Approaching to Secure Routing

ISOC Workshop @ APRICOT2014

Tomoya Yoshida
JPNIC IRR/RPKI Working Group Chair
JANOG RPKI Working Group co-Chairs
Internet Multifeed (JPNAP)
yoshdia@mfeed.ad.jp

Feb. 11,2014 ~ Feb. 12,2014

- 218.100.45.0/24 : JPNAP Tokyo II IX Prefix
  - Regularly not advertised to the Internet
- We detected **some one AS** advertise our Prefix
  - 2014/02/11 14:47:52 - 2014/02/12 5:40:41(UTC)
- Is this a fat finger or intentional?
- We checked at that time…

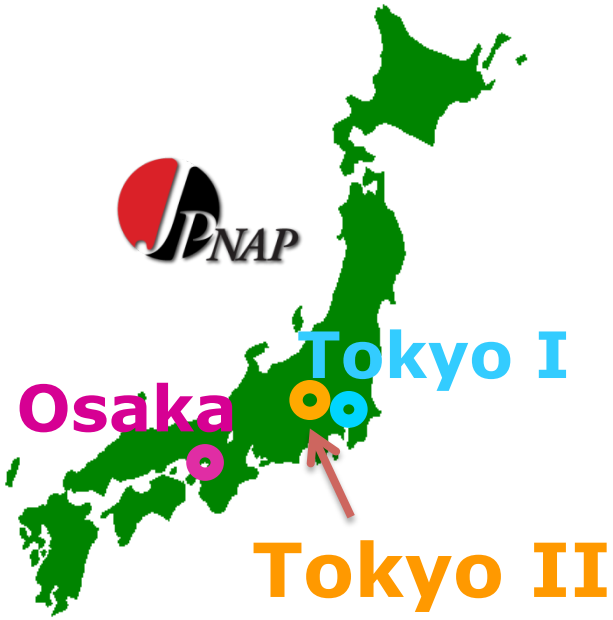

Perfixes originated from some one AS

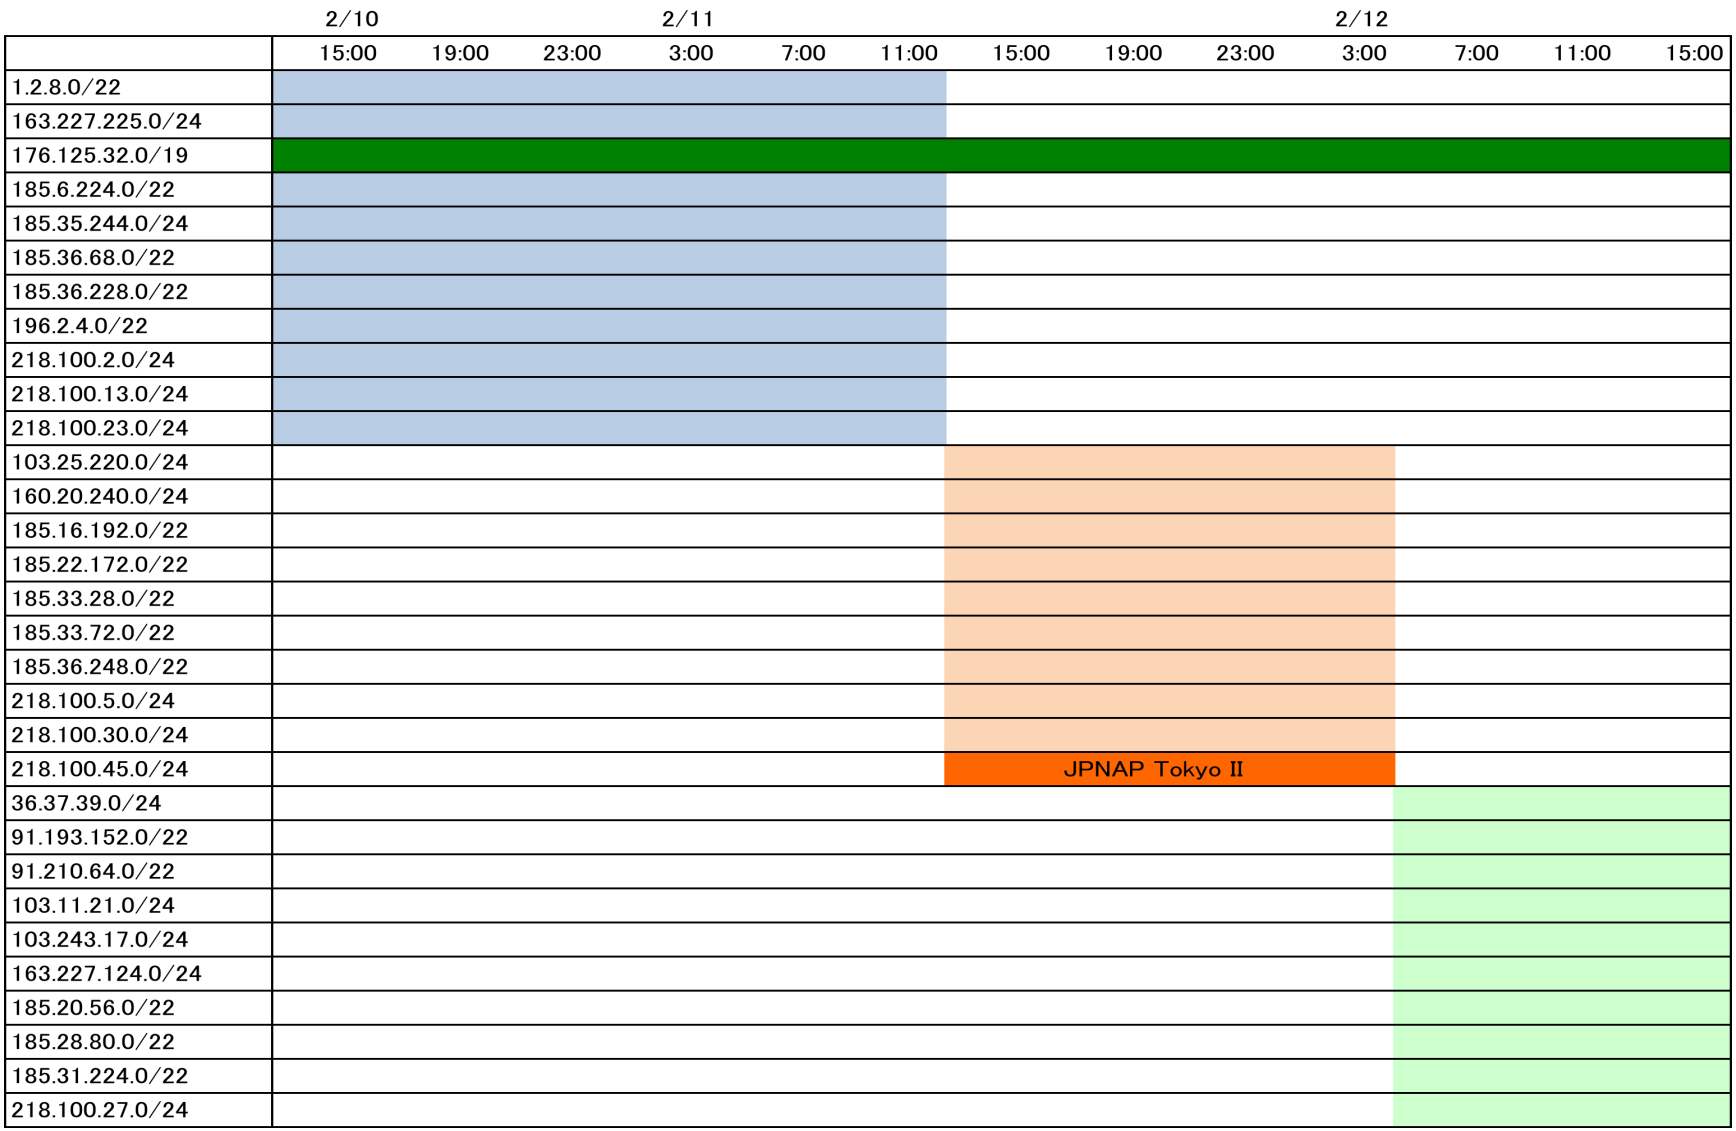

Many IX segments were hijacked

Prefix	Desc
218.100.2.0/24	Sydney IX Lan
218.100.5.0/24	OBIS-IX,Internet Exchange Point,Okayama,Japan
218.100.13.0/24	Melbourne IX Lan
218.100.23.0/24	Dunedin Peering Exchange
218.100.27.0/24	OpenIXP, Internet Exchange Point, Indonesia
218.100.30.0/24	APJII Indonesia Internet eXchange
218.100.45.0/24	JPNAP Tokyo II IX

E-mail from spamcom

spamcop.net

Help | Site Map | Text size: - +

Login

Report Spam | Filtered Email | Blocking List | Statistics | Login

SpamCop v 4.8.1.007 © 2014 Cisco Systems, Inc. All rights reserved.
Here is your TRACKING URL - it may be saved for future reference:
http://www.spamcop.net/sc?id=z5729621514zf033f7ded6df91c29bf9908db8e0d513z
Skip to Reports

```
Return-path: <Motorola@wappextil.com>
Received: from wappextil.com ([unknown] [218.100.45.34])
 by vms172083.mailsrvcs.net
 (Sun Java(tm) System Messaging Server 7u2-7.02 32bit (built Apr 16 2009))
 with ESMTP id <0N0V004K08E3TI20@vms172083.mailsrvcs.net> for
 x; Tue, 11 Feb 2014 22:27:49 -0600 (CST)
Received: by wappextil.com id hvbsaalhvj41 for <x>; Tue,
 11 Feb 2014 23:22:31 -0500 (envelope-from <Motorola@wappextil.com>)
Date: Wed, 12 Feb 2014 04:22:30 +0000
From: "Motorola 7214186" <possible@wappextil.com>
Subject: A sweet deal! Moto X. No contract. No down payment. No hassles.
X-Originating-IP: [218.100.45.34]
Message-id: <0N0V________TI20@vms172083.mailsrvcs.net>
MIME-version: 1.0
Content-type: text/html
Content-transfer-encoding: 7BIT
Original-recipient: rfc822;x
```

View entire message

Parsing header:

```
Received:  from wappextil.com ([unknown] [218.100.45.34]) by vms172083.mailsrvcs.net (Sun Java(tm) System Messaging Server 7u2-7.02 32bit (built Apr 16 2009))
with ESMTP id <0N0V004K08E3TI20@vms172083.mailsrvcs.net> for x; Tue, 11 Feb 2014 22:27:49 -0600 (CST)
```

host 218.100.45.34 (getting name) no name
Possible spammer: 218.100.45.34
Received line accepted

```
Received:  by wappextil.com id hvbsaalhvj41 for <x>; Tue, 11 Feb 2014 23:22:31 -0500 (envelope-from <Motorola@wappextil.com>)
```

no from
Ignored
218.100.45.34 not listed in cbl.abuseat.org
218.100.45.34 not listed in dnsbl.sorbs.net
218.100.45.34 is not an MX for vms172083.mailsrvcs.net
218.100.45.34 is not an MX for vms172083.mailsrvcs.net

Tracking message source: 218.100.45.34:

Routing details for 218.100.45.34
[refresh/show] Cached whois for 218.100.45.34 : tech-c@mfeed.ad.jp
Using last resort contacts tech-c@mfeed.ad.jp
Sorry, this email is too old to file a spam report. You must report spam within 2 days of receipt. This mail was received on Tue, 11 Feb 2014 22:27:49 -0600

E-mail from spamcom

spamcop.net

Help | Site Map | Text size: - +

Login

Report Spam | Filtered Email | Blocking List | Statistics | Login

SpamCop v 4.8.1.007 © 2014 Cisco Systems, Inc. All rights reserved.
Here is your TRACKING URL - it may be saved for future reference:
http://www.spamcop.net/sc?id=z5729621514zf033f7ded6df91c29bf9908db8e0d513z
Skip to Reports

```
Return-path: <Motorola@wappextil.com>
Received: from wappextil.com ([unknown] [218.100.45.34])
 by vms172083.mailsrvcs.net
 (Sun Java(tm) System Messaging Server 7u2-7.02 32bit (built Apr 16 2009))
 with ESMTP id <0N0V004K08E3TI20@vms172083.mailsrvcs.net> for
 x; Tue, 11 Feb 2014 22:27:49 -0600 (CST)
Received: by wappextil.com id hvbsaa1hvj41 for <x>; Tue,
 11 Feb 2014 23:22:31 -0500 (envelope-from <Motorola@wappextil.com>)
Date: Wed, 12 Feb 2014 04:22:30 +0000
From: "Motorola 7214186" <possible@wappextil.com>
Subject: A sweet deal! Moto X. No contract. No down payment. No hassles.
X-Originating-IP: [218.100.45.34]
Message-id: <0N0V________TI20@vms172083.mailsrvcs.net>
```

218.100.45.34 is not an MX for vms172083.mailsrvcs.net
218.100.45.34 is not an MX for vms172083.mailsrvcs.net

Tracking message source: 218.100.45.34:
Routing details for 218.100.45.34
[refresh/show] Cached whois for 218.100.45.34 : tech-c@mfeed.ad.jp
Using last resort contacts tech-c@mfeed.ad.jp
Sorry, this email is too old to file a spam report. You must report spam within 2 days of receipt. This mail was received on Tue, 11 Feb 2014 22:27:49 -0600

Agenda

- What is the mis-origination(hijacking)
- Japanese Activities for Secure Routing
- Considerations

Case-1

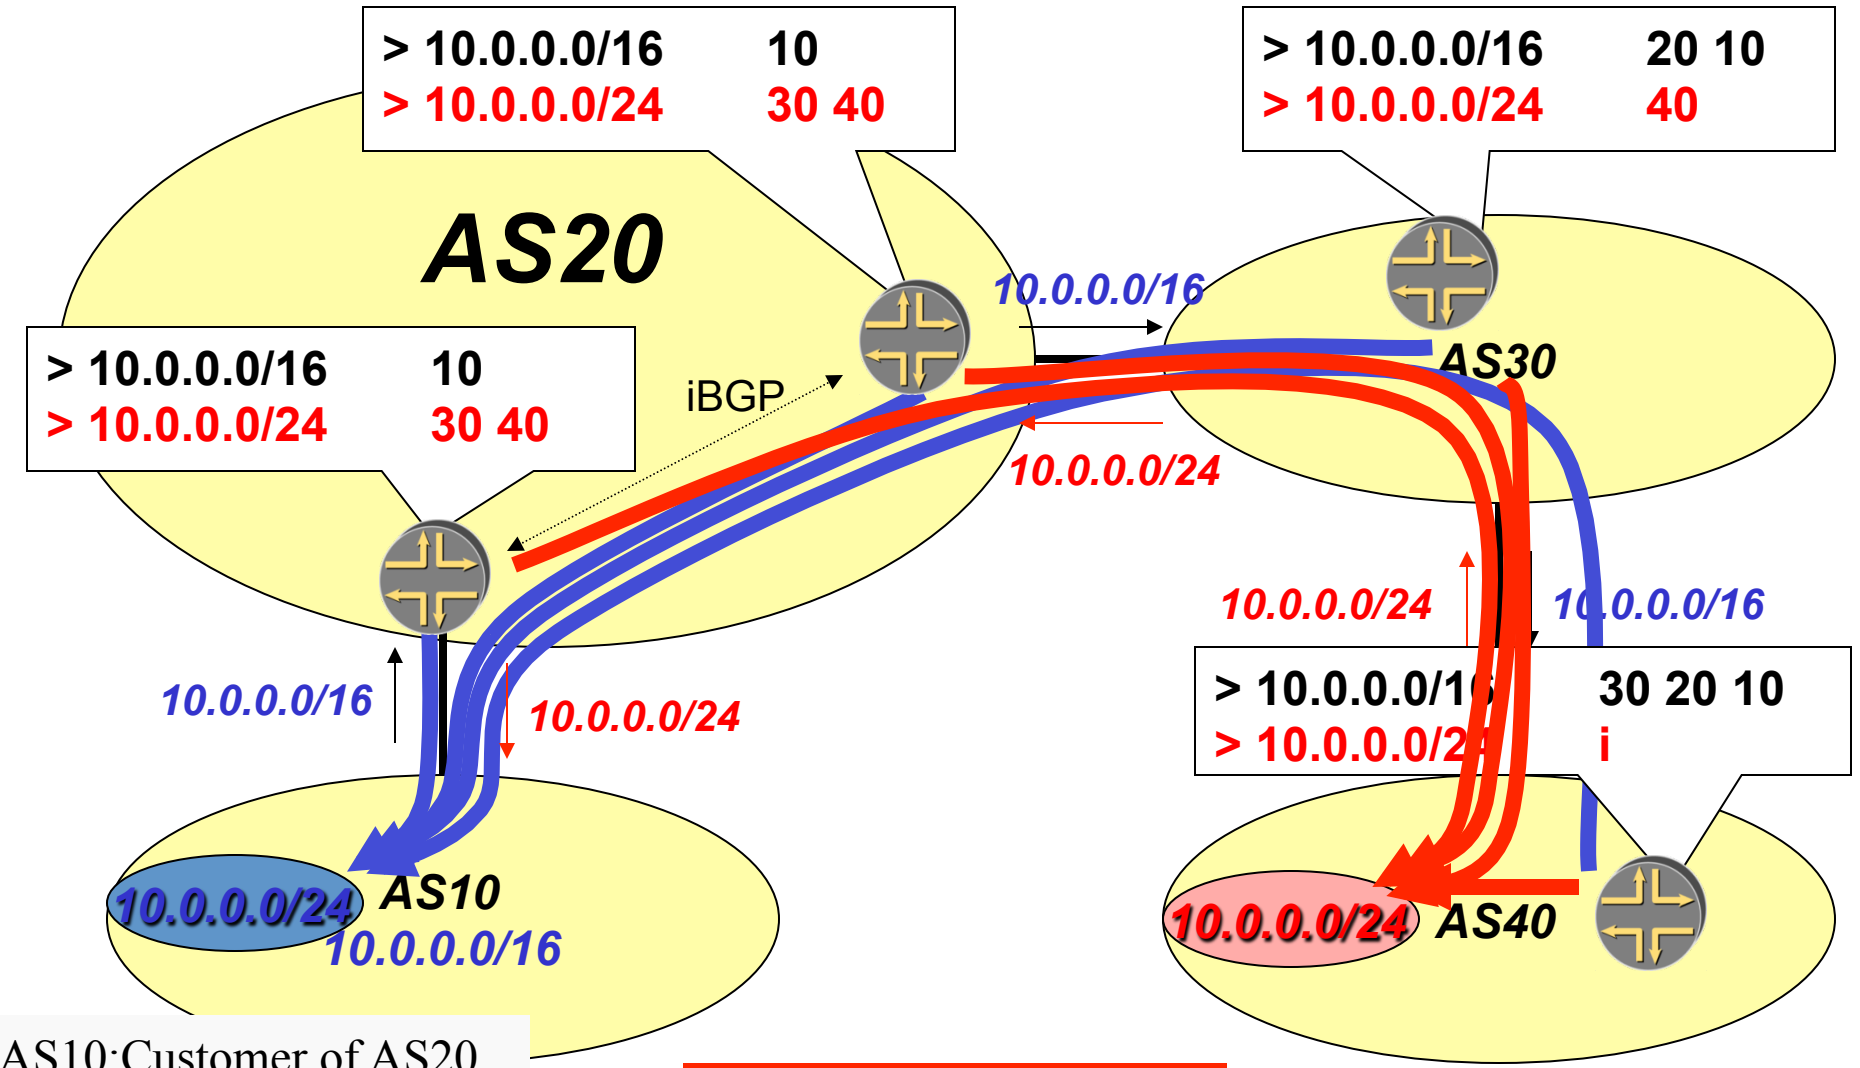

AS10:Customer of AS20
AS40:Customer of AS30
AS20 and AS30 is peering

Global Impact

Case-2

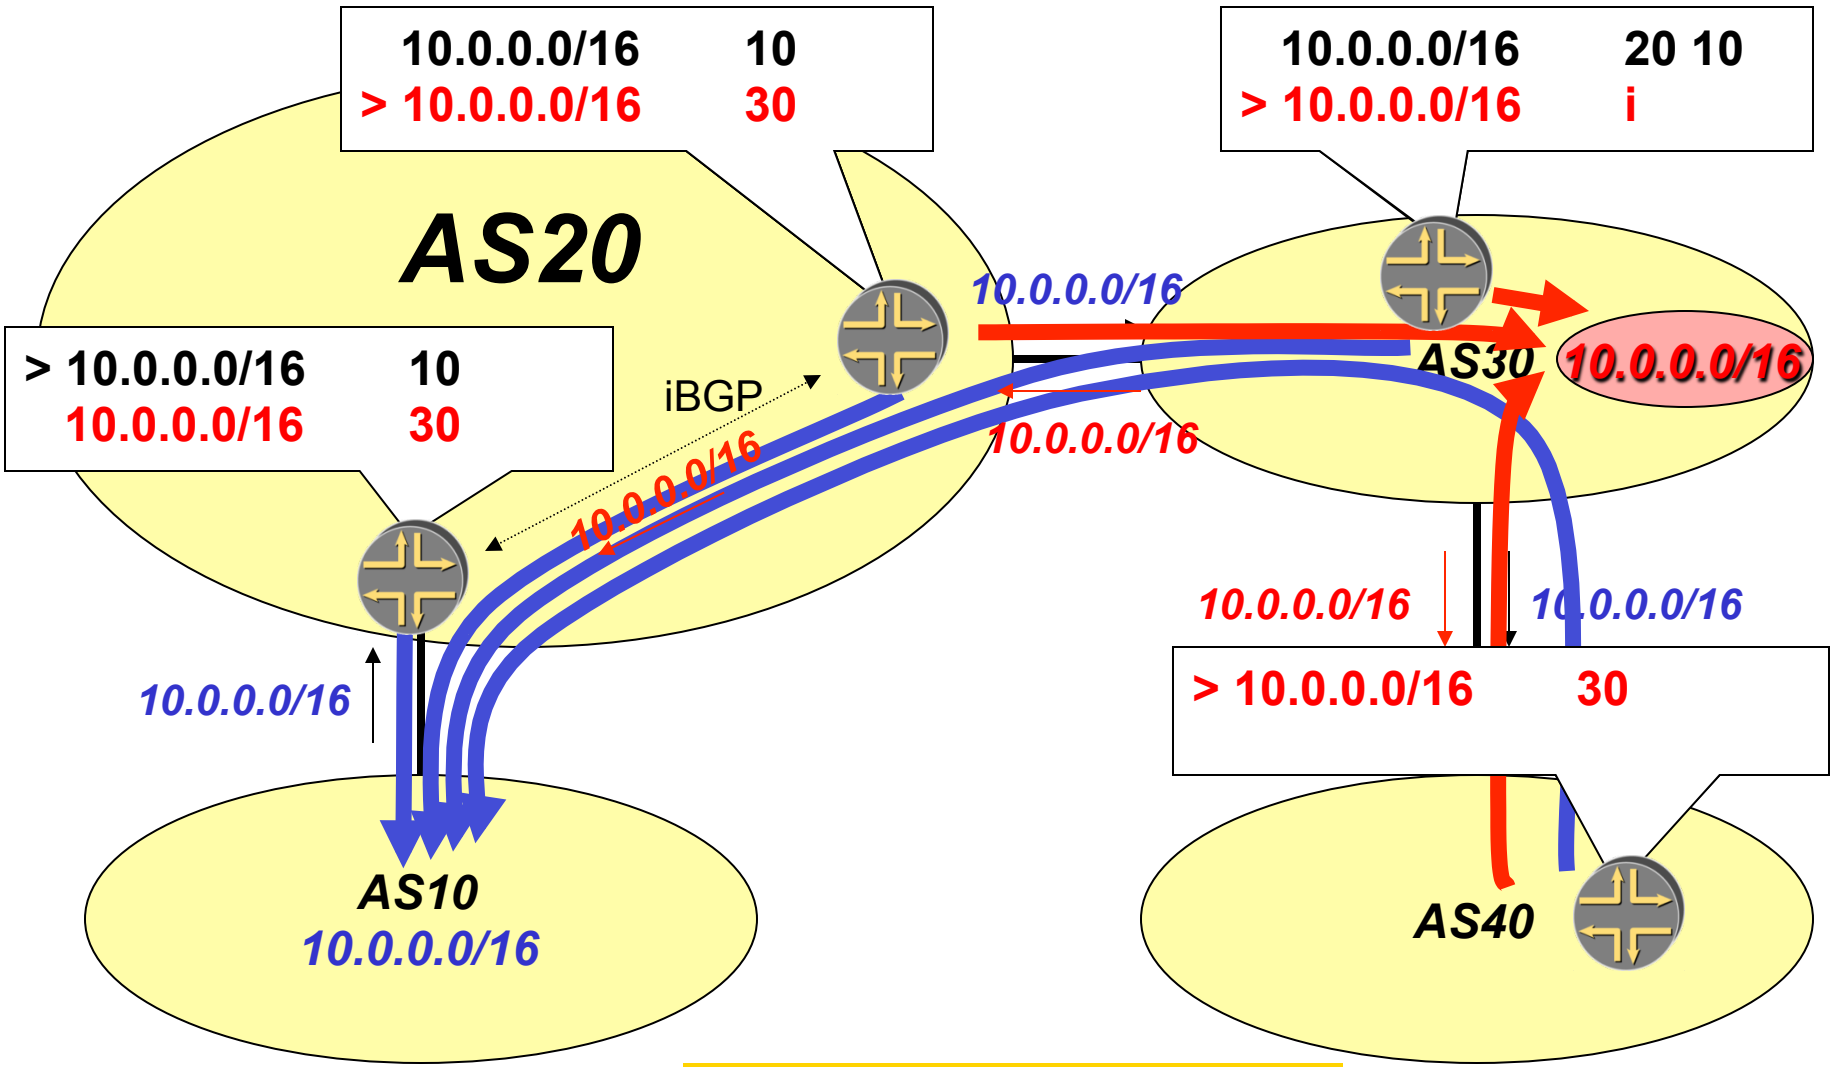

Regional Impact

mis-origination (routing hijacking)

- Two causes
  - Operational fault (Fat finger)
  - Intentional fault

- “youtube incident” in 2008 came both Fat finger and Intentional fault
  - Pakistan Telecom announced youtube prefix inside their country to divert
  - Upstreme transit ISP accidently propagated to the Internet

(1) Fat finger

1. Typing a wrong IP Address
   - Mostly 2nd Octet and 3rd Octet
   - In some cases, both IRR registration and BGP advertisement is wrong…
2. Re-distribution with your origin AS
   - Ex) IX Prefix, PrivatePeer IP, Re-redistribution
3. Advertising the prefixes using inside for Test
   - Ex) 1/8, 2/8
4. Forgetting to add no-export BGP community for black holing
5. Exchanging the Prefix information using BGP but accidently leak those…

(2) Intentional

1. Collection the packets
   - Longer prefix
   - Shorter prefix /0 /1 etc

2. Short Ribed BGP
   - Temporary advertise some prefix and SPAM at the same time

3. Cyber Terrorism

Examples observed in Japan (old days)

	Case-1	Case-2	Case-3
When	2004/6	2004/9	2006/11
Invalid Origin	Japanese ISP	Asian ISP	Asian ISP
Prefix	Longer, Invalid /24x2, /25x1, /29x1	Longer, Invalid /24x2	Same, Invalid /14x2, /17x2
Action	Asked to the Origin ISP and stopped	Asked to Upstreme ISP and stopped	NO Action (later withdrawn)
Impact	About 150minutes	About 2 days	5 minutes
			Many other routes were hijacked

What we are doing?

1. Filtering
   - Making use of JPIRR in JP

2. Minimizing the influence
   - Detection (経路奉行, ISAlarm, BGPMON etc)
   - Analysis (whois, IRR, looking glass)
   - Action (ask to peer/transit ISP using NOG contact)

JPIRR

- Internet Routing Registry
- Lunched in 2002
- Currently about 70% Japanese ASes registered to JPIRR
  - For Customer's Filter
  - For Detection System (Master IP/origin Database)

Issues using IRR

- We have many choices for IRR server
  - RADB, JPIRR, NTTCOM, etc…
- According to the IRR Server, the Data which we can see is different
  - Depends on the IRR mirroring each other
- Object Name is independent per IRR
  - Not Coordinated, sometime it happen a confliction
- Reliability of the Information
- Service spec is various

Issues using IRR

- We have many choices for IRR server
  - RADB, JPIRR, NTTCOM, etc…
- According to the IRR Server, the Data which we can see is different
  - Depends on the IRR mirroring each other
- Object Name is independent per IRR
  - Not Coordinated, sometime it happen a confliction
- Reliability of the Information
- Service spec is various

Resource Public Key Infrastructure

RPKI

Attestation of the routing information

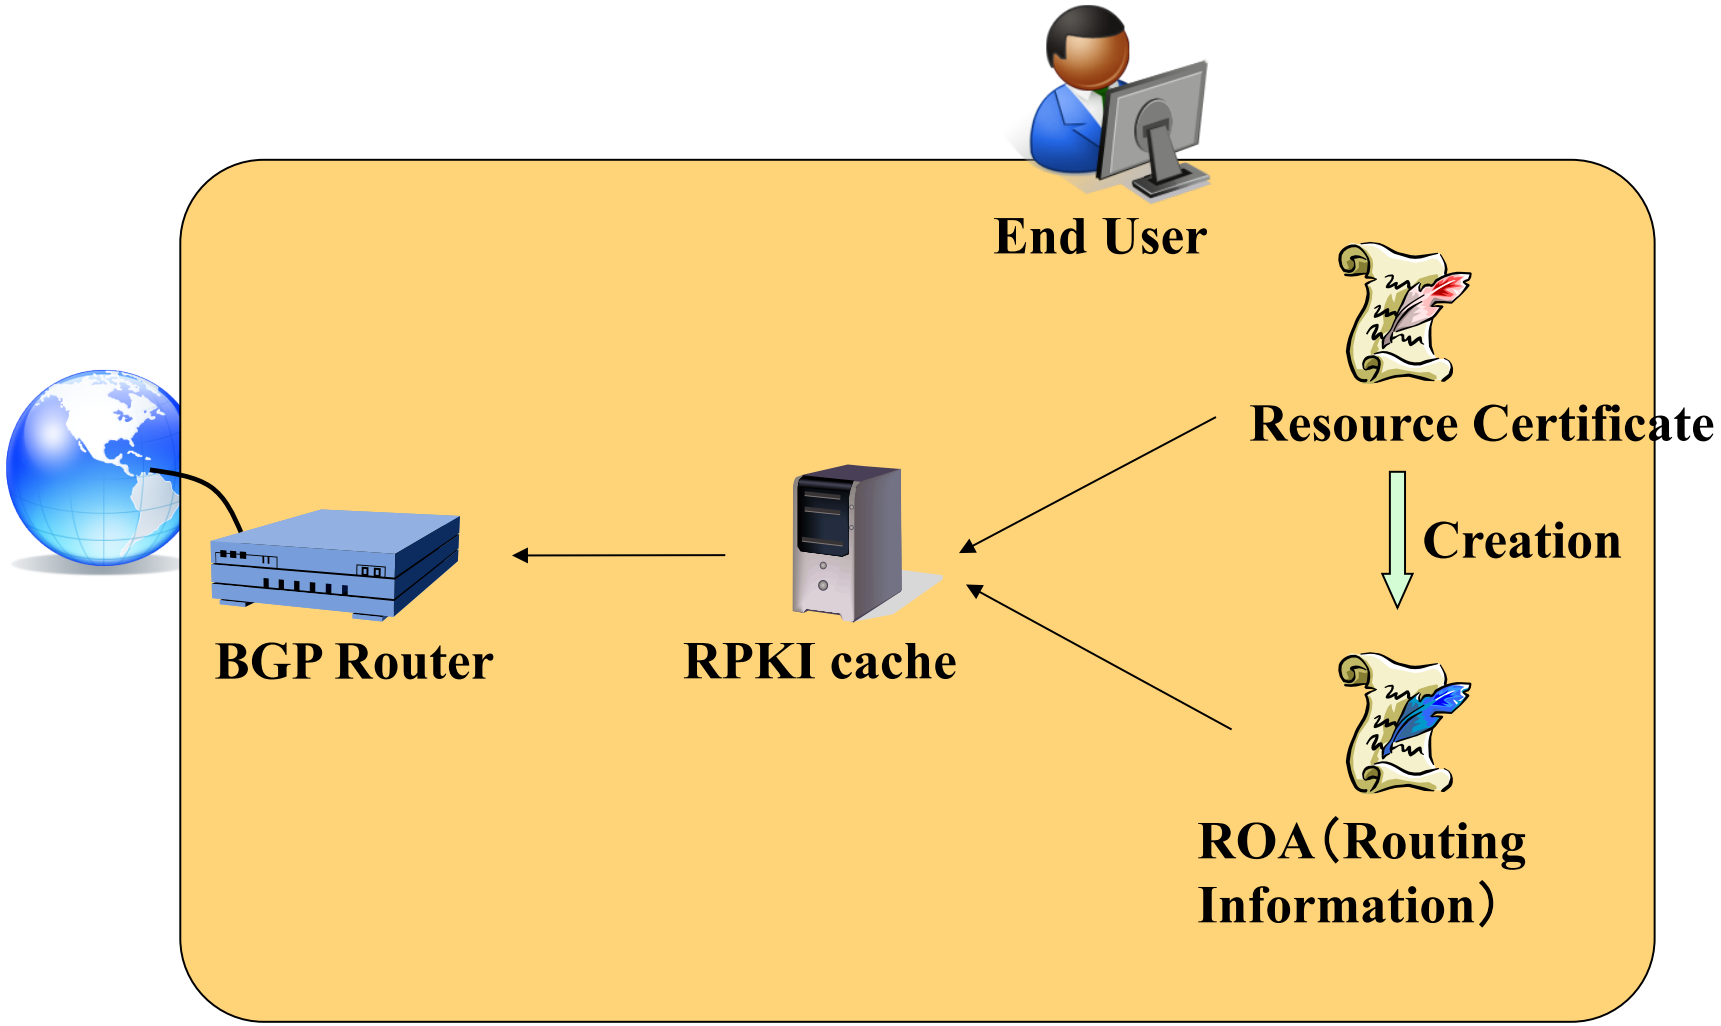

Creation of ROA

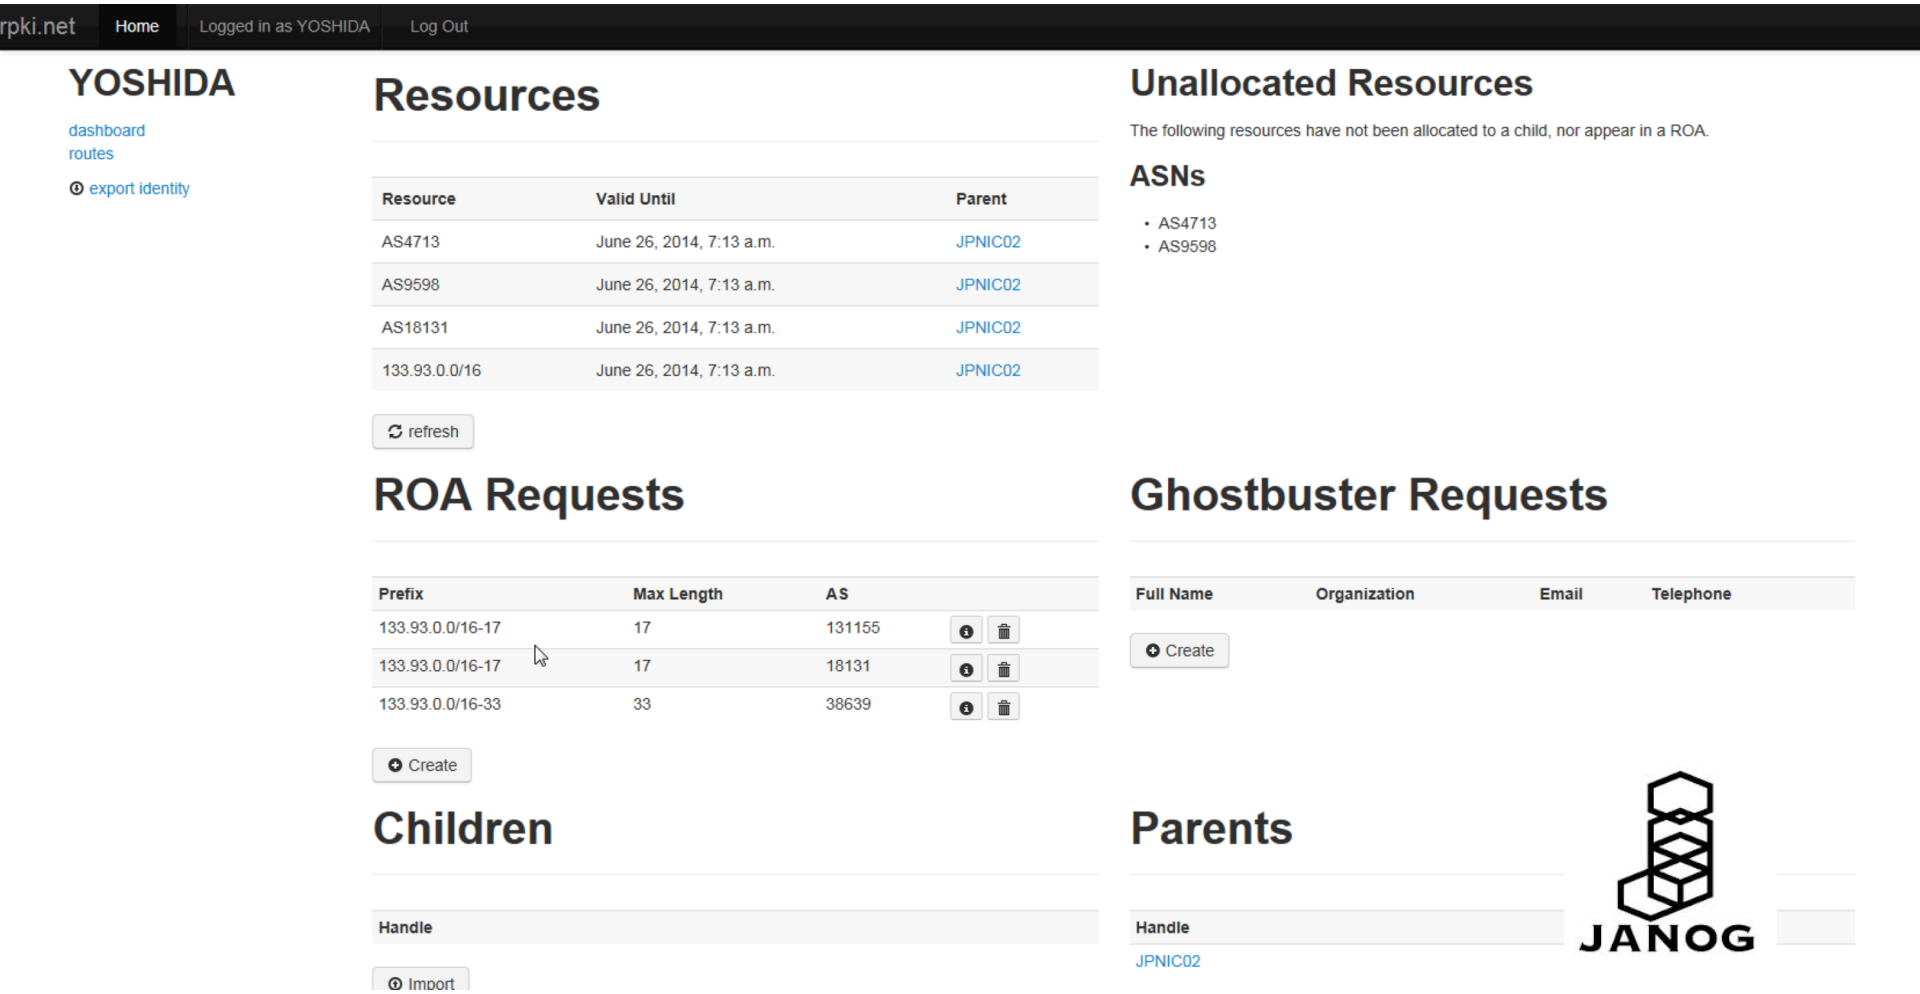

RTR Protocol(getting ROA)

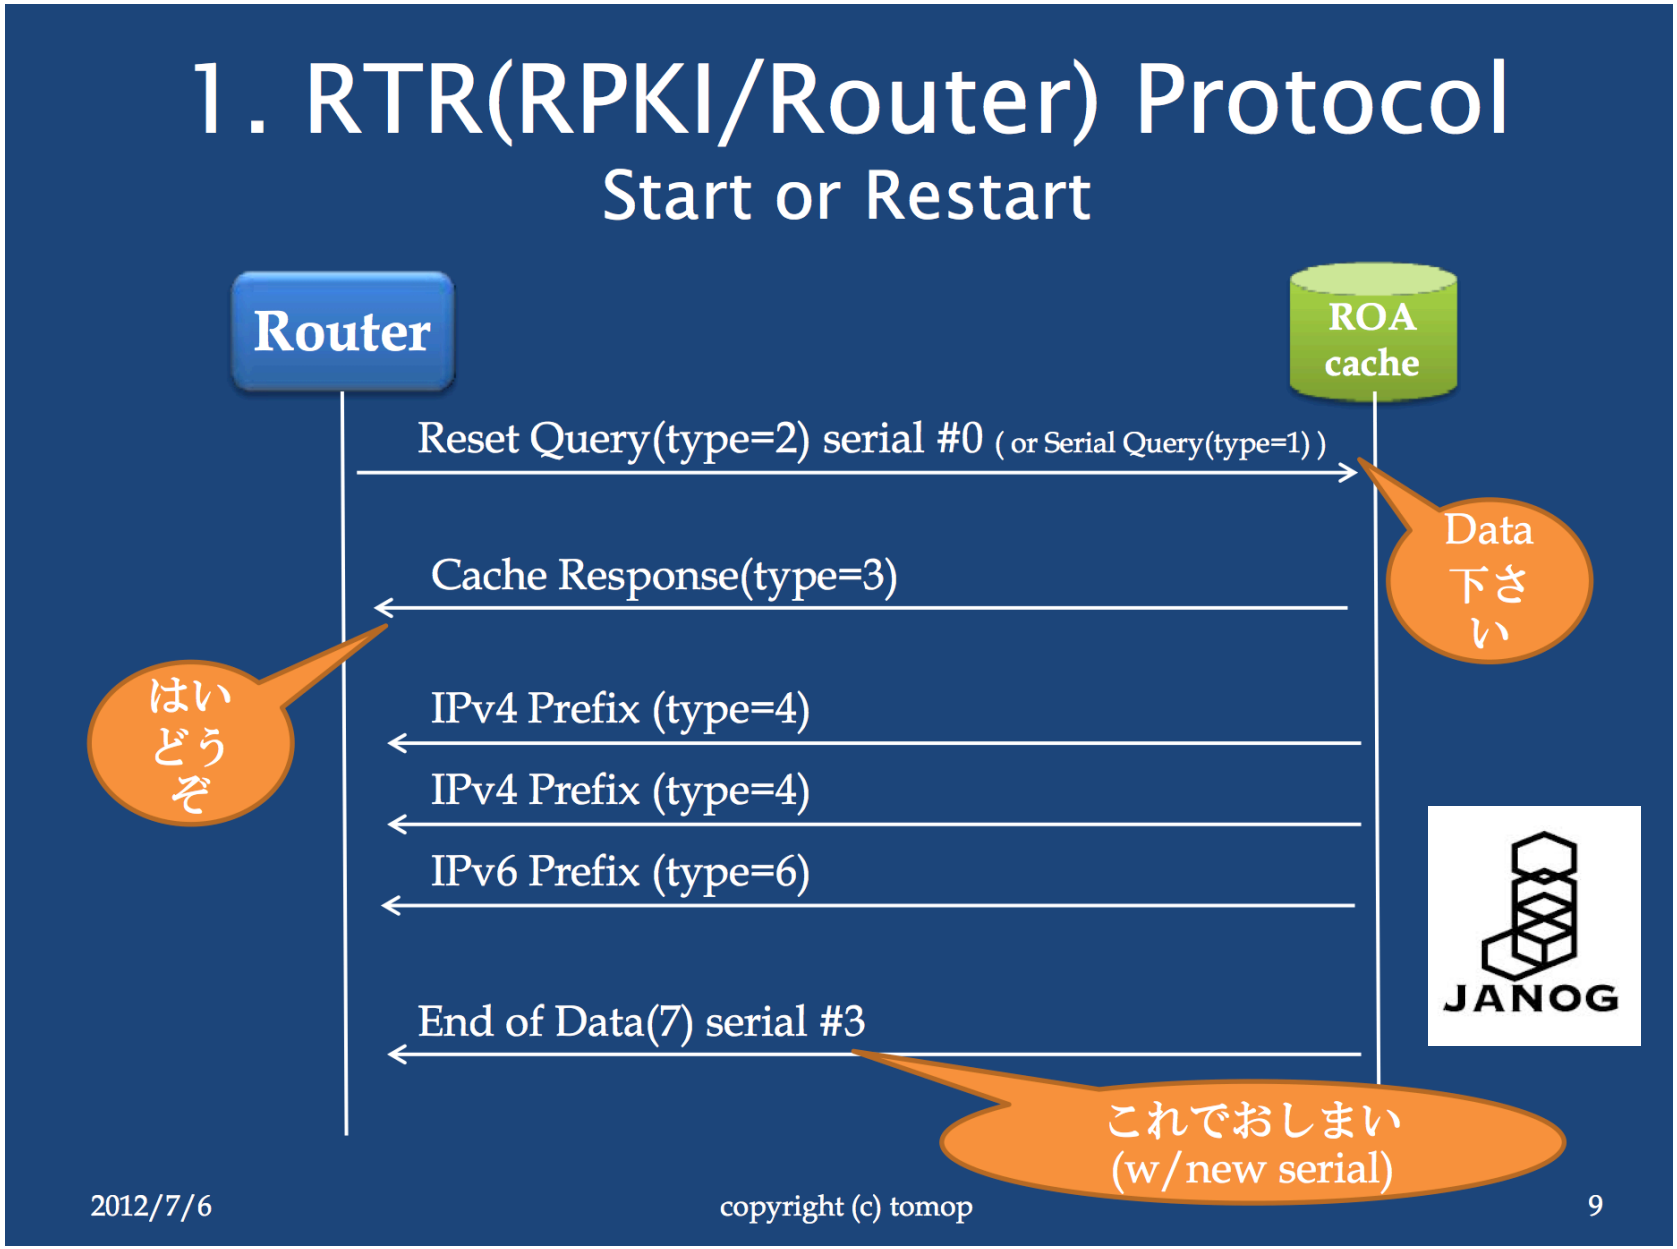

Case-1

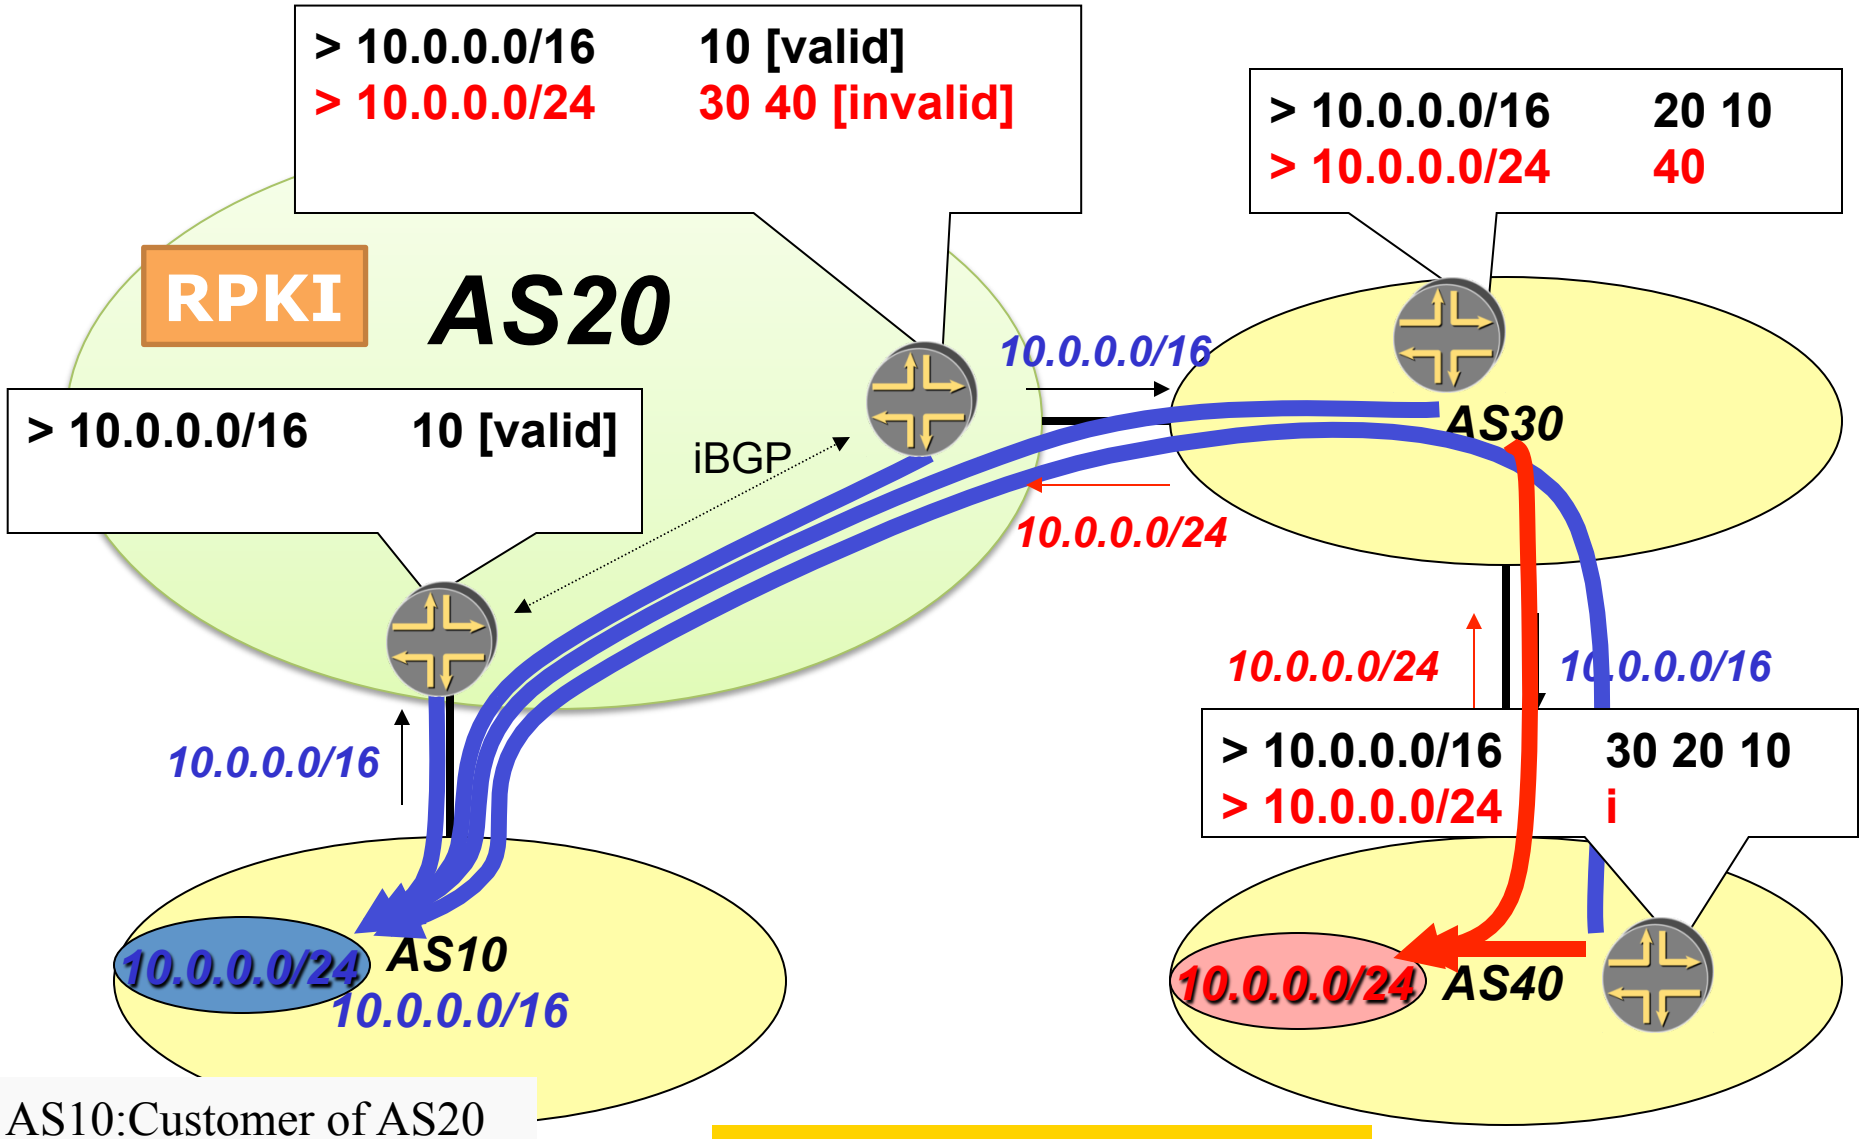

AS10:Customer of AS20
AS40:Customer of AS30
AS20 and AS30 is peering

Regional Impact

Filter of Router

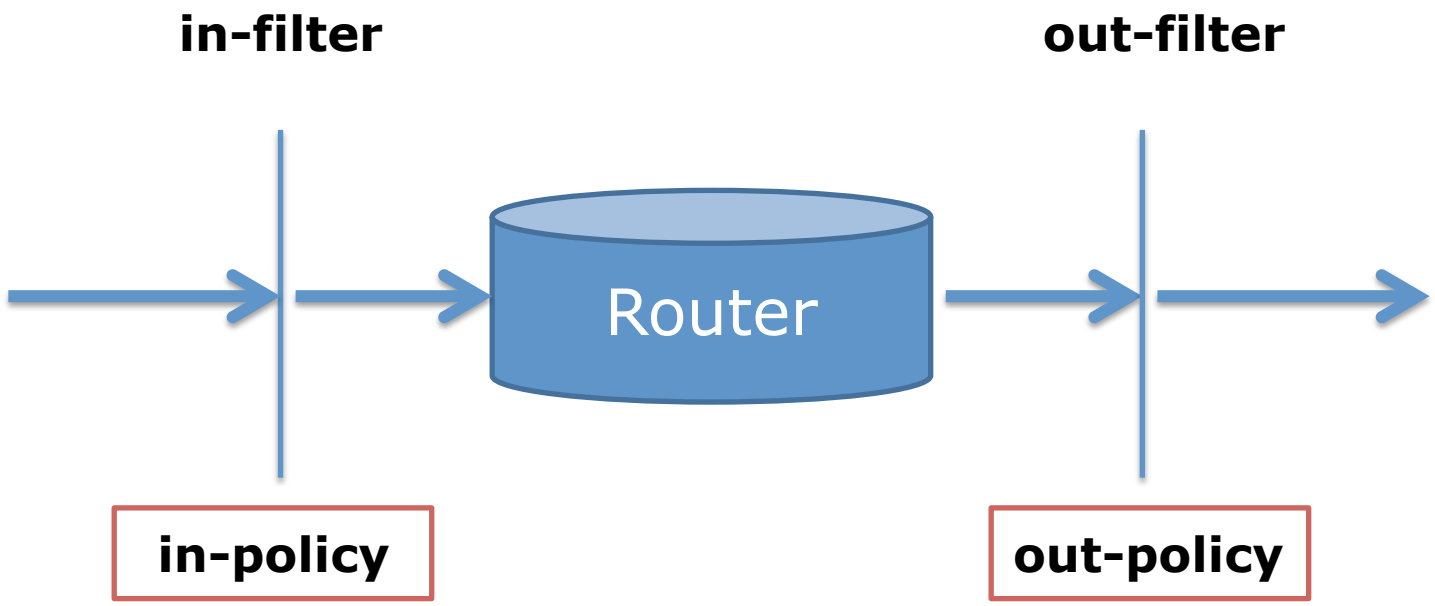

Filter of Router

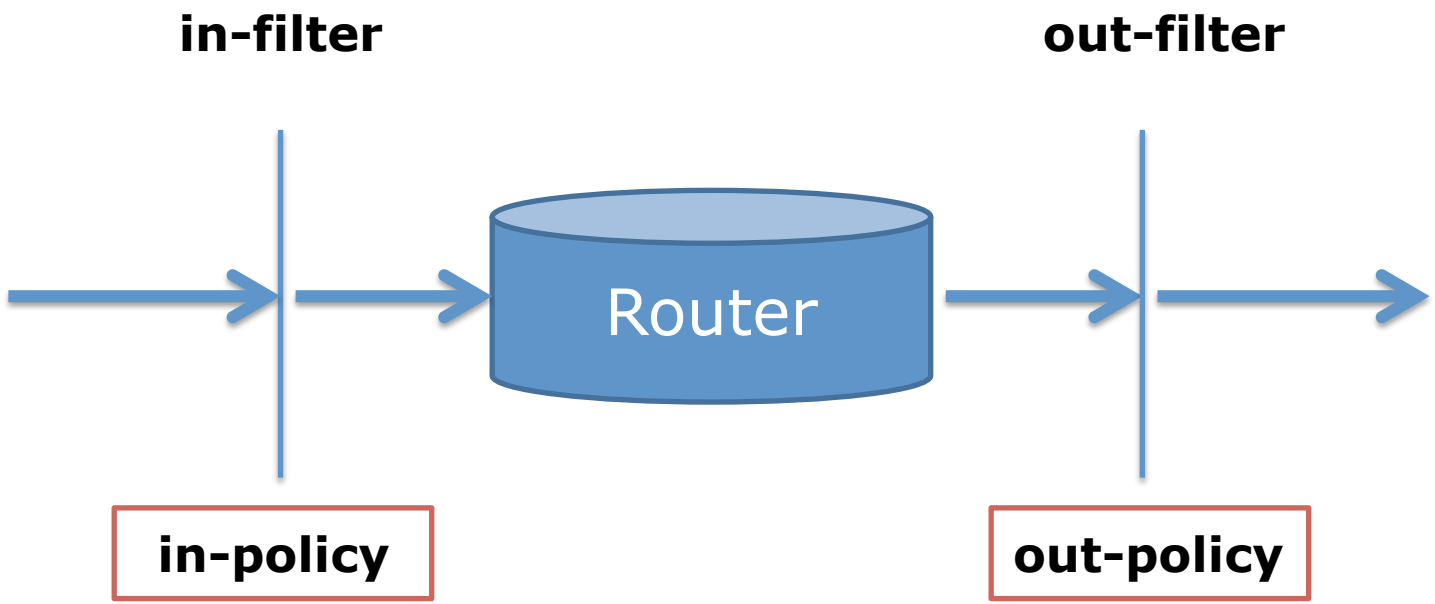

Bogon filter
Max-prefix
Prefix/AS_PATH filter
Community
LP, MED

Prepend
Community
MED
Prefix/AS_PATH filter

Filter of Router (w/RPKI)

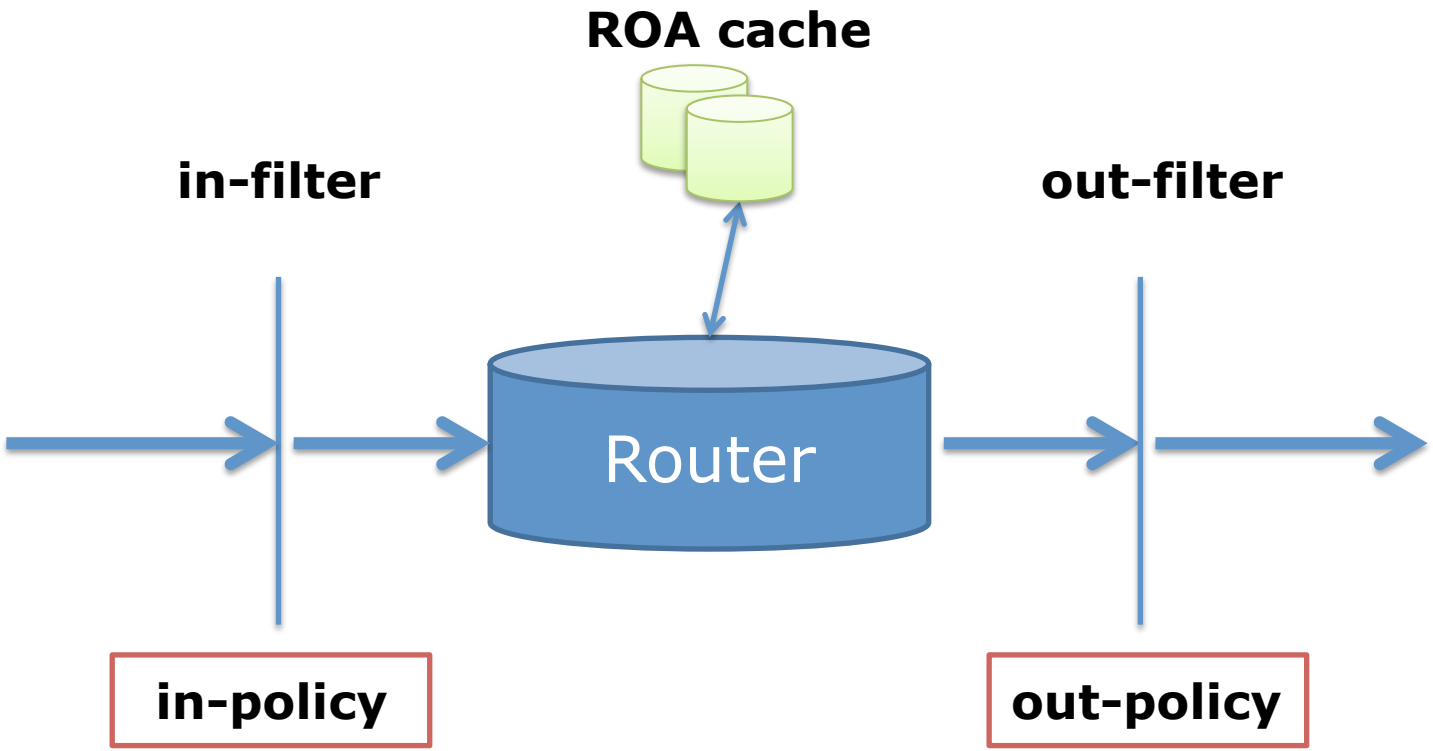

Bogon filter
Max-prefix
Prefix/AS_PATH filter
Community
LP, MED
RPKI (validation result)

Prepend
Community
MED
Prefix/AS_PATH filter
RPKI (validation result)

Impact for Routing (1)

- When RTR session goes down accidentally, validation result may be “not found” on the settings depending on the cache timer on your router
  - This means the routes which you would like to reject using the RPKI validation result cannot be rejected temporary

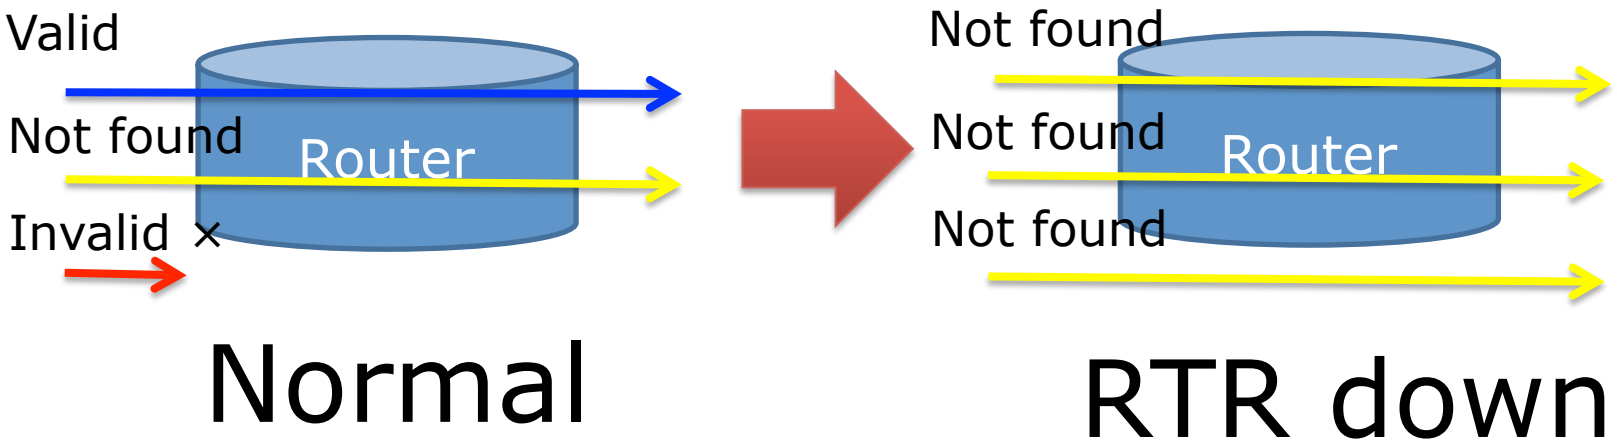

Impact for Routing (1)

- Redundancy for RTR session
  - Like a redundancy of Route Reflector
- Preparation even if the session goes down
  - ROA cache timer of the router
  - Policy rule
    - I don’t recommend to reject “not found”

Impact for Routing (2)

- When your router reboot accidentally, need to care the convergence time of RIB/FIB route and RTR
  - This means also the routes which you would like to reject using the RPKI validation result cannot be rejected temporary
  - Static filtering will not be influenced as of the RPKI

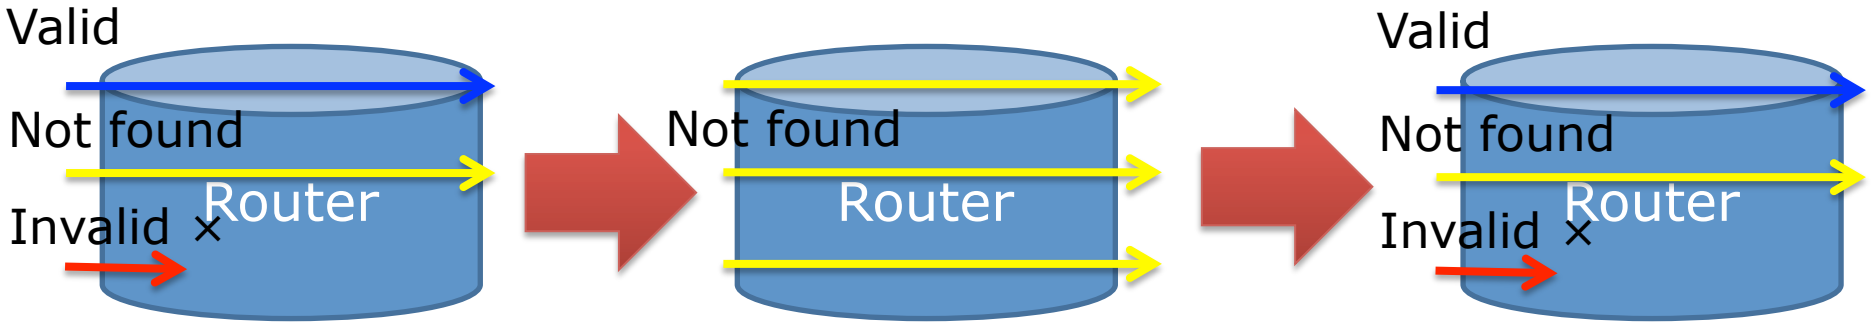

Normal

Reload To Normal

Normal

Impact for Routing (2)

- wait-for-bgp (Ref: RFC3137) like implementation may be needed
  - At the hierarchy of router ospf, you can configure “wait-for-bgp” : The router set the ospf cost “max-metric 65535” till the finishing of receiving the fullroute

- Proposal of wait-for-rpki(roa)
  - Waiting to go back to normal ospf cost till the rtr session goes back and ready to validate

Impact for Routing (3)

- If the cache DB's contents accidentally be withdrawn or cannot to be seen correctly, the result of RPKI validation may be "Not found"

Normal

DB…

Some case(multiple origin)

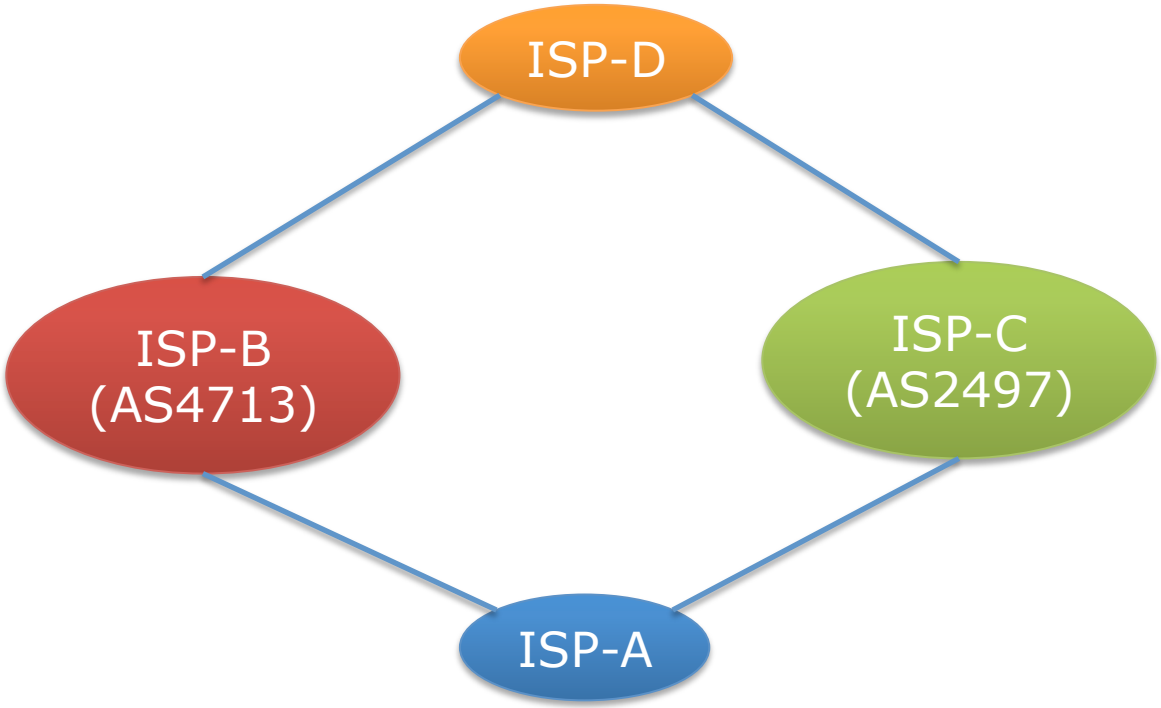

65000
(private)

Some case(multiple origin)

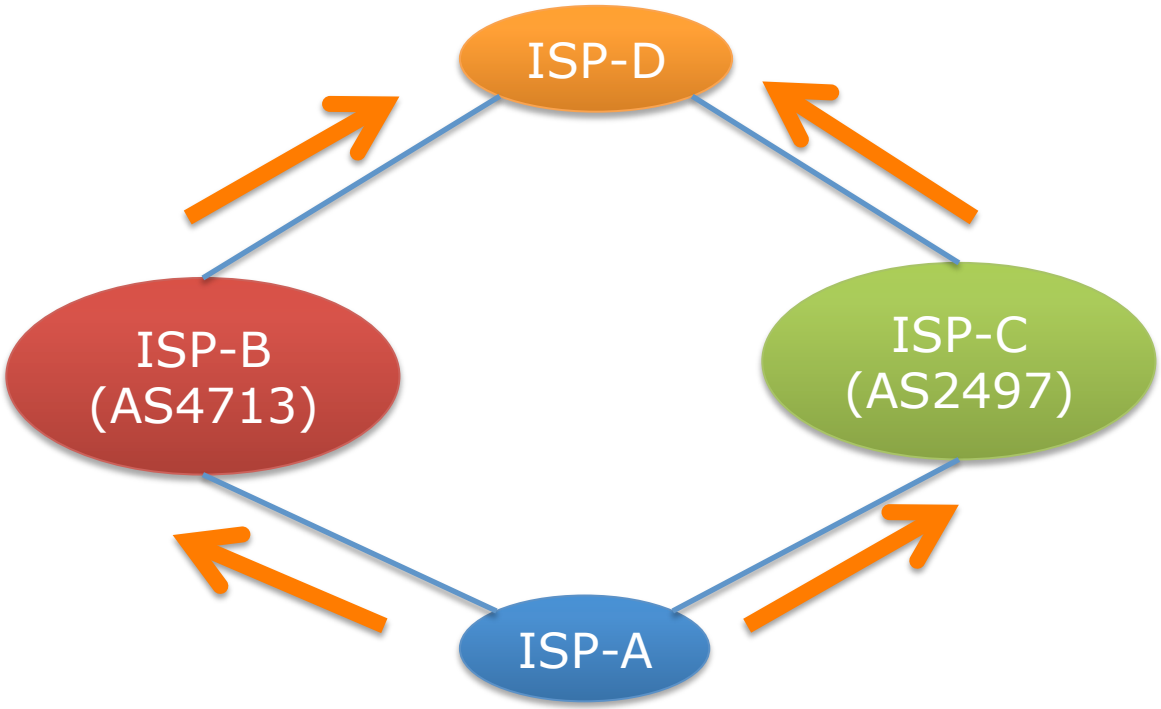

65000
(private)

Some case(multiple origin)

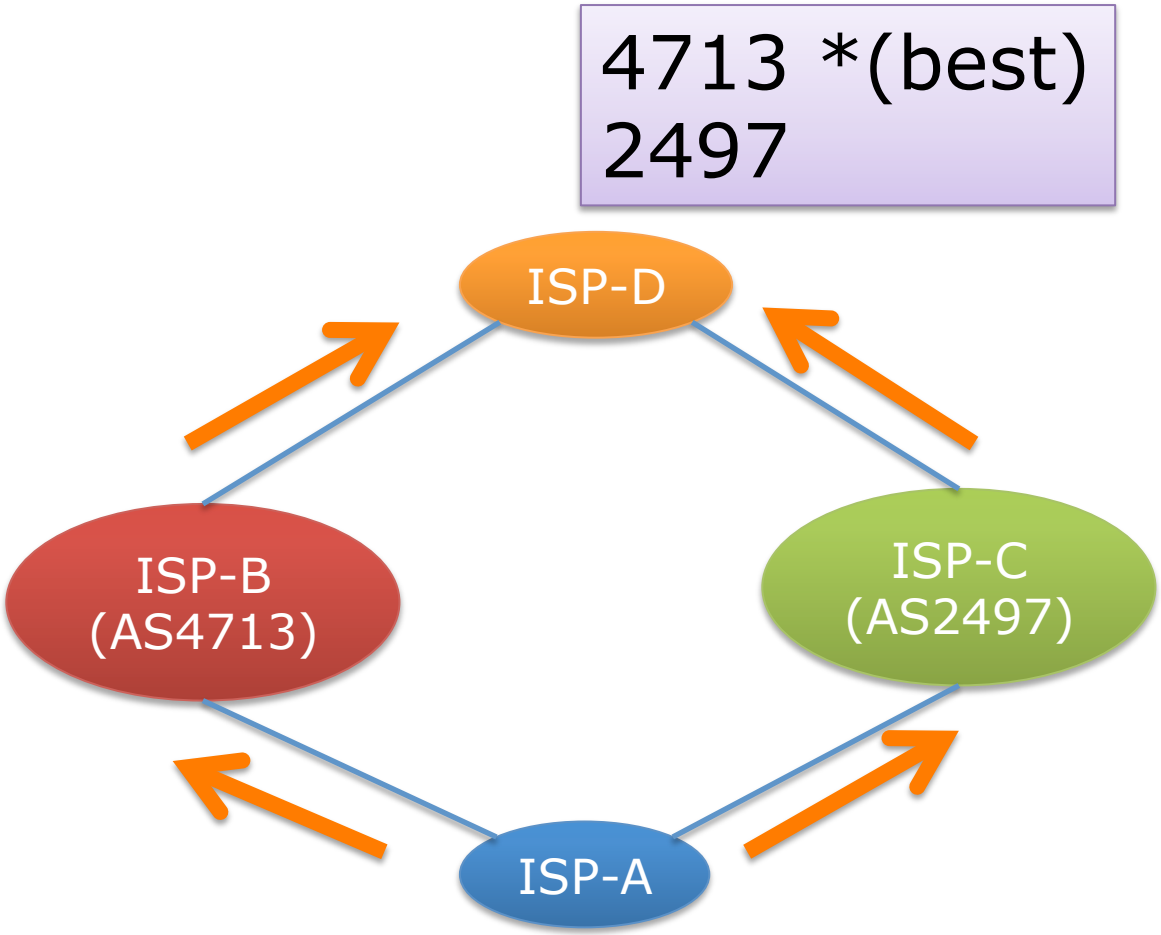

Some case(multiple origin)

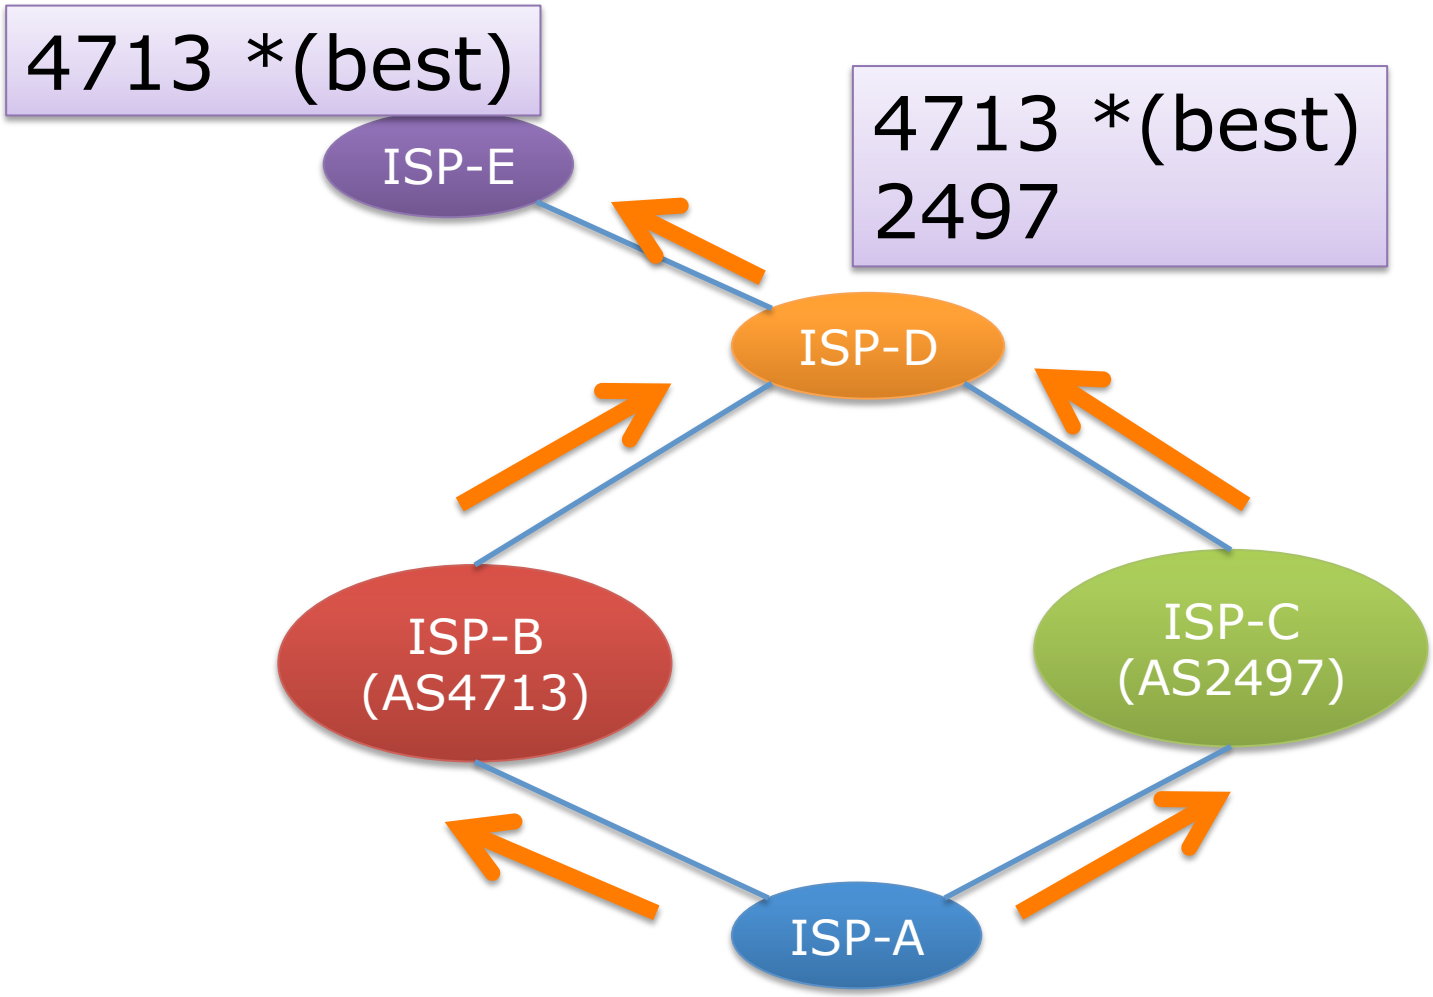

Some case(multiple origin)

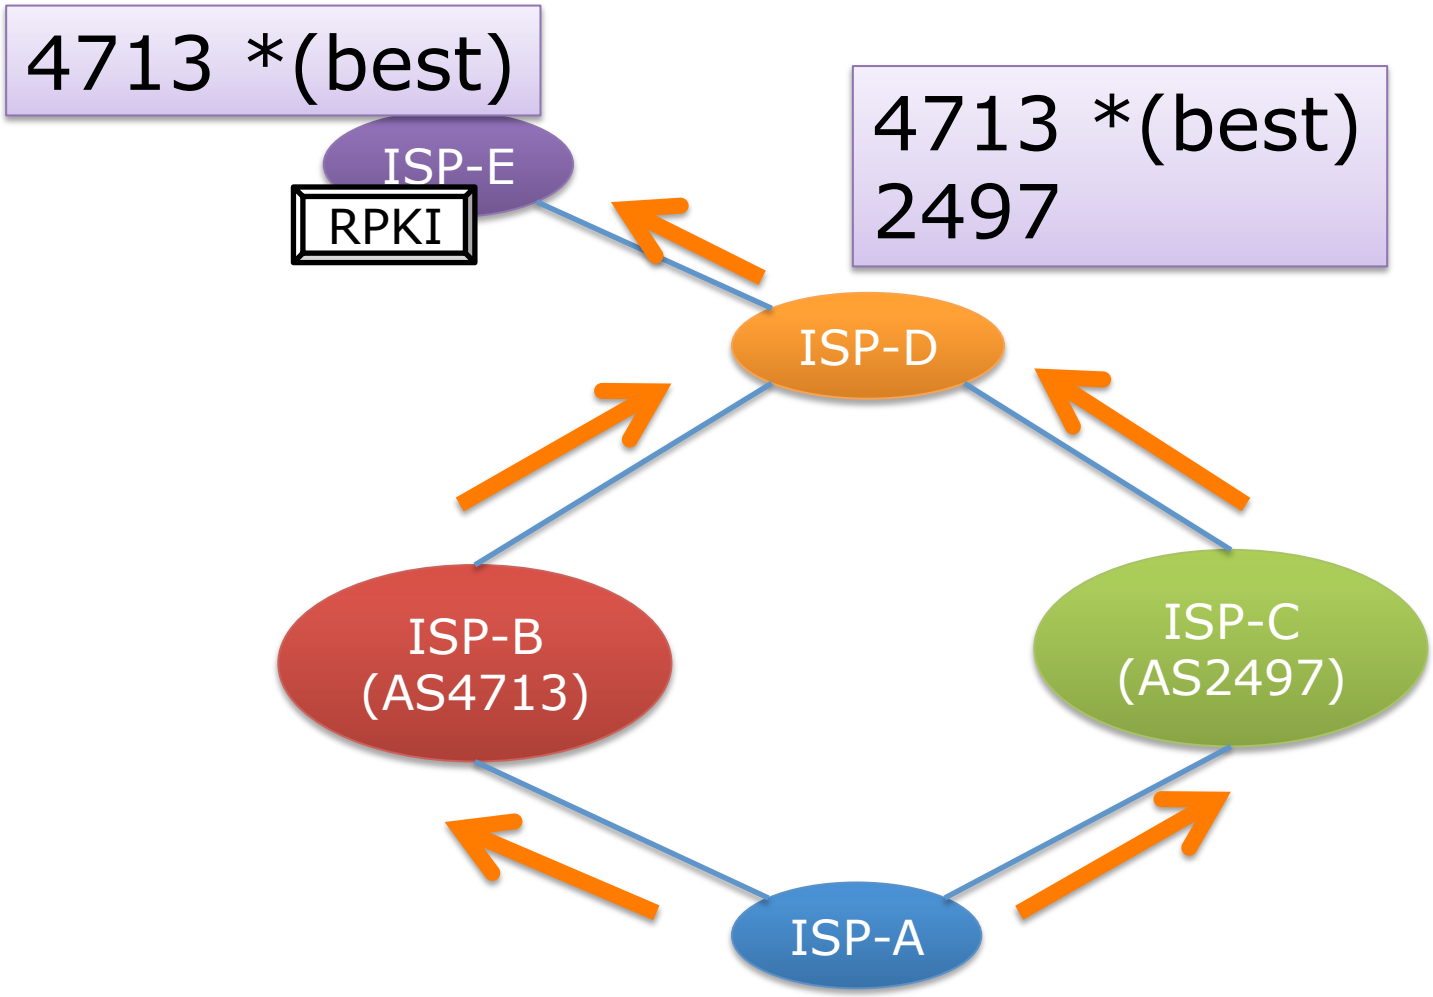

Some case(multiple origin)

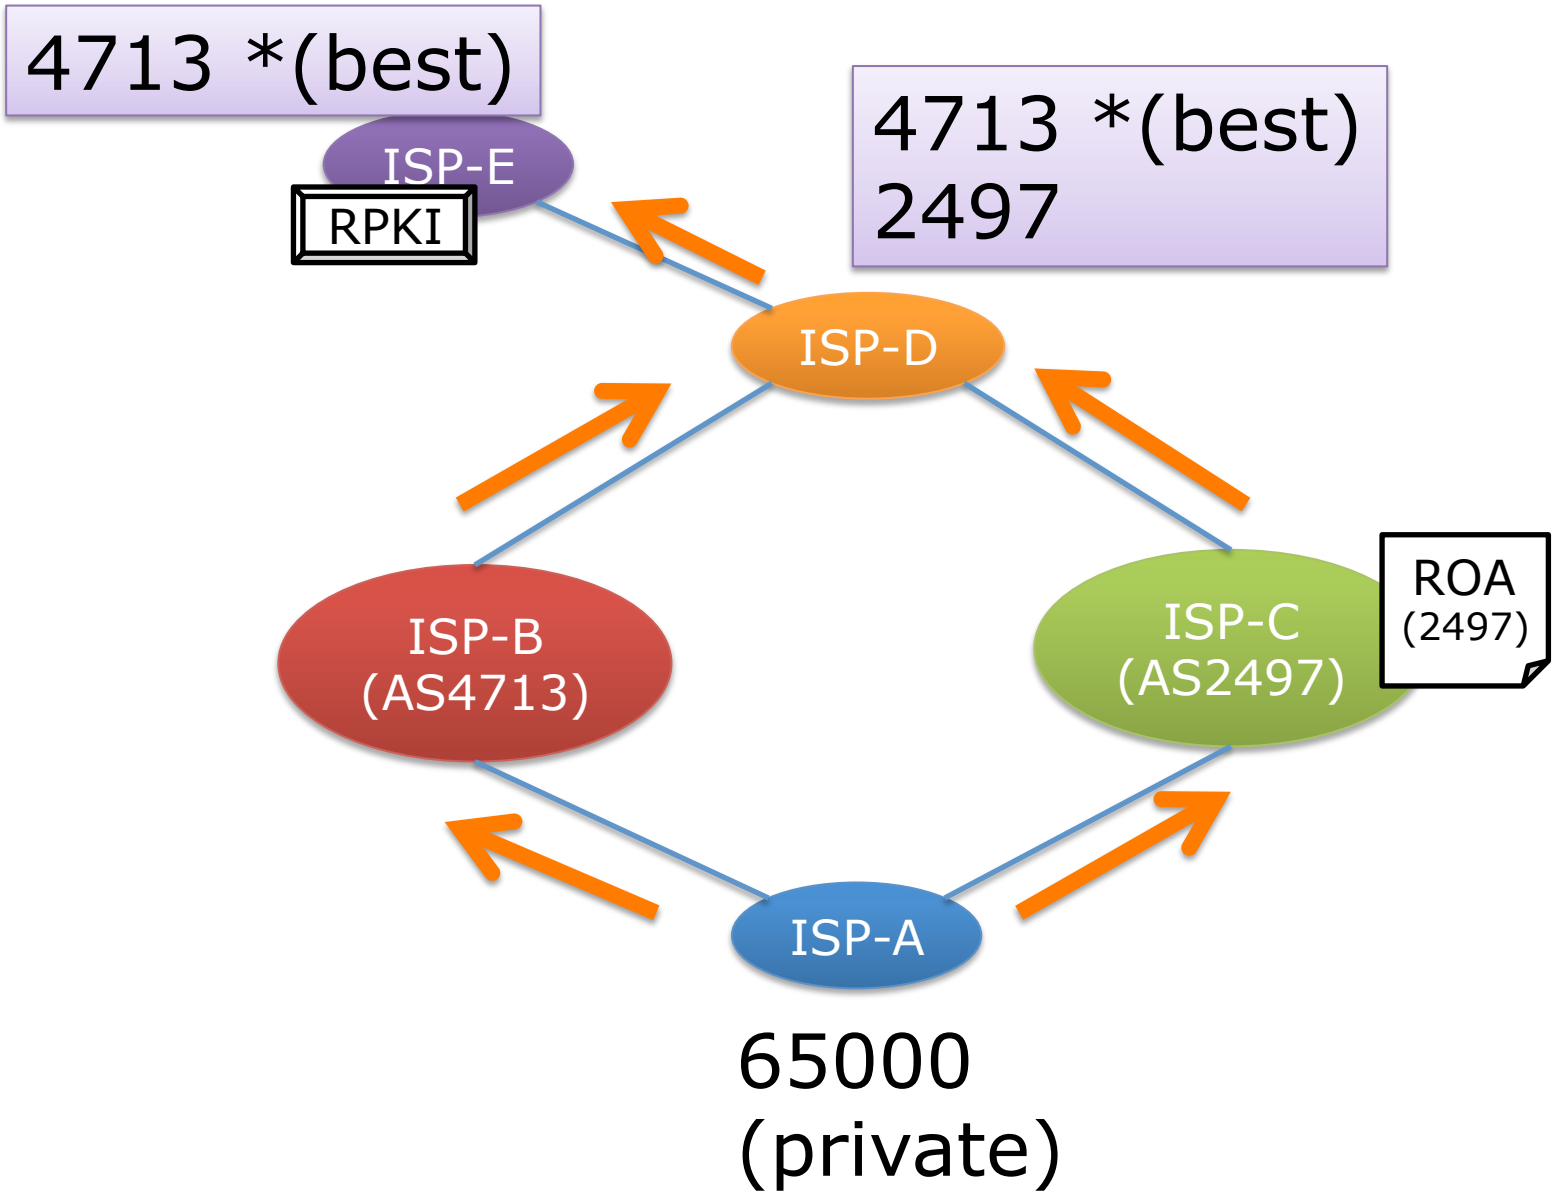

Some case(multiple origin)

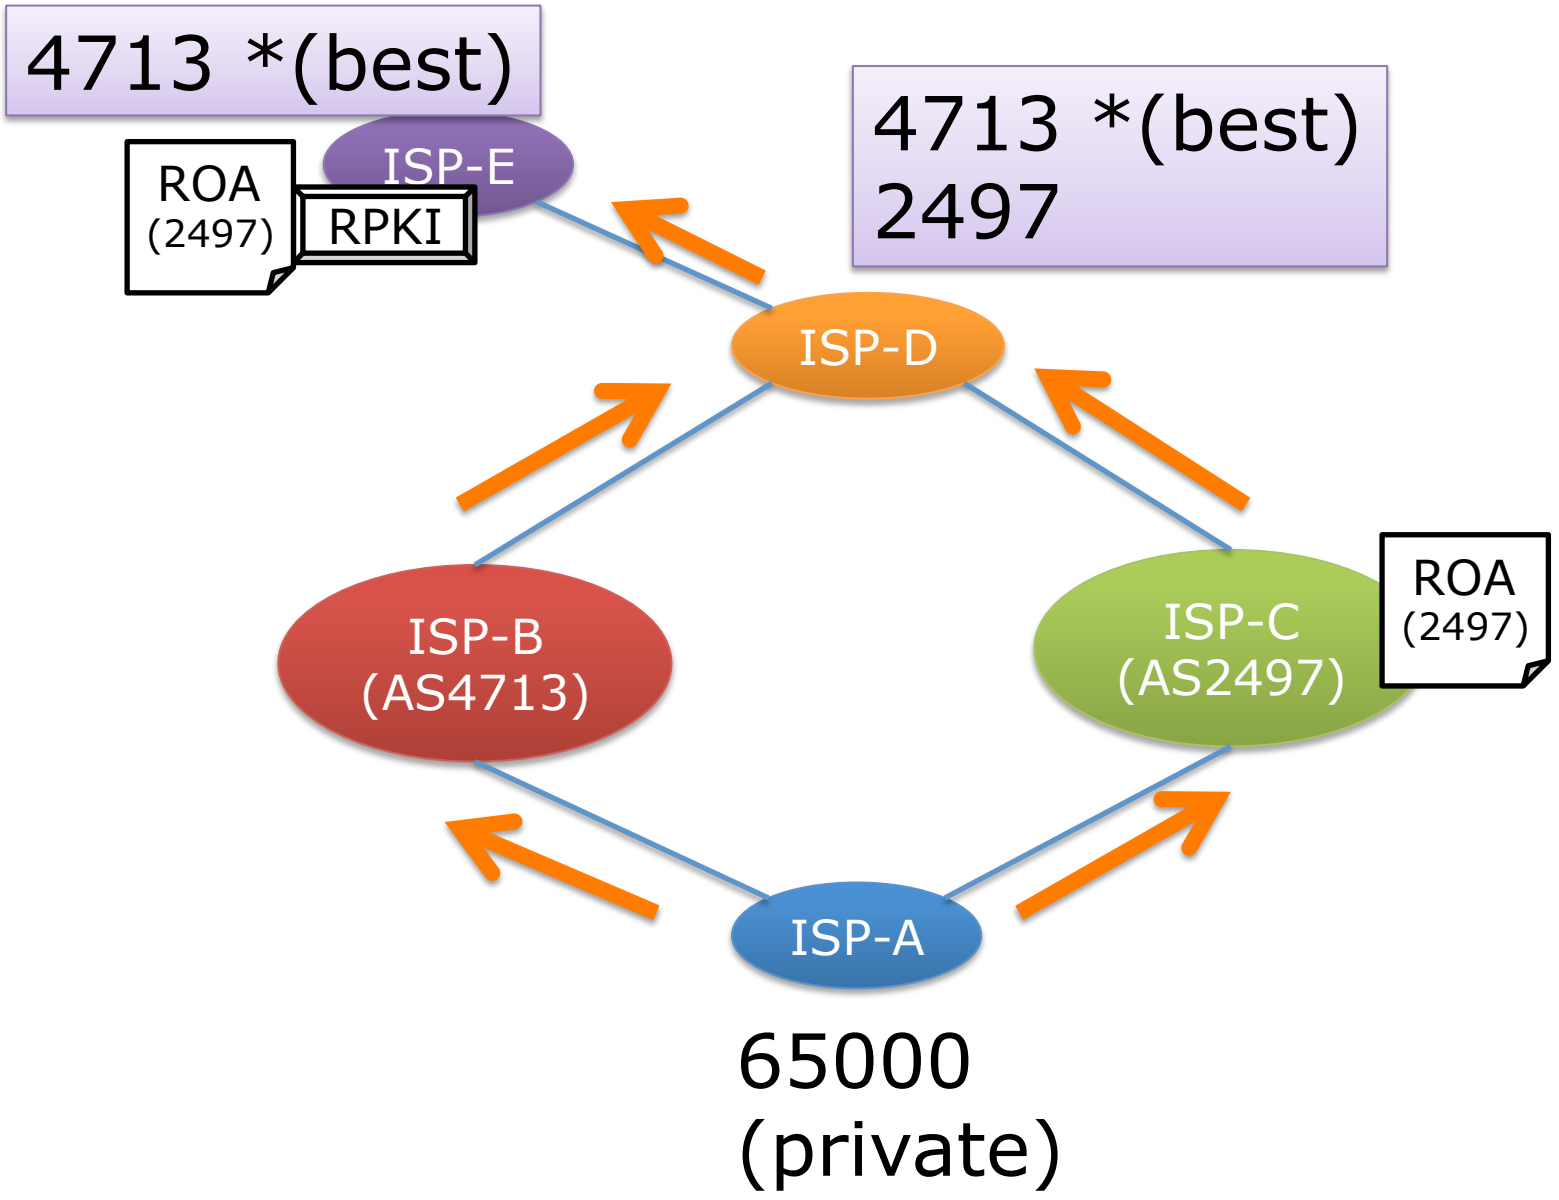

Some case(multiple origin)

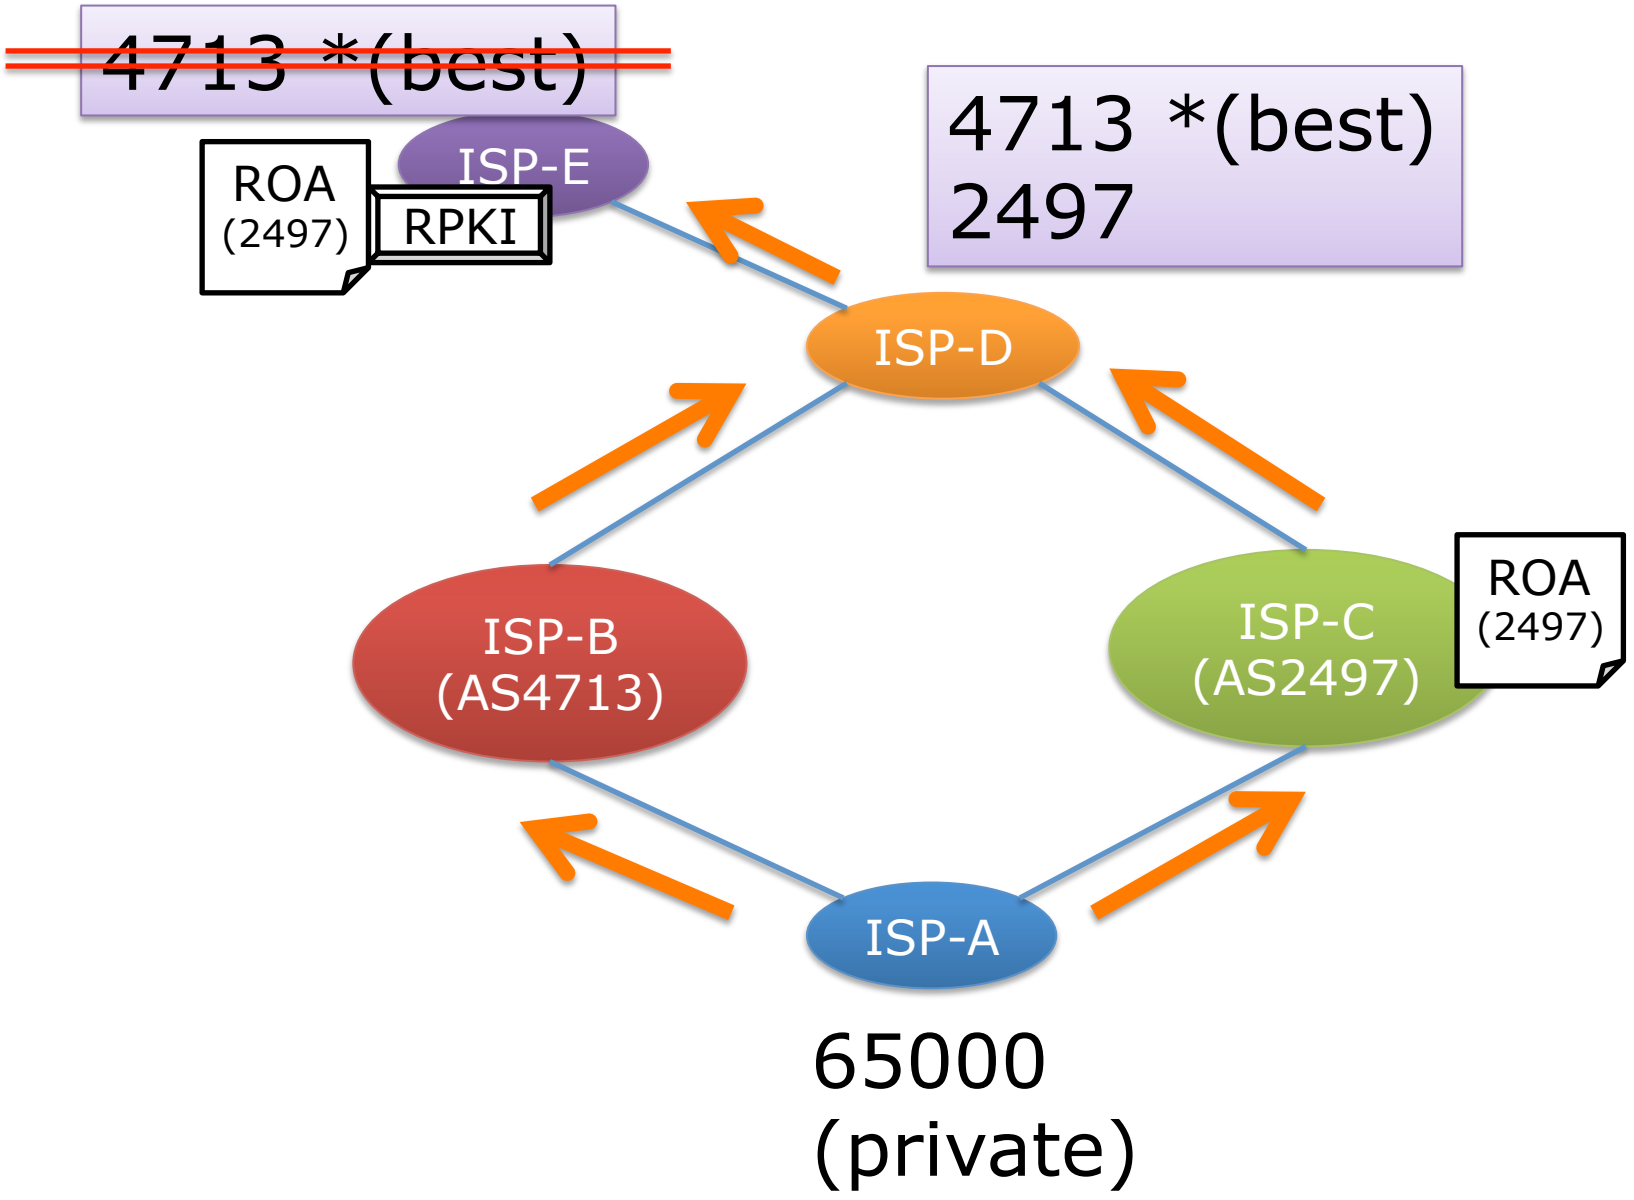

Some case(multiple origin)

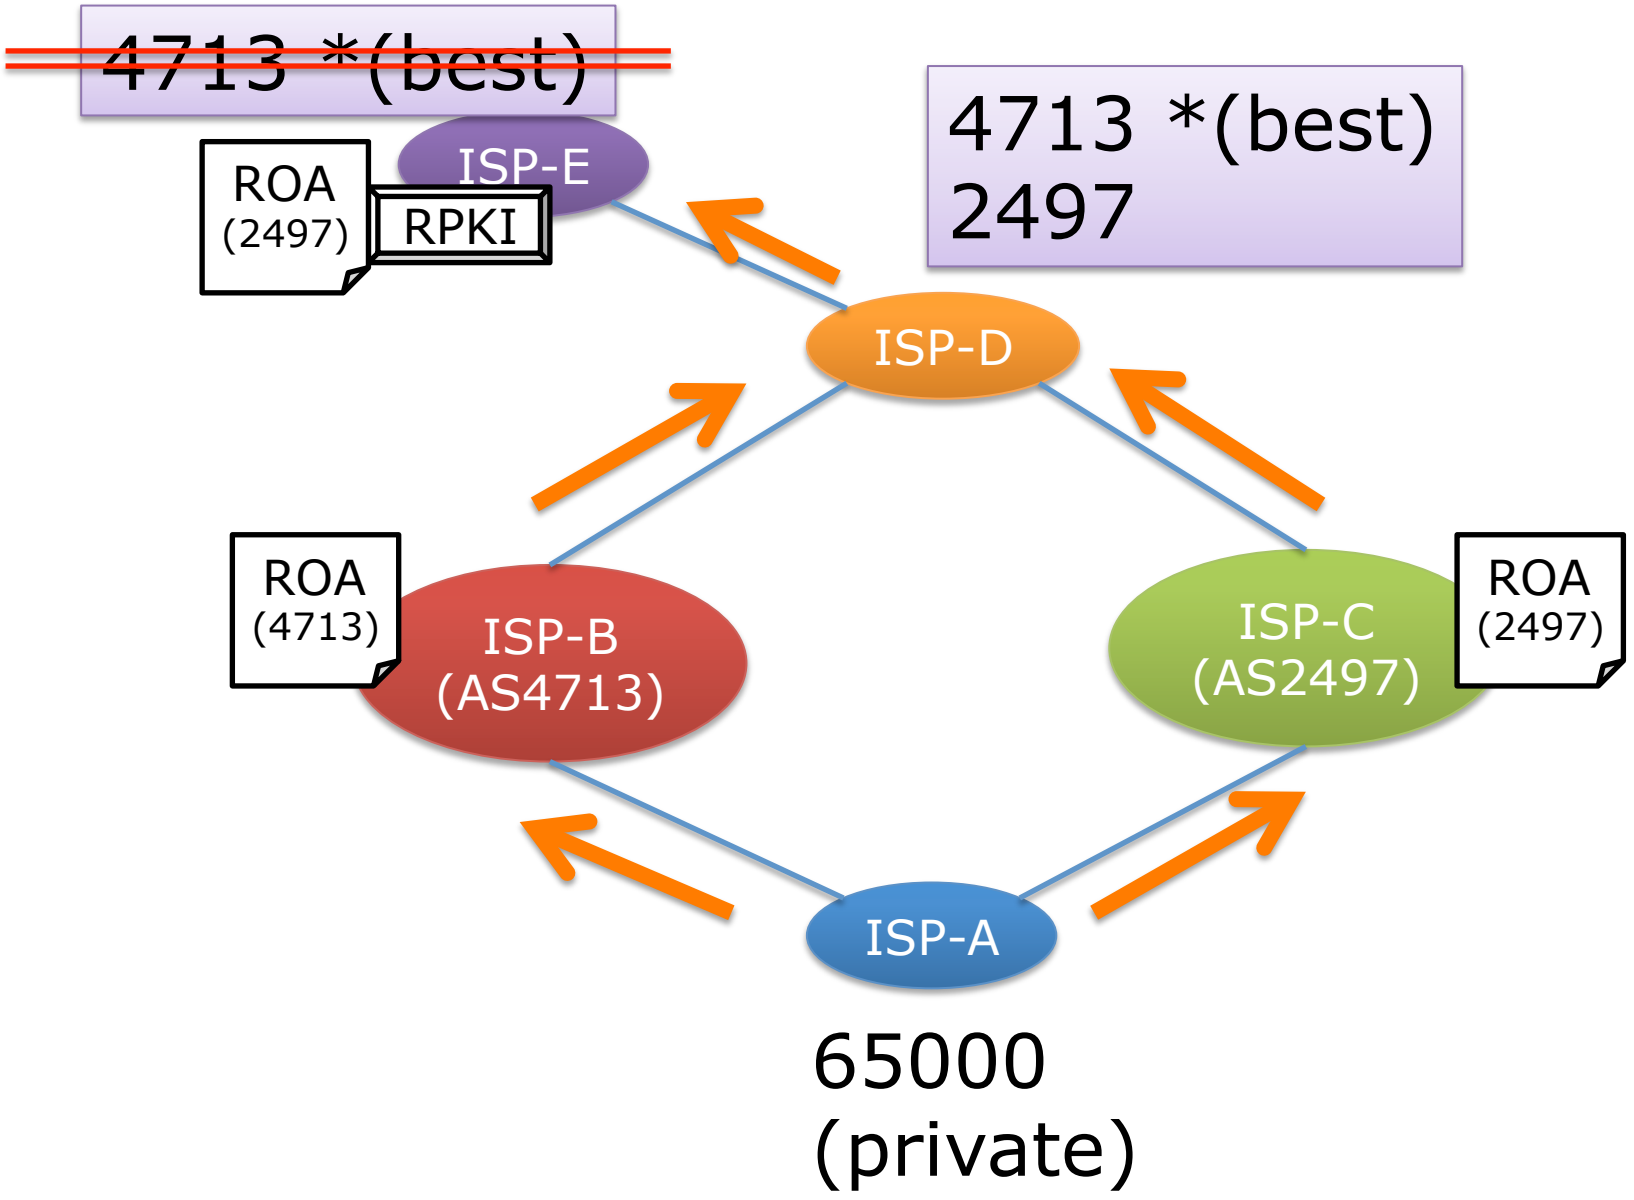

Some case(multiple origin)

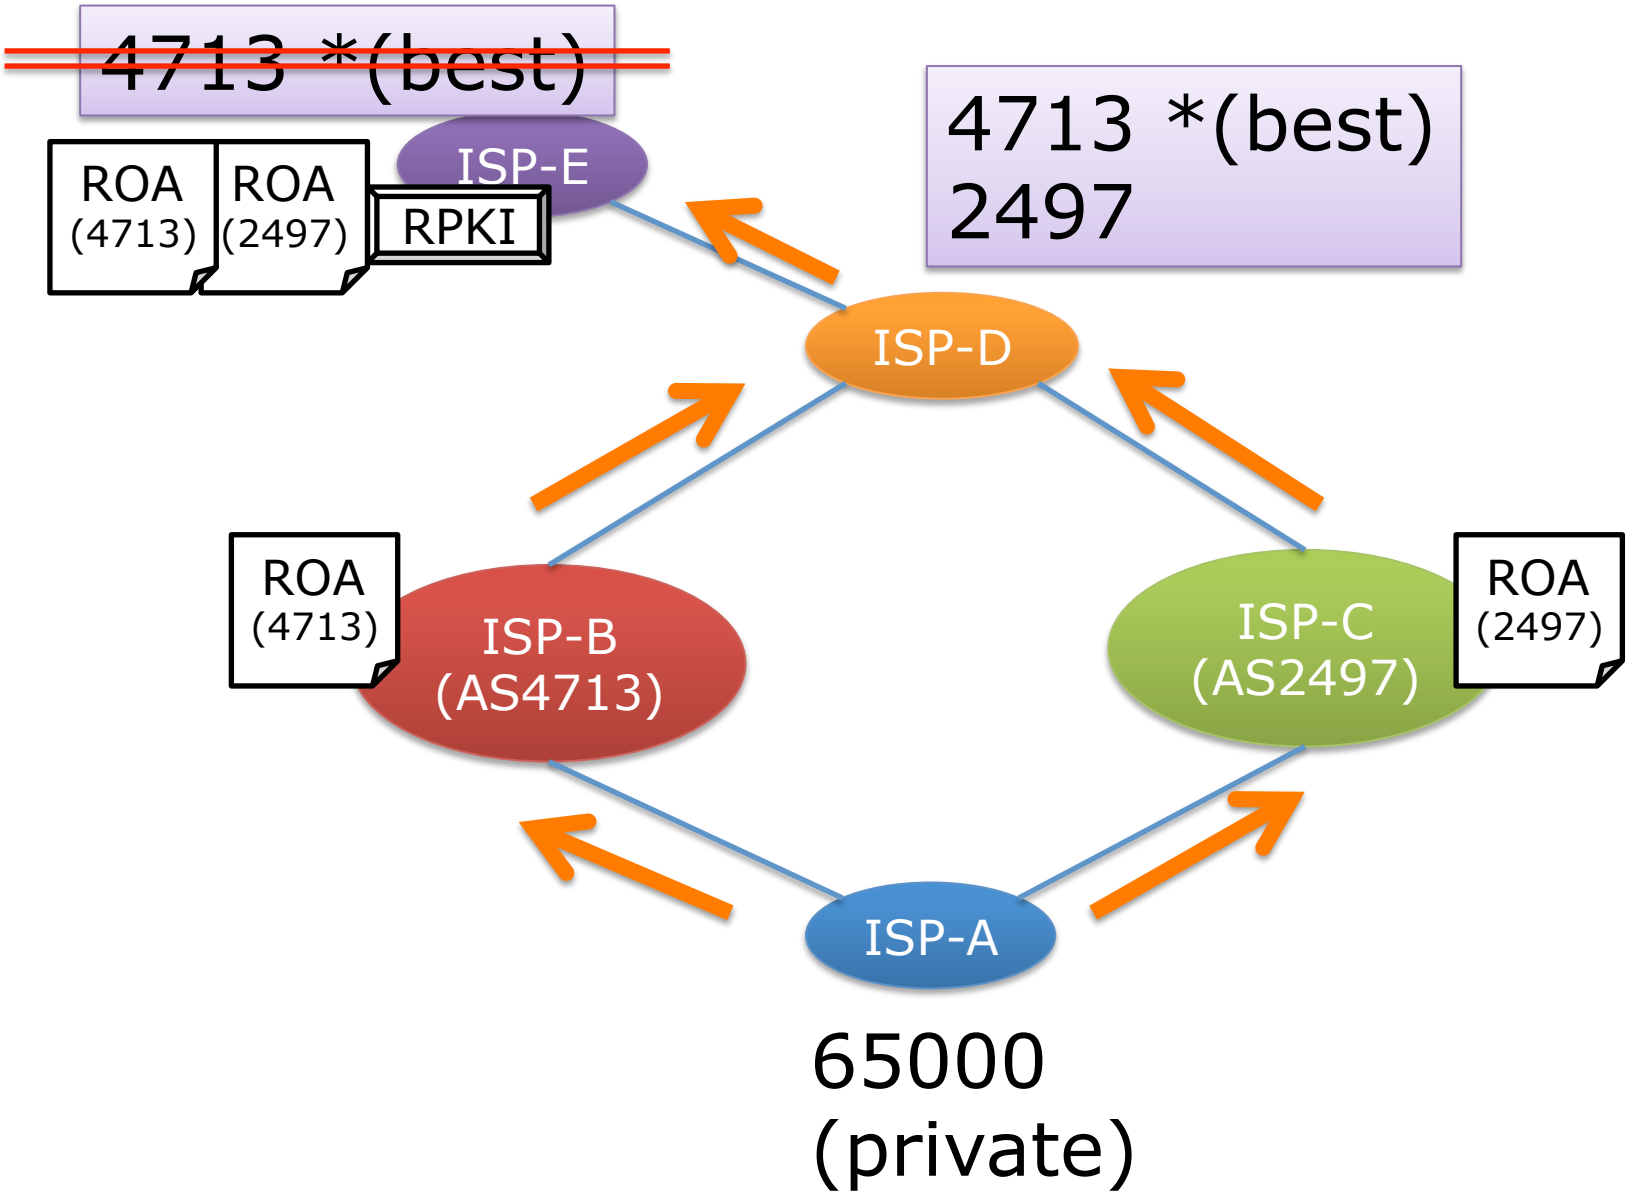

Some case(multiple origin)

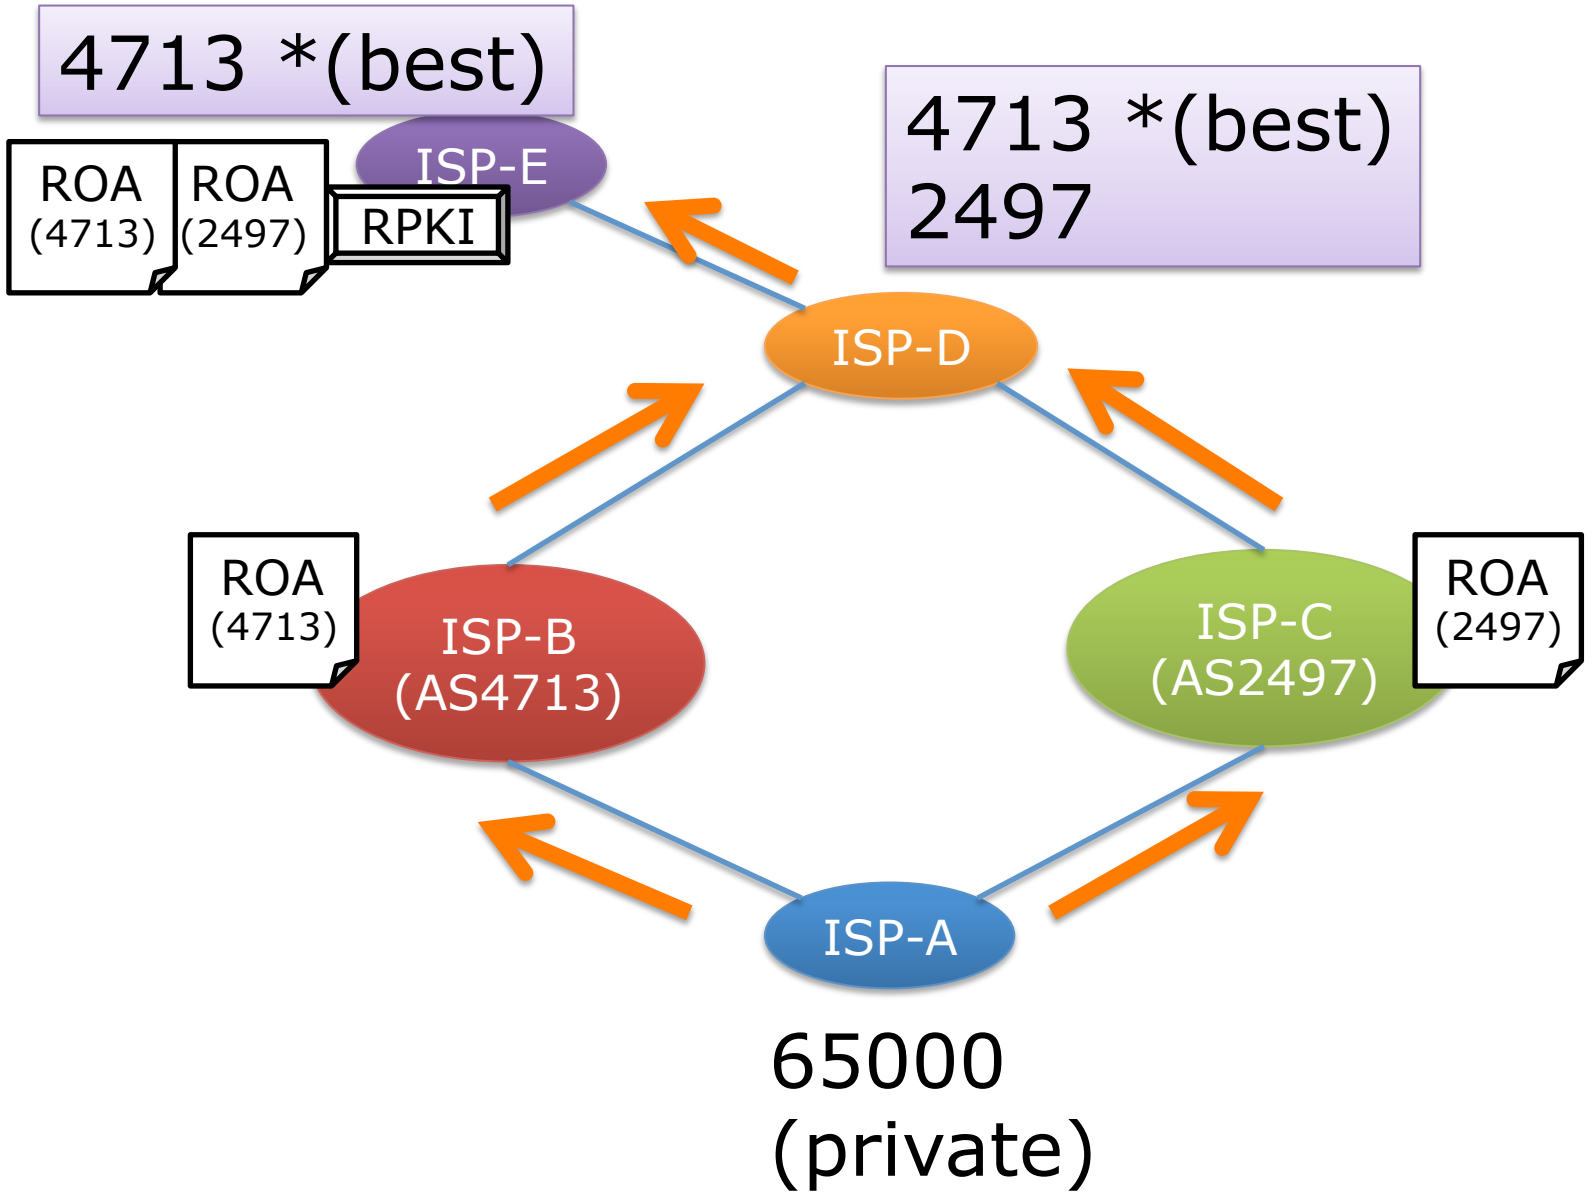

RPKI use case

- Detection (first step)
- Filtering reliable information
- Peer/Transit ISP Routing Control
  - Valid, Invalid, Not Found
- For IX route-server
- Automatically register to IRR(rpki2irr)

Japanese Activity for RPKI

RPKI(1) –Resource cert/roa providing plan–

- A RPKI plan for providing certificate and ROA for LIR has been approved in JPNIC.

2014 Apr – Jun	2014 Jul – Sep	2014 Oct – Dec	2015 Jan – Mar
Basic service design →			
	Deployment on resource management system →		
			* release

- Basic concept of the plan (followings are integrated into JPNIC's deployment issue list)
  - Interoperability
  - Useful for both LIR and network operators
  - Anomaly detection and service redundancy

Copyright © 2014 Japan Network Information Center

Japanese Activity for RPKI

RPKI(2) -IX meeting and workshops-

- A large IX in Japan has been interested in RPKI
  - JPNIC had a brief introduction of RPKI in IX's conference.
  - increased number of users in examining GUI/CA
- Discussions with BGP operators
  - Inter-domain Routing Security Workshop(IRS)
    - http://irs.ietf.to/ (Japanese only)
  - comments about threshold to deploy, rollback-able ROA store, etc

Copyright © 2014 Japan Network Information Center

Japanese Activity for RPKI

RPKI(3) –Experimental GUI/CA–

- Experimental GUI and CA with RPKI Tools
  - For LIR's trial and examination for their deployment
    - RPKI GUI and CA are kept during JPNIC's deployment term
  - 10 users from LIR includes large ISP

- Next step
  - Interoperability between NIR and APNIC
    Please contact us!
    - Email: ca-query at nic.ad.jp or taiji-k at nic.ad.jp

Copyright © 2014 Japan Network Information Center

An element and the present conditions about the RPKI spread

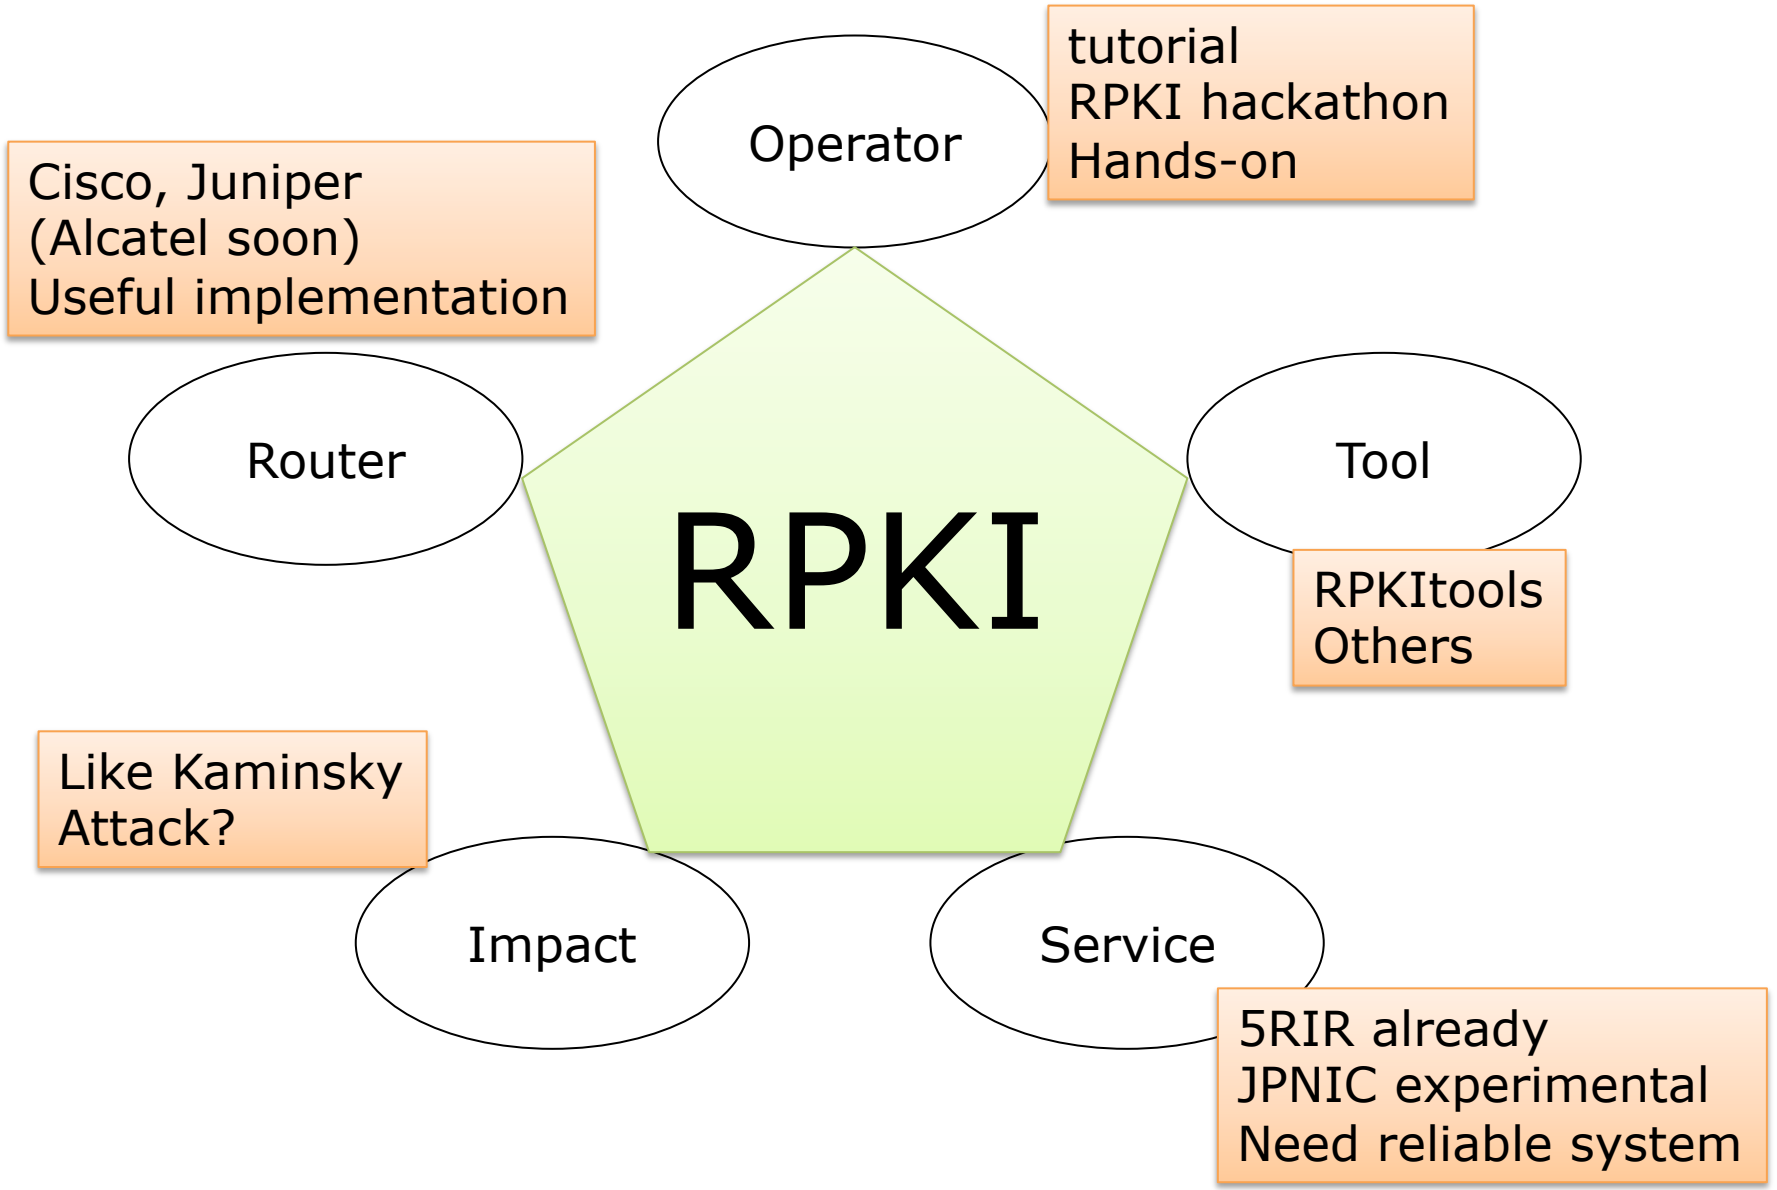

Deployment for RPKI world

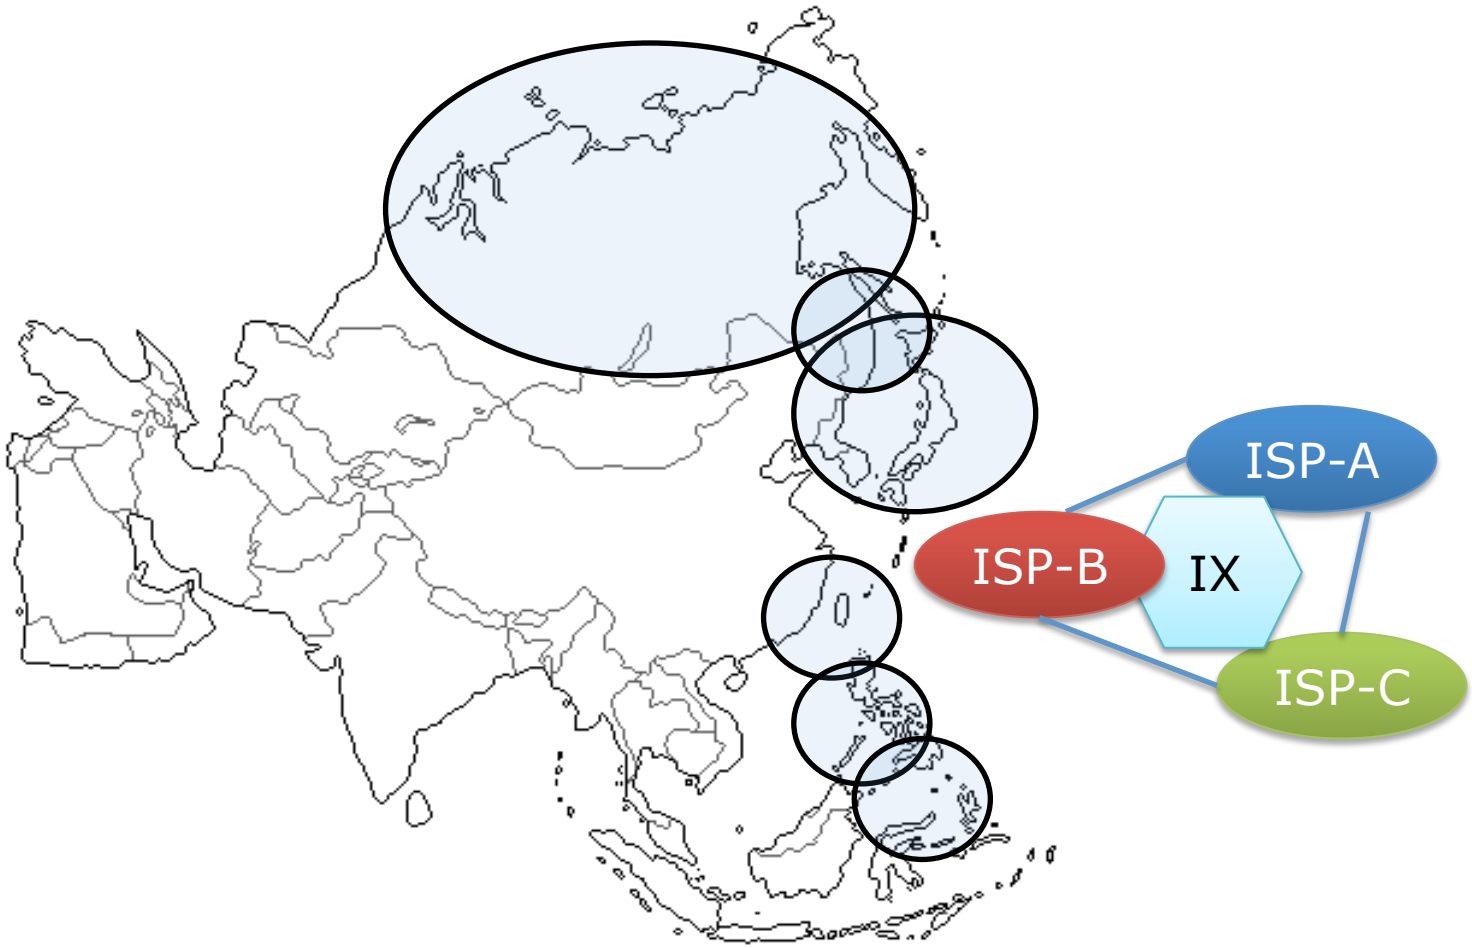

Deployment for RPKI world

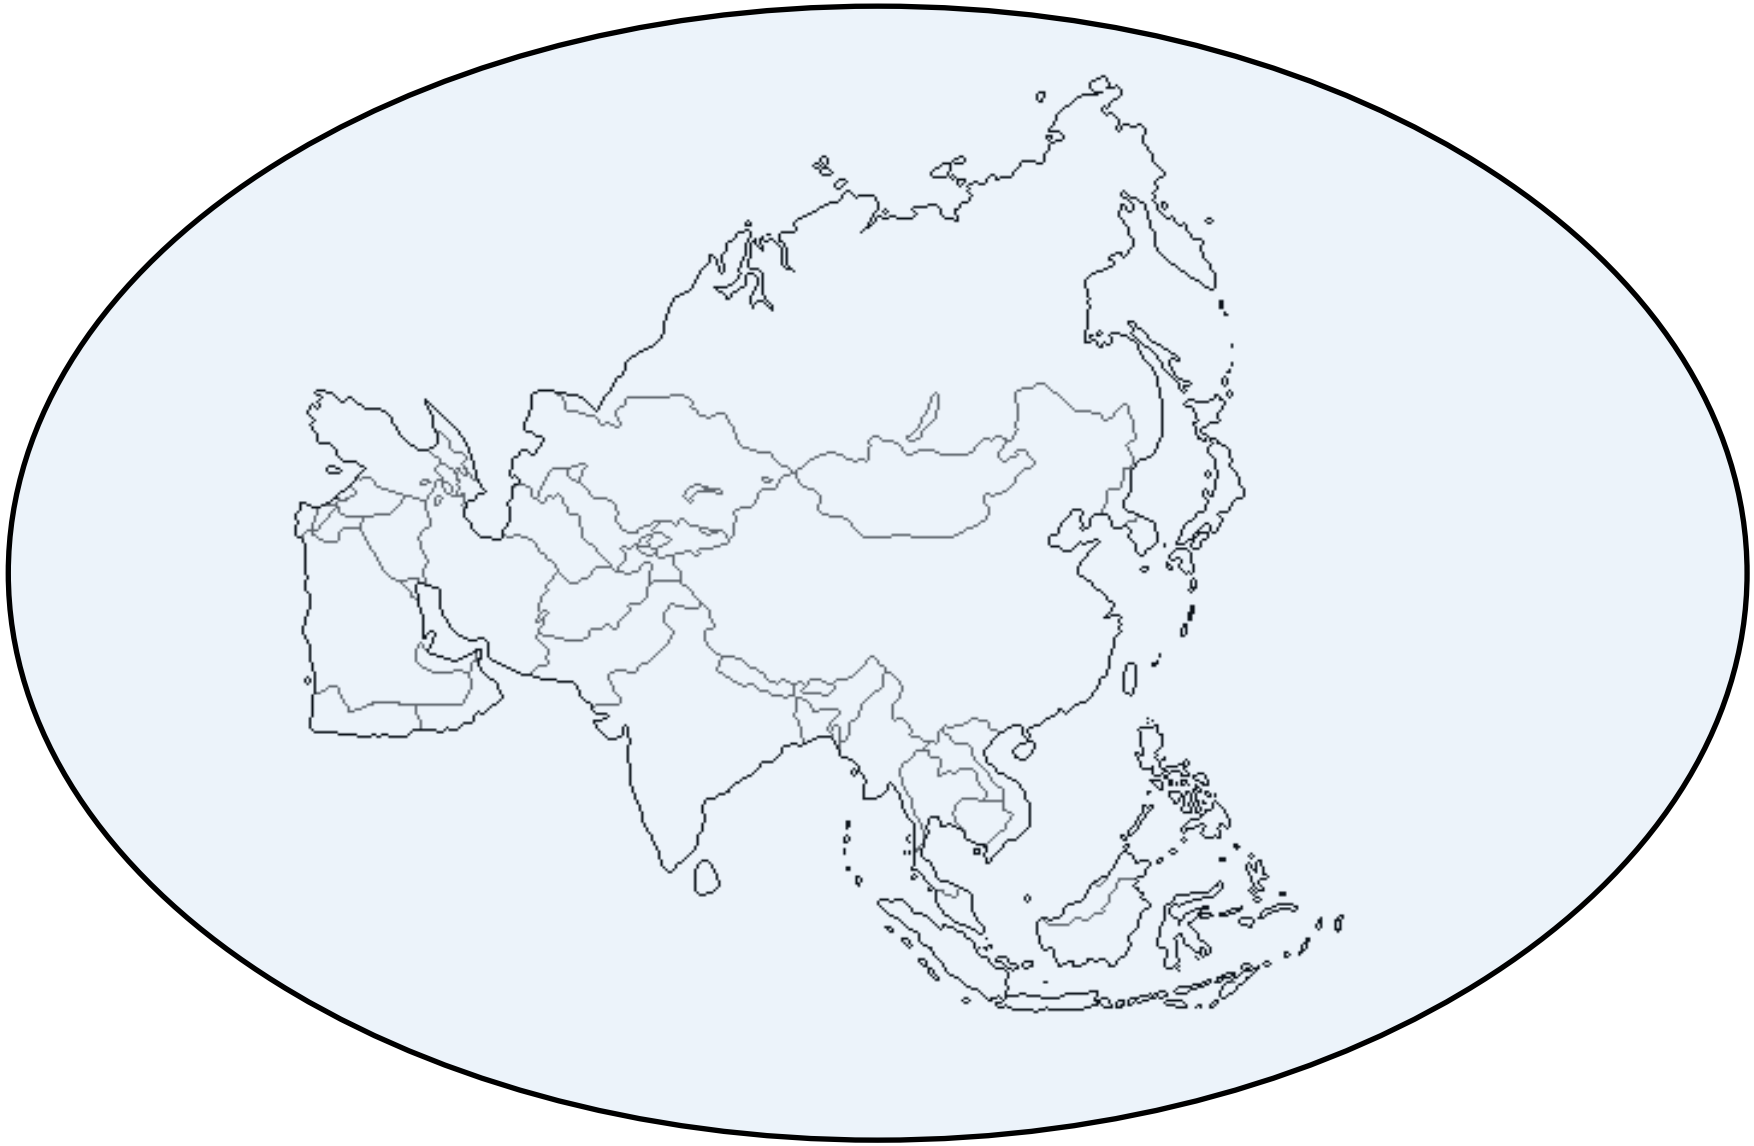

rpki.surfnet.nl

RIR	Total	Valid	Invalid	Unknown	Accuracy	RPKI Adoption Rate
AFRINIC	11709 (100%)	48 (0.41%)	49 (0.42%)	11612 (99.17%)	49.48%	0.83%
APNIC	122347 (100%)	246 (0.2%)	299 (0.24%)	121802 (99.55%)	45.14%	0.45%
ARIN	186568 (100%)	754 (0.4%)	255 (0.14%)	185559 (99.46%)	74.73%	0.54%
LACNIC	64044 (100%)	11239 (17.55%)	1181 (1.84%)	51624 (80.61%)	90.49%	19.39%
RIPE NCC	134875 (100%)	9043 (6.7%)	815 (0.6%)	125017 (92.69%)	91.73%	7.31%

RPKI hackathon, hands-on, tutorial

Almost all people successfully created ROA and experienced origin validation