

Date: 1 February 2012

Shepard: Chris Grundemann

Subject Matter Expert(s) (SME): Chris Grundemann

Status: Outline

BCOP Subject: IPv6 Peering and Transit

The content here within is intended to be original content authored by the Global Network Engineering Community (GNEC) “at-large” in an organized, democratic, “bottoms-up” approach.

Best Current Operational Practice (BCOP)

1. BCOP Summary (Appeal)

This BCOP aims to provide general IPv6 Peering and Transit guidelines that can be followed by any network operator when planning and implementing any IPv6 Peering/Transit relationship. The primary focus is on understanding BGP peering and filtering.

2. BCOP Background / History

We (the GNEC) have made many mistakes with IPv4 Peering and Transit configurations and operational practices. As operators begin turning up more and more IPv6 E-BGP sessions with peers and transit providers, there is an opportunity to do things right from the beginning. While the details of these peering and transit relationships can be varied and specific, the technical realities remain largely the same. These technical realities inform the best practices listed here.

3. BCOP

1. Establish new, IPv6-Only peering.

- Two BGP sessions between peers
 - Do NOT just exchange IPv6 NLRI over existing IPv4 peering
- IPv6 topology independent of IPv4, and vice versa
- Outages: IPv6 session goes down if IPv6 reachability is lost
- Maintenance: IPv4 and IPv6 sessions don't affect each other
- Operational Clarity

2. Filter

- Methodology not special, generally mirrors IPv4
 - Explicit, bogons/martians, maximum-prefix, prefix size, etc.
 - The devil's in the details
 - Three high-level “themes:”
 - Filtering customers
 - Filtering Peers and Transit providers
 - Filtering your own routes
- a. Filter Routes Coming from your Customers**
- Explicit filter:
 - Allow customer network(s)
 - Deny all else
 - Customer is responsible for updating you
 - Should be infrequent if at all in IPv6
 - Prefix size?
 - In IPv4 it's common to allow down to /32 (for blackholing, etc.)
 - Orders of magnitude more addresses in v6
 - If allowed in IPv6 (/128); ensure proper maximum-prefix limits
 - Filter out more specifics before announcing to Peer / Transit
 - Could filter to /48 if customer needs to leverage more specifics
- b. Filter Routes Coming from Peers & Upstreams**
- Generate prefix filters from IRR:
 - Everyone should register in an Internet Routing Registry (IRR)
 - Allows tool-based filter generation
 - As good as the data in the IRR
 - We have a chance to do this right from the beginning
 - Bogon and max-prefix
 - Some networks don't use an IRR (yell at them)
 - Bogon filter rejects the crazy stuff (reserved, small prefix, etc.)
 - A maximum-prefix limit protects against route overload
 - Better than nothing
- c. Filter your own Routes**
- Only allow you and your customer's space
 - Don't deaggregate unnecessarily
 - Only /48 and shorter (larger) prefixes
 - Use communities to identify and filter routes
 - Level of hierarchy – Filter more specifics

- Type of customer – BGP customers more likely to multihome
- Use – No need to advertise specific infrastructure prefixes
- Locational – Allows regionalization for TE, etc.

3. Use an IRR

- Register in an existing database
- Register your own database
- Synchronize your own filters with your IRR

4. Use IPAM

- Massive address space
- DNS/DHCP complexities

5. Register in the PeeringDB

4. BCOP Conclusion

1. Establish new, IPv6-Only peering.
2. Filter
 - a. Filter Routes Coming from your Customers
 - b. Filter Routes Coming from Peers & Upstreams
 - c. Filter your own Routes
3. Use an IRR
4. Use IPAM
5. Register in the PeeringDB