

## Data Communications IV

### Lab #3 – BGP Communities




A Table 3 production, in a peering relationship with Goodn3t.com


Lab 3 by Table 3's production team:




Matt Capranos



Tom Woo



Roland Szczesny



Stuart MacLean

**Table of Contents**

Executive Summary.....page 3

Overview.....page 4

Initial Router Configuration.....page 5

Communities Attribute.....page 6

Dissecting the Captured Data.....page 9

Appendix.....page 11

### **Executive Summary**

In the third lab for BGP we will be implementing communities for BGP, more specifically we will be implementing and testing the following:

- All Routers R1 prefer VLAN12 for incoming traffic directed to their loopbacks, advertise community value as AS#:12
- All Routers R2 prefer VLAN22 for incoming traffic directed to their loopbacks, advertise community value as AS#:22
- All Routers R3 prefer VLAN32 for incoming traffic directed to their loopbacks, advertise community value as AS#:32
- AS 1,2,3,and 4 prefer that VLAN12 for incoming traffic directed to their default VLAN (192.168.T#.0 / 24), advertise community value as AS#:12
- AS 5,6, and 7 prefer VLAN22 for incoming traffic directed to their default VLAN (192.168.T#.0 / 24), advertise community value as AS#:22
- AS 8,9, and 10 prefer VLAN32 for incoming traffic directed to their default VLAN (192.168.T#.0 / 24) , advertise community value as AS#:32

In order to accomplish the stated goals, configurations had to be made to the Autonomous System Border Routers on the edge of our Autonomous System (AS) 65003, this is required in order to manipulate the way Table 3's AS border routers' inbound preferred routes. An example of which is ASBR 172.16.13.3 would prefer to use VLAN 32 for all incoming traffic that is directed towards it's loopback address (3.3.3.3), and to ensure that VLAN 12 is the preferred route for any traffic directed to network 192.168.3.3

Within the report we have asked several guest speakers to assist us in discussing the configurations made on each BGP router within AS 65003 in order to accomplish advertisements using the BGP communities' attribute, additionally, the discussion of the configurations necessary on external AS border routers in order to accomplish BGP communities.

## Overview




Hello, My name is Joshua and I'm here to talk about the network configuration for the third BGP Lab, before I continue, I would like to thank Table 3 for the invite.

For Lab 3, implementing BGP Communities, we are going to simulate an Internet of 10 Autonomous systems, these AS numbered AS 65001 – 65010, will advertise routes tagged with the community attributes.

Below is a diagram showing the BGP peering relations between the various AS border routers in

the different autonomous systems.



From the above diagram we see that all AS Border Routers who's IP address is 172.16.Table#.1 will be in VLAN 12, all AS Border Routers who's IP address is 172.16.Table#.2 will be in VLAN 22, and finally all AS Border Routers whose IP address is 172.16.Table#.3 will be in VLAN 32. Within each AS the Interior Gateway Protocol that will be implemented is OSPF, and the internal network address that will be advertised will be based on 192.168.Table#.Router#, an example of which would be 192.168.3.3 which is the network being advertised by Table 3, Router 3.

Sorry, to cut your description short Josh but we must move on, thank you again.

### Initial Router Configuration


Hello, my name is JC and I'm here to show you Table 3's initial router configuration. As you can see I'm hard at work ensuring that the AS border routers at AS 65003 have the proper peering relations with their neighbors.



To help illustrate the router Configuration we are going to use AS border router Table 3 Router 3 (BGP router-id 3.3.3.3). The configurations for all the AS 65003 routers can be found in the appendix.

### Router 3.3.3.3 initial configuration

```
router bgp 65003
  bgp router-id 3.3.3.3
  network 3.3.3.3/32
  network 192.168.3.0/24
  neighbor 172.16.13.1 remote-as 65001
  neighbor 172.16.13.2 remote-as 65002
  neighbor 172.16.13.3 remote-as 65003
  neighbor 172.16.13.4 remote-as 65004
  neighbor 172.16.13.5 remote-as 65005
  neighbor 172.16.13.6 remote-as 65006
  neighbor 172.16.13.7 remote-as 65007
  neighbor 172.16.13.8 remote-as 65008
  neighbor 172.16.13.9 remote-as 65009
  neighbor 172.16.13.10 remote-as 65010
```



For the initial configuration the AS border routers 172.16.11.3, 172.16.12.3 and 172.16.13.3 within AS 65003 will be peering with all the AS border routers within their respected VLANs, for example Router 172.16.11.3 will be peering with all BGP routers within VLAN 11, router 172.16.12.3 will be peering with all BGP routers within VLAN 12 and finally router 172.16.13.3 will be peering with all BGP routers within VLAN 13.

Well JC is extremely busy and needs to get back to work, so we'll leave him be, thank you again JC.....JC, wake up JC.....

## Communities Attributes



To discuss the communities attributes configuration we have invited Steve to say a few words. Well that isn't a very nice introduction Steve.

Whatever, Capranos, I'm just doing this so you'll come and play World of Warcraft. Today, I'm here to discuss the BGP communities configuration for the AS border routers within AS 65003.

By configuring the communities attribute allows us to manipulate the AS border routers' inbound preferred routes. An example of which is ASBR 172.16.13.3 would prefer that VLAN 13 is used for all incoming traffic that is directed towards the Routers loopback address (3.3.3.3), and to ensure that VLAN 12 is the preferred route for any traffic directed to network 192.168.3.3. Below is part of the configuration with the community attributes, the full configuration can be found in the appendix.

### Community configuration

```
router bgp 65003
  bgp router-id 3.3.3.3
  network 3.3.3.3/32
  network 192.168.3.0 /24
  neighbor 172.16.13.1 remote-as 65001
  neighbor 172.16.13.1 route-map LOCAL_PREF in
  neighbor 172.16.13.1 route-map SETCOMMUNITY out
  neighbor 172.16.13.2 remote-as 65002
  neighbor 172.16.13.2 route-map LOCAL_PREF in
  neighbor 172.16.13.2 route-map SETCOMMUNITY out
  !
  access-list 1 permit 3.3.3.3
  !
  route-map LOCAL_PREF permit 10
  match community LOCAL_PREF
  set local-preference 123
  !
  route-map LOCAL_PREF permit 20
  !
  route-map SETCOMMUNITY permit 10
  match ip address 1
  set community 65003:12
  !
  route-map SETCOMMUNITY permit 20
```

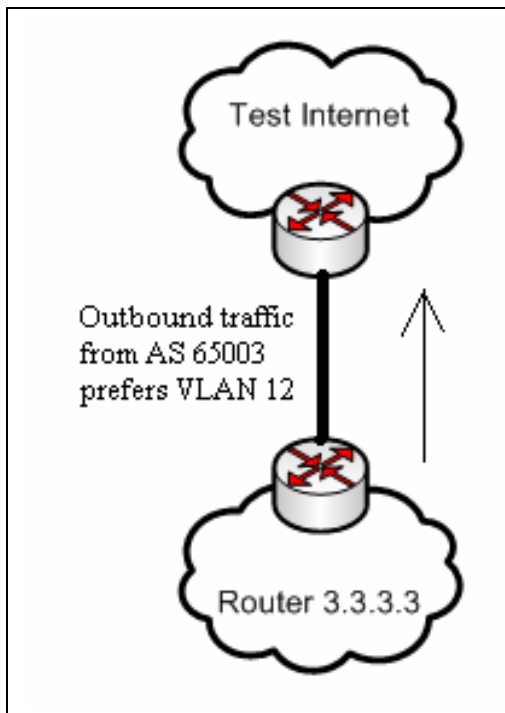
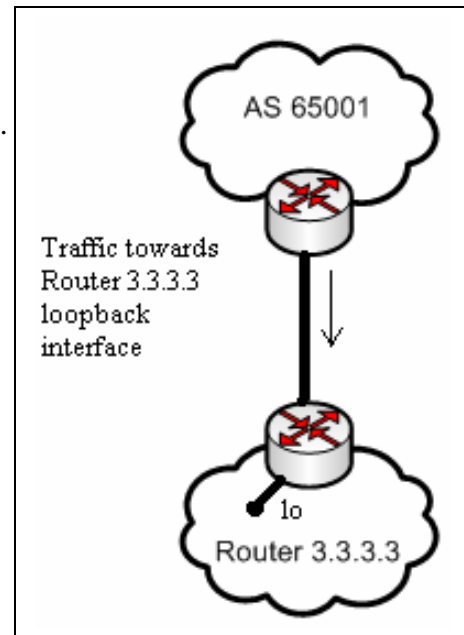
## BGP Lab 3 – BGP Communities Attributes

For traffic that is destined towards Router 3.3.3.3 loopback interface the *neighbor 172.16.13.Table3 route-map LOCAL\_PREF* in configuration was used. Route-maps which are used with BGP to control and modify routing information and to define the conditions in which the routes will be redistributed between the AS border routers in the different autonomous systems.

To give a higher local preference for any traffic that is inbound to Router 3.3.3.3 loopback interface the following route-map is used in conjunction with the neighbor command. (several lines were removed for brevity)

```
neighbor 172.16.13.table# remote-as 6500(table#)
neighbor 172.16.13.table# route-map LOCAL_PREF in
!
route-map LOCAL_PREF permit 10
match community LOCAL_PREF
set local-preference 123
```

The route-map part1 ensures that all traffic destined for Router 3.3.3.3 loopback address has a local preference of 123.



For the second phase of BGP communities must be configured to ensure that AS 1,2,3 and 4 prefer that VLAN 12 is used for incoming traffic directed to the default network located within their autonomous system. This is accomplished by using *neighbor 172.16.13.Table# route-map SETCOMMUNITY out* command, which is shown below.

```
neighbor 172.16.13.1 remote-as 65001
neighbor 172.16.13.1 route-map SETCOMMUNITY out
!
access-list 1 permit 3.3.3.3
!
route-map SETCOMMUNITY permit 10
match ip address 1
set community 65003:12
```

The above command specifies that for neighbor 172.16.13.1 use the route-map SETCOMMUNITY, that if the ip address 3.3.3.3 is match, then to advertise the community for AS65003 to use VLAN 12.

## BGP Lab 3 – BGP Communities Attributes

Thank you for that overview of the community configuration Steve, much appreciated, oh and I'm never playing World of Warcraft again..... Well maybe not until the of semester. Oh, I see Roni has something that he would like to add to the BGP community configuration overview.

Why didn't you use the "Soft-Reconfiguration" on the inbound route-maps?



Excellent questions Roni, after the initial configurations and some additional reading Table 3 modified our configurations to use the "Soft-Reconfiguration" on the inbound route-maps. The BGP Soft reconfiguration is used in order to let policies be configured and activated without restarting the BGP TCP session. Before implementing the soft-reconfiguration the BGP process on the Linux machines had to be terminated to allow new changes to be made, with soft-reconfiguration the new policies or changes to existing policies are automatically triggered and the appropriate updates are generated to enforce the new policy.

**Dissecting the captured data**



Our final speaker of the evening is Jeff Moon, taking time out of his practical exam to help dissect the captured data.

We start our analysis at the beginning, for this example we'll use router 3.3.3.2 and router 4.4.4.2. The BGP process on a Router 3.3.3.2 has not been started yet, after the process is started we watch for the three way handshake between Routers 3.3.3.2 and 4.4.4.2. The handshake is shown in the diagram below.

No. -	Time	Source	Destination	Protocol	Info
18	31.971431	172.16.12.3	172.16.12.4	TCP	36945 > bgp [SYN] Seq=0 Ack=0 win=5840 Len=0 MSS=1460 TSV=1743
19	31.971552	172.16.12.4	172.16.12.3	TCP	bgp > 36945 [SYN, ACK] Seq=0 Ack=1 win=5792 Len=0 MSS=1460 TSV
20	31.971569	172.16.12.3	172.16.12.4	TCP	36945 > bgp [ACK] Seq=1 Ack=1 win=5840 Len=0 TSV=17436384 TSER

After the three-way handshake is complete and the connection is established, an OPEN message is sent from Router 3.3.3.2,

Border Gateway Protocol
OPEN Message
Marker: 16 bytes
Length: 45 bytes
Type: OPEN Message (1)
Version: 4
My AS: 65003
Hold time: 180
BGP identifier: 3.3.3.2
Optional parameters length: 16 bytes
optional parameters
Capabilities Advertisement (8 bytes)
Capabilities Advertisement (4 bytes)
Capabilities Advertisement (4 bytes)

in this OPEN message (which is shown in the diagram to the left).

Within the BGP open message the version of the BGP router is advertised to its peer which is router 4.4.4.2, the hold down timer and the optional parameters.

After the exchange of the OPEN messages (from router 3.3.3.2 to 4.4.4.2 and 4.4.4.2 to 3.3.3.2) a series of keep alive messages are then sent, as shown in the diagram below.

24	31.972709	172.16.12.3	172.16.12.4	TCP	36945 > bgp [ACK] Seq=46 Ack=46 win=5840 Len=0 TSV=17436385 TS
25	31.972817	172.16.12.4	172.16.12.3	BGP	KEEPALIVE Message
26	31.972824	172.16.12.3	172.16.12.4	TCP	36945 > bgp [ACK] Seq=46 Ack=65 win=5840 Len=0 TSV=17436385 TS
27	31.972944	172.16.12.3	172.16.12.4	BGP	KEEPALIVE Message
28	31.973102	172.16.12.4	172.16.12.3	BGP	KEEPALIVE Message
29	31.973110	172.16.12.3	172.16.12.4	BGP	KEEPALIVE Message
30	31.973110	172.16.12.4	172.16.12.3	BGP	UPDATE Message

After approximately 13 frames (slightly less than 0.01 seconds) the first UPDATE message is sent from BGP router 172.16.12.4 to 172.16.13.4 advertising BGP community attributes. As shown in the screen capture on the next page.

## BGP Lab 3 – BGP Communities Attributes

```
Internet Protocol, Src: 172.16.12.4 (172.16.12.4), Dst
Transmission Control Protocol, Src Port: bgp (179), D
Border Gateway Protocol
  UPDATE Message
    Marker: 16 bytes
    Length: 60 bytes
    Type: UPDATE Message (2)
    Unfeasible routes length: 0 bytes
    Total path attribute length: 32 bytes
  Path attributes
    ORIGIN: IGP (4 bytes)
    AS_PATH: 65004 (7 bytes)
    NEXT_HOP: 172.16.12.4 (7 bytes)
    MULTI_EXIT_DISC: 0 (7 bytes)
    COMMUNITIES: 65004:12 (7 bytes)
      Flags: 0xc0 (Optional, Transitive, Complete)
        1... .... = Optional
        .1.. .... = Transitive
        ..0. .... = Complete
        ...0 .... = Regular length
      Type code: COMMUNITIES (8)
      Length: 4 bytes
    Communities: 65004:12
      Community: 65004:12
        Community AS: 65004
        Community value: 12
  Network layer reachability information: 5 bytes
```

Within the BGP UPDATE message under the Path Attributes tree, is the community attribute.

The advertised Community by router 172.16.13.4 to router 172.16.13.3, displays the *OPTIONAL*, *TRANSITIVE*, and *COMPLETE* flags, as shown in the capture the Community attribute being advertised is Optional and Transitive. The Community AS 65004 and its value is 12.

## BGP Lab 3 – BGP Communities Attributes

### Appendix

#### Router 172.16.11.3 configuration

```
router bgp 65003
  bgp router-id 3.3.3.2
  network 3.3.3.2/32
  network 192.168.3.0/24
  neighbor 172.16.11.1 remote-as 65001
  neighbor 172.16.11.1 route-map LOCAL_PREF in
  neighbor 172.16.11.1 soft-reconfiguration inbound
  neighbor 172.16.11.1 route-map SETCOMMUNITY out
  neighbor 172.16.11.2 remote-as 65002
  neighbor 172.16.11.2 route-map LOCAL_PREF in
  neighbor 172.16.11.2 soft-reconfiguration inbound
  neighbor 172.16.11.2 route-map SETCOMMUNITY out
  neighbor 172.16.11.4 remote-as 65004
  neighbor 172.16.11.4 route-map LOCAL_PREF in
  neighbor 172.16.11.4 soft-reconfiguration inbound
  neighbor 172.16.11.4 route-map SETCOMMUNITY out
  neighbor 172.16.11.5 remote-as 65005
  neighbor 172.16.11.5 route-map LOCAL_PREF in
  neighbor 172.16.11.5 soft-reconfiguration inbound
  neighbor 172.16.11.5 route-map SETCOMMUNITY out
  neighbor 172.16.11.6 remote-as 65006
  neighbor 172.16.11.6 route-map LOCAL_PREF in
  neighbor 172.16.11.6 soft-reconfiguration inbound
  neighbor 172.16.11.6 route-map SETCOMMUNITY out
  neighbor 172.16.11.7 remote-as 65007
  neighbor 172.16.11.7 route-map LOCAL_PREF in
  neighbor 172.16.11.7 soft-reconfiguration inbound
  neighbor 172.16.11.7 route-map SETCOMMUNITY out
  neighbor 172.16.11.8 remote-as 65008
  neighbor 172.16.11.8 route-map LOCAL_PREF in
  neighbor 172.16.11.8 soft-reconfiguration inbound
  neighbor 172.16.11.8 route-map SETCOMMUNITY out
  neighbor 172.16.11.9 remote-as 65009
  neighbor 172.16.11.9 route-map LOCAL_PREF in
  neighbor 172.16.11.9 soft-reconfiguration inbound
  neighbor 172.16.11.9 route-map SETCOMMUNITY out
  neighbor 172.16.11.10 remote-as 65010
  neighbor 172.16.11.10 route-map LOCAL_PREF in
  neighbor 172.16.11.10 soft-reconfiguration inbound
  neighbor 172.16.11.10 route-map SETCOMMUNITY out
  neighbor 192.168.3.2 remote-as 65003
  neighbor 192.168.3.2 next-hop-self
  neighbor 192.168.3.3 remote-as 65003
  neighbor 192.168.3.3 next-hop-self
!
access-list 1 permit 3.3.3.1
!
route-map inbound permit 10
  match community inbound
  set local-preference 123
!
route-map SETCOMMUNITY permit 10
  match ip address 1
  set community 65003:12
```

## BGP Lab 3 – BGP Communities Attributes

### Router 172.16.12.2 configuration

```
router bgp 65003
  bgp router-id 3.3.3.2
  network 3.3.3.2/32
  network 192.168.3.0/24
  neighbor 172.16.12.1 remote-as 65001
  neighbor 172.16.12.1 route-map LOCAL_PREF in
  neighbor 172.16.12.1 soft-reconfiguration inbound
  neighbor 172.16.12.1 route-map SETCOMMUNITY out
  neighbor 172.16.12.2 remote-as 65002
  neighbor 172.16.12.2 route-map LOCAL_PREF in
  neighbor 172.16.12.2 soft-reconfiguration inbound
  neighbor 172.16.12.2 route-map SETCOMMUNITY out
  neighbor 172.16.12.4 remote-as 65004
  neighbor 172.16.12.4 route-map LOCAL_PREF in
  neighbor 172.16.12.4 soft-reconfiguration inbound
  neighbor 172.16.12.4 route-map SETCOMMUNITY out
  neighbor 172.16.12.5 remote-as 65005
  neighbor 172.16.12.5 route-map LOCAL_PREF in
  neighbor 172.16.12.5 soft-reconfiguration inbound
  neighbor 172.16.12.5 route-map SETCOMMUNITY out
  neighbor 172.16.12.6 remote-as 65006
  neighbor 172.16.12.6 route-map LOCAL_PREF in
  neighbor 172.16.12.6 soft-reconfiguration inbound
  neighbor 172.16.12.6 route-map SETCOMMUNITY out
  neighbor 172.16.12.7 remote-as 65007
  neighbor 172.16.12.7 route-map LOCAL_PREF in
  neighbor 172.16.12.7 soft-reconfiguration inbound
  neighbor 172.16.12.7 route-map SETCOMMUNITY out
  neighbor 172.16.12.8 remote-as 65008
  neighbor 172.16.12.8 route-map LOCAL_PREF in
  neighbor 172.16.12.8 soft-reconfiguration inbound
  neighbor 172.16.12.8 route-map SETCOMMUNITY out
  neighbor 172.16.12.9 remote-as 65009
  neighbor 172.16.12.9 route-map LOCAL_PREF in
  neighbor 172.16.12.9 soft-reconfiguration inbound
  neighbor 172.16.12.9 route-map SETCOMMUNITY out
  neighbor 172.16.12.10 remote-as 65010
  neighbor 172.16.12.10 route-map LOCAL_PREF in
  neighbor 172.16.12.10 soft-reconfiguration inbound
  neighbor 172.16.12.10 route-map SETCOMMUNITY out
  neighbor 192.168.3.1 remote-as 65003
  neighbor 192.168.3.1 next-hop-self
  neighbor 192.168.3.3 remote-as 65003
  neighbor 192.168.3.3 next-hop-self
!
access-list 1 permit 3.3.3.2
!
route-map inbound permit 10
  match community inbound
  set local-preference 123
!
route-map SETCOMMUNITY permit 10
  match ip address 1
  set community 65003:12
```

## BGP Lab 3 – BGP Communities Attributes

### Router 172.16.13.3 configuration

```
router bgp 65003
  bgp router-id 3.3.3.2
  network 3.3.3.2/32
  network 192.168.3.0/24
  neighbor 172.16.13.1 remote-as 65001
  neighbor 172.16.13.1 route-map LOCAL_PREF in
  neighbor 172.16.13.1 soft-reconfiguration inbound
  neighbor 172.16.13.1 route-map SETCOMMUNITY out
  neighbor 172.16.13.2 remote-as 65002
  neighbor 172.16.13.2 route-map LOCAL_PREF in
  neighbor 172.16.13.2 soft-reconfiguration inbound
  neighbor 172.16.13.2 route-map SETCOMMUNITY out
  neighbor 172.16.13.4 remote-as 65004
  neighbor 172.16.13.4 route-map LOCAL_PREF in
  neighbor 172.16.13.4 soft-reconfiguration inbound
  neighbor 172.16.13.4 route-map SETCOMMUNITY out
  neighbor 172.16.13.5 remote-as 65005
  neighbor 172.16.13.5 route-map LOCAL_PREF in
  neighbor 172.16.13.5 soft-reconfiguration inbound
  neighbor 172.16.13.5 route-map SETCOMMUNITY out
  neighbor 172.16.13.6 remote-as 65006
  neighbor 172.16.13.6 route-map LOCAL_PREF in
  neighbor 172.16.13.6 soft-reconfiguration inbound
  neighbor 172.16.13.6 route-map SETCOMMUNITY out
  neighbor 172.16.13.7 remote-as 65007
  neighbor 172.16.13.7 route-map LOCAL_PREF in
  neighbor 172.16.13.7 soft-reconfiguration inbound
  neighbor 172.16.13.7 route-map SETCOMMUNITY out
  neighbor 172.16.13.8 remote-as 65008
  neighbor 172.16.13.8 route-map LOCAL_PREF in
  neighbor 172.16.13.8 soft-reconfiguration inbound
  neighbor 172.16.13.8 route-map SETCOMMUNITY out
  neighbor 172.16.13.9 remote-as 65009
  neighbor 172.16.13.9 route-map LOCAL_PREF in
  neighbor 172.16.13.9 soft-reconfiguration inbound
  neighbor 172.16.13.9 route-map SETCOMMUNITY out
  neighbor 172.16.13.10 remote-as 65010
  neighbor 172.16.13.10 route-map LOCAL_PREF in
  neighbor 172.16.13.10 soft-reconfiguration inbound
  neighbor 172.16.13.10 route-map SETCOMMUNITY out
  neighbor 192.168.3.1 remote-as 65003
  neighbor 192.168.3.1 next-hop-self
  neighbor 192.168.3.2 remote-as 65003
  neighbor 192.168.3.2 next-hop-self
!
access-list 1 permit 3.3.3.3
!
route-map inbound permit 10
  match community inbound
  set local-preference 123
!
route-map SETCOMMUNITY permit 10
  match ip address 1
  set community 65003:12
```

## BGP Lab 3 – BGP Communities Attributes

### Show IP BGP from router 172.16.12.4

\*Note at the time of the show IP BGP only 2 tables were running BGP routers

BGP table version is 0, local router ID is 3.3.3.2

Status codes: s suppressed, d damped, h history, \* valid, > best, i - internal,

r RIB-failure, S Stale

Origin codes: i - IGP, e - EGP, ? - incomplete

Network	Next Hop	Metric	LocPrf	Weight	Path
*> 3.3.3.1/32	0.0.0.0	0		32768	i
*> 3.3.3.2/32	0.0.0.0	0		32768	i
*> 3.3.3.3/32	0.0.0.0	0		32768	i
*> 4.4.4.1/32	172.16.12.4	0		0	65004 i
*> 4.4.4.2/32	172.16.12.4	0		0	65004 i
*> 4.4.4.2/32	172.16.12.4	0		0	65004 i
*> 192.168.3.0	0.0.0.0	0		32768	i

Total number of prefixes 7