Chair for Network Architectures and Services – Prof. Carle
Department of Computer Science
TU München

Master Course
Computer Networks
IN2097

Prof. Dr.-Ing. Georg Carle

Chair for Network Architectures and Services
Department of Computer Science
Technische Universität München
http://www.net.in.tum.de

Outline

- Interdomain Routing – cont.
  - BGP Security
  - BGP incidents
  - Prefix hijacking, AS hijacking
  - Early warning

BGP “security” Today – A Rather Sad Topic…

- BGP sessions use TCP
  - No encryption – interceptors can read everything
  - “Authentication”: accept or decline AS number in OPEN message
  - Further authentication (recommended, but optional): TCP-MD5 (RFC 2385), TCP-AO (RFC 5925)
    - TCP-AO (TCP Authentication Option): header option contains cryptographic signature of packet
    - Protects BGP sessions from spoofed TCP segments
    - TCP connections only accepted from peers with accepted signature
    - No protection against against eavesdropping, DoS attacks, …
- Defensive filtering
  - Provider knows prefixes of its (stub) AS customers:
    - Don’t accept updates for other prefixes from them
    - Don’t accept updates with other ASNs from them

BGP Routing security case study 1: How Pakistan Telecom inadvertently hijacked Youtube

- On 25 Feb 2008, users worldwide could not reach YouTube…:
- Pakistan Telecom were ordered by a Pakistani court to block access to a certain YouTube video
- Only feasible choice was to block all YouTube traffic (208.65.152.0/2**2**)
- They created an internal "black hole route" for their network:
  - Manual insertion of a new route for 208.65.152.0/2**4** into IGP
  - Packets sent via that route get discarded at the endpoint
  - Longest prefix match ➔ This route absorbs ¼ of the /22 traffic (in this case: the part containing the servers)
- Unfortunately, this black hole route slipped into eBGP…
  - … so BGP routers world-wide saw the new route and used it
- Quick remedy by Google/YouTube?
  - Announcement of even longer prefixes 208.65.152.0/25 and 208.65.152.128/25

Youtube hijacking: Assessment

- Which security mechanisms could have worked here?
- Authentication?
  - No!
  - Pakistan Telecom is a legitimate BGP speaker
  - Not known for malicious behaviour
- Defensive filtering?
  - Probably not!
  - Pakistan Telecom is not just some tiny stub AS with only one or two prefixes

BGP Routing security case study 2: How a small Czech provider terrorized the world's BGP routers

- On 16 Feb 2009, there was a world-wide surge in BGP updates.
- Small Czech provider SuproNet (AS 47868) wanted to announce their prefix with AS path prepending
- Cisco syntax: […] `as-path prepend 47868 47868 47868`
- …but they used MikroTik routers. Syntax: `bgp-prepend 3`
- 47868 cast into 8 bits: 47868 mod 256 = 252
- Result: AS path of length 252 (=unusually long)
- Path became longer as the announcement travelled through the world… and approached length 256 (=maximum)
- Many Cisco routers could not handle the long AS path and sent out invalid BGP messages
- Result = BGP session resets at their BGP neighbours
  - Remove all BGP routes learned from the crashed router
  - Accordingly, send BGP updates to neighbours

AS path terror: Assessment (1)

- So… who is to blame?
- SuproNet
  - Network administrator principle:
    Thou shalt read the documentation of your router…
  - …especially if it is about BGP
- MikroTik
  - Number was way too large
  - UI design principle:
    Thou shalt do error checking on user input!
    (If a user can enter garbage, he will do it.)
- Cisco
  - Strange input (long AS path) resulted in malformed output
  - Network software design principle:
    - Thou shalt do error checking on network input
    - Error checking on network output is a good idea

AS path terror: Assessment (2)

- Which security mechanisms could have worked here?
- Authentication?
  - No!
  - SuproNet is a legitimate BGP speaker
  - Not known for malicious behaviour
- Defensive filtering?
  - SuproNet just announced their very own prefix
- Tear down a BGP session upon receiving a malformed UPDATE (c.f. RFC 4271)
  - That's exactly what crashed those BGP sessions…

BGP Security: Suggested Mechanisms (1)

- **Origin authentication:** Only ASes that "own" a prefix can announce it
  - Can secure this cryptographically (PKI)
  - Can we outsmart this?
    - Let 10.11.12.0/24, owned by AS23, be the prefix to be hijacked
    - Rogue AS 666 can lie by announcing non-existent paths: Prefix: 10.11.12.0/24, AS path: 666 23

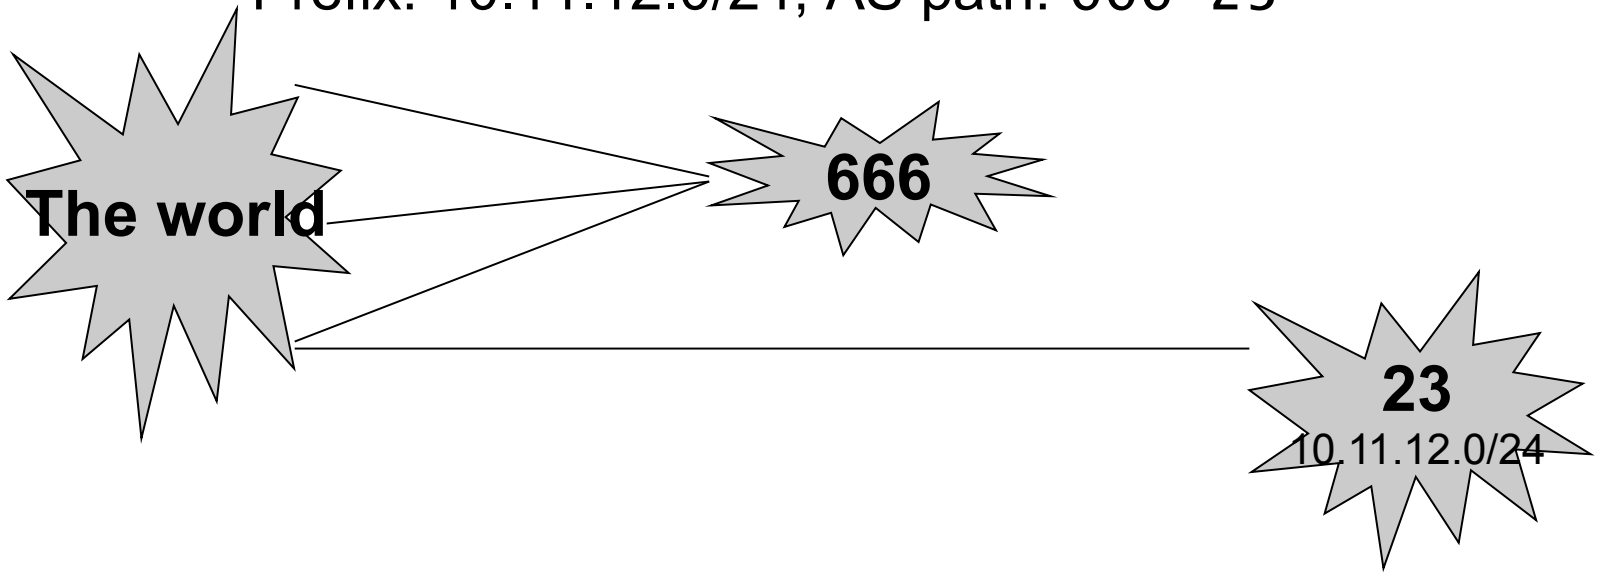

BGP Security: Suggested Mechanisms (2)

- **Secure origin authentication:** Only paths that physically exist can announce it
  - Cryptographically secured path database
  - Can we outsmart this?
    - Can announce paths that we should not see
    - Rogue AS666 knows paths 23–4711 and 4711–666 exist
    - Can announce 66 4711 23, even though it never received an announcement for prefix 10.11.12.0/24 with that path

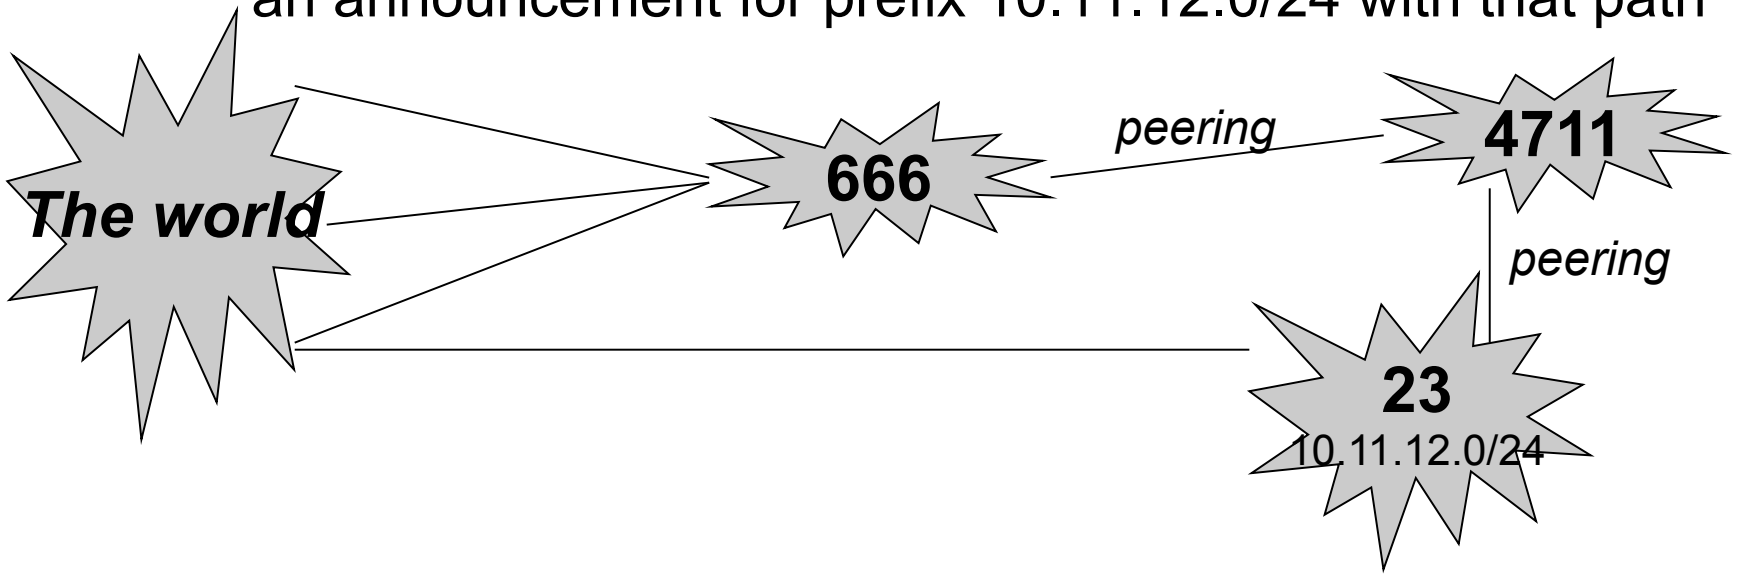

Threat: Targeted Internet Traffic Misdirection

- Credits: Josef Gustafsson
- Source: http://www.renesys.com/2013/11/mitm-internet-hijacking/
- Possible reasons for prefix „hijacking“
  - Fat-finger routing mistake
  - DoS attack
  - MiTM attack: detour traffic, inspect/modify, then forward

GlobalOneBel Incident

- Renesys reports:
  - Belarusian ISP GlobalOneBel
  - Feb 2013: Sequence of events (minutes to hours of duration)
  - Set of victim networks changing daily
  - Affected countries included the US, South Korea, Germany

AS6697 Republican Unitary Telecommunication Enterprise *Beltelecom*

27 February 2013: Traceroute from Guadalajara, Mexico to Washington, DC via Minsk		
IP	**Delay (ms)**	**Notes**
201.151.31.149	15.482	pc-gdl2.alestra.net.mx (Guadalajara, MX)
201.163.102.1	17.702	pc-mty2.alestra.net.mx (Monterrey, MX)
201.151.27.230	13.851	igmty2.alestra.net.mx (Monterrey, MX)
63.218.121.49	17.064	ge3-1.cr02.lar01.pccwbtn.net (Laredo, TX)
63.218.44.78	64.012	TenGE11-1.br03.ash01.pccwbtn.net (Ashburn, VA)
64.209.109.221	84.529	GBLX-US-REGIONAL (Washington, DC)
67.17.72.21	157.641	lag1.ar9.LON3.gblx.net (London, UK)
208.178.194.170	143.344	cjs-company-transtelecom.ethernet8-4.ar9.lon3.gblx.net (London, UK)
217.150.62.234	212.869	mskn01.transtelecom.net (Moscow, RU)
217.150.62.233	228.461	BelTelecom-gw.transtelecom.net (Minsk, Belarus)
87.245.233.198	225.516	ae6-3.RT.IRX.FKT.DE.retn.net (Frankfurt, DE)
*		no response
*		no response
129.250.3.180	230.887	ae-3.r23.nycmny01.us.bb.gin.ntt.net (New York, NY)
129.250.4.69	232.959	ae-1.r05.nycmny01.us.bb.gin.ntt.net (New York, NY)
129.250.8.158	248.685	ae-0.centurylink.nycmny01.us.bb.gin.ntt.net (New York, NY)
*		no response
63.234.113.110	238.111	63-234-113-110.dia.static.qwest.net (Washington, DC)

- Level3 (previously Global Crossing)
  - advertising a false Belarus route
  - having heard it from Russia's TransTelecom
  - who heard it from their customer, Belarus Telecom who has .customer GlobalOneBel

Icelandic Incident

- Icelandic provider Opin Kerfi (AS48685)
  - normally originates 3 prefixes
  - has no downstream AS customers
  - at 07:36:36 UTC on July 31st 2013, announcing origination routes for 597 IP networks owned by US VoIP provider
  - Faulty routes by Opin Kerfi propagated though Siminn

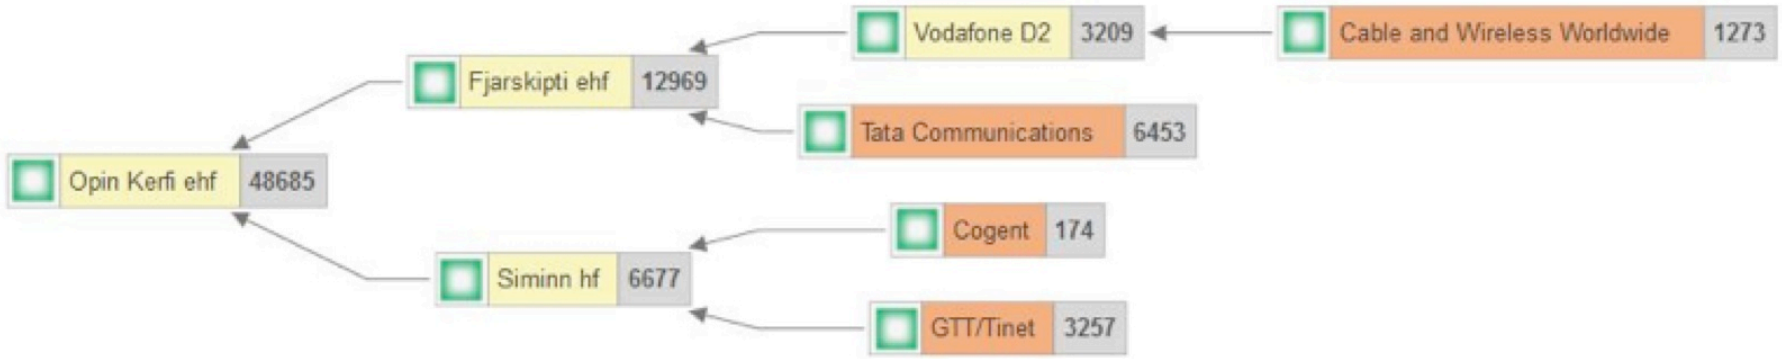

Assessment

- Assessment by Renesys
  http://www.renesys.com/2013/11/mitm-internet-hijacking/
  - „What's not known is the exact mechanism, motivation, or actors.
    We first contacted the peering team at Iceland's Síminn in July, when their traffic redirection began in earnest, highlighting some of the erroneous routes. We received no response.
    We contacted them again recently while researching this story. We were told that the problems were the result of a bug in vendor software, that the problem had gone away when patched, and that they did not believe this problem had a malicious origin. Despite repeated requests for supporting details, we received no further communication.“

Security Extensions to BGP

- Several security extensions to BGP have been proposed
  - S-BGP, psBGP, soBGP, IRV, ...
  - Resource PKI (RPKI) – c.f. subsequent slide
  - Securing BGP operation is a complex task, requiring
    - Designing a protocol with security properties
    - Agreement/adoption by many stakeholders
    - Security policies (which reaction in which case)
    - Change of operational practices

S-BGP

- Stephen Kent, Charles Lynn, and Karen Seo:
  Secure Border Gateway Protocol (S-BGP)
  IEEE JSAC, April 2000

- Secure origin authentication
- Additional attribute allows to sign a route step-by-step
- IPsec protects updates
- Can we outsmart this?
  - Rogue AS666 can still announce a “good” route but then actually use a “bad” route – or even drop the traffic

RPKI

- Resource Public Key Infrastructure (RPKI)
  - public key infrastructure (PKI) framework to secure the Internet's routing infrastructure
  - IETF WG: Secure Inter-Domain Routing (sidr)
  - Vulnerabilities addressed
    - is an Autonomous System (AS) authorized to originate an IP prefix
    - is the AS-Path represented in the route the same as the path through which the NLRI (Network Layer Reachability Information) traveled

BGP security: Further reading

- Renesys blog:
  - Posts with ‘security’ tag: www.renesys.com/blog/security/
  - Entry “Reckless driving on the Internet”
  - Entry “Longer is not always better”
  - Entry “Pakistan hijacks YouTube”

- Butler, Farley, McDaniel, Rexford:
  A survey of BGP security issues and solutions
  Proceedings of the IEEE, January 2010

- Goldberg, Schapira, Hummon, Rexford:
  How secure are secure interdomain routing protocols?
  Proceedings of ACM SIGCOMM, August 2010

Hijacking of Autonomous Systems

Johann Schlamp

How to Prevent AS Hijacking

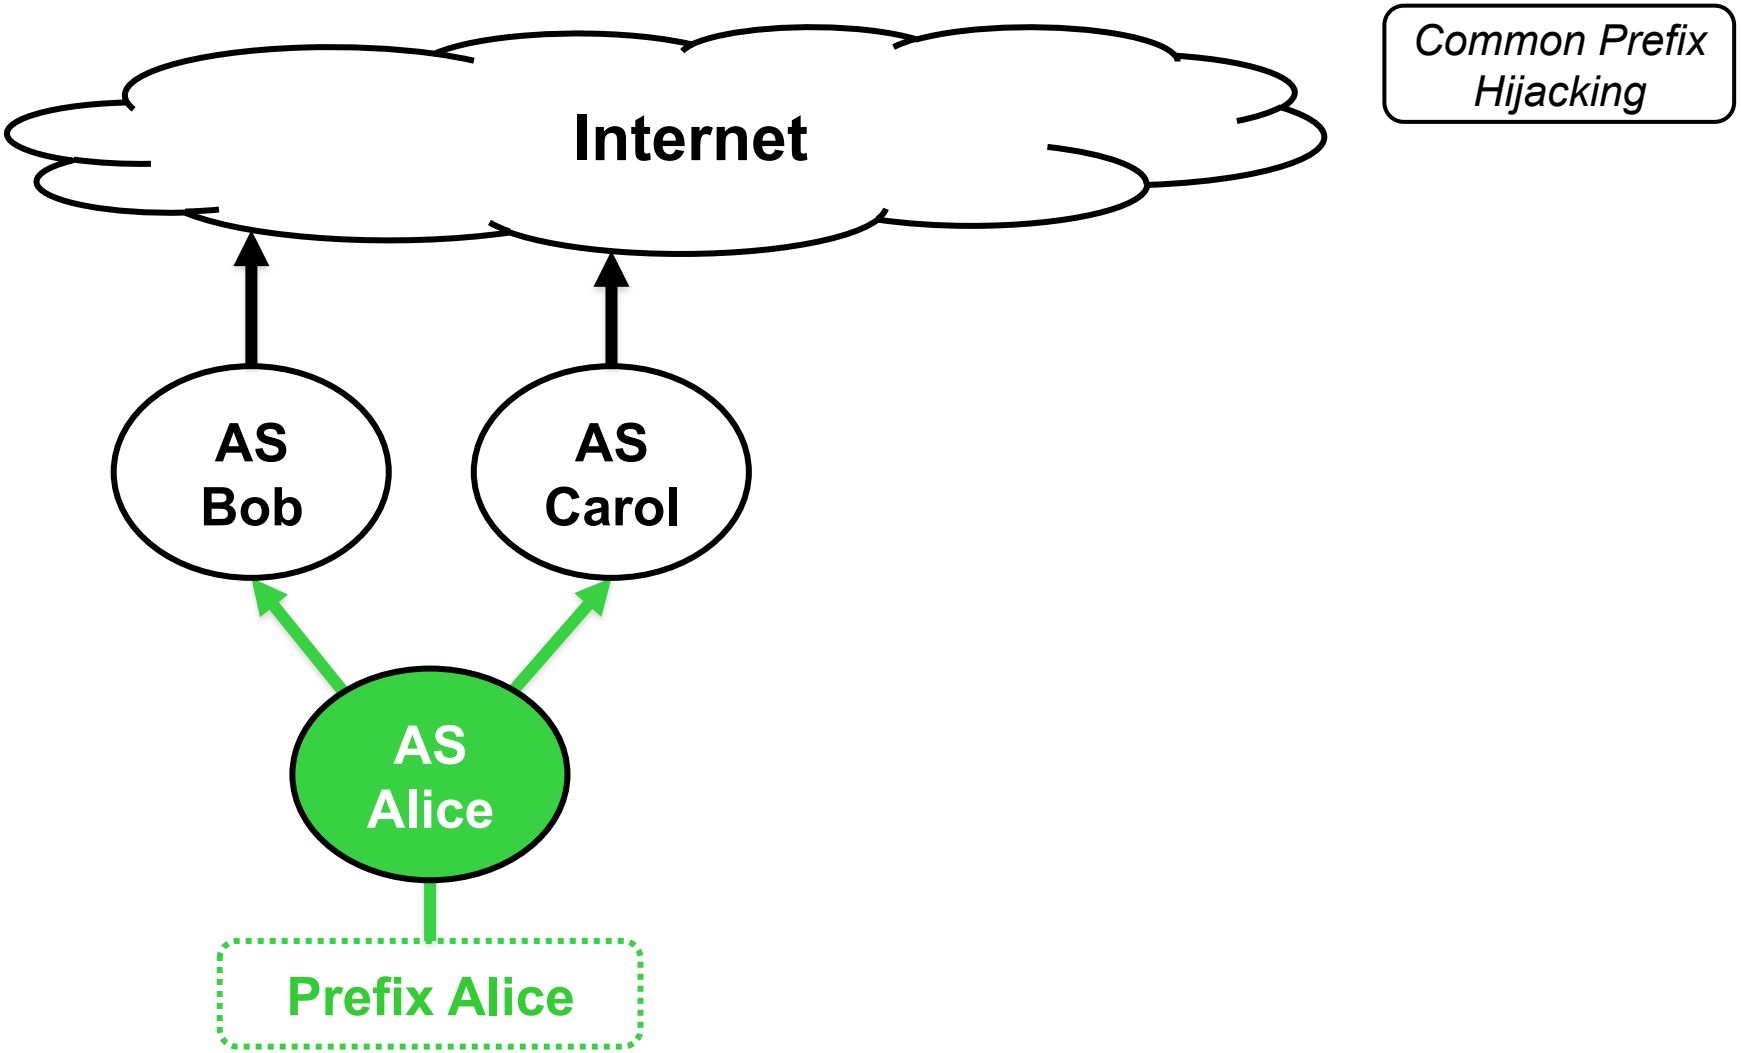

Johann Schlamp, Georg Carle, and Ernst W. Biersack. **How to prevent AS hijacking attacks**. In *Proceedings of the 2012 ACM Conference on CoNEXT Student Workshop (CoNEXT Student 2012)*, Nice, France, December 2012.

How to Prevent AS Hijacking

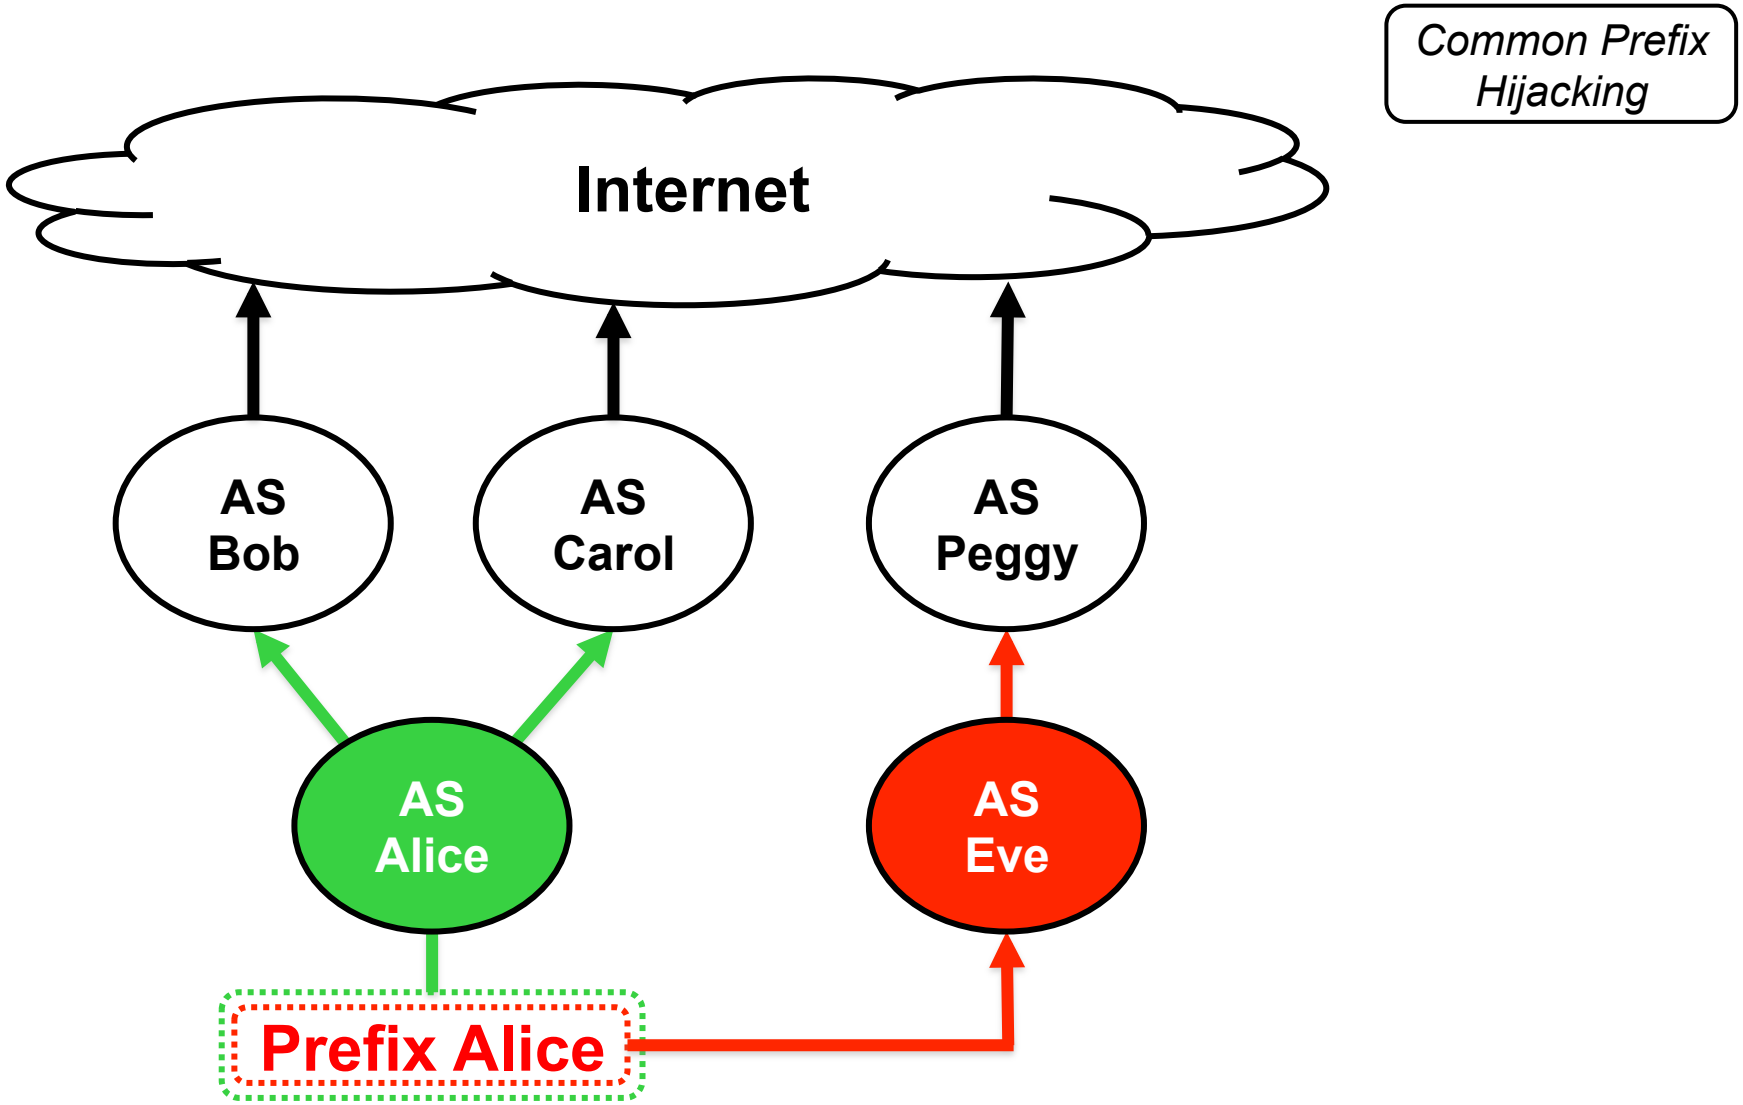

How to Prevent AS Hijacking

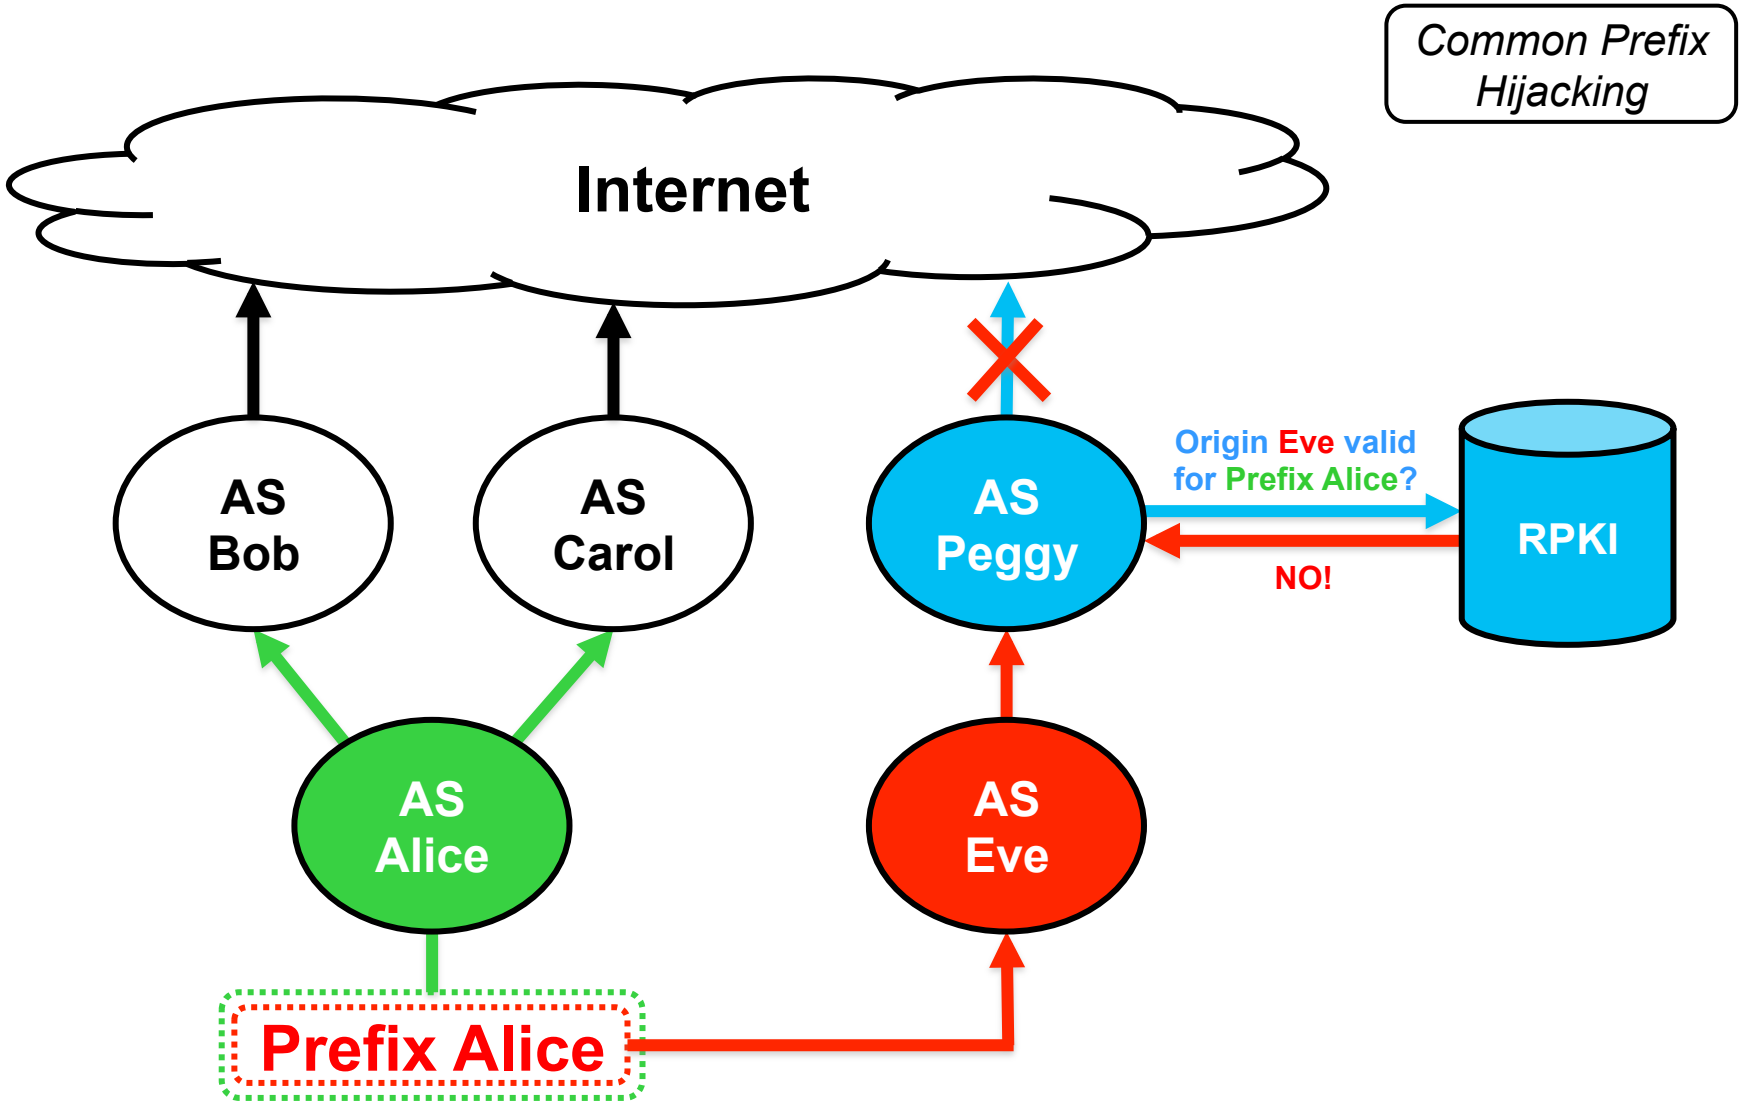

Prefix Hijacking and AS Hijacking

- Prefix hijacking
  - Hijacked prefixes originate from both the victim's AS and the attacker's AS, which is called a multi-origin AS (MOAS)
- AS hijacking

Attacks add another upstream link to the victim's AS

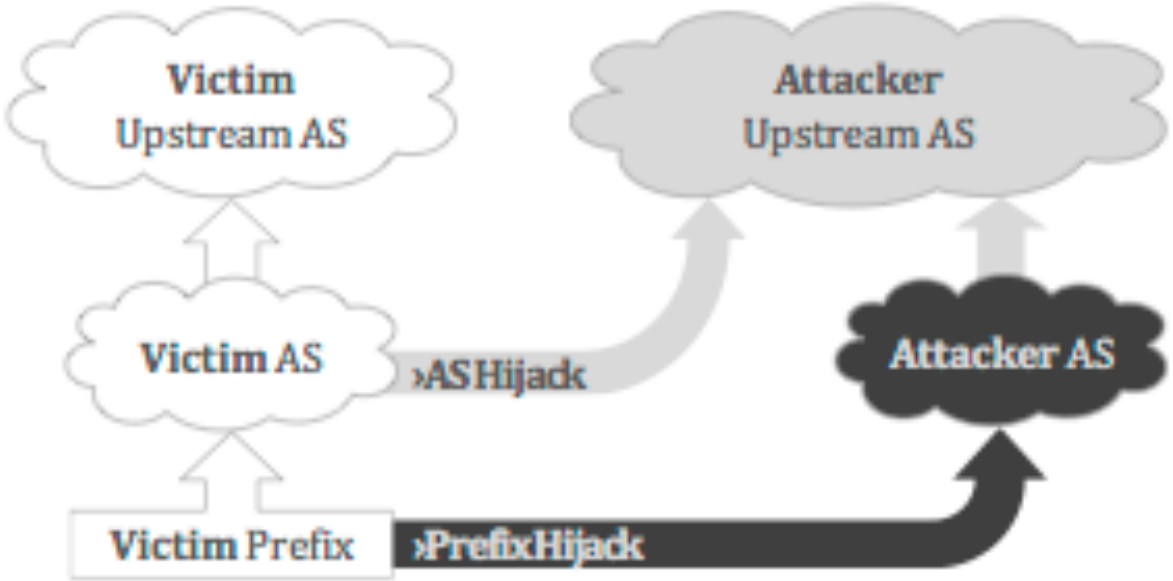

Johann Schlamp, Georg Carle, and Ernst W. Biersack. **A Forensic Case Study on AS Hijacking: The Attacker's Perspective**. *ACM Computer Communication Review (CCR)*, April 2013

Example Scenario before AS Hijacking

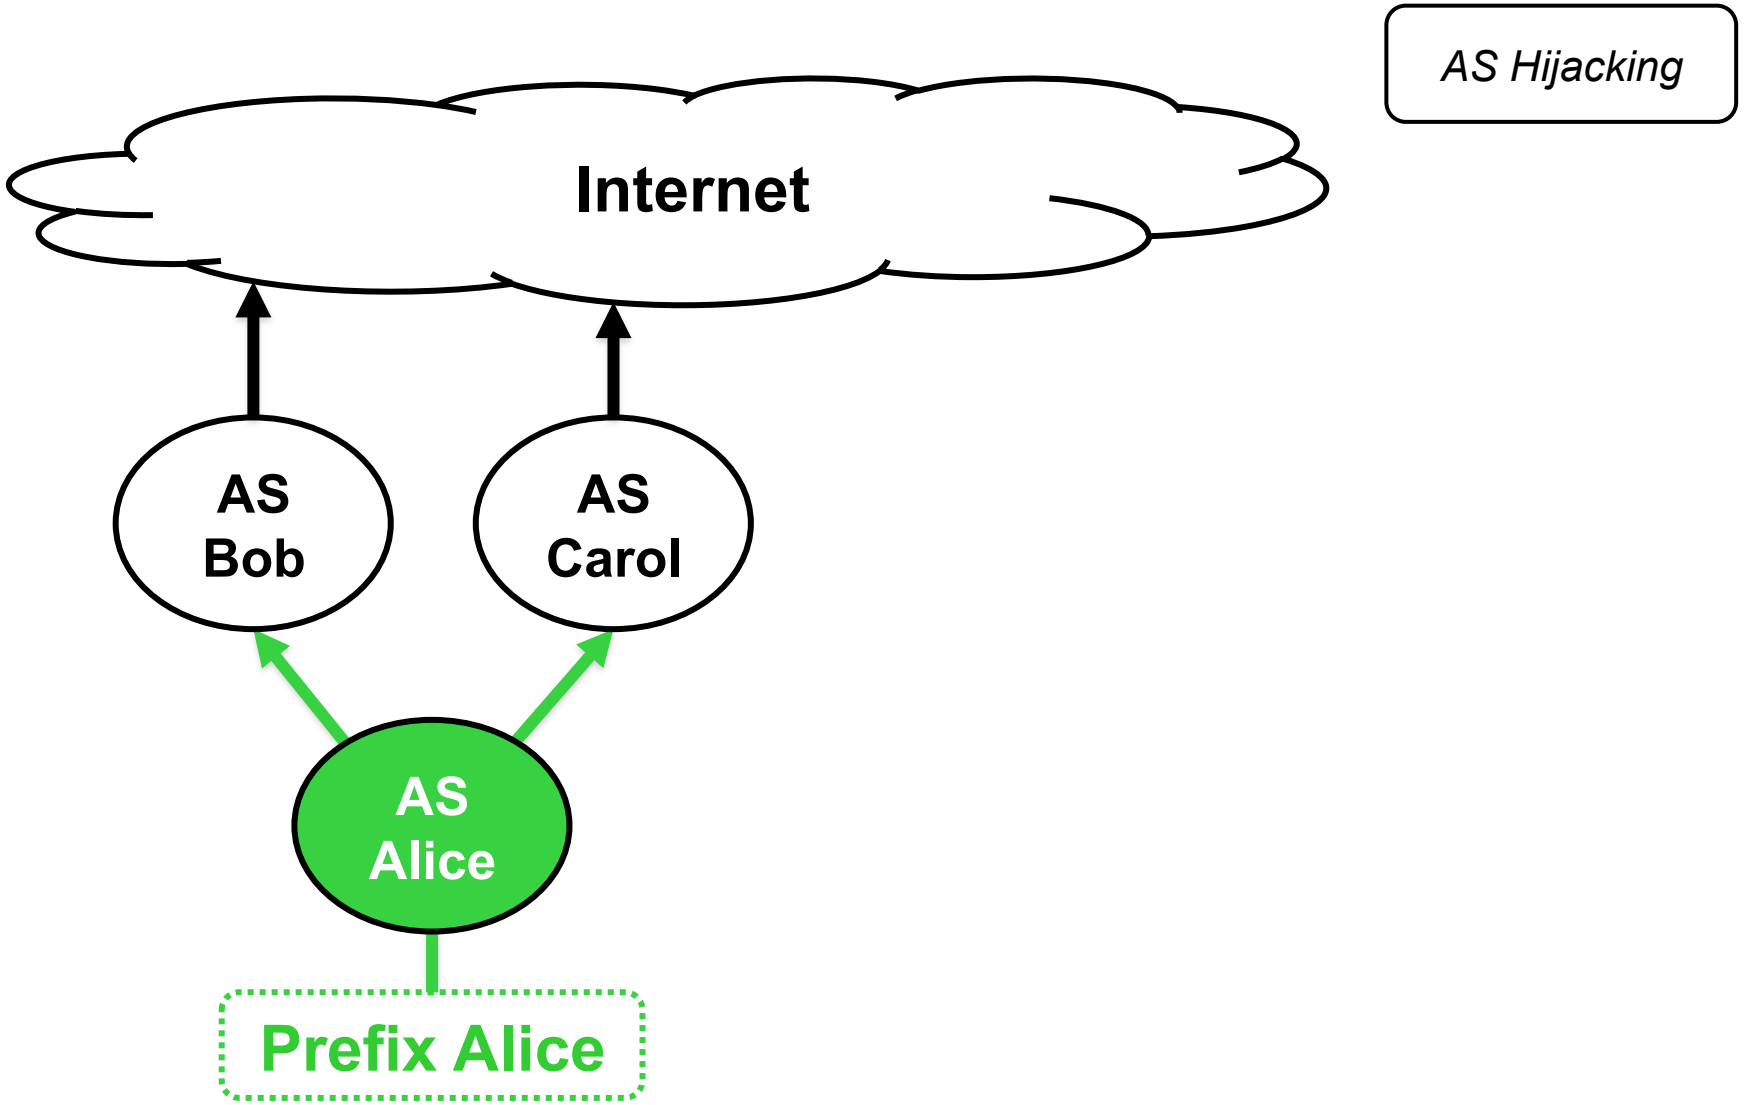

AS Hijacking Scenario

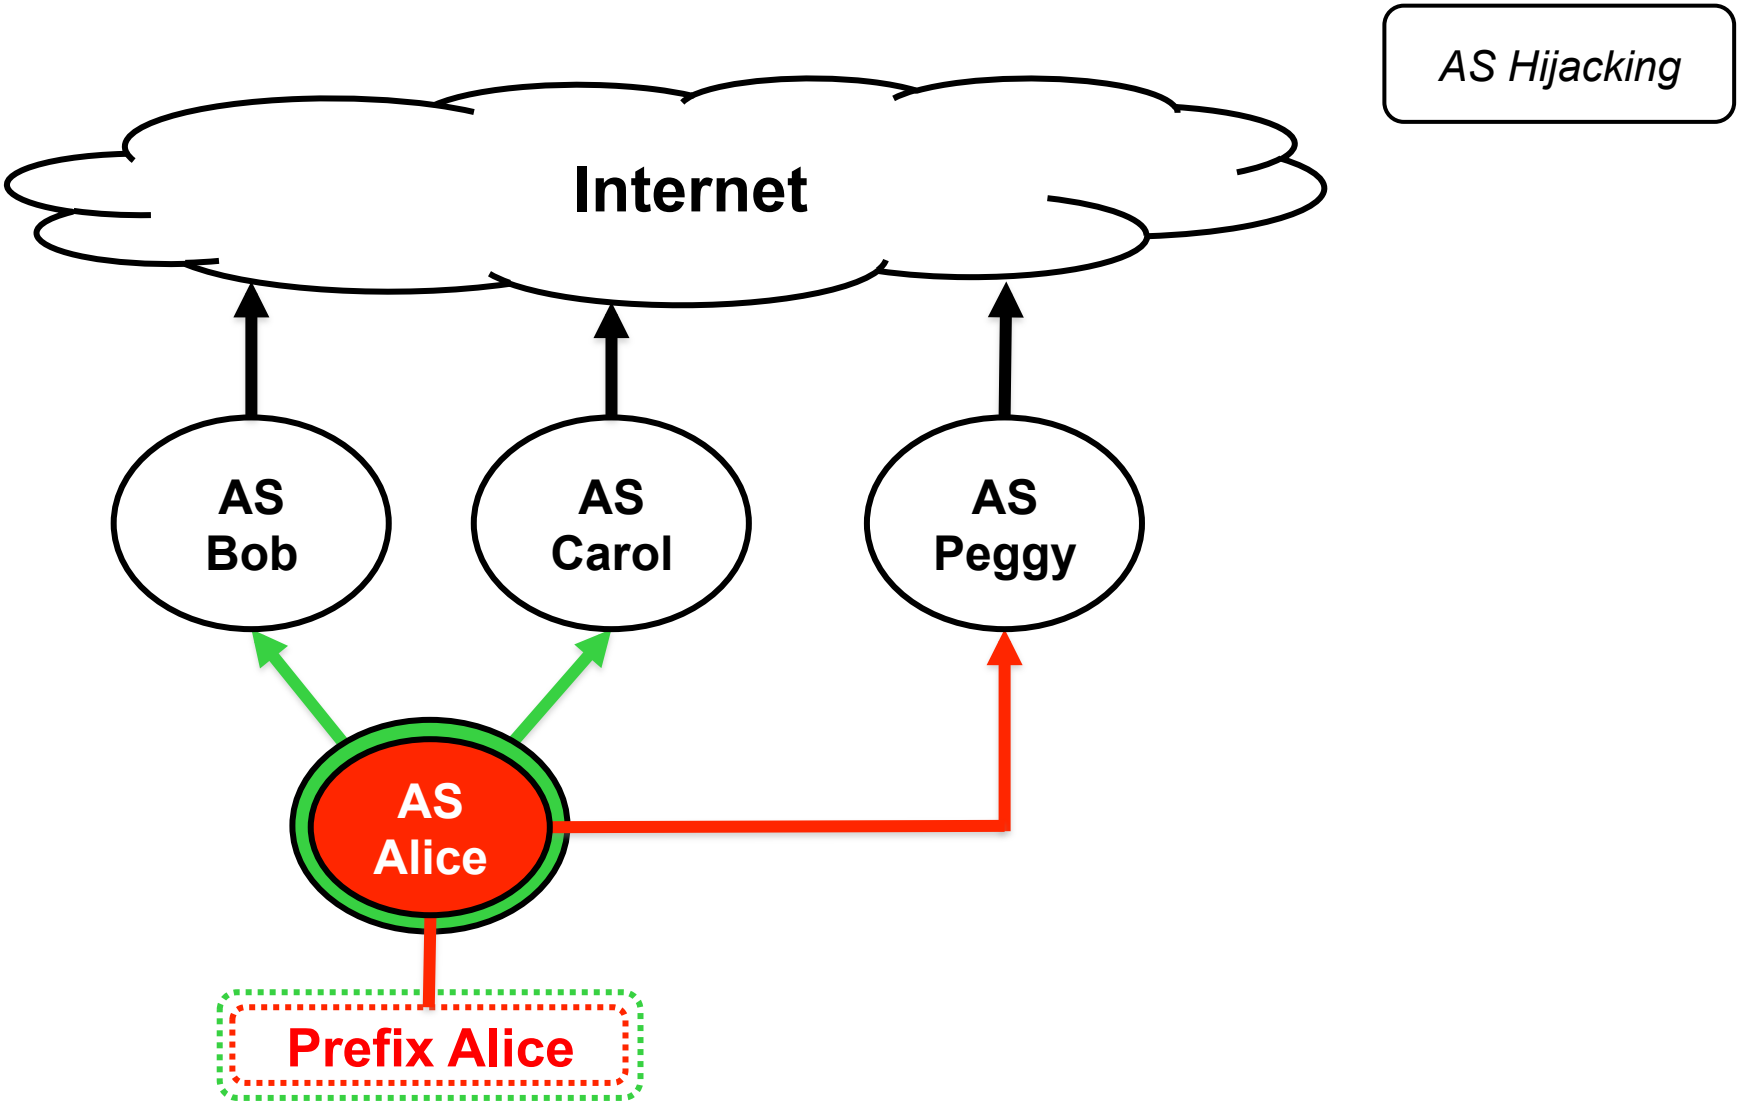

RPKI: no Protection against AS Hijacking

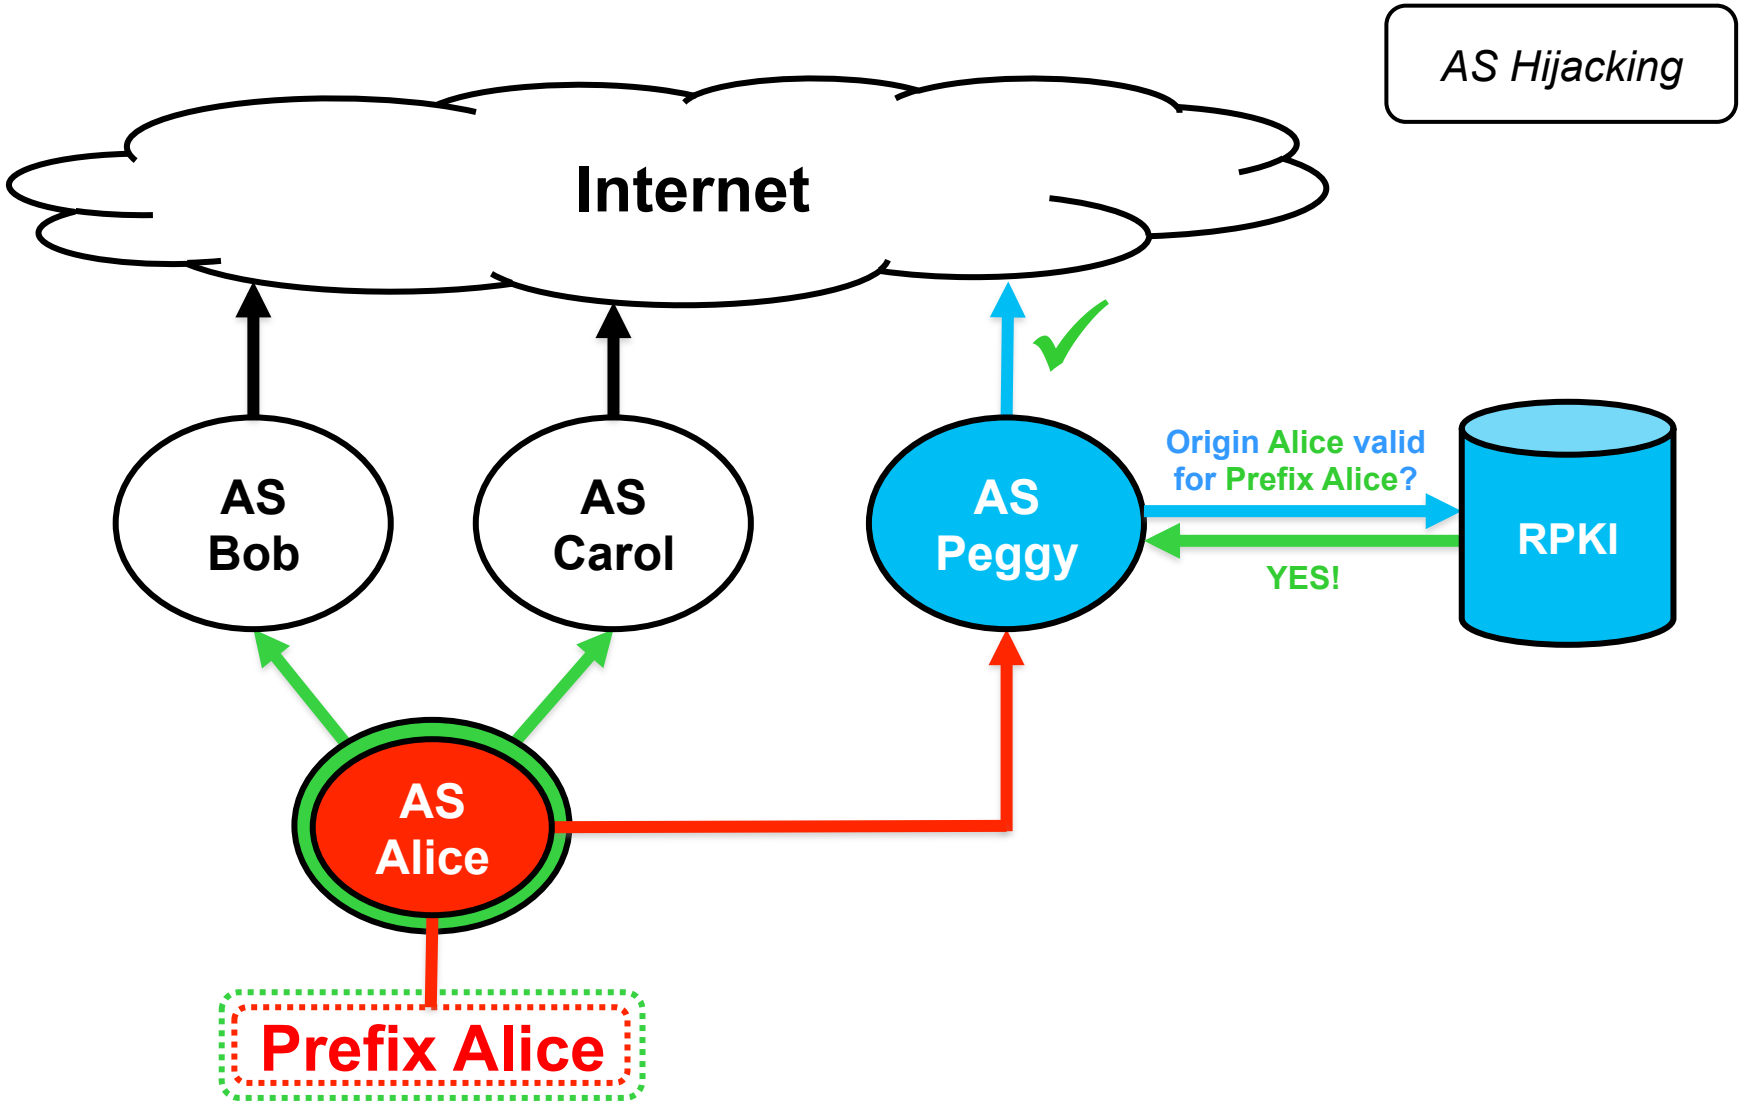

AS Hijacking: LinkTel Case

- **Common prefix Hijacking**
  - Attacker announces victim‘s (sub-) prefix
  - State of the art offers real-time detection
- **AS Hijacking**
  - Formless *letter of authorization (LoA)* is often accepted by ISPs as legitimation to advertise resources of a customer’s AS
  - More sophisticated type of attack
  - Is aimed at a long-term benefit such as over a duration of months
  - Special case of path attack
- **The “LinkTel Case“**
  - SOS mail to NANOG list from a Russian ISP
  - State of the art detection and prevention within RPKI limited
  - Little forensic evidence, but **AS Hijacking is real**

Johann Schlamp, Georg Carle, and Ernst W. Biersack. **A Forensic Case Study on AS Hijacking: The Attacker's Perspective**. *ACM Computer Communication Review (CCR)*, April 2013

AS Hijacking: LinkTel Case

The LinkTel Case

- Attack starts on **March 11, 2011** with
  - **DNS re-registration** of link-telecom.biz and abuse of mail address noc@link-telecom.biz
  - **Forged letter of authorization** to upstream provider
- First malicious BGP announcement on April 15
- Announcement of **unallocated space** on May 12
- **Hijacking of further prefixes** until June 9
- Counter-announcements by victim without effect until announcement of **more specifics**
- **Deeper analysis by utilizing flow data**

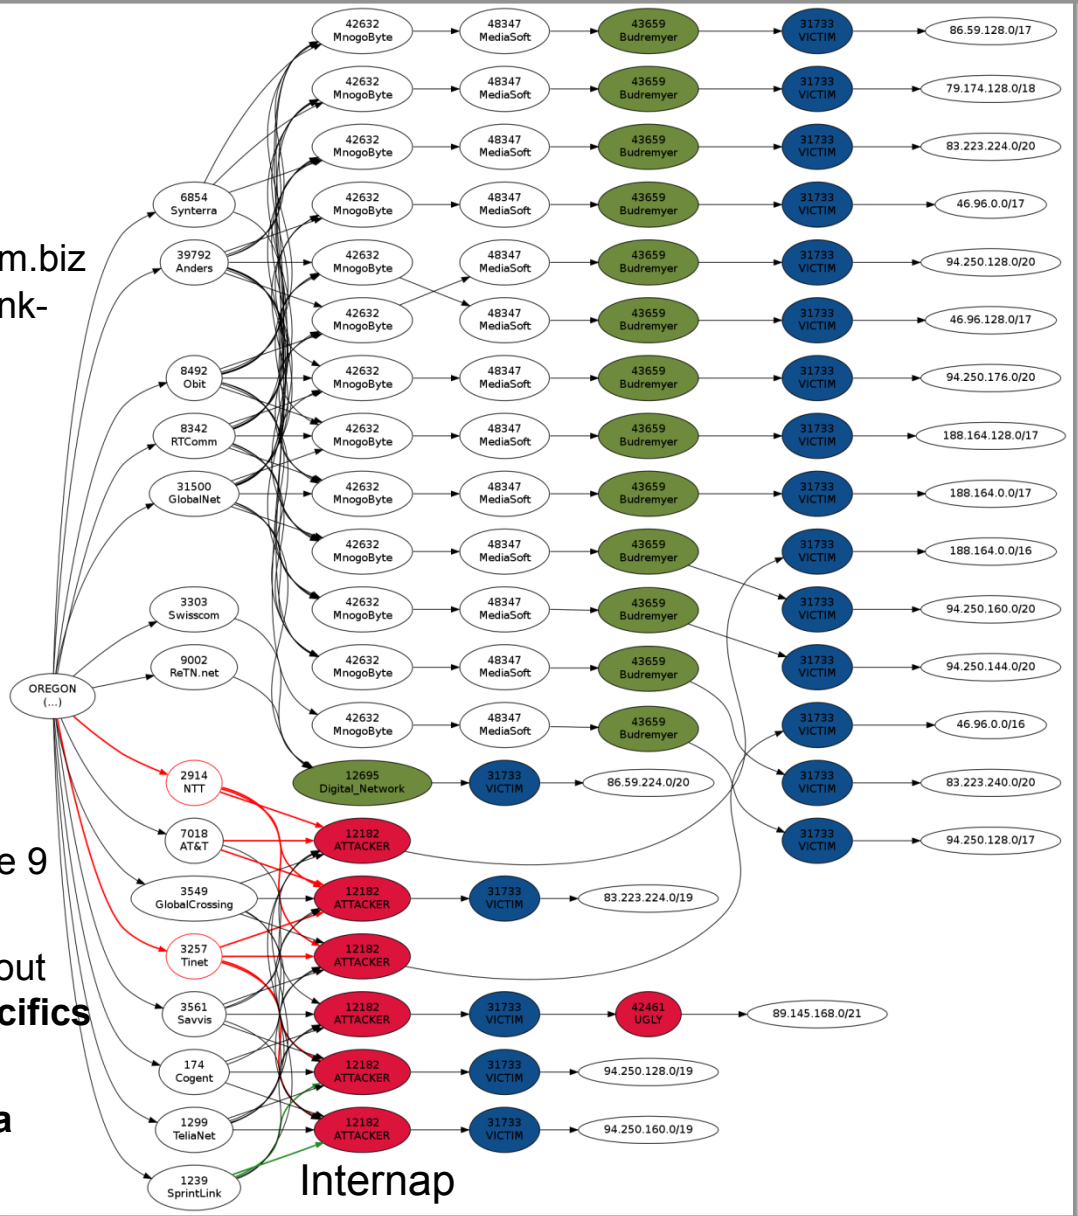

AS Hijacking: LinkTel Case

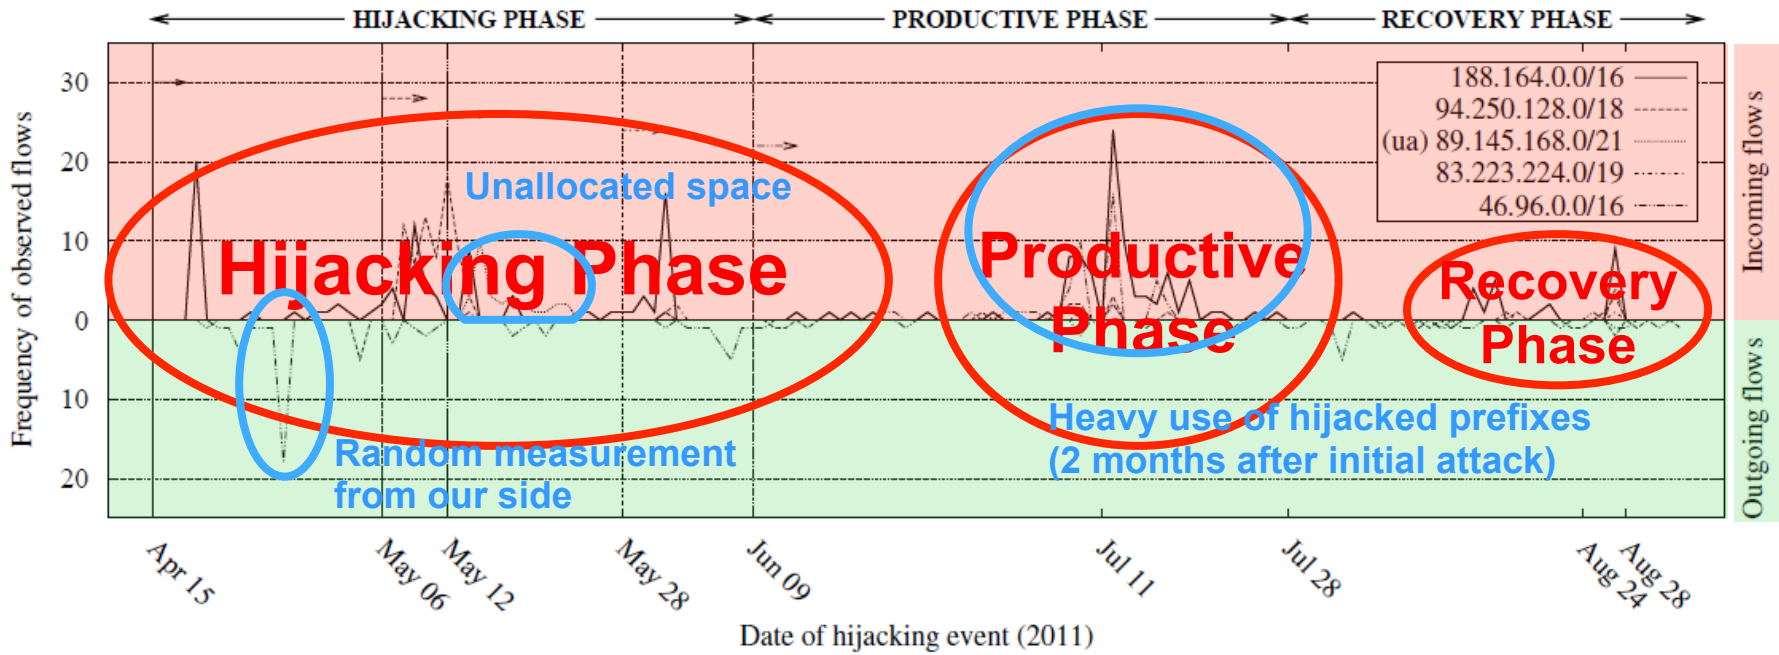

- Traffic contains web traffic, IRC and spam

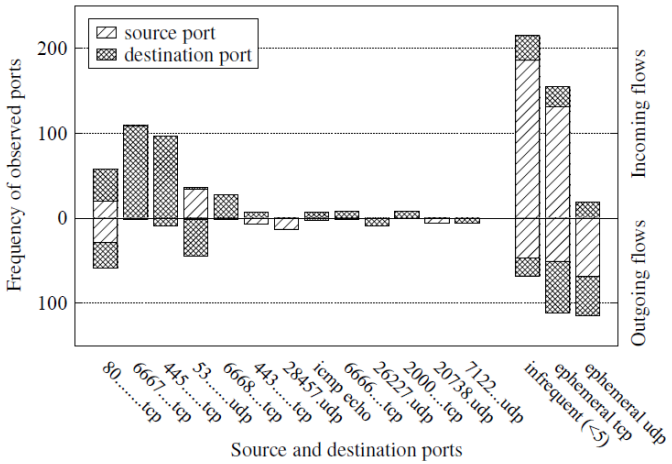

Lessons from

- **The attacker knew that...**
  - ...the victim‘s DNS domain was going to expire
  - ...the attack is going to be carried out many weeks in advance
  - ...that only a single prefix was announced by the victim
  - ...which prefixes are allocated by RIRs and which are not
  - ...when he had lost (no counter-attacks)

- **The attacker knows his steps**
  - Detailed, complex attack plan
  - Access to a variety of data sources (including DNS and RIR DBs)
  - Hand-picked target

- Further insights by **analyzing flow data** from our MWN NetFlow/IPfix observations

Design of an AS Hijacking Early Warning System

- **Escalation warning system**
  - Passive Monitoring of DNS expiry and re-registration
  - Analysis of reverse DNS and BGP activity
  - Identification of vulnerable targets
  - Integration of blacklists and active measurements possible

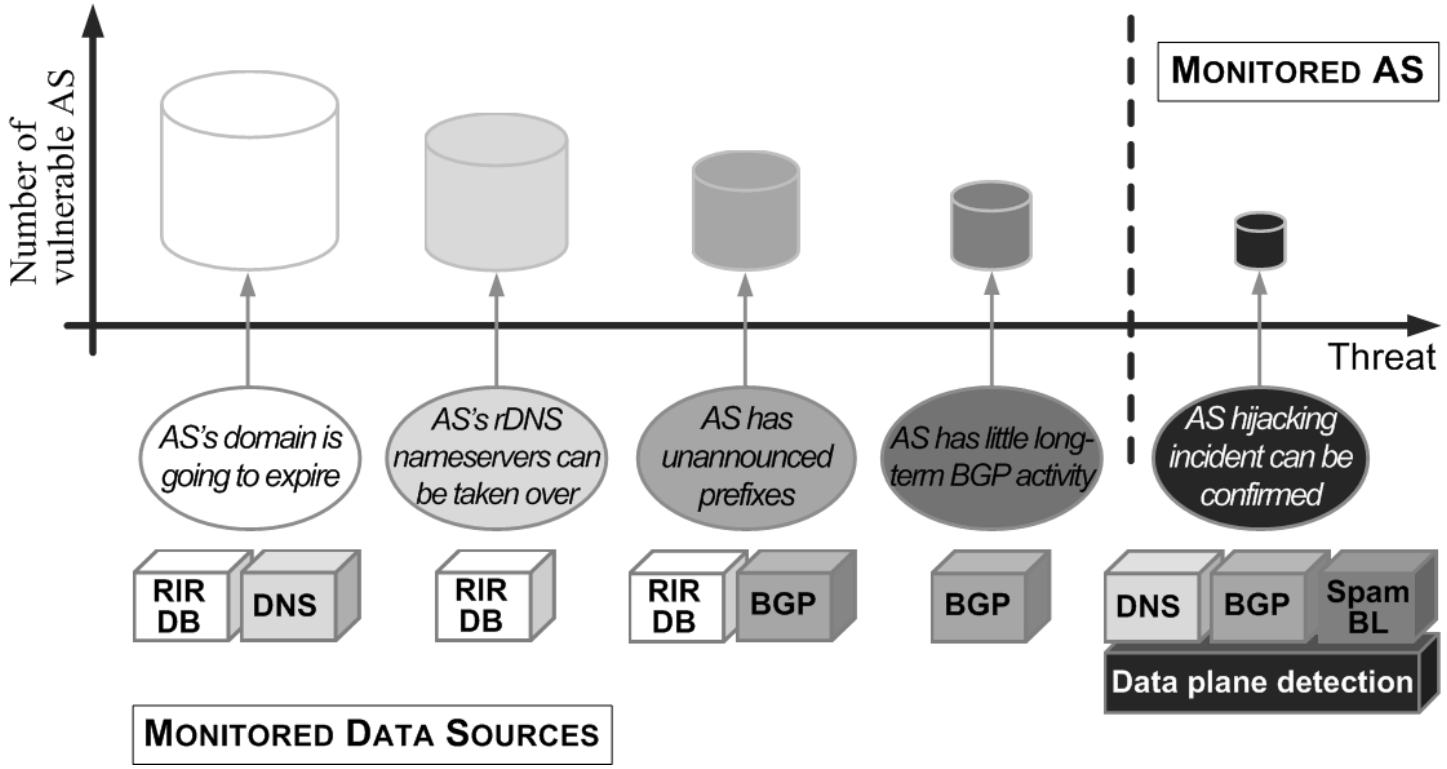

Johann Schlamp, Georg Carle, and Ernst W. Biersack. **A Forensic Case Study on AS Hijacking: The Attacker's Perspective**. *ACM Computer Communication Review (CCR)*, April 2013.