

BGP Route Aggregation

Minimizing the Size of the Routing Table

Table of Contents

What is BGP Route Aggregation	3
Benefits of BGP Route Aggregation	3
BGP Route Aggregation with Static Discard Route	6
BGP Route Aggregation with Aggregate-address Command	7
Option summary-only	8
Option suppress-map	8
Option unsuppress-map	9
Option attribute-map	10
Option advertise-map	11
Option as-set	11
Conclusion	13

What is BGP Route Aggregation

Route Aggregation (RA) also known as BGP Route Summarization is a method to minimize the size of the routing table, announcing the whole address block received from the Regional Internet Registry (RIR) to other ASes. RA is opposite to non-aggregation routing, where individual sub-prefixes of the address block are announced to BGP peers. RA reduces the size of the global routing table, decreases routers' workload and saves network bandwidth.

Benefits of BGP Route Aggregation

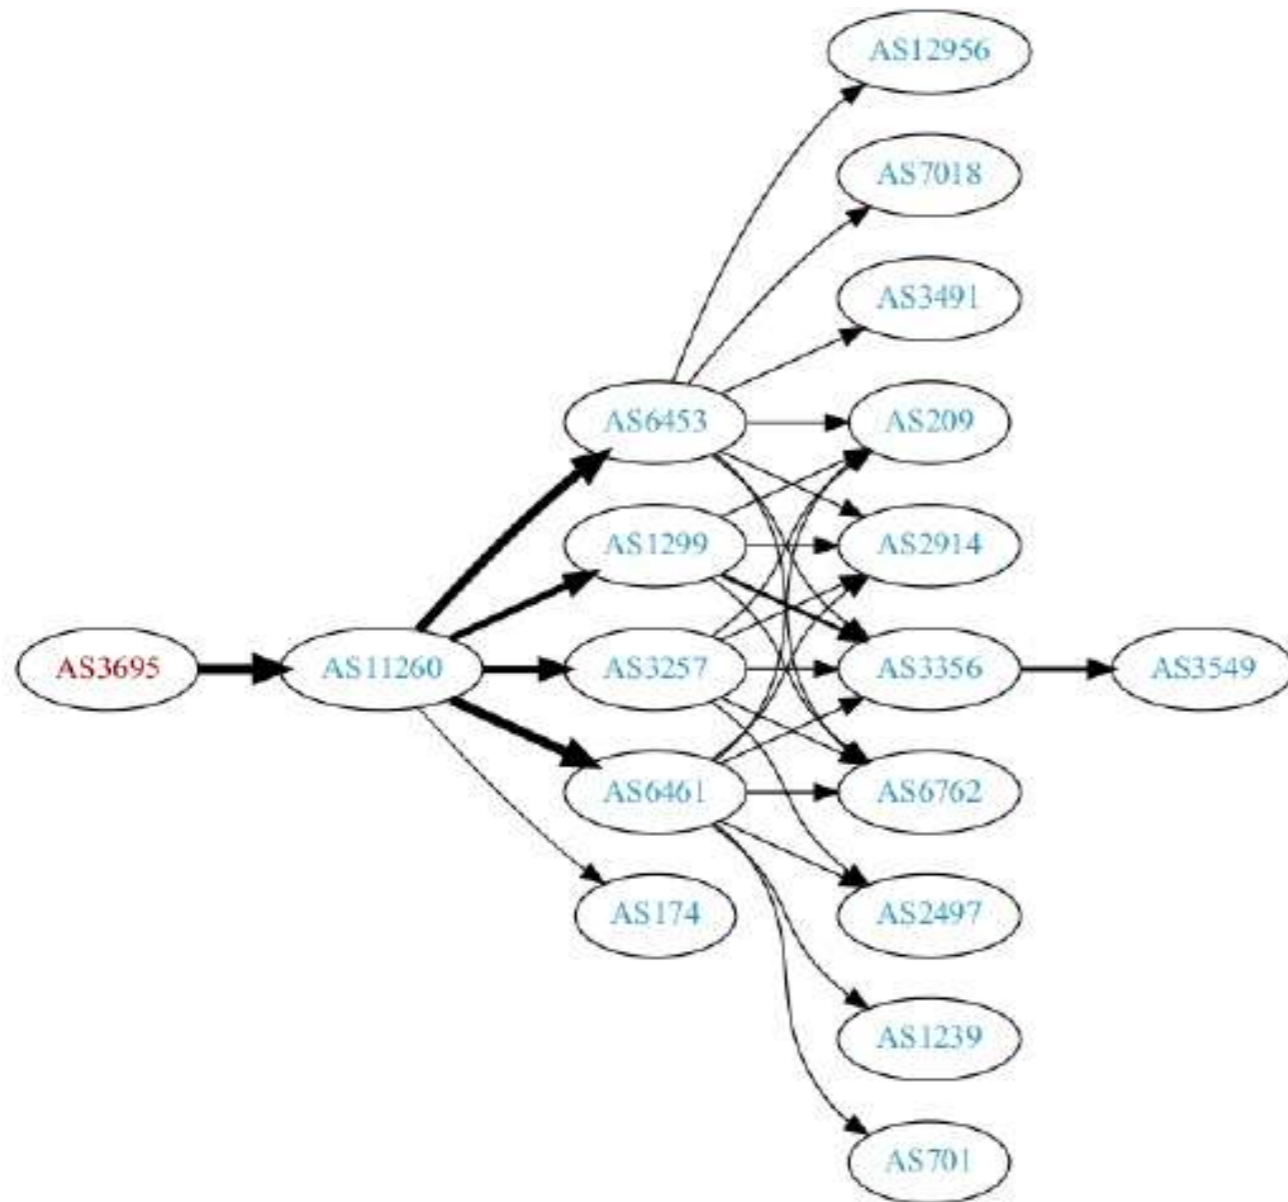
We will discuss the benefits of Route Aggregation using a real-world scenario. Let's look at SPM TELECOM (AS 3695) from Saint Pierre and Miquelon (PM) that announces an overall of 16 prefixes (Picture 1) to its Canadian upstream peer EastLink (AS 11260), (Picture 2). Each prefix is advertised with the prefix length 24.

Prefix	Country	Description
70.36.0.0/24	PM 	SPM TELECOM
70.36.1.0/24	PM 	SPM TELECOM
70.36.2.0/24	PM 	SPM TELECOM
70.36.3.0/24	PM 	SPM TELECOM
70.36.4.0/24	PM 	SPM TELECOM
70.36.5.0/24	PM 	SPM TELECOM
70.36.6.0/24	PM 	SPM TELECOM
70.36.7.0/24	PM 	SPM TELECOM
70.36.8.0/24	PM 	SPM TELECOM
70.36.9.0/24	PM 	SPM TELECOM
70.36.10.0/24	PM 	SPM TELECOM
70.36.11.0/24	PM 	SPM TELECOM
70.36.12.0/24	PM 	SPM TELECOM
70.36.13.0/24	PM 	SPM TELECOM
70.36.14.0/24	PM 	SPM TELECOM
70.36.15.0/24	PM 	SPM TELECOM

Picture 1 - Prefixes Announced by ASN 3695

Source: <https://dnslytics.com/bgp/as3695>

BGP Route Aggregation



Picture 2 - ASN 3695 Route Propagation
 Source: <https://bgpview.io/asn/3695#graph>

The <http://www.cidr-report.org> website represents a powerful aid for ISPs as it provides an option to search and compute the aggregation prefixes on a per AS basis. For instance, the site suggests AS 3695 to announce a single aggregated prefix 70.36.0.0/20 instead of 16 individual prefixes with the length of 24 (Picture 3). Once implemented, the number of prefixes announced by AS 3695 would be reduced by 93.75%. Alternatively, the prefixes can be announced as two blocks 70.36.0.0/21 and 70.36.8.0/21. In both cases, Route Aggregation decreases the size of the global routing table, reducing the number of entries in the Forwarding Information Base (FIB) which is stored in a router's Ternary Content Addressable Memory (TCAM). This memory is faster than the ordinary RAM and is used for rapid table lookups. With the fast global routing table growth exceeding 700 000 entries (Picture 4) and a finite number of FIB entries due to hardware limitation, Route Aggregation can prevent TCAM resource exhaustion for many network devices.


Rank AS	AS Name	Current	Withdw	Aggte	Annce	Redctn	%
2940	AS3695	16	16	1	1	15	93.75%

Prefix	AS Path	Aggregation Suggestion
70.36.0.0/20	4777 2497 3257 11260 3695 + Announce	- aggregate of 70.36.0.0/21 (4777 2497 3257 11260 3695) and 70.36.8.0/21 (4777 2497 3257 11260 3695)
70.36.0.0/24	4777 2497 3257 11260 3695 - Withdrawn	- aggregated with 70.36.1.0/24 (4777 2497 3257 11260 3695)
70.36.1.0/24	4777 2497 3257 11260 3695 - Withdrawn	- aggregated with 70.36.0.0/24 (4777 2497 3257 11260 3695)
70.36.2.0/24	4777 2497 3257 11260 3695 - Withdrawn	- aggregated with 70.36.3.0/24 (4777 2497 3257 11260 3695)
70.36.3.0/24	4777 2497 3257 11260 3695 - Withdrawn	- aggregated with 70.36.2.0/24 (4777 2497 3257 11260 3695)
70.36.4.0/24	4777 2497 3257 11260 3695 - Withdrawn	- aggregated with 70.36.5.0/24 (4777 2497 3257 11260 3695)
70.36.5.0/24	4777 2497 3257 11260 3695 - Withdrawn	- aggregated with 70.36.4.0/24 (4777 2497 3257 11260 3695)
70.36.6.0/24	4777 2497 3257 11260 3695 - Withdrawn	- aggregated with 70.36.7.0/24 (4777 2497 3257 11260 3695)
70.36.7.0/24	4777 2497 3257 11260 3695 - Withdrawn	- aggregated with 70.36.6.0/24 (4777 2497 3257 11260 3695)
70.36.8.0/24	4777 2497 3257 11260 3695 - Withdrawn	- aggregated with 70.36.9.0/24 (4777 2497 3257 11260 3695)
70.36.9.0/24	4777 2497 3257 11260 3695 - Withdrawn	- aggregated with 70.36.8.0/24 (4777 2497 3257 11260 3695)
70.36.10.0/24	4777 2497 3257 11260 3695 - Withdrawn	- aggregated with 70.36.11.0/24 (4777 2497 3257 11260 3695)
70.36.11.0/24	4777 2497 3257 11260 3695 - Withdrawn	- aggregated with 70.36.10.0/24 (4777 2497 3257 11260 3695)
70.36.12.0/24	4777 2497 3257 11260 3695 - Withdrawn	- aggregated with 70.36.13.0/24 (4777 2497 3257 11260 3695)
70.36.13.0/24	4777 2497 3257 11260 3695 - Withdrawn	- aggregated with 70.36.12.0/24 (4777 2497 3257 11260 3695)
70.36.14.0/24	4777 2497 3257 11260 3695 - Withdrawn	- aggregated with 70.36.15.0/24 (4777 2497 3257 11260 3695)
70.36.15.0/24	4777 2497 3257 11260 3695 - Withdrawn	- aggregated with 70.36.14.0/24 (4777 2497 3257 11260 3695)

Picture 3 - Computed Aggregation for ASN 3695

Source: <http://www.cidr-report.org/cgi-bin/as-report?as=AS3695&view=2.0>

BGP Route Aggregation



Picture 4 - Active FIB Entries

Source: <https://bit.ly/2Kbf7Ci>

Besides decreasing the routers' workload and saving network bandwidth, Route Aggregation positively impacts routing stability. Let's go back to the scenario in which Route Aggregation is not configured on a SPM TELECOM's router (AS 3695). Provider announces prefixes in their internal BGP (iBGP) to the Internet over external BGP (eBGP). The customer's prefix 70.36.0.0/24 is injected into the

internal BGP when the link to the customer is active, and gets withdrawn when the link to the customer is inactive. If the customer's link goes down, prefix 70.36.0.0/24 becomes unreachable. AS 3695 announces network withdrawal for 70.26.0.0/24 to AS 11260. East-Link's router (AS 11260) in its turn advertises the withdrawal to the BGP peers. As a result, all Internet routers having the full BGP table view remove network 70.36.0.0/24 from their FIB tables that in its turn adds to their load. In case a customer's link goes up, the prefix is reinjected in AS3695 iBGP and re-advertised to EastLink router (AS 11260). It again, adds to the load. Finally, even if a customer's link goes up, connection to the Internet is not immediately available for a customer due to the BGP propagation delays.

Now, let's say that SPM Telecom (AS 3695) announces the aggregated prefix 70.36.0.0/20 and the prefix 70.36.0.0/24 becomes unreachable due to the flapping customer's link. The network 70.36.0.0/24 is withdrawn from AS 3695 iBGP but the aggregated prefix 70.36.0.0/20 is still announced to EastLink (AS 11260). However, routing tables of the upstream routers are not impacted as there are no network updates sent from AS3695 towards its eBGP peers. Once the link goes up, the prefix 70.36.0.0/24 is re-injected into AS 3695's iBGP and connection to the Internet is immediately available for the customer. Route Aggregation improves the network stability and saves network bandwidth, as the unnecessary network updates are reduced or completely eliminated when a part of the network goes down.

BGP Route Aggregation

The next part discusses the configuration of Route Aggregation and explains parameters that can be used to manipulate the aggregate route. For the purpose of demonstration, we will use the network topology based on a real-world scenario. The network consists of three routers with eBGP configured on all of them. The prefixes 70.36.x.0/24 ($0 \leq x \leq 15$) are configured on the loopback interfaces from 0 to 15 on the router R1. To shorten the output of the show ip bgp command, we will only advertise the first two connected routes 70.36.0.0/24 and 70.36.1.0/24 with the network command. However, the aggregate prefix 70.36.0.0/20 covers overall sixteen more-specific routes in order to follow our real-world scenario.

Below is an initial configuration of all three routers.

R1

```
router bgp 3695
  bgp log-neighbor-changes
  network 70.36.0.0 mask 255.255.255.0
  neighbor 12.0.0.2 remote-as 11260
```

R2


```
router bgp 11260
  bgp log-neighbor-changes
  neighbor 12.0.0.1 remote-as 3695
  neighbor 12.0.1.1 remote-as 6453
```

R3

```
router bgp 6453
  bgp log-neighbor-changes
  neighbor 12.0.1.2 remote-as 11260
```

BGP Route Aggregation with Static Discard Route

Firstly, we are going to create an aggregate address with a static discard route 70.36.0.0/20 pointing to a null interface. This is old-fashioned method but still used. The discard static route 70.36.0.0/20 configured on a router R1 makes the router to discard any packet that matches the route. However, as long as there are more specific (longer prefix) working routes in a routing table of the router R1, packets matching these routes are not discarded. The BGP tables



Picture 5 - Network Topology

BGP Route Aggregation

of R2 and R3 routers are injected with the network command configured on R1 router, matching the static discard route. Firstly, we are going to create an aggregate address with a static discard route 70.36.0.0/20 pointing to a null interface. This is old-fashioned method but still used. The discard static route 70.36.0.0/20 configured on a router R1 makes the router to discard any packet that matches the route. However, as long as there are more specific (longer prefix) working routes in a routing table of the router R1, packets matching these routes are not discarded. The BGP tables of R2 and R3 routers are injected with the network command configured on R1 router, matching the static discard route.

```
router bgp 3695
  bgp log-neighbor-changes
  network 70.36.0.0 mask 255.255.240.0
  neighbor 12.0.0.2 remote-as 11260
!
ip route 70.36.0.0 255.255.240.0 Null0
```

The BGP table of the router R3 is shown in Picture 6.

```
R3#show ip bgp
BGP table version is 34, local router ID is 12.0.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
               x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

   Network          Next Hop           Metric LocPrf Weight Path
  *> 70.36.0.0/20    12.0.1.2           0      11260  3695  i
R3#
```

Picture 6 - BGP Table of R3

BGP Route Aggregation with Aggregate-address Command

Now, let's make the router R1 advertise the aggregate prefix 70.36.0.0/20 to its BGP neighbor R2. The aggregate address is advertised to a neighbor as long as it represents at least one part of the aggregate address in the BGP table of a router. The parts are called components or the contributing routes and represent more specific matches for the aggregated route. We will inject a single route 70.36.0.0/24 into the BGP table of R1 with the network command.

```
router bgp 3695
  bgp log-neighbor-changes
  network 70.36.0.0 mask 255.255.255.0
  aggregate-address 70.36.0.0 255.255.240.0
  neighbor 12.0.0.2 remote-as 11260
```



NOTE: A component route can be injected into the BGP table of a router via the network command, redistributed from the IGP protocols or learned by another BGP peer.

BGP Route Aggregation

The BGP table of the router R3 is shown on Picture 7. The BGP table of R3 contains the aggregated prefix along with a more-specific route 70.36.0.0/24.

```
R3#show ip bgp
BGP table version is 71, local router ID is 12.0.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
               x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

   Network          Next Hop          Metric LocPrf Weight Path
*> 70.36.0.0/24     12.0.1.2          0 11260 3695 i
*> 70.36.0.0/20     12.0.1.2          0 11260 3695 i
R3#
```

Picture 7 - BGP Table of R3

Option summary-only

By default, all more-specific routes summarized by the aggregate route are advertised. To advertise only the aggregate route and suppress the advertisement of all the component routes, the keyword `summary-only` can be used.

```
router bgp 3695
  bgp log-neighbor-changes
  network 70.36.0.0 mask 255.255.255.0
  aggregate-address 70.36.0.0 255.255.240.0 summary-only
  neighbor 12.0.0.2 remote-as 11260
```

The suppressed route 70.36.0.0/24 is shown in the BGP table of the router R1 and is marked as suppressed s>.

```
R1#show ip bgp
BGP table version is 10, local router ID is 70.36.15.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
               x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

   Network          Next Hop          Metric LocPrf Weight Path
s> 70.36.0.0/24     0.0.0.0           0          32768 i
*> 70.36.0.0/20     0.0.0.0           0          32768 i
R1#
```

Picture 8 - BGP Table of R1

R3's BGP table now contains only the aggregated route 70.36.0.0/20.

```
R3#show ip bgp
BGP table version is 4, local router ID is 12.0.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
               x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

   Network          Next Hop          Metric LocPrf Weight Path
*> 70.36.0.0/20     12.0.1.2          0 11260 3695 i
R3#
```

Picture 9 - BGP Table of R3

Option suppress-map

Suppress map defines components that should not be advertised. The following configuration advertises the aggregate route 70.36.0.0/20 along with a more specific route 70.36.1.0/24 while the route 70.36.0.0/24 is not advertised.

BGP Route Aggregation

```
router bgp 3695
  bgp log-neighbor-changes
  network 70.36.0.0 mask 255.255.255.0
  network 70.36.1.0 mask 255.255.255.0
  aggregate-address 70.36.0.0 255.255.240.0 sup-
press-map my_sup_map
  neighbor 12.0.0.2 remote-as 11260
!
ip prefix-list my_sup_list seq 10 permit 70.36.0.0/24
!
!
route-map my_sup_map permit 10
  match ip address prefix-list my_sup_list
```

The component route 70.36.0.0/24 is now marked as suppressed in the BGP table of R1 while another component route 70.36.1.0/24 is marked as active and valid route along with the aggregate route 70.36.0.0/20.

```
R1#show ip bgp
BGP table version is 17, local router ID is 70.36.15.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
               x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

   Network          Next Hop          Metric LocPrf Weight Path
s>  70.36.0.0/24    0.0.0.0           0         32768 i
*>  70.36.0.0/20    0.0.0.0           0         32768 i
*>  70.36.1.0/24    0.0.0.0           0         32768 i
R1#
```

Picture 10 - BGP Table of R1

The router R3's BGP table contains the aggregate route along with the component route 70.36.1.0/24.

```
R3#show ip bgp
BGP table version is 14, local router ID is 12.0.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
               x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

   Network          Next Hop          Metric LocPrf Weight Path
*>  70.36.0.0/20    12.0.1.2           0    11260 3695 i
*>  70.36.1.0/24    12.0.1.2           0    11260 3695 i
R3#
```

Picture 11 - BGP Table of R3

Option unsuppress-map

If a subset of the suppressed routes needs to be made available, we can unsuppress these routes on a per neighbor basis using the neighbor unsuppress-map command. The following configuration extracts a more specific route 70.36.1.0/24 from the aggregate route on the router R1 and advertises both routes to R2.

```
router bgp 3695
  bgp log-neighbor-changes
  network 70.36.0.0 mask 255.255.255.0
  network 70.36.1.0 mask 255.255.255.0
  aggregate-address 70.36.0.0 255.255.240.0 summa-
ry-only
  neighbor 12.0.0.2 remote-as 11260
```

BGP Route Aggregation

```
neighbor 12.0.0.2 unsuppress-map my_unsup_map
!  
ip prefix-list my_unsup_list seq 10 permit 70.36.1.0/24  
!  
route-map my_unsup_map permit 10  
  match ip address prefix-list my_unsup_list
```

The component routes 70.36.0.0/24 and 70.36.1.0/24 are suppressed by the aggregate route 70.36.0.0/20. The route 70.36.1.0/24 is unsuppressed for neighbor 12.0.0.2.

```
R1#show ip bgp  
BGP table version is 6, local router ID is 70.36.15.1  
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,  
               r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,  
               x best-external, a additional-path, c RIB-compressed,  
Origin codes: i - IGP, e - EGP, ? - incomplete  
RPKI validation codes: V valid, I invalid, N Not found  
  
   Network          Next Hop          Metric LocPrf Weight Path  
s>  70.36.0.0/24    0.0.0.0           0         32768 i  
*> 70.36.0.0/20    0.0.0.0           0         32768 i  
s>  70.36.1.0/24    0.0.0.0           0         32768 i  
R1#
```

Picture 12 - BGP Table of R1

The R3's BGP table displays both the aggregate route 70.36.0.0/20 and a component route 70.36.1.0/24.

```
R3#show ip bgp  
BGP table version is 30, local router ID is 12.0.1.1  
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,  
               r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,  
               x best-external, a additional-path, c RIB-compressed,  
Origin codes: i - IGP, e - EGP, ? - incomplete  
RPKI validation codes: V valid, I invalid, N Not found  
  
   Network          Next Hop          Metric LocPrf Weight Path  
*>  70.36.0.0/20    12.0.1.2          0         11260 3695 i  
*>  70.36.1.0/24    12.0.1.2          0         11260 3695 i  
R3#
```

Picture 13 - BGP Table of R3

Option attribute-map

Attributes are inherited from component routes. If we need to remove attributes or set our own attributes to the aggregate route, we will use the attribute-map. The configuration below on R1 router advertises the aggregate route 70.36.0.0/20 with the community 3695:500 towards R2 router.

```
router bgp 3695  
  bgp log-neighbor-changes  
  network 70.36.0.0 mask 255.255.255.0  
  network 70.36.1.0 mask 255.255.255.0  
  aggregate-address 70.36.0.0 255.255.240.0 as-set  
  summary-only attribute-map Set_Attribute  
  neighbor 12.0.0.2 remote-as 11260  
  neighbor 12.0.0.2 send-community  
!  
route-map Set_Attribute permit 10  
  set community 3695:500
```

BGP Table of R2 shows community 3695:500.

```
R2#show ip bgp 70.36.0.0  
BGP routing table entry for 70.36.0.0/20, version 10  
Paths: (1 available, best #1, table default)  
  Advertised to update-groups:  
    1  
  Refresh Epoch 1  
  3695, (aggregated by 3695 70.36.15.1)  
  12.0.0.1 from 12.0.0.1 (70.36.15.1)  
  Origin IGP, metric 0, localpref 100, valid, external, best  
  Community: 3695:500  
  rx pathid: 0, tx pathid: 0x0  
R2#
```

Picture 14 - BGP Table of R2

BGP Route Aggregation

Option advertise-map

If only a subset of the component routes should be used to create the aggregate attribute, the component routes can be defined by the advertise-map. The aggregate route then inherits AS_set only from these components. Consider the following configuration on R1.

```
router bgp 3695
  bgp log-neighbor-changes
  network 70.36.0.0 mask 255.255.255.0
  network 70.36.1.0 mask 255.255.255.0
  aggregate-address 70.36.0.0 255.255.240.0 summary-only advertise-map No_advertise
  neighbor 12.0.0.2 remote-as 11260
!
route-map No_advertise permit 10
  match ip address 1
!
access-list 1 permit 70.36.2.0 0.0.1.255
```

The access-list 1 matches a component route 70.36.2.0/23. The route-map No_advertise matches the access-list1. Route 70.36.2.0/23 is not populated in the BGP table of R1 as the route is not announced with the network command. Remember, the aggregate address is advertised to a neighbor as long as it is a part of the aggregate address in the BGP table of a router. For this reason, neither the aggregate route nor any component routes are advertised to the R2 router.

Option as-set

Now, delete aggregate-address command from R1 and configure R2 to advertise only the aggregate route 70.36.0.0/20. The R1 router now advertises routes 70.36.0.0/24 and 70.36.1.0/24.

```
router bgp 3695
  bgp log-neighbor-changes
  network 70.36.0.0 mask 255.255.255.0
  network 70.36.1.0 mask 255.255.255.0
  neighbor 12.0.0.2 remote-as 11260
```

The R2 router is configured to create the aggregate route 70.36.0.0/20 and suppress all the component routes.

```
router bgp 11260
  bgp log-neighbor-changes
  aggregate-address 70.36.0.0 255.255.240.0 summary-only
  neighbor 12.0.0.1 remote-as 3695
  neighbor 12.0.1.1 remote-as 6453
```

The brief inspection of R3's BGP table reveals that the aggregate route is populated there. However, the AS PATH for the aggregate route 70.36.0.0/24 is 11260 which means that the aggregate address is originating from R2.

BGP Route Aggregation

```
R3#show ip bgp
BGP table version is 19, local router ID is 12.0.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
               x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

   Network        Next Hop           Metric LocPrf Weight Path
*> 70.36.0.0/20    12.0.1.2           0             0 11260 i
R3#
```

Picture 15 - BGP Table of R3

The R1 router has installed the aggregate address 70.36.0.0/24 into its BGP table as it originates from R2 with AS_PATH 11260. If any component routes of 70.36.x.0/24 become unreachable on router R1, packets matching less specific aggregate route 70.36.0.0/20 are still sent from R1 to R2. The R2 router discards the packets as it previously removed a component route from its BGP table.

```
R1#show ip bgp
BGP table version is 20, local router ID is 70.36.15.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
               x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

   Network        Next Hop           Metric LocPrf Weight Path
*> 70.36.0.0/24    0.0.0.0            0             32768 i
*> 70.36.0.0/20    12.0.0.2           0             0 11260 i
*> 70.36.1.0/24    0.0.0.0            0             32768 i
R1#
```

Picture 16 - BGP Table of R1

The keyword `as-set` appended to an `aggregate-address` command generates AS set path information for aggregate address. AS set is a list of AS numbers collected from all component routes that are

part of the aggregate address. The option `as-set` allows AS path loop detection for BGP, as a router does not install a route into its BGP table when it detects its own AS in the AS_PATH attribute update.

```
router bgp 11260
  bgp log-neighbor-changes
  aggregate-address 70.36.0.0 255.255.240.0 summary-only as-set
  neighbor 12.0.0.1 remote-as 3695
  neighbor 12.0.1.1 remote-as 6453
```

Adding the `as-set` option ensures that the aggregate route 70.36.0.0/20 will be shown as originating from R1(AS3695).

```
R3#show ip bgp
BGP table version is 20, local router ID is 12.0.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
               x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

   Network        Next Hop           Metric LocPrf Weight Path
*> 70.36.0.0/20    12.0.1.2           0             0 11260 3695 i
R3#
```

Picture 17 - BGP Table of R3

The aggregate route 70.36.0.0/20 is removed from a R1's BGP table as it detects its own AS 3695 in the update message (Picture 18).

BGP Route Aggregation

```
R1#show ip bgp
BGP table version is 21, local router ID is 70.36.15.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
               x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

   Network        Next Hop           Metric LocPrf Weight Path
*> 70.36.0.0/24    0.0.0.0             0         32768 i
*> 70.36.1.0/24    0.0.0.0             0         32768 i
R1#
```

Picture 18 - BGP Table of R1

Conclusion

Route Aggregation is a universal method used to supersede a set of routes by a single general route, widely used for intra-domain and inter-domain routing purpose. In this guide we have explained the purpose and the benefits of BGP Route Aggregation using a real-world scenario and provided configuration steps explaining various configuration options. Carefully planned and deployed, BGP Route Aggregation reduces the number of BGP entries that have to be stored and exchanged with other BGP peers. It curbs the growth of global routing table size and conserves network bandwidth, helping to increase Internet routing stability by hiding route flaps.



This ebook was brought to you by [Noction](#).

Noction Intelligent Routing Platform enables enterprises and service providers to maximize end-to-end network performance and safely reduce infrastructure costs. The platform evaluates critical network performance metrics in real-time and responds quickly by automatically rerouting traffic through a better path to avoid outages and congestion.

Request a free trial today and see how IRP can boost your network performance.

[Start a Free Trial](#)