Identity Trail: Covert Surveillance Using DNS

`http://saikat.dyndns.org/talk.pdf`

Saikat Guha and Paul Francis

Cornell University

7th workshop on Privacy Enhancing Technologies

Identity Trail

- Track someone without them knowing
  - Using public services (DNS, DynDNS, GeoIP)
  - Used like they were meant to be used
- Exploits
  - Lack of access control in DNS
  - Information derived from IP addresses over time
- Demonstrated for over 100K hosts
- Need for a new Internet naming architecture for non-public hosts

DNS and Dynamic DNS

- DNS — Name to IP address mapping
  - All data public, privacy not considered
  - Envisioned for IP renumbering of fixed hosts
  - Occasional updates, by network admin.
- Dynamic DNS — More frequent updates
  - Envisioned for fixed hosts with DHCP addresses
  - Host updates third-party DNS server
  - Still public, privacy still not considered
  - (Ab)used by mobile hosts
    - No practical alternative for individuals

DNS: No access control

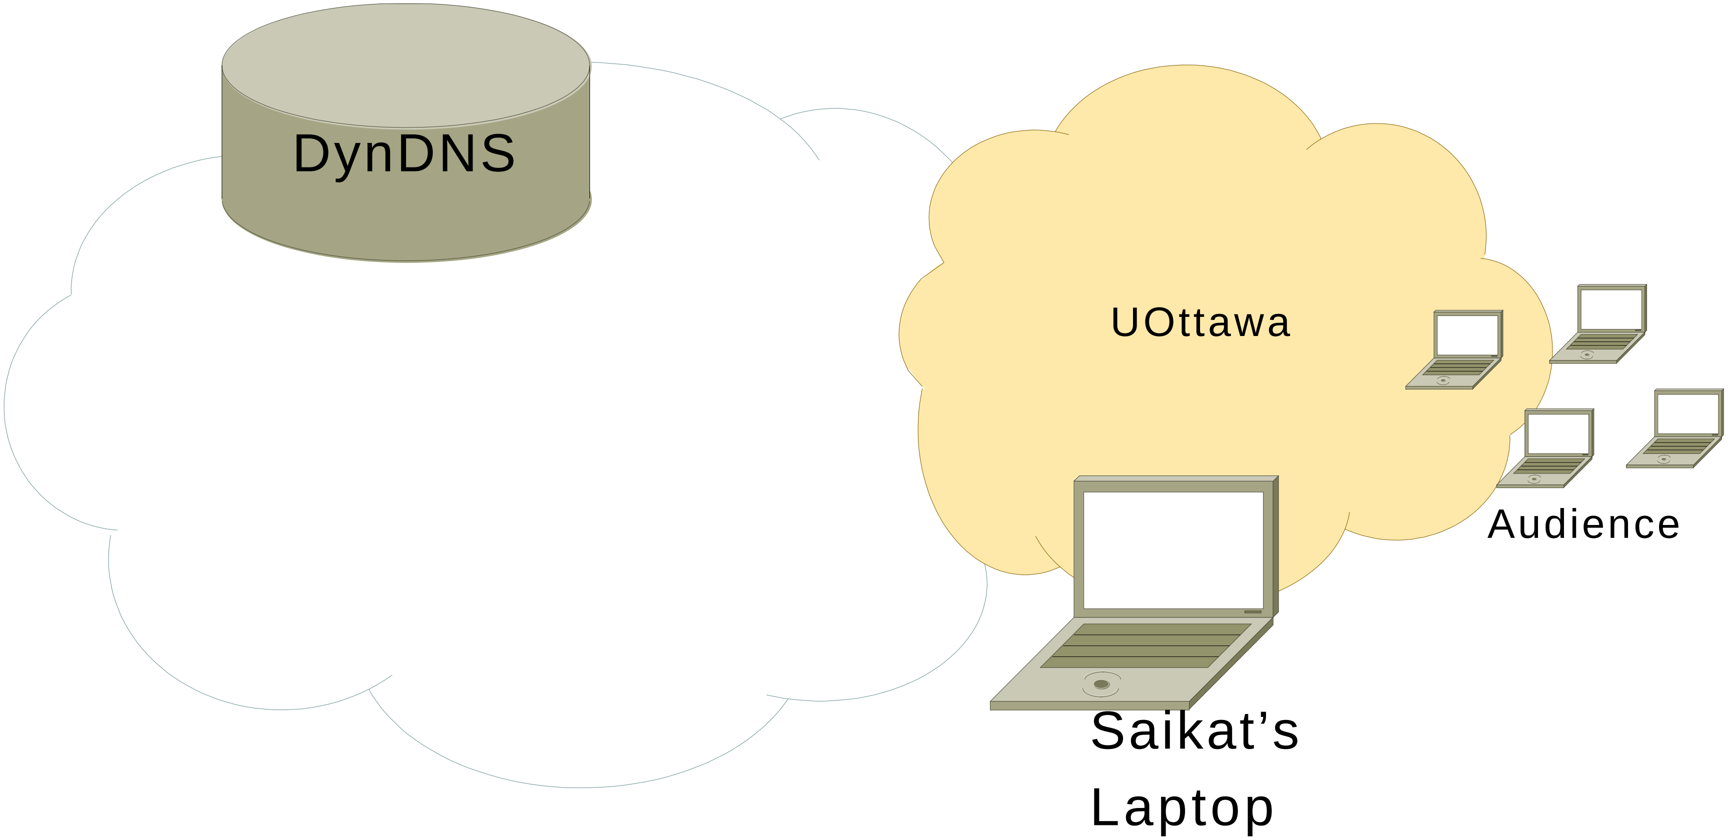

DNS: No access control

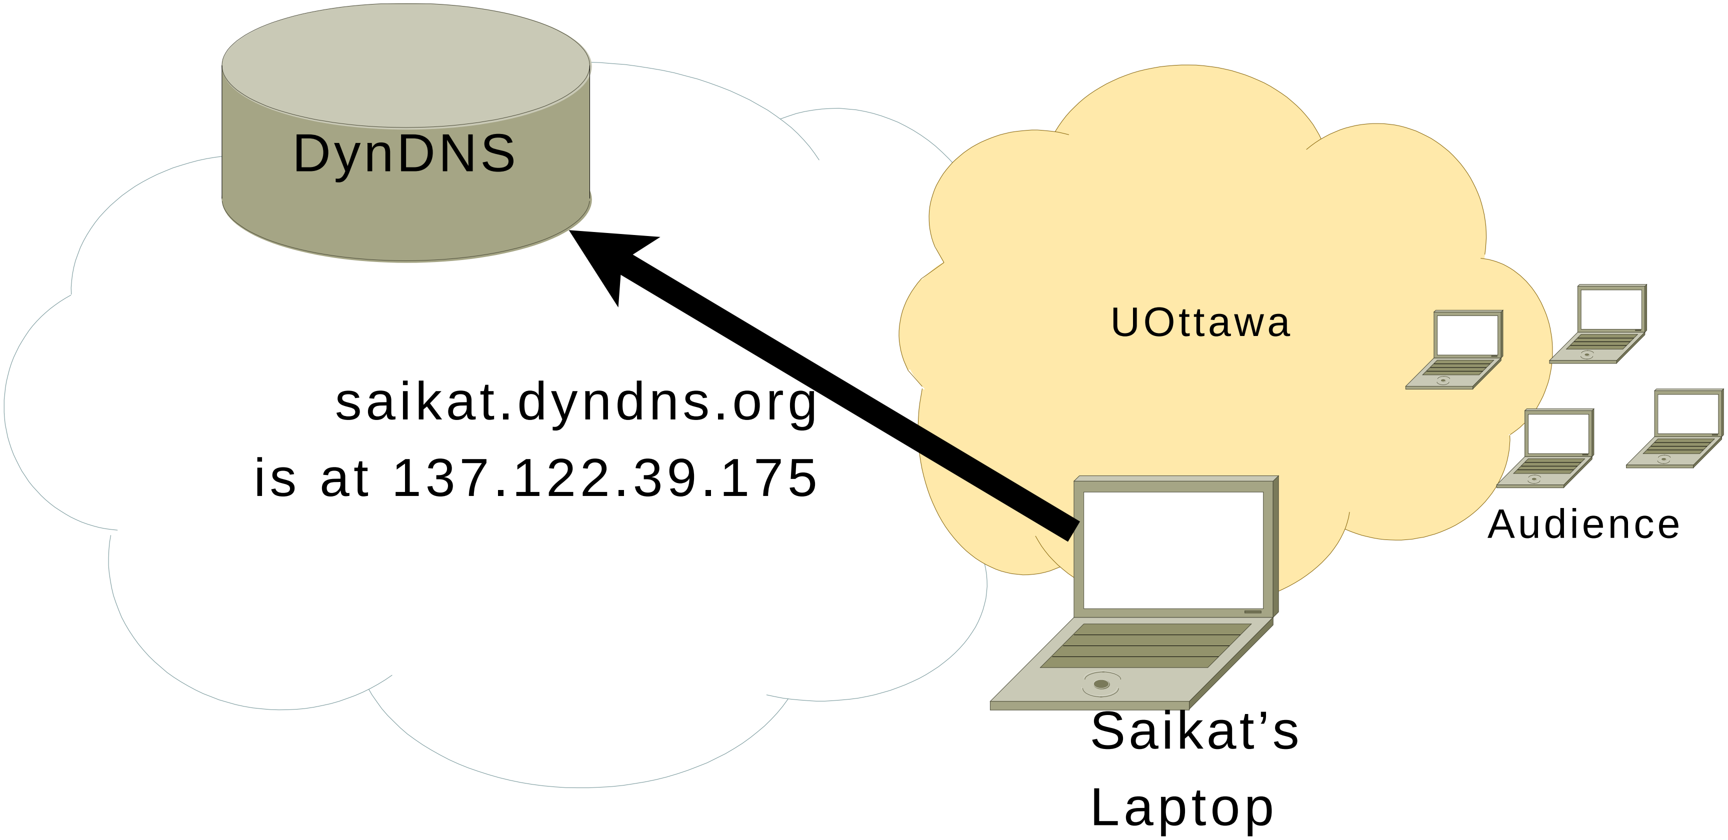

DNS: No access control

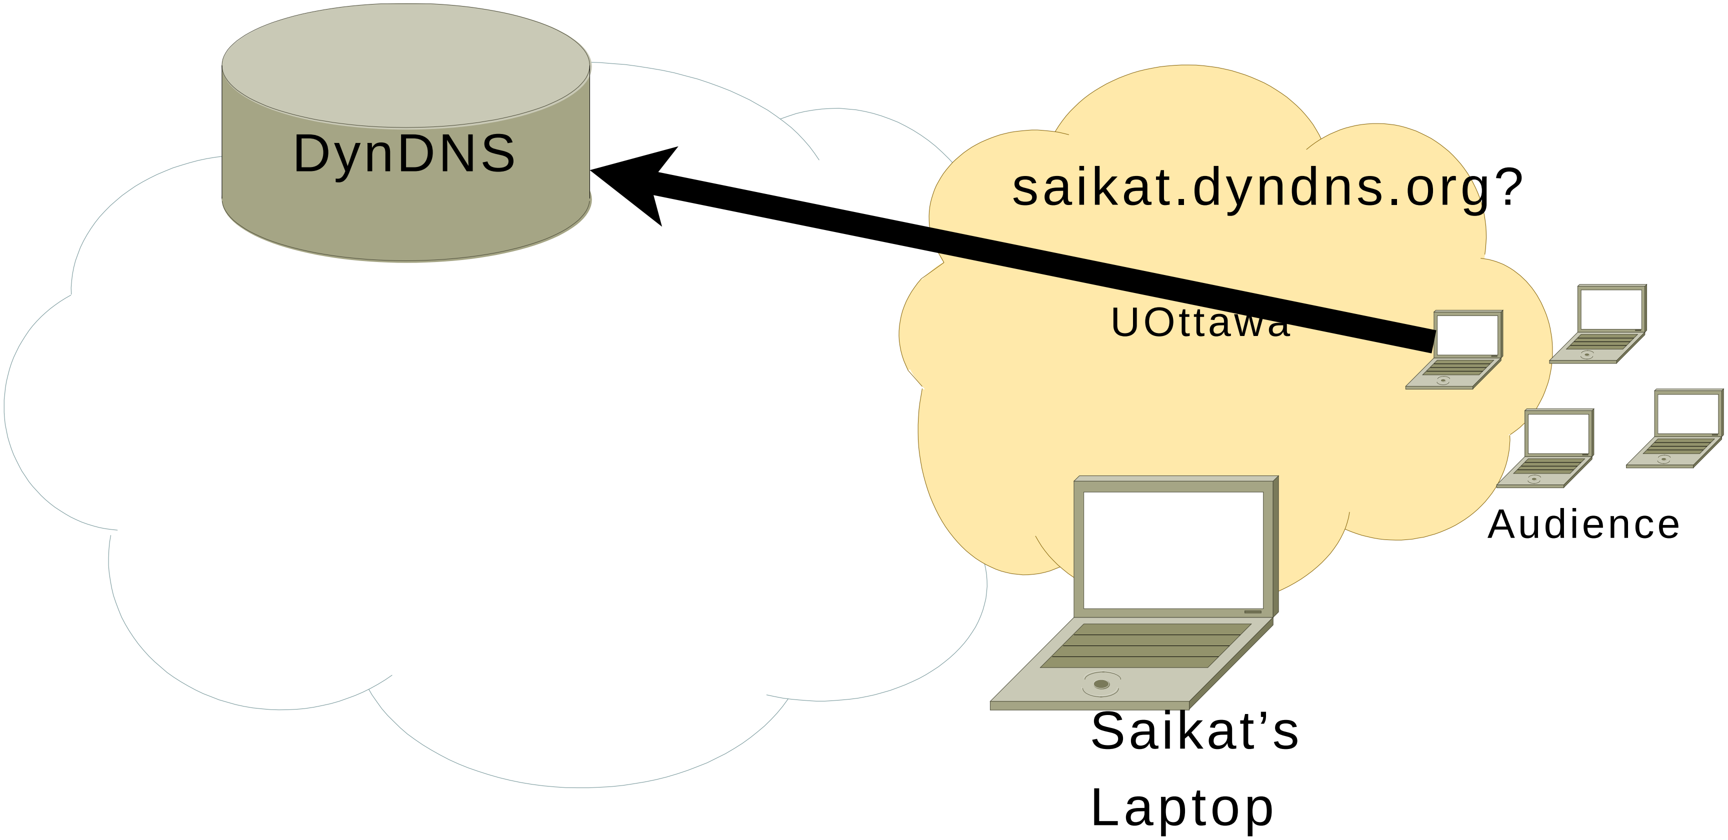

DNS: No access control

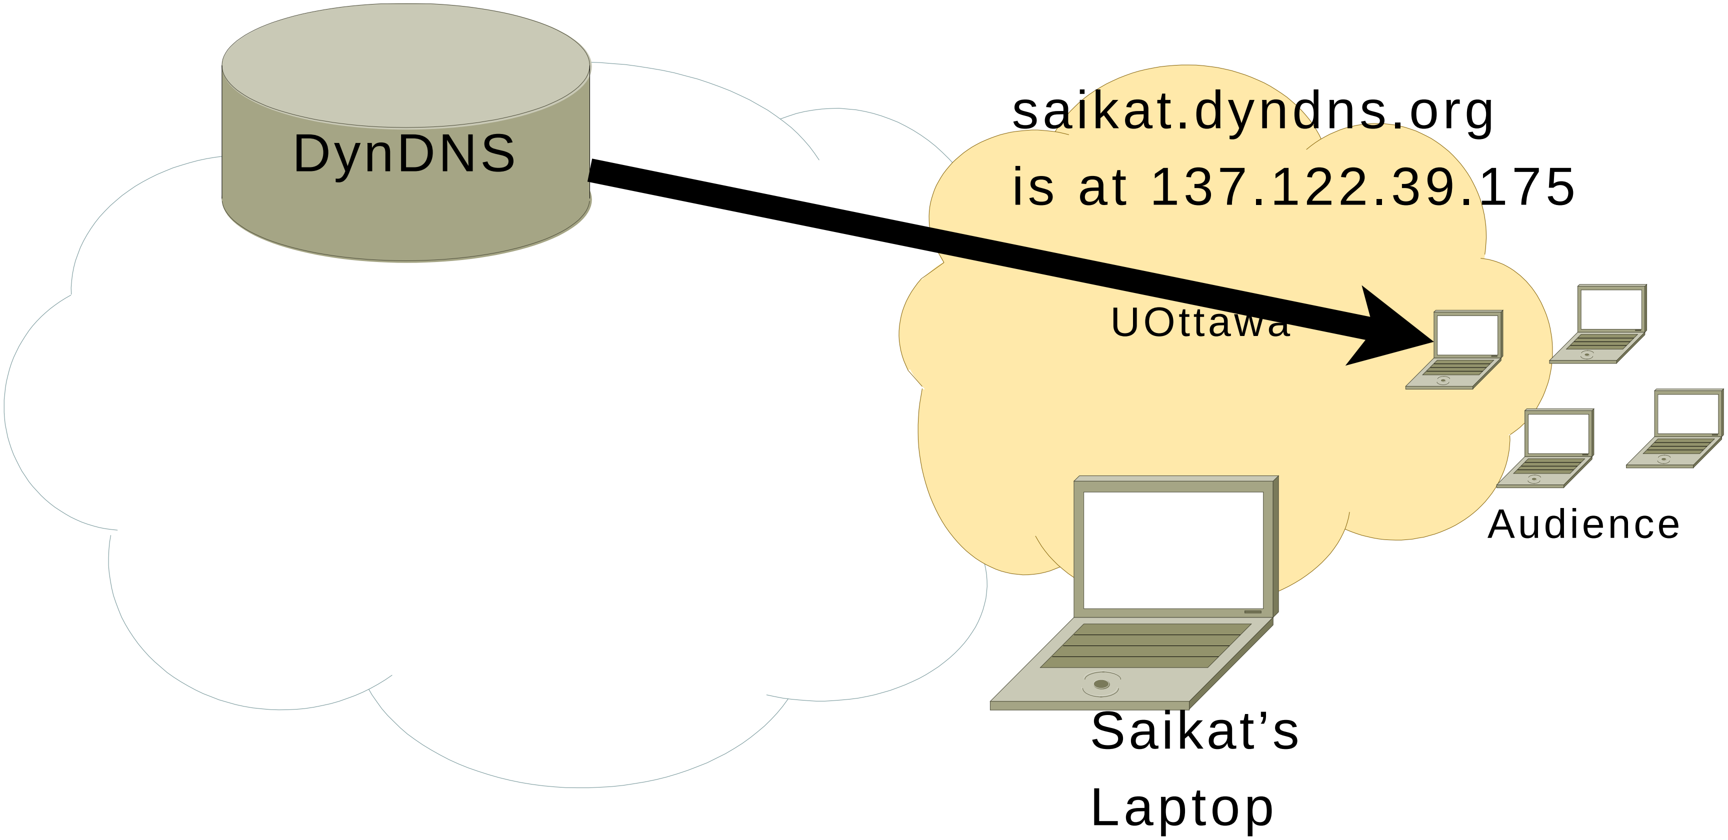

DNS: No access control

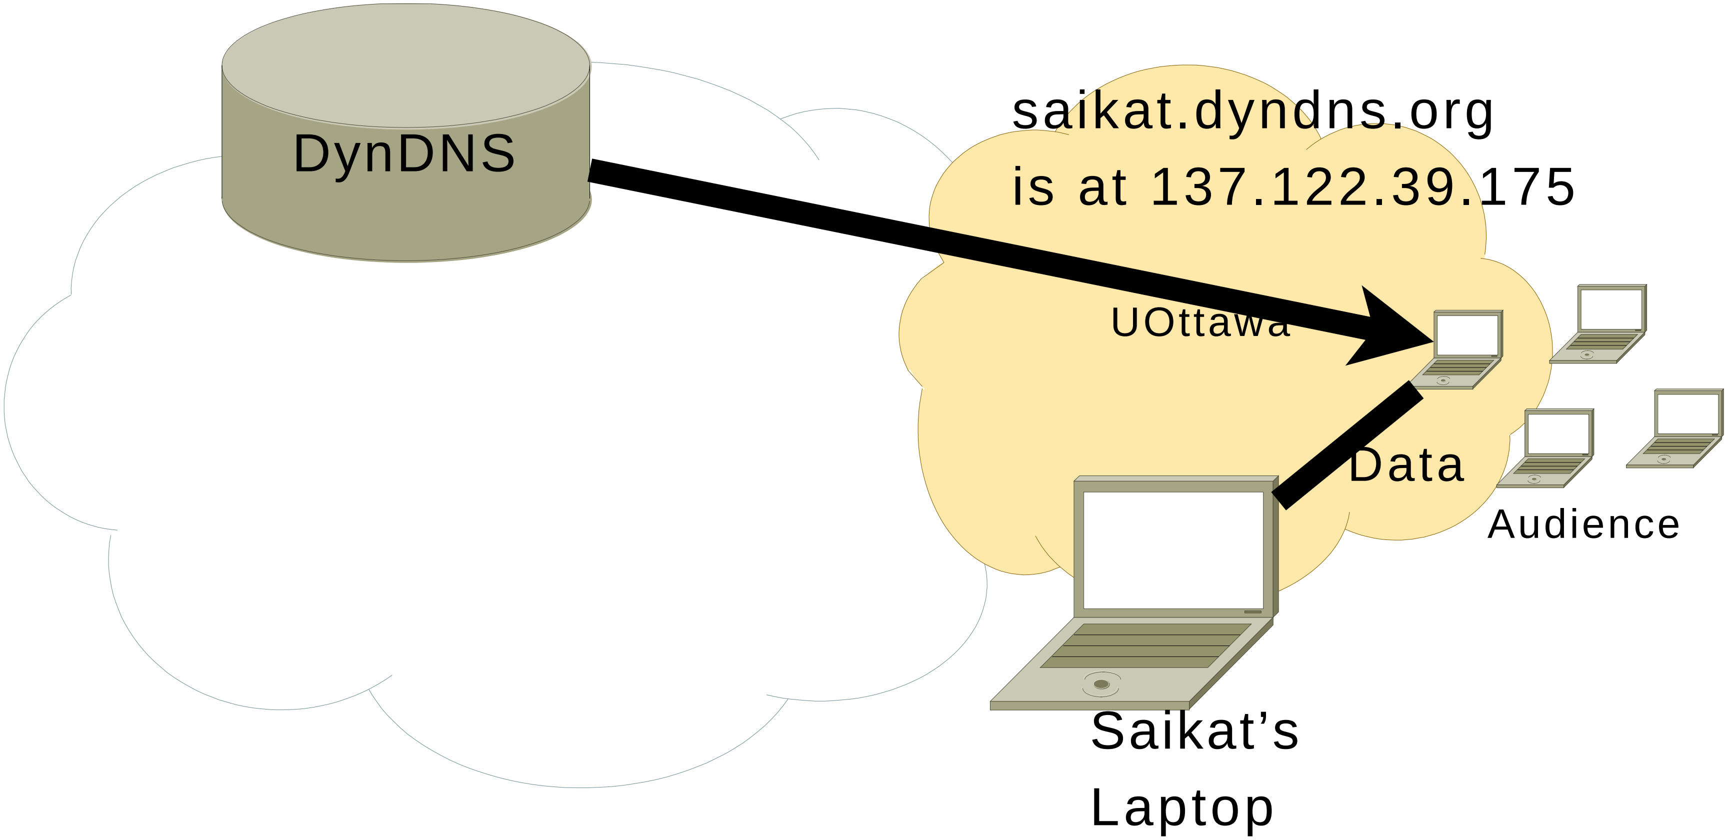

DNS: No access control

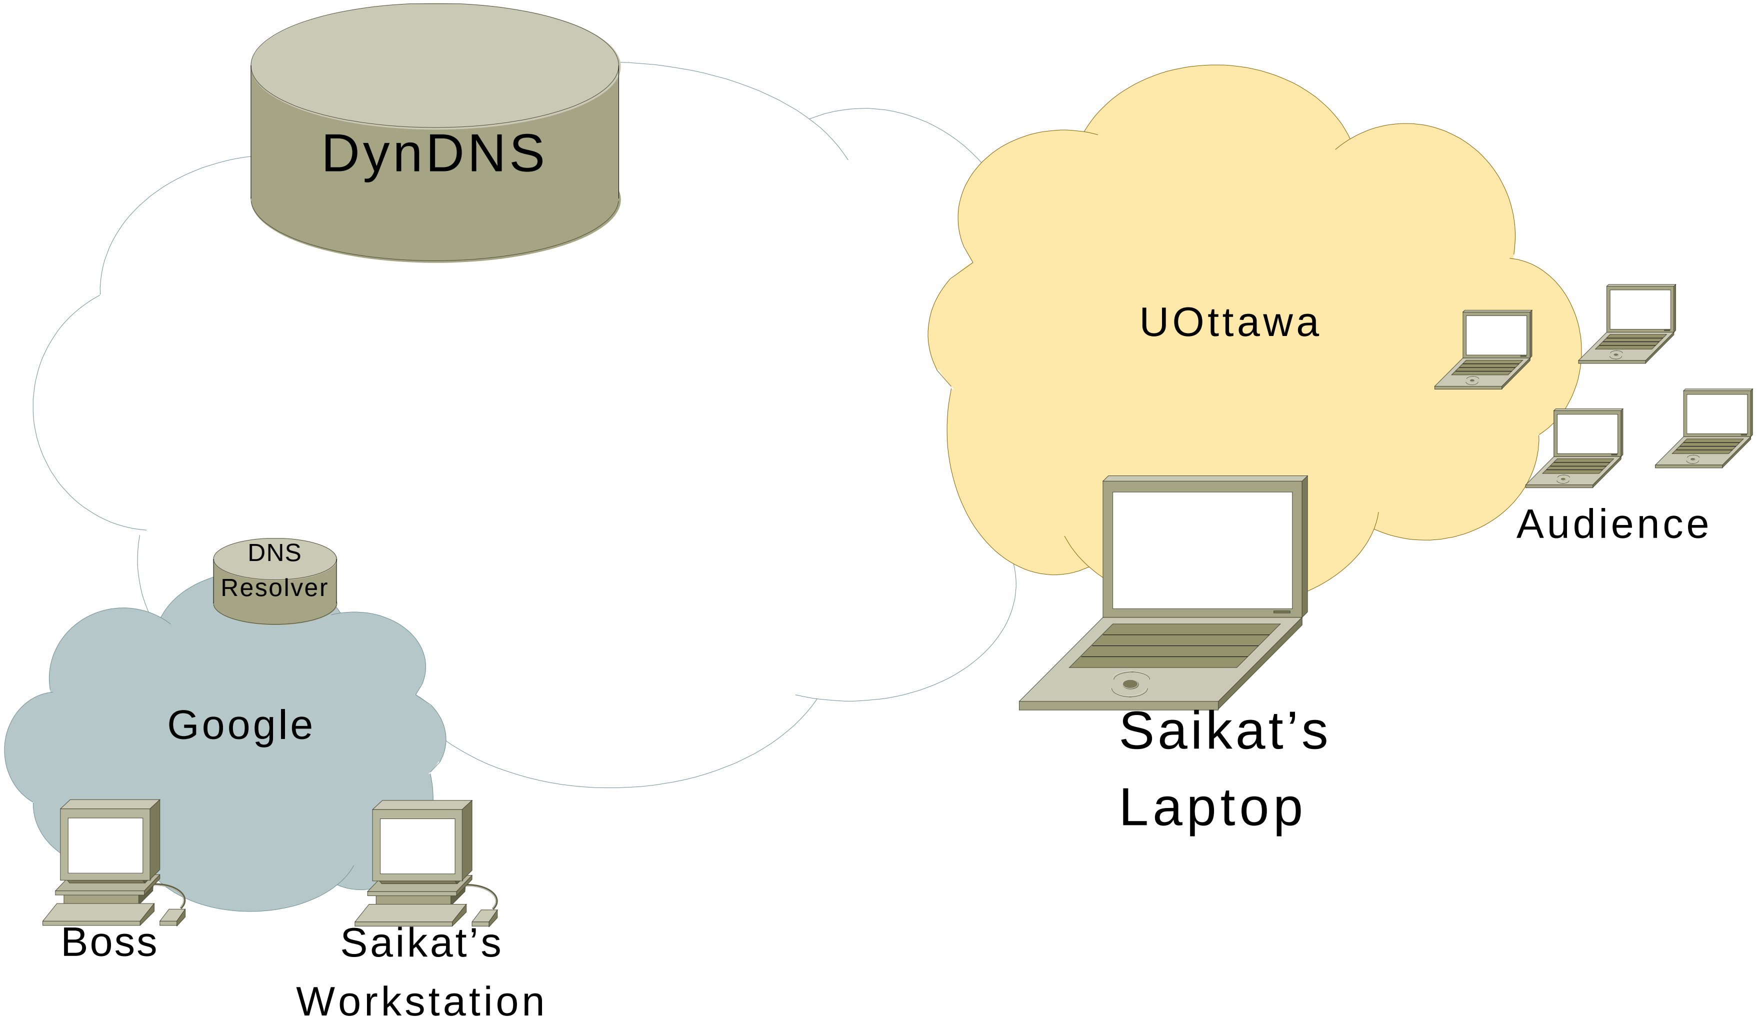

DNS: No access control

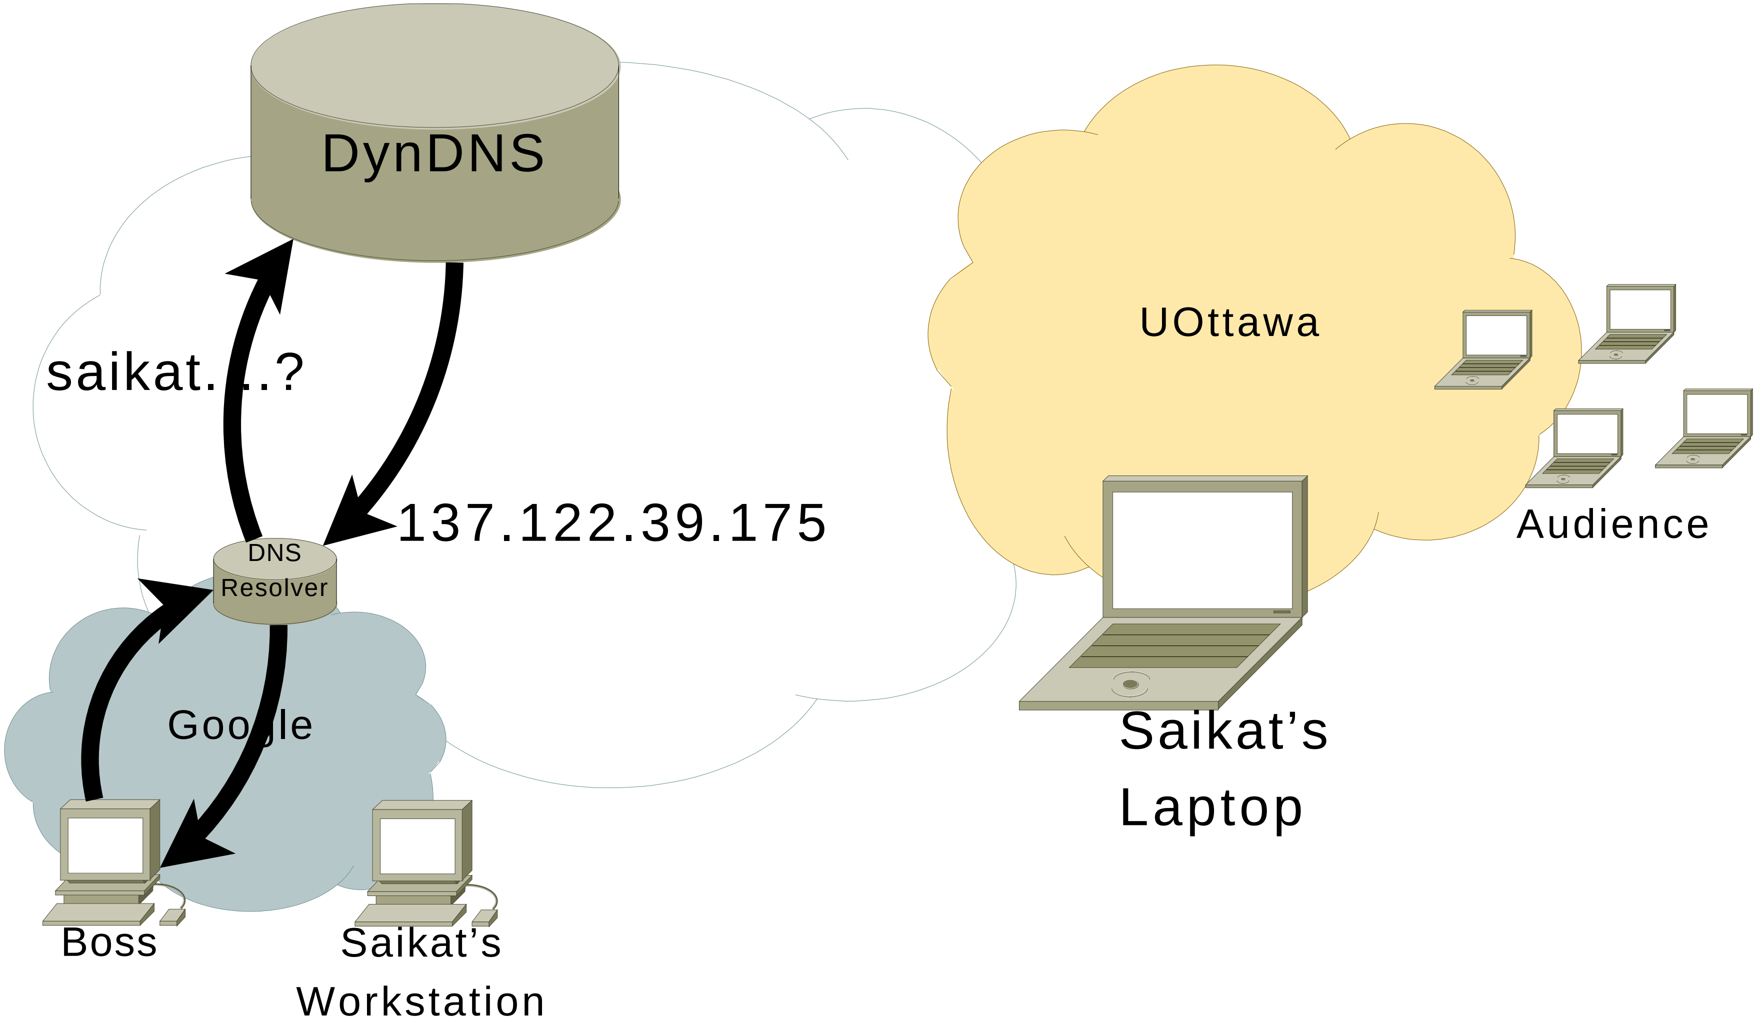

Identity Trail "Attack"

1. Find DNS hostname for victim
2. Perform DNS queries
3. Victim does not learn of query!
    - Even DynDNS doesn't know true source (recursive resolvers)
4. Geo-locate IP address
5. Create dossier over months
    - 9 lines of code. 5 minutes to write.

That simple? Yes.

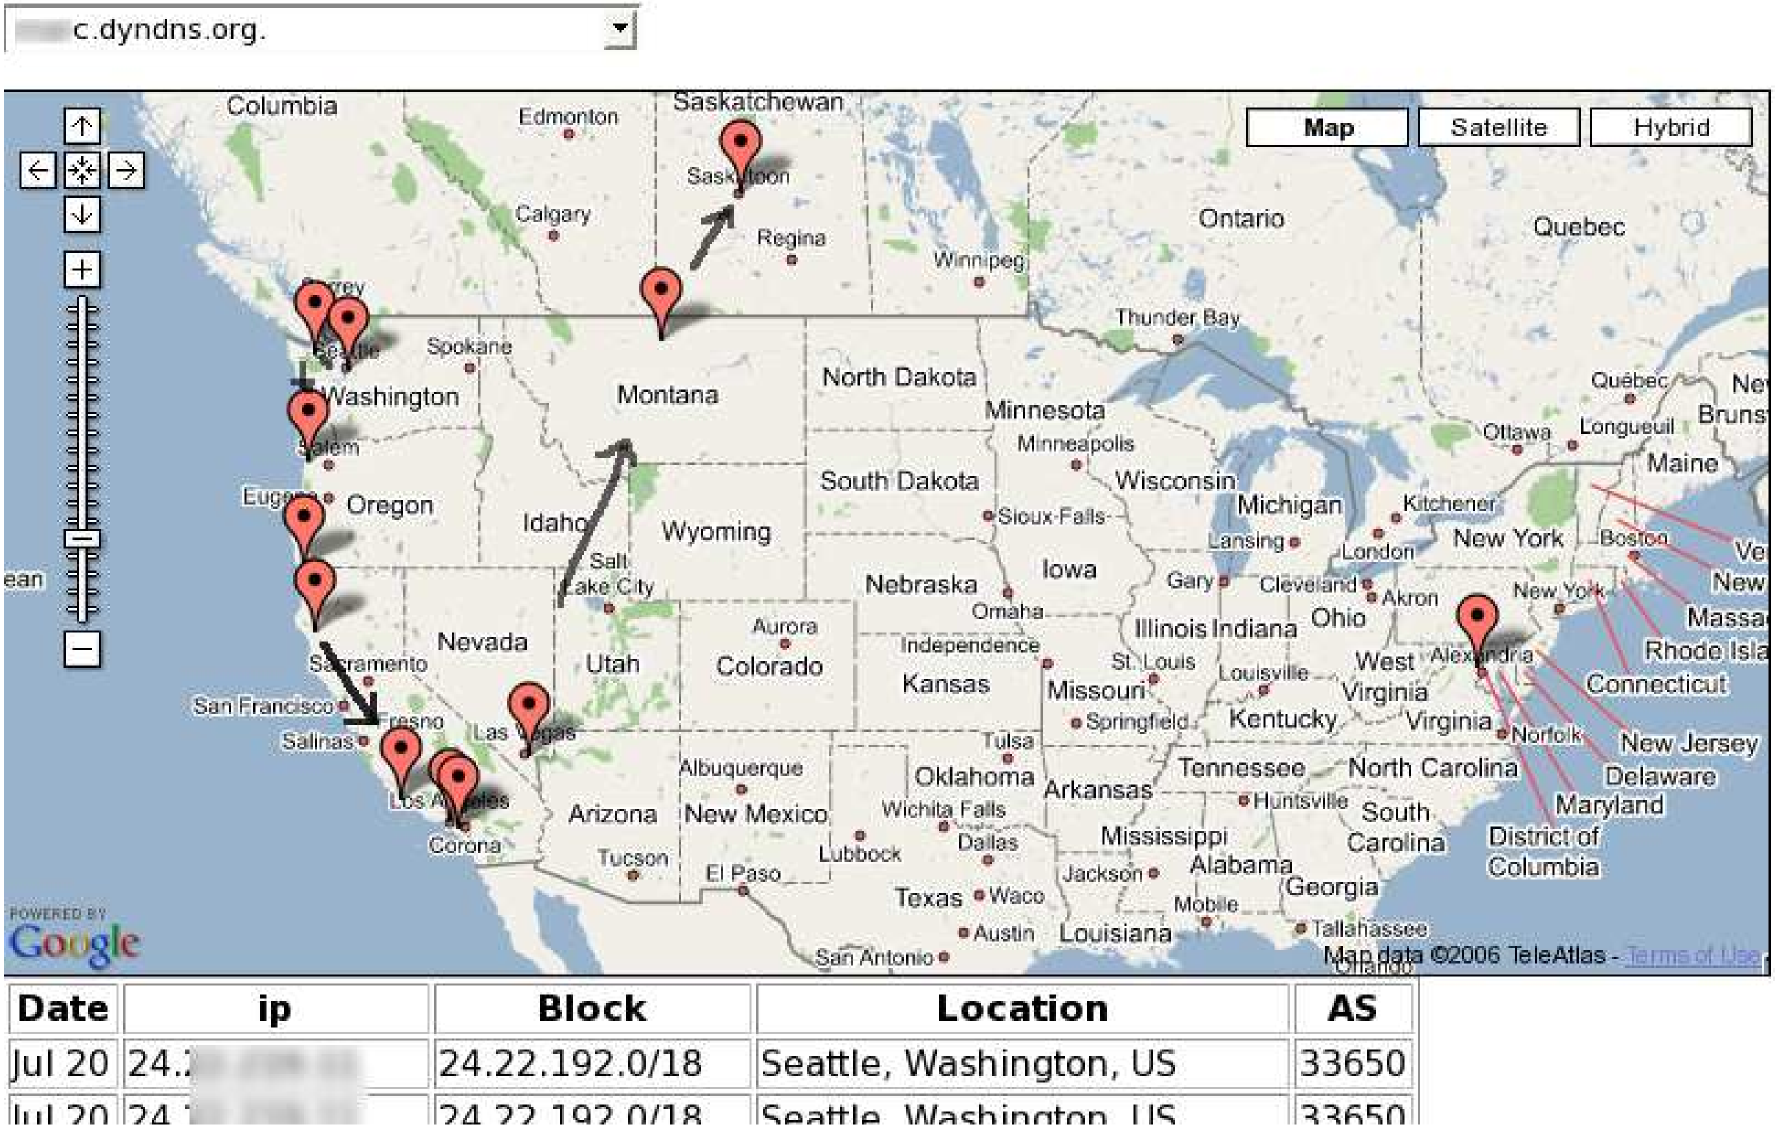

[redacted]c.dyndns.org.

Date	ip	Block	Location	AS
Jul 20	24.[redacted]	24.22.192.0/18	Seattle, Washington, US	33650
Jul 20	24.[redacted]	24.22.192.0/18	Seattle, Washington, US	33650

That simple? Yes.

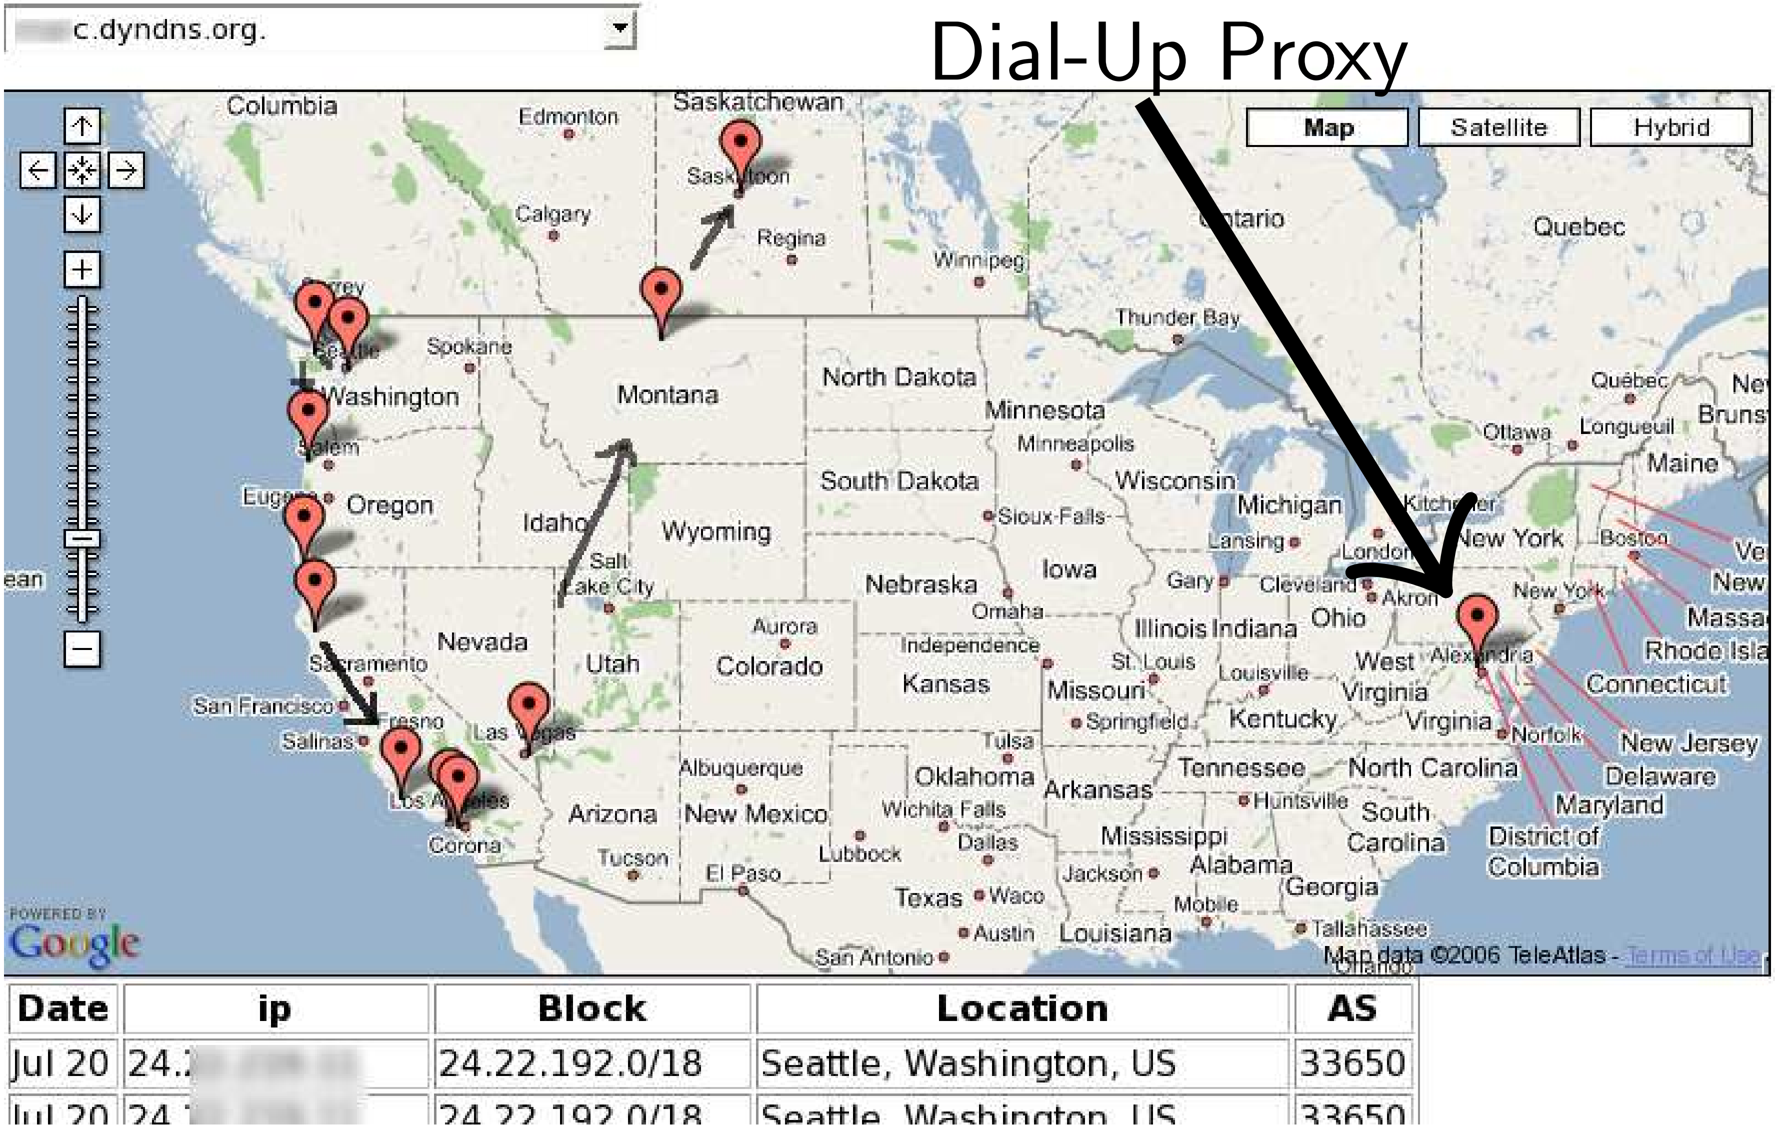

[redacted]c.dyndns.org.

Dial-Up Proxy

Date	ip	Block	Location	AS
Jul 20	24.[redacted]	24.22.192.0/18	Seattle, Washington, US	33650
Jul 20	24.[redacted]	24.22.192.0/18	Seattle, Washington, US	33650

But why use dynamic DNS for laptop?

- .Mac users:
  "...We've got the Internet, we've got .Mac, we've got my Mac at my house... I'm on the road and I need a file... when my home Mac gets a new IP address, it always tells .Mac. My mobile tells IT'S IP address to .Mac, so my notebook knows where my desktop is ..."
  — Steve Jobs, WWDC'07
- DynDNS users: tens (perhaps hundreds) of thousands mobile users

Validation: Finding Victims

- Decided to target DynDNS users
  - Real attacker model: attacker knows victim (spouse, employee)
- Google, Yahoo searches: surprisingly few (∼4K)
- Dictionary attack: many many more (∼31K)
- Nmap scan of a small number of victims
  - Services required authentication
  - Blank default web pages etc.

DynDNS hostnames rarely advertised publicly. Most likely intended for private use.

Validation: Mobility

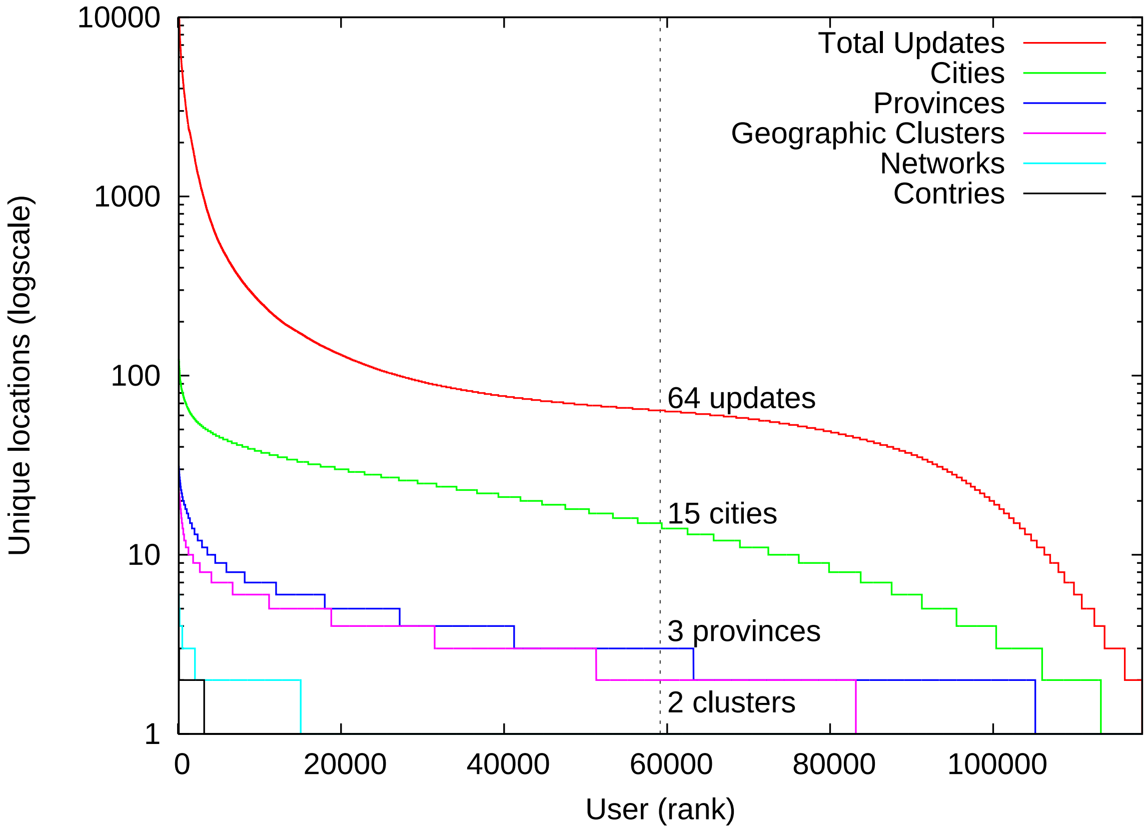

∼70% of the 125,000 DynDNS users trailed logged in from different locations. Disclaimer: data was noisy, see paper.

There exist many mobile users that want user-friendly name resolution for private services.

Validation: Accuracy

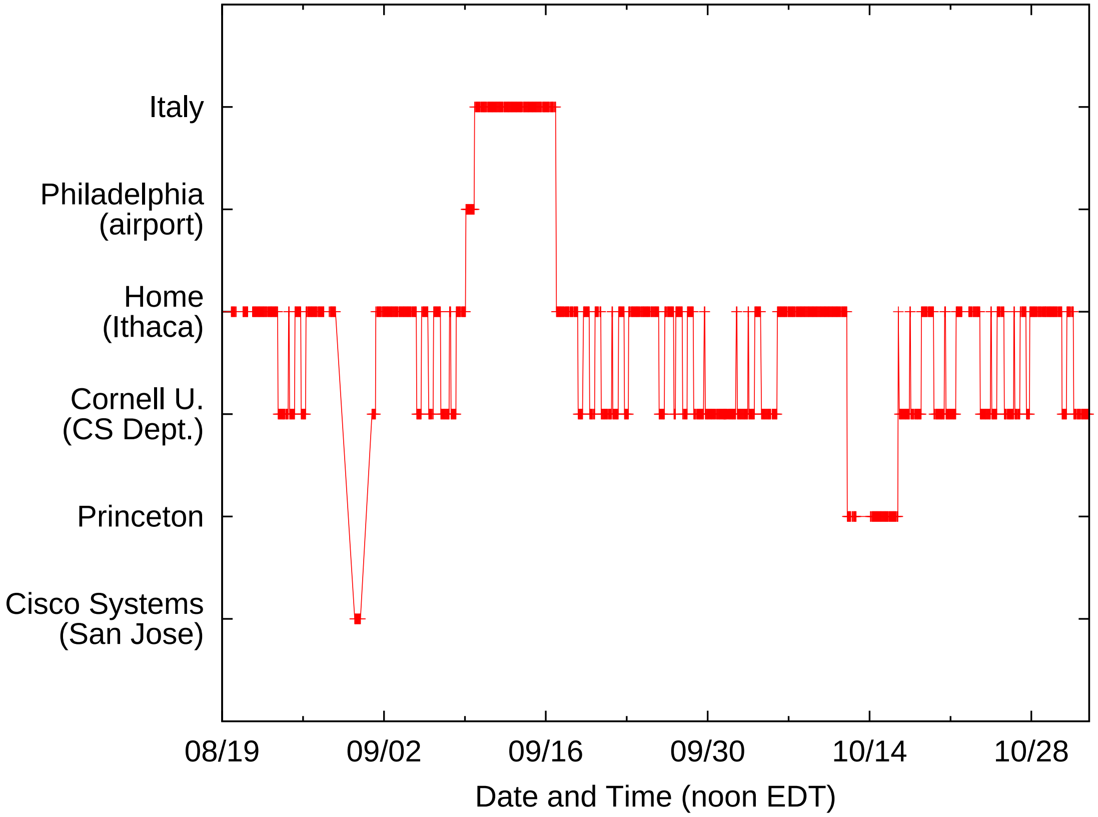

- Trailed Paul
- City-level accuracy in US (∼100 mi), province-level in Italy for GeoIP service used.
- Commute time accurate to within query interval. Some exceptions.

Reasonably good accuracy. Reconstructed travel itineraries, daily commute patterns.

Validation: Accuracy

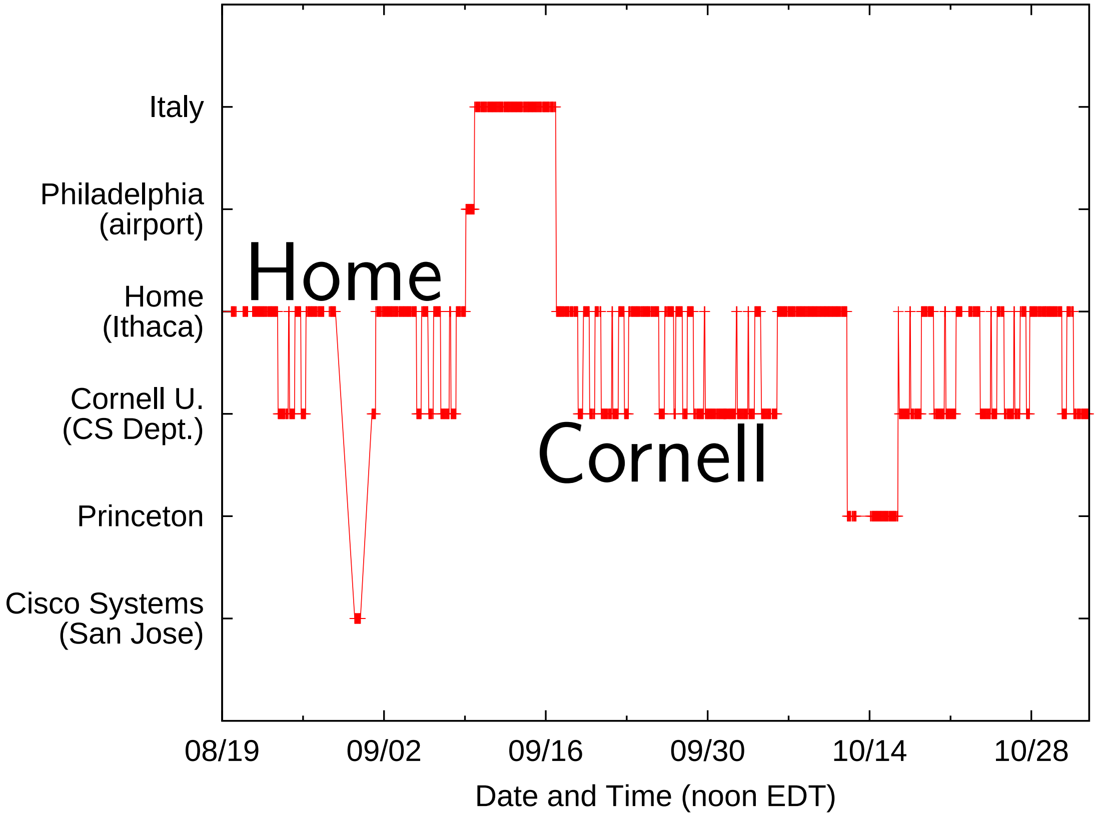

- Trailed Paul
- City-level accuracy in US (∼100 mi), province-level in Italy for GeoIP service used.
- Commute time accurate to within query interval. Some exceptions.

Reasonably good accuracy. Reconstructed travel itineraries, daily commute patterns.

Validation: Accuracy

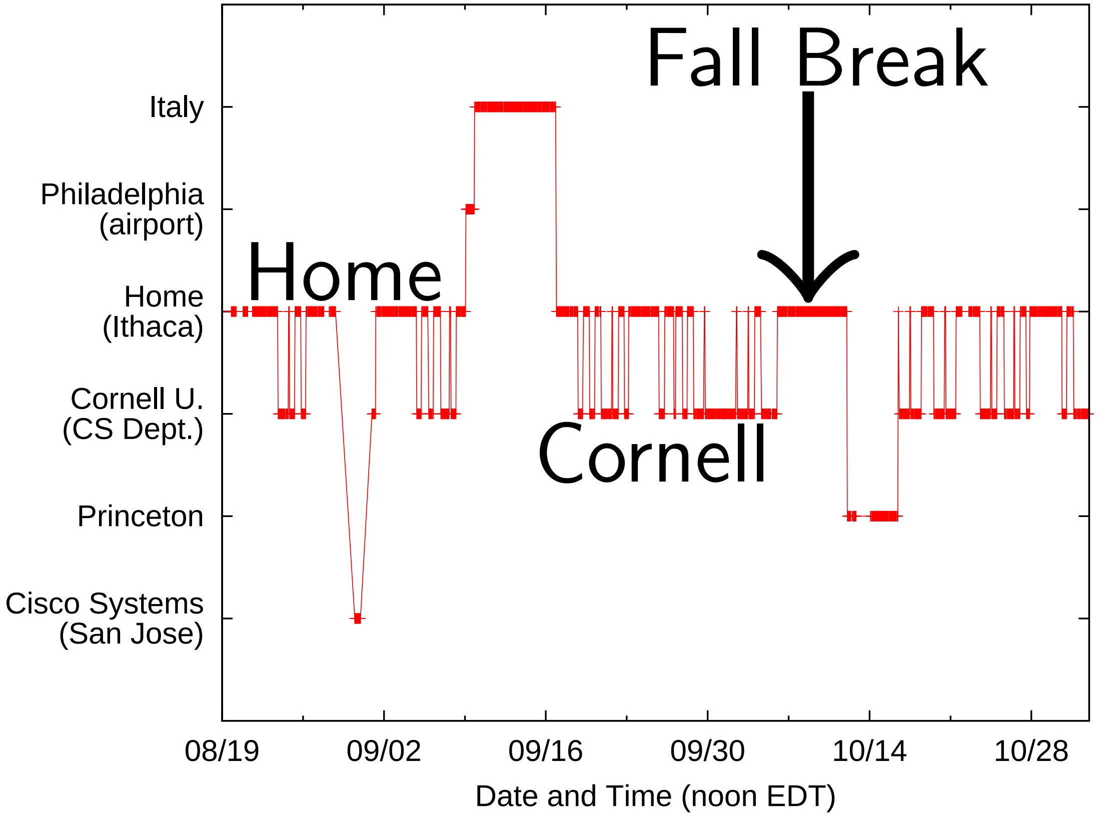

- Trailed Paul
- City-level accuracy in US (∼100 mi), province-level in Italy for GeoIP service used.
- Commute time accurate to within query interval. Some exceptions.

Reasonably good accuracy. Reconstructed travel itineraries, daily commute patterns.

Validation: Accuracy

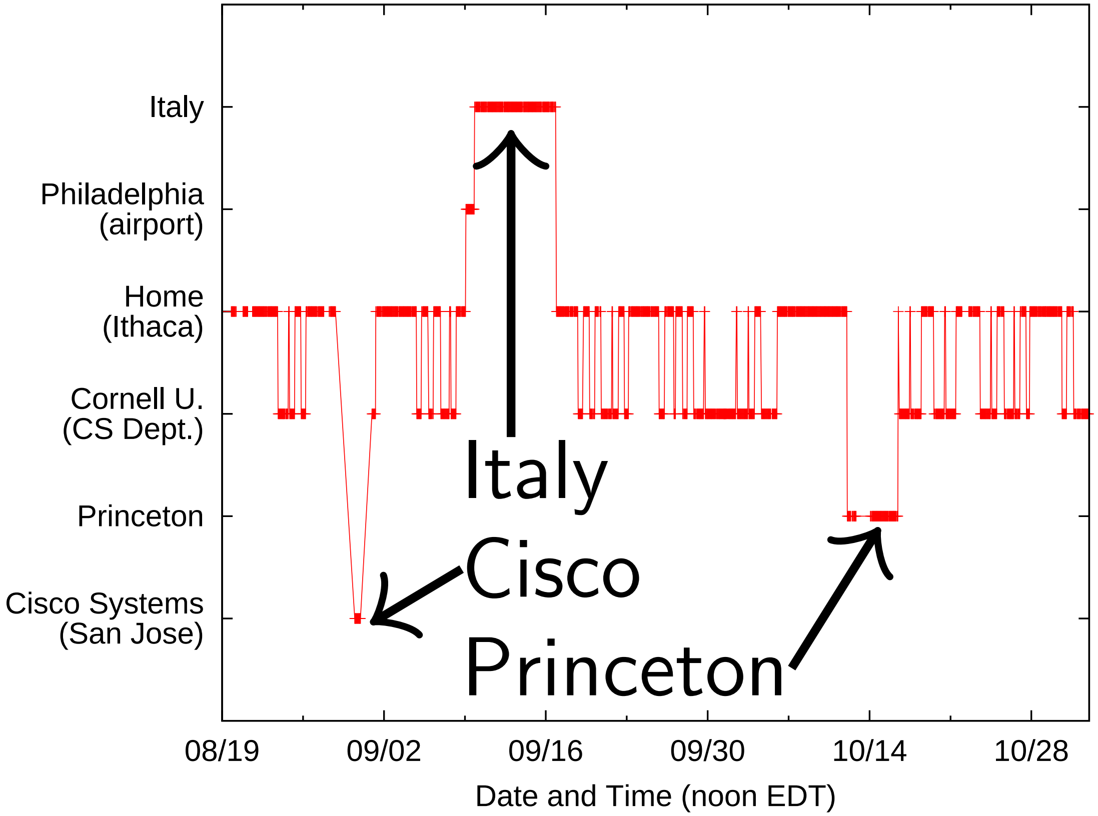

- Trailed Paul
- City-level accuracy in US (∼100 mi), province-level in Italy for GeoIP service used.
- Commute time accurate to within query interval. Some exceptions.

Reasonably good accuracy. Reconstructed travel itineraries, daily commute patterns.

Validation: Accuracy

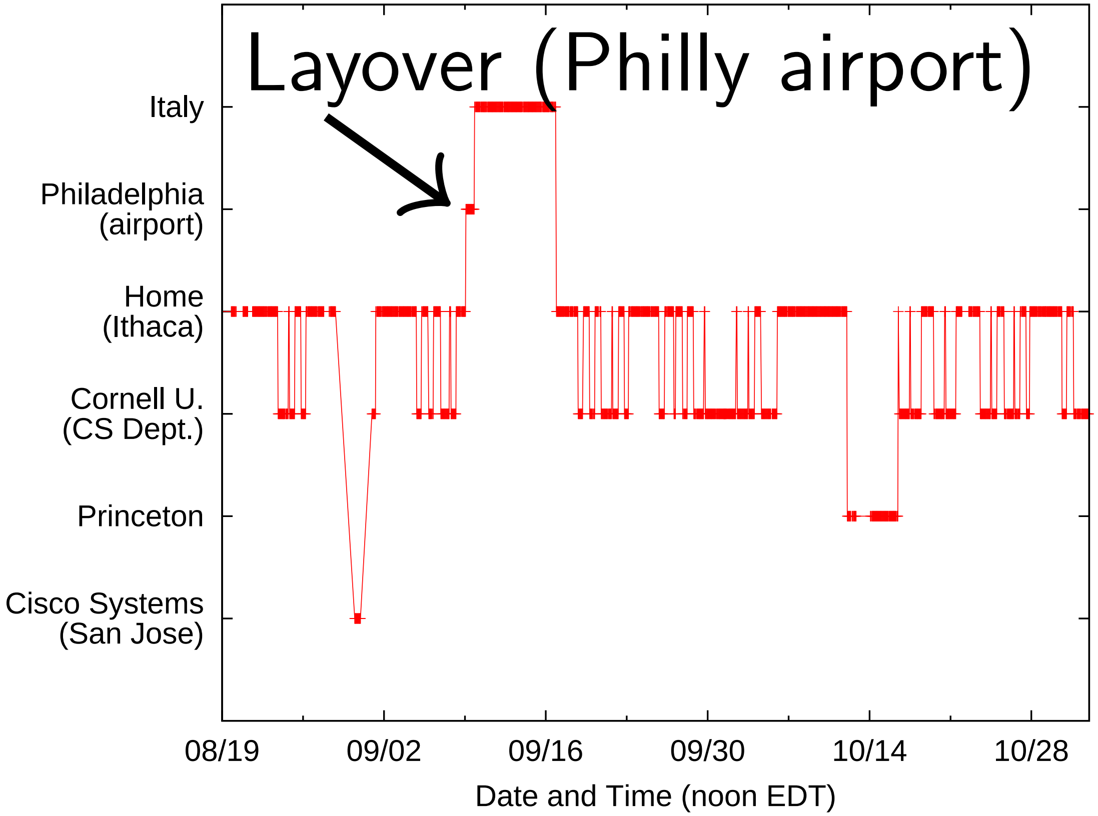

- Trailed Paul
- City-level accuracy in US (∼100 mi), province-level in Italy for GeoIP service used.
- Commute time accurate to within query interval. Some exceptions.

Reasonably good accuracy. Reconstructed travel itineraries, daily commute patterns.

End-Middle-End Name Resolution [1]

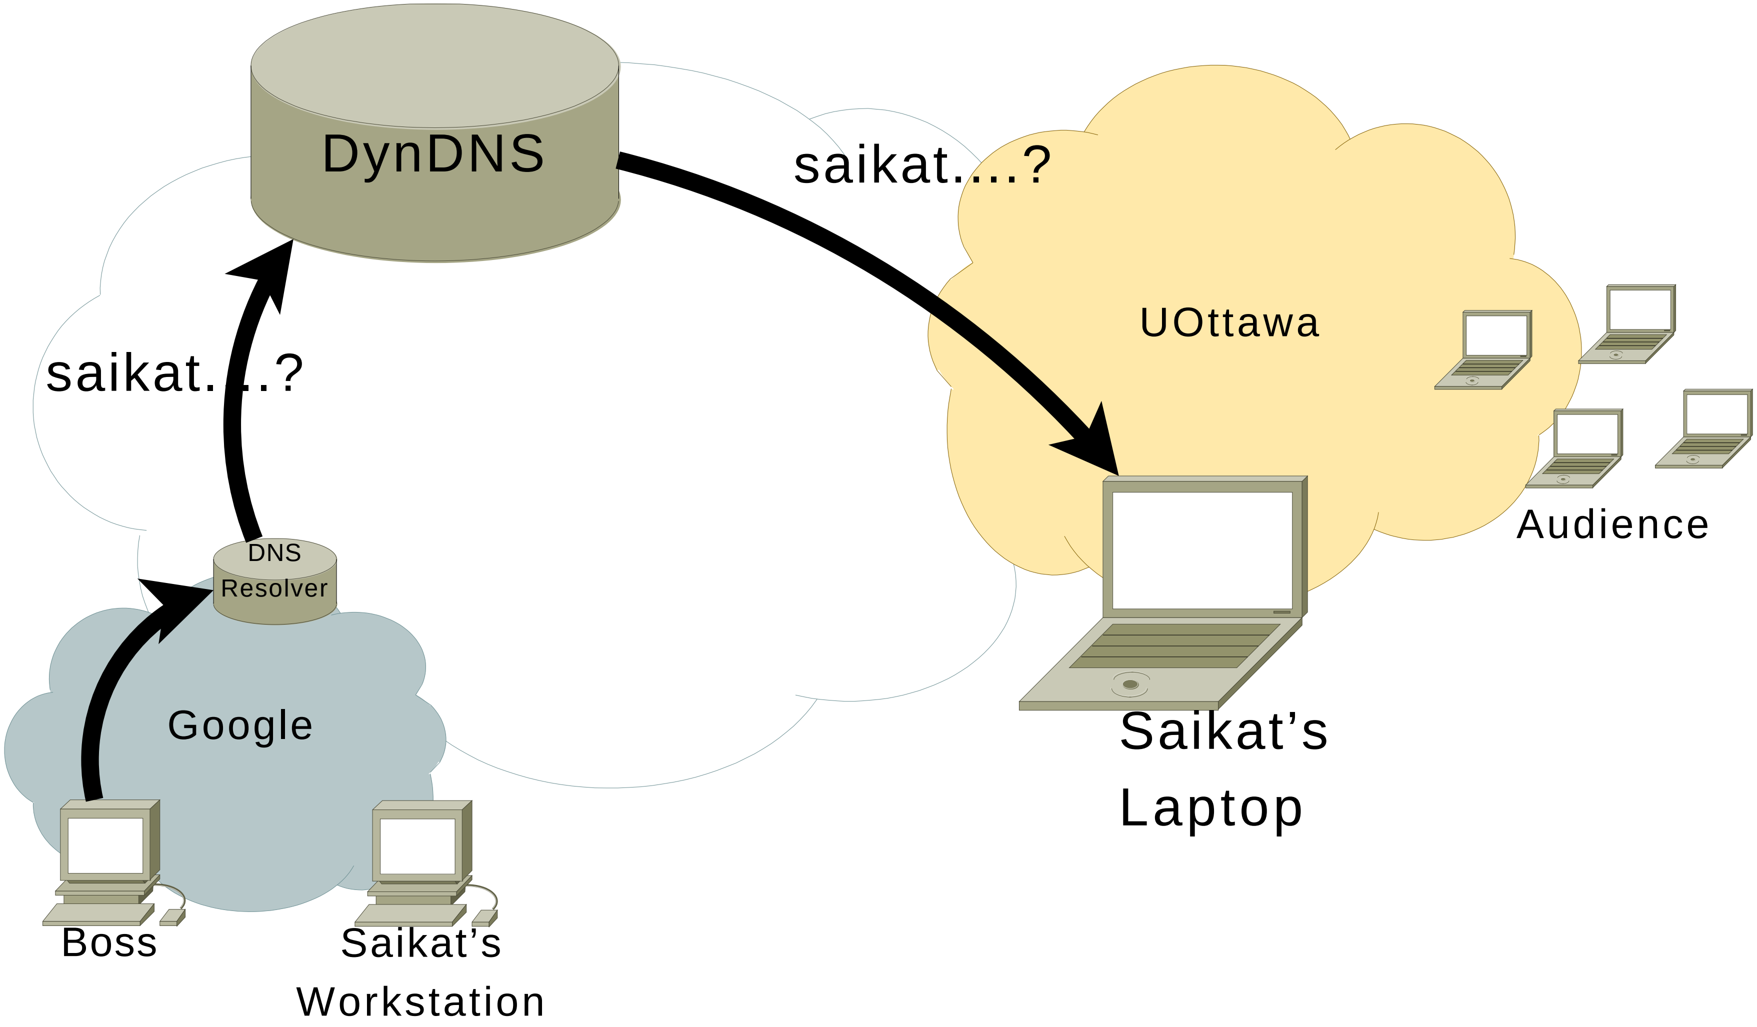

[1] Conceptually draws from ongoing EME research [SIGCOMM'07]

End-Middle-End Name Resolution [1]

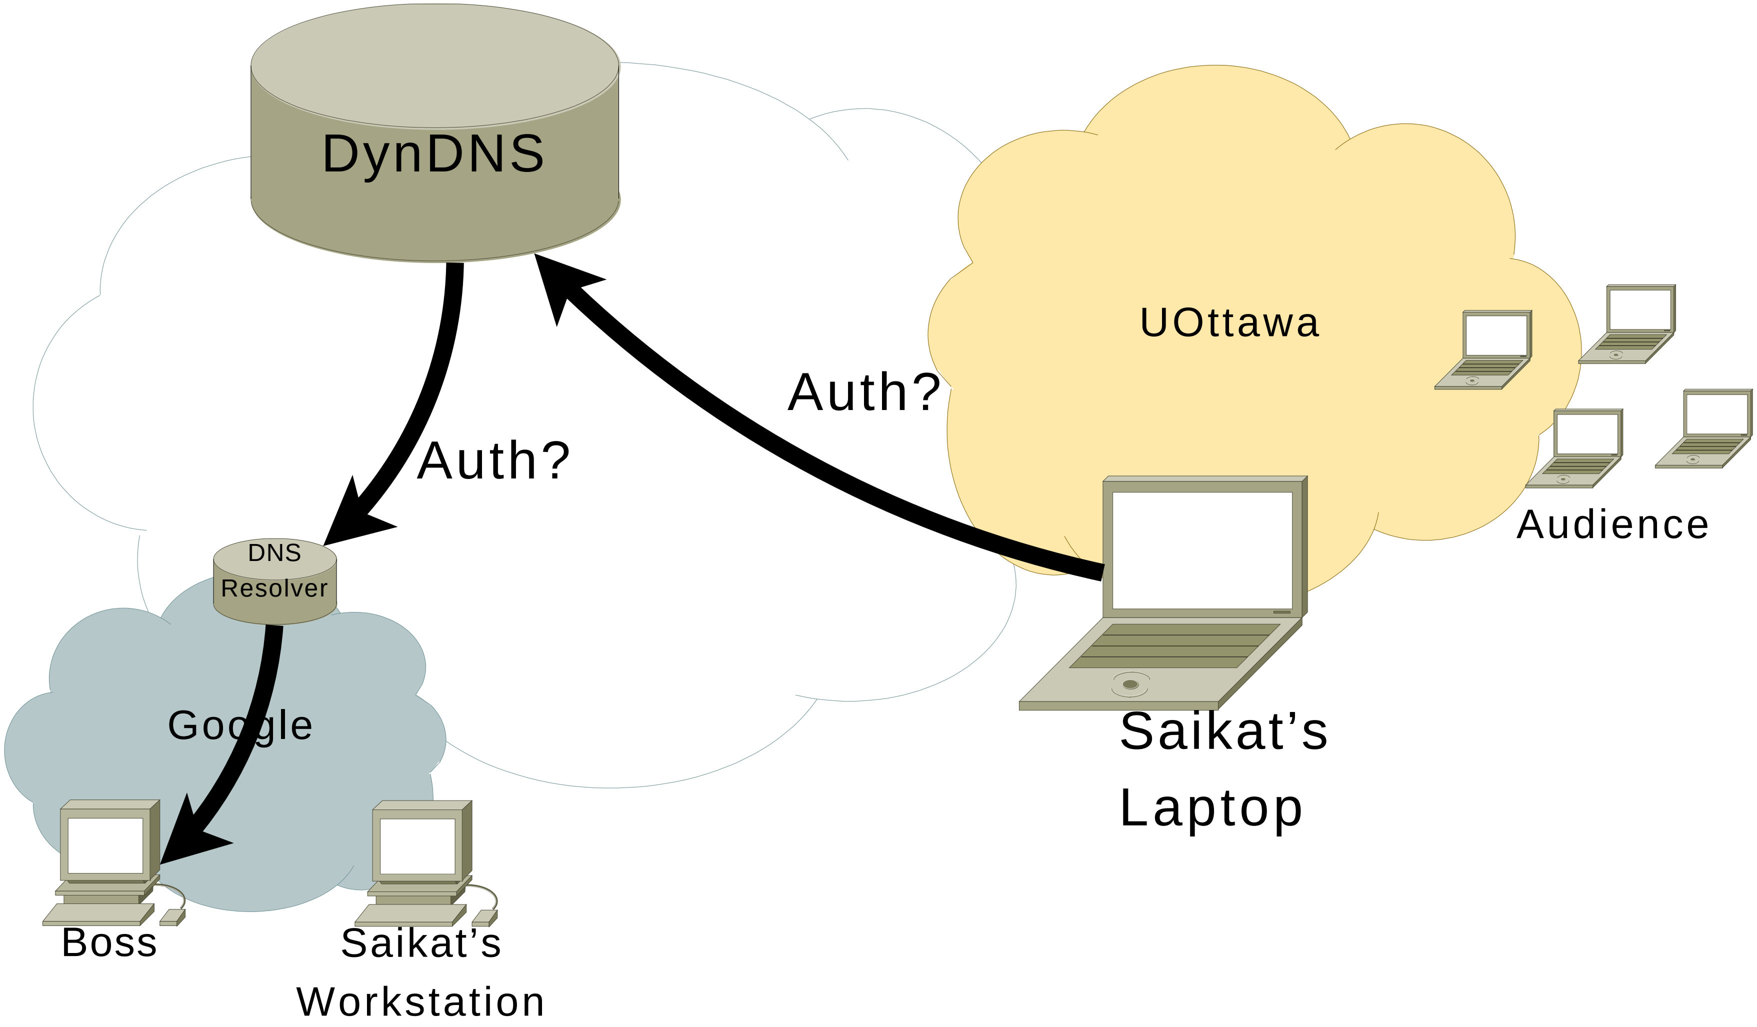

[1] Conceptually draws from ongoing EME research [SIGCOMM'07]

End-Middle-End Name Resolution [1]

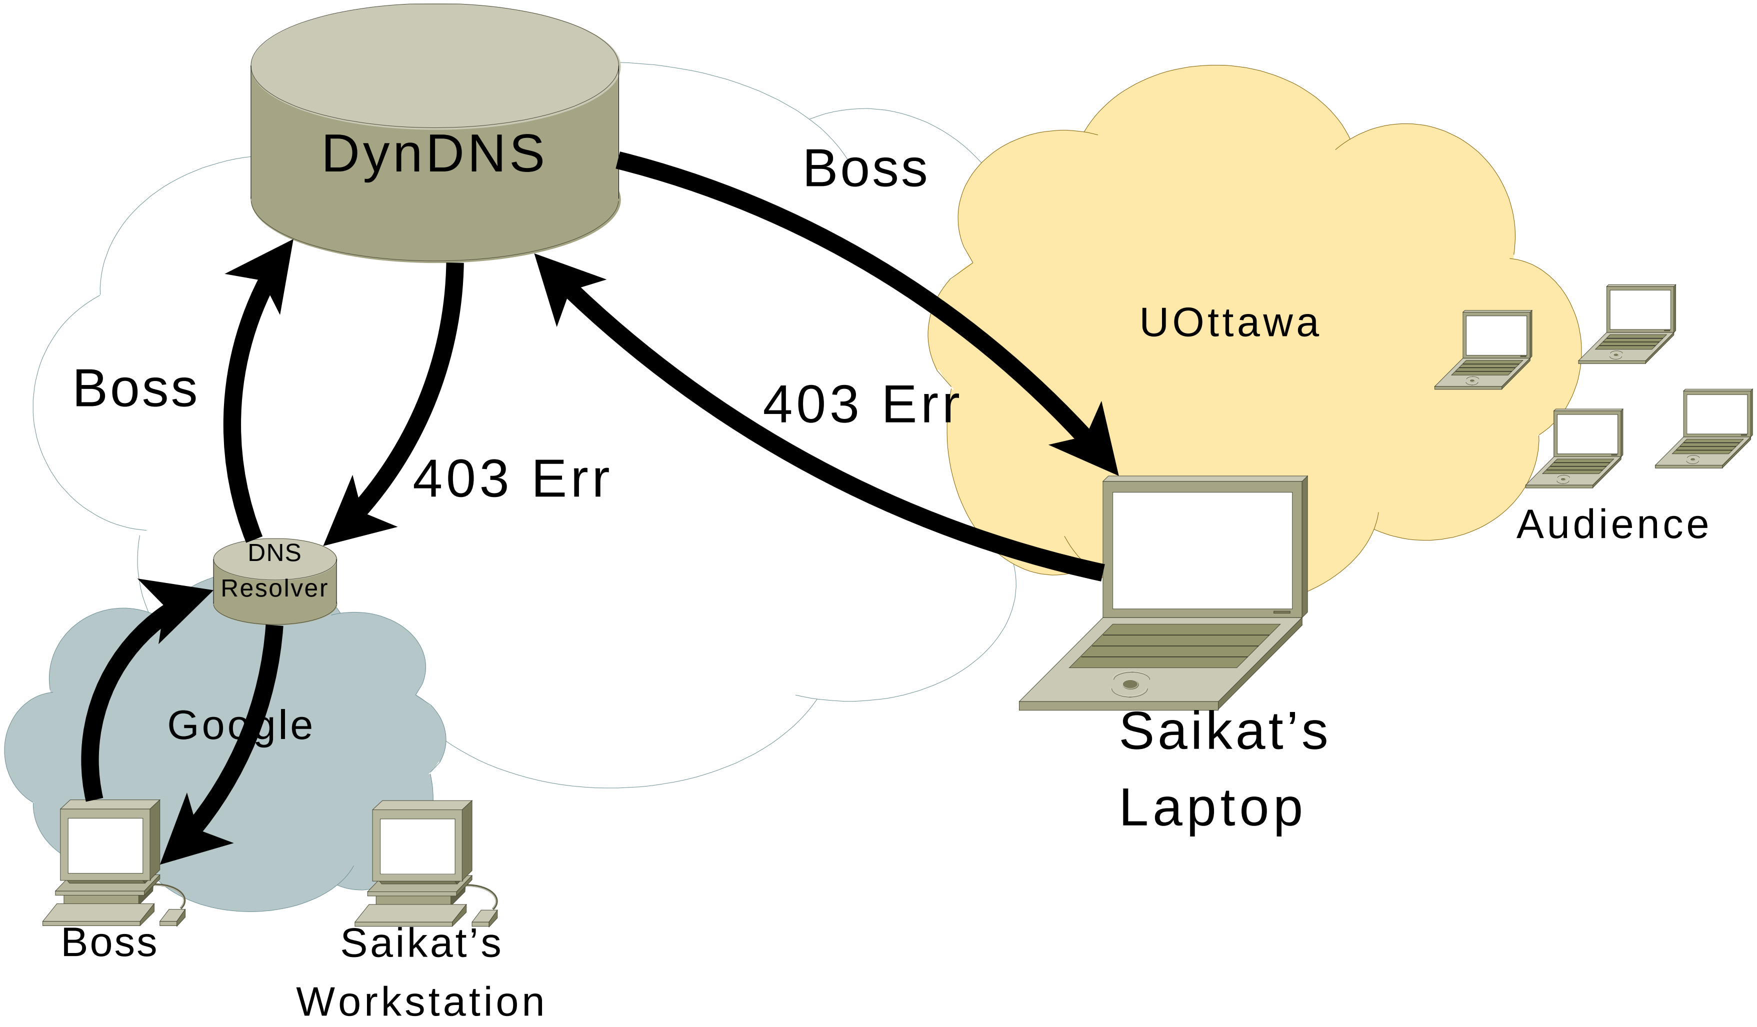

[1] Conceptually draws from ongoing EME research [SIGCOMM'07]

Summary

- Identity trail attack collects private information of mobile dynamic DNS users
- Performed covertly; demonstrated for over 100K users
- Alternative user-friendly name-resolution needed for private hosts.
- End-middle-end signaling may be a solution.

`http://nutss.net/whereissaikat`
`http://nutss.net/whereispaul`

Backup Slides: Mobility

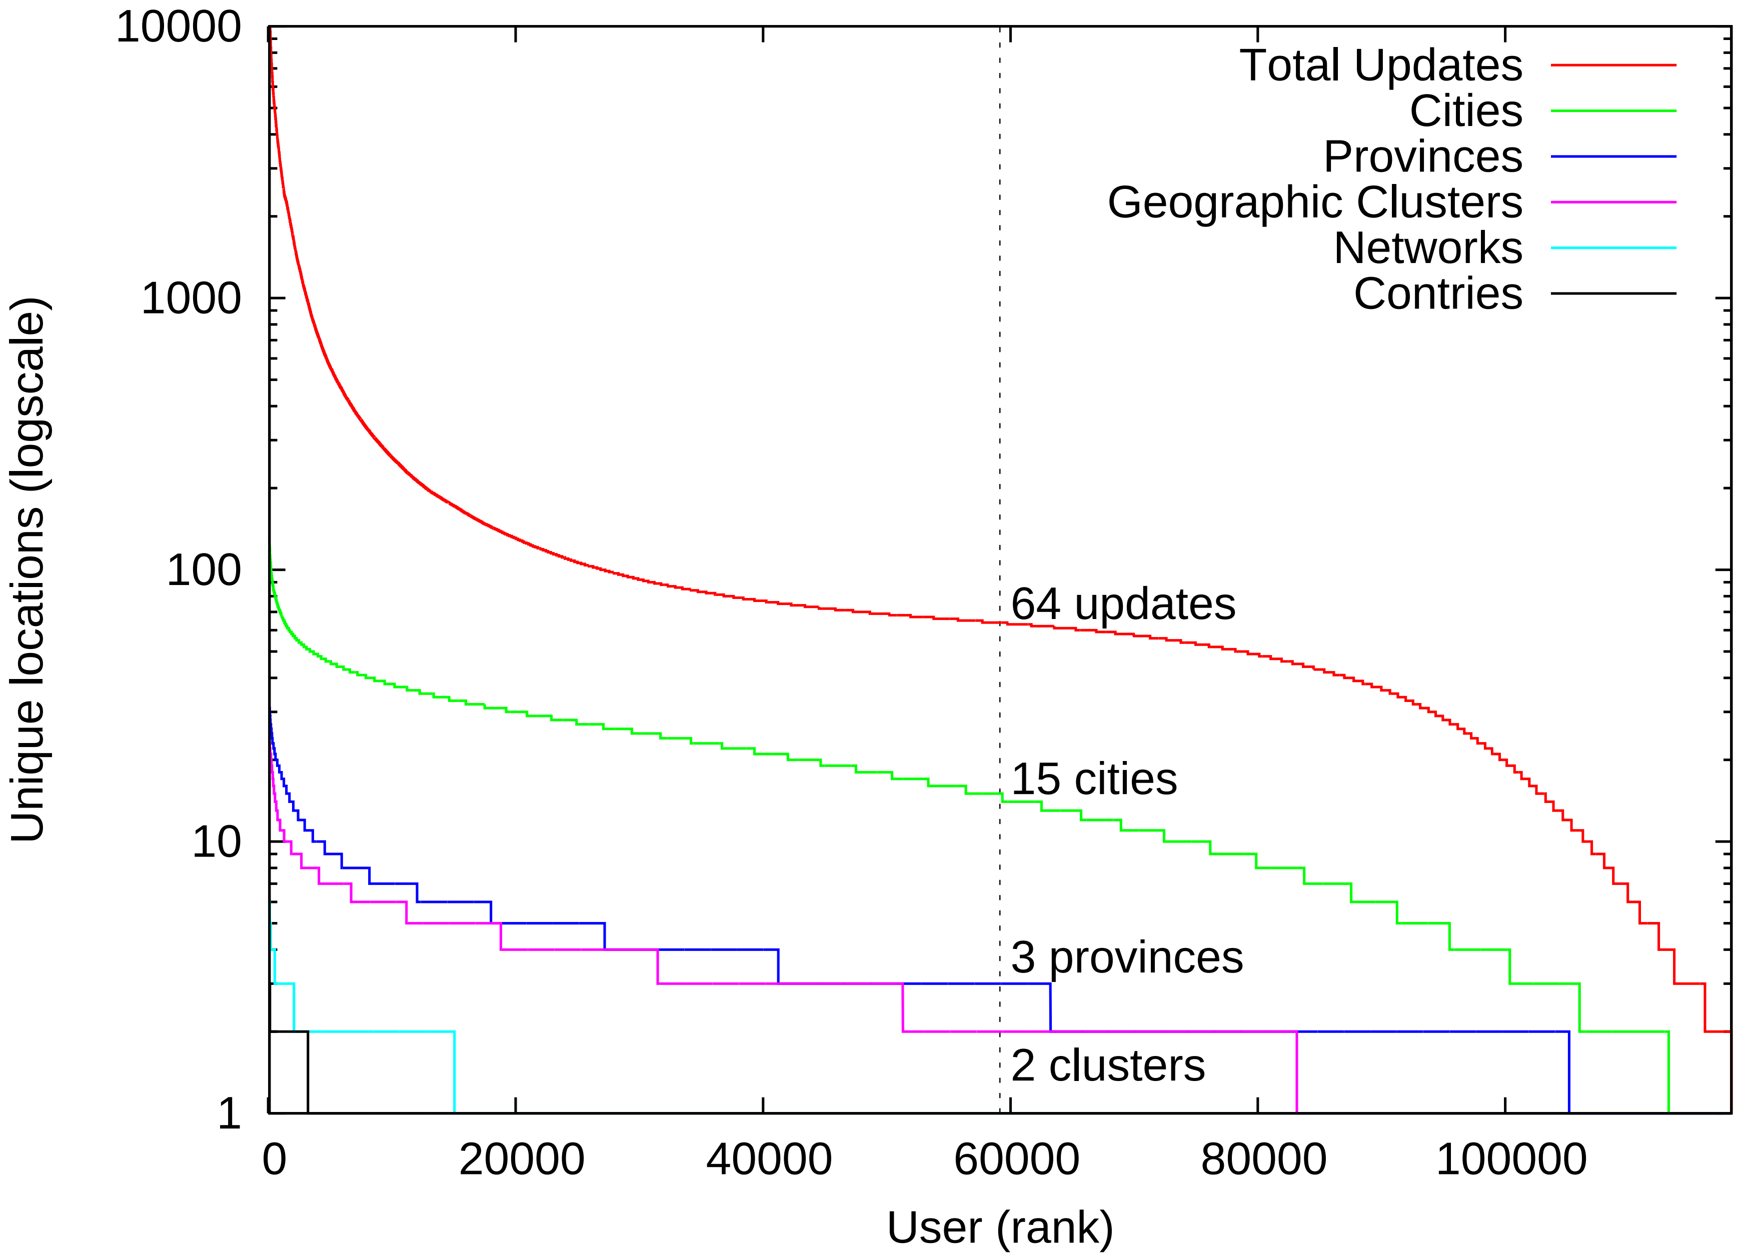

Backup Slides: Non-solutions

- Don't use DNS for mobile private hosts
  - Try `http://saikat.dyndns.org`.
    You will connect to this laptop. Without DNS need to memorize IP addresses (IPv6 even).
- Use a proxy like Akamai
  - HTTP/FTP only. No service for individuals.
- Encrypt IP addresses in DNS
  - Key management headaches