Real-time Blackhole Analysis with *Hubble*

Ethan Katz-Bassett, Harsha V. Madhyastha, John P. John, Arvind Krishnamurthy, Thomas Anderson

University of Washington

NANOG 40, June 2007

Global Reachability

- When an address is reachable from every other address
- Most basic goal of Internet, especially BGP
  - "There is only one failure, and it is complete partition" Clarke, *Design Philosophy of the DARPA Internet Protocols*
- Physical path ⇨ BGP path ⇨ traffic reaches
- Recurring *NANOG* and *Outages* postings:
  - "Can you reach me?"
  - "I'm seeing problems, what is going on?"

How often is global reachability violated?

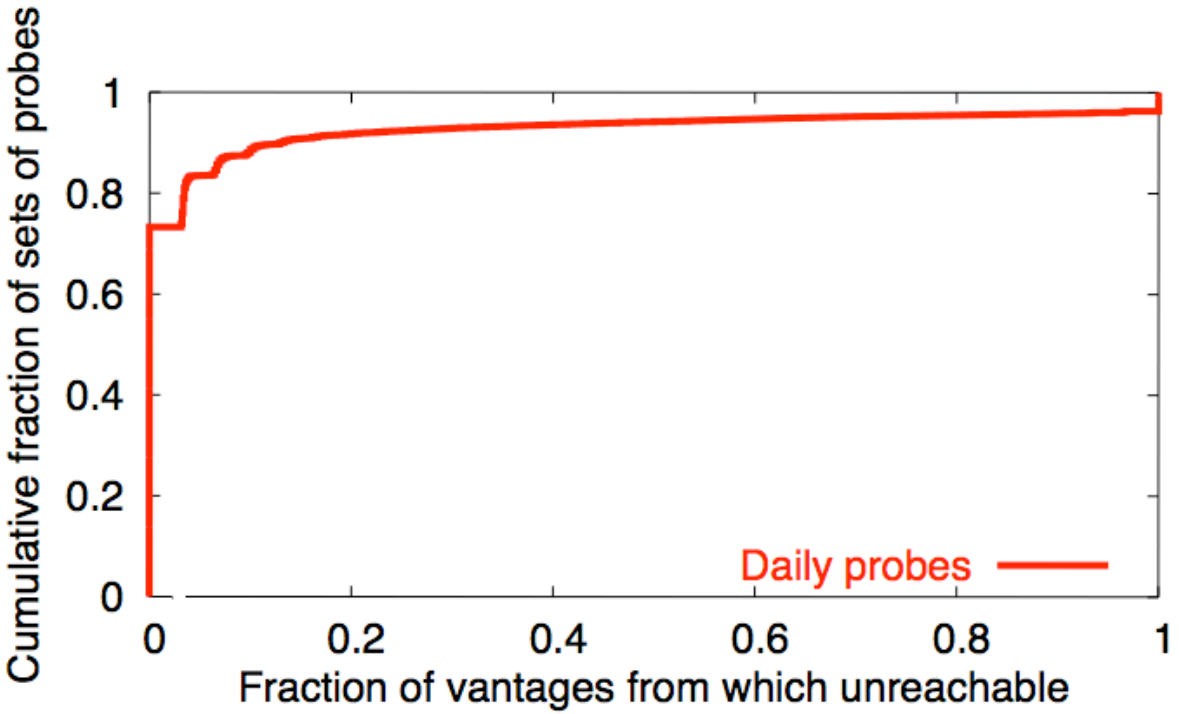

- *"Reachability problems":* >10% vantages don't reach
- On 85% of days, 10K-15k prefixes have problems
- 79,000/110,000 distinct prefixes had problems

Hubble System

Goal: in ***real-time*** on a ***global scale***, ***automatically*** monitor long-lasting reachability problems and classify causes

Approach

- Synthesis of multiple information sources
  - BGP to select targets and learn origin ASes
  - Traceroutes from distributed vantages
  - Interface alias information
  - Pings to check liveness and monitor reachability
- Historical BGP and traceroute data
  - Enables troubleshooting

Assessing a failed traceroute

1. Where is the last hop?
   - Prefix
   - Origin AS for prefix
   - Provider for origin
   - Other AS
2. Can other vantages reach destination?
3. Is the origin single or multi-homed?
4. Can we predict next hop?
   - From historical
   - From other vantages
5. If so:
   - Is it live?
   - Is it in the origin AS?
6. Does path contain loops?
   - Intra-AS
   - Inter-AS

Real-Time Automated Classification

Probed 1500 prefixes every 15 minutes for 2 weeks

Automatically classified 59136/107171 of problems as they occurred

Current simple classification approach:

- Find common cause that explains substantial number of failed traceroutes to a prefix
- Does not have to explain all failed traceroutes
- May classify a given prefix multiple ways
- Not necessarily pinpointing exact problem, may be on handoff or return path

6 classes currently

Single-homed Origin AS Down

- No probes reach single-homed Origin AS
- Some reach its provider

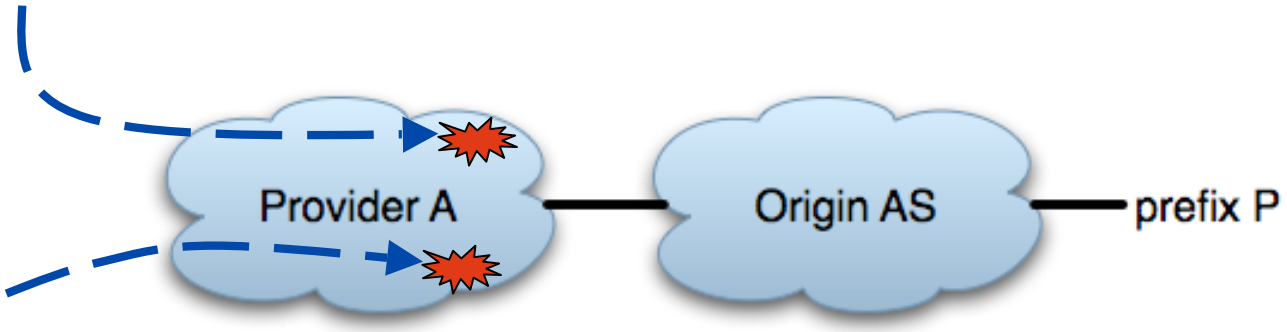

31% of classified problems (6-48% at any point in time)

Multi-homed Origin AS Down

- No probes reach multi-homed Origin AS
- Some reach its provider(s)

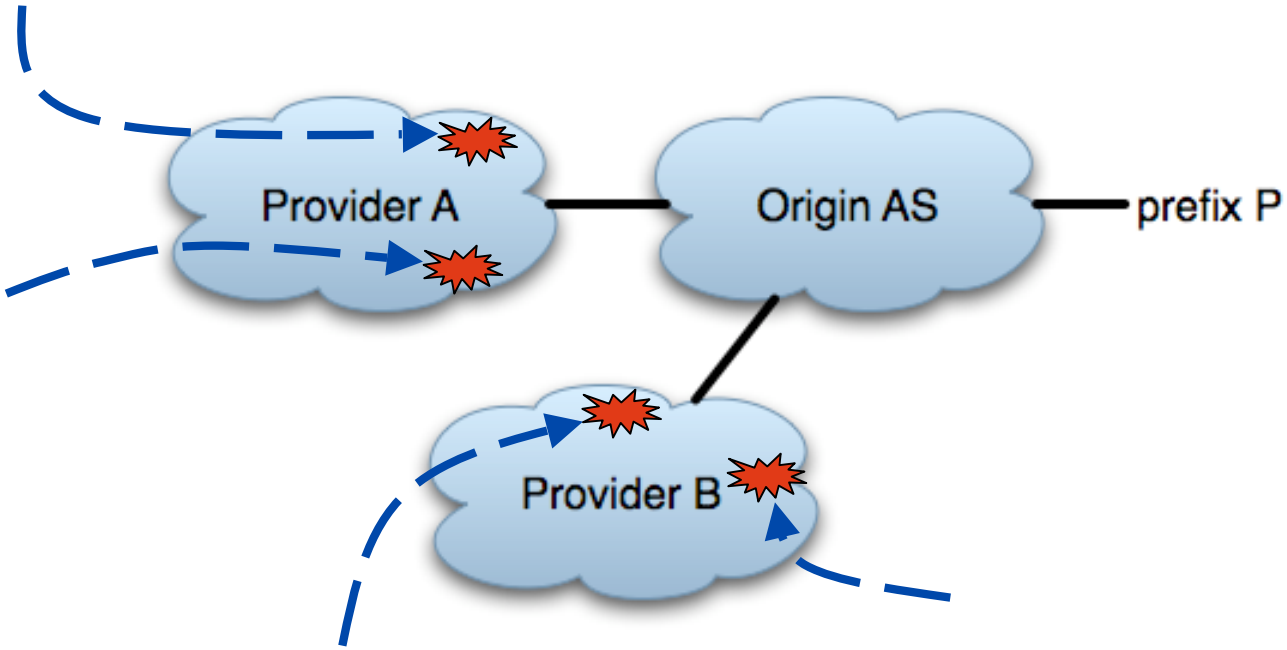

13% of classified problems (2-32% at any point in time)

Provider AS Problem for Multi-Homed

- Probes through Provider ***B*** fail to reach ***P***
- Some reach through Provider ***A***

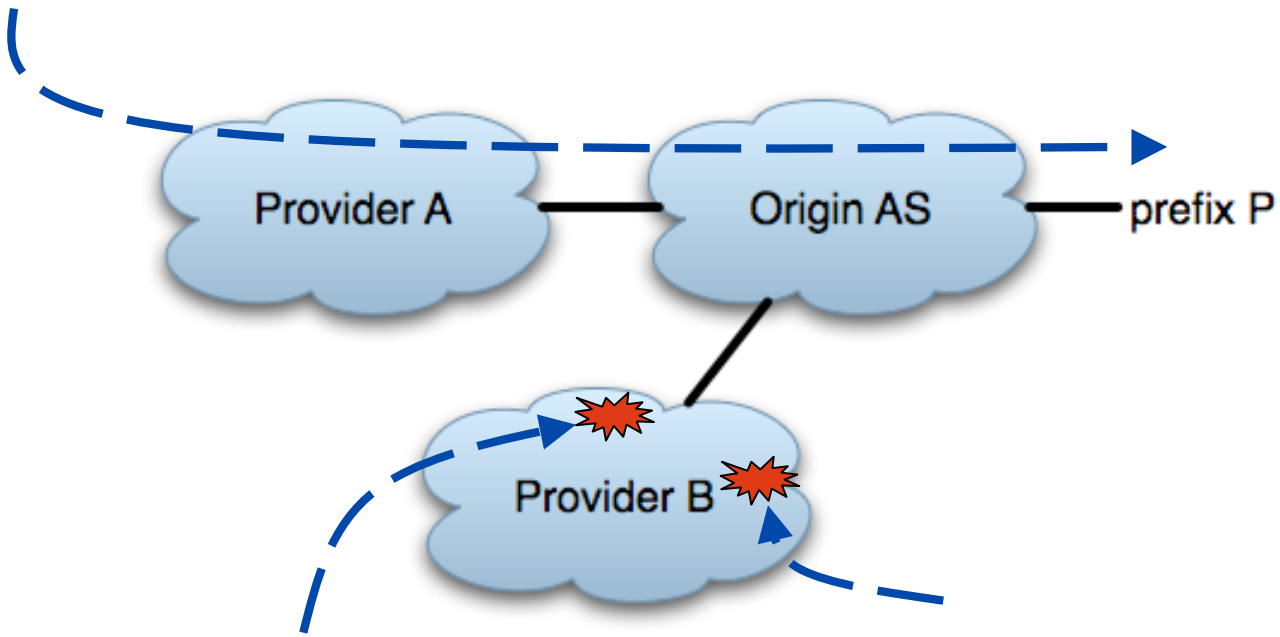

16% of classified problems (3-50% at any point in time)

Non-Provider AS Problem

- Probes through Non-Provider ***C*** fail
- Some reach through other Ases

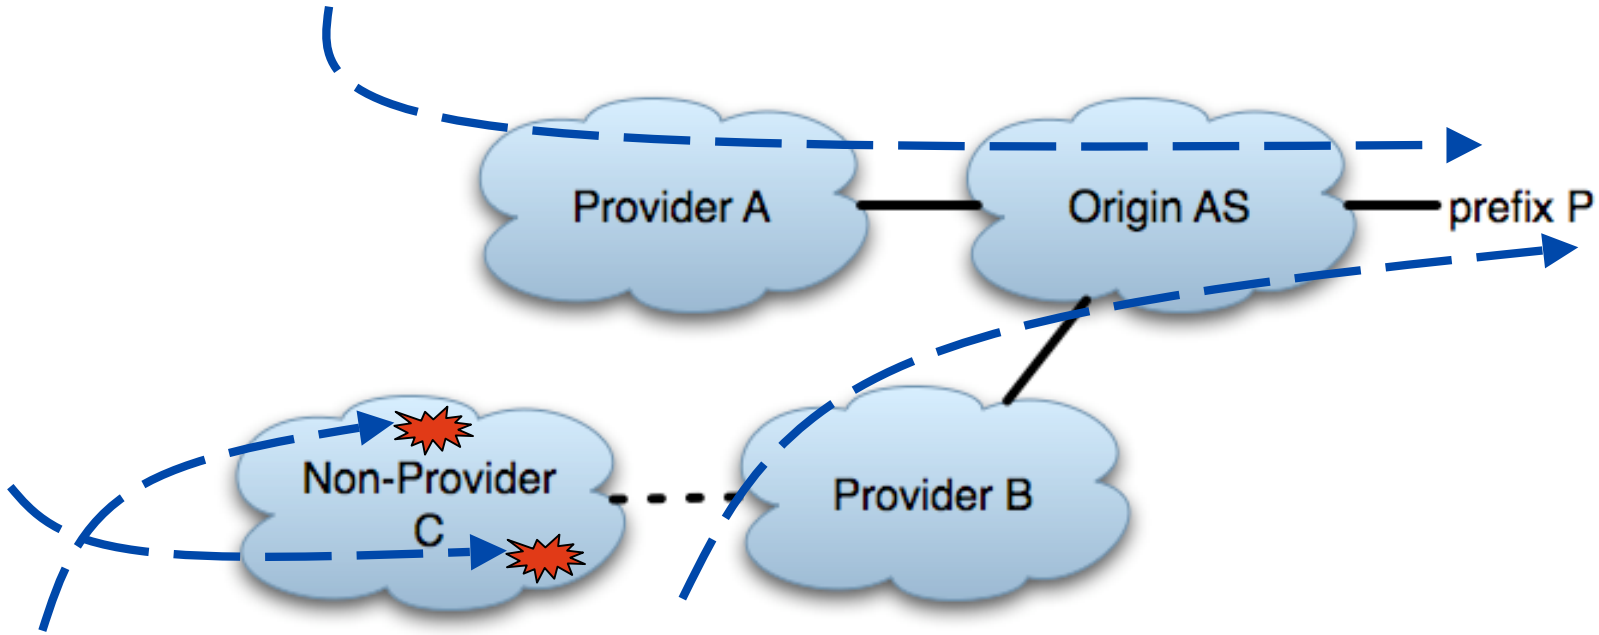

21% of classified problems (3-84% at any point in time)

Router Problem on Known Path

- Last hop router ***R*** was seen on recent paths reaching ***P***
- No probes reach ***P*** through ***R***
- Some reach through ***R'***s AS

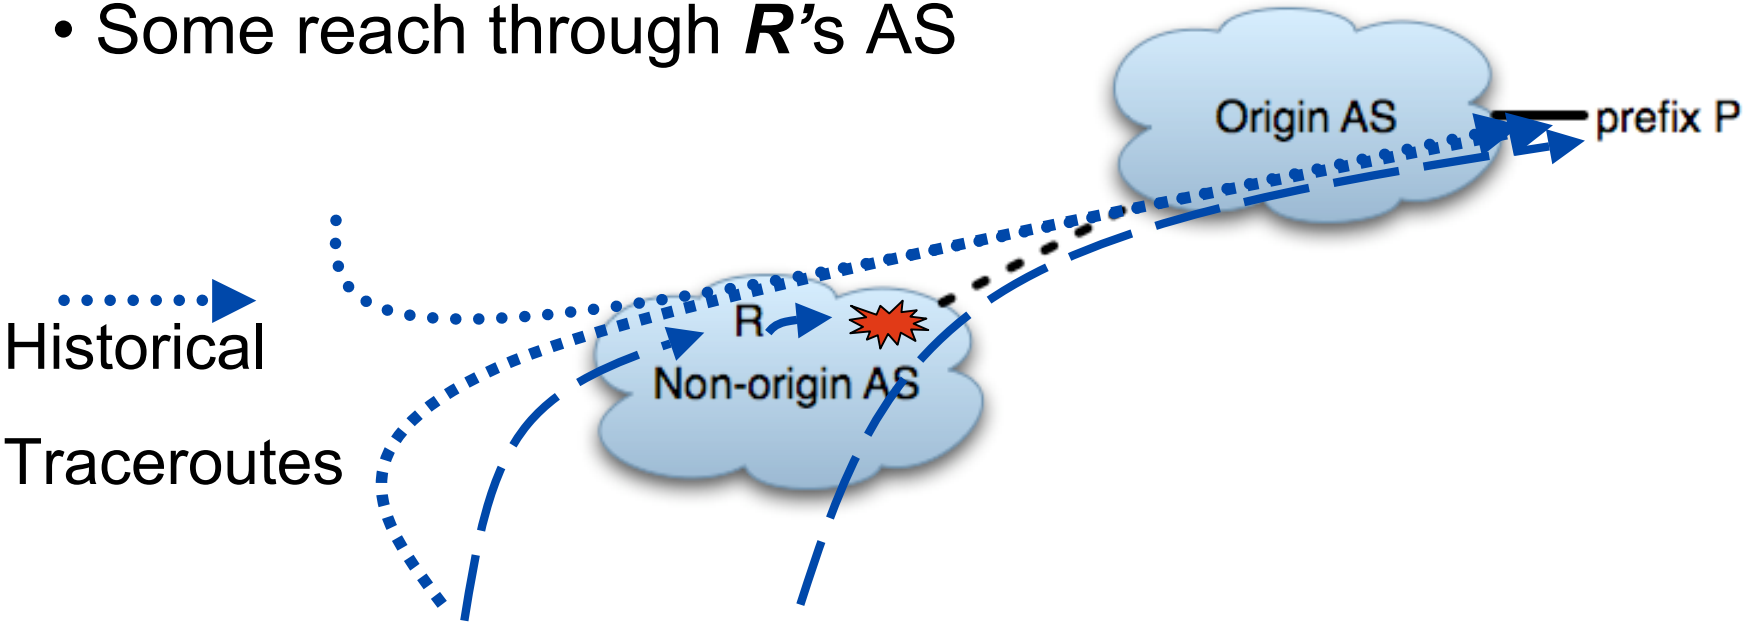

6% of classified problems (1-25% at any point in time)

Router Problem on New Path

- Last hop router ***R*** not seen on recent paths reaching ***P***
- No probes reach ***P*** through ***R***
- Some reach through ***R'***s AS

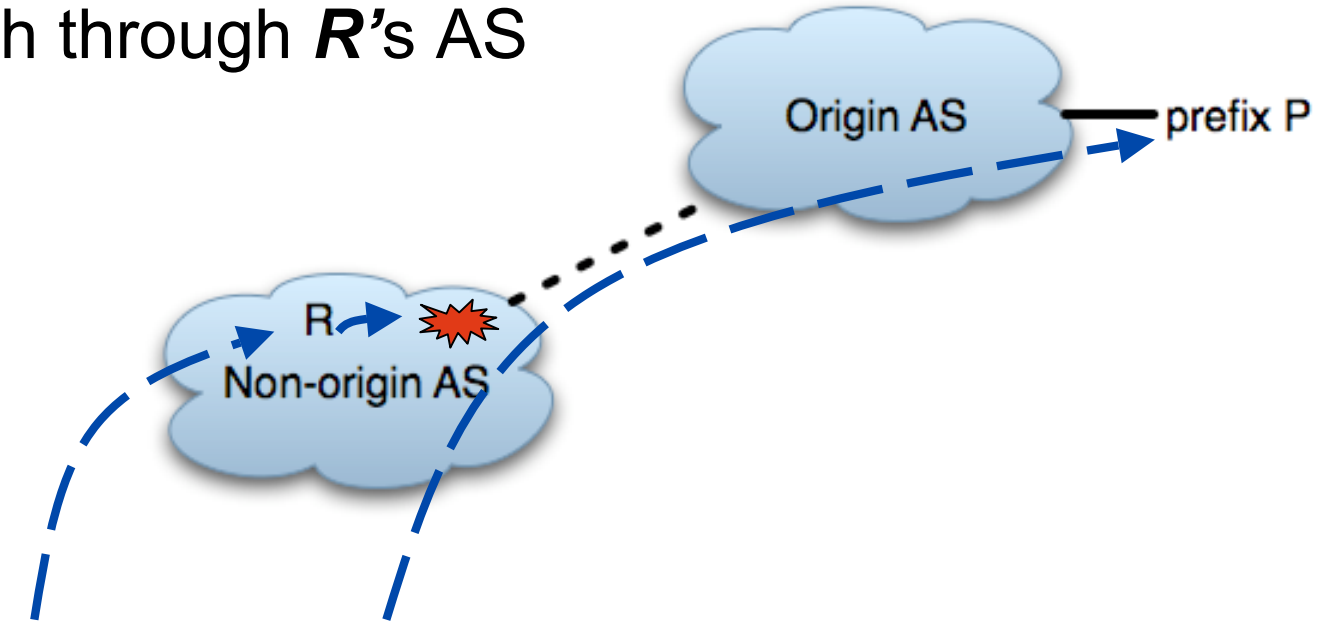

26% of classified problems (3-50% at any point in time)

Preliminary classification results

Of ones we classify: Overall (range over time)

1. Single-homed origin AS down: 31% (6-48%)
2. Multi-homed origin AS down: 13% (2-32%)
3. Provider AS problem for multi-homed origin AS: 16% (3-50%)
4. Non-provider AS problem: 21% (3-84%)
5. Router problem on old path: 6% (1-25%)
6. Router problem on new path: 26% (3-50%)

Target Selection to Reach Internet-Scale

- Maintain current and historical BGP snapshots
- Ping responsive prefixes to check reachability
- Use updates and pings to select prefixes likely to be experiencing reachability problems
  - New origins
  - BGP route changes at many vantages
  - Advertisements for newly allocated prefixes
  - Failed pings

Ongoing work. Plan to evaluate this summer.

Conclusions and Future Work

- Lots of reachability problems, some long lasting
- Historical and fine-grained data enable problem analysis
- Problems with multi-homed failover

Future:

- Hybrid probing to reach Internet-scale
- Predict availability of paths/prefixes
- Query language/ interface

NANOG and Hubble

How can we help you?

- Access to queriable real-time and historical traceroutes and reachability analysis
- Other problems or causes to look for?
- Please email ***ethan@cs.washington.edu***

How you can help us?

- Validation of specific problems to help refine our techniques