


# Gerando Rotas BGP

# BGP – Gerando rotas internas BGP



# Gerando rotas BGP

---

- Agregado
  - Aloca blocos por equipamento
  - Reduz número de rotas
  - Rota estável
- Estáticas
  - Caso não seja possível o agregado
  - Procedimento muito utilizado


# Configuração Básica iBGP

```
mexico
bgp {
  group iBGP {
    type internal;
    local-address 10.0.0.4;
    neighbor 10.0.0.5;
    neighbor 10.0.0.2;
  }
  routing-options {
    autonomous-system 65000;
  }
}
```

```
chile
bgp {
  group iBGP {
    type internal;
    local-address 10.0.0.5;
    neighbor 10.0.0.4;
    neighbor 10.0.0.2;
  }
  routing-options {
    autonomous-system 65000;
  }
}
```

```
brasil
bgp {
  group iBGP {
    type internal;
    local-address 10.0.0.2;
    neighbor 10.0.0.4;
    neighbor 10.0.0.5;
  }
  routing-options {
    autonomous-system 65000;
  }
}
```

# Injeção de um agregado no iBGP




```
routing-options {  
  /* rotas estáticas */  
  static {  
    route 192.168.1.0/24 next-hop 10.1.1.2;  
    route 192.168.2.0/24 next-hop 10.1.2.2;  
  }  
  /* rota agregada */  
  aggregate {  
    route 192.168.0.0/21;  
  }  
}
```

```
policy-statement distribui-agregado {  
  term agregado {  
    from protocol aggregate;  
    then accept;  
  }  
}  
bgp {  
  group iBGP {  
    type internal;  
    local-address 10.0.0.4;  
    export distribui-agregado;  
    neighbor 10.0.0.5;  
    neighbor 10.0.0.2;  
  }  
}
```

# Injeção de rotas individuais no iBGP

```
routing-options {  
  /* rotas estaticas */  
  static {  
    route 192.168.8.0/24 next-hop 10.1.3.2;  
    route 192.168.9.0/24 next-hop 10.1.4.2;  
  }  
}
```



```
policy-statement rotas-iBGP {  
  term rotas-especificas {  
    from {  
      protocol static;  
      route-filter 192.168.8.0/21 orlonger;  
    }  
    then {  
      next-hop self;  
      accept;  
    }  
  }  
}  
bgp {  
  group iBGP {  
    type internal;  
    local-address 10.0.0.2;  
    export rotas-iBGP;  
    neighbor 10.0.0.4;  
    neighbor 10.0.0.5;  
  }  
}
```

# Visualização das rotas

10.0.0.5



PONTO DE  
OBSERVAÇÃO

```
user@chile# run show route protocol bgp


inet.0: 16 destinations, 16 routes (16 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.0.0/21      *[BGP/170] 00:36:03, localpref 100, from 10.0.0.4
                   AS path: I
                   > to 10.0.5.4 via fe-0/0/1.0
192.168.8.0/24    *[BGP/170] 00:36:13, localpref 100, from 10.0.0.2
                   AS path: I
                   > to 10.0.2.2 via fe-0/0/2.0
192.168.9.0/24    *[BGP/170] 00:36:13, localpref 100, from 10.0.0.2
                   AS path: I
                   > to 10.0.2.2 via fe-0/0/2.0
```

# Assinante Single Homed




# Assinante Single Homed



# Configuração eBGP

```
chile
  bgp {
    group eBGP-argentina {
      type external;
      export anuncio-para-inet;
      peer-as 65001;
      neighbor 10.0.6.3;
    }
  }
  policy-statement anuncio-para-inet {
    term bloco-local {
      from {
        route-filter 192.168.0.0/20 exact;
      }
      then accept;
    }
    term rejeita {
      then reject;
    }
  }
}
```

```
argentina
  bgp {
    group eBGP-chile {
      type external;
      export simula-inet;
      peer-as 65000;
      neighbor 10.0.6.5;
    }
  }
  policy-statement simula-inet {
    term distribui-estatica {
      from protocol static;
      then accept;
    }
  }
  routing-options {
    autonomous-system 65001;
  }
}
```



# Resultado no AS 65000

```
user@chile# run show route protocol bgp terse
```

```
inet.0: 24 destinations, 24 routes (24 active, 0 holddown, 0 hidden)
```

```
+ = Active Route, - = Last Active, * = Both
```

A Destination	P Prf	Metric 1	Metric 2	Next hop	AS path
* 11.0.1.0/24	B 170	100	0	>10.0.6.3	65001 I
* 11.0.2.0/24	B 170	100	0	>10.0.6.3	65001 I
* 11.0.3.0/24	B 170	100	0	>10.0.6.3	65001 I
* 11.0.4.0/24	B 170	100	0	>10.0.6.3	65001 I
* 11.0.5.0/24	B 170	100	0	>10.0.6.3	65001 I
* 11.0.6.0/24	B 170	100	0	>10.0.6.3	65001 I
* 11.0.7.0/24	B 170	100	0	>10.0.6.3	65001 I
* 192.168.0.0/21	B 170	100		>10.0.5.4	I
* 192.168.8.0/24	B 170	100		>10.0.2.2	I
* 192.168.9.0/24	B 170	100		>10.0.2.2	I

```
user@mexico# run show route protocol bgp terse
```

```
inet.0: 27 destinations, 27 routes (20 active, 0 holddown, 7 hidden)
```


```
+ = Active Route, - = Last Active, * = Both
```

A Destination	P Prf	Metric 1	Metric 2	Next hop	AS path
* 192.168.8.0/24	B 170	100		>10.0.3.2	I
* 192.168.9.0/24	B 170	100		>10.0.3.2	I

???

# Next-hop Self

```
bgp {  
  group iBGP {  
    type internal;  
    export next-hop-self;  
    neighbor 10.0.0.4;  
  }  
}  
policy-statement next-hop-self {  
  term muda-next-hop {  
    from protocol bgp;  
    then {  
      next-hop self;  
    }  
  }  
}
```



```
user@mexico# run show route protocol bgp terse
```

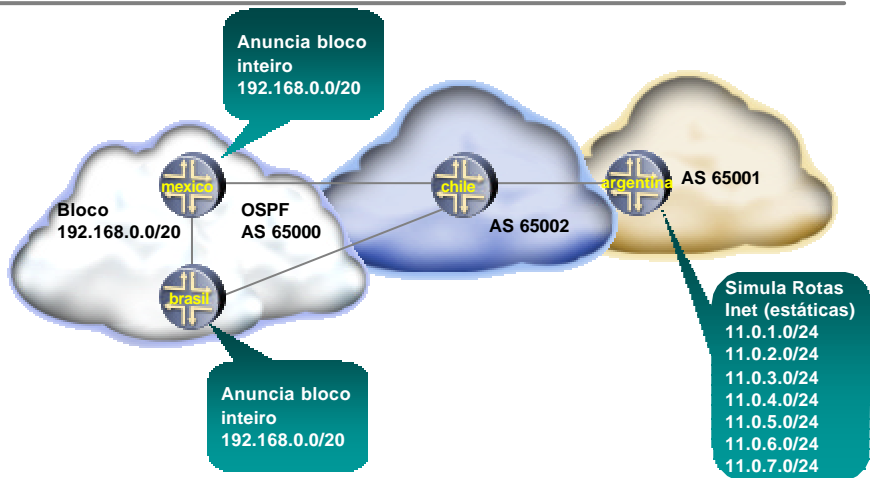
```
inet.0: 27 destinations, 27 routes (27 active, 0 holddown, 0 hidden)  
+ = Active Route, - = Last Active, * = Both
```

A Destination	P Prf	Metric 1	Metric 2	Next hop	AS path
* 11.0.1.0/24	B 170	100	0	>10.0.5.5	65001 I
* 11.0.2.0/24	B 170	100	0	>10.0.5.5	65001 I

```
...
```

# Assinante Multi Homed

# Assinante Multi Homed

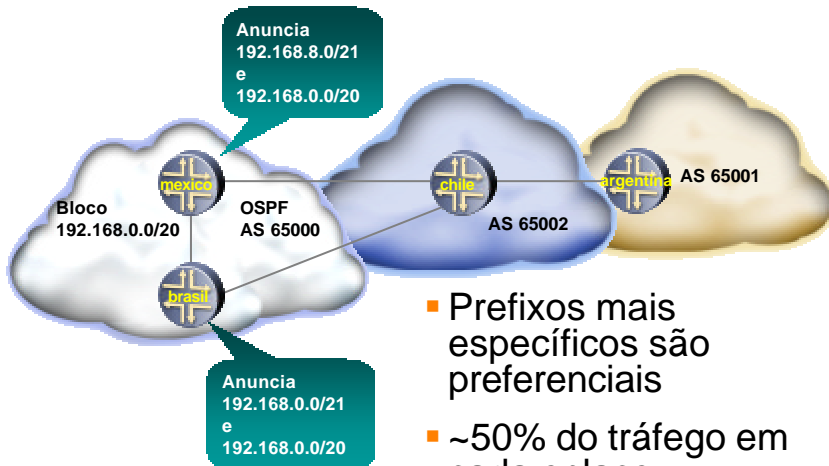


# Sem balanceamento de carga

```
user@chile# run show route 192.168/16 extensive

inet.0: 17 destinations, 18 routes (17 active, 0 holddown, 0 hidden)
192.168.0.0/20 (2 entries, 1 announced)
Path 192.168.0.0 from 10.0.2.2 Vector len 4. Val: 0 1
  *BGP      Preference: 170/-101
            Source: 10.0.2.2
            Next hop: 10.0.2.2 via fe-0/0/2.0, selected
            State: <Active Ext>
            Local AS: 65002 Peer AS: 65000
            Age: 1:15      Metric: 0
            Task: BGP_65000.10.0.2.2+4812
            Announcement bits (2): 0-KRT 1-BGP.0.0.0.0+179
            AS path: 65000 I
            Localpref: 100
            Router ID: 10.0.0.2
  BGP      Preference: 170/-101
            Source: 10.0.5.4
            Next hop: 10.0.5.4 via fe-0/0/1.0, selected
            State: <NotBest Ext>
            Inactive reason: Router ID
            Local AS: 65002 Peer AS: 65000
            Age: 1:03      Metric: 0
            Task: BGP_65000.10.0.5.4+179
            AS path: 65000 I
            Localpref: 100
            Router ID: 10.0.0.4
```

# Uma Alternativa de Balanceamento de Carga



- Prefixos mais específicos são preferenciais
- ~50% do tráfego em cada enlace
- Redundância mantida



# Com balanceamento de carga

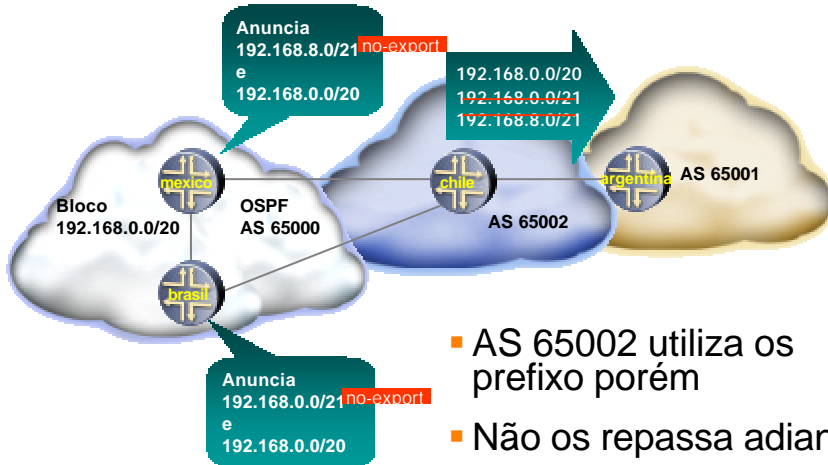
```
user@chile# run show route 192.168/16

inet.0: 19 destinations, 20 routes (19 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.0.0/20      *[BGP/170] 00:25:05, MED 0, localpref 100
                   AS path: 65000 I
                   > to 10.0.2.2 via fe-0/0/2.0
                   [BGP/170] 00:24:53, MED 0, localpref 100
                   AS path: 65000 I
                   > to 10.0.5.4 via fe-0/0/1.0
192.168.0.0/21     *[BGP/170] 00:01:23, MED 0, localpref 100
                   AS path: 65000 I
                   > to 10.0.2.2 via fe-0/0/2.0
192.168.8.0/21    *[BGP/170] 00:00:21, MED 0, localpref 100
                   AS path: 65000 I
                   > to 10.0.5.4 via fe-0/0/1.0
```

- Faz sentido para o upstream provider
- Não faz sentido para ASs além do upstream
  - Gera muitas rotas desnecessárias

# Community NO\_EXPORT



- AS 65002 utiliza os prefixo porém
- Não os repassa adiante

# Configuração no-export

```
policy-statement anuncio-inet {
  term bloco-local {
    from {
      route-filter 192.168.0.0/20 exact;
    }
    then accept;
  }
  term especifico {
    from {
      route-filter 192.168.0.0/21 exact;
    }
    then {
      community set bloqueia-especifica;
      accept;
    }
  }
  term rejeita-resto {
    then reject;
  }
}
community bloqueia-especifica members no-export;
```

```
bgp {
  group eBGP-chile {
    type external;
    export anuncio-inet;
    peer-as 65002;
    neighbor 10.0.2.5;
  }
}
```



Anuncia  
192.168.0.0/21 **no-export**  
e  
192.168.0.0/20

# Rota com no-export


```
user@chile# run show route 192/8 extensive community no-export

inet.0: 19 destinations, 20 routes (19 active, 0 holddown, 0 hidden)
192.168.0.0/21 (1 entry, 1 announced)
TSI:
KRT in-kernel 192.168.0.0/21 -> {10.0.2.2}
    *BGP      Preference: 170/-101
              Source: 10.0.2.2
              Next hop: 10.0.2.2 via fxp2.0, selected
              State: <Active Ext>
              Local AS: 65002 Peer AS: 65000
              Age: 3:50      Metric: 0
              Task: BGP_65000.10.0.2.2+4812
              Announcement bits (1): 0-KRT
              AS path: 65000 I
              Communities: no-export
              Localpref: 100
              Router ID: 10.0.0.2


192.168.8.0/21 (1 entry, 1 announced)
...
```

# Multi Homed com Múltiplos Provedores

# Multihomed com múltiplos provedores



# Função de Trânsito




```
user@mexico# run show route advertising-protocol bgp 10.0.5.5 (CHILE)

inet.0: 28 destinations, 35 routes (28 active, 0 holddown, 0 hidden)
Prefix                               Nexthop                               MED    Lc1pref  AS path
11.0.1.0/24                          Self                                   0      0        65001 I
11.0.2.0/24                          Self                                   0      0        65001 I
11.0.3.0/24                          Self                                   0      0        65001 I
11.0.4.0/24                          Self                                   0      0        65001 I
11.0.5.0/24                          Self                                   0      0        65001 I
11.0.6.0/24                          Self                                   0      0        65001 I
11.0.7.0/24                          Self                                   0      0        65001 I
192.168.0.0/20                       Self                                   0      0        I
```

# Bloqueando trânsito com communities

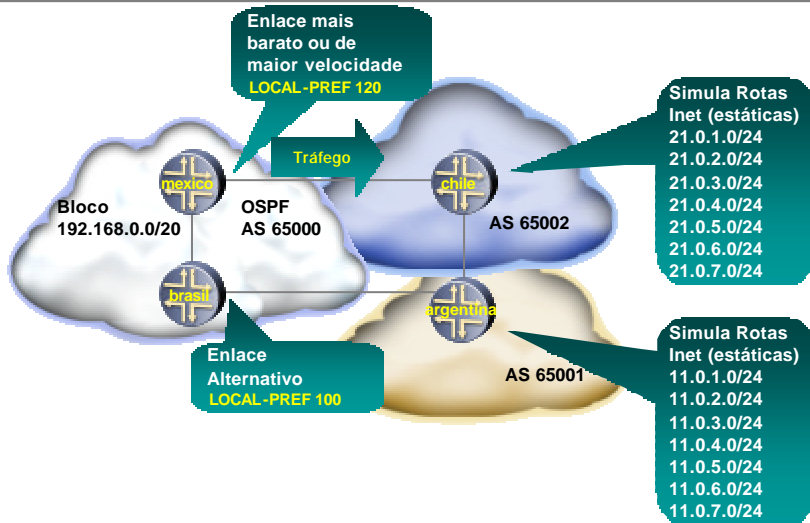
```
brasil
bgp {
  group eBGP-argentina {
    type external;
    import recebe-inet;
    export anuncio-inet;
    peer-as 65001;
    neighbor 10.0.1.3;
  }
}
policy-statement recebe-inet {
  term unico {
    then {
      community add transito;
    }
  }
}
community transito members 65000:120;
```

```
mexico
bgp {
  group eBGP-chile {
    import recebe-inet;
    export anuncio-inet;
    peer-as 65002;
  }
}
policy-statement anuncio-inet {
  term bloco-local {
    from {
      route-filter 192.168.0.0/20 exact;
    }
    then accept;
  }
  term bloqueia-transito {
    from community transito;
    then reject;
  }
}
community transito members 65000:120;
```





# Saída preferencial – Local Pref



# Impacto do Local-Pref

---

```
user@brasil# run show route protocol bgp

inet.0: 26 destinations, 40 routes (26 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

11.0.1.0/24      *[BGP/170] 00:05:44, localpref 120, from 10.0.0.4
                 AS path: 65002 65001 I
                 > to 10.0.3.4 via fe-0/0/3.0
                 [BGP/170] 02:04:45, MED 0, localpref 100
                 AS path: 65001 I
                 > to 10.0.1.3 via fe-0/0/1.0
...
21.0.1.0/24      *[BGP/170] 00:05:44, MED 0, localpref 120, from 10.0.0.4
                 AS path: 65002 I
                 > to 10.0.3.4 via fe-0/0/3.0
                 [BGP/170] 02:04:45, localpref 100
                 AS path: 65001 65002 I
                 > to 10.0.1.3 via fe-0/0/1.0
...
```

Obrigado