# RINGCENTRAL CUSTOMER CLOUDCONNECT ROUTE BGP PEERING CONFIGURATION

Customers that establish a direct connection with Ring Central must set up a routing environment whereby IP network routes to certain Ring Central address spaces direct traffic across the direct connection. Care must be taken to remove or deactivate these routes when the CloudConnect link to RingCentral becomes inactive. This can be handled using one of two mechanisms:

1. The customer sets up an e-BGP peering relationship with Ring Central and uses an interior routing protocol such as OSPF or EIGRP to redistribute the learned BGP routes through their network.
2. The customer sets up a group of static routes and uses monitoring tools to control their activation/deactivation in the RingCentral CloudConnect edge. In the Palo Alto firewall environment this is referred to as Path Monitoring. In the Cisco environment you use the IP SLA mechanism. Other platforms may have differing names for the same feature. The customer then redistributes the static route throughout their network using an interior routing protocol such as OSPF or EIGRP.

The code examples given in this document assume the following:

- RingCentral CloudConnect network link is on the 100.64.1.16/30 subnet
- Interface GigabitEthernet0/1 is used for the RingCentral CloudConnect link
- The RingCentral side of the RingCentral CloudConnect link is addressed at 100.64.1.18/30
- The customer side of the RingCentral CloudConnect link (interface GigabitEthernet0/1) is addressed at 100.64.1.17/30
- The link to the customer's interior address space is on interface GigabitEthernet0/0
- The customer BGP AS Number used for the RingCentral CloudConnect link is 65001 (if using BGP)

## E-BGP CONFIGURATION

Route announcements in e-BGP flow from RingCentral to the customer. The customer does NOT announce any routes to RingCentral except in very special conditions. RingCentral Network Engineering will advise the customer if route announcements are needed.

The following configuration shows a standard Cisco e-BGP peering session set up from an example customer's router to RingCentral.

```
! Please note that the prefix 'PFX-' is used to denote a Cisco prefix-list.
!
! ========================================
! ==> PFX-None - Block all announcements <==
! ========================================
!
no ip prefix-list PFX-None
ip prefix-list PFX-None seq 100 deny 0.0.0.0/0 le 32
!
```

```
! ============================================================================
! ==> PFX-Ringcentral-InbRt - Allow any RingCentral network or subnet to be   <==
! ==>                         announced.  InbRt refers to routes announced to <==
! ==>                         the customer.                                   <==
! ============================================================================
!
no ip prefix-list PFX-RingCentral-InbRt
ip prefix-list PFX-RingCentral-InbRt seq 800 permit 66.81.240.0/20 le 32
ip prefix-list PFX-RingCentral-InbRt seq 900 permit 80.81.128.0/20 le 32
ip prefix-list PFX-RingCentral-InbRt seq 1000 permit 103.44.68.0/22 le 32
ip prefix-list PFX-RingCentral-InbRt seq 1100 permit 104.245.56.0/21 le 32
ip prefix-list PFX-RingCentral-InbRt seq 1200 permit 185.23.248.0/22 le 32
ip prefix-list PFX-RingCentral-InbRt seq 1300 permit 192.209.24.0/21 le 32
ip prefix-list PFX-RingCentral-InbRt seq 1400 permit 199.255.120.0/22 le 32
ip prefix-list PFX-RingCentral-InbRt seq 1500 permit 199.68.212.0/22 le 32
ip prefix-list PFX-RingCentral-InbRt seq 1600 permit 208.87.40.0/22 le 32
! --- block all others explicitly
ip prefix-list PFX-RingCentral-InbRt seq 9999 deny 0.0.0.0/0 le 32
!
! ============================================================================
! ==> PFX-RingCentral-LineTest - Used to validate connection and allow tests <==
! ==>                            without routing production traffic over the <==
! ==>                            dedicated link                              <==
! ============================================================================
!
no ip prefix-list PFX-RingCentral-LineTest
! --- Allow Pingable Line Test Address
ip prefix-list PFX-RingCentral-LineTest seq 9990 permit 199.255.120.184/32
! --- Allow Network Probe Server Address
ip prefix-list PFX-RingCentral-LineTest seq 9991 permit 199.68.213.101/32
! --- Deny all others
ip prefix-list PFX-RingCentral-LineTest seq 9999 deny 0.0.0.0/0 le 32
!
! ============================================================================
! ==> RM-SetLocalPref-Primary - Used to indicate that routes coming in via this
! ==>                           link should be the primary pathway to them.
! ==> RM-SetLocalPref-Backup  - Used to indicate that routes coming in via this
! ==>                           link should be the backup pathway to them.
! ==> These local-preference values may be adjusted if needed, but should work
! ==> in almost all cases.
! ============================================================================
!
route-map RM-SetLocalPref-Primary permit 100
 set local-preference 1000
!
route-map RM-SetLocalPref-Backup permit 100
 set local-preference 900
!
!############################################################################
! ==> BGP Protocol
!
! Establish a standard e-BGP peering session.  This assumes your local AS
! number is 65001, the network link to RingCentral is on the 100.64.1.16/30
! network, Customer WAN interface address is 100.64.1.17, and RingCentral's
! interface address is 100.64.1.18.
!
! The RingCentral neighbor is created in the shutdown state so no BGP peering
! will occur unless a BGP session is manually started locally by removing the
! 'neighbor 100.64.1.18' shutdown statement.
!
router bgp 65001
 bgp log-neighbor-changes
 !
 ! Output a BGP announcement to all peers and interior that we originate the
 ! 100.64.1.16/30 network segment.
 ! This is critical.
 !
 network 100.64.1.16 mask 255.255.255.252
 !
 ! Now set up the peering relationship with RingCentral.  The customer never
 ! announces any routes to RingCentral.  The customer is initially set up to
 ! only accept an announcement of the linetest addresses so that we can safely
 ! turn up the peer for testing without impacting any production call(s)
 ! already in use.
 !
 neighbor 100.64.1.18 remote-as 40627
 neighbor 100.64.1.18 shutdown
 neighbor 100.64.1.18 soft-reconfiguration inbound
 neighbor 100.64.1.18 prefix-list PFX-RingCentral-LineTest in
 neighbor 100.64.1.18 prefix-list PFX-None out
```

```
 !
 ! The following statement should be used to make this link the PRIMARY
 ! RingCentral connection link.  If this should be a secondary link comment out
 ! this line and uncomment the next line.
 !
 neighbor 100.64.1.18 route-map RM-SetLocalPref-Primary in
 !neighbor 100.64.1.18 route-map RM-SetLocalPref-Backup in
 !
 ! For safety sake (and to prevent paranoia) have BGP automatically shut down this
 ! peering link if RingCentral tries to send too many routes.  RingCentral will NEVER
 ! send more than 200 routes - usually the number is less than 10
 !
 neighbor 100.64.1.18 maximum-prefix 200
end
```

This creates the BGP neighbor set up receive ONLY the test routes and leaves it in a disabled state. Once ready to test, run the following commands to enable the BGP peering session with RingCentral:

```
! Activate RingCentral neighbor.
!
router bgp 65001
 no neighbor 100.64.1.18 shutdown
end
```

Once the peer has been activated, run the '`show ip bgp summary`' command to see that the peering session has established properly. The 'State/PfxRcd' column shows the state of the connection or, if the peering session is up, the count of routes the neighbor is advertising.

```
!====================================================
! This shows a neighbor that is DOWN.
!
CISCO-RTR#show ip bgp summary
BGP router identifier 100.64.1.18, local AS number 65001
BGP table version is 7, main routing table version 7

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
100.64.1.18     4        40627       0       0        1    0    0 00:01:12 Idle
!
!====================================================
! and this shows a neighbor that is up and sending routes to us.
!
CISCO-RTR#show ip bgp summary
BGP router identifier 100.64.1.17, local AS number 65001
BGP table version is 8, main routing table version 8
3 network entries using 408 bytes of memory
3 path entries using 168 bytes of memory
1/1 BGP path/bestpath attribute entries using 128 bytes of memory
1 BGP AS-PATH entries using 24 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 728 total bytes of memory
2 received paths for inbound soft reconfiguration
BGP activity 11/8 prefixes, 11/8 paths, scan interval 60 secs

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
100.64.1.18     4        40627       3       4        8    0    0 00:00:08        1
```

The '`show ip bgp neighbor x.x.x.x routes`' command can be used to see the routes that have been received **AND inserted into the routing table** (controlled by 'prefix-list in' clause).

```
CISCO-RTR#sho ip bgp neighbor 100.64.1.18 routes
BGP table version is 8, local router ID is 100.64.1.17
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
```

```
              r RIB-failure, S Stale, m multipath, b backup-path, x best-external, f RT-Filter
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 199.255.120.184/32
                    100.64.1.18                            0 40627 ?

Total number of prefixes 1
```

You can issue the '`show ip bgp neighbor x.x.x.x received-routes`' command to see the routes that have been received by BGP regardless of whether they have been inserted into the routing table.

```
CISCO-RTR#sho ip bgp neighbor 100.64.1.18 received-routes
BGP table version is 8, local router ID is 100.64.1.17
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, x best-external, f RT-Filter
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 199.255.120.184/32
                    100.64.1.18                            0 40627 ?
*> 199.255.120.128/25
                    100.64.1.18                            0 40627 ?
*> 104.245.57.0/24
                    100.64.1.18                            0 40627 ?

Total number of prefixes 3
```

You can move the connection into production by issuing the following configuration commands:

```
!
router bgp 65001
 no neighbor 100.64.1.18 prefix-list PFX-RingCentral-LineTest in
 neighbor 100.64.1.18 prefix-list PFX-RingCentral-InbRt in
end
```

Follow this configuration change by issuing the following CLI command (please note that you **do NOT use** the 'neighbor' keyword):

```
CISCO-RTR#clear ip bgp 100.64.1.18 soft-reconfiguration in
```

An MPLS based CE router will automatically redistribute the routes accepted from RingCentral into the MPLS link and the rest of the customer's routers as both neighbors are e-BGP links. A non-MPLS based router may need to redistribute BGP into an interior routing protocol in order to propagate the routes to the rest of the network.

```
!
! Configure your interior routing protocol, EIGRP in this sample, to redistribute
! all learned BGP routes.
!
router eigrp 1
 network 172.16.0.0
 redistribute bgp 65001 metric 10000 100 255 1 1500
exit
```

Remember to save your work and back up the resulting configuration.

## ADVANCED BGP ROUTING CONTROL - COMMUNITIES

**Note: These values have not yet been finalized.**

Upon customer request, RingCentral will tag route advertisements with a community string to identify the originating RingCentral Data Center. This allows an advanced customer or carrier routing engineer the opportunity to utilize route-maps to fine-tune traffic flow between the customer and RingCentral.

The communities in use are:

| Community Value | Location |
|---|---|
| **Americas** | |
| 40627:9000 | SJC – San Jose, CA, USA |
| 40627:9001 | IAD01 – Vienna, VA, USA |
| 40627:9002 | DFW – Dallas/Fort Worth, TX, USA |
| 40627:9003 | ORD – Chicago, IL, USA |
| 40627:9004 | ATL – Atlanta, GA, USA |
| 40627:9005 | JFK – New York, NY, USA |
| 40627:9006 | MIA – Miami, FL, USA |
| 40627:9007 | YYZ – Toronto, CAN |
| 40627:9008 | SEA – Seattle, WA, USA |
| 40627:9009 | LAX – Los Angeles, CA, USA |
| 40627:9010 | IAD02 – Ashburn, VA, USA |
| 40627:9016 | GIG – Rio de Janeiro, BRA |
| **EMEA** | |
| 40627:9200 | AMS – Amsterdam, NLD |
| 40627:9201 | ZRH – Zurich, CHE |
| 40627:9202 | LHR – London, GBR |
| **APAC** | |
| 40627:9300 | SIN – Singapore, SGP |
| 40627:9301 | SYD – Sydney, AUS |
| 40627:9302 | NRT – Tokyo, JPN |

Be extremely careful when doing route fine-tuning. It is easy to inadvertently set up a 'routing loop'. If you are not completely sure of your configuration please contact RingCentral Customer Engineering for assistance and confirmation of your configuration.

## DYNAMIC (FLOATING) STATIC ROUTES

**Please note that the use of BGP peering is preferred as it allows RingCentral to dynamically alter traffic flow in real-time without customer interaction.**

The following Cisco configuration snippets show setting up static routes across the RingCentral CloudConnect link with IP SLA tracking to disable them if they fail. The interior routing shown is a stub of a representative EIGRP configuration that is set to redistribute all active static routes which are marked with an attached tag value of 9999.

```
! Create an SLA rule to ping out the Ring Central Direct Connect interface to the
! Ring Central line test address with DSCP marking EF.  Ping every 10 seconds.
!
ip sla 1
 icmp-echo 199.255.120.184 source-interface GigabitEthernet0/1
 tos 184
 frequency 10
!
```

```
! Schedule the test to run all the time starting now.
!
ip sla schedule 1 life forever start-time now
!
! Create a tracking object attached to this rule.
!
track 8 ip sla 1 reachability
!
! Create the routes to be marked as active ONLY if the SLA test shows the linetest
! node as reachable via the interface.  Tag these routes with an attached tag value
! of 9999 so that they will be redistributed into the interior routing protocol.
!
ip route 66.81.240.0 255.255.240.0 GigabitEthernet0/1 100.64.1.18 track 8 tag 9999
ip route 80.81.128.0 255.255.240.0 GigabitEthernet0/1 100.64.1.18 track 8 tag 9999
ip route 103.44.68.0 255.255.252.0 GigabitEthernet0/1 100.64.1.18 track 8 tag 9999
ip route 104.245.56.0 255.255.248.0 GigabitEthernet0/1 100.64.1.18 track 8 tag 9999
ip route 185.23.248.0 255.255.252.0 GigabitEthernet0/1 100.64.1.18 track 8 tag 9999
ip route 192.209.24.0 255.255.248.0 GigabitEthernet0/1 100.64.1.18 track 8 tag 9999
ip route 199.68.212.0 255.255.252.0 GigabitEthernet0/1 100.64.1.18 track 8 tag 9999
ip route 199.255.120.0 255.255.252.0 GigabitEthernet0/1 100.64.1.18 track 8 tag 9999
ip route 208.87.40.0 255.255.252.0 GigabitEthernet0/1 100.64.1.18 track 8 tag 9999
!
! Set up a route-map to restrict redistribution of static routes into EIGRP to be
! only those routes with an attached tag of 9999.
!
route-map RM-Redis-StaticToEigrp permit 10
 match tag 9999
!
! Configure your interior routing protocol, EIGRP in this sample, to redistribute
! all active static routes passed by the indicated route-map.
!
router eigrp 1
 network 172.16.0.0
 redistribute static route-map RM-Redis-StaticToEigrp metric 10000 100 255 1 1500
exit
```

This configuration will route all traffic destined for RingCentral out of the configured interface so long as the RingCentral line-test server can be pinged over the direct connection link.  Note that in actual use, RingCentral may specify a different address for testing.

Remember to save and back up the resulting configuration.

## S-NAT CONFIGURATION

Make sure that the router is set up to perform S-NAT overload on internal traffic going out through the RingCentral CloudConnect interface.

```
! From the NAT perspective Ring Central is an Outside network.
!
interface Gi0/1
 description CloudConnect Link to Ring Central
 ip address 100.64.1.17 255.255.255.252
 ip nat outside
!
! From the NAT perspective the customer's interior is an Inside network.
!
interface Gi0/0
 description To Customer Networks
 ip address x.x.x.x m.m.m.m
 ip nat inside
!
! Define an access list that matches everything except the link network addresses.
!
! Critical, everything going out the RingCentral interface EXCEPT traffic with a source
! address of the link interface should be NATted.  If you don't prevent this from being
! NATted BGP will not be able to form a peering relationship.
!
! If any other addresses are advertised on by the customer's BGP session, those networks
! should also be excluded from NATting.
!
ip access-list extended ACL-NatInside
 deny ip 100.64.1.16 0.0.0.3 any
 permit ip any any
!
! Using route-map NAT syntax so that we can support multiple NAT interfaces in the
! event that functionality is needed now or in the future.
!
route-map RM-NatRingCentral
 ! Match traffic that is supposed to be NATted
 match ip address ACL-NatInside
 ! And make sure it is going out the RingCentral Cloudconnect interface
 match interface GigabitEthernet0/1
!
ip nat inside source route-map RM-NatRingCentral interface GigabitEthernet0/1 overload
```

Extremely large customers or customers with multiple BGP links to RingCentral may need to utilize multiple Source NAT addresses and/or otherwise utilize addresses other than the RingCentral CloudConnect interface.

```
! From the NAT perspective Ring Central is an Outside network.
!
interface Gi0/1
 description CloudConnect Link to Ring Central IAD02
 ip address 100.64.1.17 255.255.255.252
 ip nat outside
!
interface Gi0/2
 description CloudConnect Link to Ring Central SJC
 ip address 100.64.0.17 255.255.255.252
 ip nat outside
!
! From the NAT perspective the customer's interior is an Inside network.
!
interface Gi0/0
 description To Customer Networks
 ip address x.x.x.x m.m.m.m
 ip nat inside
!
! Define a NAT pool with the address range to use as source addresses.  The netmask must enclose
! the pool addresses such that they are NOT the first or last address within the block.
!
! Critical, the NAT pool addresses/networks must be announced over the BGP links to RingCentral
! in order for RingCentral to know how to route return traffic.
!
ip nat pool NTPL-1 100.64.17.16 100.64.17.17 netmask 255.255.255.224
!
! Define an access list that matches everything except the link network addresses and the
```

```
! NAT Pool addresses.
!
! Critical, everything going out the RingCentral interface EXCEPT traffic with a source
! address of the link interface should be NATted.  If you don't prevent this from being
! NATted BGP will not be able to form a peering relationship.
!
! If any other addresses are advertised by the customer's BGP session, those networks
! should also be excluded from NATting.
!
ip access-list extended ACL-NatInside-All
 deny ip 100.64.1.16 0.0.0.3 any
 deny ip 100.64.0.16 0.0.0.3 any
 deny ip 100.64.17.16 0.0.0.1 any
 permit ip any any
!
! Using source access-list NAT syntax since we want the same NAT source addresses used
! regardless of the egress interface.
!
ip nat inside source list ACL-NatInside-All pool NTPL-1 overload
```

Remember to save and back up the resulting configuration.