Peering observations on security and resiliency at IXPs

Greg Hankins, AS 38016
<greg.hankins@nokia.com>

NANOG 67

Image source: http://as2914.net/

NOKIA

Agenda

- Introduction
- Research questions
- BGP configuration results

Supporters

- Contributors: DE-CIX R&D
  - Thomas King <thomas.king@de-cix.net>
  - Christoph Dietzel <christoph.dietzel@de-cix.net>

- Thanks: everyone who peers with AS 38016!

My perspective on peering

Managed BGP routing and peering at AS 4355

1998

1999

Started working for router vendors

...

...

...

Peering project with AS 18508 at PAIX

2007

PAIX Palo Alto: 118 members

...

...

...

Peering project with AS 38016 at DE-CIX

2016

DE-CIX Frankfurt: 677 members
DE-CIX New York: 113 members

Changes in the industry

Goal of this presentation: look at how peering is deployed from an operational perspective

Research questions

- Very interested in how peering is deployed from a operational configuration and technical perspective
- What features are people using and not using at an IXP?
  - And why?
  - What are the implications?
  - Especially interesting to me as a router vendor because of feature impact
- Not really interested in prefixes or path attributes
  - Other excellent research projects* are looking at these things

```
A:br1-nyc.iplabs# configure
A:br1-nyc.iplabs>config# router
A:br1-nyc.iplabs>config>router# bgp
A:br1-nyc.iplabs>config>router>bgp# info
----------------------------------------------
            enable-peer-tracking
            add-paths
                ipv4 send 1 receive
                ipv6 send 1 receive
            exit
            backup-path ipv4 ipv6
            best-path-selection
                always-compare-med strict-as zero
                deterministic-med
            exit
A:br1-nyc.iplabs>config>router>bgp#
```

*See references slide

What can you observe passively from BGP peering?

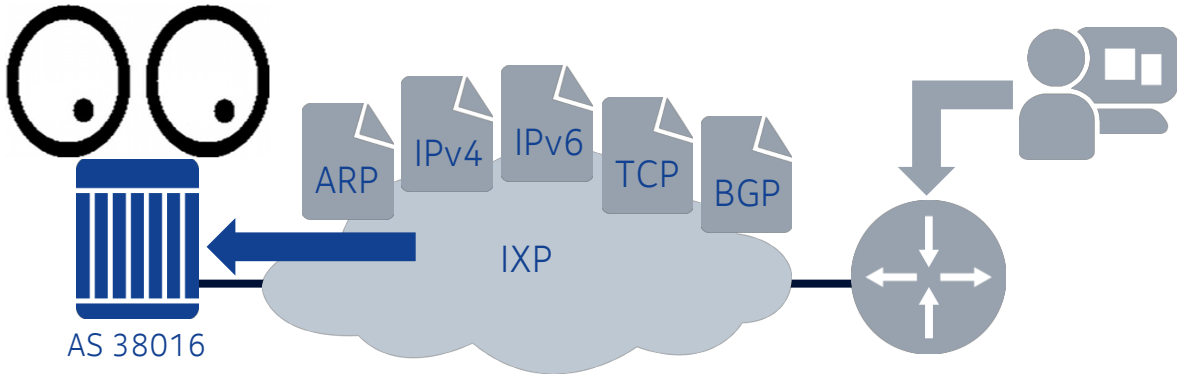

Can observe peering behavior	Cannot observe configuration
• Control plane packets	• Router configuration
• Protocols (IPv4/IPv6)	• Routing policies
• BGP capabilities	• Peering policies
• BGP routes and prefixes	• Network architecture
• BGP path attributes	• Not interested in these things

BGP peering configuration results

Configuration	2007 PAIX Palo Alto	2016 DE-CIX Frankfurt & New York	Observations
Unique BGP peers[1]	21	25	
Total BGP peers[2]	30	76	IPv6, multiple connections, route servers
IPv4 unicast BGP peers	21	39	
IPv6 unicast BGP peers	6	37	Almost 100% IPv6 peering!
IPv4 multicast BGP peers	3	0	No interdomain multicast
IPv6 multicast BGP peers	0	0	No interdomain multicast
MSDP peers	1	0	No interdomain multicast
BFD peers	0	0	Interesting!
Route server peers	0	4	Not available in 2007 at PAIX

[1]Unique BGP peers: unique peers based on peer ASN
[2]Total BGP peers: total operational BGP peers

BGP security features in use

Feature	2007	2016	Observations
Unique BGP peers	21	25	
Protection of BGP Sessions via the TCP MD5 Signature Option (RFC 2385)	16	4	Rarely used or required now, but why?
The Generalized TTL Security Mechanism (GTSM) for BGP (RFC 5082)	Not supported	0	23 with TTL = 1 2 with TTL = 64
The TCP Authentication Option (RFC 5925)	Not available	0	No interest?

Protection of BGP Sessions via the TCP MD5 Signature Option

- RFC 2385 basics
  - Defines a TCP option for carrying an MD5 digest in a TCP segment
  - The digest is a signature for that segment
  - Protects BGP against spoofed TCP segments (RST)
- People generally dislike MD5, because it causes more operational problems than it solves
  - Password negotiation
  - Password management, storage and lost passwords
- Notable decrease in use, and rarely required now
  - 2007: 16 of 21 unique peers
  - 2016: 4 of 25 unique peers

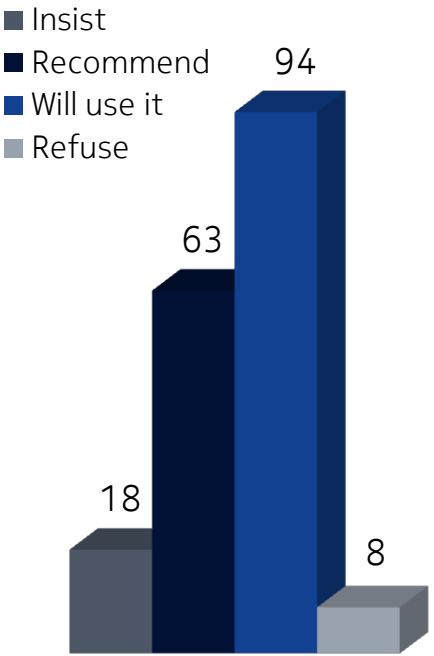

Peering Survey 2008 – MD5
169 respondents

Source: http://www.twoguys.org/~gregh/peering-survey-2008/peering-survey-2008-results.pdf

The Generalized TTL Security Mechanism (GTSM) for BGP

- RFC 5082 basics for BGP
  - EBGP peering is usually between routers that are on the same IP network, especially at an IXP
  - Use Time to Live (TTL) (IPv4) or Hop Limit (IPv6) to see if a packet came from a different IP network
  - Send TTL = 255, expect TTL = 255
  - TTL < 255 means the packet was spoofed, because routers always decrement the counter
- Requires manual negotiation and configuration
  - Commands and logic varies between BGP implementations
  - Must configure matching send and receive behavior on each router
- Not requested or in use by any peers
  - 23 of 25 unique peers sending TTL = 1
  - 2 of 25 unique peers sending TTL = 64

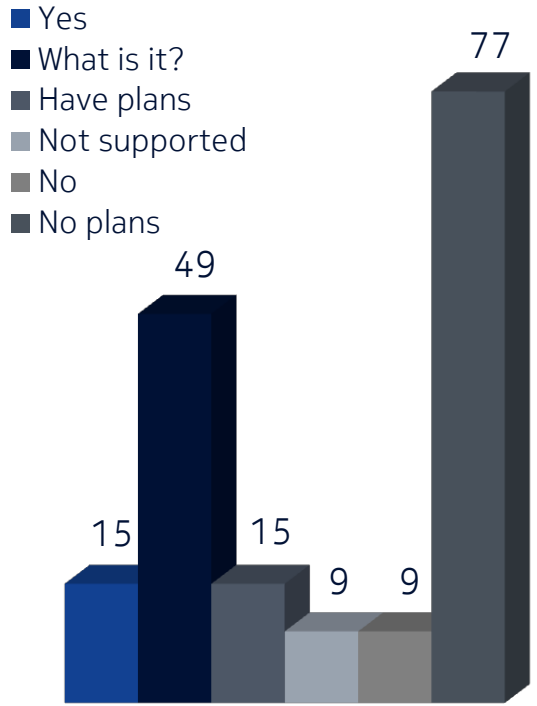

Peering Survey 2008 – GTSM
169 respondents

Source: http://www.twoguys.org/~gregh/peering-survey-2008/peering-survey-2008-results.pdf

BGP capabilities in use

Feature	2007	2016	Observations
Unique BGP peers	21	25	
Multiprotocol Extensions for BGP-4 (RFC 4760, Capability Code 1)	21	25	Router is capable of MP, probably IPv6
Multiple address families advertized	Didn't count	2	1 x "IPv4 IPv6", 1 x "IPv4 MCAST-IPv4"
Route Refresh Capability for BGP-4 (RFC 2918, Capability Code 2)	21	25	Enabled by default on most routers, but?
Old Cisco Route Refresh Capability (Capability Code 128)	21	0	Historical, not advertised anymore
Enhanced Route Refresh Capability for BGP-4 (RFC 7313, Capability Code 70)	Not available	0	Too early?
Outbound Route Filtering Capability for BGP-4 (RFC 5291, Capability Code 3)	0	0	Does it make sense at an IXP?
Graceful Restart Mechanism for BGP (RFC 4724, Capability Code 64)	2	0	Interesting!
BGP Support for Four-Octet Autonomous System (AS) Number Space (RFC 6793, Capability Code 65)	0	24	Broad support now
Advertisement of Multiple Paths in BGP (draft-ietf-idr-add-paths, Capability Code 69)	Not available	0	Too early?

Interesting observations based on protocol configurations

Category	Observations	Conclusions?
Security	• IXP network is often leaked into the global routing table • MD5 is rarely used • No TCP-AO • No GTSM	• Don't care? • Too complicated? • Trust IXP and peers? • Is TCP/BGP packet security (still) an issue? • Use policy filtering and/or external databases instead?
Resiliency	• No BFD, peer aliveness detection via slow BGP mechanism for remote link failure • No BGP graceful restart • No ADD-PATH • BGP route refresh is enabled by default, is it used?	• Don't care? • Too complicated? • Rely on stability of IXPs or peering infrastructure? • Design around IXP and peer failures?

Little interest in member protocol security and resiliency at IXPs… but why?

References

- All data as seen on Mon Apr 5 04:00:00 EDT 2016 by AS 38016
- EPF 2007 “Peering Observations” presentation
  - Not available online anymore, ask me and I’ll send it to you
- Peering Survey 2008 announcement
  - https://www.nanog.org/meetings/nanog42/presentations/peering-bof-XVII.pdf
- Peering Survey 2008 results
  - http://www.twoguys.org/~gregh/peering-survey-2008/peering-survey-2008-results.pdf
  - http://meetings.ripe.net/ripe-56/presentations/eix/Hankins-Peering_Survey_2008_Results.pdf
  - http://www.uknof.org.uk/uknof10/Davidson-Peering-survey.pdf
- EPF 2015 “Peering Observations 2007 vs. 2015” presentation
  - https://www.peering-forum.eu/system/documents/55/original/20150921_0900_greg_hankins_epf-10-peering-observations.pdf
- GPF 11.0 “Peering Observations 2007 vs. 2016” presentation
  - https://www.peeringforum.com/11.0%20presentations/Wed_4_gpf-11-peering-observations.pdf
- Some excellent research projects and tools
  - APNIC R&D: http://bgp.potaroo.net/
  - Route Views: http://www.routeviews.org/
  - RIPEstat: https://stat.ripe.net/
  - BGPmon: http://www.bgpmon.net/
  - NLNOG RING: https://ring.nlnog.net/

BGP Capabilities as seen from a real network

Job Snijders

job@ntt.net

NANOG67, Chicago

NTT Communications / AS 2914

We looked at **tens of thousands of BGP sessions,**
landing on **hundreds of devices**,
spread over **20 countries.**

% of eBGP sessions IPv4 versus IPv6

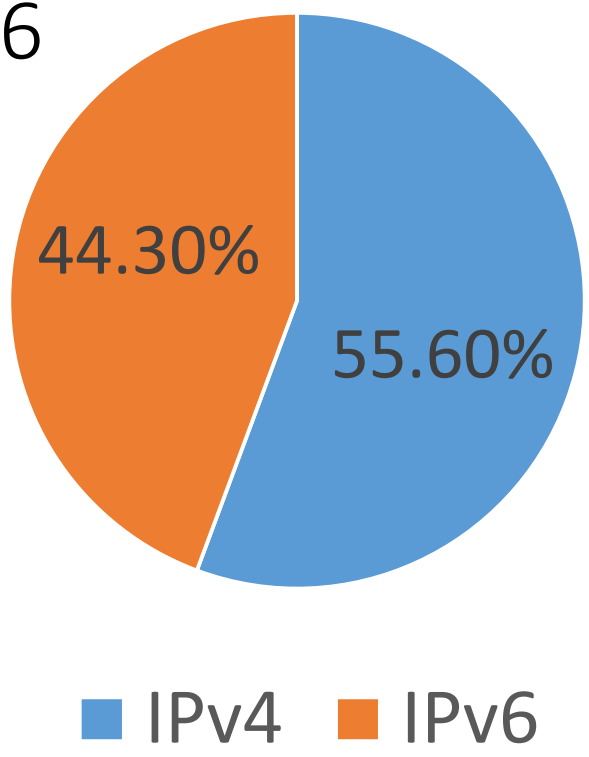

IPv4/IPv6 ASN adjacency overlap

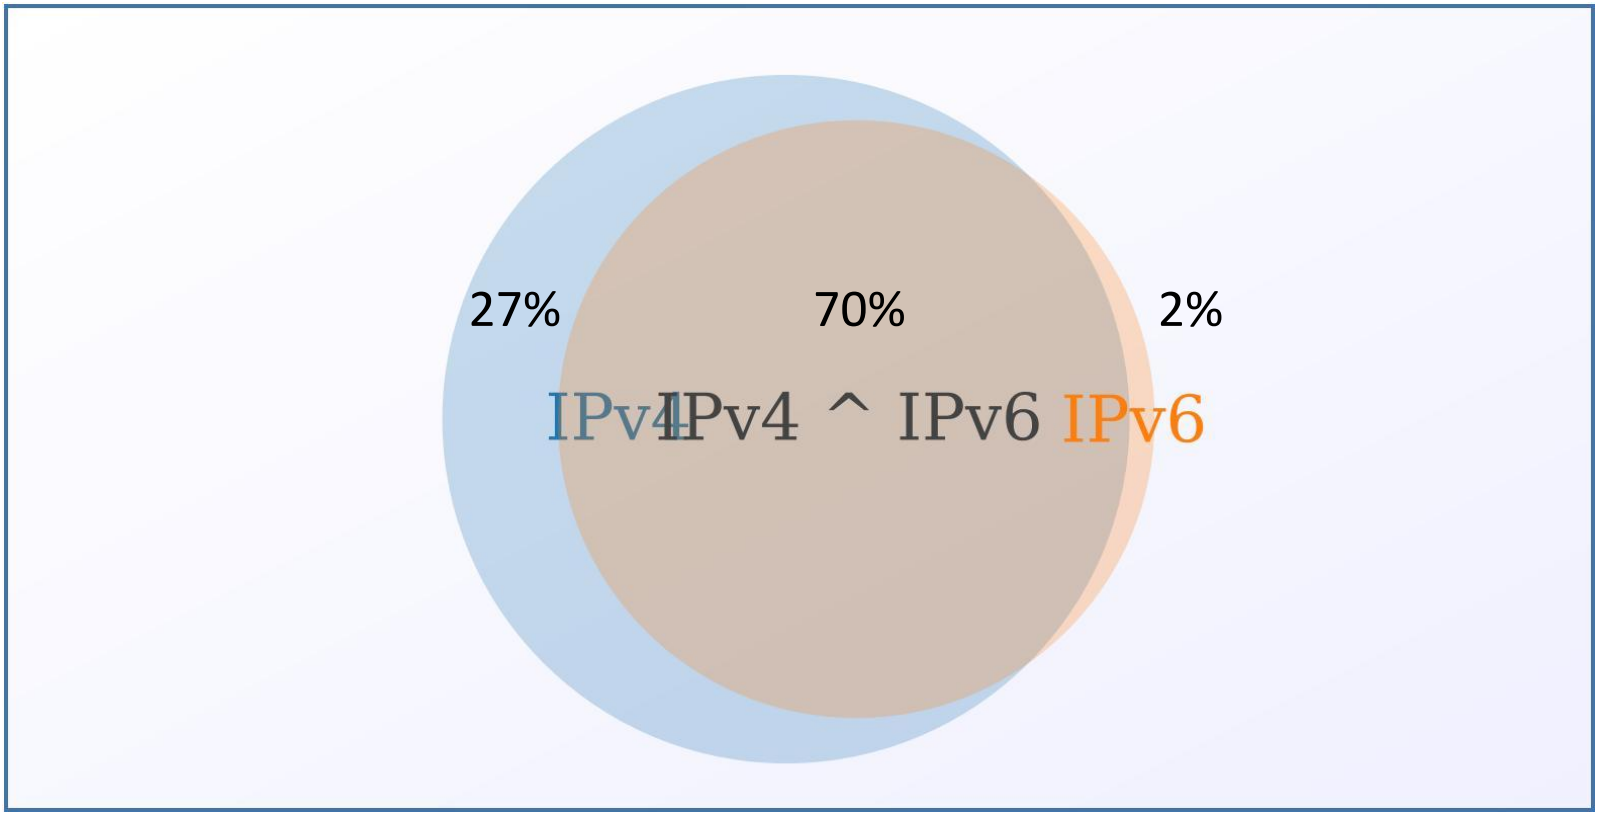

Route Refresh capability

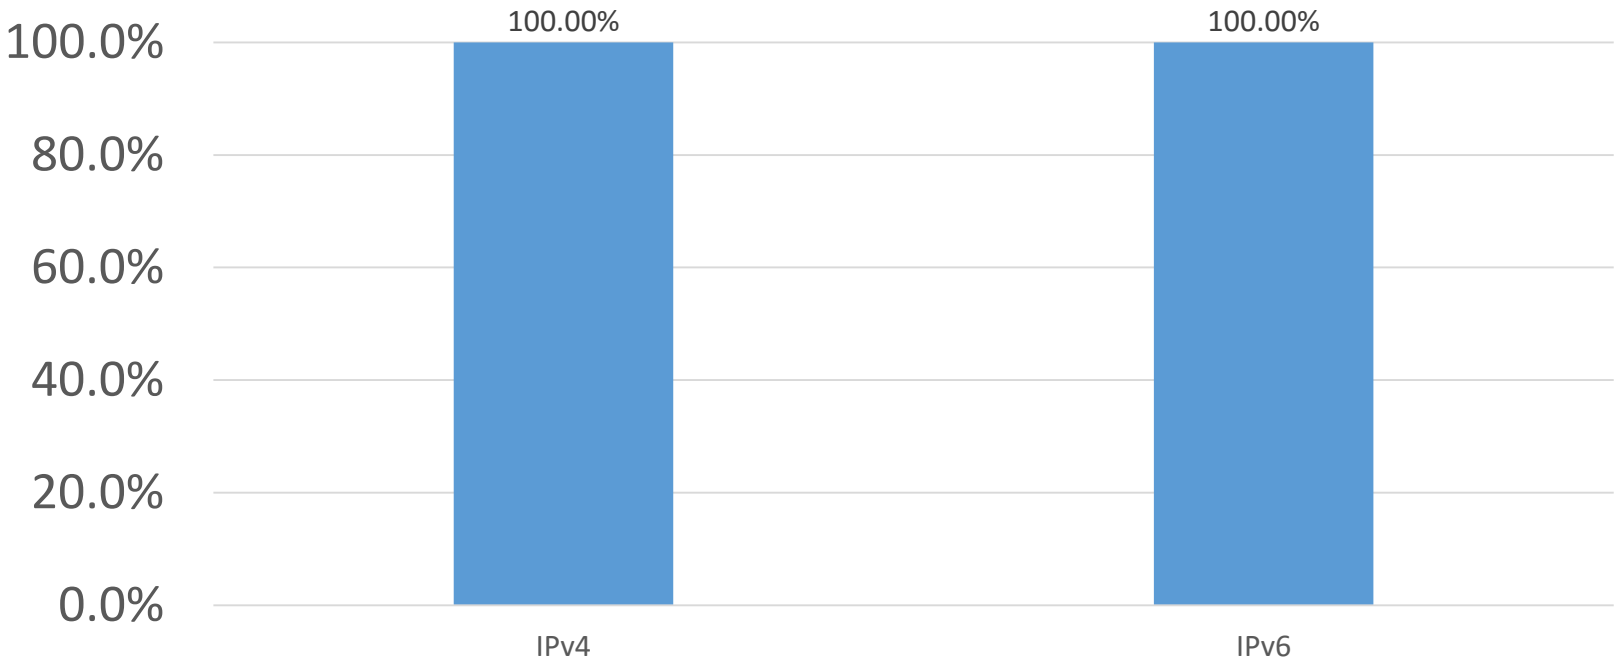

Enhanced Route Refresh capability

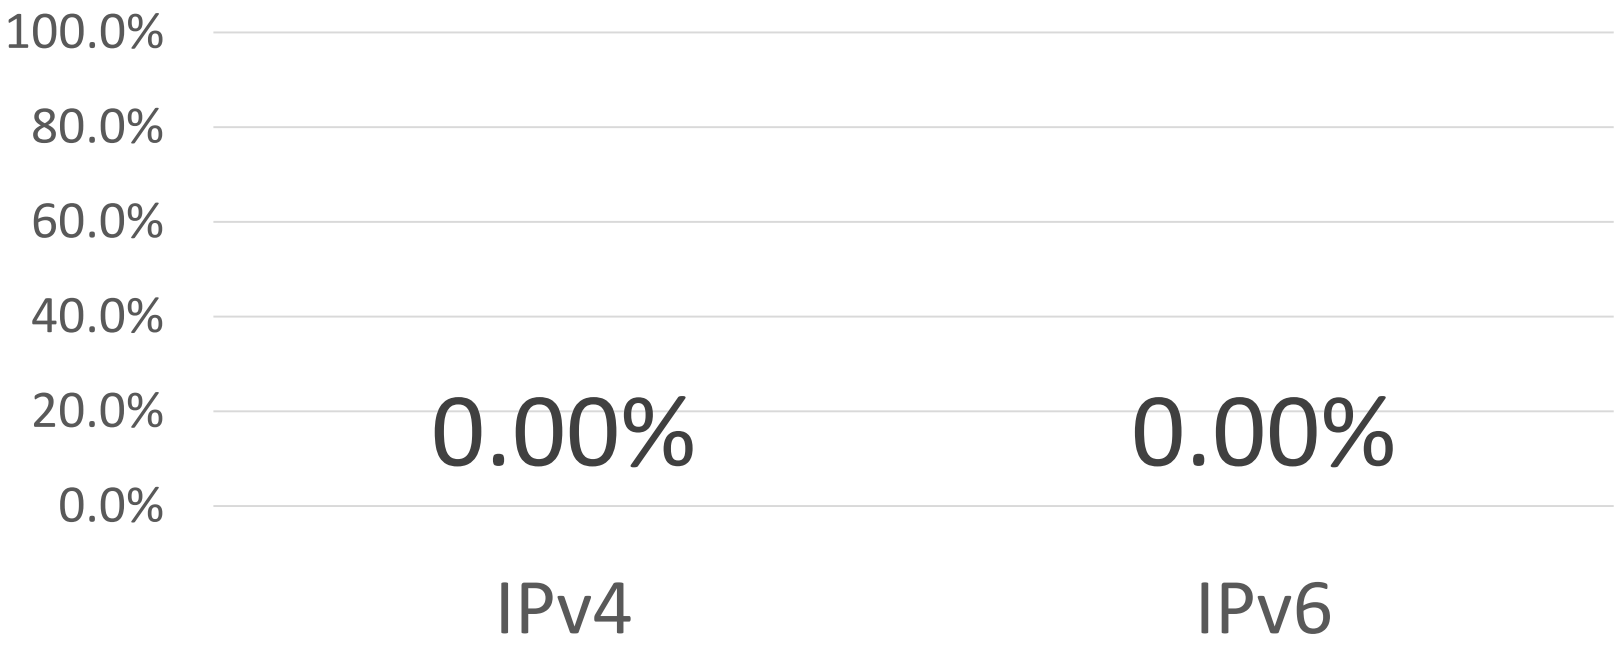

Graceful restart capability

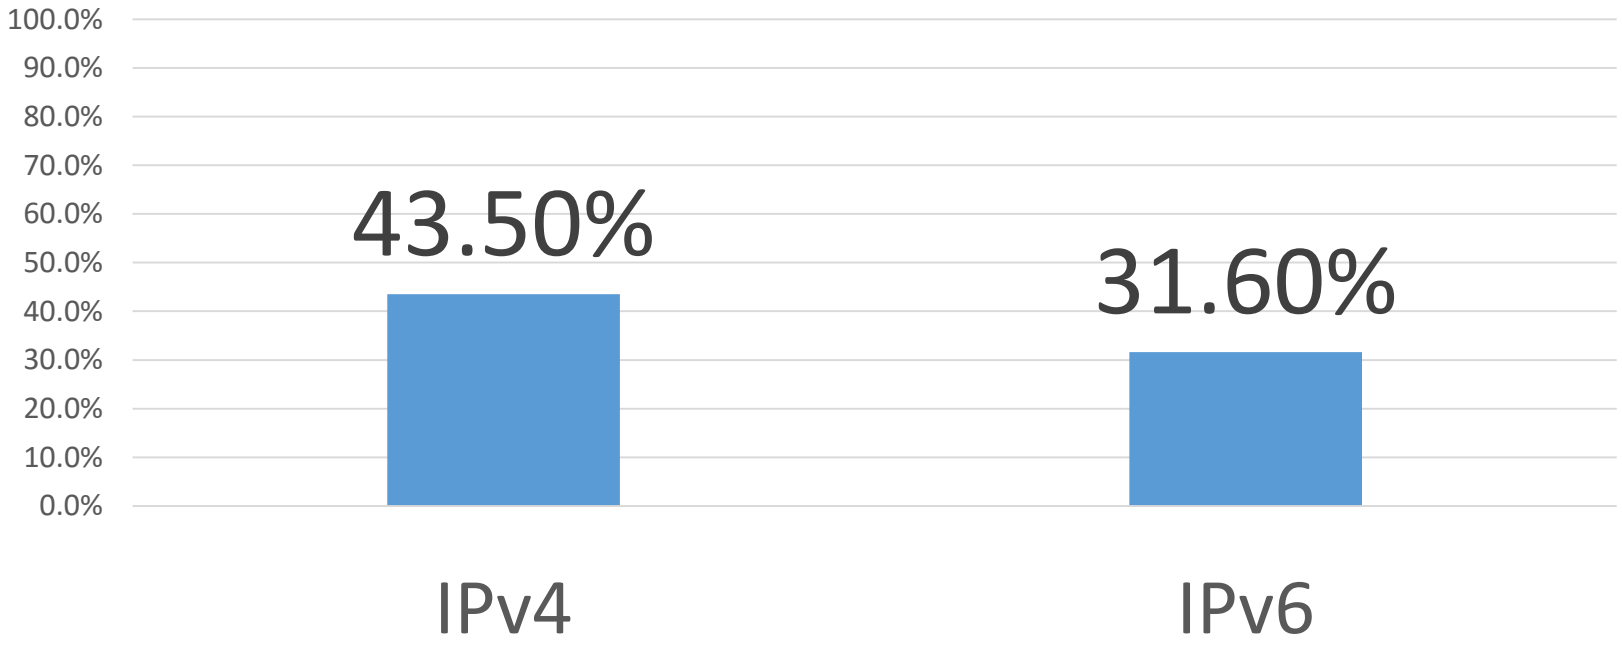

4-byte ASN capability

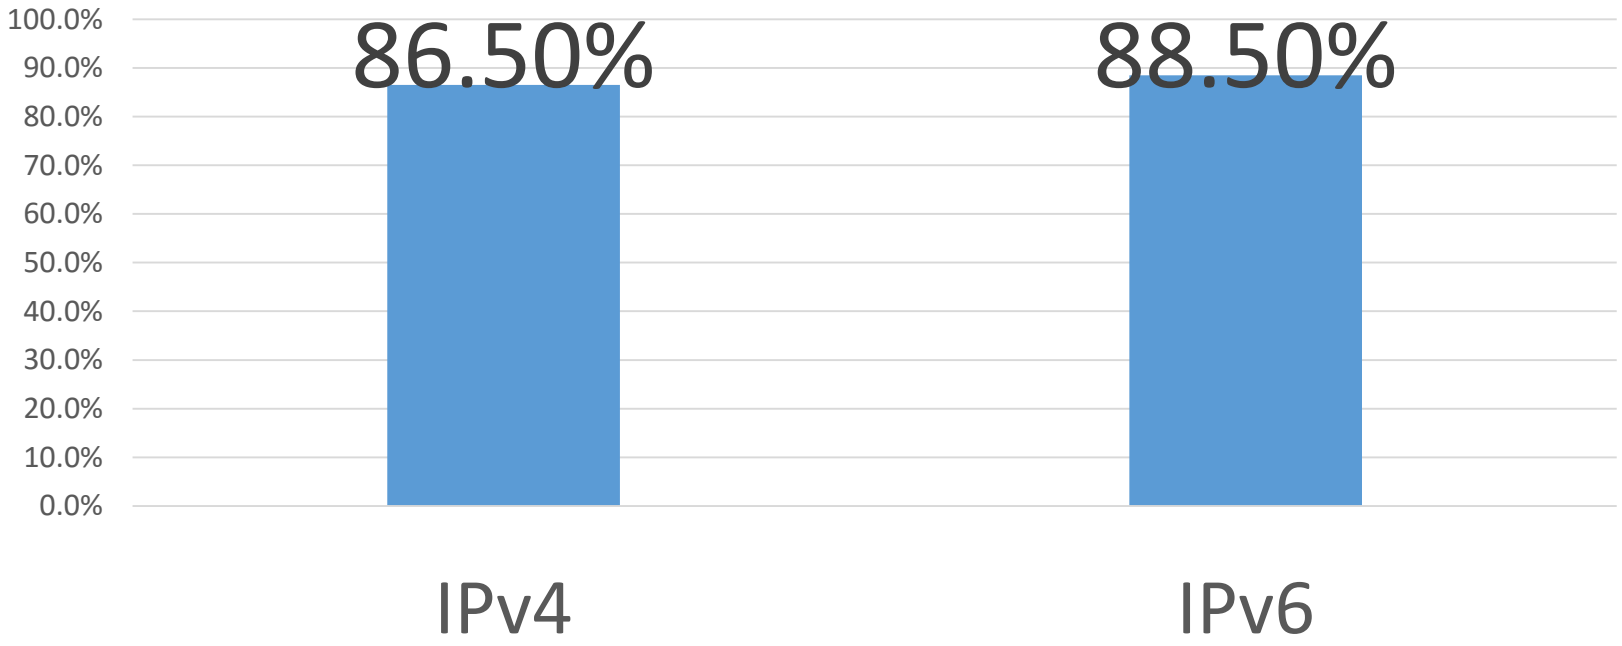

MD5

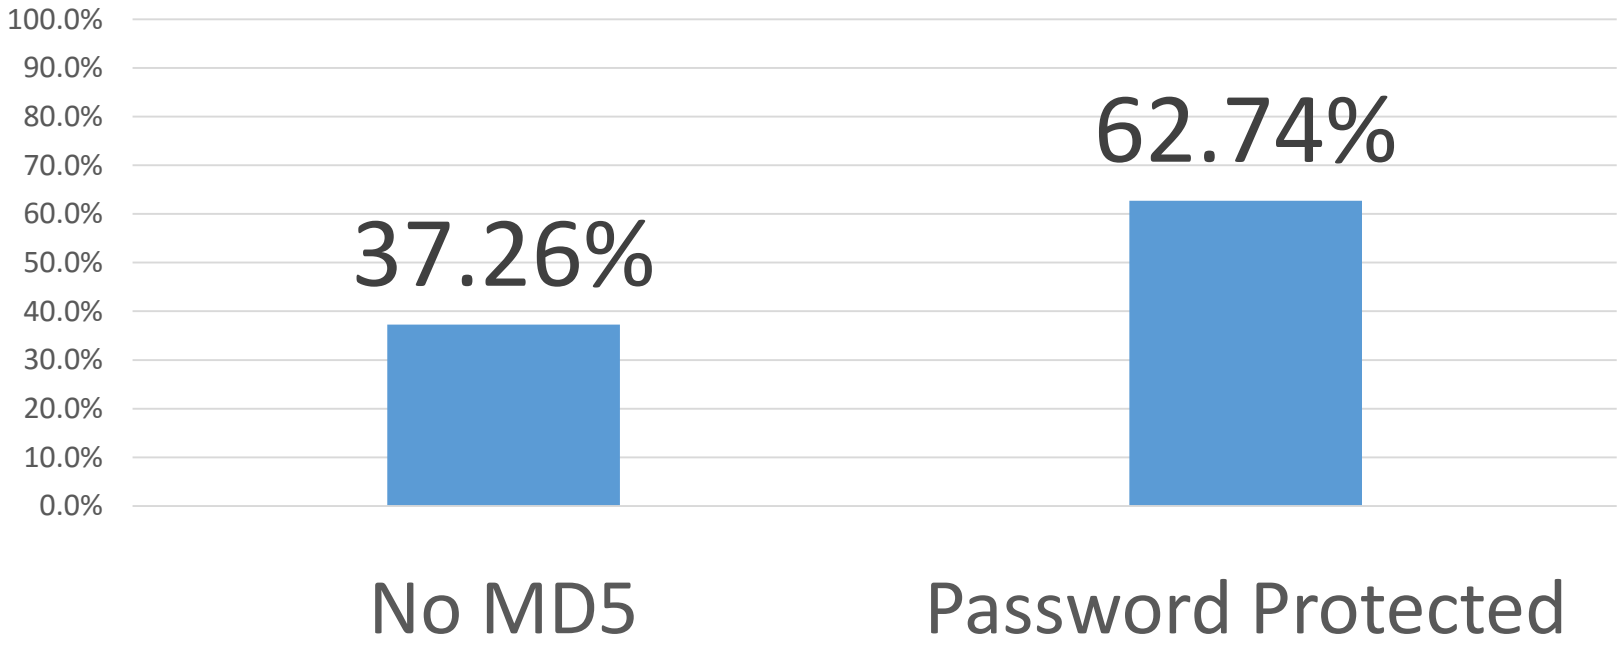
100.0%
90.0%
80.0%
70.0%
60.0%
50.0%
40.0%
30.0%
20.0%
10.0%
0.0%
37.26%
62.74%
No MD5
Password Protected

IPv4 Multicast

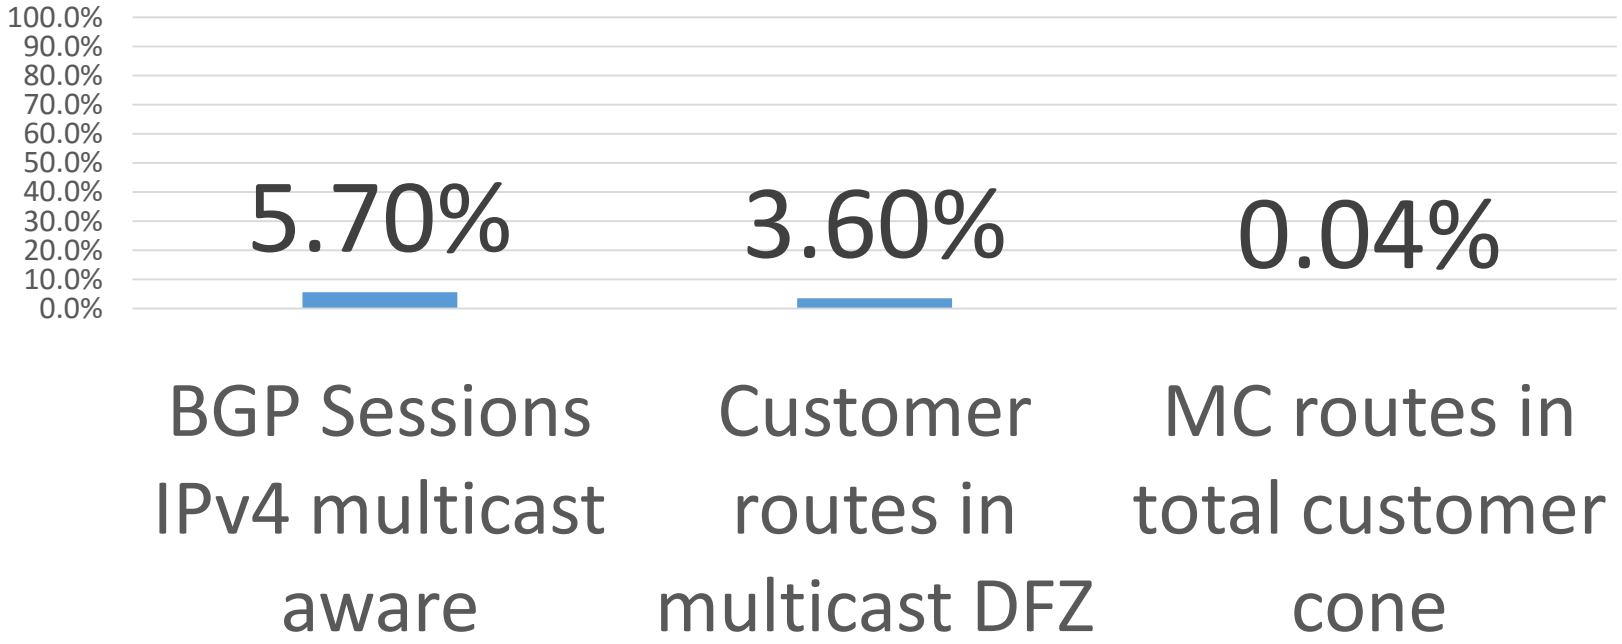

BFD, ORF ADD-PATH

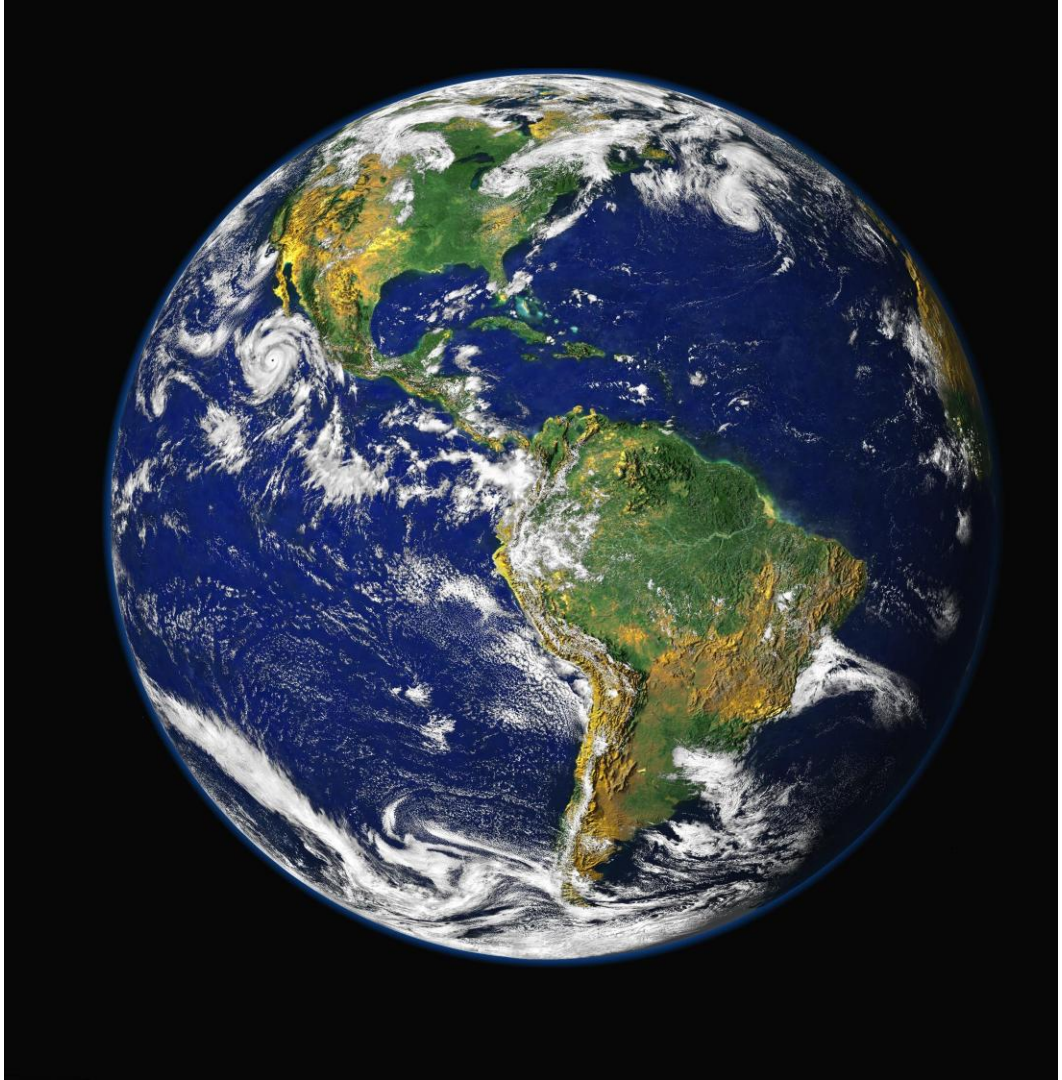

Questions?
Answers?
Comments?
Observations?