


Protection of Communication Infrastructures

Chapter 4 Routing Security

- ❑ Routing Protocol Threats
- ❑ Countermeasures in context of BGP


Acknowledgement: These slides have been compiled from various sources (see last two slides on additional references)




General Threats to Routing Protocols

- ❑ Threats to routing can be characterized according to:
 - ❑ *Threat source*:
 - Subverted link or subverted / rogue router
 - ❑ *Threat consequence* (generic):
 - Disclosure of (routing) information
 - Deception of other routers (e.g. with forged messages)
 - Disruption of normal (router) operation
 - Usurpation (= gaining control over a router's operation, e.g. by "stealing" traffic originally to be routed by that router)
 - ❑ *Threat consequence zone*:
 - Single node / part of a network / whole Internet
 - ❑ *Threat consequence period*:
 - Only during attack / for a certain period of time


(characterization mostly according to [BMV06])



- ❑ Consequences regarding the network as a whole:
 - ❑ *Network congestion*: more traffic is routed through a specific part of the network than would usually be
 - ❑ *“Blackhole”*: packets go into a certain router/region and “disappear”
 - ❑ *Looping*: traffic is forwarded along a route that loops (this causes both traffic to disappear and congestion)
 - ❑ *Partitioning*: some portion of the network believes that it is partitioned from the rest of the network when in fact it is not
 - ❑ *Frequent route changes*: resulting in unnecessary routing processing and message exchanges as well as large variations in forwarding delay
 - ❑ *Instability of the routing protocol*: convergence towards a global forwarding state is not achieved
 - ❑ *Routing overload*: routing protocol messages become a significant part of the overall transported traffic the network carries




- ❑ Consequences regarding a specific target host / network:
 - ❑ *Delay and Jitter*: traffic from / to a target host / network is routed along routes that are inferior to the route the traffic would otherwise take
 - ❑ *Cut*: some part of the network believes that there is no route to the target host / network when, in fact, there is
 - ❑ *Starvation*: the traffic destined for the target host/network is routed to a part of the network that can not deliver it
 - ❑ *Eavesdropping*: traffic is routed through some router or network that would normally not “see” this traffic, so that an attacker can eavesdrop on the traffic or at least monitor the traffic pattern
 - ❑ *Controlled delivery or Greyhole attack*: traffic is routed through a router / network so that an attacker can selectively delay, delete or modify packets destined to a target host / network



Generally Identifiable Routing Threats (1)

- ❑ **Disclosing routing information:**
 - ❑ *Deliberate exposure of routing information:* e.g., by a subverted router in order to disclose routing information
 - ❑ *Eavesdropping on routing exchanges:* different attacking technique, also leading to disclosure of routing information
 - ❑ *Traffic analysis:* by eavesdropping on forwarded data traffic, an attacker can gain insight about routing information
- ❑ **Masquerade:**
 - ❑ An entity claims the identity of a router (also called spoofing)
 - ❑ Masquerade is usually performed in order to realize further attacks
- ❑ **Interference:**
 - ❑ An attacker inhibits the exchange of routing information between routers, e.g. by delaying or deleting routing messages or receipts, breaking synchronization, etc.
 - ❑ The consequence may be (partial) disruption of routing operations




Generally Identifiable Routing Threats (2)

- ❑ **Falsification of routing information:**
 - ❑ Either by an *originator (forging)* or a *forwarder (modification)*
 - ❑ **Overclaiming:**
 - Announcing better routes / link capacity than available
 - Goals can be to attract traffic to a certain area in order to control the traffic or to mislead the traffic so that it will not be delivered at all or with higher delay
 - Consequences for the network are potential overload of single routers, increase of overall traffic load
 - ❑ **Underclaiming:**
 - Announcing inferior routes / link capacities than actually exist
 - Potential goals are to keep traffic out of certain areas of the network, e.g. in order to avoid forwarding of traffic at certain routers or to increase attractiveness of alternative routes
 - Potential consequences are that certain destinations become unreachable, and the overall traffic load in the network increases (because packets take inferior routes)



- ❑ **Resource exhaustion:**
 - ❑ E.g. by an attacker that announces frequent changes in his routing information, or triggers a router to create an excessive amount of state information which can not be handled by other routers
 - ❑ Sometimes also referred to as *overload*
 - ❑ Goal is degradation / disruption of routing protocol operation
- ❑ **Resource destruction:**
 - ❑ *Link destruction*: either physically (“cutting”) or by strong interference
 - ❑ *Node destruction*: e.g. physically or logically by exploiting weaknesses in the router software (OS, routing software)
 - ❑ Depending on the network topology, the consequences can be either of local or global scope (single network / part of network unreachable or network partitioning)




- ❑ Sometimes the following terms are used in the routing security context:
 - ❑ *Sinkhole attack*: Regardless of a specific approach, an attack is called sinkhole attack, if it attracts more traffic than the attacker has usually direct access to, e.g., by *Overclaiming* to perform a *Greyhole attack*
 - ❑ *Wormhole attack*: Wormholes are additional rogue links under control of the attacker that may either exist *physically*, e.g., by a wireless amplifier, or *virtually*, e.g., by tunneling traffic. Here, not the routing protocol itself is attacked, but the topology is changed maliciously.

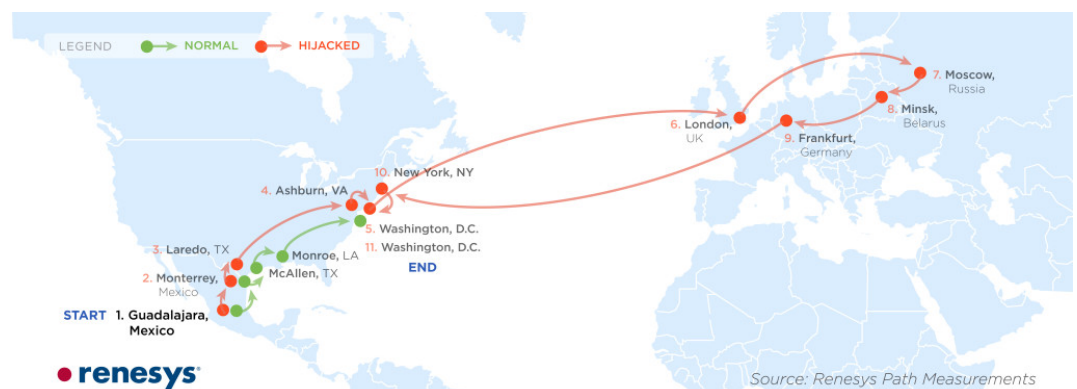
→ In the following, we will focus on Inter-AS routing threats and countermeasures, as it concerns availability of the Internet as a whole




- ❑ Inter-AS routing threats mainly concern BGP operation
- ❑ Attack Scenarios:
 - ❑ Disabling of parts of the Internet by disrupting Internet routing tables
 - ❑ Forcing multi-homed AS to use alternate paths to / from an outside AS instead of the preferred path
 - ❑ Disabling a single- or multi-homed AS
 - ❑ Creating traffic “blackholes”
- ❑ The above mentioned attack scenarios can e.g. be realized by:
 - ❑ announcing to “host” IP addresses ranges for that the attacker has no ownership
 - ❑ inserting unauthorized “prefixes” into routing table (= announcing paths for networks for which no authorization to route exists)
 - ❑ modifying or forging routing messages during transmission
 - ❑ resource destruction



- ❑ Early 2013 a Belorussian provider attracted traffic from GlobalOneBel over an uplink to Moscow
- ❑ Attacked networks changed daily but continued for a month
- ❑ Attracted traffic was forwarded to an unaffected uplink to Frankfurt
- ❑ Extremely difficult to detect: Servers and clients in Washington cannot do so, even with traceroutes!




Details: <http://www.renesys.com/2013/11/mitm-internet-hijacking>



Securing BGP Operation: Verifying Peer Messages (1)

- ❑ Forcing routers to accept only protocol messages from directly connected peers (if direct links exist):
 - ❑ Referred to as BGP TTL Security Hack (BTSH) [GHM03]
 - ❑ Idea – directly connected peer routers:
 - send routing messages with IP TTL field set to 255, and
 - accept only routing messages with IP TTL field ≥ 254
 - ❑ Messages from attackers which can only reach a target router over multiple hops will be discarded by router


- ❑ Question: why can this mechanism *not be implemented* as follows?
 - send routing messages with IP TTL set to 1, and
 - let routers in between automatically discard routing messages after one hop (so routing messages from attacker will not reach the target)



Securing BGP Operation: Verifying Peer Messages (2)


- ❑ More general approach:
 - ❑ *Generalized TTL Security Mechanism*
 - ❑ Standardized for IPv4 and IPv6 in RFC 5082 [GHM+07]
 - ❑ Routers set TTL=255, but may be multiple hops away
 - ❑ Packets are accepted depending on the distance, e.g., with TTL=253 when the router is two hops away
 - ❑ More configuration overhead, may be less secure than BTSH

- ❑ Better way: authenticate routing messages between peers
 - ❑ Protection of BGP sessions via the TCP MD5 signature option (RFC 2385)
 - ❑ Deploy IPsec between BGP peer entities (see chapter 11 in [Sch03])



- ❑ TCP MD5 Signature Option [Hef98]:
 - ❑ Goal: protect BGP exchanges between peers from spoofed TCP segments (attacker would only need to eavesdrop or “guess” correct sequence number)
 - ❑ Sender computes an MD5 hash value over each TCP segment and a secret shared with its peer entity
 - ❑ The hash value is transported in an option field
 - ❑ As all options in a TCP PDU together may not exceed 40 bytes this option has been defined to use 16 Byte long MD5 hash values (plus two bytes for TCP option information; type and length)

- ❑ Problem: MD5 is not state of the art, no automatic key negotiation / update procedure defined, leading to deployment difficulties (+ known vulnerabilities of manual key mgmt.)




- ❑ TCP Authentication Option [TMB10]:
 - ❑ Successor to TCP MD5 Signature with different algorithms
 - ❑ Better replay protection (even when TCP seq. numbers roll over)
 - ❑ Not (yet) widely deployed
 - ❑ Still no automatic key negotiation / update procedure defined
- ❑ Deployment of IPsec between peers:
 - ❑ Provides authentication and replay protection for IP packets
 - ❑ Allows for additional confidentiality
 - ❑ Leverages key management protocol that may use certificates and private keys
 - ❑ Potential problem: Low convergence speed when a router has many peers, e.g. > 1000, as key exchanges may take seconds per neighbor
- ❑ Sometimes routers may still be contacted from outside without (!) any authentication [CKV11]



- ❑ Address space “ownership” verification:
 - ❑ Who has been assigned an IP address range and has thus the right to announce this range / delegate the announcement of this range?
- ❑ Autonomous System (AS) authentication:
 - ❑ To whom has a claimed AS-number actually been assigned?
- ❑ Router authentication and authorization (relative to an AS):
 - ❑ Are the entities pretending to belong to an autonomous system authentic?
- ❑ Route and address advertisement authorization:
 - ❑ Who is allowed to announce specific address ranges / routes
- ❑ Route withdrawal authorization:
 - ❑ Who is allowed to withdraw a route?



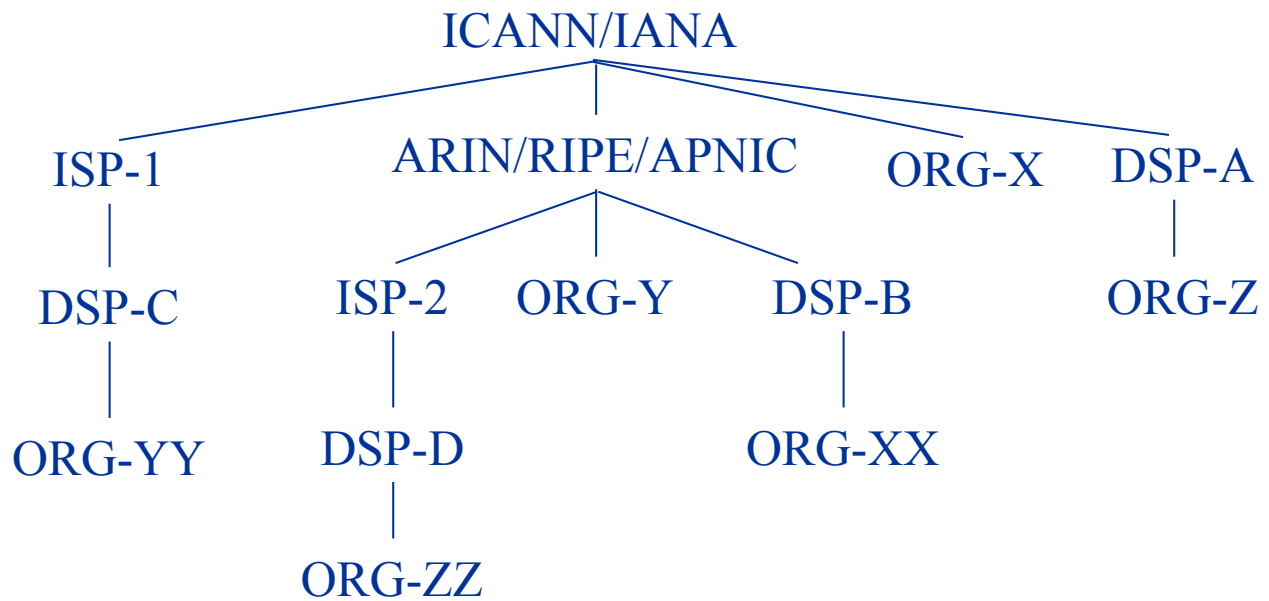
➔ Need for further security measures, one approach for this is S-BGP




- ❑ *IPsec*:
 - ❑ Provides authenticity and integrity of peer-to-peer communication with support for automated key management
- ❑ *Public Key Infrastructures (PKIs)*:
 - ❑ Secure identification of BGP speakers and of owners of AS's and of address blocks
- ❑ *Attestations*:
 - ❑ Authorization of the subject (by the issuer) to advertise specified address blocks
- ❑ *Validation of BGP UPDATES*:
 - ❑ Based on a new path attribute, using certificates and attestations
- ❑ *Distribution of security specific data*:
 - ❑ Certificates, certificate revocation lists (CRLs), attestations

(material on S-BGP partly taken from [Kent] and [Lynn99])






- ❑ Internet address space is managed hierarchically with the *Internet Assigned Numbers Authority (IANA)* acting as the root authority for assigning address ranges




- ❑ ICANN issues certificates for address space ownership to regional authorities and to entities that have direct address allocations (from IANA)
- ❑ Each of these certificates contains an extension specifying the address space being delegated, so that certificate validation is address-constrained
- ❑ Holders of address space certificates can create an address attestation, authorizing an AS (or a router) to advertise the specified address space
- ❑ Only networks that execute BGP need certificates:
 - ❑ All Internet Service Providers (ISPs) are BGP users,
 - ❑ Only about ~10% of Downstream Providers (DSPs),
 - ❑ Maybe 5% of subscribers, are BGP users


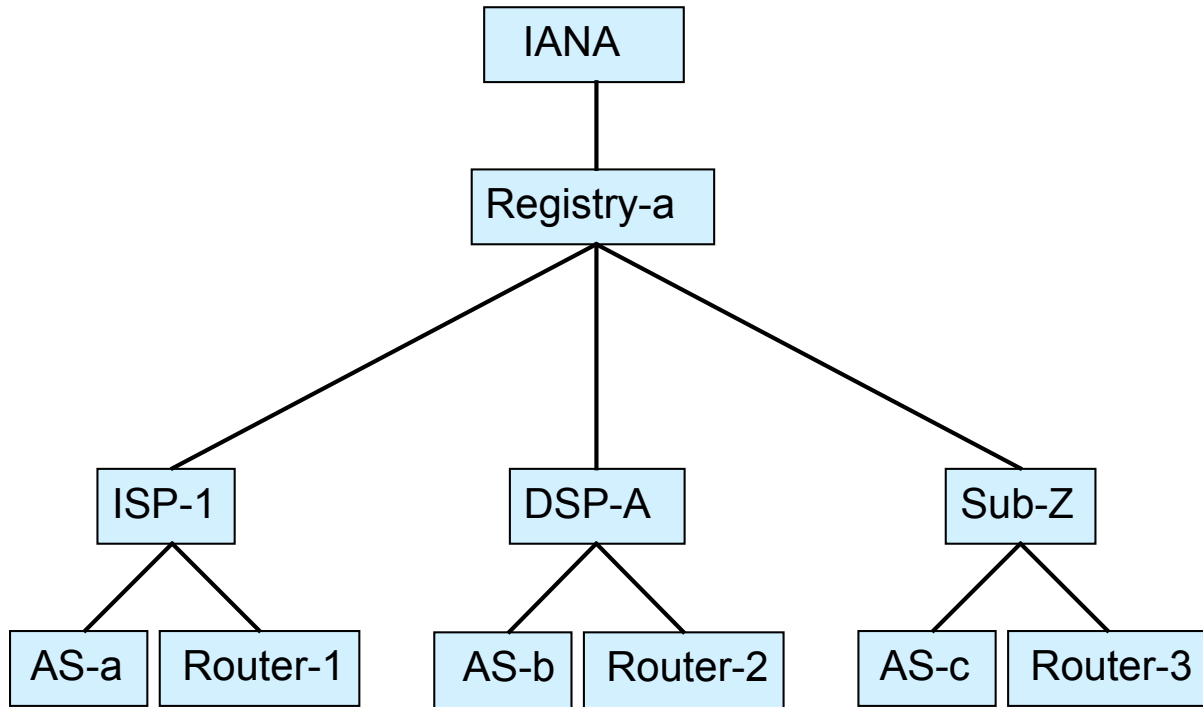


- ❑ ICANN issues certificates for AS ownership to:
 - ❑ ISPs, DSPs, and organizations that run BGP
- ❑ AS operators issue certificates to:
 - ❑ Routers as AS representatives
- ❑ Holders of AS (or router) certificates generate route attestations that:
 - ❑ authorize the advertisement of a route by a specified next hop AS
 - ❑ are used to express a secure route as a sequence of AS hops
- ❑ Securing an UPDATE:
 - ❑ A secure UPDATE consists of an UPDATE message with a new, optional, transitive path attribute for route authorization
 - ❑ This attribute consists of a signed sequence of route attestations, nominally terminating in an address space attestation
 - ❑ This attribute is structured to support both route aggregation and AS sets
 - ❑ Validation of the attribute verifies that the route was authorized by each AS along the path and by the ultimate address space owner




	Issuer	Subject	Extensions
Root Certificate	IANA	IANA	all addr
Registry Certificate	IANA	Registry	addr blocks
ISP/DSP Certificate	Registry (or IANA)	ISP/DSP	addr blocks
Subscriber Certificate	ISP/DSP (or Registry, IANA)	Subscriber	addr blocks






	Issuer	Subject	Extensions
Root Certificate	IANA	IANA	all ASes
Registry Certificate	IANA	Registry	ASes
AS Owner Certificate	Registry (or IANA)	ISP/DSP or Subscriber	ASes
AS Certificate	ISP/DSP or Subscriber	AS	
Router Certificate	ISP/DSP or Subscriber	Router*	AS, RtrId

* the subject name could be a fully-qualified DNS name



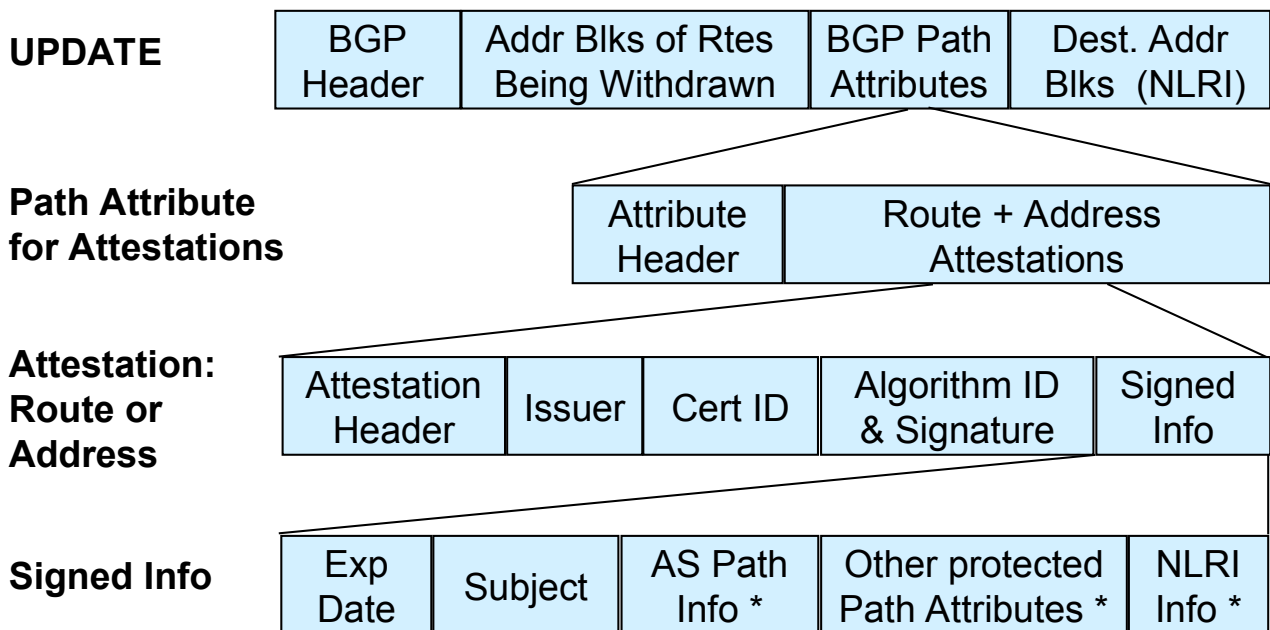

- ❑ Address Attestations:
 - ❑ Used to validate that a destination address block is being originated by an authorized AS
- ❑ Route Attestations:
 - ❑ Used to validate that an AS is authorized to use an AS Path
- ❑ Each UPDATE includes:
 - ❑ one or more Address Attestations, and
 - ❑ a set of Route Attestations
- ❑ These are carried in a new, optional, transitive BGP Path Attribute



- ❑ Address Attestations include identification of:
 - ❑ address blocks,
 - ❑ their owner's certificate,
 - ❑ AS authorized to originate (advertise) the address blocks, and
 - ❑ expiration date/time
- ❑ Indicate that the AS listed in the attestation is authorized by the owner to originate/advertise the address blocks in an UPDATE
- ❑ Digitally signed by owner of the address blocks, traceable up to the IANA via a certification path
- ❑ Used to protect BGP from erroneous UPDATES (authenticated but misbehaving or misconfigured BGP speakers)

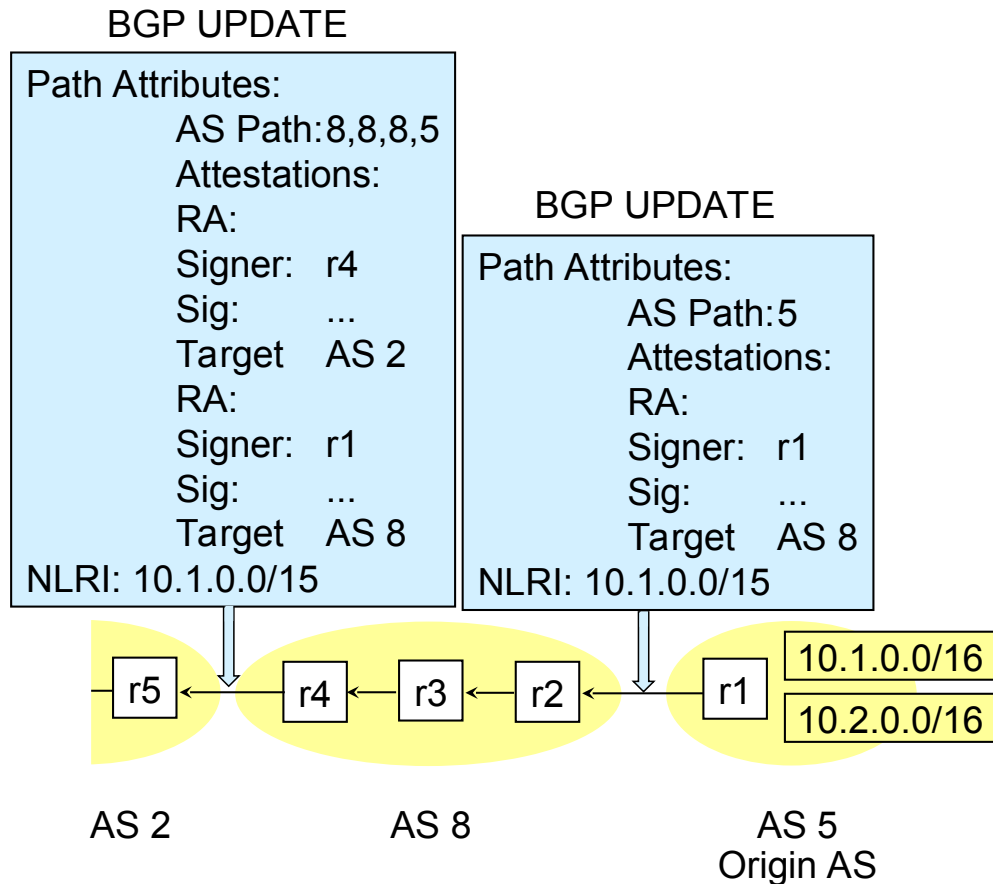


- ❑ Include identification of:
 - ❑ AS's or BGP speaker's certificate issued by the AS owner,
 - ❑ the address blocks and the AS Path (ASes) in the UPDATE,
 - ❑ the AS number of the receiving (next) neighbor, and
 - ❑ expiration date/time
- ❑ Indicate that the BGP speaker or its AS authorizes the receiver's AS to use the AS Path & NLRI in the UPDATE
- ❑ Digitally signed by owner of the BGP speaker (or its AS) distributing the UPDATE, traceable to the IANA ...
- ❑ Used to protect BGP from erroneous UPDATES (authenticated but misbehaving or misconfigured BGP speakers)




*explicit in the aggregation case, or if Path Attribute changes unpredictably






- ❑ To validate a route from AS_n , AS_{n+1} needs:
 - ❑ 1 address attestation from each organization owning an address block(s) in the network layer reachability information (NLRI),
 - ❑ 1 address allocation certificate from each organization owning address blocks in the NLRI,
 - ❑ 1 route attestation from every AS along the path (AS_1 to AS_n), where the route attestation for AS_k specifies the NLRI and the AS Path up to that point (AS_1 through AS_{k+1}),
 - ❑ 1 certificate for each AS along the path (AS_1 to AS_n) to use to check signatures on the route attestations, and
 - ❑ of course, all the relevant certificate revocation lists (CRLs) must have been verified (in case a private key was compromised and the corresponding certificate must be revoked)


- ❑ Putting certificates & CRLs in UPDATES:
 - ❑ would be redundant and make UPDATES too big
 - ❑ same is true for address attestations
- ❑ Solution – use servers for these data items:
 - ❑ replicate for redundancy & scalability
 - ❑ locate at network access points (NAPs = multiple BGP speakers interconnected with high speed LANs) for direct (non-routed) access
- ❑ Download options:
 - ❑ whole certificate/AA/CRL databases
 - ❑ queries for specific certificates/AAs/CRLs
- ❑ To minimize processing & storage overhead, network operations centers (NOCs) should validate certificates & AAs, and send processed extracts to routers
 - ❑ However, in this case trust is delegated to the NOC!



- ❑ Certificates (generation and signing done offline):
 - ❑ Disk space for storing certificates
 - ❑ CPU resources for validating certificates
- ❑ CRLs (generation and signing done offline):
 - ❑ Disk space for storing CRLs
 - ❑ CPU resources for validating CRLs
- ❑ Attestations:
 - ❑ Routing Information Base (RIB) memory space for storing attestations
 - ❑ Disk space for faster recovery from router reboot (optional)
 - ❑ CPU resources for signing and validating attestations
 - ❑ Resources for transmitting attestations (to make this a dynamic system)
- ❑ Size of the problem (June 1999):
 - ❑ ~ 5,300 AS, ~ 44,000 owners of address prefixes, ~ 7,500 BGP speakers
 - ❑ Resulting certificate database size: ~ 26Mbyte (~ 450 byte / certificate)
 - ❑ CRLs would add to this (should not be too much)




- ❑ Remaining vulnerabilities:
 - ❑ Failure to advertise route withdrawal
 - ❑ Premature re-advertisement of withdrawn routes
 - ❑ Erroneous application of local policy
 - ❑ Erroneous traffic forwarding, bogus traffic generation, etc. (not really a BGP issue, since BGP deals with routing, but not traffic forwarding)
 - ❑ Erroneous topology changes, e.g., wormholes
- ❑ (Non-)Deployment:
 - ❑ Up to now, only tests have been conducted in prototype environments
 - ❑ Discussion on S-BGP calmed down ... but raising again with BGPSEC
 - ❑ Nevertheless, S-BGP still:
 - shows the tasks to be accomplished regarding certification of IP address ownership, AS# ownership, and authorization to advertise certain routes
 - gives an impression on the scale of the effort that has to be invested in order to secure a global-scale Inter-AS routing protocol




- ❑ Drawbacks of S-BGP lead to development of different approaches
- ❑ Most important are:
 - ❑ S-A (Signature Amortisation)
 - Optimizes S-BGP signature system ➔ less computing intensive
 - ❑ SPV (Secure Path Vector) Protocol
 - only symmetric cryptography in routers ➔ less computing intensive
 - ❑ IRV (Interdomain Route Validation)
 - Offline distribution of BGP updates to validation servers
 - ➔ allows partial deployment
 - ➔ without expensive cryptographic operations
 - ❑ soBGP (secure origin BGP)
 - Web-of-Trust to authenticate routers ➔ avoids expensive PKI
 - ❑ BGPSEC
 - Another try of S-BGP authors for a better deployable S-BGP with high security standards




- ❑ Observation: messages are more often validated than signed
- ❑ S-BGP uses DSA to sign messages
- ❑ RSA (with certain parameters) can validate signatures faster than DSA
 - ❑ But higher signing effort required, as RSA keys of equal security operate in a larger field
 - ➔ Simulations showed that there is no direct gain



- ❑ Observation 2: S-BGP forwards UPDATE messages individually
 - ❑ E.g. upon forward to three of five peers, three individual packets with three signatures are generated
- ❑ Generate only a single packet and address neighboring peers in a bit vector (e.g. three bits set to one), sign once and distribute to affected peers
 - ❑ Called Signature Amortization Across Peers (S-A-P)
 - ❑ Which peers corresponds to which bit is to be set in the certificate of the peers
 - ➔ Peer associations change infrequently, so it is feasible
 - ➔ less signing effort required
 - ➔ gain of using RSA increases




- ❑ Merkle Hash Trees: Each node is defined to be the hash of its successors
- ❑ Only leaf nodes are directly hashed
- ❑ E.g. $h_{0-7} := H(h_{0-3} || h_{4-7})$, where $||$ represents the concatenation operator




- ❑ Observation 3: BGP-Routers accept new messages only after an MRAI (Minimum Route Advertisement Interval) to avoid instabilities
- ❑ Only a single signature operation is required for all packets in an MRAI, if they are buffered
 - ❑ Signature Amortization Across Output Buffers (S-A-B)
 - ❑ Constructs a Merkle Hash Tree of all messages in buffers
 - ❑ Signs the root of the tree asymmetrically
 - ❑ Signatures consist of the digital signature and all values that are required to generate the root of the tree
 - ❑ E.g. the signature of m_3 is h_{0-1} , h_2 , h_{4-7} and the RSA signature of h_{0-7}
 - ➔ less signing and verification effort required


- ❑ Note: Positive effects of S-A-P and S-A-B do not add
 - ❑ Different MRAI timers in peers
 - ❑ No further speed up
- ❑ Both lead to a (slight) decrease in S-BGPs convergence time (due to faster signature processing)
- ❑ Exploit BGP behavior to reduce signing and verification effort



- ❑ Digital Signatures are computing intensive, as they are done by asymmetric cryptography
- ❑ Lamports Signatures use only cryptographic hash functions
- ❑ Step 1:
 - ❑ Alice generates random numbers tuples $(s_{i,0}, s_{i,1})$ being her secret key, say 128 tuples ($i=1, \dots, 128$) of values, each one of length 128 bit
 - ❑ Publishes the hashes of all values as her public key $(H(s_{i,0}), H(s_{i,1}), \dots)$
- ❑ Step 2:
 - ❑ Alice signs a message m by generating $H(m)$ (output length 128 bit)
 - ❑ She selects and publishes 128 values of her secret depending on the hash to be her signature, e.g. if the value of bit i is “0” she takes $s_{i,0}$ and $s_{i,1}$ otherwise.
- ❑ Step 3:
 - ❑ Bob verifies the signature by generating $H(m)$ and checking if the hash of the published secret key values are the correct selection public key of values



- ❑ Advantage
 - ❑ Does not use asymmetric cryptography
- ❑ Disadvantages
 - ❑ Keys can only be used once, as parts of the private key are disclosed
 - ❑ Large signatures and keys
- ❑ However several improvements exist...




- ❑ A naïve Lamport Signature Scheme improvement
- ❑ Public key is root of a Merkle Hash Tree
- ❑ Secret keys are all leaf nodes
- ❑ A signature is derived as follows:
 1. Calculate $H(m) \bmod n = i$, where n is the number of secret leaves
 2. Disclose s_i and all hash values that are needed to derive the root
 3. Recipient uses hash values to calculate root value ➔ Authentication
 4. Compares position of s_i to $H(m) \bmod n$ ➔ “Integrity”
- ❑ E.g. let $H(m) = 326$, $n=8$ (corresponds to lower $\log_2 n$ bit of $H(m)$)
 - ❑ $i = 6$
 - ❑ Signature is $h_{0-3}, h_{4-5}, s_6, h_7$
- ❑ Advantage: Public key and signature require less space
- ❑ Problem: n has to be chosen large ($> 2^{80}$) enough to avoid that a new message is found for a known signature
- ❑ But $n = 2^{80}$ is infeasible!



- ❑ Idea: Publish more than one part of the private key
- ❑ HORS signature scheme selects m of the n private key parts
- ❑ Takes not only the last bits to identify a single private key, but $m \cdot \log_2 n$ bits


- ❑ Leading to $\binom{n}{m}$ combinations of parts of the key
- ❑ The chance of finding a new message that complies to an existing signature is therefore reduced to $\binom{n}{m}^{-1}$
- ❑ E.g. $\binom{8}{3} = 56$, which is way better than 8 possibilities in the hash tree mechanism
- ❑ Recommended parameters for 80-bit hash functions
 - ❑ $n = 1024, m = 20 \approx$ “113 bit security”
 - ❑ $n = 256, m = 24 \approx$ “81 bit security”



- ❑ What happens if a key is used more than once?
- ❑ Security decreases as more parts of the private key are disclosed
 - ❑ Attackers can use more parts of the private key to generate valid signature for arbitrary signatures
- ❑ Probability for an effective attack is about $\left(\frac{r \cdot m}{n}\right)^m$ (for large values of n and m), where r is the number of times a key is reused
- ❑ Therefore effective security of $n = 1024, m = 20$ decreases from 2^{-113} to 2^{-73} after issuing four signatures ($r = 4$)

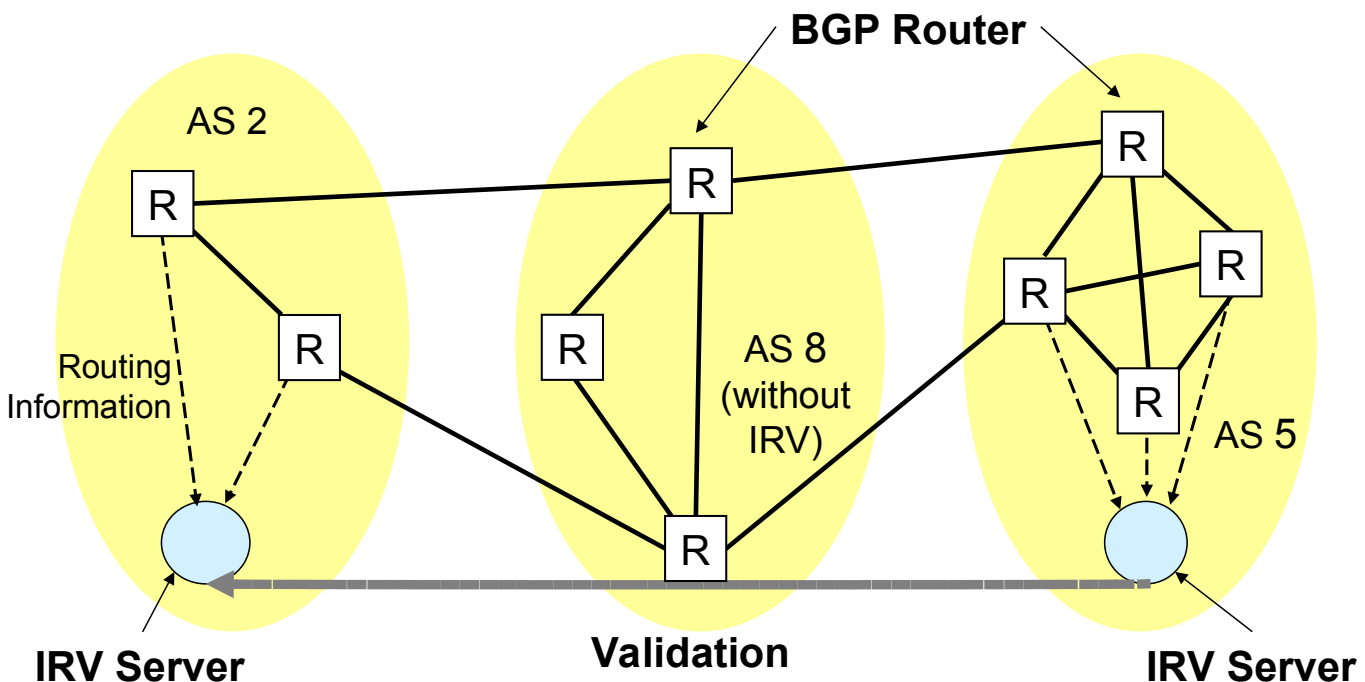


- ❑ SPV utilizes modified HORS signatures with $n = 256$, $m = 6$
- ❑ Performs 15 signatures without changing the key ($r = 15$)
- ❑ Lower effective security required as attackers can only toggle 2^{16} AS-values
- ❑ Uses a complicated mapping scheme to create „ASPATH protectors“
- ❑ In fact so complicated that
 - ❑ Severe flaws were found in SPV [RPM07]
 - ❑ Attackers can modify AS-Paths, if certain topology criteria are met
 - ❑ The required hash operations make it not faster than S-A
 - ❑ The communication overhead compared to S-BGP is 273%
- ❑ But we learned:
 - ❑ A complex system to use fast operations is not necessarily better than a simple system with slow operations




Interdomain Route Validation (IRV)

- ❑ Direct BGP extension
- ❑ Decentralized system with focus on interoperability
- ❑ Central IRV server in each participating AS



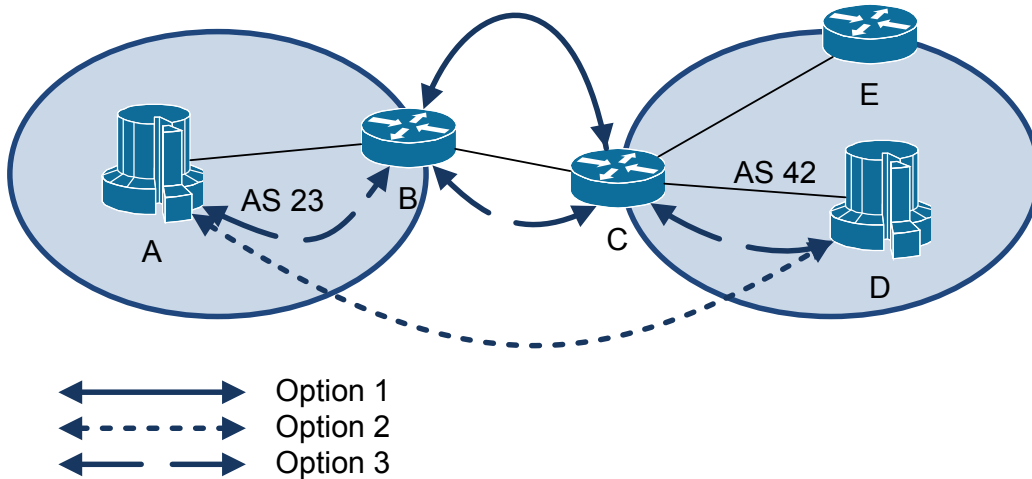
- ❑ Each IRV server represents a single AS
- ❑ AS validates received BGP UPDATE message by querying other IRV server along the AS-Path subsequently
 - ❑ IRV server are found by fields embedded into update messages or certificates in a central register
- ❑ No own authentication and access control methods, but the use in combination with IPsec and TLS is proposed
- ❑ A PKI similar to S-BGP could be deployed to support authentication
- ➔ Not entirely suited against attackers as responding IRV servers might themselves be compromised
 - ➔ No address and no route attestations for validating IRV server's answers
- ➔ Additional barrier against misconfiguration and “dumb” attackers
 - ❑ Most outages are caused by these effects
- ❑ BUT: What about robustness or a “clean reboot”?



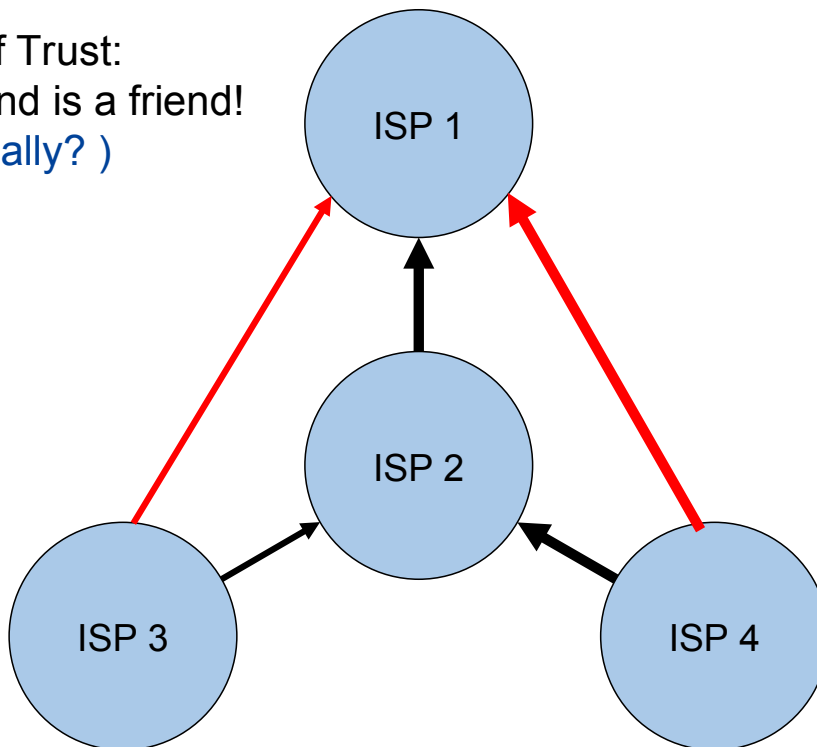
- ❑ Expired IETF Draft initiated by Cisco, based on BGP [Whi06]
- ❑ Further development objectives compared to S-BGP
 - ❑ Shall work without PKI (too expensive (\$), concentrates power to holder of root certificate)
 - ❑ Should not require working Internet routing (for fetching CRLs, etc.)
 - ❑ Incremental deployment
- ❑ Introduces new SECURITY message to transport certificates (e.g. route attestations)
- ❑ UPDATE messages are not altered and backward compatible
- ❑ Several known mechanisms
 - ❑ AuthCerts ≈ S-BGP Address Attestation
 - ❑ EntityCerts ≈ S-BGP AS Certificates
 - ❑ PolicyCerts – each AS lists all its BGP connections to neighboring ASes (functionality comparable to S-BGP Route Attestations)



- ❑ Other known mechanisms: the deployment options
 - ❑ Option 1: Edge routers perform certificate exchange and processing (like in S-BGP)
 - ❑ Option 2: Internal servers exchange certificates and validate them (like in IRV)
 - ❑ Option 3: Compromise of option 1 and 2
 - Routers exchange information
 - Processing swapped to internal servers



Web of Trust:
A friend's friend is a friend!
(→ really?)




Indirect Trust →
Direct Trust →

- ❑ Deploys Web-of-Trust instead of PKI in difference to S-BGP
 - ❑ Each ISP signs identity of other ISP it trusts
 - ❑ Indirect trust over certificate paths
 - ❑ Associations are stored aggregated in a trust database
 - ❑ Decentralized approach to security
 - ❑ Configurable “trust depth”


- ❑ Q: Why is this different from BGP peering policies?

- ❑ Security not proofed (no actual address attestation by a TTP!)
 - ➔ Colluding attackers can circumvent security by lying in PolicyCerts
 - ❑ Even less guarantees if deployed incrementally




- ❑ Standardization (partly still standardization effort) for an approach based on S-BGP principles
- ❑ Several changes
 - ❑ Split in BGPSEC [LT13] (mostly responsible for routing attestations) and Resource Public Key Infrastructure (RPKI) [LK12] (= directory responsible for Secure Origin Authentication)
 - ❑ Certificate information is replicated among distributed servers
 - ❑ Signatures are distributed by BGP UPDATES non-transitively
 - Allows for BGPSEC negotiation between routers
 - Support of “BGPSEC islands”
 - ❑ Several optimizations with regard to efficiency
 - ❑ No IANA root certificate

- ❑ Status: RPKI is rolled out, adoption is slowly progressing
 - ❑ See: <https://blog.apnic.net/2020/01/29/is-rpki-ready-for-the-big-screen/>




RPKI (State as of Late April 2014)




- Unknown
- Invalid AS
- Prefix Len Exceeded
- Valid
- Prefix Len Mismatch
- Invalid AS & Prefix Mismatch

Even for routes with RPKI information way too many false positives!




RPKI (State as of 2019)




Growth of RPKI in terms of:

1. Number of VRRP IPv4 prefixes
2. Percentage of ASes where some of their IPv4 addresses are covered by VRRPs to all ASes managed by the RIR
3. Percentage of IPv4 addresses covered by VRRPs to all assigned IPv4 addresses for the RIR

(Source: <https://blog.apnic.net/2020/01/29/is-rpki-ready-for-the-big-screen/>)



- ❑ Drawbacks of seen cryptographic approaches
 - ❑ Computation and communication intensive
 - ❑ Usually public-key infrastructures or central databases needed
 - ❑ Incremental deployment with little security gain
 - Even soBGP and IRV require a certain number of peers for gaining any security advantage
 - ❑ Do not help against routing instabilities caused by attackers, wormhole attacks etc.
- ❑ Idea: Use available information to check credibility of BGP Update messages
- ❑ (Some) interesting approaches:
 - ❑ Pretty Good BGP: Cautiously Adopting Routes
 - ❑ Topology-based Analysis
 - ❑ Stable Route Information Objects
 - ❑ Monitoring TCP Flows




- ❑ Observation: Almost half of bogus origin/prefix associations last less than 24 hours
- ❑ Idea: Treat unfamiliar routes cautiously
 - ❑ Time for a secondary validation process (manual, Internet Alert Registry, or by others)
 - ❑ Exploits natural redundancy, as other older routes still exist
- ❑ First step: identifying normal routes
 - ❑ Routers store history of known origin/prefix pairs for h days
 - ➔ Database defines normal behavior
- ❑ Second step: detect anomalous routes
 - ❑ Received route updates compared with database
 - ❑ Updates altering the normal state
 - ➔ Marked *suspicious* for s days (“*suspicious period*”)
 - ❑ After s days, suspicious routes added to the history



- ❑ Third step: avoiding suspicious routes
 - ❑ Suspicious routes get lowest possible preference
 - ❑ Routers select best trusted route (if possible)
 - ➔ False positives possible (less desirable route)
 - However, routing operates normally

- ❑ Drawback of approach: If new subprefixes are introduced (or generated by an attacker)
 - ❑ Routers will use known route to the larger address block during suspicious period
 - ➔ Leads to false positives: Potentially better path to new (valid) subprefix not used during suspicious period


- ❑ All attacks persisting longer than suspicious period are successful, as new routes are not tested.



- ❑ Observation: Internet exhibits certain structure
 - ❑ Densely connected *core* nodes (backbone)
 - ❑ *Periphery* nodes with connection to the core and at most a few direct neighbors


- ❑ Connectivity graph
 - ❑ Routers are nodes, direct links are edges
 - ❑ Can be approximated with information from route updates (combine several routers)

- ❑ Yellow and Red AS have many links (up to 1845), Blue AS have few links to other AS




- ❑ Remove core nodes from Graph
- ➔ Clusters of periphery nodes
- ❑ Routers with access to geographical data can determine the diameter of a cluster
 - ❑ Maximum geographical distance between two systems within a cluster
 - ❑ Diameter of most clusters is small (local networks connected to large providers)
 - ❑ Kruegel at al. use preprocessed information from the whois databases to determine geographical positions
 - ❑ Example excerpt of a whois record:

```
inetnum:      141.24.0.0 - 141.24.255.255
netname:      THILM-NET
descr:        Technische Universitaet Ilmenau; Rechenzentrum
descr:        Helmholtzring 9
descr:        98684 Ilmenau
country:      DE
admin-c:      KH195
```




Valid routes must satisfy constraints:

- ❑ A valid route has only one single subsequence of core nodes
 - ➔ Identify so called path modification attacks
 - ❑ In the example the sequence goes through core nodes before AS 43 and after AS 666, hence considered invalid




Valid routes must satisfy constraints:

- ❑ All consecutive pairs of periphery nodes in a route must be in a cluster or close geographical range (a 300km threshold proposed for the Internet)
 - ➔ Identify path truncation attacks
- ❑ In the example the direct link between AS 6 and AS 43 is a violation of the constraint



- ❑ Advantages:
 - ❑ Simple algorithm
 - ❑ Improves security at reasonable costs
- ❑ Drawbacks:
 - ❑ Gathering reliable geographical data is very difficult and introduces an additional requirement
 - ❑ No algorithm to dynamically update the connectivity graph and the clustering
 - ❑ Attacks within a cluster still possible
 - ❑ Who guarantees the valid geographic information in whois records?

- ❑ Routes change dynamically but are comprised of two basic *route information objects*
 1. Direct links between neighboring systems
 2. Prefix/origin associations
- ❑ Inter-domain routing
 - ❑ Stable routing infrastructure
 - ❑ Route information objects also relatively stable
- ➔ Develop a historical database to compare UPDATE messages with
- ❑ Similar to previous approaches, but with directed links (policies)
- ❑ Less computing intensive (only own data analyzed) more accurate (as shown in simulation by authors) compared to previous approaches.
- ❑ Algorithm
 - ❑ Check each directed link beginning with observer
 - ❑ First link not found returned as suspicious (i.e. potential last modification)
 - ❑ Check if prefix / origin association has been in database before (aggregations and de-aggregations considered legitimate)




Improve quality of the database by several heuristics


- ❑ Removing transient objects
 - ❑ Uptime threshold for prefix / origin associations
 - ❑ Lifetime criterion for directed links
- ❑ Considering route updates with no profit for an attacker as legitimate
 - ❑ Route only modified downstream the former (already trusted) origin and within announced address range of the former origin
 - ❑ Routers already on a trusted path are said to have no motivation to hijack or spoof routes
 - ❑ Problem: May maliciously redirect traffic for policy reasons



- ❑ Consider common practices on the Internet
 - ❑ Neighboring autonomous systems sharing a direct link often have similar address ranges and even share address ranges they announce, for an example for a common customer with two uplinks
 - Attackers may then hijack traffic for a neighbor AS
 - ❑ ASes can expand existing prefixes to some degree, as Internet registries might have assigned new addresses to the AS
 - Attackers may then hijack traffic for a similar address range
- ❑ Event-based clustering
 - ❑ Too many heuristics increase the number of false negatives
 - ❑ But few are enough
 - ❑ Messages in one cluster often share a common cause
 - ➔ If some of them are bogus, others might be as well




- ❑ Observe TCP traffic in the data plane
 - ❑ Can detect reachability problems
 - ❑ If problems are detected revert to alternative routes
- ❑ Routers observe packets in only one direction
- ➔ Prefix unreachable, if:
 - ❑ During period t no complete TCP handshake is observed
 - ❑ t is maximum of the time it takes to observe N incomplete flows to different destinations and a predefined period T
- ➔ Prefix reachable, if:
 - ❑ Complete TCP connections are observed
 - ❑ SYN followed by DATA (within timeout)
 - ❑ Indicates reachability of a prefix
- ❑ But: Incomplete TCP connections not necessary indicator of a problem
 - ❑ Destination hosts unavailable
 - ❑ Port scanners creating lots of SYN packets




- ❑ Attackers, being aware of the countermeasure can
 1. Perform TCP-SYN-Flooding to simulate unreachability
 2. Attacker can simulate complete TCP connections by sending SYN Packets followed by DATA in order to camouflage other attacks

- ❑ Possible countermeasures against those attacks
 - ❑ Router drops or delays a few (random) SYN packets and checks for retransmission
 - ❑ Router also checks that not too many SYN packets of connections it did not drop or delay are retransmitted (attacker might retransmit all)
 - ➔ False positives can be reduced with appropriate thresholds for both checks
 - ➔ Leads to lower quality of service




- ❑ Does not solve all BGP security issues, but:
 - ❑ deploys a first line of defense
 - ❑ requires low additional computation power and communication overhead
 - ❑ prevents misconfiguration (>50% of all BGP “attacks”)
 - ❑ is an example for the effective use of gathered information

- ❑ However:
 - ❑ Mechanisms have an inherent risk of being exploited by attackers



Routing Security References (1)

- [BFM10] BUTLER, K. ; FARLEY, T.R. ; MCDANIEL, P. ; REX-FORD, J.: A Survey of BGP Security Issues and Solutions. In: Proceedings of the IEEE 98 (2010), Nr. 1, S. 100–122
- [BMY06] BARBIR, A. ; MURPHY, S. ; YANG, Y.: Generic Threats to Routing Protocols. 2006. – RFC 4593, IETF, Status: Standard, <https://tools.ietf.org/html/rfc4593>
- [CKV11] CAVEDON, L. ; KRUEGEL, C. ; VIGNA, G.: Are BGP Routers Open to Attack? An Experiment. In: Open Research Problems in Network Security Bd. 6555. Springer Berlin Heidelberg, 2011, S. 88–103
- [GHM+07] GILL, V. ; HEASLEY, J. ; MEYER, D. ; SAVOLA, P. ; PI-GNATARO, C.: The Generalized TTL Security Mechanism (GTSM). 2007. – RFC 5082, IETF, Status: Standard, <https://tools.ietf.org/html/rfc5082>
- [Hef98] HEFFERNAN, A.: Protection of BGP Sessions via the TCP MD5 Signature Option. 1998. – RFC 2385, IETF, Status: Standard, <https://tools.ietf.org/html/rfc2385>
- [HPS04] Y. Hu, A. Perrig, M. Sirbu. *SPV: Secure Path Vector Routing for Securing BGP*. SIGCOMM 2004, 2004.
- [LK12] LEPINSKI, M. ; KENT, S.: An Infrastructure to Support Secure Internet Routing. 2012. – RFC 6480, IETF, Status: Proposed Standard, <https://tools.ietf.org/html/rfc6480>
- [LMS03] LYNN, C. ; MIKKELSON, J. ; SEO, K.: Secure BGP (S-BGP). 2003. – IETF, Status: Expired Internet-Draft, <https://tools.ietf.org/html/draft-clynn-s-bgp-protocol-01>
- [LT13] LEPINSKI, M. ; TURNER, S.: An Overview of BGPSEC. 2013. – IETF, Status: Internet-Draft, <https://tools.ietf.org/html/draft-ietf-sidr-bgpsec-overview-04>



Routing Security References (2)

- [Kent] S. Kent. Design and Analysis of the Secure Border Gateway Protocol (S-BGP). Presentation.
- [KMR03] KRUEGEL, C. ; MUTZ, D. ; ROBERTSON, W. ; VALEUR, F.: Topology-based detection of anomalous BGP messages. In: Recent Advances in Intrusion Detection (RAID), 2003, S. 17–35
- [Lynn99] C. Lynn. *Secure Border Gateway Protocol (S-BGP)*. Presentation at the 1999 Network and Distributed Systems Security Symposium (NDSS'99), 1999.
- [NSZ03] NICOL, D. M. ; SMITH, S. W. ; ZHAO, M.: Efficient Security for BGP Route Announcements / Dartmouth College, Computer Science. 2003 (TR2003-440). – Research Report
- [QGR07] QIU, J. ; GAO, L. ; RANJAN, S. ; NUCCI, A.: Detecting bogus BGP route information: Going beyond prefix hijacking. In: SecureComm 2007, 2007, S. 381–390
- [SRS+ 04] SUBRAMANIAN, L. ; ROTH, V. ; STOICA, I. ; SHENKER, S. ; KATZ, R.: Listen and Whisper: Security Mechanisms for BGP. In: Proc. First Symposium on Networked Systems Design and Implementation (NSDI), 2004
- [TMB10] TOUCH, J. ; MANKIN, A. ; BONICA, R.: The TCP Authentication Option. 2010. – RFC 5925, IETF, Status: Standard, <https://tools.ietf.org/html/rfc5925>
- [Whi06a] WHITE, R.: Architecture and Deployment Considerations for Secure Origin BGP (soBGP). 2006. – IETF, Status: Expired Internet-Draft, <https://tools.ietf.org/html/draft-white-sobgp-architecture-02>

