Estimating CPU Cost of BGPSEC on a Router

IETF 82 SIDR WG Meeting

Paris

March 2012

Kotikalapudi Sriram <ksriram@nist.gov>

Randy Bush <randy@psg.com>

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

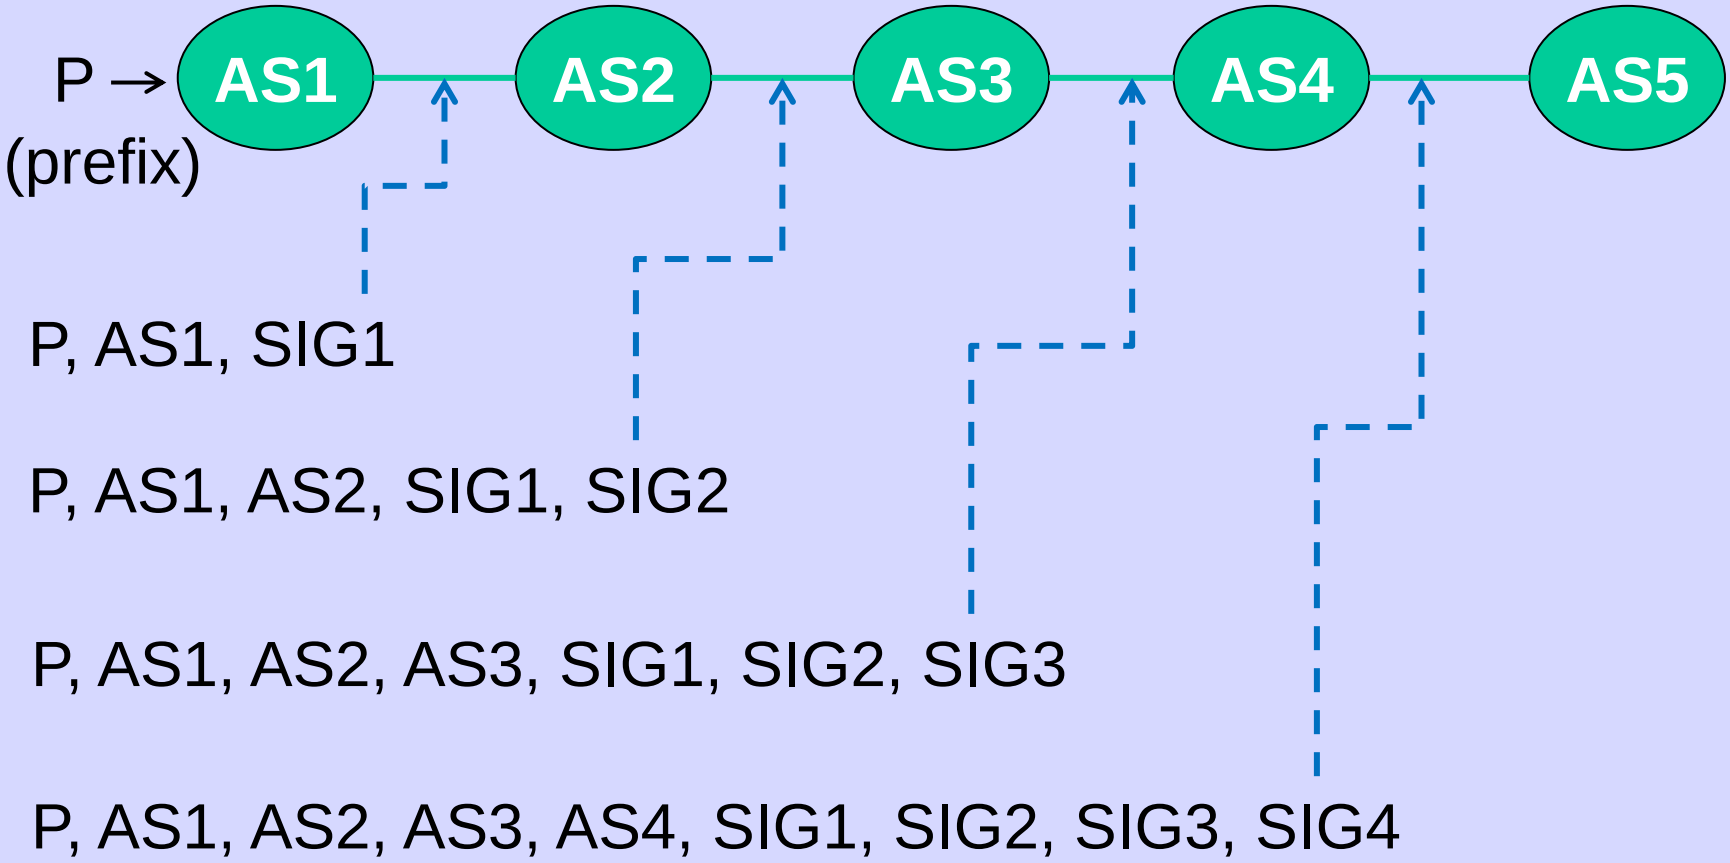

But Reality is This

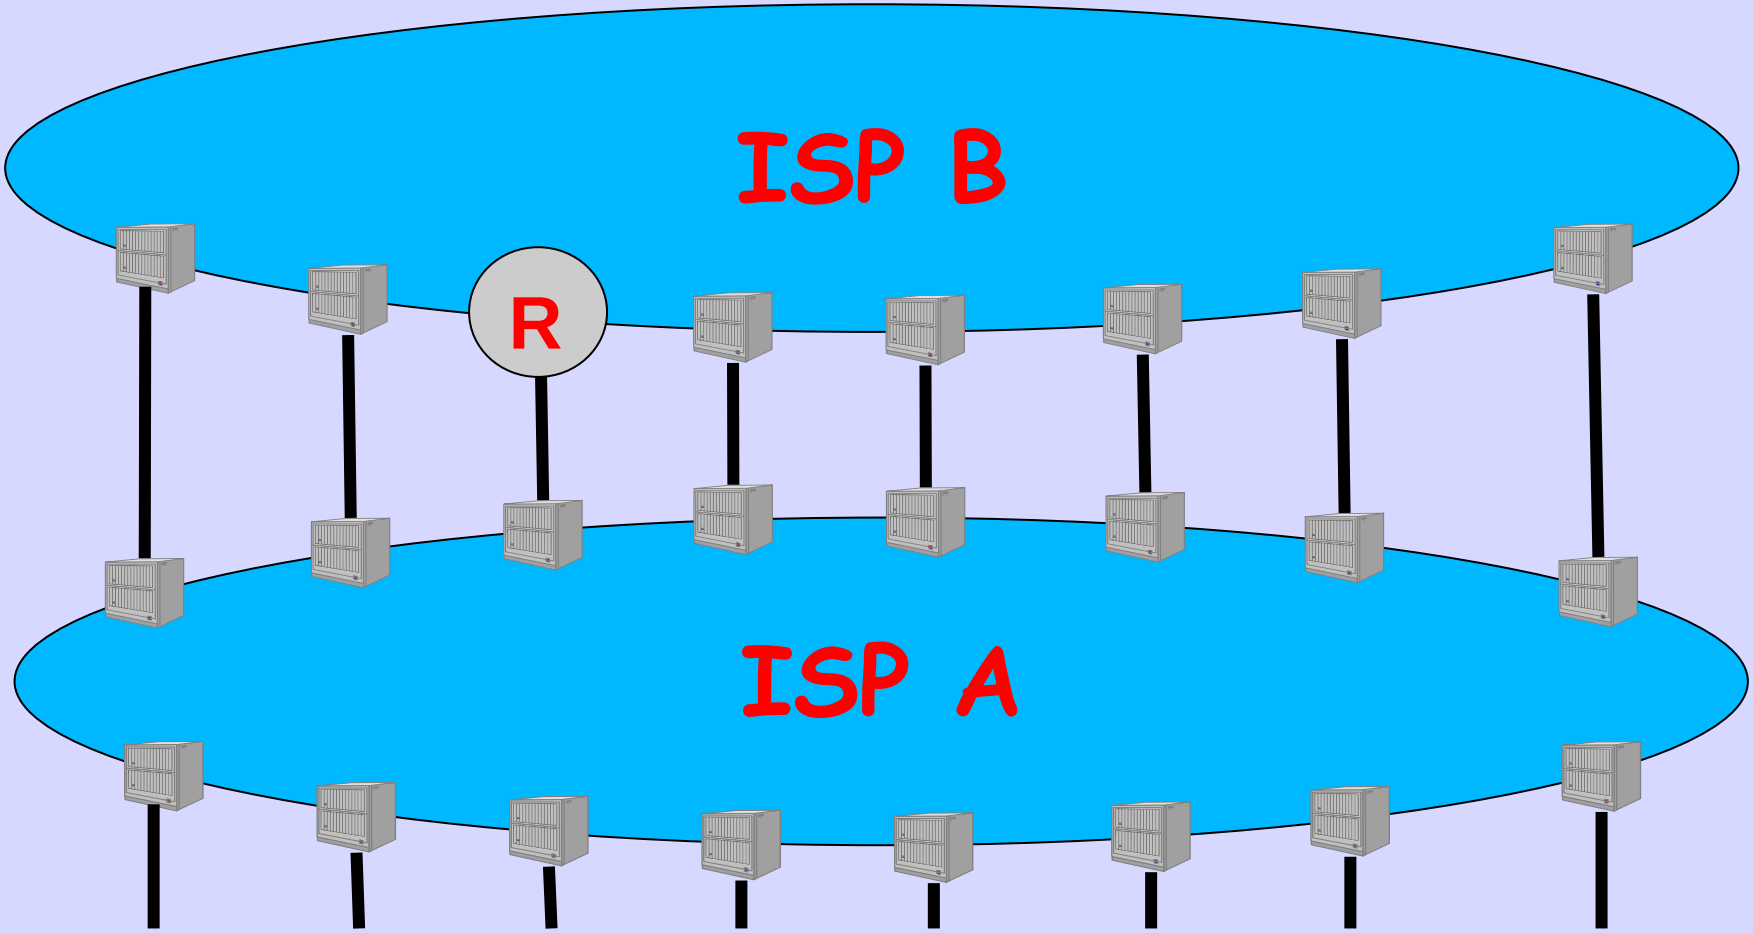

Number of Paths

- 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	

Very Large Global

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

Large Global

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

Large Global

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

Asian Regional

1	152
2	791
3	120
4	35
5	3

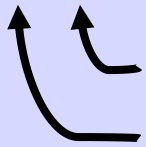

pfxs

path length at seen next upstream ISP

Yes, there are rather long tails

Yes, we removed prepending

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;	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

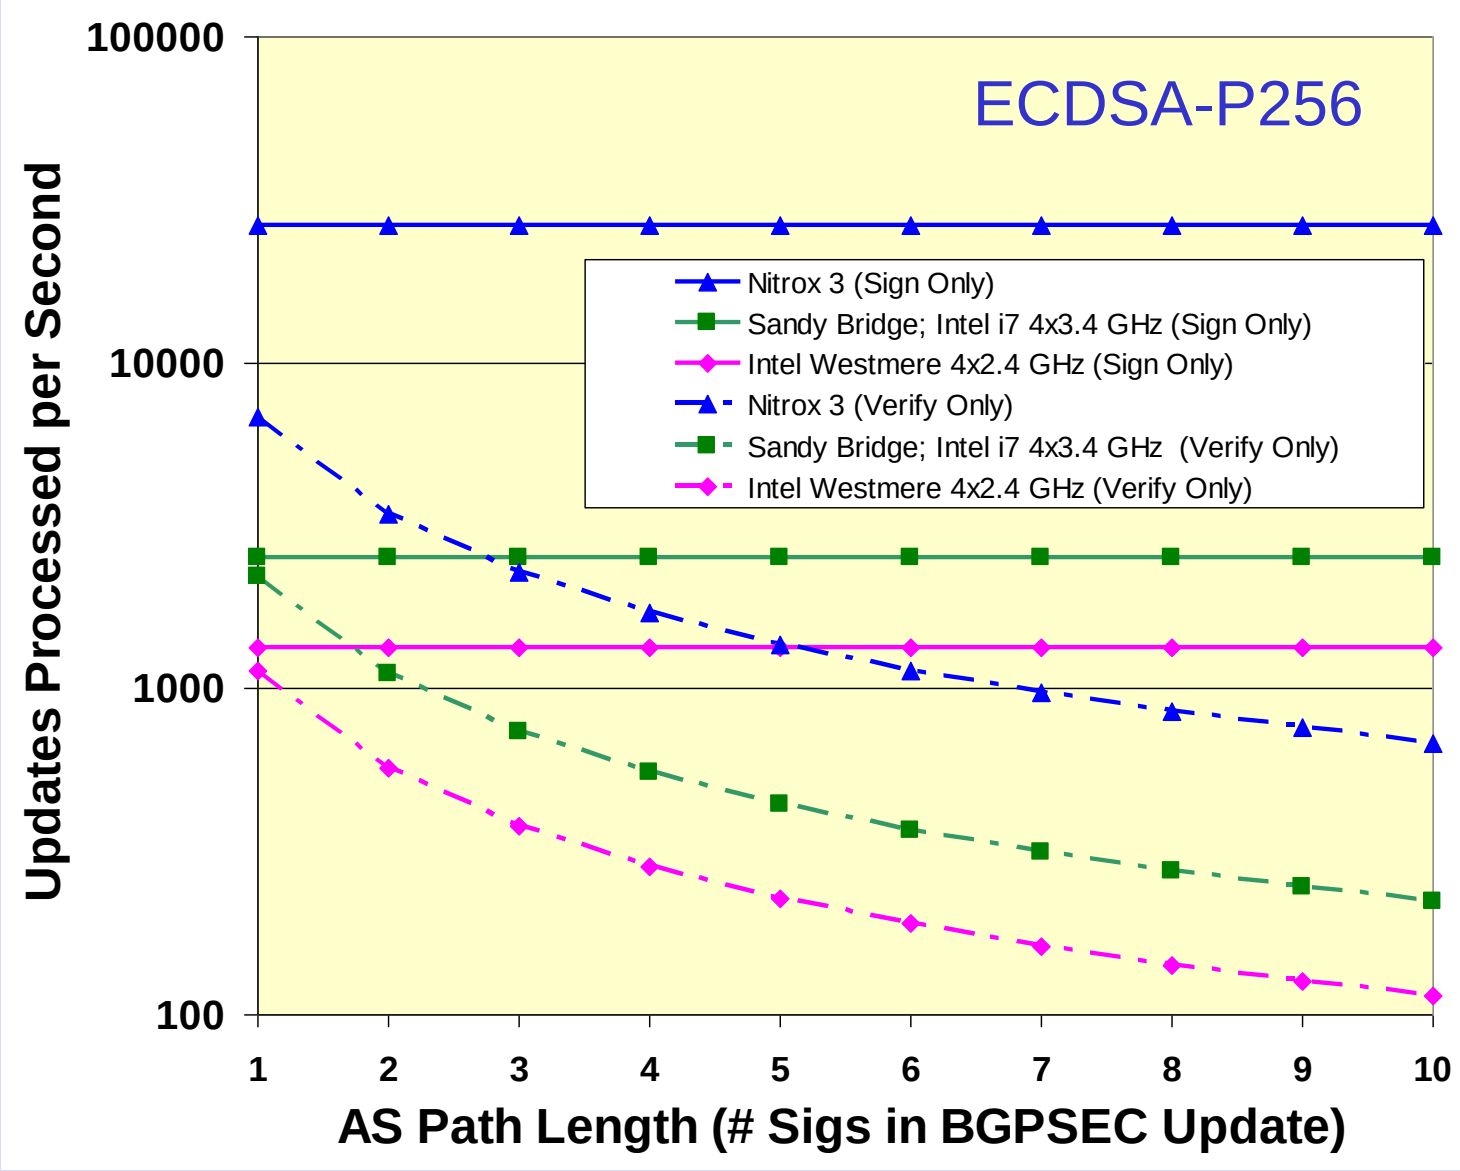

Validation Cost Model

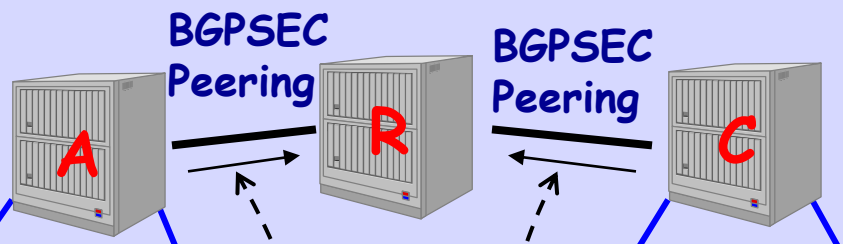

ISP A and A's Customer Cone

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

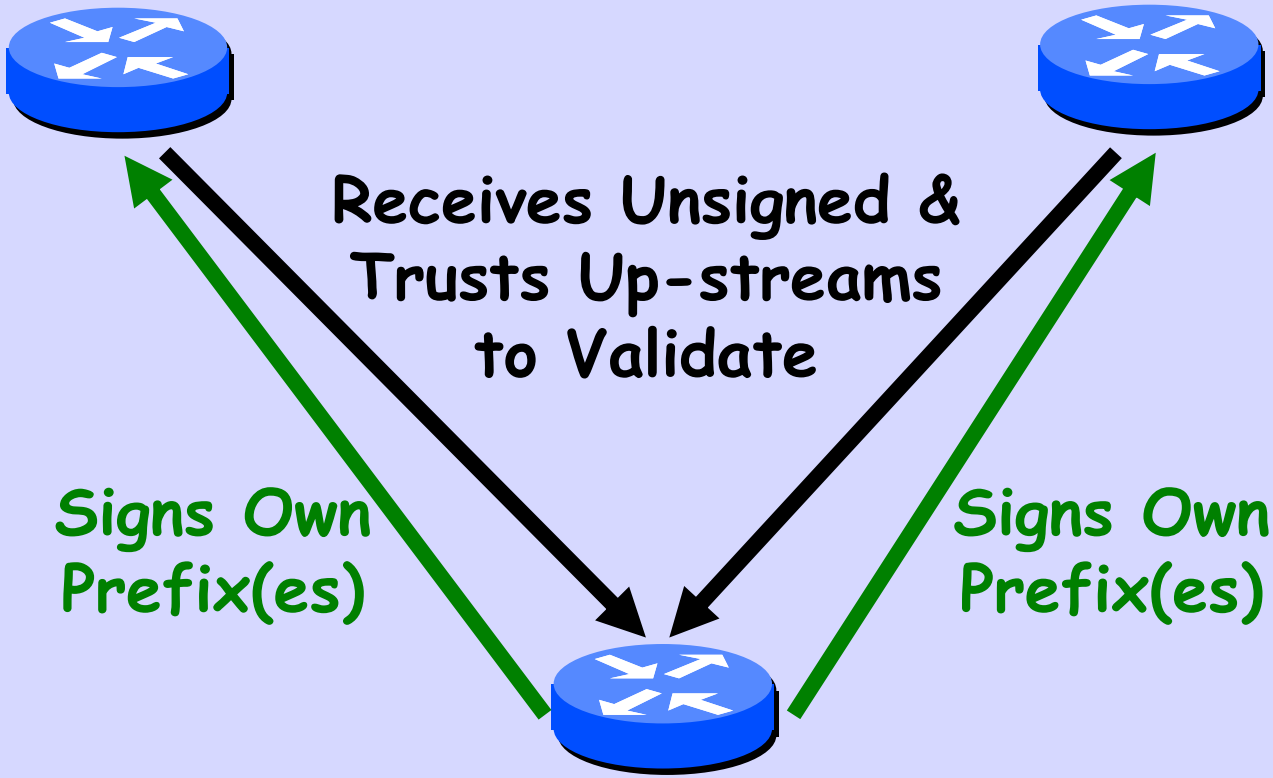

Only Needs to Have Own Private Key, No Other Crypto or RPKI Data

No Hardware Upgrade!!

Stub ASs vs. Transit

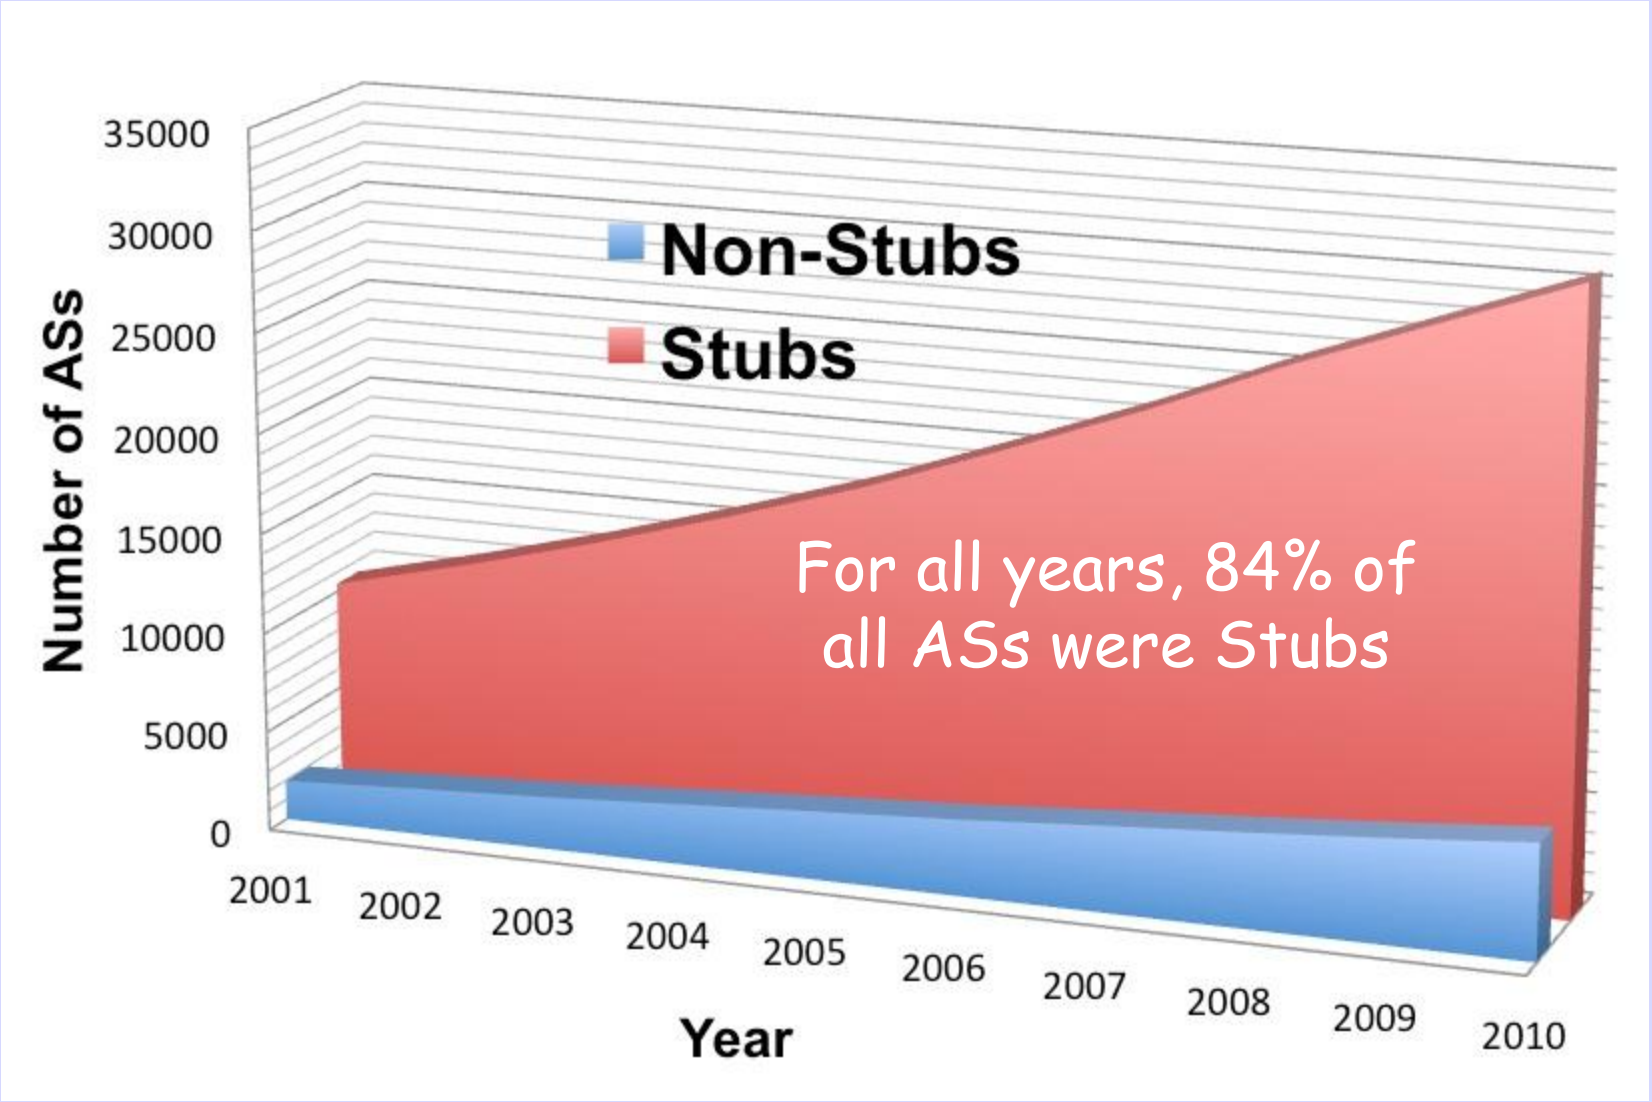

Data* on Number of Peers per Router and Number of Customers per Router for Large ISPs

ISP	Total BGP Peers	Transit (Full Table)	BGP Customers	BGP Non-Stub Customers (16%)
W	29	TBD	95	15
X	3	TBD	20	3
Y	6	TBD	12	2
Z	8	TBD	16	3

- Only non-stub customers are bi-directional BGPSEC
- 84% of customer ASes are stubs; 16% non-stub
- Router does not sign updates to stub customers

Estimated

* Source: Data collected by Randy Bush

Signing CPU Cost

- Except for W, it comes to 2-3 BGPSEC customers per aggregation router
- Say 80K routes (one fifth of current Internet) in the BGPSEC island
- Signed at 2530 sigs/sec
- If peering session with a BGPSEC customer resets, Router R needs 80,000/2530 = 32 seconds to repopulate customer's BGPSEC table

CPU for Validation and Signing

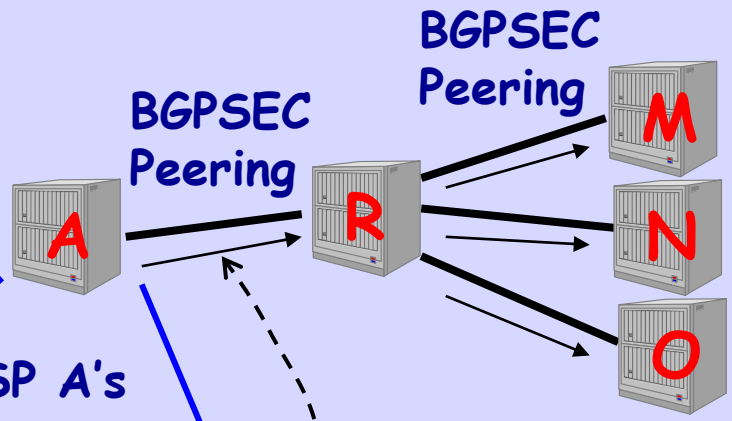

- R peers with 3 BGPSEC peers
- R's other peers are not BGPSEC aware

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 R, 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