
# Implementation of native IPv6 for RedCLARA

CLARA Network Engineering Group

August 2005

This document presents the proposal for native IPv6 routing implementation within the backbone of RedCLARA.

CLARA/NEG/2005-0013
Version-1.1

VERSION MANAGEMENT

This document outlines the IPv6 addressing and routing plan for the backbone of RedCLARA. When modifications to this document are required, it will be updated accordingly, and the new version release will be recorded in the table below.

| Version | Modification description | Date | Reviewed by |
|---|---|---|---|
| Preliminary | First draft | 07-Jul-2005 | GT-IPv6 (*) |
| 1.0 | Corrections and changes | 30-Jul-2005 | GT-IPv6 (*) |
| 1.1 | Corrections and changes | 31-Aug-2005 | Eriko Porto |
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |

* The GT-IPv6 is the CLARA Technical Forum workgroup established to discuss and provide recommendations for the CLARA-NEG regarding IPv6 infrastructure for RedCLARA. For more information about this workgroup see www.redclara.net.

**Summary**

## 1. Introduction

This document describes the configuration for Integrated Intermediate System-to-Intermediate System (IS-IS) for IPv6. IS-IS is the Interior Gateway Protocol (IGP) in use at the backbone of RedCLARA. The IS-IS protocol advertises link-state information throughout the network to create a picture of the network topology. IS-IS is an Open Systems Interconnection (OSI) hierarchical routing protocol that designates an intermediate system as a Level 1 or Level 2 device. The backbone of RedCLARA is a Level 2 domain IS-IS routing backbone. The Integrated IS-IS uses a single routing algorithm to support several network address families, such as IPv6, IPv4, and OSI. This document outlines the addressing schema for the backbone and the IPv6 routing plan for RedCLARA.

## 2. Overview of CLARA network

The CLARA organization – Cooperación Latino Americana de Redes Avanzadas – is responsible for the implementation and management of a network infrastructure that will interconnect the national academic networks (NRENs) of several Latin American countries.

The backbone of CLARA network is comprised of five main routing nodes, each one corresponding to a network PoP as depicted in Figure 1, with the five nodes connected in a ring topology. All other connections coming through LA-NRENs will have access to the backbone through some of the CLARA PoPs at the ring.

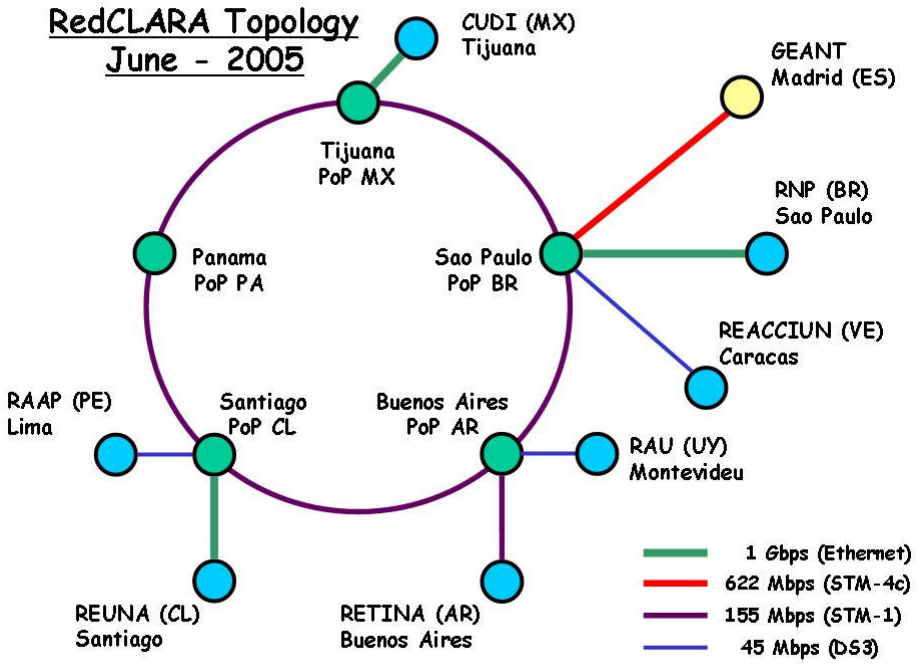


**Figure 1: Backbone of RedCLARA**

The five main IP nodes are located in Sao Paulo (BR), Buenos Aires (AR), Santiago (CL), Panama (PA) and Tijuana (MX).

## 3.
## IPv6 addressing for the backbone

The following IPv6 address block is assigned for CLARA organization:

- Address block – 2001:1348::/32
- Status – allocated
- Owner – Cooperación Latino Americana de Redes Avanzadas
- Owner ID – UY-CLAR-LACNIC
- Address – Colonia, 2066, 11200 – Montevideo – UY
- Phone – (+55 21) 3205-9660
- Block contact – ALG5

Table 1 summarizes the IPv6 addresses distribution for the equipment interfaces within the sub block allocated for the backbone. Detailed information about the whole IPv6 prefix address distribution and allocation will be available from another CLARA-NEG document to de released.

**Table 1: IPv6 addresses for the backbone**

| address space | subnet type | Hosts | Usage |
|---|---|---|---|
| 2001:1348:: /40 | | | RedCLARA |
| 2001:1348:: /48 | | | Backbone |
| 2001:1348::0000/120 | /126 | 2 | Point-to-point links between PoP routers |
| 2001:1348::0100/120 | | | Future use |
| 2001:1348::1000/120 | /128 | 1 | Routers loopback 0 interfaces |
| 2001:1348::1100/120 | /128 | 1 | Routers loopback 1 interfaces |
| | | | |
| | | | |
| 2001:1348:0001::/48 | /64 | 2 | Point-to-point links to LA-NRENs |
| 2001:1348:0002::/48 | | | Future use |
| | | | |
| 2001:1348:0003::/48 | | | |
| 2001:1348:0003:0000::/64 | /64 | 12 | Tijuana (MX) PoP Ethernet LAN |
| 2001:1348:0003:0001::/64 | | | Future use |
| 2001:1348:0003:0002::/64 | /64 | 6 | Sao Paulo (BR) PoP Ethernet LAN |
| 2001:1348:0003:0003::/64 | | | Future use |
| 2001:1348:0003:0004::/64 | /64 | 6 | Buenos Aires (AR) PoP Ethernet LAN |
| 2001:1348:0003:0005::/64 | | | Future use |
| 2001:1348:0003:0006::/64 | /64 | 6 | Santiago (CL) PoP Ethernet LAN |
| 2001:1348:0003:0007::/64 | | | Future use |
| 2001:1348:0003:0008::/64 | /64 | 6 | Panama (PA) PoP Ethernet LAN |
| | | | |
| | | | |
| | | | |

# 4.
# IS-IS for IPv6

## 4.1.
## Introduction

IPv6 enhancements to IS-IS allow IS-IS to advertise IPv6 prefixes in addition to IPv4 routes, and extensions to the IS-IS command-line interface (CLI) allow configuration of IPv6-specific parameters. IPv6 IS-IS extends the address families supported by IS-IS to include IPv6, in addition to OSI and IPv4, and supports either single-topology mode or multiple topology mode.

### 4.1.1. IS-IS Single-Topology

Single-topology support for IPv6 allows IS-IS for IPv6 to be configured on interfaces along with other network protocols (for example, IPv4 and Connectionless Network Service [CLNS]). All interfaces must be configured with the identical set of network address families. In addition, all routers in the IS-IS area (for Level 1 routing) or the domain (for Level 2 routing) must support the identical set of network layer address families on all interfaces.

When single-topology support for IPv6 is being used, either old- or new-style TLVs may be used. However, the TLVs used to advertise reachability to IPv6 prefixes use extended metrics. Cisco routers do not allow an interface metric to be set to a value greater than 63 if the configuration is not set to support only new-style TLVs for IPv4. In single-topology IPv6 mode, the configured metric is always the same for both IPv4 and IPv6.

### 4.1.2. IS-IS Multitopology

IS-IS multitopology support for IPv6 allows IS-IS to maintain a set of independent topologies within a single area or domain. This mode removes the restriction that all interfaces on which IS-IS is configured must support the identical set of network address families. It also removes the restriction that all routers in the IS-IS area (for Level 1 routing) or domain (for Level 2 routing) must support the identical set of network layer address families. Because multiple SPFs (Shortest Path First) calculations are performed, one for each configured topology, it is sufficient that connectivity exists among a subset of the routers in the area or domain for a given network address family to be routable.

## 4.2.
## IPv6 IS-IS deployment

The IS-IS routing in RedCLARA is a single Level 2 domain for the five main nodes of the backbone, and each one have a Level 2 link-state database with all the information for intra-area routing. For IPv6 implementation the backbone will adopt an IS-IS single topology mode with wide-metric style TLVs

Configuring IS-IS comprises two basic activities. The first one is to create an IS-IS routing process using protocol-independent IS-IS commands which is already set for RedCLARA. The second activity is to configure the operation of the IPv6 IS-IS protocol specific commands on the interfaces.

#### 4.2.1. IPv6 addresses

The table below shows the IPv6 addresses for the backbone interfaces and for the interfaces that connect with the LA-NRENs that already have IPv6 implemented in their networks. The procedures and policy to request and establish an IPv6 peering with RedCLARA are subject of another document from the CLARA-NEG to be released in the future.

**Table 2: IPv6 addresses for the interfaces of the main nodes**

| router | interface | destination | IP address |
|---|---|---|---|
| **clara-br** | pos-1 0 | PoP – Tijuana | 2001:1348::0005/126 |
| | pos-1 1 | PoP – Buenos Aires | 2001:1348::0026/126 |
| | loopback 0 | N/A | 2001:1348::1005/128 |
| | loopback 1 | N/A | 2001:1348::1126/128 |
| | pos-1 2 | NREN RPN (BR) | 2001:1348:0001:0001::/64 |
| | | | |
| **clara-mx** | pos-1 0 | PoP – Panama | 2001:1348::000D/126 |
| | pos-1 1 | PoP – Sao Paulo | 2001:1348::0006/126 |
| | loopback 0 | N/A | 2001:1348::100D/128 |
| | loopback 1 | N/A | 2001:1348::1106/128 |
| | giga-0 1 | NREN CUDI (MX) | 2001:1348:0001:0002::/64 |
| | | | |
| **clara-pa** | pos-1 0 | PoP – Santiago | 2001:1348::0015/126 |
| | pos-1 1 | PoP – Tijuana | 2001:1348::000E/126 |
| | loopback 0 | N/A | 2001:1348::1015/128 |
| | loopback 1 | N/A | 2001:1348::110E/128 |
| | | | |
| | | | |
| **clara-cl** | pos-1 0 | PoP – Buenos Aires | 2001:1348::001D/126 |
| | pos-1 1 | PoP – Panama | 2001:1348::0016/126 |
| | loopback 0 | N/A | 2001:1348::101D/128 |
| | loopback 1 | N/A | 2001:1348::1116/128 |
| | giga-0 1 | NREN REUNA (CL) | 2001:1348:0001:0004::/64 |
| | | | |
| **clara-ar** | pos-1 0 | PoP – Sao Paulo | 2001:1348::0025/126 |
| | pos-1 1 | PoP – Santiago | 2001:1348::001E/126 |
| | loopback 0 | N/A | 2001:1348::1025/128 |
| | loopback 1 | N/A | 2001:1348::111E/128 |
| | atm-1 0 | NREN RETINA (AR) | 2001:1348:0001:0005::/64 |

#### 4.2.1. Configuration guidelines

First is necessary to globally enable IPv6 using the **ipv6 unicast-routing global configuration** command, then it must be configured both IPv6 and IPv4 on an IS-IS interface for running the same IS-IS level, which is Level 2 for the RedCLARA routers.

For configuring Single-Topology IS-IS for IPv6 in the RedCLARA backbone router the following tasks must be completed:

1. **enable**
2. **configure terminal**
3. **router isis** *backbone*
4. **net** *49.0205.2000.0020.5xxx.00*
5. **is-type level-2-only**
6. **metric-style wide level-2**
7. **passive-interface** *type number*
8. **exit**
9. **interface** *type number*
10. **ipv6 address** *ipv6-prefix/prefix-length [eui-64]*
11. **ipv6 router isis** backbone
12. **exit**

**Notes:**
1 – Steps 1 to 5 are already configured in the backbone routers.
2 – Step 6 must be verified.
3 – Step 7 is already configured for all the interfaces that are not participating in the IS-IS routing process but need to inject the connected route to the IGP.
4 – Steps 9 and 10 must be configured for each interface with its respective IPv6 address.

#### 4.2.2. Customizing IPv6 IS-IS

The Cisco IPv6 IS-IS set of commands support a many customizations as new administrative distance, configure the maximum number of equal-cost paths, configure summary prefixes, and configure an IS-IS instance to advertise the default IPv6 route (::/0).

The initial configuration will use the defaults for these features that should meet our current requirements and features. If modifications from the defaults are needed new procedures will going to be defined by the CLARA-NEG and updated accordingly.

## 5.
## Exterior Gateway Protocol

All peering sessions between CLARA network and other ASes are established using the Border Gateway Protocol (BGP) version 4, the "de facto" EGP routing protocol in the Internet. In addition, all backbone routers inside CLARA network will run internal BGP (iBGP) peering sessions to exchange BGP routing information.

This module describes the configuration of Multiprotocol Border Gateway Protocol (BGP) for IPv6. BGP is an Exterior Gateway Protocol (EGP) used mainly to connect separate routing domains that contain independent routing policies (autonomous systems). BGP is also used within RedCLARA autonomous system and this variation is referred to as internal BGP (iBGP). Multiprotocol BGP is an enhanced BGP that carries routing information for multiple network layer protocol address families, for example, IPv6 address family and for IP multicast routes. All BGP commands and routing policy capabilities can be used with multiprotocol BGP.

### 5.1.
### Configuring IPv6 BGP Routing Process and BGP Router ID

BGP uses a router ID to identify BGP-speaking peers. The BGP router ID is 32-bit value that is often represented by an IPv4 address. By default, the Cisco IOS software sets the router ID to the IPv4 address of a loopback interface on the router.

Since the BGP process is already running in RedCLARA backbone and all routers have its own loopback interface up, there is no need to manually configure BGP Routing Process and BGP Router ID.

### 5.2.
### Configuring IPv6 MBGP Peer with peer group

By default, neighbors that are defined using the neighbor remote-as command in router configuration mode exchange only IPv4 unicast address prefixes. To exchange other address prefix types, such as IPv6 prefixes, neighbors must also be activated using the **neighbor activate** command in address family configuration mode for the other prefix types.

Peer groups are defined in router configuration mode using the **neighbor peer-group** command. To exchange other address prefix types, such as IPv6 prefixes, the peer groups must be activated using the **neighbor activate** command in the address family configuration mode for the other prefix types. Members of a peer group automatically inherit the address prefix configuration of the peer group.

IPv4 active neighbors cannot exist in the same peer group as active IPv6 neighbors. For the configuration of the routers separate peer groups for IPv4 peers and IPv6 peers must be created. For RedCLARA the IPv6 peer groups must be configured using the same names of the IPv4 peer groups but adding the **v6** suffix after the names, as represented in Table 3.

**Table 3: IPv6 peer groups**

| IPv4 peer group | IPv6 peer group |
|---|---|
| INTERNAL | INTERNALv6 |
| eCUDI | eCUDIv6 |
| eRNPv6 | eRNPv6 |
| eGEANT | eGEANTv6 |

The update source for the RedCLARA side will use the same interface being used by the IPv4 address family MBGP sessions.

IPv6 MBGP configuration example using peer groups:

1. **enable**
2. **configure terminal**
3. **router bgp** *27750*
4. **neighbor** *peer-group-name* **peer-group**
5. **neighbor** *peer-group-name* **remote-as** *autonomous-system-number*
6. **neighbor** *peer-group-name* **update-source** *interface-type*
7. **address-family ipv6 [unicast** | **multicast]**
8. **neighbor** *ipv6-address* **peer-group** *peer-group-name*
9. **neighbor** *peer-group-name* **activate**

Steps 1 to 3 are already configured for IPv4 in the backbone routers. Steps 4 to 9 must be configured for each IPv6 external peer of the backbone router, and for the internal MBGP peering sessions between the backbone routers.

### 5.3. IPv6 MBGP Advertised Routes

The routes for the internal links of the backbone should be injected at the IPv6 routing database using the **network** IPv6 address family command, and should be advertised by the MBGP to the other peers. This will be done to allow the internal interfaces to be reachable by external hosts. The PoP LAN IPv6 prefixes will also be injected at the IPv6 routing database, since this will allow the RedCLARA servers installed in the PoPs to run IPv6 experiments with external workstations.

### 5.4. IPv6 MBGP Prefix List

The following basic prefix list must be configured for the IPv6 peering sessions of the network in order to avoid propagation of bogus routes.

```
ipv6 prefix-list IPv6-BOGUS deny 2001:db8::/32 le 128
ipv6 prefix-list IPv6-BOGUS permit 2002::/16
ipv6 prefix-list IPv6-BOGUS deny 2002::/16 le 128
ipv6 prefix-list IPv6-BOGUS deny 0000::/8 le 128
ipv6 prefix-list IPv6-BOGUS deny fe00::/9 le 128
ipv6 prefix-list IPv6-BOGUS deny ff00::/8 le 128
ipv6 prefix-list IPv6-BOGUS permit 0::/0 le 48
ipv6 prefix-list IPv6-BOGUs deny 0::/0 le 128
```

Rationale:

- 2001:db8::/32 – IPv6 documentation prefix (RFC3849)
- 2002::/16 – only permits the /16 - no more-specifics
- 0000::/8 – is denied (loopback, unspecified, v4-mapped)
- FE00::/9 and FF00::/8 – multicast ranges are denied (RFC3513)
- 0::0/0 – all the rest of the IPv6 unicat address space is permitted
- 3FFE::/16 (6bone) has special treatment according to the 6bone rules

### 5.5. IPv6 Configuration Template

The following template demonstrates the basic MBGP configuration for the RedCLARA routers using peer groups under the IPv6 Address Family style.

```
!
! ========================
!  CONFIGURATION TEMPLATE
! ========================
!

router bgp 27750

!
! GENERAL CONFIGURATION
!

bgp log-neighbor-changes

!
! INTERNAL PEERS
!

neighbor INTERNALv6 peer-group
neighbor INTERNALv6 remote-as 27750
neighbor INTERNALv6 update-source Loopback0

! exclude the self address of the router from the list below

neighbor 2001:1348::1005 peer-group INTERNALv6
neighbor 2001:1348::100D peer-group INTERNALv6
neighbor 2001:1348::1015 peer-group INTERNALv6
neighbor 2001:1348::101D peer-group INTERNALv6
neighbor 2001:1348::1025 peer-group INTERNALv6

!
! EXTERNAL PEERS
! NAME is the name of the NREN or Regional Network
!

neighbor eNAMEv6 peer-group
neighbor eNAMEv6 remote-as xxxxxx
neighbor eNAMEv6 update-source interface
neighbor ipv6_address remote-as xxxxxx
neighbor ipv6_address password md5_password

!
! IPv6 ADDRESS FAMILY
!

address-family ipv6
 neighbor INTERNALv6 activate
 neighbor INTERNALv6 soft-reconfiguration inbound
 neighbor eNAMEv6 activate
 neighbor eNAMEv6 soft-reconfiguration inbound
 neighbor eNAMEv6 prefix-list IPv6-BOGON in
 neighbor eNAMEv6 prefix-list IPv6-BOGON out
 no auto-summary
 no synchronization
 bgp dampening
 network 2001:1348::/32
 exit-address-family
```