

BGP Table Version

TAC

Document ID: 116511

Contributed by Luc De Ghein, Cisco TAC Engineer.
Sep 30, 2013

Contents

Introduction

Network Diagram

Best Path

Types of Table Versions

Initial Table Version Number

Conditions for a Change in BGP Table Version

Usage of Table Version


Usage for Troubleshooting

Introduction

This document describes Table Version, which is a number used by Border Gateway Protocol (BGP) in order to track which best path changes of BGP prefixes are propagated to which BGP peers. It is a number used by the BGP software. You can view the Table Version number if you enter show commands, which helps the network administrator troubleshoot problems.

Network Diagram

This is the network diagram that is used for this article:



Best Path

A BGP prefix has one or more paths, because the BGP prefix is learned from different BGP peers and sources.

Here is an example of a BGP prefix with multiple paths. There are two paths, and the best path is the second one.

```
R1#show bgp ipv4 unicast 10.100.1.1
```

```

BGP routing table entry for 10.100.1.1/32, version 2
Paths: (2 available, best #2, table default)
  Advertised to update-groups:
    1
  Refresh Epoch 1
  5 4
    10.1.5.5 from 10.1.5.5 (10.1.5.5)
      Origin IGP, localpref 100, valid, external
      rx pathid: 0, tx pathid: 0
  Refresh Epoch 1
  4
    10.1.3.4 from 10.1.3.4 (10.100.1.1)
      Origin IGP, metric 0, localpref 100, valid, external, best
      rx pathid: 0, tx pathid: 0x0

```

Only one path is picked as the BGP best path based on the BGP Best Path algorithm. This is always the case. Refer to the BGP Best Path Selection Algorithm article for more information.

The path is either learned from a BGP peer or from a source, such as from redistribution from a routing protocol into BGP. When there is a change in the best path, the BGP must inform its peer by sending an update or a withdraw. The withdraw is sent when the last path of the BGP prefix is removed.

Here is an example where the prefix is sourced locally by the network command:

```

R4#show bgp ipv4 unicast 10.100.1.1
BGP routing table entry for 10.100.1.1/32, version 4
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 (10.1.3.4)
      Origin IGP, metric 0, localpref 100, weight 32768, valid, sourced, local, best
      rx pathid: 0, tx pathid: 0x0

```

The output shows *Origin IGP*.

Here is an example where the prefix is sourced locally by the redistribution connected command:

```

R4#show bgp ipv4 unicast 10.100.1.1
BGP routing table entry for 10.100.1.1/32, version 7
Paths: (1 available, best #1, table default)
Flag: 0x820
  Not advertised to any peer
  Refresh Epoch 1
  Local
    0.0.0.0 from 0.0.0.0 (10.1.3.4)
      Origin incomplete, metric 0, localpref 100, weight 32768, valid, sourced, best
      rx pathid: 0, tx pathid: 0x0

```

The output shows *Origin Incomplete*.

Types of Table Versions

The Table Version number is a 32-bit value, and there are four types of Table Versions:

- BGP Table Version
- Routing Information Base (RIB) Table Version
- Peer Table Version
- Prefix Table Version

These are further explained in the *Usage of Table Version* section.

Initial Table Version Number

When BGP has not learned about any prefixes yet, the global Table Version, the RIB Table Version, and the Peer Table Version are *1*, which is the starting point for the Table Version number.

The BGP command with the *summary* keyword gives you three Table Version numbers. The summary keyword can be provided for all the address families in BGP.

```
R1#show bgp ipv4 unicast summary
```

```
BGP router identifier 10.1.3.1, local AS number 1
```

```
BGP table version is 1, main routing table version 1
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
10.1.1.2	4	2	4	4	1	0	0	00:01:15	0
10.1.2.3	4	3	4	4	1	0	0	00:01:06	0
10.1.3.4	4	4	4	4	1	0	0	00:01:33	0

You can view the Prefix Table Version if you look at a prefix in the BGP table.

```
R1#show bgp ipv4 unicast 10.100.1.1/32
```

```
BGP routing table entry for 10.100.1.1/32, version 2
```

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

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

```
    1
```

```
  Refresh Epoch 1
```

```
  4
```

```
    10.1.3.4 from 10.1.3.4 (10.1.3.4)
```

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

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

You can view the Table Version if you enter the *show ip bgp internal* command:

```
R1#show ip bgp internal
```

```
Time left for bestpath timer: 964 secs
```

```
Consistency-checker not enabled
```

```
Update generation pool version 8, messages 0, in pool 0, below 00:00:24.432.
```

```
Enhanced Refresh EOR Stalepath-time disabled
```

```
Enhanced Refresh max-eor-time disabled
```

```
Total number of BGP Acceptor process: 50, Spawned count: 0
```

```
Total number of neighbors: 4
```

```
Total number of sessions : 4
```

```
  Established           : 4
```

```
  OpenConfirm          : 0
```

```
  OpenSent             : 0
```

```
  Active               : 0
```

```
  Connect              : 0
```

```
  Idle                 : 0
```

```
  Closing              : 0
```

```
  Uninitialized        : 0
```

```
Address-family IPv4 Unicast, Mode : RW
```

```
Table Versions : Current 39 Init 2 RIB 39
```

```
Start time : 00:00:18.919    Time elapsed 22:15:38.198
```

```
First Peer up in : 00:00:06.830    Exited Read-Only in : 00:01:07.966
```

```
Done with Install in : 00:01:07.967    Last Update-done in : 00:01:07.969
```

```
0 updates expanded
```

```
L3VPN Tunnel Encapsulated Paths : 0
```

```
Slow-peer detection is disabled    BGP Nexthop scan:-
```

```
  penalty: 0, Time since last run: 21:19:42.174, Next due in: none
```

```

    Max runtime : 0 ms Latest runtime : 0 ms Scan count: 2
BGP General Scan:-
    Max runtime : 1 ms Latest runtime : 0 ms Scan count: 0

BGP future scanner version: 1333
BGP scanner version: 0
Address-family IPv4 Multicast, Mode : RW
Table Versions : Current 1 Init 1 RIB 1

Start time : 00:00:18.919    Time elapsed 22:15:38.199
First Peer up in : never    Exited Read-Only in : 00:00:10.286
Done with Install in : 00:00:10.286    Last Update-done in : never
0 updates expanded
L3VPN Tunnel Encapsulated Paths : 0
Slow-peer detection is disabled    BGP Nexthop scan:-
    penalty: 0, Time since last run: never, Next due in: none
    Max runtime : 0 ms Latest runtime : 0 ms Scan count: 0
BGP General Scan:-
    Max runtime : 1 ms Latest runtime : 0 ms Scan count: 0

BGP future scanner version: 1334
BGP scanner version: 0
Address-family MVPNV4 Unicast, Mode : RW
Table Versions : Current 1 Init 1 RIB 1

Start time : 00:00:18.919    Time elapsed 22:15:38.200
First Peer up in : never    Exited Read-Only in : 00:00:10.286
Done with Install in : 00:00:10.286    Last Update-done in : never
0 updates expanded
L3VPN Tunnel Encapsulated Paths : 0
Slow-peer detection is disabled    BGP Nexthop scan:-
    penalty: 0, Time since last run: never, Next due in: none
    Max runtime : 0 ms Latest runtime : 0 ms Scan count: 0
BGP General Scan:-
    Max runtime : 1 ms Latest runtime : 0 ms Scan count: 0

BGP future scanner version: 1334
TX VPN optimization enabled.

```

Conditions for a Change in BGP Table Version

For the BGP Table Version number to change, there must be a change in the best path and a change propagated to the RIB. A change to the RIB for a BGP prefix only occurs if the prefix is in the RIB as a BGP prefix. If any other routing protocol places the prefix in the routing, then the BGP prefix is marked as a RIB-failure. In that case, even if the best path changes, the Table Version does not change.

Here is an example where the BGP Table Version does not change. The BGP prefix **10.100.1.1/32** learned from **R4** is also learned by a static route configured on **R1**. So, **R1** installs the static route in the RIB, and BGP on **R1** marks the prefix as a RIB-failure, because it is not the BGP that installs the prefix in the RIB. Any change to the BGP paths for this prefix is not propagated to the RIB. So even though there is a best path change, the BGP Table Version is not bumped, because there is no update to the RIB.

```

R1#show bgp ipv4 unicast 10.100.1.1/32
BGP routing table entry for 10.100.1.1/32, version 8
Paths: (2 available, best #1, table default, RIB-failure(17))
  Advertised to update-groups:
    2
  Refresh Epoch 2
  4
    10.1.3.4 from 10.1.3.4 (10.100.1.1)
      Origin IGP, metric 0, localpref 100, valid, external, best

```

```

    rx pathid: 0, tx pathid: 0x0
Refresh Epoch 2
5 4
  10.1.5.5 from 10.1.5.5 (10.1.5.5)
    Origin IGP, localpref 100, valid, external
    rx pathid: 0, tx pathid: 0

R1#show ip route 10.100.1.1
Routing entry for 10.100.1.1/32
  Known via "static", distance 1, metric 0 (connected)
  Routing Descriptor Blocks:
  * directly connected, via Loopback0
    Route metric is 0, traffic share count is 1

```

Usage of Table Version

When the best path changes the BGP prefixes, a few things must happen:

- The RIB must be notified.
- The BGP peers must be informed.
- The router must track which BGP peer is informed of which best path changes.

The BGP Table Version is the main number used. This number is the same as the highest Prefix Table Version of any BGP prefix for a specific address family. Assume there are five prefixes in the BGP table, with Prefix Table Versions 3, 6, 8, 10, and 18. The BGP Table Version will then be 18.

The Peer Table Version is used in order to track which peers must be informed of which prefixes for which there were changes in the best path. The Peer Table Version of each peer is checked against the Prefix Table Version of the prefixes. If the Prefix Table Version of a prefix is lower than the Peer Table Version, then BGP must send an update for that prefix to that BGP peer. For example, if the peer **10.1.1.2** has a Peer Table Version of **60**, then that peer is up to date for all prefixes with Prefix Table Version of **60** and lower. The router must send a BGP update for all prefixes with a Prefix Table Version that is higher than **60**.

Once the router updates the BGP peer for the best path changed prefixes, the router updates the Peer Table Version for this peer. This Peer Table Version value is adjusted to match the value of the highest Prefix Table Version of all the prefixes for which this BGP peer was updated. Assume the Peer Table Version was **60**, and there are two prefixes with Prefix Table Versions **61** and **62**. Once the router sends the new best paths for these two prefixes to that BGP peer, then the Peer Table Version is updated to **62**.

The Prefix Table Version is the Table Version number attached to the BGP prefix. It is changed when the best path changes for that prefix. Each time the best path changes for one BGP prefix, its Prefix Table Version is bumped, which means it is updated to be equal to the next available version number. Assume prefix **10.0.0.0/8** has Prefix Table Version **27**, and the BGP Table Version is **30**. In this case, when the best path changes for prefix **10.0.0.0/8**, its Prefix Table Version is bumped to **31**.

The RIB Table version is used in order to track if the RIB needs to be updated after BGP best path changes occur. The RIB must be informed of the BGP prefixes which have a Prefix Table Version higher than the RIB Table Version. For these prefixes, there is a RIB ADD, DELETE, or MODIFY event.

Usage for Troubleshooting

In order to know when BGP has converged, enter the **show bgp summary** command. If the Peer BGP Table Version equals the BGP Table Version, that peer has converged. If the main routing table version equals the BGP Table Version, the RIB has converged.

```
R1#show bgp ipv4 unicast summary
```

```

BGP router identifier 10.1.3.1, local AS number 1
BGP table version is 2, main routing table version 2
1 network entries using 144 bytes of memory
1 path entries using 80 bytes of memory
1/1 BGP path/bestpath attribute entries using 144 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 392 total bytes of memory
BGP activity 1/0 prefixes, 1/0 paths, scan interval 60 secs

```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
10.1.1.2	4	2	69	69	2	0	0	01:00:54	0
10.1.2.3	4	3	69	70	2	0	0	01:00:45	0
10.1.3.4	4	4	72	70	2	0	0	01:01:12	1

There can be many changes to the BGP Table Versions, and that does not always mean that something is wrong.

Assume that the router is connected to the Internet, and has the full Internet routing table. Typically, there are a few changes almost every second on the Internet BGP table. Then, the router must recalculate the best path for some prefixes, and update its RIB and its BGP peers. This is expected behavior.

Assume you clear a BGP peer (the session is reset), then the router must advertise its full BGP table to that peer. It is expected for that peer to have an increasing Table Version. The Table Version increases as the peer receives the BGP prefixes again. The sending BGP peer does not increase the Table Version for the BGP prefixes.

Here is an example. The Table Version starts with **28**.

```

R1#show bgp ipv4 unicast summary
BGP router identifier 10.1.3.1, local AS number 1
BGP table version is 28, main routing table version 281
network entries using 144 bytes of memory2 path entries using 160 bytes of memory
2/1 BGP path/bestpath attribute entries using 288 bytes of memory
2 BGP AS-PATH entries using 48 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 640 total bytes of memory
BGP activity 1/0 prefixes, 16/14 paths, scan interval 60 secs

```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
10.1.1.2	4	2	117	125	28	0	0	01:43:50	0
10.1.2.3	4	3	117	125	28	0	0	01:43:53	0
10.1.3.4	4	4	10	12	28	0	0	00:04:22	1
10.1.5.5	4	5	55	63	28	0	0	00:45:45	1

```

R1#show bgp ipv4 unicast 10.100.1.1/32
BGP routing table entry for 10.100.1.1/32, version 28

```

```

Paths: (2 available, best #1, table default)
  Advertised to update-groups:
    1
  Refresh Epoch 2
    4
    10.1.3.4 from 10.1.3.4 (10.100.1.1) <<< path from R4
      Origin IGP, metric 0, localpref 100, valid, external, best
      rx pathid: 0, tx pathid: 0x0
  Refresh Epoch 2
    5 4
    10.1.5.5 from 10.1.5.5 (10.1.5.5) <<< path from R5
      Origin IGP, localpref 100, valid, external

```

rx pathid: 0, tx pathid: 0

Conduct a hard clear for the BGP session towards R1 on the peer **10.1.3.4 (R4)**. The peer advertises only one prefix **10.100.1.1/32** towards **R1**. **10.100.1.1/32** is learned from **R4** and **R5**. The best path is the path from **R4**.

Ensure that you have **debug ip bgp internal** enabled in order to see what happens to the BGP Table Versions. You should have **debug ip bgp updates** enabled in order to see what happens when the update arrives.

R1#**debug ip bgp updates**

BGP updates debugging is on for address family: IPv4 Unicast

R1#**debug ip bgp internal**

BGP internal debugging is on

R1#**show debugging**

IP routing:

BGP internal debugging is on

BGP updates debugging is on for address family: IPv4 Unicast

R1#

%BGP-5-NBR_RESET: Neighbor 10.1.3.4 reset (Peer closed the session) <<< **BGP**

session to R4 goes down

BGP: TX IPv4 Unicast Net global 10.100.1.1/32 Changed.

BGP: TX IPv4 Unicast Net global 10.100.1.1/32 RIB done.

BGP: TX IPv4 Unicast Net global 10.100.1.1/32 Changed.

BGP: TX IPv4 Unicast Mem global 1 1 10.1.3.4 Resetting counters.

BGP: TX IPv4 Unicast Mem global 1 1 10.1.3.4 Ignoring dummy policy change.

BGP: TX IPv4 Unicast Mem global 1 1 10.1.3.4 Resetting counters.

BGP: TX IPv4 Unicast Mem global 1 1 10.1.3.4 Ignoring dummy policy change.

BGP: TX IPv4 Unicast Mem global 1 1 10.1.3.4 Changing state from ACTIVE to DOWN (session not established).

BGP: TX IPv4 Unicast Mem global 1 1 10.1.3.4 Removing from group (3 members left).

%BGP-5-ADJCHANGE: neighbor 10.1.3.4 Down Peer closed the session

%BGP_SESSION-5-ADJCHANGE: neighbor 10.1.3.4 IPv4 Unicast topology base removed from session Peer

closed the session

BGP: TX IPv4 Unicast Mem global 10.1.3.4 State is DOWN (session not established).

BGP: net global:IPv4 Unicast:base 10.100.1.1/32 RIB-INSTALL Attempting to

install. <<< **RIB gets informed**

BGP: net global:IPv4 Unicast:base 10.100.1.1/32 RIB-INSTALL Built route type:

1024, flags: 200000, tag: 5,

metric: 0 path: 1.

BGP: net global:IPv4 Unicast:base 10.100.1.1/32 RIB-INSTALL Path 1, type: DEF,

gw: 10.1.5.5, idb: N/A,

topo_id: 0, src: 1.1.5.5, lbl: 1048577, flags: 0.

BGP: net global:IPv4 Unicast:base 10.100.1.1/32 RIB-INSTALL Installing 1 paths, multipath limit 1 (from 1).

BGP(0): Revise route installing 1 of 1 routes for 10.100.1.1/32 -> 10.1.5.5

(global) to main IP table <<< **The remaining path through R5 gets installed**

in the RIB

BGP: net global:IPv4 Unicast:base 10.100.1.1/32 RIB-INSTALL Install successful.

BGP: TX IPv4 Unicast Net global 10.100.1.1/32 RIB done.

BGP: TX IPv4 Unicast Net global 10.100.1.1/32 RIB done.

BGP: TX IPv4 Unicast Tab RIB walk done version 29, added 1 topologies.

BGP: TX IPv4 Unicast Tab Executing.

BGP: TX IPv4 Unicast Wkr global 1 Cur Processing.

BGP: TX IPv4 Unicast Top global Appending nets from attr 0x9362CB4.

BGP: TX IPv4 Unicast Wkr global 1 Cur Attr change from 0x0 to 0x9362CB4.

BGP(0): (base) 10.1.1.2 send UPDATE (format) 10.100.1.1/32, next 10.1.1.1,

metric 0, path 5 4 <<< R1 sends update for 10.100.1.1/32 for Table Version 29.

(bestpath is still the one from R5, i.e. the only one R1 has at this moment)

BGP: TX IPv4 Unicast Wkr global 1 Cur Net 10.100.1.1/32 (Pxt 0x9F58FA0:0x0)

Formatted.

BGP: TX IPv4 Unicast Top global No attributes with modified nets.

BGP: TX IPv4 Unicast Top global Added tail marker with version 29.

BGP: TX IPv4 Unicast Wkr global 1 Cur Reached marker with version 29.
BGP: TX IPv4 Unicast Top global No attributes with modified nets.
BGP: TX IPv4 Unicast Wkr global 1 Cur Replicating.
BGP: TX IPv4 Unicast Wkr global 1 Cur Done (end of list), processed 1 attr(s),
1/1 net(s), 0 pos.
BGP: TX IPv4 Unicast Grp global 1 Checking EORs again (3/3).
BGP: TX IPv4 Unicast Grp global 1 Start minimum advertisement timer (30 secs).
BGP: TX IPv4 Unicast Wkr global 1 Cur Blocked (minimum advertisement interval).
BGP: TX IPv4 Unicast Wkr global 1 Cur Reached end of list.
BGP: TX IPv4 Unicast Grp global 1 Converged.
BGP: TX IPv4 Unicast Tab Processed 1 walker(s).
BGP: TX IPv4 Unicast Tab Generation completed.
BGP: TX IPv4 Unicast Top global Deleting first marker with version 28.
BGP: TX IPv4 Unicast Top global Collection reached marker 28 after 0 path
extension(s).
BGP: TX IPv4 Unicast Top global Collection done on marker 29 after 1 path
extension(s).
BGP: TX IPv4 Unicast Top global Collection done on marker 29 after 0 path
extension(s).
BGP: TX IPv4 Unicast Mem global 10.1.3.4 Policy change while no group and
member is DOWN.
BGP: TX IPv4 Unicast Mem global 10.1.3.4 Changing state from DOWN to WAIT
(pending advertised bit allocation).
BGP: TX IPv4 Unicast Mem global 1 1 10.1.3.4 Added to group (now has
4 members).
BGP: TX IPv4 Unicast Mem global 1 1 10.1.3.4 Continuing into ACTIVE state.
BGP: TX IPv4 Unicast Mem global 1 1 10.1.3.4 Refresh Start-of-rib for afi 1,
safi 1.
BGP: TX IPv4 Unicast Mem global 1 1 10.1.3.4 Full refresh requested.
BGP: TX IPv4 Unicast Mem global 1 1 10.1.3.4 Refresh has to wait for pathext
prepend.
**%BGP-5-ADJCHANGE: neighbor 10.1.3.4 Up <<< BGP session to R4 is up again.
But, R1 did not learn the prefix 10.100.1.1/32 yet from R4.**
BGP: nbr_topo global 10.1.3.4 IPv4 Unicast:base (0x63D50D0:1) rcvd Refresh
Start-of-RIB
BGP: nbr_topo global 10.1.3.4 IPv4 Unicast:base (0x63D50D0:1) refresh_epoch
is 2
BGP: TX IPv4 Unicast Top global Start pathext prepend.
BGP: TX IPv4 Unicast Tab Pathext prepend full table refresh.
BGP: TX IPv4 Unicast Tab Pathext prepend full table refresh.
BGP: TX IPv4 Unicast Top global Inserting initial marker.
BGP: TX IPv4 Unicast Top global Done pathext prepend (1 attrs).
BGP: TX IPv4 Unicast Grp global 1 Starting refresh after prepend completion.
BGP: TX IPv4 Unicast Mem global 1 1 10.1.3.4 Starting refresh (first member,
1, 0, marker).
BGP: TX IPv4 Unicast Wkr global 1 Ref Start at marker 1.
BGP: TX IPv4 Unicast Wkr global 1 Ref Unblocked
BGP: TX IPv4 Unicast Top global Collection done on marker 1 after 0 path
extension(s).
BGP: TX IPv4 Unicast Tab Executing.
BGP: TX IPv4 Unicast Wkr global 1 Ref Processing.
BGP: TX IPv4 Unicast Wkr global 1 Ref Attr change from 0x0 to 0x9362CB4.
BGP(0): (base) 10.1.1.2 send UPDATE (format) 10.100.1.1/32, next 10.1.1.1,
metric 0, path 5 4
BGP: TX IPv4 Unicast Wkr global 1 Ref Net 10.100.1.1/32 (Pxt 0x9F58FA0:0x0)
Formatted.
BGP: TX IPv4 Unicast Wkr global 1 Ref Reached marker with version 29.
BGP: TX IPv4 Unicast Wkr global 1 Ref Replicating (pending member_pos
processing).
BGP: TX IPv4 Unicast Mem global 1 1 10.1.3.4 Completed refresh.
BGP: TX IPv4 Unicast Mem global 1 1 10.1.3.4 Refresh stop.
BGP: TX IPv4 Unicast Grp global 1 Refresh complete.
BGP: TX IPv4 Unicast Wkr global 1 Ref Stop.
BGP: TX IPv4 Unicast Wkr global 1 Ref Blocked (not in list).
BGP: TX IPv4 Unicast Grp global 1 Converged.
BGP: TX IPv4 Unicast Mem global 1 1 10.1.3.4 Send EOR.

BGP: TX IPv4 Unicast Wkr global 1 Ref Suspending / blocked (member marker), processed 1 attr(s), 1/1 net(s), 1 pos.

BGP: TX IPv4 Unicast Tab Processed 1 walker(s).

BGP: TX IPv4 Unicast Tab Generation completed.

BGP: TX IPv4 Unicast Top global Deleting first marker with version 1.

BGP: TX IPv4 Unicast Top global Collection reached marker 1 after 0 path extension(s).

BGP: TX IPv4 Unicast Top global Collection done on marker 29 after 1 path extension(s).

BGP: TX IPv4 Unicast Top global Collection done on marker 29 after 0 path extension(s).

BGP(0): 10.1.3.4 rcvd UPDATE w/ attr: nexthop 10.1.3.4, origin i, metric 0, merged path4, AS_PATH

BGP(0): 10.1.3.4 rcvd 10.100.1.1/32 <<< R1 received 10.100.1.1/32 from R4 again

BGP: TX IPv4 Unicast Net global 10.100.1.1/32 Changed.

BGP: nbr_topo global 10.1.3.4 IPv4 Unicast:base (0x63D50D0:1) rcvd Refresh End-of-RIB

BGP: net global:IPv4 Unicast:base 10.100.1.1/32 RIB-INSTALL Attempting to install.

BGP: net global:IPv4 Unicast:base 10.100.1.1/32 RIB-INSTALL Built route type: 1024, flags: 200000, tag: 4, metric: 0 path: 1.

BGP: net global:IPv4 Unicast:base 10.100.1.1/32 RIB-INSTALL Path 1, type: DEF, gw: 10.1.3.4, idb: N/A, topo_id: 0, src: 1.1.3.4, lbl: 1048577, flags: 0.

BGP: net global:IPv4 Unicast:base 10.100.1.1/32 RIB-INSTALL Installing 1 paths, multipath limit 1 (from 1).

BGP(0): Revise route installing 1 of 1 routes for 10.100.1.1/32 -> 10.1.3.4 (global) to main IP table

BGP: net global:IPv4 Unicast:base 10.100.1.1/32 RIB-INSTALL Install successful.

BGP: TX IPv4 Unicast Net global 10.100.1.1/32 RIB done.

BGP: TX IPv4 Unicast Net global 10.100.1.1/32 RIB done.

BGP: TX IPv4 Unicast Tab RIB walk done version 30, added 1 topologies.

BGP: TX IPv4 Unicast Tab Executing.

BGP: TX IPv4 Unicast Tab Generation completed.

BGP: TX Member message pool under period (60 < 600).

BGP: TX IPv4 Unicast Mem global 1 1 10.1.2.3 State is ACTIVE (ready).

BGP: TX IPv4 Unicast Grp global 1 Minimum advertisement timer expired.

BGP: TX IPv4 Unicast Wkr global 1 Cur Unblocked

BGP: TX IPv4 Unicast Tab Executing.

BGP: TX IPv4 Unicast Wkr global 1 Cur Processing.

BGP: TX IPv4 Unicast Top global Appending nets from attr 0x9362D54.

BGP: TX IPv4 Unicast Wkr global 1 Cur Attr change from 0x0 to 0x9362D54.

BGP(0): (base) 10.1.1.2 send UPDATE (format) 10.100.1.1/32, next 10.1.1.1, metric 0, path 4 <<< R1 sends an update for 10.100.1.1/32 for Table Version 30 (bestpath is again the one from R4)

BGP: TX IPv4 Unicast Wkr global 1 Cur Net 10.100.1.1/32 (Pxt 0x9F58FA0:0x0) Formatted.

BGP: TX IPv4 Unicast Top global No attributes with modified nets.

BGP: TX IPv4 Unicast Top global Added tail marker with version 30.

BGP: TX IPv4 Unicast Wkr global 1 Cur Reached marker with version 30.

BGP: TX IPv4 Unicast Top global No attributes with modified nets.

BGP: TX IPv4 Unicast Wkr global 1 Cur Replicating.

BGP: TX IPv4 Unicast Wkr global 1 Cur Done (end of list), processed 1 attr(s), 1/1 net(s), 0 pos.

BGP: TX IPv4 Unicast Grp global 1 Checking EORs again (4/4).

BGP: TX IPv4 Unicast Grp global 1 Start minimum advertisement timer (30 secs).

BGP: TX IPv4 Unicast Wkr global 1 Cur Blocked (minimum advertisement interval).

BGP: TX IPv4 Unicast Wkr global 1 Cur Reached end of list.

BGP: TX IPv4 Unicast Grp global 1 Converged.

BGP: TX IPv4 Unicast Tab Processed 1 walker(s).

BGP: TX IPv4 Unicast Tab Generation completed.

BGP: TX IPv4 Unicast Top global Deleting first marker with version 29.

BGP: TX IPv4 Unicast Top global Collection reached marker 29 after 0 path extension(s).

BGP: TX IPv4 Unicast Top global Collection done on marker 30 after 1 path extension(s).

BGP: TX IPv4 Unicast Top global Collection done on marker 30 after 0 path extension(s).

BGP: TX IPv4 Unicast Tab RIB walk done version 30, added 0 topologies.

All Table Versions are at **30** now:

```
R1#show bgp ipv4 unicast summary
```

```
BGP router identifier 10.1.3.1, local AS number 1
```

```
BGP table version is 30, main routing table version 30
```

```
1 network entries using 144 bytes of memory
```

```
2 path entries using 160 bytes of memory
```

```
2/1 BGP path/bestpath attribute entries using 288 bytes of memory
```

```
2 BGP AS-PATH entries using 48 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 640 total bytes of memory
```

```
BGP activity 1/0 prefixes, 17/15 paths, scan interval 60 secs
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
10.1.1.2	4	2	127	135	30	0	0	01:52:42	0
10.1.2.3	4	3	126	136	30	0	0	01:52:45	0
10.1.3.4	4	4	12	14	30	0	0	00:06:25	1
10.1.5.5	4	5	64	73	30	0	0	00:54:37	1

```
R1#show bgp ipv4 unicast 10.100.1.1/32
```

```
BGP routing table entry for 10.100.1.1/32, version 30
```

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

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

```
    1
```

```
  Refresh Epoch 2
```

```
  4
```

```
    10.1.3.4 from 10.1.3.4 (10.100.1.1)
```

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

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

```
  Refresh Epoch 2
```

```
  5 4
```

```
    10.1.5.5 from 10.1.5.5 (10.1.5.5)
```

```
      Origin IGP, localpref 100, valid, external
```

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

In the end, on **RI**, there were two best path changes. So, the Table Version got bumped by **2**.

First, the peer **10.1.3.4** went down on **RI**. The best path changed to the path received from **R5**. The Table Version increased to the next available number, which was **29**. The Prefix Table Version was bumped to **29** as well. The RIB was updated with this new best path. The Table Version of the RIB was increased to **29**. Then, **RI** sent an update to the BGP peer **10.1.1.2** for the new best path and updated the peer Table Version to **29**. Every other peer was updated as well.

Second, once the peer **10.1.3.4** was up again, **RI** received the update for **10.100.1.1/32** from **R4** and recalculated the best path. The path from **R4** was the new best path, which caused the Table Version and the Prefix Table Version to be bumped to the next available number of **30**. Again, the RIB and all other BGP peers were updated, and the RIB and peer Table Versions were updated to **30**. The Table Version was bumped only by one each time here. However, if the other BGP prefixes underwent other changes, this Table Version would be bumped by more than one, because it jumps each time to the next available number.

If you enter the **clear ip bgp out** command for a BGP peer, that router resends its BGP prefixes to that peer. This does not cause a change in the best path on the receiving BGP peer. Hence, there is no change in the Table Version on that peer.

When you run the **debug ip bgp updates** on the receiving router, you see:

```
BGP(0): 10.1.3.4 rcvd UPDATE w/ attr: nexthop 10.1.3.4, origin i,  
metric 0, merged path 4, AS_PATH  
BGP(0): 10.1.3.4 rcvd 10.100.1.1/32...duplicate ignored
```

The received update is recognized as a duplicate, so it is ignored and no best path change occurs.

Assume you have a router with 100.000 prefixes in the BGP table, and the BGP Table Version increases by 100.000 every minute. This is not expected, and the behavior must be examined. One reason for the behavior could be that the next-hop for the BGP prefixes is flapping for all prefixes every minute.

One of the results when the BGP Table Version increases rapidly is that the process BGP Router and BGP IO are busy, which might cause a constant high router CPU.

Updated: Sep 30, 2013

Document ID: 116511
