Idealized BGPsec: Formally Verifiable BGP

NANOG / Denver

2011.06.14

Randy Bush <randy@psg.com>

for the

Informal BGPsec Design Team

BGPsec Design Team

chris morrow (google)
dave ward (juniper)
doug maughan (dhs)
doug montgomery (nist)
ed kern (cisco)
heather schiller (uunet)
jason schiller (uunet)
john scudder (juniper)
kevin thompson (nsf)
keyur patel (cisco)
kotikalapudi sriram (nist)
luke berndt (dhs)
matt lepinski (bbn)
pradosh mohapatra (cisco)
randy bush (iij)
rob austein (isc)
ruediger volk (deutsche telekom)
russ housley (vigilsec)
russ mundy (sparta)
sam weiler (sparta)
sandy murphy (sparta)
sharon goldberg (boston uni)
steve bellovin (columbia uni)
steve kent (bbn)
warren kumari (google)

http://www.ietf.org/iesg/statement/design-team.html

Assume RPKI is a Given

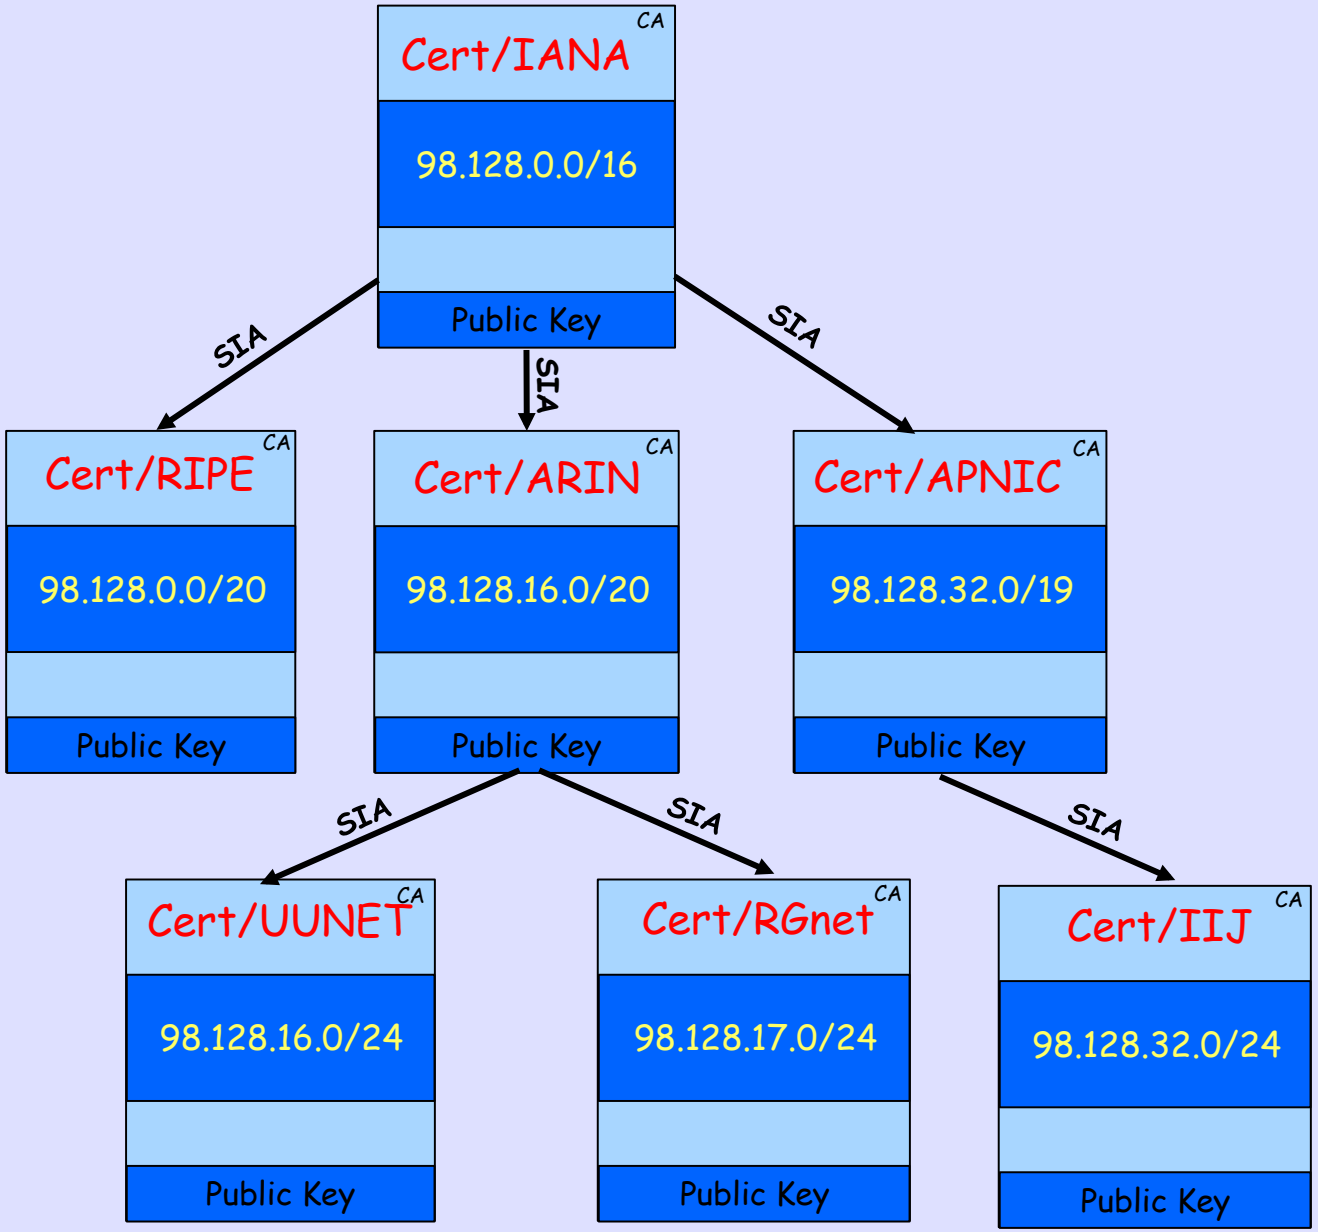

Assume ROAs are a Given

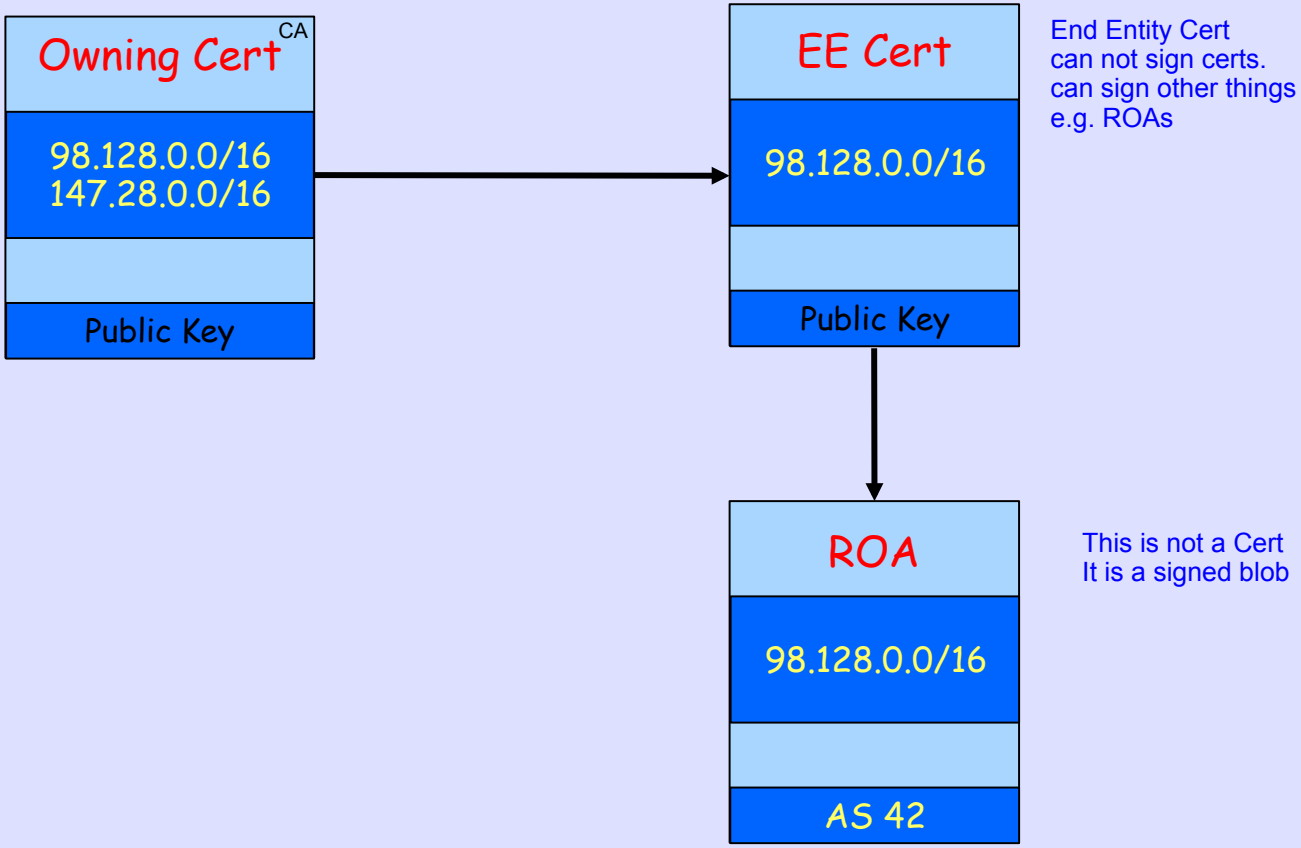

Assume RPKI-RTR Given

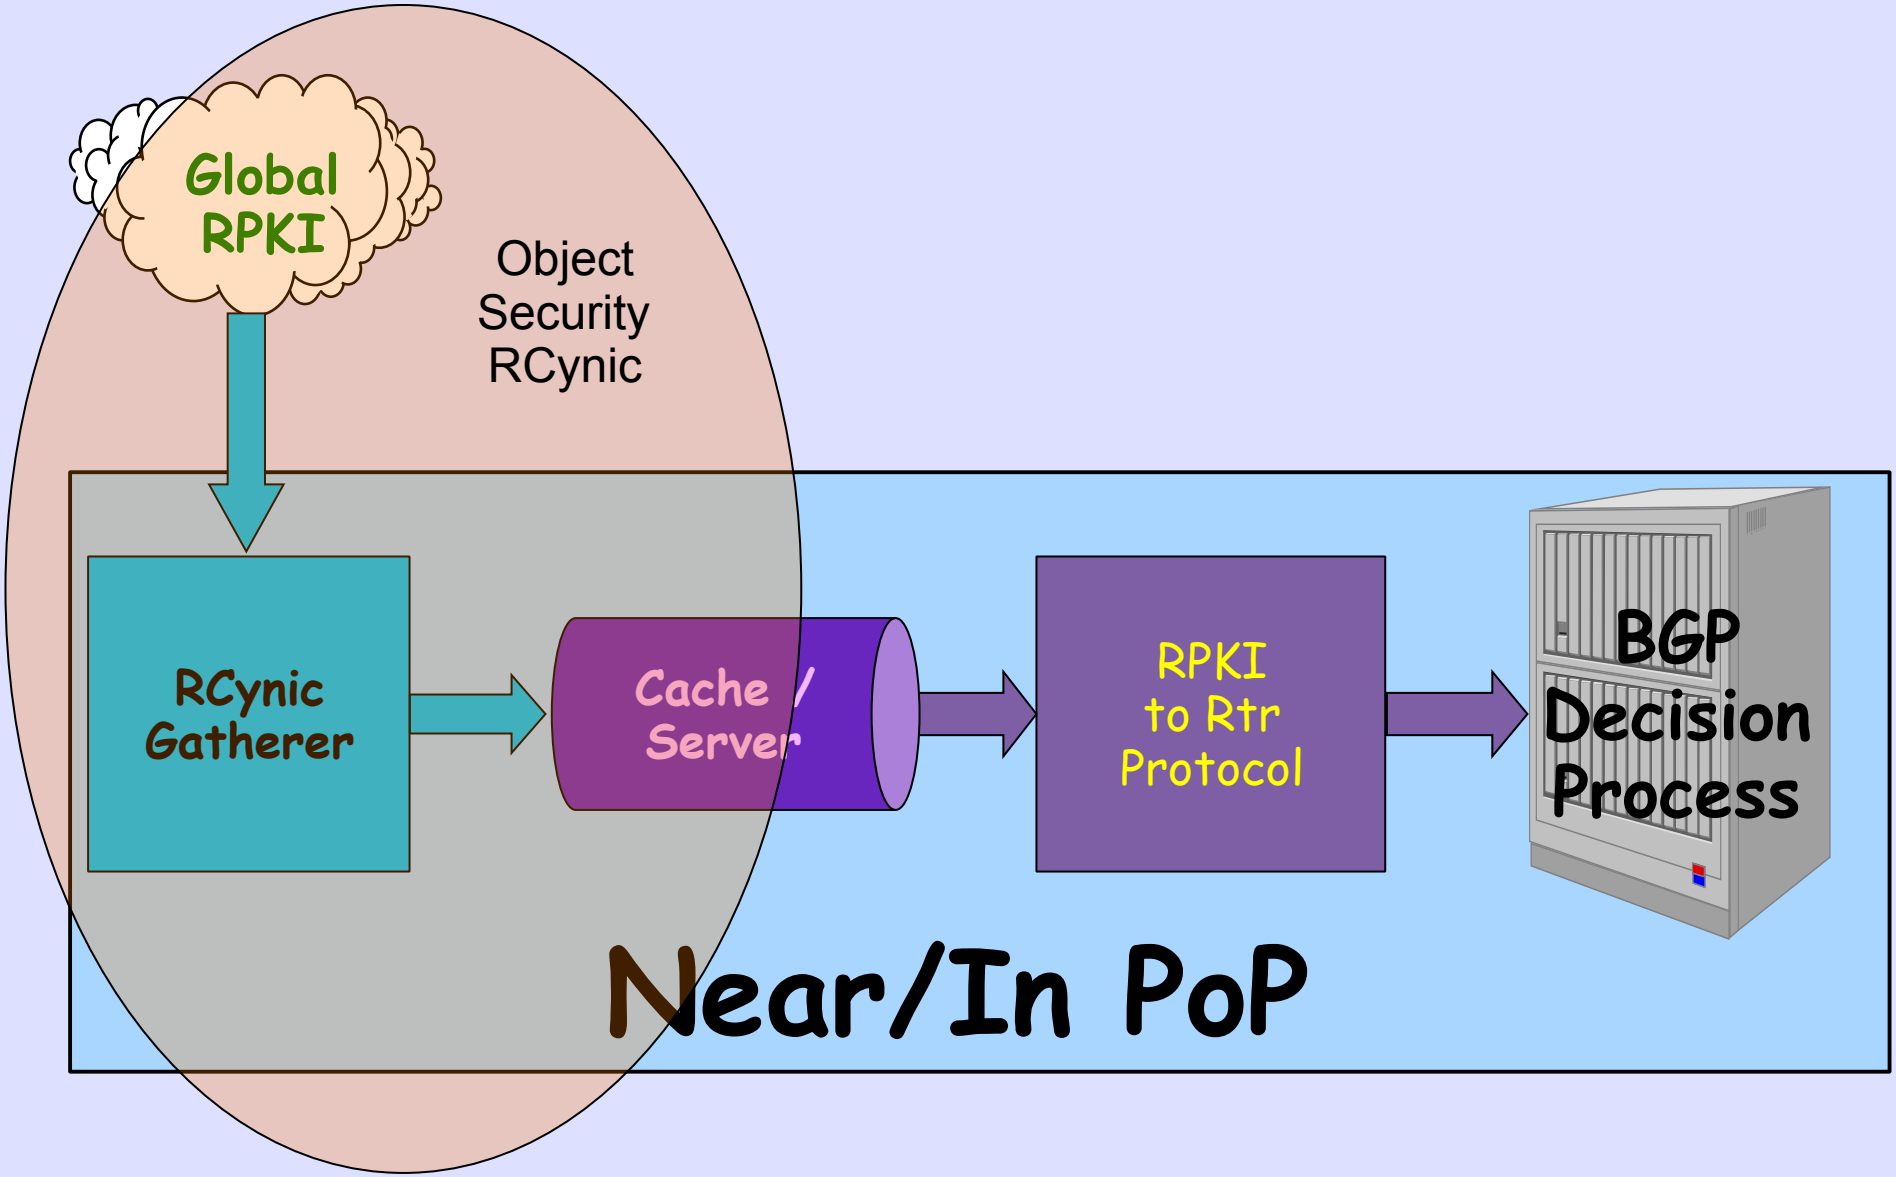

Assume Origin Validation

```
R3#sh ip bg 98.128.0.0/24
BGP routing table entry for 98.128.0.0/24, version 299
Paths: (2 available, best #2, table default)
 65000 3130
    10.0.0.1 from 10.0.0.1 (193.0.24.64)
      Origin IGP, localpref 100, valid, external
      path 680D859C RPKI State invalid
 65001 4128
    10.0.1.1 from 10.0.1.1 (193.0.24.65)
      Origin IGP, localpref 100, valid, external, best
      path 680D914C RPKI State valid
```

The Gap

RPKI-based Origin Validation provides neither cryptographic assurance (announcements are not signed), nor assurance of the AS Path of the announcement

Origin Validation is Weak

- Today's Origin Validation only stops accidental misconfiguration, which is quite useful. But ...
- A malicious router may announce as any AS, i.e. forge the ROAed origin AS.
- This would pass ROA Validation as in draft-ietf-sidr-pfx-validate.

Full Path Validation

- Rigorous per-prefix AS path validation is the goal
- Protect against origin forgery and AS-Path monkey in the middle attacks
- Not merely showing that a received AS path is not impossible
- Yes, this is S-BGP-like not SO-BGP-like

Protocol Not Policy

- Policy on the global Internet changes every 36ms
- We already have a protocol to distribute policy or its effects, it is called BGP
- We can not know intent, should Mary have announced the prefix to Bob
- But Joe can formally validate that Mary did announce the prefix to Bob
- BGPsec validates that the protocol has not been violated, and is not about intent or business policy

Path Shortening Attack

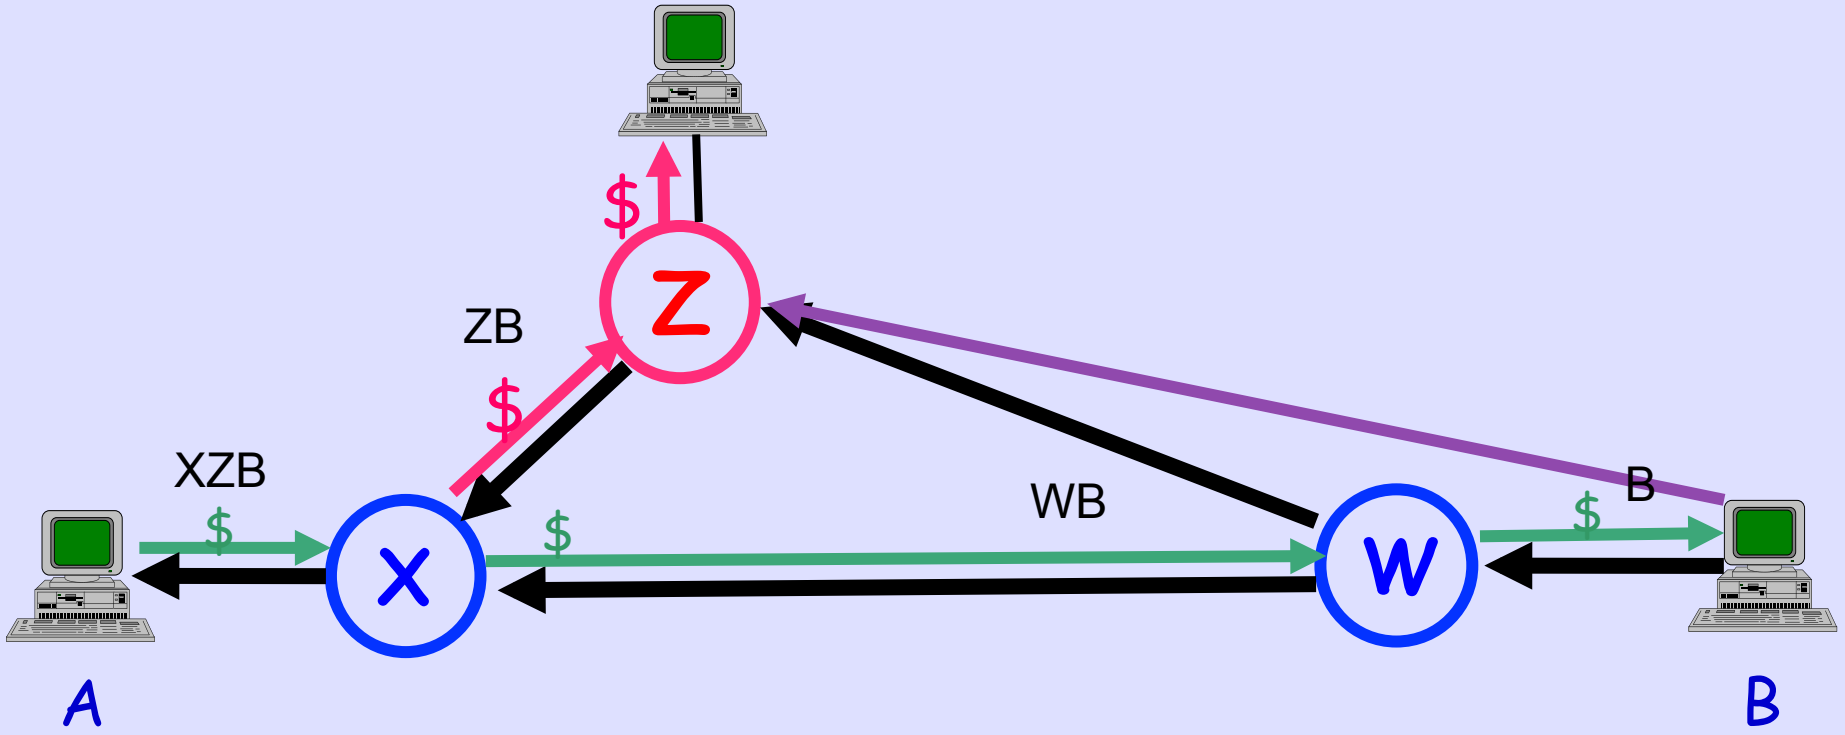

Expected Path - A->X->W->B

Diverted Path - A->X->Z->W->B

There Are Many Many Other Attacks

Forward Path Signing

AS hop N signing (among other things) that it is sending the announcement to AS hop N+1 by AS number, is believed to be fundamental to protecting against monkey in the middle attacks

Forward-Signing

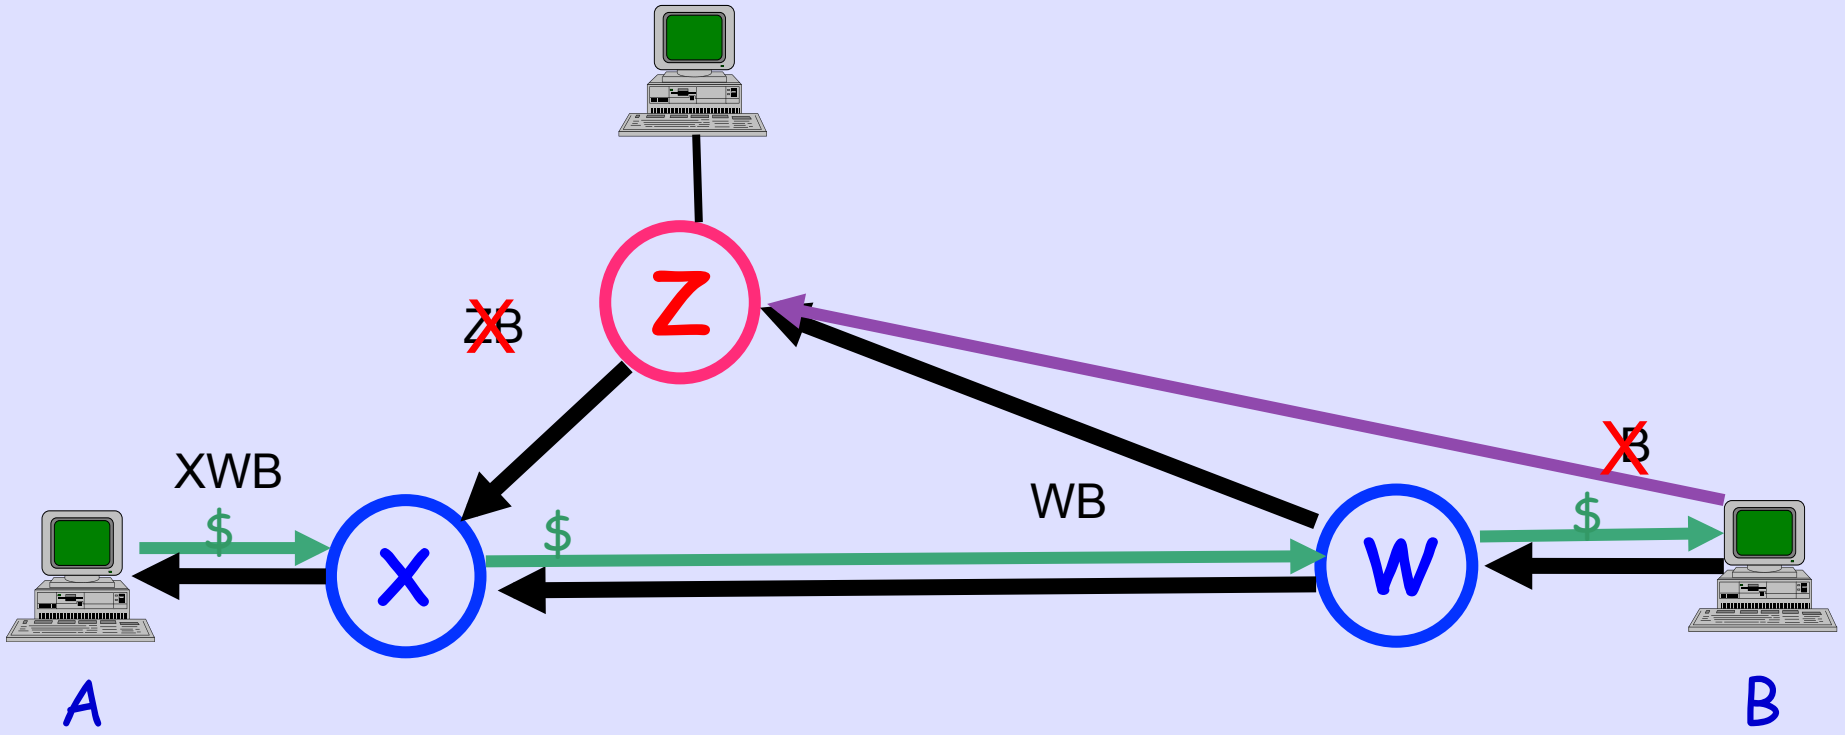

B cryptographically signs the message to W Sb(B->W)

W signs messages to X and Z encapsulating B's message

Sw(W->X (Sb(B->W))) and Sw(W->Z (Sb(B->W)))

X signs the message to A Sx(X->A (Sw(W->X (Sb(B->W)))))

Z can only sign Sz(Z->X (Sw(W->Z (Sb(B->W)))))

Capability Negotiation

- It is assumed that consenting routers will use BGP capability exchange to agree to run BGPsec between them
- The capability will, among other things remove the 4096 PDU limit for updates
- If BGPsec capability is not agreed, then only traditional BGP data are sent

Replay Attack

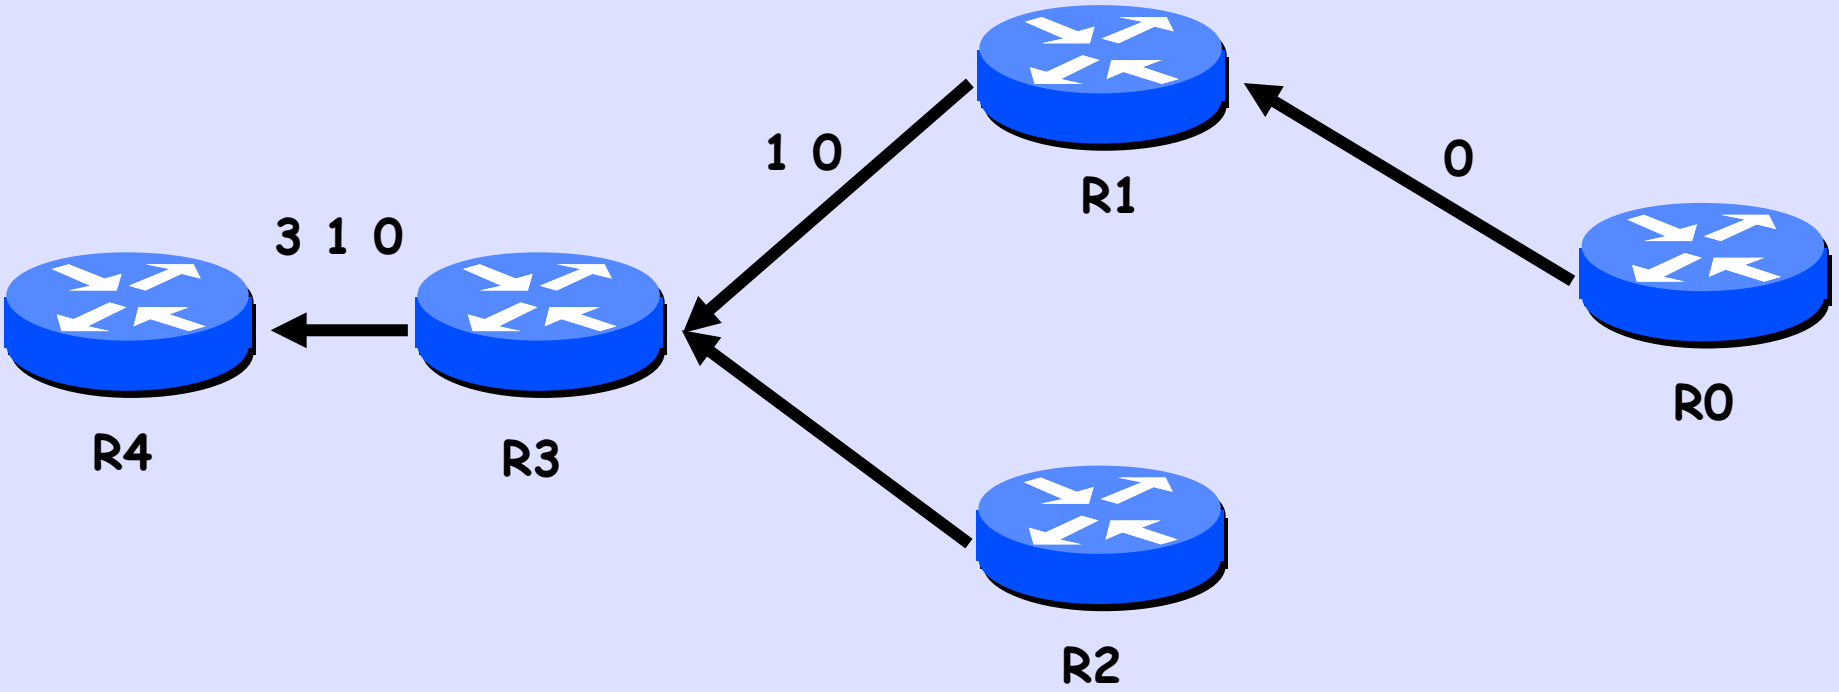

Replay Attack

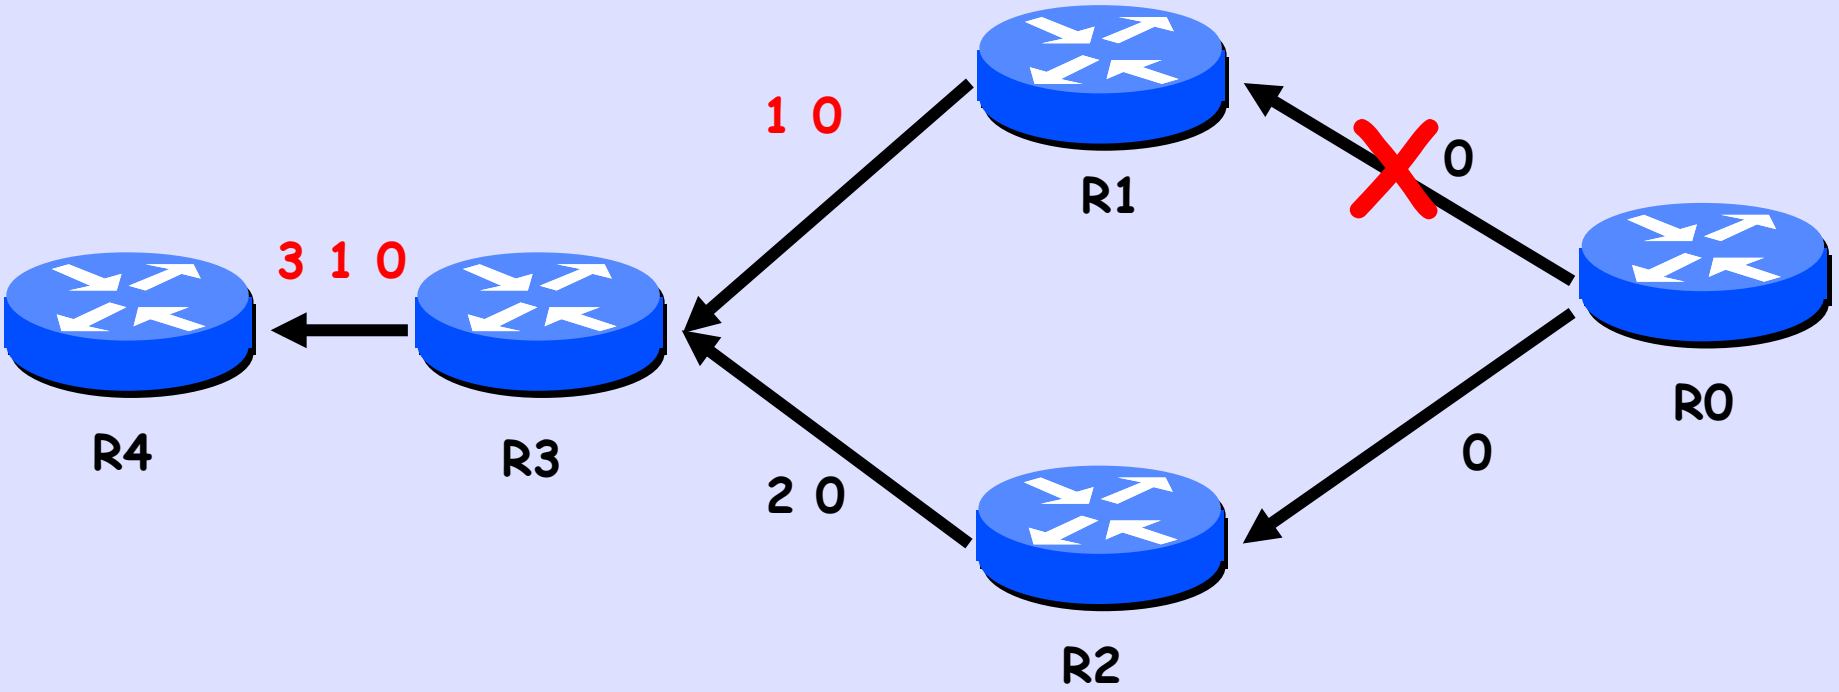

Replay Reduction

- Announcement replay is a vulnerability
- Therefore freshness is critical
- So originating announcer signs with a relatively short signature lifetime
- Origin re-announces prefix well within that lifetime, AKA ***beaconing***
- Suggested to be days, but can be hours for truly critical infrastructure

Per-Router Keys

- Needed to deal with compromise of one router exposing an AS's private key
- Implies a more complex certificate and key distribution mechanism
- A router could generate key pair and send certificate request to RPKI for signing
- Certificate, or reference to it, must be in each signed path element
- If you want one per-AS key, share a router key

Cert / Key Structure for an ISP

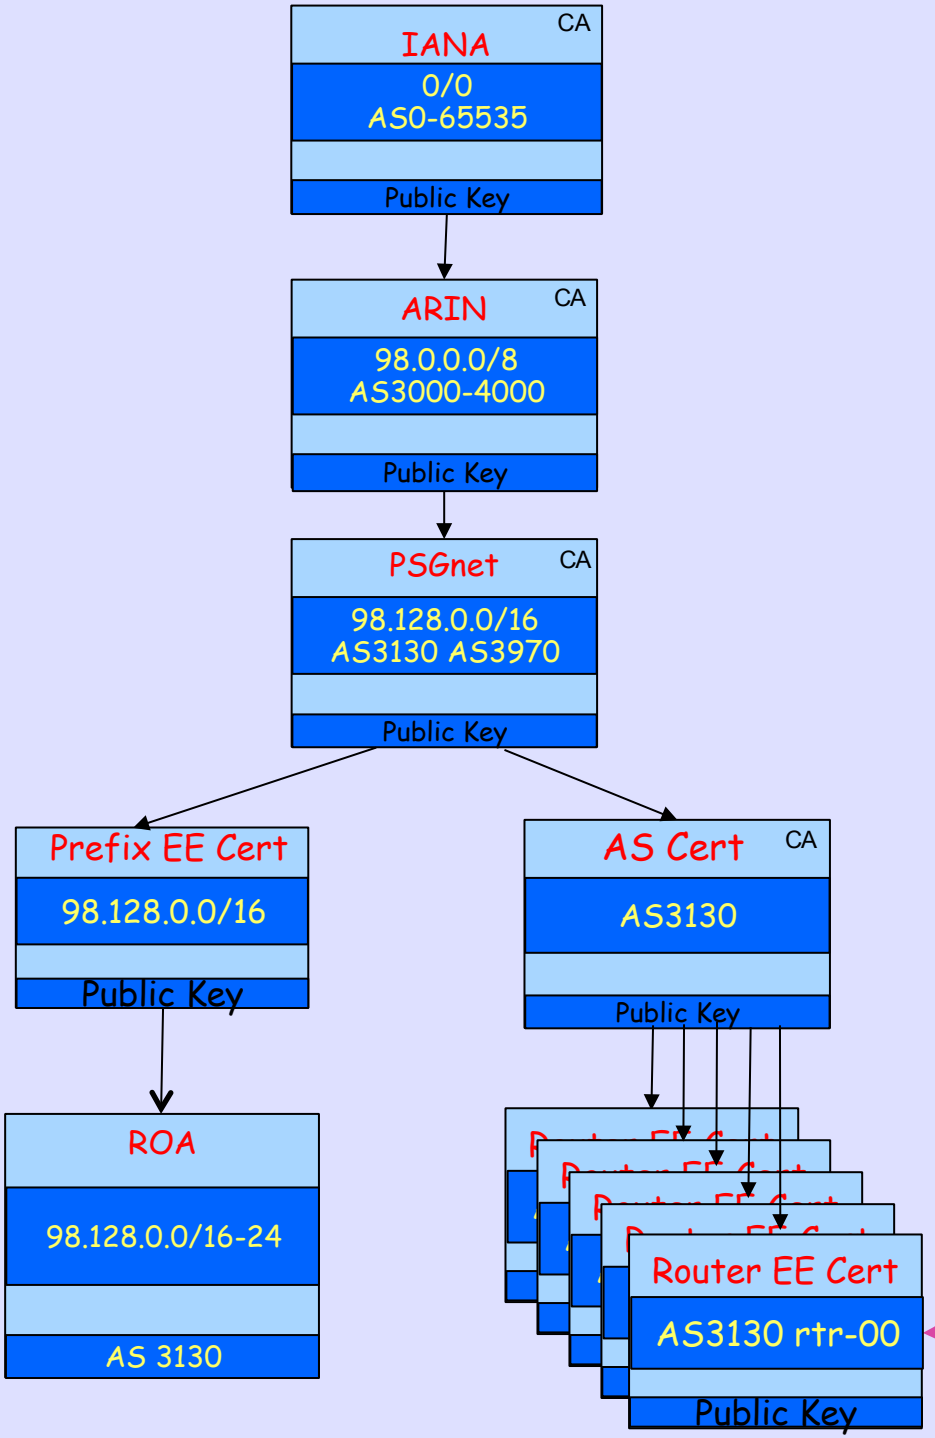

Encodes ASN and Router ID

Origination by AS0 to AS1

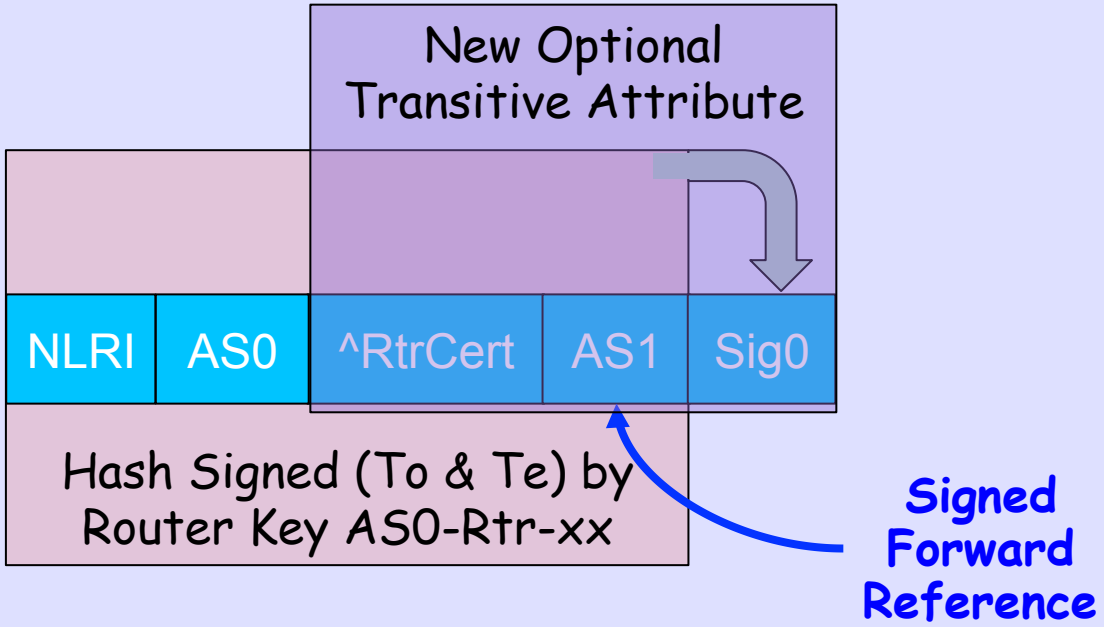

- To and Te are times of signature origination and expiration
- Signature has a well-jittered validity end time, Te, of days
- Re-announcement by origin, AKA *beaconing*, every ~(Te-To)/3
- ROA is not needed as prefix is sufficient to find it in RPKI as today

Announcement AS1 to AS2

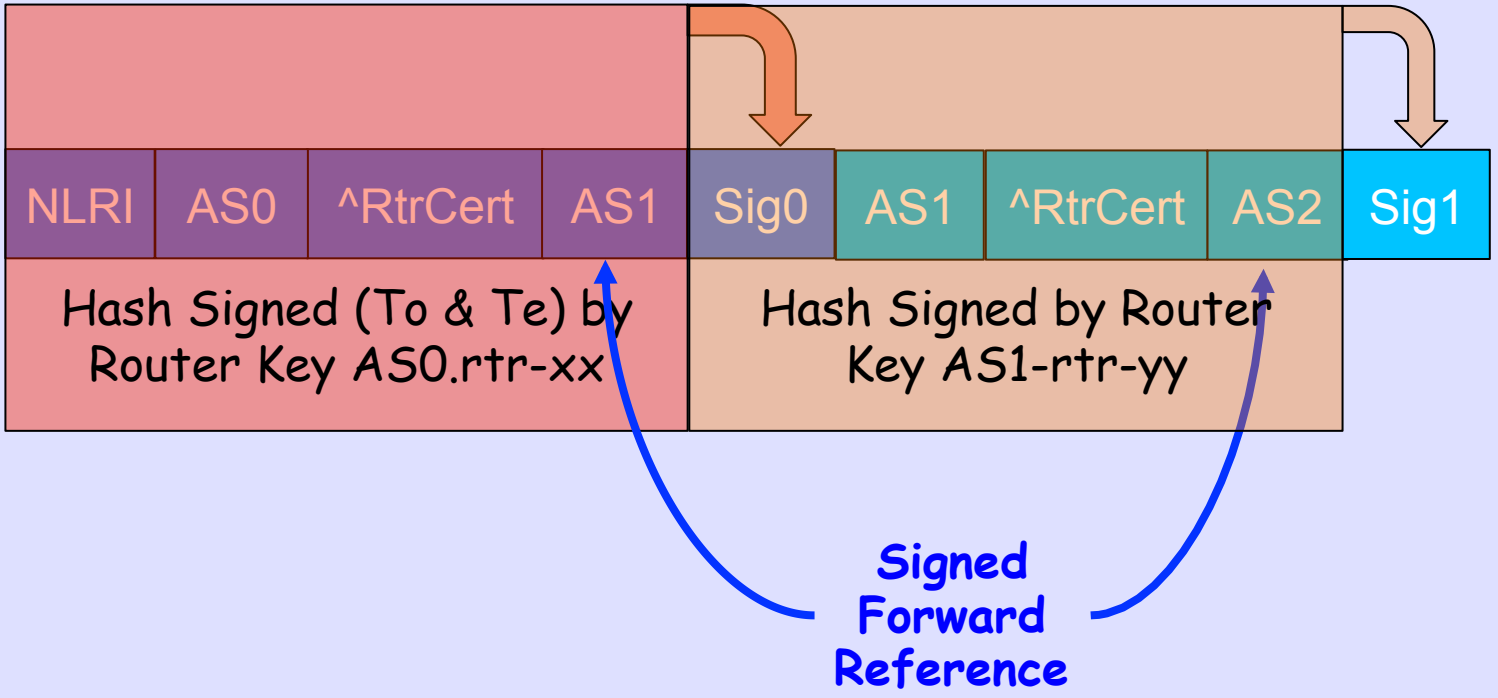

- R1 signing over R0's signature is same as signing over entire R0 announcement
- Non-originating router signatures do not have validity periods
- But when they receive a beacon announcement, they must propagate it

Only at Provider Edges

- This design protects only inter-domain routing, not IGPs, not even iBGP
- BGPsec will be used inter-provider, only at the providers' edges
- Of course, the provider's iBGP will have to carry the BGPsec information
- Providers and inter-provider peerings might be heterogeneous

Simplex End Site

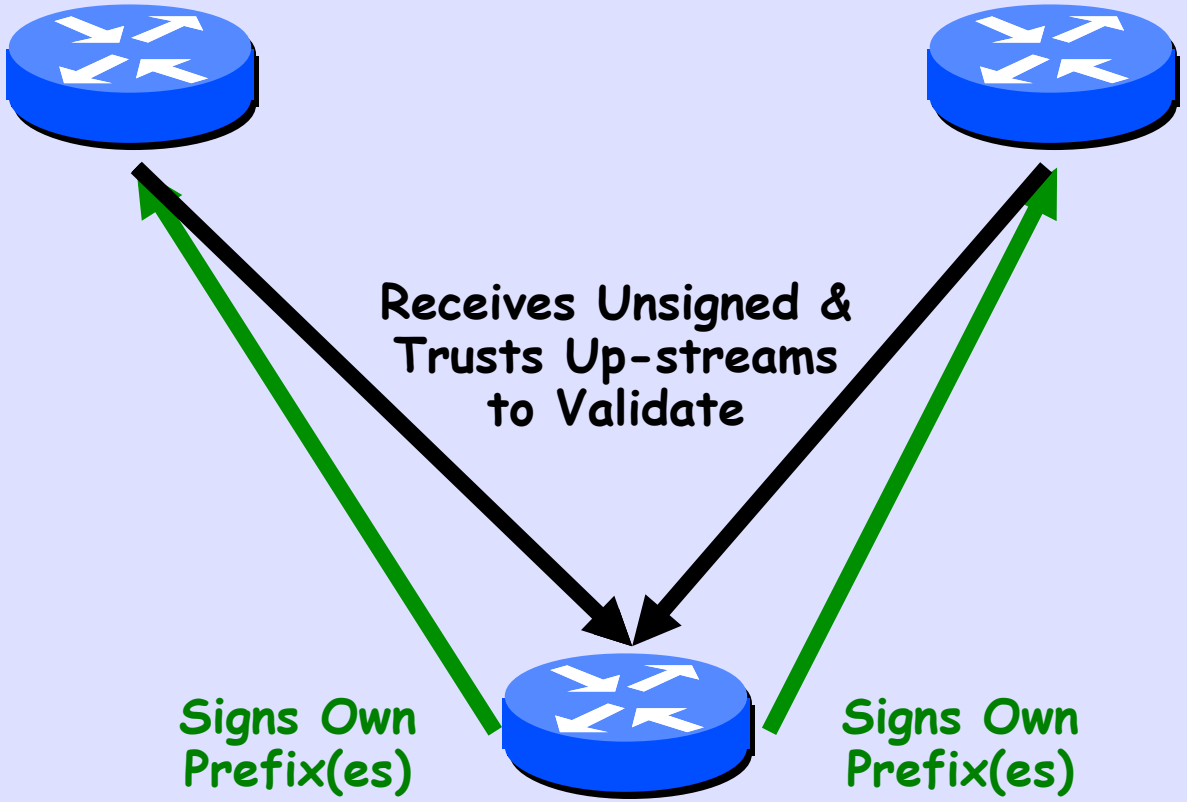

Only Needs to Have Own Private Key, No Other Crypto or RPKI Data

No Hardware Upgrade!!

Incremental Deployment

Meant to be incrementally deployable in today's Internet, and does not require global deployment, a flag day, etc.

No Increase of Operator Data Exposure

- Operators wish to minimize any increase in visibility of information about peering and customer relationships etc.
- No IRR-style publication of customer or peering relationships is needed

iBGP & Confederations

- iBGPsec speakers who are also eBGPsec speakers naturally carry BGPsec data
- Route Reflectors must be non-signing iBGPsec speakers
- Confederations are eBGP boundaries, but a subAS should not sign to coreAS as that signature would have to be removed.

Only Prefix & AS-Path

- Until clear vulnerabilities demonstrate a need for more, only the prefix and the AS path are covered by the signature.
  - Other attributes are too variable, are ephemeral, or we do not understand the security needs.
  - I.e. don't sign what we don't understand.
- NO-EXPORT etc. are over a [secured] next-hop, and thus do not need signing.]

Utterly Un-Optimized

- This design very intentionally abjures premature, in fact any, optimization in an attempt to get the semantics of the protocol correct in a simple and understandable way
- It is assumed that optimization, prepends, packing, etc. will be worked out as the design is finalized

An Example Optimization

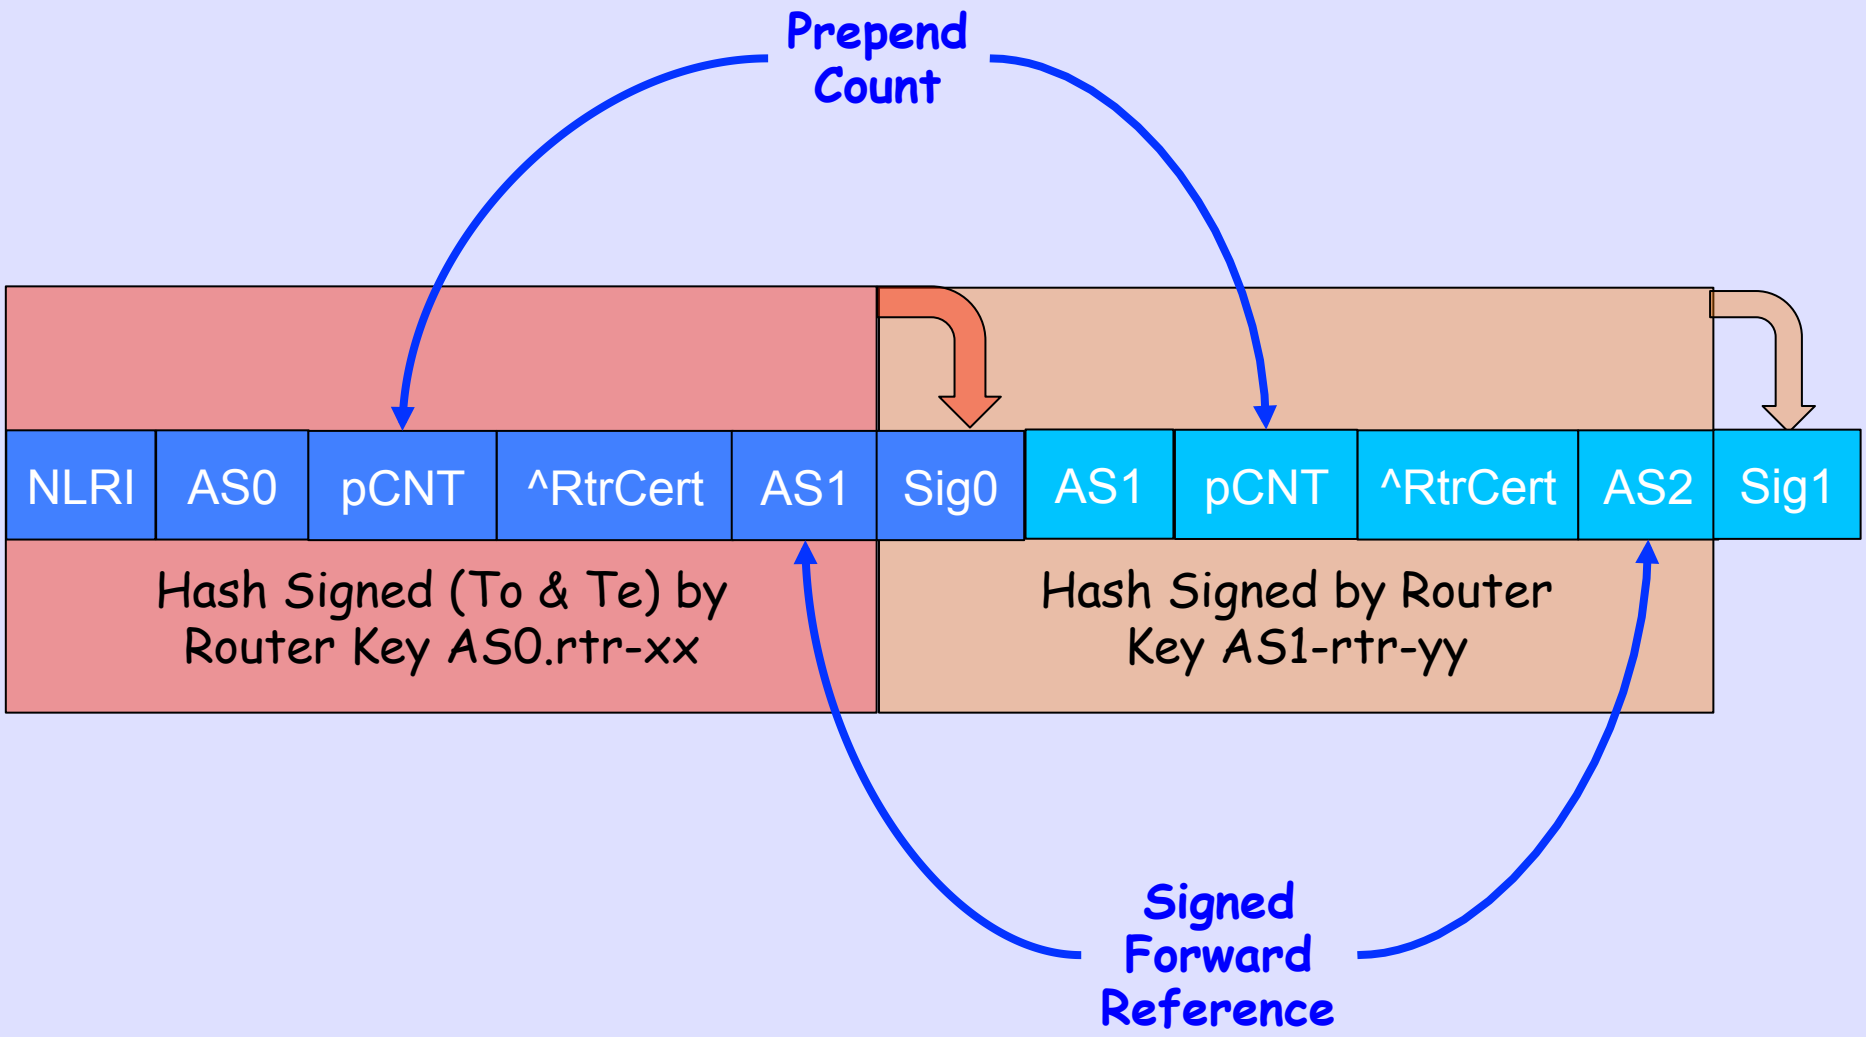

Uses Global RPKI

It is assumed that any needed global RPKI data can be delivered to routers (or ancillary devices) by augmenting the RPKI to Router Protocol described in draft-ietf-sidr-rpki-rtr-protocol, with the additional PDUs necessary to transport certificates, CRLs, etc.

Origin Validation Assumed

- We assume that prefix origin validation can be and/or is already being done by routers using ROAs from the RPKI
- We can leverage the ROA being in the router's prefix trie already, so need not include it in signed updates

Just Another BGP Decision

- The result of validation is similar to any other BGP decision
- Local policy decides what to do with the result of validity testing, a la origin validation
- And the vendors will give the ops too many knobs

Some Consequences

New Hardware Generation

It is likely that routers will have to be upgraded to use this design, likely with much more memory and probably with hardware crypto assistance. It is accepted that this means that it will be some years, O(IPv6 ASIC upgrades) before there is more than test deployment

RIB Size Estimation – – Relative Measure of Contributions due to IGP / eBGP Prefixes

- Estimate for Route Reflector (prefix path multiplier factor = 9.55)
- Contribution to RIB size due to internal prefixes (unsigned) is very small
- Signed eBGP prefixes dominate

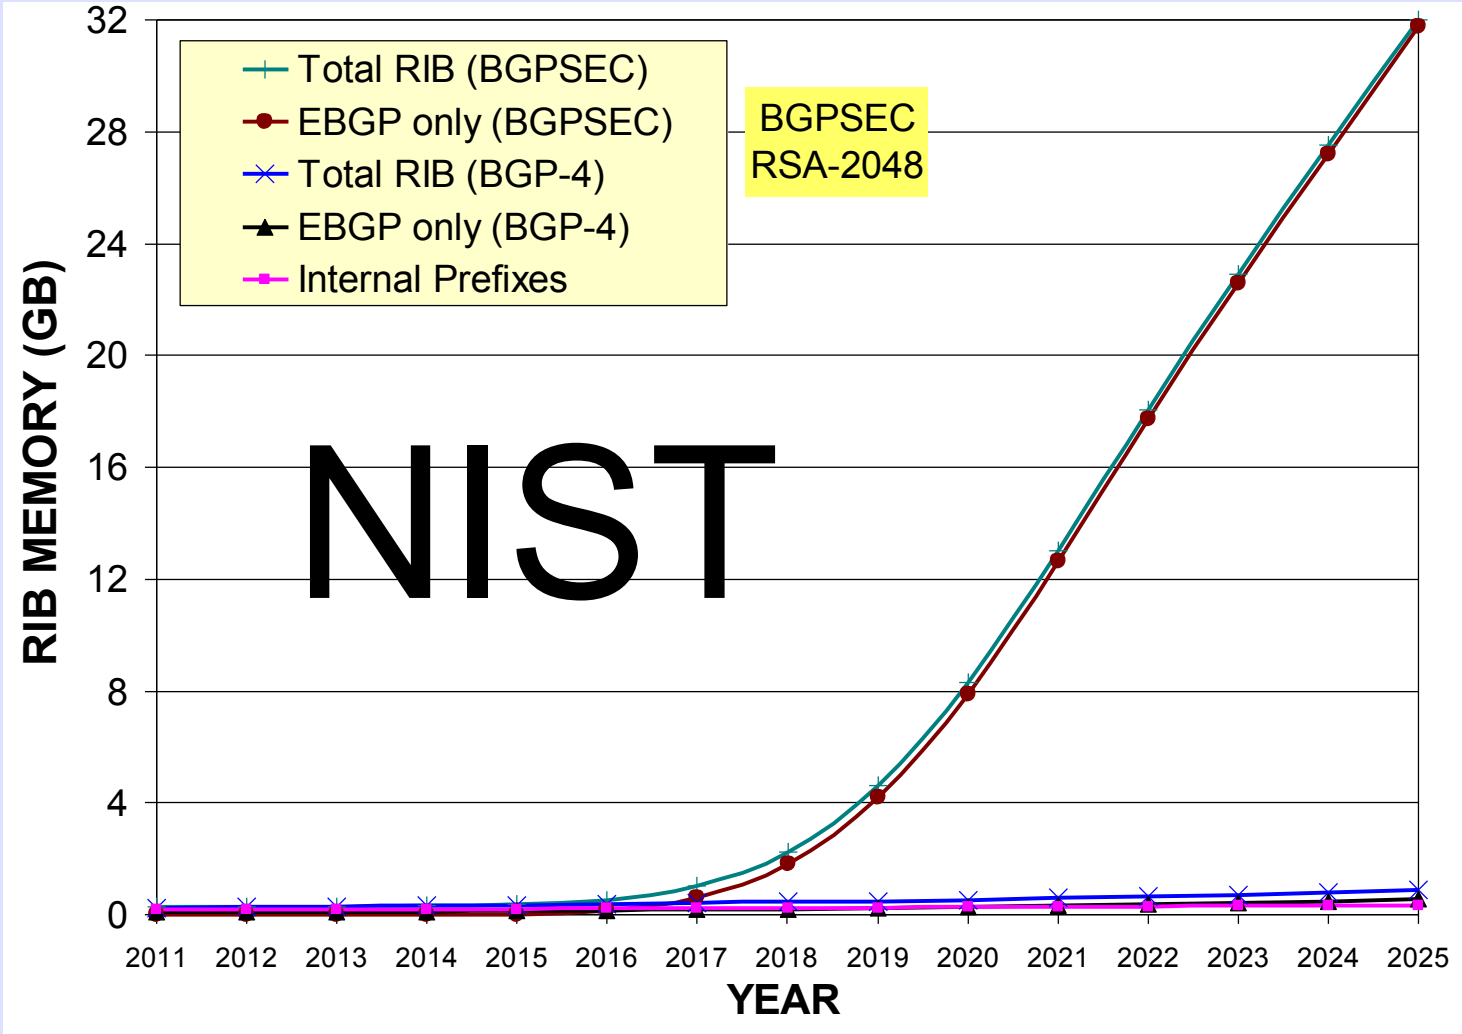

No PDU Packing

- This 'idealized' protocol has only one prefix in each announcement PDU
- Routers currently unpack prefixes from PDUs, and subsequent re-announcement repacks and reorders rather arbitrarily
- PDU optimization can be studied after the protocol semantics are solid

Route Servers

BGPsec can't forward sign across an AS-transparent route server as you do not know the peer AS

Proxy Aggregation

Proxy Aggregation, i.e.
AS-Sets, is not supported

Does Not Lock Data Plane

- It is acknowledged that rigorous control plane verification does not in any way guarantee that packets follow the control plane
- See IMC 2009 paper which shows that 70% of the ASs in the so-called 'default free zone' also have default

THIS WORK IS SPONSORED IN PART BY THE DEPARTMENT OF HOMELAND SECURITY UNDER AN INTERAGENCY AGREEMENT WITH THE AIR FORCE RESEARCH LABORATORY (AFRL).

They Take your Scissors Away and we turn them into plowshares