

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

CISCO SYSTEMS, INC.
Petitioner,

v.

UMBRA TECHNOLOGIES, LTD.
Patent Owner.

U.S. Patent No. 11,108,595
Issue Date: August 31, 2021
Title: SYSTEMS AND METHODS FOR PROVIDING A GLOBAL VIRTUAL
NETWORK (GVN)

Inter Partes Review No.: IPR2024-00270

**PETITION FOR *INTER PARTES* REVIEW OF U.S. PATENT NO.
11,108,595 UNDER 35 U.S.C. §§ 311-319 and 37 C.F.R. §§ 42.1-.80, 42.100-
.107**

Mail Stop “PATENT BOARD”
Patent Trial and Appeal Board
U.S. Patent and Trademark Office
P.O. Box 1450
Alexandria, VA 22313-1450

TABLE OF CONTENTS

I.	INTRODUCTION	1
II.	OVERVIEW.....	1
III.	GROUND FOR STANDING (37 C.F.R. § 42.104(A)).....	6
IV.	REASONS FOR THE REQUESTED RELIEF.....	6
V.	BACKGROUND	6
A.	Summary of the '595 Patent.....	6
B.	Prosecution History	13
C.	Claim Construction.....	13
D.	Priority Date of the Challenged Claims	14
E.	Person of Ordinary Skill in the Art	14
F.	State of the Art	15
1.	Communication Tunnels.....	15
2.	Selection of Communication Paths By Controller.....	17
VI.	IDENTIFICATION OF CHALLENGES.....	18
A.	Challenged Claims	18
B.	Statutory Grounds for Challenges	18
VII.	Identification of How the Challenged Claims Are Unpatentable	19
A.	Ground 1: Claims 1-14.....	19
1.	The Kommula Reference	19
2.	Detailed Application of Kommula to the Challenged Claims ..	22
a.	Claim 1.....	22
[1.0]	A method comprising:.....	22

[1.1]	receiving, by one or more computer processors, a request for a list of available servers from a network device;	23
[1.2]	retrieving, by the one or more computer processors, the list of available servers;	24
[1.3]	retrieving, by the one or more computer processors, one or more records associated with the network device;	24
[1.4]	dynamically ordering, by the one or more computer processors, the list of available servers based on the one or more records associated with the network device; and.....	26
[1.5]	transmitting, by the one or more computer processors, the ordered list of available servers to the network device.....	28
b.	Claim 2.....	28
c.	Claim 3.....	29
d.	Claim 4.....	29
e.	Claim 5.....	30
f.	Claim 6.....	32
g.	Claim 7.....	32
h.	Claim 8.....	33
[8.0]	A system composing:	33
[8.1]	a memory storing instructions; and one or more processors coupled to the memory, the one or more processors being configured to execute the instructions, the instructions when executed causing the one or more processors to perform operations comprising:	34
[8.2]	receiving a request for a list of available servers from a network device;	34
[8.3]	retrieving the list of available servers;	34

[8.4]	retrieving one or more records associated with the network device;	35
[8.5]	dynamically ordering the list of available servers based on the one or more records associated with the network device; and	35
[8.6]	transmitting the ordered list of available servers to the network device.	35
i.	Claim 9.....	35
j.	Claim 10.....	35
k.	Claim 11.....	35
l.	Claim 12.....	36
m.	Claim 13.....	36
n.	Claim 14.....	36
B.	Ground 2 –Claims 15-20	37
1.	The Hankins Reference.....	37
2.	Detailed Application of Hankins, alone or in view of Kommula to the Challenged Claims	40
a.	Claim 15.....	40
[15.0]	A network system for managing a global virtual network, comprising:.....	41
[15.1]	an access point server configured to connect to one or more endpoint devices via one or more first tunnels; and	41
	Hankins discloses that the tunnel server (red) (access point server) is configured to connect one or more gateways (green) (endpoint devices) via one or more tunnels (purple). Hankins discloses that a controller may provision one or more tunnel servers to provide access to different broadcast domains and facilitate communications between network devices communicating as part of different broadcast domains. <i>EX-1004</i> 8:63-66; 26:65-27:8. This is consistent with the disclosure of the '595 Patent	

where gateways are the endpoints of a tunnel connecting the gateways, and the gateways are the edge of their respective networks. *EX-1001* FIGs. 9, 15; *EX-1003* ¶ 130. 41

[15.2] a control server connected to the access point server via a second tunnel, the control server storing one or more records associated with each of the one or more endpoint devices,.....43

[15.3] wherein the control server is configured to receive a request for a list of available servers from a first endpoint device among the one or more endpoint devices, the list of available servers including the access point server; 45

[15.4] retrieve the list of available servers; retrieve the one or more records associated with the first endpoint device 47

[15.5] dynamically order the list of available servers for the first endpoint device based on the one or more records associated with the network device; and 48

[15.7] transmit the ordered list of available servers to the first endpoint device. 52

b. Claim 17.....53

c. Claim 18.....53

d. Claim 19.....53

e. Claim 20.....54

**VIII. THE BOARD SHOULD NOT EXERCISE ITS DISCRETION
AND DENY INSTITUTION.....54**

**A. The Board Should Not Deny Institution Under 35
U.S.C. § 314(a) 54**

IX. Mandatory Notices Under 37 C.F.R. §42.8 57

A. Real Party-in-Interest (37 C.F.R. § 42.8(b)(1)).....57

B. Related Matters (37 C.F.R. § 42.8(b)(2)).....57

1. Judicial Matters 57

3.	Related Patents.....	57
C.	Lead/Back-up Counsel (37 C.F.R. § 42.8(b)(3)):	58
D.	Notice of Service Information (37 C.F.R. § 42.8(b)(4)):	58
X.	CONCLUSION	58

Petitioner's Exhibit List

<i>Exhibit #</i>	<i>Description</i>
1001	U.S. Patent No. 11,108,595 (the '595 Patent)
1002	Prosecution history of the '595 (“’595 File History”)
1003	Declaration of Petitioner’s Expert Dr. A.L. Narasimha Reddy(“Reddy”)
1004	U.S. Pat. Appl. Publ. No. 2010/0153558 to Kommula (“Kommula”)
1005	U.S. Pat. No. 10,177,957 (“Hankins”)
1006	U.S. Pat. App. Publ. No. 2010/0275017 to Iyer (“Iyer”)
1007	Related Patents

I. INTRODUCTION

Pursuant to 35 U.S.C. §§ 311 *et seq.* and 37 C.F.R. §§ 42.1 *et seq.*, Cisco System, Inc. (“Petitioner”) hereby petitions for an *inter partes* review of U.S. Patent No. 11,108,595 (“the 595’ Patent”). Petitioner respectfully submits that Claims 1-20 (the “Challenged Claims”) of the ’595 Patent are unpatentable under 35 U.S.C. §103 in view of the prior art references discussed herein. This Petition demonstrates by a preponderance of the evidence that there is a reasonable likelihood that Petitioner will prevail with respect to at least one of these claims. Accordingly, it is respectfully requested that the Board institute an *inter partes* review of the ’595 Patent pursuant to 37 C.F.R. § 42.108.


II. OVERVIEW

The Challenged Claims are unpatentable as obvious over the prior art. The claims are directed to prior art systems and combinations of conventional components that perform conventional functions regarding managing virtual network connections, all of which was well-known in the art of virtual networks.

In one aspect, some of the claims of the ’595 Patent are directed to a high level method performed by a computer processor including the steps of (1) “receiving” a request for a list of available servers from a network device; (2) “retrieving” the list of available servers; (3) “retrieving” one or more records

associated with the network device; (4) “dynamically ordering” the list of available servers based on the one or more records; and (5) “transmitting” the ordered list of available servers to the network device, all as recited in Claim 1. *EX-1001* 32:49-61. Independent Claim 8 is directed to a generic system with memory and a processor configured to perform the same steps recited in Claim 1. *EX-1001* 33:14-31; *EX-1003* ¶28.

These high level steps are described by the prior art reference Kommula as being performed by the architecture shown in FIG. 1:




EX-1004 FIG. 1; *EX-1003* ¶ 29.

Kommula describes that the one or more processors performs each of the recited steps including providing the same example of dynamically ordering the list of available servers as described in the '595 Patent. Specifically, Kommula

describes that the list of available servers can be dynamically ordered based on the proximity of the server to the network device, just like described in the '595 Patent. *EX-1003* ¶ 30.

In another aspect, some of the claims of the '595 Patent are directed to a system in which a controller dynamically orders a list of available access point servers, which can be used to form tunnels between first and second endpoint devices. For example, FIG. 35 of the '595 Patent illustrates the network system recited in Claim 15 including a first endpoint device (green) communicating with a second endpoint device (blue) via a first tunnel (purple) connecting an access point server (red) to one or more endpoint device, and a controller (orange) connected to an access point server via a second tunnel (brown) for dynamically ordering a list of available servers for the first endpoint device:



EX-1005 FIG. 2 (annotated); EX-1003 ¶ 32.

FIG. 2 of Hankins illustrates the network system recited in Claim 15 including a first endpoint device (green) communicating with a second endpoint device (blue) via a first tunnel (purple) connecting an access point server (red) to one or more endpoint device, and a controller (orange) connected to an access point server via a second tunnel (brown) for dynamically ordering a list of available servers for the first endpoint device. EX-1003 ¶ 33.

Thus, the Challenged Claims of the '595 Patent are directed to systems and methods that were fully disclosed in the Kommula and Hankins references.

III. GROUNDS FOR STANDING (37 C.F.R. § 42.104(A))

Petitioner certifies that the Patent for which review is sought is available for post grant review and that Petitioner is not barred or estopped from requesting a post grant review of the Challenged Claims on the grounds identified herein. 37 C.F.R. § 42.104(a). This Petition is filed pursuant to 37 C.F.R. § 42.106(a).

IV. REASONS FOR THE REQUESTED RELIEF

As explained in §§ II, V and VII of this Petition and in the attached Declaration of Petitioner's Expert, Dr. A.L. Narasimha Reddy ("Reddy"), *EX-1003*, the systems and methods of managing virtual network connections as described and claimed in the 595' Patent, are obvious over the prior art to a person of ordinary skill in the art ("POSITA") at the time of the invention.

V. BACKGROUND

A. Summary of the '595 Patent

The '595 Patent is directed towards managing a network connection between an endpoint device and an access point server. The '595 Patent discloses as a problem that in the prior art there is insufficient control over the routes between an endpoint device and access point server as it relies on the policies of various middle players on the internet. *EX-1001* 2:18-36.; *EX-1003* ¶ 35.

The '595 Patent describes that it is not uncommon to connect from within one local area network (LAN) to the resources of another remote LAN using a direct connection through a wide area network (WAN), such as the internet.

However, the '595 Patent describes a problem that arises when a gateway to the Internet is located at a LAN, for example, located at the corporate headquarters, which handles traffic from a remote LAN of a subsidiary located in a different country. Traffic from the remote LAN is routed through the gateway at the headquarter LAN to the Internet. A slowdown occurs as traffic flows through the Internet back to servers in the same country as the subsidiary. Traffic must then go from the remote LAN through the WAN to the headquarters LAN and then through the Internet back to the server used by the subsidiary in the different country, then back through the Internet to the gateway of the headquarters, and then back down the dedicated line to the client device within the remote LAN. This type of connection in essence doubles or triples (or worse) the global transit time of what should take a small fraction of global latency to access a server close to the remote LAN. Thus, the prior art connection mechanisms may not take into account the context of the connection needed for a particular source endpoint device and available access point servers. *EX-1001* 1:59-2:17; *EX-1003* ¶ 36.

The '595 Patent then describes a number of specific cloud based solutions for establishing dedicated connections between endpoint devices and an access point server. *EX-1001* 2:40-57. But, the specificity of these solutions are found nowhere in the claims, which are much broader and directed to well-known

Internet connections and/or well-known network tunneling technology; *EX-1003* ¶38.

As examples, in one embodiment, the network system may include an endpoint device, an access point server, and a control server. The endpoint device and the access point server may be connected with a first tunnel, and the access point server and the control server may be connected with a second tunnel. *EX-1001* Abstract; 3:4-11; *EX-1003* ¶39.

The '595 Patent describes that in one embodiment, in order to build a tunnel, certain information is needed. *EX-1001* 10:65-11:3. For example, a server availability module examines current allocation of IP address and port combinations while taking account actual use, load, and other factors of the live production environment. It further analyzes the potential impact of assigned but standby allocations to predict the impact of those resource demands going live. When a device makes a request to the server availability module, what it is really asking for is a list of available servers, and specifically an IP and port which it could utilize to in the present or in the future to build a tunnel to the available server. Each tunnel has specific information which needs to be known to each peer pair in the tunnel building process as defined by tables under tunnel information. *EX-1001* 23:11-24; *EX-1003* ¶ 41.

The list of server availability is a list which can be utilized by all devices but for which a contextual list for one certain device is easily retrieved with contextual records relevant to it. *EX-1001* 23:33-36. The server availability module considers many factors balancing demand for resources by clients such as endpoint devices wanting to build tunnels against available resources of server access points which they would connect with. *EX-1001* 23:40-43; *EX-1003* ¶43.

The '595 Patent describes the use of a governing mechanism that ranks the available sever access points in order contextually for that specific endpoint device. The ranked list takes into account a holistic view and predictively spreading load in anticipation of demand pressures on the production environment. *EX-1001* 23:55-59; *EX-1003* ¶ 44.

This governing mechanism ensures that automatic jumps of tunnels from many endpoint devices to one specific device do not occur. For example, if all endpoint device share the same list of server availability, without an influence on how an endpoint device picks which servers to connect with, there are scenarios where too many endpoint devices connect with some access point servers and too few to others. And when the endpoint devices connected to the oversaturated access point server realize that they are receiving less than optimal connectivity, if they share the same list, by jumping to the next server on the list, they will be

causing a wobbly shifting of the problem down the line. *EX-1001* 23:44-54; *EX-1003* ¶ 45.

In one embodiment, server availability prioritizes the list of server access port IP address and port combinations based on expected best performance of the tunnels to be built while also looking at current load of available access point servers, balancing assigned lists given to other end point devices as well as other available information. *EX-1001* 11:21-26; *EX-1003* ¶ 46.

The server availability list is generated for each request received from an endpoint device and the list is shared with the endpoint devices and available access point servers *EX-1001* 12:50-59; *EX-1003* ¶ 47.

The '595 Patent describes that the server availability module's main benefit is to catalog the resources usage, IP addresses and ports of access point servers and to look at which devices, such as end point devices that would like to connect to server access points to build tunnels with them. The server availability module calculates and compiles a custom-made list of IP addresses and ports which are contextually designed to provide best options for endpoint devices to utilize. The contextual information includes connectivity tests based on current and historical conditions to weigh and rank best connectivity options for tunnels to be built. *EX-1001* 27:20-31; *EX-1003* ¶ 48.

Although the '595 Patent describes very specific embodiments of generating available access point servers to build tunnels between endpoint devices, the Challenged Claims are not so limited. For example, independent Claims 1 and 8 are method and system, respectively, broadly directed to dynamically ordering a list of available servers based on a record associated with a network device. They do not recite that the server is an access point server, do not recite selecting a server to connect to the network device or provisioning a tunnel and connecting the network device to the tunnel, nor do they recite that the recited records include contextual information regarding the network device. Claim 1 is representative:

1. A method comprising:
 - receiving, by one or more computer processors, a request for a list of available servers from a network device;
 - retrieving, by the one or more computer processors, the list of available servers;
 - retrieving, by the one or more computer processors, one or more records associated with the network device;
 - dynamically ordering, by the one or more computer processors, the list of available servers based on the one or more records associated with the network device; and
 - transmitting, by the one or more computer processors, the ordered list of available servers to the network device.

Dependent claims 2-14 add limitations regarding the server being an access point server, and the dynamic ordering being based on an anticipated demand of a server, but do not recite selecting a server to connect to the network device or provisioning a tunnel and connecting the network device to the tunnel.

Independent Claim 15, and dependent claims 16-20 are directed to a controller dynamically ordering a list of available access point servers based on records associated with an endpoint device, where the access point server is configured to connect one or more endpoints via a tunnel. Claim 15 recites:

15. A network system for managing a global virtual network, comprising:
 - an access point server configured to connect to one or more endpoint devices via one or more first tunnels; and
 - a control server connected to the access point server via a second tunnel, the control server storing one or more records associated with each of the one or more endpoint devices, wherein the control server is configured to:
 - receive a request for a list of available servers from a first endpoint device among the one or more endpoint devices, the list of available servers including the access point server;
 - retrieve the list of available servers;
 - retrieve the one or more records associated with the first endpoint device;

dynamically order the list of available servers for the first endpoint device based on the one or more records associated with the network device; and

transmit the ordered list of available servers to the first endpoint device.

B. Prosecution History

The Challenged Claims received a first Office Action allowance, subsequent the filing of a terminal disclaimer with respect to U.S. Patent No. 10,756,929, in which the examiner issued a general statement that the limitations recited in the independent claims are not found in the prior art. *EX-1002* pp. 9, 16, 40-41.

C. Claim Construction

Petitioner proposes that each claim term in the Challenged Claims be given its plain and ordinary meaning in this proceeding, and that no specific construction of any claim term is required because the prior art relied on in this Petition meets each of the claim terms under any reasonable construction.

It is noted that this interpretation is only applicable to the *inter partes* review sought herein and should not be construed as constituting, in whole or in part, the Petitioner's own interpretation of any claims for any other purposes, including any litigation. Accordingly, Petitioners expressly reserves the right to present an interpretation of a claim term in other proceedings, which is different, in whole or in part, from that presented in this Petition.

D. Priority Date of the Challenged Claims

The application that issued as the '595 Patent was filed August 24, 2020 as Application Ser. No. 17/000,997 as a continuation of Application No. 15/563,253 filed on September 29, 2017 as a U.S. national stage application of application PCT/US2016/026489 on April 7, 2016 claiming priority to provisional application No. 62/151,174 filed on April 22, 2015 and to provisional application No. 62/144,293 filed on April 7, 2015. Patent Owner bears the burden of establishing priority to the provisional applications.

The prior art cited below predates the earliest possible priority date, April 7, 2015, and thus for purposes of this Petition, this issue is moot.

E. Person of Ordinary Skill in the Art

With respect to the '595 Patent, a POSITA in the April 2015 timeframe would have been familiar with using cloud based resources offered on top of regular Internet connections to provide virtual networks and the formation of tunnels and the associated protocols as demonstrated by the prior art references and description in the State of the Art below. *EX-1003* ¶ 55. A POSITA would have gained knowledge of these concepts through a mixture of training and work experience, such as by having a Bachelor's degree in computer science and two years of experience; or by obtaining a Master's degree in computer science, but


having only one year of experience; or by having no formal education but experience in storage systems of at least four years. *EX-1003* ¶ 56.

F. State of the Art

The following section describes the state of the art for using cloud based resources on top of regular Internet connections to provide virtual networks and the formation of tunnels and the associated protocols build tunnels between an end point device and an access point server as of April 2015. *EX-1003* ¶ 58. The prior art references, and the discussions of what was known to a POSITA, provide the factual support for the general description of the state of the art at the time of the invention, provide additional motivation to modify or combine the reference. Accordingly, these references should be considered by the Board.

1. Communication Tunnels

The use of “tunnels” was well known in the art. For example, Iyer (2010) disclosed a plurality of communication paths (tunnels 1 and 2) between wireless clients, which were identified as “peer-to-peer tunnels” established between “any two clients in the mobility domain” to route traffic between the clients. Traffic between the clients pass through the tunnels and intermediate access point servers. *EX-1006* [0015], [0030]; *EX-1003* ¶ 60.



EX-1006 Fig. 4; *EX-1003* ¶ 60.

The two communication paths connecting wireless clients 500a and 500b are “form[ed]” by access nodes 400a and 400b, because access nodes 400a and 400b “establish” the tunnel 610, or tunnels 600a and 600b.

[M]obility manager 280 instructs access node 400a to establish peer-to-peer tunnel 610 between access node 400a and access 400b.

EX-1006 [0030]; *EX-1003* ¶ 61.

In addition, the access nodes 400a,b can also be connected to controller 200 via tunnels:

[T]he access nodes establish communications with the controller using tunnels.

EX-1006 [0013]; *EX-1003* ¶ 62.

2. Selection of Communication Paths By Controller

It was well known in the art to provide a controller that selected the communication path, including the selection of tunnels, between client devices. For example, Iyer discloses that the Controller 200 identifies the traffic destination and selects the tunnel through access nodes that the traffic uses. EX-1006, [0025 - 26]; *EX-1003* ¶ 63.

Access nodes 400a, 400b, 400c, 400d provide wireless and possibly wired services to clients. As shown in FIG. 1, wireless client 500a is connected to access node 400a. Access node 400b supports wireless client 500b and wired client 510b. Access node supports wireless client 500c. EX-1006 [0018]. EX-1003 ¶ 64 Iyer describes a controller establishes peer-to-peer tunnels between access nodes to route that traffic between the clients. EX-1006 [0030]; *EX-1003* ¶ 65.

This routing is performed by controller 200 using the IP addresses of clients 500a and 500b, as well as the MAC (media access controller) addresses of clients 500a, 500b and access nodes 400a and 400b.

EX-1006 [0026]; *EX-1003* ¶ 65.

When client 500a wishes to send data to client 500b, it in essence forms an IP packet with client 500b's IP address as the destination,

and with client 500a's IP address and MAC address as the source.
This information is encapsulated and sent to controller 200.

EX-1006 [0026]; *EX-1003* ¶ 65.

Controller 200 identifies the traffic destination as client 500b, and sends the traffic though tunnel 600b to access node 400b and client 500b.

Ex-1006 [0026]; *EX-1003* ¶ 65.

VI. IDENTIFICATION OF CHALLENGES

A. Challenged Claims

Claims 1-20 of the '595 Patent are challenged in this Petition.

B. Statutory Grounds for Challenges

The Challenges are set forth in detail below and summarized as follows:

Ground	Claims	Basis	Reference
1	1-14	§ 103	Kommula
2	15-20	§ 103	Hankins alone, or in view of Kommula

Ground 1: Claims 1-14 of the '595 Patent are invalid as being obvious over Kommula.

U.S. Patent Publication No. 2010/0153558 is titled "Global Server Load Balancing" and it published on June 17, 2010 from Application Ser. No. 12/635,371 filed on December 10, 2009 ("Kommula") (*EX-1004*). Kommula qualifies as prior art under AIA §102(a)(1). Kommula was not considered by the

examiner during prosecution of the '595 Patent, nor was it previously considered by the Board.

Ground 2: Claims 15-20 of the '595 Patent are invalid as being obvious over Hankins, alone, or in view of Kommula.

U.S. Patent No. 10,177,957 is titled “Connecting a Cloud Network to the Internet” and it issued on January 8, 2019 from Application Ser. No. 13/802,529 filed on March 13, 2013 (“Hankins”) (*EX-1005*). Hankins qualifies as prior art under AIA §102(a)(2). Hankins was not considered by the examiner during prosecution of the '595 Patent, nor was it previously considered by the Board.

VII. IDENTIFICATION OF HOW THE CHALLENGED CLAIMS ARE UNPATENTABLE

A. Ground 1: Claims 1-14


1. The Kommula Reference

Kommula discloses the use of a global server for load balancing cloud resources and communicates with numerous site switches that are coupled to host servers serving specific applications. *EX-1005* Abstract; *EX-1003* ¶ 68.

Kommula identifies the same contextual problem in the prior art identified in the '595 Patent—that the prior art global server does not take into consideration the host server’s locations relative to the client. *EX-1004* [0007]. Like the '595 Patent, Kommula recognizes that the geographical distance between the server and a client is a factor affecting the response time for the client’s access to the host

server. *EX-1004* [0007]. Kommula provides the example that a client from Japan could receive better response time from a host server in Japan than from a host server in New York. Further, the conventional domain name server (DNS) algorithm allows invalid IP addresses (*e.g.*, that corresponding to a downed server) to persist in a local DNS server until the invalid IP address expires. *EX-1004* [0007]; *EX-1003* ¶ 69.

As illustrated in annotated FIG. 1 below, Kommula provides a solution to the problem by using a global server load balancing switch (GSLB)(green).



EX-1004 FIG. 1; *EX-1003* ¶ 70.

In one embodiment, the global server load balancing switch provides IP addresses of available host servers (red) to a client program (blue) which requested

access to an application hosted on one of the host servers. The available host services are based on a selected set of performance metrics. Each site switch (yellow) provides the GSLB switch with current site-specific information regarding the host servers (red) associated with the site switch. In one embodiment, when a domain name server (DNS) (purple) resolves a host name in a query and returns one or more IP addresses, the GSLB switch filters the IP addresses using performance metrics compiled from the site-specific information collected from the site switches and/or other metrics. The GSLB switch then returns a ranked or weighted list of IP addresses to the requesting computer program (blue) inquirer. In one embodiment, the IP address that is estimated to provide the best expected performance for the client is placed at the top of the list. *EX-1004* [0008]; *EX-1003* ¶ 71.

Kommula explains that examples of suitable performance metrics include availability metrics (*e.g.*, a server's or an application's health), load metrics (*e.g.*, a site switch's session capacity or a corresponding preset threshold), and proximity metrics (*e.g.*, a round-trip time between the site switch and a requesting DNS server, the geographic location of the host server, the topological distance between the host server and the client program). A topological distance is the number of hops between the server and the client. Another proximity metric is the site switch's "flashback" speed (*i.e.*, how quickly a switch receives a health check

result). Yet another metric is a connection-load metric that is based on a measure of new connections-per-second at a site. The ordered list can also be governed by other policies, such as the least selected host server. *EX-1004* [0009]; *EX-1003* ¶ 72.

2. Detailed Application of Kommula to the Challenged Claims

a. Claim 1

[1.0] A method comprising:

Kommula discloses a method for using a global server load balancing switch to resolve a DNS query by a client by providing an ordered list of servers that is expected to serve the client with a high performance in a given application. *EX-1004* [0003]. As annotated below, Kommula discloses the architecture used to perform the method recited in Claim 1.




Fig. 1

EX-1004 FIG. 1 (annotated)¹; EX-1003 ¶ 73.

[1.1] receiving, by one or more computer processors, a request for a list of available servers from a network device;

Kommula discloses receiving by the GSLB switch 12 and DNS 16 (*one or more computer processors*) (orange dashed line) a query for resources from a host server (red) which is a request for a list of IP addresses of host servers (*a request for a list of available servers*) from a client program on a client machine 28 (*network device*) (dashed blue line). EX-1004 [0003]; [0008]; [0026]; Annotated FIG. 1. Specifically, Kommula discloses that client program 28 is connected to Internet 14, and communicates with local DNS server 30 (purple). When a

¹ This figure will be referred to in the claims analysis below as Annotated FIG. 1:

browser on client program 28 requests a web page, for example, using a Universal Resource Locator (URL), a query is sent to local DNS server 30 to resolve the symbolic host name (*e.g.*, *www.foundrynet.com*) to an IP address of a host server. This query is either sent to local DNS server 30's cache to retrieve the list of IP addresses, or if the time to live (TTL) of the responsive IP addresses in the cache has expired, the query is sent to GSLB switch 12 (green) to retrieve the list of IP address from DNS 16. *EX-1004* [0019]. Kommula's disclosure is consistent with the disclosure of the '595 Patent that when a device makes a request to a server, what it is really asking for is a list of available servers. *EX-1001* 23:11-24; *EX-1003* ¶¶ 75- 76.

[1.2] retrieving, by the one or more computer processors, the list of available servers;

Kommula discloses that the GSLB switch in combination with DNS 16 (*one or more computer processors*) retrieves a list of IP addresses corresponding to the resolved host name. *EX-1004* [0019], [0027]; *EX-1003* ¶ 77.

[1.3] retrieving, by the one or more computer processors, one or more records associated with the network device;

Kommula discloses several examples of records of performance metrics associated with the network device that can be retrieved by the GSLB. For example, Kommula describes that a site switch (brown) measures the round trip time between the client machine and a host server and records it in an entry of the

cache database (*one or more records associated with the network device*).

Periodically, each site switch reports the RTT database to the GSLB switch along with load conditions (*retrieving, by the one or more computer processors, one or more records associated with the network device*). The GSLB switch aggregates the RTTs reported into a proximity table indexed by network neighborhood. *EX-1004* [0135]. More specifically, Kommula discloses that the RTT is computed for the interval between the time when a client machine requests a TCP connection to a proxy server configured on a site switch, sending the proxy server a TCP synchronization (TCP SYN) packet, and the time a site switch receives from the client program a TCP acknowledgment (TCP ACK) packet. (In response to the TCP SYN packet, a host server sends a TCP SYN ACK packet, to indicate acceptance of a TCP connection; the client machine returns a TCP ACK packet to complete the setting up of the TCP connection.) The GSLB switch 12 maintains a database of RTT, which it creates and updates from data received periodically from the site switches (e.g., site switches 18A, 18B, 22A and 22B). Each site collects and stores RTT data for each TCP connection established with a client machine. *EX-1004* [0095]; *EX-1003* ¶¶ 78- 79.

Kommula further describes that the GSLB switch can look up the RTT for a client machine (*network device*) to any specific host server, based on the client's network neighborhood specified in the client's IP address. *EX-1004* [0135]. Thus,

the GSLB (*one or more computer processors*) retrieves from the proximity table the RTT for a client machine to any specific host (*one or more records associated with the network device*). *EX-1003 ¶ 80.*

Further, Kommula discloses other examples of performance metrics associated with a network device that can be retrieved by the GSLB switch from a database or memory that is remote from the GSLB. *EX-1004 [0024]*. Examples of performance metrics associated with a network device include the topological distance between the host server and the client. A topological distance is the number of hops between the server and the client. *EX-1004 [0009]; EX-1003 ¶ 81.*

[1.4] dynamically ordering, by the one or more computer processors, the list of available servers based on the one or more records associated with the network device; and

Kommula discloses this limitation. Kommula discloses that when a client requests a resolution of a host server name, the GSLB switch (*one or more computer processors*) returns one or more ordered IP addresses for the host name (*dynamically ordering the list of available servers*). The IP addresses are ordered using metrics that include the information collected from the site switches or based on other metric information (*based on the one or more records associated with a network device*). The GSLB switch places the address that is deemed “best” at the top of the list (*dynamically ordering*). *EX-1004 Abstract; EX-1003 ¶ 82.*

As discussed in limitation [1.3], one of the performance metrics associated with the network device is RTT. Kommula discloses that the IP addresses can be reordered (*dynamically ordering*) by the GSLB switch (*one or more computer processors*) based upon a round-trip time (RTT) between the site switch for the IP address and the client (*one or more records associated with the network device*). Specifically, the RTT is computed for the interval between the time when a client machine requests a TCP connection to a proxy server configured on a site switch, sending the proxy server a TCP SYN packet, and the time a site switch receives from the client program a TCP ACK packet. (In response to the TCP SYN packet, a host server sends a TCP SYN ACK packet, to indicate acceptance of a TCP connection; the client machine returns a TCP ACK packet to complete the setting up of the TCP connection.) The GSLB switch 12 maintains a database of RTT, which it creates and updates from data received periodically from the site switches (e.g., site switches 18A, 18B, 22A and 22B). Each site collects and stores RTT data for each TCP connection established with a client machine. In one embodiment, the GSLB switch favors one host server over another only if the difference in their RTTs with a client machine is greater than a specified percentage, the default specified percentage value being 10%, for example (*dynamically ordering*). EX-1004 [0095]; EX-1003 ¶¶83- 85.

[1.5] transmitting, by the one or more computer processors, the ordered list of available servers to the network device.

Kommula discloses that the list of ordered IP addresses is transmitted by the GSLB switch to the client. *EX-1004* [0095]. More specifically, Kommula discloses that the list of IP addresses of host servers is ordered (*ordered list of available servers*) by GSLB switch 12 (*one or more computer processors*) based on metrics such as RTT and sent to the client 28 (*network device*). *EX-1004* [0019]; [0095]; *EX-1003* ¶¶ 86- 87.

b. Claim 2

The method of claim 1 wherein the network device is an endpoint device.

Kommula discloses that a client program 28 running on a client machine is connected to Internet 14, and communicates with local DNS server 30 in order to establish a TCP connection to a host server running an application. *EX-1004* [0005]; [0019]; [0095]. A POSITA understood that client machine 28 is an endpoint device. *EX-1003* ¶¶ 88- 89.

To the extent that the Patent Owner argues that the client machine is not an endpoint device, a POSITA understood that the client machine 28 would communicate through an endpoint device in order to access the Internet 14. *EX-1004* FIG. 1; *EX-1006* [0018]. In this alternative, the request from the client machine would go through (and thus be received from) the endpoint device (as

network device) and be returned to the endpoint through that device (as the network device). EX-1003 ¶¶90- 91.

c. **Claim 3**

The method of claim 2, wherein at least one server among the list of available servers is an access point server.

Kommula discloses that host server may host application server programs, and POSITA understood that the host server may provide access for a wired or wireless client to access the network, and thus act as an access point server EX-1004 [0005, 0017] FIG. 1 ; EX-1003 ¶ 92.

d. **Claim 4**

The method of claim 1, further comprising anticipating a level of demand for at least one server among the list of available servers, and wherein dynamically ordering the list of available servers is further based on the anticipated level of demand.

Kommula discloses that anticipating a level of demand at a server is one of the performance metrics that can be used by the GSLB switch to order the IP address of the host servers. EX-1003 ¶ 93.

In accordance with embodiments of the invention, a weighted site or a weighted IP metric can be used. The weighted site metric can be based on a weight that determines how much traffic a site will receive with respect to other sites in comparison, while the weighted IP can be based on a weight that determines how much traffic an IP address will receive with respect to other IP addresses in comparison. An

embodiment provides an active bindings metric that is based on a measure of the number of active real servers bound to a virtual IP address residing on a GSLB site. The active bindings metric bases its selection of the best IP address on this number and prefers a VIP with the highest number of active bindings.

EX-1004 [0010]; *EX-1003* ¶ 93.

Specifically, Kommula discloses that one performance metric is a measure of new connections-per-second at a site which is used to calculate an average load for a site. *EX-1004* [0098]-[0099]. A POSITA understood that this calculation of average load is used as an estimate of the anticipated load at the site, e.g., it will continue at the average rate, so that the list of IP addresses of the servers can be order based on the site's connection load. Thus, Kommula disclosing calculating an average load using new connections-per-second at the site (*anticipating a level of demand of at least one server*) and reordering the list of IP addresses of the servers based on the average load (*dynamically reordering the list of available servers based on the anticipated level of demand*). *EX-1003* ¶¶ 94- 95.

e. Claim 5

The method of claim 1, wherein dynamically ordering the list of available servers is further based on one or more second ordered lists of available servers transmitted to one or more second network devices.

This limitation is disclosed for the same reasons as disclosed with respect to limitations [1.3], [1.4] and [1.5] but for a request from a second network device. Kommula discloses that the GSLB switch will develop a second ordered list of the IP addresses (*dynamically ordering*) based upon a round-trip time (RTT) between a site switch for the IP addresses and the second client (*one second ordered lists of available servers*) and the GSLB switch will send the list to the second client (*transmitted to one or more second network devices*). *EX-1004* [0095], [0097]; *EX-1003* ¶ 96.

Further, because Kommula retrieves a performance metric for each site, such as a site switch's session capacity threshold, the first client's use of an IP address of a selected host server, may impact that site switch's session capacity which could impact the dynamic reordering of the second ordered list for the second client, and vice versa. *EX-1004* [0025]. Thus, the dynamic ordering of a list for a first network device would be based on the list of ordered available servers provided to a second network device. *EX-1003* ¶¶ 97- 98.

A POSITA understood that the same weighted list of servers cannot be provided to all the network devices for load balancing purposes. If a server 1 is returned to one network device, returning the address of a second server 2 to a second network device would improve load balancing. Hence a POSITA

understood that the ordered list returned to a network device would depend on responses to other network devices. *EX-1003* ¶ 99.

f. **Claim 6**

The method of claim 1, wherein dynamically ordering the list of available servers comprises generating ranking information for each server in the list of available servers.

Kommula discloses that in dynamically ordering as described in limitation [1.4] Kommula can generate a ranked or weighted list of IP addresses (*generating ranking information for each server in the list of available servers*):

Under one aspect of the present invention, when an authoritative DNS server resolves a host name in a query and returns one or more IP addresses, the GSLB switch filters the IP addresses using the performance metrics compiled from the site-specific information collected from the site switches and/ or other metrics. The GSLB switch then returns **a ranked or weighted list of IP addresses** to the inquirer. In one embodiment, the IP address that is estimated to provide the best expected performance for the client is placed at the top of the list.

EX-1004 [0008]; *EX-1003* ¶ 100.

Kommula also discloses ordering “availability metrics” and RTTs which are dynamic in nature. *EX-1004* [0009, 0025]; *EX-1003* ¶ 101.

g. **Claim 7**

The method of claim 1, wherein the one or more records are retrieved based on a unique device identity associated with the network device.

As discussed with respect to limitation [1.3], the GSLB switch periodically receives RTT between a client machine and any specific host server aggregates the RTTs reported into a proximity table indexed by network neighborhood. The GSLB switch can thus look up the RTT for a client machine to any specific host server, based on the client's network neighborhood specified in the client's IP address. *EX-1004* [0135]; *EX-1003* ¶ 102. A POSITA understood that in order to maintain a RTT for each client machine to each host server, a *unique device identity* would be assigned to each client machine in order to distinguish the RTT associated with specific client devices. To the extent that Patent Owner argues that Kommula does not expressly state that the client machine has a unique device identity, a POSITA would be motivated to associate a unique device identity with the client machine to facilitate the lookup of the RTT in the proximity table associated with the client machine. *EX-1003* ¶ 103.

h. Claim 8

[8.0] A system composing:

As discussed with respect to Claim 1, Annotated FIG. 1 of Kommula shows a system for using a global server load balancing switch to resolve a DNS query by a client by providing an ordered list of servers that is expected to serve the client with a high performance in a given application. *EX-1004* [0003]; FIG. 1; *EX-1003* ¶ 104.

[8.1] a memory storing instructions; and one or more processors coupled to the memory, the one or more processors being configured to execute the instructions, the instructions when executed causing the one or more processors to perform operations comprising:

Kommula describes that FIG. 3 illustrates a block diagram showing the functional modules of the GSLB switch for performing the operations described in FIGs. 2A-2D. Kommula discloses that the instructions to be executed by the GSLB switch to control its operations are embodied in software or other machine-readable instruction stored on one or more machine-readable storage media (*memory storing instructions*). *EX-1004* [0026]; *EX-1003* ¶ 105. For example, such software to perform portions of the algorithm may be present at the GSLB switch 12 in one embodiment and executed by the switch controller 401 (*the instructions when executed causing the one or more processors to perform operations*). *EX-1004* [0026], FIGs. 2A-2D, 3; *EX-1003* ¶ 106.

[8.2] receiving a request for a list of available servers from a network device;

Kommula discloses this limitation for the same reasons as discussed for limitation [1.1]. *EX-1003* ¶ 107.

[8.3] retrieving the list of available servers;

Kommula discloses this limitation for the same reasons as discussed for limitation [1.2]. *EX-1003* ¶ 108.

[8.4] retrieving one or more records associated with the network device;

Kommula discloses this limitation for the same reasons as discussed for limitation [1.3]. *EX-1003* ¶ 109.

[8.5] dynamically ordering the list of available servers based on the one or more records associated with the network device; and

Kommula discloses this limitation for the same reasons as discussed for limitation [1.4]. *EX-1003* ¶ 110.

[8.6] transmitting the ordered list of available servers to the network device.

Kommula discloses this limitation for the same reasons as discussed for limitation [1.5]. *EX-1003* ¶ 111.

i. Claim 9

The system of claim 8, wherein the network device is an endpoint device.

Kommula discloses this limitation for the same reasons as discussed for limitation [2.0]. *EX-1003* ¶ 112.

j. Claim 10

The system of claim 9, wherein at least one server among the list of available servers is an access point server.

Kommula discloses this limitation for the same reasons as discussed for limitation [3.0]. *EX-1003* ¶ 113.

k. Claim 11

The system of claim 8, wherein the operations further comprise anticipating a level of demand for at least one server among the list of available servers, and wherein dynamically ordering the list of available servers is further based on the anticipated level of demand.

Kommula discloses this limitation for the same reasons as discussed for limitation [4.0]. *EX-1003* ¶ 114.

l. Claim 12

The system of claim 8, wherein dynamically ordering the list of available servers is further based on one or more second ordered lists of available servers transmitted to one or more second network devices.

Kommula discloses this limitation for the same reasons as discussed for limitation [5.0]. *EX-1003* ¶ 115.

m. Claim 13

The system of claim 8, wherein dynamically ordering the list of available servers comprises generating ranking information for each server in the list of available servers.

Kommula discloses this limitation for the same reasons as discussed for limitation [6.0]. *EX-1003* ¶ 116.

n. Claim 14

The system of claim 8, where in the one or more records are retrieved based on a unique device identity associated with the network device.


Kommula discloses this limitation for the same reasons as discussed for limitation [7.0]. *EX-1003* ¶ 117.

B. Ground 2 –Claims 15-20

1. The Hankins Reference

Hankins describes managing networks, including connecting a private network to the Internet and communicating signal packets from a first device in a first logical broadcast domain to another network device in another logical broadcast domains using a tunnel as if the logical broadcast domains are not separate. *EX-1005* Abstract. Hankins describes that broadcast domain refers to a set of devices occupying a network address space, in which any device is able to communicate with any other device in the broadcast domain without rerouting a transmission. *EX-1005* 6:33-40. A logical broadcast domain means that portions of the broadcast domain may not be entirely related to a particular set of physical devices and instead can be associated with a cloud implemented functionality. *EX-1005* 1:6-7; 4:12-16; 7:45-50; *EX-1003* ¶¶118-120.

Hankins discloses that signal packets are transmitted between the network devices of separate logical broadcast domains via a tunnel server. *EX-1005* Abstract; *EX-1003* ¶ 121. As shown below Hankins connects the network devices via one or more tunneling servers (**red**). A tunneling server facilitates connections between a gateway device (**green**) in the first virtual domain with another gateway device (**blue**) in a second virtual domain via a plurality of tunnels (**purple**):



EX-1005, Fig. 2; *EX-1003* ¶ 122.

Hankins identifies the same contextual problem identified in the '595 Patent—that prior art tunnel controllers did not take into consideration contextual information regarding the network, such as the location of a tunnel server's locations relative to the gateways. Like the '595 Patent, Hankins proposes that the communication channel may be enhanced by using this contextual information, such as by provisioning a tunnel server located more physically proximate to endpoints. *EX-1005* 25:7-25. Like the '595 Patent, Hankins discloses that the controller may generate a list of available tunnel servers based at least in part on a host of considerations, such as, for example, cost, capability, efficiency, compatibility, resources, proximity, latency, bandwidth (e.g., throughput),

utilization, others or combinations thereof. *EX-1005* 26:55-65; *EX-1003* ¶¶ 123, 124.

For example, the gateway 311 communicates via tunnel server one 360 or tunnel server two 370 with gateway 321 over a plurality of tunnels (purple). The tunnels are selected by the controller 350 (orange) using information obtained from the network devices (gateway, tunnel servers, etc.) *EX-1005* 24:39-57, 25:15-55. Gateway device 311 may transmit signal packets via tunnel server one 360 since tunnel server one 360 is more proximate to both gateway device 311 and gateway device 321 in this example. Hankins discloses that an additional server may be selected as a secondary route such as gateway device 311 may transmit signal packets via tunnel server two 370, in the event that tunnel server one 360 is unavailable for a variety of reasons, such as, but not limited to one or more of the following situations: a software issue, a configuration issue, a power issue, scheduled maintenance, insufficient permissions, a security issue, other issues, or combinations thereof. *EX-1005* 25:56-67; *EX-1003* ¶ 125.

Fig. 4 of Hankins generally illustrates the process and tasks of the Gateways (Green), Controller (Yellow) and Tunnel Servers (Red) in establishing the communication paths and transmitting packets.




FIG. 4

EX-1005 Fig. 4; EX-1003 ¶ 126.

2. Detailed Application of Hankins, alone or in view of Kommula to the Challenged Claims

a. Claim 15


[15.0] A network system for managing a global virtual network, comprising:

To the extent that the preamble is limiting, annotated FIG. 2 illustrates a global virtual network of network devices, including gateways, distributed in various logical broadcast domains. A logical broadcast domain means that portions of the broadcast domain may be implemented in the cloud and may not be entirely related to a particular set of devices. *EX-1005* 1:6-7; 4:12-16; 7:45-50; 13:8-26; 19:51-55. A POSITA understood that each of the various logical broadcast domains include gateways that extend the network edge into the cloud at various geographic locations and would be considered part of a *global virtual network*. This is consistent with the '595 Patent's disclosure that one of the features of the global virtual network is that the endpoint device, such as a gateway, acts as a bridge between the broader network and a local area network and brings elements of the cloud as a local node extension into the edge of the LAN. *EX-1001* 2:53-57; 10:37-43; FIG. 9; *EX-1003* ¶ 127,128.

[15.1] an access point server configured to connect to one or more endpoint devices via one or more first tunnels; and

Hankins discloses that the tunnel server (red) (*access point server*) is configured to connect one or more gateways (green) (*endpoint devices*) via one or more tunnels (purple). Hankins discloses that a controller may provision one or more tunnel servers to provide access to different broadcast domains and facilitate

communications between network devices communicating as part of different broadcast domains. *EX-1004* 8:63-66; 26:65-27:8. This is consistent with the disclosure of the '595 Patent where gateways are the endpoints of a tunnel connecting the gateways, and the gateways are the edge of their respective networks. *EX-1001* FIGs. 9, 15; *EX-1003* ¶ 129, 130.



EX-1005 FIG. 2; *EX-1003* ¶ 130.


[15.2] a control server connected to the access point server via a second tunnel, the control server storing one or more records associated with each of the one or more endpoint devices,

Hankins discloses that a controller (yellow) (*a control server*) is connected to a tunnel server (red) (*access point server*) via a *second tunnel* (yellow).

Although Hankins does not expressly state that the communication path shown in FIG. 2 between controller 350 and tunnel sever 360 or tunnel server 370 is a tunnel, a POSITA understood this path could be a tunnel. For example, Hankins discloses that a gateway can operate as a controller, and further discloses that the communication path between a gateway and a tunnel server is via a tunnel. *EX-1005* 14:41-49; 12:1-25. Further, as discussed in the State of the Art, it was known to use a tunnel between an access point and the controller. *EX-1006* [0013]. Thus, a POSITA understood that controller could be implemented in a gateway which is connected to a tunnel server via a second tunnel. *EX-1003* ¶¶ 131,132.

Hankins further discloses that a controller may provision one or more tunnel servers to facilitate communications between network devices communicating as part of different broadcast domains. *EX-1004* 8:63-66. In order to do so, the gateway devices transmit respective generated broadcast configuration (*one or more records associated with each of the one or more endpoint device*) to controller. *EX-1004* 20:24-25. Hankins explains that broadcast domain configuration information includes one or more network protocols, available

addresses, used addresses, topologies, devices used, such as switches or hubs, historical settings, such as for security, for a network protocol, etc., modifications of the foregoing, user accounts, including status, authentication, etc., security settings of a broadcast domain, workgroup or domain names, device names including status, available device features, etc., services available and/or used, status of the network devices, as well as other features. *EX-1004* 11:64-12:8; *EX-1003* ¶¶ 133-135.



EX-1005 FIG. 2; *EX-1003* ¶ 135.

In another example, Hankins discloses that a discovery device may comprise a network device capable of communicating as part of a broadcast domain to

discover features of the particular broadcast domain without human intervention to generate a broadcast domain configuration of the particular broadcast domain. A discovery device may engage in a variety of activities to accomplish this including, but not limited to any of the following: passively monitoring signal packets, actively probing other network devices, querying other network devices, querying servers, such as a domain controller or other server responsible for a directory of broadcast domain features, any combinations thereof, or other approaches. *EX-1004* 9:10-22. Thus, configuration information for a particular broadcast domain associated with a gateway may be generated using a discovery device. *EX-1003* ¶¶ 136, 137.

As shown in FIG 1, a gateway device transmits a generated broadcast domain configuration associated with the gateway to the controller 206, and the controller stores the generated broadcast domain configuration associated with the gateway device in database 207 (*the control server storing one or more records associated with each of the one or more endpoint devices*). *EX-1004* 21:7-19; *EX-1003* ¶ 138.

[15.3] wherein the control server is configured to receive a request for a list of available servers from a first endpoint device among the one or more endpoint devices, the list of available servers including the access point server;

Hankins discloses that the controller receives a request from a gateway 311 (*endpoint device*) for a list of available servers including tunnel servers (*access*

point server). For example, Hankins discloses that gateway 311 in one logical broadcast domain, may request services available from gateway 321 in another logical broadcast domain. *EX-1005* 12:9-16. The controller identifies potential tunnel servers that are available to provide a tunnel between gateway 311 to gateway 321. For example, as shown in Annotated FIG. 2, tunnel servers 360 and 370 can be used to establish communications between gateway 311 and 321. Thus, a POSITA understood that a request from gateway 311 for services available from gateway 321 is a request for available tunnel servers that can provide a tunnel between gateway 311 and 321. This is consistent with the disclosure of the '595 Patent that when a device makes a request to a server, what it is really asking for is a list of available servers, and specifically an IP and port which it could utilize to in the build a tunnel. *EX-1001* 23:11-24; *EX-1003* ¶¶ 139, 140.

Furthermore, via a controller, such as 350, a private network may be reconfigured without human intervention, for example, to improve performance and/or reduce cost. For example, if available bandwidth decreases, if latency increases, and/or if network utilization increases, controller 350 may provision additional tunnel servers or may provision tunnel servers in a manner to improve performance and/or reduce cost without human intervention. Examples of possible monitoring and/or measurement of network performance and/or cost-related parameter were previously discussed. *EX-1005* 25:45-55; *EX-1003* ¶ 141.

Hankins discloses a controller 350 that receives domain configuration information and identifies tunnel servers associated with one of the paths or channels for communications between the gateway devices. *EX-1005* 24:49-55; *EX-1003* ¶ 142.

[15.4] retrieve the list of available servers; retrieve the one or more records associated with the first endpoint device

Hankins discloses that controller *retrieves a list of available tunnel servers* based on the stored configuration of the broadcast domain of a gateway (*retrieve the one or more records associated with the first endpoint device*). For example, Annotated FIG. 2 shows an embodiment with four logical broadcast domains in communication via more than one tunnel server. Gateway devices 311, 321, 331, and 341 may generate broadcast configurations for their respective logical broadcast domains which may be used by controller to identify a list of available tunnel servers. For example, controller 350 may provision tunnel server one 360 and initiate instructions to gateway devices 311, 321, 331, and 341 to communicate with tunnel server one 360. *EX-1004* 24:47-55; *EX-1003* ¶ 143.

In addition, Hankins discloses retrieving more than one tunnel server to create multiple tunnels between gateways. For example, in addition to identifying tunnel server one 360, controller may identify tunnel server two 370 which is more physically proximate to local broadcast domain c 330 and logical broadcast domain D 340. *EX-1004* 25:14-19. Additional tunnel servers can be identified based on a

variety of factors or considerations, including, signal packets communicated via a tunnel server, geography, latency, utilization, bandwidth (e.g., throughput), availability, efficiency, cost, etc. *EX-1004* 34-45; *EX-1003* ¶¶ 144, 145.

[15.5] dynamically order the list of available servers for the first endpoint device based on the one or more records associated with the network device; and

Hankins discloses several examples of this limitation. Hankins discloses dynamically ordering a list of available tunnel servers for the gateway based on the one or more records associated with the gateway that provide contextual information. For example, gateway 311 is provided with a list of tunnel servers such as tunnel server one 360 and tunnel server two 370, based on information obtained from the network devices (gateway, tunnel servers, etc.) such as proximity of gateway device to tunnel server 360. *EX-1005* 24:39-57, 25:15-55. In addition, Hankins discloses that additional server, tunnel server two 370, may be selected as a secondary route such that gateway device 311 transmits signal packets via tunnel server two 370, in the event that tunnel server one 360 is unavailable for a variety of reasons, such as, but not limited to one or more of the following situations: a software issue, a configuration issue, a power issue, scheduled maintenance, insufficient permissions, a security issue, other issues, or combinations thereof. *EX-1005* 56-67. Thus, Hankins discloses identifying an additional tunnel server to provide a secondary route in the event that a primary

tunnel server is unavailable, which a POSITA understood discloses dynamically ordering the availability of servers. *EX-1003* ¶¶ 146- 148.

By way of another example, Hankins discloses an embodiment where broadcast domains communicate through multiple tunnel servers. Each nth packet goes to the nth tunnel server. Hankins describes that this approach can improve security, fault tolerance, redundancy and/or load balancing of tunnel servers. *EX-1005* 26:21-48. A POSITA understood that providing load balancing and fault tolerance means that a list of tunnel services would be provided in preference of order. *EX-1003* ¶¶ 149, 150.

Hankins further describes that different tunnel servers could provide different services. For example, one tunnel server may be better suited for use in file serving and the other tunnel server could be better suited for message serving. Hankins describes that the provisioning of tunnel servers may be generated based on the capability of the tunnel servers and thus a POSITA understood that Hankins orders the availability of the tunnel servers on the needs of the gateways in the broadcast domains. For example, if the needs include file service more than messaging, the tunnel server handling file service would be ranked above the tunnel server for message servicing. *EX-1003* ¶ 151.

Thus, Hankins discloses dynamically ordering the list of available servers. *EX-1003* ¶ 152.

To the extent that Patent Owner assert that Hankins does not expressly state that the available tunnel servers are dynamically ordered, Kommula disclose dynamically ordering a list of available servers to provide a communication path between network devices based on contextual information. A POSITA would be motivated to modify the identification of available tunnel servers as taught by Hankins to include dynamically ordering as taught by Kommula in order to make a more efficient use of network resources. *EX-1003* ¶ 153.

Hankins and Kommula are analogous art. Both Hankins and Kommula are directed to identifying available servers to provide a communication path through a network. *EX-1005* 12:9-16, 24:39-57, 25:15-55; *EX-1004* [0008], [0019], [0026], [0027]; *EX-1003* ¶ 154. Both Hankins and Kommula use contextual information to identify available servers such as the proximity of the server to the endpoint device or gateway. *EX-1005* 25:7-25; 26:55-65; *EX-1004* [0009]; *EX-1003* ¶ 155. Kommula expressly describes dynamically ordering the list of available servers based on the contextual information such as proximity using RTT. *EX-1004* [0009], [0095], [0135]; *EX-1003* ¶ 156.. A POSITA would be motivated to modify the teaching of Hankins to identify available tunnel servers based on contextual information by including dynamically ordering the available servers as taught by Kommula because it results in a more efficient use of network resources. For example, Hankins describes that the available servers in FIG. 2 include tunnel

server one 360 and tunnel server two 370. Dynamically ordering server one 360 ranked ahead of tunnel server two 370 because tunnel server 360 is more proximate to gateway 311 results in a shorter RTT, as taught by Kommula, for communications using tunnel server one as compared to tunnel server two, which is one of the express goals of Hankins. *EX-1003* ¶¶157.

A POSITA would be confident of the success in modifying the teachings of Hankins with Kommula for several reasons. *EX-1003* ¶ 158. First, Hankins and Kommula both teach using software implementations to identify available servers, which could be identified by accessing stored contextual data. *EX-1005* FIGs. 4, 5; *EX-1004* [0013], [0026], FIGs 2A-2D; *EX-1003* ¶ 159. Both Hankins and Kommula teach using a software implementation to access stored data to find the “best” available server, e.g. closest proximity or shortest RTT. *Id*; *EX-1003* ¶ 160. Thus, it is only a simple software change to Haskins teachings to dynamically rank the results of the contextual information retrieved from the database. *EX-1003* ¶ 161.

Further, Hankins discloses that the tunnel servers can implement network policies and do other services. *EX-1005* 8:12-37; *EX-1003* ¶ 162. Hankins further describes that the provisioning of tunnel servers are reconfigurable in real-time. *EX-1005* 18:46-63; *EX-1003* ¶ 163. Moreover, it was known that ranking of tunnel servers can be established as a network policy, as taught by Kommula. And

Hankins contemplates that tunnel servers are readily capable of executing such policies. *EX-1003* ¶ 164.

[15.7] transmit the ordered list of available servers to the first endpoint device.

Hankins in combination with Kommula discloses this limitation. Hankins discloses that the available tunnel servers are provided to gateway 311. For example, as discussed with limitation [15.4], Hankins discloses that controller can identify tunnel servers 360 and 370 for gateway 311 to communicate with gateway 321. Gateway 311 may transmit signal packets to gateway device 321 via tunnel server one 360 since tunnel server one 360 is more proximate to both gateway device 311 and gateway device 321 in this example. *EX-1003* ¶ 165.

Additionally, in response to dynamically ordering the available servers, gateway device 311 may transmit signal packets via tunnel server two 370 if tunnel server one may be unavailable, for a variety of reasons, such as, but not limited to one or more of the following situations: a software issue, a configuration issue, a power issue, scheduled maintenance, insufficient permissions, a security issue, other issues, or combinations thereof. Thus, Hankins discloses transmitting an ordered list, i.e., tunnel server 360 if available, and then tunnel server 370. *EX-1003* ¶ 166.

b. Claim 16

The network system of claim 15, wherein the control server is further configured to predict a level of demand for the access point server, and wherein dynamically ordering the list of available servers is further based on the anticipated level of demand.

Hankins as modified by Kommula discloses this limitation for the same reason as discussed with respect to Claim 4. *EX-1003* ¶ 167.

b. Claim 17

The network system of claim 15, wherein dynamically ordering the list of available servers is further based on one or more second ordered lists of available servers transmitted to one or more second endpoint devices among the plurality of endpoint devices.

Hankins as modified by Kommula discloses this limitation for the same reason as discussed with respect to Claim 5. *EX-1003* ¶ 168.

c. Claim 18

The network system of claim 15, where in dynamically ordering the list of available servers comprises generating ranking information for each server in the list of available servers.

Hankins as modified by Kommula discloses this limitation for the same reason as discussed with respect to Claim 6. *EX-1003* ¶ 169.

d. Claim 19

The network system of claim 15, the one or more records are retrieved based on a unique device identity associated with the first endpoint device.

Hankins as modified by Kommula discloses this limitation for the same reason as discussed with respect to Claim 7. *EX-1003* ¶ 170.

e. **Claim 20**

The network system of claim 15, wherein the control server receives tunnel information associated with the one or more first tunnels via the global virtual network, and wherein dynamically ordering the list of available servers is further based on the tunnel information.

Hankins as modified by Kommula discloses this limitation for the same reason as discussed with respect to Claim 4. *EX-1003* ¶ 171.

VIII. THE BOARD SHOULD NOT EXERCISE ITS DISCRETION AND DENY INSTITUTION

The Board Should Not Deny Institution Under 35 U.S.C. § 314(a)

The Board has discretion to deny institution under 35 U.S.C. § 314(a).

Related to the Director's June 2022 interim guidance regarding application of the *Fintiv* factors², (1) the Petition is particularly strong in the underlying merits, and (2) the Petitioner adopts the identical stipulation from *Sotera*, namely, that if the IPR is instituted, Petitioner will not pursue in the Related Litigation any ground raised or that could have been reasonably raised in an IPR for the challenged patent. Thus, under the Interim Guidance, the Board should not deny institution.

² Interim Procedure For Discretionary Denials In AIA Post Grant Proceedings With Parallel District Court Litigation, PTO Director's Memorandum (June 21, 2022).

Nevertheless, the Board's decision in *Sotera Wireless, Inc. v. Masimo Corp.*, IPR2020-01019, Paper 12 (Dec. 1, 2020) instructs that a holistic view of the remaining *Fintiv* factors also weighs in favor of institution.

Factor 1 is neutral because no request for stay has been filed.

Under factor 2, trial date has been set for August 10th-14, 2026. A Final Written Decision is expected in the present matter in August 2025, 12 months before trial in the Related Litigation. Thus, this factor weighs against discretionary denial.

Factor 3 favors institution. The co-pending litigation is in its early stages, and the investment in it has been minimal. The parties have not exchanged proposed claim terms or constructions, and will not complete claim construction briefing until October 2024. The Markman hearing is not until December 2024. Fact discovery is not set to close until September 2025, and expert discovery is not set to close until February 5, 2026.; *See PEAG LLC v. Varta Microbattery GMBH*, IPR2020-01214, Paper 8 at 17 (Jan. 6, 2021) (finding that since no claim construction hearing had yet been held and discovery was not completed, the little investment in the parallel proceeding weighed against discretionary denial). Thus, this factor weighs against discretionary denial.

Under factor 4, in addition to the *Sotera* stipulation that Petitioner issued to PO, PO has only identified in the Complaint a subset of the claims that are being

challenged in the present Petition. Thus, it is unlikely the district court will be addressing the validity of all of the claims challenged by this Petition.

Under factor 5, Petitioner is a defendant in the Related Litigation. Nothing within *Fintiv* suggests that the same party between the proceedings weighs in favor of discretionary denial. See *HP Inc. v. Slingshot Printing LLC*, IPR2020-01084, Paper 13 at 9 (having the “same parties as parallel proceeding” makes factor 5 “neutral”)).

Under factor 6, other circumstances weigh in favor of institution. Here, the merits of the Petition are particularly strong. For example, Kommula, along with Hankins are parallel disclosures to the ’595 Patent with respect to the subject matter of Claims 1-14 and 15-20 respectively, meaning that they are directed at the same problem and proposes the same solution as the ’595 Patent, and discloses identical architecture and the identical concepts disclosed in the ’595 Patent, but use some different words, thus demonstrating that the Petition is particularly strong in the underlying merits. The evidence presented, if unrebutted in trial, would plainly lead to a conclusion that one or more claims are unpatentable by a preponderance of the evidence.

When viewed holistically, the timing of the present Petition is reasonable, there has been relatively limited investment in the Related Litigation, and there is minimal overlap between the present IPR and the Related Litigation. Further,

coupling these *Fintiv* considerations with compelling evidence of unpatentability presented in the Petition, the efficiency and integrity of the IPR process is best served by instituting review.

IX. Mandatory Notices Under 37 C.F.R. §42.8

A. Real Party-in-Interest (37 C.F.R. § 42.8(b)(1))

The real party-in-interest in this Petition is Cisco Systems, Inc.

B. Related Matters (37 C.F.R. § 42.8(b)(2))

1. Judicial Matters

As of the filing date of this Petition and to the best knowledge of Petitioner, the '595 Patent is involved in the following litigations:

UMBRA Technologies Ltd. (UK) et al v. Cisco Systems, Inc., Case No. 1:23-cv-00903-DII (W.D. Tex.) (filed Aug. 2, 2023) (ongoing)

UMBRA Technologies Ltd. (UK) et al v. VMware, Inc., Case No. 1:23-cv-00904-DII (W.D. Tex.) (filed Aug. 2, 2023) (ongoing).

2. Administrative Matters:

As of the filing date of this Petition and to the best knowledge of Cisco Systems, the '595 Patent was subject to the following *inter partes* reviews which were not instituted:

3. Related Patents

See Exhibit *EX-1007*

C. Lead/Back-up Counsel (37 C.F.R. § 42.8(b)(3)):

Lead Counsel:

Patrick D. McPherson, USPTO Reg. No. 46,255
DUANE MORRIS LLP, 901 New York Avenue N.W., Suite 700 East
Washington, DC 20001
P: (202) 776-7800 (202) 776-7801; PDMcPherson@duanemorris.com

Back-Up Counsel: Patrick Muldoon, USPTO Reg. No. 47,343
DUANE MORRIS LLP, 901 New York Avenue N.W., Suite 700 East
Washington, DC 20001
P: (202) 776-7840; F: (202) 776-7801; PCMuldoon@duanemorris.com

Joseph English, USPTO Reg. No. 42,514
DUANE MORRIS LLP, 901 New York Avenue N.W., Suite 700 East
Washington, DC 20001
P: (202) 776-5228; F: (202) 776-7801; DJEnglish@duanemorris.com

D. Notice of Service Information (37 C.F.R. § 42.8(b)(4)):

Please direct all correspondence to lead and back-up counsel at the above addresses. Petitioner consents to electronic service at the email addresses above.

X. CONCLUSION

For the reasons set forth above, Petitioner has established a reasonable likelihood of prevailing with respect to all Challenged Claims of the '595 Patent and requests the Board institute *inter partes* review and then cancel all claims as unpatentable.

Respectfully submitted,

DUANE MORRIS LLP

BY: /Patrick D. McPherson/

Patrick D. McPherson
USPTO Reg. No. 46,255
Duane Morris LLP
901 New York Avenue N.W.
Suite 700 East
Washington, D.C. 20001

Dated: February 26, 2024

CERTIFICATION OF SERVICE ON PATENT OWNER

Pursuant to 37 C.F.R. §§ 42.6(e), 42.8(b)(4) and 42.105, the undersigned certifies that on the 26th of February 2023, a complete and entire copy of this Petition for *Inter Partes* Review of U.S. Patent No. 11,108,595 and all supporting exhibits were served via Federal Express, postage prepaid, to the Patent Owner by serving the correspondence address of record for the '595 Patent:

WILMER HALE/BOSTON
60 STATE STREET
BOSTON, MA 02109
UNITED STATES

/Patrick D. McPherson/
Patrick D. McPherson, Reg. No. 46,255
901 New York Avenue, N.W.,
Suite 700 East
Washington, D.C. 20001
P: (202) 776-7800
F: (202) 776-7801
PDMcPherson@duanemorris.com

ATTORNEY FOR PETITIONER

CERTIFICATE OF COMPLIANCE

Pursuant to 37 C.F.R. § 42.24 *et seq.*, the undersigned certifies that this document complies with the type-volume limitations. This document contains 11,048 words as calculated by the “Word Count” feature of Microsoft Word 2010, the word processing program used to create it.

Dated: February 26, 2024

By: Patrick D. McPherson
Patrick D. McPherson
Reg. No. 46,255
Duane Morris LLP
901 New York Avenue N.W.,
Suite 700 East
Washington, D.C., 20001
P: (202) 776-7800
F: (202) 776-7801
PDMcPherson@duanemorris.com