

INTERCONNECTING
CDNS
AKA “PEERING PEER-TO-
PEER”

Bruce Davie & Francois le Faucheur

Outline

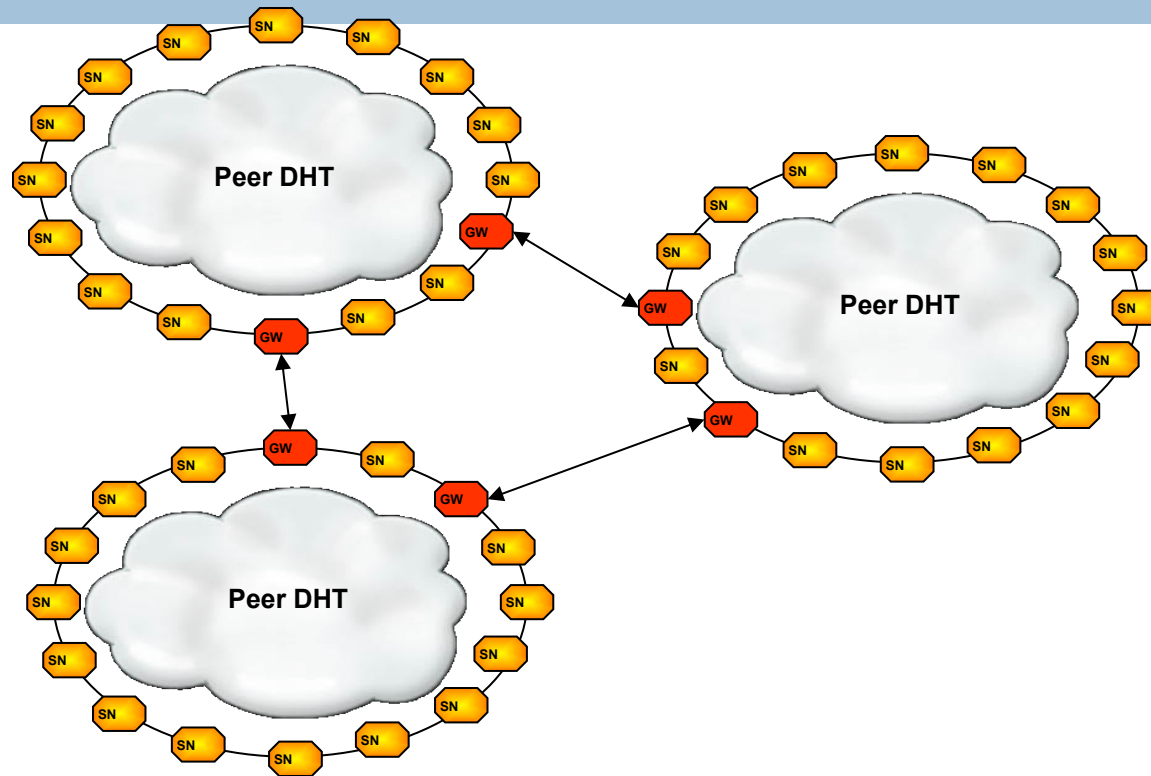


- Background: “peering peer-to-peer providers”
- CDN Interconnect Motivation
- CDNI Technical Issues
- Discussion

Why is this a P2PRG item?

- My interest in CDN Interconnect started with “peering peer-to-peer providers”
 - ▣ Balakrishnan, Shenker and Walfish paper from IPTPS 2005 proposed a model for interconnecting SP-operated DHTs
 - ▣ If you happen to build a CDN using a DHT infrastructure, then CDN interconnect looks a lot like DHT peering
 - ▣ Even without DHTs, lots can be leveraged from inter-DHT interface
- Likely need for standardization, but needs pre-standards work now

Peering DHTs



- Each AS/provider operates one DHT serving full keyspace
- Select nodes (peering gateways) can communicate across rings
- While each ring serves the entire DHT keyspace, not all content is in each ring

DHT interface



- A DHT provides a “put, get” interface
 - ▣ Put(key, value) stores value at location key
 - ▣ Get(key) returns value from location key
- This is roughly what OpenDHT provided as its API
- Also a reasonable inter-DHT interface
 - ▣ No requirement that internal implementation is a DHT
- You can build content delivery on top of this
 - ▣ use key to name the content (e.g. hash of a URI/URL) and value to store the content or a pointer to it

DHT Interconnect options



1. Broadcast Put

- When (k,v) is put into one DHT, the same (k,v) is put to all other DHTs
- Results in all descriptors being stored in all rings

2. Broadcast Get

- (k,v) is put in one DHT only
- $\text{get}(k)$ is broadcast to all DHTs
- content stored in original DHT, may be cached in others

3. Broadcast Put of Key Only

- (k,v) is put in one DHT only
- (k, DHT) is broadcast to all DHTs
- $\text{get}(k)$ can be forwarded directly to origin DHT

Towards Open Content Delivery

- Content Delivery is currently siloed into parallel, non-interoperable CDN “islands”
- A more open global Content Delivery architecture and infrastructure is desired:
 - ▣ To maximize QoE
 - ▣ To support wide range of business models (including a redistribution of revenue across involved parties that aligns better with respective costs)
- CDN Interconnect is an enabling technology for such an Open Content Delivery infrastructure

CDN Interconnect Vision



- CDN providers should be able to interconnect freely, as ISPs do today
 - ▣ Should support a wide range of “money flow” models
- Arguably, today’s big global CDNs are analogous to the walled-garden packet networks that preceded the Internet
- Hope to reap the same benefits that the Internet’s interconnection model brought to packet networks

Related Standardization Efforts

□ IETF

- Prior “CDN Internetworking” effort in IETF
- CDNI WG produced some info RFCs:
 - RFC 3466 A Model for Content Internetworking (CDI)
 - RFC 3570 Content Internetworking (CDI) Scenarios
- CDNI WG put on hold in 2003 (actual protocols not specified)

□ Open IPTV Forum (OIPF)

- CDN in scope, but left for Rel2, will probably not cover CDN Interco initially

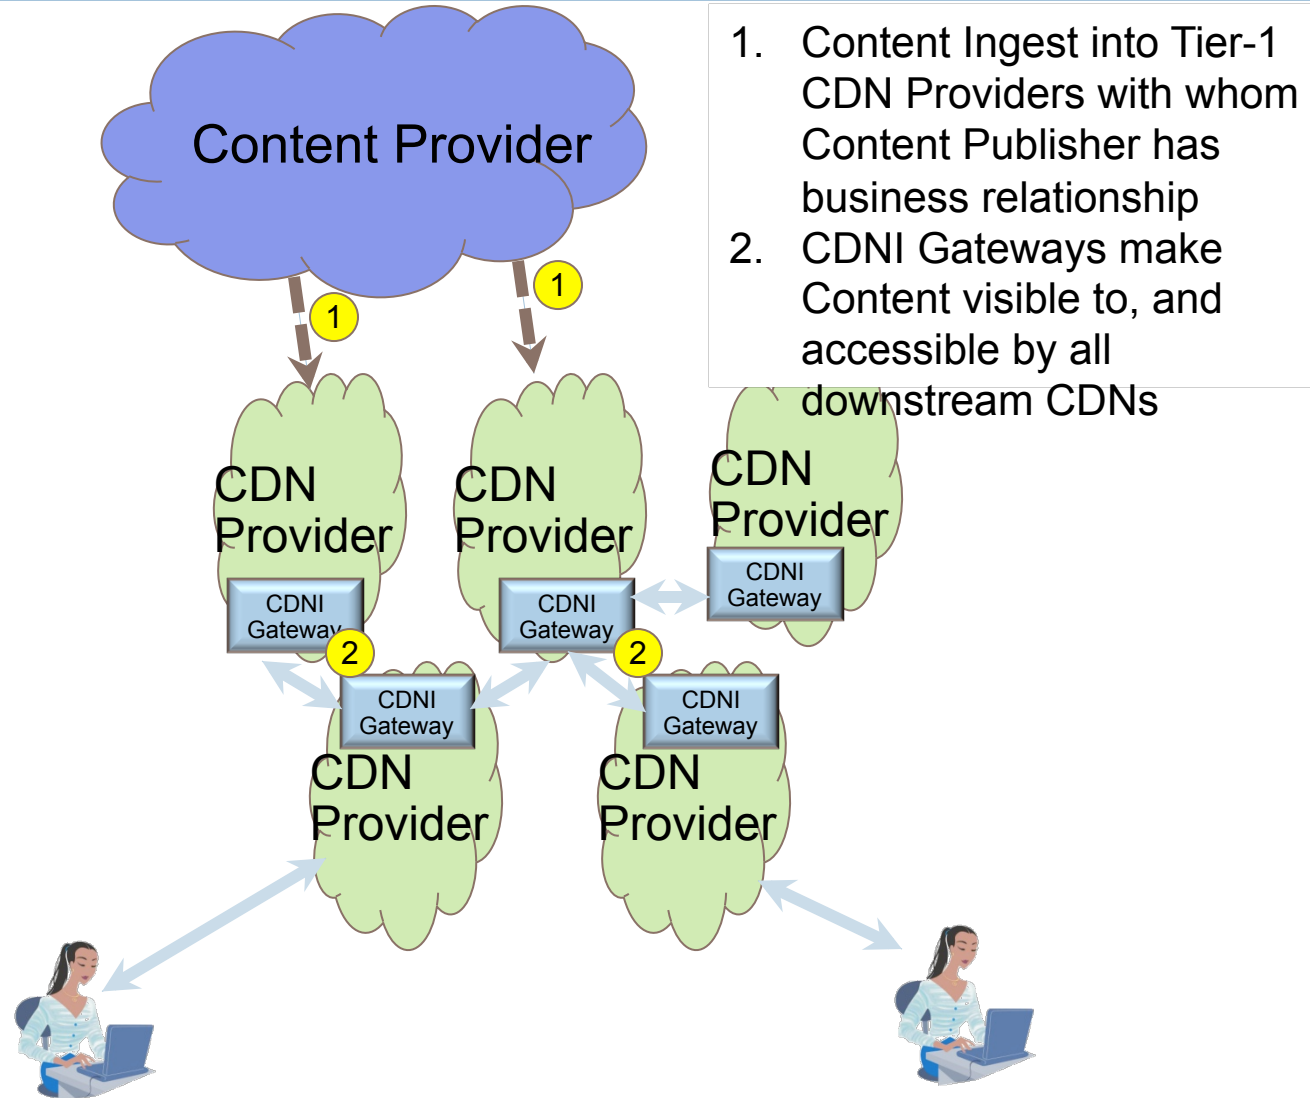
□ ETSI TISPAN

- Some work on CDN in scope for Rel 3, does not seem to cover CDNI

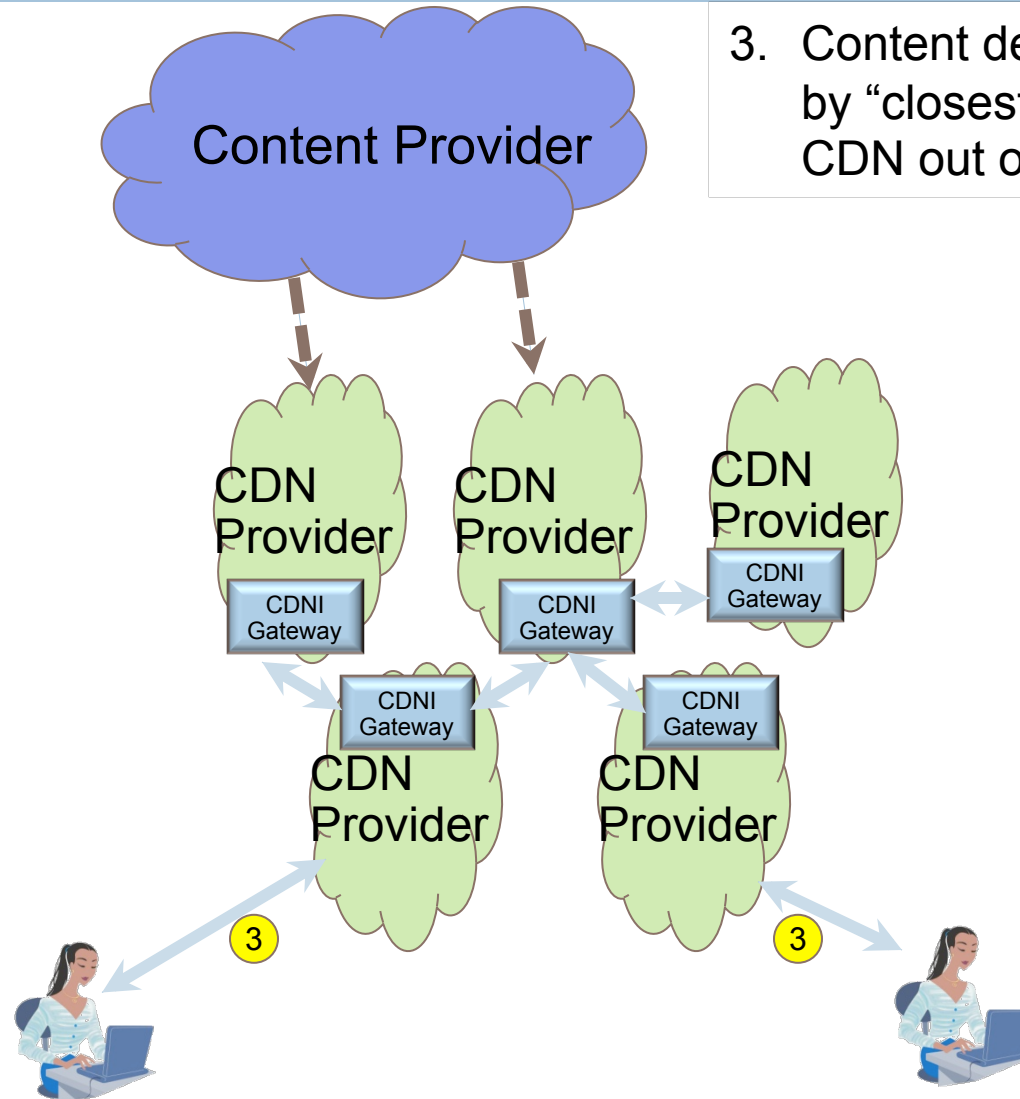
Towards Open Content Delivery

- CDN Infrastructure & Services being deployed by ISPs, telcos, Cable operators, Mobile operators,...
- Opportunity to Interconnect these CDNs to offer a compelling Open Content Delivery service
- Will allow Content Publishers to reach more users, with higher QoE, with fewer contractual relationships
- Will allow CDN operators to:
 - Monetize their infrastructure to deliver more content (e.g. from Content Publishers with whom they don't have a direct relationship)
 - Participate in a “Global” CDN
 - Act as “CDN aggregator” for Content Publishers

CDN Interconnect

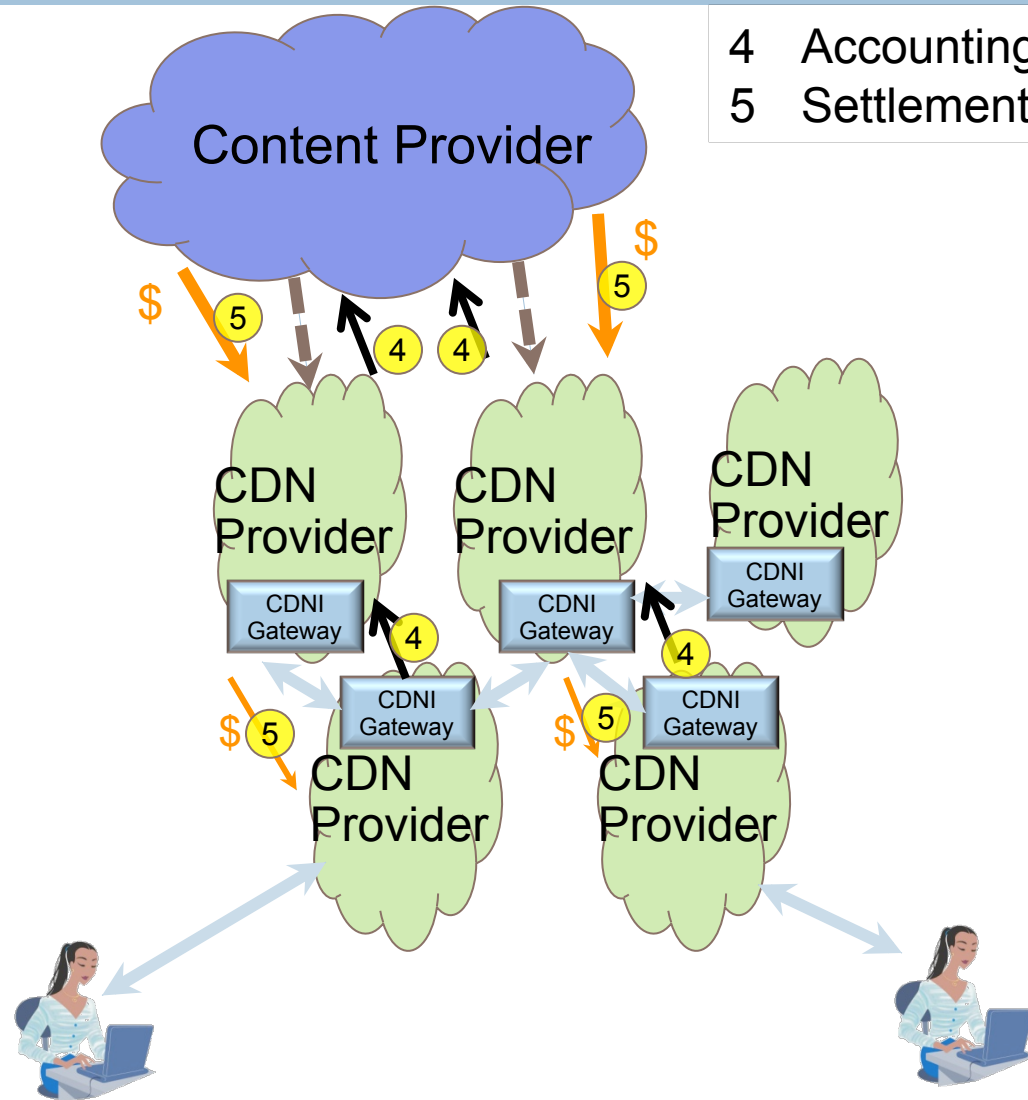


CDN Interconnect



3. Content delivered to user by "closest" downstream CDN out of many

CDN Interconnect



CDN Interconnect

Functional Components

- Request-Routing
 - ▣ How to steer user request towards right Surrogate in right CDN
- Content Distribution
 - ▣ How serving Surrogate acquires the asset through CDN Mesh
- Accounting
 - ▣ How volume of requests served by each CDN are recorded and used for settlement
- Reporting
 - ▣ How Content Publisher & CDN Providers can track serving activity (in their CDN and downstream) :
 - Near-Real time
 - Detailed Log

Request Routing

- There are a handful of ways to cause a client to fetch content from a given surrogate
 - DNS
 - HTTP redirect
 - Explicitly configured proxy
 - “transparently” intercept requests
- CDNI requires co-operation among CDN providers in this step
- Can think of this as two-phases:
 - Select the CDN
 - Select the surrogate in the CDN

Request Routing Requirements

- Content owner controls which CDN or CDNs are the “top level” CDNs
- Client needs ultimately to be directed to a “leaf” CDN that
 - ▣ Has the content, or can get it
 - ▣ Can deliver it with suitable latency
- Likely to be policies involved in CDN choice
 - ▣ e.g. use this CDN for clients in country X
- Within a given CDN, selection of the exact surrogate best done by that CDNs policies/ algorithms

Content distribution



- To get a piece of content that is stored in CDN A delivered by CDN B, those CDNs need a common name for the content
 - ▣ Is that a URL or something more specific?
 - ▣ The fact that URLs have embedded DNS names is a drawback
- CDN A either tells B that it has the content a priori (“put” model) or CDN B asks CDN A when it needs it (“get” model)
- In richly connected topology (think Internet AS graph) these puts and gets need to be routed

CDNI Accounting

- Each CDN needs to collect records (eg W3C Transaction Log) for each transaction it served incl:
 - ▣ Client IP
 - ▣ Start/stop time
 - ▣ Quality indicators (rate/resolution ...)
- CDN needs to (aggregate? and) export to PHOP CDN all records for assets associated with that PHOP CDN comprising:
 - ▣ Records for deliveries performed by that CDN (*)
 - ▣ Records for deliveries performed by downstream CDNs on behalf of that CDN

(*) with disambiguation between deliveries to an end-user vs delivery to the Downstream CDN

CDN Interconnect - Summary

- Set of technologies allowing many CDNs to operate as a “single big CDN”
- Content Publisher can leverage CDN infrastructure from all CDN Providers while only establishing relationship with 1 (or a few) Tier-1 CDN Provider(s)
- Need for standardized interfaces, redirect mechanisms, etc.
- Accounting + Settlement allows CDN Providers to get compensation proportionate to their contribution towards better delivery
 - ▣ Money can flow in multiple directions
- Should facilitate wide range of business models, not bake one in, e.g.
 - ▣ “PSTN Call Termination Model”
 - ▣ Per view, per user, per CDN
 - ▣ Settlement-free, etc.