IETF 90 Technical Plenary, Toronto, ON, CA

NETWORK TOPOLOGY AND GEOGRAPHY

Speakers

Antonio Gamba-Bari, IXmaps Project

Jane Coffin, ISOC

Amogh Dhamdhere, CAIDA

IXmaps.ca

Mapping internet routing and surveillance from a user privacy point of view

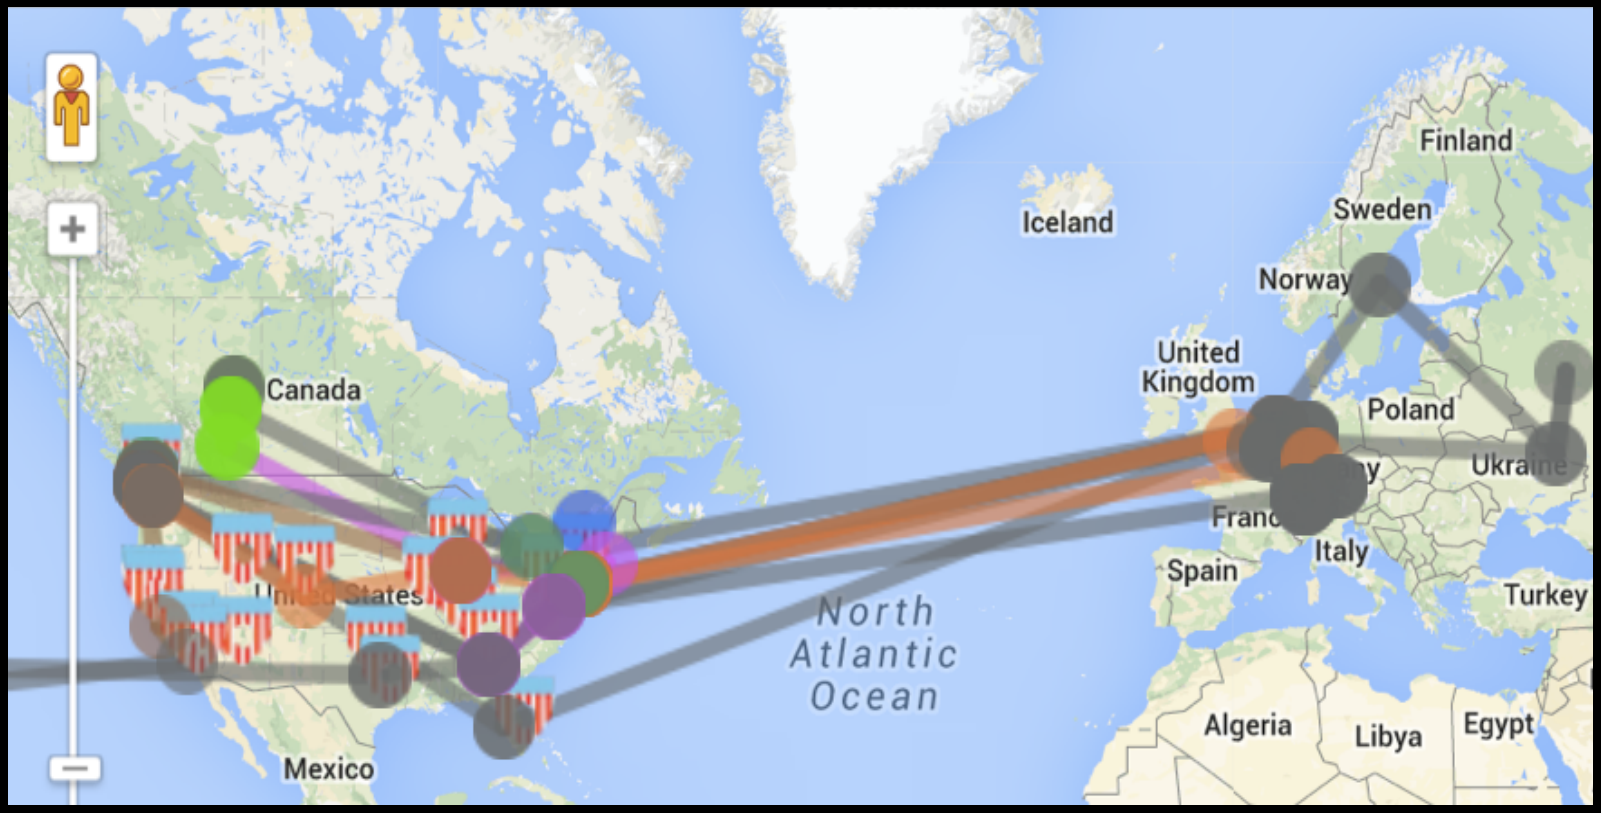

Antonio Gamba, Colin McCann,
Andrew Clement, Jonathan Obar
Faculty of Information, University of Toronto

Network topology and geography panel
Technical Plenary, IETF90
Toronto, Canada, July 21, 2014

IXmaps – Internet Exchange mapping

- Crowdsourced traceroute generation & collection
  > 30,000 TRs > 250 contributors/origins > 2,500 URLs
- Systematic geo-location of (core) routers
- Map traceroute paths via GoogleMaps/Earth
  – NSA surveillance splitter sites, carrier transparency + …
- Custom filtering of traceroutes
  – NSA interception, 'boomerang routing,' ISP, city, ….

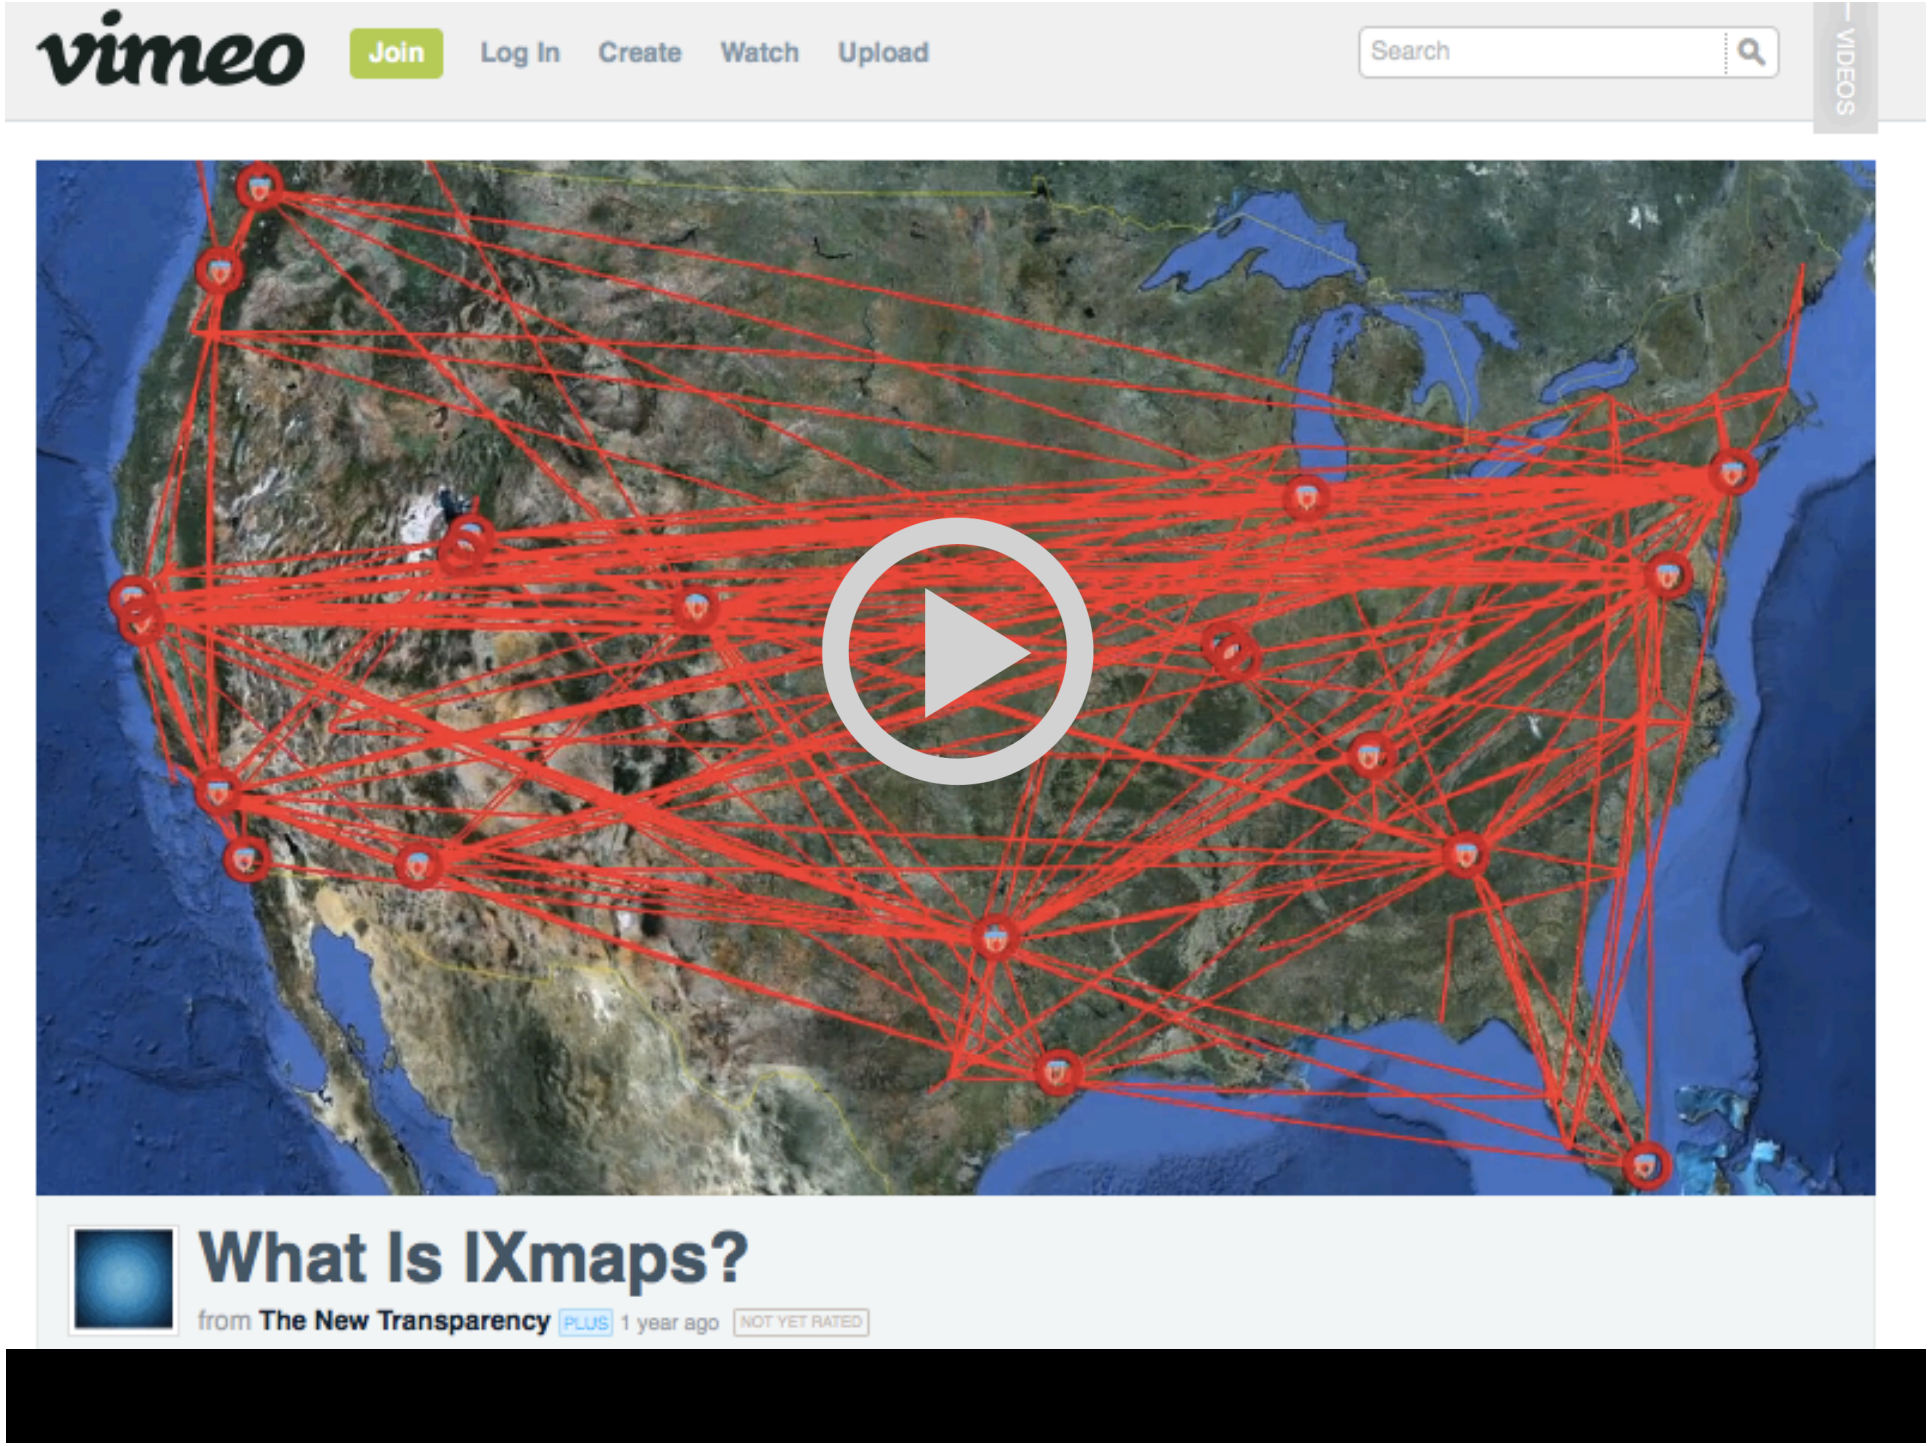
vimeo
Join
Log In
Create
Watch
Upload
Search
VIDEOS
What Is IXmaps?
from The New Transparency
PLUS
1 year ago
NOT YET RATED

Notable results/implications

US NSA interception

- Comprehensive Continental US coverage
- Shows where your packets get ‘split’ by the NSA

Canadian Boomerang routings

- ~25% Canadian traffic?
- High risk of NSA interception, depending on carriers involved

Reveals carriers/ISPs,
comparative privacy transparency assessments

- Part of wider transparency and accountability initiatives

Current development

With support of CIRA grant - see RFP

- Re-build TR generation
- Re-build geo-location

With further support?

- Internationalization
- Sustainability/FLOSS migration

See where your packets go!
(and contribute to the database)

Note: RFP for re-building traceroute generation and geo-location modules

http://IXmaps.ca

Work supported by the Social Sciences and Humanities Research Council (SSHRC) and Office of the Privacy Commissioner (OPC)

Internet Society (ISOC): Internet Exchange Point (IXP) – Global Development Work

Jane Coffin and Christian O'Flaherty (on-site)

Many other ISOCers and Partners around the world (off-site)

July 2014

Internet Society™

What is an IXP

- An Internet Exchange Point (IXP) is a **physical location** where different IP networks **meet to exchange traffic** (switch, routers, cabling, ports) with each other to keep local traffic local. **BUT they are much more than just "boxes and wires":**

- IXPs are **vital part of the Internet ecosystem**, essential for facilitating a robust domestic ICT sector

- Benefits of an Internet Exchange Point (IXP):

  - **Keeps local Internet traffic within a local infrastructure, and reduces costs** associated with traffic exchange between networks.

  - **Builds local Internet community and develops** human technical **capacity** – better net management skills and routing

  - **Improves the quality of Internet services and drive demand** in by reducing delay and improving end-user experience

  - **Convenient hub for attracting hosting key Internet infrastructures** within countries – **content is key and confidence** builds in local infra when delivery is consistent and reliable

  - **Catalyst** for overall Internet development

Measuring the Benefits and Impacts of IXPs: Kenya and Nigeria Case Study

Benefit	*KIXP*	*IXPN*	*Summary*
Latency	Reduced from 200-600 ms to 2-10 ms	Reduced from 200-400 ms to 2-10 ms	Noticeable increase in performance for end users
Local traffic exchange	1 Gbit/s peak	300 Mbit/s peak	Savings on international transit of over $1 million per year in each country
Content	Google network present locally, along with rehoming of domestic content	Same as in Kenya	Increase in usage and corresponding revenues for mobile data traffic
E-government	Kenya Revenue Authority gathers taxes online	Usage by education and research networks	Social benefits from e-government access to IXPs
Other benefits	An increasing amount of regional traffic exchanged at KIXP	Financial platforms hosted locally	Further economic benefits resulting from IXPs

- **Reduced latency and increased performance** and driving demand
- Direct **savings on international transit** ($1.5M p.a. Kenya, $1M Nigeria)
- **Facilitating e-government and education** services
- **Catalyzing local hosting and content** industry
- **Increased mobile data market** by an estimated $6 million in Kenya
- **Attracting regional traffic - KIXP**
- http://www.internetsociety.org/ixpimpact

LAC IXP Study| November 2013

- **LAC Findings:**
  - Argentina: In one city → $100.00Mbps pre IXP/ down to $40.00Mbps post IXP
  - Brazil: NIC.br | PPT Metro System 26 IXPs attracting investment/content | 600Gbps at Peak
  - Ecuador: (Pre) International transit was $100 Mbps | (post) Local traffic costs $1.00 Mbps
    - Now running RPKI
    - After CDN cache installed in Quito in 2009 -> traffic up 700%
  - **Additional Studies:**
  - **Measurement Study** in Bolivia | Raspberry Pi deployment
  - **Network efficiency** Study in Argentina | Cabase and University of Buenos Aires

LAC IXP Study can be found here: http://bit.ly/1k6NaO0

IXPs Around the World

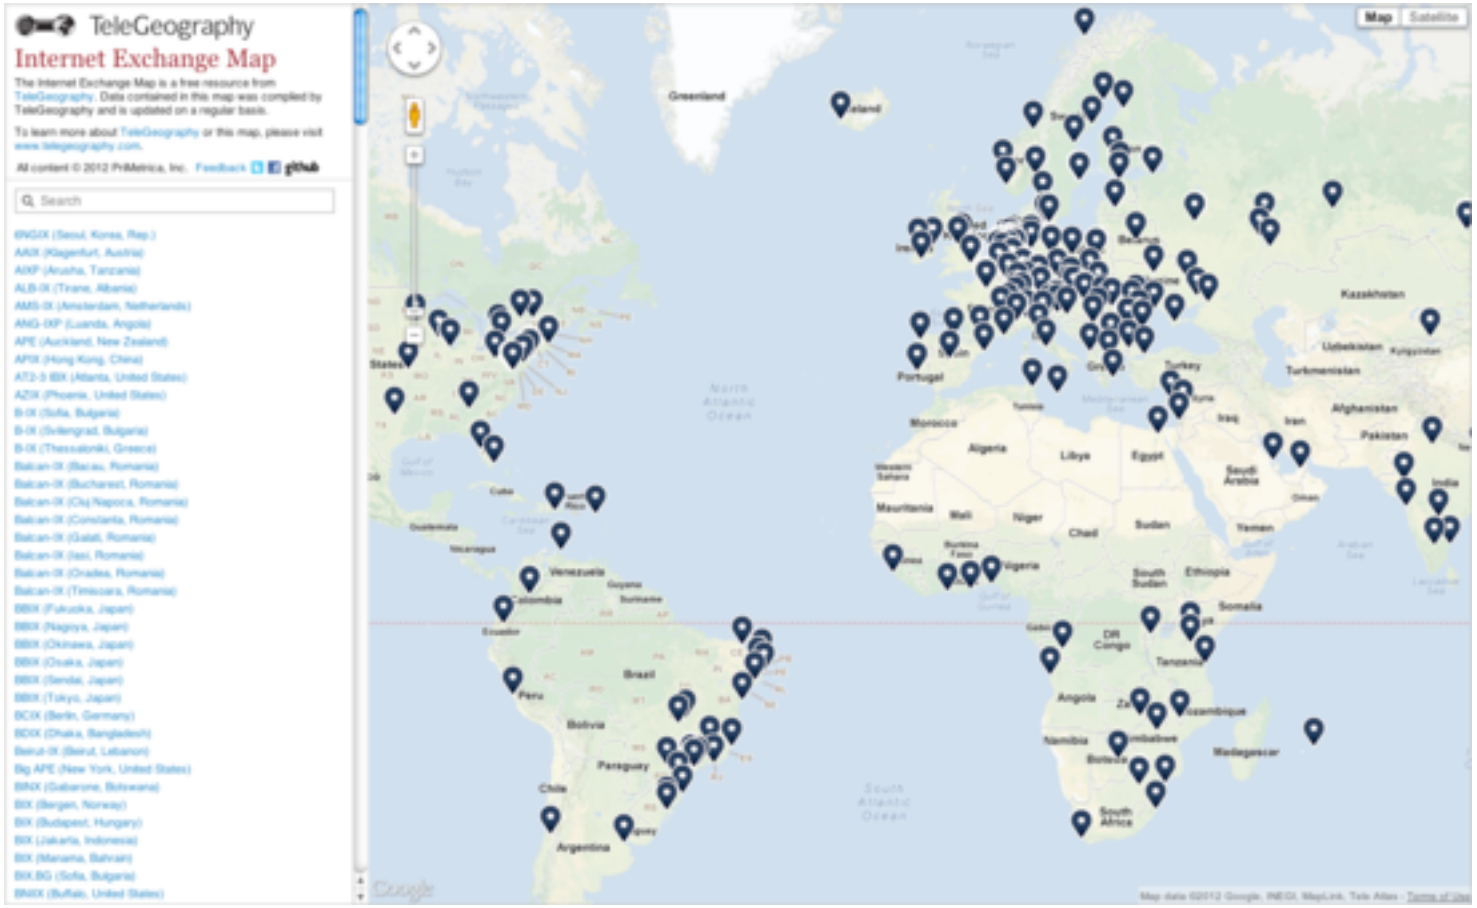

Source: TeleGeography World IX Map, http://www.internetexchangemap.com/

Other sources: www.euro-ix.net | www.ixptoolkit.org | www.pch.net

Africa: Need for Capacity Building

- **Best practices for IXPs**
  - How can we make the IXP grow and become valuable for the local and regional ecosystem?
  - What are the right business models?
- **Technical skills**
  - Routing, network management, and network efficiencies
  - Running an IXP and working with local Internet community and authorities

Photos: © Internet Society/Shoot the Earth/ Nyani Quarmyne

AXIS I & II and AfPIF

- African Union Projects | Implemented by the Internet Society
- AXIS I
  - 30 Best Practice and Community Mobilization & 30 Technical Aspects Workshops (hands-on)
  - 4 IXPs launched with partners (AfriNIC, Jaguar Networks, Lyon-IX, INEX)
- AXIS II
  - 5 Regional meetings to focus on development of Regional IXPs and Regional Internet Carriers
- AfPIF – African Peering & Interconnection Forum
  - Peering, interconnection, IXP meet-ups

LAC – the Need for Capacity Building

- Countries **that deployed IXPs 16 years** ago (Argentina, Brazil, Colombia, Ecuador, Chile) developed stronger Internet technical infrastructures and markets. Related to **market conditions and regulatory/policy environment**.

- Countries that have **more recently deployed IXPs or are in the process of deploying** IXPs - **symptom of market and regulatory/policy conditions** and a less developed Internet community (e.g., Bolivia, El Salvador, Guatemala, Honduras, Paraguay).

- Strong **incumbents**, **lack of strong** Internet technical community and infrastructure
  - **ISPs in** some countries are just in the re-selling business.

- **Pre Best Practices training** with Governments (Reg+Min)**:**
  - Help **invite companies** to initial training sessions. Partnership environment.
  - **Joint training objective** – train the Govt and Internet community
  - Faster progress → countries where the Govt does not try to regulate the entire process
  - Two **different examples using the same approach**:
    - Costa Rica: Did not mandate everything related to the IXP be regulated.
    - Bolivia: Imbedded in law and regulation. Top down. Longer process.

LAC – Capacity Building & Partnerships (cont…)

- **Intro to BGP** and **traffic engineering** using BGP (how to reflect their businesses in the network).
- **Joint training** usually with LACNIC, LACNOG, PCH, Governments, company experts – basics of architecture and how to obtain resources.
- **Equipment**: Work with local experts to identify their needs and help provide equipment (Cisco, Google Foundation):
  - **Start-up**: Difficult at the beginning (think of IXP as additional set of costs). Provide **equipment and training** and the value of the IXP becomes more apparent.
  - Later: **Easier to "level-up"** to charge (maintenance, upgrades, electricity).

LAC Partnerships have developed and….

- **Development of LAC-IX**
  - ISOC and LACNIC helped develop LAC-IX
  - ISOC working with LACNIC and partners to train in the region
- **Community Building**
  - Regional Interconnection **Forum** (within LACNIC Meeting)
  - **LAC Peering Forum** (WG within LACNOG)
- **Partners**
  - LACNIC, LAC-IX, NIC.BR, PCH, LACTLD
  - Governments: CITEL, regulators, ministries
  - Companies/Organizations: Cisco, Google Foundation

IXP Toolkit & Best Practices Project

- The Internet Society was **awarded a grant to extend its Internet exchange point (IXP)** activities in more places**.**
- **The IXP Toolkit Grant builds on the Internet Society's previous efforts and is:**
  - Creating and improving an IXP **Toolkit** | A study and Methodology to Identify Best Practices | http://www.internetsociety.org/ixptoolkitguide
  - Creating and improving an IXP "**Portal**" | www.ixptoolkit.org
  - Partnering to Conduct **Training and Hold Workshops** | Building Capacity around the World
    - Working with: Academics, Euro-IX, IXPs (INEX, Lyon-IX), LACNIC, RIPE-NCC, NSRC (in works)
    - Also working in: Asia-Pacific, Eastern Europe, Commonwealth of Independent States

IXP Toolkit & Portal | Maps

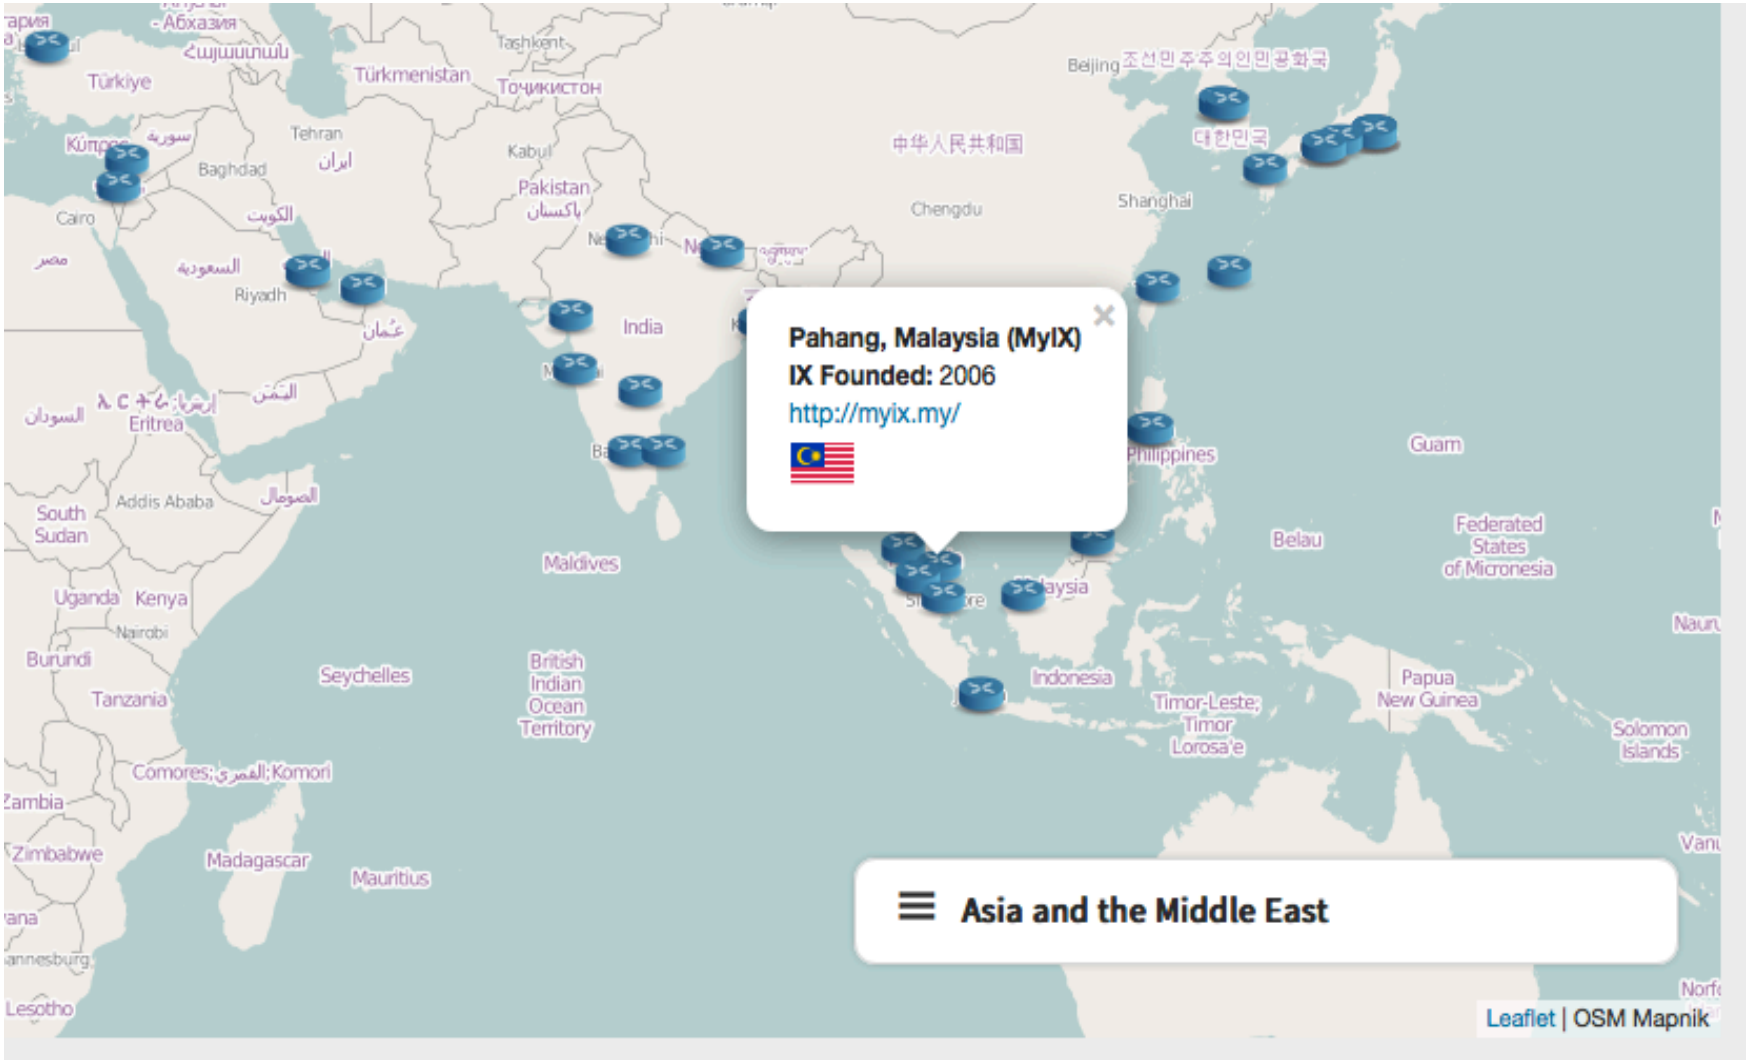

Keeping “Local Traffic Local”

- **Develop local Internet infrastructure & Ecosystem**
  - Human | people
  - Technical | equipment & training
  - Governance | evolving models
- **Snowden implications**
  - Questions from local governments about local traffic
  - IXP is **not set-up to be** a monitoring facility
  - Local content creation, local hosting, local DNS

CAIDA’s Mission Statement

The Center for Applied Internet Data Analysis (CAIDA) is an independent analysis and research group based at the University of California's San Diego Supercomputer Center. CAIDA investigates both practical and theoretical aspects of the Internet.

UCSD

infrastructure

Archipelago

http://www.caida.org/projects/ark

- CAIDA's active measurement infrastructure
- 102 monitors, growing by 1 or 2 per month
  - 37 IPv6 capable
  - 39 countries (88 cities)
  - 54 Raspberry Pis
- current projects
  - team-probing experiment to collect IPv4 and IPv6 traceroutes
  - alias resolution to get router-level topology
  - interdomain congestion measurement
  - spoofing measurement

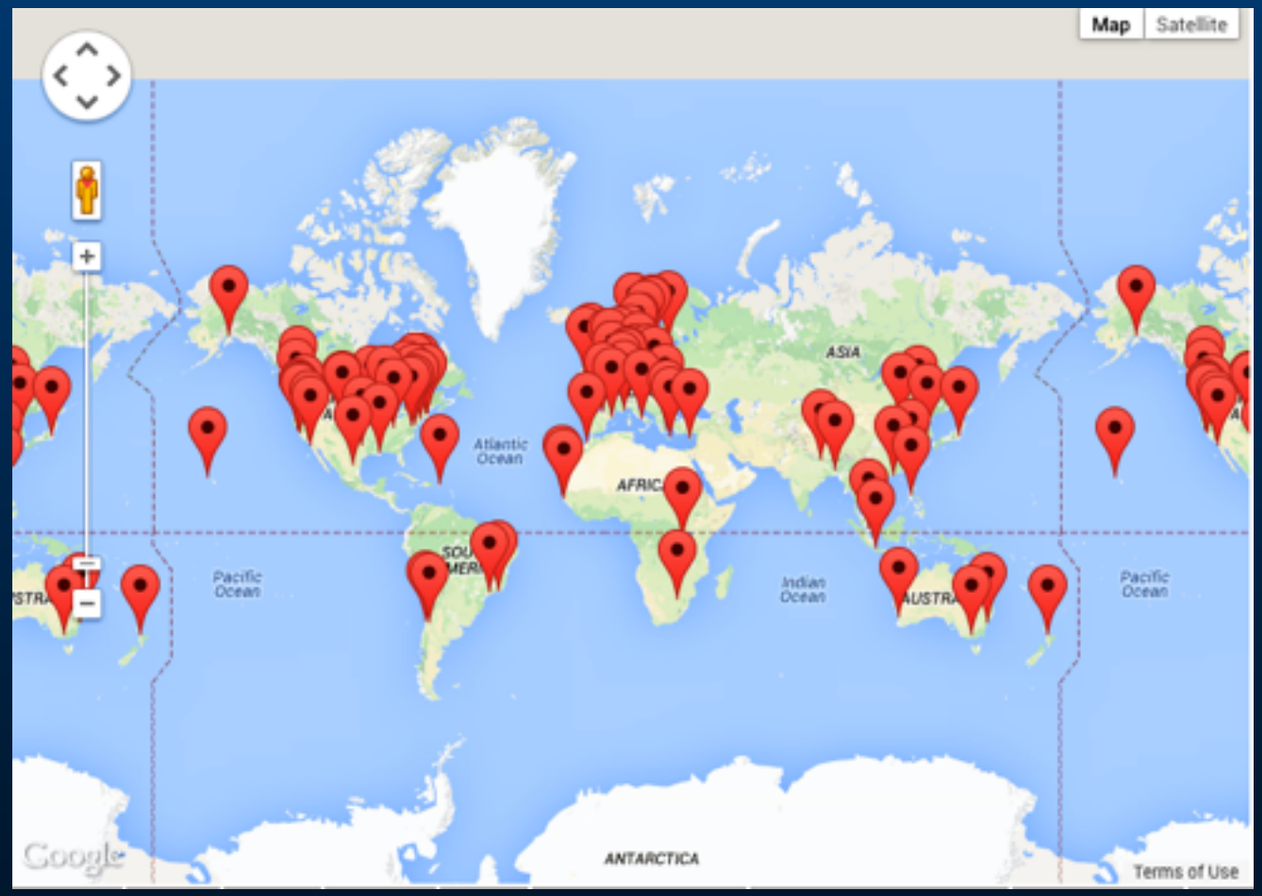

Raspberry Pi

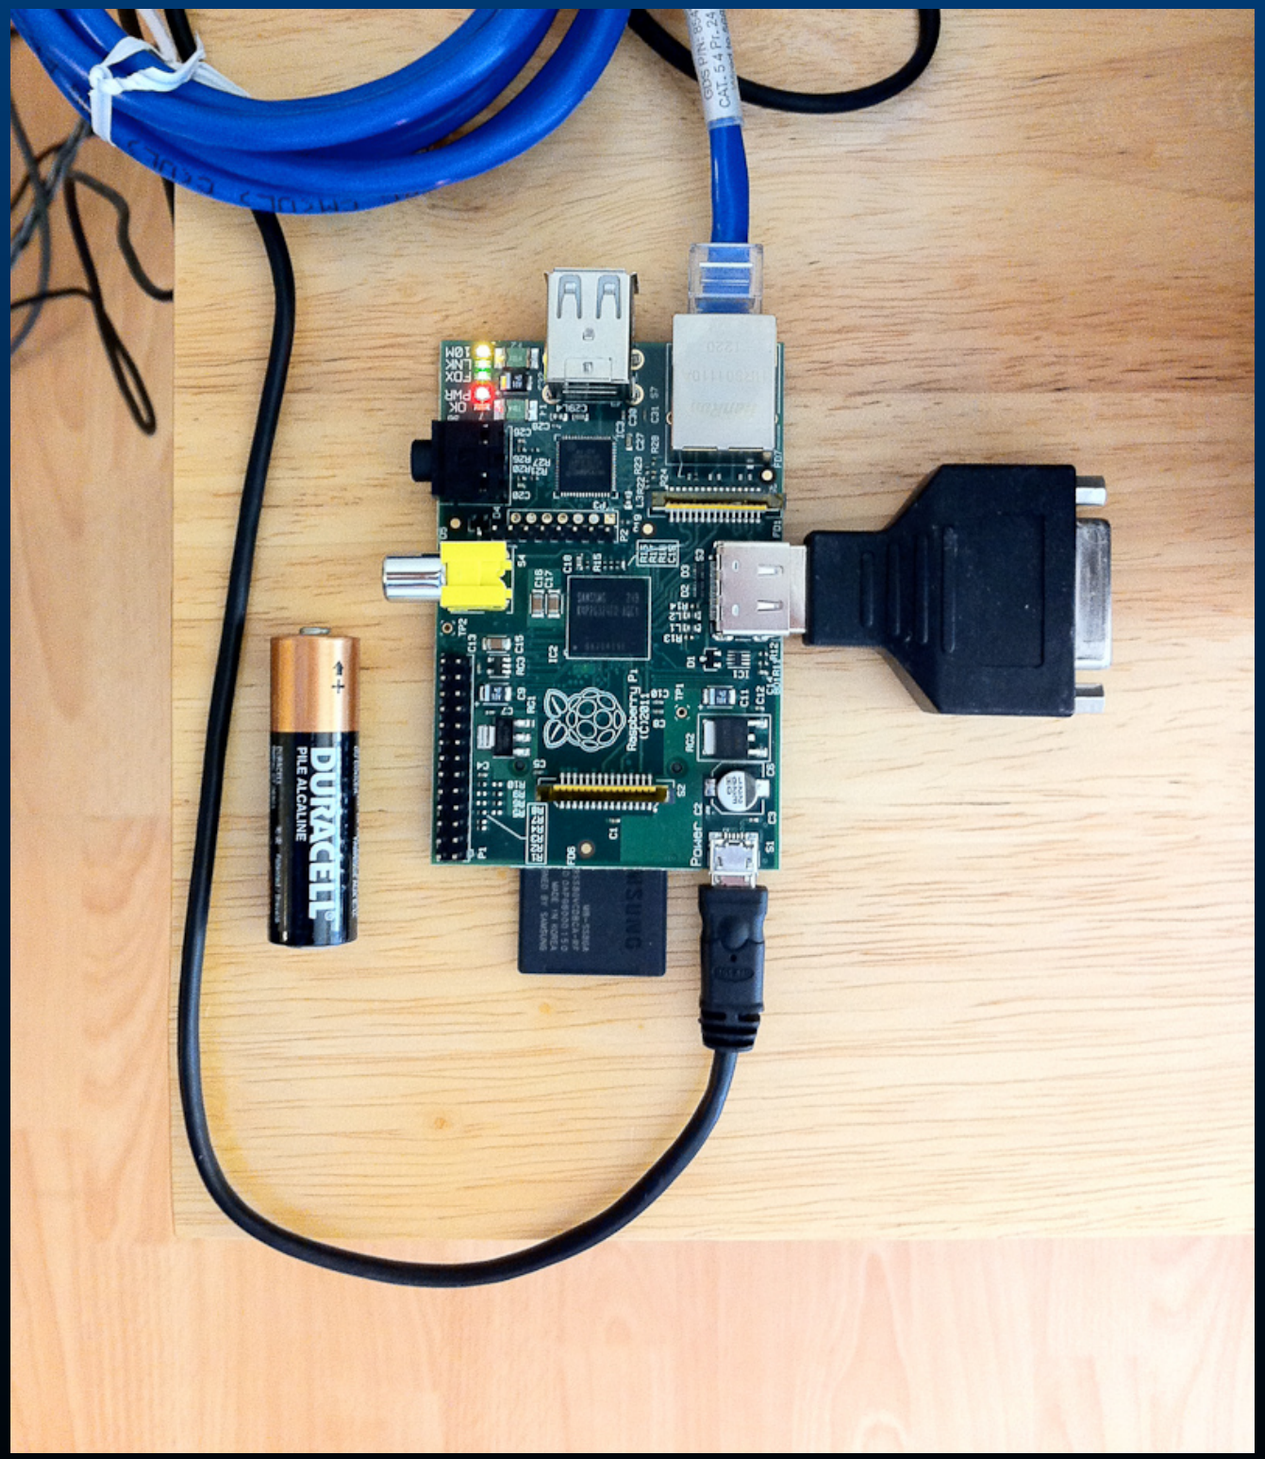

- 700 MHz ARM CPU
- 512 MB RAM
- 100 Mbps Ethernet
- 2 x USB 2.0
- SD card slot
- HDMI display output
- Cost only $35
- Always looking for new vantage points: talk to me later if you can host one!

Archipelago monitors and data

- 5.5 TB (skitter+ark compressed)
  - routed IPv4: 5.4TB since Sep 2007
  - routed IPv6: 8GB since Dec 2008

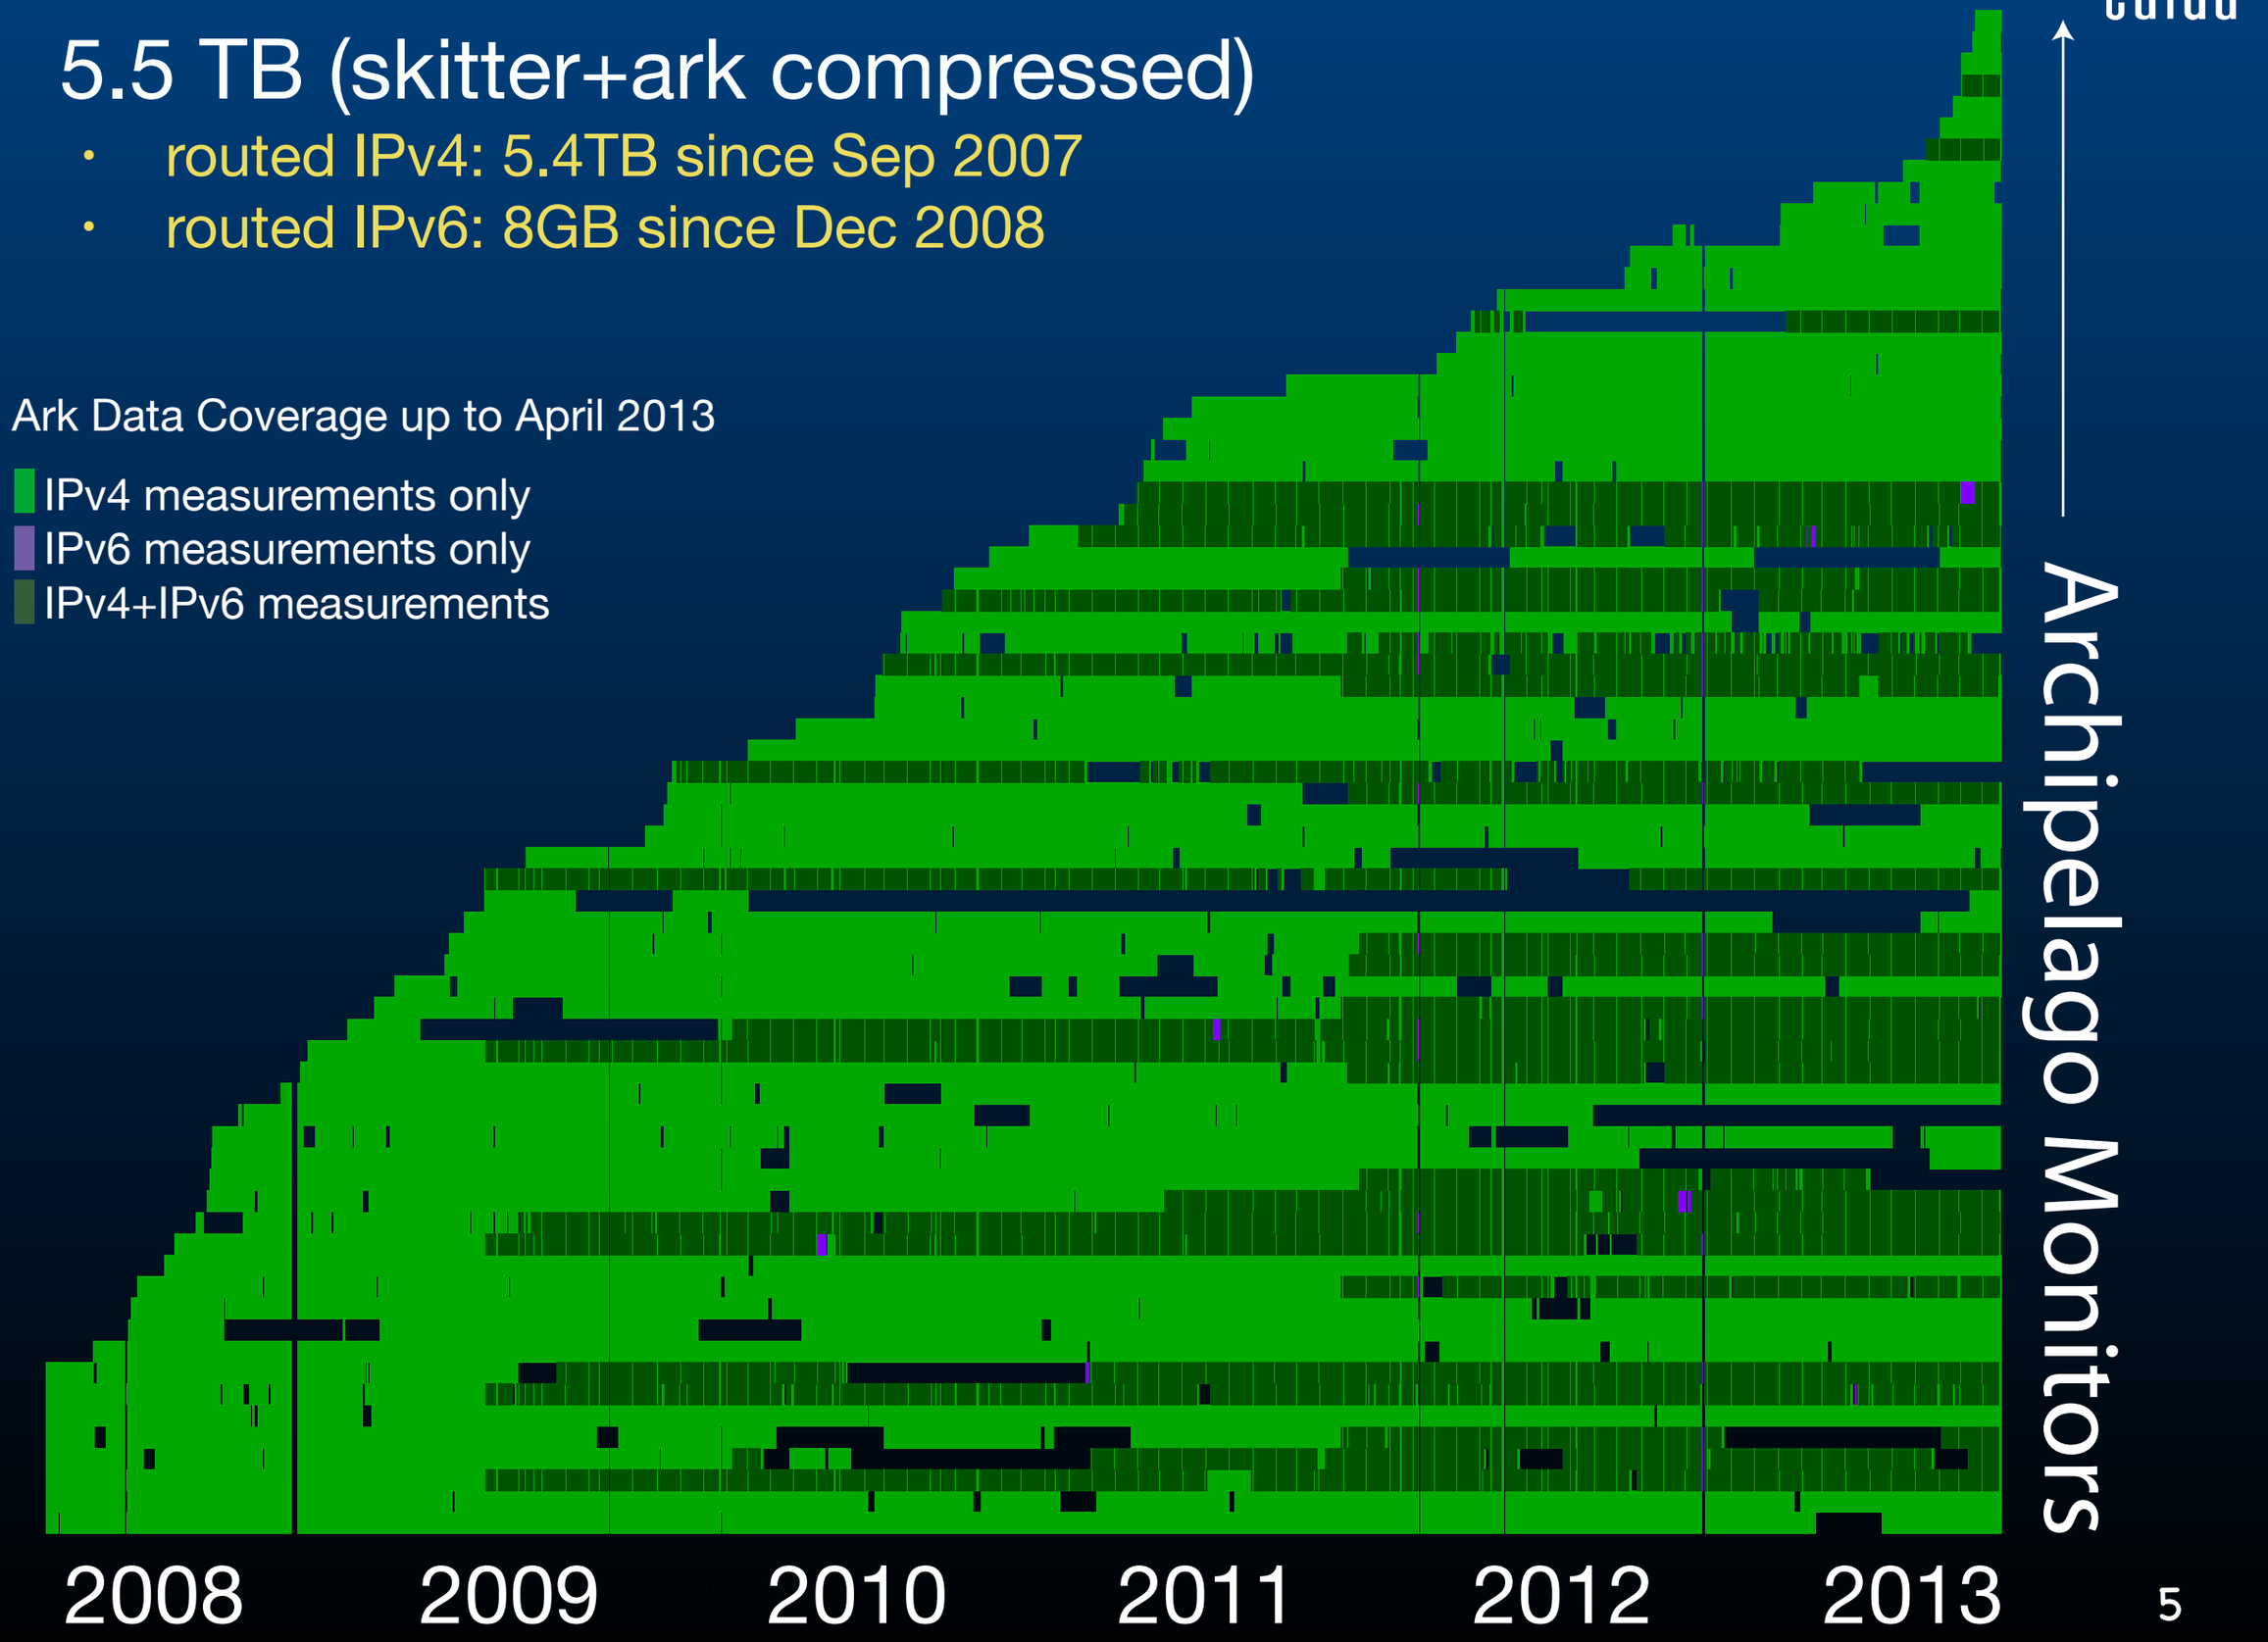

Archipelago data available to researchers

- Raw traceroute data 2007-present (IPv4 and IPv6)
  - 5.5TB of trace data

- Curated topology snapshots: Internet Topology Data Kit (ITDK), two per year
  - Router-level topology
  - Router-AS assignment
  - DNS names
  - Geolocation

- Traceroute-derived IPv4 and IPv6 AS links

Supporting rich queries on Ark data

- Goal: support rich queries on traceroute data + geolocation + annotated AS-level topology + router-level topology
- Example 1: Show all traces from a monitor in Canada to destinations in Canada that traverse at least N hops in the United States
- Example 2: Suppose we predict that a certain region will be affected by a natural disaster or political instability. Find all paths that currently traverse that region.
- Example 3: Show connectivity statistics from all monitors to all probed addresses in a given country
- Which queries would you like to see supported?

Vela: Interactive topology-on-demand

http://www.caida.org/projects/ark/vela

- Vela: interactive interface to on-demand measurements from Ark monitors, currently ping and traceroute

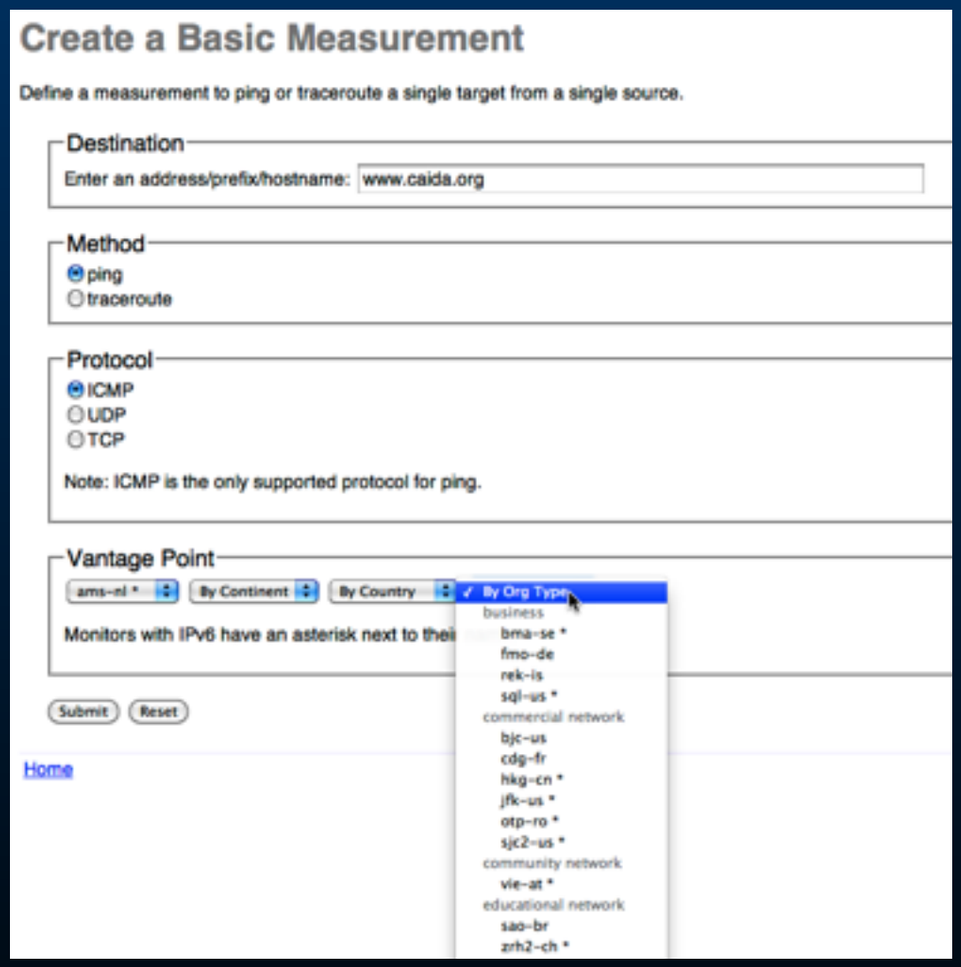

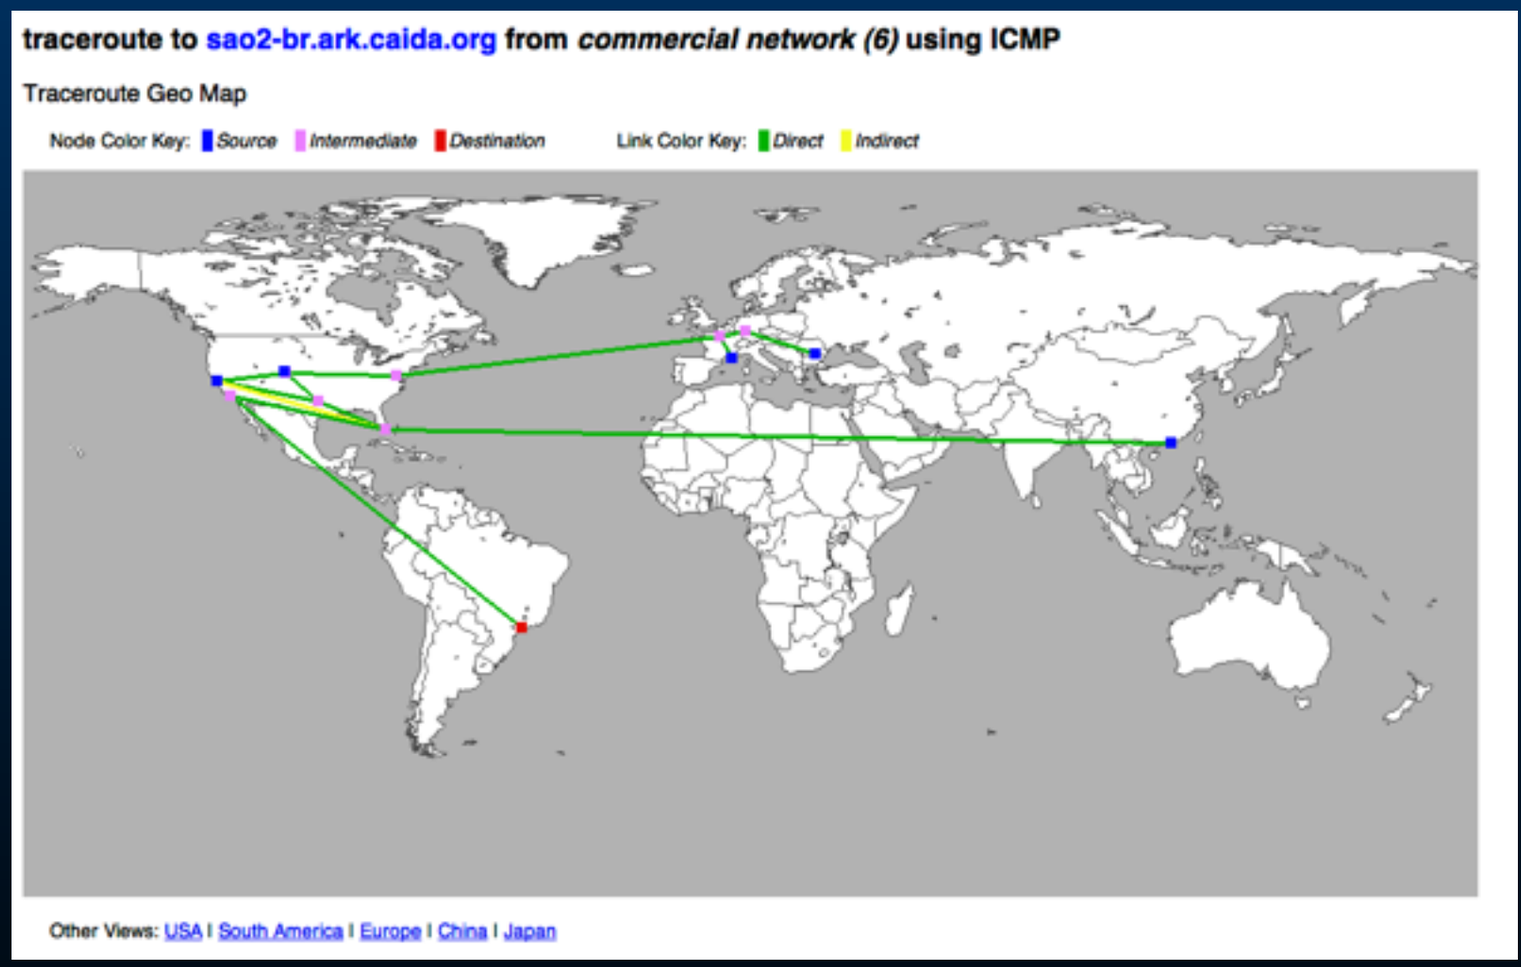

DRoP+DDec: DNS-based geolocation

ddec.caida.org

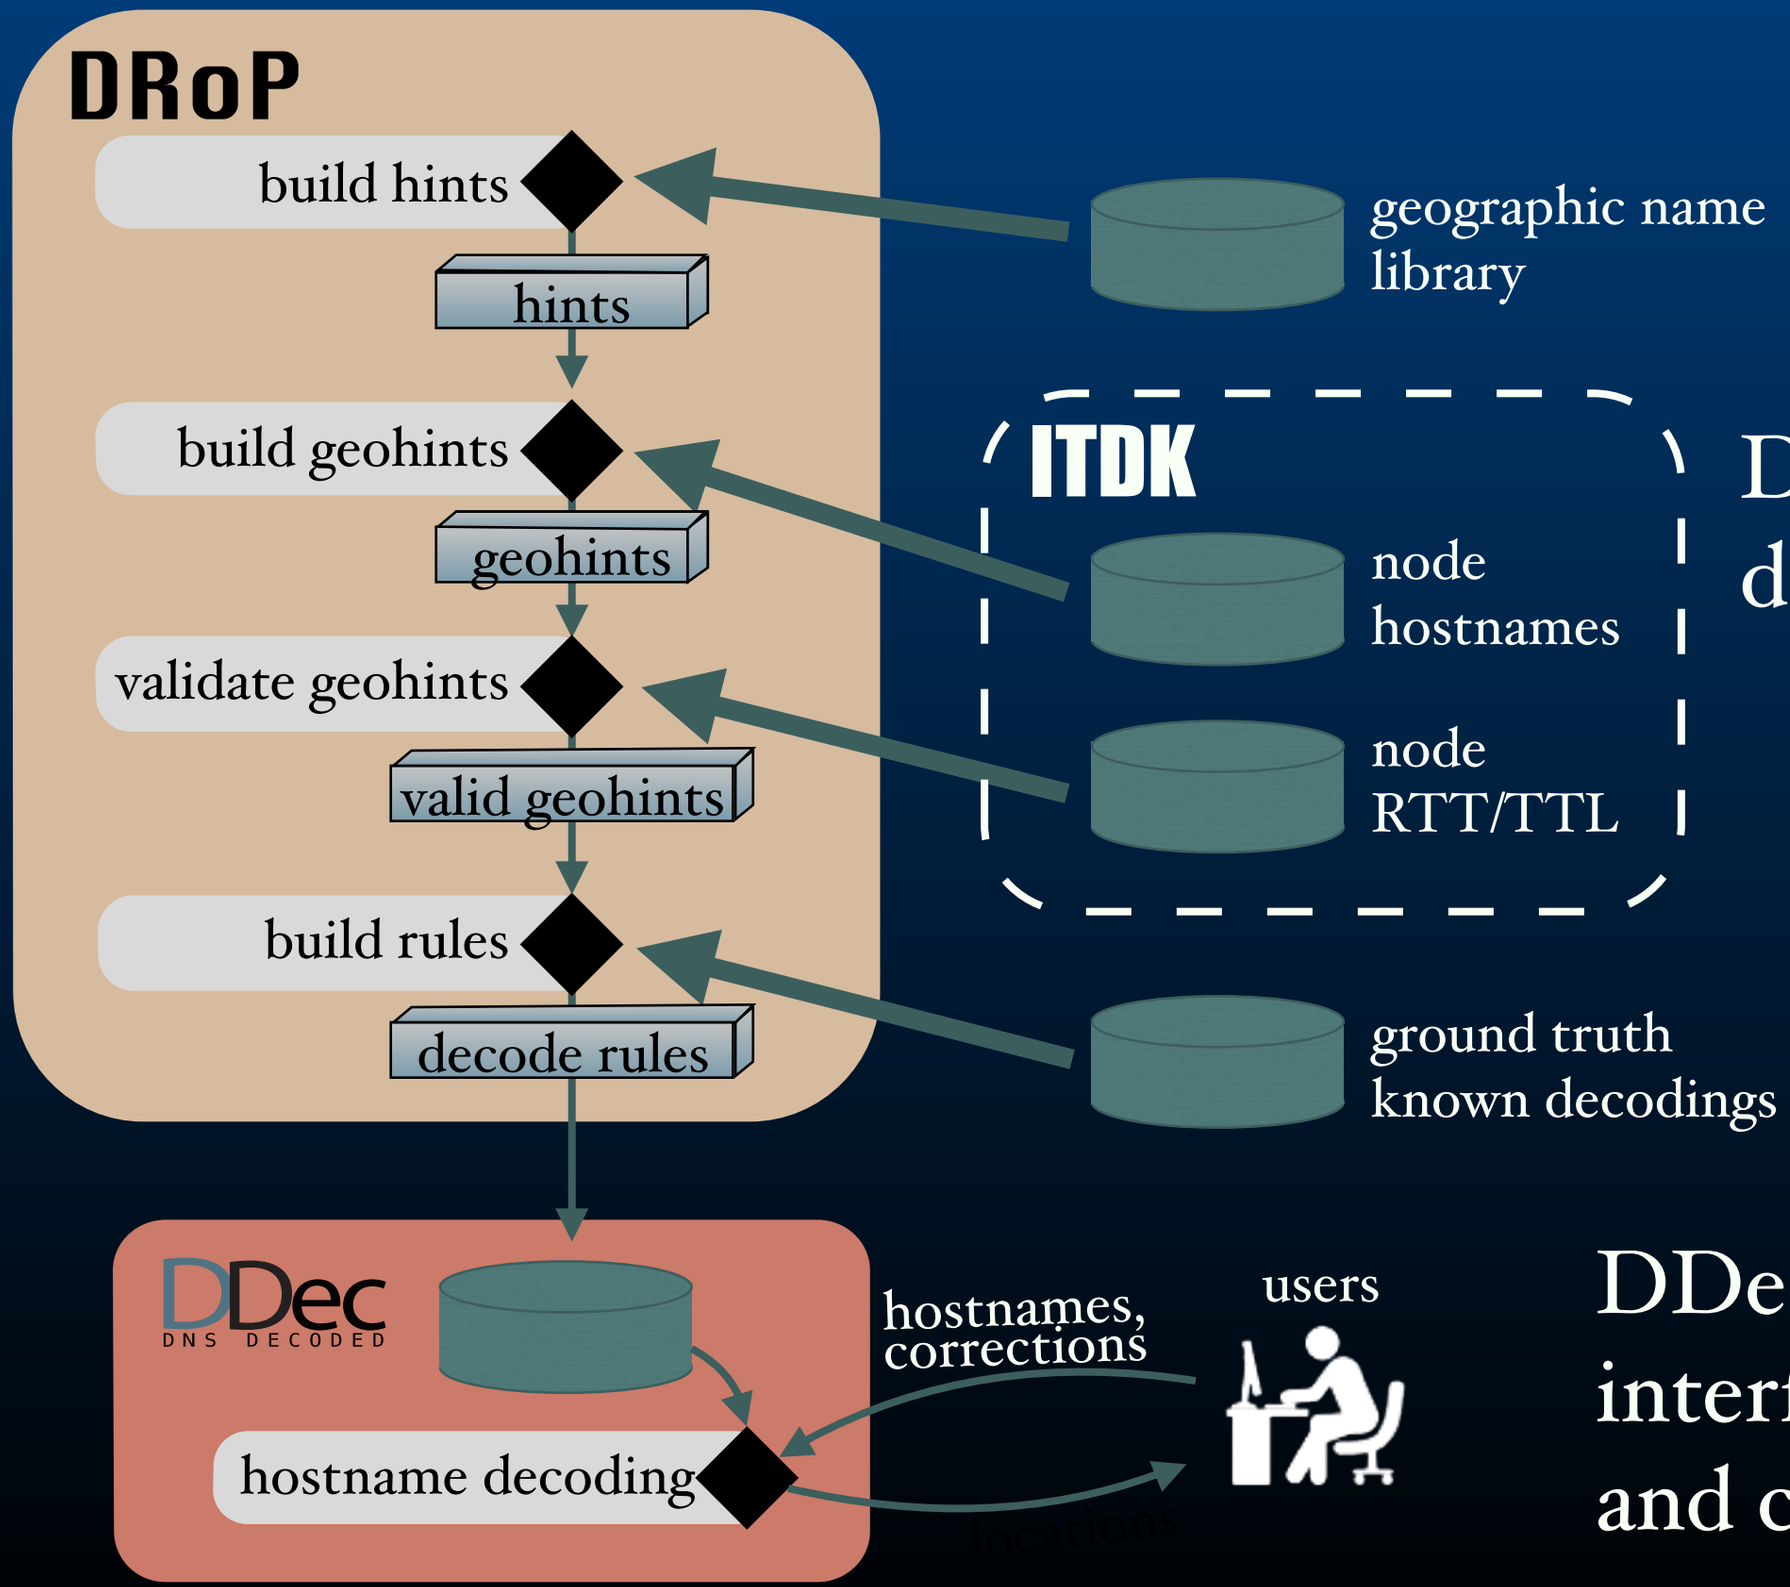

DRoP: automated dns hint detection

DDec: public interface for lookups and corrections

tools

AS-rank

as-rank.caida.org

AS business relationships, customer cones, ranking

AS Ranking | Org Ranking | Information for a single AS | Information for a single Org | Background | Data Sources | Help — AS Ranking Help

The top ASes ranked by customer cone size are displayed below.
For information about a specific AS, enter its AS name, its AS number, or the name of the Org of which the AS is a member.

Dataset: 2013-04-01 — Change dataset

Look up an AS by number or name — Search

Table shows 10 of 44086 ASes, sorted by number of ASes in customer cone — update view

AS rank	AS number	AS name	Org name	customer cone: Number of ASes	Number of IPv4 Prefixes	Number of IPv4 Addresses	Percentages of all ASes	Percentages of all IPv4 Prefixes	Percentages of all IPv4 Addresses	AS transit degree
1	3356	LEVEL3	Level 3 Communications	22,685	261,219	1,401,759,501	51%	57%	65%	3621
2	3549	LVLT-3549	Level 3 Communications	15,103	200,586	698,222,855	34%	44%	32%	3264
3	3257	TINET-BACK...	Tinet SpA	14,873	188,737	709,433,321	33%	41%	33%	942
4	174	COGENT-174	Cogent/PSI	13,594	147,701	589,730,708	30%	32%	27%	3855
5	1299	TELIANET	TeliaNet Global Network	12,722	160,514	616,234,216	28%	35%	28%	764
6	2914	NTT-COMMUN...	NTT America, Inc.	11,159	169,846	711,971,065	25%	37%	33%	888
7	6453	AS6453	TATA Communications	7,062	120,037	459,993,873	16%	26%	21%	580
8	701	UUNET	MCI Communications Services, Inc. d/b/a Verizon Business	5,402	96,864	738,082,126	12%	21%	34%	1693
9	6762	SEABONE-NET	TELECOM ITALIA SPARKLE S.p.A.	4,808	61,319	190,002,775	10%	13%	8.8%	284
10	2828	XO-AS15	XO Communications	4,118	80,165	353,394,094	9.3%	17%	16%	1047

data sources

geo	database	2013.03.02	netacquity
organization	whois	0000.00.00	JPNIC, KRNIC, LACNIC
		2012.06.29	AFRINIC, APNIC, ARIN, LACNIC, RIPE
topology	BGP	2013.04.01, 2013.04.02, 2013.04.03, 2013.04.04, 2013.04.05	ripe rrc00, rrc03, rrc04, rrc05, rrc06, rrc07, rrc10, rrc12, rrc13, rrc14, rrc15
			routeviews eqix, isc, jinx, kixp, linx, routeviews2, saopaulo, sydney, telxatl, wide
	ITDK	2012.07.23	MIDAR IFF

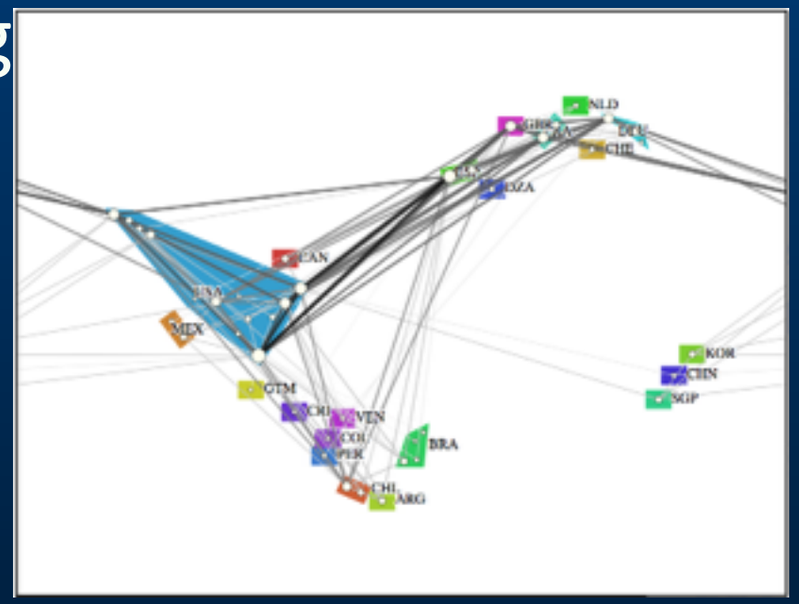

PoP-level map

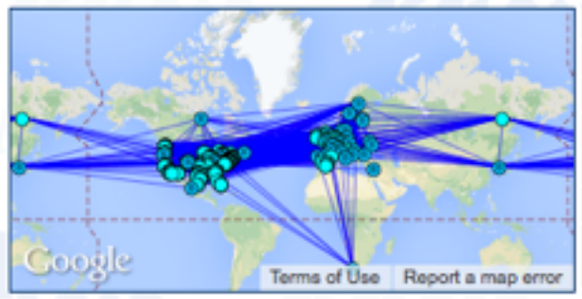

AS number:	174
AS name:	COGENT-174
Org name:	Cogent/PSI
AS rank:	4
Country:	US
Customer cone size:	13,594
AS transit degree:	3,855 (0 Provider, 1 Sibling, 65 Peer, 3,789 Customer)

Router-level map

Operator feedback

neighbor: AS rank	AS	AS name	Org name	inferred relationship type	actual relationship type
5	1299	TELIANET	TeliaNet Global Network	↑ provider	
46	11164	INTERNET2-TRANSITRAIL-CPS	National LambdaRail, LLC	↑ provider	
9	6762	SEABONE-NET	TELECOM ITALIA SPARKLE S.p.A.	↔ peer	
13	6939	HURRICANE	Hurricane Electric, Inc.	↔ peer	
15	3491	BTN-ASN	Beyond The Network America, Inc.	↔ peer	

actual relationship type: (correct) / ↓ customer / ↑ provider / ↔ peer / ∞ sibling / (remove entry)

Recent relevant research

www.caida.org/publications

- Inferring which networks peer at which IXPs using route servers

  “Inferring Multilateral Peering”, Giotsas, Zhou, Luckie, Claffy, ACM CoNEXT 2013

- Mining historical peeringDB data for colocation at IXPs, peering policies, geographical expansion

  “Using PeeringDB to understand the Peering ecosystem”, Lodhi, Larson, Dhamdhere, Dovrolis, Claffy, ACM SIGCOMM CCR 2014

- Investigating connectivity in the LACNIC region

  “LACNIC Connectivity”, Lutu, Bagnulo, Dainotti, Dhamdhere, Claffy, In progress

Recent relevant research

www.caida.org/publications

- Analysis of country-level Internet Blackouts (*BGP withdrawals, packet-filtering, satellite-signal jamming, …*)

  "Analysis of country-wide Internet Outages caused by Censorship", Dainotti et al., IMC 2011

Egypt, Jan 2011
Government orders to shut down the Internet

- Natural disasters affecting infrastructure/population (earthquakes/hurricanes)

  "Extracting Benefit from Harm..", Dainotti, Ammann, Aben, Claffy, ACM SIGCOMM CCR 2012

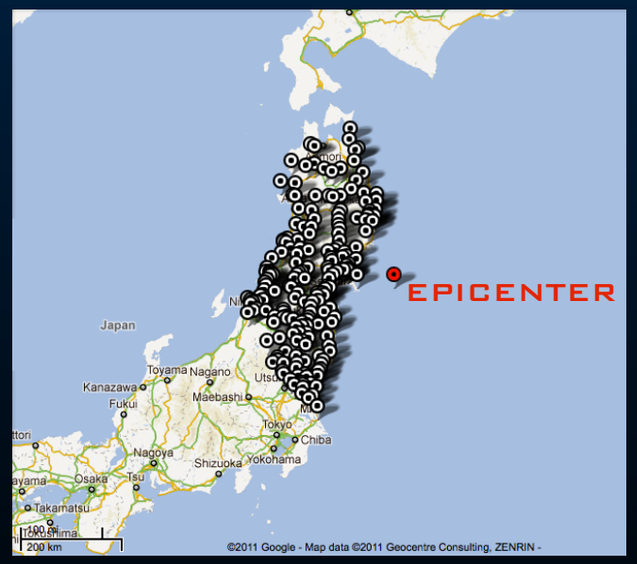

Japan, Mar 2011
Earthquake of Magnitude 9.0

Thanks!

amogh@caida.org

www.caida.org

Discussion

Antonio Gamba-Bari, IXmaps Project

Jane Coffin, ISOC

Amogh Dhamdhere, CAIDA