

Design Concern	Full Mesh IBGP	BGP Route Reflector	BGP Confederation
Scalability	Least, between each BGP node there is an IBGP session, number of session and configuration is highest	High since there is only one IBGP session on the RR client, if there are two RRs, two sessions per RR client	Medium since it requires Route Reflector inside each Sub AS for scalability
Logical Topology	Full Mesh	Hub and Spoke between Route Reflectors and the RR Clients	Point to point between Sub Ases, Hub and Spoke RR-RR client in Sub AS or Full mesh IBGP inside Sub AS
Resource Requirement	Highest	Lowest	Medium
Number of Next Hop	All available next-hops are sent thus highest resource consumption on the BGP routers.	Only one best path is sent by RR to RR Clients. Thus lowest resource consumption among other options	Depends on Full mesh IBGP or RR inside Sub AS, if full mesh all the available paths, if RR is used, then only one best path is sent
Loop Prevention	No loop since all the BGP routers direct IBGP peer	Cluster List and Originator ID are used for loop prevention	Between Sub Ases EBGP rules apply, if Sub AS sees it's own AS number in the AS path, packet is dropped
Migration from Full Mesh	No	Migration from Full Mesh IBGP is easy compare to Confederation	Migration from Full Mesh IBGP is hard compare to BGP Route Reflectors
Standard Protocol	Yes	Yes	Yes
Different IGP inside an AS	Technically Yes but running same IGP is the best practice inside an AS	Technically Yes but running same IGP is the best practice inside an AS	Definitely yes and there is no problem. One of the advantages of BGP Confederation, compare to BGP RR design. Inside an each Sub AS, different IGP can run, IGP topologies are not extended between SUB Ases
Commonly Used	Small scale environment most common	Large scale environment most common	Not very common
Path Visibility	All the available path are known by the IBGP routers. Highest level visibility	Route Reflectors only advertise best path to RR client, path visibility is worst	Depends. Inside Sub AS if Full Mesh IBGP, it is best, if RR is running, worst.
Provisioning	Hard since it requires too much configuration	Very easy, one IBGP session per RR	Medium. Worse than Route Reflector.
IPv6 Support	Yes	Yes	Yes
Default Convergence	Fastest since all IBGP routers connected directly, second best path is already known by the routers	Worst since RR adds extra convergence time and the RR clients know only the best path when RR is deployed	Depends. If Full Mesh IBGP then fast, if RR inside an AS, then slow.
MPLS VPN support	Yes	Yes	Yes but hard to manage since Inter-AS VPN operation is required