

Do we need a registry for IP geolocation information?



Matthew Lepinski
Richard Barnes
BBN Technologies

Background

- ❑ Content providers increasingly wish to tailor their content to the geographic location of the viewer
 - E.g., language, relevance, rights management
- ❑ To facilitate this goal, content providers use IP to geolocation mapping data that comes from
 - Proprietary commercial databases (e.g. MaxMind or IP2Location)
 - Mining of Whois data
 - Sophisticated Heuristic guessing
- ❑ This works quite well most of the time, but ...

The Problem



MaxMind thinks
we're here

So we get this
version of the page



The Problem

When we're
actually here

And should get this
version of the page



To: nanog@nanog.org

Subject: New netblock Geolocate wrong (Google)

I just lit up a new IP netblock (assigned directly from ARIN) and the companies that provide Geolocate databases do not have the correct location information available yet.

Specifically Maxmind thinks we are in Canada and IP2LOCATION has no data.

For the most part this is benign or at worst slightly impacting since I often get redirected to global load balance nodes up in Canada instead of locally in the North West, however, the more major issue I am running into is that **Google chooses to redirect all my users to <http://www.google.ca>**

So my questions to others are:

1. **How do I get my data updated in all of the geolocation providers databases as quickly as possible?**
2. What geolocate database does Google use (is it homegrown?) and how do I get them to update my data?

To: nanog@nanog.org

Subject: Google/Yahoo - Geo-Location Issues

Hi all.

Grateful if someone from Google and Yahoo can contact me
off-list re: some geo-location issues with their web sites,
our side of the world.

E-mail to the 'noc@' addresses seem to have > /dev/null'ed.

To: nanog@nanog.org

Subject: Geolocation contact for Bing/Microsoft?

Can someone from Bing/MS contact me about correcting Geolocation info
for some IP's. Folks are erroneously getting redirected - and I can't
find any info about how to get it fixed.

In Summary

- ❑ Things work pretty well most of the time

- ❑ But when things don't work ...
 - ISP customers are getting the wrong content
 - ISP employees are scrambling to try and find the right contact method for each content provider

- ❑ Perhaps there is a better way

A Case for Optimism

- ❑ Content providers want to deliver geographically appropriate content
- ❑ Geolocation database providers want their databases to be accurate
- ❑ End-users (almost always) want to get content that is appropriate for their geography
- ❑ ISPs want their customers to get geographically appropriate content
- ❑ ... So maybe we just need a standard way for ISPs to tell people where their networks are located

Why not make a registry?

A thick, blue, brushstroke-style underline that tapers at both ends, positioned directly beneath the text.

Let's make a registry

- ❑ We already tell people a lot about our IP address allocations
 - What organizations they're registered to
 - What ASNs will be originating them
 - Who to contact if there's trouble
- ❑ Geolocation information is just another data element
- ❑ Provide real data instead of guesses
 - ISPs can control how much information is revealed
- ❑ Complement other techniques
 - More general than GPS, more reliable than latency-based

Possibility #1: Extend WHOIS

inetnum: 169.223.0.0 - 169.223.255.255

netname: APNIC30

located: SURFERS-PARADISE-MARRIOTT

geoloc: SURFERS-PARADISE-MARRIOTT

address: 158 Ferny Avenue

address: Surfers Paradise, Queensland 4217

country: AU

Possibility #1: Extend WHOIS

☐ Positives:

- Re-uses existing databases, tools, provisioning systems
- Easy to tie into existing structures for describing IP address blocks

☐ Negatives:

- Have to update existing databases, tools, provisioning systems
- Unstructured location data format – ambiguous parsing

Possibility #2: Web Service

```
<locationRequest>
  <device><prefix>169.223.0.0/16</prefix></device>
</locationRequest>

<locationResponse>
  <presence>
    <tuple><status><geopriv><location-info><civicAddress>
      <country>AU</country>
      <A1>Queensland</A1>
      <A3>Surfer's Paradise</A3>
    </civicAddress></location-info></geopriv></status></tuple>
  </presence>
  <locationUriSet>
    <locationURI>http://example.com/apnic30loc</locationURI>
  </locationUriSet>
</locationResponse>
```

Possibility #2: Web Service

☐ Positives:

- More structured format for location info, especially for geospatial information (coordinates)
- Better support for Internationalization
- Can still bootstrap from WHOIS (add a URI)
- Built-in location URI support

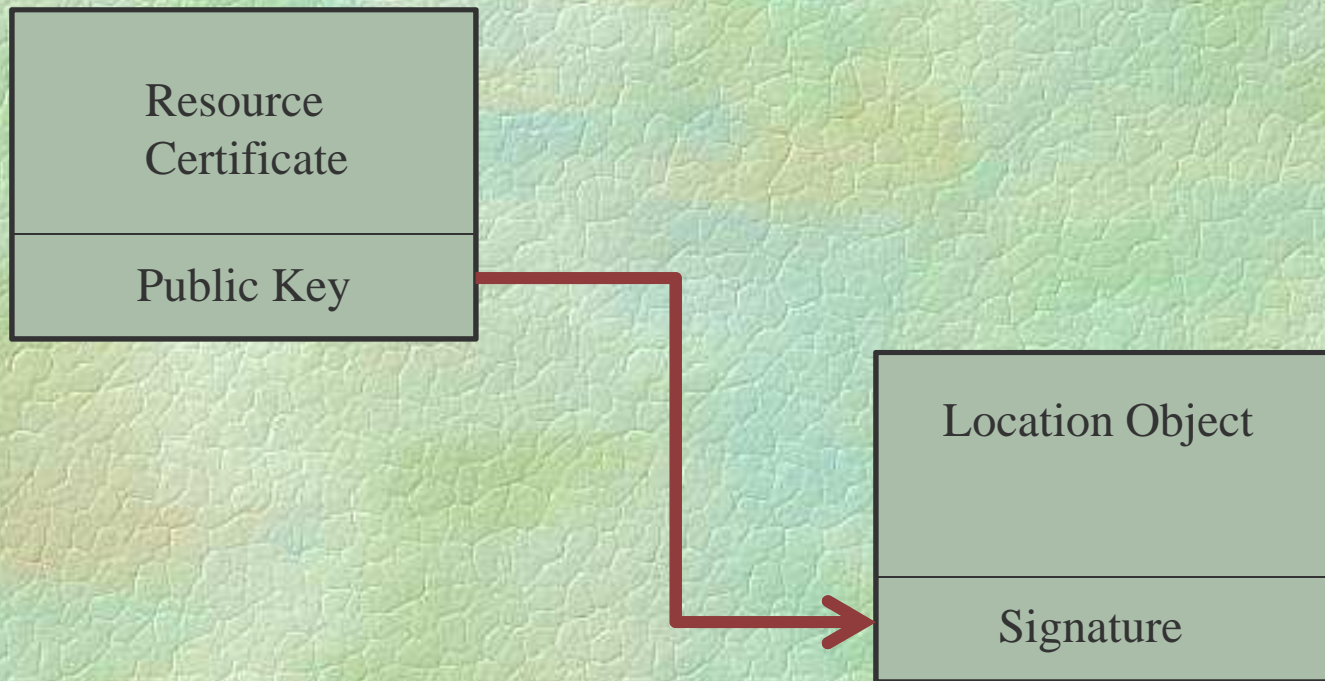
☐ Negatives:

- Much more verbose
- New database, tools, provisioning systems

Location URIs (minor digression)

- ❑ A registry web service could return either a location object or a location URI
(resource holder could decide which to use)
- ❑ A location URI is a URI that can be dereferenced to obtain a location object
 - E.g. <http://location.example.com/loc/ABCXYZ>
- ❑ Why?
 - An ISP who wishes to do so can serve location from their own server
 - Can give different answers to different requestors
E.g., more specific location for advertising partners

Possibility #3: RPKI Signed Object



- ❑ Resource certificate binds a prefix to a public key
- ❑ Key is used to validate a signed location object

Possibility #3: RPKI Signed Object

☐ Positives:

- RIRs are already building tools to support resource certificate issuance and signed object repositories
- Consumer of location data knows the data was provided by the legitimate resource holder
- Could adopt the same structured location format used in the web service example

☐ Negatives:

- Some complexity in creating and managing the signed objects

A Registry is Not a Panacea

- ❑ Registry would not replace existing location products
 - Although a registry could improve such products by giving them with a centralized source of operator-provided data
- ❑ Operator-provided data has no guarantee of accuracy
 - Although most of the data would likely be correct
 - Operators likely have good location data for their networks
 - Operators have an incentive to provide correct information
 - Even if not perfectly accurate, such data is a valuable input into the determination of IP-geolocation mappings
 - In cases where regulation calls for accurate data, additional validation would certainly be required
(e.g. tax jurisdictions)

Questions

- Is there a problem here to solve?
- Are any of the proposed solutions worth doing?
- Would you contribute data for your network?

Thank You

