

Security Topics in Networking and
Distributed Systems
CS 590D

Lecture 11: Security of BGP

Department of Computer Sciences
Purdue University

References

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



A Network of Networks

- Internet is a “network of networks”
- Autonomous System (AS): a network, a single administrative domain, can span more organizations
- How do those networks connect
 - Internet Exchange (IX)
 - Network Access Points (NAP)
 - Metropolitan Area Exchange (MAE)

Getting on the Net

- Internet Service Provider (ISP) provide access to Internet
- Several types of ISPs depending on size, transit/peering and multihoming
- **Transit**: carrying packets for other ISP
- **Peering**: two ISPs of about the same size exchange traffic
- **Multihoming**: being connected to multiple networks

Finding the Way Through the Net

- IP address: identifies a computer IPv4 - 32 bits, IPv6 - 128 bits
- Routing protocols: propagate information about routes to reach hosts (IP addresses) or networks (IP prefixes)
- Algorithms:
 - Distance vector protocols
 - Link-state protocols
 - Path vector protocols
- Relative to an AS
 - inter-routing: RIP, IGRP, OSPF,
 - intra-routing: BGP

Link-State Routing

- Each node:
 - Maintains **global view** of the network.
 - Sends periodically the current state of all links (link-state updates or advertisements) to all nodes (via flooding).
 - Notes the change and recompute its routes (use **shortest-path** – Dijkstra algorithm) to destination.
- Less bandwidth-intensive than Distance-Vector, but more complex and more computational and memory intensive.
- Examples: OSPF uses link-state routing.

Distance-Vector Routing

- Each node:
 - Maintains a vector with distances to all of the nodes.
 - Sends periodically its distance-vector to all its neighbors.
 - Updates its distance vector based on the information received from the neighbors (shortest path Bellman-Ford): for each network path, the receiving routers pick the neighbor advertising the lowest cost, then add this entry into its routing table for re-advertisement.
- Examples: RIP uses distance-vector routing.

Path Vectors

- Similar to Distance Vector protocols
- BUT: Routing updates contain an ordered list of the path of traversed “nodes”

BGP

- Path Vector Protocol: Routing updates contain an **ordered list or AS path of traversed autonomous systems and a set of network prefixes belonging to the first AS in the list** (UPDATE messages)
- BGP uses TCP to exchange routing updates
- Each BGP router receives UPDATES from its neighbors and selects one path for each prefix as the “best” and reports that path to its neighbors (before that it has to withdraw the “old” path)
- Selecting “best path”: policies, local preference, shortest AS path, other metrics

Some Statistics....

- BGP routing tables: ~ 125K address prefixes mapping to about 17-18K paths
- BGP routers: ~ 10K BGP
- How many organizations own AS? ~ 2K
- How many organization own prefixes?~ 60K
- Path length for a route:
 - Average AS path length for a route is about 3.7
 - 50% of routes have a length < 4 Ases
 - about 95% have a length < 5

UPDATE Messages (4096 bytes)

- The AS that knows how to reach directly a prefix of IP addresses will be the first AS in the list (in UPDATE messages)
- When receiving an UPDATE message, an AS updates its own routing table then propagates the update about new routes, after he added itself in the list of ASes
- Each AS along the path must be authorized by the preceding AS to advertise the prefixes contained in the UPDATE message
- A route may be withdrawn only by the neighbor AS that advertised it

Threats

- BGP does not use any security mechanisms at all
- Misconfiguration: an ISP advertised addresses they it does not know how to reach; consequence: packets will get dropped.
- Same effect if malicious attacker injects or modifies update packets, withdraws routes.

Solution

- Authentication of source of packets?
- Data integrity?

NEEDED BUT NOT ENOUGH!

...Because ...

- IP addresses are own by organizations
- An AS can span more organizations
- An AS must be authorized by the organization that owns a set of IP addresses to advertise them
- Each AS along the path must be authorized by the preceding AS to advertise the prefixes contained in the UPDATE message

WHAT IS NEEDED?

- A BGP speaker should be able to verify:
 - The owner of each prefix authorized the origin AS to advertise that prefix
 - Each subsequent AS in the path has been authorized by the preceding AS to advertise a route to that prefix
 - Neighbor withdrawing a route is the advertiser for that route

S-BGP

- Secure the BGP traffic by using IPSEC, ESP mode with no encryption (provides authentication and integrity)
- Uses Public Key Infrastructure to provide an authorization for the mapping between address space and AS
- Uses digital signatures to bind authorization information to UPDATE messages

PKI Support

- Goal: bind the relationship among Internet registries, ISPs and subscribers
- X.509 certificates are issued to ISPs and subscribers to identify “owners” of ASes and prefixes

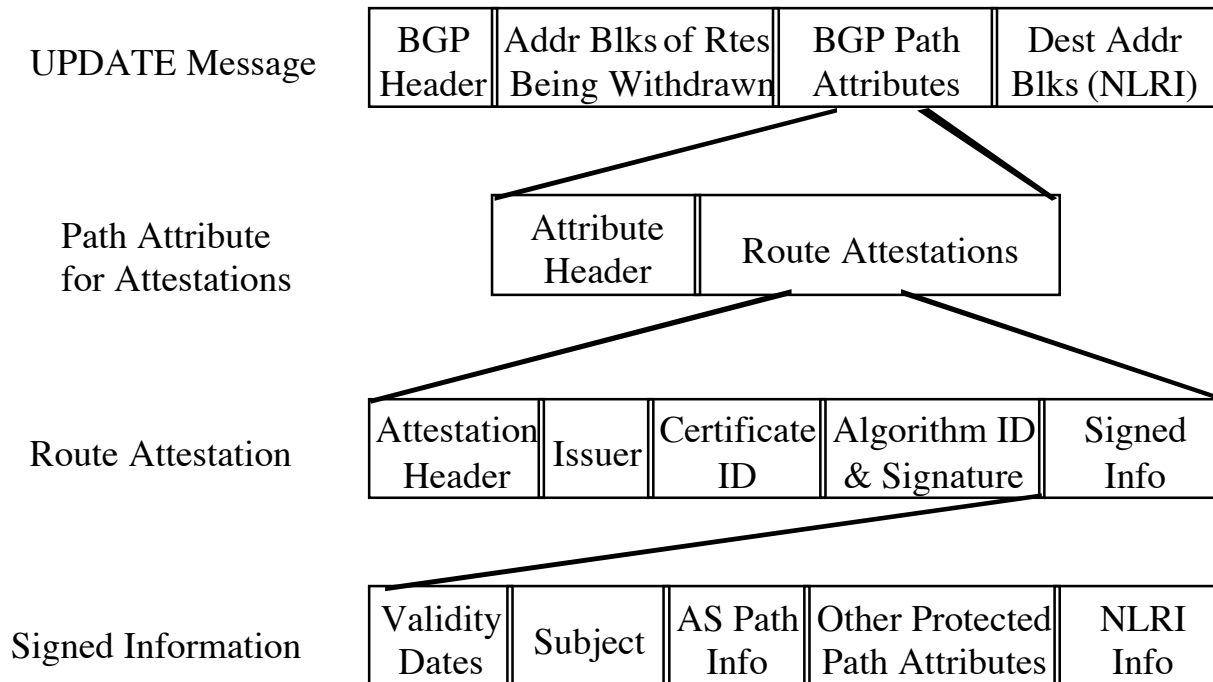
Attestations

- **Address Attestation:** issued by the “owner” of one or more prefixes, to identify the origin AS authorized to advertise the prefixes
- **Route Attestation:** is issued by a S-BGP speaker, to authorize neighbor ASes to use the route in the UPDATE message will issue

PKI Distribution

- How to distribute the certificates?
- How about the CRLs?
- On-line vs off out-of-band
- What about attestations?
 - Address attestations
 - Route attestations

S-BGP UPDATE Message



Deployment

- Processing overhead
- Bandwidth
- Storage/memory
- Backward compatibility
- Adoption
- PKI deployment