Securing the Border Gateway Protocol Using S-BGP

APNIC Open Policy Meeting
Routing SIG
Kyoto, Japan

Dr. Stephen Kent

Chief Scientist - Information Security

kent@bbn.com

BBN TECHNOLOGIES

Background

- I assume this audience knows
  - How BGP works
  - Why security for BGP is a concern
  - What are critical BGP security requirements
  - Why “trust” is not the preferred basis for determining claims about prefix holders, origination, and routes
  - The need for incremental deployment capabilities for proposed solutions
  - …
- This presentation will focus on just one proposed solution: S-BGP

Secure BGP (S-BGP)

- S-BGP is an architectural solution to the BGP security problems described earlier
- S-BGP represents an extension of BGP
  - It uses a standard BGP facility to carry additional data about paths in UPDATE messages
  - It adds an additional set of checks to the BGP route selection algorithm
- S-BGP avoids the pitfalls of transitive trust that are common in today's routing infrastructure
- S-BGP mechanisms exhibit the same dynamics as BGP, and they scale commensurately with BGP

S-BGP Design Overview

- S-BGP makes use of:
  - **IPsec** to secure point-to-point communication of BGP control traffic
  - **Public Key Infrastructure** to provide an authorization framework representing prefix holders and owners of AS #'s
  - **Attestations** (digitally-signed data) to represent authorization information
- S-BGP requires routers to:
  - **Generate** an attestation when generating an UPDATE for another S-BGP router
  - **Validate** attestations associated with each UPDATE received from another S-BGP router

IPsec for S-BGP

- S-BGP uses IPsec to protect all BGP traffic between neighbor routers
- As used here, IPsec provides cryptographically enforced data authentication, data integrity, and anti-replay features
- IPsec represents a significant improvement over the MD5 TCP checksum option used in some contexts today
  - Automated key management
  - More comprehensive security guarantees
  - Better, standards-based cryptographic protection

Prefix Allocation

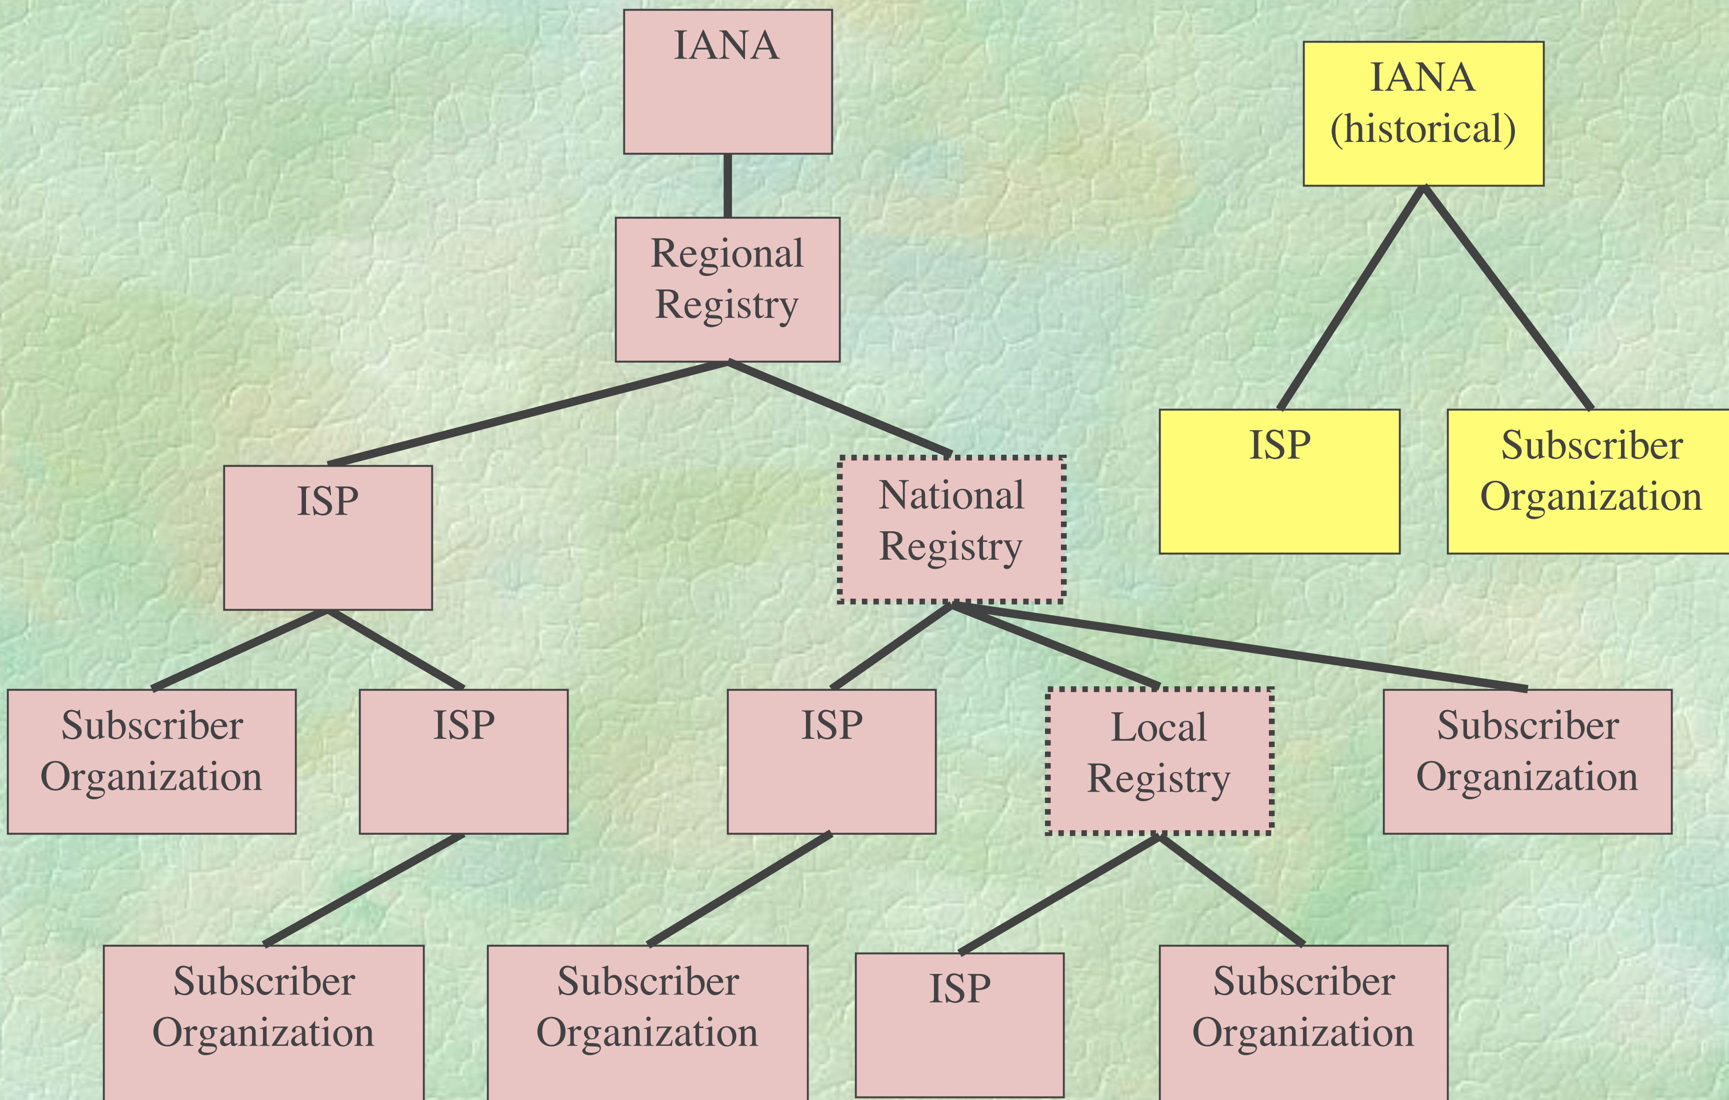

AS # Allocation Hierarchy

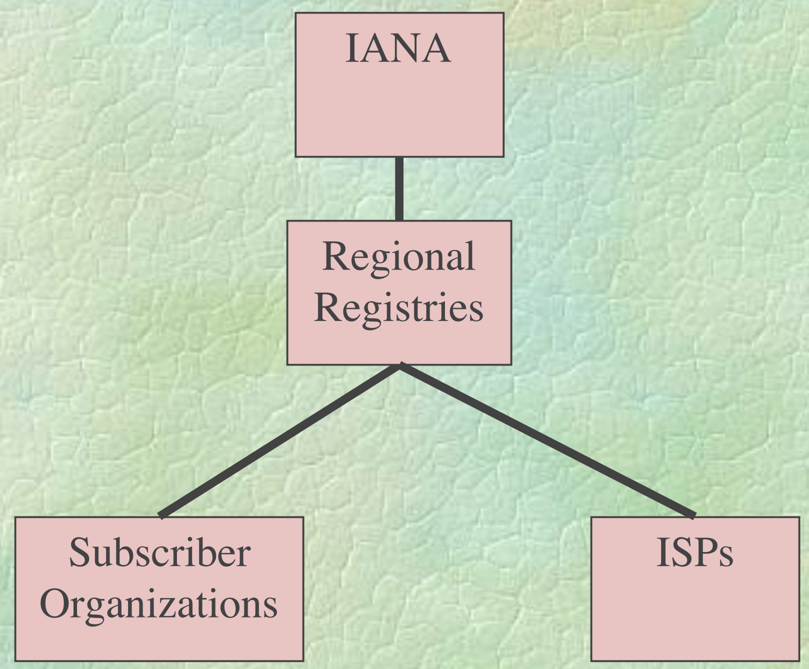

A PKI for S-BGP

- Public Key (X.509) certificates are issued to ISPs and subscribers to identify AS # owners and prefix holders, using RFC 3779 syntax
- Prefixes and public keys in certificates are used to verify authorization of address attestations
- Address attestations, AS #'s and public keys from certificates are used as inputs to verification of UPDATE messages
- The PKI does NOT rely on any new organizations that require trust; it just makes explicit and codifies the relationships among regional, national and local registries, ISPs, and subscribers

S-BGP PKI Characteristics

- S-BGP certificates do not identify ISPs per se
- Most of these certificates bind AS #'s and prefixes to public keys, not to meaningful IDs (avoids name problems re mergers, bankruptcy, …)
- Each RIR (NIR/LIR) acts as a CA to issues certificates that allocate prefixes and AS #'s
- Each ISP acts as a CA to issue certificates to each entity to which it assigns prefixes, but only if the entity executes S-BGP
- ISPs also issue certificates to their S-BGP routers, and operations personnel who interact with the S-BGP repositories

Two Types of Attestations

- An **Address Attestation (AA)** is issued by the "owner" of one or more prefixes (a subscriber or an ISP), to identify the first (origin) AS authorized to advertise the prefixes
- A **Route Attestation (RA)** is issued by a router on behalf of an AS (ISP), to authorize neighbor ASes to use the route in the UPDATE containing the RA
- These data structures share the same basic format

Simplified Attestation Formats

Attestation Type	Certificate Issuer ID	Algorithm ID & Sig Value	Signed Info

Route Attestation

(Prefix$_1$, … Prefix$_n$)
AS$_n$, AS$_{n-1}$, … As$_2$, Origin AS

Address Attestation

(Prefix$_1$, … Prefix$_n$)
Origin AS

Processing an S-BGP UPDATE

- When an S-BGP router generates an UPDATE for a recipient neighbor that implements S-BGP, it generates a new RA that encompasses the path and prefixes plus the AS # of the neighbor AS
- When an S-BGP router receives an UPDATE from an S-BGP neighbor, it:
  - Verifies that its AS # is in the first RA
  - Validates the signature on each RA in the UPDATE, verifying that the signer represents the AS # in the path
  - Checks the corresponding AA to verify that the origin AS was authorized to advertise the prefix by the prefix holder

Housekeeping for S-BGP

- Every S-BGP router needs access to all the certificates, CRLs, and address attestations so that it can verify any RA
- These data items don't belong in UPDATE messages
- S-BGP uses replicated, loosely synchronized repositories to make this data available to ISPs and organizations
- The repository data is downloaded by ISP/organization Network Operation Centers (NOCs) for processing
  - Each NOC validates retrieved certificates, CRLs, & AAs, then downloads an extracted file with the necessary data to routers
  - Avoids need for routers to perform this computationally intensive processing
  - Permits a NOC to override problems that might arise in distributing certificates and AAs, but without affecting other ISPs

S-BGP PKI Repositories

- ISPs & organizations upload their own new data, download full database, on a daily basis
- Repositories use the PKI to enforce access controls to counter DoS attacks
  - Access granted only to S-BGP users and other repositories
  - An ISP or organization is constrained to prevent overwriting data of another ISP or organization
- Major ISPs could operate repositories for themselves & their subscribers
- Internet exchanges or registries could operate repositories for other ISPs & subscribers
- Note that repositories need not be highly available, e.g., they are NOT accessed in real time by routers

S-BGP System Interaction Example

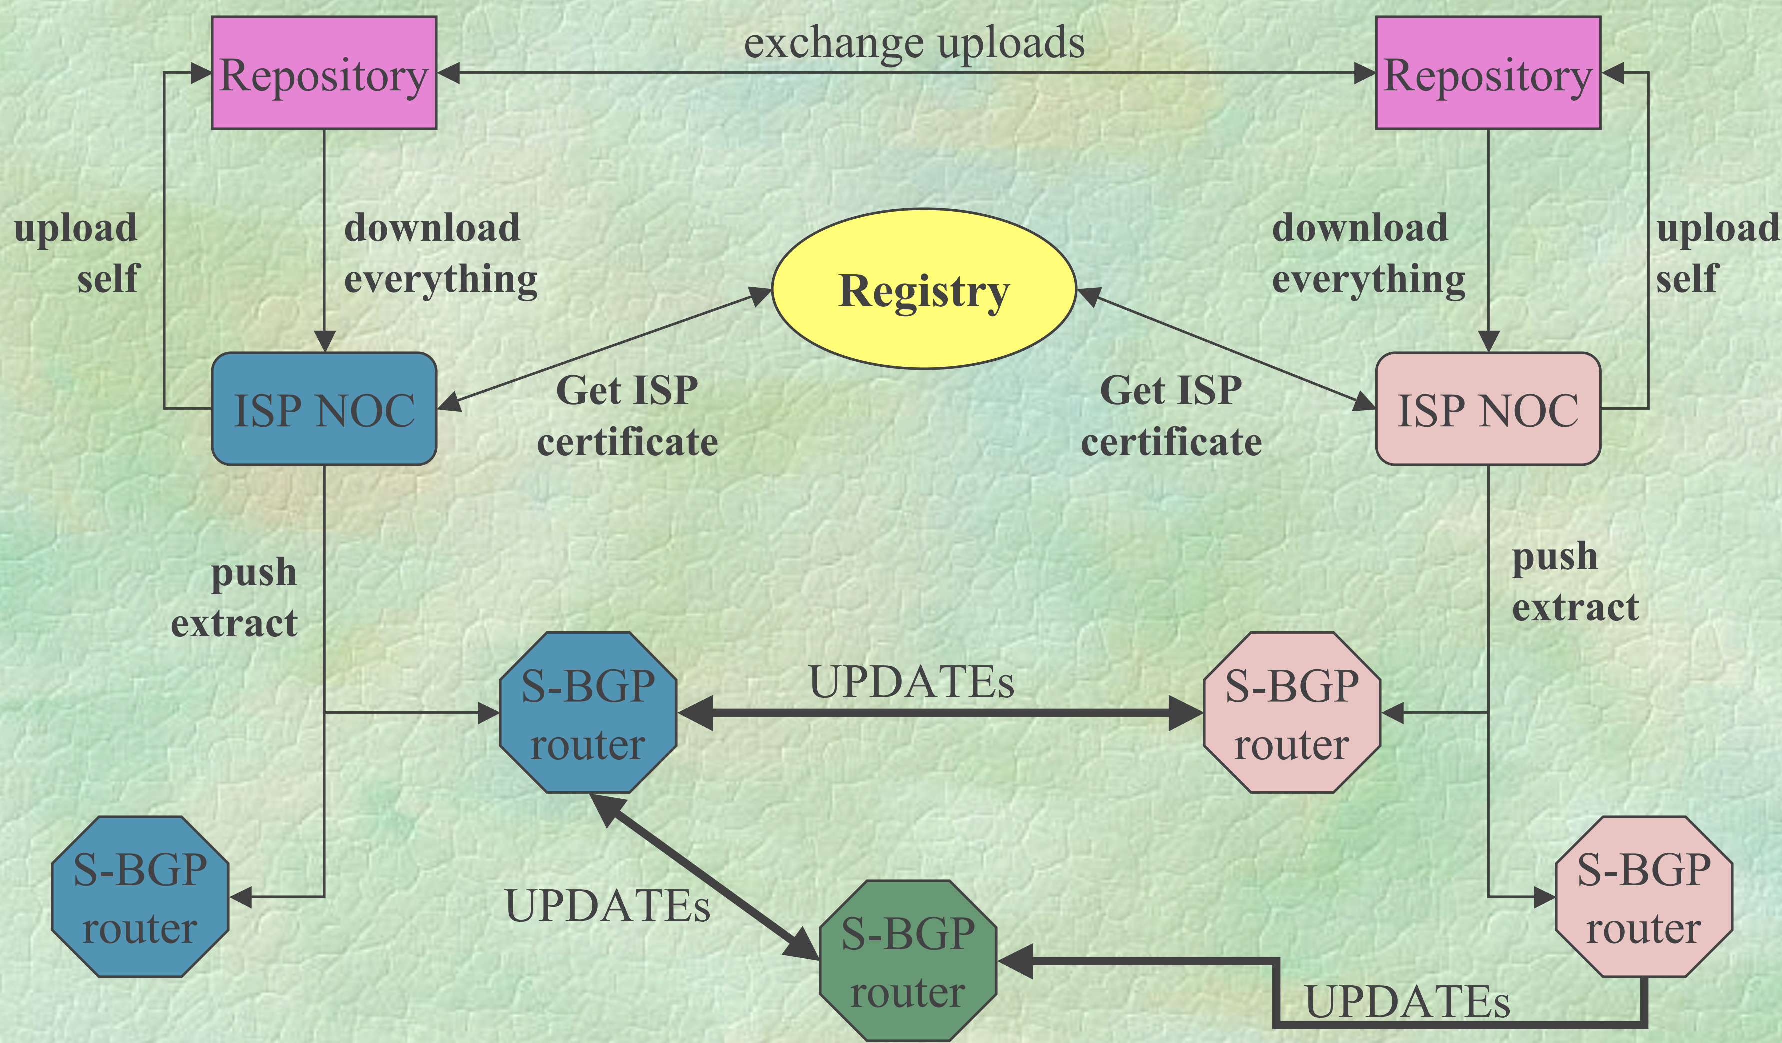

Residual Vulnerabilities

- S-BGP cannot ensure that a router withdraws a route when the only path (known to the router) for the route is withdrawn by a neighbor
- S-BGP does not ensure timeliness of UDATEs, except to the extent that RAs time out
  - This means that a router could retransmit an UPDATE after it withdrew a route, without having been authorized to re-advertise the route

What Exists Today?

- S-BGP code
  - Implemented on MRT code base
  - Includes basic policy controls for incremental deployment
- NOC Tools
  - Mini-registration authority for certificate requests
  - AA generation
  - Repository upload/download tools
  - Certificate, CRL & AA validation & extract file generation
- Repository
  - PKI-based access controls for access & uploads
  - Primitive management capabilities, no synchronization
- CA for S-BGP PKI
  - A high assurance CA on an SELinux base processes X.509 certificate requests with S-BGP private extensions

Summary

- S-BGP addresses the architectural security problems of BGP and supports verification of route changes in realtime
- The impact on daily Registry & ISP operations is minimal, although training will be needed
- The S-BGP PKI leverages existing authorization relationships and creates no new ones
- Routers may require hardware upgrades, for crypto, even if not for memory
- The security model embodies the principle of least privilege, providing containment in the face of compromise

Questions?

http://www.ir.bbn.com/projects/s-bgp

Additional Slides

Deploying S-BGP

- Router software must implement S-BGP
- Router hardware must have appropriate storage & digital signature processing capabilities
- Registries must assume CA responsibilities for address prefixes and AS # allocation
- ISPs and subscribers that execute BGP must upgrade routers, must act as CAs, and must interact with repositories to exchange PKI & AA data

Router Memory & Performance

- Routers need enough memory to hold route attestations in Adj-RIBs and Loc-RIB, plus storage for address attestation and processed certificates
- Signature generation and validation pose a modest burden in a steady state context, well within the capabilities of CPUs used for router management
- But, to accommodate surge volume during attacks, and to better protect router keys, use of a crypto accelerator is preferable
- RA validation heuristics, e.g., deferred UPDATE validation, can reduce the crypto processing burden

Deferred UPDATE Validation

- If validating every UPDATE poses too great a processing burden on a router, it can defer processing most UPDATEs
- Only if an UPDATE would result in a new Loc-RIB entry is it necessary to validate it
- Thus, a router with many peers, one that would receive the most UPDATEs, can defer validation for the vast majority of these messages
- Also, if a router filters inbound UPDATEs using local policy info, it may ignore many UPDATEs anyway!
- If validation is deferred, the router should at least check to verify that the RAs were current when the UPDATE was received