

DEFCON 16:

How to evade geospatial intrusion detection techniques

Ryan W. Trost

Agenda

- GIS
- Geospatial Intrusion Detection implementations
- Geospatial Intrusion Detection methodology
- Accuracy of IP -> lat/long translation
- Okay...so how do I beat it?
- Q&A

GIS

- GIS (Geographic Information Systems): computer based methodology to collect, store, manipulate, retrieve, display and analyze georeferenced data.
 - GoogleEarth
 - ESRI
 - Intergraph
 - GRASS

GIS

- Traditional GIS tools focused more on environmental issues
- These days thanks to GoogleEarth and/or GoogleMaps the average Internet user is starting to be exposed to the power that mapping software unleashes.
 - Track a cell phone (cell tower triangulation)
 - Track flights in mid-air!!!

Cell Phone



Track Any Flight



Where has IT security and mapping collided?

- Multiple security firms have implemented GIS tools in their products (in varying degrees)
 - MeerCat (Secure Decisions)
 - VisCat (ETRI)
 - GeoSWAT

MeerCat

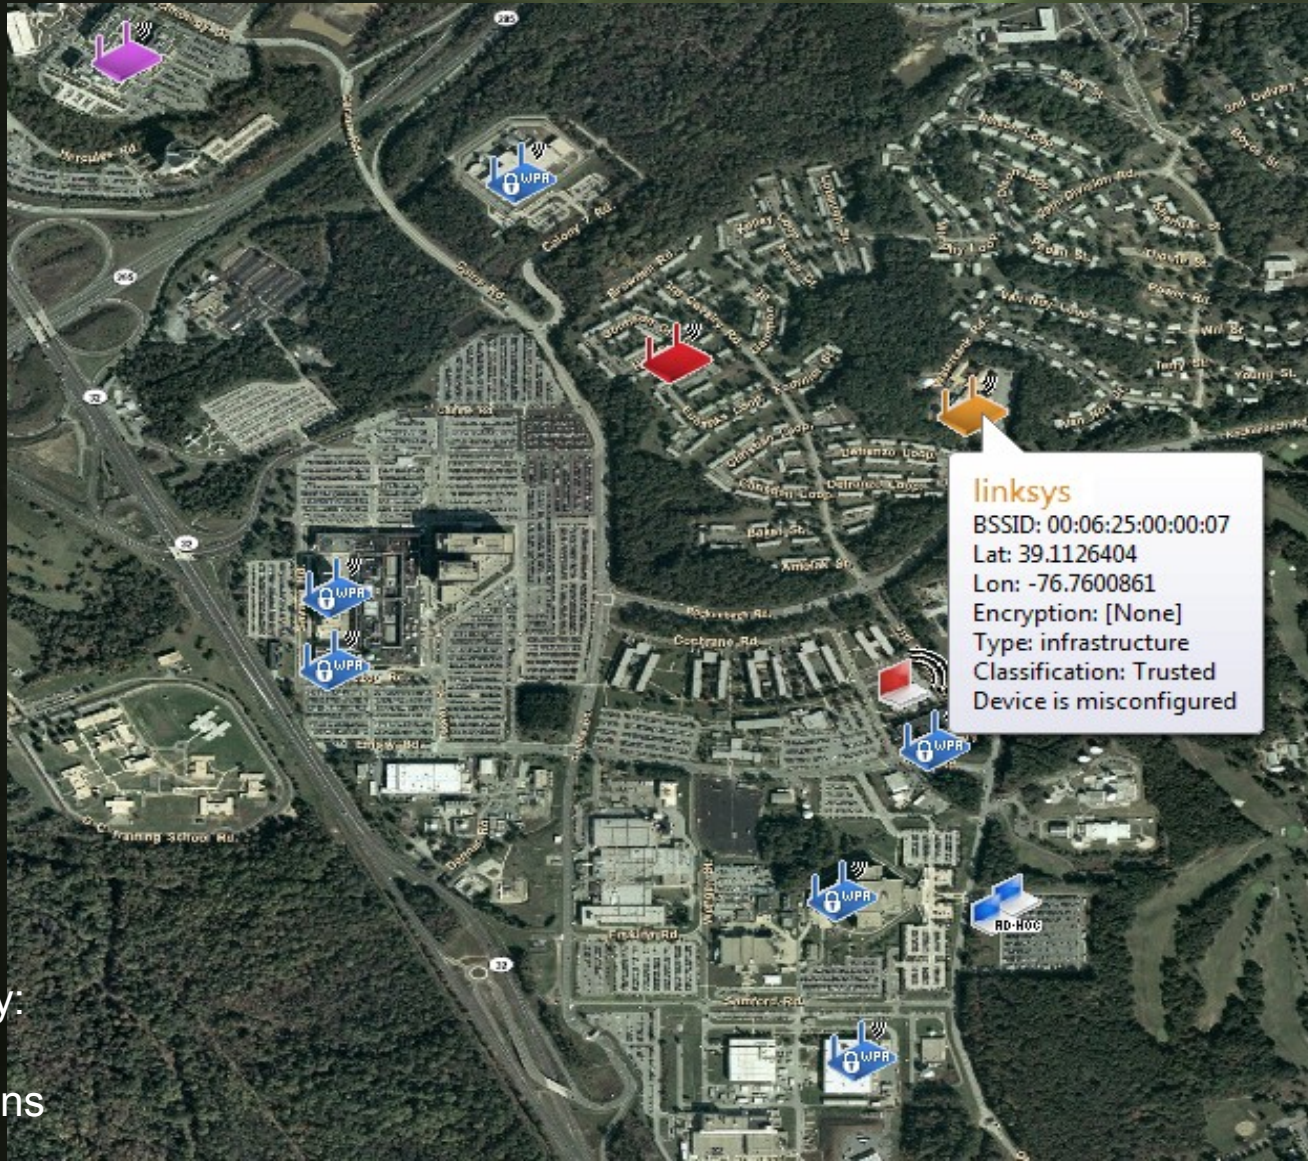
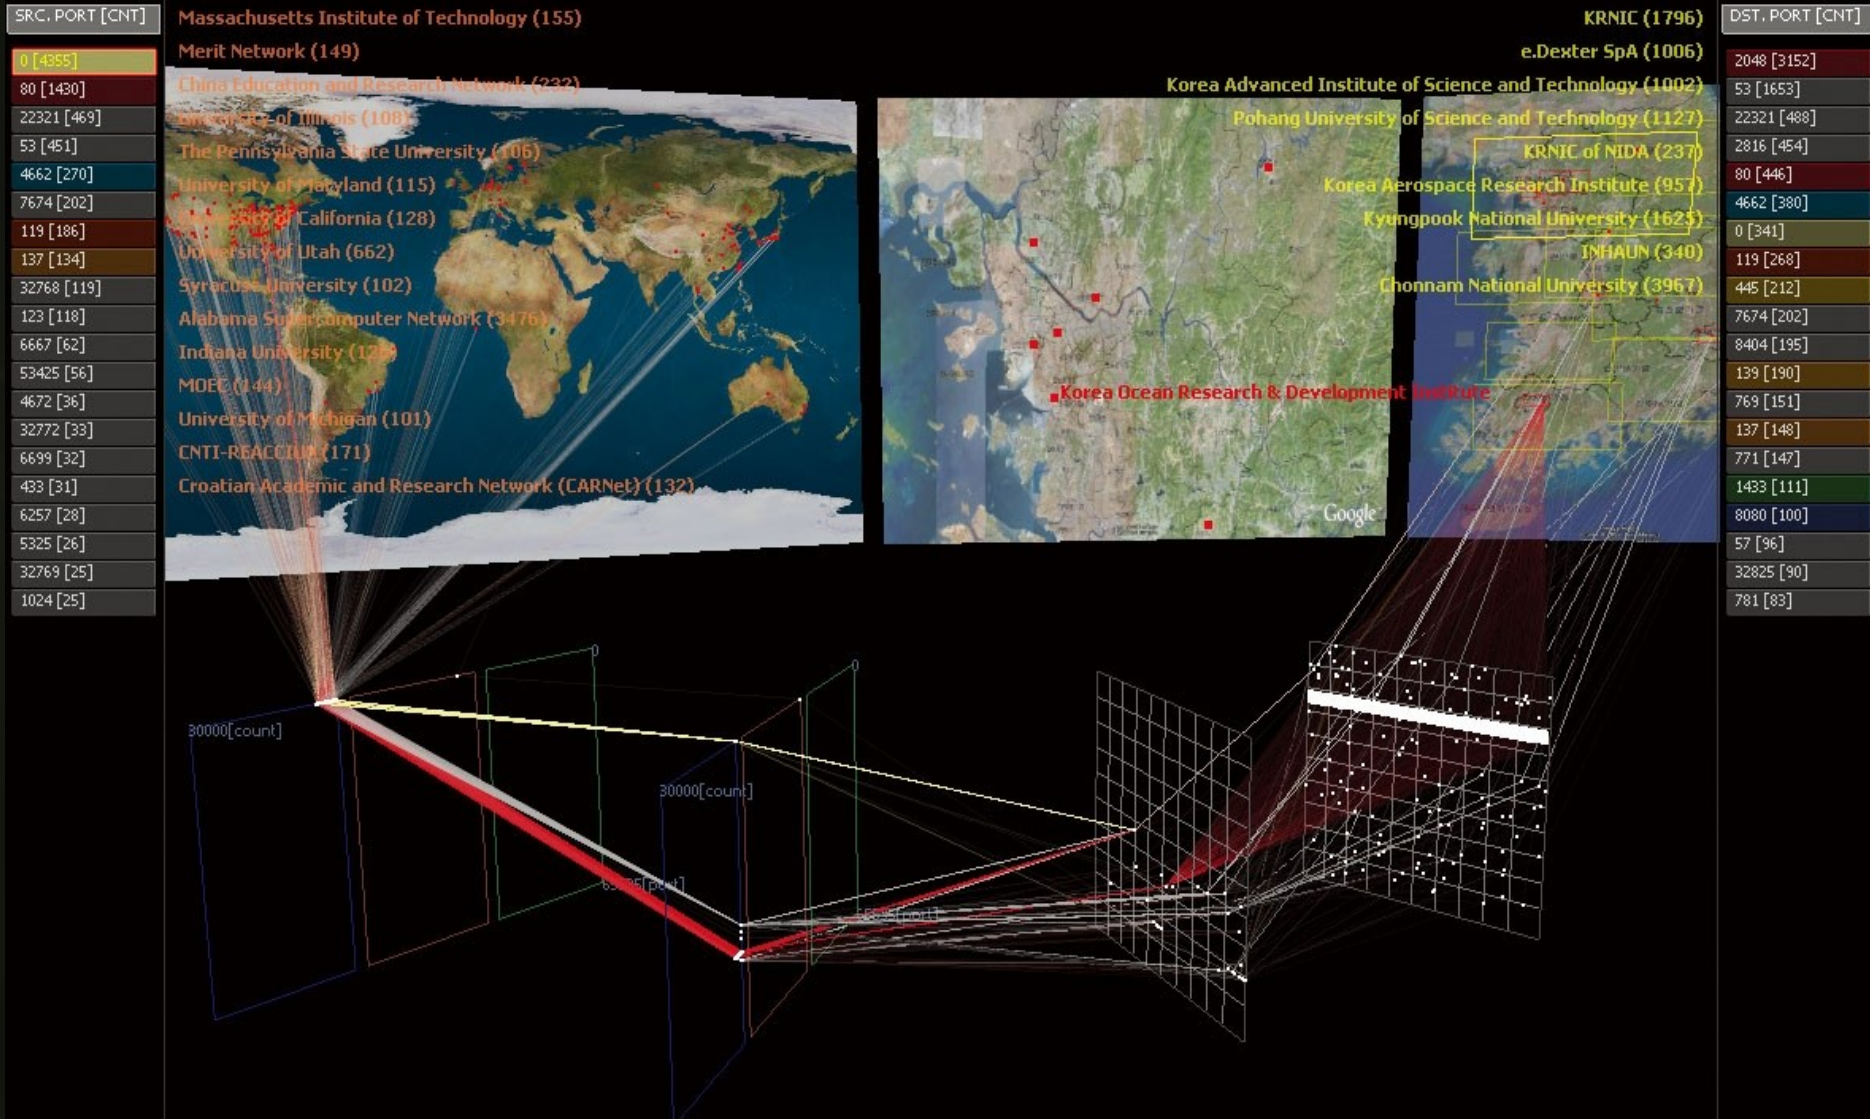
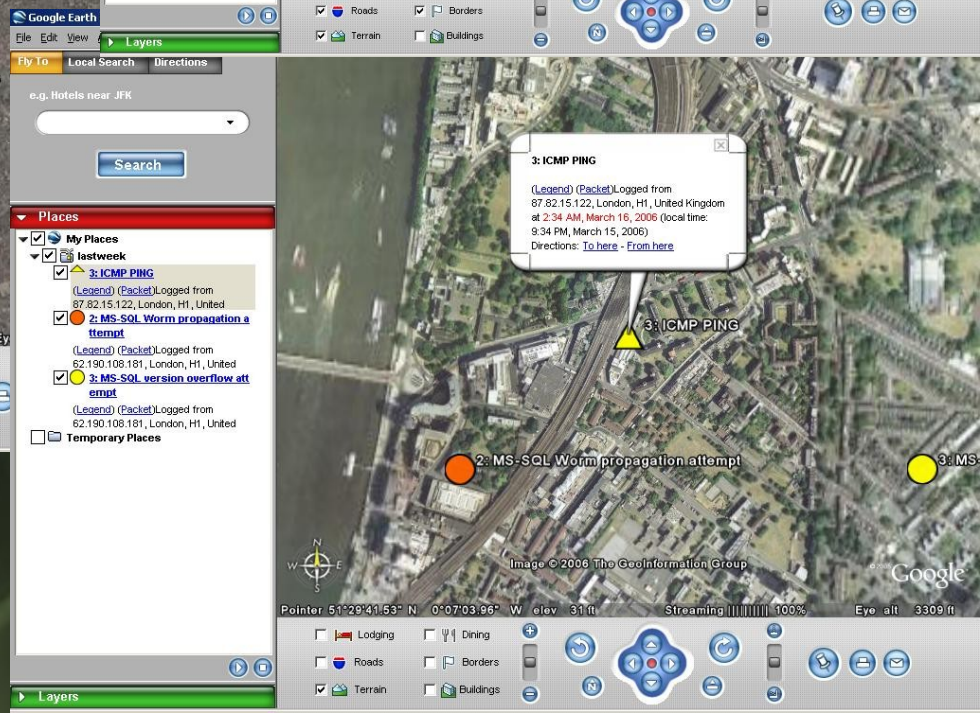
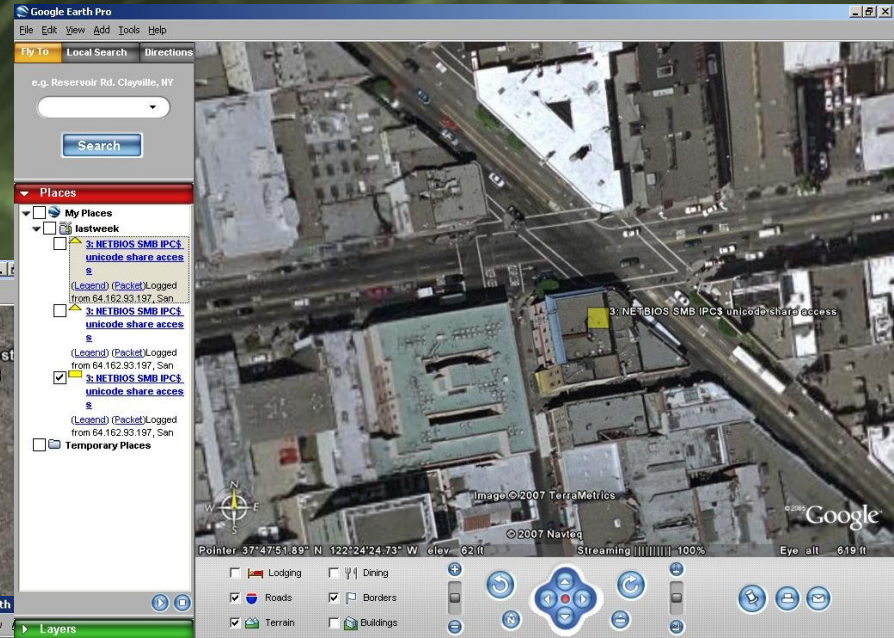
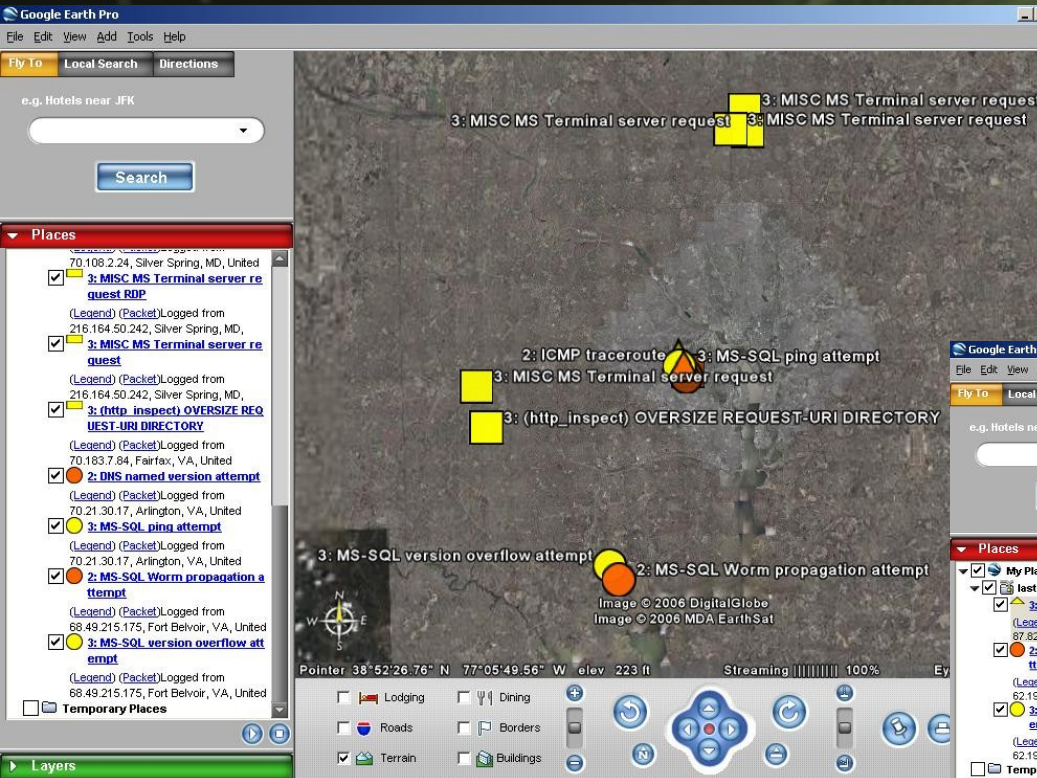


Image Courtesy:
John Goodall
Secure Decisions

VizCat



GeoSWAT



Geospatial Intrusion Detection

- Goal: find a direct correlation between externally based network alerts by plotting their source locations on a geographic map

How does it work?

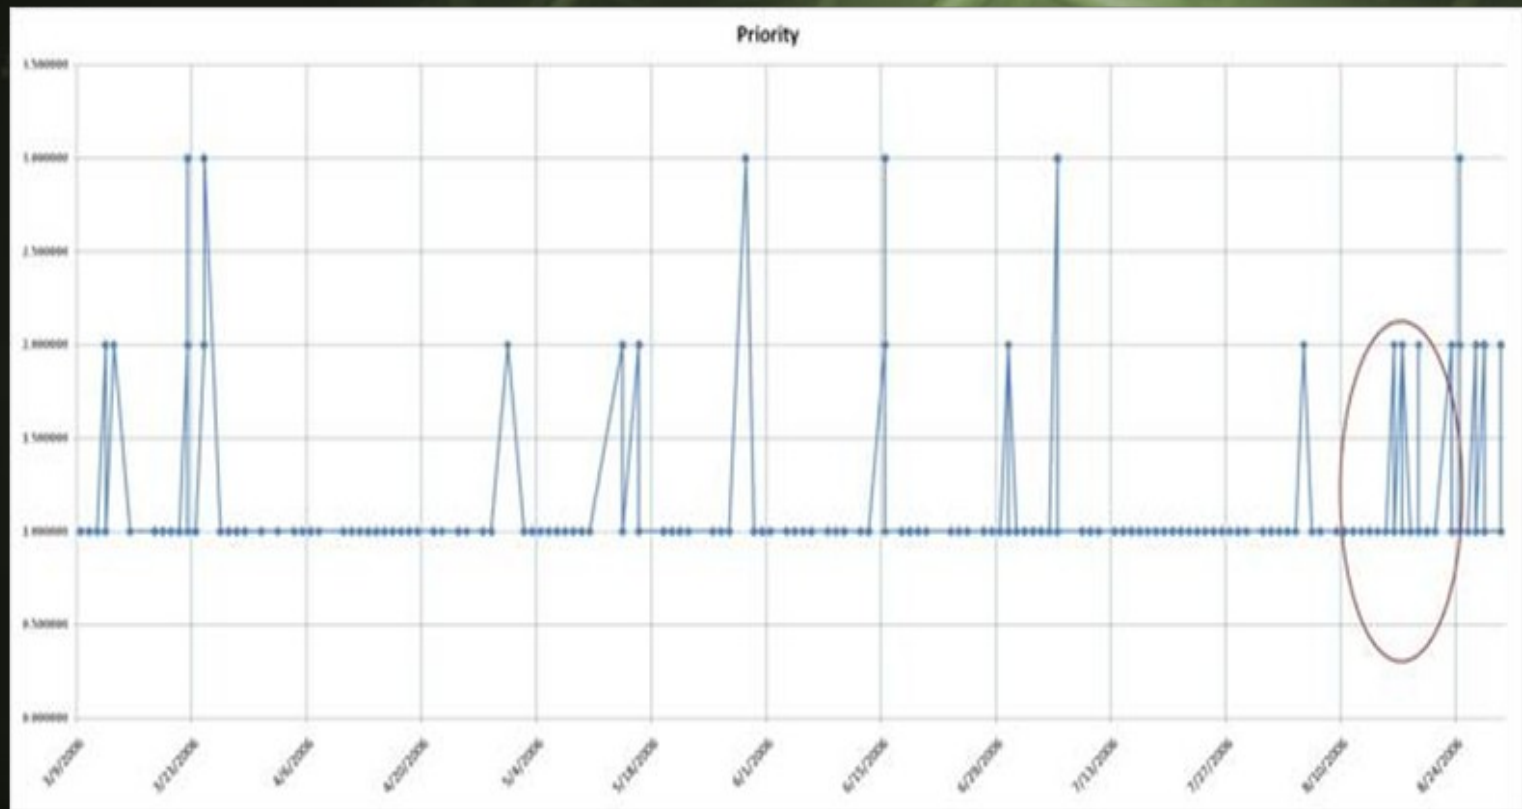
- High-level 'roadmap':
 1. Plot rolling time period (1 week/2 weeks/4 weeks)
 2. Eliminate 'friendlies' to reduce IDS false positives
 - 2.1 Geographically plot remote branches, SOHO, business partners locations by street address (very accurate in GoogleEarth)
 - 2.2 Create an IDS alert that is triggered when a customer authenticates to a website

Geospatial Intrusion Detection

3. Run a clustering algorithm on plotted data
 - There are several different clustering algorithms to choose from:
 - Poisson, nearest-neighbor, Moran's I Index, Ripley's K Function, Getis-Ord
4. Extract network alerts within identified 'hot-spot'
5. Run a weight calculating algorithm to evaluate if there is a relationship between alerts
 - Correlating elements in an alert
 - Alert severity
 - Destination ports
 - Timestamp

1. Plot rolling timeseries

14 day intervals (the slow probe theory)




2. Eliminate 'friendlies'

Map Legend

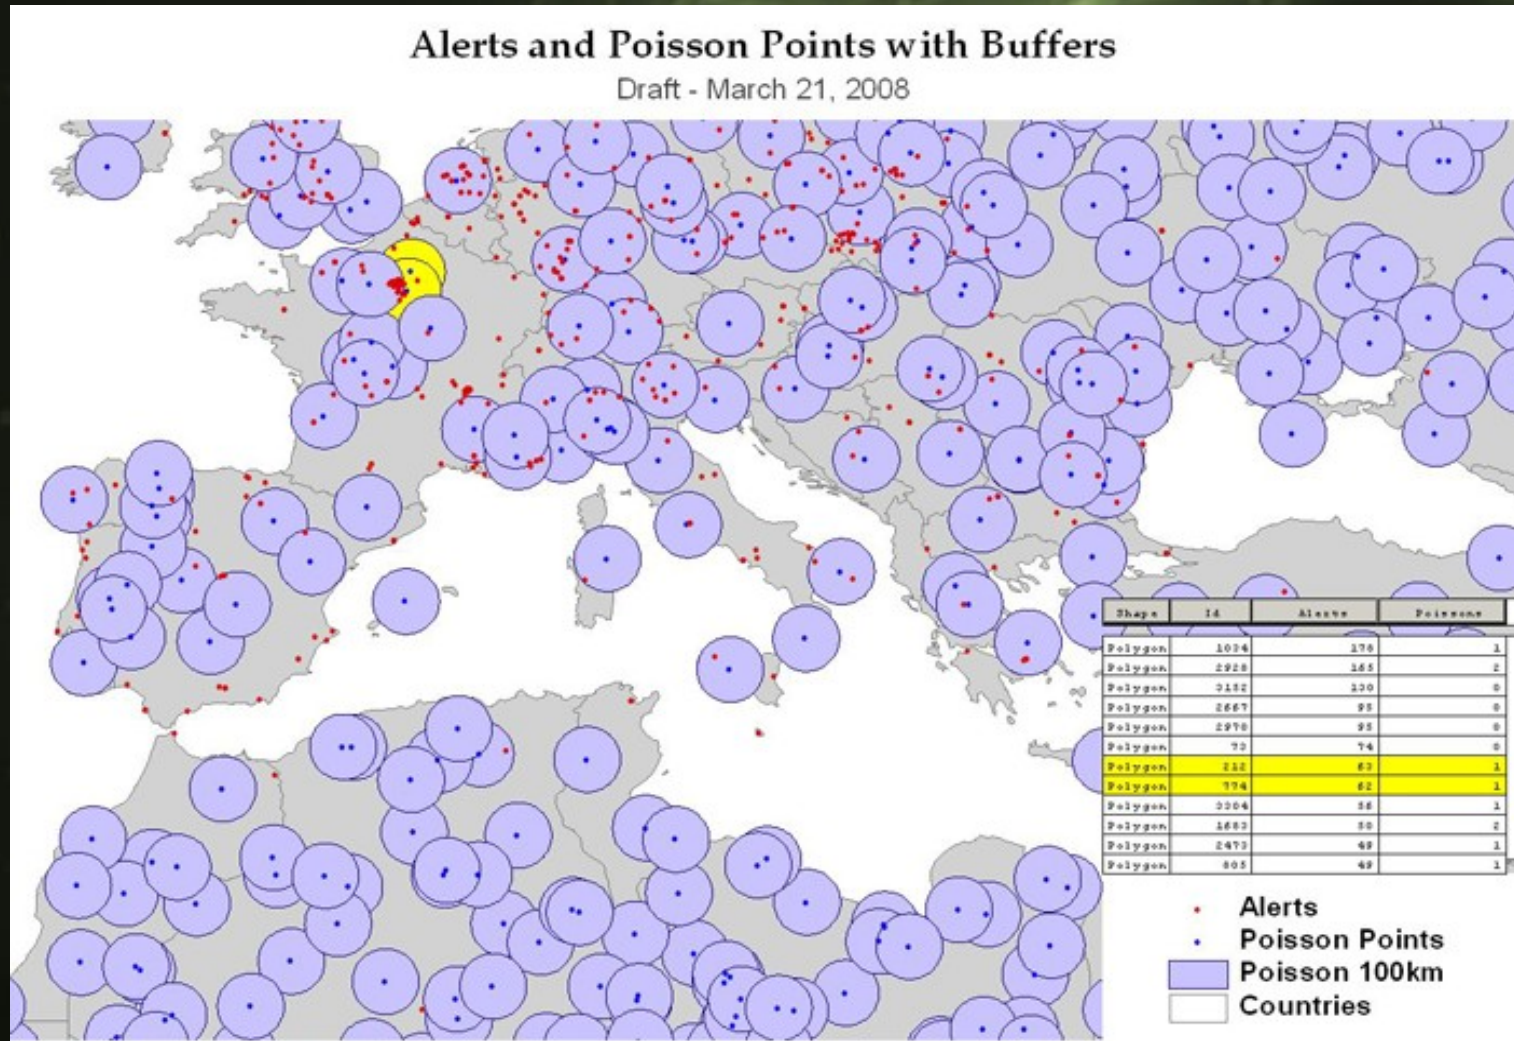
△ = Source IDS Alert

● = Remote Offices or
Telecommuting
Residences




Depending on your
risk comfort abilities
this would eliminate
~30% of potential
false positives.

3. Run a GIS clustering algorithm



4. Extract 'hot-spot' network alerts



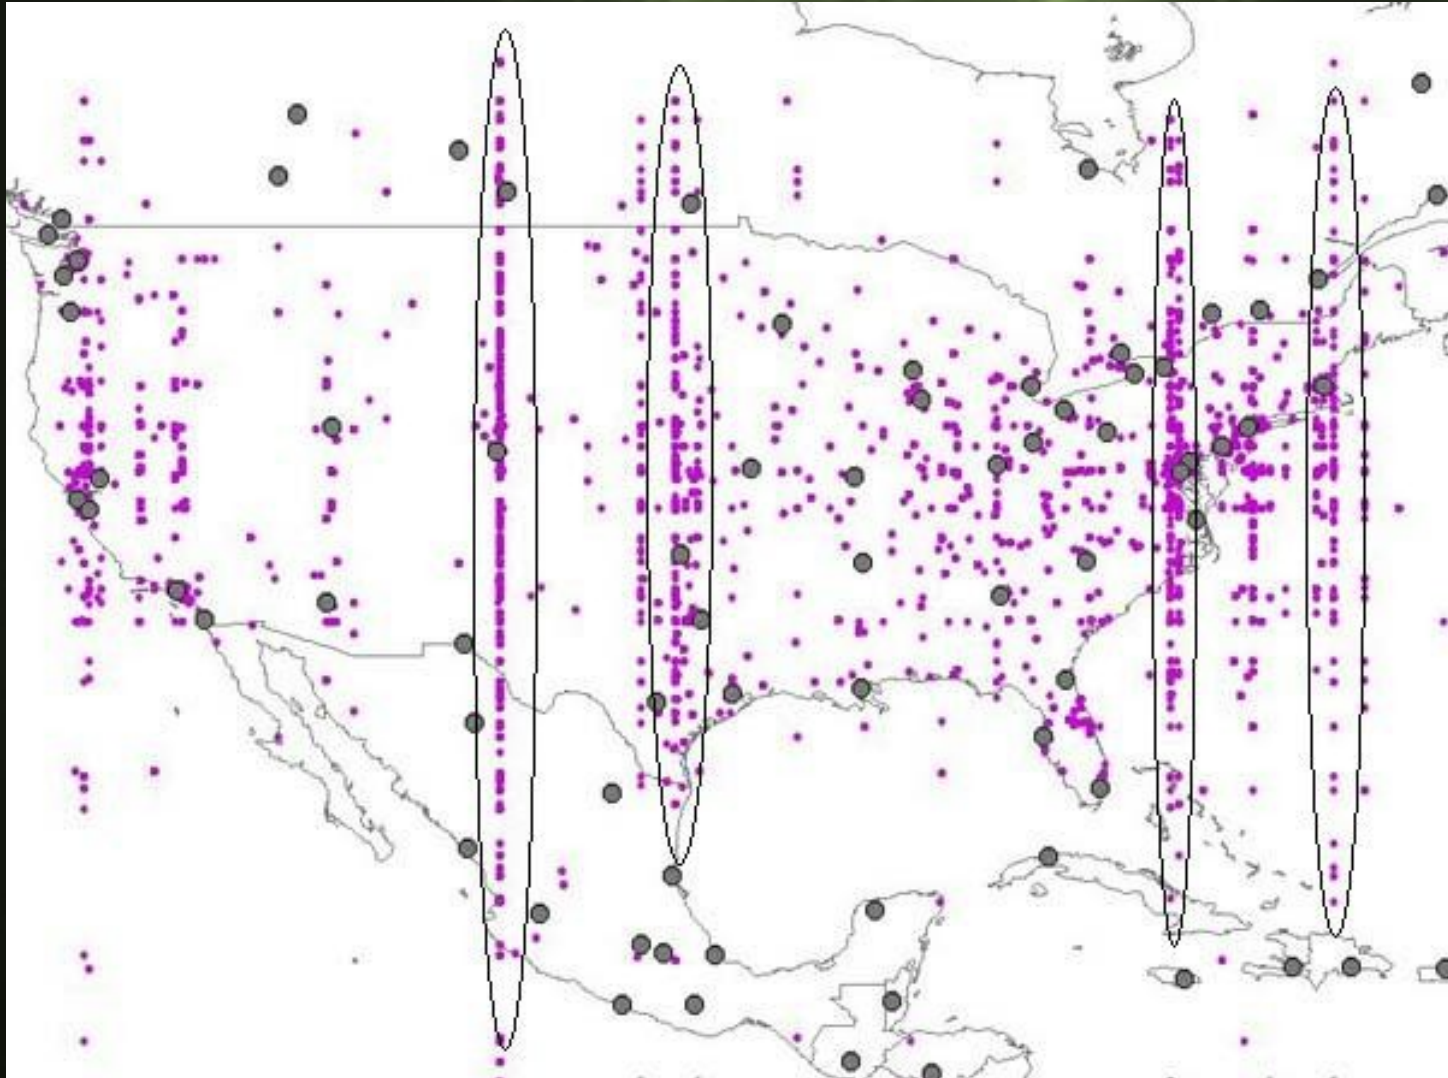
5. Run a weight calculating algorithm to evaluate if there is a relationship between alerts

- Insert image(s) later...

Accuracy translation

- Vendors
 - IP2Location
 - MaxMind
 - Quova (wireless capabilities)
- How is the translation calculated?
 - Domain scrapping
 - Compare Traceroutes
 - Strategic partnerships with ISPs
 - Strategic partnerships with downloading FTP sites

Example of 'less' accurate translation



Okay...how do I beat it?

- As with many (if not all) defenses there are always loopholes:
 - Most scanning/enumeration tools primarily do sequential scans of IP address – DO NOT USE SEQUENTIAL IPs TO ATTACK A VICTIM
 - Map remote locations and use a tool to extract neighboring IP addresses – which will hopefully get extracted when eliminating ‘friendlies’
 - Attack from random geographic locations and with varying times

Deconstruct the translation file

- Insert image of deconstructed application

Q&A

Back to the Guinness tap...