

draft-ietf-idr-bgp-nh-cost-02

Carrying next-hop cost information in BGP

Ilya Varlashkin <ilya@nobulus.com>

Robert Raszuk <robert@raszuk.net>

Keyur Patel <keyupate@cisco.com>

Manish Bhardwaj <manbhard@cisco.com>

Serpil Bayraktar <serpil@cisco.com>

IETF IDR interim 20150629

Overview

- Convey next-hop cost using BGP (typically to RR)
 - when bringing IGP to RR is difficult/impractical
 - when cost isn't IGP cost but arbitrary preference
- Building upon BGP ORR draft
- *Using BGP-LS as transport*
- Modifying best-path selection, still compatible

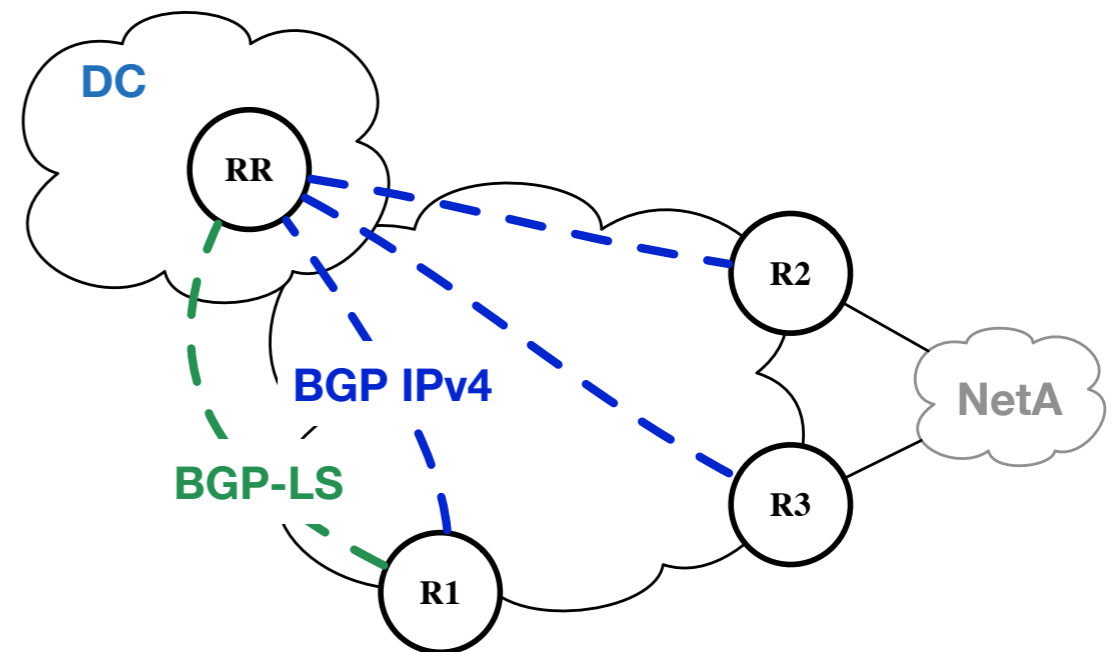
Cost Prefix Descriptor TLV

TLV Code Point	Description	Length	Value defined in:
TBD	Cost	4 bytes	Cost Value

- Can be used with NLRI Type 5 or 6
- Used together with IP Reachability SubTLV
- Means: “cost from sender to IP is X”

Operations

- R1, R2, R3 are RR clients
- R1 also speaks BGP-LS to RR
- RR uses R1 costs to calculate best paths for R1 (modified 9.1.2.2 of RFC4271)



NLRI Type=5; ProtoID=5;
SubTLV=265, 1.1.1.2/32; **SubTLV=Cost, 1**

NLRI Type=5; ProtoID=5;
SubTLV=265, 1.1.1.3/32; **SubTLV=Cost, 2**

Usage notes

- Not competing with ADDPATH - use together
- Normal Graceful Restart rules apply
- Can be used with Route Servers

Questions?