

IN THE UNITED STATES DISTRICT COURT
FOR THE SOUTHERN DISTRICT OF NEW YORK

VIACOM INTERNATIONAL INC., *et al.*,

Plaintiffs,

v.

YOUTUBE, INC., *et al.*,

Defendants.

Civil Action No. 07-cv-2103

THE FOOTBALL ASSOCIATION
PREMIER LEAGUE LIMITED, *et al.*,

Plaintiffs,

v.

YOUTUBE, INC., *et al.*,

Defendants.

Civil Action No. 07-cv-3502

DECLARATION OF MICHAEL GORDON

I, Michael Gordon, hereby declare as follows:

1. I am Chief Strategy Officer and Co-founder of Limelight Networks, Inc. (“Limelight”). I make this declaration based on my personal knowledge and belief.

2. Limelight provided Internet content delivery services to YouTube, Inc. (“YouTube”) from approximately 2005 to 2008, but YouTube does not currently use Limelight’s services.

3. Limelight has provided and continues to provide Internet content delivery services to customers that are owned or controlled by Viacom International Inc. (“Viacom Customers”).

4. For example, Limelight has provided Internet content delivery services to iFilm Corporation.

5. I am informed and believe that iFilm Corporation is owned or controlled by Viacom International Inc. (“Viacom”).

6. Limelight currently provides Internet content delivery services to MTV Networks.

7. I am informed and believe that MTV Networks is owned or controlled by Viacom.

8. In addition, Limelight has provided Internet content delivery services to Reality Digital, Inc.

9. Limelight’s US-based servers are located in its ten US points of presence. During a portion of the time YouTube subscribed to Limelight’s services, Limelight had fewer than ten US points of presence.

10. To my knowledge, no users of YouTube or of any Viacom Customer upload content directly to servers or computers owned or operated by Limelight.

11. To my knowledge, Limelight has no technical information about how users upload content to servers or computers operated by YouTube or by the Viacom Customers.

12. To access content from YouTube or from the Viacom Customers, Internet users typically request it by typing a Uniform Resource Locator (URL) into the address bar of an Internet browser, or by using a device or software program that similarly requests content identified by URLs. The user’s device or software (“the user device”) obtains the requested content through a series of steps, as follows:

13. Step 1: the URL is parsed to obtain an Internet hostname or host IP address (Internet URLs contain one or the other, usually explicitly, but sometimes through reference to a

base URL or domain). If the URL contains a host IP address Steps 2 through 7 are skipped and the user device proceeds to Step 8.

14. Step 2: the hostname is sent to a resolver, which is a software component that interacts with other components of the Internet Domain Name System (“DNS”).

15. Step 3: the resolver determines whether it already has on-hand a valid resolution of the hostname; if it does, it skips Steps 4 through 7 and the user device proceeds to Step 8. If not, it moves to Step 4.

16. Step 4: The resolver sends a hostname resolution request to the DNS proxy server it is configured to use.

17. Step 5: The DNS proxy server determines whether it already has on-hand a valid resolution of the hostname; if it does, it skips step 6 and proceeds to step 7. If it does not (and assuming for the sake of illustration that it does not have on-hand a valid resolution for any part of the hostname), then

18. Step 6: The DNS proxy server sends a hostname resolution request to an Internet DNS root name server. The root name server answers with a referral, normally in the form of one or more NS records, referring the DNS proxy server to a name server authoritative for the top-level domain contained in the hostname. The DNS proxy server then sends a hostname resolution request to the top-level domain name server, which similarly answers with a referral, also normally in the form of one or more NS records, referring the DNS proxy to a name server authoritative for the second-level domain contained in the hostname. The DNS proxy server then sends a hostname resolution request to the authoritative second-level name server, which typically answers either with one or more IP addresses (an A record) or answers with a CNAME

record. If the DNS proxy server receives a CNAME record, it restarts the resolution process from the beginning for the hostname contained in the CNAME record.

19. Step 7: Once the DNS proxy server has received a resolution containing one or more IP addresses, it returns the IP address(es) to the resolver.

20. Step 8: The user device then sends a request for the URL to the IP address (or one of the IP addresses, if it has received more than one from the DNS proxy server).

21. In order to use Limelight services for delivery of content, a publisher such as YouTube or a Viacom Customer uses a Limelight-assigned hostname, either by putting it in URLs it sends to users or by putting it in CNAME records it sends in response to requests for hostname resolutions. In either case, Limelight authoritative name servers then resolve the Limelight-assigned hostname to the IP addresses of either Limelight content servers or the IP addresses of Limelight switches (which will subsequently connect users to Limelight content servers). User devices then send content requests to Limelight content servers or switches, respectively.

22. Both YouTube and Viacom Customers used this process in order to utilize Limelight content delivery services.

23. When a Limelight content server receives a request for content, it typically determines the origin URL by deriving it from the requested URL.

24. If it is a first request, the Limelight server requests the content from the publisher via the Internet, using the origin URL.

25. If the Limelight server receives the content from the publisher, it delivers the content, and as a part of that process, it may cache the content.

26. If it does not receive the content from the publisher, the content server returns a message.

27. Under some circumstances, the Limelight content server may look in its cache prior to requesting the content to determine if the content has previously been cached.


28. If the content has been cached and has not expired, the content server delivers the cached content. Requests for content can be for a complete object, such as is normally the case for graphics and pictures, or can be for a part of an object (typically, for a specific byte range), such as is often the case for multimedia objects. Objects in the cache “age out” automatically according to a “Least Recently Used” algorithm, whereby, when space in cache is needed, the least recently used content will normally be overwritten first.

29. It is not possible for Limelight to determine the number of instances that a content object would be cached within its servers. The number of instances that an object would be cached would depend on many factors, such as, for example, the paths through which data travels, and the number of, location of, and timeframes during which user devices request the content.

30. The cache is located within servers that Limelight does not manufacture, but purchases from third parties. Limelight is not knowledgeable regarding the exact physical means by which data is electronically stored within these servers.

31. Both YouTube and Viacom Customers used this process in order to utilize Limelight content delivery services.

Dated: April 25, 2010



Michael Gordon