Interactive Access to Internet Topology Data

Lindisconde, J

Delhi, IN Bombay, IN

Bradley Huffaker and Young Hyun CAIDA SDSC/UCSD

> Gateway 2016 3 Nov. 2016

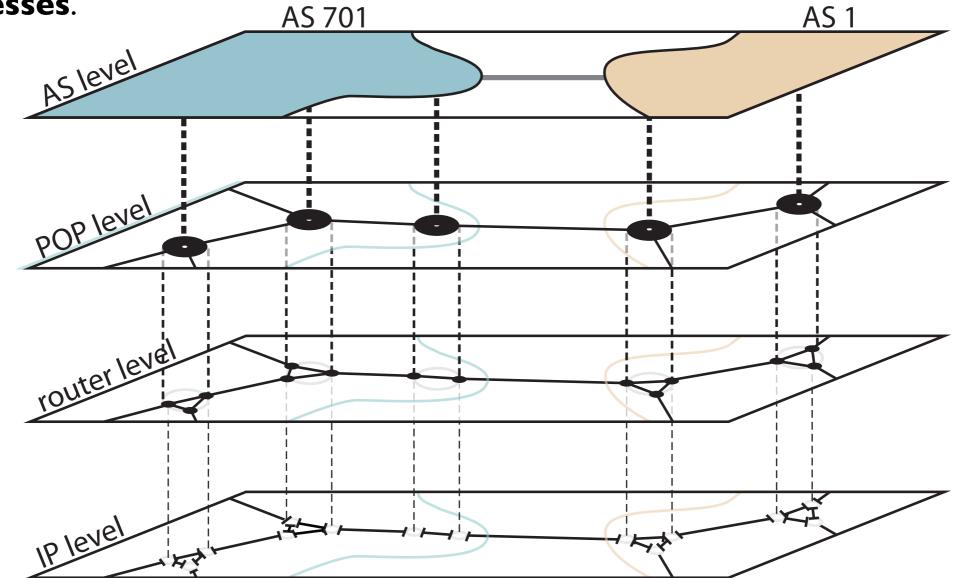
> > Honolulu, US

Layered view of the Internet

background

Internet maps can be grouped into three levels.

- 1. **IP** addresses that connect devices on to the Internet.
- 2. Routers, machines that route the traffic, interconnect via IP addresses.
- 3. Point of Presence (PoP) geographic location with
- 4. Autonomous Systems (AS) are numbers used to route groups of IP addresses.

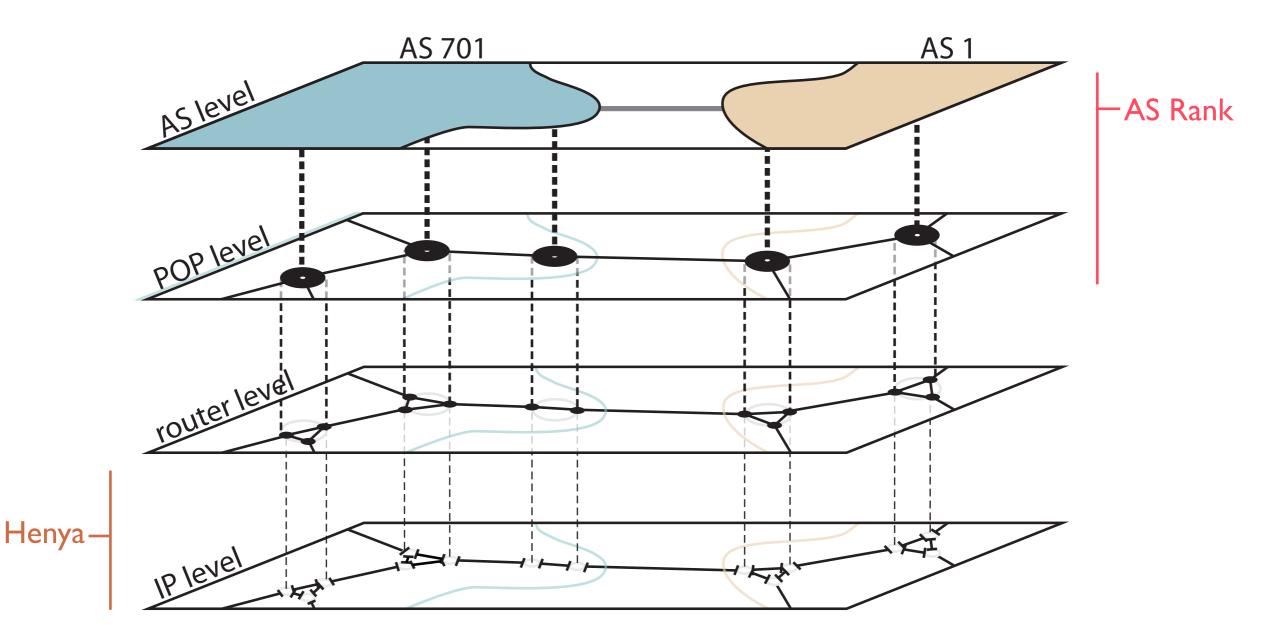


caida



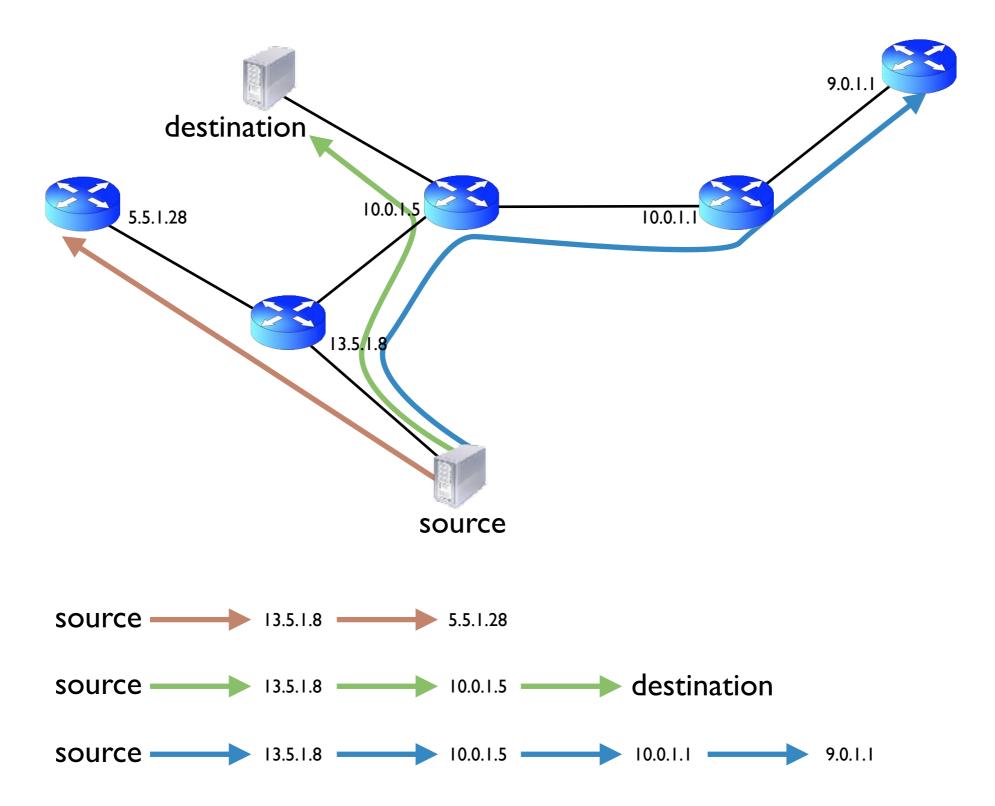
background

- Henya primarily concerned with querying IP paths
- **AS Rank** primarily concerned with querying **AS topology**



traceroute paths

background



4

calda



background

- 9+ years of CAIDA traceroute data
 - 47 billion traces in 20.3 TB of files
 - now growing by 20 billions traces/year
- useful for studying global Internet connectivity, evolution, performance, censorship, ...
- basis for higher-level Internet maps





- CAIDA's large-scale topology query system
- provides remote search of traceroute data without requiring data downloads
- built-in **analyses and visualizations**
 - for commonly occurring needs
- **responsive** enough for interactive data exploration
 - goal: query latency of 30 seconds or less

topology queries

• find occurrences of traceroute path elements

Henya

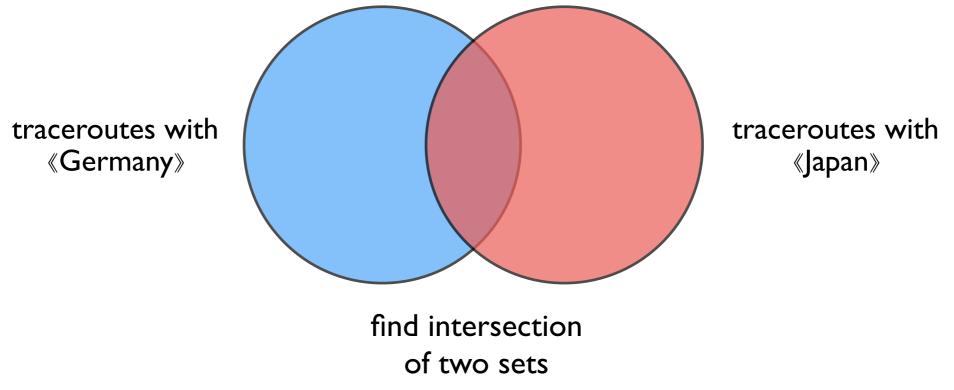
- «targets» = IP addreses (and other types)
- queries:
 - traceroutes toward «targets»
 - traceroutes containing one or more «targets»
- parameters:
 - measurement vantage points
 - data collection time periods
 - position of «targets» in path
 - hop distance between sets of «targets»

and a complexity

- the most complex case:
 - traceroutes containing **two or more** «targets»
 - precisely: traceroutes containing some hop $h_1 \in \langle targets_1 \rangle$, $h_2 \in \langle targets_2 \rangle$, …

Henya

- example: traceroutes containing hops in both «Germany» and «Japan»



- harder:
 - traceroutes with hops in «Germany or UK or France» and hops in «ATT or Level3 network» and hops in «Amsterdam Internet Exchange»



Henya

- large target sets
 - «Germany» = 92,239,360 target IP addresses
 - «Japan» = 154,025,984 target IP addresses
- multiple «targets» in a single query
 - need the **intersection** of subqueries for $\langle targets_1 \rangle$ and $\langle targets_2 \rangle$ and ...

- these challenges poorly met by existing database systems
 - relational databases not designed/optimized for multi-key searches
 - can't always use column indexes; may need to do table scans on separate columns
 - not a good fit for existing NoSQL databases
 - schema-less document stores (JSON/XML) come the closest





• implemented custom index data structures

- highly tailored and tuned to the characteristics of our data and workload
 - efficiently supports large numbers of targets and subquery intersections
- gave up generality and flexibility for speed
- built on **RocksDB** key-value store
 - persistent hash table
 - maps binary string (key) to binary string (value)
 - can also traverse keys in sorted order
 - stores both traceroute data and custom indexes

• custom query engine

- written in Python
- running on 64 cores; may use HPC facilities in future

web interface

- user-friendly GUI to query system
 - also built-in analyses and visualizations for commonly occurring needs

Henya

- lower barrier to use; reach casual users
- uses Bokeh for client-side visualizations
 - supports user interaction and offline viewing on client-side
 - data + visualization (JavaScript) loaded entirely in client browser
 - use Python to implement new visualizations on server-side
 - use library of Bokeh visualization primitives

caida

ad-hoc queries

Query Traces for IP Paths

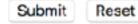
Displays traceroute paths.

Query
Target Address/Prefix/AS/Country: 192.168.0.0/24
Second Target for neigh Query:
Separate multiple targets with commas. Example: 1.2.3.4, 10.0.0.0/8, as1234, .sy
Start Date: 2016-01 End Date: 2016-02
Dates can be YYYY, YYYY-MM, or YYYY-MM-DD. End date is exclusive. Leave start/end (or both) blank for an open-ended range.
Query Method: Odest Oddr Oneigh
 dest — search by trace destination address addr — search for responding address (hop or responding destination address) neigh — search for neighboring addresses (responding hop or destination)
Target Position/Neighbor Separation: 0 0 Max Traces: 10 0 Reverse Order
positive position — hop distance relative to <i>beginning</i> of trace negative position — hop distance relative to <i>end</i> of trace neighbor separation — hop distance <i>between</i> neighboring targets
Vantage Point

Henya

Monitors with IPv6 have an asterisk next to their name.

By Continent 🔇 By Country 🔇 By Org Type 🔇



ams-nl *



ad-hoc queries

Neighbor query of 206.223.119.0/24 and as6939 from bma-se

Download JSON results

1. Traceroute to 173.218.24.1 on 2016-01-01 00:26:24

Нор	Address	Target Match	Prefix	AS	Location	RTT (ms)	
1	•						
2	*						
3	95. 143.207.173		95.143.192.0/20	49770	hudiksvall swe	5.8	1
4	MX-CORE1.internetport.se 95.143.207.229		95.143.192.0/20	49770	hudiksvall swe	5.4	I.
5	CO-RO2.internetport.se 95.143.207.186		95.143.192.0/20	49770	hudiksvall swe	5.5	
6	gige-g2-1.core1.sto1.he.net 192.121.80.162				stockholm swe	18.8	
7	v991.core1.slc1.he.net 72.52.92.81	72.52.64.0/18 (as6939)	72.52.92.0/24	6939	fremont, ca usa	30.0	
8	100ge5-2.core1.par2.he.net 72.52.92.13	72.52.64.0/18 (as6939)	72.52.92.0/24	6939	fremont, ca usa	40.2	
9	100ge10-1.core1.nyc4.he.net 184.105.81.77	1 84.104.0.0/15 (as6939)	184.104.0.0/15	6939	new york, ny usa	117.4	
10	100ge5-1.core1.chi1.he.net 184.105.223.161	184.104.0.0/15 (as6939)	184.104.0.0/15	6939	chicago, il u sa	132.2	
11	equinix-chi.suddenlink.NET 206.223.119.72	206.223.119.0/24 (A)			chicago, il usa	127.7	
12	173-219-231-169.suddenlink.net 173.219.231.169		173.216.0.0/14	19108	lufkin, tx usa	164.7	

Henya

pre-made analysis



Query Traces for RTT Time Series

Plots an RTT time series for target destinations, an RTT histogram, and a time series of target unreachability.

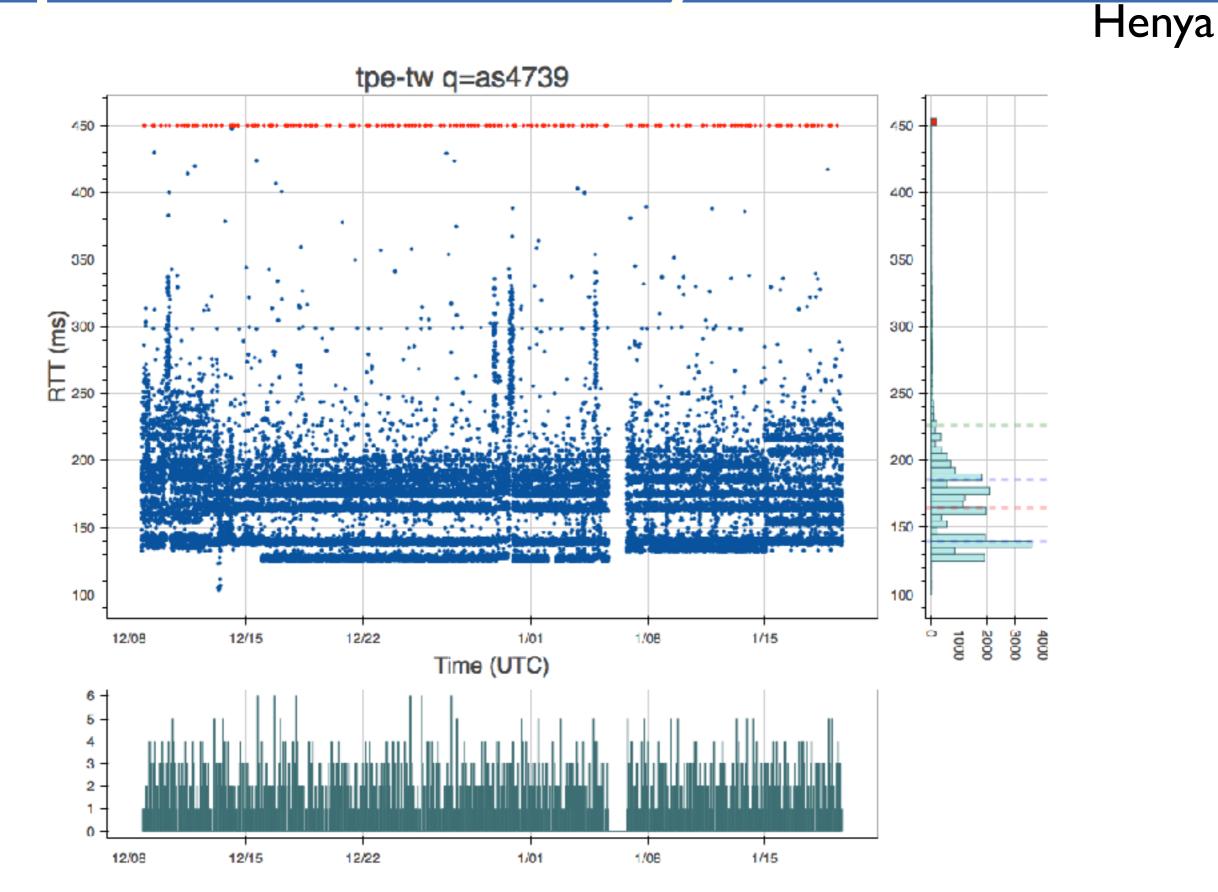
Query		
Target Address/Prefix	<th></th>	
	argets with commas. 10.0.0/8, as1234, .sy	
Start Date:	End Date:	
	Y, YYYY-MM, or YYYY-MM-DD. End date is exclusive. both) blank for an open-ended range.	
Vantago Point		
Vantage Point		
By Name ᅌ By Contir	nent ᅌ 🛛 By Country 😒 🖓 By Org Type ᅌ	

Monitors with IPv6 have an asterisk next to their name.

Submit Reset

caida

pre-made analysis



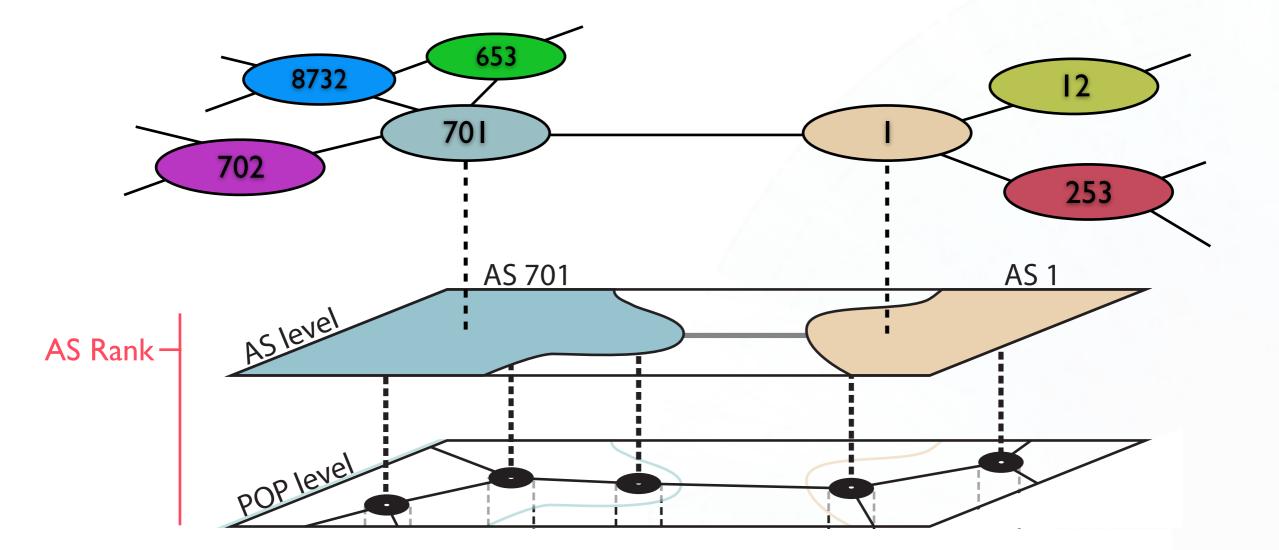
15

calda



as-rank.caida.org

• **AS Rank** primarily concerned with querying **AS topology**





as-rank.caida.org

To provide a way of examining and comparing Internet Service Provider (ISP).

• Most ISP have a single dominate AS.

For many purposes this dominate AS can be used as a proxy for a given organization.

- AS Rank provides two major views:
 - ranked ordering of ASes (global)
 - selected AS and its neighbors (local)

AS Rank highlights

as-rank.caida.org

• presents ranking for a target date

The user selects a single target data from the list of available for the AS topologies.

- annotate with values from secondary datasets
 - heterogeneous dates and values
 - annotate topology with "near in time" datasets

— primary key —	— key selected based on AS topology			
AS Topology	Organization	AS Geography	AS Classifcation	
monthly		quartly (roughly)	yearly (ish)	
selected date	matched org_id	nearest to date		

AS Rank tables

as-rank.caida.org

Dataset Tables

Dataset			
This holds the set of datasets both AS			
topol	topology and ITDK		
dataset_id internal ID			
type	"ITDK" "ASNGEO" "ASN2ORG" "ASNGEO_LINK" "ASNTOP"		
name	string		
date	date		
description	string		
state string			

ASN2ORG Tables

ASN2ORG_Organization org_id_external from AS2org file		
dataset_id internal ID		
org_id	internal ID	
org_id_external external string		
name string		
country	string	
source	string	

ASN2ORG_Asn source is not related to Sources, but from the algorithm			
dataset_id internal ID			
asn	int		
org_id	internal ID		
name	string		
source	string		

ASNTOP Tables

ASNTOP_Information				
dataset_id	internal ID			
address_family	"AF_INET",			
	"AF_IENT6"			
asn2org_dataset_id	internal ID			
number_asnes	int			
number_organizes	int			
number_prefixes	int			
number_addresses	int			

ASNTOP_AsnTopClique			
ASN which are in the top Clique			
dataset_id internal ID			
asn int			
hostname string			

	AsnPath			
	the as path is stored as "I"			
	separated ASes asn_path_id internal ID			
	asn_path	string		

AsnLink			
asn_link_id	internal ID		
asn	int		
sequence_num	int		

ASNTOP_AsnRank				
dataset_id	internal ID			
asn	internal ID			
cone_rank	int			
degree_rank	int			

ASNTOP_AsnDegree					
dataset_id	internal ID				
asn	int				
asn_transit_degree	int				
asn_degree	int				

ASNTOP_AsnCus	stomerCone		
dataset_id	internal ID		
asn	int		
number_asnes	int		
number_prefixes	int		
number_addresses	int		

ASNTOP_AsnLinkPath				
	internal ID			
asn_link_id	internal ID			
asn_path_id	internal ID			

ASNTOP_AsnType				
dataset_id	internal ID			
asn	int			
type	string			

ASNTOP_OrgRankdataset_idinternal IDorg_idinternal IDcone_rankintdegree_rankint

ASNTOP_OrgDegree						
dataset_id	internal ID					
asn_transit_degree	int					
asn_degree	int					
org_transit_degree	int					
org_degree	int					

ASNTOP_OrgCustomerCone					
dataset_id	internal ID				
org_id	internal ID				
number_asnes	int				
number_prefixes	int				
number_addresses	int				

ASNTOP_AsnLinkRelationship						
dataset_id	internal ID					
asn_link_id	internal ID					
type	"PROVIDER", "CUSTOMER", "PEER", "SILBING", "UNKNOWN"					
ASNTOP_AsnLinkAttribute						

Aonti						
additional information IPv6, etc						
dataset_id	internal ID					
asn_link_id	internal ID					
attribute	string					

ASNTYPE Tables

ASNTYPE_AsnType				
dataset_id	internal ID			
asn_link_id	int			
asn	int			
type	string			

ASNGEO Tables

ASNGEO_AsnGeo					
dataset_id	internal ID				
asn	int				
latitude	float				
longitude	float				

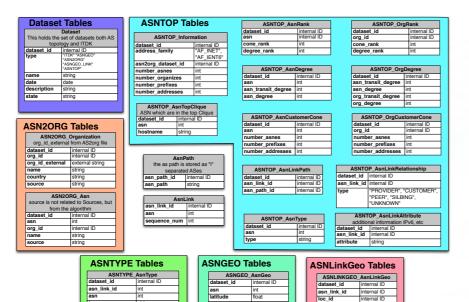
ASNLinkGeo Tables

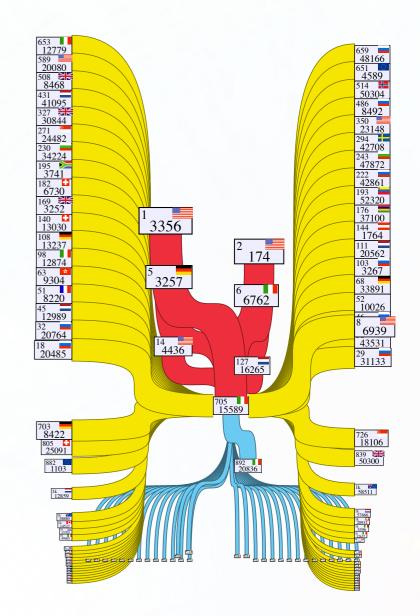
asn_link_id inte	rnal ID
	rnal ID
loc_id inte	rnal ID

caida

Calda AS Rank architecture calda as-rank.caida.org

- CGI (Perl) front end
 - cache based on CGI parameters
- visualizations
 - Perl and C scripts
 - generates SVG files
- MySQL server back end







as-rank.caida.org

organization dataset (04/01/2016)

topology dataset (06/16/2016)

AS rank	AS number	AS name	Org name	AS Type(s)	customer cone						AS transit
					Number of			Percentages of all			
					ASes	IPv4 Prefixes	IPv4 Addresses	ASes	IPv4 Prefixes	IPv4 Addresses	degree
1	3356	LEVEL3	Level 3 Communications, Inc.	(Tr/Ac)	29,494	224,970	783,401,728	53%	<mark>34%</mark>	36%	4138
2	174	COGENT-174	Cogent Communications	Tr/Ac	23,299	172,963	616,423,936	42%	<mark>26</mark> %	28%	4567
3	1299	TELIANET	TeliaSonera AB	Tr/Ac	21,954	191,391	667,346,176	40%	29%	31%	1272
4	2914	NTT-COMMUN	NTT America, Inc.	T//Ac	18,991	174,304	642,432,768	34%	26%	29%	1352
5	3257	GTT-BACKBONE	Tinet Spa	Tr/Ac	18,140	161,377	565,089,024	33%	24%	26%	1282
6	6762	SEABONE-NET	TELECOM ITALIA SPARKLE S.p.A.	Tr/Ac	14,394	123,771	329,530,624	26%	18%	15%	534
7	6453	AS6453	TATA COMMUNICATIONS (AMERICA) INC	Tr/Ac	12,300	135,127	533,133,824	<mark>2</mark> 2%	20%	24%	685
8	6939	HURRICANE	Hurricane Electric, Inc.	Tr/Ac	8,088	79,800	278,942,720	14%	12%	12%	4809
9	2828	XO-AS15	XO Communications	Tr/Ac	6,251	60,271	250,568,448	11%	9.2%	11%	1089
10	1273	CW	Cable and Wireless Worldwide plc	Tr/Ac	5,878	42,258	173,223,936	10%	6.4%	8.1%	296

type dataset (08/02/2015)

local AS view as-rank.caida.org caida geographic AS number: 3356 AS name: LEVEL3 organization Org name: Level 3 Communications, Inc. AS rank: 1 Country: US topology Customer cone size: 29,494 75 4.361 8 AS transit degree: 4,138 Sibling Map data @2016 | Terms of Use | Report a map error Provider Peer Qustomer type Type: Tr/Ac AS Relationship Table **Position in Ranking** AS Relationship (Geo) Graph AS Relationship Graph Corrections The relationship table below displays the neighbors of AS 3356, and each neighbor's inferred relationship type with AS 3356. of 4436 neighbor ASes, sorted by relationship type and AS rank, with simple of details. update view Table shows 100

neighbor					type
AS rank	AS	AS name	AS type(s)	Org name	
2	174	COGENT-174	TriAc	Cogent Communications	⊷ peer
3	1299	TELIANET	TriAc	TeliaSonera AB	↔ peer
4	2914	NTT-COMMUN	TriAc	NTT America, Inc.	↔ peer
5	3257	GTT-BACKBONE	TriAc	Tinet Spa	↔ peer
6	6762	SEABONE-NET	TriAc	TELECOM ITALIA SPARKLE S.p.A.	↔ peer
7	6453	AS6453	TriAc	TATA COMMUNICATIONS (AMERICA) INC	↔ peer
8	6939	HURRICANE	TriAo	Hurricane Electric, Inc.	⇔ peer
9	2828	XO-AS15	TriAc	XO Communications	↔ peer
12	701	UUNET	TriAo	MCI Communications Services, Inc. d/b/a Verizon Business	⇔ peər
13	6461	ABOVENET	TriAc	Abovenet Communications, Inc	↔ peer

22

top.

org.

type

org.

local AS visualizations

as-rank.caida.org

Orginame: CLOUDITALIA TELECOMUNICAZIONI G.P.A. AG mass: 715

Ĥ,

70

Provider Peer Customer

31

25:0

41095

15000

12237 101 12974

57 32821

⁴⁵ 9304

6939 2917

28 31133

#166

18106

•

Sibling

AS number: 10506

Country: If

Type: In/Ap

174

6762

52 152.65

74 au

Customer cone a za: 34

Position in Banking AS Belationship Table AS Belationship (Geo) Graph AS Belationship Graph Corrections

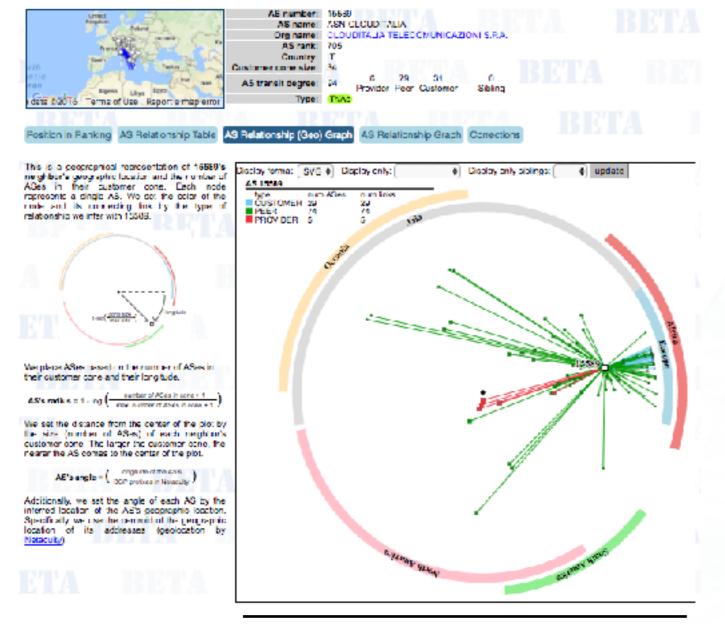
The visualization below is a color-coded, graphic representation of the relationships that AS 15589 has with other ASes.

3356

3257

AS transit degree: 94

AS name: ASN-CLOUDITALIA



metrometric contraction of the second second

Alignet

Mac Data

eountry: IT as cone: 34 degree: 92

renks 722

mak COUNTRY

PEERS

AS NUMBER

SIBLINGS

CUSTOMERS

Line Deer

AS 15589 (CLOUDITALIA TELECOMUNICAZIONI S.P.A.)

Terms of Use ... Report a map en

24487

848

47872

42861

\$2.5

6/20 7/20

3267

20562

12985

10026

43591

20764

topology & organization

topology & geography

caida





• AS rank and Hyena represent only a beginning at bring Internet data to a wider audience.

• We are eager to leverage expertise from other designers of gateways and other environments to support high performance computations.

contact: info@caida.org