NATIONAL INSTITUTE OF STANDARDS AND TECHNOLOGY
U.S. DEPARTMENT OF COMMERCE

ASPA-based BGP AS_PATH Verification and Route Leaks Solution

Kotikalapudi Sriram

US NIST

IETF Draft: https://datatracker.ietf.org/doc/draft-ietf-sidrops-aspa-verification/

Authors: A. Azimov, E. Bogomazov, R. Bush, K. Patel, J. Snijders, K. Sriram, C. Jeker

NANOG 89

October 2023

Outline of the Talk

- Brief refresh about BGP prefix hijacks, route leaks, and AS_PATH manipulations, RPKI-ROV, BGPsec
- Autonomous System Provider Authorization (ASPA)
- ASPA-based BGP AS_PATH verification & route leaks detection and mitigation

Border Gateway Protocol (BGP) Basics

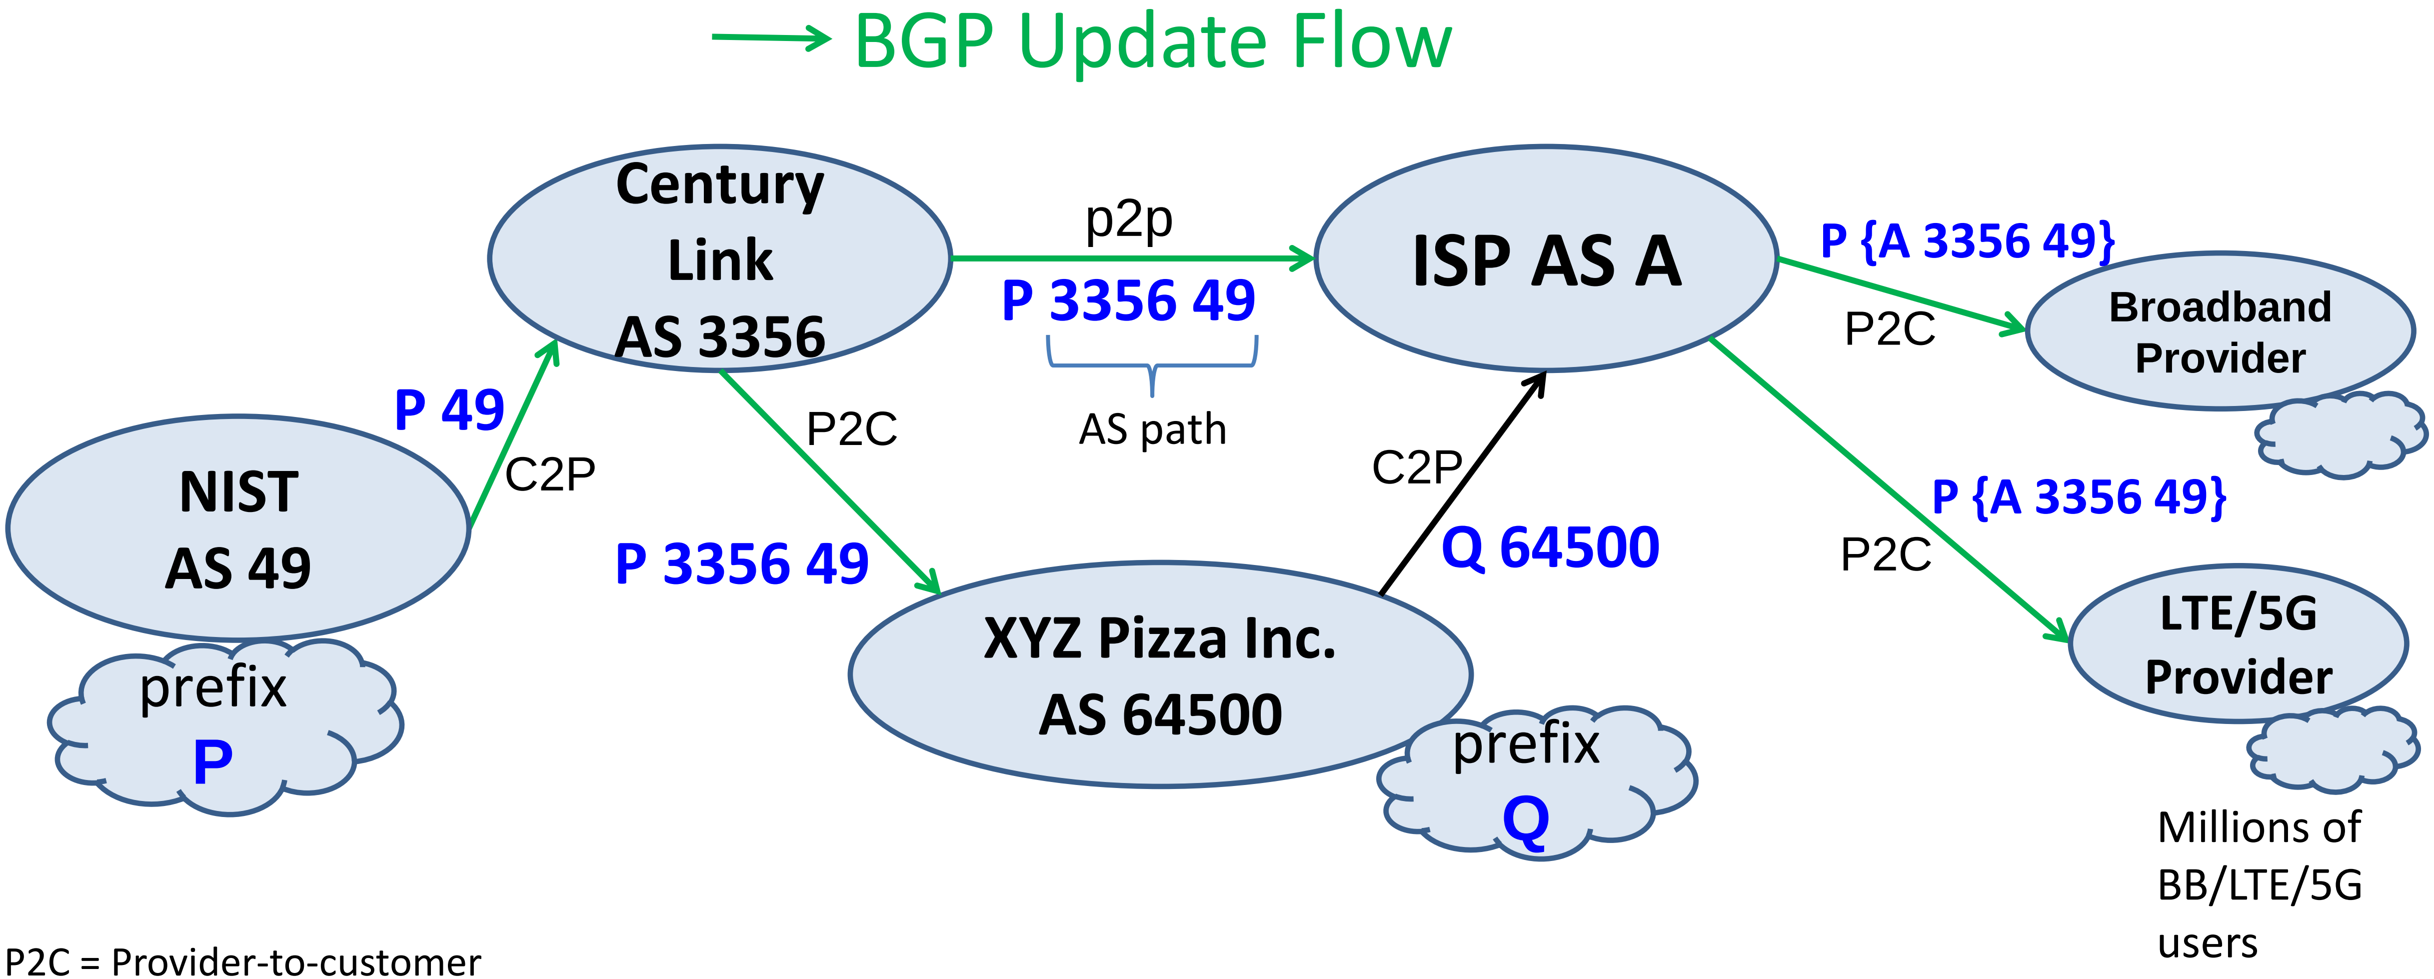

P2C = Provider-to-customer
C2P = Customer-to-provider
p2p = peer-to-peer (lateral peers)
AS = autonomous system

Note: This is only an illustration.
Not shown but update for prefix Q also propagates to all other ASes.

AS = Autonomous System

Border Gateway Protocol (BGP) Basics

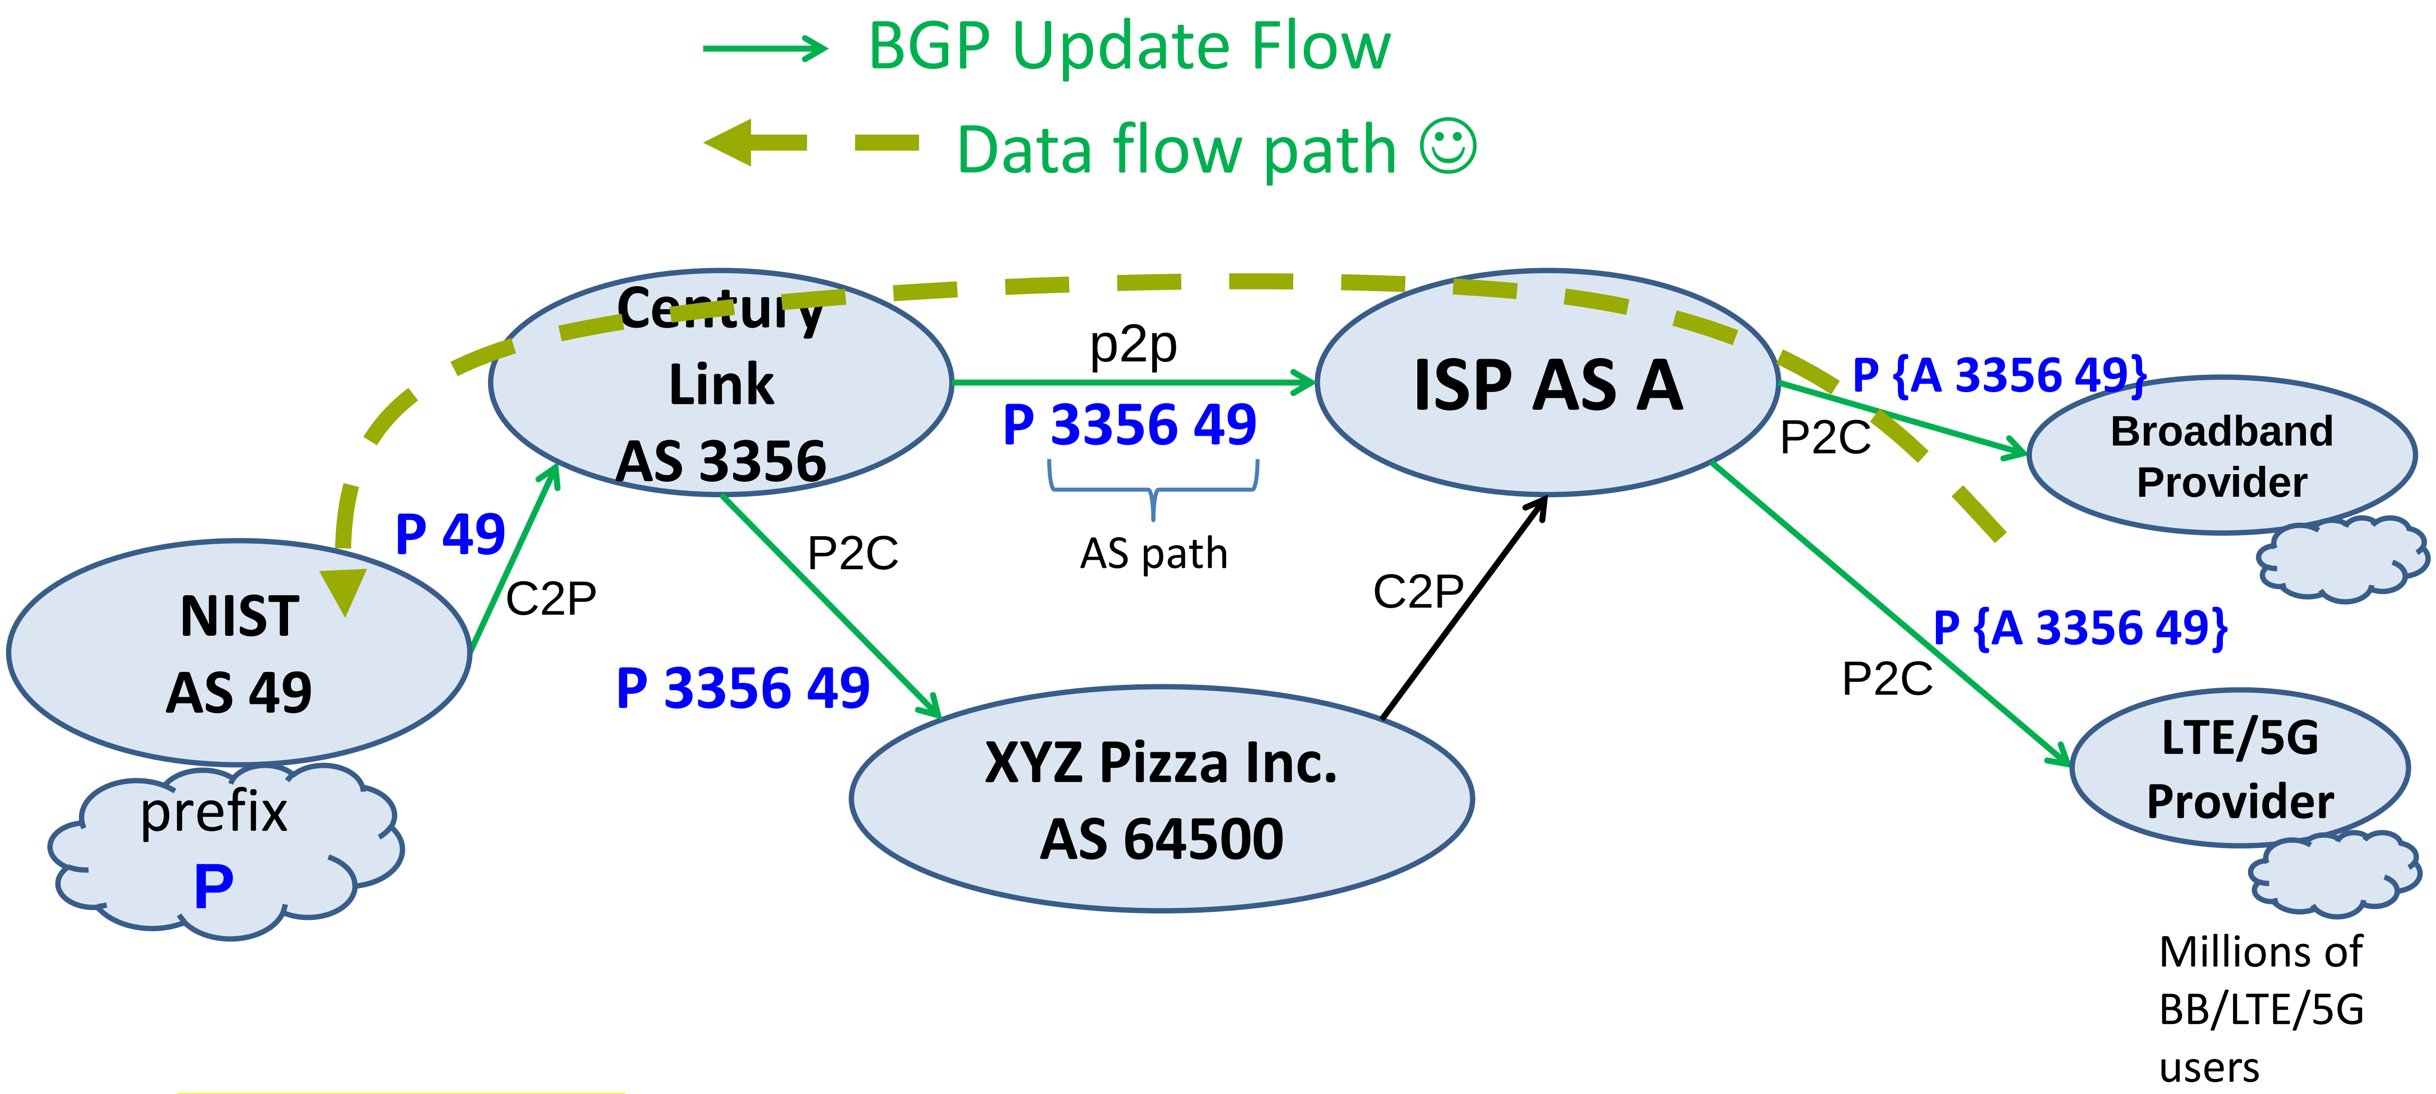

Note: This is only an illustration.

Prefix Hijack

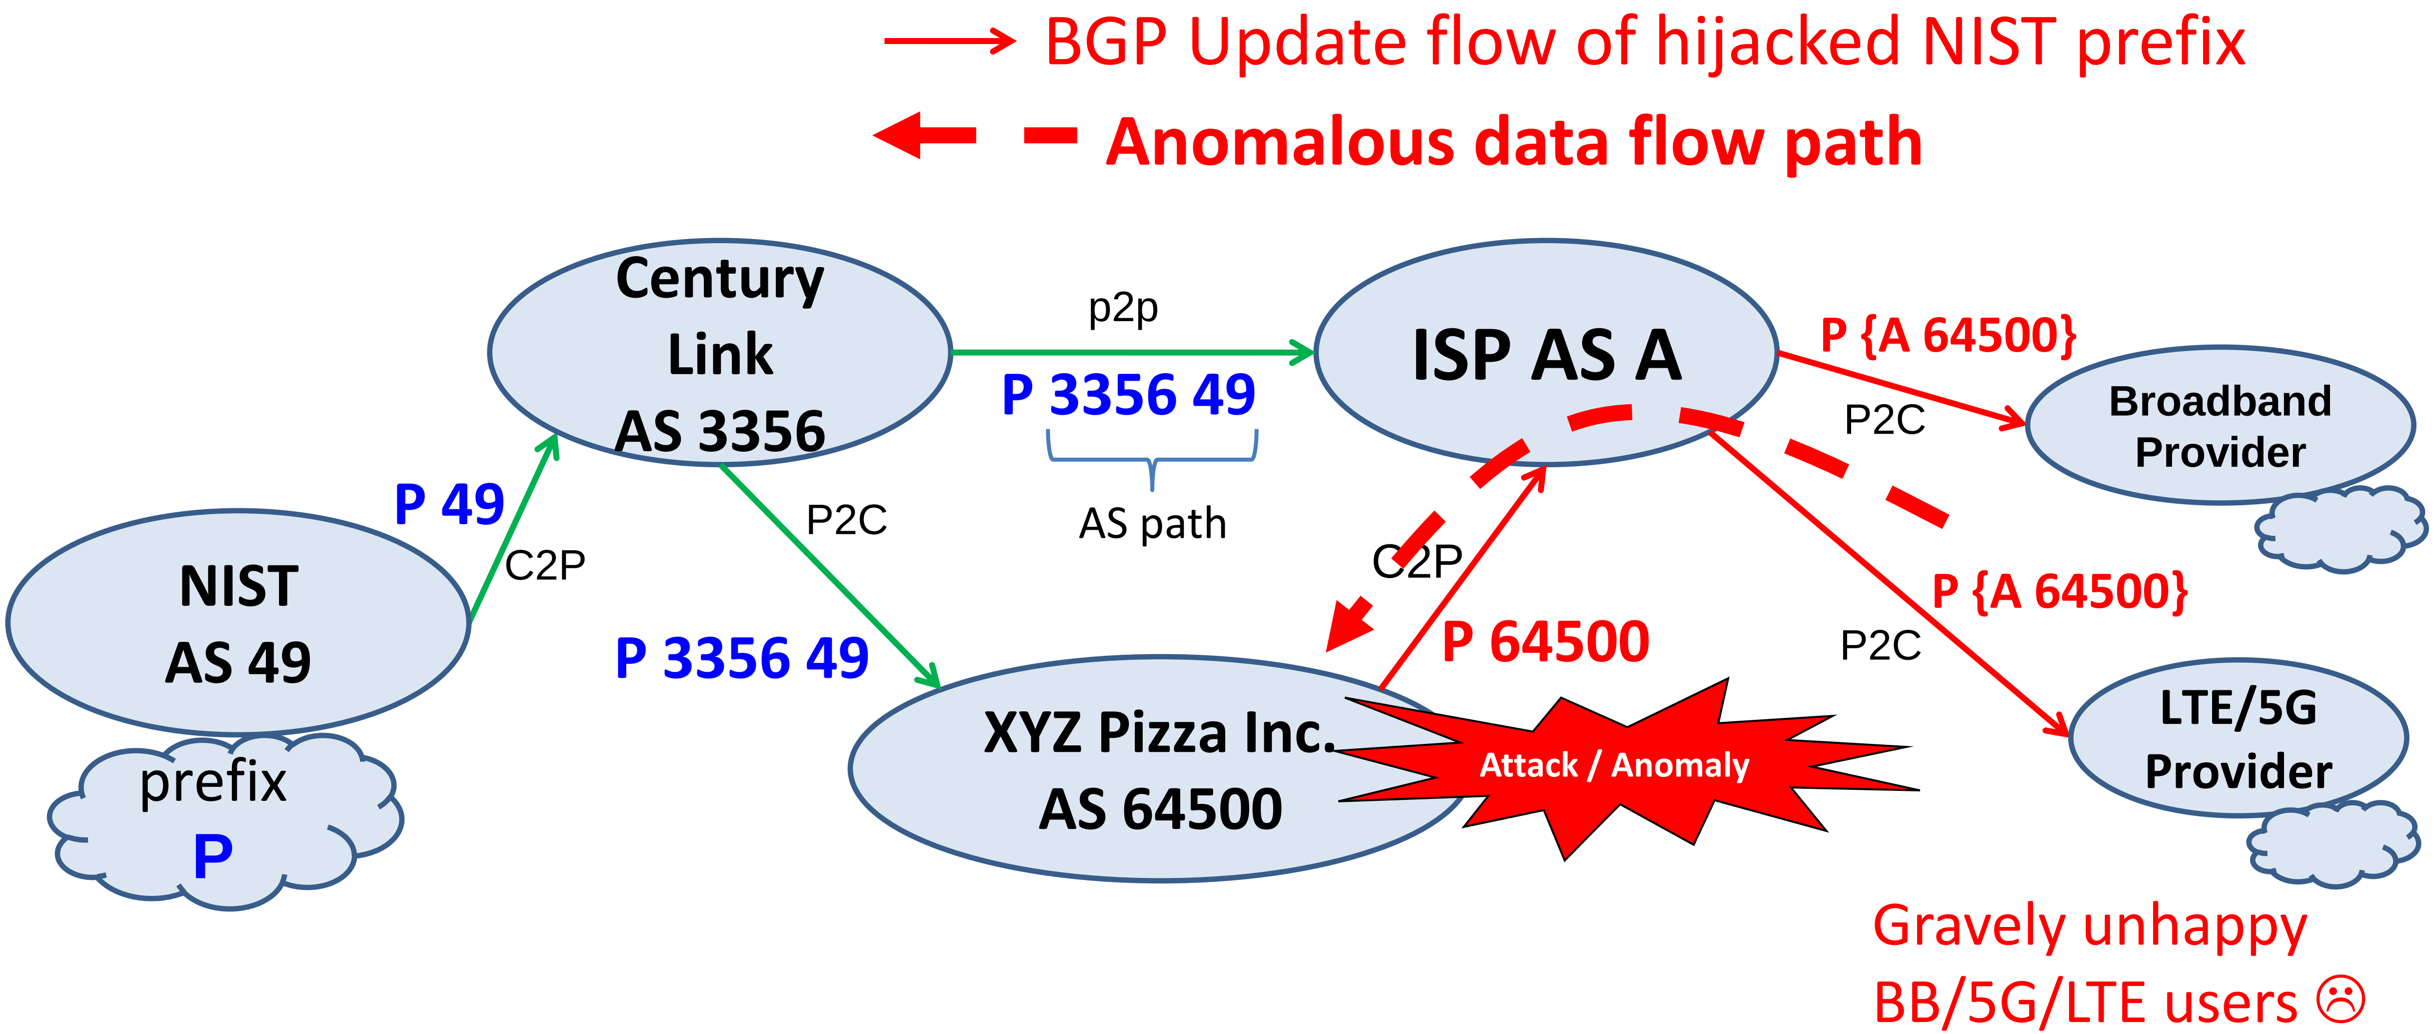

In general, ISPs prefer customer route announcements over those from other peers.

Note: This is only an illustration.

Solution for Prefix Hijacking

Resource PKI (RPKI) and Route Origin Authorization (ROA)

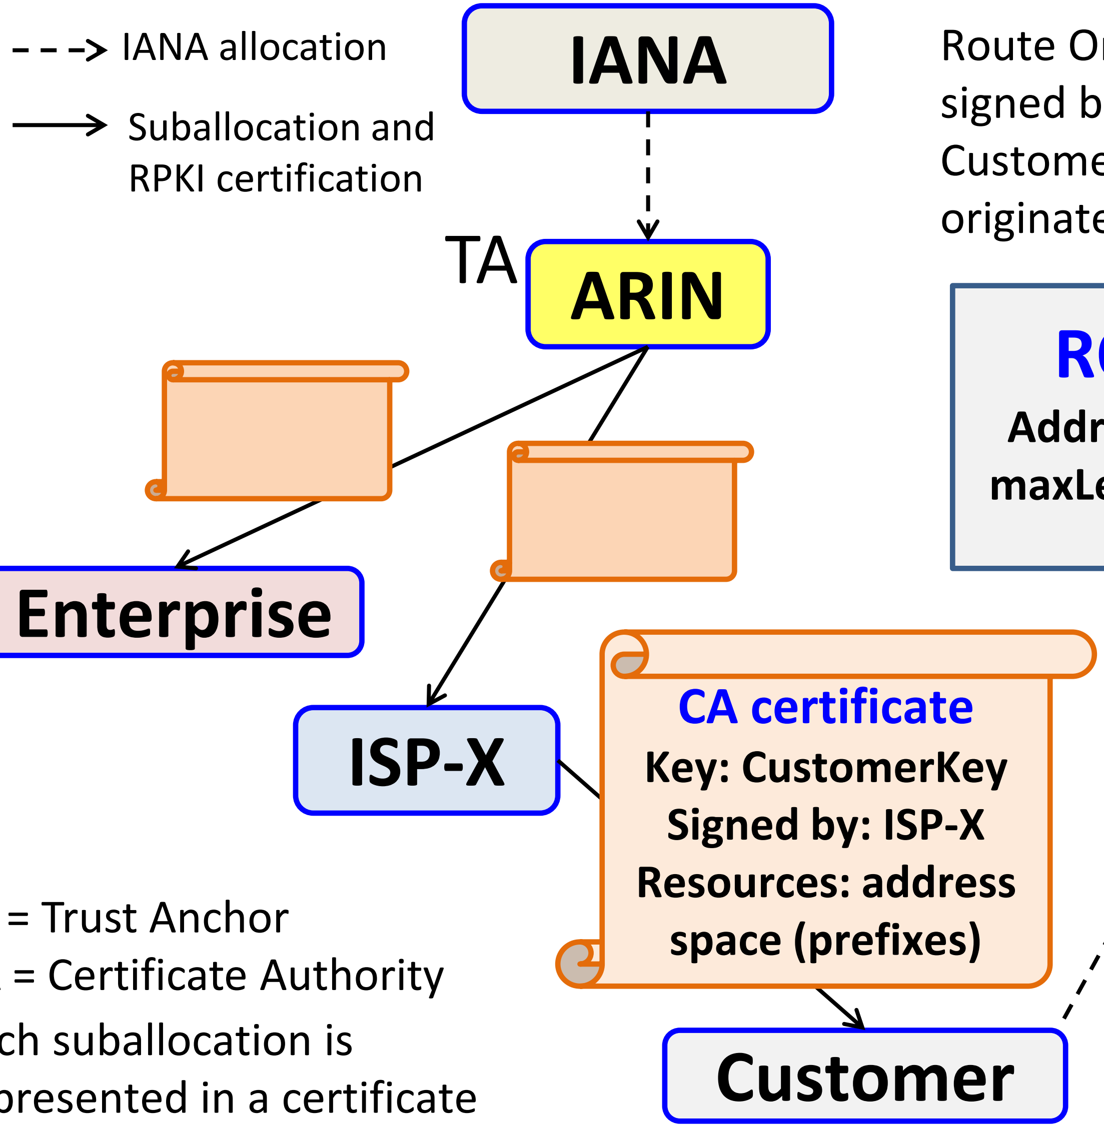

Route Origin Authorization (ROA) is signed by an ISP, Enterprise, or Customer. ROA authorizes an AS to originate one or more prefixes.

ROA (signed object)

Addresses: set of customer prefixes
maxLengths: associated with prefixes
Valid Origin: ASn

ROA is signed using the private key corresponding to an End-Entity certificate that is derived from the CA certificate.

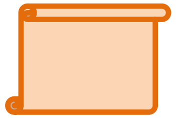

X.509 Cert

TA = Trust Anchor
CA = Certificate Authority

Each suballocation is represented in a certificate

Route Origin Validation (ROV)

ROV is performed at routers

- Routers match the prefix-origin pair in the route against the ROA data
- Determine route validity: Valid, Invalid, Unknown

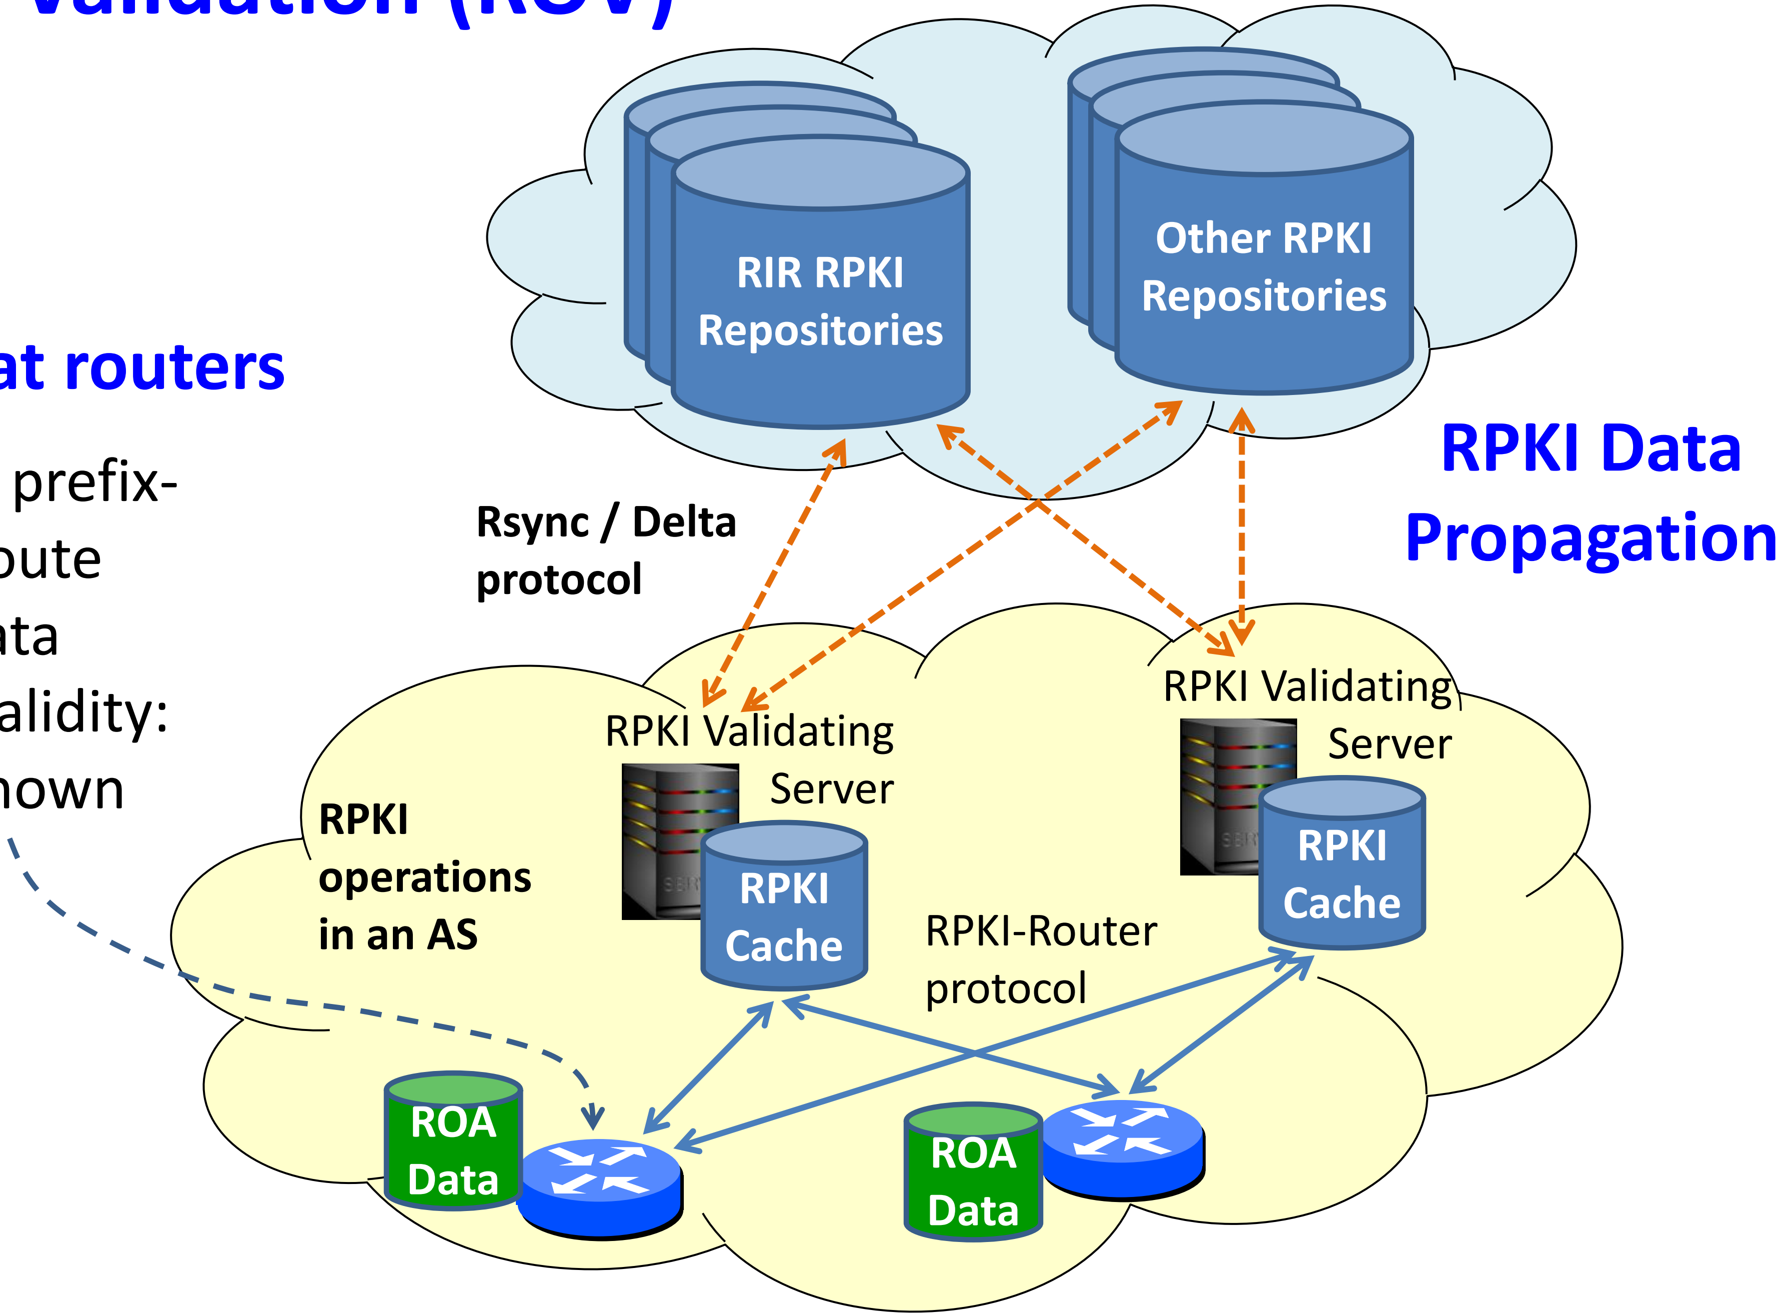

Forged-Origin Prefix Hijack

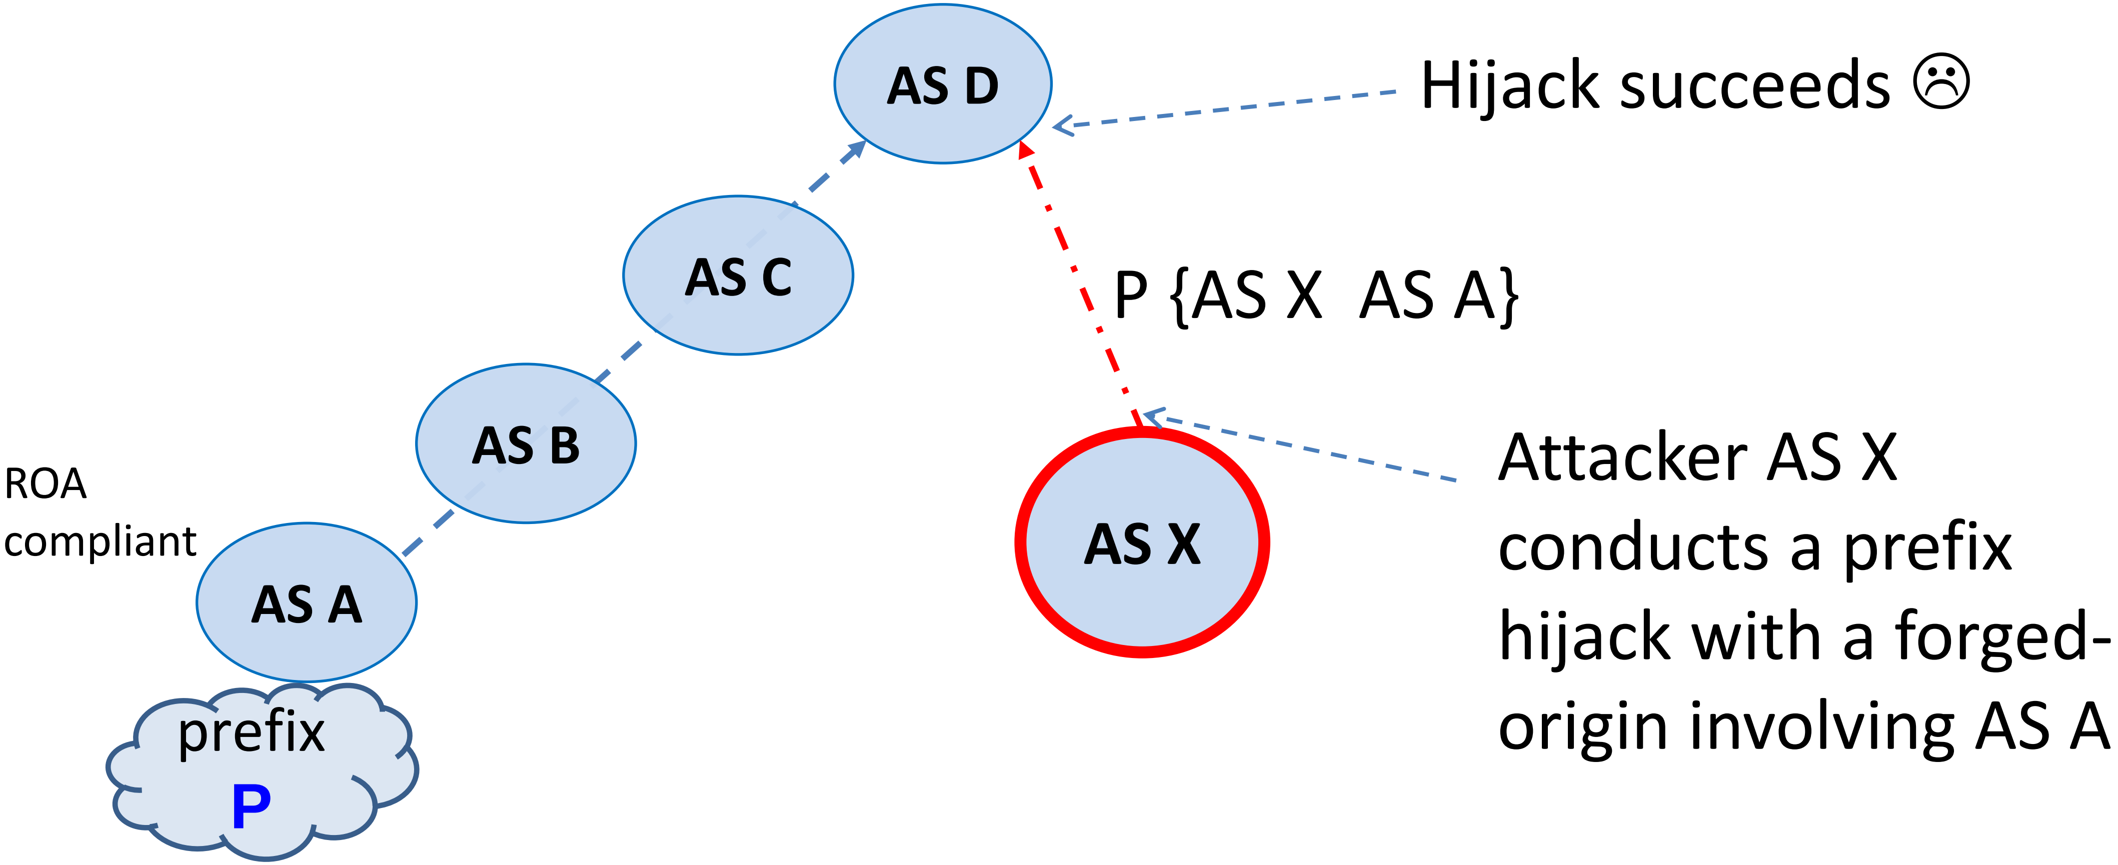

AS A has a ROA: {P, AS A}

AS Path Protection (BGPsec, RFC 8205)

Basic Principle of BGPsec AS Path Signing

Route Origin Authorization (ROA) exists that authoritatively binds the prefix P to the origin AS1

Forged-Origin Announcement will be Invalid

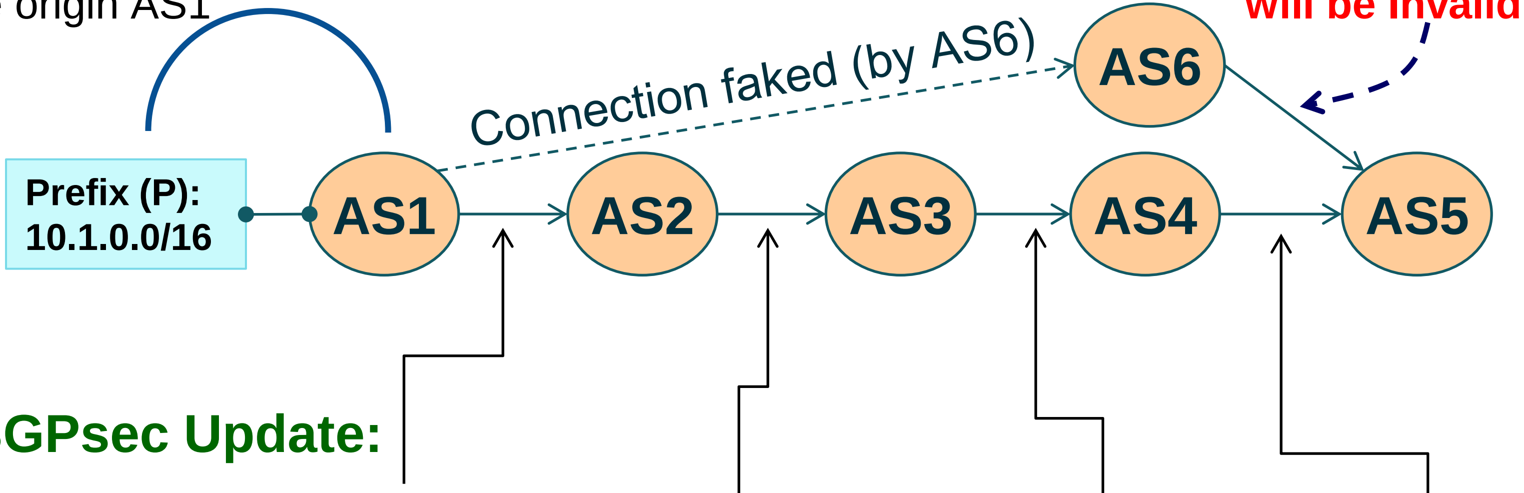

BGPsec Update:

P, {AS1, SKI1, AS2*}, {SIG1-2}

P, {AS1, SKI1, AS2, SKI2, AS3*}, {SIG1-2, SIG2-3}

P, {AS1, SKI1, AS2, SKI2, AS3, SKI3, AS4*}, {SIG1-2, SIG2-3, SIG3-4}

P, {AS1, SKI1, AS2, SKI2, AS3, SKI3, AS4, SKI4, AS5*}, {SIG1-2, SIG2-3, SIG3-4, SIG4-5}

*** Next hop AS is signed over but not included in the forwarded BGPSEC update.**

Note that if AS6 attempts to announce prefix P over a one-hop connection via AS1, it will not succeed because it never received a signed BGP announcement directly from AS1 – it can never fake being directly connected to AS1.

"BGPsec Protocol Specification", RFC 8205, https://www.rfc-editor.org/rfc/rfc8205.html

BGPsec does not solve the route leaks problem

BGPsec utilizes the same RPKI infrastructure as ROA/ROV does

SKI = Subject Key Identifier

Route Leak

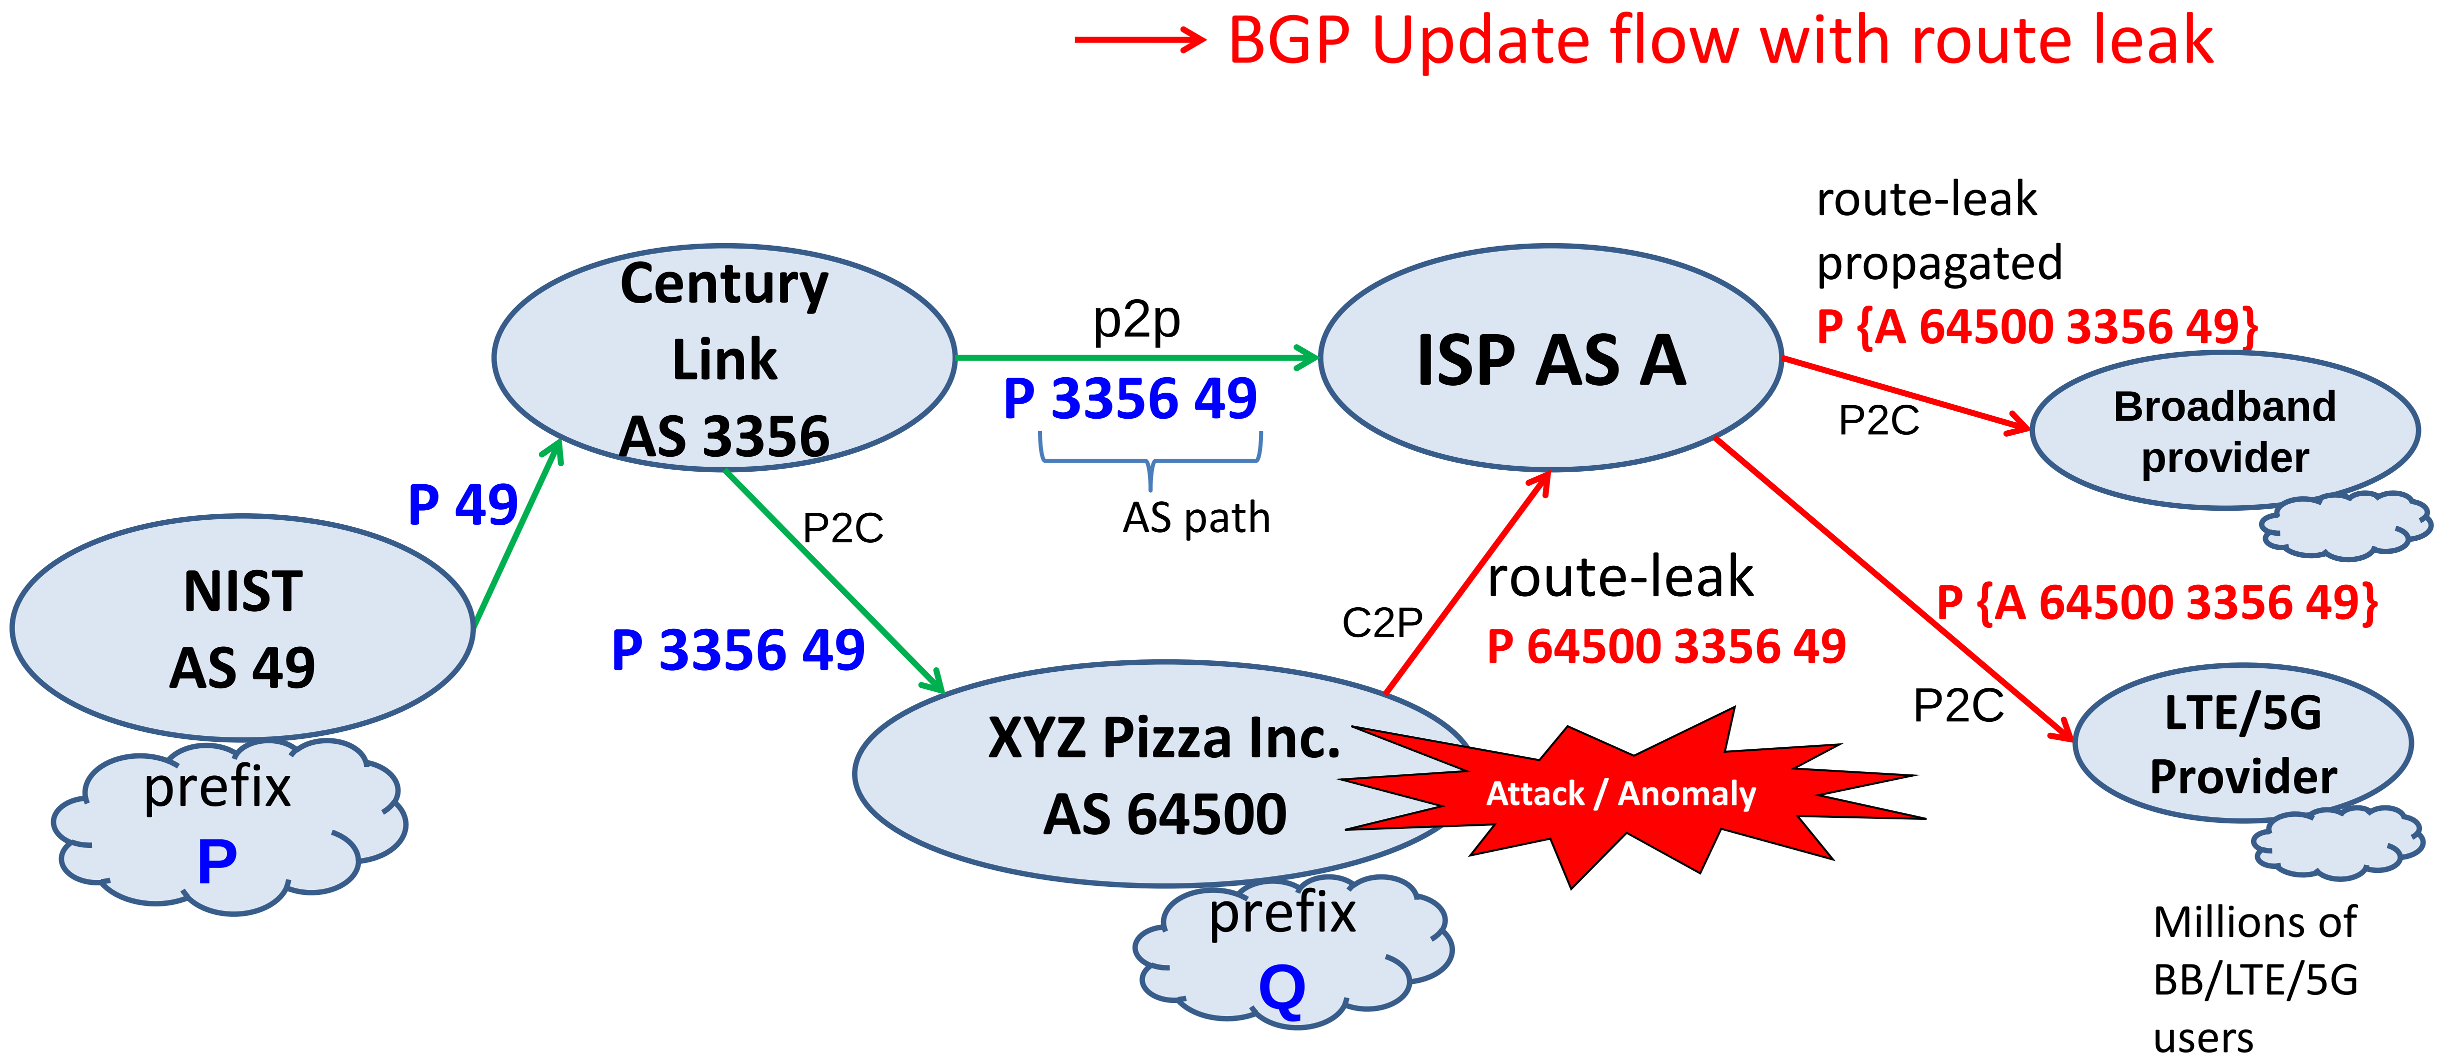

In general, ISPs prefer customer route announcements over those from other peers.

Note: This is only an illustration.

Route Leak

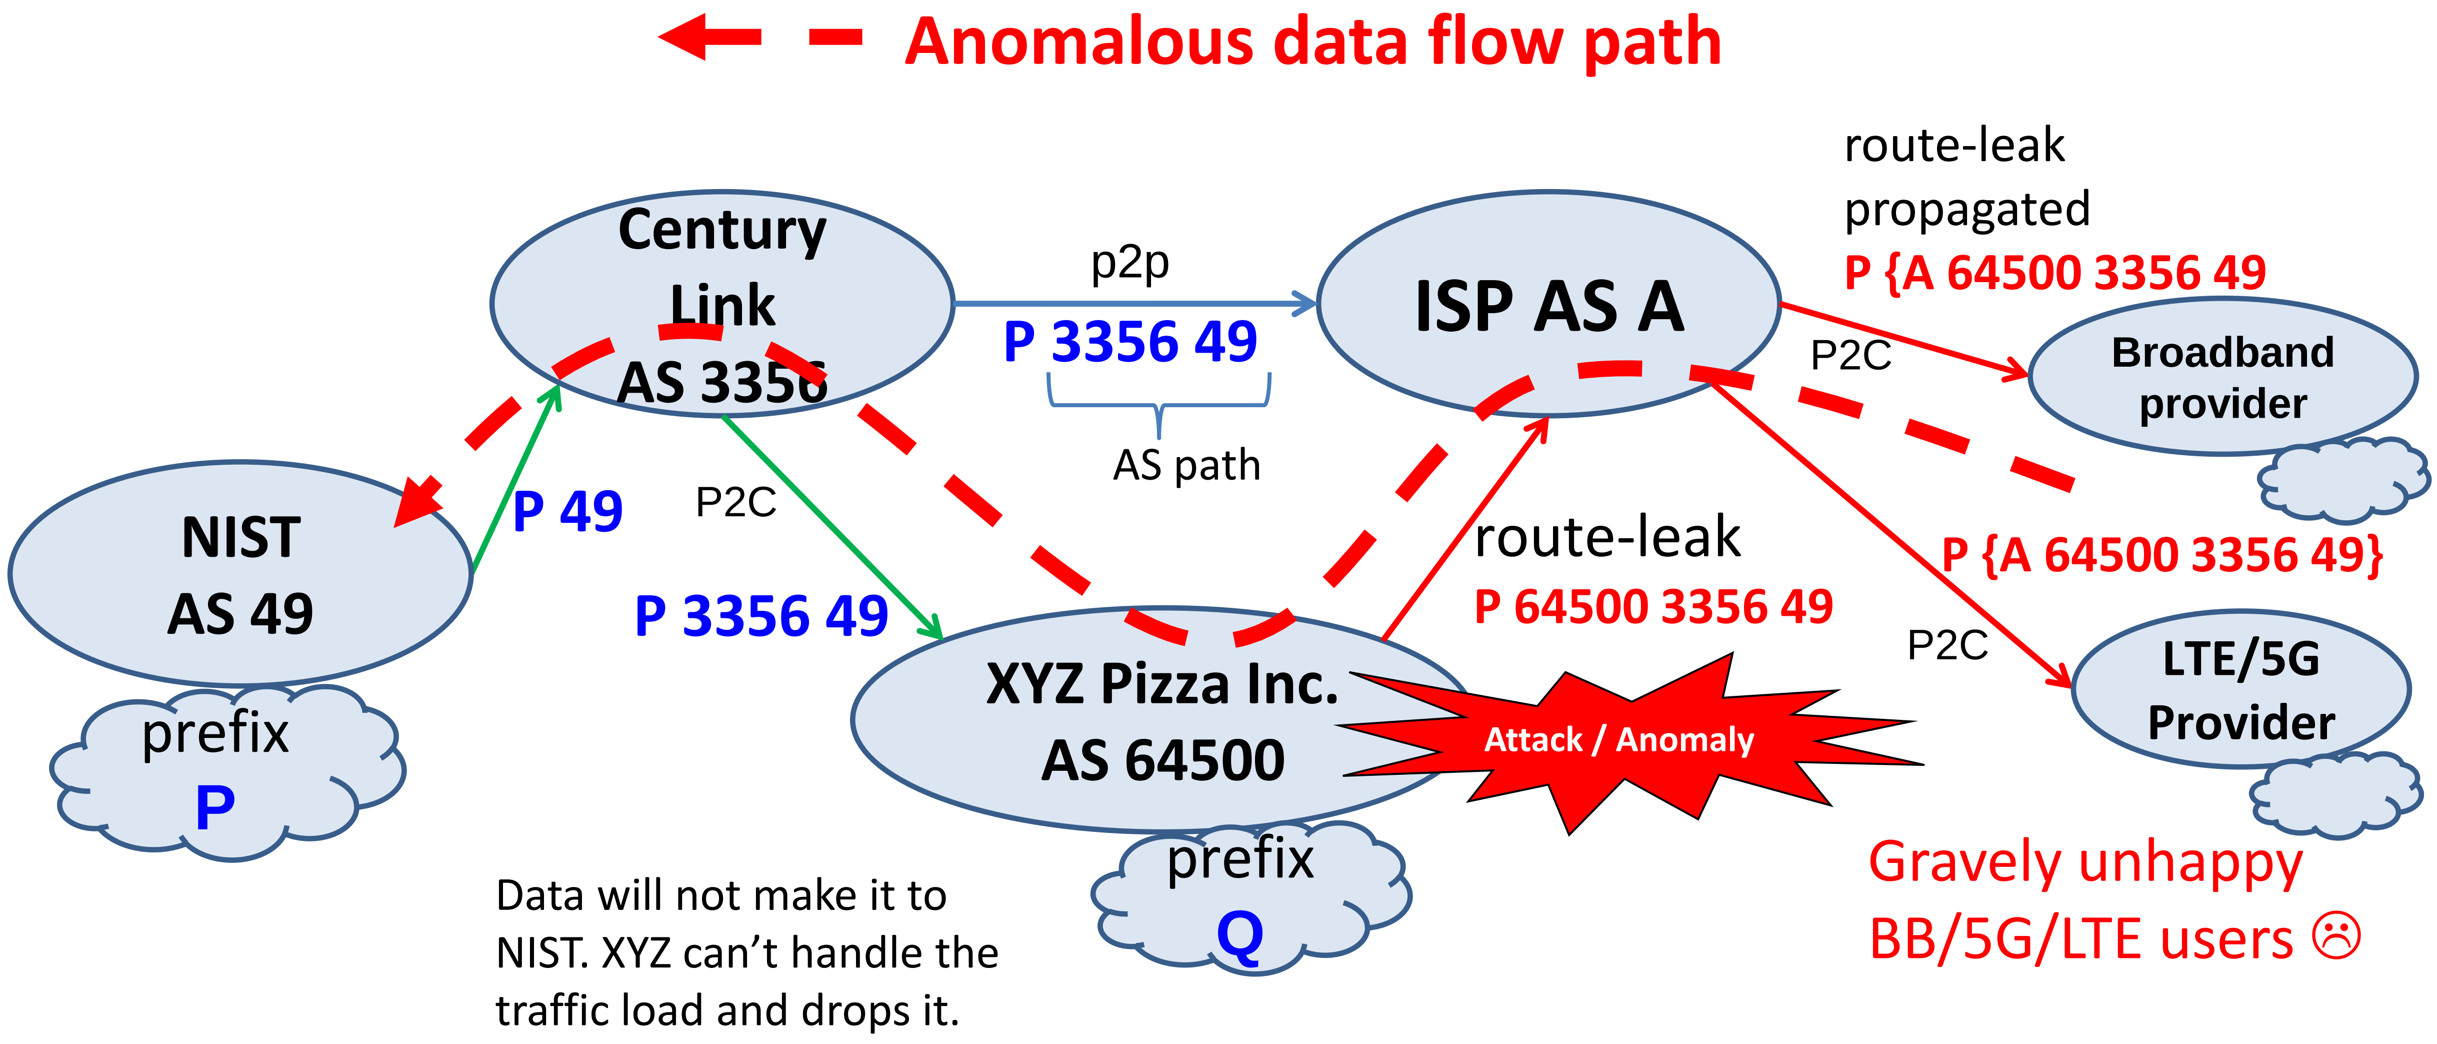

Note: This is only an illustration.

Route Leaks Occur Frequently

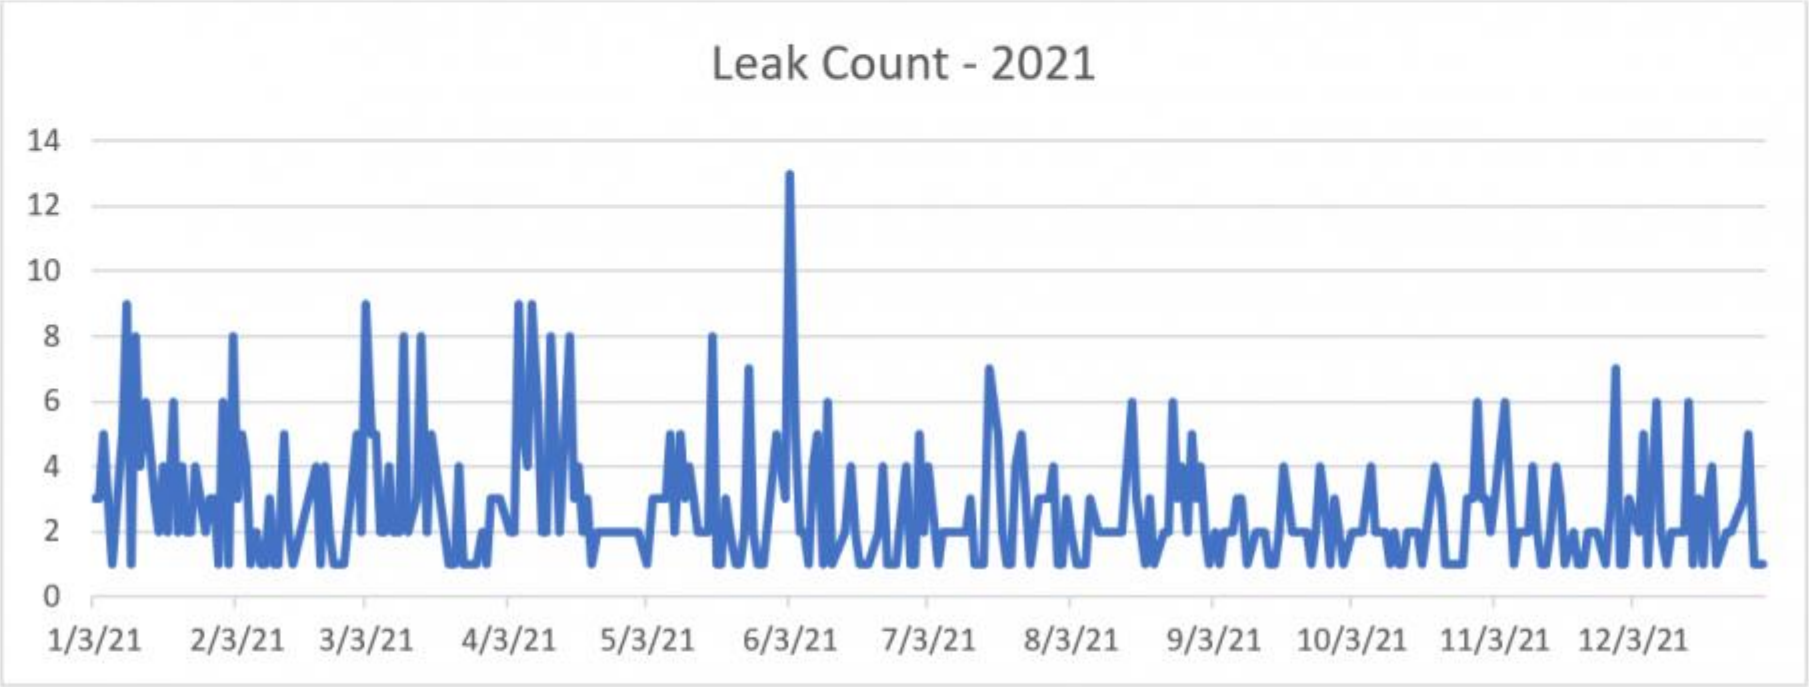

https://www.manrs.org/2022/02/bgp-security-in-2021/

- "New Year, New BGP Leaks," *Kentik Blog*. January 2023. https://www.kentik.com/blog/new-year-new-bgp-leaks/.
- "Major BGP leak disrupts thousands of networks globally," *BleepingComputer*. April 2021. https://www.bleepingcomputer.com/news/security/major-bgp-leak-disrupts-thousands-of-networks-globally/.
- D. Madory, "Large European Routing Leak Sends Traffic Through China Telecom," *MANRS*, Jun. 11, 2019. https://www.manrs.org/2019/06/large-european-routing-leak-sends-traffic-through-china-telecom/.

Example AS Path Trajectories that are Route Leaks

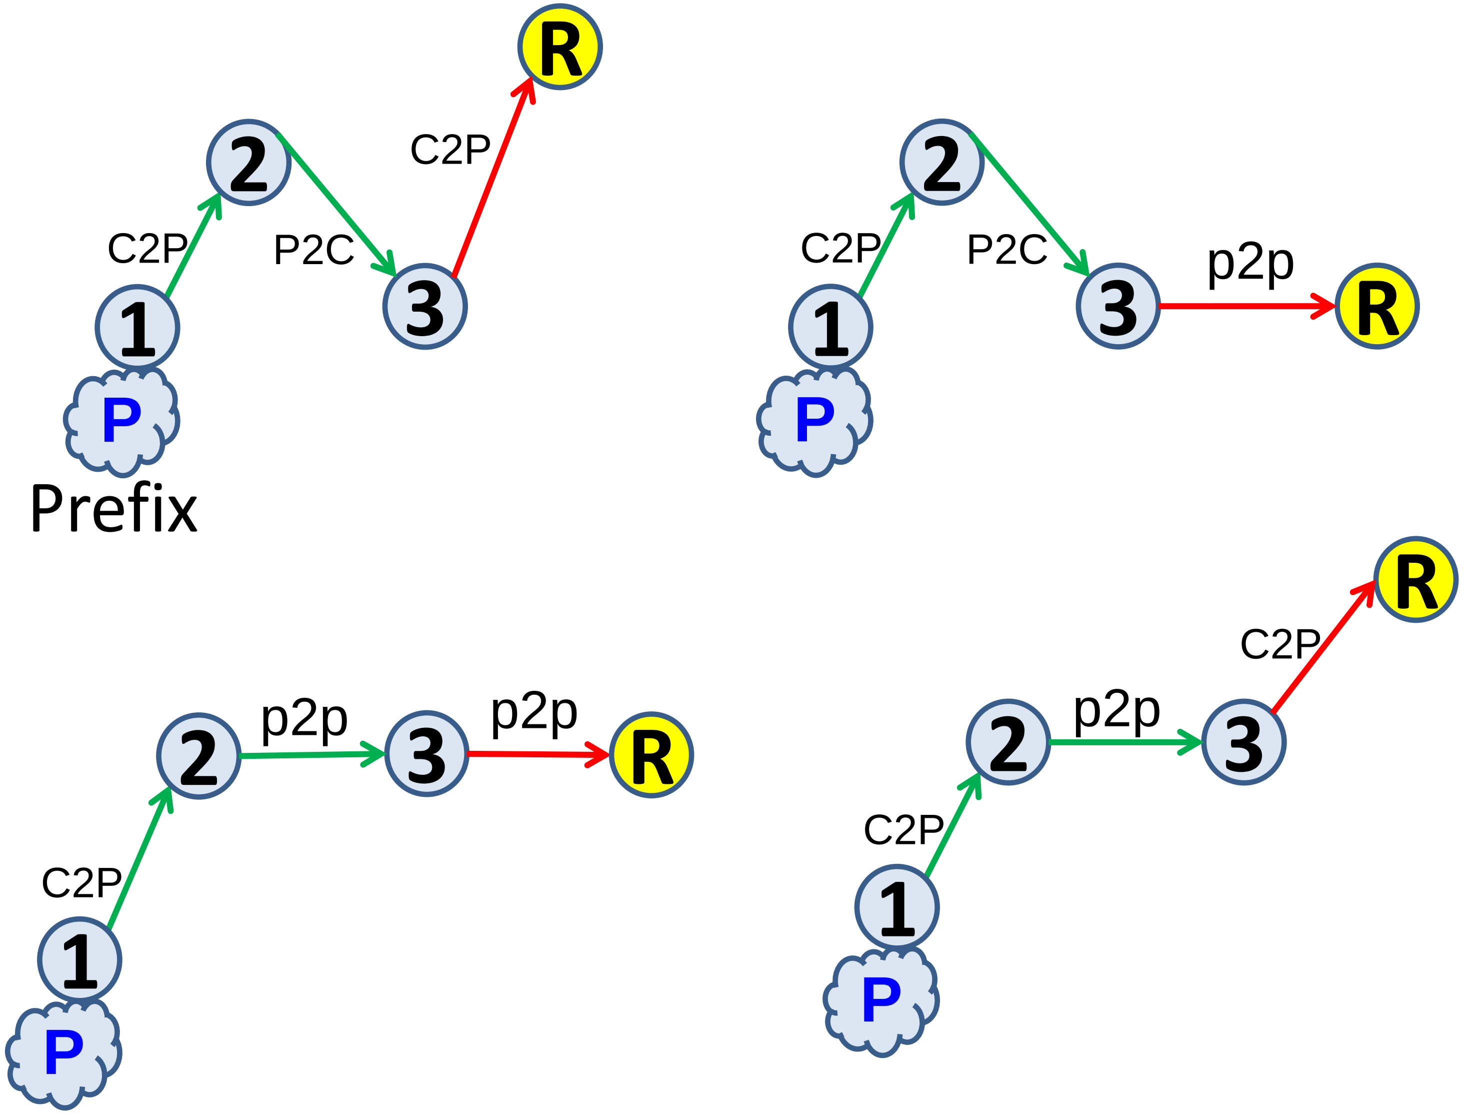

- Route leak occurs if the Update is received on a down (P2C) or lateral (p2p) hop and then forwarded on a up (C2P) or lateral (p2p) hop

Example AS Path Trajectories that are Not Route Leaks

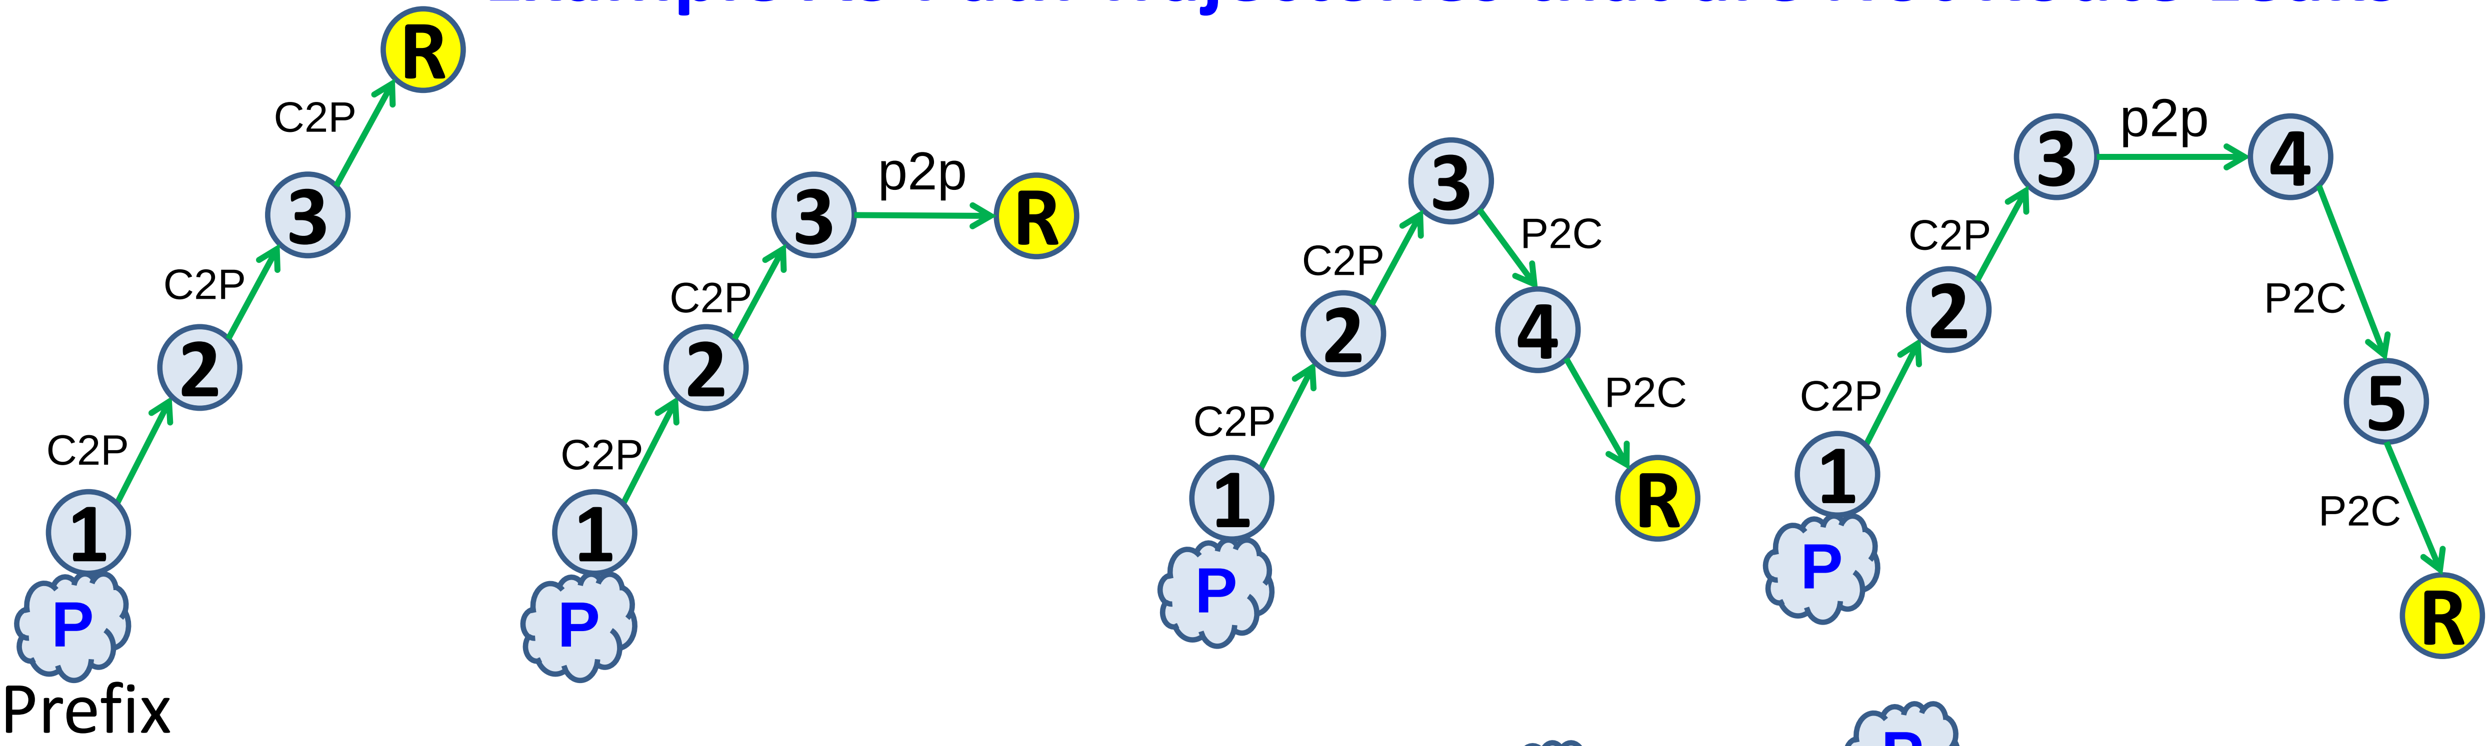

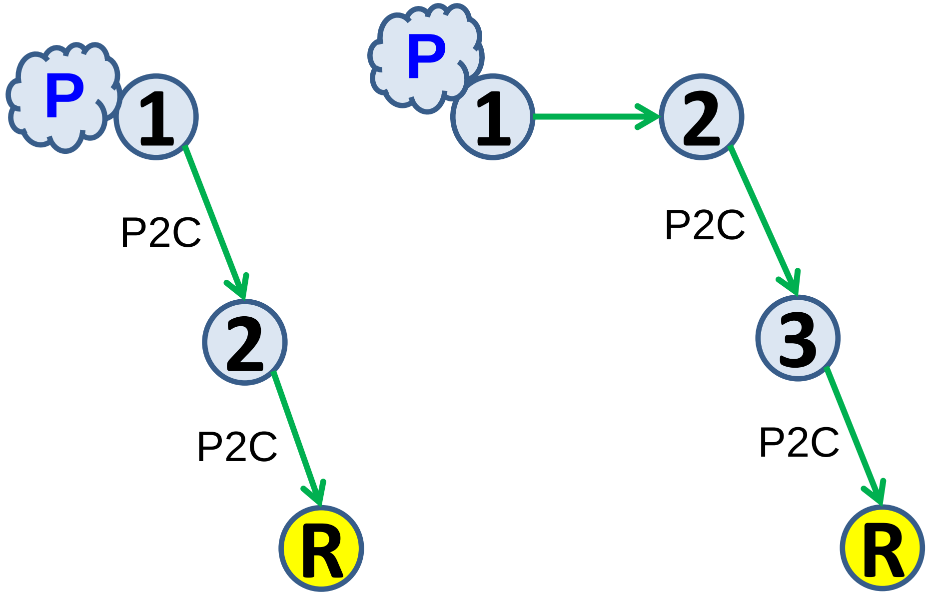

- Not a route leak: If once the BGP Update goes on a down (C2P) or lateral (p2p) hop, then all subsequent hops (if any) must be down (P2C).

P2C = Provider-to-customer
C2P = Customer-to-provider
p2p = peer-to-peer (lateral peers)

R Receiving AS

ASPA-based Solution for Mitigating BGP Route Leaks and AS_PATH Verification

IETF Drafts:

https://datatracker.ietf.org/doc/html/draft-ietf-sidrops-aspa-verification

https://datatracker.ietf.org/doc/html/draft-ietf-sidrops-aspa-profile

A helpful IETF presentation on ASPA algorithm accuracy:

K. Sriram and J. Heitz, "On the Accuracy of Algorithms for ASPA Based Route Leak Detection," IETF SIDROPS Meeting, Proceedings of the IETF 110, March 2021. https://datatracker.ietf.org/meeting/110/materials/slides-110-sidrops-sriram-aspa-alg-accuracy-01

Other IETF work related to route leak detection and mitigation:

"Route Leak Prevention and Detection Using Roles in UPDATE and OPEN Messages," IETF RFC 9234, May 2022. https://datatracker.ietf.org/doc/rfc9234/

“Methods for Detection and Mitigation of BGP Route Leaks,” https://datatracker.ietf.org/doc/draft-ietf-grow-route-leak-detection-mitigation/

ASPA: Autonomous System Provider Authorization

RPKI ASPA Object

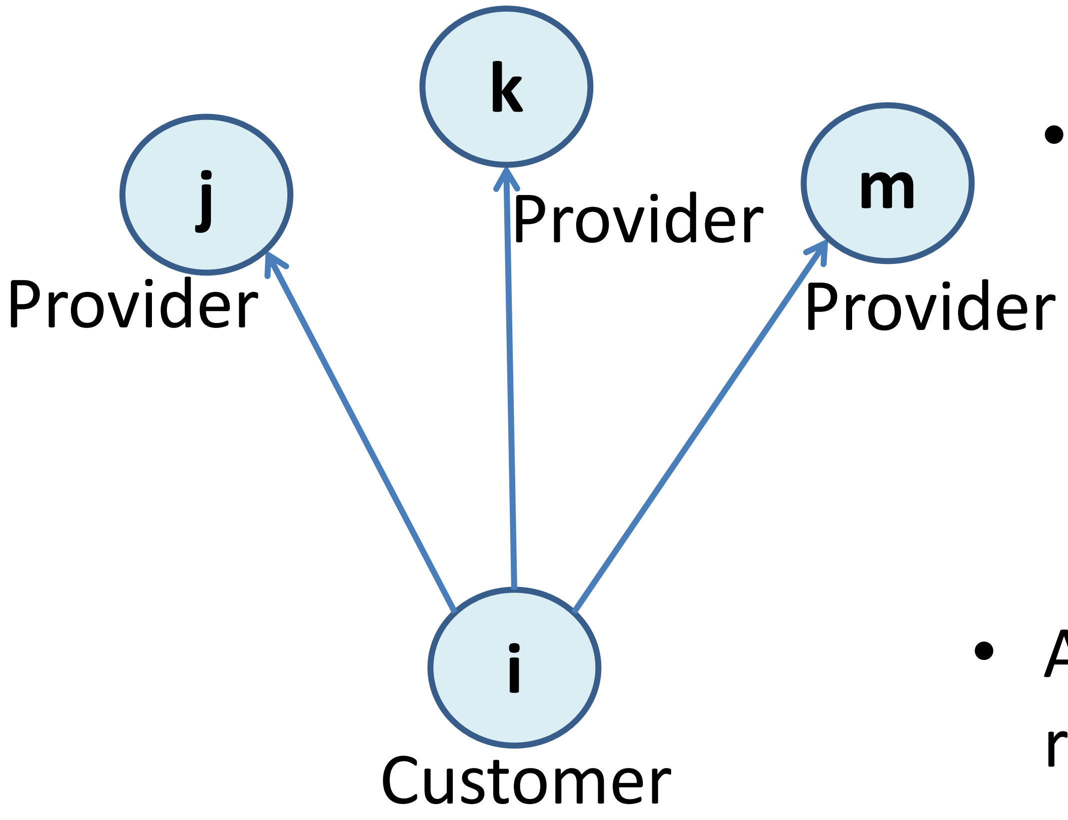

Example:

- ASPA: AS i, {AS j, AS k, AS m}

  transit providers

  AS i signs an ASPA object in the RPKI to attest that AS j, AS k, and AS m are transit providers

- ASPAs are registered/stored in the RPKI repositories

For details of ASPA registration requirements, see Section 4 in https://datatracker.ietf.org/doc/html/draft-ietf-sidrops-aspa-verification

BGP Roles and ASPAs

- Provider
- Customer
- Lateral peer
- IXP Route Server (RS)
- RS-client
- Mutual transit

- RS to RS-client relationship is like a provider to customer relationship. The RS AS is included in the RS-client AS's ASPA
- An AS having no providers registers an AS0 ASPA (i.e., ASPA containing only AS 0 as provider)
- Mutual transit ASes include each other in their ASPAs as provider

For details of ASPA registration requirements, see Section 4 in https://datatracker.ietf.org/doc/html/draft-ietf-sidrops-aspa-verification

ASPA's AS Path Anomaly Detection Capabilities

- Can detect and mitigate route leaks and improbable AS paths
- Can detect forged-origin prefix hijacks to some extent (slide 40)
- Can detect forged-path-segment prefix hijacks to some extent (slide 41)
- **Limitations**: ASPA method has limitations with regard to some forms of malicious AS path manipulations; mainly when a transit provider attacks its own customer with path manipulations (slide 43)

Route leaks involve one of four valley-free violations

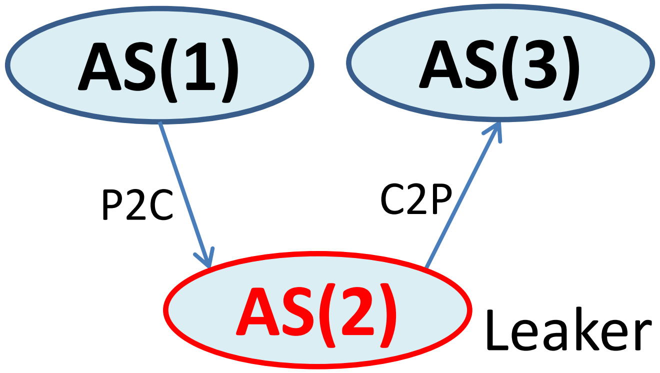

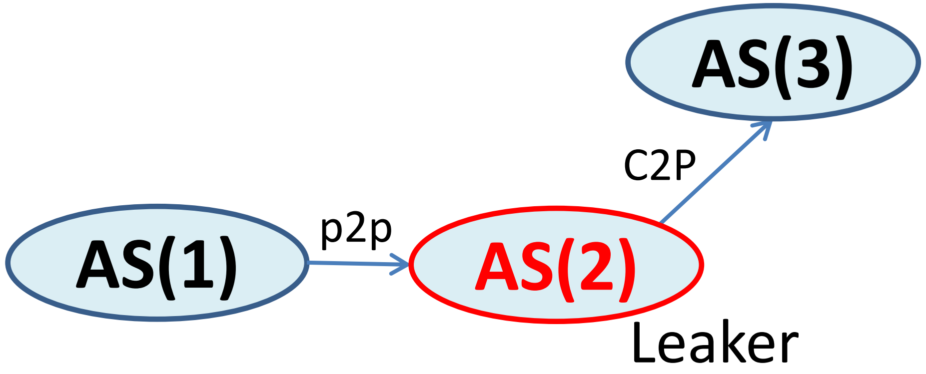

P2C = Provider to customer

C2P = Customer to Provider

p2p = lateral peers

ASPA: AS(1) {AS(5)}

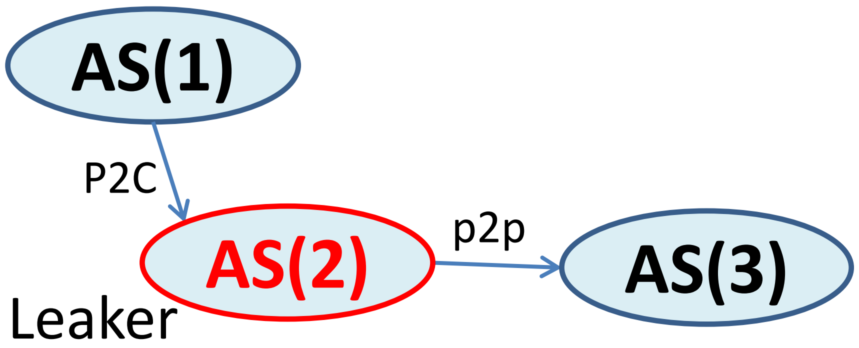

- Consider routes originated or propagated by AS(1) and received at AS(3)
- All four forms of route leaks are detected at AS(3) if AS(1) has ASPA

* Assume AS(2) is not removing AS(1) from the AS path (that then gets into the realm of AS path manipulation)

ASPA Hop Check Function

P: Provider
nP: not Provider
nA: no Attestation

Definition:

$$\text{hop(AS(i), AS(j))} = \begin{cases} \textbf{P} & \text{if AS(i) attests AS(j) is a provider} \\ \textbf{nP} & \text{if AS(i) attests AS(j) is not a provider} \\ \textbf{nA} & \text{if AS(i) does not have an ASPA} \end{cases}$$

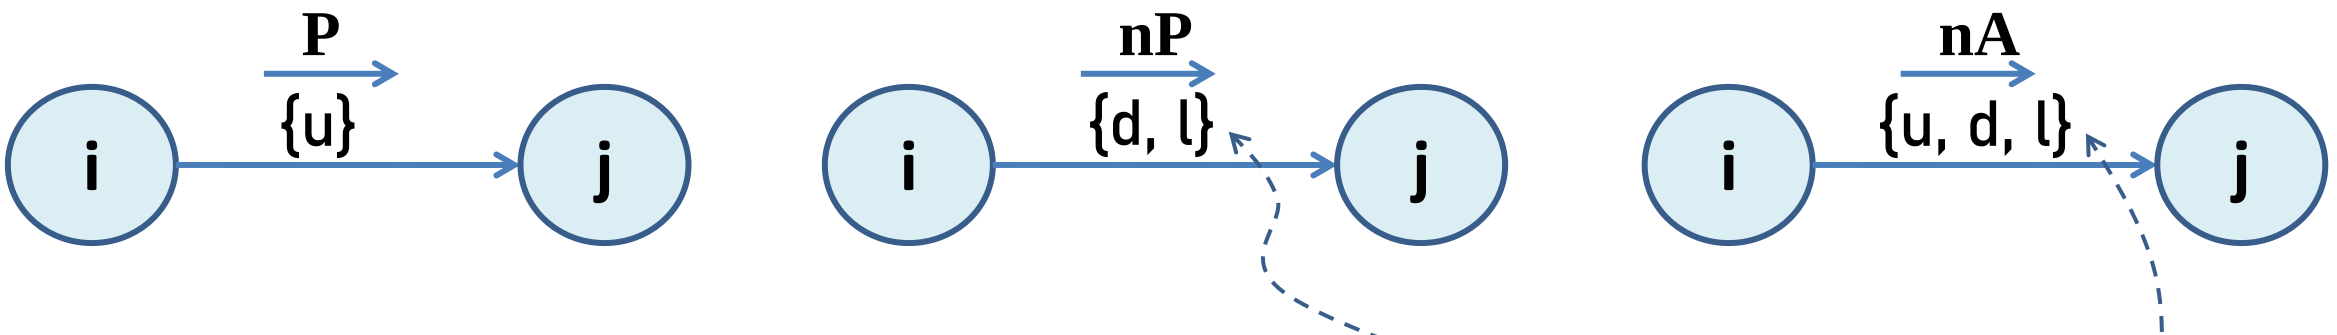

AS(i)-AS(j) peering:
u = Up (customer to provider (C2P))
d = Down (provider to customer (P2C))
l = Lateral (peer to peer (p2p))

allowed peering relations

A note about AS Path representation style

- We collapse the AS prepends. So, the AS path is represented by unique ASes such as AS(1), AS(2), …, AS(N).
- Thus AS(1) is the origin AS and AS(N) is the AS that is neighbor to the receiving/verifying AS.
- In the diagrams, for simplicity, we only show indices of ASes, i.e., AS positions. Do not mistake them for AS numbers.

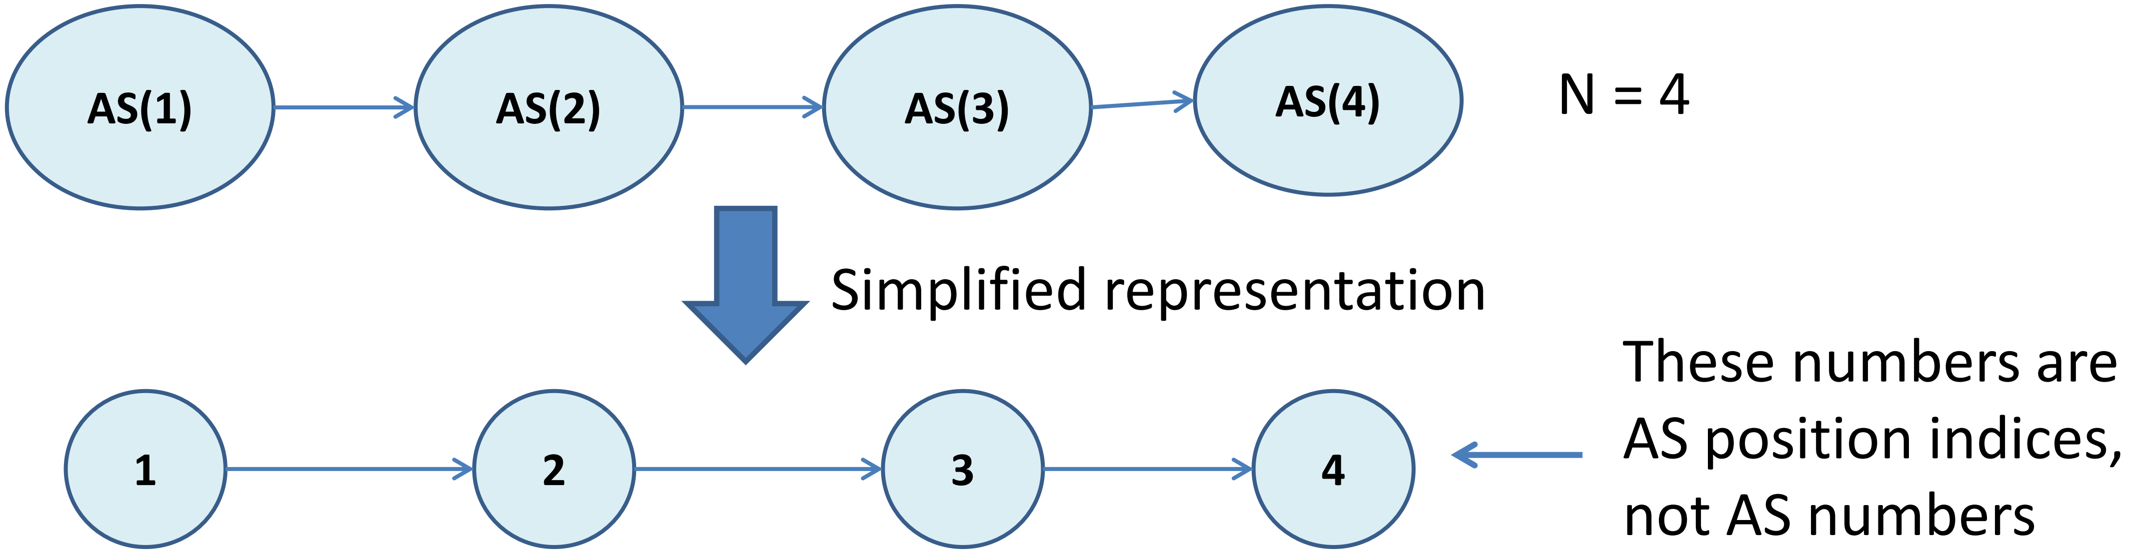

Example when Upstream AS Path is Valid

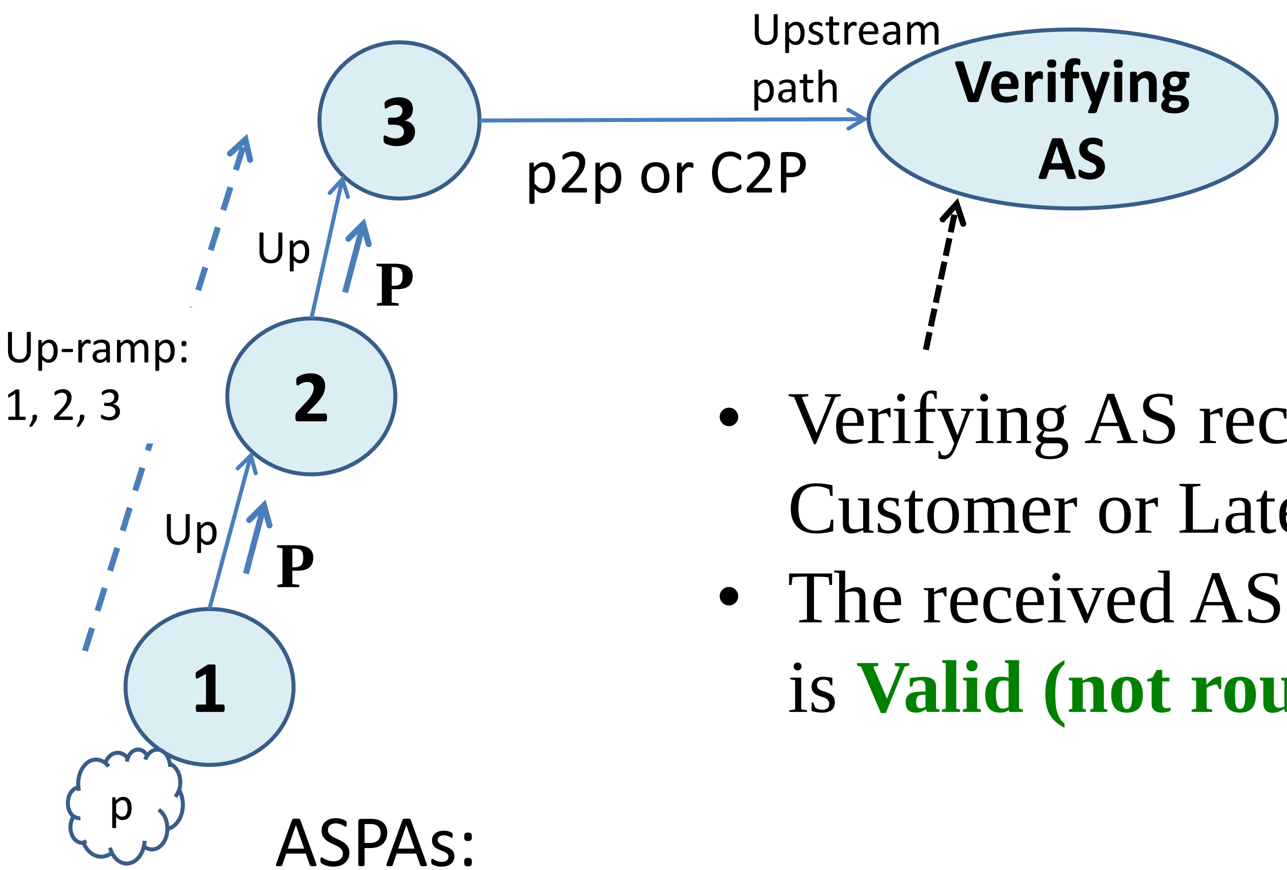

ASPA hop check:
P: Provider
nP: not Provider
nA: no Attestation

- Verifying AS receives the BGP route from a Customer or Lateral Peer;
- The received AS path {AS(3) AS(2) AS(1)} is **Valid (not route leak)**

ASPAs:
AS(1), {AS(2)}
AS(2), {AS(3)}

Example when Upstream AS Path is Invalid

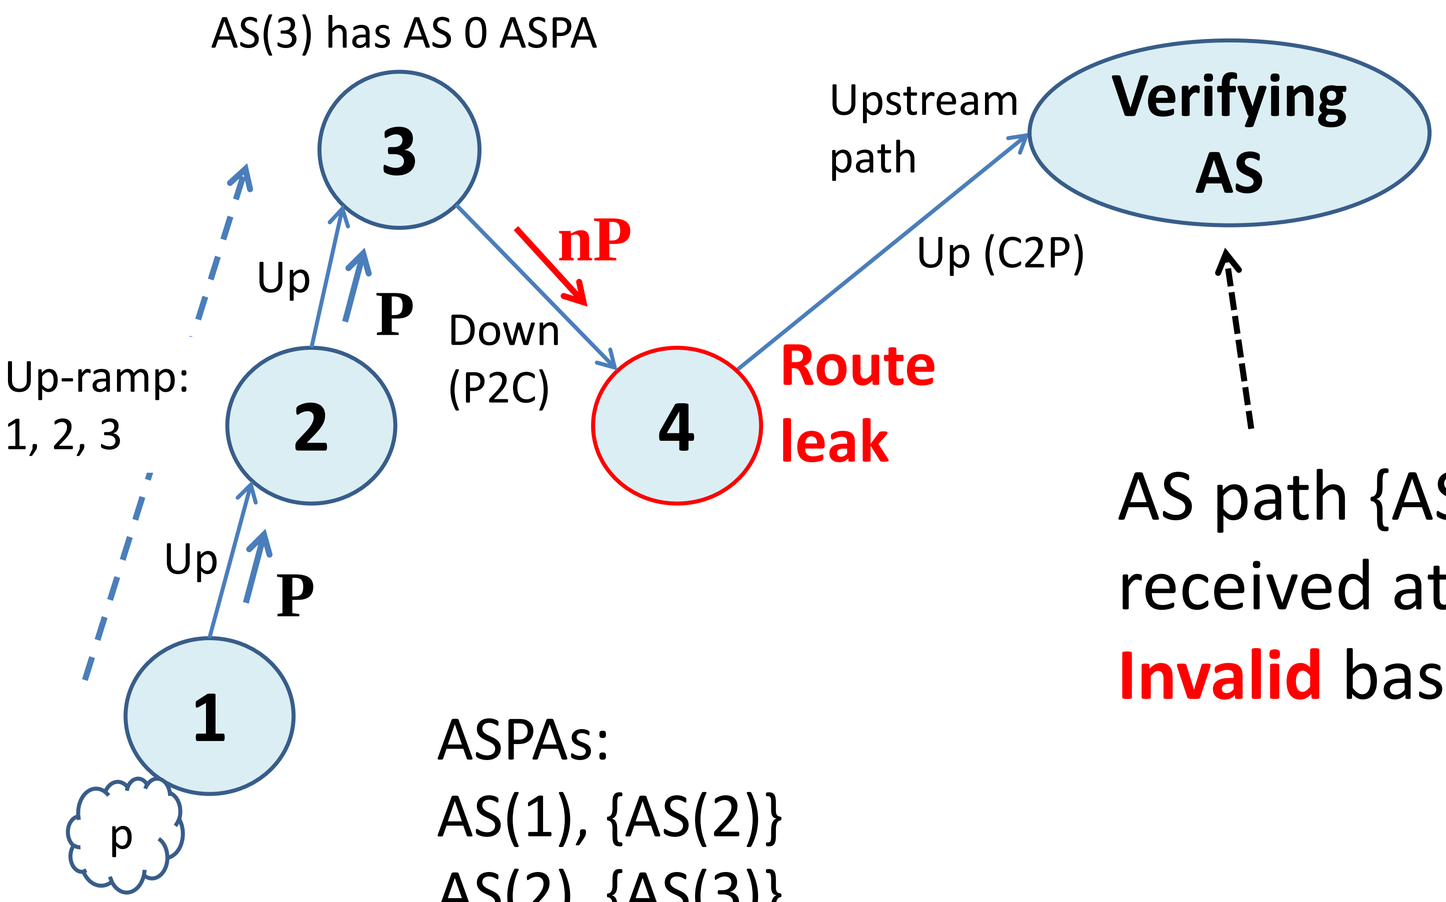

ASPA hop check:
P: Provider
nP: not Provider
nA: no Attestation

AS path {AS(4) AS(3) AS(2) AS(1)} received at the Verifying AS is **Invalid** based on ASPAs

ASPAs:
AS(1), {AS(2)}
AS(2), {AS(3)}
AS(3), {AS 0}
AS(4) is the leaker and may / may not have an ASPA…
the Verifying AS has local knowledge that AS(4) is its customer

Algorithm for Upstream AS Path Verification

- If the hop() function for each hop in the AS path is **P**, the AS path is Valid (not route leak) and return.
- Else, if the hop() function for any hop in the AS path is **nP**, the AS path is Invalid (route leak) and return.
- Else, the AS path validity is Unknown (may or may not be a route leak) and return.

ASPA Verification of Downstream AS Path: Invalid Outcome

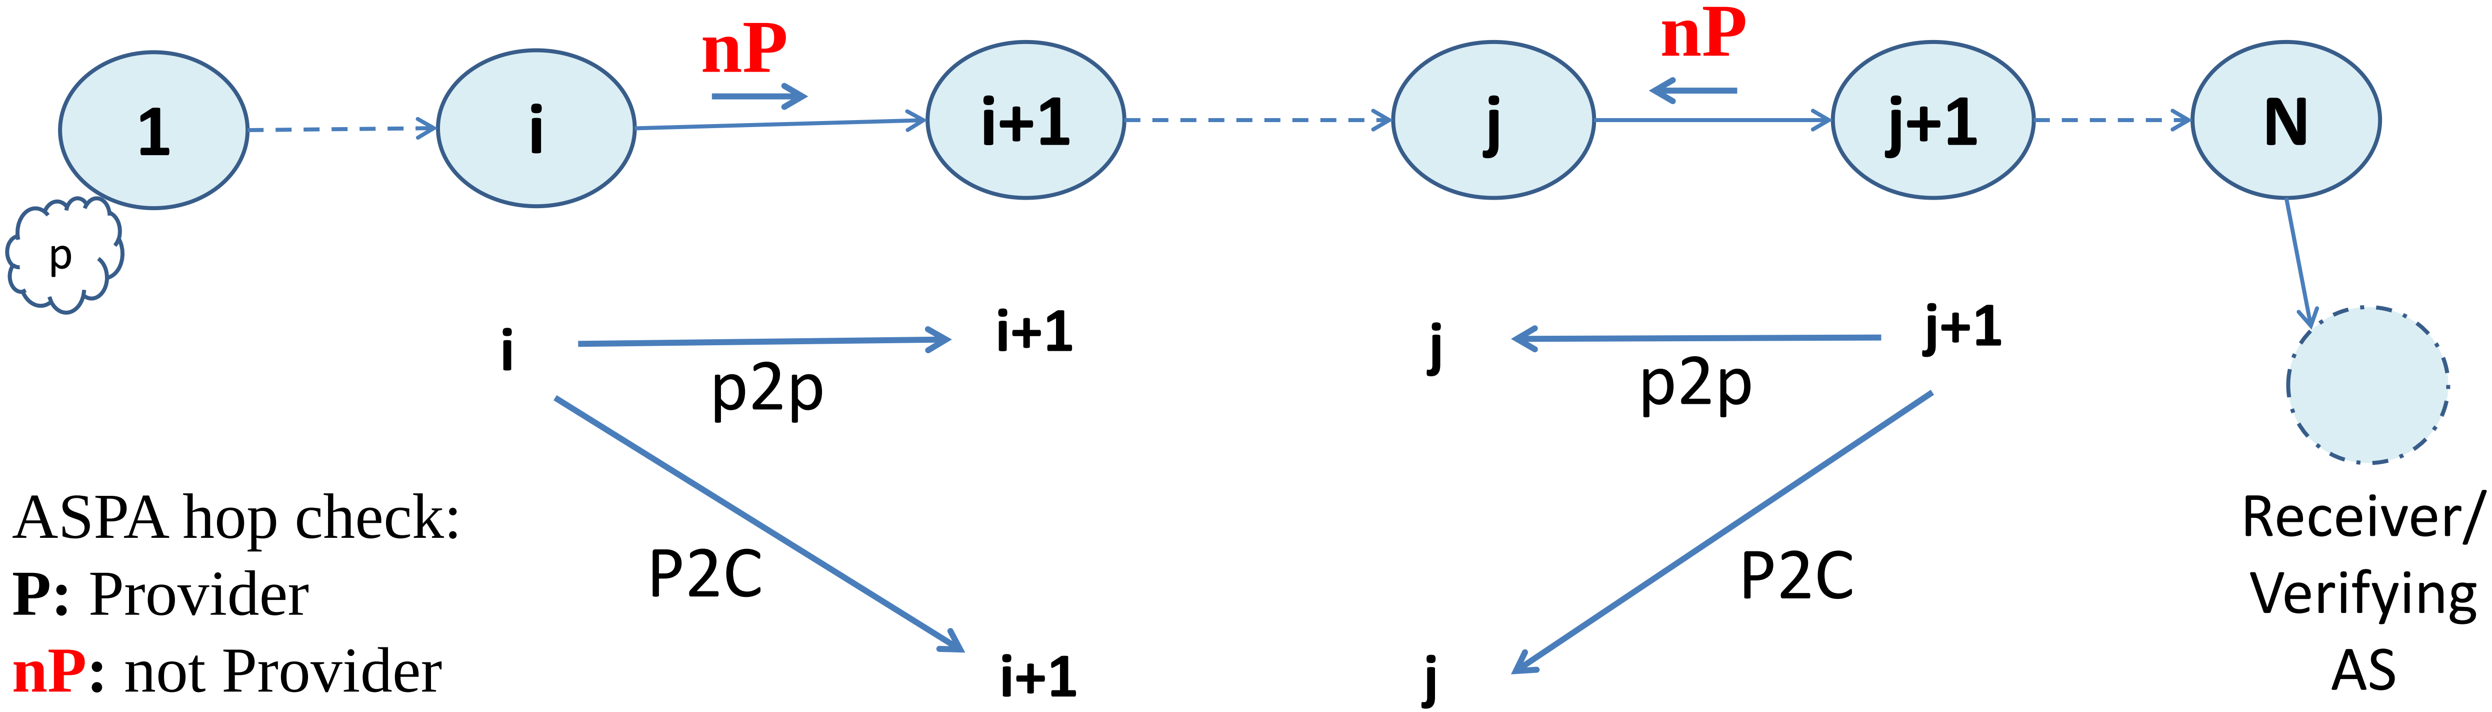

ASPA hop check:
P: Provider
nP: not Provider
nA: no Attestation

Any two hops in opposite directions are **nP** per ASPA (j > i)
(facing each other)

ASPA Verification of Downstream AS Path: Valid Outcome

The only permissible path trajectories for **Valid** outcome are an inverted V or inverted V with a one hop p2p at the apex

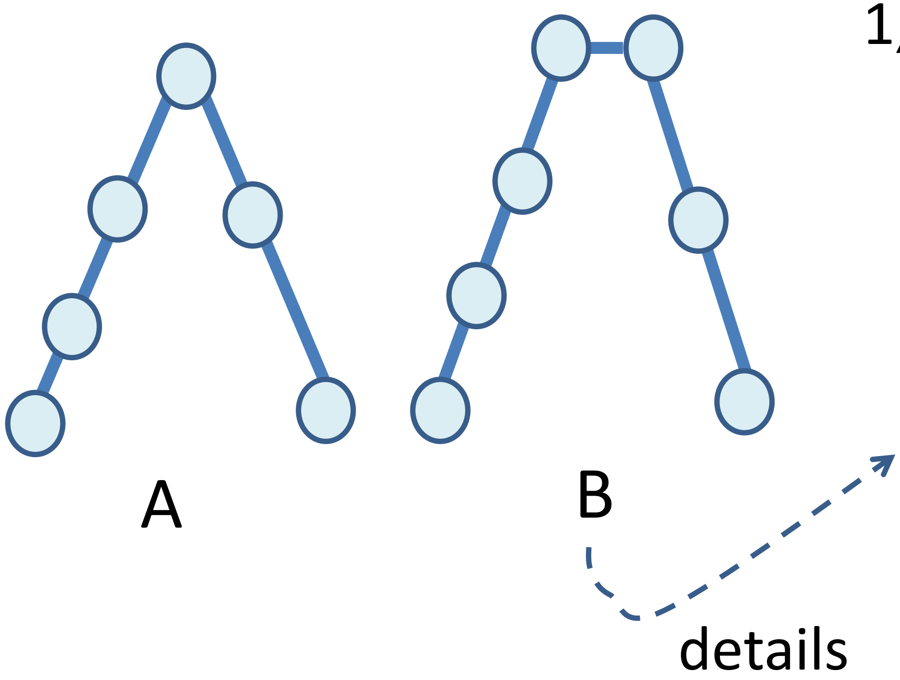

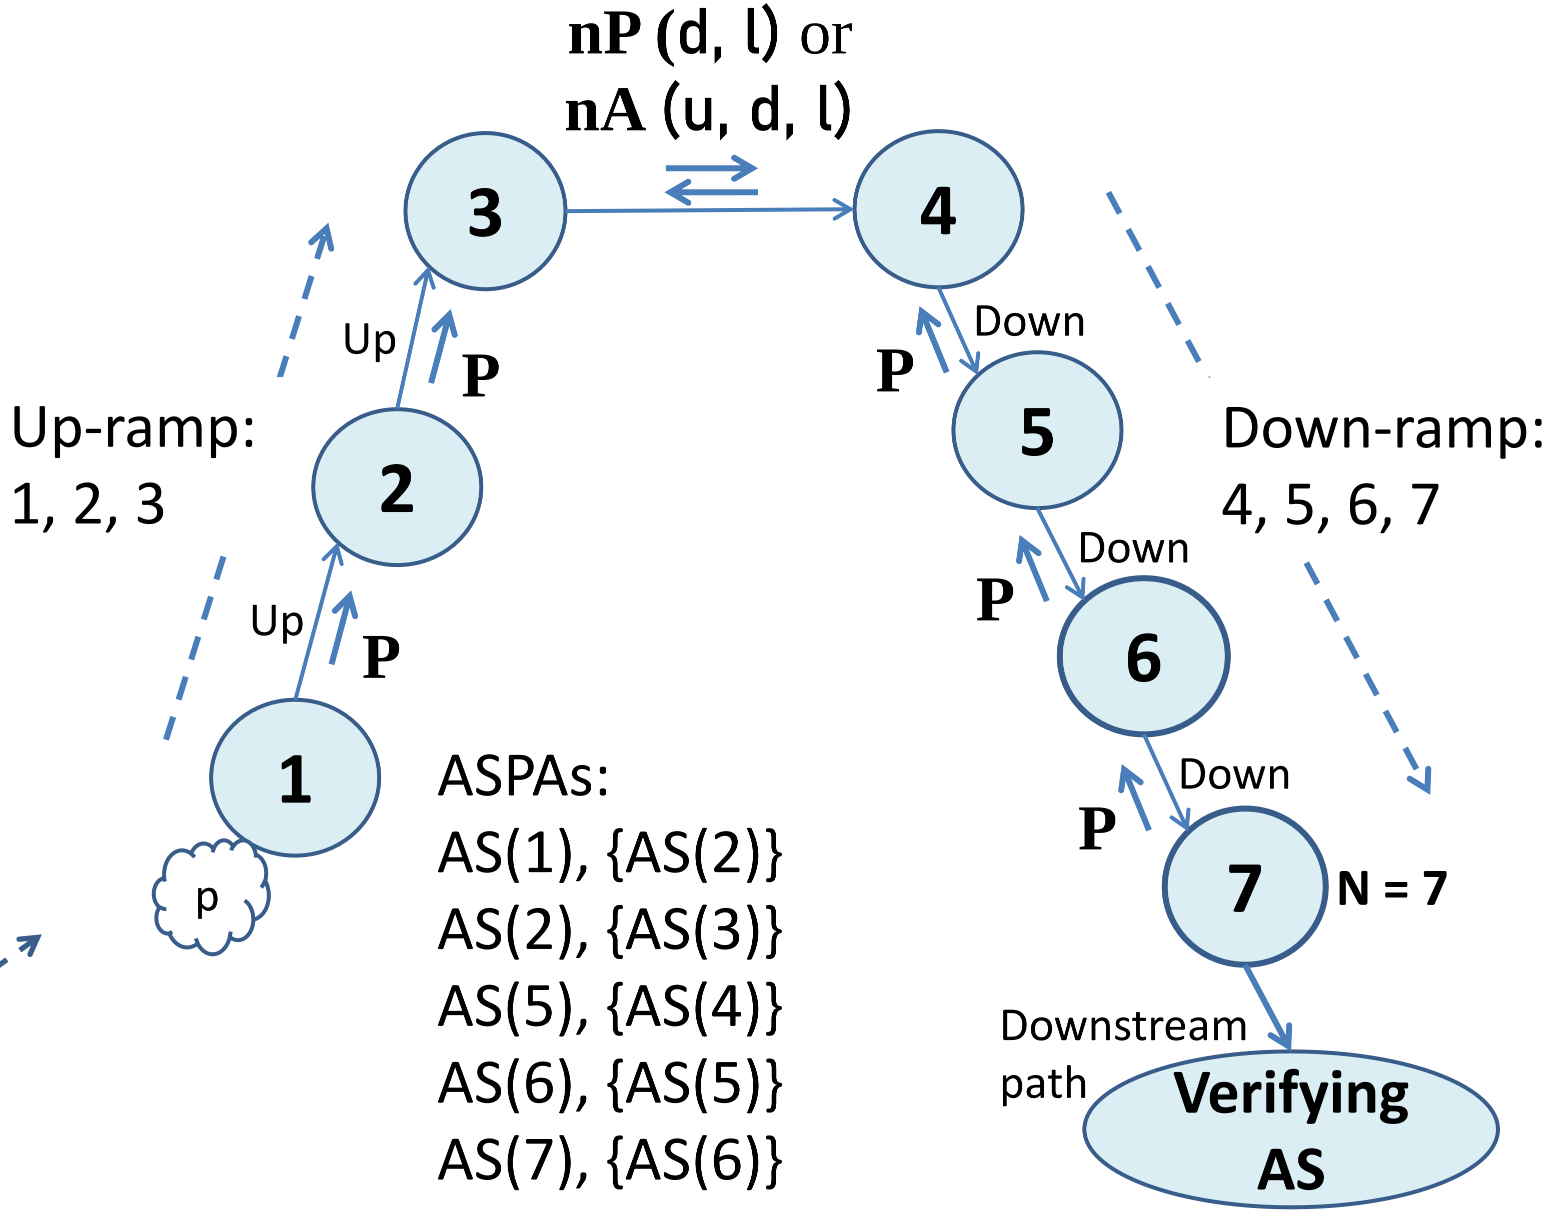

ASPA Verification of Downstream AS Path: Unknown Outcome

In partial deployment, an Unknown outcome occurs when the available ASPA's do not produce an Invalid (slide 25) or Valid (slide 26) outcome for the AS_PATH.

ASPA Verification of Downstream AS Path: Formal Algorithm Development

ASPA Verification of Downstream AS Path

(K, L) representation of downstream AS path

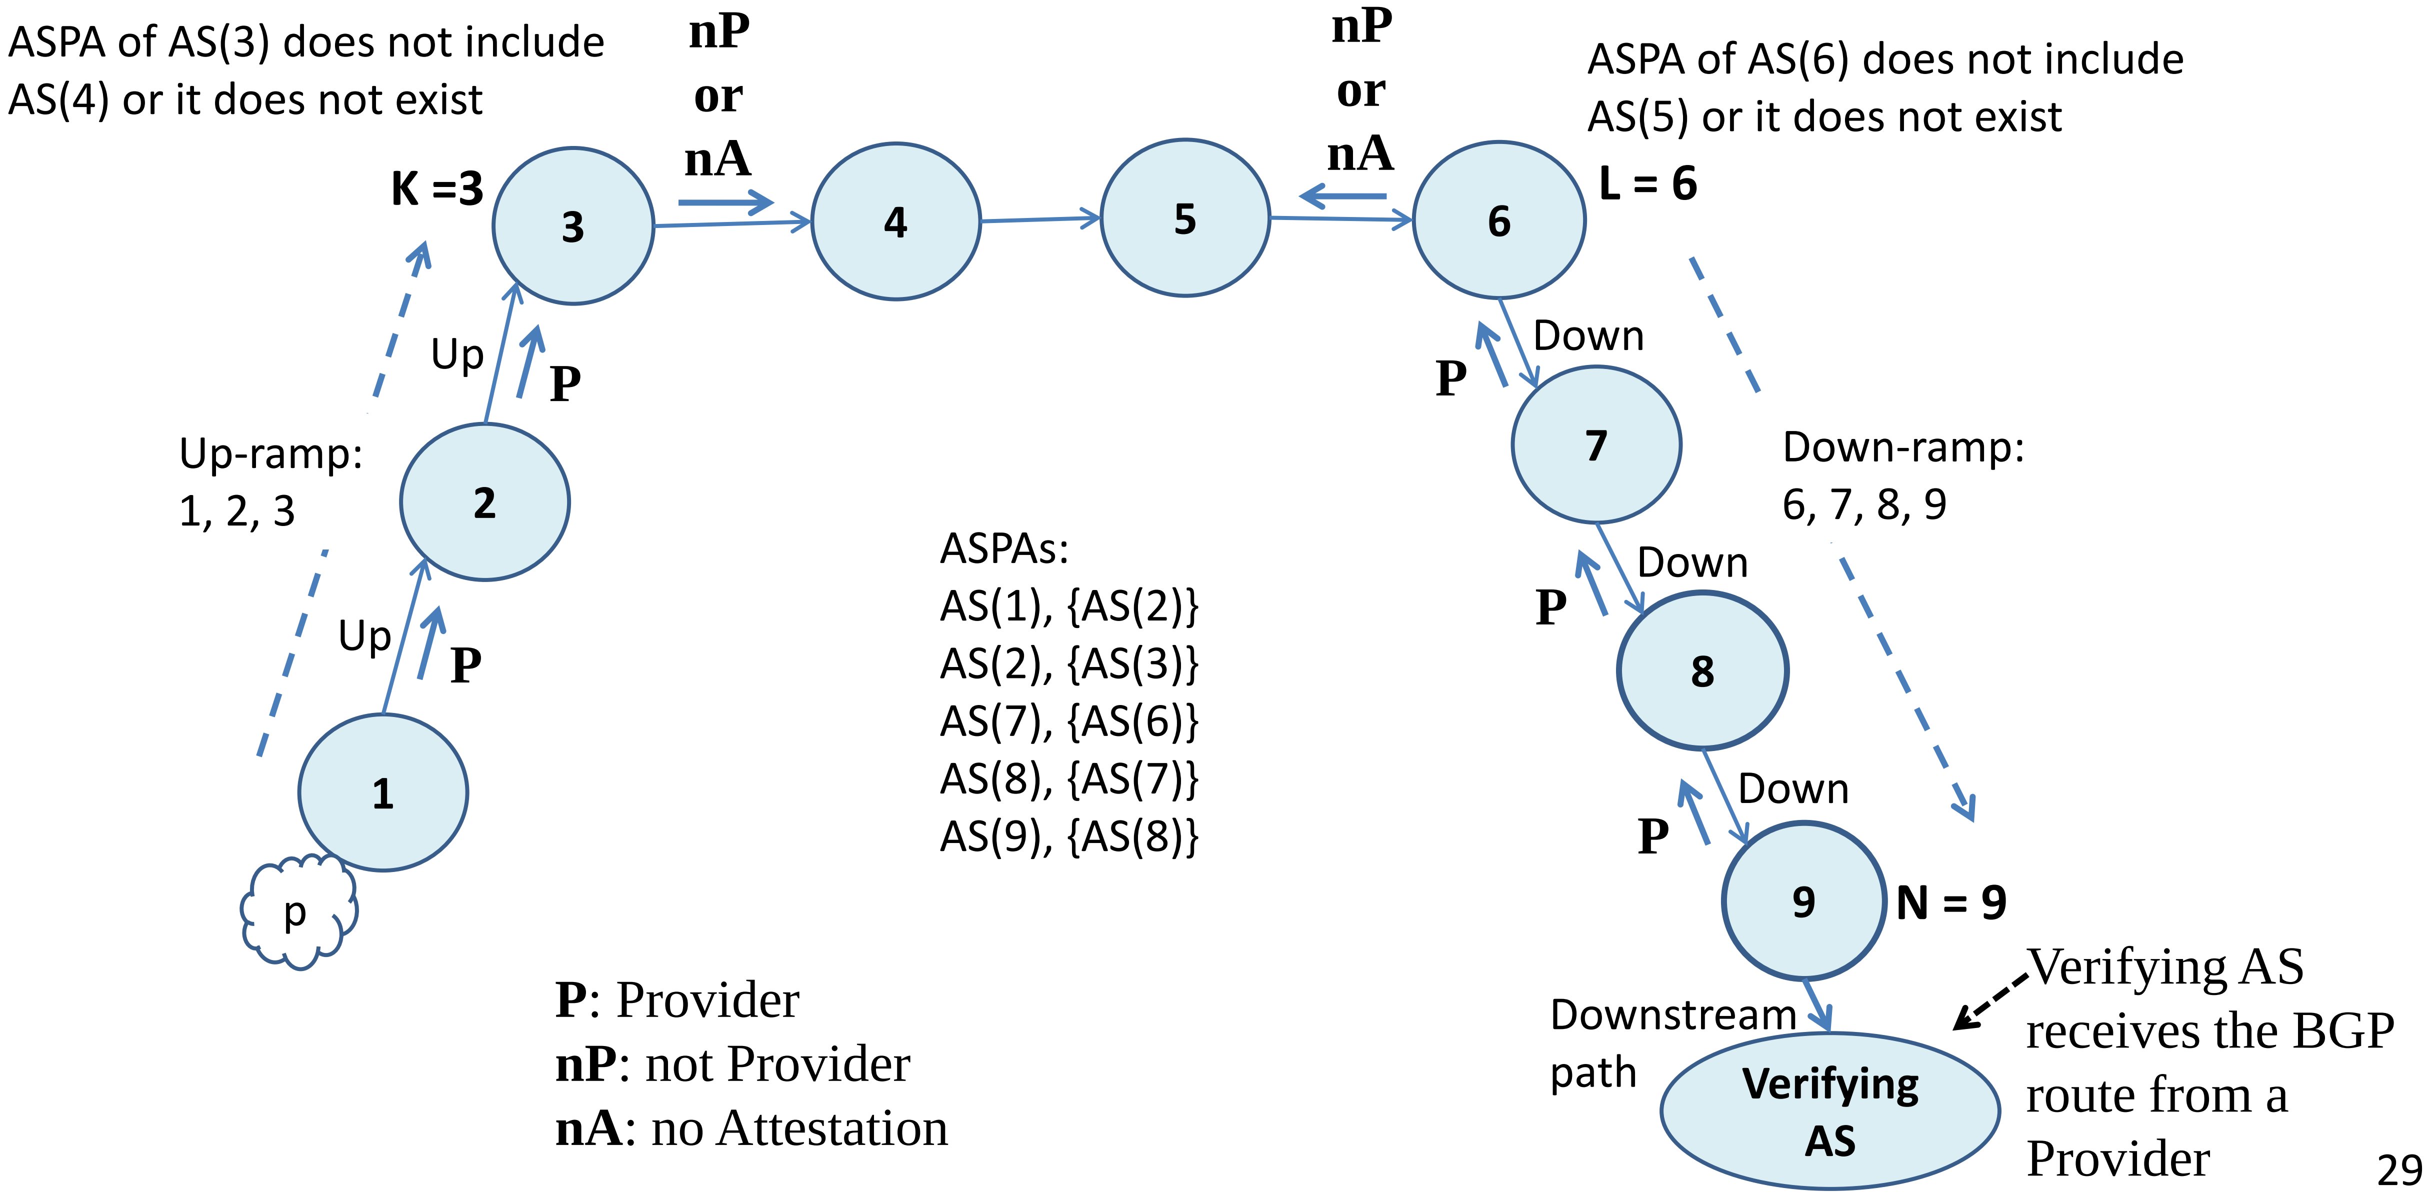

Valid downstream AS path when L – K ≤ 1

ASPA hop check:
P: Provider
nP: not Provider
nA: no Attestation

AS-AS peering:
u = Up
d = Down
l = Lateral

ASPAs:
AS(1), {AS(2)}
AS(2), {AS(3)}
AS(5), {AS(4)}
AS(6), {AS(5)}
AS(7), {AS(6)}

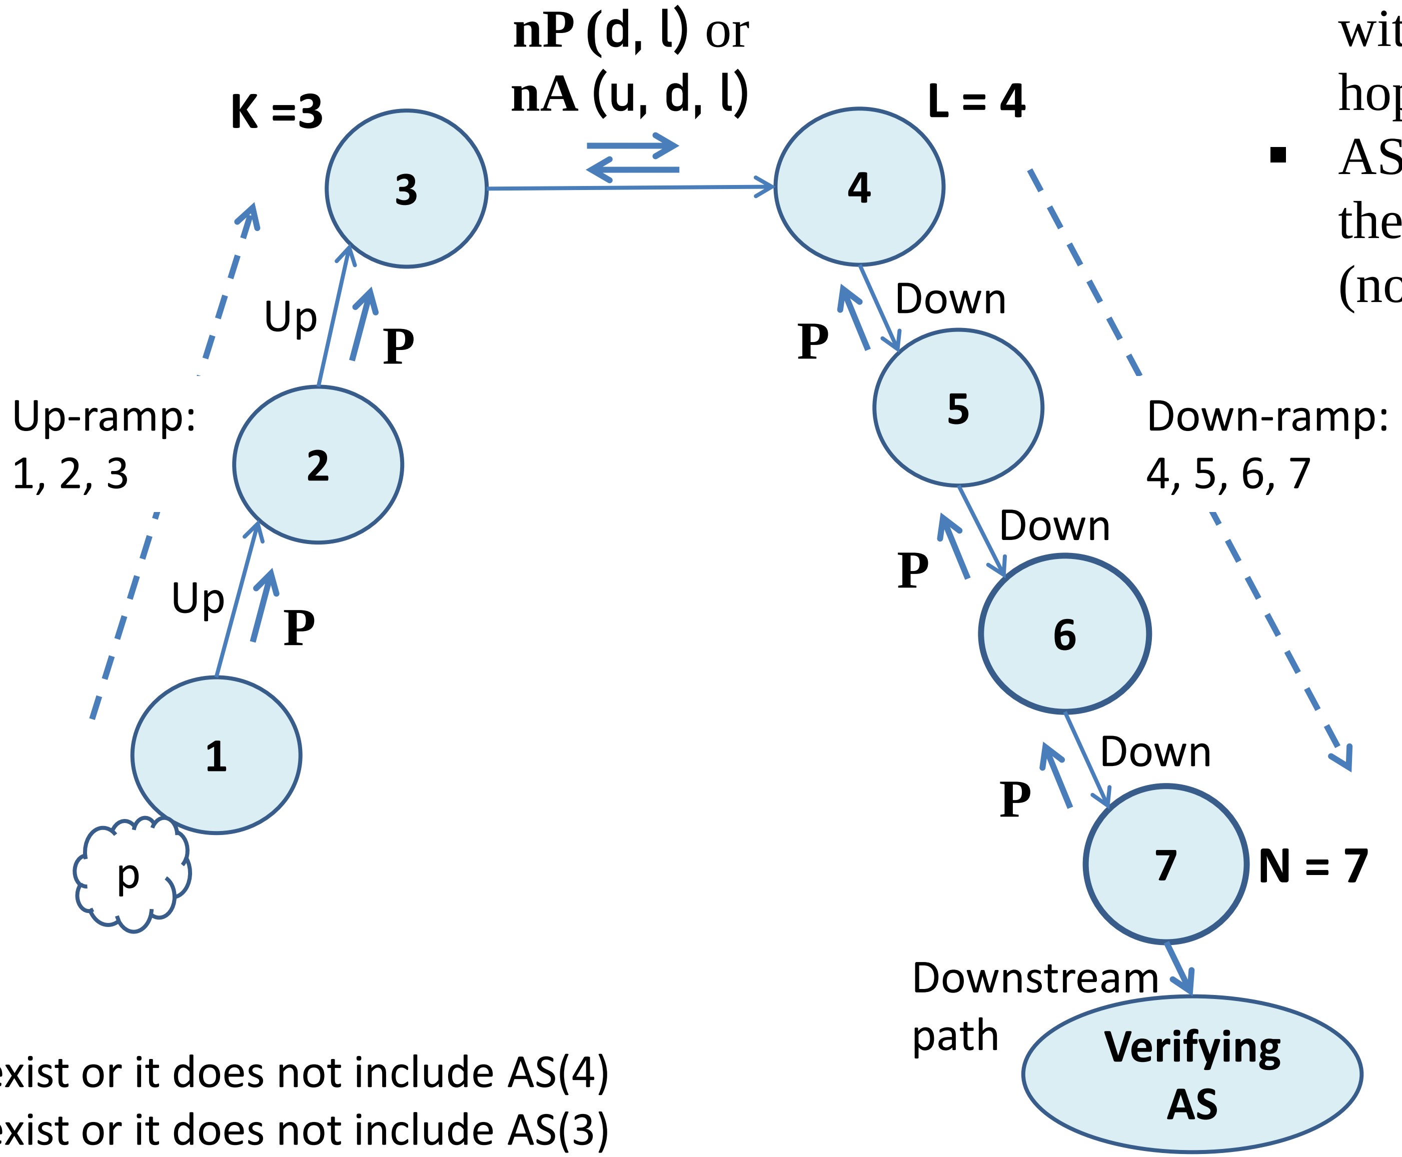

- The AS path is Valid with/without the **nP** or **nA** hop in the middle
- AS path is trivially Valid if the AS path length is 1 or 2 (no ASPA needed)

ASPA of AS(3) does not exist or it does not include AS(4)
ASPA of AS(4) does not exist or it does not include AS(3)

For L-K ≥ 2, only Invalid or Unknown are possible

Illustration for L-K = 2

ASPA hop check:
P: Provider
nP: not Provider
nA: no Attestation

AS-AS peering:
u = Up
d = Down
l = Lateral

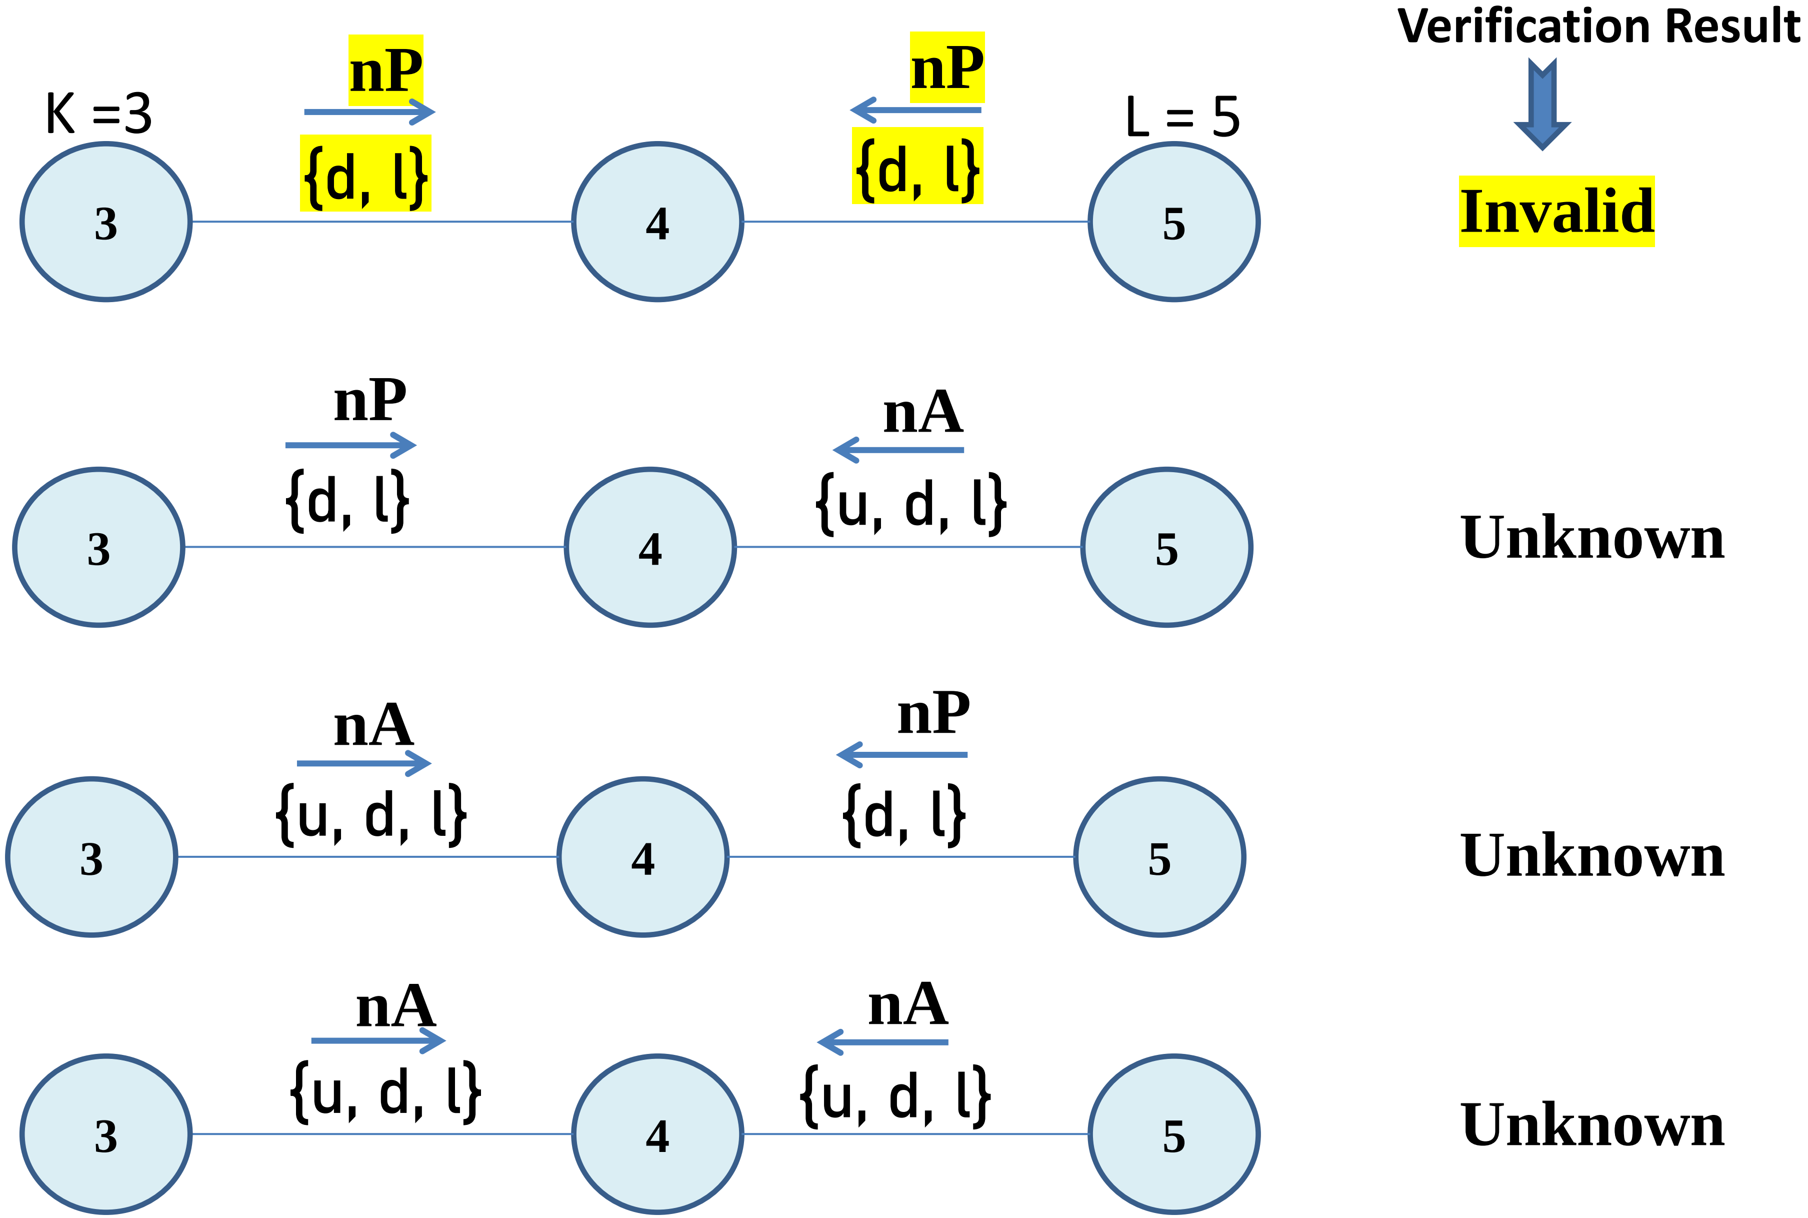

Arrows indicate direction of ASPA hop check

Theorems that help design the algorithm

Theorem 1: The downstream AS path is Valid if and only if L-K $\leq$ 1. If L-K $\geq$ 2, then the AS path can be Unknown or Invalid, but never Valid.

Theorem 2: For L-K $\geq$ 2, the validity of the whole AS path is the same as that of the partial path AS(K), AS(K+1), …., AS(L-1), AS(L). The partial path can only be either Invalid or Unknown. It is Invalid if there exist u and v (u and v in the range from K+1 to L-1) such that u $\leq$ v and hop(AS(u-1), AS(u)) = **nP** and hop(AS(v+1), AS(v)) = **nP**. Otherwise, the partial path is Unknown.

Function hop() is defined on slide 20.
Proofs exist; see next two slides; also see reference [1] below.

[1] K. Sriram and J. Heitz, "On the Accuracy of Algorithms for ASPA Based Route Leak Detection," IETF SIDROPS Meeting, Proceedings of the IETF 110, March 2021. https://datatracker.ietf.org/meeting/110/materials/slides-110-sidrops-sriram-aspa-alg-accuracy-01

Proof: For L-K ≥ 2, only Invalid or Unknown are possible

ASPA hop check:
P: Provider
nP: not Provider
nA: no Attestation

AS-AS peering:
u = Up
d = Down
l = Lateral

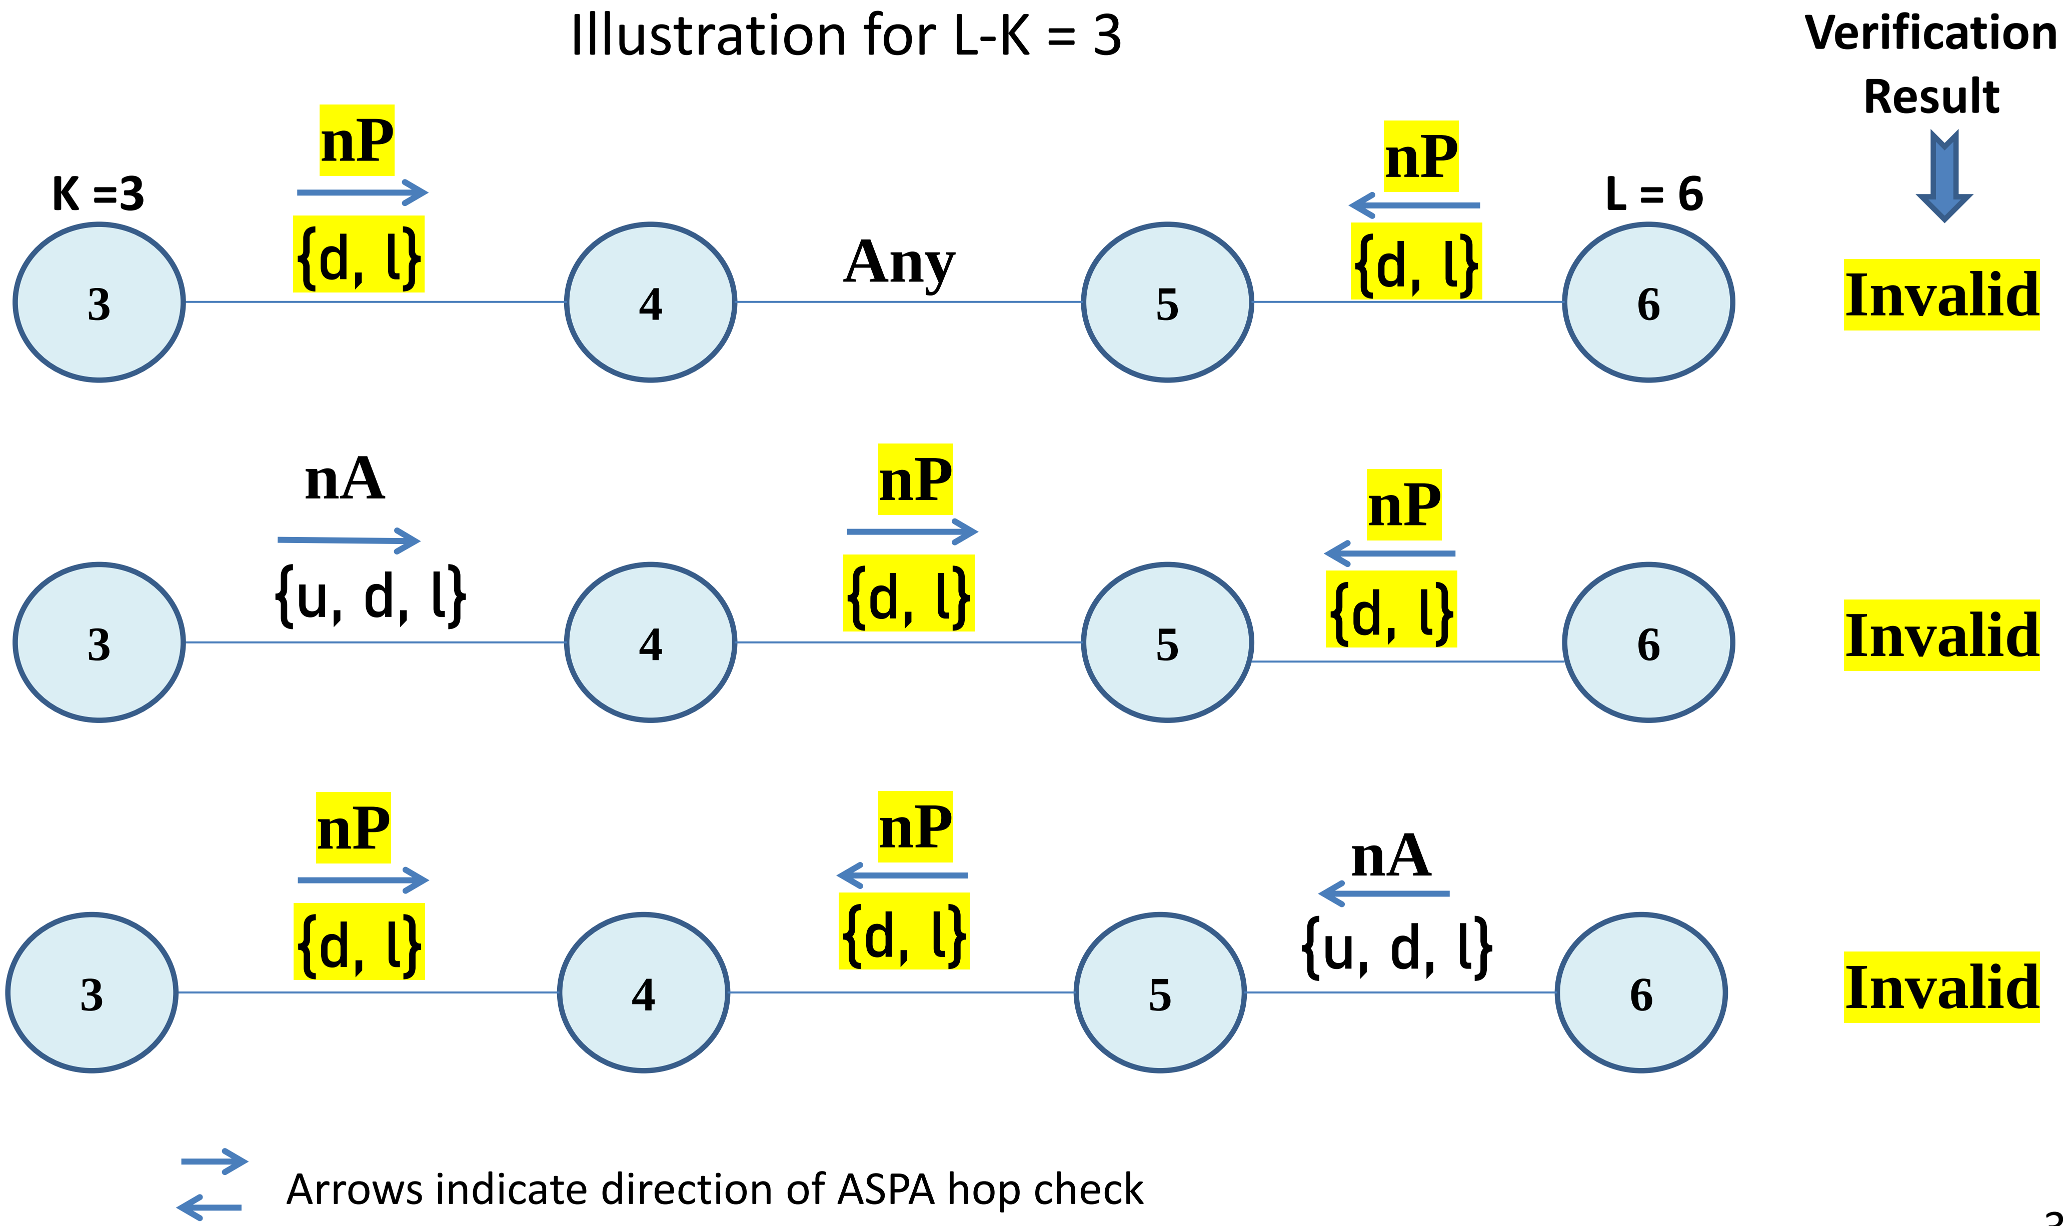

Proof: For L-K ≥ 2, only Invalid or Unknown are possible

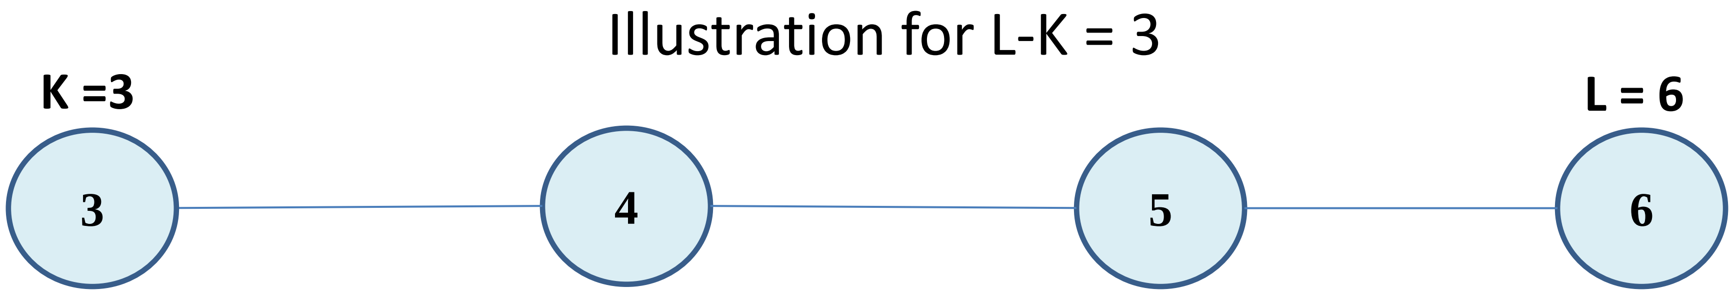

ASPA hop check:
P: Provider
nP: not Provider
nA: no Attestation

AS-AS peering:
u = Up
d = Down
l = Lateral

→ ← Arrows indicate direction of ASPA hop check

Hop 3-4	Hop 4-5	Hop 5-6	AS path
→ **nP** {d, l}	Any: **P, nP, or nA**	← **nP** {d, l}	Invalid
→ **nP** {d, l}	← **nP** {d, l}	← **nA** {u, d, l}	Invalid
→ **nP** {d, l}	← **nA** {u, d, l}	← **nA** {u, d, l}	Unknown
→ **nP** {d, l}	← **P** {u}	← **nA** {u, d, l}	Unknown
→ **nA** {u, d, l}	→ **nP** {d, l}	← **nP** {d, l}	Invalid
→ **nA** {u, d, l}	→ **nP** {d, l}	← **nA** {u, d, l}	Unknown
→ **nA** {u, d, l}	→ **nA** {u, d, l}	← **nP** {d, l}	Unknown
→ **nA** {u, d, l}	→ **nA** {u, d, l}	← **nA** {u, d, l}	Unknown
→ **nA** {u, d, l}	→ **P** {u}	← **nP** {d, l}	Unknown
→ **nA** {u, d, l}	→ **P** {u}	← **nA** {u, d, l}	Unknown

Algorithm for Downstream AS Path Verification

Crisp Description

1. If the AS path length $1 \leq N \leq 2$, then the path is trivially Valid and the procedure halts.
2. Else, now $N \geq 3$. Formulate the AS path (unique ASes) using the (K, L) representation (slide 29).

   If $L-K \leq 1$, then the AS path is Valid and the procedure halts.

   (Note: For $L-K \geq 2$, to determine whether the AS path is Invalid or Unknown, we only need to focus on the portion of the path from AS(K) to AS(L).)
3. Else, now $L-K \geq 2$.

   Consider the partial path represented by AS(K), AS(K+1), …, AS(L-1), AS(L).

   If there exist u and v in the range from K+1 to L-1 such that $u \leq v$ and
   hop(AS(u-1), A(u)) = **nP**, and
   hop(AS(v+1), A(v)) = **nP**,
   then the AS path is Invalid and the procedure halts.
4. Else, the AS path is Unknown and the procedure halts.

The detailed algorithm is in the IETF draft (ietf-sidrops-aspa-verification)

Prevention of Route Leaks at Local AS

- RFC 9234: Only to Customer (OTC) Attribute
- Add the OTC Attribute on eBGP ingress (if not already present) when a route is received from a Provider, IXP Route Server, or Lateral Peer
- If the OTC Attribute is present, do not propagate the route to a Provider, IXP Route Server, or Lateral Peer at eBGP egress
- If the OTC Attribute is not present, the route may be propagated to any type of peer at eBGP egress

ASPA Path Verification: Highlighting Some Key Properties

- These properties are early adoption incentives
- For the key properties descriptions (next 5 slides), assume that malicious AS path manipulations are not involved, especially removal of certain ASes from the AS path.
- An example of ASPA's limitation with regard to malicious AS path manipulation is on slide 43

ASPA Path Verification: Property 1

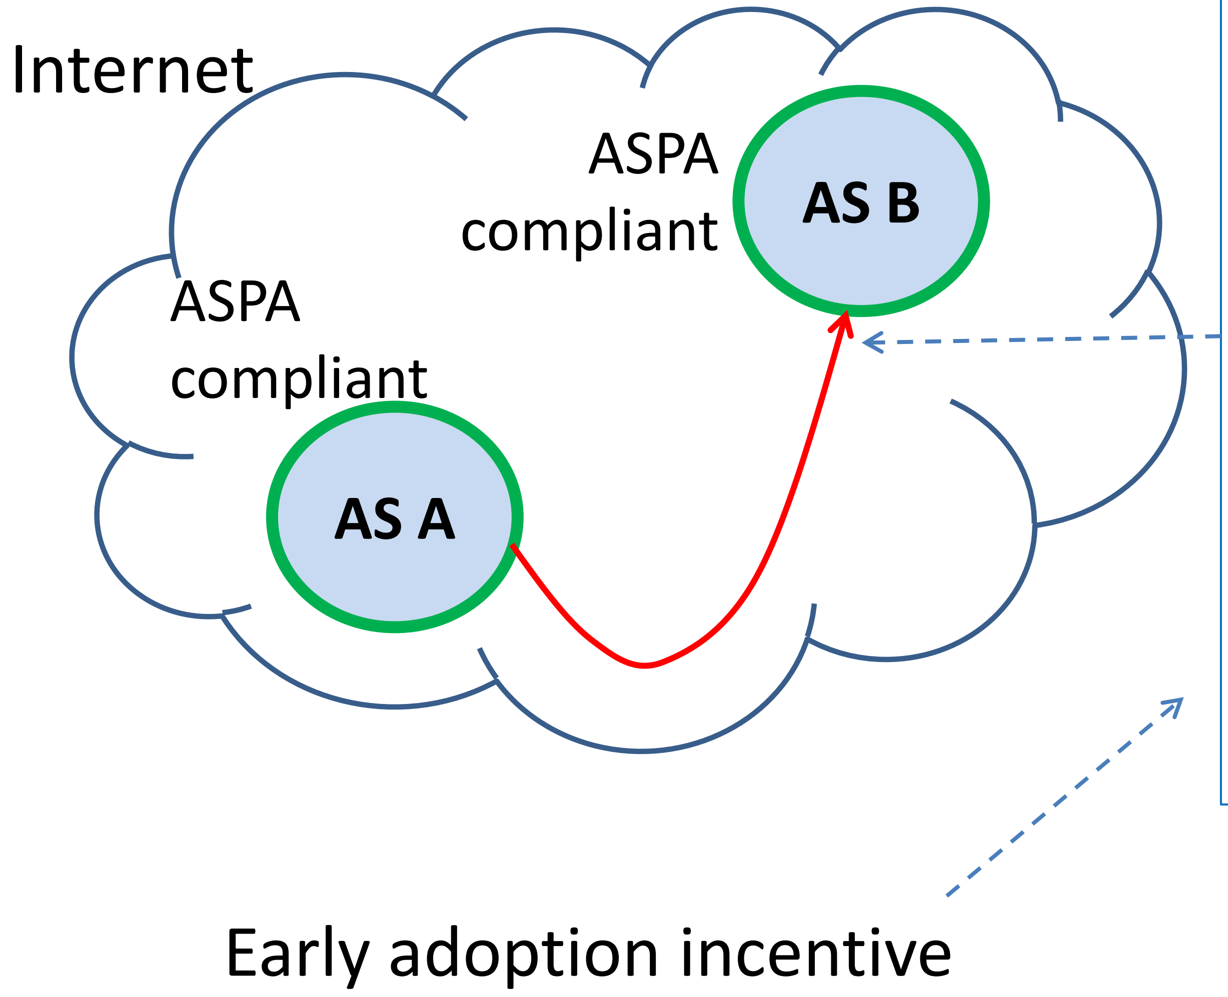

- Only two ASes A and B are doing ASPA
- AS A propagates a route to a customer or lateral peer
- AS B receives the route from a customer or lateral peer
- If the AS_PATH involves a route leak, it is always detected and mitigated at AS B

Corollary of Property 1

- In effect, if most major ISPs are ASPA compliant, the propagation of route leaks in the Internet will be severely limited.

ASPA Path Verification: Property 2

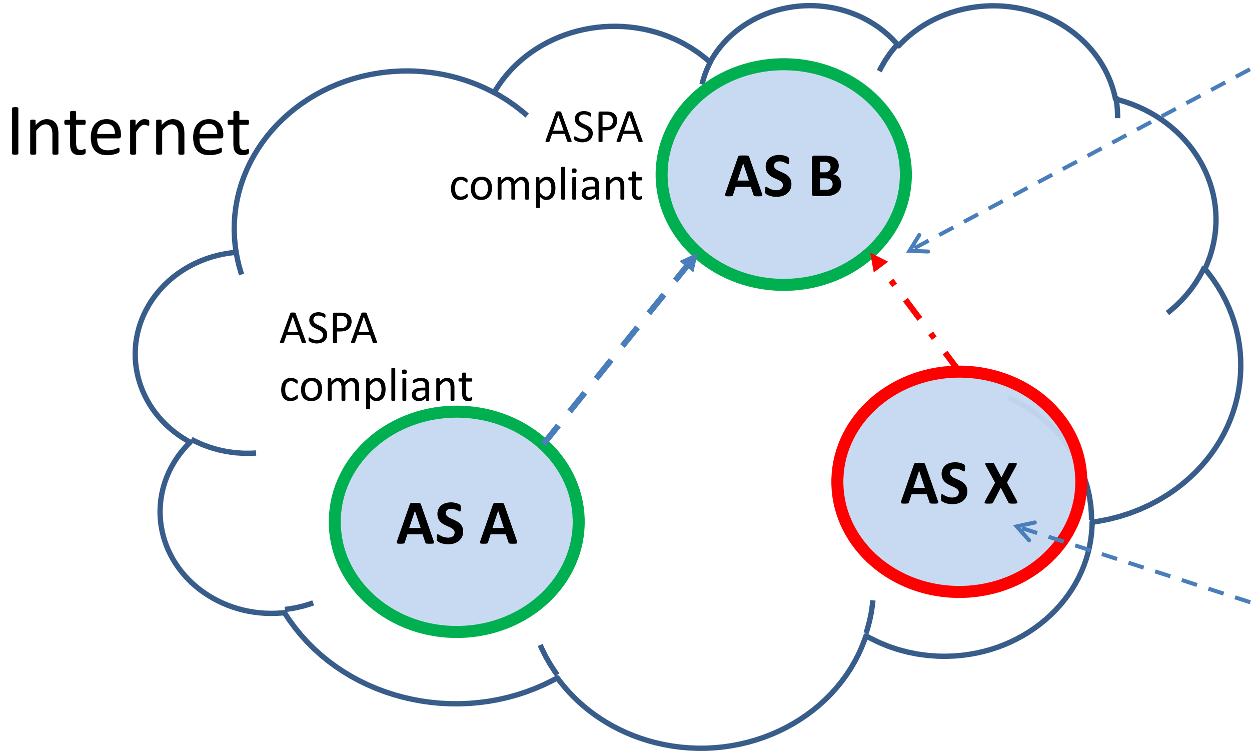

The forged-origin prefix hijack attack involving AS A is detected and mitigated at AS B

Conducts forged-origin prefix hijack involving AS A as the origin

- Only two ASes A and B are doing ASPA and ROA/ROV
- AS B receives the forged route sent by AS X (attacker) in the upstream direction (from a customer or lateral peer)

ASPA Path Verification: Property 3

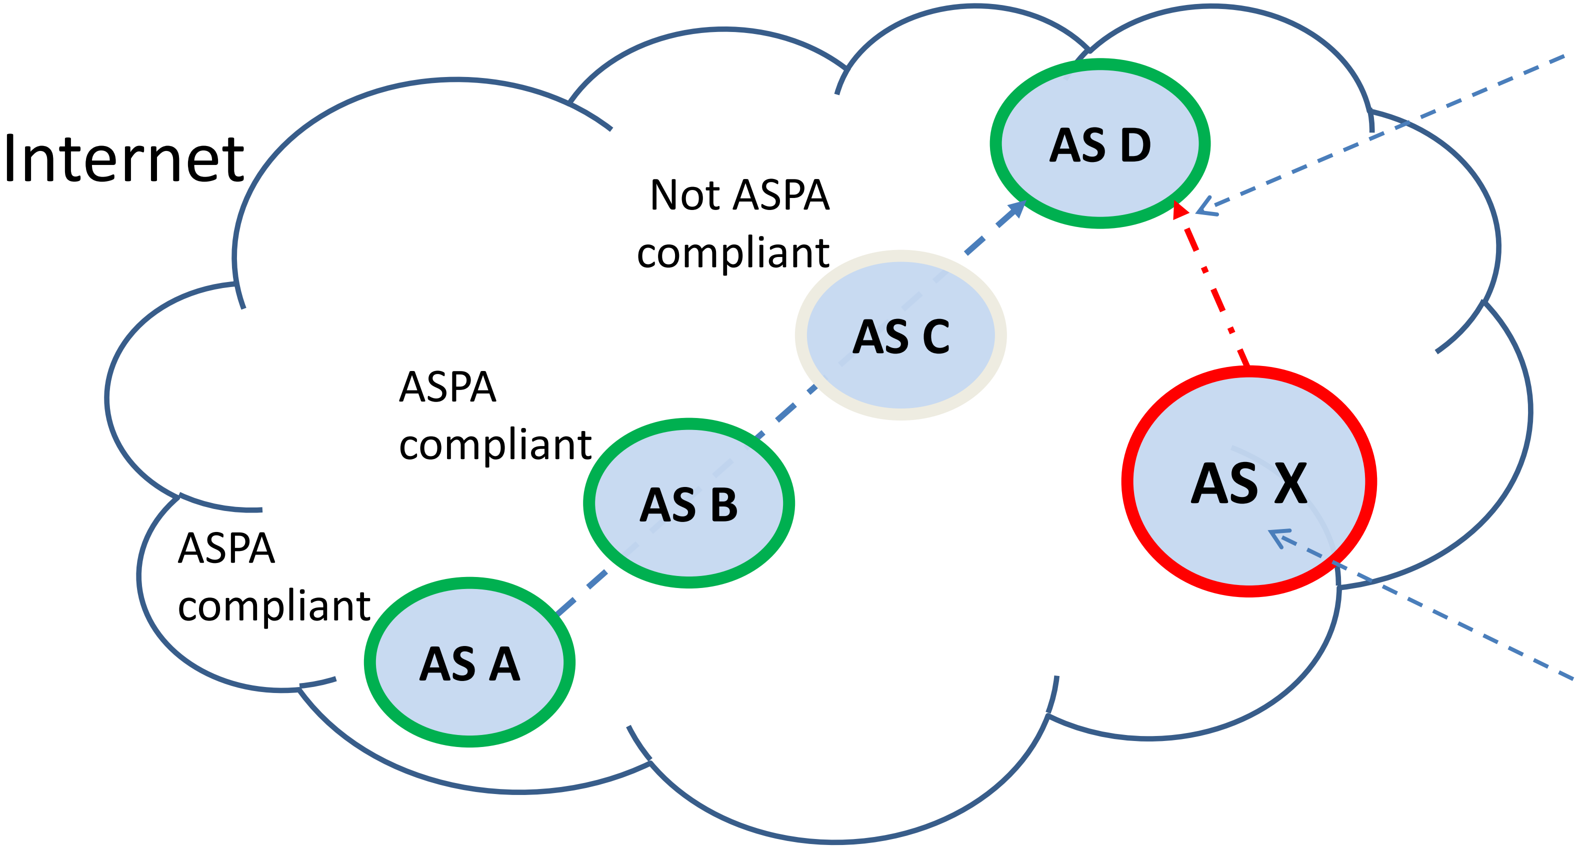

The prefix hijack with forged-path-segment involving {AS B, AS A} is detected and mitigated at AS D

Conducts a prefix hijack with forged-path-segment involving {AS B, AS A}

- AS B receives the forged route sent by AS X (attacker) in the upstream direction (from a customer or lateral peer)

ASPA Path Verification: Property 4

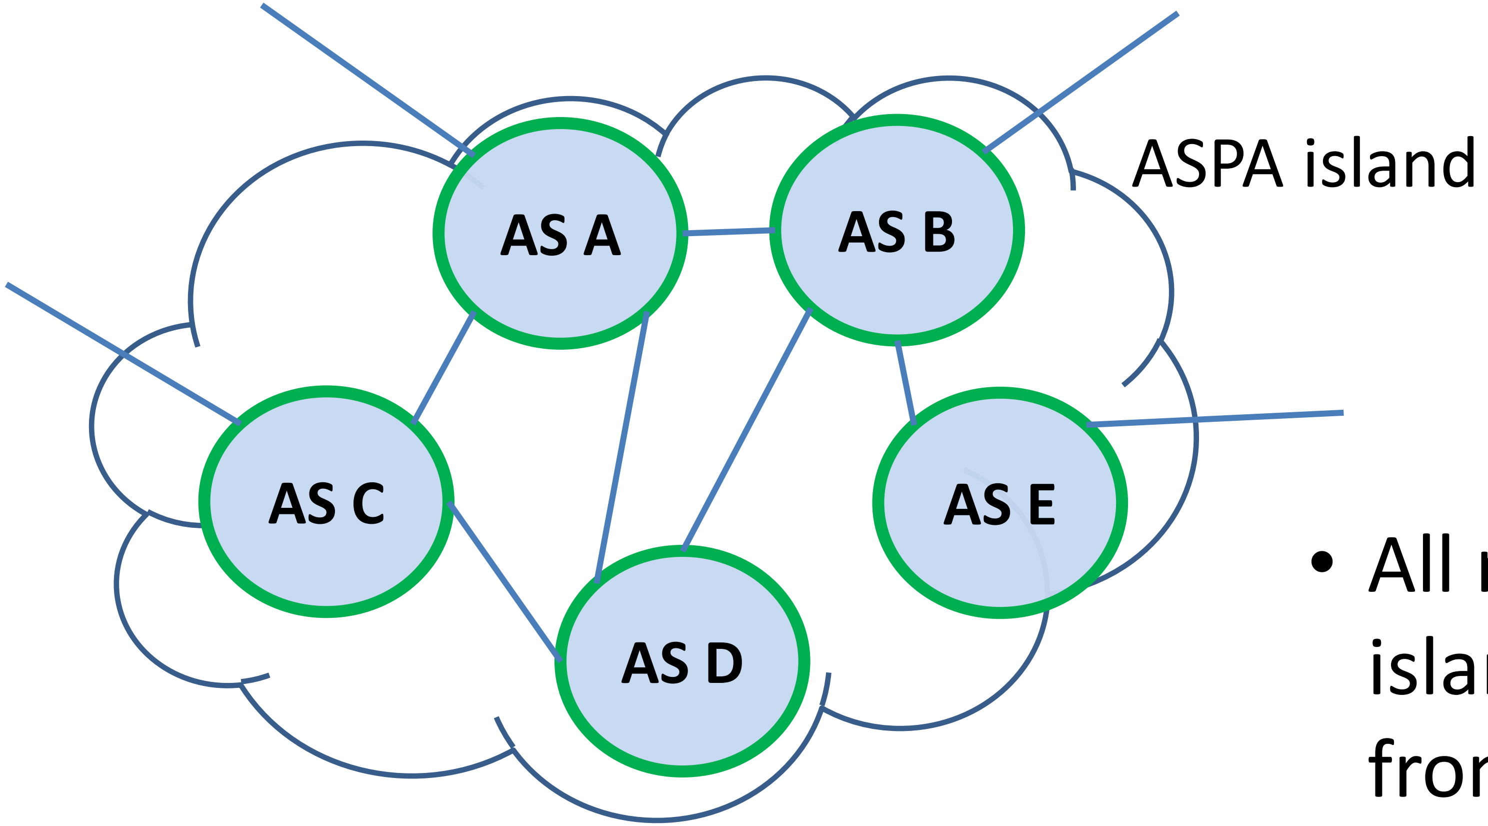

- All routes within the ASPA island are fully protected from route leaks

Shortcoming: AS_PATH maliciously shortened by a provider – undetectable

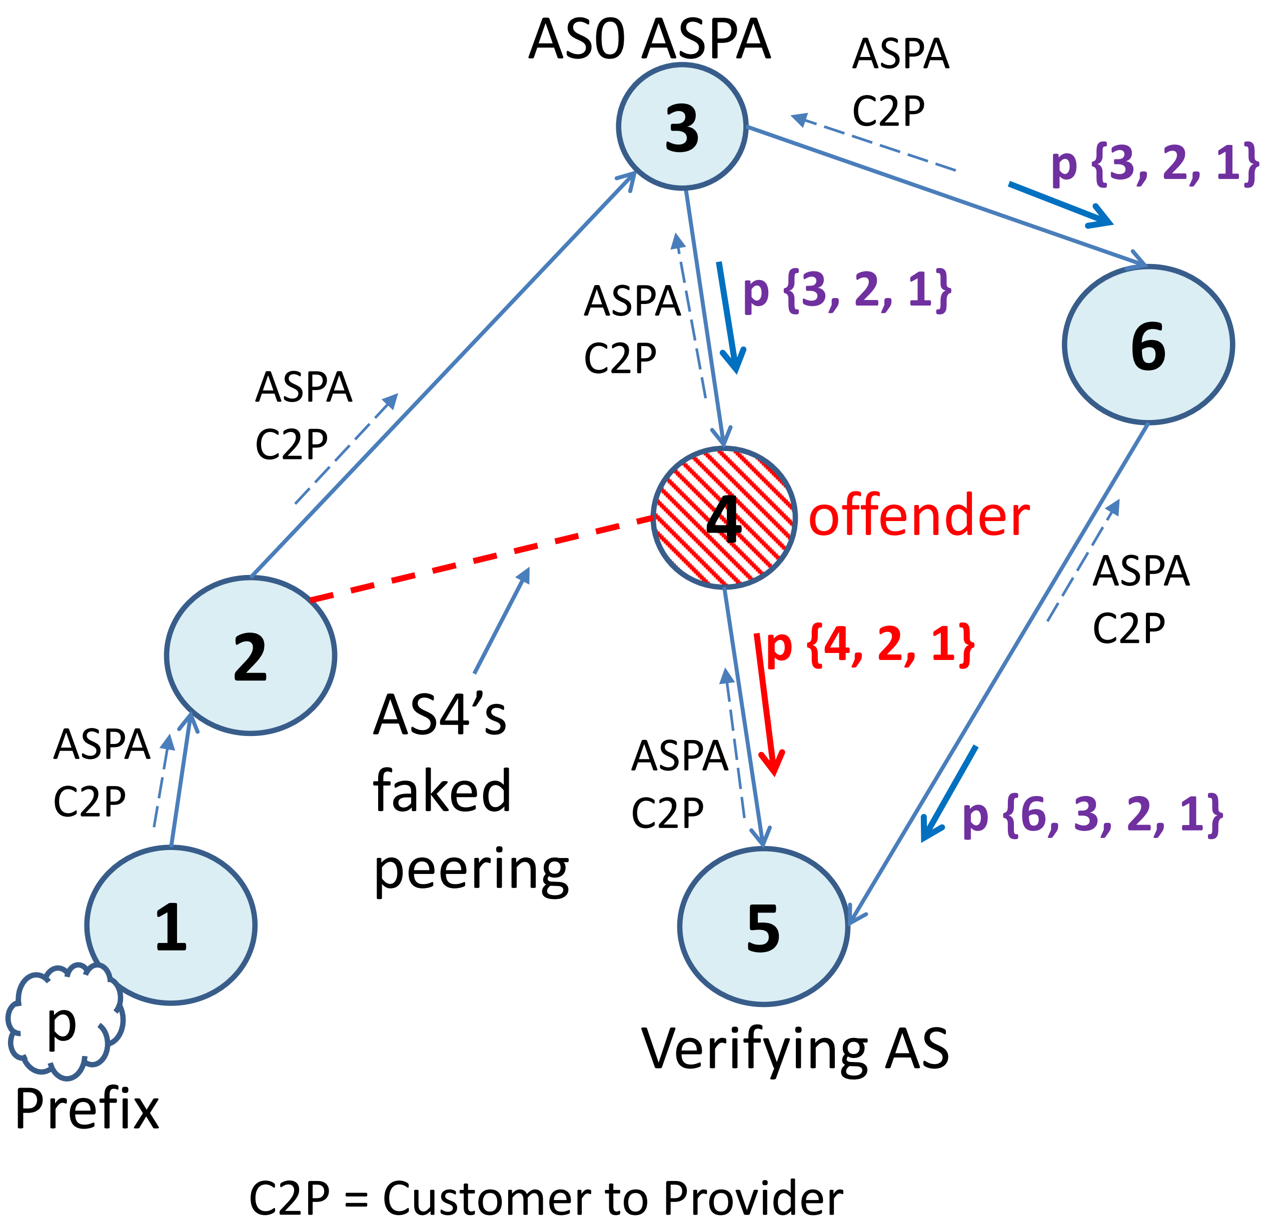

- Consider AS path verification at AS 5
- All ASes are doing ASPA
- AS 4 (provider) wants AS 5 (customer) to prefer its path
- AS 4 shortens the AS_PATH
- AS 5 chooses the manipulated shorter route via AS 4
- Since other ASes are good, if AS4 does not drop AS5's (customer's) data traffic, then the traffic still reaches the destination via a feasible and route-leak free path.

- BGPsec can provide full AS_PATH protection
- But it lacks route leak protection
- Use ASPA and BGPsec in a complementary way

Thank you

Questions?

Email: ksriram@nist.gov