

Rocky Mountain IPv6 Task Force



IPv6 experiences from a large content provider implementation

Ebben Aries
e@alterpath.net

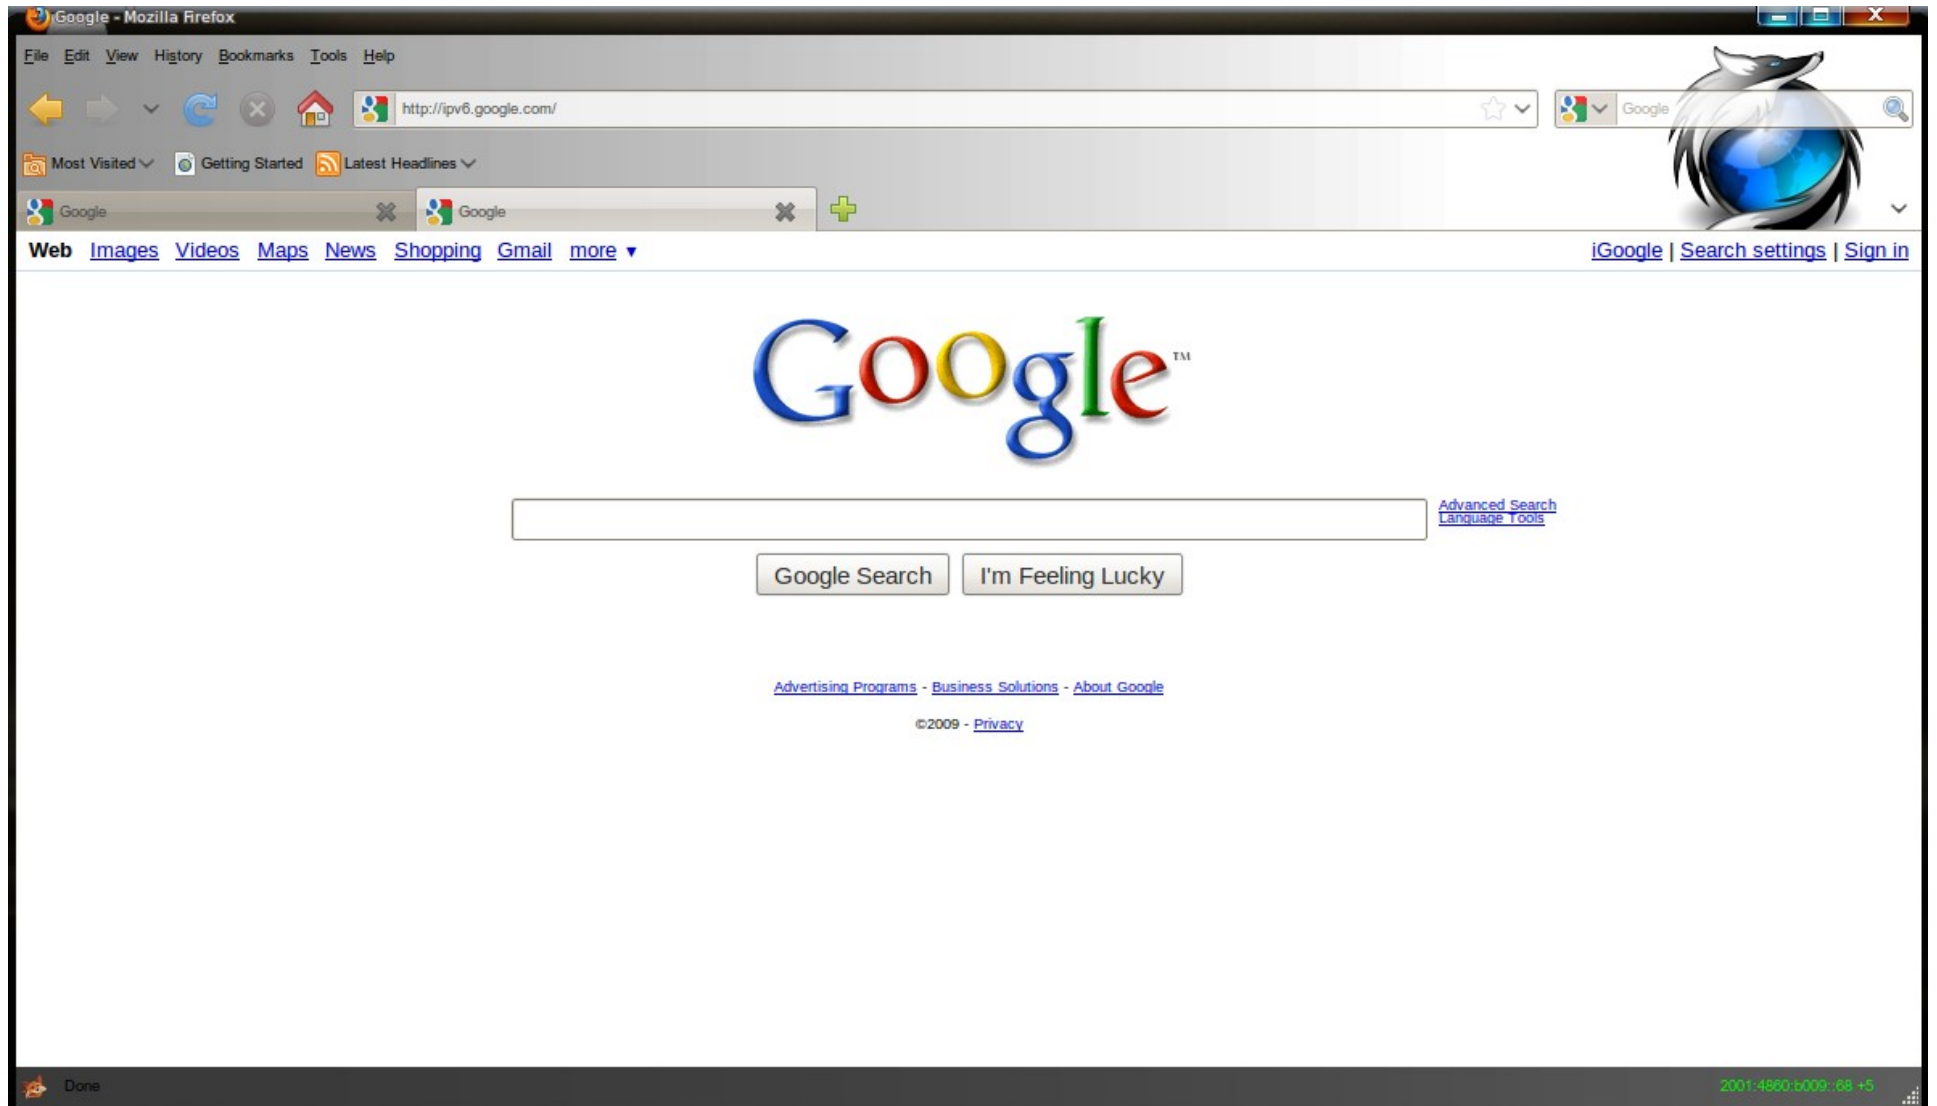
Agenda

1. Quick background/beginnings
2. ipv6.google.com
3. Google over IPv6
4. Operations and Issues encountered

Background and Beginnings

- Motivators behind rollout, Why IPv6, etc.... ?
 - Not going to get into much detail here – the usual reasons (IPv4 depletion, mobile devices, etc....)
 - http://www.ripe.net/ripe/meetings/ripe-57/presentations/Colitti-A_strategy_for_IPv6_adoption.Z8ri.pdf
 - http://www.ripe.net/ripe/meetings/ripe-56/presentations/Colitti-IPv6_at_Google.pdf
- Pilot IPv6 network started in 2007
 - Initially for developers to start development efforts
 - For most part was people's 20% projects initially
 - from 0 → production google over ipv6 (most google services) serving in ~18 mo.

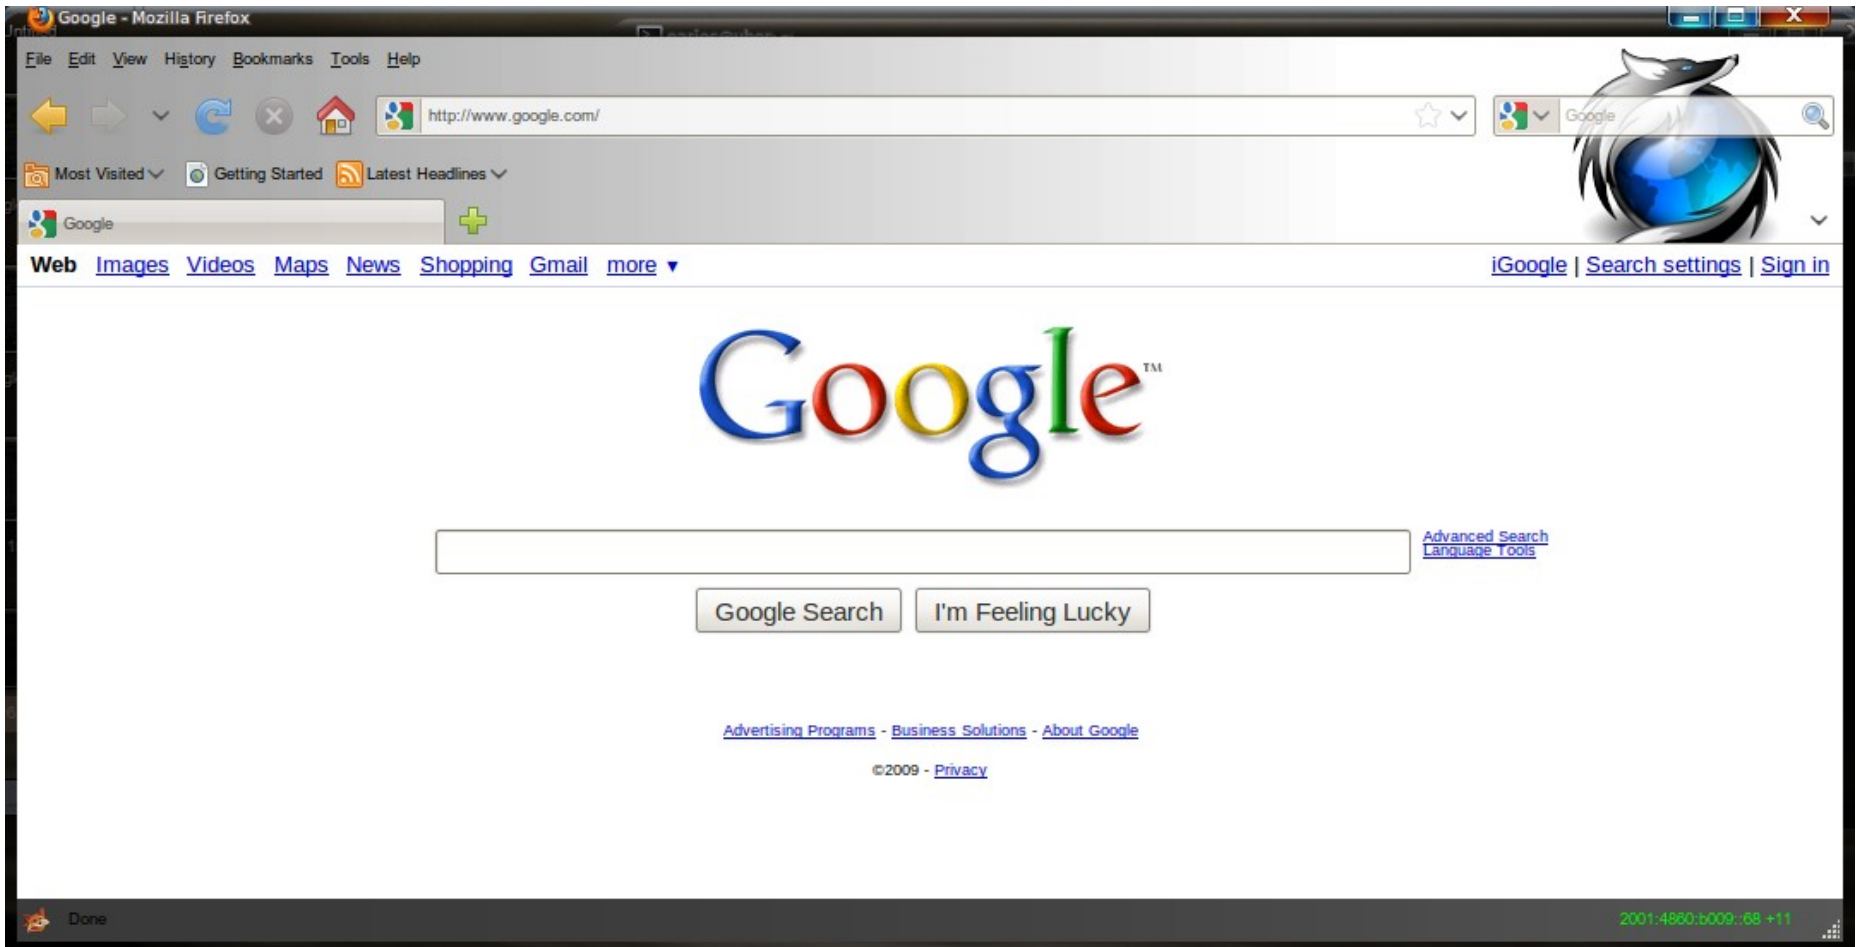
ipv6.google.com



ipv6.google.com (cont...)

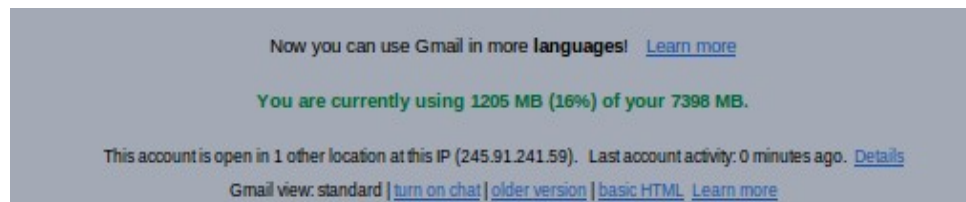
- Quick recap
 - First public service launch over IPv6 (front-end search)
- Initial launch at IETF71 during blackout session March 2008
 - only major search engine reachable during blackout
- Why not just turn on AAAA for all users, everything will work just fine right ?
 - still a portion (small) of users that will have issues – unacceptable
 - higher latency for many users to reach IPv6 networks, odds are tunneling is still happening across intermediate paths somewhere
 - lack of IPv6 geo-location databases resulting in lower performance than over IPv4
- Followed by ipv6.google.cn, ipv6.google.co.jp....

Google over IPv6



Google over IPv6 (cont...)

- <http://www.google.com/intl/en/ipv6/>
- Quick overview
 - Users now able to access handful of google services directly over IPv6 (search, email, maps, docs.....).... youtube soon
 - Trusted user program for early adopters that meet certain criteria
 - Participants must agree to terms to provide SLAs on IPv6, be willing to fix, report any bugs..... basically put forth effort in IPv6 adoption
- Since inception, no one has yet asked to be removed from the program – most say no one notices any difference, some have better user experience.



(Example of accessing gmail over IPv6)

Google over IPv6 (cont...)

```
rmv6tf - earies@uber: ~
earies@uber:~$ dig -t AAAA @74.82.42.42 www.google.com

; <<>> DiG 9.6.1-P1 <<>> -t AAAA @74.82.42.42 www.google.com
; (1 server found)
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 33741
;; flags: qr rd ra; QUERY: 1, ANSWER: 7, AUTHORITY: 0, ADDITIONAL: 0

;; QUESTION SECTION:
;www.google.com.                IN      AAAA

;; ANSWER SECTION:
www.google.com.                0       IN      CNAME   www.l.google.com.
www.l.google.com.              300     IN      AAAA    2001:4860:b009::63
www.l.google.com.              300     IN      AAAA    2001:4860:b009::6a
www.l.google.com.              300     IN      AAAA    2001:4860:b009::69
www.l.google.com.              300     IN      AAAA    2001:4860:b009::93
www.l.google.com.              300     IN      AAAA    2001:4860:b009::67
www.l.google.com.              300     IN      AAAA    2001:4860:b009::68

;; Query time: 82 msec
;; SERVER: 74.82.42.42#53(74.82.42.42)
;; WHEN: Sat Dec 5 20:19:04 2009
;; MSG SIZE rcvd: 220

earies@uber:~$
```

(Example of AAAA serving www.google.com from trusted user network over IPv4 request)

Google over IPv6 (cont...)

```
rmv6tf - earies@uber: ~
earies@uber:~$ dig -t AAAA @2001:470:20::2 www.google.com

; <<>> DiG 9.6.1-P1 <<>> -t AAAA @2001:470:20::2 www.google.com
; (1 server found)
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 23264
;; flags: qr rd ra; QUERY: 1, ANSWER: 7, AUTHORITY: 0, ADDITIONAL: 0

;; QUESTION SECTION:
;www.google.com.                IN      AAAA

;; ANSWER SECTION:
www.google.com.                604582  IN      CNAME   www.l.google.com.
www.l.google.com.              105     IN      AAAA    2001:4860:b009::6a
www.l.google.com.              105     IN      AAAA    2001:4860:b009::67
www.l.google.com.              105     IN      AAAA    2001:4860:b009::63
www.l.google.com.              105     IN      AAAA    2001:4860:b009::69
www.l.google.com.              105     IN      AAAA    2001:4860:b009::93
www.l.google.com.              105     IN      AAAA    2001:4860:b009::68

;; Query time: 76 msec
;; SERVER: 2001:470:20::2#53(2001:470:20::2)
;; WHEN: Sat Dec 5 20:22:19 2009
;; MSG SIZE rcvd: 220

earies@uber:~$
```

(Example of AAAA serving www.google.com from trusted user network over IPv6 request)

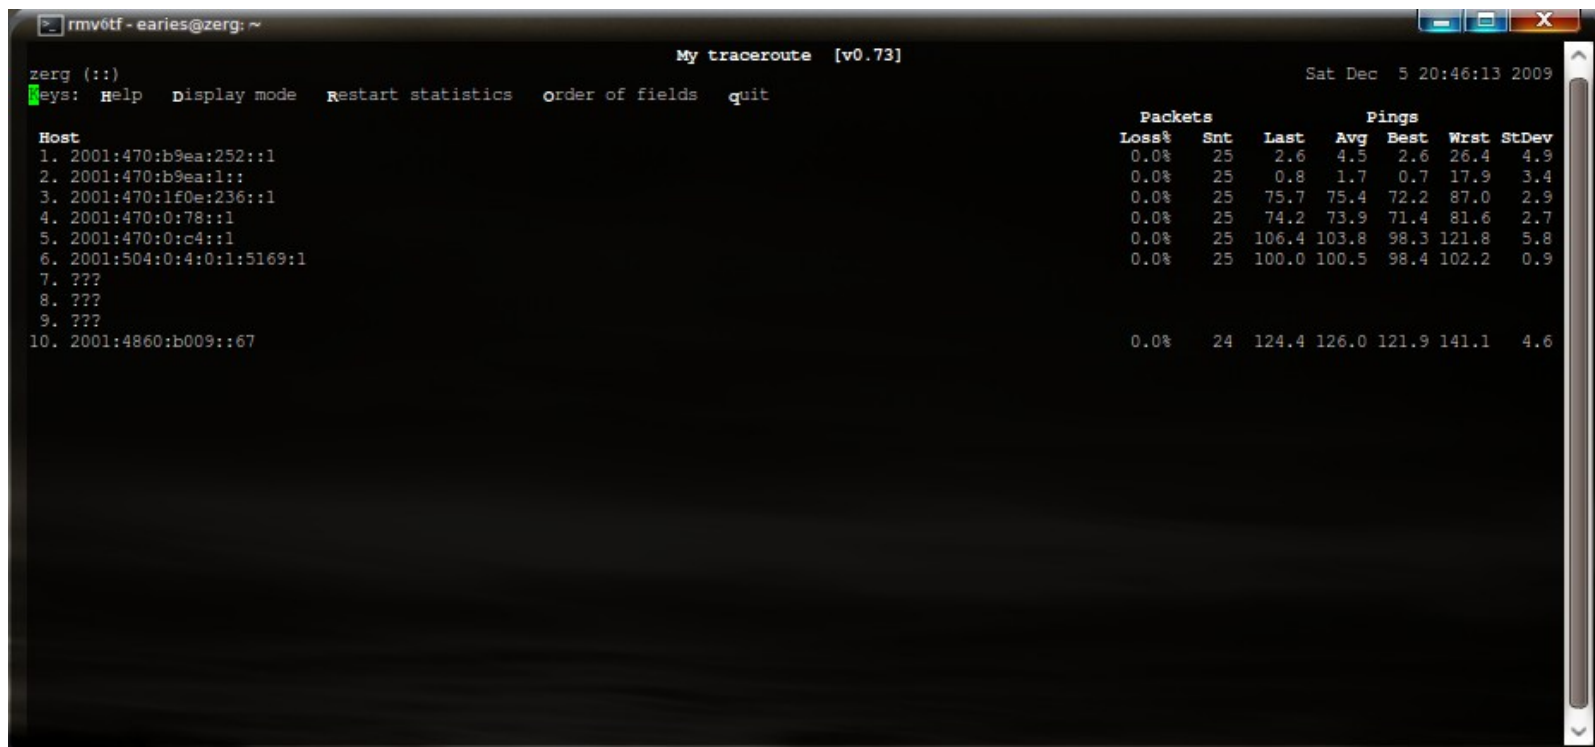
Operations

- Readiness
 - IGP Support? (OSPFv3 vs. ISIS)
 - Hardware/Software support
 - Operations training (eat your own dogfood)
 - Tunneling mechanisms (6PE, 6in4, etc...)
 - Design concerns
 - Address Allocation
 - NMS and tools support
 - etc....
- No transit – only peering
 - Peer w/ user networks directly, easier to get issues fixed
- Client Access
 - privacy-extensions

Operations (cont...)

The problem with 6PE and tunneling in general

- Intermediate control planes do not have same visibility
- When tunnels/LSPs go down, blackholing of traffic occurs
- Tunnel setup mechanisms not entirely scalable
- Traceroute interpretations as below (TTL hiding mechanisms can fix)
- Need for better TE in IPv6 (RSVP signalling over IPv6)



```
rmv6tf - earies@zerg: ~
zerg (::)
My traceroute [v0.73]
Sat Dec 5 20:46:13 2009
Keys: Help Display mode Restart statistics order of fields quit

Host
1. 2001:470:b9ea:252::1
2. 2001:470:b9ea:1::
3. 2001:470:1f0e:236::1
4. 2001:470:0:78::1
5. 2001:470:0:c4::1
6. 2001:504:0:4:0:1:5169:1
7. ???
8. ???
9. ???
10. 2001:4860:b009::67

Packets
Loss% Snt Last Avg Best Wrst StDev
0.0% 25 2.6 4.5 2.6 26.4 4.9
0.0% 25 0.8 1.7 0.7 17.9 3.4
0.0% 25 75.7 75.4 72.2 87.0 2.9
0.0% 25 74.2 73.9 71.4 81.6 2.7
0.0% 25 106.4 103.8 98.3 121.8 5.8
0.0% 25 100.0 100.5 98.4 102.2 0.9

Pings
0.0% 24 124.4 126.0 121.9 141.1 4.6
```

Operations (cont...)

```
rmv6tf - earies@uber: ~
earies@route-server> show route terse table inet6.0 aspath-regex ".* 15169$"

inet6.0: 2207 destinations, 2209 routes (2165 active, 0 holddown, 42 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1  Metric 2  Next hop          AS path
* 2001:4860::/32    B 170   1000      >2001:470:b9ea:1:: 65420 6939 15169 I
* 2404:6800::/32    B 170   1000      >2001:470:b9ea:1:: 65420 6939 15169 I
* 2620:0:1001::/48  B 170   1000      >2001:470:b9ea:1:: 65420 6939 15169 I
* 2620:0:1002::/48  B 170   1000      >2001:470:b9ea:1:: 65420 6939 15169 I
* 2620:0:1003::/48  B 170   1000      >2001:470:b9ea:1:: 65420 6939 15169 I
* 2620:0:1004::/48  B 170   1000      >2001:470:b9ea:1:: 65420 6939 15169 I
* 2620:0:1006::/48  B 170   1000      >2001:470:b9ea:1:: 65420 6939 15169 I
* 2620:0:1007::/48  B 170   1000      >2001:470:b9ea:1:: 65420 6939 15169 I
* 2620:0:1008::/48  B 170   1000      >2001:470:b9ea:1:: 65420 6939 15169 I
* 2620:0:101e::/48  B 170   1000      >2001:470:b9ea:1:: 65420 6939 15169 I
* 2620:0:101f::/48  B 170   1000      >2001:470:b9ea:1:: 65420 6939 15169 I
* 2620:0:1040::/48  B 170   1000      >2001:470:b9ea:1:: 65420 6939 15169 I
* 2620:0:1042::/48  B 170   1000      >2001:470:b9ea:1:: 65420 6939 15169 I
* 2620:0:105d::/48  B 170   1000      >2001:470:b9ea:1:: 65420 6939 15169 I
* 2620:0:105e::/48  B 170   1000      >2001:470:b9ea:1:: 65420 6939 15169 I
* 2620:0:105f::/48  B 170   1000      >2001:470:b9ea:1:: 65420 6939 15169 I
* 2620:0:1080::/48  B 170   1000      >2001:470:b9ea:1:: 65420 6939 15169 I
* 2620:0:1081::/48  B 170   1000      >2001:470:b9ea:1:: 65420 6939 15169 I
* 2620:0:1082::/48  B 170   1000      >2001:470:b9ea:1:: 65420 6939 15169 I
* 2620:0:10b0::/48  B 170   1000      >2001:470:b9ea:1:: 65420 6939 15169 I
* 2620:0:10b1::/48  B 170   1000      >2001:470:b9ea:1:: 65420 6939 15169 I
* 2a00:1450::/32    B 170   1000      >2001:470:b9ea:1:: 65420 6939 15169 I

earies@route-server>
```

(ARIN, APNIC, RIPE PI Allocations as visible via AS6939)

A Few Issues Encountered

- radix-lookup bug resulting in crashes eve of launch of ipv6.google.com (consider dedicated devices...)
- l2-port security issue due to header bit shifts triggered at AMS-IX
- load-balancer memory leaks
- Software/Vendor support
- Netflow v9 support – beware hardware/software support
- DAD – manual intervention after detection in some cases
- separate licensing
- RSVP signaling support over IPv6

Further Reading

- http://www.ripe.net/ripe/meetings/ripe-57/presentations/Colitti-A_strategy_for_IPv6_adoption.Z8ri.pdf
- http://www.ripe.net/ripe/meetings/ripe-56/presentations/Colitti-IPv6_at_Google.pdf

Google over IPv6 program

- <http://www.google.com/intl/en/ipv6/>

Misc Links

- <http://www.ipv6actnow.org/>

Thanks

- Lorenzo Colitti, Erik Kline, and the rest of the googlers involved in anything and everything IPv6

Rocky Mountain IPv6 Task Force



Questions?

Ebben Aries
e@alterpath.net