

1 Bruce S. Sostek, admitted *pro hac vice*
 bruce.sostek@tklaw.com
 2 Richard L. Wynne, Jr., admitted *pro hac vice*
 richard.wynne@tklaw.com
 3 Adrienne E. Dominguez, admitted *pro hac vice*
 adrienne.dominguez@tklaw.com
 THOMPSON & KNIGHT LLP
 4 One Arts Plaza
 1722 Routh Street, Suite 1500
 Dallas, Texas 75201
 5 Telephone: (214) 969-1700
 Facsimile: (214) 969-1751
 6

7 John V. Picone III, Bar No. 187226
 jpicone@hopkinscarley.com
 Christopher A. Hohn, Bar No. 271759
 8 chohn@hopkinscarley.com
 HOPKINS & CARLEY
 A Law Corporation
 9 The Letitia Building
 70 South First Street
 10 San Jose, CA 95113-2406

11 **mailing address:**
 P.O. Box 1469
 San Jose, CA 95109-1469
 12 Telephone: (408) 286-9800
 Facsimile: (408) 998-4790
 13

14 Attorneys for Plaintiffs
 BROADCOM CORPORATION and AVAGO
 TECHNOLOGIES INTERNATIONAL SALES
 PTE. LIMITED.
 15

16 UNITED STATES DISTRICT COURT
 17 CENTRAL DISTRICT OF CALIFORNIA
 SOUTHERN DIVISION

18 BROADCOM CORPORATION and
 19 AVAGO TECHNOLOGIES
 INTERNATIONAL SALES PTE.
 20 LIMITED,

21 Plaintiffs,

22 v.

23 NETFLIX, INC.,

24 Defendant.

Case No. 8:20-cv-00529-JVS-ADS

**FIRST AMENDED COMPLAINT
 FOR PATENT INFRINGEMENT**

DEMAND FOR JURY TRIAL

1 Plaintiffs Broadcom Corporation (“Broadcom Corp.”) and Avago
2 Technologies International Sales Pte. Limited (“Avago”) (collectively, the
3 “Broadcom Entities”) file this First Amended Complaint for Patent Infringement
4 against Defendant Netflix, Inc. (“Netflix”) and allege as follows:

5 **NATURE OF THIS ACTION**

6 1. This complaint alleges patent infringement. The Broadcom Entities
7 allege that Netflix has infringed and continues to infringe, directly and/or indirectly,
8 twelve patents: U.S. Patent Nos. 7,266,079 (the “’079 Patent”); 8,259,121 (the
9 “’121 Patent”); 8,959,245 (the “’245 Patent”); 8,270,992 (the “’992 Patent”);
10 6,341,375 (the “’375 Patent”); 8,572,138 (the “’138 Patent”); 6,744,387 (the “’387
11 Patent”); 6,982,663 (the “’663 Patent”); 9,332,283 (the “’283 Patent”), 8,548,976
12 (the “’976 Patent”); 7,457,722 (the “’722 Patent”); and 8,365,183 (the “’183
13 Patent”). Copies of these patents (collectively, the “Patents-in-Suit”) are attached
14 hereto as **Exhibits A-L**.

15 2. The Patents-in-Suit cover foundational technologies that are essential
16 to various aspects of Netflix’s video streaming service, and the systems that Netflix
17 uses to support this service.

18 3. Netflix directly infringes the Patents-in-Suit by making, using, offering
19 to sell, selling, and/or importing into the United States internet video streaming
20 technology, software, and services that practice the inventions claimed in the
21 Patents-in-Suit. Netflix directs and controls each relevant aspect of the accused
22 technology discussed herein, and benefits from the use of each feature that infringes
23 the Patents-in-Suit.

24 4. Netflix indirectly infringes the Patents-in-Suit by inducing its

1 consumer end-users to directly infringe these patents. For example, Netflix induces
2 infringement by providing software (e.g., the Netflix application) that, when used
3 by consumers or other content viewers to stream digital content to televisions,
4 personal computers, phones, tablets, video game consoles, and other devices, as
5 directed and intended by Netflix, causes those end-users to use and practice the
6 inventions claimed in the Patents-in-Suit.

7 5. The Broadcom Entities seek damages and other relief for Netflix's
8 infringement of the Patents-in-Suit.

9 **PARTIES**

10 6. Plaintiff Broadcom Corporation is a California corporation
11 headquartered at 1320 Ridder Park Drive, San Jose, California 95131. Broadcom
12 Corp. maintains offices within the Central District of California at 15101 Alton
13 Parkway, Irvine, California 92618. Broadcom Corp. is an indirect subsidiary of
14 Broadcom, Inc.

15 7. Plaintiff Avago Technologies International Sales Pte. Ltd. is a
16 corporation formed under the laws of Singapore with places of business at 1320
17 Ridder Park Dr., San Jose, California 95131 and 1 Yishun Avenue 7, Singapore
18 768923. Avago is also an indirect subsidiary of Broadcom, Inc.

19 8. Defendant Netflix, Inc. is a Delaware corporation that maintains its
20 principal place of business and global headquarters at 100 Winchester Circle, Los
21 Gatos, 95032.

22 9. Netflix maintains regular and established places of business in this
23 District, including an office at 5808 Sunset Blvd., Los Angeles, CA 90028, where
24 Netflix employs hundreds of people. According to Netflix's website, the Los

1 Angeles office “is the entertainment hub for Netflix with teams such as Content,
2 Legal, Marketing & Publicity and is located on the Sunset Bronson Studio Lot
3 where a variety of Netflix content is created.”¹

4 10. Netflix may be served through its registered agent for service of
5 process in California: CT Corporation System, 818 W. Seventh St, Suite 930, Los
6 Angeles, CA 90017.

7 11. Netflix claims to be a global leader in streaming digital video content.
8 Netflix streams videos of various types, such as films and television series, to over
9 180 million paid members in over 190 countries. Upon information and belief,
10 Netflix designs, operates, tests, manufactures, uses, offers for sale, sells, and/or
11 imports into the United States—including in the Central District of California—
12 internet video streaming software, systems, and services that generate billions of
13 dollars of revenue for Netflix each year.

14 **JURISDICTION AND VENUE**

15 12. The Broadcom Entities bring this civil action for patent infringement
16 under the Patent Laws of the United States, 35 U.S.C. § 1 et. seq., including 35
17 U.S.C. §§ 271, 281-285. This Court has subject matter jurisdiction over this action
18 pursuant to 28 U.S.C. §§ 1331 and 1338.

19 13. The Broadcom Entities’ claims for relief arise, at least in part, from
20 Netflix’s business contacts and other activities in the State of California and in this
21 District. Upon information and belief, Netflix has committed acts of infringement
22 within this District and the State of California by making, using, selling, offering
23 for sale, and/or importing into the United States and this District products, systems,

24 ¹ <https://jobs.netflix.com/locations/los-angeles-california>.

1 and services that infringe one or more claims of the Patents-in-Suit as set forth
2 herein. Further, Netflix induces others within this District to infringe one or more
3 claims of the Patents-in-Suit.

4 14. Venue is proper in this district and division under 28 U.S.C. §§
5 1391(b)-(d) and 1400(b) because Netflix has committed acts of infringement in the
6 Central District of California and has a regular and established physical place of
7 business in Los Angeles, part of the Central District. Upon information and belief,
8 Netflix employs engineers and technical professionals of many disciplines at its Los
9 Angeles facility.

10 **FACTUAL BACKGROUND**

11 15. Henry Samueli and Henry Nicholas founded Broadcom in 1991 in Los
12 Angeles, California. Since then, Broadcom has grown to be a global technology
13 company that produces category-leading semiconductor and infrastructure software
14 solutions. Among other things, Broadcom provides one of the industry's broadest
15 portfolios of highly integrated semiconductor chips that seamlessly deliver voice,
16 video, data, and multimedia connectivity in the home, office, and mobile
17 environments. From its headquarters in San Jose, California, Broadcom has
18 expanded its footprint across the United States and around the world, employing
19 thousands of individuals globally and in the United States. An overview of
20 Broadcom's history can be found on its website at:

21 <https://www.broadcom.com/company/about-us/company-history/>.

22 16. Broadcom's continued success depends in substantial part upon its
23 constant attention to research and development. Broadcom and its subsidiaries
24 spend billions of dollars on research and development for their products each year.

1 Because of this focus, Broadcom has produced a wide range of novel technologies
2 and inventions that are directed to advancements in, among other things,
3 semiconductor design and digital communications, digital content distribution,
4 enterprise and data center networking, home connectivity, set top boxes,
5 infrastructure software, and other technologies integral to business and consumer
6 settings across the United States and throughout the world.

7 17. Broadcom relies on the patent system as an important part of its
8 intellectual property program to protect the valuable technology and inventions
9 resulting from its research and development efforts. The Broadcom Entities and
10 their related entities have tens of thousands of patents in the United States and
11 abroad.

12 18. In addition to their internally developed inventions and associated
13 intellectual property, the Broadcom group of companies have acquired technology
14 and intellectual property through mergers and acquisition with other major
15 technology companies, such as the Avago family of companies, LSI, Brocade, CA,
16 Inc. (formerly known as Computer Associates International, Inc.), and Symantec's
17 enterprise business.

18 19. As explained in detail below, Netflix has built its familiar video
19 streaming business, in part, on the Broadcom Entities' patented technology. Netflix
20 relies on this technology for crucial aspects of the Netflix streaming service. This
21 includes, for example, the Netflix systems used to ensure effective and reliable
22 delivery of streaming content with minimal interruptions, to ensure the efficient,
23 effective use of Netflix server resources, and to encode Netflix streaming content in
24 a format compatible with a large percentage of the client devices (e.g., computers,

1 smart TVs, mobile phones, and videogame consoles) used to access the Netflix
2 service.

3 20. In doing so, Netflix has caused, and continues to cause, substantial and
4 irreparable harm to the Broadcom Entities. For instance, the Broadcom Entities sell
5 semiconductor chips used in the set top boxes that enable traditional cable
6 television services. Upon information and belief, as a direct result of the on-
7 demand streaming services provided by Netflix, the market for traditional cable
8 services that require set top boxes has declined, and continues to decline, thereby
9 substantially reducing Broadcom's set top box business.

10 21. For instance, it is widely reported that the rise of on-demand video
11 streaming services such as Netflix has concurrently lead to a decrease in demand
12 for traditional cable services. As an example, *Variety* reported in February 2019
13 that “[t]he five biggest U.S. pay-television providers saw their traditional subscriber
14 rolls shrink 4.2% in 2018, as they collectively lost around 3.2 million customers for
15 the year. That’s an acceleration from estimated sector-wide declines of 3.7% in
16 2017 and 2% in 2016.” The article attributes the loss in part to a migration of
17 customers to Netflix.²

18 22. Upon information and belief, Netflix could not displace traditional
19 cable television services, or could not do so as effectively, without the use of the
20 Broadcom Entities’ patented technology, which—as explained above—enable
21 critical aspects of Netflix’s systems.

22
23
24 ² <https://variety.com/2019/biz/news/cord-cutting-2018-accelerate-us-pay-tv-subscribers-1203138404/>.

1 27. The '079 Patent claims priority from U.S. Provisional Application No.
2 60/233,338, filed on September 18, 2000.

3 28. The '079 Patent was assigned to Avago, which currently holds all
4 substantial rights, title, and interest in and to the '079 Patent.

5 29. Pursuant to 35 U.S.C. § 282, the '079 Patent is presumed valid.

6 30. The '079 Patent is directed to an improvement in the functionality of
7 networked computer systems by “balancing data flow there through.”⁴ Specifically,
8 the '079 Patent’s claims describe a new approach for balancing transmission unit
9 traffic over the multiple heterogeneous links that often connect computing
10 platforms in a computer network.

11 31. The '079 Patent addresses a specific technical problem that arose in
12 the computer networking environment as the networks grew ever larger and more
13 complex, and as users sought to transmit ever greater volumes of data across these
14 networks. As the '079 Patent states, “[t]he present invention relates to
15 communications apparatus and methods, particularly to computer networking
16 apparatus and methods, and more particularly to computer networking apparatus
17 and methods for balancing data flow there through.”⁵

18 32. As the '079 Patent explains, “[a] common problem in communication
19 networks is maintaining efficient utilization of network resources, particularly with
20 regard to bandwidth, so that data traffic is efficiently distributed over the available
21 links between sources and destinations.”⁶ “Prior art solutions include apparatus and
22 methods that balance data traffic over homogeneous (same-speed) links between

23 _____
⁴ '079 Patent, 1:17-21.

24 ⁵ *Id.*

⁶ *Id.* at 1:23-27.

1 heterogeneous or homogeneous computing platforms (servers, clients, etc.).”⁷
2 However, “[i]ncreasingly, high-performance computing platforms communicate
3 with other computers, routers, switches, and the like, using multiple links which,
4 for a variety of reasons, may operate at disparate link speeds.”⁸ “For
5 example...adverse network conditions may degrade the performance of one or
6 more links, effectively presenting a heterogeneous-link-speed environment to the
7 server and its link partner(s).”⁹

8 33. Accordingly, a need existed for a means to “dynamically balance
9 transmission unit traffic in a heterogeneous-link-speed environment” in order to
10 improve the functionality of computer networks.¹⁰

11 34. The ’079 Patent claims specific, novel ways to solve these technical
12 problems by dynamically balancing data traffic in a computer networking
13 environment with heterogeneous link speeds. The claims of the ’079 Patent are
14 directed to new, improved methods and apparatuses for balancing transmission unit
15 traffic over networks links.

16 35. The methods and apparatuses described in the ’079 Patent improve the
17 functionality of a networked computer system by balancing the data traffic among
18 network links having different speeds, capabilities, and congestion levels that
19 connect the various networked elements, thereby improving the speed and
20 efficiency of data transmission within the network.

21 36. Claim 1 of the ’079 Patent reads as follows:

22 _____
23 ⁷ *Id.* at 1:27-30.

24 ⁸ *Id.* at 1:30-34.

⁹ *Id.* at 1:34-40.

¹⁰ *Id.* at 1:40-43.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24

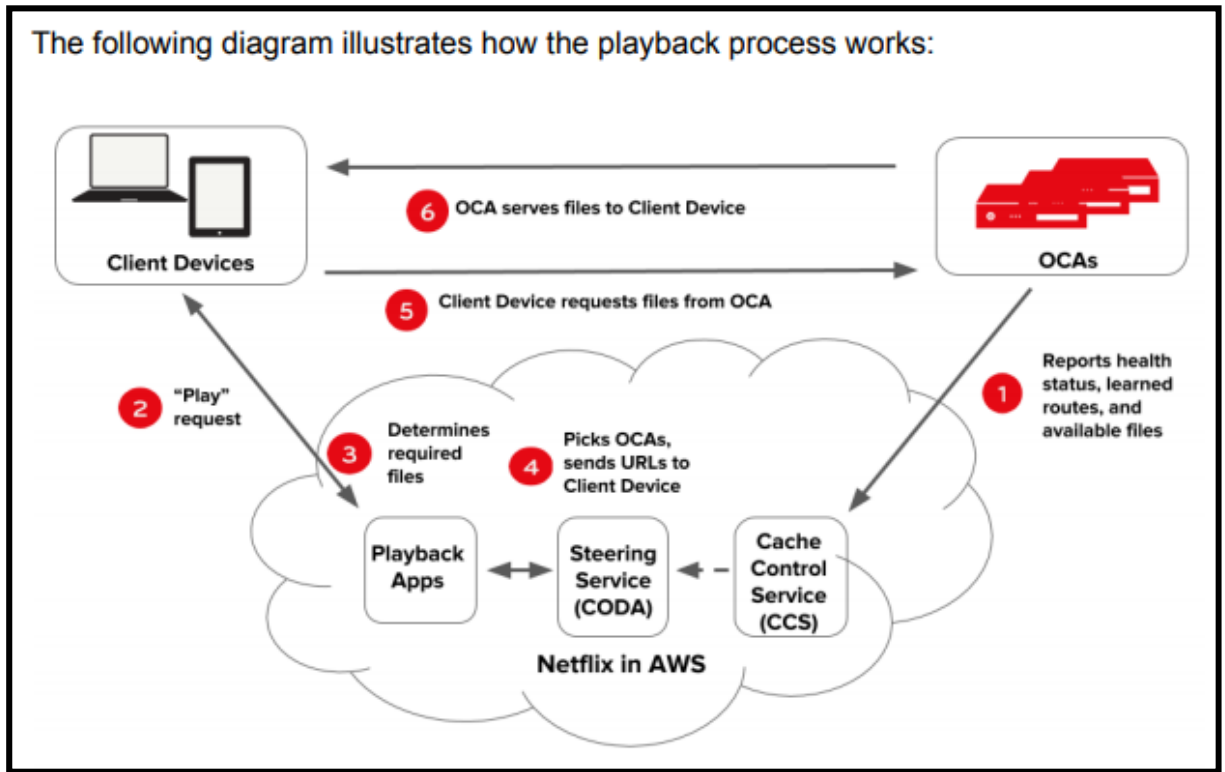
A method for balancing transmission unit traffic over network links, comprising:

- a. disposing transmission units into flows;
- b. grouping flows into first flow lists, each of the first flow lists corresponding to a selected network link;
- c. determining a traffic metric representative of a traffic load on the selected network link;
- d. responsive to the traffic metric, regrouping flows into second flow lists corresponding to the selected network link, the regrouping balancing the transmission unit traffic among the network links; and
- e. transmitting the respective second flow list over the respective selected network link.

37. Netflix directly infringes the '079 Patent by making, using, offering to sell, and/or selling in the United States its Netflix service, which utilizes the inventions claimed in the '079 Patent to balance traffic over Netflix's systems, including its content delivery network ("CDN").

38. Netflix directly infringes at least independent claim 1 of the '079 Patent at least in the exemplary manner described below.

39. Netflix utilizes the claimed "method for balancing transmission unit traffic over network links," including, for instance, in operating its CDN, which Netflix uses to stream video content to its subscribers over the internet. The Netflix CDN is illustrated in the following diagram.



Source: <https://openconnect.netflix.com/Open-Connect-Overview.pdf>.

40. The Netflix CDN comprises hardware that Netflix builds and operates. Additionally, for some aspects of its CDN, Netflix uses hardware produced and maintained by third parties, including cloud computing services purchased from third parties. However, all the infringing technologies discussed in this Complaint are developed and controlled by Netflix, including though the use of Netflix software that controls the relevant functions performed by the underlying hardware.

41. Notably, in apparent recognition that its cloud computing services could be used by a customer to infringe any number of patents involving a computer as part of a larger system or process, the AWS standard customer agreement makes clear that AWS has no liability for patent infringement claims

1 arising from infringement by combinations of AWS’s services with any other
2 product, service, software, data, content or method.¹¹

3 42. As part of its CDN, Netflix created, operates, uses, and maintains a
4 “global network of thousands of [Open Connect Appliances],” also known as
5 OCAs.¹²

6 43. As Netflix explains, “[t]he building blocks of Open Connect are our
7 suite of purpose-built server appliances, called Open Connect Appliances (OCAs).
8 These appliances store and serve our video content, with the sole responsibility of
9 delivering playable bits to client devices as fast as possible.”¹³

10 44. Within the CDN, Netflix stores the TV programs and movies that it
11 offers as a series of files. Netflix’s customers access Netflix’s video content
12 through different types of client devices, including digital televisions, desktop
13 computers, laptop computers, tablet computers, and mobile phones. These client
14 devices are produced by many different manufacturers. Each device has certain
15 capabilities and features that require media to be delivered in a specific form or
16 format. In many cases, the media format used by one device cannot be used by
17 another. Thus, Netflix must make its content available to its users in many different
18 formats.

19 45. In Netflix’s words:

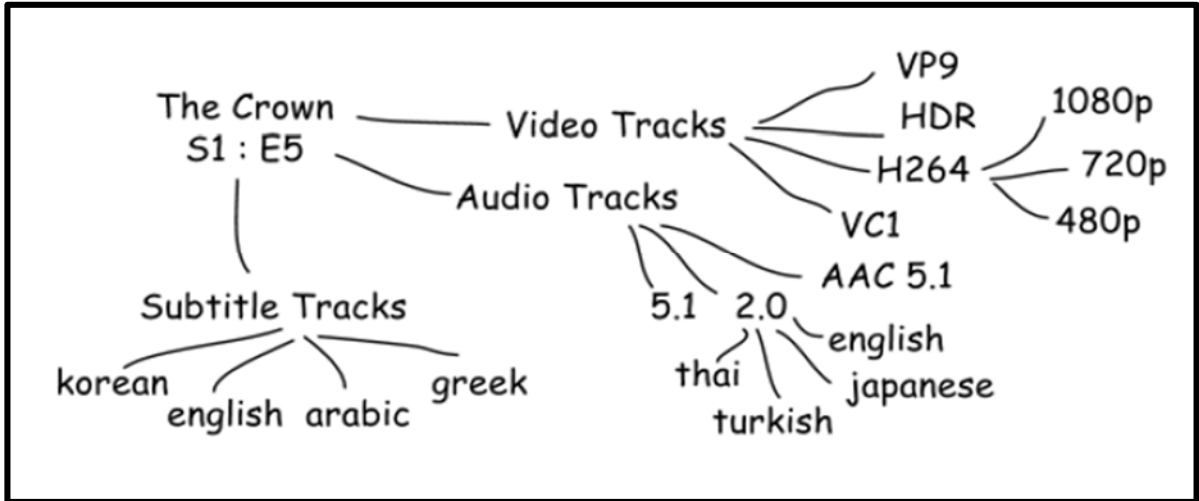
20 Every title is encoded in multiple formats, or *encoding*
21 *profiles*. For example, some profiles may be used by iOS
22 devices and others for a certain class of Smart TVs.
23 There are video profiles, audio profiles, and profiles that
24 contain subtitles. Each audio and video profile is encoded

23 ¹¹ <https://aws.amazon.com/agreement/>.

24 ¹² <https://openconnect.netflix.com/Open-Connect-Overview.pdf>.

¹³ *Id.*

1 into different levels of quality. For a given title, the
 2 higher the number of bits used to encode a second of
 3 content (bps), the higher the quality....Finally, we have
 4 audio profiles and subtitles available in multiple
 5 languages. So for each quadruple of (title, encoding
 profile, bitrate, language), we need to cache one or more
 files. As an example, for streaming one episode of The
 Crown we store around 1,200 files!”¹⁴



6
7
8
9
10
11
12
13 Source: [https://medium.com/netflix-techblog/content-popularity-for-open-](https://medium.com/netflix-techblog/content-popularity-for-open-connect-b86d56f613b)
 14 [connect-b86d56f613b](https://medium.com/netflix-techblog/content-popularity-for-open-connect-b86d56f613b).

15 46. Netflix divides these various content files “into 2 second intervals
 16 called *segments*.”¹⁵

17 47. When a Netflix subscriber initiates a playback session by selecting a
 18 movie or TV show and pressing “play,” the Netflix system “determine[s] which
 19 specific files are required to handle the playback request—taking individual client
 20 characteristics and current network conditions into account.” The Netflix system
 21 “pick[s] OCAs that the requested files should be served from [and] generates URLs

22
23 ¹⁴ [https://medium.com/netflix-techblog/content-popularity-for-open-connect-](https://medium.com/netflix-techblog/content-popularity-for-open-connect-b86d56f613b)
 24 [b86d56f613b](https://medium.com/netflix-techblog/content-popularity-for-open-connect-b86d56f613b).

¹⁵ <https://cs.uwaterloo.ca/~brecht/papers/iiswc-netflix-wload-2016.pdf>.

1 for these OCAs...” The Netflix system then “hand[s] over URLs of the appropriate
2 OCAs to the client device, and the OCA begins to serve the requested files.”

3 2. A user on a client device requests playback of a title (TV show or movie) from the
4 Netflix application in AWS.

5 3. The playback application services in AWS check user authorization and
6 licensing, then determine which specific files are required to handle the playback
7 request - taking individual client characteristics and current network conditions
8 into account.

9 4. The steering service in AWS uses the information stored by the cache control
10 service to pick OCAs that the requested files should be served from, generates
11 URLs for these OCAs, and hands the URLs over to the playback application
12 services.

13 5. The playback application services hand over URLs of the appropriate OCAs to
14 the client device, and the OCA begins to serve the requested files.

15 Source: <https://openconnect.netflix.com/Open-Connect-Overview.pdf>.

16 48. Additionally, when a Netflix subscriber initiates a playback session,
17 the Netflix system provides the “offsets of all segments.”¹⁶ The client device, under
18 the control of the Netflix application running on it, then “downloads segment
19 offsets and content from multiple content files with different bit rates, then selects a
20 starting bitrate that can be supported by available network bandwidth.”¹⁷ The client
21 device “continue[s] to download segments sequentially from the same file unless
22 network or server conditions change (which may result in switching to a different
23 bit rate) or a user event occurs (e.g., stopping or skipping to a new title position).”¹⁸
24 Thus, and as described further below, Netflix performs a “method for balancing
transmission unit traffic over network links.”

¹⁶ <https://cs.uwaterloo.ca/~brecht/papers/iiswc-netflix-wload-2016.pdf>.

¹⁷ *Id.*

¹⁸ *Id.*

1 49. The Netflix CDN also “dispos[es]” the “transmission units into flows.”
2 For example, each of the above-described “segments” of a given content file
3 constitutes a “transmission unit.” The Netflix system “disposes” these segments
4 “into flows” by causing the transmission of a series of segments of various content
5 files—each a “flow”—to each of its many subscribers during the playback process.

6 50. Upon information and belief, the Netflix system also “group[s] flows
7 into first flow lists,” each of which “correspond[] to a selected network link.” For
8 example, each of the OCAs, and each of the file locations specified by a URL
9 within each OCA, constitutes a “network link.” As explained above, Netflix
10 “group[s] flows into first flow lists” by picking the OCAs, and the specific files
11 thereon, that are used to serve a subscriber’s playback request for each of the
12 thousands, or even millions, of subscribers streaming video content from Netflix at
13 any given moment.

14 51. On information and belief, the Netflix system “determin[es] a traffic
15 metric representative of a traffic load on the selected network link.”

16 52. For example, during the playback process, which is controlled by
17 Netflix software, the client device “intelligently selects which OCA to use.”¹⁹ “It
18 does this by testing the quality of the network connection to each OCA. It will
19 connect to the fastest, most reliable OCA first.”²⁰ “The client keeps running these
20 tests throughout the video streaming process. The client probes to figure out the
21 best way to receive content from the OCA.”²¹ Thus, through the playback process,

22 _____
23 ¹⁹ <http://highscalability.com/blog/2017/12/11/netflix-what-happens-when-you-press-play.html>.

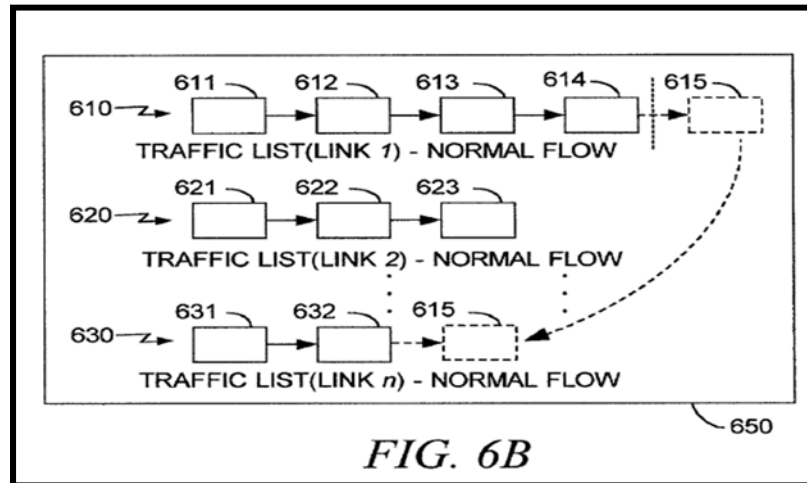
24 ²⁰ *Id.*

²¹ *Id.*

1 the Netflix system “determin[es] a traffic metric representative of a traffic load on
 2 the selected network link.”

3 53. Upon information and belief, “responsive to the traffic metric,” the
 4 Netflix system “regroup[s] flows into second flow lists corresponding to the
 5 selected network link.” For example, over the course of the playback session, the
 6 Netflix system redirects flows of content to different OCAs and to different URLs
 7 on the same OCA (i.e., different “network links”), thereby regrouping the original
 8 set of flows from a given OCA or URL to various Netflix subscribers to a second
 9 flow list corresponding to a different OCA or URL.

10 54. The “regrouping balanc[es] the transmission unit traffic upon the
 11 network links.” This element is illustrated, for example, by Figure 6B of the '079
 12 Patent.



13 Source: '079 Patent Specification, Fig. 6B.

14 55. For instance, the Netflix system moves flows of content files from one
 15 OCA or URL to another less congested or otherwise more favorable OCA or URL.
 16

17 56. Finally, the Netflix system “transmit[s] the respective second flow list
 18 over the respective selected network link.” For instance, after the Netflix system
 19
 20
 21
 22
 23
 24

1 transitions a flow of a particular content files to a new OCA or URL, it continues to
2 transmit all the flows then associated with the new OCA or URL (i.e., the “second
3 flow list”) to the many Netflix subscribers who are streaming from that OCA or
4 URL at that moment.

5 57. At least as of on or around September 26, 2019, when the Broadcom
6 Entities informed Netflix of its infringement of the ’079 Patent, and by no later than
7 the filing and service of the original complaint in this matter, Netflix has had
8 knowledge of the ’079 Patent and Netflix’s infringement thereof.

9 58. Netflix also has induced, and continues to induce, direct infringement
10 by third-parties of at least claim 1 of the ’079 Patent, at least in the exemplary
11 manner described above, by actively encouraging their use of the Netflix system, in
12 violation 35 U.S.C. § 271(b).

13 59. For example, Netflix has induced, and continues to induce, direct
14 infringement of the ’079 Patent by customers and/or end users of client devices
15 enabled with the Netflix software application and service. In light of the notice the
16 Broadcom Entities provided to Netflix of its infringement of the ’079 Patent,
17 Netflix knows that it provides and specifically intends to provide an application for
18 use on client devices that, when used as intended with the Netflix streaming service,
19 meets the limitations of claim 1 of the ’079 Patent. Netflix knows and specifically
20 intends that its end users practice the method recited in claim 1 of the ’079 Patent,
21 when using its application and service as intended.

22 60. Netflix’s knowing and willful infringement of the ’079 Patent has
23 caused and continues to cause damage to Avago, and Avago is entitled to recover
24

1 damages sustained as a result of Netflix’s wrongful acts in an amount subject to
2 proof at trial.

3 **SECOND CLAIM FOR RELIEF**

4 **(Infringement of U.S. Patent No. 8,259,121)**

5 61. The Broadcom Entities reallege and incorporate by reference the
6 allegations of paragraphs 1-60 set forth above.

7 62. The ’121 Patent, entitled “System and Method for Processing Data
8 Using a Network,” was duly and legally issued on September 4, 2012 from a patent
9 application filed on December 9, 2002, with Patrick Law, Darren Neuman, and
10 David Baer as the named inventors. A copy of the ’121 Patent is attached hereto as
11 **Exhibit B.**

12 63. The ’121 Patent claims priority from U.S. Provisional Application No.
13 60/420,151, filed on December 9, 2002.

14 64. The ’121 Patent is assigned to Avago, which currently holds all
15 substantial rights, title, and interest in and to the ’121 Patent.

16 65. Pursuant to 35 U.S.C. § 282, the ’121 Patent is presumed valid.

17 66. The claims in the ’121 Patent are directed to an improved network for
18 processing audio and visual (“A/V”) data. Specifically, the inventions described in
19 the ’121 Patent “relate[] to a network environment in an A/V system using ‘A/V
20 decoders,’ where the A/V decoders are adapted to process, decode or decompress
21 one or more input data streams.”²²

22 67. The ’121 Patent addresses a technical problem in a network processing
23 A/V data. The patent explains that, at the time, there was “no known

24 _____
²² ’121 Patent, 1:42-45.

1 methodological way to connect video processing modules in A/V systems” and that
2 “[m]ost video processing modules are connected together in an ad hoc manner.”²³
3 “As a result, such ad-hoc designs may become difficult to verify, maintain and
4 reuse. Furthermore, as more features are added to the A/V systems...it becomes
5 more difficult to design and integrate such features properly.”²⁴ Thus, there was “a
6 need for an architecture or network that provides a general model illustrating how
7 various video processing modules behave in a network environment.”²⁵

8 68. The ’121 Patent describes and claims solutions to these technical
9 problems, including specific, novel networks with various features for processing
10 A/V data.

11 69. The inventions described and claimed in the ’121 Patent have
12 applications both in the context of the ecosystem within a computer or device, and
13 in more complex computer networking environments. Indeed, as the patent
14 explains, “[m]any modifications and variations of the present invention are possible
15 in light of the above teachings. Thus, it is to be understood that, within the scope of
16 the appended claims, the invention may be practiced otherwise than as described
17 hereinabove.”²⁶

18 70. The systems and methods described in the ’121 Patent improve the
19 functionality of computer networks used for processing A/V data by providing a
20 new, advantageous approach for those networks.

21 71. Claim 1 of the ‘121 Patent is directed to:

22 _____
23 ²³ *Id.* at 1:48-51.

24 ²⁴ *Id.* at 1:51-55.

²⁵ *Id.* at 1:66-2:1.

²⁶ *Id.* at 16:26-30.

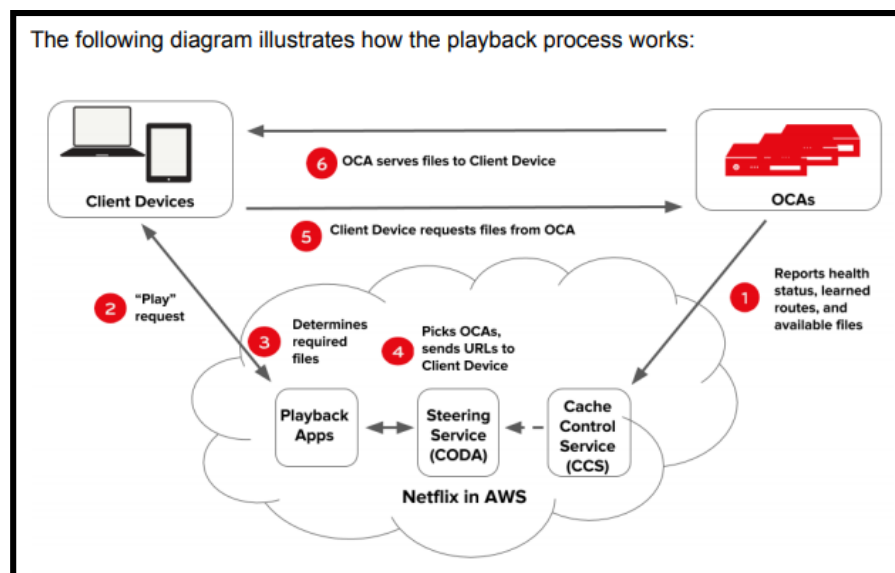
1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24

A network for processing data configured by a controller to form at least one display pipeline therein by dynamically selecting use of at least two selectable nodes from a plurality of selectable nodes and dynamically concatenating the selected at least two selectable nodes in the network together, wherein said at least one display pipeline has an independent data rate and a flow control module enables said independent data rate.

72. Netflix directly infringes the '121 Patent by making, using, offering to sell, and/or selling in the United States its Netflix service, which utilizes Netflix's CDN to process audio and visual data in a manner that uses the inventions claimed in the '121 Patent.

73. Upon information and belief, Netflix directly infringes at least independent claim 1 of the '121 Patent at least in the exemplary manner described below.

74. Netflix created, operates, and maintains a "network for processing data," namely, the CDN that Netflix uses to stream TV shows and movies to its subscribers over the internet, as illustrated in the following diagram from Netflix's website.



Source: <https://openconnect.netflix.com/Open-Connect-Overview.pdf>.

1 75. As shown in this diagram, the Netflix CDN consists of three primary
2 groups of systems: (1) the client devices (e.g., smart TVs, computers, mobile
3 phones, etc.) that subscribers use to access the Netflix service with the help
4 software applications that were developed, or partially developed, by Netflix; (2)
5 the Netflix “backend” (referred to in the diagram as “Netflix in AWS”), which
6 receives and processes requests for video content from subscribers; and (3) a
7 network of OCAs, which deliver the video content to the client devices.

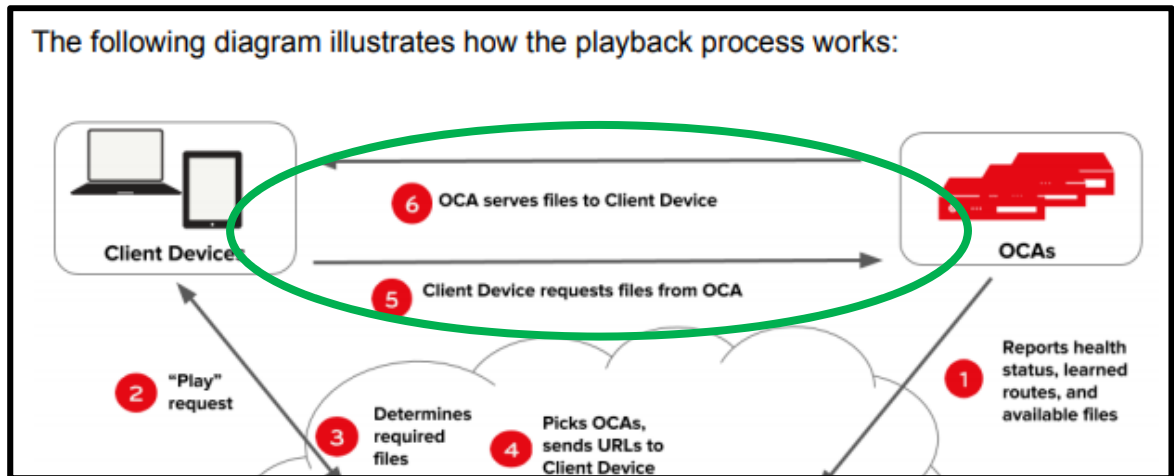
8 76. The CDN is “configured by a controller.” For example, when a
9 Netflix user requests playback of a particular title (e.g., a film or television
10 program) using a client device, various computing resources in the Netflix backend,
11 which are controlled by Netflix and run Netflix applications: (1) receive the
12 request; (2) determine which specific streaming assets are required to handle the
13 request, taking individual client characteristics and current network conditions into
14 account; (3) pick OCAs from which the requested streaming assets should be
15 streamed; and (4) provide a manifest file to the client device containing URLs that
16 specify the OCAs and files needed for the playback process. Thus, Netflix
17 “controls” the path of the delivery of video content through the network to the
18 Netflix subscribers.

- 19
- 20 2. A user on a client device requests playback of a title (TV show or movie) from the
Netflix application in AWS.
- 21 3. The playback application services in AWS check user authorization and
22 licensing, then determine which specific files are required to handle the playback
request - taking individual client characteristics and current network conditions
23 into account.
- 24

4. The steering service in AWS uses the information stored by the cache control service to pick OCAs that the requested files should be served from, generates URLs for these OCAs, and hands the URLs over to the playback application services.
5. The playback application services hand over URLs of the appropriate OCAs to the client device, and the OCA begins to serve the requested files.

Source: <https://openconnect.netflix.com/Open-Connect-Overview.pdf>.

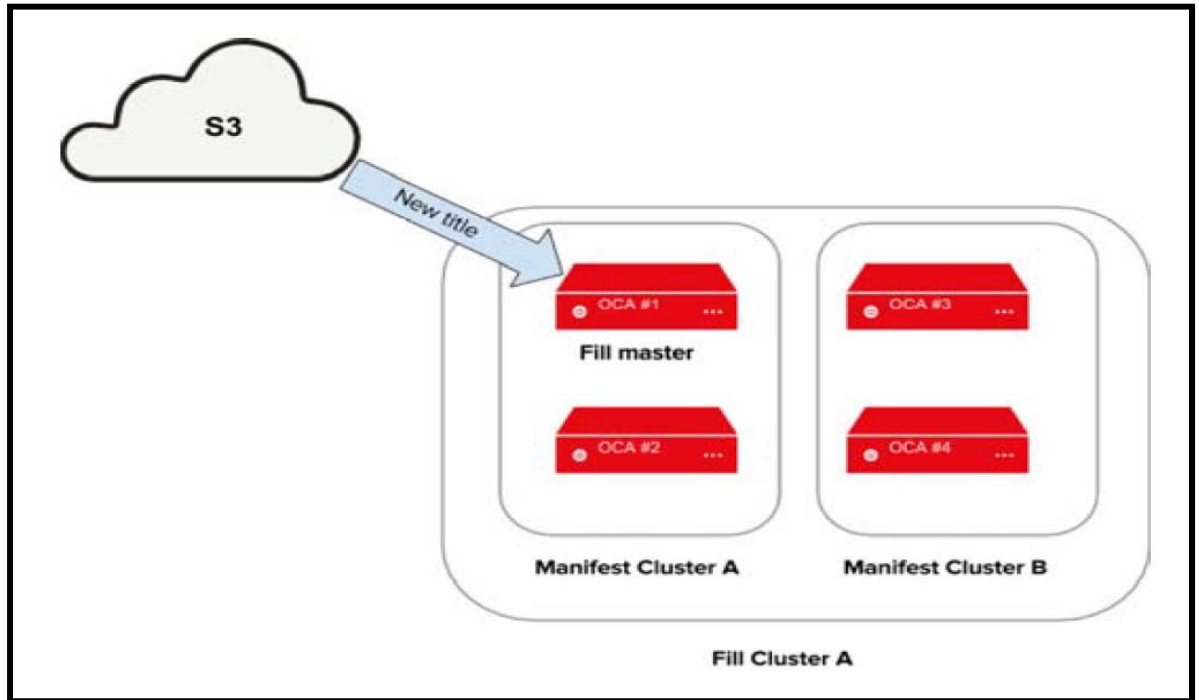
77. Through this process, Netflix’s content delivery network forms “at least one display pipeline therein.” For example, the CDN establishes a link between the client device and one or more Netflix OCAs through which the various video, audio, and other files associated with the requested title are streamed.



Source: <https://openconnect.netflix.com/Open-Connect-Overview.pdf>.

78. The “at least one display pipeline” within Netflix’s content delivery network is formed “by dynamically selecting use of at least two selectable nodes from a plurality of selectable nodes.” For instance, the content delivery network forms display pipelines between the millions of client devices streaming content from Netflix all over the world (“nodes”) and the numerous OCAs and other servers that Netflix uses to stream its content (also “nodes”).

1 79. The OCAs periodically receive new content files—which constitute
 2 video and other A/V data—from Netflix’s backend systems and from other OCAs.



3
4
5
6
7
8
9
10
11
12
13 Source: <https://medium.com/netflix-techblog/netflix-and-fill-c43a32b490c0>.

14 The above image, generated by Netflix, shows that the Netflix local OCAs are
 15 supplied with content by Netflix’s more remotely-located repository of content
 16 titles, which creates and stores copies of content that has been transcoded into
 17 various formats, as discussed above.

18 80. Upon receiving a request for content from a client device, the OCA
 19 identifies and sends the specific files requested, filtering out the requested data from
 20 the numerous content files contained on the OCA.

21 81. Netflix’s selection of nodes is “dynamic.” For example, Netflix’s
 22 “control plane services in AWS take the data that the OCAs report and use it to
 23
 24

1 steer clients via URL to the most optimal OCAs given their file availability, health,
 2 and network proximity to the client.”²⁷

3 82. As another example of Netflix’s “dynamic” selection of nodes, the
 4 Netflix system switches between content files and OCAs during the playback
 5 process in order to automatically adapt to network conditions and user behavior.
 6 This “dynamic” selection is illustrated, for instance, in Figure 1 of *Characterizing*
 7 *the Workload of a Netflix Streaming Video Server*, a technical paper published in
 8 2016 by the Institute of Electrical and Electronics Engineers (“IEEE”). The portion
 9 of this figure within the green box shows, for example, how the Netflix system
 10 switches between different video quality levels, or bit rates (shown on the vertical
 11 axis), over the course of a Netflix content streaming session (i.e., a Netflix user
 12 watching a movie or television program). The horizontal axis shows the elapsed
 13 time (in minutes) of that streaming session:




Fig. 1. Requests issued during a session

Source: <https://cs.uwaterloo.ca/~brecht/papers/iiswc-netflix-wload-2016.pdf>.

²⁷ <https://openconnect.netflix.com/Open-Connect-Overview.pdf>.

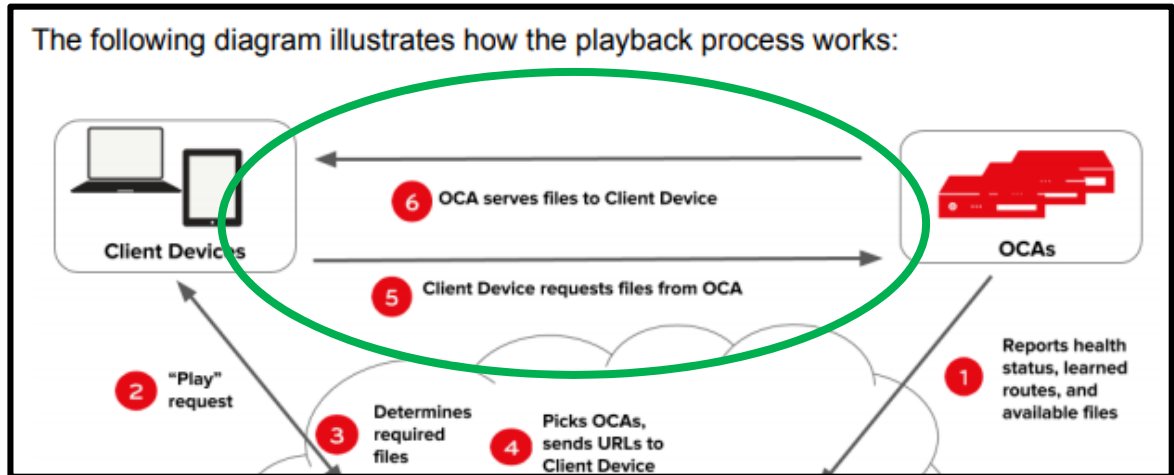
1 83. The “dynamic” selection has also been described by other third parties
2 that have analyzed the Netflix system, further illustrating how the Netflix
3 application on the client device—which Netflix controls—automatically switches
4 between URLs for video files and OCAs over the course of the streaming session:

- 5 ▪ Have you noticed when watching a video the picture quality varies? Sometimes it will
6 look pixelated, and after awhile the picture snaps back to HD quality? That’s because
7 the client is adapting to the quality of the network. If the network quality declines, the
8 client lowers video quality to match. The client will switch to another OCA when the
9 quality declines too much.

10 Source: <http://highscalability.com/blog/2017/12/11/netflix-what-happens-when-you-press-play.html>.

11
12 84. Thus, the “selection” of the “nodes” is “dynamic” because the node
13 selection can be performed instantaneously, upon the occurrence of network events,
14 and in response to changing network conditions.

15 85. The Netflix CDN “dynamically concatenate[s] the selected at least two
16 selectable nodes in the network together.” The selectable nodes are “concatenated”
17 when a connection is established between them. For example, as explained above,
18 the network establishes a connection between the client device and the best suited
19 OCA through which the client device, using instructions provided by Netflix,
20 “requests files from the OCA” and the “OCA serves files to the Client Device.”



Source: <https://openconnect.netflix.com/Open-Connect-Overview.pdf>.


86. As another example of the dynamic nature of the concatenation, the client devices within Netflix’s CDN continue to “adapt to the quality of the network...If the network quality declines, the client lowers video quality to match. The client will switch to another OCA when the quality declines too much.”²⁸

87. The “display pipeline” within the Netflix CDN “has an independent data rate.” For example, on information and belief, the data rate of the transmission of streaming content between the OCA and a given client device within the Netflix system varies from and is not dependent upon the data rate of the transmission of streaming content to the other client devices on the system.

88. The Netflix CDN includes a “flow control module” that “enables said independent data rate.” For instance, the Netflix client serves as a “flow control module” by “pacing” the transmission of the streaming content from the Netflix OCA. As illustrated below, the client device—under the control of the Netflix application operating thereon—controls the pace at which it downloads content

²⁸ <http://highscalability.com/blog/2017/12/11/netflix-what-happens-when-you-press-play.html>.

1 files from Netflix to address network issues and user events, including by varying
 2 the bit rate of the streaming video content.



bandwidth. Clients continue to download segments sequentially from the same file unless network or server conditions change (which may result in switching to a different bit rate) or a user event occurs (e.g., stopping or skipping to a new title position). Clients that are in a steady state limit the number of segments they download ahead of the playback point to avoid waste when a user event occurs. This is called *pacing* [30] and is a defining feature of HTTP streaming video clients [3]. There is no simple relationship between segments

Source: <https://cs.uwaterloo.ca/~brecht/papers/iiswc-netflix-wload-2016.pdf>.

89. Upon information and belief, Netflix controls the playback functionality of the client devices used by subscribers to access the Netflix service. As an example, for many client devices, Netflix develops the software applications used to access the Netflix service. Even in client devices for which Netflix does not develop the application, it “still has control because it controls the software development kit (SDK).” The SDK “is a set of software development tools that allows the creation of applications. Every Netflix app...plays video using the

1 SDK.” “By controlling the SDK, Netflix can adapt consistently and transparently
2 to slow networks, failed OCAs, and any other problems that may arise.”²⁹

3 90. In Netflix’s words, “[t]he SDK provides a rendering engine, JavaScript
4 runtime, networking, security, *video playback*, and other platform hooks.”

5 that runs on the metal, and a UI written in JavaScript. The SDK provides a
6 rendering engine, JavaScript runtime, networking, security, video
playback, and other platform hooks. Depending on the device, SDK

7 Source: [https://medium.com/netflix-techblog/building-the-new-netflix-
8 experience-for-tv-920d71d875de](https://medium.com/netflix-techblog/building-the-new-netflix-experience-for-tv-920d71d875de).

9 91. At least as of on or around September 26, 2019, when the Broadcom
10 Entities informed Netflix of its infringement of the ’121 Patent, and by no later than
11 the date of the original complaint in this matter, Netflix has had knowledge of the
12 ’121 Patent and that its video streaming service infringes the ’121 Patent.

13 92. In addition to direct infringement, Netflix indirectly infringes the ’121
14 Patent in violation of 35 U.S.C. 271(b) by inducing third-parties to directly infringe
15 at least claim 1 of the ’121 Patent, at least in the exemplary manner described
16 above. Netflix has induced, and continues to induce, direct infringement of the
17 ’121 Patent by customers and/or end users of infringing client devices enabled with
18 the Netflix software application and service. Netflix knows that it provides and
19 specifically intends to provide an application for use on client devices that, when
20 used as intended with the Netflix streaming service, meets the limitations of claim 1
21 of the ’121 Patent. Netflix knows and specifically intends that its customers and
22 end users practice the system recited in claim 1 of the ’121 Patent, when using its

23 _____
24 ²⁹ [http://highscalability.com/blog/2017/12/11/netflix-what-happens-when-you-
press-play.html](http://highscalability.com/blog/2017/12/11/netflix-what-happens-when-you-press-play.html).

1 application and service as intended. As explained above, the relevant aspects of
2 that use, including the “playback” process for streaming content, are controlled by
3 Netflix software on the client devices, along with other Netflix systems.

4 93. Netflix’s knowing and willful infringement of the ’121 Patent has
5 caused and continues to cause damage to Avago. Avago is entitled to recover
6 damages sustained as a result of Netflix’s wrongful acts in an amount subject to
7 proof at trial.

8 **THIRD CLAIM FOR RELIEF**

9 **(Infringement of U.S. Patent No. 8,959,245)**

10 94. The Broadcom Entities reallege and incorporate by reference the
11 allegations of paragraphs 1-93 set forth above.

12 95. The ’245 Patent, entitled “Multiple Pathway Session Setup to Support
13 QoS Services” was duly and legally issued on February 17, 2015, from a patent
14 application filed on November 25, 2008, with Jeyhan Karaoguz and James Bennett
15 as the named inventors. A copy of the ’245 Patent is attached hereto as **Exhibit C**.

16 96. The ’245 Patent is assigned to Avago, which currently holds all
17 substantial rights, title, and interest in and to the ’245 Patent.

18 97. Pursuant to 35 U.S.C. § 282, the ’245 Patent is presumed valid.

19 98. The ’245 Patent is directed to an improvement in the functionality of a
20 communication network, such as the internet.

21 99. Specifically, the ’245 Patent’s claims are directed to a novel system
22 and method for delivering content to a user through a communication network, in
23
24

1 which a network management server determines multiple routes for delivering the
2 content based on a provisioning profile for the user device.³⁰

3 100. The inventions of the '245 Patent resolve technical problems related to
4 delivering content through the then-existing internet network architecture. At the
5 time, the internet was becoming increasingly important as a commercial
6 infrastructure and there was a growing need for “massive internet based services,”
7 such as voice over internet protocol (“VoIP”), video-conferencing, and video
8 streaming. Much of the then-current internet architecture was based on the “best
9 effort” model. This architecture attempts to deliver all data traffic “as soon as
10 possible within the limits of its abilities.” However, the data packets being
11 transferred “can be dropped indiscriminately” in the event of network congestion.
12 While this approach worked well for some less time-sensitive applications, such as
13 email and FTP data transfer, it did not work as well for real-time multimedia
14 applications, such as streaming video on demand.³¹

15 101. In light of the growing use of high-bandwidth, time-sensitive
16 applications, there was a need for technologies to improve the quality of service
17 (“QoS”) of data transmissions over communication networks like the internet.

18 102. The inventions described in the '245 Patent address technical problems
19 associated with the conventional systems and methods for delivering high-quality
20 video, and other data, including by utilizing multiple routes among the available
21 routes in the communication network, thereby increasing the reliability of the data
22 transmission.³²

23 _____
³⁰ '245 Patent, Abstract, 1:44-51, 2:13-40.

24 ³¹ *Id.* at 1:19-40.

³² *See id.* at 2:29-32, 2:41-45.

1 103. The '245 Patent describes and claims specific ways to implement this
2 solution using a network management server capable of determining multiple routes
3 for delivering the data content based on a “provisioning profile,” thereby better
4 ensuring delivery of content over the communication network. The '245 Patent
5 explains how the provisioning profile may contain information relevant to the
6 delivery of the content, such as preferred service types, desired quality of service,
7 client account information, and/or client credit information.³³

8 104. As the '245 Patent explains, different sets of packets associated with
9 the data content may be transmitted over different routes amongst the multiple
10 available routes to, for instance, take advantage of paths that have less usage and
11 increase reliability. The network management server can manage and prioritize this
12 allocation of routes, which may involve the use of a primary route and/or one or
13 more secondary routes. The network management server can also enable a
14 “handoff” between the routes, such as when QoS degrades on the primary route.
15 The handoff can be seamless to the user, ensuring an uninterrupted user
16 experience.³⁴

17 105. The systems and methods described and claimed in the '245 Patent
18 improve the functionality of a computer network by providing a new, advantageous
19 approach for delivering content through the network that enables higher QoS
20 standards, such as those required for video-on-demand applications.

21 106. Claim 1 of the '245 Patent is directed to:

22 1. A method for communication, the method comprising:
23

24 ³³ *Id.* at 2:29-32.

³⁴ *Id.* at 2:41-61.

1 receiving from a user device, by a network
2 management server via a communication network,
3 a request for a service;
4 determining multiple routes for delivering content
5 associated with said requested service based on a
6 provisioning profile for said user device;
7 and delivering said content associated with said
8 requested service via said determined multiple
9 routes.

10 107. Claim 1 thus claims a novel solution for transmitting digital media
11 content over a communication network via multiple routes, using a network
12 management server and a provisioning profile. This solution was not well-
13 understood, routine, or conventional at the time of the '245 Patent because it claims
14 a new and specific improvement over the prior art. The inventions claimed in the
15 '245 Patent comprise a novel arrangement of streaming content equipment that
16 results in a better experience for the content user with fewer interruptions.

17 108. Claim 3 of the '245 Patent, which depends from claim 1, is directed to:

18 3. The method according to claim 1, wherein said
19 provisioning profile comprises preferred service types,
20 desired QoS for one or more services, client account
21 information, and/or client credit verification information.

22 109. Claim 6 of the '245 Patent, which also depends from claim 1, is
23 directed to:

24 6. The method according to claim 1, comprising
allocating via said network management server, one or
more of said determined multiple routes based on priority.

110. Thus, claims 3 and 6 further describe the invention's improved method
whereby a network management server determines multiple routes for delivering
content based on a provisioning profile in response to receiving a request for

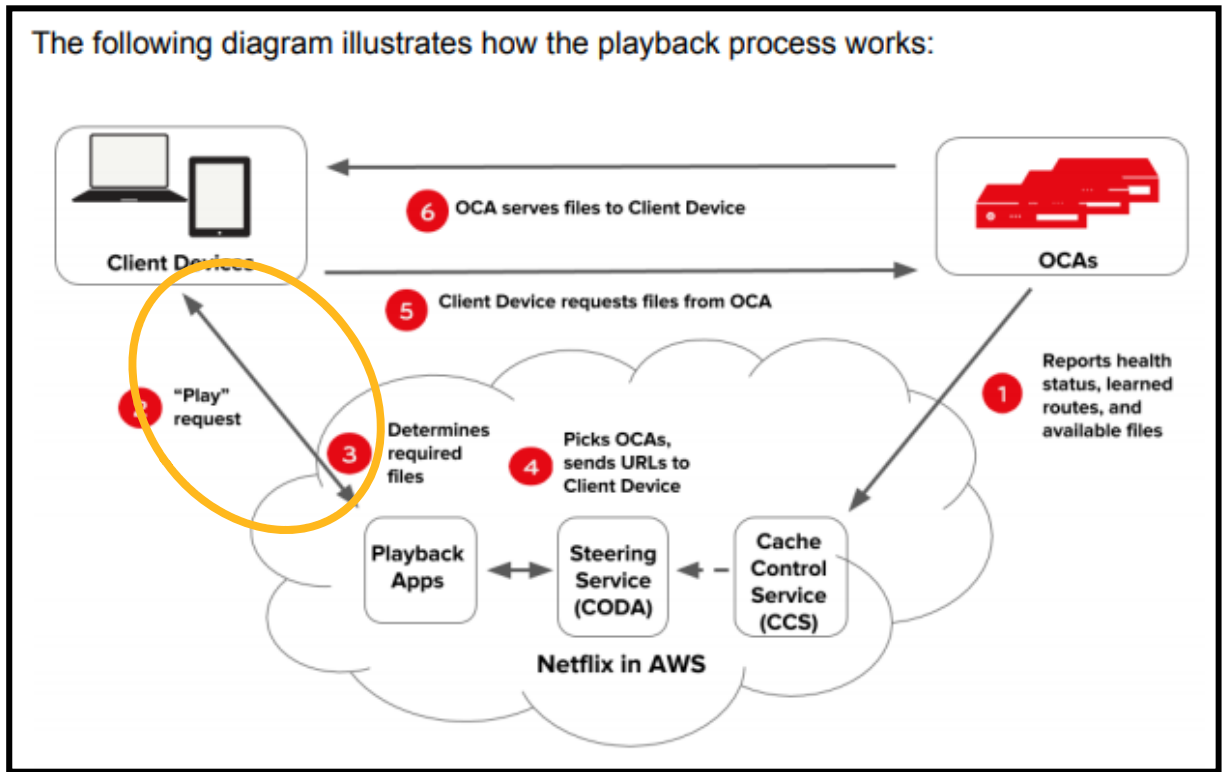
1 service from a user device, and then delivers that content via multiple routes based
2 on priority. The ordered combination of elements in each of claims 3 and 6, in
3 conjunction with the elements of the claim from which they depend, therefore recite
4 unconventional, new, and improved digital media content delivery methods that
5 were not well-understood at the time of the '245 Patent.

6 111. Netflix directly infringes the '245 Patent by making, using, offering to
7 sell, and/or selling into the United States its Netflix service, which utilizes a
8 playback system that practices the inventions claimed in the '245 Patent.

9 112. Upon information and belief, Netflix directly infringes at least claims
10 1, 3, and 6 of the '245 Patent, at least in the exemplary manner described below.

11 113. Netflix practices a “method for communication” that involves
12 “receiving from a user device, by a network management server via a
13 communication network, a request for a service.” For example, when a Netflix user
14 uses the Netflix application on their TV, computer, smartphone, or other client
15 device to view a movie or television program, the Netflix application sends a “play”
16 request to Netflix’s backend systems. This represents a “request for service” from a
17 “user device.” The Netflix backend systems then determine, using Netflix-
18 developed solutions, the optimal manner in which to deliver the requested content.

19
20
21
22
23
24



Source: <https://openconnect.netflix.com/Open-Connect-Overview.pdf>.

114. On information and belief, Netflix controls the playback process from the client side as well. For example, as explained above, Netflix develops many of the Netflix applications used by client devices to access the Netflix service, including the applications for Android and iOS devices. As for the client devices for which Netflix does not develop the application itself, it provides the software development kit, or SDK, used to create the Netflix application. On information and belief, the SDK includes the code responsible for the playback process.

115. Netflix’s control of the playback process at the client side is confirmed by third party sources:

Netflix Controls The Client

Netflix handles failures gracefully because it controls the client on every device running Netflix.

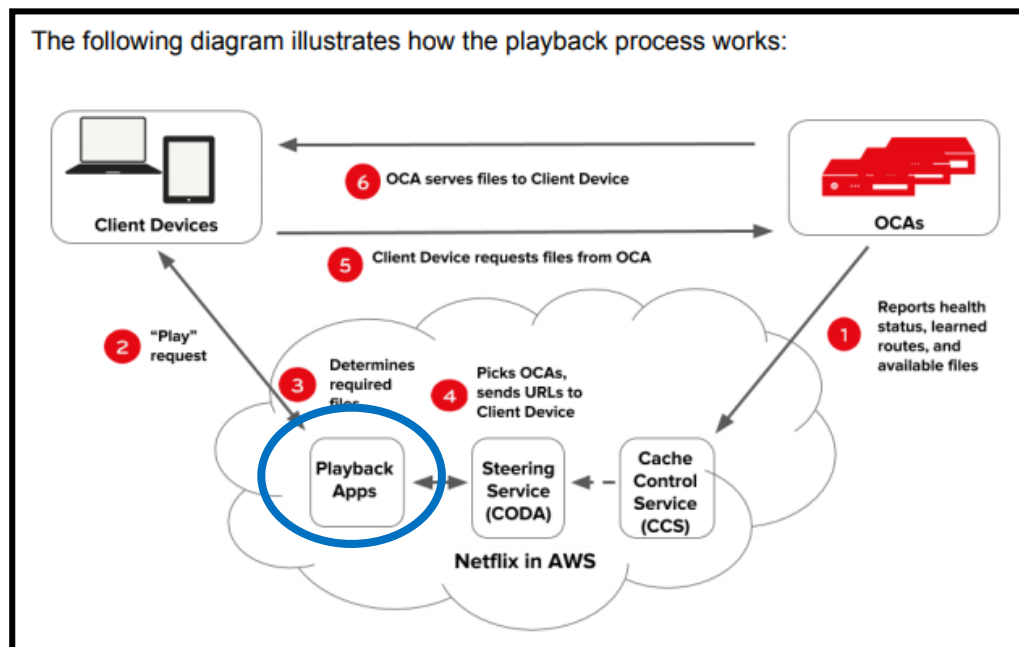
Netflix develops its Android and iOS apps themselves, so you might expect them to control those. But even on platforms like Smart TVs, where Netflix doesn't build the client, Netflix still has control because it controls the *software development kit* (SDK).

A SDK is a *set of software development tools that allows the creation of applications*. Every Netflix app makes requests to AWS and plays video using the SDK.

By controlling the SDK, Netflix can adapt consistently and transparently to slow networks, failed OCAs, and any other problems that might arise.

Source: <http://highscalability.com/blog/2017/12/11/netflix-what-happens-when-you-press-play.html>.

116. The request for service from the Netflix user's client device is received "by a network management server via a communication network." For example, as shown below, the playback request is received by the Playback Apps service within the Netflix backend, which, on information and belief, operates on one or more servers.



Source: <https://openconnect.netflix.com/Open-Connect-Overview.pdf>.

1 117. The Netflix system also “determin[es] multiple routes for delivering
2 content associated with said requested service based on a provisioning profile for
3 said user device.” For instance, upon receiving a playback request for a specific
4 title from a Netflix user, the Netflix backend systems determine the video, audio,
5 and other files needed for playback and pick the OCAs from which these files
6 should be streamed to the client device. It then generates a manifest file containing
7 the URLs for the files and OCAs, which the Netflix system sends to the client
8 device.

9 118. As Netflix explains:


- 10 4. The steering service in AWS uses the information stored by the cache control
11 service to pick OCAs that the requested files should be served from, generates
12 URLs for these OCAs, and hands the URLs over to the playback application
13 services.
- 14 5. The playback application services hand over URLs of the appropriate OCAs to
15 the client device, and the OCA begins to serve the requested files.

16 Source: <https://openconnect.netflix.com/Open-Connect-Overview.pdf>.

17 119. Through this process, the Netflix system will select as many as ten
18 different OCAs from which the requested content may be streamed:

- 19 ■ Taking into account all the relevant information, the Playback Apps service returns
20 URLs for up to ten different OCA servers. These are the same sort of URLs you use all
21 the time in your web browser. Netflix uses your IP address and information from ISPs to
22 identify which OCA clusters are best for you to use.

23 Source: <http://highscalability.com/blog/2017/12/11/netflix-what-happens-when-you-press-play.html>.



Source: Screenshot from Netflix Application.

120. On information and belief, Netflix bases its selection of OCAs and URLs, at least in part, on a “provisioning profile” for the user device used to access the Netflix service. For instance, upon receiving a playback request for a specific title, the Netflix system “checks user authentication and licensing” and takes “individual client characteristics into account” in selecting the “specific streaming assets” required to handle the playback request and the OCAs from which they should be streamed.

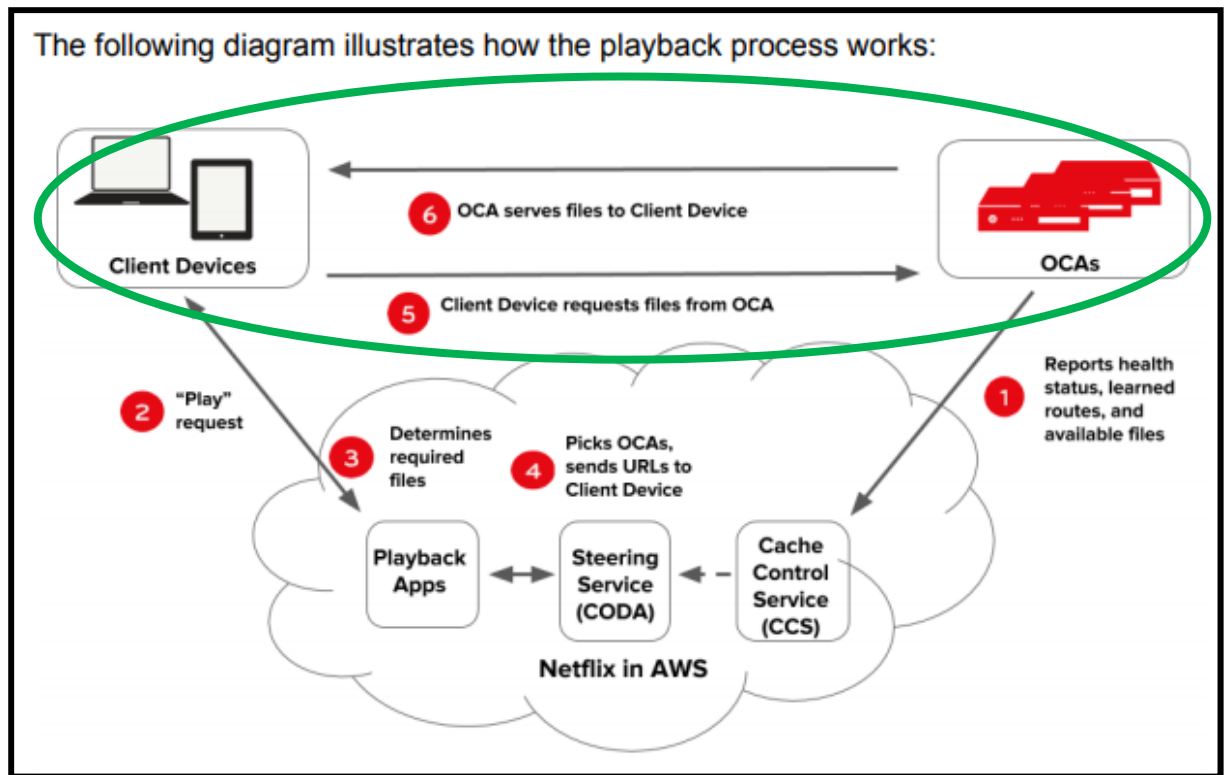
2. A user on a client device requests playback of a title from the Netflix application.
3. The playback application services check user authorization and licensing, then determine which specific streaming assets are required to handle the playback request - taking individual client characteristics and current network conditions into account.
4. The steering service uses the information stored by the cache control service to pick OCAs that the requested video assets should be streamed from, generates URLs for these OCAs, and hands the URLs over to the playback application services.
5. The playback application services hand over URLs of the appropriate OCAs to the client device, and video streaming starts.

Source: <https://openconnect.netflix.com/Open-Connect-Overview.pdf>.

121. As further explained in a study published by the IEEE, the client device—under the control of the Netflix application—“indicates the formats of the content it can play. Netflix server then sends back a manifest file based upon the

1 client request. For instance, Netflix client running on an older computer (Thinkpad
 2 T60 with Windows XP) and a newer computer (MacBook with Snow Leopard)
 3 have different capabilities and received different video downloading format and bit
 4 rates.”³⁵

5 122. The Netflix system “deliver[s] said content associated with said
 6 requested service via said determined multiple routes.” For example, the Netflix
 7 service connects to one or more of the numerous OCAs within Netflix’s CDN and
 8 streams video, audio, and other content to the user’s computer.



20 Source: <https://openconnect.netflix.com/Open-Connect-Overview.pdf>.

21 123. In doing so, the Netflix application on the client device intelligently
 22 selects which of the specific OCAs identified by Netflix to use by evaluating
 23 various factors, including the quality of the network connection to each OCA. The

24 ³⁵ <https://www.moritzsteiner.de/papers/netflix-hulu.pdf>.

1 client device connects to the fastest, most reliable OCA first, but will switch to
2 another OCA if the video quality declines too much.

- 3 ▪ The client intelligently selects which OCA to use. It does this by testing the quality of the
4 network connection to each OCA. It will connect to the fastest, most reliable OCA first.
5 The client keeps running these tests throughout the video streaming process.
- 6 ▪ The client probes to figure out the best way to receive content from the OCA.
- 7 ▪ The client connects to the OCA and starts streaming video to your device.
- 8 ▪ Have you noticed when watching a video the picture quality varies? Sometimes it will
9 look pixelated, and after awhile the picture snaps back to HD quality? That's because
10 the client is adapting to the quality of the network. If the network quality declines, the
11 client lowers video quality to match. The client will switch to another OCA when the
12 quality declines too much.

13 Source: [http://highscalability.com/blog/2017/12/11/netflix-what-happens-](http://highscalability.com/blog/2017/12/11/netflix-what-happens-when-you-press-play.html)
14 [when-you-press-play.html](https://openconnect.netflix.com/Open-Connect-Overview.pdf); [https://openconnect.netflix.com/Open-Connect-](https://openconnect.netflix.com/Open-Connect-Overview.pdf)
15 [Overview.pdf](https://openconnect.netflix.com/Open-Connect-Overview.pdf).

16 124. With regard to claim 3 of the '245 Patent, the Netflix system utilizes a
17 “provisioning profile” comprising “preferred service types, desired QoS for one or
18 more services, client account information, and/or client credit verification
19 information.” For example, the Netflix system maintains account information for
20 its users. This includes, amongst other thing, preferred playback settings, plan
21 details (e.g., standard or HD), user profile information, parental control
22 information, and credit card payment and billing information. On information and
23 belief, some or all of this account information is used by the Netflix system in
24 selecting the content files and OCAs for use in delivering the title selected by the
25 Netflix user.

26 125. With regard to claim 6, the Netflix system also allocates multiple
27 routes for delivering the streamed content “based on priority.” For instance, on

1 information and belief, the list of OCAs that Netflix provides to the client device
2 via the manifest file at the start of the playback session is ranked according to
3 priority rules established by Netflix.³⁶

4 126. As alleged above, Netflix directly infringes at least claims 1, 3, and 6
5 of the '245 Patent.

6 127. At least as of on or around September 26, 2019, when the Broadcom
7 Entities informed Netflix of its infringement of the '245 Patent, and by no later than
8 the date of the original complaint in this matter, Netflix has had knowledge of the
9 '245 Patent and its infringement thereof.

10 128. In addition to direct infringement, Netflix has induced, and continues
11 to induce, direct infringement by third-parties of at least claims 1, 3, and 6 of the
12 '245 Patent, at least in the exemplary manner described above, by actively
13 encouraging their use of the Netflix system, in violation 35 U.S.C. § 271(b).
14 Netflix has induced, and continues to induce, direct infringement of the '245 Patent
15 by customers and/or end users of playback devices enabled with the Netflix
16 software application and service. Netflix knows that it provides and specifically
17 intends to provide an application for use on playback devices that, when used as
18 intended with the Netflix streaming service, meets the limitations of claims 1, 3,
19 and 6 of the '245 Patent. Netflix knows and specifically intends that its end users
20 practice the method recited in claims 1, 3, and 6 of the '245 Patent, when using its
21 application and service as intended.

22
23
24 ³⁶ See, e.g., <http://oc.nflxvideo.net/docs/OpenConnect-Deployment-Guide.pdf>;
<https://www.moritzsteiner.de/papers/netflix-hulu.pdf>.

1 quality level than the quality level they are currently experiencing. As explained in
2 the specification, the resource allocation and utilization can be determined by
3 quality control modules and communication modules. These modules
4 communicate various capability information about the network, such as processing
5 capability, communication capability, and information access capability. The
6 quality control modules can determine whether utilizing resources of another
7 system will provide the service to the user at a higher quality level. If it so
8 determines, a distributing processing module can manage the resource allocation.³⁷

9 136. The inventions of the '992 Patent address technical problems related to
10 delivering content over unstable network environments, an issue with prior art
11 system existing at the time. In a dynamic and unstable network environment,
12 processing resources continuously toggle between available and unavailable. These
13 processing resources may offer service capabilities that are superior or inferior to
14 other resources present in the network. For example, a system providing low
15 quality audio or video service may communicate with a second system capable of
16 providing a higher quality audio or video service.³⁸

17 137. The first system providing the lower quality service needs a way of
18 both accessing the system with superior resources, utilizing those resources, and
19 ultimately delivering content from the superior resource in order to provide higher
20 quality content to the user.

21
22
23 _____
³⁷ See '992 Patent, 2:8-32.

24 ³⁸ *Id.* at 1:39-42 (describing delivery of high quality audio), 1:42-49 (describing
delivery of high quality video).

1 138. The '992 Patent, therefore, addresses the technical problem of ensuring
2 delivery of content at the highest quality level in an unstable network environment
3 by utilizing a quality-of-service based network resource allocation delivery system.

4 139. The '992 Patent claims specific ways to solve these technical problems
5 with a digital media delivery system that is capable of automatically determining
6 bandwidth capability information in multiple systems, using that information to
7 obtain digital media content at a higher quality level than the current quality, and
8 then delivering that higher quality digital media content to the user.³⁹ In this
9 manner, it improves the functionality of computer networks used to deliver media
10 content by providing a new, advantageous approach for those networks.

11 140. The digital media delivery system disclosed and claimed in the '992
12 Patent uses novel quality control modules, resource allocation modules, and
13 distributing processing modules to assess, manage, and use resources in a network
14 environment where the availability of resources may vary between systems in the
15 network.

16 141. In a two-system network environment, for example, if it is determined
17 that the second system has access to higher quality audio or video content than the
18 first system, then the first system may receive higher quality data from the second
19 system for further processing. The first system can also receive higher quality data
20 from the system for immediate delivery and presentation to the user. Alternatively,
21 the first system may direct the second system to provide the higher quality data
22 directly to the user.⁴⁰

23 _____
³⁹ *Id.* at 3:1-9:16 (describing Figure 1, the method for quality of service based
24 resource allocation and utilization in a dynamic wireless network).

⁴⁰ *Id.* at 14:42-53.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24

142. Claim 1 of the '992 Patent reads as follows:

1. In a portable system, a method for providing a digital media service to a user, the method comprising:

delivering digital media content having a current quality level to a user;

determining that a network connection with a second system is available and is characterized by a communication bandwidth that is high enough to provide the digital media content to the user at a quality level higher than the current quality level;

using the network connection to obtain the digital media content at the higher quality level from the second system; and

delivering the digital media content at the higher quality level to the user instead of the digital media content at the current quality level.

143. Claim 1 thus recites a novel solution of determining, accessing, and using the resources of another system on a dynamic network environment in order to improve digital media content quality delivered to the user by obtaining the higher quality content from the second system in a manner that was not well-understood, routine, or conventional at the time of the '992 Patent. In the prior art, among other things, the source of the content remained the same between the two systems, whereas claim 1 claims obtaining the digital media content at a higher quality level from another source, namely the second system.

144. Claims 2, 3 and 5 of the '992 Patent depend from claim 1.

145. Claim 2 of the '992 Patent reads as follows:

2. The method of claim 1, where the digital media content is video media.

146. Claim 3 of the '992 Patent reads as follows:

1 3. The method of claim 1, where the digital media content
2 is audio media.

3 147. Claim 5 of the '992 Patent reads as follows:

4 5. The method of claim 1, where the portable system
5 automatically performs, without user interaction, said
6 determining, said using, and said delivering the digital
7 media content at the higher quality level to the user.

8 148. Netflix directly infringes the '992 Patent by making, using, selling, and
9 offering to sell the Netflix service, which practices the patented invention.

10 149. Upon information and belief, Netflix directly infringes at least claims
11 1, 2, 3, and 5 of the '992 Patent, at least in the exemplary manner described below.

12 150. Netflix practices a “method for providing a digital media service to a
13 user” “[i]n a portable system.”

14 151. For example, the Netflix system provides a streaming entertainment
15 service that delivers digital video content such as TV series, documentaries, and
16 feature films to a wide variety of internet-connected devices, including mobile
17 devices such as laptops, tablets, and mobile phones.

18 152. The Netflix system “deliver[s] digital media content having a current
19 quality level to a user.” For example, the Netflix system delivers the “best video
20 quality stream” to its users “tailored to the member’s available bandwidth and view
21 device capability.” To account for variable network conditions, the Netflix
22 streaming service encodes video titles at different bit rates so that video can be
23 delivered at different quality levels. The Netflix streaming service “pre-encode[s]
24 streams at various bitrates applying optimized encoding recipes.”⁴¹

⁴¹ <https://medium.com/netflix-techblog/per-title-encode-optimization-7e99442b62a2>.

1 153. The Netflix system “determin[es] that a network connection with a
2 second system is available and is characterized by a communication bandwidth that
3 is high enough to provide the digital media content to the user at a quality level
4 higher than the current quality level.” For instance, the Netflix streaming service is
5 designed to adapt to changing network conditions so that the content can be
6 delivered and viewed at high levels of quality even when network conditions
7 become constrained. Throughout a user’s playback session, the Netflix streaming
8 service continuously monitors the network to evaluate changing conditions and
9 makes adjustments to the video that is being delivered.

10 154. The client device running the Netflix application, which is being
11 directed and controlled by Netflix code and other instructions, “intelligently selects
12 which OCA to use. It does this by testing the quality of the network connection to
13 each OCA. It will connect to the fastest, most reliable OCA first.” The client
14 device “keeps running these tests throughout the video streaming process.” “If the
15 network quality declines, the client lowers video quality to match. The client will
16 switch to another OCA when the quality declines too much.” The Netflix streaming
17 service uses the user’s IP address and information from internet Service Providers
18 to adapt to the quality of the network. The Netflix streaming service also identifies
19 which Open Connect Appliance (OCA) clusters are best for the user’s client to use.
20 It selects which OCA to use by testing the quality of the network connection to
21 each OCA and will connect to the fastest, most reliable OCA first. This process is
22 repeated throughout the user’s video streaming experience.⁴²

23 _____
24 ⁴² <http://highscalability.com/blog/2017/12/11/netflix-what-happens-when-you-press-play.html>.

1 155. Netflix “us[es] the network connection to obtain the digital media
2 content at the higher quality level from the second system” and “deliver[s] the
3 digital media content at the higher quality level to the user instead of the digital
4 media content at the current quality level.” For example, the client device
5 continues to test the quality of the network connection at each OCA throughout the
6 video streaming process. “If the network quality declines, the client lowers video
7 quality to match.” The client will switch to another OCA when the quality declines
8 too much. “The Netflix streaming service adapts to the quality of the network and
9 will adjust the video quality by switching to another OCA when the quality declines
10 too much.”⁴³ After establishing a connection to another OCA capable of streaming
11 content at a higher quality level—for example, streaming video at a higher
12 resolution—the Netflix system will provide that content at the higher quality level.

13 156. With regard to claims 2 and 3, the Netflix system provides “digital
14 media content” in the form of “video media” and “audio media.” For example,
15 during the playback process, the Netflix system streams the video and audio files
16 associated with the title being viewed.

17 157. With regard to claim 5, within the Netflix system, the “portable system
18 automatically performs, without user interaction, said determining, said using, and
19 said delivering the digital media content at the higher quality level to the user.” For
20 example, Netflix’s process of identifying and using the best OCA and bitrate is
21 generally performed by the Netflix system automatically and without user
22 interaction. As Netflix has explained, the Netflix application on the client device
23

24 ⁴³ *Id.*

1 “runs adaptive streaming algorithms which instantaneously select the best encode to
2 maximize video quality while avoiding playback interruptions due to rebuffers.”

3 158. Netflix has infringed, and continues to infringe, at least claims 1, 2, 3,
4 and 5 of the '992 Patent in the United States by making, using, offering for sale,
5 selling, and/or importing the Netflix streaming service, in violation of 35 U.S.C. §
6 271(a).

7 159. At least as of on or around September 26, 2019 when the Broadcom
8 Entities informed Netflix of its infringement of the '992 Patent, and by no later than
9 the date of the original complaint in this matter, Netflix has had knowledge of the
10 '992 Patent and that its video streaming service infringes the '992 Patent.

11 160. In addition to its own direct infringement, Netflix has induced, and
12 continues to induce, direct infringement by third-parties of at least claims 1, 2, 3,
13 and 5 of the '992 Patent, at least in the exemplary manner described above, by
14 actively encouraging their use of the Netflix system, in violation 35 U.S.C. §
15 271(b). For example, Netflix has induced, and continues to induce, direct
16 infringement of the '992 Patent by customers and/or end users of infringing
17 playback devices enabled with the Netflix software application and service. Netflix
18 knows that it provides and specifically intends to provide an application for use on
19 playback devices that, when used as intended with the Netflix streaming service,
20 meets the limitations of claims 1, 2, 3, and 5 of the '992 Patent. Netflix knows and
21 specifically intends that its end users practice the method recited in claims 1, 2, 3,
22 and 5 of the '992 Patent, when using its application and service as intended.

23 161. Netflix's knowing and willful infringement of the '992 Patent has
24 caused and continues to cause damage to Avago. Avago is entitled to recover

1 damages sustained as a result of Netflix’s wrongful acts in an amount subject to
2 proof at trial.

3 **FIFTH CLAIM FOR RELIEF**

4 **(Infringement of U.S. Patent No. 6,341,375)**

5 162. The Broadcom Entities reallege and incorporates by reference the
6 allegations of paragraphs 1-161 set forth above.

7 163. The ’375 patent, entitled “Video on demand DVD system” was duly
8 and legally issued on January 22, 2002, from a patent application filed on July 14,
9 1999, with Daniel Watkins as the named inventor. A copy of the ’375 Patent is
10 attached hereto as **Exhibit E**.

11 164. The ’375 Patent is assigned to Broadcom Corp., which holds all
12 substantial rights, title, and interest in and to the ’375 Patent.

13 165. Pursuant to 35 U.S.C. § 282, the ’375 Patent is presumed valid.

14 166. The ’375 Patent claims are directed to a new method for distributing
15 video. The ’375 Patent describes systems and methods in which the video is
16 delivered to the user on-demand using a drive server, a control server, and one or
17 more decoder devices.⁴⁴ The system can process multiple compressed video
18 streams in response to multiple user requests for video content.⁴⁵

19 167. The ’375 Patent includes embodiments described in the specification
20 referring to delivery of compressed video data originating from DVD
21 media. However, the systems and methods are not limited to compressed video
22 from a particular source and, as the patent explains, a person having ordinary skill
23

24 ⁴⁴ ’375 Patent at 1:56-60.

⁴⁵ *Id.* at 1:60-63.

1 in the art would recognize that the systems and methods described and claimed in
2 the '375 Patent can be applied to other video distribution models.⁴⁶ The Netflix
3 streaming service is an example of an application of the systems and methods
4 described in the claims.

5 168. The inventions of the '375 Patent resolve technical problems related to
6 conventional video-on-demand systems that require the use of physical connections
7 and short distances between the sources of video and the decoders or players on
8 which the end-user views video content.⁴⁷ In prior art systems, each user has a
9 dedicated video system, such as a DVD player, and decoder at the user's location.⁴⁸

10 169. The '375 Patent, therefore, addresses the technical problem of ensuring
11 delivery of compressed video content to multiple remote end user locations.⁴⁹

12 170. The '375 Patent claims specific ways to solve these technical problems
13 with a video on demand system that is centrally managed and implemented by a
14 drive server, a control server, and one or more decoder devices. Each of these
15 servers can process one or more compressed video streams in response to one or
16 more request signals initiated by a user requesting a video.

17 171. Claim 15 claims an improved method of distributing video:

18 A method for distributing video comprising the steps of:

19 (A) presenting a plurality of compressed data streams
20 with a drive server to a control server in response to one
or more first control signals;

21 (B) distributing said one or more compressed data streams
22 received from said drive server with said control server to

23 ⁴⁶ *Id.* at 5:38-43.

⁴⁷ *Id.* at 1:14-41.

24 ⁴⁸ *Id.* at 1:26-27, 33-35.

⁴⁹ *See id.* at 1:56-2:8.

1 one or more decoder devices in response to one or more
2 request signals;

3 (C) decoding at least one of said one or more compressed
4 data streams with said one or more decoders in response
5 to receiving said one or more compressed data streams
6 from said control server; and

7 (D) presenting at least one signal selected from a
8 decoded video signal and a decoded audio signal in
9 response to decoding said at least one of said one or more
10 compressed data streams, wherein at least one of said one
11 or more decoders is disposed in a separate room from said
12 control server and said driver server, wherein a first
13 portion of a selected one of said compressed data streams
14 is presented to one of said decoder devices and a second
15 portion of said selected compressed data stream is
16 presented to another of said decoder devices.

17 172. Claim 15 thus recites a novel solution involving a drive server
18 presenting multiple compressed video streams and delivering those streams to
19 multiple decoder devices in a remote location in a manner that was not well-
20 understood, routine, or conventional at the time of the '375 Patent, resulting in a
21 better video on demand system.

22 173. Netflix directly infringed the '375 Patent by making, using, offering to
23 sell, and/or selling in the United States its Netflix service, which utilizes the
24 inventions claimed in the '375 Patent to deliver streaming video content. Although
the '375 Patent is presently expired, Netflix infringed the '375 Patent prior to its
expiration as described below. Broadcom Corp. thus is entitled to damages for
Netflix's unauthorized use of the inventions described in the '375 Patent prior to its
expiration.

1 174. Upon information and belief, the Netflix streaming service directly
2 infringed at least claim 15 of the '375 Patent, at least in the exemplary manner
3 described below.


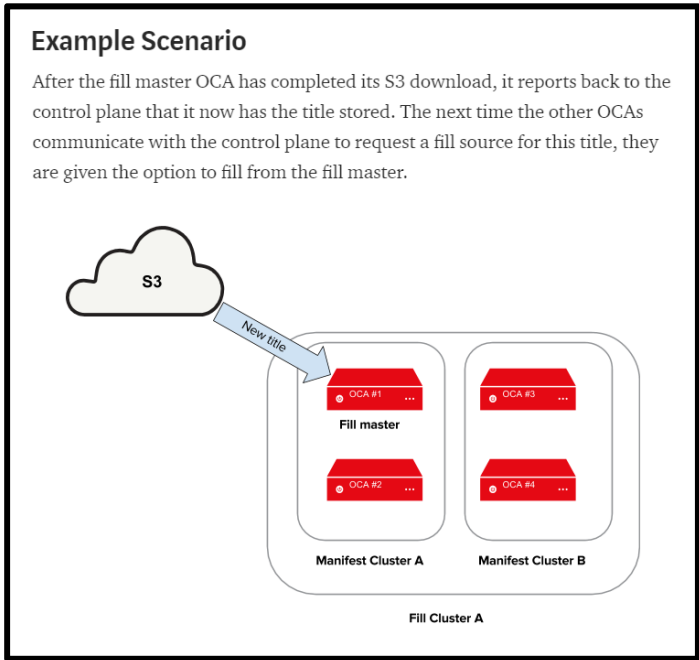
4 175. Netflix practices a “method for distributing video” and “presenting a
5 plurality of compressed data streams with a drive server to a control server in
6 response to one or more first control signals.” The Netflix streaming service is a
7 streaming entertainment service that delivers video content using the Netflix CDN,
8 OCAs, and S3 servers. For example, the Netflix video titles (“data streams”) are
9 presented from an S3 server or an OCA (“drive server”) and sent to other OCAs
10 (“control server”), which store and serve video content.

11 176. The Netflix system presents the compressed data streams to OCAs in
12 response to one or more control signals. For example, as Netflix describes, “OCAs
13 communicate at regular intervals with the control plane services, requesting (among
14 other things) a manifest file that contains the list of titles they should be storing and
15 serving to members. If there is a delta between the list of titles in the manifest file
16 and what they are currently storing, each OCA will send a request, during its
17 configured fill window, that includes a list of the new or update titles that it needs.
18 The response from the control plan in AWS is a ranked list of potential download
19 locations, a.k.a. *fill sources*, for each title.”⁵⁰

20
21
22
23
24

⁵⁰ <https://netflixtechblog.com/netflix-and-fill-c43a32b490c0>

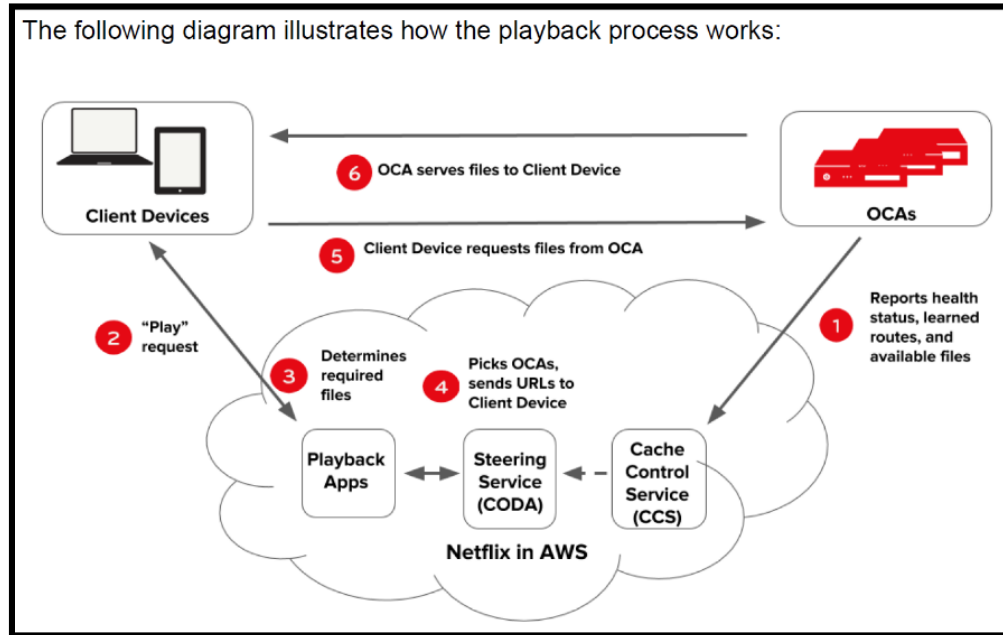
1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24



Source: <https://medium.com/netflix-techblog/netflix-and-fill-c43a32b490c0>.

177. Netflix performs the step of “distributing said one or more compressed data streams received from said drive server with said control server to one or more decoder devices in response to one or more request signals.” For example, in response to a “Play” request, the Netflix CDN delivers content to Netflix client applications (“Netflix Playback Apps”), which are installed on Netflix’s customers’ client devices, including digital televisions, desktop computers, laptop computers,

1 tablet computers, and mobile phones (“decoder devices”). The Netflix CDN is
 2 illustrated in the following diagram:



12 Source: <https://openconnect.netflix.com/Open-Connect-Overview.pdf>.

13 178. Netflix practices “decoding at least one of said one or more
 14 compressed data streams with said one or more decoders in response to receiving
 15 said one or more compressed data streams from said control server.” For instance,
 16 upon receipt of the compressed video, audio, and other content from the Netflix
 17 CDN at the client device, the content is decoded so that it can be viewed. On
 18 information and belief, the Netflix Playback Apps decode, or cause the decoding of,
 19 the content. For instance, in Netflix’s own words, “[t]he app uses several
 20 approaches for video playback on Android such as hardware decoder, software
 21 decoder, OMX-AL, iOMX”).

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24

Device Diversity

To put device diversity in context, we see almost around 1000 different devices streaming Netflix on Android every day. We had to figure out how to categorize these devices in buckets so that we can be reasonably sure that we are releasing something that will work properly on these devices. So the devices we choose to participate in our continuous integration system are based on the following criteria.

- We have at least one device for each playback pipeline architecture we support (The app uses several approaches for video playback on Android such as hardware decoder, software decoder, OMX-AL, iOMX).
- We choose devices with high and low end processors as well as devices with different memory capabilities.
- We have representatives that support each major operating system by make in addition to supporting custom ROMs (most notably CM7, CM9).
- We choose devices that are most heavily used by Netflix Subscribers.

Source: <http://techblog.netflix.com/2012/03/testing-netflix-on-android.html>.

179. On information and belief, Netflix directs and controls the playback process associated with the Netflix service from the client side through its control of the Netflix Playback Apps. Netflix develops its own Android and iOS Netflix Playback Apps for Android and Apple devices. Netflix also develops its own software development kit (SDK) to control third party development of Netflix applications on platforms like Smart TVs.

Netflix Controls The Client

Netflix handles failures gracefully because it controls the client on every device running Netflix.

Netflix develops its Android and iOS apps themselves, so you might expect them to control those. But even on platforms like Smart TVs, where Netflix doesn't build the client, Netflix still has control because it controls the *software development kit* (SDK).

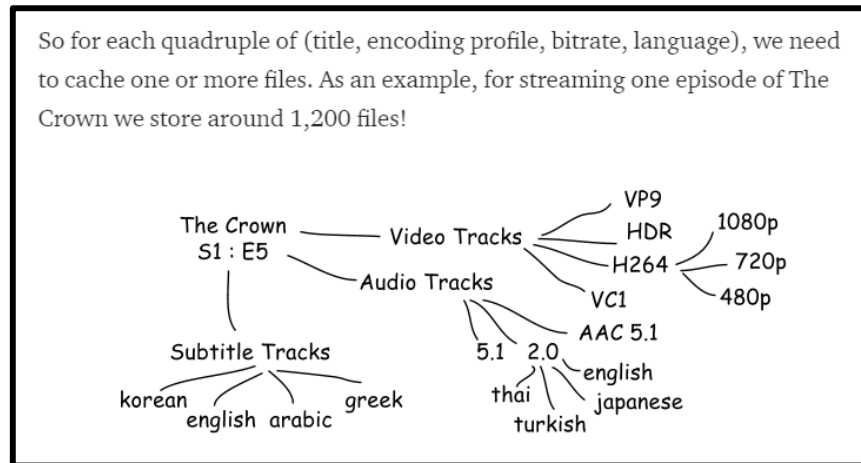
A SDK is a *set of software development tools that allows the creation of applications*. Every Netflix app makes requests to AWS and plays video using the SDK.

By controlling the SDK, Netflix can adapt consistently and transparently to slow networks, failed OCAs, and any other problems that might arise.

Source: <http://highscalability.com/blog/2017/12/11/netflix-what-happens-when-you-press-play.html>.

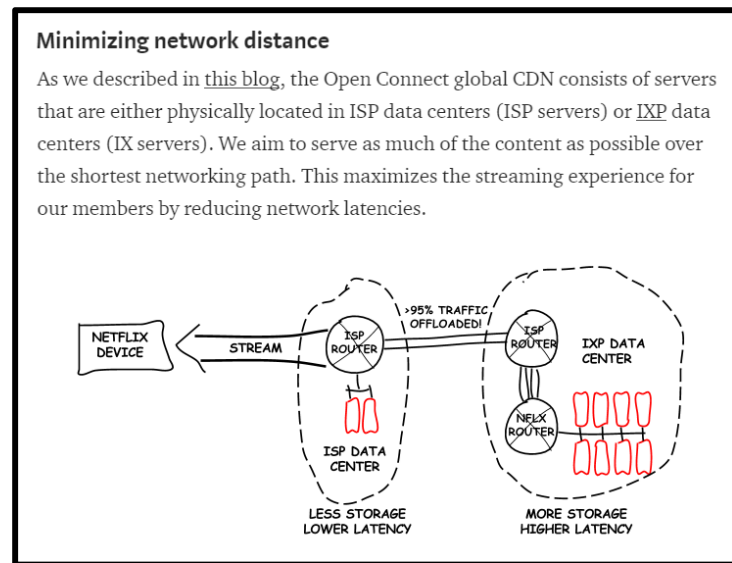
1 180. Netflix practices “presenting at least one signal selected from a
 2 decoded video signal and a decoded audio signal in response to decoding said at
 3 least one of said one or more compressed data streams, wherein at least one of said
 4 one or more decoders is disposed in a separate room from said control server and
 5 said driver server, wherein a first portion of a selected one of said compressed data
 6 streams is presented to one of said decoder devices and a second portion of said
 7 selected compressed data stream is presented to another of said decoder devices.”

8 181. For example, a Netflix client—under the control of the Netflix
 9 Playback App—presents “at least one signal selected from a decoded video signal
 10 and a decoded audio signal in response to decoding said at least one of said one or
 11 more compressed data streams” by playing the decoded video, audio, and other
 12 Netflix content on the client device. As illustrated below, Netflix content consists
 13 of video and audio data, in addition to other data types.



17 Source: [https://medium.com/netflix-techblog/content-popularity-for-open-](https://medium.com/netflix-techblog/content-popularity-for-open-connect-b86d56f613b)
 18 [connect-b86d56f613b](https://medium.com/netflix-techblog/content-popularity-for-open-connect-b86d56f613b).
 19
 20
 21
 22


1 182. Further, in the Netflix system, “at least one of said one or more
 2 decoders is disposed in a separate room from said control server and said drive
 3 server.” For example, the Netflix CDN processes video requests and audio requests
 4 using OCAs that Netflix strategically positions all over the world based on the
 5 location of Netflix’s users. The OCAs are located within internet exchange points
 6 in significant Netflix markets and are interconnected with internet service
 7 providers.⁵¹ Similarly, the S3 is located on one or more servers in data centers that
 8 are, generally speaking, far removed from the users streaming Netflix content. In
 9 contrast, Netflix users can stream Netflix content virtually anywhere with an
 10 internet connection in numerous countries around the world, including the United
 11 States. Thus, the “one or more decoders” (in the client devices) “is disposed in a
 12 separate room” from the “control server” (the OCA) and the “drive server” (the S3
 13 or OCA).



Source: <https://medium.com/netflix-techblog/content-popularity-for-open-connect-b86d56f613b>.

⁵¹ <https://openconnect.netflix.com/Open-Connect-Overview.pdf>.

1 183. Finally, within the Netflix system, “a first portion of a selected one of
 2 said compressed data streams is presented to one of said decoder devices and a
 3 second portion of said selected compressed data stream is presented to another of
 4 said decoder devices.” For example, on information and belief, when video, audio,
 5 and other data associated with Netflix content is streamed from an OCA to the
 6 numerous client devices obtaining Netflix content from that OCA, the data travels
 7 through the internet infrastructure in a compressed data stream. The relevant
 8 portions of that stream are then routed and presented to the appropriate client
 9 device.



10
 11
 12
 13
 14
 15 Source: <https://openconnect.netflix.com/Open-Connect-Overview.pdf>

16 (arrows added).

17 184. Netflix directly infringed claim 15, at least as described, when the
 18 Netflix Playback Apps, Netflix CDN, and Netflix OCAs were used to stream video
 19 content to multiple users in remote and different locations.

20 185. Netflix’s infringement of the ’375 Patent caused damage to Broadcom
 21 Corp., and Broadcom Corp. is entitled to recover damages sustained as a result of
 22 Netflix’s wrongful acts in an amount subject to proof at trial.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24

SIXTH CLAIM FOR RELIEF

(Infringement of U.S. Patent No. 8,572,138)

186. The Broadcom Entities reallege and incorporate by reference the allegations of paragraphs 1-185 set forth above.

187. The '138 Patent, entitled “Distributed Computing System Having Autonomic Deployment of Virtual Machine Disk Images,” was duly and legally issued on October 29, 2013, from a patent application filed March 30, 2007, with Jagane Sundar, Sanjay Radia, and David A. Henseler as the named inventors. A copy of the '138 Patent is attached hereto as **Exhibit F**.

188. The '138 Patent is assigned to Avago, which holds all substantial rights, title, and interest in and to the '138 Patent.

189. Pursuant to 35 U.S.C. § 282, the '138 Patent is presumed valid.

190. The '138 Patent is directed to an improvement in the functionality of a complex distributed computing system. Specifically, the '138 Patent claims a new method and system for a distributed computing environment that conforms to a multi-level, hierarchical organizational model.

191. Traditional distributed computing systems faced the significant challenge of providing an organizational structure that could handle the deployment and administration of thousands of virtual computing resources that could carry out millions of operations simultaneously. As the '138 Patent explains, an enterprise environment—such as a large business organization—often includes several business groups, and each group may have competing and variable computing requirements that necessitate separate, independent computing devices connected to each other on the network. However, the diversity of competing and variable

1 computing requirements increases the cost of distributed computing systems
2 because of the increased time and expense associated with the management of
3 resources that need to be customized to the unique computing requirements of each
4 business group.⁵²

5 192. The '138 invention solved the technical problems presented by such
6 traditional distributed computing systems by developing an infrastructure
7 management facility ("IMF") that guarantees reliable and efficient application
8 service delivery independent of the computational infrastructure. The IMF includes
9 the implementation of virtual machine managers, or "container services," capable
10 of managing other container services and virtual machines ("VMs"). The virtual
11 machines, managed by the VM managers, then appear on the network as available
12 resources as if they were independent computing resources that can be accessed by
13 various groups and utilized to suit their highly-diverse and specialized computing
14 needs.⁵³

15 193. Claim 1 recites an improved method of distributing software "images"
16 (i.e. computer programs) via a number of virtual machines to application nodes:

17 A distributed computing system comprising:

18 a software image repository comprising non-transitory,
19 computer-readable media operable to store: (i) a plurality
20 of image instances of a virtual machine manager that is
21 executable on a plurality of application nodes, wherein
22 when executed on the applications nodes, the image
23 instances of the virtual machine manager provide a
24 plurality of virtual machines, each of the plurality of
virtual machine operable to provide an environment that
emulates a computer platform, and (ii) a plurality of

⁵² '138 Patent at 1:16-33.

⁵³ *Id.* at 8:32-36, 9:65-67, 32:65-33:6.

1 image instances of a plurality of software applications that
2 are executable on the plurality of virtual machines; and
3 a control node that comprises an automation infrastructure
4 to provide autonomic deployment of the plurality of
5 image instances of the virtual machine manager on the
6 application nodes by causing the plurality of image
7 instances of the virtual machine manager to be copied
8 from the software image repository to the application
9 nodes and to provide autonomic deployment of the
10 plurality of image instances of the software applications
11 on the virtual machines by causing the plurality of image
12 instances of the software applications to be copied from
13 the software image repository to the application nodes.

14 194. Netflix directly infringes the '138 Patent by making, using, offering to
15 sell, and/or selling in the United States its Netflix service, which utilizes Netflix's
16 Titus Container Management Platform to deploy and manage virtual computing
17 resources in a manner that practices the inventions claimed in the '138 Patent.

18 195. Netflix developed the Titus Container Management Platform internally
19 and uses it in production to “power Netflix streaming, recommendations, and
20 content systems.”⁵⁴

21 196. Using the Titus Container Management Platform, Netflix is able to
22 more efficiently manage the deployment and administration of the hundreds or
23 thousands of VMs necessary for it to provide reliable, high-quality streaming
24 services to its millions of customers. For example, in 2018 Netflix reportedly
launched approximately thousands of VM managers and hundreds of thousands of
VM containers each day.

⁵⁴ <https://netflix.github.io/titus/overview/>.

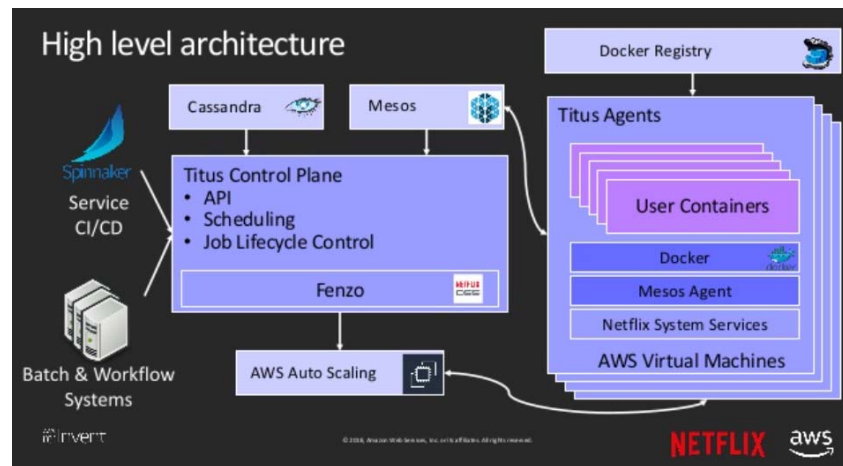
1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24

Q4 2018 Container Usage	
Common	
Jobs Launched	255K jobs / day
Different applications	1K+ different images
Isolated Titus deployments	7
Services	
Single App Cluster Size	5K (real), 12K containers (benchmark)
Hosts managed	7K VMs (435,000 CPUs)
Batch	
Containers launched	450K / day (750K / day peak)
Hosts managed (autoscaled)	55K / month

Source: https://www.slideshare.net/aspyster/qconsl18-disenchantment-netflix-titus-its-feisty-team-and-daemons?next_slideshow=1.

197. Upon information and belief, Netflix directly infringes at least independent claim 1 of the '138 Patent at least in the exemplary manner described below.

198. Netflix utilizes a hierarchical “distributed computing system,” as taught by the '138 Patent, to provide a streaming entertainment service that delivers video content and other content—such as a customer’s browsing experience, content recommendations, and payment information—over a wide geographic area through a distributed network of virtual machines. The Titus architecture is illustrated in the following diagram and description:



Titus schedules application containers to be run across a fleet of thousands of Amazon EC2 instances.

Source: <https://www.slideshare.net/aspkyer/cmp376-another-week-another-million-containers-on-amazon-ec2>; <https://medium.com/netflix-techblog/auto-scaling-production-services-on-titus-1f3cd49f5cd7>.

199. The Netflix system uses the Netflix Titus Container Management Platform (“Titus”) to manage the distributed computing environment that powers its streaming, recommendation, and content delivery systems.⁵⁵ The Titus system integrates a “software image repository,” such as Docker Registry, to store various operating system and application files (i.e., “image instances”) needed to initialize new instances of virtual machine managers, which Netflix refers to as “Titus Agents.”

⁵⁵ <https://netflix.github.io/titus>.

1 200. The Titus Agent image instances developed by Netflix are executed on
2 the “application nodes”—i.e., “cloud” servers.


3 201. For example, as Netflix explains, the Titus Agents set up and manage
4 “virtual machine” containers that are able to independently carry out discrete
5 computing tasks to ensure reliable and efficient delivery of Netflix’s streaming
6 services.⁵⁶ As in the ’138 Patent, each of these virtual machine containers emulates
7 a computer platform and is able to execute various software applications under the
8 supervision of the Titus Agents.

9 202. The Netflix video distribution system also includes “a control node
10 that comprises an automation infrastructure to provide autonomic deployment of
11 the plurality of image instances of the virtual machine manager on the application
12 nodes by causing the plurality of image instances of the virtual machine manager to
13 be copied from the software image repository to the application nodes and to
14 provide autonomic deployment of the plurality of image instances of the software
15 applications on the virtual machines by causing the plurality of image instances of
16 the software applications to be copied from the software image repository to the
17 application nodes.”

18 203. For example, the Netflix Titus system employs an automation
19 infrastructure known as “Titus Master” to be the “control node” for Netflix’s
20 distributed computing system. Like the control node of the ’138 invention, the
21 Titus Master provides “autonomic deployment” for new instances of the Titus
22 Agents (i.e., the “virtual machine managers”) on the “cloud.” The Titus Master is
23 responsible for persisting job and task information, scheduling tasks, and managing

24 _____
⁵⁶ *Id.*

1 the pool of Titus Agents and can scale the pool of Agents up or down in response to
 2 demand.⁵⁷ In doing so, the Titus Master causes the appropriate “image instances”
 3 (i.e. operating system and software application packages) to be copied from the
 4 software image repository to the Titus Agents. The Titus Master is illustrated in the
 5 following diagram and description:



Titus Master

The Titus Master is responsible for persisting job and task information, scheduling tasks, and managing the pool of EC2 Agents. The Master receives requests from Gateway instances and creates and persists job and task info in response. The Master schedules tasks onto Agents with available resources and scales the pool of Titus Agents up or down in response to demand.

Source: <https://medium.com/netflix-techblog/auto-scaling-production-services-on-titus-1f3cd49f5cd7>; <https://netflix.github.io/titus/overview/>.

204. At least as of on or around September 26, 2019, when the Broadcom Entities provided Netflix with an exemplary infringement chart for the '138 Patent, and by no later than the date of the original complaint in this matter, Netflix has had knowledge of the '138 Patent and its infringement thereof.

⁵⁷ <https://medium.com/netflix-techblog/the-evolutionof-container-usage-at-netflix-3abfc096781b>;
<https://medium.com/netflix-techblog/titusthe-netflix-container-management-platformis-now-open-source-f868c9fb5436>.

1 205. Netflix’s knowing and willful infringement of the ’138 Patent has
2 caused and continues to cause damage to Avago. Avago is entitled to recover
3 damages sustained as a result of Netflix’s wrongful acts in an amount subject to
4 proof at trial.

5 **SEVENTH CLAIM FOR RELIEF**

6 **(Infringement of U.S. Patent No. 6,744,387)**

7 206. The Broadcom Entities reallege and incorporate by reference the
8 allegations of paragraphs 1-205 set forth above.

9 207. The ’387 Patent, entitled “Method and System for Symbol
10 Binarization,” was duly and legally issued on June 1, 2004 from a patent
11 application filed on July 10, 2002, with Lowell Winger as the named inventor. A
12 copy of the ’387 Patent is attached hereto as **Exhibit G**.

13 208. The ’387 Patent is assigned to Broadcom Corp., which holds all
14 substantial rights, title, and interest in and to the ’387 Patent.

15 209. Pursuant to 35 U.S.C. § 282, the ’387 Patent is presumed valid.

16 210. The ’387 Patent generally concerns an improvement in the way a
17 computer system compresses visual and audio data. Specifically, the patent is
18 “directed to an improved method for the binarization of data in an MPEG data
19 stream.”⁵⁸

20 211. As the patent explains, MPEG refers to a family of international
21 standards developed by the Motion Picture Expert Group that specify how to
22 represent visual and audio information in a compressed digital format.⁵⁹ The
23

24 ⁵⁸ ’387 Patent, Abstract; *see also id.* at 1:7-11.

⁵⁹ *Id.* at 1:15-30.

1 MPEG formats make it possible to “represent[] a video signal with data roughly
 2 1/50th the size of the original uncompressed video, while still maintaining good
 3 visual quality.”⁶⁰ The MPEG formats achieve such high compression by taking
 4 advantage of the fact that many images in a video stream do not change
 5 significantly from picture to picture, and if they do change, the differences from one
 6 picture to the next are often simple.⁶¹ Storing and transmitting only the changes,
 7 instead of entire pictures, results in considerable savings in data transmission.⁶²

8 212. In practice, this compression technique is accomplished in a number of
 9 steps. First, pixel differences between the pictures are “transformed into frequency
 10 coefficients, and then quantized to further reduce the data transmission.”⁶³ The
 11 ’387 Patent refers to the resulting transformed and quantized coefficients as
 12 “symbols.”⁶⁴ Second, the transformed-quantized symbols are “binarized” to create
 13 binary representations of each symbol in the form of a “codeword.”⁶⁵ Third, the
 14 binarized codewords are entropy encoded “to reduce the number of bits per symbol
 15 without introducing any additional video signal distortion.”⁶⁶ The patent explains
 16 that several types of codecs⁶⁷ exist for performing the entropy encoding; “[o]ne of
 17 the most efficient of which is the family of binary arithmetic encoders (BACs).”⁶⁸

18
 19
 20 ⁶⁰ *Id.* at 1:36-39.

⁶¹ *Id.* at 3:29-37.

⁶² *Id.* at 1:51-54, 3:34-37.

⁶³ *Id.* at 1:55-58, 3:40-43.

⁶⁴ *See, e.g., id.* at 1:55-64, 4:1-4, 4:45-54, 5:10-46, Table 1.

⁶⁵ *Id.* at 4:1-4.

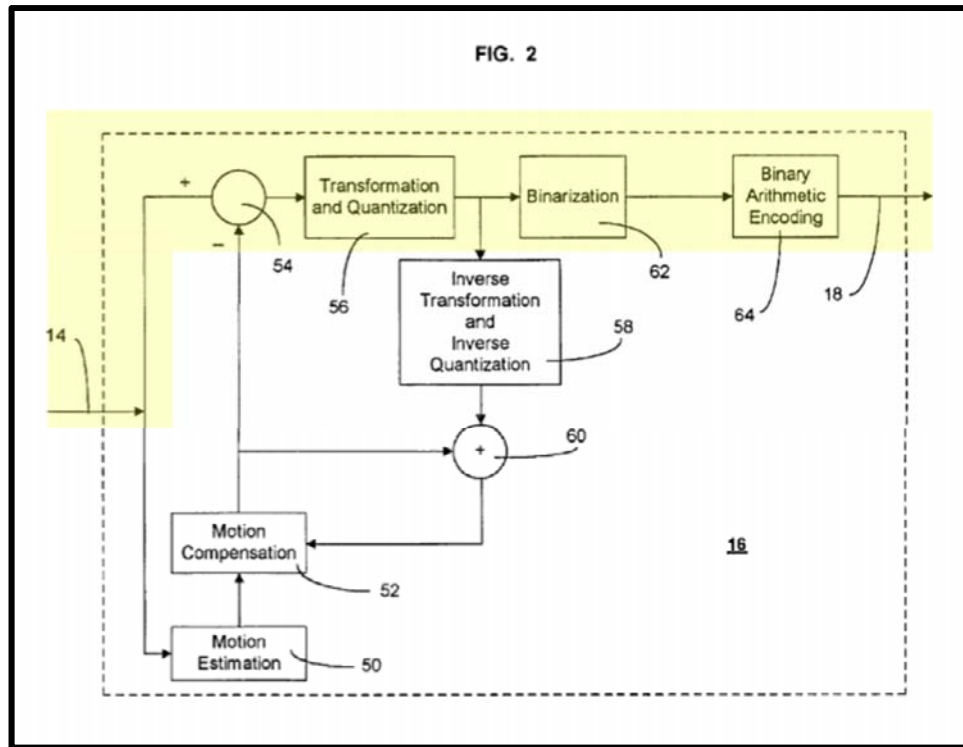
⁶⁶ *Id.* at 1:59-63, 4:26-36.

⁶⁷ A codec is a device or computer program for encoding or decoding a digital stream or signal.

⁶⁸ ’387 Patent, 4:37-39.

1 As the name implies, BACs operate only on binary valued data, which is why the
 2 symbols must be binarized before they can be entropy encoded.⁶⁹

3 213. Figure 2 of the '387 Patent—the relevant portion of which is
 4 highlighted below—is a block diagram of an encoder that carries out this encoding



16 process.

17 Source: '387 Patent, Fig. 2 (highlighting added).

18 214. Figure 2 depicts that a source video stream (14) enters the encoder (16)
 19 at the left and passes to a “combination” module (54), which assembles data related
 20 to pixel differences between pictures in the video stream.⁷⁰ The output of the
 21 combination module passes to the next module (56), where it is transformed and
 22 quantized.⁷¹ Symbols created by module 56 pass to the binarization module (62),

23 ⁶⁹ *Id.* at 4:45-49.

24 ⁷⁰ *Id.* at 3:49-53.

⁷¹ *Id.* at 3:54-57.

1 which creates binary codewords that represent the symbols. The codewords next
 2 pass to binary arithmetic encoding module (64), where they are entropy encoded.
 3 And finally, the encoded bitstream (18) exits the encoder at the right.⁷²

4 215. The '387 Patent is directed to the step in this process that occurs in the
 5 binarization module 62.⁷³ In the prior art, several techniques were available for
 6 binarizing the symbols, including, for example: unary, binary, Golomb, and exp-
 7 Golomb binarization.⁷⁴ The patent explains that those techniques each have certain
 8 strengths and weaknesses. Unary binarization, for example, generates codewords
 9 that are more easily distinguishable from one another but that can be exceptionally
 10 long. Specifically, “[u]nary binarization consists of a number of binary 1s equal to
 11 an index for a symbol followed by a zero....”⁷⁵ Thus, a symbol index value of “1”
 12 results in a 2-bit codeword, namely “10.” A symbol index value of “2” results in
 13 the 3-bit codeword “110,” and so on. Thus, each codeword is easily distinguishable
 14 from the others, but the number of binary values can be quite large—encoding a
 15 large symbol index may require tens of thousands of bits.⁷⁶

16 216. Exp-Golomb binarization, on the other hand, greatly reduces the
 17 maximum possible size of the codewords, but “it does not permit codewords with a
 18 small symbol index (other than index 0) to be uniquely distinguished from
 19 codewords with larger symbol indices.”⁷⁷

20
 21
 22 ⁷² *Id.* at 3:57-4:8.

⁷³ *Id.* at 4:1-2.

23 ⁷⁴ *Id.* at 4:50-51.

⁷⁵ *Id.* at 4:52-54.

24 ⁷⁶ *Id.* at 5:22-5:45.

⁷⁷ *Id.* at 6:19-24.

1 217. Thus, there was a need for a method and system that could exploit the
2 most valuable properties of the unary and exp-Golomb binarizations.⁷⁸ The
3 invention described in the '387 Patent meets that need:

4 The present invention provides a binarization that retains
5 the most valuable properties of the unary and exp-Golomb
6 binarizations. That is, small codewords are
7 distinguishable as with a unary binarization, while large
8 codewords have their binarization limited to a reasonable
9 length. By doing so, the present invention provides a
binarization that reduces the complexity and the
bitrate/size for compressing and decompressing video,
images, and signals that are compressed using binary
arithmetic encoding for entropy encoding.⁷⁹

10 218. The '387 Patent claims methods and systems for constructing
11 binarized codewords for digital video data (i.e., encoding) based on the index
12 values of symbols produced by the transformation-and-quantization module of an
13 encoder. Independent claim 3, for example, recites:

14 A binarization system comprising:

15 means for determining if a code symbol index value is
16 less than a threshold value

17 means for constructing a codeword using a unary
18 binarization if said code symbol index value is less than
19 said threshold value; and

20 means for constructing a codeword using a exp-Golomb
21 binarization if said code symbol index value is not less
22 than said threshold value.

23 219. The methods and systems described in the '387 Patent improve the
24 functionality of computer systems by improving the way they compress and process
video and audio data.

⁷⁸ *Id.* at 2:2-12.

⁷⁹ *Id.* at 6:26-36.

1 220. Notably, the Hon. James V. Selna of this District previously held that
2 the claims of the '387 Patent are patent-eligible under 35 U.S.C. § 101. In doing
3 so, Judge Selna concluded that these claims “do not simply use general computers
4 to perform abstract ideas; instead, the mathematical formula attempts to improve
5 the functioning of compressing and decompressing video, images, and signals.
6 Therefore, an inventive concept sufficiently transforms the nature of the
7 claims...into patent-eligible inventions.”⁸⁰

8 221. Netflix directly infringes at least claim 3 of the '387 Patent at least in
9 the exemplary manner described below.

10 222. Netflix developed, operates, and uses a “video encoding pipeline”, i.e.,
11 a series of video processing applications that operate “in the cloud.” Netflix claims
12 that it has developed, and that it uses, its own proprietary video encoding software
13 in this pipeline.

Pipeline in the Cloud

The video encoding pipeline runs EC2 Linux cloud instances. The elasticity of the cloud enables us to seamlessly scale up when more titles need to be processed, and scale down to free up resources. Our video processing applications don't require any special hardware and can run on a number of EC2 instance types. Long processing jobs are divided into smaller tasks and


Source: <https://medium.com/netflix-techblog/high-quality-video-encoding-at-scale-d159db052746>.

⁸⁰ *Broadcom Corp., et al. v. Sony Corp., et al.*, SAVC 16-1052 JVS (JCGx), p. 12 (C.D.C.A. Oct. 5, 2016).

1 223. Netflix uses its video encoding pipeline to generate encoded video
 2 files in a variety of formats, which it then uses to stream movie and TV content to
 3 its subscribers. As Netflix explains:

4 We ingest high quality video sources and generate video
 5 encodes of various codec profiles, at multiple quality
 6 representations per profile. The encodes are packaged
 7 and then deployed to a content delivery network for
 8 streaming. During a streaming session, the client requests
 9 the encodes it can play and adaptively switches among
 10 quality levels based on network conditions.⁸¹

11 224. Among other encodes, Netflix uses its video encoding pipeline to
 12 generate content files in the ITU-T H.264 format, also known as Advanced Video
 13 Coding or “AVC” (“H.264”), and in the ITU-T H.265 format, also known as High
 14 Efficiency Video Coding or “HEVC” (“H.265”).



15 Source: [https://medium.com/netflix-techblog/high-quality-video-encoding-at-](https://medium.com/netflix-techblog/high-quality-video-encoding-at-scale-d159db052746)
 16 [scale-d159db052746](https://medium.com/netflix-techblog/high-quality-video-encoding-at-scale-d159db052746).

17
 18 225. The Netflix video encoding pipeline includes a “binarization system.”
 19 For example, with regard to the H.264 format used by Netflix, the H.264
 20 documentation explains how the format employs a “concatenated unary/k-th order
 21 Exp-Golomb (UEGk) binarization process.” On information and belief, Netflix
 22
 23

24 ⁸¹ [https://medium.com/netflix-techblog/high-quality-video-encoding-at-scale-](https://medium.com/netflix-techblog/high-quality-video-encoding-at-scale-d159db052746)
[d159db052746](https://medium.com/netflix-techblog/high-quality-video-encoding-at-scale-d159db052746).

1 uses the method for UEG(k) encoding set forth in the H.264.2 reference software,
 2 which serves as an aid for the study and implementation of H.264 video coding.

3 226. On information and belief, the Netflix video encoding pipeline includes
 4 a “means for determining if a code symbol index value is less than a threshold
 5 value.” For instance, the H.264.2 reference software features the functions
 6 unary_exp_golomb_mv_encode() and unary_exp_golomb_level_encode(), which
 7 are responsible for performing unary exp-golomb encoding for various syntax
 8 elements. Both of these functions employ code that determines if a symbol index
 9 value is less than a threshold value.

```

10 static void unary_exp_golomb_mv_encode(EncodingEnvironmentPtr eep_dp,
11                                     unsigned int symbol,
12                                     BiContextTypePtr ctx,
13                                     unsigned int max_bin)
14 {
15     if (symbol==0)
16     {
17         biari_encode_symbol(eep_dp, 0, ctx );
18         return;
19     }
20     else
21     {
22         unsigned int bin = 1;
23         unsigned int l = symbol, k = 1;
24         biari_encode_symbol(eep_dp, 1, ctx++ );
25
26         while (((-l)>0) && (++k <= 8))
27         {
28             biari_encode_symbol(eep_dp, 1, ctx );
29             if ((++bin) == 2)
30                 ++ctx;
31             if (bin == max_bin)
32                 ++ctx;
33         }
34         if (symbol < 8)
35             biari_encode_symbol(eep_dp, 0, ctx);
36         else
37             exp_golomb_encode_eq_prob(eep_dp, symbol - 8, 3);
38     }
39 }

```

Source: unary_exp_golomb_mv_encode() function.

```

1      static void unary_exp_golomb_level_encode( EncodingEnvironmentPtr eep_dp,
2                                               unsigned int symbol,
3                                               BiContextTypePtr ctx)
4      {
5          if (symbol==0)
6          {
7              biari_encode_symbol(eep_dp, 0, ctx );
8              return;
9          }
10         else
11         {
12             unsigned int l=symbol;
13             unsigned int k = 1;
14
15             biari_encode_symbol(eep_dp, 1, ctx );
16             while (((--l)>0) && (++k <= 13))
17                 biari_encode_symbol(eep_dp, 1, ctx);
18             if (symbol < 13)
19                 biari_encode_symbol(eep_dp, 0, ctx);
20             else
21                 exp_golomb_encode_eq_prob(eep_dp, symbol - 13, 0);
22         }
23     }
24

```

Source: unary_exp_golomb_level_encode() function.

227. On information and belief, the Netflix video encoding pipeline includes a “means for constructing a codeword using a unary binarization if said code symbol index value is less than said threshold value.” For example, in the H.264.2 reference software, the functions unary_exp_golomb_mv_encode() and unary_exp_golomb_level_encode() perform unary binarization of various syntax elements when the symbol index value is less than a threshold value (said threshold value is 8 and 13, respectively, in the excerpts below).

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24

```

static void unary_exp_golomb_mv_encode(EncodingEnvironmentPtr eep_dp,
                                     unsigned int symbol,
                                     BiContextTypePtr ctx,
                                     unsigned int max_bin)
{
    if (symbol==0)
    {
        biari_encode_symbol(eep_dp, 0, ctx );
        return;
    }
    else
    {
        unsigned int bin = 1;
        unsigned int l = symbol, k = 1;
        biari_encode_symbol(eep_dp, 1, ctx++ );

        while (((--l)>0) && (++k <= 8))
        {
            biari_encode_symbol(eep_dp, 1, ctx );
            if ((++bin) == 2)
                ++ctx;
            if (bin == max_bin)
                ++ctx;
        }
        if (symbol < 8)
            biari_encode_symbol(eep_dp, 0, ctx);
        else
            exp_golomb_encode_eq_prob(eep_dp, symbol - 8, 3);
    }
}

```

Source: unary_exp_golomb_mv_encode() function.

```

static void unary_exp_golomb_level_encode( EncodingEnvironmentPtr eep_dp,
                                          unsigned int symbol,
                                          BiContextTypePtr ctx)
{
    if (symbol==0)
    {
        biari_encode_symbol(eep_dp, 0, ctx );
        return;
    }
    else
    {
        unsigned int l=symbol;
        unsigned int k = 1;

        biari_encode_symbol(eep_dp, 1, ctx );
        while (((--l)>0) && (++k <= 13))
            biari_encode_symbol(eep_dp, 1, ctx);
        if (symbol < 13)
            biari_encode_symbol(eep_dp, 0, ctx);
        else
            exp_golomb_encode_eq_prob(eep_dp, symbol - 13, 0);
    }
}

```

Source: unary_exp_golomb_level_encode() function.

1 228. On information and belief, the Netflix video encoding pipeline also
 2 includes a “means for constructing a codeword using an exp-Golomb binarization if
 3 said code symbol index value is not less than said threshold value.” For instance,
 4 functions unary_exp_golomb_mv_encode() and unary_exp_golomb_level_encode()
 5 in the H.264.2 reference software both call the function
 6 exp_golomb_encode_eq_prob() when the code symbol value is not less than the
 7 above-described threshold in order to construct the remainder of the codeword
 8 using exp-Golomb binarization.

```

9
10 static void exp_golomb_encode_eq_prob( EncodingEnvironmentPtr eep_dp,
11                                     unsigned int symbol,
12                                     int k)
13 {
14     for(;;)
15     {
16         if (symbol >= (unsigned int)(1<<k))
17         {
18             biari_encode_symbol_eq_prob(eep_dp, 1); //first unary part
19             symbol = symbol - (1<<k);
20             k++;
21         }
22         else
23         {
24             biari_encode_symbol_eq_prob(eep_dp, 0); //now terminated zero of unary part
25             while (k--) //next binary part
26                 biari_encode_symbol_eq_prob(eep_dp, ((symbol>>k)&1));
27             break;
28         }
29     }
30 }
  
```

17 Source: exp_golomb_encode_eq_prob() function.

18 229. On information and belief, Netflix’s video encoding pipeline also
 19 infringes claim 3 of the ’387 Patent through its use of H.265 encoding. On
 20 information and belief, the reference software associated with that format proposes
 21 a binarization process that operates in substantially the same way as described
 22 above with regard to the H.264.2 reference software. On information and belief,
 23 Netflix uses the approach to binarization proposed by the H.265.2 reference
 24 software in encoding content files to the H.265 format.

1 patent is “directed to an improved method for the binarization of data in an MPEG
2 data stream.”⁸²

3 237. As the ’663 Patent explains, MPEG refers to a family of international
4 standards developed by the Motion Picture Expert Group that specify how to
5 represent visual and audio information in a compressed digital format.⁸³ The
6 MPEG formats make it possible to “represent[] a video signal with data roughly
7 1/50th the size of the original uncompressed video, while still maintaining good
8 visual quality.”⁸⁴ The MPEG formats achieve such high compression by taking
9 advantage of the fact that many images in a video stream do not change
10 significantly from picture to picture, and if they do change, the differences from one
11 picture to the next are often simple.⁸⁵ Storing and transmitting only the changes,
12 instead of entire pictures, results in considerable savings in data transmission.⁸⁶

13 238. In practice, this compression technique is accomplished in a number of
14 steps. First, pixel differences between the pictures are “transformed into frequency
15 coefficients, and then quantized to further reduce the data transmission.”⁸⁷ The
16 ’663 Patent refers to the resulting transformed and quantized coefficients as
17 “symbols.”⁸⁸ Second, the transformed-quantized symbols are “binarized” to create
18 binary representations of each symbol in the form of a “codeword.”⁸⁹ Third, the
19 binarized codewords are entropy encoded “to reduce the number of bits per symbol
20

21 ⁸² ’663 Patent, Abstract; *see also id.* at 1:7-11.

22 ⁸³ *Id.* at 1:15-30.

23 ⁸⁴ *Id.* at 1:38-41.

24 ⁸⁵ *Id.* at 3:31-36.

⁸⁶ *Id.* at 1:51-54; 3:35-39.

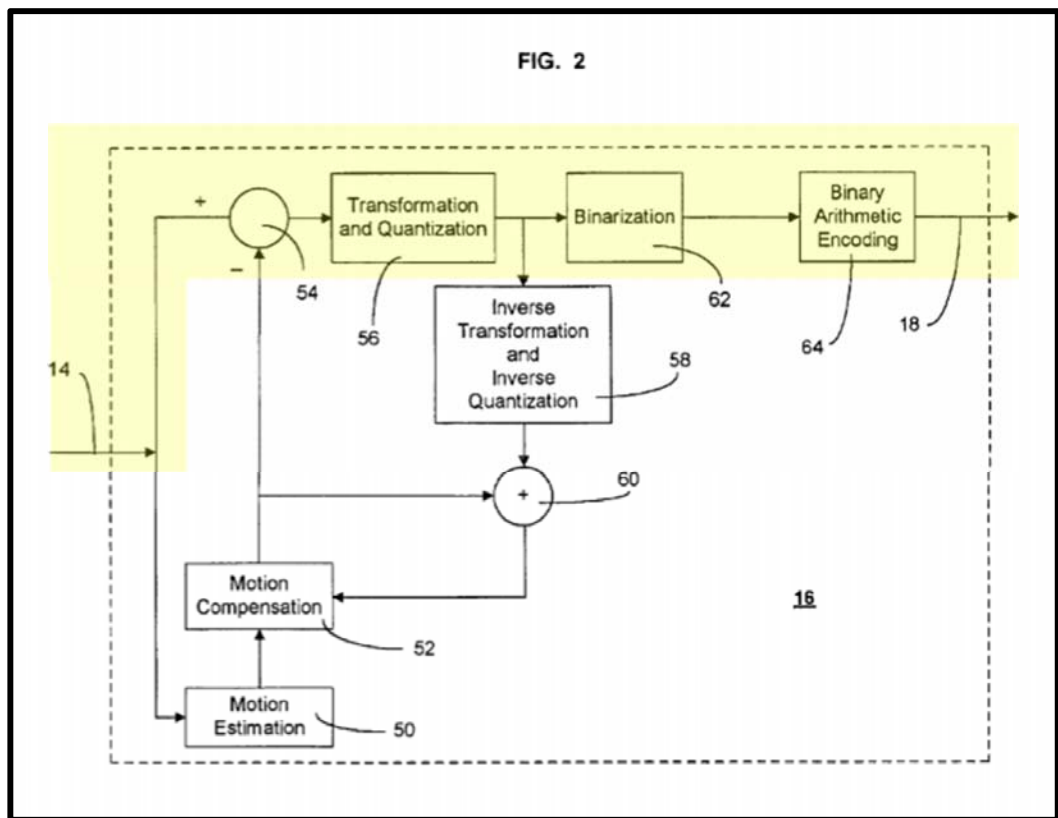
⁸⁷ *Id.* at 1:55-58; 3:41-44.

⁸⁸ *See, e.g., id.* at 1:55-64, 4:1-4, 4:45-54, 5:10-46, Table 1.

⁸⁹ *Id.* at 4:1-4.

1 without introducing any additional video signal distortion.”⁹⁰ The patent explains
 2 that several types of codecs⁹¹ exist for performing the entropy encoding, “[o]ne of
 3 the most efficient of which is the family of binary arithmetic encoders (BACs).”⁹²
 4 As the name implies, BACs operate only on binary valued data, which is why the
 5 symbols must be binarized before they can be entropy encoded.⁹³

6 239. Figure 2 of the ’663 Patent—the relevant portion of which is
 7 highlighted below—is a block diagram of an encoder that carries out this encoding
 8 process.



Source: ’663 Patent, Fig. 2 (highlighting added).

⁹⁰ *Id.* at 1:59-63, 4:26-36.

⁹¹ A codec is a device or computer program for encoding or decoding a digital stream or signal.

⁹² ’663 Patent, 4:34-36.

⁹³ *Id.* at 4:42-47.

1 240. Figure 2 depicts that a source video stream (14) enters the encoder (16)
2 at the left and passes to a “combination” module (54), which assembles data related
3 to pixel differences between pictures in the video stream.⁹⁴ The output of the
4 combination module passes to the next module (56), where it is transformed and
5 quantized. Symbols created by module (56) pass to the binarization module (62),
6 which creates binary codewords that represent the symbols. The codewords next
7 pass to binary arithmetic encoding module (64), where they are entropy encoded.
8 And finally, the encoded bitstream (18) exits the encoder at the right.⁹⁵

9 241. The ’663 Patent is directed, in part, to the step in this process that
10 occurs in the binarization module 62.⁹⁶ In the prior art, several techniques were
11 available for binarizing the symbols, including, for example: unary, binary,
12 Golomb, and exp-Golomb binarization.⁹⁷ The patent explains that those techniques
13 each have certain strengths and weaknesses. Unary binarization, for example,
14 generates codewords that are distinguishable from one another but that can be
15 exceptionally long. Specifically, “[u]nary binarization consists of a number of
16 binary 1s equal to an index for a symbol followed by a zero....”⁹⁸ Thus, a symbol
17 index value of “1” results in a 2-bit codeword, namely “10.” A symbol index value
18 of “2” results in the 3-bit codeword “110,” and so on. Thus, each codeword is
19 easily distinguishable from the others, but the number of binary values can be quite
20 large—encoding a large symbol index may require tens of thousands of bits.⁹⁹

21
22 ⁹⁴ *Id.* at 3:50-58.

⁹⁵ *Id.* at 3:57-4:8.

23 ⁹⁶ *Id.* at 4:1-2.

⁹⁷ *Id.* at 4:46-47.

24 ⁹⁸ *Id.* at 4:48-49.

⁹⁹ *Id.* at 5:22-5:45.

1 242. Exp-Golomb binarization, on the other hand, greatly reduces the
2 maximum possible size of the codewords, but “it does not permit codewords with a
3 small symbol index (other than index 0) to be uniquely distinguished from
4 codewords with larger symbol indices.”¹⁰⁰

5 243. Thus, there was a need for a method and system that could exploit the
6 most valuable properties of the unary and exp-Golomb binarizations.¹⁰¹ The
7 invention described in the ’663 Patent meets that need:

8 The present invention provides a binarization that retains
9 the most valuable properties of the unary and exp-Golomb
10 binarizations. That is, small codewords are
11 distinguishable as with a unary binarization, while large
12 codewords have their binarization limited to a reasonable
13 length. By doing so, the present invention provides a
14 binarization that reduces the complexity and the
15 bitrate/size for compressing and decompressing video,
16 images, and signals that are compressed using binary
17 arithmetic encoding for entropy encoding.¹⁰²

18 244. The ’663 Patent claims methods and systems utilizing a combination
19 of unary and exp-Golomb binarization for encoding and decoding digital video
20 data. Independent claim 12, for example, recites:

21 A method for generating a codeword from an index value
22 for digital video encoding, comprising the steps of:

23 (A) generating a first pattern in a first portion of said
24 codeword in response to said index value being at least as
25 great as a threshold;

 (B) generating a second pattern in a second portion of said
 codeword following said first portion representing an
 offset of said index value above said threshold; and

¹⁰⁰ *Id.* at 6:14-17.

¹⁰¹ *Id.* at 2:1-12.

¹⁰² *Id.* at 6:19-28.

1 (C) generating a third pattern in a third portion of said
2 codeword following said second portion representing a
value of said index value above said offset.

3 245. As in the '387 Patent, the methods and systems described in the '663
4 Patent improve the functionality of computer systems by improving the way they
5 compress and process video and audio data.

6 246. Notably, the Hon. James V. Selna of this District previously held that
7 the claims of the '663 Patent are patent-eligible under 35 U.S.C. § 101. In doing
8 so, Judge Selna concluded that claim 12, and others in the '663 Patent, are “directed
9 to improving digital video decoding” and, thus, “are not directed to abstract
10 ideas.”¹⁰³

11 247. Upon information and belief, Netflix directly infringes at least claim
12 12 of the '663 Patent at least in the exemplary manner described below.

13 248. Figure 5 of the '663 Patent, which depicts Table 3, demonstrates a
14 particular instance of the hybrid unary-exp-Golomb codes described in that patent.
15 As the '663 Patent explains, “Table 3 illustrates a binarization that is particularly
16 appropriate for the binarization of quarter pixel motion vector residual magnitudes
17 of MPEG-AVC/H.264.”¹⁰⁴ The patent describes how, upon reaching the threshold
18 at which unary to exp-Golomb switching occurs (N=64 for Table 3), the index
19 comprises three parts: (1) the initial prefix (highlighted in blue below); (2) a unary
20 representation appended to the initial prefix to form the unary prefix (highlighted in
21 red below); and (3) the exp-Golomb suffix (highlighted in green below).¹⁰⁵

22

23 ¹⁰³ *Broadcom Corp., et al. v. Sony Corp., et al.*, SAVC 16-1052 JVS (JCGx), p. 9
(C.D.C.A. Oct. 5, 2016).

24 ¹⁰⁴ '663 Patent, 6:35-40.

¹⁰⁵ *See id.* at 6:44-63.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24

Table 3 - Motion vector magnitude residual binarization.

Index	Unary Prefix	exp-Golomb Suffix
0	0	
1	10	
2	110	
...		
63	1...10	
64	1...110	0
65	1...110	1
66	1...1110	00
67	1...1110	01
68	1...1110	10
69	1...1110	11
70	1...11110	000
71	1...11110	001
72	1...11110	010
73	1...11110	011
74	1...11110	100
75	1...11110	101
...		

FIG. 5

Source: '663 Patent, Fig. 5.

249. As explained above, Netflix developed, operates, and uses a video encoding pipeline to encode its film and TV content in a variety of digital formats, including H.264 and H.265.

250. The Netflix video encoding system practices a “method for generating a codeword from an index value for digital video encoding.” For instance, with regard to the H.264 format used by Netflix, the H.264 documentation explains how the format employs a “concatenated unary/k-th order Exp-Golomb (UEGk) binarization process.” This process generates codewords from index values. On information and belief, Netflix uses the method for UEG(k) encoding set forth in

1 the H.264.2 reference software, which serves as an aid for the study and
 2 implementation of H.264 video coding.

3 251. On information and belief, the Netflix video encoding pipeline
 4 practices the step of “generating a first pattern in a first portion of said codeword in
 5 response to said index value being at least as great as a threshold.” For instance,
 6 the H.264.2 reference software features the functions
 7 unary_exp_golomb_mv_encode() and unary_exp_golomb_level_encode(), which
 8 are responsible—in part—for generating codewords from index values during the
 9 encoding process. Where the index value meets or exceeds a certain threshold (8
 10 for unary_exp_golomb_mv_encode() and 13 for
 11 unary_exp_golomb_level_encode()), these functions populate the first portion of
 12 the codeword with a pattern representing an initial prefix.

```

13 static void unary_exp_golomb_mv_encode(EncodingEnvironmentPtr eep_dp,
14                                     unsigned int symbol,
15                                     BiContextTypePtr ctx,
16                                     unsigned int max_bin)
17 {
18     if (symbol==0)
19     {
20         biari_encode_symbol(eep_dp, 0, ctx );
21         return;
22     }
23     else
24     {
25         unsigned int bin = 1;
26         unsigned int l = symbol, k = 1;
27         biari_encode_symbol(eep_dp, 1, ctx++ );
28         while (((-1)>0) && (++k <= 8))
29         {
30             biari_encode_symbol(eep_dp, 1, ctx );
31             if ((++bin) == 2)
32                 ++ctx;
33             if (bin == max_bin)
34                 ++ctx;
35         }
36         if (symbol < 8)
37             biari_encode_symbol(eep_dp, 0, ctx);
38         else
39             exp_golomb_encode_eq_prob(eep_dp, symbol - 8, 3);
40     }
41 }
  
```

24 Source: unary_exp_golomb_mv_encode() function.

```

1      static void unary_exp_golomb_level_encode( EncodingEnvironmentPtr eep_dp,
2                                                  unsigned int symbol,
3                                                  BiContextTypePtr ctx)
4      {
5          if (symbol==0)
6          {
7              biari_encode_symbol(eep_dp, 0, ctx );
8              return;
9          }
10         else
11         {
12             unsigned int l=symbol;
13             unsigned int k = 1;
14
15             biari_encode_symbol(eep_dp, 1, ctx );
16             while (((-1)>0) && (++k <= 13))
17                 biari_encode_symbol(eep_dp, 1, ctx);
18             if (symbol < 13)
19                 biari_encode_symbol(eep_dp, 0, ctx);
20             else
21                 exp_golomb_encode_eq_prob(eep_dp,symbol - 13, 0);
22         }
23     }

```

Source: unary_exp_golomb_level_encode() function.

252. On information and belief, the Netflix video encoding system also “generat[es] a second pattern in a second portion of said codeword following said first portion representing an offset of said index value above said threshold.” For example, both the unary_exp_golomb_mv_encode() and unary_exp_golomb_level_encode() functions call the function exp_golomb_encode_eq_prob() when the index value exceeds the thresholds described above. The exp_golomb_encode_eq_prob() function generates a second portion of the codeword—the unary representation appended to the initial prefix to form the unary prefix—when the index value exceeds the thresholds.

```

1  static void exp_golomb_encode_eq_prob( EncodingEnvironmentPtr eep_dp,
2                                     unsigned int symbol,
3                                     int k)
4  {
5      for(;;)
6      {
7          if (symbol >= (unsigned int)(1<<k))
8          {
9              biari_encode_symbol_eq_prob(eep_dp, 1); //first unary part
10             symbol = symbol - (1<<k);
11             k++;
12         }
13         else
14         {
15             biari_encode_symbol_eq_prob(eep_dp, 0); //now terminated zero of unary part
16             while (k--) //next binary part
17                 biari_encode_symbol_eq_prob(eep_dp, ((symbol>>k)&1));
18             break;
19         }
20     }
21 }

```

Source: exp_golomb_encode_eq_prob() function.

253. Further, on information and belief, the Netflix video encoding system “generat[es] a third pattern in a third portion of said codeword following said second portion representing a value of said index value above said offset.” For example, after the unary prefix described above, the exp_golomb_encode_eq_prob() function generates a third portion of the codeword—the exp-Golomb suffix—which captures the value of the index above the offset described above.

```

1      static void exp_golomb_encode_eq_prob( EncodingEnvironmentPtr eep_dp,
2                                          unsigned int symbol,
3                                          int k)
4      {
5          for(;;)
6          {
7              if (symbol >= (unsigned int)(1<<k))
8              {
9                  biari_encode_symbol_eq_prob(eep_dp, 1); //first unary part
10                 symbol = symbol - (1<<k);
11                 k++;
12             }
13             else
14             {
15                 biari_encode_symbol_eq_prob(eep_dp, 0); //now terminated zero of unary part
16                 while (k-->0) //next binary part
17                     biari_encode_symbol_eq_prob(eep_dp, ((symbol>>k)&1));
18                 break;
19             }
20         }
21     }

```

Source: exp_golomb_encode_eq_prob() function.

254. On information and belief, Netflix's video encoding pipeline also infringes claim 12 of the '663 Patent through its use of H.265 encoding. On information and belief, the reference software associated with that format proposes a binarization process that operates in substantially the same way as described above with regard to the H.264.2 reference software. On information and belief, Netflix uses the approach to binarization proposed by the H.265.2 reference software in encoding content files to the H.265 format.

255. At least as of on or around September 26, 2019, when the Broadcom Entities informed Netflix of its infringement of the '663 Patent, and by no later than the filing of the original complaint in this matter, Netflix has had knowledge of the '663 Patent and the infringement thereof by its encoding pipeline.

256. Netflix's knowing and willful infringement of the '663 Patent has caused and continues to cause damage to Broadcom Corp., and Broadcom Corp. is entitled to recover damages sustained as a result of Netflix's wrongful acts in an amount subject to proof at trial.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24

NINTH CLAIM FOR RELIEF

(Infringement of U.S. Patent No. 9,332,283)

257. The Broadcom Entities reallege and incorporate by reference the allegations of paragraphs 1-256 set forth above.

258. The '283 Patent, entitled "Signaling of Prediction Size Unit in Accordance with Video Coding," was duly and legally issued on May 3, 2016 from a patent application filed on June 14, 2012, with Peisong Chen, Brian Heng, and Wade Wan as the named inventors. A copy of the '283 Patent is attached hereto as **Exhibit I.**

259. The '283 Patent claims priority from U.S. Provisional Application No. 60/539,948, filed on September 27, 2011.

260. The '283 Patent is assigned to Broadcom Corp., which holds all substantial rights, title, and interest in and to the '283 Patent.

261. Pursuant to 35 U.S.C. § 282, the '283 Patent is presumed valid.

262. The '283 Patent generally concerns an improved system for encoding and decoding video content that ensures a high quality output when transmitting that content to viewers. Specifically, the patent relates to an improved method for encoding video content using a process known as binarization in order to transmit that information to users more efficiently.

263. At the time of the inventions claimed in the '283 Patent, persons engaged in developing next-generation video encoding technologies were looking to take advantage of ongoing innovations in parallel processing power and increased video resolutions, which could make performing complex encoding operations more efficient and lead to improved user experiences. Video encoding is

1 a multi-step process that begins with an input video signal, and results in an output
2 bitstream of encoded video data. As discussed above, between the input of the
3 video signal and the output of the encoded video data the encoding process includes
4 various operations that are performed on constituent portions of the input video
5 signal to create a video stream according to a particular encoding format.

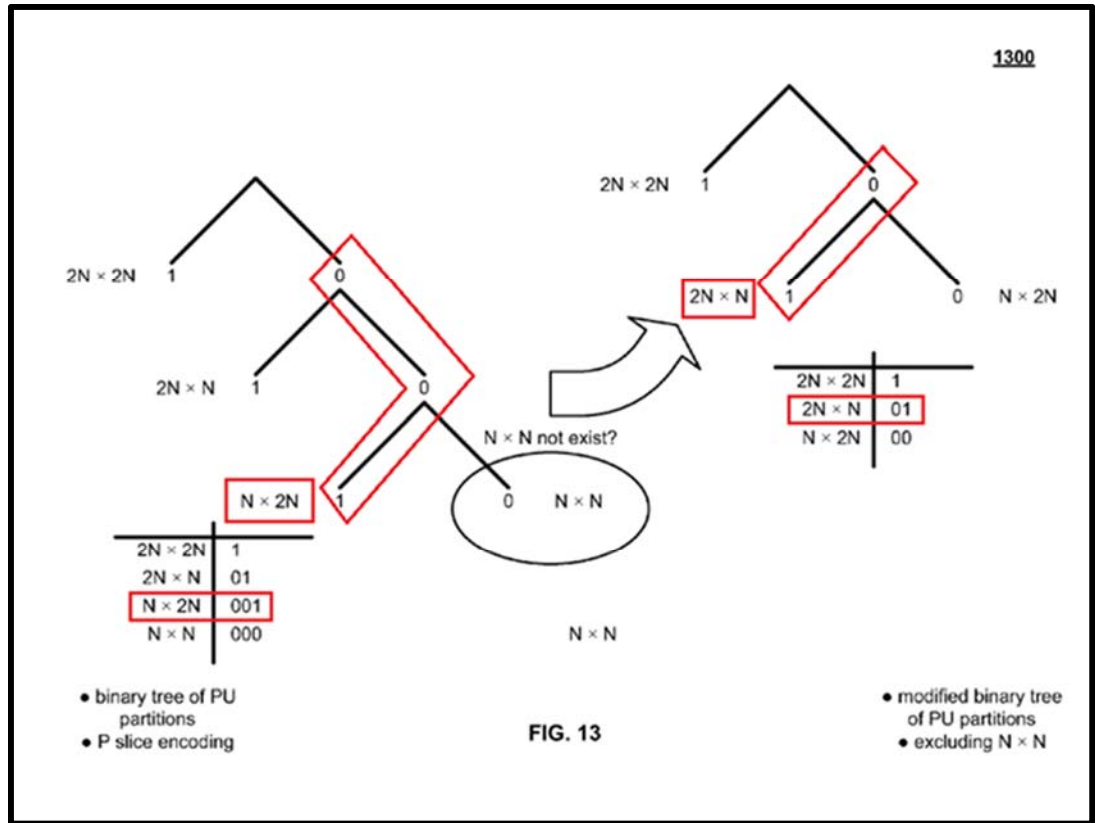
6 264. As the patent explains, among the advances at the time was the use of
7 “predictive” (P) slices and “bi-predictive” (B) slices as components of the video
8 data being processed.¹⁰⁶ P slices and B slices, in turn, are comprised of smaller
9 components, including “coding tree units” and “coding units.” As described in the
10 ’283 Patent, coding units can be “encoded” for different types of “prediction”
11 processing, namely, “inter-prediction” or “intra-prediction.”¹⁰⁷ Subsequently,
12 “prediction units” (PU) can be encoded for different “partition modes,” to be used
13 in the intra- or inter-prediction processing.

14 265. Generally speaking, prediction and partition modes are encoded using
15 binary “codewords.” These codewords can be generated using a “binary tree.” In
16 general terms, a binary tree is a data structure that can be used to represent data and
17 associate it with a corresponding bit sequence and vice versa. In this context, a
18 “binary tree” data structure is used to create a sequence of binary numbers based on
19 the selection of a “1” or a “0” at different positions (sometimes referred to as
20 “nodes”) in the binary tree, starting at the beginning (“root”). Traversing the binary
21 tree from the root to a “leaf” (an endpoint) results in a bit sequence that corresponds
22

23
24 ¹⁰⁶ ’283 Patent, at 17:24-48.

¹⁰⁷ See, e.g., *id.* at 17:9-12.

1 to a specific encoding.¹⁰⁸ Figure 13 of the '283 Patent, for example, represents
 2 binary trees as follows:



Source: '283 Patent, Fig. 13 (annotations added).

15
16 In the left side of the exemplary image above, the annotated binary tree generates a
 17 codeword of “001,” which is associated with a partition mode of “Nx2N.”¹⁰⁹

18 266. The use of P slices and B slices in an encoding protocol can introduce
 19 inefficiencies when encoding coding units of the different slice types. This is
 20 because at the time of the invention of the '283 Patent, the encoding of P and B
 21 slices required separate codewords generated using different binary trees, which
 22 resulted in extra burdens on the encoding system and higher overhead.¹¹⁰

23 ¹⁰⁸ See *id.* at Figs. 13, 14, 15.

24 ¹⁰⁹ *Id.* at Fig. 13 (annotations added).

¹¹⁰ *Id.* at 19:21-35; 20:19-30.

1 267. Thus, there was a need for a method and system that could reduce this
2 inefficiency.¹¹¹ The invention described in the '283 Patent meets that need.
3 Specifically, the '283 Patent's system and method for using "only a singular
4 codebook . . . for both processing of the B slices and P slices," may be employed to
5 provide a "very efficient implementation."¹¹² The methods and systems described
6 in the '283 Patent thus improve the functionality of computer systems by improving
7 the way they compress and process video and audio data.

8 268. The inventions described and claimed in the '283 Patent include
9 encoding using two syntax elements derived from a single binary tree that can be
10 applied to both P slices and B slices to indicate (1) whether inter-prediction or intra-
11 prediction applies to a selected coding unit, and (2) the prediction unit partition
12 mode that applies based on whether or not the selected coding unit is the smallest
13 coding unit (SCU) having a prediction unit size NxN.

14 269. The '283 Patent claims methods and systems that use a single binary
15 tree to encode coding unit (CU) prediction when processing P slices and B slices
16 for digital video data. Independent claim 1, for example, recites:

17 A video processing device comprising:
18 a video encoder configured to:
19 encode an input video signal to generate an output
20 bitstream;
21 employ a single binary tree when processing at least one P
22 slice and at least one B slice to generate the output
23 bitstream, wherein the at least one P slice is used for
 unidirectional prediction forward or behind in at least one
 frame sequence, and wherein the at least one B slice is

24 ¹¹¹ '283 Patent at 18:55-19:12.

¹¹² *Id.* at 18:60-63.

1 used for bidirectional prediction both forward and behind
2 in the at least one frame sequence;

3 employ the single binary tree to encode coding unit (CU)
4 prediction based on a selected CU that is selected from a
5 plurality of CUs when generating a first syntax element
6 for both the at least one P slice and the at least one B slice
7 that undergo entropy encoding to generate the output
8 bitstream, wherein the first syntax element specifies intra-
9 prediction processing or inter-prediction processing for
10 the selected CU; and

11 employ the single binary tree to encode prediction unit
12 (PU) partition mode based on the selected CU when
13 generating a second syntax element for both the at least
14 one P slice and the at least one B slice that undergo the
15 entropy encoding to generate to generate [sic] the output
16 bitstream, wherein the second syntax element specifies
17 the PU partition mode for the selected CU, wherein the
18 PU partition mode is based on a size $N \times N$ PU when the
19 selected CU is a smallest CU (SCU) of the plurality of
20 CUs and is based on a different size PU than the size $N \times N$
21 PU when the selected CU is another CU than the SCU of
22 the plurality of CUs, wherein N is a positive integer.

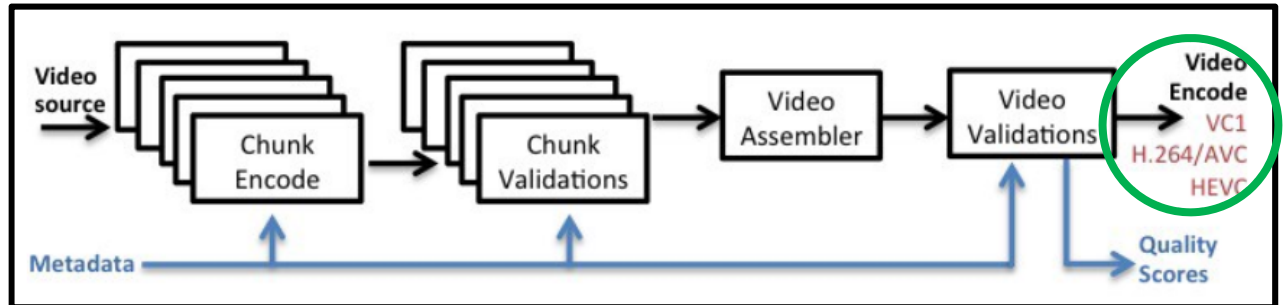
23 270. Netflix directly infringes at least claim 1 of the '283 Patent at least in
24 the exemplary manner described below.

271. Netflix developed, operates, and uses a “video encoding pipeline”, *i.e.*,
a series of video processing applications. Netflix uses its video encoding pipeline to
generate encoded video files in a variety of formats, which it then uses to stream
movie and TV content to its subscribers. As Netflix explains:

We ingest high quality video sources and generate video
encodes of various codec profiles, at multiple quality
representations per profile. The encodes are packaged
and then deployed to a content delivery network for
streaming. During a streaming session, the client requests
the encodes it can play and adaptively switches among

1 quality levels based on network conditions.¹¹³

2 272. Among other encoding formats, Netflix uses its video encoding
3 pipeline to generate content in the H.265 format, also known as High Efficiency
4 Video Coding or “HEVC” (“H.265”).



5
6
7
8
9 Source: <https://medium.com/netflix-techblog/high-quality-video-encoding-at-scale-d159db052746>.

10
11 273. On information and belief, the Netflix video encoder includes a
12 binarization system that employs a “single binary tree.” For example, the H.265
13 format employs a “binary tree” in the processing (i.e., encoding) of P slices and B
14 slices. Specifically, the H.265 format utilizes coding schemes that define a unique
15 mapping of syntax element values to sequences of binary symbols, which are
16 interpreted in terms of a binary code tree.¹¹⁴

17 274. On information and belief, the Netflix video encoder employs the
18 binary tree to generate a “first syntax element” of a coding unit, wherein the “first
19 syntax element” specifies intra-prediction processing or inter-prediction processing
20 for the selected coding unit. For example, H.265 reference software features the
21

22 ¹¹³ <https://medium.com/netflix-techblog/high-quality-video-encoding-at-scale-d159db052746>.

23 ¹¹⁴ See, e.g., *ITU-T Recommendation H.265: High Efficiency Video Coding*, at §
24 7.3.8.5 (November 2019) (available at <https://www.itu.int/rec/T-REC-H.265>). (“H.265 Recommendation”) (describing binary tree coding syntax using pseudo code).

1 parameter PredMode which specifies the “prediction mode” (i.e., intra-prediction
 2 processing or inter-prediction processing) of a selected coding unit. The function
 3 pcCU->isIntra() then returns either 1 or 0 depending on whether intra- or inter-
 4 prediction mode is selected for the particular coding unit. If the selected coding
 5 unit pcCU was encoded using intra-prediction mode (i.e., pcCU->isIntra() is
 6 TRUE) the encodeBin() function codes 1. If inter-prediction mode is selected, the
 7 encodeBin() codes 0.¹¹⁵

```
8 // supported prediction type
enum PredMode
{
9  MODE_INTER = 0,    ///< inter-prediction mode
  MODE_INTRA = 1,    ///< intra-prediction mode
  NUMBER_OF_PREDICTION_MODES = 2,
10 };
```


```
11 Void TEncSbac::codePartSize( TComDataCU* pcCU, UInt
  uiAbsPartIdx, UInt uiDepth )
{
12   PartSize eSize      = pcCU->getPartitionSize(
  uiAbsPartIdx );
  ...
  const UInt log2DiffMaxMinCodingBlockSize = pcCU-
  >getSlice()->getSPS()-
  >getLog2DiffMaxMinCodingBlockSize();
13
  if ( pcCU->isIntra( uiAbsPartIdx ) )
  {
14     if ( uiDepth == log2DiffMaxMinCodingBlockSize )
  {
15       m_pcBinIf->encodeBin( eSize == SIZE_2Nx2N? 1 : 0,
  m_cCUPartSizeSCModel.get( 0, 0, 0 ) );
  }
  }
  return;
  ...
  ...
  }
17 }
```

18 Source: H.265.2: Reference software for ITU-T H.265 High Efficiency
 19 Video Coding (December 2016) (available at [https://www.itu.int/rec/T-REC-](https://www.itu.int/rec/T-REC-H.265.2)
 20 [H.265.2](https://www.itu.int/rec/T-REC-H.265.2)).

21 275. The binary tree that the Netflix video encoder employs to encode a
 22 coding unit prediction mode using a first syntax element for coding units for both P
 23

24 ¹¹⁵ See also, H.265 Recommendation at § 7.4.9.5.

1 slices and B slices thus may be represented as follows, in which “1” codes for intra-
 2 prediction, and “0” codes for inter-prediction:



3
 4
 5
 6
 7
 8
 9
 10 276. The Netflix video encoder is configured to employ the single binary
 11 tree to encode prediction unit (“PU”) partition mode based on the selected coding
 12 unit “when generating a second syntax element for both the at least one P slice and
 13 the at least one B slice,” wherein the second syntax element specifies the PU
 14 partition mode for the selected CU. For example, the H.265 reference software
 15 specifies the parameter “PartSize” and the function “getPartitionSize” which
 16 identifies the partition mode for the selected CU (“pcCU”). This particular example
 17 shows how in the case of a $2N \times 2N$ partition, the function encodeBin() generates the
 18 binary code “1”.

```

19  /// supported partition shape
20  enum PartSize
21  {
22  SIZE_2Nx2N = 0, ///< symmetric motion partition, 2Nx2N
23  SIZE_2NxN = 1, ///< symmetric motion partition, 2Nx N
24  SIZE_Nx2N = 2, ///< symmetric motion partition, Nx2N
25  SIZE_NxN = 3, ///< symmetric motion partition, Nx N
26  SIZE_2NxN_U = 4, ///< asymmetric motion partition, 2Nx(
27  N/2) + 2Nx(3N/2)
28  SIZE_2NxN_D = 5, ///< asymmetric motion partition,
29  2Nx(3N/2) + 2Nx( N/2)
30  SIZE_nLx2N = 6, ///< asymmetric motion partition,
31  (N/2)x2N + (3N/2)x2N
32  SIZE_nRx2N = 7, ///< asymmetric motion partition,
33  (3N/2)x2N + ( N/2)x2N
34  NUMBER_OF_PART_SIZES = 8
35  };
  
```



```

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24

```

```

Void TEncSbac::codePartSize( TComDataCU* pcCU, UInt
uiAbsPartIdx, UInt uiDepth )
{
    PartSize eSize = pcCU->getPartitionSize(
uiAbsPartIdx );
    ...|
    switch(eSize)
    {
        case SIZE_2Nx2N:
        {
            m_pcBinIf->encodeBin( 1, m_cCUPartSizeSCModel.get(
0, 0, 0) );
            break;
        }
        ...
    }
}

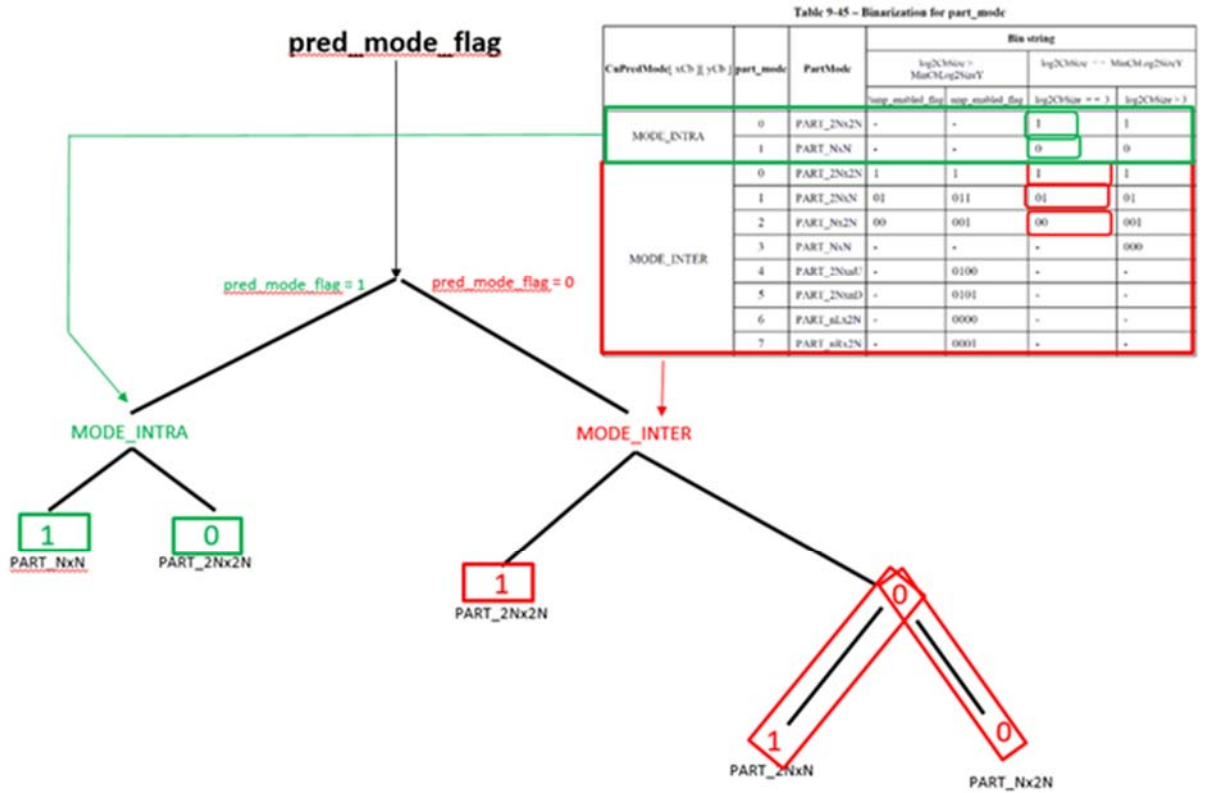
```

Source: H.265.2: Reference software for ITU-T H.265 High Efficiency Video Coding (December 2016) (available at <https://www.itu.int/rec/T-REC-H.265.2>).

277. As indicated in the Table 9-45 of the H.265 Recommendation, the partition mode for a particular coding unit is dependent on the prediction mode (CuPredMode) of the coding unit.¹¹⁶ Table 9-45 depicts the first syntax element (CuPredMode) and the second syntax element (encoding for the partition mode) for different size coding unit cases. The illustration below, in conjunction with Table 9-45, depicts the single binary tree encoding both the prediction mode and the partition mode.

¹¹⁶ See H.265 Recommendation, at Table 9-45.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24



Source: H.265 Recommendation at Table 9-45 (annotations added).

278. Thus, on information and belief, the Netflix video encoder employs the single binary tree to encode prediction unit partition mode when generating a second syntax element, wherein the second syntax element specifies the partition mode for the selected coding unit. In the table above, the coding unit example is of a size “log2CbSize” equals “3”.

279. As indicated in the image above, when the prediction mode is intra-prediction processing, the single binary tree encodes a second syntax element of either 1 or 0, which indicates a partition mode (PartMode (column 3)) of PART_2Nx2N or PART_NxN, respectively. Similarly, when the prediction mode is inter-prediction processing, the single binary tree for the example coding unit size

1 specifies a second syntax element of 1, 01, or 00, which indicates a partition mode
2 of PART_2Nx2N, PART_2NxN, or PART_Nx2N, respectively.

3 280. On information and belief, the Netflix video encoder employs a PU
4 partition mode that is “based on a size N×N PU when the selected CU is a smallest
5 CU (SCU)” and “is based on a different size PU than the size N×N PU when the
6 selected CU is another CU than the SCU.” The H.265 format supports variable
7 prediction block sizes from 64×64 down to 4×4 samples. The minimum size CU is
8 assigned a size N×N Prediction Unit (PU) as a special case.¹¹⁷ Thus, the prediction
9 unit partition mode is based on a coding unit size of N×N when the coding unit is
10 the smallest coding unit, and the prediction unit partition mode is based on a
11 different size prediction unit when the coding unit is a size other than the smallest
12 coding unit.¹¹⁸ Thus, the Netflix video encoding system, as described above,
13 practices at least claim 1 of the ’283 Patent.

14 281. Indeed, Netflix has published studies in which it has quantified the
15 benefits of H.265 encoding. Through its own testing, Netflix determined that
16 H.265 encoders can achieve equivalent subjective reproduction quality as encoders
17 that conform to H.264/MPEG-4 AVC while using approximately 50% less bit rate
18 (i.e., the amount of data “bits” transmitted per second).¹¹⁹ Thus, Netflix has
19 specifically recognized benefits achieved by the inventions of the ’283 Patent.

22 ¹¹⁷ See e.g., *Overview of the High Efficiency Video Coding (HEVC) Standard*, IEEE
23 Transactions On Circuits And Systems For Video Technology, Vol. 22, No. 12,
(December 2012).

24 ¹¹⁸ See Table 9-45 at column “PartMode.”

¹¹⁹ <https://netflixtechblog.com/a-large-scale-comparison-of-x264-x265-and-libvpx-a-sneak-peek-2e81e88f8b0f>.

1 289. The '976 Patent addresses specific technical challenges that arose in
2 the computer networking environment as web services became an increasingly
3 popular means of enabling access to software systems.

4 290. The specification of the '976 Patent explains that “[w]eb services are
5 software systems that can be identified by Universal Resource Identifiers (URI) in a
6 fashion that is analogous to the way websites may be identified by Universal
7 Resource Locators (URLs).”¹²⁰ Web services “enhance the way computers
8 communicate with users and each other”¹²¹ and, at the time, they were “quickly
9 transforming the way modern businesses interact[ed] and share[d] information.”¹²²

10 291. However, the use of web services presented a number of challenges, as
11 the '976 Patent describes. For instance, “before a software system can utilize the
12 functionality of a web service, the software system should first be able to locate and
13 connect to the web service,” a process known as “discovery and integration.”¹²³

14 292. As another example, web services may be available to a large number
15 of potential users, with some being accessible worldwide via the internet to a wide
16 variety of users and software systems operating on a wide variety of platforms.¹²⁴
17 “Because of the very large number of potential web service users and the
18 complexity of many web services, web services can push even the most capable
19 servers running the web services to their limits.”¹²⁵ Over-loaded servers may not
20

21
22 ¹²⁰ '976 Patent, 1:19-23.

¹²¹ *Id.* at 1: 26-27.

¹²² *Id.* at 1:18-19.

¹²³ *Id.* at 1:34-38.

¹²⁴ *Id.* at 1:47-59.

¹²⁵ *Id.* at 1:60-63.

1 function properly, may stop handling web service requests, and may cease
2 functioning altogether.¹²⁶

3 293. As the '976 Patent specification explains, one option for addressing the
4 issue of excess load on a server from a web service is to spread the load across
5 “multiple servers all working toward processing web service requests.”¹²⁷
6 However, this requires a system for balancing the load among the multiple
7 servers.¹²⁸

8 294. The '976 Patent further explains that the web service provider may
9 also wish to employ failover servers—“redundant or standby server[s] in the event
10 the primary server fails.”¹²⁹ These failovers “may be used to ensure the continued
11 offering of web services in any number of circumstances that may render the
12 primary server non-functional.”¹³⁰

13 295. The '976 Patent claims specific, novel ways to address these technical
14 challenges through methods, systems, and apparatuses that enable the efficient,
15 effective provision of web services using multiple servers.

16 296. The methods, systems, and apparatuses described in the '976 Patent
17 improve the functionality of a networked system of computing resources by
18 enabling the efficient, effective use of those resources to provide software services.

19 297. Claim 14 of the '976 Patent reads as follows:

20 A computer system comprising:

21 a processor; and a computer recording medium including

22 ¹²⁶ *Id.* at 1:66-2:3.

23 ¹²⁷ *Id.* at 2:14-18.

¹²⁸ *See id.* at 2:14-20.

24 ¹²⁹ *Id.* at 2:55-57.

¹³⁰ *Id.* at 2:65-67.

1 computer executable code executable by the processor for
2 connecting to a web service, the computer executable
code comprising:

3 code for selecting a web service;

4 code for selecting a server among one or more servers
5 capable of running the selected web service, the selected
6 server being selected independent of input from a
requesting application subsequent to selection of the web
service;

7 code for determining a real address for the selected web
8 service running on the selected server; and

9 code for connecting to the selected web service running
on the selected server using the determined real address.

10 298. Claim 15 of the '976 Patent reads as follows:

11 The computer system of claim 14, wherein selecting a
12 web service comprises searching a directory of web
services.

13 299. Claim 18 of the '976 Patent reads as follows:

14 The computer system of claim 14, wherein selecting a
15 server among one or more servers capable of running the
16 selected web service comprises performing a distributed
scheduling algorithm.

17 300. Claim 20 of the '976 Patent reads as follows:

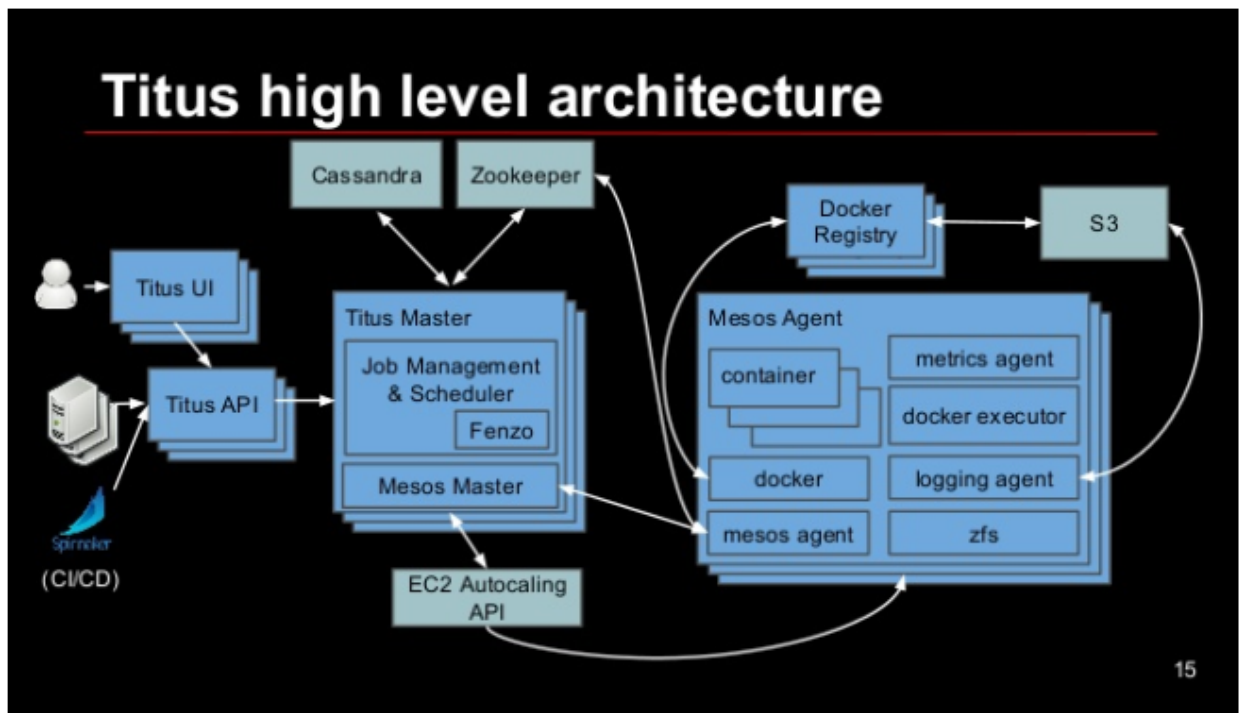
18 The computer system of claim 18, wherein the distributed
scheduling algorithm is a dynamic algorithm.

19 301. Thus, claims 14, 15, 19, and 20 are directed to a novel system for
20 managing the provision of web services using a distributed computing system
21 comprising multiple servers capable of running that web service. Upon information
22 and belief, this system was not well-known, routine, or conventional at the time of
23 the '976 Patent.

24

1 302. Upon information and belief, Netflix directly infringes at least claims
2 14, 15, 18, and 20 of the '976 Patent, at least in the exemplary manner described
3 below.

4 303. Netflix directly infringes the '976 Patent by making, using, offering to
5 sell, and/or selling in the United States its Netflix service, which relies on the
6 Titus/Fenzo container architecture (depicted below) to deploy and manage virtual
7



17 computing resources in flexible, scalable manner.

18 Source: <https://www.slideshare.net/aspyker/netflix-and-containers-titus>.

19 304. In the words of Netflix:

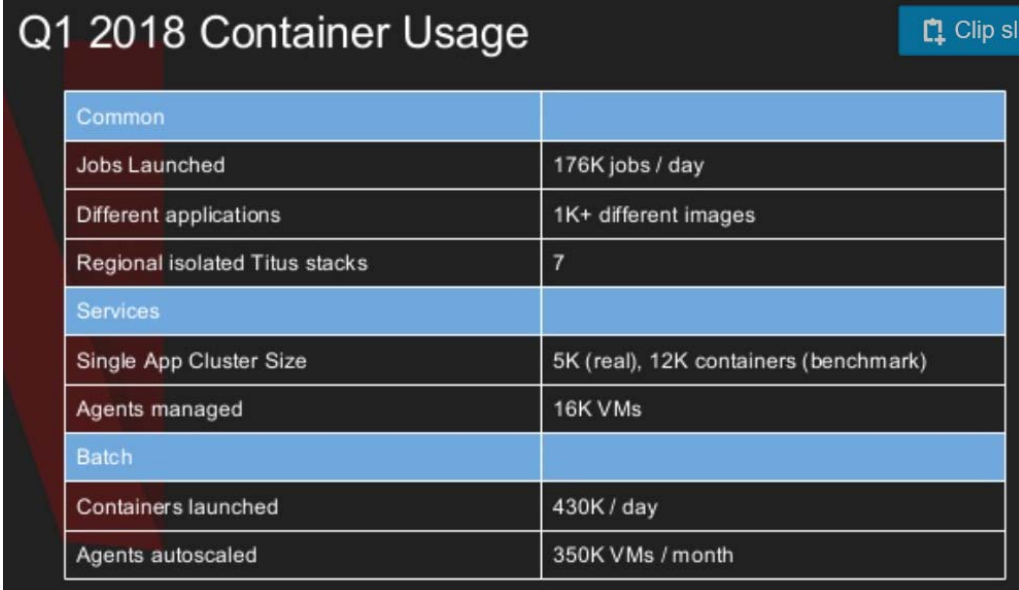
20 Titus is a container management platform that provides
21 scalable and reliable container execution and cloud-native
22 integration with Amazon AWS. Titus was built internally
23 at Netflix and is used in production to power Netflix
24 streaming, recommendation, and content systems.¹³¹

¹³¹ <https://netflix.github.io/titus/>.

1 305. According to Netflix, “Titus powers critical aspects of the Netflix
2 business, from video streaming, recommendations, and machine learning, big data,
3 content encoding, studio technology, internal engineering tools, and other Netflix
4 workloads.”¹³²

5 306. As Netflix explains, “Titus offers a convenient model for managing
6 compute resources.”¹³³ On information and belief, Netflix currently uses the Titus
7 platform to launch and manage millions of containers per week, and to host
8 thousands of applications globally across tens of thousands of virtual machines.¹³⁴

9



Q1 2018 Container Usage	
Common	
Jobs Launched	176K jobs / day
Different applications	1K+ different images
Regional isolated Titus stacks	7
Services	
Single App Cluster Size	5K (real), 12K containers (benchmark)
Agents managed	16K VMs
Batch	
Containers launched	430K / day
Agents autoscaled	350K VMs / month

10
11
12
13
14
15
16
17 Source: <https://www.slideshare.net/aspyker/container-world-2018>.

18 307. With regard to claim 14 of the '976 Patent, Netflix practices “[a]
19 computer system comprising: a processor; and a computer recording medium
20 including computer executable code executable by the processor for connecting to a
21

22 ¹³² <https://netflixtechblog.com/titus-the-netflix-container-management-platform-is-now-open-source-f868c9fb5436>.

23 ¹³³ <https://medium.com/netflix-techblog/titus-the-netflix-container-management-platform-is-now-open-source-f868c9fb5436>.

24 ¹³⁴ *Id.*

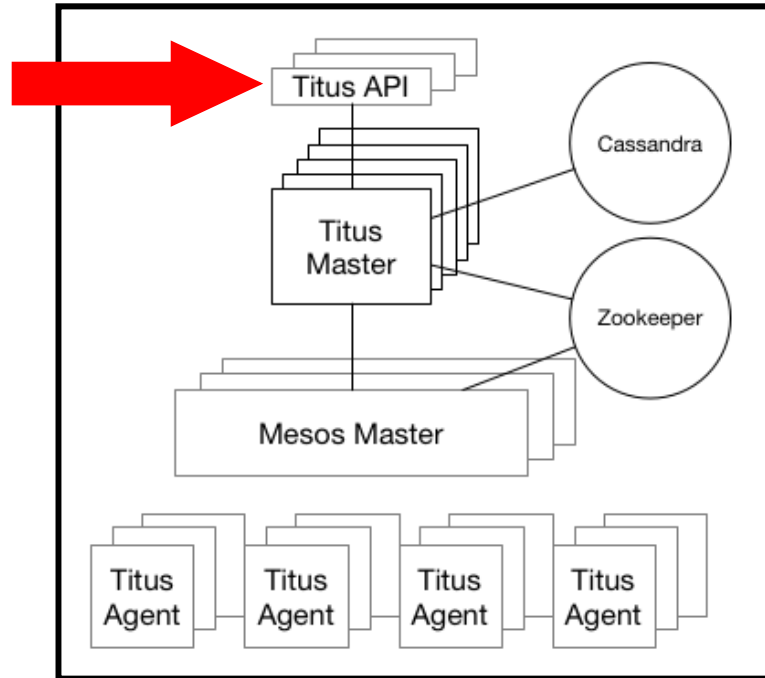
1 web service...” For example, Netflix’s Titus system allows users at Netflix to
2 connect to the platform (a web service), which deploys and manages containers—
3 virtual computing resources—to run the users’ applications and otherwise
4 accommodate their computing needs. On information and belief, the Titus software
5 (computer executable code executable by the processor for connecting to a web
6 service) runs on a plurality of servers in the cloud, where it is stored in the servers’
7 computer recording medium (e.g., computer memory) and executed by the servers’
8 processors.

9 308. The Netflix system also includes “code for selecting a web service.”
10 For instance, the Titus API, also referred to as the Titus Gateway (shown by the red
11 arrow in the diagram below), handles requests from users at Netflix, allowing users
12 to provide the details of the specific container management services they require.¹³⁵
13 As Netflix explains, “[w]ork in Titus is described by a job specification that details
14 what to run (e.g., a container image and entry point), metadata (e.g., the job’s
15 purpose and who owns it), and what resources are required to run it, such as CPU,
16 memory or scheduling constraints (e.g., availability zone balancing or host
17 affinity)”.¹³⁶

23 ¹³⁵ See <https://netflix.github.io/titus/overview/> (“The Titus Gateway is a scalable
24 API tier that handles direct requests from clients and users.”).

¹³⁶ *Id.*

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24



Source: <https://netflix.github.io/titus/overview/>.

309. The Netflix system also includes “code for selecting a server among one or more servers capable of running the selected web service.” For example, as Netflix explains, “Titus is run in production at Netflix, managing thousands of AWS EC2 instances,”¹³⁷ each of which “functions as a virtual private server.”¹³⁸ The Titus Master (identified by the orange arrow in the diagram below) “is responsible for persisting job and task information, scheduling tasks, and managing the pool of EC2 Agents”¹³⁹ (each a Titus Agent, identified by the green arrow in the diagram below). As Netflix further explains, the Titus Master “schedules tasks onto the Agents with available resources and scales the pool of Titus Agents up or down in response to demand.”¹⁴⁰

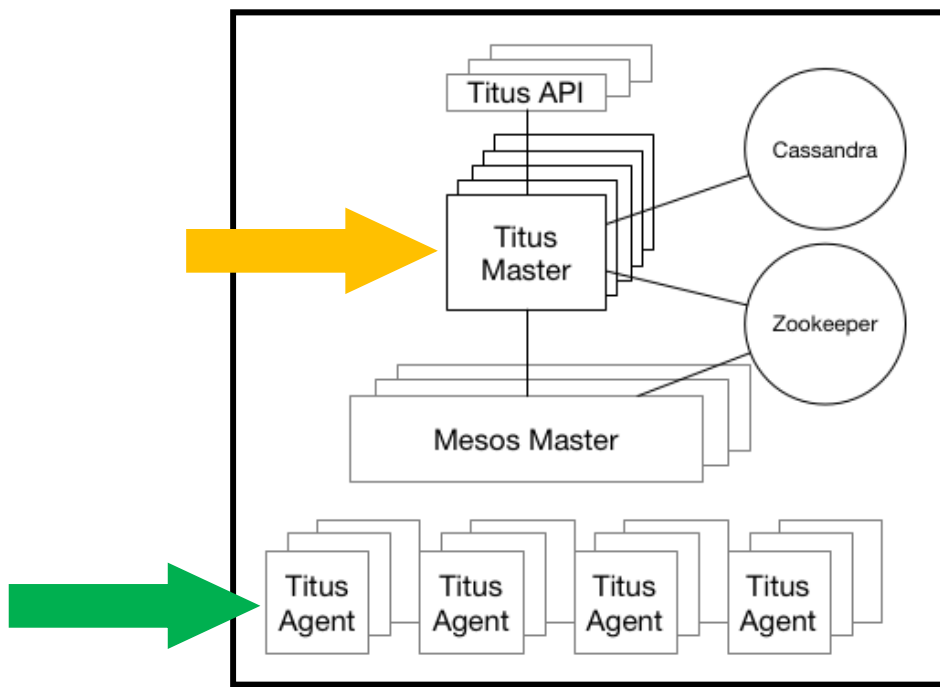
¹³⁷ <https://netflix.github.io/titus/>

¹³⁸ https://en.wikipedia.org/wiki/Amazon_Elastic_Compute_Cloud

¹³⁹ <https://netflix.github.io/titus/overview/>

¹⁴⁰ *Id.*

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24



Source: <https://netflix.github.io/titus/overview/>.

310. Specifically, the Titus Master utilizes a Netflix-developed software called Fenzo to select the server (i.e., the Titus Agent) to use in providing the service requested by the user. As Netflix explains:

The Fenzo scheduler selects potential hosts to assign tasks to based on whether the hosts have sufficient resources to successfully execute the tasks. Each task has its own resource requirements, each host offer lists its available resources, and Fenzo maintains an accounting of which resources are available and which are already assigned on each host.¹⁴¹

311. Within this process, the “selected server” is selected “independent of input from a requesting application subsequent to selection of the web service.”

For example, as Netflix explains:

One of the benefits of using containers through Titus is that it abstracts much of the machine-centric management that applications were doing in VMs. In many cases,

¹⁴¹ <https://github.com/Netflix/Fenzo/wiki/Introduction>.

1 users can tell Titus to “run this application” without
2 worrying about where or on which instance type the
container runs.¹⁴²

3 312. The Netflix system includes “code for determining a real address for
4 the selected web service running on the selected server.” For example, Netflix
5 engineers have explained that “[w]hen Titus is preparing to launch a new container,
6 it creates a network namespace for it, assigns it a specific IP address from an ENI
7 [an Elastic Network Interface], and connects the container’s network namespace to
8 the host’s using a *veth* (virtual Ethernet) interface.”¹⁴³

9 313. Finally, the Netflix system includes “code for connecting to the
10 selected web service running on the selected server using the determined real
11 address.” For example, on information and belief, the Titus system uses the IP
12 address created for the containers (as described in previous paragraph) in utilizing
13 and managing those containers.

14 314. With regard to claim 15 of the ’976 Patent, within the Titus system,
15 “selecting a web service comprises searching a directory of web services.” For
16 example, on information and belief, the Fenzo scheduler used by the Titus Master
17 tracks and maintains data on the Titus Agents and their available resources, and
18 assigns tasks to the Agents based on whether they have sufficient resources to
19 successfully execute those tasks.¹⁴⁴

20 315. With regard to claim 18 of the ’976 Patent, within the Titus system,
21 “selecting a server among one or more servers capable of running the selected web
22 service comprises performing a distributed scheduling algorithm.” For instance,

23 ¹⁴² <https://queue.acm.org/detail.cfm?id=3158370>.

24 ¹⁴³ *Id.*

¹⁴⁴ <https://github.com/Netflix/Fenzo/wiki/Introduction>.

1 Netflix publications explain how the Fenzo scheduler utilized by the Titus Master
2 “adopts a strategy of scoring hosts [i.e., Titus Agents] for a task placement, with
3 higher scores indicating better ‘fitness.’”¹⁴⁵ Fenzo selects the Titus Agent to which
4 to assign tasks “based on whether the hosts have sufficient resources to successfully
5 execute the tasks.”¹⁴⁶

6 316. With regard to claim 20 of the ’976 Patent, the “distributed scheduling
7 algorithm” described in the preceding paragraph “is a dynamic algorithm.” For
8 example, as explained in the preceding paragraph, the Fenzo scheduler takes into
9 account the resources then available on the potential Titus Agents in determining
10 whether the Agent is sufficiently “fit” to handle the task Fenzo seeks to assign.

11 317. Upon information and belief, Netflix directly infringes other claims of
12 the ’976 Patent as well, including for the reasons discussed in the preceding
13 paragraphs.

14 318. At least as of the filing and service of this Amended Complaint,
15 Netflix has had knowledge of the ’976 Patent and Netflix’s infringement thereof.

16 319. Netflix’s infringement of the ’976 Patent, which is knowing and
17 willful at least as of the filing of this Amended Complaint, has caused and
18 continues to cause damage to Avago. Avago is entitled to recover damages
19 sustained as a result of Netflix’s wrongful acts in an amount subject to proof at trial.
20
21
22
23

24 ¹⁴⁵ *Id.*

¹⁴⁶ *Id.*

1 **ELEVENTH CLAIM FOR RELIEF**

2 **(Infringement of U.S. Patent No. 7,457,722)**

3 320. The Broadcom Entities reallege and incorporate by reference the
4 allegations of paragraphs 1-319 set forth above.

5 321. The '722 Patent, entitled "Correlation of Application Instance Life
6 Cycle Events in Performance Monitoring," was duly and legally issued on
7 November 25, 2008 from a patent application filed on November 17, 2004, with
8 Tomer Shain and John A. Colgrove as the named inventors. A copy of the '722
9 Patent is attached hereto as **Exhibit K**.

10 322. The '722 Patent was assigned to Avago, which currently holds all
11 substantial rights, title, and interest in and to the '722 Patent.

12 323. Pursuant to 35 U.S.C. § 282, the '722 Patent is presumed valid.

13 324. The '722 Patent is directed to an improvement in the functionality of
14 complex distributed computer networks.

15 325. The '722 Patent addresses specific technical challenges that arose in
16 information technology and computational services as the service levels required by
17 users continued to increase, while the computer networks used to provide those
18 services became increasingly complex.

19 326. The '722 Patent explains that:

20 In the information technology (IT) departments of modern
21 organizations, one of the biggest challenges is meeting the
22 increasingly demanding service levels required by users.
23 With more and more applications directly accessible to
24 customers via automated interfaces such as the world
wide web, "normal" business hours for many enterprises
are now 24 hours a day, 7 days a week. As a result, the
importance of monitoring and maintaining the quality of

1 computational services has increased dramatically.¹⁴⁷

2 327. The '722 Patent also describes the complexities of monitoring the
3 distributed computing environments used to provide these increasingly demanding
4 services. It explains how “[m]onitoring the server side may comprise monitoring
5 one or more applications executing on a cluster of nodes.”¹⁴⁸ Clusters of nodes
6 “may comprise a utility computing environment, wherein the computational power
7 of one or more servers may be purchased as needed.”¹⁴⁹ “Alternatively, dozens or
8 hundreds of nodes may be organized into interconnected tiers of web servers,
9 application servers and databases.”¹⁵⁰ “Each node in such a system may execute
10 multiple instances of one or more applications, with each instance operable to
11 handle a different client request.”¹⁵¹

12 328. As the '722 Patent explains, in these sorts of distributed computing
13 environments, “application instances may be created or destroyed as demand
14 changes.”¹⁵² Application instances may also “migrate from node to node in
15 response to a hardware or software failure, or in response to a load-balancing
16 algorithm, for example.”¹⁵³

17 329. Existing monitoring systems did not adequately address these
18 challenges. The '722 Patent describes, for instance, how “monitoring systems may
19 be unable to track the migration of instances from node to node, thereby preventing
20 the monitoring system from presenting a complete picture of a given application

21 ¹⁴⁷ '722 Patent, 1:12-20.

22 ¹⁴⁸ *Id.* at 1:25-26.

23 ¹⁴⁹ *Id.* at 1:27-29.

24 ¹⁵⁰ *Id.* at 1:29-31.

¹⁵¹ *Id.* at 1:32-34.

¹⁵² *Id.* at 1:35-36.

¹⁵³ *Id.* at 1:37-39.

1 instance.”¹⁵⁴ Additionally, “monitoring systems may not be operable to monitor
2 data on the creation or destruction of application instances.”¹⁵⁵ Further, the ’722
3 Patent explains that then-existing software performance monitors did not “correlate
4 performance data with events, such as the creation, migration or destruction of a
5 particular application instance, and across all instances of a particular application,
6 application group, and/or technology tier.”¹⁵⁶

7 330. The ’722 Patent claims specific, novel ways to address these technical
8 challenges with a performance monitoring solution for application instances that
9 includes instance life cycle event monitoring and that correlates the performance
10 data with instance life cycle events. The claims of the ’722 Patent are directed to
11 new, improved methods and systems implementing this performance monitoring
12 solution.

13 331. The methods and systems described in the ’722 Patent improve the
14 functionality of a distributed computer system by better monitoring the
15 performance of the applications running in this environment.

16 332. Claim 17 of the ’722 Patent reads as follows:

17 A performance monitoring system, comprising:

18 a processor;

19 a memory coupled to the processor and storing program
20 instructions executable by said processor to implement:

21 collecting performance data for one or more
22 application instances, wherein the performance data
is associated with the performance of said one or
more application instances, wherein each

23 ¹⁵⁴ *Id.* at 1:39-43.

24 ¹⁵⁵ *Id.* at 1:43-45.

¹⁵⁶ *Id.* at 1:45-49.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24

application instance is a computer program
executing on a computer system;

detecting one or more instance life cycle events
associated with said one or more application
instances, wherein said one or more instance life
cycle events comprise at least one of: the creation
of at least one of said one or more application
instances, the destruction of at least one of said one
or more application instances, and the migration of
at least one of said one or more application
instances;

correlating said performance data to said one or
more instance life cycle events; and

storing the correlated performance data.

333. Claim 18 of the '722 Patent reads as follows:

The performance monitoring system of claim 17, wherein
the one or more application instances are instances of an
application, and wherein said correlating said
performance data to one or more instance life cycle events
comprises correlating the one or more instance life cycle
events to the performance of the application.

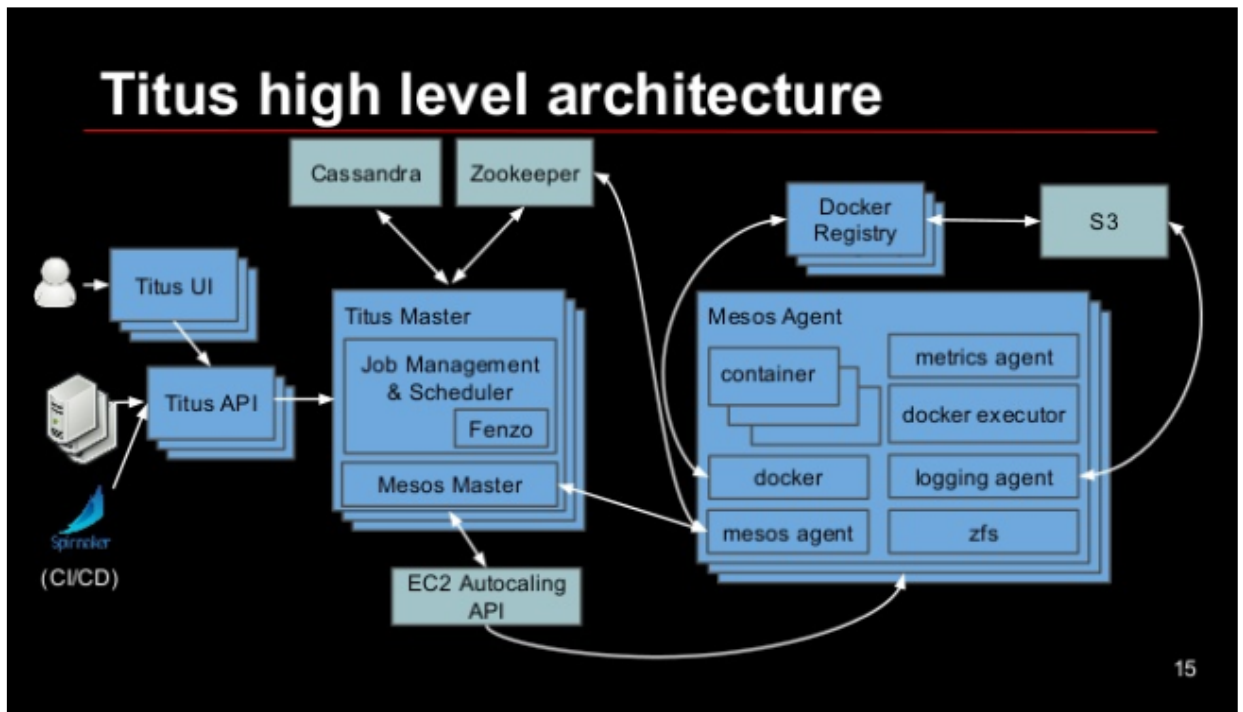
334. Claim 19 of the '722 Patent reads as follows:

The performance monitoring system of claim 18, wherein
said correlating said performance data to one or more
instance life cycle events comprises determining the
change in performance of the application as the one or
more application instances are created and destroyed.

335. Thus, claims 17, 18, and 19 are directed to a novel performance
monitoring system for a distributed computing network that includes application
instance life cycle event monitoring and that correlates the performance data with
instance life cycle events, amongst other things. Upon information and belief, this
system was not well-known, routine, or conventional at the time of the '976 Patent.

1 336. Upon information and belief, Netflix directly infringes at least claims
2 17, 18, and 19 of the '722 Patent, at least in the exemplary manner described below.

3 337. Netflix directly infringes the '722 Patent by making, using, offering to
4 sell, and/or selling in the United States its Netflix service, which relies on the
5 Titus/Atlas architecture (depicted below) to deploy, monitor, and manage virtual
6 computing resources in a flexible, efficient manner.



17 Source: <https://www.slideshare.net/aspkyer/netflix-and-containers-titus>.

18 338. In the words of Netflix:

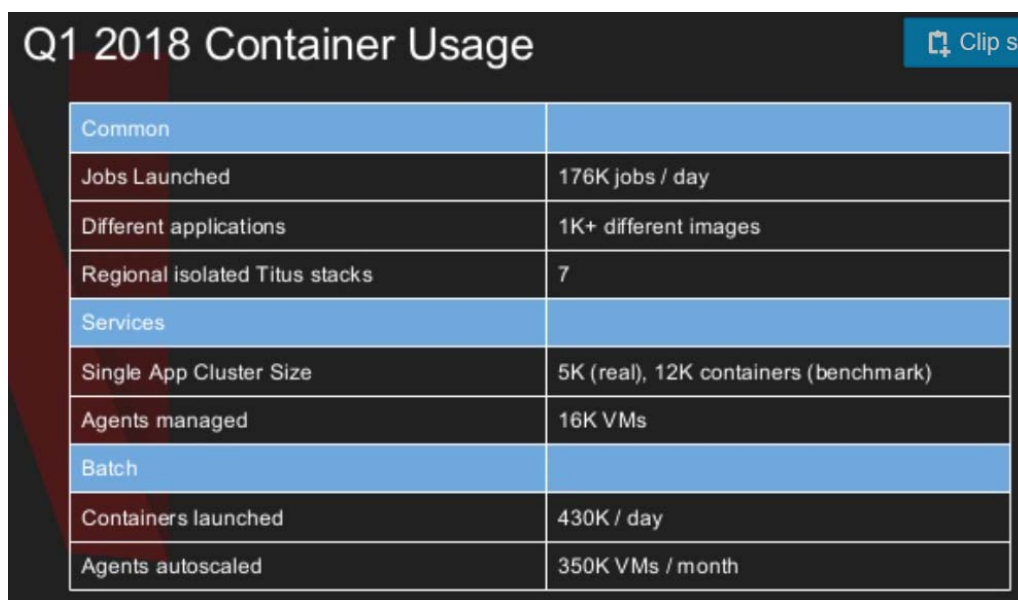
19 Titus is a container management platform that provides
20 scalable and reliable container execution and cloud-native
21 integration with Amazon AWS. Titus was built internally
22 at Netflix and is used in production to power Netflix
23 streaming, recommendation, and content systems.¹⁵⁷

24 ¹⁵⁷ <https://netflix.github.io/titus/>.

1 339. According to Netflix, “Titus powers critical aspects of the Netflix
2 business, from video streaming, recommendations, and machine learning, big data,
3 content encoding, studio technology, internal engineering tools, and other Netflix
4 workloads.”¹⁵⁸

5 340. On information and belief, Netflix currently uses the Titus platform to
6 launch and manage millions of containers per week, and to host thousands of
7 applications globally across tens of thousands of virtual machines.¹⁵⁹

8



9

Q1 2018 Container Usage	
Common	
Jobs Launched	176K jobs / day
Different applications	1K+ different images
Regional isolated Titus stacks	7
Services	
Single App Cluster Size	5K (real), 12K containers (benchmark)
Agents managed	16K VMs
Batch	
Containers launched	430K / day
Agents autoscaled	350K VMs / month

10
11
12
13
14
15

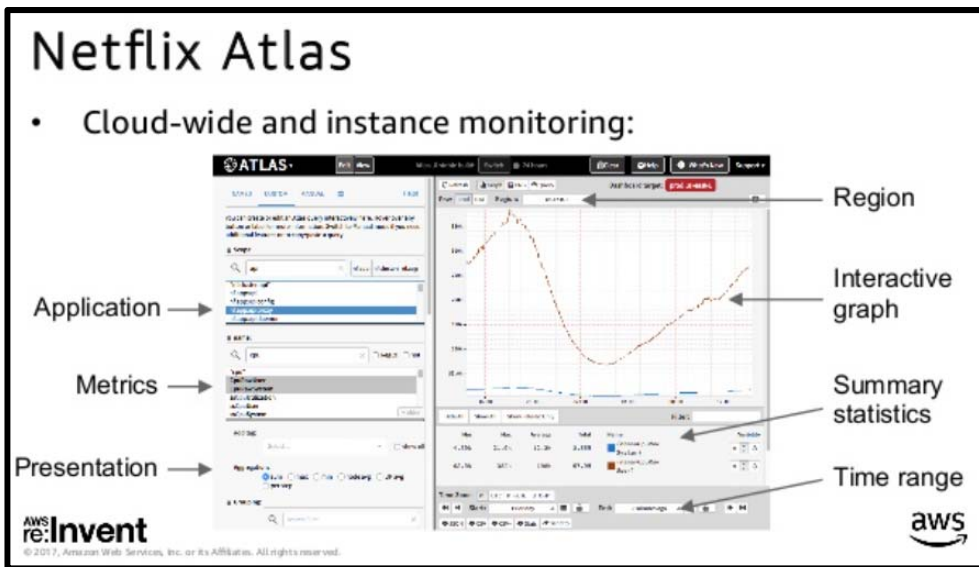
16 Source: <https://www.slideshare.net/aspkyer/container-world-2018>.

17
18 341. The Titus platform relies on Atlas, a performance monitoring tool
19 developed by Netflix, to monitor performance within the Titus system.

20
21
22
23 ¹⁵⁸ <https://netflixtechblog.com/titus-the-netflix-container-management-platform-is-now-open-source-f868c9fb5436>.

24 ¹⁵⁹ See <https://medium.com/netflix-techblog/titus-the-netflix-container-management-platform-is-now-open-source-f868c9fb5436>.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24



Source: <https://www.slideshare.net/brendangregg/how-netflix-tunes-ec2-instances-for-performance>.

342. In Netflix’s words, “Atlas was developed by Netflix to manage dimensional time series data for near real-time operational insight.”¹⁶⁰ “Atlas captures operational intelligence . . . operational intelligence provides a picture of what is currently happening within a system.”¹⁶¹

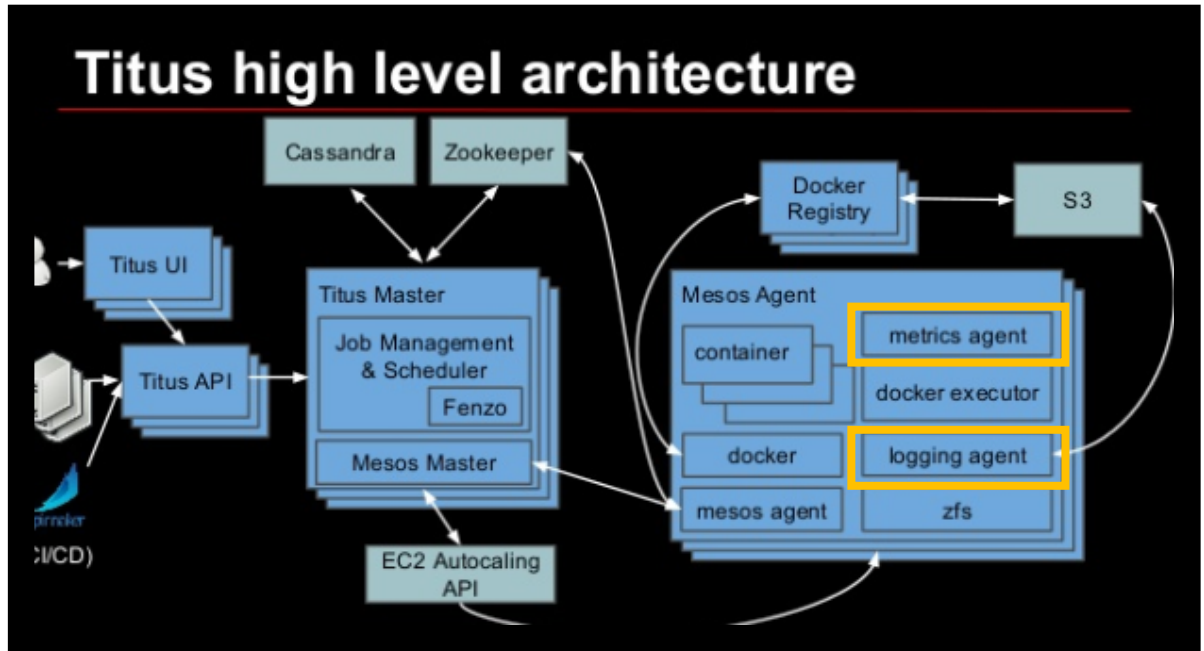
343. Netflix’s Titus system utilizes an Atlas “agent” on each Titus Agent (depicted in yellow in the figure below) that provides “container-level system metrics (e.g., CPU and memory usage) from Titus agents.”¹⁶²

¹⁶⁰ See <https://github.com/Netflix/atlas/wiki>.

¹⁶¹ *Id.*

¹⁶² See <https://queue.acm.org/detail.cfm?id=3158370>.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24

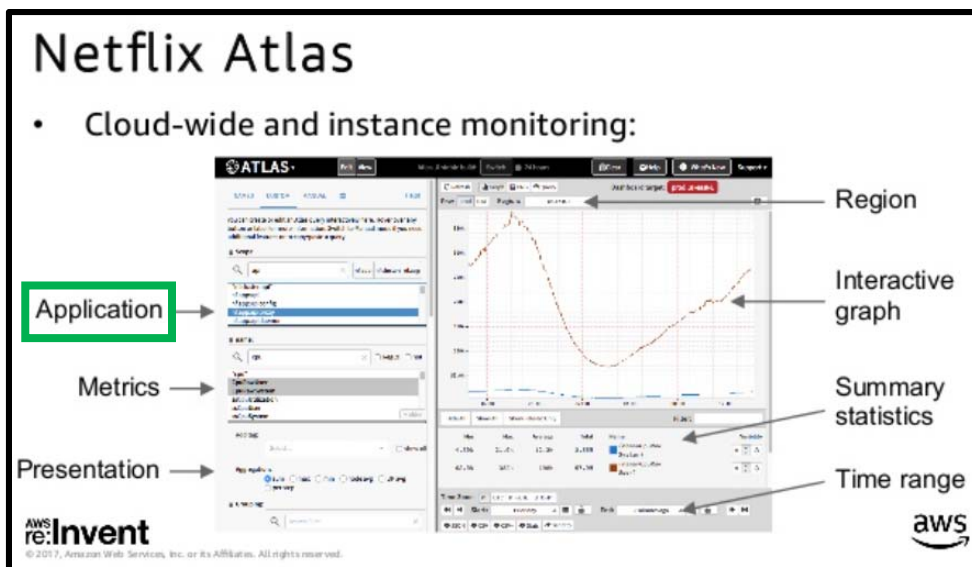


Source: <https://www.slideshare.net/aspyker/netflix-and-containers-titus>
(orange annotations added)

344. With regard to claim 17 of the '722 Patent, Netflix practices a “performance monitoring system, comprising a processor” and “a memory coupled to the processor and storing program instructions executable by said processor.” For example, Netflix’s Titus system utilizes Atlas (a performance monitoring system). On information and belief, the Atlas software (instructions executable by a processor) runs on one or more servers in the cloud, where it is stored in the servers’ memory and executable (and executed) by the servers’ processors.

345. Netflix’s system practices the step of “collecting performance data for one or more application instances, wherein the performance data is associated with the performance of said one or more application instances, wherein each application instance is a computer program executing on a computer system.” For example, Netflix’s Titus system integrates with Atlas to collect performance data on each

1 Titus Agent at the container level.¹⁶³ Each container on the Titus Agent represents
2 an application instance and is a computer program executing on a computer system.



3
4
5
6
7
8
9
10
11 Source: [https://www.slideshare.net/brendangregg/how-netflix-tunes-ec2-](https://www.slideshare.net/brendangregg/how-netflix-tunes-ec2-instances-for-performance)
12 [instances-for-performance](https://www.slideshare.net/brendangregg/how-netflix-tunes-ec2-instances-for-performance) (green annotation added).

13 346. Netflix’s system also practices the step of “detecting one or more
14 instance life cycle events associated with said one or more application instances,
15 wherein said one or more instance life cycle events comprise at least one of: the
16 creation of at least one of said one or more application instances, the destruction of
17 at least one of said one or more application instances, and the migration of at least
18 one of said one or more application instances.” For instance, containers (i.e.,
19 “application instances”) in Linux-based systems such as Titus rely on control

20
21 ¹⁶³ See, e.g., <https://queue.acm.org/detail.cfm?id=3158370> (“Similarly, the health-
22 check polling system was changed to query Titus and provide health-check polling
23 for containers in addition to VMs. The on-instance Atlas telemetry agent was
24 changed to collect and emit container-level system metrics (e.g., CPU and memory
usage) from Titus agents. Previously, it collected only metrics for the entire
host.”); <https://netflix.github.io/titus/overview/> (“The [Titus] Agent sets up on host
resources, such as storage and networking resources, and launches the containers
using Docker. The Agent monitors the task, reporting status and cleaning up
resources when it complies.”).

1 groups (or “cgroups”) to track container processes and expose metrics about CPU,
2 memory, and disk usage.¹⁶⁴ Among the metrics that may be retrieved are several
3 “‘counters’, or values that can only go up, because they represent occurrences of a
4 specific event.”¹⁶⁵ Such events may include, for example, when a container is
5 created or destroyed. For example, the container-level metric “cpuacct.stat”
6 indicates the “amount of time a process has direct control of the CPU, executing
7 process code” (i.e., “user” time) as well as the “time the kernel is executing system
8 calls on behalf of the process” (i.e., “system” time).¹⁶⁶ In other words, the
9 “cpuacct.stat” metric detects when a new container is launched and measures the
10 amount of time that a process has been running on that container.

11 As another example, cgroups can also be used to detect how much CPU,
12 memory, etc., a container (an “application instance”) has used when a container
13 exits or is terminated, how much time has elapsed since the container’s launch (i.e.,
14 container uptime), and various other metrics.¹⁶⁷ As Netflix explains, the Titus/Atlas
15 system utilizes at least the “cpuacct.stat” and container uptime metrics to detect at
16 least when a container is created or destroyed.

17
18
19
20
21
22 ¹⁶⁴ See <https://docs.docker.com/config/containers/runmetrics/>.

23 ¹⁶⁵ *Id.*

24 ¹⁶⁶ *Id.*

¹⁶⁷ *Id.*; see also <https://docs.docker.com/engine/reference/commandline/ps/>
(explaining how the “docker ps” command can be used in Docker-based container
systems such as Titus to detect when a container is created or terminated).

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24

/sys/fs/cgroups (raw)

The best source for per-cgroup metrics. e.g. CPU:

```
# cd /sys/fs/cgroup/cpu,cpuacct/docker/02a7c65f82e3f3e75283944caa4462e82f8f6ff5a7c9a...
# ls
cgroup.clone_children  cpuacct.usage_all      cpuacct.usage_sys     cpu.shares
cgroup.procs           cpuacct.usage_percpu   cpuacct.usage_user    cpu.stat
cpuacct.stat          cpuacct.usage_percpu_sys  cpu.cfs_period_us     notify_on_release
cpuacct.usage         cpuacct.usage_percpu_user  cpu.cfs_quota_us      tasks

# cat cpuacct.usage
1615816262506
# cat cpu.stat
nr_throttled 74
throttled_time 3816445175
```

total time throttled (nanoseconds). saturation metric.
average throttle time = throttled_time / nr_throttled

- <https://www.kernel.org/doc/Documentation/cgroup-v1/./scheduler/sched-bwc.txt>
- <https://blog.docker.com/2013/10/gathering-lxc-docker-containers-metrics/>

Note: `grep cgroup /proc/mounts` to check where these are mounted

These metrics should be included in performance monitoring GUIs

LISA '7

nsenter: Host -> Container top

... Given PID 28321, running top for its container by entering its namespaces:

```
# nsenter -t 28321 -m -p top
top - 18:16:13 up 36 days, 20:28, 0 users, load average: 5.66, 5.29, 5.28
tasks: 6 total, 1 running, 5 sleeping, 0 stopped, 0 zombie
%Cpu(s): 30.5 us, 1.7 sy, 0.0 ni, 65.9 id, 0.0 wa, 0.0 hi, 1.8 si, 0.1 st
KiB Mem: 65958552 total, 54664124 used, 11294428 free, 164232 buffers
KiB Swap: 0 total, 0 used, 0 free. 1592372 cached Mem

  PID USER      PR  NI  VIRT  RES  SHR  S  %CPU  %MEM    TIME+  COMMAND
  301 root        20   0 33.127g 0.023t 37564 S 537.3 38.2 40269:41 java
    1 root        20   0  21404 2236 1812 S   0.0  0.0   4:15.11 bash
 87888 root        20   0  21464 1720 1348 R   0.0  0.0   0:00.00 top
```

Note that it is PID 301 in the container. Can also see this using:

```
# grep NSpid /proc/28321/status
NSpid: 28321 301
```

LISA '7

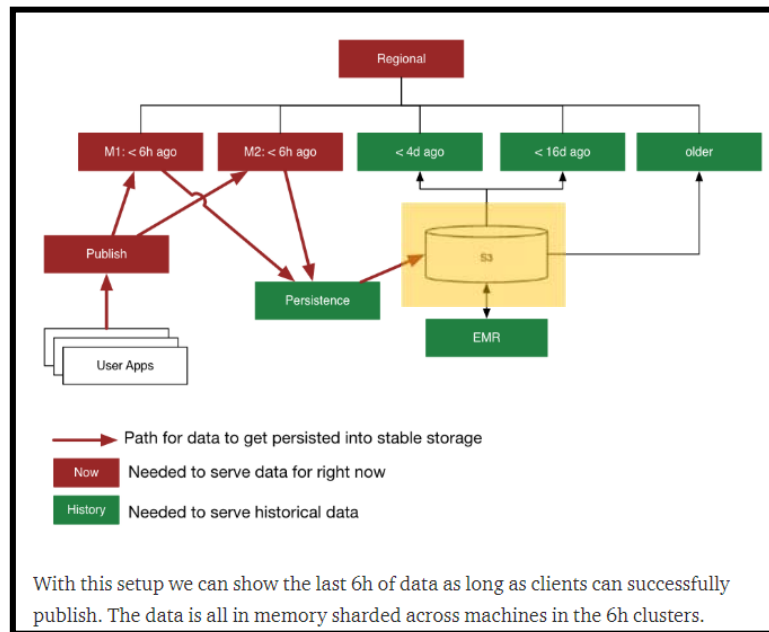
Source: <https://www.slideshare.net/brendangregg/lisa17-container-performance-analysis> (red annotations added).

347. Netflix’s system also practices the step of “correlating said performance data to said one or more instance life cycle events.” For example, as Netflix explains, “Atlas is a backend for storing and querying dimensional time series data. . . . In Atlas[,] each time series is paired with metadata called tags that allow us to query and group the data.”¹⁶⁸ Thus, the Atlas system correlates the

¹⁶⁸ See <https://github.com/Netflix/atlas/wiki/Concepts>.

1 container data collected from Titus to visualize the change in performance over
2 time, including the performance of the system before and after the creation and
3 destruction of individual containers.

4 348. Finally, the Netflix system practices the step of “storing the correlated
5 performance data.” For example, Netflix explains that “Atlas features in-memory
6 data storage, allowing it to gather and report very large numbers of metrics, very
7 quickly.”¹⁶⁹ Netflix further explains that this data is stored in one or more S3
8 “cloud” servers controlled by Netflix.



18 Source: <https://netflixtechblog.com/introducing-atlas-netflixs-primary-telemetry-platform-bd31f4d8ed9a> (orange annotation added).

19
20 349. With regard to claim 18 of the '722 Patent, Netflix practices “[t]he
21 performance monitoring system of claim 17, wherein the one or more application
22 instances are instances of an application.” For example, as previously discussed,
23 each container on the Titus Agent represents an application instance and is a

24 ¹⁶⁹ See <https://github.com/Netflix/atlas/wiki>.

1 computer program executing on a computer system. As Netflix explains, “[u]sing
2 containers, Titus runs any batch application letting the user specify exactly what
3 application code and dependencies are needed.”¹⁷⁰ “Container images [i.e.,
4 instances] provide an easy way to build application-specific images that have only
5 what the application needs.”¹⁷¹

6 350. The Netflix system practices the system of claim 18, “wherein said
7 correlating said performance data to one or more instance life cycle events
8 comprises correlating the one or more instance life cycle events to performance of
9 the application.” For example, as previously discussed, Netflix uses the “Atlas
10 telemetry agent . . . to collect and emit container-level system metrics (e.g., CPU
11 and memory usage) from Titus agents.”¹⁷² As Netflix explains, these container-
12 level metrics include such metrics as “cpuacct.stat” and container uptime (using, for
13 example, the “docker ps” command) that detect and emit data in response to the
14 occurrence of specific events such as the creation or destruction of a container.¹⁷³
15 As previously explained, Atlas correlates these container-level metrics by enabling
16 users to compare the change in a container’s performance over time.¹⁷⁴

17 351. With regard to claim 19 of the ’722 Patent, Netflix practices “[t]he
18 performance monitoring system of claim 18, wherein said correlating said
19 performance data to one or more instance life cycle events comprises determining
20 the change in performance of the application as the one or more application

21 ¹⁷⁰ <https://netflixtechblog.com/the-evolution-of-container-usage-at-netflix-3abfc096781b>.

22 ¹⁷¹ <https://queue.acm.org/detail.cfm?id=3158370>.

23 ¹⁷² *Id.*

24 ¹⁷³ See <https://www.slideshare.net/brendangregg/lisa17-container-performance-analysis>; see also, <https://docs.docker.com/engine/reference/commandline/ps/>.

¹⁷⁴ See <https://github.com/Netflix/atlas/wiki/Concepts>.

1 instances are created and destroyed.” For example, as Netflix explains, the Atlas
2 telemetry agent collects container-level metrics such as “cpuacct.stat” and container
3 uptime (using, for example, the “docker ps” command) that detect and emit data in
4 response to the occurrence of specific events such as the creation or destruction of a
5 container.¹⁷⁵ As Netflix further explains, Atlas displays these metrics along with
6 various other performance data in a time series format that determines the change in
7 performance of an individual container or group of containers over a period of
8 time.¹⁷⁶

9

10 **Namespaces**

11 Worth checking namespace config before analysis:

```

12 # ./dockerps.sh
13
14 CONTAINER   NAME           PID PATH          CGROUP          IPC             NET             NET             PID           USER          UTS
15 host
16 hl7909cd6d1  Titus-1435830-watcher 37280 /sbin/init      4026531835 4026531839 4026531840 4026532533 4026531836 4026531837 4026531838
17 dr:El6506de45 Titus-1392192-watcher 27992 /apps/spaa/spaa 4026531835 4026531354 4026533352 4026532991 4026533355 4026531837 4026533353
18 370a3f041f36 Titus-1243558-watcher 98602 /apps/spaa/spaa 4026531835 4026531290 4026533288 4026533223 4026533291 4026531837 4026533289
19 dc:7549cf6d9a Titus-1243553-watcher 97972 /apps/spaa/spaa 4026531835 4026533216 4026533214 4026533149 4026533217 4026531837 4026533215
20 dc:27769e8b9c Titus-1243546-watcher 97356 /apps/spaa/spaa 4026531835 4026533142 4026533140 4026533075 4026533143 4026531837 4026533141
21 el:8b6189dcd Titus-1243517-watcher 96733 /apps/spaa/spaa 4026531835 4026533068 4026533066 4026533001 4026533069 4026531837 4026533067
22 ab:45227dcaaf Titus-1243516-watcher 96173 /apps/spaa/spaa 4026531835 4026532920 4026532918 4026532830 4026532921 4026531837 4026532919

```

- 23 • A POC “docker ps --namespaces” tool. NS shared with root in red.
- 24 • <https://github.com/docker/docker/issues/32501>
- <https://github.com/kubernetes-incubator/cri-o/issues/868>

LISA

16 Source: <https://www.slideshare.net/brendangregg/lisa17-container-performance-analysis> (green annotations added).

17 352. Upon information and belief, Netflix directly infringes other claims of
18 the '722 Patent as well, including for the reasons discussed in the preceding
19 paragraphs.
20

21
22 ¹⁷⁵ See <https://www.slideshare.net/brendangregg/lisa17-container-performance-analysis>; see also, <https://docs.docker.com/engine/reference/commandline/ps/>
23 (explaining how to use the “docker ps” command to detect when a container is
24 created or destroyed).

¹⁷⁶ See <https://github.com/Netflix/atlas/wiki/Concepts>.

1 where multiple users require computing resources to perform various processing
2 activities.¹⁷⁷

3 361. The '183 Patent specification explains that for distributed computing
4 systems that share resources across multiple users, the various processing activities
5 submitted by users may compete with one another for system resources, resulting in
6 poor performance or system failure by the networked computers.¹⁷⁸ Further, while
7 a distributed system may have several parallel computer devices (e.g., servers)
8 available to perform the job, not every one of the available servers may be
9 appropriate or even suitable for performing the job.¹⁷⁹ “For example, some of the
10 available computer devices may not be adequately provisioned to perform the job
11 while others may be too busy to perform the job in spite of being adequately
12 provisioned.”¹⁸⁰ This too may result in poor performance or system failures within
13 the distributed computing system.

14 362. As the '183 Patent specification explains, one solution addressing this
15 problem is to form complex rule statements using, for example, “an overall
16 utilization value” (e.g., a calculated aggregate of performance metrics that may be
17 based on user-defined parameters), “individually collected utilization metrics” (e.g.,
18 performance metrics for individual devices in the distributed system), and “user-
19 defined parameters” in order to “identify the best computer device available to
20 perform a job, to make decisions about when it is appropriate to take user defined
21 actions, or to provide a new computer device for new work to be performed.”¹⁸¹

22 ¹⁷⁷ '183 Patent, 1:13-20.

23 ¹⁷⁸ *Id.*

¹⁷⁹ *Id.* at 2:33-41.

24 ¹⁸⁰ *Id.*

¹⁸¹ *Id.* at 2:58-3:2.

1 363. The invention of the '183 Patent enables distributed computing
2 systems to make real-time intelligent decisions about the state of user defined
3 computing resource pools while increasing the reliability and efficiency of the
4 system.¹⁸² As the specification explains, the ability to dynamically provision and
5 place tasks within a distributed system is especially beneficial in the context of data
6 centers that may have hundreds or even thousands of computer devices from which
7 to choose.¹⁸³

8 364. Yet another advantage of the '183 Patent's invention is that it
9 "segregate[s] server level rules from service level resource utilization," by creating
10 a separate management structure that is flexible, reactive, and can readily adapt to
11 meet the demands of virtually any distributed computing system regardless of
12 scale.¹⁸⁴

13 365. The '183 Patent claims specific, novel ways to address these technical
14 challenges and achieve at least the advantages described above using
15 unconventional methods, systems, and apparatuses that enable the reliable,
16 efficient, and dynamic provision of resources in a distributed computer system.

17 366. The methods, systems, and apparatuses described in the '183 Patent
18 improve the functionality of a distributed computer system by enabling the reliable,
19 efficient, and dynamic provision of its computing resources to better handle
20 computing tasks.

21 367. Independent claim 11 of the '183 Patent reads as follows:

22 A system for dynamic resource provisioning for job

23 _____
¹⁸² *Id.* at 2:53-58.

¹⁸³ *Id.* at 2:46-49.

¹⁸⁴ *See id.* at 2:7-9.

1 placement, comprising:
2 an interface operable to receive a request to perform a job
3 on an unspecified computer device from a client machine;
4 one or more processors operable to:
5 determine one or more job criteria for performing the job,
6 the one or more job criteria defining one or more
7 operational characteristics needed for a computer device
8 to perform the job;
9 determine one or more utilization criteria for performing
10 the job;
11 provide a list of available computer devices, the list
12 comprising a plurality of computer devices currently
13 provisioned to perform computer operations;
14 from the list of available computer devices, determine a
15 list of suitable computer devices for performing the job by
16 comparing operational characteristics for each available
17 computer device with the job criteria, the list of suitable
18 computer devices comprising one or more computer
19 devices having operational characteristics that satisfy the
20 job criteria;
21 use the utilization criteria to determine whether one or
22 more underutilized computer devices exist on the list of
23 suitable computer devices, the one or more underutilized
24 computer devices having a suitable level of utilization for
performing the job; and
if the one or more underutilized computer devices exist,
identify the one or more underutilized computer devices
to the client machine.

368. Dependent claim 12 of the '183 Patent reads as follows:

The system of claim 11, wherein the one or more processors are operable to use the utilization criteria to identify one or more underutilized computer devices by, for each computer device on the list of suitable computer devices:

1 determining a utilization value;
2 comparing the utilization value with the utilization
3 criteria; and
4 removing the each device from the list of suitable
5 computer devices if the utilization value does not satisfy
6 the utilization criteria.

369. Dependent claim 15 of the '183 Patent reads as follows:

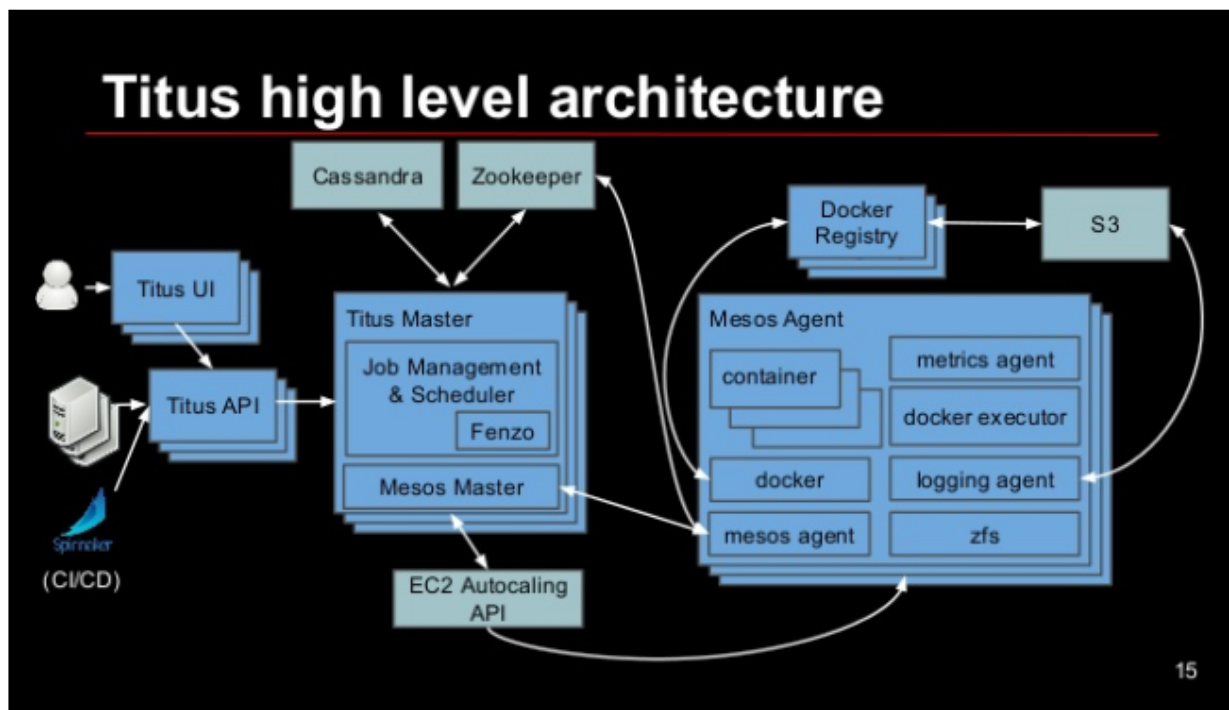
7 The method of claim 11, wherein the one or more
8 processors are configured to:
9 use the utilization criteria to determine whether one or
10 more underutilized computer devices exist on the list of
11 suitable computer devices by identifying a computer
12 device having a lowest level of activity relative to other
13 computer devices on the list of suitable computer devices;
14 and
15 identify the computer device having the lowest level of
16 utilization to the client machine.

370. Thus, claims 11, 12, and 15 are directed to a novel system for the
17 dynamic provision of computing resources within a distributed computing system
18 that, amongst other things, utilizes job criteria, operational characteristics of the
19 computing resources, and utilization criteria of the computing resources to place
20 jobs without requiring the user to specify the computer device to be used. Upon
21 information and belief, this system was not well-known, routine, or conventional at
22 the time of the '183 Patent.

371. Upon information and belief, Netflix directly infringes at least claims
23 11, 12, and 15 of the '183 Patent, at least in the exemplary manner described below.

372. Netflix directly infringes the '183 Patent by making, using, offering to
24 sell, and/or selling in the United States its Netflix service, which relies on the

1 Titus/Fenzo container architecture (depicted below) to deploy and manage virtual
2 computing resources in flexible, scalable manner.



13 Source: <https://www.slideshare.net/aspkyer/netflix-and-containers-titus>.

14 373. In the words of Netflix:

15 Titus is a container management platform that provides
16 scalable and reliable container execution and cloud-native
17 integration with Amazon AWS. Titus was built internally
18 at Netflix and is used in production to power Netflix
19 streaming, recommendation, and content systems.¹⁸⁵

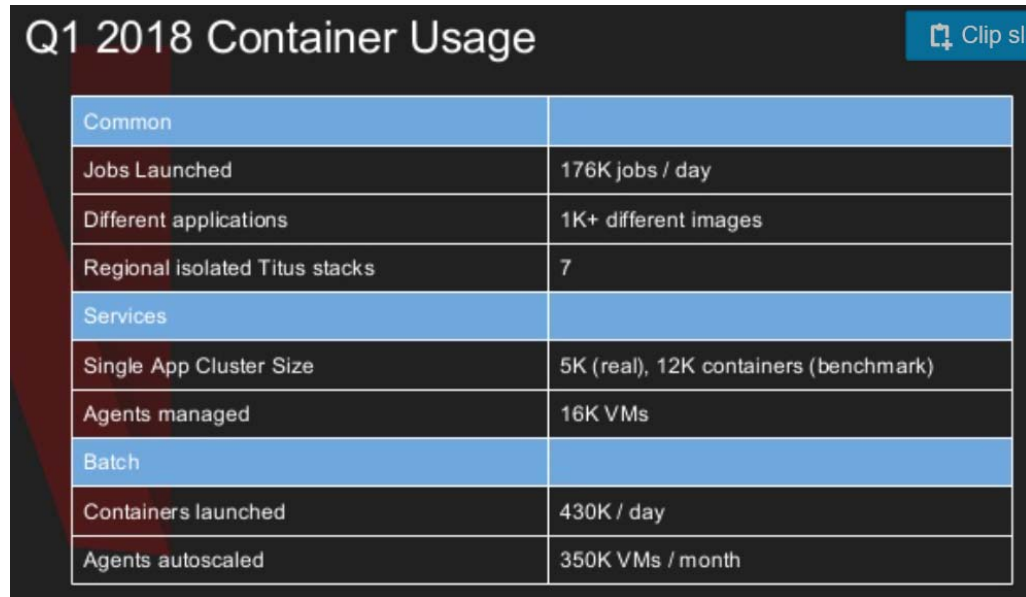
20 374. According to Netflix, “Titus powers critical aspects of the Netflix
21 business, from video streaming, recommendations, and machine learning, big data,
22 content encoding, studio technology, internal engineering tools, and other Netflix
23 workloads.”¹⁸⁶

24 ¹⁸⁵ See <https://netflix.github.io/titus/>.

¹⁸⁶ <https://netflixtechblog.com/titus-the-netflix-container-management-platform-is-now-open-source-f868c9fb5436>.

1 375. As Netflix explains, “Titus offers a convenient model for managing
2 compute [sic] resources.”¹⁸⁷ On information and belief, Netflix currently uses the
3 Titus platform to launch and manage millions of containers per week, and to host
4 thousands of applications globally across tens of thousands of virtual machines.¹⁸⁸

5



Q1 2018 Container Usage	
Common	
Jobs Launched	176K jobs / day
Different applications	1K+ different images
Regional isolated Titus stacks	7
Services	
Single App Cluster Size	5K (real), 12K containers (benchmark)
Agents managed	16K VMs
Batch	
Containers launched	430K / day
Agents autoscaled	350K VMs / month

6
7
8
9
10
11
12
13 Source: <https://www.slideshare.net/aspyker/container-world-2018>.

14
15 376. According to Netflix, “Titus is a framework [that runs] on top of
16 Apache Mesos, a cluster-management system that brokers available resources
17 across a fleet of machines.”¹⁸⁹ As Netflix explains, “Titus consists of a replicated,
18 leader-elected scheduler called Titus Master.”¹⁹⁰ The Titus Master is responsible
19 for persisting job and task information, scheduling tasks, and managing the pool of
20 EC2 virtual machine instances (“Titus Agents”).¹⁹¹

21
22 ¹⁸⁷ See <https://medium.com/netflix-techblog/titus-the-netflix-container-management-platform-is-now-open-source-f868c9fb5436>.

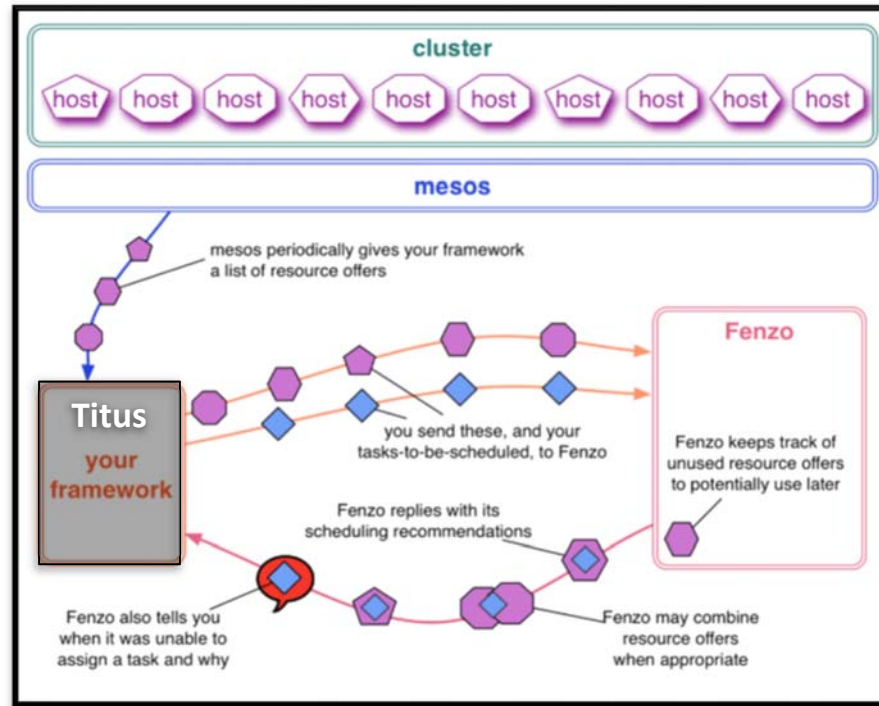
23 ¹⁸⁸ *Id.*

24 ¹⁸⁹ *Id.*

¹⁹⁰ *Id.*

¹⁹¹ *Id.*

1 377. To perform these task placement and resource scheduling functions,
2 the Titus Master relies on Fenzo, an extensible scheduler library developed and
3 used internally by Netflix.¹⁹² As Netflix explains, “Fenzo is a Java library that
4 implements a generic task scheduler for Apache Mesos frameworks.”¹⁹³ Apache
5 Mesos is an open source project for managing clusters of computer devices
6 (individually referred to as “hosts”).¹⁹⁴



7
8
9
10
11
12
13
14
15
16
17 Source: <https://github.com/Netflix/Fenzo/wiki> (gray annotation added) .

18 378. With regard to claim 11, Netflix practices a “system for dynamic
19 resource provisioning for job placement.” For example, the Titus Master uses the
20 Fenzo task scheduler “in conjunction with its agent autoscaling capabilities to grow
21 and shrink the agent pool dynamically as workload demands.”¹⁹⁵

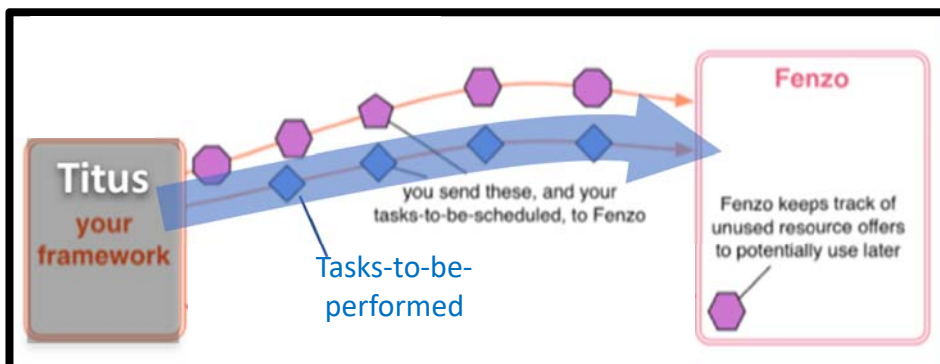
22
23 ¹⁹² See <https://queue.acm.org/detail.cfm?id=3158370>.

¹⁹³ See <https://github.com/Netflix/Fenzo/wiki>.

¹⁹⁴ See <http://mesos.apache.org/>.

¹⁹⁵ See <https://queue.acm.org/detail.cfm?id=3158370>.

1 379. Netflix also practices the step of “an interface operable to receive a
2 request to perform a job on an unspecified computer device from a client machine.”
3 For example, the Titus Master provides an “interface” that receives requests to
4 perform a job on one or more Titus Agents (unspecified computer device) from one
5 or more Gateway instances (client machine). As Netflix explains, “[t]he Master
6 receives requests from Gateway instances and creates and persists job and task info
7 in response.”¹⁹⁶ This request does not specify a computer device for performing a
8 particular job—rather, the Titus/Fenzo architecture “consider[s] only the task’s
9 required resources and scheduling constraints” for task placement.¹⁹⁷



15 Source: <https://github.com/Netflix/Fenzo/wiki> (annotations added).

16 380. Netflix also practices the step of “one or more processors.” For
17 example, on information and belief, the Titus and Fenzo software run on one or
18 more servers in the cloud, where the processors in those servers execute the
19 software.

20 381. Netflix also practices the step wherein the one or more processors are
21 operable to “determine one or more job criteria for performing the job, the one of
22 more job criteria defining one or more operational characteristics needed for a
23

24 ¹⁹⁶ See <https://netflix.github.io/titus/overview/>.

¹⁹⁷ See <https://queue.acm.org/detail.cfm?id=3158370>.

1 computer device to perform the job.” For example, each task in Titus is assigned a
2 job specification that defines the job criteria, including specific operational
3 characteristics required for the task. In Netflix’s words:

4 Work in Titus is described by a job specification that
5 details what to run (e.g., a container image and entry
6 point), metadata (e.g., the job’s purpose and who owns it),
7 and what resources are required to run it, such as CPU,
8 memory, or scheduling constraints (e.g., availability zone
9 balancing or host affinity).¹⁹⁸

10 382. Netflix also practices the step wherein the one or more processors are
11 operable to “determine one or more utilization criteria for performing the job.” For
12 example, the Fenzo scheduler operating in conjunction with the Titus Master
13 determines the utilization criteria for each task as well as for the available host
14 resources. As Netflix explains:

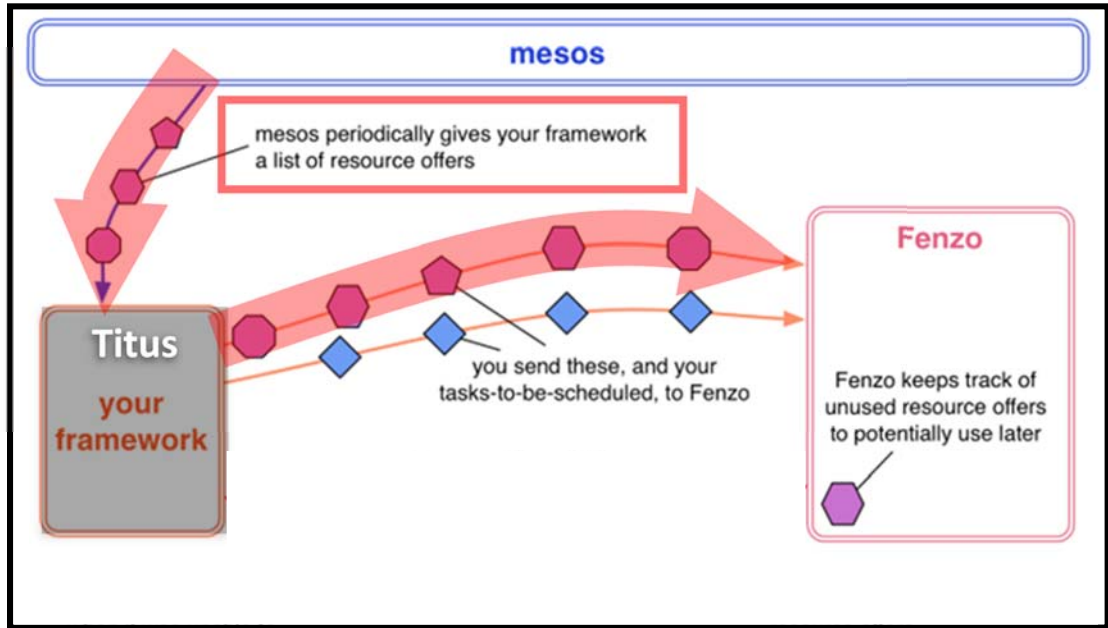
15 The Fenzo scheduler selects potential hosts [i.e., “Titus
16 Agents”] to assign tasks to based on whether the hosts
17 have sufficient resources to successfully execute the tasks.
18 Each task has its own resource requirements, each host
19 offer lists its available resources, and Fenzo maintains an
20 accounting of which resources are available and which are
21 already assigned on each host.¹⁹⁹

22 383. Netflix also practices the step wherein the one or more processors are
23 operable to “provide a list of available computer devices, the list comprising a
24 plurality of computer devices currently provisioned to perform computer
operations.” For example, Netflix explains that “Fenzo takes as its input the
resource offers that a framework [such as Titus] receives from Mesos and the tasks
that are given to the framework.”²⁰⁰

¹⁹⁸ See <https://netflix.github.io/titus/overview/>.

¹⁹⁹ See <https://github.com/Netflix/Fenzo/wiki/Introduction>.

²⁰⁰ See <https://github.com/Netflix/Fenzo/wiki>.



Source: <https://github.com/Netflix/Fenzo/wiki> (annotations added).

384. Netflix also practices the step wherein the one or more processors are operable to, “from the list of available computer devices, determine a list of suitable computer devices for performing the job by comparing operational characteristics for each available computer device with the job criteria, the list of suitable computer devices comprising one or more computer devices having operational characteristics that satisfy the job criteria.” For example, Mesos allows the Titus platform to set filters to reject resource offers that do not satisfy the desired operational characteristics in order to present a list of suitable resource offers.²⁰¹ For instance, a filter may include a constraint such that Mesos will “only offer nodes with at least R resources free” to the Titus Master.²⁰² Notably, after Mesos applies such filters to the resource offers, “the frameworks [i.e., the Titus/Fenzo

²⁰¹ Hindman, Benjamin, et al., *Mesos: A Platform for Fine-Grained Resource Sharing in the Data Center*, NSDI, (2011) at p. 3-4 (available at https://www.usenix.org/legacy/events/nsdi11/tech/full_papers/Hindman.pdf).

²⁰² *Id.*

1 architecture] still ha[s] the ultimate control to reject any resources that [it] cannot
2 express filters for and to choose which tasks to run on each node.”²⁰³

3 As another example, the Fenzo scheduler itself applies multiple iterative
4 constraints on the resource offers provided by Mesos via the Titus Master. Fenzo
5 includes various built-in constraints that compare operational characteristics (such
6 as available CPUs, disk space, memory, bandwidth, etc.) for each host to the job
7 requirements in order to identify suitable resource offers.²⁰⁴ As Netflix explains, a
8 constraint evaluator inspects a target and applies any hard constraints (i.e., those
9 that require “a simple Boolean thumbs-up or thumbs-down”) to identify a list of
10 suitable resource offers.²⁰⁵ After applying these hard constraints, the Fenzo
11 scheduler applies any soft constraints and uses any fitness calculators to further
12 narrow the list of suitable resources.²⁰⁶

13 385. Netflix also practices the step wherein the one or more processors are
14 operable to “use the utilization criteria to determine whether one or more
15 underutilized computer devices exist on the list of suitable computer devices, the
16 one or more underutilized computer devices having a suitable level of utilization for
17 performing the job.” As previously discussed, the Fenzo scheduler is configured to
18
19

20 ²⁰³ *Id.* at p. 4.

21 ²⁰⁴ See <https://github.com/Netflix/Fenzo/wiki/Constraints>; see also
<https://github.com/Netflix/Fenzo/wiki/Insights>.

22 ²⁰⁵ *Id.*; see also, [http://netflix.github.io/Fenzo/fenzo-
core/com/netflix/fenzo/ConstraintEvaluator.html](http://netflix.github.io/Fenzo/fenzo-core/com/netflix/fenzo/ConstraintEvaluator.html).

23 ²⁰⁶ See <https://github.com/Netflix/Fenzo/wiki/Fitness-Calculators> (explaining that
24 the “VMTaskFitnessCalculator” interface provided by Fenzo “does not have to
check to see that the proposed host has sufficient resources for the proposed task. It
can assume that this has already been done.”).

1 apply various hard and/or soft constraints to determine the “best fit” from the list of
2 suitable resource offers.²⁰⁷ As Netflix explains:

3 The Fenzo scheduler selects potential hosts to assign tasks
4 to based on whether the hosts have sufficient resources to
5 successfully execute the tasks. . . . But frequently there
6 are many hosts that meet these baseline qualifications.
7 You can fine-tune how the Fenzo scheduler assigns tasks
8 to compatible hosts by using fitness calculators and
9 constraints.²⁰⁸

10 One such constraint is the “Balanced Host Attribute Constraint” that is built-in to
11 Fenzo.²⁰⁹ The Balanced Host Attribute Constraint “schedules tasks so that they are
12 distributed evenly among types of hosts.”²¹⁰ As Netflix explains, when applied as a
13 soft constraint, the Balanced Host Attribute “weighs how close a host group would
14 be to having the average number of tasks if a new task were assigned to it” in order
15 to identify any underutilized resources.²¹¹ Applying this constraint causes an
16 underutilized resource to return a higher “fitness” rating for the new task.

17 386. Finally, Netflix practices the step wherein the one or more processors
18 are operable to “if the one or more underutilized computer devices exist, identify
19 the one or more underutilized computer devices to the client machine.” For
20 example, after Fenzo matches a task request with a resource offer, it provides its
21 scheduling recommendations to the Titus Master.²¹² As Netflix explains, the

22 ²⁰⁷ <https://github.com/Netflix/Fenzo/wiki/Introduction>.

23 ²⁰⁸ *Id.*

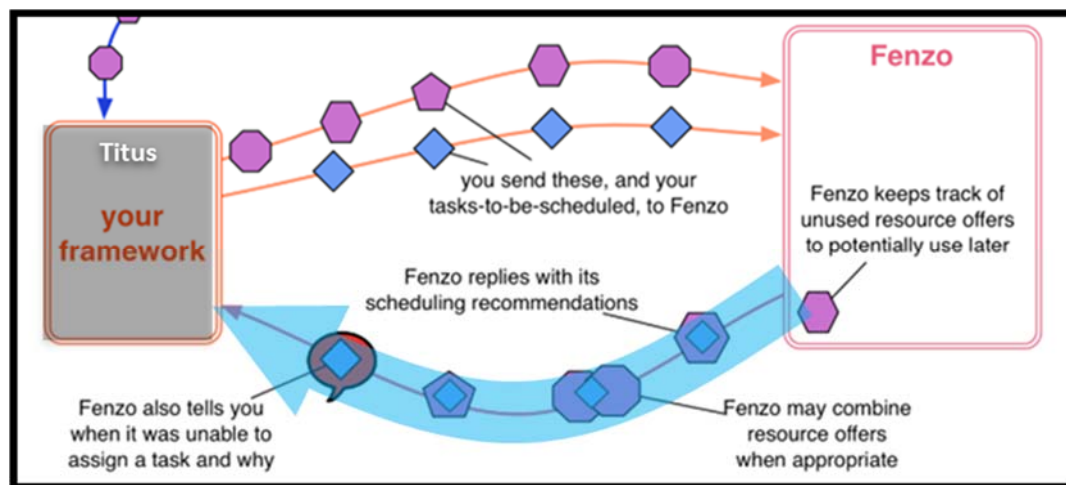
24 ²⁰⁹ *Id.*

²¹⁰ *Id.*

²¹¹ *Id.*

²¹² *See* <https://netflix.github.io/titus/overview/>.

1 Master then “schedules tasks onto Titus agents that launch containers based on the
2 task’s job specification.”²¹³



3
4
5
6
7
8
9
10 Source: <https://github.com/Netflix/Fenzo/wiki> (annotations added).

11 387. With regard to claim 12 of the '183 Patent, Netflix practices “[t]he
12 system of claim 11, wherein the one or more processors are operable to use the
13 utilization criteria to identify one or more underutilized computer devices by, for
14 each computer device on the list of suitable computer devices: determining a
15 utilization value.” For example, the Fenzo scheduler provides a built-in “Host
16 Attribute Value Constraint” that “matches a task with a host that has a particular
17 value [i.e., a utilization value] for a particular host attribute [such as CPU or disk
18 utilization].”²¹⁴

19 388. Netflix practices the system of claim 11, further comprising
20 “comparing the utilization value with the utilization criteria.” For example, as
21 Netflix explains,

22 When [the Host Attribute Value Constraint] evaluates a
23 host to see if it is appropriate for a task, it calls that

24 ²¹³ *Id.*

²¹⁴ <https://github.com/Netflix/Fenzo/wiki/Constraints>.

1 function and then checks the value it returns against the
2 value of the attribute on the host. If they match, the
3 constraint is met; otherwise (or if the host does not have
such an attribute) the constraint is not met.²¹⁵

4 389. Finally, Netflix practices the system of claim 11, further comprising
5 “removing the each device from the list of suitable computer devices if the
6 utilization value does not satisfy the utilization criteria.” For example, as Netflix
7 explains:

8 Fenzo will apply each constraint in this list to the hosts it
9 evaluates for the task. If any of these constraints fail,
10 Fenzo will disqualify that host from hosting the task and
will move on to evaluating the next host.²¹⁶

11 390. With regard to claim 15 of the '183 Patent, Netflix practices “[t]he
12 method of claim 11, wherein the one or more processors are configured to: use the
13 utilization criteria to determine whether one or more underutilized computer
14 devices exist on the list of suitable computer devices by identifying a computer
15 device having a lowest level of activity relative to other computer devices on the
16 list of suitable computer devices.” For example, as previously discussed, the
17 Balanced Host Attribute Constraint used by Fenzo “schedules a set of tasks so that
18 they are distributed evenly among types of hosts.”²¹⁷ As Netflix explains, “this
19 constraint attempts to minimize the difference between the number of co-tasks
20 running on the host type with the *most* co-tasks running on it and the number of co-
21 tasks running on the host type with the *least* co-tasks on it.”²¹⁸ This constraint may

22
23 ²¹⁵ *Id.*

²¹⁶ *Id.*

²¹⁷ *Id.*

²¹⁸ *Id.*

1 be used to identify the host having the lowest level of activity relative to other
2 suitable computer devices. As Netflix explains, the Balanced Host Attribute
3 Constraint identifies underutilized hosts by assigning each host a fitness score
4 between 0.0 (not a fit) to 1.0 (perfect fit).²¹⁹

5 391. Finally, Netflix practices the system of claim 11 wherein the one or
6 more processors are configured to “identify the computer device having the lowest
7 level of activity to the client machine.” For example, as discussed above, the Best
8 Host Attribute Constraint will assign the highest fitness score to the resource with
9 the fewest tasks. After the Fenzo scheduler matches a task request with a resource
10 offer that complies with any applicable constraints (including, for example, the
11 “Balanced Host Attribute Constraint”), it provides a recommendation to the Titus
12 Master.²²⁰ As Netflix explains, the Master then “schedules tasks onto Titus agents
13 that launch containers based on the task’s job specification.”²²¹

14 392. Upon information and belief, Netflix directly infringes other claims of
15 the ’183 Patent as well, including for the reasons discussed in the preceding
16 paragraphs.

17 393. At least as of the filing and service of this Amended Complaint,
18 Netflix has had knowledge of the ’183 Patent and Netflix’s infringement thereof.

19 394. Netflix’s infringement of the ’183 Patent, which is knowing and
20 willful at least as of the filing of this Amended Complaint, has caused and
21 continues to cause damage to Avago. Avago is entitled to recover damages
22 sustained as a result of Netflix’s wrongful acts in an amount subject to proof at trial.

23 _____
²¹⁹ *Id.*

24 ²²⁰ See <https://netflix.github.io/titus/overview/>.

²²¹ *Id.*

1
2 **PRAYER FOR RELIEF**

3 WHEREFORE, the Broadcom Entities respectfully request that the Court
4 enter a judgment in their favor and against Netflix:

5 1. Declaring that Netflix has directly infringed one or more claims of the
6 Patents-in-Suit in violation of 35 U.S.C. § 271;

7 2. Declaring that Netflix has induced infringement of one or more claims
8 of the '079, '121, '245, and '992 Patents in violation of 35 U.S.C. § 271(b);

9 3. Declaring that Netflix's infringement of the '079, '121, '245, '992,
10 '138, '387, '663, '283, '976, '722, and '183 Patents is willful and deliberate
11 pursuant to 35 U.S.C. § 284;

12 4. Enjoining Netflix from further infringing the '079, '121, '245, '992,
13 '138, '976, '722, and '183 Patents;

14 5. Ordering that the Broadcom Entities be awarded damages in an
15 amount no less than a reasonable royalty for each asserted patent arising out of
16 Netflix's infringement of the Patents-in-Suit, together with any other monetary
17 amounts recoverable, such as treble damages;

18 6. Declaring that this is an exceptional case under 35 U.S.C. § 285 and
19 awarding the Broadcom Entities their attorneys' fees and costs;

20 7. Ordering that Netflix is required to pay exemplary damages pursuant
21 to 35 U.S.C. § 284;

22 8. Awarding pre-judgment and post-judgment interest and costs against
23 Netflix; and
24

