IPv4 addressing and routing plan for CLARA network

CLARA Network Engineering Group

May 2006

This document presents the NEG proposal for IPv4 addressing and routing plan for the backbone and PoPs of CLARA network.

CLARA/NEG/2004-0001
Version-1.4

VERSION MANAGEMENT

This document outlines the IPv4 addressing and routing plan for the CLARA network backbone. 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	14-Jul-2004	Eriko Porto
1.0	Corrections and changes	02-Aug-2004	Eriko Porto
1.1	Corrections and changes	16-Aug-2004	Eriko Porto
1.2	Corrections and changes	24-Aug-2004	Eriko Porto
1.3	Change in the topology map	06-Dec-2004	Eriko Porto
1.4	Change in Loopback-1 interfaces address	01-May-2006	Eriko Porto

Summary

1. Introduction

This document outlines the addressing and routing plan for the CLARA network.

In the following text, we will give a brief description of the CLARA network and presents its topology, followed by the description of the IPv4 address allocation, within the block obtained from LACNIC, to the backbone links and PoPs of the CLARA network.

The last sections of the document describe the routing protocols to be adopted for both intra-AS and inter-AS routing, and the routing policy to be used in the backbone.

2. Overview of CLARA network

The CLARA organization – Latin American Cooperation of Advanced Networks – 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 from the LA-NRENs will have access to the backbone through one of the CLARA PoPs in the ring.

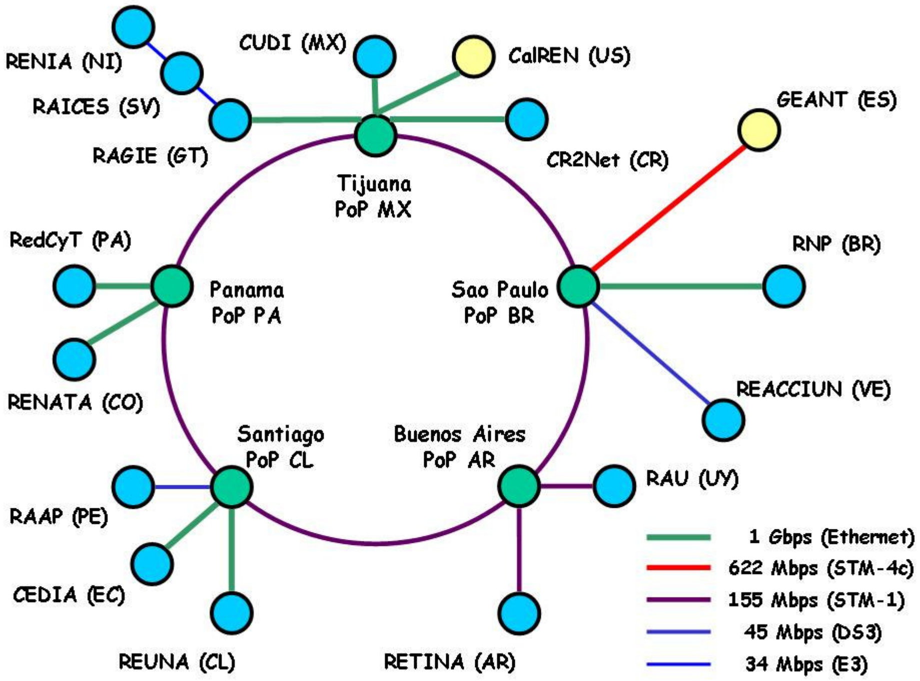

Figure 1: Backbone of CLARA network

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

3. IPv4 addressing

The following IPv4 address block is assigned for CLARA:

- Address block – 200.0.204/22
- Entity which received the Block – 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 IPv4 addresses distribution.

Table 1: IPv4 addresses distribution

address space	subnet type	Hosts	usage
200.0.204.0/22			RedCLARA assigned
200.0.204.0/24			
200.0.204.0/26	/30	2	Point-to-point links between PoP routers
200.0.204.64/26	/26	62	Future use
200.0.204.128/26	/30	2	Point-to-point links between routers and LA-NRENs
200.0.204.192/26	/26	62	Future use
200.0.205.0/24			
200.0.205.0/26	/32	1	Routers loopback interfaces
200.0.205.64/26	/26	62	Future use
200.0.205.128/26	/26	62	Future use
200.0.205.192/26	/26	62	Future use
200.0.206.0/24			
200.0.206.0/27	/27	30	Tijuana (MX) PoP Ethernet LAN
200.0.206.32/27	/27	30	Future use
200.0.206.64/27	/27	30	Future use
200.0.206.96/27	/27	30	Future use
200.0.206.128/28	/28	14	Sao Paulo (BR) PoP Ethernet LAN
200.0.206.144/28	/28	14	Future use
200.0.206.160/28	/28	14	Buenos Aires (AR) PoP Ethernet LAN
200.0.206.176/28	/28	14	Future use
200.0.206.192/28	/28	14	Santiago (CL) PoP Ethernet LAN
200.0.206.208/28	/28	14	Future use
200.0.206.224/28	/28	14	Panama (PA) PoP Ethernet LAN
200.0.206.240/28	/28	14	Future use
200.0.207.0/24		254	Reserved

It is important to notice that it is not the intention of CLARA NEG to lease blocks of IPv4 addresses to LA-NRENs or clients of LA-NRENs. The IPv4 addresses listed here are to be used exclusively at CLARA network backbone and PoPs. CLARA NEG is presuming that all clients of the network have their own address blocks and AS number obtained from a proper international organization of IP registration.

CLARA organization will not be primarily subletting any of its available prefixes for neither national networks nor clients. The CLARA network must be seeing as an interconnection network for other networks with its own self administration and policy regarding IPv4 address space and allocation.

4.
Interior Gateway Protocol

4.1.
Introduction

In order to keep up-to-date routing information about internal links inside CLARA network, an Interior Gateway routing Protocol (IGP) must be deployed. There are two main IGP routing protocols being used on today wide-area ISP networks: OSPF (Open Shortest Path First) and IS-IS (Intermediate System to Intermediate System).

These protocols have many similarities including the use of link-state algorithm, use of SPF algorithm to compute best paths, existence of costs associated to links and hierarchical design. Based on these considerations, both could be applied on CLARA network.

The IGP to be adopted in CLARA network is IS-IS based on some features of the protocol that turns it best suitable for the scenario of CLARA network operation. The CLARA network has to be prepared from the beginning of its deployment to scale for many advanced IP services that will be available for the LA-NRENs. In order to achieve this goal and provide a better platform to the integration of these services, the appropriate choice for the backbone operation is the IS-IS routing protocol.

The version of IS-IS that will be running in the backbone is the Integrated IS-IS, which was developed by the IS-IS working group of the Internet Engineering Task Force (IETF). The specification for Integrated IS-IS is described in RFC 1195. The Integrated IS-IS Routing Protocol may be used as an IGP to support IP as well as OSI. This allows a single routing protocol to be used to support pure IP environments, pure OSI environments, and dual environments.

4.2.
IS-IS deployment

The IS-IS will be deployed as a single Level 2 area for the five main nodes of the backbone, and each one will have a Level 2 link-state database with all the information for intra-area routing.

4.2.1. NSAP Addresses

NSAP is the network-layer address for CLNS packets. An NSAP describes an attachment to a particular service at the network layer of a node, similar to the combination of IP destination address and IP protocol number in an IP packet. NSAP encoding and format are specified by ISO 8348/Ad2. An NSAP address (Figure 2) has two major parts: the initial domain part (IDP) and the domain specific part (DSP). The IDP consists of a 1-byte authority and format identifier (AFI) and a variable-length initial domain identifier (IDI), and the DSP is a string of digits identifying a particular transport implementation of a specified AFI authority. Everything to the left of the system ID can be thought of as the area address of a network node.

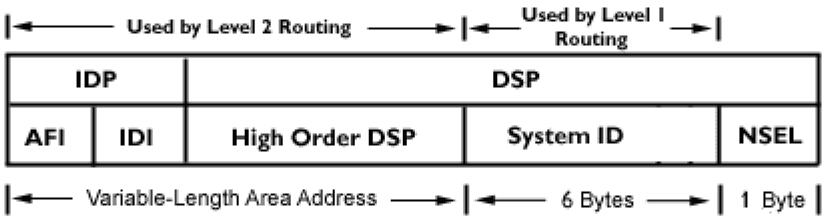

Figure 2: NSAP address

The routers will be configured with an AFI value of 49, which denotes private address space, like IP address space for private Internets as defined in RFC 1918. For the system ID part it will be used the loopback 0 interface IP address converted like the following example:

200.0.205.5 → 200.000.205.005 → 2000.0020.5005

The area ID will be set to 49.0205.

4.2.2. IPv4 Addresses

Figure 3 shows the routers interconnection diagram for the main nodes in the backbone of CLARA network.

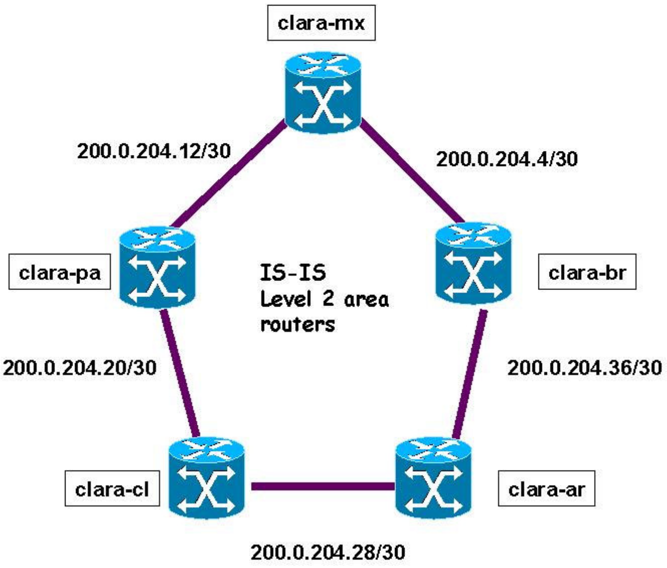

Figure 3: IS-IS topology

The IPv4 addresses that will be configured at the POS and loopback interfaces of the routers are listed in Table 2.

Table 2: IPv4 addresses for the interfaces of the main nodes

router	interface	destination	IP address
clara-br	pos-1 0	PoP – Tijuana	200.0.204.5/30
	pos-1 1	PoP – Buenos Aires	200.0.204.38/30
	loopback 0	N/A	200.0.205.5/32
	loopback 1	N/A	200.0.205.6/32
	pos-1 2	NREN RPN (BR)	200.0.204.129/30
clara-mx	pos-1 0	PoP – Panama	200.0.204.13/30
	pos-1 1	PoP – Sao Paulo	200.0.204.6/30
	loopback 0	N/A	200.0.205.13/32
	loopback 1	N/A	200.0.205.14/32
	giga-0 1	NREN CUDI (MX)	200.0.204.133/30
clara-pa	pos-1 0	PoP – Santiago	200.0.204.21/30
	pos-1 1	PoP – Tijuana	200.0.204.14/30
	loopback 0	N/A	200.0.205.21/32
	loopback 1	N/A	200.0.205.22/32
	giga-0 1	NREN RedCyT (PA)	200.0.204.137/30
clara-cl	pos-1 0	PoP – Buenos Aires	200.0.204.29/30
	pos-1 1	PoP – Panama	200.0.204.22/30
	loopback 0	N/A	200.0.205.29/32
	loopback 1	N/A	200.0.205.30/32
	giga-0 1	NREN REUNA (CL)	200.0.204.141/30
clara-ar	pos-1 0	PoP – Sao Paulo	200.0.204.37/30
	pos-1 1	PoP – Santiago	200.0.204.30/30
	loopback 0	N/A	200.0.205.37/32
	loopback 1	N/A	200.0.205.38/32
	atm-1 0	NREN RETINA (AR)	200.0.204.145/30

4.2.3. Metrics

The total cost to a destination is the sum of the costs (metrics) on all outgoing interfaces along a particular path from the source to the destination, and the least-cost paths are preferred.

Cisco IOS Software supports a 24-bit metric field, the so-called “wide metric”. Using this metric style, link metrics now have a maximum value of 16777215 (2^{24} - 1) with a total path metric of 4261412864 (254 x 2^{24}). Deploying IS-IS in the IP network with wide metrics is recommended to enable finer granularity and to support future applications such as Traffic Engineering, especially with high bandwidth links.

This extended metric will be configured in the routers using a default value since all links have the same cost. Eventually it should be necessary to adjust the costs for traffic engineering.

4.2.4. Redistribution

If it is necessary, redistribution from any other routing protocol, static configuration, or connected interfaces is allowed in Level 2 type of router. By default the metric type will be set as internal, which means that the metric of the route will compete with all other internal routes.

Metric type may be set to external. In that way the prefix will have a metric equal to the cost specified in the redistribution command plus a value of 64.5. Although the metric is increased if the metric is flagged as external on redistribution, the internal/external bit used to increase the metric is actually ignored when calculating routes unless the use of external metric is specified in the configuration. It is recommended that the metric type internal always be used when redistributing. Wide-metric TLVs do not use and discriminate between internal and external.

The links between LA-NRENs and the CLARA routers will run in passive mode as a way to inject the connected route to the IGP and simultaneously avoid undesirable adjacencies. In general redistribution of networks from other routing protocols to the IS-IS database will be strongly avoided.

4.2.5. Authentication Passwords

The IS-IS protocol, as specified in ISO 10589, provides for the authentication of LSPs through the inclusion of authentication information, as part of the LSP. Routers that want to become neighbors must exchange the same password for whatever level of authentication they are configured. A router not in possession of the appropriate password is prohibited from participating in the corresponding function (that is, it may not initialize a link, be a member of an area, or be a member of a Level 2 domain, respectively). Simple passwords are supported and will be implemented in the IGP domain.

5.
Exterior Gateway Protocol

All peering sessions between CLARA network and other ASes will be 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.

External peering sessions will be established only between CLARA network (ASN 27750) and two type of Autonomous Systems: *customer ASes*, formed by all CLARA LA-NRENs connected to CLARA network, and *upstream providers*, which actually includes GÉANT backbone and possibly Abilene in the near future. No other peering sessions will be established apart from these types of AS.

According to its devised purpose as an academic and research network for interconnection of LA-NRENs, CLARA network will not give transit between the LA-NRENs and the commodity Internet.

All prefixes announced from LA-NRENs to CLARA network will be subject to filtering policies to be applied on external BGP peering sessions. These policies can be grouped in two main classes, inbound and outbound policies, which will be described in the next sections.

5.1.
Inbound Filtering Policies

The main inbound policy covers acceptance of LA-NRENs prefixes. Each CLARA LA-NREN must declare all prefixes that need to get transit on CLARA network to the CLARA Network Engineering Group (NEG). Additionally, their prefixes must be registered in the Routing Assets Database (RADb). Only those prefixes will be accepted in the announcements coming from CLARA LA-NRENs and all other will be discarded.

The BGP AS-PATH attribute will also be considered by the filtering policies. CLARA LA-NRENs must report to the CLARA NEG their own AS number and other ASNs they give transit to. Only prefixes originated from these ASes will be accepted.

Since the above inbound policies are restrictive in nature, other common filtering practices like denying bogus routes and limiting the maximum prefix that can be accepted will not be required.

On the other hand, prefixes announced from GÉANT backbone will be accepted in general and only inbound route filtering for bogus routes will be applied.

5.2.
Outbound Filtering Policies

Once inbound filtering policies were correctly applied, there is almost no action to be taken regarding outbound prefix filtering. Thus, all BGP prefixes accepted by the CLARA network will be announced to all peering sessions with LA-NRENs and GÉANT. In any case, only ordinary route filtering for bogus routes will be implemented at the peering session.

Outbound announcements from CLARA network will only be modified through the specific BGP communities allowed to be used by LA-NRENs. BGP communities will be covered in the section 5.4.

5.3.
Network Masks

Following the common practice on international academic networks, no filtering will be done regarding netmask length on prefixes.

Therefore, LA-NRENs will be allowed to announce prefixes with generic netmasks to the CLARA network, including netmasks stricter then /24. This adds flexibility on what prefixes LA-NRENs want to get routed at CLARA network and its upstream providers. Similarly, LA-NRENs should expect announcements from CLARA network with generic netmasks up to /32.

Note: In general we don’t expect to have many of such announcements. However based in the RNP experience in administering a backbone for interconnection of academic networks, such networks involving a few hosts (usually a department laboratory of some university) may have to be announced in order to carry on some research or similar related work.

5.4.
BGP Communities

All prefixes received from CLARA LA-NRENs and GÉANT will be distributed with the BGP community attribute set for internal use within CLARA network, mainly for administration and troubleshooting of the routing policy.

Furthermore, special values for the community attribute will be available to LA-NRENs in order to allow them to influence the announcement of their prefixes outside CLARA network. LA-NRENs will be able to stop announcement of a prefix to other LA-NRENs or to GÉANT by attaching communities according to a BGP community policy (yet to be defined).

5.5.
Other BGP attributes and features

Internally, all prefixes will have the same associated Local-Preference attribute. If a LA-NREN becomes multi-homed with CLARA, new BGP communities will be provided in order to allow changes in Local Preference attribute of their prefixes within CLARA network, allowing load balancing between links.

In this situation, the Multi-Exit Discriminator (MED) attribute can also be used. MED will be accepted without modification on all CLARA network peering sessions.

BGP route dampening will be active on all routers in order to avoid BGP route flapping to be propagated between peers connected to CLARA network.

6.
Glossary of Terms

To facilitate the reading and understanding of this document, the following list of acronyms is offered:

AFI	Authority and Format Identifier
AS	Autonomous System
ASes	Autonomous Systems
ASN	Autonomous System Number
BGP4	Border Gateway Protocol version 4
DSP	Domain Specific Part
eBGP	External BGP
EGP	Exterior Gateway Protocol
iBGP	Internal BGP
IDP	Initial Domain Part
IETF	Internet Engineering Task Force
IGP	Interior Gateway Protocol
IS-IS	Intermediate System to Intermediate System
ISO	International Organization for Standardization
LACNIC	Latin America and Caribbean Network Information Center
LA-NREN	Latin America NREN
LSP	Link-State Packet
MED	Multi-Exit Discriminator
NEG	Network Engineering Group
NOC	Network Operation Center
NREN	National Research and Educational Network
NSAP	Network Service Access Point
OSI	Open Systems Interconnection
OSPF	Open Shortest Path First
PoP	Point of Presence
POS	Packet Over SONET
RADb	Routing Assets Database
RFC	Request For Comments
TLV	Type/Length/Value