CDNI Footprint Advertisement

draft-previdi-cdni-footprint-advertisement-00

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

draft-previdi-cdni-footprint-advertisement-00

- This draft proposes the use of Multiprotocol-BGP to realize the Footprint & Capabilities Advertisement component of the Request Routing Interface.

draft-previdi-cdni-footprint-advertisement-00

- Introduction:
  - When an upstream CDN (uCDN) receives a request from a user, it has to determine what is the downstream CDN (dCDN) to which the request is to be redirected.
  - Decision is based on MULTIPLE criteria
  - CDN's Footprint & Capabilities are SOME of these criteria

Footprint Information
The problem…

- If the Footprint is the set of prefixes a CDN is willing to serve
  - Typically: a bunch of IP prefixes (can be quite large)
- Then, CDNs participating in the CDNI Mesh must know each other Footprint in order to select the best dCDN during Request Routing process
- If Footprint information had to be advertised between CDNs, it will result in large amount if information flooded across the CDNI Mesh
- Requiring each CDN to advertise its Footprint information to the CDNI Mesh is unpractical and un-scalable

Footprint Information
The solution…

- Footprint Information is in fact prefixes, not different (in the format) from what available in the Internet BGP table
  - In all SP networks
- A CDN Footprint is in fact a subset of the Internet BGP-4 table
- What if a Footprint can be identified by its Autonomous System membership
  - I.e.: all prefixes originated in AS-X are considered as part of Footprint X
- Footprint granularity is the Autonomous System
  - In a first stage of CDNI it will matches well the requirements
- Footprint Information becomes: set of ASs a CDN is willing to serve
- This draft proposes a set of mechanisms through which:
  - Each CDN gets a BGP4 Internet feed from its underlying SP
  - Each prefix is mapped to its Autonomous System of origin
  - Each CDN advertise its ability to serve ‘group’ of prefixes identified by their AS number
- At this stage of the draft a Footprint is identified by its AS Number
- Future revisions may include other grouping methods

Connectivity Information

- Set of Footprints the CDN is capable/willing to serve
- Advertised into the CDNI Mesh
- Low volume information

Distinction between Footprint and CDN connectivity to the Footprint

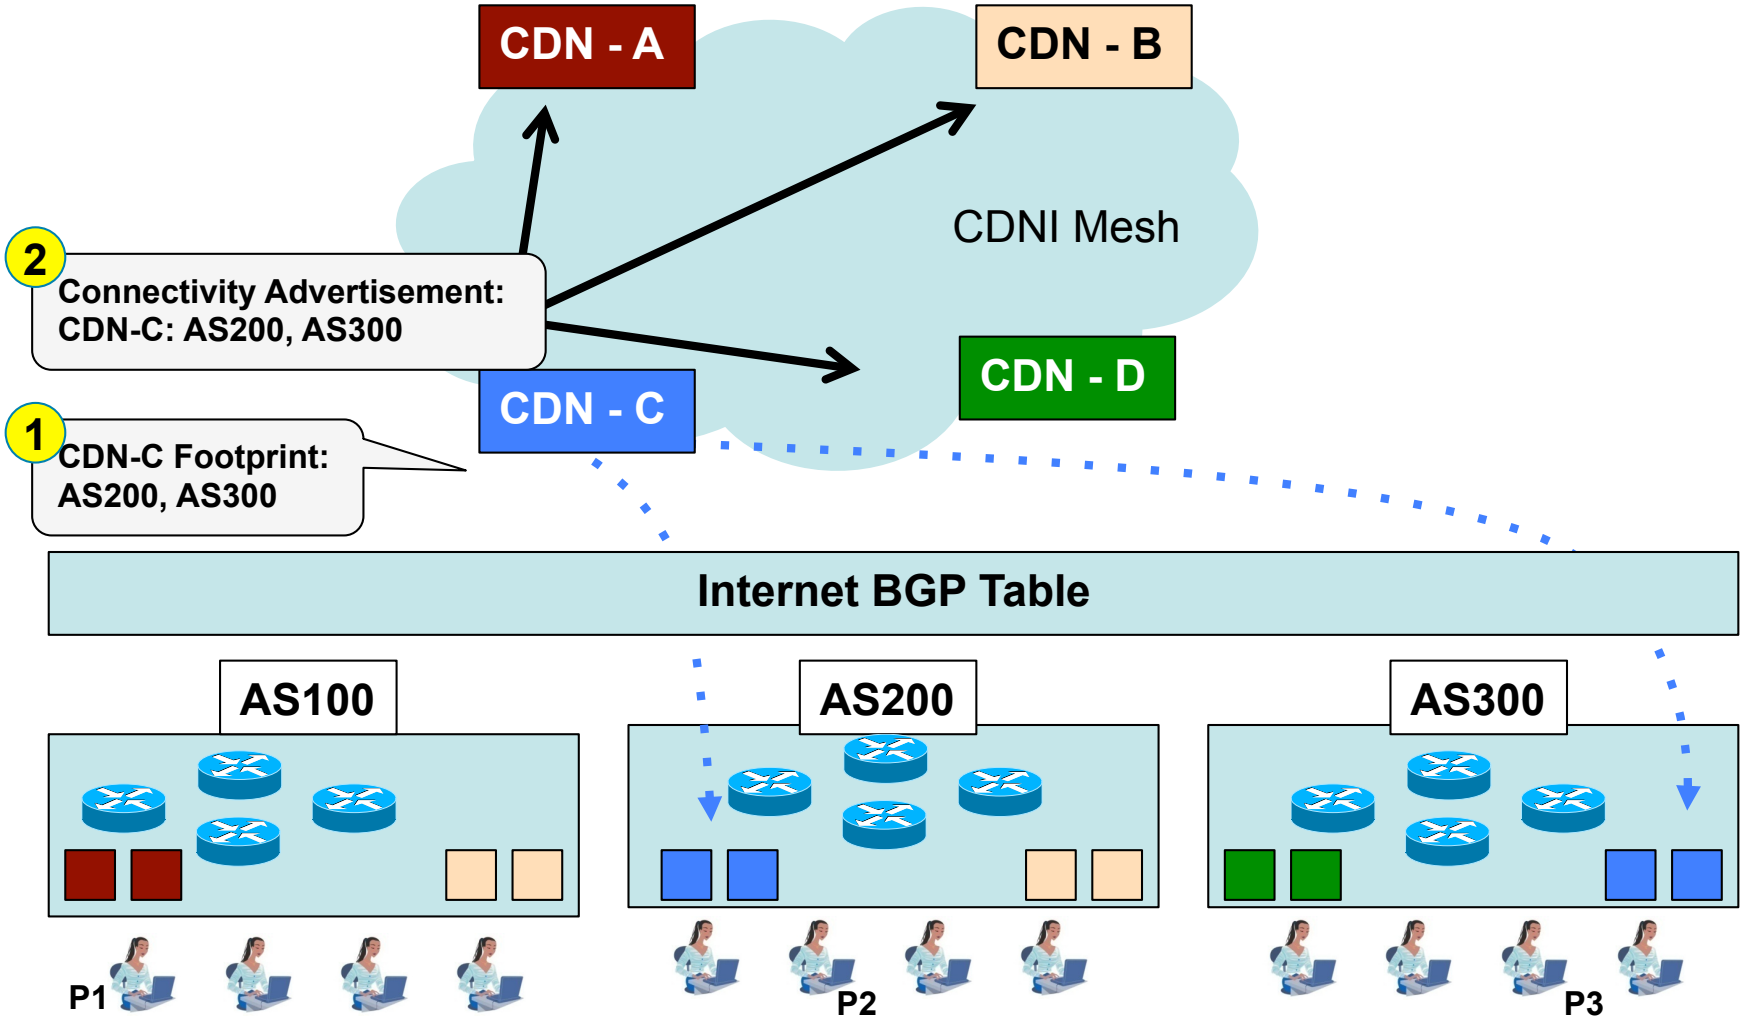

Capabilities Information

- CDN Capabilities are signaled using the same method SPs signal prefix capabilities and characteristics:
  - Standard Community Attribute
  - Extended Community Attribute
- Communities are already widely deployed in order to describe prefixes:
  - Location
  - Connectivity type
  - Peering point
  - …
- Same for CDNs
  - Multiple communities can be used for a variety of capabilities
- CDN originates MP-BGP advertisement including the set of Communities describing the CDN capabilities

Multiprotocol-BGP

- BGP4 is well known, scalable, efficient, flexible, …
- MP-BGP is an extension to BGP4 in order to carry different address families
  - Allowing isolation
- In this proposal, CDNs use MP-BGP messages to advertise their capabilities and their ability to serve given Footprints
- MP-BGP allows:
  - Advertisement of Footprint connectivity
  - Advertisement of CDN capabilities
  - Advertise how CDNs are logically inter-connected
  - Each CDN to control the policy applicable to these advertisements

Step-1: Infer Footprint from BGP-4 Database

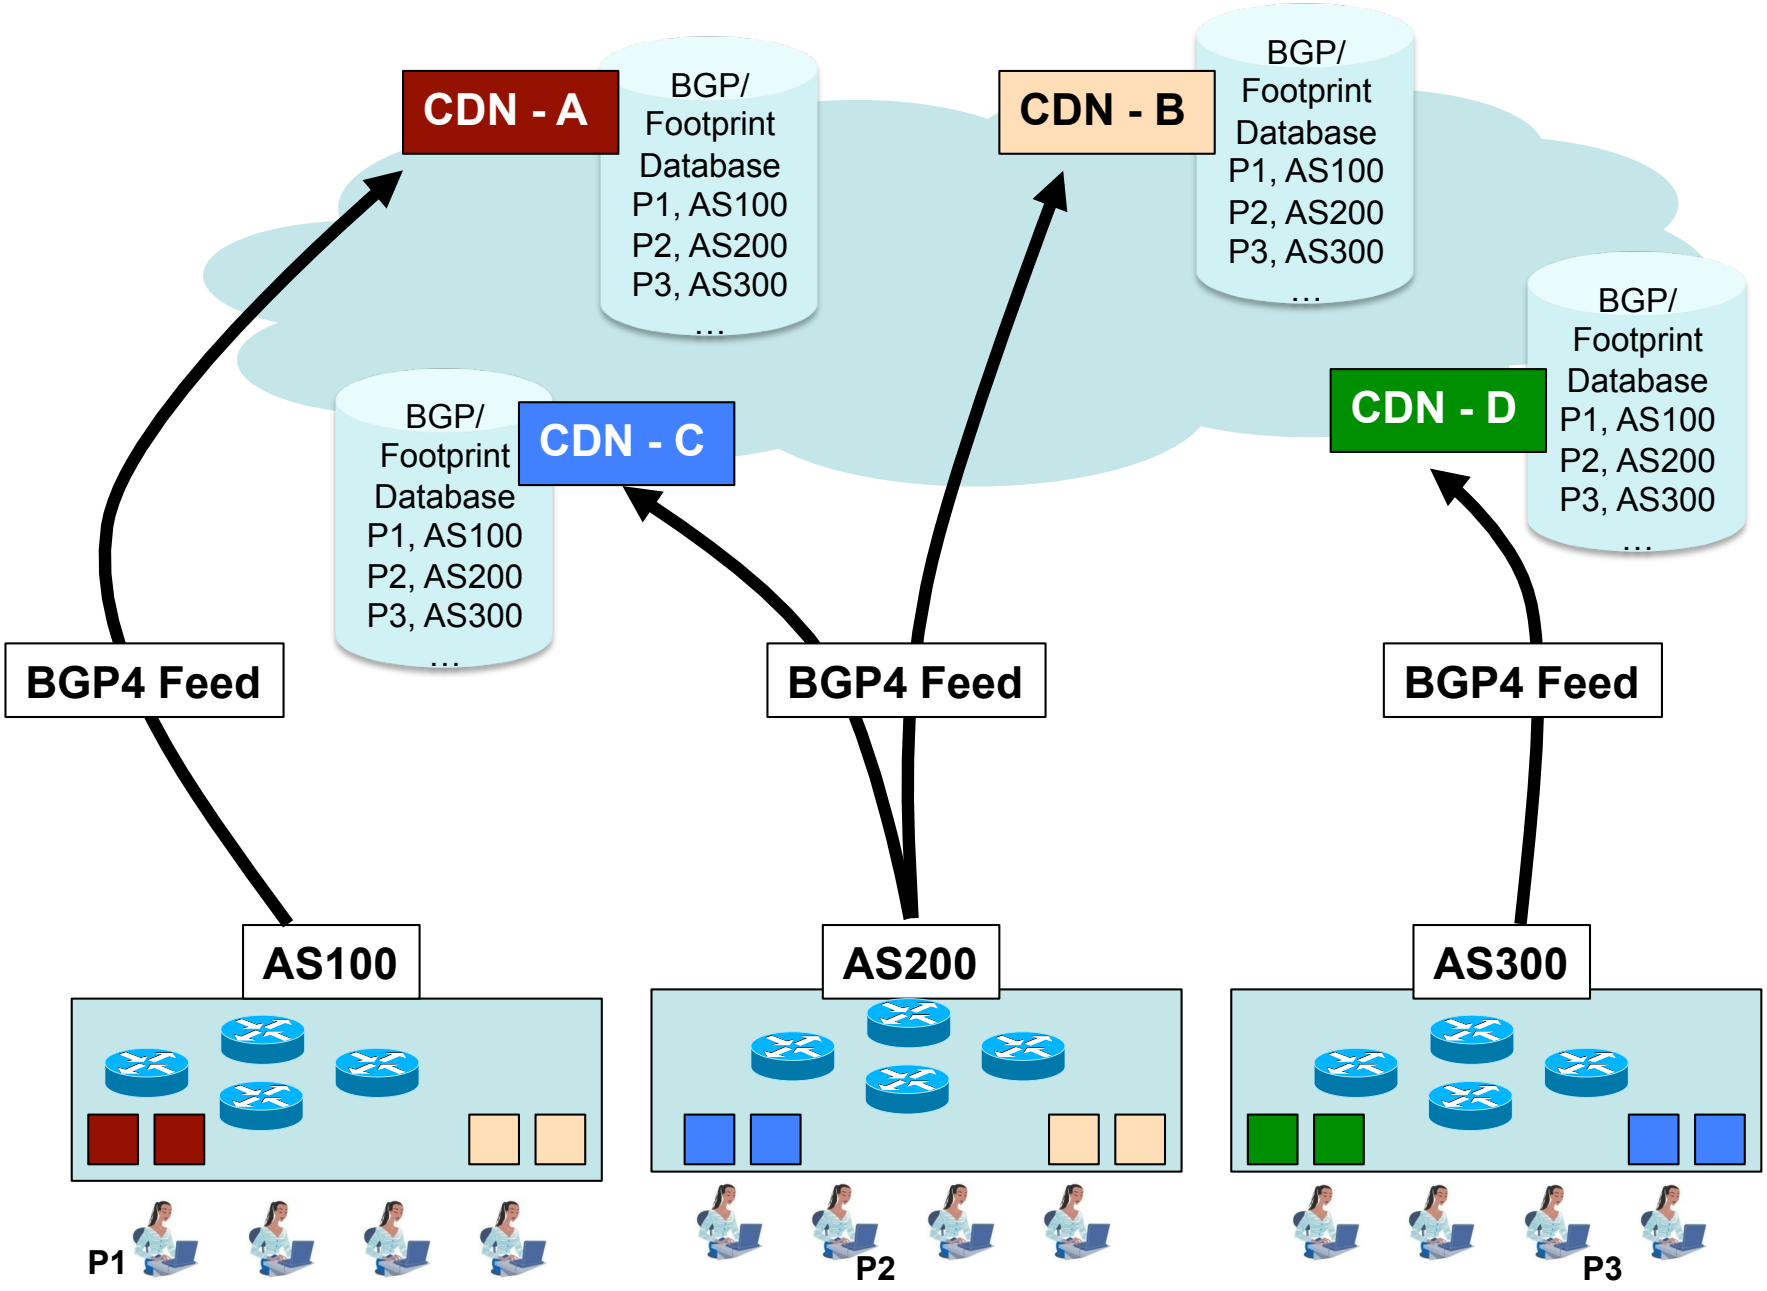

Step-2: Originate Connectivity Database

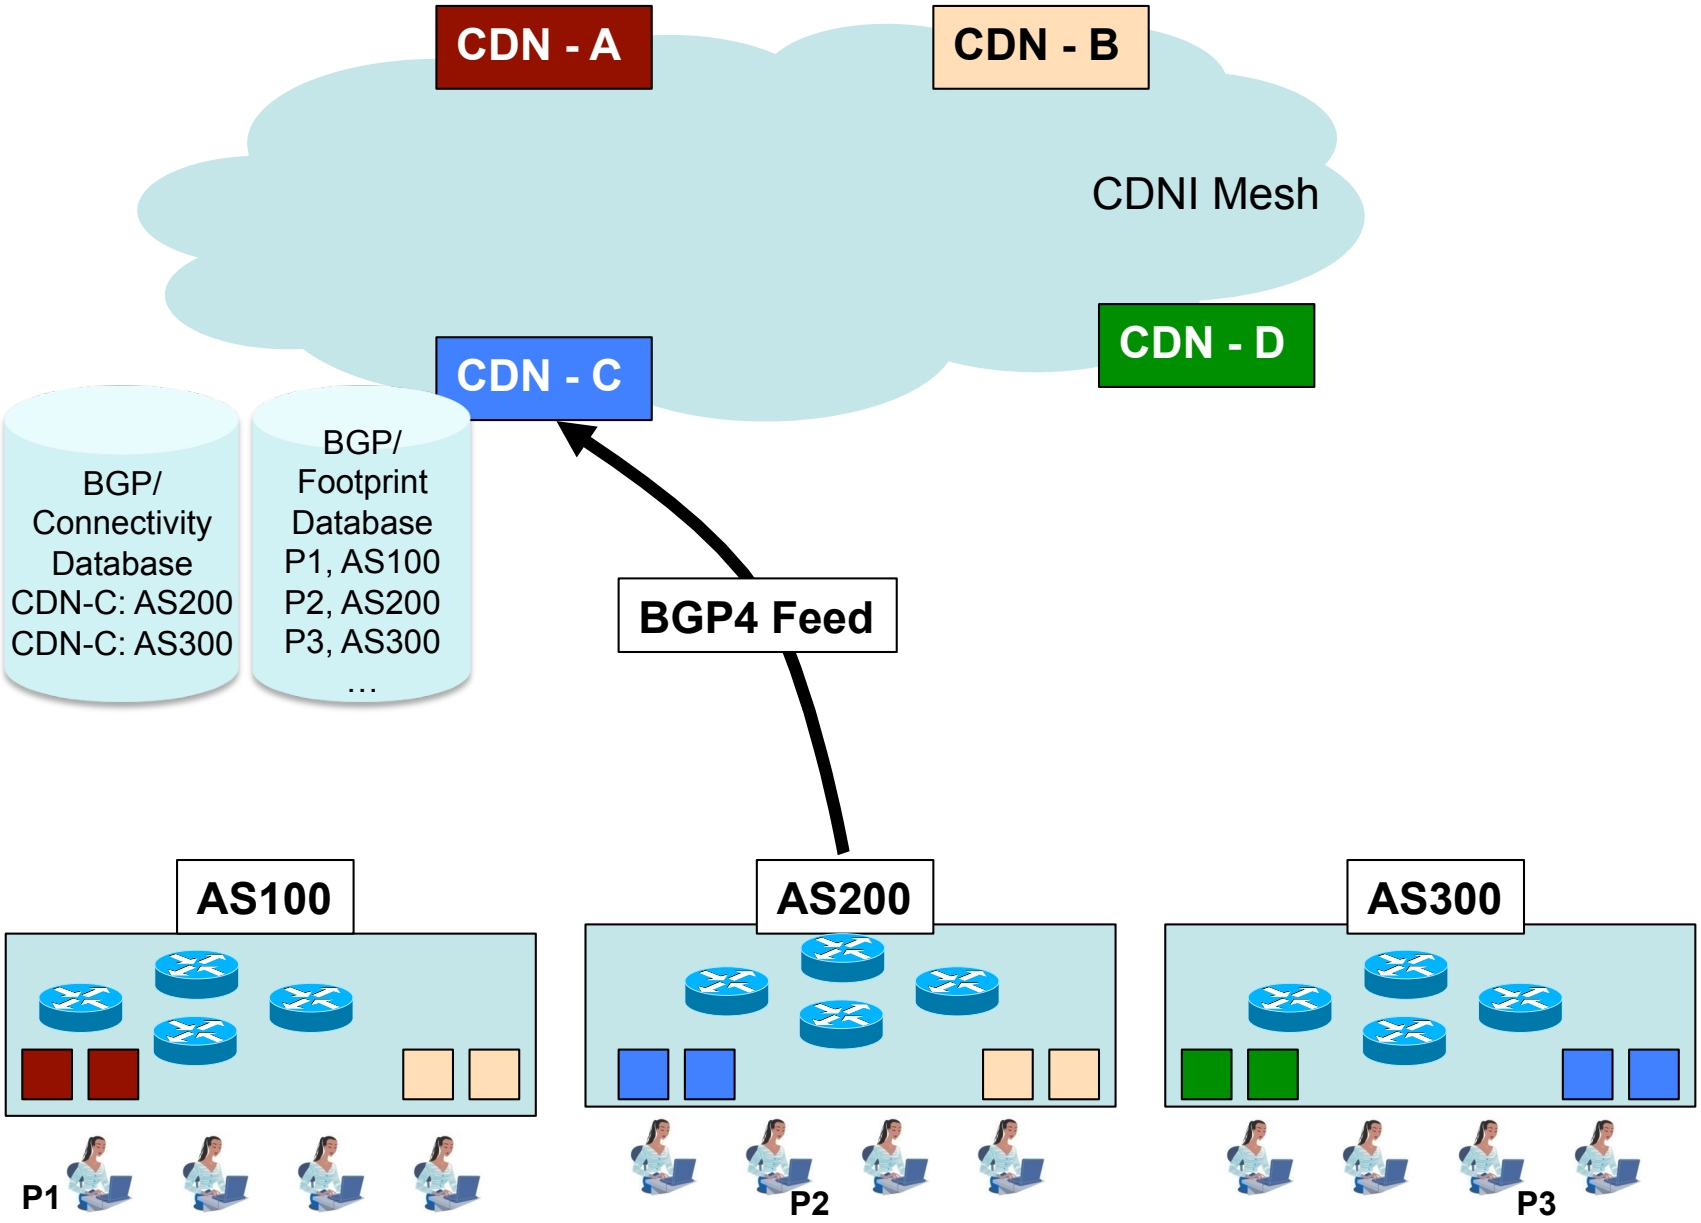

Step-3: Advertise Connectivity & Capabilities

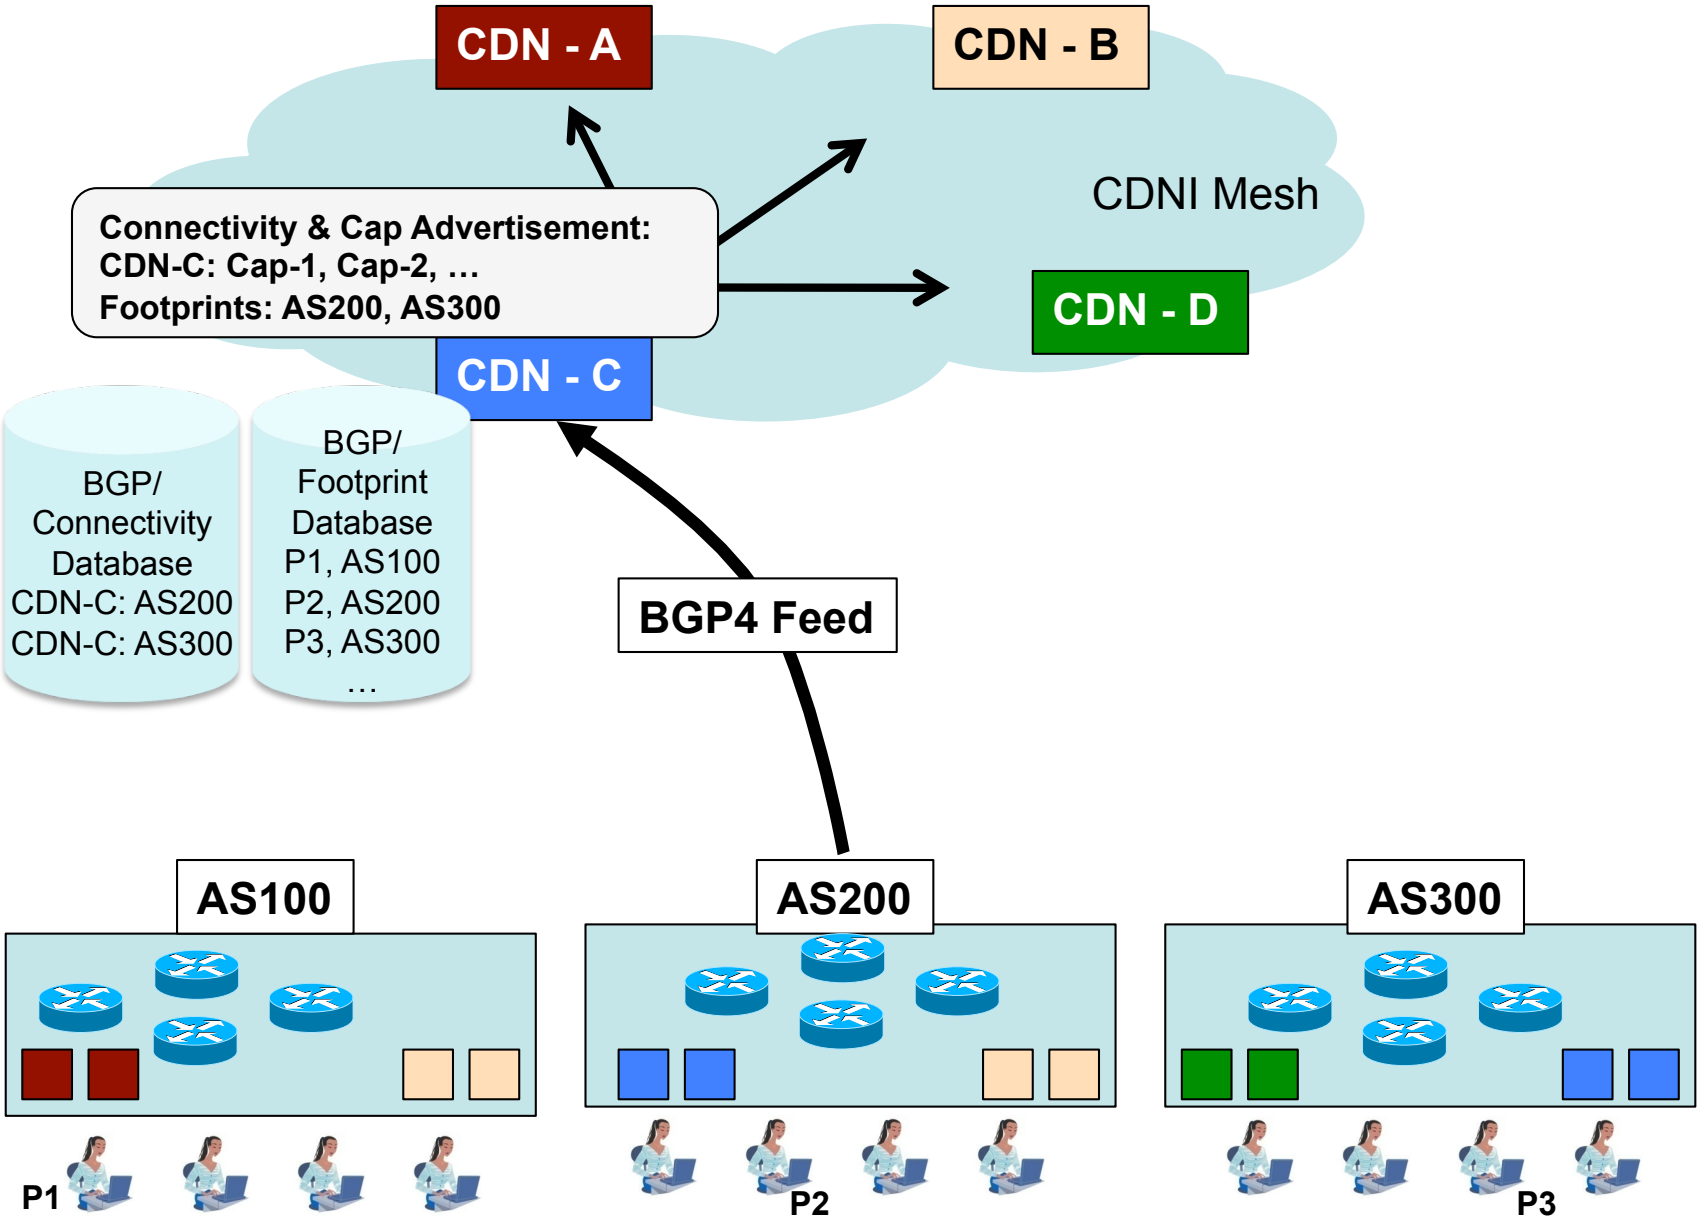

Partial CDNI Mesh

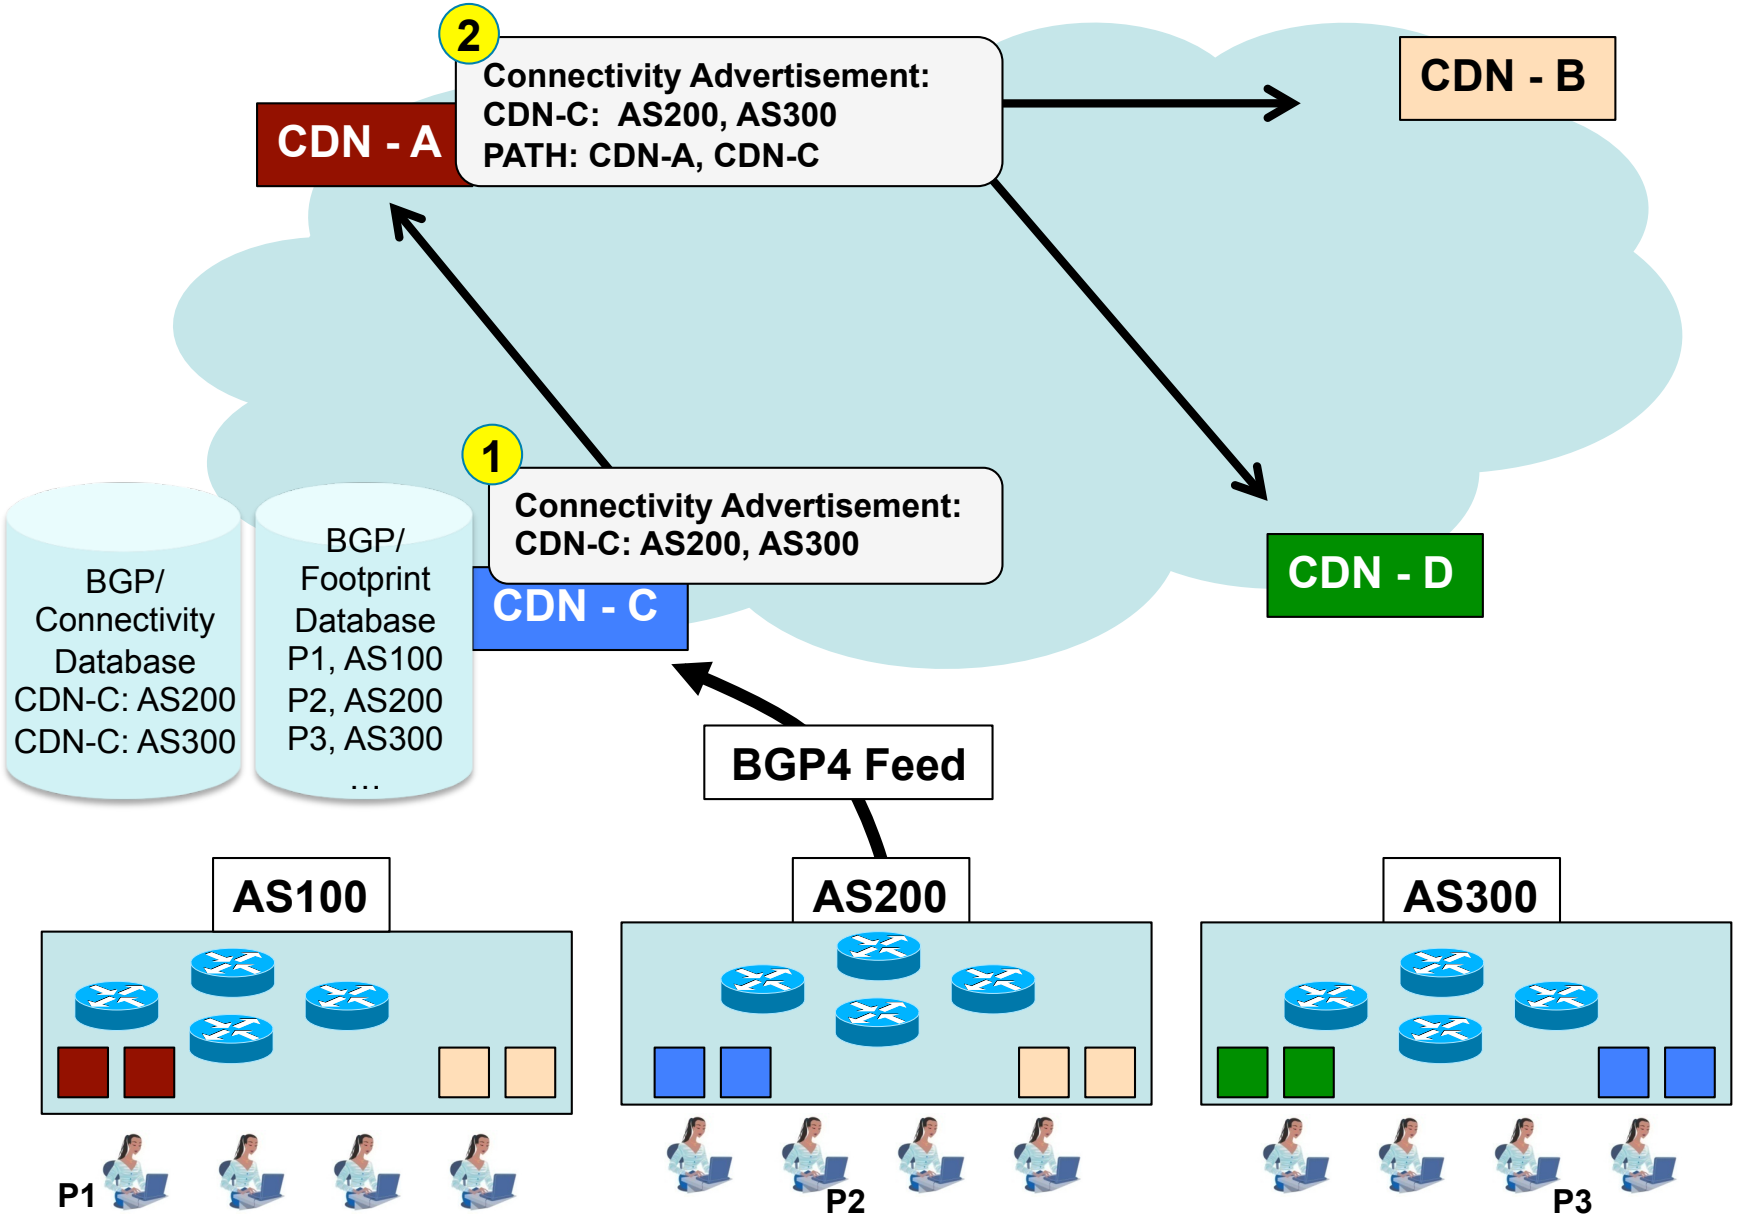

Workflow

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

To Do List

- Describe MP-BGP Message details
  - Define CDNI Address Family
- Describe scenarii where Footprint information is to be advertised
  - E.g.: when a CDN has finer granularity visibility than the one available in the BGP4 Internet table
- Define Footprint information details
  - E.g.: original AS_PATH
- Define Connectivity Information Details
  - E.g.: AS_PATH Vs. CDNI_PATH

draft-previdi-cdni-footprint-advertisement-00

Thank You

Footprint Information

- The right tool for the job: Multiprotocol-BGP
  - Proven over the years to be scalable and efficient
  - Multiprotocol allows separate Address Families: CDNI Address Family
  - Total separation between Internet-BGP and CDNI-BGP
  - BGP is already used for propagating connectivity _and_ capabilities: Communities