Secure Border Gateway Protocol (S-BGP): Real World Performance & Deployment Issues

Stephen Kent, Charles Lynn, Joanne Mikkelson, and Karen Seo

BBN Technologies

A Part of GTE

Outline

- BGP Model
- BGP security concerns & requirements
- S-BGP design
- S-BGP performance & scaling
- Conclusions

Basic BGP Model

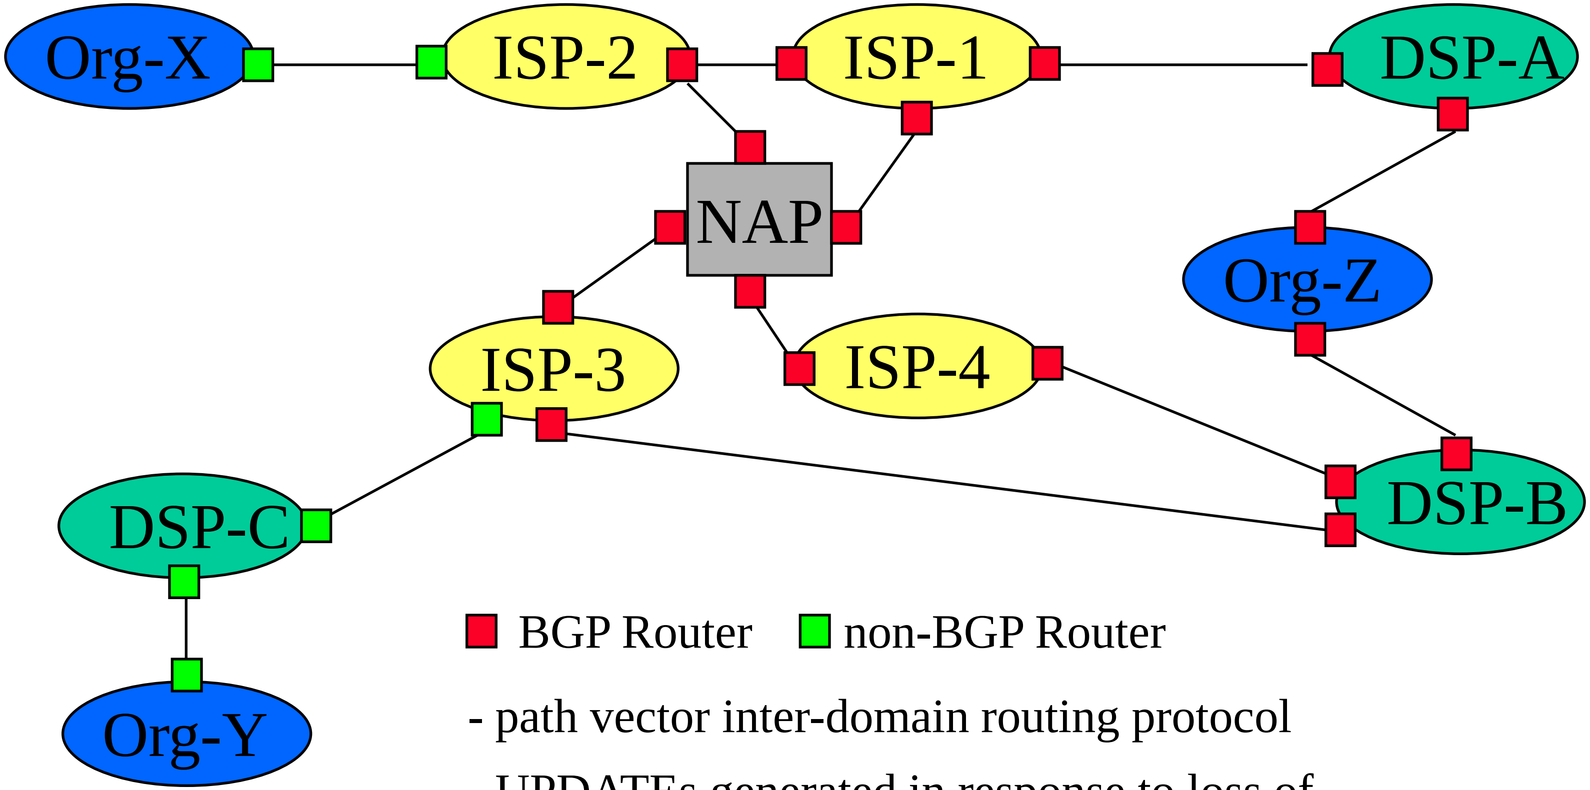

BGP Router   non-BGP Router

- path vector inter-domain routing protocol
- UPDATEs generated in response to loss of connectivity or receipt of an UPDATE from a peer router, that results in a LOCRIB change

The BGP Security Problem

- BGP is the critical infrastructure for Internet, inter-domain routing
- Benign configuration errors have wreaked havoc for portions of the Internet address space
- The current system is highly vulnerable to human errors, as well as a wide range of attacks
- At best, BGP uses point-to-point keyed MAC, with no automated key management
- Most published BGP security proposals have been pedagogic, not detailed, not deployable
- Solutions must take into account Internet topology, size, update rates, ...

Attack Model

- **BGP can be attacked in various ways**
  - active or passive wiretapping of communications links between routers
  - tampering with BGP speaker software
  - tampering with router management data en route
  - tampering with router management workstations/servers

    (the last three can result in Byzantine failures)
- **Addition of the proposed countermeasures introduces a new concern**
  - compromise of secret/private keying material in the routers or in the management infrastructure

BGP Security Requirements

- Verification of address space "ownership"
- Authentication of Autonomous Systems (AS)
- Router authentication and authorization (relative to an AS)
- Route and address advertisement authorization
- Route withdrawal authorization
- Integrity and authenticity of all BGP traffic on the wire
- Timeliness of BGP traffic

S-BGP Design Overview

- IPsec: authenticity and integrity of peer-to-peer communication, 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 UPDATEs based on a new path attribute, using certificates and attestations
- Distribution of countermeasure data: certificates, CRLs, attestations

S-BGP Residual 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)

Internet Address Space Ownership

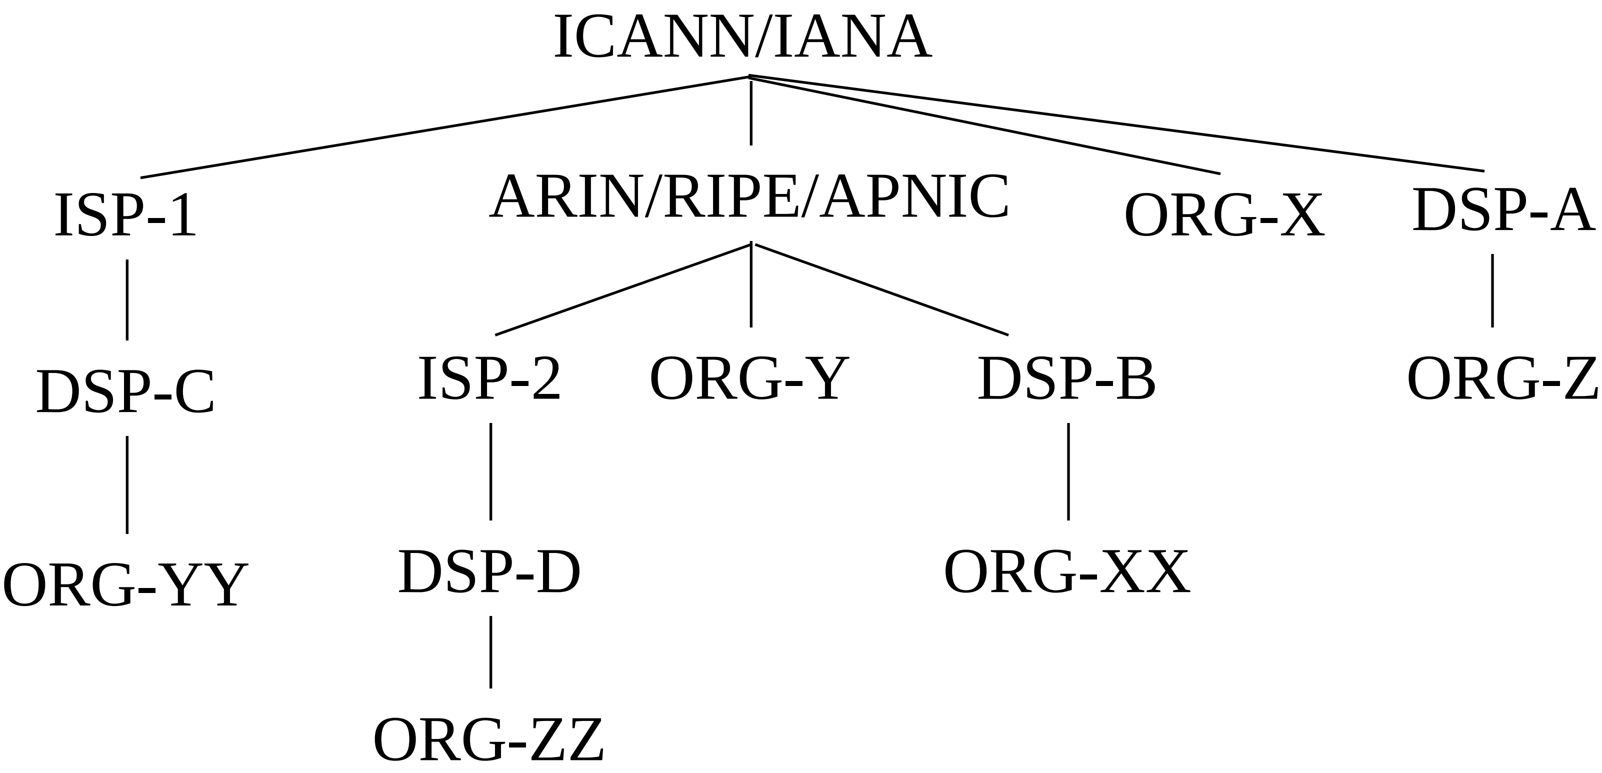

Simplified PKI for Address Blocks

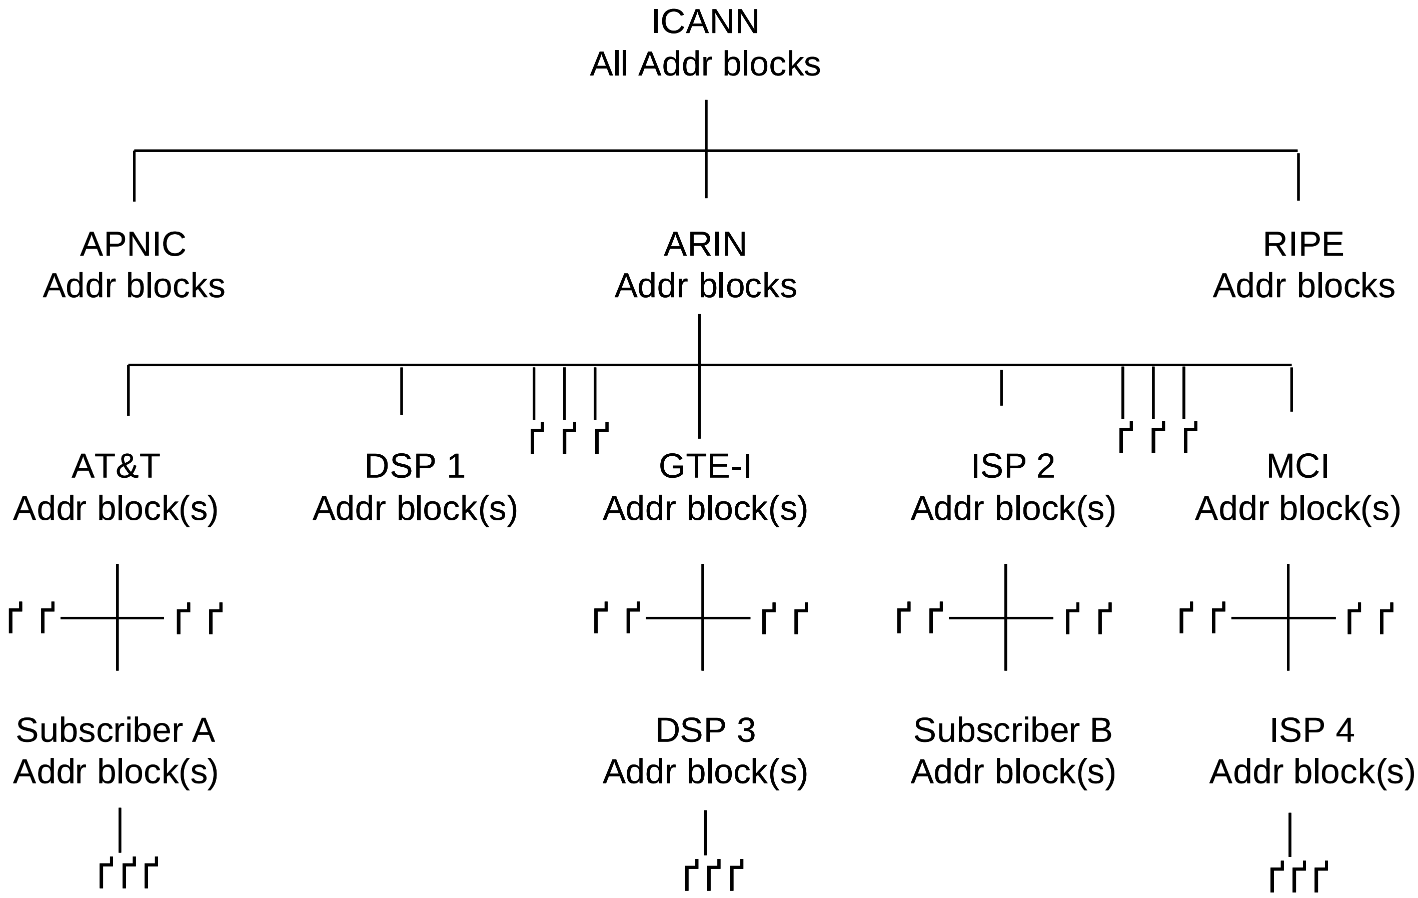

- Only networks that execute BGP need certificates
- All ISPs are BGP users, but only about ~10% of DSPs, maybe 5% of subscribers, are BGP users

PKI for Speaker ID & AS Assignment

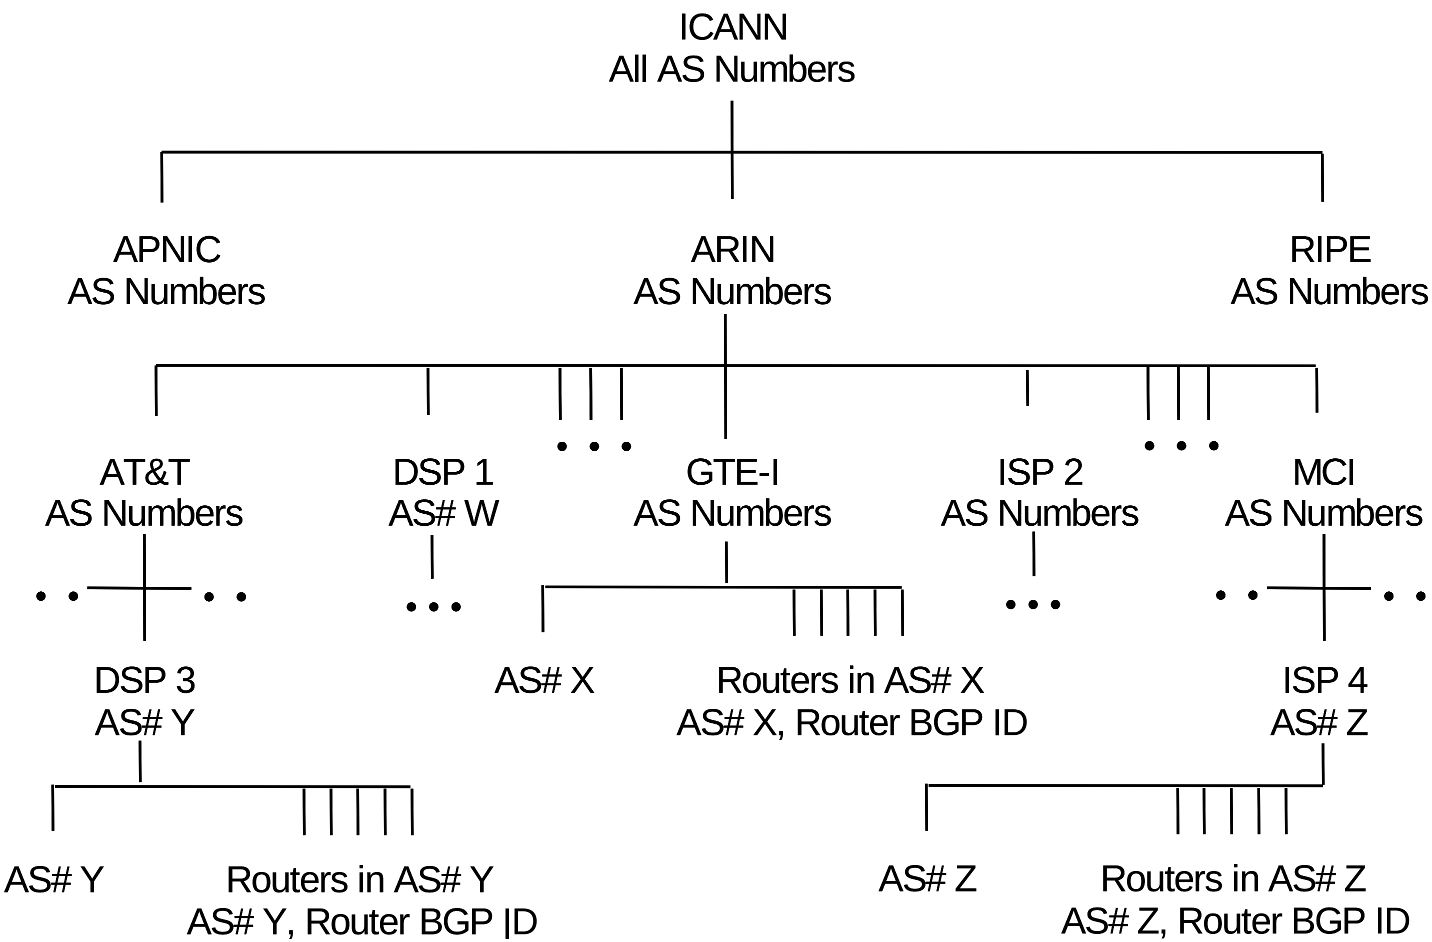

Securing UPDATE messages

- 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

An UPDATE with Attestations

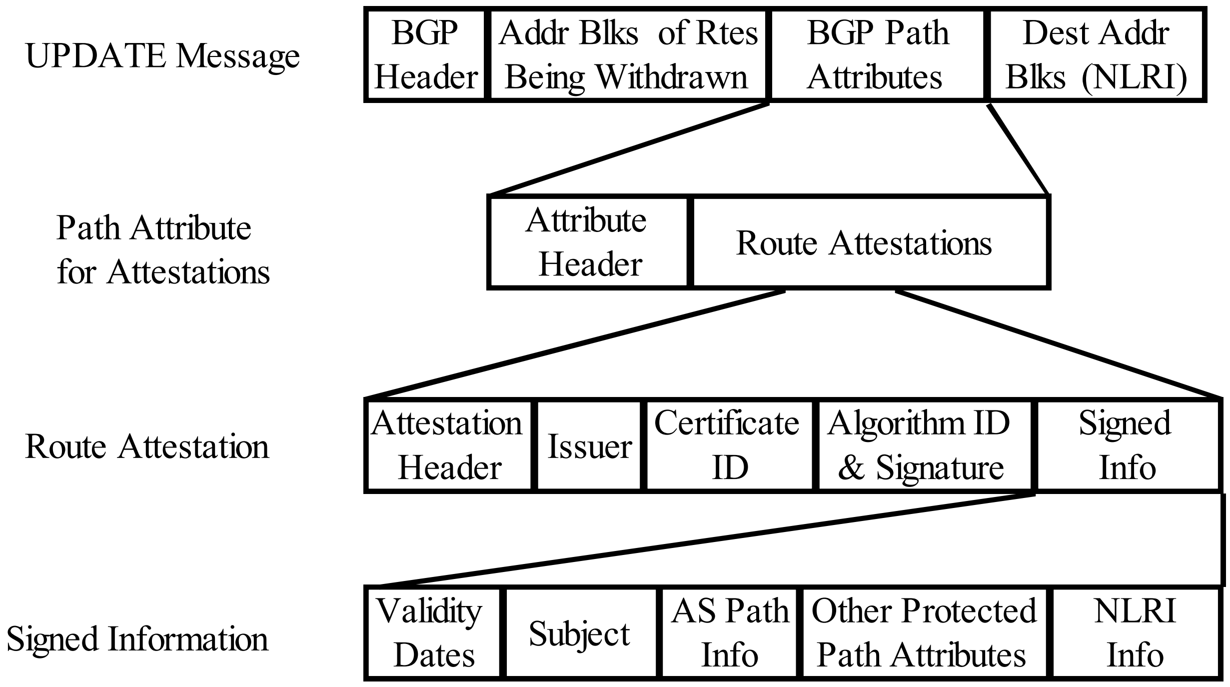

Simplified Attribute Format

BGP Hdr: Withdrawn NLRI, Path Attributes, Dest. NLRI

RA: Issuer, Cert ID, Validity, Subject, Path, NLRI, SIG

RA: Issuer, Cert ID, Validity, Subject, Path, NLRI, SIG

RA: Issuer, Cert ID, Validity, Subject, Path, NLRI, SIG

AA: Owning Org, NLRI, first Hop AS, SIG (usually omitted)

Distributing Certificates, CRLs, & AAs

- **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 NAPs for direct (non-routed) access
  - download options:
    - whole certificate/AA/CRL databases
    - queries for specific certificates/AAs/CRLs
- **To minimize processing & storage overhead, NOCs should validate certificates & AAs, and send processed extracts to routers**

Distributing Route Attestations

- Distributed with BGP UPDATEs as path attributes
- RAs have implicit encoding option to reduce size, avoid exceeding UPDATE size limit (4096b)
- Cache with associated routes in ADJ-RIBs to reduce validation overhead
- Expiration date present, but no revocation mechanism chosen yet

BGP Statistics

- ~ 1,800 organizations own AS numbers
- ~ 44,000 own address prefixes (NLRI)
- ~ 7,500 BGP speakers
- ~ 75,000 routes in an ISP BGP database
- Few AS sets (~100), little address aggregation
- Average path length (NAP perspective) is 2.6 hops; 50% of routes ≤ 2 hops, 96% ≤4 hops
- ~ 43,000 UPDATEs received each day at a BGP speaker at a NAP (30 peers)

S-BGP Storage Statistics

- ~ 58,000 certificates in database (~550b each)
- Certificate & CRL database ~35Mb
- Address attestation database ~4 Mbytes
- Extracted certificate & AA database (with data structure overhead in GateD) ~ 42Mb
- Route attestations occupy ~16 Mb per ADJ-RIB: about 64 Mb (4 peers) to 480 Mb (at NAP)
- ADJ-RIB caching for received UPDATEs increases storage requirements by about 50%, and yields about 58% validation savings

Route Attestation Overhead

- **Transmission**
  - RAs add ~450 bytes to a typical (3.6 ASes in path) UPDATE of 63 bytes, 700% overhead!
  - But UPDATEs represent a very small portion of all traffic, so steady state bandwidth for RA transmission is only ~ 1.4Kb/s
- **Processing**
  - Average of 3.6 signature validations per received UPDATE and 1 generation per emitted UPDATE
  - Peak rates ~ 18/s validation and ~5/s generation w/o caching (peak estimated as ten times average)
  - UPDATE caching reduces validation rate by ~50%
  - Start up transient would overwhelm a speaker, thus some form of NV storage or heuristic is required

Conclusions

- **The transmission and processing costs of S-BGP are not significant**
- **The proposed distribution mechanisms for certificates, CRLs, and AAs is viable**
- **Storage overhead exceeds the capacity of existing routers, but adding adequate storage is feasible, especially for ISP BGP speakers**
- **Testing and deployment issues**
  - Cisco handling of optional, transitive path attributes
  - Intra-domain distribution of S-BGP attribute
- **But deployment poses a chicken and egg problem!**