

Technical Requirements Policy for IX.br - V1.0

An **Internet Exchange Point (IXP or IX)** is a network solution typically consisting of switches and routers operating at the layer 2 level of the ISO/OSI reference model, which offers a range of services for the interconnection of Autonomous Systems (AS), called Participants.

IX.br's goal is to promote the exchange of Internet traffic through a layer 2 switch, offering Ethernet access ports and using IP protocol, with neutrality and transparency.

Note that the exchange of traffic in an IX occurs in the set of all network facilities and access channels, be they for the connection between Interconnection Points (PIX or CIX) or between Participants. To ensure that this infrastructure operates smoothly, **restrictions are imposed on the type of traffic allowed and authority is given to IX.br's Network Operation Management (NOC) to temporarily disable Participants who do not comply with the rules or cause disturbances to the functioning of the network.**

All ports of the Participants' network equipment connected to access channels should use the BGP-4 protocol (Border Gateway Protocol 4 - RFC4271), which is the standard protocol used by all Autonomous Systems on the Internet to propagate and learn from their peers the routes of the networks of other Autonomous Systems connected to it in a connection known as BGP Session. It is through this exchange information that the Internet can be viewed as a single network or a mesh of networks. Each Participant may advertise a maximum of 20 routes, however this limit may be revised if properly justified.

In an IX it is not necessary that the network equipment of a Participant establish BGP sessions with all other Participants, since route servers that store the routes advertised by each participant are used, allowing that through a single BGP session established with it, the network equipment may get the routes advertised by all other Participants, saving resources overall. At IX.br, each location has at least two redundant route servers.

PIX and CIX definitions

- A Participant connects to IX.br through a PIX or a CIX, defined below. A series of rules are adopted in these interconnection points, which imply restrictions and the control of the resources offered by the network.
- **Interconnection Point (PIX):** is a network solution used to receive traffic from Participants or from Intermediate Interconnection Points (CIX). It can connect to other PIXs or to Central PIXs, forming the network solution that makes up the Internet Exchange Point (IX). In the PIXs, Participants use Ethernet ports, or several aggregated Ethernet ports, to join the network. A PIX connects to the Central IX PIXs through dedicated optical fibre, also called unlit or dark fibre, and houses IX.br network assets. There are specific rules for the operation of a PIX that are not in the scope of this document. This simplified definition is contained herein only to facilitate, by comparison, the conceptualisation of CIXs.
- **Intermediate Interconnection Point (CIX):** is a network solution used to aggregate traffic from more than one Participant, linking them to the IX via a shared access

channel in an Interconnection Point (PIX) or in a Central Interconnection Point (Central PIX). This delivery of several Participants is made through a type of port called a **shared port**, which is a set of one or more aggregated Ethernet ports (LAG), from the viewpoint of the PIX that hosts it. Each Participant transported by the CIX will have a **virtual port** on the IX. The CIXs are thus also part of the network solution that makes up an Internet Exchange Point (IX).

Classification of the CIXs

The CIXs are classified into three types:

Type 1 - Small: Operates with up to 20 VLANs (IEEE 802.1Q), connecting to a PIX via an optical link without redundancy, i.e. by a single route.

In this case the virtual ports are defined by a set of VLANs. Typically 2 VLANs, one for IPv4 MTA (Multilateral Traffic Agreement) and the other for IPv6 MTA. Typically, thus, a Type 1 CIX can connect up to 10 Participants to the IX, the number being reduced in case the Participants use bilateral VLANs.

Type 2 - Midsize: Operates with up to 100 VLANs (IEEE 802.1Q), and must be connected to a PIX by means of a redundant optical link, which means having two non-overlapping separate routes.

Such as with the Type 1 CIX, the virtual ports are defined by a set of VLANs for each Participant. Typically a Type 2 CIX can connect up to 50 participants to the IX, the number being reduced in case the Participants use bilateral VLANs.

Type 3 - Large: Operates with Metro TAGs, tags used for Q-in-Q (IEEE 802.1ad) implementation, which allow the aggregation of a larger number of Participants. Must be connected to a PIX by means of a redundant optical link, which means having two non-overlapping separate routes. The Metro TAGs will be allocated in blocks of 100 to CIX operators and their values will be defined alongside IX.br.

Unlike in Type 1 and Type 2, in the Type 3 CIX, a Participant's virtual port is defined by a Q-in-Q tunnel. Up to 10 VLANs are allowed in the same virtual port, which include the VLANs for participation in the IPv4 and IPv6 MTAs and bilateral VLANs.

Bandwidth Limitation on a CIX:

Upon reaching peak loads of 80% of the nominal capacity of the shared port on the PIX, no new VLANs will be commissioned on the CIX, either to new Participants or current Participants. At this point, or preferably before, preventively, the CIX should provide increased capacity to the shared port to which the PIX is connected.

If case the traffic, in any way, exceeds 95% of the nominal capacity of the PIX's shared port, IX.br's NOC can disable one or more Participants' virtual ports at its own discretion, in order to bring usage down to less than 80% of the nominal capacity.

Note that, for Type 2 and Type 3 CIXs which use optical fibres from the street, the nominal capacity of the shared port will be considered the capacity without redundancy. For example, if a Type 3 PIX has 6 10G aggregated ports with two redundant 30G optical links, the nominal capacity will be considered 30G, and not 60G.

CIXs hosted at the same data centre of the PIX to which it is interconnected:

For all types of CIX, if installed on the same data centre of the destination PIX, the optical link may be by means of optical cross-connects without redundancy.

Note that in the case of bandwidth limitation, in the case of the CIX hosted at the same data centre where the PIX that is, the total capacity of the aggregated ports is considered the nominal capacity, even for Type 2 and Type 3 CIXs. For example, for 6 10G cross-connected ports, the nominal capacity will be considered 60G.

INOC-DBA Extension:

INOC-DBA is an exclusive voIP network for the Autonomous Systems, the networks that form the Internet: it provides a hotline, a quick and easy way of communication between their NOCs (Network Operations Centres) and CSIRT (Computer Security Incident Response Teams). In INOC connections are made using ASN (Autonomous System Number).

The CIX operator's NOC must have an INOC-DBA extension always online. It will be used as the main contact medium by the IX.br team, if necessary.

It is recommended that IX.br Participants interconnected via CIX also get INOC-DBA extensions. To get an extension, the Participant must access: <http://inoc.nic.br>.

Data in Registro.br:

Registro.br is responsible for the management of IP and ASN numbers in the country. It maintains a database with registration information, contacts, blocks delegation and designation. This database can be consulted using *whois*.

The CIX operator, in case it is an AS, as well as the IX.br Participants connected via CIX must maintain their data up to date on the Registro.br database.

Internet numbering resources should be managed by the numbering resources management system, which can be accessed at:

<https://registro.br/cgi-bin/nicbr/login>

It's the same management interface of the ".br" domains. To connect, one needs to provide ID and the password of the Internet resource contact person. Once the system is accessed, there will be a list of all the IP blocks and ASNs which that user/ID can manage.

Technical standards and restrictions for CIX operators

- **Ethernet frame:** the entire infrastructure is based on Ethernet II standard (or "DIX Ethernet") the LLC/SNAP (802.2) encapsulation not being allowed. Thus, traffic exchange is made by forwarding frames in the Ethernet II format from one point of the infrastructure to another.

- **Ethertypes:** only frames with the type specification field (Ether type) equal to one of the types below may go through the infrastructure:

- 0x8000 - IPv4
- 0x0806 - ARP
- 0x86DD - IPv6
- 0x8809 - LAG (802.3ad)

- **MTU equal to 1523 for the CIX operator:** to operate CIX using Q-in-Q, the 1523 bytes MTU will be default, guaranteeing the Participant 1500 bytes of data, and the rest for signalling.

- **Capability to withstand a minimum amount of MACs per VLAN:**

- in São Paulo, 2048 MACs per Participant VLAN
- in Rio de Janeiro, 512 MACs per Participant VLAN
- elsewhere, 256 MACs per Participant VLAN

- **Q-in-Q (802.1ad) support capability** for Type 3 CIX

Technical standards and restrictions for IX.br Participants

- **Ethernet framework:** the entire infrastructure is based on Ethernet II standard (or "DIX Ethernet") the LLC/SNAP (802.2) encapsulation not being allowed. Thus, traffic exchange is made by forwarding frames in the Ethernet II format from one point of the infrastructure to another.

- **Ethertypes:** only frames with the type specification field (Ether type) equal to one of the types below may go through the infrastructure:

- 0x8000 - IPv4
- 0x0806 - ARP
- 0x86DD - IPv6

- **One physical address (MAC) per VLAN:** frames sent on a certain VLAN should all have the same MAC address.

- **Number of MAC addresses per IX.br Participant:** each participant should use a single MAC physical address per network equipment, typically a router, and a maximum of two network equipment. That is, IX.br filters MAC addresses in each CIX Participant's virtual port, and a maximum of 2 physical addresses may be registered.

- **No ARP proxy:** the use of the ARP proxy is not allowed in any network equipment connected to IX.br's physical or virtual ports.
- **Neighbour discovery protocols may not be enabled:** protocols such as CDN, MNDP, etc. are not allowed on the interface of the network equipment connected to IX.br.
- **Router advertisement (RA):** The IPv6 RA must be disabled on the interface of the network equipment connected to IX.br.
- **Only unicast physical addresses:** In a CIX virtual port, only frames routed to unicast physical addresses are allowed. These two cases are exceptions:
 - Broadcast ARP packets
 - ICMPv6 multicast Neighbour Discovery packets (ND). Please note that this does not include Router Solicitation (RS) or Router Advertisement (RA) packets.
- **No broadcast or multicast:** It is no allowed to send IPv4 broadcast or multicast or IPv6 multicast packets in IX.br, except for ICMPv6 ND packets.
- **No disclosure of the IX.br network addresses:** the network address space of each IX.br location, i.e. the addresses used on the ports of routers connected to IX, must not be advertised to other networks. It is also recommended that these addresses not be announced internally in the Participants network, which implies the use of next-hop-self for the internal advertisement of routes learned via IX.
- **MTU equal to 1523 for the CIX operator:** to operate CIX using Q-in-Q, the 1523 bytes MTU will be default, guaranteeing the Participant 1500 bytes of data, and the rest for signalling.
- **1500 MTU for CIX Participant:** equipment and circuits used by a Participant interconnected via CIX must be capable of operating with a 1500 bytes MTU.

Traffic exchange within the same AS is forbidden: For a Participant connected to IX.br through more than one port, be they virtual ports in CIXs or ports in PIX, it is forbidden to exchange local AS traffic using the IX.br structure.

Obs. 1: A CIX operator can also be a Participant in the same CIX in which it operates, in which case the technical standards and restrictions for Participants also apply.

Obs. 2: It is recommended that the CIX operator inform their customers about the technical standards and constraints as specified in this document which will apply to them as Participants.