

NWAX Usage Guidelines

Send questions to info@nwax.net

About the exchange:

- Peering on the exchange can be via the route servers or bilateral or both (encouraged).
- There are only three ethertypes allowed to traverse the IX: 0x0800 (IPv4), 0x0806 (ARP), and 0x86dd (IPv6).
- Non-unicast traffic is not allowed except broadcast ARP and multicast ICMPv6 Neighbor Discovery packets. Per-neighbor timeouts that result in flooded (broadcast/multicast) packets should be set to 4 hours or as close to that as possible in the case of vendor limitations. Short timeouts may result in quarantine.
- BPDU guard is enabled on the fabric; if BPDUs are detected on member ports towards the fabric the port will temporarily be put into ERRDISABLE.
- The internet exchange will send BPDUs toward member ports. Please filter/drop/ignore them.
- The internet exchange will send LLDP toward member ports to aid in troubleshooting. Members are welcome to send LLDP toward the exchange.

Important Connection Settings

NWAX ports are all by default set to auto-negotiation since all port speeds NWAX offers (1g+) are full-duplex only. If you are connecting with multiple ports, we will have bonded them together into a LACP channel unless otherwise requested.

Please note that NWAX leverages dynamic MAC address filtering on member ports and only allows a single MAC address per member port. If your port sends frames towards the fabric from more than one MAC address, the NWAX switch will auto-disable your port. The NWAX switches are set to auto-reenable the port after five minutes. If you are seeing your port flap up and down once every five minutes, this is one of two likely causes.

NWAX also utilizes spanning-tree as a loop prevention mechanism. The NWAX core switches send BPDUs toward member ports, but if they receive *any* BPDUs from a member port they will auto-disable the port. After five minutes the port will be auto-reenabled. If you are seeing your port bouncing on five minute intervals, it is likely either due to BPDU's or MAC address violations.

Each new member gets 1 IPv4 address. If a member wants to connect two routers, a second IPv4 address can be assigned. Each port or bundle of ports may only be connected to one router (i.e., only one MAC address allowed per port). Anything beyond two IPs will be considered with justification on a case-by-case basis.

Member Portal

NWAX provides a portal for members which provides traffic graphs for your ports, peer to peer graphs (how much traffic you are sending and receiving from each member), the contact and peering details of all other members, a Peering Manager tool, documentation, support information and much more.

The NWAX IXP Manager is available at: <https://portal.nwax.net>

Participants must:

- Use BGP-4 or its successor and must set NEXT_HOP_SELF if advertising routes from other NWAX participants
- Present only one MAC address
- For IPv4 a participant's router must be configured to receive and respond to ARP packets from all NWAX participants, even those that are not direct peers.
- For IPv6, participant routers must receive and respond to ICMPv6 neighbor solicitation packets from both fe80::/10 and all NWAX participant addresses, including those that are not direct peers, directed toward fe80::/10, ff02::1:ff00:0/104, and the participant's unicast NWAX assignments
- Have at least one contact on the mailing list. To join the member mailing list please send an email to: members+subscribe@nwax.net
- Be responsive to NWAX administrators and other participants. This includes keeping contact records up to date in PeeringDB and IXP Manager.
- Have redundant and diverse external network paths in case of an issue with their paths to or through NWAX.
- Register prefixes in IRR in order for routes to be accepted by NWAX route servers.

Participants must not:

- Point default or otherwise use another participant's resources without permission
- Use ACLs that violate neighbor discovery norms. This is to prevent excess flooded packets on the fabric.
- Allow NWAX subnets to propagate externally from their network and should minimize internal propagation. If a participant's network beyond their NWAX edge router(s) can reach the NWAX subnet addresses, participants must use ACLs to prevent this.
- Sniff traffic between other participants.
- Use the IX as transport between sites.

Guidelines/helps:

- Use the route servers to simplify administration. Bilateral sessions are also encouraged.
- Try to conserve IPv4 addresses to delay the need to renumber. Each member gets one IPv4 address by default. A second port or LAG can have a second address. Additional IPv4 addresses will be issued at the NOC's discretion with justification.
- A jumbo VLAN is available. Inquire with noc@nwax.net if you want to use it.
- Email noc@nwax.net if you have any technical questions.
- IRR tutorial provided by the SIX <https://www.seattleix.net/irr-tutorial>
- MAC address changes do not require coordination. We have dynamic port security enabled which will learn whichever MAC you're using when the interface comes up.
- NLNOG reference for filtering: <https://bgpfilterguide.nlnog.net/>

Consequences of not following the rules:

- NWAX administrators may shut down the port/route server peering sessions of any participant not adhering to the rules of the exchange at their discretion. Participants may have their port(s) enabled again by curing the infraction.
- NWAX reserves the right to deny access to the exchange at its sole discretion.

Peering

NWAX does not mandate any member to peer with any other member. It is up to each member to determine their peering strategy (though we are hoping to facilitate as many connections as possible).

You will find a full list of members on the IXP Manager portal web site, along with the correct email addresses to use for peering requests. Information may also be available in PeeringDB.

When emailing other NWAX members about peering requests, please include all technical details relevant to the peering session, including your IP address, your AS number, and an estimate of the number of prefixes you intend to announce to that candidate peer.

The My Peering Manager tool in the IXP Manager will compose mails with the above details for you automatically.

About the route servers:

- Your organization must have PeeringDB and IRR records. Your PeeringDB "IRR Record" should contain simply your IRR as-set name.

- There are two route servers, and you need to establish sessions with both so that you are not affected by route server maintenance.
- Max prefix for IPv4 and IPv6 is set by the NOC. It can be increased upon request and validation. For your peering, if you use a max prefix with the route server, we recommend 200,000 for IPv4 routes and 70,000 for IPv6 routes.
- The AS-SET representing all participants at NWAX is AS-NWAX.
- The route servers update from the IRR once per hour.
- Strict filtering is performed using Internet Routing Registry (IRR) data.
- We do not use RPKI currently.
- The route servers are configured to be BGP-passive thus your side needs to be BGP-active.
- In addition to peering through the route servers, you may wish to peer bilaterally with organizations that are important to your organization.
- MD5 BGP passwords are available on request.
- Route server peering configuration examples are shown here:
<https://www.seattleix.net/route-servers>
- The NWAX route server ASN is 63028.
- IPv4 routes shorter than /8 are filtered and IPv6 routes shorter than /19 are filtered. Bogon ASNs and martian prefixes are filtered.
- The smallest supported prefixes are /24 for IPv4 and /48 for IPv6.
- Our IPv6 peering prefix is 2620:124:2000::/64
- Our IPv4 peering prefix is 198.32.195.0/24

How to connect to the route servers:

- Refer to our [PeeringDB entry](#).
- Log in to IXP Manager at portal.nwax.net and get the MD5 session password, if applicable
- Establish a peering session with .251, .252 and ::251, ::252 (AS63028)
- Make sure that your BGP session has enforce-first-as set to disabled
 - Cisco IOS/IOS XE: no bgp enforce-first-as
 - Cisco IOS-XR: bgp enforce-first-as disable (this can be done per group)
 - JUNOS: JUNOS does this by default.

Router Configuration

If you are new to internet exchanges, we would ask you to note that all members are expected to adhere to the NWAX technical requirements.

For Cisco IOS based routers, we recommend the following interface configuration commands:

```
no ip redirects
```

```
no ip proxy-arp
no ip directed-broadcast
no mop enabled
no cdp enable
udld port disable
```

If you intend to use IPv6 with a Cisco IOS based router, please also consider the following interface commands:

```
no ipv6 redirects
ipv6 nd suppress-ra
```

Connecting Switches to NWAX

Many members choose to connect their NWAX port to a layer 2 switch and then forward their peering traffic to a router virtual interface hosted elsewhere on their network. While connecting layer 2 switches to the NWAX peering LAN is permitted, incorrect configuration can cause serious and unexpected connectivity problems.

The primary concern is to ensure that only traffic from the router subinterface is presented to the NWAX port. NWAX implements per port mac address counting: if more than 1 mac address is seen on any member switch port at any time, that port will automatically be disabled for a five minute cooling off period, and your connectivity to NWAX will temporarily be lost.

This policy prevents two potential problems: firstly, it ensures that layer 2 traffic loops are prevented and secondly, it ensures that no other traffic escapes to the NWAX peering LAN which shouldn't be seen there.

If you choose to connect your NWAX port or ports to a switch, it is critically important to assign one unique VLAN for each NWAX connection. If you share an NWAX facing VLAN between multiple NWAX ports or share a NWAX-facing VLAN with any other network, your connection will automatically be shut down due to the security mechanisms implemented by NWAX.

Please also note that by default, several switch models send link-local traffic to all ports. On Cisco switches, this can be disabled using the following interface commands:

```
interface GigabitEthernetx/x
spanning-tree bpdufilter enable
no keepalive
no cdp enable
udld port disable
```

Revision History

V.1 Oct 13, 2023

V.2 Jan 9, 2024

V.2.1 Feb 22, 2024 Added rule that IX cannot be used for transport between member sites

V.3 Apr 17, 2024 Added rule about IPv4 address assignment, removed monitoring section,
other minor changes