Idealized BGPsec: Formally Verifiable BGP

JaNOG 27.5 / Tokyo

2011.04.14

Randy Bush <randy@psg.com>

for the

Informal BGPsec Design Group

Informal BGPsec Group

chris morrow (google)
dave ward (juniper)
doug maugham (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 (dt)
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)

Assume RPKI

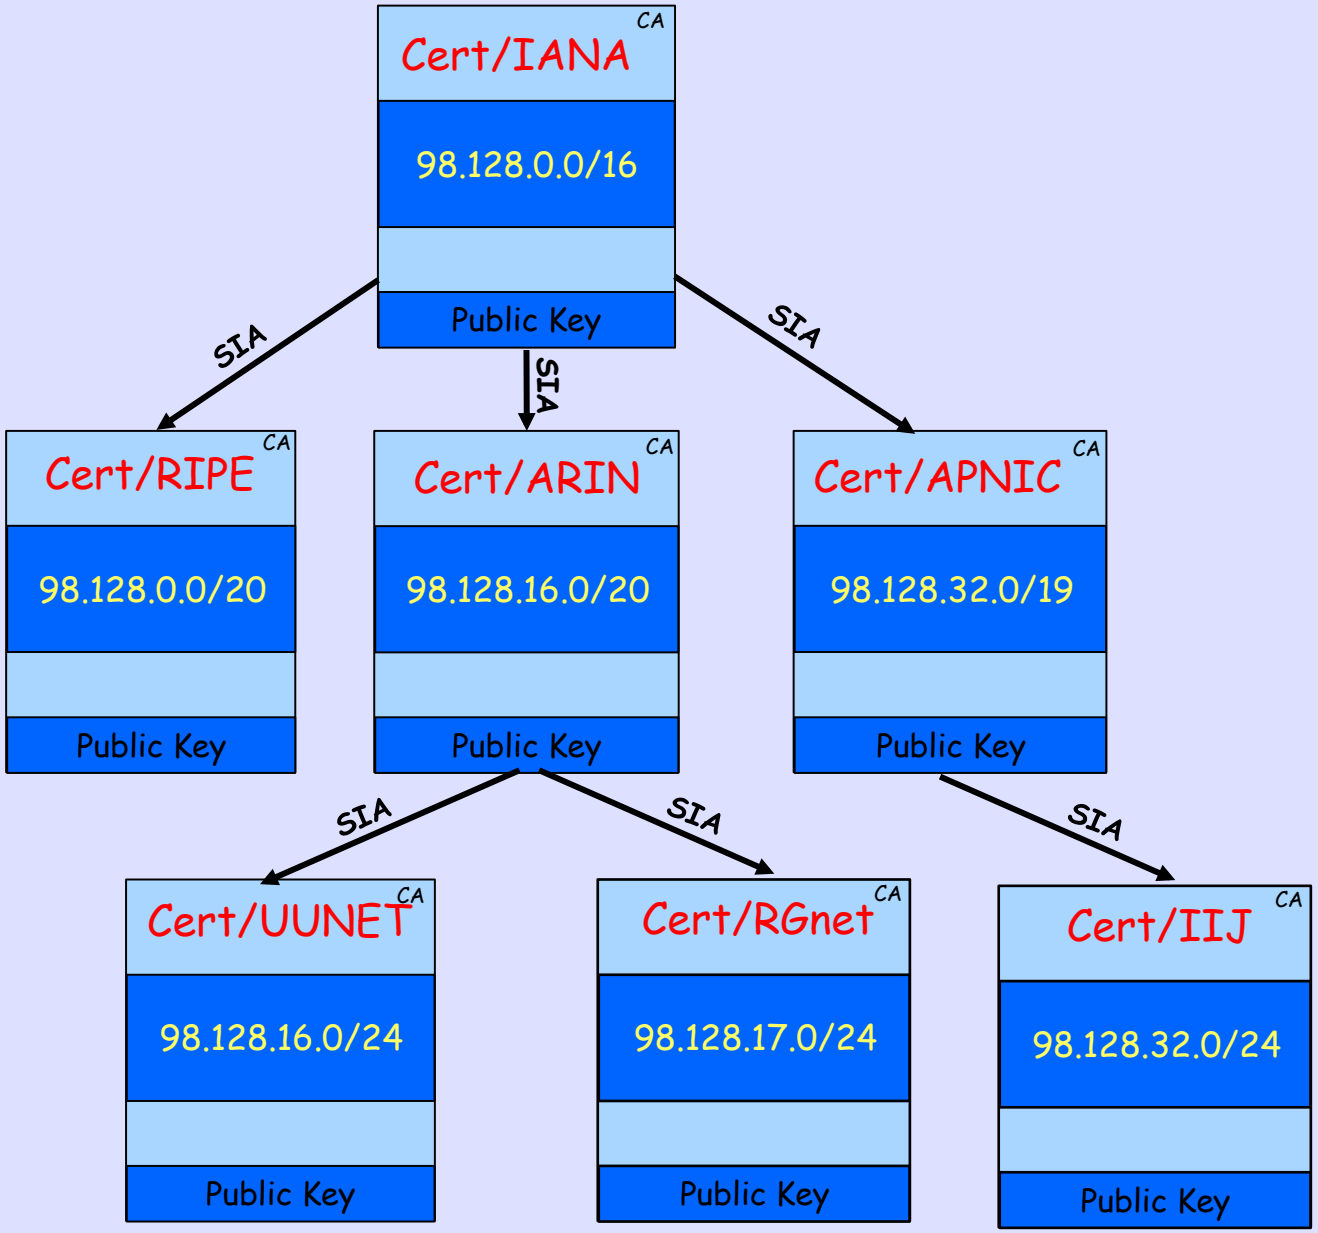
Cert/IANA
CA
98.128.0.0/16
Public Key
SIA
SIA
SIA
Cert/RIPE
CA
98.128.0.0/20
Public Key
Cert/ARIN
CA
98.128.16.0/20
Public Key
Cert/APNIC
CA
98.128.32.0/19
Public Key
SIA
SIA
SIA
Cert/UUNET
CA
98.128.16.0/24
Public Key
Cert/RGnet
CA
98.128.17.0/24
Public Key
Cert/IIJ
CA
98.128.32.0/24
Public Key

Assume RPKI-RTR

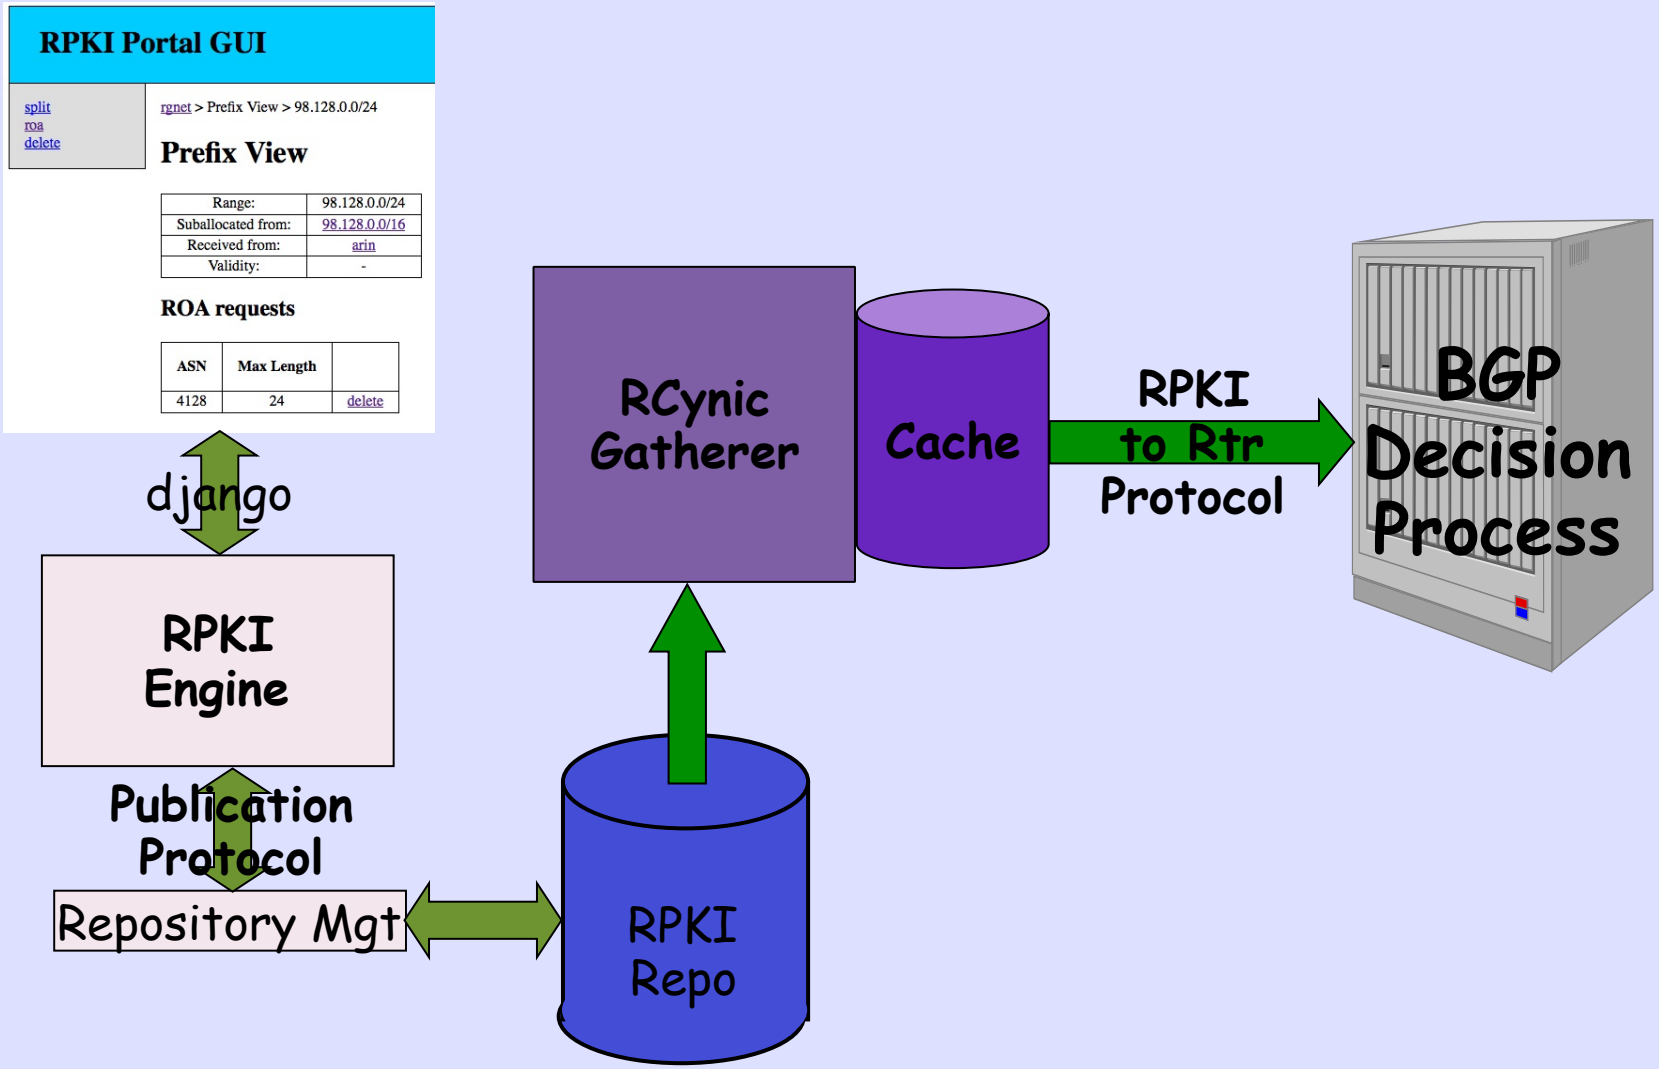

Assume Origin Validation

```
R3#sh ip bg 98.128.0.0/24
BGP routing table entry for 98.128.0.0/24, version 94
Paths: (2 available, best #2, table default)
 65000 3130
    10.0.0.1 from 10.0.0.1 (65.38.193.12)
      Origin IGP, localpref 100, valid, external
      path 6802D4DC RPKI State invalid
  65001 4128
    10.0.1.1 from 10.0.1.1 (65.38.193.13)
      Origin IGP, localpref 100, valid, external, best
      path 6802D7C8 RPKI State valid
```

Origin Validation is Weak

- Today's Origin Validation provides neither cryptographic assurance (announcements are not signed), nor assurance of the AS Path of the announcement.
- A malicious router may announce as any AS, i.e. forge the ROAed origin AS.
- This would pass ROA Validation

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

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

Path Shortening Attack

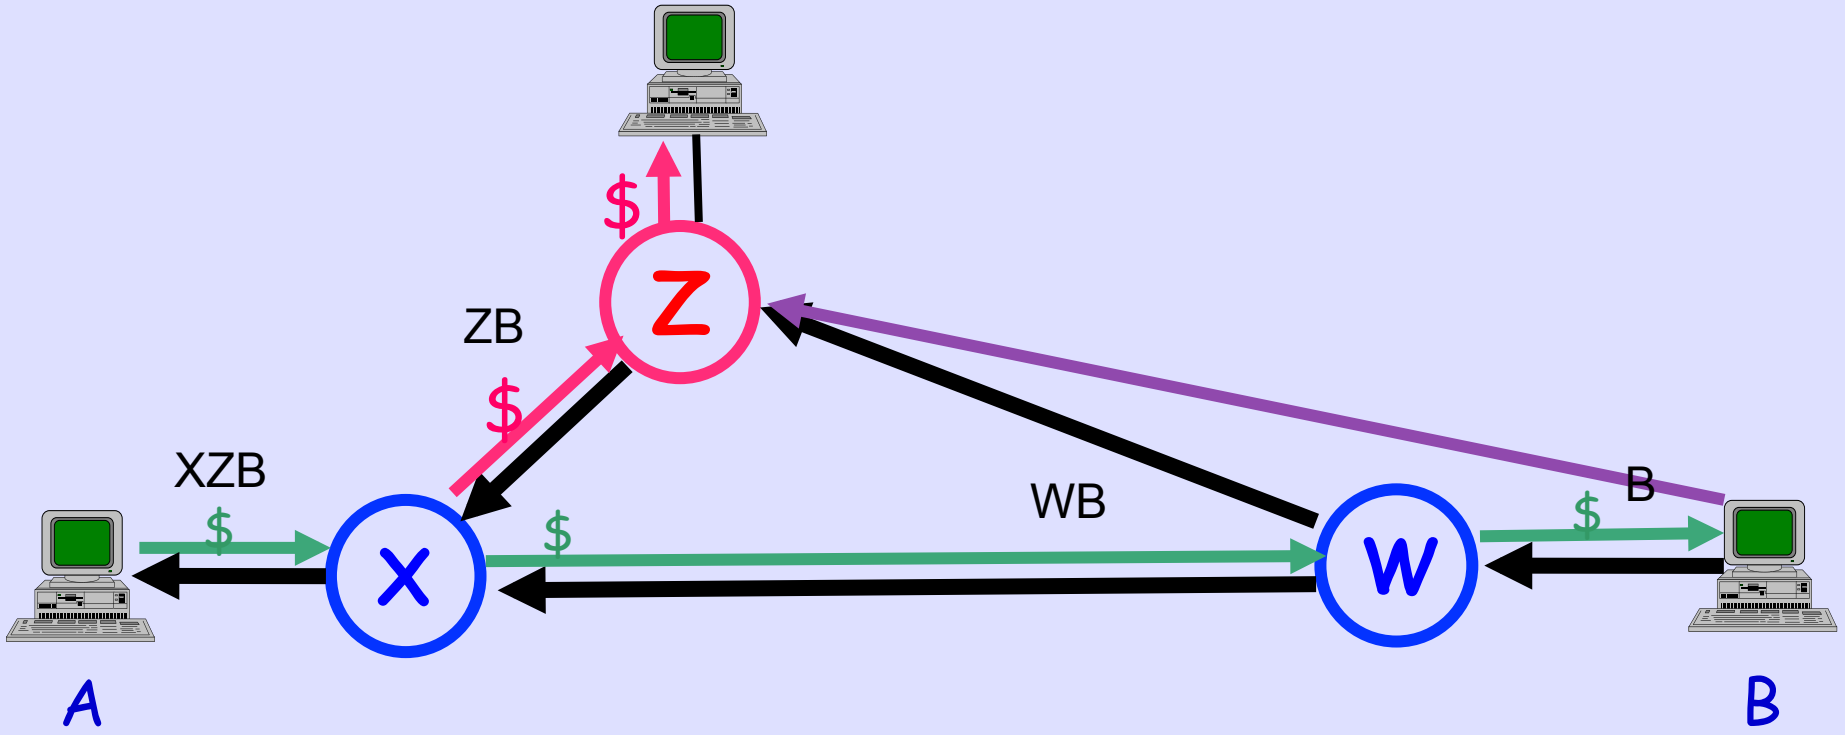

Expected Path - A->X->W->B

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

There Are Many Many Other Attacks

Forward Path Signing

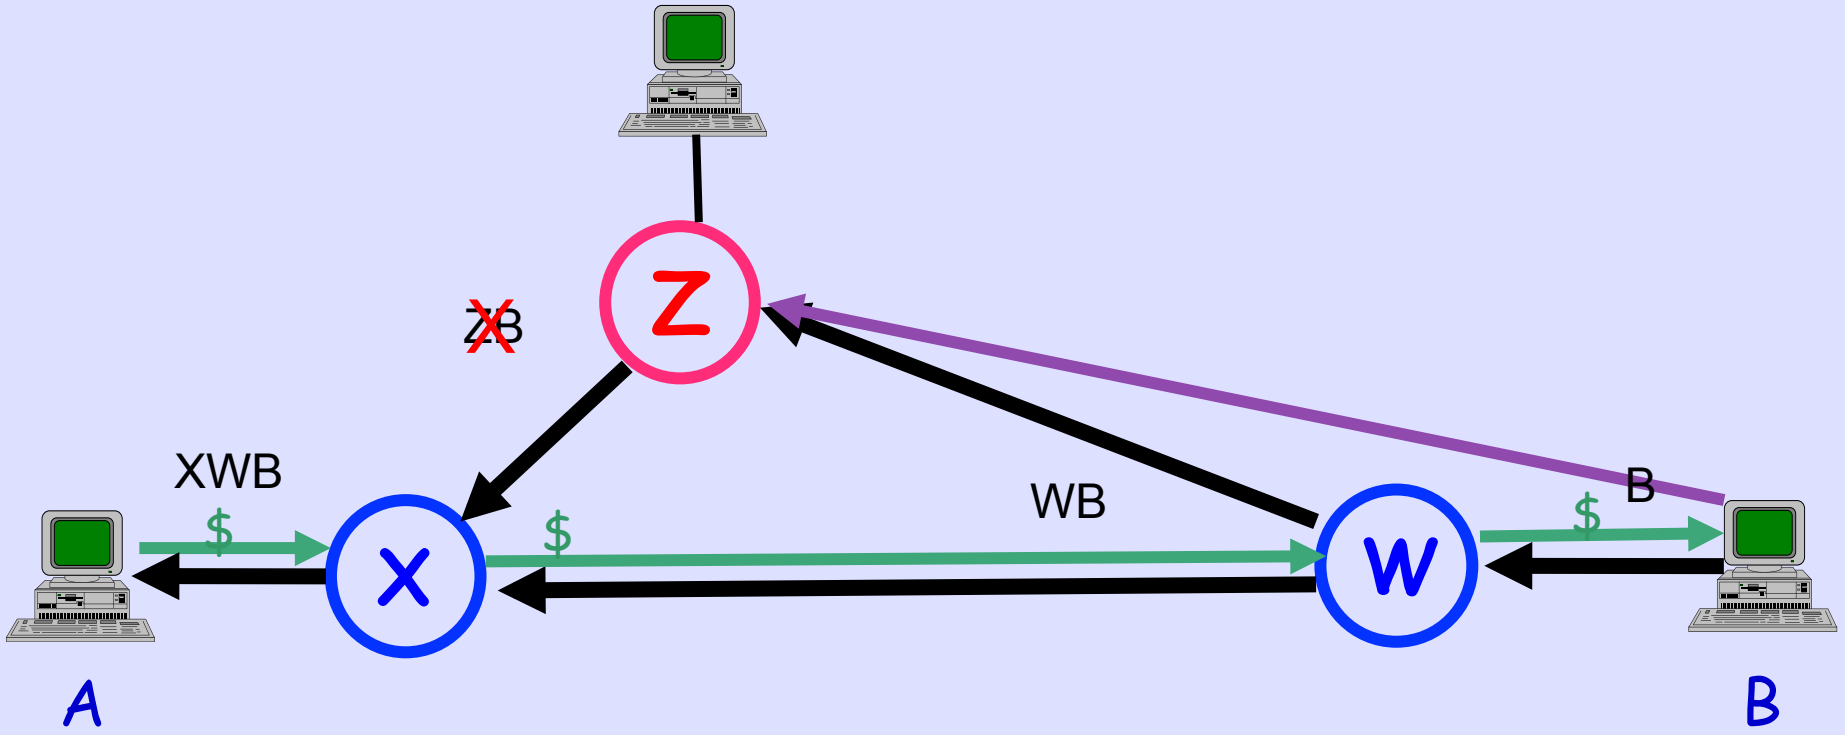

B cryptographically signs the message to W Sb(B->W)

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

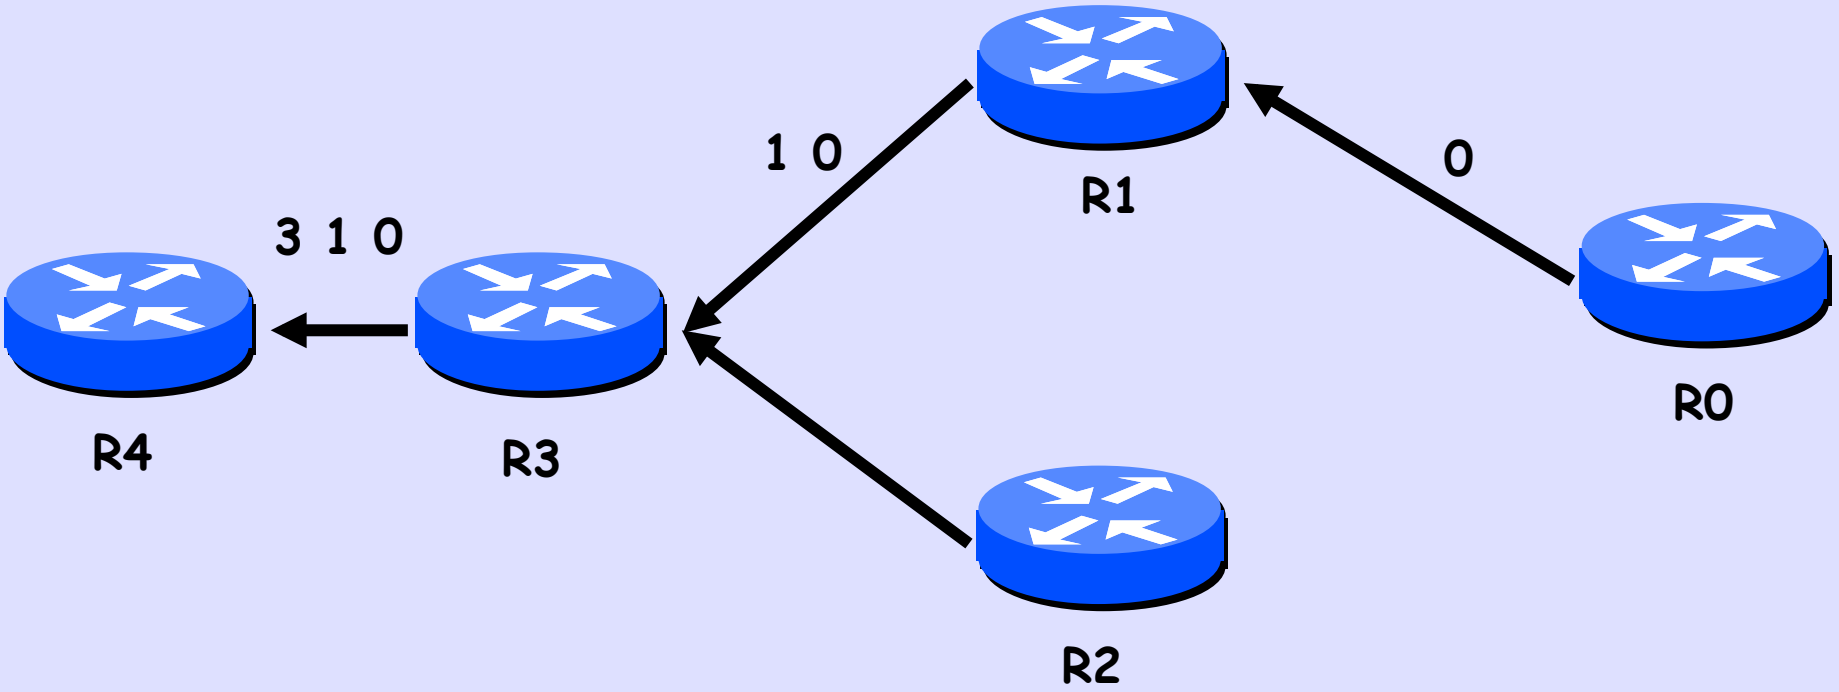

Replay Attack

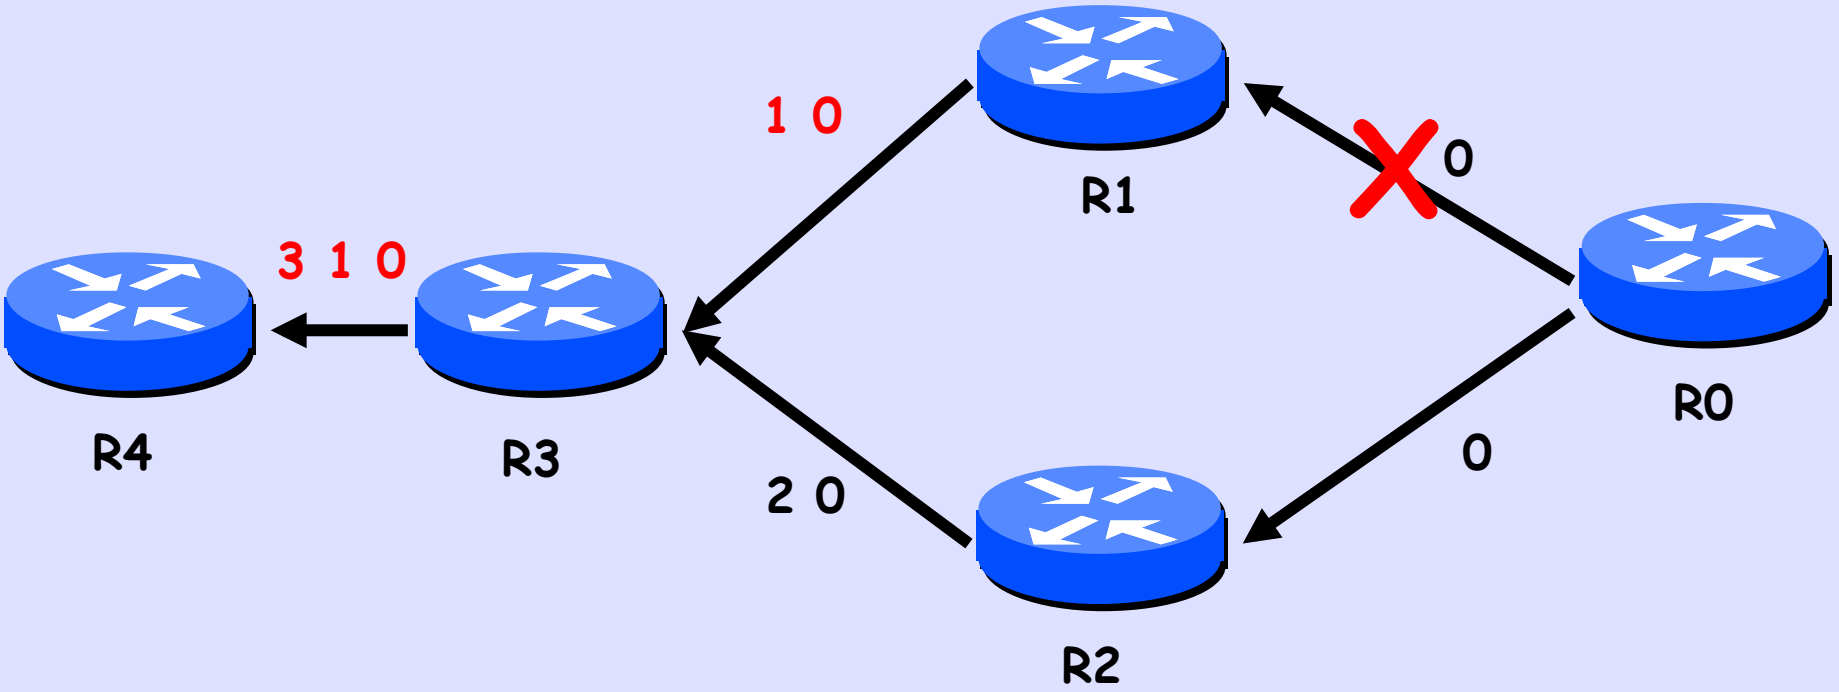

Replay Reduction

- 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

Origination by AS0 to AS1

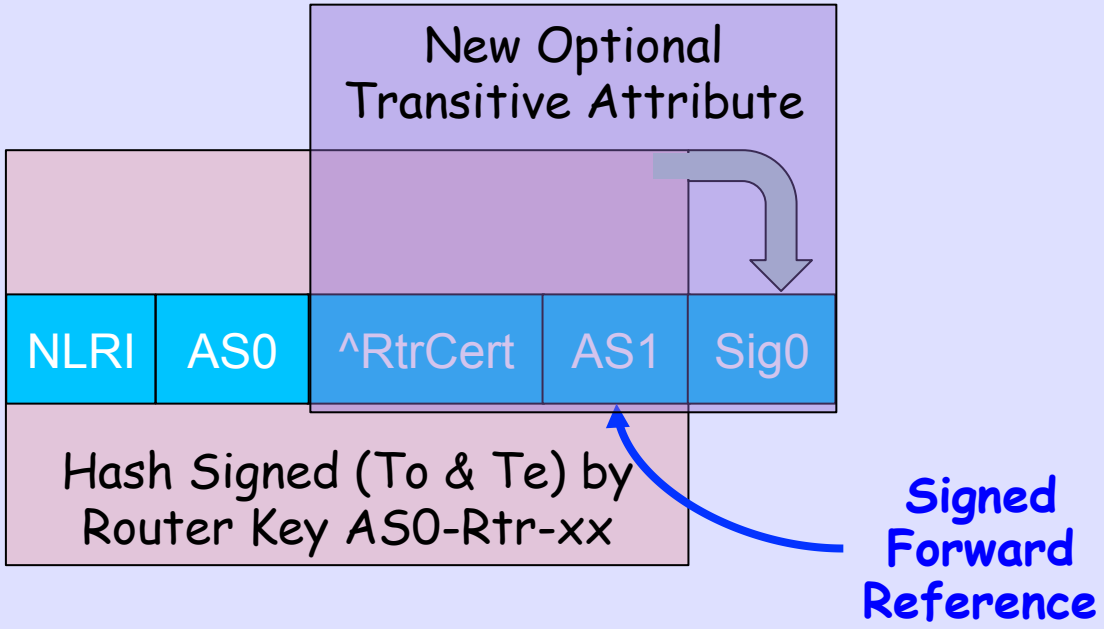

- To and Te are times of signature origination and expiration
- Signature has a well-jittered validity end time, Te, of days
- Re-announcement by origin, AKA *beaconing*, every ~(Te-To)/3
- ROA is not needed as prefix is sufficient to find it in RPKI as today

Announcement AS1 to AS2

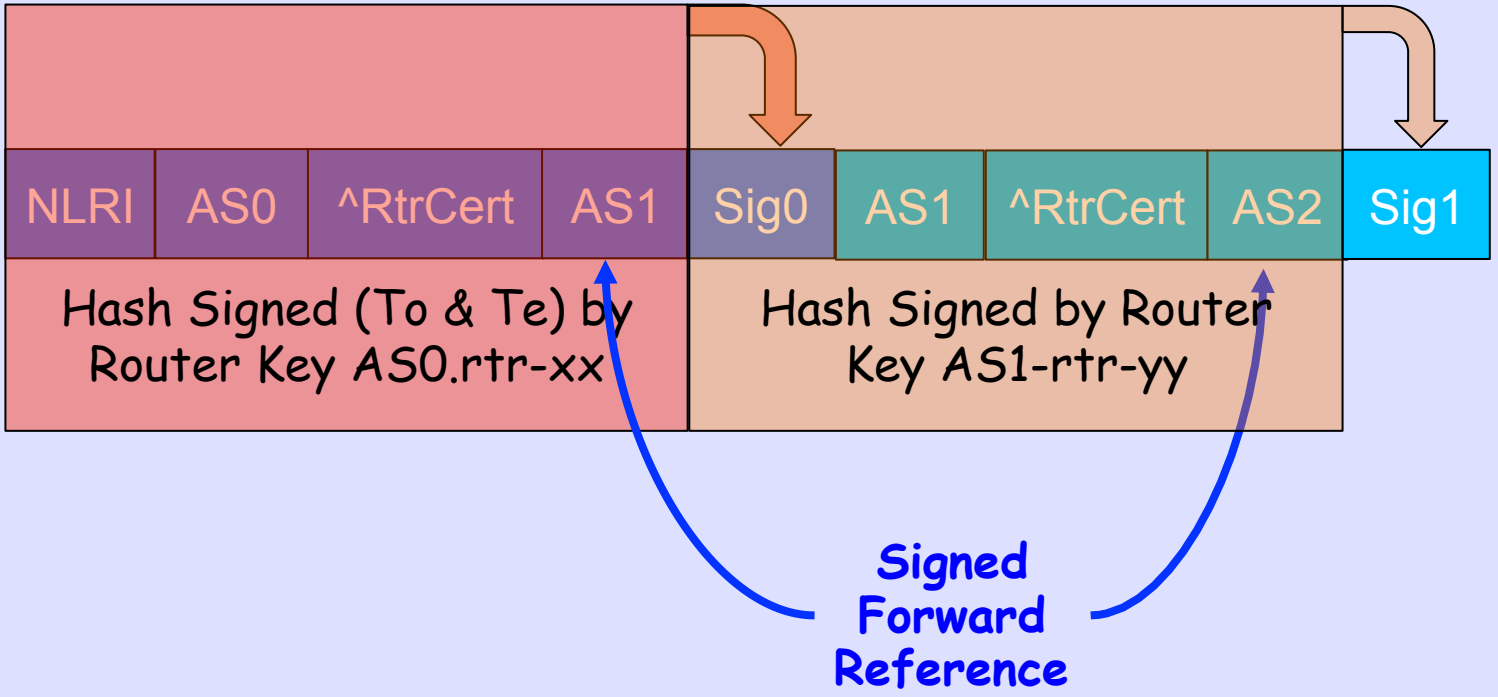

- R1 signing over R0's signature is same as signing over entire R0 announcement
- Non-originating router signatures do not have validity periods
- But when they receive a beacon announcement, they must propagate it

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

- End Site can trust up-stream's policies
- But they want their origination signed
- So they announce capability to send but do not accept signed data
- They sign announcement and beacon
- This can be done without hardware upgrade!!

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

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).

we take your scissors away and turn them into plowshares