



Publication 1064

BGP Routing Policy Manual for Advanced Users

TOWARDEX Carrier Services
Backbone Operations

APRIL 9, 2012

Prepared by
TowardEX Technology Services Corporation



USER MANUAL
CUSTOMER COMMUNICATIONS

P.O. BOX 51068
BOSTON, MA 02205
(617) 849-7278

www.towardex.com

Copyright © 2006-2012 TowardEX Technology Services Corporation.
All rights reserved.

TowardEX Technologies International, Inc.
d/b/a TOWARDEX Carrier Services
P.O. Box 51068
Boston, MA 02205

These documents have been prepared by, and remain the sole property of TowardEX Technologies, Inc. No part of the contents of this document may be reproduced or transmitted in any form or by any means without the written permission of TowardEX Technologies, Inc.

TowardEX Technologies, Inc. makes no representations or warranties with respect to the contents or use of this document, and specifically disclaims any express or implied warranties of merchantability or fitness for any particular purpose. Further, TowardEX Technologies, Inc. reserves the right to revise this publication and to make changes to its content, at any time, without obligation to notify any person or entity of such revisions or changes. Customer should always verify with its own technical resource before accepting any recommendations presented in these documents.

Limitation of Liability

In no event, whether based in contract or tort (including negligence) shall TowardEX Technologies, Inc. be liable for incidental, consequential, indirect, special or punitive damages of any kind, for loss or revenue, loss of business, or other financial loss arising out of or in connection with the sale, installation, maintenance, use, performance, failure, or interruption of its products or design documents, even if TowardEX Technologies, Inc. or its authorized reseller or authorized contractor has been advised of the possibility of such damages.

GVN-M104253

TABLE OF CONTENTS

INTRODUCTION	4
PURPOSE	4
Customer Requirements	4
KEY DIFFERENTIATORS	5
Route Propagation Control and Identification Using BGP Communities	5
Granular Route Export Controls.....	5
Proxy Communities.....	5
Remote-Triggered Blackhole Routing.....	5
Controlling Hot-Potato vs. Cold-Potato Routing for Inbound Traffic	5
BGP COMMUNITIES	6
INFORMATIONAL COMMUNITIES (ANNOUNCED)	6
Area Code “CAA”	6
Relationship Type Code “T”	7
Example Informational Communities	7
ACTION COMMUNITIES (RECEIVED)	8
Calling Code “NNNNN”	8
Request for Action Code “R”	9
Example Action Communities	10
LOCAL-PREFERENCE COMMUNITIES (RECEIVED)	11
SPECIAL-USE COMMUNITIES (RECEIVED)	11
CUSTOMER BGP POLICIES	12
ROUTE ANNOUNCEMENT POLICY	12
MAXIMUM PREFIXES POLICY	12
DUAL-HOMED CUSTOMERS USING PRIVATE AS.....	13
BGP CUSTOMER PEER TYPES	14
Full Routes	14
Full Routes with Default	14
Customer Routes	14
Customer Routes with Default	14
Default Route	14

TOWARDEX—THE PERFORMANCE NETWORK AND ENGINEERS YOU CAN RELY ON.™

INTRODUCTION

Purpose

This manual outlines the routing policies used by TowardEX Carrier Services for multi-homed customers interconnecting with the TowardEX Network (AS 27552) using Border Gateway Protocol (BGP). The manual is intended specifically for advanced users with good experience and in-depth understanding of BGP and internetwork routing concepts.

Customer Requirements

In order to have a BGP peer with TowardEX, customer must be multi-homed, either within TowardEX's network or with another Internet Service Provider (ISP) network. If the customer is multi-homed with another ISP, the customer must also have a registered Autonomous System (AS) number from an appropriate Internet Registry (i.e. ARIN).

If customer is multi-homed only to TowardEX's own network and does not have an AS number, TowardEX will assign a private AS number to the customer. If customer is utilizing MCE (Managed Customer Equipment) Service, TowardEX will handle the entire process, including assignment of private AS number and configuration of managed equipment.

Unless customer has purchased MCE Service, registration of numbering resources, including AS numbers; and configuration of customer's own equipment to engage in BGP operations, are entirely customer's own responsibility. TowardEX engineering staff will refer to Account Team for Professional Services if assistance in any of these tasks is requested.

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 Management Center (NMC) by calling 617-849-7278 or emailing 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 configuration 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 NMC to verify *prior* to making any changes. Professional Services and Managed Customer Equipment (MCE) Service are also available for customers that require such resource from TowardEX to implement its desired custom changes.

[remainder of this page intentionally left blank]

Key Differentiators

TowardEX differentiates itself from its competition by providing the best technology and value to our customers. The same can be said for advanced users of our clientele who appreciate having as much granular control as possible when it comes to engineering of their internet traffic. No other provider to date (2012) in New England provides the same level of extensive control and customization that TowardEX offers for controlling customer BGP announcements.

Below is a summary of key differences between competing ISPs and TowardEX in providing routing control options to customers using BGP communities.

Route Propagation Control and Identification Using BGP Communities

TowardEX's BGP routing policy makes extensive use of community tags that meet or exceed all of RFC1998 guidelines in providing ways for multi-homed customers to influence routing of their prefixes. Most "multi-carrier blend" ISPs in New England provide only limited support for RFC1998, or worse, none at all.

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 ports, all the way down to port location of the adjacent network.

This granular control gives extensive routing options for experienced customers where normally they would have to ask their ISP to make *manual changes* on their backbone network. Such requests can involve unacceptable downtime risks and are often rejected.

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 solves this hurdle by providing special BGP community tags 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.

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.

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

TowardEX uses cold-potato routing for all incoming traffic with its upstream transit providers in order to leverage their SLA. Some content network customers have specific traffic engineering requirements that are strictly based on latency where cold-potato routing may not be useful. 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 cold-potato routing is disabled, inbound traffic to customer prefix will use hot-potato, closest-exit routing. Its behavior will be same as traffic received from a "*Settlement-Free Interconnection*" peer.

TOWARDEX—THE PERFORMANCE NETWORK AND ENGINEERS YOU CAN RELY ON.™

BGP COMMUNITIES

BGP communities are 32-bit values which are attached to a BGP route in order to carry additional information about that route. In common usage, this 32-bit value is split into two 16-bit values separated by a colon (:), where the first half typically represents the ASN that the community "belongs to" or "targets", and the second half typically represents the proprietary data.

Informational Communities (Announced)

Informational Communities are BGP communities added by TowardEX, in order to convey information about how and where the route was initially learned by our network. These communities will always have 5 digits in the second half, and are constructed in the following format:

27552:5CAAT

C	Area Code: Continent
AA	Area Code: Control Sector (Regional Area)
T	Relationship Type Code

Area Code "CAA"

Area Code consists of "CAA" pattern as identified above. Area Code is used to associate the geographical location of a route. For Information Communities, Area Code identifies the location in which a route has entered our network. For Action Communities, Area Code allows you to specify which geographical location you want your policy change to be applied to.

When applying policies toward TowardEX's upstream transit providers, in addition to specifying which upstream AS you want to influence, you can also specify the geographic location of a port by using Area Code. You may also do the exact same thing to TowardEX customer ASes.

C Continent	AA Sector	Regional Area	NTC Control Sector
0	00	All Locations – Apply Action Globally	GLOBAL
7	..	North America – United States & Canada	FDC
7	01	Boston BOS1-5 (Boston, MA) BSN1 (Somerville, MA — Boston North) BSN2 (Cambridge, MA — Boston North)	KZBW
7	02	New York JFK1 (New York, NY) EWR1 (Newark, NJ)	KZNY
8	..	EMEA – Europe & Middle East	EUC
8	02	Amsterdam AMS1 (Amsterdam, NL)	EHAA

Note: Use of the Area Code toward some 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.

[remainder of this page intentionally left blank]

Relationship Type Code “T”

Relationship Type Code consists of “T” pattern as identified previously. Relationship Type Code is only used by Informational Communities (Announced), to convey what type of adjacent network’s relationship the route has been learned from.

T Relationship Type	Description
1	Transit The route is from an upstream provider.
2	Peer The route is from a network that is directly peering with TowardEX.
4	Customer Route has originated from a TowardEX Customer.
5	Internal TowardEX is internally originating the route from its backbone.

Example Informational Communities

The following is set of examples of Informational Communities added by TowardEX on routes being announced to you:

BGP Community	Description
27552:57011	Transit Route from: North America, Boston Center This is a transit route learned from an upstream provider port in Boston, North America (7 01).
27552:57022	Peer Route from: North America, New York Center The route is from a network that is peering with TowardEX in New York, North America (7 02).
27552:57014	Customer Route from: North America, Boston Center The route is from a TowardEX customer port located in Boston, North America (7 01).
27552:57025	Internal Route from: North America, New York Center TowardEX is internally originating this route from a backbone equipment in New York, North America (7 02).

[remainder of this page intentionally left blank]

Action Communities (Received)

Action Communities are optional communities which may be added by TowardEX customers, to control route attributes or to control how routes are advertised to other networks. These communities may also be targeted to specific peer ASNs, specific locations (by continent, region), or specific classes of neighbors (transits, peers, or customers).

Action Communities use the full 32-bit values of BGP community string. The first half 16-bit field denoting target AS number, left of the colon is defined as **Calling Code**. The second half to the right of colon always starts with the number “4” and consists of the following structure:

NNNNN:4CAAR

NNNNN	Calling Code (AS number of target network)
C	Area Code: Continent
AA	Area Code: Control Sector (Regional Area)
R	Request for Action Code

Calling Code “NNNNN”

The Calling Code can be anywhere between 1 to 5 digits in length, and is used to denote the Target or Targets – in other words, which group of peers or which specific adjacent network AS you wish to apply your action to.

Below is a list of common Calling Codes available on our network:

NNNNN Calling Code	Description
0	All Networks – Global Apply action to ALL networks connected to TowardEX.
64512	All Networks – North America Apply action to networks connected to TowardEX in North America.
64513	All Networks – Europe Apply action to networks connected to TowardEX in Europe.
27552	TowardEX Customers Apply action to TowardEX customers connected using BGP.
6461	AboveNet Transit or Paid-Peer
174	Cogent Transit or Paid-Peer
7922	Comcast Transit or Paid-Peer
3356	Level 3 Transit or Paid-Peer
2914	NTT Transit or Paid-Peer
3257	Inteliquent Transit or Paid-Peer
4436	nLayer Transit or Paid-Peer
30071	OCCAID (IPv6 Only) Transit or Paid-Peer
XXXXX	AS Number of a Directly Connected Peer Use AS number of any directly connected peer [1 through 64511].

[remainder of this page intentionally left blank]

Request for Action Code “R”

Request for Action Code consists of “R” pattern at the last digit as identified previously. Request for Action Code is used to define what specific desired action you wish to take toward the target AS or groups of peers. The following set of actions is available:

R Request for Action	Description
0	No Export Do not export the route to target.
1	Prepend 27552 Lengthen the AS_PATH to target once.
2	Prepend 27552 27552 Lengthen the AS_PATH to target twice.
3	Prepend 27552 27552 27552 Lengthen the AS_PATH to target by three times.
5	Override Inclusive Request for Action Overrides “less-specific” routing policy — ignores inclusive Request for Action and just exports the route normally as is.
<p>Proxy Communities</p> <p>The following “Request for Action” codes allow you to have TowardEX transmit a peer-specific community to the target network by proxy, on your behalf. This gives you a powerful tool to not only influence the routing of your prefixes within the TowardEX network, but also influence the routing of TowardEX’ upstream provider or a peer.</p> <p>Note that not all networks support the full set of proxy community actions listed below.</p>	
4	Instruct Upstream to Treat Route as Peer If supported by the target upstream, forward a request to treat the route as a Peering Partner learned route.
6	Instruct Upstream to No Export If supported by the target upstream, forward a request to not export the route to anybody connected to their AS.
7	Instruct Upstream to No Export to Transits If supported by the target upstream, forward a request to not export the route to their upstream ISPs (transit providers). Only works against Tier-2 target networks. Tier-1 “transit-free” networks do not support this feature.
8	Instruct Upstream to No Export to Peers If supported by the target upstream, forward a request to not export the route to their peering partners.
9	Instruct Upstream to No Export to Customers If supported by the target upstream, forward a request to not export the route to their customers using BGP.

[remainder of this page intentionally left blank]

Example Action Communities

The following is set of examples of Action Communities that customers can use. Left column describes a valid BGP community string you can set, with description on the right hand column.

BGP Community	Description
0:40000	<p>No-Export To: Everyone at Everywhere (Global No-Export)</p> <p>Do not export to all peers, customers, transits at all geographical locations, everywhere.</p> <p>Equivalent to well-known “no-export” community.</p>
174:40000	<p>No-Export To: Cogent (AS174) at Everywhere</p> <p>Do not export to Cogent (AS174) at all geographical locations.</p>
6461:47010	<p>No-Export To: AboveNet (AS6461) at Boston Area</p> <p>Do not export to AboveNet (AS6461) in Boston. Route will still be advertised to AboveNet ports in other locations.</p>
7922:47022	<p>Prepend Twice To: Comcast (AS7922) at New York Area</p> <p>Prepend “27552 27552” to AS_PATH to Comcast port in New York only. Comcast ports in other locations will see the route normally, without the prepend.</p>
27552:40003	<p>Prepend Three-Times To: TowardEX Customers at Everywhere</p> <p>Prepend “27552 27552 27552” to AS_PATH to TowardEX Customers using BGP, at all geographical locations. Route is still advertised normally to peers and transit ports.</p>
0:40000 174:40005	<p>No-Export To: Everyone at Everywhere (Global No-Export)</p> <p>Except: Cogent (AS174) at Everywhere</p> <p>Only Cogent (AS174) will see your route normally. Route will not be advertised to anyone else, including TowardEX Customers.</p>
0:40001 3356:47015 27552:47015	<p>Prepend Once To: Everyone at Everywhere (Global)</p> <p>Except: Level 3 (AS3356) at Boston Area TowardEX Customers at Boston Area</p> <p>Prepend “27552” to AS_PATH to everybody, except TowardEX Customers and Level 3 (AS3356) port in Boston area.</p> <p>Everyone else, including TowardEX Customers and Level 3 ports outside of Boston will receive the “27552” AS_PATH prepend.</p>

[remainder of this page intentionally left blank]

Local-Preference Communities (Received)

Pursuant to RFC1998, networks should provide non-transitive BGP communities to alter Local Preference values to influence BGP best-path selection.

The local-preference attribute is not transitive, so this applies only to path selection within the TowardEX network. A value of 50 will create a "last resort route" that is never used or propagated to the rest of the Internet, so long as any other path is heard for the prefix in question. This route will only become active if no other route is heard.

By default, all received customer routes are assigned a Local Preference of 300. The following communities are available to alter the Local Preference:

BGP Community	Description
27552:50	Set Local Preference to 50 Route of last resort.
27552:100	Set Local Preference to 100 Default transit route.
27552:150	Set Local Preference to 150 Less preferred than a peer, more preferred than transit.
27552:200	Set Local Preference to 200 Default peer route.
27552:250	Set Local Preference to 250 Backup customer route.
27552:300	Default Behavior (Local Preference to 300) Default customer route.

Special-Use Communities (Received)

These communities are used for special-case situations. **WARNING: Incorrect use of these communities may result in partial or total destruction of connectivity to affected prefixes.** Please use them with careful consideration. Outages or service degradations caused by customer's configuration changes may not give claims for SLA credits.

BGP Community	Description
27552:911	Discard Traffic Remote-Triggered Blackhole / Null-Route Kills connectivity to the prefix. Will also request all transits and peers to do the same, if they support it.
27552:912	Non-Transitive Null-Route for Aggregate Prefixes Internal Use Only Used internally by TowardEX to signal aggregate route announcements without instructing adjacent networks to blackhole. This is for internal use only; it will not work on customer sessions.
27552:900	Reset MED Enforce Hot-Potato Routing Bypasses cold-potato routing established with upstream transits and customer ports. All ingress traffic will take closest-exit/hot-potato route; however will result in asymmetric routing.
27552:31337	Traffic Congestion Advisory (Received and Announced) Automatically set by TowardEX and can also be used by Customers to signal that path to this route is congested or is under a degraded condition. Ports triggering Congestion Advisory are alerted to NMC and NTC Maint Ops staff for capacity review and remedial actions.

Customer BGP Policies

Route Announcement Policy

1. TowardEX propagates prefixes shorter than or equal to the following lengths to its peers and upstream transit providers:
 - a. **IPv4: /24**
 - b. **IPv6: /48**

Prefixes that are longer than the above (smaller IP blocks) are not propagated outside of the TowardEX network.

2. Customers are allowed to advertise smaller networks with prefixes longer than those listed above, for the purpose of triggering blackhole routing or for traffic-engineering purposes. **However, customers must beware of the Maximum Prefixes Policy as noted below. Exceeding the configured Maximum Prefixes will interrupt and disable your BGP session, potentially causing downtime to announced prefixes.**
3. TowardEX requires its customers to properly register their announced routes in the Internet Routing Registry (IRR). Failure to register your routes will cause TowardEX to forcibly register them by proxy, on your behalf.
 - a. TowardEX automatically updates your BGP prefix/announcement filters at 05:00 GMT every day based on your IRR entries.
 - b. TowardEX supports most major IRR databases, including RADB, RIPE, ALTDB, LEVEL3, etc.
 - c. All internet route announcements made by TowardEX are recorded under the *AS-TOWARDEX* object.
4. TowardEX honors MEDs (Multi-Exit Discriminators) from customers.
5. TowardEX does not accept announcements with private AS numbers (64512 through 65535) in their AS_PATH. However, customers that are multi-homing solely into the TowardEX network using private AS are exempted from this policy. More information is provided under the section titled “**Dual-Homed Customers using Private AS.**”
6. Announcements with invalid prefixes, including RFC 1918 private internet numbers and other reserved networks are not permitted.

Maximum Prefixes Policy

1. By default, the following Maximum-Prefixes Limits are active on all customer BGP sessions:
 - a. IPv4: 300 prefixes
 - b. IPv6: 50 prefixes
2. Please contact Customer Care if you need to advertise more prefixes, so that we can have your sessions reconfigured as appropriate with higher limits. **Exceeding the aforementioned maximum-prefixes limits will immediately interrupt and disable your BGP session, causing outage to affected prefixes.**

[remainder of this page intentionally left blank]

TOWARDEX—THE PERFORMANCE NETWORK AND ENGINEERS YOU CAN RELY ON.™

Dual-Homed Customers using Private AS

1. Customers dual-homing (or multi-homing) solely into the TowardEX network, without any other ISP providing BGP service, can use a Private AS number in their BGP configuration with TowardEX.
 - a. TowardEX uses **AS 64552** for most internally dual-homed customers using Private AS. TowardEX will assign the customer with a private AS number to use.
2. Any BGP advertisement originated from such configuration involving Private AS will have their AS_PATH masqueraded to terminate under AS 27552. Private AS numbers will not be leaked outside of the TowardEX network.
3. If the announced prefix is part of a TowardEX-allocated IP address space, TowardEX will not propagate the prefix as it is part of a larger aggregate announcement already announced by TowardEX.

[remainder of this page intentionally left blank]

BGP Customer Peer Types

Customers can request different set of routes from TowardEX. It is recommended for customers to take advantage of the announced BGP communities. However, for those who would prefer a simpler configuration, TowardEX offers the following peer types, each with different route announcement sets:

Full Routes

TowardEX will send the complete internet BGP routing table. This is the default mode of operation for customers multi-homed between TowardEX and another ISP.

Full Routes with Default

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

Customer Routes

TowardEX will only send its own internal and customer routes, as normally advertised to a peer.

Customer Routes with Default

TowardEX will only send its own internal and customer routes, but will also announce a last resort/default route. This is recommended for customers running on low-memory routers or Layer-3 switches. The default route provides redundancy for multi-homed IP transit, while the more specific customer routes provide a more efficient routing.

Default Route

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