

Maximizing Utility for Content Delivery Clouds

Mukaddim Pathan, James Broberg, and Rajkumar Buyya

Grid Computing and Distributed Systems (GRIDS) Laboratory
Department of Computer Science and Software Engineering
The University of Melbourne, Australia
{apathan, brobergj, raj}@csse.unimelb.edu.au

Abstract. A content delivery cloud, such as MetaCDN¹, is an integrated overlay that utilizes cloud computing to provide content delivery services to Internet end-users. While it ensures satisfactory user perceived performance, it also aims to improve the traffic activities in its world-wide distributed network and uplift the usefulness of its replicas. To realize this objective, in this paper, we measure the utility of content delivery via MetaCDN, capturing the system-specific perceived benefits. We use this utility measure to devise a request-redirection policy that ensures high performance content delivery. We also quantify a content provider's benefits from using MetaCDN based on its user perceived performance. We conduct a *proof-of-concept* testbed experiment for MetaCDN to demonstrate the performance our approach and reveal our observations on the MetaCDN utility and content provider's benefits for using it.

1 Introduction

Content Delivery Networks (CDNs) provide improved Web access performance to Internet end-users through multiple, geographically distributed replica servers [1]. A commercial CDN lock-in a customer, i.e. content provider, for a particular period of time under specific Service Level Agreements (SLAs) with a high monthly/yearly fees and excess data charges [8]. Thus, far from democratizing content delivery, most CDN services are often priced out of reach for all but large enterprise customers [2]. On the other hand, a commercial CDN realizes monetary penalization if it fails to meet the SLA-bound commitments to provide high quality service to end-users, due to the over-provisioning of its resources under heavy traffic surges. In the recent years, rising operational costs have led to the major consolidation in CDN market, down to a handful of key players. Moreover, the main value proposition for CDNs has shifted over time. Initial focus was on improving the user performance by decreasing response time, especially when there is an unexpectedly high load on content provider's Web sites. Recently content providers have viewed CDNs services as a way to use a shared network infrastructure to handle their peak capacity requirements, thus allowing reduced investments in their own Web site infrastructure [7].

One approach to mitigate these problems is to build a *content delivery cloud*, on top of existing *cloud services*, e.g. Amazon Simple Storage Service (S3), Nirvanix Storage Delivery Network (SDN), and Mosso Cloud Files. The use of clouds for content delivery is highly appealing as they charge customers for their utilization of storage and transfer of content, typically in order of cents per gigabyte. They offer SLA-backed performance and uptime guarantees for their services. Moreover, they can rapidly and cheaply scale-out during flash crowds [4] and anticipated increases in demand. However, unlike a fully-featured CDN, they do not provide capabilities for automatic replication, fail-over, geographical load redirection and load balancing.

¹ It extends traditional Content Delivery Networks (CDNs) model. Available at: <http://www.metacdn.org>

MetaCDN [3] realizes a content delivery cloud, providing the required features for high performance content delivery. It is an integrated overlay service, which leverages existing *storage clouds*. It allows content providers to reveal advanced content delivery services without having to build a dedicated infrastructure. When a user requests content, MetaCDN chooses an optimal replica for content delivery, thereby ensuring satisfactory user perceived experience. The ultimate goal is to improve the *utility* of MetaCDN's content delivery services. In our context, utility is the quantification of the traffic activities in MetaCDN and represents the usefulness of its replicas in terms of data circulation in its world-wide distributed network. While the notion of utility is enticing in relation to content delivery, only a few previous work [5][6][7][18] have considered it. Our work is in line with them, focusing on maximizing the utility of MetaCDN. We present a request-redirection policy based on the measured utility to ensure high content delivery performance. We also quantify a content provider's benefits from using MetaCDN. We perform a *proof-of-concept* empirical study on a global testbed to evaluate the performance of our approach and provide insights on the content-serving utility of MetaCDN. The main contributions of this paper are:

- A utility-based request-redirection policy to improve MetaCDN's utility.
- An approach to quantify a content provider's benefits from using MetaCDN.
- An experimental study on a world-wide testbed for performance evaluation.

The rest of the paper is structured as follows. Section 2 provides an overview of related work. It is followed by a brief description of MetaCDN in Section 3. The formulation of utility metric and the devised request-redirection policy is presented in Section 4. Section 5 presents a description of the testbed experiments. Next, in Section 6, empirical results are discussed. Finally, the paper is concluded in Section 7.

2 Related Work

There has been a growing interest in interconnecting provider capabilities in CDNs, such as Content Distribution Internetworking (CDI) [9], multi-provider peering [13], and peering CDNs [11]. These research efforts explore the benefits of internetworking of CDNs, with main focus on offering increased capacity, intelligent server selection, reduced cost, and better fault tolerance. In contrast, our work with the MetaCDN overlay system assumes no cooperation or peering. Rather it follows a brokering approach as in CDN brokering [10]. Our work is a logical fit in existing storage cloud deployments coupled with content delivery capabilities, such as Amazon CloudFront; VoxCAST CDN; Mosso Cloud Files, which leverages content delivery services from Limelight Networks; Nirvanix SDN, which partners with CDNetworks for content delivery; and Edge Content Network (ECN) from Microsoft, which is reported to partner with Limelight Networks for content delivery [15].

Along with the trend-shift in the content delivery space, CDN utility and pricing have gained notable attention. There has been prior work reflecting the utility computing notion for content delivery [7][18][19]. They mostly provide description of architecture, system features and challenges related to the design and development of a utility computing platform for CDNs. On the contrary, we not only provide an overview of the utility model for MetaCDN, but also quantify the perceived utility and devise a utility-based request-redirection policy. Our work is complimentary to the simulation-based evaluation of utility as described in previous work [5][6]. We differ by providing a proof-of-concept implementation for evaluating utility, devising a request-redirection algorithm, and revealing the system performance through a global testbed experiment. In addition, we determine a content provider's perceived utility using an approach complimentary to traditional CDN pricing [8].

Recent innovations such as P4P [12] enables P2P to communicate with network providers through a portal for cooperative content delivery. Our work endorses them in that the MetaCDN overlay system assists toward a systematic understanding and practical realization of the interactions between storage clouds, who provide an operational storage network and content delivery resources, and content providers, who generate and distribute content.

3 The MetaCDN Overlay

Figure 1 provides an illustration of MetaCDN. It is coupled with each storage clouds via *connectors* that provide an abstraction to hide the complexity associated with different access methodologies of heterogeneous providers. End-users can access the MetaCDN overlay either through a Web portal or via RESTful Web services. In the first case, the Web portal acts as an entry point to the system and performs application level load balancing for end-users who intend to download content that has been deployed through MetaCDN. Content providers can sign up for an account on the MetaCDN system and enter credentials for any storage cloud providers that have an account with. Upon authentication, they can utilize MetaCDN functionalities to deploy content over geographically spanned replicas from multiple storage clouds, according to their performance requirements and budget limitations.

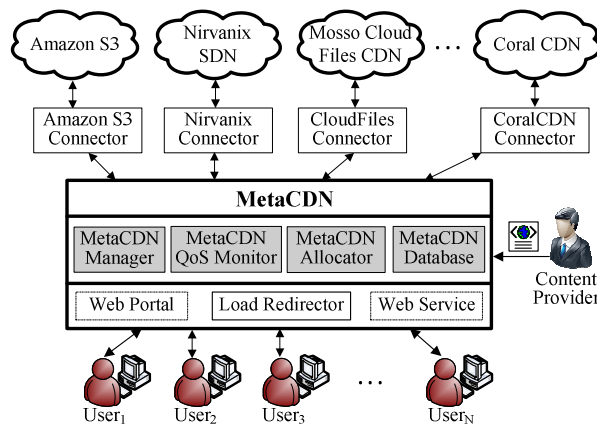


Fig. 1. An abstract view of MetaCDN.

MetaCDN provides the logic and management required to encapsulate the functionality of upstream storage providers with a number of core components. The *MetaCDN allocator* performs optimal provider selection and physical content deployment using four options, namely, *maximize-coverage*, *geolocation-based*, *cost*, and *QoS-optimized* deployment. The *MetaCDN QoS monitor* tracks the current and historical performance of participating storage providers. The *MetaCDN Manager* has authority on each user's current deployment and performs various housekeeping tasks. The *MetaCDN Database* stores information, such as user accounts and deployments, and the capabilities, pricing and historical performance of providers, which are required to ensure reliable and persistent operation of the system. Finally, the *MetaCDN Load Redirector* is charged with different redirection policies and is responsible for directing end-users to the most appropriate replica according to performance requirements. Details on functionalities, architecture, and development methodology of MetaCDN can be found in a previous work [3].

4 MetaCDN Utility

End-users experience little of the complex technologies associated with the content delivery services of MetaCDN. Content providers interact with the service in a limited number of ways, such as enabling their content to be served, viewing traffic reports, and receiving usage-based billing. Since the responsibility to ensure high performance content delivery is largely on the MetaCDN system itself, we aim at improving its content-serving utility. In this section, we formulate the utility maximization problem with quantitative expressions, and devise a utility-based request-redirection policy.

We use $R = \{r_i\}$, $i \in \{1, 2, \dots, M\}$ to denote the set of user requests, with r_i being the i -th arriving request to the MetaCDN overlay system, comprising a set of N replicas. Utility maximization in MetaCDN can be achieved from two perspectives. The first aspect is the profit maximization of MetaCDN, which is formulated as:

$$\text{maximize} \quad \sum_{r_i \in R, j \in N} U_{mcdn} x_{ij} \quad (\textit{Profit}) \quad (1)$$

where U_{mcdn} is content-serving utility of the MetaCDN overlay system.

The second aspect examines the general welfare of the content provider for using the MetaCDN infrastructure to maximize its own benefit. Each content provider obtains a perceived utility U_{CP} (benefits) for QoS-constrained content delivery to its end-users via MetaCDN. The measured utility is expressed as the fraction of processed requests (throughput) or the total valuation (weighted throughput). It can be formulated either by maximizing the following one or two measurements:

$$\text{maximize} \quad \sum_{r_i \in R, j \in N} x_{ij} \quad (\textit{Throughput}) \quad (2)$$

$$\text{Maximize} \quad \sum_{r_i \in R, j \in N} U_{CP} x_{ij} \quad (\textit{Weighted throughput}) \quad (3)$$

where U_{CP} is the content provider's perceived utility; indicator variable $x_{ij}=1$ if MetaCDN replica j serves request r_i within service requirements; and $x_{ij}=0$ otherwise.

4.1 Content-Serving Utility

We now derive the quantitative measure for MetaCDN utility U_{mcdn} and content providers perceived utility U_{CP} . Ideally for the MetaCDN overlay system, the most useful replicas are those exhibiting the highest utility. We quantify utility with a value that expresses the relation between the number of bytes of the served content against the number of bytes of the replicated content. It ranges in $[0, 1]$ and provides an indication of the traffic activities. Formally, utility of a MetaCDN replica i is:

$$u_i = (2/\pi) \times \arctan(\xi) \quad (4)$$

The main idea behind this metric is that a replica is considered to be useful (high utility) if it serves content more than it replicates, and vice versa. The parameter ξ is the ratio of the serviced bytes to the replicated bytes, i.e.

$$\xi = \textit{No of bytes serviced} / \textit{No of bytes replicated} \quad (5)$$

The resulting utility from (4) ranges in $[0, 1]$. The value $u_i=1$ is achieved if the replica only serves content, without replicating ($\xi=\textit{infinity}$). It results when a replica already has the content, and does not replicate a new copy of the content for serving successive requests. On the contrary, the value $u_i=0$ is achieved if the replica has the

content, however fails to serve ($\zeta=0$) due to service over-provisioning and/or network perturbations under heavy traffic surges. The content-serving utility U_{mcdn} of MetaCDN can be expressed as a mean value of the individual replica utilities, i.e.

$$U_{mcdn} = \sum_{n=1}^N u_i / N \quad (6)$$

MetaCDN outsources customer's (content provider) content to the replicas and it is charged by the cloud providers based on usage. Since the utility measure captures the usage of storage cloud resources, the measured value can be easily translated into a price of the offered services. The resulting price could be used to derive a content provider's benefits or perceived utility. We draw inspiration from Hosanagar et al. [8], which show that the benefits of a content provider depends on its revenue, benefit from content delivery to its end-users through a CDN, replication cost, and usage-based charges. We adopt this approach by using performance measures of end-users that belong to a content provider. We gauge the throughput for serving requests and the response time improvement by using MetaCDN over direct replica access. We also measure the replication cost and interpret the pricing of storage clouds according to the MetaCDN utility. We express a content provider's perceived utility as:

$$U_{cp} = T(X) + (R - R_d) \times X - \psi b - P(u) \quad (7)$$

where $T()$ is the weighted throughput for X requests; R and R_d respectively are the user perceived response times from direct replica access and via MetaCDN; ψ is the unit replication cost; b is the content size; $P(u)$ is the utility-based pricing function.

4.2 Utility-based Request-Redirection

To choose the optimal replica for an end-user request, the MetaCDN Load Redirector module evaluates the utility metric reflecting the state of its replicas and the network conditions between users and replica sites. The measured utility is used for a utility-based request-redirection policy (Figure 2) for MetaCDN to serve data-rich content, thus improving the content delivery performance by relieving network congestions.




Fig. 2. Utility-based request-redirection in MetaCDN.

The sequence of steps for an end-user in the East Coast of the USA to retrieve content through MetaCDN is as follows:

1. The end-user requests for a MetaCDN URL such as *http://www.metacdn.org/FileMapper?itemid=1*, and the browser attempts to resolve the base hostname *www.metacdn.org*.
2. The user DNS eventually contacts the authoritative DNS (ADNS) for that domain to resolve this request to the IP address of the closest MetaCDN gateway, e.g. *http://us.metacdn.org*.

3. The user (or browser) then makes an HTTP GET request for the desired content on the MetaCDN gateway.
4. In the case of utility-based request-redirection, the MetaCDN Load Redirector is triggered to select the highest-utility optimal replica that conforms to the specified service requirements. At this point, the MetaCDN gateway returns an HTTP redirect request with the URL of the selected replica. The following tests are performed to determine the best replica for serving user requests:
 - Is there a content replica available within required response time threshold?
 - Is the throughput of the target replica within tolerance?
 - Is the end-user located in the same geographical region as the target replica?
 - Is one of the target replicas preferred, according to user requirements or any administrative settings?
 - Is the replica utility the highest among all target sites?
 - If there is more than one replica with same highest utility, which replica site provides the fastest response time?
5. Upon receiving the URL of the selected replica, the DNS resolves its domain name and returns the associated IP address to the end-user.
6. The user sends the request for the content to the selected replica.
7. The selected replica satisfies the end-user request by serving the desired content.

If it is assumed that all candidate target replicas are available and have capacity, i.e. response time and throughput thresholds are met, the MetaCDN system checks for the continent/geographic location and administrative preference (an indicative flag used by MetaCDN manager to manually prefer or avoid a replica). While end-users are directed to the highest utility replica by the MetaCDN Load Redirector, if there is more than one candidate target replica exhibiting the highest utility, the one with the fastest response time is chosen to redirect user requests. In addition, a secondary level of internal redirection enabled by an individual cloud provider ensures that the utility-based request-redirection policy does not overload any particular replica.

While the above request-redirection policy directs users to the best responding replica, an extra feature is realized through its ability to automatically avoid failed replicas or replicas without the desired content. Bypassing occurs in the following two ways. Firstly, if a replica has the desired content, but shows limited serving capacity due to network congestions, it is reflected in its measured utility and exhibits a low value. As a consequence, the replica is not considered as a candidate for redirection. Secondly, if the replica does not have the desired content, it can not serve user requests and thus leads to a content-serving utility of 0. Hence, it is automatically discarded to be considered as a candidate replica.

5 Methodology

We conduct a *proof-of-concept* testbed experiment to determine the content delivery utility of MetaCDN, evaluate the performance of the utility-based redirection policy and measure the user perceived response time and throughput. Figure 3 provides a schematic representation of the experimental testbed and Table 1 provides a summary of the conducted experiment. The global testbed spans six continents with distributed clients, replicas and MetaCDN gateways. All client locations, except Africa, South America and South Asia, have high speed connectivity to major Internet backbones to minimize the client being the bottleneck during experiments. We used test files of size 1KB and 5MB, deployed by the *MetaCDN Allocator* module, which was instructed to maximize coverage and performance, and consequently the test files were deployed in all available replica locations of the storage cloud providers integrated to MetaCDN.

While these file sizes are appropriate for our experiments, a few constraints restrict us to use varied and/or even larger sized files. Firstly, the experiments generate heavy network traffic consuming significant network bandwidth, thus larger file trafficking would impose more strain and network congestions on the voluntary clients, which some clients may not be able to handle. Moreover, at some client locations, e.g. India and South Africa, Internet is at a premium and there are checks regarding Internet traffic so that other users in the client domain accessing Internet are not affected.

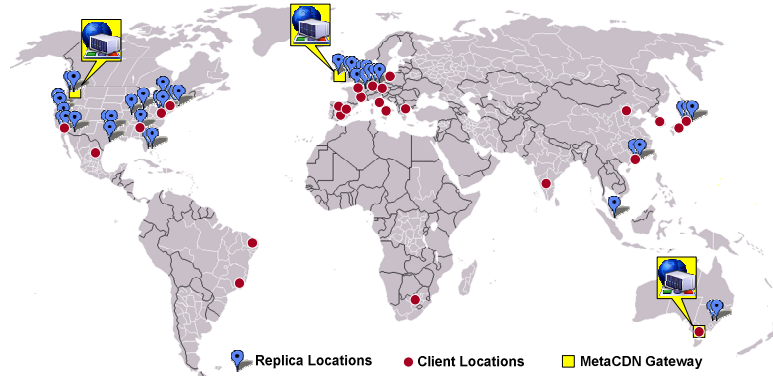


Fig. 3. Experiment testbed.

The experiment was run simultaneously at each client location over a period of 48 hours, during the middle of the week in May, 2009. As it spans two days, localized peak times (*time-of-day*) is experienced in each geographical region. The workload to drive the experiment incorporates recent results on Web characterization [4][16][17]. The high variability and self-similar nature of Web access load is modeled through heavy-tailed distributions. The experiment time comprises epochs of 2 hours, with each epoch consisting of a set of user sessions. Each session opens a persistent HTTP connection to MetaCDN and each client generates requests to it to download each test files, with a timeout of 30 seconds. Between two requests, a user waits for a *think time* before the next request is generated. The mean think time, together with number of users defines the mean request arrival rate to MetaCDN. For statistical significance, each client is bounded to generate a maximum number of 30 requests in each epoch. The files are downloaded using the UNIX utility, *wget*, with the *--no-cache* and *--no-dns-cache* options to ensure that a fresh copy of the content is downloaded each time (not from any intermediary cache) and that the DNS lookup is not cached either.

5.1 Schemes and Metrics for Comparison

The primary objectives are to measure MetaCDN utility, evaluate performance of the proposed utility-based request-redirection policy, and provide observations on how MetaCDN's content-serving ability is varied during the experiment. For performance comparison, we experiment with two other request-redirection policies, namely, *random* and *geo-redirection*. The first policy directs an end-user to a randomly picked replica, whereas the second policy takes into account user preferences and directs him/her to the closest physical replica in the specified region(s).

We measure the *response time* and *throughput* obtained from each client location. The first captures the end-to-end performance for users when downloading a 1KB test file from MetaCDN. Due to the negligible file size, the response time is dominated by DNS lookup and HTTP connection establishment time. Lower value of response time

indicates fast serviced content. The latter shows the transfer speed obtained when the 5MB test file is downloaded by the users from the MetaCDN replica infrastructure. It provides an indication of consistency and variability of throughput over time.

Table 1. Summary of the experiment.

Experiment Testbed	<i>Category</i>	<i>Value</i>	<i>Provider</i>	<i>Locations</i>
	Number of MetaCDN gateways	3	Amazon EC2 and own cluster	Asia/Australia, Europe, and North America
	Number of replicas	40	Amazon, Mosso and Nirvanix	Asia, Australia, Europe, and North America
	Number of clients (end-user nodes)	26	Voluntary	Asia, Australia, Europe, North and South America, and Africa
Experiment Details	<i>Category</i>		<i>Description</i>	
	Total experiment time		48 hours	
	Duration of an epoch		2 hours	
	Maximum user requests/epoch		30 requests from each client	
	Service timeout for each request		30 seconds	
	Test file size		1KB and 5MB	
	Content Deployment		Maximize-coverage deployment	
Request-redirection policies		Random, Geo, and Utility		
End-user Request Modeling	<i>Category</i>	<i>Distribution</i>	<i>PMF</i>	<i>Parameters</i>
	Session inter-arrival time [16]	Exponential	$\lambda e^{-\lambda x}$	$\lambda = 0.05$
	Content requests per session [4]	Inverse Gaussian	$\frac{\lambda}{\sqrt{2\pi x^3}} e^{-\frac{\lambda(x-\mu)^2}{2\mu^2 x}}$	$\mu = 3.86, \lambda = 9.46$
	User think time [17]	Pareto	$\alpha k^\alpha x^{-\alpha-1}$	$\alpha = 1.4, k = 1$

The *utility* of MetaCDN is measured according to the model in Section 4.1. A high utility value shows the content-serving ability of the system, and signifies its durability under highly variable traffic activities. To emphasize the impact of request-redirection on the measured utility, we use the *probability* that MetaCDN achieves a given level of utility as the performance metric. Finally, based on the measured observations, we determine the benefits of a content provider from using MetaCDN. Table 2 summarizes the performance indices used in the experimental evaluation.

Table 2. List of performance indices.

Performance Index	Description
Response time	The time experienced by a end-user to get serviced
Throughput	Transfer speed to download a test file by a end-user
Utility	Content-serving ability, ranges in [0, 1]
Prob (Utility achieved)	The probability or the fraction of time that the system achieves the given utility
Content provider's benefit	Surplus from using MetaCDN, expressed as a percentage

6 Empirical Results

Due to space constraints, we present results from the following eight representative clients in five continents: Paris (France), Innsbruck (Austria), and Poznan (Poland) in Europe; Beijing (China) and Melbourne (Australia) in Asia/Australia; Atlanta, GA, Irvine, CA (USA) in North America, and Rio de Janeiro (Brazil) in South America.

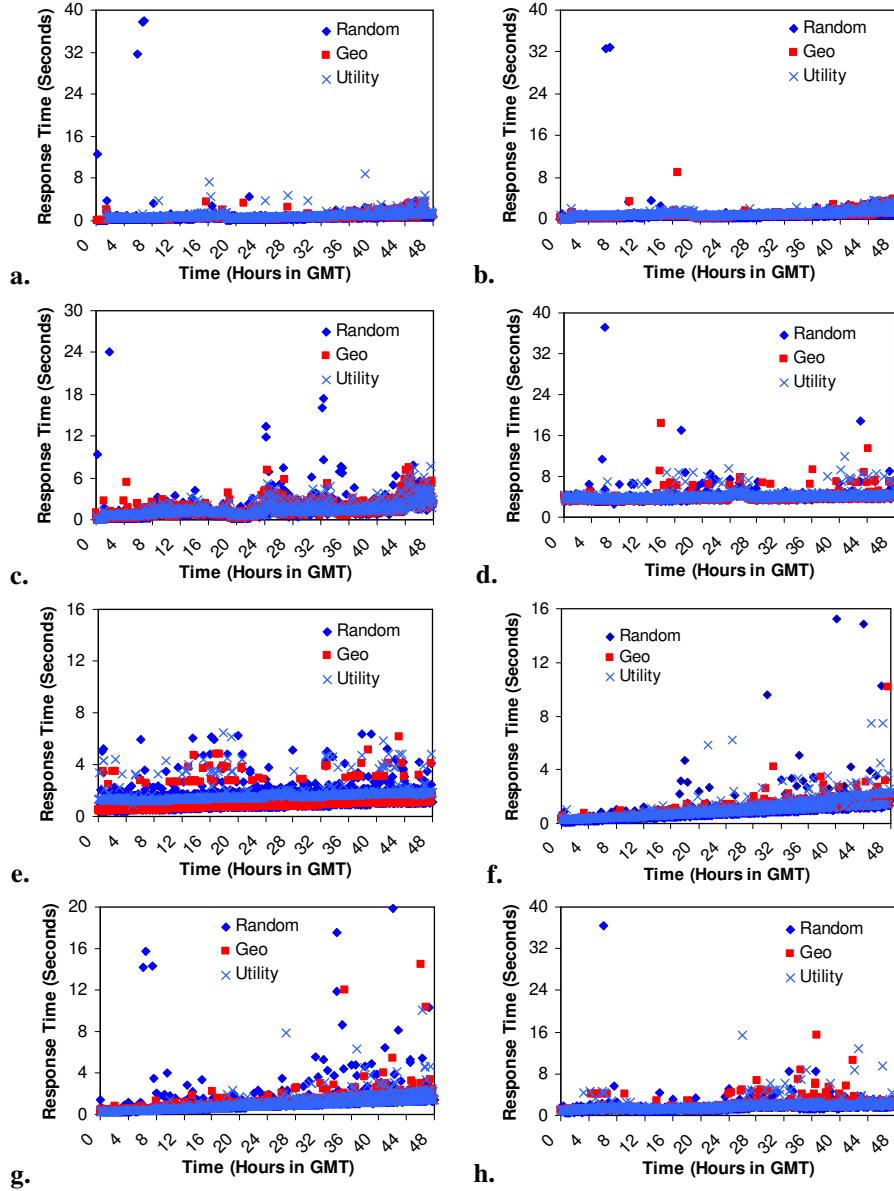


Fig. 4. Response time obtained in each client locations—(a) Paris; (b) Innsbruck; (c) Poznan; (d) Beijing; (e)Melbourne; (f) Atlanta; (g) Irvine; and (h) Rio de Janeiro.

6.1 Response Time Observations

Figure 4 shows the end-to-end response time experienced by end-users when downloading the 1 KB test file over a period of 48 hours. The measure of the response time depends on the network proximity, congestions in network path and traffic load of on the target replica server. It provides an indication of the responsiveness of the

replica infrastructure and the network conditions in the path between the client and the target replica which serves the end-user. We observe a general trend that the clients experience mostly consistent end-to-end response time. For all the request-redirection policies, the average response time in all the client locations except Beijing is just over 1 second, with a few exceptions. Notably the users in Beijing experience close to 4 seconds average response time from the MetaCDN replica infrastructure. This exception originates as a consequence of firewall policies applied by the Chinese government, which is also reported in another work [20].

At several time instances during the experiment, end-users experience increased response time. The resulting spikes are due to the sudden increases in request traffic, imposing strain on the MetaCDN replica infrastructure. Under traffic surges, the MetaCDN load redirector module activates to handle peak loads. As a consequence, user requests are often redirected to a target replica outside its authoritative domain and/or are served from an optimal distant proximity server(s), thereby, contributing to the increased response time. However, MetaCDN handles peak loads well to provide satisfactory responsiveness to end-users. This phenomenon of increased response time is more visible for the random-redirection policy. As it makes a random choice, often the target replica selection is not optimized, thus leading to highly variable response time. Especially, at several occasions, users observe more than 30 seconds response time, thus leading to service timeout. The geo-redirection policy directs request to the closest proximity server, understandably producing low response time. On the contrary, utility-redirection chooses the highest utility replica, which may not be in close proximity to an individual client location. Nevertheless, we do not find a clear winner between them in terms of response time, as they exhibit changeable performance in different client locations. The utility-redirection performs as well as geo-redirection in that we observe similar performance in all client locations except Paris and Melbourne. End-users in Paris enjoy better average response time (0.77 seconds) with geo-redirection, due to their close proximity to the Amazon, Mosso and Nirvanix nodes in Frankfurt (Germany), Dublin (Ireland), and London (UK). For Melbourne, the reason behind better performance of geo-redirection is the existence of the Mosso node in Sydney. For both of these two clients, utility-redirection directs requests to a distant replica than the closest one and results in increased response time

6.2 Throughput Observations

Figure 5 shows the average throughput obtained per two hours, when downloading content from the MetaCDN replica infrastructure. As expected, we observe that in almost all the client locations, geo-redirection results in highest throughput as the users get serviced from the closest proximity replica. However, it performs worse than random-redirection for the Irvine client. The reason is that random-redirection decision in this location most of the time selects closer proximity Amazon replica(s) with better network path than that of geo-redirection, which chooses Mosso replica. Moreover, the service capability from these two replicas and the network path between the replica and client also contribute to the observed throughput variations.

For most of the clients, except Rio de Janeiro, utility redirection performs much worse than geo-redirection. The reason is understandable, as utility-redirection emphasizes maximizing MetaCDN's utility rather than serving an individual user, thus sacrificing end-user perceived performance. For Rio de Janeiro, geo-redirection leads to closest Mosso node in the USA, whereas utility-redirection results in more utility-aware replica, which is Amazon node(s) in the USA. It could be presumed that Amazon node supersedes the Mosso node in terms of its service capability, better network path, internal overlay routing, and less request traffic strain.

It is observed that users in Poznan enjoy the best average throughput, which is 9MB/s for geo-redirection. The reason is that the client machine is in a MAN network, which is connected to the country-wide Polish optical network PIONEER with high capacity channels dedicated to content delivery traffic. Another client location with high throughput is Atlanta, which achieves speeds of approximately 6.2 MB/s for geo-redirection and 3.3 MB/s for utility-redirection, due to the existence of better network path between the client and the MetaCDN replica infrastructure. This reasoning is deemed valid, since there are Mosso nodes in the same location.

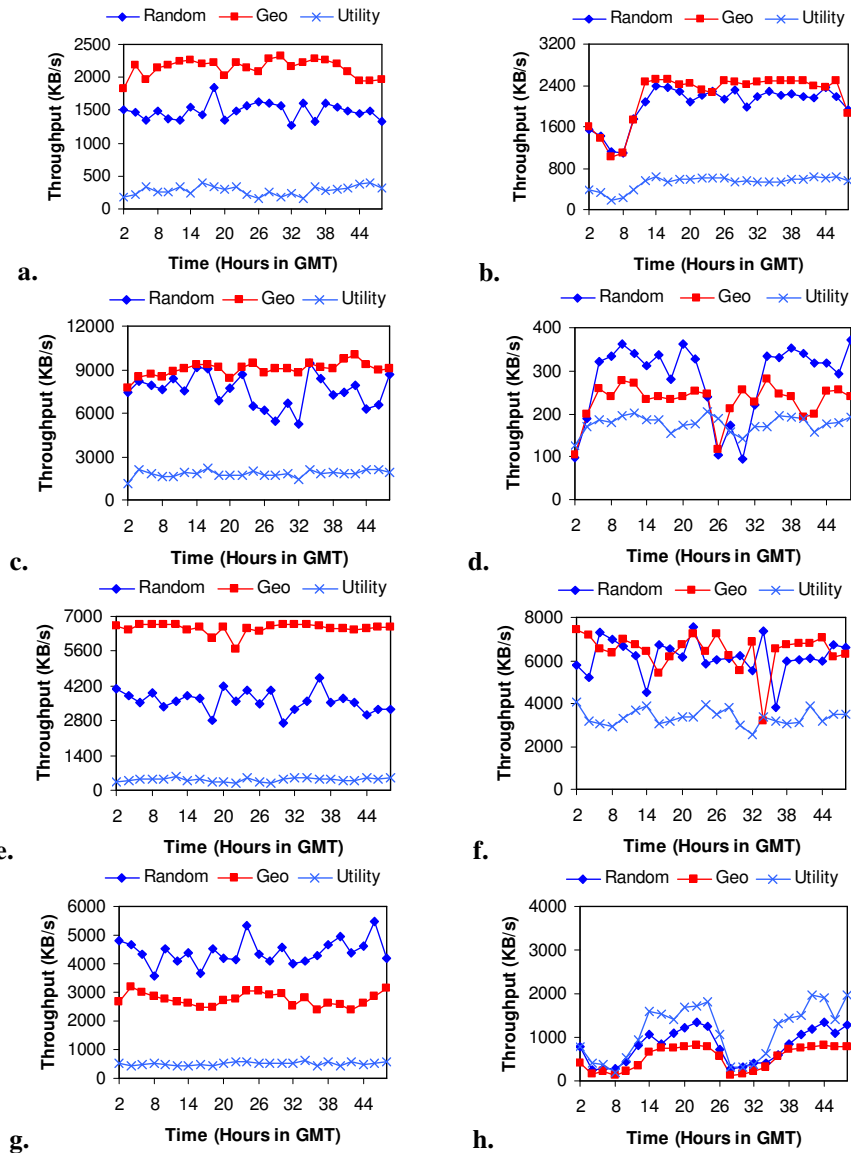


Fig. 5. Average throughput obtained in client locations—(a) Paris; (b) Innsbruck; (c) Poznan; (d) Beijing; (e) Melbourne; (f) Atlanta; (g) Irvine; and (h) Rio de Janeiro.

Alike the perceived response time, end-users in China achieves the lowest throughput among all the client locations. The underlying reason is again checks on the request traffic due to firewall policies. We put more emphasis on the results from Melbourne, which is of interest as Australia is not as highly connected as Europe or North America, depending on a small number of expensive international links to major data centers in Europe and the USA. Specifically, we observe that due to the existence of a nearby Mosso node in Sydney, the users in Melbourne experience 6.5 MB/s of throughput with geo-redirection and 3.6 MB/s for random-redirection. However, for utility-redirection the replica selections result in Amazon node(s) in the USA, thus leading to a lower but consistent average throughput of 410 KB/s.

From these observations, we come to the following decisive conclusions. Although utility-redirection outcomes sensible replica selection in terms of response time, it may not provide a high throughput performance to end-users. Nevertheless, being it is focused on maximizing the utility of the MetaCDN system; it results in high utility for content delivery. We provide sufficient results to support this claim in the next section.

6.3 MetaCDN Utility

Figure 6 shows how MetaCDN utility is varied during the testbed experiment upon replica selection for incoming content requests. Here we have used utility values averaged over three deployed MetaCDN gateways in Asia/Australia, Europe and the USA. We observe that utility-redirection produces the highest utility in the system by selecting most active replicas to serve end-users. It also improves the traffic activities in the system and contributes to uplifting MetaCDN system's content-serving ability. It should be noted that there is a warm-up phase at the beginning of the 48 hours experiment during which the replicas are populated with content requests, resulting in low utility value. This is visible during the initial hours for utility and geo-redirection.

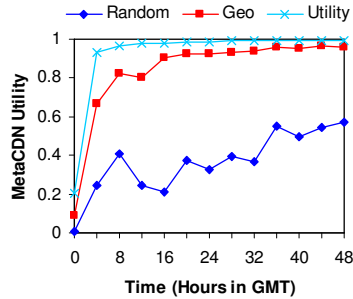


Fig. 6. MetaCDN utility over time.

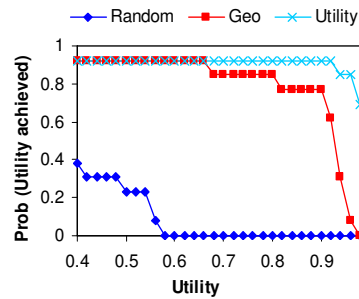


Fig. 7. Probability of achieving specified utility.

In order to emphasize the content-serving ability of MetaCDN, we now present the probability that the system can achieve required minimum utility. The intention is to show to what extent the system can maximize its own profit. Figure 7 presents the probability (or the fraction of time) that the system observes a utility above a certain utility level during the experiment. The higher the probability, the more likely it is that the specified utility level could be achieved. From the figure, it is noticeable that utility-redirection outperforms other alternatives, as it often produces over 0.95 utility for MetaCDN with a 0.85 probability. Geo-redirection performs well as it has a 0.77 probability that it can achieve 0.9 utility. Finally, random-redirection performs the worst and it can only achieve close to 0.56 utility for MetaCDN with a probability of 0.23. Therefore, a MetaCDN administrator may utilize a redirection policy apart from random, in order to maximize the system's content-serving ability.

Table 3. Content providers’ benefits based on user perceived performance.

Content Provider	End-user Location	Average Response Time	Average Direct Replica Access Time	Throughput	Perceived Utility (%)
1	Paris	0.99 seconds	0.747 seconds	475.39 KB/s	13.31
2	Innsbruck	1.03 seconds	0.955 seconds	518.67 KB/s	29.29
3	Poznan	1.52 seconds	0.667 seconds	1.80 MB/s	68.99
4	Beijing	4.17 seconds	1.337 seconds	176.54 KB/s	58.14
5	Melbourne	1.72 seconds	0.75 seconds	413.15 KB/s	25.32
6	Atlanta	1.09 seconds	0.605 seconds	3.35 MB/s	66.47
7	Irvine	1.06 seconds	0.391 seconds	504.74 KB/s	25.38
8	Rio de Janeiro	1.81 seconds	1.17 seconds	1.14 MB/s	34.02

6.4 Content Provider’s Perceived Utility

We now shed light on the content provider’s perceived utility or benefits (Section 4.1) from using MetaCDN, as reported in Table 3. For this purpose, we consider a scenario with 8 client locations belonging to 8 content providers. We use weighted throughput and normalized values of perceived response times from the client locations of respective content providers, using utility-redirection. We also measure the direct replica access time from each of the client locations. Finally, we make use of pricing information of Amazon S3, as reported earlier [3]. The perceived utilities reported in Table 3 are to be considered representative, as they could be varied depending on the heterogeneous pricing structure of different storage cloud providers. In addition, different redirection policy leads to different perceived utility for a content provider.

We observe that a content provider’s utility is heavily dependent on the throughput that its end-users receive. Therefore, content providers whose end-users benefit from high throughput also realize high quantitative benefit from using MetaCDN. As Poznan and Atlanta experience highest throughputs, the content provider’s surplus for this locations are also highest. The average throughput experienced in Paris is the lowest, thus leading to the least content provider’s surplus. It is important to note that the utility-redirection policy, being antagonistic to content provider’s utility, can still assist in resulting reasonable perceived utility for the content providers. Thus, the above results show that the MetaCDN system is helpful for content providers, even at times when it uses redirection techniques to maximize its own utility.

7 Conclusion and Future Work

MetaCDN is an integrated overlay service over existing storage clouds to provide high performance content delivery. It provides sensible performance and availability to the content providers without requiring content providers to build or manage complex content delivery infrastructure themselves. In this paper, we have presented an approach to maximize the utility for content delivery via MetaCDN. The utility metric reflects the traffic activities in the system and exhibit the usefulness of its replica infrastructure. This quantitative measure is vital as system wellness greatly affects the content delivery performance to end-users. We have used the measured utility to devise a redirection policy and quantify the benefit of a content provider for using MetaCDN. We conducted experiments in a global testbed to evaluate the performance of our approach. From the results obtained, we conclude that the utility of MetaCDN is maximized with utility-redirection with sensible replica selection and consistent average response time, however, with the cost of lower throughput in comparison to other policies. In contrast, a content provider’s benefit is improved

with improvement of the perceived throughput through MetaCDN. Therefore, a MetaCDN administrator should use a redirection policy based on the objective of either maximizing system utility or content provider's utility. The results are quite encouraging to conduct a set of future work, which include devising advanced redirection and pricing policies to benefit both MetaCDN and content providers.

Acknowledgments. We thank all participants² of the testbed experiments for their assistance with providing client nodes to run the experiments from. This work is supported by ARC and DIISR research grants.

References

- [1] G. Pallis and A. Vakali, "Insight and perspectives for content delivery networks," *Communications of the ACM*, 49(1), ACM Press, NY, USA, pp. 101-106, Jan. 2006.
- [2] D. Rayburn, "CDN pricing: Costs for outsourced video delivery," In *Streaming Media West: The Business and Technology of Online Video*, Sept. 2008.
- [3] J. Broberg, R. Buyya, and Z. Tari, "MetaCDN: Harnessing 'Storage Clouds' for high performance content delivery," *Journal of Network and Computer Applications (JNCA)*, Elsevier, Amsterdam, The Netherlands, 2009. To appear.
- [4] M. Arlitt, and T. Jin, "Workload characterization of the 1998 world cup Web site," *IEEE Network*, 14(3), pp. 30-37, 2000.
- [5] B. Mortazavi and G. Kesidis, "Model and simulation study of a peer-to-peer game with a reputation-based incentive mechanism," In *Proc. of the Information Theory and Applications (ITA) Workshop*, UC San Diego, Feb. 2006.
- [6] K. Stamos, G. Pallis, A. Vakali, and M. D. Dikaiakos, "Evaluating the utility of content delivery networks," In *Proc. of the 4th UPGRADE-CN*, ACM Press, NY, USA, Jun. 2009.
- [7] P. Gayek, R. Nesbitt, H. Pearthree, A. Shaikh, and B. Snitzer, "A Web Content Serving Utility," *IBM Systems Journal*, 43(1), pp. 43-63, 2004.
- [8] K. Hosanagar, J. Chuang, R. Krishnan, M. D. Smith, "Service adoption and pricing of content delivery network (CDN) services," *Management Science*, 54(9), pp. 1579-1593, Sept. 2008.
- [9] M. Day, B. Cain, G. Tomlinson, and P. Rzewski, "A model for content internetworking," *IETF RFC 3466*, Feb. 2003.
- [10] A. Biliris, C. Cranor, F. Douglis, M. Rabinovich, S. Sibal, O. Spatscheck, and W. Sturm, "CDN brokering," *Computer Communications*, 25(4), pp. 393-402, 2002.
- [11] M. Pathan, J. Broberg, K. Bubendorfer, K. H. Kim, and R. Buyya, "An architecture for virtual organization (VO)-based effective peering of content delivery networks," In *Proc. of the 2nd UPGRADE-CN*, ACM Press, NY, USA, Jun. 2007.
- [12] H. Xie, Y. R. Yang, A. Krishnamurthy, Y. Liu, and A. Silberschatz, "P4P: Provider portal for (P2P) applications," In *Proc. of ACM SIGCOMM*, ACM Press, NY, USA, Aug. 2008.
- [13] L. Amini, A. Shaikh, and H. Schulzrinne, "Effective peering for multi-provider content delivery services," In *Proc. of 23rd Annual IEEE Conference on Computer Communications (INFOCOM'04)*, IEEE CS Press, Los Alamitos, CA, USA, pp. 850-861, 2004.
- [14] R. Cohen, "Content delivery cloud (CDC)," In *ElasticVapor: Life in the Cloud*, Oct. 2008. <http://www.elasticvapor.com/2008/10/cloud-content-delivery-cd.html>.
- [15] R. Miller, "Microsoft building own CDN network," In *Data Center Knowledge*, Jan. 2008.
- [16] S. Floyd and V. Paxson, "Difficulties in simulating the Internet," *IEEE/ACM Transactions on Networking*, 9(4), pp. 392-403, 2001.
- [17] P. Barford and M. E. Crovella, "A performance evaluation of Hyper Text Transfer Protocols," In *Proc. of ACM Sigmetrics*, ACM Press, NY, USA, pp. 188-197, May 1999.
- [18] S. R. Subramanya and B. K. Yi, "Utility model for on-demand digital content," *IEEE Computer*, 38(6), pp. 95-98, 2005.
- [19] C. Canali, M. Rabinovich, and Z. Xiao, "Utility computing for Internet applications," In *Web Content Delivery*, vol. II, X. Tang, J. Xu, and S. T. Chanson, Eds. Springer, 2006, pp. 131-151.
- [20] H. Rahul, M. Kasbekar, R. Sitaraman, and A. Berger, "Towards realizing the performance and availability benefits of a global overlay network", In *Proc. of Passive and Active Measurement Conference*, Adelaide, Australia, Mar. 2006.

² Full details of all the participants are listed at <http://www.gridbus.org/cdn/>