



Cisco Support Community Expert Series Webcast

Border Gateway Protocol (BGP) Fundamentals and Troubleshooting

Vinit Jain, CCIE Security, Data Center, SP, and R&S

January 12, 2016

Upcoming Events

<https://supportforums.cisco.com/expert-corner/events>

Cisco Support Community
Community Directory Expert Corner Community Corner Solutions

Monica Lulus 2179

Search

Home Expert Corner Events Language: English - Contact Us Help Follow Us -

Events

Ask the Experts Webcasts

Cisco experts engage in discussions with you, our members, on specific networking issues. Each event runs for a two-week period.

Ask the Expert: Border Gateway Protocol (BGP) Fundamentals and Troubleshooting
Welcome to this Cisco Support Community Ask the Expert conversation. This is an opportunity to learn and ask questions about how to configure and troubleshoot and design network using Border Gateway Protocol with Cisco expert Vinit Jain. Ask...
Begins January 12, 2016
Add to Calendar

Troubleshooting Cisco Identity Services Engine (ISE)
Welcome to this Cisco Support Community Ask the Expert conversation. This is an opportunity to ask questions about Cisco Identity Service Engine (ISE) to Artem Tkachov and Wojciech Cecot. Ask questions from Monday December 14 to Wednesday...
Ended December 23, 2015
Read the Q&A

Ask the Expert: Cisco Unified Computing System (UCS) Blades
This session provides an opportunity to learn and ask questions about Cisco Unified Computing System in general and the blades in more specific. Sivakumar will cover questions on design, configuration and troubleshooting. The Cisco Unified...
Ended December 11, 2015
Read the Q&A

Page 1

netBrain
Automate Network Diagramming
Without Any Compromises
SEE A DEMO

LIVE Webcast
BGP Fundamentals and Troubleshooting
JAN 12, 2016 10AM PST
Vinit Jain
Register Now

ManageEngine
NetFlow Analyzer

Become an Event Top Contributor

Participate in Live Interactive Technical Events and much more
<http://bit.ly/1jll93B>

Cisco Support Community
Community Directory Expert Corner Community Corner Solutions

Home Expert Corner Top Contributors Language: English Contact Us Help Follow Us

Top Contributors

Recognition Program VIPs Spotlight Awards Hall of Fame Events Top Contributors
Expert Interviews

Cisco Designated VIPs

The Cisco Designated VIP program recognizes the top external individual contributors in Cisco's online communities, including the Cisco Support Community (CSC), Cisco Learning Network (CLN) and the Cisco Developers Network (CDN). Cisco Designated VIPs are recognized by their peers for their expertise and tireless contributions, and their abundant participation is vital to community success. With this program, Cisco formally recognizes the positive, valuable influence our top individual members exert on the communities overall.

To learn more, please visit our [FAQ](#)

CISCO DESIGNATED
cisco VIP < 2016 2015 2014 2013 2012 2011 >

 Aaron Harrison Contact Center, IP Telephony	 Aman Sci 2016 IP Telephony
 Ayodeji Oladipo Okanlawon 2016 IP Telephony	 Carlo Poggiarelli 2016 IP Telephony
 Chris Deren 2016 IP Telephony, Contact Center, Unified Communications	 Dan Lukes 2016 Small Business
 George Stefanick 2016 Wireless	 Jon Marshall 2016 LAN, WAN, Firewalling
 Jonathan Schulenberg 2016 IP Telephony	 Karsten Iwen 2016 Firewalling, VPN

Automate Network Diagramming
Without Any Compromise

LEARN MORE

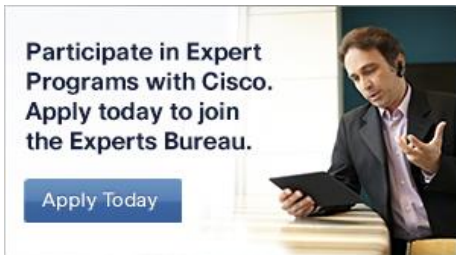
Cisco Support Community
Directory Expert Corner Solutions Community Corner

Home Experts Bureau Top Contributors Leaderboards Knowledge Sharing Voting results Panelizer Hierarchy

Experts Bureau ← Events

Use the Cisco Experts Bureau to find, connect, and follow recognized Subject Matter Experts and the programs they participate in regularly. The Experts Bureau comprises Cisco employees as well as Partners and Customers who have contributed to, or been selected for knowledge sharing programs on the Cisco Support Community, such as Webcasts, Ask the Expert Events, Facebook Forums, Tech-Talks, Meetups, and Blogs.

If you have interest in participating, apply online through this [simple form](#). After applying, a member of the Cisco Support Community team will be in contact with additional details.



<https://supportforums.cisco.com/expert-corner/top-contributors>

Rate Content



Encourage and acknowledge people who generously share their time and expertise

Now your ratings on documents, videos, and blogs count give points to the authors!!!

So, when you contribute and receive ratings you now get the points in your profile.

Help us to recognize the quality content in the community and make your searches easier. Rate content in the community.

<https://supportforums.cisco.com/blog/154746>

Cisco Support Community Expert Series Webcast

Vinit Jain

CCIE Security, Data Center SP and R&S
#22854

Vinit's Twitter – twitter.com/vinugenie

Vinit's Facebook - <https://www.facebook.com/vinu.genie>

Cisco Live Event - BRKRST-3320 (Cisco Live Berlin)
<https://cisco.rainfocus.com/scripts/catalog/cleu16.jsp?search=brkrst-3320>



Question Managers

Brian Dunn

Team Lead HTTS
team



Manish Kumar

CCIE



Ask the Expert Event following the Webcast

Now through January 22nd

<https://supportforums.cisco.com/discussion/12720596/ask-expert-border-gateway-protocol-bgp-fundamentals-and-troubleshooting>



Thank You For Joining Us Today!



If you would like a copy of the presentation slides, click the PDF file link in the chat box on the right or go to:

<https://supportforums.cisco.com/document/12746296/webcast-slides-border-gateway-protocol-bgp-fundamentals-and-troubleshooting>





Submit Your Questions Now!

Use the Q & A panel to submit your questions and the panel of experts will respond.

Please take a moment to complete the survey at the end of the webcast



Cisco Support Community Expert Series Webcast

Border Gateway Protocol (BGP) Fundamentals and Troubleshooting

Vinit Jain, CCIE Security, Data Center, SP, and R&S

January 12, 2016

Agenda

- BGP Fundamentals
 - BGP Refresher
 - BGP neighbor relationships
 - BGP Prefix advertisement
 - BGP Filtering and route policy
 - BGP Community
- BGP Troubleshooting
 - BGP Peering Issues
 - Missing BGP Routes

Polling Question 1


How big is the Internet Routing Table?

- A. ~ 250,000
- B. ~ 500,000
- C. ~ 400,000
- D. ~ 600,000

Routing with BGP

Protocol overview


- Defined in RFC 4271 – Path Vector protocol with high scalability and flexibility
- BGP is primarily used to provide inter-AS connectivity
- No periodic routing updates – scalability
- Neighbors are established over TCP port 179
 - Same AS is iBGP, different AS is eBGP
- Support for multiple address-families (MP-BGP)
 - IPv4, IPv6, VPNv4, MDT, L2VPN, etc



Routing with BGP

Protocol overview

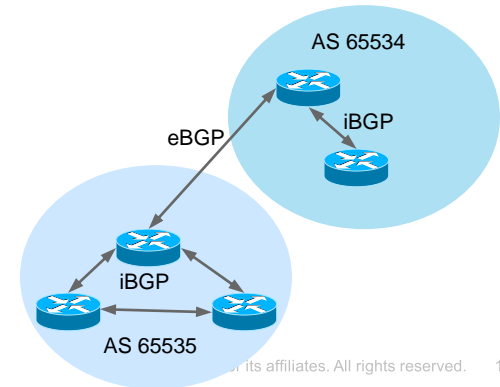
- Path Attributes are associated with each network path to control routing policy
 - Well-known (mandatory, discretionary)
 - Optional (transitive, non-transitive)
- Uses a path selection algorithm to determine which prefix is best
- AS_PATH attribute used for loop prevention



Routing with BGP


Protocol overview

- Route propagation
 - eBGP learned route is advertised to iBGP peers and eBGP peers
 - iBGP learned route is advertised to eBGP peers
- iBGP must have a full mesh of peers (Confederation or RR can be used)
- BGP holds prefixes in its own table (BGP table – show ip bgp)
 - Only the 'best' path is advertised to neighbors or installed in the RIB by default
- eBGP route adv. to iBGP peer – next-hop is the eBGP neighbor
- Message types
 - OPEN, UPDATE, NOTIFICATION, KEEPALIVE, ROUTE REFRESH
 - Keepalive timer is 60 holdtimer is 180
- May require another protocol to provide NH reachability



Routing with BGP

Neighbor Establishment



```
R2#debug ip bgp
```

```
*Mar 23 20:45:18.184: BGP: 10.23.23.3 passive went from Idle to Connect
```

```
*Mar 23 20:45:18.208: BGP: 10.23.23.3 passive went from Connect to OpenSent
```

```
*Mar 23 20:45:18.208: BGP: 10.23.23.3 passive sending OPEN, version 4, my as: 65535, holdtime 180 seconds, ID 202020
```

```
*Mar 23 20:45:18.208: BGP: 10.23.23.3 passive went from OpenSent to OpenConfirm
```

```
*Mar 23 20:45:18.208: BGP: 10.23.23.3 passive went from OpenConfirm to Established
```

```
R2#show tcp brief
```

TCB	Local Address	Foreign Address	(state)
09BAD1E0	10.23.23.2.179	10.23.23.3.21995	ESTAB
070A2910	2.2.2.2.29651	1.1.1.1.179	ESTAB

```
R2#
```

```
R3#debug ip bgp
```

```
*Mar 23 20:45:18.180: BGP: 10.23.23.2 active went from Idle to Active
```

```
*Mar 23 20:45:18.184: BGP: 10.23.23.2 active went from Active to OpenSent
```

```
*Mar 23 20:45:18.184: BGP: 10.23.23.2 active sending OPEN, version 4, my as: 65534, holdtime 180 seconds, ID 3030303
```

```
*Mar 23 20:45:18.204: BGP: 10.23.23.2 active went from OpenSent to OpenConfirm
```


```
*Mar 23 20:45:18.204: BGP: 10.23.23.2 active went from OpenConfirm to Established
```

```
R3#show tcp brief
```

TCB	Local Address	Foreign Address	(state)
09AC7AD0	10.23.23.3.21995	10.23.23.2.179	ESTAB

Routing with BGP

Neighbor Establishment



R2#show ip bgp summary

BGP router identifier 2.2.2.2, local AS number 65535
BGP table version is 1, main routing table version 1

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down
1.1.1.1	4	65535	75	75	1	0	0	01:05:04
10.23.23.3	4	65534	44	44	1	0	0	00:36:25

R2#

R3#show ip bgp summary

BGP router identifier 3.3.3.3, local AS number 65534
BGP table version is 1, main routing table version 1

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down
10.23.23.2	4	65535	45	44	1	0	0	00:36:57

R3#


Interesting fields:

- Neighbor – IP address of the BGP peer
- AS – Autonomous system of the peer
- TblVer – BGP table version advertised to that peer

- InQ/OutQ- Number of messages queued from/to the peer
- UP/Down – Time in the current state
- State/PfxRcd – Current state of the peer and number of prefixes received from the peer


Routing with BGP

Prefix advertisement – getting routes into BGP



BGP prefixes are injected by explicit configuration

- Network statement - **network <prefix> mask <mask>**
 - Prefix/mask needs to match the RIB exactly
 - Does not enable BGP on an interface like IGP's
- Redistribution – **redistribute ospf <PID>**
 - Injects prefixes from the specified protocol
 - Does not inject 0.0.0.0/0
- Aggregate route – **aggregate-address <prefix> <mask>**
 - Component route must exist in BGP
 - Aggregator attribute is added
- Default route – **default-information originate**



Routing with BGP

Prefix advertisement – getting routes into BGP

```
R2#show ip bgp 3.3.3.3
```

```
BGP routing table entry for 3.3.3.3/32, version 2
```

```
Paths: (1 available, best #1, table default)
```

```
Advertised to update-groups:
```

```
2
```

```
Refresh Epoch 1
```

```
65534
```

```
10.23.23.3 from 10.23.23.3 (3.3.3.3)
```

```
Origin IGP, metric 0, localpref 100, valid, external, best
```

```
R2#show ip bgp neighbors 1.1.1.1 advertised-routes
```

```
BGP table version is 2, local router ID is 2.2.2.2
```

```
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
*>	3.3.3.3/32	10.23.23.3	0	0	65534	i



```
R3#show ip bgp 3.3.3.3
```

```
% Network not in table
```

```
R3#conf t
```

```
Enter configuration commands, one per line. End with CNTL/Z.
```

```
R3(config)#router bgp 65534
```

```
R3(config-router)#network 3.3.3.3 mask 255.255.255.255
```

```
R3(config-router)#end
```

```
R3#show ip bgp 3.3.3.3
```

```
BGP routing table entry for 3.3.3.3/32, version 2
```

```
Paths: (1 available, best #1, table default)
```

```
Advertised to update-groups:
```

```
1
```

```
Refresh Epoch 1
```

```
Local
```


```
0.0.0.0 from 0.0.0.0 (3.3.3.3)
```

```
Origin IGP, metric 0, localpref 100, weight 32768, valid, sourced, local,
```

```
best
```

Routing with BGP

Prefix advertisement – Route Policy and filtering



BGP policies determine what prefixes are permitted or denied and can modify path attributes


- Policies are configured per-neighbor, and can be inbound or outbound
- **Neighbor <IP> distribute-list <X> <in/out>**
 - X = Access-list, filters routes in or out
- **Neighbor <IP> prefix-list <X> <in/out>**
 - X = prefix-list name, similar to distribute list (mutually exclusive)
- **Neighbor <IP> filter-list <x> <in/out>**
 - X = “ip as-path access-list”, filters by AS_PATH attribute
- **Neighbor <IP> route-map <x> <in/out>**



X= Route-map, extremely versatile, can modify attributes and filter

Routing with BGP

Prefix advertisement – Route Policy example



```
R2#show ip bgp 3.0.0.0
```

BGP routing table entry for 3.0.0.0/24, version 3

Paths: (1 available, best #1, table default)

Advertised to update-groups:

2

Refresh Epoch 3

65534

10.23.23.3 from 10.23.23.3 (3.3.3.3)

Origin IGP, metric 0, localpref 100, valid, external, best

```
R2#show ip bgp 3.0.1.0
```

BGP routing table entry for 3.0.1.0/24, version 7

Paths: (1 available, best #1, table default)

Advertised to update-groups:

2

Refresh Epoch 3

65534

10.23.23.3 from 10.23.23.3 (3.3.3.3)

Origin IGP, metric 0, localpref 200, valid, external, best

Router 2

```
router bgp 65535
  bgp router-id 2.2.2.2
  bgp log-neighbor-changes
  neighbor 1.1.1.1 remote-as 65535
  neighbor 1.1.1.1 update-source Loopback0
  neighbor 10.23.23.3 remote-as 65534
  neighbor 10.23.23.3 route-map local-pref in
  !
  ip forward-protocol nd
  !
  !
  no ip http server
  no ip http secure-server
  !
  !
  ip prefix-list lp-as65534-in seq 5 permit 3.0.1.0/24
  !
  route-map local-pref permit 10
    match ip address prefix-list lp-as65534-in
    set local-preference 200
  !
  route-map local-pref permit 20
  !
```

Routing with BGP

Using policy to influence traffic

Outbound policy to influence incoming traffic


- Neighboring AS can make their own policy which makes this challenging
- AS_PATH pre-pending – Makes a particular path look better or worse to the rest of the world

Inbound policy to influence outgoing traffic

- Weight
- Local Preference
- Community can be set incoming to allow granular control and easier policy management

Routing with BGP

BGP Next-Hop



R2#show ip bgp

BGP table version is 56, local router ID is 2.2.2.2

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
* i 3.0.0.0/24	10.45.45.5	0	100	0	65534 i
*>	10.23.23.3	0	0	65534	i
* i 3.0.1.0/24	10.45.45.5	0	100	0	65534 i
*>	10.23.23.3	0	0	65534	i
* i 3.0.2.0/24	10.45.45.5	0	100	0	65534 i
*>	10.23.23.3	0	0	65534	i
* i 3.0.3.0/24	10.45.45.5	0	100	0	65534 i
*>	10.23.23.3	0	0	65534	i
*>i 111.111.111.0/24	1.1.1.1	0	100	0	i

R3# show ip bgp

BGP table version is 59, local router ID is 3.3.3.3

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
* i 3.0.0.0/24	5.5.5.5	0	100	0	i
*>	0.0.0.0	0	32768		i
* i 3.0.1.0/24	5.5.5.5	0	100	0	i
*>	0.0.0.0	0	32768		i
* i 3.0.2.0/24	5.5.5.5	0	100	0	i
*>	0.0.0.0	0	32768		i
* i 3.0.3.0/24	5.5.5.5	0	100	0	i
*>	0.0.0.0	0	32768		i
* i 111.111.111.0/24	10.45.45.4	0	100	0	65535 i
*>	10.23.23.2	0	65535		i

Routing with BGP

BGP Next-Hop

```
R1#show ip route 3.0.2.0
% Network not in table
R1#show ip bgp 3.0.2.0
BGP routing table entry for 3.0.2.0/24, version 0
Paths: (2 available, no best path)
Flag: 0x820
Not advertised to any peer
Refresh Epoch 2
65534
10.23.23.3 (inaccessible) from 2.2.2.2 (2.2.2.2)
Origin IGP, metric 0, localpref 100, valid, internal
rx pathid: 0, tx pathid: 0
Refresh Epoch 2
65534
10.45.45.5 (inaccessible) from 4.4.4.4 (4.4.4.4)
Origin IGP, metric 0, localpref 100, valid, internal
rx pathid: 0, tx pathid: 0
```



```
R1#show ip route 10.23.23.3
% Subnet not in table
```


Routing with BGP

Neighbor x.x.x.x next-hop-self

```
R1#show ip bgp 3.0.2.0
```

```
BGP routing table entry for 3.0.2.0/24, version 5
```

```
Paths: (2 available, best #1, table default)
```

```
Not advertised to any peer
```

```
Refresh Epoch 2
```

```
65534
```

```
  2.2.2.2 (metric 11) from 2.2.2.2 (2.2.2.2)
```

```
    Origin IGP, metric 0, localpref 100, valid, internal, best
```


```
    rx pathid: 0, tx pathid: 0x0
```

```
Refresh Epoch 2
```

```
65534
```


```
  4.4.4.4 (metric 11) from 4.4.4.4 (4.4.4.4)
```

```
    Origin IGP, metric 0, localpref 100, valid, internal  
    rx pathid: 0, tx pathid: 0
```



```
#neighbor 1.1.1.1 next-hop-self
```


Routing with BGP



```
R1#show ip bgp 3.0.2.0
BGP routing table entry for 3.0.2.0/24, version 9
Paths: (2 available, best #2, table default)
Not advertised to any peer
Refresh Epoch 2
65534 65534 65534 65534
 2.2.2.2 (metric 11) from 2.2.2.2 (2.2.2.2)
   Origin IGP, metric 0, localpref 100, valid, internal
   rx pathid: 0, tx pathid: 0
Refresh Epoch 2
65534
 4.4.4.4 (metric 11) from 4.4.4.4 (4.4.4.4)
   Origin IGP, metric 0, localpref 100, valid, internal, best
   rx pathid: 0, tx pathid: 0x0
R1#
```

```
!
hostname R3
!
router bgp 65534
  bgp router-id 3.3.3.3
  bgp log-neighbor-changes
  network 3.0.0.0 mask 255.255.255.0
  network 3.0.1.0 mask 255.255.255.0
  network 3.0.2.0 mask 255.255.255.0
  network 3.0.3.0 mask 255.255.255.0
  neighbor 5.5.5.5 remote-as 65534
  neighbor 5.5.5.5 update-source Loopback0
  neighbor 10.23.23.2 remote-as 65534
  neighbor 10.23.23.2 route-map prepend out
!
!
ip as-path access-list 1 permit ^$
!
!
route-map prepend permit 10
  match as-path 1
  set as-path prepend 65534 65534 65534
!
route-map prepend permit 20
!
```

Validating prefix advertisement



```
R1#show ip bgp neighbors 2.2.2.2 routes
BGP table version is 28, local router ID is 1.1.1.1
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
*>i 3.0.0.0/24	2.2.2.2	0	200	0	65534 i
*>i 3.0.2.0/24	2.2.2.2	0	200	0	65534 i

Total number of prefixes 2

```
R2#show ip bgp neighbors 1.1.1.1 advertised-routes
BGP table version is 80, local router ID is 2.2.2.2
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
*> 3.0.0.0/24	10.23.23.3	0	200	0	65534 i
*> 3.0.2.0/24	10.23.23.3	0	200	0	65534 i

Total number of prefixes 2
R2#

Polling Question 2


Which protocol would you prefer for carrying Internet Routing Table?

- A. OSPF
- B. EIGRP
- C. BGP
- D. IS-IS


Scenario 1 - Failed BGP Peering

Problem Description

- iBGP / eBGP is not establishing
- Newly configured BGP session not coming up
- Session was up before, but not coming up now




Physical Topology



Logical Topology

Failed BGP peering

Configuration



Check

- AS Numbers
- Peering IP
- eBGP Multihop?

```
router bgp 65535
  bgp router-id 2.2.2.2
  bgp log-neighbor-changes
  neighbor 1.1.1.1 remote-as 65535
  neighbor 1.1.1.1 update-source Loopback0
  neighbor 10.23.23.3 remote-as 65534
  . . .
```


```
router bgp 65534
  bgp router-id 3.3.3.3
  bgp log-neighbor-changes
  neighbor 10.23.23.2 remote-as 65535
  . . .
```

Failed BGP peering

Reachability

```
R1# ping 2.2.2.2
Sending 5, 100-byte ICMP Echos to 2.2.2.2,
timeout is 2 seconds:
Packet sent with a source address of 10.12.12.1
!!!!
Success rate is 0 percent (0/5)
```

```
R1# ping 2.2.2.2 source loopback0
Sending 5, 100-byte ICMP Echos to 2.2.2.2, timeout is 2 seconds:
Packet sent with a source address of 1.1.1.1
.....
Success rate is 0 percent (0/5)
```



Failed BGP peering



Verify any Firewall / ACL in path for TCP port 179

```
ASA_FW# sh run access-list
```

```
access-list OUT extended permit icmp any any
access-list OUT extended permit ospf any any
access-list OUT extended permit tcp any any eq telnet
. . . . .
```

```
Rn# sh ip access-list R1_R2
```

```
permit icmp any any
permit ospf any any
permit tcp host 10.12.12.1 eq bgp 2.2.2.2
permit tcp host 10.12.12.1 2.2.2.2 eq bgp
. . . . .
```



Failed BGP peering

Verify TCP session

```
R2#sh tcp brief
```

TCB	Local Address	Foreign Address	(state)
65F19834	2.2.2.2.179	1.1.1.1.46523	ESTAB

Quick test when BGP is down

```
R1#telnet 2.2.2.2 179 /source-interface loopback 0
```

```
Trying 2.2.2.2 ...
```

```
% Destination unreachable; gateway or host down
```

```
R1#
```

- This means BGP Packets are being blocked between R1 and R2

Failed BGP peering

Blocked process in XR

- ✓ Ensure BGP process is in **Run** state.
- ✓ Check for blocked BGP or TCP process on the RP / LC using **show process blocked** command.

```
RP/0/RSP0/CPU0:ASR9010-B# show process bgp
Mon Jun  3 09:47:12.646 EST
                        Job Id: 1040
                        PID: 307494
Executable path: /disk0/iosxr-routing-
4.2.3/bin/bgp
                        Instance #: 1
                        Version ID: 00.00.0000
                        Respawn: ON
                        Respawn count: 1
Max. spawns per minute: 12
                        Last started: Tue May 28 14:35:50
2013
                        Process state: Run
                        Package state: Normal
                        Started on config: default
                        . . . .
```

Failed BGP peering

Show process bgp (contd. Output)

```
RP/0/RSP0/CPU0:ASR9010-B# show process bgp
<snip>
10 2 488K 10 Nanosleep 0:00:02:0004 0:00:00:0847 bgp
1049 13 3 488K 10 Receive 0:00:00:0811 6:36:52:0264 bgp
1049 14 3 488K 10 Condvar 14:56:55:0236 9:07:49:0890 bgp
1049 15 0 488K 10 Condvar 14:56:55:0240 25:09:49:0542 bgp
1049 16 3 488K 10 Running 0:00:00:0000 57:53:33:0110 bgp
1049 17 1 488K 10 Receive 0:00:28:0379 0:00:00:0066 bgp
1049 18 1 488K 10 Mutex 13:15:50:0870 3:31:49:0712 bgp
<snip>
```

- You can also use “show process blocked” to check the blocked processes

Failed BGP peering

Sniffer Capture

Use SPAN to get traffic to your sniffer

- monitor session 1 source interface Te2/4 rx
- monitor session 1 destination interface Te2/2

IOS-XR

- Only supported on ASR-9000
- Use ACLs to control what packets to SPAN

RSPAN

- *“RSPAN has all the features of SPAN, plus support for source ports and destination ports that are distributed across multiple switches, allowing one to monitor any destination port located on the RSPAN VLAN. Hence, one can monitor the traffic on one switch using a device on another switch.”*

Failed BGP peering

Wireshark

*Standard input [Wireshark 1.12.1 (v1.12.1-0-g01b65bf from master-1.12)]

File Edit View Go Capture Analyze Statistics Telephony Tools Internals Help

Filter: Expression... Clear Apply Save

No.	Time	Source	Destination	Protocol	Length	Info
2	2.107306000	10.1.12.2	10.1.12.1	TCP	60	51182->179 [SYN] Seq=0 win=16384 Len=0 MSS=1460
3	2.722406000	10.1.12.1	10.1.12.2	TCP	60	179->51182 [SYN, ACK] Seq=0 Ack=1 Win=16384 Len=0
4	2.784806000	10.1.12.2	10.1.12.1	TCP	60	51182->179 [ACK] Seq=1 Ack=1 Win=16384 Len=0
5	2.847206000	10.1.12.2	10.1.12.1	BGP	116	OPEN Message
6	2.878406000	10.1.12.1	10.1.12.2	BGP	118	OPEN Message, KEEPALIVE Message
7	2.987606000	10.1.12.2	10.1.12.1	BGP	75	NOTIFICATION Message
8	3.034406000	10.1.12.1	10.1.12.2	TCP	60	179->51182 [FIN, PSH, ACK] Seq=65 Ack=84 win=16
9	3.096806000	10.1.12.2	10.1.12.1	TCP	60	51182->179 [ACK] Seq=84 Ack=66 win=16320 Len=0
10	3.769608000	ca:02:0e:e0:00:00	ca:02:0e:e0:00:00	LOOP	60	Reply
11	7.780815000	10.1.12.2	10.1.12.1	TCP	60	51182->179 [FIN, PSH, ACK] Seq=84 Ack=66 win=16
12	7.827615000	10.1.12.1	10.1.12.2	TCP	60	179->51182 [ACK] Seq=66 Ack=85 win=16301 Len=0
13	9.998019000	c2:01:14:b8:00:00	c2:01:14:b8:00:00	LOOP	60	Reply
14	13.477825000	ca:02:0e:e0:00:00	ca:02:0e:e0:00:00	CDP/VTP/DTP/PAGP/LLDP	366	Device ID: R2 Port ID: FastEthernet0/0

Frame 5: 116 bytes on wire (928 bits), 116 bytes captured (928 bits) on interface 0
Ethernet II, Src: ca:02:0e:e0:00:00 (ca:02:0e:e0:00:00), Dst: c2:01:14:b8:00:00 (c2:01:14:b8:00:00)
Internet Protocol Version 4, Src: 10.1.12.2 (10.1.12.2), Dst: 10.1.12.1 (10.1.12.1)
Transmission Control Protocol, Src Port: 51182 (51182), Dst Port: 179 (179), Seq: 1, Ack: 1, Len: 62
Border Gateway Protocol - OPEN Message
Marker: ffffffff
Length: 62
Type: OPEN Message (1)
Version: 4
My AS: 100

```
0000 c2 01 14 b8 00 00 ca 02 0e e0 00 00 08 00 45 c0 .....E.  
0010 00 66 93 1b 40 00 ff 06 bb b1 0a 01 0c 02 0a 01 .f..@.....  
0020 0c 01 c7 ee 00 b3 5b 27 41 b8 a7 30 0c 23 50 18 .....['A..O.#P.  
0030 40 00 62 d8 00 00 ff ff ff ff ff ff ff ff ff @.b.....  
0040 ff ff ff ff ff ff 00 3e 01 04 00 64 00 b4 02 02 .....>.....d....  
0050 02 02 21 02 06 01 04 00 01 00 01 02 02 80 00 02 ..!.....  
0060 02 02 00 02 03 83 01 00 02 02 46 00 02 06 41 04 .....F...A.
```

File: "C:\Users\ADMINI~1\AppData\Local\Temp\..." Profile: Default

Failed BGP peering

Platform Specific Packet Capture Tools

IOS

- ✓ Embedded Packet Capture

6500 / 7600

- ✓ ELAM
- ✓ NETDR Capture
- ✓ MPA (Mini Protocol Analyzer)

ASR9000

- ✓ Network Processor Capture

Nexus (7k, 5k, 3k)

- ✓ Ethalyzer
- ✓ Elam



For your
reference only

Scenario 2 – Missing Routes

Problem Description

- Routes advertised were not learnt on peer router
- Symptoms
 - Traffic Loss / No Traffic for the prefix

Scenario 2 – Missing Routes


Update Filtering

- Types of filters
 - Prefix filters
 - AS_PATH filters
 - Community filters
 - Route-maps
- Applied in/out direction

Scenario 2 – Missing Routes

Update Filters

- Determine which filters are applied to the BGP session
 - Show ip bgp nei x.x.x.x
 - Show run | include neighbor x.x.x.x
- Examine the route and pick out the relevant attributes
 - Show ip bgp y.y.y.y
- Compare the attributes against the filters



```
R1#show ip bgp neigh 2.2.2.2 routes
```

```
Total number of prefixes 0
```

Scenario 2 – Missing Routes

Community Problems

```
R2#show run | begin bgp
router bgp 2
network 10.1.1.0 mask 255.255.255.0 route-map set-community
...
route-map set-community permit 10
    set community 2:2 1:50
```

```
R2#show ip bgp 10.1.1.0
BGP routing table entry for 10.1.1.0/24, version 1660 Paths: (1
available, best #1)
Not advertised to any peer
Local
0.0.0.0 from 0.0.0.0 (2.2.2.2)
Origin IGP, metric 0, localpref 100, weight 32768,
valid, sourced, local, best
Community 2:2 1:50
```

R1 filtering routes based on community, doesn't see anything in their BGP table

Scenario 2 – Missing Routes

Community Problems

```
R2#show run | begin bgp
router bgp 2
network 10.1.1.0 route-map set-community
neighbor 1.1.1.1 remote-as 1
neighbor 1.1.1.1 prefix-list my-agg out
neighbor 1.1.1.1 prefix-list their-agg in
!
ip prefix-list my-agg permit 10.0.0.0/8
ip prefix-list their-agg permit 20.0.0.0/8
!
route-map set-community permit 10
set community 2:2 1:50
```

- Configuration looks Okay – filters okay, route-map okay

Scenario 2 – Missing Routes

Community Problems

- R2 now advertises prefix with community to R1
- But R1 still doesn't see the prefix
 - Since nothing is wrong on R2, so turn attention to R1

```
R1#show run | begin bgp
router bgp 1
neighbor 2.2.2.2 remote-as 2 neighbor 2.2.2.2 route-map R2-in in
neighbor 2.2.2.2 route-map R1-out out
!
ip community-list 1 permit 1:150
!
route-map R2-in permit 10
match community 1
set local-preference 150
```

Scenario 2 – Missing Routes

Community Problems

- Community match on R1 expects 1:150 to be set on prefix
- But R2 is sending 1:50
Typo or miscommunication between operations?
- R2 is also using the route-map to filter
If the prefix does not have community 1:150 set, it is dropped
 - there is no next step in the route-map
- Watch the route-map rules in Cisco IOS – they are basically:
if <match> then <set> and exit route-map
 else if <match> then <set> and exit route-map
 else if <match> then <set> etc...
- Blank route-map line means match everything, set nothing

Scenario 2 – Missing Routes

Debugging with ACL

- If unable to find any config issues, try enabling debugs (conditional / filtered debugs)

```
R1#show access-list 99
```

```
Standard IP access list 99
  permit 10.1.1.0 0.0.0.255
```

```
R1#debug ip bgp 2.2.2.2 update 99
```

```
BGP updates debugging is on for access list 99 for neighbor 2.2.2.2
```

```
4d00h: BGP(0): 2.2.2.2 rcvd UPDATE w/ attr: nexthop 2.2.2.2,
  origin i, metric 0, path 12
```

```
4d00h: BGP(0): 2.2.2.2 rcvd 10.1.1.0/24 -- DENIED due to: route-map;
```

Resources

References

- BGP-4 RFC
 - <https://www.ietf.org/rfc/rfc4271.txt>
- BGP Configuration Guide
 - http://www.cisco.com/c/en/us/td/docs/ios/12_2/ip/configuration/guide/fipr_c/1cfbgp.html
- Troubleshooting BGP
 - <http://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/22166-bgp-trouble-main.html>



Submit Your Questions Now!

Use the Q & A panel to submit your questions and our expert will respond

Ask the Expert Event following the Webcast

Now through January 22nd

<https://supportforums.cisco.com/discussion/12720596/ask-expert-border-gateway-protocol-bgp-fundamentals-and-troubleshooting>



Collaborate within our Social Media

Learn About Upcoming Events



Facebook- <http://bit.ly/csc-facebook>



Twitter- <http://bit.ly/csc-twitter>



You Tube <http://bit.ly/csc-youtube>



Google+ <http://bit.ly/csc-googleplus>



LinkedIn <http://bit.ly/csc-linked-in>



Instagram <http://bit.ly/csc-instagram>



Newsletter Subscription
<http://bit.ly/csc-newsletter>

Cisco has support communities in other languages!

If you speak Spanish, Portuguese, Japanese, Russian or Chinese we invite you to participate and collaborate in your language



Spanish

<https://supportforums.cisco.com/community/spanish>

Portuguese

<https://supportforums.cisco.com/community/portuguese>

Japanese

<https://supportforums.cisco.com/community/csc-japan>

Russian

<https://supportforums.cisco.com/community/russian>

Chinese

<http://www.csc-china.com.cn>



More IT Training Videos and Technical Seminars on the Cisco Learning Network

View Upcoming Sessions Schedule
<https://cisco.com/go/techseminars>

Thank you for participating!

- Redeem your 35% discount offer by entering code: CSC when checking out:

Visit Cisco Press at:

<http://bit.ly/csc-ciscopress-2016>

Cisco Press





Please take a moment to complete the survey

Thank you for Your Time!



CISCO

TOMORROW starts here.