Distribute Service Metric by BGP

draft-lin-idr-distribute-service-metric-02

Changwang Lin(New H3C Technologies)

Huijuan Yao(China Mobile)

IETF 120

Motivation

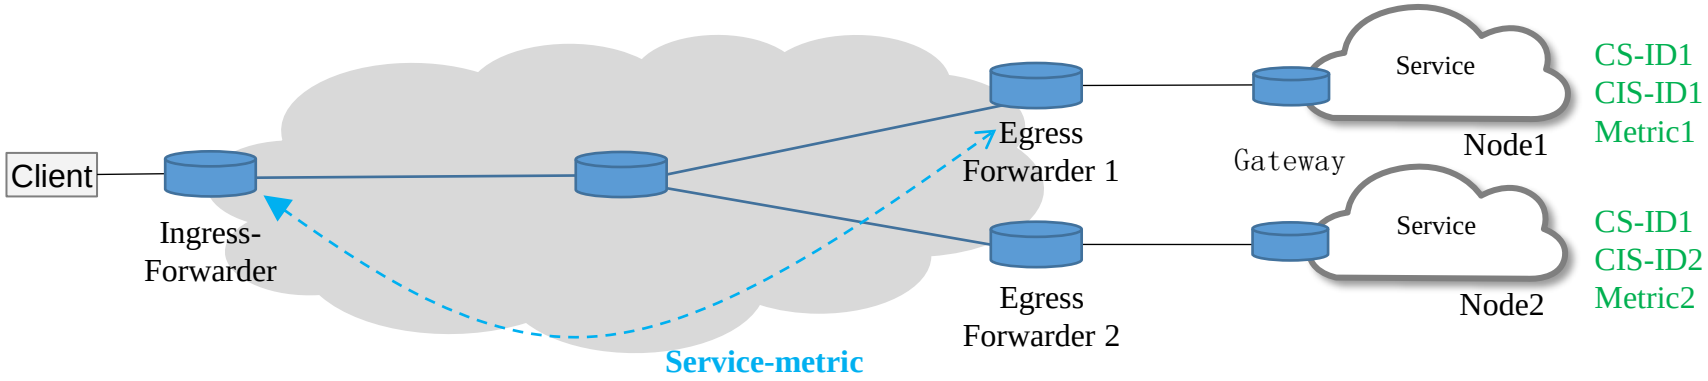

- **CATS：**
  - As described in draft-ietf-cats-framework-02, a C-SMA is required to collect service metrics.
  - Routing at the ingress-forwarder should be based on both network and service metrics.
- **Problem:**
  - Service metrics periodically change, unlike existing routing information.
  - Services are identified by CS-ID and CIS-ID, and there are limitations in passing them through the existing address families.
  - Periodic updates of service metrics need to avoid impacting existing routing.
  - Periodic updates of service metrics require efficient mechanisms to improve transmission performance.
- **Thoughts:**
  - Use extended BGP to transmit computation metrics, leveraging BGP's efficient and powerful extension capabilities.
  - Extend an independent address family to transmit computation metrics without affecting the transmission of other routing information.
  - Define flexible NLRI to efficiently achieve the transmission of service metrics.
  - Implement on-demand subscription functionality to enable the on-demand transmission of service metrics.

Extend BGP with new service metric address families

- The format of the Service Metric NLRI
  - 1 - Service Metric Register route
  - 2 - Service Metric Subscribe route
  - 3 - Service Metric Update route

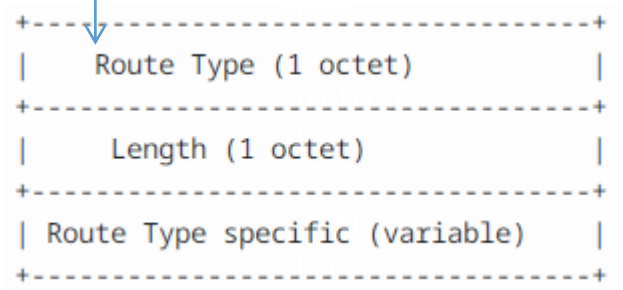

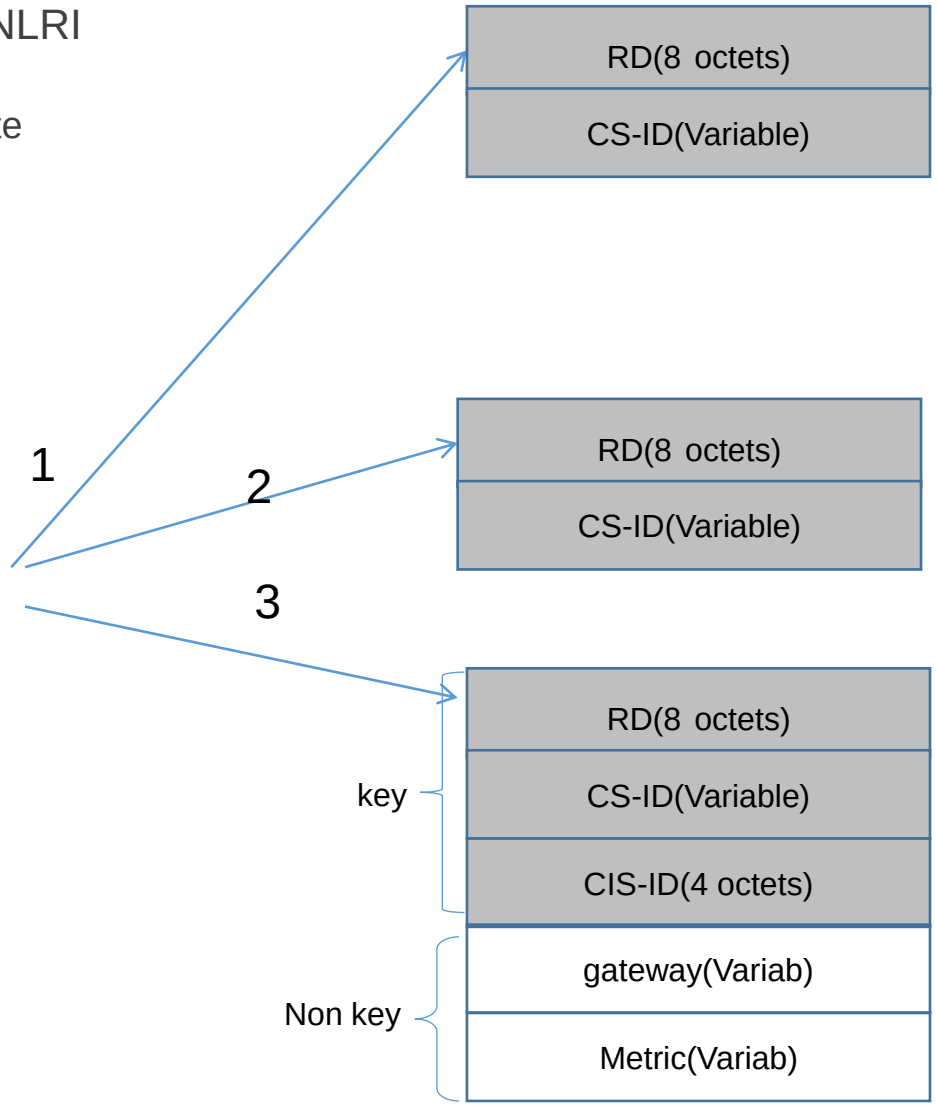

Operational Mechanism

Register

- Egress send registration route to RR
- RR forwards the registration route to all Ingress nodes.

Subscribe

- Ingress receives a service request and sends a subscription route to the RR.
- The RR forwards the subscription route to the corresponding Egress nodes.
- Egress responds with an updated route to Ingress..

Period Update

- Egress send periodic Update service metric route to subscribers.

UnSubscribe

- Upon forwarding table aging, send a withdraw subscription route to RR.
- RR notifies corresponding Egress of the route withdrawal.
- Egress no longer sends periodic route updates to the Ingress.

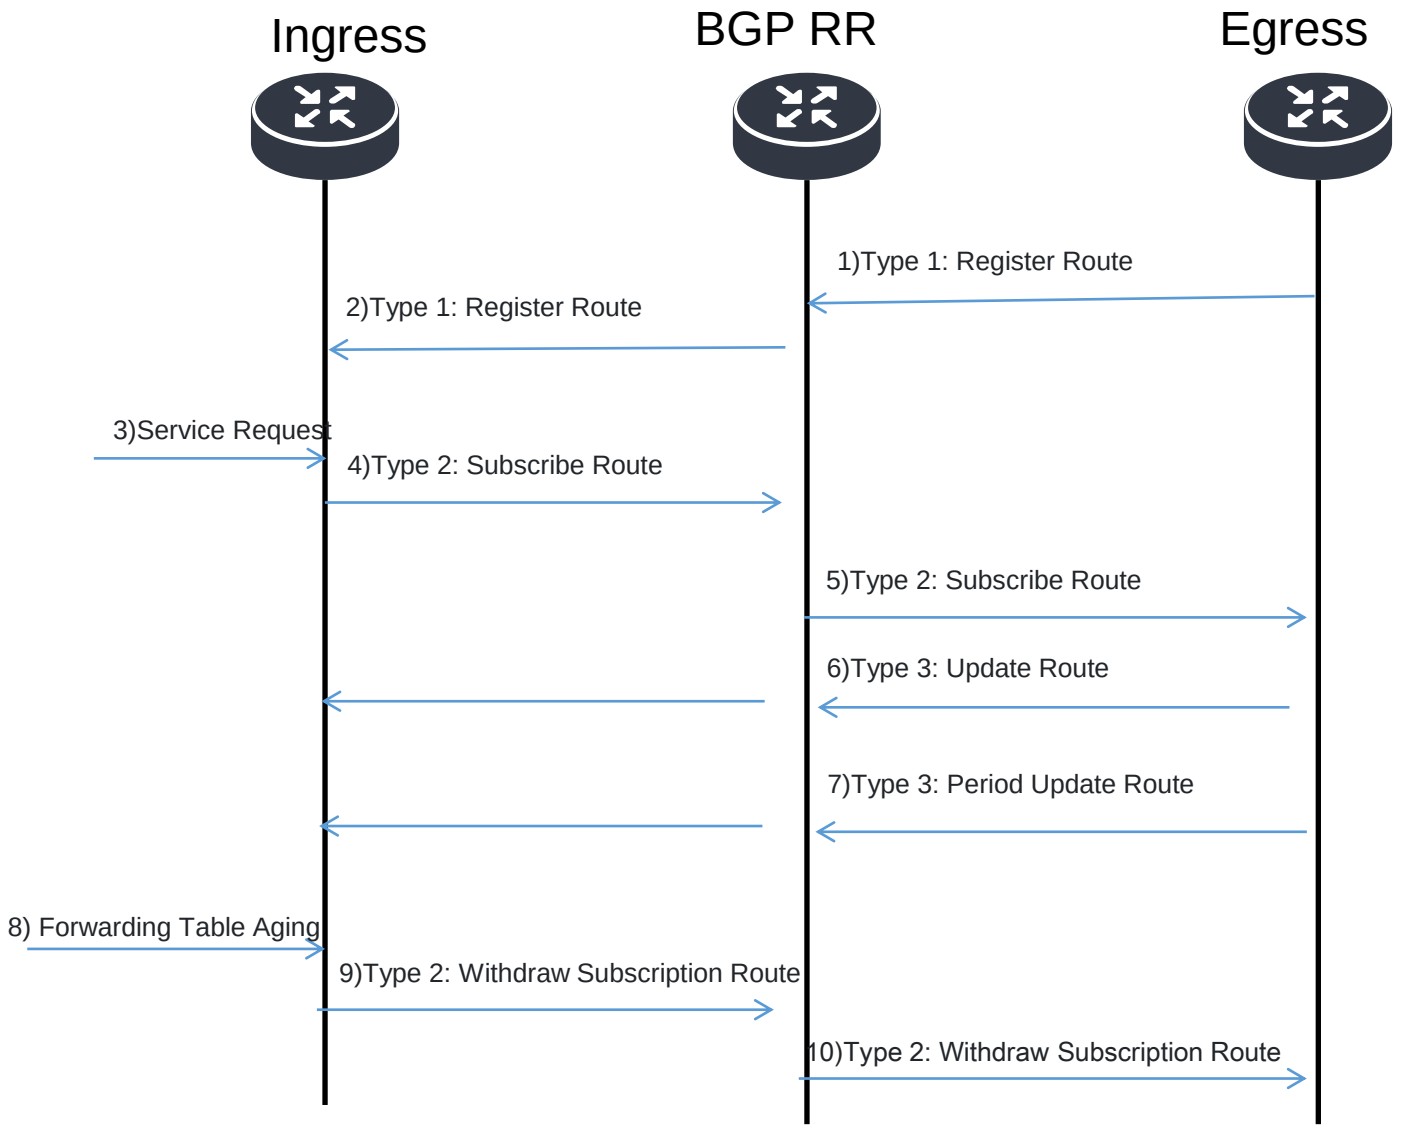

Summary

Solution:

- Defined new BGP address families to convey computing service metric information.
- Introduced variable-length NLRIs that distinguish between key and non-key parts.
- Metrics that change periodically are placed in the non-key part.
- Set the next hop for service metrics routing to 0, and placed the service-associated location address in the non-key part.
- Established mechanisms for subscribing to and publishing service metrics.

Goal:

- Minimize the impact on existing route propagation.
- Service CSID identifiers are not limited to anycast IP addresses.
- Improve the efficiency of sending periodic service metric updates.
- Send on-demand and control the propagation scope of service metrics.

Usecase

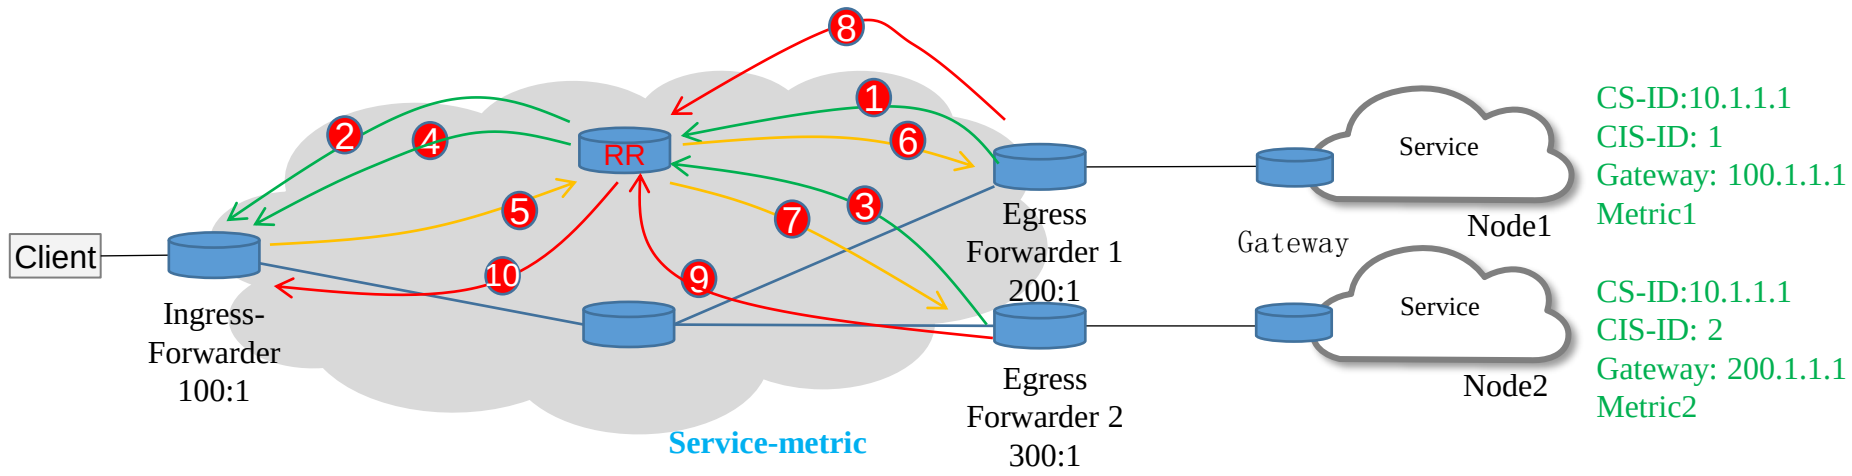

1 2 Type 1(Register) [RD: 200:1, CS-ID:10.1.1.1]

3 4 Type 1(Register) [RD: 300:1, CS-ID:10.1.1.1]

5 6 7 Type 2(Subscribe) [RD: 100:1, CS-ID:10.1.1.1]

8 Type 3(Update) [RD: 200:1, CS-ID:10.1.1.1, CIS-ID: 1,
Gateway: 100.1.1.1, Metric 1]

9 Type 3(Update) [RD: 300:1, CS-ID:10.1.1.1, CIS-ID: 2,
Gateway: 200.1.1.1, Metric 2]

10 Type 3(Update) [RD: 200:1, CS-ID:10.1.1.1, CIS-ID: 1,Gateway: 100.1.1.1, Metric 1]
[RD: 300:1, CS-ID:10.1.1.1, CIS-ID: 2, Gateway: 200.1.1.1, Metric 2]

Note: RR will bundle multiple prefixes (NLRIs) with the same attributes into a single Update packet.

Next Steps

- Any questions or comments are Welcomed.

7

Thanks