RIB Size and CPU Workload Estimation for BGPSEC

K. Sriram
(with D. Montgomery, R. Bush)

Joint IDR-SIDR Meeting, IETF 91, Honolulu, Hawaii
November 14, 2014

RIB Size Estimation

Update Size Estimation

Avg. eBGP update size (IPv4)	78	Octets
Avg. # prefixes per update	3.8	
Avg. # ASes per update	4.2	
ECDSA-P256 signature size (per AS)	64	Octets
Avg. eBGPSEC update size (IPv4)	418	Octets
Avg. eBGPSEC update size (IPv6)	430	Octets

- No prefix packing in BGPSEC updates
- One update per prefix

Measurement of Prefixes and Paths in ISP's Route Reflectors and PE Routers

Measurement data from a Large, Tier 1 ISP in 2011

	Provider Edge (PE) routers	Route Reflectors (RR)
# Unique Prefixes Observed	377,000	377,000
Total number of Prefix Routes Observed (Low)	750,000	3,100,000
Total number of Prefix Routes Observed (High)	1,100,000	3,600,000
Ratio (Total # Prefix Routes / # Unique Prefixes) (High case)	2.92	9.55

- We assume these ratios remain approximately constant.

IPv4 and IPv6 Prefix Growth & BGPSEC Adoption Rate Projections

Year	IPv4 prefixes	IPv6 prefixes	BGPSEC adoption
2016	590000	36400	2%
2017	640000	54000	6%
2018	690000	80000	15%
2019	740000	119000	30%
2020	794000	177000	50%

- IPv4 and IPv6 growth projections based on G. Huston's IEPG presentation, November 2014. http://www.iepg.org/2014-11-09-ietf91/index.html.
- BGPSEC adoption is based a Normal-distribution growth assumption.

BGPSEC RIB Size Projection

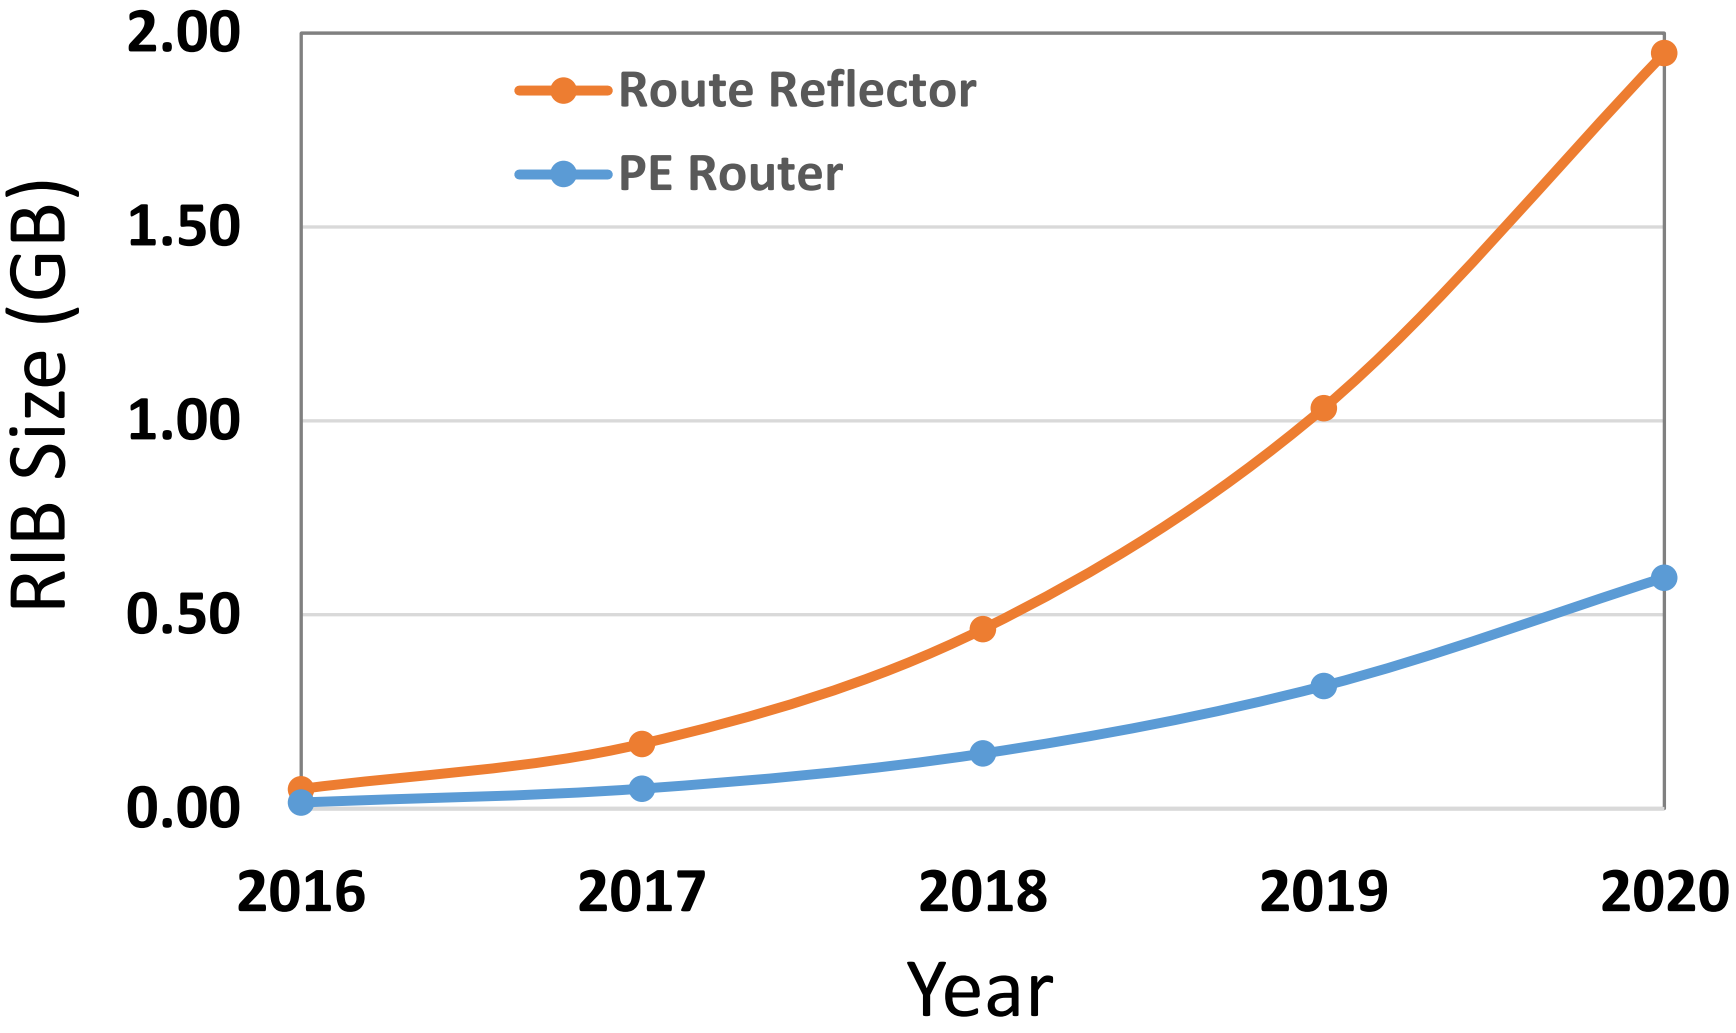

CPU Workload Estimation

(using BGPSEC island model)

BGPSEC Sign/Verify: Operations per sec Using One Core

	Operations per second	
	amd64; Westmere (206c2); 2010 Intel Xeon E5620; 1 x 2400MHz	amd64, Sandy Bridge; 2011 Intel i7-2600K; 1 x 3400MHz
ECDSA-P256 Verify	1139	2215
ECDSA-P256 Sign	1335	2530

- Source: eBACS: ECRYPT Benchmarking of Cryptographic Systems
  `http://bench.cr.yp.to/results-sign.html`

Updates Per Second

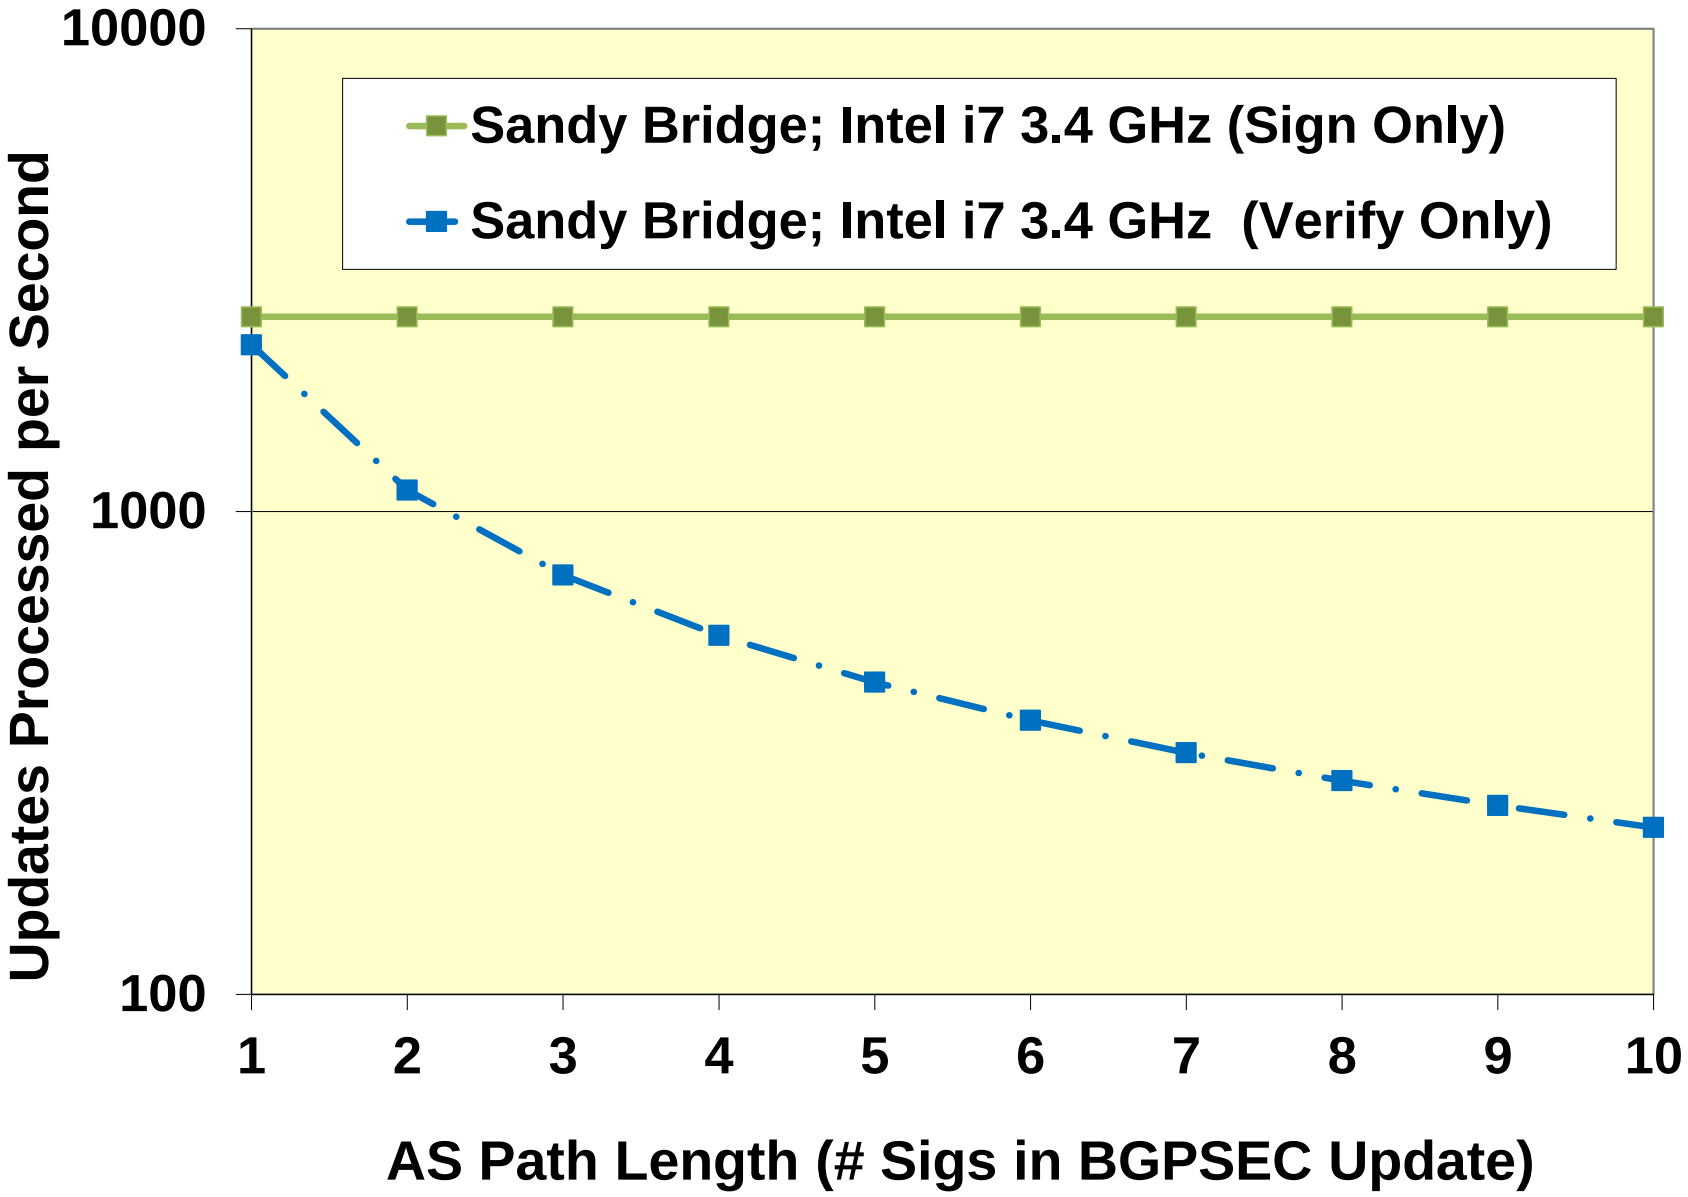

Validation Cost Model

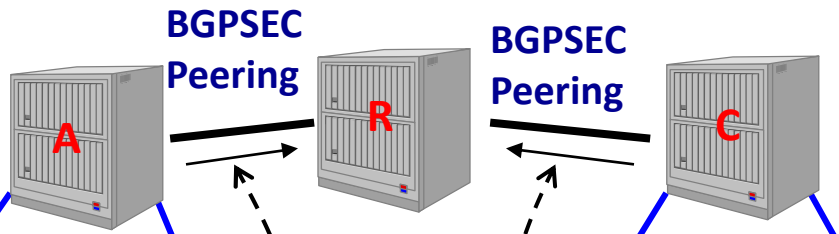

ISP A and A's Customer Cone

CPU Time on R if Session to A is Reset

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

ISP C and C's Customer Cone

CPU Time on R if Session to C is Reset

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

Need not Sign To Stubs

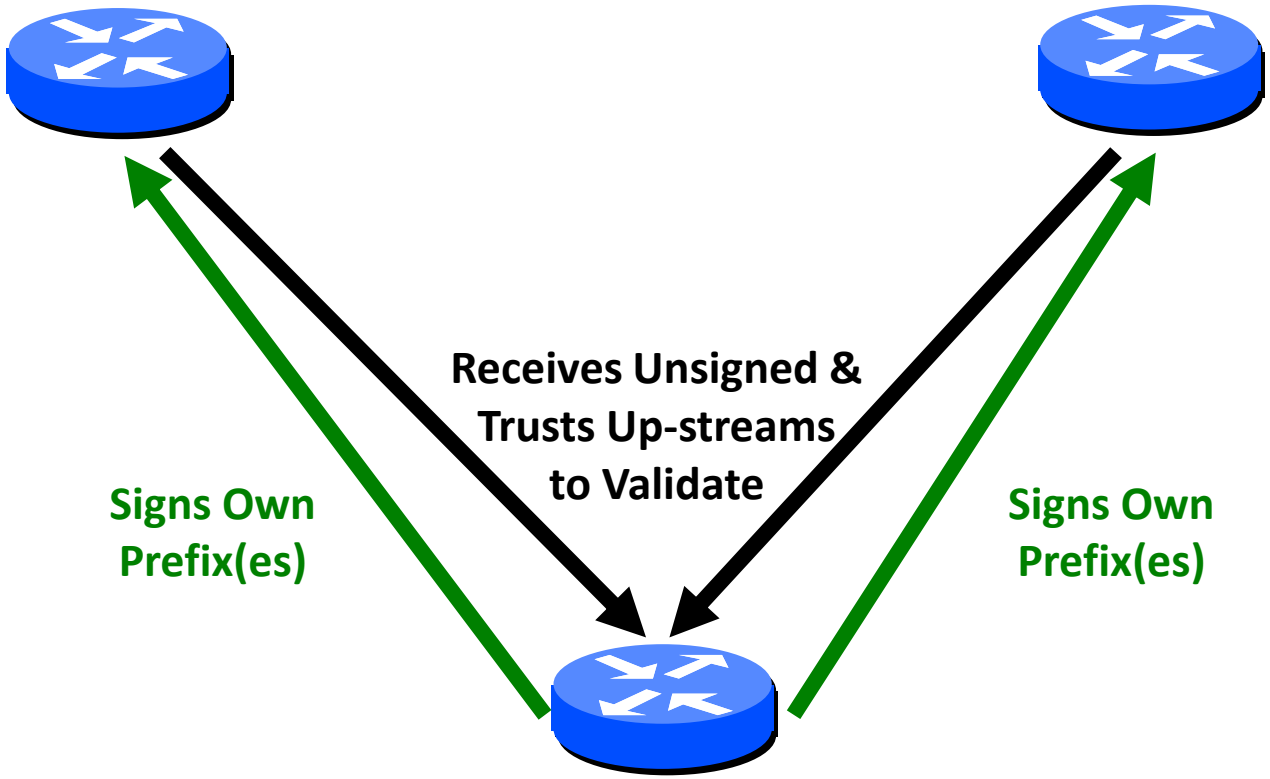

Only Needs to Have Own Private Key, No Other Crypto or RPKI Data

No Hardware Upgrade!!

What Fraction are Stub ASs?

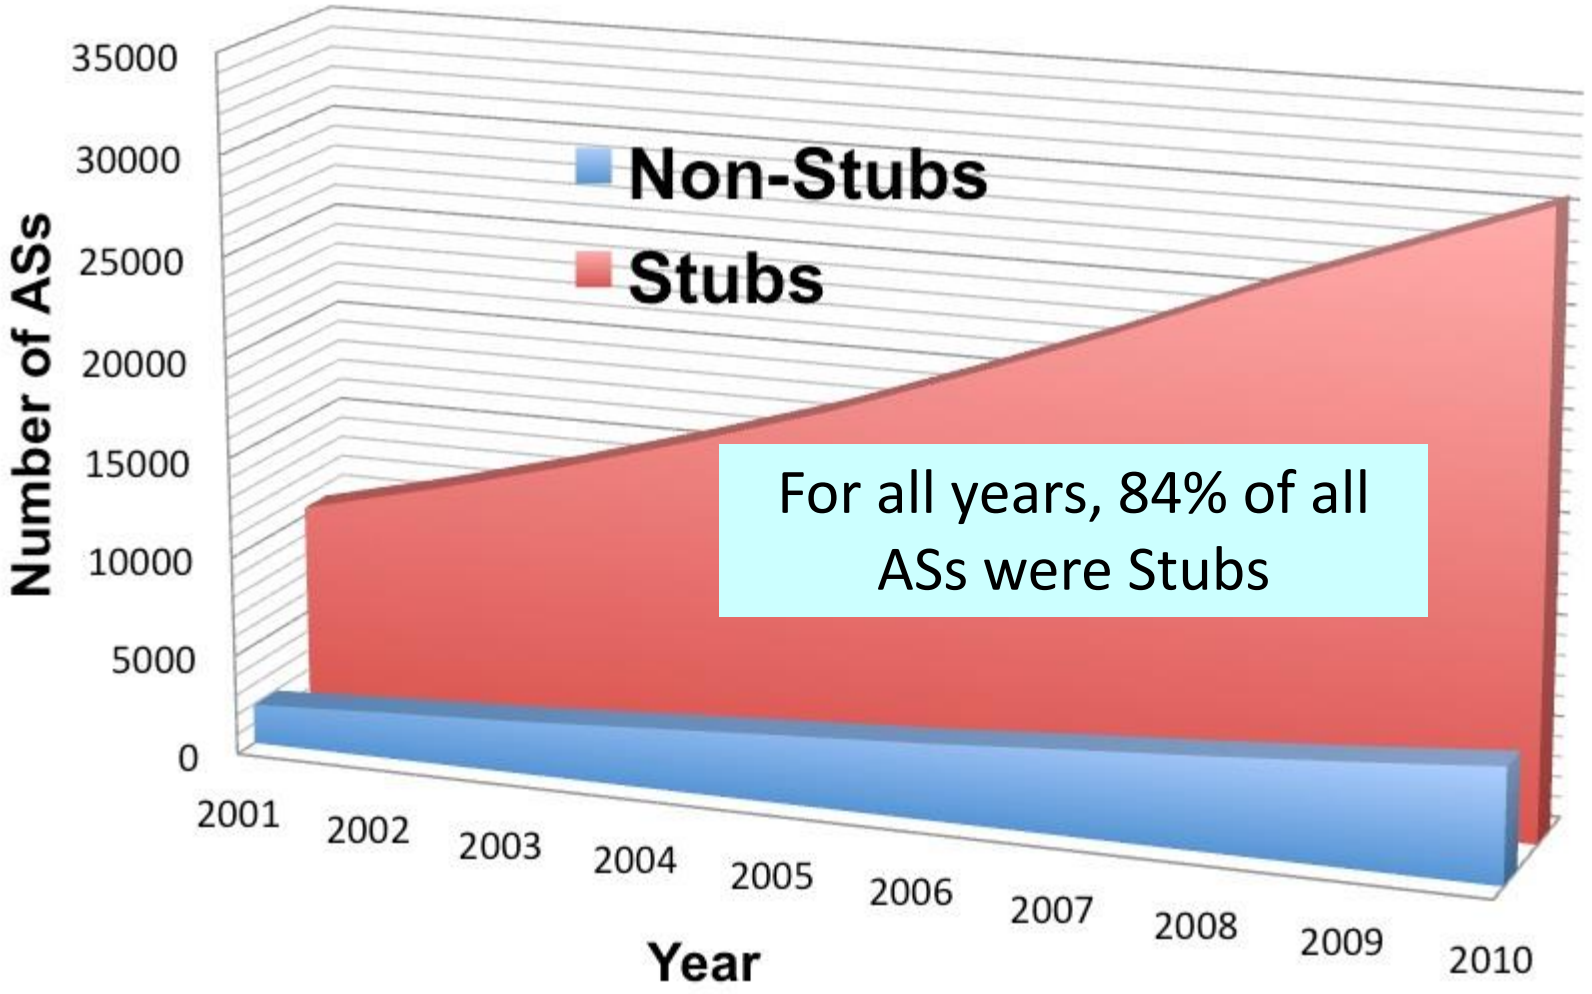

CPU for Validation and Signing

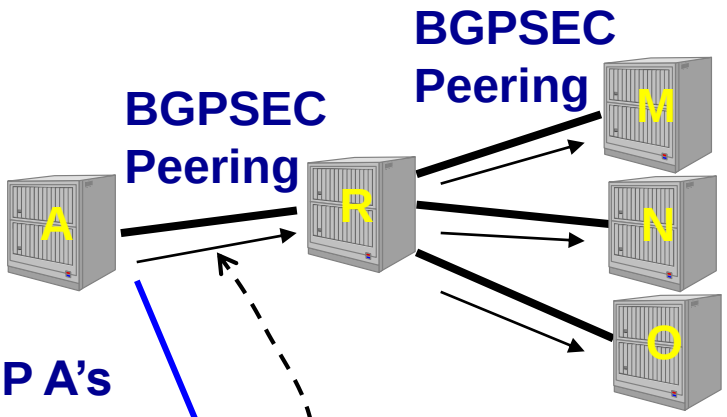

Path Len	#Pfxs	Secs
1	1353	2.21
2	21586	45.08
3	6820	17.32
4	1627	4.87
5	942	3.24
6	45	0.18
7	14	0.06
8	6	0.03
Total (seconds)		73

- R peers with 3 non-stub BGPSEC peers
- R's other peers are stub ASes

CPU workload 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

For More Details

"RIB Size Estimation for BGPSEC"

https://www.ietf.org/proceedings/81/slides/sidr-2.pdf
(SIDR meeting, IETF-81)

http://www.nist.gov/itl/antd/upload/BGPSEC_RIB_Estimation.pdf
(a few more details here)

"Estimating CPU Cost of BGPSEC on a Router"

http://www.ietf.org/proceedings/83/slides/slides-83-sidr-7.pdf
(SIDR meeting, IETF-83)

http://ripe63.ripe.net/presentations/127-111102.ripe-crypto-cost.pdf
(slightly different version presented at RIPE-63)

Thank you.

Questions?

Backup slides

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

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