Karlsruhe Institute of Technology

ESnet
ENERGY SCIENCES NETWORK

LHCONE Edge Filtering Policy and Practice

Bruno Hoeft / KIT
Michael O'Connor / ESnet

STEINBUCH CENTRE FOR COMPUTING - SCC

NSP Packet Filtering Requirements

All LHCONE Traffic is subject to the following conditions:

- Traffic injected into the LHCONE must only be originated from addresses within an LHCONE routable prefix
- Only address ranges present in the LHCONE routing table should be transported on the network

Objective: In order to maintain route symmetry and access control, each NSP will implement policy and packet filters to manage their connected customer address prefix ranges.

- Ensures that a return route exists in the LHCONE network
- Blocks spoofed packets (Similar to BCP 38)

https://twiki.cern.ch/twiki/pub/LHCONE/LhcOneVRF/LHCONEconnectionguide-1.2.pdf

NSP BGP Import Policy

Prefix Lists will be negotiated between connecting institutions and their NSP within the constraints imposed by the LHCONE AUP.

LHCONE NSPs have agreed to to configure:

1. BGP import filters
2. Source address packet filters

End sites are encouraged to implement source address filters at their edge in order to count their own unroutable LHCONE packets. NSPs will generally discard these packets without informing the site.

Connecting institutions/sites will not add prefixes to the LHCONE routing table without direct cooperation with their NSP.

The Investigation

DE-KIT Ingress Packet Filters

- Unsampled ingress filtering detected LHCONE route table misses from over 44 source locations

Private IP destinations: 10.0.0.0/8, 172.16.0.0/12 192.168.0.0/16
renater, garr, jinr-net, tanet, tein, sut-th, nben-tw, ernet-in, aarniec, kisti

ESnet

- Three months of ESnet netflow IPv4 & IPv6 sampling from July 2018 - September 2018 for the following sites and peers

aarnet	fnal	nordunet	ufl
aglt2	geant	ou	uiuc
anl	ind-gpop	pnnl	unl
ansp	internet2	rnp	uta
asgc	JGN	sinet	uwmadison
bnl	kreonet	slac	vanderbilt
caltech	lhcone_cern	tacc	
canet	lhcone_ornl	uchicago	
cernlight	mit	ucsb	
duke	net2	ucsd	

ESnet counted:
- All LHCONE ingress packets
- Unroutable source packets
- Packets with non-lhcone/missing origin ASN

* corrected for netflow sampling rate

Detailed ESnet data at:
https://twiki.cern.ch/twiki/pub/LHCONE/LhcOneVRF/LHCONE-Filterdata-10-2018.pdf

unroutable packet count @ DE-KIT

IPv4 only

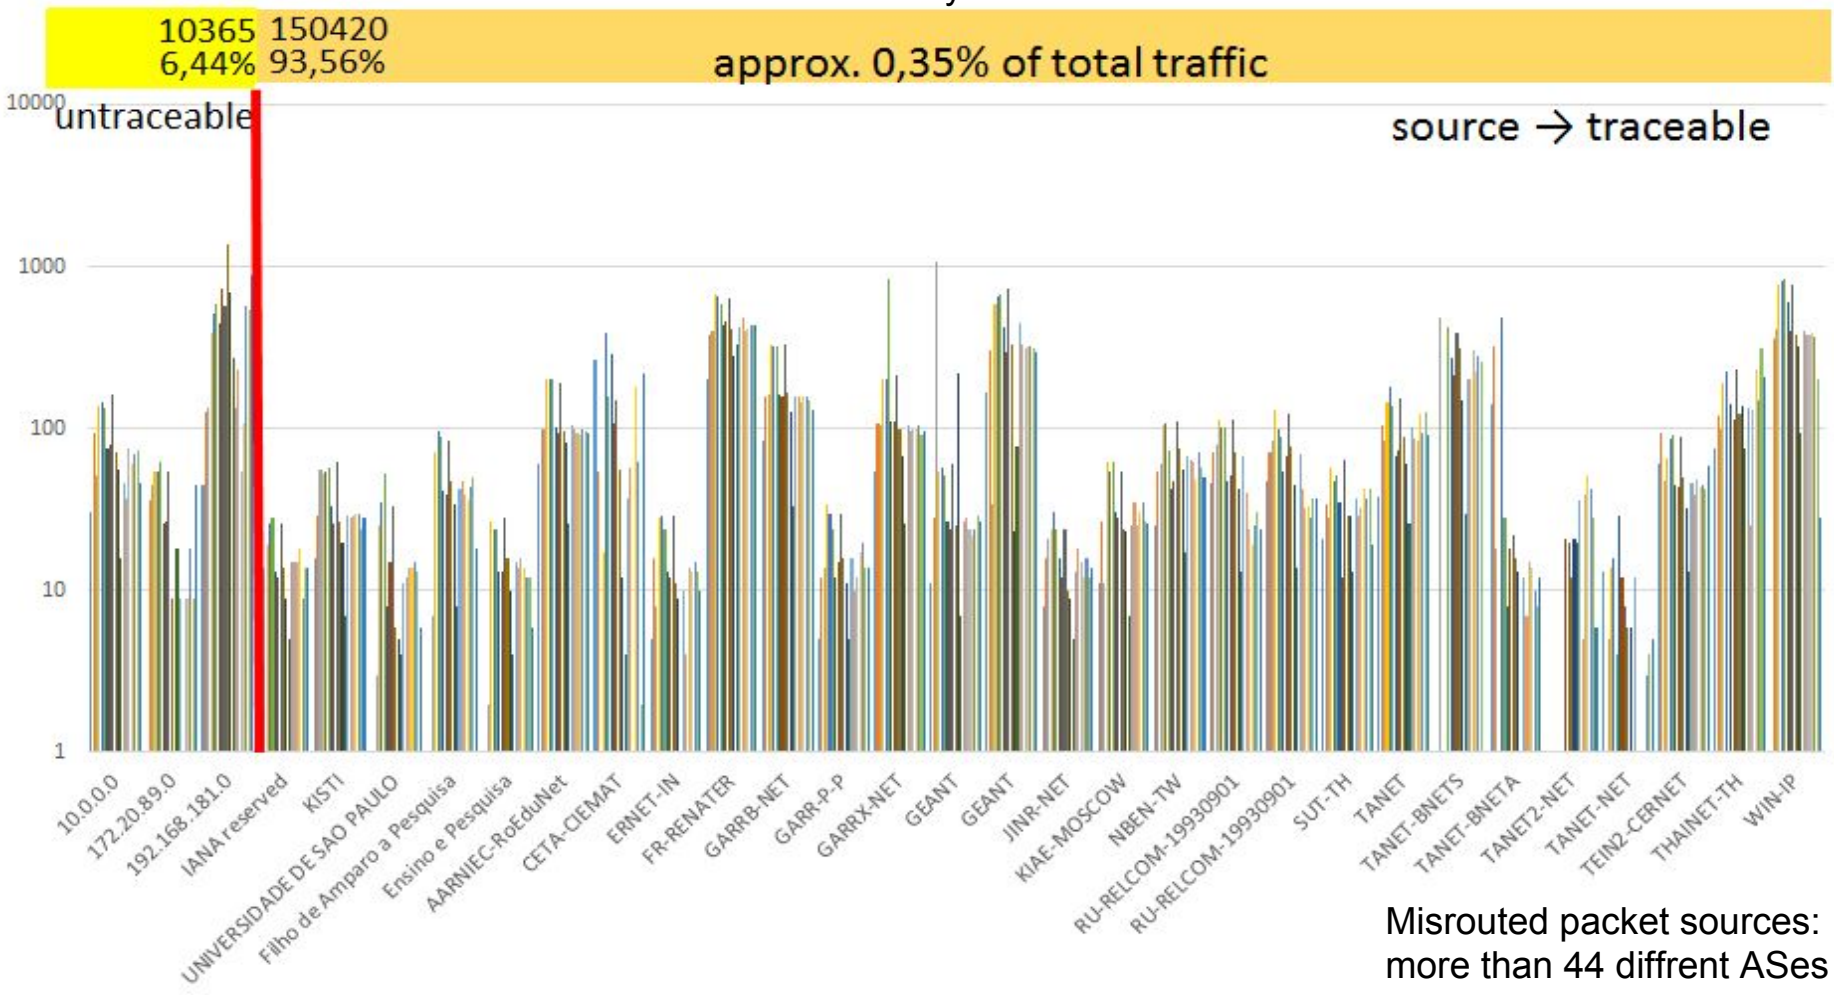

Misrouted packet sources:
more than 44 diffrent ASes

measurements over three weeks

20180912 20180913 20180914 20180915 20180916 20180917 20180918 20180919 20180920 20180921
20180922 20180926 20180927 20180928 20180929 20180930 20181001 20181002 20181003

mis-routed IPv6 packets

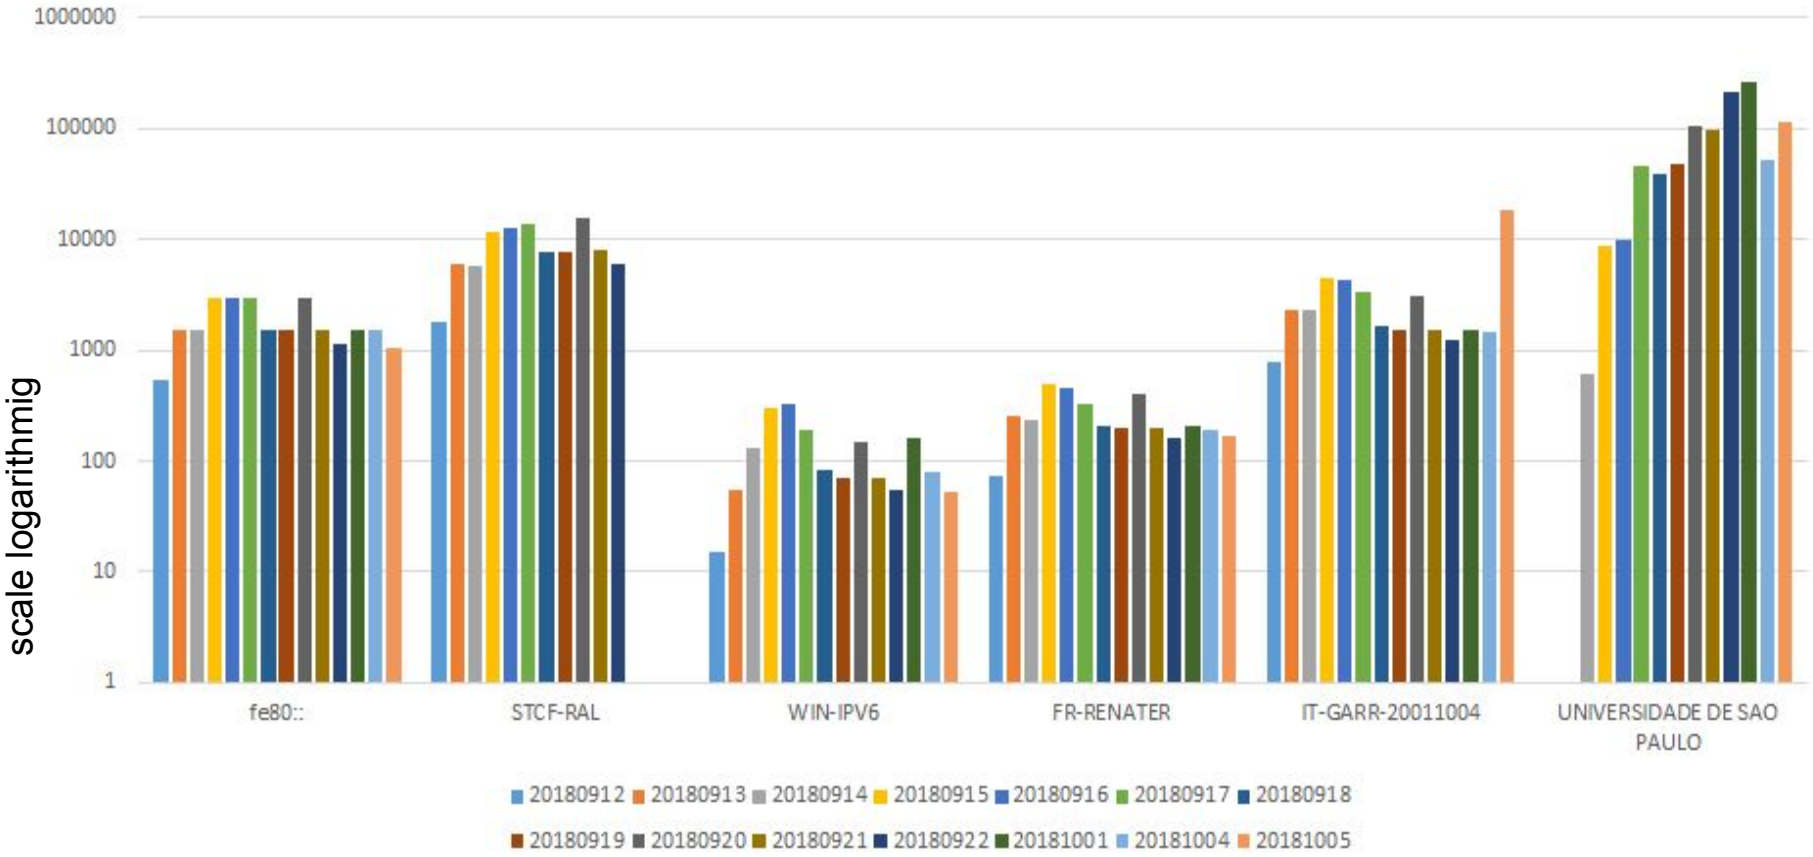

- a view sources only
- high packet rate
- while only perfsonar server at DE-KIT dual-stack

ICMP / none ICMP

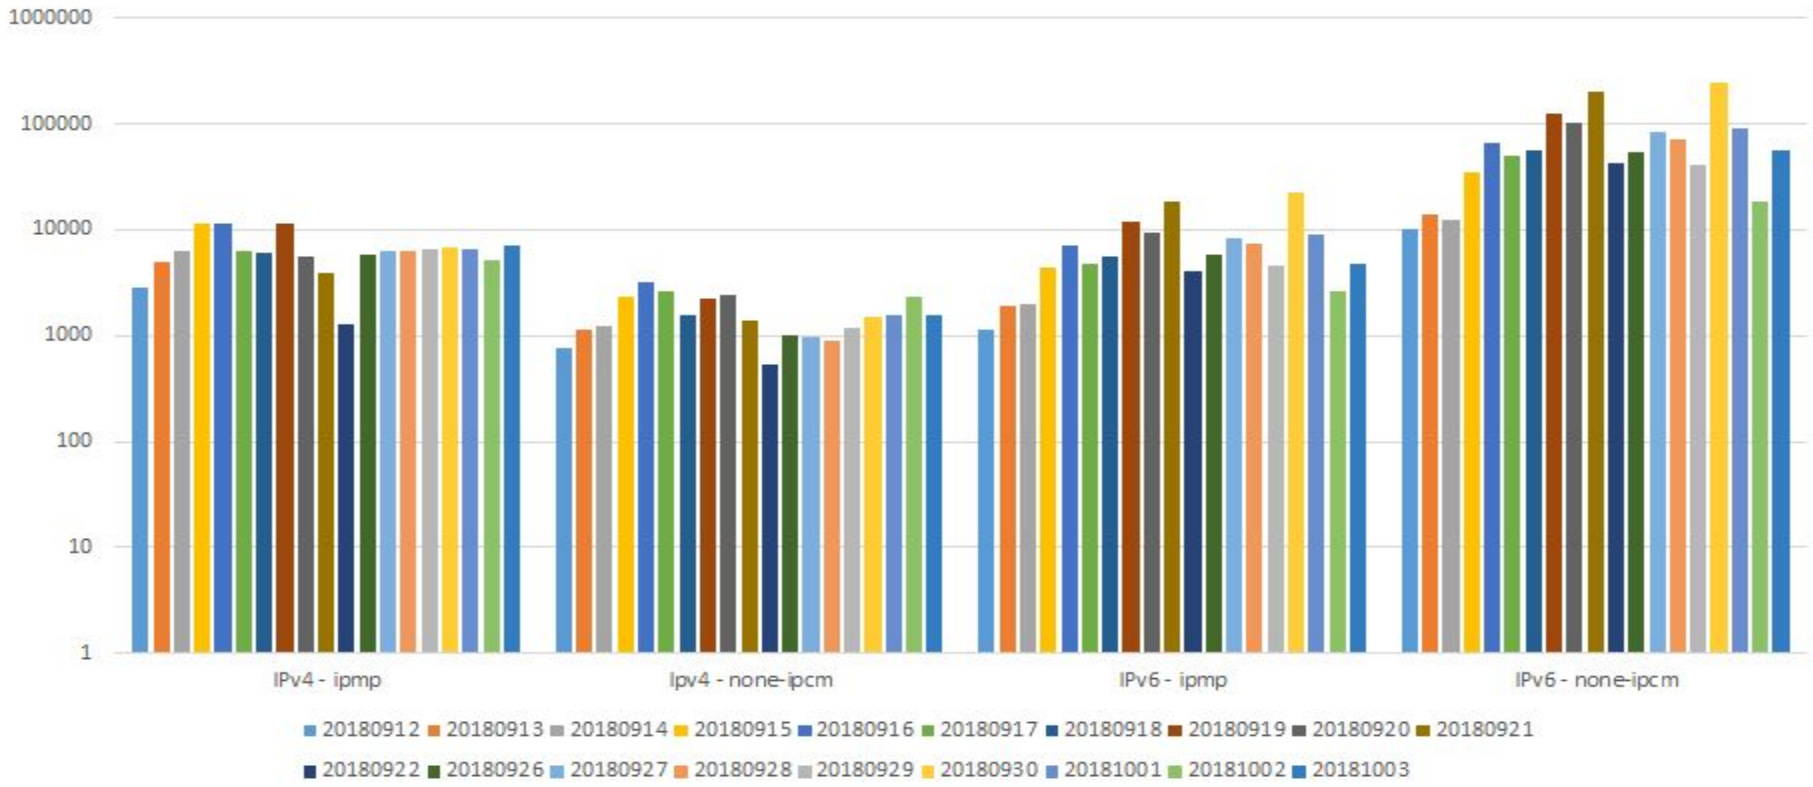

IPv4 : approx. half / half
IPv6 : none ICMP factor 1000 higher
→ further investigation necessary

ESnet monitoring

Unroutable Source Addresses by percentage

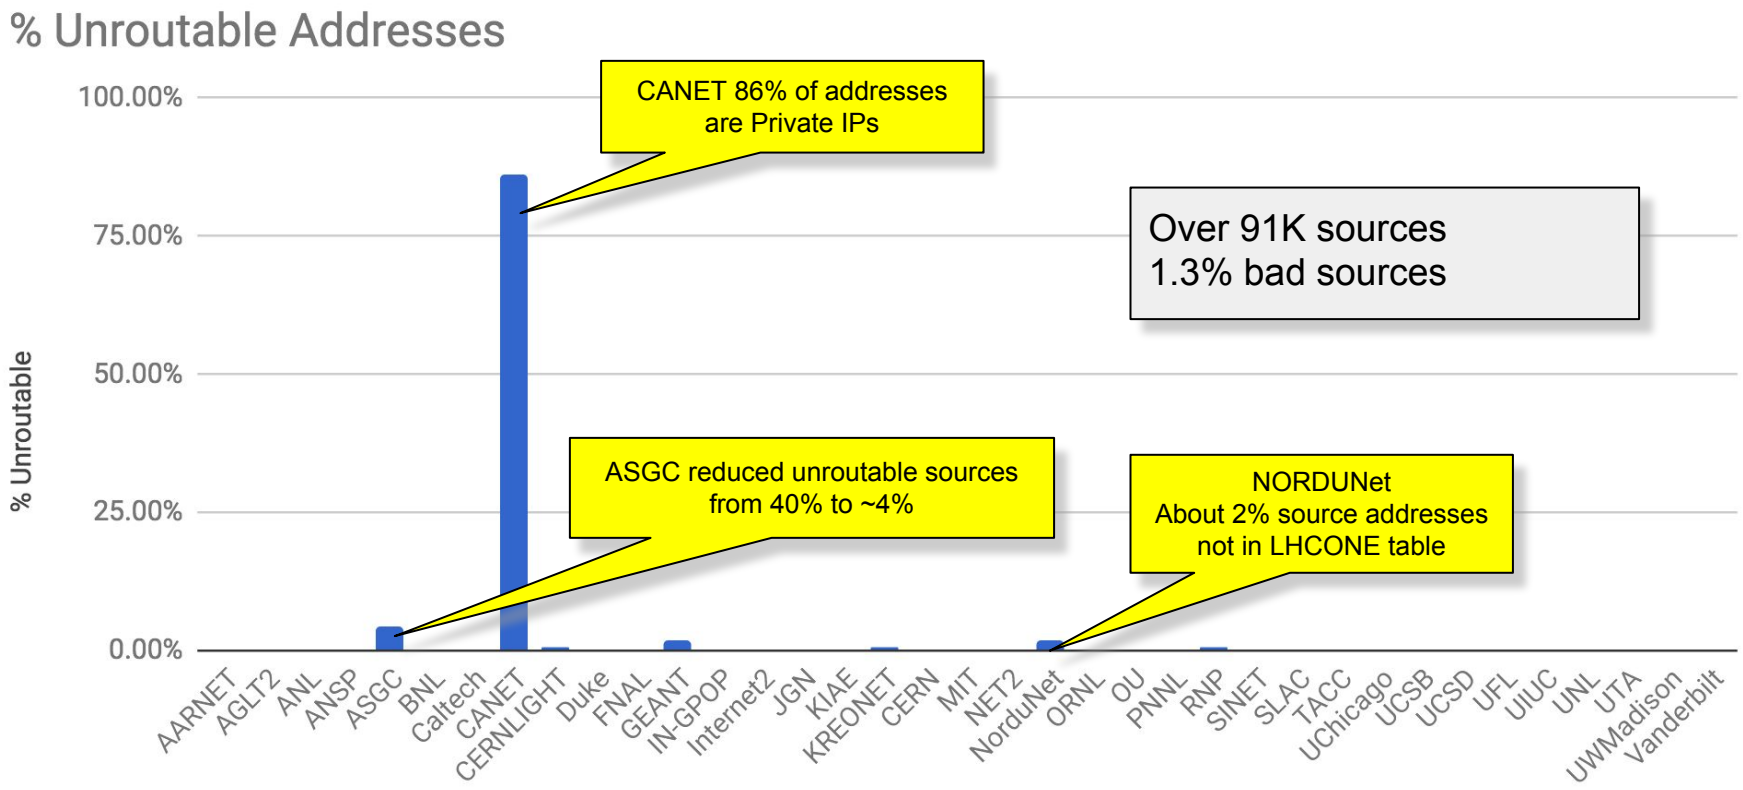

Bogon filters would block 86% of the unroutable source addresses injected by CANET into LHCONE

ESnet monitoring

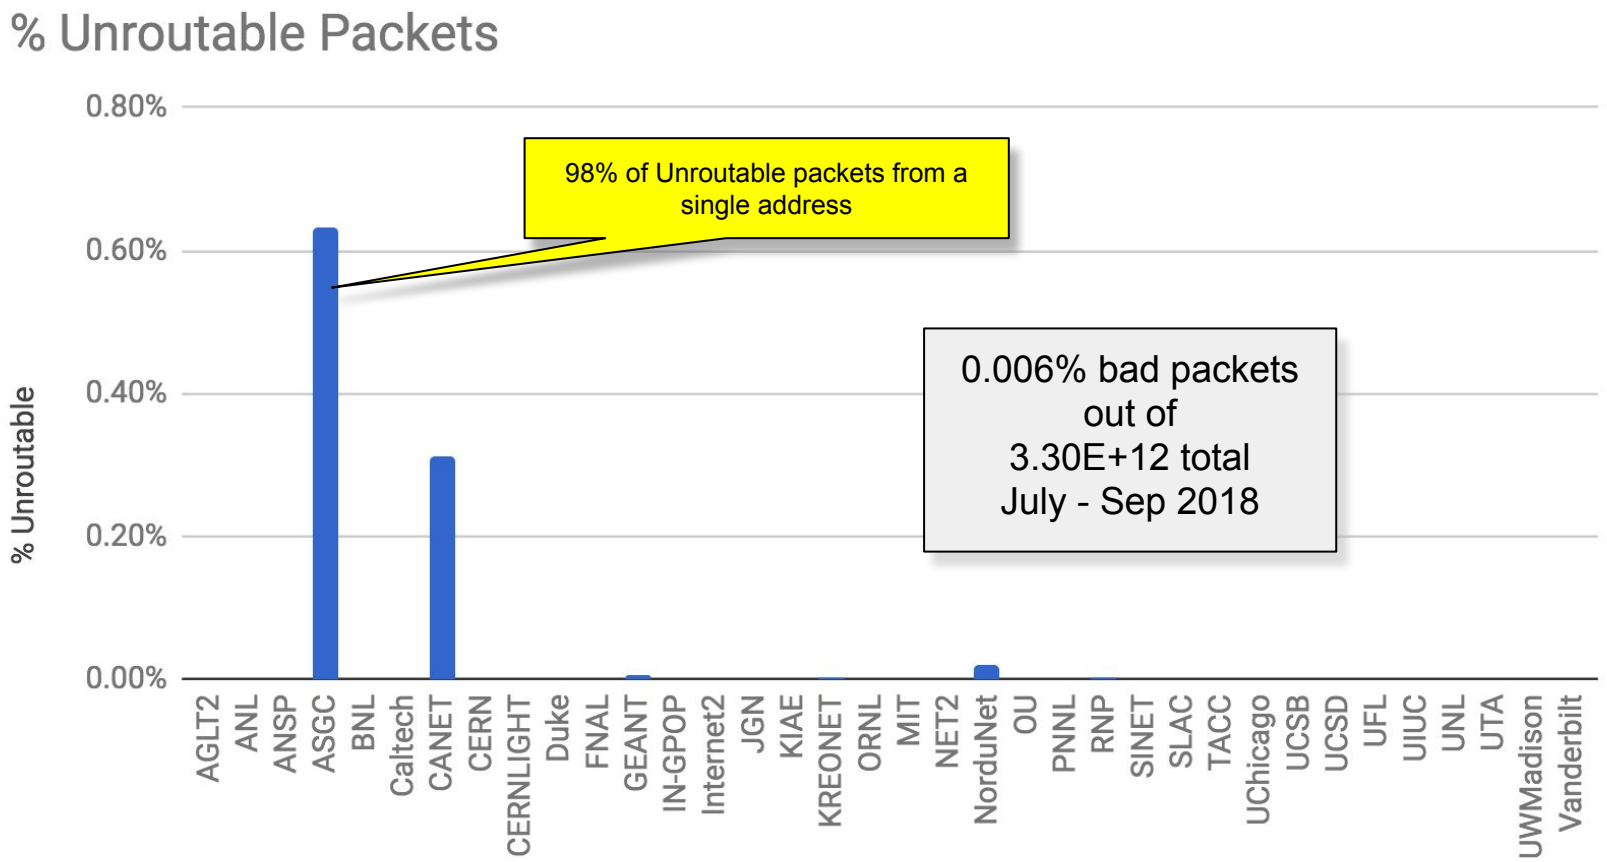

ASGC transits TANET2-TW, JGN, SINET, this community reduced their unroutable packet percentage from 44% in March 2018 to 0.63% today.

Within the NREN domain

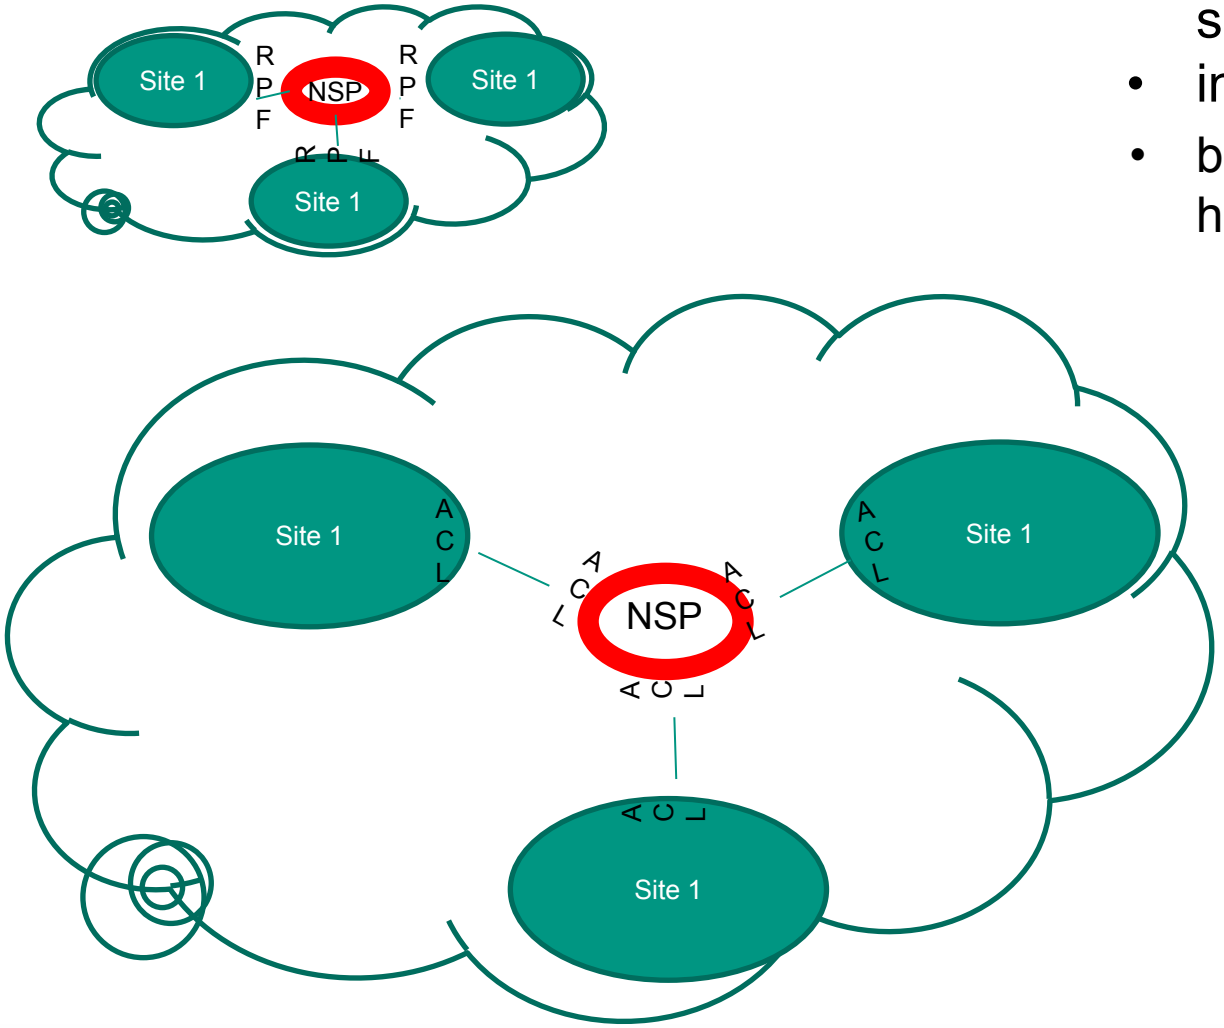

- ACL filter at connected sites
- in both directions
- but keep in mind: Is only half of the solution?
  - Verify that sites content are AUP compliant
  - Educate the connected site
  - Workout a AUP compliant configuration with the connected site

Edge Filtering Special Case

L3 Network Exchange Fabrics

An exchange is like an NSP:

- BGP import filtering
- Packet filtering
- Community based BGP filtering

An exchange is like a site:

- Require the full LHCONE table via a transit NSP
- Packet filters are configured and require maintenance

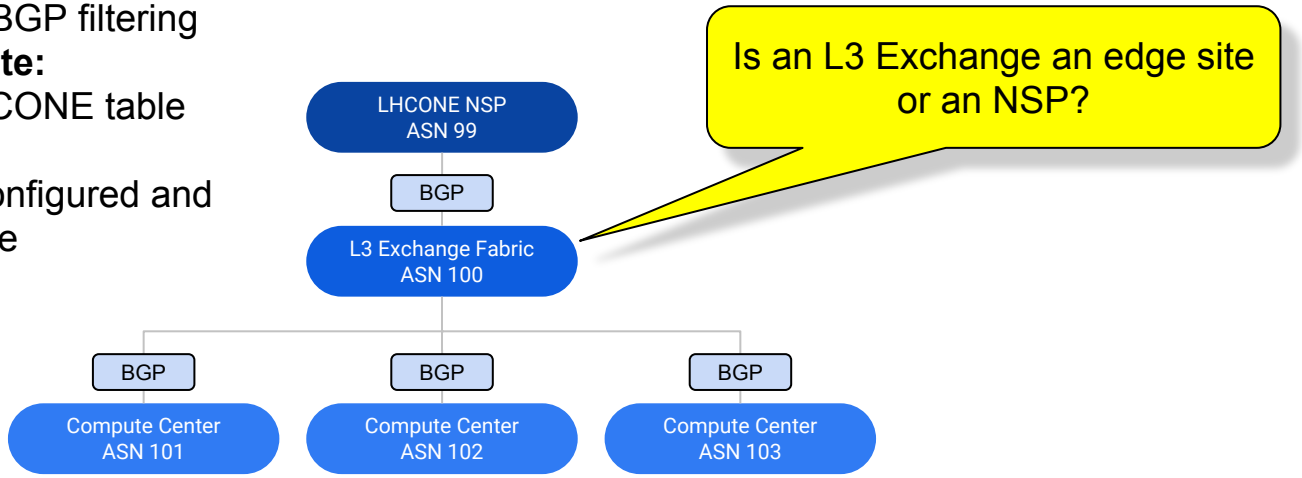

Indiana GigaPOP is a current ESnet example.
SOX is planned to be the second and will connect UFL, FSU and others.

- Will L3 Exchange Fabrics implement and maintain LHCONE specific services?
- Should there be an LHCONE defined role for these network organizations?
- How are they represented on the CERN LHCONE wiki?

Potential Courses of Action

To eliminate unroutable traffic:

Detection

- Regularly scheduled monitoring?
- Periodic NSP self run audits?

Prevention

- Edge Site filter configuration
  - RPF → too strict?
  - Templated policy & filter configuration

Information

- Regular AUP updates to address special cases
- Sharing configuration best practices

Conclusion / actions

- Fewer LHCONE unroutable source packets are being detected by ESnet since the March meeting
- Still room for improvement, particularly in the private IP ranges, which should be the easy packets to catch
- We will continue to monitor and report progress
- Exchanges are supporting LHCONE and need to be considered as an additional connection type in the LHCONE connection documents.

Questions
Suggestions
Discussion