Núcleo de Informação e Coordenação do Ponto BR

Program for a Safer Internet

Actions in the IX.br

Introduction

IX.br is present in 31 locations in Brazil, through the installation and operation of Internet Traffic Exchange Points (IXs), being an integral part of the Internet network infrastructure of Brazil, where Autonomous Systems (ASs) can exchange traffic in nearby metropolitan regions.

Two types of services are offered to participants of traffic exchange: (i) **the Multilateral Peering Agreement** (MPA) and (ii) **Bilateral Peering Agreement** (BPA). MPA participants exchange traffic with each other: as a general rule, each AS exchanges traffic with all others. In the Bilateral Peering only two ASs participate, whether or not they use an exclusive layer 2 domain (a bilateral VLAN).

The Multilateral Peering Agreement (MPA), in practice, works with a shared VLAN for exchanging IPv4 (MPAv4) traffic and another for exchanging IPv6 (MPAv6) traffic. Each IX has two or more *route servers,* which are also used in the Multilateral Peering Agreement to centralize the receipt of route announcements from all participants of the traffic exchange, allowing, with a single BGP session, the route table is loaded and maintained. Establishing BGP sessions with *route servers* is a necessary condition for joining the MPA. Most participants in a IX participate in the multilateral peering, but not all. Even non-MPA attendees may be present on MPAv4 or MPAv6 VLANs for monitoring purposes, or for other purposes.

There are cases where the participant is present in the VLAN of MPAv4 or MPAv6 and does not close BGP session with the *route server*, but closes BGP sessions directly with the router of other participants with whom he wants to exchange traffic, using the IPs provided by IX.br . That is, bilateral peering agreements can use both the common-use VLAN (MPAv4 or MPAv6) and specific VLANs (bilateral VLANs).

In this scenario, we have in each city of IX.br:

- a **private environment**, formed by Bilateral Agreements with direct exchange of traffic through VLANs, either bilateral VLANs or MPAv4 and / or MPAv6 VLANs, and
- a **shared environment** formed by participants in the MPAv4 and / or MPAv6 VLANs and with BGP sessions with the route servers.

Within the **Safer Internet Program** that NIC.br is developing with the Internet community, **this document concerns actions that will be implemented in IX.br to increase the security of the Route Servers**, where the occurrence of problems with the shared route table can seriously affect the ASs participating or not in the IXs. **These are actions on the shared environment.** Occurrences in **the private environment are not in the scope of action of NIC.br and IX.br,** but all the recommendations made in the program by a safer Internet, such as the adoption of the

actions proposed in the MANRS initiative (Mutually Agreed Norms for Routing Security - https://www.manrs.org/), should be considered and applied by the managers of the ASs involved in private relations within the IX.br.

Internet security depends crucially on the participation of all ASs. The security of the environment of an IX (Internet Exchange Point) reflects the care that each participating network adopts internally. If all the networks adopt the best practices recommended in the configuration of their equipment, surely we will have a healthier environment in the IX. Setting up a network to prevent problems from spreading, that is, controlling what goes out of a network, is much simpler and more cost-effective than protecting network entry from everything outside. If everyone protects the outputs of their networks, there will be no problems entering the IX. This is the philosophy of MANRS, a global initiative supported by the Internet Society, which provides crucial recommendations for eliminating threats caused by the most common routing problems, and aims to:

- Raise awareness and encourage action with supporters' commitment.
- Promote the culture of collective responsibility for the resilience and security of the global routing system of the Internet.
- Demonstrate the ability of the industry to address resilience and security issues with a spirit of collective responsibility.
- Provide a framework for Internet service providers (ISPs) to better understand and help solve problems related to the resiliency and security of Internet routing.

In IXs there is no way to deal with all possible problems caused by networks configured without proper protections, due to technical restrictions on the network equipment. However, **we can act on the security of the *route servers***, with measures that reduce the possibility of the occurrence of prefix hijack or route leaks that have caused both prejudice and concern to Brazilian Internet users.

This document describes the actions already in use and others that will be implemented soon. This document was the result of a broad process of consultation and interaction with the community, and takes into account the *feedback* received.

Enhancing Route Servers Security

To increase the reliability of the local MPA route table of an IX, validations must be performed on the route servers, testing the prefixes received from participants in a number of different ways to mitigate any errors, malicious actions, or in the configuration of the routers of the peering participants. From these validations the prefixes are identified with BGP communities, and finally filtered, or not, according to specific criteria.

The BGP Announcement Validation Process

The process as a whole is divided into different steps:

1. The **route server receives the prefixes**.

- Each participant has defined a maximum number of prefixes that can be announced to route servers. If this number is exceeded the BGP session is knocked over for protection for 10 (ten) minutes. If in the return the maximum number of prefixes is exceeded, the session will be felled indefinitely and the other steps of the process do not apply. To increase the maximum number of announced prefixes, the participant must open a specific support ticket on the participant's portal.

2. Ads **are validated** according to various criteria:
   - Verification of not allowed prefixes or AS PATHs not allowed.
   - Verification of the origin of the prefix.
   - Verification of the policy informed by the participant to IX.br.

   For these validations, different types of tests can be done:
   - Verification of the AS type or prefix: stub or non-stub, from the LACNIC region or from other regions.
   - Ad comparison with static tables of prefixes or non-allowed ASNs (bogons, IX.br prefixes, etc).
   - Queries to external databases such as RDAP, RPKI, IRRs, or others.
   - Consult the IX.br database (policy defined by the participants).

3. **Each prefix is marked with communities** as the result of the validations.
   - Note that for each type of validation a different BGP community to be specified in the future will be used, so it will be clear to determine for which reason a particular prefix is considered invalid.

4. The result of the validations is visible in the **Looking Glass Web**, including prefixes that will not survive the filters later in the process.
   - In Looking Glass Web it will be possible to visually identify in a simple and descriptive way if a prefix will be exported to the other participants and, in case the prefix is filtered, what validation or validations are responsible for the discard of the advertisement.
   - Note that, as is the case today, participants can mark their ads with certain BGP communities in order to control the behavior of Route Servers, allowing the export of that particular ad only to a specific group of other participants. In this case a prefix that appears on the route server may not be exported to one or more participants. These communities used by the participants themselves will remain visible the same way they are today in the Looking Glass Web, anonymizing the ASNs involved in the filters, unless the viewer is linked to the AS that made the advertisement to the Route Servers.

5. The **prefixes can be filtered or not, according to the marked communities** and different criteria.
   - Some criteria are fixed (eg, bogons are always discarded).
   - Some criteria can be chosen by the participants (eg, filtering or not prefixes in which the validation is inconclusive, that is, they do not appear in one or more databases).

6. The **prefixes that survived the filters are exported** to the other participants.

The following diagram illustrates the different steps in the BGP Ad Validation Process:

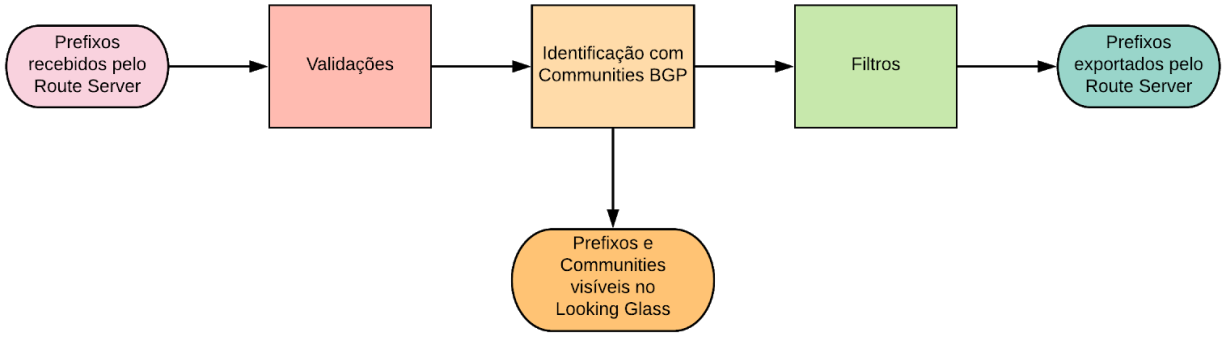

Validations, identification with BGP communities, and filters are done by Route Servers in real time as the prefixes are received. However, this real-time process is based only on the static configuration of the Route Servers. The IX.br and external databases are not consulted in real time.

Information is collected periodically, or at the time of activation. Information such as IX.br participants, prefixes announced, AS-PATHs, data reported in the IX.br portal, data from external databases such as RDAP, IRR, etc. Based on this set of information, periodic prefix lists and other elements that are part of Route Servers configurations are constructed periodically.

It is important to note then that the change of an information on an external basis, such as the designation of a prefix in the Registro.br, or the inclusion of a transit client in AS-SET published in an IRR, or even the change in information provided via portal IX.br, will not have immediate effect on the process of validating BGP ads on Route Servers. The effects will occur in up to 24 hours or less periodicity to be defined in the future by IX.br.
An important exception is the RPKI base. Route Servers are able to query a local mirror of the RPKI in real time, so that the effects of modifications will occur with slower delays.

How will be the implementation of the BGP Advertising Validation Process in IX.br ?

Below we present a series of validation criteria and filters to be implemented in the Route Servers of IX.br. For each we classify as the expected time of implementation as:

- IN USE;
- SHORT (45 days);
- MEDIUM (120 days) or
- LONG (12 to 18 months).

The **time for implementation** changes according to several criteria:

- **Technical feasibility and ease of implementation:** for example, validation of the prefixes against a list of bogons is already IN USE, it is an extremely simple and very effective validation;
- **Technical maturity or degree of adoption of the technology**: for example, validation of origin for stubs in the LACNIC area in the RDAP requires the creation of specific software, so the term for implementation is AVERAGE, while the validation of origin by IRR still requires a better maturity of issues related to security and availability for use without cost, so its implementation period is LONG.

Limit the number of prefixes advertised

Currently, the standard policy IN USE in IX.br is to accept up to 100 IPv4 prefixes and up to 100 IPv6 prefixes in their sessions with the route servers. During the activation process, or later via a ticket, that number can be increased if necessary.
If this number is exceeded the BGP session is knocked over for 10 (ten) minutes and reactivated. If the maximum number of prefixes is exceeded, the session will be knocked over again and the participant must open a ticket so that it can be manually reset with a new limit for the number of prefixes. This filter has the main objective of avoiding ads, by configuration error, of the complete BGP table, or of prefixes learned from the route servers themselves.

Disaggregation of Internet prefixes should be used sparingly. The disaggregation can be done in a studied and planned way, for traffic engineering, for example. This relatively large default maximum quantity of prefixes accepted by the IX.br route servers is by no means intended to encourage prefix unbundling. This number currently caters to both small participants, Autonomous Systems stubs, who advertise only their own prefixes, as well as most of the larger participants, who provide traffic to other Autonomous Systems and advertise a relatively large number of prefixes for this reason.

Validation

1. **Validation prefixes or ASNs not allowed in IX.br network**

    a. **prefixes Size**

Prefixes with masks between / 8 and / 24 inclusive for IPv4, or between / 3 and / 48 inclusive for IPv6 will be accepted. IPv4 prefixes with masks between / 25 (inclusive) and / 31 (inclusive) will be marked as invalid. IPv6 prefixes with masks between / 49 (inclusive) and / 127 (inclusive) will be marked as invalid.

IPv4 prefixes with masks / 32 and IPv6 prefixes with masks / 128 will be considered as prefixes for blackhole, as long as they are accompanied by the appropriate community, and are subsequently reviewed by other criteria. They will be marked as valid by the size criterion. If they are not marked with the appropriate community, they will be marked as invalid.

Summing up, for clarity, the following table with prefix classification ccording to size analysis:

Address Type	**Valid**	**Invalid**
IPv6	between / 3 and /48 /128 (if marked with community of blackhole)	</3 >/ 48 and < /128 /128 (if not marked with blackhole community)
IPv4	Between / 8 and /24 /32 (if marked with blackhole community)	</ 8 > /24 and </32 /32 (if not marked with blackhole community)

There are currently filters IN USE on the IX.br network, discarding ads with prefixes greater than /48 IPv6 and greater than /24 IPv4. In the new implementation, invalid ads will be tagged with appropriate communities visible on the Looking Glass Web for AS knowledge and tracking. Later they will be filtered. The new implementation will also allow in some cases advertisements for /64 IPv6, /128 IPv6 or /32 IPv4 as blackholes. This change will also allow the development of educational actions with the ASs, to prevent this type of situation from occurring with other peers and the generation of an indicator (KPÌ) on the subject.

As long as the validation and blackhole filters are not implemented, the /64 or /128 IPv6, and /24 IPv4 prefixes will be considered invalid by the size criteria.

Time to implement / status: SHORT.

b. Bogons or Undue Prefixes

Received announcements containing address block using reserved address space, as listed below, are marked with specific communities and then discarded.

IPv4 prefixes not allowed:

```
0.0.0.0/8 prefixlen> = 8              # 'this' network [RFC1122]
10.0.0.0/8 prefixlen> = 8             # private space [RFC1918]
100.64.0.0/10 prefixlen> = 10         # CGN Shared [ RFC1998]
127.0.0.0/8 prefixlen> = 8            # localhost [RFC1122]
169.254.0.0/16 prefixlen> = 16        # local link [RFC3927]
172.16.0.0/12 prefixlen> = 12         # private space [RFC1918]
192.0.0.0/ 24 prefixlen> = 24         # IETF Protocol Assignments
192.0.0.0/29 prefixlen> = 29          # DS-Lite [RFC6333]
192.0.2.0/24 prefixlen> = 24          # TEST-NET-1 [RFC5737]
192.88.99.0/24 prefixlen> = 24        # 6to4 Relay Anycast [RFC3068]
192.168.0.0/16 prefixlen> = 16        # private space [RFC1918]
198.18.0.0/15 prefixlen> = 15         # benchmarking [RFC2544]
198.51.100.0/24 prefixlen> = 24       # TEST- NET-2 [RFC5737]
203.0.113.0/24 prefixlen> = 24        # TEST-NET-3 [RFC5737]
224.0.0.0/4 prefixlen> = 4            # multicast
240.0.0.0/4 prefixlen> = 4            # reserved for future use
255,255 .255.255 / 32 prefixlen = 32 # Limited Broadcast [RFC0919]
```

IPv6 prefixes that are allowed (all others are allowed identified as not allowed):

```
2001: 0200 :: / 23 prefixlen accepted / 23 a / 48
2001: 0400
:: / 23 prefixlen accepted / 23 a / 48
2001: 23/48 / /
/ / /
/ / /
/ / / / / / / / / / / / 23 a / 48
2001: 1200 :: / 23 prefixlen accepted / 23 a / 48
2001: 1400 :: / 23 prefixlen accepted / 23 a / 48
2001
: : / 23 prefixlen accepted / 23 to / 48
2001: 1A00 :: / 23 prefixlen accepted / 23 to / 48
2001: 1C00 :: / 22 prefixlen accepted / 22/48
2001: 2000 :: / 20 prefixlen accepted / 20 / 48
2001: 3000 :: / 21 Prefixes accepted / 21 a / 48
2001: 3800 :: / 22 Prefixes accepted / 22 a / 48
2001: 4000 :: / 23 Prefixes accepted / 23 a / 48
2001: 23 prefixlen accepted / 23 to / 48
2001: 4400 :: / 23 prefixlen accepted / 23 to / 48
2001: 4600 :: / 23 prefixlen accepted / 23 to / 48
2001: 4800 :: / 23 prefixlen accepted / 23 to / 48
2001: 4a00 :: / 23 prefixlen accepted / 23 a / 48
```

2001: 4c00 :: / 23 pr efixlen accepted / 23 and / 48
2001: 5000 :: / 20 prefixlen accepted / 20/48
2001: 8000 :: / 19 prefixlen accepted / 19 and / 48
2001: a000 :: / 20 prefixlen accepted / 48
20/2001: b000 :: / 20 prefixlen accepted / 20 a / 48
2002: 0000 :: / 16 prefixlen accepted / 16 a / 48
2003: 0000 :: / 18 prefixlen accepted / 18 a / 48
2400: 0000 :: / 12 / 12 a / 48
2600: 0000 :: / 12 prefixlen accepted / 12 a / 48
2610: 0000 :: / 23 prefixlen accepted / 23 a / 48
2620: 0000 :: / 23 prefixlen accepted / 23 a / 48
2800: 0000::: / 12 prefixlen accepted / 12 a / 48
2a00: 0000 :: / 12 prefixlen accepted / 12 a / 48
2c00: 0000 :: / 12 prefixlen accepted / 12 a /

In addition to these, they are explicitly not allowed:

2001::/23 prefixlen >= 23	# IETF Prot Assignments [RFC2928]
2001::/32 prefixlen >= 32	# TEREDO [RFC4380]
2002::/16 prefixlen >= 32	# 6to4 [RFC3056]
2001:2::/48 prefixlen >= 48	# BMWG [RFC5180]
2001:10::/28 prefixlen >= 28	# ORCHID [RFC4843]
2001:20::/28 prefixlen >= 28	# ORCHIDv2 [RFC7343]
2001:db8::/32 prefixlen >= 32	# document range [RFC3849]

Currently there are filters IN USE in the IX.br network, discarding ads with bogons prefixes. In the new implementation, invalid ads will be tagged with appropriate communities visible on the Looking Glass Web for AS knowledge and tracking. Later they will be filtered. This change will also allow the development of educational actions with the ASs, to prevent this type of situation from occurring with other peers and the generation of an indicator (KPÌ) on the subject.

Time to implement / status: SHORT.

c. Prefixes used by IX.br

The address space used by IX.br in order to address the routers participating in the local traffic exchange, even if it uses public IPv4 and IPv6 addresses, MUST NOT be routed (it should not be announced by any participant, and is not announced by IX.br or by NIC.br).

Public IPv4 and global IPv6 addresses are used to facilitate the investigation of connectivity problem and/or routing problems (troubleshooting).

Currently, there are filters already IN USE in the IX.br network for this case, discarding the prefixes. In the new filter implementation, invalid prefix ads used by IX.br will be marked with a community reporting the irregularity, exporting the ad only to the Looking Glass Web for knowledge and tracking of the AS.

Time to implement / status: SHORT.

It is important to remember that IX.br's Technical Requirements policy already defines that: *"the address space of the network of each IX.br location, that is, the addressing used on the ports of the IX-connected routers, should not be announced to other networks. It is recommended that these addresses also not be advertised internally in the participants' network, which implies the use of next-hop-self for the internal advertisement of routes learned through IX.br*.

The dissemination of the IX.br network in the internal routing of a participant or to the Internet implies a serious security risk, allowing, for example, DDoS attacks directed to the routers of the other participants and Route Servers. If the announcement of IX.br prefixes is detected by a participant, it is notified and, if the situation persists, it may be disconnected from the IX.br network.

d. ASNs bogons or improper

Ads containing ASNs reserved (bogons) anywhere in AS-PATH will be marked with a specific community and subsequently rejected.

ASNs not allowed:

0 - RFC 7607
23456 - RFC 6793 AS_TRANS
64496 to 64511 - RFC 5398 and documentation / example ASNs
64512 to 65534 - RFC 6996 Private ASNs
65535 - RFC 7300 Last 16 bit ASN
65536 to 65551 - RFC 5398 and documentation / example ASNs
65552 to 131071 - IANA reserved ASNs
4200000000 to 4294967294 - RFC 6996 Private ASNs
4294967295 - RFC 7300 Last 32 bit ASN

Time to Deploy / Status: SHORT.

e. ASNs of traffic free networks in AS-PATH

Validation and identification with the appropriate BGP community as invalid, of the ads containing in the AS-PATH, after the participant's ASN, the ASN of networks known as traffic-free, typically the major Tier 1:

174 - Cogent
209 - Centurylink
286 - KPN
701 - Verizon
702 - Verizon
703 - Verizon
1239 - Sprint
1299 - Telia
2828 - XO
2914 - NTT
3257 - GTT Communications
3320 - Deutsche Telekom
3356 - Level 3
3491 - PCCW Global
3549 - Level 3
3561 - Centurylink
4134 - China Telecom
4323 - TWTC
4436 - GTT
5511 - Orange
6453 - Tata Communications
6461 - Zayo
6762 - Telecom Italia Sparkle
6830 - UPC
6939 - HE
7018 - AT & T
12956 - TIWS

The presence of these ASNs in AS- PATH is an indication of misconfiguration on the traffic exchange participant router. These networks typically do not hire traffic from other networks. The ASN presence of one of them in AS-PATH indicates that the participant, or a participant's transit customer, is 'providing traffic' to the participant, indicating a configuration error or other type of problem.

Several comments were made regarding this type of validation, during the consultation phase of the previous version of this document, summarized below:

- The hypothesis that the listed networks are actually "traffic free" may not be valid. Cases have been reported in which some of them hire partial traffic from local (Brazilian) operators to improve their connectivity. An in-depth study considering the route table of the various locations will be carried out to determine the current situation.
- It was suggested to include free transit networks in Brazil, which would only actually contract international traffic, such as RNP, Embratel Tim, Telefônica, Oi, Claro, etc. At first the proposal seems valid, which will be confirmed by

analyzing the route table of the various localities, as well as directly with the companies.

- Inclusion of service provider / content provider ASs such as Google, Netflix, Amazon, Microsoft, etc. in the list of ASNs. In this case we can request a formal positioning of these ASs regarding the filtering policy to be adopted in the IX.br locations.

Note: on 03/28/2018 in IX.br SP were found 90 ads containing the ASNs listed above, with a daily traffic of 36.8 TB and an average of 3.5 Gbps.

The validation and identification of prefixes in this situation will be implemented on a SHORT term. However, in-depth studies and consultations with the companies involved will be done BEFORE USING SUCH INFORMATION FOR FILTERS. If necessary, exception cases will be established for participants of IX.br who actually offer IP traffic for such networks, so that prefixes in those conditions announced by them are not marked.

The ASNs on this list will be periodically reviewed, based on the analysis of the global route table of the Internet.

Time to implement / status: SHORT.

2. Source Validation

Source validation aims to verify that the AS that originated the ad, ie the rightmost AS in the AS-PATH, really has the right to advertise that prefix.
In the following ad samples, what you want to check with source validation is whether AS 64511 has the 192.0.2.0/24 prefix assigned to it; if AS 65537 has the prefix 198.51.100.0/24 assigned to it, and if AS 65536 has the prefix 203.0.113.0/24 assigned to it.

```
Prefix              AS-PATH
192.0.2.0/24        0 64500 64499 64511 i
198.51.100.0/24     0 64500 64500 64500 65540 65536 65537 i
203.0.113.0/24      0 65536 i
```

This action promotes the use of external databases for ad validation .

As previously announced in the document dealing with the "Proposal to change route servers and treatment policies of communities BGP in IX.br", the ads received by the route servers will be marked as invalid, invalid or unknown according to searches performed in 3 types of services / databases: RDAP, IRR and RPKI.

a. RDAP

Registration Data Access Protocol (RDAP) can be viewed as an API for accessing a Whois database, such as the Registro.br

The Registro.br and LACNIC do a direct correlation between the prefixes and the ASNs assigned to the same organization. That is, by referring to the prefixes it is possible to relate them to an ASN. And by referring to an ASN, you can relate it to your prefixes.

In ARIN there is no direct correlation for all assignments. However both ASNs and prefixes are assigned to an organization. By consulting an ASN it is possible to reach the organization to which it was assigned, and then to the prefixes assigned to that organization. Similarly, by having a prefix, it is possible to discover the ASNs or ASNs assigned to the same organization.

In other RIRs, as far as the authors of this document are aware, it is not possible to establish, through public database queries, a direct relationship between the prefixes and the ASNs assigned to the same organization.

In Brazil, the updating of the database consulted via RDAP is done by Registro.br in the process of assigning the numbering resources, not requiring any type of registration by the ASN. On the other hand, it is essential that delegations or transfer of blocks between ASNs are duly informed and registered in the Registro.br, without which there may be rejection of advertisements.

The validation using Registro.br RDAP will be implemented on a SHORT term. There will be an appropriate BGP community to:

- indicate whether or not the prefix is in the base, if it does not appear in the base it will be marked as unknown
- indicate if the ASN that originates the prefix, ie the rightmost ASN in the AS-PATH, is either not the same ASN to which the prefix is assigned in the Registro.br database

Instead of direct queries to RDAP, in the case of Registro.br, one can also choose to consult the file available at the following address, updated daily, and with equivalent information:

- ftp://ftp.registro.br/pub/numeracao/origin/nicbr-asn-blk-latest.txt

The possibility of validation of origin in LACNIC and ARIN RDAP will also be studied in greater depth. Practical questions such as speed of consultations and limit of queries to the respective bases need to be evaluated, among others. A time is foreseen for implementation is LONG.

Another source of data that can be used is the files made available by LACNIC at ftp://ftp.lacnic.net/pub/stats/ containing the assignments made by other RIRs:

AFRINIC, APNIC, ARIN, RIPE/NCC and LACNIC itself. This information is made available daily. Time for implementation MEDIUM.

Note that source validation via RDAP or WHOIS will not be done immediately by announcing a new prefix for Route Servers. These validations are done through a static configuration (prefix lists and others) in Route Servers. These settings are periodically generated based on queries made to RDAP bases based on the current list of prefixes advertised to Route Servers and the list of participating Brazilian stub ASs.

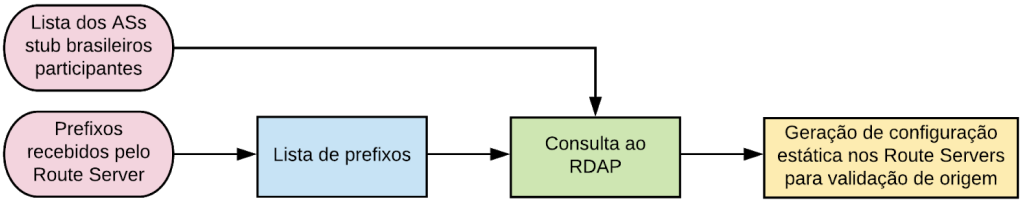

It is worth noting that for the STUB Autonomous Systems in Brazil, by definition, the AS that originates the prefix is the same AS of the participant of IX.br. That is, by definition, there is only one AS in the AS-PATH. The validation of origin made in the RDAP of Registro.br also applies to the Autonomous Systems STUB in Brazil.

b. IRR

The IRRs are databases that store routing policies described in a language called Routing Policy Specification Language (RPSL). These databases are distributed and operated by various organizations such as RIRs (Regional Internet Registry), telecom companies, etc. Some of these bases operate as paid services, such as the RADB, others are free, such as the TC and bases operated by RIRs.

In an IRR database it is possible to inform:

- wich AS originate a certain prefix (this is precisely the information to be used in the validation of the origin of the prefixes):
    - For example, when querying the prefix 200.160.0.0/24 in the IRR The following information, which shows that the originating AS is the AS22548:

```
route: 200.160.0.0/20
descr: Registro.BR Network
origin: AS22548
(...)
```

IRRs also store other types of information:

- what are the transit ASNs for a specific AS (import and export policies):

- For example, when you query AS22548 in the IRR databases you can find the following information, showing that your transits are AS 3549, 12989, 16735 and 52320:

```
aut-num: AS22548
(...)
import: from AS3549 accept ANY
import: from AS12989 accept ANY
import: from AS16735 accept ANY
import: from AS52320 accept ANY
mp-import: from AS3549 accept ANY
mp-import: from AS12989 accept ANY
mp-import: from AS16735 accept ANY
mp-import: from AS52320 accept ANY
export: to AS3549 announce AS22548 AND {200.160.0.0/20}
export: to AS12989 announce AS22548 AND {200.160.0.0/20}
export: to AS16735 announce AS22548 AND {200.160.0.0/20}
export: to AS52320 announce AS22548 AND {200.160.0.0/20} }
MP-20export: to AS3549 announce AS22548 AND {2001: 12ff :: /
32}-exportmp: to announce AS12989 AS22548 AND {2001: 12ff :: /
32}-exportmp: to announce AS16735 AS22548 AND {2001: 12ff::/32}
mp-export: to AS52320 announce AS22548 AND {2001:12ff::/32}
(...)
```

- which are the ASs advertised via BGP (traffic clients, usually) of a particular AS (AS-SET):
  - For example, when consulting AS22548 in PeeringDB you can verify that your AS-SET is called AS-NIC-BR

```
Organization: NIC.br
Also known as: Information and Coordination Center of Ponto BR
Company Website: http://nic.br
Primary ASN: 22548
IRR Registration: AS-NIC-BR
(...)
```

  - When consulting the AS-NIC-BR SET in the bases of IRR, it can be observed that the ASs announced in BGP by AS22548 are ASs 10906, 11284, 11644, etc:

```
as-set: AS-NIC-BR
descr: Nucleo de Informacao e coordination Point BR
members: AS10906
members: AS11284
members: AS11644
members: AS11752
members: AS12136
members: AS14026
members: AS14650
members: AS20121
members: AS22548
members: AS26162
members: AS53035
(...)
```

There are services that have copies (mirrors) of existing databases in order to provide a more complete view of existing IRR databases.

The information in the IRR databases is not totally reliable as regards its use for source validation, since it is possible for a third party to insert in a given database information about prefixes or ASNs that it does not manage. This reinforces the importance that AS administrators have their information published in at least one IRR. Since in this case, if someone improperly inserts invalid information about a prefix or ASN, they will at least be conflicting in a query with the correct information entered by the resource administrator. If there is no information registered by the real administrator of the resource, all that will be obtained in a query will be the information incorrectly entered by third parties, which will probably be interpreted as correct.

For IRR source validation purposes, if there is at least one IRR record indicating that the source is valid, we will consider it valid even if there are other conflicting records.

The identification of the ads as valid or not, according to the origin information found in the IRR bases, will be implemented in the LONG term. However, in-depth studies will be conducted BEFORE USING SUCH INFORMATION FOR FILTERS.

A community BGP, or appropriate BGP communities, will be used to indicate if the AS-PATH prefix has an ASN source, the right-most AS_SEQUENCE, valid according to IRR bases, or if it is invalid or unknown in themselves.

Note that source validation via IRR will not be done immediately by announcing a new prefix for Route Servers. These validations are done through a static configuration (prefix lists and others) in Route Servers. These settings are periodically generated based on queries made to IRR databases based on the current list of prefixes advertised to Route Servers.

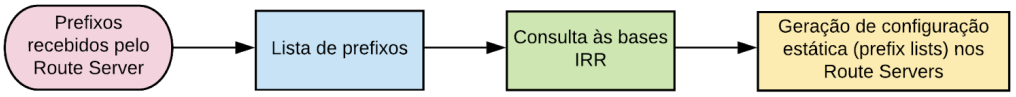

Time to implement / status: AVERAGE.

a. RPKI

RPKI (Resource Public Key Infrastructure) is a service implemented by RIRs and NIRs, through which the organization that has a prefix informs which Autonomous System, or Autonomous Systems, is authorized to announce it.

The system uses a public key certificate infrastructure to secure Internet routing through the generation of attests called Route Origination Authorization (ROAs).

The RPKI is very reliable for the validation of origin, however it is a service not yet available in Brazil and with little adhesion of the Autonomous Systems globally.

A community BGP, or appropriate BGP communities, will be used to indicate whether the prefix has a valid source ASN according to the RPKI in its AS-PATH, or whether it is invalid or unknown.

Time to implement: AVERAGE.

2. Validation of routing policy of IX.br participants

The validation of the participant's routing policy has the purpose of verifying that the routing policy informed to IX.br is consistent with the advertisements for the Route Servers.

Routing policy in this context means:

- information on whether the participant is a Brazilian AS stub or not;
- if not, the complete list of ASs announced by the participant to IX.br's Route Servers (which normally consists of participant's AS clients).

That is, the participant informs IX.br by means independent of the BGP which announcements he intends to make, and IX.br verifies that the announcements effectively made via BGP are in accordance with this information.

Each prefix received by the Route Server will be identified with an appropriate community, indicating whether it is valid or not, based on the participant's information.

a. Brazilian STUBs Autonomous Systems

ASNs connected to IX.br can be classified into two categories: stub or transit:

- An ASN stub is one that is directly connected to IX.br, announcing only its own prefixes, and there is no other ASN in AS-PATH .
- Already an ASN traffic is one that announces prefixes of other ASNs, in addition to yours, with AS-PATH being able to contain multiple ASNs.

During the quarantine stage of the activation process of the participant in IX.br, the ads received will be analyzed and the classification of the ASN will be presented as stub or transit. The initial classification may be changed at any time through the portal of the IX.br participant.

When the process is implemented, for participants already connected we will perform the analysis and classification based on the table of routes in force.

The AS stub that eventually gives transit to other ASs, must specify in the portal of the IX.br participant that its classification should change to transit.

For an AS stub, the prefix validation, made via RDAP, is in practice equivalent to validating the routing policy, since a stub by definition only announces prefixes assigned to the AS itself.

The Brazilian AS stub will be identified as such through an appropriate community or communities.

Time to implement / status: SHORT.

b. Prefixes and ASNs informed to IX.br

For the transit ASs participating in IX.br it is impossible to infer the prefixes that will be announced, besides the prefixes assigned to it. This information must come from the participant himself.

The IX.br will ask the participating traffic attendants to:

1. **Register an AS-SET on an IRR base**, which includes ALL the ASs that they wish to advertise to the Route Servers.

   For example, the AS-SET of AS22548, was registered in the IRR RADB, with the following information:

```
as-set: AS-NIC-BR
descr: Information Center and Coordination of the Ponto BR
members: AS10906
members: AS11284
members: AS11644
members : AS11752
members: AS12136
members: AS14026
members: AS14650
members: AS20121
members: AS22548
members: AS26162
members: AS53035
(...)
```

2. **Inform in the PeeringDB IRR Registration field your AS-SET**, or enter the name of the PeeringDB its AS-SET in the format specified by RFC 4012 and preferably the IRR base on which the object was registered.

   PeeringDB will be the source used to obtain the name of the AS-SET used, so it is fundamental that this register is made.

For example, in PeeringDB find the following information about the AS SET AS22548:

```
Organization: NIC.br
Also known as: Nucleo Information and Coordination Point BR
Company Website: http://nic.br
primary ASN: 22548
Registration IRR: AS-NIC-BR
(...)
```

the PeeringDB is a service that facilitates the exchange of information related to peering. Specifically, it is a database of peering networks, where each AS can inform if it does peering, in which IXs or datacenters, what is the peering policy, etc. Its use is free of charge.

With this information IX.br can validate the AS-PATHs of the ads, observing the first ASN immediately to the right of the participant's ASN.

For example, let's consider the AS64500 participant. It reported an AS-SET consisting of the ASNs: 64499, 65540, 65550. The ads on the Route Server are as follows:

```
Prefix                AS-PATH
192.0.2.0/24          0 64500 64499 64511 i
198.51.100.0/24       0 64500 64500 64500 65540 65536 65537 i
203.0.113.0/24        0 64500 65536 i
```

In this example, the first and second ads are valid, but the third is not, because AS65536 is not in AS-SET.

The advertisement will be identified with the appropriate community, or communities, to identify whether or not it is valid from the point of view of the informed routing policy, the community or communities will also identify whether or not the routing policy was informed.

Time to implement / status: AVERAGE.

3. Prefix validation for BLACK HOLE for Brazilian STUB Autonomous Systems

Announcements received with mask /128 for IPv6 or /32 prefixes for IPv4 prefixes will be identified through a community, or appropriate communities, as blackhole prefixes. They will be marked as valid if:

1. they are Brazilian ASs stub and also
2. They are valid according to RDAP (Registro.br database)

Will be marked as invalid if they are not Brazilian stub ASs or not valid using RDAP.

Time to implement / status: AVERAGE.

Summary of validations

The following table summarizes the types of validation carried out, their implementation deadlines and possible results.

Type	Subtype	Time forimplementation validation	Time for filter implementation	Possible results
Validation of prefixes or ASNs not allowed in the IX network.	Size of prefixes	IN USE	D0	Valid or not
	Prefixes bogons or improper	IN USE	D0	Valid or not
	Prefixes used by IX	IN USE	D0	Valid or not
	ASNs bogons or unduly	SHORT	D + 30 days	Valid or not
	ASNs of traffic free networks	SHORT	D + 30 days	Valid or not
Validation of origin	Prefix is at the base of the Registry.br	SHORT	D + 30 days	Yes or no
	Valid source according to RDAP in the Registry.br	SHORT	D + 30 days	Valid or not Valid
	source according to RDAP / WHOIS in LACNIC and ARIN	LONG	D + 30 days	Valid or not Valid
	source according IRR	LONG	D + 30 days	Valid, invalid or unknown
	Valid source according to RPKI	MEDIUM	D + 30 days	Valid, not valid or unknown
Routing policy validation	ASN is Brazilian stub	SHORT	D0	Yes or no
	AS-PATH validation according to AS SET of participant	LONG	D + 30 days	Valid, not valid, or AS SET not informed
Validation of prefixes for BLACK HOLE	Systems Autonomous Brazilian STUB	AVERAGE	D + 30 days	Valid or not

Filters

Based on the validations made, filters will be implemented before the ads are exported for each participant.

The diagram below outlines the operation of filters based on the validation of prefixes or non-allowed ASNs in the IX.br network:

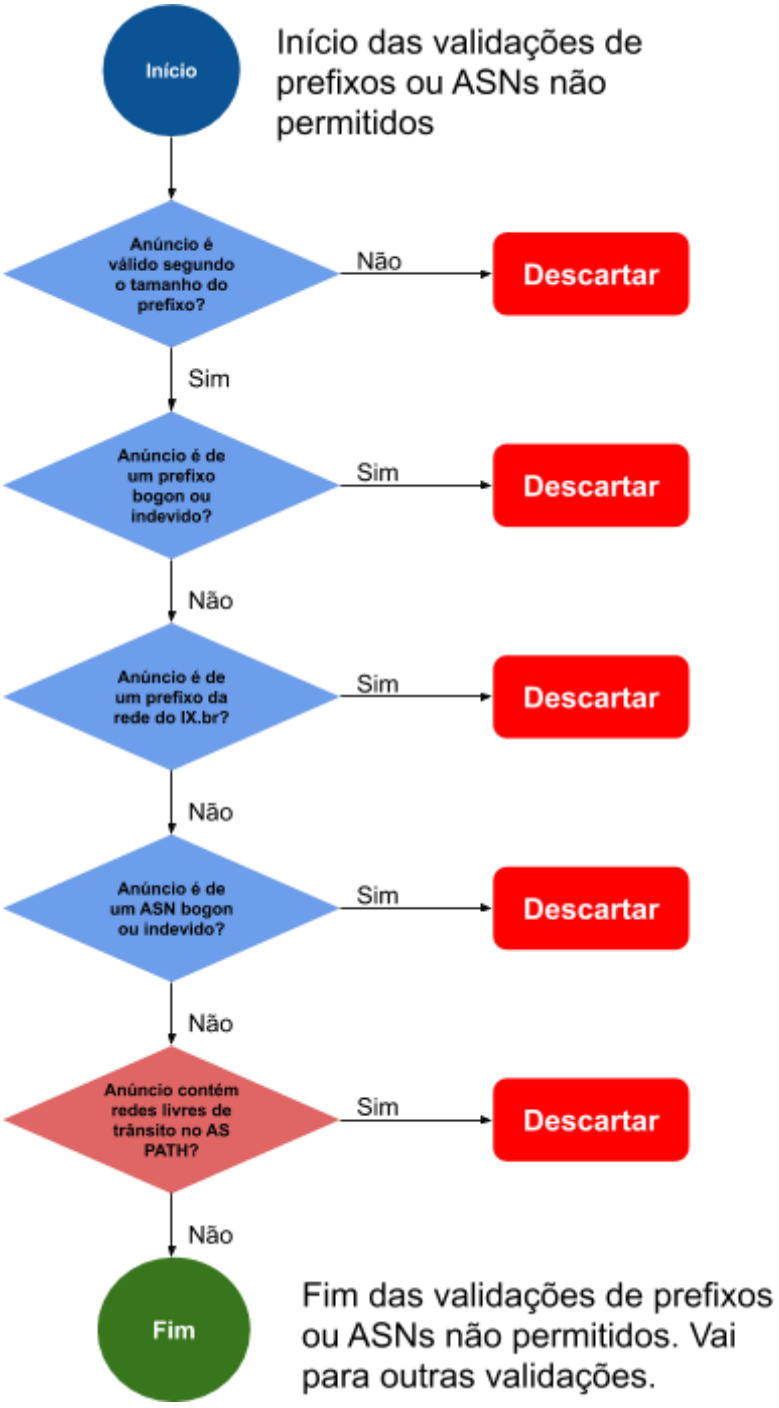

After the validation of prefixes or non-allowed ASNs in the IX.br network, the filters will be made based on the other validations. The following diagram outlines how the filters will be implemented in a SHORT term.

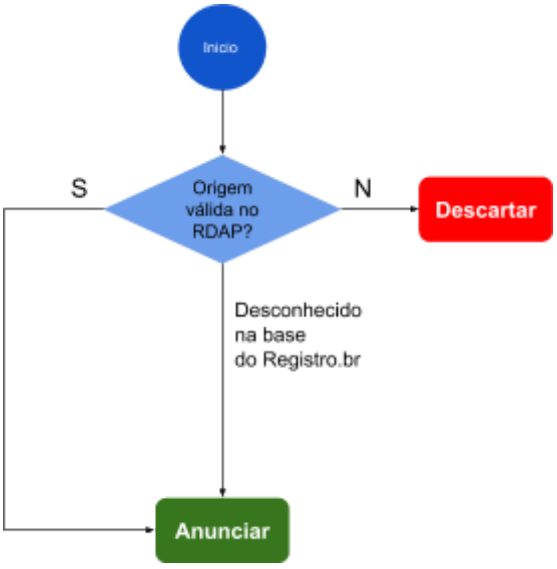

The diagram below outlines how the filters will be implemented in the AVERAGE term.

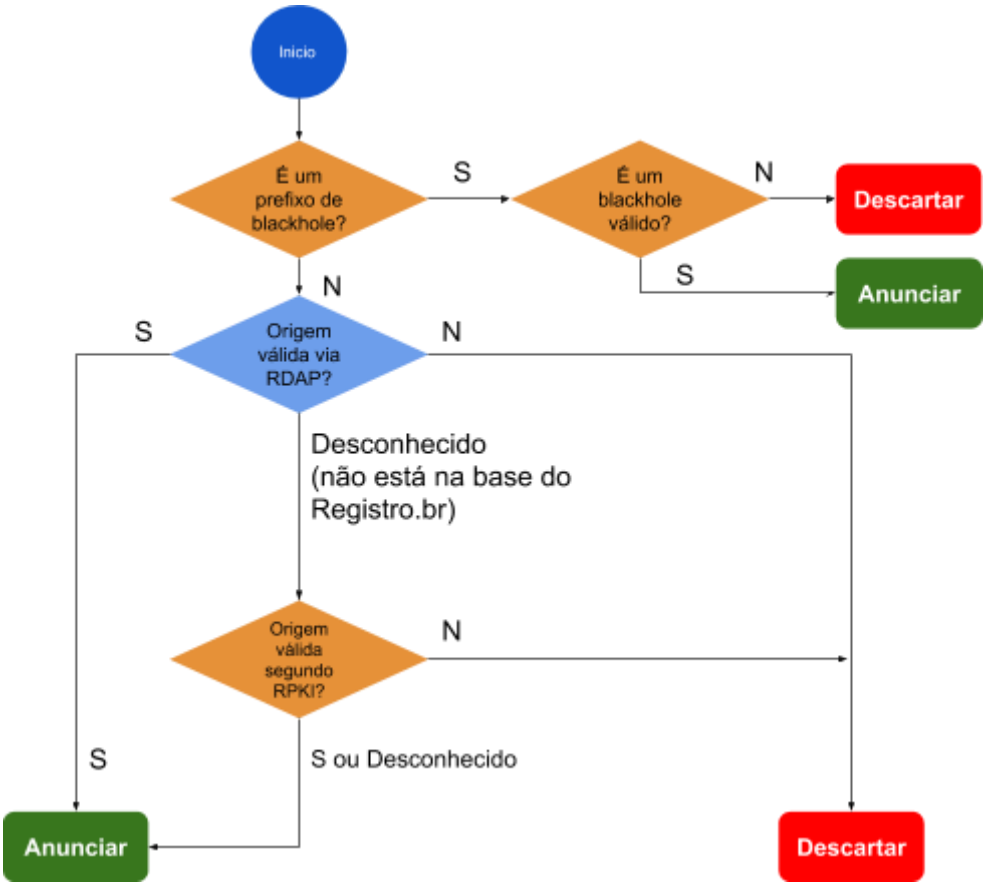

The diagram below outlines how the filters will be implemented in the LONG term.

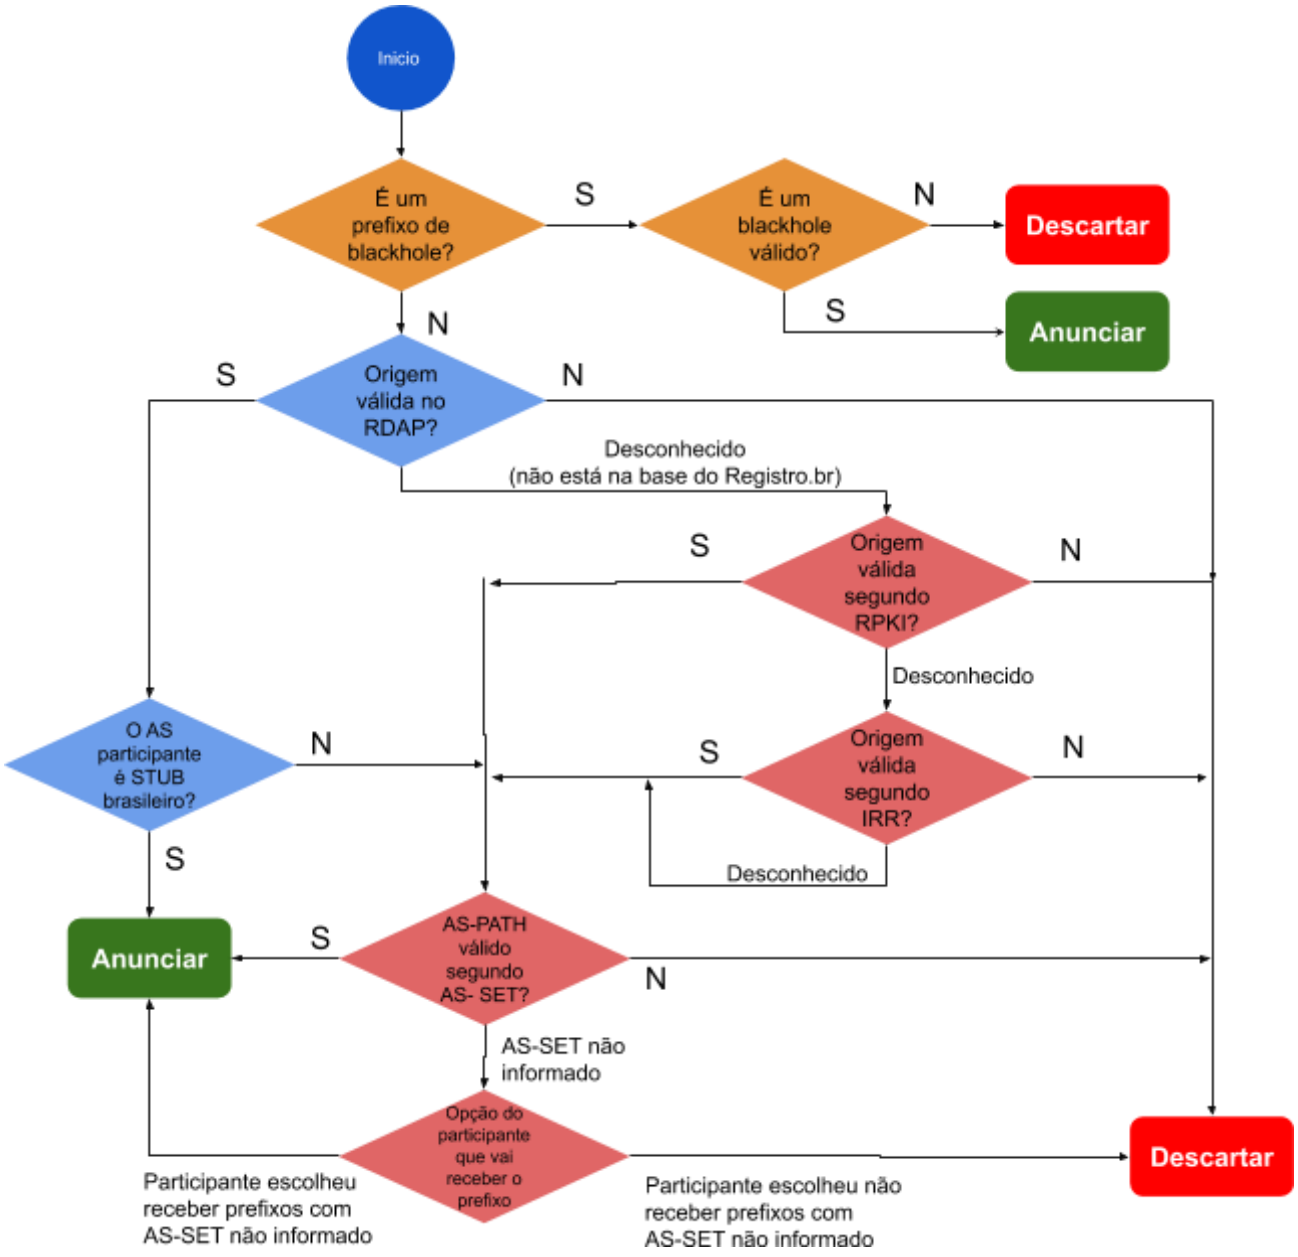

The LONG term diagram does not yet include filters using LACNIC or ARIN RDAP. Additional studies are needed to determine how they will be inserted into this flowchart.

For the AS SET filter, the participant, through the portal, will be allowed to choose whether or not to receive announcements from traffic participants and who have not informed AS SET.

It is possible that there is also the same choice for unknown origin, when the origin is not known neither via RDAP, nor RPKI, nor IRR. Additional studies, mainly on the effectiveness of RPKI validations in their current state of adoption, and IRR, are needed to determine if this option will actually exist.

Schedule of activities

1) Publication on the IX.br website of the first version of the document: 04/05/18.
2) Presentation to the community: at the 9th Regional Forum of São Paulo, on 05/17/2018, and at the GTER 45 meeting, 05/22/2018.
3) Receipt of comments via the GTER mailing list: until 06/04/18.
4) Analysis, preparation and publication on the IX.br website of the second version of the document, with the inclusion of suggestions from the community: 13/07/2018.
5) Receipt of comments via the GTER mailing list: until 07/27/2018 .
6) Publication of the final version of the document with the actions to be implemented: 08/01/2018.
7) Estimated schedule of implementation of actions:

Local	Short	Medium	Long
Aracaju	not defined	not defined	not defined
Belem	not defined	not defined	not defined
Belo Horizonte	not defined	not defined	not defined
Brasília	not defined	not defined	not defined
Campina Grande	not defined	not defined	not defined
Campinas	no defined	not defined	not defined
Caxias do Sul	not defined	not defined	not defined
Curitiba	not defined	not defined	not defined
Florianópolis	not defined	not defined	not defined
Fortaleza	not defined	not defined	not defined
Foz do Iguaçu	not defined	not defined	not defined
Goiânia	not defined	not defined	no defined
João Pessoa	not defined	not defined	not defined
Lajeado	not defined	not defined	not defined
Londrina	not defined	not defined	not defined
Maceió	not defined	not defined	not defined
Manaus	not defined	not defined	not defined

Maringa	not defined	no the defined	not defined
Natal	not defined	not defined	not defined
Porto Alegre	not defined	not defined	not defined
Recife	not defined	not defined	not defined
Rio de Janeiro	not defined	not defined	not defined
Salvador	not defined	not defined	not defined
Santa Maria	not defined	not defined	not defined
São José dos Campos	undefined	undefined	undefined
Sao Jose do Rio Preto	undefined	undefined	undefined
São Luis	undefined	not defined	not defined
São Paulo	10/16/2018	12/15/2018	15/09/2019
Teresina	not defined	not defined	not defined
Vitória	not defined	not defined	not defined

Remarks:

- São Paulo will be the base location of the new functionalities. Once approved, they will be replicated to the other locations.
- The deadlines presented are the maximum, and the implementation of functionalities may occur before the stated deadline.
- At first we are considering the implementation of all the functionalities in all the IX.br locations.
- The proposed dates are linked to the Schedule of Changes Implantation in the Route Servers of São Paulo, published in the IX.br website, Documentation area.

SUMMARY OF THE CONTRIBUTIONS RECEIVED (in the first version of this document)

1) I think they may include a "4.5-TiER-1-national" filter by listing in thelist

   RNP
   Embratel
   Tim
   Telefonica
   OI,
   Others?

2) Another is to have a "large content providers" class listing the ASNs of content / service providers such as Google, Netflix and Amazon ?? As a general rule, as far as I know, these guys are not transported by anyone up to PTT.Action

3) 4: is a functionally valid proposal, with only setback centralizes a lot of responsibility over NIC.br in order to make sure that it has the complete list of carriers t-1 (bankruptcy, new companies, etc.), and ends up being supplanted by others of the proposed actions. In the long run, this will be done, and in the long run, interrupt after the other proposals that include this type of limitation are already in place and in production.

4) Actions 5, 6 and 8: partial
   - According to RDAP proposal action 6.
   - The rest of the proposals have basically the same purpose and could be concentrated in a single method:
   - In my view the main problem of proposals like IRR is that they are decentralized systems and managed by totally independent entities, what makes it difficult to identify authority because there is no root entity (eg root servers that give authority to global DNS).
   - The RPKI proposal seems to me redundant since the security problem is more focused on authorization rather than proof of identity, unnecessarily increasing complexity.

5) Suggestions:
   The NIC.br itself should have an integrated system, accessed in a similar way as the.br and my.ix register, where ASNs can register and explicitly define which ASNs / prefixes can advertise their own ASN / prefix (s).
   In this way, having the same functionality as an IRR, but centralized and managed by NIC.br itself, guaranteeing authority over the Brazilian ASs.
   Ideally this process should be fully automated, via GUI (Portuguese + English) and optionally text commands (to facilitate the use of users' scripts), but if there is a certain difficulty in the implementation of such a system, called, similar to requests in IX.br.

   To make life easier for transit ASs, have an option to actively request another ASN if they

wish to allow the advertisement of their prefixes.
In this case, the traffic ASN would access your portal in NIC.br and click on an "I want to announce ASNs / third party prefixes" option, just fill in which ASNs and / or prefixes you want to advertise.
Once this is done, the NIC.br system automatically performs a whois on the reported resources and sends an email to the email account registered in the whois result.
This email should contain all information pertaining to the request and a URL link where you can click to confirm that it allows traffic.
If this link is clicked but they want to revoke the permission, they would then have to access their portal in the NIC.br and edit the permissions (if they have not yet registered, it would be necessary).

6) I've been talking since last year doing a dashboard, similar to icloud with all the services and tools of Nic.BR

7) "Action 4" seems a little too invasive. It is contemplated as an option for the participant to choose whether or not to receive such ads seems to me more consistent with the neutrality discourse that I have already heard to justifynegationsof other suggestions.
A broader IX.br policy would be more legal. Create a "we do not recommend this ad" policy and marking this with communities that indicate why it was inadvisable would be legal. It seems to me that I would need add-paths for this to work but IX.br chose not to use it and it seems that this is messing up. The choice, or proposal I heard years ago, to do this (thin control of what to receive) on the portal so far has not rolled and I prefer to interact only with the router connected to the IX than with it and the portal.

Comment received:
While section 4 indeed requires periodic review, it is a highly effective method to guard against large operational issues. I can assure you that any 2914 in the AS_PATH passing through any route server is either a misconfiguration, a software defect or a malicious activity. I do not think it wise to simply dismiss such 'ground truth' information.

Reply to comment:
Imagine the following scenario:
IX RS - AS4 enabled participant - not-AS4 capable downstream - AS4 capable participant

The AS Path on the IX RS could look like this, if all of them were AS4 capable:

65000 - 65001 - 65536

But, since 65001 is not AS4 capable, it will send this path upwards:

65000 - 65001 - 23456

Note that IX RS is AS4 capable, and the IX member (65000) is also AS4 capable. But, 65000 has a customer that it's not, and they in turn have a customer that is ASN greater than 65535.

So, this scenario requires no misconfigurations, and still present AS_TRANS (23456) in the path.

Reply to the comment above:
No, it is not visible on the route from a policy perspective, because the AS4_PATH attribute is used to tunnel through the AS4-capable ASN and 65536 will be the ASN visible on the route server.

The AS4 / non-AS4 transition mechanism is quite amazing: https://tools.ietf.org/html/rfc6793

8) "Action 5" does not seem to help me at all. At most this would be an attempt to protect the other participants from mistakes of an "AS Stub". I have seen errors of small, medium and large ASs being these new or old and if it is to establish relationships of confidence differentiated by AS so that something is done based on behavior and result. Good things increase trust and rights and bad things diminish them. It also seems to me that "Action 6" and "Action 8" overlap "Action 5"

   Comment received:
   I agree that besides not helping much will generate a lot of entropy and will be difficult to maintain.
   I also agree that these problems will be solved in the long run by other actions. By myself this item can be totally suppressed.

9) Action 2: I suggest increasing the scope of the filtered prefixes to those of the .br root-servers and anybr clusters in non-stub AS; in the AS stub, which includes any .br clusters of the .br, the filter will already allow only the space allocated to those AS.

10) Action 3: Will it not be necessary to accept 23456 if some AS path has a combination of routers without and with 32-bit AS support? An alternative would be to update the documentation that this support is a requirement for both the AS member and its downstreams. Still in action 3, perhaps put an example showing that the ad will be entirely rejected, and not just the AS bogon deleted. Although the text is clear, as there is equipment with this option, an example can help to not allow Doubt.

    Comment received:
    A BGP speaker that is capable of AS4, should never see AS 23456 in the AS_PATH. Any occurrence of AS 23456 visible on a 4-byte ASN capable router is either misconfiguration, or software defect. We should not reward misconfigurations by accepting these announcements.

11) Action 4: Maybe include the big Web-Scale in the same list, like Google, Netflix, Facebook, Akamai?

Comment received:
Route servers operators should only include such companies in this filter with their explicit permission.

Reply to comment:
Why would that differ from Tier-1 operators? I could see a point for doing the same for the same Tier-1 operators, but not for treating them differently.

Reply to comment above:

> [side-note: I prefer to use the term 'transit-free' instead of
> 'tier-1', because the term 'transit-free' is something we can verify
> to a degree, 1 'has no well-defined meaning.]

All the "Big Content" providers you mention, have some form of distributed CDN approach where they connect independent clusters (the islands) to the Internet and use the Internet for feeding / filling and serving cached data. In other words, parts of their ASN are not transit-free. This is a fundamental difference compared to the transit-free networks as proposed in the IX.br document, which are expected to operate as a coherent backbone.

I just want to make sure the filter, _as proposed_, has a lot of value too. This type of filter is documented in various places, such as http://bgpfilterguide.nlnog.net/guides/no_transit_leaks/ Broadening that filter without the ASN owner's consent might be trickier.

If the community decides that transit-free / transit-using networks should be treated the same, and that those networks can simply email IX.br "never allow announcements that have our ASN anywhere in the AS_PATH on your route servers", that is fine by me too I'm supportive on an approach opt-in approach, and an approach that is open to everyone that wants to use it. I'm also supportive of the proposed list as-is.

12) Action 6:
- In RDAP validation, consider the 1st. ASN AS Path independent of AS being or not IX.br. This will ensure consistency of the advertisement, even before RPKI validation is implemented. Do not do this validation for Brazilian AS-Stub for higher performance.
- In IRR validation, consider supporting not only RPSLng bases such as TC, RADB and RIPE, but also RPSL as the Registro.br.
- Still in IRR validation, it can be optional until a hijack episode occurs on that member; but from then on, mandatory.
- If the route is valid RPKI, forward even if it has invalid RDAP. This allows legitimate agreements to announce one block to another AS to be recognized.

13) Action 8: This is a potentially complex analysis of the DFZ graph; suggestion is to split this action into smaller deliverables. For example, start only with alerts to members that your AS is in the IRR import (action 6) of another or appeared in the BGP ad of another, and walk longer for more preventive validations.

14) I would suggest the issue of blackhole ads, now with RDAP validation for example, mainly for AS stub ads makes sense to have blackhole implemented as well. The biggest concern before was about ad validation, to prevent someone from putting another's block into blackhole. Now with this validation it gets easier, especially for AS stub.

    I do not have the answer to that, I just bring the question to the discussion.

Núcleo de Informação e Coordenação do Ponto BR

SUMMARY OF CONTRIBUTIONS RECEIVED (in the second version of this document)

1) One suggestion that occurred to me was to change the number of prefix limiters from the session to the participant to the exit to other participants. This would have two advantages:
   - For AS Stub, even if they do leak, this will not overturn the session with IX, because those prefixes advertised wrongly have been leaked. Any undue announcements may be handled by cases opened by the NOC of the IX with the participants as minor alarms, without however affecting operation.
   - For transit AS, this would expedite the retrieval of the service, since flapping sessions can impact the routers of these participants. In addition, one of the directions of traffic would be preserved as it would continue to receive the prefixes of other participants.

   Comment Team IX.br: maintaining the prefix limit on entry is also intended to protect route servers against malicious actions. Since we intend to mark invalid ads for export to Looking Glass Web, sending a multitude of ads with irregularities may consume the existing resources for the operation of Route Servers.

2) How complex is IX to use IRR to validate prefixes / AS-PATH?
   If the transits filtered the prefixes learned via AS-PATH, many of these problems would be avoided. Most of the major national transits already do so.

   Response to the above comment:

   - According to the document posted, the expected time for this is AVERAGE, that is, order of months.
   - Actually no, because today the prefix limit is directly in the participant's session, so the session will drop before the validation acts. The suggestion I gave was exactly to change that, and then all applicable validations would work before the prefix limit.
   - The AS-Path filtering problem is that every AS that appears on the path would need to have some AS-Path (IRR or not) registration against which to validate itself ... and that's a long way from reality. Note that the IX member having this would not suffice for validation, since actual validation depends on all AS ...

3) BCP38 ..... this alone was already helpful.
   What I want to talk about is a way for nic.br to be able to validate whether or not the ASN within the IX's are following recommendations.