

Optimizing Traffic Flow with Akamai

Eric Leung
Akamai Technologies



Agenda

Akamai Introduction

- Akamai deployment in South Asia

How Akamai map traffic

- Different types of Akamai Clusters
- Why most of the BGP Traffic Engineering techniques doesn't work with Akamai

Best practices and Recommendations

- Setup own DNS resolvers
- Maintain complete and consistent route announcements
- Do not filter traffic

Akamai deployment in South Asia

Afghanistan: -

Bangladesh: Dhaka

Bhutan: Thimphu

India: 40 cities

Maldives: -

Nepal: Kathmandu

Pakistan: 6 cities

Sri-Lanka: Colombo



How Akamai works: Some facts

Are Akamai clusters backbone connected to each other?

Most of the Akamai clusters are operated independent, and not talk to each others.


Where does the content Akamai serve come from?

Akamai operates a caching infrastructure

Each cluster has full Internet connectivity to obtain content from the original website or mid-tier servers.

How does Akamai direct users to its cluster?


Akamai maps users via a multi level DNS server architecture



How Akamai works: Cluster types


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

Public: Clusters shared by multiple networks. Role: Edge, Mid-Tier and Infrastructure
Clusters inside some public facility, connecting to multiple networks via PNI, IXs and Transit providers.



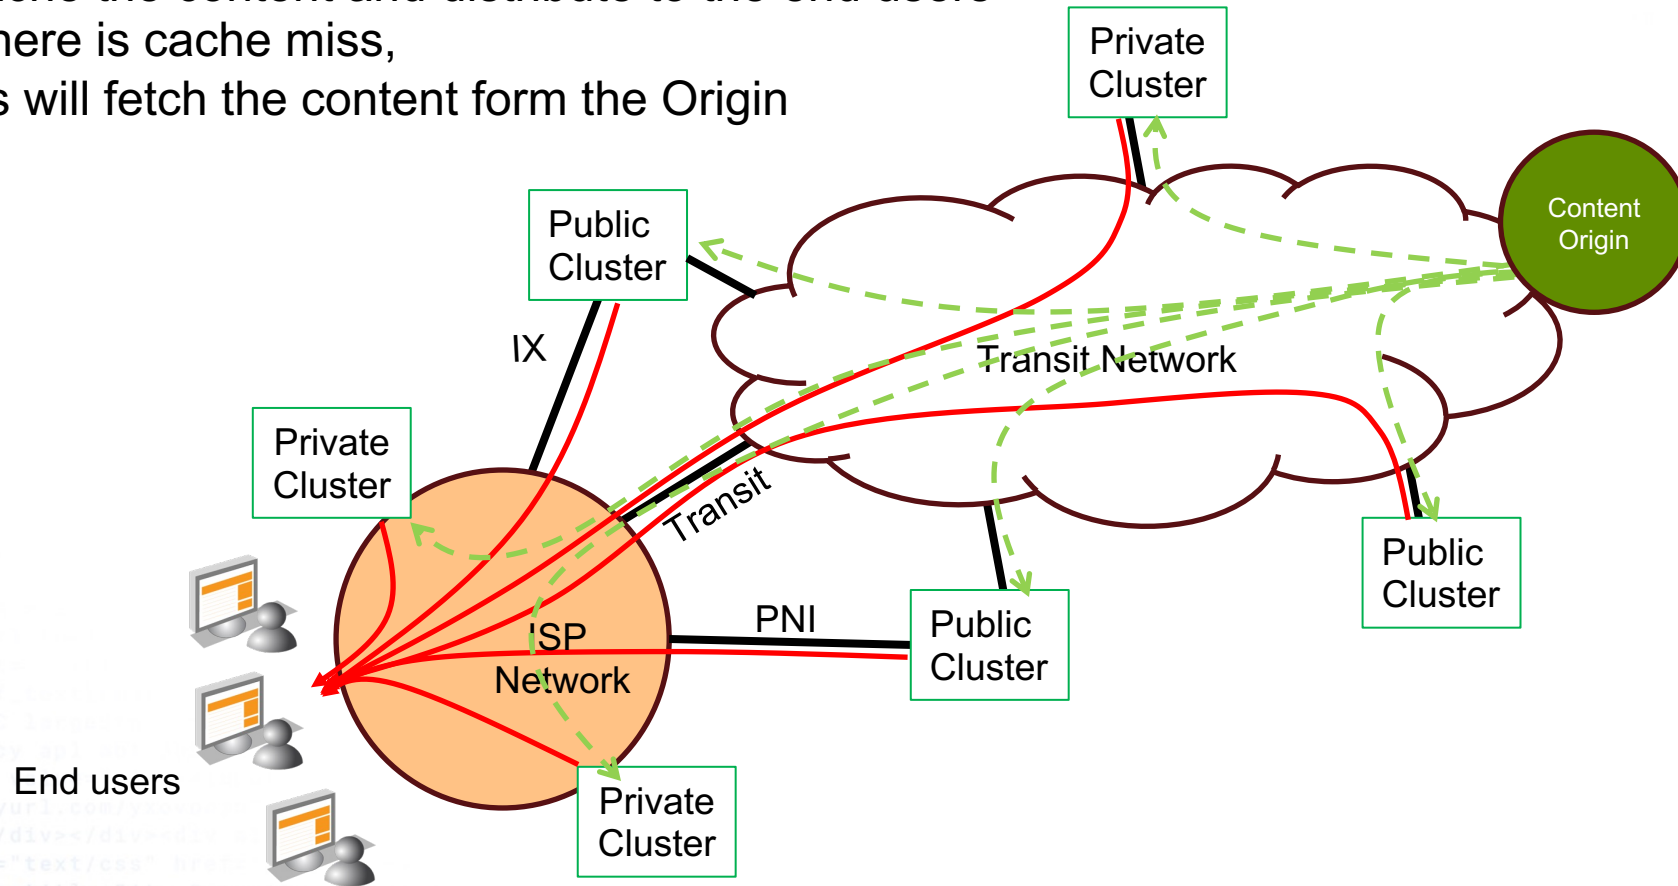
How Akamai works: Cluster roles - Edge

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



How Akamai works: Cluster roles - Edge

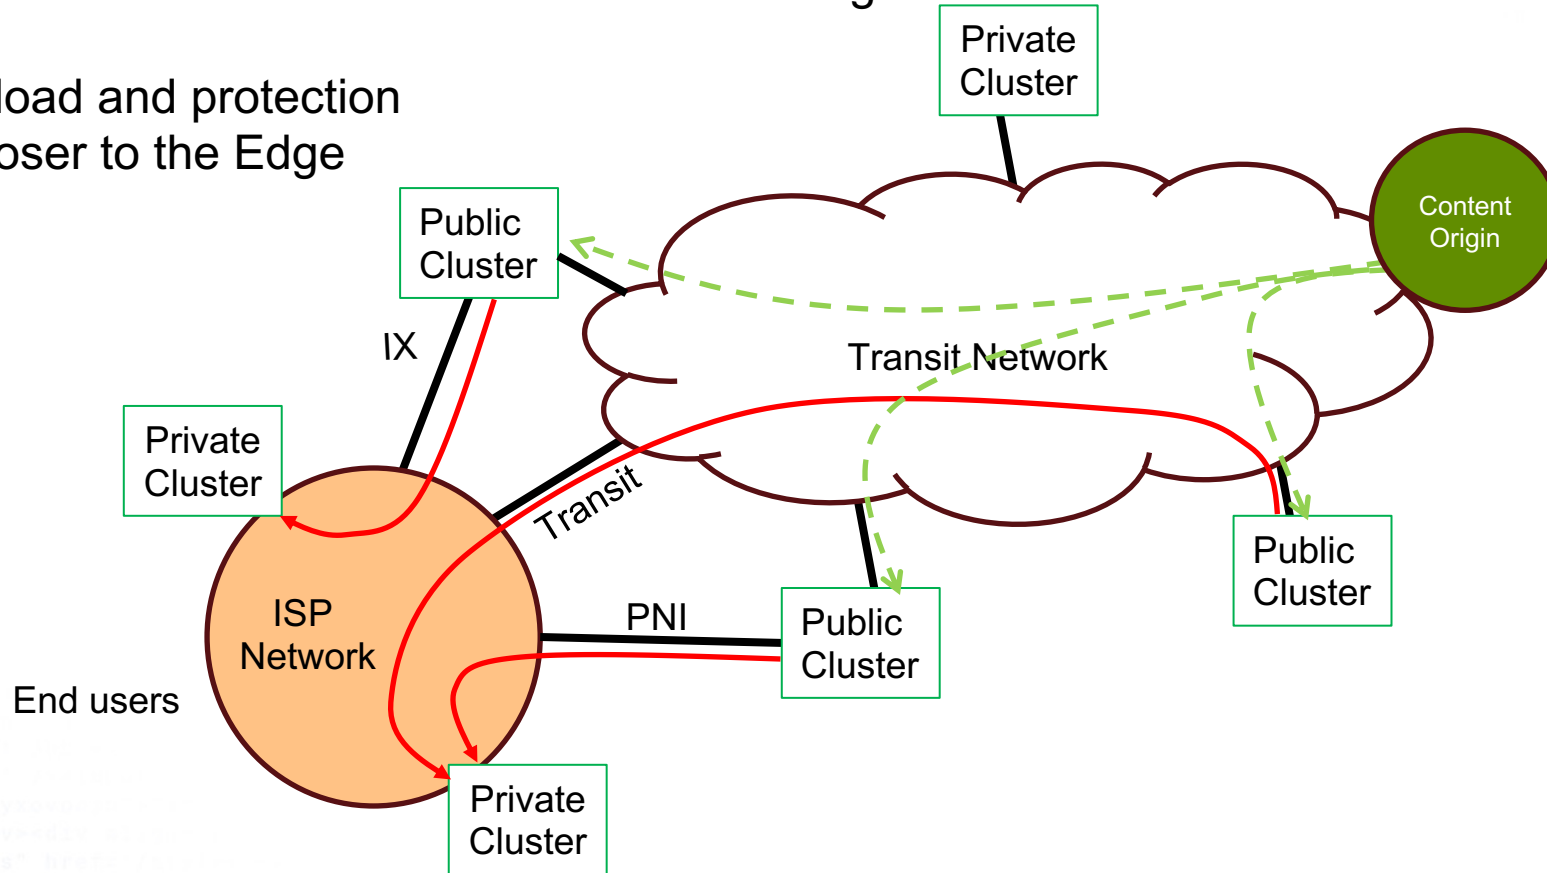
1. Edge: Cache the content and distribute to the end users when there is cache miss, clusters will fetch the content form the Origin



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 traffic to a cluster based on ...

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




Akamai does not map traffic to a cluster based on ...

AS path length, so AS Path prepending is not always working




Akamai does not map traffic to a cluster based on ...

MED value, so changing MED value is not always working



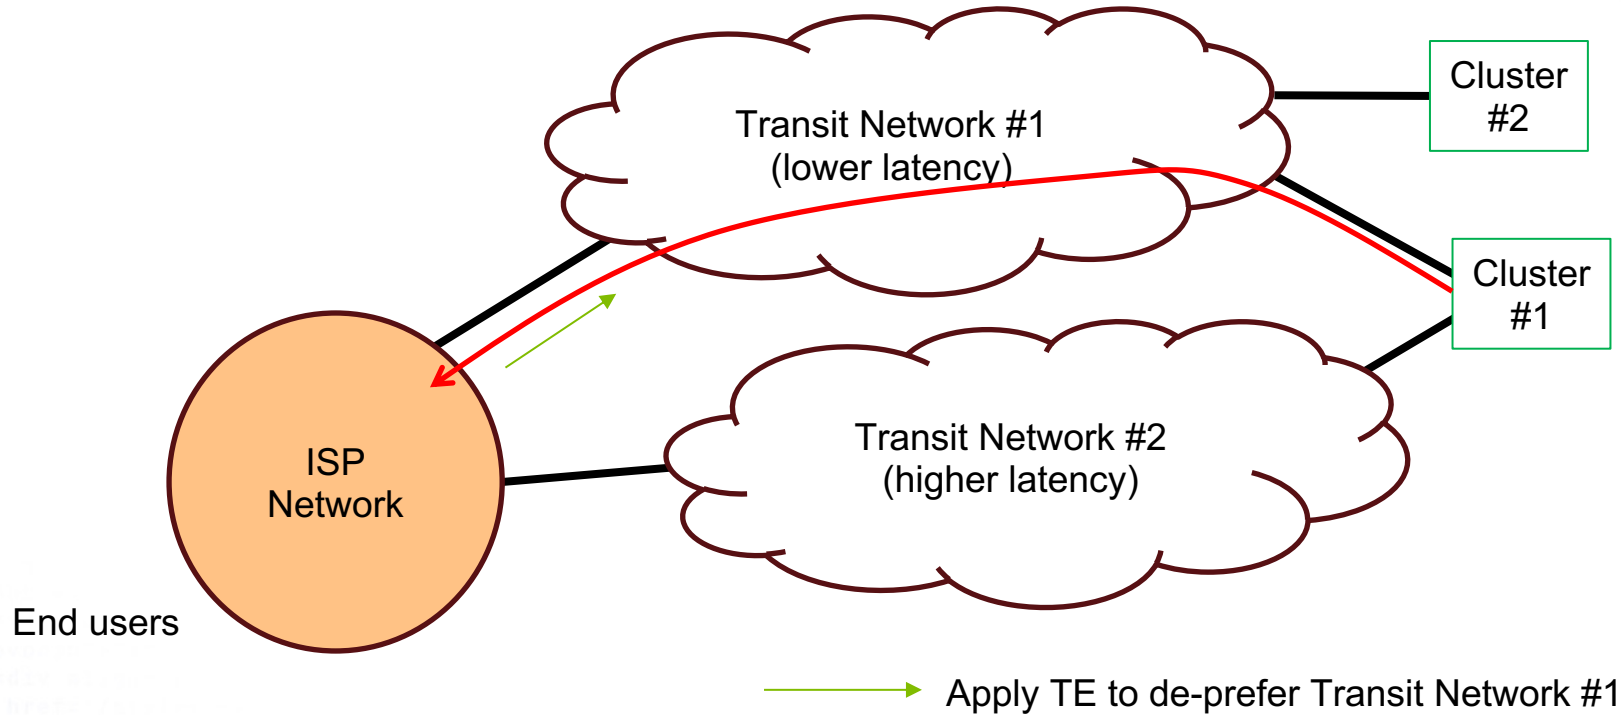
Akamai does not map traffic to a cluster based on ...

prefix length, so advertising specific route is not always working



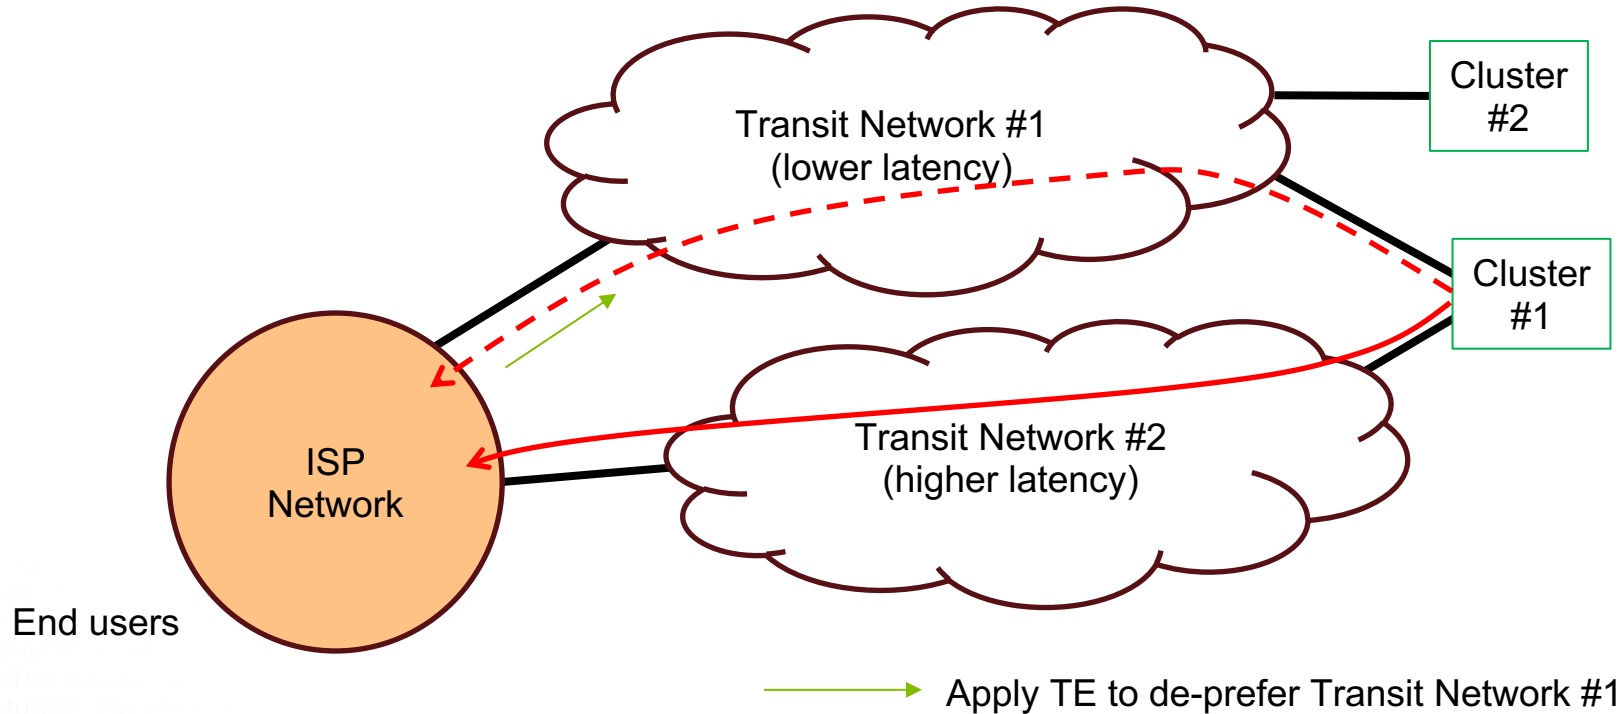
Sometime bgp based traffic engineering seems working

in case both paths are reachable from the same cluster



Sometime bgp based traffic engineering seems working

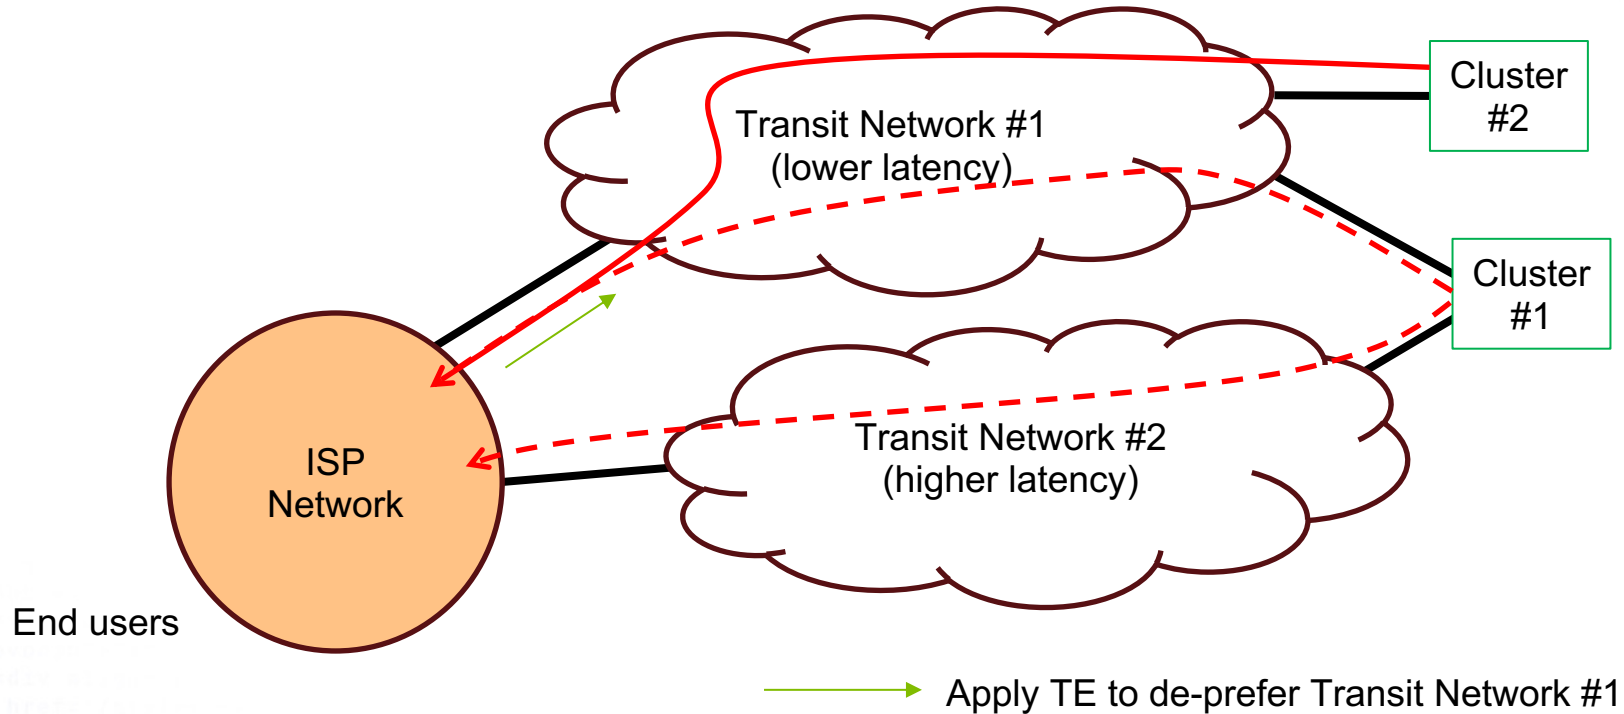
in case both paths are reachable from the same cluster, the router of that cluster will reroute the traffic as the best-path has been changed



Sometime bgp based traffic engineering seems working

in case both paths are reachable from the same cluster, the router of that cluster will reroute the traffic as the best-path has been changed

However, Akamai mapping may then change the serving cluster to #2 as its performance is now better...



Why don't these techniques have the usual effect?

- Akamai use BGP to learn routes so as to determine which clusters are eligible to serve what traffic
- Akamai uses Mapping to determine which cluster to serve the traffic
- Akamai nodes are mostly isolated; there is no backbone between them
- Akamai uses multiple criteria to choose the optimal node / server
- These criteria include standard network metrics:
 - Latency
 - Throughput
 - Packet loss

Akamai suggestions to ISPs

- Please accommodate our best practices and recommendations to enable the Akamai mapping system to make optimal decision
- Feel free to contact us if you wish to change the way traffic being sent to you
- We can work together on optimization traffic flow

Best Practices and Recommendations

- Setup own DNS resolvers
- Maintain complete and consistent route announcements
- Do not filter traffic

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

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

Maintaining 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, avoid filtering or rate-limiting ICMP to your network
- Send return traffic to Akamai closet location to maintain lowest latency

Summary

Akamai Intelligent Platform

- Highly distributed edge servers
- DNS-based mapping CDN

Optimizing Traffic Flow

- 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

Questions?

Eric Leung <taleung@akamai.com>

More information:

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