Estimating CPU Cost of BGPSEC on a Router

IETF 82 SIDR WG Meeting

Taipei

November 2011

Kotikalapudi Sriram <ksriram@nist.gov>

Randy Bush <randy@psg.com>

BGPSEC from AS0 to AS1

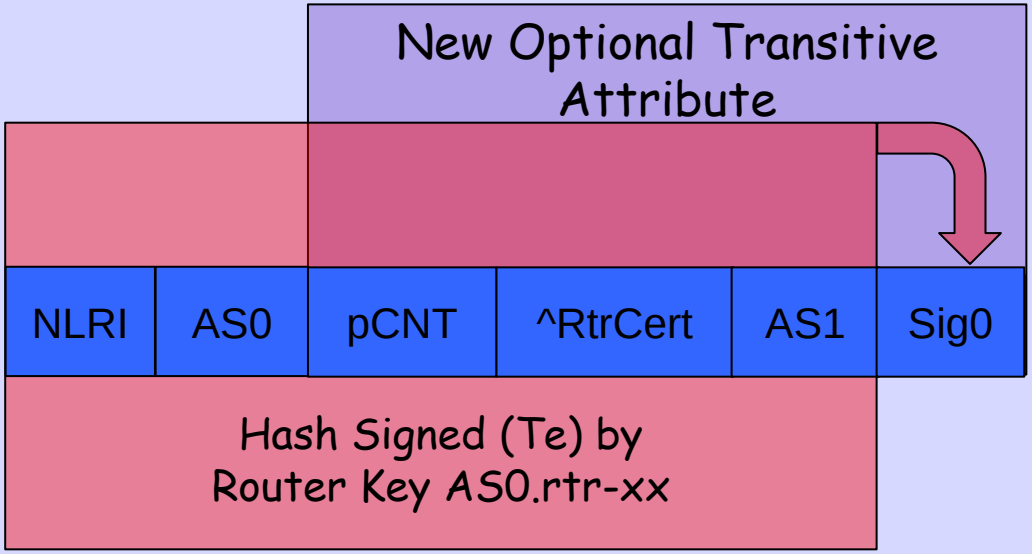

BGPSEC AS1 to AS2

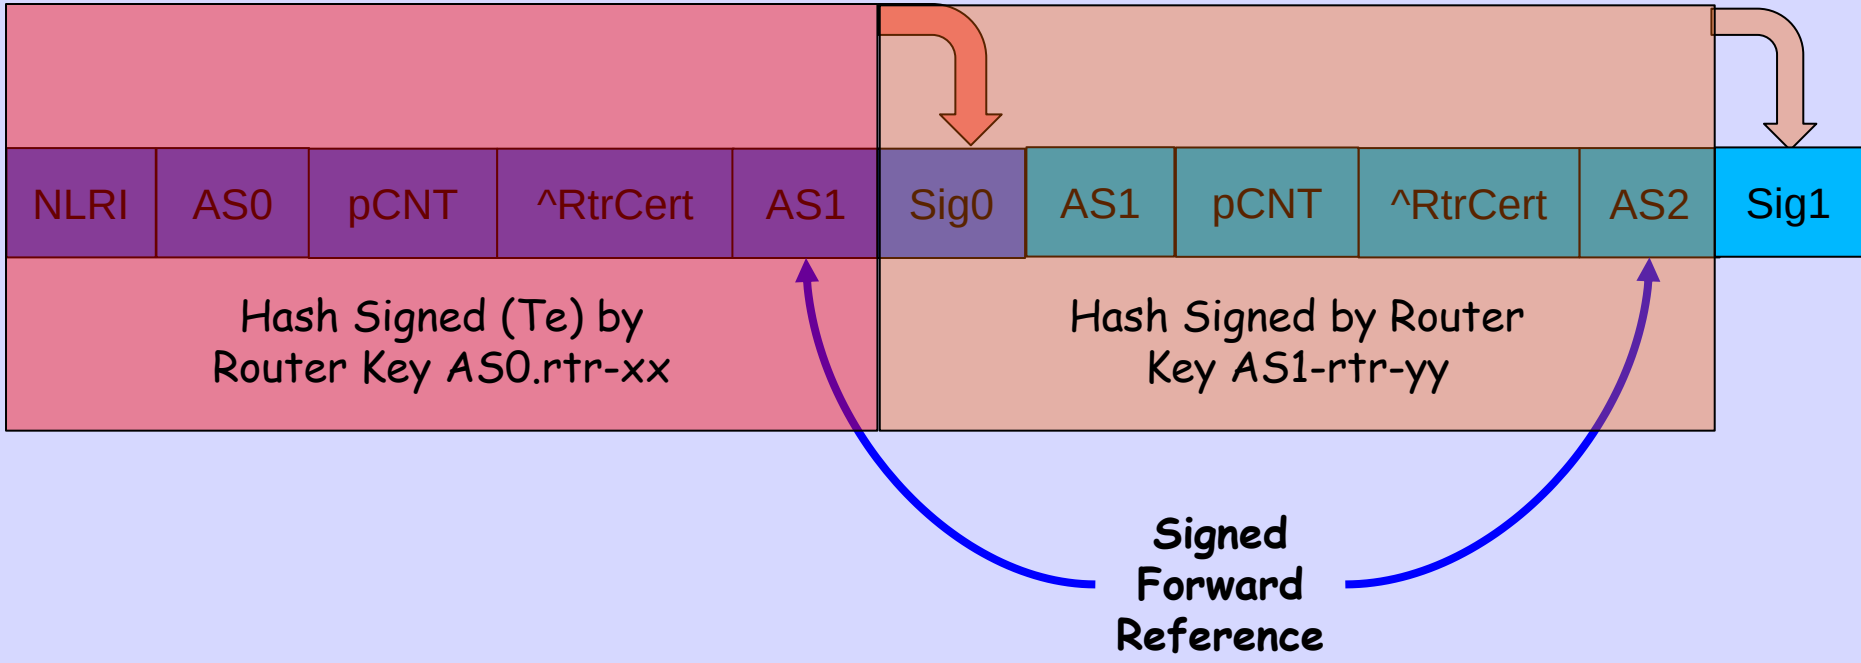

R1 signing over R0's signature is same as signing over entire R0 announcement

BGPSEC Islands

- **RPKI-Based Origin Validation** can be deployed by randomly scattered ISPs
- Each gets the benefit of origin validation
- **BGPSEC** depends on your neighbor signing
- It will deploy as islands which eventually interconnect

BGPSEC – Conceptually

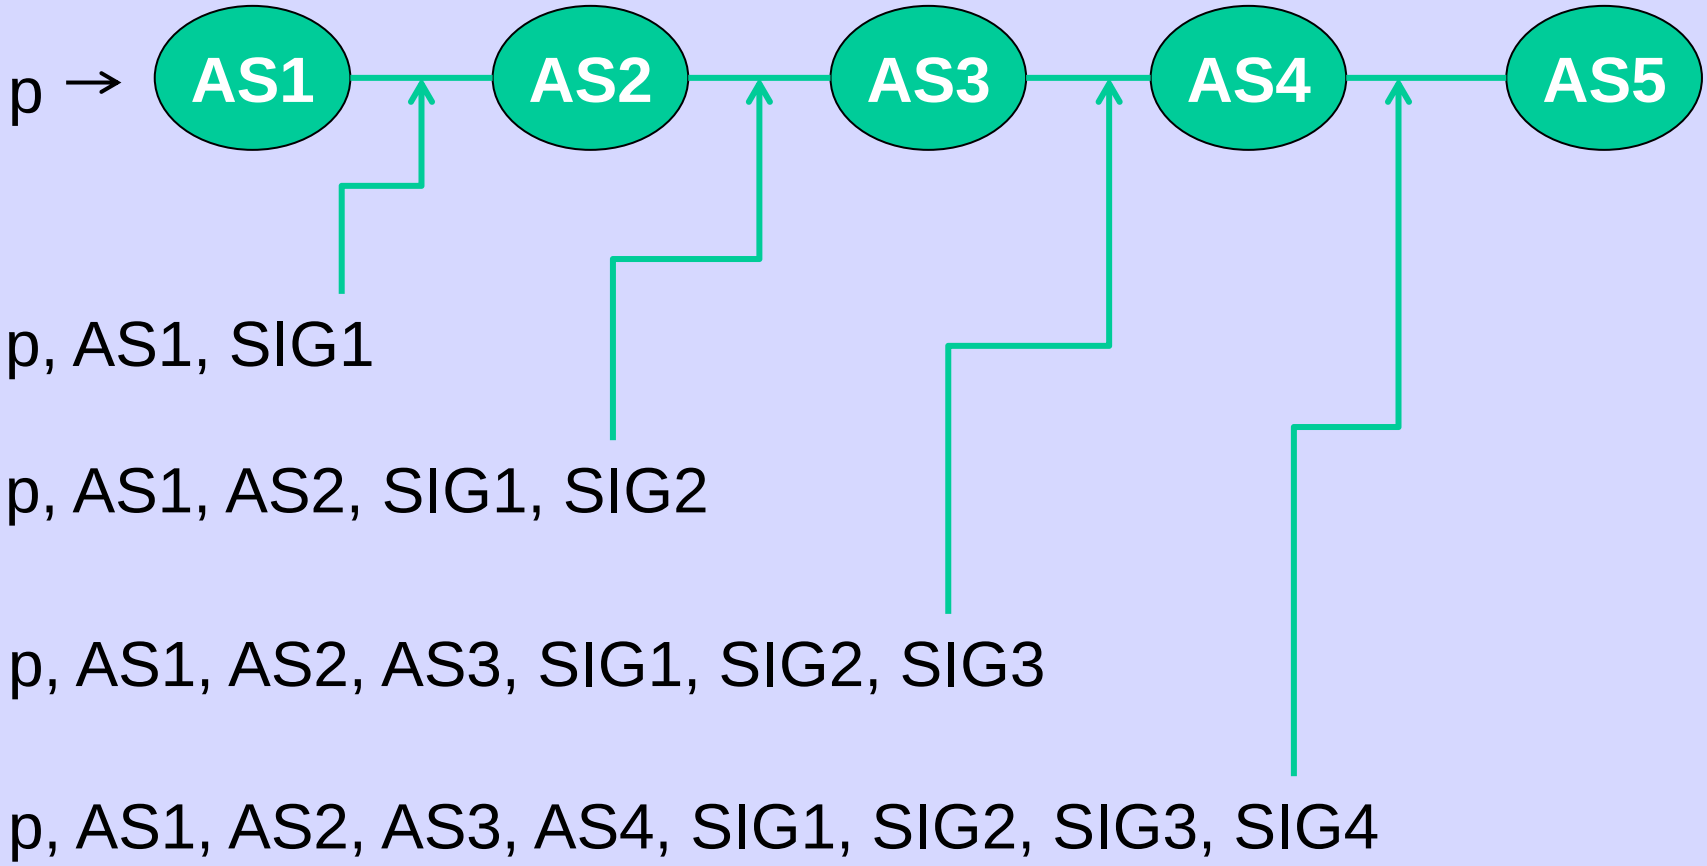

But Reality is This

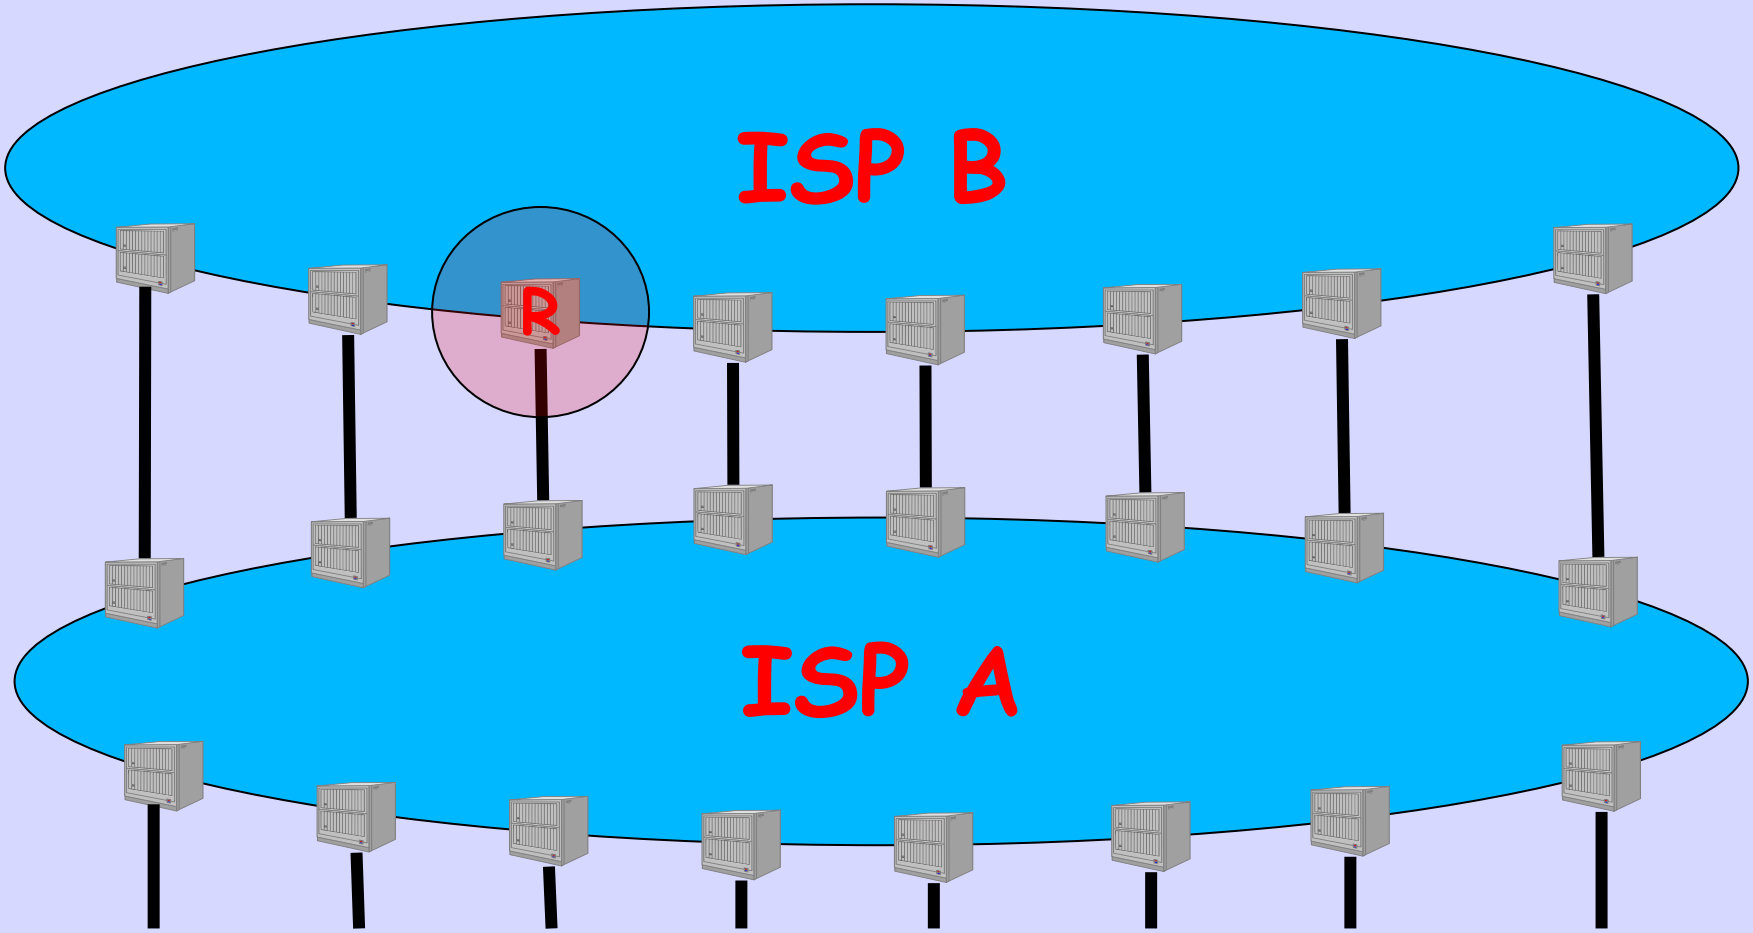

A's Customer 'Cone'

Number of Paths

- One ISP router, R, has many paths for prefix P
- All but one are from iBGP peers
- BGPSEC spec says R does not validate paths received from iBGP peers
- I.e., R has to validate only one path for each P from peer A

Some Largish ISPs Cones

Very Large Global

1	1353	--- **ISP's Own Pfx**
2	21586	--- **BGP Cust Pfx**
3	6820	--- **Cust's Cust Pfx**
4	1627	--- …
5	942	
6	45	
7	14	
8	6	

Large Global

1	443
2	8197
3	8052
4	2715
5	387
6	37
7	48
8	157
9	2

Asian Regional

1	152
2	791
3	120
4	35
5	3

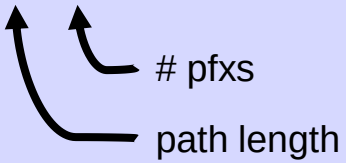

Very Large Global

1	620
2	16028
3	9434
4	2922
5	435
6	46
7	15
8	27
9	1

Large Global

1	501
2	3686
3	3603
4	816
5	45
6	9
8	1

Yes, there are rather long tails

Yes, we removed prepending

Incremental Deployment

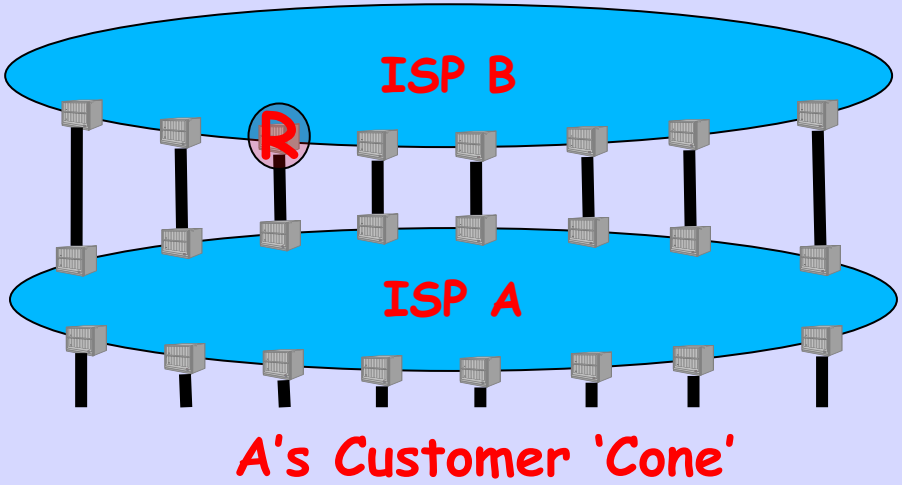

If A and B Deploy BGPSEC, What is the Load on a Router?

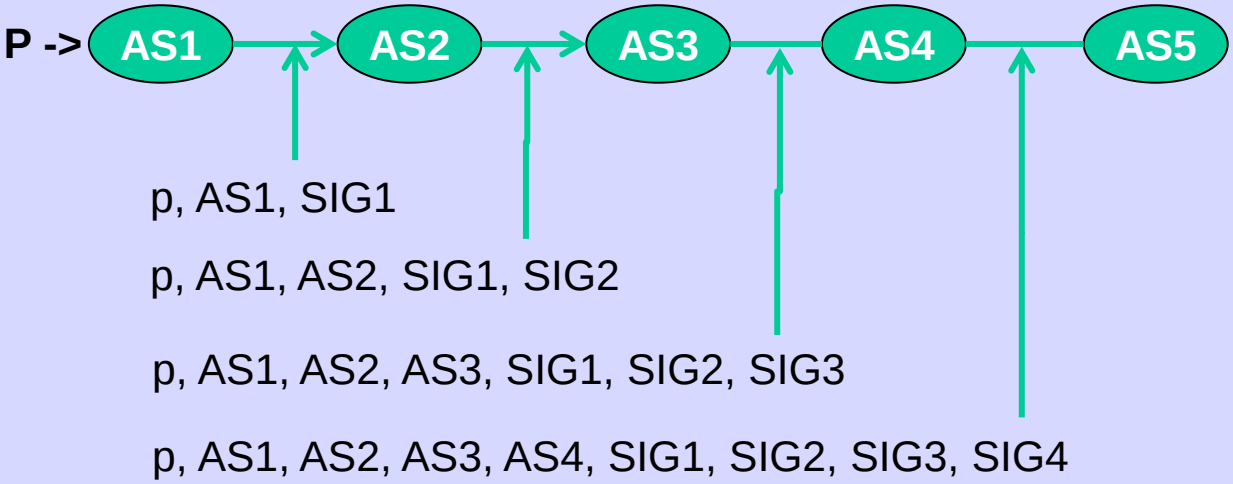

Now this Picture Makes Sense!

Cost to Sign/Validate Using One Core

	Operations per second				
	Intel Core 2 Duo, 64-bit, 3 GHz, 8GB, Linux 5.7	amd64; Westmere (206c2); 2010 Intel Xeon E5620; 4 x 2400MHz	amd64, Sandy Bridge; 2011 Intel i7-2600K; 3400MHz; threads	NITROX PX PCI-Express CN1620 - PCIe Look-aside Processor	NITROX III PCI-Express CNN3570-PCIe Look-aside Processor
ECDSA-P256 Verify	890	1139	2215	854	6832
ECDSA-P256 Sign	1100	1335	2530	3293	26344

- Source: eBACS: ECRYPT Benchmarking of Cryptographic Systems

  http://bench.cr.yp.to/results-sign.htm

- And: Cavium, Inc. (private communication)

Updates Per Second

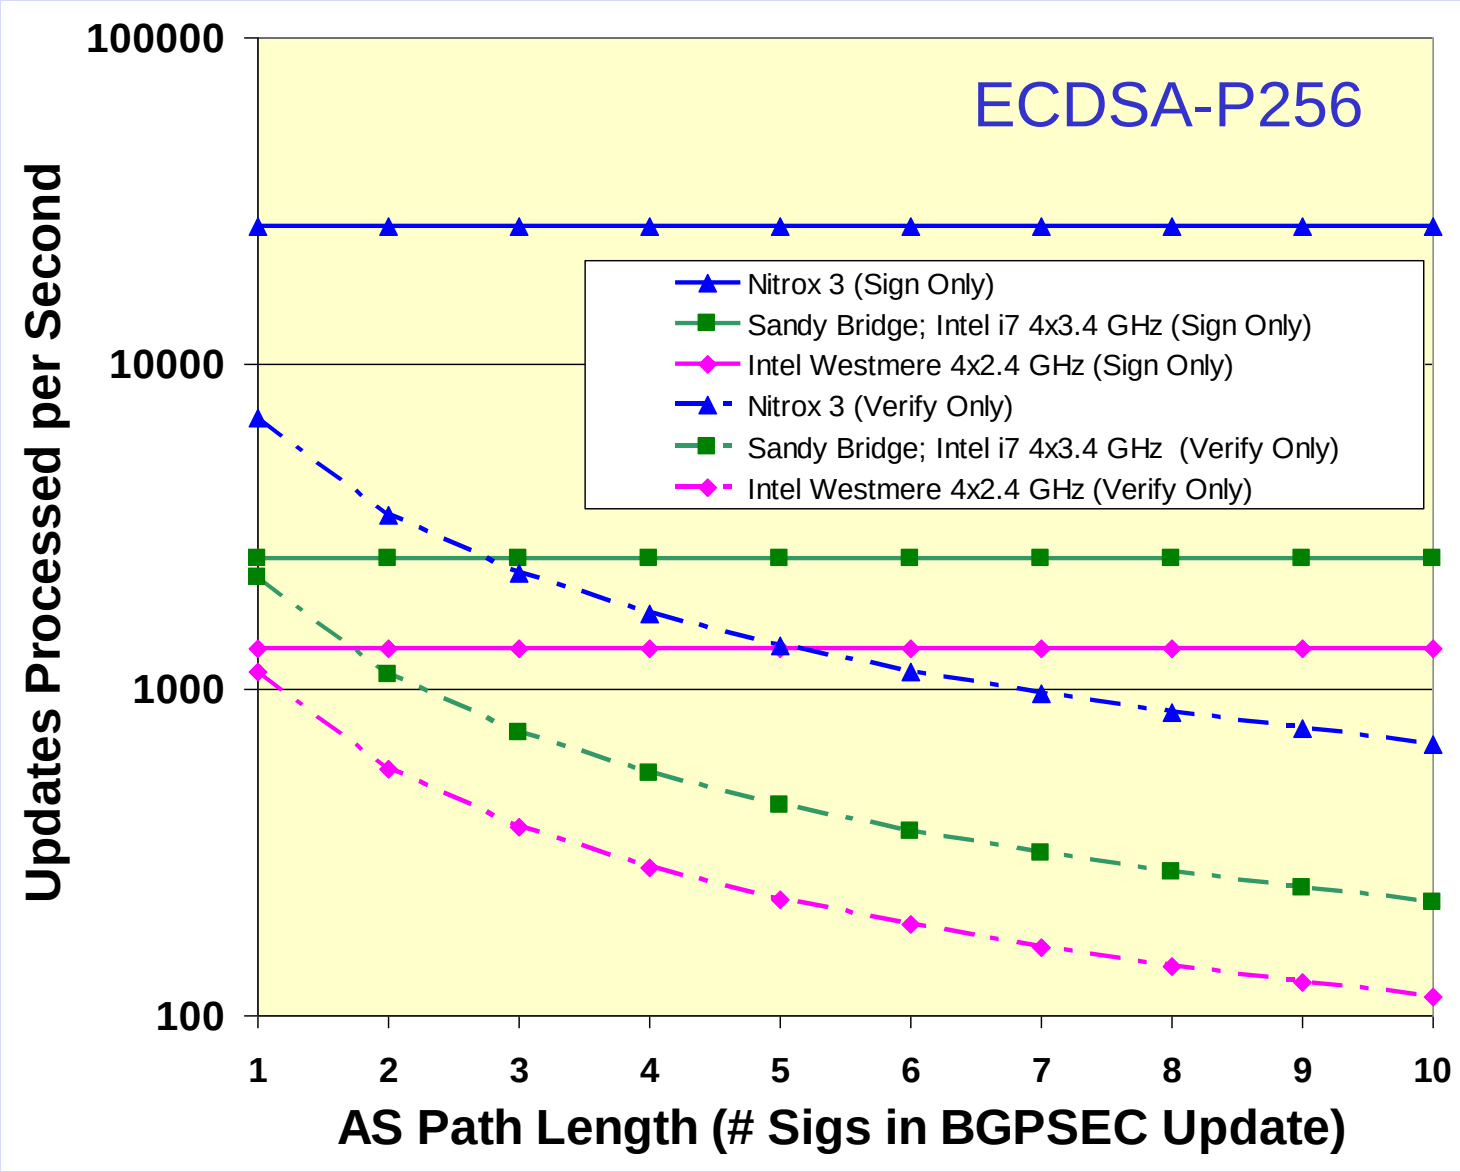

Validation Cost Model

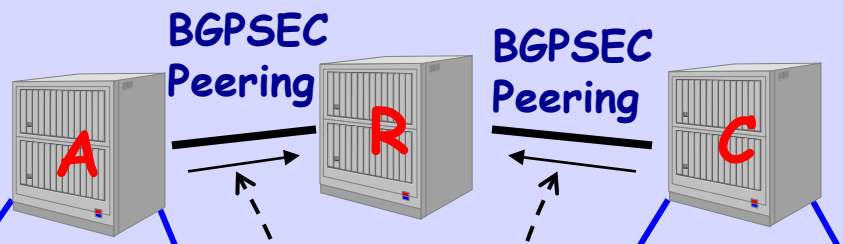

ISP A and A's Customer Cone

Path	#Pfxs	Secs
1	1353	0.61
2	21586	19.49
3	6820	9.24
4	1627	2.94
5	942	2.13
6	45	0.12
7	14	0.04
8	6	0.02
Total Seconds		34.59

CPU Time on R if Session to A is Reset

ISP C and C's Customer Cone

Path	#Pfxs	Secs
1	620	0.28
2	16028	14.47
3	9434	12.78
4	2922	5.28
5	435	0.98
6	46	0.12
7	15	0.05
8	27	0.10
9	1	0.00
Total Seconds		34.06

CPU Time on R if Session to C is Reset

Signing Cost

- You only sign once per update per peer (not dependent on AS-Path length)
- You only sign toward BGPSEC speakers

Need not Sign To Stubs

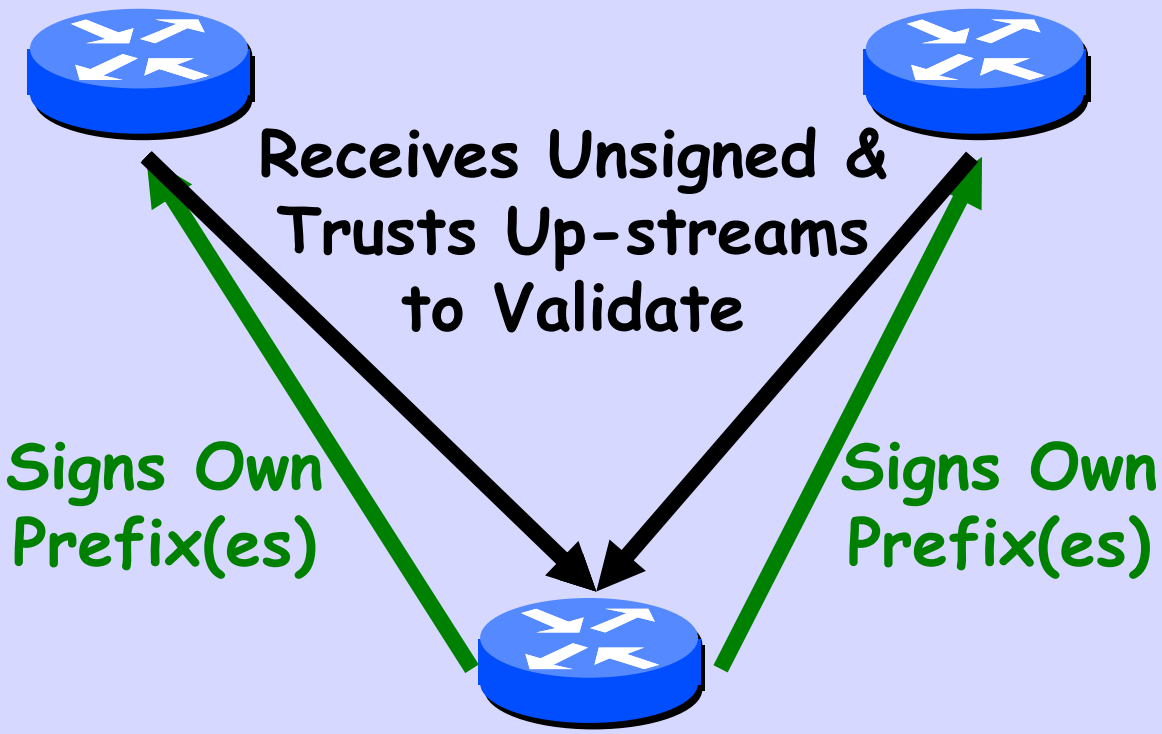

Only Needs to Have Own Private Key, No Other Crypto or RPKI Data

No Hardware Upgrade!!

Stub ASs vs. Transit

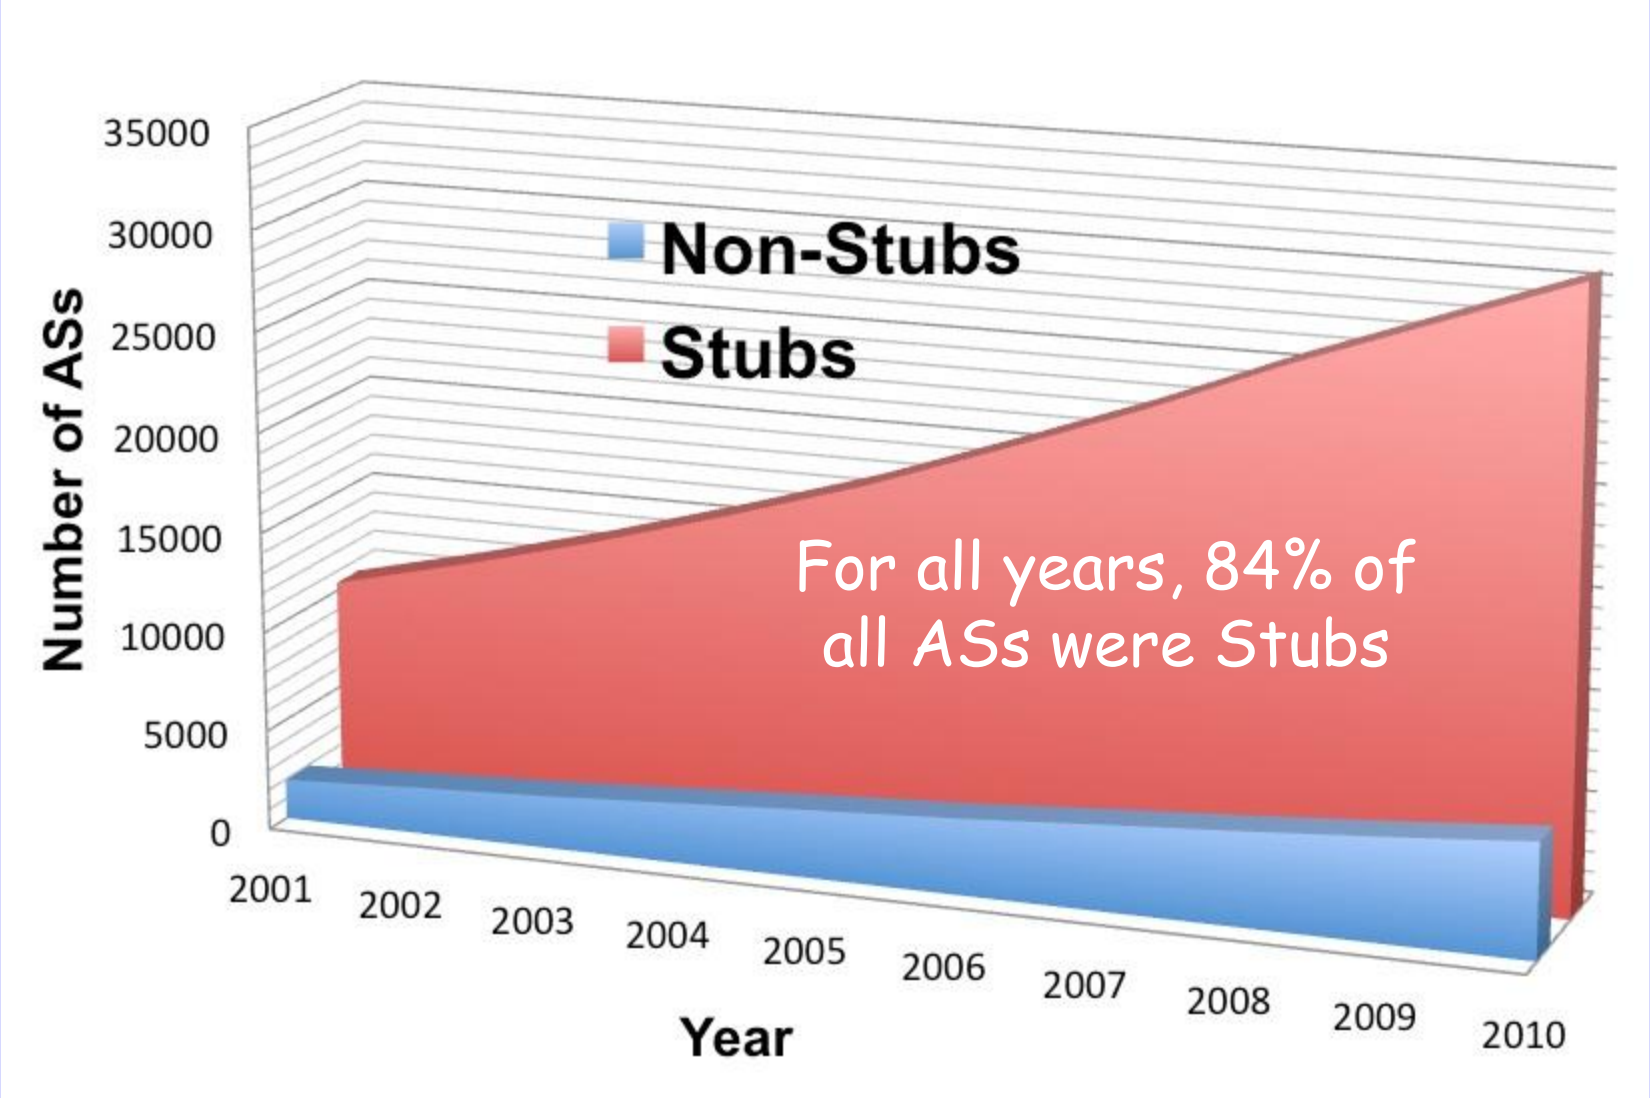

BGP Peers per Router

ISP	BGP Peers	BGP Custs
W	29	95
X	3-4	20
Y	6	12
Z	8	16

These numbers are from real ISPs, but large ones

Signing Bottom Line

- Except for W, it comes to 2-3 BGPSEC customers per aggregation router (Note: 16% are non-stub)
- Say 400k routes at 2530 sigs/sec
- (3*400000)/2530 = 475 seconds
- But this presumes the entire Internet is signed, which is a looooooooong time from now
- But W will eventually have a problem!

CPU for Validation and Signing

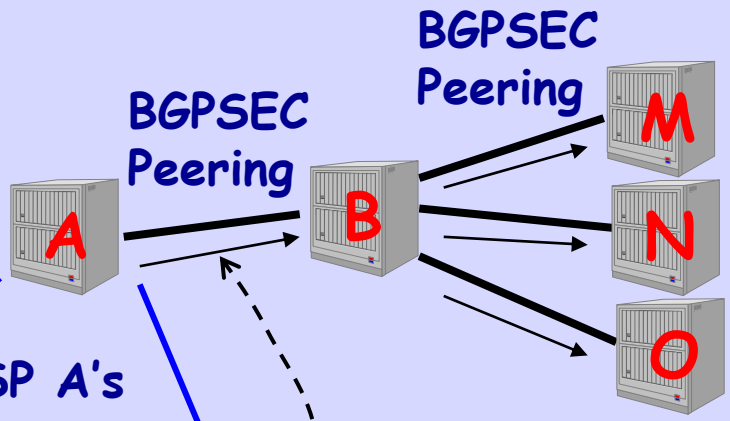

- B peers with four BGPSEC peers
- B's other peers are not BGPSEC aware

ISP A's Customer Cone

Path Len	#Pfxs	Secs
1	1353	2.21
2	21586	45.08
3	6820	17.32
4	1627	4.87
5	942	3.24
6	45	0.18
7	14	0.06
8	6	0.03
Total (seconds)		73

CPU Load on B, including Validation & Signing, if Session to A is Reset.

Summary

- CPU cost estimated for Intel Sandy Bridge i7 using only a Single-core CPU at 3.4 GHz
- The CPU cost numbers for convergence after a peering session reset look very reasonable for BGPSEC island models.