The Present and Future of the Internet Routing Tables

April 19, 2013

The Present and Future of the Internet Routing Tables

IPv6 Routing is to Become the World Standard

- To get best performance the community will use IPv6 routing

The Basics of Routing

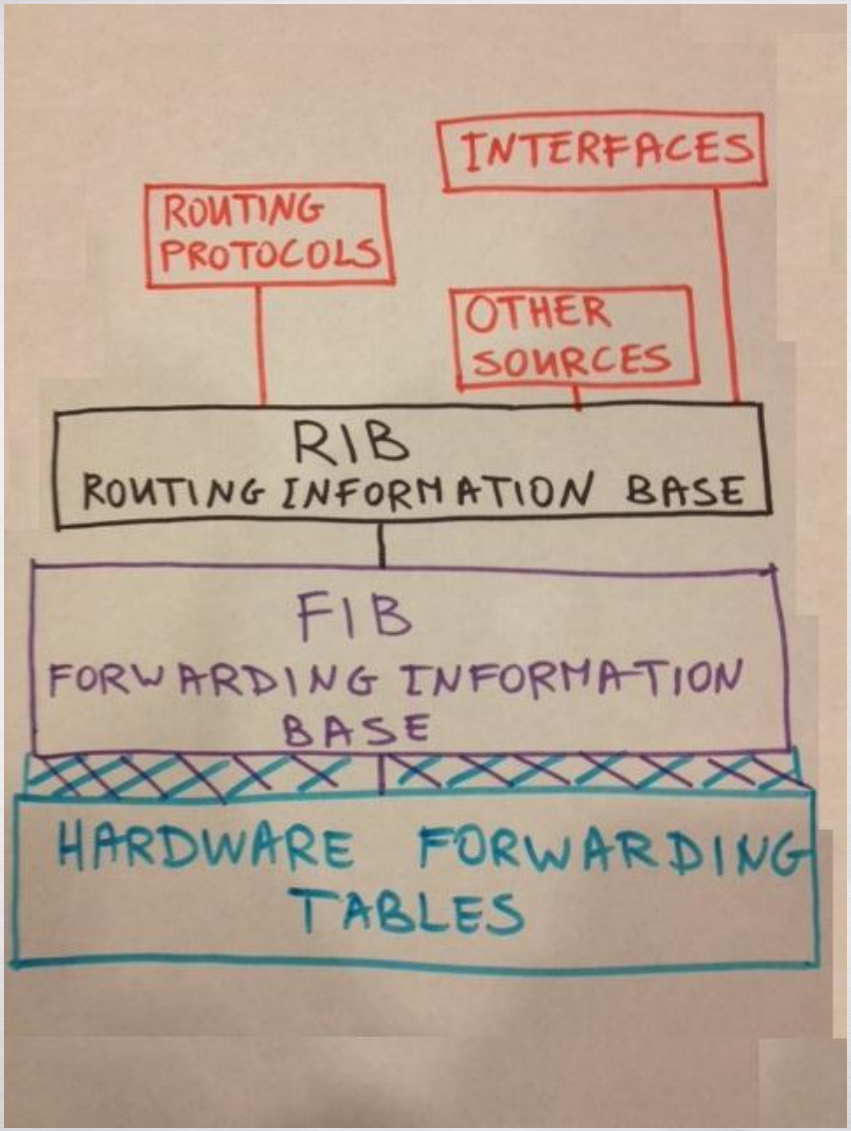

Internet Routing Tables Growth

- Sky is the limit, is it really?
  - Routing scalability is going to be one of the most important problems facing the Internet

IPv4 Routing Table Growth

- The dogs are out of the fence

- Present number of IPv4 Internet routes +450k

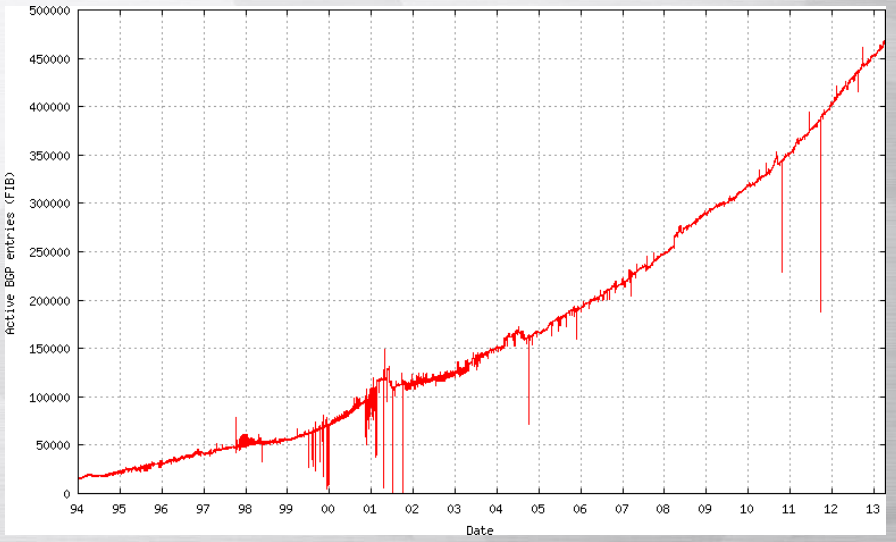

IPv6 Routing Table Growth

- A chance to train the puppy

- Present number of IPv6 Internet routes +12k

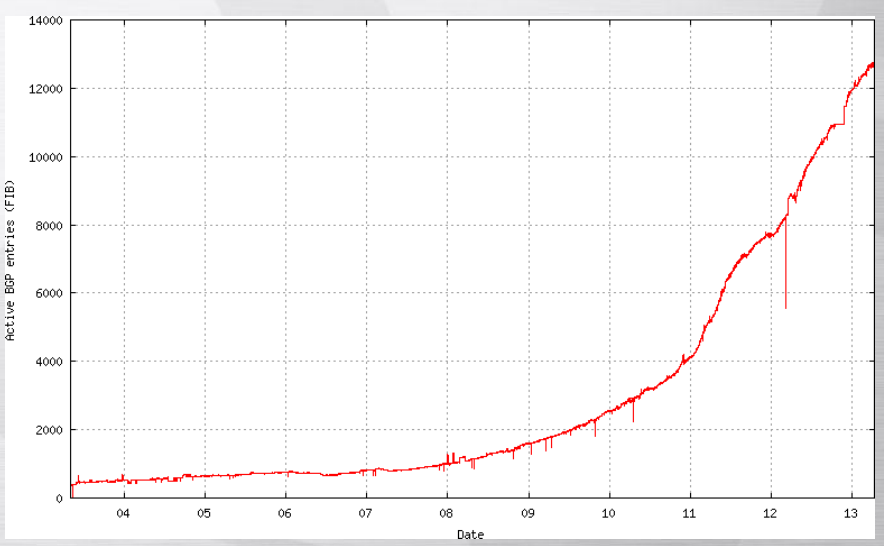

Routing Performance Observations

- IPv4 is routed over the longer route to alleviate congestion

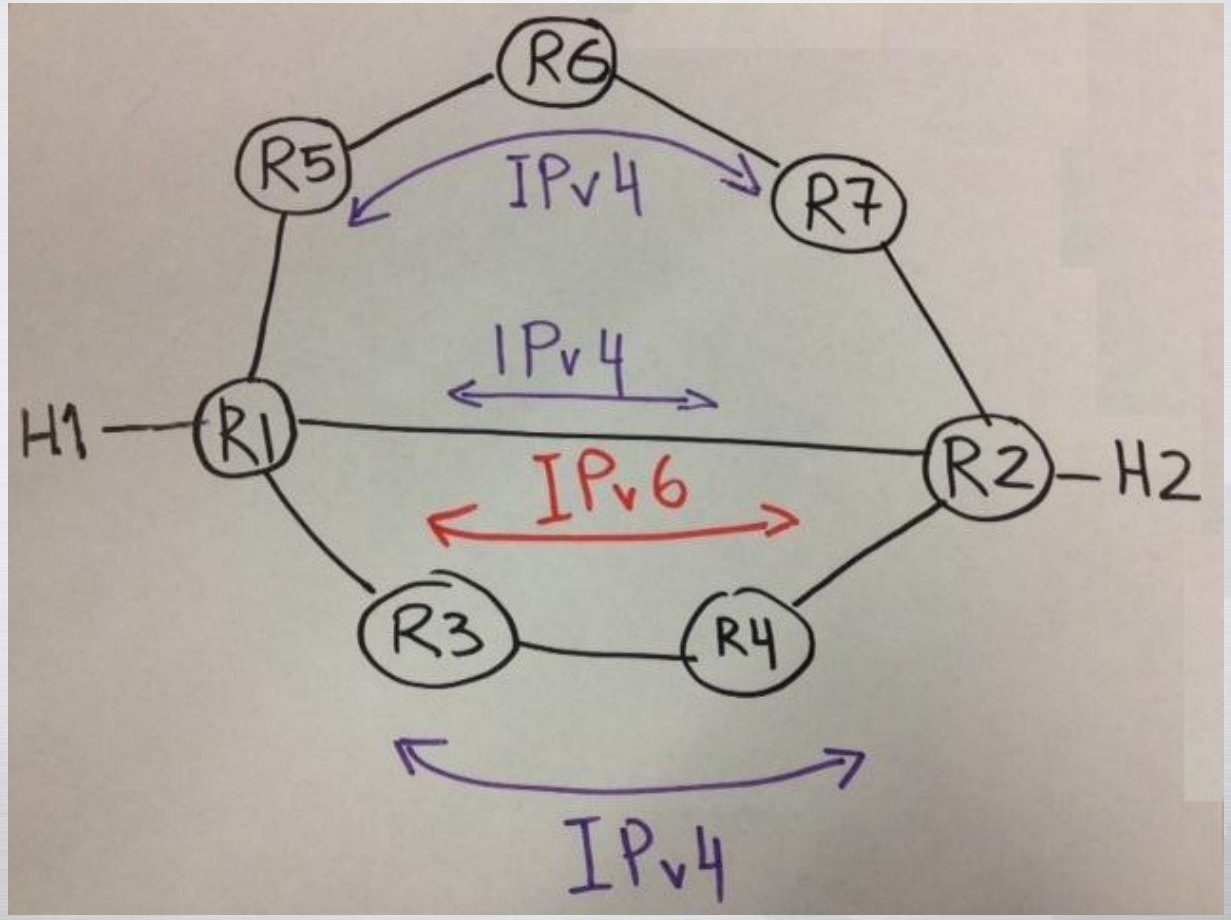

Routing Performance Observations

- Some providers do traffic engineering on IPv4 and for IPv6 they don’t (99% vs 1% of the traffic)

Routing Performance Observations

- With IPv4 you may have suboptimal routing

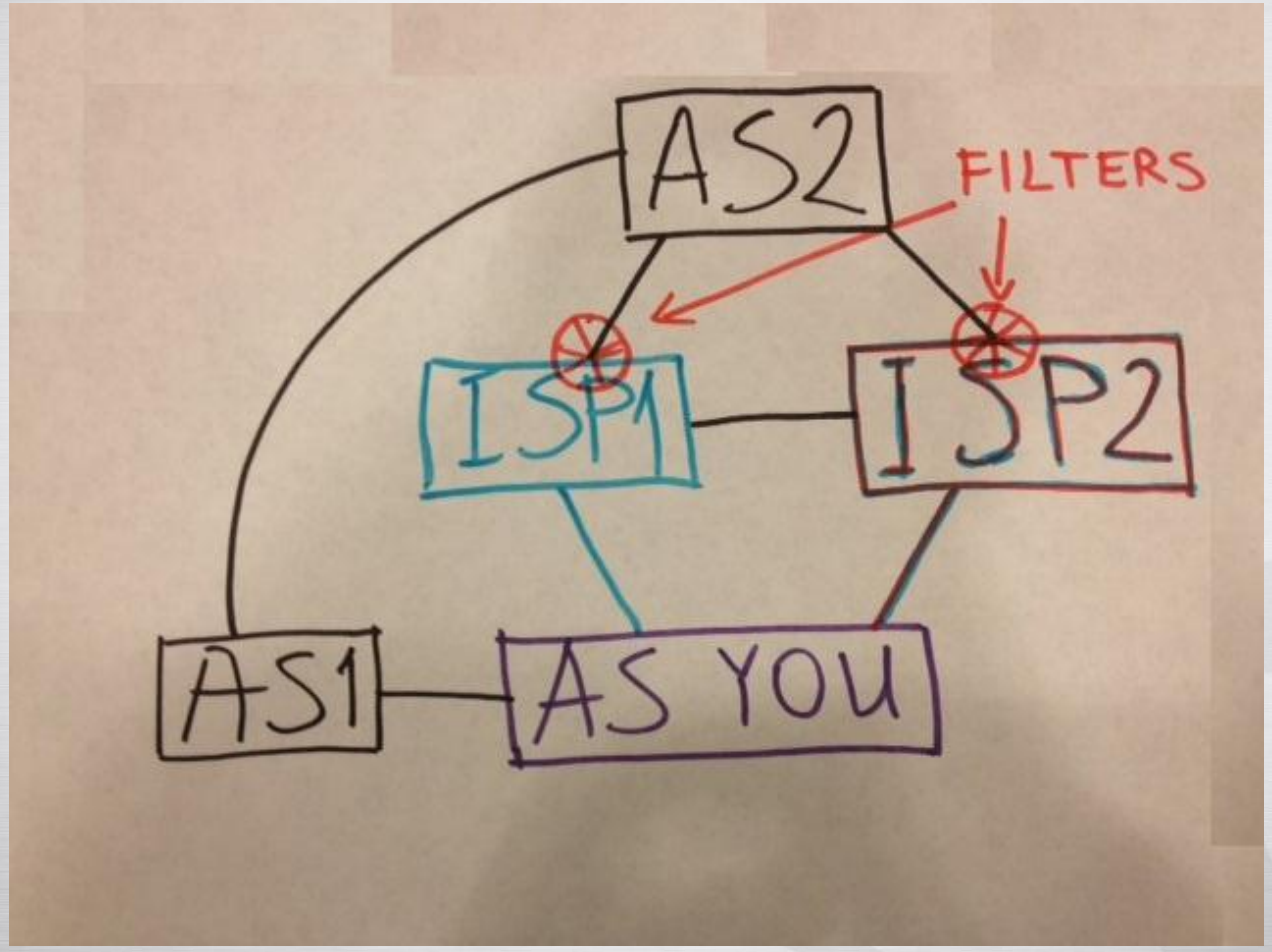

Routing Performance Observations

- Players on the Internet are starting to filter de-aggregated IPv4 prefixes

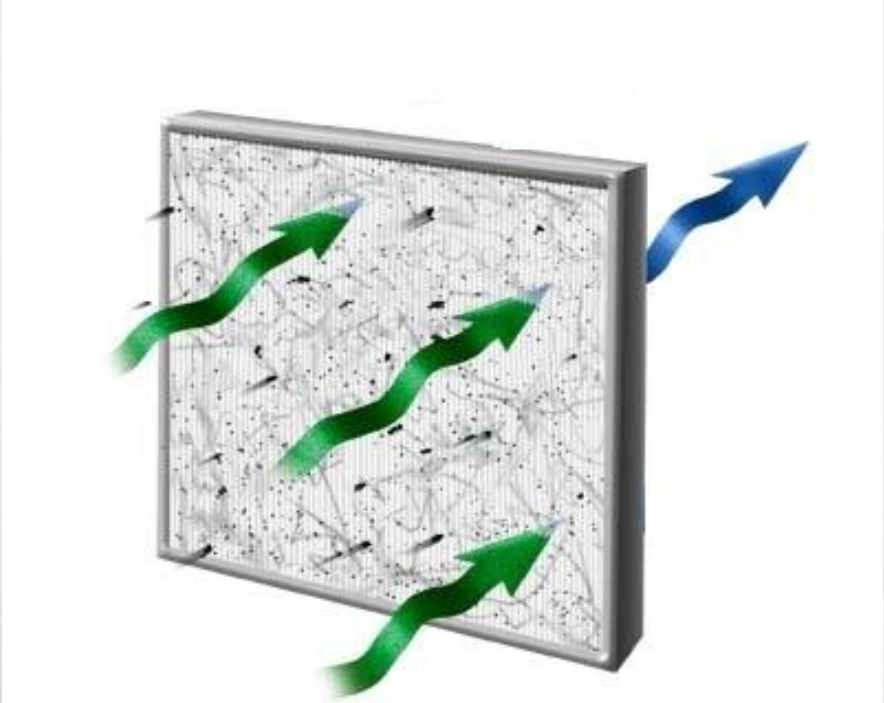

Routing Performance Observations

- Convergence with routing tables – bigger tables takes longer for CPU to process

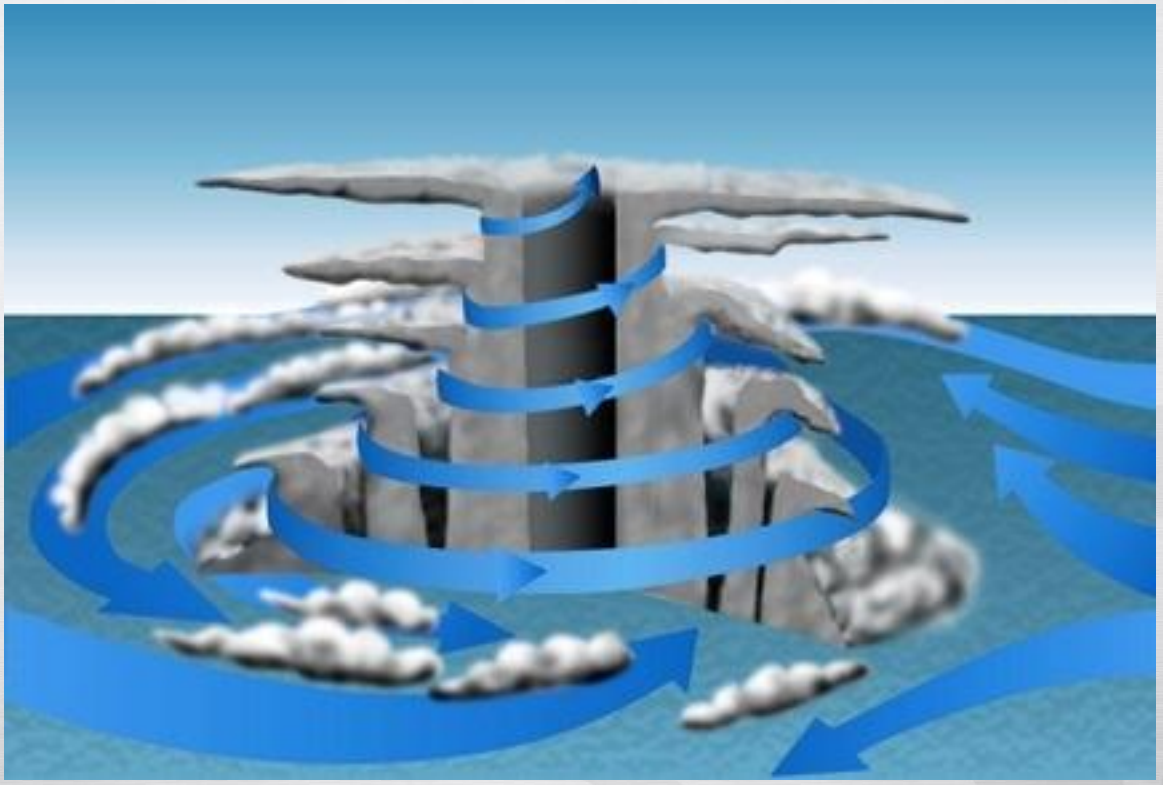

Routing Performance Observations

- On some platforms BGP takes a long time to re-converge
  - IPv4 Internet table several minutes
  - IPv6 Internet tables less than a minute

Routing Performance Observations

- You need a lot of expensive memory to hold the IPv4 Internet Tables in the forwarding plane

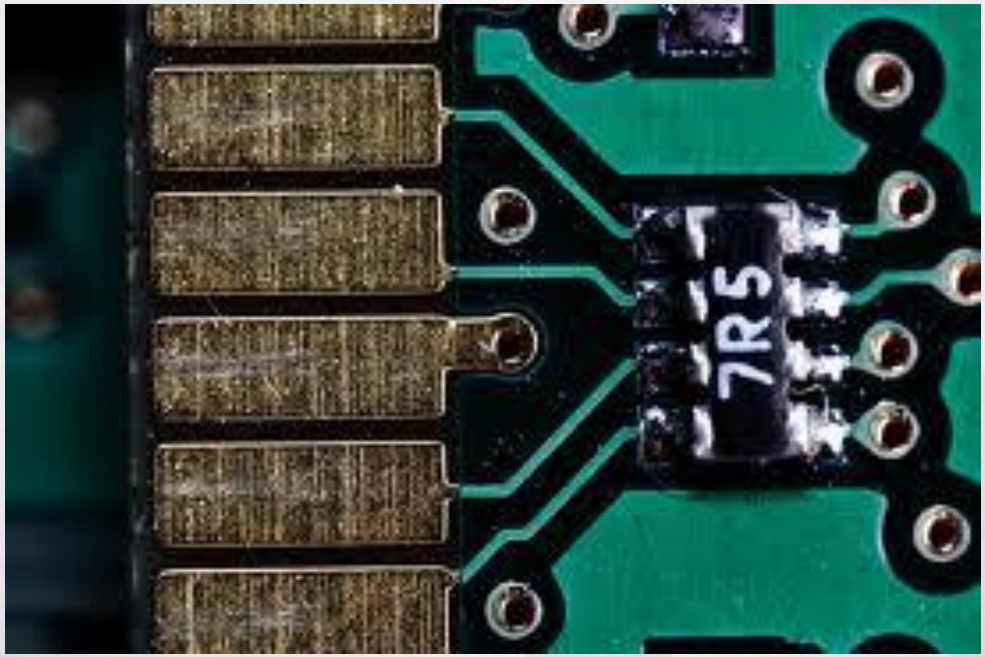

Routing Performance Observations

- IP addresses follow the topology limiting aggregation

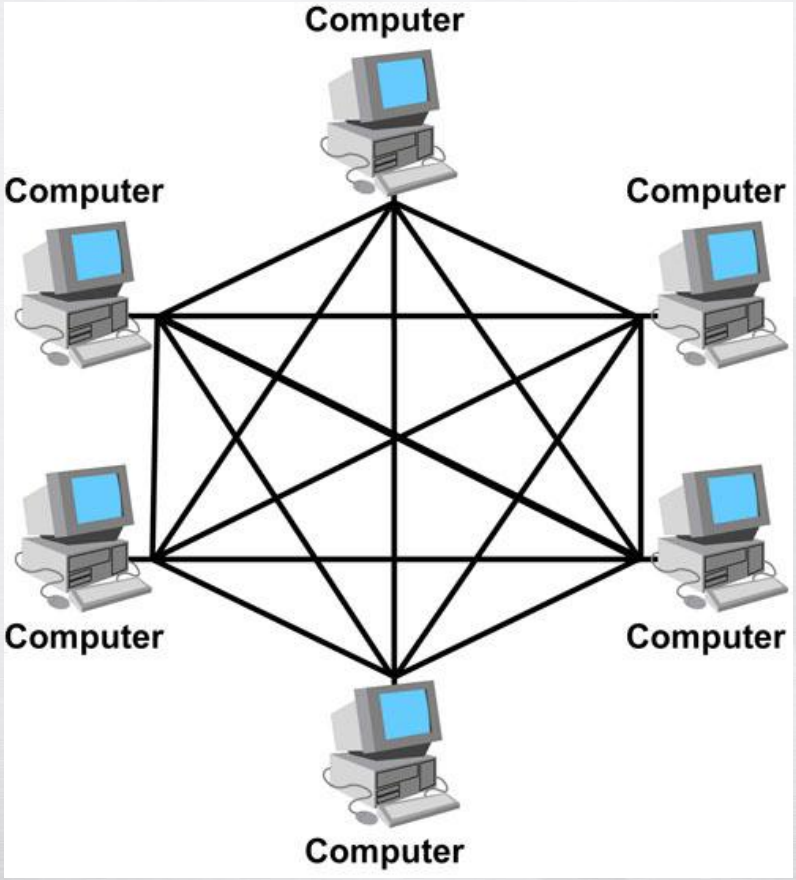

Common Platforms with Hardware Limitations

Cisco 6500

Juniper MX80

Cisco 6500 – Supervisor 720/PFC3 Architecture

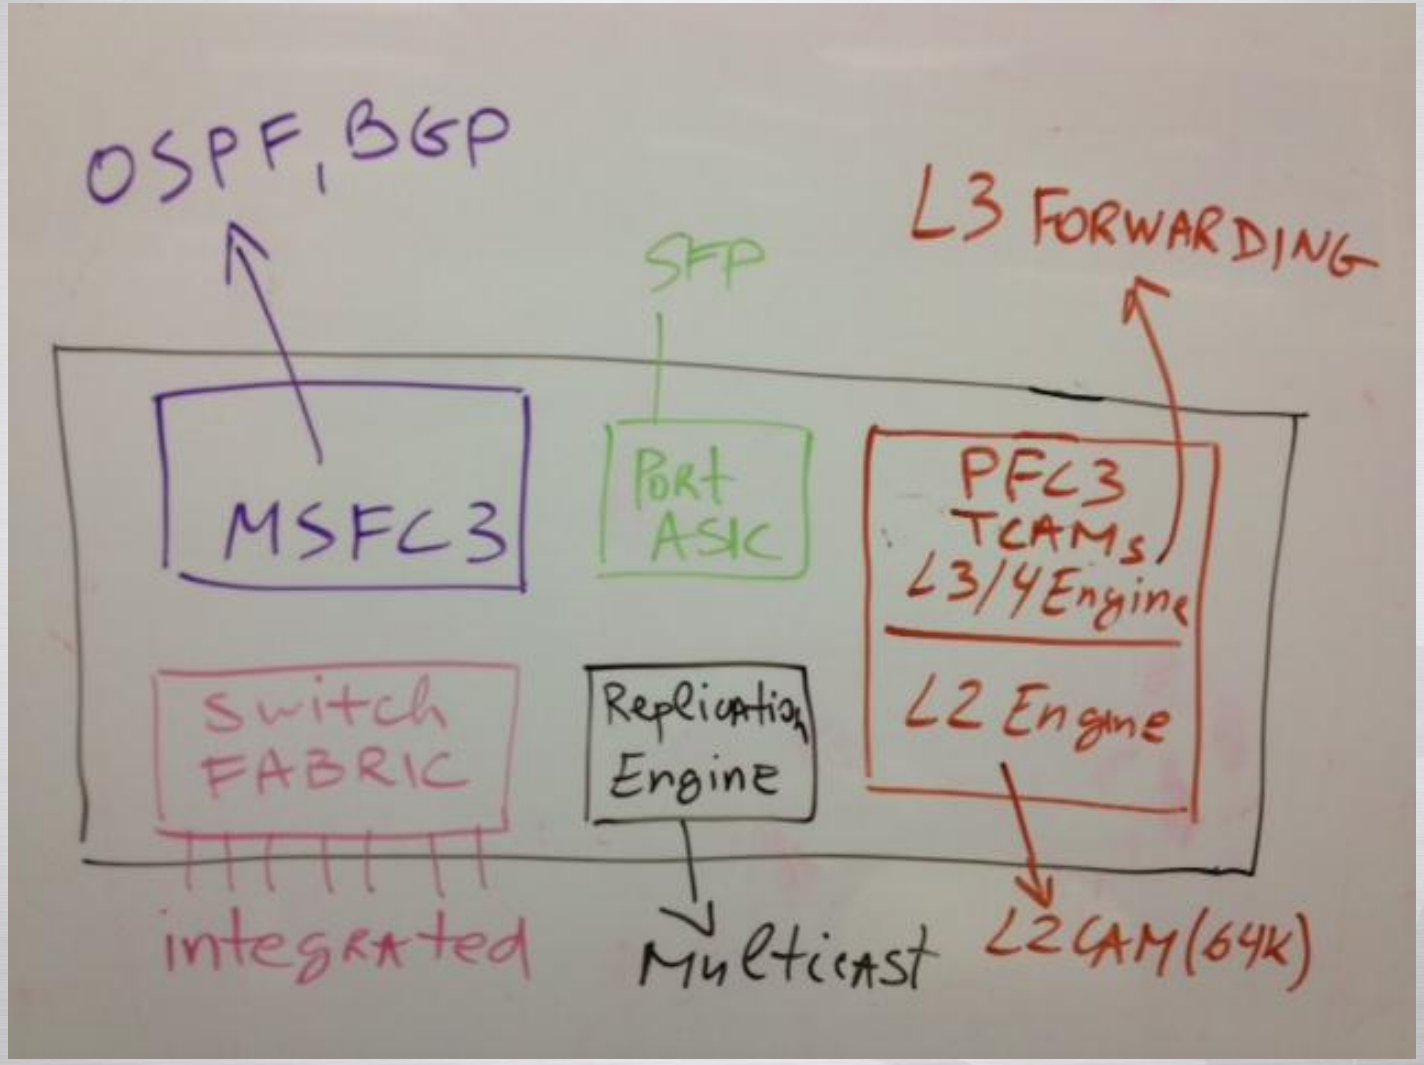

Cisco 6500 – Supervisor FIB TCAM Resources

- Default Settings

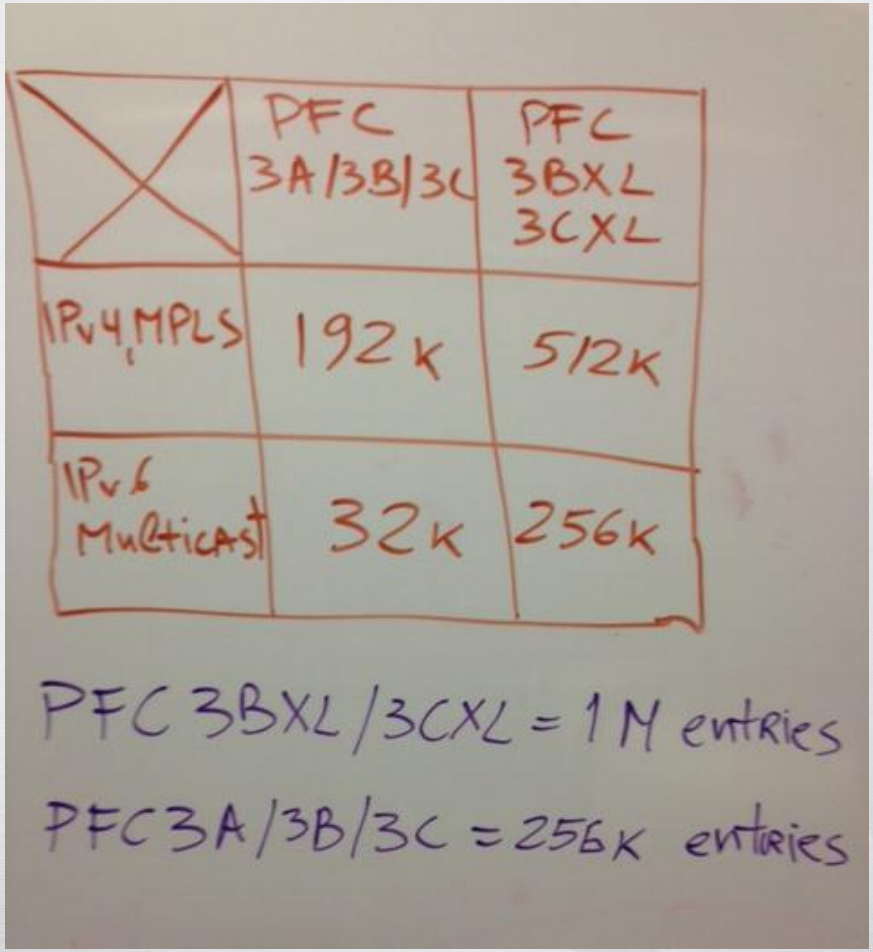

Cisco 6500 – Switching Paths

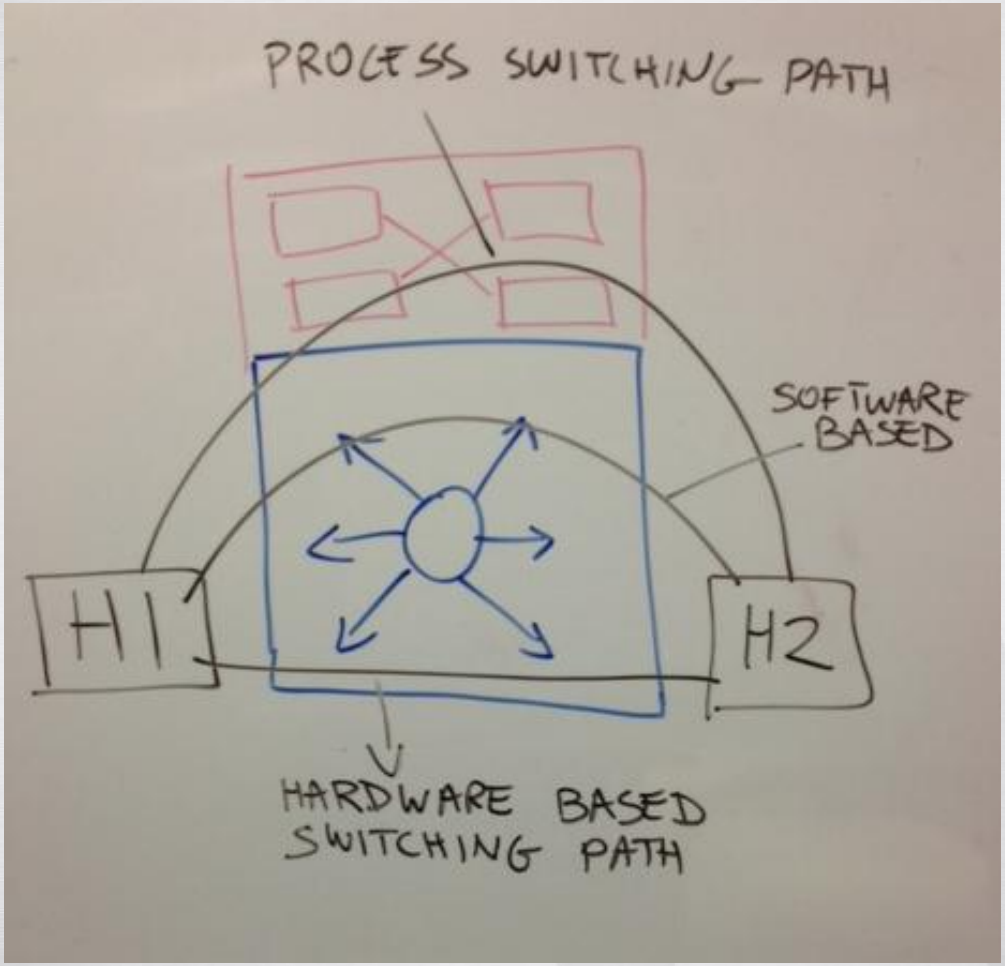

This is a polite way to say you are in trouble

- High CPU Utilization

%MLSCEF-SP-7-FIB_EXCEPTION: FIB TCAM exception, Some entries will be software switched

```
Cisco#show mls cef maximum
    FIB TCAM maximum routes :
    =========================
    IPv4 + MPLS                    - 192k (default)
    IPv6 + IP Multicast            - 32k (default)

cisco#show mls cef summary
    Total routes:                  194532
    IPv4 unicast routes:           194527
    IPv4 Multicast routes:         3
    MPLS routes:                   0
    IPv6 unicast routes:           2
```

Juniper MX80 Platform

- Juniper M80 Routing Engine (RE) and Packet Forwarding Engine (PFE)

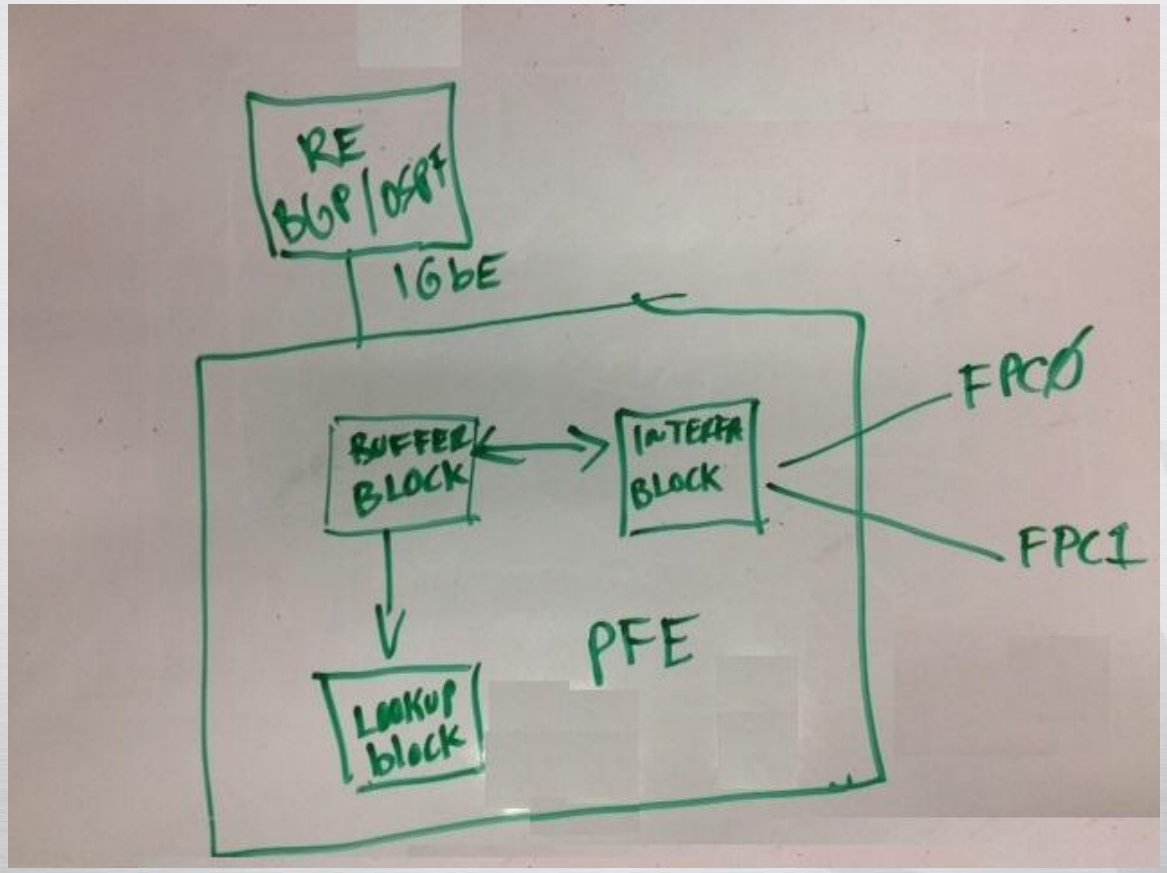

Juniper MX80 Platform

- 4,000,000 v4/v6 prefixes in RIB & 1,000,000 v4/v6 prefixes in FIB (numbers are based on IPv4 only)

```
MX80> show route summary
Autonomous system number: 1
  Router ID: 6.6.6.6

inet.0: 448956 destinations, 2766483 routes (448956 active, 1 holddown, 0 hidden)
          Direct:     3 routes,     3 active
           Local:     2 routes,     2 active
            OSPF:   1133 routes,   1125 active
             BGP: 2765338 routes, 447818 active
          Static:     8 routes,     8 active

inet6.0: 12600 destinations, 93785 routes (12600 active, 0 holddown, 0 hidden)
          Direct:     6 routes,     5 active
           Local:     4 routes,     4 active
           OSPF3:   156 routes,    156 active
             BGP:  93614 routes,  12430 active
          Static:     5 routes,     5 active
```

Juniper MX80 Platform

```
MX80> show chassis routing-engine
Routing Engine status:
  Temperature                 40 degrees C / 104 degrees F
  CPU temperature             51 degrees C / 123 degrees F
  DRAM                      2048 MB
  Memory utilization          86 percent  #*YIKES*
  CPU utilization:
    User                      1 percent
    Background                0 percent
    Kernel                    2 percent
    Interrupt                 0 percent
    Idle                     97 percent
  Model                       RE-MX80
  Start time                  2012-11-07 16:08:31 UTC
  Uptime                      154 days, 47 minutes, 18 seconds
  Last reboot reason          Router rebooted after a normal shutdown.
  Load averages:              1 minute   5 minute  15 minute
                               0.05       0.08      0.07
```

Juniper MX80 Platform

```
MX80> show chassis tfeb
TFEB status:
Slot 0 information:
  State                       Online
  Intake temperature            39 degrees C / 102 degrees F
  Exhaust temperature           59 degrees C / 138 degrees F
  CPU utilization              5 percent
  Interrupt utilization         0 percent
  Heap utilization             32 percent  #*80% could see pkt loss*
  Buffer utilization           13 percent
  Total CPU DRAM               1024 MB
  Start time:                  2012-11-07 16:11:15 UTC
  Uptime:                      154 days, 49 minutes, 4 seconds
```

Ways to Address the Growing Routing Tables Problem

- Ignore
  - Don’t care about the internet routing tables, use default route
- Improvise
  - Adjust TCAM allocations
- Selective Hearing
  - Filter de-aggregated subnets
- External Assistance
  - LISP
- Spend Cash
  - Hardware upgrades

Ignore

- Get rid of the tables, use default
  - Do you really need the full tables?

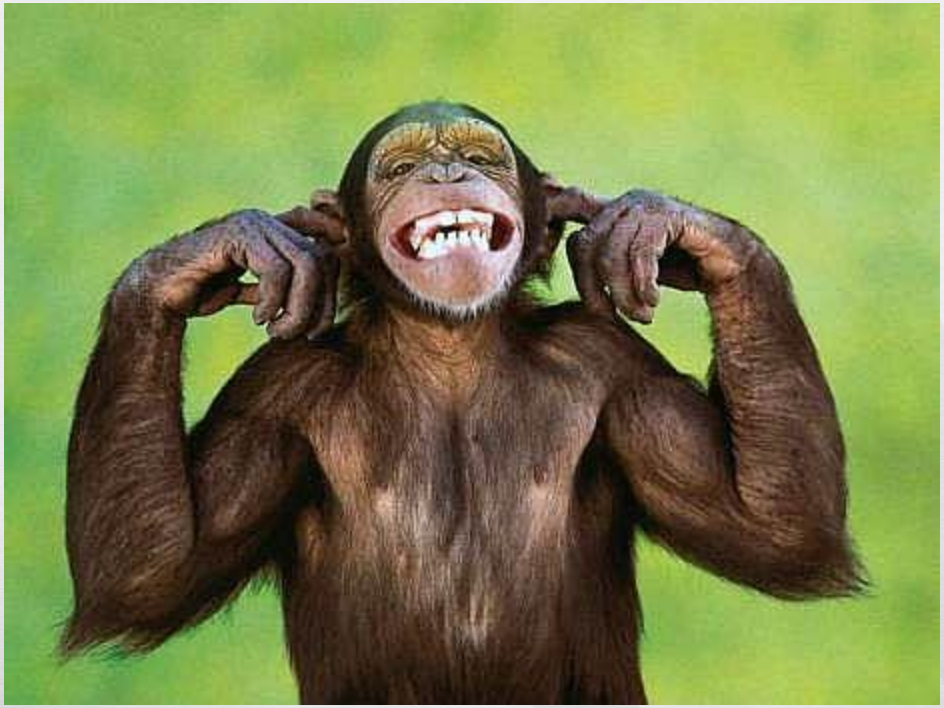

Improvise

- Change v4/v6 TCAM allocations
  - Cisco 6500 Supervisor 720/PFC3 example – manual change

```
mls cef maximum-routes ipv6 192
```

```
  Cisco6500#sh mls cef maximum-routes
FIB TCAM maximum routes :
=========================
Current :-
-------
 IPv4 + MPLS          - 624k (default)
 IPv6                 - 192k
 IP multicast         - 1k


  Cisco6500 #
```

- Juniper MX80 tables are dynamically allocated

Selective Hearing

- Filter Prefixes that are not important
  - Understand your business, your top destinations, are you doing any traffic management

External Assistance

- Locator/ID Separation Protocol (LISP)
  - Cisco is leading the standards effort

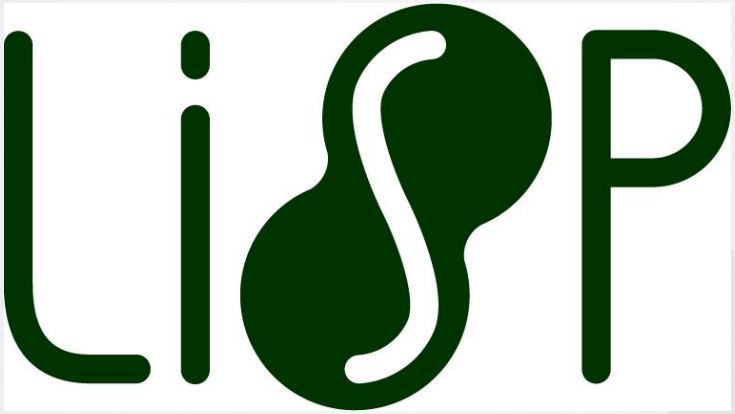

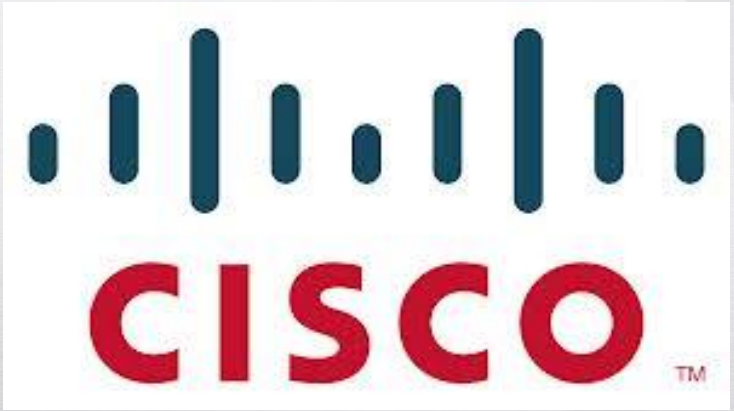

External Assistance

- LISP and DNS

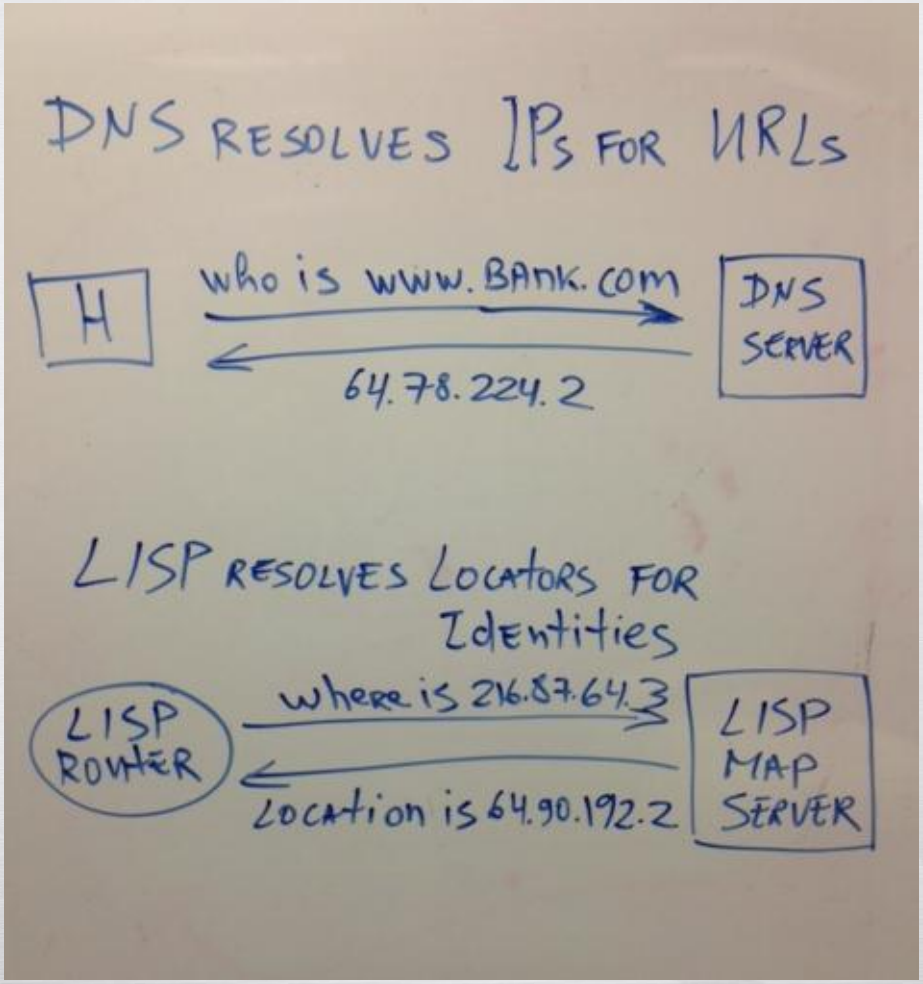

Spend Cash

- Hardware Upgrades

Cisco ASR9000

Juniper MX960

What do we think will happen with the tables?

Future of Internet Routing Tables - Prediction

- For at least five years we will continue to run dual-stack

- For most players on the Internet, carrying the full Internet routing tables will become unnecessary and it will not be done

Future of Internet Routing Tables - Prediction

- Hardware forwarding memory will become affordable

- If we follow the trend we will reach 1.2 million IPv4 routes and 100k IPv6 routes in 2023

Future of Internet Routing Tables - Prediction

- Replace prefixes based routing with ASN based routing
  - Revision of the packer header
  - MPLS on a global level

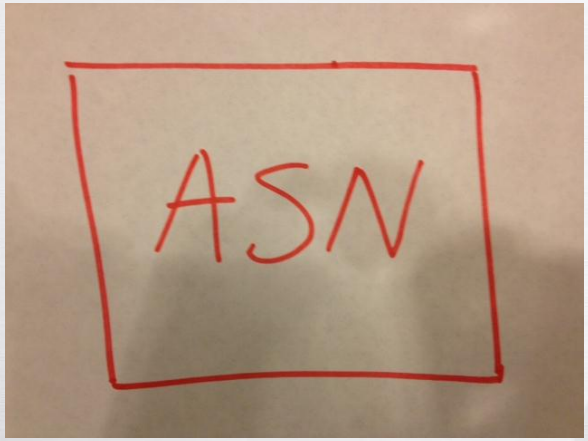

Future of Internet Routing Tables - Prediction

- IPv4 goes away
  - It is going to become very expensive to route
  - Destinations become unreachable because of routing table growth

The Present and Future of the Internet Routing Tables

- The present and future of the internet routing tables are undetermined
  - The problems are here to be solved
  - This is a chance to get it right
  - IPv6 is here to stay (no IPv8 is coming out)