Akamai CDN, IPv6 and DNS security

Christian Kaufmann
Akamai Technologies
APNIC 36
26th August 2013

Agenda

Akamai Introduction
- Who's Akamai?
- Intelligent Platform & Traffic Snapshot

Basic Technology
- Akamai mapping
- Finding the IP address
- Downloading www.example.com

Akamai & IPv6 World Launch Anniversary
- Akamai IPv6 Deployment and Observations

Secure the Internet
- Open recursors and reflection attacks
- BCP-38 and DNS server maintenance

Akamai Introduction

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:

110,000+ Servers	2,000+ Locations	1,100+ Networks	700+ Cities	83 Countries

Typical daily traffic:

- More than **2 trillion** requests served
- Delivering over **15 terabits/second**
- **15-30%** of all daily web traffic

Basic Technology

Akamai mapping

How CDNs Work

When content is requested from CDNs, the user is directed to the optimal server

- This is usually done through the DNS, especially for non-network CDNs, e.g. Akamai
- It can be done through anycasting for network owned CDNs

Users who query DNS-based CDNs be returned different A (and AAAA) records for the same hostname

This is called “mapping”

The better the mapping, the better the CDN

How Akamai CDN Work

Example of Akamai mapping

- Notice the different A records for different locations:

```
[NYC]% host www.symantec.com
www.symantec.com    CNAME  e5211.b.akamaiedge.net.
e5211.b.akamaiedge.net.  A      207.40.194.46
e5211.b.akamaiedge.net.  A      207.40.194.49

[Boston]% host www.symantec.com
www.symantec.com    CNAME  e5211.b.akamaiedge.net.
e5211.b.akamaiedge.net.  A      81.23.243.152
e5211.b.akamaiedge.net.  A      81.23.243.145
```

How Akamai CDN Work

Akamai use multiple criteria to choose the optimal server

- These include standard network metrics:
  - Latency
  - Throughput
  - Packet loss
- These also include things like CPU load on the server, HD space, network utilization, etc.

Finding the IP Address: The Akamai Way

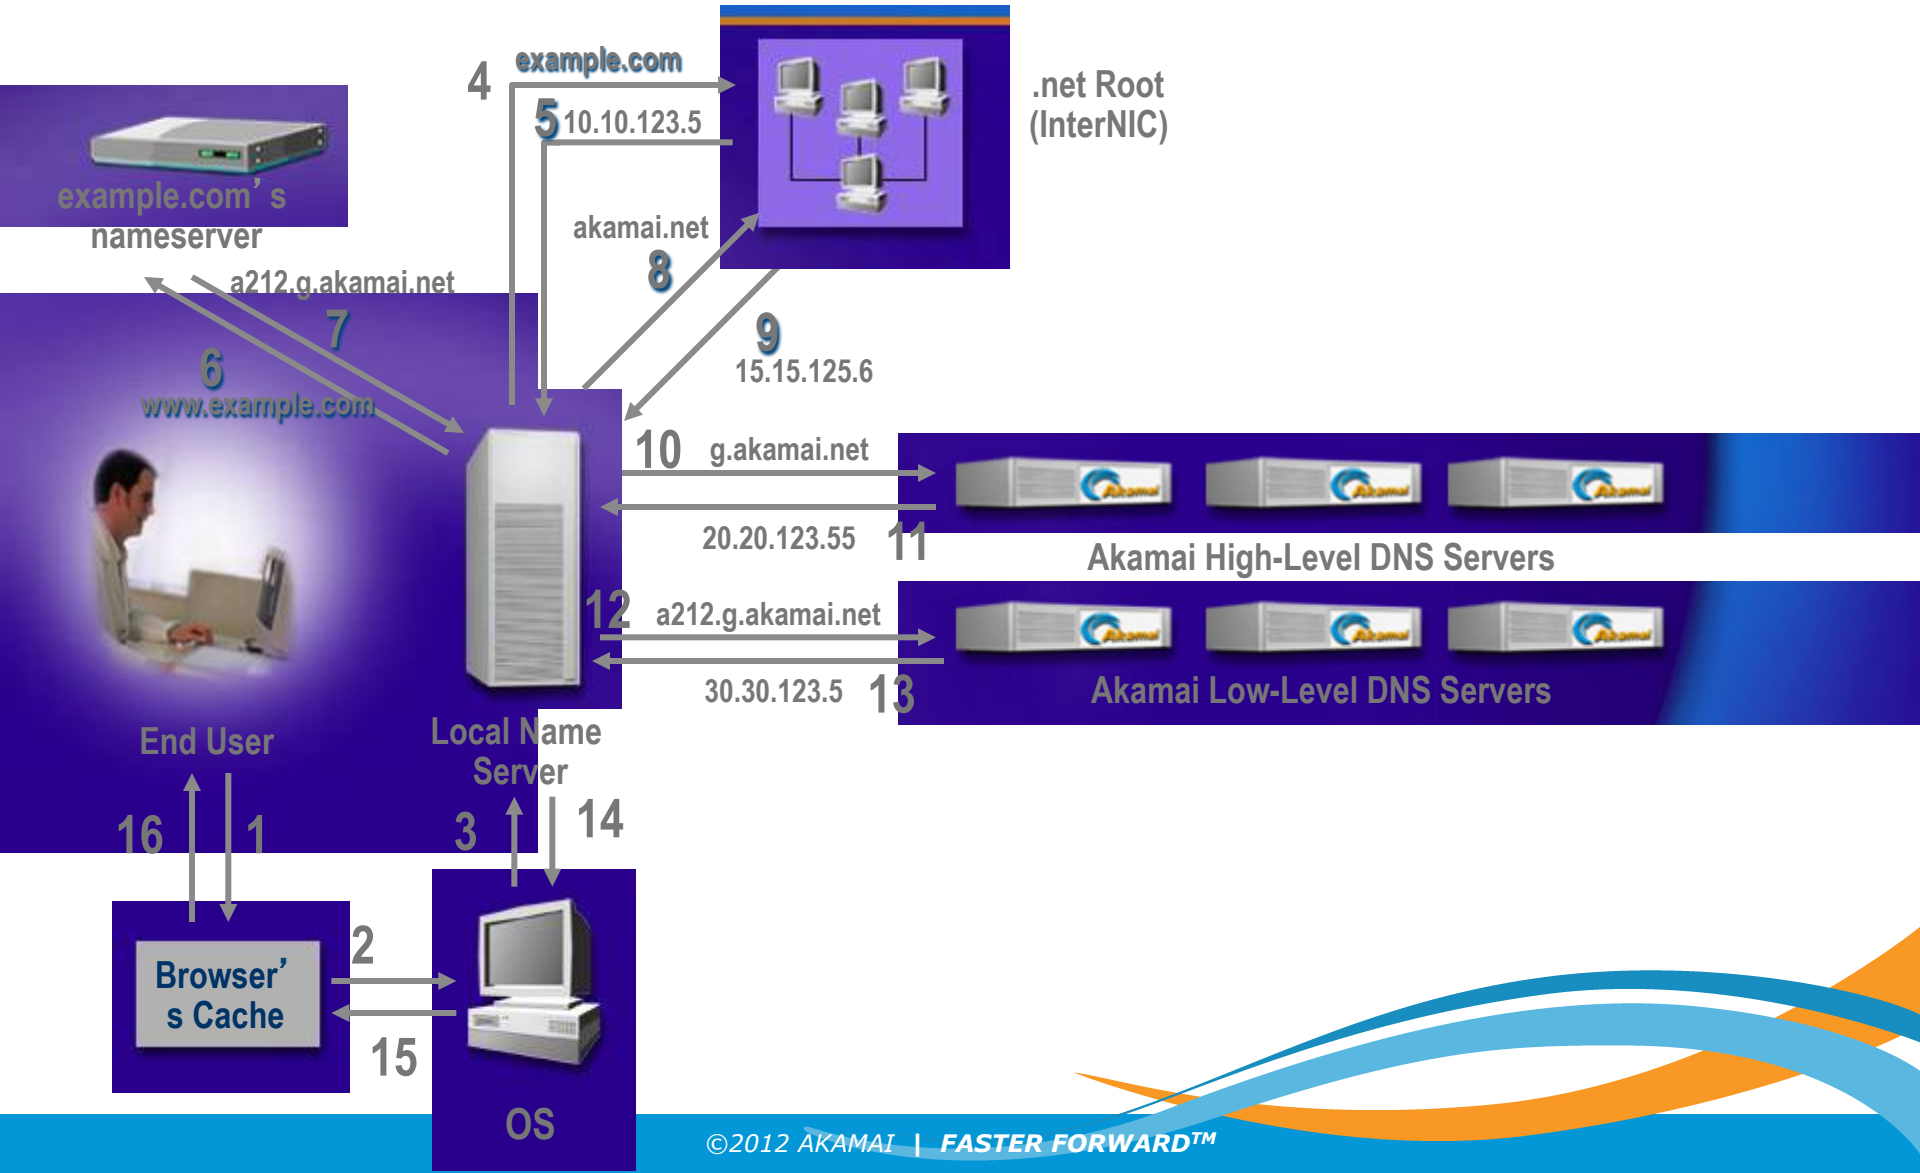

Downloading www.example.com with Akamai's EdgeSuite

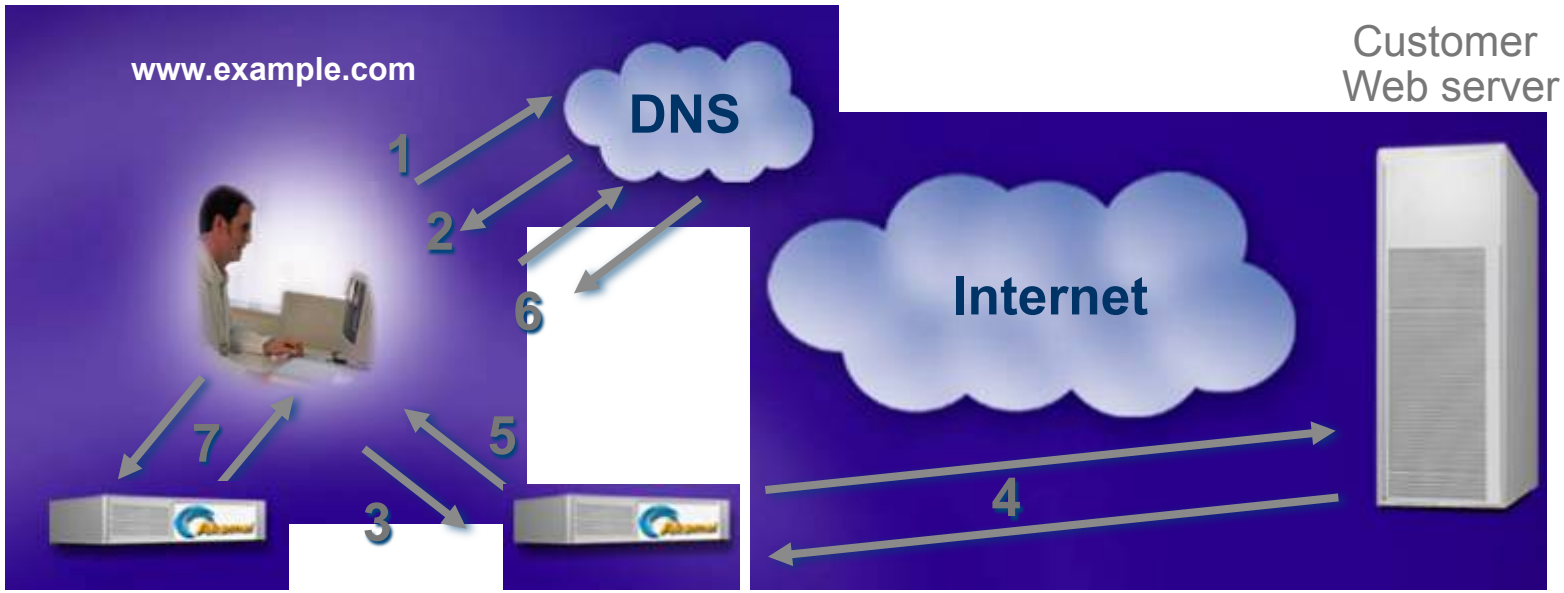

- User enters www.example.com
- 1. Browser requests IP address for www.example.com
- 2. DNS returns IP address of optimal Akamai server
- 3. Browser requests HTML
- 4. Akamai server assembles page, contacting customer Web server if necessary
- 5. Optimal Akamai server returns Akamaized HTML
- 6. Browser obtains IP address of optimal Akamai servers for embedded objects
- 7. Browser obtains objects from optimal Akamai servers

Akamai & IPv6

World IPv6 Launch Anniversary

How we enable IPv6

Dual-stacking edge servers

Customer properties can be dual-stacked

- Terminate IPv4 and IPv6 connections in server software
- Can go forward to customer origin via IPv4 *(or IPv6)*

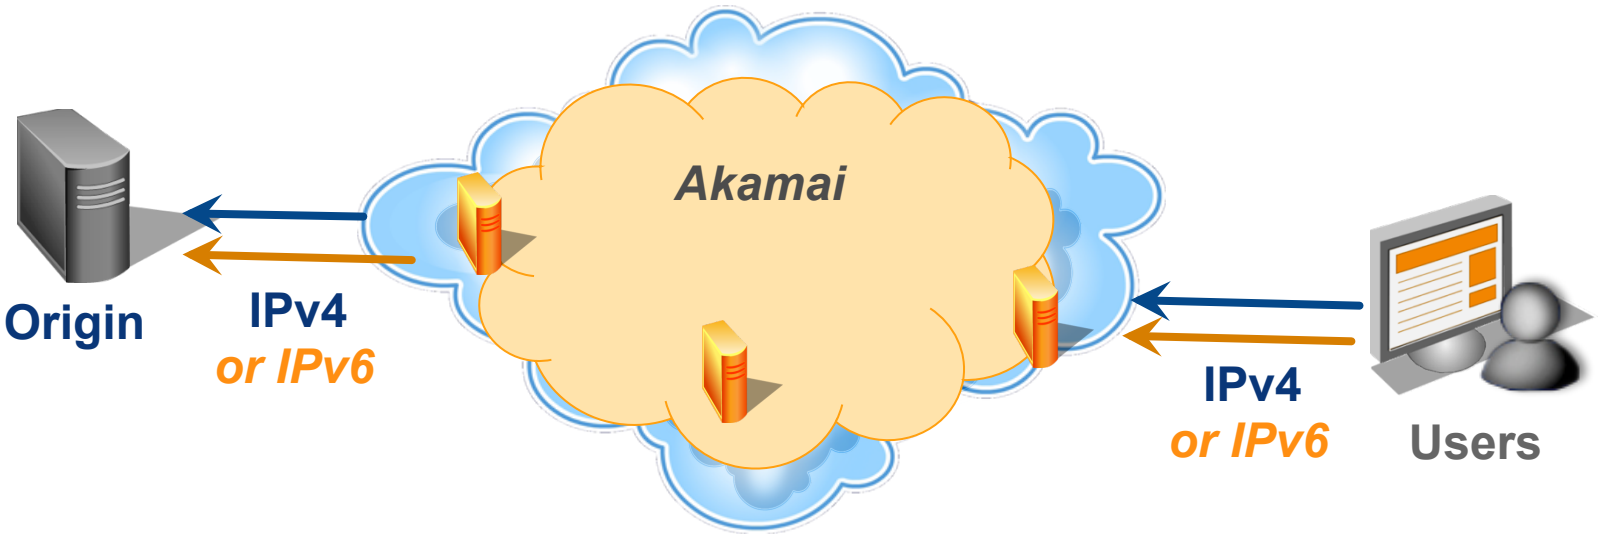

World IPv6 Launch Day: deployment status

In-production serving HTTP over IPv6 to users, tried to dual-stack every server everywhere

As of 2012-06-06, IPv6 now live in…

… over 53 countries

… over 175 cities (in all continents except Antarctica)

… over 225 networks

… over 600 Akamai server locations

… over 37,000 Akamai servers

Compare to a total of 1070 networks in 83 countries

(many network providers don't have working IPv6 yet, not all networks have full IPv6 routing table)

Current deployment status

In-production serving HTTP over IPv6 to users, tried to dual-stack every server everywhere

As of Jun 2013, IPv6 now live in…

… over 64 countries

… over 240 cities (in all continents except Antarctica)

… over 300 networks

… over 800 Akamai server locations

… over 70,000 Akamai servers

Compare to a total of 1100+ networks in 83 countries

(many network providers don't have working IPv6 yet, not all networks have full IPv6 routing table)

World IPv6 Launch Anniversary: A closer look from Akamai

IPv6 Addresses

- 2011: 280,229
- 2012: 18,899,253
- ***67x***
- 2013: 200m – 300m
- ***10x***

IPv6 Requests/Day

- 2011: 8,343,590
- 2012: 3,394,971,156
- ***460x***
- 2013: >10 billions
- ***2.5x***

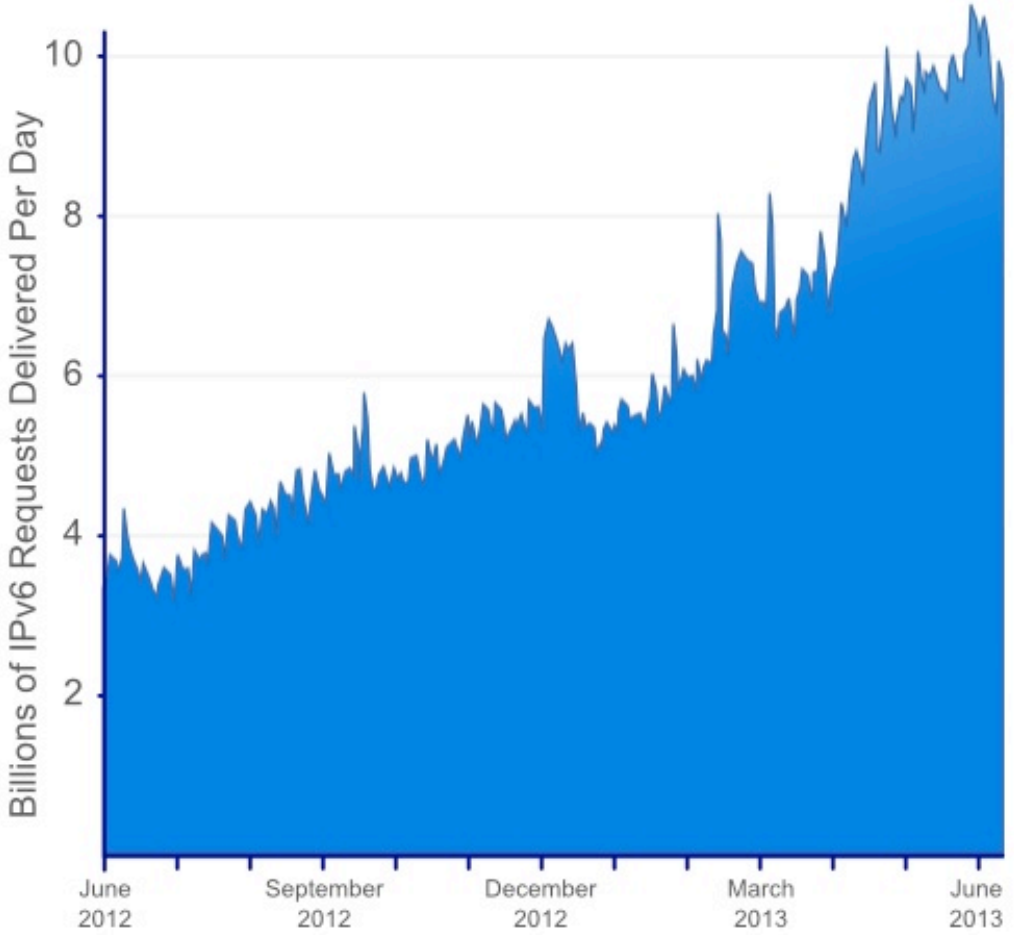

IPv6 Requests/Day on Akamai from June 2012 to June 2013

World IPv6 Launch Anniversary: Observations

Top 10 IPv6 by Geo

Country	IPv6 as % of Requests
Switzerland	10.4%
Romania	7.7%
France	4.6%
Luxembourg	3.6%
Belgium	3.3%
United States of America	3.2%
Germany	2.9%
Japan	2.1%
Peru	2.1%
Norway	1.4%

World IPv6 Launch Anniversary: Observations

Top 10 IPv6 by Network Provider

Network Operator	IPv6 as % of Requests	Primary Country
Verizon Wireless	34.9%	U.S.A.
Brutele (VOO)	29.7%	Belgium
Free/Proxad	18.9%	France
RCS & RDS	18.5%	Romania
Swisscom	15.8%	Switzerland
KDDI	9.9%	Japan
AT&T	8.4%	U.S.A.
Comcast	3.2%	U.S.A.
Deutsche Telekom AG	3.4%	Germany
Telefonica del Peru	2.6%	Peru

World IPv6 Launch Anniversary: IPv6 and Mobile

Mobile Operating System	IPv6 as % of Requests
Windows Phone OS 8	12%
BlackBerry OS 10	5.9%
Android 4.1/4.2 ("JellyBean")	10.8%
Android 4.0 ("Ice Cream Sandwich")	3.2%
Android 2.3 ("Gingerbread")	1.6%
Apple iOS 6	1.8%
Apple iOS 5	**1.4%**
Apple iOS 3/4	1.1%

- using Akamai's Mobile Browser Detection for categorization
- Within Android, there are individual device types where well over 50% of the traffic to dual-stacked websites arrived over IPv6.

World IPv6 Launch Anniversary: IPv6 and Desktop/Laptop Operating Systems

Operating System	Browser	IPv6 as % of Requests
Microsoft Windows 8		4.1%
Microsoft Windows Vista		3.3%
Microsoft Windows 7		2.5%
Microsoft Windows XP		0.5%
Mac OS X 10.5 & 10.6	Chrome & Firefox	3.4%
Mac OS X 10.5 & 10.6	Safari	3.3%
Mac OS X 10.7 & 10.8	Chrome & Firefox	3.3%
Mac OS X 10.7 & 10.8	Safari	2.1%

- Happy Eyeballs

World IPv6 Launch Anniversary: Three drivers of IPv6 growth

1. Content availability
   - Customers opting in to have their sites, content, and applications permanently available dual-stacked.
2. Availability of IPv6 from access network providers
   - IPv6 in production networks, e.g. Verizon Wireless, AT&T, and Comcast.
   - Some ISPs, Universities and Research Labs in Europe and Asia that have had IPv6 deployed
3. End-user device support
   - Recent desktop and laptop OS and client software supports IPv6
   - Many home routers / gateways start to support IPv6 recently.
   - 4G LTE smart phones.

Observations from World IPv6 Launch

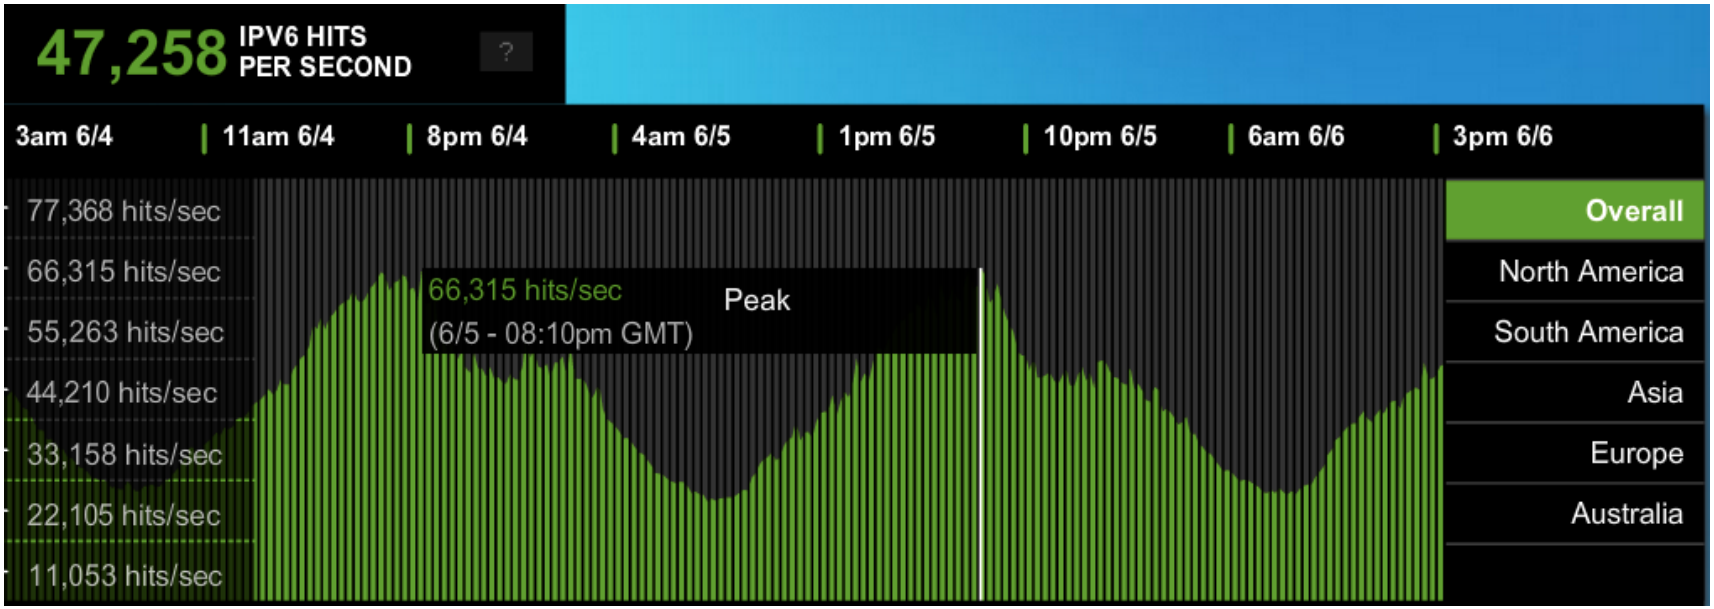

Akamai has a lot of customers on IPv6

- Over 700 US government hostnames
- Over 20 US government agencies
- 1/3 of top-30 World IPv6 Launch Day participants (by Alexa rank), etc.

Those customers who were dual-stacked before World IPv6 Launch show 0.3% to 1.5% of their traffic on IPv6

Observations from World IPv6 Launch Anniversary

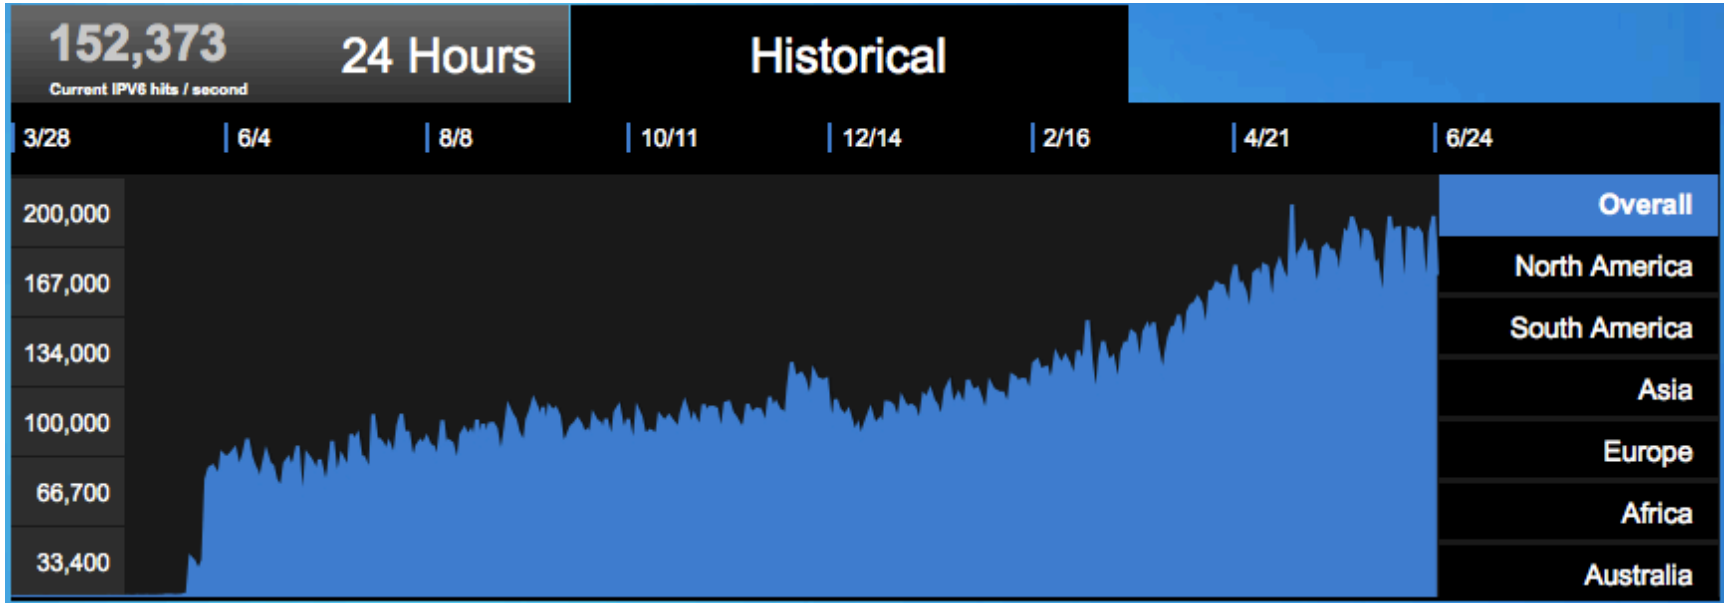

IPv6 traffic continue to growth steadily after World IPv6 Launch

- 2x customers
- 2.5x daily IPv6 requests
- 2.5x dual-stack hostnames (over 1,600 US government hostnames)
- Users upgrade their devices over the next few years
- **We really running out of IPv4!**

Akamai
FASTER FORWARD

Secure the Internet

Open recursors and DNS reflection attack

Recursors

Why?

- Exist to aggregate and cache queries
  - Not every computer run its own recursive resolver.
- ISPs, Large Enterprises run these
- Query through the root servers and DNS tree to resolve domains
- Cache results
- Deliver cached results to clients.

www.cloudflare.com

Recursors

The Problem!

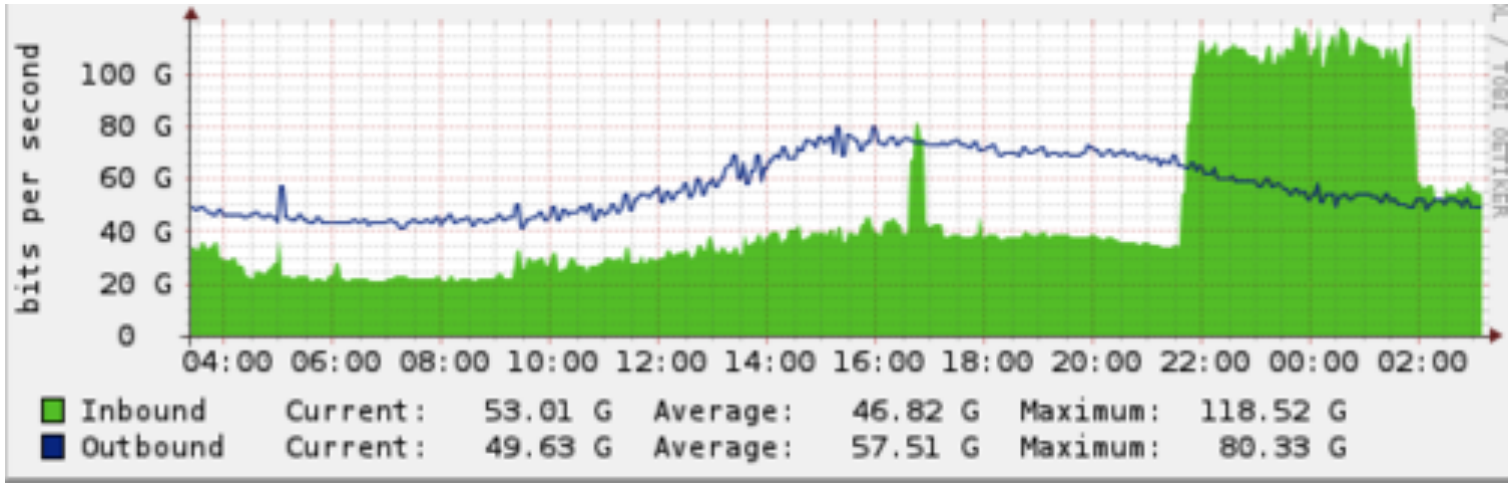

- Example of DNS Based reflection attack exceeding 70Gbit.

www.cloudflare.com

Recursors

Open / Unsecured Recursors ?

- DNS server set up for recursion
  - i.e. non-authoritative
  - Will answer for zones it is not authoritative for
  - Recursive lookups
  - Will answer queries for anyone
- Some Public Services:
  Google, OpenDNS, Level 3, etc.
  - These are “special” set-ups and secured.

www.cloudflare.com

Recursors

Say Again?

- There are hundreds of thousands of DNS Recursors.
- Many of these are not secured.
- Non secured DNS Recursors can and will be abused
- CloudFlare has seen DNS reflection attacks hit 300Gbit traffic globally.

www.cloudflare.com

Reflection Attack

- UDP Query
- Spoofed source
  - Using the address of the person you want to attack
  - DNS Server used to attack the victim (sourced address)

- Amplification used
  - Querying domains like ripe.net or isc.org
  - ~64 byte query (from attacker)
  - ~3233 byte reply (from unsecured DNS Server)
  - 50x amplification!

- Running an unsecured DNS server helps attackers!

www.cloudflare.com

Reflection Attack

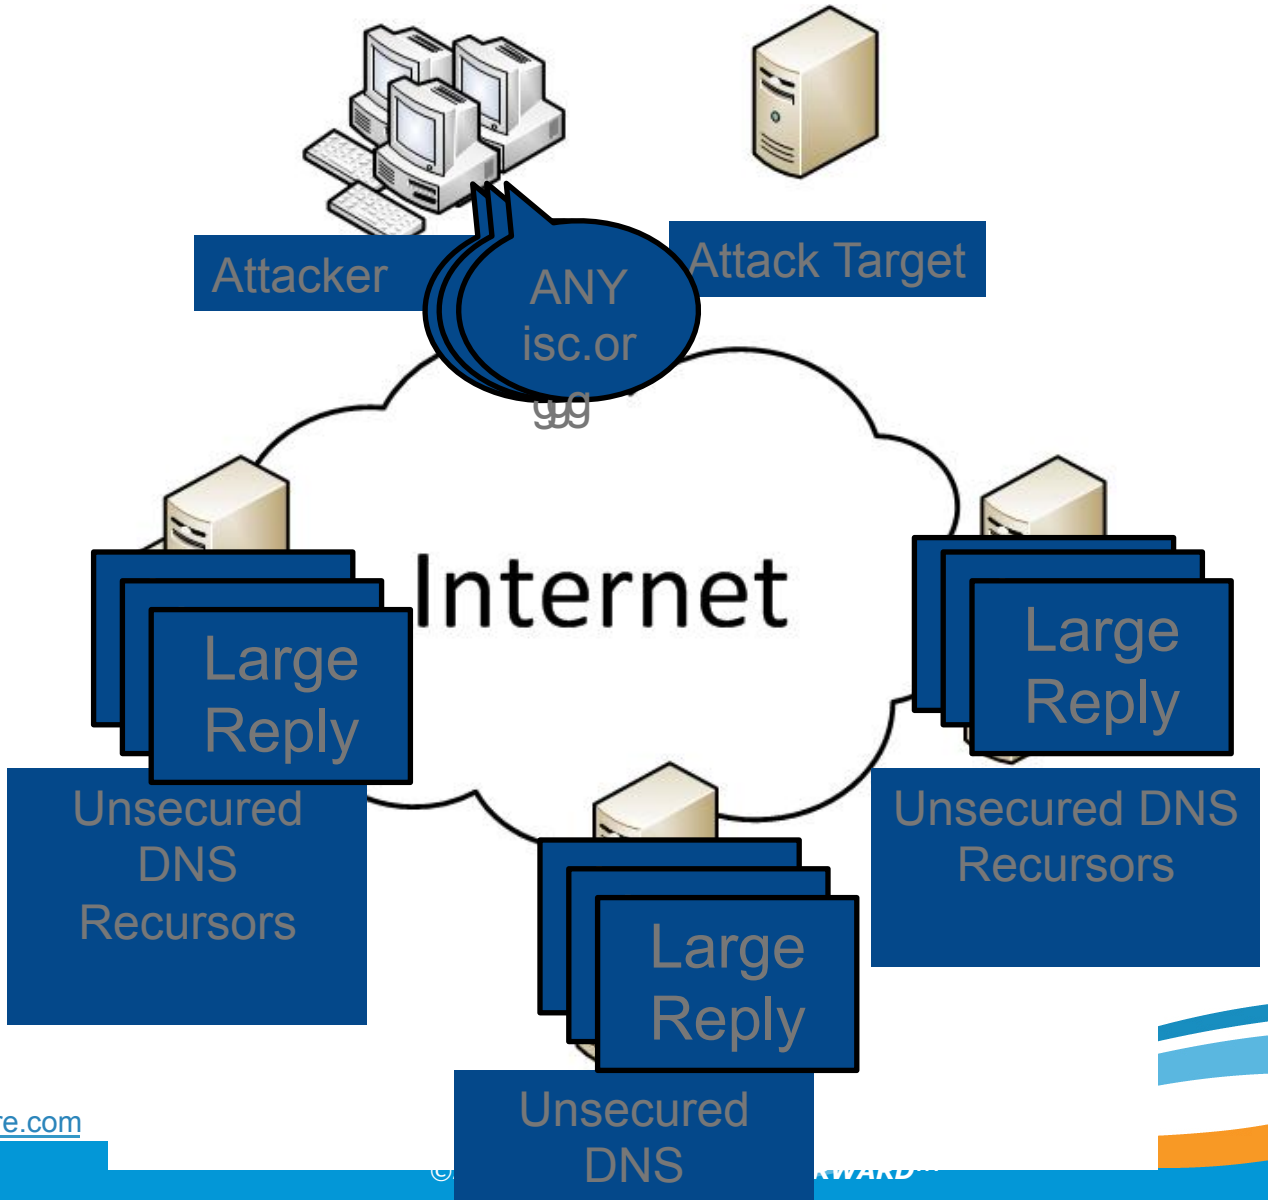

Reflection Attack

- With 50x amplification:
  - 1Gbit uplink from attacker (eg: Dedicated Servers)
  - 50Gbit attack
  - Enough to bring most services offline!

- Prevention is the best remedy.

- In recent attacks, we've seen around 80,000 open/ unsecured DNS Resolvers being used.
- At just 1Mbit each, that's 80Gbit!
  - 1mbit of traffic may not be noticed by most operators.
  - 80Gbit at target is easily noticed!

www.cloudflare.com

Where are the open Recursors?

- Nearly Everywhere!

- CloudFlare has seen DNS Reflection attack traffic from:
  - 27 out of 56 Economies in APNIC Region
  - More attacks from higher populated economies.

www.cloudflare.com

Where are the open Recursors?

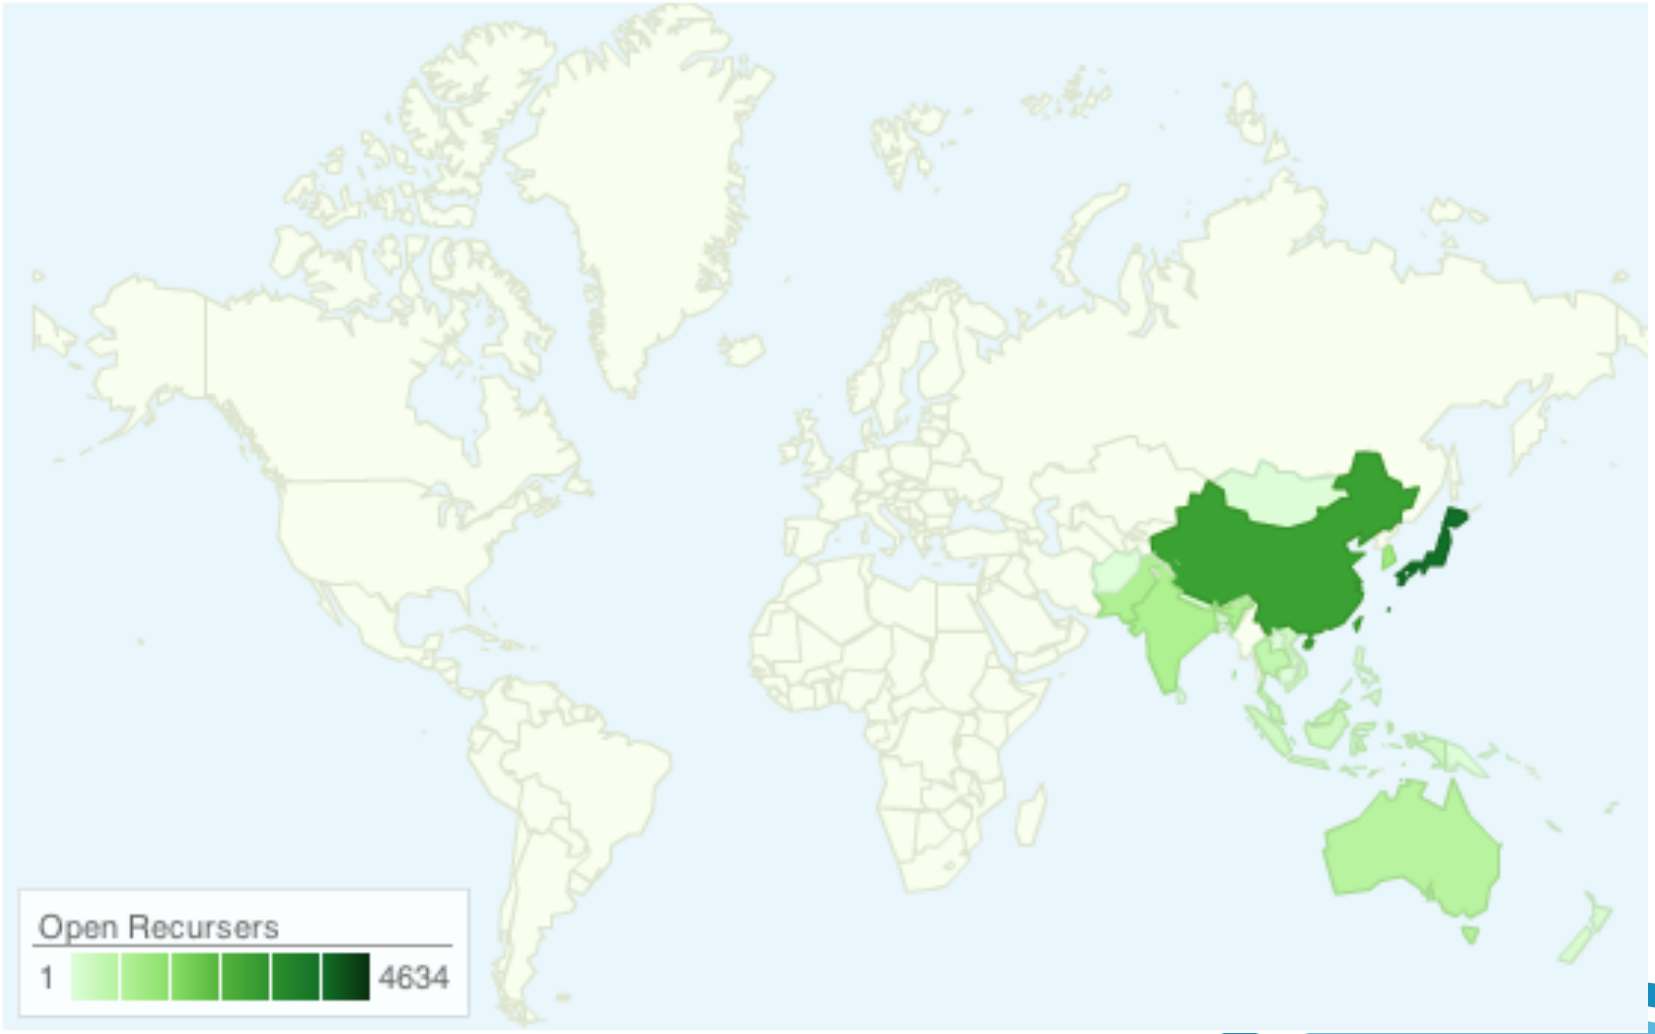

www.cloudflare.com

Where are the open Recursors?

Country	Open Recursors		Country	Open Recursors
Japan	4625		Bangladesh	103
China	3123		New Zealand	98
Taiwan	3074		Cambodia	13
South Korea	1410		Sri Lanka	7
India	1119		Nepal	7
Pakistan	1099		Mongolia	5
Australia	761		Laos	4
Thailand	656		Bhutan	2
Malaysia	529		New Caledonia	2
Hong Kong	435		Fiji	2
Indonesia	349		Maldives	2
Vietnam	342		Papua New Guinea	1
Philippines	151		Afghanistan	1
Singapore	118			

www.cloudflare.com

Where are the open Recursors?

- Where are they running?

Mostly on Servers.

~11,000	Servers profiled.
~7,500	BIND
~1600	unknown / undetermined
~900	Microsoft DNS Server
~500	dnsmasq
~200	ZyWALL DNS (a consumer internet router)

Fixing this?

Preventative Measures!

- BCP-38
  - Source Filtering.
  - You shouldn't be able to spoof addresses.
  - Needs to be done in hosting and ISP environments.
  - If the victim's IP can't be spoofed the attack will stop
  - Will also help stop other attack types
    - (eg: Spoofed Syn Flood).

Fixing this?

Preventative Measures!

- DNS Server Maintenance
  - Secure the servers!
    - Lock down recursion to your own IP addresses
  - Disable recursion
    - If the servers only purpose is authoritative DNS, disable recursion
  - Turn them off!
    - Some Packages (eg, Plesk, cPanel) have included a recursive DNS server on by default.

www.cloudflare.com

Fixing this?

Consumer Internet Routers / Modems

- Update firmware.
  - Some older firmware has security bugs
    - Allows administration from WAN (including DNS, SNMP)
- Does the feature need to be on?
  - Make sure its set up properly

Fixing this?

Information

•BCP-38:
http://tools.ietf.org/html/bcp38

•BIND:
http://www.team-cymru.org/Services/Resolvers/instructions.html

•Microsoft:
http://technet.microsoft.com/en-us/library/cc770432.aspx

•The Open Resolver Project:
http://openresolverproject.org/

Summary

- **Akamai Intelligent Platform**
  - Highly distributed edge servers
  - Akamai mapping is different than BGP routing
- **IPv6 traffic is still small today, but catching up**
  - Dual-stack approach
  - IPv4 is really running out!
- **Secure the Internet**
  - Open recursors and DNS reflection attacks
  - BCP-38 and DNS servers maintenance

Questions?

Christian Kaufmann <ck@akamai.com>

More information:

Peering: http://as20940.peeringdb.com

IPv6: http://www.akamai.com/ipv6

Acknowledgement:

Tomas Paseka <tom@cloudflare.com>