Securing the Border Gateway Protocol

Dr. Stephen Kent

Chief Scientist - Information Security

BBN TECHNOLOGIES

A Verizon Company

Outline

- BGP Overview
- BGP Security
- S-BGP Architecture
- Deployment Issues for S-BGP
- Alternative Approaches to BGP Security
- S-BGP Software Status
- Questions

Why Are We Here?

- BGP provides the critical routing infrastructure for the Internet, the basis for all inter-ISP routing
- The current system is highly vulnerable to human errors, and a wide range of malicious attacks
- Configuration errors are commonplace, perpetual
- BGP has been attacked; more attacks seem likely
- No comprehensive solutions to BGP security problems have been developed by vendors or ISPs
- Solutions will require buy-in by vendors, ISPs, & subscribers, and will likely to take years to deploy

External vs. Internal use of BGP

Routes acquired externally (from other ASes via eBGP) have to be propagated to other border routers in an AS. This is effected using iBGP, either directly or via a route server

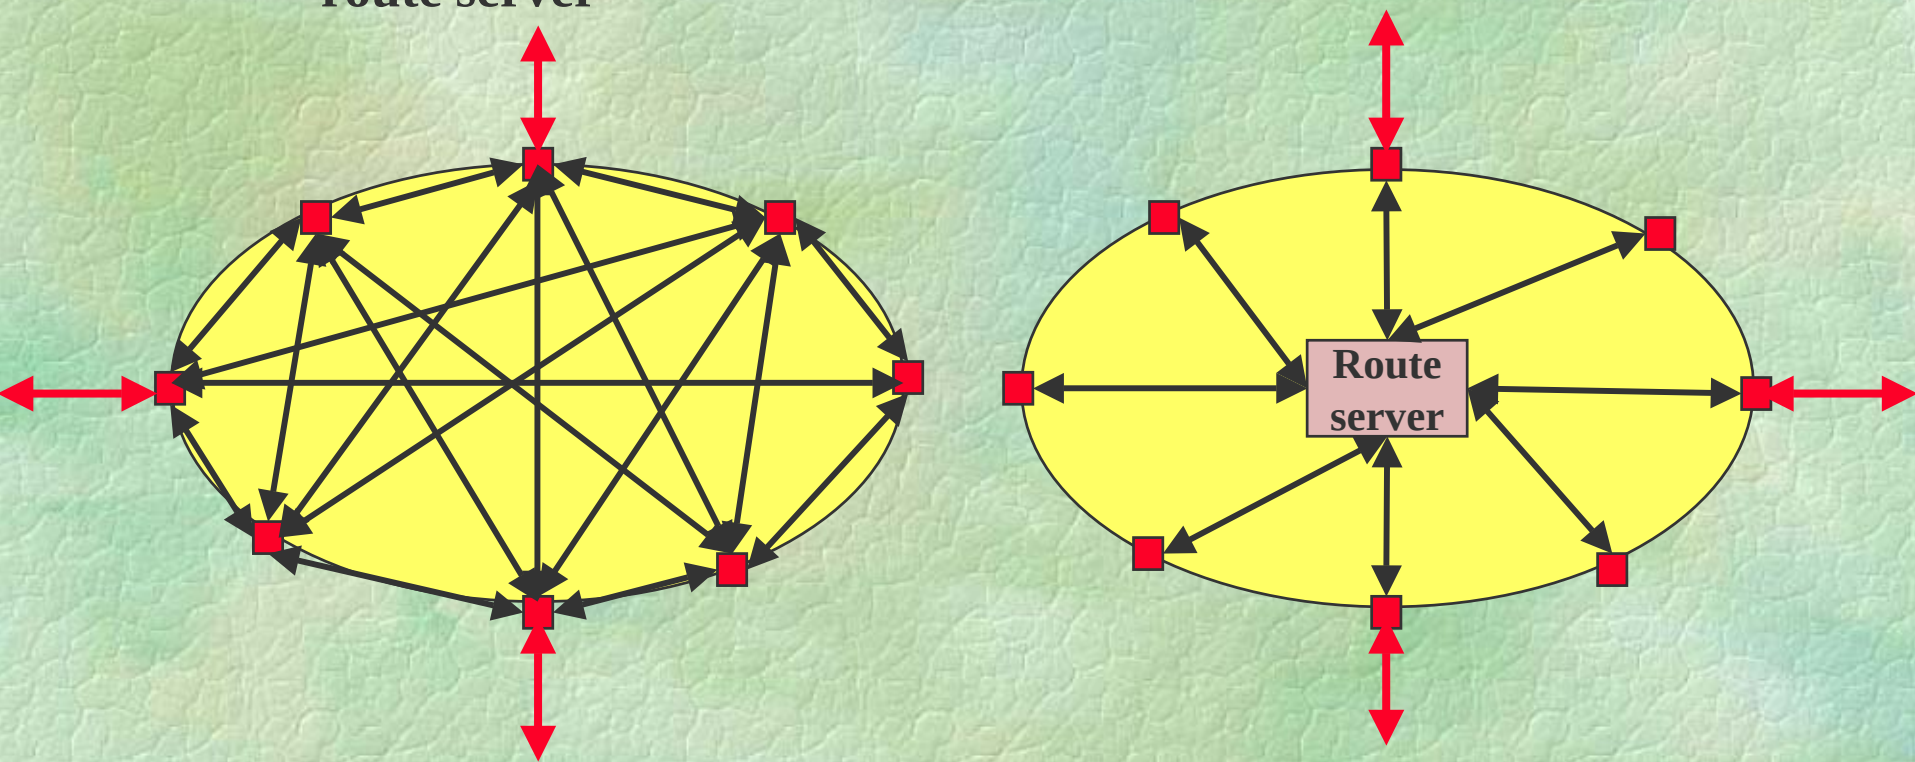

A Simplified UPDATE Message

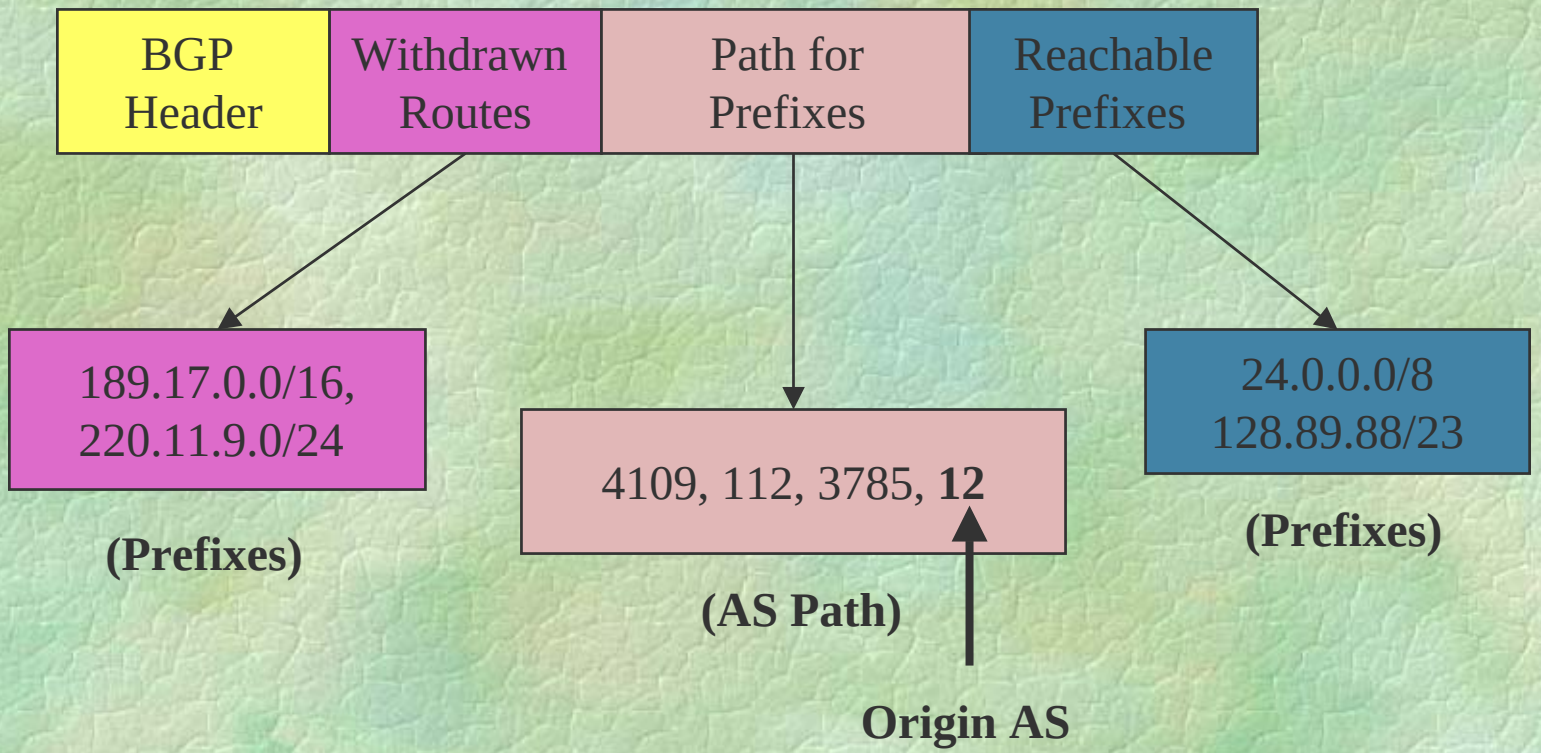

Processing an UPDATE

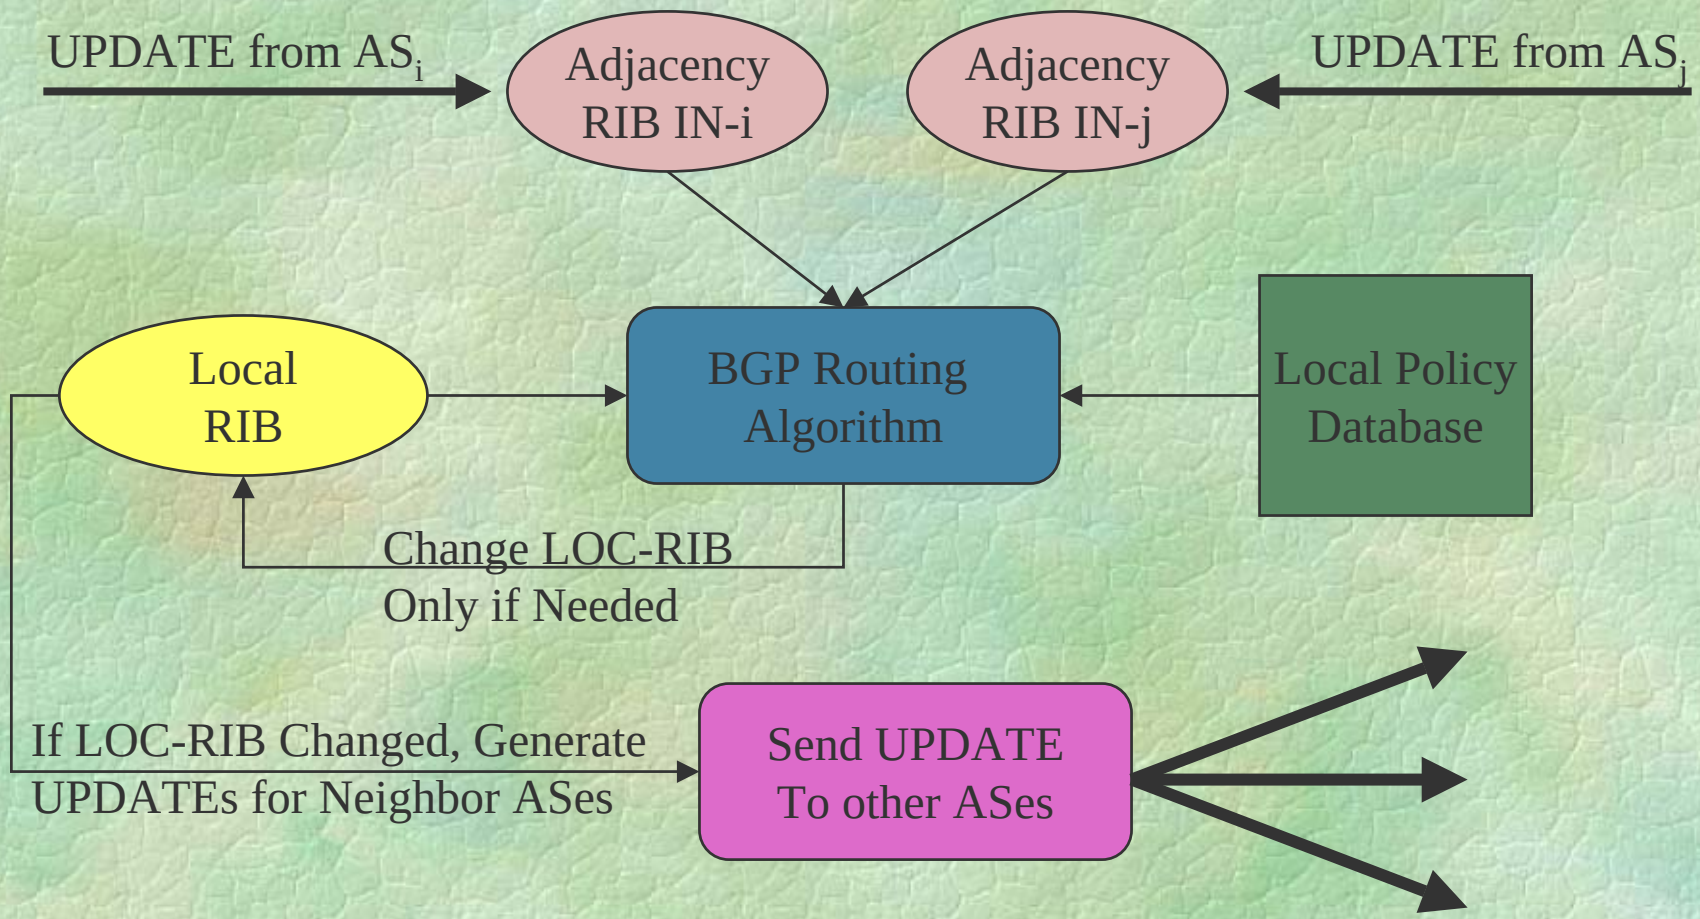

Underlying Assumption re UPDATEs

- Each AS along the path is assumed to have been authorized by the preceding AS to advertise the prefixes contained in the UPDATE message
- The first AS in the path is assumed to have been authorized to advertise the prefixes by the "owner" of the prefixes
- A route may be withdrawn only by the neighbor AS that advertised it (ADJ-RIB-IN locality)
- **If any of these assumptions are violated, BGP becomes vulnerable to many forms of attack, with a variety of adverse consequences**

Some BGP Subtleties

- The “best” route is very much influenced by local policies, which represent concerns such as business arrangements between ISPs and internal traffic engineering decisions internal to an ISP
- An AS may report different routes to different neighbors because of local policies, asymmetric routes are common
- Not all connections between ASes are visible to the Internet at large, e.g., private peering links
- Withdrawal of a route for a prefix by one AS need not result in a neighbor withdrawing the route for that prefix (since the neighbor may have an alternative route available from another source)

BGP Security

Adversary Goals for BGP Attacks

- Degrade service (locally or globally) by effecting a DoS attack against a router's implementation of BGP
- Reroute subscriber traffic (via a path it otherwise would not take) to subject that traffic to passive or active wiretapping
  - Examine/copy subscriber traffic and pass it on to the destination
  - Modify subscriber traffic and pass it on …
  - Delete selected subscriber traffic
  - Masquerade as subscribers, consuming traffic directed to them and responding on their behalf

BGP Security Problems

- The BGP architecture makes it highly vulnerable to human errors and malicious attacks against
  - Links between routers
  - The routers themselves
  - Management stations that control routers
- Most router implementations of BGP are susceptible to various DoS attacks that can crash the router or severely degrade performance
- Many ISPs rely on local policy filters to protect them against configuration errors & some forms of attacks, but creating and maintaining these filters is difficult, time consuming, and error prone

Is BGP Under Attack?

- DARPA-sponsored research has discovered that configuration errors affect about 1% of all routing table entries at any time, but these seem to be the result of configuration errors, not attacks
- Attack tools for BGP have been developed and demonstrated at hacker conferences
- Attacks against ISP routers do occur, and these attacks permit BGP attacks to be launched from compromised routers
- BGP-based attacks have been used by hackers as part of an effort to masquerade as root DNS servers

BGP Security Solution Requirements

- Security architectures for BGP should not rely on "trust" among ISPs or subscribes
  - On a global scale, some ISPs will never be trusted
  - People, even trusted people, make mistakes, and trusted people do "go bad"
  - Transitive trust in people or organizations causes mistakes to propagate (domino effect)
- Elements of security solutions must exhibit the same dynamics as the parts of BGP they protect
- The memory & processing requirements of a solution should scale consistent with BGP scaling

Principle of Least Privilege

- Each element of a system should be granted permissions consistent with the functions that the element performs, but no more
- This principle is a cornerstone of information assurance strategy
- In BGP, this translates into the notion that a security failure (or benign error) by an ISP or subscriber should not propagate to other ISPs
- Any security strategy for BGP should incorporate this “fire break” approach to containing (Byzantine) security failures or errors

Scope & Dynamics of BGP Data

	LOCAL	GLOBAL
SLOW	Install new link Operation staff changes	New prefixes or AS # allocation/assignment
FAST	Add/delete BGP router	Route change

Architecture & Implementation

- The quality of BGP router implementations must be improved to reduce the likelihood that an individual router can be crashed or that the BGP software can be subverted as a result of router compromise
- However, improvements in BGP implementations will not result in a secure routing system, absent architectural changes that address fundamental BGP security problems
- **Thus, both architectural and implementation security improvements are required to make BGP secure & robust**

BGP and Router DoS Issues

- Routers generally are unable to process management data (e.g., BGP, SNMP, etc.) at aggregate line rates, normally not a problem
- This translates into a DoS vulnerability for the processor that deals with management traffic
- This is an implementation vulnerability, but it may merit an architectural solution, given its severity and pervasiveness (not just a BGP issue)
- With regard to BGP traffic, its point-to-point relay nature may permit various solution approaches, but other management data, which is end-to-end, requires more sophisticated solutions

The Basic BGP Security Requirement

- **For every UPDATE it receives, a BGP router should be able to verify that the "owner" of each prefix authorized the first (origin) AS to advertise the prefix and that each subsequent AS in the path has been authorized by the preceding AS to advertise a route to the prefix**
- This requirement, if achieved, allows a BGP router to detect and reject unauthorized routes, irrespective of what sort of attack resulted in the bad routes
- Conversely, if a security approach fails to achieve this requirement, a BGP router will be vulnerable to attacks that result in misrouting of traffic in some fashion

Derived BGP Security Requirements

- Verification of address space “ownership”
- Verification of Autonomous System (AS) “ownership”
- Binding a BGP router to the AS(es) it represents
- Verification of UPDATEs by routers
- Route withdrawal authorization
- Integrity and authenticity of all BGP traffic on the wire (as a counter to active wiretapping attacks that could result in DoS)
- Timeliness of UPDATE propagation*

S-BGP Architecture

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 security mechanisms exhibit the same dynamics as BGP, and scale commensurate 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 address space and AS # "ownership"
  - **Attestations** (digitally-signed data) to bind authorization information to UPDATE messages
- 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 also could be used to filter all management traffic addressed to a router, thus improving security for other management protocols (but its use may pose new DoS concerns)
- IPsec represents an improvement over the MD5 TCP checksum option used in some contexts today

A PKI for S-BGP

- Public Key (X.509) certificates are issued to ISPs and subscribers to identify “owners” of AS #’s and prefixes
- 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 registries, ISPs, and subscribers

Address Allocation/Assignment

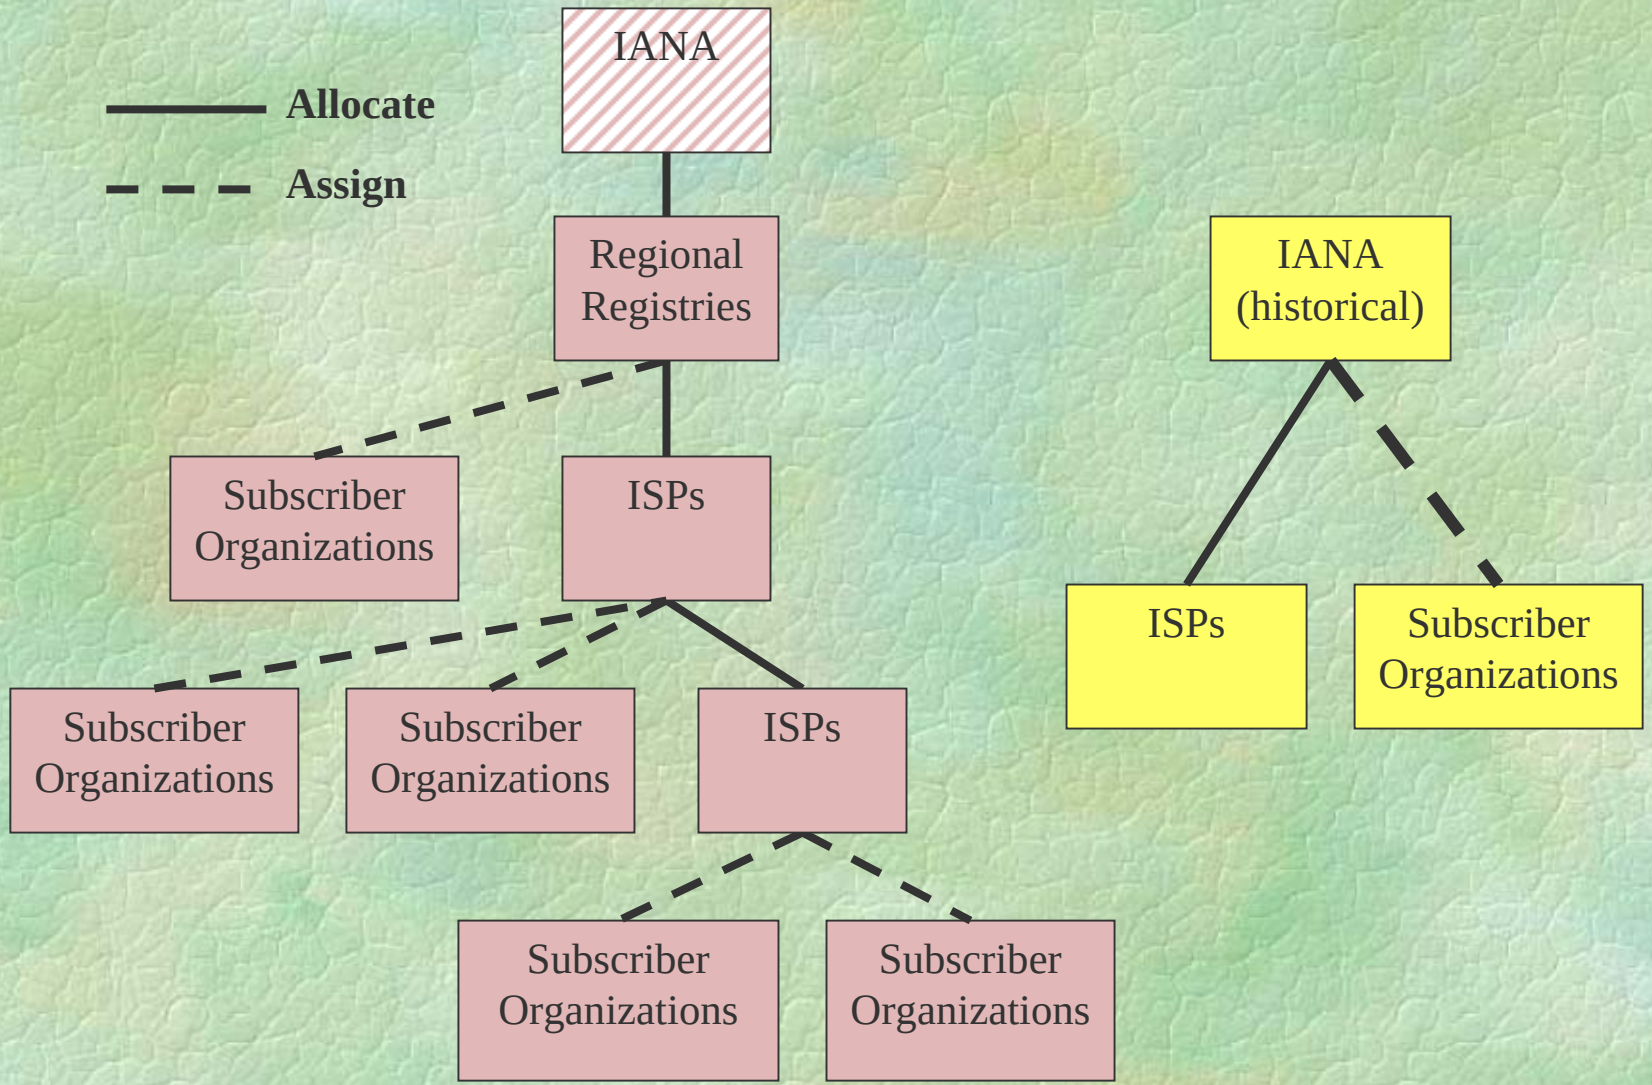

AS # Allocation/Assignment Hierarchy

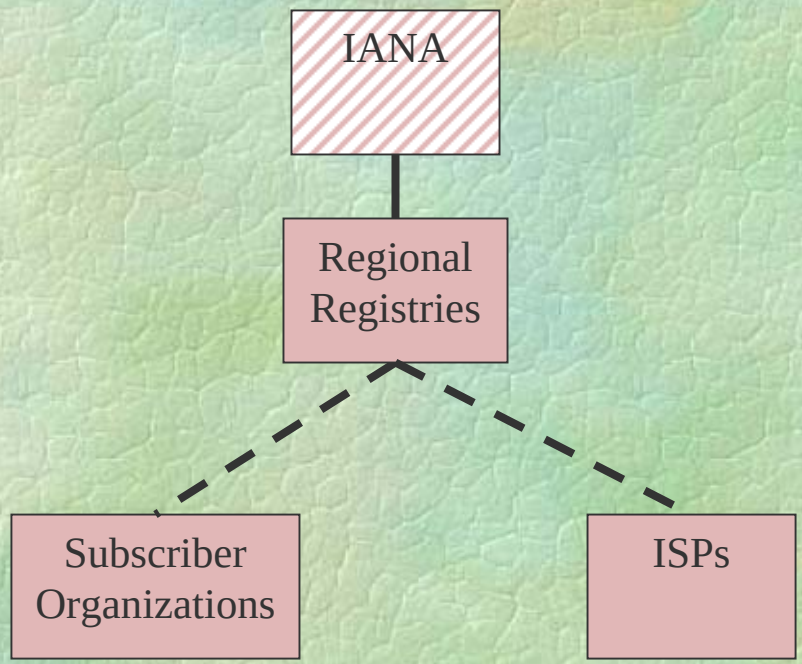

S-BGP PKI: Who's the Root?

- We could use IANA as the root
  - Exactly matches prefix and AS # allocation system
  - Infrequent operations, so not a significant operational burden
- OR, we could create a virtual root by having each RIR cross-certify one another
  - A bit more complex
  - An alternative to IANA as root model if the community is more comfortable with this approach
- In either case, the critical requirement is that the root be authoritative for prefix & AS # allocation

S-BGP PKI: Cross-Certified Root

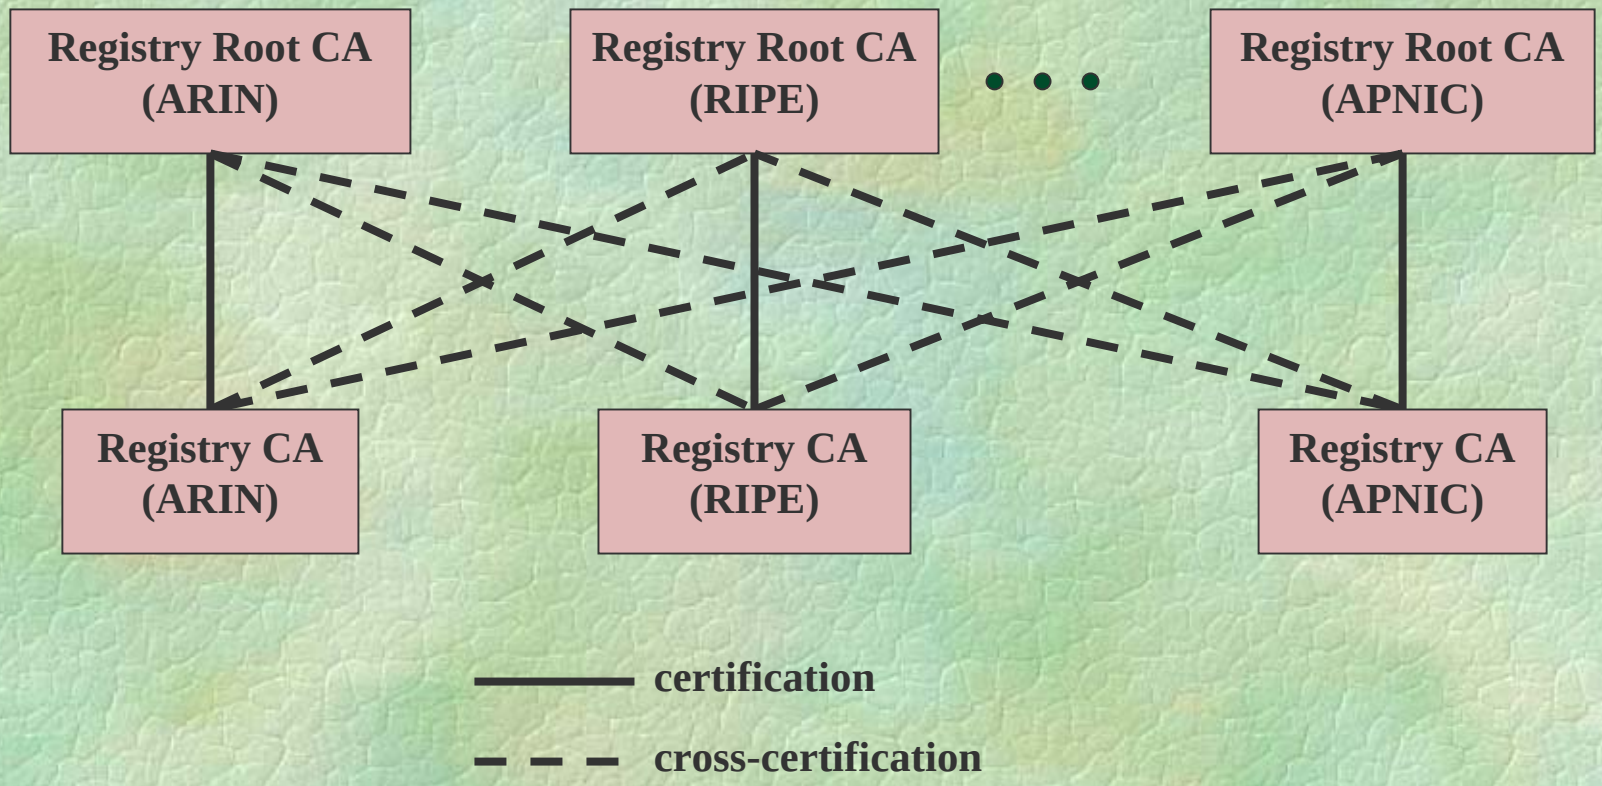

S-BGP PKI: Lower Tiers

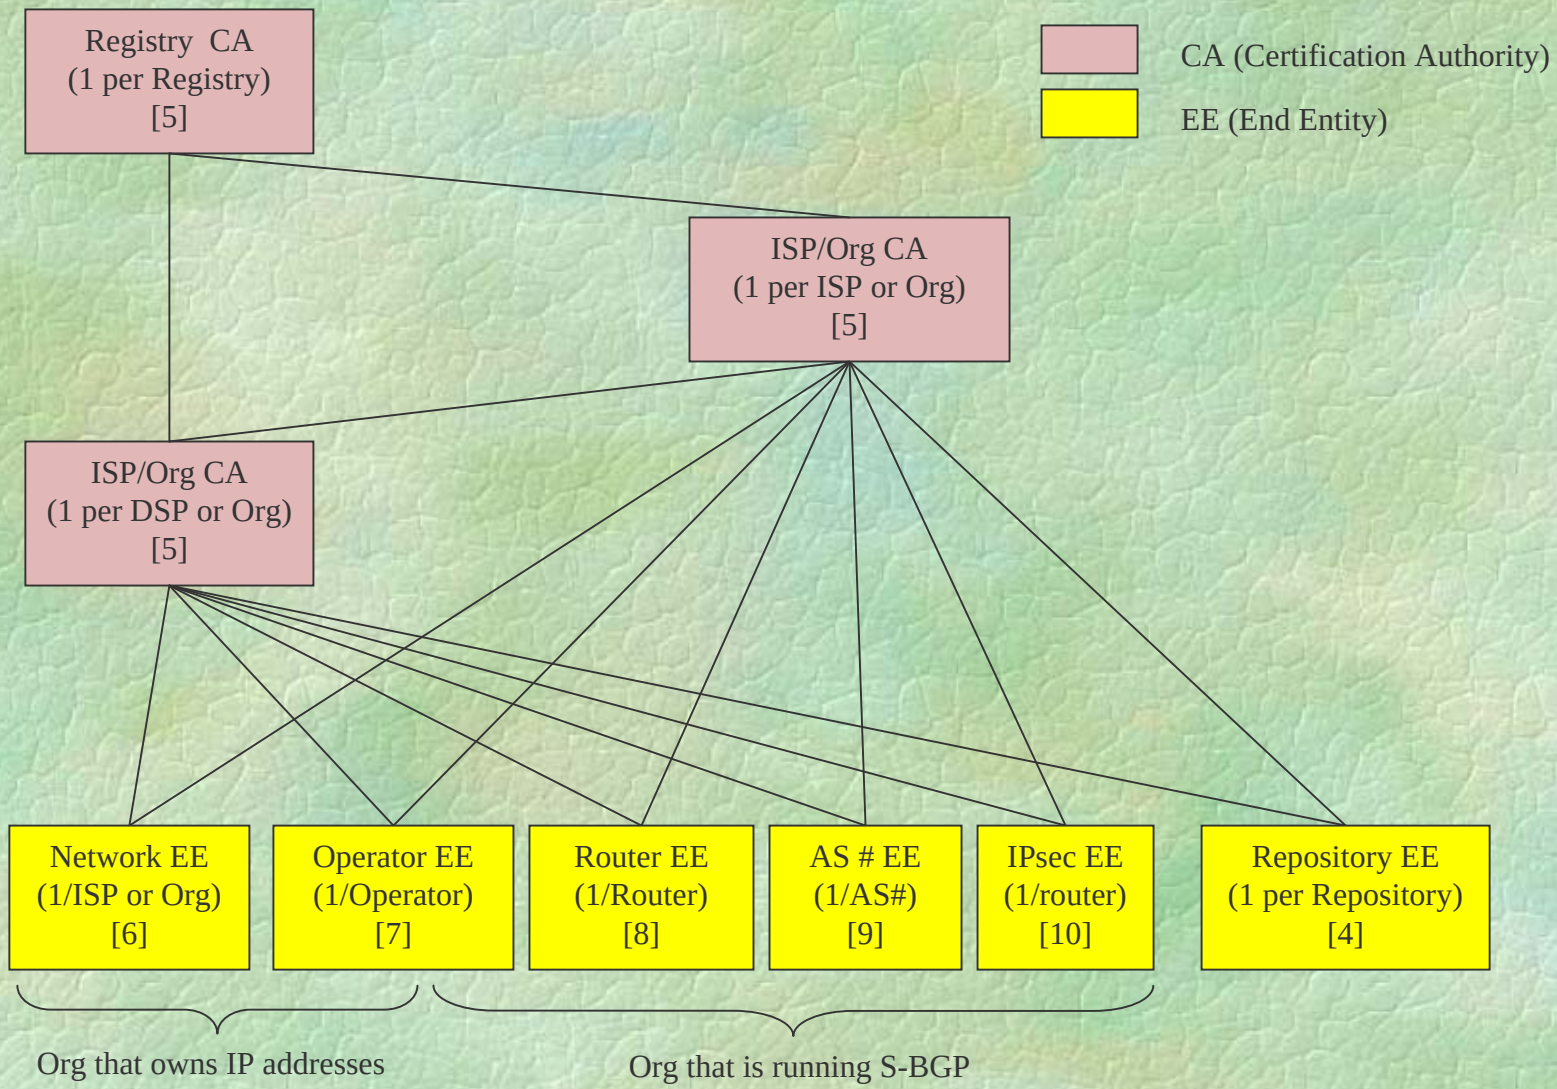

S-BGP PKI Characteristics

- 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 Regional Internet Registry acts as a CA to 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 those 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

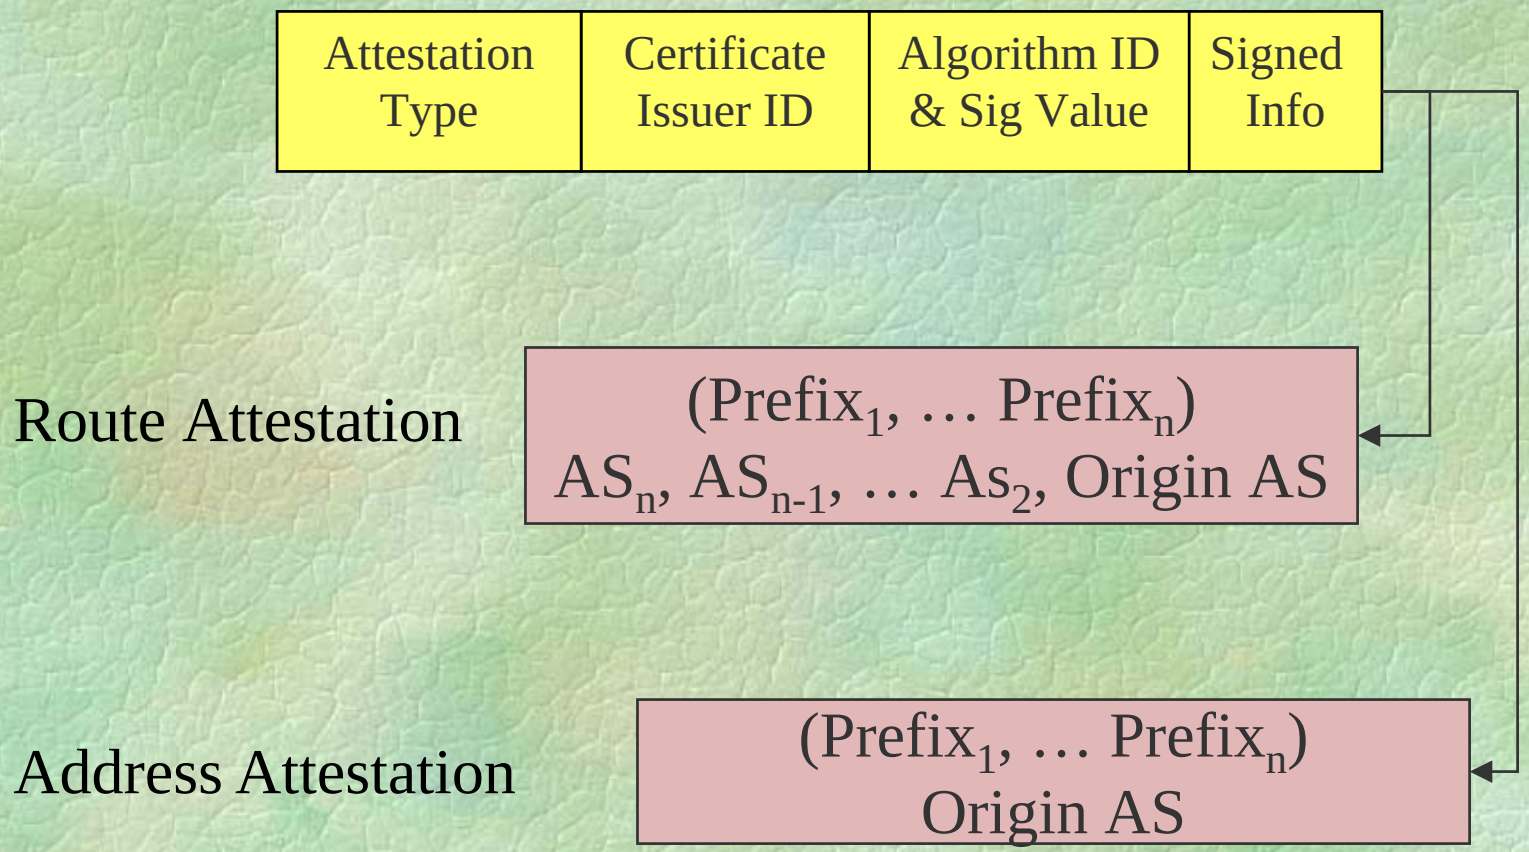

Processing an S-BGP UPDATE

- 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 "owner"

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 exchange sites could operate repositories for other ISPs & subscribers
- Open question: how to find repositories?

S-BGP System Interaction Example

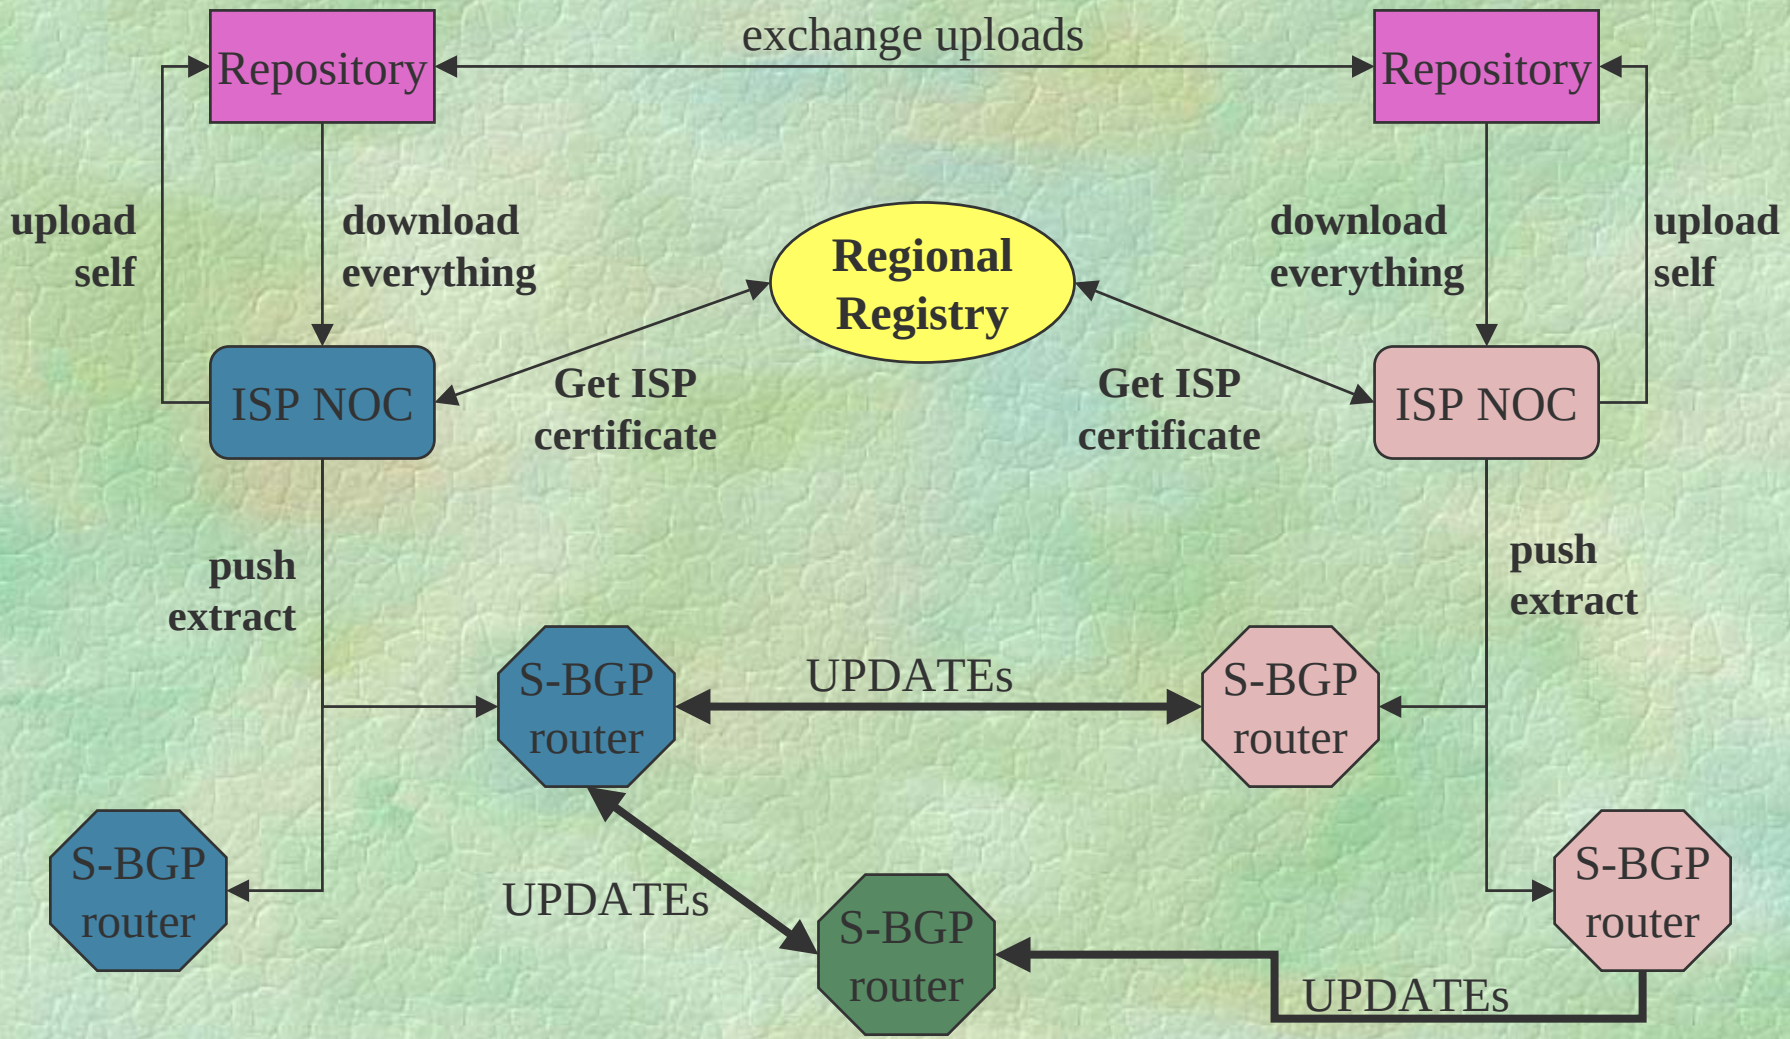

S-BGP Scaling Characteristics

- Certificates issued to ISPs and organizations for prefix and AS # allocation/assignment, and CRLs, correspond to the number of these entities executing S-BGP
- The number of AAs is comparable to the number of prefix owners
- Certificates issued to NOC staff & for IPsec are only local
- Certificates issued for validating RAs grow as the number of ASes (or routers) grows
- The number of RAs in an UPDATE is generally equal to the length of the path in the UPDATE (aggregation can cause the number of signatures to be larger, but aggregation of routes is rare)

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
- S-BGP does not address the more general problem of routers being barraged with management traffic

Deployment Issues for S-BGP

Deploying S-BGP

- S-BGP requires:
  - Router software that implements S-BGP
  - Router hardware with appropriate storage & signature processing capabilities
  - Regional registries must assume CA responsibilities for address prefixes and AS # assignment/allocation
  - ISPs and subscribers that execute BGP must upgrade routers, must act as CAs, and must interact with repositories to exchange PKI & AA data
- **S-BGP can be deployed incrementally, with the constraint that only adjacent S-BGP ASes will receive and make use of S-BGP UPDATEs**

S-BGP Deployment Impediments

- Technical
  - Insufficient memory in most routers for RAs, AAs, public keys, etc.
  - Insufficient non-volatile memory for S-BGP data (e.g., to speed up recovery after reboot)
  - Slow CPUs for management protocol processing
- Procedural
  - NOC & registry staff have to be trained
  - Operations staff have to be convinced it's a good idea
- Economic
  - ISPs cannot afford to replace/upgrade BGP routers
  - Registries cannot afford to offer CA services w/o imposing fees
  - Router vendors cannot afford to implement S-BGP software and hardware unless ISPs will buy it

Router Memory Issues

- Storage of the RAs that accompany UPDATEs, plus the AAs and certificate extracts, might require a total of ~500MB-1GB for RIBs (full deployment, moderate number of neighbors)
- This is just cheap PC memory, nothing special, but most routers have insufficient memory & most cannot be upgraded with more memory
- S-BGP also benefits from a similar amount of non-volatile storage, also generally absent from routers, to speed up recovery after a reboot
- Incremental deployment postpones the need to upgrade router memory, since fewer AAs, RAs, and PKI data would appear initially

Router Performance Issues

- 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 possible surge volume during attacks, and to better protect router keys, use of a crypto accelerator is preferable
- RA validation heuristics can reduce the CPU burden, but some heuristics increase router memory requirements
- Here too, incremental deployment minimizes the processing burden on routers, delays need for hardware upgrades

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
- If validation is deferred, the router should at least check to verify that the RAs were current when the UPDATE was received

Alternative Approaches to BGP Security

MD5 Checksum

- The MD5 checksum mechanism is a cryptographic function that replaces the usual TCP checksum in packets for BGP. It provides only link protection, analogous to use of IPsec's packet integrity function, and so does not protect against attacks that subvert routers, management stations, operator errors, etc. It lacks automated key management, which means keys are often passwords and/or are never changed. It also has crypto weaknesses that make it inferior to the IPsec integrity mechanism

RPSL

- The Routing Policy Specification Language (RPSL) provides ISPs with a standard syntax for publishing a variety of network data in routing registries. This data can encompass address and AS # allocations and assignments, plus information about local policies. A complex security model was developed for managing this data, but with integrity and authentication mechanisms of varying levels of assurance. Some of the data that would be published in a routing registry is viewed as business sensitive by ISPs. Distribution of route data via registries exhibits dynamics not consistent with the propagation of these routes in the Internet. Experience suggests that the data is usually quite stale, exacerbating the problem.

Secure Origin BGP (soBGP)

- This is a new and evolving protocol being developed by Cisco. Despite the name, the protocol encompasses path authorization as well as origin AS advertisements. soBGP might make use of repositories or it might transmit (signed) tables reflecting connectivity and peer authorization. The table data might be processed offline, like S-BGP certificate and AA processing, or it might be performed by routers. The PKI is not well defined. At this stage of its evolution, too many unspecified details of the protocol make it hard to analyze.

Some Criticisms of S-BGP

- All ISP operations personnel hate S-BGP (false)
- It's too complex (eye of the beholder?)
- It can't be deployed incrementally
- It's not an IETF standard (true, but …)
- It's not ready for prime time (maybe)
- Signature processing will overwhelm routers (probably not, certainly not with new hardware)
- It requires ISPs to publish local policy info (false)
- Operations personnel can't understand it (?)
- Repositories create new DoS vulnerabilities (not really a serious problem)

S-BGP Software Status

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 is the only concrete proposal that addresses all of the architectural security problems of BGP, and that responds to route changes in realtime
- The impact on daily RIR & ISP operations is likely to be minimal, although training will be needed
- The S-BGP PKI leverages existing authorization relationships, creates no new ones, and does not require ISPs to disclose any additional data
- Routers will require hardware upgrades for full deployment of S-BGP, an obvious $ problem
- Incremental deployment postpones the need for router upgrades, offers benefits, and is feasible

Questions?

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