

# BGP Update Unpacking a First Experiment

**IEPG / Beijing**

2010.11.07

Randy Bush <randy@psg.com>

Cristel Pelsser <cristel@iij.ad.jp>

Keyur Patel <keyupate@cisco.com>

<<http://archive.psg.com/101107.iepg-updates.pdf>>

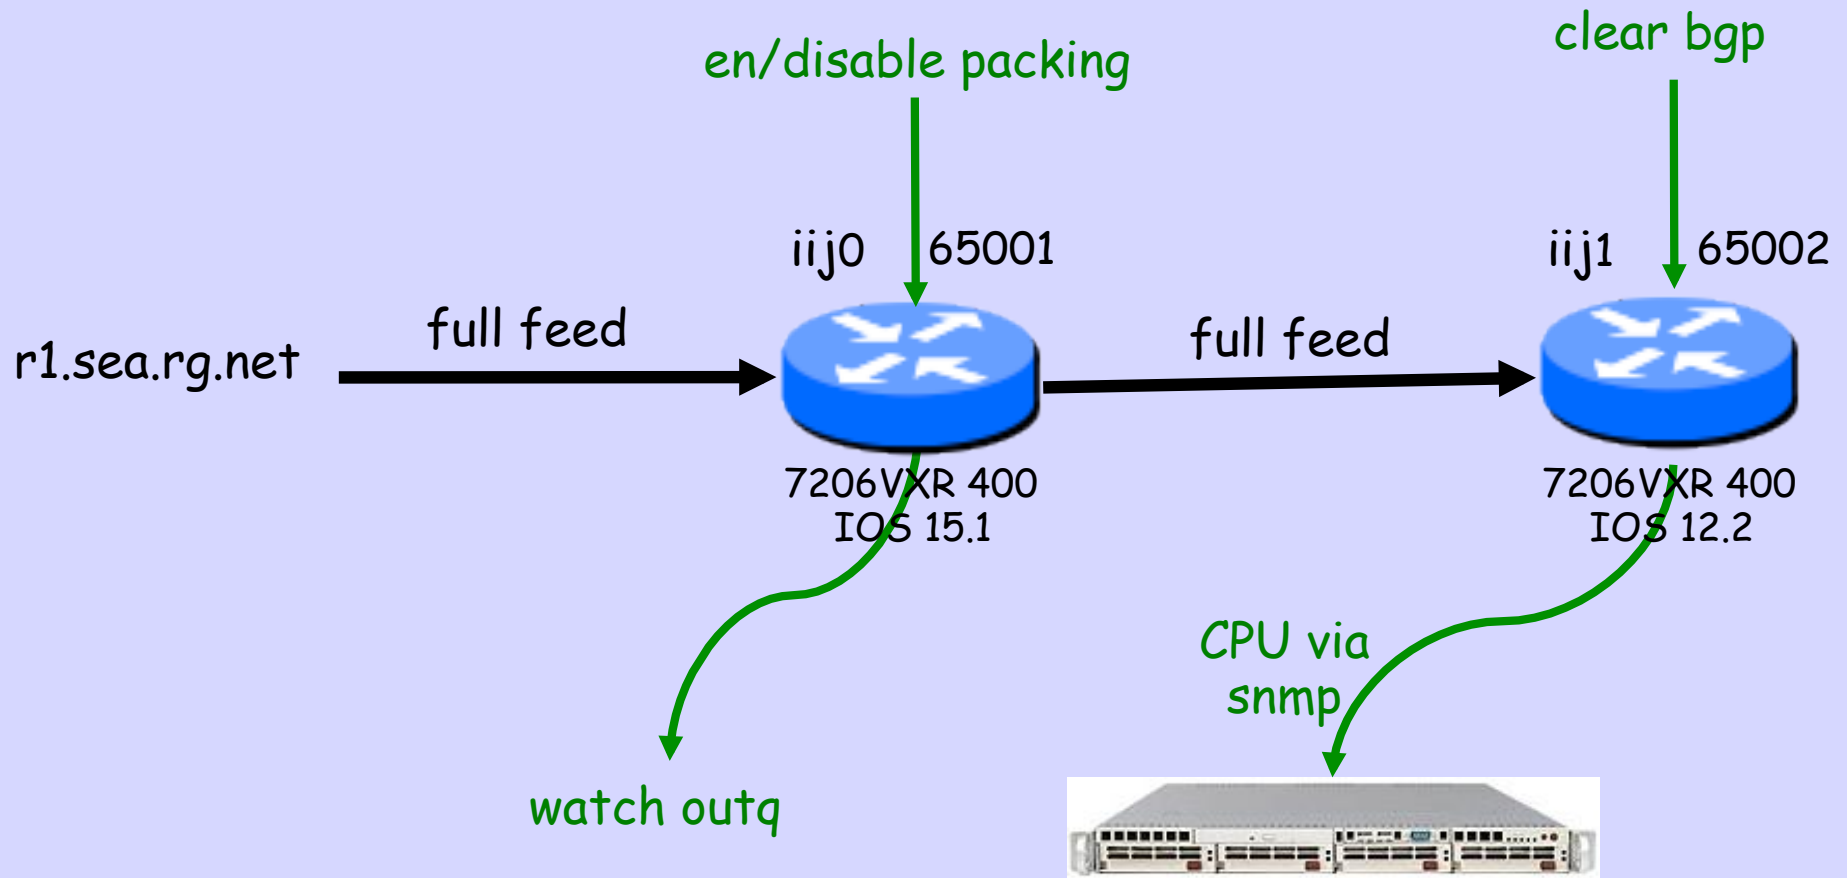
# Motivation

- BGP PDU is limited to 4096 bytes
- Pack NLRI with same attributes  
    {NLRI\*} Attributes
- If crypto or other information is added, size will be a issue
- Some crypto may force unpacking
- Will that be a big penalty?

# March 2010, Keyur Agreed to a Hack

```
retrieving revision 1.17
diff -u -4 -r1.17 iij0.bogus.com
@@ -182,8 +182,9 @@
    neighbor 202.232.15.110 description iij1 101
    neighbor 202.232.15.110 version 4
    neighbor 202.232.15.110 next-hop-self
    neighbor 202.232.15.110 send-community
+   neighbor 202.232.15.110 disable-packing
    distance mbgp 15 15 150
    no auto-summary
    !
    ip bgp-community new-format
```

# Set-Up



# Sender Queue


## Packed

202.232.15.110	65002	55121	1	0	999	00:01:29
202.232.15.110	65002	56456	2056403	0	572	00:01:31
202.232.15.110	65002	57029	2056403	0	0	00:01:33

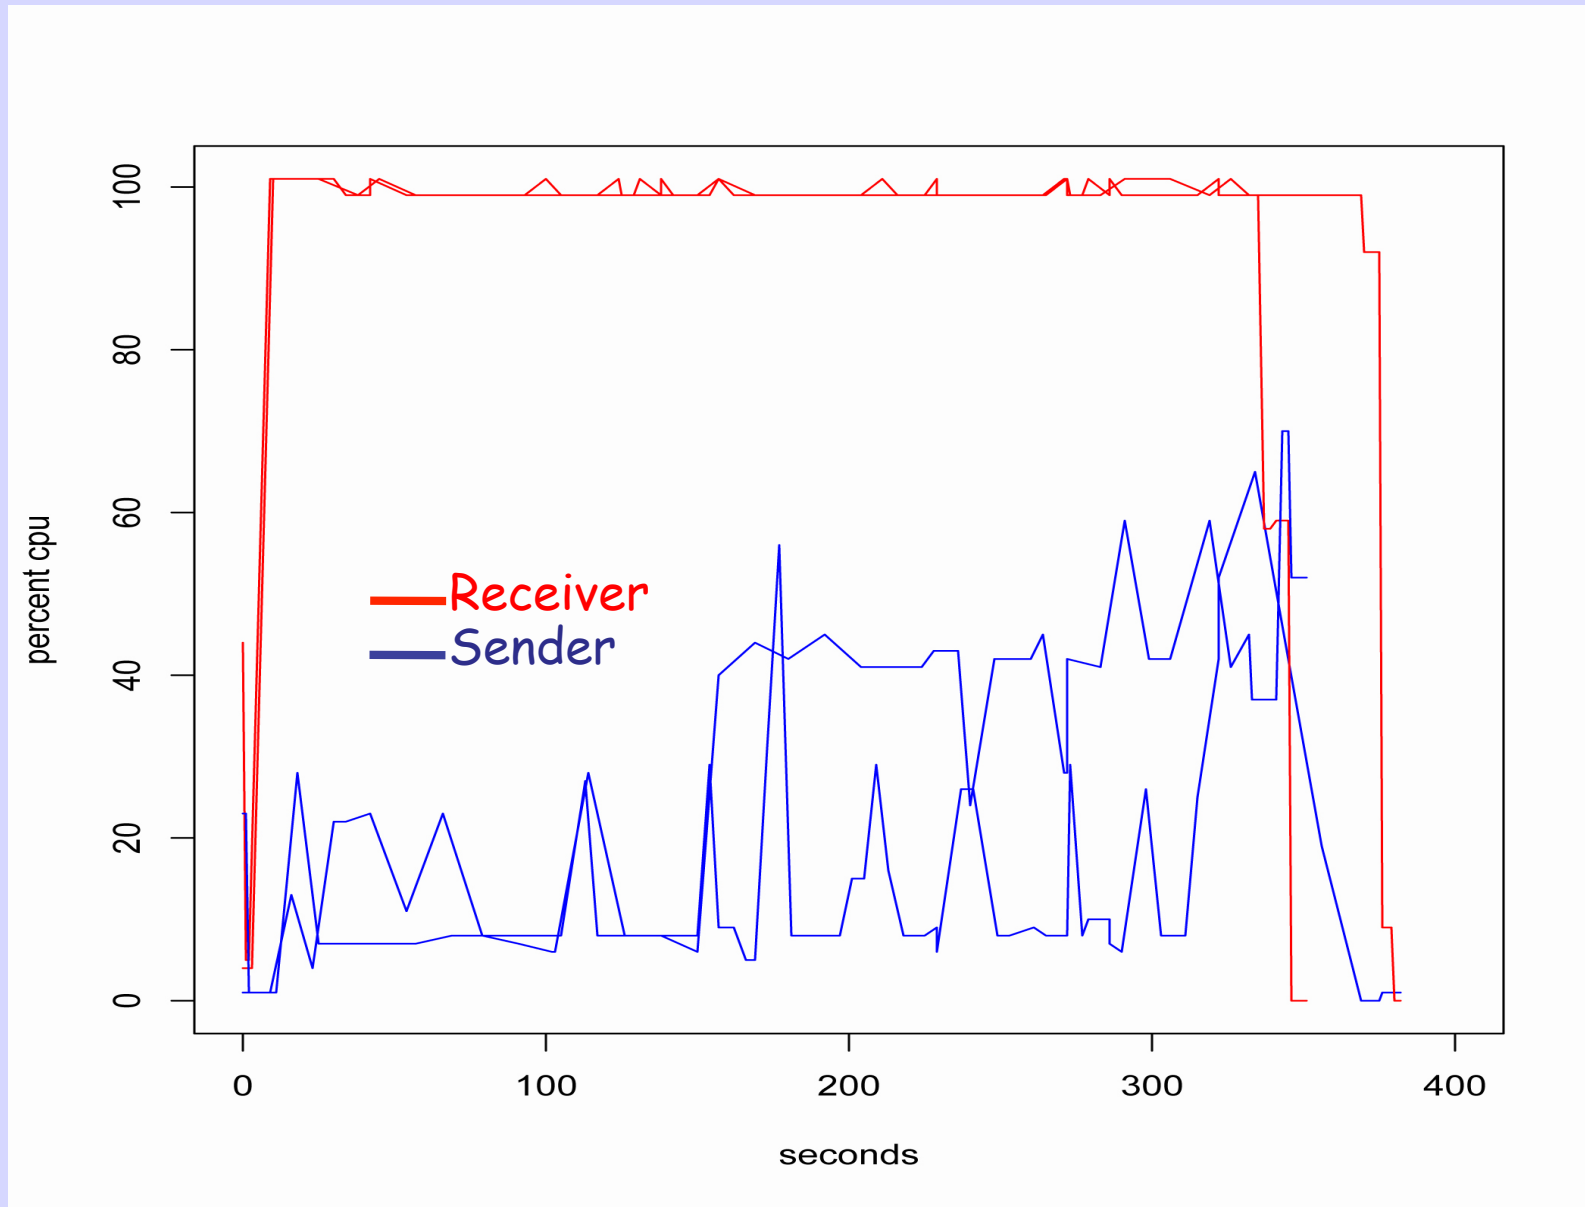
## Unpacked

202.232.15.110	65002	326220	1	0	1000	00:05:14
202.232.15.110	65002	327558	2050929	0	289	00:05:16
202.232.15.110	65002	327848	2050929	0	0	00:05:18

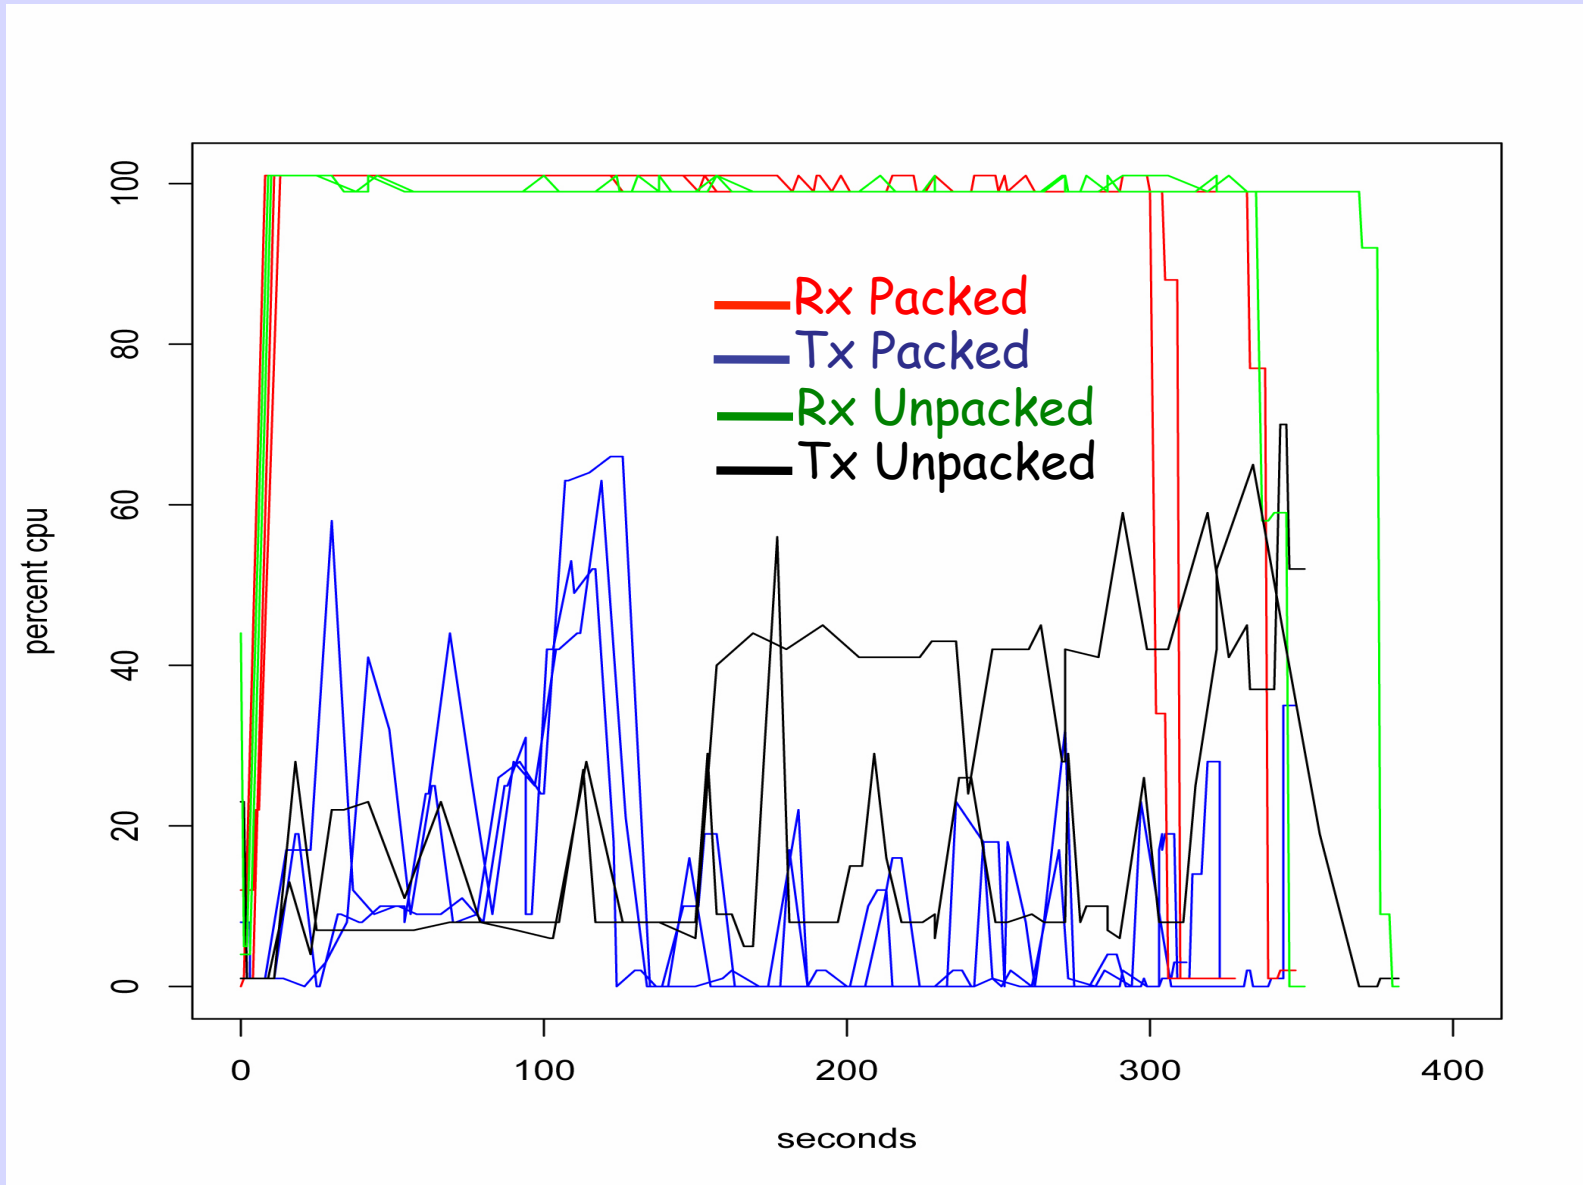
# CPU over Time - Packed



# CPU over Time - Unpacked



# CPU over Time - Combined





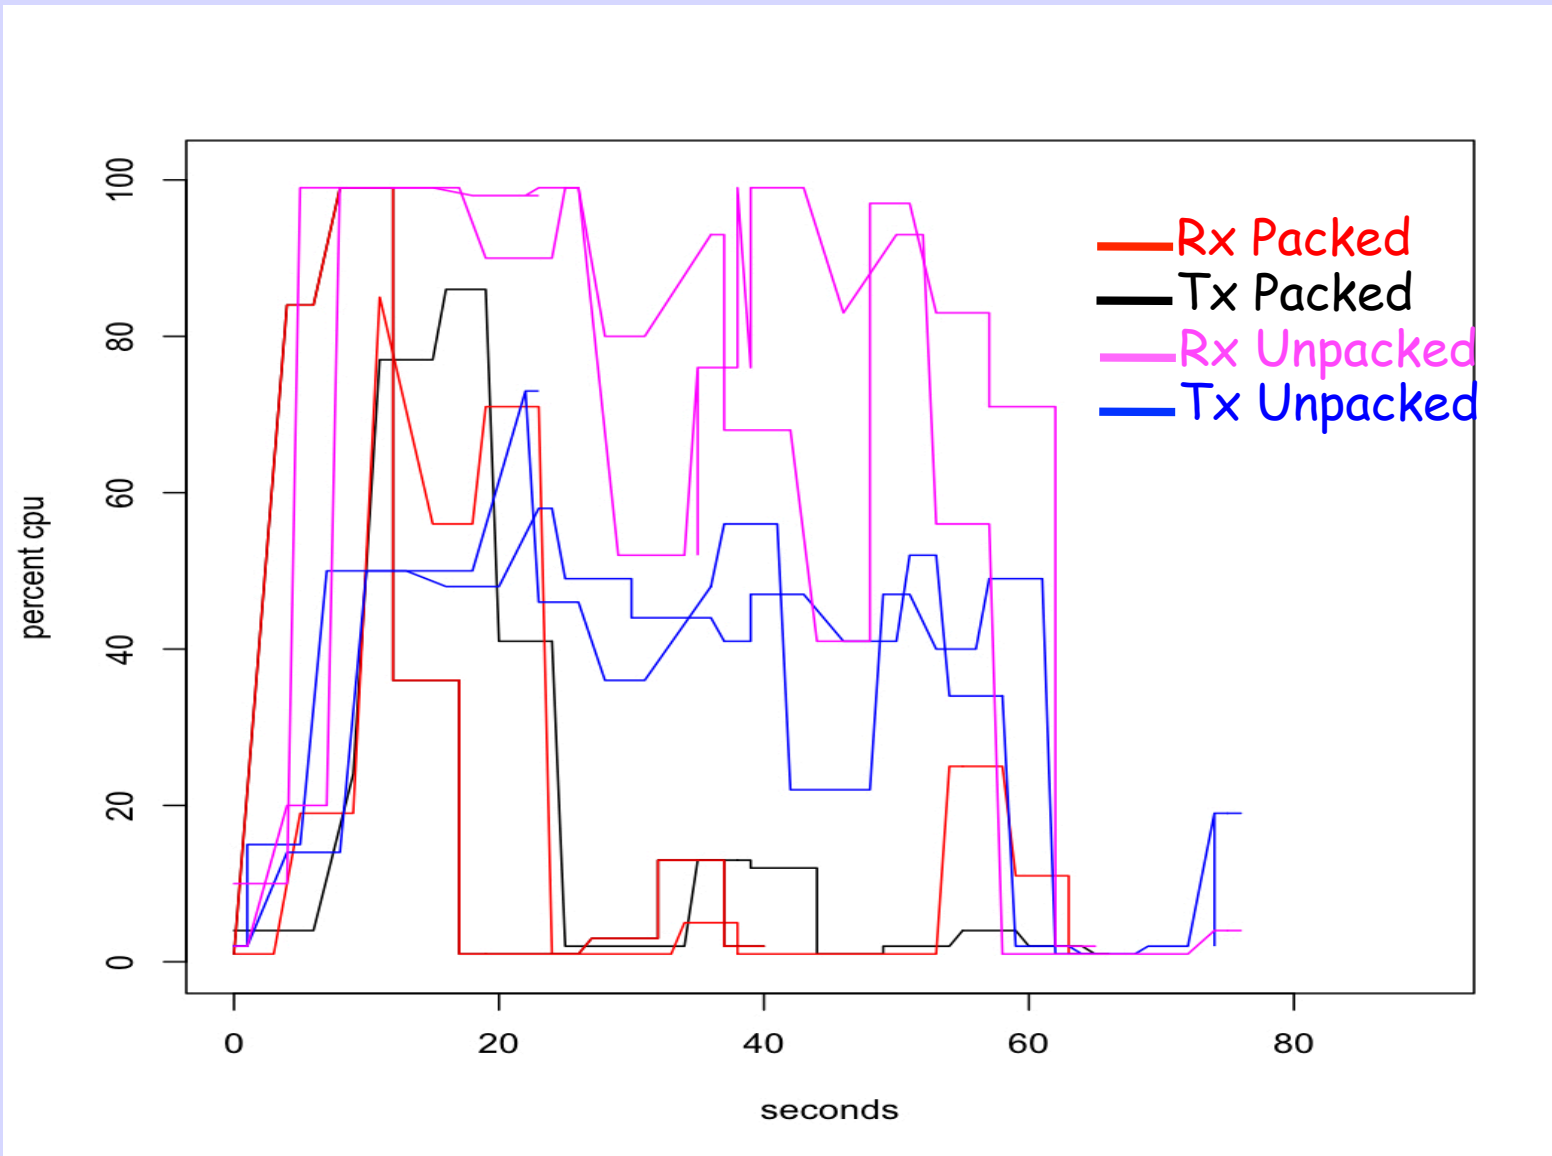
# Tentative Conclusions

- Receiver spends about the same amount of time & CPU packed or unpacked
- Tx Packed is about 1:31
- Tx Unpacked is about 5:20
- Tx Time likely biased by current Tx 'pacing' (next experiment)

# Or Is It RIB to FIB?

- Greenberg et alia found that RIB-FIB was the big cost for OSPF in 2000
- We re-ran the same experiments with route refresh inbound
- I.e. no FIB processing

# Without FIB Updating



# More Tentative Conclusions

- My Father was right, FIBs are costly
- Now we see a real difference for unpacking
- Packed about 25 seconds
- Unpacked about 60 seconds
- Still not so bad

# Early Peek at IOS XR

Platform: C12k, PRP2

IOS-XR 4.0.1

Total prefixes: 328K (57K attributes)

	<u>Packed</u>	<u>Unpacked</u>
Total bytes:	5,138,652	22,101,207

**\*\*Including RIB -> FIB install**

BGP convergence: 35s

55s

FIB convergence: 69s

65s

CPU util: 3%-35%-100%-25%-6% (60s) 3%-20%-94%-30%-4% (60s)

**\*\*Without RIB -> FIB install**

BGP convergence: 27s

43s

CPU util: 3%-30%-100%-20%-3% (26s) 4%-15%-100%-20%-3% (40s)