

Leveraging Interconnections for Performance

The Serving Infrastructure of a Large CDN

Florian Wohlfart²¹, Nikolaos Chatzis¹, Caglar Dabanoglu¹, Georg Carle², Walter Willinger³

¹Akamai Technologies, ²Technical University of Munich, ³NIKSUN Inc.



Motivation

Recent interest in large CP's serving infrastructures (or peering edge)

Motivation

Recent interest in large CP's serving infrastructures (or peering edge)

- Google and Facebook papers at SIGCOMM'17

Motivation

Recent interest in large CP's serving infrastructures (or peering edge)

- Google and Facebook papers at SIGCOMM'17
- Scalable systems for traffic engineering at the serving infrastructure

Motivation

Recent interest in large CP's serving infrastructures (or peering edge)

- Google and Facebook papers at SIGCOMM'17
- Scalable systems for traffic engineering at the serving infrastructure

Actual size and structure of serving infrastructures largely unknown

Motivation

Recent interest in large CP's serving infrastructures (or peering edge)

- Google and Facebook papers at SIGCOMM'17
- Scalable systems for traffic engineering at the serving infrastructure

Actual size and structure of serving infrastructures largely unknown

- Footprint: machines, clusters/deployments, etc.
- Connectivity fabric: peerings

Findings

- Explicit and Implicit Peerings
 - Peerings absent from public BGP datasets

Findings

- Explicit and Implicit Peerings
 - Peerings absent from public BGP datasets
- Network Prefix Deaggregation
 - /25+ network prefixes

Findings

- Explicit and Implicit Peerings
 - Peerings absent from public BGP datasets
- Network Prefix Deaggregation
 - /25+ network prefixes
- The “private” Internet
 - Heavy use of private links

Approach


- 1) We describe and quantify of Akamai's serving infrastructure

Approach

- 1) We describe and quantify of Akamai's serving infrastructure
- 2) We show how Akamai uses the connectivity fabric for its CDN service


Deployments

colocation




Deployments

colocation




hosting network




Deployments


hosting network




Deployments in a Host Network




Deployments in a Host Network




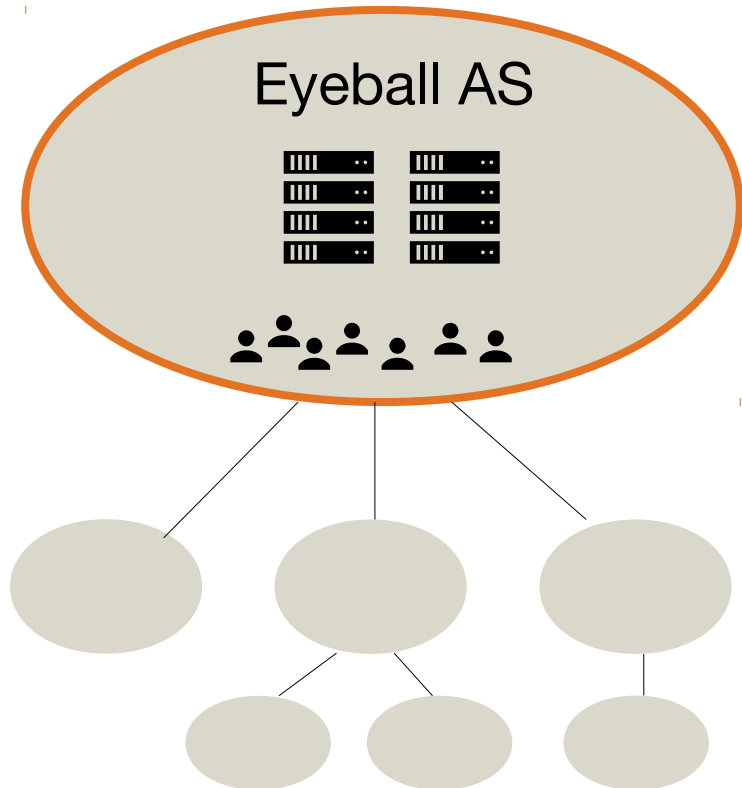
Deployments in a Host Network



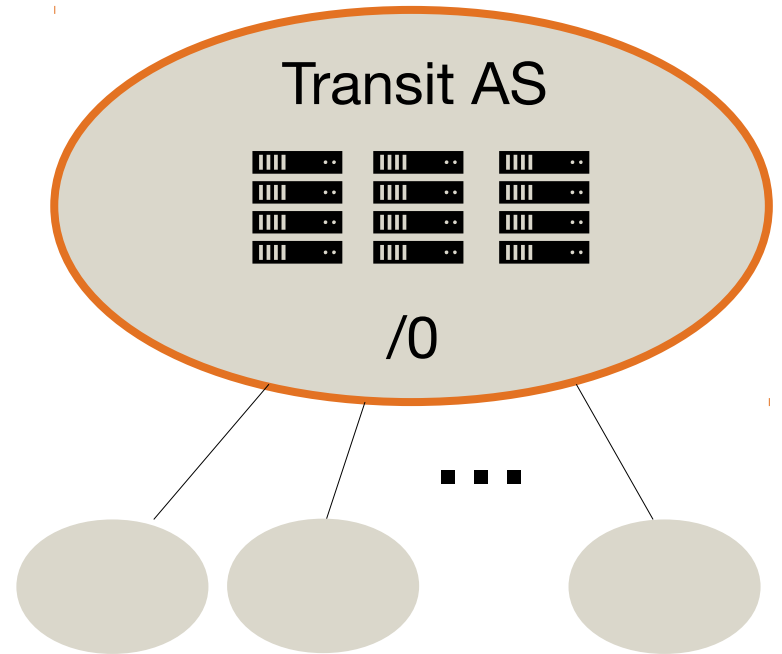

Deployments in a Host Network



Deployments in a Host Network




Deployments in a Host Network

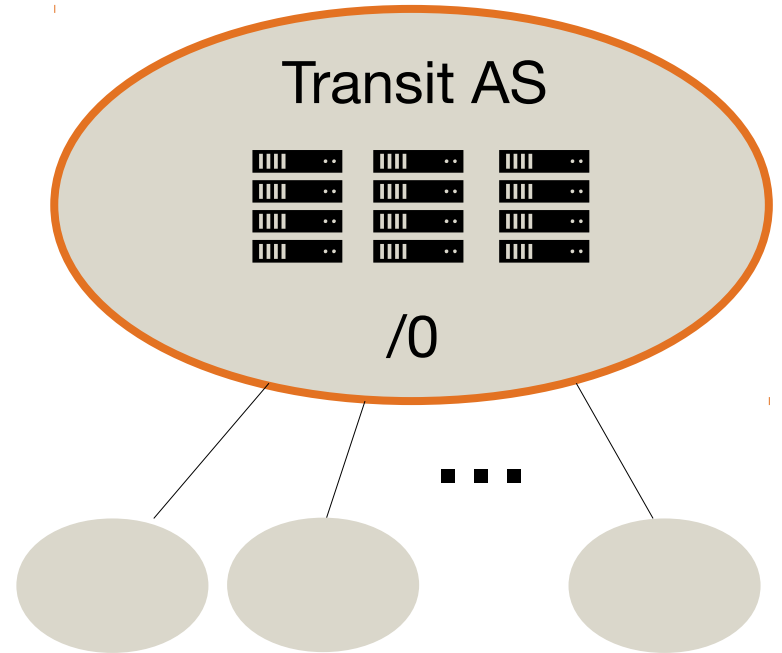


Deployments in a Host Network

Onnet (Type 1)

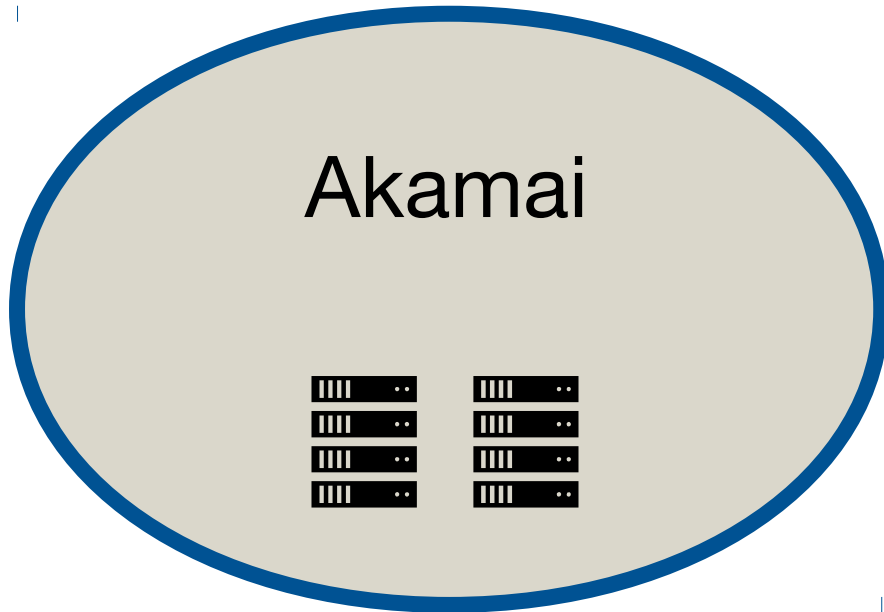


Transit (Type 2)





Deployments

colocation




hosting network




Deployments


colocation



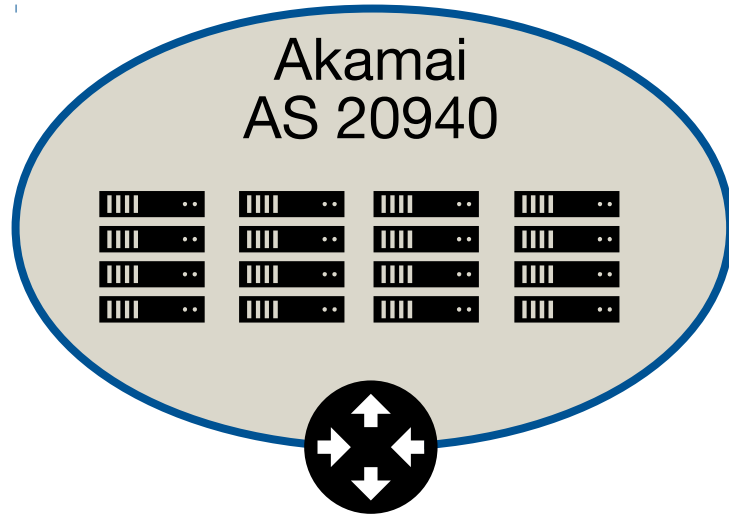
Deployments in a Colocation




Deployments in a Colocation



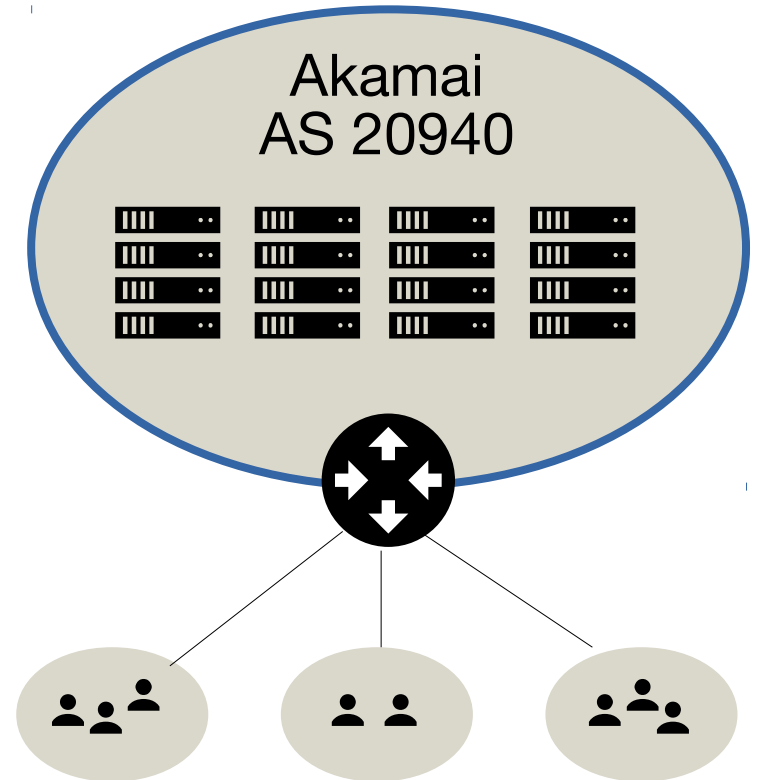
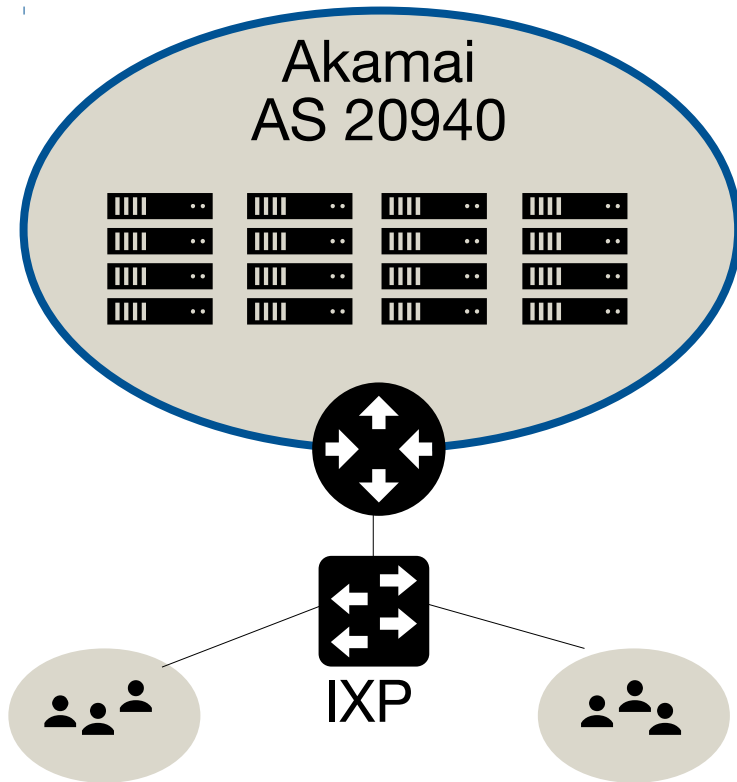
Deployments in a Colocation



Deployments in a Colocation




Deployments in a Colocation

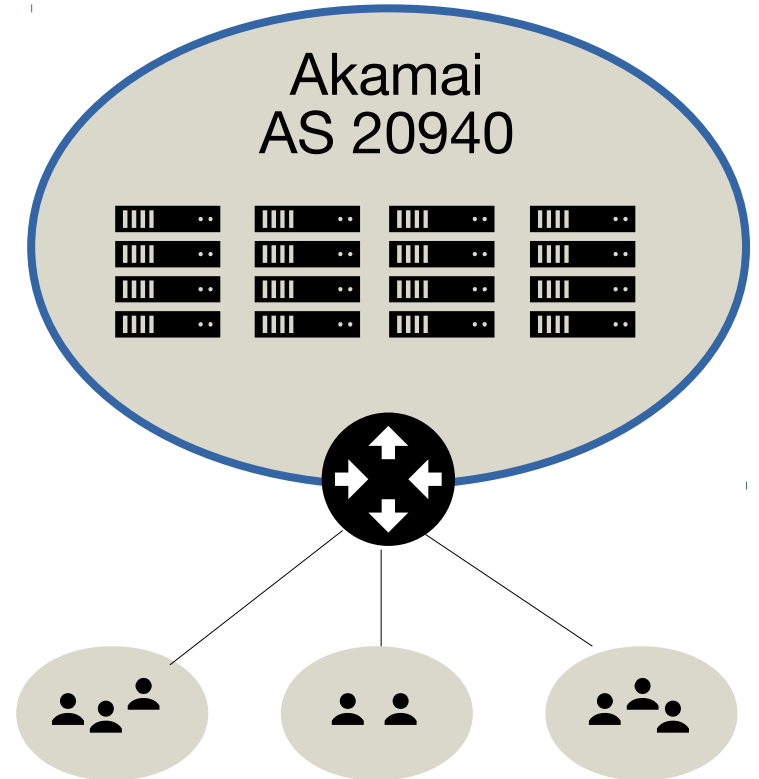


Deployments in a Colocation

IXP (Type 3)

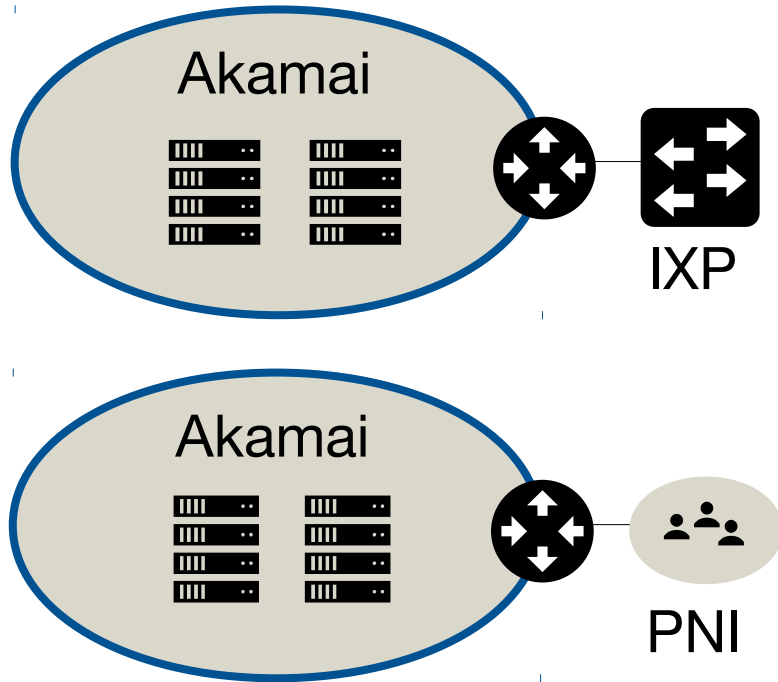


PNI (Type 4)

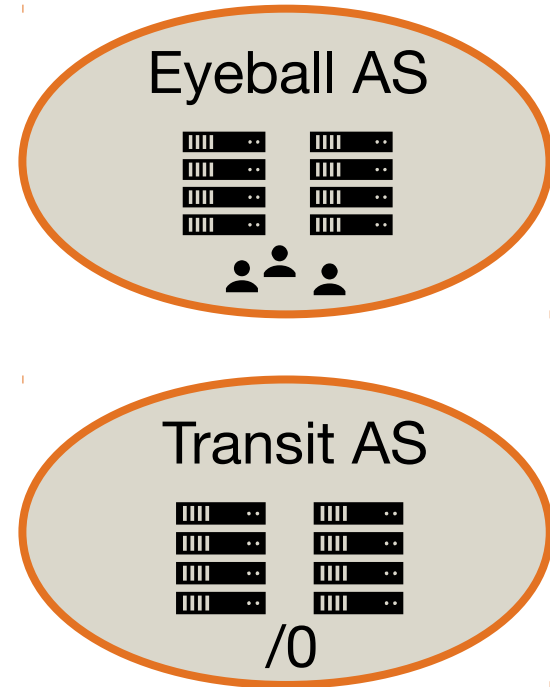


Deployments: Summary

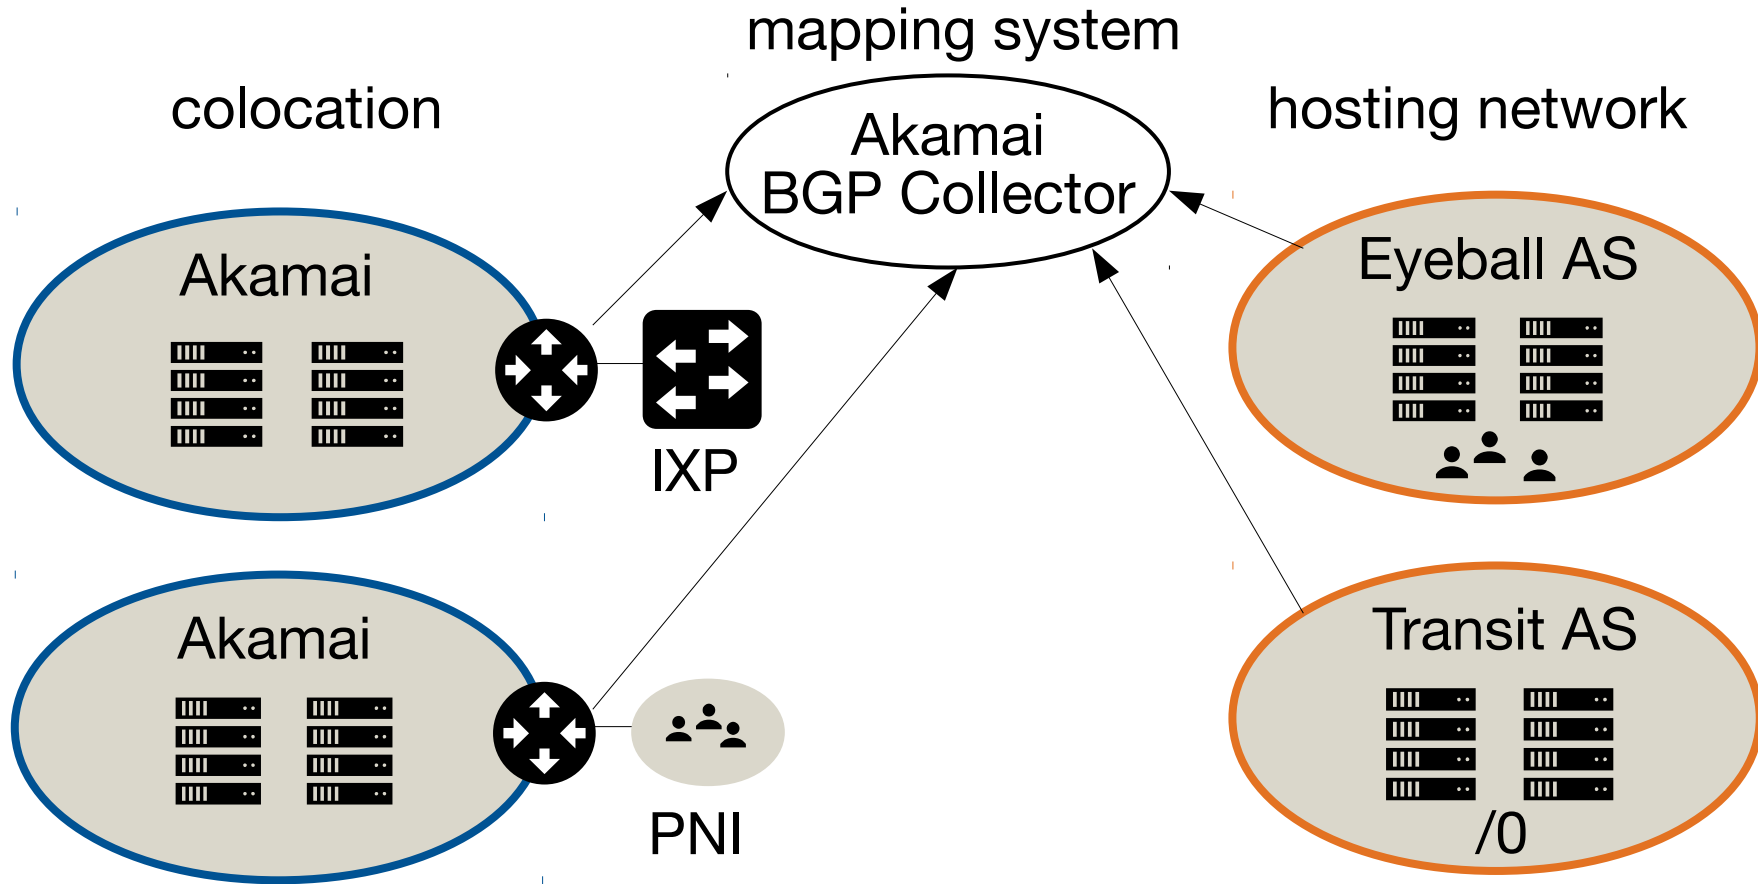
colocation



hosting network




Deployments: Summary



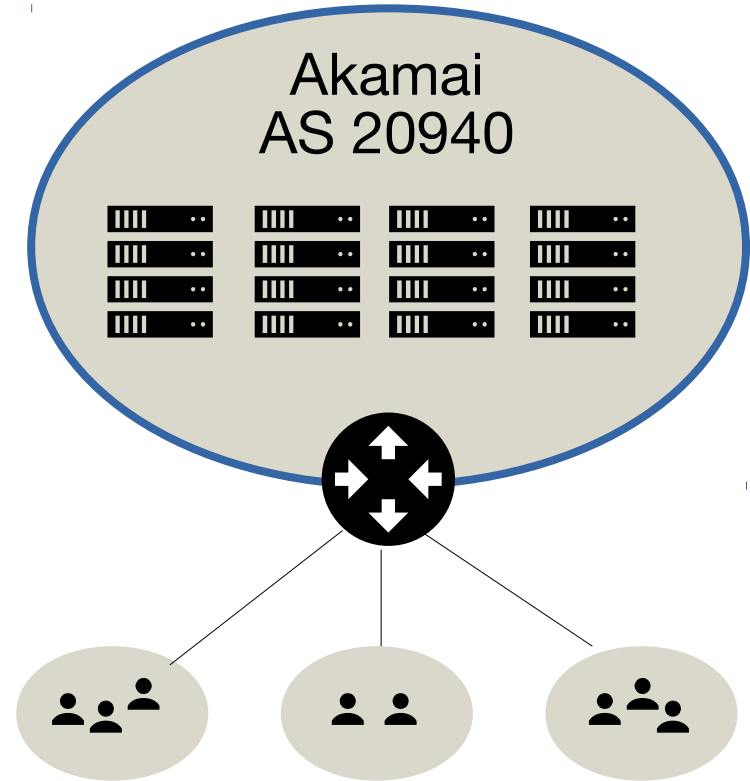
Deployments → Peerings

Explicit Peerings

IXP (Type 3)




PNI (Type 4)

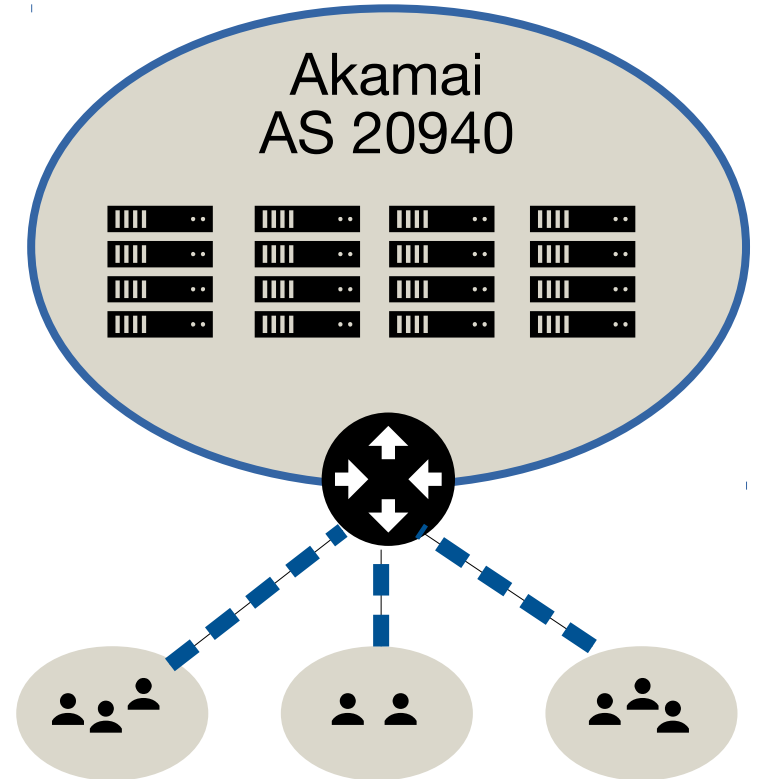


Explicit Peerings

IXP (Type 3)



PNI (Type 4)

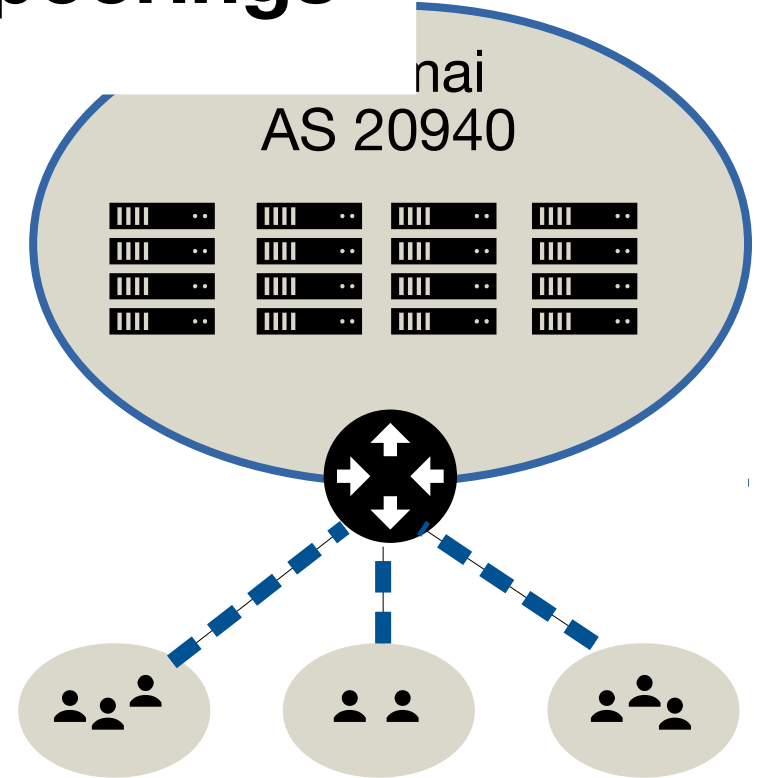



Explicit Peerings

IXP (Type 3)

6,111 explicit peerings

Type 4)




Explicit Peerings


6,111 explicit peerings

Explicit Peerings

6,111 explicit peerings




IXP (Type 3)



PNI (Type 4)


Explicit Peerings



6,111 explicit peerings

IXP (Type 3)


6,075



PNI (Type 4)

227


Explicit Peerings



6,111 explicit peerings

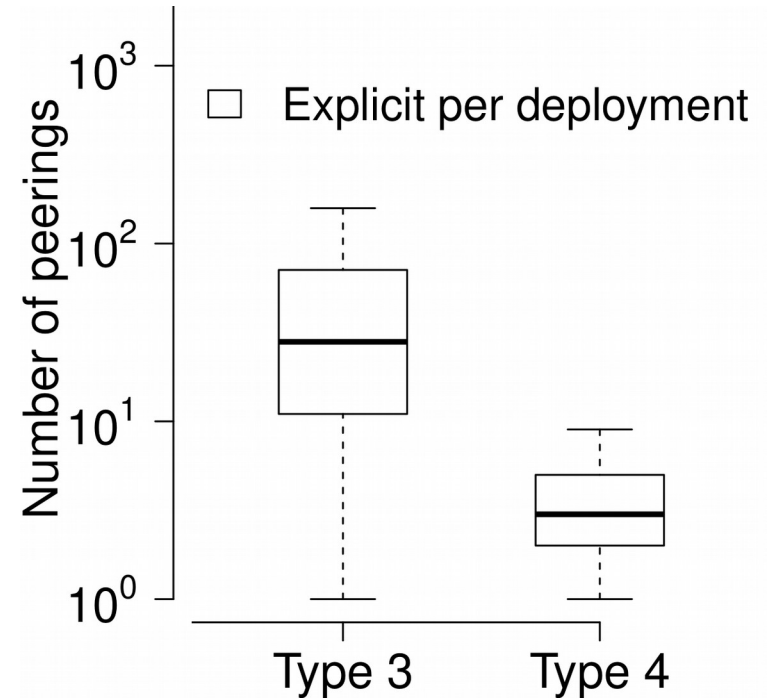
IXP (Type 3)

6,075




PNI (Type 4)

227




Implicit Peerings

Onnet (Type 1)




Transit (Type 2)




Implicit Peerings

Onnet (Type 1)





Transit (Type 2)



Implicit Peerings

Onnet (Type 1) **28,353 implicit peerings** (Type 2)

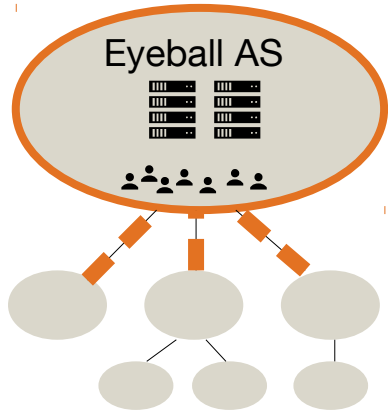


Implicit Peerings


28,353 implicit peerings

Implicit Peerings

28,353 implicit peerings




Onnet (Type 1)



Transit (Type 2)


Implicit Peerings

28,353 implicit peerings



Onnet (Type 1)

26,429




Transit (Type 2)

7,322

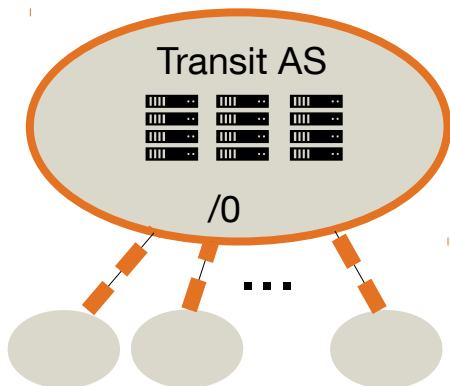
Implicit Peerings

28,353 implicit peerings




Onnet (Type 1)

26,429




Transit (Type 2)

7,322




Implicit Peerings

28,353 implicit peerings




Onnet (Type 1)

26,429



Transit (Type 2)

7,322



onnet deployments \gg # transit deployments

Peerings: Summary

- Previous work on existence of traffic flows via implicit peerings, but no quantification
- Found 8k explicit and 28k implicit peerings
- Implicit peerings are not observable from public BGP data

Peerings: Summary

- Previous work on existence of traffic flows via implicit peerings, but no quantification
- Found 8k explicit and 28k implicit peerings
- Implicit peerings are not observable from public BGP data

Up next

- Relevance of implicit and explicit peerings in terms of traffic

Network Traffic


BGP data +
data plane data

→ usage of peerings

Network Traffic

BGP data +
data plane data

→ usage of peerings

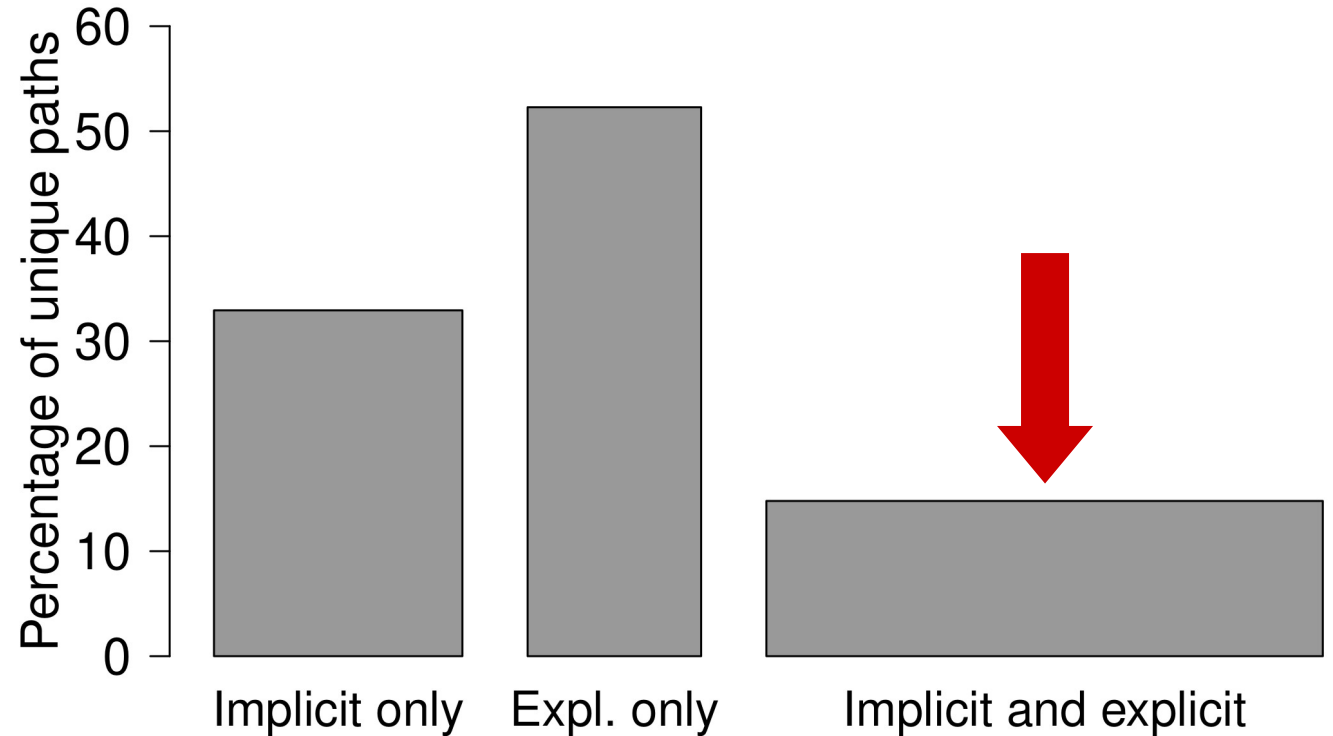


Bar width proportional to traffic volume

Network Traffic

BGP data +
data plane data

→ usage of peerings




Bar width proportional to traffic volume

Network Traffic

BGP data +
data plane data

→ usage of peerings

**Both types are
relevant for
serving traffic**




Bar width proportional to traffic volume

Deaggregated Prefixes

Comparison with
public BGP data

Deaggregated Prefixes


Comparison with
public BGP data



Deaggregated Prefixes

Comparison with
public BGP data

Learned from onnet
deployments only




Deaggregated Prefixes

Comparison with public BGP data

Learned from onnet deployments only

Akamai steers traffic based on /25+ prefix information



The “private” Internet

Food for thought

The “private” Internet

Food for thought

- A large content owner utilizes a cloud provider to store its content

The “private” Internet

Food for thought

- A large content owner utilizes a cloud provider to store its content
- A CDN retrieves the content and transports it across its private backbone network

The “private” Internet

Food for thought

- A large content owner utilizes a cloud provider to store its content
- A CDN retrieves the content and transports it across its private backbone network
- The CDN uses PNIs to deliver the content to large eyeball networks

The “private” Internet

Food for thought

- A large content owner utilizes a cloud provider to store its content
- A CDN retrieves the content and transports it across its private backbone network
- The CDN uses PNIs to deliver the content to large eyeball networks

Network researchers lose visibility into large portions of Internet traffic

The “private” Internet: Evidence


- Akamai connects with cloud providers via PNI

The “private” Internet: Evidence

- Akamai connects with cloud providers via PNI
- Private backbone network announced at NANOG 71

The “private” Internet: Evidence

- Akamai connects with cloud providers via PNI
- Private backbone network announced at NANOG 71
- Akamai serves mainly via onnet and PNI links (example showing one country)
IXPs mainly used to serve the “long tail”



Conclusion and Outlook

Akamai relies on both explicit and implicit peerings to serve traffic

Akamai steers traffic based on deaggregated network prefixes

Shift towards a “private” Internet

Conclusion and Outlook

Akamai relies on both explicit and implicit peerings to serve traffic

Akamai steers traffic based on deaggregated network prefixes

Shift towards a “private” Internet

These problems seem to be impossible to study based on public data

Can we find ways to reproduce such findings without proprietary data?