

# Routing Technologies for Non Routing Technologies

## The CDNI Case

Stefano Previdi – [sprevidi@cisco.com](mailto:sprevidi@cisco.com)

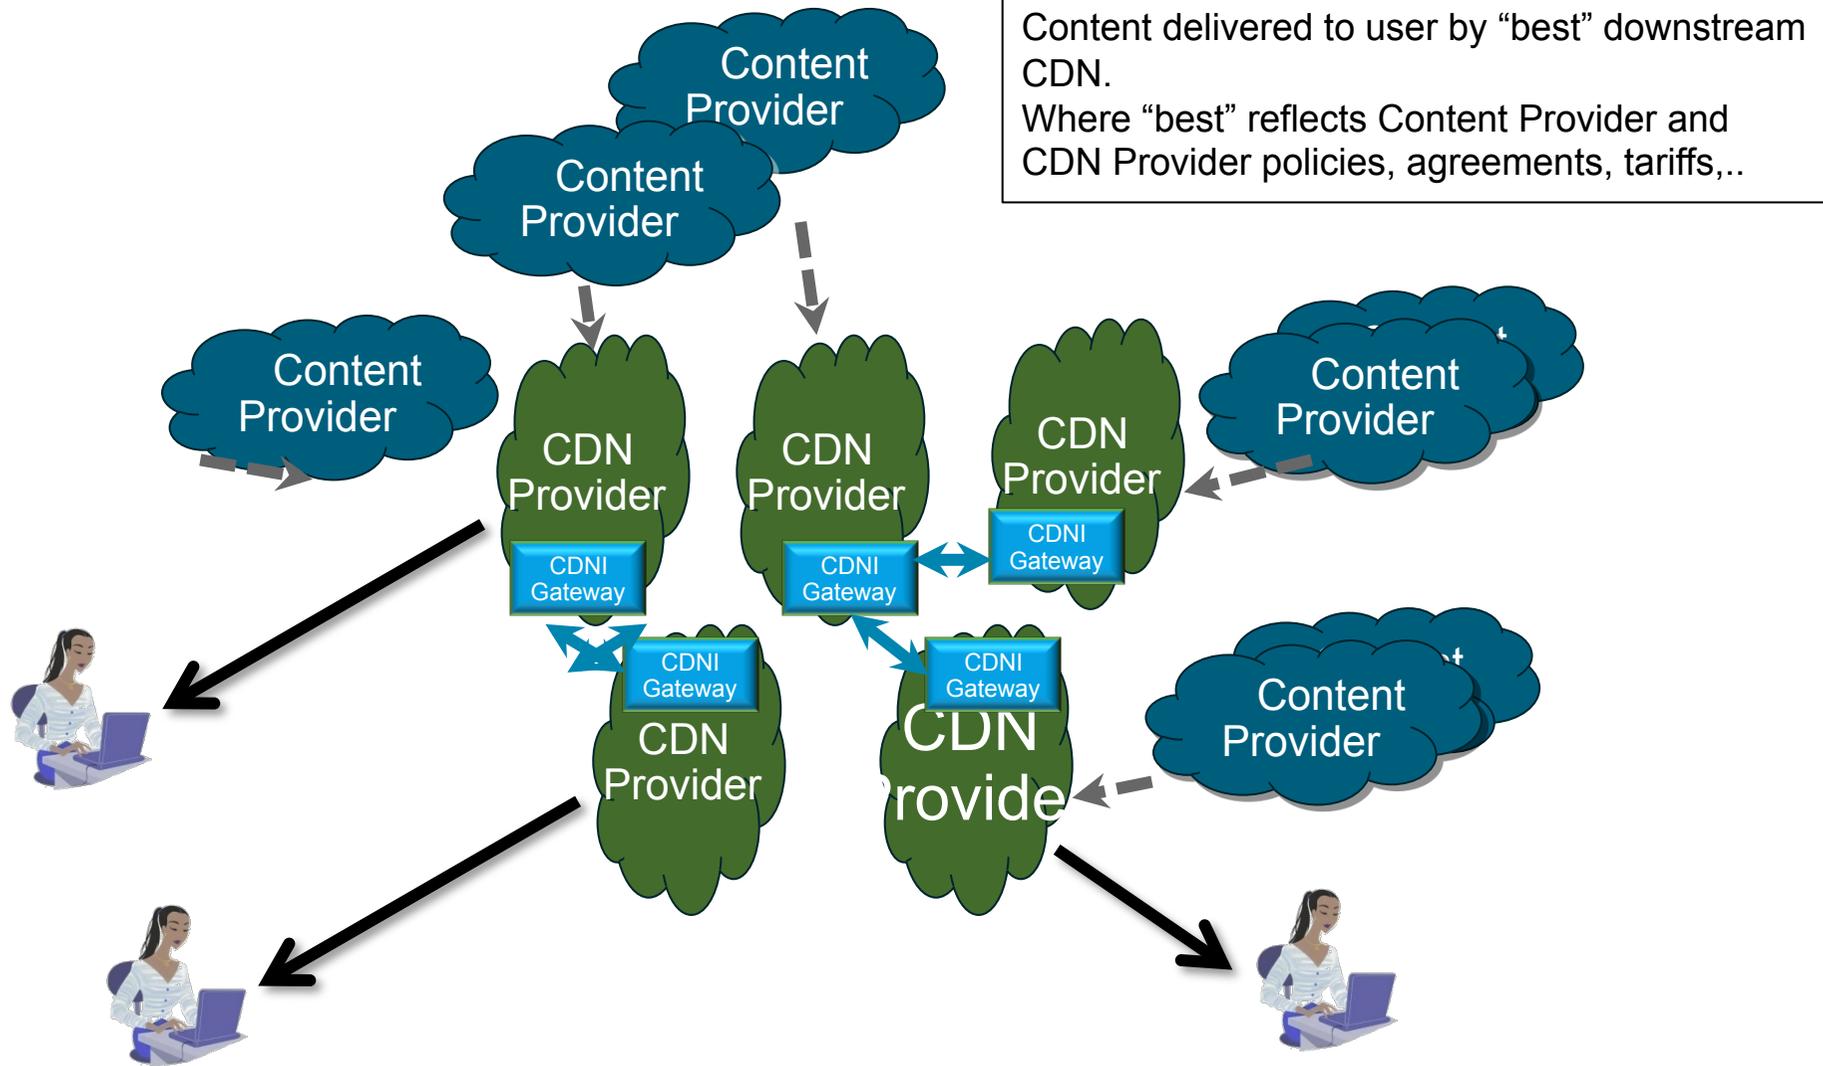
# Introduction

- Routing concepts and paradigms aim to answer a simple question:
  - Where this packet should go ?
- The question is intensively asked in layer-3
  - And brilliantly answered by routers...
- A similar (but NOT the same) question is asked in upper layers (application, content, cloud, ...)
  - E.g.: where this user should get content from ?
- Despite network layer business is “routing” and application layer business is “redirection”, it is undeniable that topology/infrastructure information does help in the process for both cases
  - E.g.: IGP shortest path, PCALC, FRR, ...
  - E.g.: Proximity Ranking, topology abstraction, ...
  - ...
- This presentation gives an overview on one case where this may be applicable

# Example: CDNI (IETF CDNI WG)

- CDNI: Content Delivery Networks Interconnect
  - CDNI is a CDN mesh/federation where each CDN has to answer a simple question:
    - Where this user should get the content from ?
- CDNI overview
  - CDNs form a mesh/federation
  - CDNs exchange information about their footprints and capabilities
  - An upstream CDN relies on a downstream CDN for content delivery
  - The downstream CDN is selected based on multiple criteria among which `_ONE_` is related to CDN footprint:
    - Where is the user ? Where is the content ? Is there anything valuable to know in the middle ?
- A CDN footprint consists on the set of prefixes the CDN is capable/willing/able to serve
  - In theory: the whole internet
  - In practice: the set of prefixes under “close” reach of CDN caches
  - “Close” is a floating value...

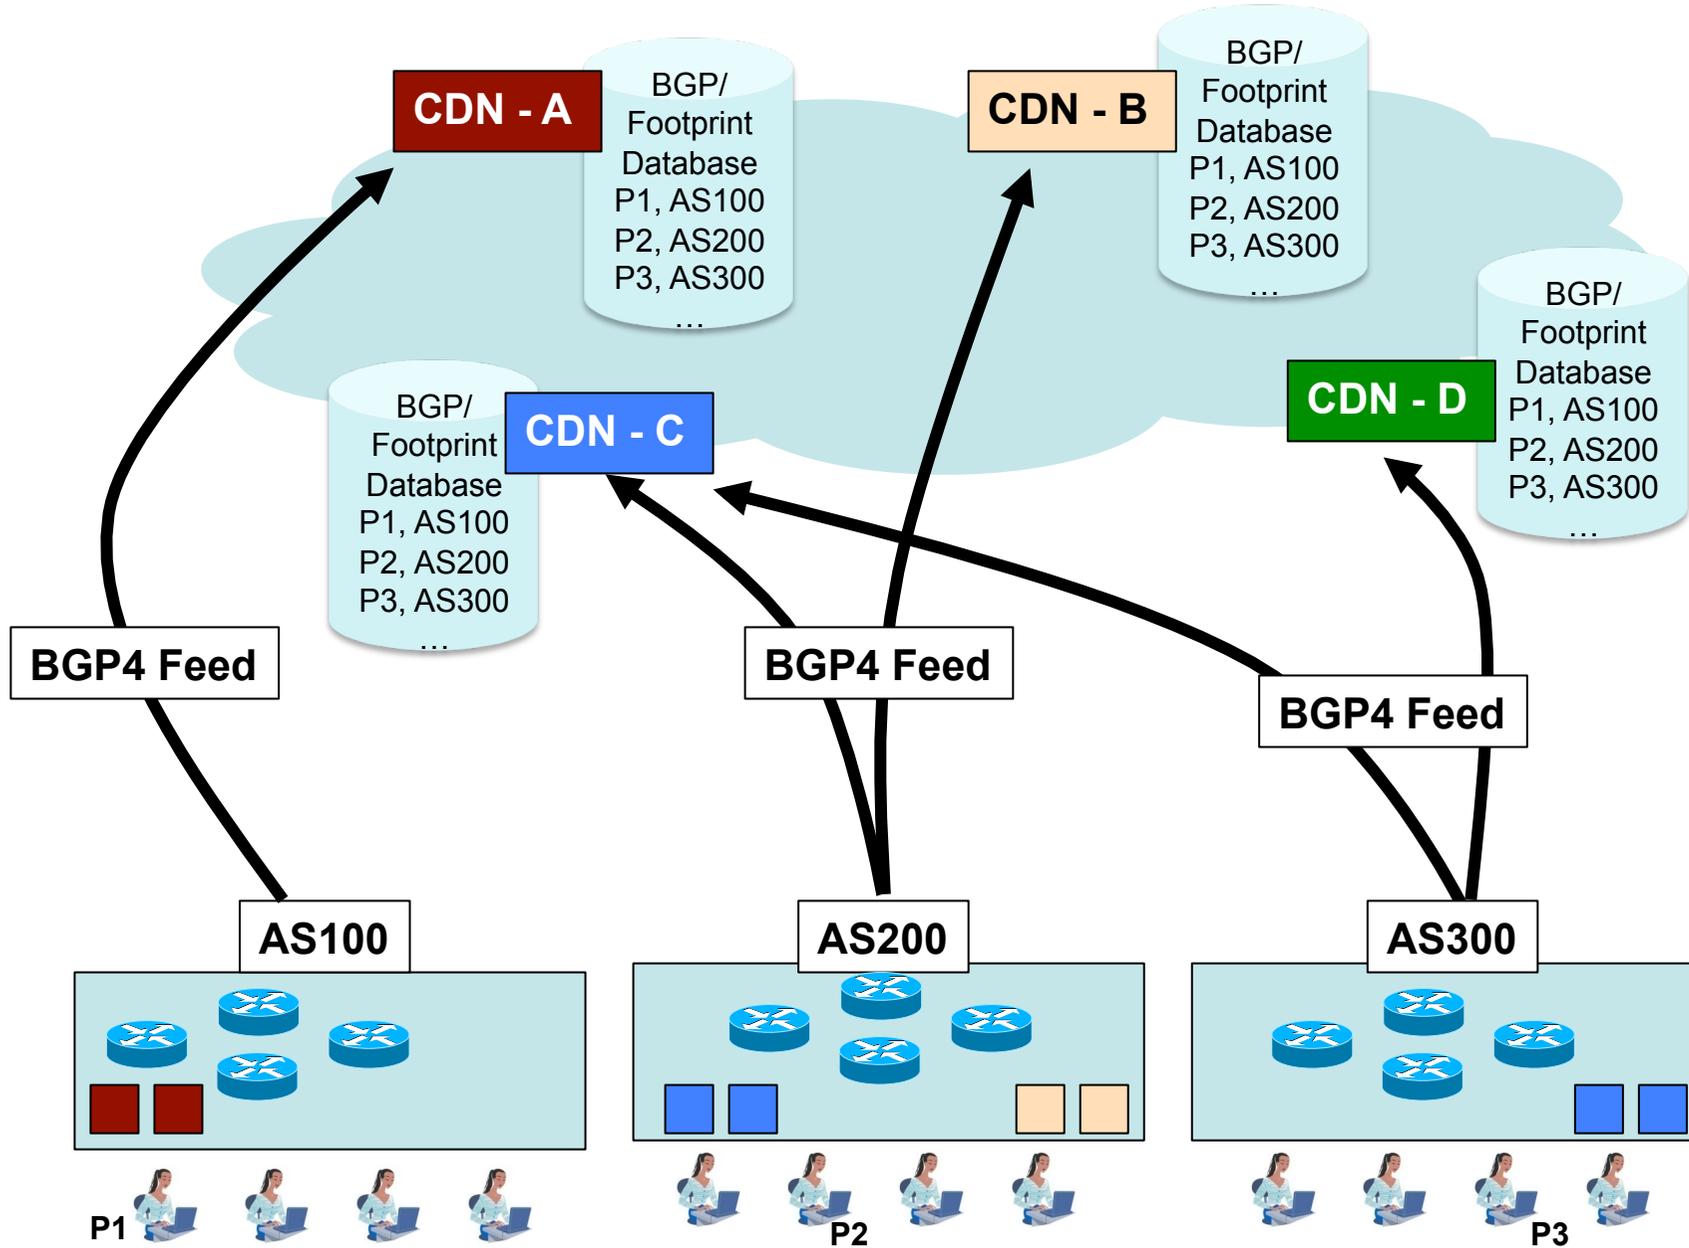
# Example: CDNI (IETF CDNI WG)



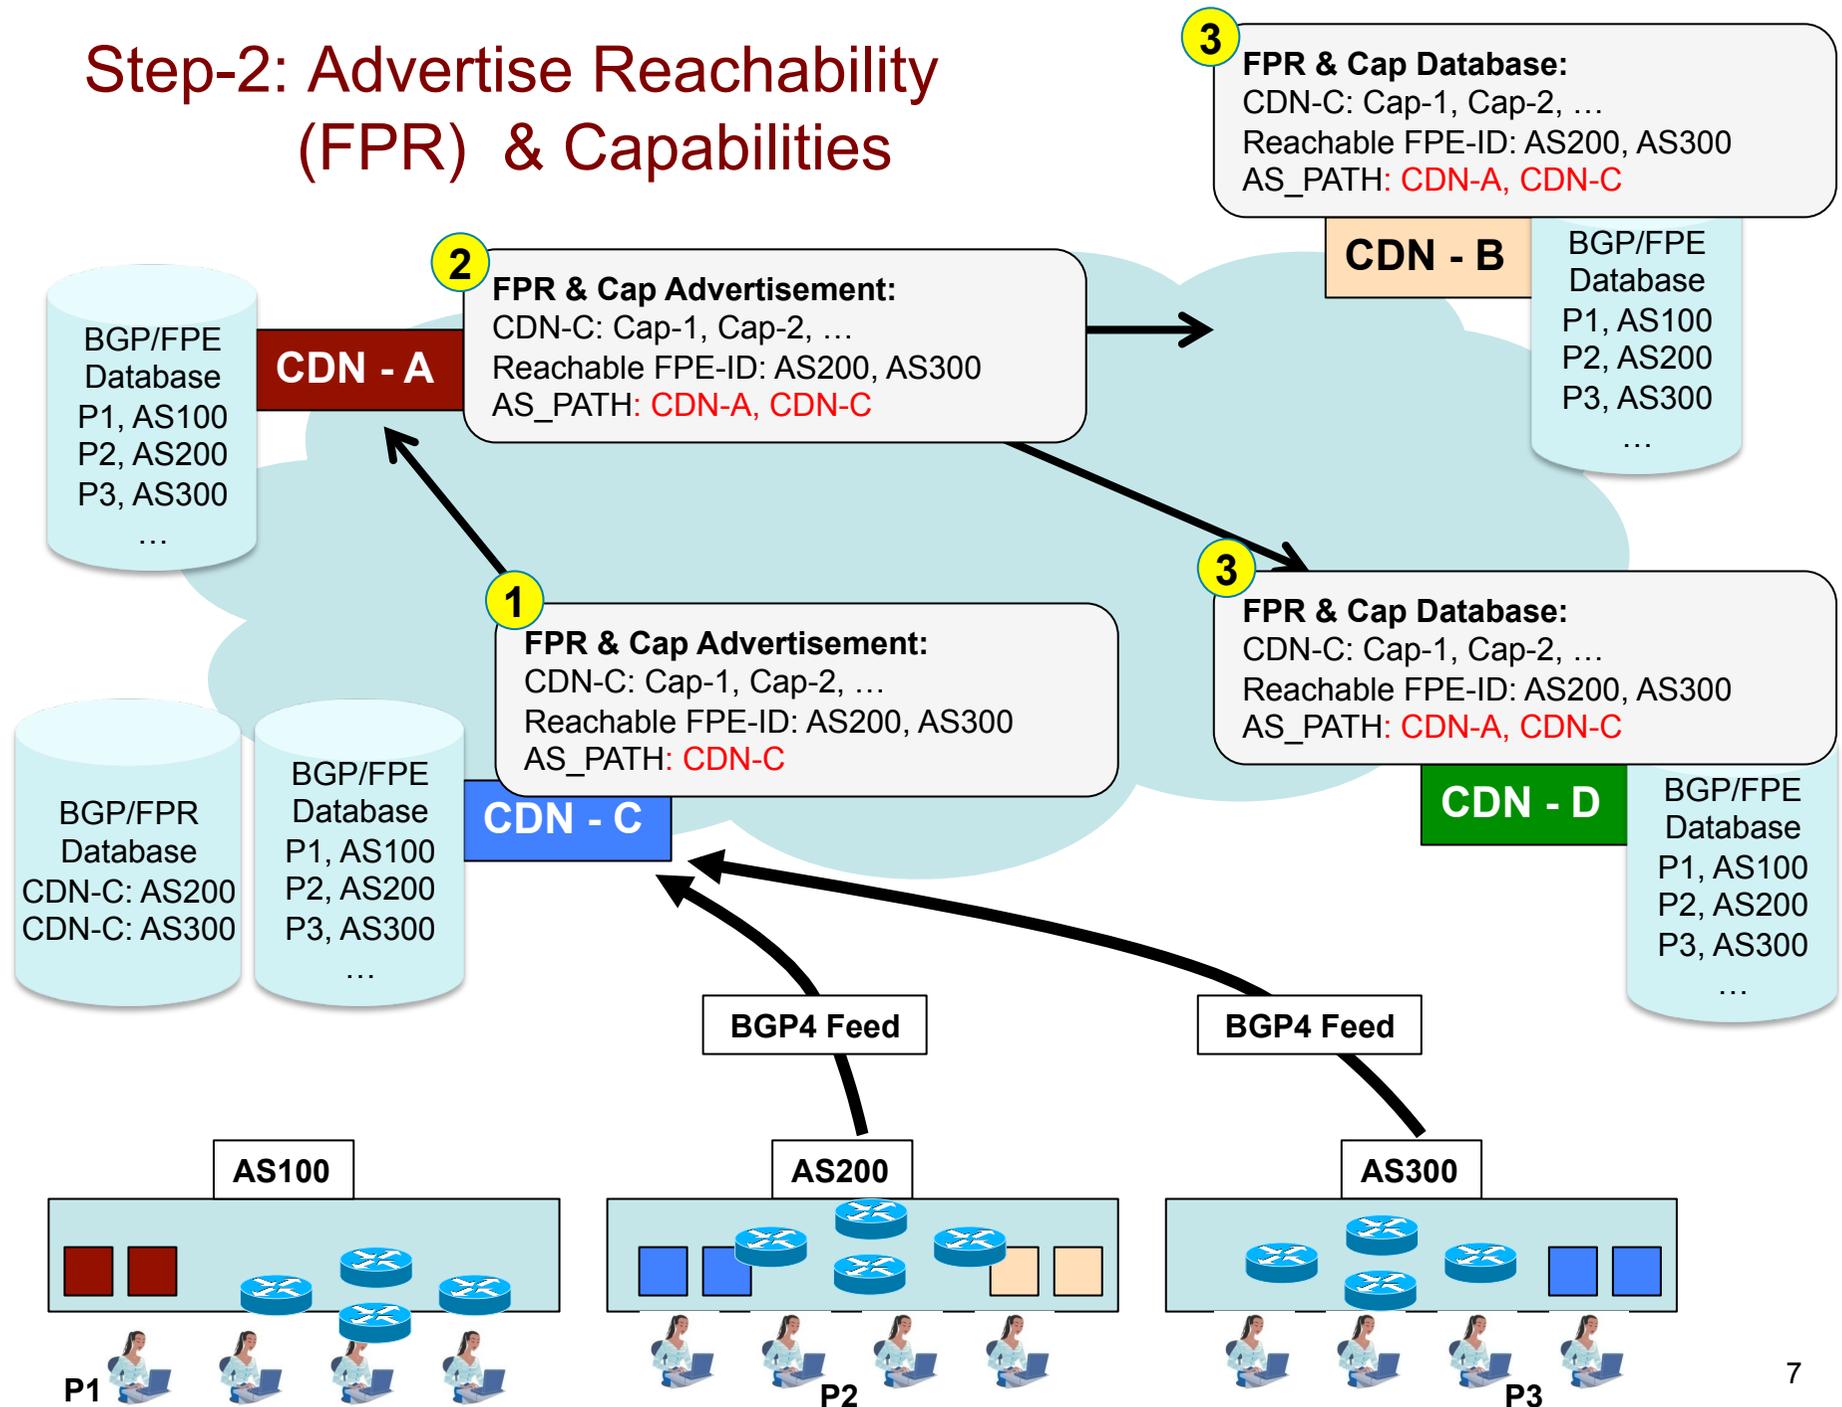
# Footprint and Footprint Reachability

- One proposal makes use of MP-BGP (with extensions)
  - Note well: proposal is under discussion in the CDNI WG and there's no consensus at this stage (go to slide 11 for details)
    - draft-previdi-cdni-footprint-advertisement
    - Under work: new attributes TBD, fix the encoding proposal, ...
- Footprint Element and Footprint Reachability concepts
  - Footprint Element (FPE): set of prefixes a CDN can “locally/best” reach
    - Can be inferred from BGP database if AS grouping is enough
    - Explicit advertisement of group of prefixes if necessary (grouping through communities would just work fine)
    - FPE gets an identifier
  - Footprint Reachability (FPR): FPE a CDN claims reachability to
    - FPR Advertisement: set of FPE Identifiers plus attributes

# 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

# What CDNI needs from (MP)-BGP

- New AF, new NLRI, new messages
- Additional selection rules
- Additional attributes... maybe
- Nothing substantially different from what has been done already
  - Yet Another BGP Extension...
- The difference is that we don't use the information for routing

# Other Examples

- BGP-LS
  - Distribution of topology information (optical, link-state, TE, ...) to whoever may have a use for it
- ALTO
  - Application Layer Traffic Optimization WG
  - Aims to provide network hint to applications
  - BGP-LS already plays a significant role in topology acquisition
    - In some implementations...

# Concerns, Issues, Experiences...

- Layers have a scope: isolation and separation... and it works pretty well: routing and application folks talk very little...

“Routing folks don’t understand we, application folks, have different requirements and their IGP/BGP/... stuff won’t work”

“Application folks don’t get that with IGP/BGP/TE/PCE/... you solve all problems”

“Don’t hack my routing protocol, you application fool”

Pick your preferred one...

# Conclusions

- Multiple use cases exist where routing technologies may help
  - CDNI, Cloud, Network Proximity, ...
- Some cases require protocol extensions
  - BGP-LS, BGP-CDNI, ISIS/OSPF Metric Extensions, GenApp, MI, ...
- Obviously:
  - (Re)using routing technology means NOT fate sharing
- Efforts on:
  - Extend routing protocols so to cope with these use cases
  - New deployment guidelines for routing for application cases

Thank You