Karlsruhe Institute of Technology

ESnet
ENERGY SCIENCES NETWORK

# LHCONE Edge Filtering Policy and Practice

**Bruno Hoeft / KIT**
**Michael O'Connor / ESnet**
**Richard Cziva / 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

## ESnet

- Three months of ESnet netflow IPv4 & IPv6 sampling from Feb. 2019 - April 2019 for the following connected sites and peers

| | | | |
|---|---|---|---|
| aarnet | duke | ornl | transpac |
| aglt2 | flr | mit | uchicago |
| anl | fnal | net2 | ucsb |
| ansp | geant | nordunet | ucsd |
| asgc | ind-gpop | ou | uiuc |
| bnl | internet2 | pnnl | unl |
| caltech | JGN | rnp | uta |
| canet | kiae | sinet | uwmadison |
| cern | kreonet | slac | vanderbilt |
| cernlight | cern | tacc | |

**ESnet counted:**

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

* corrected for netflow sampling rate

# unroutable IPv4 LHCONE packet statistics

## DE-KIT Ingress Packet Filters

- Unsampled ingress filtering detected LHCONE route table misses from 44 diffrent IPv4 source locations and 3 private address areas

**Private IP destinations: 10.0.0.0/8, 172.16.0.0/12 192.168.0.0/16**

**IPv4** *previous (Oct. 2018):*

Total : **1.044.471** (4 weeks) *160.785 (19 days)*

Packets per day: **37.302** *8.462*

- priv
- AARNIEC-RoEduNet
- GEANT
- GARRX-NET
- IANA - reserved
- TANET-NET
- T-MCU.EDU.TW-NET
- TANET-NET Taiwan
- University of Toronto
- REDIRIS
- KREONet-KR
- TANET-BNETS Taiwan
- TANET
- T-NCU.EDU.TW-NET
- TANET
- T-NTHU.EDU.TW-NET
- TANET Taiwan
- TANET-BNETA
- T-NSYSU.EDU.TW-NET

- TANET-BNETA Taiwan
- TRIUMF
- UNIVERSIDADE DE SAO PAULO
- RRC-KIAE-Moscow
- CAS-TCZ
- Indiana University
- JINR-NET
- Kasetsart University, Thailand
- TANET-B T-HCRC.EDU.TW-NET
- WIN-IP
- Inst. Nat. de Physique Nucleaire
- CZ-RELCOM-19930901
- GARR-P-P
- INFNNET-LNF
- GARRB-NET
- FR-RENATER
- Associação Rede Nacional (BR)

- Fundação Carlos Chagas Filho (BR)
- ERNET-IN
- TEIN2-CERNET
- THAINET-TH
- SUT-TH
- CANARIE
- CHINANET-FJ
- SAVECOM SAVECOM-NET
- AIT-TH
- TFN-NET TAI-SHIN-NET
- CIECCHQ-CN
- GZIN

# unroutable IPv6 LHCONE packet statistics

IPv6

| | | *previous:* |
|---|---|---|
| Total packets | : **1.376.338** (4 weeks) | *1.517.209 (19 days)* |
| Packets per day: | **35.290** | *79.853* |

33 different ipv6 sources + private (link local)

- fe80::
- RWTH Aachen
- Imperial College London
- UNIVERSIDADE DE SAO PAULO
- IJS-IPv6-NET - Ljubljana
- CERNET
- SINET-JPNIC
- UNI Michigan
- DFN WIN-IPV6
- GR-GRNET-19991208
- FR-IN2P3-LAL-ORSAY
- FR-IN2P3-LLR-PALAISEAU
- FR-CEA-SACLAY-GRILLES
- FR-IN2P3-LPNHE-PARIS
- FR-IN2P3-LPC-CLERMONT-AUBIERE
- FR-IN2P3-LAPP-ANNECY
- FR-IN2P3-CPPM-MARSEILLE
- FR-RENATER
- PNPI
- RU-ROSNIIROS-20180806
- CAS-PRG-6TCZ
- ES-REDIRIS-20010521
- UAM
- CIEMAT
- IT-GARR-20011004
- NL-GEANT-20020131
- UK-GEANT-20020131
- RoEduNet-IPv6-NET-1
- ASGC-NET
- IHEP-IPv6
- UNI of Nebraska-Lincoln
- Bcnet Vancouver
- VANDERBILT

# All counted packets (44 Sites)

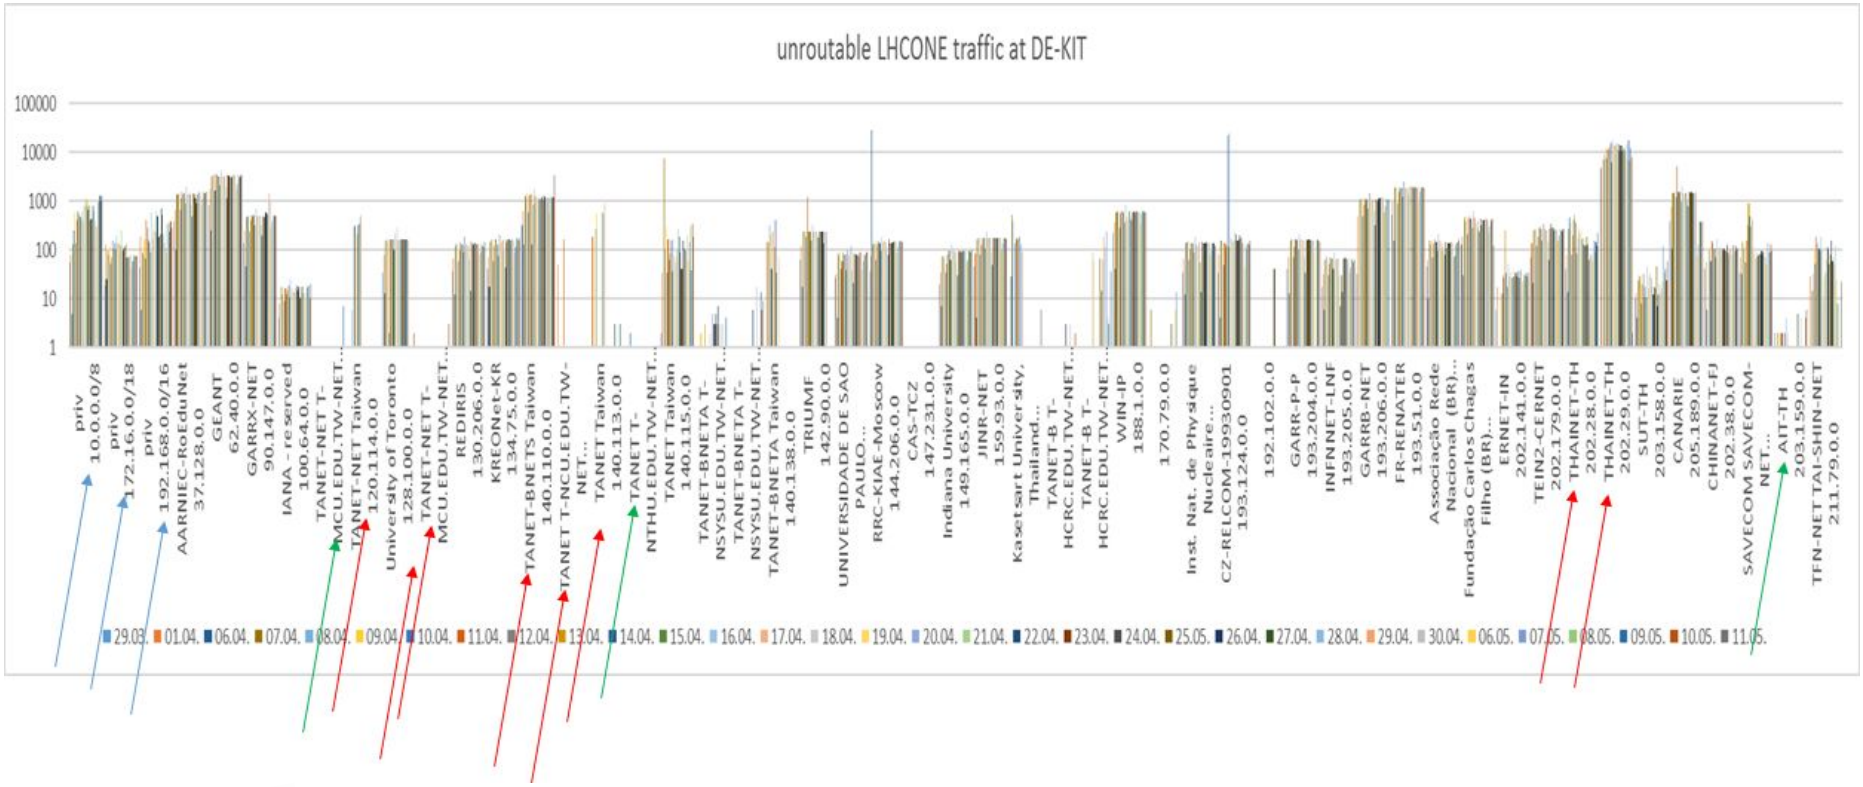


private

grouping different CIDR of one organization

removing sites with less than 1000 unroutable packets over four weeks

# unroutable IPv4 packets

Reduced to 30 different sites only

-Combined private address areas

-Removed sites with less than 1000 packets (per month)

-Pull different subnets of one site together

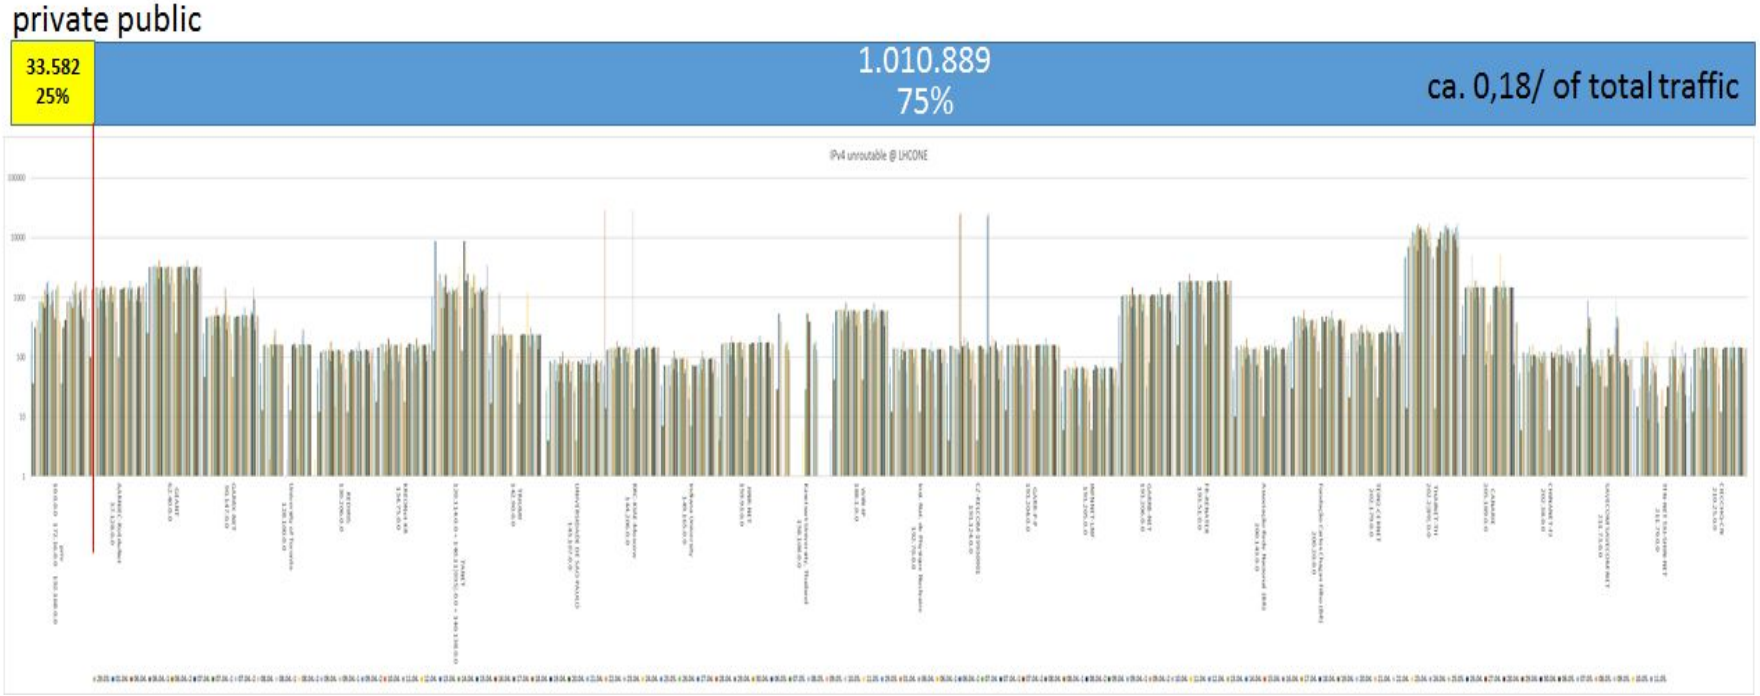


total : 1.010.889 + 33582 = 1.044.471

packts per day: **37.302**

ICMP approx. 43% = 459.508

none ICMP = 584.963

# unroutable IPv4 packets

Reduced to 30 different sites only

-Combined private address areas

-Removed sites with less than 1000 packets (per month)

-Pull different subnets of one site together

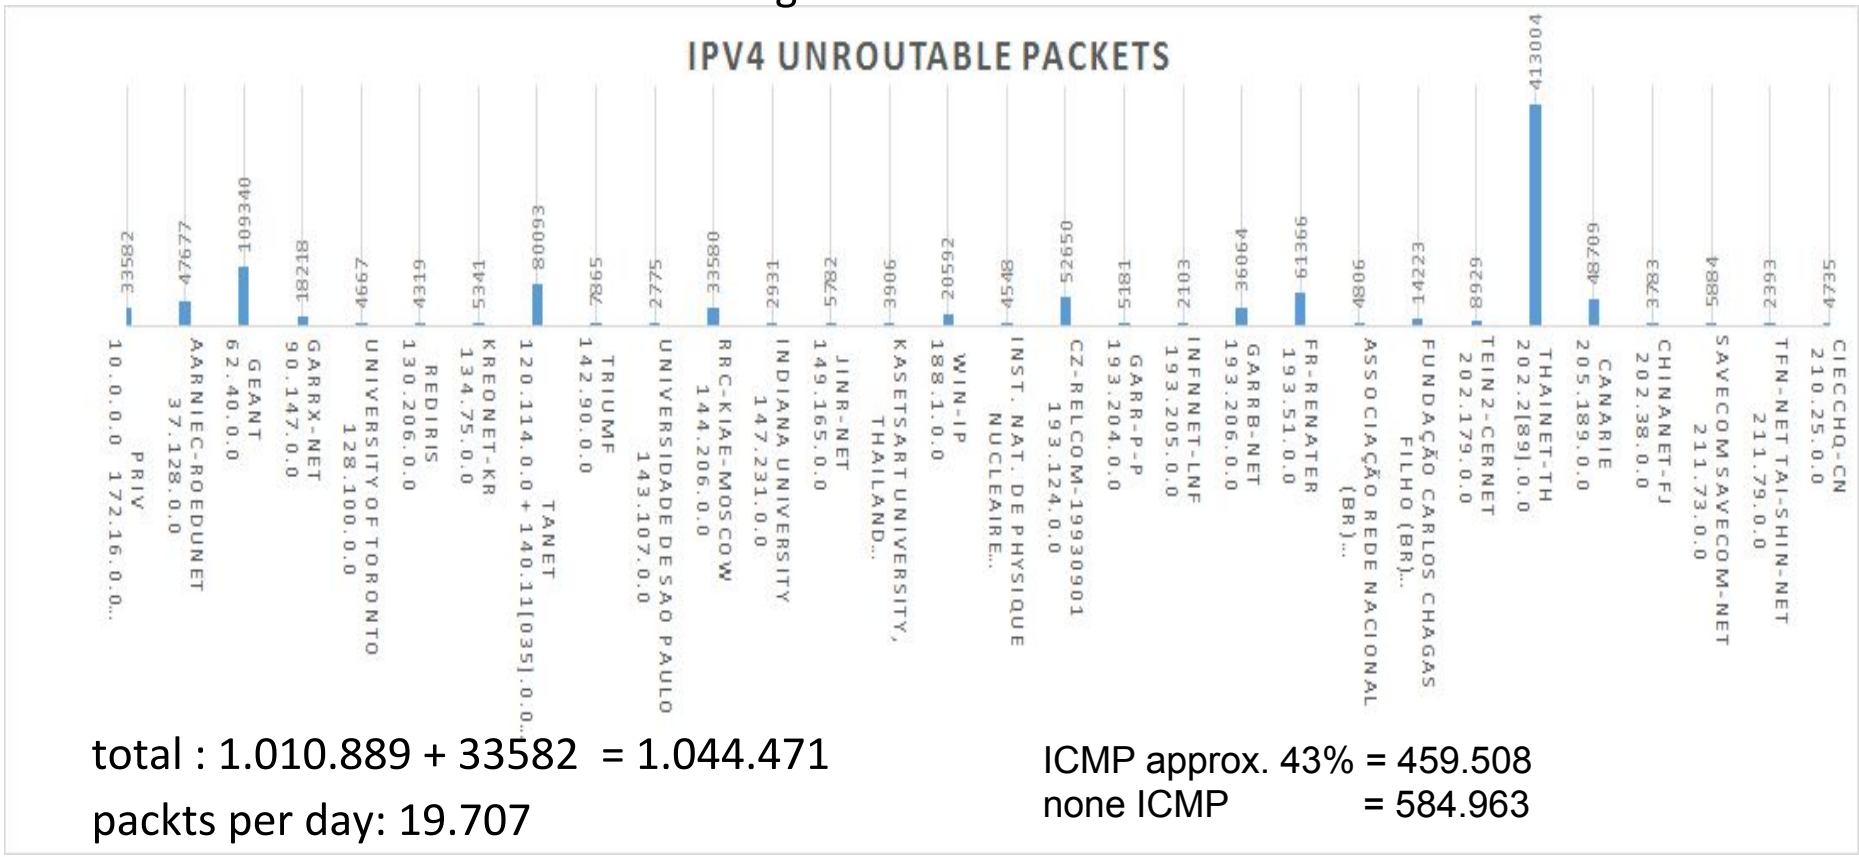


total : 1.010.889 + 33582 = 1.044.471

packts per day: 19.707

ICMP approx. 43% = 459.508

none ICMP = 584.963

# unroutable IPv4 packets

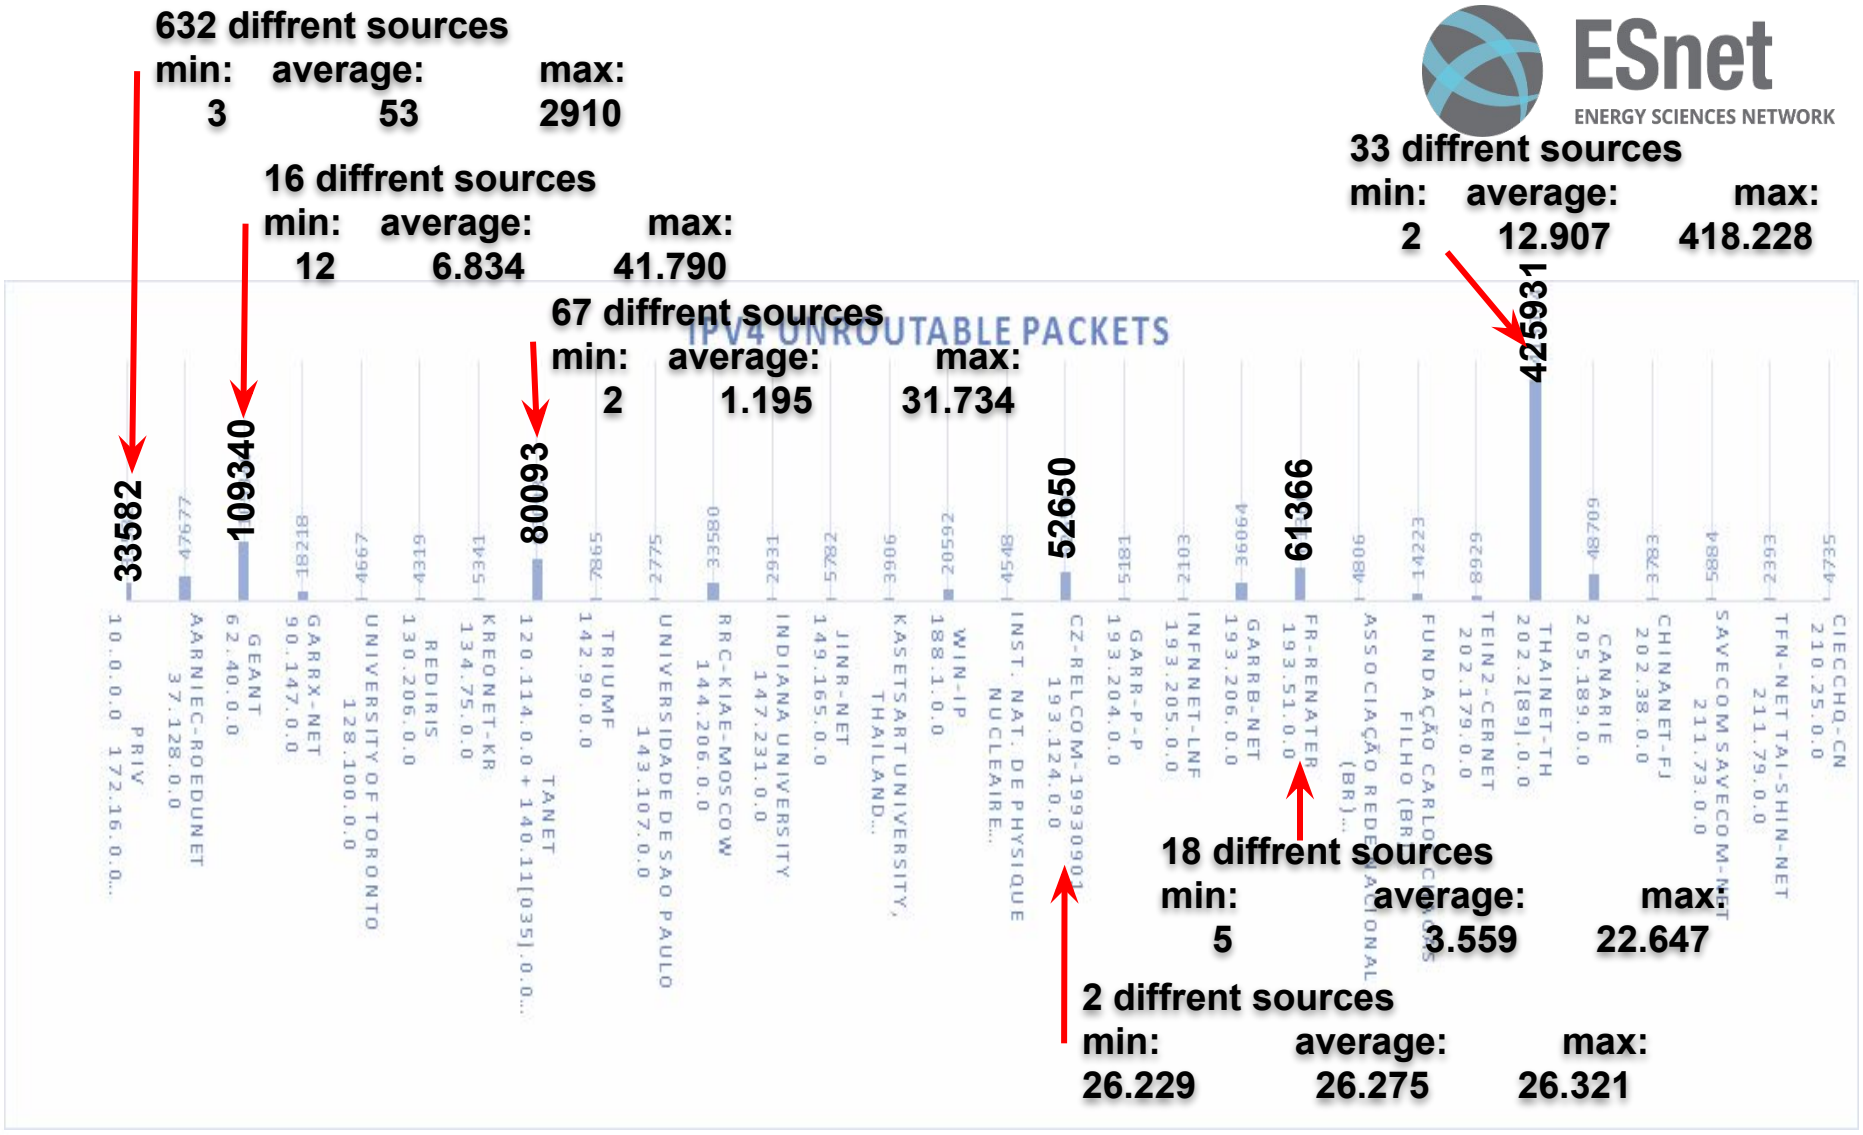

# All captured IPv6 unroutable packkts

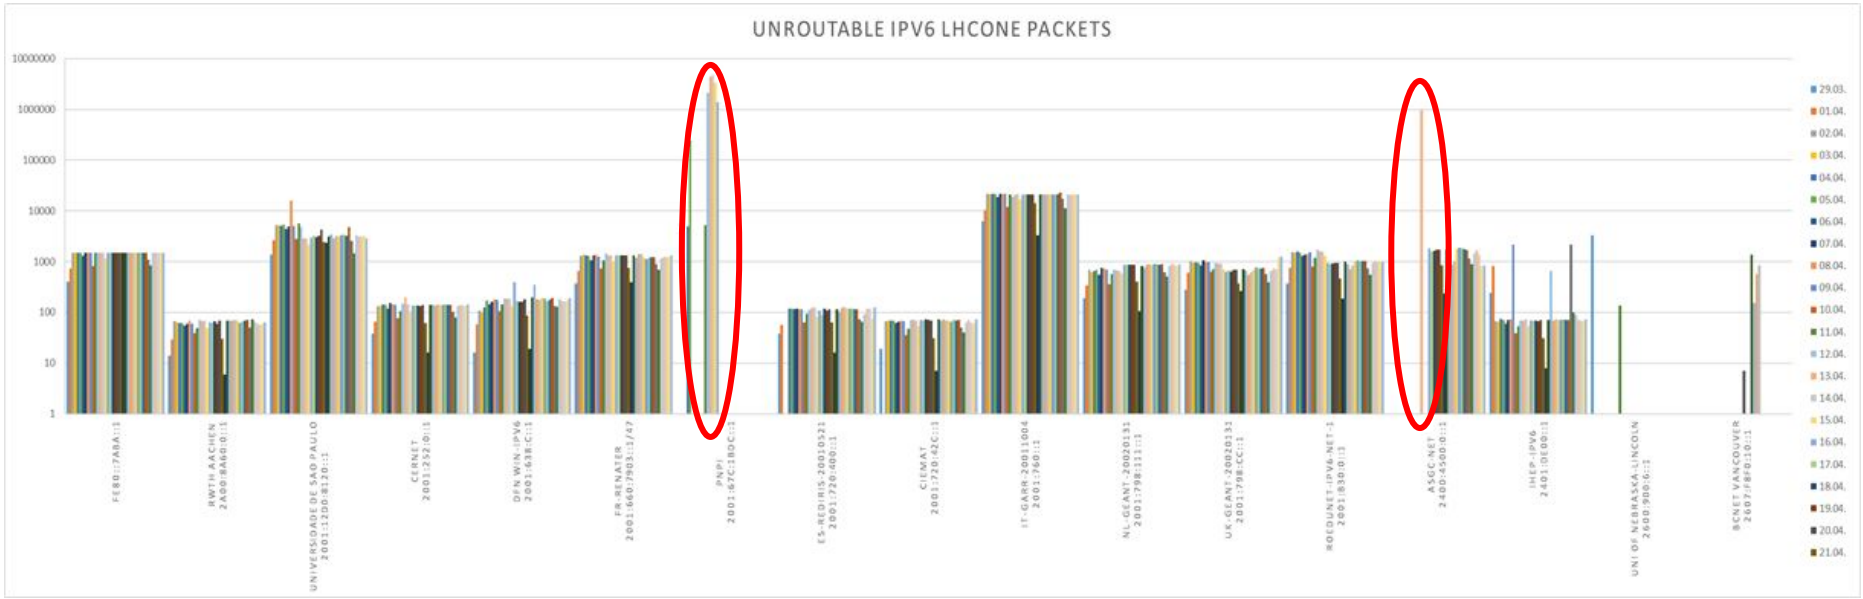


Total : 18.658.434

Per day: 478.421

18 sources

# Adjusted IPv6 Filter : unroutable LHCONE packet

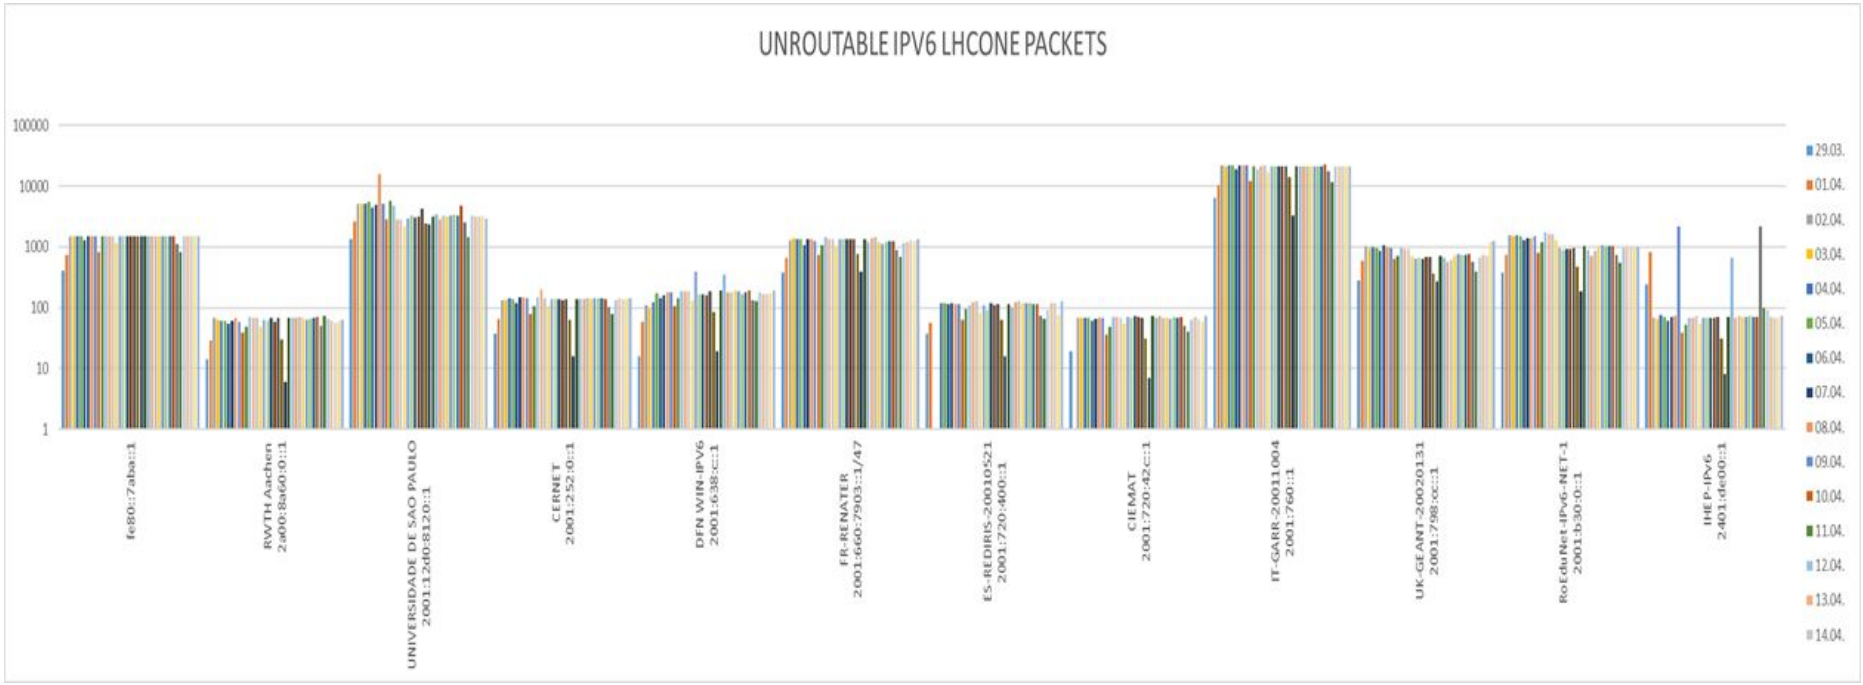


Total packets : 1.376.338

Packets per day: 35.290

12 sources

# IPv6 Filter : unroutable LHCONE packet

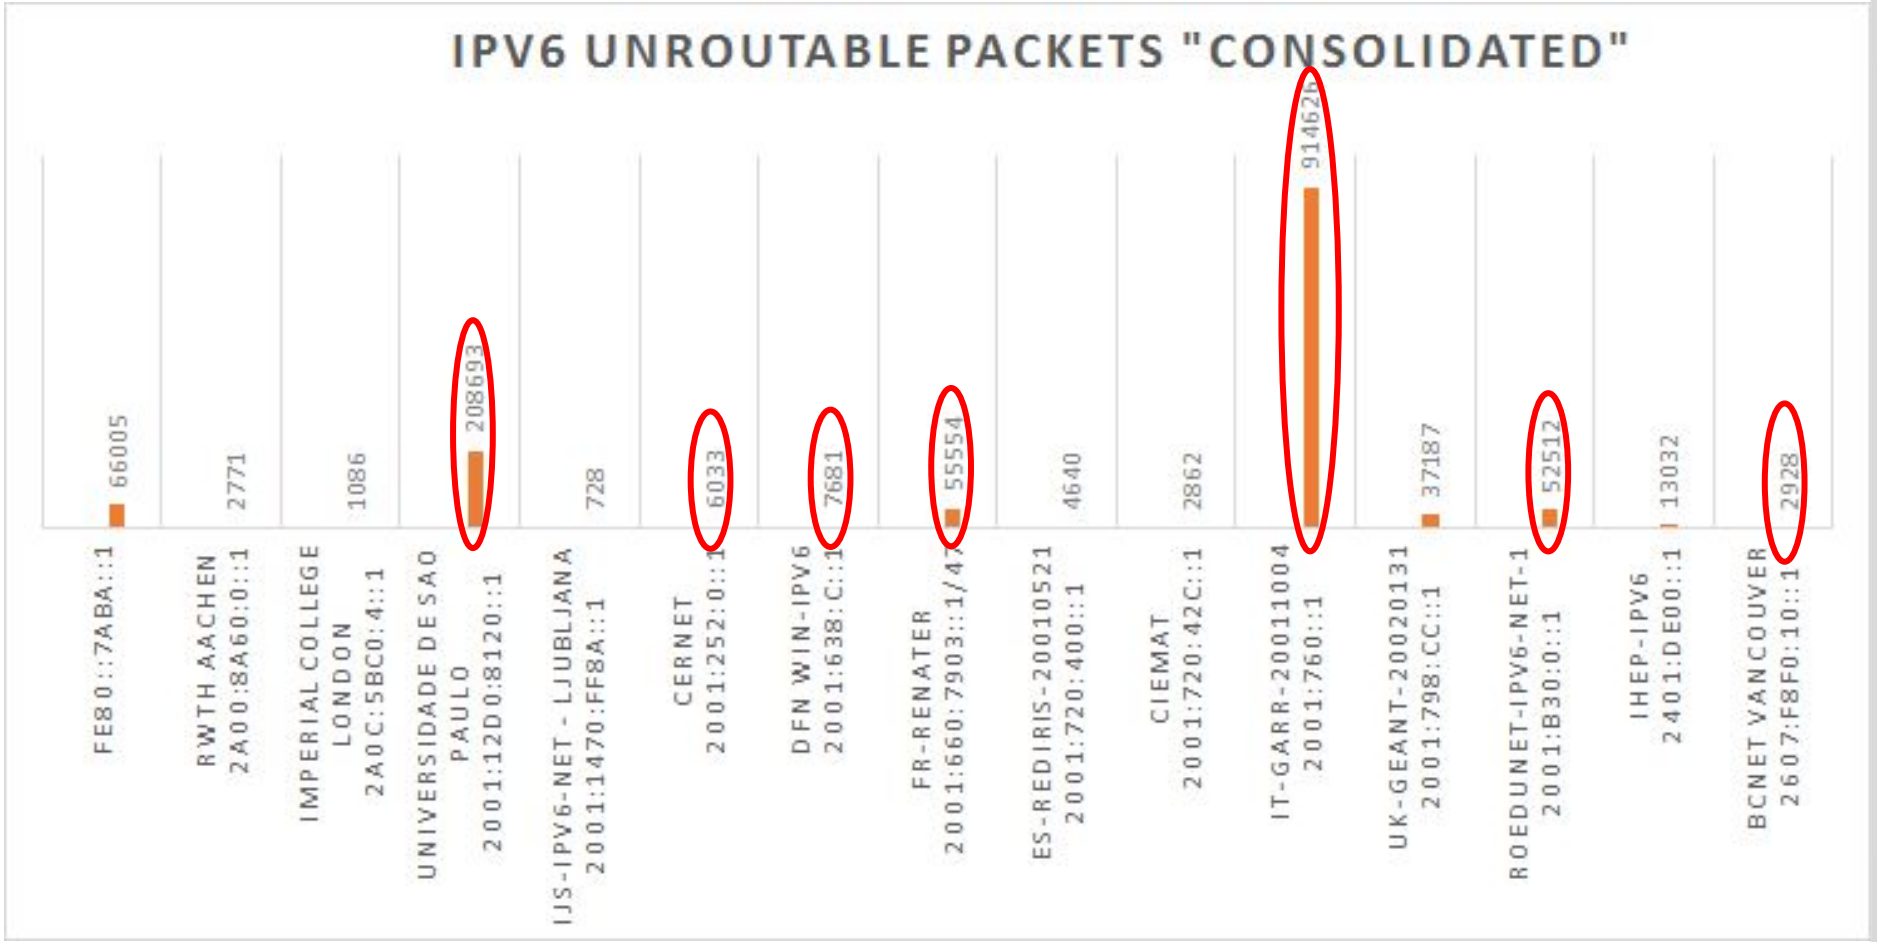

# IPv6 Filter : unroutable LHCONE packet

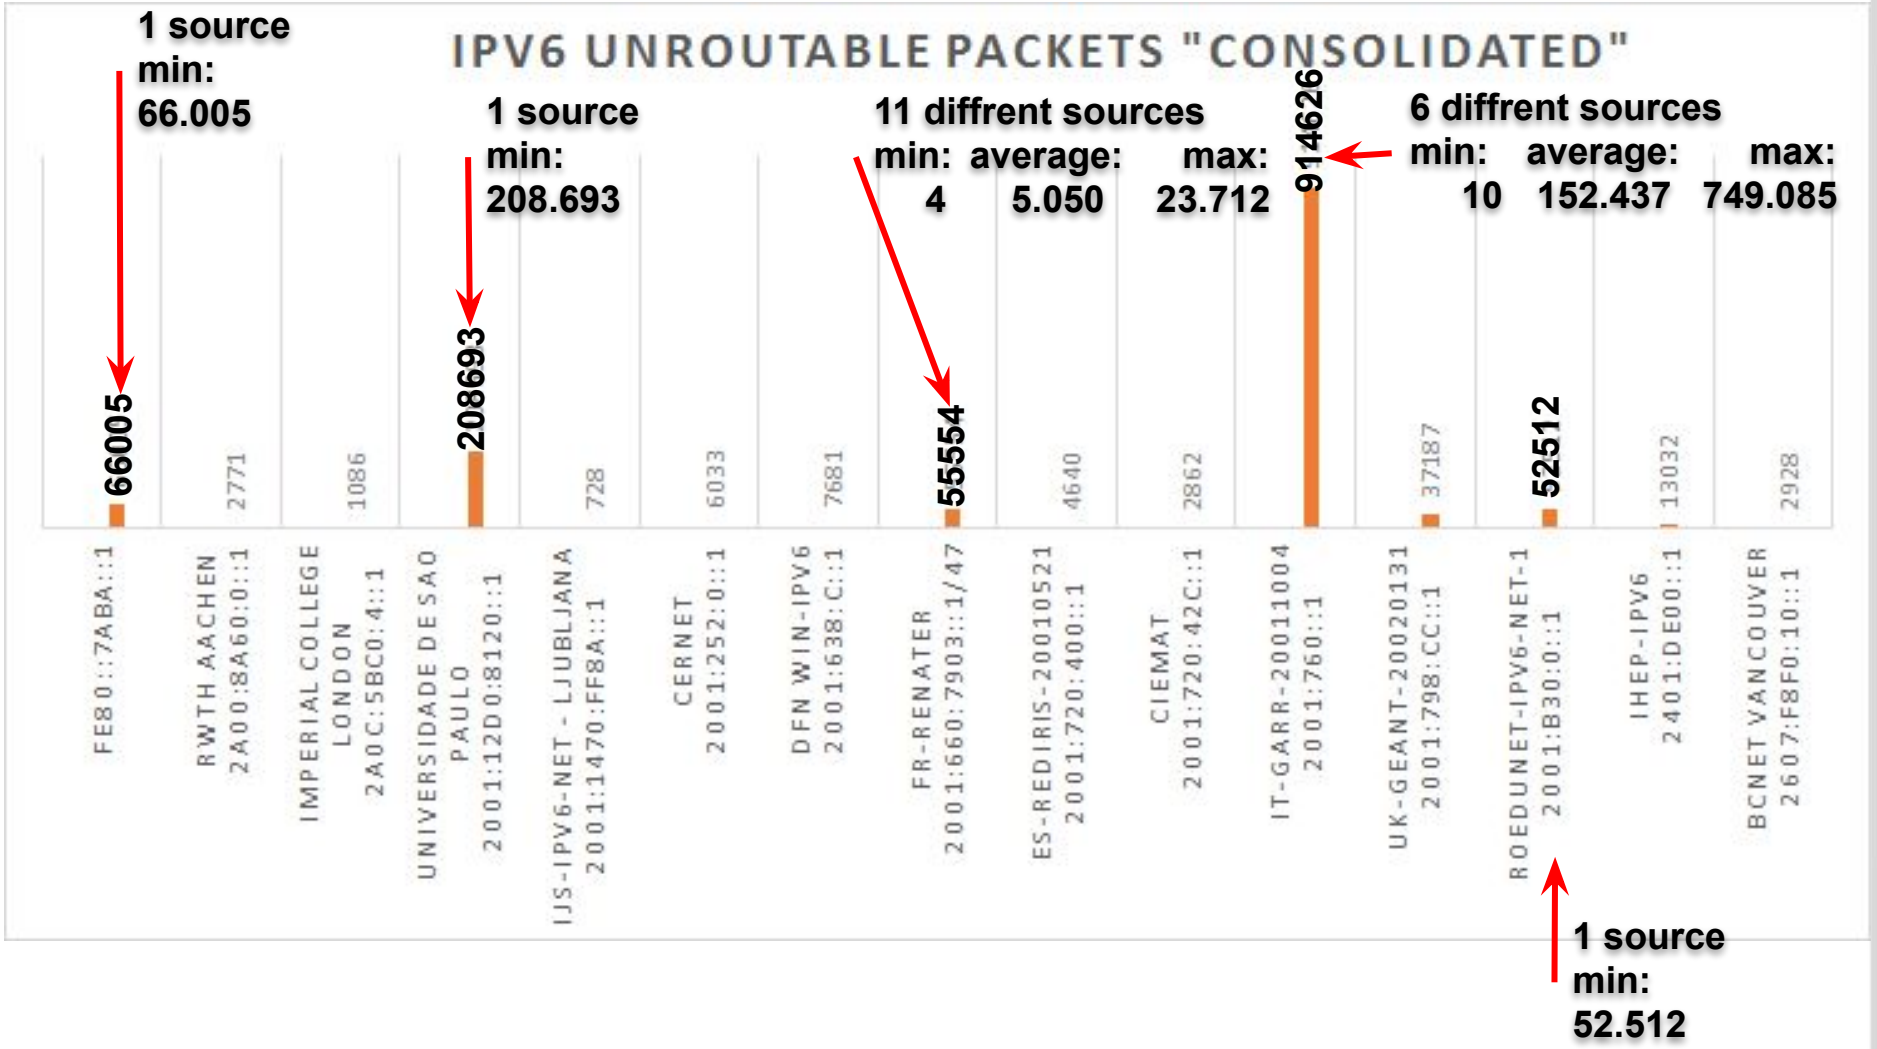

# ESnet monitoring

## Unroutable Source Addresses by percentage

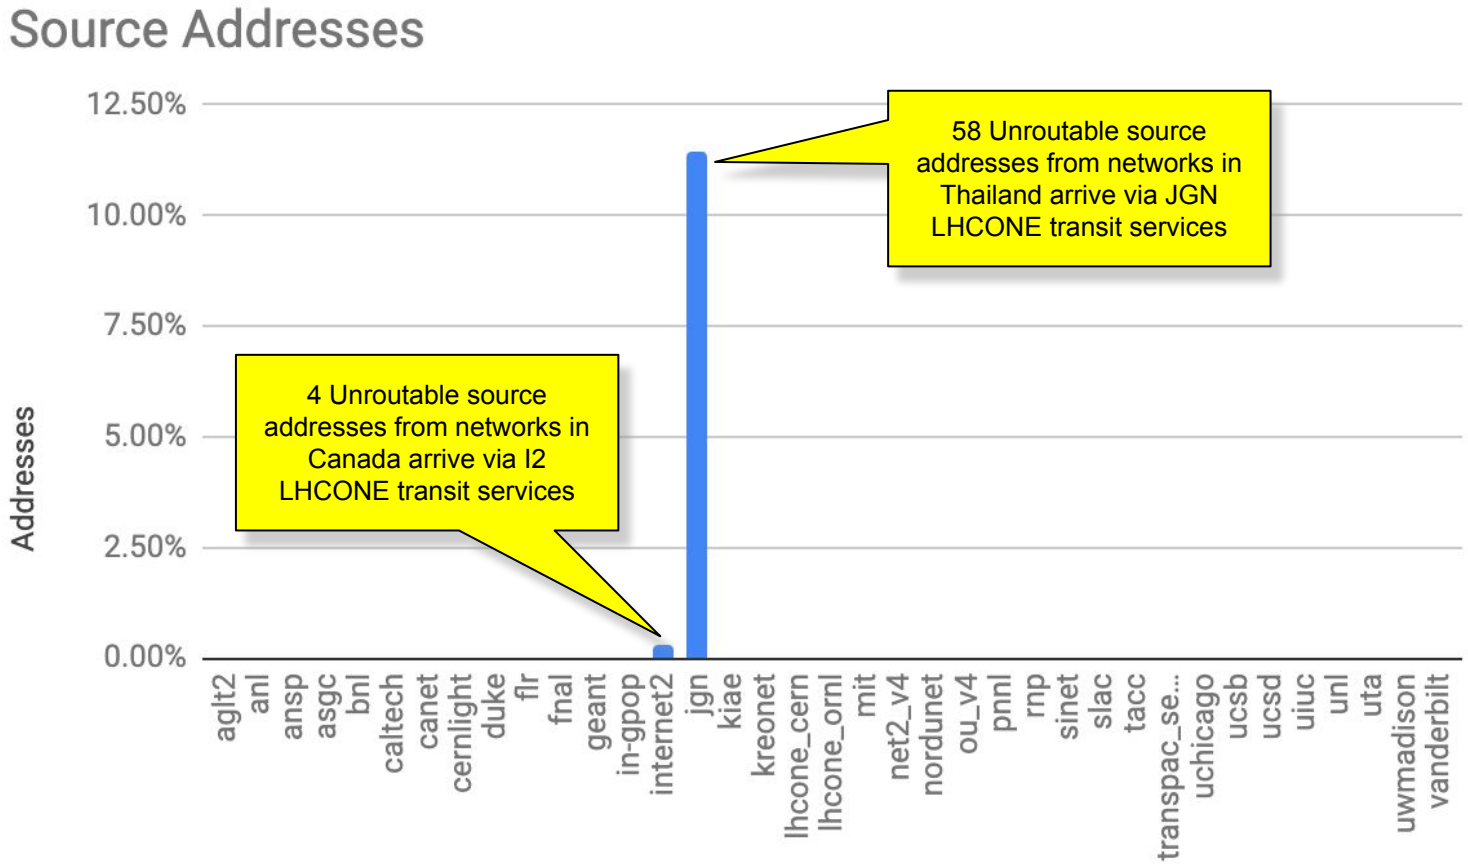


Very good results!!!

# ESnet monitoring

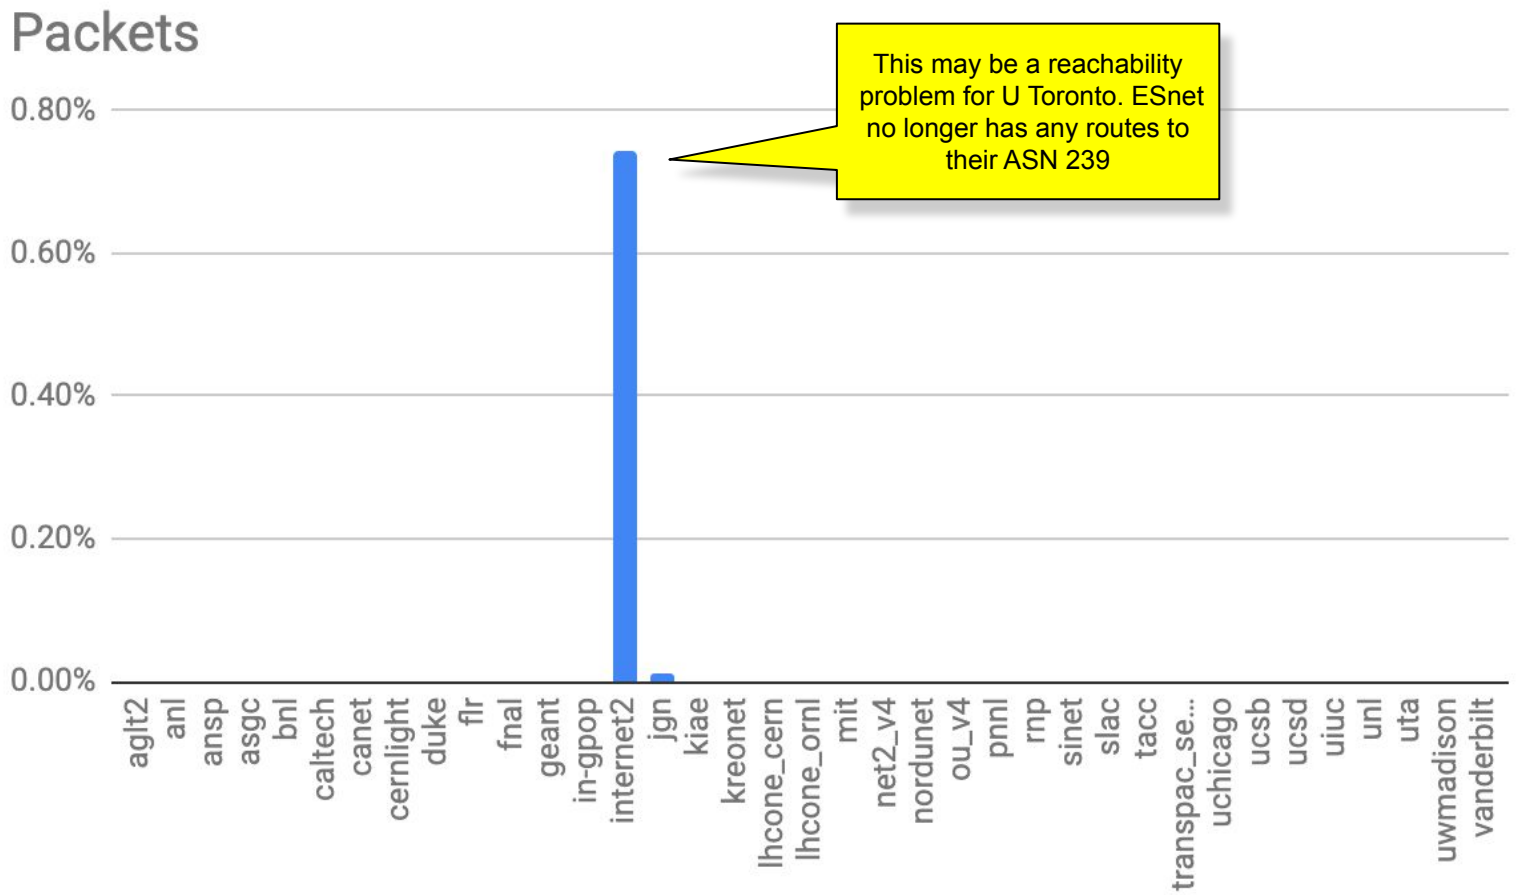


A small number of U Toronto hosts are transmitting large amounts of data on LHCONE. ESnet had routes for U Toronto in the past.

# Identifying the Source

## Example BGP prefix: 202.28.248.0/24

**LHCONE partial JGN BGP path tree**

JGN(17934) TEIN2-NORTH(24489) THAIREN(24475)

JGN(17934) TEIN2-NORTH(24489) THAIREN(24475) THAISARN(3836)

JGN(17934) TEIN2-NORTH(24489) THAIREN(24475) THAISARN(3836) NSTDA-TH-AS-AP(38296)

JGN(17934) TEIN2-NORTH(24489) THAIREN(24475) THAISARN(3836) PUBNET-TH-AS(7588)

JGN(17934) TEIN2-NORTH(24489) THAIREN(24475) UNINET-TH(836)

JGN(17934) TEIN2-NORTH(24489) THAIREN(24475) UNINET-TH(836) ERX-CHULANET(3839)

JGN(17934) TEIN2-NORTH(24489) THAIREN(24475) UNINET-TH(836) SUT-AS-AP(55545)

**Global Routing table AS path lookup:**

Chiang Mai University CMU-TH-AP(17479)

JGN(17934) TEIN2-NORTH(24489) THAIREN(24475) UNINET-TH(4621) CMU-TH-AP(17479)

# Within the NREN domain

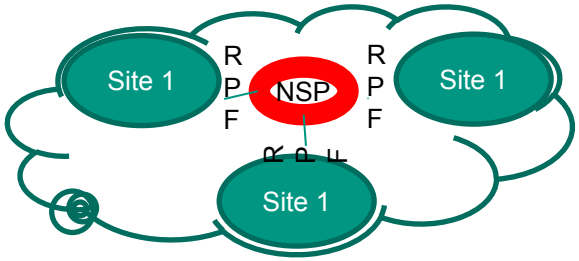


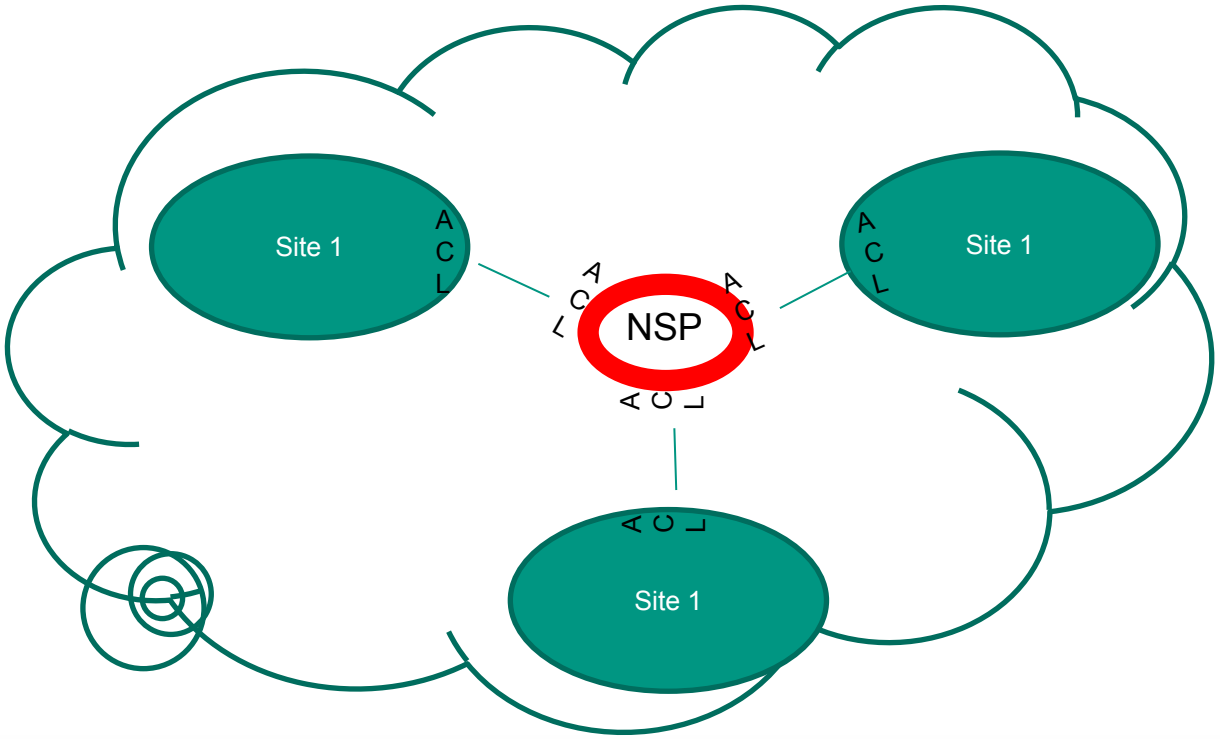


- ACL filter at connected sites
- in both directions
- but keep in mind: Is only half of the solution?
  - Verify that sites content are AUP compliant
  - Educate the connected site
  - Workout a AUP compliant configuration with the connected site

# 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

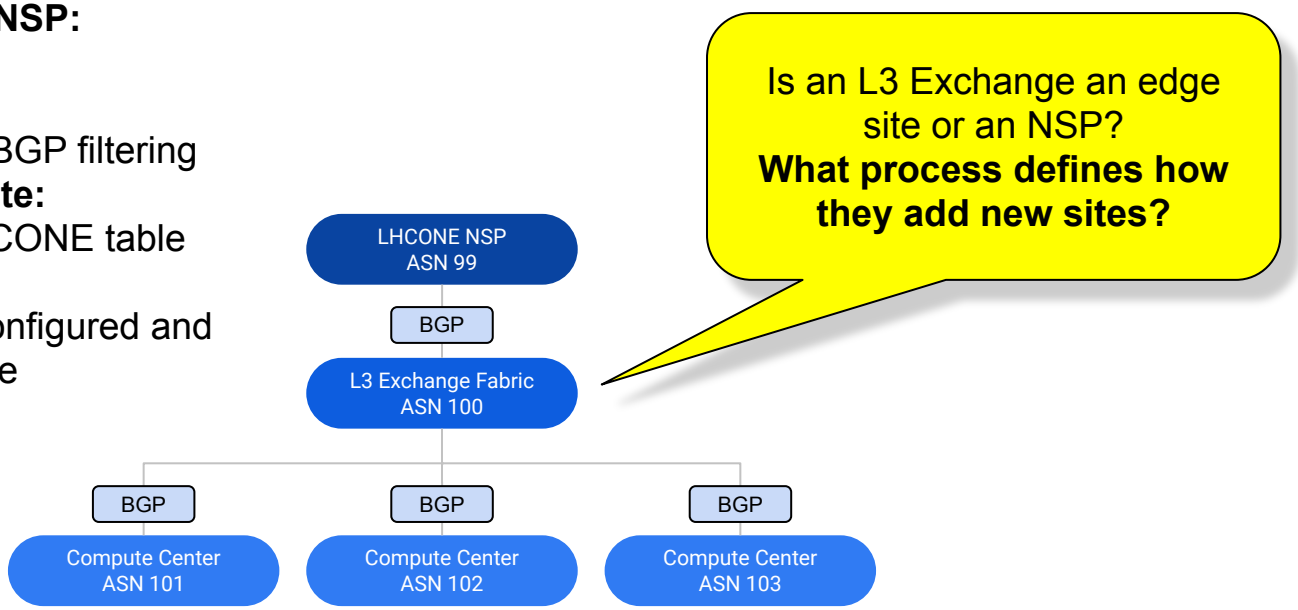


Indiana GigaPOP is a current ESnet example.
SOX is planned to be the second and will connect UFL, FSU and others.

- Will L3 Exchange Fabrics implement and maintain LHCONE specific services?
- Should there be an LHCONE defined role for these network organizations?
- Are they permitted to attach new sites?

# 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

# Project @ DE-KIT : unroutable packet -- web portal

- automate unroutable packet information gathering
  - → store into a organized and structured database (kibana)

  - visualise the data (elastic search / Kibana / grafana) with different levels
    - abstract overview
    - and zoomable into detailed view (up to source/dest. of a single packet)

this data shall be available for the LHCONE connected sites (but not for the world), one idea:

- community securing the data
- restrict access --> personal authentication enabled (via EGUgain)

project just started,
working on first results by the end of this year/

# Conclusion / actions

- LHCONE ingress filtering/control has improved dramatically since measurement began in Q1 2018!
- Growth in Asia has likely contributed to a small number of exceptions.
- Routing table inconsistency may also be a source of false positives.
- Exchanges and transit providers are being added to LHCONE, the community needs to provide better guidance to these providers.



# Questions
# Suggestions
# Discussion