OSPF Stub neighbor Draft

draft-raza-ospf-stub-neighbor-00

Faraz Shamim – Cisco

Padma Pillay-Esnault – Cisco

Khalid Raza – Viptela

Andrew Kulawiak – Bank of America

John Cavanaugh – 405 Labs

Abstract

- Enhancement to support large scale hub and spoke neighbors
- Improves scalability
- Improves convergence
- New functionality restricts to the hub only
- Concepts also applicable to the hosts in VM environment

Introduction

- In a large hub and spoke model, L3 routing are being extended to the remotes (top of the rack/Hypervisor)
- Increased number of VM routes being advertised to the hub by VM hosts
- These remotes do not need any topological info for the entire area for optimal routing as it will never become transit
- These remotes usually do not have enough capacity to become part of a larger area
- Currently: remotes receives the entire area info from the hubs even with stub with no-summary
- Large number of areas puts a burden on ABR, remotes will still receive intra-area routes
- Ideally: remote should sends its subnet info to the hub & receive only a default from the hub

Proposal

- Extend OSPF to support very large hub and spoke topologies more efficiently
- Spokes only need IP reachability to the hub router that are gateways to the rest of the network
- Spokes should not receive unnecessary information from the hub about other routers in the area
- Spokes should only receive a default route and/or aggregated prefixes from the hub
- Spokes must be able to send its own connected routing information to the hub
- Incremental deployment approach, as the changes will only be limited to the hub routers - No upgrades required to the rest of the network*
- Hub can be a normal intra area router or could be an ABR/ ASBR

LSA filtering options – Option #1

- Hub is the ABR
- Filtering is done at the ABR
- Hub sends type 3 default towards the spokes
- Spokes do not receive other areas routes (via filtering)
- Spokes still receive other spokes routes within an area
- Prefixes aggregation can be done at the ABR towards the spokes

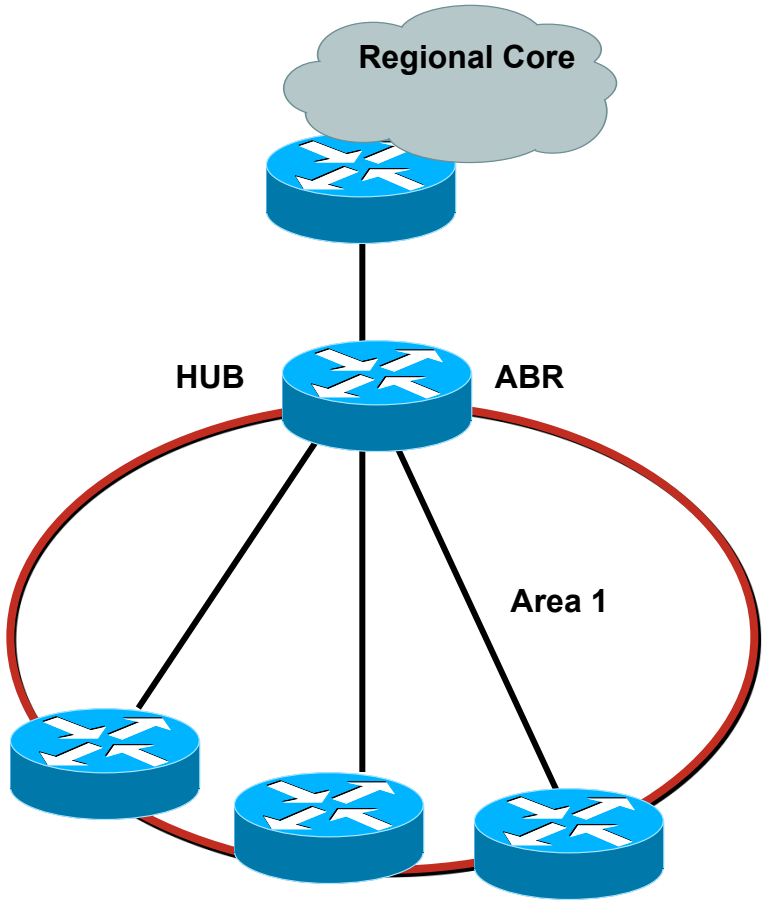

LSA filtering options – Option #2

- Hub is the ASBR
- Hub sends type 5 or type 7 default towards the spokes
- Spokes still receive other devices routes within an area
- Prefixes aggregation(external prefixes only) can be done at the ASBR towards the spokes

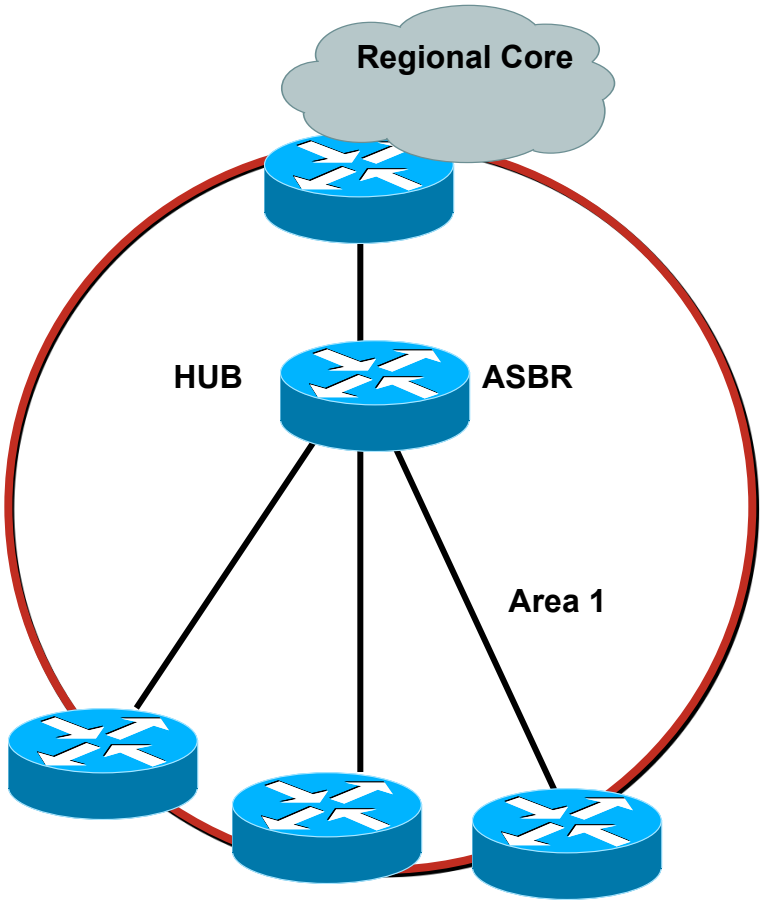

LSA filtering options – Option #3

- Hub is the neither ABR nor ASBR
- Filtering can not be done
- Spokes receives all intra area routes
- No aggregation of prefixes capability towards the spoke
- Severely impacts scalability

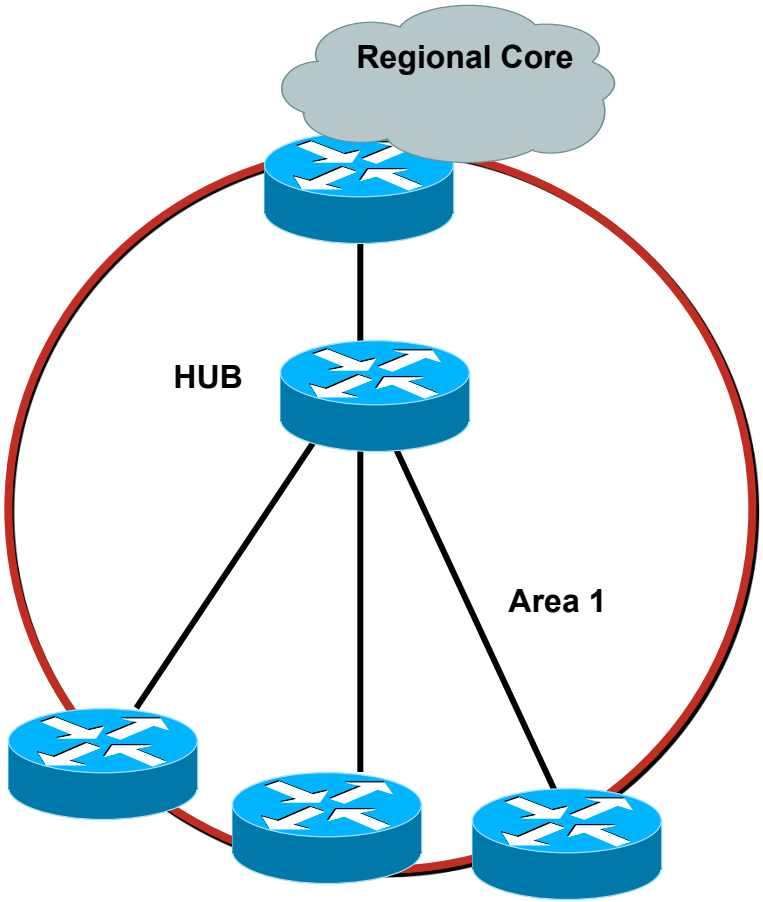

OSPF stub neighbor overview

- Local adjacency between hub and spoke - Only Hub will be aware of the local adjacency
- Local adjacency will not be advertised towards non-stub neighbors
- Hub “Local” router LSA will lists 2 links; point-to-point and a stub link (for configured range or a default prefix or both)
- New – intra area default route
- Hub router will effectively hide all the area topology from the spokes
- Demand circuit recommended for enhanced scalability for some use cases (long distance wireless network)

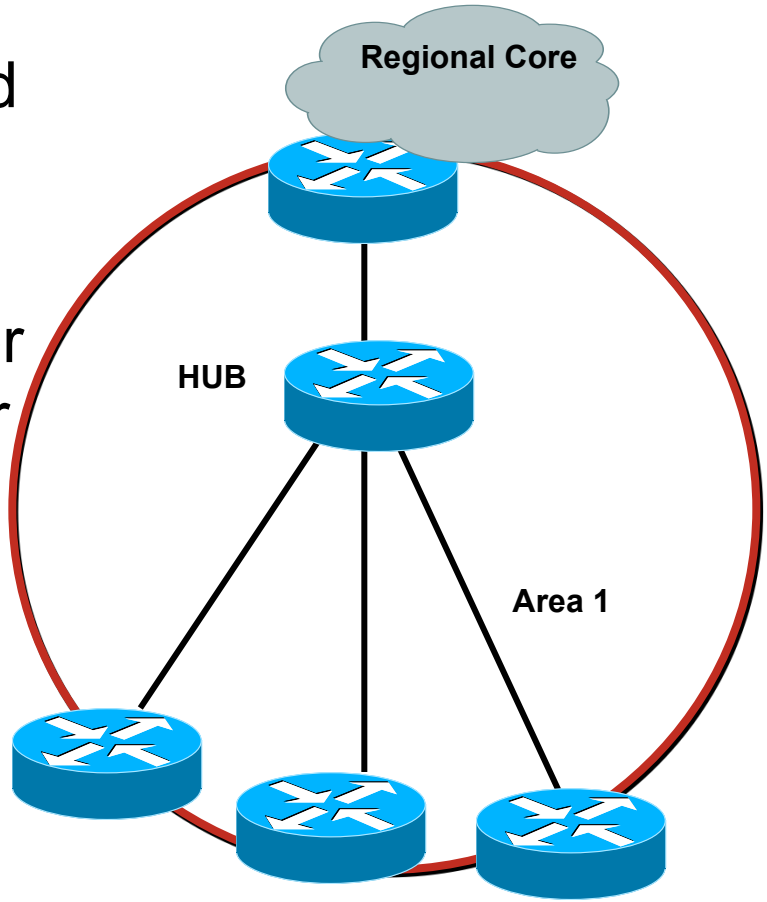

OSPF stub neighbor overview – default only

Router Link States (Area 1)

```
LS age   =   0
Options   =   (E-bit)
LS type   =   1
Link State ID = 192.0.3.1
Advertising Router   = 192.0.3.1
# links   =   2
   Link ID   =   192.0.6.1
   Link Data   =   192.0.3.1
   Type   =   1
    # TOS metrics   =   0
   metric   =   8
   Link ID   =   0.0.0.0
  Link Data   =   0.0.0.0
   Type   =   3
   # TOS metrics   =   0
   metric   =   8
```

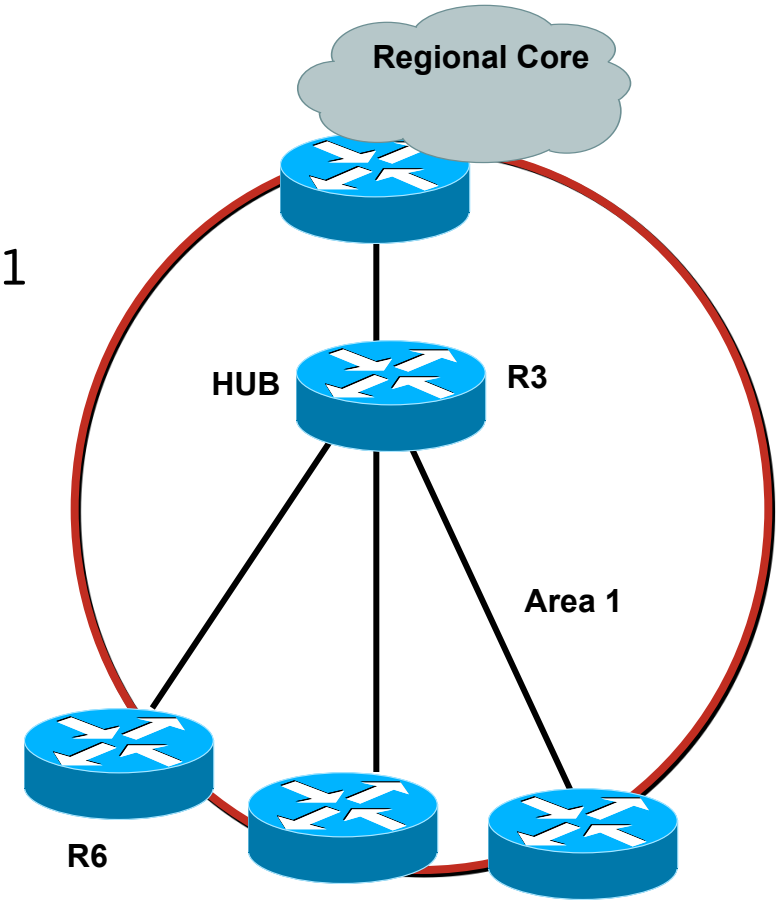

OSPF stub neighbor overview – Aggregated prefix only

Router Link States (Area 1)

```
LS age   =   0
Options   =   (E-bit)
LS type   =   1
Link State ID = 192.0.3.1
Advertising Router = 192.0.3.1
#links   =   2
   Link ID   =   192.0.6.1
   Link Data   =   192.0.3.1
   Type   =   1
   # TOS metrics   =   0
   metric   =   8
   Link ID   =   10.10.0.0
  Link Data   =   255.255.0.0
   Type   =   3
   # TOS metrics   =   0
   metric   =   8
```

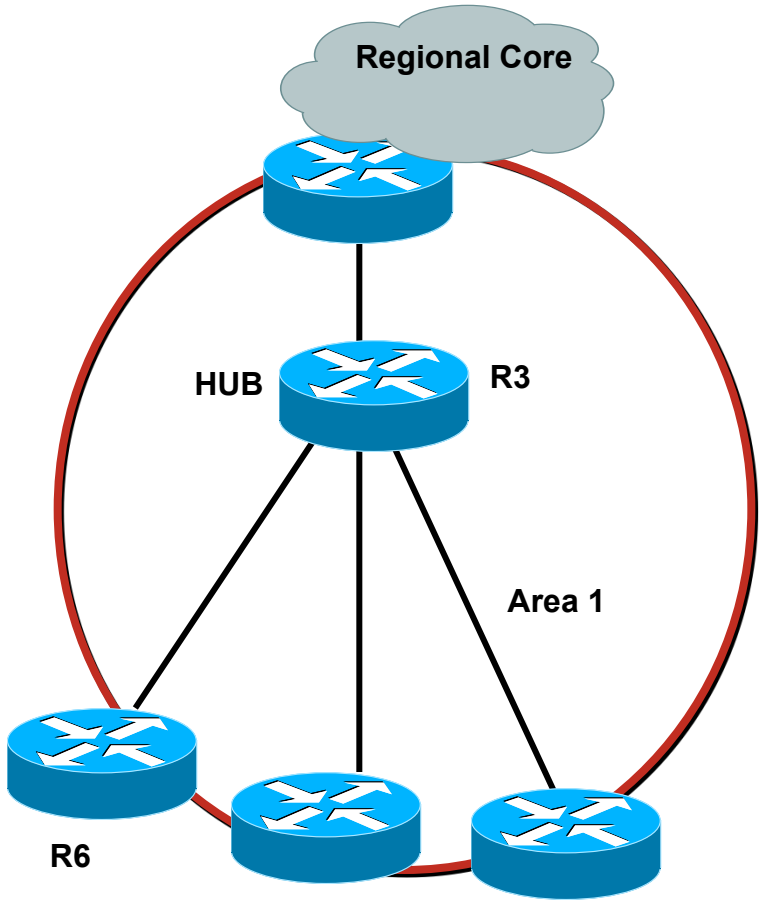

OSPF stub neighbor overview – Aggregated prefix & default route

Router Link States (Area 1)

```
LS age  =  0
Options  =  (E-bit)
LS type  =  1
Link State ID  =  192.0.3.1
Advertising Router  =  192.0.3.1
# links  =  3
    Link ID  =  192.0.6.1
    Link Data  =  192.0.3.1
    Type  =  1
     # TOS metrics  =  0
    metric  =  8
    Link ID  =  0.0.0.0
    Link Data  =  0.0.0.0
    Type  =  3
    # TOS metrics  =  0
    metric  =  8
    Link ID  =  10.10.0.0
    Link Data  =  255.255.0.0
    Type  =  3
    # TOS metrics  =  0
    metric  =  8
```

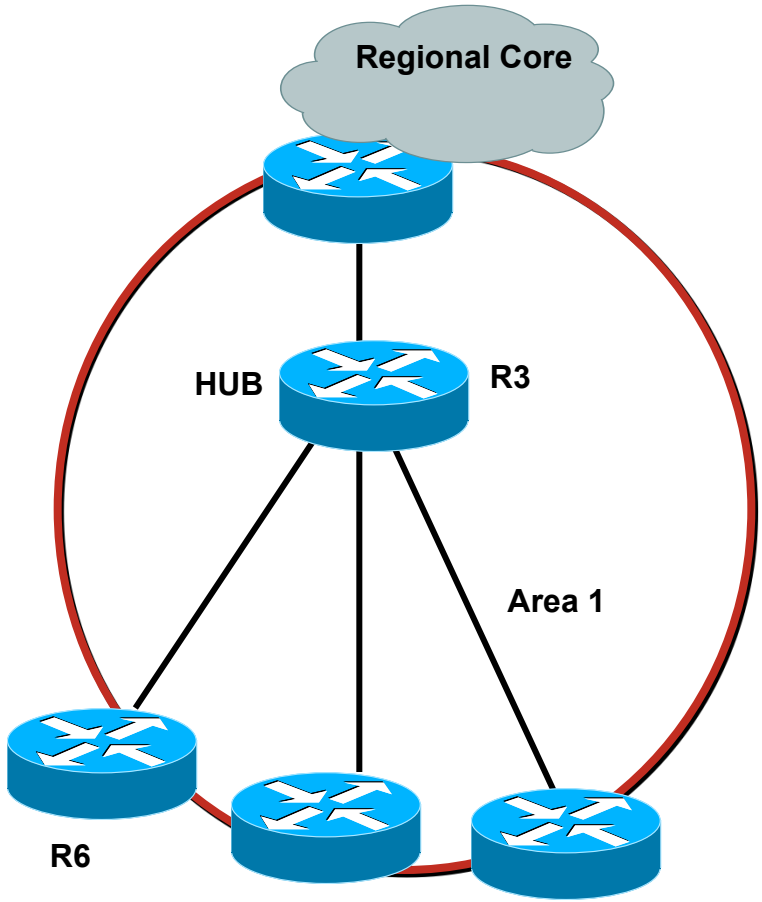

OSPF stub neighbor – DUAL homed

- In case of dual-homed spokes, both hubs MUST be configured to view the spoke as a stub neighbor
- Local Router LSA of a hub could get flooded over towards the other hub
- Hub SHOULD ignore local Router LSAs from other hub flooded by the stub neighbors
- During migration, a spoke can temporary become a transit for the second hub(not yet converted to stub neighbor)

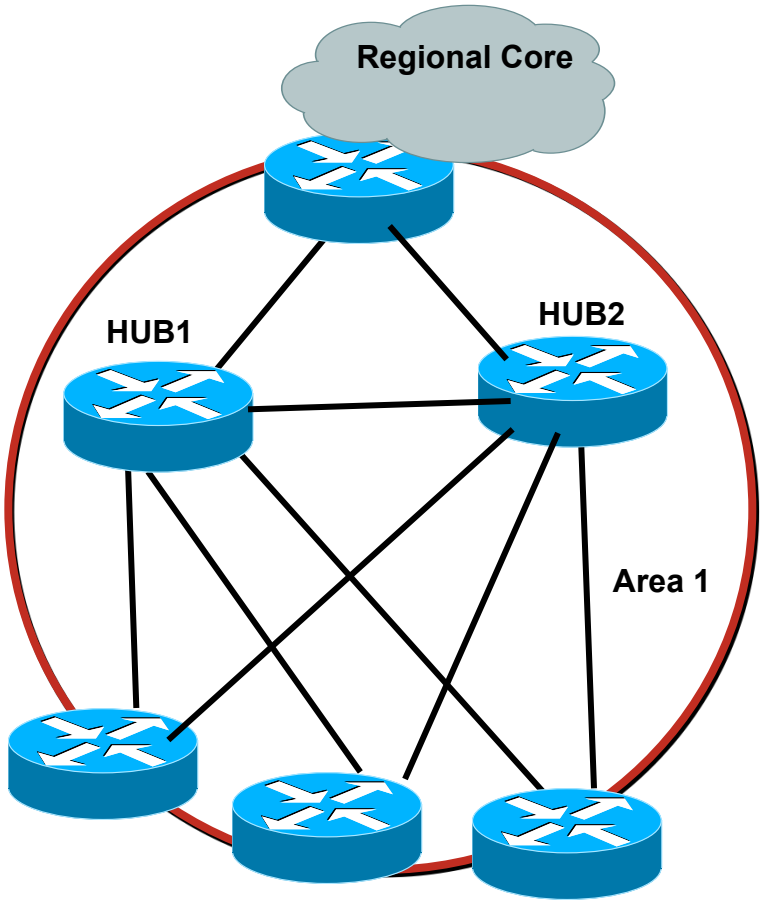

Receiving & propagation of spoke routes

- Rest of the network should not be aware of this new local adjacency
- Hub router upon receiving the route from spoke SHOULD NOT treat those routes as intra-area
- Hub that is acting as an ABR should convert all routes received via stub neighbor as if they were part of single area
- Hub can now summarize all these routes at the area boundary
- With well defined addressing all spoke routes could be summarized into a single or very few routes

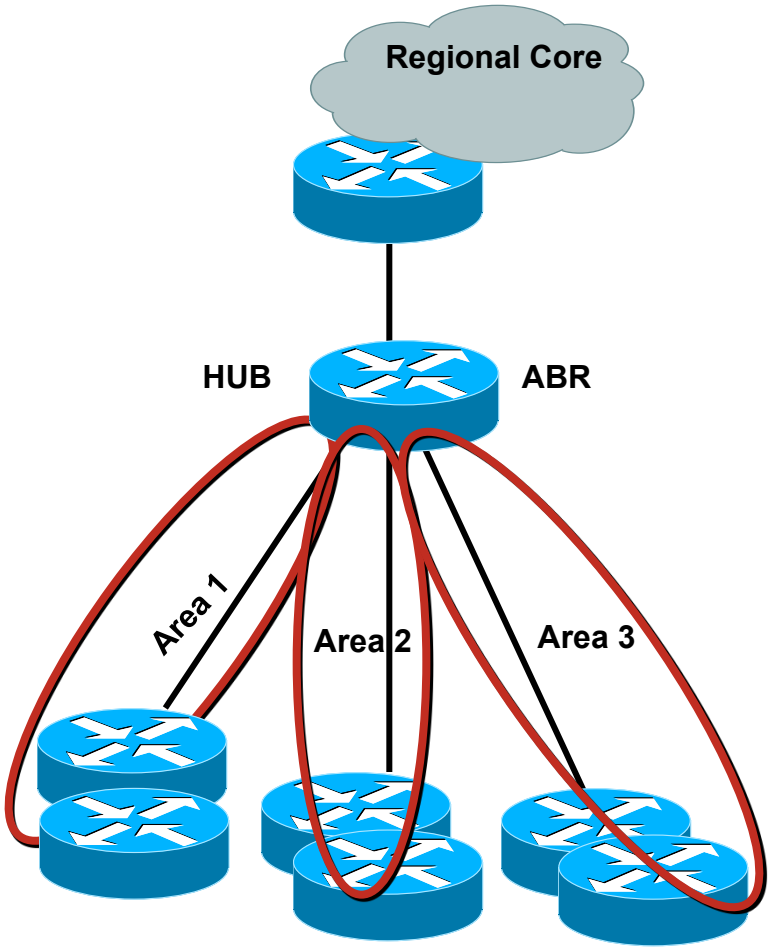

Q and A
Feedback
Next Step