Weaponizing BGP Communities: Yet another attack on routing?

”BGP Communities: Even more Worms in the Routing Can”, ACM IMC 2018

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

IETF104, Prague, March 2019

[1] Max Planck Institute for Informatics (MPII), [2] Naval Postgraduate School (NPS), [3] University of Strasbourg, [4] TU Berlin (TUB), [5] Internet Initiative Japan (IIJ)

FD: Randy grow@IETF103

Weaponizing BGP Using Communities

Florian Streibelt, Franziska Lichtblau, Robert Beverly, Cristel Pelsser, Georgios Smaragdakis, Randy Bush, Anja Feldmann

Introduction

BGP Community usage is increasing

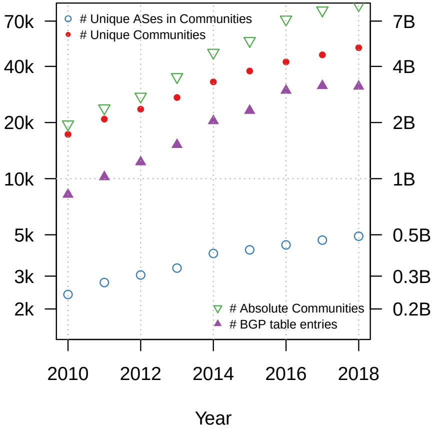

Increasing usage warrants a closer look.

BGP Community usage is increasing

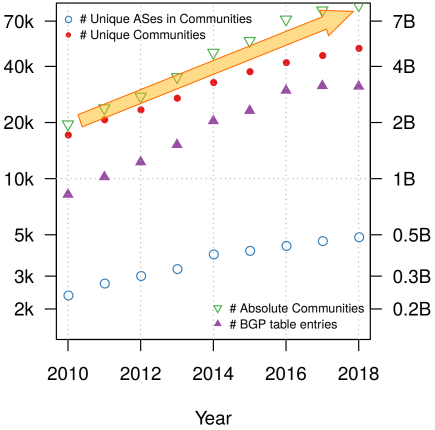

Increasing usage warrants a closer look.

BGP Communities

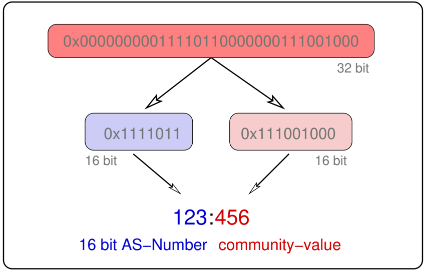

- RFC 1997: Optional Attribute in BGP message (32 bit)
- By convention written *ASN:VALUE*
- ASN can be both sender or intended 'recipient'
- It's up to the peers to agree upon 'values' used
- Every network decides on the semantics of values

BGP Large Communities

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

BGP Large Communities

BGP Large Communities

Sorry. . . as we only found a very small number of occurrences[1] we could not conduct any meaningful measurements, yet.

[1] 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 Large Communities

Update: The number of global administrators is increasing[1]

In Feb./March 2019 we see more than 120 global administrators...

[1] see `https://labs.ripe.net/Members/emileaben/bgp-large-communities-uptake-an-update`

BGP Communities: Usage (examples)

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

- 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

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 BGP announcements have at least one BGP community set, 5,659 ASes are using communities.

BGP Communities propagation

BGP Communities propagation

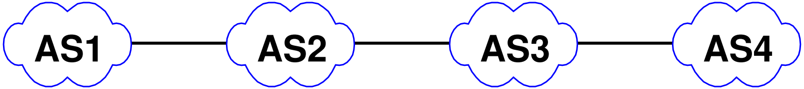

BGP Communities propagation

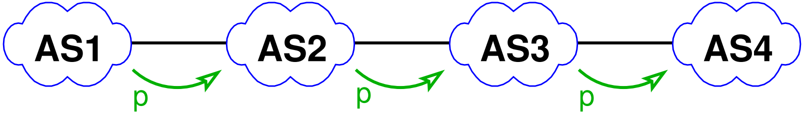

- AS1 announces prefix p

BGP Communities propagation

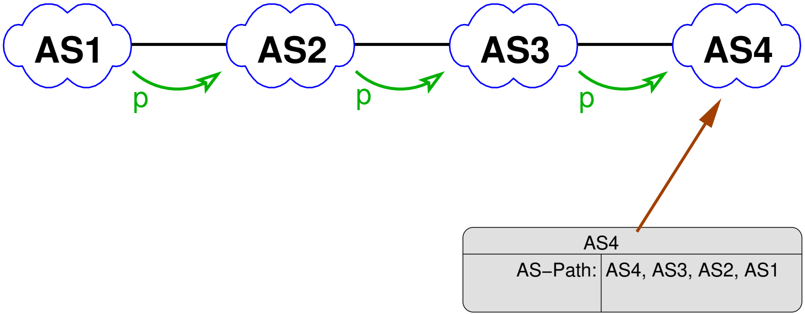

AS4	
AS–Path:	AS4, AS3, AS2, AS1

- AS1 announces prefix p, AS4 receives announcement

BGP Communities propagation

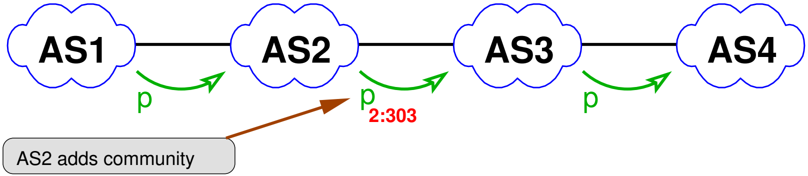

AS2 adds community

2:303 informational community of AS2

AS4	
AS–Path:	AS4, AS3, AS2, AS1

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

BGP Communities propagation

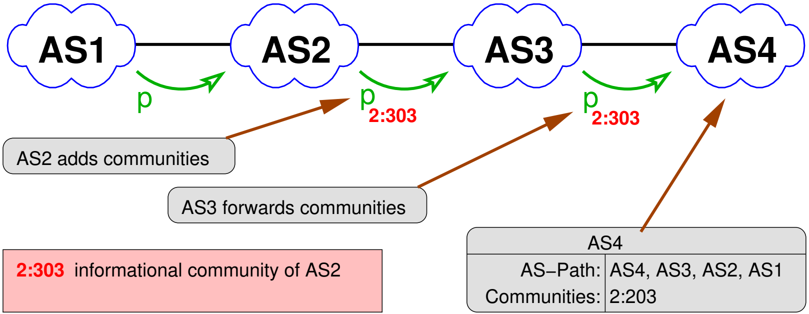

AS2 adds communities

AS3 forwards communities

2:303 informational community of AS2

AS4	
AS–Path:	AS4, AS3, AS2, AS1
Communities:	2:203

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

BGP Communities propagation

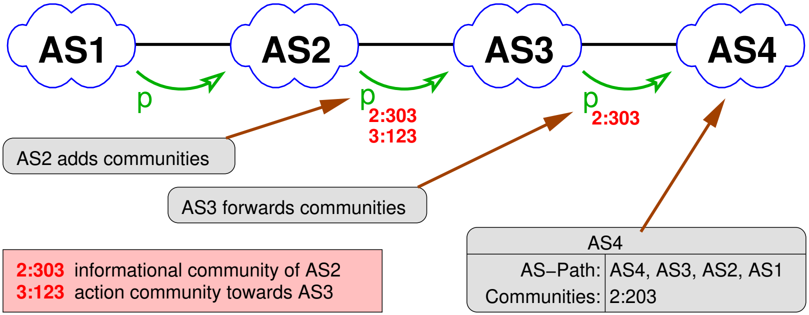

AS2 adds communities

AS3 forwards communities

AS4	
AS–Path:	AS4, AS3, AS2, AS1
Communities:	2:203

2:303 informational community of AS2
3:123 action community towards AS3

- 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

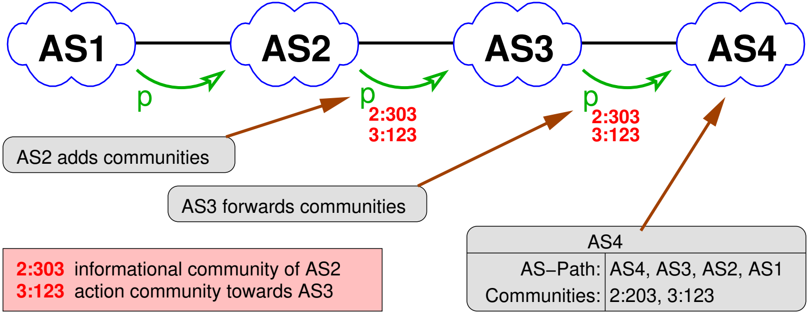

AS2 adds communities

AS3 forwards communities

2:303 informational community of AS2
3:123 action community towards AS3

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

- 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

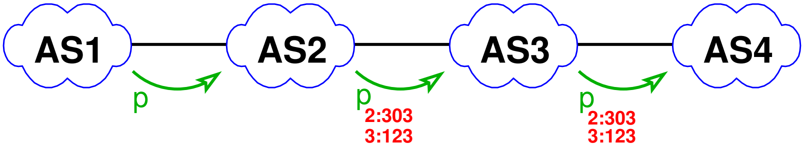

AS4	
AS–Path:	AS4, AS3, AS2, AS1
Communities:	2:203, 3:123

BGP Communities propagation

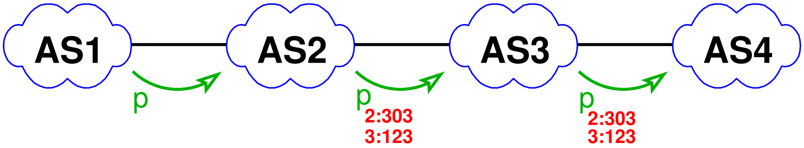

AS4	
AS–Path:	AS4, AS3, AS2, AS1
Communities:	2:203, 3:123

- We can only infer which AS added a specific community

BGP Communities propagation

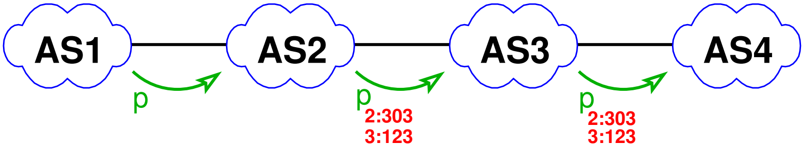

AS4	
AS–Path:	AS4, AS3, AS2, AS1
Communities:	2:203, 3:123

- 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

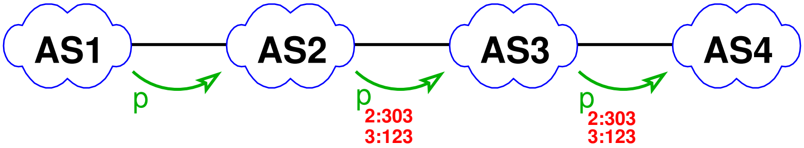

inferred travel–distance is a lower bound!

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

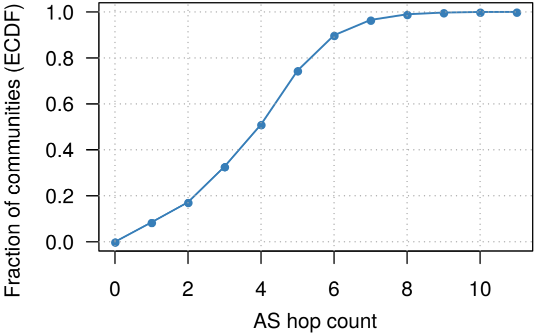

BGP Community Propagation Observations

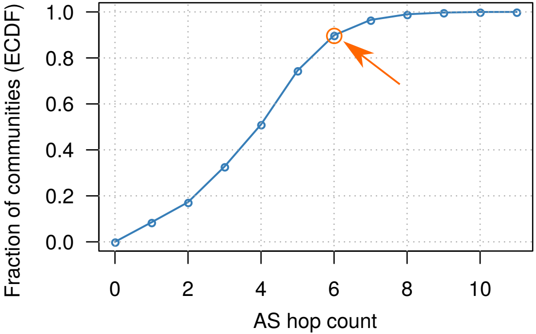

- 10% of communities have an AS hop count of more than six

BGP Community Propagation Observations

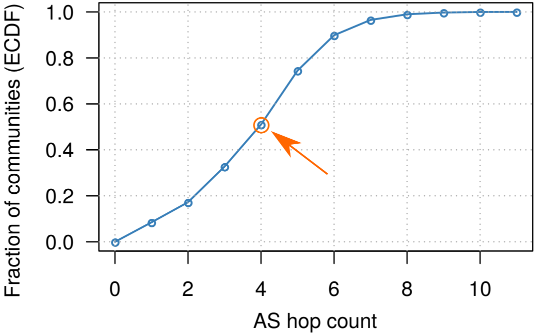

- 10% of communities have an AS hop count of more than six
- More than 50% of communities traverse more than four ASes

BGP Community Propagation Observations

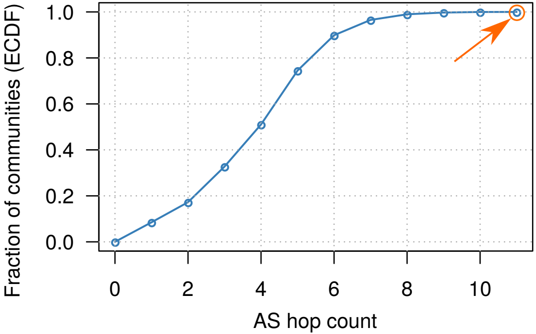

- 10% of communities have an AS hop count of more than six
- More than 50% of communities traverse more than four ASes
- Longest community propagation observed: 11 AS hops

BGP Community Propagation Behavior

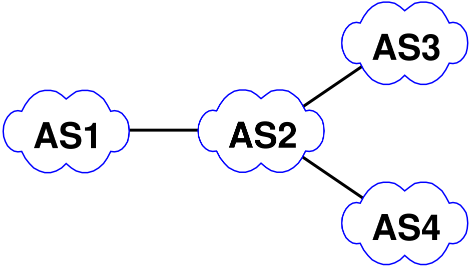

BGP Community Propagation Behavior

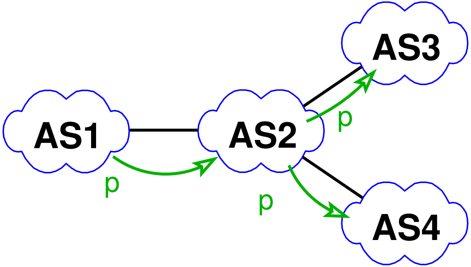

- AS1 announces prefix p

BGP Community Propagation Behavior

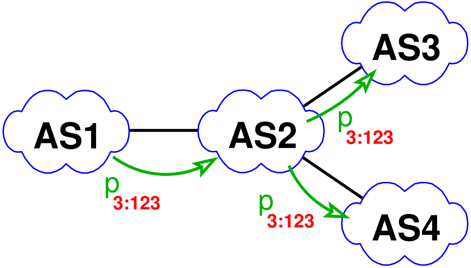

- AS1 announces prefix p, tagged with 3:123

BGP Community Propagation Behavior

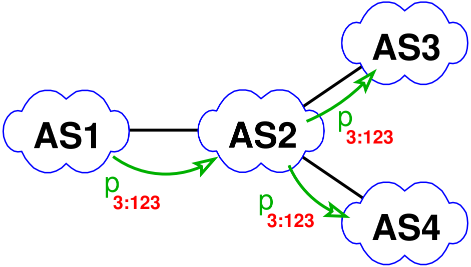

- AS1 announces prefix p, tagged with 3:123
- Community is intended for signaling towards AS3

BGP Community Propagation Behavior

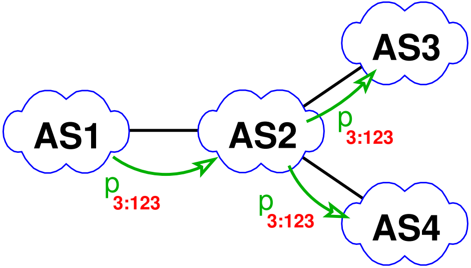

- AS1 announces prefix p, tagged with 3:123
- Community is intended for signaling towards AS3
- AS4 also receives this announcement

BGP Community Propagation Behavior

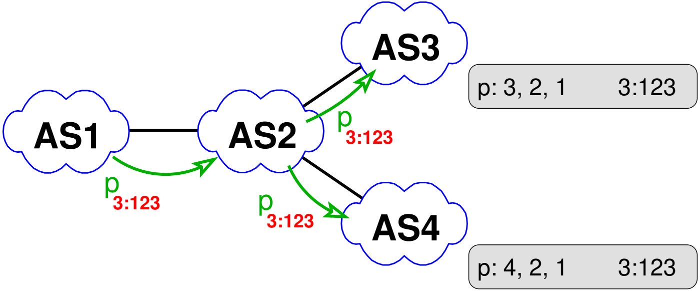

- AS1 announces prefix p, tagged with 3:123
- Community is intended for signaling towards AS3
- AS4 also receives this announcement

BGP Community Propagation Behavior

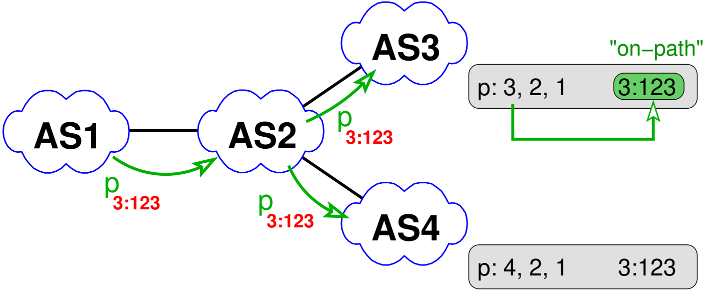

- AS1 announces prefix p, tagged with 3:123
- Community is intended for signaling towards AS3
- AS4 also receives this announcement

BGP Community Propagation Behavior

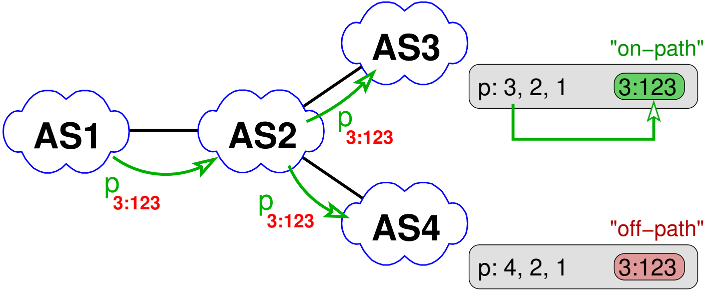

- AS1 announces prefix p, tagged with 3:123
- Community is intended for signaling towards AS3
- AS4 also receives this announcement

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

On-path versus off-path

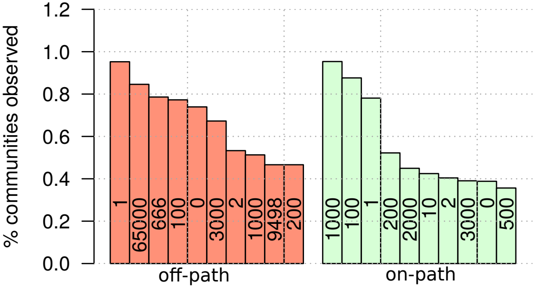

- Blackholing communities (e.g., :666) 'leaking' off path
- But AS implementing RTBH `SHOULD` add `NO_ADVERTISE` or `NO_EXPORT` (RFC7999)

Suggests ASes not implementing RTBH do not filter.

Experiments

Experimental setup

- Experiments conducted in a lab environment[2]
- Validated on the Internet

Scenarios

- Remote Triggered Blackholing (RTBH)
- Traffic redirection attack

...more in the paper.

[2] Configurations available at: `https://www.cmand.org/caas/`

RTBH: how it works

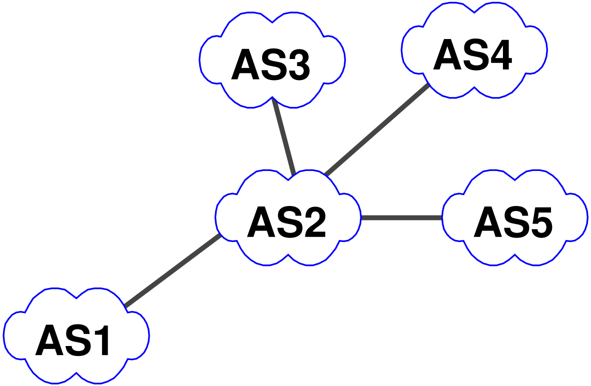

RTBH: how it works

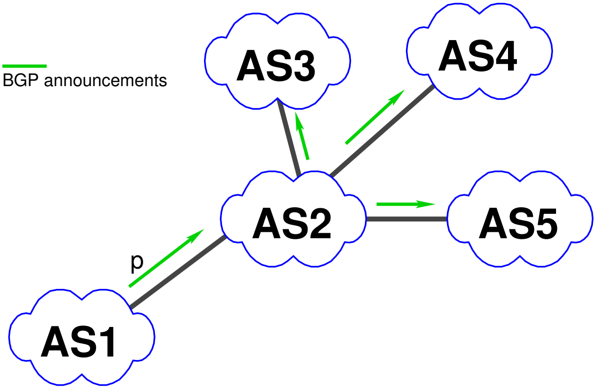

RTBH: how it works

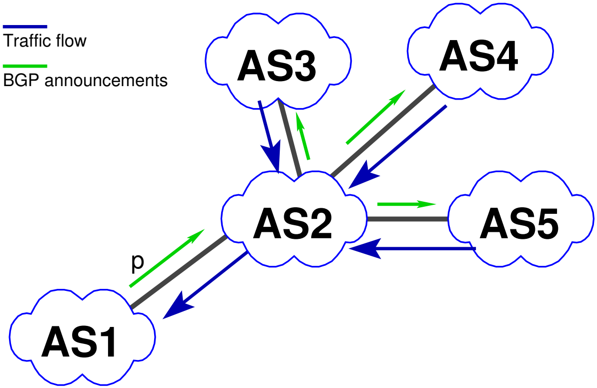

RTBH: how it works

- AS announces BH-prefix to upstream

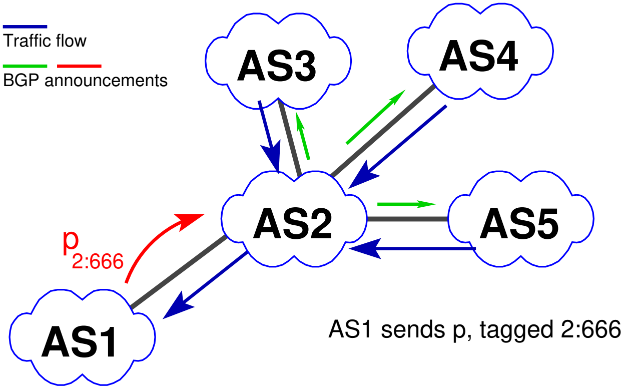

AS1 sends p, tagged 2:666

RTBH: how it works

- AS announces BH-prefix to upstream

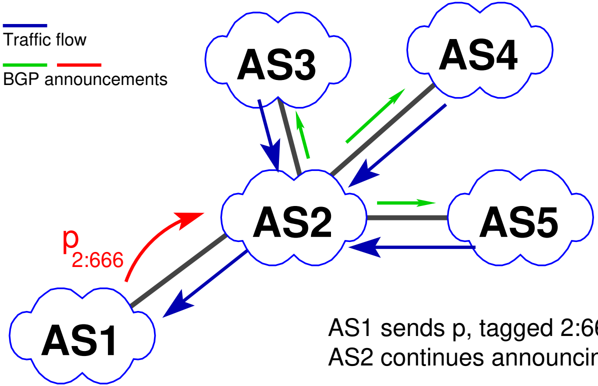

AS1 sends p, tagged 2:666
AS2 continues announcing p

RTBH: how it works

- AS announces BH-prefix to upstream

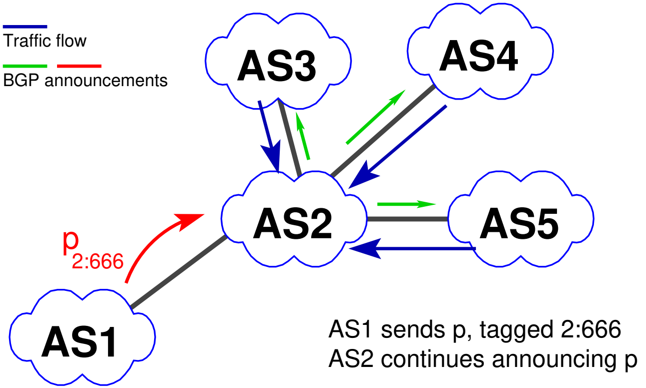

AS1 sends p, tagged 2:666
AS2 continues announcing p

RTBH: how it works

- AS announces BH-prefix to upstream

→ Provider blackholes prefix

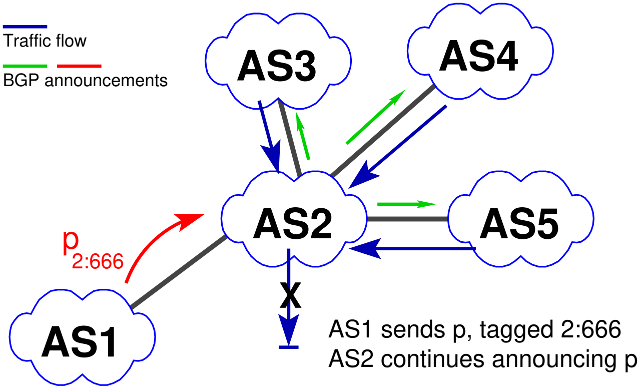

AS1 sends p, tagged 2:666
AS2 continues announcing p

RTBH: how it works

- AS announces BH-prefix to upstream

→ Provider blackholes prefix

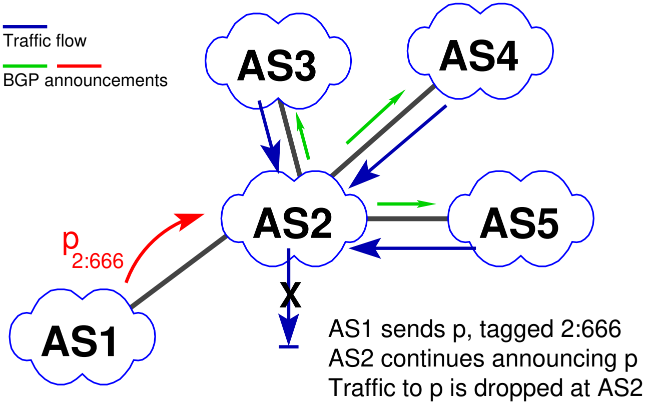

AS1 sends p, tagged 2:666
AS2 continues announcing p
Traffic to p is dropped at AS2

RTBH: how it works

- AS announces BH-prefix to upstream

→ Provider blackholes prefix

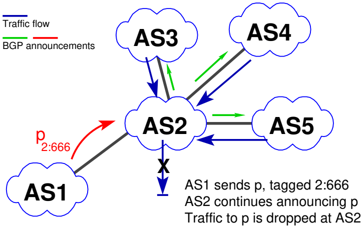
Traffic flow
BGP announcements
AS3
AS4
AS2
AS5
AS1
p
2:666
X
AS1 sends p, tagged 2:666
AS2 continues announcing p
Traffic to p is dropped at AS2

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

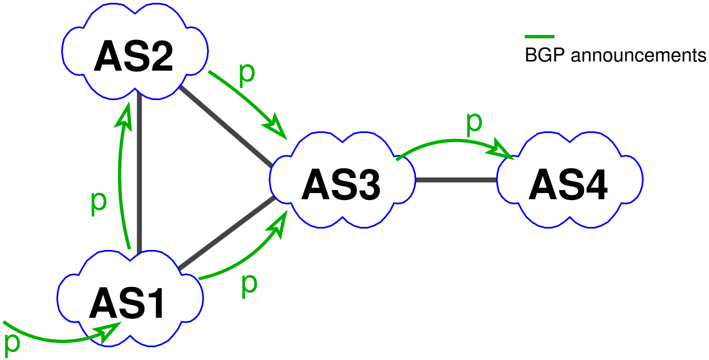

RTBH: how it should not work

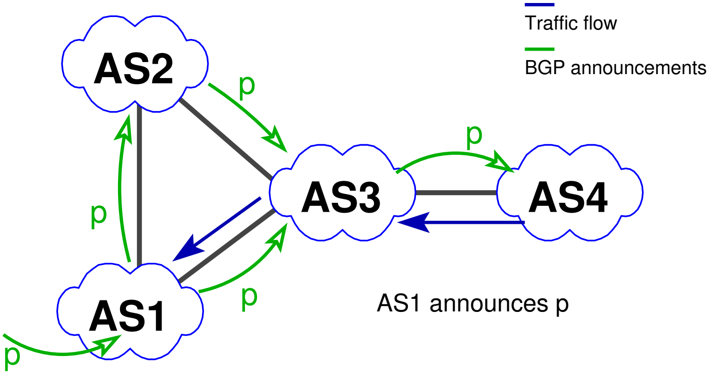

AS1 announces p

RTBH: how it should not work

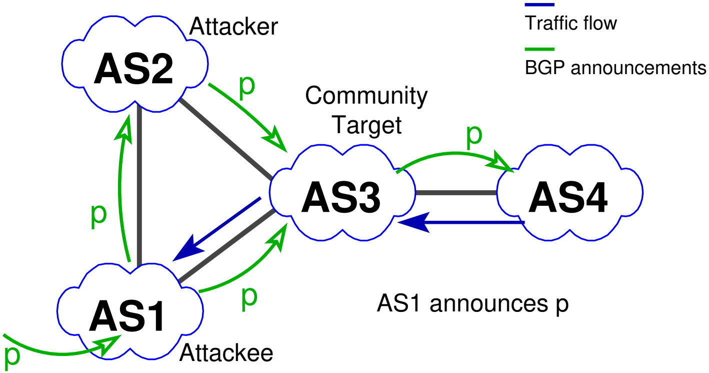

AS1 announces p

RTBH: how it should not work

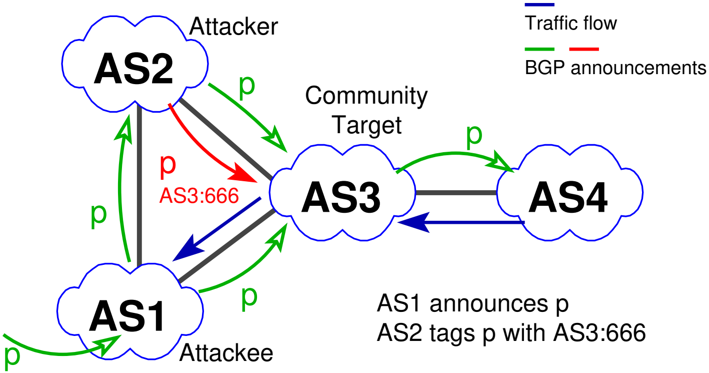

AS1 announces p
AS2 tags p with AS3:666

RTBH: how it should not work

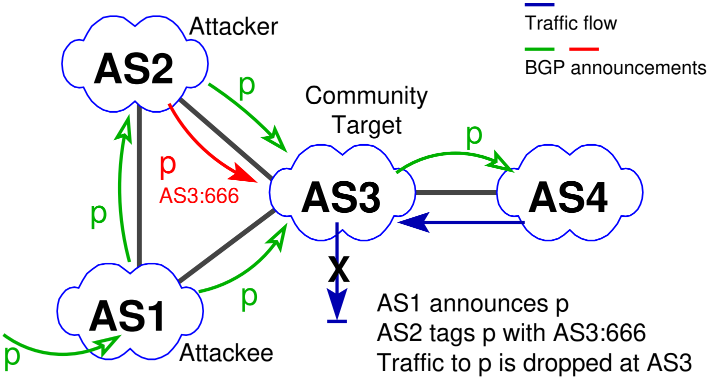

AS1 announces p
AS2 tags p with AS3:666
Traffic to p is dropped at AS3

RTBH: how it should not work

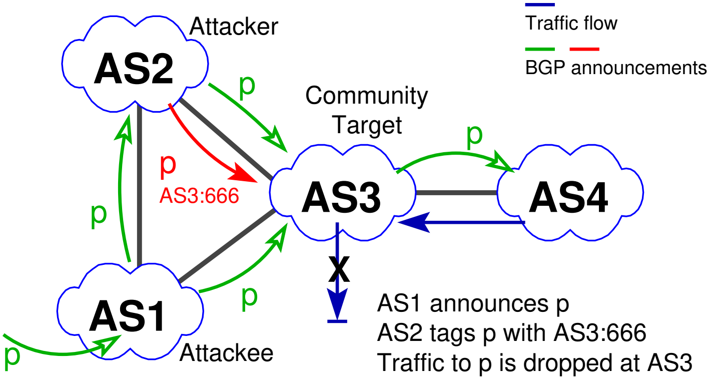

AS1 announces p
AS2 tags p with AS3:666
Traffic to p is dropped at AS3

- 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)

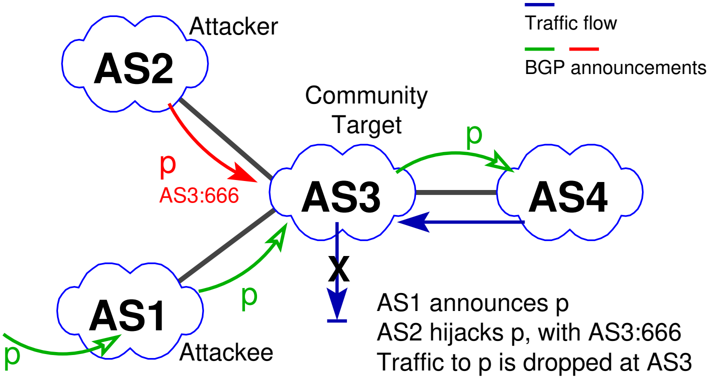

AS1 announces p
AS2 hijacks p, with AS3:666
Traffic to p is dropped at AS3

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

RTBH: Attack confirmed

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[3]
- `NO_ADVERTISE` or `NO_EXPORT` often is ignored / not set
- Problem: No validation for origin of community

[3] we found configuration guides with that bug

RTBH: Attack Mitigation

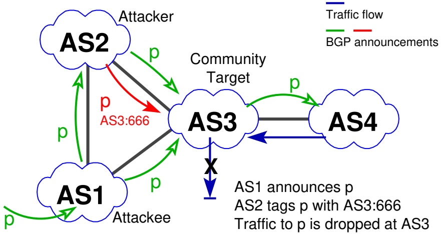

AS1 announces p
AS2 tags p with AS3:666
Traffic to p is dropped at AS3

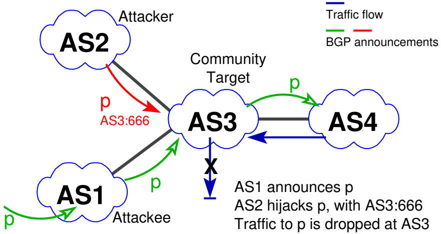

AS1 announces p
AS2 hijacks p, with AS3:666
Traffic to p is dropped at AS3

Mitigation

- RTBH Provider should check for best path
- Accept Blackholing announcement only if that peer is currently on the best path

Checkout talk at IEPG by Job Snijders yesterday!

Traffic redirection attack

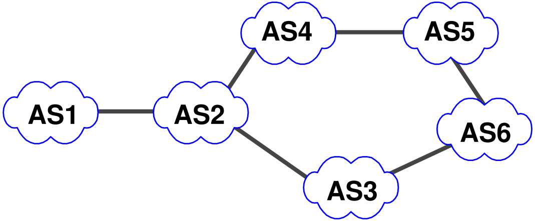

Traffic redirection attack

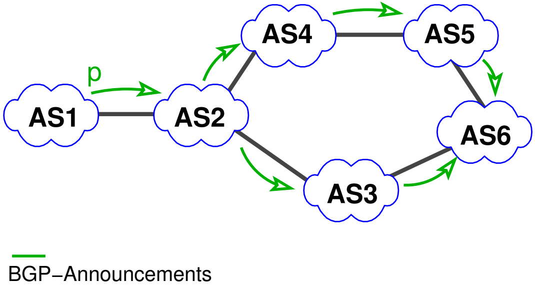

Traffic redirection attack

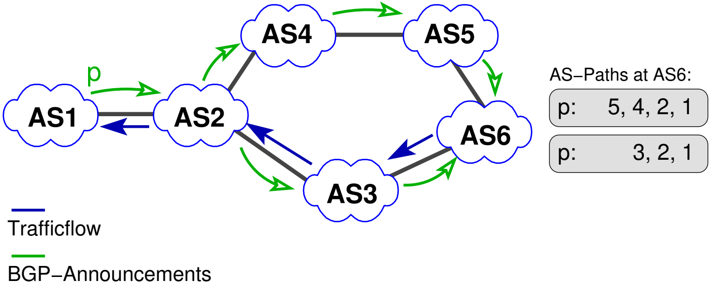

AS−Paths at AS6:

p:	5, 4, 2, 1
p:	3, 2, 1

Trafficflow

BGP−Announcements

Traffic redirection attack

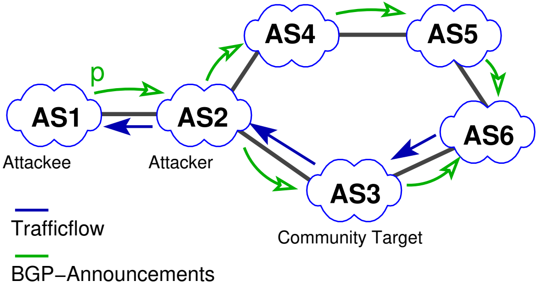

AS−Paths at AS6:

p: 5, 4, 2, 1

p: 3, 2, 1

Traffic redirection attack

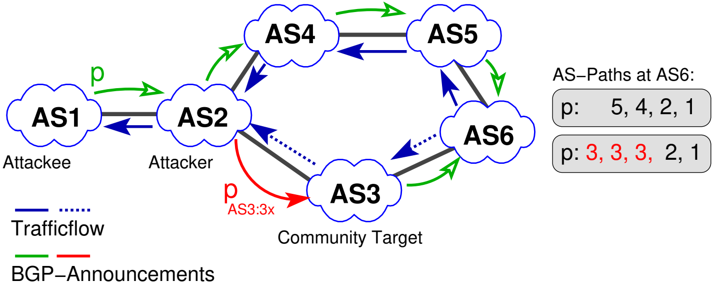

p

AS1

AS2

AS4

AS5

AS6

AS3

Attackee

Attacker

p

AS3:3x

Community Target

Trafficflow

BGP–Announcements

AS–Paths at AS6:

p: 5, 4, 2, 1

p: 3, 3, 3, 2, 1

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

Traffic redirection attack

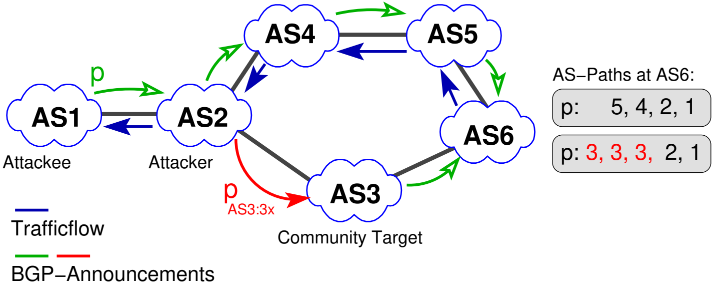

AS–Paths at AS6:

p: 5, 4, 2, 1

p: 3, 3, 3, 2, 1

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

Traffic redirection attack

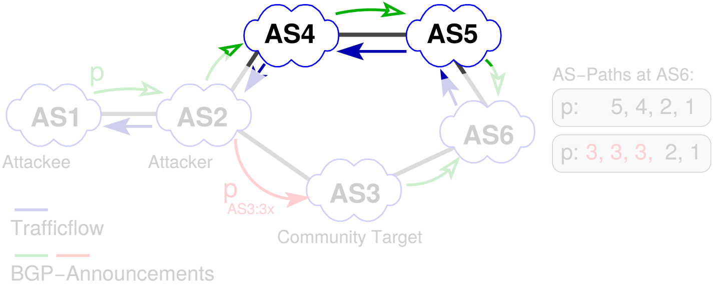

- Attacker AS2 uses community to add path-prepending in AS3
- AS6 routes traffic towards prefix p via AS5, AS4

Traffic redirection attack

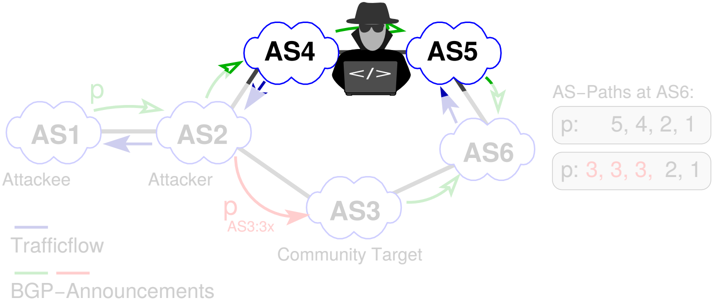

- Attacker AS2 uses community to add path-prepending in AS3
- AS6 routes traffic towards prefix p via AS5, AS4
    - Network tap?

Traffic redirection attack

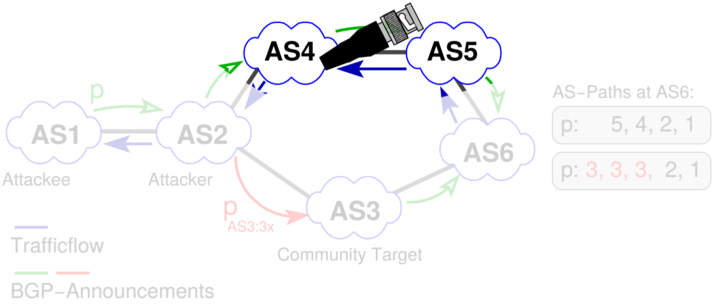

- Attacker AS2 uses community to add path-prepending in AS3
- AS6 routes traffic towards prefix p via AS5, AS4
    - Network tap?
    - Slow/Congested link?
    - ...

Communities Confirmed In Attacks

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

What now?

Discussion

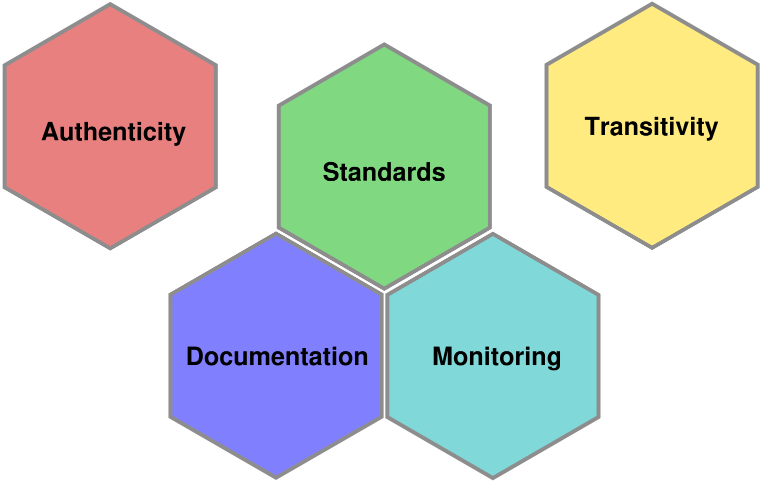
Authenticity
Standards
Transitivity
Documentation
Monitoring

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

Authenticity

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!

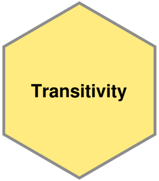

Transitivity

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

- Notation of ”*ASN:value*” is just convention
- No defined semantics: values can mean anything
- Used both for signaling and triggering of actions
- There are limited standardized communities
- Many ASes do not implement these

Standards

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

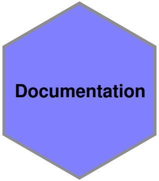
Documentation

Documentation is limited and fragmented.

Discussion: Standards in Documentation

- DTAG internally developed a system for ”community structuring”
- Translates string representation to communities (short + large)
- Example: `tag.origin.country.DE`
- Allows definition of parameters to communities
- Documents communities and parameters
- Working code, used in production
- System is documented in an Internet-Draft style document

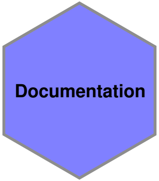
Documentation

Is this a way for standardizing documentation?

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 communties, helps debugging

BGP Communities: The Problem

- BGP communities are the only feasible way to realize signaling between ASes
- Secure usage requires good **operational knowledge** and **diligence**
- Overcomplex security mechanisms around their short comings are not the solution

BGP Communities: The Problem

- BGP communities are the only feasible way to realize signaling between ASes
- Secure usage requires good **operational knowledge** and **diligence**
- Overcomplex security mechanisms around their short comings are not the solution
- While people in this room probably know what they are doing:
  Based on experience we should not rely on that globally...

Do we need less fragile protocols and mechanisms?

Summary

- Communities are widely in use
- Used to realize policies

But:

- Heavily relies on mutual trust between peers:
- 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.

BGP Communities: Even more Worms in the Routing Can

BGP Communities: Even more Worms in the Routing Can

Florian Streibelt, Max Planck Institute for Informatics
Franziska Lichtblau, Max Planck Institute for Informatics
Robert Beverly, Naval Postgraduate School
Anja Feldmann, Max Planck Institute for Informatics
Cristel Pelsser, University of Strasbourg
Georgios Smaragdakis, TU Berlin
Randy Bush, Internet Initiative Japan

ABSTRACT

1 INTRODUCTION

Get the preprint version at:

`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