A PKI for IP Address Space and AS Numbers

Dr. Stephen Kent

Chief Scientist - Information Security

BBN
TECHNOLOGIES

Presentation Outline

- Why a PKI?
- PKI background
- Address & AS number allocation system
- The proposed PKI
  - Structure
  - Names and chaining
  - Certificate & CRL profile features
  - “Shadow” certificates and signed objects
  - Repository issues

Why A PKI?

- Proposals for improving BGP security rely on a secure infrastructure that attests to address space and AS number holdings by ISPs and subscribers
- A PKI is a natural way to satisfy this requirement
- The proposed PKI provides a first step towards improved BGP security; it can be used to help:
  - Detect bogus route origination info in UPDATEs
  - ISPs avoid “social engineering” attacks that attempt to trick them into issuing bogus routes
  - Securely communicate requests from one ISP to another to change route filters

PKI Terminology

- Certificate: a digitally-signed data structure; typically an X.509 public key certificate (PKC), the certificate standard adopted by the IETF and employed in SSL/TLS, IPsec (IKE), S/MIME, and many other security protocol standards
- Certification Authority (CA): an entity that issues (signs) certificates, aka an Issuer
- Subject: an entity to whom a certificate is issued; for a PKC, the subject is the holder of the private key corresponding to the public key in the certificate
- End entity (EE): a certificate subject that does not issue certificates, i.e., does not act as a CA
- Relying party (RP): an individual or organization that takes actions based on using a public key from a certificate

More PKI Terminology

- Trust anchor (aka root): a public key and associated data used as a reference for validating certificates
  - A trust anchor is often represented as a self-signed certificate, but it need not be
- Certification path: a series of certificates between a trust anchor and a certificate being validated, linked by subject/issuer name
- Certificate validation: the process of determining that a certificate is valid
  - creating a certificate path between a certificate and a trust anchor
  - verifying the signature on each certificate in the path
  - checking the revocation status of each certificate in the path
- PKI: a set of procedures, policies, and technical measures employed to manage (issue, renew, revoke, publish) certificates

X.509 Certificate (v3)

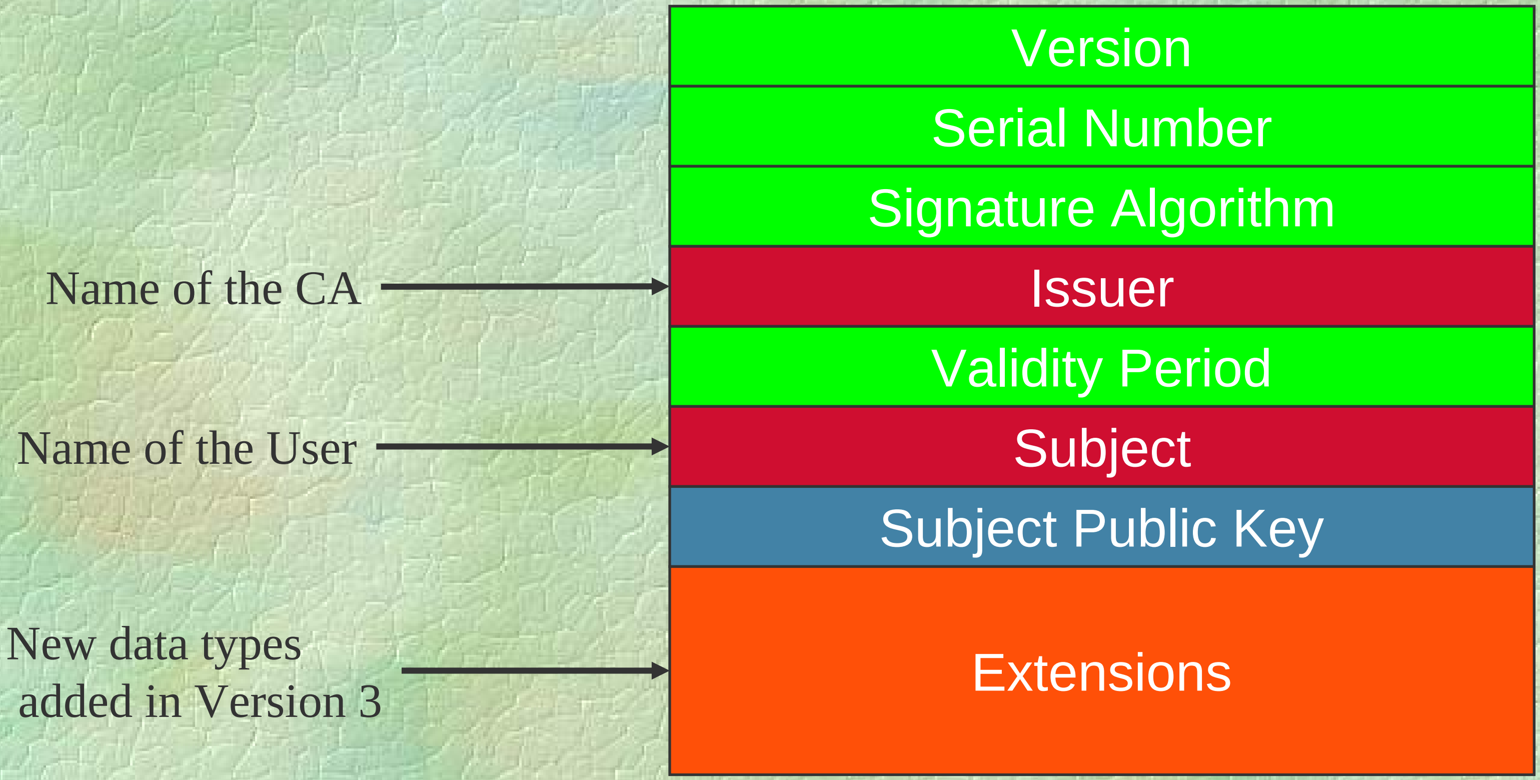

Version

Serial Number

Signature Algorithm

Name of the CA → Issuer

Validity Period

Name of the User → Subject

Subject Public Key

New data types added in Version 3 → Extensions

What Does the PKI Look Like?

- The PKI consists of two major parts:
  - X.509 certificates that attest to address space and AS number holdings
  - A repository system for these certificates, CRLs, and other signed objects that are globally useful
- The PKI makes use of the existing address space and AS number allocation system
- This PKI also embodies the “principle of least privilege,” which constrains the impact of errors or security compromise at each entity in the PKI, relative to the authorization of that entity

Matching the PKI to Reality

- The intent in this PKI is to issue certificates that attest to resource holdings by registries, ISPs, and subscribers (where appropriate)
- Because the allocation of these resources is done via a simple, hierarchic scheme, the PKI should parallel this scheme
- Each entity that participates in the allocation process should act as a CA, issuing certificates to match the resource allocation records of that entity

Address Allocation Hierarchy

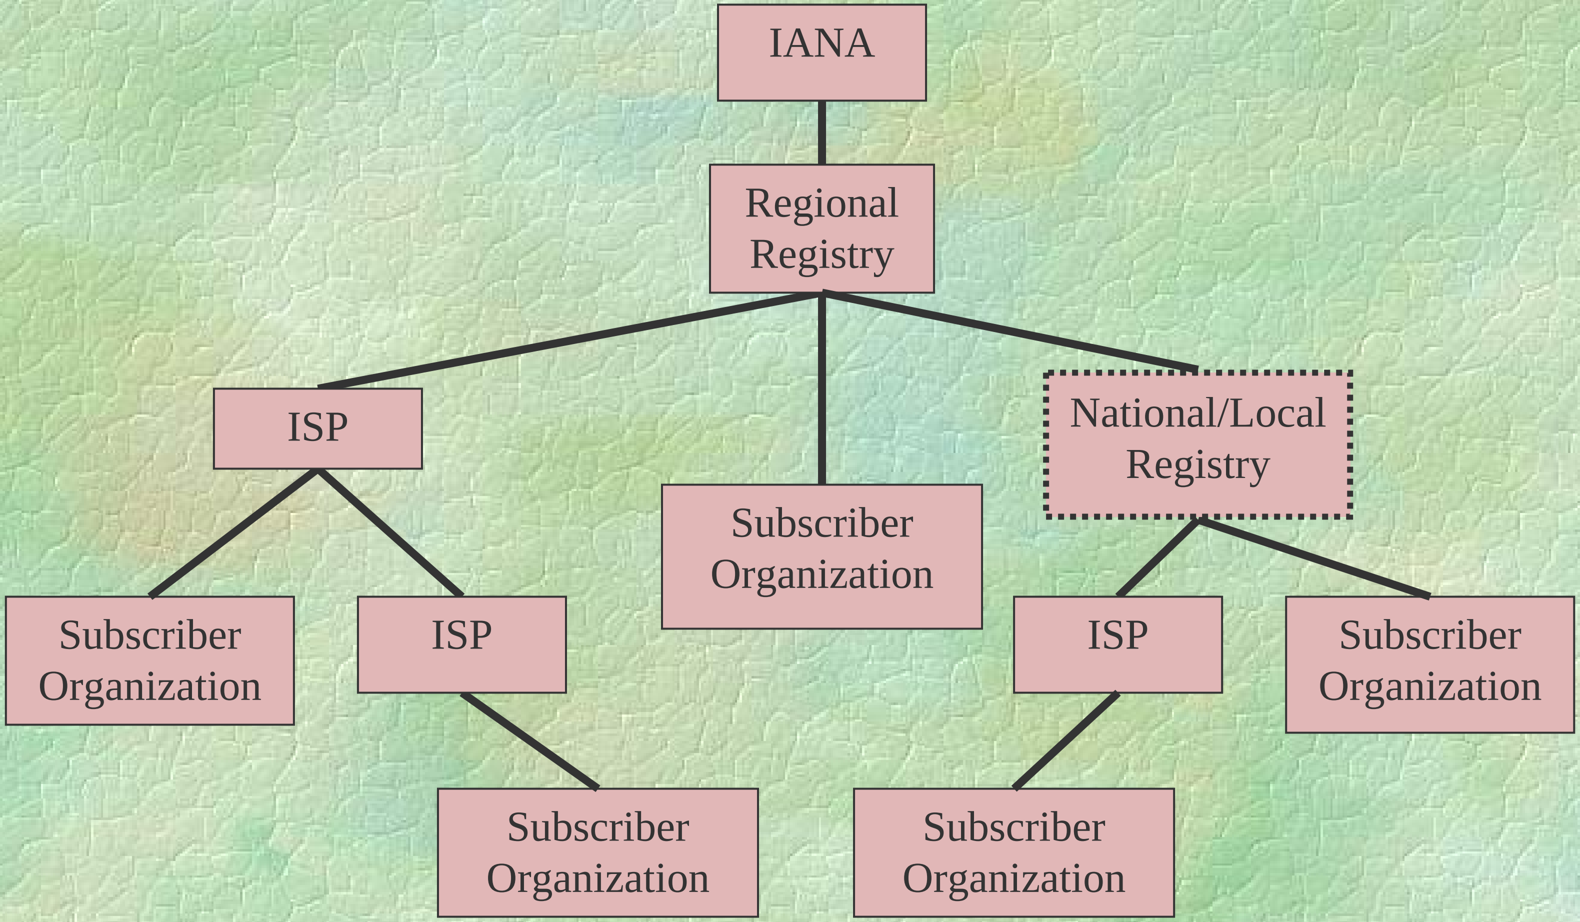

AS Number Assignment Hierarchy

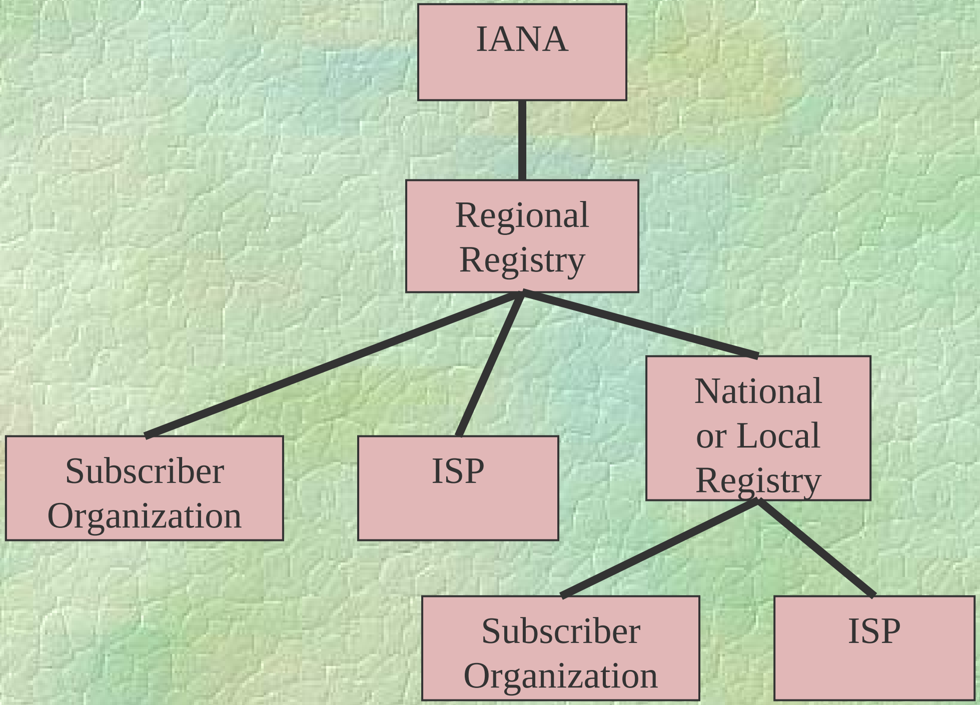

IANA

Regional Registry

Subscriber Organization

ISP

National or Local Registry

Subscriber Organization

ISP

How Will the PKI Work?

- The root issues certificates to the 5 RIRs, and each RIR issues certificates to national/local registries (if applicable) and to ISPs and subscribers
- ISPs issue certificates to downstream providers and to subscribers
- Each organization issues certificates that match the address space (and AS number) allocations in its records
- All resource holders are CAs
- The PKI uses two X.509 extensions (defined by RFC 3779) to represent address and AS number allocation data
- Each certificate path represents sub-allocation by the organizations noted above, a subset constraint that can be verified by ISPs downloading these certificates

PKI Top Tier Example (APNIC)

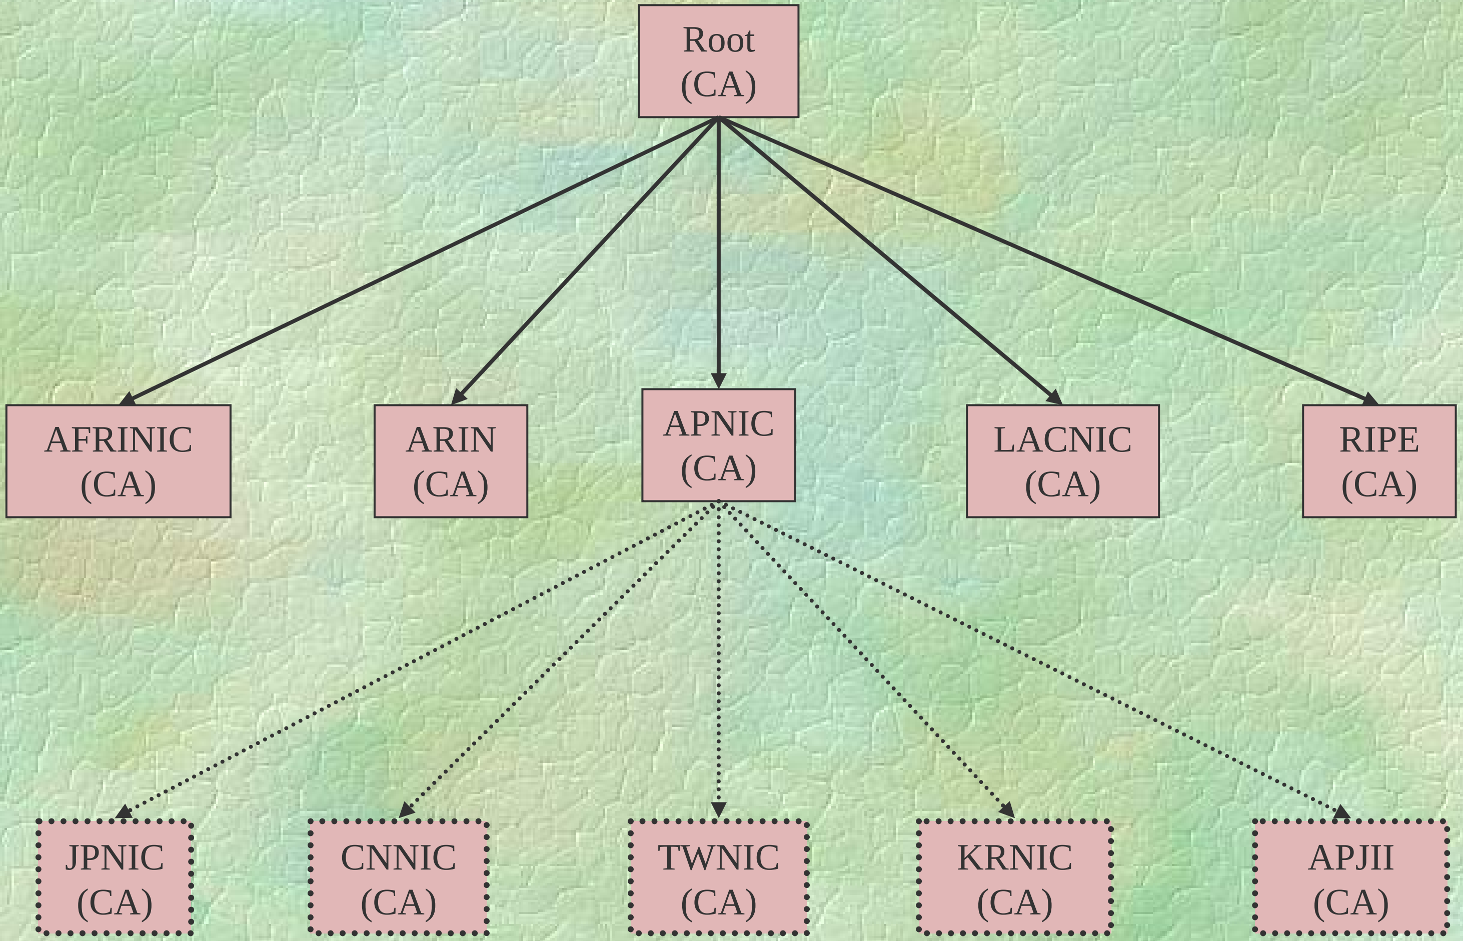

PKI Additional Details Example

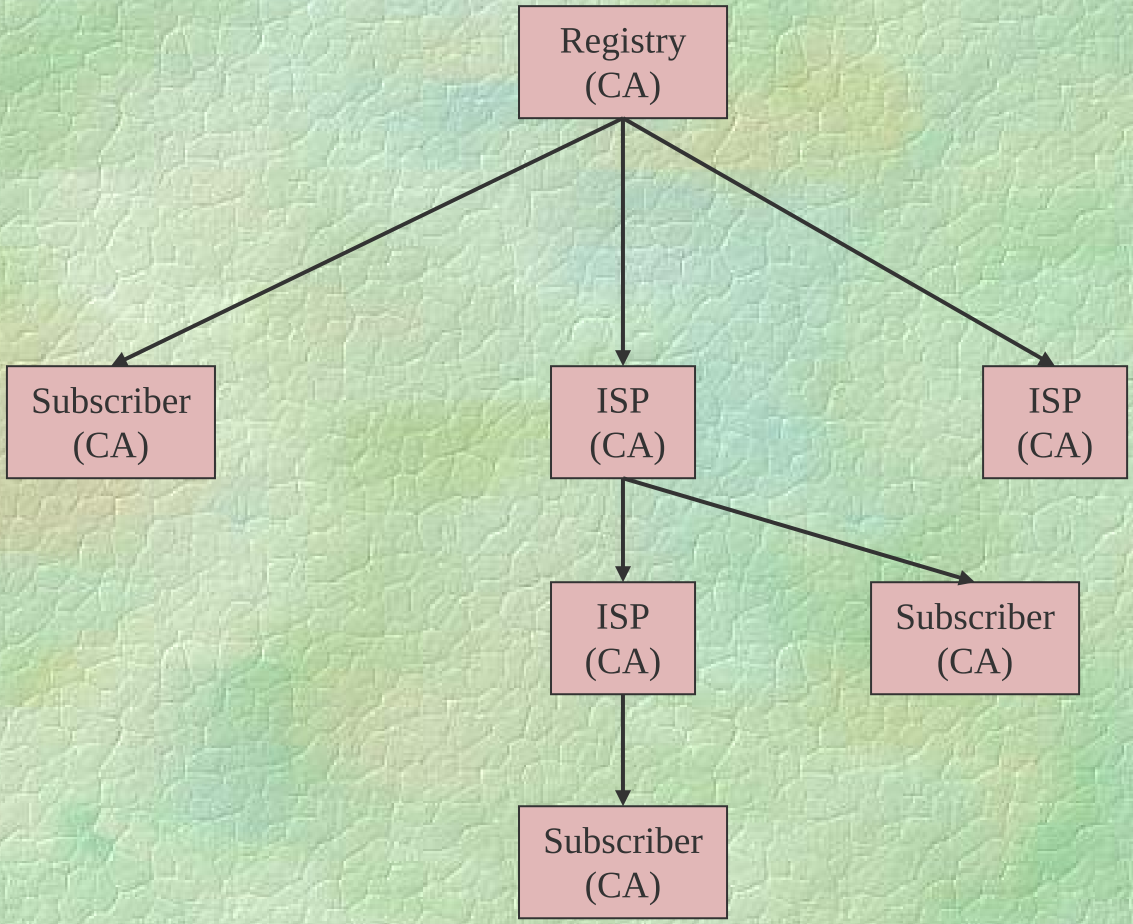

Certificate Chain Example

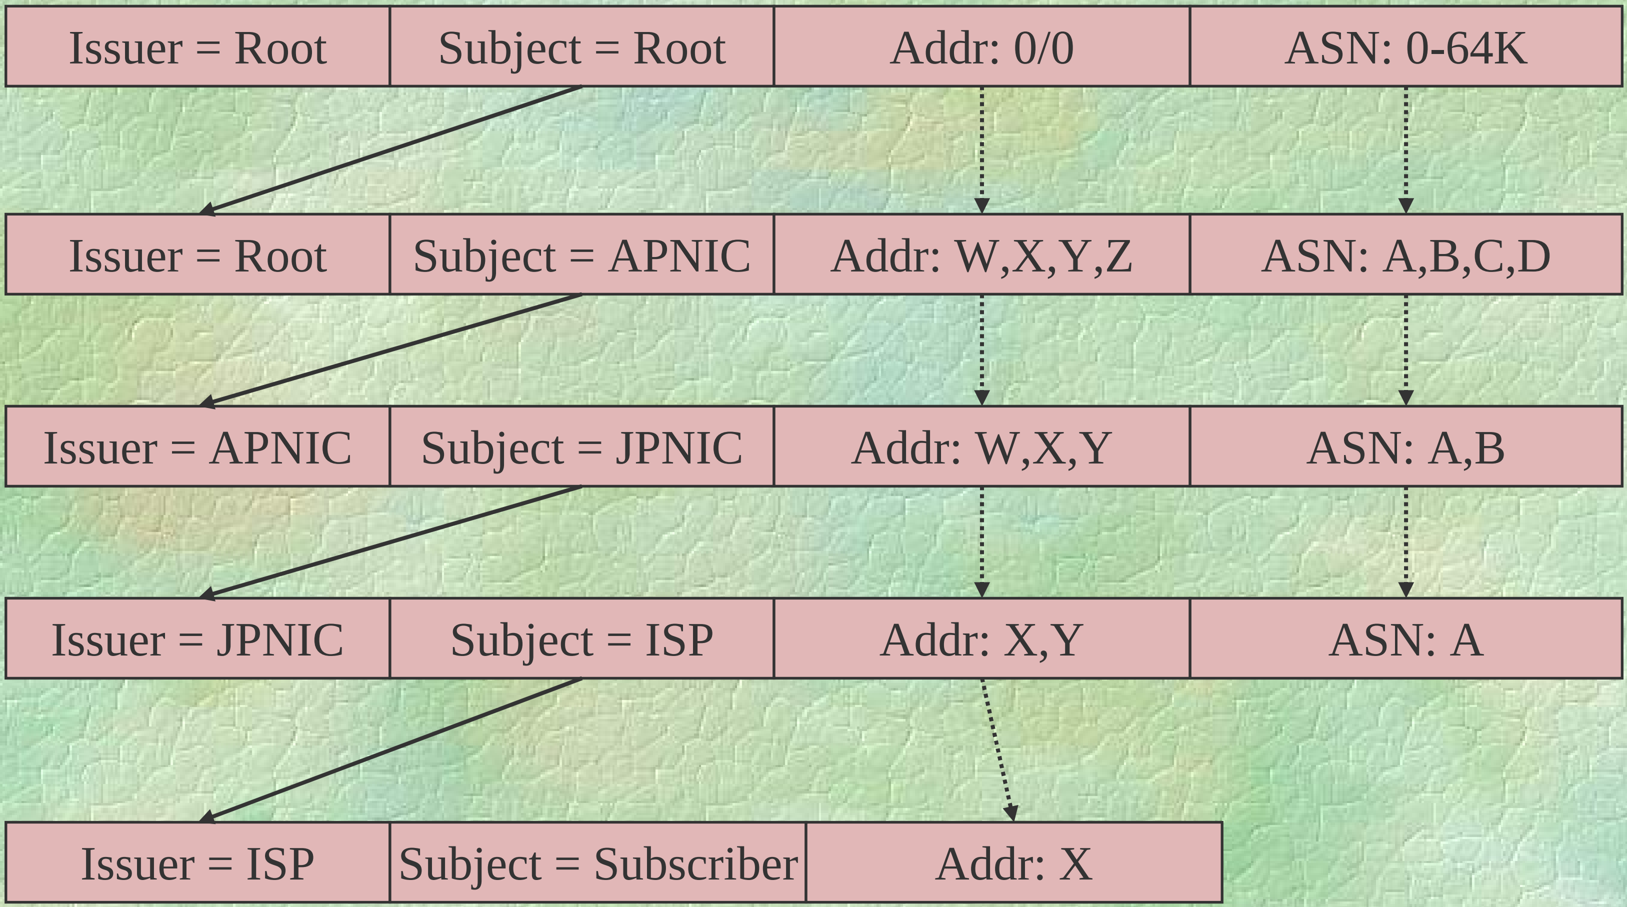

Issuer = Root	Subject = Root	Addr: 0/0	ASN: 0-64K
Issuer = Root	Subject = APNIC	Addr: W,X,Y,Z	ASN: A,B,C,D
Issuer = APNIC	Subject = JPNIC	Addr: W,X,Y	ASN: A,B
Issuer = JPNIC	Subject = ISP	Addr: X,Y	ASN: A
Issuer = ISP	Subject = Subscriber	Addr: X	

Names in Certificates

- Because the intent of the PKI is to enable digital signing of objects that express authorization, is it not necessary for these certificates to contain meaningful names!
- This is a big departure from most PKI designs, but it is appropriate for this context, and it helps avoid liability issues for CAs when dealing with lots of subjects
- It may be appropriate and easy to use meaningful names for the top tiers (registries), as they are not the resource users
- But, we should not include the registry name as a prefix for an ISP (or subscriber) subject name, since then we cannot assign the same name to the same organization if that organization holds resources from two different registries!

Some Name Examples

- RIR CA name
  - C = AU, O = APNIC, OU = Resource Registry CA
- NIR CA name
  - C= JP, O = JPNIC, OU = Resource Registry CA
- ISP or subscriber CA name
  - CN = FC3209809267

Certificate Extensions for this PKI

- Basic Constraints
  - Marks the certificate as for a CA (vs. an EE)
- Certificate Policy
  - Marks the certificate as being restricted to use with this PKI
- Address space & AS Number (RFC 3779)
  - Lists of address and AS number ranges (two extensions)
- Key Usage
  - Indicates how the public key may be used
- Key Identifiers
  - IDs used as pointers to help select the right CA certificate
- Authority Information Access
  - Pointers to the Issuer’s certificate and OCSP server
- CRL Distribution Point
  - Used to locate the CRL for the certificate in question

“Shadow” Certificates

- Good PKI practice says that a CA should NOT sign objects other than certificates and CRLs
- So, we introduce an end-entity (EE) certificate under each ISP & each subscriber CA, and use the corresponding private key to sign objects such as route origination authorizations (ROAs), route filter charge requests, etc.
- We call this EE certificate a “shadow” certificate, since it can represent exactly the same authorization as the CA certificate under which it is issued
- Indirection helps manage revocation of signed objects, i.e., to revoke a signed object before it expires, revoke the shadow certificate (issue a CRL with that certificate)
- A CA can create as many shadow certificates as it wants, as a means of restricting the authorization associated with each

PKI with Shadow Certificates

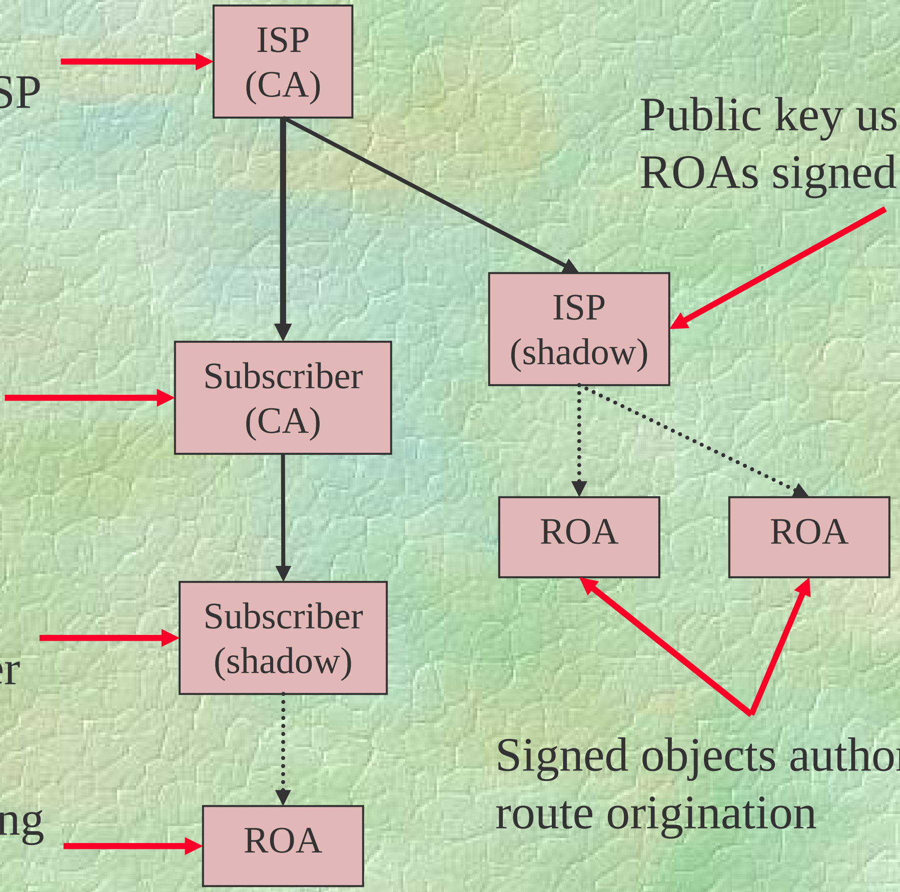

Public key used to verify certificates issued by the ISP

Public key used to verify ROAs signed by the ISP

Public key used to verify certs issued by the subscriber

Public key used to verify ROAs signed by the subscriber

Signed objects authorizing route origination

Signed object authorizing route origination

Signed Object Example: ROA

- A ROA consists of the following elements
  - Address prefixes: one of more prefixes, corresponding to the NLRI that the ROA signer authorizes for origination by one or more ISPs, enumerated below
  - AS numbers: the ISP(s) authorized to originate routes to the above list of prefixes
  - Validity Interval: start and end time & date defining the interval for which the ROA is valid
  - Signature list: a list of one or more pairs of the following data
    - Certificate pointer: data to help a verifier locate the shadow certificate needed to verify this ROA (subject key ID)
    - Signature: a digitally signed hash of the first 3 items, plus the hash algorithm and digital signature algorithm OID

ROA Format

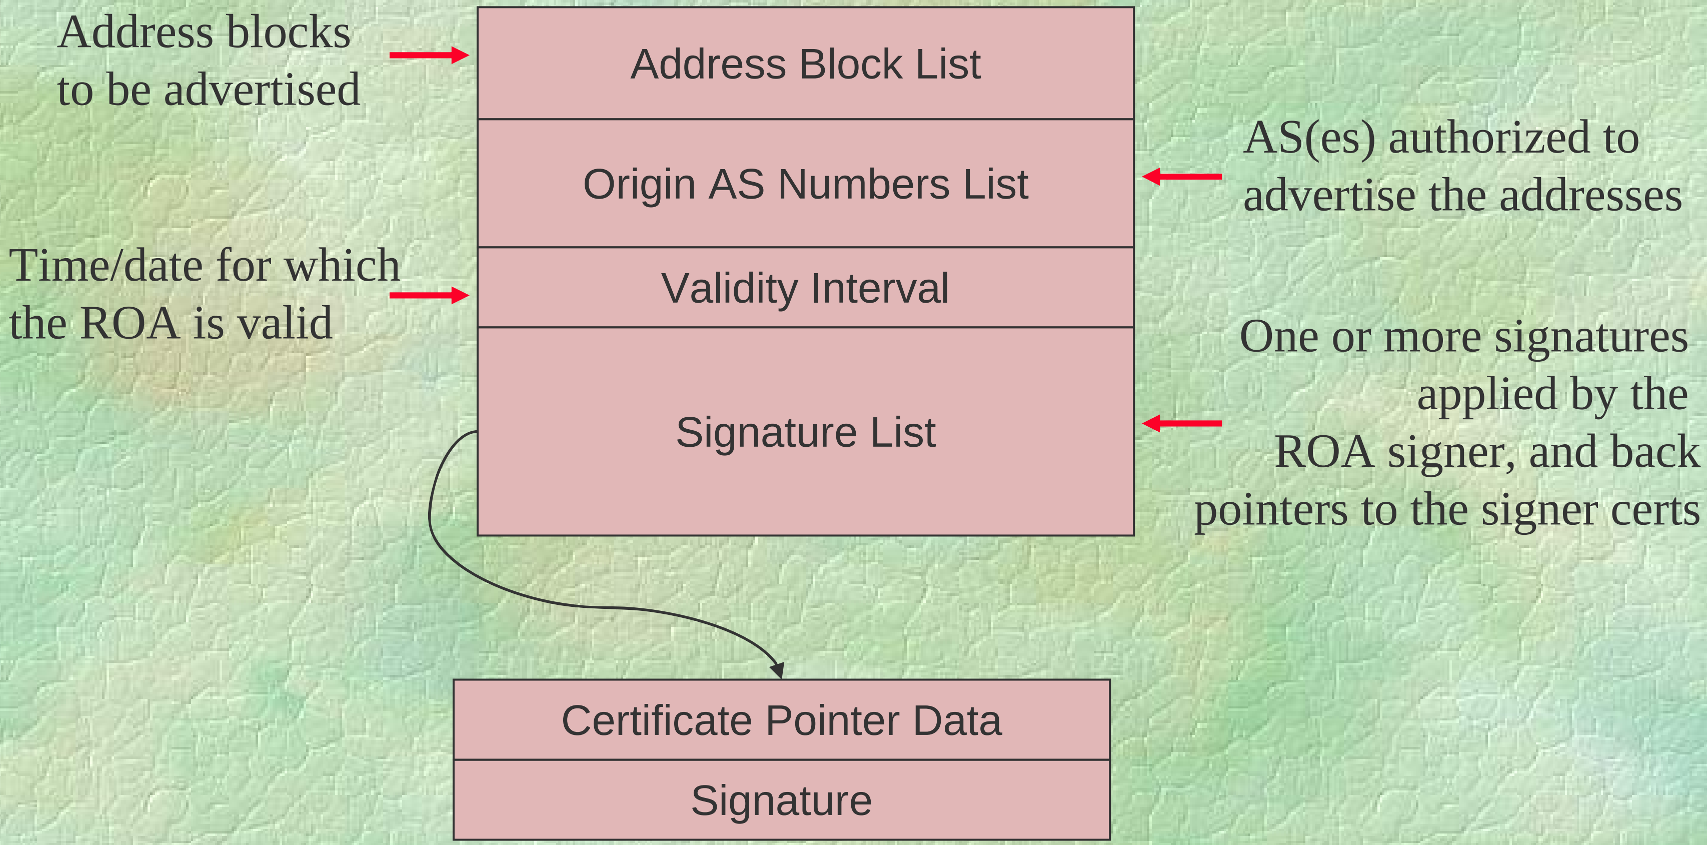

Address blocks
to be advertised

Address Block List

Origin AS Numbers List

AS(es) authorized to
advertise the addresses

Time/date for which
the ROA is valid

Validity Interval

Signature List

One or more signatures
applied by the
ROA signer, and back
pointers to the signer certs

Certificate Pointer Data

Signature

Generating a ROA

- An ISP (or subscriber) generates one ROA for each prefix for which it wants to authorize route origination
  - If all prefixes of the resource holder are to be advertised by the same ISPs, and came from one source, then the entity signs one ROA with all the prefixes and all the AS numbers of the holder
  - If the resource holder wants to authorize some ISPs to originate some prefixes, and other ISPs to originate other prefixes, the holder signs one ROA for each set of addresses to be independently originated
  - If a resource holder has addresses from different sources, it can sign the ROA multiple times, using the private key associated with each relevant shadow certificate

Repository Issues

- This repository is unusual in that ALL of the data is signed and verifiable via certificate path validation
- Most repositories used for certificates and CRLs, e.g., LDAP, assume searching & selective retrieval of entries
  - Users search the repository for specific entries and retrieve the entries, or selected data items within an entry
- For route origination authorization, the retrieval model is very different
  - EVERY ISP would request ALL changed entries since the last time it processed ROAs
  - For many ISPs, this would be a daily retrieval request
- This suggests a different repository model
  - While the repository system need not be very highly available, it needs to be fairly robust
  - The large user population (all ISPs) should not have to query a very large number of sites

A Repository Model

- One approach is a repository model analogous to the whois database system, one repository per RIR, serving all the RIR members and the entities to whom sub-allocations have been made
- ISPs & subscribers upload their own new data, download repository changes, on a daily basis
- Each ISP will need to contact each RIR repository to gather all the data need to verify ROAs
- Repositories can use the PKI to enforce access controls to counter DoS attacks
  - Upload access granted only to PKI users, authenticated via shadow certificates issued to operations personnel
  - An ISP or subscriber is automatically prevented from overwriting data of another ISP or subscriber

Frequency of Repository Updates

- A resource holder needs to upload changed certificates, CRLs, and globally useful signed objects (e.g., ROAs) to the repository
  - Certificates will usually change infrequently, only when new allocations are received
  - An ISP decides how frequently to issue its own CRL, so he control the update frequency for that data item
  - ROAs change only when allocations change, or when origination authorization changes, presumably not too often
- For smaller ISPs and subscribers, changes to the data will be very infrequent
- For large ISPs, daily updates will probably suffice

Using the PKI

- Route filter generation procedure
  - Download all the (changed) repository data: certificates, CRLs, and ROAs
  - Verify the certificate paths
  - Use shadow certificates to verify ROAs
  - Construct a table of authorized origin ASes and address prefixes from the validated ROAs
- Securing route origination requests
  - Subscriber (or downstream ISP) sends a ROA to the ISP that it wants to advertise its prefix, e.g,, via S/MIME
  - ISP verifies the ROA and that the sender is the subscriber in question
  - ISP can now accept request from user with confidence

Status

- Test certificates and CRLs are being generated
- A draft CP for the PKI has been written
- A draft CPS for registries and one for ISPs has been written
- APNIC is developing software to support the PKI, and is running a trial for their region
- The RIRs have met to discuss this proposal, and are working to refine details of the design

Questions