

# BGP Communities:

## A measurement study

@RIPE77, Amsterdam

---

Florian Streibelt<sup>1</sup>, 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>

Oct. 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)

# BGB-Communities: A weapon for the Internet!


RIPE 77 / Amsterdam

2018.10.16

# Introduction


---

# BGP Community usage is increasing



Increasing usage warrants a closer look.

# BGP Community usage is increasing



Increasing usage warrants a closer look.

- Optional Attribute in BGP message (32 bit field)
- Defined in RFC 1997
- By convention written *ASN:VALUE*
- ASN can be both sender or intended 'recipient'
- It's up to the peers to agree upon 'values' used

# BGP Large Communities

- Defined by RFC 8092 (usage recommendations ins RFC 8195)
- 12 byte attribute
- Enable networks with 4-byte ASNs to use communities
- The first 4 byte contain the ASN of the "global administrator"



**Sorry... as we only found a very small number of occurrences<sup>1</sup> we could not conduct any meaningful measurements, yet.**

---

<sup>1</sup>283 individual large communities by 51 global administrators over the whole month of April 2018 at all available route collectors at RIPE/RIS, Routeviews, Isolario and PCH



# 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

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

# What This Talk Is About

Given the **increasing popularity** of BGP communities and the ability to **trigger actions** as well as **relay information**, the first question that comes to the mind of an Internet measurement researcher is. . .

# What This Talk Is About



What could possibly go wrong?

# Propagation behavior



# Propagation behavior

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

**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

---

## Dataset


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 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
- Validated on the Internet

## Scenarios


- Remote Triggered Blackholing (RTBH)
- Traffic redirection attack

...for others see our paper.

## RTBH: how it works




## RTBH: how it works






## RTBH: how it works




# RTBH: how it works

- AS announces BH-prefix to upstream



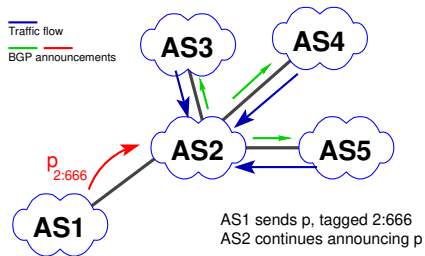
# RTBH: how it works

- 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

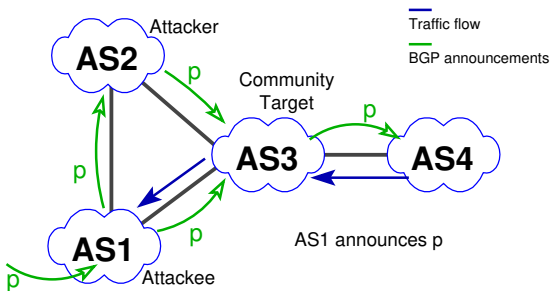




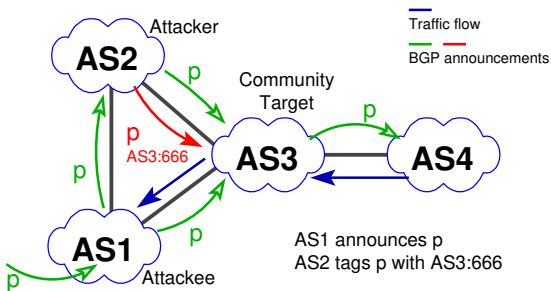
## RTBH: how it should not work



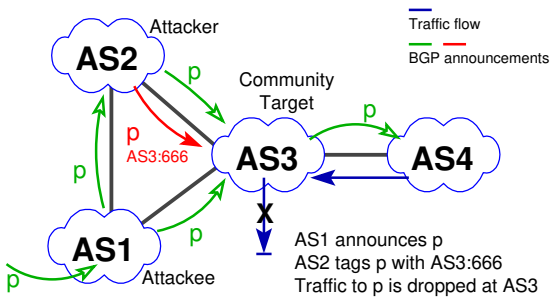
## RTBH: how it should not work



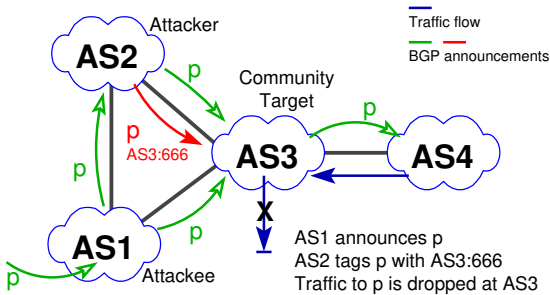
## RTBH: how it should not work



## RTBH: how it should not work

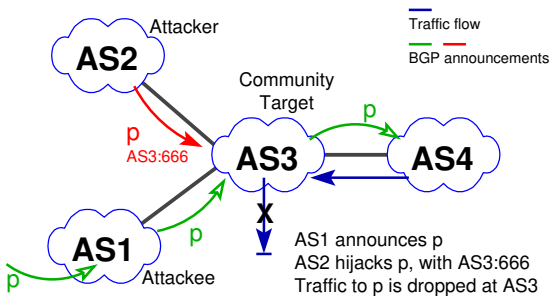


# RTBH: how it should not work



- 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-prepend in AS3
- AS6 routes traffic towards prefix p via AS5, AS4

# Traffic redirection attack




- Attacker AS2 uses community to add path-prepend 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
  - 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?
  - ...



## Discussion: What now?

---



## BGP Communities Shortcomings Summarized

- Notation of "*ASN:value*" is just convention
- No defined semantics: values can mean anything
- Used both for signaling and triggering of actions
- No cryptographic protection
- Attribution is impossible
- Large Communities have, in principle, similar limitations



# BGP Communities: The Problem

- BGP Communities as they are used are not necessarily broken
- Secure usage requires good **operational knowledge** and **diligence**

# BGP Communities: The Problem

- BGP Communities as they are used are not necessarily broken
- Secure usage requires good **operational knowledge** and **diligence**
- While people in this room probably know what they are doing:  
Based on experience we do not rely on that globally. . .

**Do we need less fragile protocols and mechanisms?**

# Recommendations

- Filter incoming Informational Communities for your ASN
- Publish community documentation, to enable others to filter
- Monitor and log received communities to track abuse
- Talk to your Downstreams, so they filter  
Action Communities for your ASN on ingress if necessary
- Provide a looking glass (that shows communities!)

## Discussion: Authenticity

- Communities can be modified, added, removed by every AS
- No attribution is possible
- No cryptographic protection
- 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.**

There is no easy way to find meaning of a community:

- Some ASes document in the whois
- Some ASes document on their website
- Some ASes provide documentation only to customers
- Some ASes do not provide any documentation

**Documentation is limited and fragmented.**

# Summary

- Communities are widely in use
- Foundation of many policies

But:

- Relies heavily on mutual trust in capabilities
- No authenticity/security in place
- Attribution is impossible
- Hard to detect attacks
- While our prefix hijacks were reported,  
no one reported our community attacks

**It's unknown if there are other unnoticed attacks.**



Get the preprint at:

[https://people.mpi-inf.mpg.de/~fstreibelt/preprint/  
communities-imc2018.pdf](https://people.mpi-inf.mpg.de/~fstreibelt/preprint/communities-imc2018.pdf)

Published at ACM IMC 2018

<https://conferences.sigcomm.org/imc/2018/>



Contact:

Florian Streibelt <fstreibelt@mpi-inf.mpg.de>

Images:

Unicorn illustrations: Telegram stickers by Darya Ogneva:

<https://tlgrm.eu/stickers/BornToBeAUnicorn>

The Spanish Inquisition: by Miki Montllo

[http://miquelmontllo.blogspot.com/2013/10/  
the-spanish-inquisition-wallpaper.html](http://miquelmontllo.blogspot.com/2013/10/the-spanish-inquisition-wallpaper.html)