Segment Routing - Segment Types

Global Segment - All SR enabled nodes in the SR Domain support the instruction that is associate with a Global Segment. So each node installs the instruction in its forwarding table.

Local Segment - Only the node that originates a Local Segment supports the instruction and thus is the only node to install it in its forwarding table.

A **SID (Segment Identifier)**, as the name suggests, identifies a Segment. The format of a SID can be:
> An MPLS Label
> An IPv6 address (SRv6)
> An index value in an MPLS label space - if this method is used, the SRBG (which is included in the SID-advertisement) is combined with the index reach the label.

Segment Operations
PUSH - Insert new instruction at the head of the Segment List. If MPLS labels are used a new label, representing that Segment, is PUSHED onto the stack.
CONTINUE - The active segment is not complete. If MPLS labelling is used the label is SWAPPED for the label that next hop uses for this SID. If the same SRGB is used within the domain, this swap operation will be to the same label. The continue action will invoke PHP if the forwarding node is the penultimate node to the originator of the SID, and the "PHP-off" and "Explicit-Null" flags were unset in the SID advertisement (the default is for both to be unset).
NEXT - The active segment is complete so the next one in the list becomes the active segment. If MPLS labelling is used the top most label will be POPPED.

Segment*	Type & Significance**	Instruction***	Details	Example
IGP Prefix Segment	Prefix Segment with Global Significance	Steer Traffic across the ECMP-aware shortest path to the prefix associated with that Segment	A Prefix-SID can only be associated with a single prefix. A Prefix-SID can be assigned to any prefix, not just a /32 or /128 host prefix.	lo0: 6.6.6.6/24 SID 16006. Incoming Active Segment: 16006 Action: CONTINUE Next hop: R2 and R4 (ECMP)
IGP Node Segment	Prefix Segment with Global Significance	Steer Traffic across the ECMP-aware shortest path to the prefix associated with that Segment	A Node-SID is a special type of Prefix-SID. It's forwarding is identical to a Prefix-SID. The Node-SID is associated with a prefix that identifies that host and must be a /32 or /128. The N-Flag in a SID advertisement identifies it as a Node SID.	lo0: 6.6.6.6/32 SID 16006. Incoming Active Segment: 16006 Action: CONTINUE Next hop: R2 and R4 (ECMP). Incoming Active Segment: 16005 Action: CONTINUE Next hop: R4. lo0: 5.5.5.5/32 SID 16005
IGP Anycast Segment	Prefix Segment with Global Significance	Steer Traffic across the ECMP-aware shortest path to the prefix associated with that Segment. There could be multiple nodes advertising that Segment.	Traffic is forwarded to the nearest and anycast node. If multiple nodes within the "Anycast Set" have the same cost then packets will be load balanced. There are inherent redundancy advantages as with any Anycast setup. Using an Anycast SID is useful if you want steer traffic via a set of transit nodes on the way to a destination. The N-flag will not be set when advertising Anycast-SIDs.	Anycast-SID 16024. Incoming Active Segment: 16024 Action: CONTINUE Next hop: R2. Incoming Active Segment: 16024 Action: CONTINUE Next hop: R2 and R4 (ECMP). lo0: 1.1.1.24/32 On both R2 and R4 are advertised with Anycast SID 16024
(Layer 3) IGP Adjacency Segment	Adjacency Segment with Local Significance to the advertising node	Steer Traffic out of the link(s) of the IGP adjacency associated with this segment, regardless of lowest cost	These are assigned locally by nodes for their own adjacencies. These are advertised throughout the SR domain but only the originating node will create a forwarding entry for it. This is typically done to forward traffic out of an interface towards a destination if the shortest path to that destination is not out of that interface (if it was, the Prefix-SID would simply be used). There can be more than one Adj-SID associated with a link, such that each Adj-SID has a different property linked with it	Adj-SID 17021. Adj-SID 17023. Adj-SID 17024. Incoming Active Segment: 17021 Action: NEXT Outgoing interface: to R1. Incoming Active Segment: 17024 Action: NEXT Outgoing interface: to R4. Incoming Active Segment: 17023 Action: NEXT Outgoing interface: to R3
Layer 2 Adjacency Segment	Adjacency Segment with Local Significance to the advertising node	Steer Traffic out of the referenced Layer 2 bundle member, typically a physical interface that is part of a LAG	IGP adjacencies are formed over layer 3. LAG protocols bundle several layer 2 interfaces together, over which one single IGP adjacency is formed. A layer 2 Adjacency-SID is associated with a specific link of a layer 2 bundle. For example, a layer 2 Adjacency-SID can be used to send traffic over a specific interface of an EtherChannel, even if only one IGP adjacency is formed over that EtherChannel	L2 Adj-SID 17123. L2 Adj-SID 17223. L2 Adj-SID 17323. Adj-SID 17023. Bundle Po1. Incoming Active Segment: 17123 Action: NEXT Outgoing interface: upper port to R3. Incoming Active Segment: 17223 Action: NEXT Outgoing interface: middle port to R3. Incoming Active Segment: 17323 Action: NEXT Outgoing interface: lower port to R3. Incoming Active Segment: 17023 Action: NEXT Outgoing interface: Po1 (using ECMP)
Group Adjacency Segment	Adjacency Segment with Local Significance to the advertising node	Steer Traffic out of the link(s) of the grouped IGP adjancenc(ies) associated with this Segment, in an ECMP manner	This is used in cases whereby there are multiple adjacencies between nodes that are not bundled together via a LAG protocol, but you still wish to group ports together.	Group Adj-SID 17123. Group Adj-SID 17223. 3 separate Adjacencies No bundle. Incoming Active Segment: 17123 Action: NEXT Outgoing interface: upper port to R3. Incoming Active Segment: 17223 Action: NEXT Outgoing interface: ECMP over the lower two ports to R3
BGP Prefix Segment	BGP Prefix Segment with Global Signficance - understood by all nodes in the SR BGP Domain	Steer traffic along the ECMP-aware BGP multi-path to the prefix associated with this Segment	A BGP prefix-SID is associated with a BGP prefix in the same way an IGP Prefix-SID is associated with an IGP prefix (OSPF or IS-IS)	R3 and R4 have BGP sessions to R1 and R2. Incoming Active Segment: 16004 Action: CONTINUE Next hop: R1 and R2 (ECMP). lo0: 4.4.4.4/32 SID 16004
BGP Anycast Segment	BGP Prefix Segment with Global Signficance - understood by all nodes in the SR BGP Domain	Steer traffic along the ECMP-aware BGP multi-path to the prefix associated with this segment. There could be multiple nodes advertising this Segment.	Traffic is forwarded to the nearest anycast node via the BGP learned prefix. If multiple nodes within the "Anycast Set" have the same cost then packets will be load balanced. There are inherent redundancy advantages as with any Anycast setup. This is useful to identify a group of transit nodes, for example spine nodes running BGP in a DC, to steer the traffic through on the way to some other destination. The N-flag will not be set when advertising Anycast-SIDs.	R4 and R5 have BGP sessions to R1, R2 and R3. Incoming Active Segment: 17023 Action: CONTINUE Next hop: R2 and R3 (ECMP). BGP Anycast-SID 17023. lo0: 1.1.1.23/32 On both R2 and R3 are advertised with BGP Anycast SID 17023. lo0: 5.5.5.5/32 SID 16005. If R4 receives a segment list of (17023, 16005) the packet will transit R2 and R3 using ECMP before being sent to R5
BGP Peer Node Segment	BGP Peer Segment with Local Significance	Steer Traffic to the specific BGP peer node via ECMP multi-path towards that peer router	BGP Peer Node Segments specify a neighbor and is signalled using BGP Link State. This Segment will override the traditional BGP decision making process. It applies to both eBGP and iBGP. The BGP Peer Node-SID is associated with a prefix that identifies that BGP Peer.	AS 1. BGP-LS Peerings. Control Server. Incoming Active Segment: 16002 Action: NEXT Next hop: R2. Peer Node SID 16002. Peer Node SID 16003. BGP-LS is used to communicate the SIDs assigned by R1 to the central control server. AS 2. AS 3. BGP Peerings
BGP Peer Adjacency Segment	BGP Peer Segment with Local Significance	Steer Traffic to the specific BGP peer node via the specified interface towards that peer	BGP Peer Adjacency Segments specify a path or link to a specific neighbor and is signalled using BGP Link State. For single-hop BGP session, a BGP Peer Adjacency SID is not created since the BGP Peer Node SID provides the same functionality. This segment will override the traditional BGP decision making process. It applies to both eBGP and iBGP.	AS 1. BGP-LS Peerings. Control Server. Incoming Active Segment: 16031 Action: NEXT Next hop: R3 (via left link). Peer Adj SID 16031. Peer Adj SID 16032. BGP-LS is used to communicate the SIDs assigned by R1 to the central control server. AS 2. AS 3. eBGP Multihop Peering Between Loopbacks. A loopback to loopback peering is used to illustrate different links to what is the same neighbor (interface to interface would require two BGP sessions making the assigned segments BGP Peer SIDs)
BGP Peer Set Segment	BGP Peer Segment with Local Significance	Steer Traffic to the specific BGP peer node(s) that are members of the set via ECMP BGP multi-path	BGP Peer Set Segments specify a group of neighbors, the resulting instruction of which is to ECMP towards those neighbors. This is similar to the Anycast BGP SIDs but refers to neighbors and not BGP prefixes learned from neighbors (as it the case with BGP Anycast-SIDs) This segment will override the traditional BGP decision making process. It applies to both eBGP and iBGP.	AS 1. BGP-LS Peerings. Control Server. Incoming Active Segment: 17123 Action: NEXT Next hop: R2 and R3 (ECMP). This segment is only significant to R1 and refers to neighbors not links. BGP-LS is used to communicate the SIDs assigned by R1 to the central control server. AS 2. AS 3. BGP Peerings. Peer set (R2, R3): Label 17123

* SID and Segment are used interchangeably in this table. So a Prefix Segment is a Prefix SID.
** Adjacency SIDs could be configured globally but it is usually not recommended since doing so would mean keeping unnecessary forwarding state
*** The term "shortest path" could really say "shortest path according to the applied algorithm". The default is the simple SPF algorithm but SR Architecture allows for the use of other algorithms, such as Strict SPF.

by Steven Crutchley