



Archipelago

Measurement Infrastructure

Updates

Young Hyun
CAIDA

ISMA 2011 AIMS-3 Workshop
Feb 10, 2011

Introduction

- * Archipelago (Ark) is CAIDA's active measurement infrastructure
 - * in production since Sep 2007
- * focusing on
 - * easy development and rapid prototyping
 - * dynamic and coordinated measurements
 - * measurement services (service-oriented architecture)
- * please see AIMS'09 talk for greater details

Architecture

- * measurement nodes (“monitors”) located worldwide
- * standard rack-mounted servers
- * many thanks to the organizations hosting Ark boxes
- * special thanks for finding hosting sites:
 - Emile Aben (RIPE)
 - Sebastian Castro Avila (.nz Registry Services)
 - Hyunchul Kim (Seoul National University)

Monitor Deployment



* 54 monitors in 29 countries (13 new since AIMS-2)

<u>Continent</u>		<u>Organization</u>		
21	North America	27	academic	} 70% academic or research
3	South America	11	research network	
19	Europe	1	military research	
2	Africa	9	commercial network	} 30% commercial
7	Asia	5	network infrastructure	
2	Oceania	1	community network	

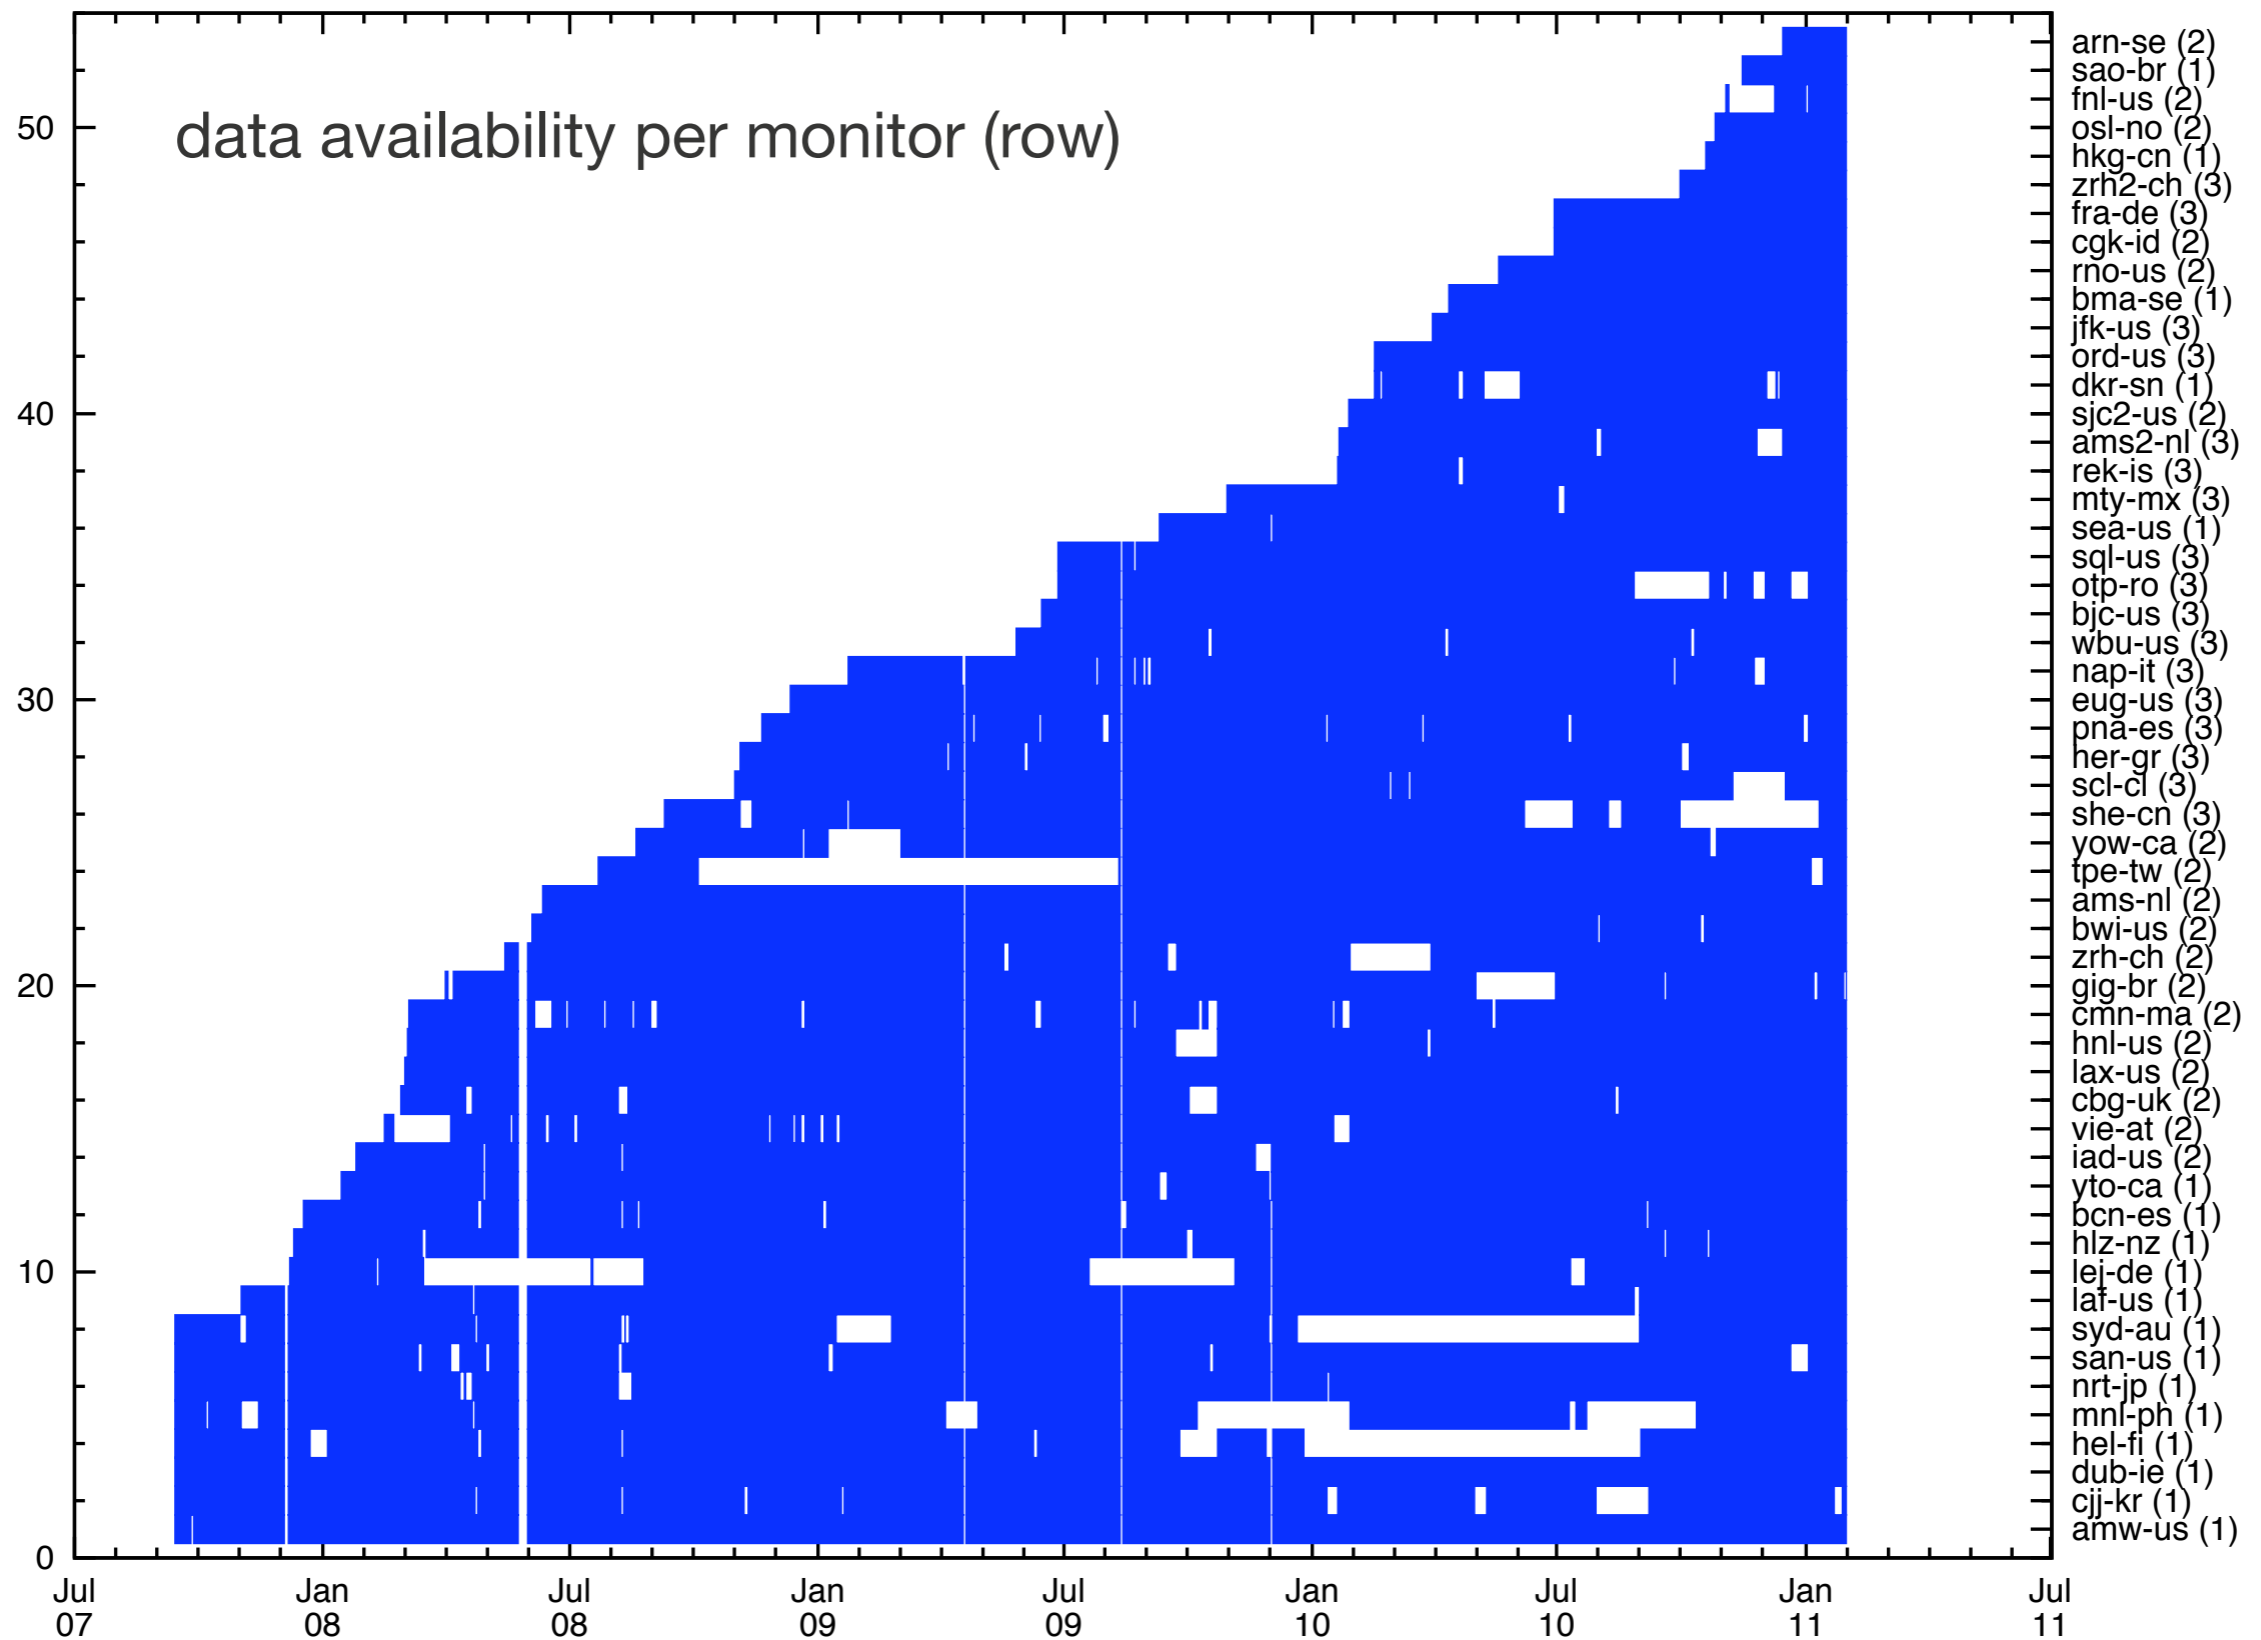
Measurements

- * IPv4 Routed /24 Topology (and AS Links)
- * IPv6 Topology
- * DNS Names & Query/Response Traffic
- * Alias Resolution

IPv4 Routed /24 Topology

- * ongoing large-scale topology measurements
 - * ICMP Paris traceroute to every routed /24 (9.04 million)
 - about 138 /8-equivalents of routed space (as of Dec 2010)
 - * running *scamper*
 - written by Matthew Luckie of WAND, University of Waikato
- * dynamically divide up the measurement work among members of monitor teams
 - * 3 teams active
 - * 17/18-member team probes every /24 in 2 days at 100pps
 - only one monitor probes each /24 per cycle (== one pass through all /24's)

IPv4 Routed /24 Topology



IPv4 Routed /24 Topology

- * collected from Sep 2007 to Jan 2011 (41 months):
 - * 10.1 billion traceroutes; 4.0 TB data
 - * 1312 cycles
- * collecting every month now:
 - * ~432 million traceroutes; ~173 GB data
- * IPv4 topology data is key input into other datasets
 - * e.g., AS links and alias resolution

IPv6 Topology

- * ongoing large-scale IPv6 measurements
 - * 9.0 million traces since Dec 2008
- * 16 monitors
 - * 5 in US, 7 in Europe, 3 Asia, 1 Oceania
- * ICMP Paris traceroute to every routed prefix
 - * each monitor probes a random destination in every routed prefix in every cycle
 - 3,972 prefixes \leq /48 (as of Dec 2010)
 - * probing rate intentionally reduced to 2 days per cycle

Alias Resolution

- * goal: determine which interfaces belong to the same router
- * MIDAR
 - * RadarGun-inspired approach
 - find addresses that share an IP ID counter
 - * paper coming soon (~Mar 2011)
 - * three runs in Jan, Apr, and July 2010:

	addr	aliases (pairs)
Jan	1.1M	425k
Apr	1.5M	1.32M
Jul	1.9M	1.68M

Developments

- * RADclock deployments
 - * collaboration with Julien Ridoux and Darryl Veitch
 - * highly accurate (sub-ms) software-based time synchronization
 - far better than NTP, especially with 200ms RTT to time server
 - * deployed RADclock on 23 monitors + Ark servers

<http://www.synclab.org/radclock/>

Developments

- * experimental on-demand topology measurements
- * remote programmatic access to perform on-demand traceroute/ping measurements from any Ark monitor
 - controlled 3rd-party access to Ark infrastructure without login account

Developments

- * exploring scalable query system for topology data
 - * want easier and faster access to data
 - traceroute and ITDK data
 - * support queries from researchers in the community
 - alternative to downloading raw data (too much to download)
 - * experimenting with Tokyo Cabinet and Dystopia
 - Tokyo Cabinet: non-SQL datastore
 - Tokyo Dystopia: full-text search engine
 - * future: Hadoop, cloud computing, other non-traditional datastores

<http://fallabs.com/tokyodystopia/>

Developments

* topostats

- * programs that calculate 31 graph statistics
- * handles millions of nodes for many statistics
- * example stats:
 - average neighbor degree
 - assortative coefficient
 - mean clustering, clustering coefficient
 - top clique size
 - node coreness, core size, fringe size
 - distance, eccentricity, radius, node/edge betweenness

www.caida.org/tools/utilities/topostats

Collaborations

* Rob Beverly

- * MIT Spoofer Project: added IPv6 support
- * Beverly, et. al, “Primitives for Active Internet Topology Mapping: Toward High-Frequency Characterization,” in IMC 2010
 - conducted on-demand traceroute measurements

* Matthew Luckie

- * using Ark monitors for various topology measurements
- * Luckie, et. al, “Measured Impact of Crooked Traceroute,” in ACM SIGCOMM CCR, Jan 2011


Collaborations

- * Benoit Donnet and Pascal Mérindol
 - * mrinfo measurements
 - * Mérindol, et. al, “MERLIN: MEasure the Router Level of the INternet,” in submission.

Future Work

- * release mper, Marinda, MIDAR, and other Ark software under GPL
- * deploy RADclock on all monitors
- * improve infrastructure to allow more collaborators to use Ark

Thanks!



For more information or to request data:

www.caida.org/projects/ark

For questions, or to offer hosting: ark-info@caida.org