

LANDER Project Update March 2011

John Heidemann and Christos Papadopoulos
joint work with Xue Cai,
Maureen Dougherty (co-PIs), Xun Fan,
Zi Hu, Lin Quan, Yuri Pradkin
USC/ISI, CSU, USC/ITS

21 March 2011

Copyright © 2010 by John Heidemann and Christos Papadopoulos
Release terms: CC-BY-NC 3.0 unported



LANDER Project Status / Mar. 2011



1

Overview

- data collection and infrastructure
 - 10Gb/s collection
 - new datasets
 - IRBs
- data hosting and distribution
 - address visualization
 - dataset distribution
 - metadata wiki
- research
 - outreach
 - collaboration
 - network outages & prefix usage



LANDER Project Status / Mar. 2011

2

Overview

- **data collection and infrastructure**
 - 10Gb/s collection
 - new datasets
 - IRBs
- data hosting and distribution
 - address visualization
 - dataset distribution
 - metadata wiki
- research
 - outreach
 - collaboration
 - network outages & prefix usage



LANDER Project Status / Mar. 2011

3

Data Collection: Review

- review:
 - Los Nettos: 2 commercial peerings (Verio and Level3) and Internet2
 - ServePath: 1 commercial peering in San Jose
 - FRGP: 2 commercial peerings: Level3 and Qwest/Comcast



LANDER Project Status / Mar. 2011

4

Data Collection: New

- deploying 10Gb/s at Los Nettos and FRGP
 - both sites on-line this month
 - plan swap-over at USC this week
 - delays purchasing and debugging new hardware
 - upgrading Los Nettos optical link's transceivers
 - faulty motherboard and wrong raid card at FRGP
- BGP feeds
 - pulling ISI, USC, and CSU with BGPmon
 - expect to archive snapshots every 2 hours, plus full feed
 - not currently planning to export this data



LANDER Project Status / Mar. 2011

5

Collection Infrastructure Plans

- definite plan: 10Gb/s to production use
 - will happen this month
- problem: disk throughput
 - see ~40MB/s headers in
 - need to touch data ~3x (2 anon + 1 user)
 - fileservers tops out around 100MB/s
- tentative plans
 - replace NFS with parallel file system
 - Hadoop is a candidate



LANDER Project Status / Mar. 2011

6

Old Datasets (review)

- Internet address censuses
 - pings of all allocated addresses
 - internet_address_survey_it28w-20090914
- Internet address surveys
 - we pick 1% of the /24 subnets, pinged frequently
 - internet_address_survey_reprobing_it28w-20090914
- general anonymized packet headers
 - anon packet headers from a regional network
 - lander_sample-20080903
- anonymized attack traffic (packet headers)
 - curated packet headers w/known attacks
 - attack-tcpsyn-20061106
- artificial attacks over real background traffic
 - UniformAttack_Traces_Generated_20070821-20041202
- other paper-specific datasets
 - p2p traffic detection, TCP SYNs, etc.

LANDER Project Status / Mar. 2011 7

Dataset Generation (since Dec. 2010): censuses and related

- censuses
 - internet_address_census_it37w-20101124
 - internet_address_census_it38w-20110112
 - and it37c and it38c from Colo. State
- and corresponding surveys (4 datasets)
- and corresponding hitlists (2 datasets)
- total of 10 new census/survey/hitlists datasets

LANDER Project Status / Mar. 2011 8

IRB'ing LANDER

- IRB'ing all LANDER activities (in progress)
- since Dec. 2010
 - AS/org mapping: approved
 - anycast discovery: proposed as and declared NHSR
 - (Not Human Subjects Research)
 - census taking: IRB review, but declared NHSR
 - next, planned as NHSR: ping survey taking, then routing outages
- all students and staff at USC and staff at CSU now IRB-trained

LANDER Project Status / Mar. 2011 9

Overview

- data collection and infrastructure
 - 10Gb/s collection
 - new datasets
 - IRBs
- **data hosting and distribution**
 - **address visualization**
 - **dataset distribution**
 - **metadata wiki**
- research
 - outreach
 - collaboration
 - network outages & prefix usage

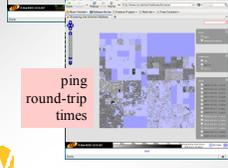
LANDER Project Status / Mar. 2011 10

Address Visualization

this work primarily by AMITE Project



address responsiveness



ping round-trip times



address allocation (now with data from all 5 RIRs)

experimental visualizations have now been publically released

LANDER Project Status / Mar. 2011 11

Address Visualization: the Video

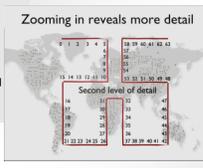
this work primarily by AMITE Project

Our prober sends a ping asking "Are you there?"

Destination responds saying "Yes I am here"

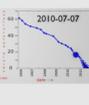
Successful response is assigned **GREEN**

Zooming in reveals more detail



Second level of detail

The graph shows unallocated address blocks over time



<http://www.isi.edu/ant/address/video/>

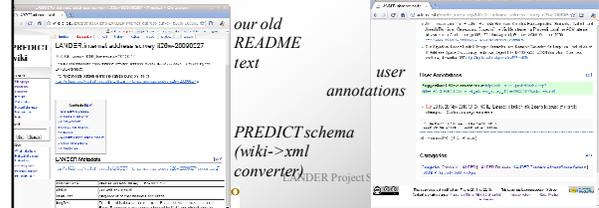
LANDER Project Status / Mar. 2011 12

Dataset Distribution (since Dec. 2010)

- 32 datasets given out to 5 people since July 2010
 - Dec. 2010: 5 dataset to 2 people (PREDICT)
 - Jan. 2011: 10 datasets to 2 people (PREDICT)
 - Feb. 2011: 55 datasets to 3 people (PREDICT)
 - Mar. 2011 (so far): 7 datasets to 1 person (PREDICT)
- observations
 - continued international requests, including one from China
 - (although no actual international distributions this period)

MediaWiki for Metadata (review)

- use standard MediaWiki
 - same as wikipedia
- using it for *all* metadata and annotations
- two extensions:
 - ParseFunctions: fancy templating
 - Lockdown: control editing of namespaces



our old README text

PREDICT schema (wiki->xml converter)

user annotations

Metadata Wiki Status

- wiki continued internal success
 - added support to document dataset categories in wiki
- drew external interaction
 - back and forth about datasets
 - moved from e-mail to wiki
 - assisted in documentation

Overview

- data collection and infrastructure
 - 10Gb/s collection
 - new datasets
 - IRBs
- data hosting and distribution
 - address visualization
 - dataset distribution
 - metadata wiki
- **research**
 - outreach
 - collaboration
 - network outages & prefix usage

Outreach: Video and Papers

- ISI did a press release about the video
 - picked up in The Register and Wired
- papers
 - “Dynamics of Prefix Usage at an Edge Router”, Gadkari, Massey, Papadopoulos, to appear at PAM 2011
 - “Low-Rate, Flow-Level Periodicity Detection”, Bartlett, Heidemann, Papadopoulos, to appear at Global Internet 2011
 - “IP Reachability Differences: Myths and Realities”, Yan, Say, Sheridan, Oko, Papadopoulos, Pei, Massey, to appear at Global Internet 2011.

Collaboration

- with PCH and AMITÉ
 - evaluating using PCH to carry out geolocation
- with DETER
 - joint effort to enable DDoS trace replay in DETER
 - LANDER is collecting traces of DDoS + background traffic
 - DETER will develop trace replay/topology workbench tool
- using BPGMon (Massey, ColoState)

Research: Evaluating Outages

[Lin Quan & Heidemann]

- question:
 - how reachable is the Internet?
 - what do outages look like?
- most prior work used BGP
- can we use our Internet surveys?
 - pings every 11 minutes
 - covering ~1% of /24s for 2 weeks

Outage Detection: Methodology

1. start with a *survey* (green: up, red: error, back: no reply)



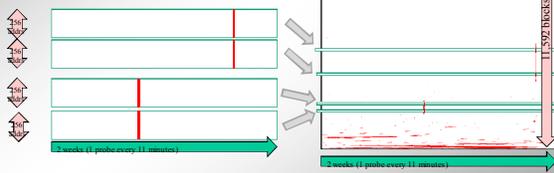
2. evaluate typical responsiveness (% green)
3. look for sudden drops



Outages Over Internet Sample

survey gives ~10k random 24/s

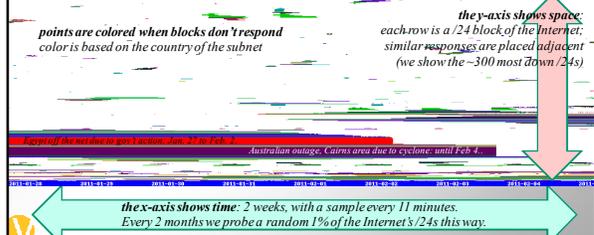
1. find outages in each
2. plot, grouping by outage similarity



Egyptian Internet Outage—Feb. 2011

Network outages: Jan. 28 to Feb. 5 (UTC)

internet_address_survey_reprobing_it38c-20110127



Outage Detection: Next Steps

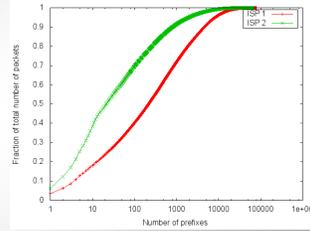
- validating results in routing data
 - installed BGPmon; got local BGP feeds
- metrics to evaluate Internet health
- longitudinal study
 - we have data since early 2007
 - one point: Libya didn't show up (too small)

Research: Dynamics of Prefix Usage at an Edge Router

- Question: current research looks at reduced forwarding tables (FIBs) at routers
 - Routing tables are getting very big
 - FIB memory is expensive (20x DRAM)
 - IPv6 may make things much worse
- But no good metrics to understand prefix dynamics
- Our work: understand prefix dynamics, towards potential FIB caching

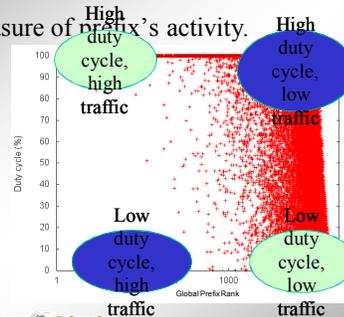
Old Metric: Global Rank

- Rank prefixes according to global rank in 24-hours
- Store most popular prefixes in FIB.
- But how about prefix dynamics?
- Example from our trace: prefix with global rank 4 (online backup service) is only active for less than an hour!



New Metric: Duty Cycle

- Measure of prefix's activity.



Implications of Duty Cycle for FIB Cache Performance

- Prefixes with high duty cycles will always be in the cache.
 - Most of the top globally popular prefixes have a duty cycle between 90 – 100%.
 - 4000 prefixes account for 88% of all packets.
- FIB caches can be relatively small and still achieve high hit rates.

New metric: Mean Rank Difference

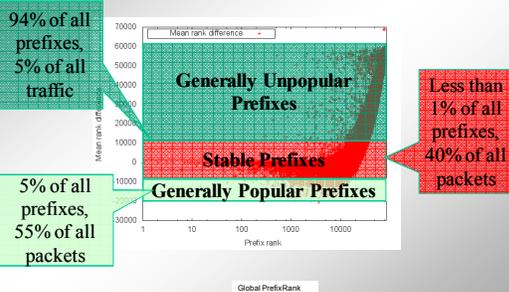
- Defined as :

$$\frac{\sum(\text{Global Rank} - \text{Interval Rank})}{\text{Number of intervals with at least 1 packet}}$$

Number of intervals with at least 1 packet

- Measure of a prefix's "busy-ness".
 - Mean rank difference of close to 0 => prefix maintains its rank.
 - ve mean rank difference => prefix is less popular in some intervals compared to its global popularity
 - +ve mean rank difference => prefix is more popular in some intervals compared to its global popularity
- How much would the cache churn due to prefix dynamics?

Mean Rank Difference Results



Implications of Mean Rank Difference for FIB Cache Performance

- Stable and generally popular prefixes are only 6% of the prefixes, but account for 95% of all traffic.
 - Should not be evicted from the cache.
 - Will miss only 5% of traffic.
- Generally unpopular prefixes will contend for cache slots.

Conclusions

- hosting and providing infrastructure
 - 10Gb/s going live
 - positive results on metadata wiki
 - IRB'ing in progress
- analysis of data
 - new collaborations
 - started study of network outages
 - published study of router prefix dynamics
- <http://www.isi.edu/ant/lander/>

