soBGP

ftp://ftp-eng.cisco.com/sobgp/index.html

soBGP

- Problem Overview (BGP Security)
- Design Constraints
- Validating Keys (Entities)
- Validating Authorization
- Validating the AS Path
- soBGP Operation
- Deploying soBGP

Route Origin

- Validating which AS originated the advertisement doesn't prevent any attacks (although it can provide a fingerprint proving the source of the attack).
- This attack could be prevented if AS65500 could discover if AS65501 is authorized to advertise 10.1.1.0/24 or not.

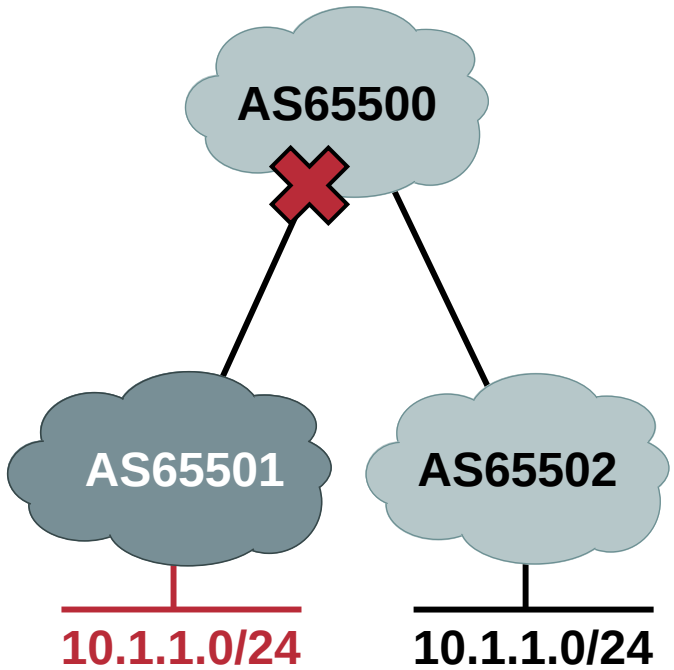

AS Path

- **Validating the AS Path is also important—*but what can we actually validate?***
- **AS65500 receives just one advertisement from AS65501.**
- **The apparent path is AS65501, AS65503, AS65504.**
- **AS65500 is actually choosing AS65501, AS65502, AS65504.**
- **Routing is based on IP Address, not AS Path.**

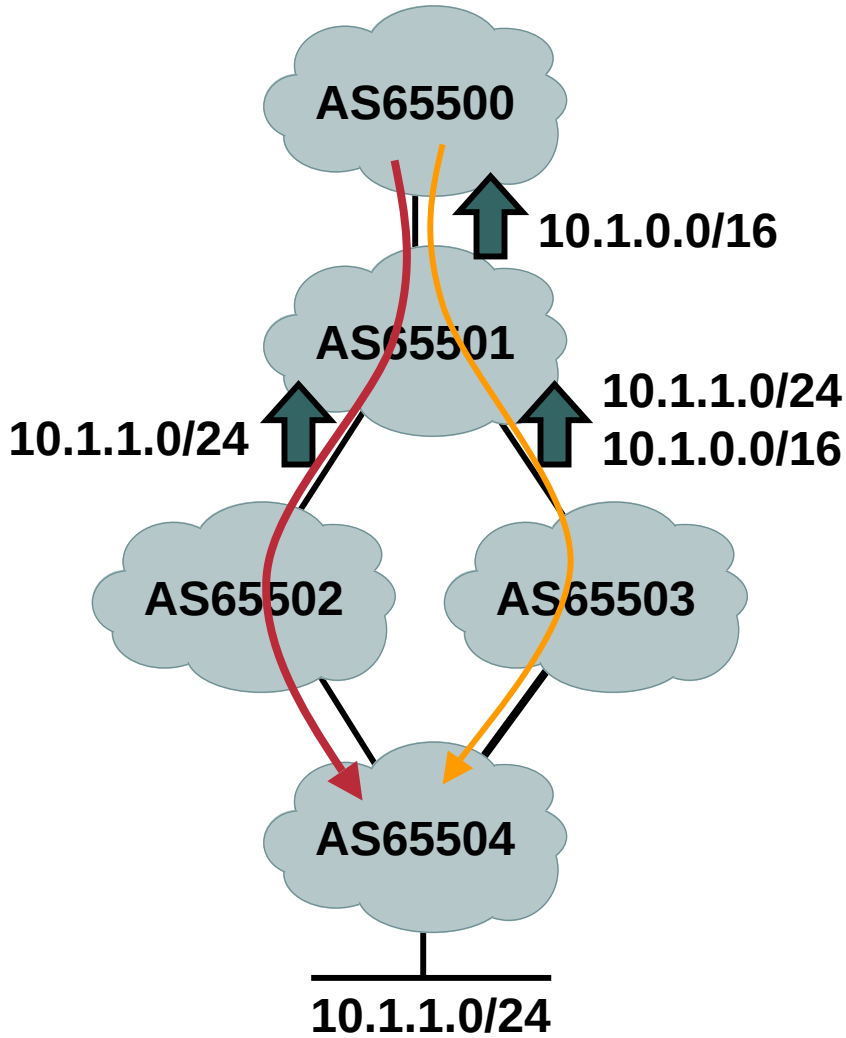

AS Path

- This can occur for various reasons:
  - Aggregation
  - Filtering
  - Longest Match Bounding
- We cannot validate the path the traffic will take.
- Validating the path the updates take is not a goal.
- From AS65500's point of view, validating that AS65501 has at least one valid path to 10.1.1.0/24 *is* a goal, however.

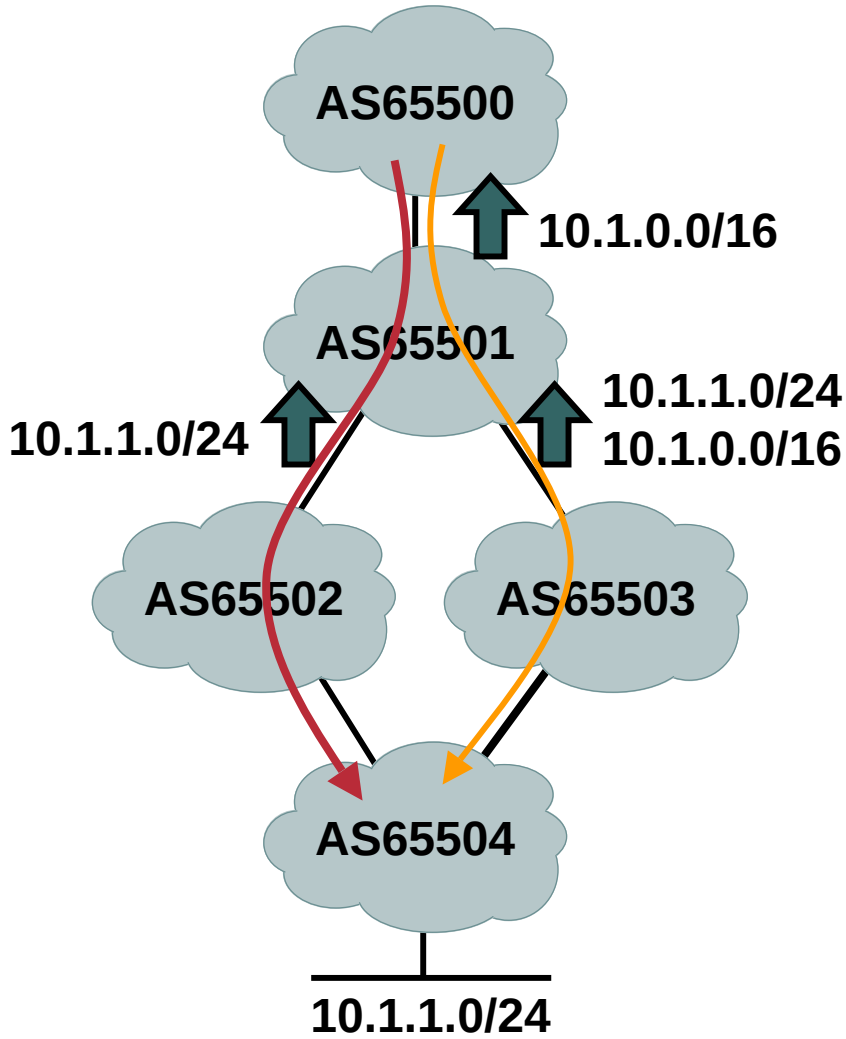

Secure Origin BGP (soBGP) Operation

- **We have two goals:**
  - **Validating an AS is authorized to originate a prefix.**
  - **Verifying a peer which is advertising a prefix has at least one valid path to the destination.**

Design Constraints

- **Must *not* rely on a central authority of any type.**
- **Must be incrementally deployable (it must provide some level of security without the participation of every AS).**
- **Must allow deployment flexibility (on box or off box cryptography, etc.).**

Design Constraints

- **Should not rely on routing to secure routing (No external database connection on system initialization).**
- **Flexibility should be provided to allow operators to configure the level of security vs. overhead and convergence speed.**
- **Minimize impact to current implementations of the BGP protocol.**

soBGP Doesn’t Protect

- **The BGP Transport Connection**
  - **There are other mechanisms designed for this (IPSec)**
- **BGP Attribute Validity**
  - **Most attributes indicate local policy (Local Pref, MED)**

Validating Keys

- Each participant (entity) in the internetwork creates a public/private key pair.
- Each participant then has the public key/AS pair signed by a trusted third party.
- The public key/AS pairing can be validated using the signer's public key.
- This signed certificate is called an *EntityCert.*

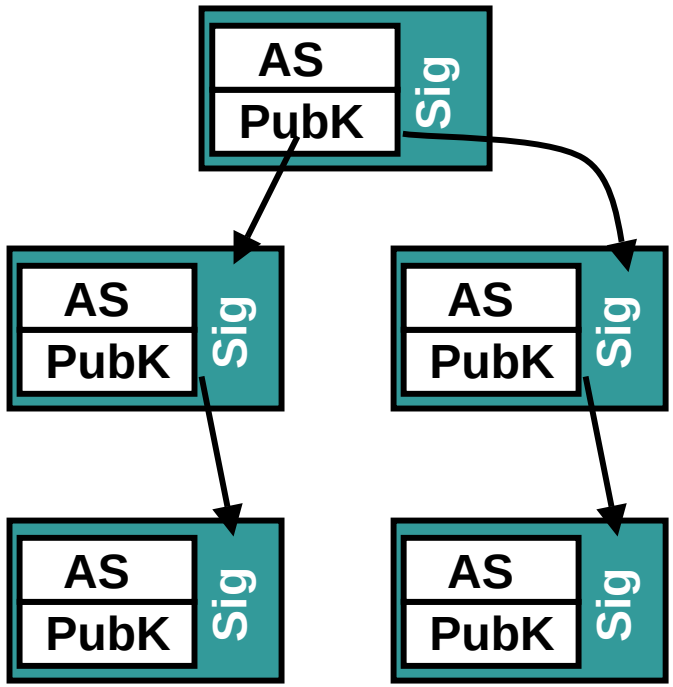

Validating Keys

- Various trusted third parties may sign these EntityCerts:
  - Authority which issued the AS number
  - Commercial authority
  - Any universally known and trusted party in the Internet domain

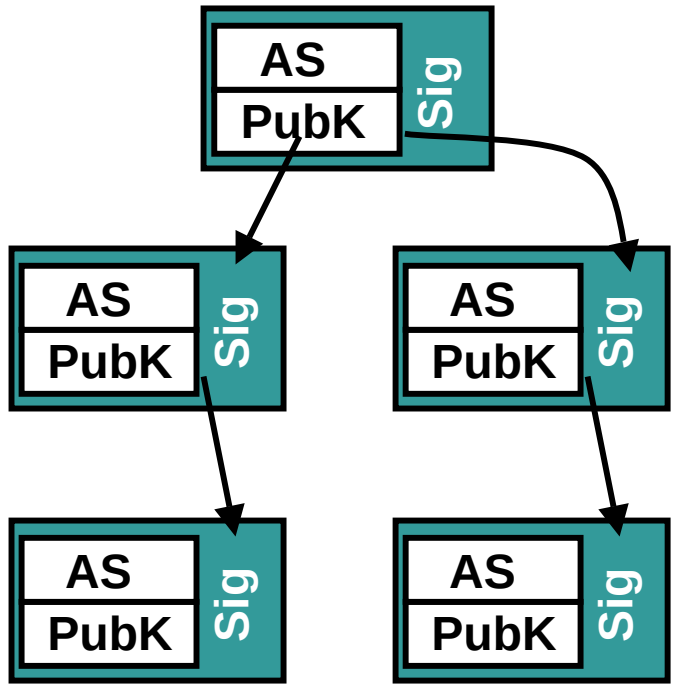

Validating Keys

- **Some small number of keys are accepted as valid at the root, and manually configured on devices running soBGP.**
- **The remainder may be learned and validated using the public keys learned through other validated EntityCerts.**
- **Thus, the EntityCerts form a web of trust through which the public key of each AS in the internetwork may be trusted.**

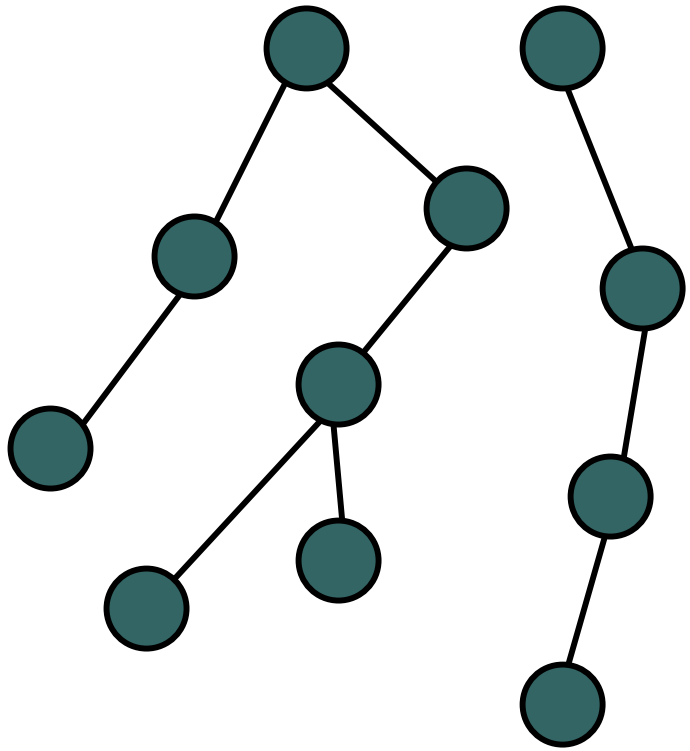

Validating Authorization

- If an AS authorizes another AS to advertise a block of prefixes, it issues a certificate signed with its private key, indicating this authorization.
- This is called an *AuthCert.*
- The public key of each signer, learned from the EntityCerts, can be used to validate each AuthCert.
- The AuthCerts form a tree (or trees) of authorization, with shorter prefixes at the top of the tree.

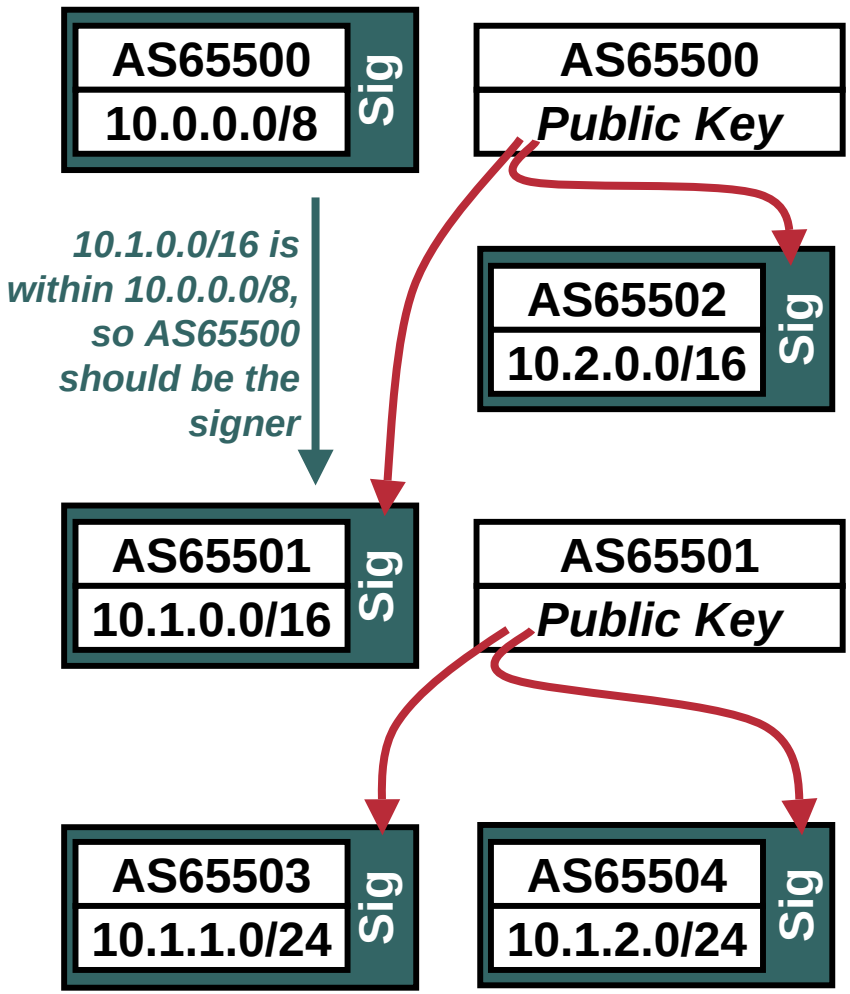

Validating Authorization

- Each AS also builds a certificate which contains policy information, and signs it with its private key.
- This is called the *PolicyCert.*
- One policy which can be included in this certificate is the maximum prefix length within any given address space.

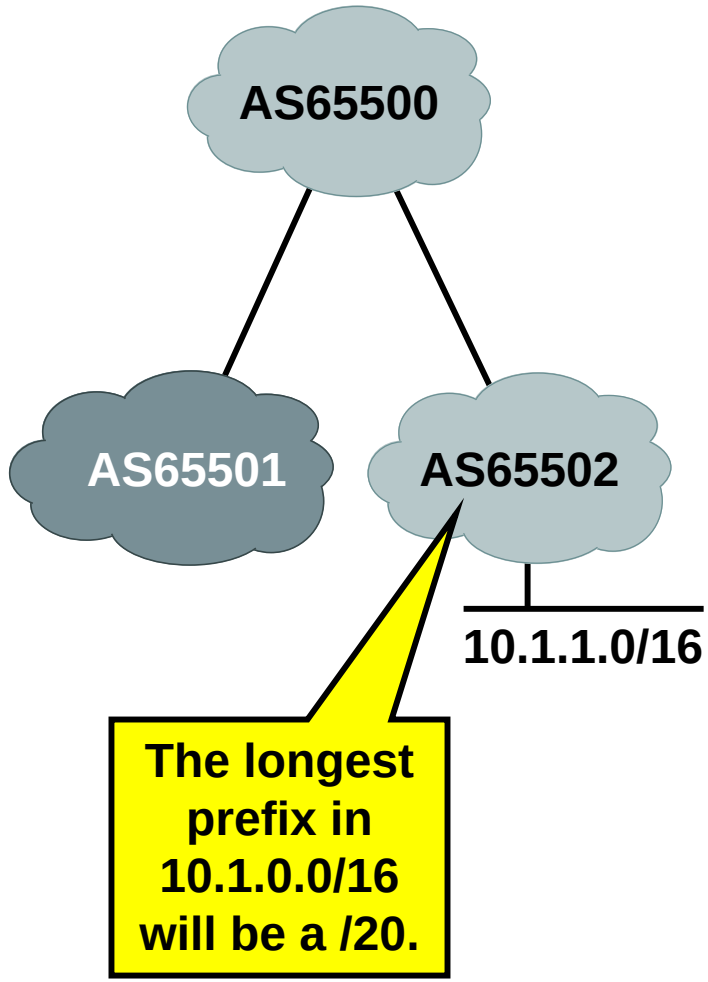

Validating Authorization

- If AS65500 receives an advertisement for 10.1.1.0/24, in any form, from AS65501, it can safely discard the advertisement, since the AS authorized to advertise prefixes within that address space has stated the longest prefix length will be a /20.
- This helps to protect against longer prefix match attacks.

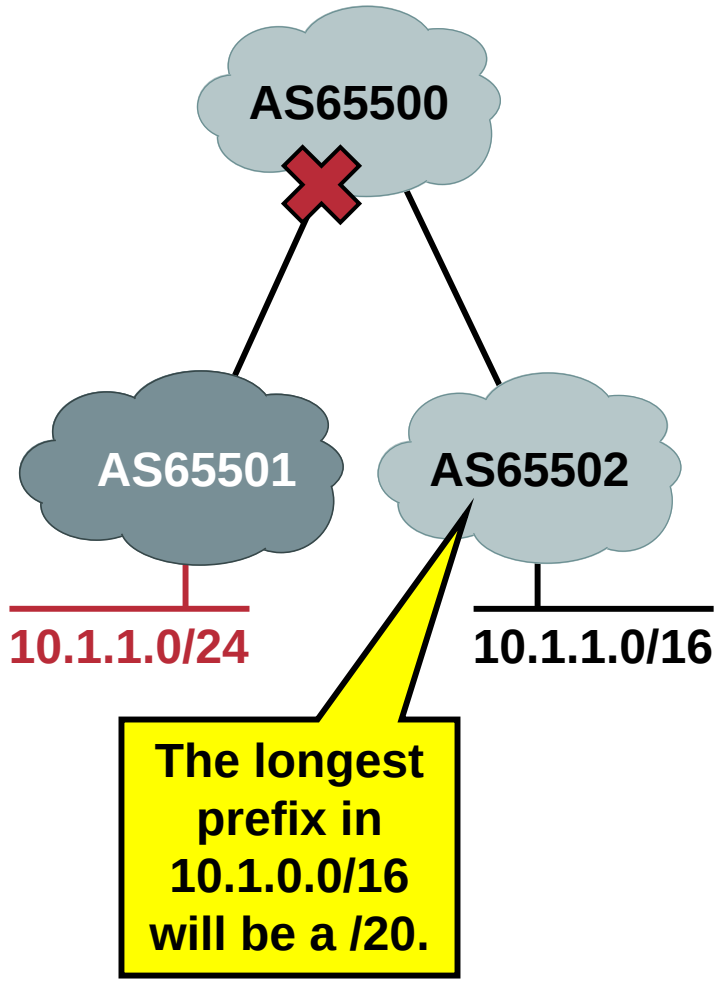

Validating the AS Path

- **Within the PolicyCert, each AS also advertises a list of its peers, signed using its private key, throughout the entire internetwork.**

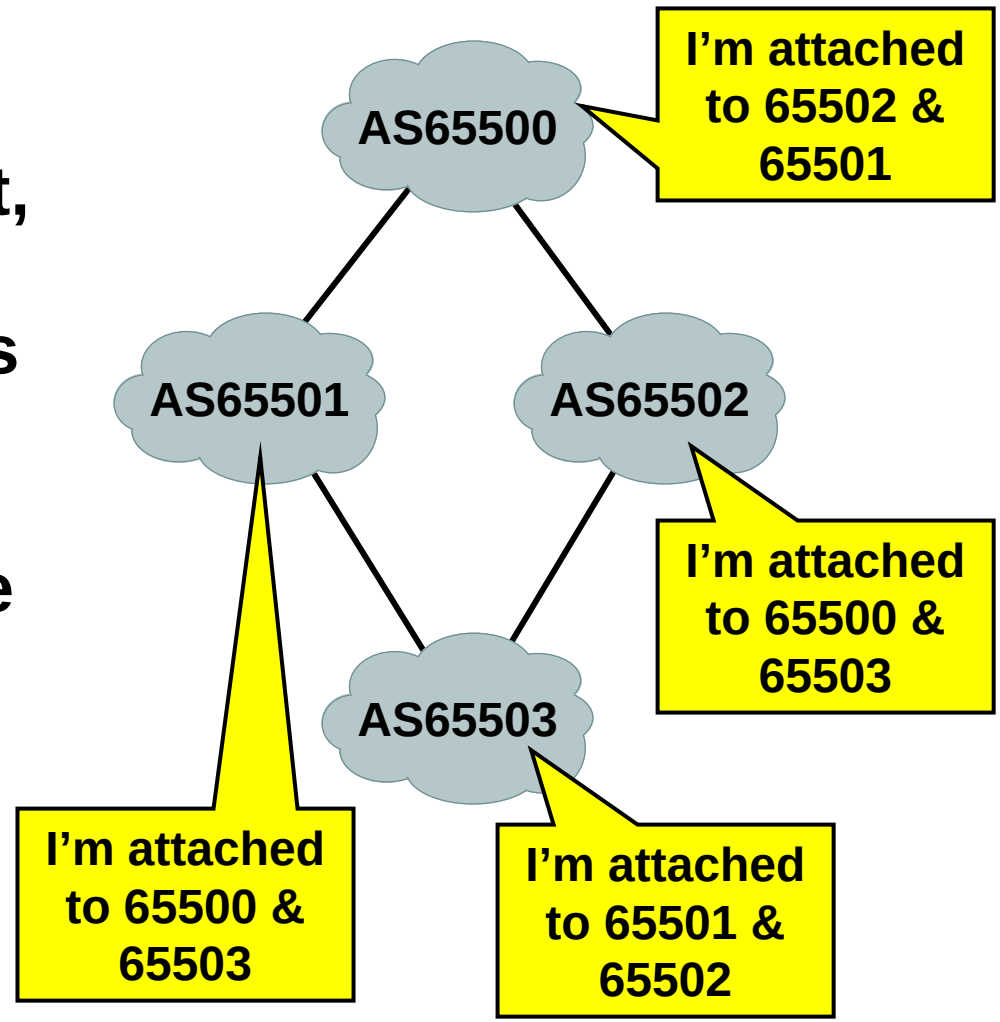

Validating the AS Path

- From this information, a graph of the internetwork topology can be built showing all valid AS Paths.
- Two way connectivity checks are used; two AS' must advertise they are connected to each other for the link to be considered valid.

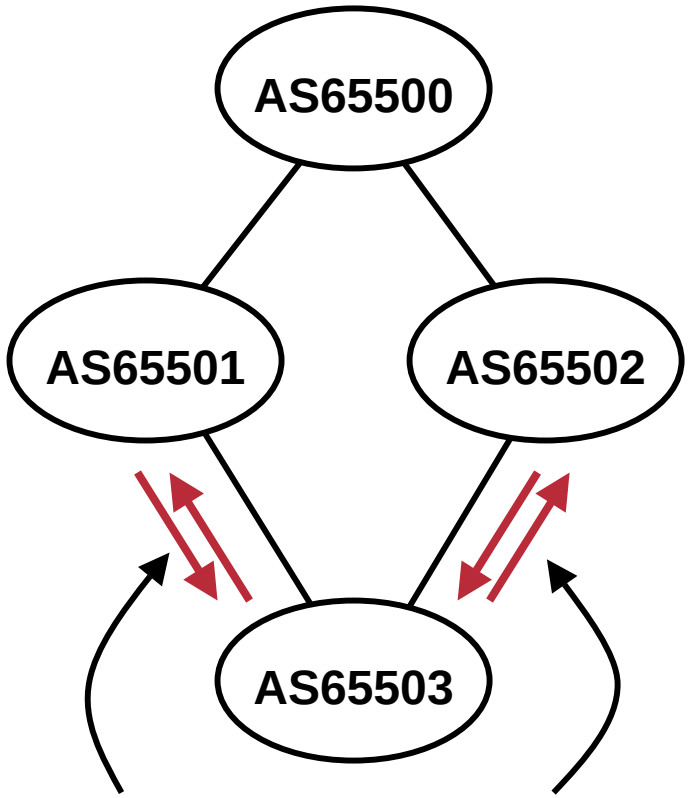

AS65503 advertises a connection to AS5501; AS 65501 advertises a connection to AS65503; the link is valid

AS65503 advertises a connection to AS5501; AS 65501 advertises a connection to AS65503; the link is valid

Validating the AS Path

- If an AS tries to cut hops out of the path to appear to be a better path, the two way connectivity check will fail, so the path is marked as invalid.
- If AS65502 attempts to make its path to AS65504 shorter by cutting AS65503 out of the path, AS65500 will be able to detect the alteration in the AS Path.

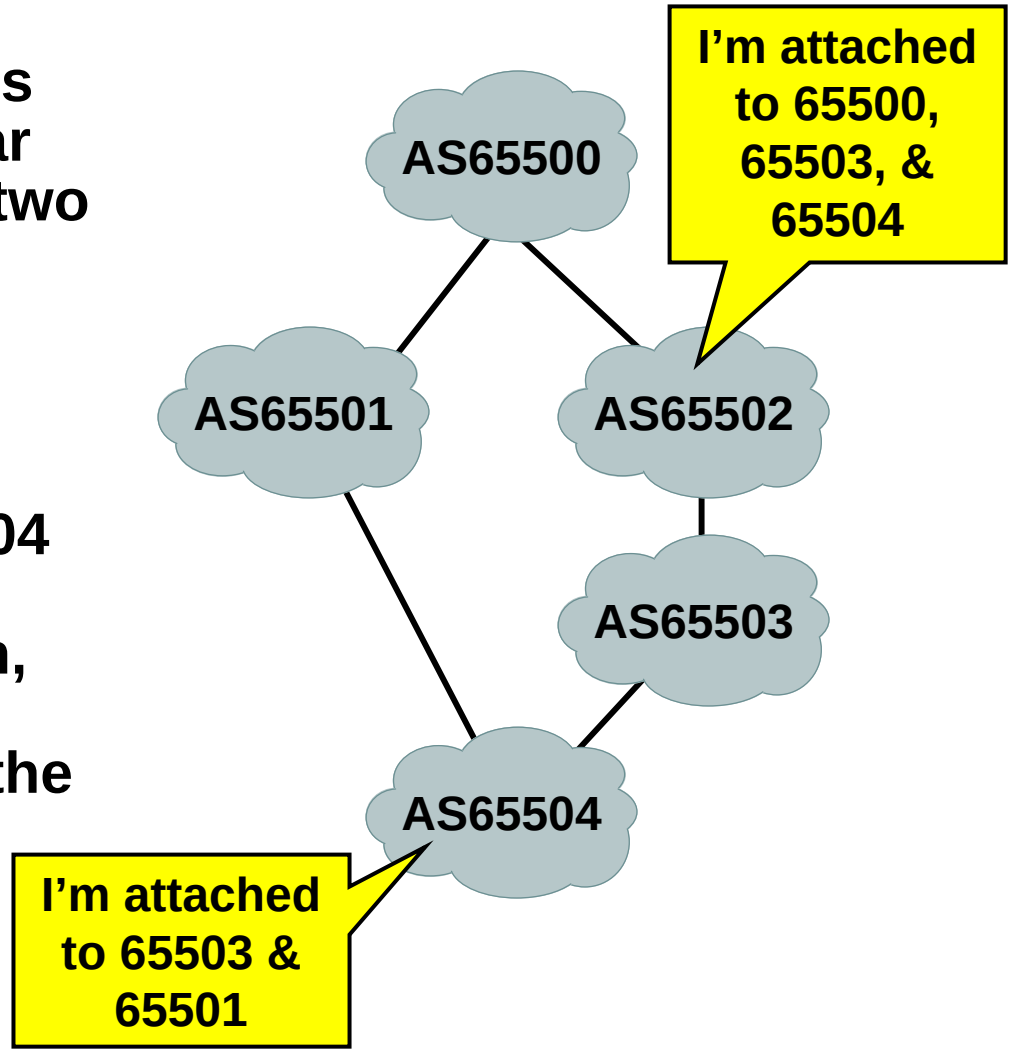

Validating the AS Path

- If an AS attempts to spoof by claiming an AS is attached to it, the two way connectivity check will fail, and the path will be discarded.
- If AS65502 attempts to claim AS65503 is behind it (to advertise AS65503's prefixes), AS65500 will be able to detect the bad AS Path from AS 65503's PolicyCert.

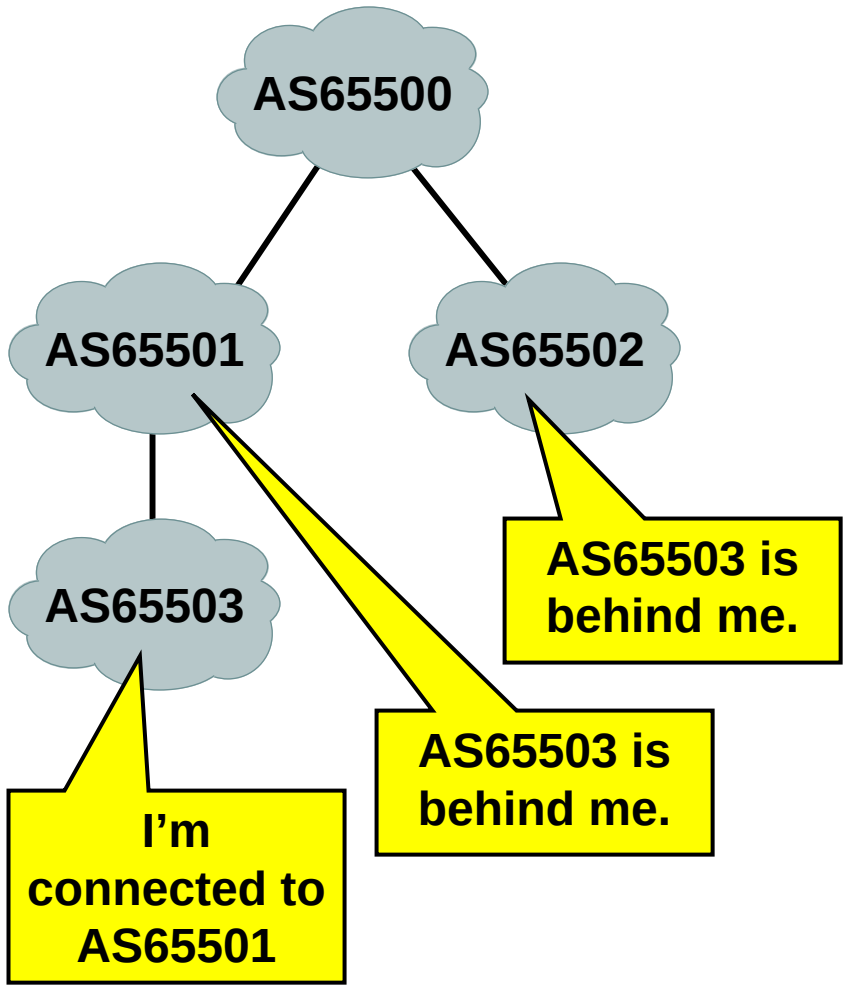

Validating the AS Path

- An AS may advertise a peer as a nontransit, which prevents dual homed customers from being able to transit traffic intentionally or unintentionally.
- It's also possible to hide connectivity and maintain security in some cases.

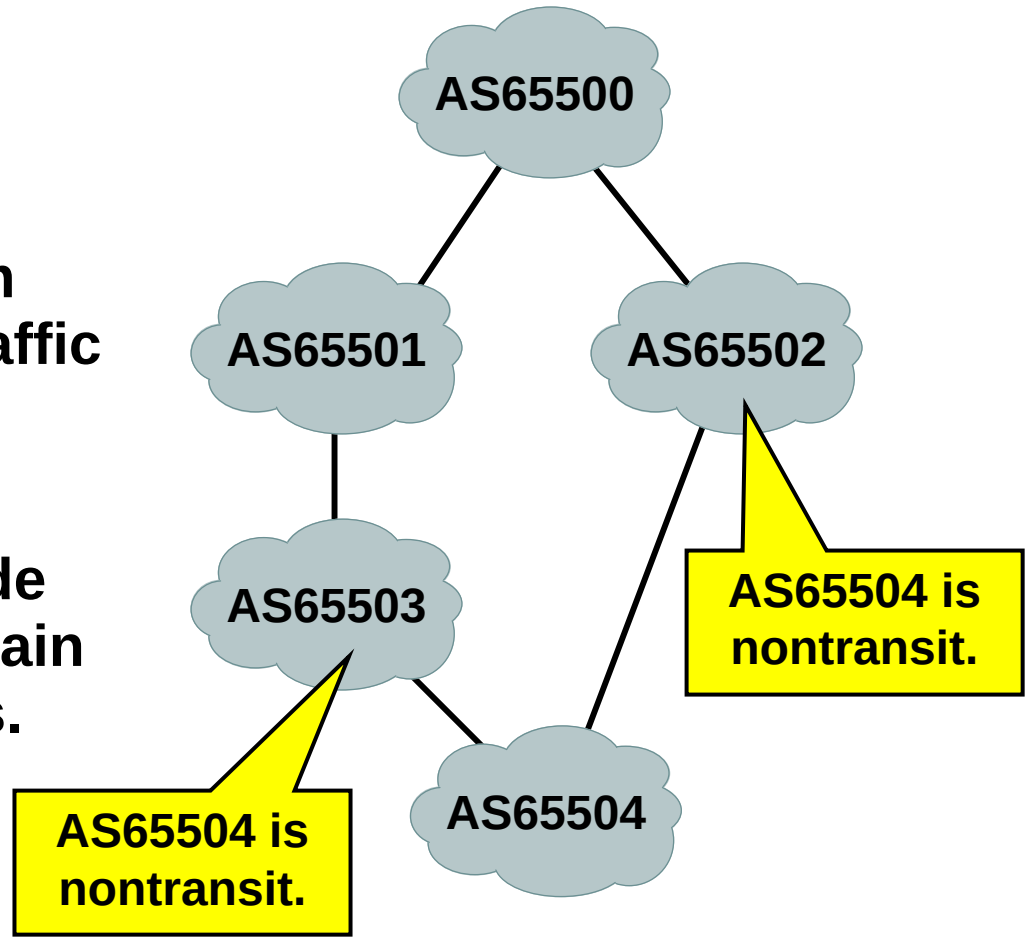

soBGP Operation

- **Certificate Transport**
- **Certificate Processing**
- **Update Processing**

Certificate Transport

- **Certificates are transported in a new BGP message type, the SECURITY message.**
  - **Certificates are carried within TLVs**
  - **Expandable to other security related information**
- **Negotiated at session startup**
  - **Certificates may be exchanged before routing**
  - **Routing may be exchanged before certificates**
  - **Certificates only may be exchanged**

Certificate Transport

- **The SECURITY message type also provides requests**
  - **Security messages may be filtered for various reasons**
  - **The Request message provides the ability to readvertise all security information or just a subset**

Certificate Processing

- **Begin with a set of manually entered EntityCerts which are implicitly trusted to be correct.**
- **As EntityCerts are received, the signer AS is checked against the list of validated keys.**

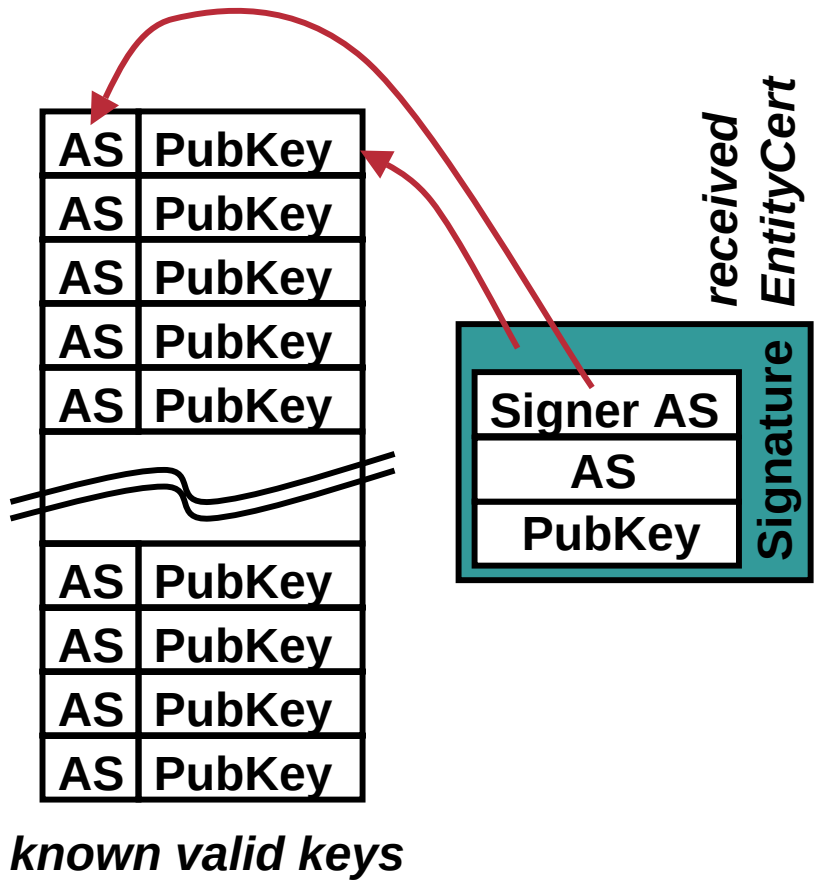

Certificate Processing

- If the signer's key is present, the signature is checked, and the public key added to the list of validated keys.
- The result is a list of the autonomous systems in the internet, each with an associated, validated public key.

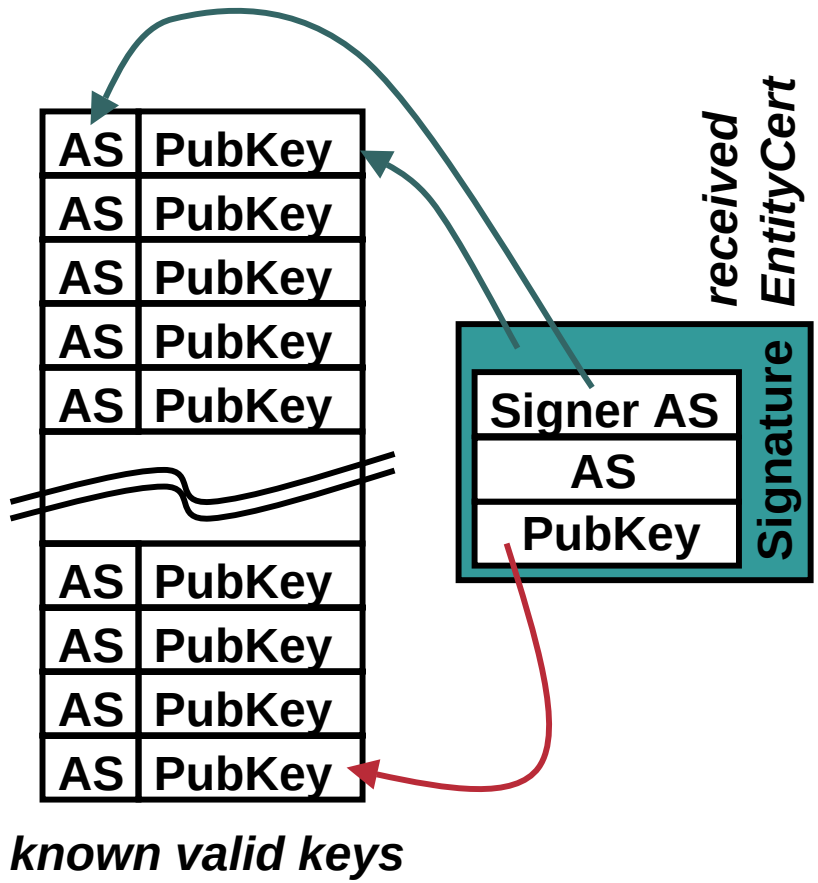

Certificate Processing

- As PolicyCerts are received, their signatures are checked against the database of known valid public keys.
- If the signature validates:
  - Any policies are entered into the policy database.
  - The list of attached peers is added to the database from which the topology graph is built.

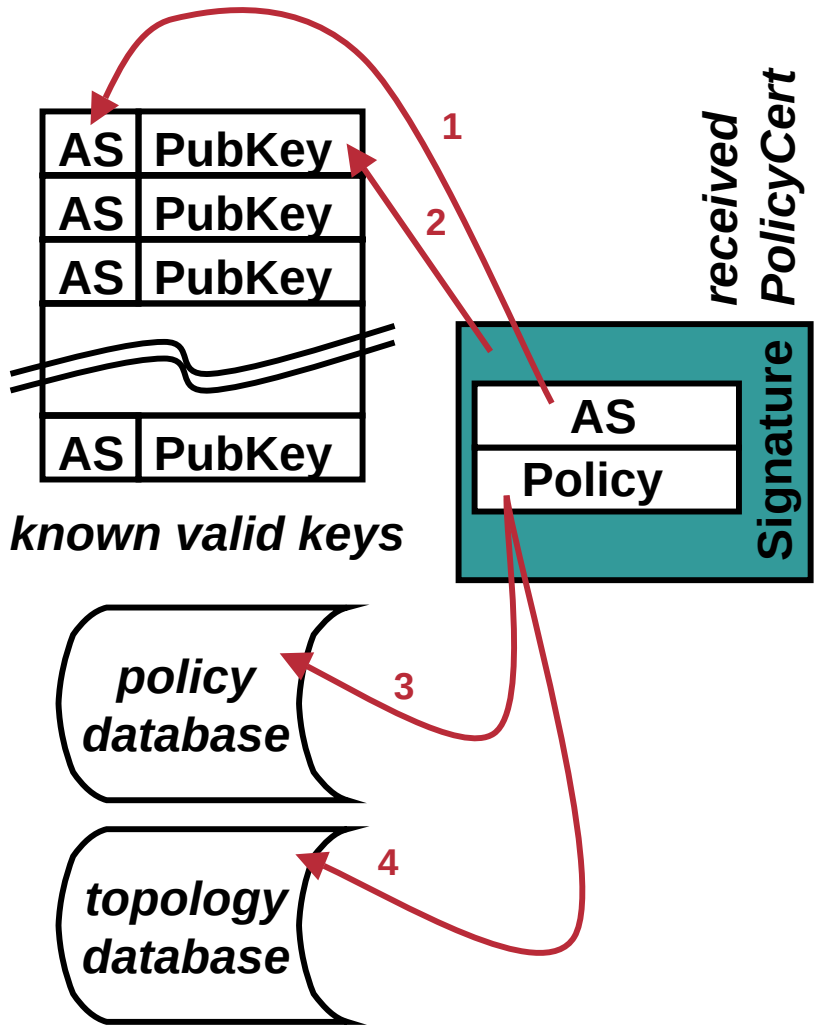

Certificate Processing

- **Each address block listed with a policy is examined in the AuthCert database, and any policies required applied.**

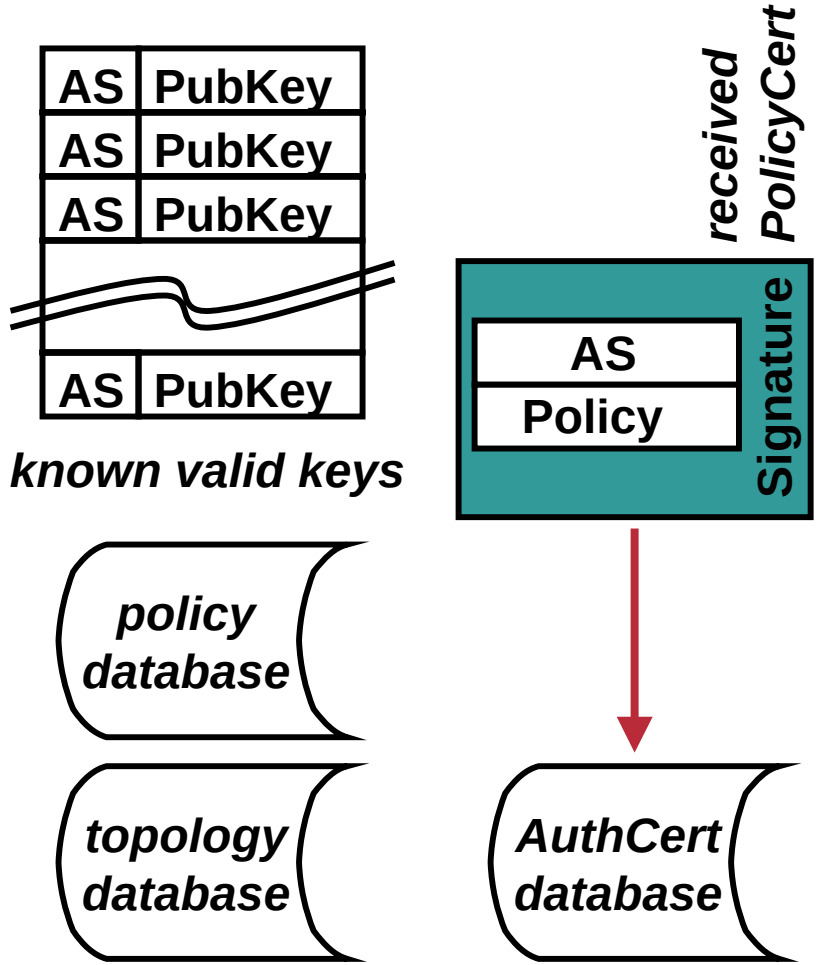

Certificate Processing

- As AuthCerts are received, the authorizing AS is looked up in the list of known valid keys.
- If it is found, the public key is used to validate the signature on the AuthCert.
- If the AuthCert validates, the policy database is checked, and any policies applied.
- The address block and origin AS are then added to the AuthCert database.

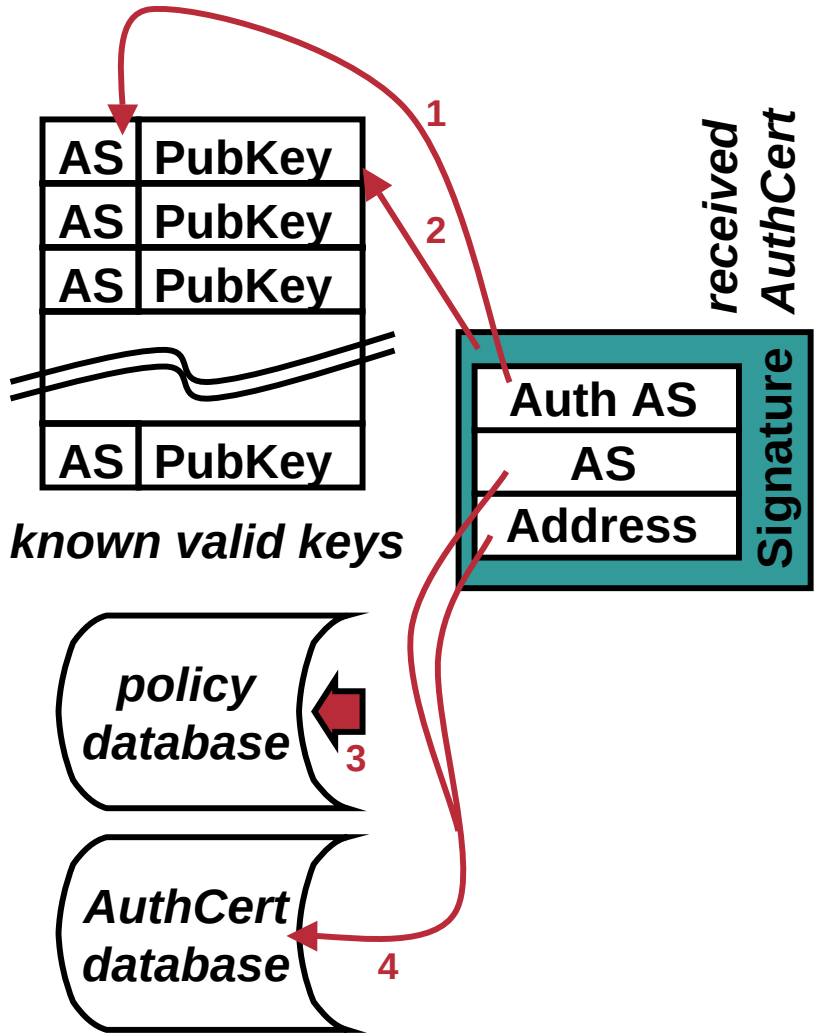

Update Processing

- There are no per-update crypto operations needed (enhances scalability)
- Received updates can be processed in a number of different ways:
  - The update may be validated before being placed in the BGP table.
  - The update may be placed in the BGP table, then validated before installation in the RIB.
  - The update may be installed in the RIB (and advertised), then validated (and withdrawn if needed).

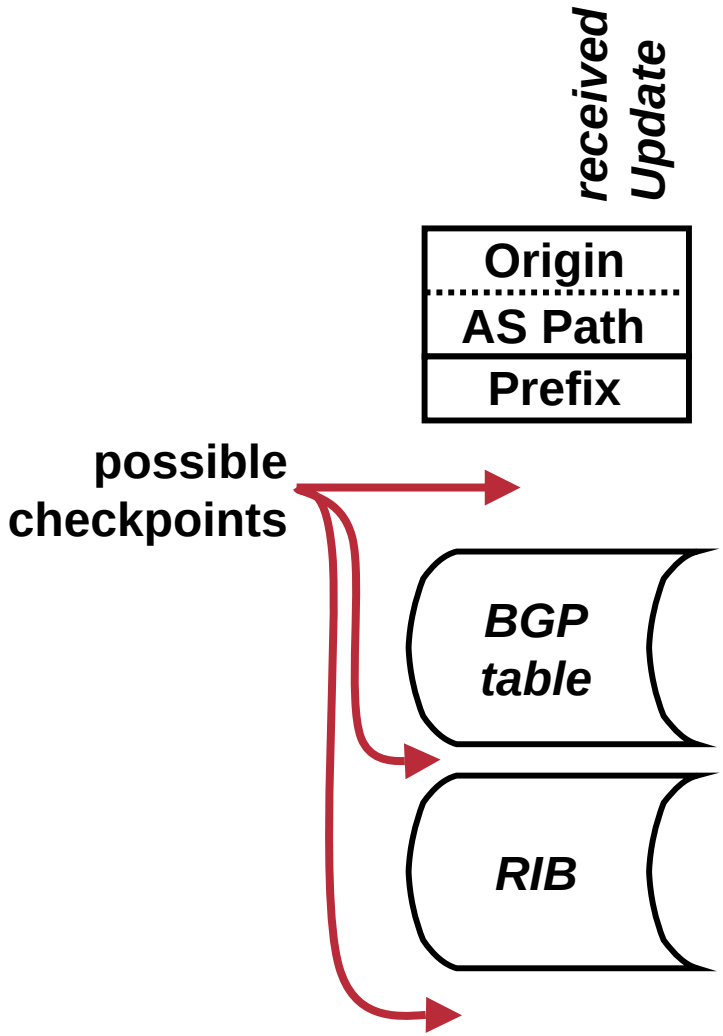

Update Processing

- The first hop in the AS Path (the origin AS) is looked up in the AuthCert database. The prefix must fall within the range of addresses this AS is authorized to originate.
- Any policies which are tied to the entry in the AuthCert database are checked, and actions taken as needed.

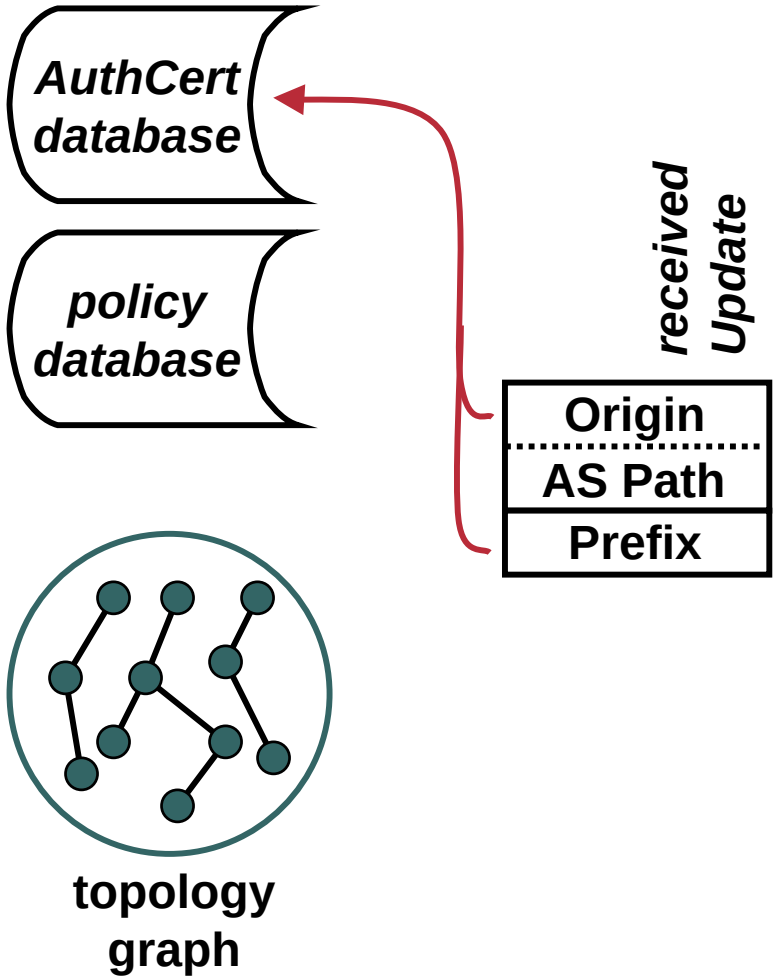

Update Processing

- The AS Path is then checked against the topology graph built earlier.
- If the AS Path is found to be invalid, the route may be discarded.

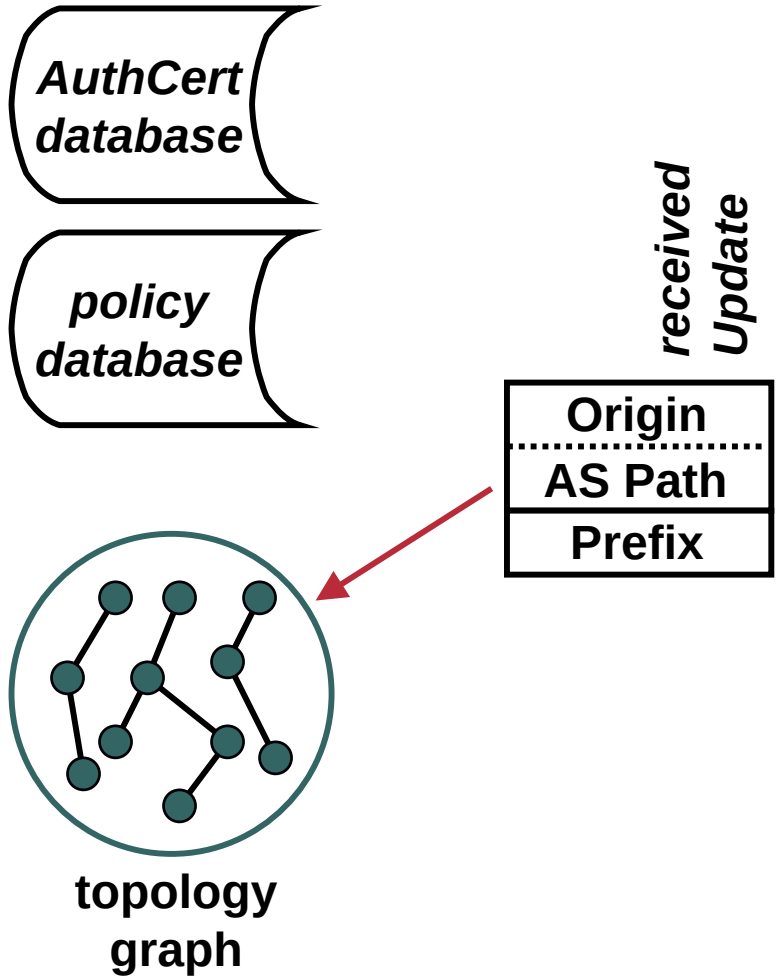

soBGP Deployment Considerations

- **Deployment Options**
- **Incremental Deployment**
- **Aggregation**

Deployment Options

- **The most straightforward deployment option is:**
  - **Exchange certificates at all eBGP peering points (AS edges).**
  - **Process the certificates, and build the required soBGP tables at each eBGP speaker.**
- **Each eBGP speaker must then be capable of running the cryptographic processes needed to process certificates.**

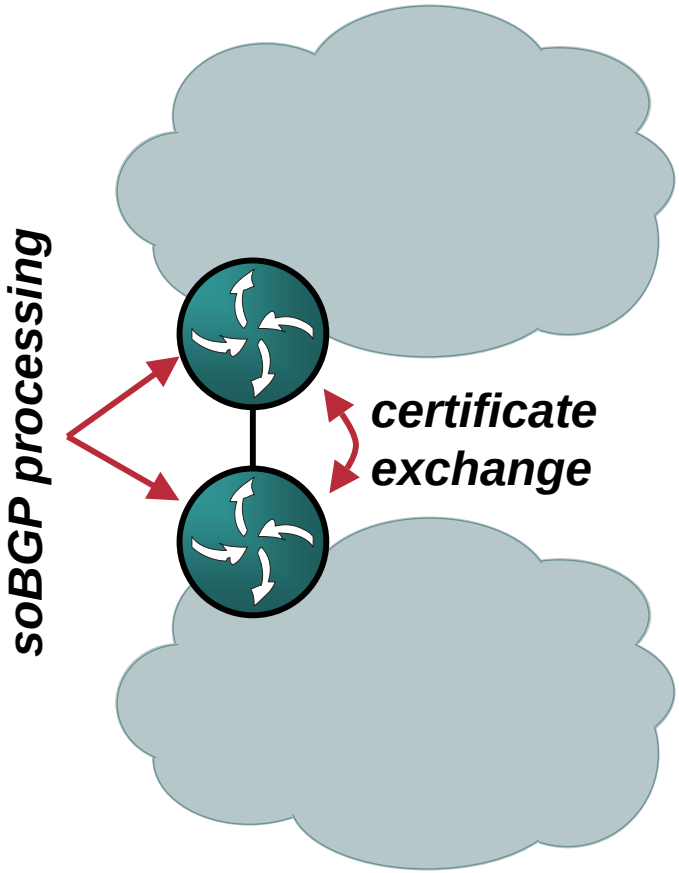

Deployment Options

- Certificates can also be exchanged at the AS edge, and “shuttled,” using iBGP connections, to a server within the AS.
- These servers then perform all certificate processing, and build the necessary databases.
- The edge routers then consult these servers, to validate received updates.

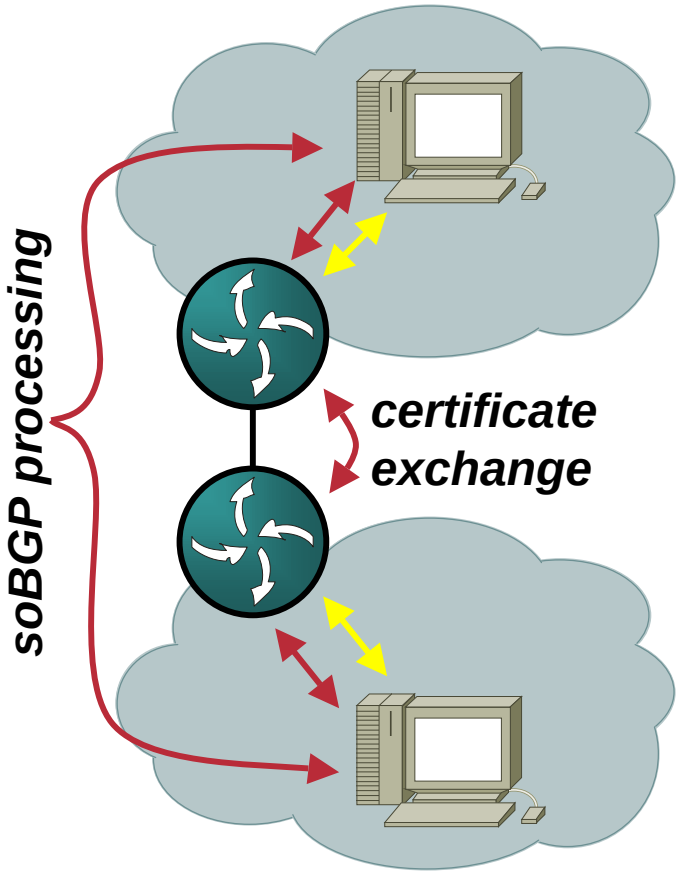

Deployment Options

- **Certificates can also be exchanged, using multihop eBGP directly between the soBGP servers in each AS.**

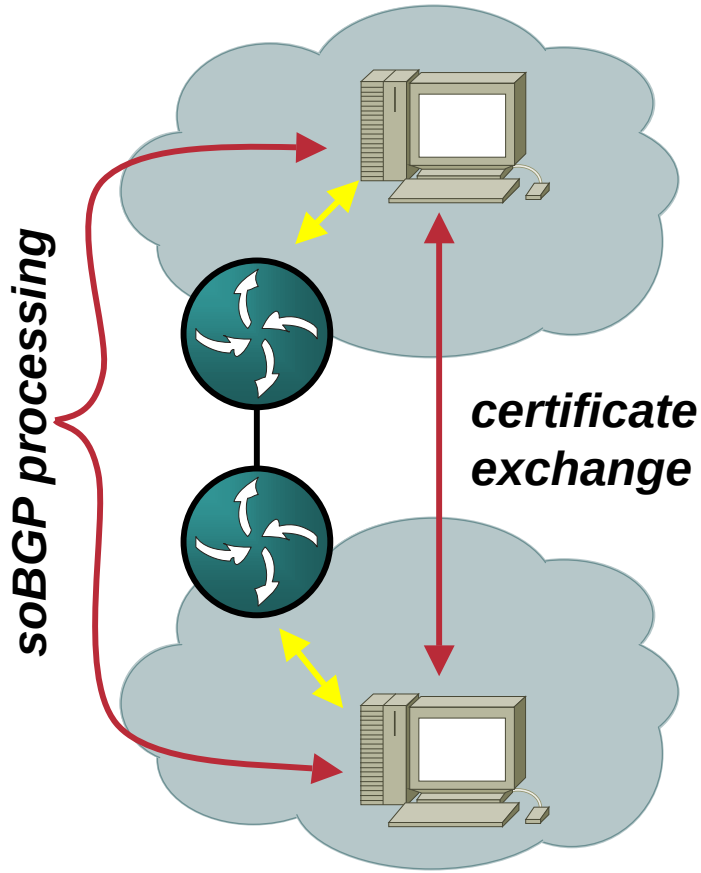

Deployment Options

- Certificates may be generated by one AS, and advertised by another AS.
- This could be the case when a customer wants a third party to advertise their certificates for high availability reasons.
- It doesn't matter who injects the certificates into the routing system, as long as the same process is used to validate them.

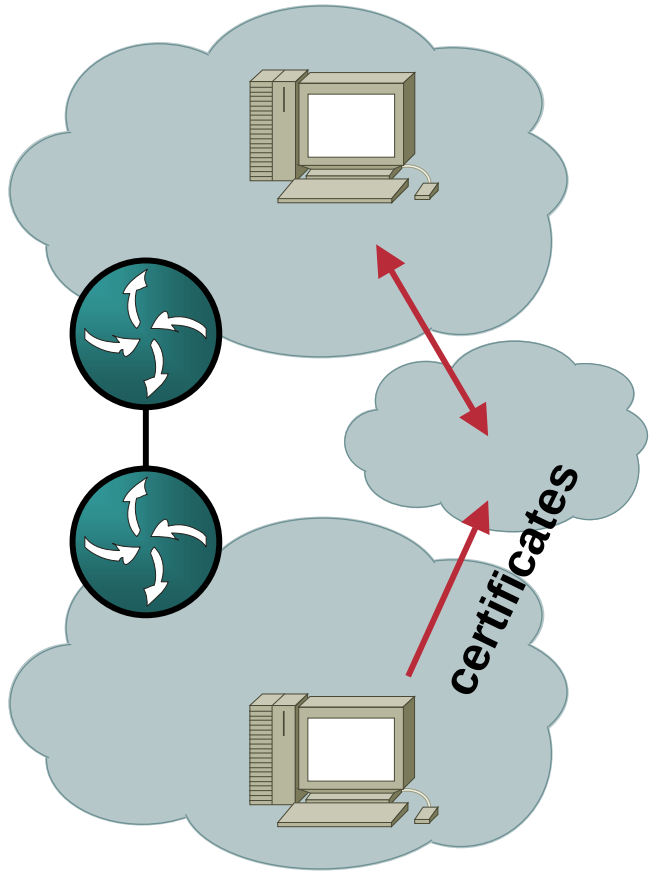

Incremental Deployment

- **Incremental deployment is a large hurdle for any security system.**
- **There's no way to have a "flag day," after which all AS' must be running the security system, in a large internetwork.**
- **soBGP allows incremental deployment—*but the amount of security provided is proportional to the completeness of the deployment!***

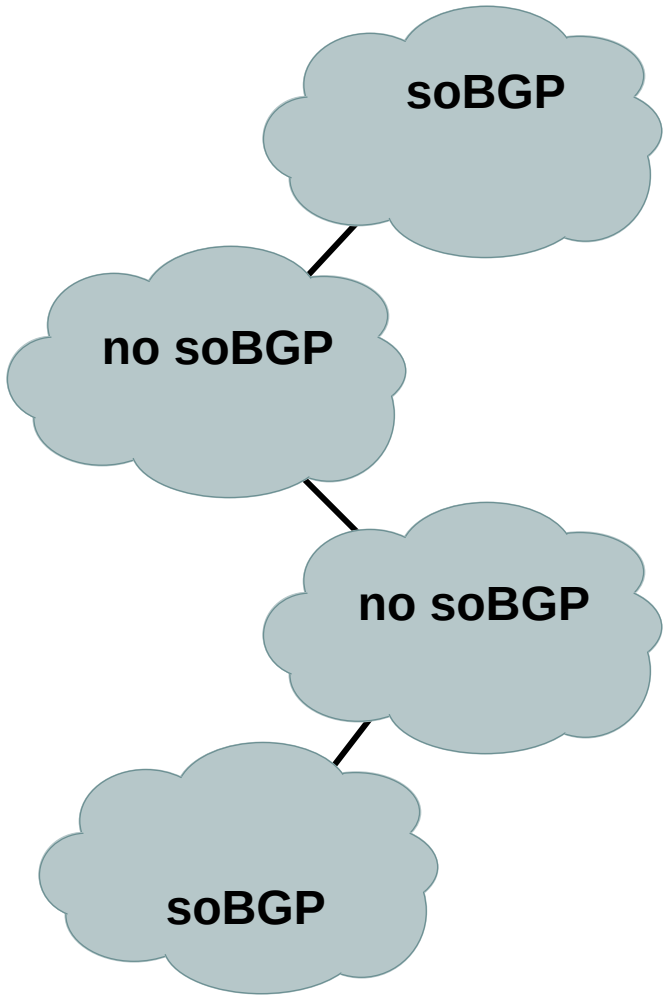

Incremental Deployment

- **The two autonomous systems which would like to run soBGP can exchange their certificates directly through eBGP multihop sessions, or through some other mechanism.**

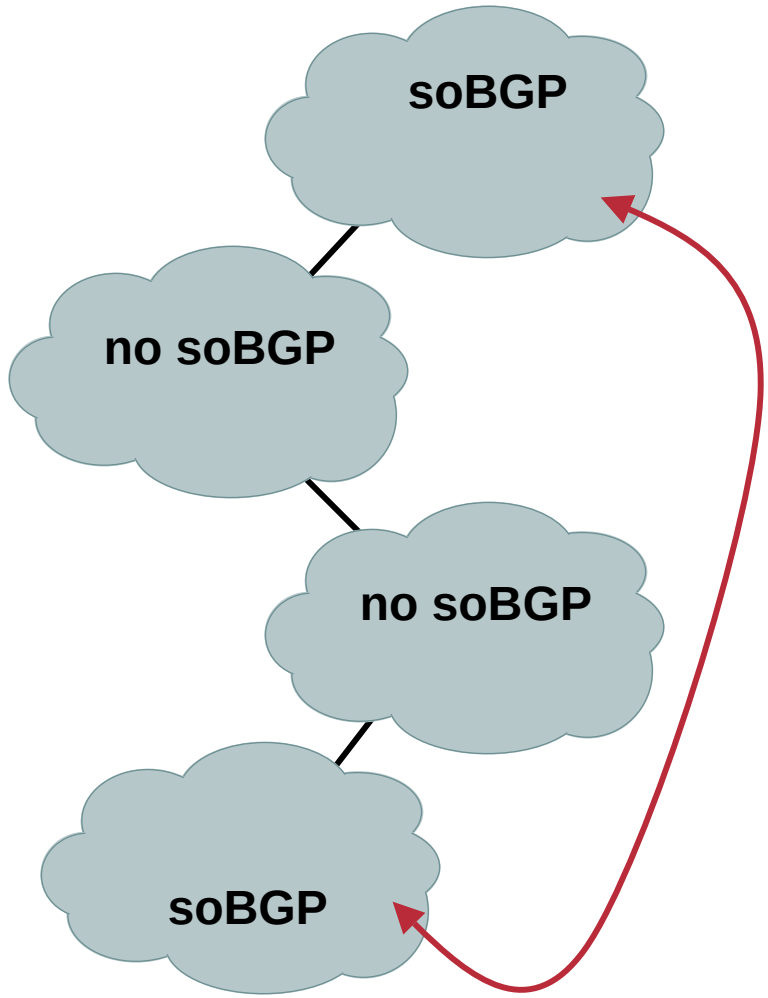

Incremental Deployment

- They are able to validate the second hop in the AS Path, using the connectivity advertised in the PolicyCerts.
- As more of the AS' participate, more of the path can be validated.

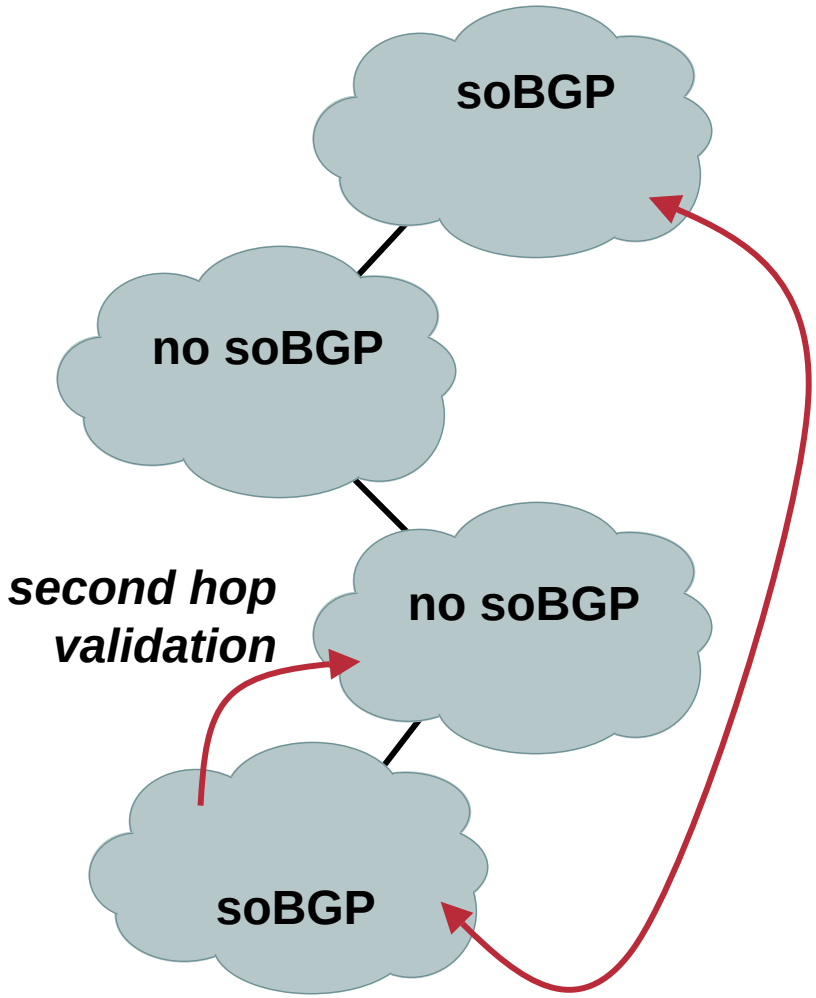

Aggregation

- **Aggregation is a problem for any mechanism that uses the AS_PATH to authenticate information**
- **The problem can be avoided by restricting AS' to only aggregate for prefixes that they are authorized to originate**

soBGP Deployment Examples

- **Unauthorized Prefix Advertisement**
- **Spoofed Peering Session**
- **Unauthorized Transit**

Unauthorized Prefix Advertisement

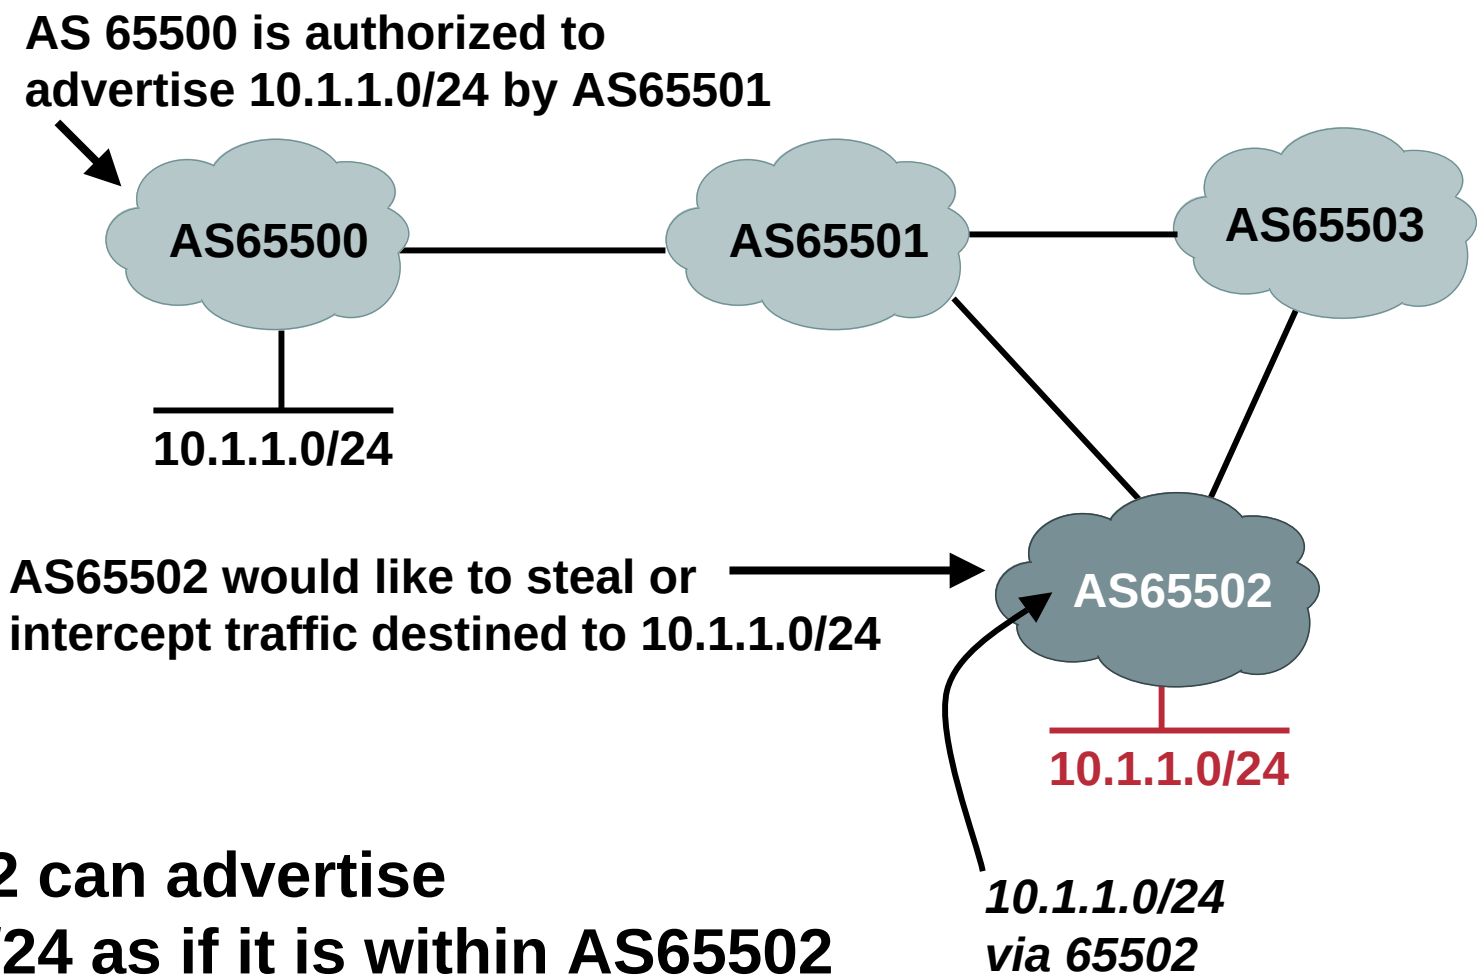

AS65502 can advertise 10.1.1.0/24 as if it is within AS65502

Unauthorized Prefix Advertisement (cont'd)

soBGP Solution:

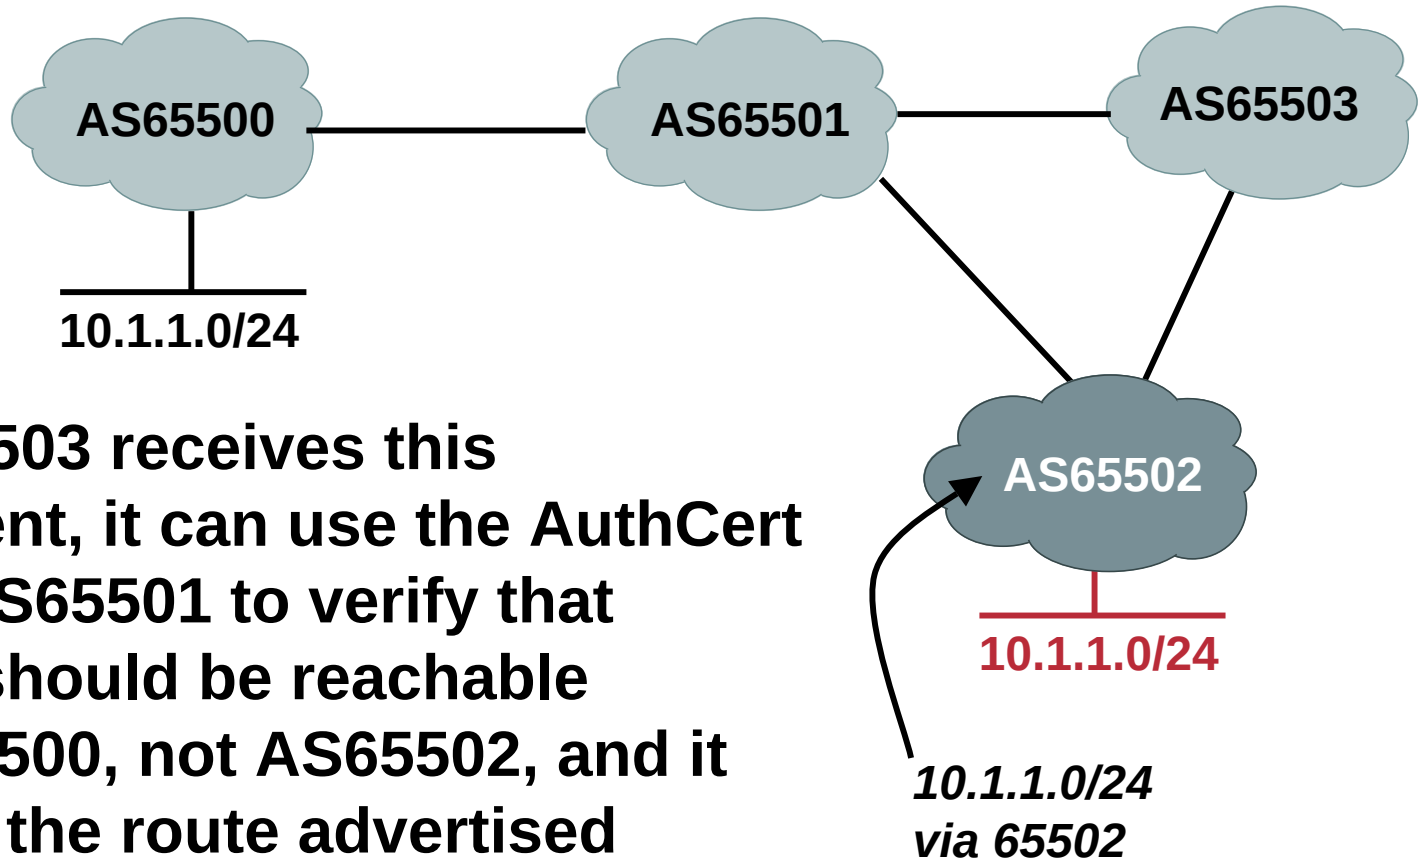

When AS65503 receives this advertisement, it can use the AuthCert issued by AS65501 to verify that 10.1.1.0/24 should be reachable within AS65500, not AS65502, and it can discard the route advertised by AS65502

Spoofed Peering Session

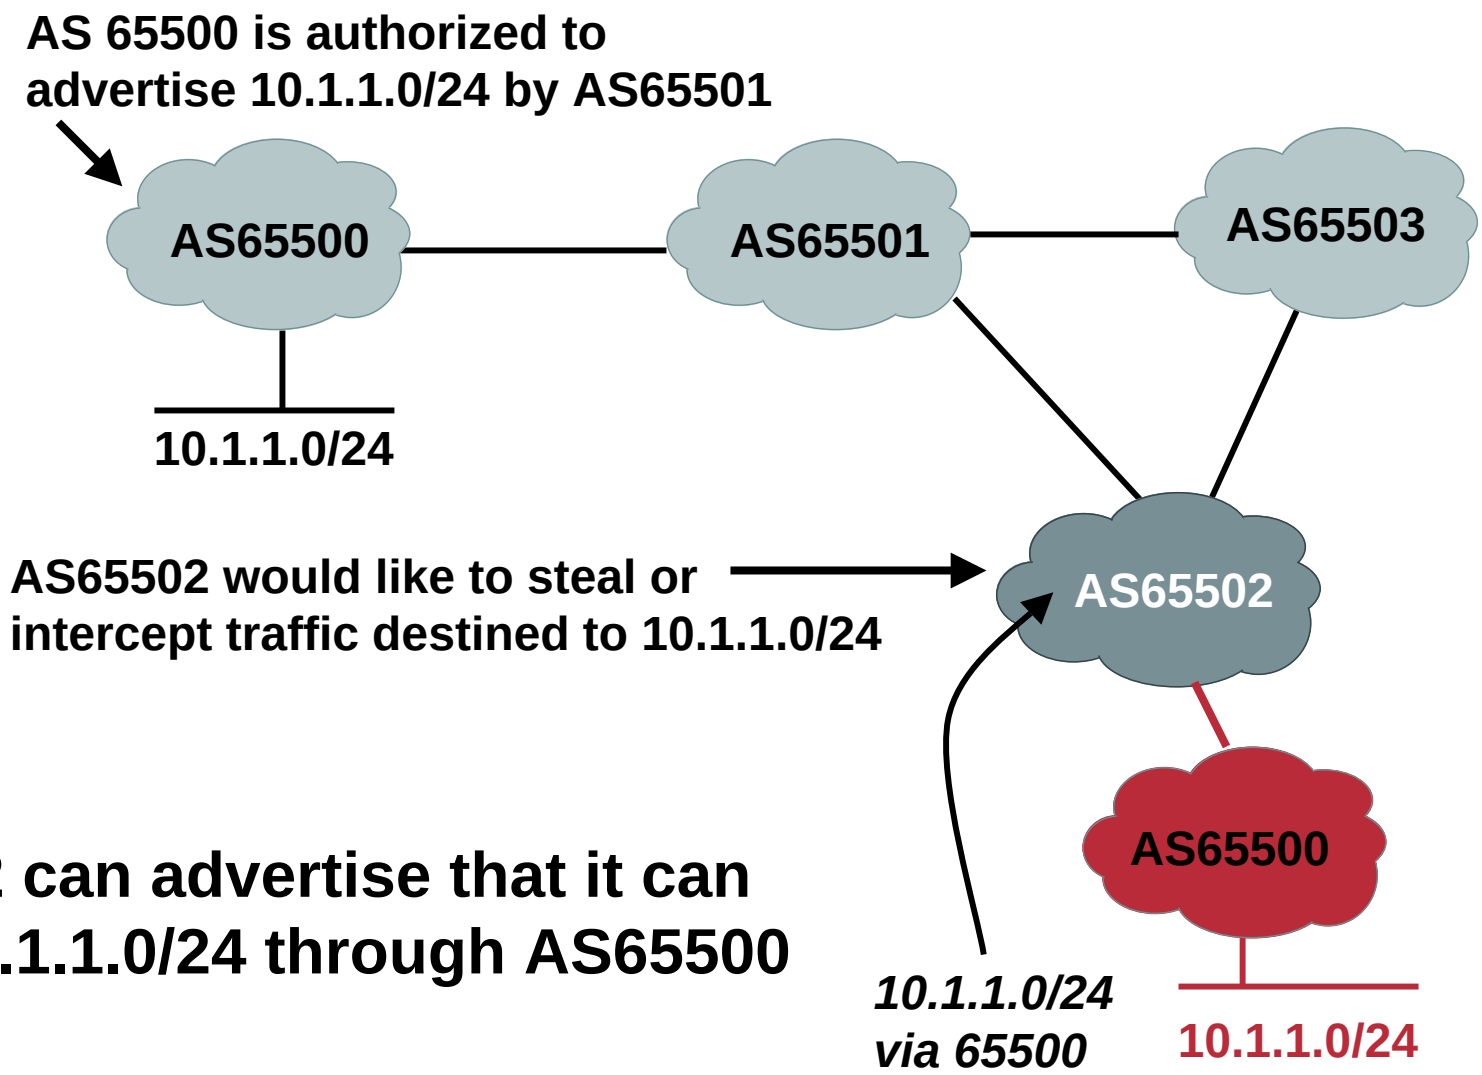

AS65502 can advertise that it can reach 10.1.1.0/24 through AS65500

Spoofed Peering Session (cont'd)

soBGP Solution:

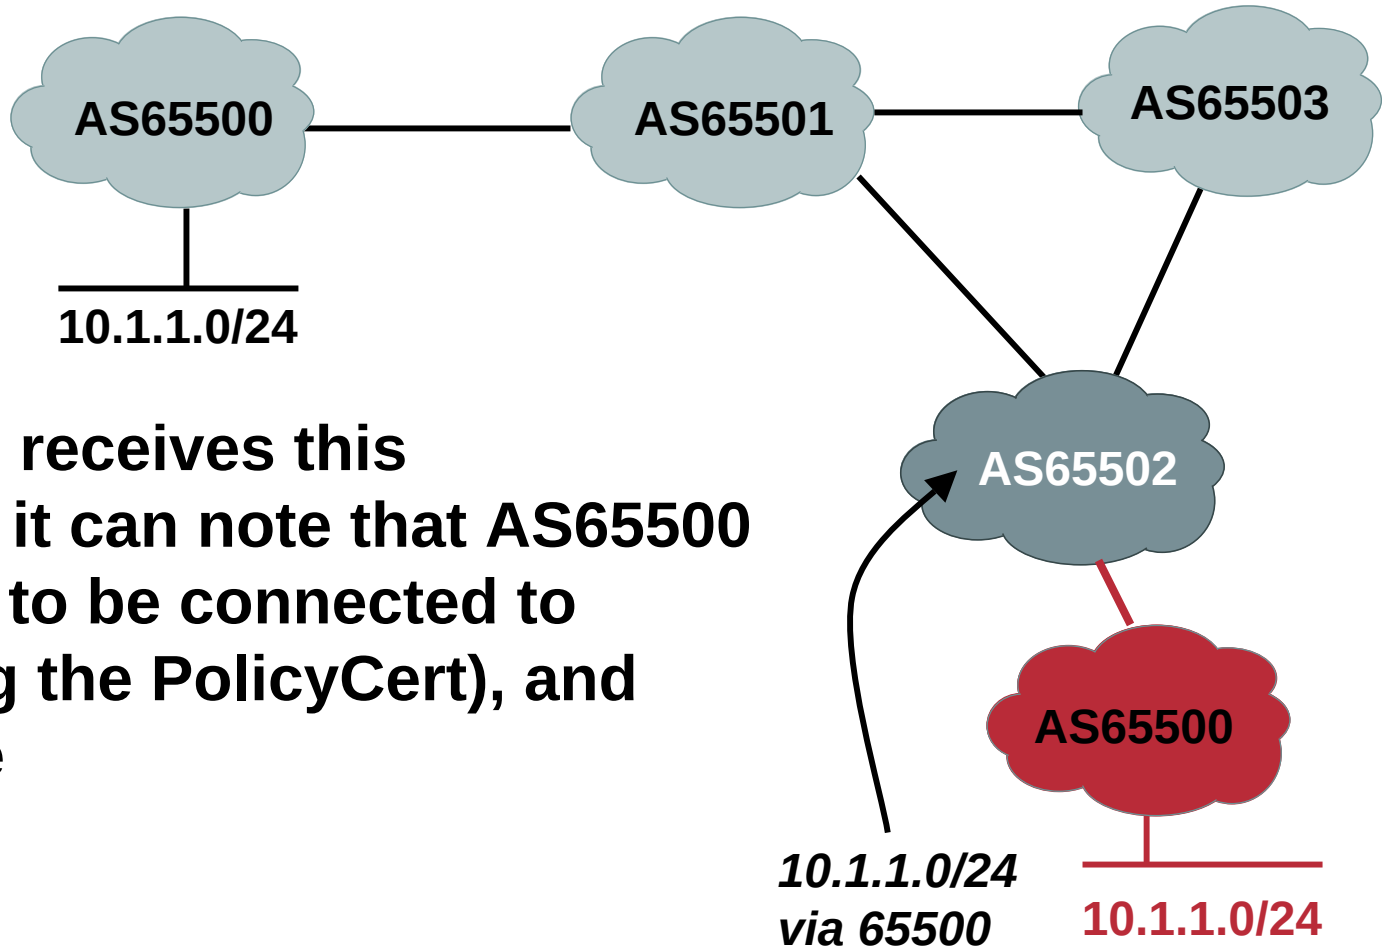

When AS65503 receives this advertisement, it can note that AS65500 does not claim to be connected to AS65502 (using the PolicyCert), and reject the route

Unauthorized Transit

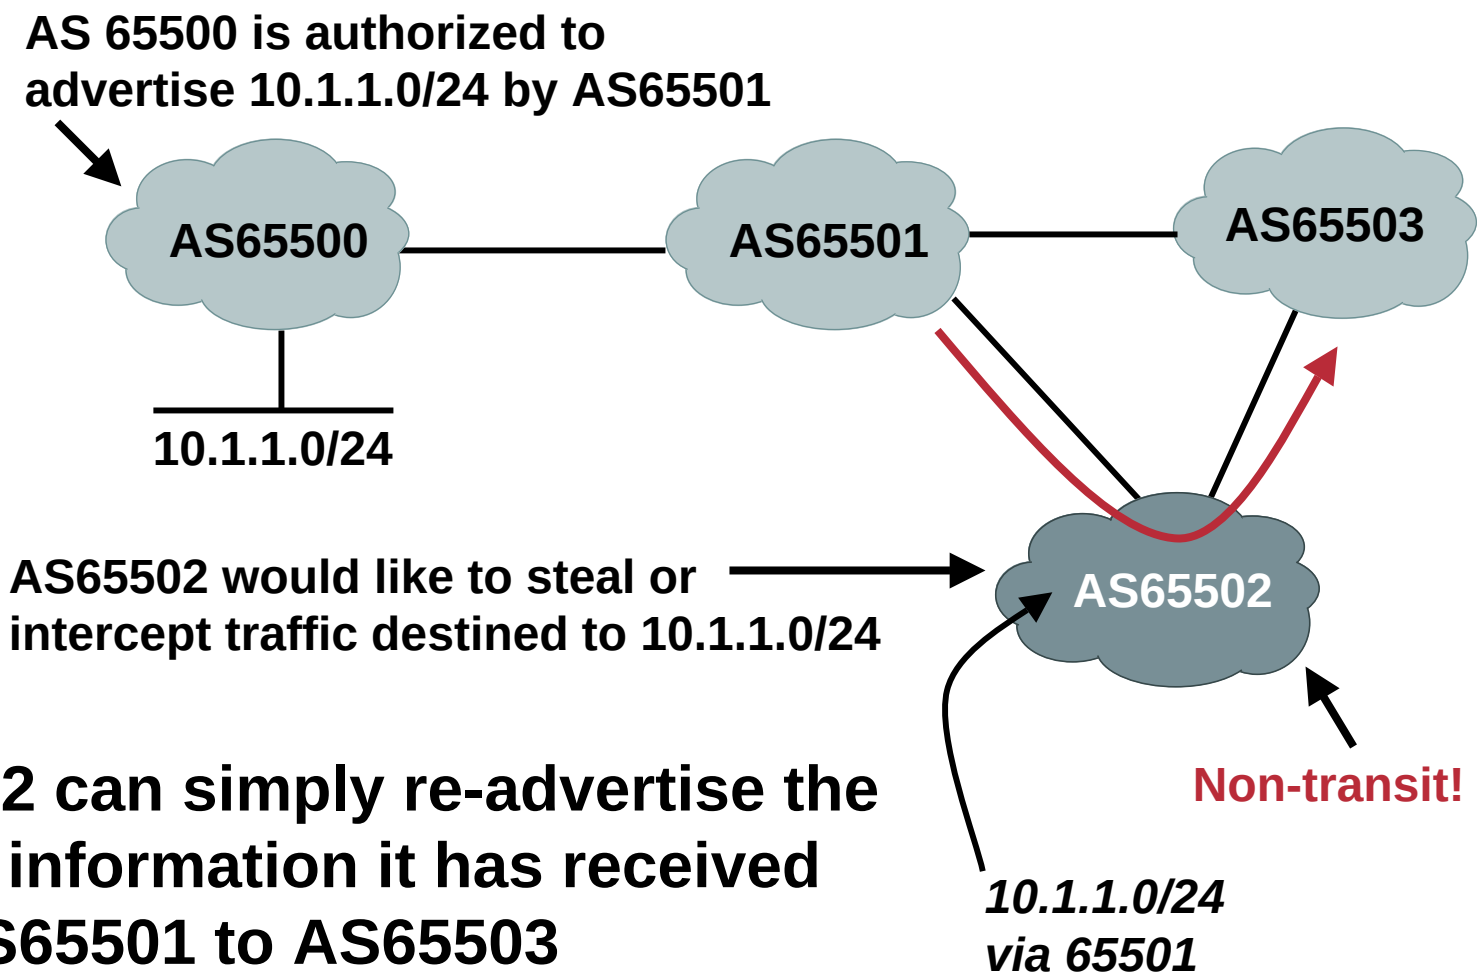

AS65502 can simply re-advertise the routing information it has received from AS65501 to AS65503

Unauthorized Transit (cont'd)

soBGP Solution:

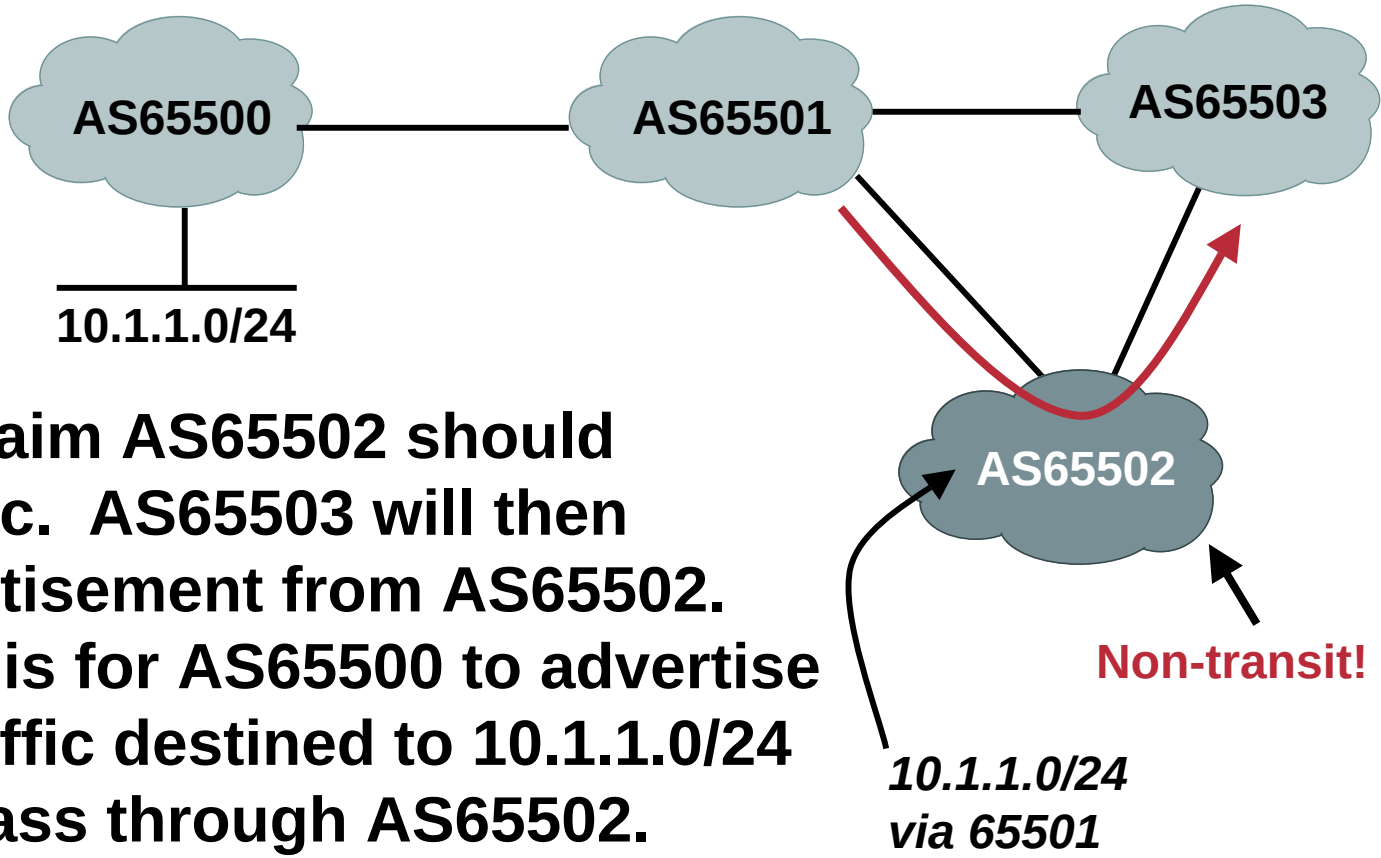

AS65501 can claim AS65502 should not transit traffic. AS65503 will then reject the advertisement from AS65502. Another option is for AS65500 to advertise a policy that traffic destined to 10.1.1.0/24 should never pass through AS65502. AS65503 will reject the advertisement

For More Information

soBGP:

ftp://ftp-eng.cisco.com/sobgp

The mailing list is open, archives are available, draft participation is encouraged!

Routing Protocols Security:

http://www.rpsec.org