

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

DELL TECHNOLOGIES INC. AND DELL INC.

Petitioners

v.

CLOUD BYTE LLC

Patent Owner

U.S. Patent No. 9,900,249

**DECLARATION OF MARK CROVELLA, PH.D. IN SUPPORT OF
PETITION FOR *INTER PARTES* REVIEW OF U.S. PATENT NO. 9,900,249**

TABLE OF CONTENTS

I.	Background and Qualifications	11
II.	Materials Considered.....	14
III.	Relevant Legal Standards	16
A.	Prior Art.....	17
B.	Person of Ordinary Skill in the Art	18
C.	Anticipation.....	19
D.	Obviousness	19
IV.	Technical Background.....	24
A.	Source Routing.....	24
B.	Autonomous Systems (ASs) and Inter-AS Routing.....	28
V.	Technical Analysis of the '249 Patent.....	32
A.	Overview of the Patent Specification.....	32
B.	Overview of the Prosecution History.....	39
C.	The Patent Claims	48
D.	Person of Ordinary Skill in the Art	62
E.	Claim Construction	64
VI.	Prior Art References	65
A.	Shimizu (EX1005)	65
B.	RFC-4271 (EX1006).....	73
VII.	The Challenged Claims Are Obvious in Light of Shimizu	74
A.	Independent Claim 1	75
1.	[1pre] A packet forwarding node, comprising:	75

TABLE OF CONTENTS

(continued)

2.	[1a] a storage unit including a non-transitory computer readable medium configured to store a packet forwarding table used for forwarding a packet based on an identifier which identifies a link between nodes on a packet forwarding path; and.....	77
3.	[1b] at least one processor configured to execute program instructions to provide a forwarding unit configured to:.....	89
4.	[1c] receive an incoming packet including a plurality of the identifiers;	91
5.	[1d] determine a direction to forward the incoming packet;	94
6.	[1e] determine whether to forward to an external network;	96
7.	[1f] add to path or link information when the incoming packet is not forwarded to the external network;	102
8.	[1g] remove path or link information header when the incoming packet is forwarded to the external network; and.....	104
9.	[1h] forward the incoming packet by using a path or link that corresponds to the identifier from among the plurality of the identifiers.	106
B.	Dependent Claims 2-5.....	110
1.	Claim 2: The packet forwarding node according to claim 1, wherein the storage unit stores the table used for forwarding the packet based on the identifier which uniquely identifies the link within a predetermined domain.	110
2.	Claim 3: The packet forwarding node according to claim 1, wherein the storage unit stores the table used for forwarding the packet based on the identifier which uniquely identifies the link within a predetermined local domain.	115
3.	Claim 4: The packet forwarding node according to claim [1], wherein the storage unit stores the table used for forwarding the packet based on the identifier which	

TABLE OF CONTENTS

(continued)

	uniquely identifies the link within the packet forwarding node.....	117
4.	Claim 5: The packet forwarding node according to claim 1, wherein the storage unit stores the table used for forwarding the packet based on the identifier which unique identifies the link among the packet forwarding nodes.....	118
C.	Independent Claim 6	119
1.	[6pre] A packet forwarding node, comprising:	119
2.	[6a] a storage unit including a non-transitory computer readable medium configured to store a packet forwarding table used for forwarding a packet based on an identifier which identifies an interface of a packet forwarding node; and.....	119
3.	[6b] at least one processor configured to execute program instructions to provide a forwarding unit configured to	120
4.	[6c] receive an incoming packet including a plurality of identifiers, ...	120
5.	[6d] determine a direction to forward the incoming packet, ...	120
6.	[6e] determine whether to forward to an external network, ...	120
7.	[6f] add to path or link information when the incoming packet is not forwarded to the external network,.....	121
8.	[6g] remove path or link information header when the incoming packet is forwarded to the external network, and.....	121
9.	[6h] to forward the incoming packet by using a path or link that corresponds to the identifier from among the plurality of identifiers.	121
D.	Dependent Claims 7-10.....	122

TABLE OF CONTENTS

(continued)

1.	Claim 7: The packet forwarding node according to claim 6, wherein the storage unit stores the table used for forwarding the packet based on the identifier which uniquely identifies the link within a predetermined domain.	122
2.	Claim 8: The packet forwarding node according to claim 6, wherein the storage unit stores the table used for forwarding the packet based on the identifier which uniquely identifies the link within a predetermined local domain.	122
3.	Claim 9: The packet forwarding node according to claim 6, wherein the storage unit stores the table used for forwarding the packet based on the identifier which uniquely identifies the link within the packet forwarding node.	123
4.	Claim 10: The packet forwarding node according to claim 6, wherein the storage unit stores the table used for forwarding the packet based on the identifier which uniquely identifies the link among the packet forwarding nodes.	123
E.	Independent Claim 11	123
1.	A packet forwarding node, comprising:	123
2.	a memory storing a program;.....	124
3.	at least one processor configured to execute the program to provide:	124
4.	a communication unit configured to communicate with a server for receiving a plurality of identifiers, each of which identifies a link between nodes on a packet forwarding path; and.....	125
5.	a forwarding unit configured to:	127
6.	receive an incoming packet including the plurality of the identifiers, ...	127
7.	determine a direction to forward the incoming packet, ...	128

TABLE OF CONTENTS

(continued)

8.	determine whether to forward to an external network,.....	128
9.	add to path or link information when the incoming packet is not forwarded to the external network,.....	128
10.	remove path or link information header when the incoming packet is forwarded to the external network, and.....	128
11.	forward the incoming packet by using a path or link that corresponds to an identifier from among the plurality of identifiers.	129
F.	Dependent Claims 12-15.....	129
1.	Claim 12: The packet forwarding node according to claim 11, wherein the forwarding unit forwards the packet based on the identifier which uniquely identifies the link within a predetermined domain.	129
2.	Claim 13: The packet forwarding node according to claim 11, wherein the forwarding unit forwards the packet based on the identifier which uniquely identifies the link within a predetermined local domain.....	130
3.	Claim 14: The packet forwarding node according to claim 11, wherein the forwarding unit forwards the packet based on the identifier which uniquely identifies the link within the packet forwarding node.....	130
4.	Claim 15: The packet forwarding node according to claim 11, wherein the forwarding unit forwards the packet based on the identifier which uniquely identifies the link among the packet forwarding nodes.....	131
G.	Independent Claim 16	132
1.	A packet forwarding node, comprising:	132
2.	a communication unit configured to communicate with a server for receiving a plurality of identifiers, each of which identifies an interface of the packet forwarding node; and	132

TABLE OF CONTENTS

(continued)

3.	at least one processor configured to execute program instructions to provide a forwarding unit configured to	132
4.	receive an incoming packet including a plurality of identifiers, ...	133
5.	determine a direction to forward the incoming packet, ...	133
6.	determine whether to forward to an external network, ...	133
7.	add to path or link information when the incoming packet is not forwarded to the external network, ...	133
8.	remove path or link information header when the incoming packet is forwarded to the external network, and	134
9.	forward the incoming packet by using a path or link that corresponds to an identifier from among the plurality of identifiers.	134
H.	Dependent Claims 17-20	134
1.	Claim 17: The packet forwarding node according to claim 16, wherein the forwarding unit forwards the packet based on the identifier which uniquely identifies the link within a predetermined domain.	134
2.	Claim 18: The packet forwarding node according to claim 16, wherein the forwarding unit forwards the packet based on the identifier which uniquely identifies the link within a predetermined local domain	135
3.	Claim 19: The packet forwarding node according to claim 16, wherein the forwarding unit forwards the packet based on the identifier which uniquely identifies the link within the packet forwarding node	135
4.	Claim 20: The packet forwarding node according to claim 16, wherein the forwarding unit forwards the packet based on the identifier which uniquely identifies the link among the packet forwarding nodes	136

TABLE OF CONTENTS

(continued)

I.	Independent Claim 21	136
1.	A server, comprising:.....	136
2.	a memory storing a computer program;.....	137
3.	at least one processor configured to execute the computer program to provide:	137
4.	a first unit configured to determine a plurality of identifiers, each of which identifies a link between nodes on a packet forwarding path; and.....	138
5.	a second unit configured send information elements to a packet forwarding node which is capable of	139
6.	receiving an incoming packet including a plurality of identifiers, ...	140
7.	determining a direction to forward. the incoming packet,.....	141
8.	determining whether to forward to an external network, ...	141
9.	adding to path or link information when the incoming packet is not forwarded to the external network,.....	142
10.	removing path or link information header when the incoming packet is forwarded to the external network, and.....	142
11.	forwarding the incoming packet by using a path or link, that corresponds to an identifier from among the plurality of identifiers.	143
J.	Dependent Claims 22-27	143
1.	Claim 22: The packet forwarding node according to claim 1, wherein the storage unit stores the table used for forwarding the packet based on the identifier which identifies the link between adjacent nodes on a packet forwarding path.....	143

TABLE OF CONTENTS
(continued)

2.	Claim 23: The packet forwarding node according to claim 6, wherein the storage unit stores the packet forwarding table used for forwarding the packet based on the identifier which identifies the link between adjacent nodes on a packet forwarding path.	144
3.	Claim 24: The packet forwarding node according to claim 11, further comprising a storage unit that stores a table used for forwarding the packet based on an identifier which identifies the link between adjacent nodes on a packet forwarding path.	145
4.	Claim 25: The packet forwarding node according to claim 16, further comprising a storage unit that stores a table used for forwarding the packet based on an identifier which identifies the link between adjacent nodes on a packet forwarding path.	145
5.	Claim 26: The server according to claim 21,.....	146
6.	Claim 27: The packet forwarding node according to claim 1, wherein when the forwarding unit receives the incoming packet including the identifier for the link connecting nodes on the packet forwarding path, the forwarding unit forwards the incoming packet based on the identifier included the received incoming packet from among the plurality of the identifiers.	147
VIII.	The Challenged Claims Are Rendered Obvious by the Combination of Shimizu and RFC-4271	148
A.	Motivation to Combine Shimizu with RFC-4271.....	149
B.	Claim Limitation: [determine/determining] whether to forward to an external network;	151
IX.	Secondary Considerations	152
X.	Availability for Cross-Examination	153

TABLE OF CONTENTS
(continued)

XI.	Right to Supplement	153
XII.	Conclusion	153

I, Mark Crovella, Ph.D., declare as follows:

1. I make this declaration based upon my own personal knowledge and, if called upon to testify, would testify competently to the matters stated herein.

2. I have been asked by the petitioners to provide technical assistance in connection with the *inter partes* review of U.S. Patent No. 9,900,249 (which I will refer to as the “’249 Patent”). This declaration is a statement of my opinions on issues related to the patentability of claims 1-27 (which I will refer to as the “challenged claims”) of the ’249 Patent.

3. My compensation is not based on the content of my opinions or the resolution of this matter. I have no financial interest in the outcome of this proceeding.

I. BACKGROUND AND QUALIFICATIONS

4. In formulating my opinions expressed in this declaration, I have relied on my knowledge, training, and experience in the relevant field, which I will summarize briefly. Additional details are provided in my curriculum vitae, which I understand has been marked as Exhibit 1004.

5. I am a Professor in the Department of Computer Science at Boston University. From 2013 to 2018, I served as Chair of the Department. Since 2020, I also have been a Professor in the Faculty of Computing and Data Sciences, and I serve as the Chair of Academic Affairs there.

6. I earned an undergraduate degree in Biology from Cornell University, which I received in 1982. I received a master's degree in Computer Science from the University of Buffalo in 1989. I received a Ph.D. in Computer Science from the University of Rochester in 1994. The subject of my Ph.D. thesis was "Performance Predication and Tuning of Parallel Programs."

7. From 1982 to 1984, I worked as a computer programmer for the State of Colorado. From 1984 to 1994, I was employed at Calspan Corporation, a research and development firm in Buffalo, NY, where I rose to the level of Senior Computer Scientist. My work at Calspan focused on development of experimental software and large-scale simulation software in support of contracts between Calspan and the U.S. Department of Defense.

8. In 1994, I joined the faculty of Boston University, initially as an Assistant Professor of Computer Science. I was promoted to the rank of Associate Professor in 2000 and became a full Professor in 2006. I served as Chair of the Department from 2013 to 2018.

9. I have been pursuing research in the area of computer networking since 1994. I have conducted research in a variety of areas related to the Internet and the World Wide Web. Among other areas, I have studied the efficient design of Web servers and content distribution systems; I have studied the statistical properties of Internet traffic; and I have made extensive measurements of Internet

infrastructure and the behavior of Internet protocols.

10. For many years, I have taught a number of courses that have covered topics relevant to the '249 Patent, including Computer Networking. In that class, I teach the students about network routing and switching technology (past and present) and other technologies relevant to the challenged patent.

11. From 2007 to 2009, I was the Chair of ACM SIGCOMM (the Special Interest Group in Computer Communication), the main professional organization for scientists in the field of computer networking. I have been elected a Fellow of the IEEE and also of the ACM.

12. I am co-author of Internet Measurement: Infrastructure, Traffic, and Applications (Wiley Press, 2006), which is the first book written on the subject. I am the author of over two hundred research papers; according to Google Scholar, my work has been cited by over 30,000 other publications. Also, according to Google Scholar, as of May 2018, only 11 other researchers worldwide were more cited in the area of computer systems. I am a past editor of principal journals in the field of networking: Computer Communication Review, IEEE/ACM Transactions on Networking, Computer Networks, and IEEE Transactions on Computers.

13. I hold nine U.S. Patents derived from my Internet related research.

14. I also have considerable experience working with Web protocols and

distributed systems in industry. I co-founded a company that developed protocols for content delivery in the Web, which through a chain of acquisitions became part of Network Appliance, Inc. At Network Appliance, I served for two years as a Technical Director working on content distribution and Web caching systems.

15. I have also served as Chief Scientist for Guavus, Inc., a company founded by one of my Ph.D. students, that commercialized work that we did analyzing the properties of Internet traffic. Guavus was subsequently acquired and is now a division of Thales, Inc.

16. I am very familiar with the subject matter of the challenged patent. Based on my substantial academic and professional experiences, I consider myself an expert in the technology fields pertaining to the '249 Patent.

II. MATERIALS CONSIDERED

17. In preparation of this declaration, I have considered the materials discussed in this declaration, including for example the '249 Patent and its prosecution history. In particular, I have reviewed and considered the following documents, among others, in connection with my analysis:

- EX1001: The '249 Patent;
- EX1002: The prosecution history of U.S. Patent Application No. 14,244,701, which led to the issuance of the '249 Patent;
- EX1005: U.S. Patent Application Publication No. 2006/0227779

(“Shimizu”);

- EX1006: Request for Comments 4271, “A Border Gateway Protocol 4 (BGP-4)” (“RFC-4271”);
- EX1007: Request for Comments 791, “Internet Protocol”;
- EX1008: Request for Comments 2026, “The Internet Standards Process – Revision 3”;
- EX1009: RFC Editor – Publication Process;
- EX1010: Kurose and Ross, “Computer Networking: A Top-Down Approach” (4th ed. 2008);
- EX1011: Black, “TCP/IP and Related Protocols” (McGraw Hill, 1992);
- EX1012: U.S. Patent No. 6,956,854;
- EX1013: U.S. Patent No. 8,724,638;
- EX1014: U.S. Patent No. 7,389,359;
- EX1015: U.S. Patent No. 7,660,259;
- EX1016: U.S. Patent No. 7,411,948;
- EX1017: U.S. Patent No. 5,274,631;
- EX1018: U.S. Patent No. 7,417,990;
- EX1019: U.S. Patent No. 7,693,158;
- EX1020: U.S. Patent No. 5,889,777;

- EX1021: U.S. Patent No. 6,411,506;
- EX1022: U.S. Patent No. 6,950,895;
- EX1023: McKeown, “Fast Switched Backplane for a Gigabit Switched Router,” Business Communications Review (Dec. 1997);
- EX1024: Partridge, et al, “A 50-Gb/s IP Router,” IEEE/ACM Transactions on networking (June 1998);
- EX1025: Complaint for Patent Infringement;
- EX1026: Cloud Byte’s Preliminary Infringement Contentions with attached Exhibit 10; and
- EX1030: Internet Archive Declaration of Mina Ching for RFC-4271.
- EX1033: Cloud Byte’s Disclosure of Proposed Claim Constructions and Extrinsic Evidence with attached Exhibit A

18. My analysis and opinions are further based on my education, training, experience, and knowledge in the relevant field.

III. RELEVANT LEGAL STANDARDS

19. I am not an attorney. For the purposes of this declaration, I have been informed by counsel for the petitioners about certain aspects of the law that are relevant to forming my opinions. My understanding of the relevant law is as follows.

A. Prior Art

20. I understand that for an invention claimed in a patent to be found patentable, it must be (among other things) new and not obvious based on what was known before the invention was made.

21. I understand that the information used to evaluate whether a claimed invention was new and not obvious when made is generally referred to as “prior art.” I understand that prior art includes all patents and printed publications that existed before the earliest filing date of the patent. This includes foreign language material. I also understand that a patent is prior art if it was filed before the effective filing date of the claimed invention and that a printed publication is prior art if it was publicly available before the effective filing date.

22. I understand that in this *inter partes* review proceeding, the petitioners have the burden of proving that the challenged claims are unpatentable in light of the prior art by a preponderance of the evidence. I understand that a preponderance of the evidence is evidence sufficient to show that a fact is more likely true than not true.

23. I understand that patent claims in an *inter partes* review are interpreted by applying the same standard that applies in district court litigation. After the claims are construed, they are then compared to the prior art.

24. I understand that in this *inter partes* review proceeding, the prior art

that may be used to evaluate whether the challenged claims are new and not obvious is limited to patents and printed publications. My analysis, which is set out in detail below, compares the challenged claims to patents and printed publications that I have been informed constitute prior art to the claims.

B. Person of Ordinary Skill in the Art

25. I understand that my assessment of the challenged claims and the teachings of the prior art as well as my analysis and opinions herein must be undertaken from the perspective of what would have been known or understood by a person having ordinary skill in the art, reading the challenged patent on its priority date and in light of the specification and file history of the patent. I may refer to such a person as a “POSITA.”

26. I further understand that in determining the level of ordinary skill in the art, I may consider factors including:

- (a) the type of problems encountered in the art or field of invention,
- (b) prior art solutions to those problems,
- (c) the rapidity with which innovations are made,
- (d) sophistication of the technology, and
- (e) the educational level of active workers in the field.

27. I understand that a person of ordinary skill in the art, or POSITA, is not a specific real individual but is instead a hypothetical individual having the

qualities reflected by the factors above. This hypothetical person has knowledge of all prior art in the relevant field and takes from each reference what it would teach to a person having the skills of a POSITA.

C. Anticipation

28. I understand that a patent claim may be “anticipated” if each element of that claim is present either explicitly or inherently in a single prior art reference, and that the elements should be arranged in the prior art reference as in the claim. I understand that for a claimed limitation to be inherently present, the prior art reference need not expressly disclose the limitation so long as the claimed limitation necessarily flows from a disclosure in the reference.

D. Obviousness

29. I understand that prior art can render a patent claim unpatentable where subject matter falling within the scope of the claim would have been “obvious” to a person of ordinary skill in the art. I understand that the following standards govern the determination of whether a patent claim is rendered obvious in light of the prior art. I have applied these standards in my evaluation of whether the challenged claims are obvious in light of the prior art.

30. I understand that a claimed invention is not patentable if it would have been obvious to a person of ordinary skill in the field of the invention at the time the claimed invention was made. Even if all the elements of a claim were not

found in a single prior art reference, the claim is not patentable if the differences between the subject matter in the prior art and the subject matter in the claim would have been obvious to a POSITA at the time the application was filed. Prior art disclosing a method or device that falls within the scope of a claim can render that claim obvious even if other, different methods or devices might also fall within the scope of the claim.

31. I understand that a determination as to whether a claim would have been obvious should be based on four factors (though not necessarily in the following order): (i) the level of ordinary skill in the art at the time the application was filed; (ii) the scope and content of the prior art; (iii) the differences between the claimed invention and the prior art; and (iv) any “objective factors” indicating obviousness or non-obviousness that may exist in a particular case.

32. I understand that an obviousness analysis should not be based on hindsight but must be performed using the perspective of a person of ordinary skill in the relevant art as of the effective filing date of the patent claim.

33. I understand that secondary indicia of non-obviousness may include: (1) a long felt but unmet need in the prior art that was satisfied by the invention of the patent; (2) commercial success of processes covered by the patent; (3) unexpected results achieved by the invention; (4) praise of the invention by others skilled in the art; (5) taking of licenses under the patent by others;

(6) deliberate copying of the invention; (7) failure of others to find a solution to the long felt need; and (8) skepticism by experts. I also understand that there must be a relationship between any such secondary considerations and the claimed invention. I further understand that contemporaneous and independent invention by others is a secondary consideration supporting an obviousness determination.

34. I understand that the teachings of two or more prior art references may be combined in the manner disclosed in a challenged claim if such a combination would have been obvious to one having ordinary skill in the art. In determining whether a combination would have been obvious, the following exemplary rationales may support a conclusion of obviousness:

- combining prior art elements according to known methods to yield predictable results;
- simple substitution of one known element for another to obtain predictable results;
- use of a known technique to improve similar devices (methods, or products) in the same way;
- applying a known technique to a known device (method, or product) ready for improvement to yield predictable results;
- “obvious to try” – choosing from a finite number of identified, predictable solutions, with a reasonable expectation of success;

- known work in one field of endeavor may prompt variations of it for use in either the same field or a different one based on design incentives or other market forces if the variations are predictable to one of ordinary skill in the art;
- some teaching, suggestion, or motivation in the prior art that would have led one of ordinary skill to modify the prior art reference or to combine prior art reference teachings to arrive at the claimed invention; and
- common sense.

35. I understand that the obviousness analysis need not seek out precise teachings directed to the specific subject matter of the challenged claim but instead can take account of the ordinary innovation and experimentation in the relevant field that does no more than yield predictable results.

36. I understand that, in assessing whether there was an apparent reason to modify or combine known elements as claimed, it may be necessary to look to interrelated teachings of multiple patents and printed publications, the effects of commercial demands, and the background knowledge of a person of ordinary skill in the art. I further understand that any motivation that would have applied to a person of ordinary skill in the art, including motivation from common sense or derived from the problem to be solved, is sufficient to explain why references

would have been combined.

37. I understand that modifications and combinations suggested by common sense are important and should be considered. Common sense suggests that familiar items can have obvious uses beyond the particular application being described in a prior art reference, that if something can be done once it would be obvious to do it multiple times, and that in many cases a person of ordinary skill in the art can fit the teachings of multiple patents together in an obvious manner to address a particular problem. I also understand that the prior art does not need to be directed to solving the same problem that is addressed in the challenged patent claims.

38. I understand that a person of ordinary skill in the art is also a person of ordinary creativity. In many fields, it may be that there is little discussion of obvious techniques, modifications, and combinations, and it may be the case that market demand, rather than scientific research or literature, will drive a new design. When there is market pressure or design need to solve a particular problem and there are a finite number of identified, predictable solutions, a person of ordinary skill in the art has a good reason to employ the known options. If this leads to the expected success, then it is likely the product of ordinary skill and common sense as opposed to patentable innovation. I understand that if a combination was obvious to try, that may show that it was obvious and therefore

unpatentable. That a particular combination of prior art elements was obvious to try suggests that the combination was obvious even if no one made the combination.

IV. TECHNICAL BACKGROUND

39. The '249 Patent is generally related to routing data packets through network devices. The following information would have been known to a person of ordinary skill in the art prior to the '249 Patent and is useful for understanding the patent and prior art.

A. Source Routing

40. By 2009, the Internet relied on the Internet Protocol (IP) for addressing and forwarding data packets. Request for Comments 791, entitled "Internet Protocol," set out the technical standard for IP routing that was commonly in use by the time the '249 Patent was filed. EX1007.

41. A Request for Comments, or RFC, is essentially a technical standard or protocol used by networking professionals. RFCs are prepared and distributed under a formalized publication process overseen by the Internet standards governing body called the Internet Engineering Task Force (IETF). EX1008. I have personal experience working with and teaching about RFCs, including the IP standards and many others. For example, I discuss dozens of RFCs in my book *Internet Measurement*. I understand that once RFCs are published, the IETF does

not make changes to any published RFC. The “Publication Process” webpage of the RFC Editor, which publishes RFCs on behalf of the IETF, states that “We DO NOT make changes to RFCs once they have been published.” EX1009, 4. This is consistent with my personal experience with RFCs. Thus, it is my understanding, and I believe, that RFCs available today through the IETF website or the RFC Editor website are substantively identical to the originally published versions. Further, a person of ordinary skill in the art in the 2000s would have been familiar with RFCs in general, and the IP standards in particular, and such a person would have known how to obtain copies of RFCs from the IETF website or the RFC Editor website.

42. The IP protocol applies to the network-layer of the Internet protocol stack. A network-layer packet is commonly referred to as a datagram. Each datagram includes a header and data payload. The header includes information used for routing the packet from its source to its destination.

43. The most common method for forwarding an IP datagram is referred to as “destination routing.” In this approach, when a switch in the network receives the packet, the switch uses the destination address to identify the appropriate port on the switch on which to forward the packet. The switch maintains a forwarding table that is used to associate each particular destination with a port. The associated port is used to forward the packet in the direction of its

destination. Each such forwarding decision along the path is referred to as a “hop.” Destination routing is flexible and robust, but because each router only selects the next outgoing link on which to forward the packet no single router controls the particular path a packet will take through a network.

44. The IP protocol specifies another technique referred to as “source routing” in which the particular path a data packet will take through the network is specified in the packet’s header. The following figure from “TCP/IP and Related Protocols” illustrates a packet header including route data through several IP addresses.




Figure 5.5 Source Routing.

EX1011, 28.

45. RFC-791 describes two different types of source routing: loose

source routing and strict source routing. EX1007, 22. RFC-791 describes the “loose source and record route (LSSR) option,” which “provides a means for the source of an internet datagram to supply routing information to be used by the gateways in forwarding the datagram to the destination, and to record the route information.” EX1007, 24. The route data is “a series of internet addresses” in the header of the datagram specifying certain specific nodes the packet must traverse. EX1007, 24-25. “This option is a loose source route because the gateway or host IP is allowed to use any route of any number of other intermediate gateways to reach the next address in the route.” EX1007, 25. Thus, loose source routing specifies certain nodes along the path but not every hop. EX1011, 29 (“[L]oose source routing, gives the IP modules the option of using intermediate hops to reach the addresses obtained in the source list as long as the datagram traverses the nodes listed.”).

46. Strict source routing, on the other hand, specifies every hop along a path. “This option is a strict source route because the gateway or host IP must send the datagram directly to the next address in the source route through only the directly connected network indicated in the next address to reach the next gateway or host specified in the route.” EX1007, 26; EX1011, 29 (“[S]trict source routing requires that the datagram travel only through the networks whose addresses are indicated in the source list.”).

B. Autonomous Systems (ASs) and Inter-AS Routing


47. The '249 Patent describes a routing technique used in a private network that is distinct from “external” networks. A person of ordinary skill in the art would have understood one embodiment of the private network to be an “autonomous system” or AS.

48. An AS is a group of network nodes, such as switches and routers, grouped together and “typically under the same administrative control (e.g., operated by the same ISP or belonging to the same company network).” EX1010, 167; see also EX1006, 4 (“The classic definition of an Autonomous System is a set of routers under a single technical administration, using an interior gateway protocol (IGP) and common metrics to determine how to route packets within the AS, and using an inter-AS routing protocol to determine how to route packets to other ASes.”).

49. Network nodes within the same AS run the same routing algorithm and have information about each other. Different ASs may, and commonly do, have different routing algorithms. “It will be necessary, of course, to connect ASs to each other, and thus one or more of the routers in an AS will have the added task of being responsible for forwarding packets to destinations outside the AS; these routers are called gateway routers.” EX1010, 167.

50. The following figure from “Computer Networking” provides a

simplified example of three ASs: AS1, AS2, and AS3.



EX1010, 168. In the figure, “AS1 has four routers—1a, 1b, 1c, and 1d—which run the intra-AS routing protocol used within AS1” so “each of these four routers knows how to forward packets along the optimal path to any destination within AS1.” EX1010, 167. As illustrated, AS2 and AS3 each have three routers. “Note that the intra-AS routing protocols running in AS1, AS2, and AS3 need not be the same” and that “routers 1b, 1c, 2a, and 3a are all gateway routers.” EX1010, 167.


51. “The routing algorithm running within an autonomous system is called an intra-autonomous system routing protocol.” EX1010, 167. (This is also called an “Interior Gateway Protocol.” EX1006, 5). In the example shown, router 1b uses the intra-autonomous system routing protocol applicable to AS1 when forwarding to another node in AS1 (routers 1a, 1c, 1d). But when router 1b forwards packets to router 2a in AS2, it must use an “inter-AS routing protocol.”

EX1010, 169. The inter-AS routing protocol allows a gateway router in one AS, using one intra-AS routing protocol, to communicate with another AS that may be using a different intra-AS routing protocol.

52. To illustrate the point, consider the example of an organization such as a commercial enterprise like IBM arranging to obtain Internet service from an Internet service provider like Verizon. The network run by IBM (one AS) and the network run by Verizon (another AS) may have different protocols and policies in place, which would prevent connecting the networks together unless an inter-AS routing protocol is used to form a bridge between them.

53. As explained in “Computer Networking”: “Since the inter-AS routing protocol involves communication between two ASs, the two communicating ASs must run the same inter-AS routing protocol. In fact, in the Internet all ASs run the same inter-AS routing protocol, called BGP4....” EX1010, 169.

54. The following figure from “Computer Networking” illustrates the same three ASs employing the BGP4 protocol.



EX1010, 179. The connection and messages sent between nodes are called BGP sessions. “[A] BGP session that spans two ASs is called an external BGP (eBGP) session, and a BGP session between routers in the same AS is called an internal BGP (iBGP) session.” EX1010, 179-180.

55. BGP4 provides a mechanism allowing connection between two networks managed by independent organizations (i.e., separate ASs). In my example, the IBM network would establish an external BGP connection (eBGP) with the Verizon network. Creating this connection would cause the switches inside the IBM AS to be configured with knowledge of the eBGP link between the IBM AS and the Verizon AS. The switches inside the Verizon AS would likewise be configured with knowledge of the eBGP link. The routing protocols in the IBM network and in the Verizon network would thereby be equipped to forward packets to the external link to the other network when appropriate.

V. TECHNICAL ANALYSIS OF THE '249 PATENT

A. Overview of the Patent Specification


56. The '249 Patent is entitled "Communication System, Forwarding Node, Path Management Server, Communication Method, and Program."

EX1001, 1. The patent issued based on an application filed on April 3, 2014, which was a continuation of an earlier U.S. application filed on July 5, 2011, which itself was a continuation of an international application filed on September 13, 2010. EX1001, 1-2. The patent claims priority to Japanese Patent Application No. 2009-212221 filed on September 14, 2009. EX1001, 1.

57. The '249 Patent discloses "[a] packet forwarding node [that] includes a storage unit to store a packet forwarding table used for forwarding a packet based on an identifier which identifies a link between nodes on a packet forwarding path and a forwarding unit to receive an incoming packet including a plurality of the identifiers, and to forward the incoming packet based on an identifier corresponding to the packet forwarding node." EX1001, 1 (Abstract).

58. Figure 20 of the '249 Patent illustrates a prior art "network configuration using IP (Internet Protocol)." EX1001, 1:30-31.

FIG. 20



EX1001, Fig. 20. In such a network, “[w]hen a forwarding node 200 receives an IP packet transmitted by a communication node 100, it determines a forwarding destination of the IP packet” and forwards the packet toward the destination.


EX1001, 1:34-37. “In determining the forwarding destination of the IP packet, the forwarding node 200 uses a routing table held internally.” EX1001, 1:40-42. “The routing table is a table indicating which network destination packet should be transmitted, th[r]ough which interface, to a forwarding node that takes on the next forwarding processing....” EX1001, 1:42-47. The “[f]orwarding nodes repeat this, and ultimately forward the IP packet to a communication node that is the destination.” EX1001, 1:37-39.

59. The '249 Patent asserts that there is a problem with IP networks because “packet forwarding depends on the routing table of the respective forwarding nodes, and paths cannot be completely controlled.” EX1001, 2:1-5; see

also 2:5-9 (“there is a problem in that exact path control is not possible”).

60. Figure 1 is a “diagram showing a communication system according to the first exemplary embodiment of the present invention.” EX1001, 6:62-64.

FIG. 1




EX1001, Fig. 1. The illustrated communication system includes “communication nodes 100a and 100b, boundary forwarding nodes 300a and 300b, internal forwarding nodes 400, and a path management server 500 as shown.” EX1001, 6:64-67. The figure shows “data transfer network 600,” which “is a network for performing forwarding processing of a packet by a method of the present invention, and an external network 700,” which “is a network for performing packet forwarding processing by a method different from the network 600” (e.g., an IP network). EX1001, 7:1-5; see also 7:9-10 (“Here, a description is given with the external network 700 as an IP network.”). “The external network 700 is

connected to the data transfer network 600 via the boundary forwarding nodes 300a and 300b.” EX1001, 7:10-13; see also 7:23-29 (“The boundary forwarding nodes 300a and 300b are disposed between the data transfer network 600 and the external network 700....”).

61. When the boundary forwarding nodes 300a and 300b receive a packet from the communication nodes 100a or 100b, the boundary forwarding nodes 300a and 300b “assign a header including forwarding path information...and in addition, forward the packet in question to an internal forwarding 400 in the data transfer network 600 based on the forwarding path information in the header.” EX1001, 7:23-31. Once the packet reaches the edge of data transfer network 600, “the boundary forwarding nodes 300a and 300b remove the header from the received packet, and thereafter transmit the packet to communication nodes 100a and 100b of the external network 700.” EX1001, 7:32-39.

62. Figure 4 depicts the “format for a path information header used in the present exemplary embodiment.” EX1001, 9:20-21.

FIG. 4



EX1001, Fig. 4. In the figure, “‘Direction’ indicates a forwarding direction. For example, ‘1’ indicates a regular (forward) direction with respect to forwarding, and ‘2’ indicates a reverse direction with respect to forwarding.” EX1001, 9:21-25. “In ‘Current Offset’, offset information is set for local ID information to be referred to when forwarding....” EX1001, 9:25-27. “In ‘Local ID #n’, the local ID, which is to be referred to as a destination by the boundary forwarding node 300 of the nth hop or the internal forwarding node 400, is set.” EX1001, 9:38-41.


63. When “the ‘Direction’ field of the path information header of a received packet is ‘1’, that is, when a forwarding in a regular direction is indicated,” a packet forwarding unit uses “the ‘Current Offset’ of the path information header” to identify “a corresponding local ID” and determines the “communication interface 310” for that destination. EX1001, 9:58-67. The offset is increased to indicate the next local ID and the packet “is transmitted from the

determined communication interface 310.” EX1001, 9:65-67. In other words, the sequence of local IDs in the path information header dictates the sequence of forwarding nodes the packet will traverse through data transfer network 600.

64. “On the other hand, in a case where the ‘Direction’ is ‘2’, that is, when forwarding a reverse direction is indicated, reference is made to the ‘Current Offset’ of the path information header” to identify the “local ID that is one ahead of the corresponding local ID” and “the communication interface 310 that is the destination is determined.” EX1001, 9:67-10:6. The “Current Offset” amount is reduced and the “packet is transmitted from the determined communication interface 310.” EX1001, 9:67-68. Again, the sequence of local IDs in the path information header dictates the sequence of forwarding nodes the packet will traverse through data transfer network 600. This is simply “source routing” (which I discussed above) through a private network.

65. The forwarding path information (in Figure 4) can be “acquired from the path management server 500” or “may be generated from information held within a boundary [of] forwarding node 300.” EX1001, 7:40-45.

66. Figure 10 is a “flowchart representing details of forwarding processing....” EX1001, 15:11-12. “[F]irst the packet forwarding unit 320 checks the ‘Direction’ field of the path information header and decides whether to forward in a regular (forward) direction or to forward in a reverse direction (step S200).”

FIG. 10

EX1001, Fig. 10. If “the ‘Direction’ field indicates forwarding in a regular direction” (S200=Y), “a decision is made as to whether or not the destination of the packet in question is the external network 700 (step S201).” EX1001, 15:18-22. If the packet is not being forwarded to the external network (S201=N), the “local ID...is held” and “the number of bytes of 1 hop are added to the ‘Current Offset’ of the path information....” EX1001, 15:31-37. “The packet forwarding unit 320 uses the local ID that is held and performs forwarding processing of the packet in a regular direction (step S206).” EX1001, 15:38-40. The local ID is used to

determine the communication interface 310 “using the information of the forwarding table...and the packet is forwarded from the communication interface 310 in question.” EX1001, 15:40-44.

67. When “the value of the ‘Direction’ field indicates forwarding in a reverse direction...a decision is made as to whether or not the forwarding destination of the packet in question is the external network 700 (step S203).” EX1001, 15:45-49. If the packet is not being forwarded to the external network (S203=N), “1 hop amount is subtracted from the value of ‘Current Offset’ of the path information header (step 204)” and the “local ID, which is referred to by the ‘Current Offset’ after the subtraction, is held.” EX1001, 16:1-6. “The packet transfer unit 320 uses the local ID that is held and performs forwarding processing of the packet in a reverse direction (step S206).” EX1001, 16:7-9. The local ID is used to determine the communication interface 310 “using the information of the forwarding table...and the packet is forwarded from the communication interface 310 in question.” EX1001, 16:9-15.

68. “When a decision is made to forward to the external network 700” (S201=Y or S203=Y), “processing of removing the path information header is performed by the header operation unit 330 (step S205).” EX1001, 16:16-20.

B. Overview of the Prosecution History

69. The ’249 Patent issued based on Patent Application No. 14,244,701,

which was filed on April 3, 2014. EX1002, 2-101. The application contained 21 claims, including 5 independent claims (i.e., application claims 1, 6, 11, 16, and 21). EX1002, 60-63.

70. On December 16, 2014, the examiner rejected all application claims. EX1002, 148-163. Among other things, the examiner rejected the claims based on U.S. Patent No. 7,522,603 (“Vasseur”). EX1002, 157-161.

71. On March 16, 2015, the applicant responded by making minor amendments to the claims and arguing that Vasseur does not anticipate the claims. The applicant asserted (in a section titled “The Claimed Invention”) that “[t]he claimed invention (e.g., as defined by exemplary claim 1) is directed to a packet forwarding node, comprising a storage unit to store a packet forwarding table used for forwarding a packet based on an identifier which identifies a link between nodes on a packet forwarding path; and a forwarding unit to receive an incoming packet including a plurality of the identifiers, and to forward the incoming packet based on the identifier corresponding to the packet forwarding node.” EX1002, 182 (emphasis in original). The applicant also argued that Vasseur, which discussed MPLS routing, does not disclose the claimed invention because “in MPLS, only one identifier (label of MPLS) is included in one packet and a plurality of identifiers are not included in one packet as in the claimed invention....” EX1002, 184 (emphasis in original).

72. On April 15, 2015, the examiner again rejected all application claims based on Vasseur. EX1002, 195-206.

73. On September 15, 2015, the applicant filed a response adding several new application claims (i.e., 22-26) and asserting various arguments to distinguish the prior art. EX1002, 325-335. The applicant argued that the application claims were not anticipated by Vasseur because “[t]he claimed invention (e.g., as defined by exemplary claim 1) is directed to a packet forwarding node, including a storage unit to store a packet forwarding table used for forwarding a packet based on an identifier which identifies a link between nodes on a packet forwarding path, and a forwarding unit to receive an incoming packet including a plurality of the identifiers, and to forward the incoming packet based on the identifier corresponding to the packet forwarding node.” EX1002, 332 (emphasis in original). The applicant further argued that while “[i]n Vasseur, a packet being used in MPLS (Multi-Protocol Label Switching), includes a **single label**,” the application claims required “a packet includes **a plurality of identifiers**.” EX1002, 332 (emphasis in original). The applicant argued that “in MPLS of Vasseur, merely one identifier (label of MPLS) is included in one packet and so a plurality of identifiers are not included in one packet as in the claimed invention....” EX1002, 332-333 (emphasis in original).

74. On October 6, 2015, the examiner issued an Office Action rejecting

all application claims. EX1002, 338-339. Among other things, the examiner construed several claims terms as subject to 35 U.S.C. § 112 ¶6 (pre-AIA). In particular, the examiner applied Section 112 ¶6 to the terms “a storage unit” and “a communication unit.” EX1002, 341-347. The examiner invited the applicant to amend the claims to avoid Section 112 ¶6 if the applicant “does not intend to have the claim limitation(s) treated under 35 U.S.C. 112(f) or pre-AIA 35 U.S.C. 112, sixth paragraph...” EX1002, 341-347. In addition, the examiner again rejected application claims 1-21 based on Vasseur and rejected the new claims 22-26 based on the combination of Vasseur (which the examiner dubbed “Vasseur-1”) and U.S. Patent Publication No. 2007/0280102 (“Vasseur-2”). EX1002, 358-361.

75. On February 1, 2016, the applicant filed a response. EX1002, 413-426. The applicant amended the claims to avoid the examiner’s construction of certain terms under Section 112 ¶6. EX1002, 414-419, 421-422. The applicant also argued that Vasseur-1 does not disclose the claimed invention because it has no “teaching or suggestion that the incoming packet actually includes a plurality of the identifiers.” EX1002, 423. According to the applicant: “In Vasseur, a packet, being used in MPLS (Multi-Protocol Label Switching), includes a single label. On the other hand, in the present application, a packet includes a plurality of identifiers. ... Therefore, in MPLS of Vasseur, merely one identifier (label of MPLS) is included in one packet and so a plurality of identifiers are not included in

one packet as in the claimed invention showing also that the forwarding table in Vasseur and the forwarding table in the claimed invention are not the same.”

EX1002, 424 (emphasis in original).

76. On March 3, 2016, the examiner issued another office action rejecting the application claims. EX1002, 432-461. The examiner began by confirming that following the applicant’s claim amendments, the terms “storage unit” and “communication unit” were not subject to Section 112 ¶6. EX1002, 435-438. The examiner stated that the claim terms “first unit” and “second unit” in application claim 21 would be subject to Section 112 ¶6 absent amendment or further argument. EX1002, 438. In addition, the examiner again rejected application claims 1-21 and 27-28 based on Vasseur-1 and rejected claims 22-26 based on the combination of Vasseur-1 and Vasseur-2. EX1002, 443-458.

77. The applicant filed a notice of appeal, and on August 3, 2016, the applicant filed an appeal brief. EX1002, 509-524. The application argued that “[i]n Vasseur ’603 [i.e., Vasseur-1], a packet, being used in MPLS (Multi-Protocol Label Switching), includes **a single label**. On the other hand, in the present application, a packet includes a plurality of identifiers. ... Therefore, in MPLS of Vasseur ’603, merely one identifier (label of MPLS) is included in one packet and so a plurality of identifiers are not included in one packet as in the claimed invention showing also that the forwarding table in Vasseur ’603 and the

forwarding table in the claimed invention are not the same.” EX1002, 514 (emphasis in original). In addition, the applicant argued that the claims “clarify that the communication unit communicates with a server for receiving a plurality of identifiers, each of which identifies a link between nodes on a packet forwarding path.” EX1002, 516. The applicant argued that “[t]he claims clearly clarify the type of identifier as each of the plurality of identifiers identifies a link between the nodes” and “[n]o such disclosure is identically made in Vasseur ’603,” which “relates to MPLS which includes a single label.” EX1002, 516-517.

78. On November 22, 2016, the examiner issued an office action in response to the applicant’s appeal brief and again rejected all application claims. EX1002, 633-666. The examiner construed a number of terms as subject to Section 112 ¶6 and also identified a number of indefinite terms in the application claims. EX1002, 634-640. The examiner also rejected the applicant claims under 35 U.S.C. § 101. EX1002, 641-645. The examiner also rejected the application claims as obvious over the combination of U.S. Patent No. 7,496,096 (“Dong”) and U.S. Patent Application Publication No. 2011/0149973 (“Rothenberg”). EX1002, 646-665.

79. On February 22, 2017, the applicant filed a response. EX1002, 681-694. The applicant amended the claims to avoid the examiner’s construction of certain terms under Section 112 ¶6 and to address the indefiniteness rejections.

EX1002, 682-689. The applicant argued that the amendments also addressed the examiner's Section 101 rejection. EX1002, 689. The applicant also argued that the combination of Dong and Rothenberg would not teach or suggest every feature of the claimed invention. For instance, "[t]he combination fails to teach or suggest (e.g., claim 1), '*a forwarding unit configured to receive an incoming packet including a plurality of the identifiers, and to forward the incoming packet by using a link that corresponds to the identifier from among the plurality of the identifiers.*'" EX1002, 690 (emphasis in original). The applicant argued: "The claimed invention has a technical feature that when the packet forwarding node receives a packet including identifier for links connecting nodes on a packet forwarding path, the packet forwarding node forwards this packet based on the identifier included [in] the received packet. That is, in the present application, the packet itself has [the] technical feature" (e.g., Figure 4). EX1002, 690. Dong, on the other hand, "discloses a generation method for [a] routing table regarding a network where IP path and MPLS path are mixed." EX1002, 690. In other words, the applicant argued that the application claims require a plurality of identifiers carried in the data packet itself and that Dong does not disclose a plurality of identifiers in the packet. The applicants argued that Rothenberg "is silent as to the above mentioned limitations...." EX1002, 692.

80. On April 24, 2017, an examiner-initiated interview was held.

EX1002, 738.

81. On May 31, 2017, the examiner issued an office action. EX1002, 708-737. The examiner began by confirming that applicant's claim amendments had avoided Section 112 ¶6 and addressed the indefiniteness rejections. EX1002, 711-712. The examiner again rejected most of the application claims under Section 101. EX1002, 712-716. The examiner also rejected the application claims based on the combination of Dong and Rothenberg. EX1002, 717-736.

82. Another interview was conducted on July 25, 2017, and electronic mail correspondence was exchanged on July 27 and 28, 2017, addressing proposed claim amendments.

83. On August 31, 2017, the applicant filed a response making a number of amendments to the independent application claims. EX1002, 750-759. The applicant noted and agreed with the examiner that the claims "do not invoke" Section 112 ¶6. EX1002, 757. For example, the applicant amended application claim 1 to recite that the "forwarding unit" is configured to "determine a direction to forward the incoming packet," "determine whether to forward to an external network," "add the number of bytes to path/link information when the incoming packet is not forwarded to the external network," and "remove path/link information header when the incoming packet is forwarded to the external network." EX1002, 750. Similar amendments were made to the other independent

claims. EX1002, 751-755. The applicant argued that the claim amendments addressed the Section 101 rejection. EX1002, 758. The applicant also argued that the cited prior art references did not disclose the amended application claims. EX1001, 758-759.

84. On October 2, 2017, the examiner issued a Notice of Allowance. EX1002, 763-771. The examiner again confirmed that the claim terms do not invoke Section 112 ¶6. EX1002, 770. The examiner also stated that the cited prior art references did not disclose the limitations added through the most recent claim amendments. EX1002, 770-771.

85. On November 9, 2017, the applicant filed a response to the Notice of Allowance making several more amendments to the claims. EX1002, 791-798. For example, the applicant amended claim 1 to change the limitation reciting “add the number of bytes to path/link information when the incoming packet is not forwarded to the external network” to delete the words “the number of bytes.” EX1002, 792. The applicant amended the other independent claims in the same way. EX1002, 793-794, 796-797. In addition, the applicant changed every occurrence of “path/link” to “path or link.” EX1002, 792-797.

86. On November 30, 2017, an examiner-initiated interview was held and further amendments after allowance were discussed.

87. On December 27, 2017, the examiner issued a Corrected Notice of

Allowability making the applicants amendments and adding a few more. EX1002, 815-823.

88. The '249 Patent issued on February 20, 2018. EX1001, 1.

C. The Patent Claims

89. The '249 Patent contains 27 claims, including 5 independent claims (i.e., 1, 6, 11, 16, 21).

90. Claim 1 reads as follows:

[1pre] A packet forwarding node, comprising:

[1a] a storage unit including a non-transitory computer readable medium configured to store a packet forwarding table used for forwarding a packet based on an identifier which identifies a link between nodes on a packet forwarding path; and

[1b] at least one processor configured to execute program instructions to provide a forwarding unit configured to:

[1c] receive an incoming packet including a plurality of the identifiers;

[1d] determine a direction to forward the incoming packet;

[1e] determine whether to forward to an external network;

[1f] add to path or link information when the incoming packet is not forwarded to the external network;

[1g] remove path or link information header when the incoming packet is forwarded to the external network; and

[1h] forward the incoming packet by using a path or link that

corresponds to the identifier from among the plurality of the identifiers.

I have separately identified the limitations (1pre, 1a, 1b, etc.) for ease of reference in my analysis below.

91. Each of the five independent claims are similar to the others. The following table compares the language of independent claims 1 and 6, identifying differences in claim language with bold and underlining.

Claim 1	Claim 6
A packet forwarding node, comprising:	A packet forwarding node, comprising:
a storage unit including a non-transitory computer readable medium configured to store a packet forwarding table used for forwarding a packet based on an identifier which identifies <u>a link between nodes on</u> a packet forwarding <u>path</u> ; and	a storage unit including a non-transitory computer readable medium configured to store a packet forwarding table used for forwarding a packet based on an identifier which identifies <u>an interface of</u> a packet forwarding <u>node</u> ; and
at least one processor configured to execute program instructions to provide a forwarding unit configured to:	at least one processor configured to execute program instructions to provide a forwarding unit configured to
receive an incoming packet including a plurality of <u>the</u> identifiers;	receive an incoming packet including a plurality of identifiers,
determine a direction to forward the	determine a direction to forward the

Claim 1	Claim 6
incoming packet ₂	incoming packet ₂
determine whether to forward to an external network ₂	determine whether to forward to an external network ₂
add to path or link information when the incoming packet is not forwarded to the external network ₂	add to path or link information when the incoming packet is not forwarded to the external network ₂
remove path or link information header when the incoming packet is forwarded to the external network ₂ and	remove path or link information header when the incoming packet is forwarded to the external network ₂ and
forward the incoming packet by using a path or link that corresponds to the identifier from among the plurality of <u>the</u> identifiers.	<u>to</u> forward the incoming packet by using a path or link that corresponds to the identifier from among the plurality of identifiers.

92. The following table compares the language of independent claims 1 and 11, identifying differences in claim language with bold and underlining.

Claim 1	Claim 11
A packet forwarding node, comprising: <u>a storage unit including a non-transitory computer readable medium configured to store a packet forwarding table used for forwarding a packet based on an</u>	A packet forwarding node, comprising:

Claim 1	Claim 11
<u>identifier which identifies a link between nodes on a packet forwarding path; and</u>	
	<u>a memory storing a program;</u>
at least one processor configured to execute program <u>instructions</u> to provide...	at least one processor configured to execute <u>the</u> program to provide:
	<u>a communication unit configured to communicate with a server for receiving a plurality of identifiers, each of which identifies a link between nodes on a packet forwarding path; and</u>
a forwarding unit configured to:	a forwarding unit configured to:
receive an incoming packet including <u>a</u> plurality of the identifiers;	receive an incoming packet including <u>the</u> plurality of the identifiers;
determine a direction to forward the incoming packet;	determine a direction to forward the incoming packet;
determine whether to forward to an external network;	determine whether to forward to an external network;
add to path or link information when the incoming packet is not forwarded to the external network;	add to path or link information when the incoming packet is not forwarded to the external network;

Claim 1	Claim 11
remove path or link information header when the incoming packet is forwarded to the external network ₂ ; and	remove path or link information header when the incoming packet is forwarded to the external network ₂ ; and
forward the incoming packet by using a path or link that corresponds to <u>the</u> identifier from among the plurality of <u>the</u> identifiers.	forward the incoming packet by using a path or link that corresponds to <u>an</u> identifier from among the plurality of identifiers.

93. The following table compares the language of independent claims 1 and 16, identifying differences in claim language with bold and underlining.

Claim 1	Claim 16
A packet forwarding node, comprising:	A packet forwarding node, comprising:
<u>a storage unit including a non-transitory computer readable medium configured to store a packet forwarding table used for forwarding a packet based on an identifier which identifies a link between nodes on a packet forwarding path; and</u>	
	<u>a communication unit configured to communicate with a server for receiving a plurality of identifiers,</u>

Claim 1	Claim 16
	<u>each of which identifies an interface of the packet forwarding node; and</u>
at least one processor configured to execute program instructions to provide a forwarding unit configured to:	at least one processor configured to execute program instructions to provide a forwarding unit configured to
receive an incoming packet including a plurality of the identifiers;	receive an incoming packet including a plurality of identifiers,
determine a direction to forward the incoming packet;	determine a direction to forward the incoming packet,
determine whether to forward to an external network;	determine whether to forward to an external network,
add to path or link information when the incoming packet is not forwarded to the external network;	add to path or link information when the incoming packet is not forwarded to the external network,
remove path or link information header when the incoming packet is forwarded to the external network; and	remove path or link information header when the incoming packet is forwarded to the external network, and
forward the incoming packet by using a path or link that corresponds to <u>the</u> identifier from among the plurality of <u>the</u> identifiers.	forward the incoming packet by using a path or link that corresponds to <u>an</u> identifier from among the plurality of identifiers.

94. The following table compares the language of independent claims 1

and 21, identifying differences in claim language with bold and underlining.

Claim 1	Claim 21
A <u>packet forwarding node</u> , comprising:	A <u>server</u> , comprising:
<u>a storage unit including a non-transitory computer readable medium configured to store a packet forwarding table used for forwarding a packet based on an identifier which identifies a link between nodes on a packet forwarding path; and</u>	
	<u>a memory storing a computer program;</u>
at least one processor configured to execute program <u>instructions</u> to provide <u>a forwarding unit configured to:</u>	at least one processor configured to execute <u>the computer</u> program to provide:
	<u>a first unit configured to determine a plurality of identifiers, each of which identifies a link between nodes on a packet forwarding path; and</u>
	<u>a second unit configured send information elements to a packet</u>

Claim 1	Claim 21
	<u>forwarding node which is capable of</u>
receive <u>e</u> an incoming packet including a plurality of <u>the</u> identifiers ₂ ;	receiv <u>ing</u> an incoming packet including a plurality of identifiers ₂ ;
determine <u>e</u> a direction to forward the incoming packet ₂ ;	determin <u>ing</u> a direction to forward the incoming packet ₂ ;
determine <u>e</u> whether to forward to an external network ₂ ;	determin <u>ing</u> whether to forward to an external network ₂ ;
add to path or link information when the incoming packet is not forwarded to the external network ₂ ;	add <u>ing</u> to path or link information when the incoming packet is not forwarded to the external network ₂ ;
remove <u>e</u> path or link information header when the incoming packet is forwarded to the external network ₂ ; and	remov <u>ing</u> path or link information header when the incoming packet is forwarded to the external network ₂ ; and
forward the incoming packet by using a path or link that corresponds to <u>the</u> identifier from among the plurality of <u>the</u> identifiers.	forward <u>ing</u> the incoming packet by using a path or link ₂ that corresponds to <u>an</u> identifier from among the plurality of identifiers.

95. Many of the dependent claims are also quite similar. The following table compares the language of dependent claims 2 and 7, identifying differences in claim language with bold and underlining.

Claim 2	Claim 7
The packet forwarding node according to claim <u>1</u> , wherein the storage unit stores the table used for forwarding the packet based on the identifier which uniquely identifies the link within a predetermined domain.	The packet forwarding node according to claim <u>6</u> , wherein the storage unit stores the table used for forwarding the packet based on the identifier which uniquely identifies the link within a predetermined domain.

96. The following table compares the language of dependent claims 2 and 12, identifying differences in claim language with bold and underlining.

Claim 2	Claim 12
The packet forwarding node according to claim <u>1</u> , wherein the <u>storage unit stores the table used for forwarding</u> the packet based on the identifier which uniquely identifies the link within a predetermined local domain.	The packet forwarding node according to claim <u>11</u> , wherein the <u>forwarding unit forwards</u> the packet based on the identifier which uniquely identifies the link within a predetermined domain.

97. The following table compares the language of dependent claims 12 and 17, identifying differences in claim language with bold and underlining.

Claim 12	Claim 17
The packet forwarding node according to claim <u>11</u> , wherein the forwarding unit forwards the packet based on the	The packet forwarding node according to claim <u>16</u> , wherein the forwarding unit forwards the packet based on the

Claim 12	Claim 17
identifier which uniquely identifies the link within a predetermined domain.	identifier which uniquely identifies the link within a predetermined domain.

98. The following table compares the language of dependent claims 3 and 8, identifying differences in claim language with bold and underlining.

Claim 3	Claim 8
The packet forwarding node according to claim <u>1</u> , wherein the storage unit stores the table used for forwarding the packet based on the identifier which uniquely identifies the link within a predetermined local domain.	The packet forwarding node according to claim <u>6</u> , wherein the storage unit stores the table used for forwarding the packet based on the identifier which uniquely identifies the link within a predetermined local domain.

99. The following table compares the language of dependent claims 3 and 13, identifying differences in claim language with bold and underlining.

Claim 3	Claim 13
The packet forwarding node according to claim <u>1</u> , wherein the <u>storage unit stores the table used for forwarding</u> the packet based on the identifier which uniquely identifies the link within a predetermined local domain.	The packet forwarding node according to claim <u>11</u> , wherein the <u>forwarding unit forwards</u> the packet based on the identifier which uniquely identifies the link within a predetermined local domain.

100. The following table compares the language of dependent claims 13

and 18, identifying differences in claim language with bold and underlining.

Claim 13	Claim 18
The packet forwarding node according to claim <u>11</u> , wherein the forwarding unit forwards the packet based on the identifier which uniquely identifies the link within a predetermined local domain.	The packet forwarding node according to claim <u>16</u> , wherein the forwarding unit forwards the packet based on the identifier which uniquely identifies the link within a predetermined local domain.

101. The following table compares the language of dependent claims 4 and 9, identifying differences in claim language with bold and underlining.

Claim 4	Claim 9
The packet forwarding node according to claim <u>I</u> , wherein the storage unit stores the table used for forwarding the packet based on the identifier which uniquely identifies the link within the packet forwarding node.	The packet forwarding node according to claim <u>6</u> , wherein the storage unit stores the table used for forwarding the packet based on the identifier which uniquely identifies the link within the packet forwarding node.

102. The following table compares the language of dependent claims 4 and 14, identifying differences in claim language with bold and underlining.

Claim 4	Claim 14
The packet forwarding node according to claim <u>I</u> , wherein the <u>storage unit</u>	The packet forwarding node according to claim <u>11</u> , wherein the <u>forwarding</u>

Claim 4	Claim 14
<u>stores the table used for forwarding</u> the packet based on the identifier which uniquely identifies the link within the packet forwarding node.	<u>unit forwards</u> the packet based on the identifier which uniquely identifies the link within the packet forwarding node.

103. The following table compares the language of dependent claims 14 and 19, identifying differences in claim language with bold and underlining.

Claim 14	Claim 19
The packet forwarding node according to claim <u>11</u> , wherein the forwarding unit forwards the packet based on the identifier which uniquely identifies the link within the packet forwarding node.	The packet forwarding node according to claim <u>16</u> , wherein the forwarding unit forwards the packet based on the identifier which uniquely identifies the link within the packet forwarding node.

104. The following table compares the language of dependent claims 5 and 10, identifying differences in claim language with bold and underlining.

Claim 5	Claim 10
The packet forwarding node according to claim <u>1</u> , wherein the storage unit stores the table used for forwarding the packet based on the identifier which uniquely identifies the link among the packet forwarding nodes.	The packet forwarding node according to claim <u>6</u> , wherein the storage unit stores the table used for forwarding the packet based on the identifier which uniquely identifies the link among the packet forwarding nodes.

105. The following table compares the language of dependent claims 5 and 15, identifying differences in claim language with bold and underlining.

Claim 5	Claim 15
The packet forwarding node according to claim <u>1</u> , wherein the <u>storage unit stores the table used for forwarding</u> the packet based on the identifier which unique identifies the link among the packet forwarding nodes.	The packet forwarding node according to claim <u>11</u> , wherein the <u>forwarding unit forwards</u> the packet based on the identifier which uniquely identifies the link among the packet forwarding nodes.

106. The following table compares the language of dependent claims 15 and 20, identifying differences in claim language with bold and underlining.

Claim 15	Claim 20
The packet forwarding node according to claim <u>11</u> , wherein the forwarding unit forwards the packet based on the identifier which uniquely identifies the link among the packet forwarding nodes.	The packet forwarding node according to claim <u>16</u> , wherein the forwarding unit forwards the packet based on the identifier which uniquely identifies the link among the packet forwarding nodes.

107. The following table compares the language of dependent claims 22 and 23, identifying differences in claim language with bold and underlining.

Claim 22	Claim 23
The packet forwarding node according	The packet forwarding node according

Claim 22	Claim 23
to claim <u>1</u> , wherein the storage unit stores the table used for forwarding the packet based on the identifier which identifies the link between adjacent nodes on a packet forwarding path.	to claim <u>6</u> , wherein the storage unit stores the <u>packet forwarding</u> table used for forwarding the packet based on the identifier which identifies the link between adjacent nodes on a packet forwarding path.

108. The following table compares the language of dependent claims 22 and 24, identifying differences in claim language with bold and underlining.

Claim 22	Claim 24
The packet forwarding node according to claim <u>1</u> , <u>wherein the</u> storage unit stores <u>the</u> table used for forwarding the packet based on <u>the</u> identifier which identifies the link between adjacent nodes on a packet forwarding path.	The packet forwarding node according to claim <u>11</u> , <u>further comprising a</u> storage unit <u>that</u> stores <u>a</u> table used for forwarding the packet based on <u>an</u> identifier which identifies the link between adjacent nodes on a packet forwarding path.

109. The following table compares the language of dependent claims 22 and 25, identifying differences in claim language with bold and underlining.

Claim 22	Claim 25
The packet forwarding node according to claim <u>1</u> , <u>wherein the</u> storage unit stores <u>the</u> table used for forwarding the	The packet forwarding node according to claim <u>16</u> , <u>further comprising a</u> storage unit <u>that</u> stores <u>a</u> table used for

Claim 22	Claim 25
packet based on <u>the</u> identifier which identifies the link between adjacent nodes on a packet forwarding path.	forwarding the packet based on <u>an</u> identifier which identifies the link between adjacent nodes on a packet forwarding path.

110. In addition, dependent claim 26 depends from claim 21, and dependent claim 27 depends from claim 1.

D. Person of Ordinary Skill in the Art

111. As discussed above, I have been informed and understand that the level of ordinary skill in the relevant art at the time of the claimed invention is relevant to inquiries such as the meaning of claim terms, the meaning of disclosures in prior art references, and the reasons that a person of ordinary skill in the art might have combined the teachings of certain prior art references.

112. In my opinion, a person of ordinary skill in the art with regard to the '249 Patent would have had an undergraduate degree (or equivalent) in electrical engineering, computer science, or a comparable subject and approximately two to three years of additional academic or professional work experience in a technical field with exposure to computer networking and switching/routing technology. A higher level of education could substitute for less industry experience and more industry experience could substitute for the specific level of education. In any event, such a person would have been knowledgeable

about then-existing network switching and routing technologies and techniques.

113. In reaching my opinion about the level of ordinary skill in the art, I have considered the types of problems, and techniques for addressing them, discussed in the '249 Patent and in the prior art references cited on its face. The patent discusses the problem of routing data packets through a network and presumes a certain level of background understanding. For example, the '249 Patent discusses IP (Internet Protocol) networks without specifically describing them and states that “[s]ince the communication nodes 100a and 100b are similar to a general IP node, a detailed description thereof is omitted,” which indicates that a person of ordinary skill in the art would already understand IP networks and switching/routing technology. The cited prior art discussed during prosecution also addresses problems in the technical field of network switching/routing. The problems and technical solutions discussed in the '249 Patent and the prior art cited on the face of the patent support my opinion regarding the level of ordinary skill in the art.

114. From my work in academia and industry in the 2000s and into the 2010s, I am familiar with the rapidity with which innovations were being made in the art and the education level of active works in the relevant field. Thus, my personal experience also informs and supports my opinion regarding the level of ordinary skill in the art.

115. In view of my educational background, as well as my decades of academic and professional experience (as discussed above), I was a person of more than the ordinary level of skill in the art in the 2009-2010 timeframe. Although my qualifications and experience exceed those of the hypothetical person of ordinary skill in the art defined above, my analysis and opinions regarding the '249 Patent have been based on the perspective of a person of ordinary skill in the art in 2009-2010.

116. To the extent that a person of ordinary skill in the art may be defined slightly differently than I define it above, or the qualifications have minor substantive differences, I would still be at least a person of ordinary skill in the art based on my academic and industry background described in this declaration. I have considered whether my opinions expressed in this declaration would change if the level of ordinary skill in the art were varied by a few years or varied somewhat with respect to subject matter. My opinions would not change in light of such minor variations.

E. Claim Construction

117. I have been advised and understand that patent claims in an *inter partes* review are interpreted by applying the same standard that applies in federal district court proceedings, which I understand is referred to as the *Phillips* standard.

118. I understand that under the *Phillips* standard, claim terms are typically given their ordinary and customary meaning as would have been understood by a POSITA at the time of the invention based on the language of the claims, the patent specification, and the prosecution history. I understand that these sources of evidence are often referred to as “intrinsic evidence” or the “intrinsic record” of the patent. I understand that under the *Phillips* standard, extrinsic evidence such as expert or inventor testimony, dictionary definitions, and prior art publications can also be considered but are given less weight than the intrinsic evidence.

119. In my opinion, the terms of the challenged claims can be understood based on their plain and ordinary meaning in light of the intrinsic evidence without further specific construction.

VI. PRIOR ART REFERENCES

120. I am informed that each of the following references constitutes prior art to the '249 Patent.

A. Shimizu (EX1005)


121. Shimizu is a published U.S. patent application that was published on October 12, 2006. Shimizu is entitled “Network Based Routing Scheme.” EX1005, 1.

122. The Abstract states that “[a] network organization and routing scheme are provided.” EX1005, 1 (Abstract). Shimizu further states: “These techniques

enable a network to control the paths that frames take through the network. A device transmitting a frame may add tags to the frame in order to define a particular path that the frame should take through the network. For example, the tags may indicate a link-by-link path between switches in the network.”

EX1005, 1 (Abstract); see also [0004] (“According to particular embodiments, these techniques describe a routing scheme and network organization that enable a private network to control the path frames take through the private network.”), [0006] (“For example, these techniques may allow a device to specify a particular path that a frame should take th[r]ough a private network.”), [0020] (“In operation, devices can control the path of frames through switch network 12 using tags in the frames.”).

123. Figure 1 of Shimizu “illustrates a system, indicated generally at 10, that uses a network based routing scheme to control the path of frames through a switch network.” EX1005, [0014].



EX1005, Fig. 1. The system “includes a switch network 12, a management device 14, network devices 16, and internet 18, and an external network 20.”

EX1005, [0014]. “As illustrated, switch network 12 includes switches 22 organized into groups 24.” EX1005, [0015]. Shimizu explains that “switches 22 may include switches, routers, or any appropriate network devices capable of interconnecting network elements.” EX1005, [0016].

124. Shimizu states that in some embodiments “switch network 12 is a private network that allows devices to specifically determine the path of a frame through the network.” EX1005, [0015]. Shimizu also states that “devices can

control the path of frames through switch network 12 using tags in frames.”

EX1005, [0020]. Shimizu describes two types of tags: “inbound tags and outbound tags.” EX1005, [0020]. “Frames may include header information indicating that the frame contains these tags.” EX1005, [0020].


125. Shimizu’s “[i]nbound tags may be used, for example, when a frame enters switch network 12 to ensure that responses to the frame are sent from switch network 12 on the same port that originally received the frame.” EX1005, [0020]. Shimizu states that “[w]hen an inbound frame enters switch network 12 at a specific port, the receiving switch 22 may add an inbound tag indicating the group/switch/port that received the frame.” EX1005, [0020]. Further, in some embodiments, “each switch 22 receiving an inbound frame adds an inbound tag indicating the group/switch/port that received the frame.” EX1005, [0021]. In this way, “the path a frame has taken through switch network 12 can be known by an examination of the inbound tags in the frame....” EX1005, [0021]; see also [0027] (“In certain embodiments, one or more of switches 22 add inbound tags to network management messages so that management device 14 knows the port the network management message entered switch network 12 and/or the precise route that each network management message took to reach the management device 14.”).

126. “Outbound tags may be used to determine the particular path a frame should take through switch network 12.” EX1005, [0022]. To use outbound tags,

“a transmitting device may add one or more tags to the header of the frame.”


EX1005, [0022]. “Each outbound tag may enable switch network 12 to transport the frame along a set portion of switch network 12, such as over a link between two switches 22 or across one group 24.” EX1005, [0022]. A sequence of “outbound tags in the frame can specify the path(s) of the frame, link by link through switch network 12.” EX1005, [0023]. “Each receiving switch 22 simply identifies the port(s) specified by the current outbound tag, removes the current outbound tag from the frame, and then forwards the frame on the specified port(s).” EX1005, [0023].

127. Figure 3A “illustrates the path 50 of a frame....” EX1005, [0041]. As shown, “path 50 includes five different hops (1-5)” starting at management device 14, proceeding through four switches, and exiting the switch network to external network 20. EX1005, [0042]. In this example, the frame is proceeding in an “outbound” direction. See EX1005, [0022] (“[M]anagement device 14 and network devices 16 can transmit frames with outbound tags in order to define a desired frame path.”).



EX1005, Fig. 3A. Shimizu describes using a plurality of tags in the frame header to dictate the sequence of hops (i.e., the path) through the switch network.

128. For example, in some embodiments, “the frame is structured to include vector based tags” as illustrated in Figure 3B. EX1005, [0042].



EX1005, Fig. 3B. “[T]he frame is structured with a vector based tag indicating the next switch 22 in path 50.” EX1005, [0043]. “As illustrated, a frame includes a destination address 52, a source address 54, one or more (or no) tags 56, and a payload 58.” EX1005, [0043]. The illustrated tags direct the frame along path 50 illustrated in Figure 3A.

129. Shimizu explains that “[t]he frame corresponding to hop 1 [in Figure 3B] represents the structure of the frame sent from management device 14 to switch X.A.” EX1005, [0044]. “During hop 1, the frame contains destination address 52, source address 54, four tags 56, and payload 58. Switch X.A receives the frame and extracts tag 56a. Because tag 56a contains the value 5, switch X.A

sends the frame on port 5 after removing tag 56a from the frame. So, during hop 2, the frame contains destination address 52, source address 54, three tags 56, and payload 58.” EX1005, [0044]. The frame proceeds through hops 3, 4, and 5 of path 50 in a similar manner and exits to external network 20. EX1005, [0044].

130. Shimizu also describes a frame that “is structured to include port ID based tags” as illustrated in Figure 3C. EX1005, [0042]. Again, the illustrated tags direct the frame to proceed along path 50 illustrated in Figure 3A.




FIG. 3C

EX1005, Fig. 3C. In this embodiment, “the frame is structured with a port ID based tag indicating the next group 24 in path 50” (i.e., a group of switches within switch network 12 as shown in Figure 1). EX1005, [0046]. The frame “includes destination address 52, source address 54, one or more (or no) tags 60, and payload 58.” EX1005, [0046]. “Tags 60 operate [to] control the path of frames through switch network 12.” EX1005, [0046].

131. Shimizu explains that, as in the prior example, “[t]he frame

corresponding to hop 1 [as illustrated in Figure 3C] represents the structure of the frame sent from management device 14 to switch X.A.” EX1005, [0047]. “During hop 1, the frame contains destination address 52, source address 54, three tags 60, and payload 58. Switch X.A receives the frame and extracts tag 60a. Tag 60a specifies port X.C.4: port 4 of switch X.C.” EX1005, [0047]. “Switch X.A first compares the group identification of tag 60a, ‘X,’ to its group identification” and “[f]inding there is a match, switch X.A then compares the switch identification of tag 60a, ‘C,’ to its switch identification.” EX1005, [0047]. Because the switch identification (“C”) does not match, “switch X.A looks to its routing information to determine the appropriate outbound port for a frame destined for switch X.C.” EX1005, [0047]. In other words, switch X.A uses its routing tables to identify the appropriate outbound port. “Switch X.A identifies port 5 as the correct port and sends the frame on port 5 towards switch X.C.” EX1005, [0047]. In this example, “switch X.A does not remove tag 60a from the frame....” EX1005, [0047]. “During hop 2, ...[s]witch X.C receives the frame and extracts tag 60a.” EX1005, [0048]. Switch X.C then “compares the group identification of tag 60a, ‘X,’ to its group identification.” EX1005, [0048]. Because there is a match, “switch X.C then compares the switch identification to tag 60a, ‘C,’ to its switch identification” and finds a match. EX1005, [0048]. “Because there is a match, switch X.C evaluates the port identification” (i.e., 4), “removes tag 60a, and then sends the

frame on port 4 towards switch Y.C.” EX1005, [0048]. The frame proceeds through hops 3, 4, and 5 of path 50 in a similar manner and exits to external network 20. EX1005, [0049]-[0050].

B. RFC-4271 (EX1006)

132. RFC-4271 is a Request for Comments published by the IETF in January 2006. It is entitled “A Border Gateway Protocol 4 (BGP-4).” EX1006, 1. The document “discusses the Border Gateway Protocol (BGP), which is an inter-Autonomous System routing protocol.” EX1006, 1.

133. “The primary function of a BGP speaking system is to exchange network reachability information with other BGP systems.” EX1006, 1. “BGP-4 provides a set of mechanisms for supporting Classless Inter-Domain Routing (CIDR),” which “include support for advertising a set of destinations as an IP prefix, and eliminating the concept of network ‘class’ within BGP.” EX1006, 1.

134. RFC-4271 was the specification for BGP4 in effect in the 2009-2010 timeframe when the ’249 Patent was filed. The BGP4 protocol itself had been defined at least since the publication of RFC-1654 in July 1994. I personally taught the key elements of the BGP4 protocol, as described in RFC-4271, in my courses beginning in 1996. Thus, I am personally familiar with RFC-4271 (EX1006) and can confirm that it was readily available to a person of ordinary skill in the art on the IETF website following its publication in 2006 and certainly

before 2009. Further, as I stated above, once RFCs are published, the IETF does not make changes to any published RFC. EX1009, 4. This is consistent with my personal experience going back well before RFC-4271 was published in 2006. Thus, it is my understanding, and I believe, that RFC-4271 available today through the RFC Editor website is the same as the version published in 2006.

VII. THE CHALLENGED CLAIMS ARE OBVIOUS IN LIGHT OF SHIMIZU

135. It is my opinion that Shimizu discloses and renders obvious the challenged claims of the '249 Patent for at least the reasons discussed below.

136. A person of ordinary skill in the art would have naturally considered Shimizu because it is in the same field of endeavor as the '249 Patent, namely network communications. EX1005, 1 (Abstract), [0002]; EX1001, 1 (Abstract), 1:9-15, 4:14-67. Further, like the '249 Patent, Shimizu addresses the problem of controlling the path of data packets through a private network. EX1005, [0004], [0006]-[0007]; EX1001, 2:1-9, 5:1-10. Indeed, Shimizu and the '249 Patent address the problem in the same way by using source routing, i.e., by adding a series of identifiers, or tags, to the header of a data packet to dictate each hop along a path through the network.

137. Further, Shimizu identifies advantages of the disclosed routing technique that are the same types of advantages identified in the '249 Patent. Shimizu states that “[i]n certain embodiments, network routing tables maintain

routing information based on switch/port identifiers,” which “enables smaller table sizes and simpler look-up mechanisms, since switches 22 in one group 24 need not maintain routing information for switches 22 in other groups 24.” EX1005, [0019]; see also [0006] (“In some embodiments, appropriate routing can be enabled with smaller routing tables that require less memory in a switch.”). These benefits of Shimizu’s technique address the same types of problems identified in the ’249 Patent. EX1001, 3:20-4:20.

138. Shimizu provides a sufficiently detailed description to permit a person of ordinary skill in the art to implement the aspects of Shimizu’s network routing technique discussed below with regard to the challenged claims with a reasonable expectation of success and without undue experimentation. It would have been within the skill level of a person of ordinary skill in the art to implement Shimizu’s techniques based on the disclosure.

A. Independent Claim 1

1. [1pre] A packet forwarding node, comprising:

139. Shimizu discloses a *packet forwarding node*, as recited in the preamble of claim 1. (Note that for ease of reference I will use italics to indicate claim language throughout my analysis.) In particular, Shimizu discloses a “[s]witch network 12 [that] includes switches 22....” EX1005, [0015]. “Switches 22 represent networking devices that route frames within networks

and/or between different networks. According to particular embodiments, each switch 22 represents a network device with multiple ports capable of switching network traffic among the ports.” EX1005, [0016]. Each switch 22 is a *packet forwarding node* as recited in the claim.

140. As Shimizu explains, “[w]ithin networks switches route frames between different network elements.” EX1005, [0003]. The term “frame” is another name for a data packet. In particular, a “frame” is a data packet being transmitted through Layer 2 switches using, for example, the Ethernet protocol. Shimizu states that “[f]or example, switches 22 may each represent a chip-based switching device capable of switching frames among multiple ports based on standard Ethernet protocols.” EX1005, [0016]. A person of ordinary skill in the art would know that when an IP packet, or datagram, is transmitted via the Ethernet protocol (or any link-layer protocol) it is called a “frame.” EX1010, 75 (“In this book, we’ll refer to link-layer packets as frames.”), 77-78 (“The datagram is then passed to the link layer, which (of course!) will add its own link-layer header information and create a link-layer frame. Thus, we see that at each layer, a packet has two types of fields: header fields and a payload field. The payload is typically a packet from the layer above.”).

141. Shimizu expressly describes switches in switch network 12 *forwarding frames (packet[s])*. For example, Shimizu states that “[s]witch Z.C

may add an inbound tag before ***forwarding the frame*** on the port indicated by routing information maintained on switch Z.C, possibly ***forwarding the frame*** to switch Y.C.” EX1005, [0021]. Similarly, “[s]witch Y.C may ***forward the frame*** on the port indicated by its routing information....” EX1005, [0021]. Shimizu states that “[e]ach receiving switch 22 simply identifies the port(s) specified by the current outbound tag, removes the current outbound tag from the frame, and then ***forwards the frame*** on the specified port(s).” EX1005, [0023]; see also [0024] (“switch Y.B ***forwards the frame*** on that port”), [0028] (“switches 22 may ***forward*** management messages based on type or other suitable factors”), [0037] (“switch Z.A strips the first tag from the frame and ***forwards the frame*** on the port identified by the tag”), [0058] (“Switch 22 determines, from the routing information, on which port the ***frame should be forwarded....***”). Each switch 22 in switch network 12 is a *packet forwarding node*.

142. Thus, Shimizu discloses a *packet forwarding node* as recited in the preamble of claim 1 and explained further below.

2. **[1a] a storage unit including a non-transitory computer readable medium configured to store a packet forwarding table used for forwarding a packet based on an identifier which identifies a link between nodes on a packet forwarding path; and**


143. Shimizu discloses this limitation.

144. In Shimizu, each switch 22 (*packet forwarding node*) has a

forwarding table. Shimizu states that “network routing information is stored in network **routing tables**...” EX1005, [0019]. For example, Shimizu states: “In particular embodiments, switches within a private network are divided into groups, and a specific port of a specific one of the switches can be uniquely referenced by a tag that indicates a group, a switch, and a port identifier. This can permit switches within a group to maintain private **routing tables** defined for only other switches in the group. In some embodiments, appropriate routing can be enabled with smaller **routing tables** that require less memory in a switch.” EX1005, [0006]; see also [0019] (“The network **routing table** for each constituent switch 22 in a particular group 24 may only provide routing information for other switches 22 in that particular group 24.”).

145. A POSITA would understand that the “routing tables” in Shimizu are *forwarding tables* because they are tables with the information used by switches 22 to forward frames (*packets*) from one switch to the next through the network. A POSITA would have used the terms “routing table” and “forwarding table” interchangeably in this context. EX1010, 95-97, 263-264; EX1012, 7:41-59; EX1013, 1 (Abstract), 1:26-37, 2:66-3:37, 6:5-20; EX1014, 2:25-60; EX1015, 1:17-2:44, 4:24-51, 6:19-43, Figs. 1, 3.


146. Shimizu illustrates two exemplary routing tables (each *a packet forwarding table*) in Figures 2B and 2C.



EX1005, Figs. 2B, 2C. Shimizu states that Figure 2B “illustrates one example of how routing information can be stored by switches 22. In the illustrated embodiment, switch Z.A stores routing information as a network routing table 34.” EX1005, [0035]. Shimizu’s network routing table 34 (illustrated in Figure 2B) is a *packet forwarding table*. Similarly, Figure 2C “illustrates an example of group routing information that may be stored and maintained by switch Z.A.” EX1005, [0039]. Shimizu states that “switch Z.A maintains a table 40 providing routing information for group Y, as a whole, and group X, as a whole.” EX1005, [0039]. Shimizu’s group routing table 40 (illustrated in Figure 2C) is also a *packet forwarding table*.

147. Figure 2A “illustrate[s] one of groups 24, group Z, employing a network based routing scheme....” EX1005, [0033]. Figure 2A illustrates a

portion of the switch network illustrated in Figure 1.



EX1005, Fig. 2A. The sample routing tables illustrated in Figures 2B and 2C would reside on switch Z.A in group Z. EX1005, [0033].

148. Figure 2A illustrates “an embodiment of one of groups 24, group Z, and its constituent switches 22: switch Z.A, switch Z.B, switch Z.C, and switch Z.D.” EX1005, [0033]. “The constituent switches 22 are connected by intra-group links 30 and external links 32” where “[i]ntra-group links 30 connect constituent switches 22 to other constituent switches 22, and external links 32 connect constituent switches 22 to switches 22 in other groups 24 and network devices 16m,n” as well as to “a variety of networks, components, elements, or devices” outside the group such as “other network devices 16, management device 14, internet 18, and external network 20.” EX1005, [0033].

149. As shown in Figure 2A, switch Z.A has six links “designated by ports 0 through 5.” EX1005, [0034]. “Network routing table 34 [in Figure 2B] includes a switch column 36 indicating a switch or device that should receive a frame, and an outbound port column 38 showing the corresponding outbound port on which a frame should be sent.” EX1005, [0035]. Each outbound port on the switch designates a link between that switch and other switches and devices. EX1005, [0034].

150. Using Shimizu’s figures to illustrate a hypothetical example, assume that a frame arrives at switch Z.A with header information indicating that the frame should be forwarded to switch Z.B, which is in the same group 24 (group Z). In that case, switch Z.A would use network routing table 34 illustrated in Figure 2.B (below left) to determine that the correct outbound port for forwarding the frame is port 1.

34 NETWORK ROUTING TABLE IN SWITCH Z.A

36 SWITCH/DEVICE	38 OUTBOUND PORT
B	1
C	2
D	2
N	4
P	5
Q	3
MANAGEMENT DEVICE	0

FIG. 2B




FIG. 2A

EX1005, Figs. 2B and 2A (highlighted). As shown in Figure 2A (above right), switch Z.A would transmit the frame using output port 1 thereby forwarding it to switch Z.B. Thus, Shimizu's *packet forwarding table* (network routing table 34) is *used for forwarding a packet* (frame) as recited in limitation 1a. Further, the routing table is *used for forwarding a packet based on an identifier* (switch/device value and/or outbound port value) *which identifies a link between nodes on a packet forwarding path* (link between switch Z.A and switch Z.B in my example).

151. Shimizu explains that “network routing table 34 for switch Z.A only provides routing information for other switches 22 in group Z, devices connected to switch Z.A, and management device 14.” EX1005, [0036]. In contrast, group routing table 40, which is also stored on switch Z.A, specifies the particular outbound port for frames destined for other groups. EX1005, [0039] (“[S]witch

Z.A maintains a table 40 providing routing information for group Y, as a whole, and group X, as a whole.”).

40

GROUP ROUTING IN SWITCH Z.A

GROUP	OUTBOUND PORT
X	0
Y	1

FIG. 2C




FIG. 2A

EX1005, Figs. 2C and 2A (highlighted).

152. Again, using Shimizu’s figures to illustrate, assume that a frame arrives at switch Z.A with header information indicating that the frame should be forwarded to switch X.B, which is in a different group 24 (group X). In that case, switch Z.A would use group routing table 40 illustrated in Figure 2.C (above left) to determine that the correct outbound port for forwarding the frame is port 0. As shown in Figure 2A (above right), switch Z.A would transmit the frame using output port 0 thereby forwarding it to switch X.B. Thus, Shimizu’s *packet forwarding table* (group routing table 40) is *used for forwarding a packet* (frame) as recited in limitation 1a. Further, the routing table is *used for forwarding a*

packet based on an identifier (group value and/or outbound port value) *which identifies a link between nodes on a packet forwarding path* (link between switch Z.A and switch X.B in my example).

153. Shimizu's routing tables (each *a packet forwarding table*) are stored on each switch 22 (*packet forwarding node*), as a person of ordinary skill would know. Switches store routing tables in order to use them to route network traffic. Indeed, Shimizu states that switch 22 "stores" the routing tables. EX1005, [0035] (Figure 2B "illustrates one example of how routing information can be **stored by switches 22**. In the illustrated embodiment, **switch Z.A stores** routing information as a network routing table 34."), see also [0019] ("One embodiment of selected routing information **stored by switches 22** is further explained below with reference to FIGS. 2A-C."), [0036] ("Routing information is illustrated and described as being **stored in table 34...**"), [0038] ("In particular embodiments, routing information **stored in switches 22** identifies a specific outbound port on which any and all network management messages are to be forwarded."), [0039] (Figure 2C "illustrates an example of group routing information that may be **stored** and maintained by switch Z.A."), [0053] ("switch 22 accesses routing information **stored on switch 22**"), [0058] ("In some embodiments, **switch 22 stores** routing information in a tabular form.").

154. Shimizu also discloses that the routing tables are stored in "memory"

in switches 22. EX1005, [0006] (“In some embodiments, appropriate routing can be enabled with smaller routing tables that require less **memory in a switch**.”); [0023] (“In another embodiment, vector based tags may reference **memory** locations, such as VLAN table entries, that indicate one or more outbound ports.”), [0045] (“As noted above, each vector based tag may include a bit vector, a reference to a bit vector stored **in memory**, or any other data structure suitable to indicate one or more outbound ports.”), 14 (claim 9: “**a memory** operable to maintain network routing information and standard routing information”).

155. A POSITA would understand Shimizu’s reference to “memory” in switches 22 to refer to *a non-transitory computer readable medium* such as RAM (Random Access Memory), SDRAM (Synchronous Dynamic Random Access Memory), or similar memory devices, which were commonly used in switches before 2009. See, e.g., EX1014, 3:22-23 (“FIG. 1 illustrates a layer 2 switch....”), 5:37-40 (“It should be understood that either of discovery list 110, forwarding CAM 106, and outgoing port LUT 108 can be any type of memory, including RAM, SDRAM, or the like.”), Fig. 1; EX1015, 3:19-23 (“FIG. 1 is a block diagram illustrating exemplary components of a layer 2 forwarding device....”), 3:29-33 (“FIG. 3 is a block diagram of exemplary components of an alternative implementation of layer 2 forwarding device....”), 3:47-49, 3:8-12 (“Exemplary computer-readable media suitable for implementing the methods and systems

described herein include chip memory devices, disk storage devices, such as optical or magnetic discs....”), Figs. 1 and 3; EX1017, 1:61-64 (“FIG. 1 illustrates a conventional multiport bridge architecture for network 10. Multiport bridge 8 is comprised of system bus 12 coupled to a system processor 11, memory 9, and two-port bridges 16, 17, and 18.”), Figs. 1 and 2; EX1018, 5:41-45 (“The Layer 2 switch includes...a memory 100-4, ...”), Fig. 3; EX1023, 4-5, Fig. 2; EX1024, 3.

156. A POSITA would have understood that Shimizu’s reference to “memory” could also indicate non-volatile memory such as a flash memory, NVRAM (Non-Volatile Random Access Memory), or similar devices, which would have provided the advantage of permitting faster initial configuration and restart. EX1016 (“A default configuration is stored in a Flash memory or other type of non-volatile memory...within switch 10 that ensures that switch 10 can be quickly connected to a network and can pass traffic with little or no user intervention, and preserving configuration in case of a power outage to switch 10.”). Again, such “memory” constitutes *a non-transitory computer readable medium*.

157. Further, the ’249 Patent describes using the same type of memory disclosed by Shimizu. The ’249 Patent states that “the recording unit 370, the recording unit 450, and the path information recording unit 540 can be realized by a device that can record information, such as a semiconductor memory, a hard disk

drive or the like.” EX1001, 14:35-38. Shimizu discloses storing routing tables (*packet forwarding table[s]*) using a semiconductor memory or similar storage device. EX1005, [0016], (switches 22 are “chip-based”). Therefore, Shimizu discloses *a storage unit including a non-transitory computer readable medium* in the same way that the ’249 Patent does.

158. Thus, Shimizu discloses switches 22 including *a storage unit including a non-transitory computer readable medium configured to store a packet forwarding table* as recited in the claim.

159. Further, as I discussed above, the routing tables in Shimizu are *used for forwarding a packet based on an identifier* as recited in the claim. Shimizu provides further support for my opinion in this regard. Shimizu states that a “device transmitting a frame may add tags to the frame in order to define a particular path that the frame should take through the network. For example, the tags may indicate a link-by-link path between switches in the network.” EX1005, Abstract. Shimizu states:


In the first mode of operation, a first tag in the frame is read. Whether the first tag’s switch identification corresponds to the switch is determined. If the switch identification does not correspond to the switch, the routing information is accessed to determine the output port corresponding to the switch identification, and the frame is transmitted on the determined one of the ports. If the switch identification corresponds to the switch, the frame is transmitted on

the port corresponding to the first tag's port identification.

EX1005, [0005]. Shimizu describes several examples of devices controlling the path of frames through the switch network using “tags” in the frames. EX1005, [0020] (“In operation, devices can control the path of frames through switch network 12 using tags in the frames.”), [0043] (“Tags 56 operate [to] control the path of frames through switch network 12.”), [0056] (“In the illustrated embodiment, port ID based tags are used to control the path of frames through switch network 12.”).

160. Thus, Shimizu describes a frame (*packet*) that is forwarded according to information in the tag. Shimizu's tags include the same types of information I discussed above regarding tables 34 and 40, such as switch/device value, group value, and outbound port value. The tags are another form of routing information and constitute *identifier[s]* for the same reasons.

161. Further, Shimizu's *identifier* used for forwarding the packet *identifies a link between nodes on a packet forwarding path*. For example, Figure 3A below “illustrate[s] a path of an outbound frame....” (i.e., path 50). EX1005, [0041].



EX1005, FIG. 3A.

162. As illustrated, “path 50 includes five different hops (1-5).” EX1005, [0042]. Shimizu states that “[i]n some embodiments, the frame is structured to include vector based tags. The structure of a frame with vector based tags is shown in FIG. 3B for each hop (1-5) of path 50. In particular embodiments, the frame is structured to include port ID based tags. The structure of a frame with port ID based tags is shown in FIG. 3C for each hop (1-5) of path 50.” EX1005, [0042]. In Shimizu, each “hop” along path 50 is *a link between nodes on a packet forwarding path* specified by *identifiers* used to obtain forwarding information in the routing tables (*packet forwarding tables*).

163. Thus, Shimizu discloses *a storage unit including a non-transitory computer readable medium configured to store a packet forwarding table used for forwarding a packet based on an identifier which identifies a link between nodes on a packet forwarding path* as recited in claim 1.

3. [1b] at least one processor configured to execute program instructions to provide a forwarding unit configured to:

164. Shimizu discloses this limitation.

165. Shimizu states that “[s]witches 22 represent networking devices that route frames within networks and/or between different networks” and that “each switch 22 represents a network device with multiple ports capable of switching network traffic among the ports.” EX1005, [0016]. “For example, switches 22

may each represent a chip-based switching device capable of switching frames among multiple ports based on standard Ethernet protocols.” EX1005, [0016]. A person of ordinary skill in the art would have known that Shimizu’s reference to a “chip-based switching device” means a switch with *at least one processor* used to perform the routing technique disclosed by Shimizu. A POSITA would certainly not think that Shimizu’s switches 22 operated without a processor because there is nothing in Shimizu to suggest such a thing.

166. A POSITA would have known that by 2009 switches and other routing devices, as disclosed in Shimizu, included *at least one processor*. See, e.g., EX1014, 3:22-23 (“FIG. 1 illustrates a layer 2 switch....”), 3:57-64 (“Switch 100 includes...a central processor (CPU) 112....”), Fig. 1; EX1015, 3:29-33 (“FIG. 3 is a block diagram of exemplary components of an alternative implementation of layer 2 forwarding device....”), 1:45-2:44, 6:23-26 (“each I/O module includes its own processor”), Fig. 3 (CPUs 337-339); EX1018, 5:41-45 (“The Layer 2 switch includes...a processor 100-3....”), Fig. 3; EX1023, 4-5, Fig. 2; EX1024, 3-4, Fig. 2. A POSITA would also have known that switches and other routing devices used a *processor to execute program instructions* to provide forwarding functionality (*at least one processor configured to execute program instructions to provide a forwarding unit*). See, e.g., EX1014, 2:32-60, 4:19-36, 5:42-52; EX1015, 1:45-58, 6:19-27; EX1019, 4:53-5:15, Fig. 2; EX1023, 4-5,

Fig. 2; EX1024, 3-4, Fig. 2.

167. Indeed, Shimizu recites a claim including “[l]ogic for routing frames, the logic encoded in media and operable when executed to” perform Shimizu’s routing techniques. EX1005, 9 (claim 16). A POSITA would understand Shimizu to disclose *program instructions* (logic) executed by *at least one processor* to provide Shimizu’s disclosed functionality.

168. Thus, Shimizu discloses and suggests to a POSITA that switches 22 include *at least one processor*. Further, a POSITA would have known that the routing techniques disclosed by Shimizu are executed using each switch’s *at least one processor* that is *configured to execute program instructions to provide a forwarding unit configured to* perform each of the functions recited in the following limitations for the reasons I discuss below.

4. [1c] receive an incoming packet including a plurality of the identifiers;

169. Shimizu discloses this limitation.


170. Shimizu’s switches 22 *receive an incoming packet*. Shimizu states that the disclosed “method for routing frames comprises **receiving frames** using a group of ports of a switch.” EX1005, [0005]; see also [0020] (“switch 22 will add an inbound tag only to **frames that switch 22 receives** from outside switch network 12”), [0021] (“each switch 22 **receiving an inbound frame**”), [0023] (“switch X.A may **receive a frame** with vector based outbound tags”; “[t]his frame

would then be **received by switch** Y.A”), [0024] (“switch Y.B may **receive a frame**”), [0035] (“switch Z.A **receives a frame**”), [0037] (“For example, switch Z.A may **receive a frame** with port ID based outbound tags.”), [0044] (“Switch X.A **receives the frame** and extracts tag 56a.”), [0047] (“Switch X.A **receives the frame** and extracts tag 60a.”), [0048] (same), [0049] (“Switch Y.C **receives the frame** and extracts tag 60b.”), [0050] (“Switch Z.C **receives the frame** and extracts tag 60c.”), [0053] (“Switch 22 **receives a frame** at step 82.”), [0056] (“Switch 22 **receives a frame** at step 102.”), Figs. 4, 5. Shimizu’s Figure 5 is “a flowchart illustrating a method of processing **a frame received at a switch** in the switch network.” EX1005, [0013], Fig. 5.

171. Further, Shimizu describes switches receiving *packet[s]* (frames) that *include a plurality of the identifiers* (tags). Shimizu states that “[e]ach frame includes a source address, a destination address, and a payload. At least some portion of the frames also include **one or more tags**, where each tag specifies a group identification, a switch identification, and a port identification.” EX1005, [0005]; see also [0020] (“In operation, devices can control the path of frames through switch network 12 **using tags in the frames.**”), [0022] (“Outbound **tags** may be used to determine the particular path a frame should take through switch network 12.”).

172. Shimizu’s Figure 3B illustrates the “structure of a frame with vector


based tags....” EX1005, [0042]. The illustrated “frame is structured with a vector based tag indicating the next switch 22 in path 50” (of Figure 3A). EX1005, [0043].



EX1005, Fig. 3B (annotated). For instance, during hop 1, the frame contains four tags 56 (i.e., 56a, 56b, 56c, and 56d). EX1005, [0044]. “Tags 56 operate [to] control the path of frames through switch network 12.” EX1005, [0043].

Shimizu’s tags 56 constitute *a plurality of identifiers*.

173. Shimizu’s Figure 3C illustrates the “structure of a frame with port ID based tags....” EX1005, [0042]. The illustrated “frame is structured with a port ID based tag indicating the next group 24 in path 50” (of Figure 3A). EX1005, [0046].



EX1005, Fig. 3C (annotated). For instance, during hop 1, the frame contains three tags 60 (i.e., 60a, 60b, and 60c). EX1005, [0047]. “Tags 60 operate [to] control the path of frames through switch network 12.” EX1005, [0046]. Shimizu’s tags 60 constitute *a plurality of identifiers*.

174. Thus, Shimizu’s switches 22 *receive an incoming packet including a plurality of the identifiers*.

5. **[1d] determine a direction to forward the incoming packet;**

175. Shimizu discloses this limitation.

176. Indeed, Shimizu expressly states that “[m]anagement device 14 may structure a frame so that various switches 22 in switch network 12 know **which particular direction** to route the frame.” EX1005, [0041]. Thus, Shimizu

expressly discloses that switches 22 *determine a direction to forward the incoming packet*.

177. Further, Shimizu describes switches 22 distinguishing between “inbound” frames and “outbound” frames. EX1005, [0020] (“In some embodiments, two types of tags exist: inbound tags and outbound tags.”). Inbound frames enter the switch network 12 from external network 20 or Internet 18 (i.e., from outside switch network 12). EX1005, [0020]. Switches 22 add inbound tags to inbound frames. EX1005, [0021]. Outbound frames originate in the network (e.g., from management device 14), and “[o]utbound tags may be used to determine the particular path a frame should take through switch network 12.” EX1005, [0022]. The switch uses the outbound tags to identify the next hop in the path through the network. EX1005, [0023]-[0024]. Thus, as a POSITA would have understood, switches 22 *determine a direction to forward the incoming packet* (as inbound or outbound) in order to correctly process the frame’s tags.

178. In addition, Shimizu discloses that switches 22 within switch network 12 “may provide for special handling of received management messages” to control routing them to the management device. EX1005, [0027]. In this example, switch 22 *determine[s] a direction to forward the incoming packet* toward (i.e., in the direction of) management device 14.


179. Thus, in several different ways, Shimizu discloses that switches 22 *determine a direction to forward the incoming packet.*

6. [1e] determine whether to forward to an external network;


180. Shimizu discloses this limitation.

181. Shimizu expressly identifies *an external network*. For example, Figure 1 illustrates switch network 12 connected to “an internet 18, and an external network 20.” EX1005, [0014]. Both internet 18 and external network 20 are *an external network* because they are outside (i.e., external to) switch network 12.

182. In Figure 1, switch Z.C is connected to external network 20 and switch Y.B is connected to internet 18.



EX1005, Fig. 1 (highlighted). Those same two switches are also connected to other switches within switch network 12 as illustrated in the highlighted figure below.



EX1005, Fig. 1 (highlighted). At least switches Z.C and Y.B distinguish between packets that are forwarded to *an external network* (internet 18 or external network 20) and a switch 22 within switch network 12. Thus, at least switches Z.C and Y.B *determine whether to forward to an external network* because, as a POSITA would know, those switches employing standard network protocols would need to process frames being forwarded to external networks differently than frames being forwarded to another switch in switch network 12.


183. As a POSITA would have known, *an external network* (like external network 20 or internet 18) is distinguished from switch network 12 because the external network either uses a different address space or uses a separate routing instance or both. And both of these would require the switch to be configured to determine that one or more specific ports on the switch are connected to an external network. For example, Shimizu states that internet 18 “may be any group of networks connected together, such as a world wide collection of computers and computer networks that can use TCP/IP network protocols to facilitate data transmission and exchange.” EX1005, [0030]. This means a switch 22 that is connected to internet 18 is configured to receive a frame from another switch 22 in switch network 12 and forward it using a different protocol (e.g., TCP/IP) to internet 18. The switch 22 would not forward the frame using this protocol if the frame were being forwarded to another switch 22 in switch network 12. Thus, the switch 22 *determine[s] whether to forward to an external network* or not.

184. Shimizu also states that “external network 20 may include one or more networks such as a telecommunications network, a satellite network, a cable network, local area networks (LANs), wide area networks (WANs), public switched telephone networks, and/or any other appropriate networks.” EX1005, [0030]. Again, any switch 22 that is connected to one of these types of external networks would receive a frame from another switch 22 in Shimizu’s switch

network 12 and forward it using the appropriate protocol and format for whatever specific type of external network will receive the frame. In order to ensure that the frame is forwarded using the appropriate protocol for the external network, the switch 22 determines that in fact it is forwarding to such an external network. Again, the switch 22 *determine[s] whether to forward to an external network* or not. This would have been obvious to a POSITA based on Shimizu's description. See, e.g., EX1005, [0030] ("While not expressly illustrated, external network 20 may separately employ a network based routing scheme, similar to that described for switching network 12, for communications within external network 20."), [0031] ("Although not illustrated, it is understood that internet 18 and external network 20 can each also be connected to other networks or other devices, components, or elements. While system 10 includes two networks, internet 18 and external network 20, any number and types of networks 16 may be connected to switch network 12."), [0033] ("Although not illustrated, external links 32 can connect constituent switches 22 to a variety of networks, components, elements, or devices, including, for example, other network devices 16, management device 14, internet 18, and external network 20.").


185. In addition, Shimizu describes "[o]utbound tags [that] may be used to determine the particular path a frame should take through switch network 12." EX1005, [0022]. Figure 3A illustrates path 50 of a frame through switch

network 12 and then to external network 20. EX1005, [0041].



In the illustrated example, switch Z.C receives the frame on hop 4 and then forwards it to external network 20 on hop 5.

186. Figure 3B “shows the corresponding structure of a frame through the various hops (1-5) of path 50” where the frame employs vector based tags. EX1005, [0043]. Switch Z.C receives the frame from hop 4 including tag 56d containing the value 3. EX1005, [0044].




EX1005, Fig. 3B (highlighted). “Switch Z.C removes tag 56d from the frame and sends the frame out port 3, after which it is received by external network 20.”

EX1005, [0044]. Thus, the switch *determine[s] whether to forward to an external network* based on the relevant tag value. In this example, switch Z.C determines that the frame is to be forwarded to external network 20 because tag 56d has the value of 3 indicating the switch’s port connected to the external network. If tag 56d had a different value, switch Z.C would make a different determination, e.g., that the frame is to be forwarded to another switch 22 in switch network 12 (i.e., switch Y.C, Z.B, or Z.D in Figure 1).

187. Figure 3C “shows the corresponding structure of a frame through the various hops (1-5) of path 50” where the frame employs port ID based tags.

EX1005, [0046]. This time switch Z.C receives the frame from hop 4 including tag 60c containing the value Z.C.3 (indicating group Z, switch C, port 3). EX1005, [0050].



EX1005, Fig. 3C (highlighted). “Switch Z.C removes tag 60c and sends the frame on port 3...” EX1005, [0050]. “The frame, during hop 5 contains destination address 52, source address 54, and payload 58.” EX1005, [0050]. “After hop 5, the frame is received by external network 20.” EX1005, [0050]. Thus, the switch *determine[s] whether to forward to an external network* based on the relevant tag value. In this example, switch Z.C determines that the frame is to be forwarded to external network 20 because tag 60c has the value of Z.C.3 indicating the switch’s port connected to the external network. If tag 60c had a different value, switch Z.C would make a different determination, e.g., that the frame is to be forwarded to another switch 22 in switch network 12 (i.e., switch Y.C, Z.B, or Z.D in Figure 1).

188. In both examples illustrated in Figures 3B and 3C, switch Z.C uses tags in the frame to *determine whether to forward to an external network*.

7. [1f] add to path or link information when the incoming packet is not forwarded to the external network;

189. Shimizu discloses this limitation.

190. Shimizu describes adding tags (*identifiers*) to a frame traveling in an inbound direction. In Shimizu, “[i]nbound tags may be used, for example, when a frame enters switch network 12 to ensure that responses to the frame are sent from switch network 12 on the same port that originally received the frame.” EX1005, [0020]. “When an inbound frame enters switch network 12 at a specific port, the receiving switch 22 may ***add an inbound tag*** indicating the group/switch/port that

received the frame.” EX1005, [0020]; see also [0006] (“Also, in particular embodiments, ***a tag can be added*** to a frame entering a private network to ensure that responses to the frame leave the private network on the exact same port.”). Additionally, in some embodiments, “each switch 22 receiving an inbound frame ***adds an inbound tag*** indicating the group/switch/port that received the frame.” EX1005, [0021]. “In particular embodiments, a receiving switch 22 ***adds an inbound tag*** that simply indicates the port on which the receiving switch 22 received the frame, and each subsequent switch 22 that receives the frame may ***add another inbound tag*** that indicates the port on which that switch 22 received the frame.” EX1005, [0021].

191. Inbound frames are—as described by Shimizu—frames that originate from an external network and route through switch network 12, so inbound tags are added by switches 22 as they forward inbound frames to other switches 22 (i.e., *when the incoming packet is not forwarded to the external network*).

192. Shimizu also describes adding tags to “network management messages” that are routed to management device 14 in switch network 12. Shimizu states that “[i]n some embodiments, management device 14 receives network management messages from network devices 16, internet 18, external network 20, and/or switches 22.” EX1005, [0027]. “To control routing of management messages to management device 14, switches 22 within switch

network 12 may provide for special handling of received management messages.”

EX1005, [0027]. “[O]ne or more of switches 22 **add inbound tags** to network management messages so that management device 14 knows the port the network management message entered switch network 12 and/or the precise route that each network management message took to reach the management device 14.”

EX1005, [0027]; see also [0054] (“[S]witch 22 determines whether switch 22 received the frame from an external link in step 92. If so, switch 22 **adds a tag** to the frame specifying the input port on which switch 22 received the frame at step 94.”), claims 4, 12, 19. These tags are added by switches 22 as they forward the management message to other switches 22 in switch network 12 and ultimately to management device 14 (i.e., *when the incoming packet is not forwarded to the external network*).

193. Thus, Shimizu’s switches 22 *add to path or link information (tags) when the incoming packet is not forwarded to the external network.*


8. [1g] remove path or link information header when the incoming packet is forwarded to the external network; and

194. Shimizu discloses this limitation.

195. Shimizu discloses that when the switch 22 at the edge of switch network 12 forwards an outbound frame to an *external network* (e.g., external network 20 or internet 18), the switch removes the last outbound tag and thereby *remove[s] path or link information header when the incoming packet is forwarded*

to the external network.


196. For example, Shimizu states (with regard to the frame illustrated in Figure 3B) that “[s]witch Z.C receives the frame and extracts the last tag 56: tag 56d” (i.e., value 3). EX1005, [0044]. “Switch Z.C **removes tag 56d** from the frame and sends the frame out port 3, after which it is received by external network 20.” EX1005, [0044].



EX1005, Fig. 3B (annotated). In hop 4, the outbound frame includes tag 56d (value 3). The switch removes tag 56d (*remove path or link information header*) and then forwards the frame in hop 5 to external network 20 (*when the incoming packet is forwarded to the external network*).

197. As another example, Shimizu states (with regard to the frame illustrated in Figure 3C) that “[s]witch Z.C receives the frame and extracts tag 60c.” EX1005, [0050]. The switch “**removes tag 60c** and sends the frame on port 3” so that “[t]he frame, during hop 5 contains destination address 52, source

address 54, and payload 58.” EX1005, [0050]. “After hop 5, the frame is received by external network 20.” EX1005, [0050].



EX1005, Fig. 3C (annotated). In hop 4, the outbound frame includes tag 60c (value Z.C.3). The switch removes tag 60c (*remove path or link information header*) and forwards the frame in hop 5 to external network 20 (*when the incoming packet is forwarded to the external network*).


9. [1h] forward the incoming packet by using a path or link that corresponds to the identifier from among the plurality of the identifiers.

198. Shimizu discloses this limitation.

199. Shimizu states that “[i]n operation, devices can control the path of frames through switch network 12 using tags in the frames.” EX1005, [0020]. Each tag is an *identifier*, and the frame can contain multiple tags or a *plurality of the identifiers*. The switch uses each tag (*the identifier from among the plurality of the identifiers*) to transmit the frame (*forward the incoming packet*) as indicated by

the tag (*by using a path or link that corresponds to the identifier*).


200. For example, Figure 3B illustrates the “structure of a frame through the various hops (1-5) of path 50” (shown in Figure 3A). EX1005, [0043]. The outbound frame “is structured with a vector based tag indicating the next switch 22 in path 50.” EX1005, [0043]. “During hop 1, the frame contains destination address 52, source address 54, four tags 56, and payload 58.” EX1005, [0044]. The “four tags 56” constitute *the plurality of the identifiers*.



EX1005, Fig. 3B (annotated).

201. The switches 22 use the four tags 56 to “control the path” of the frame through switch network 12. EX1005, [0043]. For example, in hop 1, “[s]witch X.A receives the frame and extracts tag 56a,” which has value 5. EX1005, [0044]. “Because tag 56a contains the value 5, switch X.A sends the frame on port 5 after

removing tag 56a from the frame” as shown in Figure 3A. EX1005, [0044].




EX1005, Fig. 3A (highlighted). In this example, switch X.A *forward[s] the incoming packet* received on hop 1 *by using a path or link* (hop 2 from switch X.A to switch X.C) *that corresponds to the identifier* (tag 56a) *from among the plurality of the identifiers* (tags 56a, 56b, 56c, 56d).

202. In hop 2, “the frame contains destination address 52, source address 54, three tags 56, and payload 58.” EX1005, [0044]. “After hop 2, switch X.C receives the frame and extracts tag 56b” (*the identifier from among the plurality of the identifiers*). EX1005, [0044]. In Figure 3B, “tag 56b contains the value 4.” EX1005, [0044]. “Switch X.C removes tag 56b from the frame and sends the frame out port 4” over hop 3. EX1005, [0044]. Thus, switch X.A *forward[s] the incoming packet* received on hop 2 *by using a path or link* (hop 3 from switch X.C to switch Y.C) *that corresponds to the identifier* (tag 56b) *from among the plurality of the identifiers* (tags 56b, 56c, 56d). The switches 22 forward the frame in a similar manner over hops 4 and 5. EX1005, [0044].

203. Shimizu describes a similar technique with reference to Figure 3C, which also illustrates the “structure of a frame through the various hops (1-5) of

path 50” (shown in Figure 3A). EX1005, [0046]. This outbound frame “is structured with a port ID based tag indicating the next group 24 in path 50.”

EX1005, [0046]. “During hop 1, the frame contains destination address 52, source address 54, three tags 60, and payload 58.” EX1005, [0047]. The “three tags 60” constitute *the plurality of the identifiers*.




EX1005, Fig. 3C (annotated).

204. The switches use the three tags 60 to “control the path” of the frame through switch network 12. EX1005, [0046]. For example, in hop 2, “[s]witch X.C receives the frame and extracts tag 60a,” which has value X.C.4. EX1005, [0048]. “Switch X.C first compares the group identification of tag 60a, ‘X,’ to its group identification” and “[f]inding there is a match, switch X.C then compares the

switch identification of tag 60a, 'C,' to its switch identification.” EX1005, [0048].

“Because there is a match, switch X.C evaluates the port identification, which is 4.” EX1005, [0048]. “Switch X.C first removes tag 60a, and then sends the frame on port 4 towards switch Y.C” over hop 3 in Figure 3A. EX1005, [0048].



EX1005, Fig. 3A (highlighted). In this example, switch X.C *forward[s] the incoming packet* received on hop 2 *by using a path or link* (hop 3 between switch X.C and switch Y.C) *that corresponds to the identifier* (tag 60a) *from among the plurality of the identifiers* (tags 60a, 60b, 6c). The switches 22 forward the frame in a similar manner over each hop in path 50. EX1005, [0047]-[0050].

205. Thus, Shimizu’s switches 22 *forward the incoming packet by using a path or link that corresponds to the identifier from among the plurality of the identifiers*.

B. Dependent Claims 2-5


- 1. Claim 2: The packet forwarding node according to claim 1, wherein the storage unit stores the table used for forwarding the packet based on the identifier which uniquely identifies the link within a predetermined domain.**

206. Shimizu discloses the limitations of claim 1 for the reasons I

explained for that claim. Shimizu also discloses the additional limitation recited in claim 2.

207. Shimizu discloses that *the storage unit stores the table used for forwarding the packet based on the identifier which...identifies the link* between nodes on a packet forwarding path for the reasons I explained for limitation 1a. See Section VII.A.2, above.

208. Shimizu also discloses that the identifier *uniquely identifies* the link as recited in claim 2. For example, Shimizu's Figure 2A illustrates the switches 22 in group Z. EX1005, [0033]. Shimizu explains that, as illustrated, "switch Z.A has six intra-group links 30 and external links 32, designated by ports 0 through 5." EX1005, [0034]. Each port is connected to another switch or device, as shown below.




EX1005, Fig. 2A; see also Fig. 1. Thus, each *identifier* indicating the appropriate outbound port *uniquely identifies the link* between nodes on a packet forwarding path.

209. Shimizu states that switch network 12 “may be organized into one or more groups of switches, with each group and switch within a group having ***unique identifiers*** so that any given switch or port in the network may be referred to by a hierarchical identification, such as a group/switch/port identification.”

EX1005, 1 (Abstract); see also [0006] (“In particular embodiments, switches within a private network are divided into groups, and a specific port of a specific one of the switches can be ***uniquely referenced by a tag*** that indicates a group, a switch, and a port identifier.”), [0018] (“each group 24 is assigned a ***unique identifier*** within switch network 12” and “each switch 22 in a particular group 24 is assigned a ***unique identifier*** within the particular group”), [0018] (“Accordingly, within switch network 12, a specific port may be ***uniquely referenced*** by use of a group/switch/port identifier.”), [0036] (“In particular embodiments, the division of switch network 12 into groups 24 of switches 22 allows any specific port in switch network 12 to be ***uniquely referenced*** by use of a group/switch/port identifier.”).

210. Additionally, Shimizu discloses that the identifier identifies *the link within a predetermined domain* as recited in claim 2. Indeed, Shimizu does so in two different ways.

211. First, Shimizu discloses switch network 12, which is *a predetermined domain*. “Switch network 12 includes switches 22 and represents any suitable collection of hardware, software, and controlling logic that interconnect switches 22 and elements coupled to switch network 12.” EX1005, [0015].




EX1005, Fig. 1 (annotated).

212. “In some embodiments, switch network 12 is a private network that allows devices to specifically determine the path of a frame through the network.” EX1005, [0015]; see also [0004] (“[T]hese techniques describe a routing scheme and network organization that enable a private network to control the path [of] frames take through the private network”), [0006] (“[T]hese techniques may allow a device to specify a particular path that a frame should take th[r]ough a private

network.”).

213. Shimizu states that “within a private network, the switches can use private network routing protocols in place of Ethernet routing for selected network traffic.” EX1005, [0006]. The private switch network 12 is connected to, but distinct from, internet 18 and external network 20 either of which may employ different routing techniques. EX1005, [0030]. A person of ordinary skill in the art would have recognized Shimizu’s private switch network 12 as *a predetermined domain* set up and administered by a particular entity.

214. Second, Shimizu discloses groups 24 within switch network 12. Groups 24 are each *a predetermined domain*. Shimizu states that in some embodiments “switches within a private network are divided into groups, and a specific port of a specific one of the switches can be uniquely referenced by a tag that indicates a group, a switch, and a port identifier.” EX1005, [0006].



EX1005, Fig. 1 (annotated). “This can permit switches within a group to maintain private routing tables defined for only other switches in the group.” EX1005, [0006]. A POSITA would have recognized that each of Shimizu’s groups 24 within private switch network 12 is *a predetermined domain* set up and administered by the entity responsible for private switch network 12.

2. **Claim 3: The packet forwarding node according to claim 1, wherein the storage unit stores the table used for forwarding the packet based on the identifier which uniquely identifies the link within a predetermined local domain.**

215. Shimizu discloses the limitations of claim 1 for the reasons I explained for that claim. Shimizu also discloses the additional limitation recited in claim 3.

216. Shimizu discloses that *the storage unit stores the table used for*

forwarding the packet based on the identifier which...identifies the link between nodes on a packet forwarding path for the reasons I explained for limitation 1a. See Section VII.A.2, above. The identifier *uniquely identifies* the link for the reasons I explained for claim 2. See Section VII.B.1, above.

217. Additionally, Shimizu discloses that the identifier identifies *the link within a predetermined local domain* as recited in claim 3. As discussed for claim 2, Shimizu discloses switch network 12 including groups 24. Both switch network 12 and each of groups 24 constitute a *predetermined...domain* as I explained for claim 2. See Section VII.B.1, above.

218. Shimizu refers to switch network 12 as a “private” network. EX1005, [0004] (“[T]hese techniques describe a routing scheme and network organization that enable a **private network** to control the path frames take through the private network.”), [0006] (“[T]hese techniques may allow a device to specify a particular path that a frame should take th[r]ough **a private network**.”), [0015] (“In some embodiments, **switch network 12 is a private network** that allows devices to specifically determine the path of a frame through the network.”). Shimizu also discloses “private network routing protocols” for routing frames through the private network. EX1005, [0006] (“[W]ithin a private network, the switches can use **private network routing protocols** in place of Ethernet routing for selected network traffic.”).

219. A POSITA would have known that Shimizu's private switch network 12 with private routing protocol would commonly be implemented as a *local* network such as a LAN (Local Area Network). Thus, such a person would understand Shimizu to disclose that the identifier identifies *the link within a predetermined local domain* as recited in claim 3.

3. Claim 4: The packet forwarding node according to claim [1], wherein the storage unit stores the table used for forwarding the packet based on the identifier which uniquely identifies the link within the packet forwarding node.

220. Shimizu discloses the limitations of claim 1 for the reasons I explained for that claim. Shimizu also discloses the additional limitation recited in claim 4.

221. Shimizu discloses that *the storage unit stores the table used for forwarding the packet based on the identifier which...identifies the link* between nodes on a packet forwarding path for the reasons I explained for limitation 1a. See Section VII.A.2, above. The identifier *uniquely identifies* the link for the reasons I explained for claim 2. See Section VII.B.1, above.

222. Additionally, Shimizu discloses that the identifier identifies *the link within the packet forwarding node* as recited in claim 4. Shimizu describes switches 22 using an *identifier* to forward a frame using a particular outbound port on the switch as I explained for limitation 1h. See Section VII.A.9, above. Thus,

Shimizu discloses using an identifier to specify *the link within the packet forwarding node* (i.e., the link by which the frame is transferred from the inbound port to the outbound port of the switch).

4. Claim 5: The packet forwarding node according to claim 1, wherein the storage unit stores the table used for forwarding the packet based on the identifier which unique identifies the link among the packet forwarding nodes.

223. Shimizu discloses the limitations of claim 1 for the reasons I explained for that claim. Shimizu also discloses the additional limitation recited in claim 5.

224. Shimizu discloses that *the storage unit stores the table used for forwarding the packet based on the identifier which...identifies the link* between nodes on a packet forwarding path for the reasons I explained for limitation 1a. See Section VII.A.2, above. The identifier *unique[ly] identifies* the link for the reasons I explained for claim 2. See Section VII.B.1, above. The identifier identifies *the link among the packet forwarding nodes* for the reasons I explained for limitation 1a. See Section VII.A.2, above. There is no material difference for purposes of my analysis between the phrase in limitation 1a (“identifies a link between nodes on a packet forwarding path”) and the similar phrase in claim 5 (“identifies the link among the packet forwarding nodes”). Shimizu discloses claim 5 for the reasons I explained for limitation 1a and claim 2.

C. Independent Claim 6

1. [6pre] A packet forwarding node, comprising:

225. The preamble of claim 6 is identical to the preamble of claim 1. See Paragraph 91, above (comparing claims 1 and 6). The preamble of claim 6 is disclosed for the reasons I discussed for 1pre. See Section VII.A.1, above.

2. [6a] a storage unit including a non-transitory computer readable medium configured to store a packet forwarding table used for forwarding a packet based on an identifier which identifies an interface of a packet forwarding node; and

226. This limitation of claim 6 is quite similar to the corresponding limitation in claim 1 (limitation 1a). Claim 1 refers to an identifier that “identifies a link between nodes on a packet forwarding path,” while this limitation of claim 6 refers to an identifier that “identifies an interface of a packet forwarding node.” See Paragraph 91, above (comparing claims 1 and 6).

227. As I explained for claim 1, Shimizu’s *identifier* specifies an outbound port on switch 22 for *forwarding a packet*. See Sections VII.A.1 (1pre), VII.A.2 (1a), and VII.A.9 (1h), above. Each outbound port designates a link between the switch and another network node. EX1005, [0034]. Each outbound port is also *an interface of a packet forwarding node*—it is the interface between the switch and other network nodes (e.g., switches 22, management device 14, etc.) and external networks (e.g., internet 18, external network 20, etc.). Thus, this limitation of claim 6 is disclosed for the reasons I discussed for limitation 1a. See

Section VII.A.2, above.

3. [6b] at least one processor configured to execute program instructions to provide a forwarding unit configured to

228. This limitation of claim 6 is identical to the corresponding limitation in claim 1 (limitation 1b). See Paragraph 91, above (comparing claims 1 and 6). This limitation of claim 6 is disclosed for the reasons I discussed for limitation 1b. See Section VII.A.3, above.

4. [6c] receive an incoming packet including a plurality of identifiers,

229. This limitation of claim 6 is virtually identical to the corresponding limitation in claim 1 (limitation 1c). See Paragraph 91, above (comparing claims 1 and 6). Limitation 1c refers to “a plurality of the identifiers” while this limitation refers to “a plurality of identifiers.” The minor difference is immaterial for my analysis. This limitation of claim 6 is disclosed for the reasons I discussed for limitation 1c. See Section VII.A.4, above.

5. [6d] determine a direction to forward the incoming packet,

230. This limitation of claim 6 is identical to the corresponding limitation in claim 1 (limitation 1d). See Paragraph 91, above (comparing claims 1 and 6). This limitation of claim 6 is disclosed for the reasons I discussed for limitation 1d. See Section VII.A.5, above.

6. [6e] determine whether to forward to an external network,

231. This limitation of claim 6 is identical to the corresponding limitation

in claim 1 (limitation 1e). See Paragraph 91, above (comparing claims 1 and 6). This limitation of claim 6 is disclosed for the reasons I discussed for limitation 1e. See Section VII.A.6, above.

7. [6f] add to path or link information when the incoming packet is not forwarded to the external network,

232. This limitation of claim 6 is identical to the corresponding limitation in claim 1 (limitation 1f). See Paragraph 91, above (comparing claims 1 and 6). This limitation of claim 6 is disclosed for the reasons I discussed for limitation 1f. See Section VII.A.7, above.

8. [6g] remove path or link information header when the incoming packet is forwarded to the external network, and

233. This limitation of claim 6 is identical to the corresponding limitation in claim 1 (limitation 1g). See Paragraph 91, above (comparing claims 1 and 6). This limitation of claim 6 is disclosed for the reasons I discussed for limitation 1g. See Section VII.A.8, above.

9. [6h] to forward the incoming packet by using a path or link that corresponds to the identifier from among the plurality of identifiers.

234. This limitation of claim 6 is virtually identical to the corresponding limitation in claim 1 (limitation 1h). See Paragraph 91, above (comparing claims 1 and 6). Limitation 1h recites “forward the incoming packet,” while this limitation recites “to forward the incoming packet.” Limitation 1h refers to “the plurality of the identifiers” while this limitation refers to “the plurality of identifiers.” These

minor differences are immaterial for my analysis. This limitation of claim 6 is disclosed for the reasons I discussed for limitation 1h. See Section VII.A.9, above.

D. Dependent Claims 7-10

- 1. Claim 7: The packet forwarding node according to claim 6, wherein the storage unit stores the table used for forwarding the packet based on the identifier which uniquely identifies the link within a predetermined domain.**

235. Shimizu discloses the limitations of claim 6 for the reasons I explained for that claim. Shimizu also discloses the additional limitation recited in claim 7.

236. The limitation recited in claim 7 is identical to the limitation recited in claim 2. See Paragraph 95, above (comparing claims 2 and 7). Claim 7 is disclosed for the reasons I discussed for claim 2. See Section VII.B.1, above.

- 2. Claim 8: The packet forwarding node according to claim 6, wherein the storage unit stores the table used for forwarding the packet based on the identifier which uniquely identifies the link within a predetermined local domain.**

237. Shimizu discloses the limitations of claim 6 for the reasons I explained for that claim. Shimizu also discloses the additional limitation recited in claim 8.

238. The limitation recited in claim 8 is identical to the limitation recited in claim 3. See Paragraph 98, above (comparing claims 3 and 8). Claim 8 is disclosed for the reasons I discussed for claim 3. See Section VII.B.2, above.

3. **Claim 9: The packet forwarding node according to claim 6, wherein the storage unit stores the table used for forwarding the packet based on the identifier which uniquely identifies the link within the packet forwarding node.**

239. Shimizu discloses the limitations of claim 6 for the reasons I explained for that claim. Shimizu also discloses the additional limitation recited in claim 9.

240. The limitation recited in claim 9 is identical to the limitation recited in claim 4. See Paragraph 101, above (comparing claims 4 and 9). Claim 9 is disclosed for the reasons I discussed for claim 4. See Section VII.B.3, above.

4. **Claim 10: The packet forwarding node according to claim 6, wherein the storage unit stores the table used for forwarding the packet based on the identifier which uniquely identifies the link among the packet forwarding nodes.**

241. Shimizu discloses the limitations of claim 6 for the reasons I explained for that claim. Shimizu also discloses the additional limitation recited in claim 10.

242. The limitation recited in claim 10 is identical to the limitation recited in claim 5. See Paragraph 104, above (comparing claims 5 and 10). Claim 10 is disclosed for the reasons I discussed for claim 5. See Section VII.B.4, above.

E. Independent Claim 11

1. **A packet forwarding node, comprising:**

243. The preamble of claim 11 is identical to the preamble of claim 1. See

Paragraph 92, above (comparing claims 1 and 11). The preamble of claim 11 is disclosed for the reasons I discussed for 1pre. See Section VII.A.1, above.

2. a memory storing a program;

244. Shimizu discloses this limitation.

245. As I explained for limitation 1a, Shimizu's switches 22 contain memory devices such as RAM, SDRAM, NVRAM, etc. See Section VII.A.2, above. Switches 22 also include at least one processor as I explained for limitation 1b. See Section VII.A.3, above. As person of ordinary skill in the art would have known that the switch processor's fundamental purpose is to execute program instructions stored in memory. Such a person would have known that the routing technique disclosed by Shimizu is performed by one or more processors on switch 22 executing program instructions stored in the switch's memory. Thus, Shimizu discloses to a POSITA that each switch 22 (*packet forwarding node*) includes *memory storing a program* to perform Shimizu's routing techniques.

3. at least one processor configured to execute the program to provide:

246. As I explained for the prior limitation, Shimizu's switches 22 contain memory devices and at least one processor as I have discussed above. See Sections VII.A.2 and VII.A.3, above. A POSITA would have known that the routing techniques disclosed by Shimizu are performed by one or more processors on switches 22 executing program instructions stored in the switch's memory.

Thus, Shimizu discloses to a POSITA that each switch 22 (*packet forwarding node*) includes *at least one processor configured to execute the program* in memory to perform Shimizu's routing techniques. The program performs the routing technique corresponding to the following limitations of claim 11 as I discuss below.

247. Further, this limitation of claim 11 is virtually identical to the corresponding limitation in claim 1 (limitation 1b). See Paragraph 92, above (comparing claims 1 and 11). Limitation 1b refers to "program instructions" while this limitation refers to "the program." The minor difference is immaterial for my analysis. This limitation of claim 11 is disclosed for the reasons I discussed for limitation 1b. See Section VII.A.3, above.

4. a communication unit configured to communicate with a server for receiving a plurality of identifiers, each of which identifies a link between nodes on a packet forwarding path; and

248. Shimizu discloses this limitation.

249. Shimizu explains that management device 14 can be used to configure switches 22 in switch network 12. EX1005, [0026] ("Management device 14 may control the operation of switch network 12 and/or the interoperation between different switches 22 within switch network 12. In particular embodiments, management device 14 performs an initial configuration of switch network 12."). A person of ordinary skill in the art would understand that this configuration

includes loading routing tables on switches 22 in switch network 12. The routing tables include *a plurality of identifiers, each of which identifies a link between nodes on a packet forwarding path* as I explained for limitations 1a and 1c. See Sections VII.A.2 and VII.A.4, above. Shimizu's management device 14 constitutes *a server sending a plurality of identifiers* to switches 22.

250. Further, Shimizu's "management device 14 and network devices 16 can transmit frames with outbound tags in order to define a desired frame path." EX1005, [0022]. The outbound tags correspond to *a plurality of identifiers, each of which identifies a link between nodes on a packet forwarding path* as I explained for limitations 1c and 1h. See Sections VII.A.4 and VII.A.9, above. Like management device 14, a network device 16 can be *a server*. EX1005, [0029] ("Network devices 16 represent any device connectable to switch network 12. For example, network devices 16 may include personal computers, servers, file storage, network printers, and other network enabled devices."). Thus, Shimizu's management device 14 and network devices 16 disclose *a server sending a plurality of identifiers* to switches 22.

251. The fact that Shimizu discloses *a server sending a plurality of identifiers* to switches 22 informs a person of ordinary skill in the art that Shimizu's *packet forwarding node[s]* (switches 22) contain *a communication unit configured to communicate with a server for receiving a plurality of identifiers*.

Switches 22 could not receive configuration and other information from management device 14 and network devices 16 without such *a communication unit*.

5. a forwarding unit configured to:

252. The limitation is part of language in claim 11 that corresponds to limitation 1b. That limitation of claim 1 recites “at least one processor configured to execute program instructions to provide a forwarding unit configured to...” Combining this limitation with the language of claim 11 addressed in Section VII.E.3, above, recites “at least one processor configured to execute the program to provide...a forwarding unit configured to...” See Paragraph 92, above (comparing claims 1 and 11). The minor difference in wording is immaterial for my analysis. This limitation of claim 11 is disclosed for the reasons I discussed for limitation 1b. See Section VII.A.3, above.

6. receive an incoming packet including the plurality of the identifiers,

253. This limitation of claim 11 is virtually identical to the corresponding limitation in claim 1 (limitation 1c). Paragraph 92, above (comparing claims 1 and 11). Limitation 1c refers to “a plurality” of identifiers while this limitation refers to “the plurality.” The minor difference is immaterial for my analysis. This limitation of claim 11 is disclosed for the reasons I discussed for limitation 1c. See Section VII.A.4, above.

7. determine a direction to forward the incoming packet,

254. This limitation of claim 11 is identical to the corresponding limitation in claim 1 (limitation 1d). Paragraph 92, above (comparing claims 1 and 11). This limitation of claim 11 is disclosed for the reasons I discussed for limitation 1d. See Section VII.A.5, above.

8. determine whether to forward to an external network,

255. This limitation of claim 11 is identical to the corresponding limitation in claim 1 (limitation 1e). Paragraph 92, above (comparing claims 1 and 11). This limitation of claim 11 is disclosed for the reasons I discussed for limitation 1e. See Section VII.A.6, above.

9. add to path or link information when the incoming packet is not forwarded to the external network,

256. This limitation of claim 11 is identical to the corresponding limitation in claim 1 (limitation 1f). Paragraph 92, above (comparing claims 1 and 11). This limitation of claim 11 is disclosed for the reasons I discussed for limitation 1f. See Section VII.A.7, above.

10. remove path or link information header when the incoming packet is forwarded to the external network, and

257. This limitation of claim 11 is identical to the corresponding limitation in claim 1 (limitation 1g). Paragraph 92, above (comparing claims 1 and 11). This limitation of claim 11 is disclosed for the reasons I discussed for limitation 1g. See Section VII.A.8, above.

- 11. forward the incoming packet by using a path or link that corresponds to an identifier from among the plurality of identifiers.**

258. This limitation of claim 11 is virtually identical to the corresponding limitation in claim 1 (limitation 1h). Paragraph 92, above (comparing claims 1 and 11). Limitation 1h refers to “the identifier from among the plurality of the identifiers,” while this limitation of claim 11 refers to “an identifier from among the plurality of identifiers.” The minor difference is immaterial for my analysis. This limitation of claim 11 is disclosed for the reasons I discussed for limitation 1h. See Section VII.A.9, above.

F. Dependent Claims 12-15

- 1. Claim 12: The packet forwarding node according to claim 11, wherein the forwarding unit forwards the packet based on the identifier which uniquely identifies the link within a predetermined domain.**

259. Shimizu discloses the limitations of claim 11 for the reasons I explained for that claim. Shimizu also discloses the additional limitation recited in claim 12.

260. Shimizu discloses a *forwarding unit* that *forwards the packet based on the identifier which...identifies the link* between nodes on a packet forwarding path for the reasons I explained for limitations 1b-1h. See Sections VII.A.3-VII.A.9, above. The identifier *uniquely identifies* the link for the reasons I explained for claim 2. See Section VII.B.1, above. Similarly, the identified *link* is *within a*

predetermined domain for the reasons I explained for claim 2. See Section VII.B.1, above; see also Paragraph 96, above (comparing claims 2 and 12).

2. Claim 13: The packet forwarding node according to claim 11, wherein the forwarding unit forwards the packet based on the identifier which uniquely identifies the link within a predetermined local domain.

261. Shimizu discloses the limitations of claim 11 for the reasons I explained for that claim. Shimizu also discloses the additional limitation recited in claim 13.

262. Shimizu discloses a *forwarding unit* that *forwards the packet based on the identifier which...identifies the link* between nodes on a packet forwarding path for the reasons I explained for limitations 1b-1h. See Sections VII.A.3-VII.A.9, above. The identifier *uniquely identifies* the link for the reasons I explained for claim 2. See Section VII.B.1, above. The identified *link* is *within a predetermined local domain* for the reasons I explained for claim 3, which recites the same phrase. See Section VII.B.2, above; see also Paragraph 99, above (comparing claims 3 and 13).

3. Claim 14: The packet forwarding node according to claim 11, wherein the forwarding unit forwards the packet based on the identifier which uniquely identifies the link within the packet forwarding node.

263. Shimizu discloses the limitations of claim 11 for the reasons I explained for that claim. Shimizu also discloses the additional limitation recited in

claim 14.

264. Shimizu discloses a *forwarding unit* that *forwards the packet based on the identifier* for the reasons I explained for limitations 1b-1h. See Sections VII.A.3-VII.A.9, above. The *identifier...uniquely identifies the link within the packet forwarding node* for the reasons I explained for claim 4, which recites the same phrase. See Section VII.B.3, above; see also Paragraph 102, above (comparing claims 4 and 14).

4. Claim 15: The packet forwarding node according to claim 11, wherein the forwarding unit forwards the packet based on the identifier which uniquely identifies the link among the packet forwarding nodes.

265. Shimizu discloses the limitations of claim 11 for the reasons I explained for that claim. Shimizu also discloses the additional limitation recited in claim 15.

266. Shimizu discloses a *forwarding unit* that *forwards the packet based on the identifier which...identifies the link* between nodes on a packet forwarding path for the reasons I explained for limitations 1b-1h. See Sections VII.A.3-VII.A.9, above. The identifier *uniquely identifies* the link for the reasons I explained for claim 2. See Section VII.B.1, above. Similarly, the identified *link* is *among the packet forwarding nodes* for the reasons I explained for claim 5, which recites the same phrase. See Section VII.B.4, above; see also Paragraph 105, above (comparing claims 5 and 15).

G. Independent Claim 16

1. A packet forwarding node, comprising:

267. The preamble of claim 16 is identical to the preamble of claim 1. See Paragraph 93, above (comparing claims 1 and 16). The preamble of claim 16 is disclosed for the reasons I discussed for 1pre. See Section VII.A.1, above.

2. a communication unit configured to communicate with a server for receiving a plurality of identifiers, each of which identifies an interface of the packet forwarding node; and

268. Shimizu discloses this limitation, which is similar to the “communication unit” limitation of claim 11. Shimizu discloses *a communication unit configured to communicate with a server for receiving a plurality of identifiers* for the reasons I explained for that corresponding limitation of claim 11. See Section VII.E.4, above. The only difference between these limitations is that claim 11 refers to an identifier that “identifies a link between nodes on a packet forwarding path,” while this limitation of claim 16 refers to an identifier that *identifies an interface of the packet forwarding node*. As I explained for claim 6 (limitation 6a), Shimizu discloses identifiers that identify *an interface of the packet forwarding node* (as well as a link between nodes). See Section VII.C.2, above.

3. at least one processor configured to execute program instructions to provide a forwarding unit configured to

269. This limitation of claim 16 is identical to the corresponding limitation in claim 1 (limitation 1b). Paragraph 93, above (comparing claims 1 and 16). This

limitation of claim 16 is disclosed for the reasons I discussed for limitation 1b. See Section VII.A.3, above.

4. receive an incoming packet including a plurality of identifiers,

270. This limitation of claim 16 is identical to the corresponding limitation in claim 1 (limitation 1c). Paragraph 93, above (comparing claims 1 and 16). This limitation of claim 16 is disclosed for the reasons I discussed for limitation 1c. See Section VII.A.4, above.

5. determine a direction to forward the incoming packet,

271. This limitation of claim 16 is identical to the corresponding limitation in claim 1 (limitation 1d). Paragraph 93, above (comparing claims 1 and 16). This limitation of claim 16 is disclosed for the reasons I discussed for limitation 1d. See Section VII.A.5, above.

6. determine whether to forward to an external network,

272. This limitation of claim 16 is identical to the corresponding limitation in claim 1 (limitation 1e). Paragraph 93, above (comparing claims 1 and 16). This limitation of claim 16 is disclosed for the reasons I discussed for limitation 1e. See Section VII.A.6, above.

7. add to path or link information when the incoming packet is not forwarded to the external network,

273. This limitation of claim 16 is identical to the corresponding limitation in claim 1 (limitation 1f). Paragraph 93, above (comparing claims 1 and 16). This

limitation of claim 16 is disclosed for the reasons I discussed for limitation 1f. See Section VII.A.7, above.

8. remove path or link information header when the incoming packet is forwarded to the external network, and

274. This limitation of claim 16 is identical to the corresponding limitation in claim 1 (limitation 1g). Paragraph 93, above (comparing claims 1 and 16). This limitation of claim 16 is disclosed for the reasons I discussed for limitation 1g. See Section VII.A.8, above.

9. forward the incoming packet by using a path or link that corresponds to an identifier from among the plurality of identifiers.

275. This limitation of claim 16 is virtually identical to the corresponding limitation in claim 1 (limitation 1h). Paragraph 93, above (comparing claims 1 and 16). Limitation 1h refers “the identifier from among the plurality of the identifiers,” while this limitation refers to “an identifier from among the plurality of identifiers.” The minor difference is immaterial for my analysis. This limitation of claim 16 is disclosed for the reasons I discussed for limitation 1h. See Section VII.A.9, above.

H. Dependent Claims 17-20

1. Claim 17: The packet forwarding node according to claim 16, wherein the forwarding unit forwards the packet based on the identifier which uniquely identifies the link within a predetermined domain.

276. Shimizu discloses the limitations of claim 16 for the reasons I

explained for that claim. Shimizu also discloses the additional limitation recited in claim 17.

277. The limitation recited in claim 17 is identical to the limitation recited in claim 12. See Paragraph 97, above (comparing claims 12 and 17). Claim 17 is disclosed for the reasons I discussed for claim 12. See Section VII.F.1, above.

2. Claim 18: The packet forwarding node according to claim 16, wherein the forwarding unit forwards the packet based on the identifier which uniquely identifies the link within a predetermined local domain.

278. Shimizu discloses the limitations of claim 16 for the reasons I explained for that claim. Shimizu also discloses the additional limitation recited in claim 18.

279. The limitation recited in claim 18 is identical to the limitation recited in claim 13. See Paragraph 100, above (comparing claims 13 and 18). Claim 18 is disclosed for the reasons I discussed for claim 13. See Section VII.F.2, above.

3. Claim 19: The packet forwarding node according to claim 16, wherein the forwarding unit forwards the packet based on the identifier which uniquely identifies the link within the packet forwarding node.

280. Shimizu discloses the limitations of claim 16 for the reasons I explained for that claim. Shimizu also discloses the additional limitation recited in claim 19.

281. The limitation recited in claim 19 is identical to the limitation recited

in claim 14. See Paragraph 103, above (comparing claims 14 and 19). Claim 19 is disclosed for the reasons I discussed for claim 14. See Section VII.F.3, above.

4. Claim 20: The packet forwarding node according to claim 16, wherein the forwarding unit forwards the packet based on the identifier which uniquely identifies the link among the packet forwarding nodes.

282. Shimizu discloses the limitations of claim 16 for the reasons I explained for that claim. Shimizu also discloses the additional limitation recited in claim 20.

283. The limitation recited in claim 20 is identical to the limitation recited in claim 15. See Paragraph 106, above (comparing claims 15 and 20). Claim 20 is disclosed for the reasons I discussed for claim 15. See Section VII.F.4, above.

I. Independent Claim 21

1. A server, comprising:

284. Shimizu explains that “[m]anagement device 14 may control the operation of switch network 12 and/or the interoperation between different switches 22 within switch network 12.” EX1005, [0026]. For instance, “management device 14 performs an initial configuration of switch network 12. The initial configuration separates switches 22 into groups 24, assigns a switch identification to each switch 22, assigns a group identification to each group 24, and distributes to each switch 22 the switch identification and the group identification corresponding to that particular switch.” EX1005, [0026]. A

POSITA would have understood from Shimizu's description of the functions performed by management device 14 that it is a *server* computer. Further, Shimizu expressly states that network device 16 can be a *server*. EX1005, [0029] ("Network devices 16 represent any device connectable to switch network 12. For example, network devices 16 may include personal computers, **servers**, file storage, network printers, and other network enabled devices."). Thus, Shimizu's management device 14 and network devices 16 disclose *[a] server* as recited in the preamble of claim 21.

285. Thus, Shimizu's management device 14 and network devices 16 each disclose the claimed *server*.

2. a memory storing a computer program;

286. A person of ordinary skill in the art would have known that a server computer includes memory (e.g., RAM, SDRAM, disk drive, etc.) and that one of the most basic operations of a server involves a processor executing *a computer program* stored in *memory* on the computer. EX1020, 4:10-5:12; EX1021, 1 (Abstract), 1:23-2:13, 8:63-9:34; EX1022, 4:22-6:24.

3. at least one processor configured to execute the computer program to provide:

287. A person of ordinary skill in the art would have known that a server computer includes *at least one processor* and that one of the most basic operations of a server involves a processor executing *a computer program*. A POSITA would

know that the functions performed by Shimizu's management device 14 and network devices 16 (*server*) require a processor on the device *to execute [a] computer program*. EX1020, 4:10-5:12; EX1021, 1 (Abstract), 1:23-2:13, 8:63-9:34; EX1022, 4:22-6:24.

4. a first unit configured to determine a plurality of identifiers, each of which identifies a link between nodes on a packet forwarding path; and

288. Shimizu explains that management device 14 can be used to configure switches 22 in switch network 12. EX1005, [0026] (“Management device 14 may control the operation of switch network 12 and/or the interoperation between different switches 22 within switch network 12. In particular embodiments, management device 14 performs an initial configuration of switch network 12.”). A person of ordinary skill in the art would understand that this configuration includes loading routing tables on switches 22 in switch network 12. EX1005, [0026]. The routing tables include *a plurality of identifiers, each of which identifies a link between nodes on a packet forwarding path* as I explained for limitations 1a and 1c. See Sections VII.A.2 and VII.A.4, above. Shimizu's management device 14 constitutes a *server* sending *a plurality of identifiers* to switches 22. A POSITA would understand that in order for management device 14 to configure routing tables on switches 22, it would include *a first unit configured to determine a plurality of identifiers* (i.e., it would first determine the *plurality of*

identifiers and then configure routing tables on switches 22 using those *identifiers*).

289. In addition, Shimizu’s “management device 14 and network devices 16 can transmit frames with outbound tags in order to define a desired frame path.” EX1005, [0022]. The outbound tags correspond to *a plurality of identifiers, each of which identifies a link between nodes on a packet forwarding path* as I explained for limitations 1c and 1h. See Sections VII.A.4 and VII.A.9, above. Shimizu’s management device 14 and network devices 16 each disclose *a server sending a plurality of identifiers* to switches 22. A POSITA would understand that in order for such a device to send *a plurality of identifiers*, it would include *a first unit configured to determine a plurality of identifiers* (i.e., it would first determine the *plurality of identifiers* to include as outbound tags in the frame).

5. a second unit configured send information elements to a packet forwarding node which is capable of

290. Shimizu’s management device 14 configures the switches 22 in switch network 12. EX1005, [0026]. A person of ordinary skill in the art would understand that this configuration includes loading routing tables on switches 22 in switch network 12. EX1005, [0026]. The routing tables include *information elements* including identifiers specifying links between nodes in the packet forwarding path as I explained for limitation 1a. See Section VII.A.2, above. Thus, Shimizu’s management device 14 constitutes a *server sending information elements* to switches 22. A POSITA would understand that in order for

management device 14 to send *information elements* to switches 22, it would include *a second unit configured [to] send information elements to a packet forwarding node* (switch 22).

291. In addition, Shimizu’s “management device 14 and network devices 16 can transmit frames with outbound tags in order to define a desired frame path.” EX1005, [0022]. The outbound tags are *information elements* corresponding to identifiers identifying a path through the network as I explained for limitations 1c and 1h. See Sections VII.A.4 and VII.A.9, above. Shimizu’s management device 14 and network devices 16 each disclose *a server* sending *information elements* to switches 22. A POSITA would understand that in order for such a device to send *information elements* to switches 22, it would include *a second unit configured [to] send information elements to a packet forwarding node* (switch 22).

292. Shimizu’s *packet forwarding node* (switch 22) *is capable of* performing the following steps recited in claim 21 for the reasons I explained for claim 1 (limitations 1c-1h). See Section VII.A.4-VII.A.9, above.

6. receiving an incoming packet including a plurality of identifiers,

293. This limitation of claim 21 is virtually identical to the corresponding limitation in claim 1 (limitation 1c). Paragraph 94, above (comparing claims 1 and 21). Limitation 1c recites “receive an incoming packet,” while this limitation

refers to “receiving an incoming packet.” The minor difference is immaterial for my analysis. Shimizu discloses *a packet forwarding node which is capable of...receiving an incoming packet including a plurality of identifiers*, as recited in claim 21, for the reasons I discussed for limitation 1c. See Section VII.A.4, above.

7. determining a direction to forward the incoming packet,

294. This limitation of claim 21 is virtually identical to the corresponding limitation in claim 1 (limitation 1d). Paragraph 94, above (comparing claims 1 and 21). Limitation 1d recites “determine a direction,” while this limitation refers to “determining a direction.” The minor difference is immaterial for my analysis. Shimizu discloses *a packet forwarding node which is capable of...determining a direction to forward the incoming packet*, as recited in claim 21, for the reasons I discussed for limitation 1d. See Section VII.A.5, above.

8. determining whether to forward to an external network,

295. This limitation of claim 21 is virtually identical to the corresponding limitation in claim 1 (limitation 1e). Paragraph 94, above (comparing claims 1 and 21). Limitation 1e recites “determine whether to forward,” while this limitation refers to “determining whether to forward.” The minor difference is immaterial for my analysis. Shimizu discloses *a packet forwarding node which is capable of...determining whether to forward to an external network*, as recited in claim 21, for the reasons I discussed for limitation 1e. See Section VII.A.6, above.

9. adding to path or link information when the incoming packet is not forwarded to the external network,

296. This limitation of claim 21 is virtually identical to the corresponding limitation in claim 1 (limitation 1f). Paragraph 94, above (comparing claims 1 and 21). Limitation 1f recites “add to path or link information,” while this limitation refers to “adding to path or link information.” The minor difference is immaterial for my analysis. Shimizu discloses *a packet forwarding node which is capable of...adding to path or link information when the incoming packet is not forwarded to the external network*, as recited in claim 21, for the reasons I discussed for limitation 1f. See Section VII.A.7, above.

10. removing path or link information header when the incoming packet is forwarded to the external network, and

297. This limitation of claim 21 is virtually identical to the corresponding limitation in claim 1 (limitation 1g). Paragraph 94, above (comparing claims 1 and 21). Limitation 1g recites “remove path or link information header,” while this limitation refers to “removing path or link information header.” The minor difference is immaterial for my analysis. Shimizu discloses *a packet forwarding node which is capable of...removing path or link information header when the incoming packet is forwarded to the external network*, as recited in claim 21, for the reasons I discussed for limitation 1g. See Section VII.A.8, above.

11. forwarding the incoming packet by using a path or link, that corresponds to an identifier from among the plurality of identifiers.

298. This limitation of claim 21 is virtually identical to the corresponding limitation in claim 1 (limitation 1h). Paragraph 94, above (comparing claims 1 and 21). Limitation 1h recites “forward the incoming packet,” while this limitation refers to “forwarding the incoming packet.” Also, limitation 1h states that the path or link “corresponds to the identifier from among the plurality of the identifiers,” while this limitation states that the path or link “corresponds to an identifier from among the plurality of identifiers.” These minor differences are immaterial for my analysis. Shimizu discloses *a packet forwarding node which is capable of...forwarding the incoming packet by using a path or link, that corresponds to an identifier from among the plurality of identifiers*, as recited in claim 21, for the reasons I discussed for limitation 1h. See Section VII.A.9, above.

J. Dependent Claims 22-27

1. Claim 22: The packet forwarding node according to claim 1, wherein the storage unit stores the table used for forwarding the packet based on the identifier which identifies the link between adjacent nodes on a packet forwarding path.

299. Shimizu discloses the limitations of claim 1 for the reasons I explained for that claim. Shimizu also discloses the additional limitation recited in claim 22.

300. Shimizu discloses a *storage unit that stores the table used for*

forwarding the packet based on the identifier which identifies the link between...nodes on a packet forwarding path for the reasons I explained for limitation 1a. See Section VII.A.2, above. Shimizu also discloses that the identified *link* is *between adjacent nodes* as recited in claim 22. For example, Figure 1 illustrates switch network 12 with links connecting *adjacent nodes* (switches 22). EX1005, Fig. 1. Similarly, Figure 2A illustrates a group 24 within switch network 12 showing links *between adjacent nodes* (switches 22). EX1005, Fig. 2A.

2. **Claim 23: The packet forwarding node according to claim 6, wherein the storage unit stores the packet forwarding table used for forwarding the packet based on the identifier which identifies the link between adjacent nodes on a packet forwarding path.**

301. Shimizu discloses the limitations of claim 6 for the reasons I explained for that claim. Shimizu also discloses the additional limitation recited in claim 23.

302. Shimizu discloses a *storage unit* that *stores the packet forwarding table used for forwarding the packet based on the identifier which identifies the link between...nodes on a packet forwarding path* for the reasons I explained for limitation 1a. See Section VII.A.2, above. Shimizu also discloses that the identified *link* is *between adjacent nodes* for the reasons I discussed for claim 22. See Section VII.J.1, above; see also Paragraph 107, above (comparing the language

of claims 22 and 23).

3. **Claim 24: The packet forwarding node according to claim 11, further comprising a storage unit that stores a table used for forwarding the packet based on an identifier which identifies the link between adjacent nodes on a packet forwarding path.**

303. Shimizu discloses the limitations of claim 11 for the reasons I explained for that claim. Shimizu also discloses the additional limitation recited in claim 24.

304. Shimizu discloses *a storage unit that stores a table used for forwarding the packet based on an identifier which identifies the link between...nodes on a packet forwarding path* for the reasons I explained for limitation 1a. See Section VII.A.2, above. Shimizu also discloses that the identified *link* is *between adjacent nodes* for the reasons I discussed for claim 22. See Section VII.J.1, above; see also Paragraph 108, above (comparing the language of claims 22 and 24).

4. **Claim 25: The packet forwarding node according to claim 16, further comprising a storage unit that stores a table used for forwarding the packet based on an identifier which identifies the link between adjacent nodes on a packet forwarding path.**

305. Shimizu discloses the limitations of claim 16 for the reasons I explained for that claim. Shimizu also discloses the additional limitation recited in claim 25.

306. Shimizu discloses *a storage unit that stores a table used for forwarding the packet based on an identifier which identifies the link between...nodes on a packet forwarding path* for the reasons I explained for limitation 1a. See Section VII.A.2, above. Shimizu also discloses that the identified *link* is *between adjacent nodes* for the reasons I discussed for claim 22. See Section VII.J.1, above; see also Paragraph 109, above (comparing the language of claims 22 and 25).

5. Claim 26: The server according to claim 21,

307. Shimizu discloses the limitations of claim 21 for the reasons I explained for that claim. Shimizu also discloses the additional limitation recited in claim 26.

- a. further comprising a storage unit to store a packet forwarding table used for forwarding the incoming packet based on the identifier which identifies the link between nodes on the packet forwarding path,**

308. This limitation of claim 26 is essentially the same as limitation 1a. This limitation of claim 26 is disclosed for the reasons I discussed for limitation 1a. See Section VII.A.2, above.

- b. wherein the storage unit stores the table used for forwarding the packet based on the identifier which identifies a link between adjacent nodes on a packet forwarding path.**

309. Shimizu discloses *a storage unit that stores the table used for forwarding the packet based on the identifier which identifies a link*

between...nodes on a packet forwarding path for the reasons I explained for limitation 1a. See Section VII.A.2, above. Shimizu also discloses that the identified *link* is *between adjacent nodes* for the reasons I discussed for claim 22. See Section VII.J.1, above.

6. **Claim 27: The packet forwarding node according to claim 1, wherein when the forwarding unit receives the incoming packet including the identifier for the link connecting nodes on the packet forwarding path, the forwarding unit forwards the incoming packet based on the identifier included the received incoming packet from among the plurality of the identifiers.**

310. Shimizu discloses the limitations of claim 1 for the reasons I explained for that claim. Shimizu also discloses the additional limitation recited in claim 27.

311. Shimizu discloses that *the forwarding unit receives the incoming packet including the identifier for the link connecting nodes on the packet forwarding path* for the reasons I explained for limitation 1a, 1b, and 1c. See Sections VII.A.2-VII.A.4, above.

312. Shimizu discloses that *the forwarding unit forwards the incoming packet based on the identifier included [in] the received incoming packet from among the plurality of the identifiers* for the reasons I explained for limitation 1h. See Section VII.A.9, above.

313. Further, Shimizu discloses that *when the forwarding unit receives the*

incoming packet... the forwarding unit forwards the incoming packet... For example, Shimizu states:

[S]witch X.A may receive a frame with vector based outbound tags. The first tag identifies port 4. In order to transmit the frame to the appropriate next switch 22, switch X.A removes the first tag and sends the frame on port 4. This frame would then be received by switch Y.A, and switch Y.A repeats a similar process in order to transmit the frame to the appropriate next switch 22.

EX1005, [0023]. In this example, *when the forwarding unit of switch X.A receives the incoming packet (frame) including the identifier for the link connecting nodes on the packet forwarding path (vector based outbound tags), the forwarding unit forwards the incoming packet based on the identifier included [in] the received incoming packet from among the plurality of the identifiers (tag identifying port 4).*

VIII. THE CHALLENGED CLAIMS ARE RENDERED OBVIOUS BY THE COMBINATION OF SHIMIZU AND RFC-4271

314. Every independent claim of the '249 Patent recites a “packet forwarding node” that “determin[es] whether to forward to an external network.” EX1001, 30:25 (claim 1), 30:60-61 (claim 6), 31:30-31 (claim 11), 31:65-66 (claim 16), 32:35-36 (claim 21).

315. As I explained above for limitation 1e, Shimizu discloses and renders obvious to a POSITA the limitation requiring that the *forwarding node* be

configured to *determine whether to forward to an external network*. In addition, Shimizu can be combined with the Border Gateway Protocol (BGP-4) standard set out in RFC-4271 (EX1006), which expressly teaches configuring switches to determine whether a packet is being forwarded to an external network.

A. Motivation to Combine Shimizu with RFC-4271

316. A person of ordinary skill in the art would have recognized that Shimizu's switch network 12 can be an autonomous system (AS) whose routing function is controlled by a single organization as I discussed above in the Technical Background. See Section IV.B, above. Shimizu's switch network 12 is connected to internet 18 and external network 20, which both are or include other ASs using different policies and protocols. Shimizu's switch network 12 would employ an inter-AS routing protocol to communicate with these external networks. EX1005, [0030], [0033].

317. By 2009, BGP4 was the standard inter-AS routing protocol that would have been used for a switch network such as the one disclosed by Shimizu. As stated in "Computer Networking," which was published in 2008, "[t]he Border Gateway Protocol version 4, specified in RFC 4271...is the *de facto* standard inter-AS routing protocol in today's Internet." EX1010, 178; see also 169 ("in the Internet all ASs run the same inter-AS routing protocol, called BGP4"), 179 ("BGP is an absolutely critical protocol for the Internet—in essence, it is the protocol that

glues the whole thing together”). A person of ordinary skill in the art would have recognized the benefit of employing the BGP4 protocol because it was the industry standard that permitted one AS to communicate with another. As stated in “Computer Networking,” BGP was “an absolutely critical protocol for the Internet—in essence, it is the protocol that glues the whole thing together....” EX1010, 179. A person of ordinary skill in the art would have been motivated to implement Shimizu’s switches 22 with the BGP4 protocol so that Shimizu’s switch network would communicate with internet 18 and external network 20 using the industry standard inter-AS routing protocol.

318. A POSITA would have known that because Shimizu’s switch network 12 connects to internet 18 and external network 20, those connections would be advantageously implemented using the BGP4 protocol of RFC-4271. Further, a POSITA would have implemented the BGP4 protocol in Shimizu’s switch network 12 because it was a well-known industry standard for the Internet. A POSITA would have been aware of BGP4 and would likely have experience studying and working with it. Indeed, many switches by 2009 would have come pre-configured to comply with the BGP4 protocol making its implementation quite simple for a POSITA. Thus, a POSITA could have implemented Shimizu’s switch network employing BGP4, as described in RFC-4271, without need for undue experimentation and with a reasonable likelihood of success.

B. Claim Limitation: [determine/determining] whether to forward to an external network;

319. RFC-4271 describes how routing works in a network that uses BGP4. It shows that network nodes maintain knowledge of destinations that are accessible via external networks. The routing device, which RFC-4271 refers to as a “BGP speaker,” includes a Routing Information Base (RIB) that stores routing information consisting of three distinct parts. EX1006, 10.

320. RFC-4271 describes the three distinct parts of the Routing Information Base on the device as follows:

- a) Adj-RIBs-In: The Adj-RIBs-In stores routing information learned from inbound UPDATE messages that were received from other BGP speakers. Their contents represent routes that are available as input to the Decision Process.
- b) The Loc-RIB contains the local routing information the BGP speaker selected by applying its local policies to the routing information contained in its Adj-RIBs-In. These are the routes that will be used by the local BGP speaker. The next hop for each of these routes MUST be resolved via the local BGP speaker’s Routing Table.
- c) Adj-RIBs-Out: The Adj-RIBs-Out stores information the local BGP speaker selected for advertisement to its peers. The routing information stored in the Adj-RIBs-Out will be carried in the local BGP speaker’s UPDATE message and advertised to its peers.

EX1006, 10. RFC-4271 states that, “[i]n summary, the Adj-RIBs-In contains unprocessed routing information that has been advertised to the local BGP speaker

by its peers; the Loc-RIB contains the routes that have been selected by the local BGP speaker's Decision Process; and the Adj-RIBs-Out organizes the routes for advertisement to specific peers (by means of the local speaker's UPDATE messages)." EX1006, 10-11.

321. RFC-4271 further states that "[r]outing information that the BGP speaker uses to forward packets (or to construct the forwarding table used for packet forwarding) is maintained in the Routing Table. The Routing Table accumulates routes to directly connect networks, static routes, routes learned from the IGP protocols, and routes learned from BGP." EX1006, 11.

322. Thus, when a routing device in a BGP-enabled network receives a packet whose destination address is part of an external network (i.e., a router that was learned from an inbound UPDATE message), the routing device distinguishes it from other packets not destined for an external network and uses BGP routing data to forward the packet toward the external network. When the routing device is an edge router (e.g., switch Z.C in Shimizu's switch network 12), the device forwards the packet over the link to the external network.

IX. SECONDARY CONSIDERATIONS

323. I have considered the impact of secondary considerations of obviousness, but I am not aware of any relevant information related to secondary considerations of obviousness that would impact my opinions at this time. To the

extent the Patent Owner comes forward with any such information, I will consider it.

X. AVAILABILITY FOR CROSS-EXAMINATION

324. In signing this declaration, I recognize that the declaration will be filed as evidence in a contested case 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 of me, I will appear for cross-examination within the United States during a time as mutually agreed between the parties.

XI. RIGHT TO SUPPLEMENT

325. I reserve the right to supplement my opinions in the future to respond to any arguments that the Patent Owner raises. This declaration represents only those opinions that I have formed to date. I reserve the right to revise, supplement, and/or amend my opinions stated herein based on new information that becomes available to me and on my continuing analysis of the materials already provided.

XII. CONCLUSION

326. For the reasons discussed in this declaration, it is my opinion that the challenged claims of the '249 Patent are unpatentable in view of the prior art I have discussed above.

Acknowledgement of Oath

I 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 1101 of Title 18 of the United States Code.

Respectfully submitted,

A handwritten signature in black ink, appearing to read "Mark Crovella". The signature is written in a cursive, flowing style.

Mark Crovella, Ph.D.

Dated: July 9, 2025