



XI'AN, CHINA
20 - 30 August 2013

ULA in the wild

APNIC



APNIC **36**
CONFERENCE



XI'AN, CHINA
20 - 30 August 2013

George Michaelson

Research Scientist



ULA in the wild

ULAs defined

- IANA allocation **fc00::/7**
- RFC4193, 2005
 - “approximate counterpart of RFC1918 for IPv6”
 - Not intended to be globally routed
- Two subforms:
 - fc00::/8 “centrally assigned”
 - No registry currently operating formally
 - fd00::/8 “locally assigned”
 - Random throw against time, EUI.164 MAC address
 - Goal: unique /48 unlikely to collide with any other consumer (future net mergers, local routing)

Not intended to be globally routed

- Do they leak?
 - Are there places we can see ULA as src address in IPv6 packets on the wire, outside the local context of use?
- Does knowledge of them leak?
 - Are there places we can see ULA referenced as payload in some other transaction?

Not intended to be globally routed

- Do they leak?
 - Are there places we can see ULA as src address in IPv6 packets on the wire, outside the local context of use?
 - **YES**
- Does knowledge of them leak?
 - Are there places we can see ULA referenced as payload in some other transaction?
 - **YES**

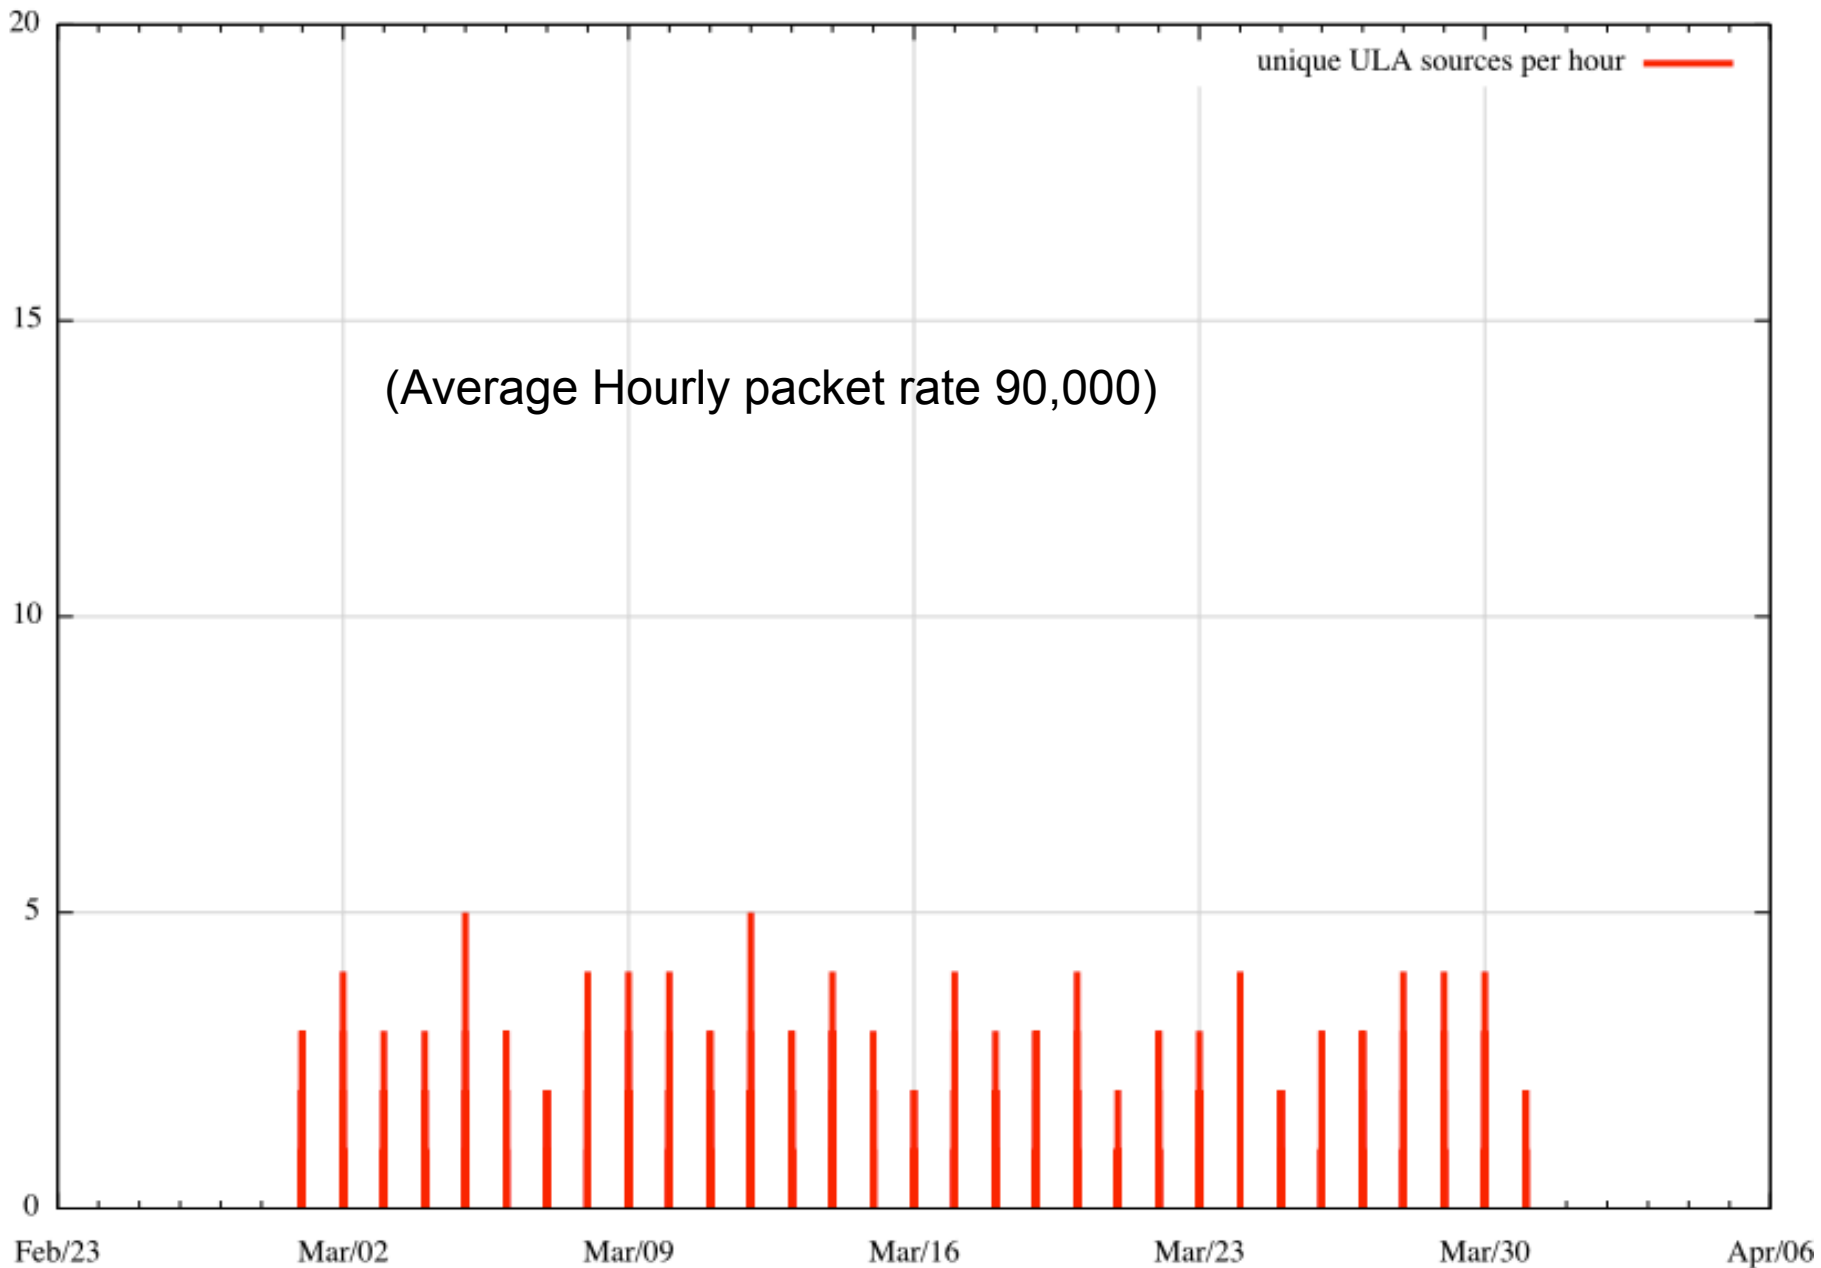
Not intended to be globally routed

- Do they leak?
 - Are there places we can see ULA as src address in IPv6 packets on the wire, outside the local context of use?
 - **YES** *but a very little compared to rfc1918*
- Does knowledge of them leak?
 - Are there places we can see ULA referenced as payload in some other transaction?
 - **YES** *quite a lot, and widely distributed*

Do they leak?

Do they leak into routing?

- Capture of 2400::- Comb hourly pcap files for unique ULA instances



But...

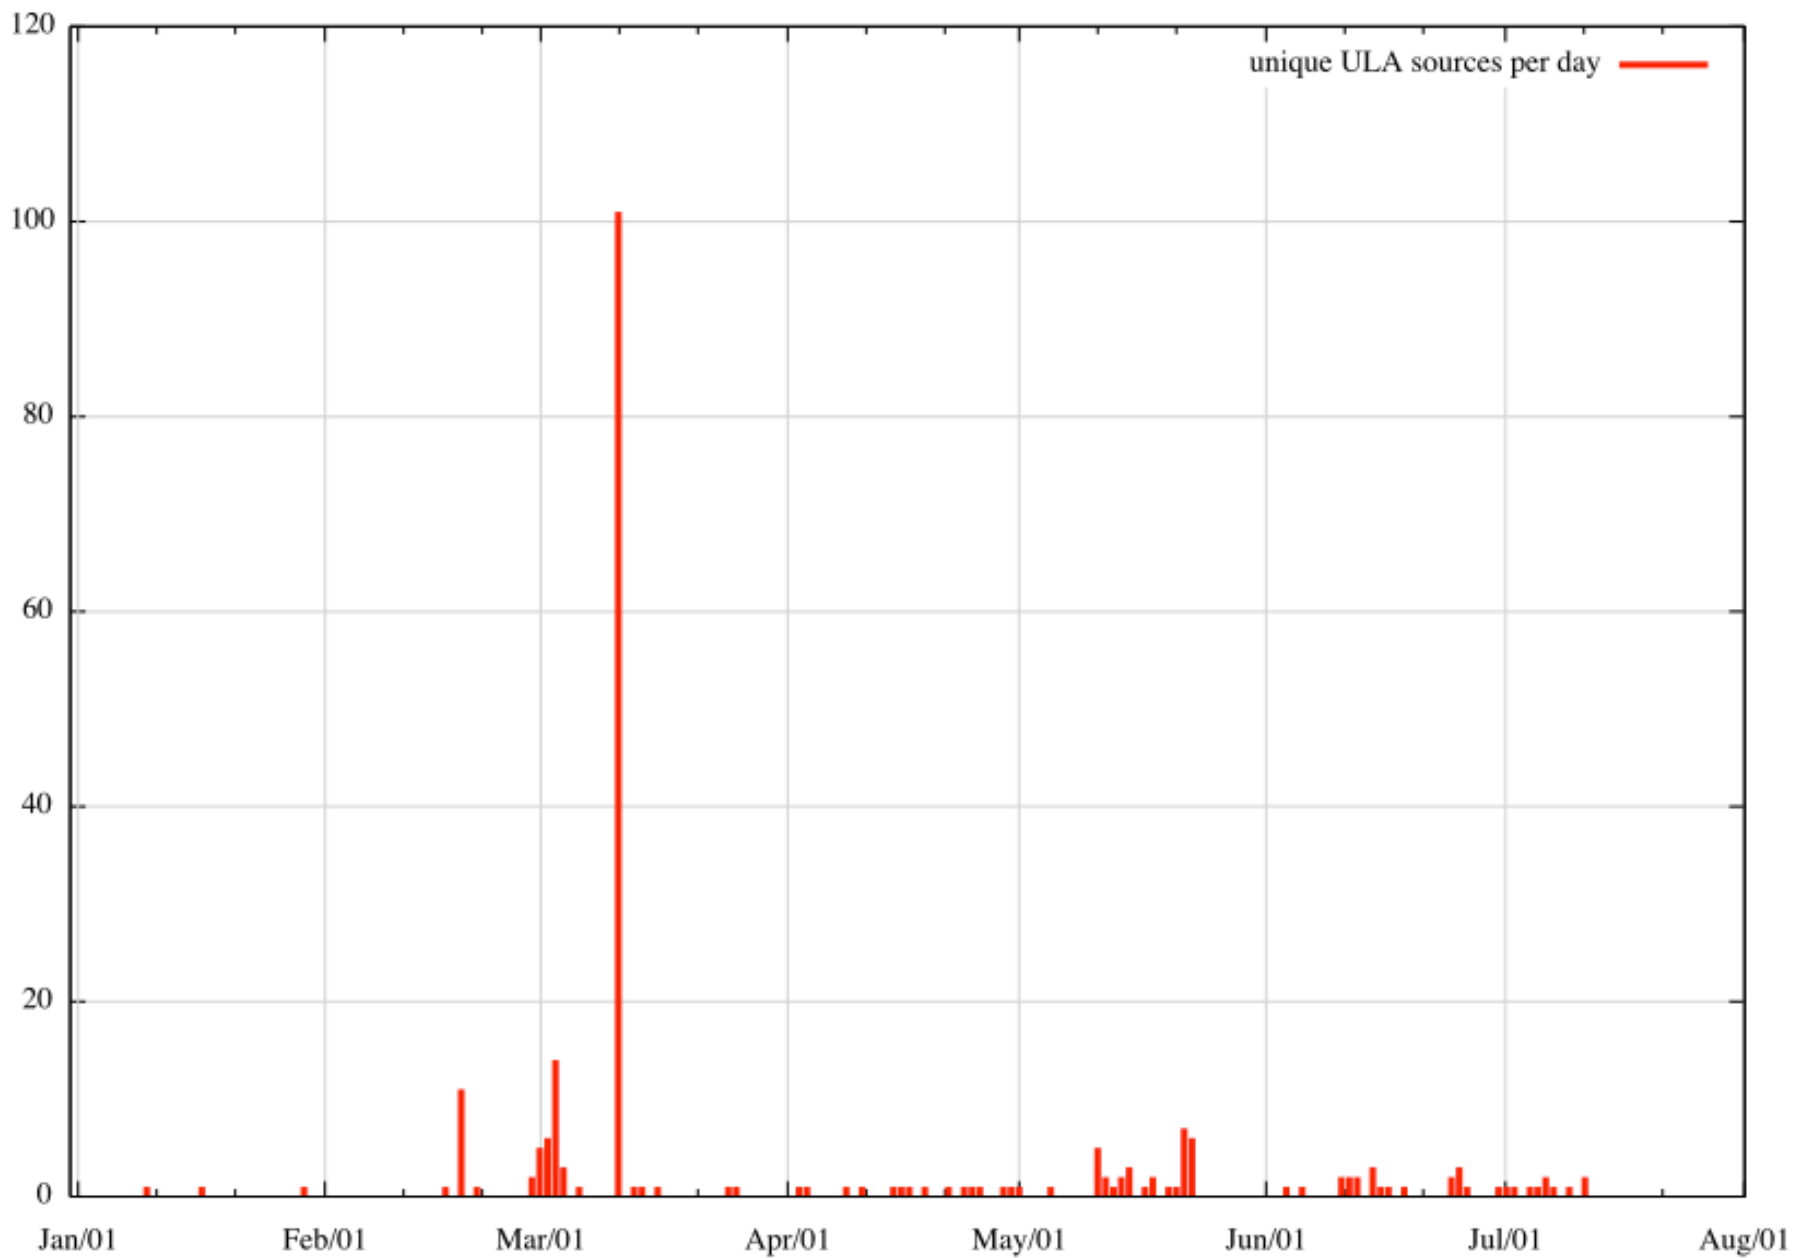
- That was ‘coerced’ packets with a covering announcement
 - They wouldn’t normally have wound up in the public routing view
- They almost universally relate to one ISP in Indonesia, and are therefore not widespread

Do they leak into routing?

- In-addr.arpa DNS delegation
 - One of 6 listed NS for in-addr.arpa, ip6.arpa
 - portspan 24/7 used to feed DiTL, DSC graphs
 - Count/collect unique src, dst per 24h
 - 1) do we see any ULA as src addresses?

Do they leak into routing?

- In-addr.arpa DNS delegation
 - One of 6 listed NS for in-addr.arpa, ip6.arpa
 - portspan 24/7 used to feed DiTL, DSC graphs
 - Count/collect unique src, dst per 24h
 - 1) do we see any ULA as src addresses?
 - YES



Low level leakage

- 1-2 ULA prefixes seen per day as source
 - Compared to 50,000 unique IPv6 sources of query
- Occasional peaks
 - Routing slips, acquired default?
- Low level background noise, few DNS questions per src

Does knowledge of them leak?

- How about the 'payload' of DNS reverse questions?

Does knowledge of them leak?

- How about the 'payload' of DNS reverse questions?
- Hang on

Does knowledge of them leak?

- How about the 'payload' of DNS reverse questions?
- Hang on
 - Why would anyone do reverse-lookup for IPv6 let alone ULA?

SMTP Received-Via

SMTP Received-Via

```
Received: from ia-mailgw.apnic.net (ia-mailgw.apnic.net. [2001:dd8:a:3::243])  
    by mx.google.com with SMTP id wn4si1461945pbc.175.2013.06.20.18.02.16  
    for <ggmichaelson@gmail.com>;  
    Thu, 20 Jun 2013 18:02:18 -0700 (PDT)
```

- Every email received over IPv6 by an SMTP server appears to add a trace line which depends on gethostbyaddr() call
- If your local SMTP is bound over IPv6 and you are using ULA, its going to do a DNS call.
- If you haven't delegated ULA locally in DNS, it goes out into the wide blue yonder
- Only a 'for instance'
 - SSHD, LPR, other daemons may well log, as may dhcpv6 or ACLs or ...

Reverse DNS, one of 6 NS of in-addr.arpa and ip6.arpa

- 350,000,000 queries per day, on the ip6.arpa and in-addr-arpa Nameserver instance we run.
- 500,000 queries into ULA space are currently being seen per day
 - Down from 1,000,000 back in 2011
 - 0.14%
- By contrast global unicast IPv6 query load is now consistently above 1,000,000 queries per day.
- Queries for ULA from > resolvers in 4000 distinct ASN

Top 25 ULA query sources,ip6.arpa

| ASN | Name | ASN | Name |
|------|---|-------|--|
| 174 | COGENT Cogent/PSI | 4802 | ASN-IINET iiNet Limited |
| 209 | ASN-QWEST-US NOVARTIS-DMZ-US | 6327 | SHAW - Shaw Communications Inc. |
| 577 | BACOM - Bell Canada | 6621 | HNS-DIRECPC - Hughes Network Systems |
| 701 | UUNET - MCI Communications Services Inc. d/b/a Verizon Business | 6830 | LGI-UPC Liberty Global Operations B.V. |
| 786 | JANET The JNT Association | 6939 | HURRICANE - Hurricane Electric Inc. |
| 1221 | ASN-TELSTRA Telstra Pty Ltd | 7018 | ATT-INTERNET4 - AT&T Services Inc. |
| 2516 | KDDI KDDI CORPORATION | 7132 | SBIS-AS AS for SBIS-AS |
| 2907 | SINET-AS National Institute of Informatics | 7922 | COMCAST-7922 - Comcast Cable Communications Inc. |
| 3320 | DTAG Deutsche Telekom AG | 9299 | IPG-AS-AP Philippine Long Distance Telephone Company |
| 3356 | LEVEL3 Level 3 Communications | 15169 | GOOGLE - Google Inc. |
| 3462 | HINET Data Communication Business Group | 17506 | UCOM UCOM Corp. |
| 4134 | CHINANET-BACKBONE No.31 Jin-rong Street | 22773 | ASN-CXA-ALL-CCI-22773-RDC - Cox Communications Inc. |
| 4713 | OCN NTT Communications Corporation | | |

Yea but..

- Some of these sources are obviously providing DNS service (8.8.8.8) and its likely they are the visible 'front' DNS query for a back-end system.
 - But it still implies quite widespread use of ULA behind these ASN
- Over 4000 ASN seen with some level of ULA query.

Both kinds of ULA

Country and Western

- 870,000 ULA sample

A little bit Country

- 870,000 ULA sample
 - 8400 in 'centrally managed' space (0.96%)

Mainly Western

- 870,000 ULA sample
 - 8400 in 'centrally managed' space (0.96%)
 - Remainder in 'self assigned' space (99%)

Country and Western

- 870,000 ULA sample
 - 8400 in ‘centrally managed’ space (0.96%)
 - Remainder in ‘self assigned’ space (99%)
- How ‘wisely’ do people consume this space?

Centrally managed fc00::/8

- Of 8400 in 'centrally managed' space

Centrally managed fc00::/8

- Of 8400 in 'centrally managed' space
 - Over 7000 in fc00::
 - There are $2^{32} / 48$ in this /8...

Centrally managed fc00::/8

- Of 8400 in 'centrally managed' space
 - Over 7000 in fc00::
 - There are $2^{32} / 48$ in this /8...
 - Over 2000 in fc00:0000:0000:
 - 'I didn't do any random throw. I just took the bottom'

Centrally managed fc00::/8

- Of 8400 in 'centrally managed' space
 - Over 7000 in fc00::
 - There are $2^{32} / 48$ in this /8...
 - Over 2000 in fc00:0000:0000:
 - 'I didn't do any random throw. I just took the bottom'
 - When pruned to the /48 equivalent, there appear to be around 250 distinct ULA from this sample.

This is not very wise

- There is no central registry function at this time
- Usage includes fc00:1111 and fc00:1234
- suggests that the choice of /48 is not driven by a strong registry process.
 - more likely is either self-assigned, and so is at risk of colliding
 - or else is a ‘first come first served’ registry service which offers uniqueness within the constraints of how people ask for a ULA at that time.

Self Assigned fd00::/8

- 167,000 unique /48 in the sample

Self Assigned fd00::/8

- 167,000 unique /48 in the sample

| Prefix | Count | Prefix | Count |
|----------------|--------|----------------|-------|
| fd00:6587:52d7 | 198825 | fdf1:6dfc:0828 | 361 |
| fdb2:2c26:f4e4 | 10867 | fdfe:7dc7:2e19 | 337 |
| fd00:0000:0000 | 8360 | fd7f:29be:fce4 | 334 |
| fd8c:215d:178e | 5597 | fdfe:1729:7999 | 333 |
| fdbd:0000:0000 | 4540 | fd37:3dd1:7688 | 330 |
| fd0d:edc3:e12a | 948 | fde8:e968:28e7 | 329 |
| fd1e:6d3c:942b | 684 | fd55:faaf:e1ab | 318 |
| fdc2:c837:3301 | 591 | fdb6:4c6e:d6fa | 309 |
| fd5e:35a9:696b | 470 | fd8f:8349:a712 | 300 |
| fdf1:a35e:8d33 | 469 | fd3d:848e:24be | 294 |
| fddb:7f1c:d199 | 407 | fd14:fad0:2c06 | 289 |
| fd29:41d0:f8c9 | 375 | fdbb:1cb5:bb90 | 285 |
| fd25:81be:cd4f | 363 | | |

Self Assigned fd00::/8

- Two naughty cases, with high levels of usage.
- Majority case is to use the random assignment method
- High levels of usage being seen
- Informal registry service available at sixxs

Seen any from sixxs?

- ‘spin the wheel service’ for your EUI.164
 - <http://www.sixxs.net/tools/grh/ula/>
- 3000 ULA listed on their ‘whois’ service
 - 20 seen in this capture. Top 3:

| ULA prefix | Who | Count from 870,000 |
|----------------|----------------|--------------------|
| fd8c:215d:178e | IBM | 5597 |
| fd0d:edc3:e12a | Hughes SE Lab | 948 |
| fde9:7537:6abe | Techno hosting | 58 |

Seen any collisions?

Seen any collisions?

- No

Seen any collisions?

- **No** but contextually, hard to prove because the ASN seen asking the question may vary but its no indication it's a different entity using the same ULA
- Observing the use of the algorithm, it looks unlikely at this level of activity
 - Simple check: 0/1 bias in assigned /48
 - Basically 50:50, slight bias to ones may be from date element in the algorithm.

How do the /128 assign?

- over half the ULA seen appear to be using ff:fe structured MAC addresses for the /128
- By comparison, use of non-privacy mode in global unicast has dropped off significantly
 - Either the processes behind ULA don't enable temporary/privacy mode
 - Or the time when ULA intrude into gethostbyname() the address selected isn't privacy mode
 - Or Privacy mode hasn't spun up yet when ULA is used

Summary

Summary

- ULA usage appears widespread geographically
- ULA usage appears to be stable
- Some 'unwise' use of fc00::/8 and fd00::/8 but most assignments honour the unique/random assignment model
- Very little leakage into global routing in this measurement
 - 1-2 instances per DAY seen in 50,000 unique IPv6

Hang on.....

Whats causing all these ULA lookups

- Around half of all the ULA seen used FF:FE structured address assignment model, via stateless autoconfiguration

We know the vendor of the MAC



How do the /128 assign?

- over half the ULA seen appear to be using ff:fe structured MAC addresses for the /128
- By comparison, use of non-privacy mode in global unicast has dropped off significantly
 - Either the processes behind ULA don't enable temporary/privacy mode
 - Or the time when ULA intrude into gethostbyname() the address selected isn't privacy mode
 - Or Privacy mode hasn't spun up yet when ULA is used

We know the vendor of the MAC



We can
Analyse this
By IEEE
Assigned
Vendor code

How do the /128 assign?

- over half the ULA seen appear to be using ff:fe structured MAC addresses for the /128
- By comparison, use of non-privacy mode in global unicast has dropped off significantly
 - Either the processes behind ULA don't enable temporary/privacy mode
 - Or the time when ULA intrude into gethostbyname() the address selected isn't privacy mode
 - Or Privacy mode hasn't spun up yet when ULA is used

Unique OUI assigned vendor codes

| Vendor | Count | % |
|----------------|---------|-------|
| APPLE | 144,927 | 98.15 |
| IBM | 950 | 0.64 |
| VMWARE | 479 | 0.32 |
| <i>various</i> | 1,290 | 0.87 |

Unique OUI assigned vendor codes

| Vendor | Count | % |
|----------------|---------|-------|
| APPLE | 144,927 | 98.15 |
| IBM | 950 | 0.64 |
| VMWARE | 479 | 0.32 |
| <i>various</i> | 1,290 | 0.87 |

This is overwhelmingly about APPLE devices using ULA.

So what does APPLE do which uses ULA?

Lets turn on back-to-my-mac



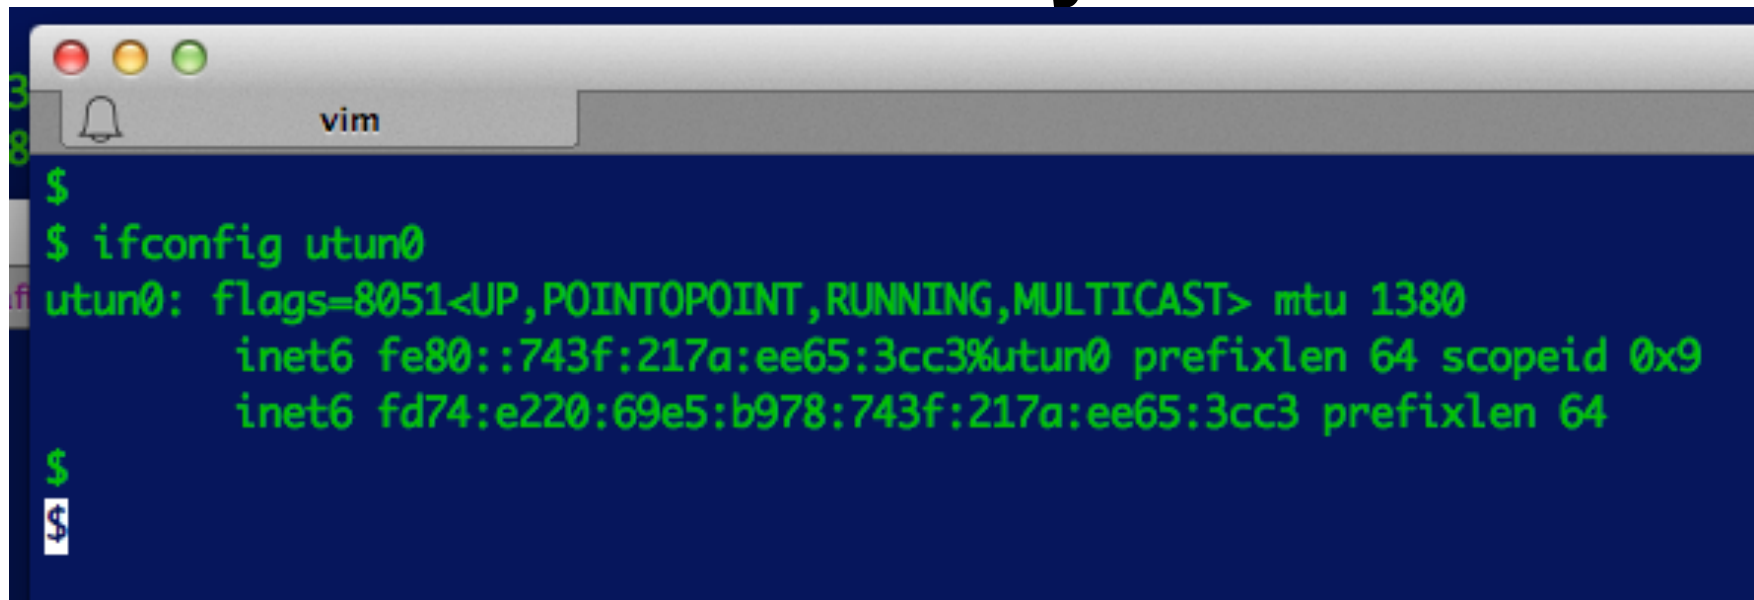
Lets turn on back-to-my-mac



Lets turn on back-to-my-mac

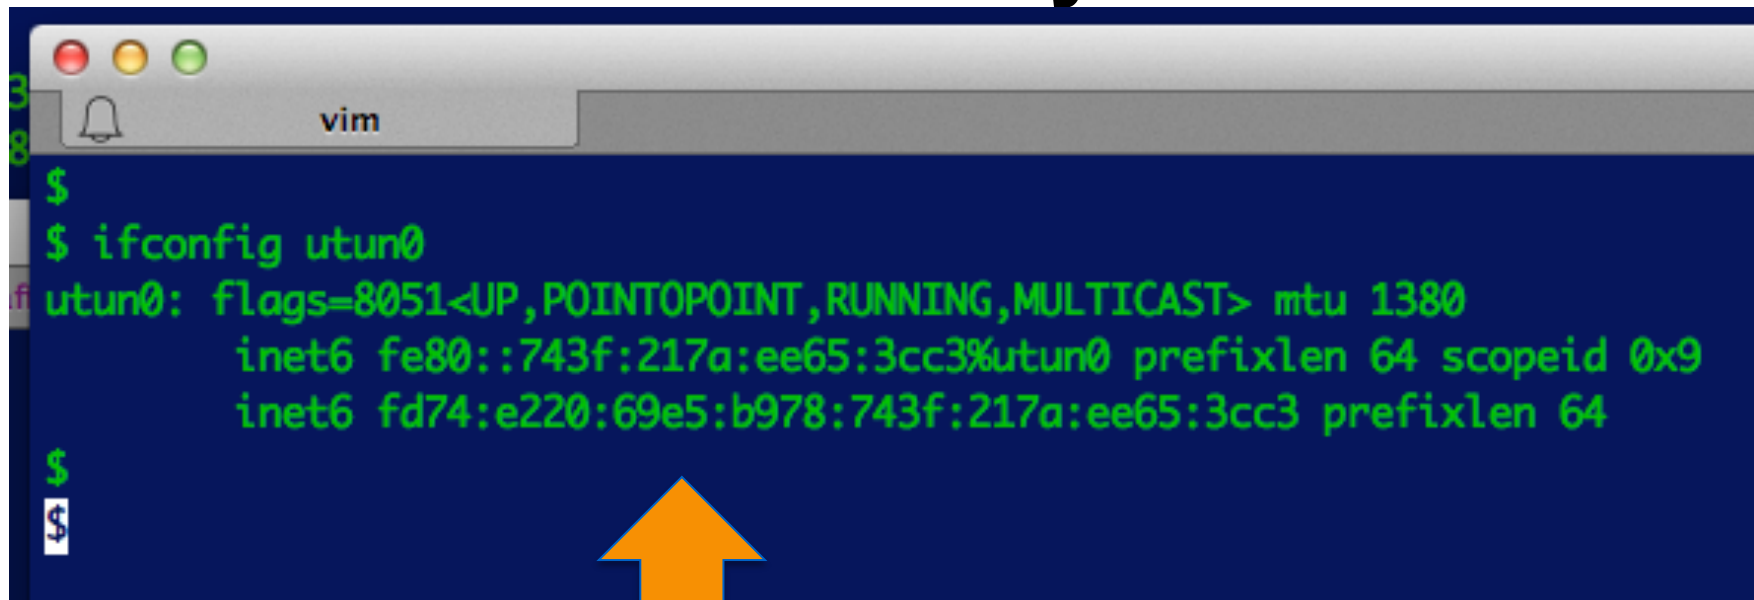


Lets turn on back-to-my-mac



```
$  
$ ifconfig utun0  
utun0: flags=8051<UP,POINTOPOINT,RUNNING,MULTICAST> mtu 1380  
    inet6 fe80::743f:217a:ee65:3cc3%utun0 prefixlen 64 scopeid 0x9  
    inet6 fd74:e220:69e5:b978:743f:217a:ee65:3cc3 prefixlen 64  
$  
$
```

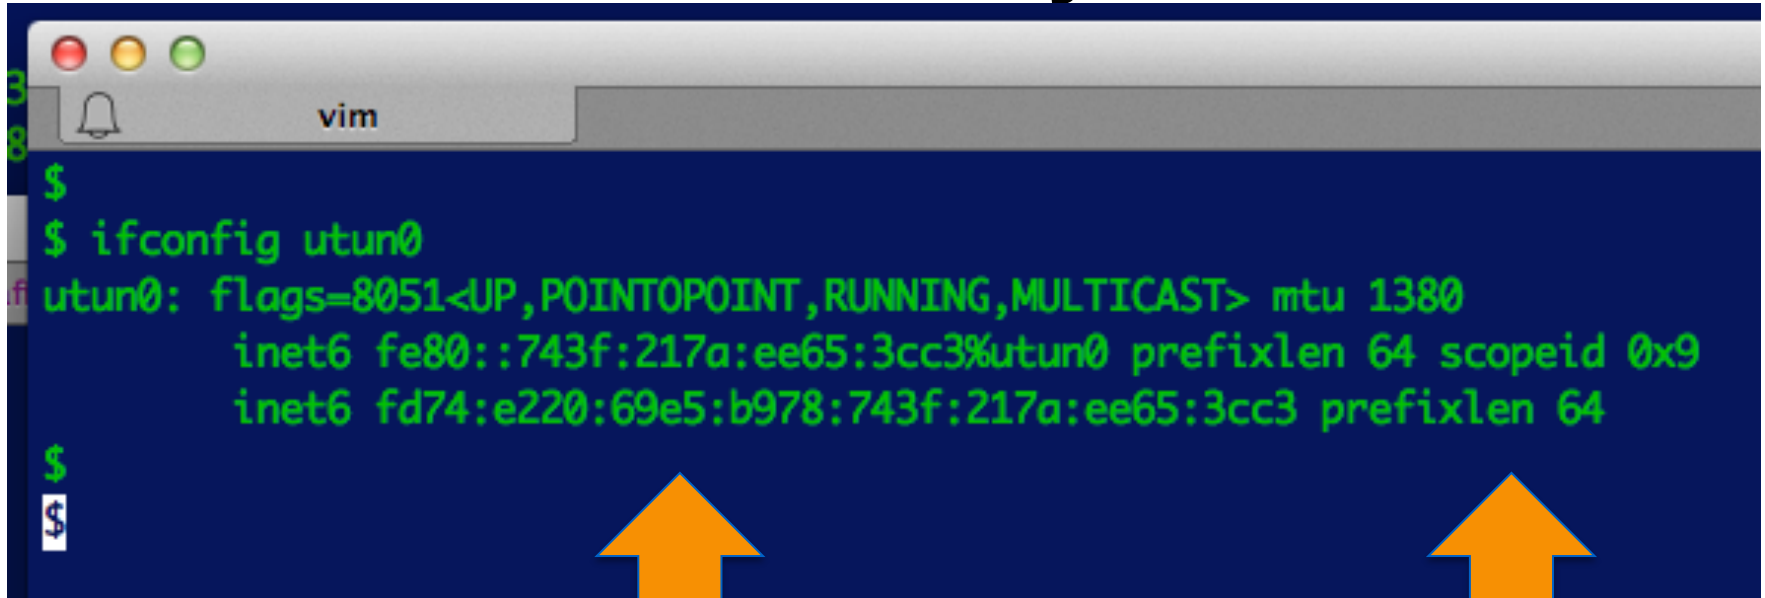
Lets turn on back-to-my-mac



```
$  
$ ifconfig utun0  
utun0: flags=8051<UP,POINTOPOINT,RUNNING,MULTICAST> mtu 1380  
    inet6 fe80::743f:217a:ee65:3cc3%utun0 prefixlen 64 scopeid 0x9  
    inet6 fd74:e220:69e5:b978:743f:217a:ee65:3cc3 prefixlen 64  
$  
$
```

That's a ULA.

Lets turn on back-to-my-mac



```
$  
$ ifconfig utun0  
utun0: flags=8051<UP,POINTOPOINT,RUNNING,MULTICAST> mtu 1380  
    inet6 fe80::743f:217a:ee65:3cc3%utun0 prefixlen 64 scopeid 0x9  
    inet6 fd74:e220:69e5:b978:743f:217a:ee65:3cc3 prefixlen 64  
$  
$
```

That's a ULA.

Its Unscoped

Back-to-my-Mac uses ULA!

- Back to my Mac is an IPv6 based service
- It creates ad-hoc ESP tunnels from your Mac
 - back into Apple.
- The IPv6 endpoint address is drawn from the ULA self-assigned upper-pool
- The IPv6 endpoint address is not scoped
 - There are explicit routes bound to the tunnel so traffic does not leak
- What about INADDR_ANY address selection?
 - Its possible other services running on the Mac acquire knowledge of the address, and use it 'by mistake' triggering reverse-DNS checks for logging.

And the rest?

- 50% of the ULA seen are Apple/Back-to-My-Mac
- Parallels Desktop also appears to use ULA in their virtual interfaces
- The rest are unknown purpose, unknown architecture (using privacy addresses)
- With around 3000 entries in sixxs.net registry, there clearly is a community of people exploring ULA, for local IPv6 or as part of homenet or other experimental deployments

Thanks to Jen Linkova, Google Network Engineer, Sydney for information about Mac, Parallels

ULA here to Stay

- Originally we thought that there was no need for RFC1918 equivalents in IPv6:
 - we would all use provider-based addressing
 - multi-addressing would work
 - and renumbering would be easy
- But as things have turned out folk **do** want a consistent, stable, **internal** address structure, independent of external provider prefixes.
- ULAs have a role in worldwide IPv6 deployment.