Observation and Analysis of BGP Behavior Under Stress

— A Study of BGP's Reaction to the Nimda Worm Attack

Lan Wang, Dan Pei, Lixia Zhang, UCLA
Xiaoliang Zhao, Daniel Massey, Allison Mankin, USC/ISI
Randy Bush, IIJ, Felix Wu, UC Davis
http://fniisc.nge.isi.edu/

Problem Statement

- Investigate BGP behavior during Nimda worm attack in Sept. 2001
  - Renesys Corp.: “Global routing instability …” [1]

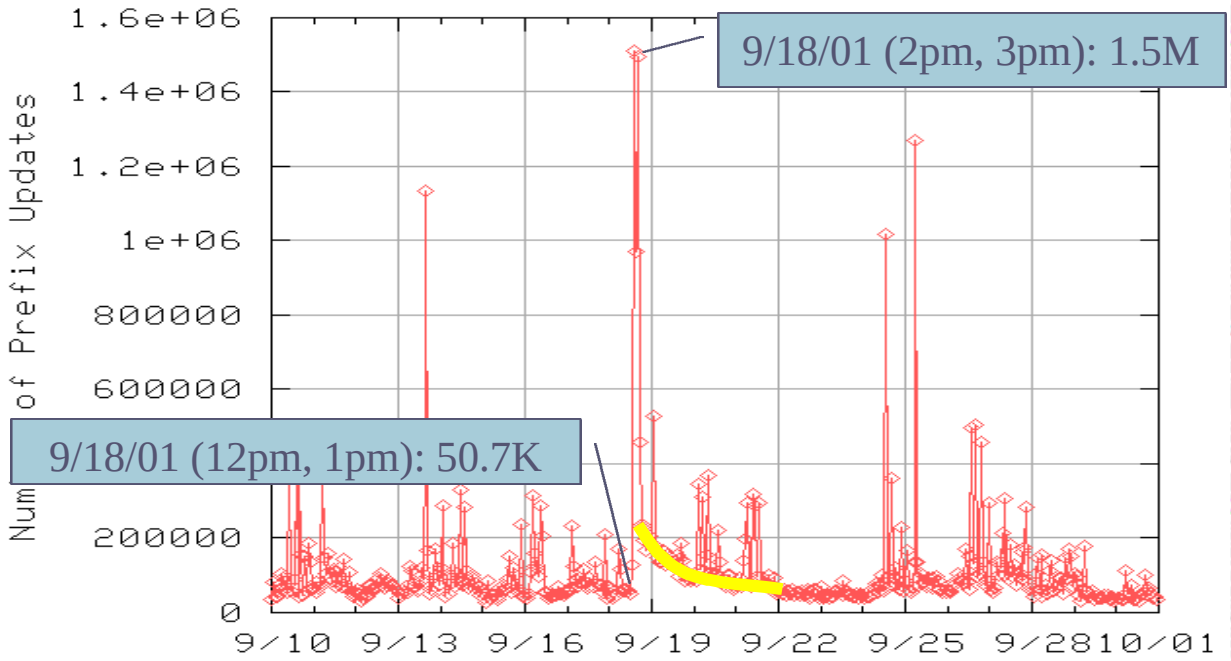

Hourly BGP Update Count at RIPE NCC rrc00

What caused the spike?

- Worm? SANS Institute: “the (Nimda) activity jumped dramatically at approximately 13:00 GMT (on Sept. 18, 2001) and then proceeded to taper off in the following hours …”
- BGP implementation bugs?
- BGP protocol design issues?
- Misconfiguration?
- Other causes?

Data

- BGP update messages collected by RIPE NCC (rrc00) from 9/10/01 to 9/30/01
- 12 BGP Peers
  - 3 US peers: AT&T, Verio, Global Crossing
  - 4 peers in Netherlands: RIPE NCC, SURFnet, Tiscali, KPNQwest
  - 2 peers in Switzerland: CERN, Nextra
  - 1 in Germany: Global Access
  - 1 in London: Global Crossing
  - 1 in Japan: NSPIXP2

Identifying What Caused the Spike

- Classify BGP update messages
- Identify the causes of each class

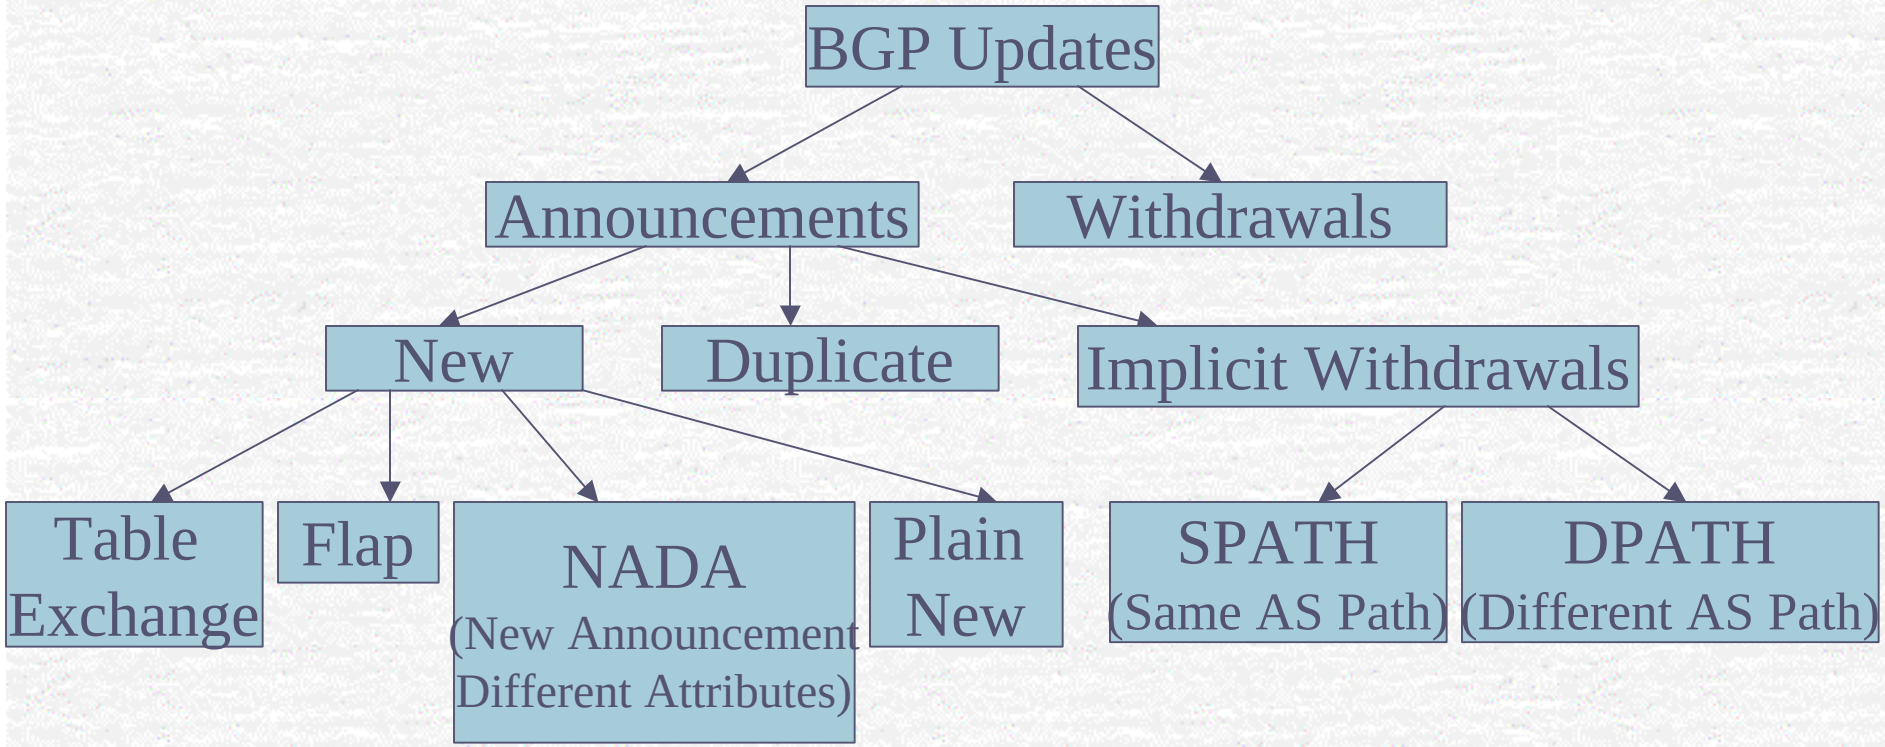

Classification Example

Time	Type	Prefix	ASPATH	Community Attribute
t0	session reset			
	New Ann -> Table Exchange			
t1	A	200.16.216.0/24	3549 701 12956 4926	3549:2256 3549:30840
	Withdrawal			
t2	W	200.16.216.0/24		
	New Ann -> NADA			
t3	A	200.16.216.0/24	3549 6453 8297 12956 4926	3549:2246 3549:30840
	Duplicate			
t4	A	200.16.216.0/24	3549 6453 8297 12956 4926	3549:2246 3549:30840
	Implicit Withdrawal -> DPATH			
t5	A	200.16.216.0/24	3549 701 12956 4926	3549:2826 3549:30840
	Implicit Withdrawal -> SPATH			
t6	A	200.16.216.0/24	3549 701 12956 4926	3549:2725 3549:30840

Daily BGP Update Count

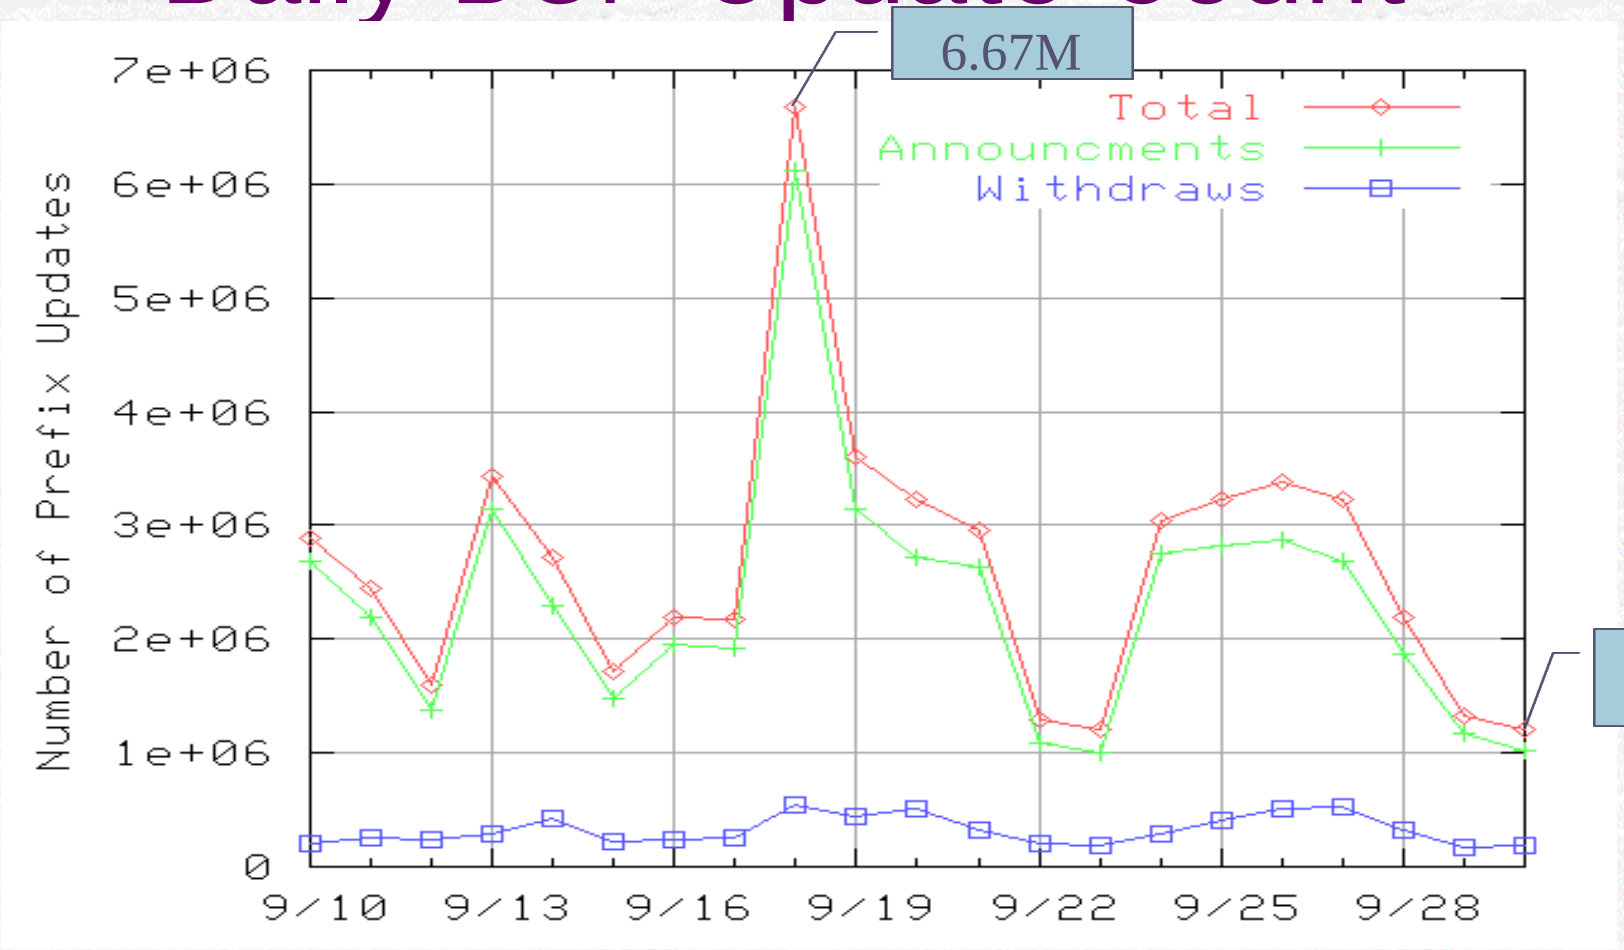

Announcements: 87.3% of the total updates on average (91.7% on 9/18/01)

Announcements Classification

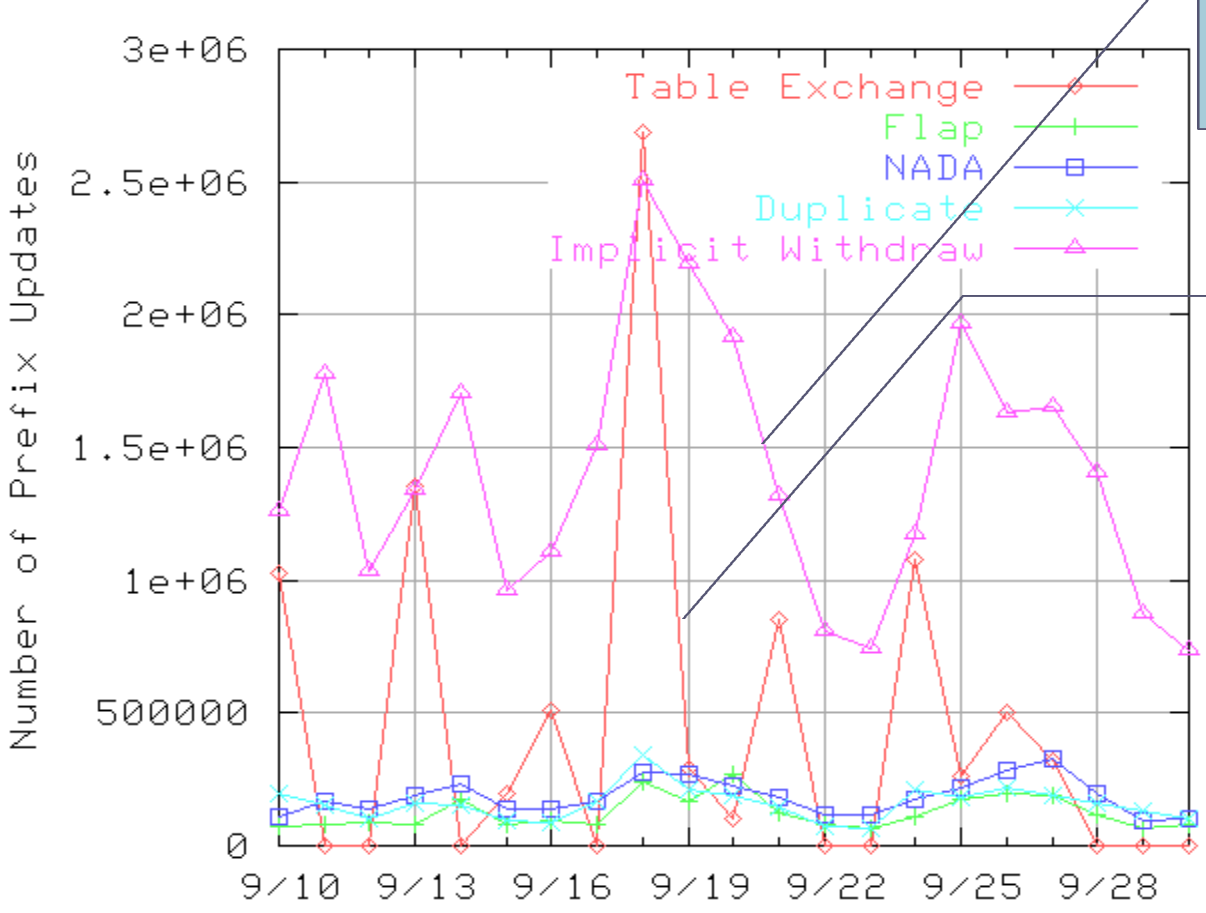

Implicit Withdrawals: largest component (40.9% ~ 81.2%) except on 9/18/01

Table Exchanges: largest component on 9/18/01, second largest on 11 days, 0 on 9 days

BGP Updates on Worm Day

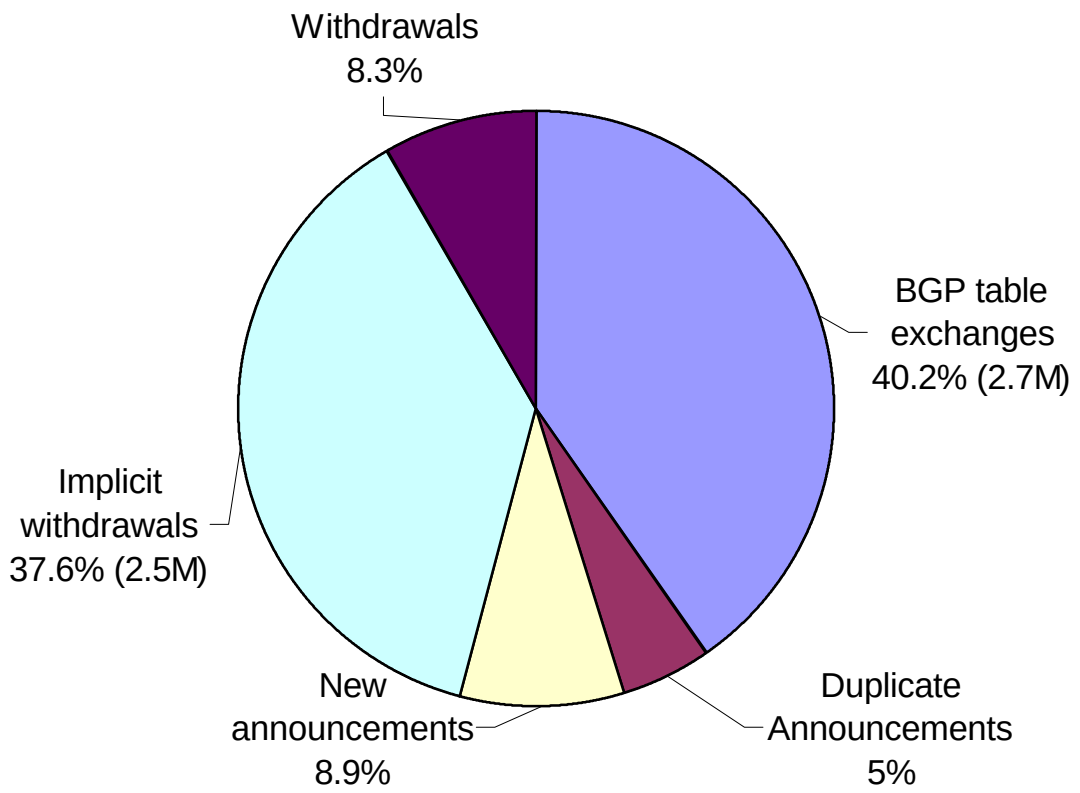

What caused the BGP table exchanges and duplicate announcements?

What caused the implicit withdrawals?

BGP Table Exchanges

2.7M BGP table exchanges on 9/18/01

•30 session resets between the monitoring point and peer routers

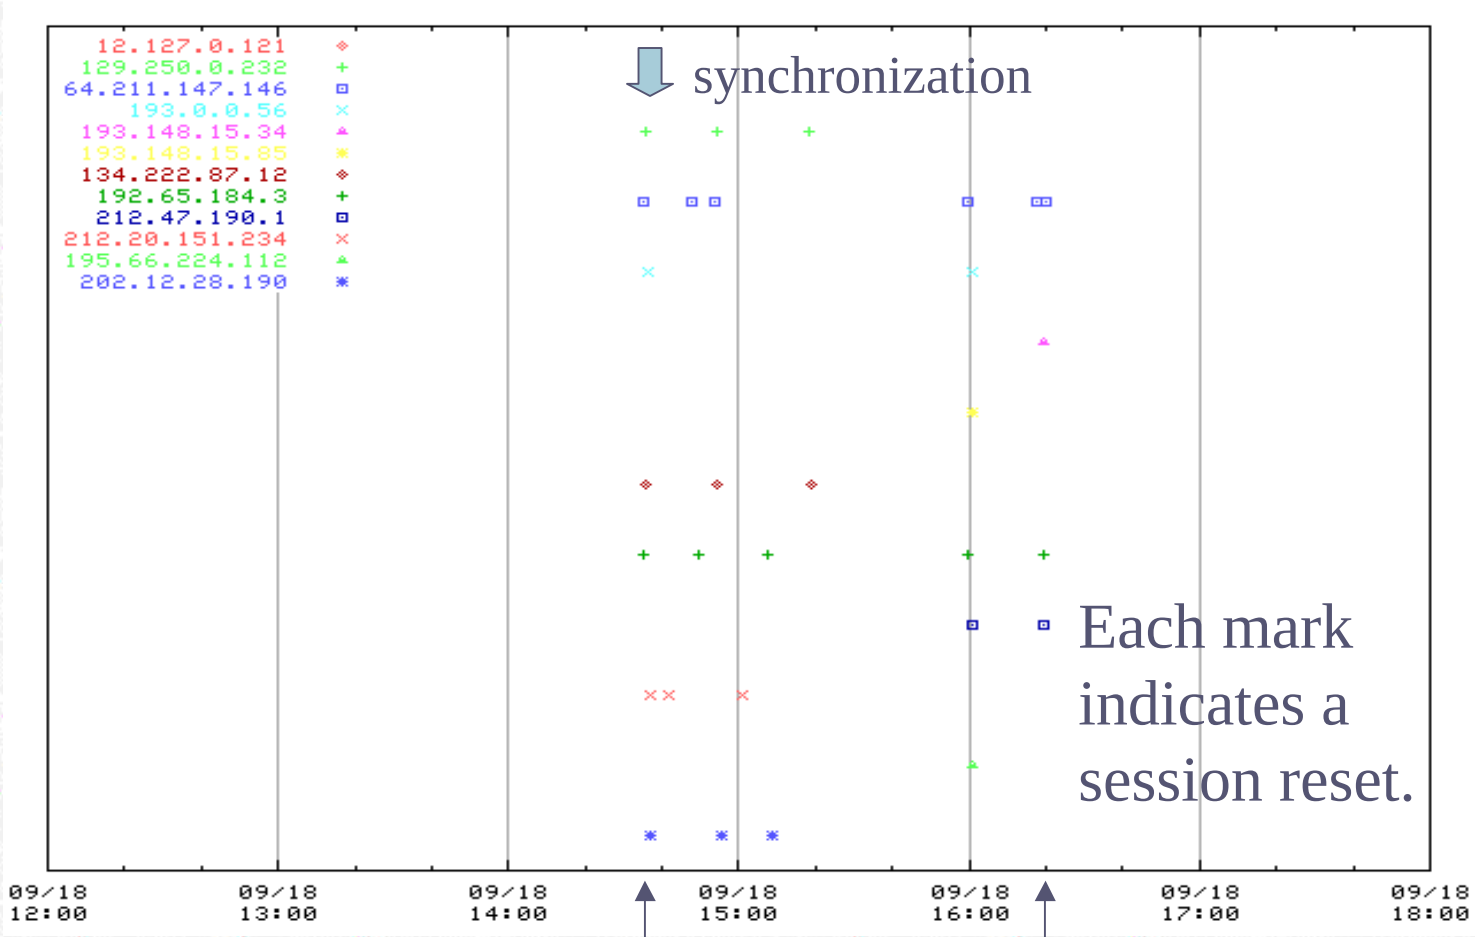

All resets on 9/18/01 occurred between 2:20pm and 4:40pm

What caused the session resets?

All the session resets recovered very quickly

➔ transient problems: congestion or routing problems …

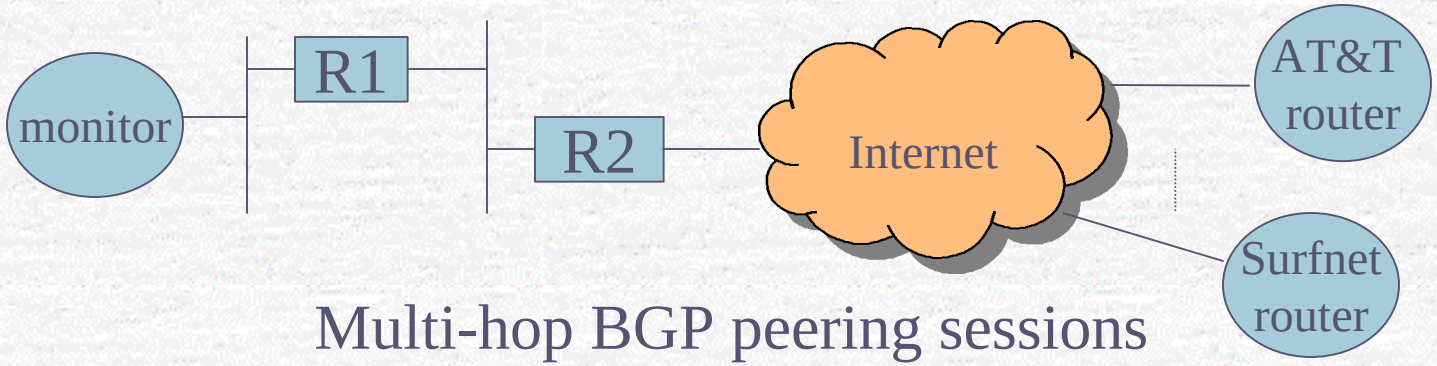

Multi-hop BGP peering sessions

The monitoring process was affected by the worm!

Hourly BGP Update Count (Table Exchanges Removed)

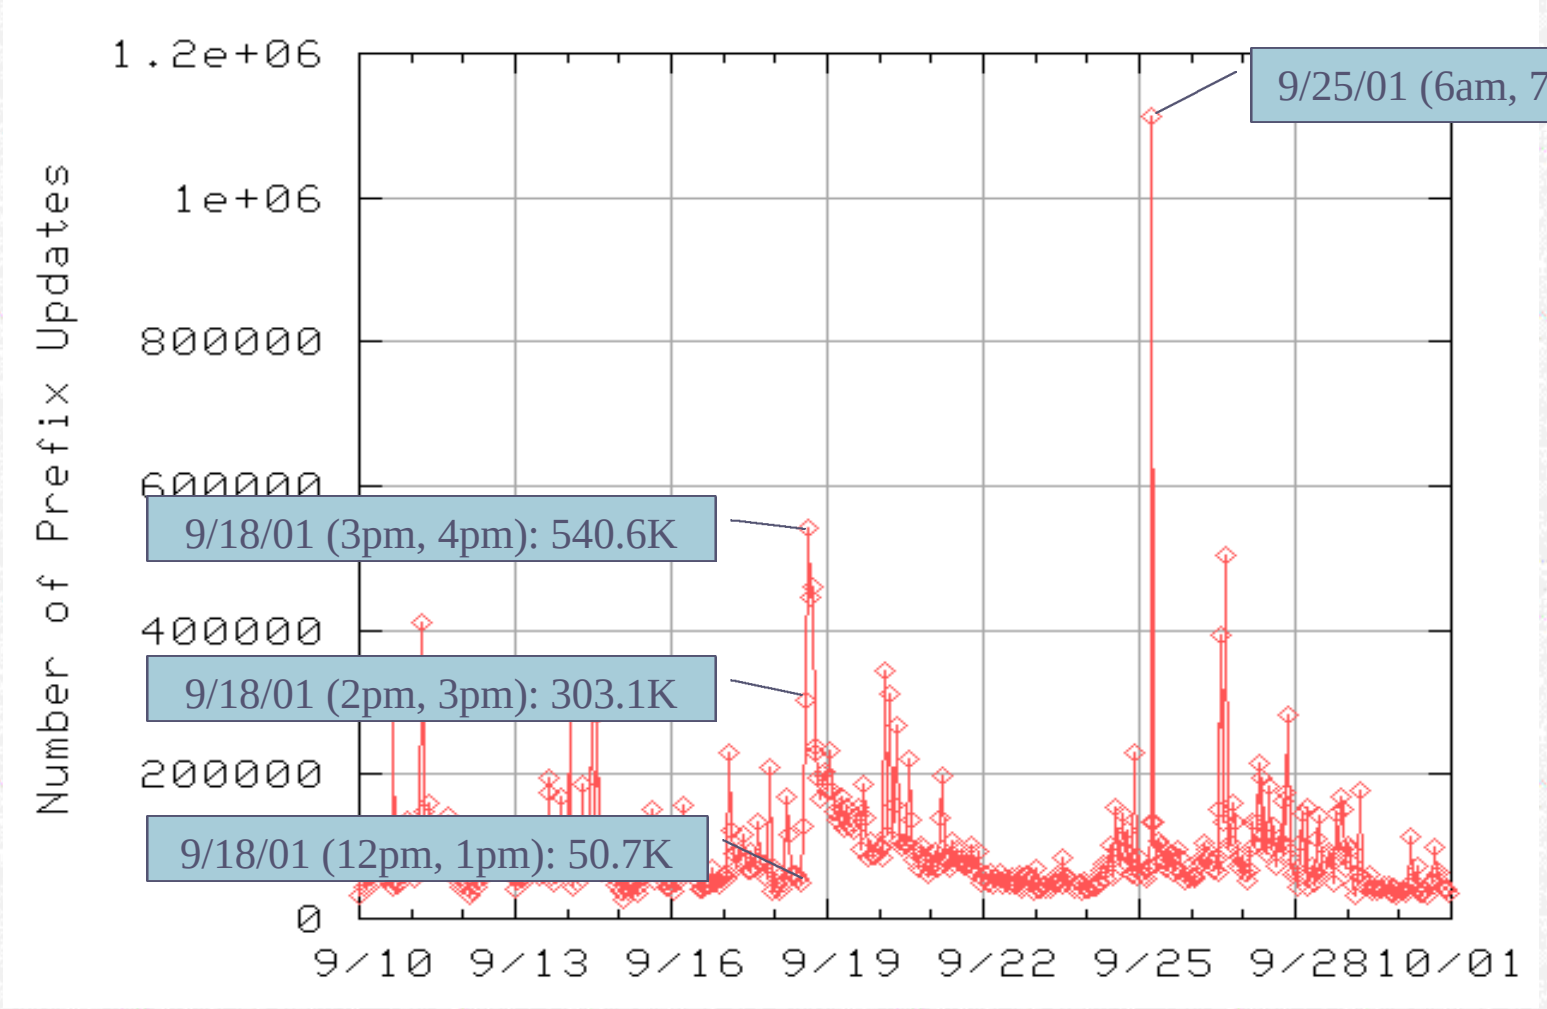

Duplicate Announcements

- 4% ~ 10% of all the updates, but represent 31% of AT&T's updates
- Most likely an implementation problem
  - Triggered by changes to non-transitive attributes

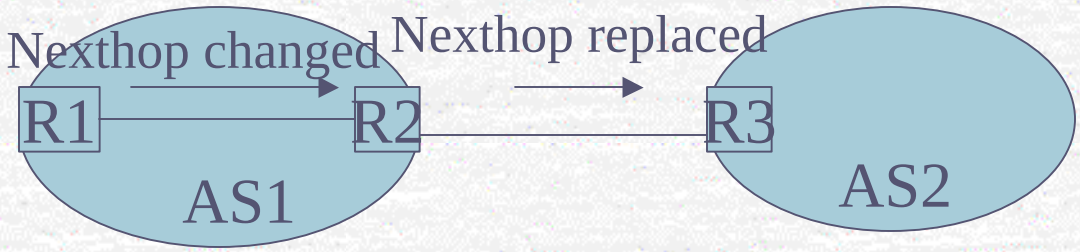

  - Little saving in the implementation increases the overall system overhead

SPATH Implicit Withdrawals

~22% of the implicit withdrawals are SPATH (i.e. the same AS path)

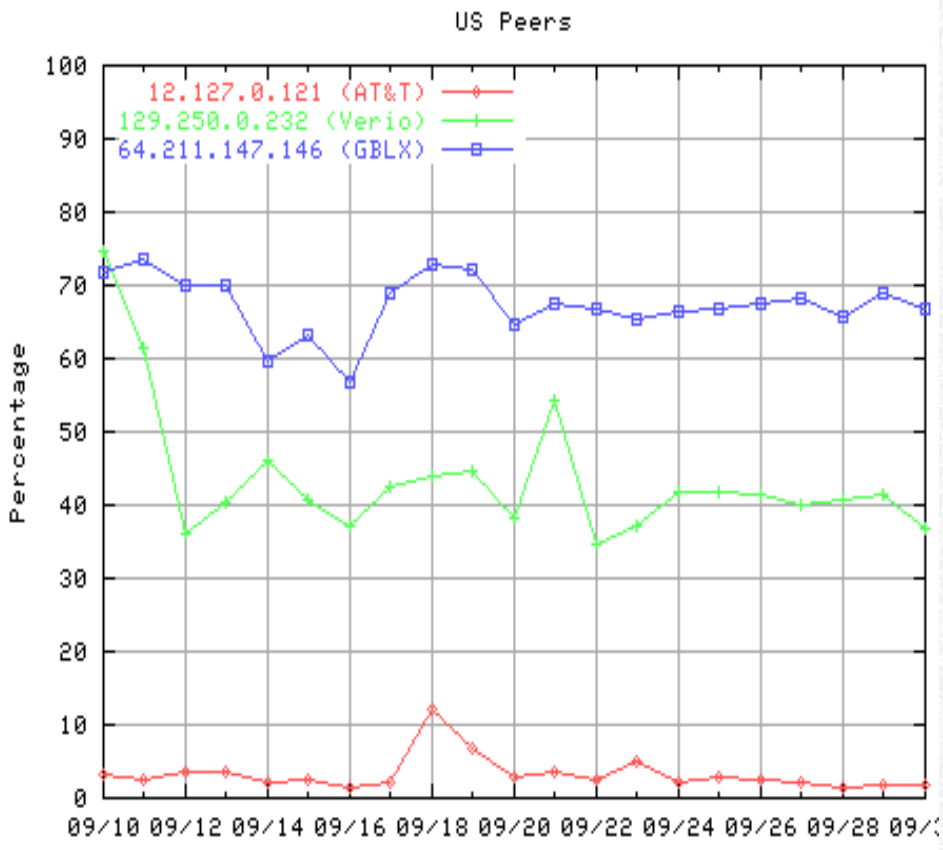

- ~40% of Verio's and ~70% of GBLX's implicit withdrawals are SPATH.
- SPATH doesn't reflect topology changes at the AS level.
- Most likely caused by internal network instability or policy changes.

Hourly BGP Update Count
(Table Exchanges, Duplicates, SPATH removed)

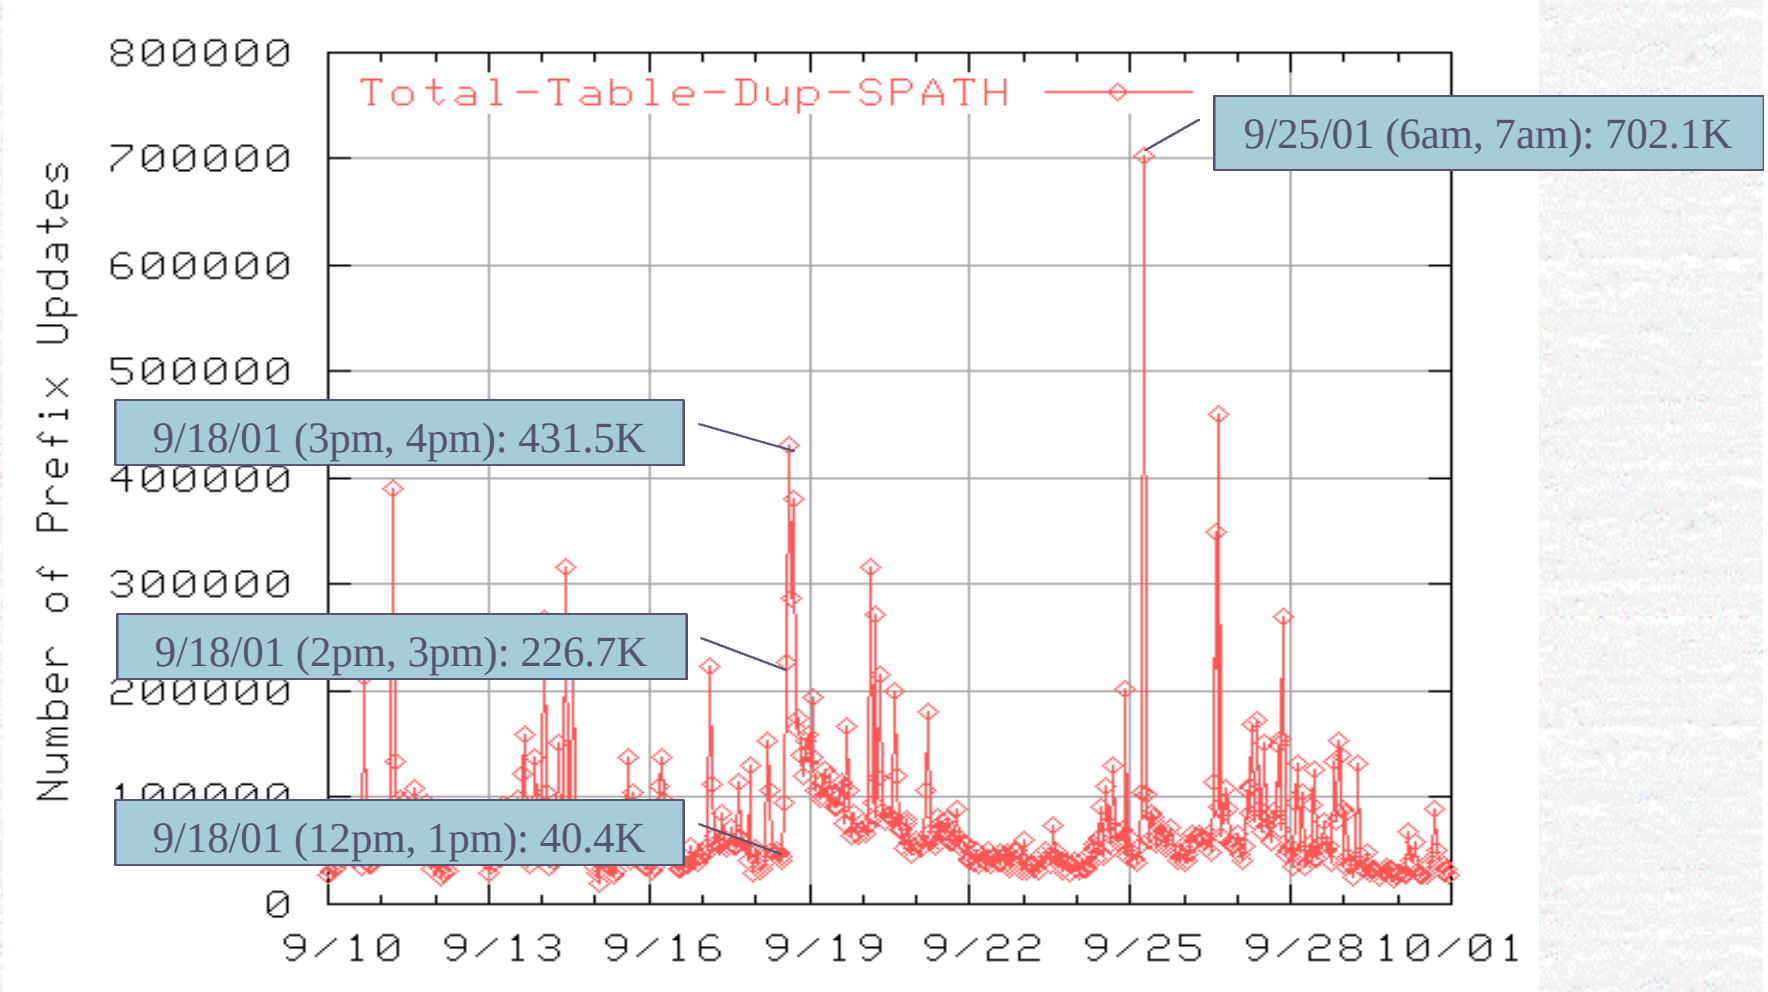

Spikes in DPATHs

Spike in DPATHs (Surfnet)

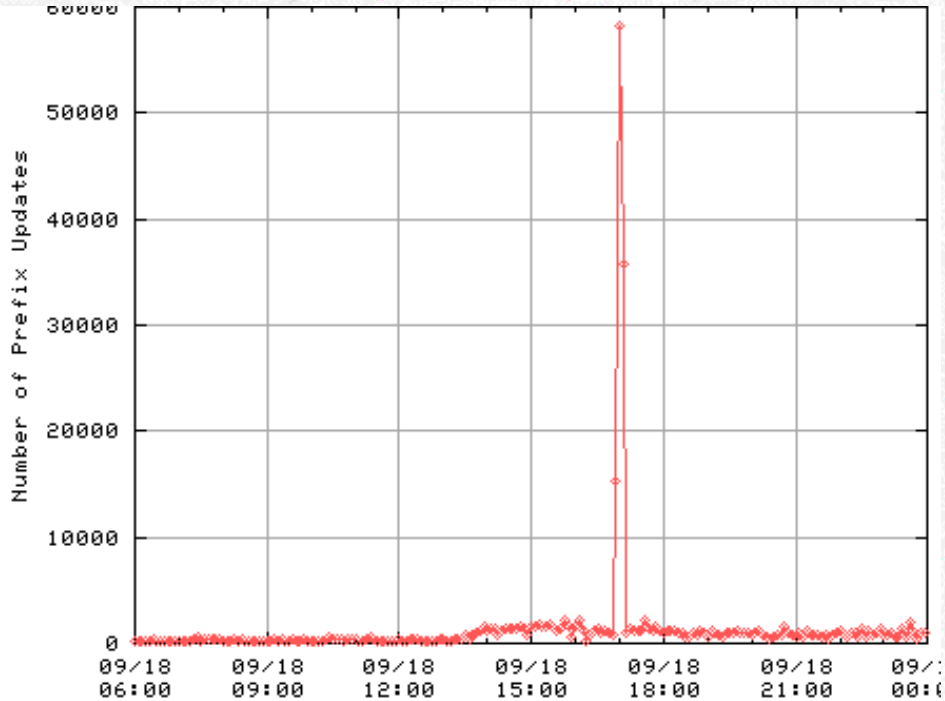

Spike Removed (Surfnet)

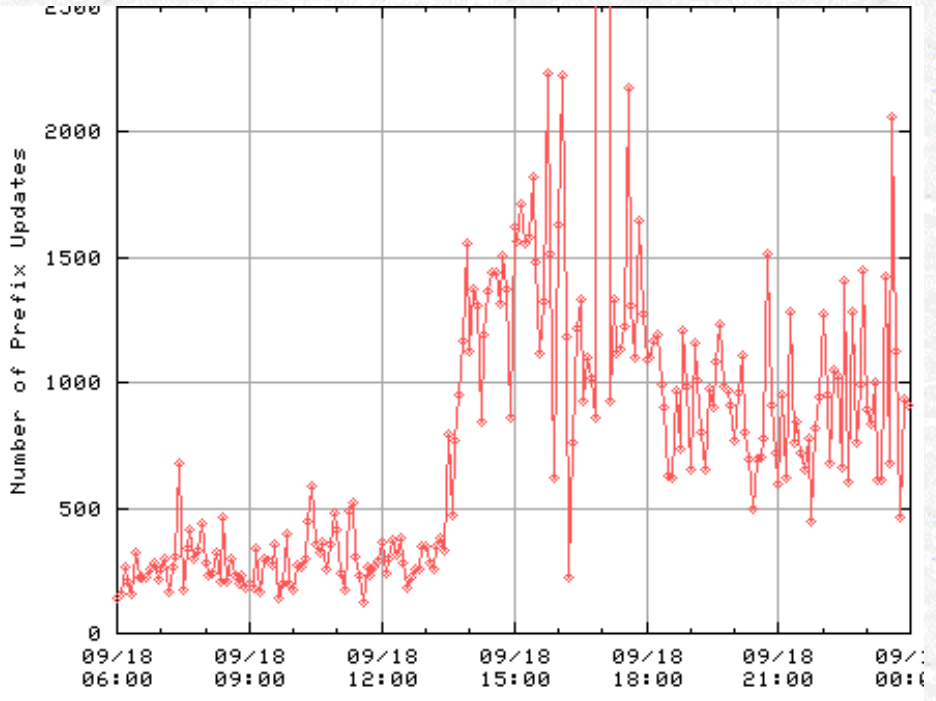

Possible explanation: the operational session between Surfnet and one of its peers went down temporarily, causing many routes to change their path and then immediately return to the original path

DPATH Distribution

AT&T on Sept. 18, 2001

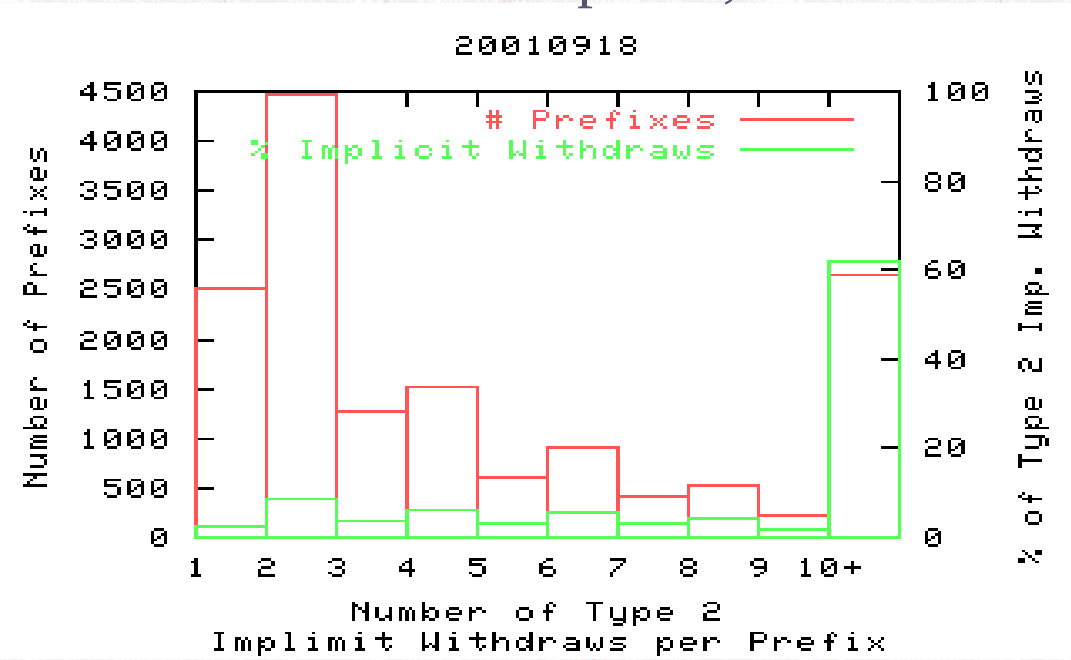

Watch the last bin: this group of highly unstable prefixes (~2500) contributed disproportionately (>60%) to the total;

Was every prefix unstable during the worm attack? No. Only 14.4% of the prefixes in AT&T's routing table had DPATH changes (similar for several other peers).

Conclusions

- Excessive session resets at the monitoring point contributed most to the spike.
  - a monitoring artifact
  - BGP does not work well over multi-hop connections;
- A substantial amount of BGP updates don't reflect AS-level topology changes.
- Worm attack may have seriously affected a small group of networks and the intermittent reachability to these networks propagated globally.

Related Work

- Labovitz et al studied BGP pathologies [2,3]
- Two experimental studies on BGP's behavior under congestion
  - Malan and Jahanian [4]: TCP failed to deliver keep-alive messages and BGP session broke
  - Shaikh et al [5]: the expected lifetime of BGP sessions decreases as congestion increases
- Maennel and Feldmann analyzed BGP traffic characteristics to produce realistic BGP traces [6]

References

[1] J. Cowie and A. Ogielski, “Global Routing Instabilities During Code Red II and Nimda Worm Propagation”, NANOG 23

[2] C. Labovitz, et al, “Internet Routing Instability,” SIGCOMM ’97

[3] C. Labovitz, et al, “Origins of Internet Routing Instability,” INFOCOM ’99

[4] G. R Malan and F. Jahanian, “An Extensible Probe Architecture for Network Protocol Performance Measurement,” SIGCOMM ’98

[5] A. Shaikh, et al, “Routing Stability in Congested Networks: Experimentation and Analysis,” SIGCOMM 2000

[6] O. Maennel and A. Feldmann, “Realistic BGP Traffic for Test Labs”, SIGCOMM 2002

BGP Session Resets from Sept. 01, 2001 to Sept. 30, 2001