



Mailing Address:  
PO Box 8905  
Vancouver, WA 98668-8905

Physical Address:  
4400 NE 77<sup>th</sup> Avenue  
Vancouver, WA 98662

## Customer BGP Routing Policy

---

Revision 1 (05/16/2002)

### I. Introduction

Before being approved for BGP routing, customers must review the policies outlined below. These policies haven't been implemented to ensure top-rate network performance and reduce the possibility of routing problems while still allowing customers to control routing behavior to the fullest extent possible.

ELI reserves the right to add, delete, or modify any of the policies contained herein at any time, without notice. Please find the most recent revision of this document at: [http://www.eli.net/technical/bgprouting\\_policy.shtml](http://www.eli.net/technical/bgprouting_policy.shtml).

Question or comments regarding this policy should be forwarded to [bgp4@eli.net](mailto:bgp4@eli.net). Responses will be processed within two business days.

### II. General

Peers must be multi-homed to run BGP, either multi-homed within ELI's network only, or multi-homed with ELI and one or more other providers.

ELI will not static route any netblocks to customers running BGP unless those netblocks are part of ELI aggregates and are too small to be announced individually (prefix length greater than /24).

All customers are encouraged to aggregate their contiguous network announcements to the highest CIDR boundary.

Customers must have a registered public AS number, the only exception is when customers who are multi-homed to ELI only and have requested a private-AS.

ELI will use only BGP version 4.

Outbound network announcements should be filtered by applying explicit prefix and AS-path filters, not exclusive (i.e. AS-filters should list your AS and AS paths you would like to announce and deny all others instead of denying just the AS's of external peers you may have).

ELI reserves the right to require a customer to use a MD5 authentication key for BGP sessions.

ELI will not run BGP with customers having a circuit line rate of less than 256 Kbps and then will only send partial routes. If requesting full routes, connected line rate must be at a minimum of 512 Kbps.

All customers wishing to run BGP with ELI must submit a BGP routing application. ELI will process the application within 2 business days of submission. The online application can be found at: [http://www.eli.net/technical/bgprouting\\_form.shtml](http://www.eli.net/technical/bgprouting_form.shtml)

ELI will employ best known practices to establish, maintain, and troubleshoot BGP sessions with all BGP4 compliant router vendors. However, ELI makes no warranty that it can establish and maintain a BGP session with any customer provided equipment due to vendor incompatibility.

ELI is not responsible for customer equipment configurations and problems arising from customer equipment mis-configurations. It is recommended that all customers deploying BGP be skilled in configuring and troubleshooting the protocol.

ELI does not policy route or alter any of its BGP configurations, including filter policies or communities, for any individual customer. This ensures the ELI network is maintained in a robust fashion and reduces mean-time-to-repair for BGP related trouble.

All customer network announcements are given a default local-preference of 100. Refer to Section V for information on how this may be adjusted.

ELI will use EBG-P-Multihop for load-balancing multiple circuits only. Under no other circumstances will ELI peer via EBG-P-Multihop. Customers must provide their own IP address for their Loopback Interface. For more information on load-balancing via EBG-P using Cisco routers please see: <http://www.cisco.com/warp/public/459/13.html#A5.1> When applying for BGP peering please indicate the desire to use EBG-P-Multihop in the comments field.

### **III. Filtering**

ELI filters ingress on AS-path and prefix. For details on how to submit, view, and update your current AS-path and prefix filters, please see Section VIII.

ELI prefix filters for public AS customers are built such that any prefix filter entry up to /24 is allowed by default. ELI feels it is the role of the Autonomous System administrator to responsibly aggregate. All network announcements whose prefix length is greater than /24 will not be accepted by ELI nor propagated to its peers.

ELI prefix filters for private AS customers are built such that network announcements with prefix length up to /30 are accepted. Contiguous networks between /24 and /30 must be aggregated to the highest CIDR boundary and will be enforced with exact prefix filter entries.

ELI as-path filters are configured by default to allow unlimited as-path prepends as well as unlimited as-path prepends for any AS's the customer has allowed transit.

ELI will not accept default routes or illegal RFC address space to include RFC 1918 and LINK-LOCAL addresses.

ELI will accept network announcements originated only from the customer AS and AS's for which the customer is intentionally providing transit. ELI will not accept leaking of full Internet routes.


If any of the above filter policies are violated and not corrected in a timely manner after customer notification, ELI may administratively shutdown the BGP session until the issue is remedied. Customers continually leaking the above mentioned routes will be converted from BGP to static routing.

#### IV. Multi Exit Discriminators

ELI allows customers to send Multi Exit Discriminators (MEDs). By default, ELI will add its own internal MEDs to facilitate closest exit routing. If this action is not desired, ELI allows this behavior to be overridden by the customer sending a "NO-IMEDS" community (See Communities, Section V).

ELI internal MEDs are set up such that the continental United States has been divided into 6 IBGP regions: *North West*, *South West*, *North Central*, *South Central*, *North East*, and *South East* (see region map below).

- Announcements learned from a router in the same region as the origination router are given an additive metric of 1000.
- Announcements learned from a router in a region either North or South of the origination router's region are given an additive metric of 1500.
- Announcements learned from router in a region either East or West of the origination router's region are given an additive metric of 2000.
- If an announcement crosses multiple region boundaries either North-South or East-West, metrics are added for every region boundary crossed.



#### V. Communities

ELI allows customers to manipulate their network announcements by sending BGP communities. Customers can send communities to adjust ELI local preference, modify export behavior, modify ELI internal MEDs behavior, and perform selective AS-path prepending. Customers can use various combinations of these communities to control their inbound traffic flows.

ELI strips any unrecognized BGP communities received from the customer. ELI does not propagate any BGP communities to non-customers, originated by ELI or otherwise.

- Local Preference Manipulation

ELI's Backbone Local Preference Policy uses the following values:

65, 75, 80, 85	ELI Peers
100	ELI BGP Customers (default)
100	ELI Internal and static routed customers

Customers may send the following communities to modify BGP local preference given to networks announced to ELI:

Community	ELI Local Preference	Effect on ELI Route Preference
5650:60	60	<i>Last Resort</i> -- Only use this route when alternate providers no longer have this route.
5650:70	70	<i>Backup #2</i> – Allows the customer to specify that the circuit on which this route is learned is to be a backup to another circuit with ELI (since default is set to 100). Also causes the same route to be preferred if learned from and ELI Bilateral peer over the route marked with this community.
5650:90	90	<i>Backup #1</i> – Allows the customer to specify that the circuit on which this route is learned is to be a backup to another circuit with ELI
5650:110	110	<i>Preferred</i> – This route is preferred over all others

- Export Manipulation

Customers may send the following communities to modify ELI default export behavior, which is to export, all announced routes (assuming they have passed ELI filters) to all BGP peers including customers:

Community	Export Behavior
5650:640	Do not export to any BGP peers or customers (similar to well known community 'no-export')
5650:666	Do not export to any BGP peers, but do export to ELI BGP customers

- Internal MED Manipulation

Customers may send the "NO-IMEDS" community to prevent ELI from applying additive regional internal MEDs that aid in closest exit routing decisions. Customer may want to do this if they have multiple connections to ELI and have an established MED policy, which they do not want ELI internal MEDs to influence.

Community	MED Action
5650:630	Do not apply internal metrics

- Selective Prepend Manipulation

Selective prepends allow the customer to have ELI perform AS-path prepending to the Internet on the customer's behalf. This allows increased control over inbound traffic flows more than simple AS-path prepending performed by the customer. With traditional AS-path prepending, the padded AS-path is seen identically by all ELI peers. With selective as-path prepending, ELI allows customers to have ELI prepend only to certain interesting peers. ELI current routing policy permits this feature with approximately 20 of its peers (large and/or high traffic networks). Selective AS-path prepending also supplies a mechanism to allow customers to prevent export to these peers individually.

Because publicly providing a community/peer cross-reference table would violate peering contracts ELI has with certain of its peers, customers are asked to privately request the Selective Prepend Supplement document via e-mail to [bgp4@eli.net](mailto:bgp4@eli.net). Customers will be expected to provide valid ELI circuit ID as well as peering IP address and AS number if possible. Once the request is received, the Selective Prepend Supplement will promptly be forwarded to the requestor.

ELI sends communities to customers by default. Customers may match on these communities and apply various internal policies to manipulate outbound traffic flows.

- Region Marking

ELI uses BGP communities to mark what region a particular route is learned. The region numbers coincide with ELI's IBGP regions. Customers that have multiple diverse connections to ELI may use these region communities to match and apply an internal policy to ensure that the optimal path is used. ELI region origination communities are as follows:

Community	Region of Origin
5650:501	Announced from a router in Region 1 (North West)
5650:502	Announced from a router in Region 2 (South West)
5650:503	Announced from a router in Region 3 (North Central)
5650:504	Announced from a router in Region 4 (South Central)
5650:505	Announced from a router in Region 5 (North East)
5650:506	Announced from a router in Region 6 (South East)

- Prefix Origin Marking

Additionally ELI uses BGP communities to mark which type of network a particular announcement was learned. These communities are additive to the Region marking communities, in other words any network announced by ELI should have at minimum both a Region marking community and a Network marking community. Network origination marking communities are as follows:

Community	Network of Origin
5650:900	Route originated from ELI core (ELI aggregate CIDR blocks)
5650:920	Static routed customers of ELI using non-ELI IP space
5650:930	Route learned from an ELI bilateral peer
5650:950	Route learned from an ELI full peer
5650:980	Connected route redistributed from an public exchange point from which ELI peers
5650:990	Route learned from an ELI BGP customer

### Dampening

ELI maintains a reasonable route dampening policy. ELI may apply dampening to customer's network announcements in cases of heavy update/withdrawal activity. This policy is standard across all gateway routers and is non-negotiable. In case of dampening, customers may contact ELI to have their dampened routes cleared by issuing a trouble-ticket with the ELI's NOC. This does not include external paths outside the ELI network.

Values used in ELI's Cisco Gateway Routers for route dampening:

Penalty for route flap: 1000\*  
 Suppress value: 2050  
 Max Suppress Duration: 80 minutes  
 Half-Life Decay: 20 minutes  
 Reuse value: 450

\* Cisco defines a route flap as an up/down transition combined where the penalty is counted on the down transition.

## VI. Announcement Types

ELI offers several types of route announcements for the customer to choose. The characteristics for each announcement type are described below. Customers are asked to choose which type of announcement they desire when filling out the BGP routing application. The customer may request a different announcement type at any time after the session is established by following the procedures listed in Section VIII.

- Default Route Only

ELI sends only a default route to the customer (0.0.0.0/0). This is different than a static default that the customer may set themselves in that this default route learned via BGP inherits all of the dynamic characteristics that BGP provides. Choosing this option is useful when low memory use is a priority on the customer router. Note that choosing only a default route may lead suboptimal routing and packet delivery.

- Partial Routes

ELI sends to the customer ELI (AS5650) routes and ELI EBGP customer routes. At present, they number approximately 750. This option is useful when memory constraints on the customer router are important however the customer still wants the ability to make semi-informed routing decisions. Obtaining partial routes without a default is usually not useful because it may limit reachability to the entire Internet. For this reason, obtaining only partial routes is not very common and is normally recommended the customer receive partial plus default routes.

- Partial plus Default Routes

ELI sends to the customer ELI (AS5650) routes and ELI EBGP customer routes as well as a default route. This option is useful when memory constraints on the customer router are important however the customer still wants the ability to make semi-informed routing decisions based on ELI announced networks. Obtaining the default route in addition to the partial route ensures reachability to other networks not specifically advertised with partial routes.

- Full Routes

ELI sends the customer full Internet routes which make up explicit network announcements to every destination on the Internet. At present, full Internet routes number approximately 105,000. Obtaining full Internet routes is useful for customers where memory utilization is not an issue on their router as full Internet routes may require 128 MB of memory or more. Obtaining full Internet routes allows for fully-informed routing decisions and optimal packet delivery.

- Full plus Default Routes

ELI sends the customer full Internet routes as well as a default route. Obtaining full Internet routes is useful for customers where memory utilization is not an issue on their router. In addition, they receive a default route, which allows routing of packets where an explicit route does not exist. Some customers prefer this because it may offer a safety net of sorts. This may not always work however because it may simply result in moving the lack of reachability from the customer's network to ELI's network. For this reason, choosing full routes plus a default is not very common.

## VII. Support and Maintenance

Requests for the following should be made via e-mail to support@eli.net:

- Changes to inbound filters

Please indicate the specifics of the AS-path or prefix filter change.

- Snapshots of existing filters

Indicate that you would simply like a snapshot of presently configured filters.

- Changes to announcement type

Indicate the new announcement type you desire.

In your requests, please include the following information:

Company Name  
ELI Circuit Identification Number  
Technical Contact Name and Phone Number  
Preferred E-mail address  
AS Number  
Interface connected IP address or Peering IP address (if different)

Response time for any of the requests above will be processed in two business days. Any questions regarding previously submitted changes should be directed to support@eli.net; other questions or concerns regarding BGP should be directed to bgp4@eli.net.