Programme ANR VERSO

Projet VIPEER

Ingnierie du trafic vido en intradomaine base sur les paradigmes du Pair Pair

Dcision n° 2009 VERSO 014 01 06 du 22 dcembre 2009
T0 administratif = 15 Novembre 2009
T0 technique = 1er Janvier 2010

Deliverable 1.4
Final state of the art on business models for content distribution

Auteurs:
F. Guillemin, S.Moteau , Y. Le Louedec (Orange Labs)
A. Gravey (Telecom Bretagne)

Edited by: F. Guillemin (Orange Labs)

March 2013

Telecom Bretagne ; Eurecom ; INRIA ; France Telecom ; NDS ; ENVIVIO

Abstract

We review in this final state of the art document the ecosystem of video distribution as at the end of the Vipeer project. During the lifetime of this project, some major events have occurred. Today, Youtube is owned by Google and streaming traffic downloaded by end users has exploded in networks, giving rise to major crises between ISPs, Content Providers and Transit networks. Vipeer has proposed an original solution to remedying the saturation of peering links and quality degradation experienced by end users. Vipeer advocates for a collaboration between Content Providers and network operators and for the operation by the network operator of a distributed CDN to assist the Content Provider in the delivery of content. By performing trace driven simulations based on traffic measurements from Orange networks, we have shown that caching is an efficient solution to decreasing the amount of resources necessary on some critical links to transport content up to the end user. To be economically efficient for the network operator, the deployment of such a caching infrastructure should enable an improvement of the quality of experience of the end user so that more users subscribe a connectivity offer by the network operator. Otherwise, this would incur some additional cost via the increase of the mean bit rate consumed by the end user.

Keywords: caching, business model, trace driven simulation

Contents

Chapter 1

Introduction

Since the launch of the Vipeer project in late 2009, the situation in the composition of traffic carried by networks, the relationships between networks (in particular between ISP and Transit networks) and the ecosystem of Content Distribution Networks has changed a lot. The driving factor in this rapid evolution of the economical and technical environment is due to the ever growing weight of content dissemination in networks, especially due to video downloads by end users.

To illustrate this phenomenon, we show in Figure 1.1 the application breakdown on Gigabit Ethernet link in Bordeaux; similar traffic breakdowns can be observed in other ADSL areas in France. This figure clearly shows the prevalence of streaming in today's networks. Streaming traffic is mainly composed by Youtube traffic. This was not the case a few years ago but since the closing of MegaUpload in January 2012, Youtube has become the major contributor of streaming traffic in Orange's networks.

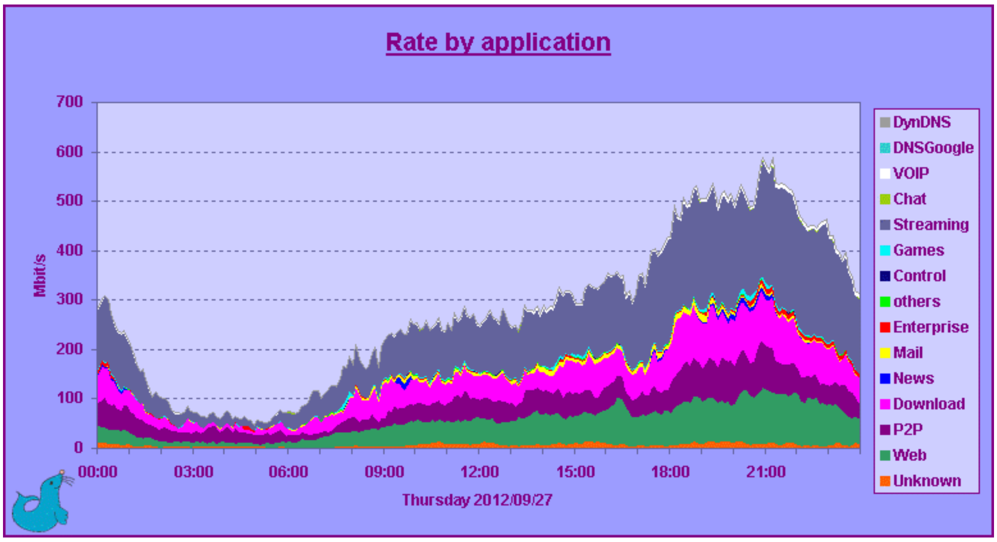

Figure 1.1: Typical network link utilization on a Gigabit Ethernet link in Orange's backbone network in France - application breakdown in Bordeaux.

The importance of streaming in the composition of traffic (about 300 Mbit/s during peak hours in Bordeaux) has a big impact on network performance since it leads to the saturation of peering links, which is turns degrades the end-to-end quality of other applications (e.g., web surfing). This situation has to led in the past few years to a deep modification in the relationships between all players involved in the content delivery chain. The key evolutions are reviewed in Section 2.

In the VIPEER project, caching has been advocated as an efficient solution to remedy peering link saturation and quality degradation. To verify this assumption (especially, for User Generated Content which is petted as being highly volatile), we present in Section 3 some trace driven simulations by using traffic measurements from Orange IPA networks. We notably prove that caching is efficient even for UGC content delivered by Youtube. While caching can be very efficient to reduce the amount of traffic of those files which are the most popular, caching may however have adverse effects on the business relationships as discussed in Chapter 4.

With the generalization of caching systems, new opportunities are offered to operators to create new business via CDN interconnection. While in VIPEER we have proposed a solution where the caching system is operated by the operator in isolation, some new possibilities can be envisaged with regard to the cooperation and the interconnection of caching systems. This point is notably addressed by the OCEAN project; some key principles recommended by this project are presented in Chapter 5.

Some concluding remarks are presented in Chapter 6.

Chapter 2

Recent evolution of the content delivery chain

2.1 The historical peering landscape

The traditional landscape of the content delivery chain is as follows: On the basis of the fundamental granularity based on Autonomous System (AS), one AS can play one or several out the following four different roles:

- the Internet Service Provider (ISP), which offers Internet access to residential customers,
- the Internet Backbone Provider (IBP), which offers transit to ISPs in order to ensure global Internet connectivity
- the Content Provider (CP), which offers various content over the Internet to residential customers
- the Content Distribution Network provider (CDN), which delivers content over the Internet on behalf of CP.

Less than 10 years ago, video content distribution had not yet reached its current part of Internet traffic (see Introduction), and the commercial relationships between AS were fairly simple:

- ISPs had to pay IBPs for transit in order to provide full Internet connectivity to their customers, or alternatively "peered" at no cost with other ISPs;
- CPs had to pay IBPs for transit in order to be able to deliver their content over the Internet. Alternatively, they relied on CDNs to deliver their content;
- CDNs had to pay IBPs for transit in order to deliver content over the Internet on behalf of their CP customers;
- IBPs interconnected in various manners, depending on their respective sizes and spans. Typically, similar IBPs (in terms of size and span), who exchanged roughly the same amount of traffic, peered at no cost. A small IBP would on the other hand pay a larger IBP for transit. All types of intermediate situations were also observed, where traffic was segmented between on-net and off-net, and where only traffic exceeding a given previously accepted incoming to outgoing ratio would have to be paid for.

In order to offer transit to its customers, an ISP has to choose either to peer with other AS, or to buy transit from IBPs.

Similarly, in order to deliver its content, a CP has to choose between buying distribution services from CDNs and ensuring its own distribution either by peering with interested ISPs, or by buying transit from IBPs.

2.2 Cost Computations

The above choices are met by the various actors based on several criteria, one of them being the respective costs of the options.

2.2.1 Transit versus Peering

It is important to be able to compare the respective costs of peering and transit. Both costs may depend on the considered regional area, although they seem to be rather similar in America and Western Europe for large cities [1].

Usually, the cost of transit computed at an IP interconnection point (IXP) is purely an operational cost; the IBP measures the actual transit rate and charges for the 95th percentile of the achieved rate [8]. The transit cost paid to the IXP is highly granular, as it is based on monitored amounts of traffic.

However, delivering traffic to and backhauling traffic from an IXP is *not* free, as the network operator has to invest in transport capacity and in equipment to be located at the IXP. Moreover, the ISP has to participate in the operational cost of the IXP, which can be translated in terms of a monthly rental fee which covers a global co-location fee, some rack space and peering ports.

As peering corresponds usually to a free exchange of traffic at IXPs between two ASes, the cost of peering is reduced to the cost of transport to and from the IXP, the cost of the equipment installed at the IXP and the monthly rental fee to the IXP.

Both transport and IXP co-location costs present a very coarse granularity. For example, transport can be provided by a global 10Gbit/s link and the equipment located in the IXP (the router and switches) can also be shared by all traffic carried between the ISP and the IXP, whether this traffic is handled by peering or by transit.

2.2.2 Direct distribution versus CDN based distribution

A CP has 3 possibilities for distributing its content.

1. the CP can buy transit from an IBP which delivers the content over the Internet; video servers are located close to the CP;

2. the CP can peer directly with its major ISP destinations (the so-called *eyeball* providers) in which case he also may locate some of its servers close to these destinations in order to both limit the amount of peered traffic and to decrease latency;

3. the CP can buy video service distribution from one or several CDNs, which distribute their own servers over the Internet.

Cost computation in cases 1 and 2 are very similar to the ones described in the previous section, although in case of the CP the outbound to inbound ratio is likely to be very large, which may hinder the negotiation of free peering relationships with ISPs as it is shown in the next section.

Cost computation in case 3 is similar to the computation of transit cost, as it is based on a measure of the peak rate of distributed content. Of course, the unit price per Mbit/s is higher since it includes both transit and caching costs.

2.3 Recent peering disputes

2.3.1 The Comcast-Level3 dispute

Comcast is a large US cable operator with more than 20 Million subscribers. As an ISP, Comcast, till late 2010, was using Level3 as IPB in order to allow traffic from its own customer to transit to the Internet. As Level3 was playing the role of a transit provider to Comcast, the latter had a commercial agreement with Level3 under which it paid Level3 for this transit service. Actually, the peering agreement between Comcast and Level3 was possibly a bit more complex : real peering for on-net traffic (i.e. traffic exchanged between Level3 and Comcast that originated from one of the 2 networks), and a commercial relationship for Comcasts transit traffic.

Netflix, as a CP, used Akamai and Limelight, 2 CDN companies, to deliver its VoD service over various ISPs. Both CDNs had commercial relationships with Comcast in order to deliver the video traffic to the Netflix users who were using Comcast as an ISP. Under these commercial relationships, both CDNs paid Comcast for locating their servers deep inside Comcasts network.

Level3 started its own CDN service and succeeded in getting Netflix signing with Level3 for CDN service.

What this meant for Comcast was first loosing the business it had with the previous CDNs. It also meant having to support all Netflix generated traffic over its interconnections with Level3. Depending on its routing policies, this could indeed mean a large increase in supporting this traffic: instead of being served by servers deep inside its network, and close to the users, the videos would all have to transit through its peering links with Level3, and moreover be supported from this interconnecting points high in the network to all users deep in the network. Some figures can illustrate part of the cost: roughly 20 supplementary 10 Gbit/s access points would be needed, each point corresponding to 50k$ CAPEX, and 25k$ OPEX [13].

The dispute was thus presented by the two parties as follows : Comcast explained that the conditions regarding its free peering with Level3 had to be revisited, as the ratio between ingoing and outgoing traffic (expected to reach 5:1) would now largely exceed the current 2:1 ratio accepted as a reasonable measure of roughly equal incoming and outgoing traffic. It thus demanded Level3 to negotiate a commercial agreement where Level3 paid for this increase in incoming traffic.

Level3 refused this argument, stating that Comcast had to pay for this increased transit, as Level3 operated as a transit operator whereas Comcast operated as an ISP. Any supplementary cost could be redistributed directly on Internet users. Level3 also accused Comcast to violate the newly FCC passed regulation regarding Net Neutrality, as by charging Level3, it indirectly charged Netflix, thus being able to favor its own in-house VoD service.

Two months later, the FCC ruled that the dispute was indeed purely commercial, and that Net Neutrality was not the issue. The FCC then did not rule against Comcast, which was clearly in its favor, as it was indeed in a stronger commercial position than Level3: Netflix customers on Comcast would not be able to receive a good quality delivery of their video services unless Level3 increased its peering capacity with Comcast. Netflix would thus be in a position to revisit its commercial agreement with Level3 and possibly revert to using another CDN provider.

One notes here that the dispute is based on the exact roles of both ComCast and Level3:

- Comcast, as a huge ISP, is large enough to impose peering relationships with IBPs, and presents itself as an IBP (although it is also clearly a dominant ISP);
- Level3, in order to increase its own revenues, presents itself both as a CDN (as it distributes Netflix content) and as an IBP (by claiming that Comcast should fund its increased transit requirements).

2.3.2 The Cogent-Orange dispute

In May 2011, Cogent, a large IBP seized the French Competition Authority on its dispute with France Telecom which operates both as an ISP (Orange) and as an IBP (Open Transit). Both Cogent and France Telecom are very large, and cover the whole Internet world; they are therefore considered as *Transit free, or level 1* IBP and are interconnected as peers.

However, Cogent is used by many CP, which translates into a large outgoing to incoming ratio with France Telecom, significantly larger than the 2.5 limit used by France Telecom in the framework of its peering agreements. Therefore, when Cogent asked France Telecom to increase its peering capacity, the latter asked the former to renegotiate its peering agreement, i.e. to accept a *commercial* peering where Cogent would have to pay for the transit of its outgoing traffic in excess of the 2.5 outgoing to incoming ratio. This is the major part of the dispute, as Cogent refused these terms, stating that France Telecom therefore limited the access to the Orange subscribers its CP clients.

The French Competition Agency rejected Cogent's claim in September 2012 [4], stating that the dispute was commercial and that France Telecom did not have to provide more interconnection capacity for free.

We note here that the dispute is based both on the exact role of France Telecom (who is at the same time an ISP and an IBP) and on the evolution of peering relationships between IBPs that are either *content-heavy* as Cogent or *subscribers-heavy* as Open Transit.

2.3.3 The Free-Google dispute

A dispute similar to the previous one is observed between the French ISP Free and Google who distributes the Youtube content. Google and Free are currently peering but French customers have been complaining for several years that the QoS delivered by Free regarding Youtube distribution is severely degraded at peak hour.

On the one hand, Google claims that Free refuses to increase its peering capacity, whereas Free would like Google to partly fund this peering capacity increase. None of the actors have seized the French Competition Authority, and it has been left to a consumer's association to carry over this dispute to the Authority, which has not yet stated.

This last dispute is typical of a latent conflict between CP and ISP:

- CPs increase offered traffics and obtain more revenues, but refuse to directly participate to network investments, claiming that ISPs can directly charge the final customers in order to recover the increased CAPEX and OPEX;
- ISPs wish to revisit existing peering agreements passed between CPs and ISPs in order to compensate for large incoming to outgoing traffic ratios.

2.4 A consolidating Ecosystem

Video traffic (or more generally of real-time entertainment services on the Internet) is today a large, and increasing part of Internet traffic. According to [3] Internet video represents in 2012 more than 50% of residential traffic, and a sharp increase of busy-hour Internet traffic which is forecasted to almost quintuple between 2011 and 2016. In parallel, ISPs observe all over the world a rapid increase of Internet traffic volumes with a growing proportion of video traffic.

This predominance of video traffic has changed the peering landscape as follows:

- Direct peering between CPs/CDNs and ISPs is often implemented. For example, at the end of 2009, more than 60% of Google traffic was already distributed to ISPs through direct peering relationships [6]; avoiding transit, both Google and many ISPs could thus decrease interconnection costs. As a consequence, the transit market has shrunk although Internet traffic constantly increases.
- Transit and CDN costs have plummeted in the last few years [8]: from 50$ per Mbit/s in 2006, transit tariffs have dropped to 5$ per Mbit/s in 2010 and less than 2,5$ per Mbit/s in 2012. A large degree of consolidation has thus been observed, as only the largest actors could subsist.
- ISPs have to constantly increase the capacity of all peering links, whether direct peering links to CPs/CDNs, or links to IBPs in order to facilitate Internet transit. This has led some majors ISPs to build their own interconnection architecture (we have seen above that this is the case for both Comcast and France Telecom). This has also been the root of ferocious disputes as outlined above. Lastly, this paves the way to *in-network caching* and *CDN interconnection* which are both proposed in the VIPEER architecture.

Chapter 3

Efficiency of caching inside an ISP

3.1 Measurements

To illustrate the technical efficiency of caching in networks, notably Orange's networks, we analyze in the chapter the characteristics of Youtube traffic in Orange IP backbone network. This network connects residential customers as well as small and medium size enterprises to the Internet. Various technologies can be used at the access, namely FTTH , ADSL, xDSL, and in the backhaul (ATM as well as GigaEthernet). To perform measurements we use passive probes (named Otarie). Equipped with packet processing capabilities, these probes can monitor traffic flows of up to 50.000 customers. These probes are located sufficiently high in the network (namely between a BAS and the first routers of the IP network) so as to observe a significant number of customers in each monitored ADSL area.

The probes were configured to detect YouTube traffic and to generate a single line record per video viewed. This record includes an anonymized customer identifier, the server address, timestamps of the starting and terminating times of the transmission of a video flow, the associated volume in packets and bytes. The video is identified by a 64 bits ID present in YouTube HTTP request which seems to be a constant and unique identifier for a given video file.

One major difficulty when observing Youtube traffic in an ISP network is that the video can be transmitted over several TCP connections, even in the case of classical progressive HTTP download. In fact, client viewers are configured to download files in pieces via several TCP connections. This situation is not surprising in the case of adaptive streaming (Microsoft's Silverlight or Apple's HLS) where files are segmented into chunks transmitted over a few TCP connections. The segmentation of files into pieces or chunks may raise some issues for caching, especially for progressive download because a cache should act as some kind of proxy in order to be sure to get the complete video file. Tracking a unique TCP connection and storing the content transmitted over this connection is not sufficient to get complete video files. Moreover, different viewers may not follow the same download pattern. The difficulty of caching Youtube content may moreover become more complex in the future if content is encrypted. These facts has to be taken into account when implementing a caching policy by an ISP. In practice, a collaboration between content providers and ISPs would be more profitable for both parties. On the one hand, ISPs could store safe content and content providers would be sure that content is delivered with a high level of quality.

In order to remedy the problem of file segmentation, some pretreatment has been performed on traffic traces. We have specifically aggregated those pieces of Youtube content with the same ID carried by several TCP connections between the same IP addresses and with starting times sufficiently close to each other. This procedure allows us to reconstruct the information on individual video files and has provided results partially consistent with those obtained in the past few years for Youtube video files observed from the edge [5].

In the following, we consider measurements from 3 ADSL areas, namely in Lyon, Bordeaux and Paris. In addition, we present results for traffic aggregated by the 12 Otarie probes installed in the French backbone.

3.2 Popularity curves

In this section, we illustrate a salient feature of Youtube files. More precisely, we exhibit the fact that a large number of files are viewed only a few times (once or twice) but there is a chunk of files which are massively viewed. The popularity curve of those latter files can be well approximated by a Zipf law as illustrated in Figure 3.1 for Bordeaux; the same phenomenon has been observed in other ADSL areas (Lyon and Paris). In this figure, we have computed the popularity of Youtube files in one day, one week and two weeks and we have approximated the popularity curve by a truncated Pareto function of the form

$$f(x) = \mathbb{1}_{\{x_{\min} \leq x \leq x_{\max}\}} \frac{\kappa}{x^{\alpha}}. \tag{3.1}$$

Of course, the value of the parameters κ and α as well as the values of $x_{\min}$ and $x_{\max}$ may change, depending on the duration of the measurements and the location where measurements are performed. When $\alpha < 1$, the range $[x_{\min}, x_{\max}]$ is necessarily finite.

We can note from Figure 3.1 that except for the most popular files (the top 10 files), the popularity curves are rather stable in time and can be well approximated by a truncated Pareto function as in Equation (3.1) with coefficients

$$\kappa = 0.024004 \quad \text{and} \quad \alpha = 0.886537$$

These coefficients have been estimated by using the popularity curve of the seventh day. The Pareto approximation is valid for those files with ranks between 10 and 1000.

Now with regard to volumes, we have distinguished between files viewed only once and those viewed more than twice. To compute volumes, we have considered for each file the maximum volume observed during the observation period. This assumption is made to compute the volume, which would ideally be transmitted if all files were completely viewed, thus ignoring interruptions due to various factors, such as quality impairments. This is to estimate the volume of data which should be cached. The bias is that we take into account the maximum volume associated with the file even if the file exists in various resolutions.

Results for Bordeaux are reported in Table 3.1. It clearly appears that files which are viewed only one give rise to huge volume of data in two weeks. But, caching those files viewed more than twice may lead to an important reduction of consumed bandwidth.

3.3 Performance of caching

To investigate further the relevance of caching we have simulated by using data from Lyon, Bordeaux, and Paris the behavior of a cache server with 1 TB of storage capacity. The results

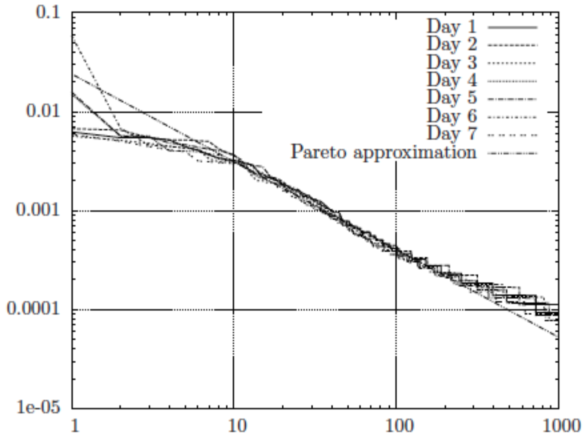

(a) One week.

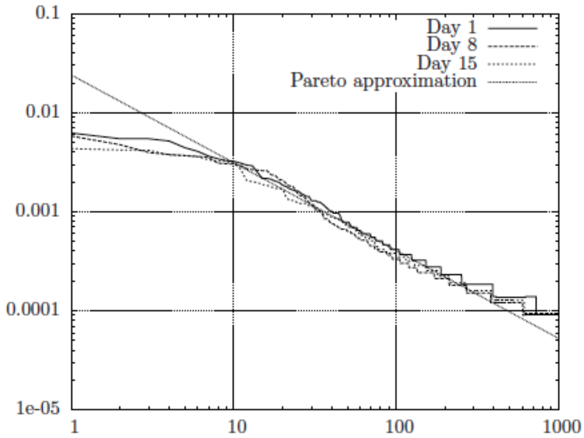

(b) The first days of three consecutive weeks.

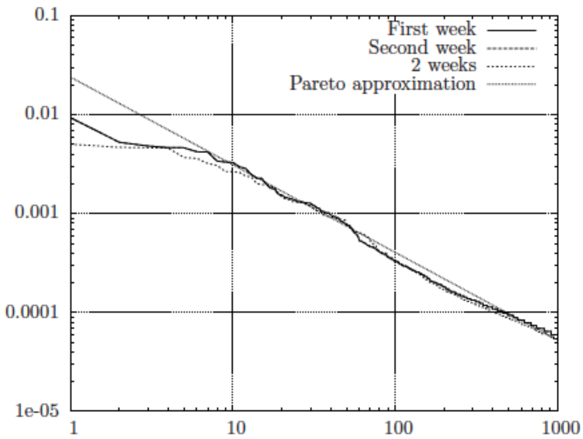

(c) Aggregation over the first and second weeks and over the two weeks.

Figure 3.1: Popularity of Youtube video files in Bordeaux.

Table 3.1: Statistics of downloads in Bordeaux.

	number	number of downloads	volume
Day 7			
Files	23,071	33,351	2,427.6 GB
More than twice	1,193	9,315 (28 %)	701.6 GB
Only once		19,720 (59.1 %)	1247.9 GB
First week			
Files	110,106	202,717	6239 GB
More than twice	10,510	90,203 (44.5 %)	2736.6 GB
Only once		86,678 (42.7 %)	2512 GB
Two weeks			
Files	208,289	425,266	9,354.8 GB
More than twice	24,291	215,158 (50.6 %)	4,462.7 GB
Only once		157,888 (37 %)	3,486.2 GB

are gathered in Table 3.2. We have given in this table the hit ratio in terms of files (i.e., the probability that a file is in the cache); the value is small because we have considered all files, including those which are viewed only once. When we consider those files which are viewed more than twice, we obtain the conditional hit ratio equal to 27.2 %, which is also rather weak. However, when we consider the byte hit ratio (i.e., the number of bytes served by the cache), we obtain rather large values, namely 74.67 % and 74.65 % for Bordeaux and Lyon, respectively. For Paris, we even obtain a larger byte hit ratio. This means that for Bordeaux, 104.1 TB have been served by the cache, and 119.6 TB in Lyon and 187.5 TB in Paris. We hence see that with a cache with medium size of 1 TB and running a basic Least Recently Used (LRU) replacement policy, we can achieve huge savings in terms of consumed bandwidth.

Table 3.2: Hit ratio and byte hit ratio for a cache of 1 TB in Bordeaux, Lyon, and Paris.

Location	Hit ratio	Byte hit ratio	Total volume
Bordeaux	5.8 %	74.67 %	139.4 TB
Lyon	5.9 %	74.65 %	160.2 TB
Paris	5.9 %	78.5 %	238.9 TB

The large byte hit ratios can be explained by file request pattern. In fact, those files, which are the most popular, are massively downloaded by end users. To illustrate this phenomenon, we have reported in Table 3.3 the ten most popular files in Bordeaux. We have given the hit ratio for the file, the number of times the file has been served by the cache, the number of times the files has been viewed and the size in bytes of the file. We see that these 10 most popular files represent 30.5 TB (21 % of the volume of all downloads). Hence, caching them already represent a significant gain.

Because many files are viewed in bursts, the cache memory is able to store all content and popular files are not pushed out of the cache by those files which are not popular (viewed only

Table 3.3: Hit ratio and byte hit ratio for a cache of 1 TB in Bordeaux

hit ratio	# served by CDN	# request	file volume
0.999833	11941	11943	1.2
0.999812	10650	10652	0.4
0.999429	3502	3504	0.4
0.999154	2363	2365	1.9
0.998985	1969	1971	0.7
0.998777	1633	1635	1.4
0.998751	1599	1601	0.5
0.998577	1403	1405	1.4
0.998569	1396	1398	0.7
0.998279	1160	1162	0.6

once or twice). In fact, filtering files does not improve the byte hit ratio even if it yields better (file) hit ratios. To illustrate this phenomenon, we have simulated a cache where files enter only if they are viewed more than twice. To implement this principle, we had to use a sliding widow to reduce the number of file identifiers to maintain. We have precisely implemented a table for recording those files which were active in a week. If a file is not viewed more than twice during a week, then the identifier is pushed out of the table. If a file is viewed more than twice in a week, the file is stored in the cache and the identifier is removed from the filter table. The results are reported in Table 3.4.

Table 3.4: Hit ratio and byte hit ratio for those files seen more than twice with a cache of 1 TB in Bordeaux, Lyon and Paris.

Location	Hit ratio	Byte hit ratio	Cached content
Bordeaux	73 %	67 %	12.4 %
Lyon	75 %	65 %	12.3 %
Paris	75 %	64 %	10.2 %

By letting enter the cache only those files which are viewed more than twice drastically increases the (file) hit ratio for this type of file. This ratio indeed jumps from 27 % to 73 % in Bordeaux. The counter part is that those same files have to wait before entering the cache. This leads to a small decrease of the byte hit ratio and then of the bandwidth saved by using the cache. Even this decrease is small in percentage, it corresponds to a non negligible proportion of the global volume. For instance, 91.9 TB are served by the filtered cache against 104.4 TB by a basic LRU cache out of a total amount of 160.2 TB in Bordeaux. The figures in Lyon are 103.6TB for the filtered cached against 119.6 TB by the basic LRU cache out of 160.2 TB.

3.4 Conclusion

By considering measurements from Orange's IP backbone network in France, we have shown that caching can efficiently reduce the amount of bandwidth needed to transmit files which are very popular. By considering measurements from Bordeaux ADSL are we have in addition proved that it is worth of distributing caching. As a matter of fact, contrary to what may have been imagined, Youtube files are massively downloaded by those end users connected to a same ADSL area. A cache of 1 TB can reduce by 75 % the amount of consumed bandwidth. The counterpart is that distributed cache servers will mainly cache the same content. In Table 3.5 we have accumulated up to a limit of 1 TB those files which are the most popular in Bordeaux, Lyon and Paris. We see that those cache servers contain up to 80 % the same files. But even caching the same content leads to saving bandwidth since files have not to be transferred from one cache server to another. Coordination between cache yields only a marginal gain.

Table 3.5: Content shared among servers located at different locations.

Location	Global	Bordeaux	Lyon	Paris
Global		831.5GB	847.3 GB	735.3 GB
Bordeaux	831.5 GB	166.9Gb	758.3 GB	686.1 GB
Lyon	847.3 GB	758.3 GB	174.5 GB	678.4 GB
Paris	735.3 GB	686.1 GB	678.4 GB	247GB

Chapter 4

Business aspects of caching

4.1 Introduction

As recalled in the previous chapter, the battle around content distribution has take new directions. The main reason is that many ISP or transit networks are overflowed by traffic related to content distribution, leading to unbalanced peering between networks. So far, the relationships between two transit networks wad based on free peering by assuming that traffic is balanced between both of them. But, the fact that some transit networks connect big content providers like Google leads to unbalanced peering, implying in turn that one transit network has to continually update its transmission links in order to carry the volume of traffic delivered by the other.

Beyond transit networks, the content distribution ecosystem has changed because of the ever growing importance taken by Content Delivery Networks (CDN), notably Akamai. In parallel, Content Providers are deploying their own CDN, for instance Google. Today, Google is able to peer (and pay) with an ISP connecting end users either directly or via a transit network owned by an ISP (for instance in the case of Orange). This shows that the ecosystem of content distribution is rapidly changing and new economic models have to be introduced.

In [7], an economic model is proposed on the basis of three main actors:

- Content ISP providing hosting and network access for end users and commercial companies that offer contents (e.g., Google, Yahoo!, Cogent, Akamai, etc.);
- Transit ISP offering transit to other ISPs (Level3, QWest, Global Crossing, etc.);
- Eyeball ISP offering Internet access to end users consuming contents (Orange, Free, etc. in France).

The business relationships of these actors is then studied by introducing the Shapley value used in game theory. The authors then establish new relationships between the various actors, which are not so far from the ones proposed by the Vipeer project in Deliverable D1.1.

Let us first recall the business model adopted by the Vipeer project and recalled in Figure 4.1. This business model is richer than the one considered in [7] as it describes in more details the ecosystem. In particular, in case of mobile terminals, the role of the IAP cannot be neglected since it generates a cost for the ISP (cf. the recent Free/Orange mobility agreement). This basic business model has been augmented by taking the more and more pervasive presence of CDN.

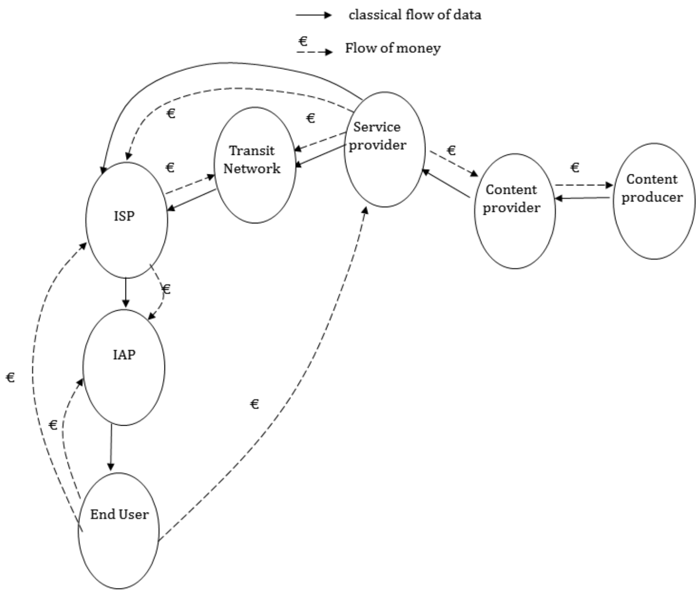

Figure 4.1: Vipeer business model with no CDN.

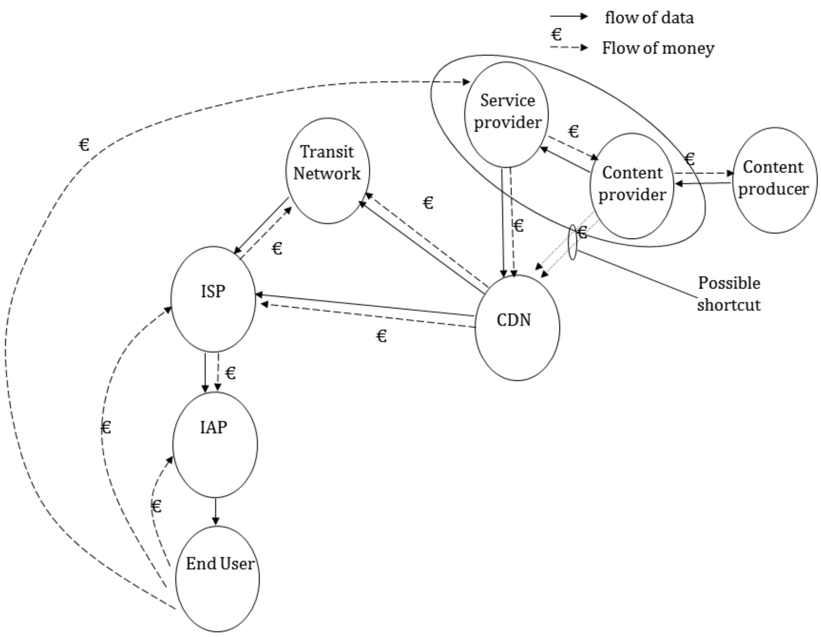

Figure 4.2: Vipeer business model with a CDN.

4.2 Formulation of a business model

4.2.1 General Internet model

Let us consider a given ISP, with customers labeled from 1 to N. It is well known that all users connected to an ISP have not the same consumption in terms of bandwidth. There are customers with very high activity (heavy hitters); they represent a rather small fraction of the total number of customers (e.g., 20 %) but generate the greatest part of traffic (about 80 %). .

While connectivity offers were dependent on bit rates a few years ago, now the connectivity is flat, only modulated by service packages and the type of access (FTTH or ADSL). Assume that ρ_n denote the revenue for the operator for a customer and χ_n the cost incurred by the nth customer for the operator. The quantity $\rho_n - \chi_n$ is the net profit for the operator and the quantity

$$\frac{1}{N}\sum_{n=1}^{N}(\rho_n - \chi_n)$$

is the average net profit.

The cost χ_n can be decomposed as the sum of fixed cost $\overline{\chi}_n$ and a variable cost $\hat{\chi}_n$ depending on the activity of the nth customer. Let β_n be the bandwidth consumed by the nth customer. The average

$$\frac{1}{N}\sum_{n=1}^{N}\beta_n$$

is the average bit rate and is used to dimension the network. The variable cost $\hat{\chi}_n$ depends on the consumed bandwidth β_n. More precisely, assume that $\hat{\chi}_n = \chi(\beta_n)$ for some function $\chi(\beta)$, which is increasing in variable β. In the following, we take $\chi(\beta) = c.\beta$, where c is the marginal cost of offering a unit of bandwidth to the customer.

Even if it may not always be true in practice, let us assume that operator's revenue ρ for the nth customer is an increasing function $\rho(\beta_n)$ of the consumed bandwidth β_n. In practice, those customers with subscription to high access bit rates (e.g., FTTH customers) pay more and generate more traffic than other customers with lower access bit rates (e.g., standard ADSL customers). It can however be observed that some highly active ADSL customers consume more bandwidth than FTTH customers.

The net revenue per user of the operator by offering Internet access to its customers is then

$$r = \frac{1}{N}\left(\sum_{n=1}^{N}(\rho(\beta_n) - c\beta_n) - \sum_{n=1}^{N}\overline{\chi}_n\right) \tag{4.1}$$

and the he global net revenue for the operator is

$$R_0 = N.r. \tag{4.2}$$

The operator has different leverages to modify the above economic equations:

- Modify the distribution of β so that the quality of experience of customers is improved and they are inclined to upgrade their offer or more customers subscribe.
- Charge content providers.

Let us consider the last possibility. There is a trend at the European level pressing Content Providers to pay for the volume of data they deliver to operators; this is already in force for Orange and Google via peering agreements (Google peers with OTI). As illustrated in Figure 1.1, half of traffic is due to streaming applications during peak hours. The majority is Youtube traffic. This implies at the individual level that a fraction η of consumed bandwidth β is due to Content Providers. If content providers are charged γ per unit of bandwidth, then Equation (4.1) becomes, without changing the other parameters,

$$r'' = \frac{1}{N}\sum_{n=1}^{N}(\rho_n - \overline{\chi}_n) + (\gamma\eta - c)\mathbb{E}(\beta), \tag{4.3}$$

and the global revenue is increased by γB where B is the bandwidth delivered by a Content Provider (e.g., 500 GBit/s between Google and Orange).

4.2.2 Introduction of caching

We have seen in Chapter 3 that caching can drastically reduce the amount of bandwidth needed on peering links to transmit to content. With caching, a fraction κ of bandwidth can be saved on peering links (e.g., $\kappa = .75$). But transmission capacities on peering links will be made available for other transactions; recall from Chapter 3 that a huge number of files are seen only one or twice and that the transmission of many files is interrupted because of quality impairments. Hence, with the introduction of caching by a network operator, one can predict:

1. An increase of the consumed bandwidth β per customer. As a matter of fact, for those files which are very popular, caching will improve the quality of those files which are very popular. These files will completely be viewed by end users and consequently their transmission will consume more bandwidth. In addition, bandwidth will be freed on peering links; there will be an opportunity for end users of downloading files which are rarely viewed and which were likely not completely downloaded because of the competition with popular files. As reported in Chapter 3 those files contribute a volume comparable with that of popular files. This phenomenon is very likely because the catalog of User Generated Content is potentially infinite, which is not the case of VoD offered by operators with finite catalog size.

   If cache servers are placed sufficiently close to end users the bandwidth consumed in the backbone will be reduced but freed bandwidth may be used for downloading additional files. Hence, the introduction of caching will increase the consumed bandwidth by a factor δ so that the consumed bandwidth per customer becomes $(1+\delta)\beta$ in the backbone, resulting in a network cost equal $c(1+\delta)\mathbb{E}(\beta)$ per customer.

2. An increase of the cost for the operator due to the installation and the operation of the caching system. Let χ_c denote the marginal cost induced by caching.

3. The improvement of quality may impact the number of customers and more users may be inclined to subscribe to higher bit rate offers, for instance FTTH. Let N' denote the new number of customers and ρ'_n the new revenue for the nth customers. It is expected that is $N' > N$ and $\rho'_n > \rho_n$ for a large number of customers.

With caching the net profit for th eoperator will become

$$r' = \frac{1}{N}\sum_{n=1}^{N}(\rho'_n - \overline{\chi}_n) - \chi_c - c(1+\delta)\mathbb{E}(\beta)$$

and the global revenue of the operator

$$R' = N'r'.$$

The introduction of caching can be profitable for the network operator only if $R' > R$ in the case the operator does not charge the content provider or if $R' > R''$ in the case the operator charges the content provider.

Chapter 5

Interconnection of CDN systems

5.1 Scenarios

CDN service providers and Telcos need CDN interconnection technology to federate CDNs deployed in different countries, to scale up their CDN capacity and their CDN service portfolio. Furthermore, CDN interconnection brings direct value to content providers, which tend to adopt multi-CDN delivery strategies to secure the distribution of their content, and to take advantage of the CDN market competition, both on price and quality of service (QoS). Such content providers select several CDN service providers and use specific systems, such as mediation platforms, to apply content distribution and request routing policies towards the selected CDNs. Interconnecting the transparent caching systems, which have been deployed by network operators to ease the pressure of the OTT traffic over their networks, with the CDNs that generate this OTT traffic is also considered. Such interconnection would optimize network resource utilization and improve the quality of experience for the end-user. IETF has detailed the most relevant use cases for short and medium-term CDN_i deployments in [10]. These include:

1. Footprint extension use cases: for geographic extension, interconnection of affiliates' CDNs, control of incoming traffic, or handling of nomadic users' access to CDNized content.
2. Offload use cases: to handle overload and dimensioning issues, and to increase resiliency to content delivery and content acquisition failures.
3. CDN capability use cases, which permit a CDN to extend its device and network technology support, to provide technology and vendor interoperability, and which enable QoE and QoS improvement.

5.2 CDN_i framework

For the sake of standardization efficiency, IETF specifies in a first stage only a simple framework including five interfaces: control, logging, request routing, metadata, and acquisition [9]. We argue that this framework is overly simplified: the interfaces' scope is too large to permit a clear definition of the roles and responsibilities of every interface. In addition, all aspects related to the business interface, which are required to establish any type of network interconnection, are completely out of scope for IETF, whereas they have a major influence on the design of all other interfaces. A key contribution of this task consists in presenting a framework that covers all aspects of CDN_i operations and is ready for future use cases, which go beyond the features currently considered in standardization bodies such as IETF [9] [10], ETSI/TISPAN [11] [12] or ATIS [2]. Figure 5.1 depicts this framework and we describe it below.

1. Business Interface: Operations to establish, update, and tear down CDN_i agreements (including negotiation of Service Level Agreements (SLA), billing, etc.). This interface interconnects CDN operators, not CDNs: its description is required to ensure any CDN_i system runs properly.

2. CDN_i Control Interface: Operations to initialize, parameterize, and tear down subsequent CDN_i interfaces, which are instantiated for every CDN_i agreement. Once the control interface has established a CDN_i interface, all runtime control over CDN_i behavior is under the purview of the instantiated interface. This is the same definition as in [9].

3. CDN_i Routing Interface: Operations to exchange routing information that enables CDN selection for user requests.

4. CDN_i Downstream Resource Identifier Signaling (DRIS) Interface: Operations to exchange all information required to forge a response which redirects the content-request to the selected downstream CDN. We call Downstream Resource Identifier(s) this set of information and CDN_i request redirection stratey a protocol-dependent (e.g., DNS or HTTP) method to redirect a content-request from an upstream CDN to a downstream CDN.A CDN_i request redirection strategy is implemented in two parts:

   - Signaling messages exchanged directly between the upstream and the downstream CDN: this is the CDN_i DRIS Interface.
   - Messages exchanged "in-band" via the end-user and/or its local DNS by leveraging features of applicative protocols such as HTTP and DNS.

   We have specified the logic of both parts and their interactions. Once the CDN selection is achieved for a given user content-request, the upstream CDN has decided whether the requested content is to be delivered by itself or the content-request is to be redirected to any downstream CDN. If the upstream CDN selects a given downstream CDN, it generates a response by using the information provided by the selected downstream CDN over the CDN_i DRIS interface. This response redirects the content-request towards the selected delivery CDN.

5. CDN_i Content Request Requirements Control (CRRC) Interface: Set of operations to communicate the requirements that govern how the content-requests from the end-user agents must be handled by interconnected CDNs. Examples of CDN_i Content Request Requirements include content-request geo-blocking directives, content-request validity windows, content-requests authorization mechanisms (a.k.a. access control mechanisms), and content-request denial (a.k.a. purge directives).

6. CDN_iUpstream Resource Identifier Signalling (UPRIS) Interface: Operations to exchange all information required to forge a request for content acquisition from a given

surrogate within the downstream CDN to one given acquisition source from the upstream CDN. We call "Upstream Resource Identifier(s)" this set of information.

7. CDN_i Acquisition Interface: Operations to exchange content between the CDNs, for example in case of cache-miss in any downstream CDN.

8. CDN_i Inner Metadata Interface: Operations to exchange between the CDNs content metadata that do not depend, neither from the CDNs, nor from the content-requests, for example the validity or expiration dates of a piece of content, the content's name and size, etc.

9. CDN_i Logging Interface: Operations that allow interconnected CDNs to exchange relevant activity logs.

Figure 5.1 illustrates the role of these interfaces, considering an example content-request from an end-user agent redirected by CDN-a to CDN-b. The content-request is received by CDN-b CDN-b checks first whether the content-request is valid and authorized, and what it has to do with that request. All information to achieve that is provided over the CDN_i CRRC Interface (double lined red arrow). Once CDN-b has checked the content-request is authorized, it consults its routing table(s), which are fed by the CDN_i routing interface (blue dotted arrow), to decide whether it is CDN-b or any downstream CDN of CDN-b that will deliver the content. Once CDN-b has decided the content-request is to be redirected to CDN-d, the process in charge of preparing the redirection request calls the CDN_i DRIS process to get the information required to forge the redirection request. The end-user follows the redirection request (e.g. a DNS CNAME resource record or an HTTP 307 message) and sends its content-request now to CDN-d instead of CDN-b. Then CDN-d iterates similar operations:

- Checking request authorization,
- consulting its routing tables,
- deciding that the content will be delivered by one of its surrogates,
- preparing a redirection request with as location that surrogate.

The surrogate receives the request. In case the URL is signed and depending on the chosen URL signing system, the surrogate has to fetch either local information (key, etc.) to check the signed URL or information provided by the CRRC interface. Once the surrogate has checked the content-request is authorized, it checks whether it has a copy of the requested content.

- If it does not (cache miss), it activates the process for preparing inter-CDN content acquisition request. This process calls the CDN_i UPRIS process from CDN-d, which interacts with a CDN_i UPRIS process from CDN-b to get all information required to forge a content acquisition request. Then, the surrogate from CDN-d sends the request for content acquisition to the selected content acquisition source from CDN-b, which answers with the requested content.

- If the surrogate has a copy of the content but wants to validate this content, it may use the CDN_i Inner metadata interface to exchange information such as content expiration and last modification dates.

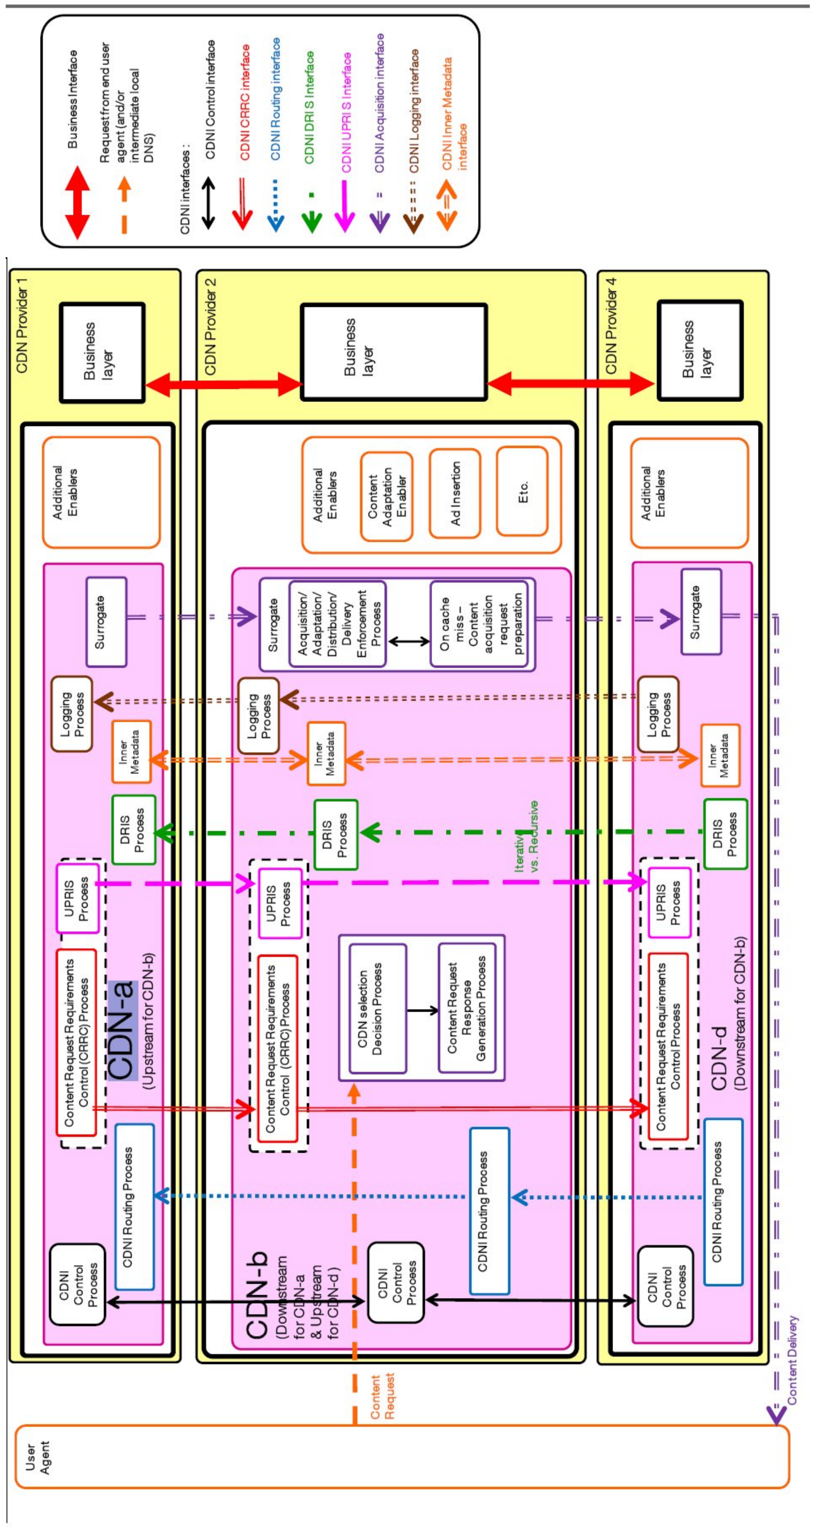

Figure 5.1: Proposed CDN_i framework.

Finally, the surrogate delivers the content to the end-user.
Note that some interfaces are unidirectional by design. For instance, the upstream CDN does not provide the downstream CDN with any CDN_i routing information or any CDN_i Downstream Resource identifier(s). Similarly, the downstream CDN does not provide the upstream CDN with any CDN_i CRRC information, any CDN_i Upstream Resource identifier(s), or any deliverable content.
We claim that the presented approach provides a more adequate framework for CDN Interconnection than the ones proposed so far in standardization bodies, mainly at IETF CDN_i WG for several reasons.
First, it is better adapted to the CDN_i deployment scenarios envisioned at short and longer terms than any proposal made so far about CDN_i request routing and CDN_i Metadata. Second, it is exhaustive: it describes all the interfaces and processes involved, including the CDN_i control interface and the bootstrapping process of all CDN_i interfaces. Third, it also provides an adequate generic framework for any task shared between CDNs (not only for distributing content between CDNs, but also to manage classes of services, and to distribute between two CDNs tasks such as transcoding, ad insertion, URL rewriting, transparent caching-CDN interconnection, etc.). In summary, it goes much beyond the current framework of all standardization and research work made public so far on CDN Interconnection.

Chapter 6

Conclusion

We have shown that caching is an efficient solution for reducing the amount of consumed bandwidth on some critical links in a network. From an economical point of view, the efficiency of caching is more questionable. It is profitable for a network operator to deploy caching only if the quality perceived by end users is indeed much improved so that more customers subscribe, possibly to higher bit rate offers. An adverse effect of caching would to increase the mean consumed bandwidth per user and imply the need for updating access networks, which is a huge cost for the network operator.

Some content providers, such as NetFlix, are currently proposing to deploy in ISP networks open source caching systems so that NetFlix can operate a CDN in an ISP network, or to directly peer with ISPs. This is not in line with the Vipeer solution, which recommends to deploy a caching system operated by the network operator in his own network. The solution proposed by NetFlix to operate its own CDN in the operator's network is highly dangerous for the network operator, which can then play no role in content delivery and becomes a dummy pipe provider.

Bibliography

[1] IP transit price declines steepen. http://www.telegeography.com/products/commsupdate/articles/2012/08/02/ip-transit-price-declines-steepen/, August 2012.

[2] ATIS-0200003. CDN interconnection use case specification and high level requirements. 2011.

[3] CISCO. The zetabyte era. May 2012.

[4] Autorité de la Concurrence. Décision n 12-d-18 du 20 septembre 2012 relative à des pratiques mises en oeuvre dans le secteur des prestations d'interconnexion réciproques en matière de connectivité internet. 2011.

[5] P. Gill, M. Arlitt, Z. Li, , and A. Mahanti. Youtube traffic characterization: a view from the edge. In *In Proceedings of the 7th ACM SIGCOMM conference on Internet measurement, IMC 07*, pages 15–28, 2007.

[6] C. Labovitz. Internet traffic and content consolidation. Keynote presented as IETF 77 Plenary. Anaheim, CA., March 2010.

[7] R.T.B. Ma, D.M. Chiu, J.C.S. Lui, V. Misra, and D. Rubenstein. On cooperative settlement between content, transit, and eyeball Internet service providers. *IEEE/AC ToN*, 19(3), June 2011.

[8] William B. Norton. A business case for peering in 2010. http://drpeering.net/white-papers/A-Business-Case-For-Peering.php.

[9] L. Peterson and B. Davie. Framework for CDN interconnection. draft-ietf-cdni-framework. In *Work in progress*, 2012.

[10] RFC6770. Use cases for content delivery network interconnection.

[11] ETSI TS 102 990 v0.0.12. Media content distribution (MCD); CDN interconnection, use cases and requirements. 2012.

[12] ETSI TS 182 032 v0.0.6. CDN interconnection architecture. 2012.

[13] Joseph W. Waz and Lynn R. Charytan. Comcast's letter to FCC on Level3. http://corporate.comcast.com/comcast-voices/comcasts-letter-to-fcc-on-level-3, November 2010.