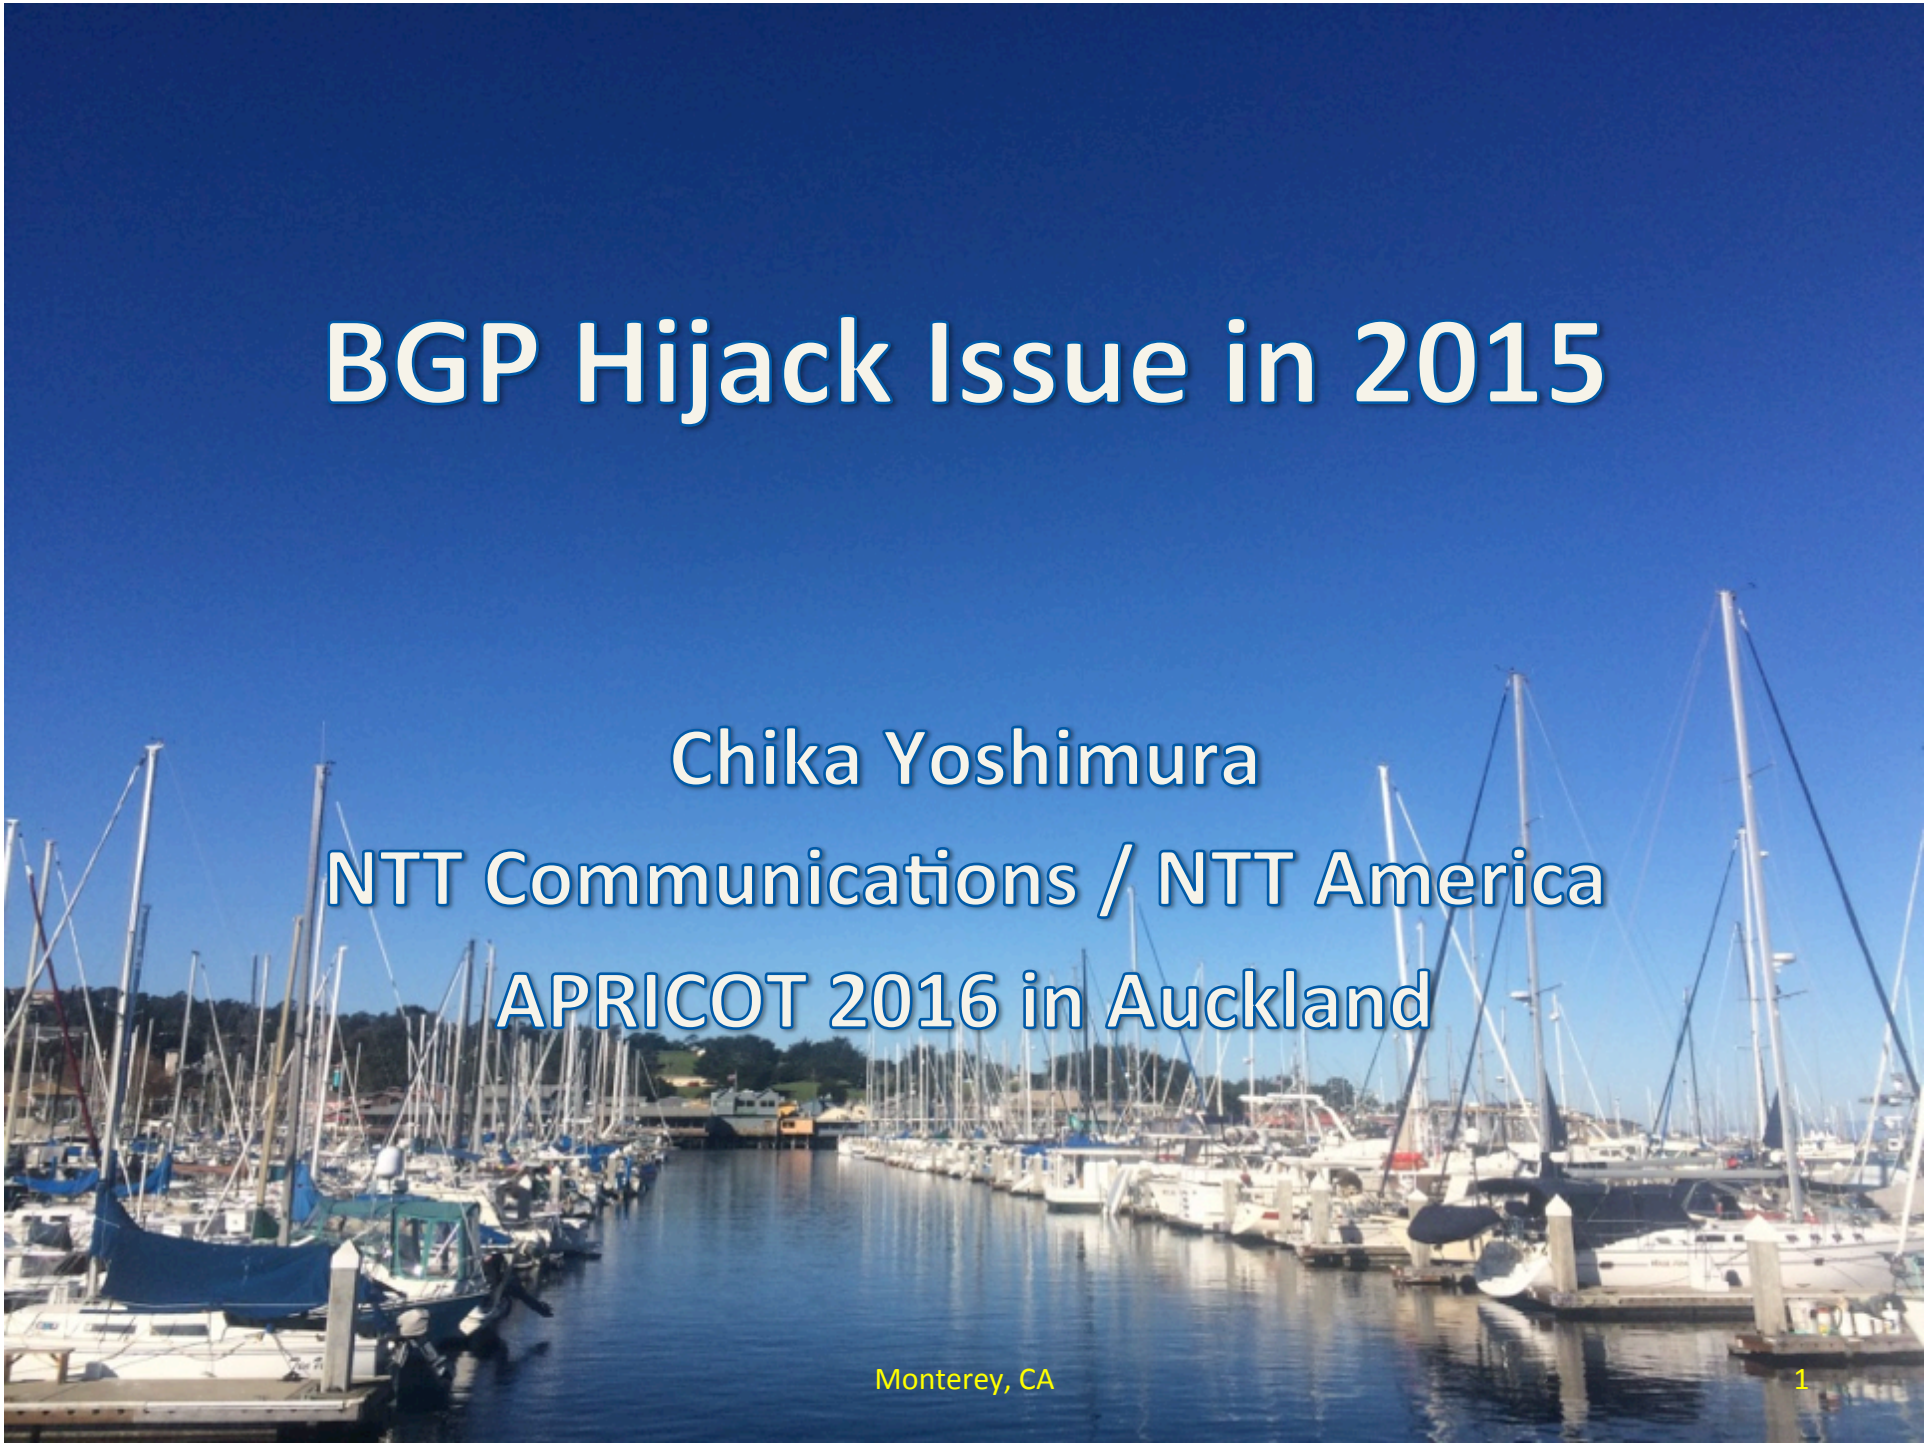

BGP Hijack Issue in 2015

Chika Yoshimura
NTT Communications / NTT America

APRICOT 2016 in Auckland

Who am I?

- Chika Yoshimura
- NTT Communications
  - Network Engineer in AS2914
  - Based in San Jose
  - 3 yrs in NTT-GIN (AS2914)
  - 10 yrs in NTT-OCN (AS4713)


November 06 2015…

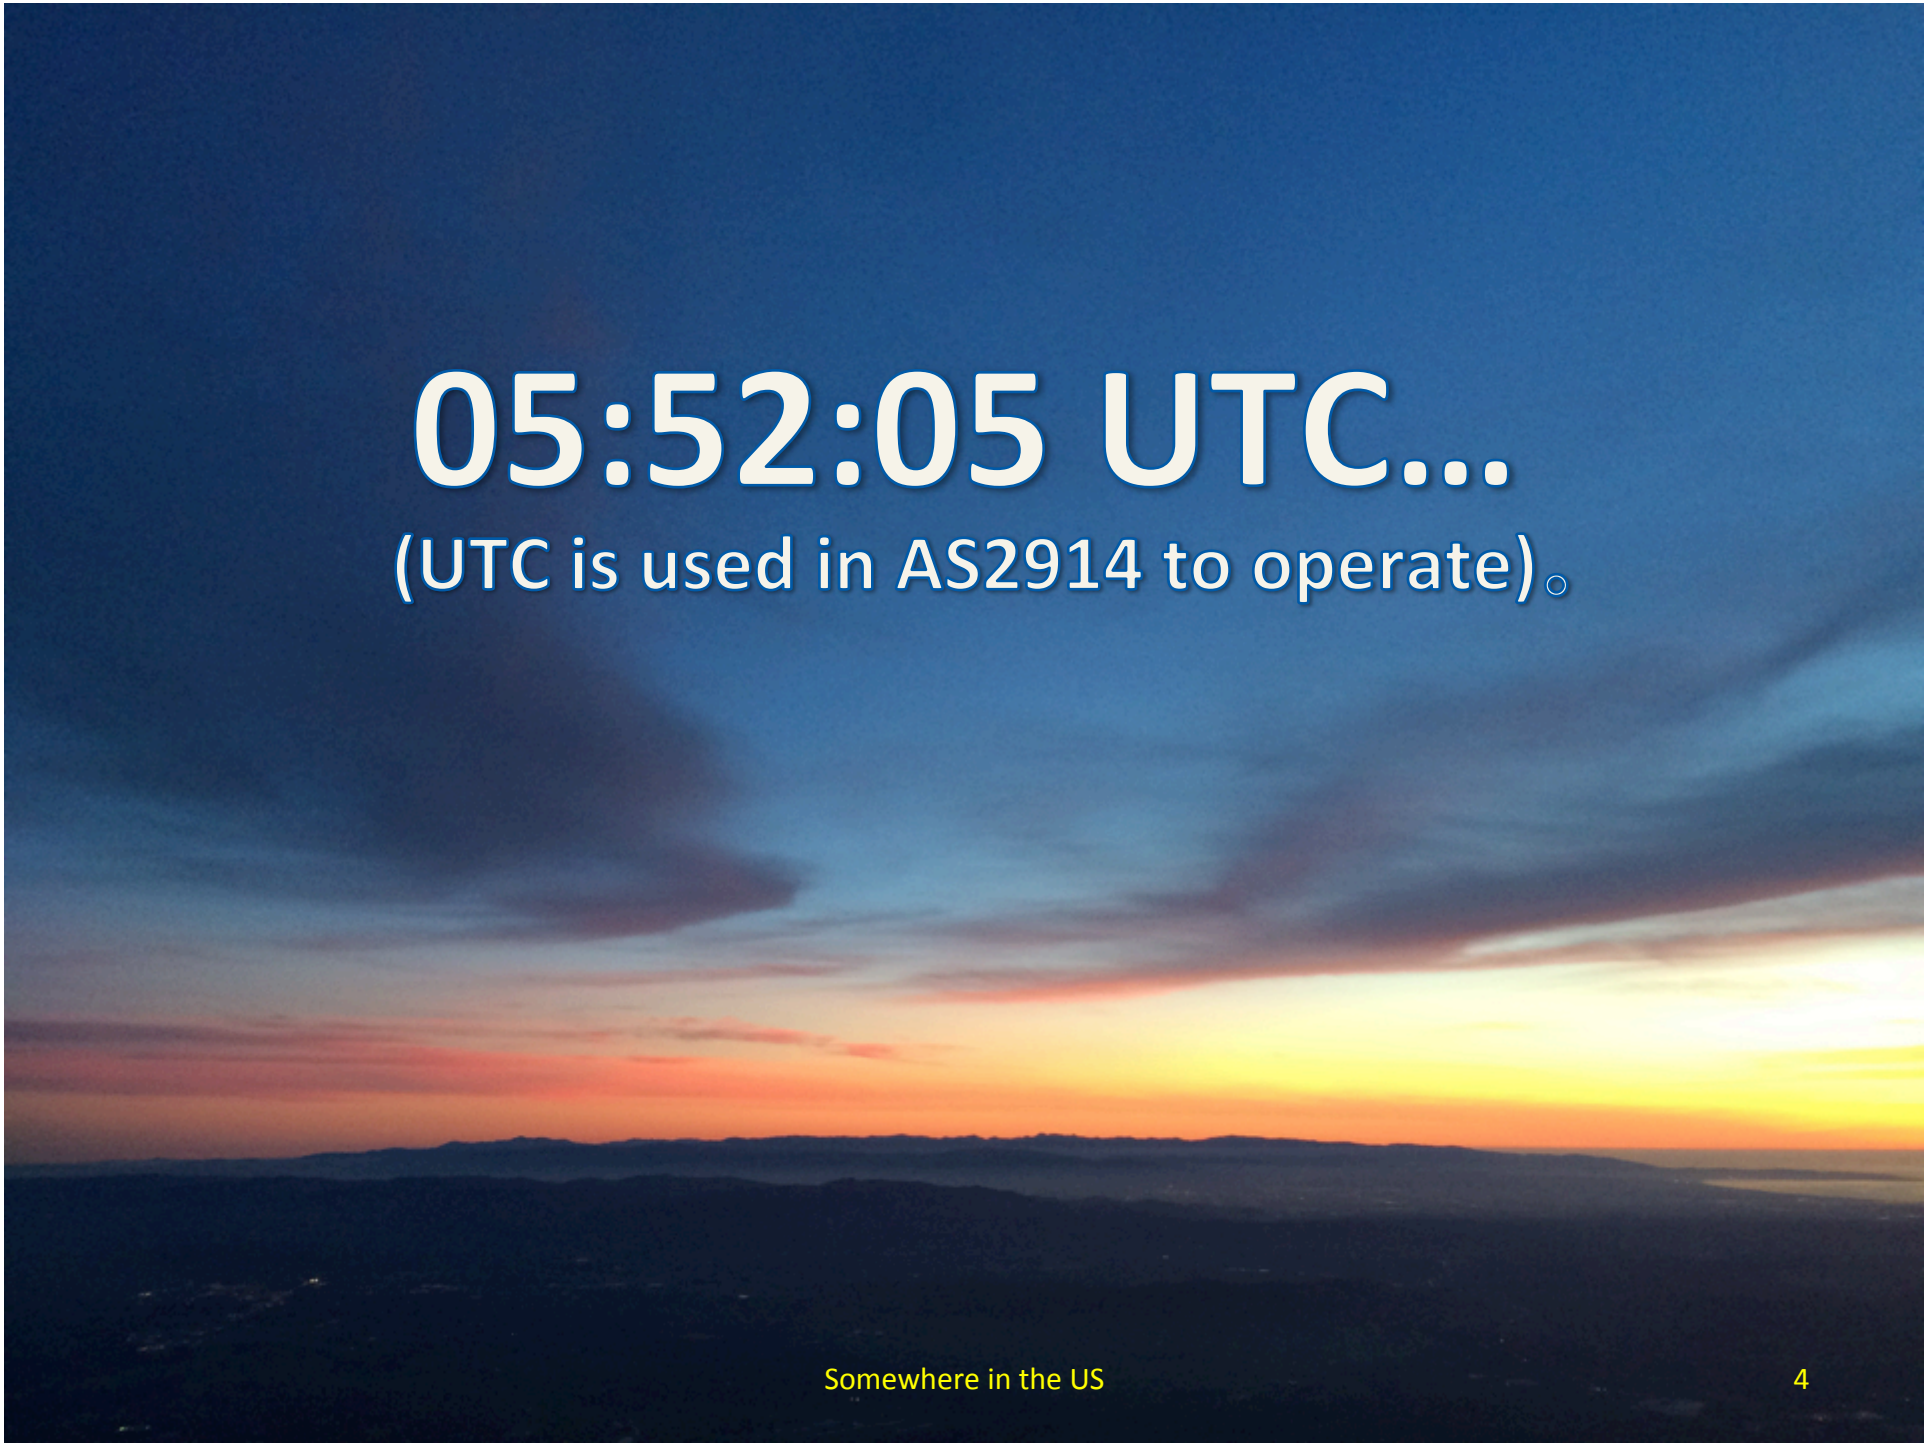
05:52:05 UTC…
(UTC is used in AS2914 to operate)。

A Huge BGP Hijack Issue Occurred!

BGPMON Now part of OpenDNS

HOME BLOG ABOUT US PRODUCTS AND SERVICES NEWS AND PRESS CLIEN

Large scale BGP hijack out of India

Posted by Andree Toonk - November 6, 2015 - Hijack - 1 Comment

BGP hijacks happen every day, some of them affect more networks than others and every now and then there's a major incident that affects thousands of networks. Our monitoring systems keep an eye out for our users and if you would like to have a general idea of what's going on in the world of BGP incidents, keep an eye on BGPstream.com. Earlier today we detected one of those major incidents that affected thousands of networks.

Starting at 05:52 UTC, AS9498 (BHARTI Airtel Ltd.) started to claim ownership for thousands of prefixes by originating them in BGP. This affected prefixes for over two thousand unique organizations (Autonomous systems).

Our systems detected origin AS changes (hijacks) for 16,123 prefixes. The scope and impact was different per prefix but to give you an idea, about 7,600 of these announcements were seen by five or more of our peers (unique peers ASns) and 6,000 of these were seen by more than 10 of our peers.

One of the reasons this was so widespread is because large networks such as AS174 (Cogent Communications) and AS52320 (GlobeNet Cabos Submarinos VZLA) accepted and propagated these prefixes to their peers and customers.

The BGPlay visualization below shows an example hijack for a prefix normally announced by AS39891 Saudi Telecom Company JSC.

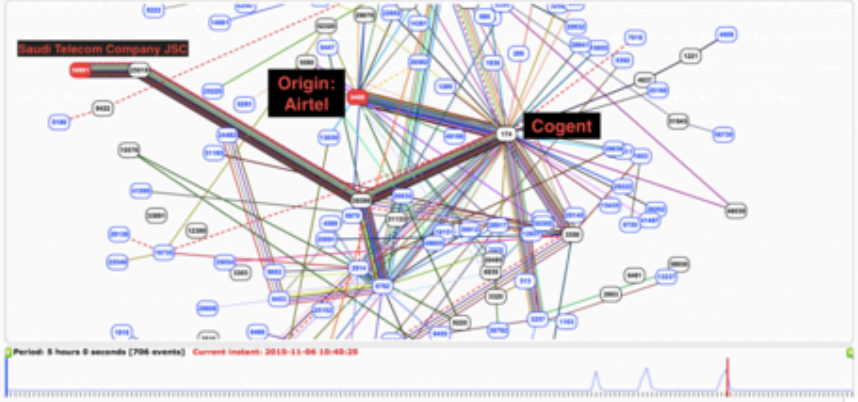

http://www.bgpmon.net/large-scale-bgp-hijack-out-of-india/

BGP Hijack

- To advertise prefix(es) from third AS which is not related to Orign AS (BGP Origin AS is disguised)
- For instance, 2.16.65.0/24
  - AS2914's prefix
  - If an AS except AS2914 advertises as its own prefix ≒ BGP hijack
- Other ASes which receive the hijacked prefix might believe it's legit
- Then traffic toward the hijacked prefix will go to the disguised Origin
- Not so rare
  - 2015/08/01-2015/12/31 : more than 850 hijack issues occur (per BGPStream)

Typical Root Causes

- Malicious Root Causes
- Non-malicious Root Causes
  - Mis-operations
  - (ex) leaking IGP prefixes to EGP
  - (ex) leaking testing prefixes to EGP
  - BGP filtering mistakes are most likely
- (FYI) Multiple origin
  - To advertise a prefix from more than 2 Origin Ases
  - Not a BGP hijack

BGP Hijack Issue on Nov 6

- Per BGPMon
  - 2015/11/06 05:52 – 14:40 UTC
  - AS9498 (Bharti Airtel) advertised:
    - 16123 prefixes (more than 2000ASes)
  - Hijacked ASes:
    - AS3257/GTT, AS4755/Tata Communications etc
    - AS2914/NTT Communications (Yes it's us!)

http://www.bgpmon.net/large-scale-bgp-hijack-out-of-india/

Root Cause of the Hijack Issue

- Still Unknown
  - AS2914 contacted AS9498
    - No response about a root cause
  - BGPMon doesn't have info of root cause
  - No info on the NANOG ML
- From what I can guess from the actual hijacked prefixes…
  - They might have missed BGP prefix filters?

Actual situation in AS2914 on Nov 6

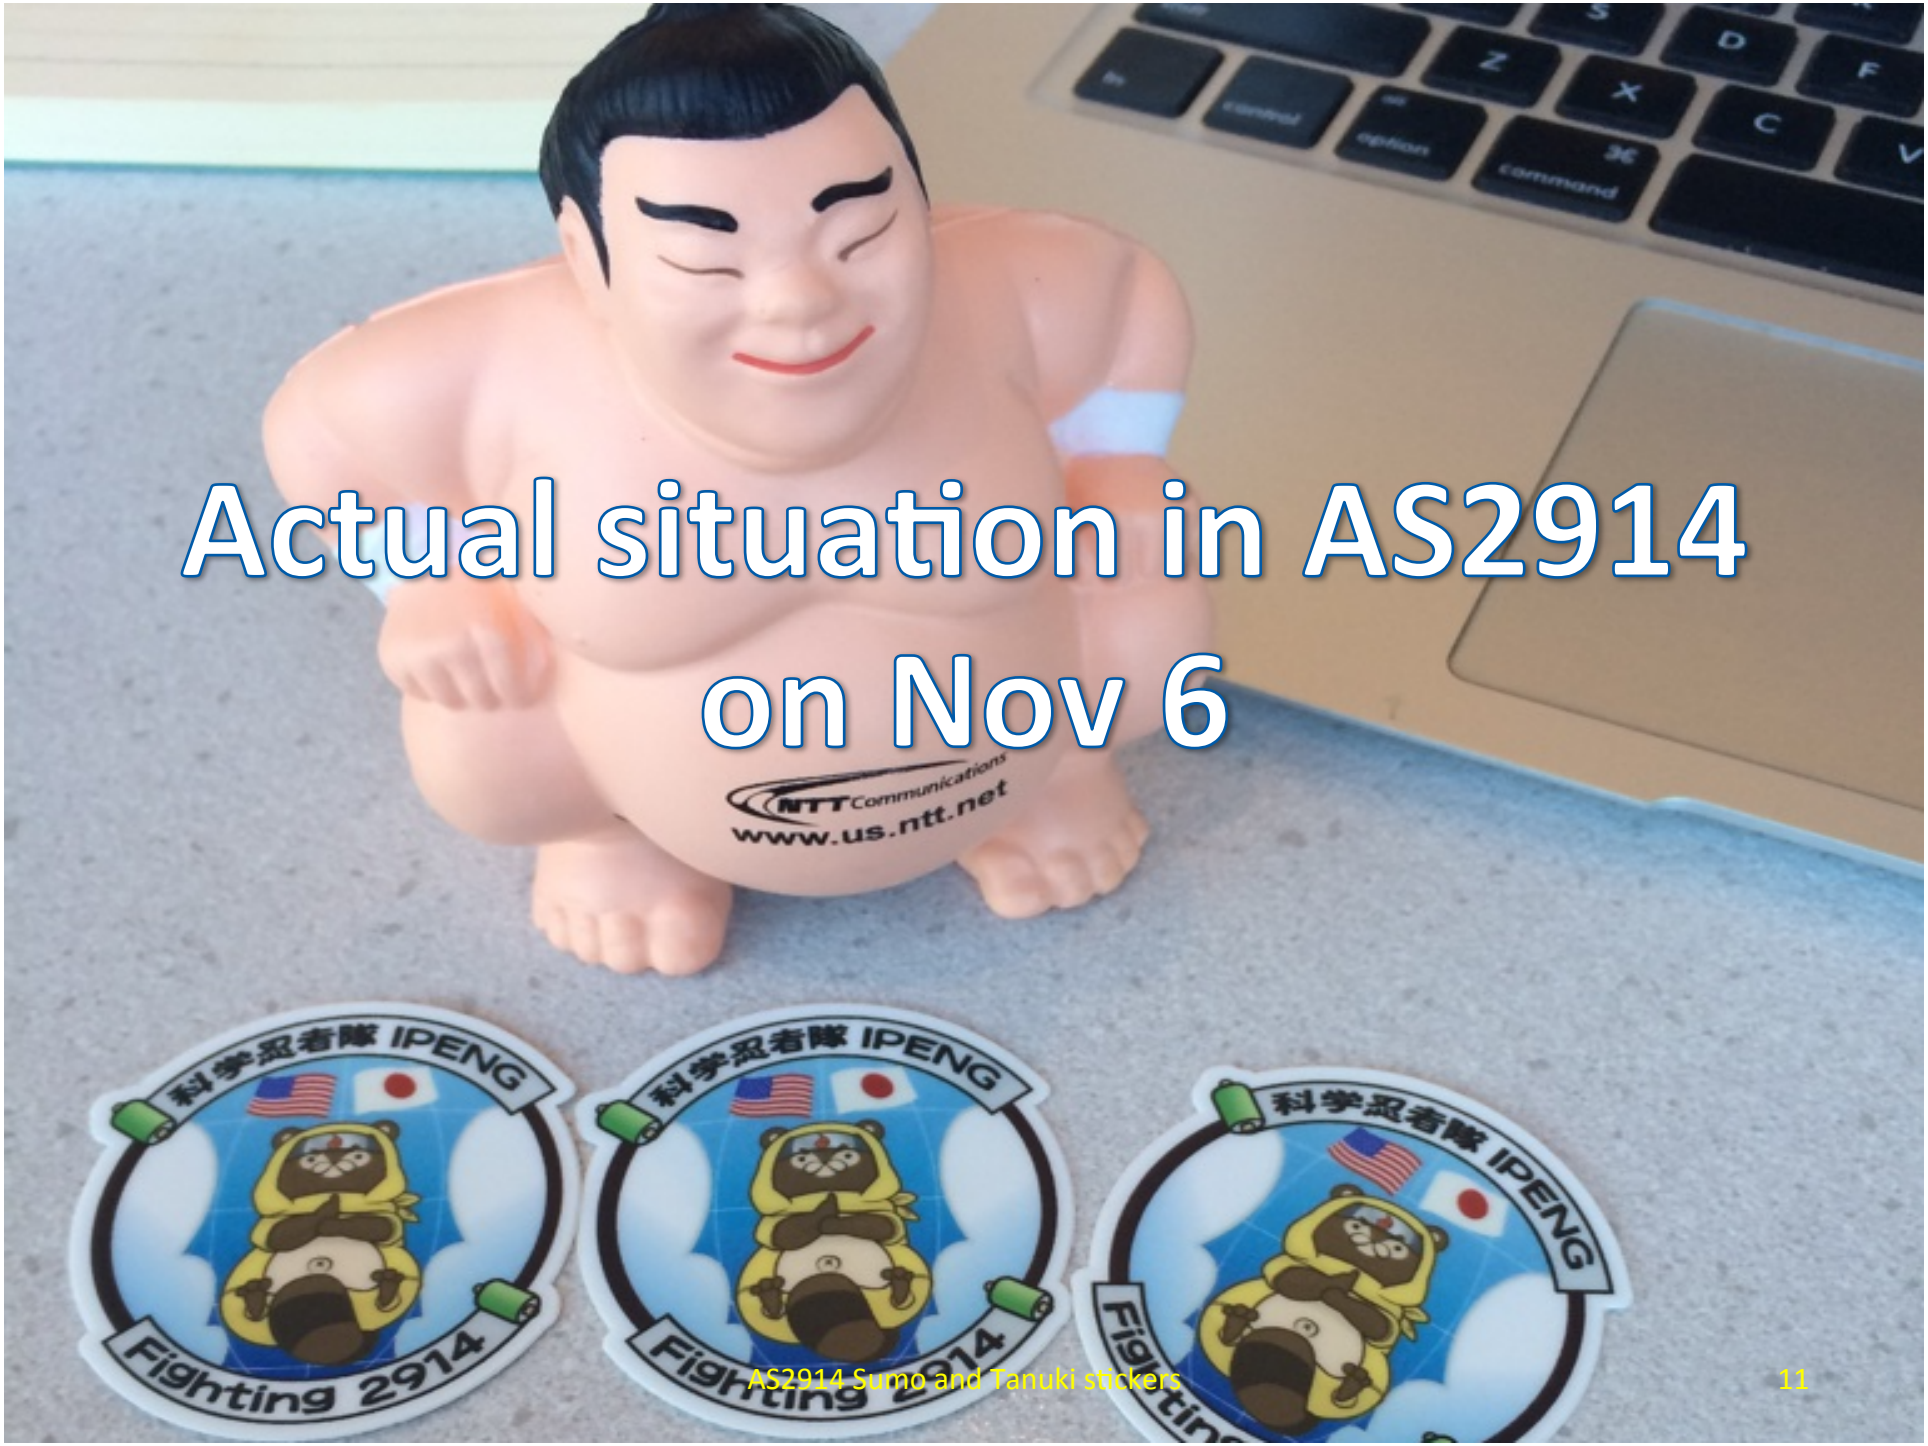

AS2914 Sumo and Tanuki stickers

AS2914 Operational Timestamp

- Nov 06 found our prefixes were hijacked
- Nov 06 AS2914 NOC sent an e-mail to AS9498
- Nov 07 AS2914 NOC sent another e-mail to AS9498
- Nov 07 AS9498 responded
  - No info about the root cause
- Nov 07 AS2914 NOC sent one more e-mail to AS9498 asking a root cause.
  - No response
- Started analyzing affected prefixes with BGPMon

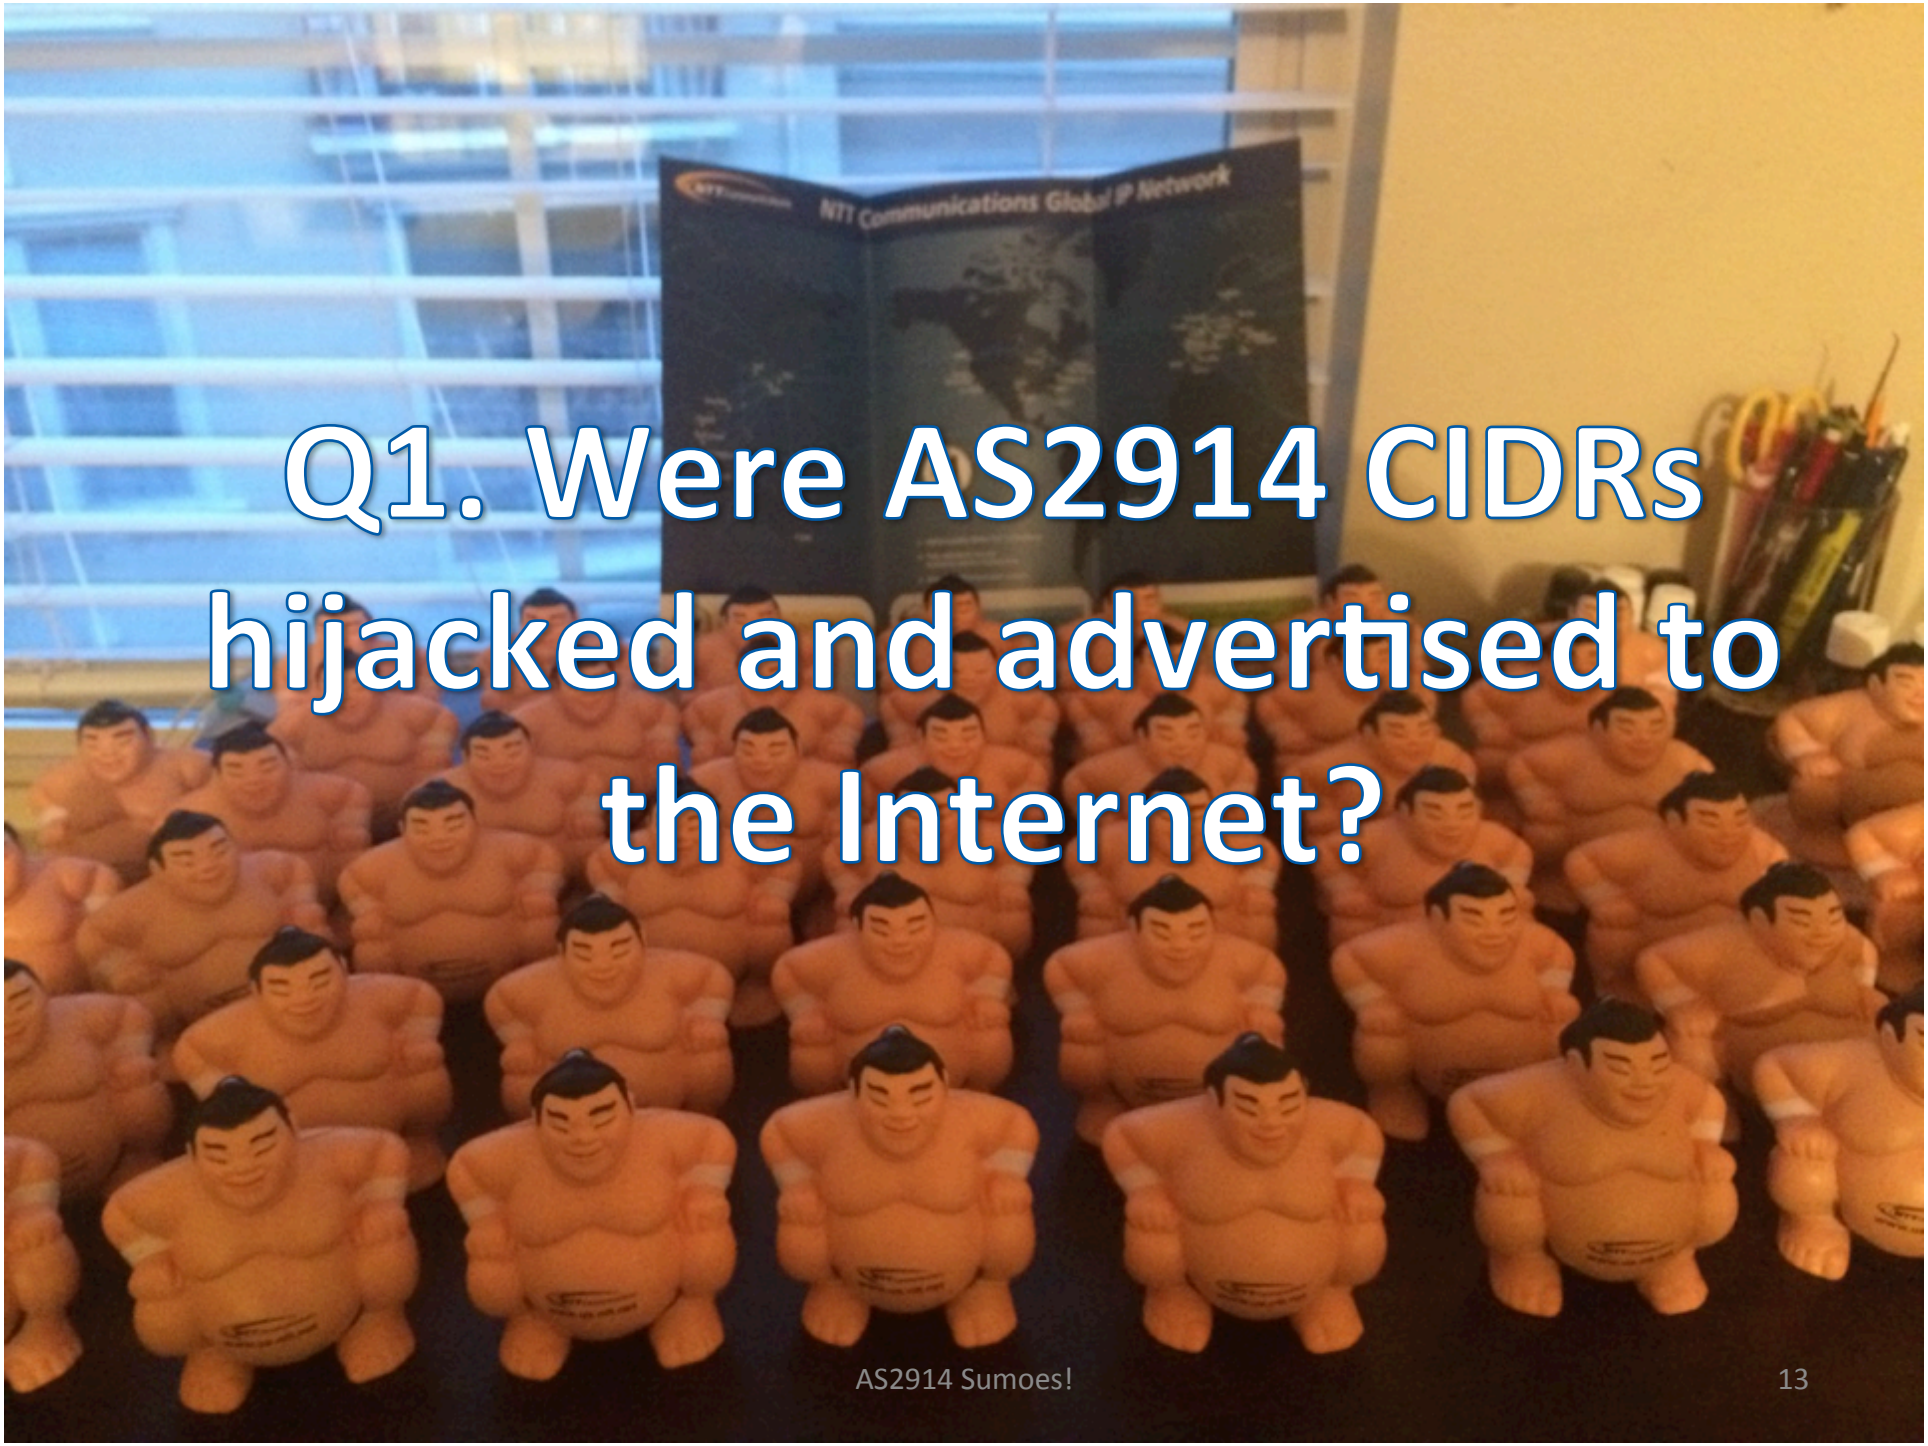

Q1. Were AS2914 CIDRs hijacked and advertised to the Internet?

Yes, our prefixes were hijacked

- 300 prefixes of AS2914 were hijacked and advertised to the Internet
- AS2914 generally doesn't allocate our CIDR to customers
  - That's why there was no significant impact to our services

(A part of) Hijacked Prefixes – Per BGPMon

announced_prefix	base_as	src_AS	start_time	Peer_count
2.16.65.0/24	2914	9498	2015-11-06 05:52:14	68
2.16.110.0/23	2914	9498	2015-11-06 05:52:20	49
2.17.196.0/22	2914	9498	2015-11-06 05:52:15	47
5.158.208.0/21	2914	9498	2015-11-06 05:52:19	37
2.21.16.0/20	2914	9498	2015-11-06 05:52:15	33
23.55.208.0/20	2914	9498	2015-11-06 05:52:26	10
23.67.64.0/22	2914	9498	2015-11-06 05:52:26	10
23.55.80.0/20	2914	9498	2015-11-06 05:52:26	10
23.38.110.0/23	2914	9498	2015-11-06 05:52:26	10
23.11.192.0/22	2914	9498	2015-11-06 05:52:23	10
23.4.32.0/20	2914	9498	2015-11-06 05:52:20	10
23.11.196.0/22	2914	9498	2015-11-06 05:52:23	10

No significant impact? Really?

- Whether there's an impact due to a BGP hijack issue depends on what services we do with the prefixes
- IP Whole Sales (like us) generally don't use our own prefixes
  - Because customers already have their own AS and prefixes
- Consumer Services use our own prefixes for customers
  - so there'll be a large impact when the prefixes are hijacked

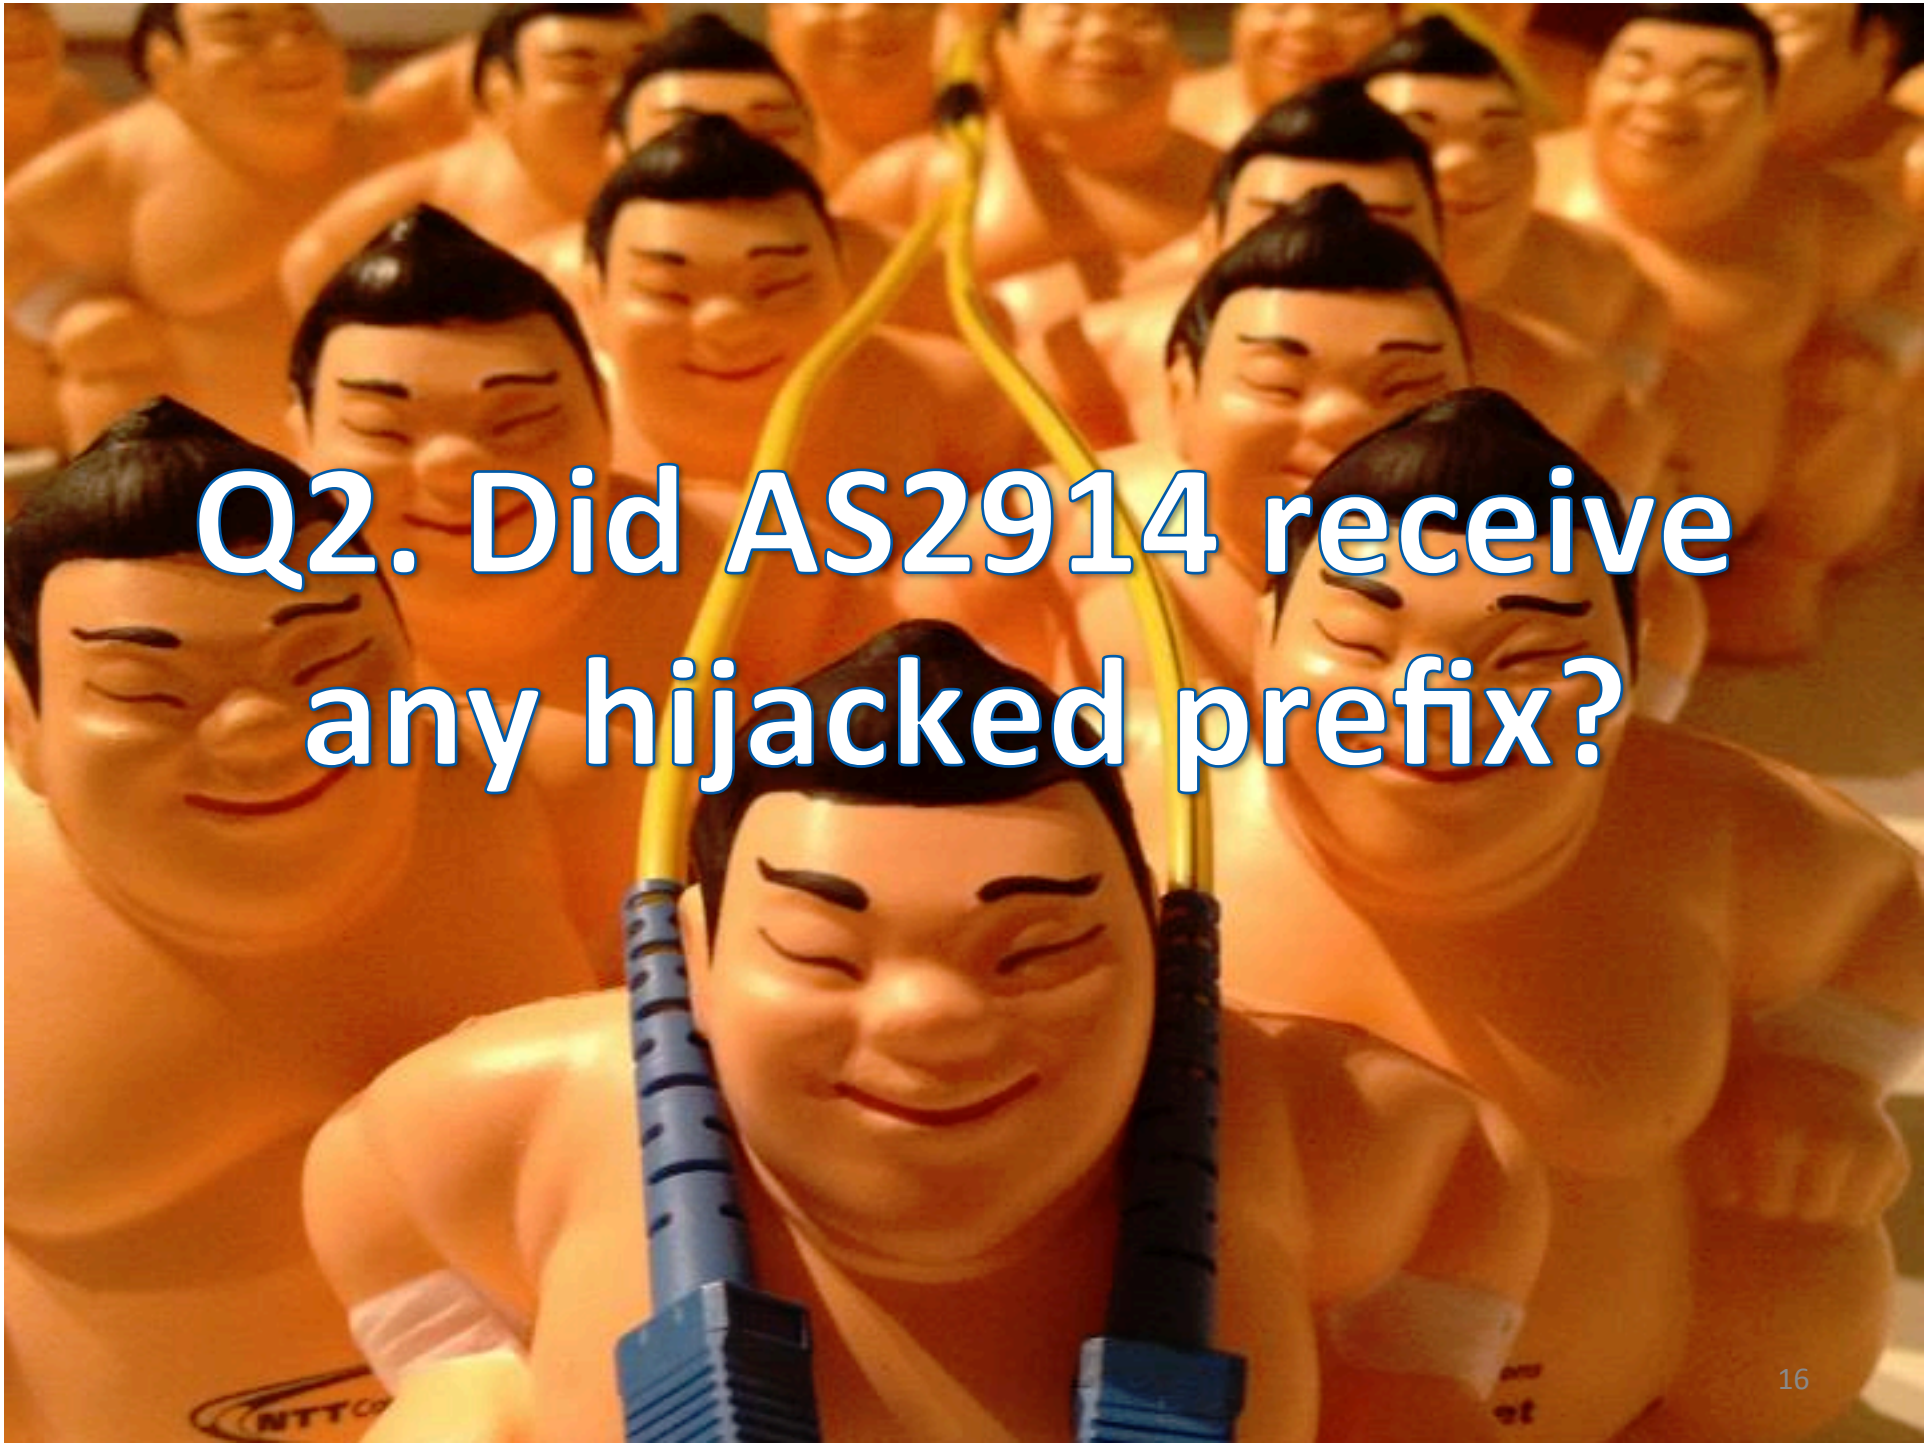
Q2. Did AS2914 receive
any hijacked prefix?

Yes, we received some of them

- Duration: 2015/11/06 05:52:05 - 14:37:41 UTC
- 4513 prefixes received (IPv4: 4512, IPv6: 1)
- Mainly received from Peer ASes
  - We don't have any upstream AS
  - There's a strict BGP prefix filter for downstream ASes
  - There's a rough BGP filter for Peer Ases
- Didn't receive our own prefixes (AS2914's prefixes)

BGP Updates of Hijacked Prefixes

- Roughly 3 peaks during the duration
- Started hijacking 1.x.x.0 first, then 2.x.x.0, then 5.x.x.0….
- Any CPU issue due to the many BGP updates? -> We didn't face this time

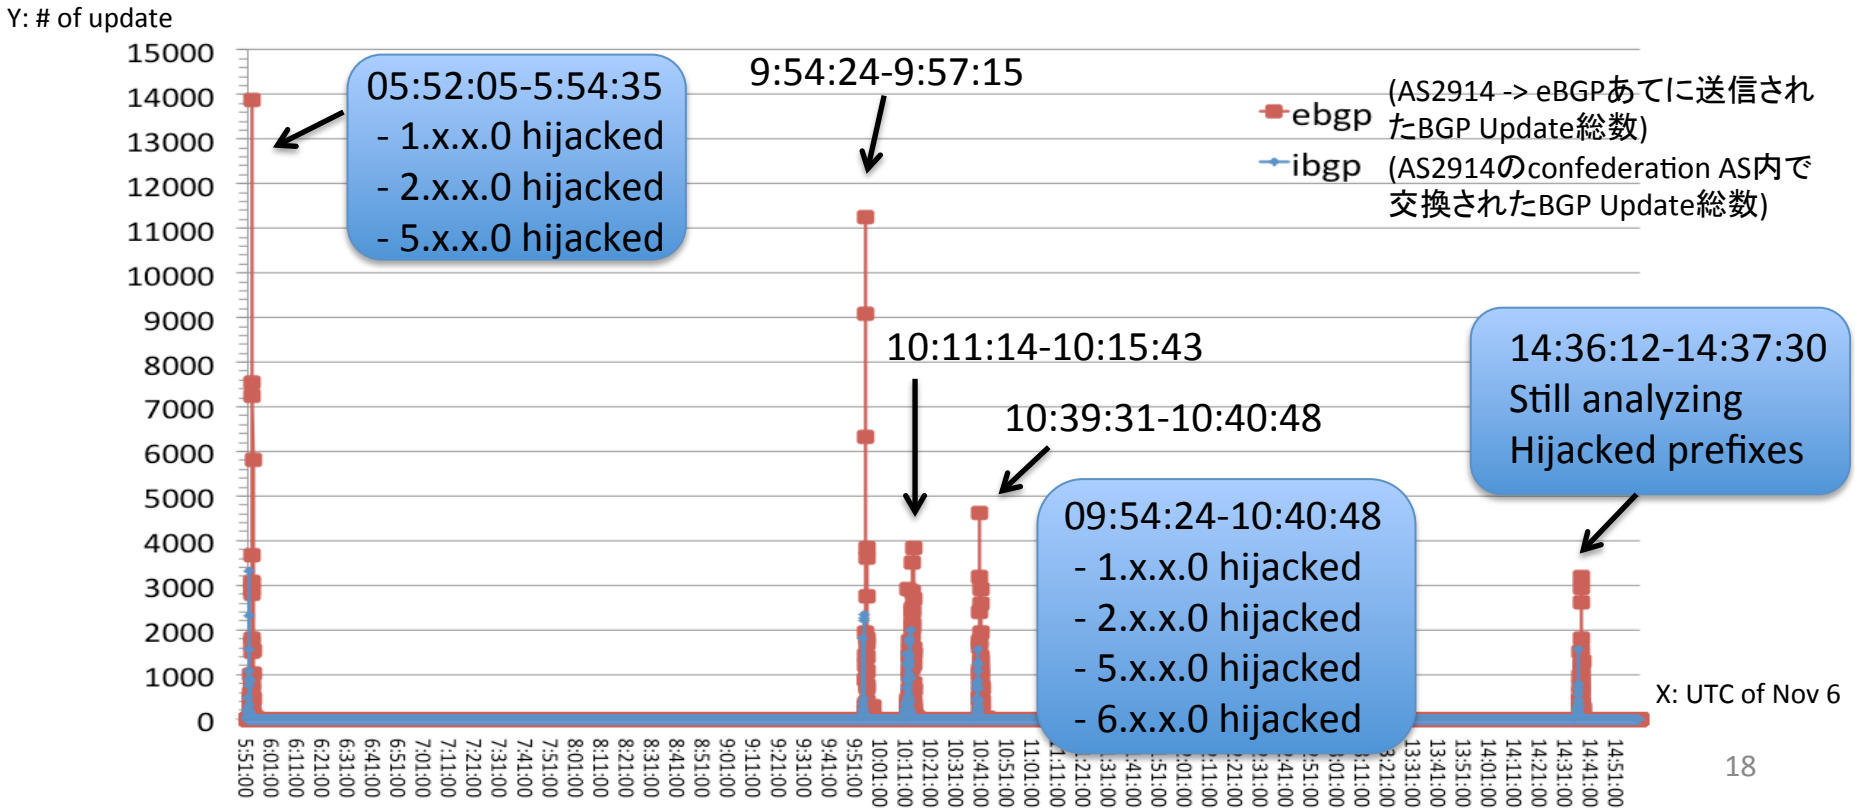

Hijacked Prefix Ranges

- Simple prefixes
  - 1.0.x.0/24
  - 2.0.x.0/24
- same subnet mask as IRR
  - Still analyzing
- Probably BGP filter mistakes?
- Probably route leaking?
  - Received from EGP
    - → distribute to IGP
    - → distribute to EGP again
- This data is just what we saw inside AS2914 so there were more hijacked prefixes

range	# of hijacked prefix
1.x	1331
2.x	175
5.x	1771
6.x	34
8.x	858
12.x	229
14.x	8
23.x	1
24.x	2
27.x	96
61.x	1
64.x	1
125.x	1
177.x	4
2c0f:fe90::	1
total	4513

What ASes the Hijacked Prefixes Belong to?

- Most of them aren't AS2914's customer
  - otherwise customer but not advertised prefixes to AS2914
- Their prefixes need to be received from other ASes
  - Mainly from peers

ASN	Name	Country
39891	Saudi Telecom Company	SA
24378	Total Access Communication	TH
12586	GHOSTnet	DE
18403	The Corporation for Financing & Promoting Technology	VN
35819	Etihad Etisalat Company	SA
4788	TM Net	MY
38266	Vodafone Essar	IN
23089	Hotwire Communications	US
45083	Beijing CheeryZone Scitech	CN
21299	2DAY Telecom	KZ

Where did the hijacked prefixes come from?

- We received the hijacked prefixes from our peer ASes (mainly Tier1 Ases)

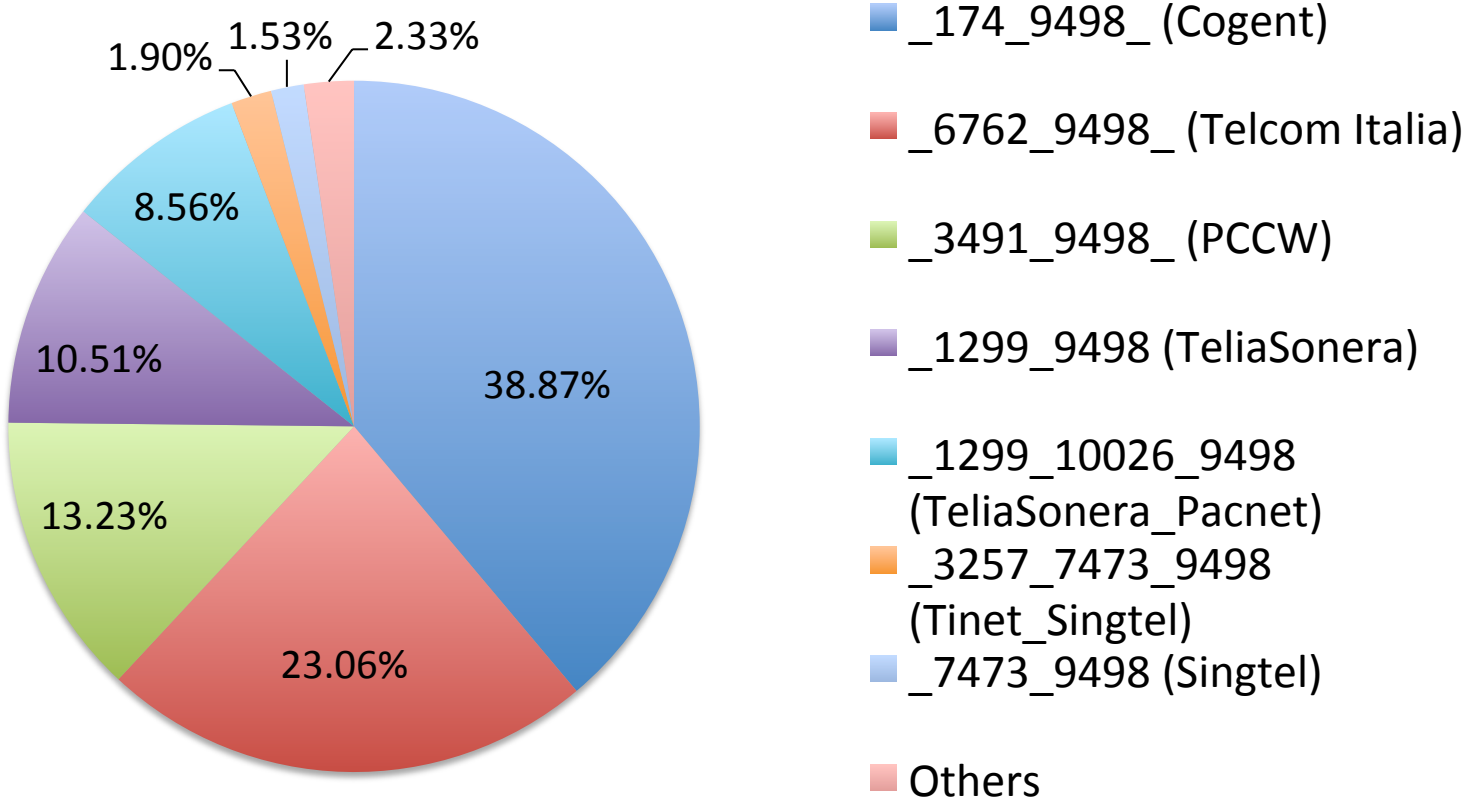

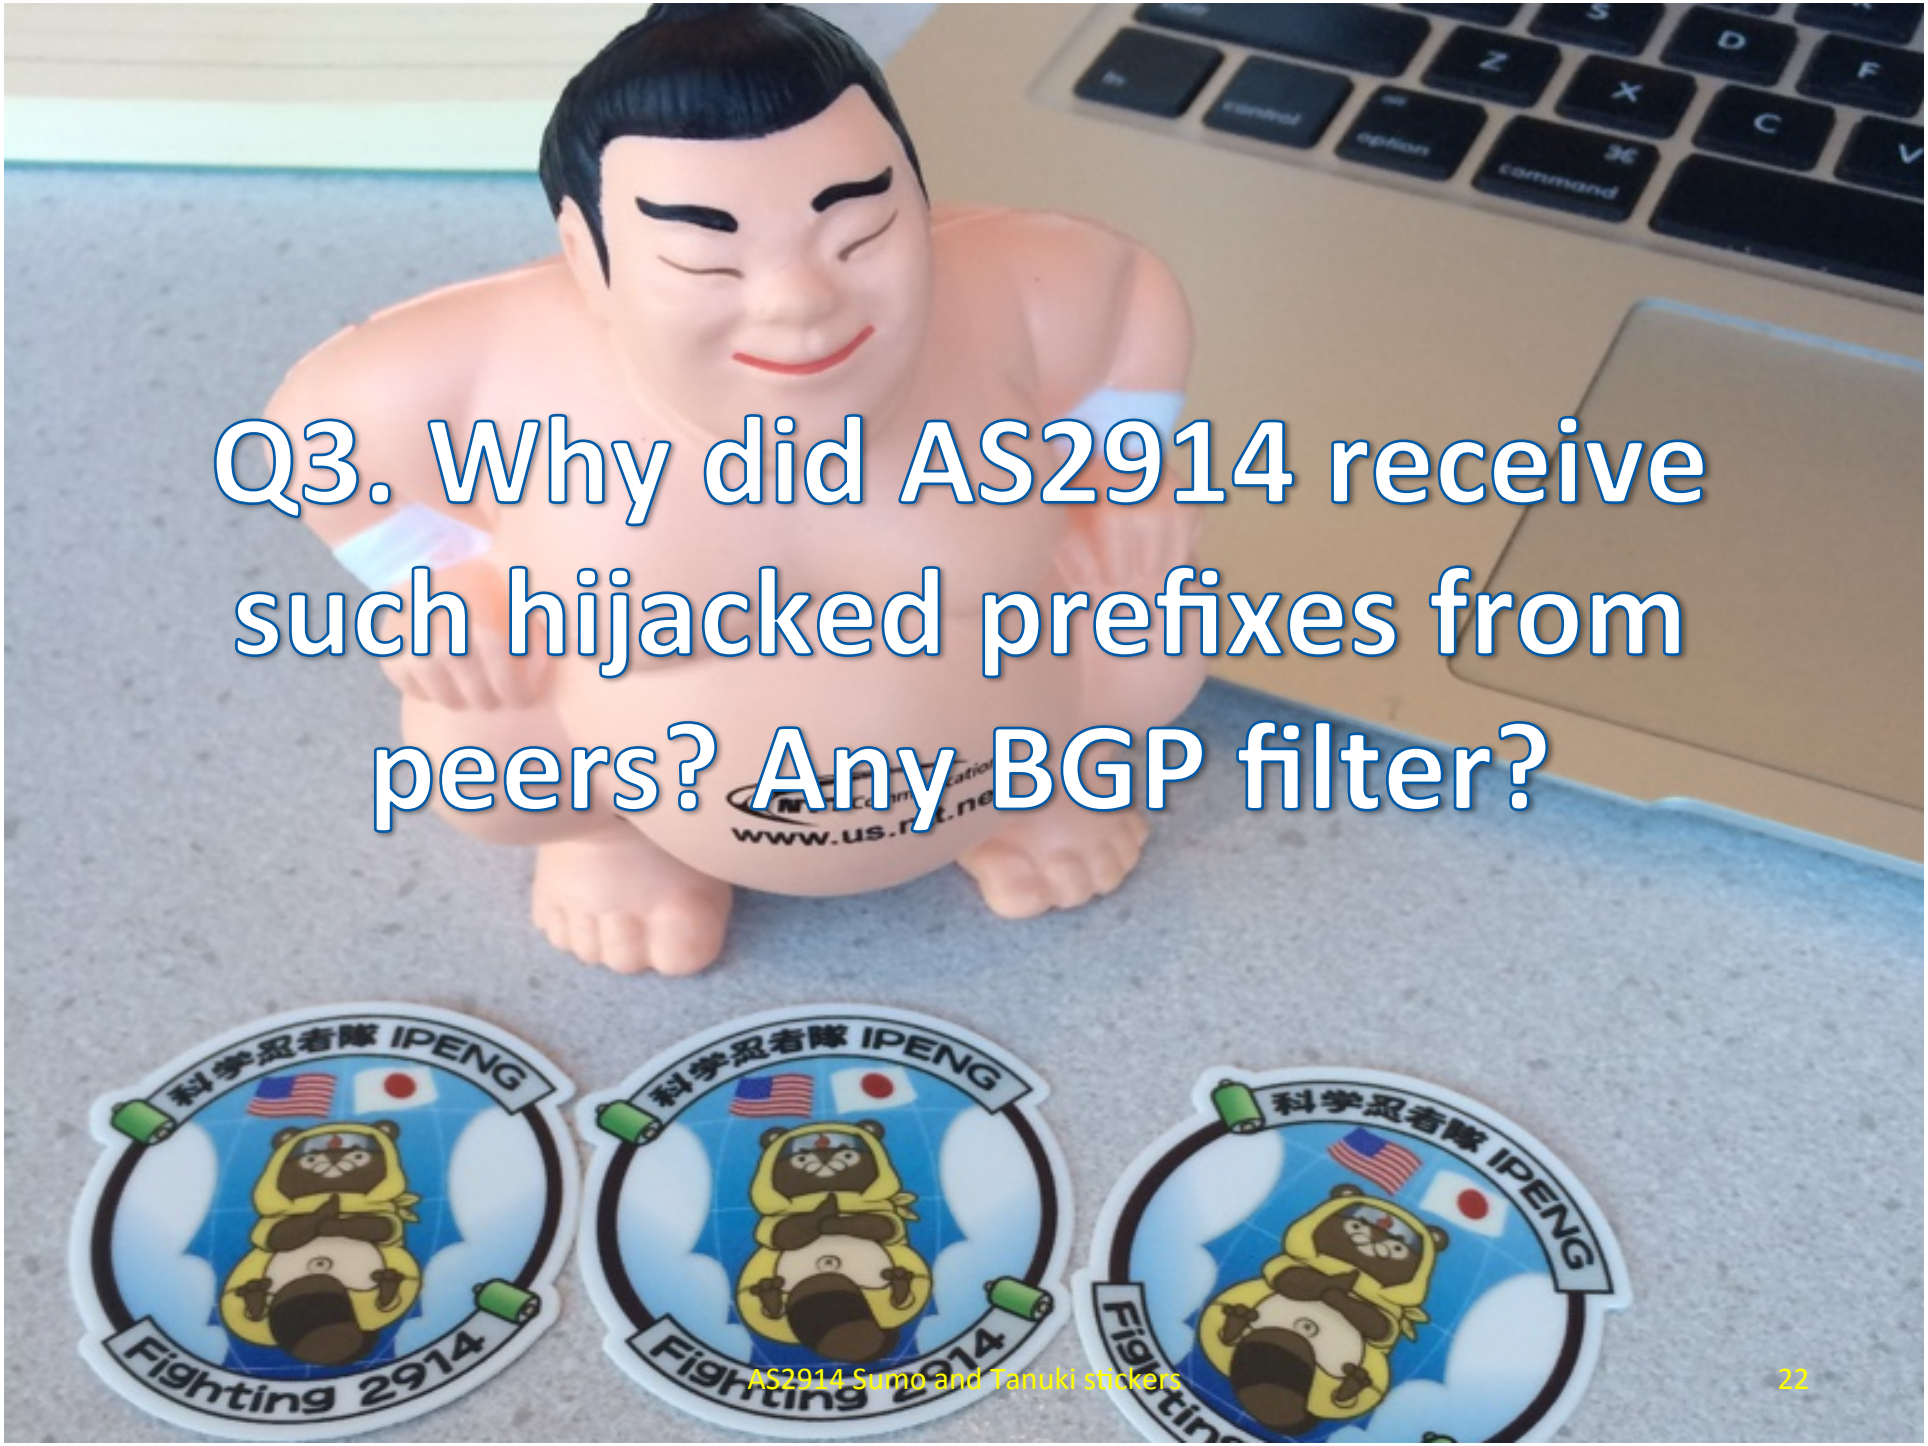

AS2914 Sumo and Tanuki stickers

What import BGP Filter do we apply?

- To Peer
  - Bogon etc
  - uRPF
  - Max prefix filter
  - (a kind of) AS path filter
  - It's not realistic to apply a strict BGP filter to Peers (Tier1 ASes) because they advertise almost of the full BGP table prefixes

What import BGP Filter do we apply?(cont)

- (FYI) To Customers
  - uRPF
  - Max prefix filter
  - Prefix filter (based on IRR)

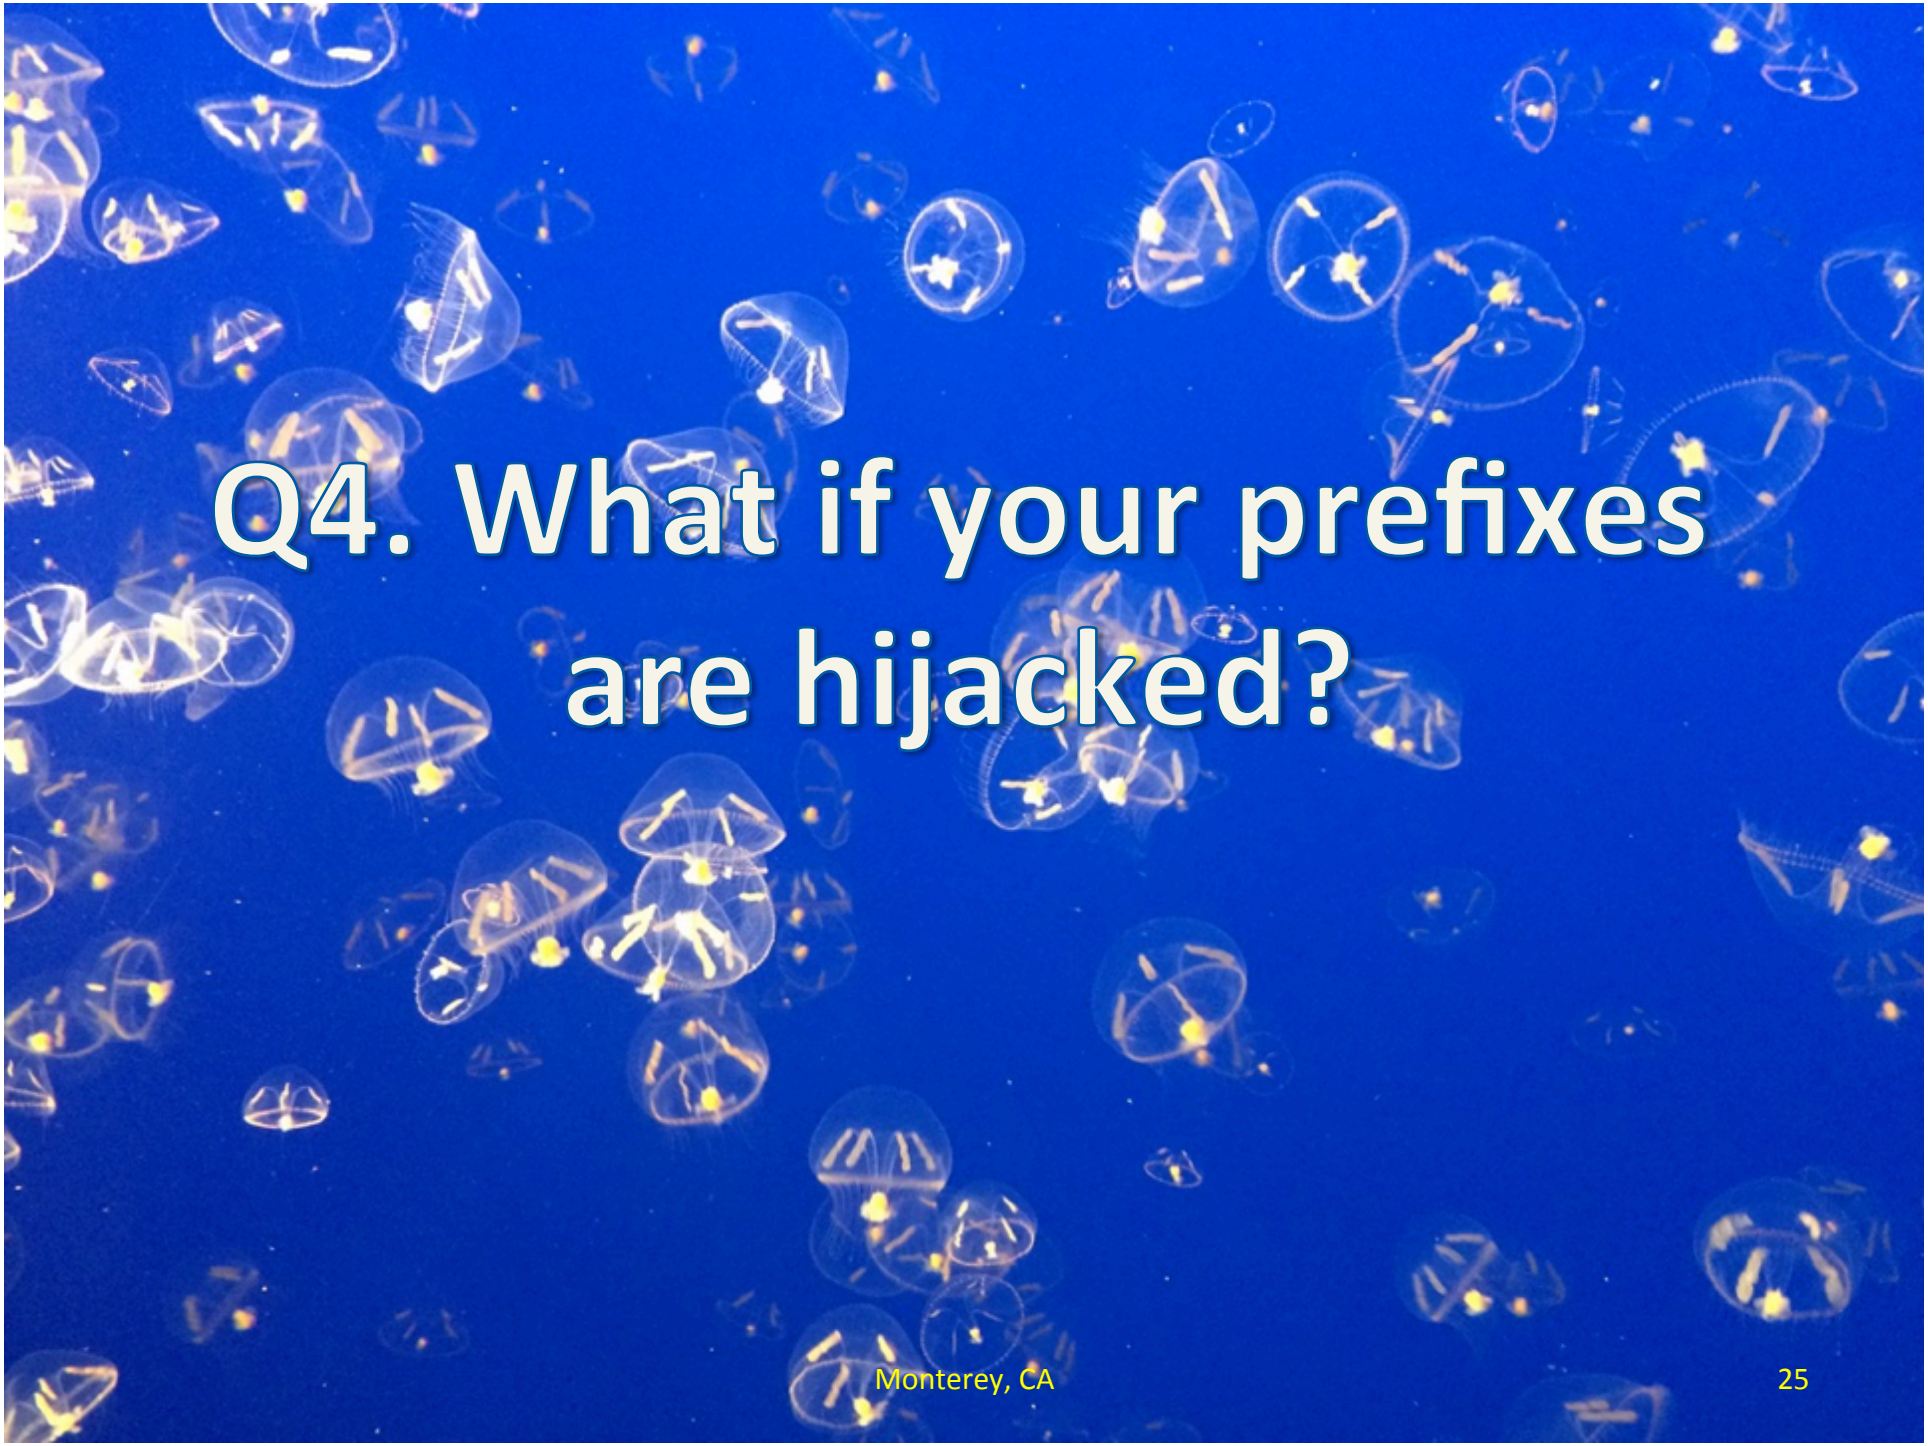
Q4. What if your prefixes are hijacked?

Advertise more specific prefix(es)

- When your prefix 10.0.0.0/16 is hijacked
  - Advertise /17
  - If other ASes accept the /17, traffic comes to you
  - If other Ases don't, it doesn't ☹
- ASes likely filter(ed) IPv4 /25 or longer and IPv6 / 64 or longer
- We're better to accept more specific masks
  - up to /28
  - IPv4 allocated mask getting more specific after IPv4 exhaustion
  - ARIN: allocates /24 - /28 from 23.128.0.0/10


Preventive Solutions
Against BGP Hijack Issues

Apply BGP Filters

- Not to receive/leak hijacked prefixes
- Not leak re-distributed prefixes to other protocols
  - EGP -> IGP -> EGP
- Not leak any prefixes used in test envirommment
- However, strict BGP filtering sometimes not match (ex. to upstream AS, to Peer)

BGP Origin Validation

- (Almost) Ultimate Solution
- Issue ROA so that other AS can validate your prefixes
- Introduce BGP Origin Validation so that your AS accepts legit prefixes
- Origin Validation can't be done by only one AS
  - ROA for each prefix is needed

Conclusion

- We experienced a huge BGP hijack issue on Nov 06
- No info of the root cause so far
- Minimum impact on AS2914
  - Whether we see a service impact depends on what prefixes are hijacked (or what service is done by using hijacked prefixes)
- Ideal idea is that every AS including Tier1 need not to advertise hijacked prefixes by using Origin Validation and BGP filtering etc.


Thank you!
Wild Sea Otter here!