SimplifiedIP

BGP Routing Policy Information

Advanced Network Services Manual

AS 27552

JANUARY 23, 2008

VERSION 1.3.1

TOWARDEX TECHNOLOGIES, INC.
INTERNET SERVICES
NTC NETWORK TECHNOLOGIES
P.O. BOX 2538
ACTON, MA 01720-6538
V(617) 451-4051 F(432) 225-3784
www.towardex.com
www.twdx.net

TowardEX Technologies, Inc.
P.O. Box 2538
Acton, MA 01720-6538

GVN-U254517

About This Manual…

This Advanced Network Services Manual provides useful guidelines and information to customers engaging in BGP routing with TowardEX's network infrastructure. The Manual is particularly useful for advanced customers who wish to gain granular control in routing of its Internet addresses through TowardEX's network backbone.

WHO NEEDS IT?

TowardEX provisioning center always provides Internet customers engaging in BGP operations with this manual so that advanced BGP routing policies and need-to-know details for customers are readily available.

*It is **not** required for customers to read and understand contents of this manual, but information is provided for **experienced** users who would like to perform advanced traffic engineering on their network through TowardEX's infrastructure.*

For customers purchasing **Redundant IP Transit** product, this manual may still be provided to you for reference, however it is mostly not going to be useful for you unless you are multihomed with another Internet service provider with your own sovereign ***Autonomous System (AS)*** number and minimum of a /24 contiguous IP address space. Customers purchasing Redundant IP Transit product without a public AS number are assisted by TowardEX's own technical resource to implement BGP routing between their network and TowardEX backbone, thus this manual is not required for those customers.

Customers with advanced and experienced technical resource in their MIS/IT department that is familiar with BGP routing and Internet backbone operations may greatly benefit from information provided throughout this manual.

WARNING

This manual assumes that customer has extensive BGP routing experience and knowledge. It is ***NOT*** recommended for customers to implement any changes on their network based out of this manual without fully understanding the operational impact. Customers are encouraged to always contact TowardEX's Network Support Center at ip-admin@twdx.net should they have ANY questions prior to implementing any changes in BGP routing configuration.

TowardEX is in no way responsible for any damages, service interruption or operational impact caused by customer's own misconfiguration or implementation of routing policies. No SLA credits will be honored for Downtime caused by customer's configuration work on its own equipment. If you have questions or are unsure, please do not hesitate to contact our support center to verify ***prior*** to making any changes. Integration Services are also available for customers that require such resource from TowardEX to implement its desired custom changes.

Customer BGP Routing Policy Overview

TowardEX is one of few Internet service providers in the region that provide most extensive documented traffic engineering and control options to our BGP customers. Below is a brief summary of key differences between competing providers and TowardEX in providing routing control options using BGP communities to our customers.

1. BGP Communities-based Route Propagation Control & Identification using RFC1998

TowardEX's BGP routing policy makes extensive use of community tags that exceed all of RFC1998 specifications in providing ways for multihomed customers to influence routing of their prefixes.

2. Granular Route Export Controls

Most providers' BGP community options allow you to influence the propagation of your BGP prefix as it is being announced through *their* network. TowardEX raises the bar by giving our customers the option of influencing propagation of BGP prefixes as they pass through our individual peering partners and upstream transit providers. This gives extensive routing options for experienced customers where normally they would have to ask their ISP to make *manual changes* on their backbone network. Worse, many ISPs will often reject or delay request for manual changes on core backbone network due to strict Change Control Policies in place.

3. Proxy Communities

One of the big hurdles in buying BGP transit from small ISPs is that it is usually difficult to leverage or make use of community sets provided by large Tier-1 carriers located upstream from your ISP. Large carriers often have extensive peering coverage both nationally and internationally, that their BGP communities are far more effective than that of a typical small ISP. TowardEX provides special BGP community tags for its customers to use to overcome these hurdles. Known as *Proxy Communities*, customers can leverage this feature to instruct TowardEX's network to send specific traffic engineering community tags to our upstream transit providers.

4. Remote-Triggered Blackhole Routing

TowardEX supports remote-triggered blackhole routing technique using BGP communities. This can be useful for BGP customers to temporarily discard traffic destined to a particular host to save bandwidth in the event of a DoS attack.

5. Controlling Hot-Potato vs. Cold-Potato Routing Behavior for Inbound Traffic

TowardEX uses cold-potato routing for all incoming traffic with its upstream transit providers in order to leverage their SLA and quality network for packet delivery. Some content network customers have specific traffic engineering requirements that are strictly based on latency where cold-potato routing may not be useful for them. In such cases, customers can use a special community tag to disable cold potato routing of its inbound traffic, received from TowardEX's upstream transit providers. When disabled, inbound traffic to customer prefix will behave as if it is being received from a "*Settlement-Free Interconnection"* (peering) partner, by using traditional hot-potato routing.

TowardEX Internet Services
Customer BGP Routing Policy (Last updated 1/23/2008)

This document contains information on TowardEX's BGP routing policies, including BGP communities sent and received by our network. As part of the internet community, it is important for networks to publicly document these policies for customers, peers and prospective customers.

Receivable: Per-Peer Export Control Communities

Export control communities are those sent by our customers to control distribution of their routes throughout the internet, beyond the TowardEX network.

The export control format is **NNNNN:4AAAR**, where "NNNNN" is the **Calling Code** -- the calling code is used to define which group of peers or which specific adjacent AS you wish to apply your policy to. The first digit of the second component is always "4" then followed by variables of the "AAA", for **Area Code**, and "R" for **Request for Action Code**, which are defined in the following tables.

The list of **Area Codes** is available below.

Calling Code [NNNNN:4AAAR]	
0	ALL
27552	TowardEX Internal Backbone and its connected BGP Customers
174	Cogent Transit
6461	AboveNet Transit
65000	Boston MXP Public Peers
XXX	To apply your policy against a particular peer of TowardEX, place it's AS number.

Request for Action Code [NNNNN:4AAAR]	
0	Do Not Announce
1	Prepend 27552
2	Prepend 27552 27552
3	Prepend 27552 27552 27552
5	Override less-specific routing policy: Ignores inclusive action requests and just announces the route normally as is. See **TIP** below for usage example.
Proxy Communities	
4	*If supported by our upstream, forward a request to treat the route as a peering partner learned route.*
6	*If supported by our upstream, forward a request to not announce to anybody.*
7	*If supported by our upstream, forward a request to not announce to transits.*
8	*If supported by our upstream, forward a request to not announce to peers.*
9	*If supported by our upstream, forward a request to not announce to customers.*

TIP: You can use the Request for Action Code (RAC) of "5" – "Override less-specific routing policy" to refine your routing policy to a more granular level. For example, if you want to export a route only to Cogent and TowardEX customers but nobody else, you can set three communities at once in a route-map like this:

174:40005 27552:40005 0:40000

The RAC of "5" on Cogent (174) and TowardEX:Customers (27552) prevents any Request for Action codes from taking effect on them. Notice the 0:40000 community tag at the end, it means Do Not Announce (RAC 0) to any neighboring AS (Calling Code 0) at all locations (Area Code 000).

What happens with the above three communities set is that there is a global, inclusive request for action code, called "0:40000" which means "do not announce to anybody." However, because there is "Override less-specific routing policy" bit set toward Cogent (174) and TowardEX:Customers (27552), both Cogent and TowardEX Customers peer groups will still see the route, while everyone else will not. This allows you to fine-tune your routing policy to a more granular level, allowing you to set exemptions toward certain peers.

NOTE: Do not set community attributes that are contradicting one another. Setting community tags that are in conflict of what you are trying to accomplish will result in undesired results. For example, let's take the following example. A customer configures a route-map that sets two communities listed below at once:

27552:40005 27552:47010

The first community tag of 27552:40005 specifies "ignore less specific routing-policy requests in all locations". The latter tag of 27552:47010 specifies "Do Not Advertise prefix to TowardEX customers in New England (701) area."

Clearly both communities are contradicting each other. Former community (27552:40005) negates possibility of any communities being able to set "Do Not Announce" condition on the prefix, including the latter community tag of 27552:47010. So in this case, 27552:47010 community tag has no effect, an undesired result.

Receivable: Routing Policy Communities	
Routing policy communities are those sent by our customers to influence TowardEX's routing decision; such as how the network will treat customer's route in terms of preference amongst its peers, other customers and transits.	
Routing Policy Communities	
27552:50	Set local preference to 50 (last resort route).
27552:100	Set local preference to 100 (transit route).
27552:150	Set local preference to 150 (depreferenced peering route).
27552:200	Set local preference to 200 (normal peering route).
27552:250	Set local preference to 250 (backup customer route).
27552:300	Set local preference to 300 (normal customer route, this is the default).

Receivable: Special-use Communities	
These communities are used in exceptional, debugging or emergency cases and are not normally used. **Incorrect use of these communities may result in partial or total destruction of connectivity to the affected routes. Please use them with careful consideration. Outages or service degradation caused by customer's configuration changes may not give claims for SLA credits.** **IMPORTANT!** **Open Caveat TD1156** Using Special-use Communities will override any Per-Peer Export Control Communities that are set on a route. If you utilize any of the listed Special-use Communities, any applicable per-peer export control communities that may be present on the route will be ignored.	
Exceptional Special-use Communities	
27552:911	Discard/drop all traffic destined to marked route. Will also request all peers and transits to do the same, if they support such community. May be useful during Denial of Service attacks.
27552:950	Bypass label-switched traffic-engineering; disables MPLS hop-hiding activity. (**WARNING**: Experimental)
27552:901	Bypass cold potato routing on inbound traffic; Set exported Multi Exit Discriminators (MULTI_EXIT_DISC) "MEDs" to 0 to all BGP peers (customers, peering partners, transit providers). Enforces hot potato on inbound traffic, but it will cause asymmetric routing of return traffic.
27552:900	Reserved for internal NOC use only. Do not use this community tag unless specifically instructed by TowardEX engineering. Incorrectly using this community may drop traffic to affected routes causing a service outage.

Announced Communities

Every route entering the TowardEX network is tagged with a community to assist in classifying and identifying its type and origin. TowardEX uses these communities to enforce its routing policies and for informational purposes. These communities are also advertised to customer BGP sessions to help our customers in achieving their own traffic engineering practice.

The format is **27552:5AAAT**, where "AAA" is again, the **Area Code**, significant to where the route was learned from. And "T" stands for **Type Code** which is defined in the following table.

Type Code [NNNNN:5AAAT]	
1	Transit route
2	Peer route
4	Customer route
5	Internally originated route

Community Area Codes

Area Code is used to associate the general geographic location to a route. For Export Control communities, Area Code can be used to specify which geographical location you wish to have your policy propagated to.

When applying policies toward TowardEX's transit upstreams, you may not only specify which upstream AS you want to influence, but also the geographic location of a port by using Area Code. You can also do the exact same thing to TowardEX customer ASes. Please do note however, use of the Area Code toward peers is NOT supported. If you wish to influence a peer, you should always set **000** as your Area Code, which means "all possible locations." The reason why we do not support Area Code for peers is because most **Settlement-Free Interconnection Agreements** require consistent route propagation at all locations.

For Announced Communities, the Area Code is used to classify where a route has been learned from.

Area Codes are currently not necessary on TowardEX's network at this point in time. However, these codes are designed to provide scalability when we expand our network footprint to other regions in the future. A commonly used Area Code is "000" for all locations.

Area Code [NNNNN:(4|5)AAAT]	
000	All Locations
7..	North America
701	New England POPs

Local Preferences

Localpref	Description
50	Depreferenced transit or last resort route
100	Normal transit route
150	Depreferenced peering route
200	Normal peering route
250	Backup customer route
300	Normal customer route
400	Normal internal route

Appendix A: Understanding TowardEX's Outbound Route Announcement Policy

When customer announces a route via BGP, the TowardEX Network performs several actions and sanity checks before handing off that route to our upstream transit providers or other neighboring Autonomous Systems (AS) such as peering partners and other BGP customers.

These actions are called "Outbound Route Announcement Policy" and BGP customers can take advantage of it by using BGP communities. Using BGP communities, customers can request our network to perform certain actions on customer's announced BGP route, including modification and deletion before handing off the route to external AS.

Outbound Route Announcement BGP Community Types

TowardEX employs several groups of different BGP communities that are associated in performing various tasks. The following community types are available:

Per-Peer Export Control Communities:
These communities are used to set a specific instruction on how to announce a customer route to neighboring AS. These communities can be applied on either per-peer basis or globally toward all peers.

Special-use Communities:
These communities can be used to perform certain tasks that are unusual or otherwise exceptional (for example, blackhole routing). These communities are effective globally. When set, the community will apply toward every peer/neighboring AS connected to our network.

Order of Operations for Outbound Route Announcement

The order in which different community types and other various actions are taken on an outbound route announcement is described below. It may be helpful to gain understanding of this order of operations to effectively utilize different BGP community types.

Step 1. Trust Relationship Check
Checks and intercepts any routes that are not supposed to be leaking out to Internet, including deaggregates longer than /24 (/48 in IPv6), private Internet numbers (i.e. RFC1918) and unauthorized route announcements that do not have an identifier community tag.

Step 2. Override Less-specific or Global Routing Policy
Routes that are set with "Override less-specific routing policy" community tag (i.e. XXX:40005) are matched and announced straight out normally, as if there are no community tags set on them. Policy statement/route-map operation is then terminated and no other communities are evaluated.

Step 3. Process routes with Do Not Announce (no-export) Community Set
Routes that are set with Do Not Announce communities (i.e. 0:40000) are matched and intercepted from being announced to external networks. Policy statement/route-map operation is then terminated and no other communities are evaluated.

Step 4. Process routes with Special-use Communities
Routes that are set with special-use communities (i.e. 27552:911) are processed and announced accordingly to actions defined in the supplied community. Policy statement/route-map operation is then terminated and no other communities are evaluated.

Step 5. Process routes with communities specifying AS Prepend Requests
Routes that are set with communities specifying Request for Action Codes related to AS prepending operation (i.e. 0:40003 to prepend 27552 three times to everyone) are matched and processed accordingly. In the event that there are multiple communities specifying prepend requests, only one of them will be matched and acted upon. Policy statement/route-map operation is then terminated and no other communities are evaluated.

Step 6. Process routes with communities specifying Proxy Community Requests
Routes that are set with communities specifying Request for Action Code related to Proxy Community operations (i.e. XXX:40004, XXX:40006, etc) are matched and processed accordingly. In the event that there are multiple communities specifying proxy requests, only one of them will be matched and acted upon. Policy statement/route-map operation is then terminated and no other communities are evaluated.

Appendix B: BGP – General Policies for Customers

1. TowardEX propagates IPv4 prefixes shorter than or equal to /24 to its peers and upstream transits. Likewise for IPv6, prefixes shorter than or equal to /48 are propagated.

2. Customers may advertise up to maximum prefix length (/32 for IPv4, /128 for IPv6) – however, only prefixes matching under the above Rule #1 will be distributed to other Internet networks. Customers may wish to advertise these longer prefixes for traffic engineering or blackhole routing purposes.

3. Customers must beware of maximum-prefix limit on their BGP session. By default, the max-prefix limit for IPv4 is 300, and for IPv6 is 50. Exceeding these limits will result in automatic shutdown of the BGP session. In such case, customer must contact the NOC to revive its affected BGP sessions. TowardEX is not responsible for any material outages or damages incurred due to customer's tripping of the max-prefix limit.

    a. Customers who legitimately announce more routes than the default max-prefix limit will be placed on a higher limit to satisfy their needs.

4. TowardEX does not accept AS_PATHs containing private numbers (64512-65535). For customers engaging in BGP without using a public AS number, TowardEX will assign a private AS number. However, all announcements to Internet made by the customer using private AS will be incorporated under AS27552 for Origin AS.

5. TowardEX requires its customers to register proper entries in the Internet Routing Registry (IRR) – most importantly, the "route" object for IPv4, and "route6" object for IPv6. Failure to properly register routes announced by customer may cause TowardEX to forcibly register them by proxy on customer's behalf.

    a. This is because all customer BGP sessions are automatically managed by the central management program, without requiring manual human intervention. This autonomous management of customer BGP sessions is done by looking up the IRR database. When a route is manually added to the customer's list of permitted prefixes (i.e. prefix-list, route-filter, etc) by hand, at customer's urgent request, such route will be automatically removed by the management software within 24 hours, unless it is registered in an IRR database.

    b. A recommended easy-to-use online based tool for registering IRR objects is available at http://www.irrweb.com.

    c. All internet route announcements made by TowardEX are recorded under the *AS-TOWARDEX* object.

6. TowardEX honors MEDs (Multi Exit Discriminators) from customers.

Appendix C: BGP – Customer Peer Types

Customers can request different set of routes from TowardEX. It is recommended that customers take advantage of the announced BGP communities, however for those who would appreciate simplicity; TowardEX offers these different peer types, each with different route-sets, below.

Full Routes

TowardEX will send the customer the complete BGP internet routing table. This is recommended configuration for customers purchasing IP transit.

Full Routes with Default

TowardEX will send the customer the complete BGP internet routing table, but will also announce a last resort (0.0.0.0/0 for IPv4, ::0/0 for IPv6) default route.

Customer Routes

TowardEX will only send our own internal and customer routes, as we would send to a peering partner.

Customer Routes with Default

TowardEX will only send our own internal and customer routes, but will also announce a last resort (0.0.0.0/0 for IPv4, ::0/0 for IPv6) default route. This is recommended for customers running on low-memory routers seeking to be multihomed. The default route provides redundancy for multihoming purposes, while more specific customer routes provide a more efficient routing.

Default Route

TowardEX will only send a last resort (0.0.0.0/0 for IPv4, ::0/0 for IPv6) default route.