



BGP Frequently Asked Questions

Version: 1824

Copyright 2007-2010 ImageStream Internet Solutions, Inc., All rights Reserved.

Table of Contents

BGP Frequently Asked Questions.....	1
What is BGP?.....	1
What implementation of BGP do ImageStream routers use?.....	1
What is "multi-homing"?.....	1
What are the benefits of using BGP?.....	1
Does using BGP make it easier to change providers?.....	2
Do I have to use BGP with all of my upstream providers?.....	2
Where can I get an Autonomous System Number?.....	2
What is the Network Layer Reachability Information (NLRI)?.....	2
What is a BGP route?.....	2
How does BGP pick the best route?.....	3
How do I force Quagga to choose an OSPF route over an EBGp route?.....	3
What ImageStream routers support BGP4 and how much memory do they require?.....	3
How can I tell if BGP4 is enabled on my ImageStream router?.....	4
How can I enable BGP4 on the router?.....	4
What is the best source of information about ImageStream's BGP4?.....	5
Why must all BGP peers in an Autonomous System be peered with each other?.....	5
What is used as the BGP Router ID on ImageStream routers?.....	5
Do I need to define my static routes in my interface configuration file (wan.conf) and in the BGP configuration?.....	5
How long must a route exist before BGP will forward it?.....	5
Why do I have to inform my providers manually about the networks that I am announcing?.....	5
In the route precedence, where are Local (Directly Connected) routes in the precedence?.....	6
What third-party products are interoperable with BGP4 on an ImageStream router?.....	6

BGP Frequently Asked Questions

What is BGP?

The Border Gateway Protocol (BGP) is an exterior routing protocol. BGP version 4, described in RFC 1771, and further defined in version 5 of the BGP4 Internet Draft of January 1997, was designed for routing between Autonomous Systems. BGP Autonomous System Confederations information can be found in RFC 1965. BGP Route Reflection information can be found in RFC 1966. The BGP COMMUNITIES attribute is described in RFC1997. Although it can be used internally, BGP is primarily used to route between networks, rather within them (such as RIP, OSPF, IS-IS).

What implementation of BGP do ImageStream routers use?

ImageStream routers provide BGP via either Quagga or GateD. While both packages are available, ImageStream recommends Quagga as the BGP daemon of choice. This document currently addresses questions from a GateD specific manner.

What is "multi-homing"?

"Multi-homing" means having multiple connections to the Internet. There are two basic types of multi-homed organizations:

- ◇ An organization which has connections to provider at more than one point of presence, and has internal links between these sites;

or

- ◇ An organization which has a connection to provider and at least one other Internet access provider.

If you have only a single connection to the Internet there is not a lot to be gained from using BGP. With only one gateway, your router already knows how to send all traffic not destined to one of your networks. Everyone should be conscious of keeping the Internet's routing table as small as possible. This can be best addressed by aggregating addresses using CIDR blocks. As a single-homed connection, your provider will likely choose to aggregate your networks with the networks of other single-homed customers and thus announce a single CIDR block under their ASN in order to reduce the routing table size.

What are the benefits of using BGP?

The single word answer is "control". When you decided to multi-home it was for a reason, maybe extra capacity, redundancy, or tariffing, so in order to actually perform any of these functions you need to control how your networks are announced to your providers. If you don't exchange routes with them but rely on them announcing your routes you could find your traffic flowing down an expensive low capacity link with the other cheaper, higher capacity links mostly idle!

Does using BGP make it easier to change providers?

Each Tier 1 backbone provider has its own procedures to handle customers that wish to use BGP. These procedures may be easier than the procedures used for statically-routed customers, but generally each will require some sort of route registration, so there is no real time or effort saved. However if you are using BGP to more than one provider, you can migrate routes between the providers without their intervention (assuming any route registration has already been performed). In order to migrate the routes, you should start announcing the routes to the new provider and confirm that the new path is visible via a "looking glass". You may need to prepend your AS on the old path to make that path "longer". Having confirmed that the new provider is successfully announcing the routes you can then terminate the announcement to the old provider.

Do I have to use BGP with all of my upstream providers?

For the best results, you will find it advantageous to use BGP, but it isn't strictly necessary. In theory, your routes should originate from the same AS number, so you should announce your routes via BGP even if you don't listen to all announcements being sent to you. For example, you could use a static route to default to a provider rather than accept the routes they are sending via BGP.

Where can I get an Autonomous System Number?

You can get an Autonomous Number from The American Registry for Internet Numbers (ARIN). ARIN is a non-profit organization established for the purpose of administration and registration of Internet Protocol (IP) numbers to the geographical areas previously managed by Network Solutions, Inc. (InterNIC). The URL for ARIN's Registration Services is <http://www.arin.net/regserv.html> One of the following conditions must be satisfied before receiving an Autonomous System Number:

- ◇ Unique Routing Policy: Please explain how your routing policy is different from your provider.
- ◇ Multi-homed Site: If your organization is currently multi-homed, please explain exactly how you are connected to the Internet. Please include the gateway protocol used, the IP network addresses that make up your AS, and technical point of contact information of your upstream providers/peers. The point of contact information must include full name, and email address for verification.

The URL for the Autonomous System Number application form is <http://www.arin.net/regserv/templates/asntemplate.txt>.

For a complete list of all registered ASes on the Internet, see this link: <ftp://rs.arin.net/netinfo/asn.txt>.

What is the Network Layer Reachability Information (NLRI)?

Sets of destinations are the Network Layer Reachability Information (NLRI). A destination is a network address prefix, with a prefix length. For example, the NLRI for the private network 192.168.20.0 would be 192.168.20.0/24.

What is a BGP route?

The combined bundle of a path to a destination Autonomous System, the attributes of the path used to reach the Autonomous System, and the NLRI located at that Autonomous System is called a BGP

route.

How does BGP pick the best route?

Quagga performs the following series checks in order to determine the best route.

- ◇ The route is ignored if the next hop is unreachable.
- ◇ The route with the largest weight.
- ◇ The route with the largest Local Preference.
- ◇ The route that is originated on the local router.
- ◇ The route with the shortest AS path.
- ◇ The route with the lowest origin code.
- ◇ The route with the lowest Multi Exit Discriminator metric.
- ◇ The route that was learned from a EBGp peer vs a IBGP peer.
- ◇ The route that can be reached through the closest IGP neighbor.
- ◇ The route learned from the lowest BGP Router Id.

How do I force Quagga to choose an OSPF route over an EBGp route?

Quagga uses an administrative distance in addition to metric and local preference to choose the best path to a network. The administrative distance is the biggest factor with this choice. Each routing protocol has a default distance.

Route Source	Default Administrative Distance
Connected interface	0
Static	1
EBGP	20
OSPF	110
RIP	120
IBGP	200

To force an EBGp route to use an internal routing protocol such as OSPF use the **network ip subnet backdoor** command. This command will set the administrative distance for the subnet to 200, allowing internal routing protocols to make the routing decision.

```
router bgp 40587
 network 205.159.243.0/24 backdoor
!
```

This will force Quagga to change the distance for 205.159.243.0/24 to 200 and choose any OSPF or RIP routes over the BGP route.

What ImageStream routers support BGP4 and how much memory do they require?

All ImageStream routers with at least 1 GB memory will support BGP4 with at least two full routing views or feeds.

How can I tell if BGP4 is enabled on my ImageStream router?

Use the Configuration and Update Menu:

1. At the Login: prompt, enter root.
2. At the Password: prompt, enter your password.
3. Choose Option 1 (Configuration and Update Menu) from the Main Menu.
4. Select Option 6 (Dynamic routing configuration).
5. From the Dynamic routing configuration menu, you can see the status of both GateD and Quagga/Zebra. If either is running, the option will show "(running)". Conversely, if either is not currently running, the menu option will display "(stopped)".

How can I enable BGP4 on the router?

Enable BGP using Quagga

Use the Configuration and Update Menu:

1. At the Login: prompt, enter root.
2. At the Password: prompt, enter your password.
3. Choose Option 1 (Configuration and Update Menu) from the Main Menu.
4. Select Option 6 (Dynamic routing configuration).
5. Select Option 2 (Dynamic routing (Quagga/Zebra))
6. To enable Quagga, select option 2 (Enable Quagga service on boot)
7. Select Option 1 (Enable Quagga on boot) and Option 2 (Enable bgpd on boot)
8. Select 0 to go back to Dynamic routing (Quagga/Zebra)
9. To start Quagga, select option 4 (Start Quagga service)
10. Select Option 1 (Start quagga) and Option 2 (Start bgpd)
11. Select 0 to go back to Dynamic routing (Quagga/Zebra)

Enable BGP using GateD (not recommended)

Use the Configuration and Update Menu:

1. At the Login: prompt, enter root.
2. At the Password: prompt, enter your password.
3. Choose Option 1 (Configuration and Update Menu) from the Main Menu. Select Option 3 (Service Configuration).
4. From the Service Configuration menu, you can see the status of the Dynamic routing (gated) package. If gated is running, the option will show "(running)".
5. To start gated, choose Option 3 (Dynamic routing) and choose Option 5 (Start gated) from the Dynamic routing menu. You will also want to select Option 3 (Enable gated on boot) to ensure that BGP session will be established at boottime. Ensure that you have configured gated prior to attempting to start it. Unless it is properly configured, gated will not start.

What is the best source of information about ImageStream's BGP4?

The Quagga home page provides the most comprehensive up to date information on the recommended routing daemon for ImageStream routers.

Why must all BGP peers in an Autonomous System be peered with each other?

A fundamental rule of BGP is that all internal peers in the same Autonomous System must all fully talk to each other. That is:

- ◇ If you have peers, A, B, C in an Autonomous System, then there has to be three peers: An A/B peer, a B/C peer, and an A/C peer. Add a fourth, D, and then you need A/D, B/D, and C/D peers.
- ◇ The point being, from the BGP spec: If a router B, learns a route (path attributes + destination) from internal peer A, it assumes that all other internal peers hear it too, so it won't forward the information onto any other internal peer. It is possible to sidestep this requirement through the use of either BGP route reflectors or BGP confederations.

What is used as the BGP Router ID on ImageStream routers?

The IP address of the Ethernet0 interface is typically used. But the router will not assume this; it must be explicitly set in the BGP configuration.

Do I need to define my static routes in my interface configuration file (wan.conf) and in the BGP configuration?

Quagga

No, Quagga will not remove those routes at runtime.

GateD

Yes, you must specify any static routes in both places. If you do not specify the static routes in the main gated configuration file (gated.conf), gated will remove those routes from the routing table at runtime.

How long must a route exist before BGP will forward it?

A route must exist for 30 seconds before BGP will forward it.

Why do I have to inform my providers manually about the networks that I am announcing?

It is security problem for a BGP peer to unilaterally accept all routes from external peers. Thus providers filter out all announcements from you except for what you explicitly tell them about.

In the route precedence, where are Local (Directly Connected) routes in the precedence?

Local (directly connected) routes are the most highly preferred routes, with static routes being the next highly preferred. Route precedence can be adjusted within the configuration by changing the preferences on different route types.

Local (Direct) routes>Static routes>Dynamic routes

What third-party products are interoperable with BGP4 on an ImageStream router?

ImageStream's BGP4 implementation has been tested for interoperability with the following products, among others:

- ◇ Cisco IOS
- ◇ Nortel/Bay Networks OS
- ◇ Juniper Networks OS
- ◇ gated public and commercial versions
- ◇ zebra / quagga

Any RFC-compliant BGP4 equipment should interoperate with ImageStream routers. Please contact ImageStream Technical Support if you have BGP interoperability questions with specific equipment.