



Let's Peer!

Guidelines and best practices

Talia Mares - Edge Strategy Manager @ Meta
Jul-2024

Happy Private Peering Workflow

Agree



How we will peer

- When, where
- Ports: how many, type
- Confirm we are both ready

Plan & Prepare



Get Ready

- PeeringDB profile
- Ports available?
- Know how to configure
- Ports, cables & panel ready


Share details



Good LOA saves time

- Full demarc details
- Using datacenter format
- Include fiber/connector type
- Patch your port to the panel


Connect



Turn on the light

- Request XC
- Provide Cross ID
- Establish physical layer
- Test connectivity
- Check for errors

Configure BGP



Turn up sessions

- Configure LAG
- Enable BGP
- Check prefix limits
- Send prefixes
- Undrain

Celebrate!



We did it!



Before peering – Guidelines

Agree




Plan & Prepare



Send LOA



Connect



Configure BGP



How we will peer

- When, where
- Ports: how many, type
- Confirm we are both ready

Get Ready

- PeeringDB profile
- Ports available
- Know how to configure
- Ports, cables & panel ready

Good LOA saves time

- Full demarc details
- Using datacenter format
- Include fiber/connector type
- Patch your port to the panel

Before peering – Guidelines

- Meta ASN: **32934**
- MTU: 1500 bytes
- We do not support VLANs
- Interfaces:
 - 10G-LR
 - 100G-LR4
 - 100G-LR1
 - 100G-ER4Lite (30 Km)
 - Soon, 400G-LR4

Check peering guideline pdf.

Before peering – Guidelines

Peeringdb profile

The screenshot shows the PeeringDB profile for Meta. The profile is a Platinum Sponsor. It includes a search bar at the top with the text "Search here for a network, IX, or facility." and buttons for "Registeror" and "Login". Below the search bar are links for "Advanced Search" and "Legacy Search", and a language selector set to "English (English)".

The profile is divided into three main sections:

- Meta (Platinum Sponsor):** This section contains a table of organization details. The "Organization" field is "Meta". The "Also Known As" field lists "Facebook, Instagram, WhatsApp, Messenger, Workplace, Oculus, Portal, Novi". The "Long Name" field is empty. The "Company Website" is "https://www.meta.com/". The "ASN" is "32934". The "IRR as-set/route-set" is "RADB::AS-FACEBOOK". The "Route Server URL" and "Looking Glass URL" are empty. The "Network Types" is "Content". The "IPv4 Prefixes" and "IPv6 Prefixes" are both "100". The "Traffic Levels" is "100+Tbps". The "Traffic Ratios" is "Heavy Outbound". The "Geographic Scope" is "Global". The "Protocols Supported" are "Unicast IPv4", "Multicast", "IPv6", and "Never via route servers". The "Last Updated" is "2024-02-01T05:42:34Z". The "Public Peering Info Updated" is "2024-02-01T05:42:28Z". The "Peering Facility Info Updated" is "2023-12-11T17:39:58Z". The "Contact Info Updated" is "2022-10-07T18:44:31Z". The "Notes" field contains "Please submit Peering Requests at: https://www.facebook.com/peering" and "For peering operational issues please contact: noc@meta.com". The "RIR Status" is "ok". The "RIR Status Updated" is "2024-01-24T04:50:42Z".
- Public Peering Exchange Points:** This section contains a table of public peering exchange points. The table has columns for "Exchange Az", "ASN", "Speed", and "RS Peer". The "Exchange Az" column is filtered by "IPV4". The table lists several exchange points, including "1-IX UA", "AKL-IX (Auckland NZ)", "AMS-IX", "AMS-IX Chicago", "AMS-IX Hong Kong", and "AMS-IX Hong Kong".
- Interconnection Facilities:** This section contains a table of interconnection facilities. The table has columns for "Facility Az", "ASN", "Country", and "City". The "Facility Az" column is filtered by "ASN". The table lists several facilities, including "165 Halsey Meet-Me Room", "AIMS Kuala Lumpur", and "Kuala Lumpur".

- Have an **updated peeringdb profile**
- Fields required for us:
 - Prefix limits
 - Contacts
 - Locations
- Create a new account at www.peeringdb.com/register

If you need help creating a peeringdb profile please refer to the guides in <https://docs.peeringdb.com/>

Before peering – Guidelines

At the Data Center

- What is a cross-connect (XC)?
 - Connection inside a Data Center, in this case **ELEA Digital BSB2**
- How does it work at **ELEA Digital BSB2**?
 - You will provide an LOA document with the demarcation details
 - Meta will request Cross Connects to the ISP demarcation inside the datacenter.
 - In the future and when adding more ports, we may ask you to request the additional XCs instead.

Before peering – Guidelines

At the Data Center

- LOA: how to issue it?
 - Can be issued by email, by the authorized ISP personnel. Include:
 - location (address of the facility, floor / room)
 - rack number
 - port
 - connector type
- Verify you have the appropriate interfaces and cables
- The ISP is responsible for any connections inside their cage/rack. A XC will not solve this

Physical connection

Agree




Plan & Prepare



Send LOA



Connect



Configure BGP



Turn on the light


- Request XC
- Provide Cross ID
- Establish physical layer
- Test connectivity
- Check for errors

Physical connection

- ISP will make sure that
 - their router interface(s) is **admin up**
 - interface is sending optical power
 - cabling between ISP router to their patch panel is completed
 - share optical power levels over email
- If ISP sent the LOA, Meta sends a request to the datacenter for a XC
- The datacenter / colocation provider runs the XC



partner




Physical connection

- After requesting the cross connect, we will
 - have the colo provider test the physical link, and confirm whether there is light or not at the other end. If the light levels are bad, the colo provider can clean the fibers and re-test, regardless of whether the peer is connected or not on their side.
 - share optical power levels over email
 - share Cross ID / Completion ID, as provided by the datacentre / colo provider.



partner




Physical connection

- Test physical link end to end
 - Share optical power levels to make sure they are OK
 - Meta will provide test IP address from the 169.254.0.0/16 prefix
 - In advance ISP must remove filters to accept prefix 169.254.0.0/16
 - Run ping tests for each port
 - Check interface log for errors
- Configure interface: LAG, **Both IPv4 and IPv6**, MTU



partner



BGP & Peering best practices

Agree




Plan & Prepare




Send LOA



Connect



Configure BGP



Turn up sessions

- Configure LAG
- Enable BGP
- Check prefix limits
- Send prefixes
- Undrain

BGP & Peering best practices

Configure BGP

- Have your PeeringDB profile updated
- set prefix limits -> talk to your peer and communicate how many routes you will send and how many will you accept.
- advertisement policy:
 - with peers: Use more specific routes
 - with transit: Use summarized routes
 - Consistent advertisements -> if multiple LACPs, send exact same routes for load balance and back-up
 - Meta accepts prefixes size up to /29 for IPv4 and /64 for IPv6 in peering
- (optional) use BGP communities to identify routes:
 - to make decisions or prioritize traffic based on communities
 - It can be informational: geographic origination, type of relationship (peer, customer)

BGP & Peering best practices

Other bgp attributes:

- as-path & prepends
 - We consider latency first, consider as-path only in cases latency is close between 2 paths
 - **Do not use prepends in peering session**
 - If you have to, please use only 1 prepend (public peering, transit, etc).
- May we use MD5? Please don't! Check [MANRS](#) instead for guidelines on routing security.
- How about MED? We do not support it. MED will be ignored.
- Traffic Engineering?
 - If you don't want traffic for a specific prefix at a specific location, please do not advertise that prefix to us in that location
 - If necessary, we prefer you use BGP communities to signal prioritization

BGP & Peering best practices

- How to do prioritization
 - If there are MNA caches in your network, remember we monitor them via IP. If we detect connectivity degradation, cache will be drained -> make sure to prioritize routes from/to MNA
 - If there's multiple peering connections & wish to prioritize routes, consider using Meta BGP community signaling (32934:10000 to 32934:10012)
- What happens if i have private IPv4 addresses and NAT in my network?
 - We do not support this in peering. Don't advertise private IP space - it will be ignored
- Does Meta use or enforce RPKI?
 - Meta uses RPKI, and is planning to enforce it (reject invalid routes) in upcoming months.
 - Update your ROAs!

After Peering is set up

Manage expectations



Happy Peering

What will happen to traffic?

- Can take up to 24 hours to move traffic
- 99% of traffic will be in peering, with some exceptions
- Peering will fail. POPs will fail. Caches will fail. **Keep headroom for failure scenarios**

After Peering is set up

Support

Options to open tickets:

- Network Partner Portal access -> best way to create a ticket, many teams checking tickets at the same time
- noc@meta.com
- Meta NOC works 24 x 7 x 365

TALK TO US

Peering > peering@meta.com

Caching > mna@meta.com

Issues? > noc@meta.com

