

Characterizing a Meta-CDN

Oliver Hohlfeld, Jan R uth, Konrad Wolsing, Torsten Zimmermann[†]

[†]Authors in alphabetical order

Communication and Distributed Systems, RWTH Aachen University
{hohlfeld, rueth, wolsing, tzimmermann}@comsys.rwth-aachen.de

Abstract. CDNs have reshaped the Internet architecture at large. They operate (globally) distributed networks of servers to reduce latencies as well as to increase availability for content and to handle large traffic bursts. Traditionally, content providers were mostly limited to a single CDN operator. However, in recent years, more and more content providers employ multiple CDNs to serve the same content and provide the same services. Thus, switching between CDNs, which can be beneficial to reduce costs or to select CDNs by optimal performance in different geographic regions or to overcome CDN-specific outages, becomes an important task. Services that tackle this task emerged, also known as CDN broker, Multi-CDN selectors, or Meta-CDNs. Despite their existence, little is known about Meta-CDN operation in the wild. In this paper, we thus shed light on this topic by dissecting a major Meta-CDN. Our analysis provides insights into its infrastructure, its operation in practice, and its usage by Internet sites. We leverage PlanetLab and Ripe Atlas as distributed infrastructures to study how a Meta-CDN impacts the web latency.

1 Introduction

Content Delivery Networks (CDNs) have become a key component of the web [7,8]. Their ongoing quest to serve web content from nearby servers has flattened the hierarchical structure of the Internet [14] and promises lower latencies, while their distributed nature promises high availability. These benefits led to a wide adoption of CDNs for web content delivery that is manifested in high traffic shares: for example, more than half of the traffic of a North American [13] or a European [20] Internet Service Provider (ISP) can be attributed to few CDNs only. Despite these benefits, customers of a single CDN are bound to its cost model and performance figures—a limitation that is solved by multihoming content on different CDNs and subsequently serving it from the CDN that currently offers better performance and/or lower costs.

To better utilize content-multihoming, *Meta-CDNs* [12] enable content providers to realize custom and dynamic routing policies to direct traffic to the different CDNs hosting their content; A concept also known as CDN-Selector [23] and that is related to auction-based CDN brokers [16,17]. Request routing is performed by the Meta-CDN according to *custom routing logic* defined by content-providers (i.e., the customers of a Meta-CDN and CDNs). This routing logic can be informed by a broad range of factors, including CDN cost models or measured CDN performance. Content providers can thus utilize a Meta-CDN to reduce costs or to optimize performance, e.g., by implementing

custom logic to direct traffic to a CDN that currently offers better performance and/or lower cost (e.g., at certain geographic regions or times). Since the routing approach employed by the Meta-CDN customers is unknown to the involved CDNs, directed traffic and thus generated revenue gets harder to predict. In particular, since decisions can be based on active performance measurements by the Meta-CDN, a (single) delivery of bad performance by the probed CDN can result in rerouting traffic to a competing CDN and thus losing revenue. Thus, while Meta-CDNs can offer cost and performance benefits to content providers, they also challenge CDN business models. Concerning Internet-users, performance-based routing decisions can yield better Internet performance and benefit end-users while cost-based decisions can have other effects (as for any server selection approach run by CDNs). While the concept is known and related work covering service specific implementations, i.e., Conviva's streaming platform [4,10,16], exists, the empirical understanding of a generic Meta-CDN and its operation in practice is still limited. We posit that this understanding is necessary.

In this paper, we thus shed light on the Meta-CDN operation by dissecting the Cedexis Meta-CDN as a prominent example that is used by major Internet companies such as Microsoft (Windows Update and parts of the XBox Live Network), Air France, and LinkedIn [1]. Given its current adoption, understanding its functionality and its usage by customers provides a first step towards understanding currently unknown implications of Meta-CDNs on Internet operation. We thus investigate the infrastructure and services powering this Meta-CDN and provide insights about its operation in practice. We analyze for *what* kind of services, e.g., media, API backends, or bulk data transfers, customers utilize Cedexis and *how* different CDNs are employed. We further investigate how the infrastructure deployed by Cedexis impacts the overall request latency performance in a PlanetLab and Ripe Atlas measurement. Specifically, our contributions are as follows:

- i) We characterize Cedexis, as a representative generic Meta-CDN, present its operation principles, and further analyze and classify its customer base. Moreover, we illustrate which CDNs are used by the customers.
- ii) We utilize globally distributed vantage points, i.e., Ripe Atlas, PlanetLab, Open DNS Resolvers and a small deployment of probes behind home user DSL connections, to obtain a *global* view on Cedexis. Based on these measurements, we analyze the deployed infrastructure of Cedexis, and are further able to investigate if the selection process varies based on the location. In addition, we find cases of suboptimal routing in terms of latency.

2 Background and Related Work

To achieve high availability, content and service providers typically employ CDN operators and utilize their already deployed and geographically distributed infrastructures [7,8]. In addition to increased availability, end-users profit from the distributed nature of these CDNs when retrieving content from close-by servers, reducing the overall latency.

Multiple works from academia and industry have investigated these infrastructures, the operation principles as well as the performance of deployed CDNs [7,8,18,19]. Besides understanding and measuring CDN infrastructures, researchers have utilized CDN routing techniques to derive network conditions [21]. In addition, approaches that optimize the routing of user requests to the respective servers within a CDN, as well as, optimized anycast load balancing have been proposed [9,11]. Poese et al. [20] present

and analyze the impact of an ISP recommendation service providing insights about the current network state, e.g., topology, load, or delay, to the CDN, which in turn bases its server selection on the returned information. The idea and concept of the presented approach is revisited by Frank et al. [12] and, among other features, enables a CDN to allocate server resources within the ISP on-demand when necessary.

To further ensure the availability of content, customers may use multiple CDN deployments. Other reasons to utilize more than one CDN provider may be cost efficiency, e.g., different prices to serve content at different times or due to traffic volume contracts. However, with multiple locations at different CDNs serving the same service or content, either the customer or an additional service has to choose between the actual CDN when a user requests a service or content [12,15,23]. With respect to *multi-homed* content, i.e., content that is distributed by multiple CDNs, Lui et al. [15] present one of the first frameworks optimizing performance and cost of the resulting CDN assignment. In the case of video streaming, Conviva [4] uses a recommendation system that is utilized by the video player software [10] for the CDN selection. Besides Conviva, commercial solutions that offer to act as the CDN selector in more generic settings, e.g., websites or services, exist [1,5,6], however, there is currently little to no understanding of their infrastructures, customers, and the effects on the global CDN landscape. With respect to Meta-CDNs and especially Cedexis, Xue et al. [23] are the first to provide brief performance figures about the selected CDNs, focusing on deployments in China. We aim at more broadly characterizing Cedexis as a whole while looking at their infrastructure and performance on a global scale. Nevertheless, we find, similar to Xue, partly suboptimal performance in terms of latency, yet, we acknowledge that routing decisions may have other goals than latency. This argument is reinforced by Mukerjee et al. [16], which is closest to our work. They analyze the effect of brokers, i.e., CDN selectors, on CDNs, characterize potential problems and propose a new interface between these brokers and CDNs. While a closer interaction may improve certain aspects, it remains open whether a Meta-CDN such as Cedexis does, in fact, harm a CDN's profitability. We cannot confirm a raised concern that a broker *might* prefer certain CDNs in certain regions, as we find similar CDN choices worldwide.

The goal of this paper is to extend the currently limited understanding of Meta-CDN operation by characterizing Cedexis as a prominent example of a generic Meta-CDN. Exemplified by understanding its overall deployment, customers, and the effects Cedexis, we aim to provide a first step towards a better understanding of Meta-CDNs in general.

3 Characterizing a Meta-CDN

The general motivation behind a Meta-CDN is to enable custom and dynamic routing of requests to content that is multi-homed in different Content Distribution Infrastructures (CDIs). A CDI can involve any infrastructure ranging from simple (cloud-hosted) servers to complex CDNs [20]. Multihoming content on different CDIs enables content providers to optimize for availability, performance, or operational costs. By utilizing a Meta-CDN, content providers can realize *custom* routing logic to direct traffic to the available CDIs. Such custom routing logic can be motivated by CDIs that offer better performance and/or lower costs in certain geographic regions, at certain times of the day, or request volumes. We refer to an infrastructure that enables routing between CDIs with customer-provided routing logic as a *Meta-CDN*, a concept that is also referred to as Multi-CDN

```

$ dig www.accorhotels.com
...

;; QUESTION SECTION:
;www.accorhotels.com.                IN      A

;; ANSWER SECTION:
① www.accorhotels.com.                624     IN      CNAME   ② 2-01-2770-000c.cdx.cedexis.net.
2-01-2770-000c.cdx.cedexis.net.     20      IN      CNAME   ③ www.accor.com.edgekey.net.
www.accor.com.edgekey.net.         270     IN      CNAME   e5809.b.akamaiedge.net.
e5809.b.akamaiedge.net.            20      IN      A       104.104.185.148

```

Fig. 1: Exemplary output of `dig` resolving a customer domain managed by Cedexis. The requesting user is redirected by the Cedexis authoritative DNS with a CNAME to the CDN selected to handle the request. The Cedexis CNAME contains the customer ID ② (C_{ID}) and a per-customer application ID ③ (App_{ID}).

selector [23] and has similarities to auction-based CDN brokers [16]. Since the individual routing approaches employed by content providers at the Meta-CDN are unknown to the involved CDIs, directed traffic and thus generated revenue gets harder to predict. In particular, since decisions can be based on active performance measurements by the Meta-CDN, a (single) delivery of bad performance by the probed CDI can result in rerouting traffic to a competing CDI and thus losing revenue. While the effects of Meta-CDN operation are relevant to Internet operation, little is known about Meta-CDNs.

To elucidate Meta-CDN operation, we start by characterizing Cedexis as a prominent example. We base this characterization on showing *i*) its operational principles to select and routing between CDIs based on the Cedexis site (Section 3.1) and *ii*) its current use in the Internet by analyzing its customer base in Section 3.2 based on our measurements. Both perspectives provide a first understanding of the principle mechanisms with which Meta-CDNs can influence content distribution and their current deployment in the wild.

3.1 Operation Principles

Like most CDNs, Cedexis employs a DNS-based redirection scheme, similar to Akamai [18], to redirect the requesting user to the CDI selected for content delivery. This redirection scheme is based on Canonical Name (CNAME) records which transfer the requesting user between the different authoritative name servers (NS). We exemplify this scheme in Figure 1. Starting at the original domain (①), the user gets transferred to the Cedexis NS, which then selects a final CDI. The configured static Cedexis CNAME includes a Cedexis customer ID (C_{ID}) (②) and configuration specific (App_{ID}) (③). Both identifiers enable the Cedexis NS to perform customer-specific request routing, once the NS is contacted by the client’s DNS resolver for name resolution. Similar to classic CDNs, routing can be subject to a user’s location, e.g., identified by DNS resolver IP or EDNS0 client subnet extension. The Cedexis NS then points to the selected CDI, which can be an IP address in an A resource record or another CNAME, e.g., pointing to CDN (Akamai in this particular example). The selected CDI can then repeat this process to select the final server handling the request or point to another CDI. By realizing routing in the DNS, Cedexis redirects requesting users to a CDI *before* a connection to the CDI is established. This way, it is not involved in the actual content delivery itself and thus does not alter the performance or security properties provided by the selected CDI.

CDI Selection Options. The above-stated request routing approach can be arbitrarily dynamic, i.e., the user to CDI mapping in the DNS can change at any time, only limited by the cacheability of their DNS records. This aspect is utilized to enable content providers to realize custom routing logic within Cedexis using three components: *i) Openmix* enables Cedexis customers to configure individual routing logic. Customers can choose between `optimal RTT` to select the CDI with the lowest RTT to the requesting user, `round-robin` balancing between all CDIs configured by a customer, `throughput` to select the CDI with the highest throughput, a `static routing` or by executing `customer provided code`. This routing logic can be informed by *ii) Radar*, a large database for decision making based on active browser-based CDN measurements performed by website visitors, and *iii) Fusion*, to retrieve data from CDNs. Every customer has the ability to configure site-specific behavior (i.e., *Apps*) which results in different `AppIDs` in the DNS. This way, a customer can configure different routing profiles for `downloads.domain.tld` and for `images.domain.tld`. We next describe the two data sources and the Openmix platform to realize custom routing decisions.

Radar. Cedexis provides a community-driven CDN performance measurement platform called Radar. Radar employs active measurements performed within the web browser of visitors of Cedexis-managed websites. The in-browser measurements require the Cedexis customers to embed a Javascript in their website. Once visited, the web browser triggers the website's `onLoad` event, the embedded Javascript waits for a user configurable timeout (default 2 s) and starts requesting probe instructions from Cedexis. Users can configure private probes (i.e., to estimate their own performance) and can choose to activate community probes (i.e., enabling Cedexis to measure other infrastructures). Probes can serve different purposes, latency or throughput measurements, which are expressed through a small 43 Byte file (latency) or 100 kB file (throughput) that is fetched via HTTP. After having performed the measurements, the Javascript reports the obtained values back to Cedexis such that they can later be used to guide performance decisions in the CDN selection process and to inform site operators about their site's performance.

Fusion. While Radar enables to realize performance-based routing decisions, Fusion enables realize decisions on statistics directly from a user's CDNs. CDNs typically offer statistics about traffic shares, quotas, budgets, performance and other KPIs in their web interfaces. By accessing these, Cedexis enables their customers to not only draw performance decisions but also business decisions (e.g., *will I hit my quota soon?*).

Openmix. Both Radar and Fusion data are available to customers in the Openmix platform. Openmix enables customers to customize the DNS resolution process by providing custom Javascript code that is executed in the DNS resolution step. Within this code, customers can define their subsequent candidate CDI choices and request measurement data (e.g., availability, latency, throughput) for these. While the system is powered by many probes, only singles values are returned suggesting that Cedexis preprocesses the measurement data, yet we were unable to find more information on this process. Thus, Openmix is used to realize customer-specific routing decision performed within the DNS resolution, i.e., directing traffic to a target CDI via a DNS CNAME.

Takeaway. *Cedexis offers its customers to realize site-specific, fine-granular, and dynamic traffic routing to CDIs, e.g., based on customer-provided code and informed by rich measurement data. The performed traffic routing is hard to predict (e.g., for CDIs).*

Type Share	Type Share	Type Share	Service Share
Business 17.7%	Unknown 8.1%	Social 1.6%	Web 62.7%
IT 12.1%	Goods 5.6%	CDN 1.6%	Unknown 15.6%
News 11.3%	Automotive 5.6%	Patents 0.8%	Assets 12.9%
Gambling 11.3%	Advertising 3.2%	Banking 0.8%	Media 5.4%
Shopping 8.1%	Streaming 2.4%		API 2.3%
Gaming 8.1%	Television 1.6%		Bulkdata 1.1%

(a) Classification

(b) Services

Table 1: Cedexis customer information obtained from manual inspection of websites served for different customer IDs. Please note that customers may operate multiple services, e.g., multiple brands of one holding company.

3.2 Customers

Before we analyze the infrastructure and configuration of Cedexis, we want to shed light on their customer base (in anonymous form). We are interested in which companies and businesses leverage this additional service on top of classical CDN infrastructures.

DNS Measurement Methodology. Our approach is twofold. First, we leverage the encoded customer and application IDs in the CNAME structure (see step ② & ③ in Figure 1) to **enumerate customers**. Applications are used by customers to define different routing profiles that map to the available CDIs, so e.g., a customer may have one profile that is optimized for latency, and another for throughput. Conveniently, App IDs start at 1. Thus our approach is to simply enumerate customers by resolving all $2-01-(C_ID)-(App_ID).cdx.cedexis.net$ domains. As customer and application ID each have 4 hex characters, we would need to probe $> 2.5 \text{ B } (16^8)$ domains. To scale-down our DNS resolution, we only enumerate the first 256 application IDs for each customer, resulting in resolving roughly 16 M domains.

Domain Lists. Second, we probe domain lists to study the usage of the enumerated CNAMEs in the wild and to discover application IDs beyond the enumerated 256 IDs. We thus resolve the A of `domain.tld` and the A `www.domain.tld` records for all domains in the *i)* `.com/.net` (obtained by Verisign), *ii)* `.org` (obtained from PIR), *iii)* `.fi` (obtained from Ficora), *iv)* `.se/.nu` (obtained from IIS), *v)* `.new` gTLD zones (obtained from ICANN’s Centralized Zone Data Service), *vi)* obtained from our passive DNS probe, and *vii)* the Alexa Top 1M list. We additionally include the Cisco Umbrella Top 1M list [3], which is based on the most frequent queries to OpenDNS resolvers and additionally contains subdomains, e.g., `images.domain.tld`. Depending on the size of the list, we perform daily or weekly resolutions for four weeks in August 2017 and extract all domains which have a CNAME pointer containing `*.cedexis.net`.

Customer List. We combine both data sets to a customer list that will form the basis for probing Cedexis globally in Section 4. The list contains all customer application tuples, of which 84 % were discovered in the enumeration step, 11.2 % were discovered in both the enumeration and in the domain lists, and 4.8 % solely in domain lists. The reasons for the latter are application IDs larger than 256, which were not part of our enumeration. Out of all customers, 55 (20) have only 1 (2) application(s) configured. We also observe 1 customer having 84 configured. By resolving the domain lists, we find 4609 (sub-)domains pointing to 16 % of all discovered (customer, application) tuples. The remaining 84 % were not hit when resolving our domain lists. For 62.7 % of all

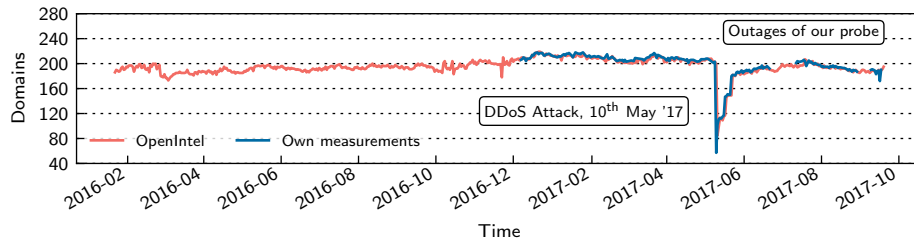


Fig. 2: Domains utilizing Cedexis in the Alexa 1M over time. The drop in May '17 was caused by a DDoS Attack on Cedexis [2]. Unfortunately, the measurement probe located at our chair experienced two outages. However, the overlap of both scans motivates the further use of the OpenIntel data set.

customer application IDs, we only find a single domain pointing to it. We find 31 (6) application IDs managing more than 10 (100) domains, respectively.

We show the popularity of Cedexis over time among Alexa-listed domains in Figure 2. We set up our own regular DNS resolutions in December 2016 and further show regular Alexa resolutions performed by OpenINTEL [22] in the Netherlands for the same resource records. First, we observe that both data sets overlap, suggesting that both are suitable for monitoring. Minor fluctuations in the number of domains per day can mainly be attributed to fluctuations in the Alexa listed domains. Second, we observe an outage of Cedexis in May 2017 which was caused by a DDoS attack on their infrastructure [2]. The outage motivated some customers to remove CNAME pointers to Cedexis in favor of pointing to operational CDNs instead, causing a drop of > 120 domains in Figure 2.

Customer Classification. We next classify the discovered customers to highlight the variety of Cedexis customer' profiles. To base this on an open classification scheme, we first tried to match customer domains against the Alexa Web Information Service API. However, Alexa classifications exist only for 17% of the queried domains and some classifications do not reflect the web pages' content. To obtain a broader picture, we instructed a single human classifier to visit each web site and categorize it according to an evolving set of categories. We show the resulting categorized web site content in Table 1(a). The table shows that Cedexis is used by a broad range of customers. We further classify the used service in Table 1(b). The table shows that most customers use Cedexis for general web content delivery. This includes few but large bulk download services, e.g., `www.download.windowsupdate.com`. This is in contrast to, e.g., Conviva which is dedicated to video delivery.

Takeaway. *Cedexis is utilized by a number of (large) web services. Decisions taken by Cedexis have the potential to impact larger bulks of Internet traffic.*

4 A Global View on Cedexis

As Cedexis customers can realize routing decisions based on (network) location, we next take a global view on its customers by using globally distributed active measurements.

Measurement Setup. We base our measurements on 35 PlanetLab nodes located in 8 countries, 6 custom Raspberry Pi (RPi) probes in 6 distinct German ISPs, and RIPE Atlas probes. We chose PlanetLab and custom RPis in addition to RIPE Atlas since they enable us to deploy custom software to perform frequent DNS resolutions. As we do

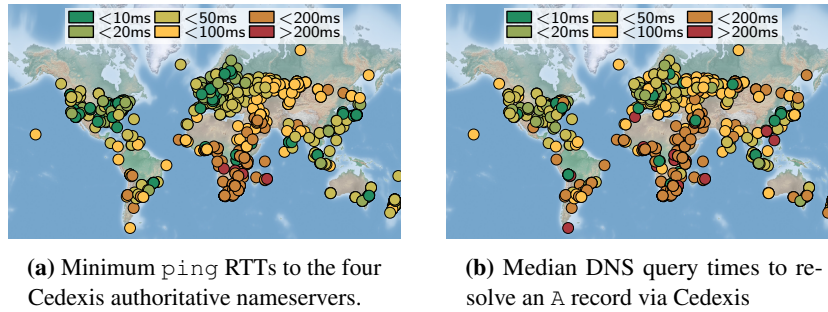


Fig. 3: RTTs and DNS Query times obtained from ≈ 870 responsive RIPE Atlas Probes performing pings and DNS A Record requests.

not include PlanetLab nodes located in Germany in our set, we refer to our deployed RPIs when mentioning *DE* in figures or plots. We selected only few PlanetLab nodes with high availability to repeatedly measure always from the same vantage points. For our measurement, we instruct the Planet Lab and our RPI nodes to resolve each domain every 15 min and subsequently measure the latency to the resulting IPs. Moreover, we keep track of all CNAME redirections to CDNs that we observe over the course of the measurement and also resolve these. This way, we learn the set of configured CDNs for every probed domain.

4.1 Infrastructure

Authoritative DNS Deployment. Cedexis core functionality is based on a distributed infrastructure of authoritative name servers managing `*.cedexis.net`. We find four servers configured in the DNS in our measurements and in the `.net` zone file. We remark that a larger number exists which we found by enumerating their naming pattern. However, they currently appear to be unused, i.e., not included in the `.net` zone and are not discovered by our active DNS measurements.

To obtain a better understanding of its DNS infrastructure, we measure the ICMP echo (ping) latency to their authoritative name servers from ≈ 870 responsive (out of 1000 selected) RIPE Atlas probes. We repeated this measurement 30 times using the same set of probes and show the minimum RTT in Figure 3(a). Based on these latency figures, we infer that Cedexis operates DNS servers located in North-America, Europe, and (probably) Asia and South America. By analyzing individual latencies and manual traceroutes per server-IP (not shown), we observe latencies of < 10 ms from multiple regions (e.g., US-East, US-West, Europe, Hongkong) to the *same* DNS server IP. Since these low latencies are lower than required by the speed of light between the respective regions, it suggests that the probed server-IPs are served using *anycast* routing.

Since the additional indirection step through Cedexis contributes latency, we next measure the DNS resolution time of only their authoritative name servers. In this step, we resolve the domain of a Cedexis customer from all authoritative name servers from the same RIPE Atlas probes, again repeated 30 times. To limit the resolution to only involve the Cedexis DNS, we chose a customer domain which directly returns an A record instead of redirecting to another CDN. This is especially important if the lifetime of the

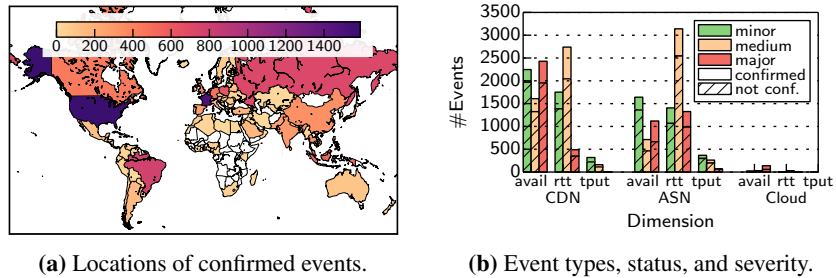


Fig. 4: Cedexis events reported from 9th October 2017 to 8th January 2018.

DNS records is short (which we analyze in Section 4.2) and clients need to frequently contact the Cedexis DNS over and over again. We show the median DNS query time in Figure 3(b) and observe that the DNS query times follow the previously measured ping latencies. Nevertheless, we observe few regions in which Cedexis appears to be uncovered as they involve high DNS resolution latencies, e.g., Latin America or Africa. **Radar Community Probes.** Cedexis customers can realize routing decisions that are based on active measurements of current service performance probed by visitors of Cedexis managed websites (Radar platform, see Section 3.1). Understanding this data is interesting since it can influence routing decisions. As the Radar data is not publicly available, we instead analyze live feeds of network events detected by Radar and published at <https://live.cedexis.com>. The events report three classes of metrics: *latency*, *throughput*, and *availability* for *ISP*, *CDN*, and *cloud* infrastructures. An event can be a latency in- or decrease, an outage, or a change in throughput. A CDN/cloud event is detected if it was reported by visitors from 5 different ASes. Likewise, an AS event is detected if it concerns 5 CDNs or clouds (see live.cedexis.com). Each event can be classified by severity into *minor*, *medium*, and *major*. Apart from being used in their decision-making process, this data allows monitoring the reported infrastructures. We thus monitored the data feed from October 9, 2017, to January 8, 2018.

Reported events further provide information *where* visitors of Cedexis-managed sites are located. This is based on fact that Cedexis customers embed Javascript measurement code into their sites that reports performance figures to the Radar platform. While the number of events is likely uncorrelated with the number of website visitors, it at least indicates the presence of a visitor in the reported AS or country. Therefore, we show the distribution of the number of events per-country in Figure 4(a). We observe almost no events in Africa, suggesting that Cedexis customers do not have a large user base in Africa, which also coincides with the suboptimal DNS deployment there. While we see events in almost every country, most events are reported in Central Europe, North America, Brazil, and Russia.

We next analyze the reported events by their type, shown in Figure 4(b). The figure shows the number of events per event type categorized to availability (avail), latency (rtt), and throughput (tput) for CDNs, AS, and cloud providers. Every bar is divided in the amount of confirmed and unconfirmed events. We observed that an event is marked as *confirmed* when it was reported for at least 9 minutes and the rolling variance of measurements from the last 5 hours exceeds an event and severity level specific threshold: e.g., a latency increase of 100% - 200% for a CDN is considered as minor, while an

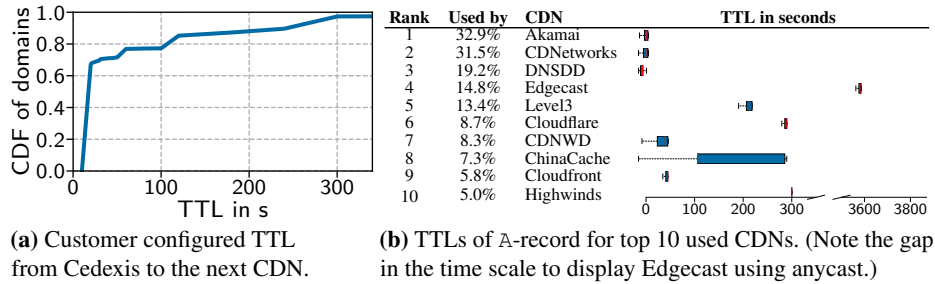


Fig. 5: DNS TTLs experienced among Cedexis-enabled domains. For a) mappings from Cedexis to the subsequent entry and b) the CDNs used for the final delivery.

increase between 200% and 500% is considered as medium severity. We find most of the reported events to concern CDN, followed by ASes. The high amount of major availability events concern CacheFly CDN outages during our measurement period.

Takeaway. We observe visitors of Cedexis-managed sites from almost every country. Yet, its anycast DNS platform is suggested to be based in the US, Europe, and Asia. Users in other countries can be subject to higher DNS query latencies.

4.2 How Customers utilize Cedexis

DNS TTL. The DNS Time To Live (TTL) defines the time a record can be cached by DNS resolver and thus the timespan between Cedexis balancing decisions. A small TTL allows more frequent switches at the cost of more frequent DNS queries to the Cedexis DNS infrastructure. This query latency can be significant, depending on the DNS resolver location.

Figure 5(a) depicts the CDF of the TTLs for the validity of the CNAME-mappings from Cedexis to the subsequent entity (see 2nd CNAME in Figure 1) for all customer domains. We did not observe country-specific settings. Around 67% of all domains have configured a TTL of at most 20 s, indicating a rather short time scale enabling rapid reactions to changes. The next 30% are within 300 s, denoting an already moderate reaction time while around 3% have configured higher TTLs. Higher TTLs *can* hint to non-latency-based, but rather throughput or cost-based optimizations.

To compare these configurations to TTLs deployed by CDNs, we show the A-record TTLs for the top 10 CDNs in Figure 5(b). To the right of every CDN, the figure shows the boxplot of TTLs observed for the A-records of all resolutions we performed. We see that the top 3 CDNs use a short TTL in the range of most Cedexis CNAMEs, whereas Edgecast has a lifetime of one hour (probably due to their use of anycast).

DNS Resolution Time. When employing Cedexis, an additional step in DNS resolution is required to enable CDN balancing. Figure 6(a) compares the latency for resolving (from our Planet Lab sites) a mapping at Cedexis, in case of multi-staged CDNs a further CNAME redirect (CDN) and the final resolution of the A-record. We observe that Cedexis performs similarly to the other CDNs. However, while this hints at a good DNS deployment for our vantage points, it also means that using Cedexis inflates the latency of a DNS lookup. Given the on average short TTLs, users will often incur an additional added latency when using Cedexis-enabled websites.

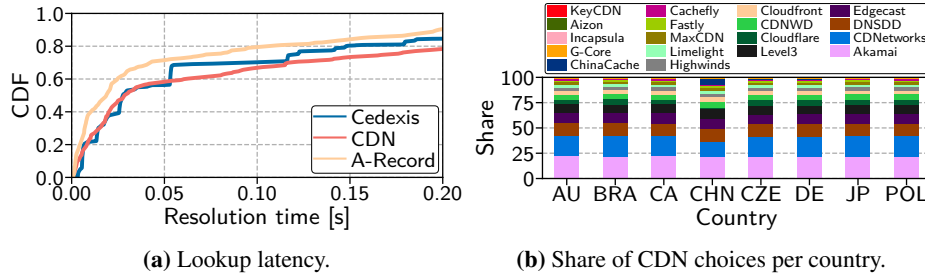


Fig. 6: Lookup latency in DNS resolution and final CDN choices of Cedexis customers.

CDN usage of Customers. We next check how many CDNs are being used by Cedexis customers. Figure 6(b) illustrates the overall CDN selection frequency for every PlanetLab and RPi node location over the period of one month for all discovered Cedexis domains. We find that domains are usually only using one or two CDNs while there are few that use more. This finding is consistent between geolocations. Only when looking at the how often which CDN is actually selected (note: Figure 5(b) did only show the share of domains using the CDN) we see a small geographic difference in China (CHN). Here ChinaCache is selected more often than in other geolocations, nevertheless, apart from this, all geolocations behave similarly. This finding contrasts a finding on the Conviva network [16] showing a bias in which some CDNs are selected more often than others in specific countries. In summary, we do not observe that Cedexis customer set country-specific routing decisions.

To extend our global view beyond the PlanetLab nodes, we next resolve the discovered Cedexis domains from open DNS resolvers obtained from `public-dns.info`. To avoid overloading (low-power) devices on user-premises (e.g., home routers), we exclude all resolvers whose DNS names indicate access lines (e.g., “pppoe”, “dial-up”, or “dsl”). We further only select resolvers with an availability $> 89\%$. In total, this leaves us with 1998 resolvers in 161 countries, out of which 67 *never* successfully responded. We resolve all Cedexis customer domains using all resolvers every two hours for four days. Subsequently, we group the reported results by continent and compare the top selected CDN. We observe that 66.9% always chose the same CDN in every continent. For the remaining, we observe disagreement, i.e., different CDN are chosen on each continent: 30.4% have two and 2.7% three CDNs present. We compare the complete CDN choices in countries of our PlanetLab nodes to their mapping results and observe similar distributions as in Figure 6(b) (not shown).

Takeaway. *Most Cedexis domains configure short TTLs to enable frequent switches. We observe that most domains indeed balance between few CDNs. Switches pose a challenge to each CDN since traffic gets harder to predict.*

4.3 Latency Perspective

We next take a latency perspective on Cedexis choices, i.e., comparing the latency of the chosen CDN to all configured CDNs for every customer domain. Thus, we measure the latency to every CDN IP by performing ICMP pings. We chose ICMP pings over more realistic HTTP requests since the pings do not generate accountable costs for the probed customers but remark that the ping latency can differ from actual HTTP response

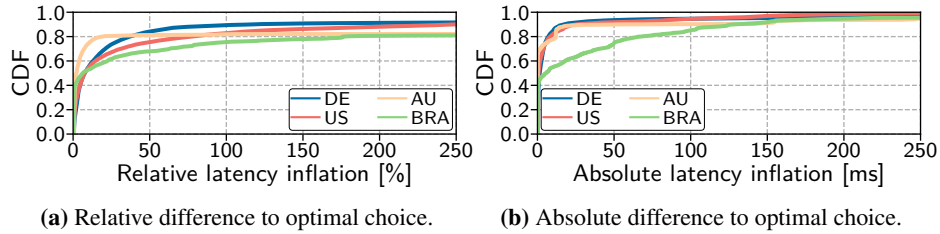


Fig. 7: Ping difference to the *fastest* CDN (RTT) when the choice was not optimal.

latencies. Yet, the ping latency can be a reasonable performance and distance indicator for the different CDN caches.

Figure 7(a) shows the relative latency inflation for cases where Cedexis did not choose the latency-optimal CDN. We observe that around 50% of all resolutions are only marginally worse than the optimal choice regardless of the geographic location (selection shown). The curves then start to flatten indicating increased latency inflation. We observe two groups, where around 10% (20%) of the choices exhibit a latency increase of over 250%. The observed increase can be in the range of few ms for nearby servers that would still deliver good performance. Therefore, we show the absolute difference of all sub-optimal decisions in Figure 7(b). We observe that $\approx 50\%$ of all decisions indeed only differ in a couple of milliseconds, indicating almost no noticeable difference. Apart from the nodes in Brazil and Japan (not shown), around 90% of all choices are even within a 20 ms difference.

We remark that it is difficult to assess the quality and correctness of CDN choices as we do not know the routing metrics that are employed by Cedexis customers. Our measurements are motivated from the perspective of an end-user, who is interested in performance, not in potentially monetary business decisions.

Takeaway. All available CDNs would deliver good latency figures in most tested cases, suggesting that the choice of CDN performed by Cedexis would not largely impact end-user experience.

5 Conclusion

This paper presents a broad assessment of a Meta-CDN deployment, exemplified by dissecting Cedexis as a representative generic platform. By this, we enrich the existing literature describing the Meta-CDN concept with an empirical assessment of a large operator. We shed light on Cedexis customers, technology, and performance. Cedexis DNS deployment, even though using anycast, appears to be focussing on Europe, North America and parts of Asia indicated by high latencies in other regions. We find customers to mostly configure short TTL values enabling fast reactions and we indeed observe that most domains balance between few CDNs. By assessing ping latencies to all available CDNs, we observe that most available CDNs offer good performance from our distributed probe platform. However, we also find a range of sub-optimal latency choices which can indicate routing metrics other than latency.

These unpredictable routing decisions implemented by customers using a Meta-CDN pose a challenge to CDN operators since inbound traffic gets much harder to predict. In particular routing decisions can be based on active measurements by the Meta-CDN—

thus, bad performance can result in rerouting traffic and thus losing revenue. Studying Meta-CDNs and their consequences thus pose an interesting angle for future work.

Acknowledgments

This work has been funded by the DFG as part of the CRC 1053 MAKI. We would like to thank Dean Robinson (Univ. Michigan) for his early contributions to the Alexa analysis, our shepherd Ignacio Castro and the anonymous reviewers.

References

1. Cedexis. <https://www.cedexis.com/>.
2. Cedexis Blog: DDoS Attack. www.cedexis.com/blog/ddos-attack-details/.
3. Cisco umbrella list of top 1M domains. <http://s3-us-west-1.amazonaws.com/umbrella-static/index.html>.
4. Conviva. <https://www.conviva.com/>.
5. Dyn DNS Traffic Steering. <http://dyn.com>.
6. FORTINET FortiDirector. <http://xdn.com/fd>.
7. V. K. Adhikari, Y. Guo, F. Hao, V. Hilt, and Z. L. Zhang. A Tale of Three CDNs: An Active Measurement Study of Hulu and Its CDNs. In *IEEE INFOCOM Workshops*, 2012.
8. M. Calder, A. Flavel, E. Katz-Bassett, R. Mahajan, and J. Padhye. Analyzing the Performance of an Anycast CDN. In *ACM IMC*, 2015.
9. F. Chen, R. K. Sitaraman, and M. Torres. End-User Mapping: Next Generation Request Routing for Content Delivery. In *ACM SIGCOMM*, 2015.
10. F. Dobrian, V. Sekar, A. Awan, I. Stoica, D. Joseph, A. Ganjam, J. Zhan, and H. Zhang. Understanding the Impact of Video Quality on User Engagement. In *ACM SIGCOMM*, 2011.
11. A. Flavel, P. Mani, D. Maltz, N. Holt, J. Liu, Y. Chen, and O. Surmachev. FastRoute: A Scalable Load-Aware Anycast Routing Architecture for Modern CDNs. In *USENIX NSDI*, 2015.
12. B. Frank, I. Poese, Y. Lin, G. Smaragdakis, A. Feldmann, B. Maggs, J. Rake, S. Uhlig, and R. Weber. Pushing CDN-ISP collaboration to the limit. *ACM SIGCOMM CCR*, 43(3), 2013.
13. A. Gerber and R. Doverspike. Traffic Types and Growth in Backbone Networks. In *IEEE OFC/NFOEC*, 2011.
14. C. Labovitz, S. Iekel-Johnson, D. McPherson, J. Oberheide, and F. Jahanian. Internet Inter-domain Traffic. In *ACM SIGCOMM*, 2010.
15. H. H. Liu, Y. Wang, Y. R. Yang, H. Wang, and C. Tian. Optimizing Cost and Performance for Content Multihoming. In *ACM SIGCOMM*, 2012.
16. M. K. Mukerjee, I. N. Bozkurt, B. Maggs, S. Seshan, and H. Zhang. The Impact of Brokers on the Future of Content Delivery. In *ACM HotNets*, 2016.
17. M. K. Mukerjee, I. N. Bozkurt, D. Ray, B. M. Maggs, S. Seshan, and H. Zhang. Redesigning CDN-Broker Interactions for Improved Content Delivery. In *ACM CoNEXT*, 2017.
18. E. Nygren, R. K. Sitaraman, and J. Sun. The Akamai Network: A Platform for High-performance Internet Applications. *SIGOPS OS Rev.*, 44(3):2–19, Aug. 2010.
19. J. S. Otto, M. A. Sánchez, J. P. Rula, and F. E. Bustamante. Content Delivery and the Natural Evolution of DNS: Remote Dns Trends, Performance Issues and Alternative Solutions. In *ACM IMC*, 2012.
20. I. Poese, B. Frank, B. Ager, G. Smaragdakis, and A. Feldmann. Improving Content Delivery Using Provider-aided Distance Information. In *ACM IMC*, 2010.
21. A.-J. Su, D. R. Choffnes, A. Kuzmanovic, and F. E. Bustamante. Drafting Behind Akamai: Inferring Network Conditions Based on CDN Redirections. *IEEE/ACM ToN*, 17(6), 2009.
22. R. van Rijswijk-Deij, M. Jonker, A. Sperotto, and A. Pras. A High-Performance, Scalable Infrastructure for Large-Scale Active DNS Measurements. *IEEE JSAC*, 34(6), 2016.
23. J. Xue, D. Choffnes, and J. Wang. CDNs Meet CN An Empirical Study of CDN Deployments in China. *IEEE Access*, 5, 2017.