IN THE UNITED STATES DISTRICT COURT
FOR THE DISTRICT OF COLORADO

Civil Action No. 1:19-cv-3709

SOUND VIEW INNOVATIONS, LLC,
a Delaware corporation,

Plaintiff,

v.

JURY TRIAL DEMANDED

SLING TV L.L.C., a Colorado corporation,
and

Defendants.

COMPLAINT FOR PATENT INFRINGEMENT

Plaintiff Sound View Innovations, LLC ("Sound View"), for its Complaint for Patent Infringement against Sling TV L.L.C. ("Sling"), alleges as follows:

INTRODUCTION

1. Sound View is an intellectual property licensing company with a patent portfolio including approximately 300 active U.S. Patents. Those patents were developed by researchers at Alcatel Lucent ("Lucent") and its predecessors. Lucent was home to the world-renowned Bell Laboratories, which has a long and storied history of innovation. Researchers at Lucent's Bell Laboratories developed a wide variety of key innovations that have greatly enhanced the capabilities and utility of computer systems and networks. This has resulted in benefits such as better and more efficient computer networking, computer security, and user experiences.

2. Patents enjoy the same fundamental protections as real property. Sound View, like

any property owner, is entitled to insist that others respect its property and to demand compensation from those who take that property for their own use. Sling has used, and continues to use, Sound View's patents without authorization. Moreover, despite Sound View's attempt to negotiate, Sling refuses to take a license though it continues to use Sound View's property.

NATURE OF THE CASE

3. This action arises under 35 U.S.C. § 271 for Defendant's infringement of Sound View's United States Patent Nos. 6,708,213 (the "'213 patent"), 6,757,796 (the "'796 patent"), 6,725,456 (the "'456 patent") (collectively the "Patents-In-Suit").

THE PARTIES

4. Plaintiff Sound View is a Delaware limited liability company with its principal place of business at 2001 Route 46, Waterview Plaza, Suite 310, Parsippany, New Jersey 07054.

5. On information and belief, Defendant Sling TV L.L.C. is a limited liability company organized under the laws of the State of Colorado, with a principal place of business at 9601 S. Meridian Blvd., Englewood, Colorado 80112. Sling TV L.L.C. may be served through its registered agent for service, Timothy Allen Messner, 9601 S. Meridian Blvd., Englewood, Colorado 80112.Sling TV L.L.C., among other things, maintains the streaming website www.sling.com and offers the Sling TV streaming service.

JURISDICTION AND VENUE

6. This action arises under the patent laws of the United States, including 35 U.S.C. § 271 *et seq*. The jurisdiction of this Court over the subject matter of this action is proper under 28 U.S.C. §§ 1331 and 1338(a).

7. Venue is proper in this Court pursuant to 28 U.S.C. §§ 1391(b) and (c) and 1400(b)

as to Sling TV L.L.C., at least because Sling TV L.L.C. resides in this judicial district.

8. This Court has personal jurisdiction over Sling TV L.L.C., because it, among other things: is incorporated under the laws of Colorado, and has placed services that practice the claims of the Patents-in-Suit into the stream of commerce with the knowledge, or reasonable expectation, that actual or potential users of such services were located within this judicial district.

THE PATENTS-IN-SUIT

9. Sound View incorporates by reference the preceding paragraphs as if fully set forth herein.

A. The '213 Patent

10. The '213 patent, titled "Method for Streaming Multimedia Information Over Public Networks," was duly and properly issued by the USPTO on March 16, 2004. A copy of the '213 patent is attached hereto as Exhibit A.

11. Sound View is the owner and assignee of the '213 patent and holds the right to sue for and recover all damages for infringement thereof, including past infringement.

12. The '213 patent generally relates to streaming multimedia data (*e.g.*, audio and video data) over the Internet and other networks, and, more specifically, to methods to improve caching of streaming multimedia data from a content provider over a network to a client's computer.

13. At the time of the invention of the '213 patent, multimedia data could either be downloaded by the client or streamed over the network to the client. Streaming eliminated the need for the client to wait for the downloading to complete before watching or listening to the multimedia data. However, with conventional unicast connections, streaming posed problems: to

content providers in that server load increased linearly with the number of clients; to Internet service providers in that streaming caused network congestion problems; and to clients in that streaming often resulted in high start-up latency and unpredictable playback quality.

14. Conventional caching systems attempted to address network congestion, but these were unsuitable for streaming multimedia data: (1) video files were typically too large to be cached in their entirety, so only a few streams could be stored at a cache; (2) breaking video files into smaller pieces was not feasible because the caching systems would treat different chunks from the same video object independently; and (3) streaming multimedia has temporal characteristics, like the transmission rate, while conventional caching was only capable of handling static web objects.

15. The inventors of the ’213 patent solved those discrete computer-based problems and improved upon conventional caching techniques by providing a novel architecture and method for supporting high quality live and on-demand streaming multimedia on network systems using helper servers.

16. The techniques described in the ’213 patent advantageously reduce server and network loads by employing helper servers with dynamic data transfer rate control to overcome arrival time and range heterogeneity in client requests, thereby improving the quality perceived by end users making requests for streaming media objects.

17. The ’213 patent has been recognized with the 2013 Edison Patent Award in Multimedia Technology for inventing “fundamental concepts and techniques to design content distribution networks and caching systems originally built for text and images to better support streaming media over the Internet.” A press release regarding the award is attached as Exhibit B.

B. The '796 Patent

18. The '796 patent, titled "Method and System for Caching Streaming Live Broadcasts Transmitted Over a Network," was duly and properly issued by the USPTO on June 29, 2004. A copy of the '796 patent is attached hereto as Exhibit C.

19. Sound View is the owner and assignee of the '796 patent and holds the right to sue for and recover all damages for infringement thereof, including past infringement.

20. The '796 patent generally relates to real-time multimedia applications, and more specifically, to methods for decreasing the playback delay at a client computer of a live streaming broadcast transmitted over a network.

21. At the time of the invention of the '796 patent, live broadcasting of streaming multimedia over the Internet (including through movie broadcasts, television, sports, talk and music radio, business events, seminars, and tutorials) was becoming increasingly popular.

22. Streaming data involves sending a continuous transmission of data from the server to a client. At the client computer, received data is buffered in a cache memory and continuously processed as soon as, or soon after, being received by the client. The client computer creates a multimedia output from the received multimedia data. The advantage of streaming is that the client computer does not have to wait until all data is downloaded from the server before some of the data is processed and the multimedia output is created.

23. Because multimedia applications involve transferring large amounts of information, such systems place a considerable load on the resources of the network, server, and client. As more people accessed network-based multimedia applications, there was an increased demand for longer, more complicated, more flexible multimedia applications.

24. Multicast technology was developed for scaling live broadcasts. However, one problem that such technology did not address was that of start-up latency, *i.e.*, the delay between the client requesting multimedia playback and the beginning of the playback on the client.

25. The techniques described in the '796 patent solve that discrete computer-based problem and improve upon prior caching systems to better support the live broadcasting of streaming multimedia over the Internet and other network systems. In particular, the '796 patent provides novel methods for supporting high quality live streaming multimedia broadcasts on a network by using helper servers which operate as caching and streaming agents inside the network to enhance caching and reduce playback delay without sacrificing perceived playback quality. To allow the client's buffer to be filled faster (and thus allow playback to start faster), a playout history buffer is allocated and maintained at the helper server in response to a client request for a particular live streaming media broadcast. The playout history buffer operates as a moving window of fixed size that advances with the live broadcast stream, storing the last few seconds of the datastream. An advantage of utilizing playout history buffers is that as subsequent client requests are received at the helper server for a live streaming media broadcast which is currently being stored in a previously allocated playout history buffer in response to a former request, each subsequent request can be serviced directly from the playout history buffer thereby reducing start up latency. An advantage in streaming data packets to each client is realized by virtue of having some number of them pre-stored in the playout history buffer. When a request is received at the helper server, the stored packets are immediately available for distribution to the requesting client.

26. Servicing subsequent requests from the playout history buffer prevents the need to individually service each subsequent request from the content server as a unicast datastream, which

reduces network congestion by redirecting requests to the helper server. Also, the playout history buffer (which may be considered a form of short term dynamic cache) allows the cached data to be made immediately available to a requesting client to fill the client's playout buffer as rapidly as possible.

C. The '456 Patent

27. The '456 patent, titled "Methods and Apparatus for Ensuring Quality of Service in an Operating System," was duly and properly issued by the USPTO on April 20, 2004. A copy of the '456 patent is attached hereto as Exhibit D.

28. Sound View is the owner and assignee of the '456 patent and holds the right to sue for and recover all damages for infringement thereof, including past infringement.

29. The '456 patent generally relates to computer systems, and more particularly to techniques for providing a desired quality of service ("QoS") for an application running in a computer system.

30. At the time of the invention of the '456 patent, in a typical computer system multiple applications would contend for the same physical resources, such as a central processing unit, memory, and disk or network bandwidth. Conventional time-sharing operating systems could achieve acceptably low response time and high system throughput in some environments, but several trends made resource management techniques of conventional time-sharing operating systems increasingly inappropriate. For example, many workloads began including real-time applications like multimedia, which required that requests be processed within certain performance bounds. Also, a trend towards distributed client-server architectures increased the importance of fairness, i.e., preventing certain clients from monopolizing system resources.

31. The aforementioned trend towards client-server architectures made it necessary to manage resources hierarchically. For example, web servers and other user-level servers often needed mechanisms for processing client requests with specified QoS and/or fairness bounds. However, conventional time-sharing operating systems did not provide such mechanisms.

32. Then-existing proportional share schedulers did not provide satisfactory solutions to many problems that arose in their adoption in operating systems. For example, proportional share schedulers were proposed without an application programming interface ("API"), since they were not implemented and were evaluated only analytically or in simulations. As a further example, proportional share schedulers that were implemented used an API limited to a given scheduler and resource. As yet another example, proportional share schedulers that simply added resource reservations to conventional objects such as files or sockets did not provide correct sharing semantics, as such proportional share schedulers allowed those objects to be shared inappropriately by different users. As yet another example, proportional share schedulers did not propose how a parent process running on an operating system could limit the resource reservations used by its children processes. Finally, proportional share schedulers would hold resource reservations for processes that terminated abnormally, causing the reserved resource to become permanently unavailable.

33. The inventors of the '456 patent provided a technical solution for ensuring a desired QoS for an application running on an operating system.

34. Using the techniques for providing a desired QoS claimed by the '456 patent is particularly useful because it allows selected applications to isolate their performance and the performance of their corresponding client(s) from CPU, memory, disk, or network traffic

overloads caused by other applications. Such a capability is increasingly important for real-time, multimedia, Web, and distributed client-server applications as demands on network resources grow.

BACKGROUND FACTS

35. [redacted]

36. [redacted]

37. [redacted]

38. [redacted]

39. [redacted]

40. [redacted]

41. [redacted]

[redacted]

42. [redacted]

43. [redacted] Sling continues to knowingly, intentionally, and willfully infringe Sound View's patents so as to obtain significant benefits without paying any compensation to Sound View. Sound View thus has no other choice but to seek relief through litigation.

COUNT ONE
INFRINGEMENT OF THE '213 PATENT

44. Sound View incorporates by reference the preceding paragraphs as if fully set forth herein.

45. The '213 patent is valid and enforceable.

46. [redacted]

47. A content delivery network, also called a content distribution network (CDN), is a network of connected computers that delivers internet content, such as streaming video, to end users. When a service, such as Sling, uses a CDN, the content comes from an "origin server" and is replicated on numerous "edge servers." When an end user requests particular content, the CDN provides the content from an edge server near to the end user. This arrangement has numerous benefits, such as: faster response time (lower latency) because the content is served from a nearby

edge server, instead of a potentially distant origin server; greater throughput because the edge server will be less loaded than a single origin server would be; and greater availability because the multiplicity of servers allows for a request to be failed over to another server if an edge server crashes.

48. Sling provides and has provided streaming services, including at least Sling TV (the "Sling '213 Services"), to allow users to watch streaming video. Sling provides streaming video services to its users utilizing content delivery networks, including at least Akamai Technologies Inc. ("Akamai"), Fastly Inc. ("Fastly"), and Level 3 Communications, LLC ("Level 3") (collectively, "the CDNs"). The Sling '213 Services provide video that is encoded using certain protocols, including the HTTP Live Streaming ("HLS") protocol and the MPEG-DASH protocol.

49. HLS is an HTTP-based media streaming communications protocol. It works by breaking the overall stream into a sequence of small HTTP-based file downloads; each download is one short chunk that is part of an overall potentially unbounded transport stream. As the stream is played, the client may select from a number of different alternate chunks containing the same material encoded at a variety of data rates.

50. MPEG-DASH is an adaptive bitrate streaming technique that enables high quality streaming of media content over the Internet delivered from conventional HTTP web servers. Similar to HLS, MPEG-DASH works by breaking the content into a sequence of small HTTP-based file segments, each segment containing a short interval of playback time of content that is potentially many hours in duration, such as a live broadcast of a sports event. The content is made available at a variety of different bit rates, with alternative segments encoded at different bit rates covering aligned short intervals of playback time.

51. The CDNs each support Sling's delivery of video content to users using MPEG-DASH and/or HLS. Moreover, each of the CDNs openly advertises and promotes the use of those protocols to deliver video content to users.

52. Knowing that each of the CDNs supports the delivery of video content using MPEG-DASH and/or HLS, and directing and controlling such support, Sling delivers video streams to its users, including the Sling '213 Services, using at least the CDNs by transcoding videos into MPEG-DASH segments with different bit rates, and providing those segments to each of the CDNs. The CDNs store those MPEG-DASH segments in caches, and send them to Sling users who request to view the video files.

53. Sling contracts or has contracted with each of the CDNs, so that when at least certain Sling users request a video stream, the request is routed to one of the edge servers of the CDN, which receives the request. The edge server then allocates a local buffer to store portions of the stream.

54. Sling had and has the ability to configure and/or customize aspects of the operation of each of the CDNs in delivering content to its users. Moreover, on information and belief, Sling can and has configured and/or customized aspects of the operation of each of the CDNs in delivering content to its users. For example, Sling can customize the operation of the Akamai CDN through configuration tools, such as Akamai's Luna Control Center. As a further example, Sling can customize the operation of the Fastly CDN through configuration tools, such as the Fastly Control Panel. As a further example, Sling can customize the operation of the Level 3 CDN through configuration tools, such as Level 3 CDN Portal.

55. At least through contracting with Akamai and configuring and/or customizing

aspects of the operation of the Akamai CDN, Sling has knowledge of the operations of the Akamai CDN and the steps the Akamai systems will perform in order to deliver content to Sling's users. Sling thus knowingly causes and specifically intends for Akamai to perform those steps, or directs and controls Akamai's performance of these steps by means of at least its contractual relationship with Akamai and by configuring and customizing Akamai's CDN.

56. For example, utilizing Akamai's CDN requires storing segments in a local buffer on an edge server, and at least by entering into a contractual relationship with Akamai, Sling knowingly intends for Akamai to do so, or directs and controls Akamai (either implicitly or explicitly) to do so. Sling intends for, or directs, the Akamai edge server to request the MPEG-DASH or HLS segments from a datacenter cache, store them in the local buffer, and send them to Sling users who view the video. Further, Sling intends for, or directs, the edge server to store data in the buffer so that its end users can receive content with a lower latency.

57. While the Akamai edge server sends the requested segments to the user, it concurrently requests the next few segments in the stream from the datacenter cache or from the cache of another server. By doing so, the content can be streamed smoothly without pauses for buffering. Akamai advertises this process as "pre-fetching." Sling intends for and contracts with Akamai to use pre-fetching so that its users can receive content without pauses for buffering. Sling and other customers have the ability to configure the size of the segments to be fetched in the Akamai system. The Akamai CDN, as configured and customized by Sling, also allows Sling users to receive content without pauses for buffering by allowing end users to send byte range requests to the edge server.

58. While the content is being played back by an MPEG-DASH or HLS client, the

client automatically selects the next segment to download and play based on current network conditions. The streaming server then provides the requested alternate segment, resulting in the server adjusting the data rate. Sling intends for and controls the Akamai CDN to adjust the data rate by directing, controlling, and/or inducing Akamai to provide the content on its CDN at different data rates.

59. As a further example, at least through contracting with Fastly and configuring and/or customizing aspects of the operation of the Fastly CDN, Sling has knowledge of the operations of the Fastly CDN and the steps the Fastly systems will perform in order to deliver content to Sling's users. Sling thus knowingly causes and specifically intends for Fastly to perform those steps, or directs and controls Fastly's performance of those steps by means of at least its contractual relationship with Fastly and by configuring and customizing Fastly's CDN.

60. For instance, utilizing Fastly's CDN requires storing segments in a local buffer on an edge server, and at least by entering into a contractual relationship with Fastly, Sling knowingly intends for Fastly to do so, or directs and controls Fastly (either implicitly or explicitly) to do so. Sling intends for, or directs, the Fastly edge server to request the MPEG-DASH or HLS segments from a datacenter cache, store them in the local buffer, and send them to Sling users who view the video. Further, Sling intends for, or directs, the edge server to store data in the buffer so that its end users can receive content with a lower latency.

61. While the Fastly edge server sends the requested segments to the user, it concurrently requests the next few segments in the stream from the datacenter cache or from the cache of another server. By doing so, the content can be streamed smoothly without pauses for buffering. Sling intends for and contracts with (or has contracted with) Fastly to deliver content

in this manner so that its users can receive content without pauses for buffering. Sling and other customers have the ability to configure the size of the segments to be fetched in the Fastly system. The Fastly CDN, as configured and customized by Sling, also allows Sling users to receive content without pauses for buffering by allowing end users to send byte range requests to the edge server.

62. While the content is being played back by an MPEG-DASH or HLS client, the client automatically selects from the alternatives the next segment to download and play based on current network conditions. The streaming server then provides the requested alternate segment, resulting in the server adjusting the data rate. Sling intends for and controls the Fastly CDN to adjust the data rate by directing, controlling, and/or inducing Fastly to provide the content on its CDN at different data rates.

63. As a further example, at least through contracting with Level 3 and configuring and/or customizing aspects of the operation of the Level 3 CDN, Sling has knowledge of the operations of the Level 3 CDN and the steps the Level 3 systems will perform in order to deliver content to Sling's users. Sling thus knowingly causes and specifically intends for Level 3 to perform those steps, or directs and controls Level 3's performance of those steps by means of at least its contractual relationship with Level 3 and by configuring and customizing Level 3's CDN.

64. For instance, utilizing Level 3's CDN requires storing segments in a local buffer on an edge server, and at least by entering into a contractual relationship with Level 3, Sling knowingly intends for Level 3 to do so, or directs and controls Level 3 (either implicitly or explicitly) to do so. Sling intends for, or directs, the Level 3 edge server to request the MPEG-DASH or HLS segments from a datacenter cache, store them in the local buffer, and send them to Sling users who view the video. Further, Sling intends for, or directs, the edge server to store data

in the buffer so that its end users can receive content with a lower latency.

65. While the Level 3 edge server sends the requested segments to the user, it concurrently requests the next few segments in the stream from the datacenter cache or from the cache of another server. By doing so, the content can be streamed smoothly without pauses for buffering. Sling intends for and contracts with Level 3 to deliver content in this manner so that its users can receive content without pauses for buffering. Sling and other customers have the ability to configure the size of the segments to be fetched in the Level 3 system. The Level 3 CDN, as configured and customized by Sling, also allows Sling users to receive content without pauses for buffering by allowing end users to send byte range requests to the edge server.

66. While the content is being played back by an MPEG-DASH or HLS client, the client automatically selects from the alternatives the next segment to download and play based on current network conditions. The streaming server then provides the requested alternate segment, resulting in the server adjusting the data rate. Sling intends for and controls the Level 3 CDN to adjust the data rate by directing, controlling, and/or inducing Level 3 to provide the content on its CDN at different data rates.

67. Sling directly infringes one or more claims of the ’213 patent (including at least claim 16) under 35 U.S.C. § 271(a), literally and/or under the doctrine of equivalents, at least by directing and/or controlling at least the performance of the claimed steps by the CDNs to infringe the ’213 patent to deliver the Sling ’213 Services.

68. For example, Sling has directly infringed, and continues to directly infringe, claim 16 of the ’213 patent under 35 U.S.C. § 271(a), literally and/or under the doctrine of equivalents, at least by directing and/or controlling Akamai to deliver the Sling ’213 Services. For example,

Sling has directly infringed, and continues to directly infringe, claim 16 of the '213 patent under 35 U.S.C. § 271(a) literally and/or under the doctrine of equivalents, at least by directing and/or controlling Akamai (through at least contracting with Akamai and customizing the Akamai CDN) to infringe claim 16 by using a method of reducing latency in a network having a content server which hosts streaming media ("SM") objects (such as videos) which comprise a plurality of time-ordered segments (such as HLS or MPEG-DASH segments) for distribution over said network through a plurality of helpers ("HSs") (such as Akamai cache or edge servers) to a plurality of clients (such as users of the Sling '213 Services). Further:

a. Sling directs and/or controls Akamai, at least via its contract with Akamai and/or its configuration and customization of Akamai's CDN, to receive a request for an SM object from one of said plurality of clients (such as a user of one of the Sling '213 Services requesting to watch a hosted video) at one of said plurality of helper servers (such as by directing and/or controlling one of the Akamai cache or edge servers to receive such a request from a user of one of the Sling '213 Services to watch a hosted video);

b. Sling directs and/or controls Akamai, at least via its contract with Akamai and/or its configuration and customization of Akamai's CDN, to allocate a buffer at one of said plurality of HSs to cache at least a portion of said requested SM object (such as by directing and/or controlling Akamai to allocate a local buffer to store portions of the stream as HLS or MPEG-DASH segments at the Akamai cache or edge servers);

c. Sling directs and/or controls Akamai, at least via its contract with Akamai and/or its configuration and customization of Akamai's CDN, to download said portion of said requested SM object to said requesting client, while concurrently retrieving a remaining portion of

said requested SM object from one of another HS and said content server (such as by directing and/or controlling the Akamai cache or edge server to pre-fetch the next segment of video content by requesting the next HLS or MPEG-DASH segments in the stream from the datacenter cache, and/or by directing and/or controlling the Akamai cache or edge server to be capable of receiving a byte range request in order to download a segment of a requested video stream to a client while concurrently downloading the next segments from another server); and

d. Sling directs and/or controls Akamai, at least via its contract with Akamai and/or its configuration and customization of Akamai's CDN and/or its provision of content encoded at multiple bitrates, to adjust a data transfer rate at said one of said plurality of HSs for transferring data from said one of said plurality of helper servers to said one of said plurality of clients (such as by directing and/or controlling Akamai to provide alternate segments encoded at different data rates to the client to accommodate the current network conditions (*e.g.*, the client's current bandwidth), such that providing the requested alternate segment results in an adjusted data rate).

69. As a further example, Sling has directly infringed, and continues to directly infringe, claim 16 of the '213 patent under 35 U.S.C. § 271(a), literally and/or under the doctrine of equivalents, at least by directing and/or controlling Fastly to deliver the Sling '213 Services. For example, Sling has directly infringed, and continues to directly infringe, claim 16 of the '213 patent under 35 U.S.C. § 271(a) literally and/or under the doctrine of equivalents, at least by directing and/or controlling Fastly (through at least contracting with Fastly and customizing the Fastly CDN) to infringe claim 16 by using a method of reducing latency in a network having a content server which hosts SM objects (such as videos) which comprise a plurality of time-ordered

segments (such as HLS or MPEG-DASH segments) for distribution over said network through a plurality of HSs (such as Fastly cache or edge servers) to a plurality of clients (such as users of the Sling '213 Services). Further:

a. Sling directs and/or controls Fastly, at least via its contract with Fastly and/or its configuration and customization of Fastly's CDN, to receive a request for an SM object from one of said plurality of clients (such as a user of one of the Sling '213 Services requesting to watch a hosted video) at one of said plurality of helper servers (such as by directing and/or controlling one of the Fastly cache or edge servers to receive such a request from a user of one of the Sling '213 Services to watch a hosted video);

b. Sling directs and/or controls Fastly, at least via its contract with Fastly and/or its configuration and customization of Fastly's CDN, to allocate a buffer at one of said plurality of HSs to cache at least a portion of said requested SM object (such as by directing and/or controlling Fastly to allocate a local buffer to store portions of the stream as HLS or MPEG-DASH segments at the Fastly cache or edge servers);

c. Sling directs and/or controls Fastly, at least via its contract with Fastly and/or its configuration and customization of Fastly's CDN, to download said portion of said requested SM object to said requesting client, while concurrently retrieving a remaining portion of said requested SM object from one of another HS and said content server (such as by directing and/or controlling the Fastly cache or edge server to pre-fetch the next segment of video content by requesting the next HLS or MPEG-DASH segments in the stream from the datacenter cache, and/or by directing and/or controlling the Fastly cache or edge server to be capable of receiving a byte range request in order to download a segment of a requested video stream to a client while

concurrently downloading the next segments from another server); and

d. Sling directs and/or controls Fastly, at least via its contract with Fastly and/or its configuration and customization of Fastly's CDN and/or its provision of content encoded at multiple bitrates, to adjust a data transfer rate at said one of said plurality of HSs for transferring data from said one of said plurality of helper servers to said one of said plurality of clients (such as by directing and/or controlling Fastly to provide alternate segments encoded at different data rates to the client to accommodate the current network conditions (e.g., the client's current bandwidth), such that providing the requested alternate segment results in an adjusted data rate).

70. As a further example, Sling has directly infringed, and continues to directly infringe, claim 16 of the '213 patent under 35 U.S.C. § 271(a), literally and/or under the doctrine of equivalents, at least by directing and/or controlling Level 3 to deliver the Sling '213 Services. For example, Sling has directly infringed, and continues to directly infringe, claim 16 of the '213 patent under 35 U.S.C. § 271(a) literally and/or under the doctrine of equivalents, at least by directing and/or controlling Level 3 (through at least contracting with Level 3 and customizing the Level 3 CDN) to infringe claim 16 by using a method of reducing latency in a network having a content server which hosts SM objects (such as videos) which comprise a plurality of time-ordered segments (such as HLS or MPEG-DASH segments) for distribution over said network through a plurality of HSs (such as Level 3 cache or edge servers) to a plurality of clients (such as users of the Sling '213 Services). Further:

a. Sling directs and/or controls Level 3, at least via its contract with Level 3 and/or its configuration and customization of Level 3's CDN, to receive a request for an SM object from one of said plurality of clients (such as a user of one of the Sling '213 Services requesting to

watch a hosted video) at one of said plurality of helper servers (such as by directing and/or controlling one of the Level 3 cache or edge servers to receive such a request from a user of one of the Sling '213 Services to watch a hosted video);

b. Sling directs and/or controls Level 3, at least via its contract with Level 3 and/or its configuration and customization of Level 3's CDN, to allocate a buffer at one of said plurality of HSs to cache at least a portion of said requested SM object (such as by directing and/or controlling Level 3 to allocate a local buffer to store portions of the stream as HLS or MPEG-DASH segments at the Level 3 cache or edge servers);

c. Sling directs and/or controls Level 3, at least via its contract with Level 3 and/or its configuration and customization of Level 3's CDN, to download said portion of said requested SM object to said requesting client, while concurrently retrieving a remaining portion of said requested SM object from one of another HS and said content server (such as by directing and/or controlling the Level 3 cache or edge server to pre-fetch the next segment of video content by requesting the next HLS or MPEG-DASH segments in the stream from the datacenter cache, and/or by directing and/or controlling the Level 3 cache or edge server to be capable of receiving a byte range request in order to download a segment of a requested video stream to a client while concurrently downloading the next segments from another server); and

d. Sling directs and/or controls Level 3, at least via its contract with Level 3 and/or its configuration and customization of Level 3's CDN and/or its provision of content encoded at multiple bitrates, to adjust a data transfer rate at said one of said plurality of HSs for transferring data from said one of said plurality of helper servers to said one of said plurality of clients (such as by directing and/or controlling Level 3 to provide alternate segments encoded at

different data rates to the client to accommodate the current network conditions (e.g., the client's current bandwidth), such that providing the requested alternate segment results in an adjusted data rate).

71. In addition or in the alternative, Sling has induced infringement, and continues to induce infringement, of one or more claims of the '213 patent under 35 U.S.C. § 271(b), literally and/or under the doctrine of equivalents. Sling has actively, knowingly, and intentionally induced (and continues to induce) infringement of the '213 patent by making, using, offering for sale, selling, supplying, maintaining, and/or supporting the Sling '213 Services; by contracting with the CDNs and customizing the CDNs with the specific intent to cause the CDNs to perform the steps claimed in the '213 patent to deliver video data, including the Sling '213 Services, to Sling's users, and with the knowledge that such actions infringe the '213 patent.

72. For example, at least through repeated correspondence from Sound View, Sling knows that at least Akamai, Fastly, and Level 3 perform the claimed methods of the '213 patent to deliver the Sling '213 Services, and that Sling induces the infringement of each of those CDNs. (*See* Exhibit E, incorporated herein by reference.) Moreover, Sling specifically intends that infringement, at least by continuing to contract with and utilize the Akamai, Fastly, and Level 3 CDNs to offer the Sling '213 Services; configuring the Akamai, Fastly, and Level 3 CDNs to perform the claimed methods of the '213 patent; and by encouraging and facilitating its infringement through the use of the Sling '213 Services by Sling's users and/or the creation and dissemination of documentation related to the Sling '213 Services, including by, for example, encouraging and instructing its agents and contractors, such as Akamai, Fastly, and Level 3, to provide video to Sling's users through the Sling '213 Services, causing the performance of the

claimed methods with the knowledge that such actions infringe the ’213 patent.

73. For example, Sling intends for and induces Akamai to infringe claim 16 to deliver the Sling ’213 Services by using a method of reducing latency in a network having a content server which hosts SM objects (such as videos) which comprise a plurality of time-ordered segments (such as HLS or MPEG-DASH segments) for distribution over said network through a plurality of HSs (such as Akamai cache or edge servers) to a plurality of clients (such as users of the Sling ’213 Services). Sling further intends for and induces Akamai to:

a. receive a request for an SM object from one of said plurality of clients (such as a user of one of the Sling ’213 Services requesting to watch a hosted video) at one of said plurality of helper servers (such as one of the Akamai cache or edge servers, with knowledge that Akamai’s cache or edge servers will receive such a request from a user of one of the Sling ’213 Services to watch a hosted video);

b. allocate a buffer at one of said plurality of HSs to cache at least a portion of said requested SM object (such as by inducing Akamai to allocate a local buffer to store portions of the stream as HLS or MPEG-DASH segments at the Akamai cache or edge servers, with knowledge that Akamai’s CDN will allocate such a buffer at one of the Akamai cache or edge servers to store portions of the stream as HLS or MPEG-DASH segments);

c. download said portion of said requested SM object to said requesting client, while concurrently retrieving a remaining portion of said requested SM object from one of another HS and said content server (such as the Akamai cache or edge server pre-fetching the next segment of video content by requesting the next HLS or MPEG-DASH segments in the stream from the datacenter cache, with knowledge that Akamai’s cache or edge servers will pre-fetch the next

segment of video by requesting the next HLS or MPEG-DASH segment in the stream from the datacenter cache, and/or by directing and/or controlling the Akamai cache or edge server to be capable of receiving a byte range request in order to download a segment of a requested video stream to a client while concurrently downloading the next segments from another server); and

d. adjust a data transfer rate at said one of said plurality of HSs for transferring data from said one of said plurality of helper servers to said one of said plurality of clients (such as providing alternate segments encoded at different data rates to the client to accommodate the current network conditions (*e.g.*, the client's current bandwidth), and then providing the requested alternate segment resulting in an adjusted data rate, with knowledge that the Akamai CDN will provide alternate segments encoded at different data rates to the client).

74. As a further example, Sling intends for and induces Fastly to infringe claim 16 to deliver the Sling '213 Services by using a method of reducing latency in a network having a content server which hosts SM objects (such as videos) which comprise a plurality of time-ordered segments (such as HLS or MPEG-DASH segments) for distribution over said network through a plurality of HSs (such as Fastly cache or edge servers) to a plurality of clients (such as users of the Sling '213 Services). Sling further intends for and induces Fastly to:

a. receive a request for an SM object from one of said plurality of clients (such as a user of one of the Sling '213 Services requesting to watch a hosted video) at one of said plurality of helper servers (such as one of the Fastly cache or edge servers, with knowledge that Fastly's cache or edge servers will receive such a request from a user of one of the Sling '213 Services to watch a hosted video);

b. allocate a buffer at one of said plurality of HSs to cache at least a portion of

said requested SM object (such as by inducing Fastly to allocate a local buffer to store portions of the stream as HLS or MPEG-DASH segments at the Fastly cache or edge servers, with knowledge that Fastly's CDN will allocate such a buffer at one of the Fastly cache or edge servers to store portions of the stream as HLS or MPEG-DASH segments);

c. download said portion of said requested SM object to said requesting client, while concurrently retrieving a remaining portion of said requested SM object from one of another HS and said content server (such as the Fastly cache or edge server pre-fetching the next segment of video content by requesting the next HLS or MPEG-DASH segments in the stream from the datacenter cache, with knowledge that Fastly's cache or edge servers will pre-fetch the next segment of video by requesting the next HLS or MPEG-DASH segment in the stream from the datacenter cache, and/or by directing and/or controlling the Fastly cache or edge server to be capable of receiving a byte range request in order to download a segment of a requested video stream to a client while concurrently downloading the next segments from another server); and

d. adjust a data transfer rate at said one of said plurality of HSs for transferring data from said one of said plurality of helper servers to said one of said plurality of clients (such as providing alternate segments encoded at different data rates to the client to accommodate the current network conditions (e.g., the client's current bandwidth), and then providing the requested alternate segment resulting in an adjusted data rate, with knowledge that the Fastly CDN will provide alternate segments encoded at different data rates to the client).

75. As a further example, Sling intends for and induces Level 3 to infringe claim 16 to deliver the Sling '213 Services by using a method of reducing latency in a network having a content server which hosts SM objects (such as videos) which comprise a plurality of time-ordered

segments (such as HLS or MPEG-DASH segments) for distribution over said network through a plurality of HSs (such as Level 3 cache or edge servers) to a plurality of clients (such as users of the Sling '213 Services). Sling further intends for and induces Level 3 to:

a. receive a request for an SM object from one of said plurality of clients (such as a user of one of the Sling '213 Services requesting to watch a hosted video) at one of said plurality of helper servers (such as one of the Level 3 cache or edge servers, with knowledge that Level 3's cache or edge servers will receive such a request from a user of one of the Sling '213 Services to watch a hosted video);

b. allocate a buffer at one of said plurality of HSs to cache at least a portion of said requested SM object (such as by inducing Level 3 to allocate a local buffer to store portions of the stream as HLS or MPEG-DASH segments at the Level 3 cache or edge servers, with knowledge that Level 3's CDN will allocate such a buffer at one of the Level 3 cache or edge servers to store portions of the stream as HLS or MPEG-DASH segments);

c. download said portion of said requested SM object to said requesting client, while concurrently retrieving a remaining portion of said requested SM object from one of another HS and said content server (such as the Level 3 cache or edge server pre-fetching the next segment of video content by requesting the next HLS or MPEG-DASH segments in the stream from the datacenter cache, with knowledge that Level 3's cache or edge servers will pre-fetch the next segment of video by requesting the next HLS or MPEG-DASH segment in the stream from the datacenter cache, and/or by directing and/or controlling the Level 3 cache or edge server to be capable of receiving a byte range request in order to download a segment of a requested video stream to a client while concurrently downloading the next segments from another server); and

d. adjust a data transfer rate at said one of said plurality of HSs for transferring data from said one of said plurality of helper servers to said one of said plurality of clients (such as providing alternate segments encoded at different data rates to the client to accommodate the current network conditions (e.g., the client's current bandwidth), and then providing the requested alternate segment resulting in an adjusted data rate, with knowledge that the Level 3 CDN will provide alternate segments encoded at different data rates to the client).

76. Sound View has been and continues to be damaged by Sling's infringement of the '213 patent and is entitled to recover from Sling the damages sustained by Sound View as a result of Sling's wrongful acts in an amount adequate to compensate Sound View for Sling's infringement subject to proof at trial.

77. [redacted] Sling nonetheless committed—and continues to commit—acts of direct and indirect infringement despite knowing that its actions constituted infringement of the valid and enforceable '213 patent, despite a risk of infringement that was known or so obvious that it should have been known to Sling, and/or even though Sling otherwise knew or should have known that its actions constituted an unjustifiably high risk of infringement of that valid and enforceable patent. Sling's conduct in light of these circumstances is egregious.

78. Sling's infringement of the '213 patent was and is knowing, deliberate, and willful, entitling Sound View to increased damages under 35 U.S.C. § 284 and to attorney fees and costs incurred in prosecuting this action under 35 U.S.C. § 285.

COUNT TWO
INFRINGEMENT OF THE '796 PATENT

79. Sound View incorporates by reference the preceding paragraphs as if fully set forth herein.

80. The '796 patent is valid and enforceable.

81. [redacted]

82. Sling provides and has provided live streaming services, including at least Sling's streaming service Sling TV (the "Sling '796 Services"), to allow users to watch live streaming video.

83. The CDNs, including Akamai, Fastly, and Level 3 each support Sling's delivery of video content to users using MPEG-DASH and/or HLS. Moreover, each of the CDNs openly advertises and promotes the use of those protocols to deliver video content to users. Knowing that each of the CDNs supports the delivery of video content using MPEG-DASH and/or HLS, and directing or controlling such support, Sling delivers the Sling '796 Services to its users using at least the Akamai, Fastly, and Level 3 CDNs by transcoding videos into MPEG-DASH and/or HLS segments.

84. Sling contracts or has contracted with each of the CDNs, so that when at least certain Sling users request the Sling '796 Services video stream, the request is routed to one of the edge servers of the CDN, which receives the request. Moreover, on information and belief, Sling

can and has configured and/or customized aspects of the operation of each of the CDNs in delivering content to its users. For example, Sling can customize the operation of the Akamai CDN through configuration tools, such as Akamai's Luna Control Center. As a further example, Sling can customize the operation of the Fastly CDN through configuration tools, such as the Fastly Control Panel. As a further example, Sling can customize the operation of the Level 3 CDN through configuration tools, such as Level 3 CDN Portal.

85. For example, at least through contracting with Akamai and configuring and/or customizing aspects of the operation of the Akamai CDN, Sling has knowledge of the operations of the Akamai CDN and the steps the Akamai systems will perform in order to deliver content to Sling's users. Sling thus knowingly causes and specifically intends for Akamai to perform those steps, or directs and controls Akamai's performance of these steps by means of at least its contractual relationship with Akamai and by configuring and customizing Akamai's CDN.

86. For example, Sling contracts with Akamai knowing that when at least certain Sling users request the Sling '796 Services live stream, the request is routed to an Akamai edge server, which receives the request, and that the Akamai edge server allocates a local buffer to store portions of the stream. Sling contracts with Akamai also knowing that when a second user requests the same video stream, the Akamai edge server will provide the stream from the same local buffer, because Akamai's edge servers serve the second request from the same local buffer because doing so saves space and bandwidth. Sling's contract with Akamai thus implicitly or explicitly directs and controls Akamai to serve a second request for the same stream from the same local buffer. Because the Akamai edge server already has the requested stream in a local buffer, it takes less time to send it to the second user.

87. As a further example, at least through contracting with Fastly and configuring and/or customizing aspects of the operation of the Fastly CDN, Sling has knowledge of the operations of the Fastly CDN and the steps the Fastly systems will perform in order to deliver content to Sling's users. Sling thus knowingly causes and specifically intends for Fastly to perform those steps, or directs and controls Fastly's performance of those steps by means of at least its contractual relationship with Fastly and by configuring and customizing Fastly's CDN.

88. For instance, Sling contracts or has contracted with Fastly knowing that when at least certain Sling users request the Sling '796 Services live stream, the request is routed to a Fastly edge server, which receives the request, and that the Fastly edge server allocates a local buffer to store portions of the stream. Sling contracts with Fastly also knowing that when a second user requests the same video stream, the Fastly edge server will provide the stream from the same local buffer, because Fastly's edge servers serve the second request from the same local buffer because doing so saves space and bandwidth. Sling's contract with Fastly thus implicitly or explicitly directs and controls Fastly to serve a second request for the same stream from the same local buffer. Because the Fastly edge server already has the requested stream in a local buffer, it takes less time to send it to the second user.

89. As a further example, at least through contracting with Level 3 and configuring and/or customizing aspects of the operation of the Level 3 CDN, Sling has knowledge of the operations of the Level 3 CDN and the steps the Level 3 systems will perform in order to deliver content to Sling's users. Sling thus knowingly causes and specifically intends for Level 3 to perform those steps, or directs and controls Level 3's performance of these steps by means of at least its contractual relationship with Level 3 and by configuring and customizing Level 3's CDN.

90. For example, Sling contracts with Level 3 knowing that when at least certain Sling users request the Sling '796 Services live stream, the request is routed to a Level 3 edge server, which receives the request, and that the Level 3 edge server allocates a local buffer to store portions of the stream. Sling contracts with Level 3 also knowing that when a second user requests the same video stream, the Level 3 edge server will provide the stream from the same local buffer, because Level 3's edge servers serve the second request from the same local buffer because doing so saves space and bandwidth. Sling's contract with Level 3 thus implicitly or explicitly directs and controls Level 3 to serve a second request for the same stream from the same local buffer. Because the Level 3 edge server already has the requested stream in a local buffer, it takes less time to send it to the second user.

91. Sling directly infringes one or more claims of the '796 patent (including at least claim 27) under 35 U.S.C. § 271(a), literally and/or under the doctrine of equivalents, at least by directing and/or controlling at least the performance of the claimed steps by Akamai, Fastly, and Level 3 to infringe the '796 patent to deliver the Sling '796 Services.

92. For example, Sling has directly infringed, and continues to directly infringe, claim 27 of the '796 patent under 35 U.S.C. § 271(a), literally and/or under the doctrine of equivalents, at least by directing and/or controlling Akamai (through at least contracting with Akamai and customizing the Akamai CDN) to infringe claim 27 by using, in a network having a content server (such as a web content server) which hosts a plurality of live SM broadcast objects (such as live video) for distribution over said network through a plurality of HSs (such as Akamai's edge servers) to a plurality of clients (such as Sling's users), a method of reducing start-up latency associated with distributing said plurality of live SM broadcast objects from said content server

and said plurality of HSs to said plurality of clients. Further:

a. Sling directs and/or controls Akamai, at least via its contract with Akamai and/or its configuration and customization of Akamai's CDN, to receive a first request for one of said plurality of live SM broadcast objects at one of said plurality of HSs (such as by directing and/or controlling Akamai to receive a first request from a Sling user to watch a live video at one of Akamai's edge servers);

b. Sling directs and/or controls Akamai, at least via its contract with Akamai and/or its configuration and customization of Akamai's CDN, to service said first request from a non pre-configured playout history ("PH") buffer (such as by directing and/or controlling Akamai to contact a content server, retrieve and cache the requested MPEG-DASH or HLS segments at the Akamai edge server in a local buffer, and deliver the requested content to the client) at a first data rate;

c. Sling directs and/or controls Akamai, at least via its contract with Akamai and/or its configuration and customization of Akamai's CDN, to receive a second request for said one of said plurality of live SM broadcast objects at said one of said plurality of HSs (such as by directing and/or controlling Akamai to receive a second request for the same MPEG-DASH or HLS segments at the Akamai edge server); and

d. Sling directs and/or controls Akamai, at least via its contract with Akamai and/or its configuration and customization of Akamai's CDN, to partially service said second request from said non pre-configured PH buffer (such as by directing and/or controlling Akamai to deliver the requested MPEG-DASH or HLS segments to the client from the same local buffer

on the Akamai edge server) at a second data rate, wherein said second data rate is higher than said first data rate.

93. As a further example, Sling has directly infringed, and continues to directly infringe, claim 27 of the '796 patent under 35 U.S.C. § 271(a), literally and/or under the doctrine of equivalents, at least by directing and/or controlling Fastly (through at least contracting with Fastly and customizing the Fastly CDN) to infringe claim 27 by using, in a network having a content server (such as a web content server) which hosts a plurality of live SM broadcast objects (such as live video) for distribution over said network through a plurality of HSs (such as Fastly's edge servers) to a plurality of clients (such as Sling's users), a method of reducing start-up latency associated with distributing said plurality of live SM broadcast objects from said content server and said plurality of HSs to said plurality of clients. Further:

a. Sling directs and/or controls Fastly, at least via its contract with Fastly and/or its configuration and customization of Fastly's CDN, to receive a first request for one of said plurality of live SM broadcast objects at one of said plurality of HSs (such as by directing and/or controlling Fastly to receive a first request from a Sling user to watch a live video at one of Fastly's edge servers);

b. Sling directs and/or controls Fastly, at least via its contract with Fastly and/or its configuration and customization of Fastly's CDN, to service said first request from a non pre-configured PH buffer (such as by directing and/or controlling Fastly to contact a content server, retrieve and cache the requested MPEG-DASH or HLS segments at the Fastly edge server in a local buffer, and deliver the requested content to the client) at a first data rate;

c. Sling directs and/or controls Fastly, at least via its contract with Fastly and/or its configuration and customization of Fastly's CDN, to receive a second request for said one of said plurality of live SM broadcast objects at said one of said plurality of HSs (such as by directing and/or controlling Fastly to receive a second request for the same MPEG-DASH or HLS segments at the Fastly edge server); and

d. Sling directs and/or controls Fastly, at least via its contract with Fastly and/or its configuration and customization of Fastly's CDN, to partially service said second request from said non pre-configured PH buffer (such as by directing and/or controlling Fastly to deliver the requested MPEG-DASH or HLS segments to the client from the same local buffer on the Fastly edge server) at a second data rate, wherein said second data rate is higher than said first data rate.

94. As a further example, Sling has directly infringed, and continues to directly infringe, claim 27 of the '796 patent under 35 U.S.C. § 271(a), literally and/or under the doctrine of equivalents, at least by directing and/or controlling Level 3 (through at least contracting with Level 3 and customizing the Level 3 CDN) to infringe claim 27 by using, in a network having a content server (such as a web content server) which hosts a plurality of live SM broadcast objects (such as live video) for distribution over said network through a plurality of HSs (such as Level 3's edge servers) to a plurality of clients (such as Sling's users), a method of reducing start-up latency associated with distributing said plurality of live SM broadcast objects from said content server and said plurality of HSs to said plurality of clients. Further:

a. Sling directs and/or controls Level 3, at least via its contract with Level 3 and/or its configuration and customization of Level 3's CDN, to receive a first request for one of said plurality of live SM broadcast objects at one of said plurality of HSs (such as by directing

and/or controlling Level 3 to receive a first request from a Sling user to watch a live video at one of Level 3's edge servers);

b. Sling directs and/or controls Level 3, at least via its contract with Level 3 and/or its configuration and customization of Level 3's CDN, to service said first request from a non pre-configured PH buffer (such as by directing and/or controlling Level 3 to contact a content server, retrieve and cache the requested MPEG-DASH or HLS segments at the Level 3 edge server in a local buffer, and deliver the requested content to the client) at a first data rate;

c. Sling directs and/or controls Level 3, at least via its contract with Level 3 and/or its configuration and customization of Level 3's CDN, to receive a second request for said one of said plurality of live SM broadcast objects at said one of said plurality of HSs (such as by directing and/or controlling Level 3 to receive a second request for the same MPEG-DASH or HLS segments at the Level 3 edge server); and

d. Sling directs and/or controls Level 3, at least via its contract with Level 3 and/or its configuration and customization of Level 3's CDN, to partially service said second request from said non pre-configured PH buffer (such as by directing and/or controlling Level 3 to deliver the requested MPEG-DASH or HLS segments to the client from the same local buffer on the Level 3 edge server) at a second data rate, wherein said second data rate is higher than said first data rate.

95. In addition or in the alternative, Sling has induced infringement, and continues to induce infringement, of one or more claims of the '796 patent under 35 U.S.C. § 271(b), literally and/or under the doctrine of equivalents. Sling has actively, knowingly, and intentionally induced (and continues to induce) infringement of the '796 patent by making, using, offering for sale,

selling, supplying, maintaining, and/or supporting the Sling '796 Services; by contracting with the CDNs and customizing the CDNs with the specific intent to cause the CDNs to perform the steps claimed in the '796 patent to deliver video data, including the Sling '796 Services, to Sling's users, and with the knowledge that such actions infringe the '796 patent.

96. For example, at least through repeated correspondence from Sound View, Sling knows that at least Akamai, Fastly, and Level 3 perform the claimed methods of the '796 patent, and that Sling induces the infringement of each of those CDNs. (*See* Exhibit E, incorporated herein by reference.) Moreover, Sling specifically intends that infringement, at least by continuing to contract with and utilize the Akamai, Fastly, and Level 3 CDNs to offer the Sling '796 Services; configuring or customizing the Akamai, Fastly, and Level 3 CDNs to perform the claimed methods of the '796 patent; and by encouraging and facilitating its infringement through the use of the Sling '796 Services by Sling's users and/or the creation and dissemination of documentation related to the Sling '796 Services, including by, for example, encouraging and instructing its agents and contractors, such as Akamai, Fastly, and Level 3 to provide video to Sling's users through the Sling '796 Services, causing the performance of the claimed methods with the knowledge that such actions infringe the '796 patent

97. For example, Sling intends for and induces Akamai to infringe claim 27 to deliver the Sling '796 Services by using, in a network having a content server (such as a web content server) which hosts a plurality of live SM broadcast objects (such as live video) for distribution over said network through a plurality of HSs (such as Akamai's edge servers) to a plurality of clients (such as Sling's users), a method of reducing start-up latency associated with distributing said plurality of live SM broadcast objects from said content server and said plurality of HSs to

said plurality of clients, said method comprising:

a. receiving a first request for one of said plurality of live SM broadcast objects (such as a Sling user requesting to watch a live video) at one of said plurality of HSs (such as the Akamai edge servers);

b. servicing said first request from a non pre-configured PH buffer (such as by contacting a content server, retrieving and caching the requested MPEG-DASH or HLS segments at the Akamai edge server in a local buffer, and delivering the requested content to the client) at a first data rate;

c. receiving a second request for said one of said plurality of live SM broadcast objects at said one of said plurality of HSs (such as receiving a second request for the same MPEG-DASH or HLS segments at the Akamai edge server); and

d. partially servicing said second request from said non pre-configured PH buffer (such as by delivering the requested MPEG-DASH or HLS segments to the client from the same local buffer on the Akamai edge server) at a second data rate, wherein said second data rate is higher than said first data rate.

98. As a further example, Sling intends for and induces Fastly to infringe claim 27 to deliver the Sling '796 Services by using, in a network having a content server (such as a web content server) which hosts a plurality of live SM broadcast objects (such as live video) for distribution over said network through a plurality of HSs (such as Fastly's edge servers) to a plurality of clients (such as Sling's users), a method of reducing start-up latency associated with distributing said plurality of live SM broadcast objects from said content server and said plurality of HSs to said plurality of clients, said method comprising:

a. receiving a first request for one of said plurality of live SM broadcast objects (such as a Sling user requesting to watch a live video) at one of said plurality of HSs (such as the Fastly edge servers);

b. servicing said first request from a non pre-configured PH buffer (such as by contacting a content server, retrieving and caching the requested MPEG-DASH or HLS segments at the Fastly edge server in a local buffer, and delivering the requested content to the client) at a first data rate;

c. receiving a second request for said one of said plurality of live SM broadcast objects at said one of said plurality of HSs (such as receiving a second request for the same MPEG-DASH or HLS segments at the Fastly edge server); and

d. partially servicing said second request from said non pre-configured PH buffer (such as by delivering the requested MPEG-DASH or HLS segments to the client from the same local buffer on the Fastly edge server) at a second data rate, wherein said second data rate is higher than said first data rate.

99. As a further example, Sling intends for and induces Level 3 to infringe claim 27 to deliver the Sling '796 Services by using, in a network having a content server (such as a web content server) which hosts a plurality of live SM broadcast objects (such as live video) for distribution over said network through a plurality of HSs (such as Level 3's edge servers) to a plurality of clients (such as Sling's users), a method of reducing start-up latency associated with distributing said plurality of live SM broadcast objects from said content server and said plurality of HSs to said plurality of clients, said method comprising:

a. receiving a first request for one of said plurality of live SM broadcast objects (such as a Sling user requesting to watch a live video) at one of said plurality of HSs (such as the Level 3 edge servers);

b. servicing said first request from a non pre-configured PH buffer (such as by contacting a content server, retrieving and caching the requested MPEG-DASH or HLS segments at the Level 3 edge server in a local buffer, and delivering the requested content to the client) at a first data rate;

c. receiving a second request for said one of said plurality of live SM broadcast objects at said one of said plurality of HSs (such as receiving a second request for the same MPEG-DASH or HLS segments at the Level 3 edge server); and

d. partially servicing said second request from said non pre-configured PH buffer (such as by delivering the requested MPEG-DASH or HLS segments to the client from the same local buffer on the Level 3 edge server) at a second data rate, wherein said second data rate is higher than said first data rate.

100. Sound View has been and continues to be damaged by Sling's infringement of the '796 patent and is entitled to recover from Sling the damages sustained by Sound View as a result of Sling's wrongful acts in an amount adequate to compensate Sound View for Sling's infringement subject to proof at trial.

101. [redacted]

[redacted] Sling nonetheless committed—and continues to commit—acts of direct and indirect infringement despite knowing that its actions constituted infringement of the valid and enforceable '796 patent, despite a risk of infringement that was known or so obvious that it should have been known to Sling, and/or even though Sling otherwise knew or should have known that its actions constituted an unjustifiably high risk of infringement of that valid and enforceable patent. Sling's conduct in light of these circumstances is egregious.

102. Sling's infringement of the '796 patent was and is knowing, deliberate, and willful, entitling Sound View to increased damages under 35 U.S.C. § 284 and to attorney fees and costs incurred in prosecuting this action under 35 U.S.C. § 285.

COUNT THREE
INFRINGEMENT OF THE '456 PATENT

103. Sound View incorporates by reference the preceding paragraphs as if fully set forth herein.

104. The '456 patent is valid and enforceable.

105. Sling has used software known as Apache Hadoop YARN ("Yarn") in its data systems. For example, Sling's use of Yarn has been openly advertised by Sling and its employees, and includes, without limitation, management of cluster resources via Yarn (the "Sling Yarn Services").

106. Yarn is the architectural center of Hadoop that allows multiple data processing engines such as interactive SQL, real-time streaming, data science and batch processing to handle data stored in a single platform. Yarn provides resource management and a central platform to

deliver consistent operations, security, and data governance tools across Hadoop clusters.

107. Yarn splits up the functionalities of resource management and job scheduling into separate daemons, by having a global ResourceManager ("RM") and per-application ApplicationMaster ("AM"). The RM is the ultimate authority that arbitrates resources among all the applications in the system. The per-application AM is, in effect, a framework specific library and is tasked with negotiating resources from the RM and working with the NodeManager(s) to execute and monitor the tasks. Yarn provides the ability to preempt certain applications in order to make room for other more time-sensitive or higher priority applications.

108. Within Yarn, the fundamental unit of scheduling is a queue. The capacity of each queue specifies the percentage of cluster resources that are available for applications submitted to the queue. Yarn uses a hierarchy of queues wherein each leaf (child) queue is tied to a single parent queue. Parent queues contain more parent queues or leaf queues but do not themselves accept any application submissions directly. Child queues live under a parent queue and accept applications.

109. A user may launch an application on Yarn using the YarnClient, ApplicationSubmissionContext, Cluster Applications API, and/or Resource APIs. New clients define all the information needed by the RM to launch the AM, which includes the application ID, name, queue, and priority information. ApplicationSubmissionContext can be used to, among other things, define the container in which the AM will be launched and run. It defines all required information needed to run the application, including resources and environmental settings. ApplicationSubmissionContext includes resource requirements such as memory and vCores. Moreover, helper APIs convert values obtained from the environment into objects.

110. Additionally, Yarn's Cluster Reservation Submit API can be used to submit reservations. When the reservation is made, the user can use the reservation ID used to submit the reservation to get access to the resources by specifying it as part of a Cluster Submit Applications API. The Cluster Submit Applications Object includes a resource object, which includes memory and vCore requirements for each container.

111. Yarn's RM includes a Fair Scheduler and Capacity Scheduler, which allow assigning guaranteed minimum shares to queues. When an API submits a reservation, it is validated by the RM, which returns a reservation ID and creates reservable queues. RM's schedulers then provide containers, giving a user guaranteed access to the required resources, as identified by objects, in accordance with capacity and fairness sharing protocols.

112. Yarn's Fair Scheduler and Capacity Scheduler guarantee minimum resource reservations, e.g., memory and/or vCores, to queues. If a queue's minimum share is not satisfied, it will be offered available resources before any other queue under the same parent. Fair Scheduler uses hierarchical queues, such that queues are sibling queues when they have the same parent. Associated with each queue is a weight, which determines the amount of resources a queue deserves in relation to its sibling queues. This amount is known as Steady FairShare, which is calculated at the queue level. For the root queue, the Steady FairShare is equal to all the cluster's resources. The Steady FairShare is calculated such that the minimum amount of resources associated with the parent queue is at least equal to the sum of the minimum resources associated with each of the parent's children.

113. Sling has infringed one or more claims of the '456 patent under 35 U.S.C. § 271(a), literally and/or under the doctrine of equivalents, by making, using, selling, and/or offering for

sale in the United States, and/or importing into the United States, products and/or methods encompassed by those claims, including for example, by making, using, selling, offering for sale, and/or importing systems and platforms that include or use the Sling Yarn Services.

114. For example, Sling infringed at least claim 13 by using a method of ensuring a particular quality of service for an application in a computer system, the method comprising the steps of:

a. utilizing an application programming interface of an operating system to establish one or more quality of service guarantees that correspond to a reference to an object (such as the YarnClient, ApplicationSubmissionContext, Cluster Applications API, Resource, and/or Cluster Reservation Submit APIs) ; and

b. providing a particular quality of service to a request in accordance with the one or more quality of service guarantees that correspond to one or more object references used in the request (such as through use of Yarn's Fair Scheduler and/or Capacity Scheduler); and

c. wherein the quality of service guarantees comprise resource reservations, each specifying a portion of a resource set aside for exclusive use by one or more processes (such as memory and vCores); and

d. wherein the resource reservations are organized hierarchically such that each resource reservation r may have at most one parent and one or more siblings and children, and associated with r is a weight that specifies how r shares the resources of r's parent with r's siblings (such as the hierarchical queues used by Yarn's Fair Scheduler and Capacity Scheduler); and

e. wherein associated with each resource reservation r is a minimum amount of resources that r receives from its parent p, such that the minimum amount of resources associated with p is at least equal to the sum of the minimum amount of resources associated with each of p's children (such as the Steady FairShare of resources).

115. Sound View has been damaged by Sling's infringement of the '456 patent and is entitled to recover from Sling the damages sustained by Sound View as a result of Sling's wrongful acts in an amount adequate to compensate Sound View for Sling's infringement subject to proof at trial.

RELIEF REQUESTED

Wherefore, Sound View respectfully requests that this Court enter judgment against Sling as follows:

a) that Sling has infringed each of the Patents-in-Suit;

b) that Sling's infringement of the '213 and '796 patents is and/or has been willful;

c) that Sound View be awarded damages in accordance with 35 U.S.C. § 284, including treble damages and, if necessary to adequately compensate Sound View for Sling's infringement, an accounting;

d) that this case is exceptional under 35 U.S.C. § 285;

e) that Sound View be awarded the attorney's fees, costs, and expenses that it incurs in prosecuting this action; and

f) that Sound View be awarded further relief at law or in equity as the Court deems just and proper.

DEMAND FOR JURY TRIAL

Sound View demands a trial by jury on all claims and issues so triable.

Dated: December 30, 2019

By: */s/ Kathryn Bi*

DESMARAIS LLP
Alan S. Kellman
Edward Geist
Kathryn Bi
230 Park Avenue
New York, NY 10169
Tel: (212) 351-3400
Fax: (212) 351-3401
akellman@desmaraisllp.com
egeist@desmaraisllp.com
kbi@desmaraisllp.com

Counsel for Plaintiff Sound View Innovations, LLC