

The path to 100+ IXes & wish-list for VNIX

Kams Yeung
Akamai Technologies
VNIX-NOG 2022
7th Oct. 2022



Agenda

- Akamai Introduction
- The path to 100+ IXes
- Wish-List for VNIX

Akamai Introduction

- Who's Akamai?
- Akamai Intelligent Platform

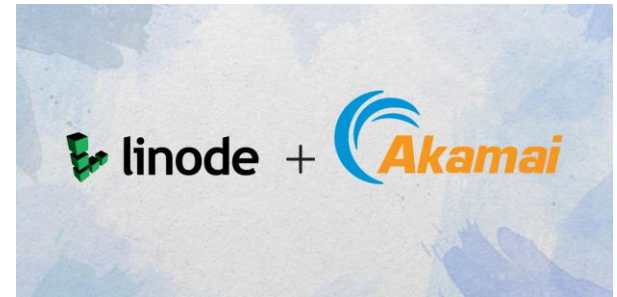


Akamai Overview

Who is Akamai?

Akamai is a leading provider of a Cloud platform, which delivers, accelerates and secure content and applications over the Internet. Our key differentiator is our highly distributed (intelligent) platform, made up of approximately 355,000 servers in more than 135 countries and nearly 1,300 networks

- Publicly traded: (NASDAQ: AKAM)
- Found: August 1998
- Headquarters: Cambridge, MA, USA
- 40+ worldwide offices, including Europe and Asia
- 9,000+ employees worldwide



The Akamai Intelligent Platform

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

The Akamai Intelligent Platform:

355,000+
Servers

4,000+
Locations

1,300+
Networks

800+
Cities

135+
Countries



Typical daily traffic:

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

An infographic with a dark background. At the top, it says "250 Tbps" in large orange letters. Below that, it asks "are you ready to binge watch?" with a play button icon. At the bottom, it states: "At 250 Tbps, you could download every feature film ever made in 37 seconds. It would take 51 years of nonstop viewing to watch those films." A globe is visible in the background, and the Akamai logo is in the bottom right corner.

The path to 100+ IXes

- Where have we started?
- By the Numbers: Connected IXes, Traffic, Locations, Technology
- The Value Proposition of Peering
- Why do we join an IX?
- Design Considerations

Introduction


- Akamai is the first CDN to pass traffic at over 100 IX locations.
 - This took 15 years! (2000 – 2015)
- What/Why/How/When do we go to IXes?
- How does this fit into general trends in localization of traffic?
- What's next?

Where have we started?

- First IX in Americas: **PAIX** (~2000)
- First IXes in EMEA
 - UK: **LINX** (2001)
 - Continental Europe: **DECIX & AMSIX** (2001), **ESPANIX** (2002)
- First IX in APJ: **JPIX** (Early 2002)

By the Numbers: Connected IXes

of Connected IXes: 194



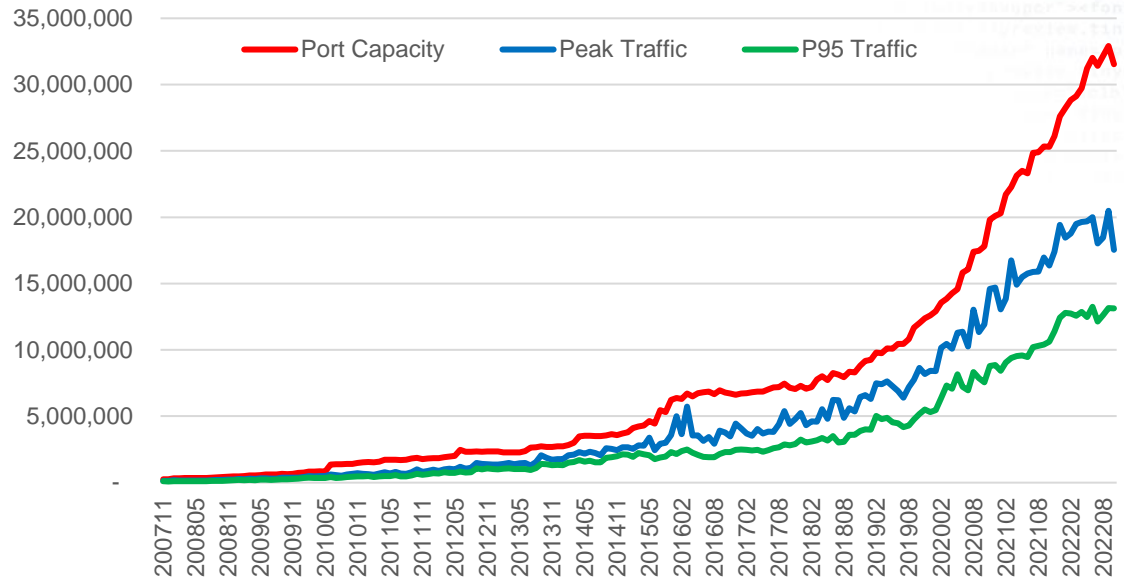
By the Numbers: Traffic

32Tb of *public* peering capacity

17Tb of peak traffic


13Tb of P95 traffic

Traffic




By the Numbers: Locations

of Locations
50+ countries, 100+ cities



By the numbers: Technology

Link Speed / Cluster



The Value Proposition of Peering

- Peering typically costs less than transit
- The peering link is a predictable/fixed cost
- Lower distance = higher performance and reliability
- High-bandwidth applications **require** low RTTs.
 - A stable 15 Mbps stream (typical of 4K video) **requires** < 34 ms RTT
 - And that's if there's no loss or retransmits.

Why do we join an IX?

- Reduces our costs
 - Less transit
 - Because we're a CDN, it typically means less transit for the IX members as well.
- Better overall performance for peers
 - Latency, topological diversity, overall throughput
- Geographical coverage/capacity
- Encourage the provider ecosystem in an area

Design Considerations

- One vs. Multiple clusters in the same LAN
 - Multiple clusters instead of complex component redundancy
 - Same number of peers on both clusters
 - Avoid issues with a single large cluster reallocating too much traffic
 - Always running active/active: we don't have idle spares
- Multiple locations instead of single locations
 - Put clusters in different physical locations to avoid physical fate-sharing, where possible.

Route servers: pros and cons

- Route servers are a convenient way to jumpstart entry into an IX
 - Exchange routes with numerous IX members, with just 1 or 2 BGP sessions.
 - Preview how many routes are announced by members, for capacity planning.
 - *As a rule, we will peer with route servers when we enter an IX*

- Route servers also have drawbacks.
 - Fate Sharing for all BGP sessions
 - The setup and feature set is different at each IX

In general, we will attempt to turn up sessions with other IX members as soon as we can.

IX vs. PNI: tradeoffs

- IX: shared bandwidth, low per-session costs
 - One-to-many relationship, turning up BGP sessions is cheaper & faster
 - **However**, IX fabric bandwidth is large, but members' links will be comparatively small
 - No control of the bandwidth or utilization of other peers until there is packet loss
- PNI: dedicated bandwidth, higher per-session costs
 - Reserved bandwidth/capacity, but cross connects and router ports cost money.
 - Limited resiliency: interruption of a single PNI is more likely during maintenance
 - A PNI only connects to a single cluster

At some point, it's cheaper to allocate a PNI than it is to take the traffic over the IX.

When do we help to “build” and IX?

- Multiple ways to “build” in IX.
 - Provide equipment/expertise for local talent that’s short on funds.
 - Create critical mass for an IX by being a “content anchor”.
- Growth potential of an area that would have lower costs with more cooperation
 - Adding in a content anchor gives them a reason to be in the same place.
 - Better peering = lower costs = more players and room for regional players to expand.
 - In addition, they get high-speed, low-latency access to content.

De-peering and leaving IXes

- We de-peered Networks
 - Because of broken route announcements
- Leaving IX's
 - Costs
 - Because they became dysfunctional (KleyerIX, FreeIX)

How Akamai Operates at an IX

- Inconsistent Routes
 - We only announce the local clusters
 - You don't need all Akamai routes
- Best-effort delivery
 - No guarantees about what or how much traffic will be sent
- No SLA on any specific cluster
 - We can't guarantee where any particular content will be sent from
- High traffic outbound vs. inbound, because we're a CDN

Issues

- Peers announcing the world
 - Usually gets noticed and handled quickly
 - Our mapping system can detect networks being far away
- IX route servers prepend their own AS
 - Makes route selection more difficult than it has to be.
- Bad BGP Traffic Engineering
 - Incomplete announcements, overlapping routes.
 - We might end up serving you over your transit.
 - **If you need to engineer traffic, *please talk to us*.** We have better tools than just BGP.

What does this all mean?

- Currently connected to 194 IXEs, and more to go.
- 400GE interface under testing / trial
- Getting traffic closer to the user is a more general trend
 - More local peering even in well-connected countries
- **If you're at common IXEs with us, we will peer with you.**

peering@akamai.com

Wish-List for VNIx

- Technical Perspective
- Operational Perspective
- Commercial Perspective



Technical Perspective

- **Use of Global ASN and IP address blocks**
 - Currently VNIX only allow ASN and IP address blocks registered under VNNIC
- **Could use VNIX connection for all traffic**
 - Currently only allow domestic traffic, do not allow international traffic
 - IX should not restrict traffic type
- **Do no need to advertise all IP address blocks to VNIX**
 - Not all networks have backbone, anycast prefixes
- **Support of 100G interface**
 - Only 1G and 10G interfaces are supported at this moment

Operational Perspective

- **Upgrade port when make sense**
 - Currently need to use 10G when traffic > 3Gbps
 - Members should upgrade their own infrastructure as own needs
- **Members could choose who to peer selectively**
 - Should not require mandatory MPLA peering
 - Anycast prefixes have specific peering targets
- **Language support**
 - English as one of communications channel
 - Documentation in English

Operational Perspective (Con't)

- **More members participation**
 - Promote more overseas networks / OTT / Cloud / CDN to join
- **PeeringDB update**
 - Up-to-date information on PeeringDB, e.g., Local Facilities
 - Encourage members to update their peering information
- **Support RPKI**
 - Enhance network routing security
- **Public Looking Glass**
 - Traffic evaluation

Commercial Perspective

- **Organizations do not need to be a member of VNNIC**
 - Currently connecting members must be a member of VNNIC
 - Needs local entity to register as VNNIC member
- **Lower price and flat rate for port charge**
 - Lower cost for 1G port (VND16,000,000 = USD670)
 - Currently 10G port is charged based-on usage
- **Extension to other popular data centers**
 - Lower cost of access

Summary

Akamai Intelligent Platform

- Highly distributed edge servers

Path to 100+ IXes

- Akamai journey to reach to 100 Ixes
- Akamai and IXes

Wish-List for VNIX

- Technical / Operational / Commercial perspectives

Questions?

Kams Yeung <kams@akamai.com>

More information:

PeeringDB: <https://www.peeringdb.com/asn/20940>

Peering Request: <peering@akamai.com>



End

Thank You! Cảm ơn bạn!