

XO™ BGP POLICY

This document describes the Border Gateway Protocol (BGP) policy implemented by XO Communications to its BGP customers; it gives a description of how we handle BGP connections, what announcements we are willing to accept, how we treat those received announcements in relation to those heard from Internet peers, and tools provided to our customers to manipulate how we treat their announcements. Also, it gives some example configurations for customer routers.

This document assumes the reader is familiar with the BGP protocol; therefore, it should not be considered as a guide that teaches the BGP protocol.

What is BGP and Do I Need It?

BGP is an exterior gateway protocol (EGP) used to exchange routing information between different autonomous systems (AS). This network layer reachability information (NLRI) is maintained in a database on the router, and exchanged with other routers; the NLRI is used to create a graph of AS connectivity, which allows BGP to:

- remove routing loops,
 - enforce policy decisions, and
 - forward traffic to the shortest path at the AS level.
- It is the protocol used to create the Internet backbone.

The decision to use BGP is best discussed with an XO Sales Engineer because it depends on the customer network design and needs of the customer. The most basic criteria establishing the need for BGP is a network with two (or more) upstream provider connections where the customer wishes to announce his network address space to the Internet through these providers for the purpose of provider redundancy and / or shortest path routing.

Single provider homed networks do not usually need BGP, and can achieve circuit redundancy or load sharing without running BGP.

XO Announcement Options

Customers buying BGP have options as to what routes we will send them; these options consist of:

- Full Internet Routes – All routes XO receives from its peers, customers and originating from our own network.
- XO Only Routes – Routes only originated by XO or received from our BGP customers.
- Default Only – A single default route.

A customer who wishes to receive a default route in addition to either full XO routes or full Internet routes needs to request it through the order process.

For these outbound announcements to the customer, XO sends Multi-Exit Discriminator (MED) attribute and unaltered community attributes. These MEDs help give the customer a better routing decision, which reflects those destinations on the XO network that may be closer to the POP where the customer is connected. XO sets MED to the following values:

- “1” for 2828 aggregates
- “2” for routes tagged with the local market community, and
- “3” for everything else.

Customer Announcement Restrictions

Customers are limited in what XO will accept as valid routes. Any non-compliant routes are discarded. Customers should strictly filter their announced routes to only those authorized by XO, rather than depend on XO customer controls.

The following limitations are applied to XO BGP customer announcements:

- Prefix Limits – Every prefix should either be assigned by XO provisioning / pre-sales / IP admin group or a registry (ARIN, RIPE, APNIC, etc.) directly. If the block is out of another ISP's aggregate, our Best Current Practice requires a letter or email from a representative of that ISP, authorizing XO to accept and readvertise the prefix(es) from the customer's AS / connection. XO does not accept prefixes shorter than a /8 or longer than a /24.
- AS-Path Limits – XO will only accept routes from the customer AS and any downstream AS using that customer network for transit services. XO rejects any route not explicitly allowed. Direct request for addition of a customer downstream AS and / or prefix needs to be made via an email to CustomerRoutingUpdates@xo.com.
- Community Attribute – Customer announcement communities are limited to the two “well known” values of “no-export” and “no_advertise”. All other communities – with the exception of customer signaled, private communities (described further in this document) – are deleted from the update. Other than the deletion of the unauthorized community, the routes are treated like any normal advertisement.



While XO filters on customer-announced routes, it is the responsibility of the customer to ensure he is only sending approved routes. XO reserves the right to protect its network from both authorized and unauthorized announcements that may result in any negative impact to the XO network and the Internet at large. These actions may include – but are not limited to – maximum-prefix peer shut down mechanisms, route filtering or peer administrative shut down until the problem is corrected.

Customer Routing: Customer Choice

Customers can control their routing through a variety of methods. Some of these are special community-based controls implemented by XO, while others are standard BGP controls.

XO Community Controls

Local Preference Community

The general XO routing policy is to prefer our customer's BGP and static routes over "external" (non-customer) BGP routes.

For our BGP customers, XO has implemented a powerful, flexible system whereby we allow our customers to control – via predefined communities – the XO treatment of their announcements. These communities control XO local-preference (local-pref) of customer-announced routes within our network.

The most powerful attribute is local-preference.

In BGP, the higher the local-pref, the more preferred the route on the network. This affects XO backbone route selection only, and not the route selection made by XO peers and customers, because local-preference is non-transitive. This means it does not get passed beyond the XO network. The customer has a range of six local-preferences to which they can set their routes.

The following table details the local-prefs, including the actual value, the community accepted from the customer to set the local-pref, and a description of what that local-pref will do in the network.

Customer Local-pref Options		
Local-pref	Community	Description
70	2828:1507	Lower (less preferred) than all other routes on network, incl all public & private peer routes
80	2828:1508	Same as public and private peers.
90	2828:1509	
100	none	Default for all customer routes, BGP and static
100	2828:1510	Explicitly set customer BGP announcements to 100
110	2828:1511	Higher (more preferred) than all default customer BGP and static routes
120	2828:1512	Highest (most preferred) local-pref that a customer can specify

AS-Path Controls

The next attribute that customers can use to control their routing is via AS path length.

The most common use of AS-path prepending is in a direct peering route-map on the customer end where they can prepend their own AS to their routes as they announce the routes to AS2828.

XO also accepts communities that prepend AS 2828 either one time or three times to the route as it leaves the 2828 network.



Currently, there are three general groups on which the customer can set 2828 prepends:

Communities that Change Customer Announcements to Certain Peers at AS2828 Border				
Provider Name or Group	"Don't Advertise to" Community	"Prepend Once" Community	"Prepend Twice" Community	"Prepend Thrice" Community
All Peers	2828:1000	2828:1100	2828:1200	2828:1300
Sprint	2828:1003	2828:1103	2828:1203	2828:1303
Cable & Wireless	2828:1004	2828:1104	2828:1204	2828:1304
UUNet	2828:1006	2828:1106	2828:1206	2828:1306
Level3	2828:1007	2828:1107	2828:1207	2828:1307
ATT	2828:1008	2828:1108	2828:1208	2828:1308

Origin Attribute

Another standard option for customers to control XO treatment of their advertisements involves the use of the origin attribute.

XO does not alter origin code on routes customers send us. Routes set with the IGP origin code are more preferred than routes tagged with the origin EGP code, which are more preferred than origin-Incomplete routes.

DoS Discard Community

XO provides a mechanism allowing customers to trigger the XO backbone to discard all traffic entering the XO network from external sources destined for any routes specifically announced by a BGP customer with a special community.

XO customers may announce this route to XO to mitigate the over utilization of their uplink or general harm caused by such an attack until they open a ticket (if required) and XO technicians help stop the attack.

Other Accepted Communities	
Community	Description
2828:1650	Used to blackhole traffic
no-export/no-advertise	Don't announce outside of XO

MED Controls

Another standard option for customers to control XO treatment of their advertisements involves the use of the MED attribute.

The lower the MED attribute set on a route, the more preferred the route. XO currently accepts MED from its customers.

For more details on (Cisco) BGP path selection and attributes, see [Cisco's Web site](#).

Sample Customer Configurations

The following configurations are samples, provided for educational purposes only.

Customers managing their own CPE, running BGP to XO Communications, should understand the abilities of the protocol, and manage their router based on their own needs.

Inbound controls could be stricter and many other variations could place different configuration requirements.

The examples below show simple and complex configurations, but they could become more complex depending on the customer's needs.

For further guidance, see the [Cisco Web site](#).



1. General BGP Configuration (to be used in all configuration examples)

```
router BGP <AS number>
no synchronization
BGP log-neighbor-changes
BGP deterministic-med
no auto-summary
network n.n.n.n m.m.m.m summary-only
```

("n" is network address and "m" is network mask of aggregate network you want to advertise)

This anchors your aggregate network in case the more specific goes away rather than follow either a default to the provider or follow a larger aggregate announced by the provider who gave you the space. Without it, you can have unreachable traffic for these specific routes ping-pong on the wan. It also needs to be an exact match for your network statement.

```
ip route n.n.n.n m.m.m.m null 0
```

Use default ONLY in the case where you accept limited routes from one provider, and want that provider to take some traffic and be a back up for provider offering full routes.

```
ip route 0.0.0.0 0.0.0.0 <wan interface>
```

2. Simple Multi-Home – Accept All Internet Routes

```
router BGP <Local AS Number>
neighbor x.x.x.x remote-as 2828
neighbor x.x.x.x description XO Communications:AS2828:BGP@xo.com:800-575-6398
neighbor x.x.x.x send-community
neighbor x.x.x.x filter-list 1 out
neighbor x.x.x.x soft-reconfiguration inbound
neighbor y.y.y.y remote-as <Provider #2 AS#>
neighbor y.y.y.y description Provider#2
neighbor y.y.y.y send-community
neighbor y.y.y.y filter-list 1 out
neighbor y.y.y.y soft-reconfiguration inbound
```

Only send routes that originate in your own AS.

```
ip as-path access-list 1 permit ^$
```

3. Simple Multi-Home – Accept Default Only

```
router BGP <Local AS Number>
neighbor x.x.x.x remote-as 2828
neighbor x.x.x.x description XO Communications:AS2828:BGP@xo.com:800-575-6398
neighbor x.x.x.x send-community
neighbor x.x.x.x filter-list 1 out
neighbor x.x.x.x soft-reconfiguration inbound
neighbor x.x.x.x prefix-list DEFAULT-ONLY in
neighbor y.y.y.y remote-as <Provider #2 AS#>
neighbor y.y.y.y description Provider#2
```



```
neighbor y.y.y.y send-community
neighbor y.y.y.y filter-list 1 out
neighbor y.y.y.y soft-reconfiguration inbound
```

```
route-map LOCAL-ONLY-OUT permit 10
match as-path 10
```

Only send routes that originate in your own AS.

```
ip as-path access-list 1 permit ^$
```

```
ip prefix-list DEFAULT-ONLY seq 100 permit 0.0.0.0/0
```

4.Simple Multi-Home – Advanced Outbound XO Community Options

```
router BGP <Local AS Number>
neighbor x.x.x.x remote-as 2828
neighbor x.x.x.x description XO Communications
neighbor x.x.x.x send-community
neighbor x.x.x.x remove-private-AS
neighbor x.x.x.x soft-reconfiguration inbound
neighbor x.x.x.x prefix-list DEFAULT-ONLY in
neighbor x.x.x.x route-map XO-LOCAL-ONLY-OUT out
neighbor y.y.y.y remote-as <Provider #2 AS#>
neighbor y.y.y.y description Provider#2
neighbor y.y.y.y send-community
neighbor y.y.y.y remove-private-AS
neighbor y.y.y.y soft-reconfiguration inbound
neighbor y.y.y.y route-map <Provider#2>-LOCAL-ONLY-OUT out
```

More complicated or specific stanzas could be added, but this would cover the basic tools provided for customers to manipulate their XO announcements.

```
route-map XO-LOCAL-ONLY-OUT permit 10
match as-path 1
match ip address prefix-list XO-PREF-70
set community 2828: 1507
```

```
route-map XO-LOCAL-ONLY-OUT permit 20
match as-path 1
match ip address prefix-list XO-PREF-80
set community 2828: 1508
```

```
route-map XO-LOCAL-ONLY-OUT permit 30
match as-path 1
match ip address prefix-list XO-PREF-120
set community 2828: 1512
```

```
route-map XO-LOCAL-ONLY-OUT permit 40
match as-path 1
match ip address prefix-list XO-PUB-PRIV-NOADV
set community 2828: 1000 2828:1001
```

```
route-map XO-LOCAL-ONLY-OUT permit 50
match as-path 1
```



```
match ip address prefix-list XO-PREPEND-ONCE
set community 2828: 1100 2828:1101
```

```
route-map XO-LOCAL-ONLY-OUT permit 60
match as-path 1
match ip address prefix-list XO-PREPEND-THRICE
set community 2828: 1200 2828:1201
```

```
route-map XO-LOCAL-ONLY-OUT permit 80
match as-path 1
match ip address prefix-list XO-NO-EXPORT
set community no-export
```

```
route-map XO-LOCAL-ONLY-OUT permit 90
match as-path 1
match ip address prefix-list XO-STD-OUT
```

```
route-map XO-LOCAL-ONLY-OUT deny 200
```

```
Only send routes that originate in your own AS.
ip as-path access-list 1 permit ^$
```

```
ip prefix-list DEFAULT-ONLY seq 100 permit 0.0.0.0/0
ip prefix-list XO-PREF-70 seq 10 permit x.x.x.x/x
ip prefix-list XO-PREF-80 seq 10 permit x.x.x.x/x
ip prefix-list XO-PREF-120 seq 10 permit x.x.x.x/x
ip prefix-list XO-PUB-PRIV-NOADV seq 10 permit x.x.x.x/x
ip prefix-list XO-PREPEND-ONCE seq 10 permit x.x.x.x/x
ip prefix-list XO-PREPEND-THRICE seq 10 permit x.x.x.x/x
ip prefix-list XO-NO-EXPORT seq 10 permit x.x.x.x/x
ip prefix-list XO-STD-OUT seq 10 permit x.x.x.x/x
```

