

Inter-AS Cost

draft-van-beijnum-idr-iac-00

idr @ IETF-71, March 12, 2008
Φιλαδέλφεια, PA, US

Iljitsch van Beijnum

The problem

- No good inter-AS metric in BGP!
- So no easy inbound traffic engineering
 - local pref: local
 - MED: only survives one hop
 - origin: not updated in transit (or usable)
- Only the AS path length

AS path length

- Increases with every hop
 - this is what we need
- But: AS hierarchy too flat
 - prepending once can move most of your traffic to another link

Inter-AS Cost

- "IAC"
- Conceptually:
 - a metric that is increased at every AS hop
- But:
 - need to be compatible with existing implementations

Solution

- $IAC_{local} = 16 * AS_path_length + IAC\ attr$
- On eBGP sessions, IAC attribute may be:
 - decremented by upto 14
 - incremented by upto 112
- Net effect: increase of 4 - 256 for every AS hop
- So much more granular than AS path length

Result

- Now:
 - paths over ISPs A and B both 3 hops
 - prepend A: almost all traffic over B
- With IAC:
 - paths over A and B both ± 48
 - IAC - I on B: slightly more traffic over B

More compatibility

- Don't want routers to suddenly select completely different paths after update
- solution: multiplier value
 - IAC attr * 0 = no effect
 - IAC attr * 1 = full effect
 - 0..1 = partial effect
- Deploy gradually in existing networks

Deployment

- Optional transitive attribute
- So survives IAC-ignorant AS hops
- Just no increment/decrement
- Existing increment/decrement resurfaces in downstream IAC-supporting ASes
 - so support in the middle isn't necessary

What do you think?