UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

GOOGLE LLC,
Petitioner

v.

SANDPIPER CDN, LLC,
Patent Owner

Case IPR2025-00806
U.S. Patent No. 8,645,517

DECLARATION OF DR. PRASHANT SHENOY IN SUPPORT OF PATENT OWNER'S RESPONSE

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

Sandpiper EX2017
Google v. Sandpiper
IPR2025-00806

TABLE OF CONTENTS

I. INTRODUCTION 1
II. LIST OF DOCUMENTS CONSIDERED 6
III. OVERVIEW OF THE STATE OF THE ART 8
IV. THE '517 PATENT 10
A. Background of the '517 Patent 10
B. Technical Overview of the '517 Patent 11
V. PROSECUTION HISTORY 16
VI. LEGAL STANDARDS 18
A. Person of Ordinary Skill in the Art ("POSA") 19
B. Claim Construction 19
C. Obviousness 20
D. Written Description 22
VII. ARGUMENT 23
A. Hasan is non-analogous to the '517 patent 23
1. Hasan is not from the same field of endeavor as the '517 patent 24
2. Hasan is not reasonably pertinent to the particular problem to which the '517 patent relates 30
B. Grounds 1 and 2 33
1. Petitioner and Dr. Lin fail to identify any "policy" in Beaumont 33
a. The claimed policy is a rule 34
b. Petitioner and Dr. Lin do not identify the claimed policy in either Beaumont or Hasan 37
c. Beaumont-Hasan does not teach using a GUI to set or obtain policies for direction of network traffic. 40
d. Beaumont-Hasan does not teach providing said at least one policy to the name server network. 45

2. Beaumont-Hasan does not teach the claimed GUI that uses a decision tree 46

a. Beaumont-Hasan's GUI does not support the setting of policies using a decision tree representing rules 46

i. The references do not teach this element individually.....47

ii. The combination of Beaumont and Hasan does not teach or suggest the claimed GUI or decision tree 50

b. Beaumont-Hassan's GUI does not support the setting of policies using a decision tree that includes branch nodes specifying decision criteria..... 52

i. Neither Beaumont nor Hasan teaches or suggests a decision tree that includes branch nodes specifying decision criteria 53

ii. The combination of Beaumont and Hasan does not teach or suggest a decision tree that includes branch nodes specifying decision criteria 56

3. A POSA would not have been motivated to combine Beaumont and Hasan..... 59

a. A POSA would not have combined Beaumont with Hasan at all..... 59

b. A POSA would not have combined Beaumont and Hasan as proposed by Petitioner and Dr. Lin. 61

C. Grounds 3A and 3B.....65

1. Overview of Ben-Shaul..... 65

a. Ben-Shaul's Edge Servers 65

b. Ben-Shaul's DNS..... 69

c. Coordination of the Edge Server and Edge DNS 74

2. The cited prior art does not teach or suggest the claimed decision tree..... 75

a. The cited prior art does not teach or suggest a decision tree that includes resource nodes specifying answers to be provided in response to a DNS request 76

b. The cited prior art does not teach or suggest a decision tree that includes branch nodes specifying decision criteria to be applied 84

c. The cited prior art does not teach or suggest a decision tree at all 87

3. The cited prior art does not teach or suggest providing said at least one policy to the name server network 90

4. Ben-Shaul does not disclose a domain name server adapted to provide adaptive policy-based domain name service 91

5. A POSA would not have combined Ben-Shaul and Hasan 93

VIII. CONCLUSION 1

I, Prashant Shenoy, Ph.D., do hereby declare as follows:

I. INTRODUCTION

1. I am over the age of twenty-one and otherwise competent to make this Declaration. This Declaration is based on my personal knowledge as an expert in the relevant field. I could testify competently if asked to do so.

2. I have been retained by Sterne, Kessler, Goldstein & Fox, P.L.L.C. as a technical expert witness on behalf of Sandpiper CDN, LLC to provide this Declaration concerning technical subject matter relevant to the inter partes review of U.S. Patent No. 8,645,517 ("'517 patent") in IPR2025-00806.

3. I am being compensated for the time spent on this matter at my standard consulting rate of $650.00/hour. My compensation is not dependent on, and in no way affects, the substance of my statements in this Declaration. My findings are based on my education, experience, and background in the fields discussed below. I reside in Amherst, Massachusetts, USA.

4. In formulating my opinions, I have relied upon my experience in the relevant art, and I have considered each document cited herein, in light of the general knowledge in the art as of September 28, 2001, the effective filing date of the '517 patent. However, my opinions herein would not change even if the effective filing date of the '517 were October 26, 2007, as argued by Petitioner.

5. I will cite to the specification of a U.S. patent using the following format: (Patent, Col: Line Number(s)). For example, the citation to the ’517 patent, 1:1-10 points to the ’517 patent specification at column 1, lines 1-10.

6. I am familiar with the technology at issue and the state of the art at the time the application leading to the ’517 patent was filed, including content delivery networks, web protocols, network protocols, and proxy servers.

7. I have been asked to provide my technical review, analysis, insights, and opinions regarding the references. I reserve the right to supplement this Declaration in response to additional evidence that may come to light.

8. Over the years, I have led research and industry collaboration projects in areas such as storage and file systems, Web and Internet systems, content distribution, streaming media, database systems, cloud computing, virtualization, RFID and sensor networks, and energy systems. During that time, I have consulted with both large companies and startups in a technical advisory role related to product architecture, IP issues, and applied research.

9. I am a tenured professor in the College of Information and Computer Sciences at the University of Massachusetts Amherst ("UMass"), where I currently hold the position of Distinguished Professor of Computer Science. I also served as an Associate Dean in the College of Information and Computer Sciences from 2016 to 2025. I have been a faculty member at UMass since 1998.

10. I received my Bachelor of Technology (B.Tech.) degree in 1993 from the Department of Computer Science and Engineering at The Indian Institute of Technology ("IIT"), Mumbai, India. IIT awarded me the Institute Medal for being the top-ranking undergraduate student. In 1994, I received my Master of Science (M.S.) degree in Computer Sciences from The University of Texas at Austin. In 1998, I received my Doctor of Philosophy (Ph.D.) also from the Department of Computer Sciences at The University of Texas at Austin. My dissertation, which was in the area of file systems, was recognized by the department as the Best Doctoral Dissertation of 1998-1999.

11. I have been involved in the research and development of distributed computing systems for more than 30 years. At UMass, I perform research and teach in the areas of distributed systems, operating systems, computer networking, data management, and multimedia systems, among others.

12. A major focus of my research has been the design of large-scale distributed computing systems. Over the years, I have led research and industry collaboration projects in areas such as computer networking, distributed systems, storage and file systems, Web and Internet systems, content distribution, streaming media, database systems, cloud computing, virtualization, RFID and sensor networks, and energy systems. During that time, I have consulted with both large

companies and startups in a technical advisory role related to product architecture, IP issues, and applied research.

13. For example, starting in the early 1990s, my research focused on the development of storage and file systems. During that time, I developed multimedia file system and data management technologies that were designed for efficient storage and retrieval of large amounts of video and audio data. I collaborated with AT&T Bell Laboratories to design a new media server for video streaming. These efforts led to resulted in U.S. Patent No. 6,079,028 on fault-tolerant architectures for continuous media servers.

14. In the late 1990s, I led the design of a new operating system called QLinux, that was based on the open-source Linux operating system. As part of this research, I developed resource management and scheduling techniques that provide isolation and service guarantees to applications. This technology provided the underpinnings for fair resource management and isolation techniques used in OS virtualization. It was adopted by a startup company for providing resource isolation in its web hosting products.

15. In the early 2000s, I led several projects on web systems and cloud computing. For example, I led research projects on web technologies such as web caching and scaling web applications. I also led research projects on cloud computing platforms (which were known as web hosting platforms at the time)

involving virtualization, dynamic resource management, replication, fault tolerance, and storage. My group at UMass collaborated with technology companies such as IBM, AT&T, Sprint, Intel, and many others as part of this research.

16. My research on virtualization has involved design of technologies for virtual machines and container resource management and their deployment in cloud computing systems and data centers. I am co-inventor on U.S. Patent No. 8,473,557, jointly assigned to AT&T and UMass, on virtual machine migration.

17. Over the years, I have authored or co-authored more than 300 articles in peer-reviewed journals and conference proceedings that have been cited over 27,000 times. Seventeen papers authored by myself and members of my research group on this project won awards for their technical contributions. Two papers have won Test-of-Time awards, an award given to a publication that has had the most impact over a ten-year period since its publication.

18. I have served as an Editor for six leading journals and have served as the Technical Program Chair for twenty conferences and workshops in my field. I founded the ACM Special Interest Group on Energy Systems and also served on steering committee of conferences and journals.

19. I am also a Fellow of the ACM, IEEE, AAAS, and AAIA, four professional societies in computing and science.

20. I am a named inventor on three U.S. Patents and one published application (Nos. 2014/0040206, 8,510,596, 8,473,557, and 6,079,028). These documents are generally related to distributed and operating systems, virtualization, and cloud data centers.

21. As a faculty member at UMass, I have also taught undergraduate and graduate courses in operating systems, distributed systems, networking, and related topics for more than twenty five years.

22. These courses teach concepts related to operating systems, virtual machines, containers, cloud computing, file systems, and distributed systems. They also involve hands-on course projects on these topics.

23. My *Curriculum Vitae* is provided as Exhibit 2016, which contains further details on my education, experience, publications, and other qualifications to render an expert opinion.

II. LIST OF DOCUMENTS CONSIDERED

Exhibit No.	Description
1001	U.S. Patent No. 8,645,517 to Stolorz et al. ("'517 Patent")
1002	Prosecution History of U.S. Application No. 12/978,537
1003	Declaration of Dr. Bill Lin ("Lin")
1004	Curriculum Vitae of Dr. Bill Lin
1005	U.S. Patent Pub. No. 2002/0010798 to Ben-Shaul et al. ("Ben-Shaul")

Exhibit No.	Description
1006	U.S. Patent No. 7,082,464 to Hasan et al. ("Hasan")
1007	U.S. Patent Pub. No. 2002/0169890 to Beaumont et al. ("Beaumont")
1008	U.S. Patent Pub. No. 2002/0078233 to Biliris et al. ("Biliris")
1009	Balachander Krishnamurthy et al., On the Use and Performance of Content Distribution Networks, Proceedings of the 1st ACM SIGCOMM Workshop on Internet measurement, November 2001 ("Krishnamurthy")
1010	Zhuoqing Morley Mao et al., A Precise and Efficient Evaluation of the Proximity between Web Clients and their Local DNS Servers, Proceedings of the 2002 USENIX Annual Technical Conference, June 2002 ("Mao")
1011	U.S. Provisional Patent App. No. 60/325,177 to Stolorz et al. ("Provisional App.")
1012	U.S. Patent Pub. No. 2003/0065762 to Stolorz et al. ("2002 Application")
1013	U.S. Patent Pub. No. 2008/0215718 to Stolorz et al. ("2007 CIP")
1014	Prosecution History of U.S. Application No. 10/259,497
2001	"Cloud CDN: content delivery network," *Google Cloud*, accessed at https://cloud.google.com/cdn (last accessed July 2, 2025)
2002	Prosecution History of U.S. Application No. 10/259,497
2003	Prosecution History of U.S. Application No. 11/932,162
2004	*Sandpiper CDN, LLC v. Google LLC*, 2:24-cv-03951, ECF No. 1 (C.D. Cal. May 10, 2024)

Exhibit No.	Description
2005	*Sandpiper CDN, LLC v. Google LLC*, No. 2:24-cv-03951, Invalidity Contentions (April 22, 2025)
2006	*Sandpiper CDN, LLC v. Google LLC*, No. 2:24-cv-03951, ECF No. 82 (C.D.Cal. June 24, 2025)
2007	*Sandpiper*, No. 2:24-cv-03951, ECF No. 28 (C.D. Cal. Sept. 16, 2024)
2008	FAQs for Interim Processes for PTAB Workload Management, accessed at https://www.uspto.gov/patents/ptab/faqs/interim-processes-workload-management
2009	*Gestures Tech. Partners LLC. v. Apple Inc. et al.,* Petition for Writ of Certiorari, Case. No. 24- (Filed June 11, 2025)
2010	*Intentionally Left Blank*
2011	*Intentionally Left Blank*
2012	*Intentionally Left Blank*
2013	Interim Processes for PTAB Workload Management Memorandum (March 26, 2025)
2014	National Judicial Caseload Profile – United States District Courts (March 2025)
2016	Curriculum Vitae of Dr. Prashant Shenoy

III. OVERVIEW OF THE STATE OF THE ART

24. Content delivery networks (CDNs) provide critical services that enable content providers to quickly deliver online content to millions of consumers simultaneously over the Internet. In the early 1990s, the World Wide Web saw increasing adoption, becoming a household staple. This mass adoption led to data

congestion issues from the ever-growing number of users seeking to simultaneously access Internet content.

25. A typical computer server in the 1990s could only handle a limited number of simultaneous connections before becoming overloaded. Moreover, signals take time to traverse routers and move through content-congested internet cables. Consumers living far from the physical server(s) hosting content experienced sluggish load times and high latency due to problems such as overloaded servers, congested network segments, and geographic separation.

26. The availability of CDN servers around the world allowed service providers to replicate appropriate content from customers' origin servers to appropriate CDN servers, and transparently direct end-users requesting that content to the "best" CDN server to deliver that content, all while providing their customers with control over their content and user experience. This service and its architecture was quickly adopted by many others in the industry.

27. End-users would connect to an optimal edge server—one that that was closer to them or that had available capacity. Distributing content across a network of servers alleviated data congestion issues and allowing consumers to connect to edge servers located near them reduced latency. Systems and methods for propagating data from origin servers to edge servers and storing them at edge servers for serving user requests are known as "caching."

28. In the early-to-mid 2000s, demand for CDNs exploded. Content providers like Google and YouTube use CDN technology to meet their enormous data streaming needs, as well as for video search results. EX2001.

IV. THE '517 PATENT

A. Background of the '517 Patent

29. U.S. Patent Number 8,645,517 ("the '517 patent") is entitled "Policy-Based Content Delivery Network Selection." As the title suggests, a novel feature of the invention includes a particular kind of policy for directing network traffic in an innovative way. The '517 patent claims benefit via four non-provisional applications in two family chains to U.S. Provisional Application No. 60/325,177 filed September 28, 2001.

30. When using CDN technology, challenges arise related to network traffic associated with requests for content. For example, companies that want to provide their content to end users using a CDN face issues regarding directing network traffic to serve end users' requests. As one example, issues arise when trying to deliver content from multiple servers at geographically separated locations with suitable end-user experiences under high traffic loads.

31. The '517 patent addresses specific needs in the art via specific combinations, which were not conventional at the time of the invention of the '517 patent, and the '517 patent describes how specific technical solutions are achieved.

For instance, content providers must address how to handle requests from end users, in some cases in accordance with policy constraints faced by the content providers, even as network conditions affecting the Internet are subject to change.

32. The inventors of the ’517 patent developed methods to address these issues, including methods in which a server network implements policy-based traffic direction. For example, embodiments of the invention in the ’517 patent provide a specific graphical user interface for obtaining at least one policy for the direction of network traffic, such as using a decision tree with resource and branch nodes. The various criteria that can be used may relate to IP addresses or geographic zones, with answers consisting of IP addresses and CNAMEs, for example.

B. Technical Overview of the ’517 Patent

33. The ’517 patent provides solutions for “configurable, adaptive, global traffic control and management in networks such as the Internet.” EX1001, 1:36-39. Specifically, the ’517 patent is directed to “a framework wherein resources of a content provider may be delivered to clients from different domains.” EX1001, Abstract. To this end, the specification and claims describe a policy-based “DNS service.” EX1001, 3:16-20, claim 1 (reciting a “name server network comprising at least one domain name server” that provides “adaptive policy-based domain name service” and “specif[ies] one or more answers to be provided in response to a

Domain Name Service (DNS) request").

34. The '517 patent describes the process for delivering network (e.g., internet) data to end users with respect to Figure 1. First, a user at a client device 12 "enters a [URL] into a browser 114 on client 112." EX1001, 3:60-62. In response, a DNS resolver 116 on the client 112 "queries a particular name server (e.g., name server 118-1)" to determine the IP address that corresponds to the domain associated with that URL. EX1001, 4:35-44. The name server determines and "returns [] the IP (Internet Protocol) address… of a server… in the subscriber server network" that hosts the requested content associated with the domain. EX1001, 4:35-44.

35. As the '517 patent explains, multiple servers may host the requested content. *See* EX1001, 4:39-42, 4:60-67. The '517 patent describes and claims policies for determining which server will receive the network requests. For instance, "[t]he determination of the particular IP address returned to the resolver 116 may be based on a number of factors, including the resolver's location (e.g., as determined from the resolver's IP address) and various policies (e.g., subscriber policies 120, other policies 122) in a policy database 124." EX1001, 4:44-49. When the browser receives the returned IP address, it is able to direct its content query to the server having that address "in order to obtain the desired resource." EX1001, 4:49-51.

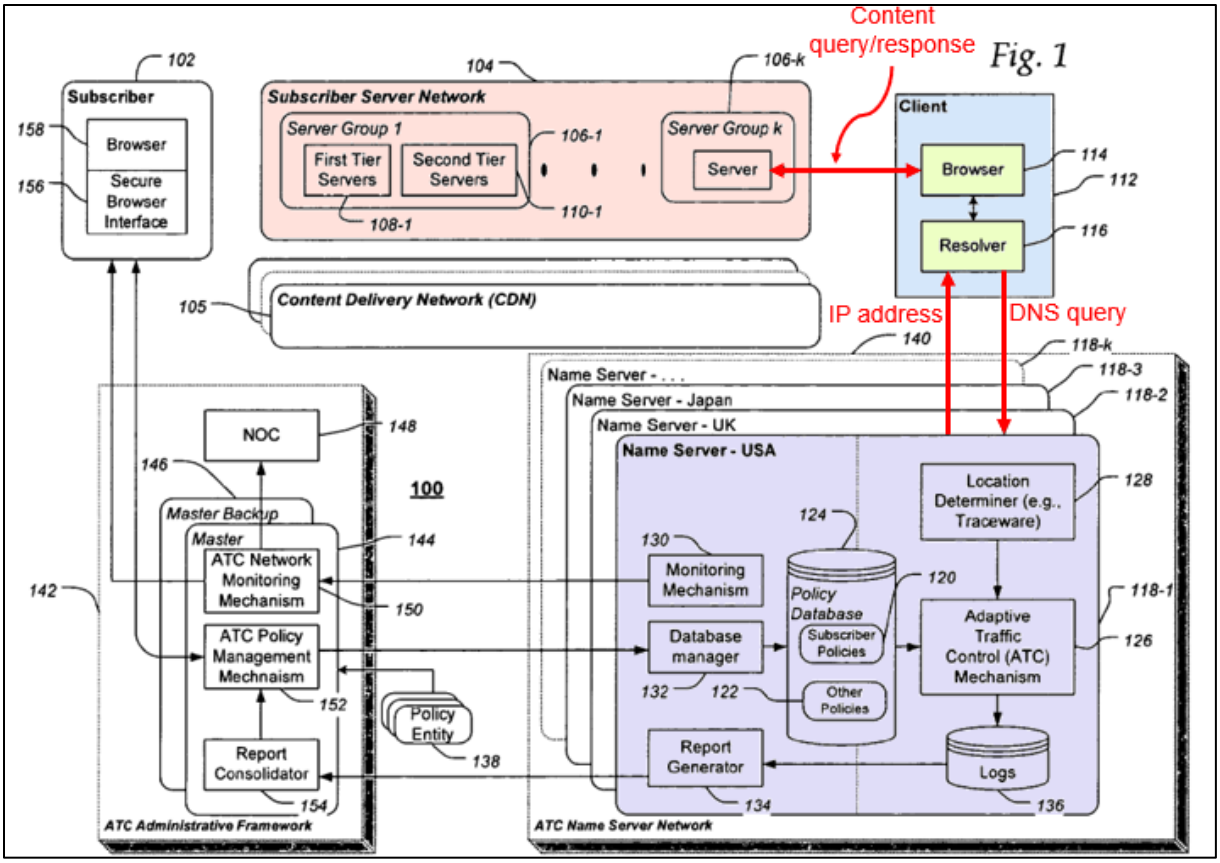

EX1001, FIG. 1. (annotated)

36. More particularly, the '517 patent discloses policies that are applied in response to receiving DNS requests in order to resolve the requests. It also discloses a graphical user interface (GUI) for setting these policies. The policies "may be hierarchically constructed to form a decision tree," also referred to as a "policy hierarchy." EX1001, 8:55-56. The decision tree comprises "rules that the user (content provider/subscriber) wants to apply" in resolving DNS requests. EX1001, 9:1-3. The "leaves" or "resource nodes" of the decision tree "specif[y] the answers (IP addresses or CNAMEs) that the user wants [] to provide in response to a DNS request." EX1001, 9:3-6. The "branches" or "branch nodes"

"specif[y] the various decision criteria that the user wants to apply." EX1001, 9:6-8. Thus, the decision tree represents a policy that is applied in response to DNS requests.

37. "FIG. 10(b) shows an exemplary decision tree embedded in a set of geographical policies that guide how the traffic should be directed to a set of servers." EX1001, 25:8-10. A DNS server applies the policy represented by the decision tree in response to receiving a DNS request from an end-user. This decision tree includes a branch node that specifies decision criteria for splitting "between the United States and the rest of the world." EX1001, 25:10-13. That is, upon receiving the DNS request, the DNS server determines whether the requester is located in the United States, and if so, applies one set of policies, and if not, applies another set of policies (i.e., "Splits on Country"). The decision tree includes another branch node specifying decision criteria for splitting based on time zone in the United States. EX1001, 25:16-20. That is, if the requestor is located in the United States, the DNS server further determines the requestor's time zone and applies different policies accordingly.

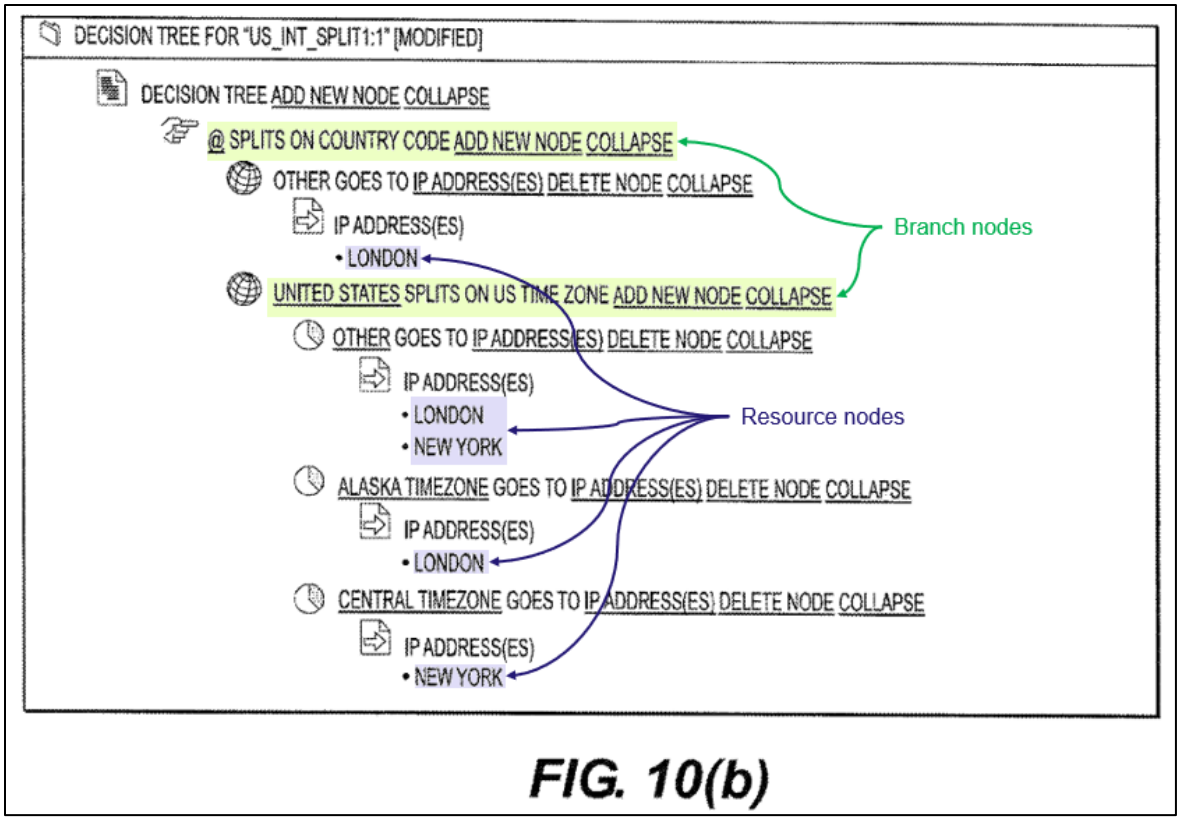

FIG. 10(b)

EX1001, FIG. 10(b). (annotated)

38. The decision tree also includes "resource nodes" specifying answers to be provided in response to the DNS request. The resource nodes include IP addresses of servers named "London" and "New York." When the DNS server reaches the leaves of the tree, after all decision criteria has been applied, it selects the IP address of the server represented by that resource node. EX1001, 25:14-25. For instance, if the requestor was located in the United States in the Central Time Zone, the DNS server would return the IP address of the New York server, but if the requestor was located in France, the DNS server would return the IP address for the London server. EX1001, 25:14-25.

39. As shown in Figure 10(b), the GUI for setting DNS policies enables

users to specify rules, including decision criteria (e.g., "Split on Country") for responding to requests in conjunction with corresponding answers for responding to requests (e.g., London, New York servers). EX1001, 25:8-25.

V. PROSECUTION HISTORY

40. The application that led to the '517 patent was filed on December 24, 2010. During prosecution, it was subject to multiple Office actions, adding to the substantive discussions concerning prior art across the entire family of applications where the '517 patent claims priority to two lines of applications that also received substantive examinations.

41. For example, the application was initially rejected in a Non-Final Office Action dated March 11, 2013. In this Non-Final Office Action, claims 2–5, 10–13, 18 and 19 were rejected as being anticipated by U.S. Patent No. 6,484,261 to Wiegel. Claims 9 and 17 were similarly rejected under as being unpatentable over Wiegel view of U.S. Patent Application Publication 2002/0145975 to MeLampy. Claims 6, 7, 14, 15, 20 and 21 were rejected over Wiegel in view of U.S. Patent No. 6,092,178 to Jindal. Claims 8 and 16 were rejected over Wiegel in view of U.S. Patent No. 6,052,788 to Wesinger. EX1002, 64-72.

42. Patent Owner overcame those rejections by explaining the novel features of the claims that were not present in the primary reference, Wiegel. Specifically, Wiegel did not disclose "having one or more resource nodes specify

one or more answers to be provided in response to a Domain Name Service (DNS) request." EX1002, 55. This unique node structure and function could not be found in the prior art. The Office then allowed the claims. *Id.*, 40.

43. The prosecution history of the '517 patent further shows that many of the references cited in the Petition were considered and even applied in rejections in predecessor applications. In particular, the Ben-Shaul reference cited in the Petition, *see* Pet., 2, 39–74, was applied in rejections during the prosecution of U.S. Application No. 10/259,497 ("'497 application") and U.S. Application No. 11/932,162 ("'162 application") in the priority chain. *See* EX2002; EX2003.

44. Patent Owner overcame the rejection in the '497 application by including the novel features of the claims indicated as allowable subject matter by the Examiner ("[c]laims 14, 15 . . . are objected to as being dependent upon a rejected base claim, but would be allowable if rewritten in independent form including all of the limitations of the base claim and any intervening claims") that were not present in the Ben-Shaul reference. *See* EX2002, 239, 290 ("[c]laim 10 has been amended to include the subject matter of objected-to claim 15," indicated as allowable subject matter by the Examiner, namely "a monitored server, wherein an availability of the monitored server is dynamically monitored" and/or "a managed server, wherein . . . a load share or a shed fraction of the managed server is defined dynamically"). Ben-Shaul's lack of dynamic monitoring and server

management also distinguish it from the policy and policy-based features of the '517 patent claims.

45. Patent Owner also overcame that rejection in the related '162 application by explaining that Ben-Shaul did not disclose or suggest "any such dependent relationship between the receipt of a client request for a domain name resolution, on one hand, and the retrieval of policies from a name server's policy database and the selection of a network based on those retrieved policies, on the other." *See* EX2003, 1316.

46. The Examiner then allowed the claims based on the amendments of the above features. *See* EX2003, 1331 ("[T]he '120 reference does not disclose the Applicant's invention as claimed in detail, such as receipt of a client request to resolve a domain name at a name server, the policy database being on the name server and selecting either the CDN or subscriber server. The '120 reference appears to be drawn to formatting content of a web page rather than load sharing/balancing of a network."). The specific role of the policy and policy-based features in the '517 claims similarly are not taught in Ben-Shaul.

VI. LEGAL STANDARDS

47. In formulating my opinions set forth in this Declaration, I applied the following legal principles.

A. Person of Ordinary Skill in the Art ("POSA")

48. I understand that patent law analyses are performed from the viewpoint of a POSA. I understand that a POSA is a hypothetical person who is presumed to be aware of all pertinent art, thinks along conventional wisdom in the art, and is a person of ordinary creativity. For purposes of this declaration, I have applied Petitioner's articulation of a POSA. *See* Pet., 8, EX1003, ¶¶50–52 ("A POSITA at the time of the invention would have had at least a bachelor's degree in computer science, electrical engineering, or a related field, and at least two years of work or research experience in the field of content delivery management or networks. Work experience can substitute for formal education and additional formal education can substitute for work experience"). It is my opinion that the definition of a POSA would not materially affect the analysis required for upholding the patentability of the challenged claims.

B. Claim Construction

49. I understand that in an IPR, patent claim terms are given their ordinary and customary meaning as understood by a POSA at the time of the invention, in view of the patent's specification and its prosecution history—unless the patent explicitly defines the claim term. I also understand that when a patent explicitly defines a claim term in the specification, the patent's definition controls.

50. I understand that a dependent claim contains all limitations of the claim from which it depends.

C. Obviousness

51. I understand that an obviousness analysis involves comparing a claim to the prior art to determine whether the claimed subject matter would have been obvious to a POSA in view of the prior art and general knowledge in the prior art at the time of the patent's priority date.

52. I understand that obviousness must be assessed from the viewpoint of a POSA at the time of the invention. I further understand that to establish obviousness, a party must perform the following factual inquiries: (a) determining the scope and content of the prior art; (b) ascertaining the differences between the claimed invention and the prior art; and (c) resolving the level of skill in the art. I understand that, in determining the scope and content of the prior art and ascertaining the differences between the claimed invention, a patent challenger must specify where each element of the claim is found in the prior art.

53. With respect to determining the proper scope of prior art to examine, I understand that for a prior-art reference to be properly used in an obviousness ground, the prior-art reference must be analogous art to the claimed invention. I have been told that a reference is analogous art to the claimed invention only if: (1)

the reference is from the same field of endeavor as the claimed invention; or (2) the reference is reasonably pertinent to the particular problem solved by the inventor.

54. I understand that factors relevant to determining the proper field of endeavor include the inventor's explanations of the subject matter (including the patent specification) and the claimed invention's structure and function. And I further understand that to be reasonably pertinent to the particular problem solved by the inventor, a prior-art reference must logically commend itself to the inventor's attention in considering his or her problem.

55. I understand that one way of establishing obviousness is by establishing that a POSA would have had both (i) a reason to modify or combine the teachings of the prior art to achieve the claimed invention and (ii) a reasonable expectation of success in doing so. I understand that the reason to combine prior-art references can come from a variety of sources, not just the prior art itself or the specific problem the patentee was trying to solve. And I understand that the references themselves need not provide a specific hint or suggestion of the alteration needed to arrive at the claimed invention; the analysis may include recourse to logic, judgment, and common sense available to a person of ordinary skill that does not need to be explicit in any reference.

56. I understand that a "reasonable expectation of success" is assessed in view of the prior art and general knowledge in the art from the viewpoint of a

POSA at the time of the patent's priority date, and it does not require an absolute certainty of success. Furthermore, I understand that determining whether the prior art would have suggested to those of ordinary skill in the art to carry out the claimed invention with a reasonable expectation of success must be done after determining that the claimed elements are present in the prior art.

57. I also understand that, before reaching a conclusion that the claimed invention would have been obvious, one must consider any objective evidence of non-obviousness if it is available. The objective evidence of non-obviousness can include evidence of commercial success attributable to the claimed invention, evidence of industry praise for the claimed invention, evidence of a long-felt need that was solved by the claimed subject matter, evidence that others copied the claimed subject matter, or evidence that the claimed subject matter achieved an unexpected, superior result relative to the closest prior art. I understand that such evidence must have a nexus, or causal relationship, to the claimed subject matter beyond what was available in the prior art and must be commensurate in scope with the patent claim(s) at issue.

D. Written Description

58. I understand that written description of a subject matter is assessed from the viewpoint of a POSA. I understand that to satisfy written description, the disclosures in the application must convey to a POSA that the inventors had

"possession" of the claimed subject matter. I understand that the inventor need not physically possess the described subject matter to have "possession" in the context of written description. I understand that possession for written description purposes is determined by the sufficiency of the disclosure in the application taken together, and in the context of, the state of the art at the time of the filing. I understand that the state of the art at the time of filing is determined by the general knowledge of a POSA, as evidenced by the published literature at the time.

59. I understand that the written description inquiry proceeds the same way whether one is analyzing the patent specification for determining the validity of the claims or analyzing a priority document for determining priority benefit.

VII. ARGUMENT

60. For the reasons provided below, it is my opinion that the cited prior art does not teach or suggest each and every element of any challenged claim. It is therefore my opinion that the Board should uphold the patentability of all challenged claims. For convenience, I focus my arguments herein on the elements of independent claim 1. But because independent claims 1, 9, and 17 are all substantially similar, my arguments herein apply equally to all claims.

A. Hasan is non-analogous to the '517 patent

61. As I explained above, I understand that, to be considered analogous art, Hasan must be (1) from the same field of endeavor as the '517 patent, or (2) if

not from the same field of endeavor, reasonably pertinent to the particular problem with which the ’517 patent inventors were involved. In my opinion Hasan is not from the same field of endeavor, nor is it reasonably pertinent to the particular problem of the ’517 patent.

1. Hasan is not from the same field of endeavor as the ’517 patent

62. The ’517 patent is directed to “a framework wherein resources of a content provider may be delivered to clients from different domains.” EX1001, Abstract. The specification and claims describe “a scenario in which DNS service is provided by a third party on behalf of a content or applications provider.” EX1001, 3:16-20, claim 1. For instance, independent claim 1 recites “an adaptive traffic control name server network [that] implements policy-based traffic direction” and includes a “domain name server… adapted to provide adaptive policy-based domain name service.” EX1001, claim 1. Claim 1 also recites “at least one policy for direction of network traffic” and a “decision tree” that “specif[ies] one or more answers to be provided in response to a Domain Name Service (DNS) request.” *Id*. The ’517 patent’s field of endeavor is thus DNS services for directing network traffic.

63. As I explained above, the ’517 patent describes the process for resolving DNS requests with respect to Figure 1. First, a user at a client device 12 “enters a resource locator, e.g., a Universal Resource Locator (‘URL’), into a

browser 114 on client 112." EX1001, 3:60-62. In response, a DNS resolver 116 on the client 112 "queries a particular name server (e.g., name server 118-1)" to determine the IP address that corresponds to the domain associated with that URL. EX1001, 4:35-44. The name server determines and "returns [] the IP (Internet Protocol) address… of a server… in the subscriber server network" that hosts the requested content associated with the domain. EX1001, 4:35-44.

64. As the '517 patent explains, multiple servers may host the requested content. *See* EX1001, 4:39-42, 4:60-67. The '517 patent describes and claims policies for determining which server will receive the network requests. For instance, "[t]he determination of the particular IP address returned to the resolver 116 may be based on a number of factors, including the resolver's location (e.g., as determined from the resolver's IP address) and various policies (e.g., subscriber policies 120, other policies 122) in a policy database 124." EX1001, 4:44-49. When the browser receives the returned IP address, it is able to direct its content query to the server having that address "in order to obtain the desired resource." EX1001, 4:49-51.

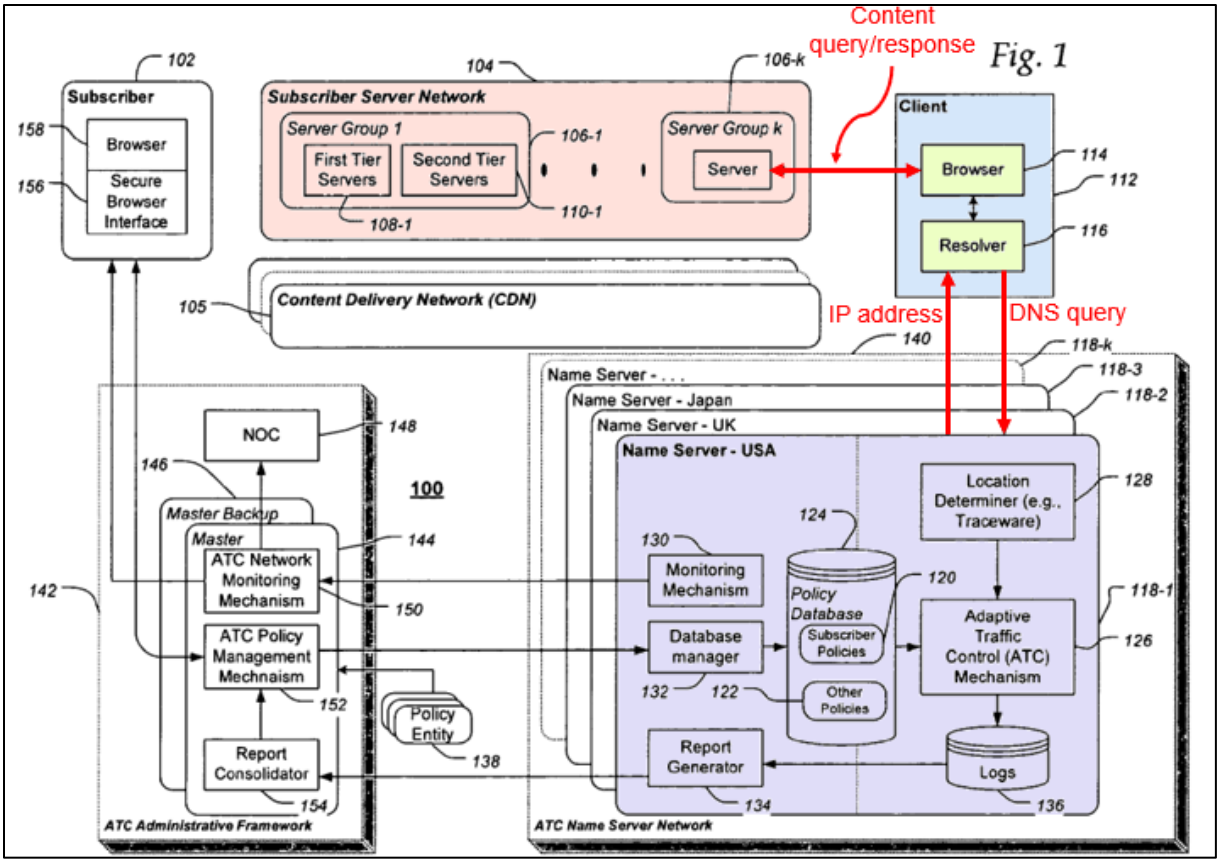

EX1001, FIG. 1. (annotated)

65. The ’517 patent thus discloses content delivery policies for resolving DNS requests to particular servers. The policies govern the direction of network traffic between end user client devices and servers hosting requested content. EX1001, 1:36-39. Indeed, the ’517 patent discloses that the claimed name server network “is responsible for responding to DNS requests, including processing DNS requests, applying the ATC policies 200 to select one or more servers from the subscriber server network (or a different designated network), and replying with the IP address(es) of the selected server(s).” EX1001, 16:43-48.

66. Hasan does not disclose a DNS system. It mentions DNS only one

time to explain that "DNS settings" are "implemented by other systems." EX1006, 17:57-63. Hasan is instead directed to a "virtual management system for a network facility, such as a data center, or any facility having a plurality of components which can be organized as objects for presentation in a virtualized environment." EX1006, Abstract. It describes its field of endeavor as the "management of network services in a data center, and in particular to management of services, subscribers, devices, log servers, and facilities using a common, virtualized management system." EX1006, 1:5-10 (describing the "Field of the Invention").

67. Hasan's virtual management system "includes a management topology presenting devices, facilities, subscribers, log servers, and services as objects to an administrative interface." EX1006, 3:46-50. It provides "centralized management to a data center having at least one facility and a plurality of devices configurable to provide network services to subscribers." EX1006, 4:9-12. It allows "a single system administrator [to] configure individual devices, log servers, or services for a subscriber, the facility, or for the entire data center" via a graphical user interface (GUI). EX1006, 6:28-31. The GUI provides a "view" of the data center, which "is a hierarchical layout of the data center, its facilities, subscriber, devices and/or log servers." EX1006, 8:19-22. The view allows a data center administrator to edit the properties of the data center objects. EX1006, 8:30-47. An example view is shown in Figure 6A.

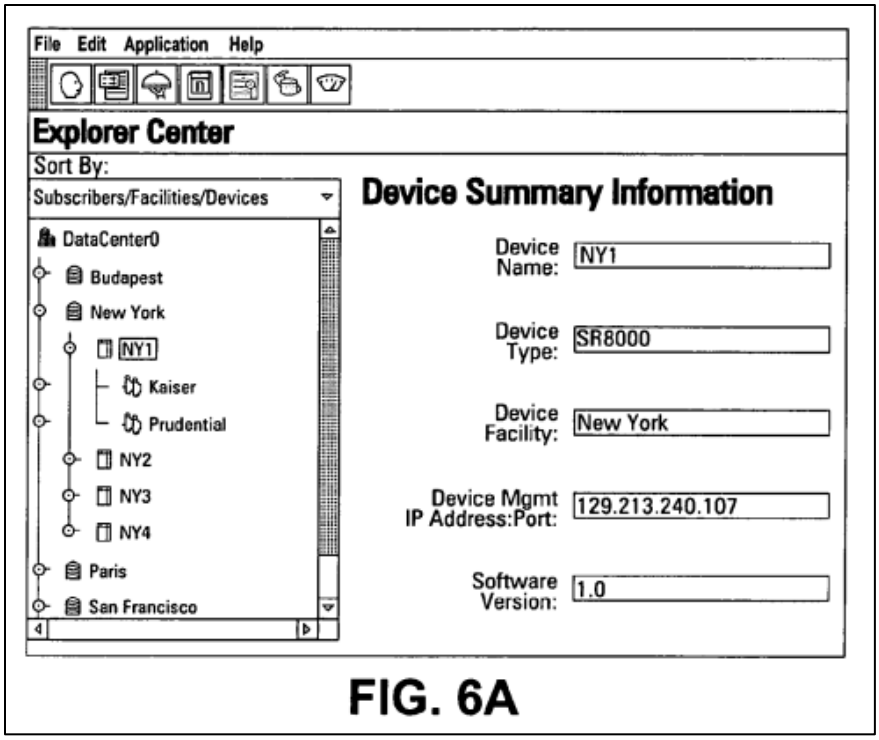

FIG. 6A

EX1006, FIG. 6A.

68. The view of Figure 6A "allows the configuration or viewing of attributes related to a specific device within a single data center facility regardless of the subscribers, whose traffic is being supported or services provided to the subscribers." EX1006, 9:1-4. Figure 6A "also indicates how many devices reside in the data center facility and how these devices will be used, say, for failover or multiple device load balancing." EX1006, 9:6-9. In other words, the view allows administrators to view the datacenter layout to perform "various management capabilities" on objects depicted in the view. EX1006, 8:32-47.

69. "Management of network services in a data center," as taught by Hasan, is unrelated to DNS services for directing network traffic, as taught by the '517 patent. While DNS is concerned with routing end-user requests to the

appropriate server hosting the requested content, Hasan's data center management is concerned with administrative services for the data center facilities and equipment. EX1006, 1:12-33. It is concerned with managing the infrastructure of the data center. EX1006, 1:12-33. To be clear, data centers can host and serve data to end users, but Hasan does not disclose any particular techniques for doing so. What's more, unlike the '517 patent, Hasan does not teach any techniques for directing end-user requests to a data center (i.e., DNS services). *See* EX1006, 17:57-63.

70. Although Hasan's virtual management system enables administrators to edit data center policies, these policies are not DNS policies for directing network traffic. *See* EX1006, 6:62-63, FIG. 4. In fact, as explained above, Hasan states that its system does not implement "DNS settings." EX1006, 17:57-63 ("DNS settings" are "implemented by other systems"). Rather, Hasan merely discloses a number of undefined, generic policies for data center administration, such as, for instance, "administrative policies", "security policies", "policy information", "data configuration polic[ies]", "subscriber policies", "VPN policies", "access control policies", "[f]irewall policies". EX1006, 6:55-57, 7:14-19, 7:52-55, 12:62-63, 13:7-9, 13:14-16, 14:52-55. It does not describe any specific policy; it only lists types of policies that can be implemented.

71. Although some devices managed by the virtual management system

are configured to implement network services, such as routing, firewall, VPN, etc., these are not DNS services. Moreover, Hasan does not disclose techniques for implementing such network services. *See, e.g.,* EX1006, 2:46-55, 9:38-43. For instance, Hasan generically describes "service bundles" that include "routing, firewall, NAT, VPN, PKI" services. EX1006, 9:38-41. But Hasan is not concerned with the implementation of such services. It is instead concerned with administering equipment that may be configured to provide these services. The solutions provided by Hasan are not solutions for implementing network services, such as those I described above—they are solutions for enabling administrators to manage and configure data center equipment. Hasan provides a general framework for such management via its virtual management system. *See* EX1006, 5:41-62.

72. Thus, Hasan is directed to the administration of data center facilities. It is not concerned with DNS services for directing end-user network requests to the appropriate servers. In my opinion, Hasan is not in the same field of endeavor as the claimed invention of the '517 patent.

2. Hasan is not reasonably pertinent to the particular problem to which the '517 patent relates

73. It is also my opinion that Hasan is not reasonably pertinent to the particular problem addressed by the '517 patent. The '517 patent explains that "[a]s the volume of Internet traffic grows, providers of web content and applications increasingly need to deliver content from multiple servers at widely-

separated locations in order to sustain a good end-user experience under high traffic loads." EX1001, 1:42-48. In particular, one of the "difficult challenges" of providing a good end-user experience on the growing internet is "how to control the way in which requests from end-users are distributed to each server according to important content provider policy constraints imposed for economic, contractual or other reasons." EX1001, 1:52-55. The '517 patent thus addresses the problem of how to distribute network requests from end users to the appropriate servers in a manner that provides a good end-user experience.

74. As I explained above, the '517 patent solves this problem by describing and claiming specific policies for directing network requests to the appropriate servers. EX1001, 1:58-60. These policies govern the server selection process based on factors such as user location, server availability, load sharing between servers, the state of the network, etc. EX1001, 1:60-67. The policies implement "an extremely fine-grained level of network Internet traffic control to providers of Internet content and applications, enabling them to dramatically improve the end-user experience (measured by speed of request resolution, associated download time, and the availability of servers) over that provided by conventional web servers and mirrored server farms." EX1001, 2:1-7. That is, the '517 patent is focused on solving problems relating to the direction of network traffic in geographically dispersed content delivery networks, such as the "speed of

request resolution, associated download time, and the availability of servers." *Id*.

75. As I explained above, Hasan is not concerned with delivering traffic to end user devices. It is directed to data center administration. The particular problem to which Hasan relates is "[t]he difficulty in administering a data center." EX1006, 2:59-62. Hasan explains that "the more subscribers one has in the data center, the more resources must be committed to administration." EX1006, 2:65-67. It further explains that "[t]his can become exceedingly difficult where changes must physically occur at the cage or via individual devices, especially in multi-facility data centers." EX1006, 2:67-3:3.

76. Hasan provides solutions "to alleviate some of these management problems." EX1006, 3:22-24. As explained above, its solution provides "a virtual management system for a network facility, such as a data center, or any facility having a plurality of components which can be organized as objects for presentation in a virtualized environment." EX1006, 3:39-43. Hasan explains that, "[o]nce organized into the virtualized management system of the present invention, management of the data center is streamlined and less user intensive." EX1006, 3:43-45.

77. "The users of the virtualized network management system are [] administrators," not end users. EX1006, 15:39-41. As I explained above, Hasan is not concerned with end-users, and it does not disclose any techniques for

delivering content to end-users or for resolving end-user DNS requests. Instead, Hasan is concerned with helping "administrators" to configure and manage data centers. It provides an "administrative interface" that enables administrators to view and configure data center objects, such as "devices, subscribers, facilities, log servers and data centers." EX1006, 3:46-56, 5:60-62. It is not concerned with delivering traffic to end users.

78. Hasan leaves "DNS settings" to be "implemented by other systems." EX1006, 17:57-63. This instructs a POSA not to look in Hasan for solutions to problems relating to DNS or resolving user network requests to appropriate servers. Rather, a POSA would have understood Hasan addresses problems relating to data center infrastructure. It addresses problems relating to the establishment and management of data centers. It does not address the same problems as the '517 patent claims.

79. Thus, Hasan is not reasonably pertinent to the problem addressed by the '517 patent.

B. Grounds 1 and 2

1. Petitioner and Dr. Lin fail to identify any "policy" in Beaumont

80. Petitioner and Dr. Lin do not identify any policy in Beaumont. The Petition points solely to Beaumont's "configurable preference table" as representing the claimed at least one policy. *See* Pet., 15 (Beaumont's DNS server

"directs requests (traffic) according to a configurable preference table (policies)"); *see also* EX1003, ¶72. In my opinion, this configurable preference table is not a policy, as claimed by the '517 patent, which must be a course, rule, or principle of action.

a. The claimed policy is a rule.

81. Independent claims 1, 9, and 17 (all independent claims) recite various features involving a "policy," specifically:

> A [(1) method/(9)administrative agent/(17)system], operable in a framework in which an adaptive traffic control name server network **implements policy-based traffic direction**, the name server network comprising at least one domain name server comprising hardware in combination with software and constructed and adapted **to provide adaptive policy-based domain name service**, the method comprising:
>
> . . .
>
> (B) **using said GUI to obtain at least one policy for direction of network traffic**, wherein the GUI supports the setting of said at least one policy using a decision tree representing rules; and
>
> (C) **providing said at least one policy to the name server network**, wherein the decision tree comprises one or more resource nodes, and one or more branch nodes, wherein the one or more resource

nodes specify one or more answers to be provided in response to a Domain Name Service (DNS) request, and wherein . . .

EX1001, 31:24-56.

82. The '517 patent is clear that a "policy" refers to a course, rule, or principle of action. *See* EX1001, 1:43-2:7 ("This invention solves these and other problems by providing a means to distribute network (e.g., Internet) traffic according to a configurable set of rules."); EX2015, ¶59. I understand that the PTAB agreed with this interpretation in its Institution Decision. Paper 14, 19 (the "understanding of a policy as a rule is consistent with the disclosure in the '517 patent that 'a policy may be created by generating a decision tree representing the ATC rules").

83. In the context of the claims, it is a course, rule, or principle of action to direct traffic of the underlying servers. EX1001, 1:43-2:7 ("These rules together provide an extremely fine-grained level of network Internet traffic control to providers of Internet content and applications … ."). The policy or policy-based features make decisions based on the geographical locations and the capacities of subscriber network servers to manage and direct traffic. *See* EX1001, 6:33-44 ("For example, the subscriber may set up policies to direct traffic with respect to considerations such as the geographical locations and the capacities of the underlying servers in the subscriber's network … ."), 7:13-21 ("Decisions are

based on location of the end-user or an approximate thereof, e.g., using the IP address of an end-user's resolver … ."), 7:22-28 ("The subscriber can explicitly specify the amount of traffic (load share) to be directed to each of their servers within a defined server set in the subscriber server network … .").

84. Importantly, the "policy" or "policy-based" features are structurally and functionally different from the "geographical locations and the capacities of the underlying servers," EX1001, 6:33-36, that the policy features use. In other words, the policy governs the process, while the geographical locations and the capacities of the underlying servers are the inputs that the policy uses to determine traffic direction. *See* EX1001, 6:33-44, 15:28-35.

85. The '517 patent further illustrates what it means by "policy" and "policy-based" course, rule, or principle of action through several examples. *See* EX1001, 7:13-8:25. One example is a "geographic policy," in which decisions are based on the location of the end-user or an approximate thereof, *e.g.*, using the IP address of an end-user's resolver. *See* EX1001, 7:13-21. In this example, the policy or policy-based course, rule, or principle of action is interpreted as "if the IP address is within the U.S., the request may be directed to a 'domestic' group of servers in the subscriber server network 104, otherwise it may be directed to 'international' servers in the network." *Id*.

b. Petitioner and Dr. Lin do not identify the claimed policy in either Beaumont or Hasan

86. Claim elements [1.b] and [1.c] require a "policy for direction of network traffic." It is my opinion that Petitioner and Dr. Lin have not identified the claimed policy in the cited prior art.

87. The Petition and Dr. Lin point solely to Beaumont's "configurable preference table" as representing the claimed at least one policy. *See* Pet., 15 (Beaumont's DNS server "directs requests (traffic) according to a configurable preference table (policies)"); EX1003, ¶72. This preference table is not a policy, nor does it contain policies. It is merely a data structure that stores data. EX1007, ¶¶23, 32. It does not specify any rules, such as via a conditional and a result. *See* EX1001, 8:55-56, 15:28-32, 16:5-7. Beaumont's preference table merely stores information, such as locations or statuses of a server, such as delivery server address, destination domain, estimated distance, number of times being selected per hour, health, and time-to-live. *See* EX1007, FIG. 3. This is just data stored in a "data structure." EX1007, ¶23.

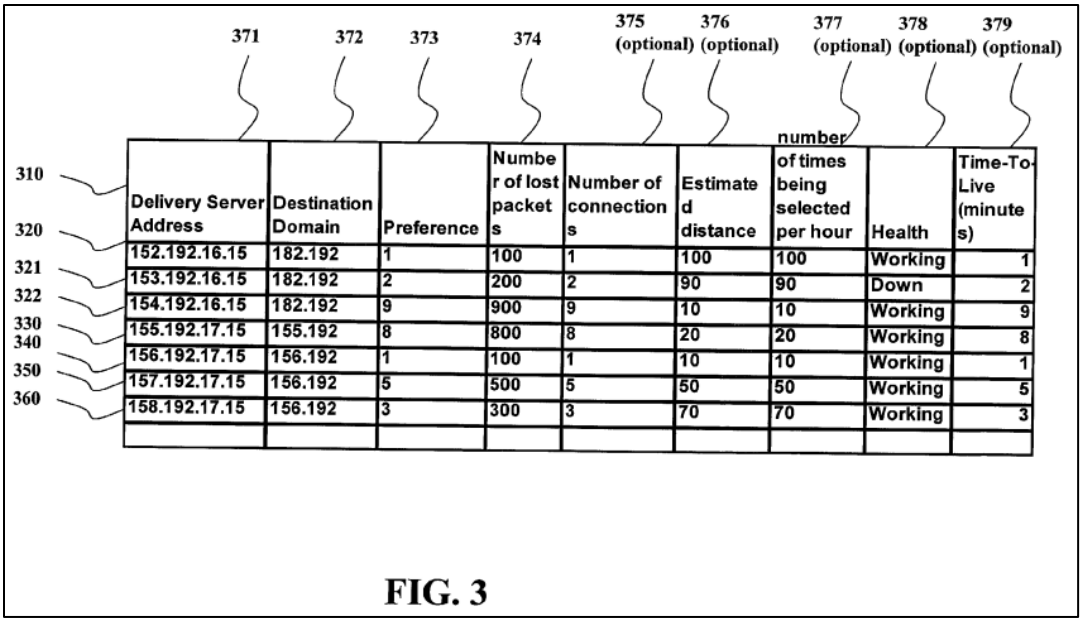

FIG. 3

EX1007, FIG. 3.

88. The preference table does not provide the claimed course, rule, or principle of action required for a policy. It does not store software or any other logic for implementing such course, rule, or principle of action. It does not govern any process or direct traffic to the content delivery servers as claimed. It is a listing of information. This information could be used as input to a policy—but neither Petitioner nor Dr. Lin point to any such policy in Beaumont.

89. In the ’517 patent claims, information is used to navigate a decision tree by following the policy rules. *See* EX1001, 7:13-21, 9:13-25. The information in Beaumont’s preference table could be used by a policy to determine which group of servers to direct the request, but Beaumont’s preference table itself does not expressly or inherently disclose any policy. For instance, Beaumont discloses “rule[s] for selecting the best content delivery server” and a “selecting module”

that implements software instructing Beaumont's DNS to "retrieve[] from preference database 453 a preference number associated with each retrieved mapped address and the requester identifier," and "select[] an address based on the retrieved preference numbers." EX1007, ¶¶31-32. To the extent Beaumont discloses a policy, it is specified by these rules. The rules provide instructions or regulations that govern the DNS's behavior in selecting content delivery servers. Conversely, the preference table only stores data that can be used as inputs to these rules.

90. The data stored in the preference table do not constitute policies either. For instance, the preference numbers stored in the table are merely integers representing user preference, and the destination domains stored in the preference table merely represent data specifying destination domains associated with certain servers. These values do not constitute rules or instructions that govern server selection. The system implements separate software-based "rules" for this purpose. The preference numbers and destination domains act as inputs for the rules, but they are separate and distinct from the rules themselves. EX1007, ¶¶31-32. Indeed, as I have explained, the rules are specified in software (e.g., in the selecting module), not in the preference table. The preference table does not provide any action or rule of directing network traffic, nor does the data stored therein.

91. Indeed, a rule, or policy, must specify some instruction or regulation

governing server selection. The preference table data does not itself provide any such instruction—because it is merely data. The system does not and cannot select a server without the underlying rules. The software-based rules instruct the system as to how to select a server. The preference table data (e.g., preference numbers) at best represents data used in applying the underlying rules. *See* EX1007, ¶¶31-32.

c. Beaumont-Hasan does not teach using a GUI to set or obtain policies for direction of network traffic.

92. Claim element [1.b] recites "using said GUI to obtain at least one policy for direction of network traffic, wherein the GUI supports the setting of said at least one policy using a decision tree representing rules." The cited prior art does not teach or suggest this element.

93. Neither Beaumont nor Hasan teach a GUI that is used to set or obtain the claimed policy for direction of network traffic, as recited in element [1.b]. Beaumont does not teach a GUI at all, let alone one for setting policies. And although Hasan teaches a GUI for editing generic data center administration policies, none of its policies are for directing network traffic in response to DNS requests, as claimed.

94. Beaumont does not disclose a GUI for obtaining or setting policies. At best, Beaumont discloses an interface for manually inputting "preference numbers" that are stored in its preference table. EX1007, ¶¶31, 34. But as I have explained, these preference numbers are not policies or rules—they are merely integers. They

represent data stored in a data structure that can be used by Beaumont's rules for selecting content delivery servers. The preference values themselves do not provide or represent any logic or rule for selecting a server, they are merely parameters that serve as inputs for Beaumont's "rules for selecting the best content delivery server." *See* EX1007, ¶31.

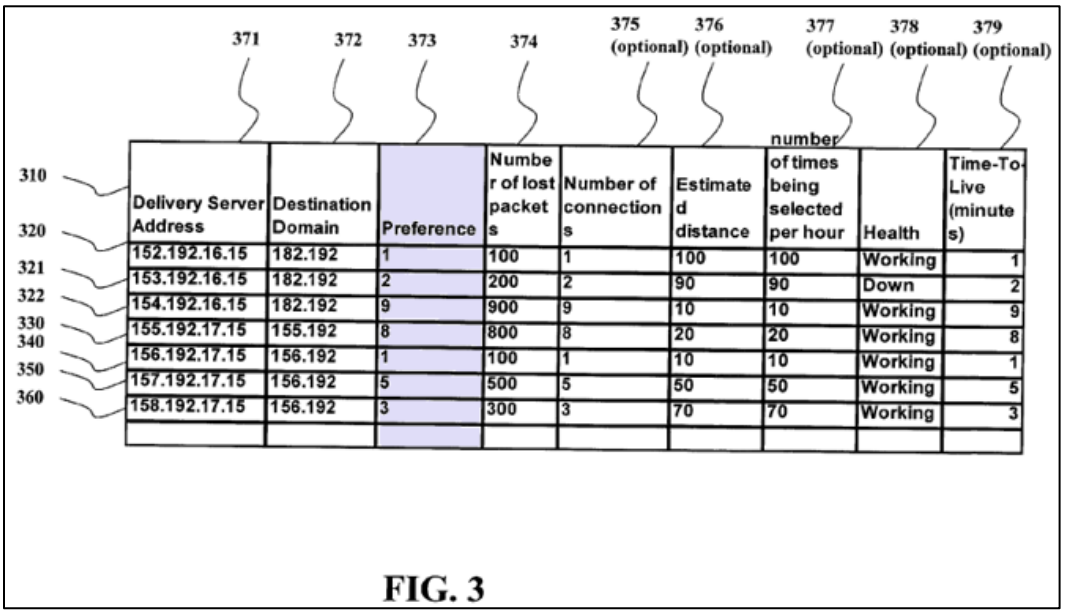

FIG. 3

EX1007, FIG. 3.

95. To the extent Beaumont discloses policies at all, these "rules for selecting the best content delivery server," not the preference numbers, represent policies. But the rules are not set or obtained through a user interface. Rather, they are merely implemented as software via the "selecting module 414" that is "recorded in a secondary memory… before being loaded into the main memory." EX1007, ¶¶32-33. A user cannot change these rules using a user interface. Beaumont teaches no mechanism to do so. It does not disclose a user interface for

inputting policies.

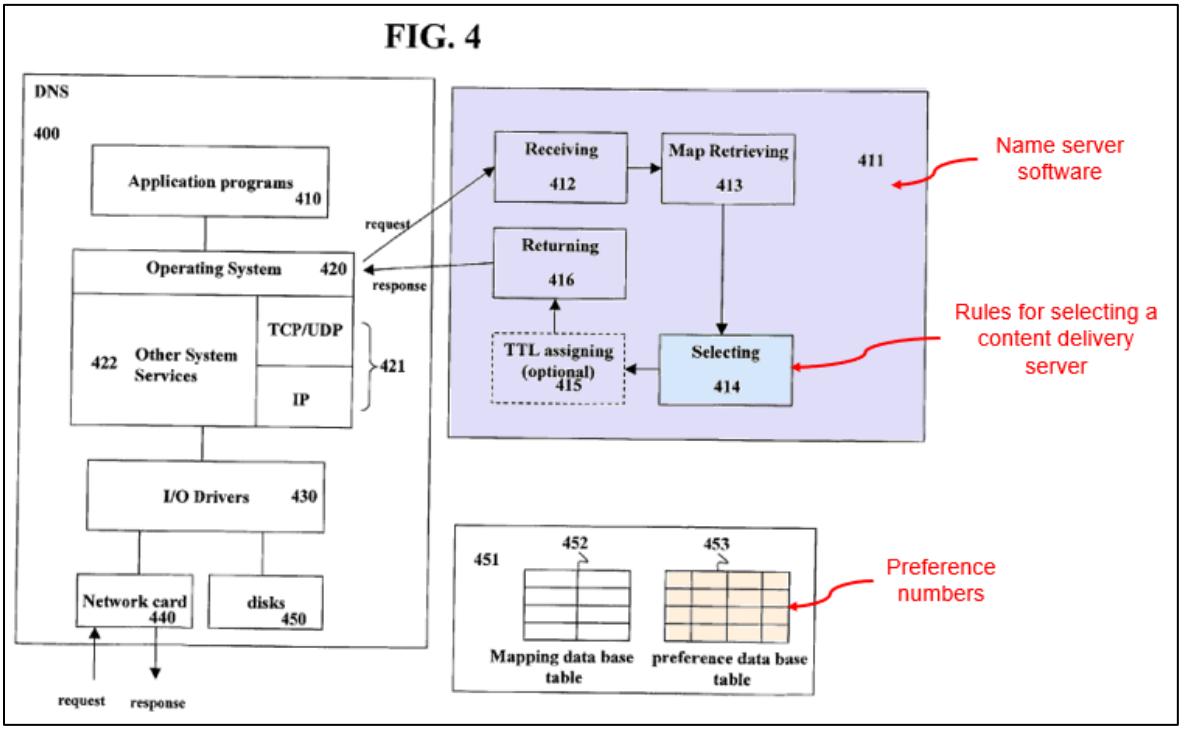

EX1007, FIG. 4. (annotated)

96. What's more, inputting preference values does not affect the underlying rules specified via software. Beaumont discloses a rule of selecting the content delivery server having the "highest preference number." EX1007, ¶32. When these preference numbers are manually input, the user can influence which servers will be selected by giving greater preference to certain servers, but the underlying rule is unchanged—the rule is still to select the server with the highest preference number. The preference number merely acts as an input to the rule. The preference number is just data. The rule represents an instruction governing the conduct of the DNS in response to a DNS request. The preference value on its own does not govern any such conduct. Without the rule specified by software, the DNS

would have no way of selecting a content delivery server.

97. Hasan does not fix Beaumont's deficiencies. The Petition does not rely on Hasan for any disclosure of policies or policy-based features. Nor does Hasan disclose the claimed policy in any event. Although Hasan's "policy editor" is a GUI for setting policies, it does not set or obtain policies for direction of network traffic, as claimed. Hasan is not concerned with DNS. It does not disclose policies for directing network traffic. Rather, it merely discloses a number of undefined, generic policies for data center administration, such as, for instance, "administrative policies", "security policies", "policy information", "data configuration polic[ies]", "subscriber policies", "VPN policies", "access control policies", "[f]irewall policies". EX1006, 6:55-57, 7:14-19, 7:52-55, 12:62-63, 13:7-9, 13:14-16, 14:52-55. It does not define or otherwise elaborate on these policies. It does not describe the structure or logic of the policies. It merely lists undefined types of policies that may be helpful in implementing data center administration tasks. Thus, Hasan does not disclose a GUI for setting or obtaining the claimed policies.

98. As neither Beaumont nor Hasan teach the claimed GUI, it is my opinion that their combination likewise does not teach or suggest this element. Petitioner and Dr. Lin's proposed combination is premised on Beaumont's manually entered "preference numbers." Pet., 13-14 ("In view of Beaumont's

teaching of manually entered rule preferences [], a POSITA would have been motivated to implement Hasan's [] teachings in Beaumont"); EX1003, ¶69. According to Petitioner, the resulting combination allows users to "conveniently configure *preferences for Beaumont's DNS servers*." Pet., 14. Petitioner alleges that, in the combined GUI, "selecting a server would launch a menu for assigning the preferences shown in Fig. 3 of Beaumont." Pet., 18, *see also* 20 ("Hasan's GUI would enable users to specify [Beaumont's] preferences"); *see also* EX1003, ¶77.

99. But again, these preferences are merely values that represent user preference. They are separate and distinct from the underlying *rules* that use the preferences. At best, Petitioner and Dr. Lin's combination includes a GUI for inputting these preference values. It does not include a GUI used "to obtain at least one policy for direction of network traffic," or that "supports the setting of said at least one policy," as claimed. Indeed, because Beaumont's rules for selecting a content delivery server are set in software, these policies are obtained and set prior to, and completely independent from, the manual entry of Beaumont's user preferences. *See* EX1007, ¶¶31-32. A POSA would have understood that Beaumont's software-based rules are set and stored in memory prior to the system's deployment, and would not be changeable by users of Beaumont's system. *See* EX1007, ¶33.

100. Thus, because Beaumont only teaches manual entry of user

preferences, and Hasan only teaches manual entry of generic data center administration policies that are unrelated to the claims, their combination at best teaches or suggests a GUI for manually inputting preference numbers. The combined Beaumont-Hasan system thus does not teach or suggest a GUI for setting or obtaining the specifically claimed DNS policies.

d. Beaumont-Hasan does not teach providing said at least one policy to the name server network.

101. Claim element [1.c] recites "providing said at least one policy to the name server network." The cited prior art does not teach or suggest this element.

102. In my opinion, neither Beaumont nor Hasan teaches providing said at least one policy to the name server network, as claimed in element [1.c]. The independent claims require "using said GUI to obtain at least one policy for direction of network traffic" (element [1.b]) and "providing [the same] at least one policy to the name server network" (element [1.c]). The combination of Beaumont and Hasan does not teach these elements.

103. Petitioner and Dr. Lin point to Beaumont's preference table as the alleged policy that is provided to the name server network. *See* Pet., 18; EX1003, ¶72. But again, this preference table is not a policy as claimed. Thus, even to the extent the data stored in this preference table (e.g., the preference values) are provided to Beaumont's DNS server, this does not satisfy this claim element.

104. To the extent Beaumont's "rules" for selecting content delivery

servers are considered policies, they do not satisfy this element either. This element requires that the at least one policy that was obtained and set using the GUI be provided to the name server network. But as explained above, Beaumont-Hasan's rules for selecting content delivery servers are not set or obtained using a GUI. Beaumont-Hasan does not teach a policy that was both obtained and set using a GUI and also provided to the name server network. Beaumont-Hasan thus fails to teach "providing said at least one policy to the name server network," as claimed.

2. Beaumont-Hasan does not teach the claimed GUI that uses a decision tree

105. Petitioner and Dr. Lin's combination does not teach or suggest the claimed GUI that "supports the setting of said at least one policy using a decision tree representing rules," as recited in element [1.b]. The decision tree must "represent[] rules," and it also must comprise "one or more branch nodes" that "specify one or more decision criteria to be applied." Neither Beaumont nor Hasan teaches these elements, alone or in combination.

a. Beaumont-Hasan's GUI does not support the setting of policies using a decision tree representing rules

106. Claim element [1.b] recites that "the GUI supports the setting of said at least one policy using a decision tree representing rules." The cited prior art does not teach or suggest this element.

107. Neither Hasan nor Beaumont, alone or in combination, discloses a

GUI that supports the setting of policies using a decision tree representing rules, as recited in element [1.b]. Petitioner and Dr. Lin argue that Beaumont-Hasan's "combined GUI" teaches this element "because it is presented in a 'familiar tree structure' and would allow users to set server preferences and rules." Pet., 18; EX1003, ¶¶76-77. I disagree. Beaumont-Hasan does not teach a decision tree at all, let alone one that is used by a GUI to support the setting of a policy.

i. The references do not teach this element individually

108. The '517 patent describes the decision tree as a "policy hierarchy." EX1001, 8:56-59; *see also* Pet., 4 ("These policies are constructed as a hierarchical decision tree"). It states that "[adaptive traffic control (ATC)] policies [] may be hierarchically constructed to form a decision tree," and that "a policy may be created by generating a decision tree representing the ATC rules that the user (content provider/subscriber) wants to apply." EX1001, 8:55-56, 9:1-3. The independent claims require a GUI that uses such a decision tree to support the setting of a policy. An example of such GUI using a decision tree is provided in FIG. 10(b). EX1001, 25:8-25.

109. Hasan does not teach a "decision tree" as claimed. Hasan's GUI merely presents "a hierarchical layout of [a] data center, its facilities, subscriber[s], devices and/or log servers." EX1006, 8:19-22. Hasan refers to this hierarchical layout as a "view." EX1006, 8:19-22. The "view" is "normally shown as the

familiar tree structure." EX1006, 8:22-25. It represents "an organization of data objects" and is used to enable an administrator to "perceive[] the virtualization topology" of the data center. EX1006, 6:54-56, 8:19-22.

110. This view is not a decision tree. That it is presented in a "tree structure" does not make it one. It is merely provided as an interface so that "an administrator can view" and configure the data center objects. EX1006, 8:50-67. Indeed, Petitioner and Dr. Lin acknowledge that the "tree structure" of the view is implemented merely to "organize information." Pet., 14; EX1003, ¶71. Hasan's GUI merely enables administrators to configure data center objects. Although it includes a "policy editor," it does not explain how this tool is used, nor does it describe any specific policy that can be edited. *See* EX1006, 6:62-63, 7:17-19.

111. Hasan's GUI thus merely represents a hierarchy of data center objects. It is not a hierarchy of *policies*, as described by the '517 patent. In Hasan, a user navigates to and selects an object in the GUI to perform "various management capabilities" *on that object*, such as viewing or editing the properties of the object. EX1006, 8:30-34. In other words, Hasan's GUI at most allows users to edit policies for respective data center objects. The policies themselves are not hierarchical. They are not related nor do they depend on one another. Indeed, as I have explained, Hasan does not describe the structure or underlying logic of any policy. It merely lists *types* of policies that could be implemented for the various

data center objects.

112. Moreover, Hasan's GUI does not represent rules, nor does it have any underlying logic for implementing a decision tree. The structure of the topology (e.g., the tree structure) does not itself represent a policy hierarchy, nor does it represent any rules, as required by the claims. Rather, the tree merely provides a convenient "view" of the data center objects.

113. Hasan's view enables users to select data center objects (e.g., devices) and input policies for the respective objects, but the view itself does not represent a decision tree. Indeed, Hasan discloses many different views showing different hierarchies of the data center objects (e.g., different objects represented as roots and branches in the hierarchy). *See, e.g.,* EX1006, 9:44-10:24, FIGs. 6(a)-(e). The user can also arrange the data center object hierarchies in various different ways. EX1006, 10:39-40, *see also* 7:48-51. Changing these hierarchies does not change the underlying policies for these objects because the views themselves do not represent rules. They merely organize data center objects in different manners, independent of the underlying policies for these objects.

114. Beaumont does not fix Hasan's deficiencies. It does not disclose any GUI, let alone the claimed GUI. Although Beaumont teaches manual entry of preference values, Petitioner acknowledges that it "does not describe a specific interface" for doing so. Pet., 14; *see also* EX1003, ¶70. Beaumont also does not

disclose the claimed decision tree. Although its preference table can be implemented as a "tree," EX1007, ¶23, this is not the decision tree as claimed. Rather, the tree is a "data structure" that represents a different arrangement of the preference table data. This tree data structure stores the same data as the preference table. Like the preference table, it does not store software, rules, or instructions that govern the content delivery selection process. And just as the preference table does not represent rules as claimed, the preference tree would also not represent such rules, nor would it represent a policy hierarchy. *See* Section VII.B.1 *supra*. Thus, even to the extent Beaumont does suggest a GUI for inputting preference values, such GUI does not "support[] the setting of said at least one policy using a decision tree representing rules," as claimed.

ii. The combination of Beaumont and Hasan does not teach or suggest the claimed GUI or decision tree

115. As neither Hasan nor Beaumont teach the claimed GUI, their combination doesn't either. Petitioner and Dr. Lin argue that the "combined GUI" of Beaumont-Hasan teaches the claimed decision tree "because it is presented in a 'familiar tree structure' and would allow users to set server preferences and rules." Pet., 18; EX1003, ¶¶76-77. But the claimed decision tree requires more than this. As I have explained, it requires a policy hierarchy.

116. As I explained above, Petitioner's combination amounts to using

Hasan's GUI to generate a topology of content delivery servers and to input Beaumont's preference values for the respective servers. *See* Pet., 14, 18, 21. Specifically, Petitioner and Dr. Lin argue that "the combined GUI would [] allow grouping content delivery servers by IP address block or geographic zone," and that "selecting a server would launch a menu for assigning the preferences shown in Fig. 3 of Beaumont." Pet., 21, 18; *see also* EX1003, ¶¶77, 89.

117. But, as with Hasan's GUI, this merely amounts to a topological layout that enables users to input respective preference values for each server. The view presented in the GUI does not represent rules. In the combined system, the "rules" for selecting the best content delivery server are represented in the name server software, as taught by Beaumont. EX1007, ¶¶31-32. The GUI thus merely represents a topology of servers, as taught by Hasan. EX1006, 8:19-29. The user is unable to affect these rules using the GUI. Indeed, Petitioner and Dr. Lin do not argue that the combined GUI would enable users to specify the underlying decision criteria to be applied for DNS policies (e.g., select servers based on destination domain of the requestor). They only argue the ability to "group[] content delivery servers" in certain manners. Pet., 21; *see also* EX1003, ¶¶81, 89. But as I have explained, grouping content delivery servers using a GUI does not impact the underlying rules or policies associated with those content delivery servers. Neither Beaumont nor Hasan teaches such functionality.

118. Thus, even under Petitioner and Dr. Lin's own argument, the combined GUI would merely enable administrators to group content delivery servers based on the destination domain of the requestors, and then to input preference values for the respective servers. But this tree structure itself does not represent rules. It only represents a layout of content delivery servers arranged in a tree-like manner. Again, the structure of the topology (e.g., the tree structure) does not itself represent any rules. Rather, the tree merely provides a convenient "view" of the data center objects.

b. Beaumont-Hassan's GUI does not support the setting of policies using a decision tree that includes branch nodes specifying decision criteria

119. Claim element [1.d] recites "wherein the decision tree comprises one or more resource nodes, and one or more branch nodes, wherein the one or more resource nodes specify one or more answers to be provided in response to a Domain Name Service (DNS) request, and wherein the one or more branch nodes specify one or more decision criteria to be applied." The cited prior art does not teach or suggest this element.

120. Even if Beaumont-Hasan's GUI taught a decision tree, it does not teach a decision tree that includes branch nodes that specify decision criteria to be applied in response to DNS requests, as recited in element [1.d]. Neither Beaumont nor Hasan teaches such a GUI, nor does their combination.

i. Neither Beaumont nor Hasan teaches or suggests a decision tree that includes branch nodes specifying decision criteria

121. Claim element [1.d] recites that "the decision tree comprises… one or more branch nodes…[that] specify one or more decision criteria to be applied." The cited references do not teach or suggest this element. To the extent Beaumont teaches decision criteria, it does not teach a GUI. And although Hasan teaches a GUI, it does not teach a GUI that uses a decision tree that specifies decision criteria.

122. First, regarding Beaumont, Petitioner and Dr. Lin cite Beaumont's "server health (e.g., working or down)" and "destination domains (e.g., 182.192)" as the claimed decision criteria. Pet., 19; EX1003, ¶79. Petitioner argues that Beaumont suggests a "decision criterion" in its depiction of rows R2, R3, and R4 in the configurable preference table (Beaumont's Figure 2) because they "are mapped to the same destination domain, 182.192." Pet., 19-20; *see also* EX1003, ¶81. But the information in the preference table is not decision criteria. The elements of rows R2, R3, and R4 of Beaumont's preference table are—at most—listings of the locations, capabilities, or statuses of a server. *See* EX1007, FIG. 2.

123. This is consistent with the '517 patent, which describes the decision criteria as the underlying rule applied via the policy. The decision criteria represent the actual rule, not just the parameters upon which the rule is applied. For instance,

example decision criteria include "*Split on* World zone" and "*Split on* country." EX1001, 9:17-25. Here, the decision criterion is to branch, or split, when navigating the decision tree based on the location of the requesting user (e.g., relative to a world zone or country). The decision criteria include the specific action to take (e.g., split) along with the parameters for which that action should be applied (e.g., world zone, country).

124. In Beaumont, the underlying decision criteria for selecting groups of servers is distinct from the information in the preference table. That an entry of the preference table includes a destination domain parameter does not itself provide any criteria for selecting a server. Beaumont's underlying logic, separate and distinct from the preference table, implements the criteria used to select the servers. For instance, Beaumont discloses "rule[s] for selecting the best content delivery server." EX1007, ¶31. The rules are not stored in the preference table. The rules are implemented in software. They instruct the system to retrieve information stored in the preference tables and select servers based on this information. EX1007, ¶31. For instance, these rules instruct the system to "retrieve[] the domain information of the local DNS from the request," and "use[] the retrieved domain as the destination domain in the preference database table." EX1007, ¶23, *see also* ¶32. The data in the preference table merely represents data that is retrieved based on the rules in order to select a delivery server.

125. Accordingly, the preference table listings do not provide decision criteria for selecting and grouping different content delivery servers. Beaumont's rules for selecting content servers are not configurable by a user via a GUI, they are stored in the underlying software implemented by Beaumont's system. Indeed, for the sake of argument, even if Beaumont's selection for rows R2, R3, and R4 was based on destination domains, then that selection *instruction* would need to be explicitly specified in the "one or more branch nodes" associated with the preference table (Beaumont's Figure 2), as claimed in '517 patent. *See* Pet., 20; *compare* EX1007, FIG. 2, *with* EX1001, FIG. 10(b). Beaumont's configurable preference table does not have this teaching or suggestion in any way.

126. What's more, even if Beaumont's destination domains (and server health) were decision criteria as claimed, Beaumont does not allow a user to input such decision criteria to be applied for DNS requests. It allows only manual entry of the preference values. *See* EX1007, ¶¶31, 34. Beaumont's destination domains and server health are not input by a user. Thus, Beaumont does not disclose a GUI, or any interface, supporting the specification of such decision criteria.

127. Hasan does not teach the claimed decision criteria either, nor does Petitioner or Dr. Lin allege that it does. Hasan does not teach DNS policies and indeed does not teach decision criteria for any policy. Thus, neither reference teaches a GUI that supports the setting of policies using a decision tree that

includes branch nodes specifying decision criteria to be applied.

ii. The combination of Beaumont and Hasan does not teach or suggest a decision tree that includes branch nodes specifying decision criteria

128. Despite neither reference teaching a decision tree with branch nodes specifying decision criteria (element [1.d]), Petitioner and Dr. Lin argue that the combined Beaumont-Hasan GUI teaches this element because, in the combined GUI, "content delivery servers would be grouped under the same destination domain and the destination domain would thus act as a decision criterion." Pet., 20; *see also* EX1003, ¶81. To the extent Petitioner and Dr. Lin argue that the act of grouping the content delivery servers using the GUI specifies the decision criteria, I disagree.

129. As explained, Petitioner's combination amounts to using Hasan's GUI to generate a topology of content delivery servers, and to input data stored in Beaumont's preference table (i.e., preference values). *See* Pet., 13-14, 20-21; *see also* EX1003, ¶¶77, 81, 89. Again, Petitioner does not argue that the combined GUI would enable users to specify the underlying decision criteria to be applied for DNS policies (e.g., select servers based on destination domain of the requestor). It only argues the ability to "group[] content delivery servers" in certain manners. Pet., 21; *see also* EX1003, ¶¶81, 89. But grouping servers in a hierarchy does not support the specification of the underlying decision criteria to be applied. Neither

reference teaches specification of such criteria, and neither does their combination.

130. In the combined system, as in Beaumont, the underlying decision criteria are specified in software and would not be affected using the GUI. *See* EX1007, ¶32. That is, arranging the servers in different manners using the GUI would have no effect on the underlying decision criteria to be applied in resolving DNS requests. As I have explained, Beaumont does not enable the user to specify such decision criteria using an interface. Nor does Hasan. Indeed, Hasan discloses many different views showing different hierarchies of the data center objects (e.g., different objects represented as roots and branches in the hierarchy). *See, e.g.,* EX1006, 9:44-10:24, FIGs. 6(a)-(e). The user can also arrange the data center object hierarchies in various different ways. EX1006, 10:39-40, *see also* 7:48-51. Changing these hierarchies (i.e., grouping the data center objects in different manners) does not change any underlying decision criteria for policies (or rules) in Hasan because the views themselves do not represent rules. They merely organize data center objects in different manners. Similarly, in the combined system, changing content delivery server hierarchies using the combined GUI would not change the underlying decision criteria or rules for selecting content delivery servers, because such decision criteria and rules are specified in software, not using the GUI.

131. In other words, setting or modifying content delivery server

hierarchies using the combined GUI would merely affect the presentation of these servers in the GUI. It would not have any effect on the underlying decision criteria for selecting servers in response to a DNS request. This is consistent with Hasan, which merely teaches "views" that present a topological layout of the data center. *See, e.g.,* EX1006, 6:54-56, 8:19-25. Just as Hasan's GUI does not specify any decision criteria for policies, the combined GUI would not either.

132. Petitioner and Dr. Lin argue that, in the combined GUI, "higher-level branches" are branch nodes that "represent decision criteria," such as destination domains. Pet., 19; EX1003, ¶79. But as I explained above, that a topological node in the GUI represents a destination domain does not mean that that node represents *decision criteria* to be applied to DNS requests. The claims require the GUI to support the specification of decision criteria. The combined GUI does not include any such functionality.

133. Petitioner and Dr. Lin further argue this element is met because "Beaumont's system first identifies working servers sharing the destination domain of the client (182.192) and then selects a server from within that set to handle the request based on a distance or other preferences." Pet., 19; *see also* EX1003, ¶80. But Beaumont does not allow the user to specify such decision criteria via an interface. It is implemented via software. *See* EX1007, ¶32. Given that Hasan does not teach such functionality for its GUI either, the combined GUI would not allow

users to specify such decision criteria.

134. Thus, grouping content delivery servers via the combined GUI as alleged does not support specification of the claimed decision criteria. As explained above, the combined GUI's "tree-like structure" merely organizes information. It does not itself represent rules or support specification of decision criteria to be applied for DNS requests.

3. A POSA would not have been motivated to combine Beaumont and Hasan

135. A POSA would not have sought to combine Beaumont and Hasan at all, much less in the manner proposed by Petitioner and Dr. Lin. *See* Pet., 13-14; EX1003, ¶¶68-71.

a. A POSA would not have combined Beaumont with Hasan at all

136. In my opinion, a POSA would not have combined Beaumont with Hasan. Beaumont is directed to a "content delivery system" comprising a domain name system (DNS) server that "receive[s] a name [e.g., domain name] from a requesting client (requester)" and selects the corresponding server that "is best to the requester." EX1007, Abstract, ¶5. Beaumont implements this DNS server using a "mapping database" that "generally comprises a table" in which "each entry… maps a name to an address." EX1007, ¶22. Upon receiving a DNS request, the server maps the name received with the request against the mapping database to

"retrieve[] the addresses identifying [] content delivery servers" mapped to the requested name. EX1007, ¶22. The DNS server then "selects one of those working mapped addresses (mapped content delivery servers) to deliver the content" by accessing a "preference database" that "preference number[s]" as to which server should be selected. EX1007, ¶¶23, 31. The preference numbers are "entered manually" by the user. EX1007, ¶31.

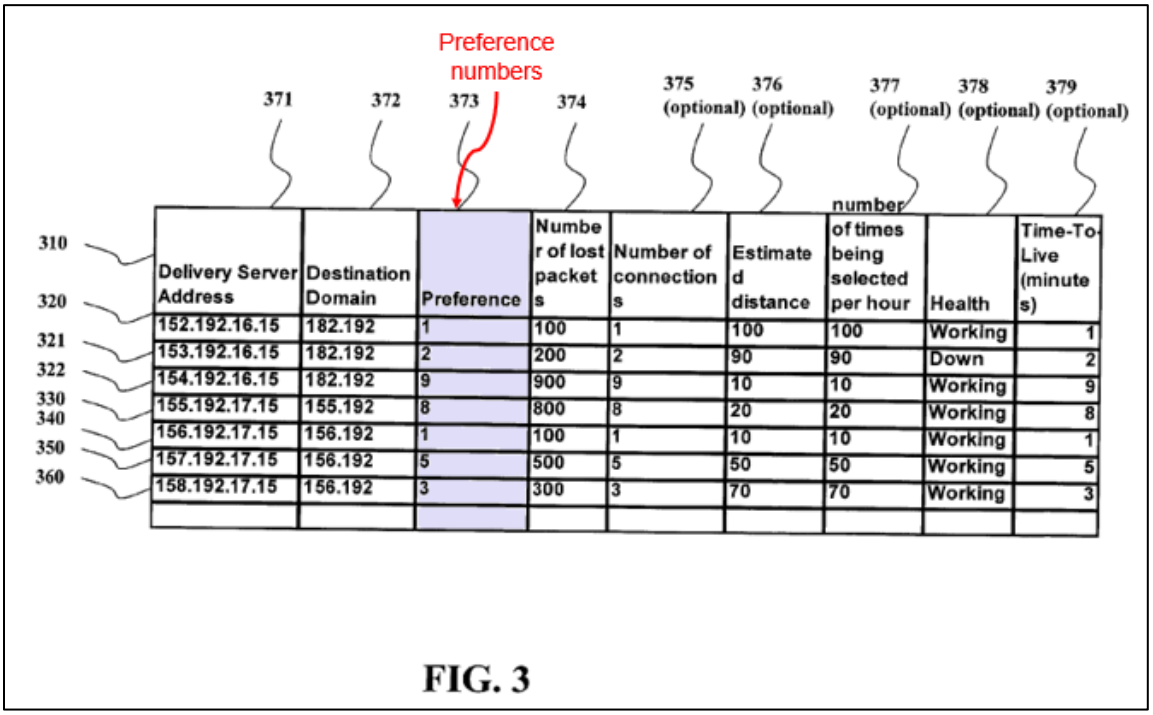

FIG. 3

EX1007, FIG. 3. (annotated)

137. Conversely, as I explained in Section VII.A, Hasan is not directed to DNS systems and does not address content delivery to users. Instead, Hasan is directed to data center administration. Hasan also does not address the problem to which Beaumont relates. In particular, Beaumont addresses problems with conventional DNS systems that implemented a "round robin" server selection

process, stating that, in such systems, "the best content delivery server might not be selected causing unnecessary delivery delay." EX1007, ¶¶4-5. Beaumont addresses this problem "by having a domain name server (DNS) receive a name from a requesting client (requester) and select one of the addresses mapped to the name, that is best to the requester." EX1007, ¶5.

138. In my opinion, a POSA would not have combined Beaumont with such a disparate reference as Hasan. Indeed, to the extent a POSA wished to modify Beaumont, one would have sought out references teaching content delivery and/or DNS techniques. For the reasons I have provided above, a POSA would have recognized that Hasan is unrelated to Beaumont's field of endeavor and the problem to which it relates. A POSA would have further recognized that Hasan's teachings are for a completely different purpose than Beaumont. Thus, in my opinion, a POSA would not have combined these references.

b. A POSA would not have combined Beaumont and Hasan as proposed by Petitioner and Dr. Lin.

139. Even if a POSA would have combined Beaumont and Hasan, it is my opinion that one would not have done so in the manner proposed by Petitioner and Dr. Lin. *See* Pet., 13-14; EX1003, ¶¶68-71. This combination amounts to using Hasan's GUI to generate a topology of content delivery servers and to input Beaumont's preference values for the respective servers. *See* Pet., 14, 18, 21; EX1003, ¶¶68-71, 77, 89. Specifically, Petitioner argues that "the combined GUI

would thus allow grouping content delivery servers by IP address block or geographic zone," and that "selecting a server would launch a menu for assigning the preferences shown in Fig. 3 of Beaumont." Pet., 21, 18; *see also* EX1003, ¶¶77, 89.

140. A POSA would not have combined Beaumont and Hasan in this manner. Beaumont does not allow users to input policies for resolving DNS requests. As explained, it teaches only user input of preference values. *See* EX1007, ¶¶31, 34. It does not allow users to specify groups of content delivery servers or to otherwise input any rules for resolving requests. A POSA would not have sought to implement such teachings in Beaumont at all, let alone based on the teachings of Hasan.

141. Aside from the preference values, the information used in the server selection process is gathered automatically and dynamically based on changing network conditions (e.g., network performance, round-trip-times, receipt of DNS requests). *See, e.g.,* EX1007, ¶¶26, 27, 31. This is not information that can be input by a user, because it reflects dynamic network conditions. *See, e.g.,* EX1007, ¶¶26, 27, 31.

142. What's more, as explained, Hasan does not disclose a GUI for configuring DNS services. A POSA would not have looked to Hasan to implement a GUI for configuring DNS services. Although Hasan teaches a GUI for creating a

data center topology, it does not teach creating server groups for purposes of resolving DNS requests. It merely presents a "topology," or "organization of data objects," in a data center for data center administration purposes. EX1006, 5:54-56.

143. A POSA also would not have modified Beaumont to enable users to specify actual DNS policies (i.e., rules) using Hasan's GUI. That is, even if a POSA would have combined Beaumont and Hasan to enable manual entry of Beaumont's preference numbers via Hasan's GUI, a POSA would not have been motivated to implement Beaumont's DNS *rules* (or policies) via Hasan's GUI. Neither reference teaches a GUI for inputting DNS rules or policies. Beaumont teaches only a manual entry of preference numbers, and Hasan's system does not include DNS services. *See* EX1007, ¶31; EX1006, 17:57-63.

144. Beaumont teaches "rule[s] for selecting the best content delivery server." EX1007, ¶31. To the extent Beaumont's system includes policies, the policies are specified in these rules. In one instance, the rule can instruct a DNS server to select a content delivery server based on the manually entered preference numbers. EX1007, ¶¶31, 34. The underlying rule, or policy, for selecting the server is distinct from the preference number itself. Beaumont does not teach a manual entry of rules or policies—only preference numbers.

145. And although Hasan teaches a GUI for editing policies, it does not

teach a GUI for editing DNS policies. As explained, Hasan discloses editing generic, undefined policies for data center administration. Hasan's management system does not provide DNS services. EX1006, 17:57-63.

146. Petitioner asserts that, in the combined Beaumont-Hasan GUI, "content delivery servers would be grouped under the same destination domain and the destination domain would thus act as a decision criterion." Pet., 20. In my opinion, a POSA would not have sought to enable users to specify any decision criteria for Beaumont's rules for selecting a content delivery server through the act of grouping content delivery servers in certain ways using Hasan's GUI. Neither reference teaches such functionality. Beaumont specifies its rules in software, not via a user interface. EX1007, ¶32. And although Hasan teaches a user interface for inputting policies, arranging data center facilities in different manners using its GUI does not impart any decision criteria for such policies. *See* Section VII.B.2, *supra*. A POSA would not have sought to implement such functionality in Beaumont in view of Hasan because neither reference teaches it.

147. Neither reference teaches a GUI for implementing DNS policies at all, let alone the specific policies recited in the claims. A POSA would not have sought to implement a GUI for inputting Beaumont's rules for selecting the best content delivery server. And even if one did, one would not have looked to Hasan for such teachings because Hasan does not disclose a GUI for inputting DNS policies.

C. Grounds 3A and 3B

148. In my opinion, Petitioner's Grounds 3A and 3B do not render obvious any challenged claim. As the '517 patent's independent claims are substantially similar, the below arguments relating to Grounds 3A and 3B apply to all claims.

1. Overview of Ben-Shaul

149. Ben-Shaul is directed to a "content and application delivery system" that "allows content providers to directly control the delivery of content." EX1005, Abstract. The system includes an "origin site… consisting of one or more origin Web servers" and "a set of subordinate edge servers which are geographically distributed and reside in the 'edges' of the internet." EX1005, ¶¶35-36. The system also includes an "edge DNS server" that performs DNS services in the system. EX1005, ¶38.

a. Ben-Shaul's Edge Servers

150. Ben-Shaul's edge servers "provide[] a variety of unique services on behalf of" their origin site, such as "caching, mirroring, content localization, service prioritization and application delivery." EX1005, ¶214. Edge servers coordinate with their origin web servers to download, or cache, content from the web servers. When a user requests content from an origin site, the edge server "intercepts" the request "to provide added targeted services relating to content." EX1005, ¶220. That is, end-user requests to the origin site are redirected to the appropriate edge server, and the edge server provides the requested content, along

with "added target services," back to the user.

151. Content providers "can control the behavior of a remote edge server by defining service profiles using a 'front edge' tool." EX1005, ¶215. An edge server service profile is "a collection of service and management directives that control the behavior of the edge server by enabling or disabling various services and by configuring these services with respect to service attributes, such as the target content, target time and target customers." EX1005, ¶220. Ben-Shaul also discloses a content delivery markup language (CDML) editor that enables users to specify such edge server directives. The edge server directives generally describe rules for downloading content from origin web servers and for delivering content to end users. *E.g.,* EX1005, ¶¶170, 196-208. Ben-Shaul describes a number of edge server directives in "Listings" 1 through 12. *See generally* EX1005, ¶¶193-209. Edge servers apply these directives to implement the functions specified therein.

152. For instance, Listing 1 describes a directive instructing an edge server to fetch a fresh copy of a web page from its origin server "as soon as [the] page expires in the edge server." EX1005, ¶193. The edge server executes this directive to download fresh copies of web pages that can be served to requesting users. This enables the edge server to maintain in its storage web pages that are up to date.

Listing 1

```
if (page expired)
then if (page was modified since cached)
then fetch new copy of page.
```

EX1005, ¶193.

153. As another example, Listing 2 describes a directive instructing edge servers in "certain regions [to] download a given URL [from the origin server] only after five o'clock p.m. local time." EX1005, ¶196. This directive instructs edge servers in the specified regions (e.g., Western Europe, middle East) to perform an action (e.g., "download" a document) (e.g., "news.html") only after the specified time (e.g., "17:00").

Listing 2

```
<?xml version="1.0"?>
<!DOCTYPE cdml SYSTEM "CDML.dtd" >
<document url-name="news.html">
 <temporal time="17:00" action="download">
 <regional region="Western Europe"/>
 <regional region="middle East"/>
 </terrtporal>
</document>
```

EX1005, ¶196.

154. As yet another example, Listing 7 describes "management directives" for an edge server group (i.e., having a "group_id" of "121"). Listing 7 instructs edge servers in the group to add specific domains to its configuration (i.e., <add rm domain action"Add" list=www.versedge.com,www.intra.versedge.com>), and to

prefetch data (i.e., <cache_prefetch type="zip">) corresponding to these domains from the origin server at a certain time. These directives allow the edge servers to obtain website data from the specified domains (e.g., www.versedge.com, www.intra.versedge.com) so that they can serve this data to end-users.

Listing 7

```
<?xml version="1.0" standalone"yes" ?>
<!DOCTYPE mgmt (View Source for full doctype . . . )>
<mgmt version="1111">
<mgmt_set version="001" group_id="121">
<directive id="1212" eid="12121EEEE">
<add_rm_domain action"Add" list="www.versedge.com,www.intra.versedge.com">
<discrete Repeat"5">
<time_element Time_Domain="GMT" Start_Date="01/01/2001" Start_Time"12:34"
/>
</discrete>
</add_rm_domain>
</directive>
<directive id="121212344" eid="121212FFFF">
<cache_prefetch tvpe="zip">
<url value="http://www.versedge.com/cdml/top.zip" />
<discrete Repeat="0">
<time_element Time_Domain="GMT" Start_Date="01/01/2000" Start_Time="00:00"
/>
</discrete>
</cache_prefetch>
</directive>
</mgmt_set>
</mgmt>
```

EX1005, ¶203.

155. Ben-Shaul's edge server directives are thus executed by edge servers to govern edge server conduct. For instance, edge server "service and management directives [] control the behavior of the edge server." EX1005, ¶220, *see also* ¶170 ("service directive" is an "request for service from one or more edge [servers]"). These directives govern edge server conduct in, for instance, downloading content

and delivering content to end users upon request.

b. Ben-Shaul's DNS

156. As explained above, Ben-Shaul also discloses an edge DNS (or EdgeDNS) server that provides DNS services within Ben-Shaul's system. EX1005, ¶41. "The EdgeDNS is responsible to return the IP address of either the edge server or another server for resources that whose URL resides in these domains." EX1005, ¶41. The EdgeDNS is a "conventional DNS name server" configured to route user requests to the appropriate server. EX1005, ¶359. In some embodiments, the EdgeDNS is adapted to implement "dynamic cache replacement" techniques that enable it to route user requests to the edge server, instead of the origin server, when the edge server hosts the requested content. EX1005, ¶¶360-361. This cache replacement technique enables the EdgeDNS to replace its resource record entry so that certain domains map to an edge server rather than the origin server. *E.g.,* EX1005, ¶¶361, 423.

157. In a representative embodiment discussed with respect to Figure 4, the EdgeDNS caches resource records mapping certain domains to the IP address of the EdgeDNS's corresponding edge server. EX1005, ¶176. A client regional DNS server 22 writes this mapping to the EdgeDNS. *Id*. In particular, "the client regional DNS server 22 conducts a periodic write of a regional domain name (a.regional.znn.com) DNS resolution with a binding to a local edge server IP

address." *Id*. "The writing is done into the resolution cache of the EdgeDNS 38." *Id*. This operation modifies the EdgeDNS's mapping for the regional domain name, switching the mapping to the edge server IP address from the origin server IP address. When a user requests content from one of these domains, the request is first routed to the client regional DNS server 22, which passes the request to the EdgeDNS. *Id*. Upon receiving the request, the EdgeDNS accesses its stored resource records and identifies from its cache the IP address (e.g., of the edge server) that is mapped to the domain. The EdgeDNS then returns the IP address to the requestor via the client regional DNS server 22.

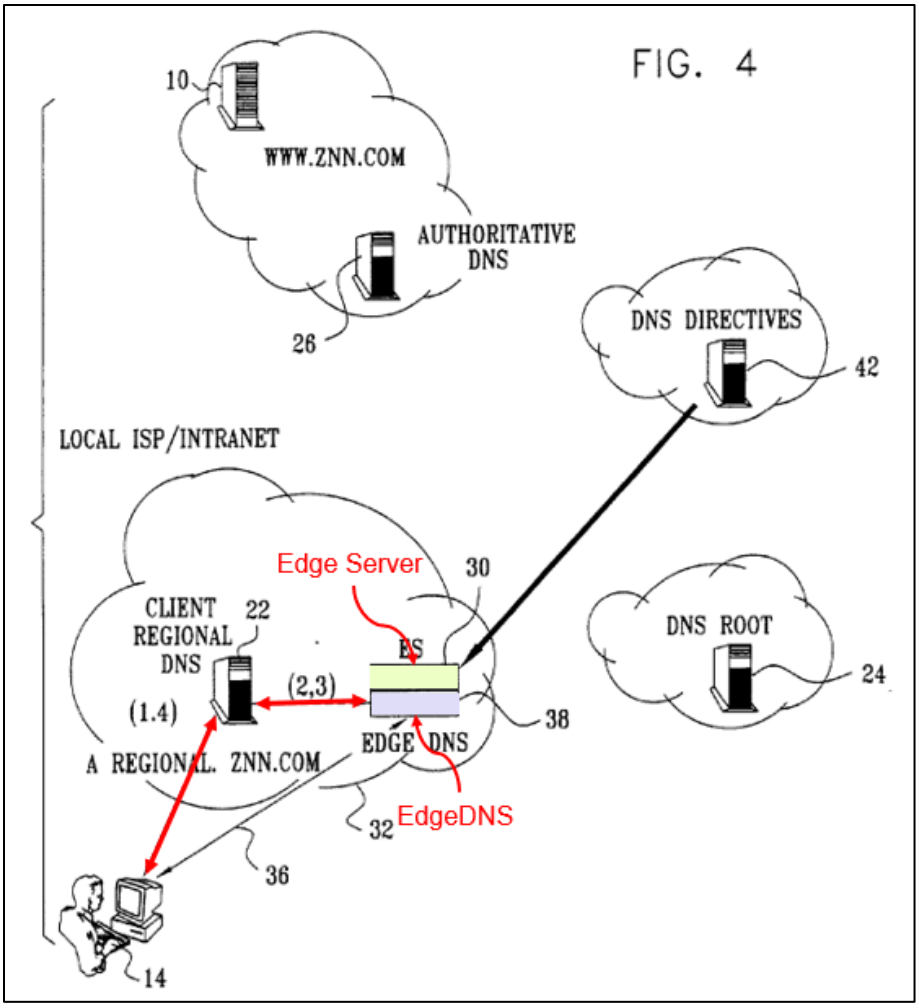

EX1005, FIG. 4. (annotated)

158. The EdgeDNS thus stores resource records that map IP addresses to domains. EX1005, ¶¶176, 373-374. It accesses these resource records to resolve DNS requests. *E.g.,* EX1005, ¶176. FIG. 15 depicts a representative method for responding to DNS requests by the EdgeDNS. When the EdgeDNS receives a DNS query for a domain, it first determines whether it is "authoritative" for that domain (i.e., whether it holds the actual resource records mapping the IP address for that domain). EX1005, ¶398. If so, the EdgeDNS retrieves and returns the IP address for that domain. If not, the EdgeDNS checks its cache for the requested resource

record, and either returns the IP address in the event of a cache hit or obtains the data from another DNS server in the event of a cache miss. EX1005, ¶¶398-399.

159. If the EdgeDNS's corresponding edge server is responsible for the requested domain, the EdgeDNS will have a resource record that maps the domain to the edge server IP address, and the EdgeDNS will return that IP address. If the edge server is not responsible, the EdgeDNS will return another server's IP address, such as the origin server.

160. Ben-Shaul's Listing 14 provides pseudo-code of a similar EdgeDNS technique for responding to DNS requests. EX1005, ¶362. As shown, upon receiving a DNS request, the EdgeDNS retrieves the relevant cache entries for the requested domain and returns either the edge server IP address (if the entry is a cache replacement (CLCR) entry) or the origin server IP address (if the entry is not a CLCR entry). EX1005, ¶362.

Listing 14

```
get cache entries relevant to the query type and domain
for each cached entry do
    if in bypass mode
      if entry type is CLCR
          ignore it
      else
          add entry
    if not in bypass mode
      if entry type is CLCR
          add entry and ignore all other types
      else if no CLCR entries found
          add entry
finally, if no entries found
    issue a normal DNS query to an external server
    after response arrives, send it back to the requester
    (while modifying the TTL value to max-ttl allowed)
```

EX1005, ¶362.

161. In Listing 14, the EdgeDNS distinguishes between "normal" DNS entries that map to the origin server and "canonical local cache replacement (CLCR)" DNS entries that map to the edge server. EX1005, ¶361. Upon receiving a DNS request, the EdgeDNS determines if the corresponding DNS entry in its cache is a CLCR record and responds accordingly with the appropriate IP address. EX1005, ¶¶361-362.

162. Thus, the EdgeDNS provides conventional DNS services supplemented using the dynamic cache replacement technique. Just as in conventional DNS servers, the EdgeDNS stores resource records mapping IP addresses to servers (e.g., edge servers or origin servers). In response to receiving a

DNS request, the EdgeDNS merely accesses these resource records to determine the IP address mapped to the requested domain. As with conventional DNS servers, if the EdgeDNS does not have a resource record for the requested domain, it queries another DNS server for this information.

163. Ben-Shaul teaches that the EdgeDNS may be modified relative to conventional DNS to implement the cache replacement techniques I have described above. EX1005, ¶¶358-360. But importantly, when resolving DNS requests, the EdgeDNS generally provides conventional DNS services by determining which servers are mapped to the requested domain as specified in its stored resource records. It does not apply any location-based policies when resolving these requests, nor does it consider any information regarding edge server groupings. The EdgeDNS simply accesses its stored resource records and returns the IP address specified in the appropriate record.

c. Coordination of the Edge Server and Edge DNS

164. Ben-Shaul's EdgeDNS and edge servers coordinate to determine the appropriate servers to which to direct traffic in response to DNS requests. For instance, an edge server can instruct its corresponding EdgeDNS to replace its stored DNS records so that certain domains are mapped to the edge server IP address instead of the origin server address. *E.g.,* EX1005, ¶¶360-361, 387-388, 414, 423.

165. The edge server does this by sending a "REPLACE" request to the EdgeDNS. *Id.* These messages are DNS messages adapted to use a non-standard REPLACE opcode that instructs the EdgeDNS to replace certain resource records. EX1005, ¶¶379, 386. For instance, these messages can include "Add/Replace" instructions, which "adds or replaces the local IP mapping of a domain name, thus hiding the original remote IP mapping." EX1005, ¶387, *see also* ¶414.

166. Upon receiving this instruction from the edge server, the EdgeDNS updates its resource records accordingly so that those domains map to the edge server IP address rather than the origin server address. EX1005, ¶¶361-362, 414. When a subsequent end-user DNS request for those domains is received, the EdgeDNS returns the edge server IP address, as specified in the modified resource records. EX1005, ¶362.

2. The cited prior art does not teach or suggest the claimed decision tree

167. Claim element [1.b] recites that "the GUI supports the setting of said at least one policy using a decision tree representing rules." Claim element [1.d] recites that "the decision tree comprises one or more resource nodes, and one or more branch nodes, wherein the one or more resource nodes specify one or more answers to be provided in response to a Domain Name Service (DNS) request, and wherein the one or more branch nodes specify one or more decision criteria to be applied." That is, the decision tree includes one or more resource nodes that

specify one or more answers to be provided in response to a DNS request and one or more branch nodes that specify one or more decision criteria to be applied. Petitioner's cited prior art does not teach a decision tree at all, let alone one with the required branch or resource nodes.

a. The cited prior art does not teach or suggest a decision tree that includes resource nodes specifying answers to be provided in response to a DNS request

168. Element [1.d] recites that "the decision tree comprises one or more resource nodes… [that] specify one or more answers to be provided in response to a Domain Name Service (DNS) request." The cited prior art does not teach or suggest this element.

169. The cited prior art does not teach a decision tree that includes *resource nodes that specify answers to be provided in response to a DNS request*, as recited in element [1.d]. Petitioner and Dr. Lin argue that three different GUIs teach or suggest the claimed decision tree: (1) Ben-Shaul's front edge tool, (2) the "Ben-Shaul single GUI," and (3) the combined Ben-Shaul-Hasan GUI. *See* Pet., 48, 50-52; *see also* EX1003, ¶¶162-168. But none of these GUIs teach a decision tree with resource nodes specifying answers to be provided in response to a DNS request.

170. Regarding the Ben-Shaul GUIs, Petitioner and Dr. Lin argue that the "hierarchical map" provided by these GUIs corresponds to the claimed decision

tree representing rules. *See* Pet., 48, 50 (arguing that the Ben-Shaul single GUI teaches a decision tree "for the same reasons the front edge tool does"); EX1003, ¶¶163, 165. Petitioner argues that the "leaves" of the hierarchical map specify answers "because an IP address for a selected server is provided as an answer to a DNS request." Pet., 53; *see also* EX1003, ¶175. But these leaves do not specify answers to DNS requests.

171. Ben-Shaul's GUIs, including the front edge tool and the Ben-Shaul single GUI, do not enable users to specify policies for responding to DNS requests. They are used to write edge server policies—they are not used to write EdgeDNS policies. *See* EX1005, ¶¶170, 211 (front edge tool is for "monitoring and managing edge servers"), 221 (front edge tool defines "edge server profiles"), 234, 350 ("CDML instruction files for the edge server"). In Ben-Shaul, edge servers are not domain name servers, as recited in claim 1, and they do not provide DNS services. *See* Section VII.C.1. Policies written for edge servers are not DNS policies because the edge servers do not implement DNS functionality. Rather, as I have explained, the EdgeDNS implements the DNS in Ben-Shaul.

172. In particular, Ben-Shaul defines the front edge tool as a "tool for defining service profiles." EX1005, ¶170. As I explained above, a service profile is "a collection of service and management directives that control the behavior of the *edge server*." EX1005, ¶220, *see also* ¶215 ("Content providers and possibly end-

users can control the behavior of a remote edge server by defining service profiles using a 'front edge' tool"). The front edge tool is also used "for purposes of monitoring and managing the *edge servers*." EX1005, ¶211. Nowhere does Ben-Shaul teach that the front edge tool is used to specify DNS policies implemented by the EdgeDNS.

173. Petitioner and Dr. Lin's proposed "Ben-Shaul single GUI" combines the front edge tool with Ben-Shaul's "CDML editor." Pet., 49; EX1003, ¶¶164-165. But the CDML editor is not used to specify policies for responding to DNS requests either. CDML stands for "content delivery markup language." EX1005, ¶37. It is an XML-based language used to author content delivery policies in Ben-Shaul's system. But these content delivery policies are not executed by the EdgeDNS, nor do they include policies for resolving DNS requests. CDML policies are delivered to and executed by edge servers. EX1005, ¶¶182, 188. As I explained above, Ben-Shaul provides example CDML directives in Listings 1-12. *See* EX1005, ¶¶193-208.

174. Petitioner and Dr. Lin assert that CDML policies include DNS redirection policies. Pet., 48-49 (citing Ben-Shaul, ¶39); EX1003, ¶164. But these policies are not executed by the EdgeDNS, nor are they for resolving DNS requests. At the outset, Ben-Shaul does not describe such redirection policies other than to mention their existence. *See* EX1005, ¶39. Regardless, as I have explained

above, the EdgeDNS does not execute CDML policies. Indeed, the EdgeDNS implements Berkeley Internet Name Daemon (BIND) software for responding to DNS requests, not CDML. *See* EX1005, ¶¶9-14, 170, 379-382. As I have explained, the EdgeDNS is a "conventional DNS Name Server" adapted to implement cache replacement. It does not execute CDML policies. Indeed, Ben-Shaul does not teach that it does.

175. Furthermore, I explained above the EdgeDNS's disclosed techniques for responding to DNS requests. These techniques are not implemented using policies implemented by a GUI. Thus, as neither the front edge tool nor Petitioner's proposed single GUI is used to write policies for responding to DNS requests, it is my opinion that neither of these GUIs specify answers to be provided in response to a DNS request as claimed.

176. As I have explained, the EdgeDNS resolves requests using conventional DNS adapted to implement a cache replacement technique. EX1005, ¶¶360-365. In particular, as explained above, Ben-Shaul discloses that the EdgeDNS responds to DNS requests by retrieving the resource records corresponding to the requested domain, and returning the IP address specified by the resource records. *See* EX1005, ¶¶176, 177, 362, 397-399, FIG. 15, Listing 14; *see also* Section VII.C.1 *supra*. Nowhere does Ben-Shaul disclose that any of these techniques for resolving DNS queries are specified in policy specified using a user

interface.

177. Petitioner and Dr. Lin argue that the hierarchical map's "edge server 'leaves' specify one or more answers to be provided in response to a [DNS] request' because an IP address for a selected server is provided as an answer to a DNS request." Pet., 53; *see also* EX1003, ¶175. They further argue that, in the example hierarchical map shown in Figure 10, "the IP address for the server 'Haifa' ('one or more answers') may be returned based on its location in Israel ('one or more decision criteria')." Pet., 53-54; EX1003, ¶175. I disagree.

178. Ben-Shaul does not consider edge server location or groups in response to receiving a DNS request. As explained, upon receiving a DNS request for a domain from an end-user, the EdgeDNS simply checks its stored resource records to identify the server (e.g., edge server or origin server) mapped to that domain. *See* Section VII.C.1 *supra*. If such a record exists, the EdgeDNS returns the IP address specified in that record. If not, the EdgeDNS queries another DNS server for the IP address. For the reasons I provided above, the EdgeDNS does not apply any policy that considers edge server grouping or location in responding to DNS requests. It does not apply any policy governing the selection of a server based on the location of that server or its inclusion in a particular group.

179. Petitioner and Dr. Lin cite Ben-Shaul paragraph 41 to support their assertion that location is considered in resolving DNS requests. *See* Pet., 53-54;

EX1003, ¶175. This paragraph merely states: "[t]he EdgeDNS is responsible to return the IP address of either the edge server or another server for resources that whose URL resides in these domains." EX1005, ¶41. It says nothing about edge server location or the hierarchical map, such as the one shown in Figure 10.

180. Petitioner and Dr. Lin also cite Ben-Shaul paragraph 44 to support their assertion that edge servers are selected based on their "association with the group" specified in the hierarchical map. Pet., 53; *see also* EX1003, ¶174. But this paragraph describes *edge server* functionality for providing "differentiate[d]" content to users in different regions or groups. It does not describe techniques for resolving DNS requests. For instance, this paragraph explains that edge servers may give different priorities to different resources, or that edge servers may make certain resources visible to end-users or organizations in different regions. *See* EX1005, ¶44, *see also* ¶197 (discussing "content-based differentiation" performed by the edge server). These are not EdgeDNS functions, and they do not describe techniques for resolving DNS requests. As I have explained, the edge servers do not implement DNS services, and policies for providing this differentiated content are thus not DNS policies as claimed.

181. Indeed, this regional and group differentiation refers to the "delivery" of content to users, which is an edge server function. The EdgeDNS does not deliver content to users, it only directs users to the servers that can deliver the

content. *See* EX1005, ¶¶41, 170. The EdgeDNS further does not know what content is stored at the edge server or what content is requested by the user. Again, its purpose is to resolve a request for a particular domain, not for particular content. In this differentiation technique, edge servers give different priorities to different resources. Again, this is not EdgeDNS functionality because the EdgeDNS is not concerned with *resources* stored at the edge server, it is only concerned with domains.

182. I agree that "an IP address for a selected server is provided as an answer to a DNS request." *See* Pet., 53; EX1003, ¶175. But the hierarchical map leaves do not specify or represent these answers. They do not govern which IP addresses are returned. They merely specify edge servers that are members of a group or in a particular location. The hierarchical map merely provides a "hierarchical view" of the edge servers that "allows the user to select and manipulate the hierarchies in an easy and intuitive manner." EX1005, ¶313. It is an "interface used for browsing the edge server map." EX1005, ¶326. It does not specify answers to DNS requests.

183. Indeed, the leaves of the hierarchical map merely represent edge servers. The front edge tool is used for "purposes of monitoring and managing the edge servers." EX1005, ¶211. The hierarchical map provides users with "an easy an intuitive manner" in which to do so. EX1005, ¶313. The edge servers in this

map thus represent edge servers for purposes of managing and configuring these edge servers. They do not specify answers to be provided in response to a DNS request. Indeed, these servers are not IP addresses or CNAMEs (answers) to be provided in response to DNS requests, as required by the claims. That the servers have IP addresses does not mean that the servers in the hierarchical map represent such answers.

184. Thus, neither of Ben-Shaul's GUIs (front edge tool or single GUI) teaches or suggests this element.

185. Hasan does not fix Ben-Shaul's deficiencies for similar reasons provided above for the Beaumont-Hasan combination. As I explained, Hasan's GUI merely provides "a hierarchical layout of [a] data center, its facilities, subscriber[s], devices and/or log servers." EX1006, 8:19-22. It is not a decision tree, it is merely "an organization of data objects" that is used to enable an administrator to "perceive[] the virtualization topology" of the data center." EX1006, 6:54-56, 8:19-22. Thus, for the reasons provided in Section VII.A regarding Hasan, Hasan does not teach a decision tree specifying answers to be provided in response to DNS requests.

186. And as Hasan is unconcerned with DNS, it doesn't cure Ben-Shaul's deficiencies regarding the DNS aspects of this element either. That is, Hasan merely provides a GUI that organizes data center objects in a hierarchy. It does not

disclose any DNS policies for responding DNS requests. The combination of Ben-Shaul and Hasan thus does not teach or suggest a decision tree that includes resource nodes specifying answers to be provided in response to a DNS request.

187. For instance, Petitioner argues, that in the combined GUI, "the policy setter could select an object, e.g., 'Jerusalem' and edit the policy applicable to that server (or group of servers)." Pet., 51; *see also* EX1003, ¶167. But this amounts to the same functionality of the Ben-Shaul GUIs. The combined GUI still does not enable specification of DNS policies and thus does not specify answers for responding to DNS requests. Neither reference teaches such functionality and their combination thus does not either.

b. The cited prior art does not teach or suggest a decision tree that includes branch nodes specifying decision criteria to be applied

188. Element [1.d] recites that "the decision tree comprises… one or more branch nodes… [that] specify one or more decision criteria to be applied. The cited prior art does not teach or suggest this element.

189. The cited prior art does not teach a decision tree that includes *branch nodes specifying decision criteria to be applied*, as recited in element [1.d]. Neither Ben-Shaul nor Hasan, alone or in combination, teaches this element. As explained, Petitioner and Dr. Lin point to the Ben-Shaul GUIs' "hierarchical map" and also to the Ben-Shaul-Hasan GUI as teaching the claimed decision tree. *See* Pet., 48, 50,

52; EX1003, ¶¶163-168. But none of these GUIs teach a decision tree with branch nodes specifying decision criteria to be applied.

190. Regarding Ben-Shaul, Petitioner and Dr. Lin argue that the "internal nodes" of Ben-Shaul's hierarchical map are branch nodes that specify decision criteria to be applied because "'they represent logical aggregations of nodes' based on a particular decision criteria," and because "policy application hinges on which of the groups is implicated." Pet., 53, 57; *see also* EX1003, ¶¶174, 185. Petitioner and Dr. Lin also allege that "[a] DNS request is thus directed to an edge server based on that server's association with the group." Pet., 53; *see also* EX1003, ¶174. I disagree for similar reasons as explained in Section VII.C.2.a *supra*.

191. As I explained, Ben-Shaul's GUIs, including the front edge tool and the Ben-Shaul single GUI, do not enable users to specify policies for providing answers in response to DNS requests. As I have also explained, the EdgeDNS does not consider edge server groupings or location when responding to DNS requests. The EdgeDNS does not apply any policy that considers edge server grouping or location in responding to DNS requests. The hierarchical map does not specify decision criteria to be applied in response to DNS requests.

192. Moreover, Ben-Shaul's hierarchical map does not include branch nodes specifying decision criteria to be applied in *any* policy. That users can write edge server directives that apply to groups of edge servers does not mean that the

hierarchical maps of servers themselves represent rules or decision criteria to be applied. Ben-Shaul states that "[a]dministrators of the system 80 can construct hierarchies of edge servers, *which are later used by profile editors in order to apply a particular service directive to a collection of edge servers*." *See* EX1005, ¶309. The directives are distinct from the groups themselves, and are applied to groups independently from, and subsequent to, the formation of the groups using the hierarchical map. The user can choose to apply directives to server groups, but the groups themselves do not specify any decision criteria. *See* EX1005, ¶309.

193. Indeed, the hierarchical map merely provides a "hierarchical view" of the edge servers that "allows the user to select and manipulate the hierarchies in an easy and intuitive manner." EX1005, ¶313. It is an "interface used for browsing the edge server map." EX1005, ¶326. These nodes in this map do not specify any decision criteria to be applied in a policy.

194. Thus, neither of Ben-Shaul's GUIs supports the setting of policies using a decision tree that includes branch nodes specifying decision criteria to be applied.

195. Hasan does not fix Ben-Shaul's deficiencies. As explained above, Hasan's GUI merely provides "a hierarchical layout of [a] data center, its facilities, subscriber[s], devices and/or log servers." EX1006, 8:19-22. Hasan refers to this hierarchical layout as a "view." EX1006, 8:19-22. This view is not a decision tree,

it is merely "an organization of data objects" that is used to enable an administrator to "perceive[] the virtualization topology" of the data center." EX1006, 6:54-56, 8:19-22. Thus, for the reasons provided in Section VII.A regarding Hasan, Hasan does not teach a decision tree as claimed.

196. And as Hasan is unconcerned with DNS, it doesn't fix Ben-Shaul's deficiencies regarding the DNS aspects of this element either. That is, Hasan merely provides a GUI that organizes data center objects in a hierarchy. It does not disclose any DNS policies for responding DNS requests. The combination of Ben-Shaul and Hasan thus does not teach or suggest GUI that supports the setting of policies using a decision tree that includes branch nodes specifying decision criteria to be applied.

197. For instance, Petitioner and Dr. Lin argue, that in the combined GUI, "the policy setter could select an object, e.g., 'Jerusalem' and edit the policy applicable to that server (or group of servers)." Pet., 51; *see also* EX1003, ¶167. But this amounts to the same functionality of the Ben-Shaul GUIs. The combined GUI still does not enable specification of DNS policies and thus does not specify answers for responding to DNS requests. Neither reference teaches such functionality and their combination thus does not either.

c. The cited prior art does not teach or suggest a decision tree at all

198. Element [1.b] recites that "the GUI supports the setting of said at least

one policy using a decision tree representing rules." The cited prior art does not teach or suggest this element.

199. Not only does the cited prior art fail to teach the decision tree having specific branch and resource nodes, it also fails to teach a decision tree at all (*see* element [1.b]). Neither Ben-Shaul nor Hasan, alone or in combination, teaches or suggests a decision tree representing rules. As explained, Petitioner and Dr. Lin point to the Ben-Shaul GUIs' "hierarchical map" and the Ben-Shaul-Hasan GUI as teaching the claimed decision tree. *See* Pet., 48, 50, 52; EX1003, ¶¶163-168. But none of these GUIs teach a decision tree that "represent[s] rules."

200. Petitioner and Dr. Lin argue that the hierarchical map provided by the Ben-Shaul GUIs, including the front edge tool and the Ben-Shaul single GUI, is a decision tree representing rules "because directives apply based on inclusion of edge servers within the groups." Pet., 48; *see also* EX1003, ¶163. I disagree. Ben-Shaul's hierarchical map merely provides a "hierarchical view" of the edge servers that "allows the user to select and manipulate the hierarchies in an easy and intuitive manner." EX1005, ¶313. It is an "interface used for browsing the edge server map." EX1005, ¶326. It does not represent rules.

201. That users can apply directives to edge server groups does not mean that the hierarchical map itself is a decision tree that represents rules. As explained in Section VII.B.2, the '517 patent describes the claimed decision tree representing

rules as a "policy hierarchy." EX1001, 8:56-59. Conversely, Petitioner and Dr. Lin admit that the hierarchical map merely represents "hierarchies of edge servers." Pet., 52; EX1003, ¶173. That this hierarchy of edge servers can be used "for effecting directives" does not make it a decision tree representing rules as claimed. *See* Pet., 52; *see also* EX1003, ¶173. Rather, it merely represents a topology of edge servers to which policies can be applied.

202. Ben-Shaul teaches that the edge server groups "simplify the application of policies applicable to a set of related edge servers." EX1005, ¶309. But as explained above, the edge server groups in the hierarchical map do not represent any rules. Ben-Shaul states that "[a]dministrators of the system 80 can construct hierarchies of edge servers, *which are later used by profile editors in order to apply a particular service directive to a collection of edge servers*." EX1005, ¶309 (emphasis added). The directives are distinct from the groups themselves, and are applied to groups independently from, and subsequent to, the formation of the groups using the hierarchical map. The user can choose to apply directives to server groups, but the groups themselves do not represent any rules. *See* EX1005, ¶309.

203. Hasan does not fix Ben-Shaul's deficiencies. It does not teach the claimed decision tree for the reasons described above in Section VII.B.2. Moreover, Petitioner's proposed combined Ben-Shaul-Hasan GUI is redundant to

Ben-Shaul's disclosed GUIs given that Ben-Shaul already discloses a hierarchy of servers.

204. According to Petitioner and Dr. Lin, the Ben-Shaul-Hasan combined GUI would merely enable a user to "select an object" such as a server or group of servers, "and edit the policy applicable to that server (or group of servers)." Pet., 51; *see also* EX1003, ¶167. But, as with the Hasan GUI, this merely amounts to a topological layout that enables users to input respective policies for servers. The Ben-Shaul-Hasan GUI does not represent a policy hierarchy; it is at best a hierarchy of content delivery servers. It merely arranges content delivery servers in a tree structure. The claimed GUI must do more than allow users to set policies for content delivery servers, it must support the specification of policy using a decision tree that itself represents rules. The Beaumont-Hasan GUI does not teach or suggest such a decision tree.

3. The cited prior art does not teach or suggest providing said at least one policy to the name server network

205. Element [1.c] recites "providing said at least one policy to the name server network." The cited prior art does not teach or suggest this element.

206. The cited prior art does not teach or suggest a policy that is obtained using a GUI and provided to a name server network as claimed. The independent claims require "using said GUI to obtain at least one policy for direction of network traffic" (element [1.b]) and "providing [the same] at least one policy to the

name server network" (element [1.c]). Ben-Shaul-Hasan does not teach these elements.

207. Petitioner and Dr. Lin identify only Ben-Shaul's "DNS directives" as the policies that are provided to the name server network. *See* Pet., 52 (citing Ben-Shaul, ¶177, which discloses "DNS directives"); *see also* EX1003, ¶170. But these DNS directives are not obtained using a GUI. Indeed, Ben-Shaul only mentions DNS directives twice, stating only that they are used for "configuration of the EdgeDNS 40 and its coordination with the client regional DNS server 22" and that they "originate from one or more policy control servers 42." EX1005, ¶177, *see also* ¶176. There is no evidence that these DNS directives are obtained or set using a GUI, as claimed.

208. Thus, Ben-Shaul-Hasan does not teach these elements.

4. Ben-Shaul does not disclose a domain name server adapted to provide adaptive policy-based domain name service

209. Element [1.pre.ii] recites "at least one domain name server comprising hardware in combination with software and constructed and adapted to provide adaptive policy-based domain name service." The cited prior art does not teach or suggest this element.

210. Petitioner fails to show that Ben-Shaul discloses a domain name server "adapted to provide adaptive policy-based domain name service," as recited in element [1.pre.ii]. Petitioner and Dr. Lin point to Ben-Shaul's EdgeDNS as

corresponding to the claimed domain name server. Pet., 44; EX1003, ¶154. Petitioner and Dr. Lin then allege that the EdgeDNS provides adaptive policy-based domain name service because "requests are handled dynamically based on network conditions," and cite only to Ben-Shaul paragraph 65 for support. Pet., 46; EX1003, ¶155. Paragraph 65 describes conditions for when an edge server may pre-fetch content from its corresponding origin server. It does not describe domain name service functionality, nor does it describe functionality of the EdgeDNS.

211. Edge server content pre-fetching is unrelated to the DNS functionality of the system. It is an interaction between the edge server and its origin server for determining when the edge server should download content from the origin server. Indeed, Ben-Shaul explicitly states that the "*edge server* may pre-fetch content." EX1005, ¶65, *see also* ¶¶60, 427, 482. It is not domain name service, nor is it performed by the EdgeDNS (alleged domain name server). The EdgeDNS is responsible for mapping domains to the appropriate server IP address, and for handling end-user DNS requests by "return[ing] the IP address of either the edge server or another server for resources that whose URL resides in these domains." EX1005, ¶41. It is not responsible for pre-fetching of content by the edge server, and even if it was, this would not be "adaptive policy-based domain name service."

212. Indeed, Ben-Shaul's EdgeDNS does not provide *adaptive* policy-based domain name service. It simply returns IP addresses that are specified in its

resource records, or that are provided to it by other DNS servers. *See* Section VII.C.1 *supra*. It does not apply adaptive policies in response to receiving DNS requests, nor does it respond to DNS requests (i.e., provide domain name service) based on "changing network conditions and/or changing user directives," as alleged by Petitioner. *See* Pet., 74; *see also* EX1003, ¶152. As explained above, it has fixed rules for responding to DNS requests and for configuring its internal resource records. It simply follows these rules to identify the appropriate IP address to return. It does not apply adaptive policies.

213. It is thus my opinion that Grounds 3A and 3B do not render obvious any challenged claim.

5. A POSA would not have combined Ben-Shaul and Hasan

214. A POSA would not have sought to combine Beaumont and Hasan at all, much less to arrive at the challenged claims.

215. As explained in Section VII.A, Hasan is not analogous art to the '517 patent because it is not directed to delivering content to end users over a network. For similar reasons, it is also not analogous to Ben-Shaul. Accordingly, a POSA would not have sought to combine Ben-Shaul and Hasan.

216. As explained in Section VII.C.1, Ben-Shaul is directed to a "content and application delivery system" that "allows content providers to directly control the delivery of content." EX1005, Abstract. The system includes an "origin site…

consisting of one or more origin Web servers" and "a set of subordinate edge servers which are geographically distributed and reside in the 'edges' of the internet." EX1005, ¶¶35-36. The system also includes an "edge DNS server" that performs DNS services in the system. EX1005, ¶38.

217. A POSA would not have sought to combine Ben-Shaul with Hasan. Hasan is not directed to DNS systems and does not address content delivery to users. Hasan also does not address the problem to which Ben-Shaul relates. In particular, Ben-Shaul addresses problems with conventional "content delivery" systems, such as problems related to the "centralized control" provided by such systems, their scalability issues, the lack of control these systems give to content providers, their inability to provide dynamic content, etc. *See* EX1005, ¶¶18-24.

218. Although Hasan teaches a GUI for editing policies, it does not teach a GUI for editing content delivery policies. As explained, Hasan discloses editing generic, undefined policies for data center administration. Hasan's management system does not provide content delivery or DNS policies. EX1006, 17:57-63. It leaves DNS to be "implemented by other systems." EX1006, 17:57-63. A POSA would not have looked to Hasan for such teachings because Hasan does not disclose a GUI for inputting content delivery policies.

219. To the extent a POSA wished to modify Ben-Shaul, one would not have looked to such a disparate reference as Hasan. A POSA would not have

sought to combine these references at all, much less to arrive at the claimed invention.

*****

220. For the above reasons, it is my opinion that Petitioner's Grounds 3A and 3B do not render obvious any challenged claim.

VIII. CONCLUSION

221. In signing this declaration, I recognize that the declaration will be filed as evidence in an *inter partes* review before the Patent Trial and Appeal Board of the United States Patent and Trademark Office. I also recognize that I may be subject to cross-examination in the case and that cross-examination will take place within the United States. If cross-examination is required, I will appear for cross-examination within the United States during the time allotted.

222. I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true, and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code.

Executed on this 8th day of February 2026, in Los Angeles, California.

_____

Prashant Shenoy

_____

Dr. Prashant Shenoy