

# BGP Communities: Even more Worms in the Routing Can

ACM IMC 2018, Boston, MA, USA

---

Florian Streibelt<sup>1</sup> <fstreibelt@mpi-inf.mpg.de>,  
Franziska Lichtblau<sup>1</sup>, Robert Beverly<sup>2</sup>, Cristel Pelsser<sup>3</sup>,  
Georgios Smaragdakis<sup>4</sup>, Randy Bush<sup>5</sup>, Anja Feldmann<sup>1</sup>

Nov 1, 2018

<sup>1</sup> Max Planck Institute for Informatics (MPII), <sup>2</sup> Naval Postgraduate School (NPS),

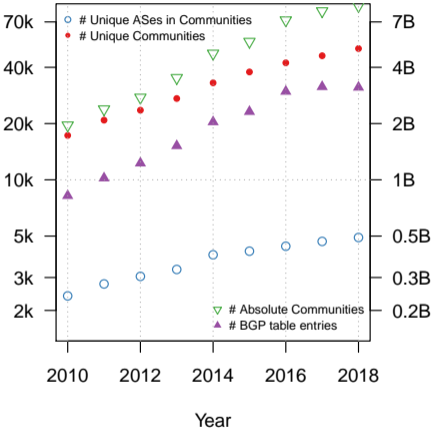
<sup>3</sup> University of Strasbourg, <sup>4</sup> TU Berlin (TUB), <sup>5</sup> Internet Initiative Japan (IIJ)

# Introduction

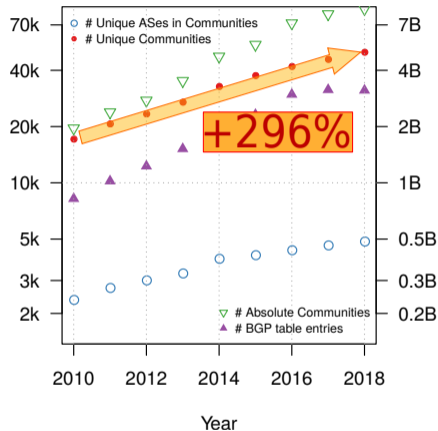
---

- We provide an analysis of BGP community propagation on the Internet
- We show that BGP communities (as used by operators to realize traffic management) can be used as attack vector
- We verify this via experiments in the lab as well as in the wild
- We provide some hints on the secure usage of BGP communities

# BGP Community usage is increasing




## BGP Community usage is increasing




Increasing usage warrants a closer look.

# BGP (Border Gateway Protocol)

## BGP (Border Gateway Protocol)




## BGP (Border Gateway Protocol)



- AS1 announces prefix p




## BGP (Border Gateway Protocol)



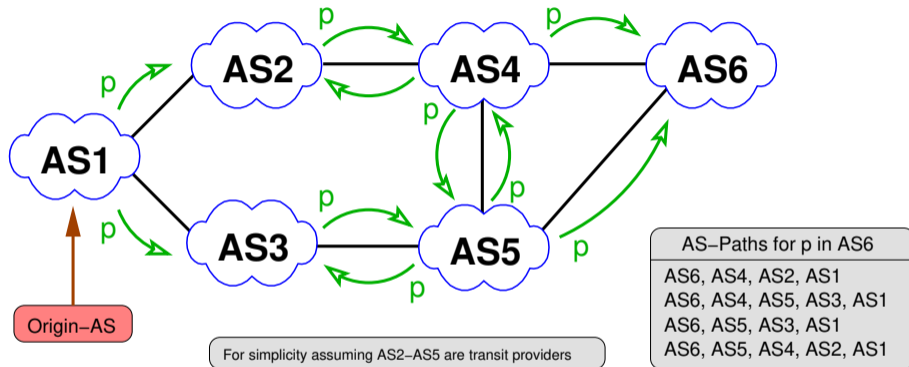
- AS1 announces prefix p, upstreams pickup p

# BGP (Border Gateway Protocol)




- AS1 announces prefix p, upstreams pickup p
- AS6 receives first announcements for p

# BGP (Border Gateway Protocol)



- AS1 announces prefix p, upstreams pickup p
- AS6 receives first announcements for p
- eventually AS6 sees multiple available paths for p


# BGP (Border Gateway Protocol)



## BGP

- BGP communicates reachability information
- Announcement messages also carry various attributes
- One of these attributes are BGP-Communities

# BGP Communities



- RFC 1997: Optional Attribute in BGP message (32 bit)
- By convention written *ASN:VALUE*
- ASN can be both sender or intended 'recipient'
- Every network decides the semantics behind the values
- New standard: Large Communities (96 bit), not yet widely deployed

# BGP Communities: Usage

## Informational Communities (Passive Semantics)

- Location tagging
- RTT tagging

## Action Communities (Active Semantics)

- Remote triggered blackholing
- Path prepending
- Local pref/MED
- Selective announcements

**Used by operators to realize policies.**

**Without documentation, you can not tell if a community is active or passive!**

## BGP Communities As Attack Vector?

Given the **increasing popularity** of BGP communities and the ability to **trigger actions** as well as **relay information**, one question arises:

To which extent can BGP communities be leveraged for attacks?

## Propagation behavior

- RFC 1997: Communities as a transitive optional attribute
- RFC 7454: Scrub own, forward foreign communities
- 14% of **transit** providers propagate received communities (2.2k of 15.5k)
- Ratio seems small, but AS graph is highly connected

**Still many people do not expect communities to propagate that widely.**



## Potential (for) misuse

- Propagated communities might trigger actions multiple AS-hops away
- No way of knowing if intended or not, e.g., for traffic management
- But are there also unintended consequences?

**Our assessment is that there is a high risk for attacks!**

## Observations

---


BGP updates and table dumps of April 2018 from publicly available BGP Collector Projects: RIPE RIS, Routeviews, Isolario, PCH.

BGP messages	38.98 bn
IPv4 prefixes	967,499
IPv6 prefixes	84,953
Collectors	194
AS peers	2,133
Communities	63,797


**More than 75% of all BGP announcements have at least one BGP community set, 5,659 ASes are using communities.**



## BGP Communities propagation




## BGP Communities propagation




- AS1 announces prefix p

## BGP Communities propagation



- AS1 announces prefix p, AS4 receives announcement

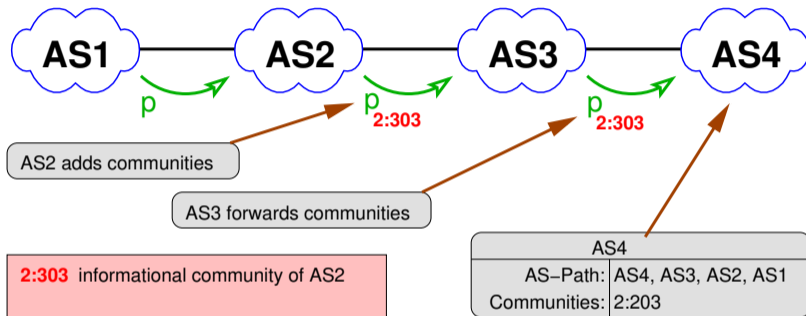
## BGP Communities propagation



- AS1 announces prefix *p*, AS4 receives announcement
- Informational community *2:303* is added by AS2

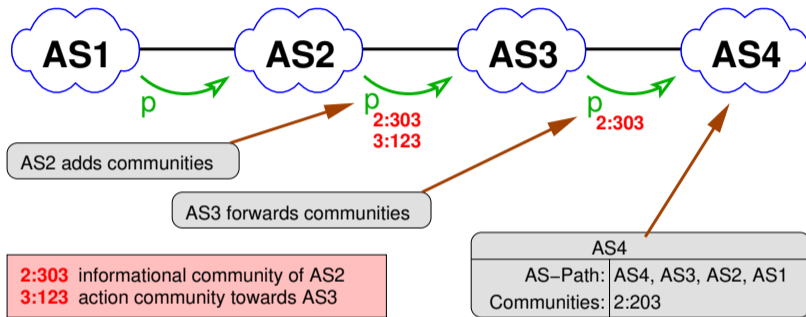


## BGP Communities propagation



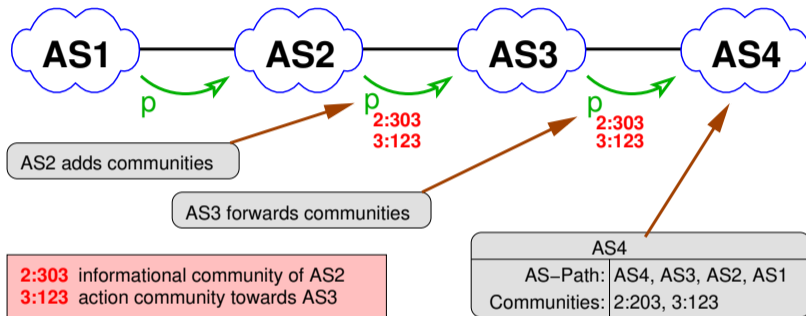
- AS1 announces prefix p, AS4 receives announcement
- Informational community 2:303 is added by AS2

## BGP Communities propagation




- AS1 announces prefix p, AS4 receives announcement
- Informational community 2:303 is added by AS2
- AS2 also adds action community 3:123 for AS3

## BGP Communities propagation




- AS1 announces prefix *p*, AS4 receives announcement
- Informational community *2:303* is added by AS2
- AS2 also adds action community *3:123* for AS3
- Both communities are forwarded by AS3 to AS4

## BGP Communities propagation



AS4	
AS-Path:	AS4, AS3, AS2, AS1
Communities:	2:203, 3:123


## BGP Communities propagation



AS4	
AS-Path:	AS4, AS3, AS2, AS1
Communities:	2:203, 3:123

- We can only infer which AS added a specific community


## BGP Communities propagation



AS4	
AS-Path:	AS4, AS3, AS2, AS1
Communities:	2:203, 3:123

- We can only infer which AS added a specific community
- We assume that a community *n:value* was added by AS *n*

# BGP Communities propagation



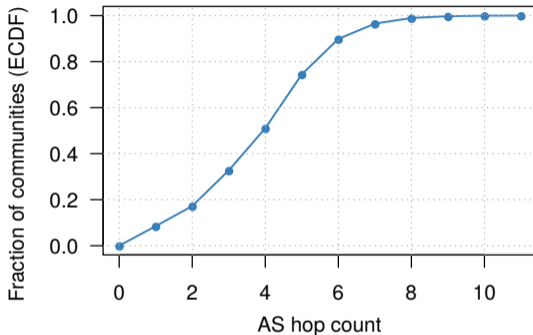
inferred travel-distance is a lower bound!

2:303 traversed at least two AS-links  
3:123 traversed at least one AS-link

AS4	
AS-Path:	AS4, AS3, AS2, AS1
Communities:	2:203, 3:123

- We can only infer which AS added a specific community
- We assume that a community *n:value* was added by AS *n*
- This gives a **lower bound** for the 'travel distance'
- In above example we calculate AS-hop-count 1 for 3:123

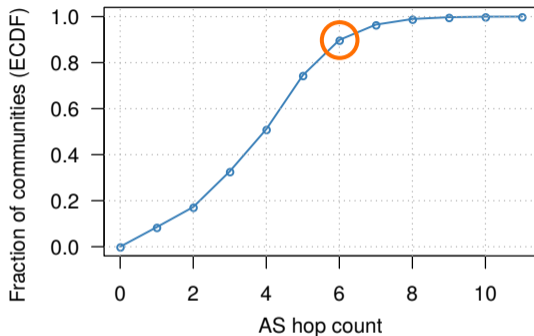
## BGP Community Propagation Observations



- 10% of communities have a AS hop count of more than six
- More than 50% of communities traverse more than four ASes
- Longest community propagation observed: 11 AS hops

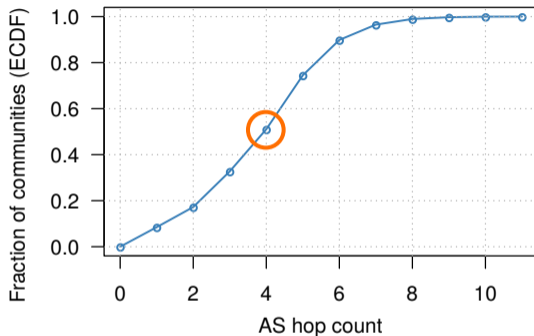


## BGP Community Propagation Observations




- 10% of communities have a AS hop count of more than six
- More than 50% of communities traverse more than four ASes
- Longest community propagation observed: 11 AS hops

## BGP Community Propagation Observations




- 10% of communities have a AS hop count of more than six
- More than 50% of communities traverse more than four ASes
- Longest community propagation observed: 11 AS hops

## BGP Community Propagation Behavior




## BGP Community Propagation Behavior




- AS1 announces prefix p

## BGP Community Propagation Behavior




- AS1 announces prefix  $p$ , tagged with 3:123

## BGP Community Propagation Behavior




- AS1 announces prefix  $p$ , tagged with 3:123
- Community is intended for signaling towards AS3

## BGP Community Propagation Behavior



- AS1 announces prefix  $p$ , tagged with 3:123
- Community is intended for signaling towards AS3
- AS4 also receives this announcement


## BGP Community Propagation Behavior



- AS1 announces prefix  $p$ , tagged with 3:123
- Community is intended for signaling towards AS3
- AS4 also receives this announcement




# BGP Community Propagation Behavior



- AS1 announces prefix p, tagged with 3:123
- Community is intended for signaling towards AS3
- AS4 also receives this announcement

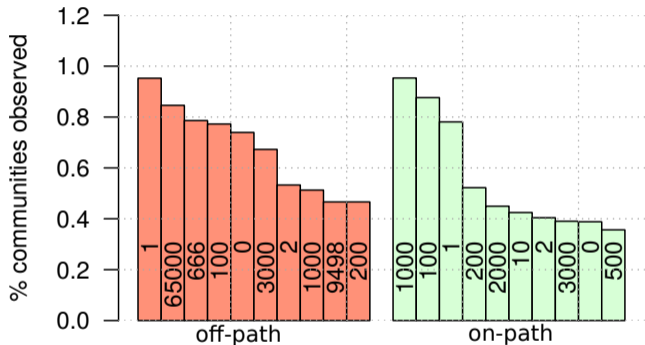
# BGP Community Propagation Behavior



- AS1 announces prefix  $p$ , tagged with  $3:123$
- Community is intended for signaling towards AS3
- AS4 also receives this announcement

**Off-path: ASN from community is not on the observed AS-path at AS4.**

## On-path versus off-path



- Blackholing communities (e.g., :666) 'leaking' off path
- But AS implementing RTBH SHOULD add NO\_ADVERTISE or NO\_EXPORT (RFC7999)

**Suggests ASes not implementing RTBH do not filter.**

# Experiments

---

# Experimental setup

- Experiments conducted in a lab environment<sup>1</sup>
- Validated on the Internet


## Scenarios

- Remote Triggered Blackholing (RTBH)
- Traffic redirection attack

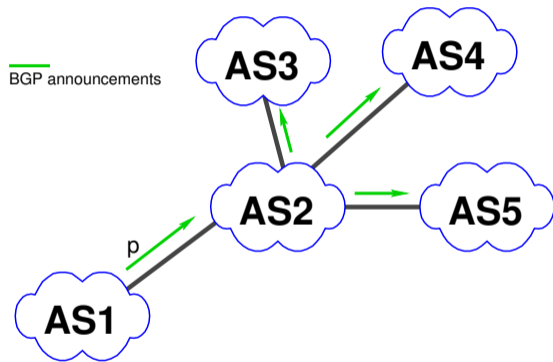
...more in the paper.

---


<sup>1</sup>Configurations available at: <https://www.cmand.org/caas/>



## RTBH: How It Works




# RTBH: How It Works






# RTBH: How It Works




- AS announces BH-prefix to upstream

# RTBH: How It Works




- AS announces BH-prefix to upstream

# RTBH: How It Works



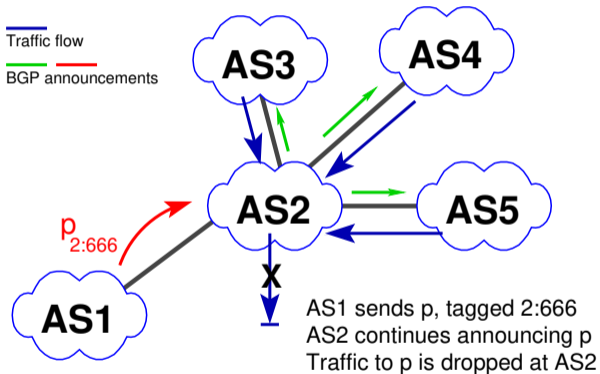
- AS announces BH-prefix to upstream

# RTBH: How It Works




- AS announces BH-prefix to upstream
- Provider blackholes prefix

# RTBH: How It Works



- AS announces BH-prefix to upstream
- Provider blackholes prefix

# RTBH: How It Works




- AS announces BH-prefix to upstream
- Provider blackholes prefix


## Safeguards

- Provider should check customer prefix before accepting RTBH
- Customer may only blackhole own prefixes
- Different policies for Customers/Peers
- On receiving RTBH, add `NO_ADVERTISE` or `NO_EXPORT` (RFC7999)

## RTBH: How It Should Not Work

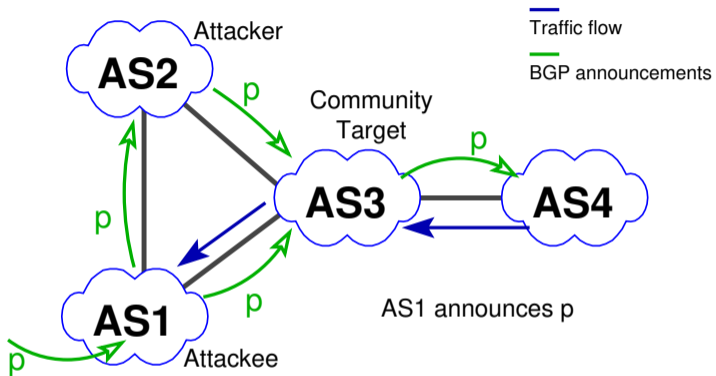


## RTBH: How It Should Not Work

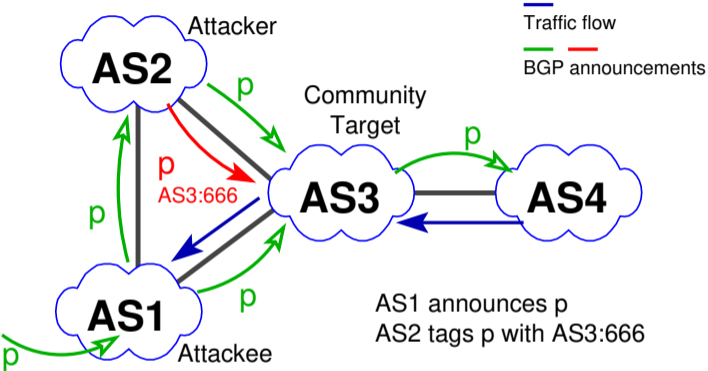




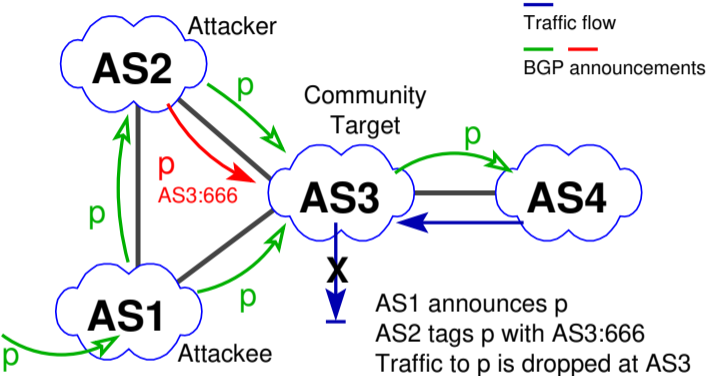
# RTBH: How It Should Not Work




# RTBH: How It Should Not Work



# RTBH: How It Should Not Work

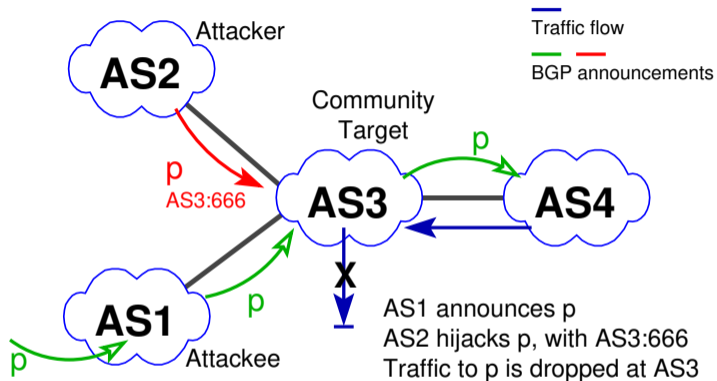


## RTBH: How It Should Not Work



- AS on 'backup' path adds RTBH-community
- Provider blackholes prefix
- Not only traffic traversing AS2 is dropped

## RTBH: How It Should Not Work (with hijack)



- Hijacker announces RTBH
- Prefix filters circumvented due to misconfiguration
- Provider blackholes prefix

**Attack confirmed to work on the Internet, works multi hop and is hard to spot**


Triggering RTBH is possible for attackers because, e.g.,:

- BH prefix is more specific, accepted via exception
- Providers check BH community before prefix filters<sup>2</sup>
- NO\_ADVERTISE or NO\_EXPORT often is ignored / not set
- Problem: No validation for origin of community


---

<sup>2</sup>we found configuration guides with that bug

# Traffic Redirection Attack




# Traffic Redirection Attack



— BGP-Announcements




# Traffic Redirection Attack



# Traffic Redirection Attack




# Traffic Redirection Attack




- Attacker AS2 uses community to add path-prepending in AS3

# Traffic Redirection Attack




- Attacker AS2 uses community to add path-prependings in AS3
- AS6 routes traffic towards prefix p via AS5, AS4

# Traffic Redirection Attack




- Attacker AS2 uses community to add path-prepending in AS3
- AS6 routes traffic towards prefix p via AS5, AS4

# Traffic Redirection Attack



- Attacker AS2 uses community to add path-prependings in AS3
- AS6 routes traffic towards prefix p via AS5, AS4
  - Network tap?

# Traffic Redirection Attack



- Attacker AS2 uses community to add path-prepending in AS3
- AS6 routes traffic towards prefix p via AS5, AS4
  - Network tap?
  - Slow/Congested link?
  - ...

### Attack on 10 July 2018

"For about 30 minutes, these hijack prefixes weren't propagated very far. Then they were announced again at 23:37:47 UTC for about 15 minutes but to a larger set of peers — 48 peers instead of 3 peers in the previous hour.


**It appears a change of BGP communities from 24218:1120 to 24218:1 increased the route propagation."**

Source: <https://dyn.com/blog/bgp-dns-hijacks-target-payment-systems/>



## Discussion

---



## Discussion: Authenticity


- Communities can be modified, added, removed by every AS
- No attribution is possible
- No cryptographic protection (RPKI does not help)
- Still operators rely on their 'correctness'
- Large communities partially improve the situation



**How can we achieve authenticity, or at least attribution?**

## Discussion: Transitivity

- Communities can help in debugging
- Easy, low overhead communication channel
- Widely in use, but often only 1-2 hops
- But: High risk of being abused!



**Are fully transitive communities still worth the clear risk?**

## Discussion: Monitoring

- There is no global state in BGP
- Route collectors only see the 'end-result'
- Inferring modifications between origin-AS and collector: almost impossible
- The meaning of a particular community can not be known
- No universal way for attribution of changes



**Monitoring communities to detect abuse is extremely difficult.**

## Discussion: Standards

- There are limited standardized communities
- Many AS do not implement these
- Is the lack of standardized communities a problem?
- Are standards doing harm, by helping attackers?
- Security by obscurity never works



**Standardization is necessary.**

## Discussion: Documentation

- Communities are individually defined by the ASes
- Documentation, if available, is scattered over whois, websites, customer-portals, ...
- Not in machine-readable format, often natural language
- Automated parsing can work for limited scope/fixed applications
- Parsing for general purpose applications is not feasible



**Documentation is limited and fragmented.**

## Communities Shortcomings

- Semantics loosely defined, no authenticity
- Secure usage requires good **operational knowledge** and **diligence**
- Attacks are possible and indeed already happening



## Communities Shortcomings

- Semantics loosely defined, no authenticity
- Secure usage requires good **operational knowledge** and **diligence**
- Attacks are possible and indeed already happening

## Future Work

- Attack detection
- Attribution
- Distributed realtime monitoring?
- Protocol improvements for BGP?

# Appendix

---

## Recommendations for Operators

- AS should filter incoming Informational Communities carrying their ASN
- Agreements with Downstreams might be needed, e.g., to filter Action Communities
- Publicly documenting Communities used is key to enable other AS to filter
- Monitoring/Logging received communities for tracking abuse
- Providing public looking glasses, showing communities, helps debugging