

IXP Technical Requirements OIX-1

Version 3.1, June 16, 2025

Questions: ix-group@oix.org

This standard defines the technical requirements for an IXP to be certified by OIX. The purpose of the requirements is to provide publicly available information on what the participants of a certified IXP can expect, and not to describe in detail how the IXP is designed, built or operated. IXPs can comply with the OIX standards to serve different communities with different requirements, from a small single datacenter IXPs serving a local community to a large IXP located in multiple datacenters in a metro area.

The keywords used throughout the document are as defined in RFC 2119.

The IXP **SHOULD** follow the Best Current Operational Practices for an Internet Exchange, posted at <https://github.com/Open-IX/BCOP>

Table of Contents

IXP Technical Requirements OIX-1	1
Table of Contents.....	2
Forward.....	3
Definition of Internet Exchange - IX.....	4
Definition of an Internet Exchange Point - IXP.....	4
Definition of Community Supported IXP.....	6
Services.....	7
Minimal Service Offering.....	7
Public Exchange VLAN (IX).....	7
Additional Service Offering.....	7
Private VLAN (PVLAN).....	7
Physical Interface.....	7
Traffic Forwarding / Fabric Protection.....	8
Layer 2.....	8
Layer 3.....	8
Customer Interface.....	8
Conforming to MANRS.....	9
Infrastructure.....	10
Switching Platform.....	10
IP Address Space.....	10
Route Server.....	11
Topology.....	12
Operations.....	13
NOC.....	13
Monitoring.....	13
Statistics.....	13
Website.....	13
Pricing.....	14
Miscellaneous.....	14
Appendix.....	15
Sample Layer 2 Filtering.....	15
Sample Layer 4 BGP Filtering.....	16
Sample Layer 4 BGP Culling Filter.....	16
Sample IXP Port Configuration.....	17

Forward

This update is to provide a single OIX-1 Specification for “Regular” and “Community Supported” IXP submissions and to also include some sample configurations for IXP Fabric Switches.

Please feel free to comment and make any suggestions.

Thanks,

Paul Emmons

ix-group@oix.org
Chair IX Committee
June 16, 2025

Definition of Internet Exchange - IX

An IX is the entity that operates one or more switching fabrics at one or more locations for the purpose of peering (interconnection).

Definition of an Internet Exchange Point - IXP

From IX-F at <https://www.ix-f.net/ixp-definition.html> - An Internet Exchange Point (IXP) is a network facility that enables the interconnection of more than two independent Autonomous Systems, primarily for the purpose of facilitating the exchange of Internet traffic.

An IXP provides interconnection only for Autonomous Systems.

An IXP does not require the Internet traffic passing between any pair of participating Autonomous Systems to pass through any third Autonomous System, nor does it alter or otherwise interfere with such traffic.

“Autonomous Systems” has the meaning given in BCP6/RFC4271 , “A Border Gateway Protocol BGP4”.

“Independent” means Autonomous Systems that are operated by organizational entities with separate legal personality.

Explanatory Notes

1. An Internet Exchange Point is a technical facility. This is distinct from the organization that provides that facility, which might be termed an *IXP operator*.

2. An IXP is distinct from an Internet access network or a transit network/carrier.
3. The function of an IXP is to interconnect networks. An IXP does not provide network access or act as a transit provider/carrier. An IXP also does not provide other services unrelated to interconnection (although this does not preclude an IXP operator from also providing unrelated services).
4. An IXP exists to interconnect networks that are technically and organizationally separate.
 - a. Without qualification the term “network” is too flexible and fails to identify the degree or kind of separation required. Once interconnected, separate networks are arguably part of the same network: the entire Internet is often considered a network of networks.
 - b. To resolve this terminological problem we employ the term “Autonomous System”, which is the standard technical definition of a technically stand-alone network.
5. The network operators whose networks are interconnected in an IXP are sometimes collectively termed “*IXP participants*”, which generalizes the relationship between those entities and the IXP operator; IXP participants **MAY** be members of the IXP operator, customers of the IXP operator, or some other relationship.
6. An IXP is a facility where numerous participants interconnect (at least three); this distinguishes Internet Exchanges from bilateral network interconnection, in which one network connects to one another.

We note that within general practice IX and IXP are largely interchangeable.

To be certified an IXP **MUST** be operating with at least 3 globally routable ASes.

The IXP **SHOULD** also be listed in a publicly available resource Such as IXPDB, PeeringDB or similar

The IXP’s ASN and IP resources **MUST** be verifiable with the RIR and assigned to the IXP

If an IXP discontinued its operations or has any material changes it **MUST** notify OIX within 30 days.

Definition of Community Supported IXP

To achieve “Community Supported” Status for this certification:

1. Is not owned or controlled, operated or managed by a colocation provider or service provider; and
2. Has the appropriate government recognition as an organized entity, Community Association, Unincorporated Association, Trade Group or similar organization such that the organization benefits the member / participant “Community” (hereinafter “Community Supported”); and
3. Has either:
 - a. no recurring charges for connected networks (pass-through / cost-recovery charges exempted); or
 - b. a maximum yearly revenue of \$100,000 USD (or local equivalent) as updated by the OIX Board.
4. A community IXP shall recertify on an annual basis

Services

Minimal Service Offering

The IXP **MUST** provide the minimum services described below. This also allows the IXP operator to provide additional services, or methods of interconnection.

Public Exchange VLAN (IX)

A switch platform which allows any-to-any interconnection. Customer interfaces with Ethernet frames tagged for the public exchange VLAN **MUST** be forwarded in accordance with the traffic rules indicated in this document.

Additional Service Offering

The IXP **MAY** provide additional services, as long as they are described on a publicly available website of the IXP.

Private VLAN (PVLAN)

A private switch platform, whereby any two or more parties may consent to interconnect through either the same physical port that delivers their access to the Public Exchange VLAN or alternatively dedicated physical port(s). If a PVLAN service is offered, in case there are exactly two parties in the private VLAN the connection **MUST** be delivered guaranteed congestion free. In case of more than two parties the service **MAY** be provided on a best efforts basis.

Physical Interface

The IXP **MUST** offer IEEE 802.3 Ethernet connectivity on a common switch infrastructure. Service offerings **MAY** be available at any IEEE defined rate, including IEEE 802.3ad or IEEE 802.1AX link aggregation of any of these rates.

The complete service offering **MUST** be described on a publicly available website. The information provided **MUST** contain: link rate and physical media (copper, fiber and fiber type). The information **SHOULD** describe how each port type is connected to the fabric.

The IXP **SHOULD NOT** allow for physical port speed and duplex negotiation.

The IXP **MUST** protect the IXP fabric from user misconfiguration, broadcast / unknown and multicast storms, and port flapping.

Traffic Forwarding / Fabric Protection

Layer 2

The IXP **MUST** forward frames with the following Ethertypes:

- 0x0800 IPv4
- 0x86dd IPv6

Valid frames with Ethertype 0x86dd may be suppressed on the Public Exchange VLAN using snooping, or alternate methods used to implement IPv6 Neighbor Discovery.

If there is no provision to handle ARP in any other way, the IXP **MUST** forward frames with the following Ethertype:

- 0x0806 ARP

If the IXP has reason to limit certain traffic, the IXP **MUST** publish on a publicly available website what traffic is not allowed and or not forwarded on the exchange platform.

If the IXP applies a MAC address locking mechanism on a participants port, then the IXP **MUST** make known to customers the process to update MAC addresses.

Layer 3

The IXP **SHOULD** provide a method to ensure that BGP Peering Sessions only originate from the IPv4 and IPv6 Ranges utilized by the IXP.

Access Lists

Per interface access lists as described in BCP214 **SHOULD** be implemented on each member / participant facing port.

Customer Interface

The IXP **MUST** provide a clear demarcation point between the IXP services and the customer. This can be either directly on the exchange or via a common demarcation point available to the participants.

Conforming to MANRS

The IXP **SHOULD** participate in Mutually Agreed Norms for Routing Security (MANRS) <https://www.manrs.org/> and **SHOULD** encourage its members / participants to participate in MANRS.

Infrastructure

Switching Platform

The IXP switching platform **MUST** have backplane capacity to sufficiently handle the aggregate traffic of all customer facing ports, without oversubscription. If individual switching elements contain multiple switch fabric modules, the same conditions **MUST** apply during single component failures.

The IXP **MUST** run any inter-switch links congestion free.

The IXP **SHOULD** have redundant power feeds fed from discrete sources (A and B) for all exchange infrastructure. If the IXP does not have redundant power feeds on any components, it **MUST** describe where not on a publicly available website.

If the IXP does not have full path diversity between two discrete switching elements in different physical locations, this **MUST** be described on the IXPs publicly available website.

The IXP **SHOULD** describe on a publicly available website the switching platform and the redundancy measures implemented to overcome single component failures.

IP Address Space

In order to be independent of any of the connected parties, the IP space used on the “Public Exchange VLAN” **MUST** be Provider Independent space or other IP space directly assigned by a RIR for the purpose of operating an IXP. This applies to both IPv4 and IPv6. The IXP operator is responsible for obtaining address space from the respective RIR, as well as providing all material for justification, documentation, and applicable fees as required by the RIR.

Route Server

If a route server service is offered then it **MUST** support both IPv4 and IPv6, and 16-bit and 32-bit ASNs. The AS number used for the route server implementation **MUST** be a unique AS number assigned by one of the RIRs. The AS **MAY** be utilized by multiple IX fabrics operated by the same entity. The IXP **SHOULD** use a 16 bit ASN.

The Route Server **will** operate in such a way:

1. That the Route Server AS Number **MUST** not be shown in the AS routing path.
2. That the NEXT-HOP ip address and AS is the originating member / participant.
3. That rather than honoring the NO-EXPORT community, it **MUST** pass that community onto the connected member / participant.
4. That the originated IP announcements:
 - a. **MUST** be filtered for reserved or special IP address ranges and AS numbers.
 - b. **MUST** limit the number of BGP prefixes each member / participant may originate.
 - c. **MUST** disallow the propagation of a default route.
 - d. **SHOULD** be filtered against the originating networks IRR entries.
 - e. **SHOULD** be filtered against invalid RPKI origination.
 - f. **SHOULD** honor the request of networks not to have announcements in the route servers that the originating networks do not intend to have on route servers (e.g. PeeringDB's "Never-via-Route-Server").
 - g. **SHOULD** ensure that paths with well known transit networks are not propagated.
 - h. **SHOULD** limit the use of excessive AS path length.
 - i. **SHOULD** disallow routes that are:
 - i. Shorter than /8 for IPv4 and /19 for IPv6.
 - ii. Longer than /24 for IPv4 and /48 for IPv6. except for /32 or /128 for Black Holing purposes.

The IXP **SHOULD** have at least two route servers. The route servers **SHOULD** be in diverse locations, and connected to discrete switching elements. If the IXP does not have route servers in diverse locations, or does not have two route servers connected to discrete switching elements, it **MUST** be described on the IXP's publicly available website.

The IXP **SHOULD** have redundant power feeds fed from discrete sources (A and B) for the Route Servers. If the Route Servers do not have redundant power feeds, it **MUST** describe where not on a publicly available website.

The IXP **SHOULD** publish the Route Server setup on a publically website. If not then the Route Server setup **MUST** be made available to members / participants.

Route Servers filtering policies **MUST** be disclosed to its members / participants.

A looking glass **SHOULD** be made available to the members / participants to verify announcements.

Topology

If the IXP operates at more than one location, the IXP **MUST** indicate what paths are redundant and **SHOULD** indicate capacity of the links.

Operations

NOC

The IXP **MUST** publish a telephone number, email address or any other means that provides immediate access to technical support, on a website available to its participants, on how to contact operational staff that is capable of managing the IXP infrastructure. The access method **MUST** be available 24x7, note this does not mean staff needs to be available 24x7, but the IXP **MUST** publish staff hours.

The IXP **MUST** provide and publish a procedure to announce service affecting maintenance to its participants.

Monitoring

The IXP **MUST** monitor the exchange platform for performance degradation and service affecting events.

The IXP **MUST** provide a procedure to inform its participants on performance degradation and service affecting events.

Statistics

The IXP **MUST** publish on a publicly available website the participants on the peering platform and the relevant AS numbers.

The IXP **MUST** publish on a publicly available website the total sum of all incoming and outgoing traffic in bps from all connected networks on the public peering VLAN. The traffic sum **MUST** include the traffic on customer facing ports only and **MUST** be made up of 5 min average traffic measurements. A distinction **MUST** be made between the traffic on the public peering VLAN and any other interconnection service.

Website

The IXP **MUST** have available and maintain a publicly available website where at least the subjects mentioned in this document **MUST** be addressed.

Pricing

The IXP **MUST** disclose any pricing, terms and conditions (if any) on its website.

Miscellaneous

The IXP **MUST** publish and maintain an accurate entry for a peering contact and configuration directory such as <https://www.peeringdb.com> or <https://ixpdb.euro-ix.net/> or any regional IXP organizations. This entry **MUST** contain a list of all facilities that the IXP maintains a point of presence.