Global Routing Instabilities

during Code Red 2 and Nimda Worm Propagation

Jim Cowie and Andy Ogielski
Renesys Corporation

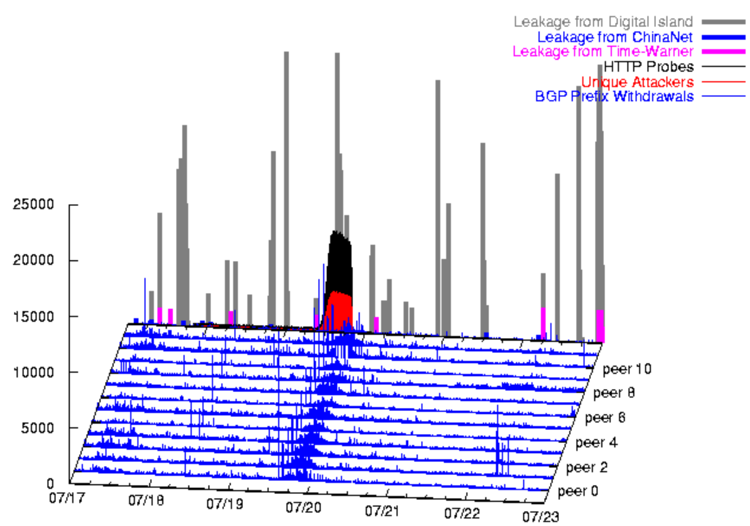

23 October 2001

renesys

Outline

Catastrophic instabilities are an expected behavior of large engineered systems (John Doyle, Caltech)

1. **Define** global routing instability
2. **Analyze** raw BGP message traffic from 150 peers (all RIPE RRCs).
3. **Paint** a picture of instabilities caused by:

   - Microsoft worms

   - router misconfigurations

   - …..?

renesys

Focus: RIPE rrc00 collection point

EBGP peers from around the world

AS		peer IP
13129	Global Access	212.20.151.253
1103	SURFnet	193.148.15.34
513	CERN	192.65.184.3
3333	RIPE NCC	193.0.0.56
286	KPN Qwest	134.222.87.12
4777	APNIC Tokyo Servers	202.12.28.190
9177	Nextra	212.47.190.1
4608	Telstra	203.37.255.126
3257	Tiscali	193.148.15.85
3549	Global Crossing UK	195.66.224.112
3549	Global Crossing USA	206.251.0.85
2914	Verio	129.250.0.232
7018	AT&T Internet4	12.127.0.121

renesys

BGP message traffic rate

received by a single BGP router from 12 major peers.

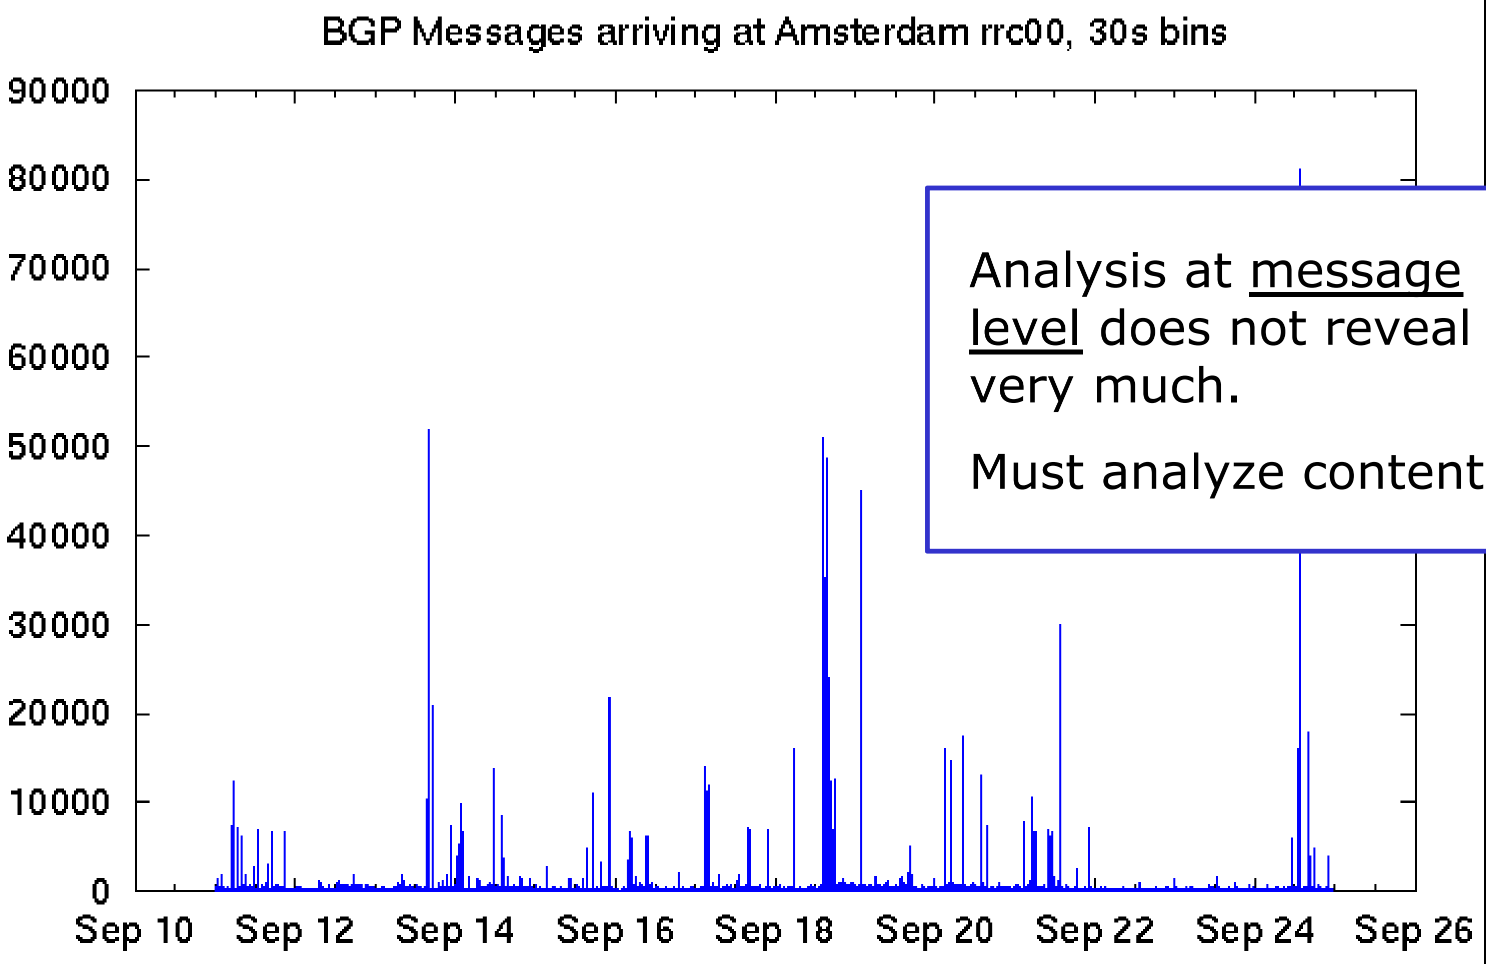

Analysis at <u>message level</u> does not reveal very much.

Must analyze content.

renesys

A view on content of the same messages

Number of prefix announcements in 30 sec intervals

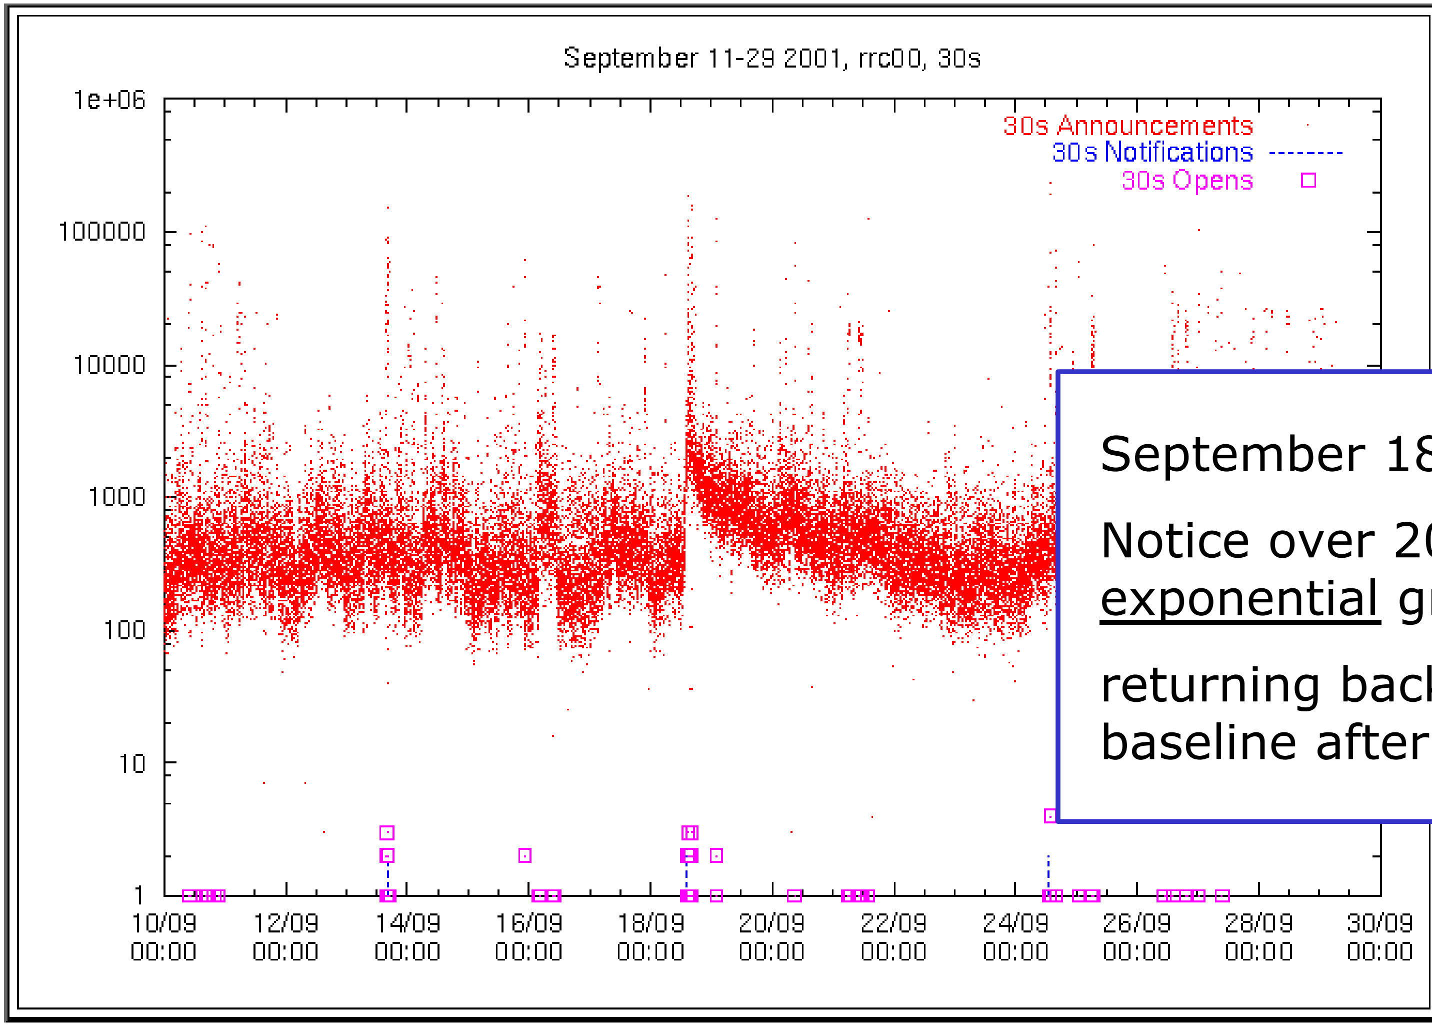

September 18:

Notice over 20-fold exponential growth

returning back to baseline after 4 days!

renesys

Analysis exposes correlations:

Behavior across…

- peers
- peering points
- origin ASs
- prefixes
- prefix length
- route lifetimes

renesys

Prefix announcements by peer

RIPE NCC, September 10 - 22, 15-min intervals

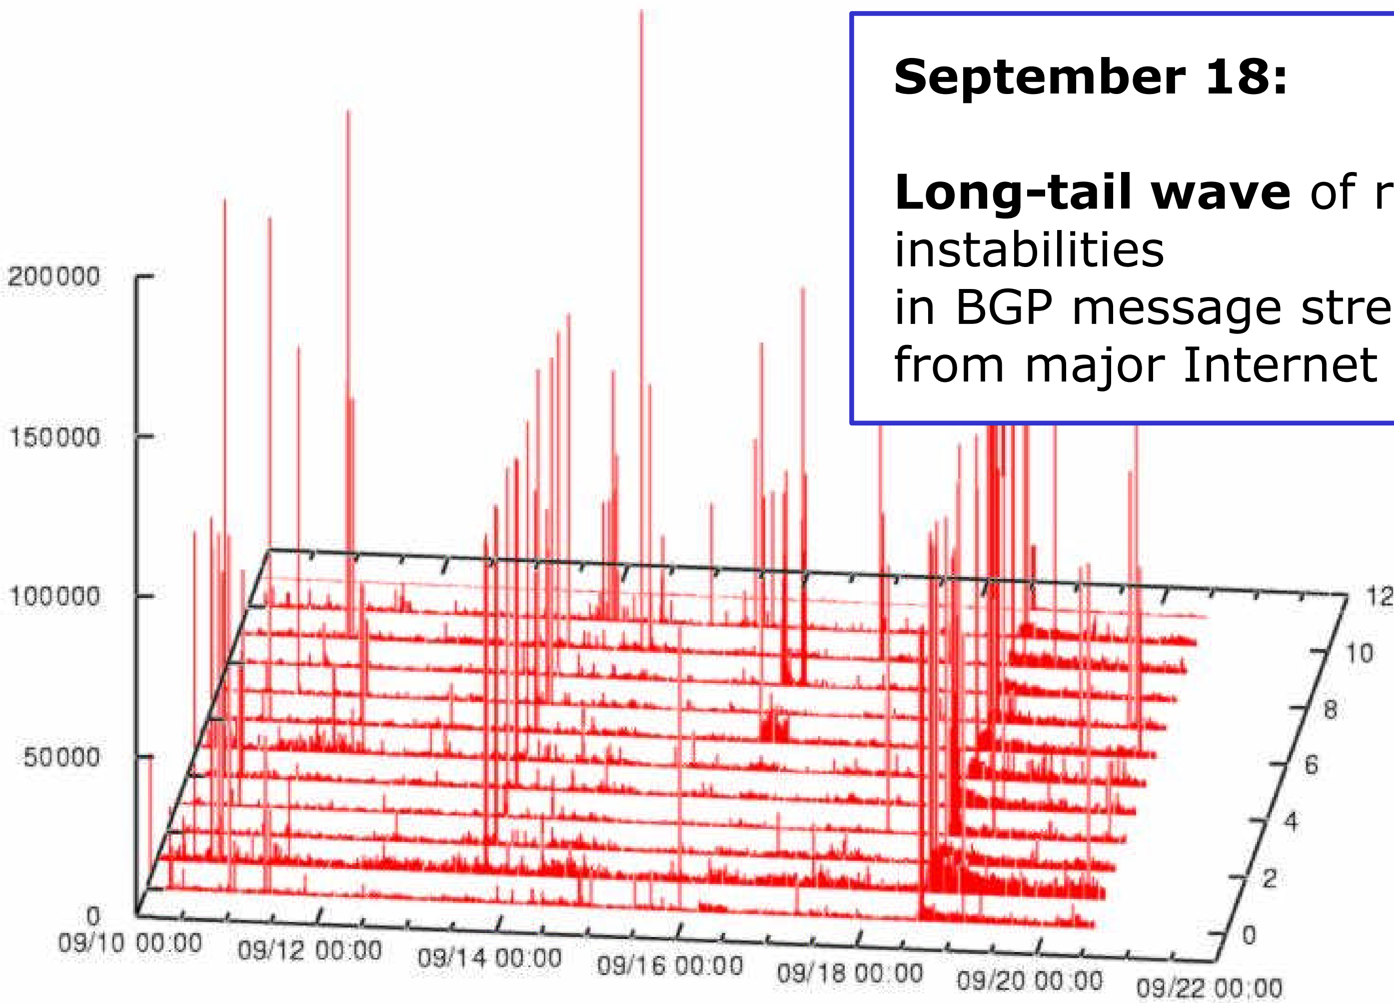

September 18:

Long-tail wave of routing instabilities
in BGP message streams
from major Internet providers

renesys

September 18 Nimda worm attack

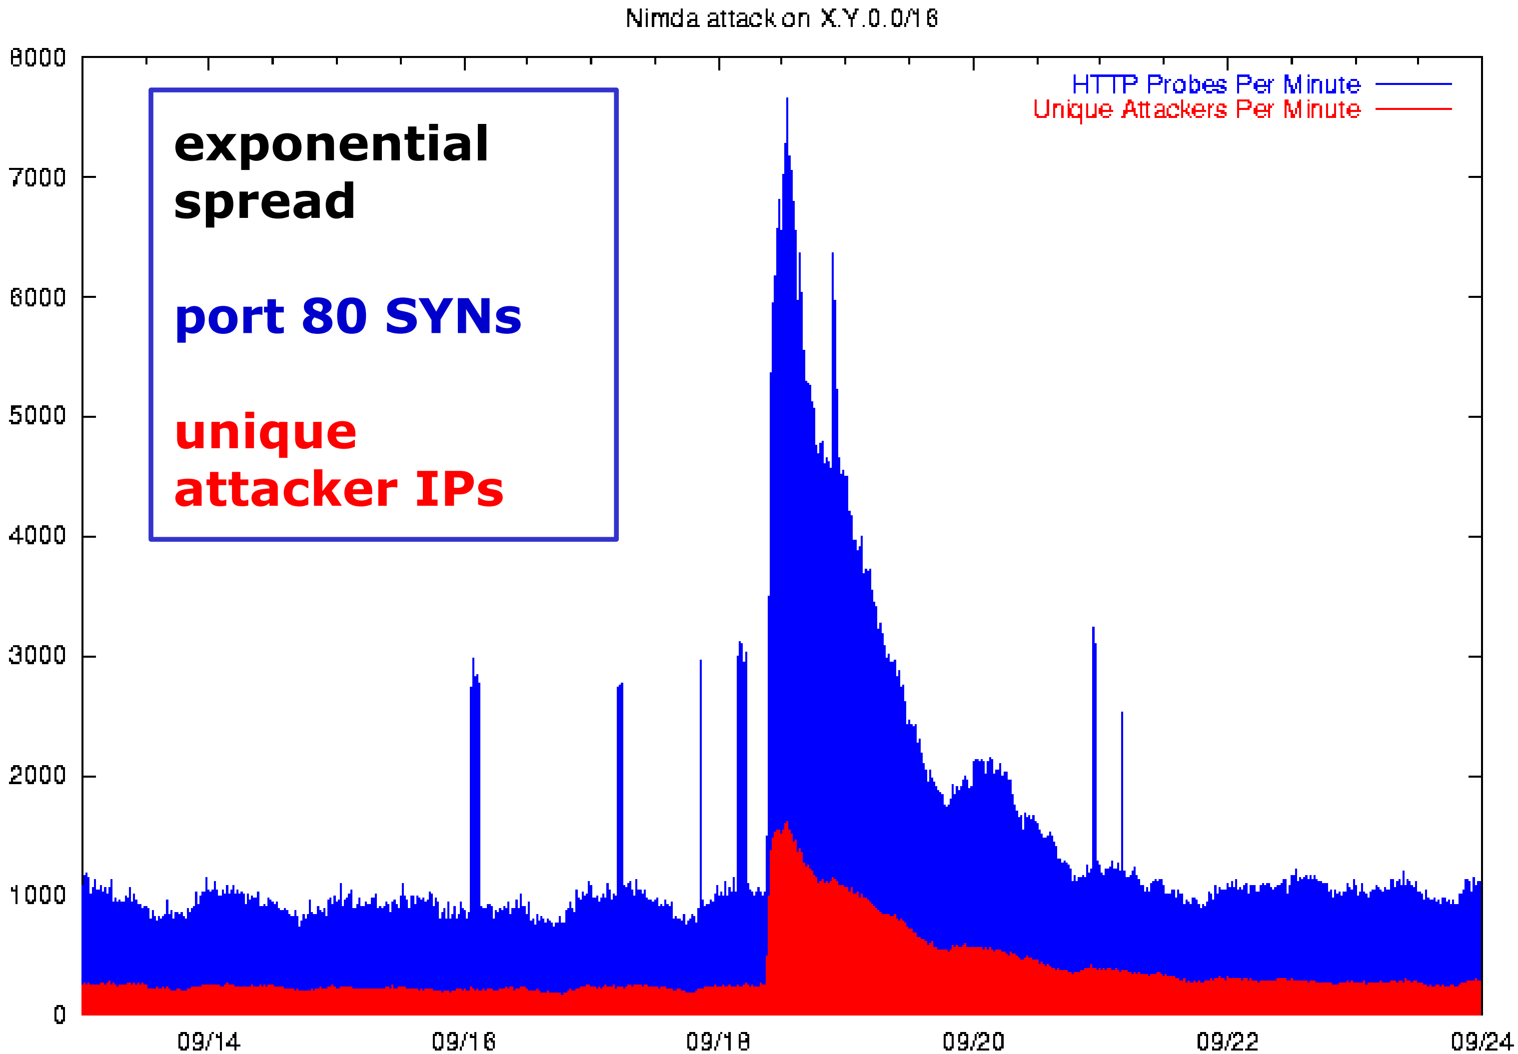

renesys

Nimda probes **burn routers' CPU cycles…**

Inset plot shows highly correlated router cpu utilization … *in a different net*

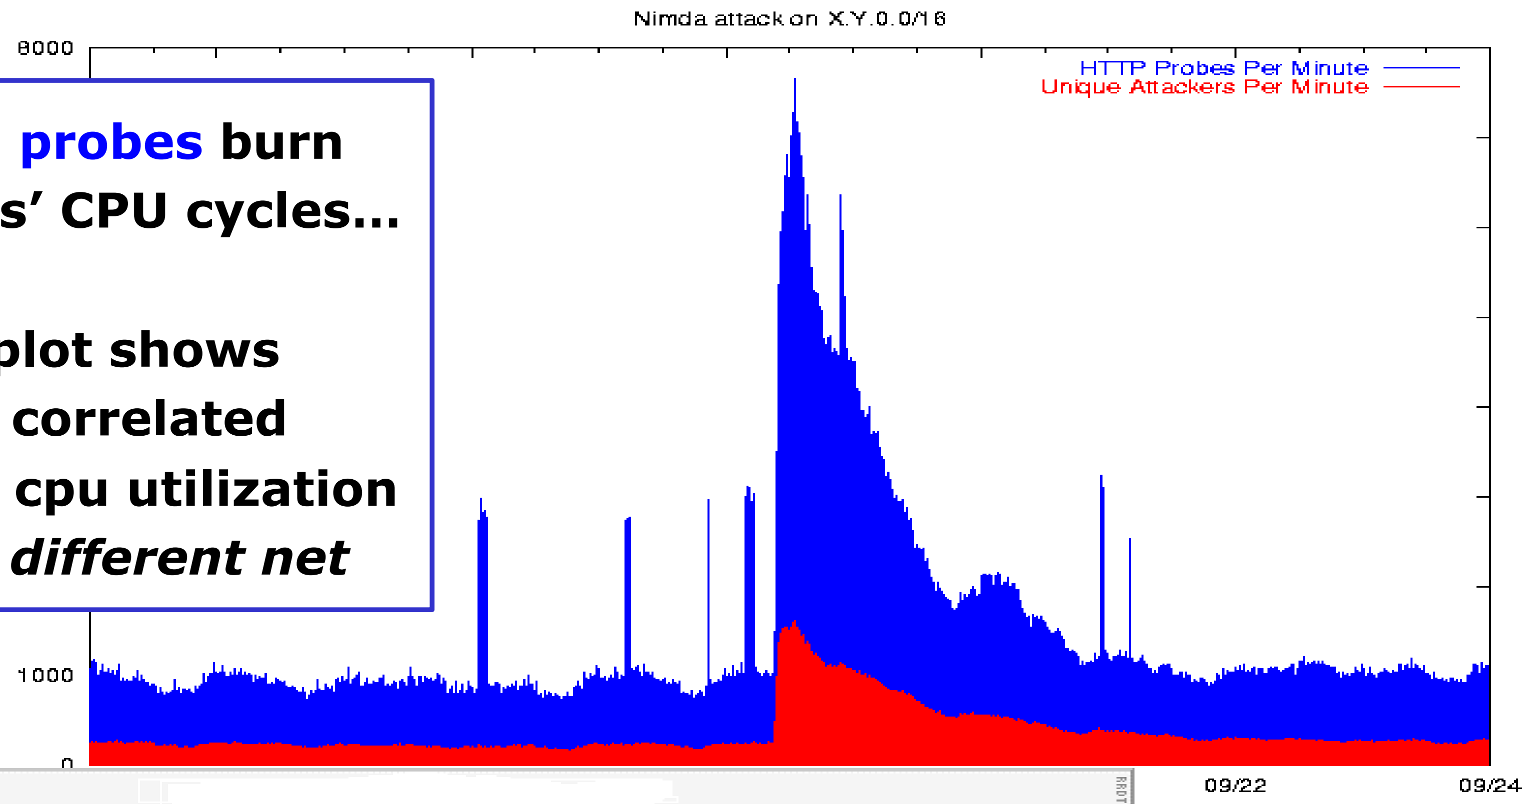

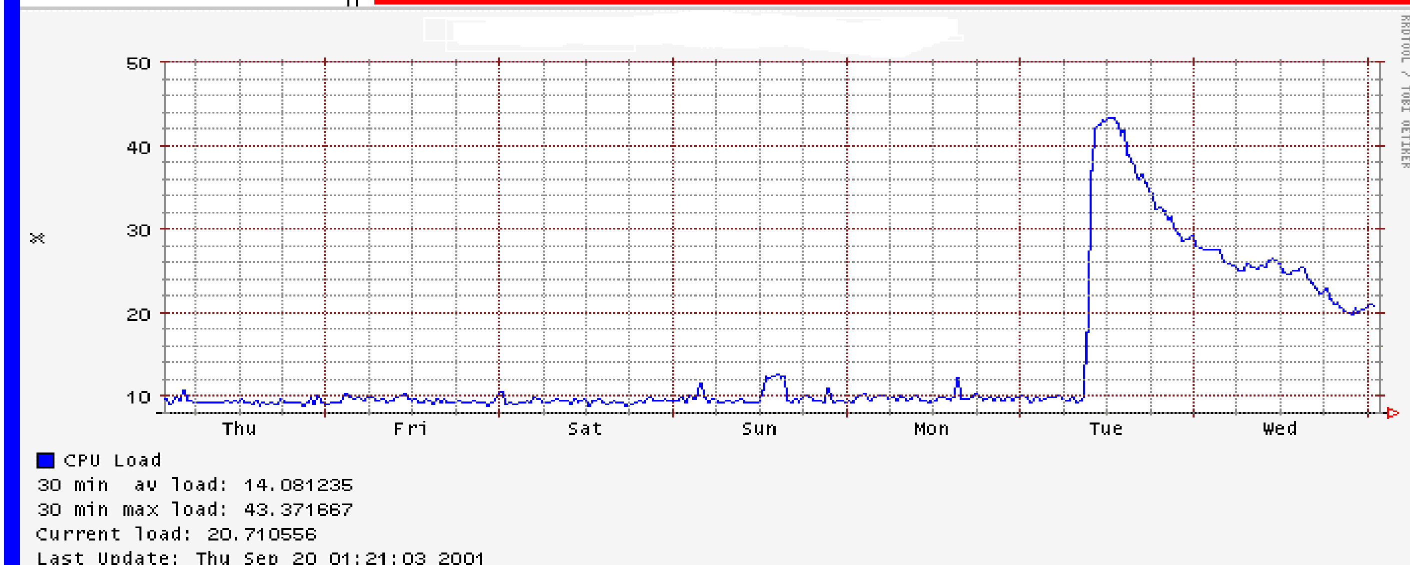

renesys

September 18 BGP event correlates in time with Nimda worm attack

Smaller events: leakage of reserved AS numbers

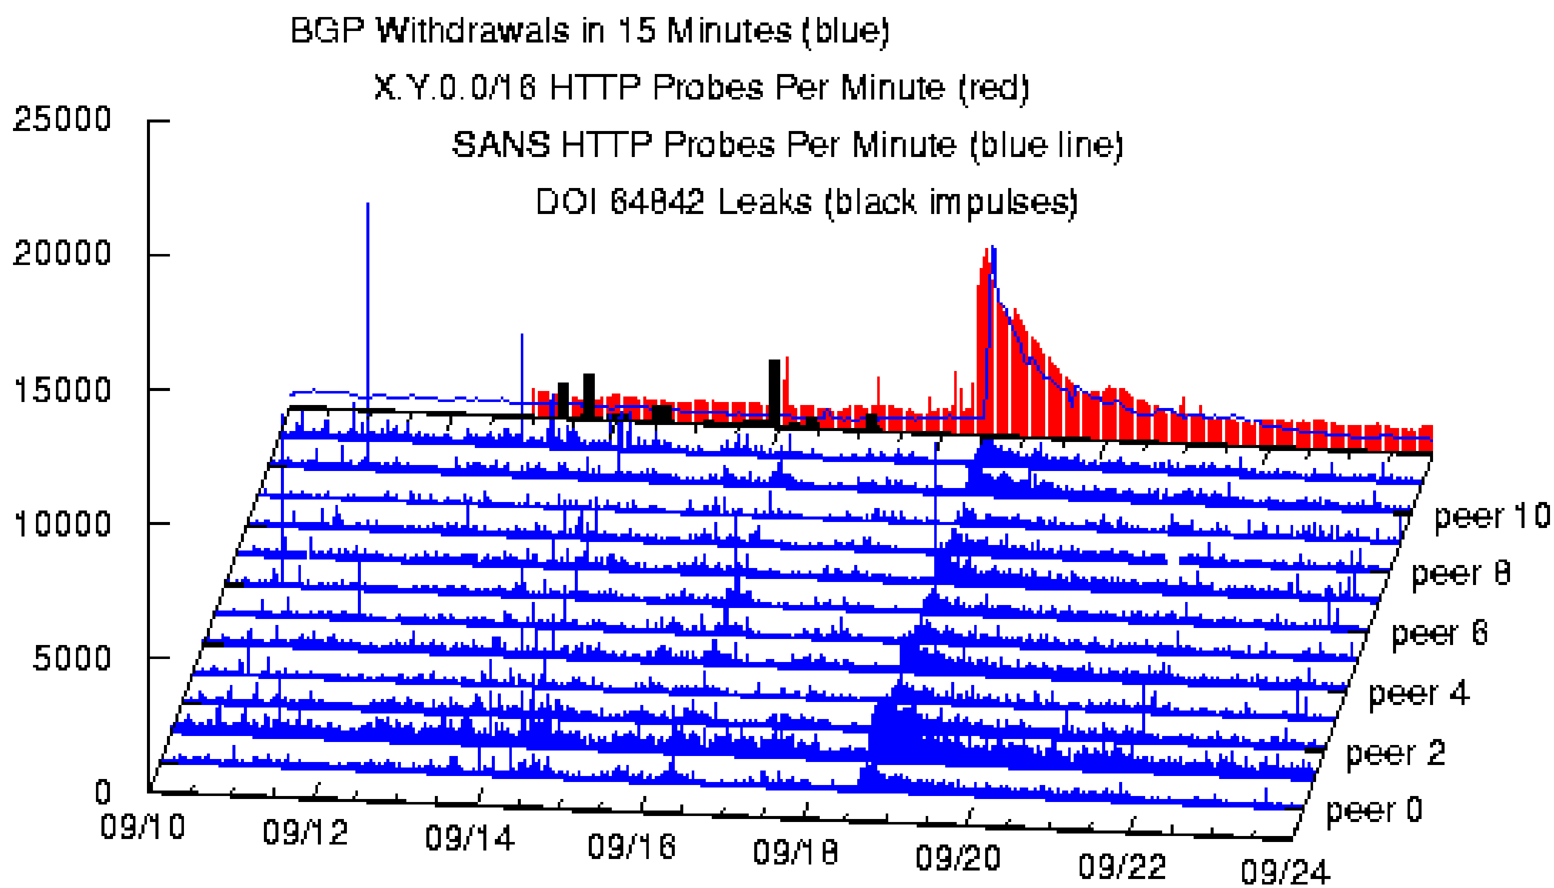

renesys

Global Internet Routing Instabilities

Qualitative definition

rate	duration	diversity
High rates of route changes: • magnitude • acceleration • variance	Very long times: • long relative to baseline noise • long relative to expected routing table convergence time	Seen at many observation points: • many external BGP peers • many exchanges • Intra-AS networks Seen in high diversity of routing traffic content: • number of prefixes • number of routes

renesys

One peer unstable: not a global instability

October 20 rrc00 announcements– AS 1103 unstable

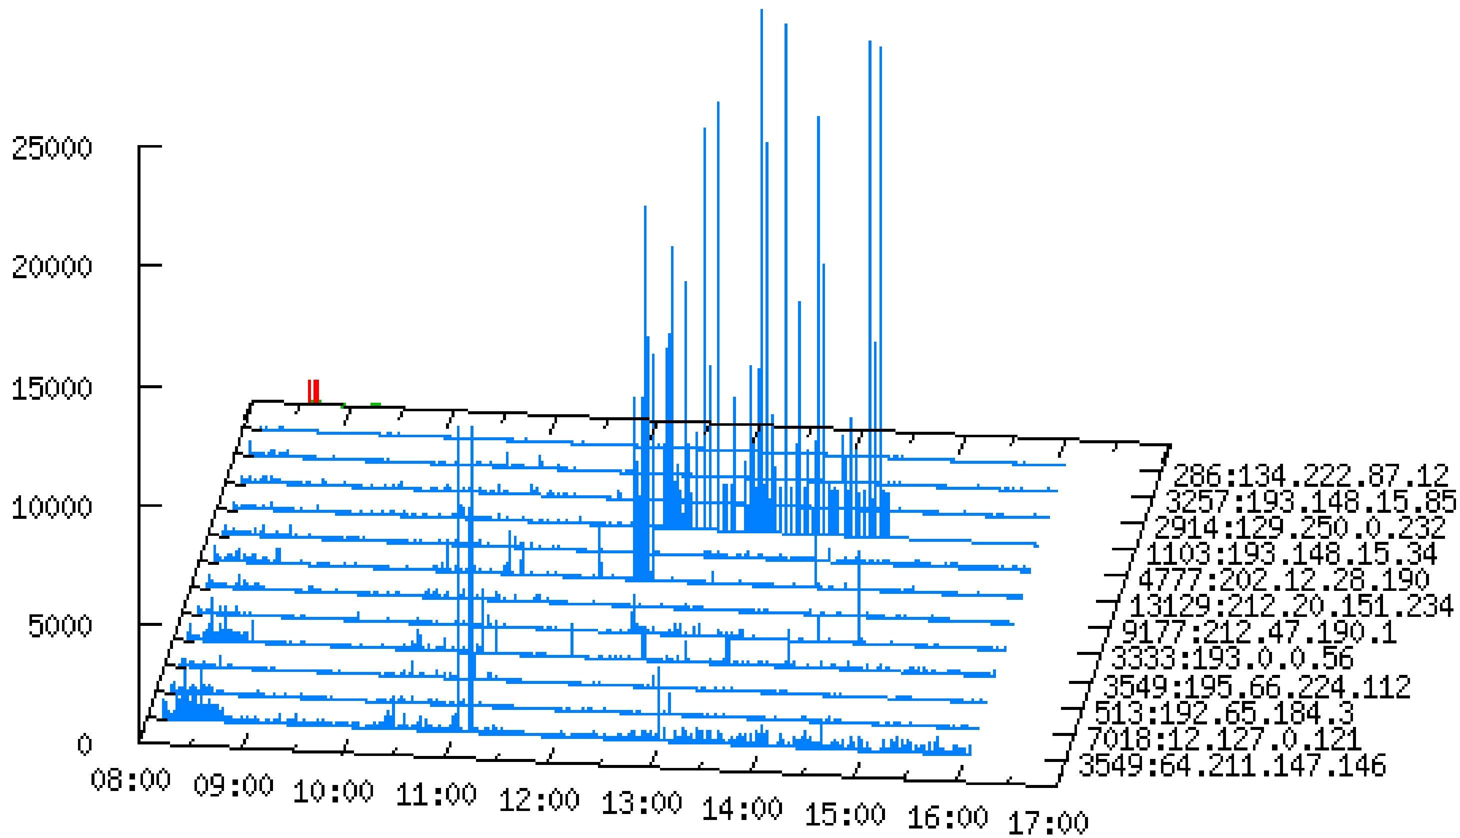

renesys

Global Internet Routing Instabilities

operational definition

rate	duration	diversity
Exponential growth of rate of prefix announcements and withdrawals	**Hours to days**	• almost all prefixes churning • from most large ISP peers

renesys

Worm story # 2:

Code Red v2 attack

prefix announcement rate in 30 sec intervals

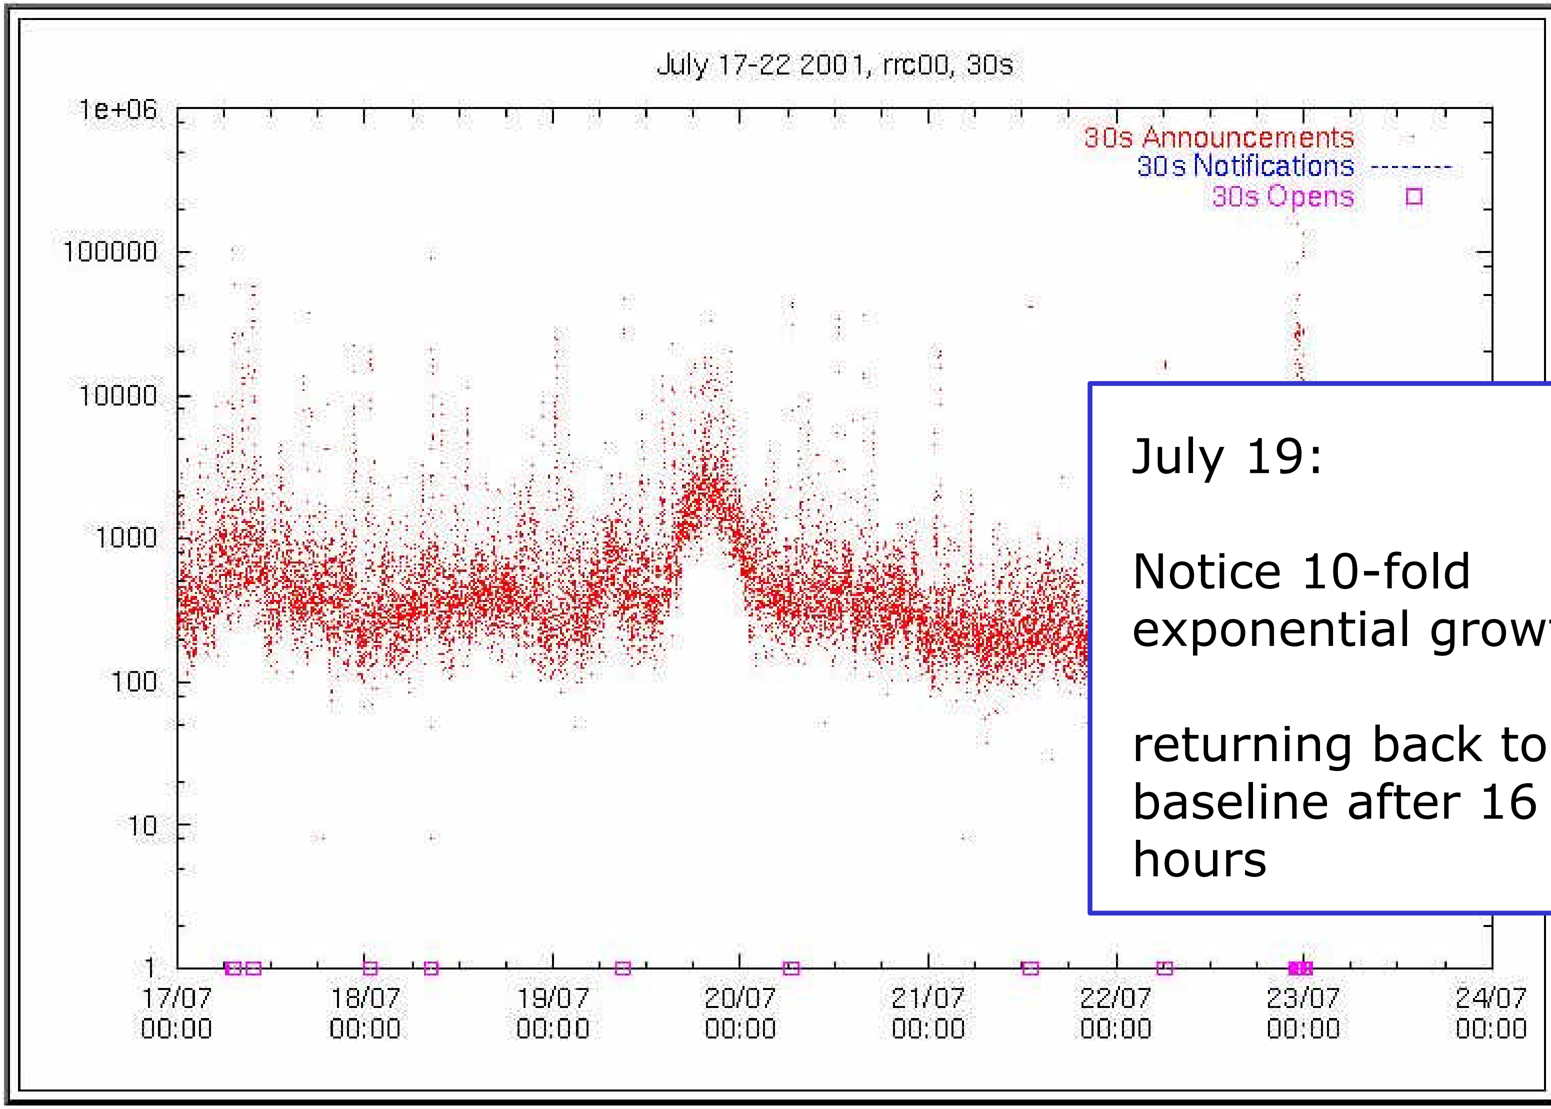

July 19:

Notice 10-fold exponential growth

returning back to baseline after 16 hours

renesys

July 19 Code Red II worm attack

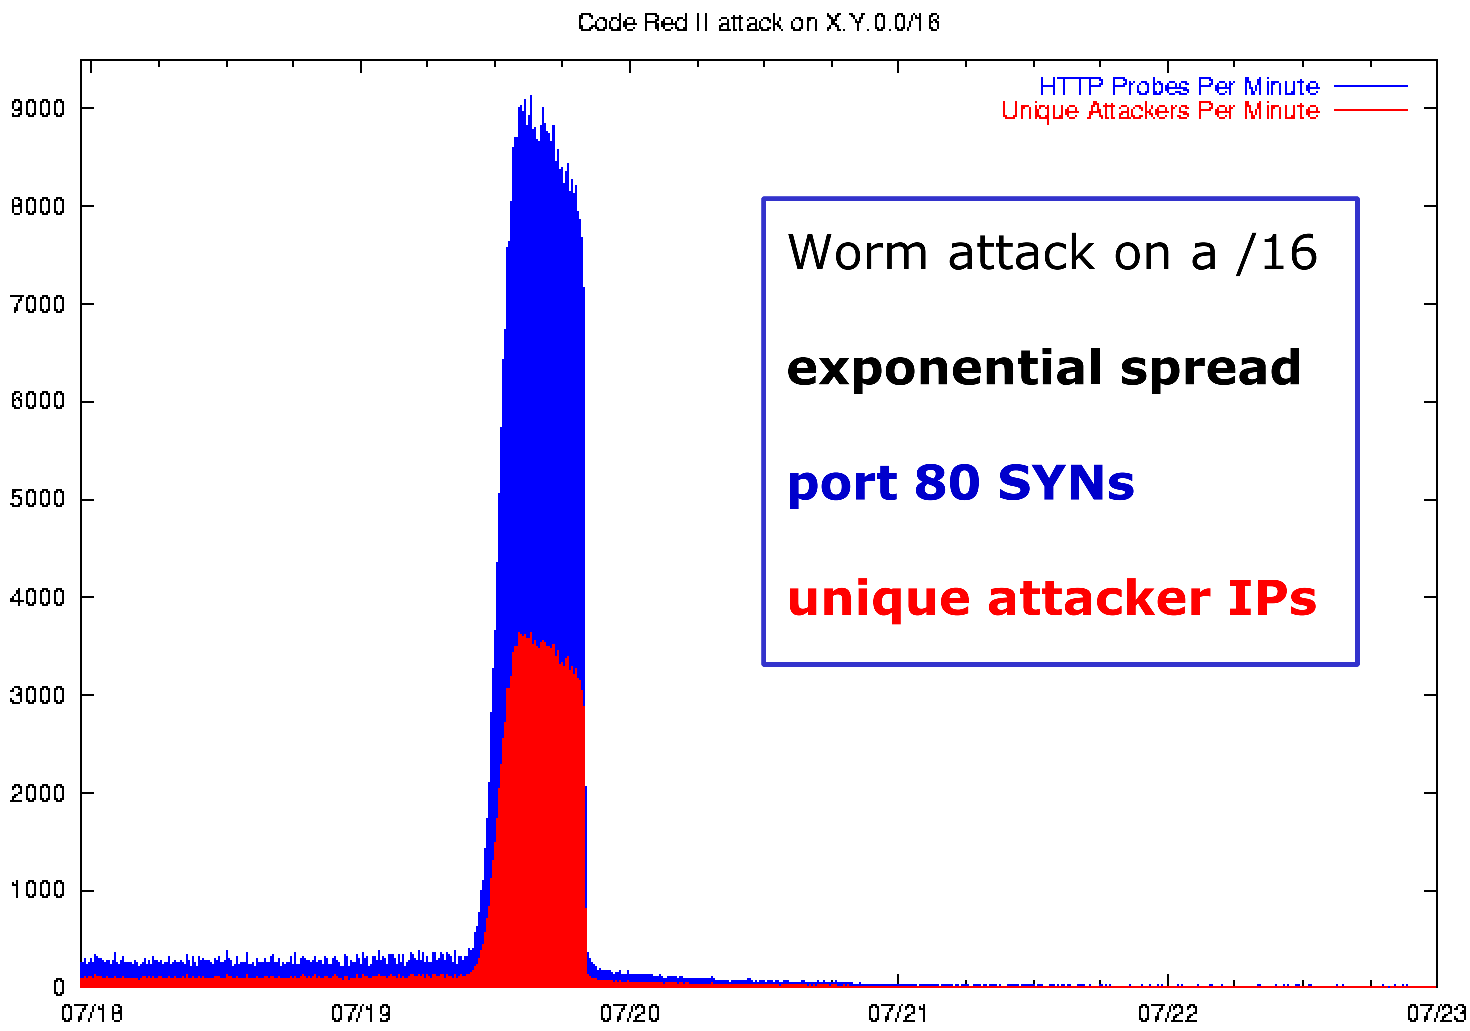

Worm attack on a /16

exponential spread

port 80 SYNs

unique attacker IPs

renesys

July 19 BGP storm correlates in time with Code Red II worm attack

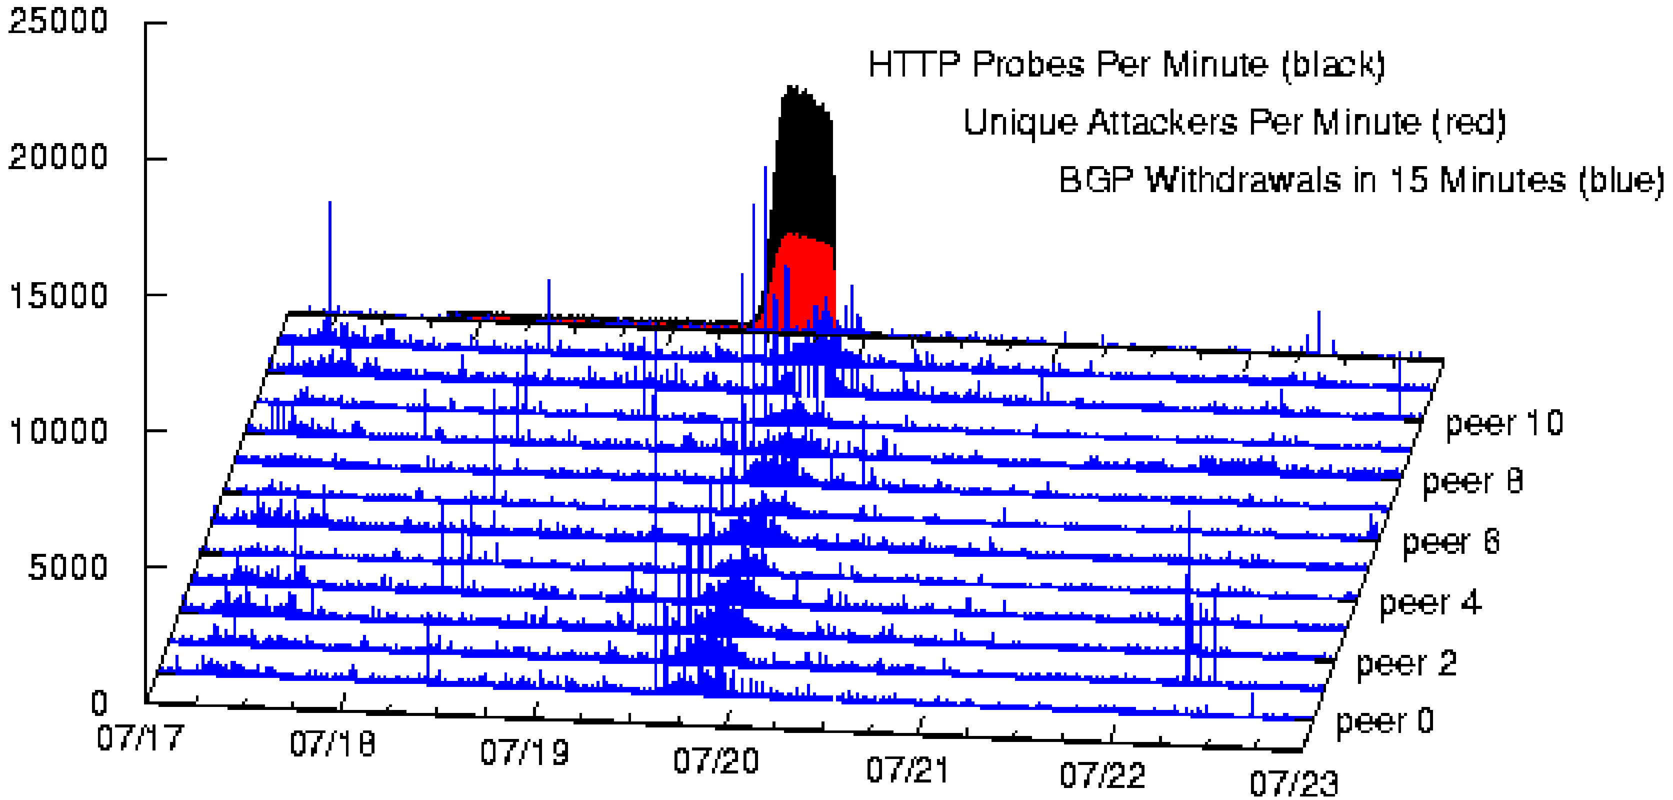

renesys

Results of detailed analysis:

Nimda and Code Red triggered long-term BGP instabilities unlike any localized network failure:

- **no suspect peers** – all major peers
- **no suspect prefixes** – most prefixes churn
- **no suspect routes** – most routes churn

renesys

worm-induced BGP instabilities

Do not look like effects of link failures between multiply-connected major Internet providers (Internet core).

Cable cuts, Baltimore tunnel fire, September 11 **did not** create global instabilities.

Cable cuts between core providers affect route changes that are **localized** between affected providers.

Worm-induced BGP events seem to arise from BGP connectivity failures at very many locations: edge?

Possible causes of BGP session failures

Why BGP routers can fail:

router CPU overload

router out of memory, cache overflows

router software bugs

Possible worm traffic causes: **thanks for emails!**

traffic intensity

traffic diversity (# flows)

HTTP servers in routers (mngmt interfaces)

failures in network gear (DSL routers,…)

IGP (Intra-AS) flapping and routing failures

proactive disconnection of networks

renesys

Preliminary analysis - summary

Worm traffic diversity causes: most likely?

extreme scan rate -> extremely many flows -> router CPU/memory, NAT problems, ARP storms.

Routing traffic causes: likely?

-- extremely high rate of BGP updates – router CPU/memory

Worm traffic intensity causes:

-- loss of BGP messages (presumably at the edge)

congestion unlikely

renesys

Misconfiguration instabilities

common BGP events in the Internet core

0. Misconfigured AS starts announcing a private (confederation) ASpath:

```
%BGP-6-ASPATH: Invalid AS path xxx 3300 (64603) 2008 received
    from x.x.x.x: Confederation AS-path found in the middle
```

1. *Certain* routers **ignore** but **propagate** the malformed route
2. Other, RFC-compliant routers **close** & **reopen** the BGP sessions.
3. The combination may propagate wildly
4. Instability ends only when the original leak is plugged.

"… we have the stick now. unfortunately, we also have a vendor who ignores sticks."

(Randy Bush)

Smaller BGP events: cascading router failures

Initiated by local leakage of malformed route announcements (ASPATH)

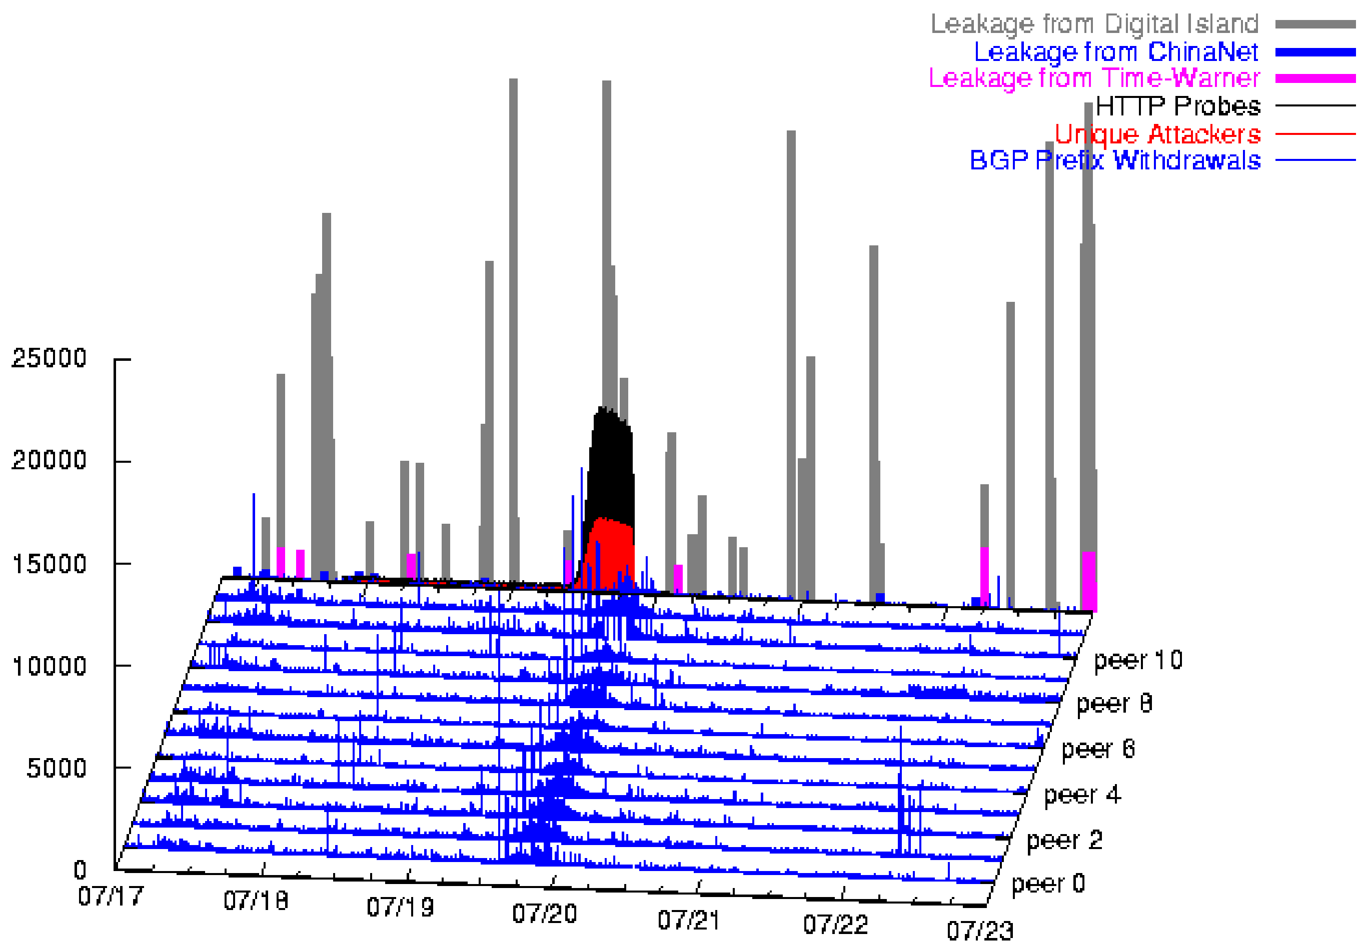

renesys

June 20 – 30 BGP instabilities

Responsible ASPATH segment: **523** - **64520** - **721**

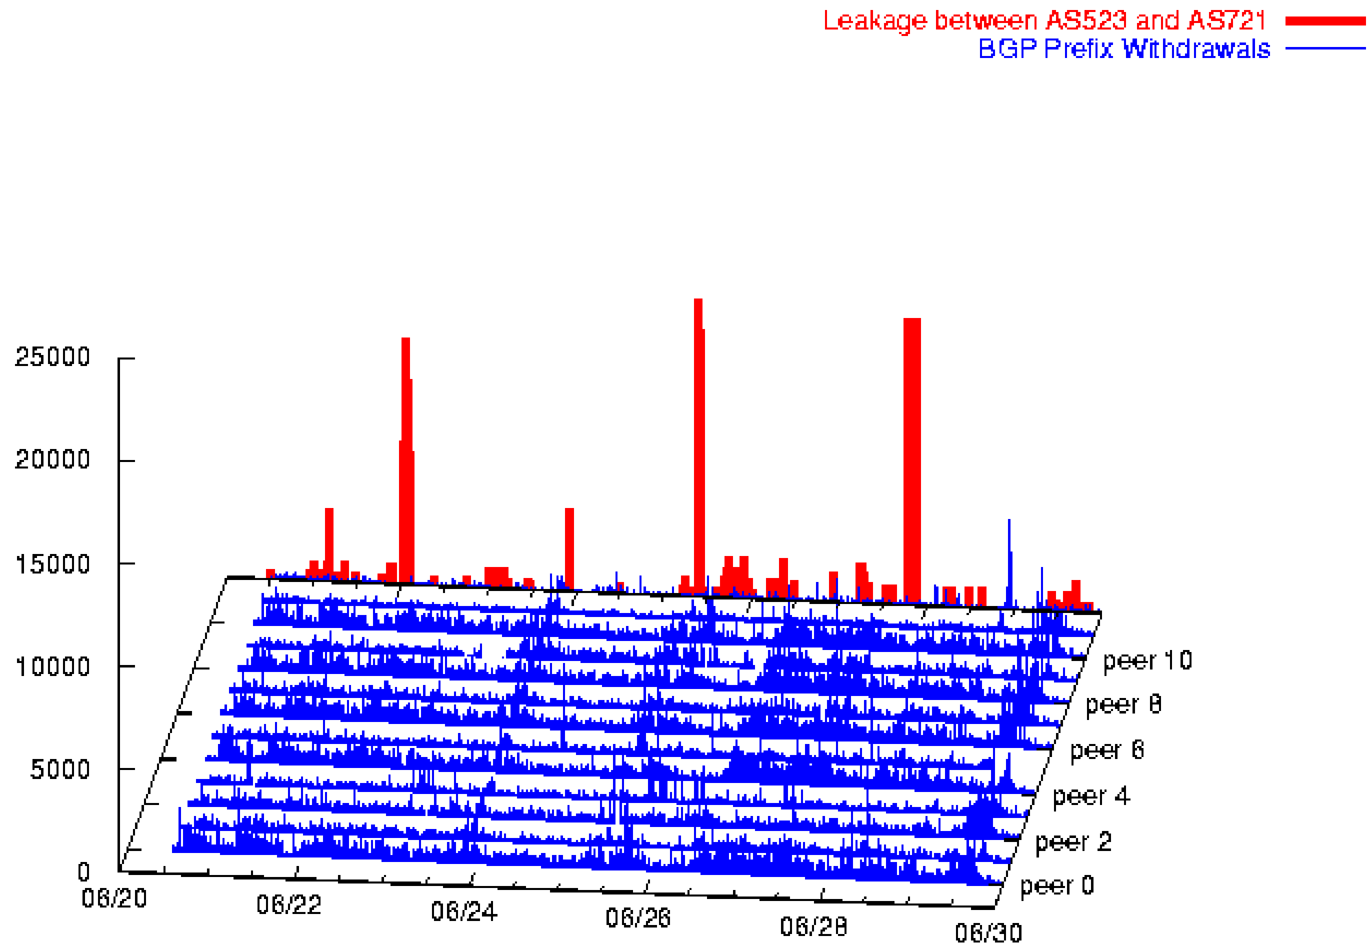

renesys

October 6 - 15 BGP instabilities, rrc00

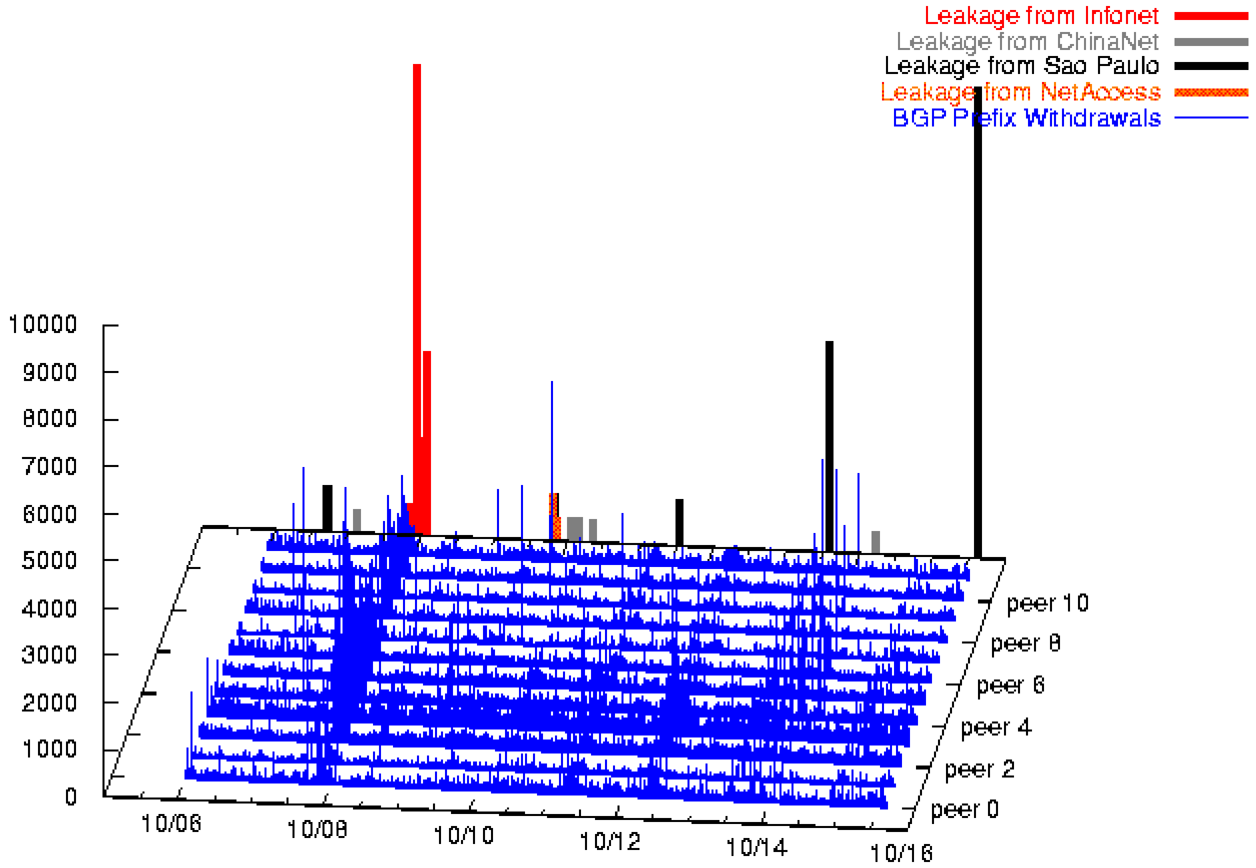

renesys

We barely scratched the surface…

1. Globally correlated BGP instabilities are common
2. Some causes are understood a bit – ASPATH oddities
3. Others are unexpected & disturbing (Microsoft worms)

renesys

Credits

- Early analysis with BJ Premore and Yougu Yuan at Renesys.

- Raw BGP msg data courtesy of RIPE RIS. Special thanks to Henk Uijterwaal (RIPE).

- Worm traffic data from several /16 networks courtesy of Vicki Irwin (SANS Institute), Ken Eichman (CAS), Vern Paxson (ACIRI).

- Thanks to many network operators and administrators for detailed case stories and observations on *Major Vendors'* router misbehaviors.

- Thanks to Tim Griffin (AT&T) and Dave Donoho (Stanford) for discussions.

renesys

Andy

ato@renesys.com

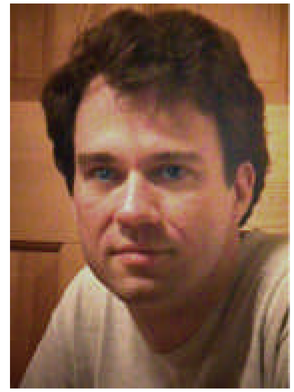

Jim

cowie@renesys.com

...we want to talk to anyone interested in contributing more multi-hop EBGP feeds for research ... *silent peering*.

Forget the chocolates and tee-shirts...

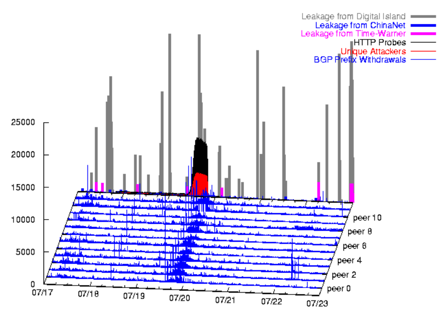

... we trade raw data for global instability alerts

renesys