Updated BGP Operations and Security

Revisiting BCP194

Tobias Fiebig

Max-Planck Institut für Informatik

November 3, 2023

Recap: BCP194 / RFC7454

- “BGP Operations and Security”
- From Feb. 2015
- Best practices for:
  - Protecting BGP sessions (TCP layer)
  - Protecting BGP speakers (general network security)
  - Protecting BGP (or rather: What to (not) import from whom)

Issues in BCP194

- Makes some borderline normative (and problematic) recommendations on announcing IXP peering LANs
  - *“ In that case, any IXP member SHOULD make sure it has a route for the IXP LAN prefix or a less specific prefix on all its routers and that it announces the IXP LAN prefix or the less specific route (up to a default route) to its downstreams.”*
- New attacks, defenses, and terminologies emerged since 2015

Issues in BCP194

- Makes some borderline normative (and problematic) recommendations on announcing IXP peering LANs
  - *“ In that case, any IXP member SHOULD make sure it has a route for the IXP LAN prefix or a less specific prefix on all its routers and that it announces the IXP LAN prefix or the less specific route (up to a default route) to its downstreams.”*
- New attacks, defenses, and terminologies emerged since 2015

General sentiment: *Somebody* should *somehow* do *something*. Maybe.

Revisiting BCP194: Approach

Currently, the idea for revisiting BCP194 is:

- Edit from the original towards a version that ‘contains everything’ that could be changed
- The current version takes a very ‘maximum’ approach to many things, with the objective of providing a foundation for discussion
- It takes a wider perspective on ‘security’, i.e., also looks at reliability / operating BGP reliably, even if no one is attacking things

Terminology

Issues

- BCP194 uses some terminology with different meaning, e.g., peer (other BGP speaker to which a BGP speaker has a BGP session) and peer (other AS with which only one's own cone is exchanged)
- Terminology has progressed a bit (see RFC9234 and the ASPA draft)
- Some things could use more exact definitions
- BCP194 codifies what a Tier-1 operator is; Which might be a bit of a ... thing.

Terminology: Changes

Changes

- Unify terminology (BGP neighbor; Peer only for topological peers)
- Import terminology from RFC9234 and the ASPA draft (including ‘mutual upstream’)
- Add more definitions where necessary
- Remove reliance on Tier-1 notion

‘New’ Attacks

Issues

- Since 2015, a couple of attack vectors showed up;
- There are some PMTUD attacks for off-path route injection
- There is the GRT-blowup with deaggregation

‘New’ Attacks: Changes

Changes

- Discuss a few more of these attacks and associated defenses
- **Discussion Point:** Introduced a wider (per neighbor AS across all sessions with the neighbor) prefix limit

IXP LANs

Issues

- The notion of handling IXP LANs in BCP194 might be seen ‘a bit controversial’ in IXP circles
- There are new use-cases of IXPs (‘upstream’ route servers, upstream via IXPs)
- uRPF breaks things

IXP LANs: Changes

Changes

- Clarify that IXPs decide whether their IXP LAN(s) should be advertised
- Removed BGP example in the Appendix
- Noted additional IXP Lan cases (upstream via etc.) and next-hop filtering needs

Filtering Types

Issues

- In BCP194, the main sectioning item is filtering types (prefixes, AS-path etc.)
- In the end, everything is ultimately about which NLRI gets imported (given it has a specific prefix/next-hop/AS path...)

Filtering Types: Changes

Changes

- Focus more on the general frame of importing NLRI
- Split prefix filter based on source of lists (IRR vs. more 'static' lists)

New Things to Filter

Issues

- We got many more things to filter for/with

New Things to Filter: Changes

Changes

- Add ASPA and BGP roles/OTC
- Note use of community-based filtering, also discussing leak-potential of static filters being used[a]
- Emphasize outbound filtering
- Reference RFC9319
- Add large-communities/scrubbing
- Add max-prefix limits and global limits

[a] Yes, been there, seen that

New Things to Filter: Changes *(cont.)*

Changes (that might go a bit far)

- Add note on suggested shortest prefix lengths (/8 (v4) and /16 (v6))
- Add note on ASPATH max-length filtering
- Suggest iBGP filtering

Rule Generation & Ordering

Issues

- BCP194 lacks reflections on the algorithmic complexity of rules[a]
- BCP194 does not discuss the impact of sequential processing on filter import[b]
- Filter generation is not discussed in-depth (and BCP194 still recommends IRRtool)

[a] Yes, got bitten by that.
[b] Yep, that one too.

Rule Generation & Ordering: Changes

Changes

- Discuss impact of order on filter computation time and import caveats
- Discuss pathways for generating filters in resource constraint environments
- Recommend BGPq4 instead of IRRtool

Community Scrubbing

Issues

- BCP194 lacks authoritative terminology on what one can do with harmful transitive attributes.

Terminology: Changes

Changes

- Added a `MAY` there

Current Todo Items

- Unify terminology around route/prefix/NLRI, which is currently a bit mixed up
- Use of `MUST` vs. `SHOULD`; Other contemporary drafts use stronger language; From a general standpoint I'd argue `MUST` might be better under the premise of 'to follow best practices this `MUST` be done' with the caveat of 'best practices `SHOULD`' be followed; Still, currently it still uses BCP14 `SHOULD` for all points
- Putting in a point on having preparations in place to filter specific BGP options at the border in case one's infrastructure topples when a specific option is sent
- Figure out if this should be update or obsolete
- Fix nits in the abstract

Comments so Far

- Add note on honoring GSHUT on IXPs
- Discuss use of lpref on IXP RS
- Add note on need for individual judgement (my network, my rules)
- Discussing MED based oscilation (RFC7964)
- Expand community filtering (scrub extended (RFC4360) and long lists (+100))
- Make attribute scrubbing more specific (only scrub what is specifically known to cause harm now)

Discussion Points

Discussion Point 1

- iBGP filtering: Is this going to far?

Discussion Point 2

- Global max prefix limits: Too unreasonable? Maybe more about monitoring?

Discussion Point 3

- What else is missing/too much/should change?

Discussion Point 4

- Should I `SHOULD` or should I `MUST`?