

# Integrating Routing with Content Delivery Networks

Brian Field, Jan van Doorn, Jim Hall  
Comcast

[See IEEE Infocom NOMEN 2012 workshop  
for additional details

<http://www.ieee-infocom.org/2012/nomen/>]

# Current CDN Technology Issues

## Current CDN Technology Issues:

- Vendor cache's aren't open
- Origins aren't considered part of CDN

## Useful CDN features:

- Single/open/unified control plane protocol:
  - Caches and origins
- Caches network aware:
  - Child cache->parent cache
  - Parent cache->origin
- CDN dynamically self-configures based on asset "hotness"

# CDN “hot” asset problem drives solution

- From CDN resource perspective, want to treat “foobar.com” differently from “foobar.com/olympics/”
- How do we solve this in the network space?
  - Aggs and more specifics
    - `24.0.0.0/16 -> 24.0.11.16`
    - `24.0.0.0/24 -> 24.0.1.4`
- Apply to content domain:
  - `foobar.com -> 24.0.3.17` !Cache's IP
  - `foobar.com/olympics -> 24.0.141.15` !Cache's IP

# MP-BGP for Intra-CDN signaling

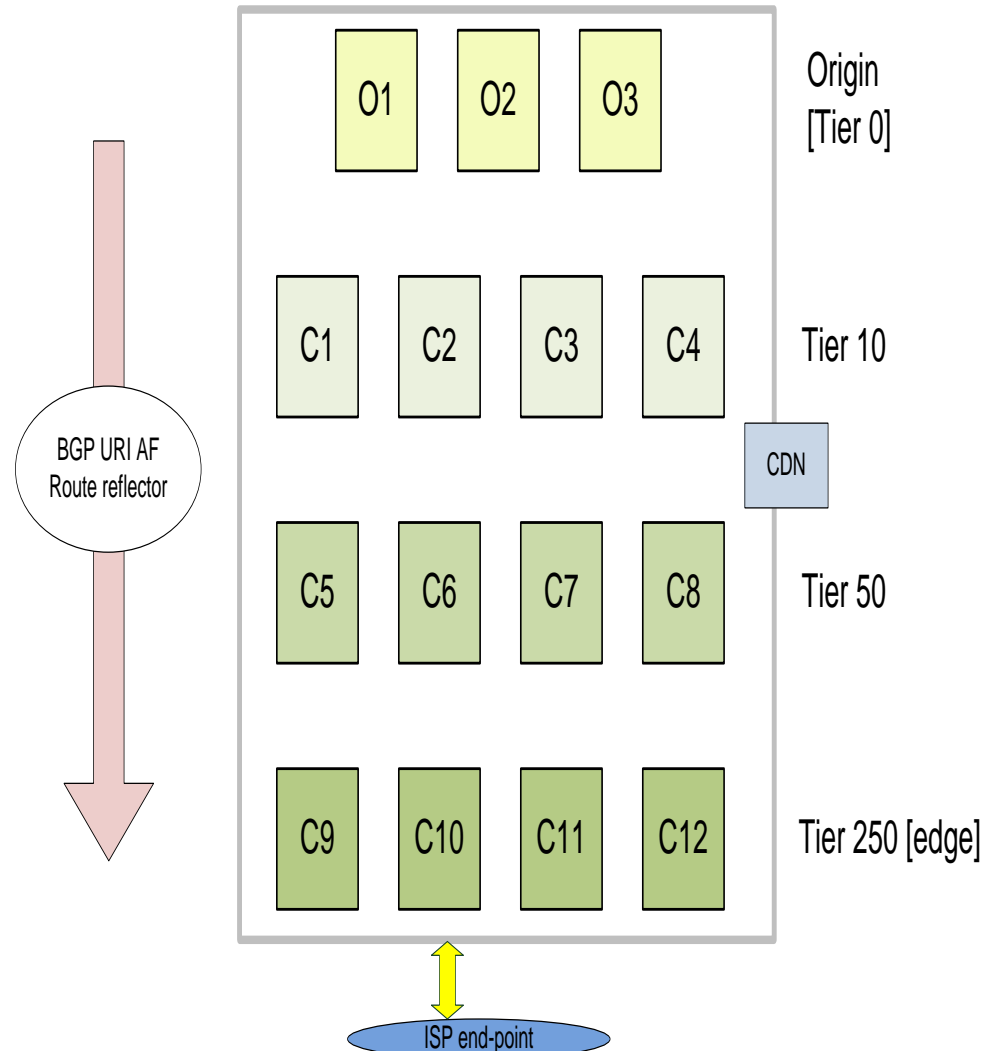
## [BGP URI AF is control plane inside CDN]

### URI AF NLRI:

- URI -> Cache/Origin NH IP
- “foobar.com” -> 24.0.32.11
- “foobar.com/olympics” -> 24.0.191.43

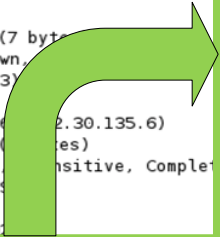
### Operation:

- Origins announce URI->NH
- CDN (caches)
  - process/select/propagate announcement down through CDN caches updating NH
- Caches perform standard prefix selection algorithm (on URIs)
- CDN is single ASN.




# BGP URI – working code on caches and origins

```
Frame 103: 162 bytes on wire (1296 bits), 162 bytes captured (1296 bits)
Ethernet II, Src: JuniperN_78:2f:f0 (00:22:83:78:2f:f0), Dst: Vmware_7e:01:5c (00:0c:29:7e:01:5c)
Internet Protocol Version 4, Src: 172.30.134.38 (172.30.134.38), Dst: 172.30.135.2 (172.30.135.2)
Transmission Control Protocol, Src Port: b2n (1179), Dst Port: 45934 (45934), Seq: 2764, Ack: 1325, Len: 96
Border Gateway Protocol
  UPDATE Message
    Marker: 16 bytes
    Length: 96 bytes
    Type: UPDATE Message (2)
    Unfeasible routes length: 0 bytes
    Total path attribute length: 68 bytes
  Path attributes
    ORIGIN: IGP (4 bytes)
      Flags: 0x40 (Well-known, Transitive, Complete)
      Type code: ORIGIN (1)
      Length: 1 byte
      Origin: IGP (0)
    AS_PATH: empty (3 bytes)
      Flags: 0x40 (Well-known, Transitive, Complete)
      Type code: AS_PATH (2)
      Length: 0 bytes
      AS path: empty
    NEXT_HOP: 172.30.135.6 (7 bytes)
      Flags: 0x40 (Well-known, Transitive, Complete)
      Type code: NEXT_HOP (3)
      Length: 4 bytes
      Next hop: 172.30.135.6 (172.30.135.6)
    COMMUNITIES: 64599:102 (4 bytes)
      Flags: 0xc0 (Optional, Transitive, Complete)
      Type code: COMMUNITIES (5)
      Length: 4 bytes
      Communities: 64599:102
    MP_REACH_NLRI (47 bytes)
      Flags: 0x80 (Optional, Non-transitive, Complete)
      Type code: MP_REACH_NLRI (14)
      Length: 44 bytes
      Address family: URI [Comcast experimental] (33)
      Subsequent address family identifier: URI to Next Hop (URI-AF) (128)
      Next hop network address (4 bytes)
      Subnetwork points of attachment: 0
      Cache layer reachability information (35 bytes)
        http://asn64599-2.com tier 100 -> 172.30.135.2 (33 bytes)
          len: 33
          tier: 100
          Next Hop Length: 4
          Next Hop: 172.30.135.6
          Number of Parent Caches: 1
          Parent Cache Length: 4
          Parent Cache: 172.30.135.2
          URI: http://asn64599-2.com (21 bytes)
      Network layer reachability information: 5 bytes
    Network layer reachability information: 5 bytes
```

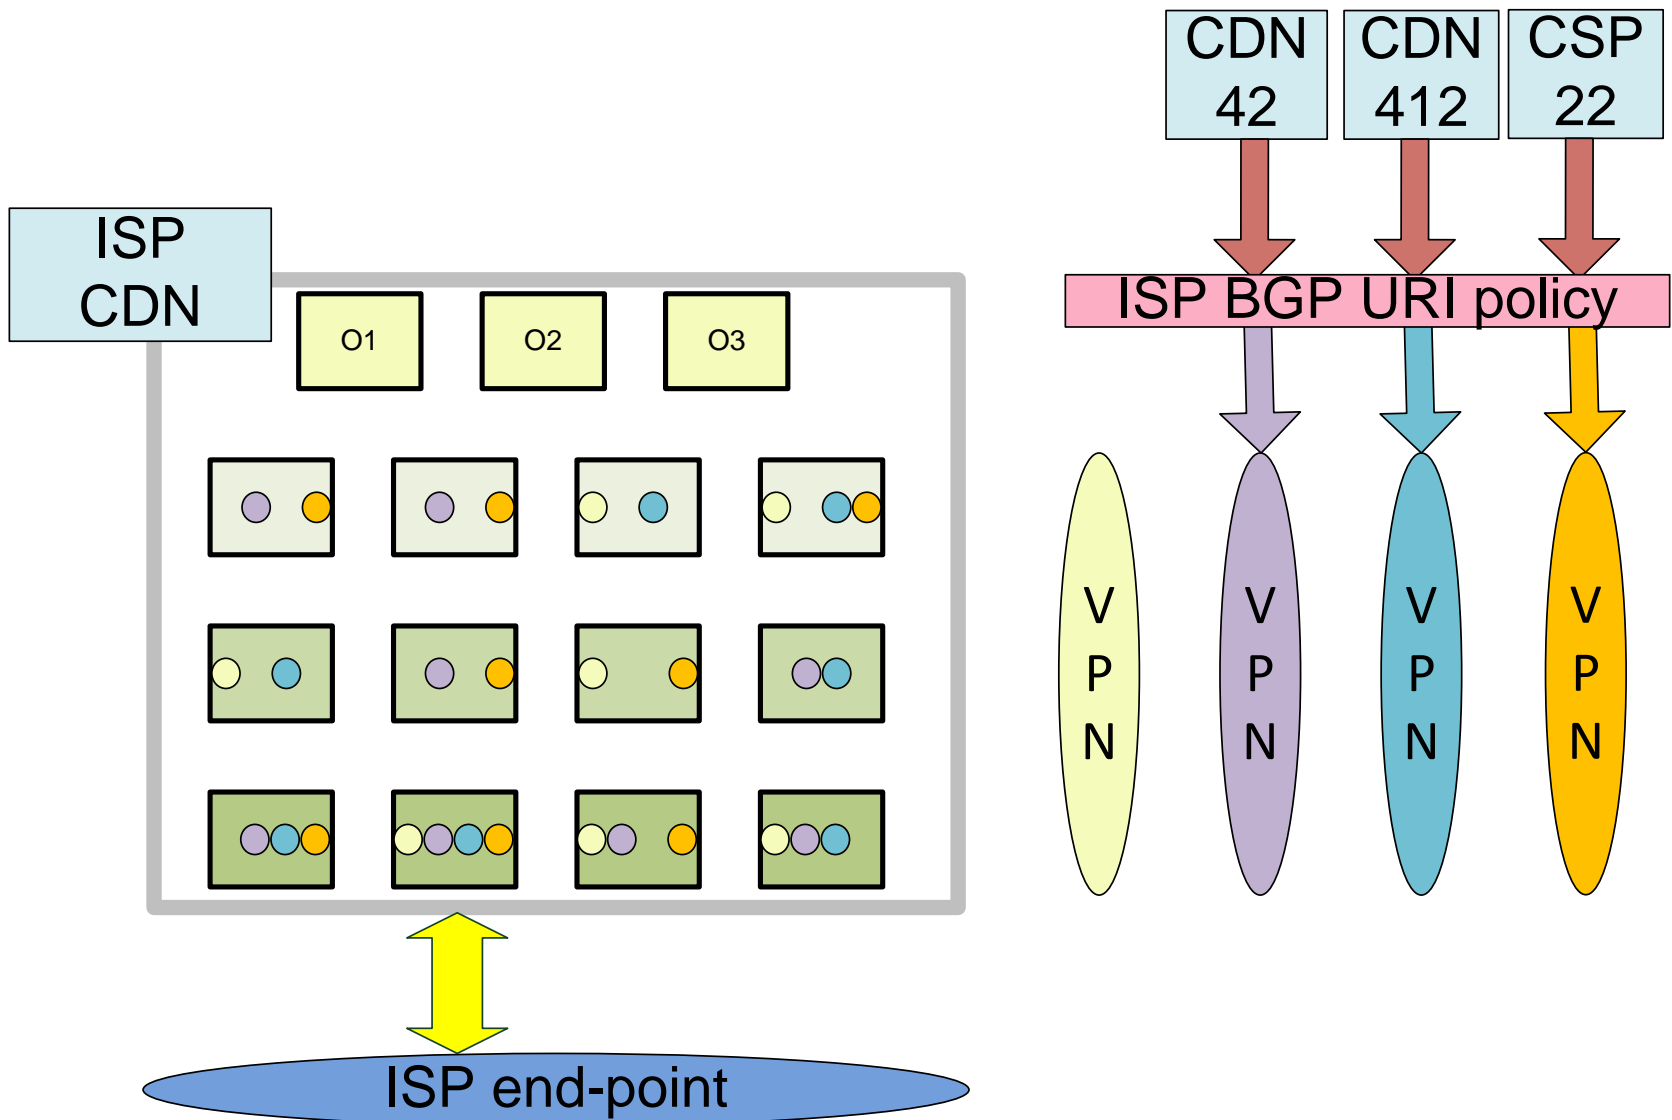


```
MP_REACH_NLRI (47 bytes)
  Flags: 0x80 (Optional, Non-transitive, Complete)
  Type code: MP_REACH_NLRI (14)
  Length: 44 bytes
  Address family: URI [Comcast experimental] (33)
  Subsequent address family identifier: URI to Next Hop (URI-AF) (128)
  Next hop network address (4 bytes)
  Subnetwork points of attachment: 0
  Cache layer reachability information (35 bytes)
    http://asn64599-2.com tier 100 -> 172.30.135.6 (33 bytes)
      len: 33
      tier: 100
      Next Hop Length: 4
      Next Hop: 172.30.135.6
      Number of Parent Caches: 1
      Parent Cache Length: 4
      Parent Cache: 172.30.135.2
      URI: http://asn64599-2.com (21 bytes)
    Network layer reachability information: 5 bytes
```

# Inter-CDN signaling via eBGP URI



# CDN VPNs ["multi-tenant"]



# Summary

## Proposal:

- iBGP URI:
  - Intra-CDN/Origin signaling
- eBGP URI:
  - Inter-CDN signaling
- Leverage all the existing BGP know-how and apply it to CDN
  - Skills, knowledge, policy logic, features, tools, capacity planning, code bases, etc.

## Benefits:

- ISP manages how content flows into their infrastructure from upstream CDNs
- Enables ISP to build own CDN in open vendor eco-system
- Enable ISP and Content sources to partner and develop cost effective delivery system via CDN VPNs



**Thanks!**