

# Measuring IPv6 Adoption

Jakub Cxyz, University of Michigan

Mark Allman, International Computer Science Institute

Jing Zhang, University of Michigan

Scott Iekel-Johnson, Arbor Networks


Eric Osterweil, Verisign Labs

Michael Bailey, University of Michigan and University of Illinois

**SIGCOMM 2014**  
**Chicago, IL, USA**  
**August 17-22, 2014**

# Why Study IPv6 Adoption Now?

- Internet continues growing
- IPv4 space shrinking...
- IPv4 exhaustion events:
  - IANA: February 2011
  - Asia/Pacific: April 2011
  - Europe: September 2012
  - Latin America: June 2014
- IPv6 Community Flag Days
  - 2011 & 2012



(Image source: Geoff Huston, <http://www.potaroo.net/tools/ipv4>)

# Our Study

- **Goal:** a systemic “big picture” of IPv6 adoption
  - Trading off depth for breadth
  - Are there cross-perspective insights?
- **Multi-perspective:** 10 datasets
- **Multi-year:** 2-10 years
- **Multi-aspect:** 12 metrics
- **Findings: IPv6 adoption**
  - varies by where you measure (region)
  - varies by what you measure
  - recently made a qualitative jump

# Data Analyzed

- **Existing/Public Datasets:**
  - RIR allocation
  - Route Views BGP, RIPE-RIS BGP
  - Google.com clients,
  - Verisign zone files,
  - CAIDA Ark RTT
- **New Datasets:**
  - **Traffic:** Arbor Networks global traffic
  - **Naming:** Verisign .com/.net queries via IPv4, via IPv6
  - **Content:** Testing data of Alexa top-10K sites

# Metrics

## Prerequisite IP Functions

- Address Allocation
- Address Advertisement
- Topology
- DNS Name servers
- DNS Resolvers
- DNS Queries
- Server Readiness
- Client Readiness


## Operational Characteristics

- Traffic Volume
- Application Mix
- Transition Technologies
- Performance (RTT)

*“IPv6 adoption” = level relative to IPv4*

# **METRICS (PREREQUISITE)**

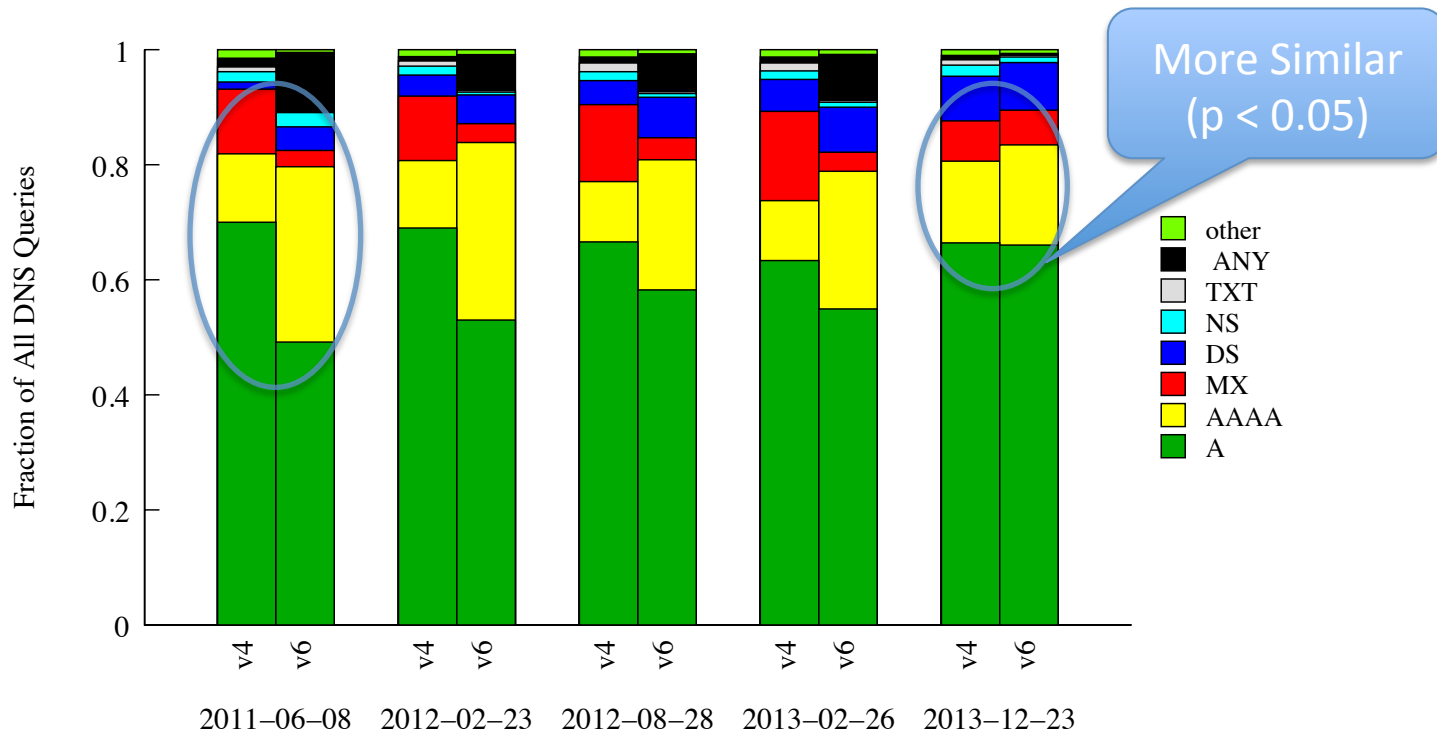
# Prefix Allocation



IANA  
Exhaustion


# Naming: Domains & Record Types

- Queries from .com/.net; IPv4 & IPv6 name servers
  - Five day-long packet samples over 2.5 years
  - IPv6 DNS users query similar **domains** as IPv4
  - Query **types** are converging over this time period:

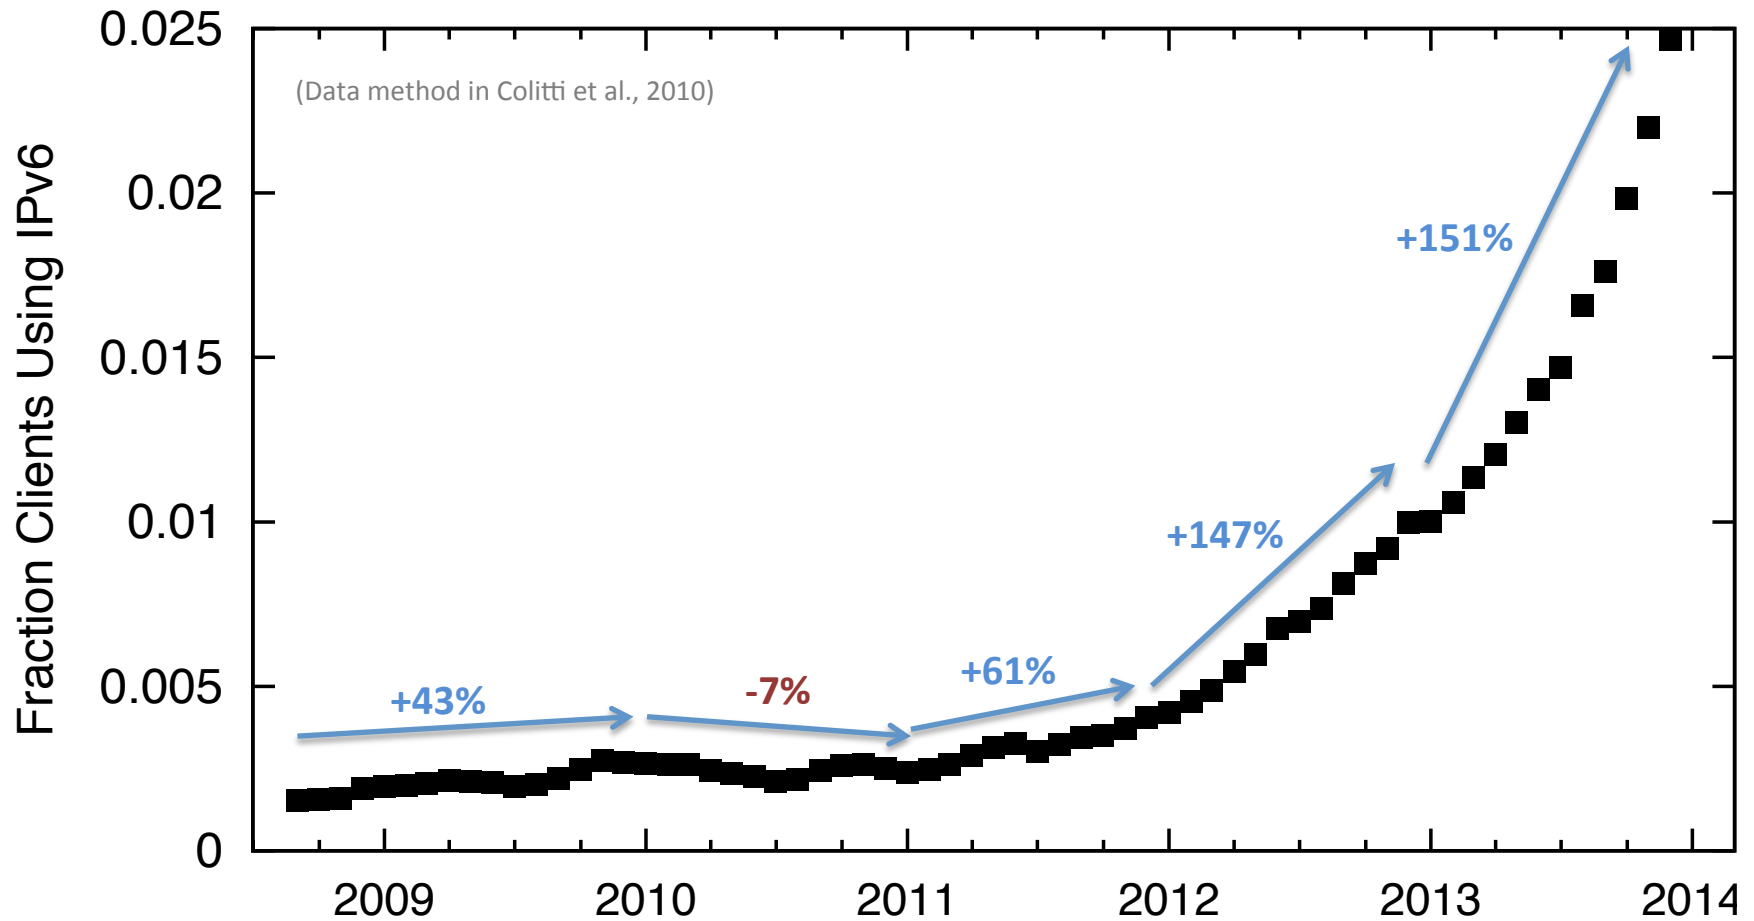




# Server Readiness: Alexa Top Domain Reachability



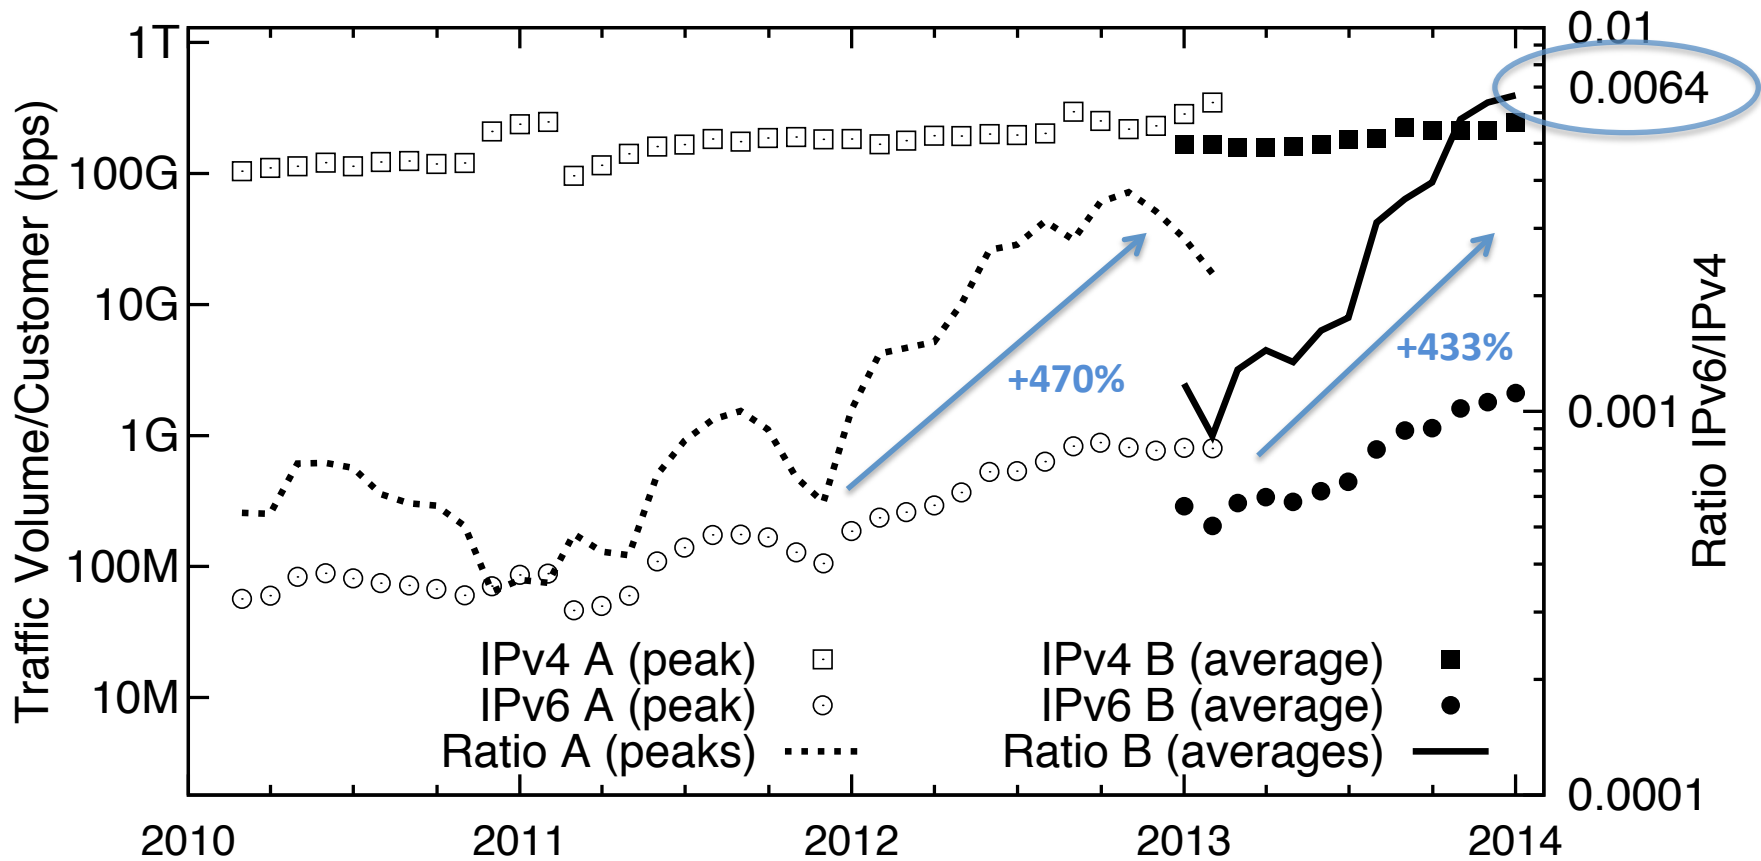
# Client Readiness: visitors to google.com



# **METRICS (OPERATIONAL)**

# Global Traffic

- Arbor Networks global provider netflow data
  - 260 service providers (Dataset B) ~ 1/3 – 1/2 of all inter-AS traffic

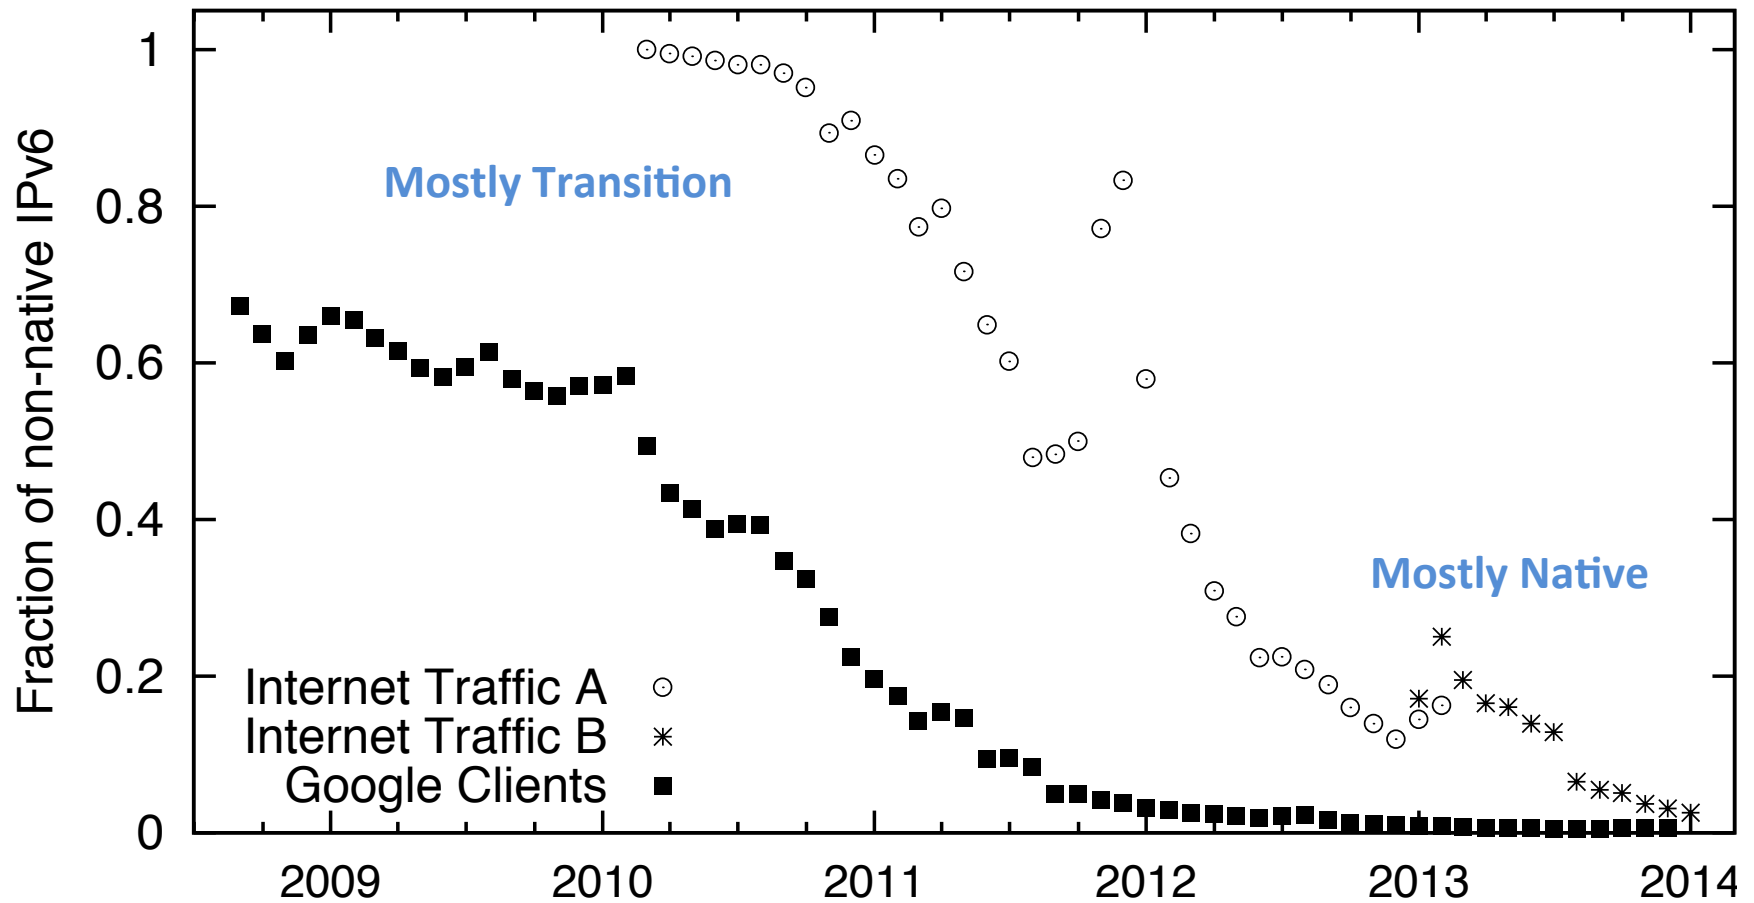


# Application Mix (% of IPv6)

User content {

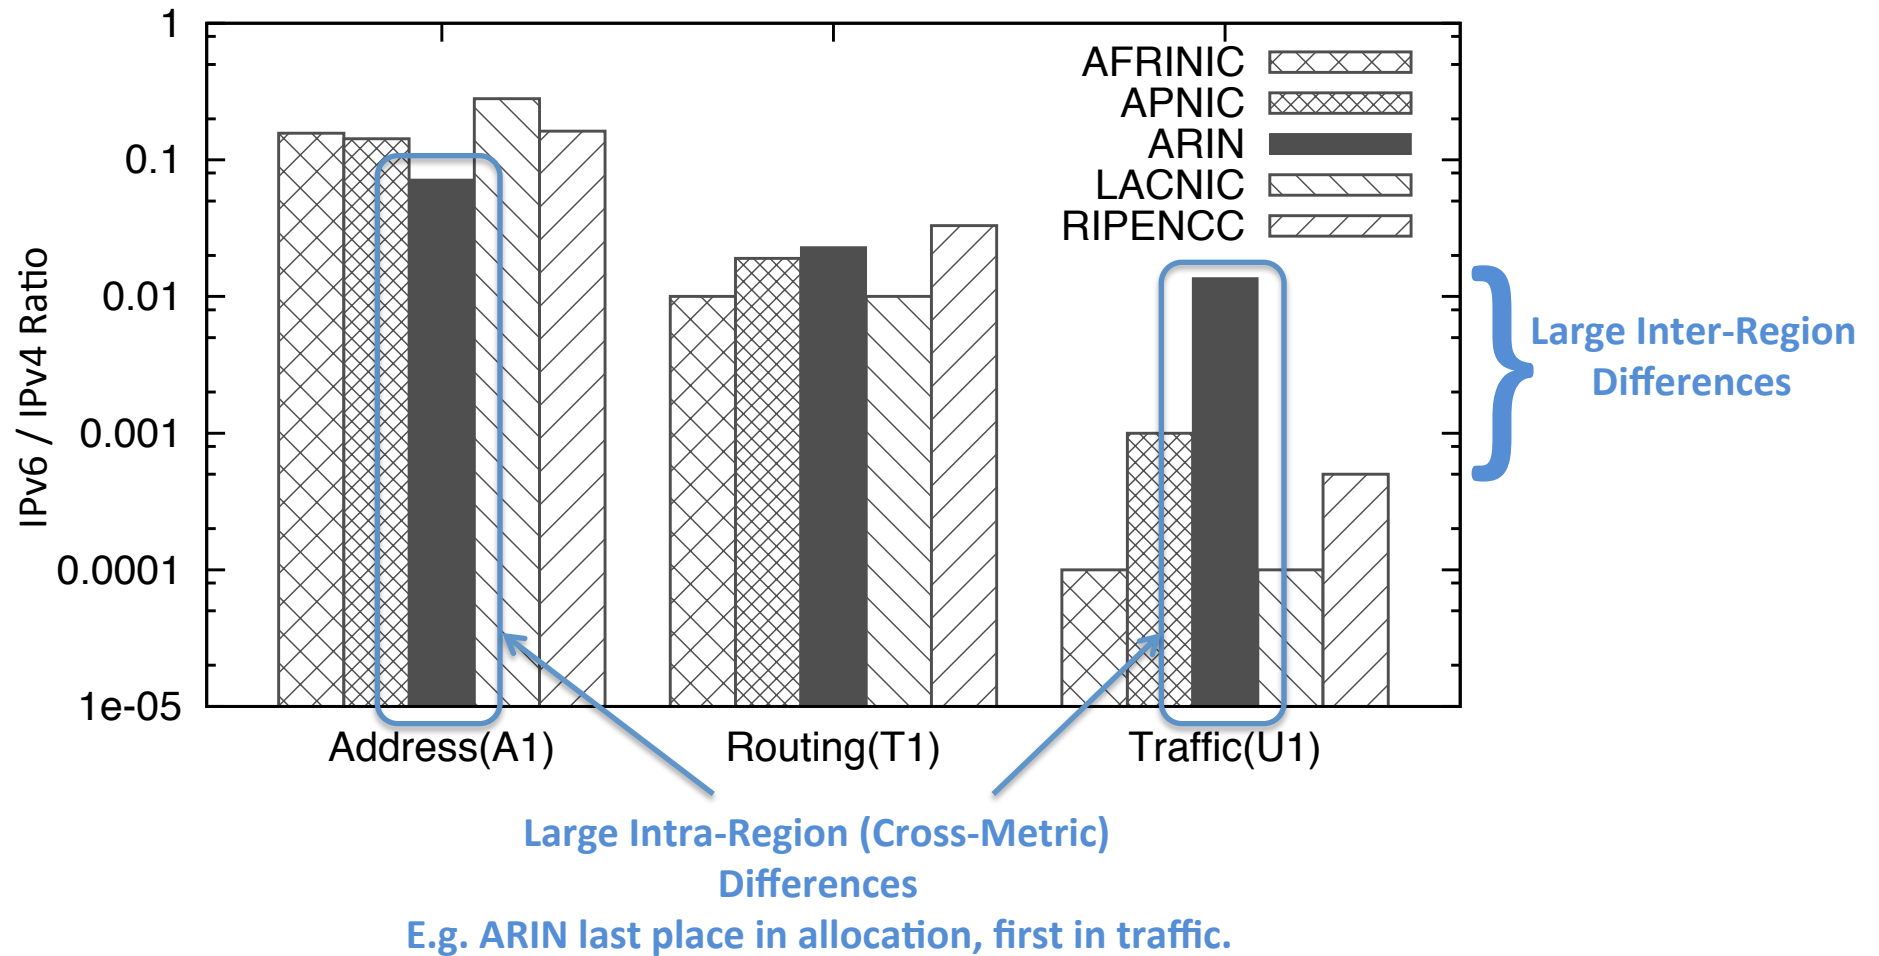
Application	Dec 2010	Apr–Dec 2013	
	IPv6	IPv6	IPv4
<b>HTTP</b>	5.61	82.56	60.61
<b>HTTPS</b>	0.15	12.66	8.59
<b>DNS</b>	4.75	0.33	0.22
<b>SSH</b>	0.56	0.27	0.20
<b>Rsync</b>	20.78	0.13	0.00
<b>NNTP</b>	27.65	0.00	0.25
<b>RTMP</b>	0.00	0.00	2.74
<b>Other TCP</b>	*	1.66	4.08
<b>Other UDP</b>	*	0.27	2.82
<b>Non-TCP/UDP</b>	*	2.11	20.21

# IPv6 Transition Technologies (Teredo + 6to4)




# CONCLUSIONS

# Conclusion 1: Regions Differ





# Conclusion 2: Perspective Matters



# Conclusion 3: IPv6 is Real!


Operational Aspect Measured	IPv6 Status at End of:		
	2010	2013	
IPv6 Percent of Internet Traffic	0.03%	0.64%	← 20x growth!
1-yr. Growth vs. IPv4 (*Mar-2010 – Mar-2011)	-12%*	+433%	
Content's Portion of Traffic (HTTP+HTTPS)	6%	95%	← 15x growth!
Native IPv6 Packets vs. All IPv6	9%	97%	← Traffic Flipped
Native IPv6 Google Clients	78%	99%	
Performance: 10-hop RTT <sup>-1</sup> vs. IPv4	75%	95%	← Nearly on-par

# Thank You!

Questions?

# BACKUP SLIDES

# Internet Devices and Users Continue to Increase




Source: Cisco VNI, 2014

The percentages in parentheses next to the legend denote the device share for the years 2013 and 2018, respectively.


# Dataset Summary

Dataset	Metrics	Time Period	Recent Scale
RIR Address Allocations	A1	Jan 2004 – Jan 2014	≈18K allocation snapshots (5 daily)
Routing: Route Views	A2, T1	Jan 2004 – Jan 2014	45,271 BGP table snapshots
Routing: RIPE	A2, T1	Jan 2004 – Jan 2014	
Google IPv6 Client Adoption	R2, U3	Sep 2008 – Dec 2013	millions of daily global samples
Verisign TLD Zone Files	N1	Apr 2007 – Jan 2014	daily snapshots of ≈2.5 million A+AAAA glue records (.com & .net)
CAIDA Ark Performance Data	P1	Dec 2008 – Dec 2013	≈10 million IPs probed daily
Arbor Networks ISP Traffic Data	U1, U2, U3	Mar 2010 – Dec 2013	≈33-50% of global Internet traffic; 2013 daily median: 50 terabits/sec (avg.)
Verisign TLD Packets: IPv4	N2, N3	Jun 2011 – Dec 2013	4 global sites, 5 of 13 gTLD NS letters (.com/.net), ≈4.5Bn queries/day
Verisign TLD Packets: IPv6	N2, N3	Jun 2011 – Dec 2013	15 global sites, both gTLD NS letters (.com/.net) w/IPv6, 647M queries/day
Alexa Top Host Probing	R1	Apr 2011 – Dec 2013	10,000 servers probed twice/month

# Prefix Advertisement

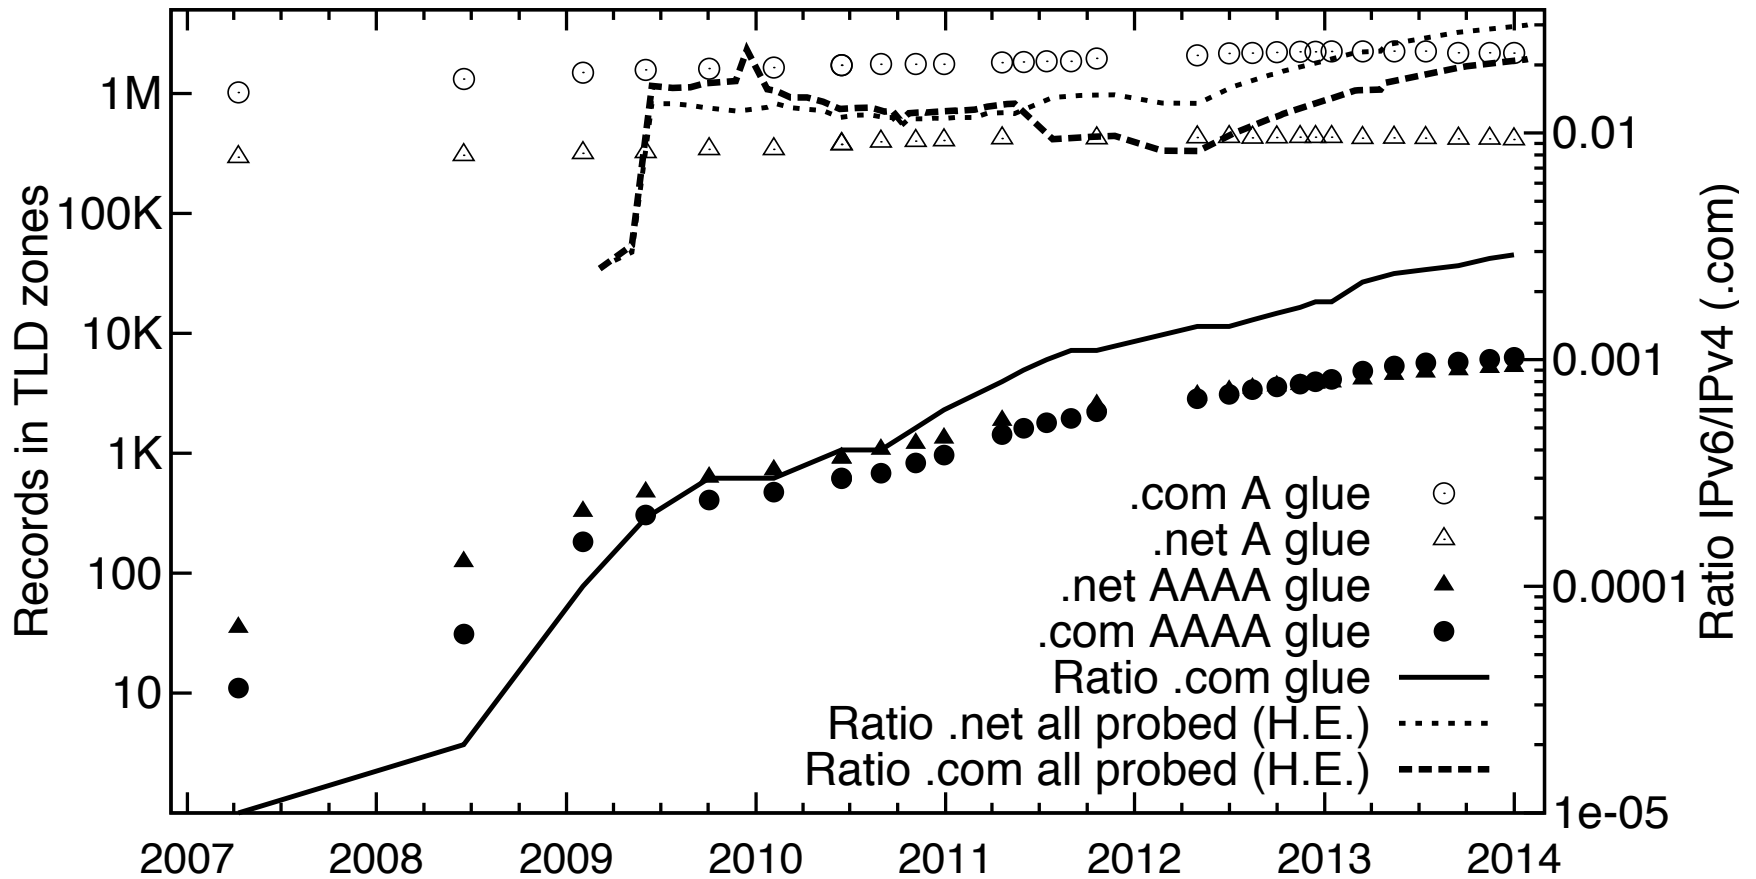


# AS Centrality







# DNS: .com & .net Zones



# IPv4 .com/.net TLD A and AAAA Query Rank Correlation (Spearman's $\rho$ )


Domain Lists	2011-06-08	2012-02-23	2012-08-28	2013-02-26	2013-12-23
<b>4.A : 6.A</b>	0.65	0.73	0.70	0.70	0.57
<b>4.AAAA : 6.AAAA</b>	0.69	0.80	0.82	0.74	0.68
<b>4.A : 4.AAAA</b>	0.32	0.32	0.35	0.34	0.42
<b>6.A : 6.AAAA</b>	0.29	0.23	0.20	0.26	0.32

 Within type: **Strong**  
 Across type: **Weak**

# Naming: Domains

- Queries from .com/.net; IPv4 & IPv6 name servers
  - Five day-long samples over 2.5 years
- Four sets of top 100k domains:
  - For both IPv4 and IPv6 packets (user populations)
  - Within each, for domains queried by A and AAAA
- Finding: IPv4 to IPv6 populations correlate strongly for the same query type {A,AAAA}
  - e.g. Spearman's  $\rho$  of 0.7 for IPv4 A versus IPv6 A
- So, IPv6 DNS users query similar domains as IPv4

# Performance (using 10- and 20-hop RTT)



# Projections

