

LACNIC 31

8 de Maio de 2019

Automatização de Listas de Prefixos em Peering BGP
Como fazer e sua importância

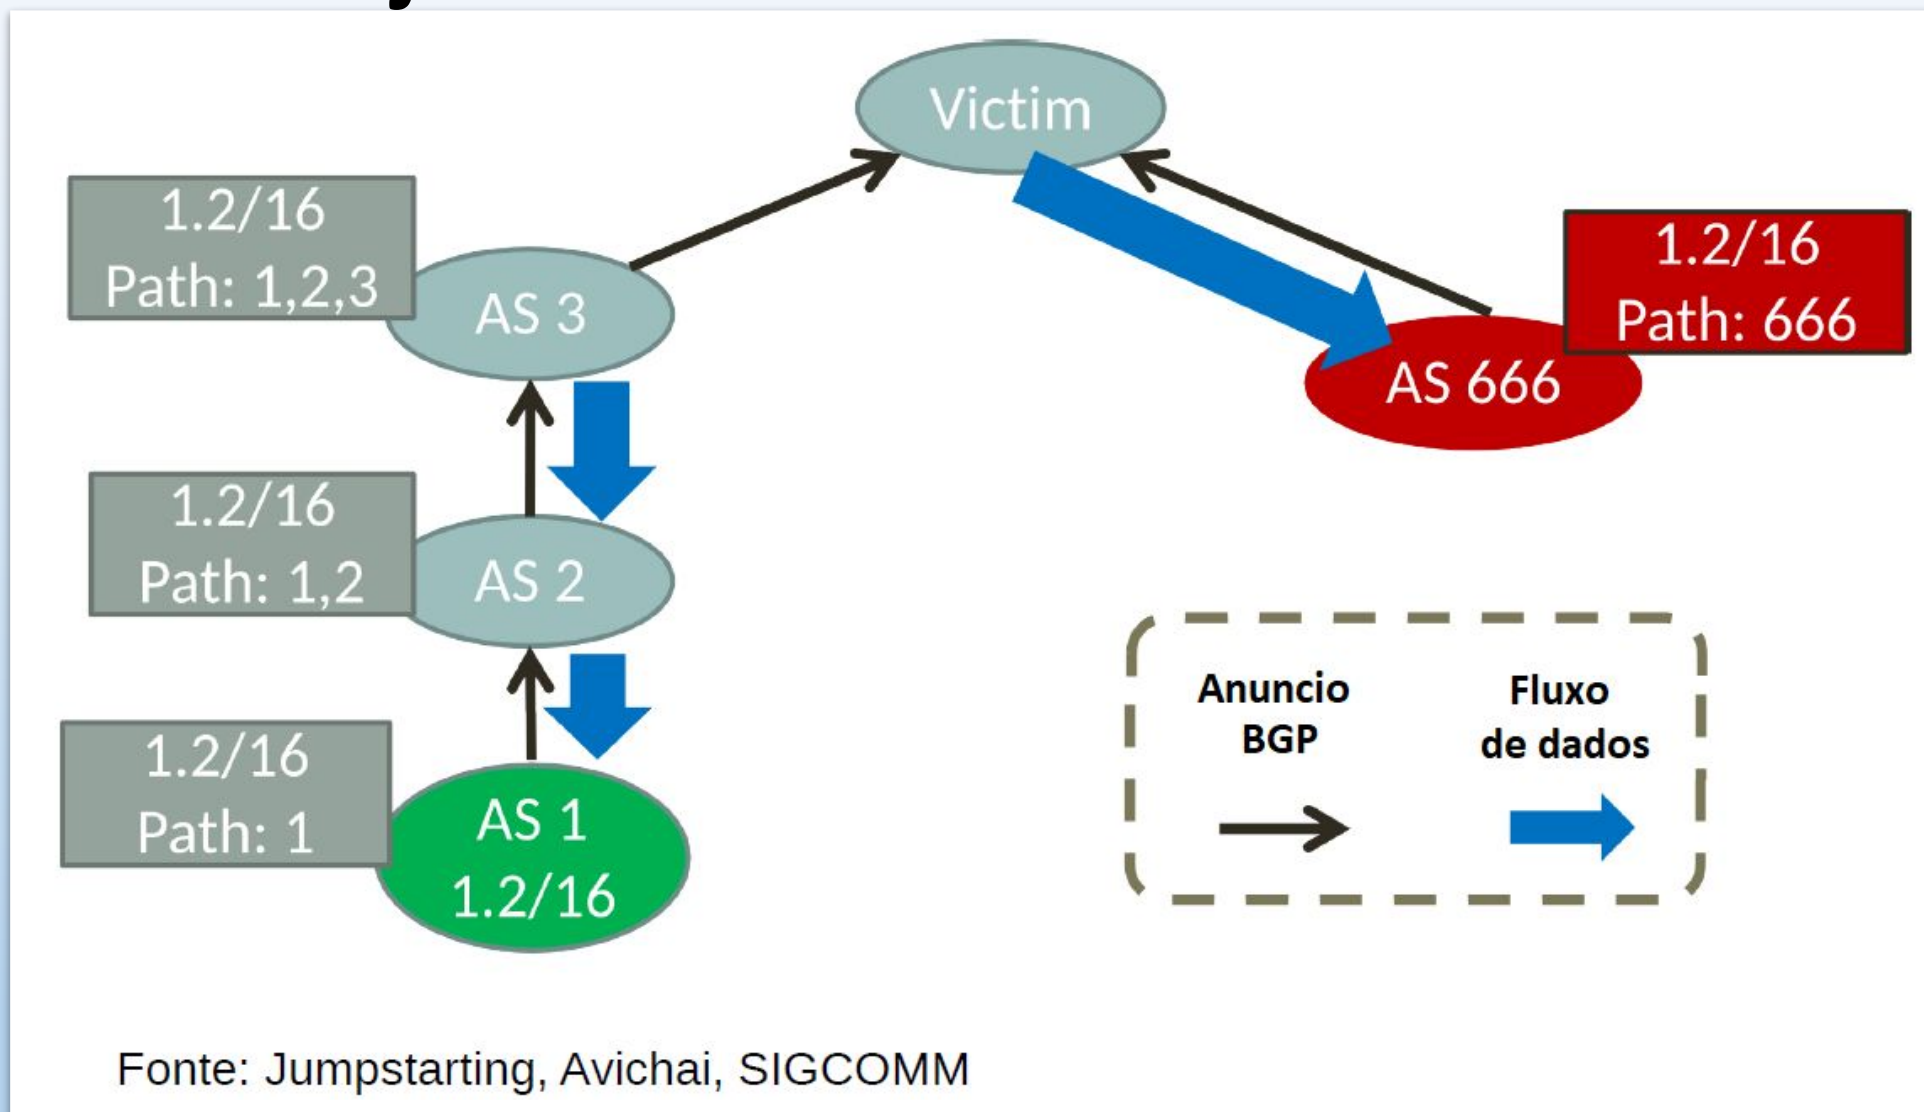
Douglas Fernando Fischer

- Engenheiro de Controle e Automação
- Atua na área de redes de telecomunicações desde 1999
- Trabalhou como engenheiro de pré-vendas e implantação em integradores de tecnologia
- Consultor na área de redes e servidores no segmento corporativo e provedores de Internet
- BPF – <http://brasilpeeringforum.org/>
- Tretísta com fins produtivos nas horas vagas

Intenções dessa Apresentação?

- Linguagem menos rebuscada - Pequenos ISPs Brasileiros.
- Coletânea de outros trabalhos.
- Route-Leak e BGP-Hijack.
- Casos Recentes que impactaram o Brasil.
- “Dedo-Gordo” vs “Má Fé”
- Oque você tem a ver com isso? – Parte 1
- Como proteger-se?
- Proteções que ajudam mas não resolvem.
 - Número máximo de prefixos.
 - AS-Path e Expressões regulares.
 - Listas de prefixos manuais
- IRR e automatização de Filtros.
- Os problemas do IRR e como estão sendo resolvidos.
- Os filtros de Prefixos do IX.BR.
- Oque você tem a ver com isso? – Parte 2

Route-Leak e BGP-Hijack.



Italo Valcy italovalcy@ufba.br – GTS 29 – Foz do Iguaçu – 26/05/2017

<https://youtu.be/x8GMBYwinKk?t=9460> –

<ftp://ftp.registro.br/pub/gts/gts29/05-Seguranca-BGP.pdf>

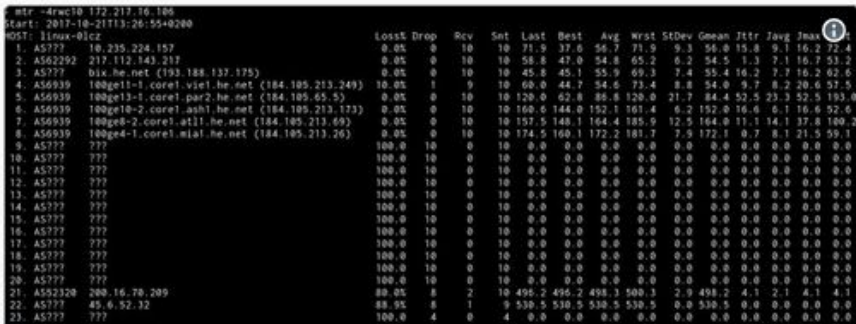
LACNIC31 – Maio/2019 – Automação Prefix-Lists BGP –

Casos Recentes que impactaram o Brasil

Today's BGP leak in Brazil

Posted by Andree Toonk - October 21, 2017 - News and Updates - No Comments

Earlier today several people noticed network reachability problems for networks such as Twitter, Google and others. The root cause turned out to be another BGP mishap.



	Loss%	Drop	Rcv	Sit	Last	Best	Avg	Wrst	StDev	Mean	Jtr	Javg	Jmax
1. AS777 10.225.224.157	0.0%	0	10	10	71.9	37.6	56.7	71.9	9.3	56.0	15.0	9.1	16.2
2. AS62292 217.112.143.217	0.0%	0	10	10	58.8	42.0	54.8	65.2	6.2	54.5	1.3	7.1	16.7
3. AS777 10.225.224.157	0.0%	0	10	10	45.8	45.1	55.9	69.3	7.4	55.4	16.2	7.7	16.2
4. AS6939 100ge1-1.core1.vie1.he.net (104.105.213.249)	10.0%	1	9	10	60.0	44.7	54.6	73.4	8.8	54.0	9.7	8.2	20.6
5. AS6939 100ge1-1.core1.par2.he.net (104.105.65.5)	0.0%	0	10	10	120.0	62.8	88.8	120.0	21.7	84.4	52.5	23.3	52.5
6. AS6939 100ge1-2.core1.ash1.he.net (104.105.213.173)	0.0%	0	10	10	100.0	144.0	152.1	161.4	6.2	152.0	16.0	6.1	16.6
7. AS6939 100ge1-2.core1.atl1.he.net (104.105.213.69)	0.0%	0	10	10	157.5	148.1	164.4	185.9	12.5	164.0	11.1	14.1	37.8
8. AS6939 100ge1-1.core1.mia1.he.net (104.105.213.26)	0.0%	0	10	10	174.5	160.1	172.2	181.7	7.5	172.1	0.7	8.1	21.5
9. AS777 777	100.0%	10	0	10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
10. AS777 777	100.0%	10	0	10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
11. AS777 777	100.0%	10	0	10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
12. AS777 777	100.0%	10	0	10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
13. AS777 777	100.0%	10	0	10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
14. AS777 777	100.0%	10	0	10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
15. AS777 777	100.0%	10	0	10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
16. AS777 777	100.0%	10	0	10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
17. AS777 777	100.0%	10	0	10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
18. AS777 777	100.0%	10	0	10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
19. AS777 777	100.0%	10	0	10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
20. AS777 777	100.0%	10	0	10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
21. AS52320 200.16.70.209	80.0%	8	2	10	495.2	495.2	495.3	500.3	2.9	495.2	4.1	2.1	4.1
22. AS777 45.6.52.32	88.9%	8	1	9	530.5	530.5	530.5	530.5	0.0	530.5	0.0	0.0	0.0
23. AS777 777	100.0%	4	0	4	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0



Some Google services seem to have been hijacked for roughly 15 minutes. Seen anything? @atoonk @bgpmon @bgpstream
MTR: xor.meo.ws/P0SYOU7j-4Ftjl...

27 9:32 AM - Oct 21, 2017

23 people are talking about this

Between 11:09 and 11:27 UTC traffic for many large CDN was rerouted through Brazil. Below an example for the Internet's most famous prefix 8.8.8.0/24 (Google DNS)
At 2017-10-21 11:09:59 UTC, AS33362, an US based ISP saw the path towards Google's 8.8.8.0/24 like this:

33362 6939 16735 263361 15169

April 26th, 2018

BGP hijacks - Malicious or Mistakes?

A few days ago several cybersecurity resources [reported details](#) of an entirely malicious traffic redirection that combined DNS, and BGP hijacking. The primary goal of this attack was to steal money from different cryptocurrency wallets and services. Moreover, it was successful, since Amazon did not detect it in time.
Today, on April 26, another significant incident happened that seems to be also unnoticed by the majority of players.

An [AS267286](#), registered almost two years ago, stayed invisible until the event we are going to cover below when it announced 28 prefixes to the outer world. Among those 28 separate announcements **sixteen** were /8 prefixes (6,25% of IPv4 address space). This initial announcement was accepted by ASNs that belong to China Telecom ([AS4134](#), [AS4809](#)), which in its turn propagated it to Tier1 carriers and thus helped to spread it all over the world.

A spread of /8 prefixes on their own does not always affect end-user services or applications. To redirect traffic using /8 prefix, several conditions are necessary:

- The receiving AS has only partial view: it is connected to IX(es) and accepts all routes from that source, but accepts only default routes from upstream providers.
- The /8 is distributed through IX, while legitimate more specific routes are not present there.

With high probability, we can state that those /8 prefixes were distributed at São Paulo IX, the biggest IX in Brazil. Furthermore, several other networks were affected, of the size ranging from /24 and up to /16, belonging to several companies, including Equinix and Incapsula. There were already 5 waves with the same set of prefixes:

Casos Recentes que impactaram o Brasil

BGPmon.net
@bgpmon

Following

looking into BGP leak incident involving @google prefixes, AS37282 out of Nigeria and China Telecom.

3:40 AM - 13 Nov 2018

54 Retweets 48 Likes



MainOne
@Mainoneservice

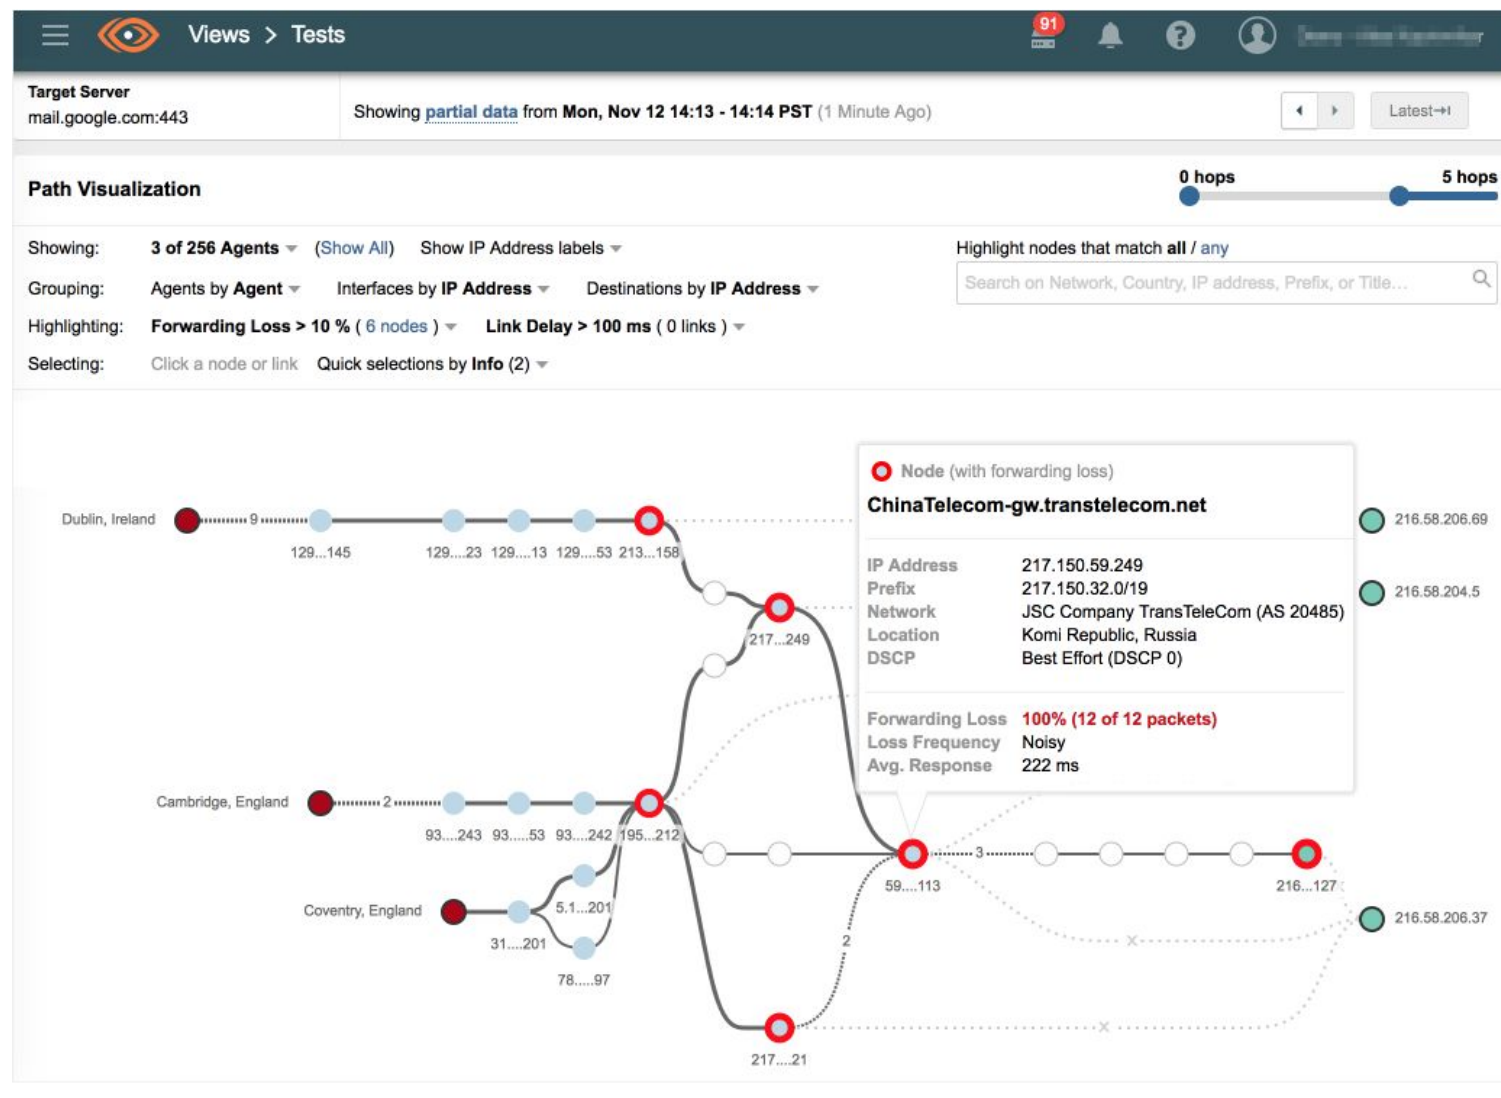
Follow

Replying to @bgpmon @Google

We have investigated the advertisement of @Google prefixes through one of our upstream partners. This was an error during a planned network upgrade due to a misconfiguration on our BGP filters. The error was corrected within 74mins & processes put in place to avoid reoccurrence

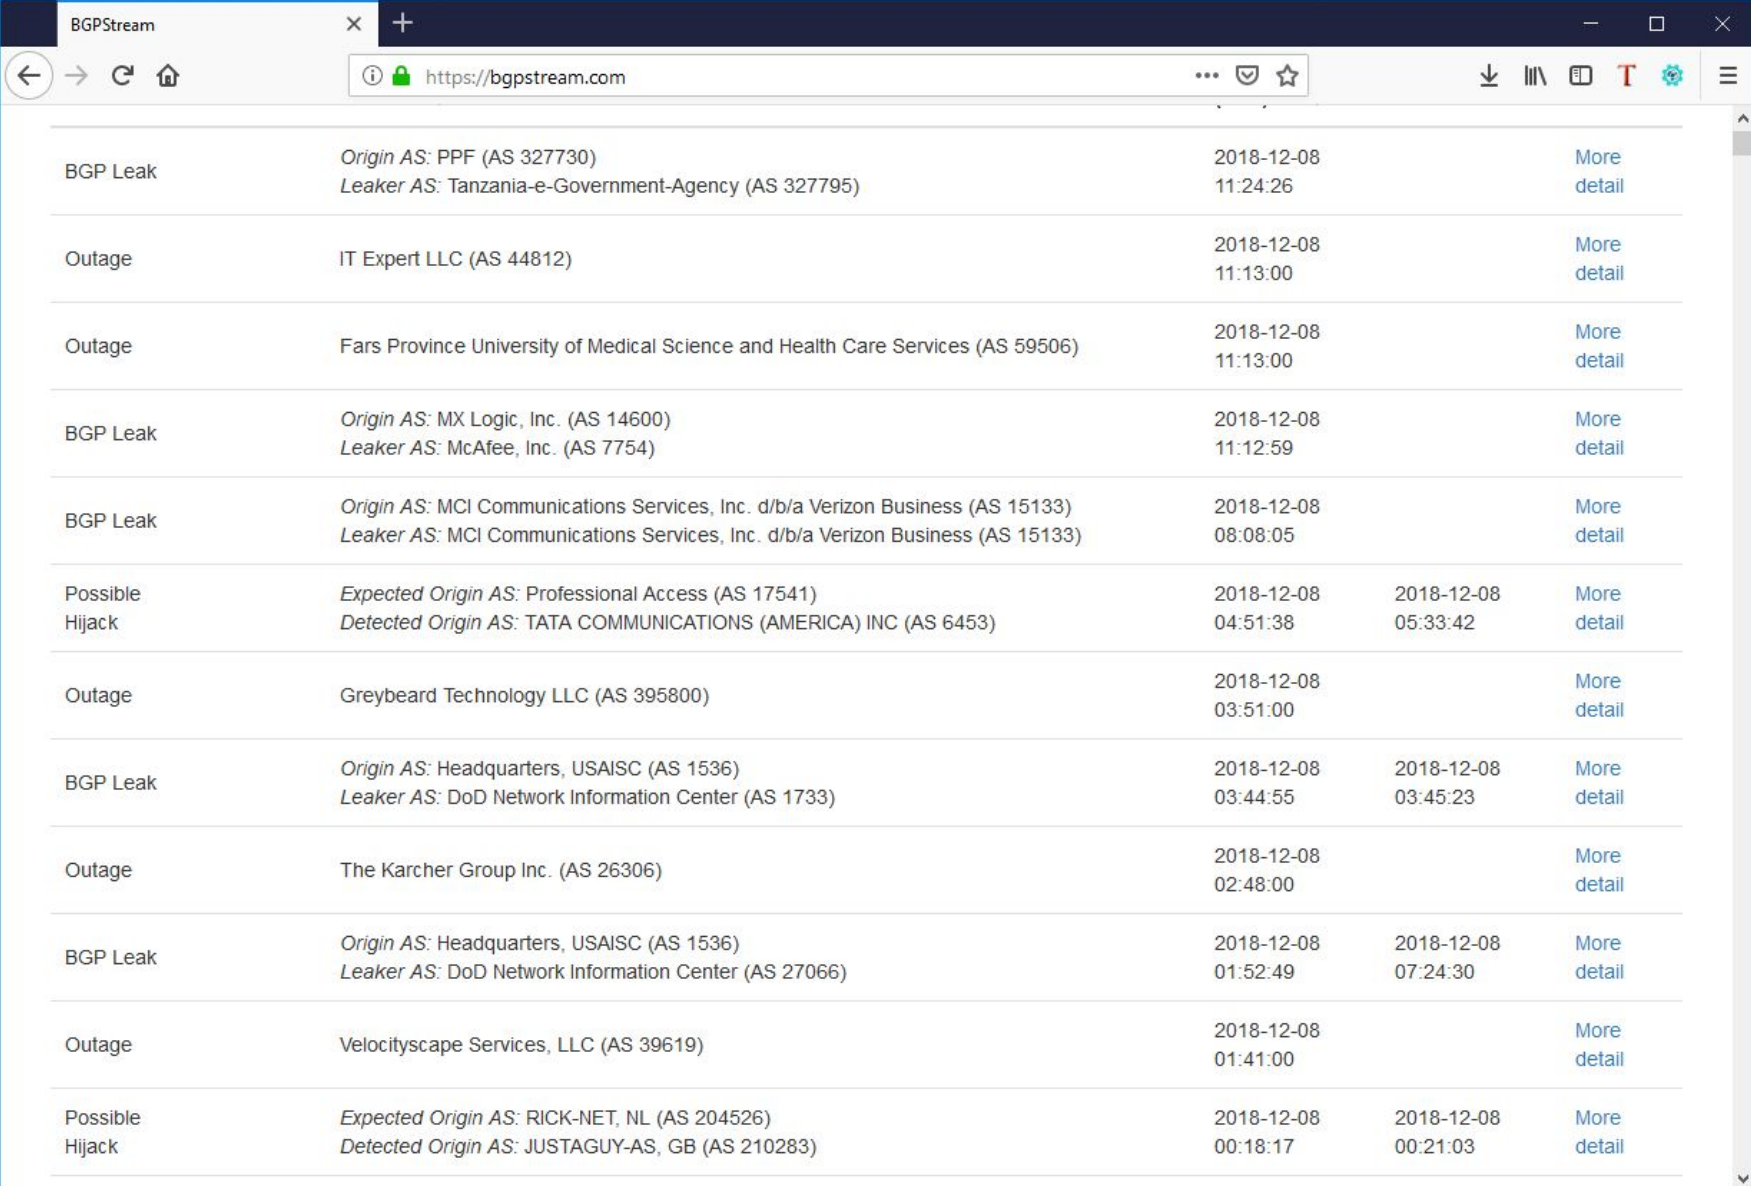
5:29 PM - 13 Nov 2018

38 Retweets 50 Likes



<https://blog.thousandeyes.com/internet-vulnerability-takes-down-google/>

Como acompanhar os Hijacks e Route-Leaks?



The screenshot shows the BGPStream website interface. The browser address bar displays 'https://bgpstream.com'. The main content area is a table listing various network events. The table has columns for event type, origin/leaker information, and timestamps. Some rows include a 'More detail' link.

BGP Leak	Origin AS: PPF (AS 327730) Leaker AS: Tanzania-e-Government-Agency (AS 327795)	2018-12-08 11:24:26		More detail
Outage	IT Expert LLC (AS 44812)	2018-12-08 11:13:00		More detail
Outage	Fars Province University of Medical Science and Health Care Services (AS 59506)	2018-12-08 11:13:00		More detail
BGP Leak	Origin AS: MX Logic, Inc. (AS 14600) Leaker AS: McAfee, Inc. (AS 7754)	2018-12-08 11:12:59		More detail
BGP Leak	Origin AS: MCI Communications Services, Inc. d/b/a Verizon Business (AS 15133) Leaker AS: MCI Communications Services, Inc. d/b/a Verizon Business (AS 15133)	2018-12-08 08:08:05		More detail
Possible Hijack	Expected Origin AS: Professional Access (AS 17541) Detected Origin AS: TATA COMMUNICATIONS (AMERICA) INC (AS 6453)	2018-12-08 04:51:38	2018-12-08 05:33:42	More detail
Outage	Greybeard Technology LLC (AS 395800)	2018-12-08 03:51:00		More detail
BGP Leak	Origin AS: Headquarters, USAISC (AS 1536) Leaker AS: DoD Network Information Center (AS 1733)	2018-12-08 03:44:55	2018-12-08 03:45:23	More detail
Outage	The Karcher Group Inc. (AS 26306)	2018-12-08 02:48:00		More detail
BGP Leak	Origin AS: Headquarters, USAISC (AS 1536) Leaker AS: DoD Network Information Center (AS 27066)	2018-12-08 01:52:49	2018-12-08 07:24:30	More detail
Outage	Velocityscape Services, LLC (AS 39619)	2018-12-08 01:41:00		More detail
Possible Hijack	Expected Origin AS: RICK-NET, NL (AS 204526) Detected Origin AS: JUSTAGUY-AS, GB (AS 210283)	2018-12-08 00:18:17	2018-12-08 00:21:03	More detail

<https://bgpstream.com/>

<https://twitter.com/bgpstream>

“Dedo-Gordo” vs “Má Fé”

- Se existe qualquer possibilidade de ser um equívoco, só deixe para supor má fé em **último caso!**
- Se você trabalha diariamente com BGP em nível médio há 2 anos ou mais, e nunca cometeu um erro de configuração, por favor levante a mão...

• Leak – Dedo-Gordo

- Erro de configuração
- Bugs de software
 - Mecanismo de Roteamento
 - Automação de Configurações

• Hijack - Má Fé

- Concorrência
- Ataques dirigidos
- Posseiro de IPs não utilizados

O que você tem a ver com isso? – Parte 1

- Você gostaria ser acordado às 03:00am porque alguém, em algum lugar do mundo, foi fazer um rearranjo de rotas, e acabou sequestrando os prefixos do seu ASN e deixando os seu clientes sem conectividade?
- Você gostaria de ser acusado de negligência porque um de seus clientes de Trânsito cometeu algum equívoco de configuração?
- 2 Regras para vida em comunidade
 - (Vale também para BGP, Spoofing, e muitas outras coisas)
 - Inversão dos papéis
 - Extrapolação

Como se proteger de Hijacks e Route-Leaks?

A parte boa da resposta é que existem algumas técnicas que pode te ajudar!

- Número máximo de prefixos
- AS-Path
- Prefix-List
 - Criados Manualmente
 - Baseados e IRR
- BGPSEc
- RPKI


A parte triste da resposta é que o que protege os seu prefixos são as ações dos outros ASNs. Ou seja:

- a) Se os outros forem omissos, **você dança!**
- b) Se você for omissos, os outros dançam, e **você fica com má reputação!**

Filtros de BGP – Aonde e em que direção?

- Attachment points
- Directionality

“One man’s ebgp-out is another man’s ebgp-in.”
– ancient Dutch proverb



Job Snijders job@ntt.net – LACNIC 30 – Rosário – 26/09/2018 – <https://youtu.be/SOA5EONXDKU>
https://www.lacnic.net/innovaportal/file/3208/1/lacnic30_snijders_routing_policy_architecture.pdf

Proteções que ajudam mas não resolvem

Número máximo de prefixos

What happens when limits are applied in pre-policy during a full table leak:




Job Snijders job@ntt.net – LACNIC 30 – Rosário – 26/09/2018 – <https://youtu.be/SOA5EONXDKU>
https://www.lacnic.net/innovaportal/file/3208/1/lacnic30_snijders_routing_policy_architecture.pdf

Proteções que ajudam mas não resolvem

Número máximo de prefixos

What happens when limits are applied post-policy



Job Snijders job@ntt.net – LACNIC 30 – Rosário – 26/09/2018 – <https://youtu.be/SOA5EONXDKU>
https://www.lacnic.net/innovaportal/file/3208/1/lacnic30_snijders_routing_policy_architecture.pdf

Proteções que ajudam mas não resolvem

AS-Path e Expressões Regulares

- Apostila do Rinaldo Vaz sobre BGP – Seção 4.4 – Entendendo o BGP REGEXP
- https://drive.google.com/file/d/1u_44sw8OKx4pe-iJLLO_RuSZzZ_p4g3w/view
- Ajuda a testar coisas como:
 - O ASN de origem do Prefixo
 - Quantidade de Prepends
 - Número máximo de Saltos de ASN no AS-Path
 - Se passou por algum ASN que não se deseja.
 - Se o First-AS é realmente o do ASN do peering
 - Recurso padrão em quase todas Engines BGP
 - Não muito útil em Sessões com Route-Servers de IX
- REGEXP costumam ser monothread.
- Se forem mal escritas ou mal utilizadas espancam a CPU.
- AS-Path é completamente “reescrevível”.
 - Logo, **não podemos confiar!**

Listas de prefixo

Filtros de entrada

Examples

This example shows how to configure a prefix list and apply it to a Border Gateway Protocol (BGP) peer:

```
switch# configure terminal
switch(config)# ip prefix-list allowprefix 10 permit 192.0.2.0 eq 24
switch(config)# ip prefix-list allowprefix 20 permit 209.165.201.0 eq 27
switch(config) router bgp 65536:20
switch(config-router)# neighbor 192.0.2.1/16 remote-as 65536:20
switch(config-router-neighbor)# address-family ipv4 unicast
switch(config-router-neighbor-af)# prefix-list allowprefix in
switch(config-router-neighbor-af)#
```

https://www.cisco.com/c/m/en_us/techdoc/dc/reference/cli/nxos/commands/bgp/ip-prefix-list.html

```
/routing bgp peer
add instance=Internet name=MyPeer in-filter=192.0.2.1 remote-address=192.0.2.1 remote-as=65536

/routing filter
add action=accept chain=allowprefix prefix=192.0.2.0/24
add action=accept chain=allowprefix prefix=209.165.201.0/27
add action=discard chain=allowprefix comment="Explicit Deny"
```

Listas de prefixo – O Tal “Cone de AS”

Filtros de entrada

- Quais prefixos devem estar nos filtros de entrada?

Imaginemos um Exemplo


Digamos que cada um desses ASNs tenha um prefixo que seja 10.0.<ASN>.0/16.

Como ficariam os filtros de entrada nos peers de downstream dos ASNs abaixo?

AS4

AS2

AS1



```
#Configs do ASN 4
#####
/routing bgp peer
add instance=Internet name=AS8-v4 in-filter=PeerFilter-AS8-v4-IN remote-address=?.?.?.? remote-as=8
```

```
/routing filter
add action=accept chain=PrefixList-AS8-v4 prefix=10.0.8.0/16 prefix-length=16-24
```


```
add action=accept chain=PeerFilter-AS8-v4-IN match-chain=PrefixList-AS8-v4
add action=discard chain=PeerFilter-AS8-v4-IN comment="Explicit Deny"
```

```
#Configs do ASN 2
#####
/routing bgp peer
add instance=Internet name=AS4-v4 in-filter=PeerFilter-AS4-v4-IN remote-address=?.?.?.? remote-as=4
add instance=Internet name=AS5-v4 in-filter=PeerFilter-AS5-v4-IN remote-address=?.?.?.? remote-as=5
```

```
/routing filter
add action=accept chain=PrefixList-AS4-v4 prefix=10.0.4.0/16 prefix-length=16-24
add action=accept chain=PrefixList-AS5-v4 prefix=10.0.5.0/16 prefix-length=16-24
add action=accept chain=PrefixList-AS8-v4 prefix=10.0.8.0/16 prefix-length=16-24
add action=accept chain=PrefixList-AS10-v4 prefix=10.0.10.0/16 prefix-length=16-24
```

```
add action=accept chain=PeerFilter-AS4-v4-IN match-chain=PrefixList-AS4-v4
add action=accept chain=PeerFilter-AS4-v4-IN match-chain=PrefixList-AS8-v4
add action=discard chain=PeerFilter-AS4-v4-IN comment="Explicit Deny"
```

```
add action=accept chain=PeerFilter-AS4-v4-IN match-chain=PrefixList-AS5-v4
add action=accept chain=PeerFilter-AS4-v4-IN match-chain=PrefixList-AS10-v4
add action=discard chain=PeerFilter-AS4-v4-IN comment="Explicit Deny"
```



#Configs do ASN 1

#####

/routing bgp peer

add instance=Internet name=AS2-v4 in-filter=PeerFilter-AS2-v4-IN remote-address=?.??.??.? remote-as=2

add instance=Internet name=AS3-v4 in-filter=PeerFilter-AS3-v4-IN remote-address=?.??.??.? remote-as=3

/routing filter

add action=accept chain=PrefixList-AS2-v4 prefix=10.0.2.0/16 prefix-length=16-24

add action=accept chain=PrefixList-AS3-v4 prefix=10.0.3.0/16 prefix-length=16-24

add action=accept chain=PrefixList-AS4-v4 prefix=10.0.4.0/16 prefix-length=16-24

add action=accept chain=PrefixList-AS5-v4 prefix=10.0.5.0/16 prefix-length=16-24

add action=accept chain=PrefixList-AS6-v4 prefix=10.0.6.0/16 prefix-length=16-24

add action=accept chain=PrefixList-AS7-v4 prefix=10.0.7.0/16 prefix-length=16-24

add action=accept chain=PrefixList-AS8-v4 prefix=10.0.8.0/16 prefix-length=16-24

add action=accept chain=PrefixList-AS9-v4 prefix=10.0.9.0/16 prefix-length=16-24

add action=accept chain=PrefixList-AS10-v4 prefix=10.0.10.0/16 prefix-length=16-24

add action=accept chain=PeerFilter-AS2-v4-IN match-chain=PrefixList-AS2-v4

add action=accept chain=PeerFilter-AS2-v4-IN match-chain=PrefixList-AS4-v4

add action=accept chain=PeerFilter-AS2-v4-IN match-chain=PrefixList-AS8-v4

add action=accept chain=PeerFilter-AS2-v4-IN match-chain=PrefixList-AS5-v4

add action=accept chain=PeerFilter-AS2-v4-IN match-chain=PrefixList-AS10-v4

add action=discard chain=PeerFilter-AS2-v4-IN comment="Explicit Deny"

add action=accept chain=PeerFilter-AS3-v4-IN match-chain=PrefixList-AS3-v4


add action=accept chain=PeerFilter-AS3-v4-IN match-chain=PrefixList-AS6-v4

add action=accept chain=PeerFilter-AS3-v4-IN match-chain=PrefixList-AS7-v4

add action=accept chain=PeerFilter-AS3-v4-IN match-chain=PrefixList-AS10-v4

add action=accept chain=PeerFilter-AS3-v4-IN match-chain=PrefixList-AS9-v4

add action=discard chain=PeerFilter-AS3-v4-IN comment="Explicit Deny"



Proteções que ajudam mas não resolvem

Prefix-Lists mantidas manualmente.

“Ain Douglas... Você está sendo bobo, feio, e chato!
Como assim prefix-list manuais não resolvem?
Não viu ali? Deu um trabalhinho, mas resolveu...”



É Besmo
Abiguinho?

Listas de prefixo – O Tal “Cone de AS”

Filtros de entrada


- Mantendo as prefix-lists **manualmente**

Agora imaginemos um novo Exemplo

Digamos que cada um desses ASNs 15-20 prefixos iniciais.

Digamos que a cada 3 meses haja uma saída ou uma entrada de prefixo de cada um desses prefixos.

Digamos que a cada mês cada um desses ASNs muda alguma dessas relações de Peering ou Trânsito.



Proteções que ajudam mas não resolvem

Prefix-Lists mantidas manualmente.

“Poxa, seria legal se tivesse um jeito de cada ASN cadastrar em algum lugar os prefixos que vão anunciar, e também se tivesse um jeito de saber com quem cada um desses ASNs conversa...

Aí quem sabe a gente poderia criar uma forma de automatizar essas tarefas de operação de filtros.”

Um tal de IRR

IRR - O que é?

- é um sistema global de bases de dados que armazenam e compartilham informações sobre políticas de roteamento.
- estende os “whois” tradicionais (registros de numeração).



Um tal de IRR

What is the Internet Route Registry?

- A distributed database of route and route-related information.
- Objects are defined in the Route Policy Specification Language (RPSL - RFC 2622, RFC 4012)
- The objects in the database are publicly available for service providers and other users to utilize for various purposes

The Three Essential Route Registry Objects

Maintainer

Defines the person or group responsible for updating route registry objects

Route

Defines an route/AS Number relationship

AS-SET

Defines your customer cone
(Customers that peer with you)

Brian Foust brian@ntt.net – IX Forum 9 – São Paulo – 07/12/2015

<http://ix.br/ixforum/9/slides/ixbr9-irr.pdf> – <https://youtu.be/gvQzOEbWMaY?t=2h56m16s>

Um tal de IRR

```
route: 200.15.0.0/16
descr: NTT Communications - NTTB-200-015
origin: AS2914
remarks: this is non-portable space, no exceptions
remarks: contacts per RFC2142:
remarks: Abuse / UCE reports abuse@ntt.net
remarks: Security issues security@ntt.net
mnt-by: MAINT-NTTCOM-BB
changed: brian@ntt.net 20151118
source: NTTCOM
```

Optional: Remarks

Required

Route

Description

Origin
Number

Maintainer

Last Update

Route
Registry

Maintainer Object Attributes

```
$ whois -h rr.ntt.net MAINT-NTTCOM-BB
mntner: MAINT-NTTCOM-BB
descr: NTT Communications Global IP Network
maintainer
admin-c: JH636-ARIN
tech-c: JH636-ARIN
upd-to: ip-eng-reports@us.ntt.net
mnt-nfy: ip-eng-reports@us.ntt.net
auth: MD5-PW XXXXXX
remarks: contacts per RFC2142:
remarks: Abuse / UCE reports abuse@ntt.net
remarks: Security issues security@ntt.net
notify: ip-eng-routing@us.ntt.net
mnt-by: MAINT-NTTCOM-BB
changed: tboudreau@us.ntt.net 20151028
source: NTTCOM
```

Required

Maintainer
Handle

Description

Admin
Contact

Auth Error
Recipient

Authentication
Scheme

Last Update

Route
Registry DB

Tech Contact

Remarks

Brian Foust brian@ntt.net – IX Forum 9 – São Paulo – 07/12/2015

<http://ix.br/ixforum/9/slides/ixbr9-irr.pdf> – <https://youtu.be/gvQzOEbWMaY?t=2h56m16s>

Já viram essas telinhas?

The screenshot shows the 'Manutenção ASN' (ASN Maintenance) page on the Registro.br website. The page has a dark blue header with the Registro.br logo and navigation links: 'Sobre Domínios', 'Tecnologia', 'Ajuda', 'Quem Somos', 'Contato', and 'Registre'. Below the header, there's a breadcrumb trail: 'Home > Sistema > Manutenção ASN'. The main content area has three tabs: 'DOMÍNIOS', 'TITULARIDADES', and 'NUMERAÇÃO' (which is active). Under the 'NUMERAÇÃO' tab, there's a section titled 'ASN NNNN'. Below this title, a note states: 'Os campos as-in e as-out devem ser fornecidos obedecendo a sintaxe da RFC 1786'. There are two input fields: 'AS-IN' and 'AS-OUT', both currently empty.

The screenshot shows the 'Whois' page on the Registro.br website for AS22548. The page has a dark blue header with the 'Whois' title and a search bar containing the URL 'https://registro.br/2/whois?qr=AS22548#lresp'. The main content area is divided into three sections: 'Blocos', 'AS-in', and 'AS-out'. The 'Blocos' section lists two IP blocks: '200.160.0.0/20' and '2001:12ff::/32'. The 'AS-in' section lists five BGP peers with their respective AS numbers and policies: 'from AS2914 100 accept ANY', 'from AS3549 100 accept ANY', 'from AS12989 100 accept ANY', 'from AS16735 100 accept ANY', and 'from AS52320 100 accept ANY'. The 'AS-out' section lists five BGP peers with their respective AS numbers and policies: 'to AS2914 announce AS22548', 'to AS3549 announce AS22548', 'to AS12989 announce AS22548', 'to AS16735 announce AS22548', and 'to AS52320 announce AS22548'.

IRR Explorer

<http://irrexplorer.nlnog.net/search/22548>

AS Number: 22548

irrexplorer.nlnog.net/search/22548

AS Number: 22548

Prefixes

prefix	bgp	radb	advice
200.128.0.0/9		7465,14650	Not seen in BGP, but (legacy?) route-objects exist, consider clean-up
200.160.0.0/20	22548	22548	Looks good: BGP origin consistent with AS in route-objects
2001:12ff::/32	22548	22548	Looks good: BGP origin consistent with AS in route-objects

Showing 1 to 3 of 3 entries
Offending prefixes are only found if initial prefix sets is smaller than 1000

Included in the following macros:

as_macro	nttcom	radb	level3	tc	altdb	ripe
AS-CTBC-CUSTOMERS1		✓				
AS-DNS-BR		✓				
AS-EMBRATELNET	✓	✓				
AS-GNT-CUSTOMERS		✓				
AS-HURRICANE		✓				
AS-HURRICANE-DT		✓				
AS-HURRICANEV6		✓				

Ferramentas para automatização

IRR – Ferramentas de Apoio

IRRToolSet^[29]


- conjunto de ferramentas destinadas à manipulação de políticas de roteamento

IRR Power Tools^[30]

- monitora e gerencia objetos IRR

BGPQ3^[31]

- versão simplificada do rtconfig



29

Herbert Faleiros herbert@faleiros.eti.br – GTER 31

São Paulo – 13/05/2011 – <ftp://ftp.registro.br/pub/gter/gter31/03-IRR-SCW.pdf>

Ferramentas para automatização – bgpq3

For Cisco we can use aggregation (-A) flag to make this prefix-filter more compact:

```
user@host:~>bgpq3 -A1 eltel AS20597
no ip prefix-list eltel
ip prefix-list eltel permit 81.9.0.0/20
ip prefix-list eltel permit 81.9.32.0/20
ip prefix-list eltel permit 81.9.96.0/20
ip prefix-list eltel permit 81.222.128.0/20
ip prefix-list eltel permit 81.222.192.0/18
ip prefix-list eltel permit 85.249.8.0/21
ip prefix-list eltel permit 85.249.224.0/19
ip prefix-list eltel permit 89.112.0.0/18 ge 19 le 19
ip prefix-list eltel permit 89.112.4.0/22
ip prefix-list eltel permit 89.112.64.0/19
ip prefix-list eltel permit 217.170.64.0/19 ge 20 le 20
```

<https://github.com/snar/bgpq3>

Generating named Juniper prefix-filter for AS20597 :

```
user@host:~>bgpq3 -J1 eltel AS20597
policy-options {
  replace:
    prefix-list eltel {
      81.9.0.0/20;
      81.9.32.0/20;
      81.9.96.0/20;
      81.222.128.0/20;
      81.222.192.0/18;
      85.249.8.0/21;
      85.249.224.0/19;
      89.112.0.0/19;
      89.112.4.0/22;
      89.112.32.0/19;
      89.112.64.0/19;
      217.170.64.0/20;
      217.170.80.0/20;
    }
}
```


Ferramentas para automatização – bgpq3

“Mas e o mititiki?”



Douglas Fischer <fischerdouglas@gmail.com>

para Grupo ▾

Resolvi reler com mais atenção a documentação do BGPQ3, e encontrei a parte de User-Defined Format.

<https://github.com/snar/bgpq3#user-defined-format>

Acredito que deve estar bem fácil de resolver só com esse recur

```
arp@arpoador:~/bgpq3$ bgpq3 -F "add action=accept chain=PrefixList-LuckyNet-V4-Automated prefix-length=%1-24 \\n" as3254
```

```
add action=accept chain=PrefixList-LuckyNet-V4-Automated prefix
```

```
add action=accept chain=PrefixList-LuckyNet-V4-Automated prefix
```

```
add action=accept chain=PrefixList-LuckyNet-V4-Automated prefix
```

```
add action=accept chain=PrefixList-LuckyNet-V4-Automated prefix
```

```
add action=accept chain=PrefixList-LuckyNet-V4-Automated prefix
```

```
add action=accept chain=PrefixList-LuckyNet-V4-Automated prefix
```

```
add action=accept chain=PrefixList-LuckyNet-V4-Automated prefix
```


Montei um exemplo MUITO SIMPLES de filtragem(apenas prefixos) c

Cenário hipotético onde a Lucky Net (as3254) fosse Downstream d

seg, 23 de jul 12:13



MikroTik Support #48



cdavid14 wants to merge 4 commits into snar:master from cdavid14:master



Conversation 0



Commits 4



Checks 0



Files changed 4



cdavid14 commented on 13 Dec 2018

Some functions for MikroTik support added.



1



cdavid14 added some commits on 13 Dec 2018



Mikrotik Dev(part 1)



Mikrotik Dev(part 2)



README Update



README Update 2

<https://eng.registro.br/pipermail/gter/2018-July/074807.html>

Ferramentas para automação

Colocando essas prefix-list pra dentro do seu router.

- SSH + BASH + Cron
 - Funciona, porém é engessado, difícil de manter e é feio.
- Ansible – <https://docs.ansible.com/>
 - Trabalhoso de iniciar, depende dos módulos para cada sintaxe.
- Napalm – <https://github.com/napalm-automation/napalm>
 - Acabei de conhecer, mas aparenta ser algo muito funcional.
- Problemas com o RouterOS
 - Resend involuntário quando se altera os filtros.

Relembrando do nosso desafio

Agora imaginemos um novo Exemplo

Digamos que cada um desses ASNs 15-20 prefixos iniciais.

Digamos que a cada 3 meses haja uma saída ou uma entrada de prefixo de cada um desses prefixos.

Digamos que a cada mês cada um desses ASNs muda alguma dessas relações de Peering ou Trânsito.



IRR – Quais são os problemas

IRR – what is broken what can be fixed?

- Some IRRdbs do not perform validation
 - Meaning that virtually anyone can create virtually any route/route6 object and sneak those into the prefix-filters
- Eleven relevant IRRs not validating: RIPE, NTTCOM, RADB, ALTDB, ARIN IRR, BBOI, BELL, LEVEL3, RGNET, TC, CANARIE
- Two solutions:
 - Lock the database down (RIPE / RIPE-NONAUTH)
 - Filter on the mirror level


Job Snijders job@ntt.net – LACNIC 30 – Rosário – 26/09/2018
https://www.lacnic.net/innovaportal/file/3135/1/lacnic30_snijders_routing_security_roadmap.pdf

Herbert Faleiros herbert@faleiros.eti.br

GTER 31 – São Paulo – 13/05/2011

<ftp://ftp.registro.br/pub/gter/gter31/03-IRR-SCW.pdf>

IRR – e no Brasil?



- cerca de 80% dos prefixos nacionais anunciados tem IRR^[28]
- destes, 80% deles são proxies
- apenas 60% são válidos

IRR – Ações para resolver os problemas

ARIN community also recognized this is an issue

- Consultation at [NANOG](#) and through [ARIN-Consult](#) mailing list
- https://www.arin.net/vault/resources/routing/2018_roadmap.html
- <https://teamarin.net/2018/07/12/the-path-forward/>

“Improve, or kill it”

Job Snijders job@ntt.net – LACNIC 30 – Rosário – 26/09/2018
https://www.lacnic.net/innovaportal/file/3135/1/lacnic30_snijders_routing_security_roadmap.pdf

ALMOST SOLVED

RIPE NWI-5 proposal & implementation

- RIPE NCC’s IRR previously allowed anyone to register any non-RIPE-managed space if it had not yet been registered. *DANGER*
- The “RPSL” password & maintainer was used for this

SOLVED

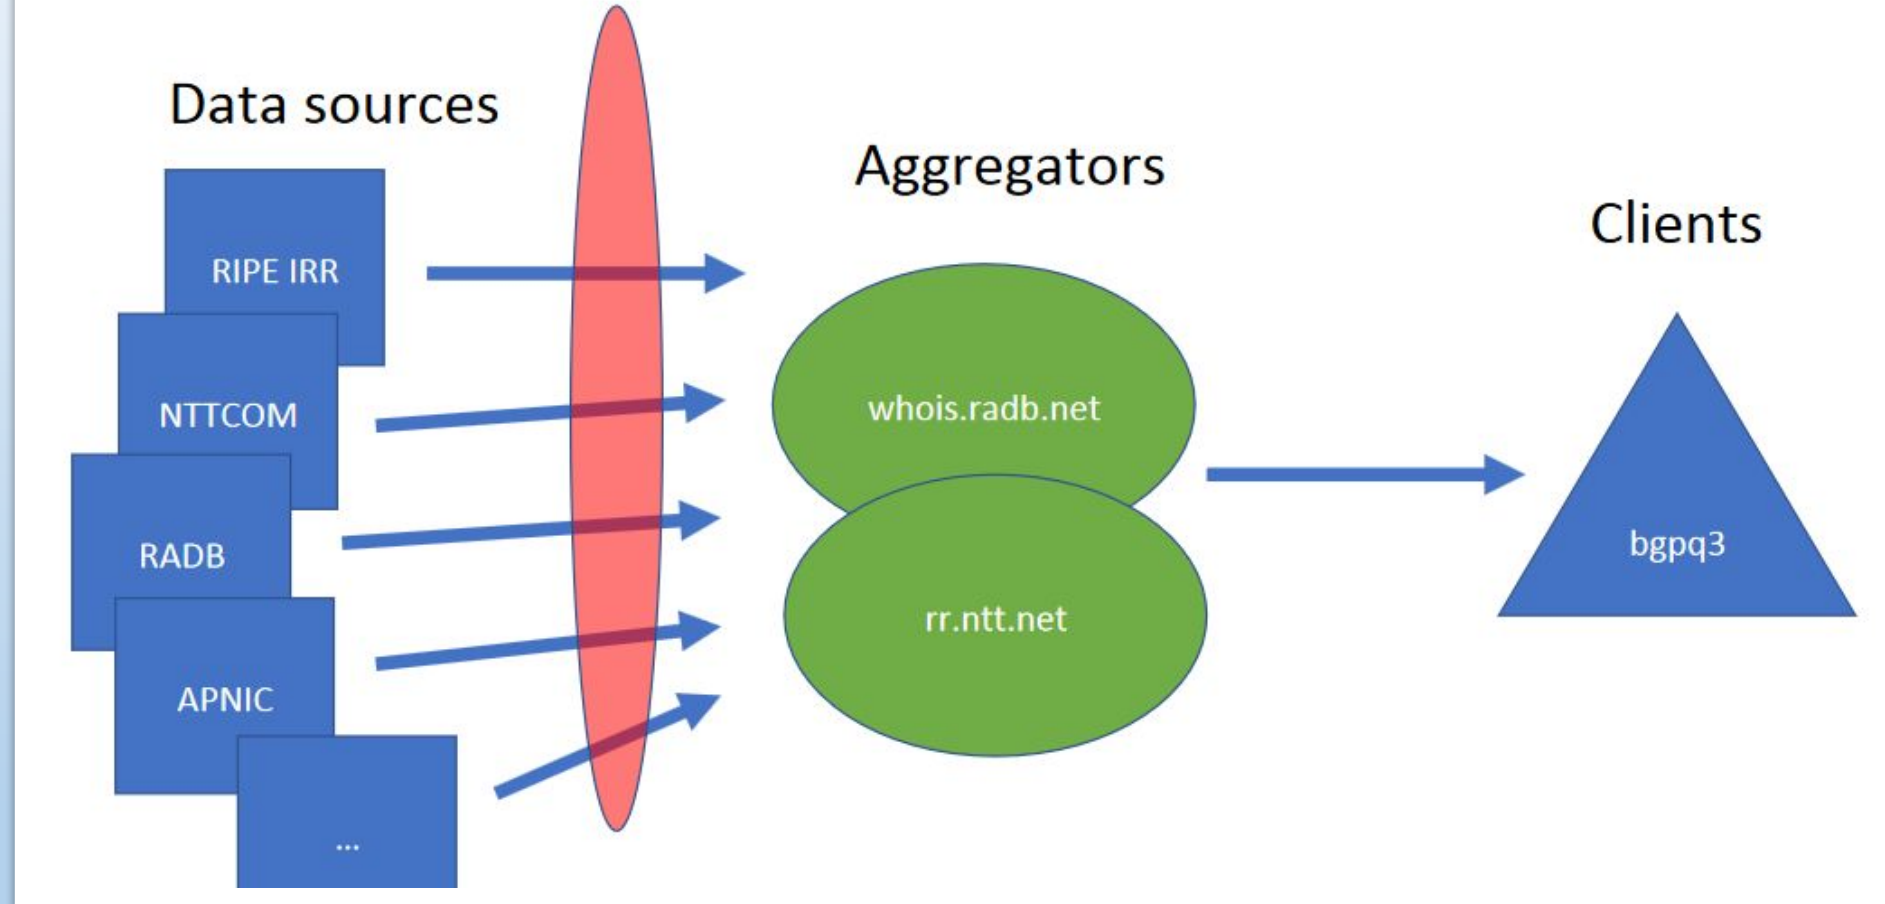
Three steps were taken:

- Cannot register non-RIPE-managed space any more
- All non-RIPE space moved to separate “RIPE-NONAUTH” database
- Route/route6 ASN authorization rules have been improved

More info: <https://www.ripe.net/manage-ips-and-asns/db/impact-analysis-for-nwi-5-implementation>

IRR – Ações para resolver os problemas

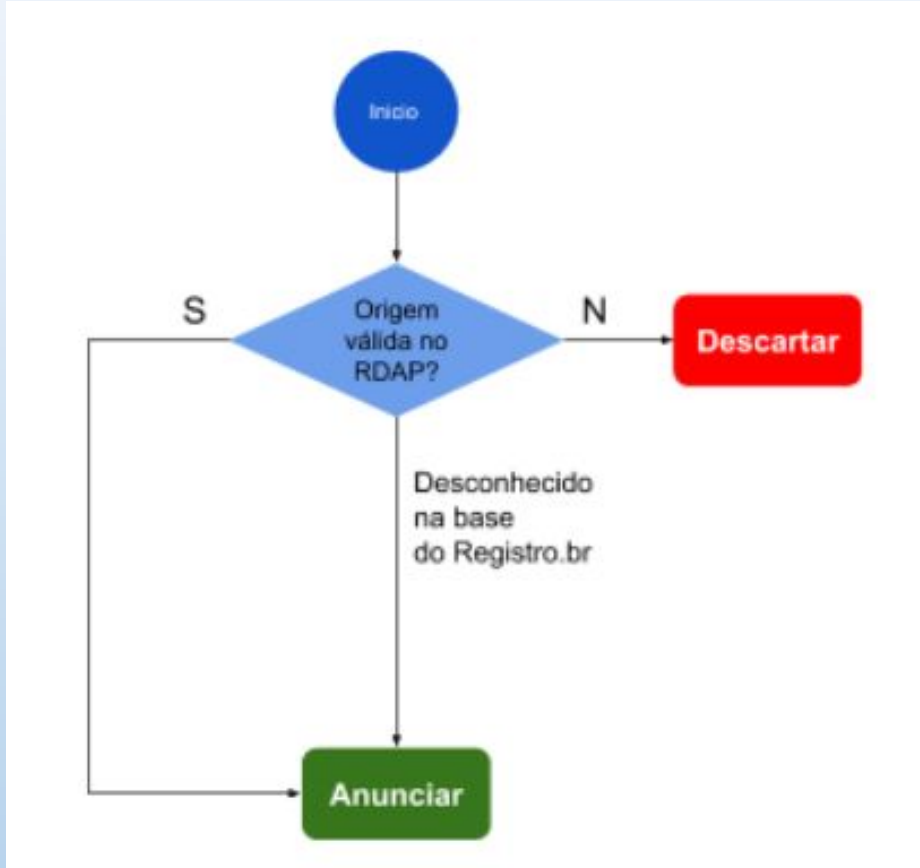
RPKI filter at the aggregators



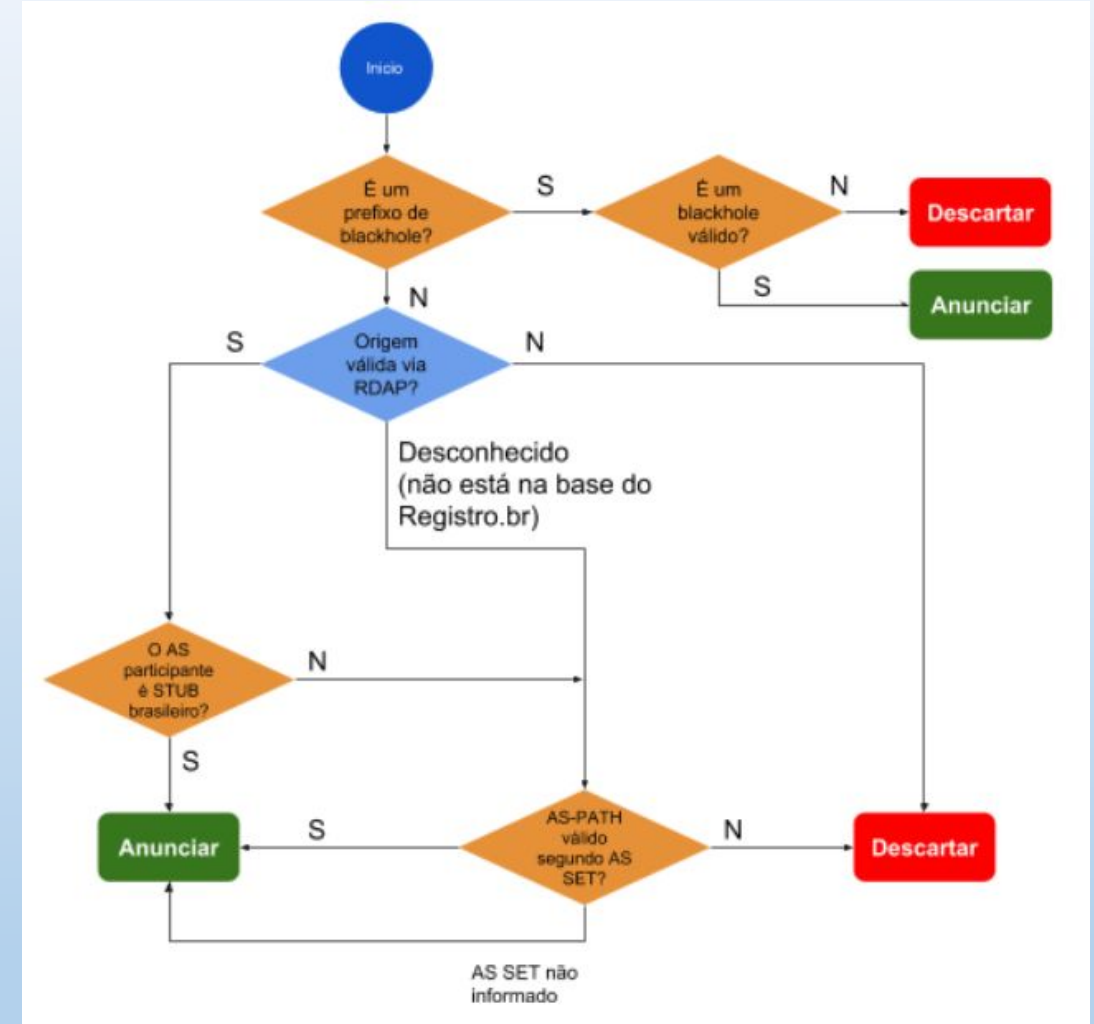
Job Snijders job@ntt.net – LACNIC 30 – Rosário – 26/09/2018

https://www.lacnic.net/innovaportal/file/3135/1/lacnic30_snijders_routing_security_roadmap.pdf

Filtros no IX.BR <http://ix.br/doc/solicitacao-comentarios-acoes-seguranca-ix-br-20180713.pdf>



os filtros serão implementados a **MÉDIO** prazo.

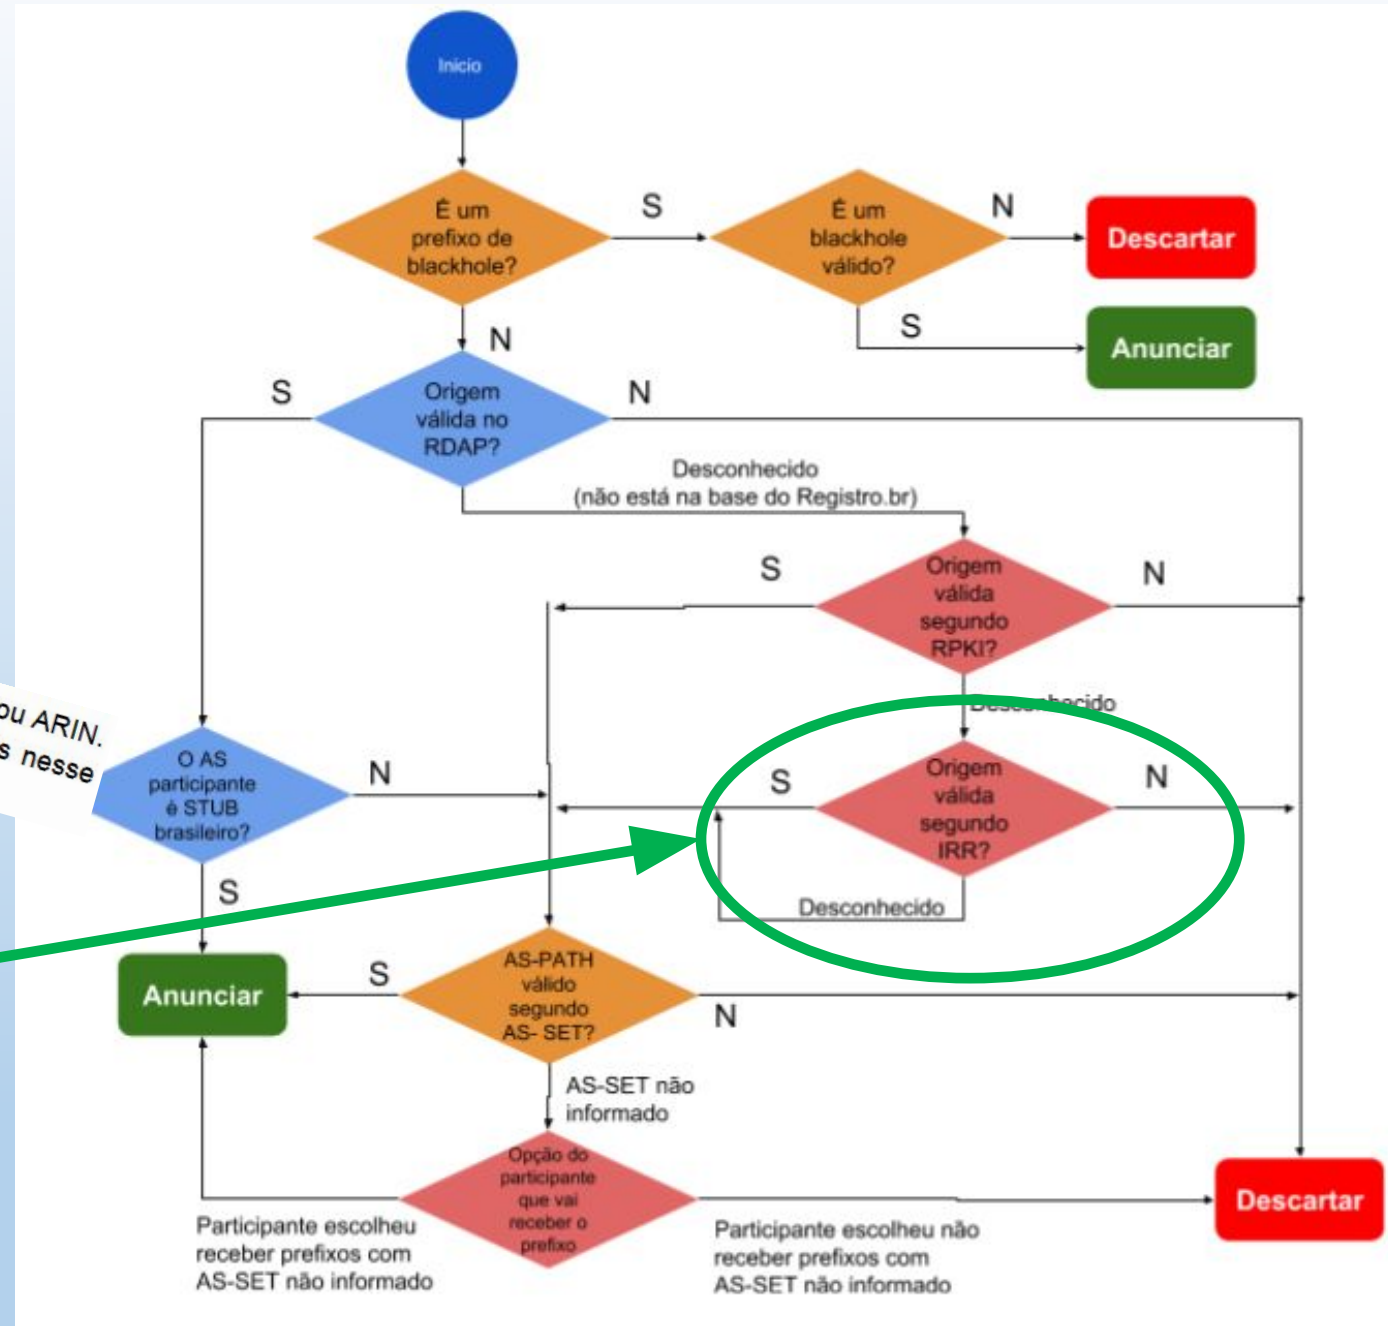


os filtros serão implementados a **LONGO** prazo.

Filtros no IX.BR



O diagrama de **LONGO** prazo não contempla ainda os filtros utilizando RDAP do LACNIC ou ARIN. São necessários estudos adicionais para determinar a exata forma como serão inseridos nesse fluxograma.




O que você tem a ver com isso? – Parte 2

- Se você é cliente de trânsito
 - Faça seu cadastro numa base IRR.
 - RADB
 - Excelente Back-End, excelente Front-End, ótimo atendimento.
 - IRR default para consultas do BGPq3 e do IRRToolset.
 - Não é “di grátis”.
 - TC (bgp.net.br)
 - Excelente Back-End, excelente Back-End, ótimo Back-End.
 - É “di grátis”.
 - Não faz proxied.
 - Cadastre sua Routing Policy no Site do Registro.

Oque você tem a ver com isso? – Parte 2

- Se você é fornecedor de trânsito
 - Defina e **publique** uma politica de roteamento
 - Tenha suporte ao básico de Communities
 - Aceite BlackHole, e faça o descarte corretamente.
 - Marque com communities as origens de onde aprendeu os prefixos que está repassando (Ex.: Trânsito Nacional, trânsito internacional, IX, CDN)
 - Utilize um mecanismo automático de criação de listas de prefixos.
 - Aplique com rigor filtros de eBGP-in, eBGP-out, iBGP-in, iBGP-out.
 - Exija que seu clientes tenham registros de IRR.
 - Ajude-os a fazer se necessário.



Perguntas?

“Você tem que ser o que você realmente é.

Pois se você não for quem você é, afinal quem é você?”

Doug Funnie