

Compatible Systems Setup Guides: BGP Configuration Guide

Document ID: 17612

- Introduction**
- Prerequisites**
 - Requirements
 - Components Used
- BGP General Configuration**
- BGP Peer Configuration**
 - Sample Peer Configuration
- BGP Route Advertisement Policy**
- BGP Networks**
- BGP Aggregate Configuration**
- IP Routing Protocol Redistribution**
 - Redistributing Static Routes into BGP
- BGP Route Map Configuration**
 - BGP Routing Mapping Rules
- Summary of BGP Route Selection Process**
- IP Route Filters and BGP**
- BGP Console Commands**
 - Show BGP rtcount
 - Show BGP Routes
 - Show BGP Peers
 - Show BGP Networks
 - Show BGP Stats
 - Show BGP Timers
 - Show BGP Mem
 - Show BGP Config
 - Show BGP Aggregates
 - BGP Disable
 - Reset BGP Peer
- BGP Quick Start Guide**
- BGP Debug Options**
- BGP RFC References**
- Related Information**

Introduction

The Border Gateway Protocol (BGP) is an exterior gateway protocol that allows Autonomous Systems to exchange routing information with each other. An Autonomous System is a set of routers under a single technical administration.

Autonomous System (AS) numbers are assigned by the American Registry for Internet Numbers. For further information, see their Web site. It includes a full listing of all assigned AS numbers under the Documentation section.

American Registry for Internet Numbers [↗](#)

It is possible, but not encouraged, to apply for an AS number to run BGP if an installation is single-homed.

However, a separate AS number is required for a multi-homed site where more than one ISP is used. This is because a single-homed installation could be considered internal to the ISP, whereas a multi-homed site cannot.

Routers that exchange BGP information are called BGP Peers. A router may have both external peers in other AS'es, and internal peers within its own AS. A peer is considered external if its AS number differs from the router's own AS number.

Routers establish BGP sessions using the TCP protocol. Upon startup of a new BGP session, BGP peers will exchange their full routing tables, and then only incremental updates are sent as the routing table changes.

This configuration guide describes the configuration options that are available with BGP running on Compatible Systems routers.

Prerequisites

Requirements

There are no specific requirements for this document.

Components Used

This document is not restricted to specific software and hardware versions.

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, make sure that you understand the potential impact of any command.

BGP General Configuration

The BGP protocol is enabled in the **BGP General** configuration section. BGP is enabled globally for the router rather than per interface, as RIP and OSPF are. BGP is **Off** by default. To enable BGP, you must set the **BGPEnabled** parameter to **On**.

```
[ BGP General ]
```

```
BGPEnabled      = Off      Enable or disable the BGP protocol
BGPAS           = ""      Autonomous system number of this router
BGPLocPref      = 100     BGP local preference, default is 100
BGPUseIPRFiltrs = False   Use IP Route Filters, default is False
```

The Autonomous System (AS) number of this router is set here. The **BGPAS** number must be provided; if it is not, BGP will not be enabled.

The local preference attribute **BGPLocPref** is exchanged among routers in the same AS, and is an indication about which path is preferred to exit the AS; a path with a higher local preference is more preferred. The default of 100 will be used if no **BGPLocPref** is specified.

BGP uses BGP Route Maps to filter routes and set attributes. More information on these are available in the BGP Peer Config and BGP Route Map sections of this document. The user has the option to use IP Route Filters instead of BGP Route Maps. The value of **BGPUseIPRFiltrs** will be checked for each peer which has no BGP Route Maps defined, and if it is TRUE, the IP Route Filters will be checked for that peer. Note that IP Route Filters are global to the router, whereas BGP Route Maps can be made specific to each peer.

BGP Peer Configuration

The **BGP Peer List** contains the list of configured peers for this router. The router will not establish a BGP connection with any router not on this list. If there is no **BGP Peer List**, BGP will not be enabled even if **BGPEnabled** is set to **On** in the **BGP General** section.

```
[ BGP Peer List ]  
  
BGPPeer = On/Off IPAddress ASNumber PeerConfigID
```

The **On|Off** parameter configures the start-up state of the router with respect to the peer; it determines whether the router will automatically try to establish a BGP session with the peer at startup. If this parameter is set to **Off**, the router will not establish a BGP session with the peer until you issue the **BGP Enable** command. Note that this will not change the start-up state; the next time you boot the router, the peer will come up in the **Off** state until you enable it.

You can configure BGP so that all peers are **Off** at startup. If **BGPEnabled = On** in the **BGP General** section, you will be able to dynamically enable selected peers after router startup.

The router will contact the peer using the **IPAddress** given in the configuration list. The **IPAddress** and **ASNumber** of the peer must be provided. The router must have the network of the supplied IP address in its routing table in order for the session to be established. The router determines if a peer is internal or external from the AS number of the peer, since internal peers have the same AS number as the router itself.

Each **BGP Peer List** entry may contain an optional **PeerConfigID**, which specifies the number of the **BGP Peer Config** section where various peer-specific BGP configuration items may be set. A **BGP Peer Config** section may be used for more than one peer only if all the same parameters are desired.

```
[ BGP Peer Config "SectionID" ]    Section ID is a character string  
  
InputRouteMap = ""      Name of input Route Map to be used for this peer  
OutputRouteMap = ""     Name of output Route Map to be used for this peer  
NextHopSelf = False    Next hop is this router  
EBGPMultihop = False   External peer not directly connected  
PeerWeight = 100      Neighbor weight  
PeerRetryTime = 30     Retry time in seconds  
PeerHoldTime = 180    Configured hold time in seconds  
BGPUseLoopback = False Use router LoopbackAddress with this peer  
AdvertiseDefault = False Advertise default route to this peer
```

Note that the **InputRouteMap** and **OutputRouteMap** are specified separately. The parameters which may be set and checked are different for input and output routes (see BGP Route Map section for details).

If **NextHopSelf** is set to **TRUE**, the router will advertise itself as the next hop to the routes it advertises to this peer.

External peers are required to be directly connected, unless **EBGPMultihop** is set to **TRUE**. If this parameter is set to **TRUE**, the router must have a route to the non-directly connected external peer in order to establish a connection.

The **PeerWeight** parameter is an internal rating assigned to the peer by the administrator; it is not advertised to other routers. Peers with a higher weight are preferred when multiple routes exist to the same destination.

The **BGP Retry Time** allows the administrator to set the amount of time between retries to establish a connection to configured peers which have gone down for some reason. If a peer is down but its state is set to **On**, the router will continually try to contact the peer every **PeerRetryTime** seconds. The minimum accepted

PeerRetryTime is 10 seconds.

The Hold Time is negotiated with the peer, so the configured **PeerHoldTime** will not necessarily end up being the actual hold time used by the peers. The peers will use the smaller of the two hold times proposed. The hold time must be either zero or at least 3 seconds. If the negotiated Hold Time interval is zero, then periodic KEEPALIVE messages will not be sent.

If no **PeerWeight**, **PeerHoldTime** or **PeerRetryTime** are provided, the defaults will be used. The default **PeerWeight** is 100, the default **PeerHoldTime** is 180 seconds, and the default **PeerRetryTime** is 30 seconds.

If a **LoopbackAddress** is specified in the **IP Loopback** section, **BGPUseLoopback** may be set to TRUE. In that case, the router will use its Loopback Address as the IP source in TCP packets to that peer rather than a specific IP address of one of its interfaces. Note, however, that the peer must know how to send packets to that address via normal IP routing procedures. If the address is not on a subnet already known to the peer, it must be added via a static route. The loopback address is normally only used for internal peers, since external peers are usually directly connected.

The router's default route is not advertised to a peer unless the parameter **AdvertiseDefault** is set to TRUE for that peer.

Sample Peer Configuration

This is a sample peer configuration:

```
[ BGP Peer List ]
BGPPeer = On   198.41.11.213   100   Peer1
BGPPeer = On   205.14.128.1    110   Peer2

[ BGP Peer Config "Peer1" ]
InputRouteMap      = bgpin1
OutputRouteMap     = bgpout1
PeerHoldTime       = 180
PeerRetryTime      = 65
PeerWeight         = 1000

[ BGP Peer Config "Peer2" ]
InputRouteMap      = bgpin2
OutputRouteMap     = bgpout1
PeerHoldTime       = 180
PeerRetryTime      = 45
PeerWeight         = 2000
```

In the **BGP Peer List** and **BGP Peer Config** Peers 198.41.11.213 and 206.14.128.2 use **BGP Peer Config 1**, and Peer 205.14.128.1 uses **BGP Peer Config 2**.

BGP Route Advertisement Policy

The default for BGP is to NOT advertise routes. This is to prevent inadvertent advertisement of routes out on the Internet.

To get routes advertised, you have to configure something: The BGP Networks list, IP Route Redistribution, BGP Route Maps, or IP Route Filters.

To get external routes advertised, use BGP Route Maps or IP Route Filters. To get internal routes advertised, use the BGP Networks List or IP Route Redistribution.

Each of these configuration sections is described below.

BGP Networks

The **BGP Networks** section defines a list of routes that the administrator wishes to advertise as originating inside the AS. These may be directly connected routes, static routes, RIP routes, or OSPF routes.

The router compares the entries in the BGP Networks list with its IP routing table, and will not advertise a route in the Networks list that it cannot find in its IP routing table. Therefore, if you want to advertise local networks which are not in the router's own IP routing table, you will need to add static routes.

Note that the only way to get directly connected routes advertised into BGP is to include them on the Network List. OSPF or RIP routes can be advertised into BGP using the **IP Route Redistribution** section. Static routes can be advertised into BGP using the redistribute flag on each configured static route.

The optional mask parameter tells the router how many bits of the IP routing table entry to match against the LocalNet address. This is not necessarily the actual mask of the network you wish to advertise. For instance, suppose the router has subnets 198.41.9.32, 198.41.9.64, and 198.41.9.96, all with mask 255.255.255.224. To get BGP to advertise one 198.41.9.0/24 network, your **BGP Networks** would look like this:

```
[ BGP Networks ]
LocalNet = IP address [mask]

[ BGP Networks ]
LocalNet = 198.41.9.32 255.255.255.255
```

The router will match only the 198.41.9.32 entry due to the mask you supplied with the LocalNet. It will advertise the network as 198.41.9.0/24, since it automatically truncates subnet masks more specific than Class C. However, if you provided a mask of 255.255.255.0, you would end up advertising the 198.41.9.0/24 net three times, since all three of your subnets would match the LocalNet entry. This truncation is not same as aggregation, and only applies to internal networks, and only to masks more specific than Class C. To get route aggregation, use the BGP Aggregates section.

BGP Aggregate Configuration

The **BGP Aggregates** section contains networks which are to be aggregated before being advertised to external peers. The router's IP routing table must contain networks which are a subset of the aggregate in order for the aggregate to be advertised; only the aggregate, and not the individual routes, will be advertised to external peers. Internal peers will receive the individual routes if they originated outside the AS; internal peers do not exchange internal routes via BGP.

It is not necessary to have an aggregate list for internal subnets of Class C networks (see BGP Networks section above). But if you have several class C's (or greater) that can be combined with a single mask to a supernet, aggregation can be used.

```
[ BGP Aggregates ]
AddrAndMask = [IPAddr] [IPMask]

IP Routing Table Entries
198.41.8.0      255.255.255.0
198.41.9.0      255.255.255.0
198.41.10.0     255.255.255.0
198.41.11.0    255.255.255.0

[ BGP Networks ]
LocalNet = 198.41.8.0 255.255.252.0
```

```
[ BGP Aggregates ]
AddrAndMask = 198.41.8.0 255.255.252.0
```

The single route 198.41.8.0/22 will be advertised to BGP external peers. Without the **BGP Aggregates** entry, the four networks would be advertised separately. All four of the networks would match the mask provided in the **BGP Networks** section, but they would not automatically be aggregated.

IP Routing Protocol Redistribution

Another way to specify RIP and OSPF routes to be imported into BGP is by using route redistribution. The default is for all routing redistribution to be disabled.

It is possible to redistribute BGP routes into RIP and OSPF, but it is not recommended unless you are only accepting a small number of BGP routes. Care must be taken with appropriate filters when doing things like importing BGP routes into OSPF and then exporting OSPF routes into BGP.

Note: The number of routes supported will also depend on the amount of memory the router has.

```
[ IP Route Redistribution ]

BGPtoOSPF      Redistribute BGP routes to OSPF
                Syntax: [True|False] [Metric]
BGPtoRIP       Redistribute BGP routes to RIP
                Syntax: [True|False] [Metric]
RIPtoBGP       Redistribute RIP routes into BGP
OSPFtoBGP      Redistribute OSPF routes into BGP
```

Redistributing Static Routes into BGP

A static route may be redistributed into BGP by using the redistribute flag when configuring the route in the **IP Static** section:

```
[ IP Static ]
198.41.16.0 255.255.255.0 198.41.9.65 1 Redist=BGP
```

BGP Route Map Configuration

BGP Route Maps are very similar to IP Route Filters, except:

- They are specific to BGP
- They can be specified on a per-peer basis
- They allow BGP attributes to be set on incoming and outgoing routes in addition to filtering routes

Route maps are used only by the BGP protocol, and are not associated with a particular interface. The **BGP Peer Config** section specifies the route maps, if any, to be applied to the peer. Input route maps and output route maps are specified separately.

BGP routes known to the router will be advertised unless denied by a route map or a route filter. Static, IGP, and directly connected routes will not be advertised unless specified in the BGP Networks section or by route redistribution.

No input routes will be accepted by the router unless a BGP Route Map or IP Route Filter has been defined. If you really want everything, a "permit 0.0.0.0" will do it. The router checks BGP route maps first, and if the route is denied, the IP route filters will not be checked even if **BGPUseIPRFIters** is True.

```
[ BGP Peer Config 2 ]
InputRouteMap          = bgpin2
OutputRouteMap         = bgpout2
```

IP Route Filters may be used with BGP instead of **BGP Route Maps**. The matching conditions are more limited, and various parameters such as community, local preference, and weight cannot be set with **IP Route Filters**.

The **BGP Route Map** *name* is a special section of the configuration, meaning that there are no keywords to document. Each section contains a complete filter set uniquely identified by the *Name* portion of the section name. Multiple sections may exist, each with a unique name. The name must be 15 characters or less.

BGP Routing Mapping Rules

This section details the parameters and modifiers pertinent to BGP Route Mapping Rules.

```
action route [direction] [out | in modifiers]
permit | deny IP Address out | in
```

The **action**, **route** and **direction** are required parameters. **In** and **out** modifiers are optional.

Action – Permit or Deny

This specifies the action to be taken when a route meets the condition of the rule.

Route – IP Address of Network

The IP Address is specified in the same fashion as described for IP Route Filters; that is, in normal dotted decimal notation, as a factorized address, a hexadecimal number, or with an optional /bits field. See the IP Route Filter manual page for details.

[Direction]

An **in** or **out** parameter must be supplied. This specifies the direction for which the rule is applied.

These modifiers apply if the direction is in:

- **ipaddr** IP Address of peer
- **srcas** route has this source AS number
- **hasas** this AS number is contained in AS path
- **nhop** route has this next hop
- **comm** this community is contained in attribute list
- **setpref** set preference to this value
- **setwt** set weight to this value

The **ipaddr** | **hasas** | **srcas** | **comm** | **nhop** modifiers limit input rules to routes originating from the designated IP address, AS number, community, or next hop. Only one of these five arguments is expected here. **hasas** means that the rule will be applied if the AS path contains the specified AS number anywhere in the AS path; **srcas** means that the rule will be applied only if the route originated in the specified AS.

The **setpref** modifier allows the preference to be set on incoming routes. If an **ipaddr**, **hasas**, **srcas**, **comm**, or **nhop** is supplied, the preference will only be set for routes that match that condition.

The **setwt** modifier allows the weight to be set on incoming routes. If an **ipaddr**, **hasas**, **srcas**, **comm**, or **nhop** is supplied, the weight will only be set for routes that match that condition.

These modifiers apply if the direction is out:

- **ipaddr** IP Address of peer
- **toas** AS number of peer
- **srcas** source AS number of the route
- **origin** protocol the route came from
- **setnhop** set next hop attribute
- **setmed** set multi-exit discriminator attribute
- **setasp** prepend an AS path to the current path
- **setcomm** set a new community list, discarding old
- **addcomm** prepend a community list to existing one

The **ipaddr** | **toas** modifiers limit output rules to routes going to the designated IP address or AS number. Only one argument is expected here. If the router only has one peer in a given AS, then **ipaddr** or **toas** will accomplish the same result. If the router has multiple peers within a neighboring AS, use the IP address of the peer to limit the rule to just that peer, or use the AS number to apply the rule to every peer in the AS.

The **srcas** modifier limits output rules to routes originating from the designated AS number.

The **origin** protocol modifier limits output rules to routes originating from the designated protocol. BGP can advertise direct, static, RIP, OSPF, or other BGP routes from its own IP routing table to peers.

The **setnhop** modifier allows the next hop to be set on the outgoing route.

The **setmed** modifier allows the multi-exit discriminator to be set on the outgoing route.

The **setasp** modifier allows the specified AS list to be prepended to the outgoing AS path attribute. Up to 6 AS numbers may be entered.

The **setcomm** modifier allows a community list to be set on the outgoing route. The parameters can either be up to 6 community numbers, or one of the special communities: "noexport", "noadv", or "noexpsub". These are the three "well-known" communities defined in RFC 1997, BGP Communities Attribute: NO_EXPORT, NO_ADVERTISE, and NO_EXPORT_SUBCONFED.

The **addcomm** modifier allows a community list to be prepended on the outgoing route. The parameters can be up to 6 community numbers.

Examples

In BGP Route Map **mymapin**, route 192.61.5.0 will be permitted in if the Community Attribute contains the community 200, and the preference will be set to 100. In line two, all other routes from Community 200 will also be accepted, but the preference will be set to 300. Routes that do not contain Community 200 will be denied.

In BGP Route Map **mymapout**, all direct routes specified in the BGP Networks section will be allowed out to AS number 200, and the MED will be set to 10. In the second line, all routes will be allowed out to AS number 300, but the Community value will be set to **noadv** (NO_ADVERTISE).

```
[ BGP Route Map "mymapin" ]
  permit 192.61.5.0 in comm 200 setpref 100
  permit 0.0.0.0 in comm 200 setpref 300

[ BGP Route Map "mymapout" ]
  permit 0.0.0.0 out toas 200 origin direct setmed 10
  permit 0.0.0.0 out toas 300 setcomm noadv
```

Summary of BGP Route Selection Process

Route maps help the administrator influence the route selection process, since BGP uses weight, preference, and MED, among other things. BGP uses the following criteria, in the order presented, to select its best route for a destination:

- The most preferred path is the path with the largest weight.
- If the weights are the same, select the path with the largest local preference.
- If the preferences are the same, select the path that has the shortest AS path length.
- If all paths have the same AS path length, select the path with the lowest MED.
- If the paths have the same MED, select the path from the BGP peer with the lowest Router ID.

IP Route Filters and BGP

The user has the option of using **IP Route Filters** with BGP instead of **BGP Route Maps**; however, **IP Route Filters** do not provide the ability to set BGP attributes as described in the **BGP Route Map** section. If an **InputRouteMap** has been defined for a peer, the IP Route Filters will be ignored for input routes even if the **BGPUseIPRFltrs** parameter has been set to TRUE in the **BGP General** section. Likewise, if an **OutputRouteMap** has been defined for a peer, the IP Route Filters will be ignored for output routes.

For BGP, an additional parameter has been added to IP route filtering, and this is filtering based on AS path. A BGP route contains information concerning each Autonomous System (AS) that it has traversed. Route 199.41.13.0, originating in AS 500, would have two AS paths to reach R1: [200,300,500] and [400,600,500].

In the following example, **IP Route Filter *bgpin*** applies to Router R1. All routes originating from AS 300 will be filtered out, and all routes originating from AS 400 will be permitted.

IP Route Filter *bgpout* allows 192.62.16.0 to be advertised to R2, and 192.62.17.0 to be advertised to R4. The IP addresses of R2 and R4 could be used instead of AS numbers in **bgpout**.


IP Route Filter *bgp600* illustrates the use of the **contains** keyword. This filter would deny any incoming routes that contained AS 600 anywhere in their AS path.

Note the final line in the route filters to prevent unintended filtering of RIP and OSPF routes:

```
[ IP Route Filter "bgpin" ]
deny 0.0.0.0 in via bgp from 300
permit 0.0.0.0 in via bgp from 400
permit 0.0.0.0 in via rip ospf

[ IP Route Filter "bgpout" ]
permit 192.62.16.0 out via bgp to 200
permit 192.62.17.0 out via bgp to 400
permit 0.0.0.0 out via rip ospf

[ IP Route Filter "bgp600" ]
deny 0.0.0.0 in via bgp contains 600
permit 0.0.0.0 in via rip ospf
```



| AS400 | | AS600 | | AS500 |

You cannot, however, do the following with AS filtering, because the AS filter applies to the origin of the route. Say that Router R1 is receiving an advertisement about route 199.41.13.0 from both its peers R2 and R4, and that the route originates in AS 500. The AS path for the route from R2 is therefore [200,300,500], and the AS path for the same route from R4 is [400,600,500].

```
[ IP Route Filter "does not work as intended" ]
  deny 199.41.13.0 in via bgp from 200
  permit 199.41.13.0 in via bgp from 400
```

Although the syntax is correct, the above filter would merely cause the route to be rejected; it would not match the filter in line 2 because its source AS number is 500, not 400. To accomplish the purpose intended by the above, you can use the IP addresses of the peers R2 and R4:

```
[ IP Route Filter "bgpin" ]
  deny 199.41.13.0 in via BGP from "R2's IP address"
  permit 199.41.13.0 in via BGP from "R4's IP address"
```

BGP Console Commands

There are several show commands for BGP, and commands to enable/disable BGP or reset BGP connections:

show bgp rtcount	BGP Routing Entry Counts
show bgp routes	Display BGP Routing Entries
show bgp peers	Display the list of BGP Peers and current status
show bgp timers	BGP Peer timer information
show bgp mem	BGP Database Memory Allocation
show bgp config	BGP configuration information
show bgp stats	BGP peer uptime and packet exchange statistics
show bgp networks	Display list of internal networks to be advertised
show bgp aggregates	Display BGP routes to be aggregated
bgp disable	Disable BGP connection to all peers or 1 specified peer Usage: { ALL IP Address }
bgp enable	Enable BGP connection to all peers or 1 specified peer Usage: { ALL IP Address }
bgp reset peer	Reset BGP connection to all peers or 1 specified peer Usage: { ALL IP Address }

Show BGP rtcount

This command displays a summary of the number of routes in the BGP Routing database. With BGP, this is useful if there are a very large number of routes and you want to know how many, but not print them all out.

```
BGP Test> sho bgp rt

BGP Routing Database Entries      In Use      Added      Removed
In IP routing table:              51548       78694      27146
BGP route heads:                  51548       78702      27154

IP Routing Table Entries: 51561
```

Show BGP Routes

The **show bgp routes** command, with no arguments, displays the best route in the BGP routing database for each destination. A sample excerpt is shown below.

The BGP routing database may contain routes that are not in the router's IP routing table; a BGP route will not be present in the IP routing table if the router did not have an entry for the next hop of that route.

```
bgptest>sho bgp ro
```

```
BGP Best Routes List
```

	Network/Mask	Bits	Pref	Weight	Next Hop	AS Path
1	128.128.0.0	/16	100	100	199.45.133.101	3404 1 1
2	129.129.0.0	/16	100	100	199.45.133.101	3404 1 1239 1673 1133 559
3	130.130.0.0	/16	100	100	199.45.133.101	3404 1 1 5727 7474 7570
4	131.131.0.0	/16	100	100	199.45.133.101	3404 1 1 1236
5	134.134.0.0	/16	100	100	199.45.133.101	3404 1 1239 1760 4983
6	135.135.0.0	/16	100	100	199.45.133.101	3404 3561 3561 4293
7	139.139.0.0	/16	100	100	199.45.133.101	3404 1 1239 568 1913 1569
8	140.140.0.0	/16	100	100	199.45.133.101	3404 1 1239 7170 374
9	141.141.0.0	/16	100	100	199.45.133.101	3404 1 1239 3739 3739 3739
10	142.142.0.0	/16	100	100	199.45.133.101	3404 3561 3561 577 549 808
11	147.147.0.0	/16	100	100	199.45.133.101	3404 3561 3561 5400 2856
12	149.149.0.0	/16	100	100	199.45.133.101	3404 1 1 3749
13	150.150.0.0	/16	100	100	199.45.133.101	3404 3561 3561 3786 6068
14	151.151.0.0	/16	100	100	199.45.133.101	3404 1 1239 174
15	152.152.0.0	/16	100	100	199.45.133.101	3404 1 1 286 1891
16	155.155.0.0	/16	100	100	199.45.133.101	3404 1 701 702 8413 1913 1564
17	158.158.0.0	/16	100	100	199.45.133.101	3404 3561 3561
18	161.161.0.0	/16	100	100	199.45.133.101	3404 1 1239 174
19	164.164.0.0	/16	100	100	199.45.133.101	3404 1 701 7633
20	165.165.0.0	/16	100	100	199.45.133.101	3404 1 701 5713

The show command may also be invoked with a specific route, in which case it will display all the paths for that route.

```
BGP 2600>sho bgp ro 129.129.0.0
```

```
BGP routing table entry for 129.129.0.0/16
```

```
Paths: (in order of preference, best first)
```

```
AS path 11129 3404 1239 1673 1133 559
```

```
Next hop 198.41.11.1 from peer 198.41.11.17 (RtrID 198.41.11.17)
```

```
Origin IGP, localpref 100, weight 100
```

```
AS path 12345 11129 3404 1239 1673 1133 559
```

```
Next hop 198.41.11.1 from peer 198.41.11.201 (RtrID 198.41.11.201)
```

```
Origin IGP, localpref 100, weight 100
```

If just an IP address is entered, the most specific route will be displayed. To display a less specific route with the same IP address, enter the mask also.

BGP routes are displayed using CIDR notation: Network/Mask Bits, rather than Route/Mask.

The preference and weight may be set using **BGP Route Maps**. If they are not, the default Local Preference and Weight values will be used.

The complete AS path is shown, with the source AS being the one farthest to the right. Each AS which passes the route on will prepend its own AS to the AS path attribute.

An IP Routing Table excerpt for the **show ip routing** command with BGP routes is shown below. For BGP, the Metric is the path length, just like for RIP. Most BGP routes are IGP, which means they originated in an interior gateway protocol. The other possibilities are EGP (exterior gateway protocol) or Incomplete (usually means a static route).

```
bgptest> sho ip ro dynamic bgp
```

```
Dynamic Routes:
Destination      Mask      Gateway      Metric    Uses   Type   Src/TTL  Interface
3.0.0.0          FF000000  198.41.11.1  5         0     BGP    INC      Ether0
6.0.0.0          FF000000  198.41.11.1  6         0     BGP    INC      Ether0
9.2.0.0          FFFF0000  198.41.11.1  6         0     BGP    IGP      Ether0
9.20.0.0         FFFF8000  198.41.11.1  6         0     BGP    INC      Ether0
12.0.0.0         FF000000  198.41.11.1  5         0     BGP    IGP      Ether0
12.2.97.0        FFFFFFF00  198.41.11.1  6         0     BGP    IGP      Ether0
12.2.183.0       FFFFFFF00  198.41.11.1  4         0     BGP    IGP      Ether0
12.4.164.0       FFFFFFF00  198.41.11.1  5         0     BGP    IGP      Ether0
12.5.164.0       FFFFFFF00  198.41.11.1  5         0     BGP    IGP      Ether0
12.5.252.0       FFFFFFFE00 198.41.11.1  6         0     BGP    IGP      Ether0
12.6.42.0        FFFFFFFE00 198.41.11.1  6         0     BGP    IGP      Ether0
12.7.214.0       FFFFFFFE00 198.41.11.1  11        0     BGP    IGP      Ether0
12.8.188.0       FFFFFC00  198.41.11.1  5         0     BGP    IGP      Ether0
12.8.188.0       FFFFFFF00  198.41.11.1  5         0     BGP    INC      Ether0
12.8.189.0       FFFFFFF00  198.41.11.1  5         0     BGP    INC      Ether0
12.8.191.0       FFFFFFF00  198.41.11.1  5         0     BGP    INC      Ether0
12.10.14.0       FFFFFFFE00 198.41.11.1  5         0     BGP    INC      Ether0
12.10.152.0      FFFFFF800  198.41.11.1  5         0     BGP    IGP      Ether0
12.10.231.0      FFFFFFF00  198.41.11.1  6         0     BGP    IGP      Ether0
12.11.134.0     FFFFFFFE00 198.41.11.1  5         0     BGP    IGP      Ether0
```

Show BGP Peers

The **show bgp peers** command displays the configured BGP peers of this router, with information about the AS number of the peer, the Router ID, the IP address, the TCP socket number, the Enable Status, and the BGP connect state.

```
bgptest>sho bgp peers
```

```
=====
                        BGP PEER STATUS
-----
Int  AS      Router      IP          TCP   Enable  BGP
Ext  Number  ID          Address     Socket Status  State
-----
Ext  23456   0.0.0.0     198.14.13.18  0     Off     IDLE
Ext  34567   198.41.11.6 198.14.12.6  82     On      ESTABLISHED
Int  11129   0.0.0.0     198.41.11.17  0     Off     IDLE
Int  11129   0.0.0.0     198.41.11.2  0     On      ACTIVE
=====
```

Int/Ext indicates whether this is an internal or external peer. (An internal peer has the same AS number as the router itself.) The AS Number of the peer is configured in the BGP Peer List.

The **Router ID** is not known until the peer contacts the router, so if the connect state is **IDLE**, **ACTIVE**, or **CONNECT**, this parameter might be 0. The Router ID is usually the IP address of one of the peer's interfaces, and may or may not be the same as the IP address.

The **Enable Status** indicates whether the router will currently accept a connection request from this peer. The peer can be brought up as enabled by setting the peer to **On** in the BGP Peer List. Also, the peer can be dynamically enabled or disabled by the **BGP Peer Enable** and **BGP Peer Disable** commands. When the Enable Status is Off, the BGP State is always **IDLE**.

The BGP connect states are: **IDLE**, **ACTIVE**, **CONNECT**, **OPENSENT**, **OPENCONFIRM**, and **ESTABLISHED**. The connect state is established by active negotiations between the peers. In the **IDLE** state, the router will not accept connections from the peer. This state is entered briefly after a connection has timed out, to prevent too-rapid up-and-down transitions of peers. In the **ACTIVE** state, the router is

listening on its server port for connection requests from the peer. In the **CONNECT** state, the router has sent out an active TCP connection request to the peer. In the **OPENSENT** and **OPENCONFIRM** states, the two peers are exchanging preliminary packets in order to establish their BGP session. If the exchanges are successful, the peers will enter the **ESTABLISHED** state. The peers must continue to exchange periodic **KEEPALIVE** packets to remain in the established state, unless the negotiated hold time is 0.

BGP communicates with its peers via TCP. Therefore, further information about BGP sessions can be obtained with the "show os tcp" command. The TCP states are not the same as BGP states, but are the standard TCP states (**LISTEN**, **SYNSENT**, **SYNRCVD**, **ESTABLISHED**, **FINWAIT1**, **FINWAIT2**, **CLOSEWAIT**, **LASTACK**, **CLOSING**, **TIMEWAIT**). BGP uses port 179 to listen for BGP connection attempts.

```
bgptest>sho os tcp
=====
                        TCP SESSION INFORMATION
-----
Num  Session Type      State      Socket  Local  Remote  Remote
-----
     1  SERVER (TELNET)    LISTEN     80      23     0       0.0.0.0
     2  SERVER (BGP)      LISTEN     81      179    0       0.0.0.0
     3  ACTIVE (BGP)      ESTABLISH  82      20001  179     198.41.9.2
-----

13 free TCBS out of 16.
=====
```

Show BGP Networks

The **show bgp networks** command displays the list of internal networks to be advertised to external BGP peers.

```
bgptest>sho bgp networks

BGP NETWORKS:  2
Address          Mask
198.41.11.0      255.255.255.0
209.14.128.0    255.255.255.0
```

Show BGP Stats

The **show bgp stats** command displays statistics about packet types received from and sent to BGP peers, and the current uptime of the peer.

```
BGP Test>sho bgp stats

Open messages:          Received      Sent
Keepalive messages:    4069         4124
Notify messages:       0            0

BGP External Peer 198.41.11.6 state ESTABLISHED
  6 peer sessions, current uptime 2 days 16 hours 40 minutes 19 secs
  0 updates received
  78791 updates sent, last at 6 secs
BGP Internal Peer 198.41.9.2 state ESTABLISHED
  1 peer sessions, current uptime 2 days 20 hours 42 minutes 28 secs
  88791 updates received, last at 7 secs
  0 updates sent
```

Show BGP Timers

The **show bgp timers** command displays the current time in seconds left on each timer associated with each peer. If the peer is in ESTABLISHED state, this will be the KEEPALIVE timer and the HOLD timer. If the peer is in ACTIVE state, this will be the CONNECT timer. If the peer is in IDLE state but enabled, this will be the AUTO ENABLE timer. If the peer is IDLE and disabled, no timers are active until the **bgp peer enable** command is issued.

```
BGP Test>sho bgp timers
```

```
=====
                        BGP TIMERS
-----
Peer Address      Status   State      Timers
-----
198.41.9.2        Enabled  ESTABLISHED  Send KEEPALIVE pkt: 2 secs
                  HOLD timer expires: 121 secs
198.14.13.2       Enabled  ACTIVE       Next CONNECT attempt: 16 secs
199.13.12.3       Enabled  IDLE         AUTO ENABLE: 112 secs
198.41.9.3        Disabled IDLE         No timers active
=====
```

When a peer is in ESTABLISHED state, the Keepalive timer indicates how many seconds until the router will send another KEEPALIVE packet to the peer. The Hold Timer indicates how many seconds until the Hold Timer for the peer will expire. The Hold Timer is set every time the router receives either an UPDATE or a KEEPALIVE packet from the peer. If the Hold Timer expires, the router will declare the peer down, transition the peer to IDLE state, and set the Auto Enable timer.

The Connect and Auto Enable Timers both indicate how many seconds remain until the router will once again try to contact the peer. The Connect timer is used when the peer is in ACTIVE state; in this state, the router will accept an incoming connection request from the peer before the Connect time expires. The Auto Enable timer is used when the peer is in IDLE state; in this state, the router will NOT accept a connection request from the peer until the Auto Enable time has expired. When the Auto Enable time expires, the peer will transition back into the ACTIVE state.

The purpose of the Auto Enable timer is to prevent peer sessions from going up and down at too fast a rate. Once a peer session has been interrupted for some reason, the peer is held down for a short period before a new session will be allowed.

Show BGP Mem

The **show bgp mem** command displays detailed dynamic memory usage information for BGP.

```
BGP Test>sho bgp mem
```

```
ROUTING DATABASE DYNAMIC MEMORY USAGE
-----
Memory Block      Allocs      Deallocs      Size (bytes)
-----
ip radix nodes                    1976180
ip routing entries                4332132
bgp ip routes          78709        27149
bgp routes             78717        27157        2062400
bgp int change         0            0
bgp aggregates         0            0
bgp agg paths          0            0
bgp timers             12           0            384
-----
Peer 198.41.9.2
```

bgp path entries	78728	27168	1443680
bgp transmit queues	0	0	0
bgp PA strings	28151	21181	1784320
bgp PA hdr entries	28151	21181	529720
bgp rejected routes	0	0	0
bgp rej entries	0	0	0
bgp history entries	0	0	0

Total Size			12128816

Show BGP Config

This command displays the **Router ID** of the router, the parameters set in the **BGP General** section, the route redistribution status, and the peer configuration parameters. Note that the **Router ID** of the router for BGP is the same as for OSPF, the largest IP address of the router's IP interfaces.

```
bgptest>sho bgp config
```

```

BGPEnabled          Yes
Router ID           205.14.128.2
BGP AS Number       100
BGP Local Preference 100
Use IP Route Filters Yes
Route Selector Server No

Redistribute RIP routes into BGP is disabled
Redistribute OSPF routes into BGP is disabled
Redistribute BGP routes into OSPF is disabled
Redistribute BGP routes into RIP is disabled

BGP Peer 205.14.128.1
  Configuration ID  1
  Startup State     Inactive
  AS Number         110
  Peer Weight       2000
  Next Hop Self     No
  Cfg Hold Time    180
  Retry Time        45
  Use Loopback      No
  Advertise Default Yes
  Input Route Map   rmapin
  Output Route Map  rmapout
BGP Peer 198.41.11.213
  Configuration ID  2
  Startup State     Active
  AS Number         100
  Peer Weight       1000
  Next Hop Self     No
  Cfg Hold Time    180
  Retry Time        65
  Use Loopback      No
  Advertise Default No
  Input Route Map   None
  Output Route Map  None

```

The peer **Startup State** indicates whether the router will attempt to establish a session with the peer upon power-on. If this is set to **Inactive**, the peer may be enabled with the **BGP Enable** command. The peer will once again be inactive at the next router restart, however.

Note that the first peer has **BGP Route Maps** defined, whereas the second peer does not. Since **Use IP Route Filters** has been set to **Yes**, they will be used for the second peer, but not the first peer.

Show BGP Aggregates

The **show bgp aggregates** command displays the routes which the administrator has configured to be aggregated to external peers. Aggregation will only occur when an instance of the route appears in the IP routing table.

```
bgptest>sho bgp agg

      BGP AGGREGATES:
      195.41.0.0/16
```

BGP Disable

This command discontinues a BGP session with a selected peer, or with all peers.

```
BGP disable all
      OR
BGP disable 205.14.128.1
```

Reset BGP Peer

This command resets a session with a selected BGP peer, or with all peers.

```
Reset BGP Peer all
      OR
Reset BGP Peer 205.14.128.1
```

BGP Quick Start Guide

Here is a very simple configuration to get BGP up and running. This assumes that you only have one exit point from your AS, and will therefore be using a static default route for your outgoing packets.

1. Enable BGP and specify your AS number in the BGP General Section.

```
[ BGP General ]

BGPEnabled = On
BGPAS = your AS number
```

2. Specify the IP address and AS number of your BGP peer, in this case your ISP's BGP router.

```
[ BGP Peer List ]

BGPPeer = On peer IP address peer AS number
```

3. Specify a network list for the internal networks you want advertised outside your AS.

```
[ BGP Networks ]

LocalNet = first IP address mask
LocalNet = second IP address mask
```

BGP Debug Options

For code versions with debugging available, there are five BGP debug commands: **BGPPKT**, **BGPDB**, **BGPCON**, **BGPKEEP**, and **BGPTXQ**. **BGPPKT** provides information about exchange of BGP update packets. **BGPFDDB** provides database update information. **BGPCON** provides information concerning the status of BGP sessions with peers. **BGPKEEP** provides information about when KEEPALIVE packets have

been sent or received. **BGPTXQ** provides information about the sending of update packets to peers in ESTABLISHED state.

```
sys debug flags BGPPKT
  sys debug flags BGPCON
  sys debug flags BGPFDB
  sys debug flags BGPKEEP
  sys debug flags BGPTXQ
```

BGP RFC References

```
rfc2283 -- Multiprotocol Extensions for BGP-4.
  T. Bates, R. Chandra, D. Katz, Y. Rekhter.
  February 1998. (Status: PROPOSED STANDARD)
rfc2042 -- Registering New BGP Attribute Types.
  B. Manning.
  January 1997. (Status: INFORMATIONAL)
rfc1998 -- An Application of the BGP Community Attribute in
  Multi-home Routing.
  E. Chen & T. Bates.
  August 1996. (Status: INFORMATIONAL)
rfc1997 -- BGP Communities Attribute.
  R. Chandra, P. Traina & T. Li.
  August 1996. (Status: PROPOSED STANDARD)
rfc1965 -- Autonomous System Confederations for BGP.
  P. Traina.
  June 1996. (Status: EXPERIMENTAL)
rfc1863 -- A BGP/IDRP Route Server alternative to a full mesh routing.
  D. Haskin.
  October 1995. (Status: EXPERIMENTAL)
rfc1774 -- BGP-4 Protocol Analysis.
  P. Traina, Editor.
  March 1995. (Status: INFORMATIONAL)
rfc1773 -- Experience with the BGP-4 protocol.
  P. Traina.
  March 1995. (Status: INFORMATIONAL)
rfc1771 -- A Border Gateway Protocol 4 (BGP-4).
  Y. Rekhter & T. Li.
  March 1995. (Status: DRAFT STANDARD)
rfc1745 -- BGP4/IDRP for IP---OSPF Interaction.
  K. Varadhan, S. Hares, Y. Rekhter.
  December 1994. (Status: PROPOSED STANDARD)
```

Related Information

- **Technical Support & Documentation – Cisco Systems**

[Contacts & Feedback](#) | [Help](#) | [Site Map](#)

© 2008 – 2009 Cisco Systems, Inc. All rights reserved. [Terms & Conditions](#) | [Privacy Statement](#) | [Cookie Policy](#) | [Trademarks of Cisco Systems, Inc.](#)

Updated: Jan 22, 2007

Document ID: 17612
