CDNI Footprint Advertisement

draft-previdi-cdni-footprint-advertisement-01

Authors: Stefano Previdi <sprevidi@cisco.com>
Francois Le Faucheur <flefauch@cisco.com>
Allan Guillou <allan.guillou@sfr.com>
Jan Medved <jmedved@cisco.com>

Terminology

- **Footprint (FP)**: The exhaustive set of Prefixes a CDN is can/is willing/is able to serve.
- **Footprint Element (FPE)**: Arbitrary set of prefixes with attributes. They can be implicit (inferred from BGP) or explicit (advertised).
- **Implicit Footprint Element Advertisement (FPE Implicit Advertisement)**: Footprint Element information derived from the BGP database.
- **Footprint Element Advertisement (FPE-Adv)**: MP-BGP Message used by a CDN in order to advertise or withdrawn Footprint Elements and their attributes.
- **Footprint Reachability Advertisement (FPR-Adv)**: MP-BGP Message used by a CDN in order to advertise or withdrawn Footprint reachability information.
- **Capability Advertisement (CAP-Adv)**: MP-BGP Message used by a CDN in order to advertise or withdrawn capability information.

Multiprotocol-BGP

- BGP is well known, scalable, loop-free, topology agnostic, efficient and flexible tool used already in multiple contexts
  - IPv4/IPv6 routing
  - Different VPN flavors: L2/L3
  - Multicast (v4/v6)
  - Multilayer topology advertisement (from optical layer up to application layer)
- The use of BGP in this proposal does NOT imply to share any fate with the BGP running in the network layer
  - Same tool, different context
  - Deployment strategy: likely a CDNI-BGP overlay
  - No disclosure of any internal data

Multiprotocol-BGP

- Case 1: SP operated CDN (focus)
  - CDN footprint is known in SP BGP database and likely advertised to Internet
  - Footprint likely aggregated on AS number
  - Still exceptions exist
- Case 2: OTT CDN
  - CDN Footprint spans multiple SPs/ASs
- Both cases addressed by two new MP-BGP messages
  - Footprint Element Advertisement (implicit or explicit)
  - Footprint Reachability Advertisement

Footprint Element

- Footprint Element Advertisement (FPE-Adv)
  - Implicit:
    - A Footprint Element (FPE) inferred from one or more prefixes known in the internet BGP database
    - Example: all prefixes belonging to a given AS.
    - Advantage: no signaling, highly scalable
  - Explicit
    - When a footprint is a sub/partial/super set of what exists in the BGP Internet database then an explicit FPE advertisement is necessary
    - Example: CDN-A advertises ability to reach a subset of prefixes located in a given AS
- FPE have an Identifier
  - Implicit: the AS number (extensible to other criteria)
  - Explicit: any ID number
- Note: most of SP operated CDNs will have their footprint inferred

Footprint Reachability

- Footprint Reachability Advertisement (FPR-Adv)
  - CDN claim reachability/willing to serve a given FPE
    - Specified through FPE Identifier
  - One message for the whole FPE
  - FPR-Adv contains information about how CDN can effectively reach the FPE
    - Origin_AS_PATH
    - Useful when multiple CDNs claim same reachability
- FPR Messages follow BGP propagation rules
  - Propagation path is recorded
  - Loop free propagation mechanisms
  - Gives information about the CDNI Mesh
- Support any CDNI Mesh topology
  - Flat full/partial mesh, hierarchical/cascade, …

Capabilities Information

- Current stage of CDN Capabilities definition doesn't allow to specify an encoding format
- Multiple options are still possible, from simplest to more elaborated:
  - Standard Community Attribute
  - Extended Community Attribute
  - Ad-Hoc BGP NLRI with the appropriate format defined by CDNI-WG
  - First requirement: no protocol standardization for new capabilities definition
- NOTE WELL: Capability granularity can go down to user prefix

Step-1: Infer Footprint from BGP-4 Database

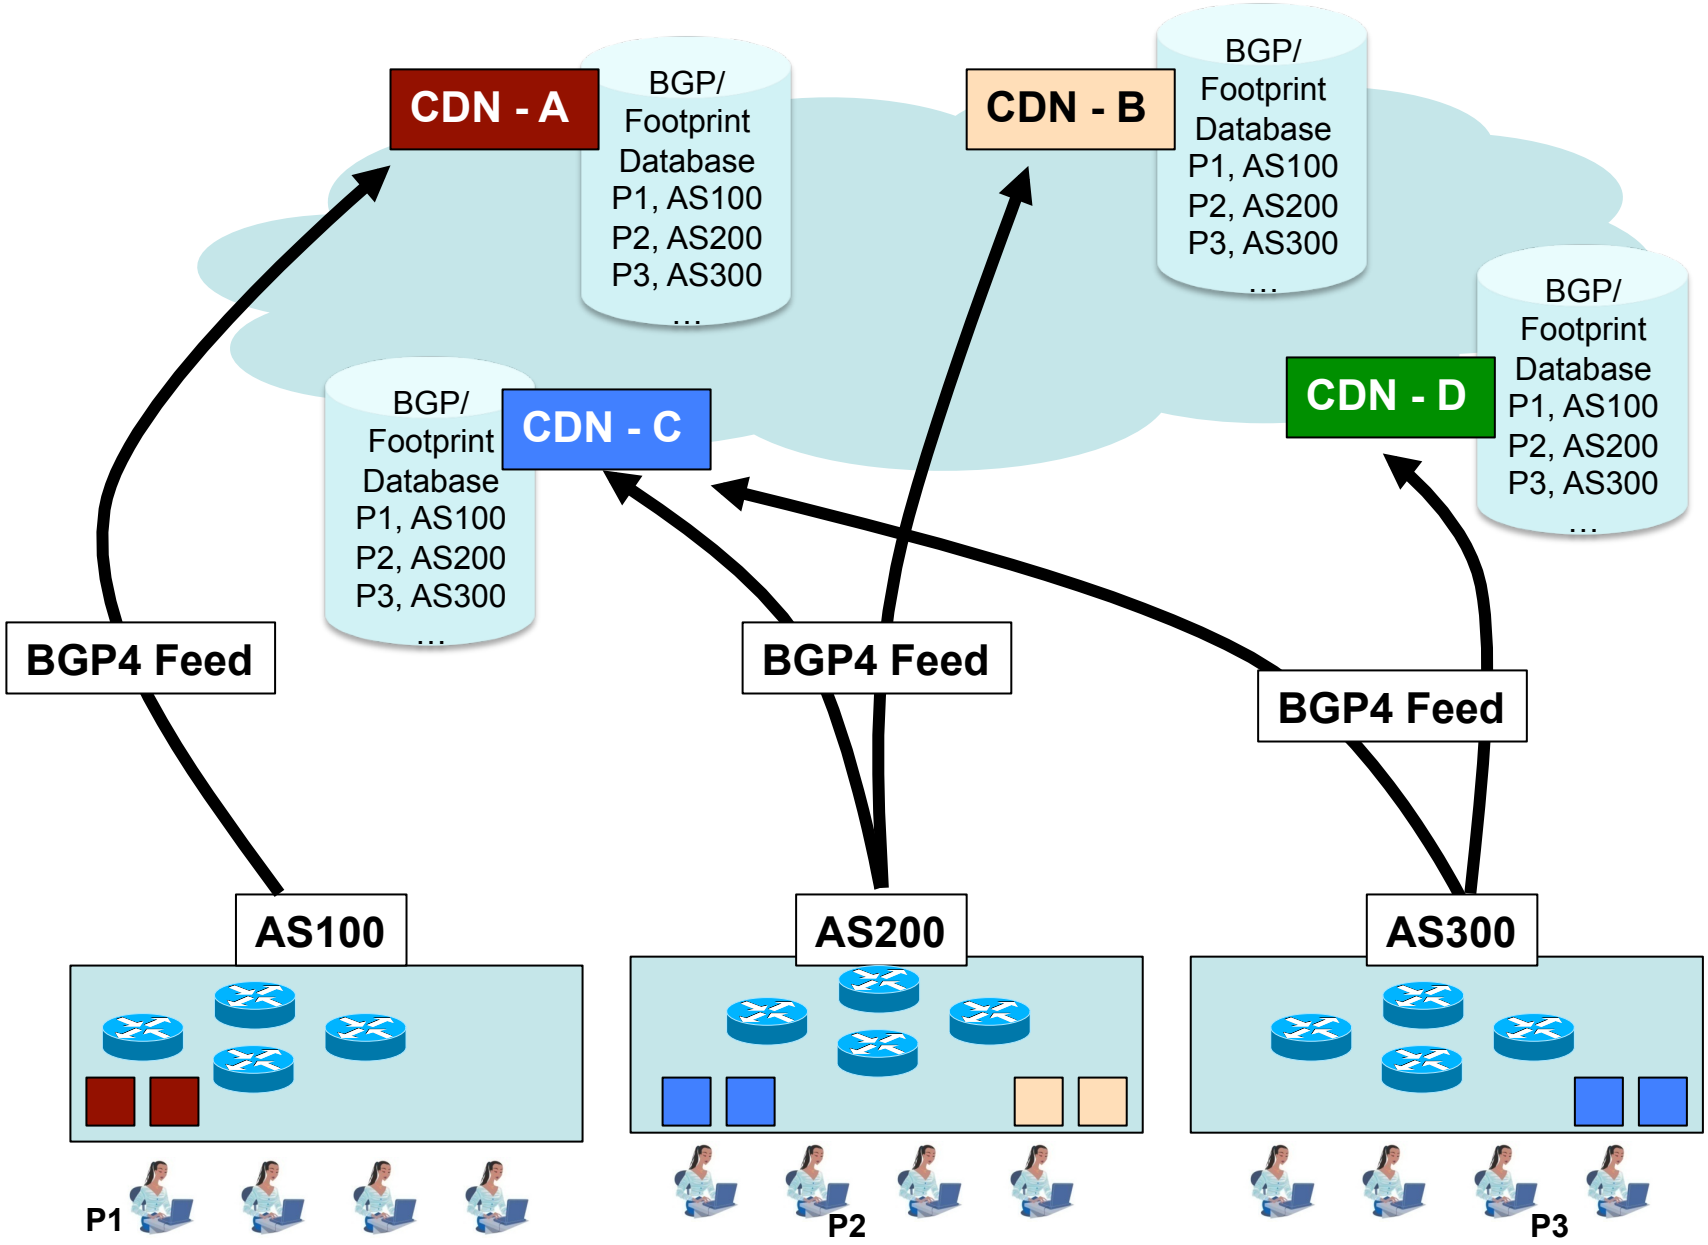

Step-2: Advertise Reachability (FPR) & Capabilities

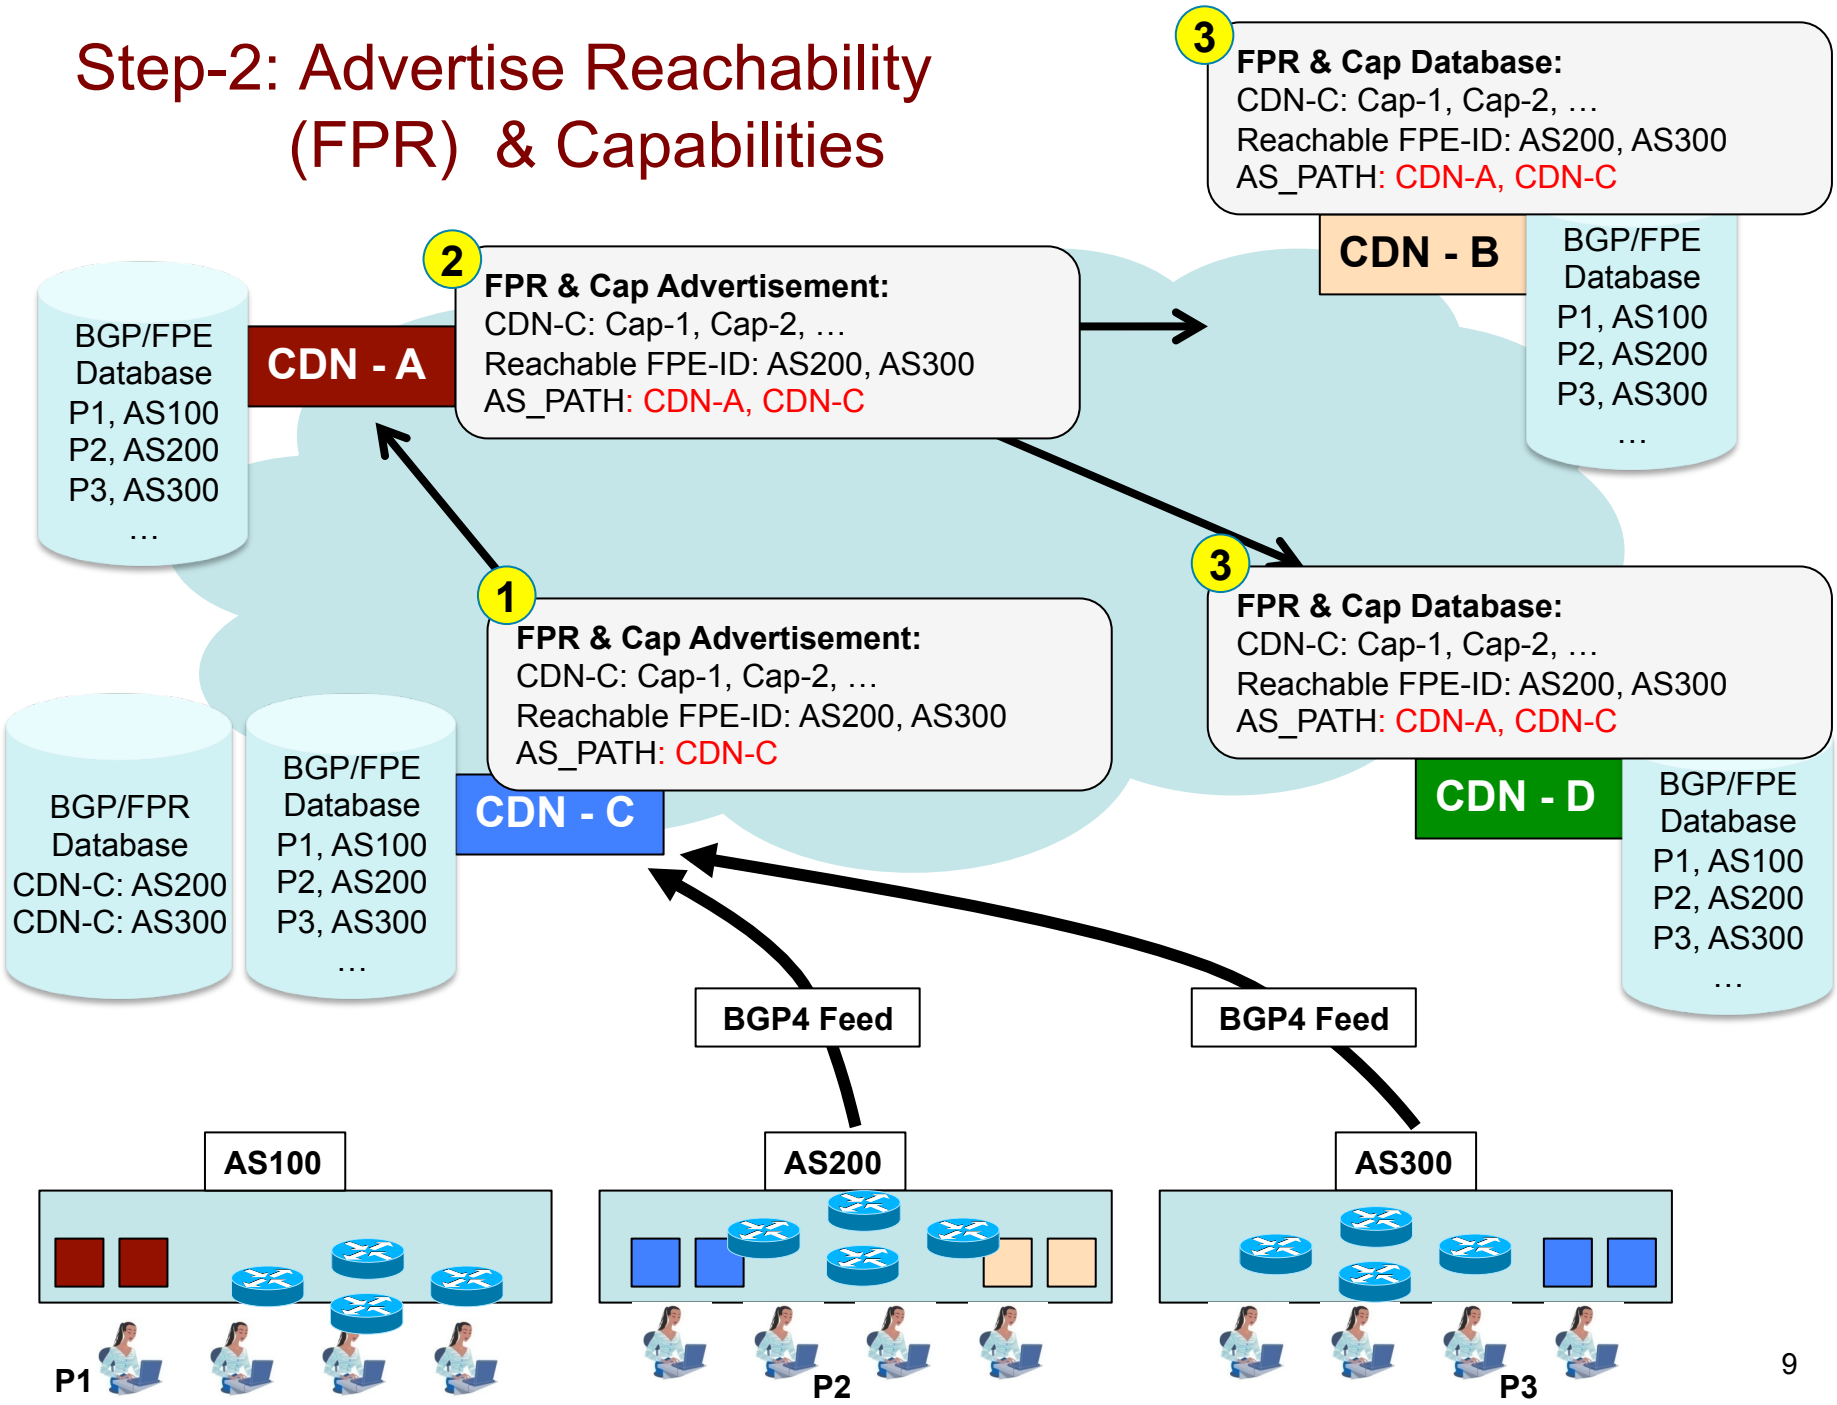

Workflow

- When an upstream CDN (uCDN) receives a request from a user, it has to determine the downstream CDN (dCDN) the request is to be redirected to:
  - Determine which footprint the user belongs to
    - Lookup in Footprint Elements Database
  - Determine dCDN claiming connectivity to user Footprint
    - Lookup in Footprint Reachability Database
  - Apply selection rules

To Do List

- Fix inconsistencies in NLRI Encoding
- Allow other grouping methods: AS, Community, location, through use of existing BGP attributes
- The same mechanism applies to other context (e.g.: clouds) which may lead into a more generic mechanism
- Capability advertisement
  - Format, encoding, granularity, …
- Btw, implementation exists

draft-previdi-cdni-footprint-advertisement-01

Thank You