


IPv6 Issues on the AMS-IX Peering LAN

Ariën Vijn
arien.vijn@ams-ix.net



Ether Type Distribution - daily

RRDTOOL / TOBI OETIKER




	Current	Average	Maximum	Minimum
other	0.0%	0.0%	0.0%	0.0%
ARP	0.0%	0.0%	0.0%	0.0%
IPv6	0.2%	0.3%	0.5%	0.2%
IPv4	99.8%	99.7%	99.8%	99.5%

Copyright (c) 2011 AMS-IX B.V.

Updated: Mon May 2 15:00:01 2011 CET


Ether Type Distribution - daily



	Current	Average	Maximum	Minimum
other	0.0%	0.0%	0.0%	0.0%
ARP	0.0%	0.0%	0.0%	0.0%
IPv6	0.2%	0.3%	0.5%	0.2%
IPv4	99.8%	99.7%	99.8%	99.5%

Ether Type Distribution - daily

RRDTOOL / TOBI OETIKER




	Sun 00:00	Sun 12:00	Sun 24:00	Mon 00:00
other	0.0%	0.0%	0.0%	0.0%
ARP	0.0%	0.0%	0.0%	0.0%
IPv6	0.2%	0.3%	0.5%	0.2%
IPv4	99.8%	99.7%	99.8%	99.5%

Copyright (c) 2011 AMS-IX B.V.

Updated: Mon May 2 15:00:01 2011 CET

Total IPv6 Traffic - daily

RRDTOOL / TOBI OETIKER



■ in


Cur = 1.8 Gbps
Avg = 2.1 Gbps
Max = 2.9 Gbps
Min = 1.2 Gbps

Copyright (c) 2011 AMS-IX B.V.

Updated: Mon May 2 15:15:01 2011 CET

Total IPv6 Traffic - yearly

RRDTOOL / TOBI OETIKER



■ in

Cur = 1.7 Gbps
Avg = 1.6 Gbps
Max = 2.6 Gbps
Min = 888.4 Mbps


Copyright (c) 2011 AMS-IX B.V.

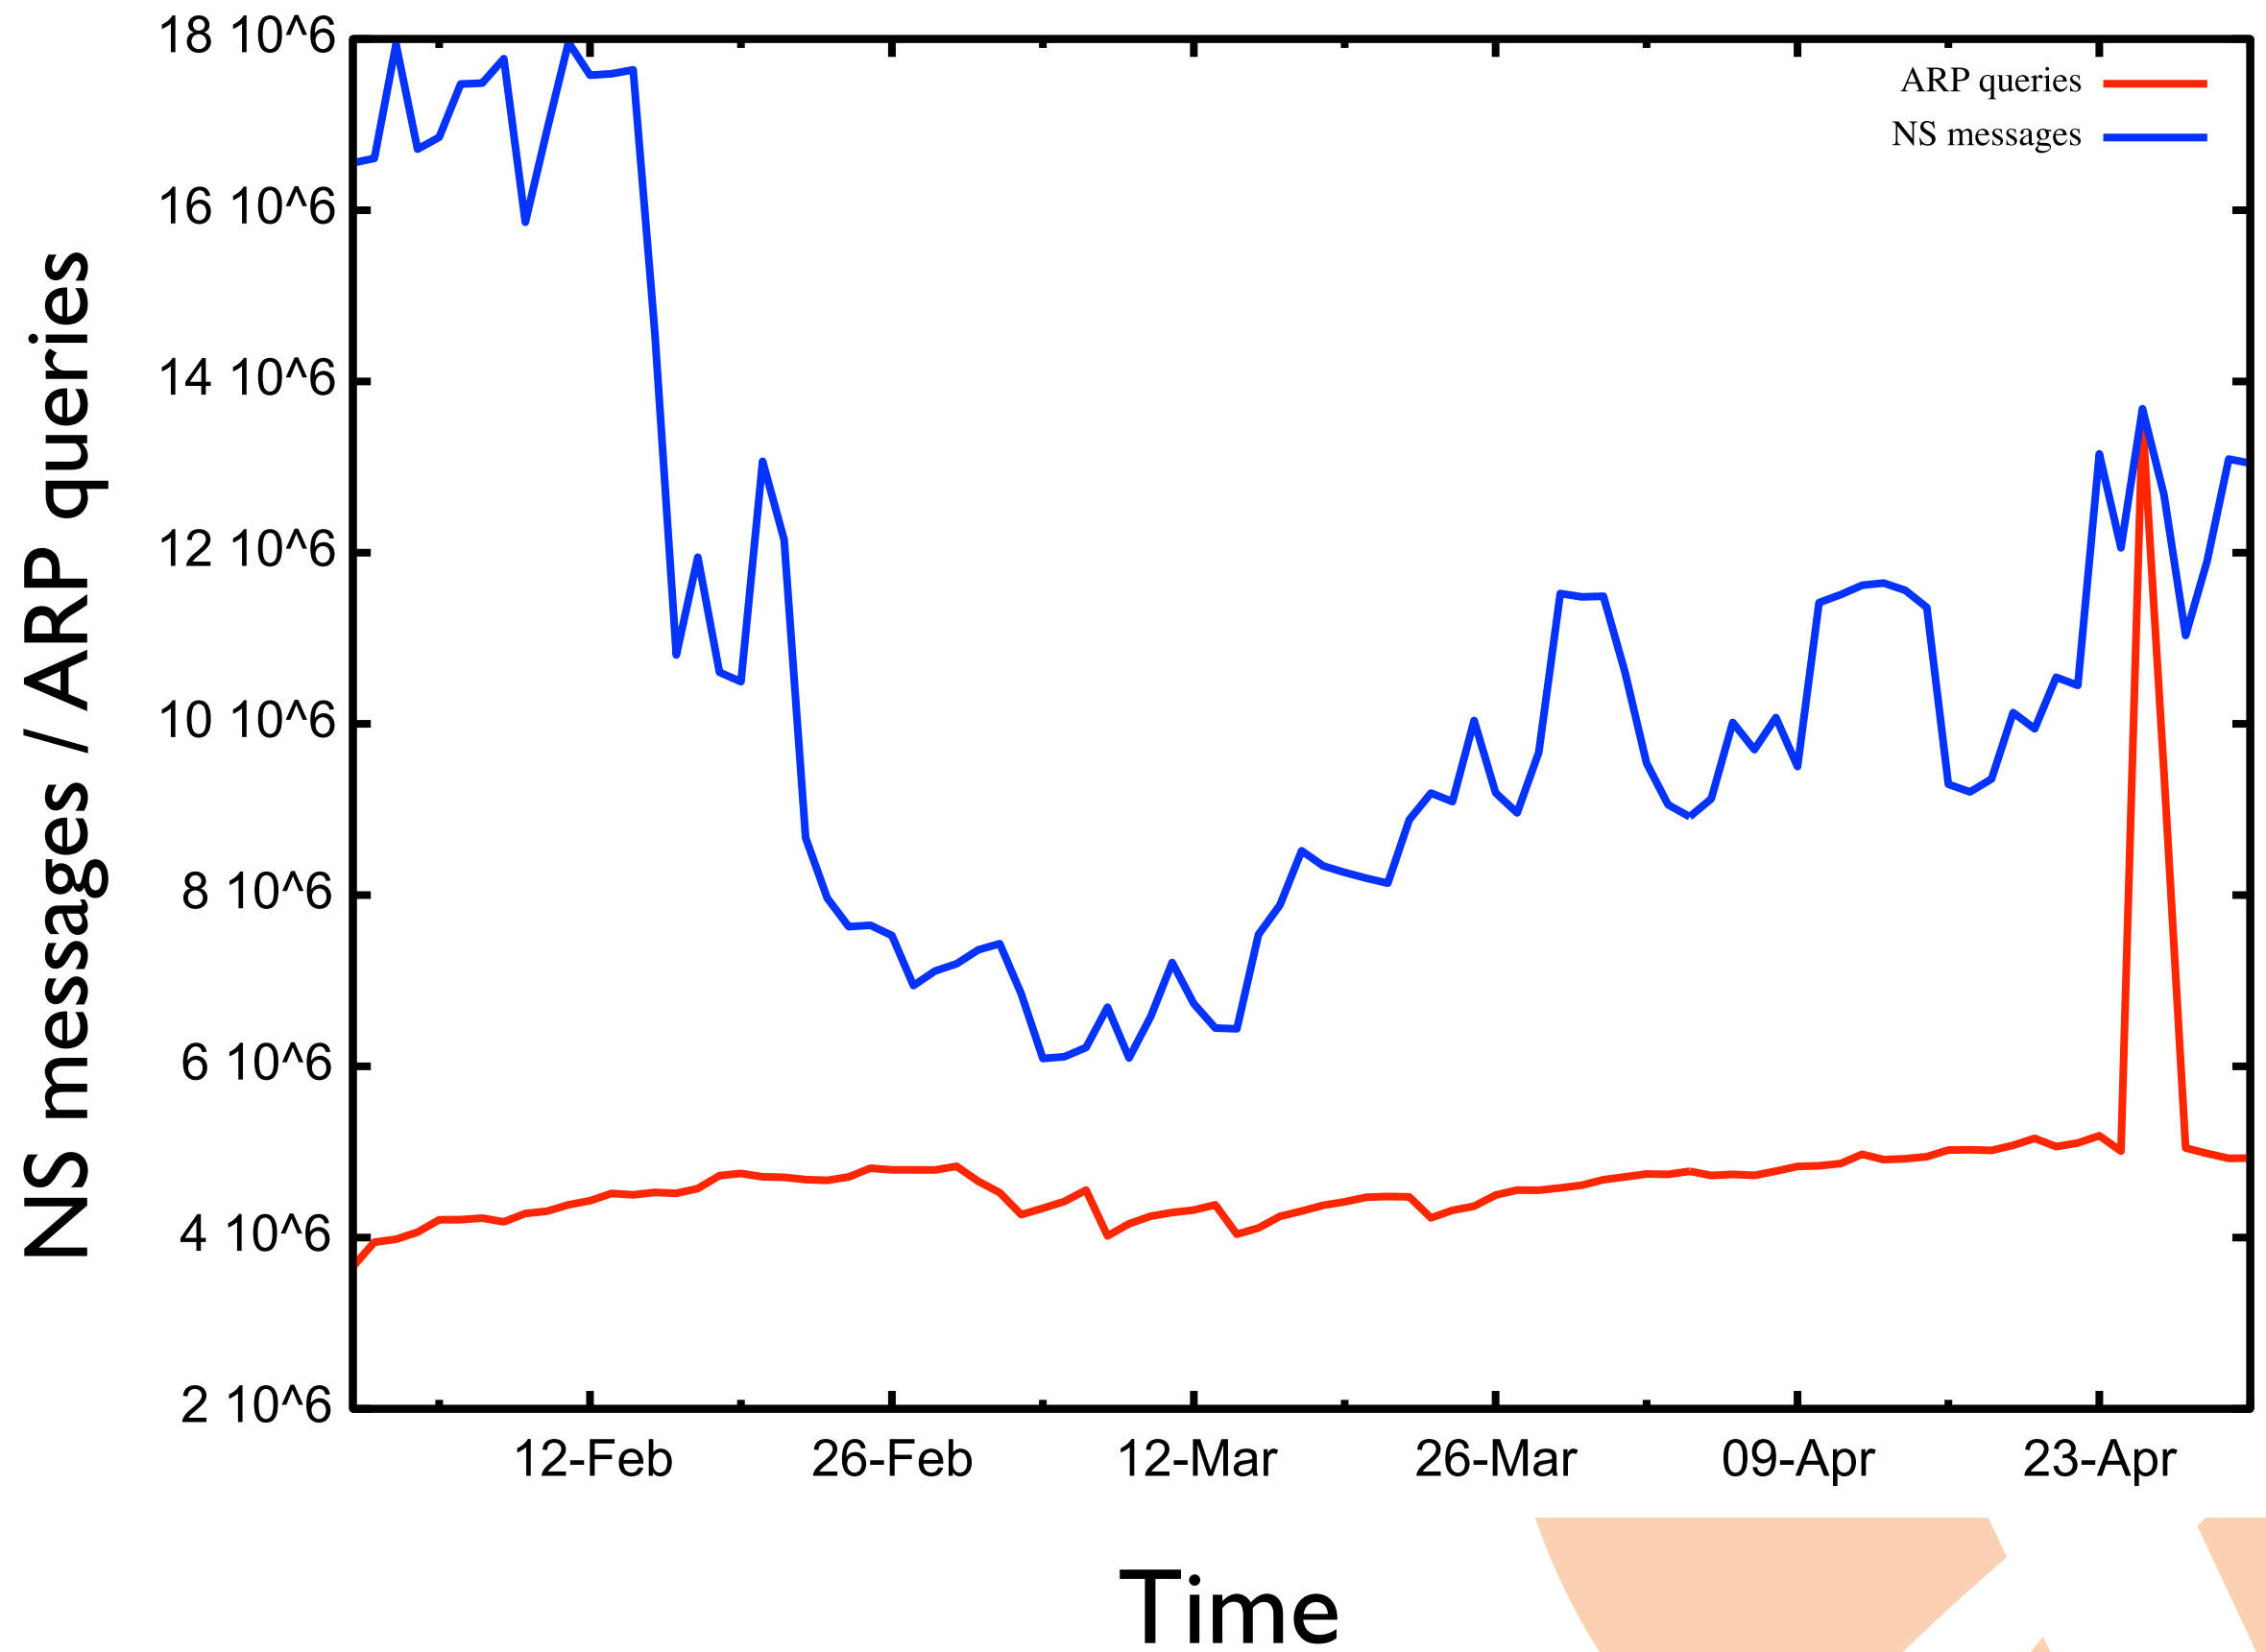
Updated: Mon May 2 15:15:01 2011 CET

Agenda

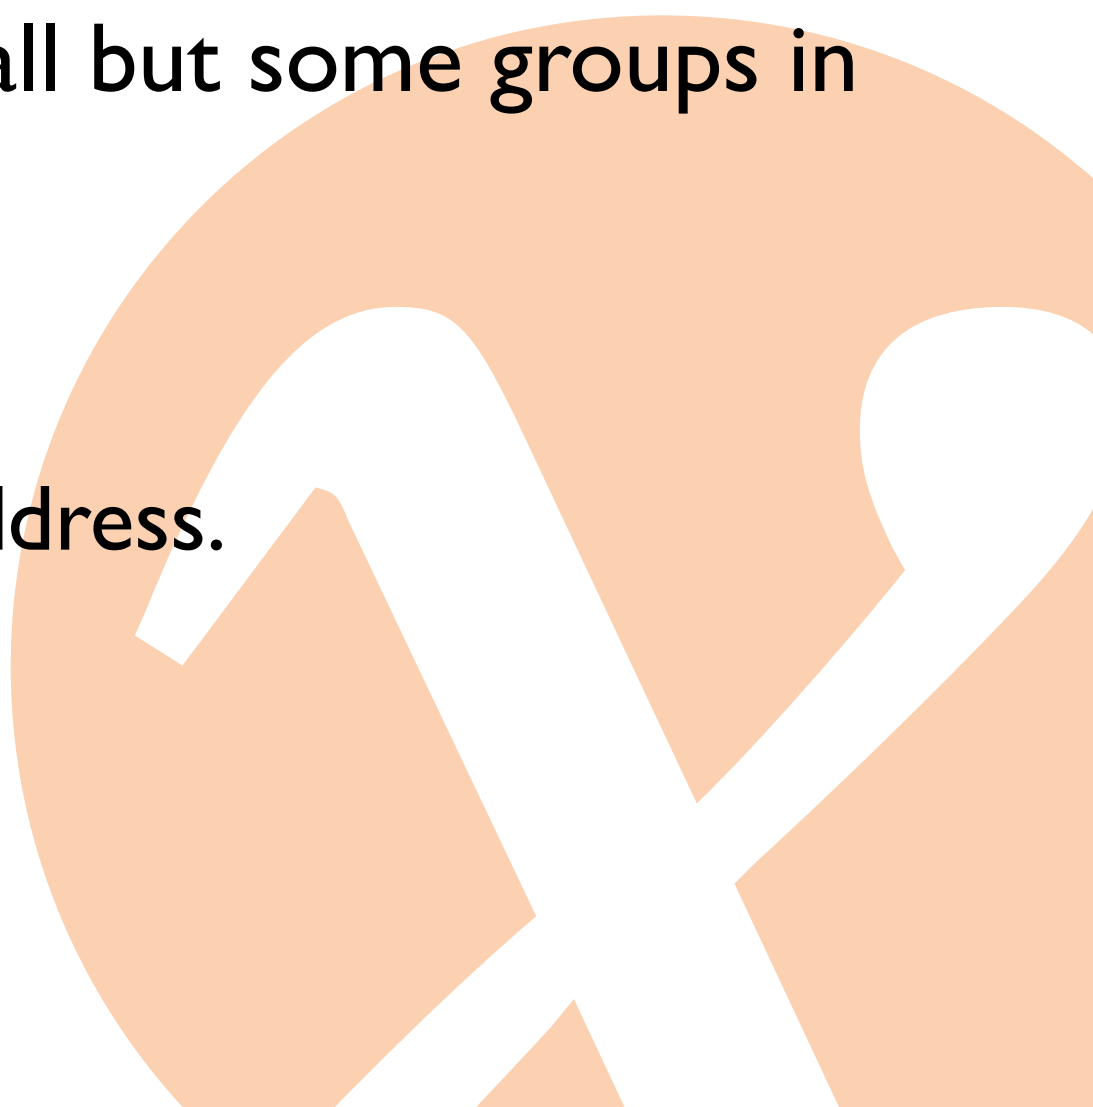
- Operational issues
 - Statistics on NS messages
 - Analysis
 - Router bugs
 - Conclusion







NS messages

- Is this really a lot?
 - Divided in ~400 multicast groups.
 - Routers should be able to ignore all but some groups in hardware.
 - 1 group for the link local address.
 - 1 group for the globally unicast address.
 - All nodes multicast addresses.
- 

NS messages

- Cisco GSR user complaint.
- Under IOS, all ND/NS messages are processed by the routing processor.
- This even caused BGP drops!
- Solution:
 - IPv6 filter.



NS messages

- Statistics showed significant more NS for certain addresses.
- These addresses did not react on NS messages.
 - Multicast
- Address is reachable after setting a static entry in neighbor cache.
 - Unicast only.



ICMPv6 filtering

- Too restrictive incoming ICMPv6 filtering.
 - ICMP filtering tradition that comes from IPv4.
- IPv6 does not work without ICMPv6.

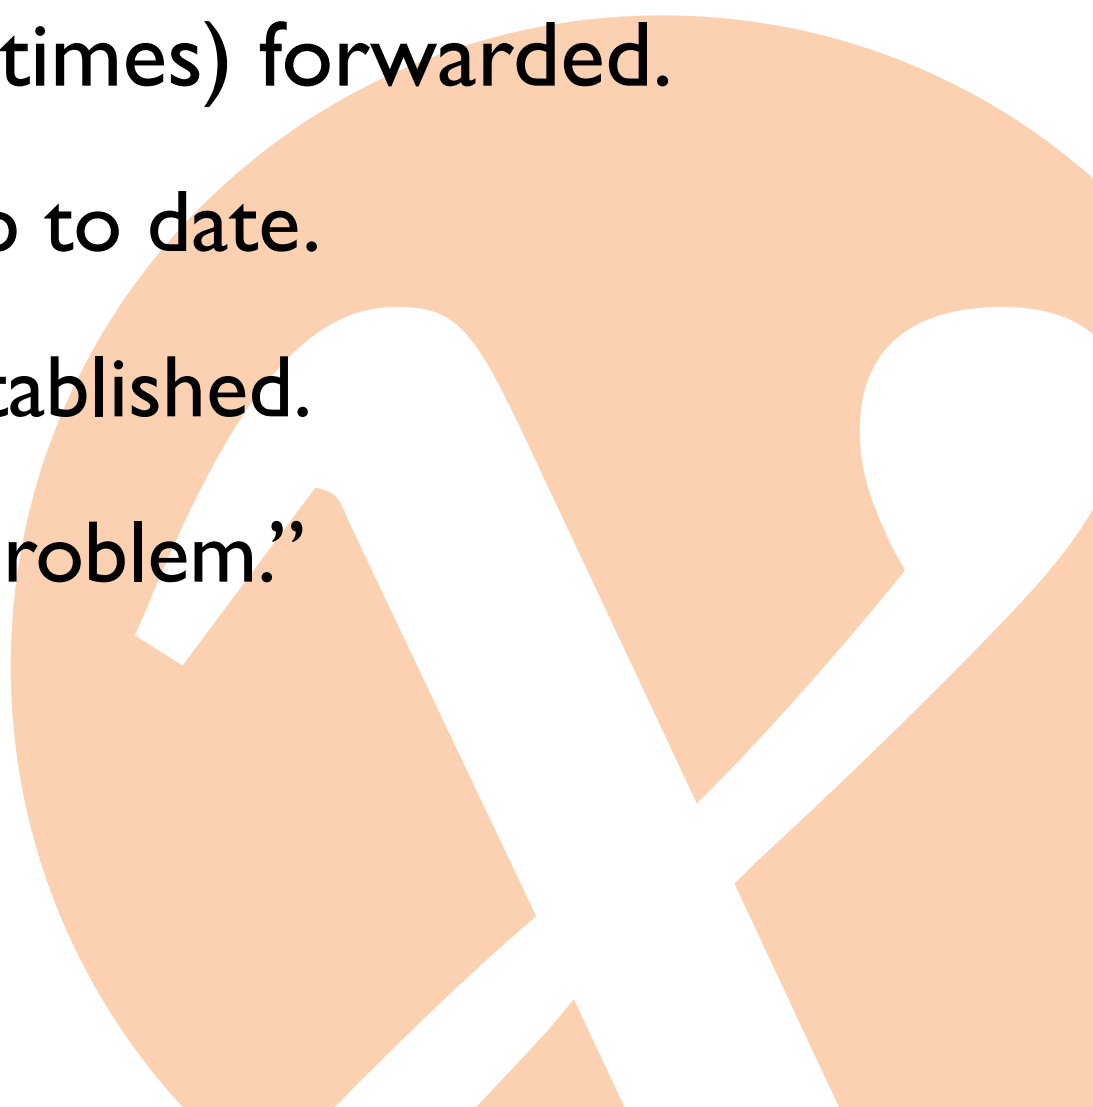


MLD

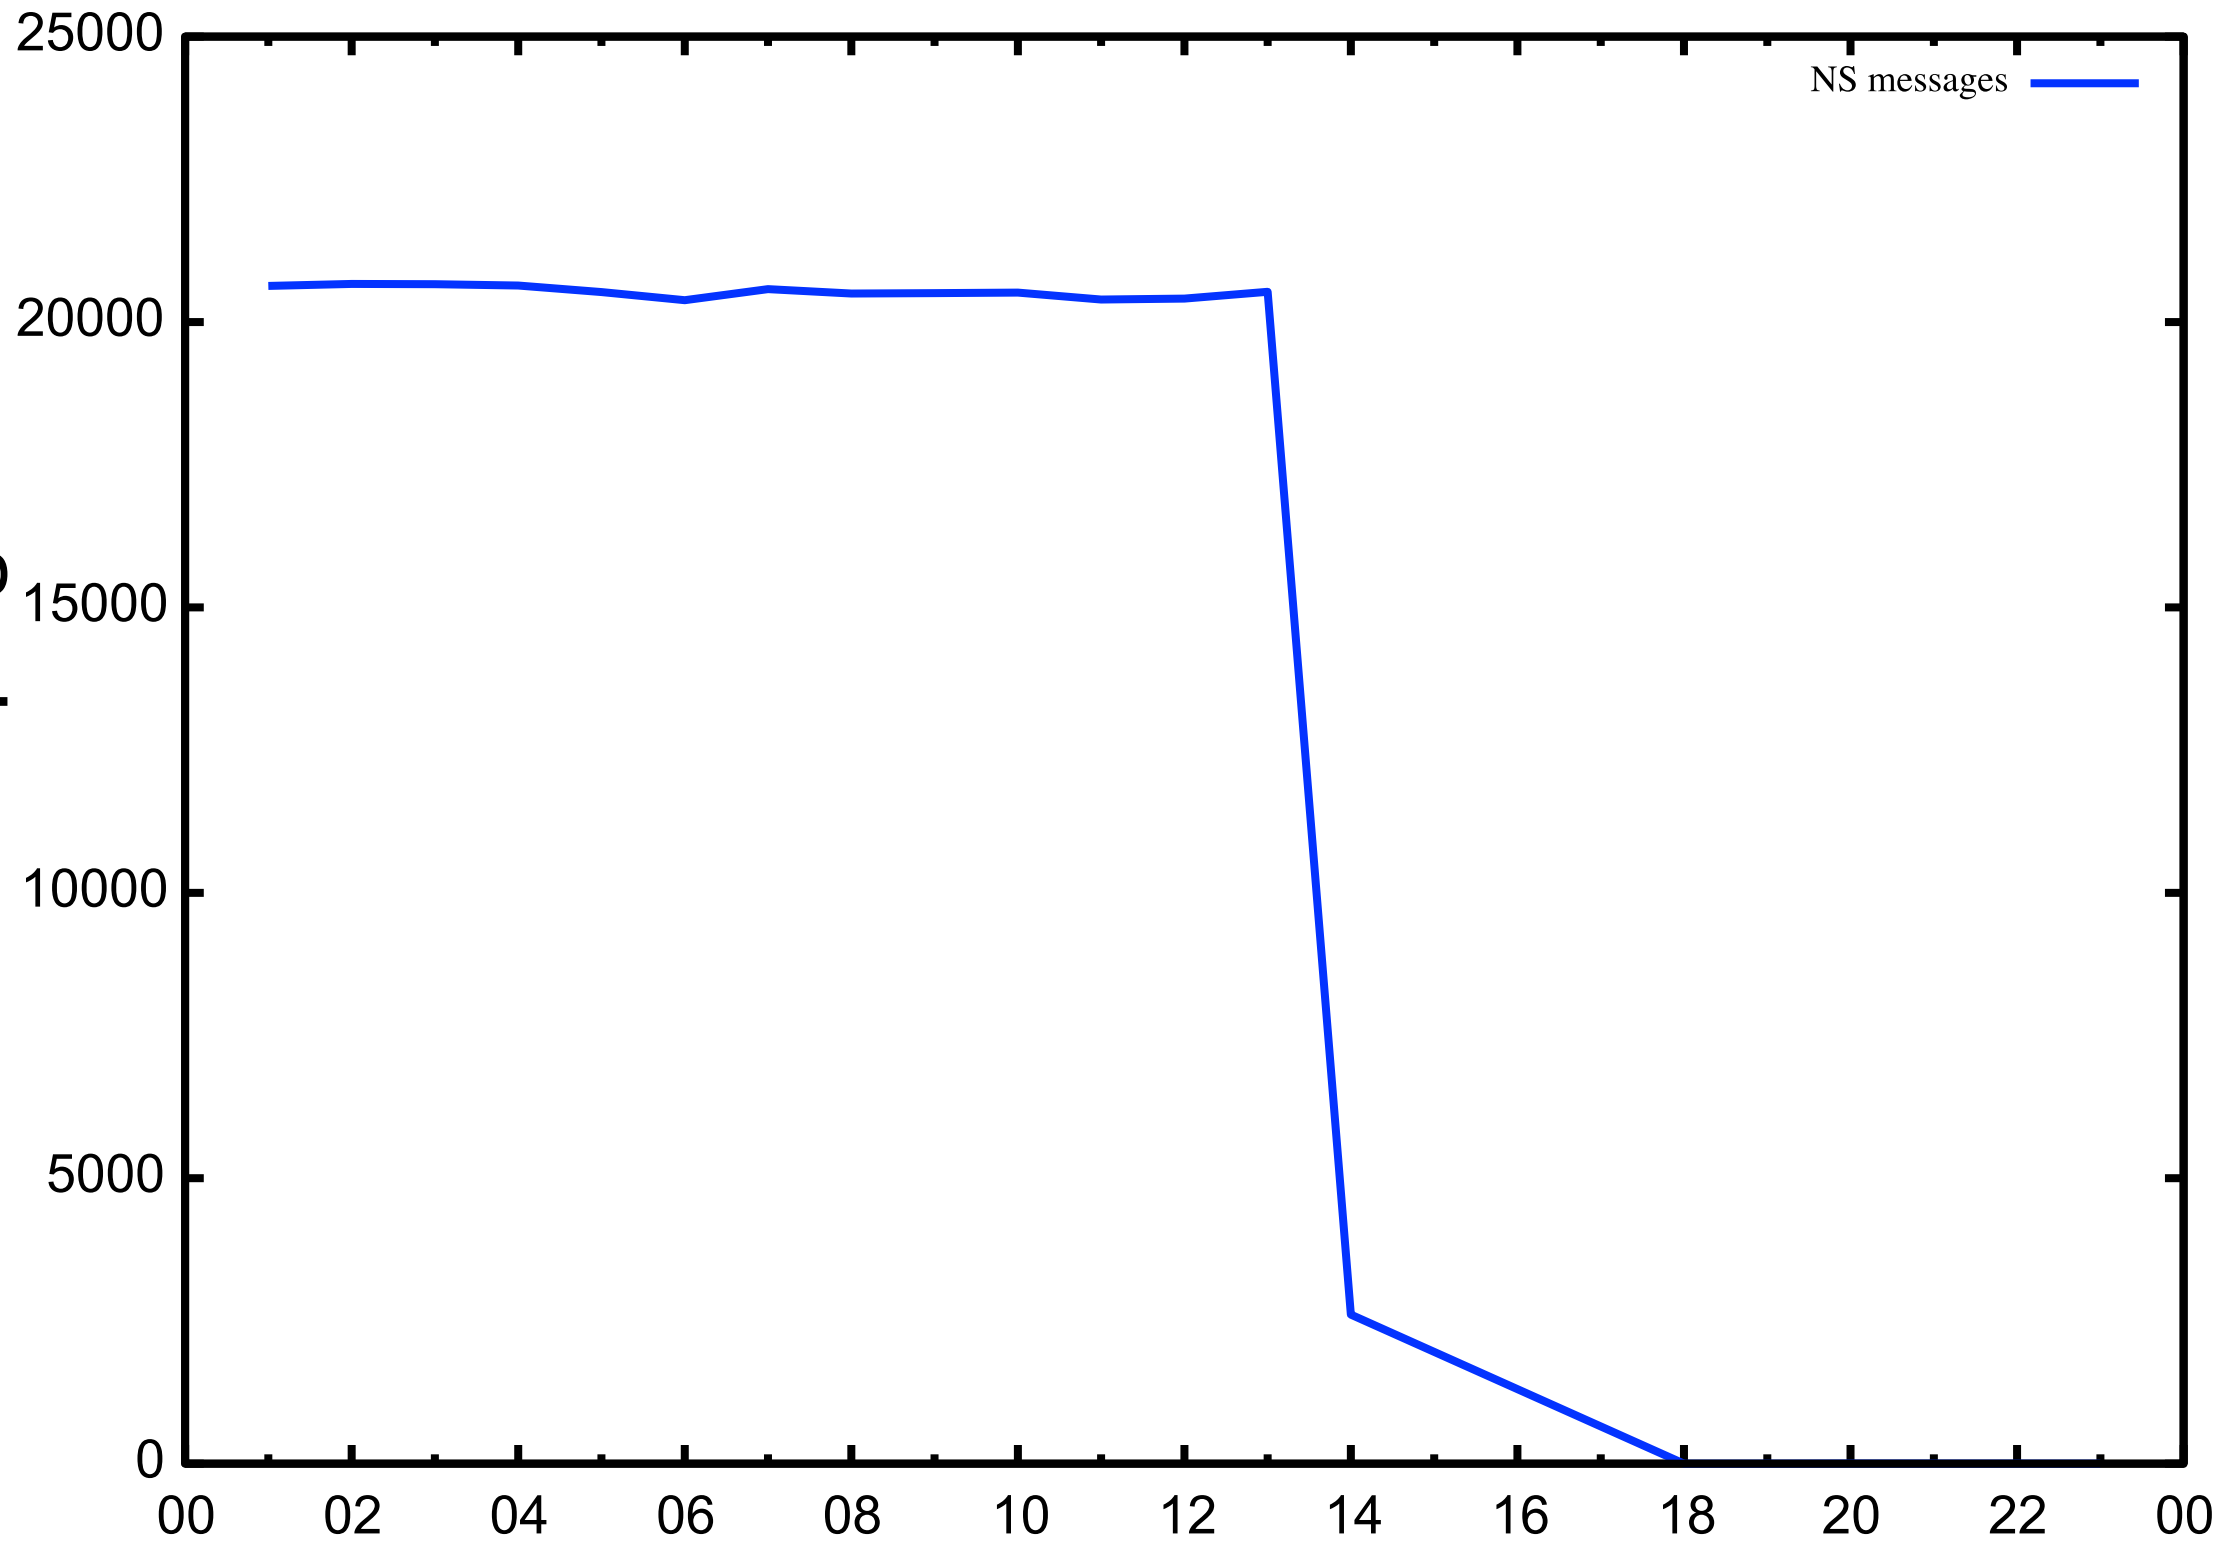
- One MLD query leads to many MLD listener reports.
 - The latter can't always be switched off.
- Default on:
 - Cisco
 - Linux kernels 2.6.26 and up
 - no `sysctl` to switch it off
 - `iptables`



MLD snooping


- Cisco 7600/6500 MLD snooping block NS messages.
 - Multicast groups not in listener reports.
 - Outgoing NS messages are (sometimes) forwarded.
 - Unicast keeps neighbor caches up to date.
 - BGP session gets and remains established.
 - “All sessions are up, there is no problem.”
 - Route server peers
 - Next-hop does not get resolved.
- 

NS messages for router
with mld snooping on/off



Time

NS messages from router with mld snooping on/off



Time

Agenda

- Operational issues
 - Statistics on NS messages
 - Analysis
 - Router bugs
 - Conclusion



Mostly Harmless?

- Garbage frames when IPv6 is enabled.
 - Cisco GSR.
 - CSCta73585
 - Corrupted NS messages.



Mostly Harmless?

- Juniper various JunOS
- Raw packets.
- DA:6c05.6ccc.0014 SA:0601.2001.07f8

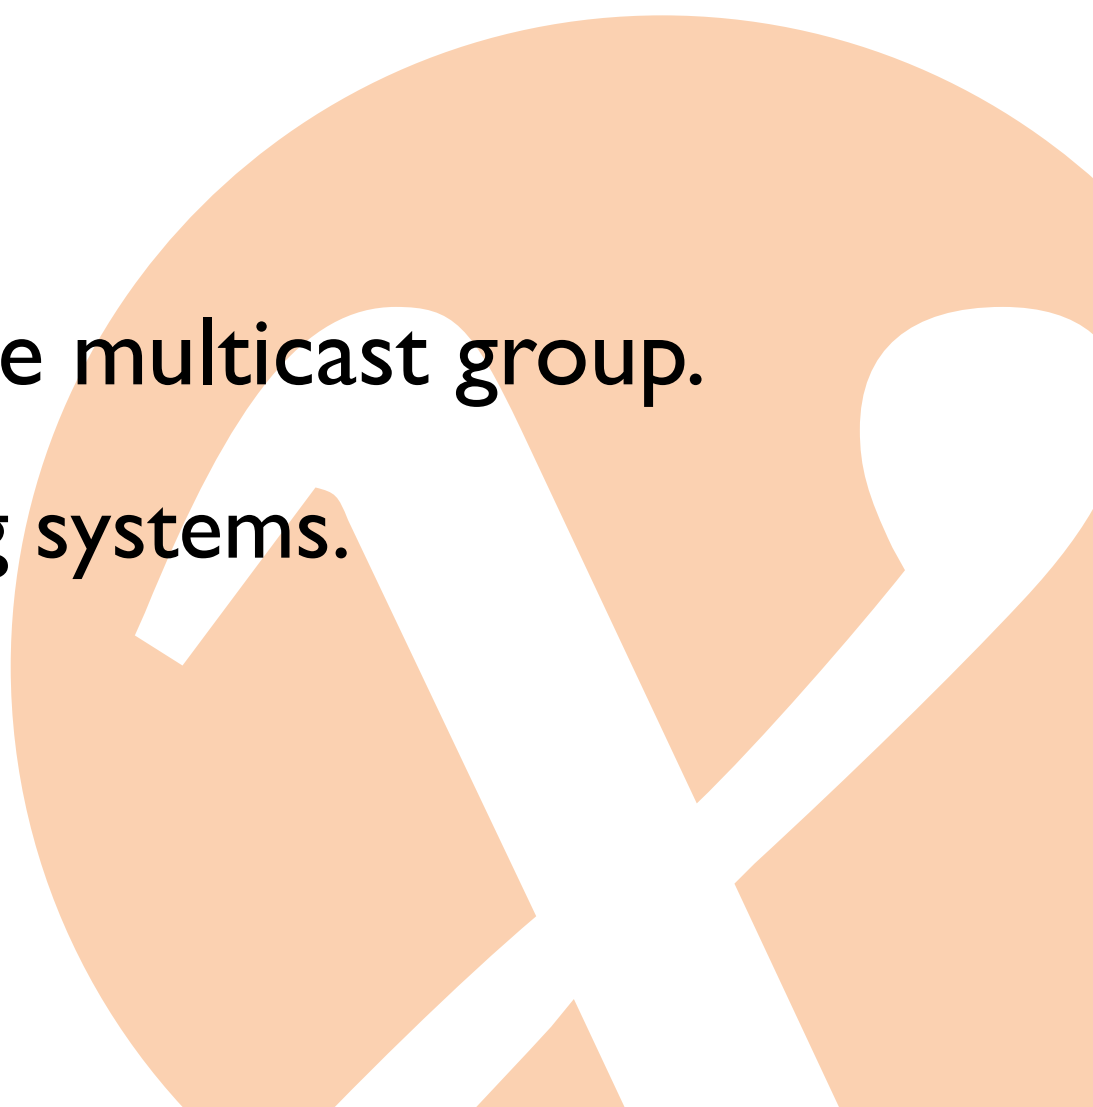
Version: 6,
Traffic class: c0,
Flowlabel: 56ccc,
Payload Length: 14,
Next header: 6 (TCP),
Hop Limit: 1
Source address: 2001.07f8...



Mostly Harmless?



Harmful

- JunOS 10.4 (R1 to R3), 11.1
 - Answers any NS message that it receives, when the target address is in its neighbor cache.
 - ND spoofer.
 - Target address must be in the same multicast group.
 - The case for many IXP numbering systems.
- 


Agenda

- Operational issues
 - Statistics on NS messages
 - Analysis
 - Router bugs
 - Conclusion



Conclusion

- ND/NS messages can be blocked
 - ICMPv6 filtering.
 - MLD snooping.
- Robustness of protocol can be deceiving.
 - “All our BGP sessions are established”
- Still bugs
 - Not so harmless.



JUST DO IT.

