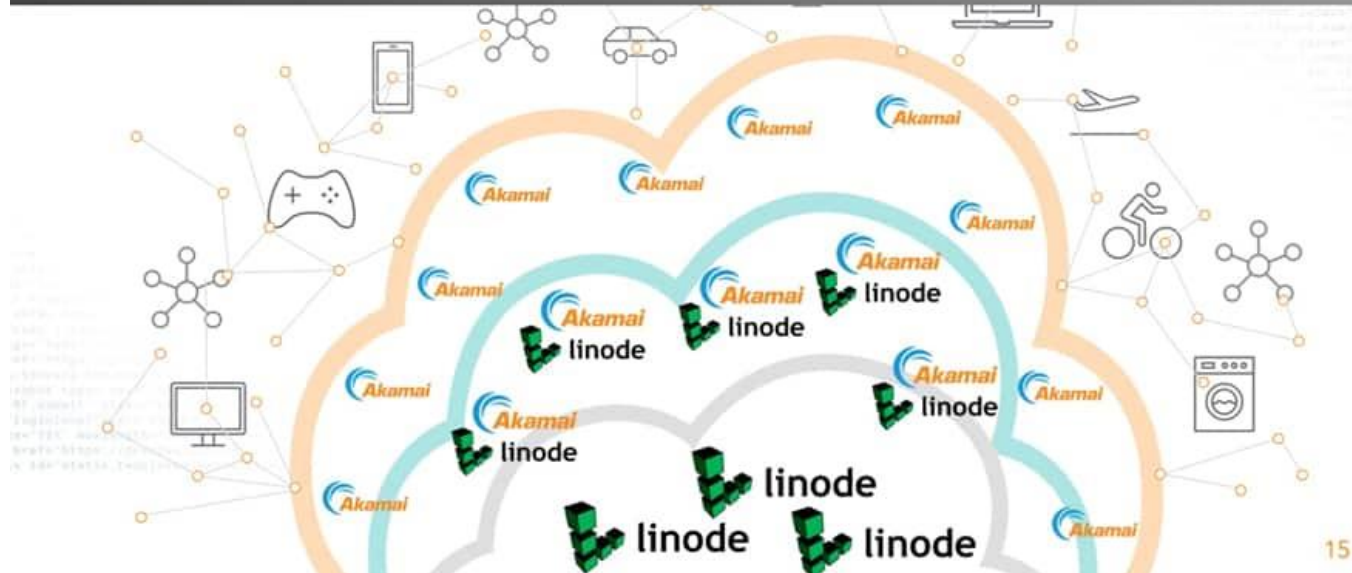
# Akamai Introduction



## Who is Akamai?

**Akamai Connected Cloud is a massively distributed edge and cloud platform that keeps experience close to users – and threats farther away**

**Akamai and Linode:** The world's most distributed compute platform – from cloud to edge – making it easier for developers and businesses to **build, run, and secure** applications.



# The Akamai Connected Cloud

The world's largest on-demand, distributed computing platform delivers all forms of web content and applications

## Content Delivery Network (CDN) Edge Deployment:

**340,000+**  
Servers

**4,000+**  
Locations

**1,200+**  
Networks

**750+**  
Cities

**130+**  
Countries

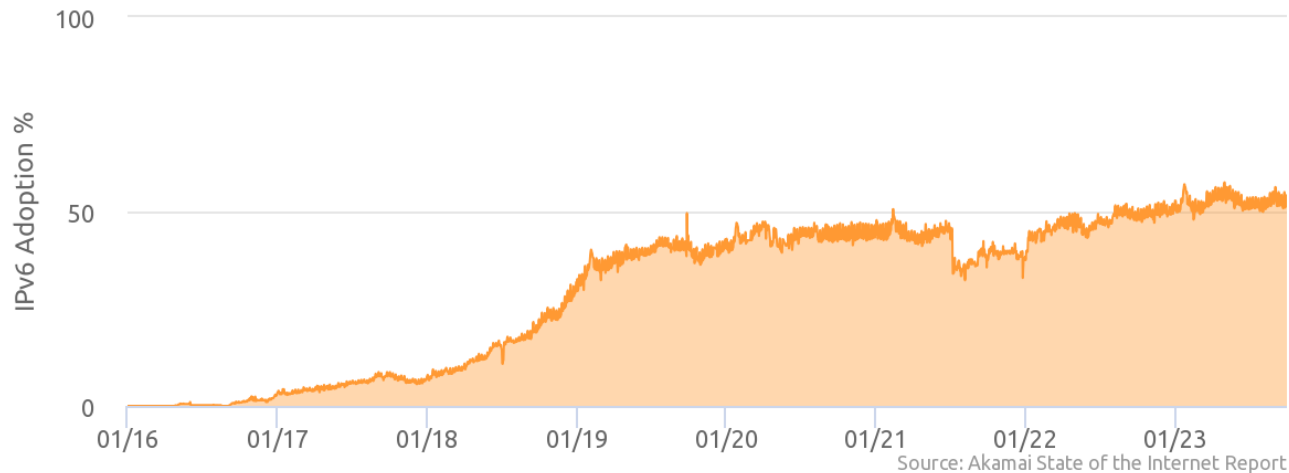


### Typical daily traffic:

- More than **3 trillion** requests served
- Delivering over **250 Terabits/second**

# Akamai Connected Cloud deployment in Vietnam

- Inside major ISP networks
- Located in 3 cities
  - Hanoi
  - Da Nang
  - Ho Chi Minh City



- Vietnam IPv6 adoption: 54.2%



# How Akamai works?

- Some facts
- Cluster types – Private and Public
- Cluster roles – Edge and Mid-tier

# How Akamai works: some facts

Are Akamai clusters connected to each other?

Most of the Akamai clusters are operated independent, and not talk to each others  
-> Peer with Akamai in one single cluster would not get all the content

Where does the content Akamai serve come from?

Akamai operates a caching infrastructure

Some content customers uses multi-tiers cache layer

Each cluster has different Internet connectivity to obtain content from the origin

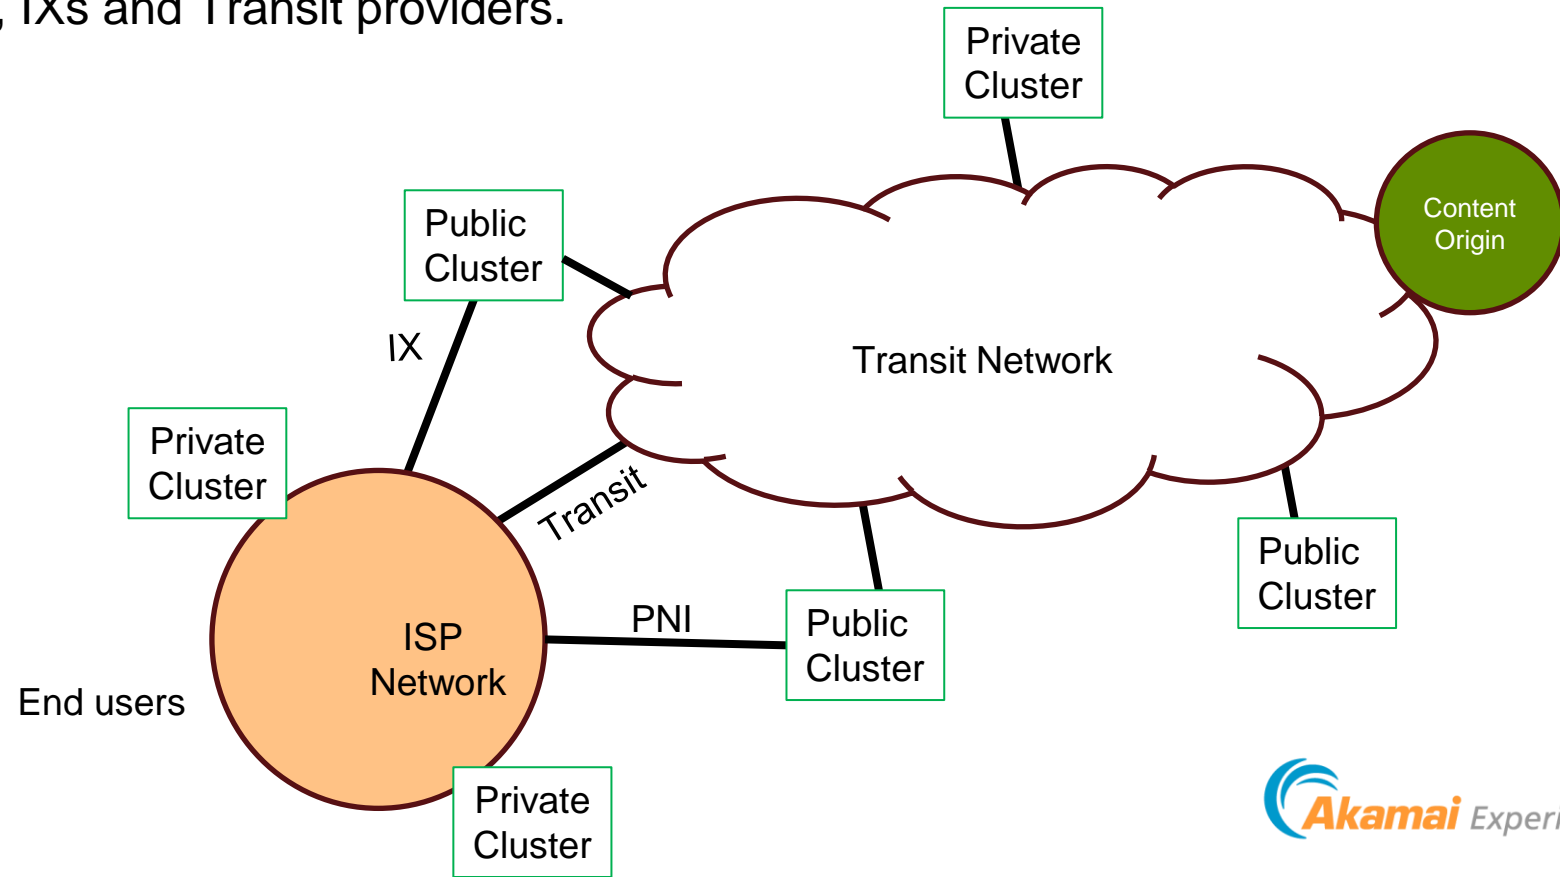
How does Akamai direct users to its cluster?

Akamai maps users based-on client's DNS, various network factors and other attributes

# How Akamai works: Cluster types

Private: Clusters dedicated to specific networks or their downstream ISPs. Role: Edge  
Clusters inside the network partners: AANP - Akamai Accelerated Network Partner

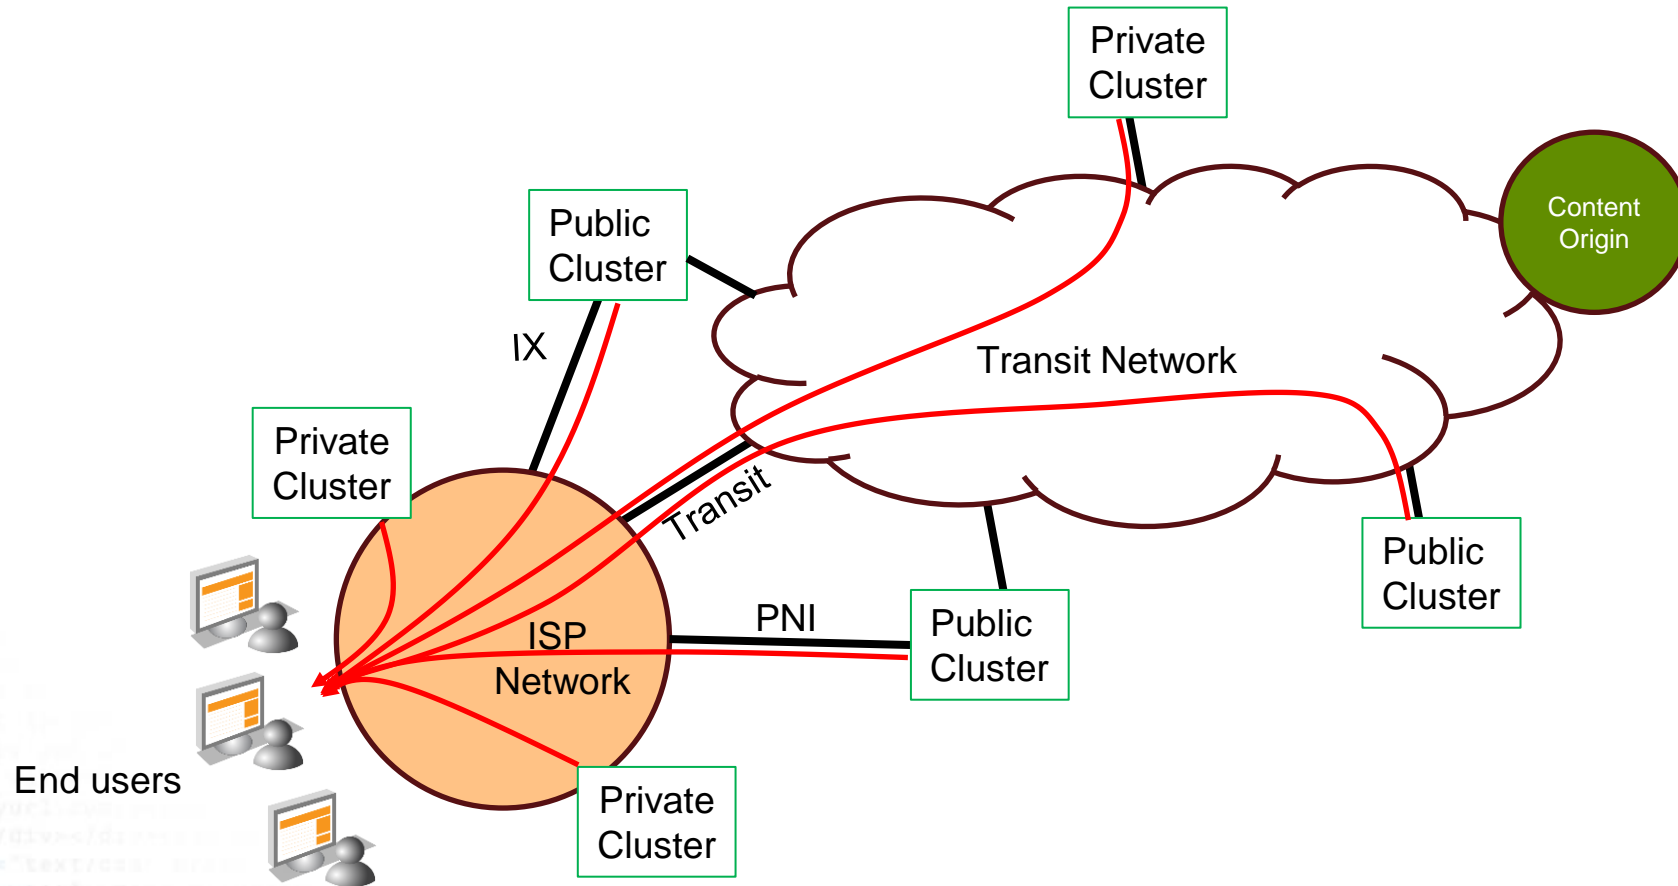
Public: Clusters shared by multiple networks. Role: Edge, Mid-Tier and Infrastructure  
Clusters inside some public facility (e.g., Mega-i, Equinix HK1, Global Switch, Equinix SG1), connecting to multiple networks via PNI, IXs and Transit providers.





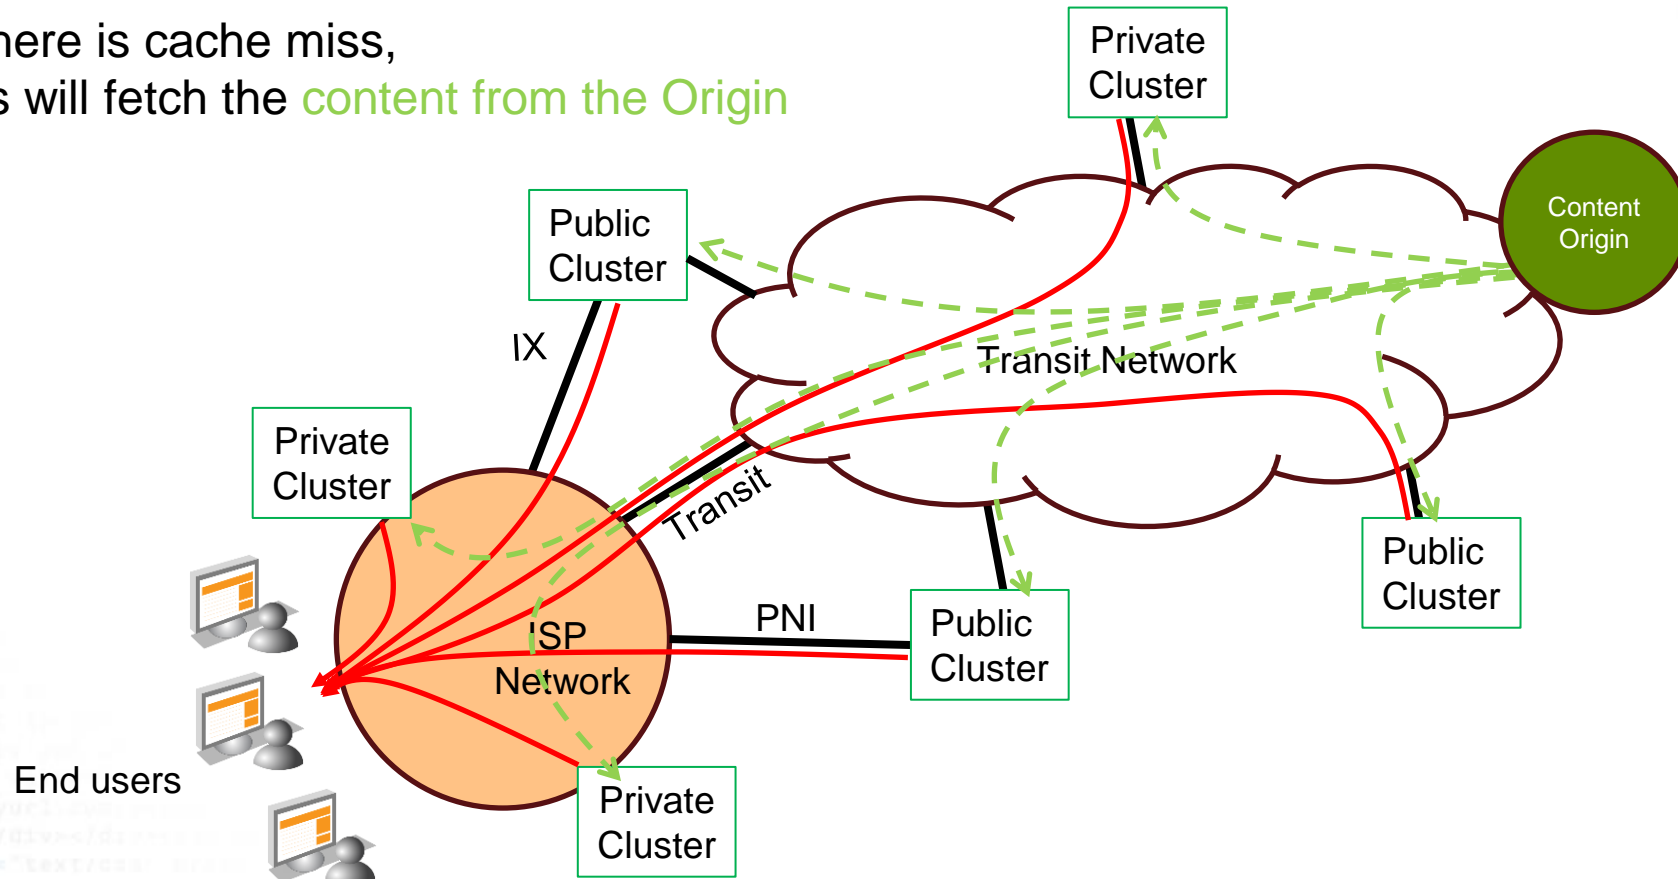
# How Akamai works: Cluster roles - Edge

1. Edge: Cache the content and distribute to the end users



# How Akamai works: Cluster roles - Edge

1. Edge: Cache the content and distribute to the end users when there is cache miss, clusters will fetch the content from the Origin

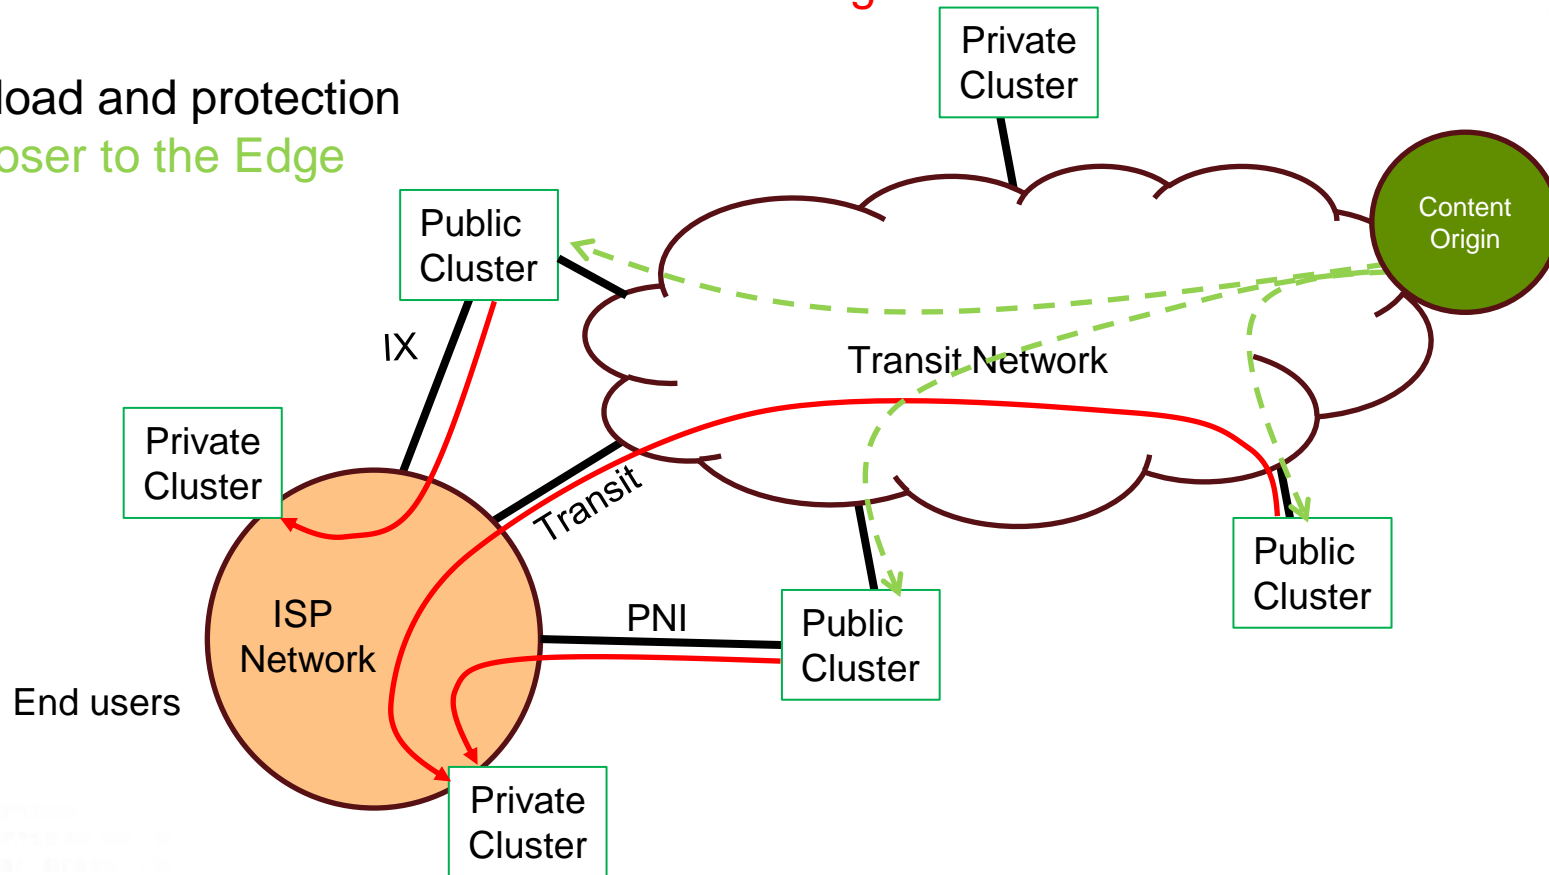


# How Akamai works: Cluster roles – Mid-Tier

2. Mid-Tier: Parent Cache the content and distribute to Edge clusters

Improve origin offload and protection

Bring the origin closer to the Edge

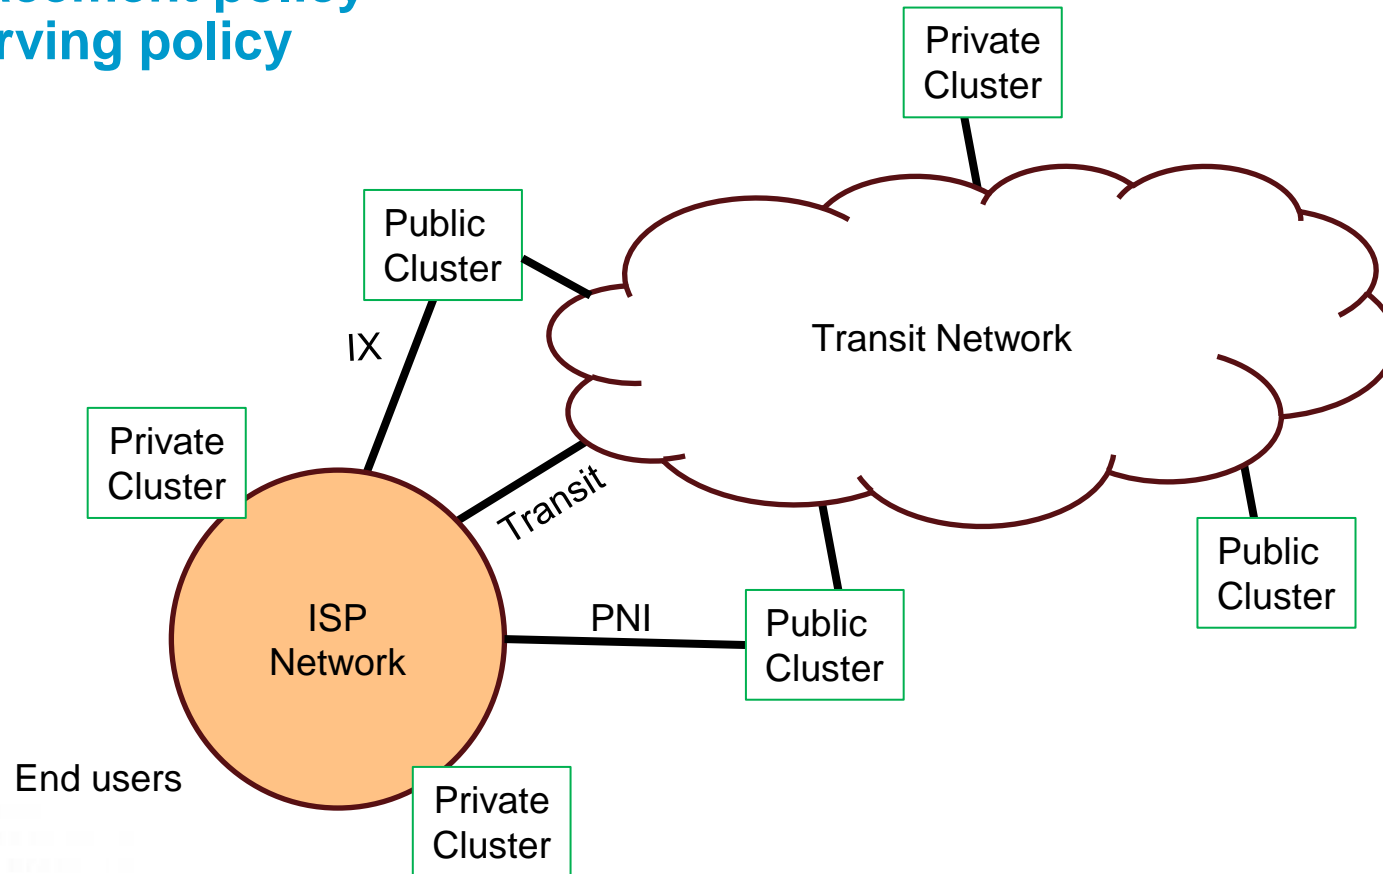


# Why most of the BGP Traffic Engineering techniques don't work with Akamai?

- AS Path Prepending
- MED
- More/less Specific Route advertisement

# Akamai maps end users demand based on ...

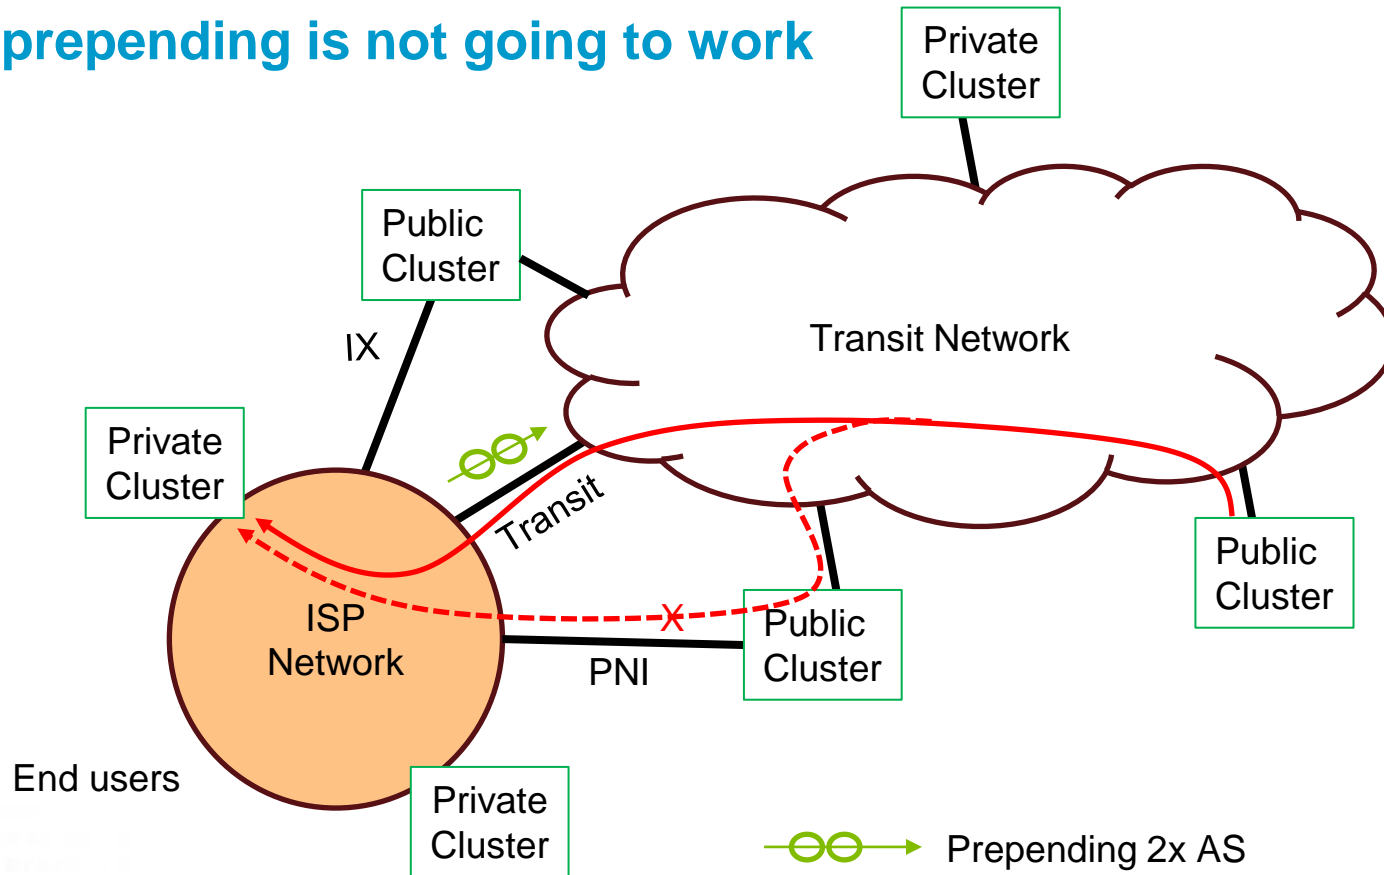
1. Network performance (Latency, packet drops, link utilizations)
2. Server capacity
3. Cache placement policy
4. Cache serving policy



# Akamai map end users demand based on ...

but not AS path length

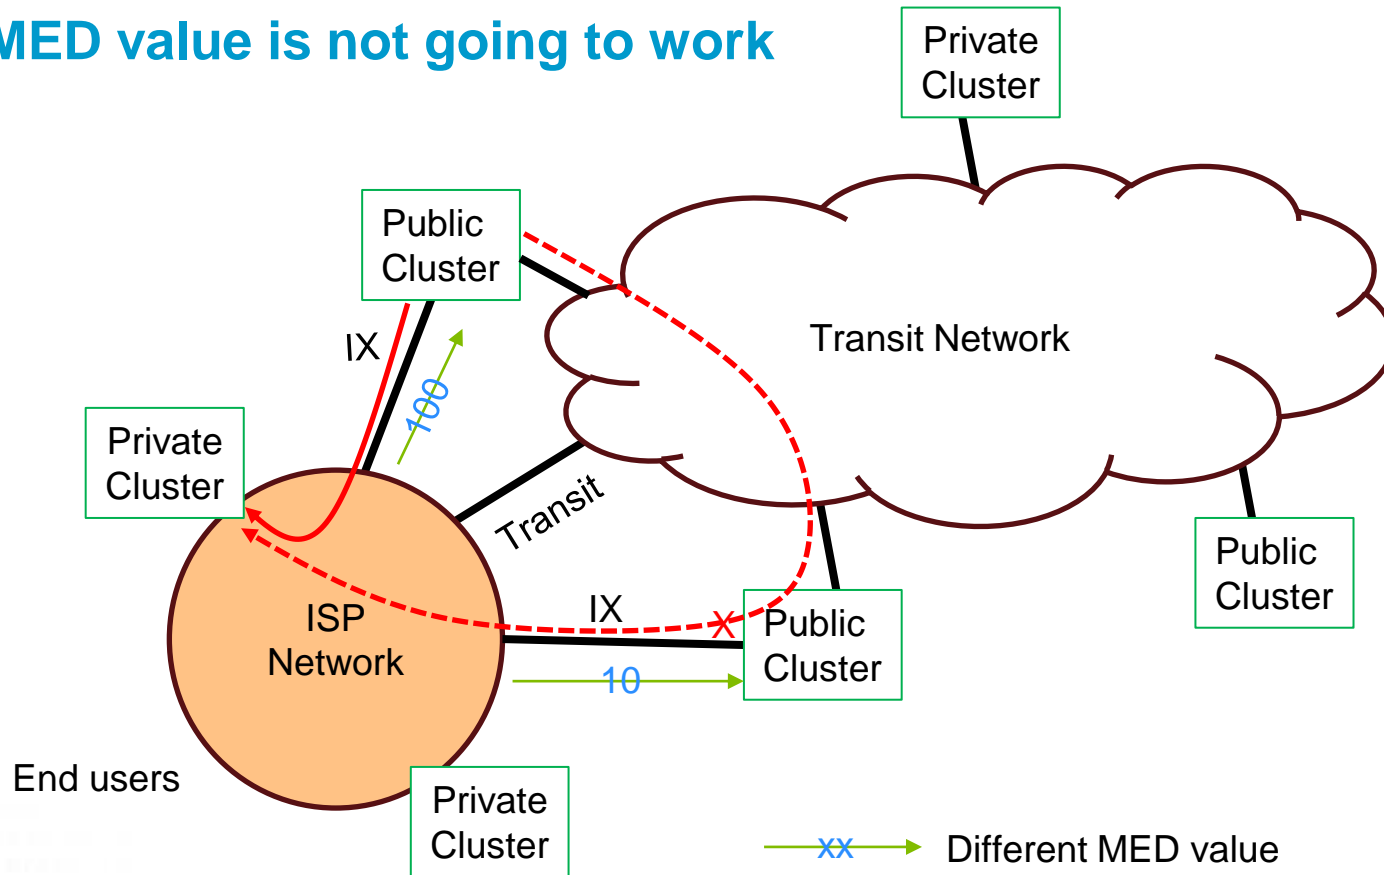
So AS Path prepending is not going to work



# Akamai map end users demand based on ...

but not MED value

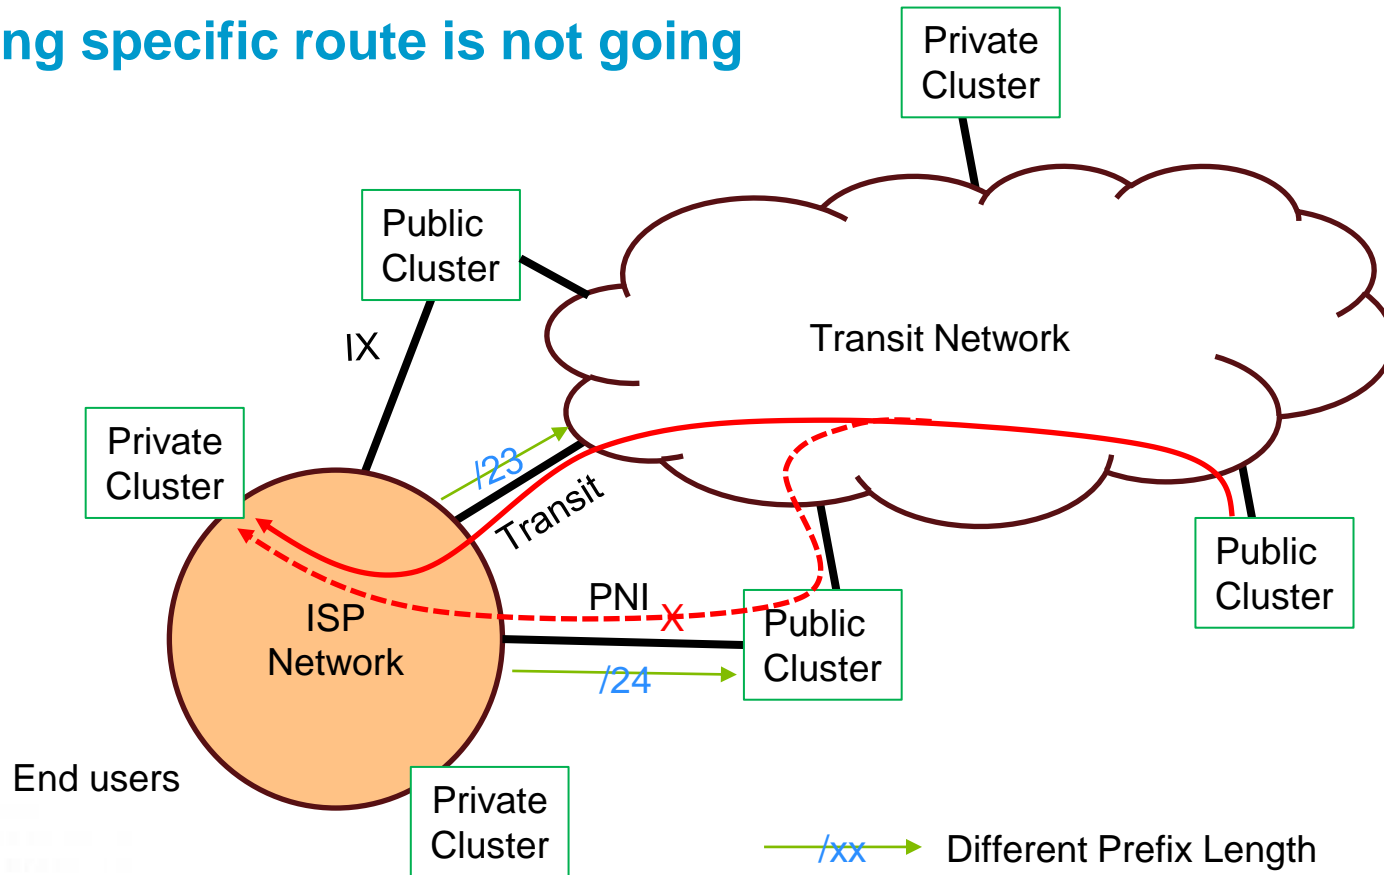
So change MED value is not going to work



# Akamai map end users demand based on ...

but not prefix length

So advertising specific route is not going to work





# Why doesn't these have the usual effect?

- Akamai uses Mapping, on top of the BGP routing
- Akamai Mapping is different from BGP routing
- Akamai nodes are mostly islands, there are no backbone between them
- Akamai uses multiple criteria to choose the optimal node / server
- These include standard network metrics:
  - Latency
  - Throughput
  - Packet loss

# Our suggestions

- Talk to us if we are sending too much / too few traffic on your preferred link(s)
- Accommodate our best practices and recommendations
- We can work together for traffic engineering

# Best Practices and Recommendations

- Setup own DNS resolvers
- Maintain complete and consistent route announcements
- Do not filter traffic
- Avoid CGNAT and enable IPv6

# Setup own DNS resolvers

Akamai CDN map traffic based-on DNS resolvers external IPs

- Use anycast IPs for user-facing DNS resolver IPs
- Use different external IPs for users in different locations
- Setup ACL to only allow your own users to use your DNS resolvers
- Do not modify Akamai hostnames TTL value

If not possible to setup your own DNS resolvers, then

- Use Google DNS (8.8.8.8 / 8.8.4.4, 2001:4860:4860::8888 / 2001:4860:4860::8844)
- Use OpenDNS (208.67.222.222 / 208.67.220.220, 2620:119:35::35 / 2620:119:53::53)
- Akamai support EDNS Client Subnet (ECS) for Google DNS and OpenDNS
- Publish GeoFeed IP location information in RFC8805 format

Maintain good Internet connectivity to your DNS resolvers

- Akamai may use your DNS resolvers external IPs for performance monitoring
- Alternatively, you may provide Akamai with your desire IPs for performance monitoring

# Maintain Complete and Consistent Route Announcements

Announce complete prefixes to Akamai

- Includes both DNS and end user IPs
- Akamai map traffic based-on DNS to the optimal node, then send user traffic from there
- Inform your downstreams to announce all prefixes to you

If not possible to announce all prefixes, then

- Akamai may block your whole ASN prefixes, to avoid suboptimal performance

Maintain consistent route announcement to your peers / upstream providers

- Akamai may send overflow traffic from your upstream providers

# Do not filter traffic

## Carry traffic that you announce

- If you promised to carry the traffic of an IP block (e.g., /20), you should not have any holes (e.g., /24) or drop any part of the traffic
- Akamai routers may not have the full Internet routing table
- The end user's connectivity will be impacted!!!

## Performance monitoring

- Akamai uses IPs in your network as performance monitoring
- If possible, do not filter / rate-limit ICMP to your network
- Send return traffic to Akamai closet location to maintain lowest latency

# Avoid CGNAT and enable IPv6

## Avoid the use of CGNAT

- When possible, try to avoid using CGNAT, improve performance
- If necessary
  - Use the standard CGNAT IP address block 100.64.0.0/10 [RFC 6598]
  - Place Akamai nodes outside of CGNAT
- Akamai uses client IPs for different purposes, e.g., global traffic management

## Enable IPv6

- If you have an Akamai cache node
  - assign IPv6 address block to Akamai clusters
  - send your IPv6 prefixes to Akamai BGP collectors
- If you have peering with Akamai, enable both IPv4 and IPv6 sessions
- Akamai has made IPv4+IPv6 dual-stack the default for new customer configurations

# Summary

## **Akamai Connected Cloud**

- Highly distributed edge servers
- DNS-based mapping CDN

## **Traffic Engineering**

- Typical BGP traffic engineering techniques doesn't work
- Collaborate with Akamai for traffic engineering

## **Best practices and Recommendations**

- Setup your own DNS resolvers
- Maintain complete and consistent route announcements
- Do not filter traffic
- Avoid CGNAT and enable IPv6



# Questions?

Kams Yeung [kams@akamai.com](mailto:kams@akamai.com)

More information:

Peering: <https://as20940.peeringdb.com>

Cảm ơn!  
(Thank You!)