Large Route Leak Detection

Qing Ju, Varun Khare, Beichuan Zhang
University of Arizona

Route Leak/Prefix Hijack

- An unauthorized network announces prefixes of other networks.
  - Prefix owner: the destination of the traffic.
  - Attacker: the blackhole of the traffic.
  - Other networks: the source of the traffic.

- Both the prefix owner (traffic destination) and other networks (traffic source) are victims.

Current Practice

- Only prefix owner deals with leak/hijack.
  - A monitoring system, such as Cyclops, MyASN, BGPMon, sends alerts to the prefix owner.
  - Prefix owner decides which one is a real incident.
  - Prefix owner contacts attacker or his upstream ISP to stop the attack.

- Problem: the whole process takes time, during which data traffic is vulnerable.
  - E.g., the YouTube case took 2 hours to resolve. In the meantime users experienced YouTube outage.

Different parties in a leak/hijack incident

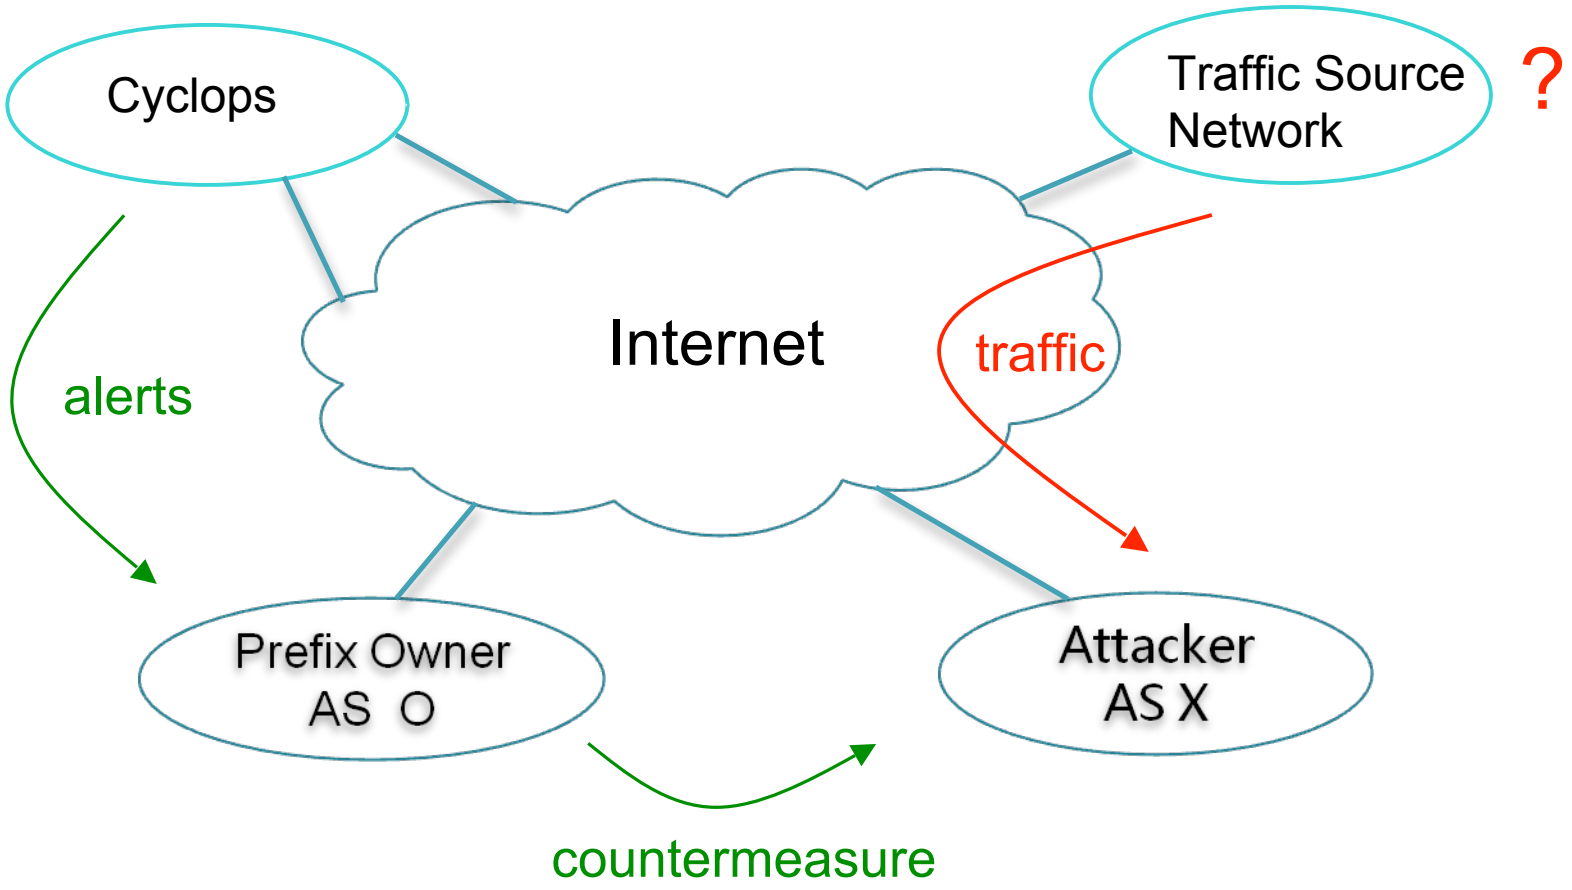

Protect My Traffic

- How do networks other than the prefix owner protect their traffic before the attack is resolved?
  - **Identify and drop false routing announcements.**
- It is very difficult to *accurately* identify *all* false routing announcements without authoritative knowledge from the prefix owner.
  - There are many legit origin changes.
- There are cases relatively easier to detect.
  - Improve upon what we have now.

Large Route Leaks (LRL)

- Sometimes a network hijacks prefixes of multiple other networks, likely due to misconfiguration.
  - More often than you thought or reported on NANOG list.

- Our goal is to *automatically* detect these incidents.
  - Without help from prefix owner.
  - Try to minimize false positives.
    - We may miss some incidents, but what we report are highly likely to be real incidents.

- So that networks (non prefix owners) can respond to these attacks quickly to protect their traffic.

Detecting Large Route Leaks

- Basic observation:
  - When an AS announces a prefix of another network, it is difficult to tell whether this is legit or not.
  - When an AS announces prefixes of ***many*** different networks at the ***same time***, it is very likely that this is a hijack/leak.

- Basic approach:
  - Get all origin changes from BGP routing updates.
  - Find all *suspicious* origin changes.
  - Correlate the suspicious origin changes along time as well as attacker AS to identify LRL events.

Narrowing Down Suspicious Events

- The raw BGP data contains way too many origin changes, and most of them are legit.
- We filter out the following ones.
  I. The AS has announced the prefix for more than one day in the past year.
  II. The AS has announced a super-prefix for more than one day in the past year.
  III. The AS has a stable inter-domain link connected to the AS that normally announces the prefix or its super-prefix.
  IV. WHOIS says that both new and old origin ASes belong to the same organization.
  V. IXP prefixes.
- This filtering does not have to be perfect. It just reduces the noise in the later results.

Identifying LRL Incidents

- After the previous step of filtering, if an AS still announces prefixes that are normally announced by N different networks, we say this AS has an *offense value* of N.
  - N is mostly 1 or 2 for the vast majority of events.
- We set N=10 as the threshold to become an LRL incident.

Distribution of Offense Values

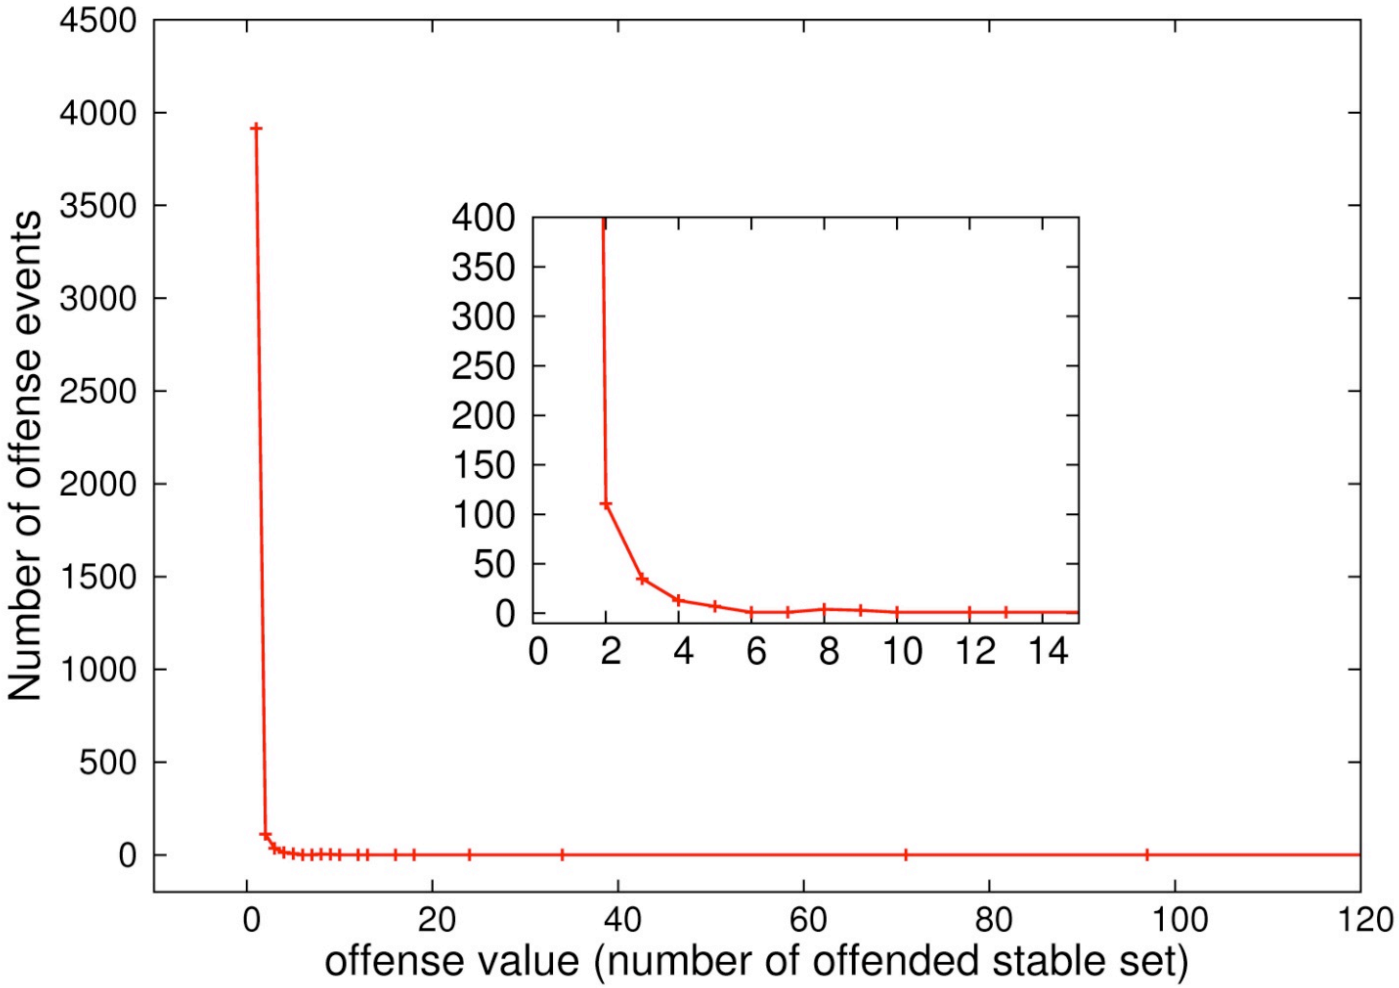

- N=10 is chosen as the threshold for LRL.

Number of LRL Incidents Detected

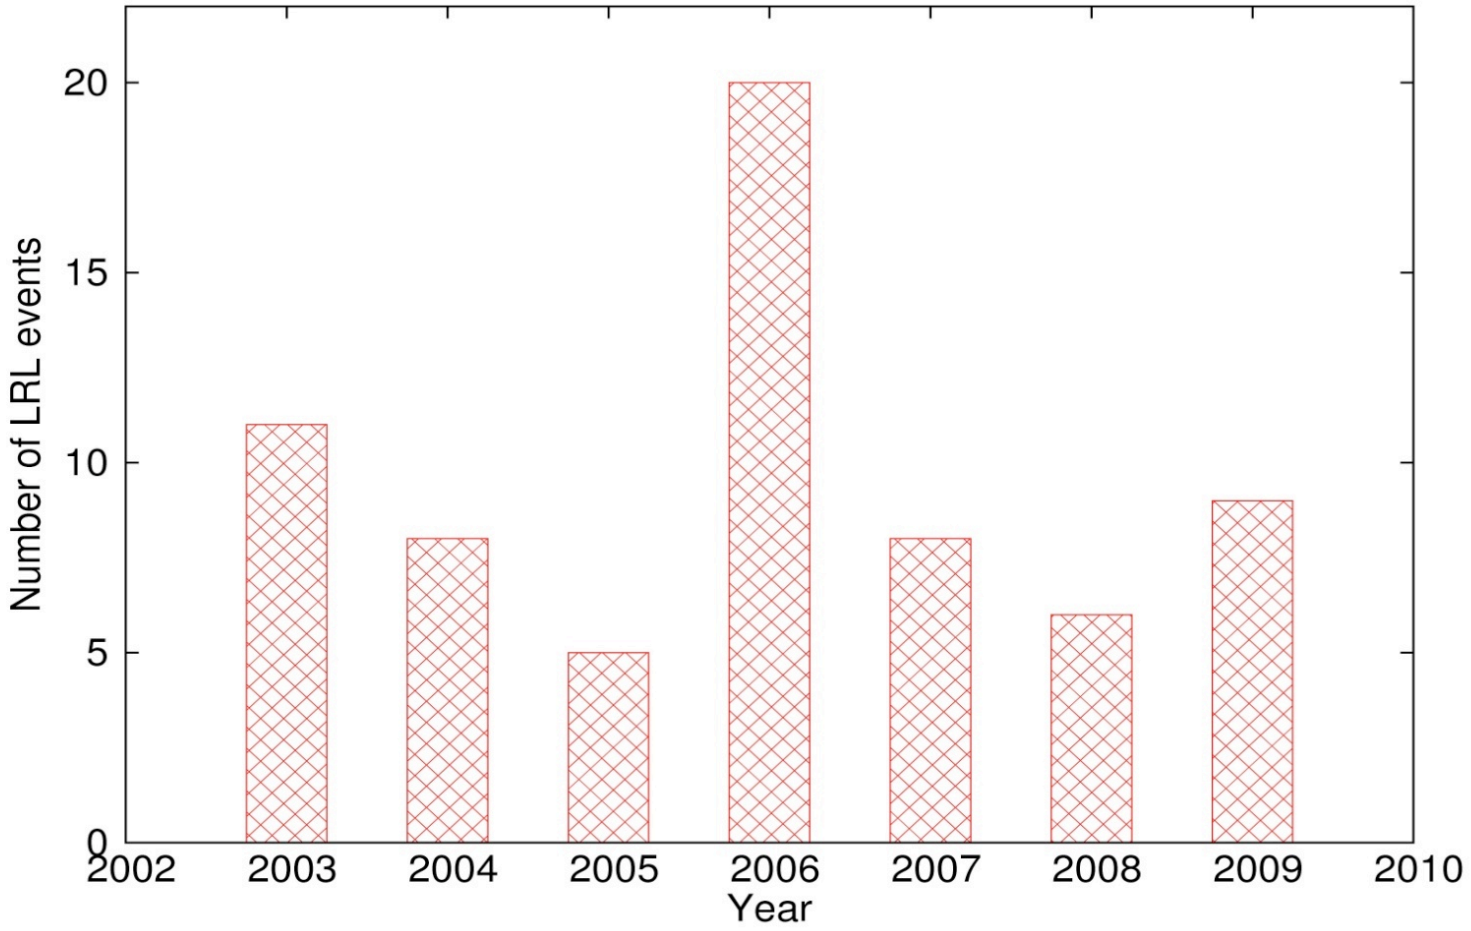

- RouteViews Oregon collector data, 2003-2009.

How Accurate and Useful Is It?

- Email to victim networks to confirm.
- All 9 incidents in 2009 and 6 incidents in 2008 have been confirmed as real route leaks/hijacks.
- Only a full table leak in 2008 was reported on NANOG list. None of the other 14 incidents was reported.
- Even many victim networks were not aware of them
- Though we do not catch all leaks/hijacks, what we are able to catch are still very useful information for operators, especially those who are not the prefix owner.

Nine Incidents Detected in 2009

DATE	ASN	OFFENSE VALUE	AS NAME	DURATION	Country
02/14	8895	34	KACST/ISU	1.96 hours	Saudi Arabia
04/07	36873	13	VNL1-AS	9.98 mins	Nigeria
05/05	10834	97	Telefornia	3.06 hours	Argetina
07/12	29568	16	Comtel Supernet	23.45 mins	Romania
07/22	8997	170	OJSC NorthWest Telecom	59 secs	Russia
08/12	4800	12	Lintasarta-AS-AP	32 secs	Indonesia
08/13	4800	71	Lintasarta-AS-AP	7.82 hours	Indonesia
12/04	31501	18	SPB-Teleport	68 secs	Russia
12/15	39386	24	Saudi Telecom	62 secs	Saudi Arabia

A Case Study

- On February 14th, 2009
  - AS 8895 (KACST/ISU, Saudi Arabia) originated 243 prefixes belonging to 34 Saudi ASes for about 2 hours.
  - A total of 41 out of 43 Routeviews Oregon monitors observed it.
  - Confirmed by a victim Saudi ISP operator via email.

- What happened:
  - AS 8895 used to be the upstream provider for many local ISPs before its customers switching to Saudi Telecom (AS39386)
  - But due to misconfiguration, AS 8895 announced prefixes of many ex customers.

A Case Study (cont.)

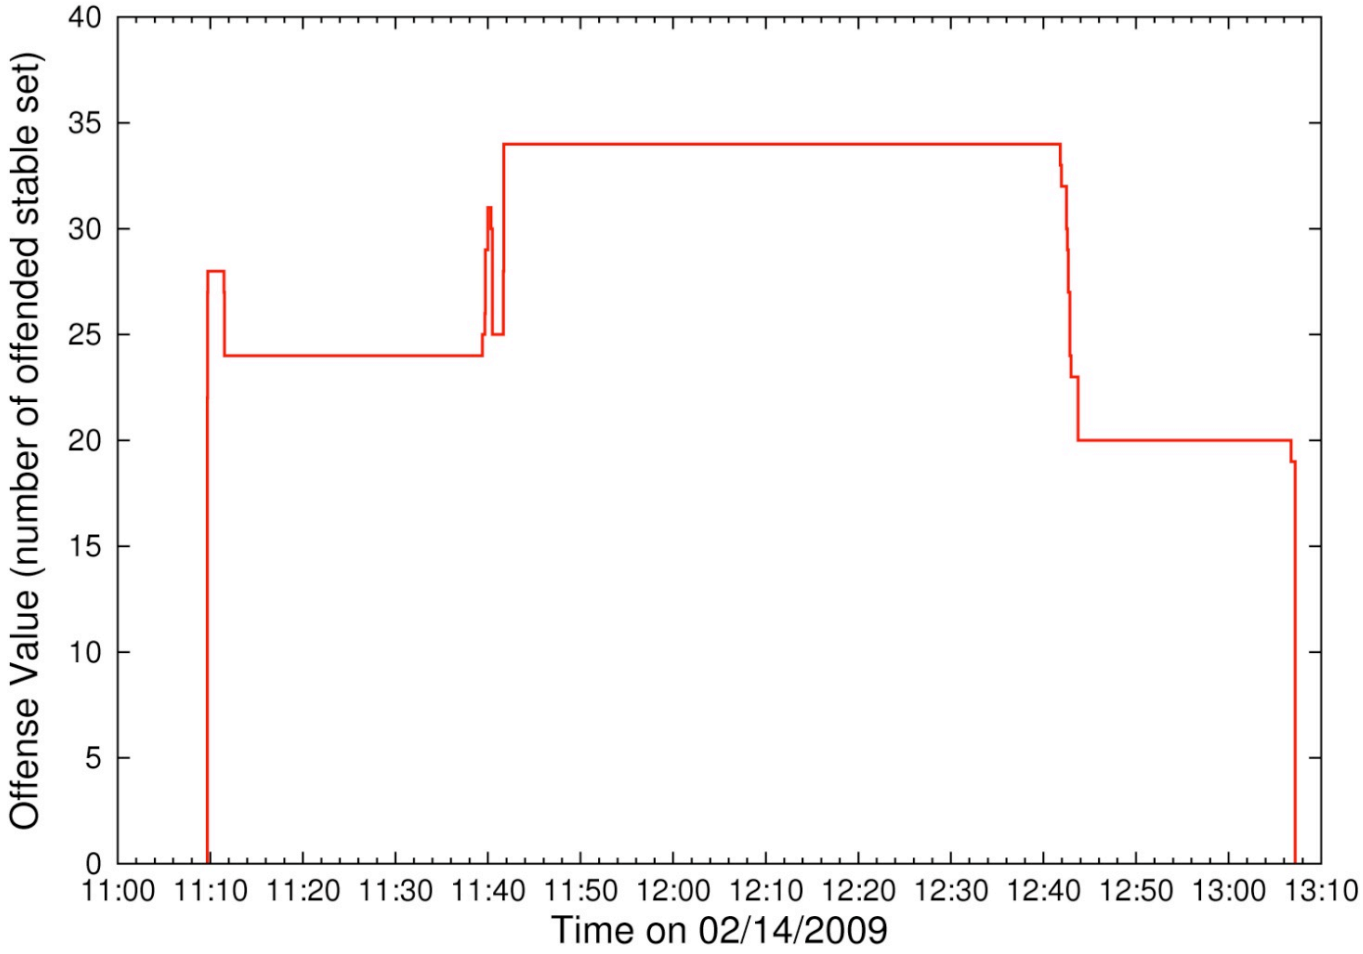

- Offense value was near zero in entire 2009 except February 14th, when the leak happened.

The Duration of LRL Incidents

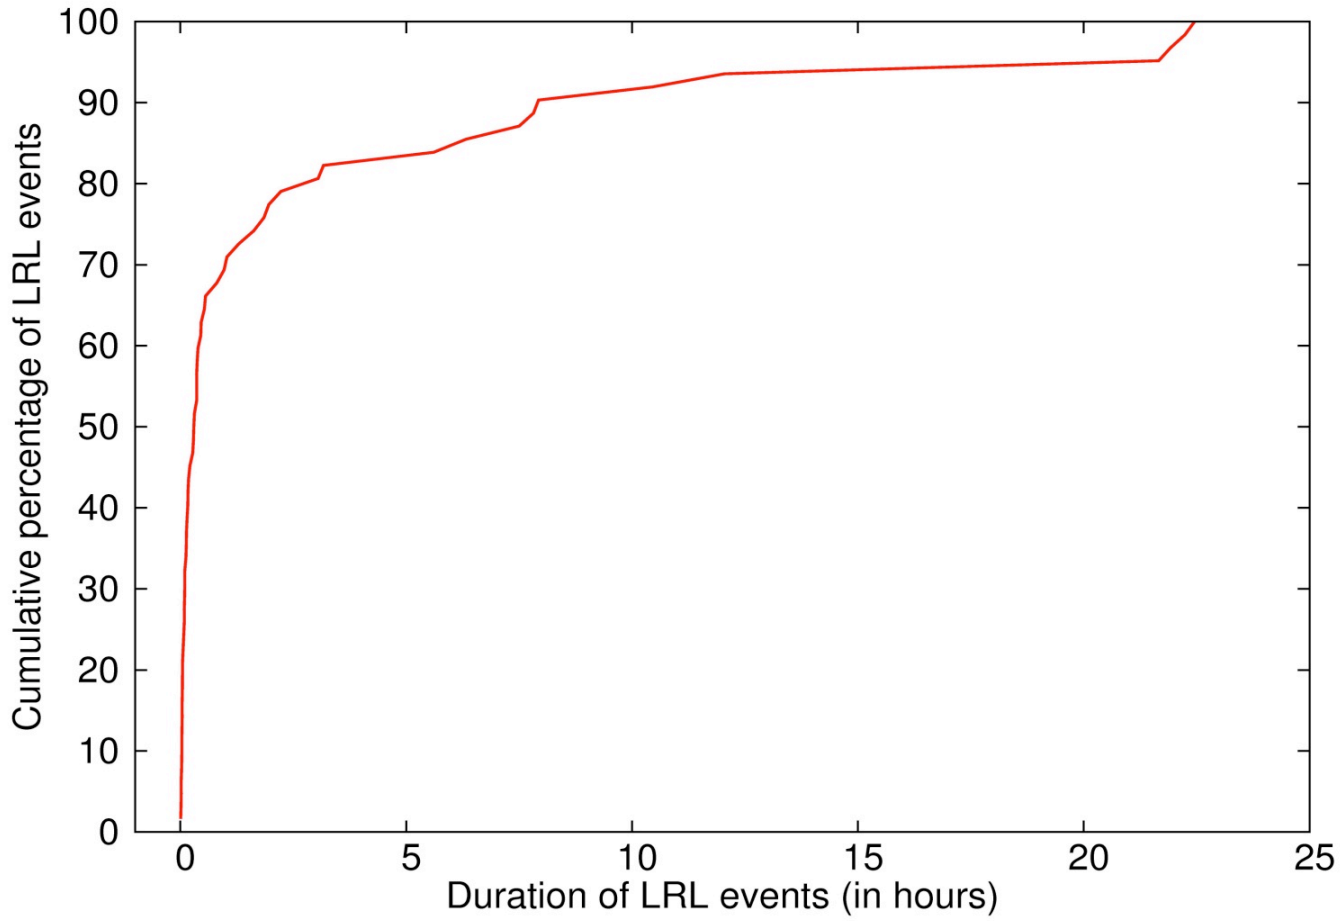

- Most LRL incidents are short, but still 20% of them lasted more than 3 hours.

The Number of Prefixes Offended

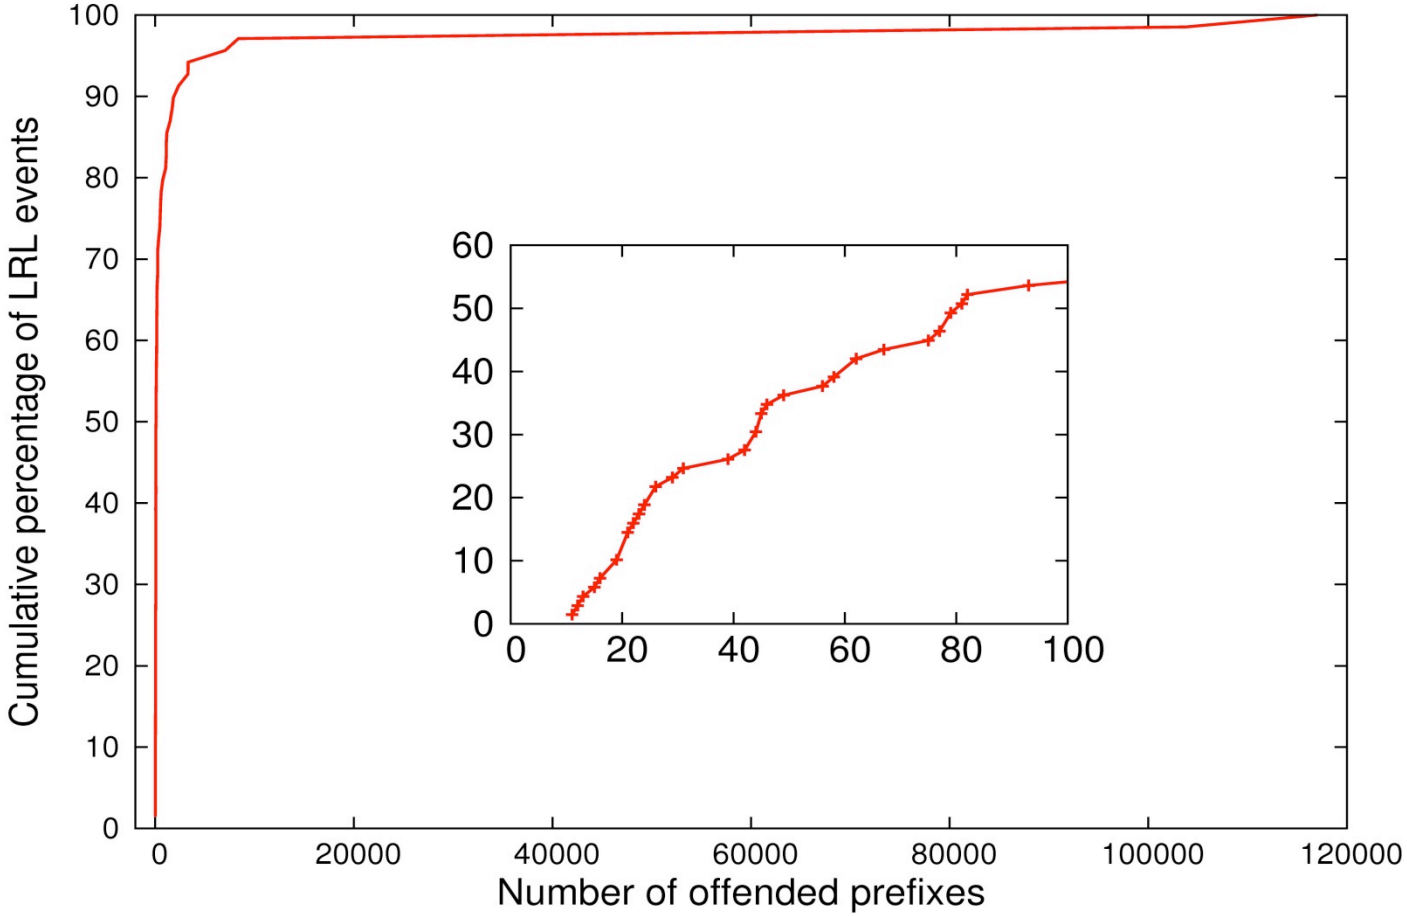

- Most LRL incidents affected tens of prefixes. The medium is 76 prefixes.

Percentage of Monitors Affected.

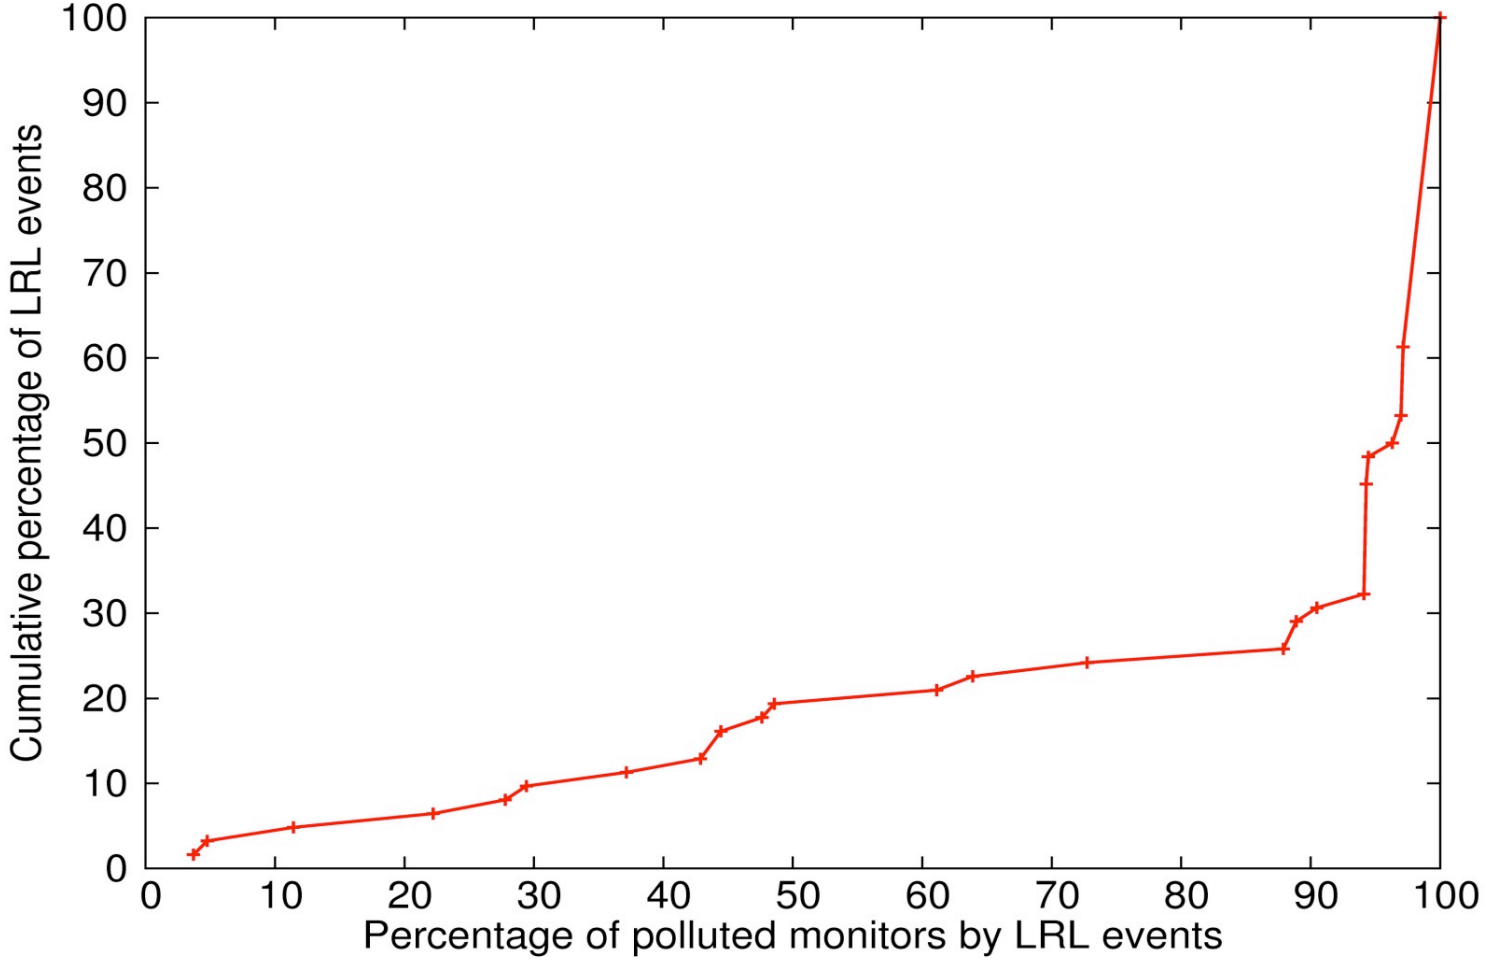

80% of the LRL incidents polluted more than 60% of the monitors from RouteViews Oregon collector.

Comparison with Pretty Good BGP

- Same goal
  - protecting data traffic by non-prefix owner networks before the attack is resolved.
- Complimentary approaches
  - PGBGP: block all new origins for 24 hours
    - No false negative, but many false positives.
    - Only block when there is an alternative path available.
  - LRL detection
    - No or very small false positives, may have many false negatives.
    - Only trigger a small number of alerts that are highly likely real attacks, making it possible to react automatically or very quickly.

Potential Deployment Scenarios

- Operating in the NOC of individual networks
  - Receive live BGP updates from border routers or public source like RouteViews, and generate alerts.
  - Can have multiple levels of thresholds for different actions, e.g.,
    - A high threshold for automatic response.
    - A medium threshold for manual intervention.

- Incorporated into monitoring systems like Cyclops
  - Registered users can receive LRL alerts in addition to alerts regarding their own prefixes.

On-going Work

- Improving the detection algorithm.
- Running the detection with real-time BGP data feed from RouteViews.
- Incorporating into monitoring systems like Cyclops.

Thanks!