Content Delivery Networks: How Big is Big Enough?*

Sipat Triukose, Zhihua Wen, and Michael Rabinovich
Case Western Reserve University
{sipat.triukose}{zhihua.wen}{michael.rabinovich}@case.edu

ABSTRACT

The central question addressed in this paper is whether a content delivery network (CDN) needs to deploy its servers in a large number of locations to achieve its current levels of performance. Our study indicates that a relatively small number of consolidated data centers might provide similar performance to end-users.

1. INTRODUCTION

Two main approaches to architect a CDN have emerged over the years, co-location and network-core approaches. The co-location approach, exemplified by Akamai, aims at creating presence at the edge of as many networks and network locations as possible. The network core approach, exemplified by AT&T and Limelight, uses large data centers near the main network backbones, and typically results in many fewer locations.

Two main "selling points" of a CDN service are that they supply on-demand capacity to content providers and improve performance of accessing the content from user perspective because they deliver the content from a nearby location. Consequently, CDNs often cite the size of their infrastructure in terms of the number of servers and data centers, implying that there is a direct link between the size and the performance benefits. This raises an important question: how many data centers is enough?

This paper attempts to answer the above question by examining Akamai performance. We chose Akamai because it is the dominant CDN provider, both in terms of the market share and size. Our general approach is to investigate how performance of Akamai-accelerated content delivery would suffer if it was done from fewer locations.

2. METHODOLOGY

We briefly describe the methodology we followed in our study.

2.1 Edge Server Discovery

Our study required measuring download performance from a large number of Akamai edge servers. To discover these servers, we perform DNS resolution of hostnames from a number of Akamai-accelerated URLs. We harvested these hostnames from the Web sites of 95 Akamai customers and utilized the DipZoom measurement platform [1] to resolve these hostnames from 472 vantage points world-wide repeatedly over a few months period. As the result, we discovered almost 12,000 Akamai edge servers, of which 10,231 servers, or over 30% of the total 34,000 servers claimed by Akamai during the study period, were pingable.

*This material is based upon work supported by the National Science Foundation under Grant No. CNS-0721890.

2.2 Overriding CDN's Edge Server Selection

As we will see shortly, our study involves comparing the download performance from Akamai-selected server and from an Akamai server of our choosing. We found that to trick an arbitrary edge server into processing our request, it is sufficient to simply include the HTTP host header that would have been submitted with a request using the proper DNS hostname. One can verify this technique by using curl - a command-line tool for HTTP downloads. For example, the following invocation will successfully download the object from Akamai edge server 206.132.122.75 by supplying the expected host header through the "-H" command argument:

```
curl -H Host:ak.buy.com \
     "http://206.132.122.75/db_assets
     /large_images/093/207502093.jpg"
```

2.3 Data Center Consolidation

Our goal is to compare the performance of the current Akamai platform with a hypothetical consolidated platform with fewer data centers. We group edge servers into consolidated data centers based on the estimated network distance between the servers as follows.

First, we estimate the pair-wise network distances between all the 10,231 pingable Akamai edge servers using a recently proposed dynamic triangles method [2]. Second, we then group nearby servers into a predefined number of big clusters using a hierarchical clustering (in the complete linkage mode) algorithm to the resulting distance matrix. We refer to the consolidated data centers as *big clusters*.

3. RESULTS

To judge the performance of the hypothetical consolidated Akamai platform, we study the performance implications of replacing the download from Akamai-selected server S with the download from the center of the big cluster to which S belongs. The center is represented by a randomly selected available server from the five closest servers to the center of the big cluster[1]. In other words, we assume that all clients that would have been sent to any server in a given big cluster in the existing platform, will be served from the cluster center in the consolidated case.

To avoid the bias from changing network conditions, we perform each pair of downloads in the immediate succes-

[1] We choose from five closest servers instead of always using the centroid node to reduce an undue effect of a single server on the entire experiment.

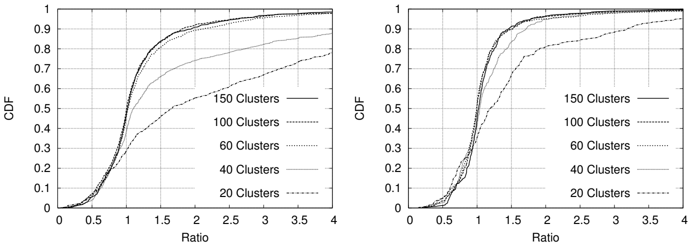

(a) All clients (b) Residential speed clients

Figure 1: The performance of a consolidated Akamai platform with different number of data centers.

sion. Further, we pre-request each object from either server to ensure the object is downloaded from the server cache in each case , thus excluding the possibility of cache-miss penalty. We compare the performance of the existing and consolidated Akamai platforms by the throughout ratio of the corresponding downloads.

Figure 1(a) presents the download throughput ratios of the existing-to-consolidated configurations. For each configuration, we obtained these ratios by downloading an outsourced Amazon[2] object of size 150K, 100K, 50K, and 10K from 412 measurement points world-wide. We were able to control precisely the target download size by selecting a large object and requesting it with range HTTP requests of the specified size.

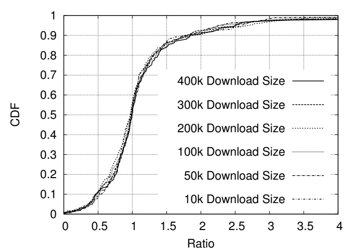

Figure 2: The performance of a consolidated Akamai platform with different target object sizes (60 Data Centers).

As seen in Figure 1(a), consolidating edge servers into 150, 100, and 60 data centers does not cause noticeable performance degradation. Only when we get down to 40 and 20 data centers, does the original platform start outperforming the consolidated configuration – 60% and 70% of the time for 40 and 20 data centers, respectively. Furthermore, as Figure 2 shows, these results are largely independent of the target object size.

The above experiment reflects our mix of measurement points, which are skewed towards well-connected hosts. To see the performance implication for users with different connectivity, we consider the performance of consolidated configurations separately for measurement points with different download bandwidth. Specifically, we group the measurement points by the maximum download throughput observed in the course of the experiment of Figure 1(a).

The results, shown in Figure 3, indicate that the existing Akamai configuration with large number of data centers favors well-connected users. For measurement points with over 6 Mbps bandwidth, the existing configuration outperformed consolidated configurations once they get down to 60 data centers. However, the less the bandwidth of the measurement points the less the advantage of the existing configuration. The question is how the consolidation may affect typical residential users.

[2]Amazon was Akamai customer during our study period.

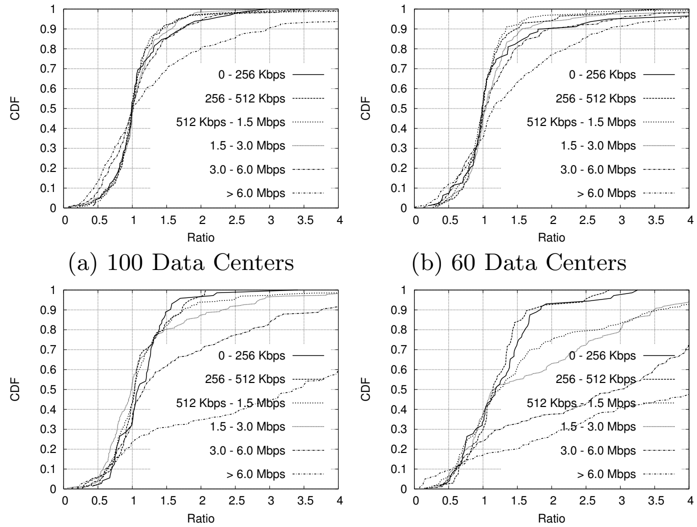

(a) 100 Data Centers (b) 60 Data Centers

(c) 40 Data Centers (d) 20 Data Centers

Figure 3: The performance of a consolidated Akamai platform with different download speeds.

Thus, we compare, in Figure 1(b), the performance of existing vs. consolidated configurations for all measurement points with bandwidth below 1.5 Mbps, which is typical DSL users. Figure 1(b) shows that these clients would not see noticeable performance difference if the servers were further consolidated into 40 data centers.

Overall, we conclude that one could consolidate Akamai platform to 60 data centers, and for typical residential users, even to 40 data centers without noticeable performance penalty.

Note that we do not claim that the network-core approach to content delivery is better or worse than co-location approach. A number of factors affect the choice between the two approaches, including non-technical factors such as business agreements between the CDN and the network operators concerned. Our study merely suggests that, in their considerations, the system designers can vary the number of data centers in a wide range, and in particular limit them to a relatively small number, without compromising the overall performance.

4. CONCLUSION

We present a large-scale performance study of the Akamai CDN platform. The main purpose was to evaluate the performance implications of consolidating a large number of points of presence maintained by Akamai into a smaller number of data centers. We found that quite significant consolidation is possible without appreciably degrading the platform performance.

5. REFERENCES

[1] DipZoom - Deep Internet Performance Zoom. http://dipzoom.case.edu.

[2] Z. Wen and M. Rabinovich. Network distance estimation with dynamic landmark triangles. In *ACM SIGMETRICS'08*, pages 433–434, 2008.