

Traffic engineering on a multihoming environment

Fernando García
fernando.garcia@tecnocom.es

Tecnocom

The logo for Tecnocom, featuring the word "Tecnocom" in a bold, blue, sans-serif font. Below the text is a stylized orange and red swoosh that underlines the name.

What's this about

- Solution to a request from a real customer
- Real solution in a real scenario

The customer...

- Two E3 lines to different carriers (A & B)
- He wants redundancy
- But he wants to get optimal usage of the lines
 - Optimal: both lines up to maximum capacity

Me...

- Must try to be a good netcitizen:
 - DON'T flap routes
 - DON'T deaggregate
- Let's try




The
customer lets
us play with
his network!!!
aka: real life test


Tecnocom

Real traffic

A



B



Best path in eBGP


- Highest weight
- Highest Local Pref
- Network/redistributed over local aggregate
- Shorter AS-PATH
- Lowest origin type
- Lowest MED
- eBGP over iBGP
- Lowest IGP metric
- Older and random alike

Input/Output

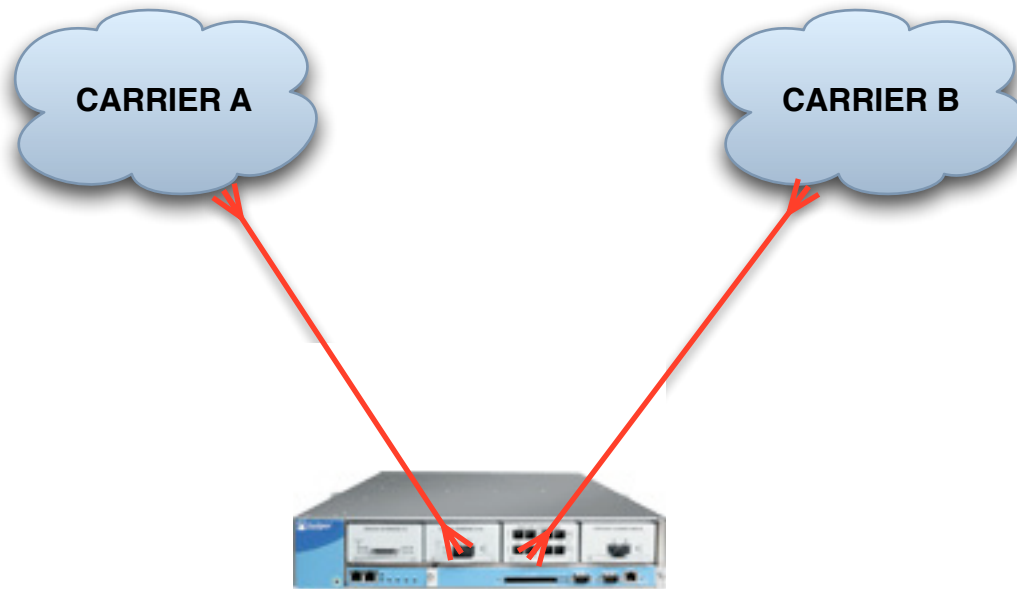
- Not always the same
- Output easier to “influence”
- Sometimes different policies and/or in each direction
- We’ll focus on input traffic

Measuring method

I. Stable situation




Measuring method




1. Stable situation
2. Reset circuit with A

Measuring method




1. Stable situation
2. Reset circuit with A

Measuring method




1. Stable situation
2. Reset circuit with A
3. Recover circuit with A

Measuring method




1. Stable situation
2. Reset circuit with A
3. Recover circuit with A

Measuring method




1. Stable situation
2. Reset circuit with A
3. Recover circuit with A
4. Measure

Measuring method




1. Stable situation
2. Reset circuit with A
3. Recover circuit with A
4. Measure
5. Reset circuit with B

Measuring method




1. Stable situation
2. Reset circuit with A
3. Recover circuit with A
4. Measure
5. Reset circuit with B

Measuring method




1. Stable situation
2. Reset circuit with A
3. Recover circuit with A
4. Measure
5. Reset circuit with B
6. Recover circuit with B

Measuring method




1. Stable situation
2. Reset circuit with A
3. Recover circuit with A
4. Measure
5. Reset circuit with B
6. Recover circuit with B

Measuring method




1. Stable situation
2. Reset circuit with A
3. Recover circuit with A
4. Measure
5. Reset circuit with B
6. Recover circuit with B
7. Measure

Starting point



After resetting
conection "A"



After resetting
conection "B"


Solution #1: Prepend

- Add your prefix several times
- Router A:


```
policy-statement prepend-carrier {  
  term 1 {  
    from {  
      route-filter 192.0.2.0/24 orlonger;  
    }  
    then as-path-prepend "65533";  
  }  
}
```

- Too much influence


Solution #1: Prepends



Original



After resetting
conection "A"



After resetting
conection "B"

Solution #2: Deaggregate


- Using more than one prefix
 - Several prefixes assigned by the RIR/LIR
 - Deaggregated prefixes
- Announce in each path with different prepend (or other attribute)
- Complex to apply and... not good netizenship

Solution #2: Deaggregate


- Router A

```
policy-statement prepend-carrier {  
  term 1 {  
    from {  
      route-filter 192.0.2.0/25 orlonger;  
    }  
    then as-path-prepend "65533";  
  }  
  term 2 {  
    from {  
      route-filter 192.0.2.128/25 orlonger;  
    }  
    then as-path-prepend "65533 65533";  
  }  
}
```



Solution #2: Deaggregate



Original



After resetting
conection "A"



After resetting
conection "B"


Solution #3: Communities

- Depend on carrier policies
- Usually allow to do a prepend on the carrier borders
- With my customer, the announced communities of the carrier didn't work


Solution #4: Origin

- Origin: Internal, external, incomplete
- Only checked when AS paths are equal
 - This happens frequently
- Set the origin of announces to one peer with “internal”
- Set the origin of announces to the other peer to “incomplete”


Solution #4: Origin



Original



After reseting
conection "A"



After reseting
conection "B"

Our solution

- Tuning:
 - Set the origin
- Finer tuning:
 - Selectively apply origin to some prefixes
 - Not deaggregated, but disjunct prefixes assigned by RIPE

General solution

- Your mileage may vary
- Gross adjust:
 - Prepend
- Fine tuning:
 - Set the origin
- Finer tuning:
 - Selectively apply prepend and/or origin to some prefixes

Questions?

