



Global routing policy and BGP Communities use

CLARA Network Engineering Group
April 2007

This document presents the Global routing policy adopted by RedCLARA regarding its BGP peering strategy, prefixes announcement rules and communities definition for better control the backbone routing.

VERSION MANAGEMENT

This guide outlines the global routing policy adopted by RedCLARA. When new procedures are required or other changes made, it will be updated accordingly, and the new version release will be recorded in the table below.

Version	Modification	Date	Reviewed by
preliminary	First draft	15-Oct-2006	Eriko Porto
1.0	Corrections and changes	01-Mar-2007	Eriko Porto
1.1	Corrections and changes	12-Apr-2007	Eriko Porto

Summary

1. Introduction	4
2. Global Routing	4
2.1. Global routing main rules.....	4
2.2. Caveats	5
3. BGP communities	5
3.1. Community values	5
3.1.1. Basic Communities	6
3.1.2. NRENs Communities	6
3.1.3. ITNs Communities	6
3.1.4. Connectors Communities.....	7
3.1.5. Blocked Communities	7
3.1.6. Local-preference manipulation.....	7
Annex A – BGP templates	8

1. Introduction

Through its new international agreements, RedCLARA has now the ability to establish BGP peering sessions with several regional networks, using its access links from the ALICE and the WHREN-LILA projects, and some of these peering sessions have multi-homed configuration (possibility to access the same network through different connections).

RedCLARA recognizes three basic types of peering: ITN (International Transit Network), Connector (regional network) and LA-NREN (National R&E Networks in Latin America). Each one of these types of established sessions to exchange traffic needs different treatment, and requires specific configurations to assure the best use of the resources and optimal routing.

The purpose of this document is to describe the general policy that will be adopted by RedCLARA for exchange traffic with its peering networks, and the management of essential BGP attributes to accomplish the policy thereby defined with optimal routing performance.

2. Global Routing

The global routing policy adopted by RedCLARA is geographic proximity – LA-NRENs will preferentially use the “nearest” exit to exchange traffic with other international networks. *Nearest* in this case means the path with less delay measured between the entry point in the backbone (RedCLARA PoP) and the existing alternative exit points for a multi-homed scenario.

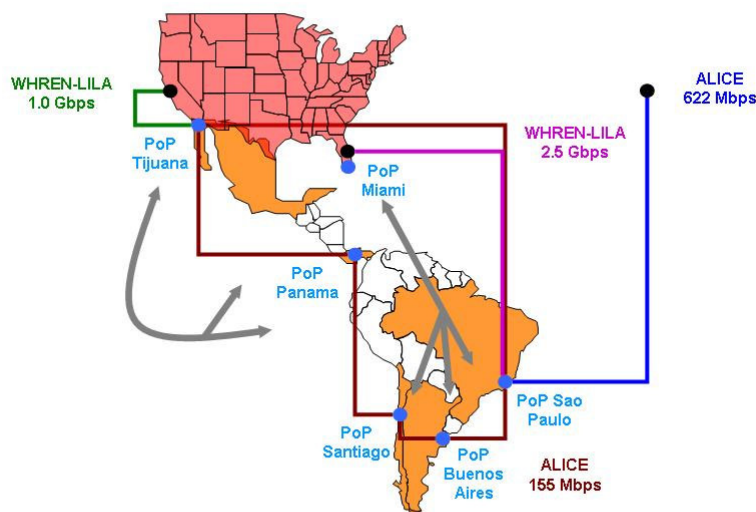


Figure 1 – Geographic proximity routing

Figure 1 shows the configuration of the routing policy based in the current topology of the network. NREN Inbound traffic in the PoPs of Tijuana and Panama, which should be routed to a network accessible via both west and east links, will preferentially exit through the west link, in contrast, NREN Inbound traffic in the PoPs of Santiago, Buenos Aires and Sao Paulo, which should be routed to a network accessible via both west and east links, will preferentially exit through the east link.

This is the default configuration and behavior for normal operation of the network. For some special cases where the researchers may need a singular routing behavior to meet their application specific requirements, a special configuration may be applied. These requests must be submitted in advance to the NEG for feasibility study and approval.

2.1. Global routing main rules

RedCLARA global routing policy is going to be regulated by the following set of rules:

- Traffic flow will follow geographic proximity based on the backbone topology and total round trip delay between the ingress point of the traffic in the backbone, and all possible egress points where multi-homed scenario applies;
- LA-NRENs will preferentially use the exit with the least path delay to exchange traffic with other international networks;
- The default behavior will be achieved by means of combined IGP metrics and the BGP decision process, which prefer ebgp paths over ibgp paths with the same AS-PATH, origin types and MED values;
- For the previous rule to work, all BGP MED values received from multi-homed connectors or ITNs will be set to zero;
- All internal and received prefixes will be marked to one or more BGP communities as described in section 3 of this document;
- NRENs receive all prefixes unless otherwise established;
- ITNs and Connectors receive only NREN prefixes, unless there is a specific agreement for transit;
- Special filters using community-lists and prefix-lists will be applied when necessary to achieve the behavior intended to the network traffic flow;
- To avoid asymmetries in the outbound traffic and maintain coherence with the policy adopted here, local-preference manipulation by means of special communities published by RedCLARA peers (where available), and as-prepend will be applied.

2.2.

Caveats

Normally RedCLARA will not be sending the communities defined in section 3 to its peers. If one NREN or other peer wishes to receive RedCLARA communities for troubleshooting purposes, this should be requested to the NOC in order to adjust the BGP configurations for sending this attribute to the requesting peer. By default the only community sent to peers is the one defined by the peer itself to manipulate its internal local-preference, when applicable.

3.

BGP communities

The BGP community path attribute is an optional transitive attribute of variable length as defined in RFC 1771 (BGP-4). The attribute consists of a set of four octet values, each of which specifies a community. The community attribute values are encoded using an AS number in the first two octets, with the remaining two octets defined by the AS.

As defined in RFC 1997 (BGP Communities Attribute), a community is a group of destinations (i.e. prefixes) that share some common attribute. Each destination can belong to multiple communities. All prefixes with the community attribute belong to the communities listed in the attribute.

Communities are particularly suitable for route filtering design, troubleshooting and BGP attributes manipulation. Moreover, they can provide mechanisms for BGP customer prefixes management. This section describes the usage of the community attribute within RedCLARA – AS 27750. All the existing prefixes in the BGP mesh of RedCLARA backbone should be tagged by communities formatted according to the directives described in this document.

3.1.

Community values

The community attribute values shall be encoded using an autonomous system number in the first two octets as described in RFC 1997, therefore for RedCLARA the communities will be of the type 0x6C66xxxx, varying from 0x6C660000 to 0x6C66FFFF. In decimal notation RedCLARA communities should be represented as 27750:y, where “y” is a decimal value ranging from 0 to 65535.

The only exception of this rule is the following well-known communities that have global significance, which will be implemented in RedCLARA following the recommendations of RFC 1997, and will be used to tag particular prefixes when applicable:

- **NO_EXPORT (0xFFFFF01)** – all routes received carrying communities attribute containing this value MUST NOT be advertised outside a BGP confederation boundary

- **NO_ADVERTISE (0xFFFFF02)** – all routes received carrying communities attribute containing this value MUST NOT be advertised to other BGP peers.
- **NO_EXPORT_SUBCONFED (0xFFFFF03)** – all routes received carrying communities attribute containing this value MUST NOT be advertised to external BGP peers (this includes peers in other members autonomous systems inside a BGP confederation).

3.1.1. Basic Communities

The basic community values identify the category of the prefix present in the BGP routing table. Every prefix present in the routing table must belong and be correctly marked to one of the communities described in Table 1

NREN prefixes are usually transit prefixes, if for some agreement reason any prefix or prefixes belonging to one NREN should not be given transit, it will be marked with the community NREN_NO_TRANSIT and will be handled by the specific policy rules.

ITN, CONNECTOR and COMERCIAL prefixes are no-transit prefixes by default, if for some agreement reason any prefix or prefixes belonging to one ITN, CONNECTOR or COMERCIAL peering should be given transit, it will be marked with the community TRANSIT and will be handled by the specific policy rules.

Table 1 – Basic communities

27750:64000	INTERNAL_PREFIXES
27750:64100	NREN
27750:64109	NREN_NO_TRANSIT
27750:64200	ITN
27750:64209	ITN_TRANSIT
27750:64300	CONNECTOR
27750:64309	CONNECTOR_TRANSIT
27750:64400	COMERCIAL
27750:64409	COMERCIAL_TRANSIT
27750:64500	BACKUP_ROUTES
27750:64600	IPv6_ROUTES

3.1.2. NRENs Communities

These communities are used to identify the NREN originating the prefix. It will follow the convention of representing the NREN ASN in the two last octets of the community value: **27750:<nren_asn>**, as represented in Table 2. All NRENs prefixes MUST also belong to the more generic NREN prefix community defined in the section 3.1.1 of this document (the 27750:64100 community)

Table 2 – NREN community examples

27750:1916	prefixes received from Brazilian NREN – RNP
27750:18592	prefixes received from Mexican NREN – CUDI
27750:11340	prefixes received from Chilean NREN – REUNA
27750:3597	prefixes received from Argentinean NREN – RETINA

Note: NRENs without a proper ASN assigned by a Regional Internet Registry and using private ones, MUST NOT be tagged by this rule, instead they will only use the more generic NREN community described in section 3.1.1 of this document (the 27750:64100 community).

3.1.3. ITNs Communities

These communities classify the ITN originating the prefix. It will follow the convention of representing the ITN ASN in the last two octets of the community value: **27750:<itn_asn>**, as represented in Table 3. ITN prefixes MUST also belong to the more generic ITN prefix community defined in the section 3.1.1 of this document (the 27750:64200 community).

Table 3 – ITN community examples

27750:20965	GEANT
27750:11537	ABILENE

3.1.4. Connectors Communities

These communities identify the Connector originating the prefix. It will follow the convention of representing the Connector ASN in the two last octets of the community value: **27750:<connector_asn>**, as represented in Table 4. All Connectors prefixes MUST also belong to the more generic Connector prefix community defined in the section 3.1.1 of this document (the 27750:64300 community)

Table 4 – Connector community examples

27750:24	NISN
27750:20080	AMPATH
27750:293	ESNet
27750:7575	AARNet

3.1.5. Blocked Communities

These are special communities used to identify prefixes that should not be advertised to specific groups. These communities are coarser than the NO_ADVERTISE and NO_EXPORT communities, and gives to RedCLARA the ability to better troubleshooting and control the routing inside the AS. The values of these communities are represented in Table 5.

Table 5 – Blocked communities

27750:64900	BLACK_HOLE
27750:64901	BLOCK_TO_NREN
27750:64902	BLOCK_TO_ITN
27750:64903	BLOCK_TO_CONNECTOR
27750:64904	BLOCK_TO_COMERCIAL

3.1.6. Local-preference manipulation

The communities defined in Table 6 are to be used for local-preference manipulation by RedCLARA peers. Any RedCLARA's BGP peer may use these specific BGP communities to influence RedCLARA's setting of its local-preference for the prefixes received from them. The community may cause the preference to be lower or higher than the default preferences, and may be useful in cases where the peer has more than one point for exchanging traffic with RedCLARA (multi-homed scenario), and would like to influence our internal BGP choice of paths to send traffic back to the peer. Refer to the Table 6 for the community to use and its respective effect.

RedCLARA's default local-preference for the prefixes received is set to 100, if the peer sends to RedCLARA through one of its connection a community with the value 27750:260 for some prefixes for instance, RedCLARA will set the local-preference for those prefixes at that path to 260.

Table 6 – Communities for local-preference manipulation

BGP local-preference	Community	Meaning
100	N/A	Default local-preference
80	N/A	Internal low preference
120	N/A	Internal high preference
50	27750:65050	Set low local-preference
150	27750:65150	Set high local-preference

Note: Only the community values specified in Table 6 are supported for local-preference manipulation, different values will be automatically ignored and discarded

Annex A – BGP templates

Below are provided the configuration templates for the Cisco routers planned to enforce the policy described here in this document.

Cisco 12006 – clara-br router

```
!  
router bgp 27750  
address-family ipv4  
  redistribute static route-map static-to-bgp  
  neighbor INTERNAL activate  
  neighbor INTERNAL next-hop-self  
  neighbor INTERNAL soft-reconfiguration inbound  
  neighbor INTERNAL send-community  
...  
  neighbor eGEANT prefix-list MARTIANS in  
  neighbor eGEANT prefix-list MARTIANS out  
  neighbor eGEANT route-map GEANT-IN in  
  neighbor eGEANT route-map ITN-OUT out  
...  
  neighbor eRNP prefix-list MARTIANS in  
  neighbor eRNP prefix-list MARTIANS out  
  neighbor eRNP route-map RNP-IN in  
...  
  neighbor eREACCIUN prefix-list MARTIANS in  
  neighbor eREACCIUN prefix-list MARTIANS out  
  neighbor eREACCIUN route-map REACCIUN-IN in  
...  
  neighbor eAMPATH prefix-list MARTIANS in  
  neighbor eAMPATH prefix-list MARTIANS out  
  neighbor eAMPATH route-map AMPATH-IN in  
  neighbor eAMPATH route-map CONN-OUT out  
...  
  neighbor eABILENE prefix-list MARTIANS in  
  neighbor eABILENE prefix-list MARTIANS out  
  neighbor eABILENE route-map ABILENE-IN in  
  neighbor eABILENE route-map ITN-OUT out  
...  
exit-address-family  
...  
!--- community list 10 - permit only NRENS and internal prefixes  
!--- access-list 1 & 2 - standard lines  
!--- one access list to filter prefixes from each ITN (10-GEANT & 20-Abilene)  
!--- one access lis to all connectors (40)  
  
ip bgp-community new-format  
ip community-list 10 permit 27750:64100  
ip community-list 10 permit 27750:64000  
ip community-list 10 deny  
ip as-path access-list 1 permit .*  
ip as-path access-list 2 permit ^$  
ip as-path access-list 10 permit _20965_  
ip as-path access-list 20 permit _11537_  
ip as-path access-list 30 permit _20080_  
ip as-path access-list 40 permit _2153_  
ip as-path access-list 40 permit _101_  
ip as-path access-list 40 permit _293_  
ip as-path access-list 40 permit _24_  
ip as-path access-list 40 permit _7575_  
ip as-path access-list 40 permit _19401_  
...  
access-list 10 permit 200.0.204.0 0.0.3.255  
...  
  
!--- map NREN prefixes to global NREN community (27750:64100)  
!--- and to the specific community of the NREN  
  
route-map RNP-IN permit 10  
  set community 27750:1916 27750:64100  
!  
route-map REACCIUN-IN permit 10  
  set community 27750:20312 27750:64100  
!
```



```

!--- Abilene prefixes received from GEANT are set to the backup
!--- community (27750:64500) with lower local preference

route-map GEANT-IN permit 10
  match as-path 20
  set community 27750:20965 27750:64200 27750:64500
  set local-preference 80

!--- all other prefixes from GEANT

route-map GEANT-IN permit 20
  set community 27750:20965 27750:64200
!
route-map AMPATH-IN permit 10
  set community 27750:20080 27750:64300
!

!--- GEANT prefixes received from Abilene are set to the backup
!--- community (27750:64500) with lower local preference

route-map ABILENE-IN permit 10
  match as-path 10
  set metric 0
  set community 27750:11537 27750:64200 27750:64500
  set local-preference 80
!
route-map ABILENE-IN permit 20
  set metric 0
  set community 27750:11537 27750:64200
!

!--- CONNECTORS and ITNS receive only NREN prefixes (community-list 10)

route-map ITN-OUT permit 10
  match community 10
!
route-map CONN-OUT permit 10
  match community 10
!

!--- Set internal prefixes to the correct community and origin

route-map static-to-bgp permit 10
  match ip address 10
  set origin igp
  set community 27750:64000
!

```

Cisco 12006 – clara-mx router

```

!
router bgp 27750
  address-family ipv4
    redistribute static route-map static-to-bgp
    neighbor INTERNAL activate
    neighbor INTERNAL next-hop-self
    neighbor INTERNAL send-community
    neighbor INTERNAL soft-reconfiguration inbound
  ...
  neighbor eCUDI prefix-list MARTIANS in
  neighbor eCUDI prefix-list MARTIANS out
  neighbor eCUDI route-map CUDI-IN in
  ...
  neighbor eABILENE prefix-list MARTIANS in
  neighbor eABILENE prefix-list MARTIANS out
  neighbor eABILENE route-map ABILENE-IN in
  neighbor eABILENE route-map ITN-OUT out
  ...
  neighbor eRAGIE prefix-list MARTIANS in
  neighbor eRAGIE prefix-list MARTIANS out
  neighbor eRAGIE route-map RAGIE-IN in
  ...
  neighbor eRAICES prefix-list MARTIANS in
  neighbor eRAICES prefix-list MARTIANS out
  neighbor eRAICES route-map RAICES-IN in
  ...
  neighbor eCR2NET prefix-list MARTIANS in

```

```

neighbor eCR2NET prefix-list MARTIANS out
neighbor eCR2NET route-map CR2Net-IN in
...
neighbor eRENIA prefix-list MARTIANS in
neighbor eRENIA prefix-list MARTIANS out
neighbor eRENIA route-map RENIA-IN in
...
neighbor ePACIFIC-WAVE prefix-list MARTIANS in
neighbor ePACIFIC-WAVE prefix-list MARTIANS out
neighbor ePACIFIC-WAVE route-map METRIC2ZERO in
neighbor ePACIFIC-WAVE route-map CONN-OUT out
...
neighbor eCalREN prefix-list MARTIANS in
neighbor eCalREN prefix-list MARTIANS out
neighbor eCalREN route-map CalREN-IN in
neighbor eCalREN route-map ITN-OUT out
...
neighbor 207.231.241.1 route-map PNWGP-IN in
neighbor 207.231.241.4 peer-group ePACIFIC-WAVE
neighbor 207.231.241.4 route-map AARNet-IN in
neighbor 207.231.241.14 peer-group ePACIFIC-WAVE
neighbor 207.231.241.14 route-map NLR-IN in
neighbor 207.231.246.2 peer-group ePACIFIC-WAVE
neighbor 207.231.246.2 route-map ESNet-IN in
neighbor 207.231.246.3 peer-group ePACIFIC-WAVE
neighbor 207.231.246.3 route-map NISN-IN in
...
exit-address-family
...
ip bgp-community new-format
ip community-list 10 permit 27750:64100
ip community-list 10 permit 27750:64000
ip community-list 10 deny
ip as-path access-list 1 permit .*
ip as-path access-list 2 permit ^$
ip as-path access-list 10 permit _20965_
ip as-path access-list 20 permit _11537_
ip as-path access-list 30 permit _20080_
ip as-path access-list 40 permit _2153_
ip as-path access-list 40 permit _101_
ip as-path access-list 40 permit _293_
ip as-path access-list 40 permit _24_
ip as-path access-list 40 permit _7575_
ip as-path access-list 40 permit _19401_
...
access-list 10 permit 200.0.204.0 0.0.3.255
...
route-map RAICES-IN permit 10
  set community 27750:64100
!
route-map AARNet-IN permit 10
  set community 27750:7575 27750:64300
!
route-map NISN-IN permit 10
  set community 27750:24 27750:64300
!
route-map ITN-OUT permit 10
  match community 10
!
route-map CUDI-IN permit 10
  set community 27750:18592 27750:64100
!
route-map CalREN-IN permit 10
  match as-path 20
  set local-preference 80
  set community 27750:2153 27750:64300 27750:64500
!
route-map CalREN-IN permit 20
  set community 27750:2153 27750:64300
!
route-map CR2Net-IN permit 10
  set community 27750:2146 27750:64100
!
route-map ABILENE-IN permit 10
  match as-path 10
  set metric 0
  set local-preference 80

```

```

set community 27750:11537 27750:64200 27750:64500
!
route-map ABILENE-IN permit 20
set metric 0
set community 27750:11537 27750:64200
!
route-map ESNet-IN permit 10
set community 27750:293 27750:64300
!
route-map CONN-OUT permit 10
match community 10
!
route-map RAGIE-IN permit 10
set community 27750:64100
!
route-map PNWGP-IN permit 10
set community 27750:101 27750:64300
!
route-map static-to-bgp permit 10
match ip address 10
set origin igp
set community 27750:64000
!
route-map NLR-IN permit 10
set community 27750:19401 27750:64300
!
route-map RENIA-IN permit 10
set community 27750:64100
!

```

Cisco 12006 – clara-cl router

```

!
router bgp 27750
address-family ipv4
redistribute static route-map static-to-bgp
neighbor INTERNAL activate
neighbor INTERNAL next-hop-self
neighbor INTERNAL send-community
neighbor INTERNAL soft-reconfiguration inbound
...
neighbor eREUNA prefix-list MARTIANS in
neighbor eREUNA prefix-list MARTIANS out
neighbor eREUNA route-map REUNA-IN in
...
neighbor eRAAP prefix-list MARTIANS in
neighbor eRAAP prefix-list MARTIANS out
neighbor eRAAP route-map RAAP-IN in
...
neighbor eCEDIA prefix-list MARTIANS in
neighbor eCEDIA prefix-list MARTIANS out
neighbor eCEDIA route-map CEDIA-IN in
...
exit-address-family...
...
ip bgp-community new-format
ip as-path access-list 1 permit .*
ip as-path access-list 2 permit ^$
...
access-list 10 permit 200.0.204.0 0.0.3.255
...
route-map RAAP-IN permit 10
set community 27750:64100
!
route-map CEDIA-IN permit 10
set community 27750:27841 27750:64100
!
route-map REUNA-IN permit 10
set community 27750:11340 27750:64100
!
route-map static-to-bgp permit 10
match ip address 10
set origin igp
set community 27750:64000
!

```

Cisco 12006 – clara-ar router

```
!  
router bgp 27750  
address-family ipv4  
redistribute static route-map static-to-bgp  
neighbor INTERNAL activate  
neighbor INTERNAL next-hop-self  
neighbor INTERNAL send-community  
neighbor INTERNAL soft-reconfiguration inbound  
...  
neighbor eRETINA prefix-list MARTIANS in  
neighbor eRETINA prefix-list MARTIANS out  
neighbor eRETINA route-map RETINA-IN in  
...  
neighbor eRAU prefix-list MARTIANS in  
neighbor eRAU prefix-list MARTIANS out  
neighbor eRAU route-map RAU-IN in  
...  
exit-address-family  
...  
ip bgp-community new-format  
ip as-path access-list 1 permit .*  
ip as-path access-list 2 permit ^$  
...  
access-list 10 permit 200.0.204.0 0.0.3.255  
...  
route-map RETINA-IN permit 10  
set community 27750:3597 27750:64100  
!  
route-map RAU-IN permit 10  
set community 27750:1797 27750:64100  
!  
route-map static-to-bgp permit 10  
match ip address 10  
set origin igp  
set community 27750:64000  
!
```

Cisco 12006 – clara-pa router

```
!  
router bgp 27750  
address-family ipv4  
redistribute static route-map static-to-bgp  
neighbor INTERNAL activate  
neighbor INTERNAL next-hop-self  
neighbor INTERNAL send-community  
neighbor INTERNAL soft-reconfiguration inbound  
...  
neighbor eREDCYT prefix-list MARTIANS in  
neighbor eREDCYT prefix-list MARTIANS out  
neighbor eREDCYT route-map RedCyT-IN in  
...  
neighbor eRENATA prefix-list MARTIANS in  
neighbor eRENATA prefix-list MARTIANS out  
neighbor eRENATA route-map RENATA-IN in  
...  
exit-address-family  
...  
ip bgp-community new-format  
ip as-path access-list 1 permit .*  
ip as-path access-list 2 permit ^$  
...  
access-list 10 permit 200.0.204.0 0.0.3.255  
...  
route-map RedCyT-IN permit 10  
set community 27750:20312 27750:64100  
!  
route-map RENATA-IN permit 10  
set community 27750:27817 27750:64100  
!  
route-map static-to-bgp permit 10  
match ip address 10  
set origin igp  
set community 27750:64000  
!
```