

The background of the entire image is a dark gray circuit board pattern with white lines representing traces and components. A central horizontal band is a lighter gray gradient.

nic.br

Brazilian Network
Information Center

egi.br

Brazilian Internet
Steering Committee

registro.br cert.br cetic.br ceptro.br ceweb.br ix.br

nic.br egi.br

ceptro.br

Communities: uma ferramenta importante para engenharia de tráfego

ceptro.br nic.br cgi.br

Agenda

- Motivação
- Communities
 - Standard, Extended e Large Communities
 - Communities informativas
 - Communities de ação
- Casos de uso

Motivação



c|net Search CNET | Reviews | News | Video | How To | Deals | US Edition

CNET • Tech Culture • How Pakistan knocked YouTube offline (and how to make sure it never happens again)

How Pakistan knocked YouTube offline (and how to make sure it never happens again)

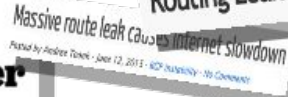
Posted by Andrew Tanzi • November 6, 2015 • Hijack • 1 Comment



0 MARCH 11, 2015 | COMMENTS (0) | VIEWS (1724) | ENGINEERING, IPTV, LATENCY, PERFORMANCE, SECURITY

Routing Leak briefly takes down Google

DOUG MAZOUZ



Massive route leak causes internet slowdown

Posted by Andrew Tanzi • June 12, 2011 • BGP, SECURITY • 39 Comments



0 JUNE 11, 2015 | COMMENTS (0) | VIEWS (411) | SECURITY, PERFORMANCE, SECURITY

Global Collateral Damage of TMnet leak



DDoS Attacks Storm Linode Servers Worldwide

BY DOUGLAS BONDERUD • JANUARY 5, 2016



0 OCTOBER 14, 2015 | COMMENTS (0) | VIEWS (161) | PERFORMANCE, SECURITY

Global Impacts of Reception

| Event type | Country | ASN | Start time |
|------------|---------|---------------------------------------------------------------------------------------------------------------|------------------------|
| BGP Leak | | Origin AS: PO box 1511 Phonyway road - Nayyetha district (AS 12100) Leaker AS: Veeba Corporation (AS 7550) | 2016-01-13 12:25:47 |
| BGP Leak | | Origin AS: Linex.net EOOD (AS 8268) Leaker AS: Traffic Broadband Communications Ltd. (AS 48452) | 2016-01-13 12:11:26 |



0 MARCH 13, 2015 | COMMENTS (0) | VIEWS (4787) | SECURITY

UK traffic diverted through Ukraine

DOUG MAZOUZ



On-going BGP Hijack Targets Palestinian ISP



BGP hijack incident by Syrian Telecommunications

Posted by Andrew Tanzi • December 9, 2014 • Hijack • 2 Comments



0 JANUARY 21, 2015 | COMMENTS (0) | VIEWS (348) | SECURITY

The Vast World of Fraudulent Routing

DOUG MAZOUZ




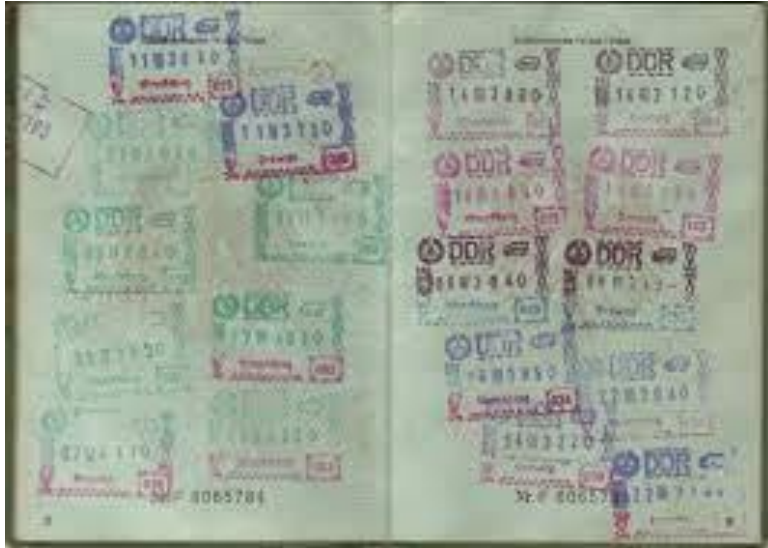
CSO Most read: | Home | Data Protection | Cyber Attacks/Espionage

TODAY'S TOP STORIES

DDoS attack on BBC may have been biggest in history

Communities

Communities



Communities

- Atributo adicionado a rota no BGP
 - Transitivo e Opcional
- Uma rota pode ter várias communities distintas
- Existem 3 tipos
 - Standard Communities (ou regular)
 - Extended Communities
 - Large Communities

Standard Communities

- Definido na RFC1997 (ano de 1996) - mais antigo
- Amplamente utilizado!
- Número de 32 bits (integer)
 - 2 campos de 2 octetos
 - Formato
 - 16 bits :16 bits
 - Para evitar repetições
 - Primeiros 16 bits normalmente são ASNs
 - Número de AS : Número de Community
 - E quando o ASN tem 32 bits?

Standard Communities Bem Conhecidas

- Lista da IANA
- Alguns roteadores já possuem a definição por padrão
- Mas precisa se implementar a sua ação
- Standard
 - <https://www.iana.org/assignments/bgp-well-known-communities/bgp-well-known-communities.xhtml>
 - Exemplo:
 - Blackhole
 - Graceful shutdown

Extended Communities

- Definido na RFC4360 (ano de 2006)
- Número de 64 bits (integer)
 - 1 ou 2 octetos para identificar o Tipo
 - Tipos registrados pela IANA
 - <https://www.iana.org/assignments/bgp-extended-communities/bgp-extended-communities.xhtml>
 - Formato diversos
 - (Tipo)16 bits : (ASN) 16 bits : (Local) 32 bits
 - (Tipo)16 bits : (IPv4) 32 bits : (Local) 16 bits
 - (Tipo)16 bits : (Local) 48bits

Large Communities

- Definido na RFC8092 (ano de 2017) - novo
- Crescendo em implantação - muitos ainda não utilizam.
- Número de 96 bits (integer)
 - 3 campos de 4 octetos
 - Formato
 - 32 bits : 32 bits : 32 bits
 - ASN : Função : Parâmetro
 - Comum - Meu ASN : Ação : outro ASN

Large Communities

BGP Speakers

| Vendor | Software | Status | Details |
|------------------|------------|---------|--------------------------------|
| Arista | EOS | ✓ Done! | EOS 4.21.3F |
| Cisco | IOS XE | ✓ Done! | IOS XE 17.4 |
| Cisco | IOS XR | ✓ Done! | IOS XR 6.3.2 |
| cz.nic | BIRD | ✓ Done! | BIRD 1.6.3 (commit) |
| Extreme | NetIron | ✓ Done! | NetIron 06.3.00 |
| Extreme | SLX-OS | ✓ Done! | SLX-OS 18r.2.00_v1 |
| ExaBGP | ExaBGP | ✓ Done! | PR482 |
| FreeRangeRouting | frr | ✓ Done! | Issue 46 (commit) |
| Juniper | Junos OS | ✓ Done! | Junos OS 17.3R1 |
| MikroTik | RouterOS | ✓ Done! | ROSv7 |
| Nokia | SR OS | ✓ Done! | SR OS 16.0.R1 |
| nop.hu | freeRouter | ✓ Done! | |
| OpenBSD | OpenBGPD | ✓ Done! | OpenBSD 6.1 (commit) |
| OSRG | GoBGP | ✓ Done! | PR1094 |
| rtbrick | Fullstack | ✓ Done! | FullStack 17.1 |
| Quagga | Quagga | ✓ Done! | Quagga 1.2.0 (875) |
| Ubiquiti | EdgeOS | Planned | Feature Requested (maybe 2.0?) |
| VyOS | VyOS | ✓ Done! | Vyos 1.2.0 |

<http://largebgpcommunities.net/implementations/>

Como usar Communities?

- Duas categorias
 - **Marcação Informativas**
 - Servem para trazer alguma referência sobre a rota
 - Ex: De que país veio a rota
 - **Marcação de Ação**
 - Servem para influenciar nas políticas de roteamento
 - Ex: aumente o local preference de determinada rota

Communities Informativas

- Geralmente o foco em transmitir:
 - Onde a rota foi aprendida
 - Rotas continentais, nacionais, regionais ...
 - Como a rota foi aprendida
 - Trânsito, peer, cliente, interna ...
- Ela pode ser usada para tomar alguma ação
 - Por você
 - Por um trânsito
 - Por um outro

Communities Informativas

- Não existe um padrão definido
- Cada AS faz a sua regra

Recomendações

- Não confunda as communities Informativas com as de Ação!
 - Deixe fácil de identificar cada uma delas
 - Ex: Informativas 5 dígitos e Ação 4 dígitos
- Não deixe que outros te enviem as suas communities informativas
 - Filtre elas antes que tenha problemas!

Communities de ação

- Principais funções
 - Manipular os atributos de BGP
 - AS-PATH
 - Local-Preference
 - Multi-Exit Discriminator (MED)
 - Next-Hop Address
 - Entre outros
 - Influenciar nas rotas repassadas
 - Repassa ou não (export)

Communities de ação

- Muitas vezes o repasse de rotas pode ser escolhido
 - Geografia
 - Não repassar as rotas internacionais
 - Relacionamento
 - Não repassar as rotas para o IX
 - Para um AS específico
 - Não repassar as rotas para o ASXXX
 - Ou outra regra que queira inventar e divulgar
 - Não deixe de publicar.

Recomendações

- Publique e atualize a documentação sobre as suas communities
 - Ajuda debug de problemas
 - Auxilia os filtros
- Encoraje seus clientes a filtrarem as suas communities de ação na entrada se necessário
- Provenha um Looking Glass que mostre as communities

IX.br possui communities

Traffic engineering

| function | standard | extended | large |
|-----------------------------------|----------------|------------------------|------------------|
| not announce to ASN | 65000:dest-asn | (ro rt):65000:dest-asn | 65000:0:dest-asn |
| export only to ASN | 65001:dest-asn | (ro rt):65001:dest-asn | 65001:0:dest-asn |
| add one prepend | 64601:dest-asn | (ro rt):64601:dest-asn | 64601:0:dest-asn |
| add two prepend | 64602:dest-asn | (ro rt):64602:dest-asn | 64602:0:dest-asn |
| add three prepend | 64603:dest-asn | (ro rt):64603:dest-asn | 64603:0:dest-asn |
| graceful shutdown | 65535:0 | - | - |
| do not announce to Afrinic (ASNs) | 65002:0 | (ro rt):65002:0 | 65002:0:0 |
| do not announce to Apnic (ASNs) | 65002:1 | (ro rt):65002:1 | 65002:0:1 |
| do not announce to Arin (ASNs) | 65002:2 | (ro rt):65002:2 | 65002:0:2 |
| do not announce to Lacnic (ASNs) | 65002:3 | (ro rt):65002:3 | 65002:0:3 |
| do not announce to Ripe (ASNs) | 65002:4 | (ro rt):65002:4 | 65002:0:4 |
| do not announce to Brasil (ASNs) | 65002:5 | (ro rt):65002:5 | 65002:0:5 |
| export to Afrinic (ASNs) | 65003:0 | (ro rt):65003:0 | 65003:0:0 |
| export to Apnic (ASNs) | 65003:1 | (ro rt):65003:1 | 65003:0:1 |
| export to Arin (ASNs) | 65003:2 | (ro rt):65003:2 | 65003:0:2 |
| export to Lacnic (ASNs) | 65003:3 | (ro rt):65003:3 | 65003:0:3 |
| export to Ripe (ASNs) | 65003:4 | (ro rt):65003:4 | 65003:0:4 |
| export to Brasil (ASNs) | 65003:5 | (ro rt):65003:5 | 65003:0:5 |
| do not announce to IXP | 65004:65XXX | (ro rt):65004:65XXX | 65004:0:65XXX |

| | | | |
|----------------------------------|-----------|-------------------|-------------|
| do not announce to rtt > 10ms | 65010:10 | (ro rt):65010:10 | 65010:0:10 |
| do not announce to rtt > 50ms | 65010:50 | (ro rt):65010:50 | 65010:0:50 |
| do not announce to rtt > 100ms | 65010:100 | (ro rt):65010:100 | 65010:0:100 |
| do not announce to rtt > 150ms | 65010:150 | (ro rt):65010:150 | 65010:0:150 |
| do not announce to rtt > 200ms | 65010:200 | (ro rt):65010:200 | 65010:0:200 |
| do not announce to rtt > 250ms | 65010:250 | (ro rt):65010:250 | 65010:0:250 |
| do not announce to rtt = unknown | 65010:999 | (ro rt):65010:999 | 65010:0:999 |
| | | | |
| one prepend to rtt > 10ms | 64611:10 | (ro rt):64611:10 | 64611:0:10 |
| one prepend to rtt > 50ms | 64611:50 | (ro rt):64611:50 | 64611:0:50 |
| one prepend to rtt > 100ms | 64611:100 | (ro rt):64611:100 | 64611:0:100 |
| one prepend to rtt > 150ms | 64611:150 | (ro rt):64611:150 | 64611:0:150 |
| one prepend to rtt > 200ms | 64611:200 | (ro rt):64611:200 | 64611:0:200 |
| one prepend to rtt > 250ms | 64611:250 | (ro rt):64611:250 | 64611:0:250 |
| one prepend to rtt = unknown | 64611:999 | (ro rt):64611:999 | 64611:0:999 |
| | | | |
| two prepend to rtt > 10ms | 64612:10 | (ro rt):64612:10 | 64612:0:10 |
| two prepend to rtt > 50ms | 64612:50 | (ro rt):64612:50 | 64612:0:50 |
| two prepend to rtt > 100ms | 64612:100 | (ro rt):64612:100 | 64612:0:100 |
| two prepend to rtt > 150ms | 64612:150 | (ro rt):64612:150 | 64612:0:150 |
| two prepend to rtt > 200ms | 64612:200 | (ro rt):64612:200 | 64612:0:200 |
| two prepend to rtt > 250ms | 64612:250 | (ro rt):64612:250 | 64612:0:250 |
| two prepend to rtt = unknown | 64612:999 | (ro rt):64612:999 | 64612:0:999 |
| | | | |
| three prepend to rtt > 10ms | 64613:10 | (ro rt):64613:10 | 64613:0:10 |
| three prepend to rtt > 50ms | 64613:50 | (ro rt):64613:50 | 64613:0:50 |
| three prepend to rtt > 100ms | 64613:100 | (ro rt):64613:100 | 64613:0:100 |
| three prepend to rtt > 150ms | 64613:150 | (ro rt):64613:150 | 64613:0:150 |
| three prepend to rtt > 200ms | 64613:200 | (ro rt):64613:200 | 64613:0:200 |
| three prepend to rtt > 250ms | 64613:250 | (ro rt):64613:250 | 64613:0:250 |
| three prepend to rtt = unknown | 64613:999 | (ro rt):64613:999 | 64613:0:999 |

| | | | |
|----------------------------------|-----------|-------------------|-------------|
| do not announce to loss > 2% | 65011:2 | (ro rt):65011:2 | 65011:0:2 |
| do not announce to loss > 10% | 65011:10 | (ro rt):65011:10 | 65011:0:10 |
| do not announce to loss = unknow | 65011:999 | (ro rt):65011:999 | 65011:0:999 |
| | | | |
| one prepend to loss > 2% | 64621:2 | (ro rt):64621:2 | 64621:0:2 |
| one prepend to loss > 10% | 64621:10 | (ro rt):64621:10 | 64621:0:10 |
| one prepend to loss = unknown | 64621:999 | (ro rt):64621:999 | 64621:0:999 |
| | | | |
| two prepend to loss > 2% | 64622:2 | (ro rt):64622:2 | 64622:0:2 |
| two prepend to loss > 10% | 64622:10 | (ro rt):64622:10 | 64622:0:10 |
| two prepend to loss = unknown | 64622:999 | (ro rt):64622:999 | 64622:0:999 |
| | | | |
| three prepend to loss > 2% | 64623:2 | (ro rt):64623:2 | 64623:0:2 |
| three prepend to loss > 10% | 64623:10 | (ro rt):64623:10 | 64623:0:10 |
| three prepend to loss = unknown | 64623:999 | (ro rt):64623:999 | 64623:0:999 |

Informatives

| function | community | extended | large |
|---------------------------|------------------|-------------------------|-------------------|
| Origin ASN | rs-asn:peer-asn | (ro rt):rs-asn:peer-asn | rs-asn:0:peer-asn |
| IXP location (XX = DDD) | rs-asn:65XXX | (ro rt):rs-asn:65XXX | rs-asn:0:65XXX |
| ASN - RTT 0.001ms < 10ms | rs-asn:64661 | (ro rt):rs-asn:64661 | rs-asn:660:1 |
| ASN - RTT 10ms < 50ms | rs-asn:64662 | (ro rt):rs-asn:64662 | rs-asn:660:2 |
| ASN - RTT 50ms < 100ms | rs-asn:64663 | (ro rt):rs-asn:64663 | rs-asn:660:3 |
| ASN - RTT 100ms < 150ms | rs-asn:64664 | (ro rt):rs-asn:64664 | rs-asn:660:4 |
| ASN - RTT 150ms < 200ms | rs-asn:64665 | (ro rt):rs-asn:64665 | rs-asn:660:5 |
| ASN - RTT 200ms < 250ms | rs-asn:64666 | (ro rt):rs-asn:64666 | rs-asn:660:6 |
| ASN - RTT > 250ms | rs-asn:64667 | (ro rt):rs-asn:64667 | rs-asn:660:7 |
| ASN - RTT = unknown | rs-asn:64669 | (ro rt):rs-asn:64669 | rs-asn:660:9 |
| ASN - LOSS 0% | rs-asn:64671 | (ro rt):rs-asn:64671 | rs-asn:670:1 |
| ASN - LOSS 0.001% < 2% | rs-asn:64672 | (ro rt):rs-asn:64672 | rs-asn:670:2 |
| ASN - LOSS 2% < 10% | rs-asn:64673 | (ro rt):rs-asn:64673 | rs-asn:670:3 |
| ASN - LOSS > 10% | rs-asn:64674 | (ro rt):rs-asn:64674 | rs-asn:670:4 |
| ASN - LOSS = unknown/100% | rs-asn:64679 | (ro rt):rs-asn:64679 | rs-asn:670:9 |
| ASN from Afrinic | rs-asn:64680 | (ro rt):rs-asn:64680 | rs-asn:680:0 |
| ASN from Apnic | rs-asn:64681 | (ro rt):rs-asn:64681 | rs-asn:680:1 |
| ASN from Arin | rs-asn:64682 | (ro rt):rs-asn:64682 | rs-asn:680:2 |
| ASN from Lacnic | rs-asn:64683 | (ro rt):rs-asn:64683 | rs-asn:680:3 |
| ASN from Ripe | rs-asn:64684 | (ro rt):rs-asn:64684 | rs-asn:680:4 |
| ASN from Brazil | rs-asn:64685 | (ro rt):rs-asn:64685 | rs-asn:680:5 |

| function | standard | extended | large |
|------------------------------------|--------------|----------------------|----------------|
| BH announce | 65535:666 | (ro rt):65535:666 | 65535:616:666 |
| Confirmed BH | rs-asn:666 | (ro rt):rs-asn:666 | rs-asn:616:666 |
| registro.br invalid | rs-asn:65110 | (ro rt):rs-asn:65110 | rs-asn:100:0 |
| registro.br valid | rs-asn:65111 | (ro rt):rs-asn:65111 | rs-asn:100:1 |
| registro.br unknown | rs-asn:65112 | (ro rt):rs-asn:65112 | rs-asn:100:2 |
| IRR - prefix present in AS-SET | rs-asn:65121 | (ro rt):rs-asn:65121 | rs-asn:200:1 |
| IRR - prefix not present in AS-SET | rs-asn:65122 | (ro rt):rs-asn:65122 | rs-asn:200:2 |
| RPKI invalid | rs-asn:65130 | (ro rt):rs-asn:65130 | rs-asn:300:0 |
| RPKI valid | rs-asn:65131 | (ro rt):rs-asn:65131 | rs-asn:300:1 |
| RPKI unknown | rs-asn:65132 | (ro rt):rs-asn:65132 | rs-asn:300:2 |
| RIRs invalid | rs-asn:65140 | (ro rt):rs-asn:65140 | rs-asn:400:0 |
| RIRs valid | rs-asn:65141 | (ro rt):rs-asn:65141 | rs-asn:400:1 |
| RIRs unknown | rs-asn:65142 | (ro rt):rs-asn:65142 | rs-asn:400:2 |
| Invalid prefix length | rs-asn:65190 | (ro rt):rs-asn:65190 | rs-asn:190:0 |
| Bogon prefix | rs-asn:65191 | (ro rt):rs-asn:65191 | rs-asn:190:1 |
| Bogon asn | rs-asn:65192 | (ro rt):rs-asn:65192 | rs-asn:190:2 |
| Transit free | rs-asn:65193 | (ro rt):rs-asn:65193 | rs-asn:190:3 |
| Never via RS | rs-asn:65194 | (ro rt):rs-asn:65194 | rs-asn:190:4 |
| IXPs prefixes | rs-asn:65195 | (ro rt):rs-asn:65195 | rs-asn:190:5 |

| | |
|---------------------------|--------------|
| ATM/MLPA prefixes | drop inbound |
| Resending our routes back | |
| Source != BGP | |
| AS-PATH length | |
| Drop IXP ASN | |
| First ASN != peer-asn | |
| next_hop != peer-ip | |

| | | | |
|----------------|--------------|----------------------|---|
| as-set invalid | rs-asn:65150 | (ro rt):rs-asn:65150 | - |
| as-set valid | rs-asn:65151 | (ro rt):rs-asn:65151 | - |
| | | | |
| stub-br | rs-asn:65180 | (ro rt):rs-asn:65180 | - |
| non-stub-br | rs-asn:65181 | (ro rt):rs-asn:65181 | - |

Locations ID

| IXP location | Origin Location | Do not announce to this IXP |
|-------------------|-----------------|-----------------------------|
| Aracaju/SE | 26162:65079 | 65004:65079 |
| Belém/PA | 26162:65092 | 65004:65092 |
| Belo Horizonte/MG | 26162:65031 | 65004:65031 |
| Boa Vista/RR | 26162:65095 | 65004:65095 |
| Brasília/DF | 26162:65061 | 65004:65061 |
| Campina Grande/PB | 26162:65183 | 65004:65183 |
| Campinas/SP | 26162:65019 | 65004:65019 |
| Campo Grande/MS | 26162:65067 | 65004:65067 |
| Cascavel/PR | 26162:65145 | 65004:65145 |
| Caxias do Sul/RS | 26162:65054 | 65004:65054 |
| Cuiabá/MT | 26162:65065 | 65004:65065 |
| Curitiba/PR | 26162:65041 | 65004:65041 |
| Florianópolis/SC | 26162:65048 | 65004:65048 |
| Fortaleza/CE | 26162:65085 | 65004:65085 |
| Foz do Iguaçu/PR | 26162:65045 | 65004:65045 |
| Goiânia/GO | 26162:65062 | 65004:65062 |
| João Pessoa/PB | 26162:65083 | 65004:65083 |
| Lajeado/RS | 26162:65051 | 65004:65051 |
| Londrina/PR | 26162:65043 | 65004:65043 |

| | | |
|---------------------|-------------|-------------|
| Maceió/AL | 26162:65082 | 65004:65082 |
| Manaus/AM | 26162:65092 | 65004:65092 |
| Maringá/PR | 26162:65044 | 65004:65044 |
| Natal/RN | 26162:65084 | 65004:65084 |
| Palmas/TO | 26162:65063 | 65004:65063 |
| Porto Alegre/RS | 26162:65051 | 65004:65051 |
| Recife/PE | 26162:65081 | 65004:65081 |
| Rio de Janeiro/RJ | 26162:65021 | 65004:65021 |
| Salvador/BA | 26162:65071 | 65004:65071 |
| Santa Maria/RS | 26162:65055 | 65004:65055 |
| S. J. dos Campos/SP | 26162:65012 | 65004:65012 |
| S. J. Rio Preto/SP | 26162:65017 | 65004:65017 |
| São Luís/MA | 26162:65098 | 65004:65098 |
| São Paulo/SP | 26162:65011 | 65004:65011 |
| Teresina/PI | 26162:65086 | 65004:65086 |
| Vitória/ES | 26162:65027 | 65004:65027 |

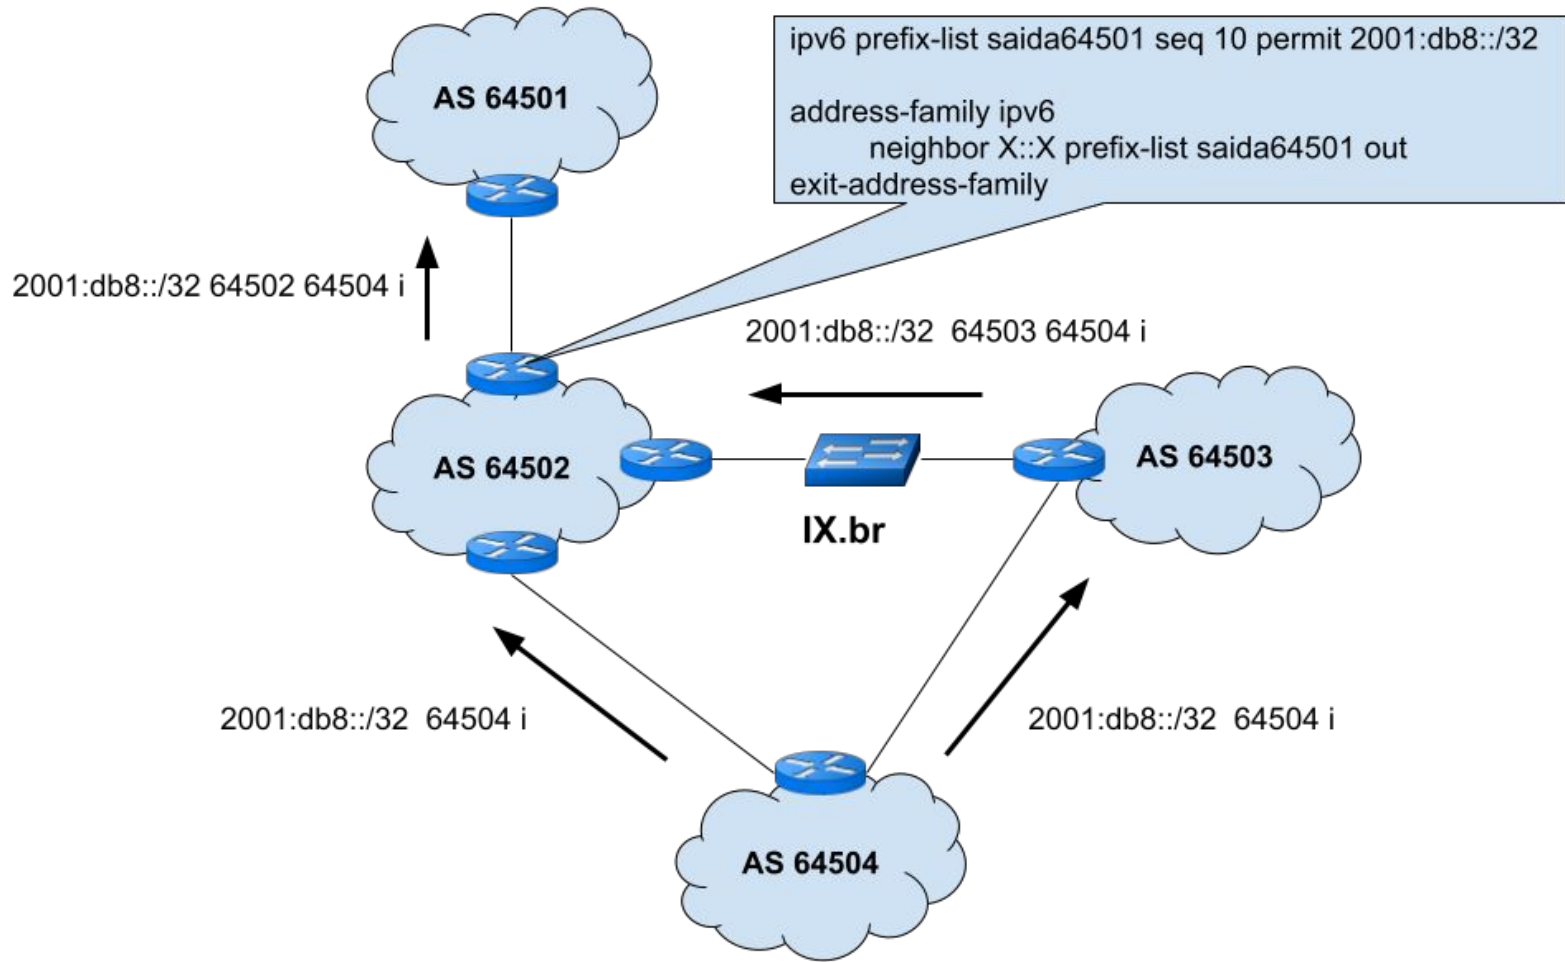
Casos de estudo

Caso I

Evitando vazamento de Rotas

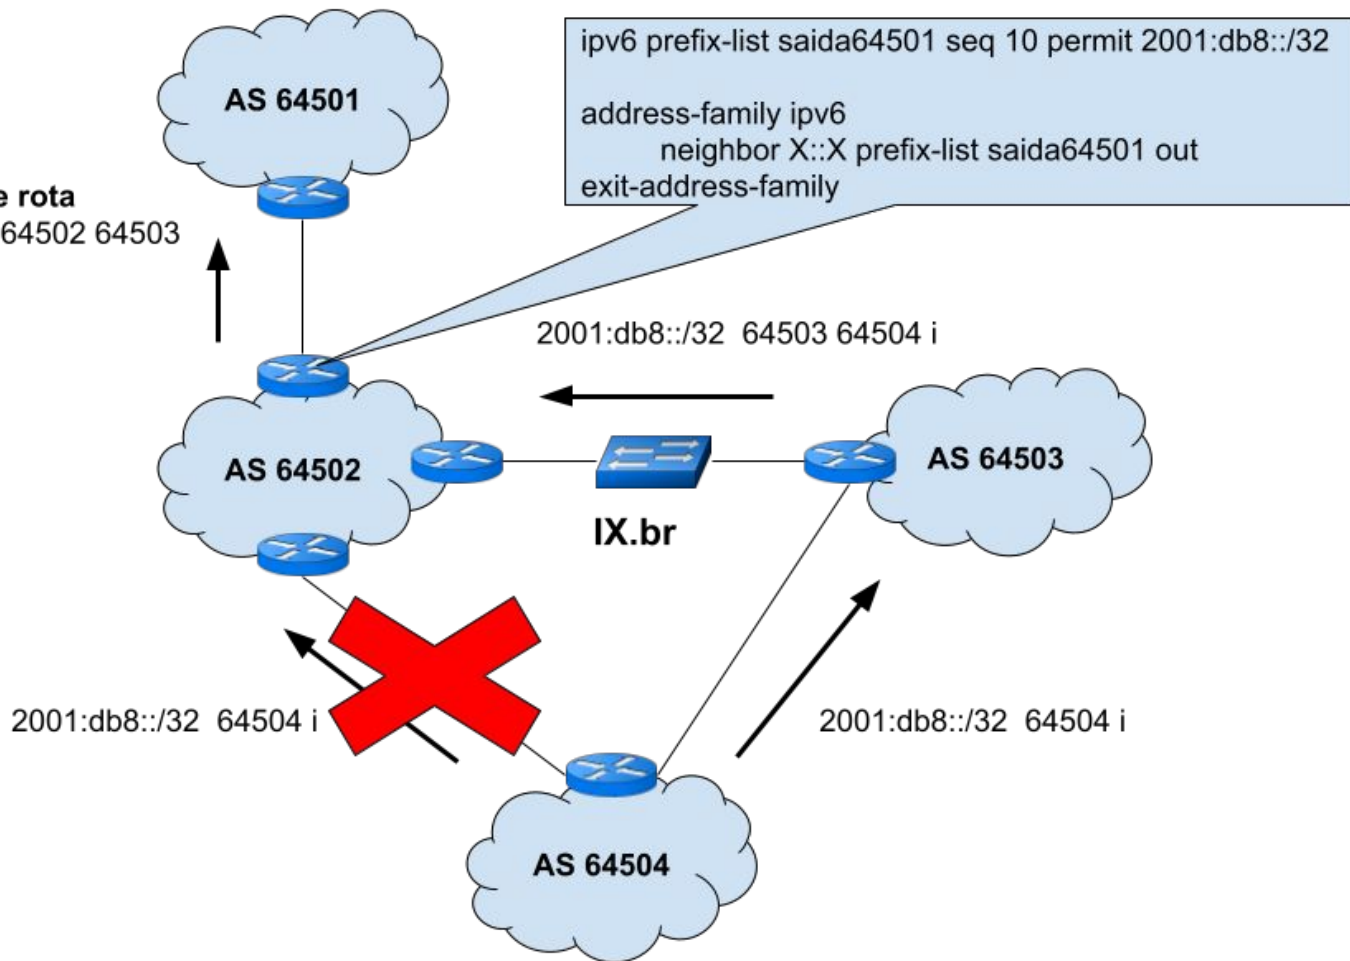
Vazamento de Rotas

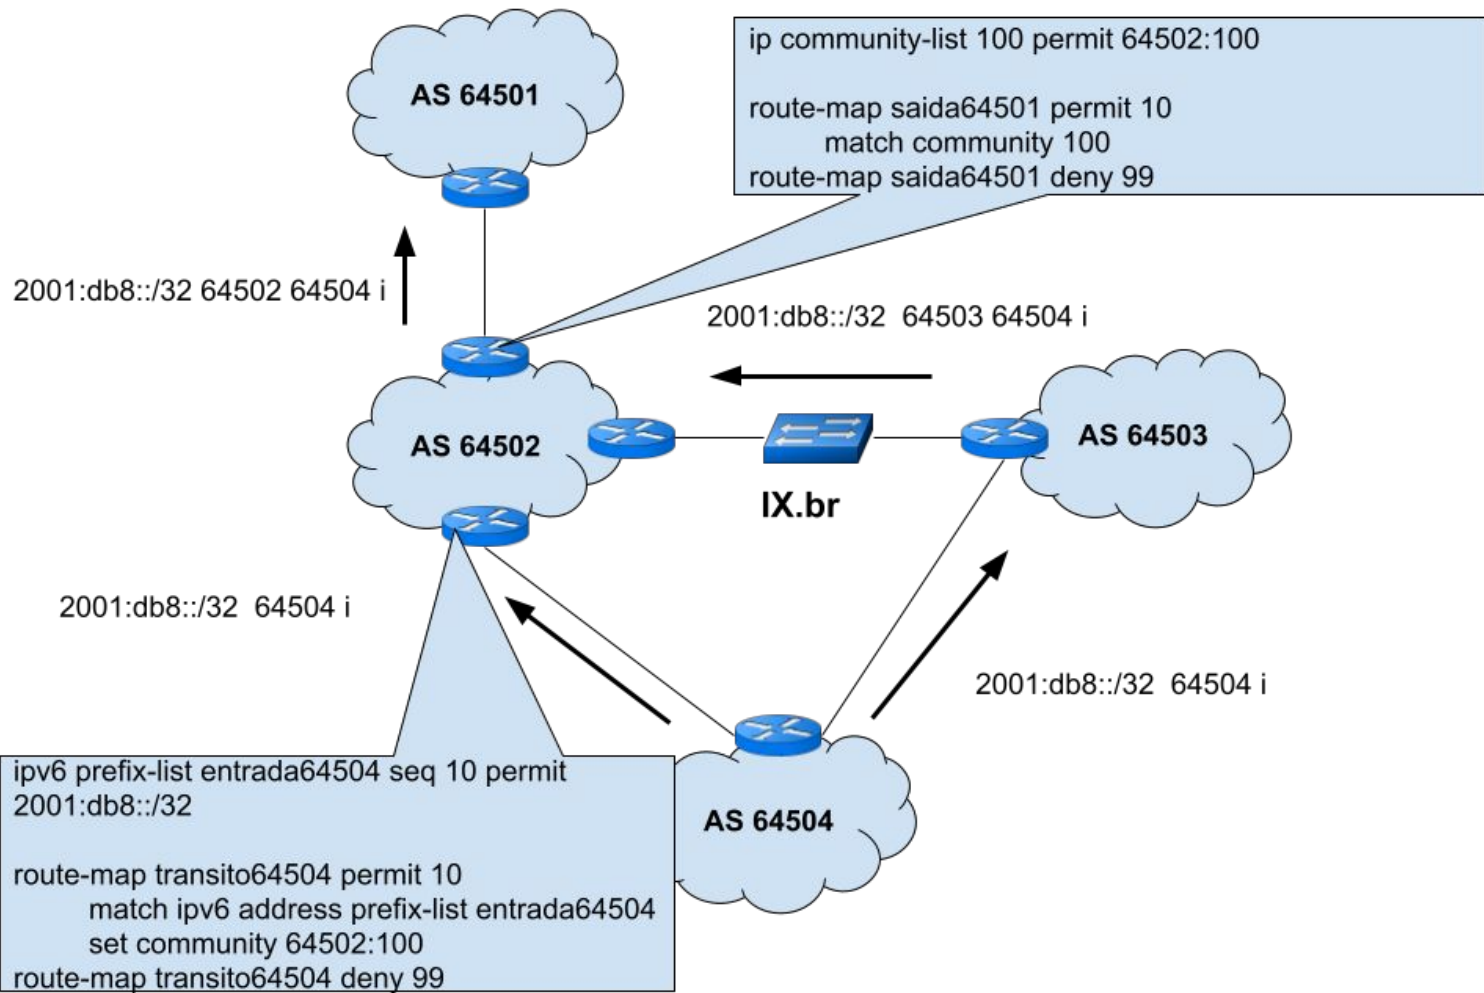
- Acontece quando uma rota é repassada e não deveria ser.
- O MANRS ensina como resolver isso!
 - Filtros
 - Prefixo
 - AS-path
 - Mas precisa saber aplicar corretamente.
 - <https://www.manrs.org/>
- Communities não são essenciais mas podem ajudar.



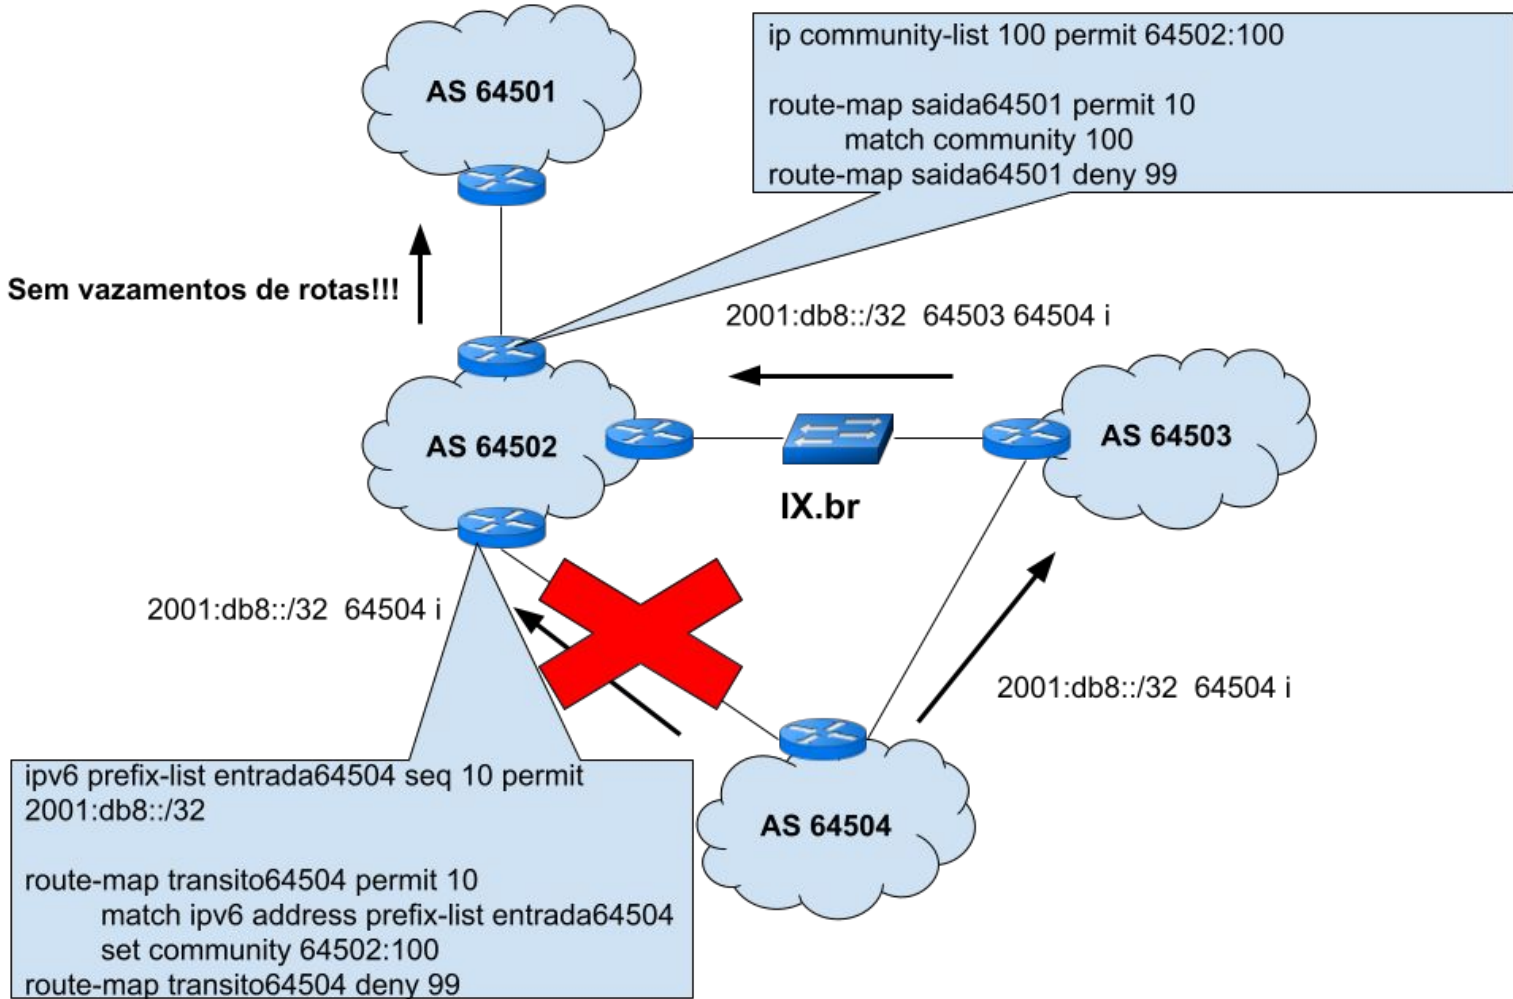
Vazamento de rota

2001:db8::/32 64502 64503
64504 i





Sem vazamentos de rotas!!!

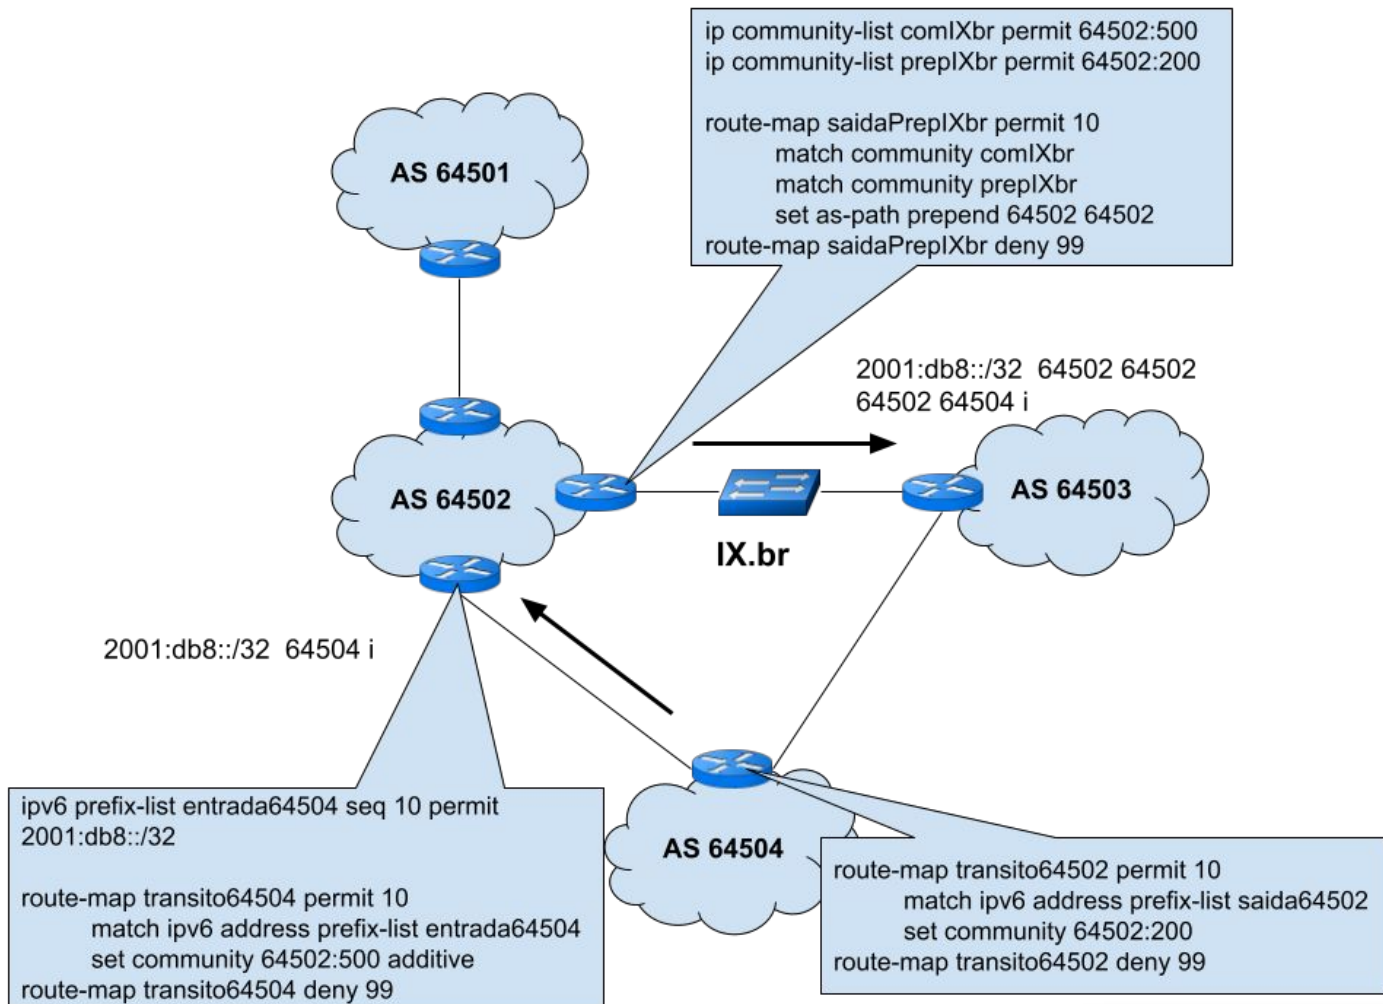


Caso 2

Pedindo para o trânsito piorar a rota no IX.br

Pedindo para o Trânsito

- Muitos trânsitos fornecem uma documentação de communities que se pode utilizar
- As vezes é possível influenciar como a sua rota é enviada para a saída dele
 - Mas nem sempre!
 - Lembre que eles possuem links com custos diferentes

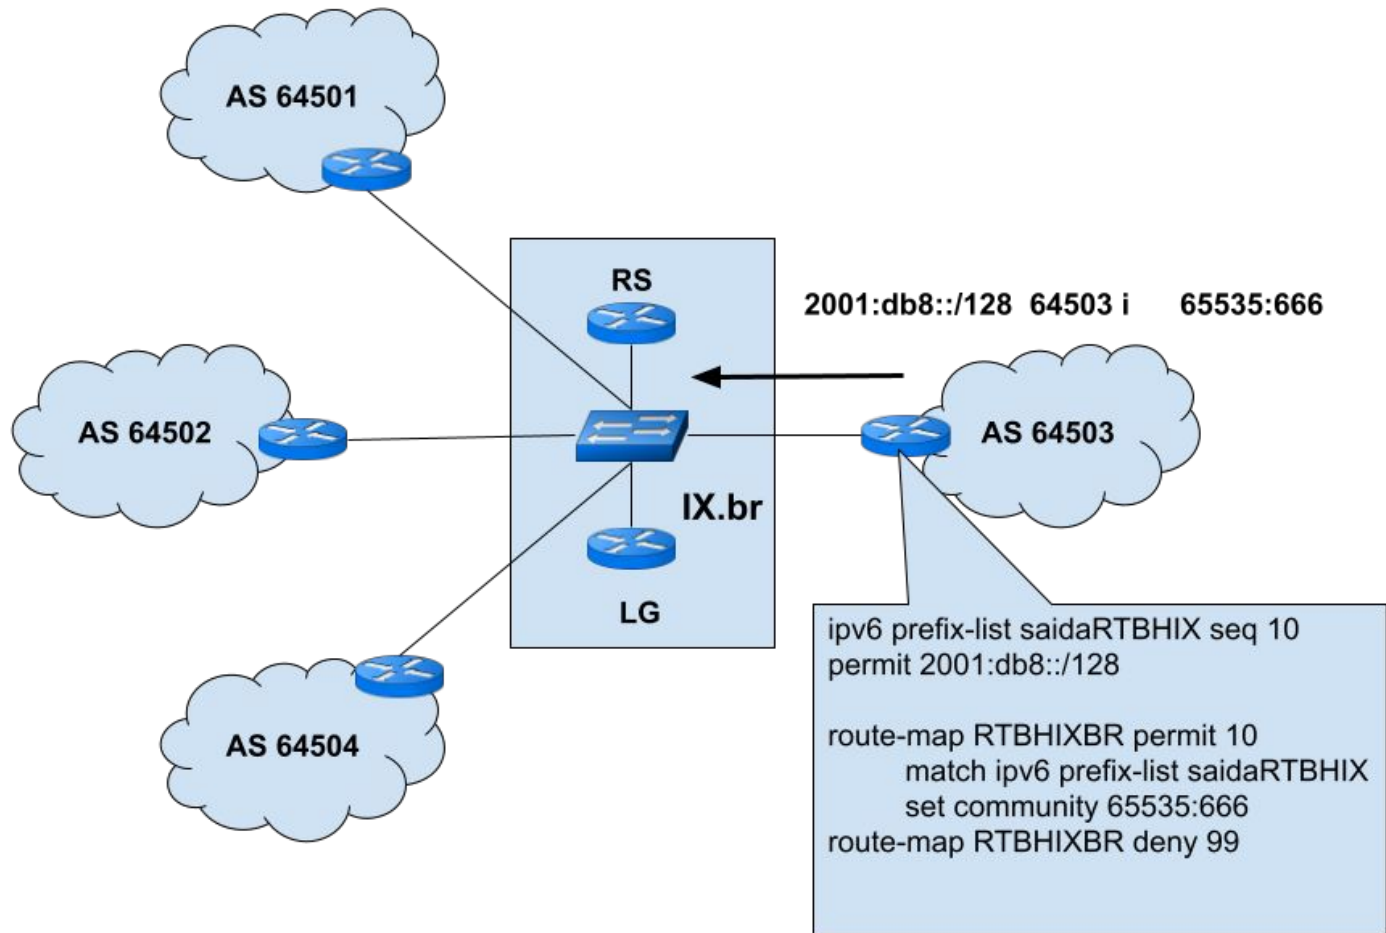


Caso 4

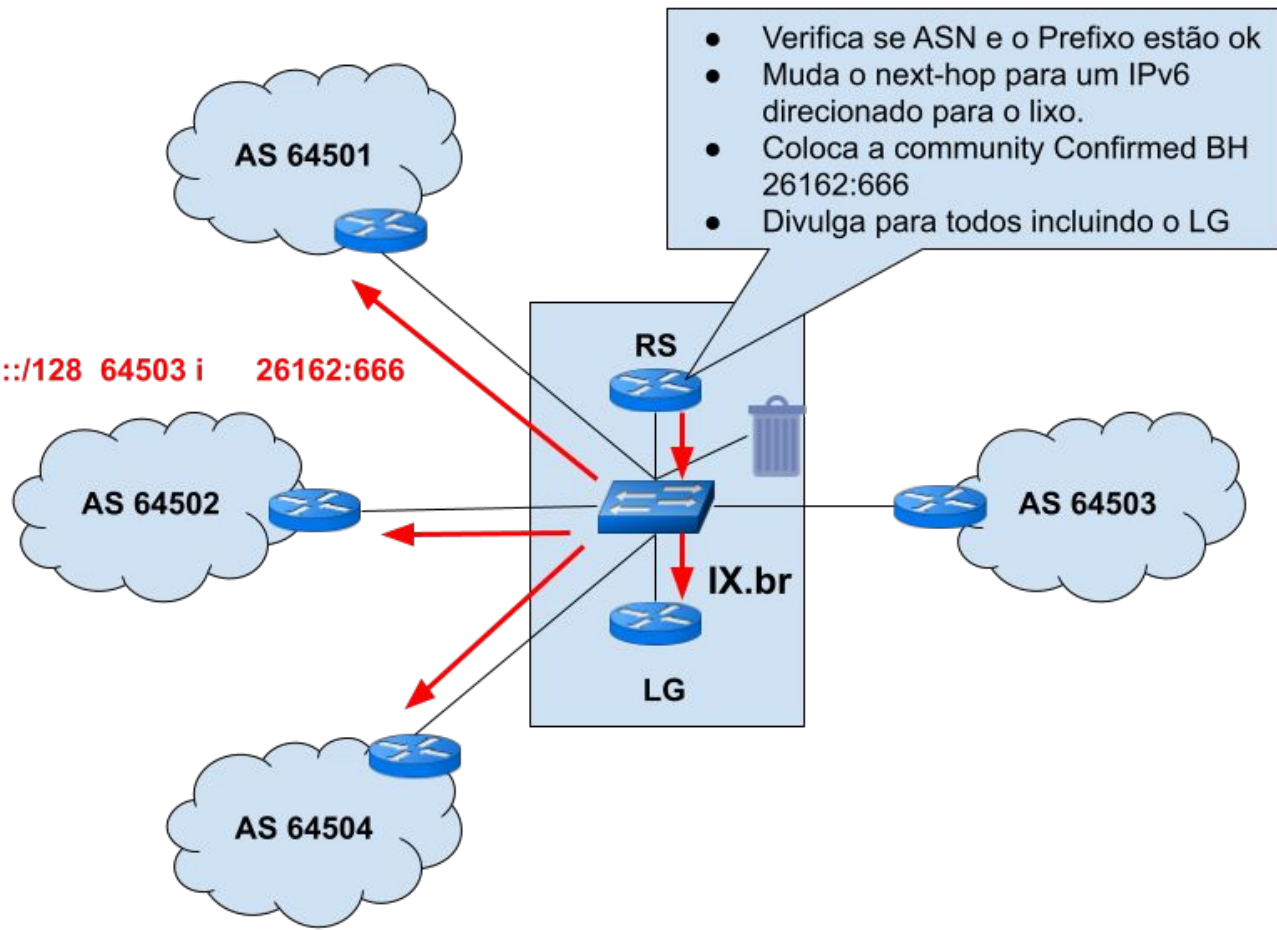
RTBH no IX.br

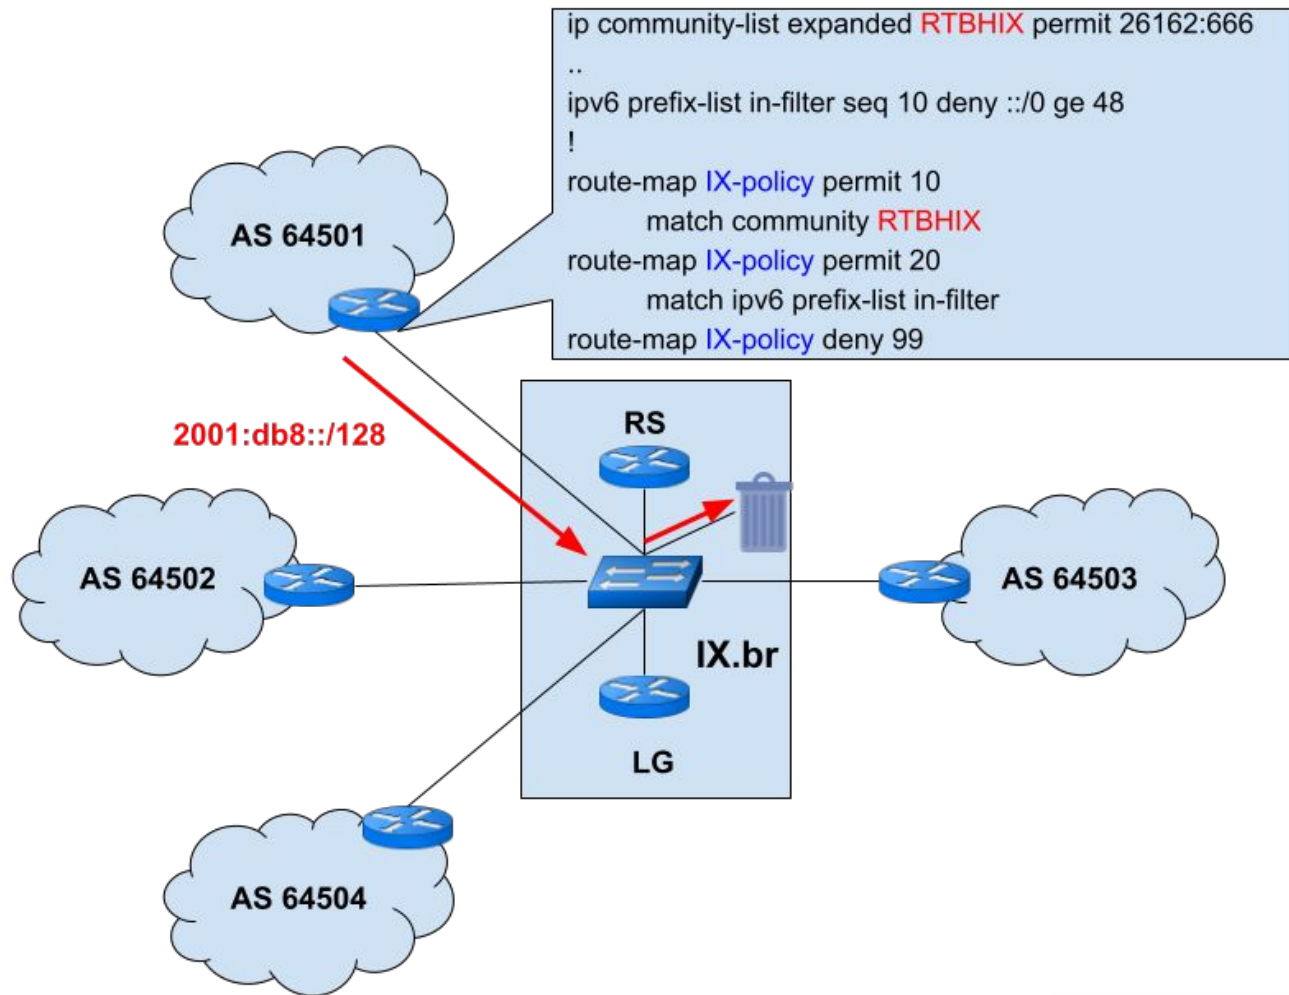
No IX.br

- Para se utilizar o Blackhole é preciso
 - Ser um Sistema Autônomo registrado no Registro.br(whois)
 - Usar rotas /32 em IPv4 ou /128 em IPv6
 - Ser o detentor da Rota (as-path com origem seu ASN)
 - Não pode ter intermediários (as-path não pode ter outro ASN além do seu)
 - Usar **65535:666** ou **(ro|rt):65535:666** ou **65535:666:0**



2001:db8::/128 64503 i 26162:666



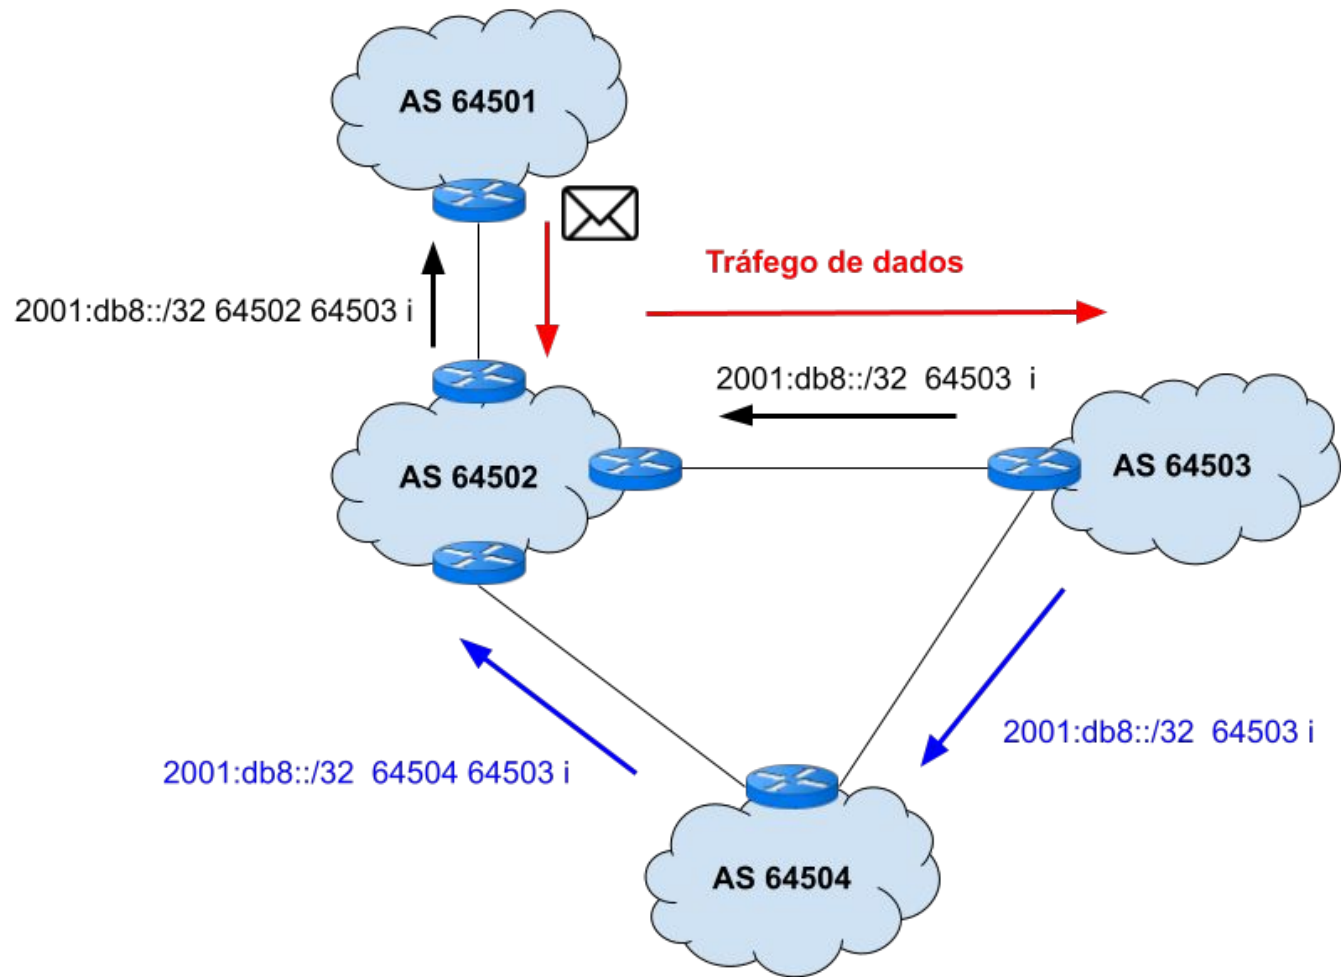


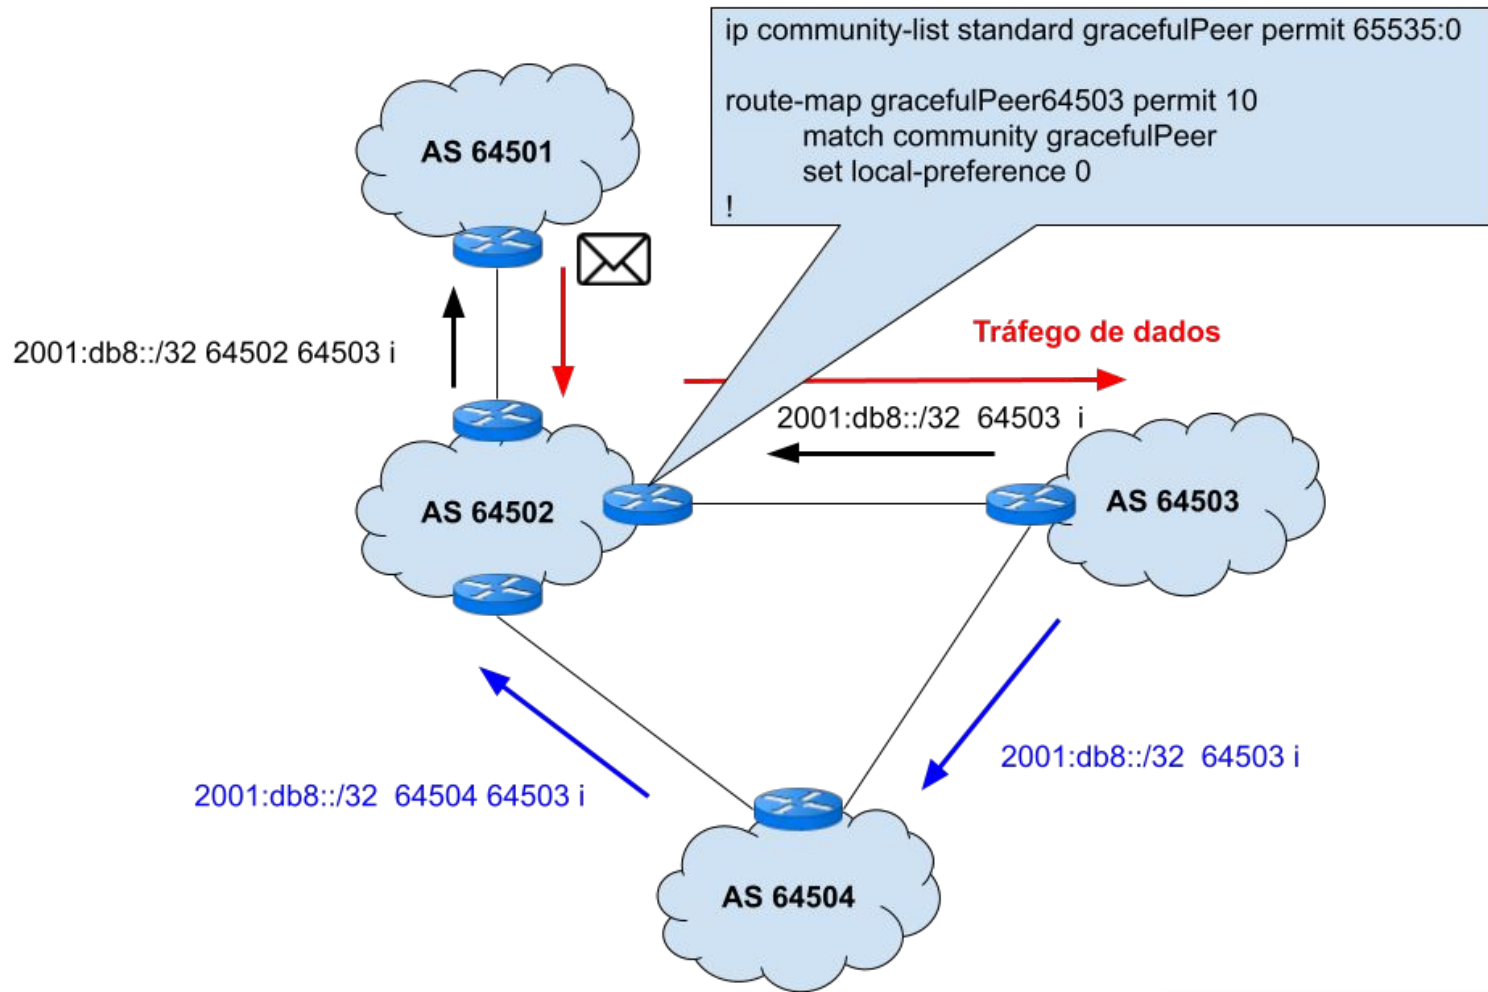
Caso 3

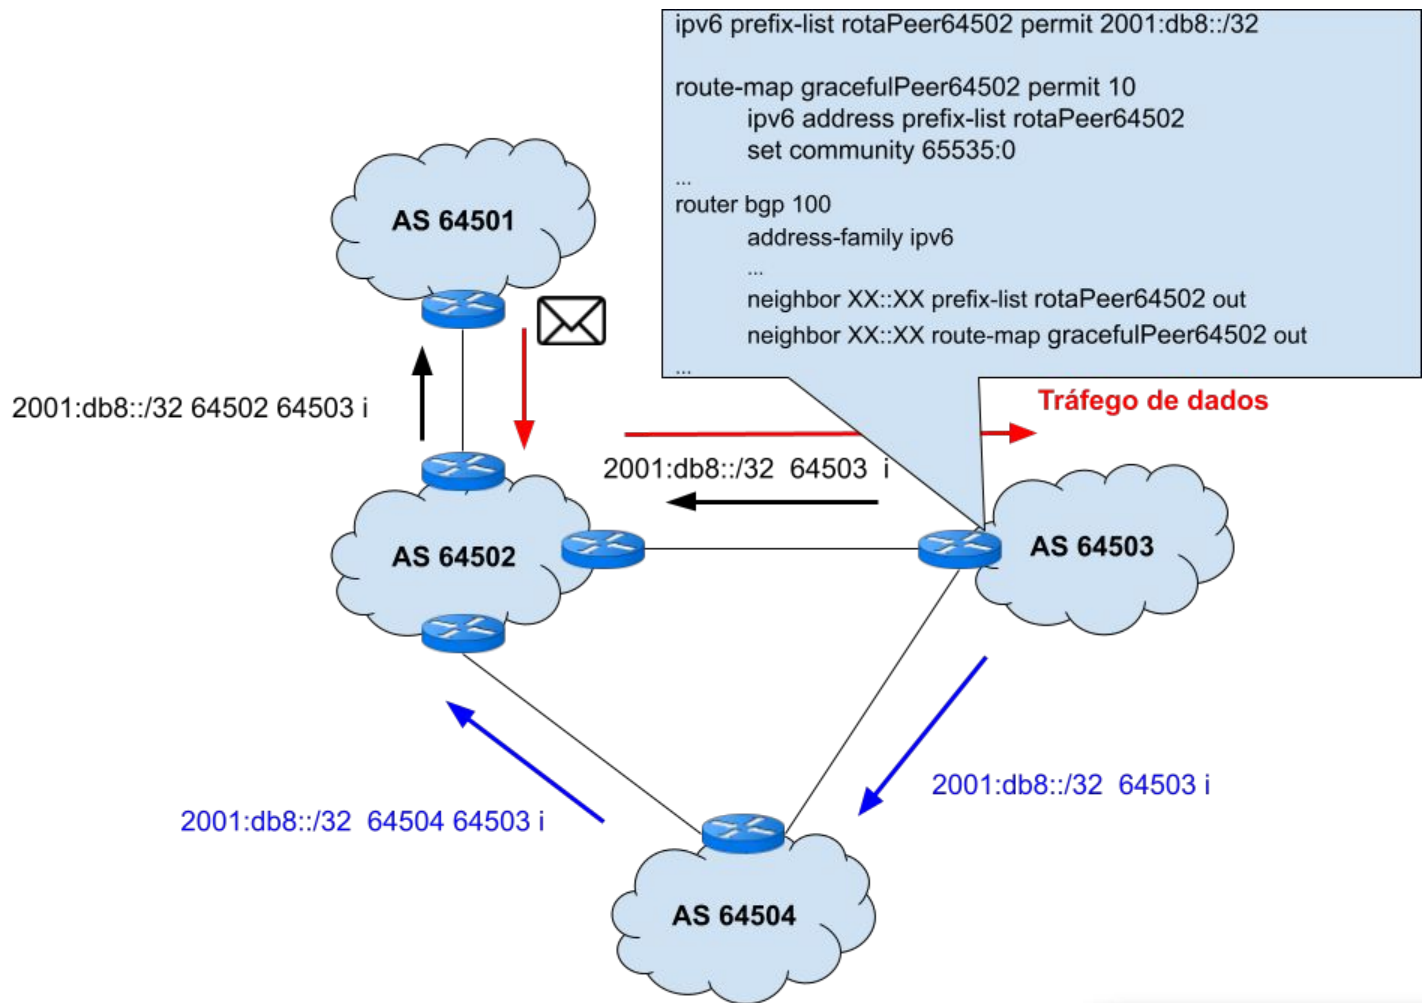
Graceful Shutdown

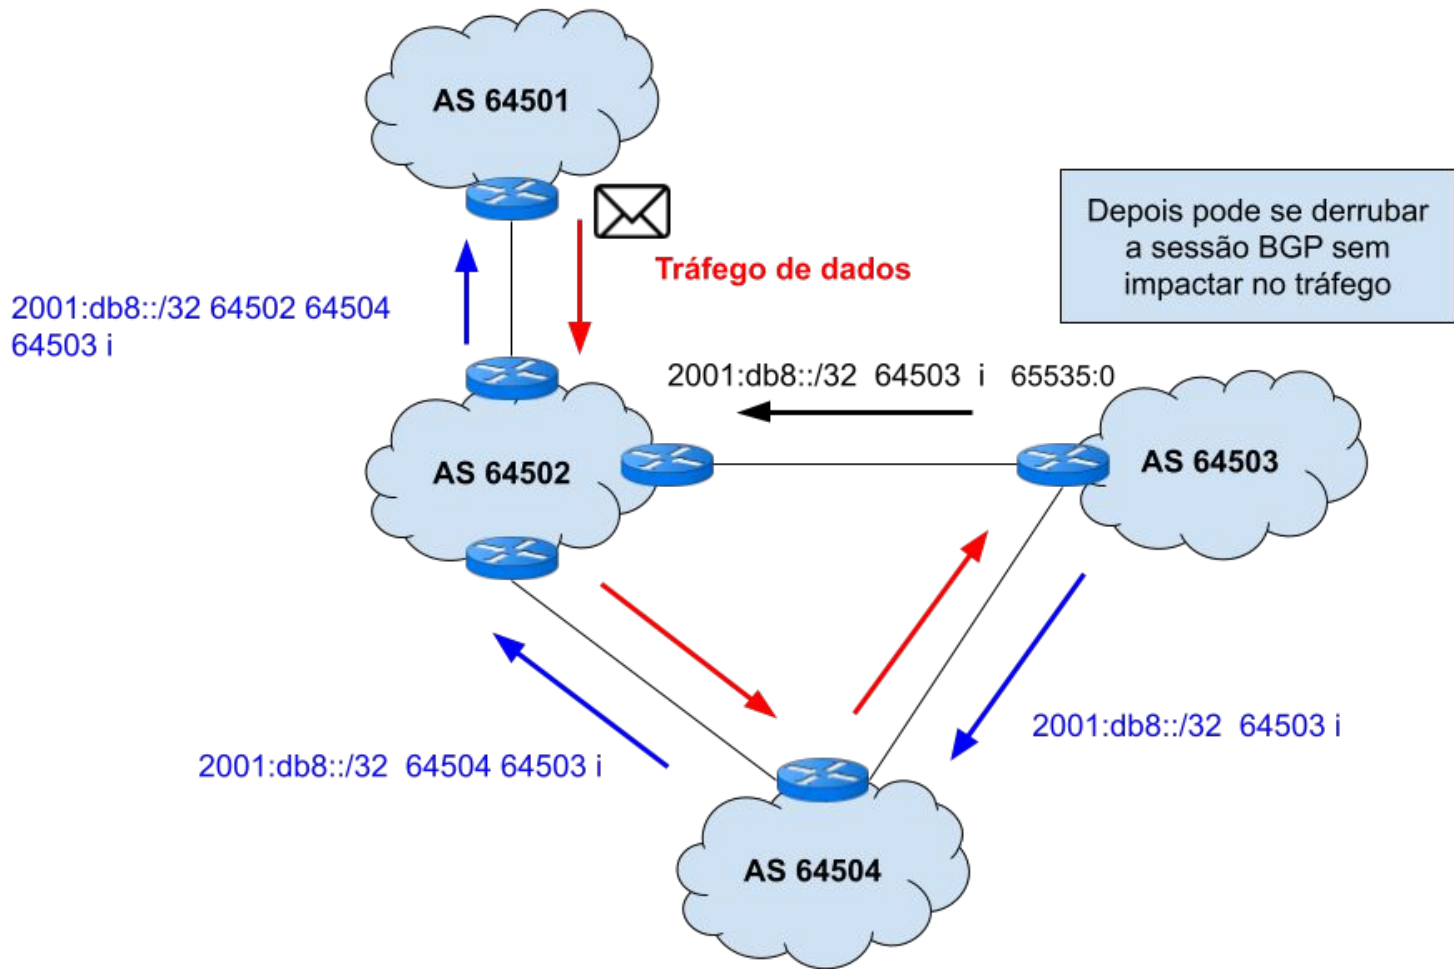
Graceful Shutdown

- Uma maneira de desligar a sessão BGP sem impactar o tráfego
 - Serve somente se tiver mais de um link
 - Se avisa o vizinho para ele preferir as rotas por outro caminho (zera o local-preference)
 - Depois que o tráfego mudar de sentido
 - Desliga-se a sessão BGP
- O IX.br implementa também.










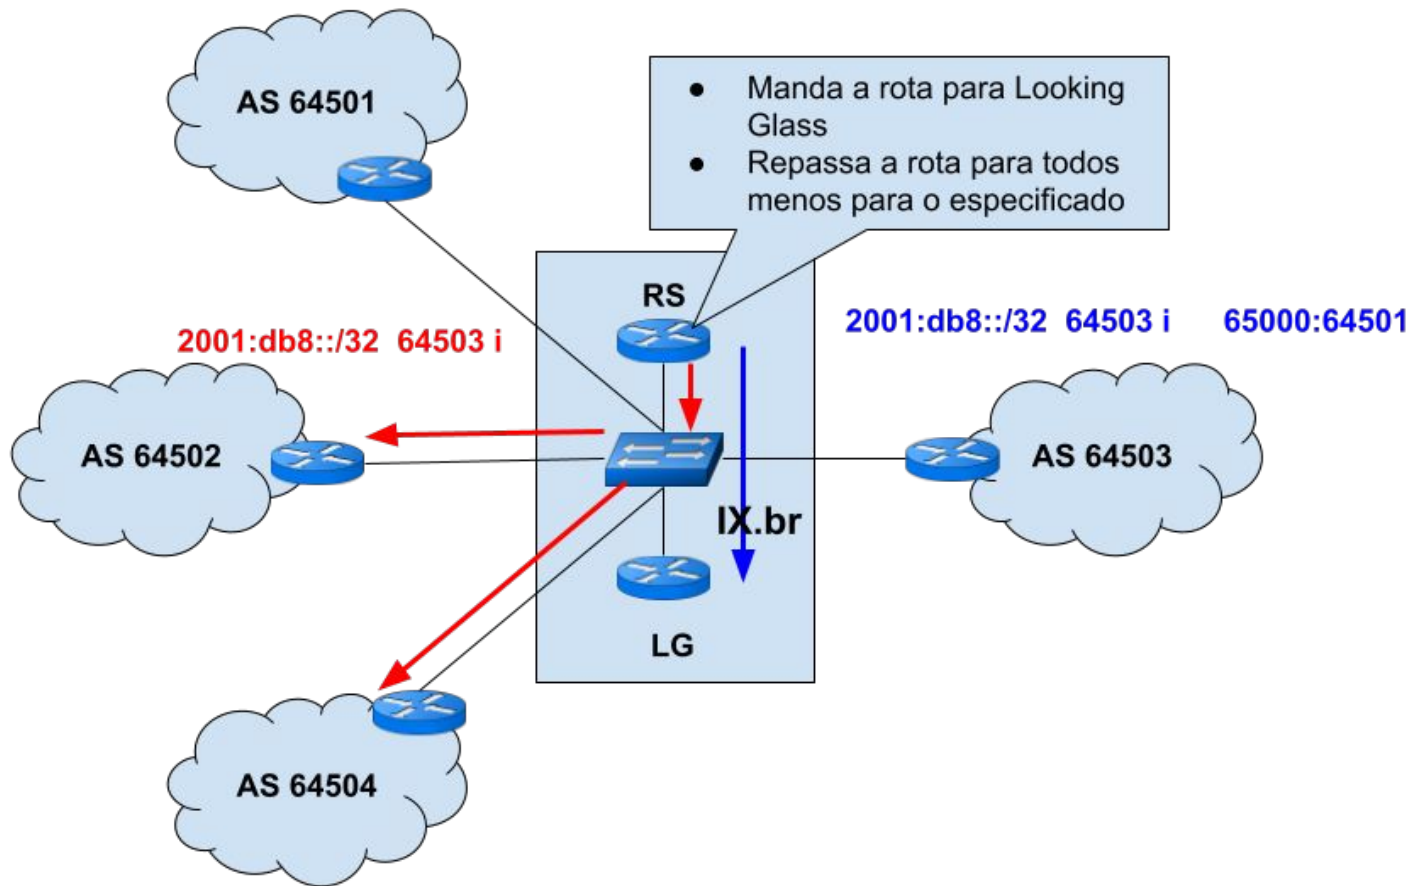
Caso 4

IX.br No Export

IX.br No Export

- Uma forma de não enviar as rotas para um determinado participante do IX.br
- Fica visível no Looking Glass



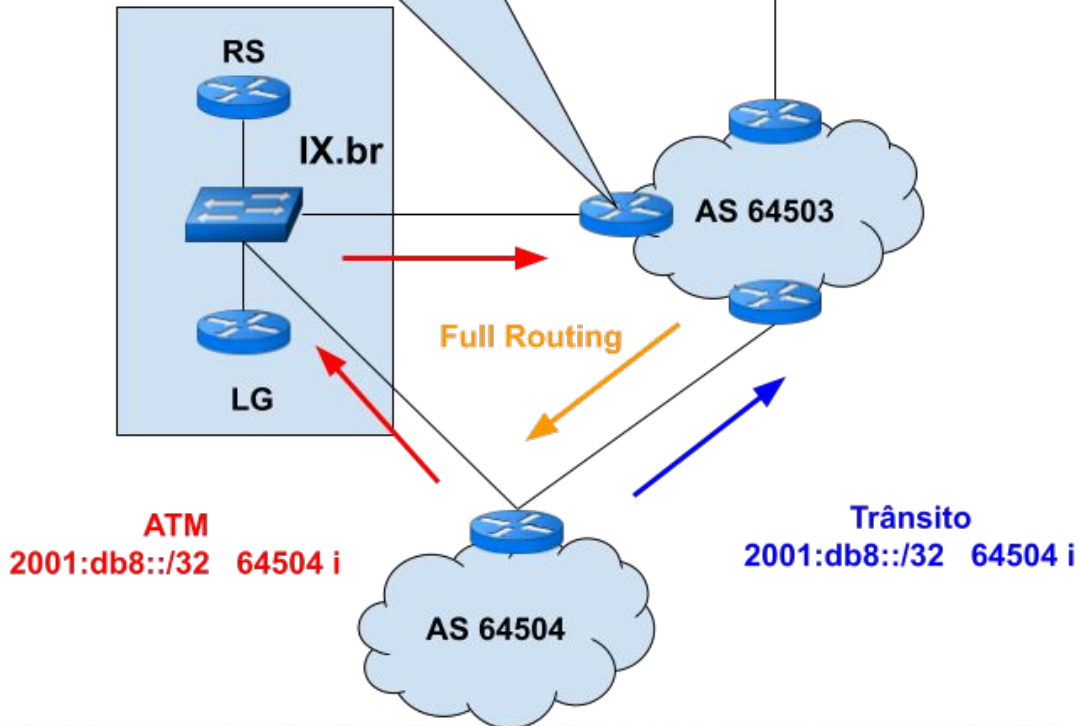


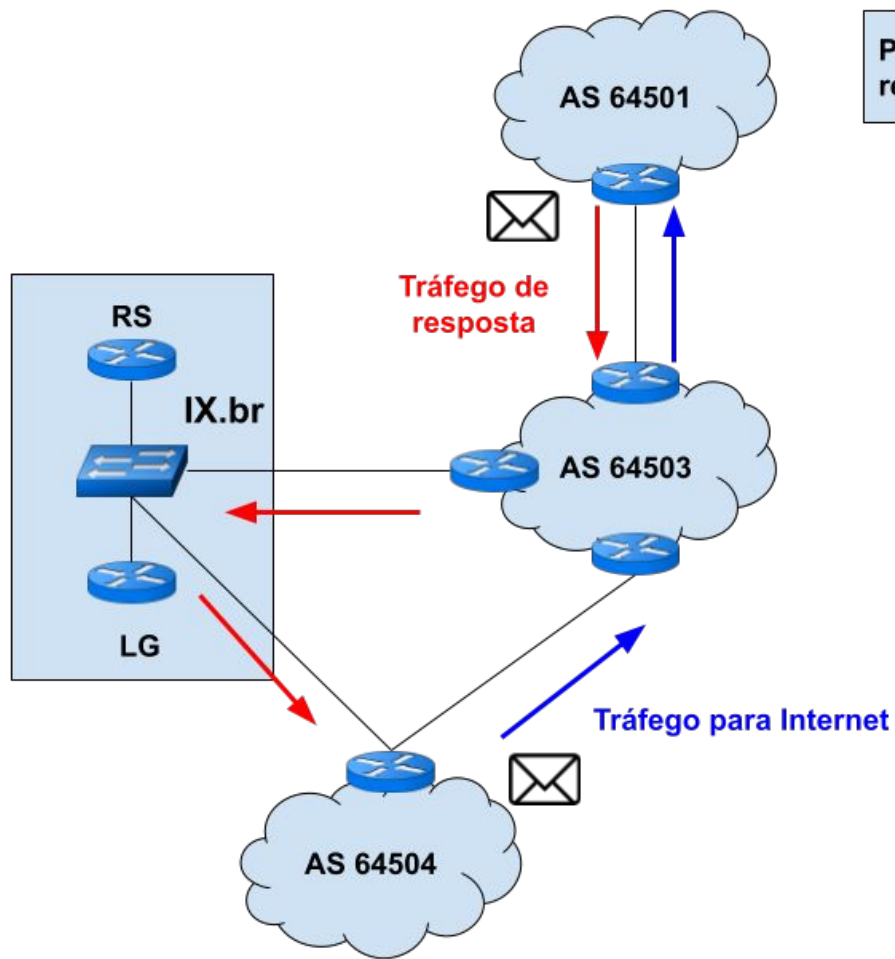
Caso 5

IX.br No Export - evitando Assimetria

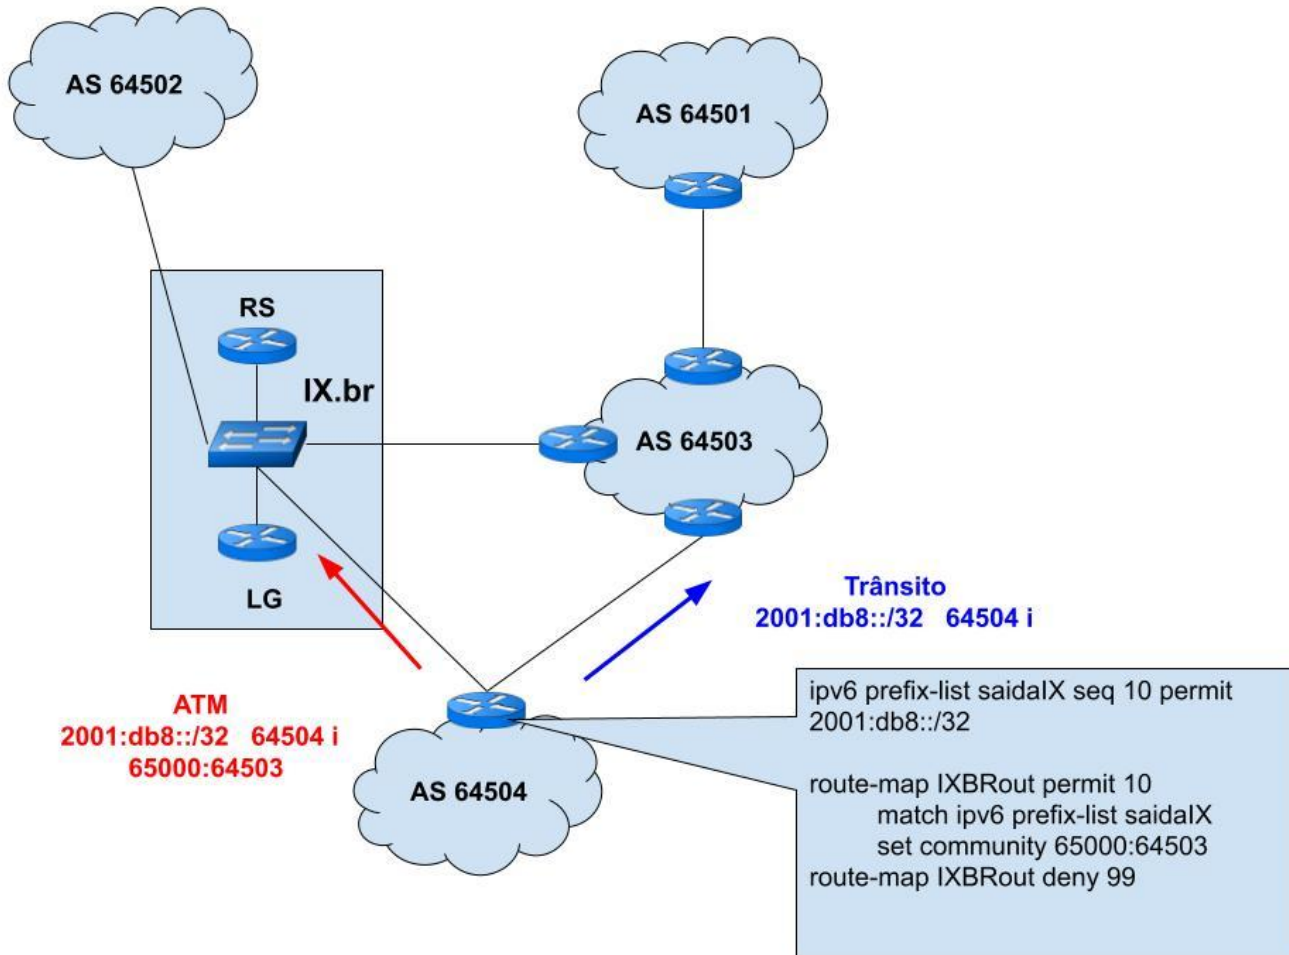
```
route-map entradaIXBR permit 10
set local-preference 150
route-map entradaIXBR deny 99
```

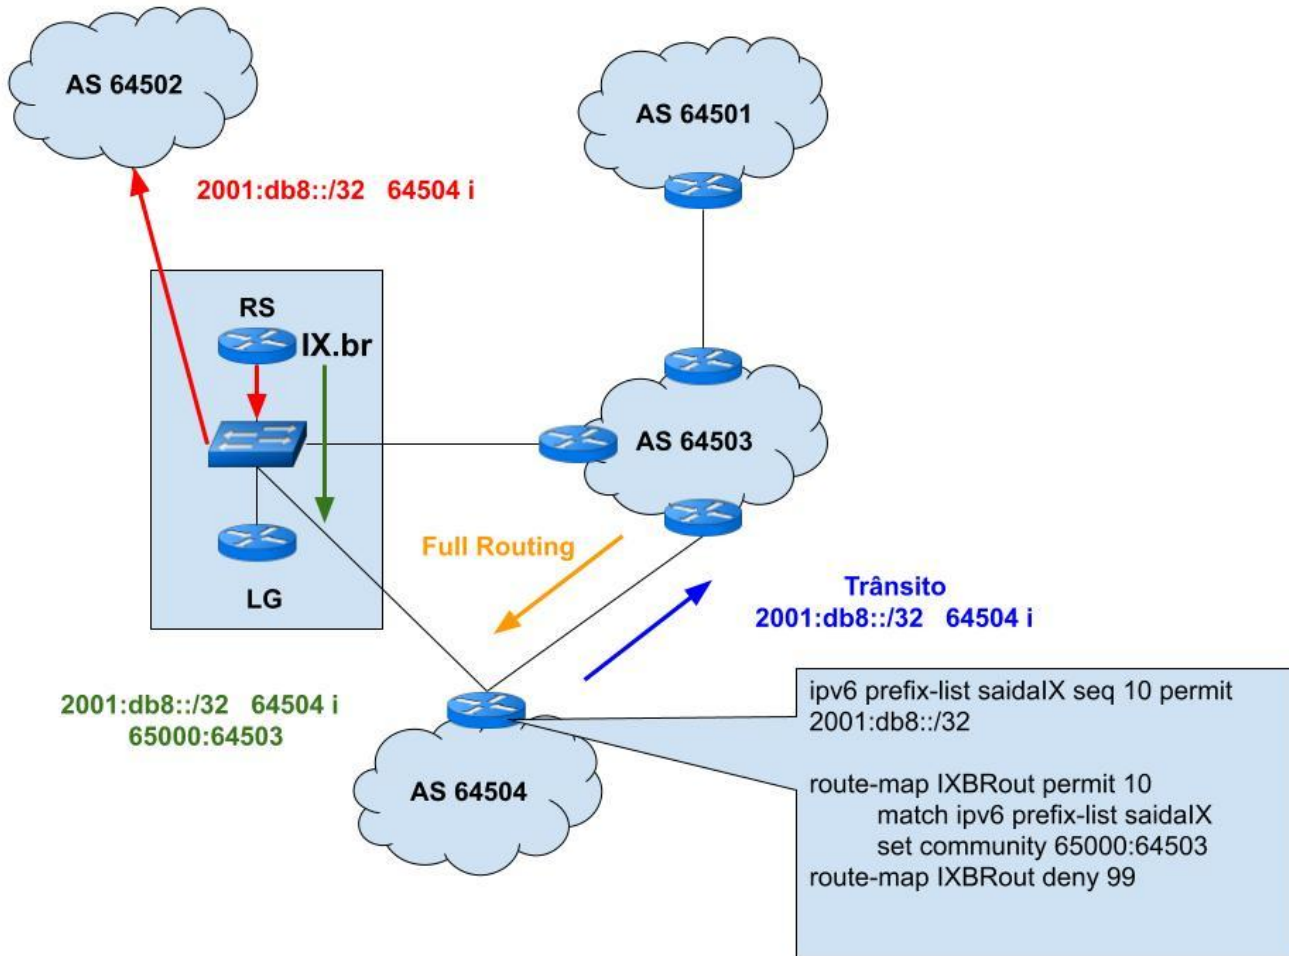
Ele prefere a rota do cliente
vinda por dentro do IX.br
mas divulga o full routing
pelo link de trânsito.

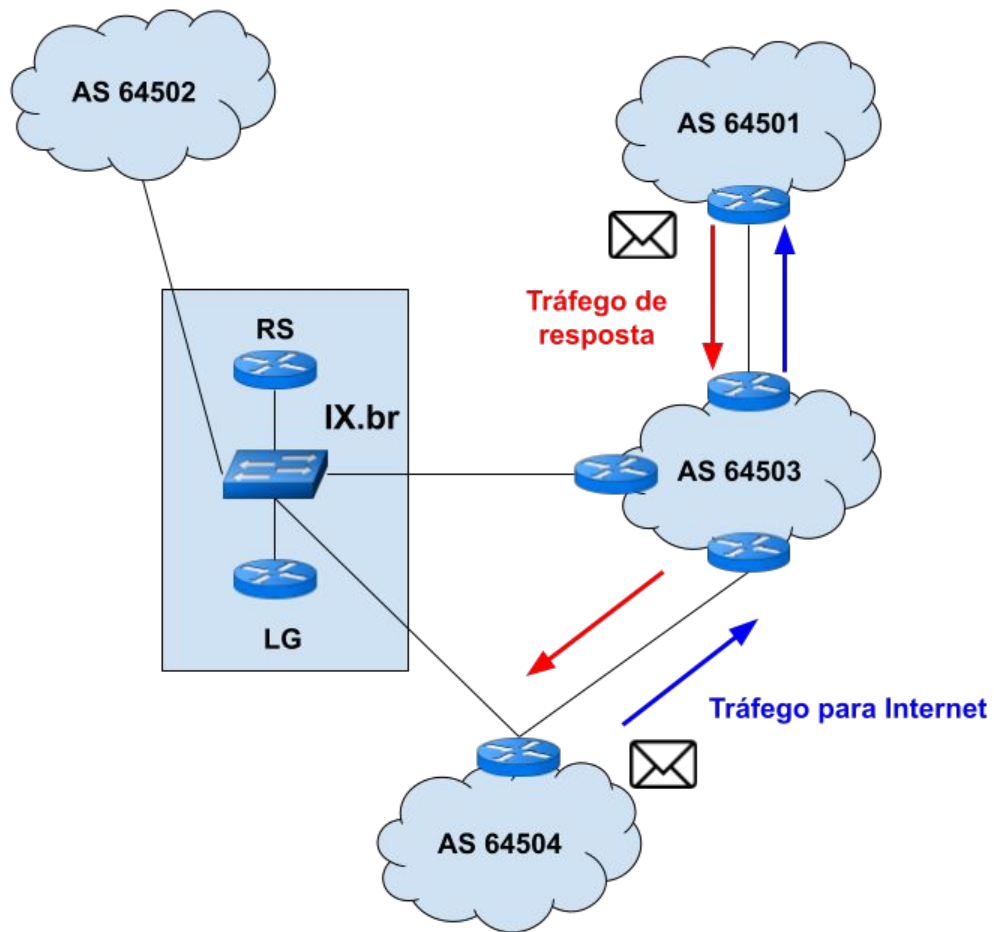




Pode causar assimetria na rede!!!





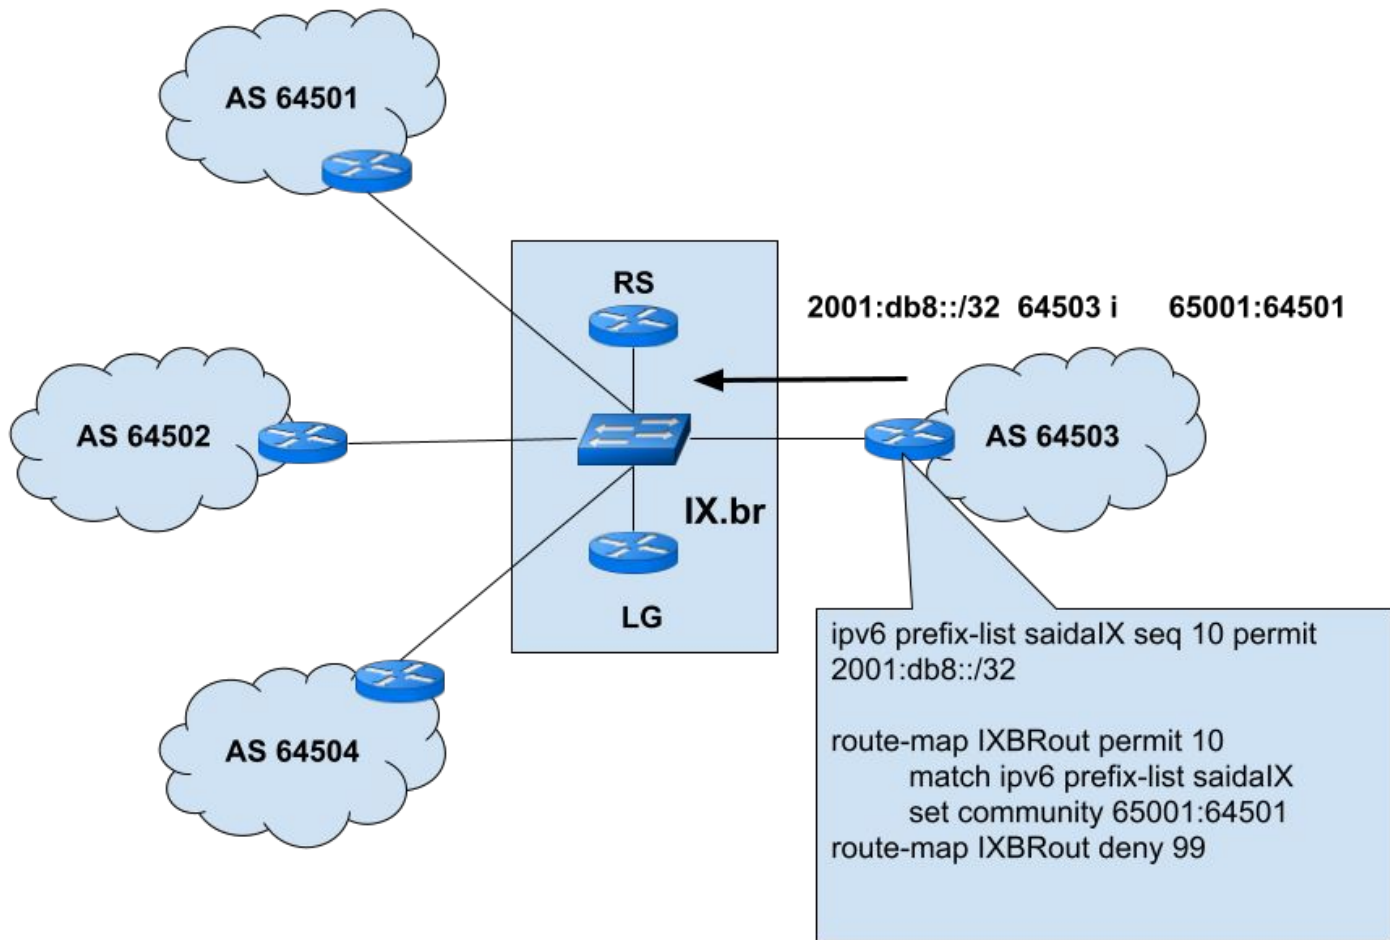


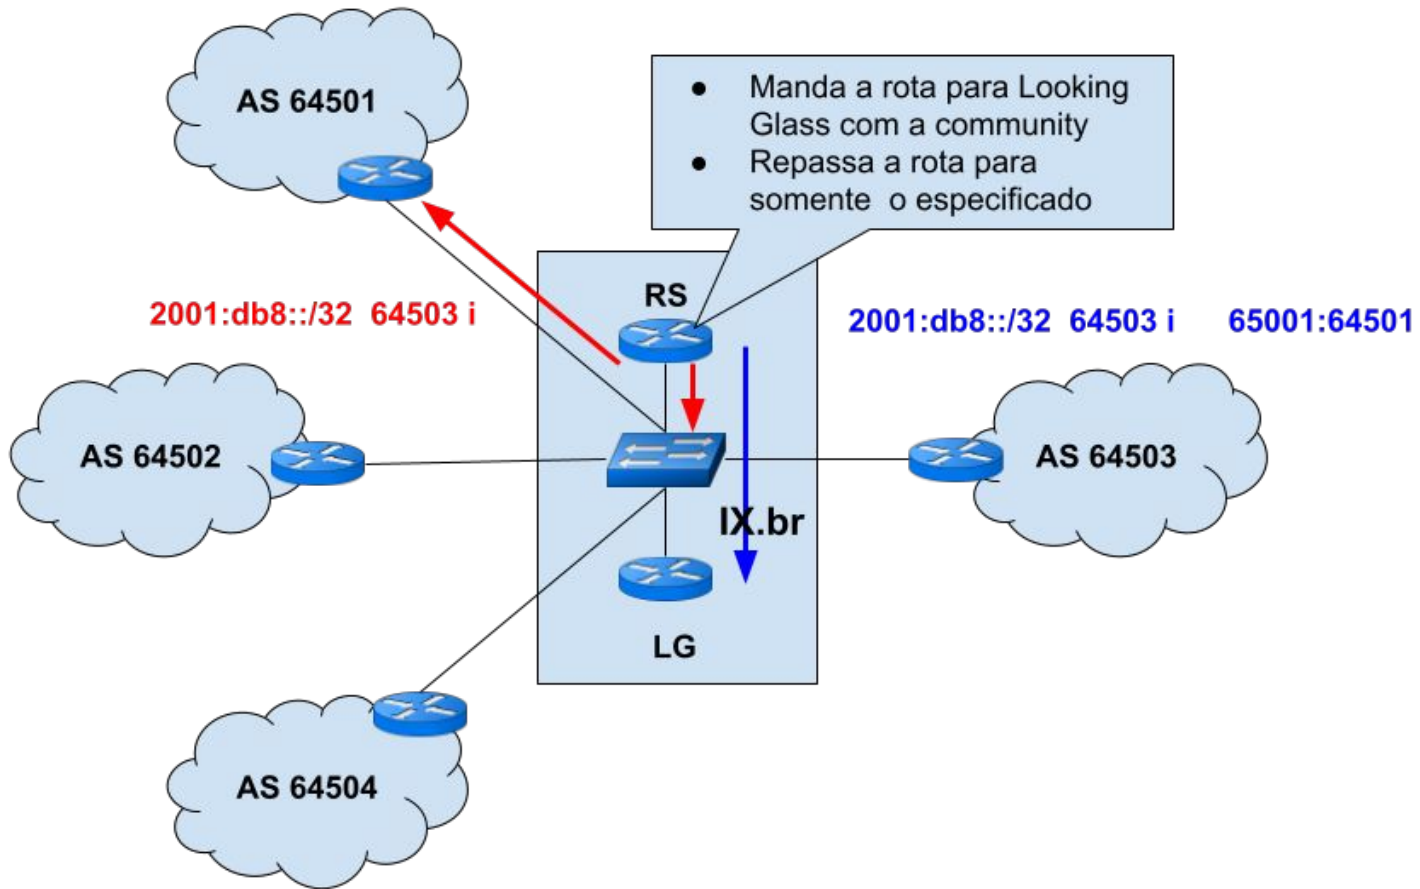
Caso 6

IX.br Only Export

IX.br Only Export

- Uma forma de enviar as rotas para um determinado participante do IX.br - Oposto do No export
- Às vezes pode ser mais interessante escolher a dedo para quem enviar as rotas do que remover da lista de todos os participantes quem não deve receber.
 - São Paulo tem muitos participantes.





Informações extras

- Muitas operadoras divulgam no seu próprio site
 - BPF compilou num documento
 - https://wiki.brasilpeeringforum.org/w/Lista_de_Communities_BGP
- Utilizem o Peering DB
 - O IX.br utiliza informações de lá - Never via Route Server <https://www.peeringdb.com/ix/171>
- Utilizem Looking Glass para entender as communities

Obrigado!!!

Tiago Jun Nakamura

@ cursosceptro@nic.br

@ ipv6@nic.br

nic.br **cgi.br**

www.nic.br | www.cgi.br