

**IN THE UNITED STATES DISTRICT COURT
FOR THE EASTERN DISTRICT OF TEXAS
MARSHALL DIVISION**

OptimNet LLC,

Plaintiff,

v.

Cisco Systems, Inc.,

Defendant.

CIVIL ACTION NO. 2:25-cv-00935-JRG

JURY TRIAL DEMANDED

**DEFENDANT CISCO SYSTEMS, INC.'S INVALIDITY CONTENTIONS
AND SUBJECT MATTER ELIGIBILITY CONTENTIONS**

I. INTRODUCTION

Pursuant to the Court’s First Amended Docket Control Order (Dkt. 29), Local Patent Rules 3-3 and 3-4, and the Court’s Standing Order Regarding Subject Matter Eligibility Contentions dated July 25, 2019, Defendant Cisco Systems, Inc. (“Defendant” or “Cisco”) provide their Invalidity Contentions and Subject Matter Eligibility Contentions demonstrating the invalidity and ineligibility of the patent claims asserted by Plaintiff OptimNet LLC (“Plaintiff” or “OptimNet”) in their Infringement Contentions served on December 3, 2025. The Asserted Claims are as follows:

Asserted Patent	Asserted Claims
U.S. Pat. No. 8,116,623	1-27
U.S. Pat. No. 8,638,807	13-16
U.S. Pat. No. 8,780,887	1-10
U.S. Pat. No. 9,313,101	1-6
U.S. Pat. No. 10,020,961	1-13

Defendant reserves the right to serve amended or supplemental invalidity contentions in the event that Plaintiff amends its Infringement Contentions and/or based on new information learned during the course of discovery, which is currently ongoing.

II. RESERVATIONS

A. General Reservations

Consistent with the Local Patent Rule 3-6 and the Court’s Standing Order Regarding Subject Matter Eligibility Contentions, Defendant reserves the right to amend these Invalidity Contentions and Subject Matter Eligibility Contentions. The information and documents that Defendant is producing herewith are provisional and subject to further revision. Defendant reserves the right to amend or supplement these disclosures and the associated document production should Plaintiff: 1) provide any information that they failed to provide in their Infringement Contentions and related disclosures; 2) amend their Infringement Contentions or

related disclosures in any way; or 3) attempt to rely upon any information at trial, in a hearing, or during a deposition which they failed to provide in their Infringement Contentions or related disclosures.

The information provided shall not be deemed an admission regarding the scope of any claims or the proper construction of those claims or any terms contained therein. Nothing contained in these Invalidity Contentions and Subject Matter Eligibility Contentions should be understood or deemed to be an express or implied admission or contention with respect to the proper construction of any terms in the Asserted Claims, or with respect to the alleged infringement of those claims.

B. Forthcoming Discovery

Furthermore, because fact discovery is still in its early stages in this case, Defendant reserves the right to revise, amend, and/or supplement the information provided herein, including by identifying, charting, and relying on additional references, should Defendant's further search and analysis yield additional information or references, consistent with this Court's Docket Control Order (Dkt. 29), the Local Rules (including the Local Patent Rules), and the Federal Rules of Civil Procedure. Defendant reserves the right to present additional items of prior art located during the course of discovery or via further investigation.

Furthermore, Defendant may seek documents from third parties in discovery to demonstrate such third parties' earlier dates of invention under pre-AIA 35 U.S.C. § 102(g) and/or prior art offers for sale and public uses of products embodying the claimed invention under pre-AIA 35 U.S.C. § 102(a)-(b) or post-AIA 35 U.S.C. § 102(a). Defendant therefore reserves the right to supplement these Invalidity Contentions with further information and documentation, if and when it becomes available.

Plaintiff has not yet produced any invalidity contentions from any prior or concurrent matters involving the Asserted Patents. Defendant incorporates by reference, and intends to rely on, each reference, disclosure, or argument served in any prior or concurrent matter involving the Asserted Patents but not yet produced to Defendant by Plaintiff.

C. Claim Constructions

Defendant reserves the right to revise their ultimate contentions concerning the invalidity and/or ineligibility of the Asserted Claims, which may change depending upon any findings as to the priority date of those claims, and/or positions that Plaintiff or expert witnesses may take concerning infringement and/or invalidity/eligibility issues. Defendant does not waive the right to contest any claim constructions or to take positions during claim construction proceedings that have yet to occur that may be inconsistent with the Invalidity Contentions or Eligibility Contentions herein. Consequently, Defendant also reserves the right to amend or supplement these Invalidity Contentions and/or Eligibility Contentions in the event that a claim is construed differently at some point in the future.

Defendant's Invalidity Contentions and Eligibility Contentions are not adoptions or admissions by Defendant as to the accuracy of Plaintiff's allegations, admissions, or positions, or of any courts' prior claim constructions. Accordingly, these contentions are made in the alternative, are not necessarily intended to be consistent with each other, and should not be otherwise construed. Defendant expressly reserves the right to take positions with respect to future claim construction or infringement issues that are inconsistent with, or even contradictory to, the claim construction or infringement positions expressed or implied in the Invalidity Contentions and Eligibility Contentions set forth herein.

D. Plaintiff's Infringement Contentions

Plaintiff's Infringement Contentions and related disclosures are deficient in numerous respects. For example, Plaintiff has failed to specifically identify where each element of each Asserted Claim is found within each purported Accused Product. Defendant specifically reserves the right to modify, amend, or supplement these contentions as Plaintiff modifies, amends, or supplements its Infringement Contentions and related disclosures, and/or produces documents in discovery.

E. The Intrinsic Record

Defendant further reserves the right to rely upon applicable industry standards and prior art cited in the file histories of the Asserted Patents and any related U.S. and foreign patent applications as invalidating references or to show the state of the art. Defendant further intends to rely on the patent applicants' admissions concerning the scope of the prior art relevant to the Asserted Patents found in, inter alia: the patent prosecution histories for the Asserted Patents and any related patents and/or patent applications, reissue applications or reexaminations; any deposition testimony of the named patent applicants on the Asserted Patents; and the papers filed and any evidence submitted by Plaintiff in connection with this litigation as well as prior litigation and PTAB proceedings or other USPTO proceedings.

F. Contextual Evidence

Defendant's claim charts cite to particular teachings and disclosures of the prior art as applied to features of the Asserted Claims. However, persons having ordinary skill in the art generally may view an item of prior art in the context of his or her experience and training, other publications, literature, products, and understanding. As such, the cited portions are only examples, and Defendant reserves the right to rely on uncited portions of the prior art references, other publications and expert testimony, and the knowledge of those of ordinary skill in the art as aids

in understanding and interpreting the cited portions, as providing context thereto, and as additional evidence that the prior art discloses a claim limitation or the claimed subject matter as a whole. Defendant further reserves the right to rely on uncited portions of the prior art references, other publications, and testimony, including expert testimony, and the knowledge of those of ordinary skill in the art to establish bases for combinations of certain cited references that render the Asserted Claims obvious. The references discussed in the claim charts attached hereto may disclose the elements of the Asserted Claims explicitly and/or inherently, and/or they may be relied upon to show the state of the art in the relevant time frame. The suggested obviousness combinations are provided in the alternative to anticipation contentions and are not to be construed to suggest that any reference included in the combinations is not by itself anticipatory, nor are they admissions that obviousness is not to be tested against the prior art as a whole.

G. Priority Dates

Defendant reserves the right to present additional items of prior art or further support for current prior art should discovery or further investigation show that the Asserted Claims are not entitled to the priority dates Plaintiff identifies in the Infringement Contentions. For example, discovery may show that one or more of the Asserted Claims of the Asserted Patents are not entitled to the benefit of filing dates of foreign priority applications Plaintiff asserts in its Infringement Contentions:

<u>U.S. Patent No.</u>	<u>Priority Date</u>
8,116,623	6/25/2007
8,638,807	6/24/2009 ¹
8,780,887	10/29/2010

¹ Defendant believes Plaintiff intended 7/24/2009, as there is no 6/24/2009 priority claim on the face of the '807 Patent.

<u>U.S. Patent No.</u>	<u>Priority Date</u>
9,313,101	4/18/2012
10,020,961	12/27/2013

H. Incorporation by Reference

Any pending or future requests for *Ex Parte* Reexamination or petitions for *Inter Partes* Review of one or more of the Asserted Patents are herein incorporated by reference in their entirety, including all of the arguments, prior art, and other evidence disclosed therein.

III. INVALIDITY

Based on Plaintiff's Infringement Contentions, Defendant identify the following items of prior art that anticipate and/or render obvious the Asserted Claims of the Asserted Patents. The referenced claim charts are attached to these contentions at Exhibits A-1 to A-5, B-1 to B-4, C-1 to C-11, D-1 to D-10 and E-1 to E-11. The patents, patent applications, and publications disclosed herein are prior art under at least pre-AIA 35 U.S.C. §§ 102(a), (b), (e), (f), or (g) (for the '623, '807, '887, and '101 Patents) or post-AIA 35 U.S.C. § 102(a) (for the '961 Patent). The AIA status of the patents is based on Defendant's current understanding of the effective priority date that Plaintiff asserts each patent is entitled and as noted above, discovery may show otherwise.

The claim charts attached hereto identify where each element of the Asserted Claims is found in the given prior art. Defendant reserves the right to revise their claim charts to rely on any of these references to prove the invalidity of the Asserted Claims in a manner consistent with the Federal Rules of Civil Procedure, the Court's Local Rules (including the Local Patent Rules), and this Court's Orders.

A. '623 Patent

1. Prior Art Patents and Publications

Reference	Issue/Publication Date
U.S. Patent No. 6,333,799 (“Bala”)	December 25, 2001
U.S. Patent No. 7,711,270 (“Ooi”)	May 4, 2010
U.S. Patent No. 5,724,167 (“Sabella”)	March 3, 1998
U.S. Patent Application Publication No. 2006/0115210 (“Nakagawa”)	June 1, 2006
U.S. Patent No. 7,715,713 (“Iyer”)	May 11, 2010

To the extent the Asserted Claims of the '623 Patent are not anticipated by the prior art references identified above and in the claim charts attached as Exhibits A-1 to A-5, the Asserted Claims are rendered obvious by those references, either alone or in combination with other prior art references as identified below and in the attached claim charts and/or the knowledge of one skilled in the art.

In *KSR International Co. v. Teleflex Inc.*, 500 U.S. 398 (2007), the U.S. Supreme Court rejected the Federal Circuit’s rigid “teaching, suggestion, or motivation” requirement in favor of a flexible, functional approach in which an explicit finding of a “motivation” to combine prior art references is not required to establish obviousness. The Supreme Court held that it is sufficient that a combination of elements was “obvious to try” holding that, “[w]hen there is a design need or market pressure to solve a problem and there are a finite number of identified, predictable solutions, a person of ordinary skill has good reason to pursue the known option within his or her technical grasp.” *Id.* at 402; *see also Dystar Textilfarben GmbH v. C.H. Patrick Co.*, 464 F.3d 1356, 1368 (Fed. Cir. 2006) (explaining that when the “combination of references results in a product or process that is more desirable, for example because it is stronger, cheaper, cleaner,

faster, lighter, smaller, more durable, or more efficient,” there exists a motivation to combine prior art references even when there is no explicit suggestion in the references themselves “[b]ecause the desire to enhance commercial opportunities by improving a product or process is universal—and even commonsensical”); *LeapFrog Enters., Inc. v. Fisher-Price, Inc.*, 485 F.3d 1157 (Fed. Cir. 2007) (applying *KSR* and holding that “one of ordinary skill in the art of children’s learning toys would have found it obvious to combine the Bevan device with the SSR to update it using modern electronic components in order to gain the commonly understood benefits of such adaption, such as decreased size, increased reliability, simplified operation, and reduced cost”).

Defendant contends that where the Asserted Claims are anticipated by a particular reference they are also rendered obvious by that same reference alone, or in combination with the other references, as show in the table below:

Primary Reference	In Combination with One or More of
Bala	Ooi
Bala	Sabella
Bala	Nakagawa
Bala	Iyer
Ooi	Bala
Ooi	Sabella
Ooi	Nakagawa
Ooi	Iyer
Sabella	Bala
Sabella	Ooi
Sabella	Nakagawa
Sabella	Iyer
Nakagawa	Bala
Nakagawa	Ooi
Nakagawa	Sabella
Nakagawa	Iyer

Defendant further contends that one of ordinary skill in the art, at the time of the alleged invention of the Asserted Claims, would have been motivated to combine the references disclosed herein in such a way to reach the alleged inventions. The teaching, suggestion, or motivation to

combine these references, although not required, is explicitly or implicitly found in one or more of the following: the knowledge or common sense of one of ordinary skill in the art; the prior art references themselves and/or the prior art as a whole, including interrelated teachings of multiple prior art references; the subject matter acknowledged as prior art in the '623 Patent; the nature of the problem to be solved and the existence of similar improvements in similar applications; design incentives and other market forces, including the advantages of creating a superior and more desirable product and the effects of demands known to the design community or present in the marketplace; the ability to implement the alleged invention as a predictable variation of the prior art; improvements in similar devices; the interrelated teachings of multiple prior art references; any needs or problems known in the field addressed by the '623 Patent; and the number of identified, predictable solutions to the problem allegedly addressed by these patents. In addition, the simultaneous (and/or prior) inventions described above, and elsewhere in these contentions, is evidence that motivation to combine the concepts described in the various prior art references did, in fact, exist, and they were, in fact, combined. The combinations evidencing teachings, suggestions, and/or motivations to combine the prior art references in a way that renders the Asserted Claim obvious are merely exemplary. A person of ordinary skill would have access to the materials found in the cited references and attached exhibits and would have at least the ordinary creativity and skill to combine the cited references in ways not explicitly recited herein.

Further, Defendant hereby identifies additional motivation and reasons to combine the cited art: one of skill in the art would have been motivated to combine these references, because these references relate to common objectives and subject matter. The identified prior art references relate to common objectives and share commonalities in terms of their general subject matter as well as the types of equipment, products, and approaches used, including their descriptions of ring-

based optical transport networks and associated equipment and methods, including optical add/drop multiplexers (OADMs), wavelength-selective switches (WSSs), optical cross-connects (OXC), and hybrid optical/electrical cross-connects used to interconnect multiple rings, perform per-wavelength switching, provide wavelength conversion when needed, and implement protection switching, which are incorporated by reference herein as knowledge of a person of ordinary skill in the art (POSITA). *See, e.g.*, Ooi '270 at 9:16–37, 10:3–10, 10:42–43, FIG. 8; Sabella '167 at 3:3–30, 6:57–7:3, 8:62–9:10; Bala '799 at 8:49–57, 9:8–15, FIGs. 3, 7; Nakagawa '210 at ¶[0016], ¶[0018], ¶[0075]–[0076]; and Iyer '713 at 6:21–28, 11:61–67, 6:61–7:14, 15:43–49, 9:11–19, 9:29–58.

Further, the Asserted Claims of the '623 Patent contain merely a duplication of concepts known in the prior art and having no patentable significance given that no unexpected results occur (i.e., the patentee merely combines known prior art elements according to known methods to yield predictable results). For example, the '623 Patent itself admits that: (a) Figures 4 and 5 depict prior art structures and methods, (b) conventional inter-ring transmission relied on optical-electrical-optical (O/E/O) conversion at hub nodes, (c) conventional implementations of multi-casting via electrical switches caused bottlenecks and increased costs, (d) optical-only multi-casting led to wavelength collisions and signal degradation, (e) ring topologies, 50 ms restoration, and 1:1 and 1+1 protection schemes were already known, and (f) wavelength selective switches (WSS) were known devices with defined functionality.

FIG. 4, FIG. 5, and associated written disclosures. '623 Patent, 5:15-61.

“In a conventional optical-electrical-optical (O/E/O) information process, i.e., electrically processing all information after optical to electrical conversion in each network node, problems arise, in that an electrical information process itself causes a bottleneck as data throughput between

nodes increases, and a higher cost is incurred for the electrical information process that increases in proportion to the data throughput.” ’623 Patent, 1:27-34.

“FIG. 3 illustrates a point-to-point communication method for data transmission between different ring networks. As shown in FIG. 3, data is transmitted via a hub node A (311) and a node E (321) to a ring network 320 (Ring 2). In this case, a cross-connection is required in order to efficiently link data transmitted between the node A (311) and the node E (321). To this end, in the disclosed conventional method, data is transmitted between two hub nodes by performing optical-electrical-optical conversion at each hub node.” ’623 Patent, 4:58-67.

“However, in this case, as described above, in the conventional optical-electrical-optical (O/E/O) information process, i.e., in electrically processing all information after optical to electrical conversion, the electrical information process itself causes a bottleneck as data throughput increases between the nodes, and the cost for the electrical information process increases in proportion to the data throughput.” ’623 Patent, 5:1-7.

“FIG. 4 is a view of an optical cross-connect apparatus for explaining an exemplary method of cross-connecting different ring networks. For the optical cross-connect apparatus of FIG. 4, in which different ring networks are cross-connected according to a conventional method, assuming input 1 is an input port of the ring network (Ring 1) and output N is an output port of the ring network (Ring 2), data can be transmitted from the ring network (Ring 1) to the ring network (Ring 2) by using a wavelength λ_2 . However, when data is transmitted from a 1:1 protection ring to a 1+1 protection ring, a single input signal has to be copied into two signals in a hub node and then transmitted due to the nature of the ring network operation method. However, this operation is actually impossible. In addition, when data is transmitted from a 1+1 protection ring to a 1:1 protection ring, only one of the two input data input from 1+1 protection ring has to be transmitted

to 1:1 protection ring. However, this function cannot be implemented with the conventional structure. In addition, another drawback of the conventional structure is that, even when 1+1 ring networks are cross-connected, network failures simultaneously occurring in the two ring networks cannot be overcome.” ’623 Patent, 4:15-37.

“FIG. 5 illustrates a structure of an optical cross-connect apparatus having a multi-casting function. FIG. 5 relates to a conventional technology of implementing a multi-casting function in which a single signal is copied into multiple signals and transmitted along different paths. In such an optical cross-connect apparatus as illustrated in FIG. 5, the signal copy function is implemented using an electrical switch instead of an optical switch. However, in this case, in order to use the electrical switch, all data has to undergo optical to electrical and electrical to optical conversion, and thus the electrical information process itself causes a bottleneck as data throughput increases, and the costs incurred for the electrical information process increases in proportion to the data throughput.” ’623 Patent, 4:38-51.

“In addition, a conventional method of implementing a multi-casting function with a signal copy function by using an electrical switch, instead of an optical switch, has problems, in that an electrical information process causes a bottleneck as data throughput increases, and a higher cost is incurred for the electrical information process that increases in proportion to the data throughput.” ’623 Patent, 1:51-57.

“Furthermore, in the case of using only an optical multi-casting process, problems arise, such as a collision of wavelengths, which occurs due to the use of the same wavelength for data transmission between different ring networks, or signal deterioration due to long-distance signal transmission.” ’623 Patent, 1:58-62.

“In general, network topologies implemented using an optical transmission apparatus are classified into a point-to-point type, a linear type, a ring type, and a mesh type. Ring networks are in widespread use due to their simple topology and their ability of fast restoration/protection within 50 ms. Methods of recovering failures in ring networks are largely classified into 1:1 protection and 1+1 protection.” ’623 Patent, 3:64-4:3.

“A wavelength selective switch (WSS) 610 is a device which allows a portion having an arbitrary wavelength from among a plurality of wavelengths of signals input through a number of input ports to selectively pass one output port. In addition, the magnitude of an input signal can be adjusted according to wavelengths, and it is possible to determine whether to block or transmit the signal.” ’623 Patent, 5:64-6:3.

Additional evidence that there would have been a motivation to combine the prior art references identified above includes the interrelated teachings of multiple prior art references; the effects of demands known to the design community or present in the marketplace; the existence of a known problem for which there was an obvious solution encompassed by the Asserted Claims; the existence of a known need or problem in the field of the endeavor at the time of the alleged invention(s); and the background knowledge that would have been possessed by a person having ordinary skill in the art. For example, the prior art references are directed to solving the same problem. Thus, a skilled artisan seeking to solve this problem, would have looked to these cited references alone or in combination. Accordingly, the motivation to combine the teachings of the prior art references disclosed herein is found in the references themselves and: (1) the nature of the problem being solved, (2) the express, implied and inherent teachings of the prior art, (3) the knowledge of persons of ordinary skill in the art, (4) the prior art as of the priority date is generally

within the same field of endeavor, or (5) the predictable results obtained in combining the different elements of the prior art.

To the extent Plaintiff contends that any reference contains multiple distinct embodiments, it would be obvious to combine elements of the distinct embodiments. A person would be motivated to make such a combination because the elements are found in the same reference and the reference as a whole is directed to the same topic or topics.

Numerous prior art references, including those identified above, reflect common knowledge and the state, scope, and content of the prior art before the alleged priority dates of the Asserted Patents. *See Graham v. John Deere Co.*, 383 U.S. 1, 35-36 (1966).

Examples of technical features that were well-known in the prior art as of the priority date of the Asserted Claims include the following.

Ring network topologies and interconnection of two or more ring networks, with traffic flowing over designated working and protection paths, and with dual-ring interworking to avoid single points of failure, were all well known. For example, Ooi describes optical nodes that connect multiple ring networks and switch among first and second transmission paths per ring using WSS-based OXC structures, with separate planes for working and protection lines and supervision/OSC messaging for fault handling.

The OXC SW 531 includes 2 input-side WSSs (1x2) and 2 output-side WSSs (2x1) for connecting the working line 500 in the ring network C and the working line 510 in the ring network D. The OXC SW 532 includes 2 input-side WSSs (1x2) and 2 output-side WSSs (2x1) for connecting the protection line 505 in the ring network C and the protection line 515 in the ring network D. Thus, in the optical node 530, the OXC SW 531 for the working lines is separated from the OXC SW 532 for the protection lines.

Ooi at 9:16–37.

In general, optical signals transmitted over a ring network include two elements: an optical signal (λ_1 to λ_n) with a plurality of wavelengths flowing over a transmission

signal channel, and an optical signal flowing over an optical supervisory channel (OSC) including AIS.

Ooi at 6:31–45.

Further, Iyer explains ring interworking and per-wavelength protection across interconnected rings, including OCh/SPRing↔OCh/DPRing dual-ring interworking with drop-and-continue and service selection so failures in one ring do not impact the other.

When two rings are interconnected, Dual Ring Inter-working prevents single points of failure by interconnecting the rings at two nodes. The drop and continue feature, in combination with path selectors/service selectors, enables the DRI protection service.

Iyer at 8:25–31.

It was conventional and well-known to provide protected rings in both 1+1 and 1:1 modes. For example, Ooi shows OUPSR (1+1) where a sending terminal uses an optical coupler to bridge signals onto working and protection lines and a receiver selects, and OSPPR (1:1-style two-working-line operation) with coordinated switching between first and second working lines.

In a typical OUPSR system, a sending terminal includes an optical coupler 33', while a receiving terminal includes a per-channel optical switch 33. For example, the optical coupler 33' of the OADM node 30a (sending terminal) branches and transmits an optical signal from the transponder 31 thereof to both of the working line 10 and the protection line 20 via the OADM switch units 32a and 32b.

Ooi at 1:41–52.

Likewise, Iyer details OCh/DPRing dedicated protection with permanent bridge and destination selection for 1+1 and coordinated head-/tail-end switching using OSC for 1:1, as well as OCh/SPRing shared protection and extra-traffic behavior. OCh/DPRing rings dedicate one fiber for working wavelength channels and the other fiber for corresponding protection wavelength channels in the counter-propagating directions. Traffic is either permanently bridged at the head-end and sent along both fibers, hence the term '1+1 protection,' or bridged (or switched) at both the head and tail end only in the event of a protection switch, hence the term '1:1 protection.'

Iyer at 6:2–13.

In 1:1-OCh/DPRING, the signal is not permanently bridged to both lightpaths at the source/head end. Hence, in the event of a protection switch, coordination is required between the head end and tail end nodes to bridge or switch the signal onto the protection lightpath. This coordination can be achieved using the optical supervisory channel (OSC), a wavelength signal dedicated to inter-nodal communication.

Iyer at 6:21–28.

Per-wavelength add/drop and switching at ring nodes using OADMs with WSSs or AWGs was likewise standard. For example, Ooi discloses WSS-based optical switch units with per-wavelength control via a supervision unit and OSC.

The optical switch unit 300 includes a plurality of wavelength-selective switches (WSS) described later. An output port of each of input-side WSSs is thru-connected or cross-connected, under control of the supervision/control unit 310, to an input port of any one of output-side WSSs.

Ooi at 6:50–60.

Likewise, Iyer describes a flexible OADM and an Optical Switching Fabric that provides per-wavelength switching under software control, including O-APS coordination during protection events.

Each node in the ring is a flexible OADM, which means that is capable of adding/dropping individual wavelengths out of the total wavelengths transported in the ring. Which wavelength each node in the ring adds, drops or passes through is completely programmable by software. This is accomplished by an Optical Switching Fabric (OSF) element that provides per-wavelength switching.

Iyer at 11:11–17.

Electrical cross-connects and hybrid cross-connect architectures were also common where signal regeneration or wavelength translation was needed. For example, Bala teaches a hybrid wavelength-interchanging cross-connect that demultiplexes optical signals, routes some transparently in the optical domain and others through a digital cross-connect for switching and wavelength conversion, and shows practical placements in inter-ring nodes.

Assume that each fiber 94, 96 carries 8 WDM signals and that the first four wavelengths λ_1 , λ_2 , λ_3 , λ_4 are electronically switched through the twelve-fold switching modules 60 while the last four wavelengths λ_5 , λ_6 , λ_7 , λ_8 remain in the optical domain. This combination of electronic and optical switching is referred to as hybrid switching.

Bala at 9:60–10:8.

Sabella '167 likewise explains that operators hand off from the optical layer to electronic digital cross-connects only when lower-rate time-division demultiplexing or service interfacing is needed, and it enumerates the electronic interfaces that motivate such selective O/E/O:

Demultiplexing and routing at lower/slower bitrate transmission hierarchies is accomplished by switching the signals from the optical layer to electronic/digital cross-connects (DXCs) 14 using electro-optical transceivers (not shown) which rearrange the data. For example, the DXCs perform time division demultiplexing of signal channels carried in one traffic channel and perform routing, switching, and other digital processing functions at considerably slower speeds. A number of potential network applications are shown interfaced by OXC's to the optical network 11 including a user interface to the electronic broadband network 20, interface with a private automatic branch exchange (PABX) 18 via an add/drop multiplexer (ADM) 16, access to the broadband ISDN (B-ISDN), homes, and interface to smaller networks like local area network (LAN) 22.

Sabella '167 at 5:11–25.

To reinforce the complementary roles of photonic and electronic switching and the rationale for minimizing unnecessary O/E/O, Sabella '167 expressly contrasts high-speed optical routing with lower-rate electronic processing and explains the throughput and efficiency advantages of keeping transit traffic in the optical domain:

While photonic switching in the space and wavelength domains may be advantageously employed to increase capacity in existing and broadband communication systems, photonics is not particularly well suited to perform data processing and data storage functions which can be more optimally performed by the electronic layer. With this combined node architecture, the 'best' features of both (1) optical and (2) electronic technologies are utilized: (1) relatively simple optical routing and switching of general traffic channels at very high speeds and (2) more complicated electronic routing, switching, and processing of individual signal channels at slower speeds. By using photonic and electronic techniques in this kind of complementary manner, the total traffic throughput of the network is considerably increased. Moreover, optical cross-connects transparently route large

blocks of traffic not only at high speeds but also with high efficiency because opto-electronic conversions are not required for straight forward traffic transmission through the path node over the transport network.

Sabella '167 at 5:26–44.

Ooi likewise acknowledges legacy electrical inter-ring nodes while teaching WSS-based optical nodes to avoid those drawbacks.

FIG. 11 is a diagram of a plurality of ring networks in which an optical signal is converted to an electric signal. A node 45 between ring networks G and H converts an optical signal transmitted over optical fibers 10 and 20 in the ring networks G and H to an electric signal. An electric switch 46 of the node 45 switches wavelength, route, and the like, based on the converted electric signal.

Ooi at 2:66–3:10.

It was well known to have hybrid optical and electrical cross-connects, where the use of the electrical cross-connect was selectively used only when needed to avoid or minimize the known drawbacks of unnecessary O/E/O conversions (including delay and additional processing power). For example,

In one aspect of the invention, the incoming multi-wavelength signal received on an input port is optically demultiplexed into its constituent signals. Some of them are sent on transparent paths to an output port while others are sent on opaque paths to an electronic switching circuit, wherein they can be both switched and wavelength translated.

Bala at 5:9-15.

Assume that each fiber 94, 96 carries 8 WDM signals and that the first four wavelengths λ_1 , λ_2 , λ_3 , λ_4 are electronically switched through the twelve-fold switching modules 60 while the last four wavelengths λ_5 , λ_6 , λ_7 , λ_8 remain in the optical domain. This combination of electronic and optical switching is referred to as hybrid switching.

Bala at 9:60-65.

The input optical switches 130 can also switch any wavelength λ_1 , λ_2 , λ_3 , λ_4 , λ_5 , λ_6 , λ_7 , λ_8 to bypass the electronic digital switching circuitry, thus providing a transparent pathway and avoiding any formatting constraints of the electronic circuitry.

Bala at 10:64-11-1.

The switching system of FIG. 4 provides the functionality required of the wavelength-interchanging cross-connect 20 of FIG. 2. However, it is opaque in the sense that no optical path extends through the DCS 30 and associated detector and laser arrays 38, 42. The opacity between networks has the advantage that it circumvents physical layer constraints exhibited by extending the all-optical network, including transparent wavelength conversion, through too many stages of amplification and wavelength conversion. That is, the DCS 30 acts as a regenerator, with both the advantages of maintaining signal form and the disadvantages of the attendant complexity and cost.

Bala at 7:14-24

In general, optical signals transmitted over a ring network include two elements: an optical signal (λ_1 to λ_n) with a plurality of wavelengths flowing over a transmission signal channel, and an optical signal flowing over an optical supervisory channel (OSC) including AIS.

Ooi at 6:36-40.

FIG. 11 is a diagram of a plurality of ring networks in which an optical signal is converted to an electric signal. A node 45 between a ring networks G and H converts an optical signal transmitted over optical fibers 10 and 20 in the ring networks G and H to an electric signal. An electric switch 46 of the node 45 switches wavelength, route, and the like, based on the converted electric signal. The above configuration has various problems such that a transmission capacity is limited, cost and size of the node 45 is increased, a signal format is fixed, and so on.

Ooi at 2:66-3:3.

Compared to the high speeds at which data can be transported over optical links, e.g., 2.5 Gbits/s, the electronic path layer nodes operate at much slower rates, e.g., 155 Mbit/s.

Sabella at 1:35-38.

Significant benefits of this network structure are that optical cross-connect nodes route very large amounts of data through the transport network at the path layer without requiring opto-electronic conversion and the losses associated with such conversions. In addition, high speed transmission over the optical path layer is 'transparent' to the network. Not only can traffic be routed through the optical cross-connect at rates much faster than through electronic cross-connects but also the optical layer transmission rate can be easily increased, e.g., from 2.5 Gbits/s to 10 Gbits/s, without impacting the optical or digital layers that make up the basic path layer. In this way, the network can be upgraded to a

much higher traffic transport rate through the optical cross-connect nodes without modifying the cross-connect node structure.

Sabella at 1:50-64.

Broadcast and multicast in the optical domain were well known, including by way of splitting and replication. For example, Ooi uses optical couplers to branch traffic for 1+1 bridging.

The optical signal is also branched by the optical coupler 523' to be simultaneously transmitted to the OADM switch unit 522b for the protection line. In other words, the optical signal flows into both the working line 500 and the protection line 505.

Ooi at 8:34-44.

Further, Nakagawa expressly describes multicast and broadcast from a central station to multiple OADM nodes on a ring and across interconnected rings via hub nodes, with NMS-synchronized control of node switches.

Furthermore, setting the same drop wavelength in each OADM node enables multicast communication in which a single transmission signal is received at a plurality of points, and broadcast communication in which a single transmission signal is received by all of the nodes.

Nakagawa at ¶[0008].

Sabella '167 also teaches that passive optical splitters such as star couplers were routinely used to split multi-wavelength signals, and it emphasizes their low cost and passive reliability compared to active space switches:

“The modular optical cross-connect node architecture includes plural optical fiber input links each containing plural wavelength channels. Each fiber link is connected to an input port of an optical coupler such as (but not limited to) a star coupler. In contrast to an expensive space switch which is an ‘active’ switching device, the star coupler is inexpensive and a passive device.”

Sabella '167 at 2:56-62.

Further, Sabella includes a star-coupler architecture that broadcasts all inputs to all outputs, providing wavelength-domain multicast and broadcast.

The wavelength channels coming from the input fiber links 1, 2, and 3, as well as optical wavelength channels locally generated by electro-optical transmitters 54

connected to digital cross-connect (DXC) 52, are combined by an optical coupler 60, such as a star coupler, that broadcasts the combined signal to each of the coupler's output ports.

Sabella at 7:12–20.

Per-wavelength connection control and management to configure protection paths, WSS states, and routing under software control was also well known. For example, Sabella describes a multi-layer management system to reconfigure networks and set up protection paths.

The operating system 100 allows an operator to reconfigure the network, set up protection paths, and supervise status. Configuration management automatically establishes channels between the end terminals in the network by allocating appropriate wavelength channels and setting up connections automatically using a routing algorithm that calculates and automatically selects free paths and channels.

Sabella at 8:62–9:10.

Further, Ooi details supervision/control units and OSC signaling for failure detection and switching.

A wavelength separating filter 301 and an optical signal receiver 311 are provided on a transmission path on an input side of the supervision/control unit 310. A wavelength coupling filter 302 and an optical signal transmitter 312 are provided on a transmission path on an output side of the supervision/control unit 310. In general, optical signals transmitted over a ring network include two elements: an optical signal (λ_1 to λ_n) with a plurality of wavelengths flowing over a transmission signal channel, and an optical signal flowing over an optical supervisory channel (OSC) including AIS.

FIG. 3B is a graph of an optical signal C1 input into the OADM nodes 120a to 120f and the optical node 130. The optical signal C1 is separated by the wavelength separating filter 301 into an optical signal C2 (λ_1 to λ_n) on the transmission signal channel and an optical signal C3 on the OSC.

Ooi at 6:31-45

Likewise, Iyer explains OSF control coordinated by O-APS messaging for ring protection.

The OSF element at each node is configured (e.g., by the O-APS protocol) to change the state of each wavelength involved in the protection switching event, from add/drop to through or vice versa.

Iyer at 12:12–19.

Wavelength conversion was a routine, well-known capability, both optically and electrically. For example, Ooi discloses wavelength conversion functions integrated with WSS-based mux/demux blocks and tunable transponders.

For example, Ooi teaches that wavelength converter 803 includes optical branching filters (demultiplexers: DEMUX) 804, tunable transponders (Tu-TRP) 805, wavelength non-depending wave couplers (multiplexers: MUX) 806, and 1x2 WSSs 807. *See* Ooi at 10:25–52.

Further, Sabella describes SOA-based four-wave-mixing converters for single channels or combs/set of wavelength channel.

The FWM SOA wavelength converter can be used to wavelength convert a single wavelength channel or an entire comb/set of wavelength channels by an amount Q . The FWM conversion process is independent of the signal or its spectrum.

Sabella at 8:54–61.

Finally, multi-degree cross-connect apparatus using WSS arrays to connect multiple rings and multiple working/protection lines was established practice, well known to a POSITA. For example, Ooi shows OXC switch units built from $1 \times N$ and $N \times 1$ WSS elements to interconnect working and protection lines of multiple rings and explains that scaling is achieved by increasing the number of WSSs; the architecture explicitly separates working and protection planes.

As shown in FIGS. 7A and 7B, the WSSs for working lines are connected to only working lines, while that for protection lines are connected to only protection lines. In an optical node connecting N ring networks, each of the OXC SWs includes $2N$ input-side WSSs and $2N$ output-side WSSs. Each of the input-side WSSs includes 1 input port and N output ports ($1 \times N$). Each of the output-side WSSs includes N input ports and 1 output port ($N \times 1$).

Ooi at 10:3–10.

Even if the number of networks to be connected to the optical node 130 is increased, the increase can be handled easily by only increasing the number of WSSs.

Ooi at 10:59–11:6.

Collectively, these references demonstrate that ring networks, protected 1+1 and 1:1 operation, hybrid optical and electrical cross-connects, broadcast/multicast cross-connects, including with optical splitting, WSS-based per-wavelength switching, wavelength conversion, and software-controlled inter-ring cross-connection were all part of the general knowledge in the art, well known to a POSITA, within the skill of the routine, ordinary toolkit well before the Asserted Claims. *See* Sabella '167; Bala '799; Ooi '270; Nakagawa '210; Iyer '713.

Defendant is not aware of any secondary considerations supporting non-obviousness of the Asserted Claims of the '623 Patent. To the extent Plaintiff seeks to introduce any alleged secondary considerations in the future, Defendant reserves the right to respond at that time.

For at least the reasons described herein as well as in the combinations of the prior art identified above and charted in the claim charts found at Exhibits A-1 to A-5, it would have been obvious to one of ordinary skill in the art to combine any of a number of prior art references, including any combination of those identified above.

Defendant further identifies the commercial embodiments of the inventions described in the prior art charted in Exhibits A-1 to A-5, to the extent those systems or services were on sale or in public use before the priority dates of the '623 Patent.

A person having ordinary skill in the art would have been motivated to combine any of the references charted in attached Exhibits A-1 to A-5 for at least the reasons described above. A person having ordinary skill in the art at the time of filing of the '623 Patent would also have understood the references listed above, alone or in combination, to contain explicit and/or implicit teaching, suggestion, and/or rationales to combine them, including as further described above.

To the extent one of the above references does not disclose a specific limitation, it would have been obvious to combine that reference with one or more of the other references. A person of ordinary skill in the art at the time of the alleged invention would have combined the references because each generally relates to similar technology e.g., ring-based optical interconnection of working and protection rings; OADM/OXC nodes with WSS-based per-wavelength switching; digital cross-connect regeneration and wavelength translation when required; optical multicast/broadcast via couplers, star couplers, or rejection/add filtering; and ring protection and interworking across multiple rings), and the references could be combined, or substituted, according to known methods that would yield predictable results. Moreover, to the extent the '623 Patent identifies any needed improvement in the art, the references provide known techniques that would yield the required improvement, and given the finite number of techniques identified, it would have been obvious to try those approaches with a reasonable expectation of success.

The combinations of references disclosed in these contentions are exemplary and are not intended to be exhaustive. Additional obviousness combinations of the references identified here are possible, and Defendant may rely on such combination(s) in this litigation. In particular, Defendant is currently unaware of Plaintiff's allegations with respect to the level of skill in the art and the qualifications of the typical person of ordinary skill in the art. Defendant is also unaware of the extent, if any, to which Plaintiff may contend that limitations of the claims at issue are not disclosed in the prior art identified by Defendant, and the extent to which Plaintiff will contend that elements not disclosed in the '623 Patent specifications would have been known to persons of skill in the art. And Defendant does not yet know how the Court will construe terms in the Asserted Claims of the '623 Patent. Defendant is also continuing their investigation of the large universe of prior art to identify potential prior art systems, publications related to those systems, and third

parties that may have information about those systems. Plaintiff may also be in possession of prior art that Defendant may receive as discovery progresses in this case. Defendant reserves the right to amend and supplement these contentions to identify other prior art and combinations rendering the Asserted Claims of the '623 Patent anticipated and/or obvious.

2. Invalidity Under 35 U.S.C. § 112

Defendant identifies the following claim limitations that render the Asserted Claims of the '623 Patent invalid based on lack of enablement and written description under pre-AIA 35 U.S.C. § 112 ¶ 1 and/or are indefinite under pre-AIA 35 U.S.C. § 112 ¶ 2. All claims that depend from a claim containing a limitation identified below are invalid by way of their dependence. Defendant reserves the right to modify these contentions or further develop them as discovery progresses, including through expert disclosure and discovery.

a) Indefiniteness

The following claim elements are indefinite and, accordingly, the claims in which they appear are invalid:

- “splitting an input multi-wavelength optical signal of the first ring network into at least one portion each having a specific wavelength, by using an optical coupler “ (Claims 1, 14, 27)
- “the wavelength selective switch selectively passes at least one of the portions of the input multi-wavelength optical signal through at least one output of the second ring network “ (Claims 1, 14, 27)
- “wherein [a multi-casting function is provided / *the system includes a multi-casting function / the apparatus includes a multi-casting function*] both in an optical domain and in an electrical domain so that if the input signal does not require signal

reproduction or wavelength conversion, it is cross-connected in the optical domain, and if the input signal requires signal reproduction or wavelength conversion, it is cross-connected in the electrical domain” (Claims 1, 14, 27)

- “a portion having a specific wavelength” (claim 4, 5, 9, 10, 11, 12, 13, 17, 22, 23, 24, 25, 26)
- “portions having specific wavelengths” (claims 6, 18, 19)

b) Lack of Enablement and Written Description

Defendant contends that the specification (including any specification to which priority is sought) fails to describe the full scope of the claimed inventions in sufficient detail such that one skilled in the art could reasonably conclude that the inventor had possession of the claimed invention, and/or lack enablement, including:

- “cross-connecting an input working ring and an input protection ring of a first ring network to an output working ring and an output protection ring of a second ring network” (Claims 1) and “a ring network connection unit connecting an input working ring and an input protection ring of a first ring network to an output working ring and an output protection ring of a second ring network” (Claim 14)

For example, to the extent it is Plaintiff’s assertion in its infringement allegations that these limitations of a “ring network connection unit connecting” or step of “cross-connecting” “an input working ring and an input protection ring of a first ring network to an output working ring and an output protection ring of a second ring network” encompass connections that are not like the type of cross-connecting of pairs of working and protection input and output rings disclosed, then Defendant contends that the specification (including any specification to which priority is sought) fails to describe the full scope of the claimed inventions in sufficient detail such that one skilled in

the art could reasonably conclude that the inventor had possession of the claimed invention, and/or fails to enable the full scope of the claimed invention. That is, the specification discloses only specific input and output ring cross-connection techniques, solutions and structures.

- “splitting an input multi-wavelength optical signal of the first ring network into at least one portion each having a specific wavelength, by using an optical coupler;” (Claims 1, 14, 27)

For example, Defendant contends that the specification (including any specification to which priority is sought) fails to describe the full scope of the claimed inventions in sufficient detail such that one skilled in the art could reasonably conclude that the inventor had possession of an invention involving splitting an input multi-wavelength signal where the result is at least one portion each having a specific wavelength, by using an optical coupler, and/or fails to enable an such an invention. Furthermore, to the extent it is Plaintiff’s assertion in its infringement allegations that these limitations encompass monitoring optical power, splitting signals such as supervisory/control signals or channels, operation of optical switches, operation of wavelength splitters, or other operations, techniques or structures not like the splitting by an optical coupler as disclosed, then Defendant contends that the specification (including any specification to which priority is sought) fails to describe the full scope of the claimed inventions in sufficient detail such that one skilled in the art could reasonably conclude that the inventor had possession of the claimed invention, and/or fails to enable the full scope of the claimed invention.

- “transmitting the input multi-wavelength optical signal to at least one wavelength selective switch” (Claim 1) and “a connection control unit transmitting the input multi-wavelength optical signal to at least one wavelength selective switch” (Claims 14, 27)

For example, to the extent it is Plaintiff's assertion in its infringement allegations that these limitations encompass operations, techniques or structures that involve inputting a multi-wavelength optical signal that is less than or different to the input multi-wavelength optical signal recited in the prior claim element ("an input multi-wavelength optical signal of the first ring network") or otherwise not like the transmitting of such signal to a wavelength selective switch as part of cross-connecting as disclosed, then Defendant contends that the specification (including any specification to which priority is sought) fails to describe the full scope of the claimed inventions in sufficient detail such that one skilled in the art could reasonably conclude that the inventor had possession of the claimed invention, and/or fails to enable the full scope of the claimed invention.

- the wavelength selective switch selectively passes at least one of the portions of the input multi-wavelength optical signal through at least one output of the second ring network," (Claims 1, 14, 27)

For example, Defendant contends that the specification (including any specification to which priority is sought) fails to describe the full scope of the claimed inventions in sufficient detail such that one skilled in the art could reasonably conclude that the inventor had possession of an invention involving the wavelength selective switch selectively passing *the* portions of the input multi-wavelength optical signal through at least one *output of the second ring network*, and/or fails to enable such an invention. Furthermore, to the extent it is Plaintiff's assertion in its infringement allegations that these limitations encompass operations, techniques or structures that involve passing a portion of multi-wavelength optical signal that is less than or different to the portion generated in the prior claim element ("splitting an input multi-wavelength optical signal of the first ring network into at least one portion each having a specific wavelength, by using an

optical coupler”) or otherwise not like the passing by a WSS of such signal to an output of the second ring network as part of cross-connecting as disclosed, then Defendant contends that the specification (including any specification to which priority is sought) fails to describe the full scope of the claimed inventions in sufficient detail such that one skilled in the art could reasonably conclude that the inventor had possession of the claimed invention, and/or fails to enable the full scope of the claimed invention.

Defendant also contends that the specification (including any specification to which priority is sought) fails to describe the full scope of the claimed inventions in sufficient detail such that one skilled in the art could reasonably conclude that the inventor had possession of an invention involving the wavelength selective switch passing optical signals through at least one *output of the second ring network*, and/or fails to enable such an invention.

- “wherein [a multi-casting function is provided / *the system includes a multi-casting function* / *the apparatus includes a multi-casting function*] both in an optical domain and in an electrical domain so that if the input signal does not require signal reproduction or wavelength conversion, it is cross-connected in the optical domain, and if the input signal requires signal reproduction or wavelength conversion, it is cross-connected in the electrical domain” (Claims 1, 14, 27)

For example, Defendant contends that the specification (including any specification to which priority is sought) fails to describe the full scope of the claimed inventions encompassed by this claim element, which is functional in nature and disembodied from particular structures of the claim, in sufficient detail such that one skilled in the art could reasonably conclude that the inventor had possession of the full scope of the claimed invention and/or fails to enable the full scope of the claimed invention. Furthermore, to the extent it is Plaintiff’s assertion in its infringement

allegations that these limitations encompass operations, techniques or structures unlike the solutions disclosed, then Defendant contends that the specification (including any specification to which priority is sought) fails to describe the full scope of the claimed inventions in sufficient detail such that one skilled in the art could reasonably conclude that the inventor had possession of the claimed invention, and/or fails to enable the full scope of the claimed invention.

B. '807 Patent

1. Prior Art Patents and Publications

Reference	Issue/Publication Date
Chinese Patent Application Publ. No. 101272310 ("Long")	September 24, 2008
U.S. Patent Application Publ. No. 2003/0165119 ("Hsu")	September 4, 2003
Amendment 1 to Recommendation ITU-T G.8032/Y.1344 ("Amendment1 G.8032 2009")	April 2009
Recommendation ITU-T G.8032/Y.1344 ("G.8032 2008")	June 2008
IEEE 802.1Q 2005, IEEE Standard for Local and Metropolitan Area Networks ("IEEE 802.1Q 2005")	May 19, 2006
Amendment1 G.8032 2009, G.8032 2008, IEEE 802.1Q 2005 (collectively, "Prior Art G.8032")	April 2009
Ethernet Ring Protection with Managed FDB Using APS Payload ("Im")	November 19, 2007

To the extent the Asserted Claims of the '807 Patent are not anticipated by the prior art references identified above and in the claim charts attached as Exhibits B-1 to B-4, the Asserted Claims are rendered obvious by those references, either alone or in combination with other prior art references as identified below and in the attached claim charts and/or the knowledge of one skilled in the art.

In *KSR International Co. v. Teleflex Inc.*, 500 U.S. 398 (2007), the U.S. Supreme Court rejected the Federal Circuit's rigid "teaching, suggestion, or motivation" requirement in favor of a flexible, functional approach in which an explicit finding of a "motivation" to combine prior art

references is not required to establish obviousness. The Supreme Court held that it is sufficient that a combination of elements was “obvious to try” holding that, “[w]hen there is a design need or market pressure to solve a problem and there are a finite number of identified, predictable solutions, a person of ordinary skill has good reason to pursue the known option within his or her technical grasp.” *Id.* at 402; *see also Dystar Textilfarben GmbH v. C.H. Patrick Co.*, 464 F.3d 1356, 1368 (Fed. Cir. 2006) (explaining that when the “combination of references results in a product or process that is more desirable, for example because it is stronger, cheaper, cleaner, faster, lighter, smaller, more durable, or more efficient,” there exists a motivation to combine prior art references even when there is no explicit suggestion in the references themselves “[b]ecause the desire to enhance commercial opportunities by improving a product or process is universal—and even commonsensical”); *LeapFrog Enters., Inc. v. Fisher-Price, Inc.*, 485 F.3d 1157 (Fed. Cir. 2007) (applying *KSR* and holding that “one of ordinary skill in the art of children’s learning toys would have found it obvious to combine the Bevan device with the SSR to update it using modern electronic components in order to gain the commonly understood benefits of such adaption, such as decreased size, increased reliability, simplified operation, and reduced cost”).

Defendant contends that where the Asserted Claims are anticipated by a particular reference, they are also rendered obvious by that same reference alone, or in combination with the other references, as show in the table below:

Primary Reference	In Combination with One or More of
Long	Hsu
Long	Amendment1 G.8032 2009
Long	G.8032 2008
Long	IEEE 802.1Q 2005
Long	Prior Art G.8032
Long	Im
Prior Art G.8032	Long
Prior Art G.8032	Hsu
Prior Art G.8032	Im

Primary Reference	In Combination with One or More of
Amendment1 G.8032 2009	Long
Amendment1 G.8032 2009	Hsu
Amendment1 G.8032 2009	Im
Amendment1 G.8032 2009	G.8032 2008
Amendment1 G.8032 2009	IEEE 802.1Q 2005
G.8032 2008	Long
G.8032 2008	Hsu
G.8032 2008	Im
G.8032 2008	Amendment1 G.8032 2009
G.8032 2008	IEEE 802.1Q 2005
IEEE 802.1Q 2005	Long
IEEE 802.1Q 2005	Hsu
IEEE 802.1Q 2005	Im
IEEE 802.1Q 2005	Amendment1 G.8032 2009
IEEE 802.1Q 2005	G.8032 2008
Hsu	Long
Hsu	Amendment1 G.8032 2009
Hsu	G.8032 2008
Hsu	IEEE 802.1Q 2005
Hsu	Prior Art G.8032
Hsu	Im
Im	Long
Im	Amendment1 G.8032 2009
Im	G.8032 2008
Im	IEEE 802.1Q 2005
Im	Prior Art G.8032
Im	Hsu

Defendant further contends that one of ordinary skill in the art, at the time of the alleged invention of the Asserted Claims, would have been motivated to combine the references disclosed herein in such a way to reach the alleged inventions. The teaching, suggestion, or motivation to combine these references, although not required, is explicitly or implicitly found in one or more of the following: the knowledge or common sense of one of ordinary skill in the art; the prior art references themselves and/or the prior art as a whole, including interrelated teachings of multiple prior art references; the subject matter acknowledged as prior art in the '807 Patent; the nature of the problem to be solved and the existence of similar improvements in similar applications; design

incentives and other market forces, including the advantages of creating a superior and more desirable product and the effects of demands known to the design community or present in the marketplace; the ability to implement the alleged invention as a predictable variation of the prior art; improvements in similar devices; the interrelated teachings of multiple prior art references; any needs or problems known in the field addressed by the '807 Patent; and the number of identified, predictable solutions to the problem allegedly addressed by these patents. In addition, the simultaneous (and/or prior) inventions described above, and elsewhere in these contentions, is evidence that motivation to combine the concepts described in the various prior art references did, in fact, exist, and they were, in fact, combined. The combinations evidencing teachings, suggestions, and/or motivations to combine the prior art references in a way that renders the Asserted Claim obvious are merely exemplary. A person of ordinary skill would have access to the materials found in the cited references and attached exhibits and would have at least the ordinary creativity and skill to combine the cited references in ways not explicitly recited herein.

Further, Defendant hereby identifies additional motivation and reasons to combine the cited art: one of skill in the art would have been motivated to combine these references, because these references relate to common objectives and subject matter. The identified prior art references relate to common objectives and share commonalities in terms of their general subject matter as well as the methods, technologies, and approaches used, including their descriptions of ethernet ring protection and management of ethernet ring networks, and virtual local area network (VLAN) based bridges. *See, e.g.*, Long at [0001] and [0007]; Hsu at [0002] and [0008]; Prior Art G.8032 at Amd1, pp. 1-2; Prior Art G.8032 at 2008, pp. 1, 6; Prior Art G.8032 at 2005, p. 1; Im at p. 1.

Further, the Asserted Claims of the '807 Patent contain merely a duplication of concepts known in the prior art and having no patentable significance given that no unexpected results occur

(i.e., the patentee merely combines known prior art elements according to known methods to yield predictable results). For example, the '807 Patent itself admits that: (a) Ethernet ring networks include Ethernet ring nodes having two ports, each forming a connection with an adjacent ring node, (b) Filtering database tables are created for each ring node in an Ethernet ring to prevent the formation of infinite loops, (c) Ring protection link (RPL) owners are positioned adjacent to RPL owner nodes and logically block one predetermined link in an Ethernet ring network, and (d) When a link failure occurs in Ethernet ring networks, the failed link is physically blocked and the RPL owner node removes logical RPL blocking to restore connectivity using ring automatic protection switching (R-APS) messages, and (e) A bridge can be based on an industry standard, such as IEEE 802.1D.

An Ethernet ring network is constituted by Ethernet bridge nodes (hereinafter, referred to as 'ring nodes') providing a plurality of Ethernet ports. Two (one is East: hereinafter, referred to as 'E' and the other is West: hereinafter, referred to as 'W') of the Ethernet ports of each ring node are used to form a connection with an adjacent ring node and the rest of the ports are connected to an Ethernet network outside of the ring network. The number of ring ports of a node receiving a sub ring may be increased by one depending on the number of received sub rings.

In the Ethernet ring network, each of the ring nodes should create a filtering database table so as to prevent the delivery of an Ethernet frame from forming an infinite loop. Such loop prevention may be implemented by logically blocking one predetermined link called a ring protection link (RPL) in the Ethernet ring network. The Ethernet ring network has a ring protection link (RPL) owner node in order to provide such logical blocking. Herein, the delivery of the infinite loop of the ring is prevented by logically blocking an Ethernet port connected with the ring protection link among Ethernet ports of the ring protection link owner node. When both sides of the ring protection link need to be logically blocked, a node position at the opposite side to the ring protection link owner node of the ring protection link may be called a ring protection link (RPL) node.

Meanwhile, when a link failure occurs in the Ethernet ring network, the link with the failure provides physical blocking and when the link with the failure is blocked in order to prevent an error, the ring protection link owner node prevents the delivery of the infinite loop of the ring and connects all nodes that belong to the ring network by rapidly removing logical blocking. For the ring blocking and ring management, and state delivery, an R-APS message is used.

‘807 Patent, 1:23-55.

Further, a bridge described in the present invention includes an Ethernet bridge based on IEEE 802.1D and a provider bridge (PB), a provider backbone bridge (PBB), and a provisioned provider backbone bridge (PBB-TE or PBT) adopting traffic engineering, acquired by the Ethernet bridge based on IEEE 802.1D. Furthermore, bridge function processing may be applied similarly, even to a switch providing MPLS label switching, and in the present invention it is briefly called a bridge.

‘807 Patent, 5:56-64.

Additional evidence that there would have been a motivation to combine the prior art references identified above includes the interrelated teachings of multiple prior art references; the effects of demands known to the design community or present in the marketplace; the existence of a known problem for which there was an obvious solution encompassed by the Asserted Claims; the existence of a known need or problem in the field of the endeavor at the time of the alleged invention(s); and the background knowledge that would have been possessed by a person having ordinary skill in the art. For example, the prior art references are directed to solving the same problem. Thus, a skilled artisan seeking to solve this problem, would have looked to these cited references alone or in combination. Accordingly, the motivation to combine the teachings of the prior art references disclosed herein is found in the references themselves and: (1) the nature of the problem being solved, (2) the express, implied and inherent teachings of the prior art, (3) the knowledge of persons of ordinary skill in the art, (4) the prior art as of the priority date is generally within the same field of endeavor, or (5) the predictable results obtained in combining the different elements of the prior art.

To the extent Plaintiff contends that any reference contains multiple distinct embodiments, it would be obvious to combine elements of the distinct embodiments. A person would be

motivated to make such a combination because the elements are found in the same reference and the reference as a whole is directed to the same topic or topics.

Numerous prior art references, including those identified above, reflect common knowledge and the state, scope, and content of the prior art before the alleged priority dates of the Asserted Patents. *See Graham v. John Deere Co.*, 383 U.S. 1, 35-36 (1966).

Non-Limiting examples of technical features that were well-known in the prior art as of the priority date of the Asserted Claims include:

- *Managing Ethernet ring networks of VLAN-based bridges*

Managing Ethernet ring networks of VLAN-based bridges was well-known in the art prior to the alleged invention of the '807 Patent. For example, Prior Art G.8032 (B-1) provides:

Ethernet ring: An Ethernet ring is a collection of Ethernet ring nodes forming a closed loop, whereby each node is connected to two adjacent nodes via a duplex communications facility.

3.2.2 Ethernet ring node: An Ethernet ring node is a network element which implements at least the following functionalities:

- One Ethernet connection function (ETH_C) with a dedicated Ethernet flow forwarding function (ETH_FF) for forwarding R-APS control traffic.
- Two ring ports, including ETHDi/ETH Adaptation function at the ring maintenance entity group level (MEL).
- Ethernet ring protection (ERP) control process controlling blocking and unblocking of traffic over the ring ports.

Prior Art G.8032 (B-1), 2008 p. 3.

6 Introduction

This Recommendation specifies protection switching mechanisms and protocol for Ethernet layer network (ETH) Ethernet rings. Ethernet rings can provide wide-area multipoint connectivity more economically due to their reduced number of links. The mechanisms and protocol defined in this Recommendation achieve highly reliable and stable protection; and never form loops, which would fatally affect network operation and service availability.

Each ring node is connected to adjacent nodes participating in the same ring, using two independent links. A ring link is bounded by two adjacent nodes and a port for a ring link is called a ring port. The minimum number of nodes on a ring is two.

Prior Art G.8032 (B-1), 2008 p. 5.

In Figure 9-5, there are two interconnected rings. Ring ERP1 is composed of nodes A, B, C and D and by the links between these nodes. Ring ERP2 is composed of nodes C, D, E and F and by the links C-to-F, F-to-E, E-to-D. The link between D and C is used for traffic of rings ERP1 and ERP2. On their own, ERP2 links do not form a closed physical loop. A closed loop may be formed by the ring links of ERP2 and the link between interconnection nodes that is controlled by ERP1; ERP2 is a sub-ring. Node A is RPL Owner Node for ERP1, Node E is RPL Owner Node for ERP2. The nodes (Nodes A and E) are responsible for blocking the traffic channel on the RPL for ERP1 and ERP2, respectively. There are no restrictions to which link on a ring may be defined as RPL. For example, the RPL of ERP1 could be defined as the link between node C and node D.

Prior Art G.8032 (B-1), Amd1 p. 2.




Figure 9-5 – Ethernet ring interconnection architecture – Normal condition (multi-ring/ladder network)

Prior Art G.8032 (B-1), Amd1 FIG. 9-5.

VLANs aim to offer the following benefits:

- a) VLANs facilitate easy administration of logical groups of stations that can communicate as if they were on the same LAN. They also facilitate easier administration of moves, adds, and changes in members of these groups.
- b) Traffic between VLANs is restricted. Bridges forward unicast, multicast, and broadcast traffic only on individual LANs that serve the VLAN to which the traffic belongs.

- c) As far as possible, VLANs maintain compatibility with existing bridges and end stations.
- d) If all Bridge Ports are configured to transmit and receive untagged frames (3.39), bridges will work in plug-and-play IEEE Std 802.1D mode. End stations will be able to communicate throughout the network.

Prior Art G.8032 (B-1), IEEE 2005 p. 2.

8.1.4 Traffic segregation

A Bridge can filter frames to confine them to LANs that belong to the VLAN to which they are assigned and, thus, define the VLAN's maximum extent (7.3). The functions that support the use and maintenance of information for this purpose are

- a) Configuration of a PVID for each Port, to associate a VID with untagged and priority-tagged received frames (6.7.1), and parameters for Protocol VLAN Classification (6.8) if implemented;

- b) Configuration of Static VLAN Registration Entries (8.8.2);

- c) Configuration of the Enable Ingress Filtering parameter to enable or disable application of Static VLAN Registration Entries to received frames.

Prior Art G.8032 (B-1), IEEE 2005 p. 39. *See also:*

Embodiments of the present invention provide a method and a device for automatic protection switching in an Ethernet ring network. The method mainly includes: upon detecting a failure or failure recovery in a link between nodes, determining, by a node, an Ethernet ring instance affected or not affected by the failure or failure recovery; and sending a protection switching message carrying information of the Ethernet ring instance affected or not affected by the failure or failure recovery. By using the embodiments of the present invention, automatic protection switching for an Ethernet ring network can be achieved while reducing the number of Operation, Administration and Maintenance (OAM) layers and conserving virtual local area network (VLAN) resources.

Long (B-2), Abstract.

To support large network configurations consisting of connected rings, the invention provides methodologies and data fields in the control packet protocol for coordinating control between connected rings. When two rings are connected through a shared link formed between two shared bridges, e.g., bridges which have ports belonging to the two connected rings, at least one of the rings carries the control packets being forwarded around the other ring. In some embodiments, the ring having the higher priority as between the two connected rings carries the control packets of the lower priority ring. The control packets are preferably marked as native to the lower priority ring, and thus foreign to the higher priority ring, before they are forwarded on to the higher priority ring.

Hsu (B-3), [0010].

In this paper, we present the recent development of Ethernet ring protection and propose a new class of an Ethernet ring protection scheme so-called managed FDB that utilizes an extended APS payload to carry effective protection topology information. If a failure is sensed at an Ethernet ring node, the node generates a flip address list (FAL) from its FDB, which are associated with the failed port. Then, the node transmits an APS signaling frame including the FAL. The rest nodes then reconstruct their own FDB to provide protected MAC forwarding in the new network topology. The managed FDB provides protection which reduces the overshoots phenomenon and thus relieves the over-provisioning requirement. Our worst-case analysis suggests that the managed FDB technique may increase network utilization up to several times over the conventional FDB flush mechanism.

Im (B-4), p. 1.

- *Recognizing ring ports of an ethernet ring node/assigning a ring ID to an Ethernet ring*

Recognizing ring ports of an ethernet ring node was well-known at the time of the alleged invention. It was also known that a ring identifier (ring ID) could be assigned to a unique Ethernet ring. Prior Art G.8032 makes clear that “logical rings shall be identifiable for management purposes.” Prior Art G.8032 (B-1), 2008, pg. 25. Additionally, for example, Long (B-2) provides:

To distinguish different Ethernet ring instances, each Ethernet ring instance is configured with an Ethernet ring instance identifier (Ringlet). On each Ethernet ring instance, a normal blocking port is configured. With this blocking port, customer service traffic can be prevented from forming a loop. Different Ethernet ring instances are typically configured with different normal blocking ports to optimize corresponding traffic routing.

Long (B-2), [0049].

Then, the embodiments of the present invention need to establish and maintain, in each node, information about the Ethernet ring instances carried on the link where each node port is located. This can be saved in the manner shown in Table 1. The table saves the correspondence between various ports of the node and Ethernet ring instances. It describes which Ethernet ring instances' service traffic is carried on the link where each port is located under the current condition. The example in Table 1 indicates that the service traffic of Ethernet ring instances Ringlet_1, Ringlet_2, and Ringlet_3 is carried on the westbound port of the ring, while the eastbound port carries the service traffic of Ethernet ring instances Ringlet_1 and Ringlet_3.

Long (B-2), [0052].

Table 1:

Port	Ethernet Ring Instance Identifier
WEST	Ringlet_1
WEST	Ringlet_2
WEST	Ringlet_3
EAST	Ringlet_1
EAST	Ringlet_3

Long (B-2), Table 1. *See also:*

Embodiments of the present invention are now described in detail with reference to the drawings in the Figures. As described in more detail, the present invention provides techniques and a control packet protocol for supporting rapid fault recovery while avoiding loops in ring topologies. Generally, and for purposes of the description of the preferred embodiments, a ring topology is a network configuration formed when Layer 2 bridges are placed in a circular fashion with each bridge having two and only two ports belonging to a specific ring. FIGS. 1A-1C show different examples of ring topologies. FIG. 1A shows a single ring 8 having bridges 10 connected by paths 12. Each bridge in the ring 8 in FIG. 1A has two ports 10 *a* and 10 *b* belonging to the ring. FIG. 1B shows two adjacent rings, 8 *a* and 8 *b*, with a single bridge 14 having two ports belonging to each ring.

Hsu (B-3), [0041].

The RPDUs contain, among other fields, a ring identifier which identifies the ring to which the RPDUs applies. Every ring has an identifier assigned by the network administrator. This field is not strictly necessary for a single ring, but is useful in allowing rings to be connected and controlled separately, as explained in greater detail below. If the ring is operational and all its links enabled, the blocking port on the master bridge should receive each RPDUs having the ring identifier within a set delay time. The blocking port thus awaits each RPDUs, step 36. The master bridge continues sending RPDUs after waiting hello seconds from transmission of the previous RPDUs, steps 38, 34.

Hsu (B-3), [0045].

In Figure 9-5, there are two interconnected rings. Ring ERP1 is composed of nodes A, B, C and D and by the links between these nodes. Ring ERP2 is composed of nodes C, D, E and F and by the links C-to-F, F-to-E, E-to-D. The link between D and C is used for traffic of rings ERP1 and ERP2. On their own, ERP2 links do not form a closed physical loop. A closed loop may be formed by the ring links of ERP2

and the link between interconnection nodes that is controlled by ERP1; ERP2 is a sub-ring. Node A is RPL Owner Node for ERP1, Node E is RPL Owner Node for ERP2. The nodes (Nodes A and E) are responsible for blocking the traffic channel on the RPL for ERP1 and ERP2, respectively. There are no restrictions to which link on a ring may be defined as RPL. For example, the RPL of ERP1 could be defined as the link between node C and node D.

Prior Art G.8032 (B-1), Amd1, p. 2.




Figure 9-5 – Ethernet ring interconnection architecture – Normal condition (multi-ring/ladder network)

Prior Art G.8032 (B-1), Amd1, FIG. 9-5.

R-APS messages of this sub-ring that are forwarded over its R-APS virtual channel are broadcasted or multicasted over the interconnected network. For this reason, the broadcast/multicast domain of the R-APS virtual channel could be limited to the necessary links and nodes. For example, the R-APS virtual channel could span only the interconnecting rings or sub-rings that are necessary for forwarding R-APS messages of this sub-ring. Care must be taken to ensure that the local R-APS messages of the sub-ring being transported over the R-APS virtual channel into the interconnected network can be uniquely disambiguated from those of other interconnected ring R-APS messages. This can be achieved, for example, using separate VIDs for R-APS virtual channels of different sub-rings.

Prior Art G.8032 (B-1), Amd1, p. 5.

- *Determining whether an Ethernet ring node is a ring protection link (RPL) owner node or an RPL neighbor node*

Determining whether an Ethernet ring node is an RPL owner node or an RPL neighbor node was well-known at the time of the alleged invention of the '807 Patent. For example, several references expressly define an RPL and an RPL owner node, depict the placement of an RPL, an RPL owner node, and RPL neighbor node, and describe their respective roles and functions in managing an Ethernet ring network. *See, e.g.,* Prior Art G.8032 (B-1), Amd1 pp. 1-3, Amd1 FIG. 9-5, Amd1 FIG. 9-6, Amd1 FIG. 9-7, 2008 p. 3, 2008 p. 7, 2008 pp. 30-36, 2008 FIG. 9-1, 2008 FIG. 9-2; Long (B-2), [0049], [0086], FIG. 3, FIG. 4; Im (B-4), FIG. 1, FIG. 4, p. 2. *See also* Hsu (B-3) [0013-0016] (describing a designated ring port convertible between a blocking and a non-blocking state).

- *Registering VLAN identifiers in filtering databases of ring ports*

It was well known in the art before the alleged invention of the '807 Patent to register virtual local area network (VLAN) identifiers in filtering databases of ring ports. For example, Prior Art G.8032 (B-1) teaches filtering frames using filtering database entries that include VLAN identifiers:

8.6.3 Frame filtering

The Forwarding Process takes filtering decisions, i.e., reduces the set of potential transmission Ports (8.6.1), for each received frame on the basis of

- a) Destination MAC Address;
- b) VID;
- c) The information contained in the Filtering Database for that MAC Address and VID;
- d) The default Group filtering behavior for the potential transmission Port (8.8.6); in accordance with the definition of the Filtering Database entry types (8.8.1, 8.8.3, and 8.8.4). The required behavior is summarized in 8.8.6, 8.8.8, Table 8-4, Table 8-5, and Table 8-6.

Each of the Reserved MAC Addresses specified in Table 8-1 shall be permanently configured in the Filtering Database in VLAN-aware Bridges. The Filtering

Database Entries for Reserved MAC Addresses shall specify filtering for all Bridge Ports and all VLANs. Management shall not provide the capability to modify or remove entries for Reserved MAC Addresses.

Prior Art G.8032 (B-1), IEEE 2005 p. 48.

Static Filtering Entries and Group Registration Entries comprise

- k) A MAC Address specification;
- l) A VLAN Identifier (VID);
- m) A Port Map, with a control element for each outbound Port to specify filtering for that MAC Address specification and VID.

Prior Art G.8032 (B-1), IEEE 2005 p. 53.

If a frame is classified into a VLAN containing a given outbound Port in its member set (8.8.9), forwarding or filtering through that Port is determined by the control elements of filtering entries for the frame's destination MAC Address and for VLANs with the same VID or Filtering Identifier (FID, 8.8.7) as the frame's VLAN.

Prior Art G.8032 (B-1), IEEE 2005 p. 62.




Figure 8-17—Ingress/egress control information in the forwarding path

Prior Art G.8032 (B-1), IEEE 2005 FIG. 8-17.

For example, for the Ethernet ring network shown in FIG. 2, FIG. 3 shows a schematic structural diagram of Ethernet ring instance 1, and FIG. 4 shows a schematic structural diagram of Ethernet ring instance 2. Ethernet ring instance 1 is used to protect VLAN_2 and VLAN_3, and its blocking port is set on node A; Ethernet ring instance 2 is used to protect VLAN_5, and its blocking port is set on node C. Ethernet ring instance 1 and Ethernet ring instance 2 each have a different Ethernet ring instance identifier.

A correspondence between the aforementioned Ethernet ring instance identifiers and VLAN protection groups is configured on each node. Generally, a group of service VLANs protected by one Ethernet ring instance have the same ring access point.

Then, the embodiments of the present invention need to establish and maintain, in each node, information about the Ethernet ring instances carried on the link where each node port is located. This can be saved in the manner shown in Table 1. The table saves the correspondence between various ports of the node and Ethernet ring instances. It describes which Ethernet ring instances' service traffic is carried on the link where each port is located under the current condition. The example in Table 1 indicates that the service traffic of Ethernet ring instances Ringlet_1, Ringlet_2, and Ringlet_3 is carried on the westbound port of the ring, while the eastbound port carries the service traffic of Ethernet ring instances Ringlet_1 and Ringlet_3.

Long (B-2), [0050-0052].

Table 1:

Port	Ethernet Ring Instance Identifier
WEST	Ringlet_1
WEST	Ringlet_2
WEST	Ringlet_3
EAST	Ringlet_1
EAST	Ringlet_3

Long (B-2), Table 1.

- *using filtering databases to perform traffic channel blocking and R-APS channel blocking*

It was also well known in the art before the alleged invention of the '807 Patent to use filtering databases to perform traffic channel blocking and ring automatic protection switching (R-APS) channel blocking. Traffic channel blocking and R-APS channel blocking were defined and extensively taught in the ITU-T G.8032/Y.1344 standards, as well as other references. *See, e.g.*, Prior Art G.8032 (B-1), Amd1 pp. 1-7, Amd1 FIG. 9-5, 2008 p. 11, 2008 FIG. 10-1; Long (B-2), [0002-0004], [0010-0031], [0055], FIG. 1; Im (B-4), pp. 1-2. Using filtering databases to perform

traffic channel blocking and R-APS channel blocking was also well known in the art. For example, Prior Art G.8032 provides:

9.4 Traffic channel blocking

Blocking traffic is supported by excluding the connection point from the ETH_FF functions for one or more VLAN IDs of the (service) traffic channel. This is equivalent to VID filtering as defined in [IEEE 802.1Q], c8.13.10. This results in blocking the transmission and reception of traffic on one ring port.

Prior Art G.8032 (B-1), 2008 p. 11.

9.5 R-APS channel blocking

R-APS channel VLAN traffic forwarding is always blocked at the same ports where traffic channel is blocked. It is supported by excluding the connection point from the ETH_FF function for the LAN ID of the R-APS traffic and is equivalent to perform VID filtering as defined in IEEE [802.1Q], c8.13.10. This only prevents R-APS messages received at one ring port from being forwarded to the other ring port. This, however, does not prevent R-APS messages, locally generated at the ERP control process, from being transmitted over both ring ports, and also allows R-APS messages received at each port to be delivered to the ERP control process.

Prior Art G.8032 (B-1), 2008 p. 11.

The Filtering Database supports queries by the Forwarding Process as to whether frames received by the Forwarding Process, with given values of destination MAC Address parameter and VID, are to be forwarded through a given potential transmission Port (8.6.1, 8.6.3, 8.6.4). It contains filtering information in the form of filtering entries that are either

- a) Static, and explicitly configured by management action; or
- b) Dynamic, and automatically entered into the Filtering Database by the normal operation of the bridge and the protocols it supports.

Prior Art G.8032 (B-1), IEEE 2005 p. 52. *See also:*

In this paper, we present the recent development of Ethernet ring protection and propose a new class of an Ethernet ring protection scheme so-called managed FDB that utilizes an extended APS payload to carry effective protection topology information. If a failure is sensed at an Ethernet ring node, the node generates a flip address list (FAL) from its FDB, which are associated with the failed port. Then, the node transmits an APS signaling frame including the FAL. The rest nodes then reconstruct their own FDB to provide protected MAC forwarding in the new network topology. The managed FDB provides protection which reduces the overshoots phenomenon and thus relieves the over-provisioning requirement. Our worst-case analysis suggests that the managed FDB technique may increase network utilization up to several times over the conventional FDB flush mechanism.

Im (B-04), p. 1.

- *performing a topology update of filtering databases*

Performing a topology update of filtering databases was well known in the art before the alleged invention of the '807 Patent. For example, Prior Art G.8032 discloses:

10.2.1 Protection switching – Link signal failure

A ring with no SF request will have a topology with the traffic channel blocked at the RPL and unblocked in all other ring links. In this situation, the detection of a signal-fail condition in a ring link will trigger protection switching as follows:

- A node detecting an SF condition on one of its ring ports will block the traffic channel and R-APS channel on the failed ring port.
- The node detecting an SF condition will transmit an R-APS message indicating SF over both ring ports. The R-APS SF message informs other ring nodes of the SF condition and that the traffic channel is blocked on one port. R-APS (SF) message shall be continuously transmitted by the node detecting the SF condition while this condition persists.
- Assuming the node was in an idle state before the SF condition occurred, upon detection of this SF condition, the node will trigger a local FDB Flush.
- A node accepting an R-APS (SF) message, without any local higher priority requests, will unblock any blocked port which does not have an SF condition. This action will unblock the traffic channel over the RPL.
- The node accepting an R-APS (SF) message, without any local higher priority requests, will stop transmission of R-APS messages.
- Assuming the node was in idle state before the SF condition occurred, a ring node accepting an R-APS (SF) message without a DNF indication will perform a Flush FDB operation.
- While a node remains in the protection state, the acceptance of subsequent R-APS (SF) messages does not re-trigger an FDB flush.

Prior Art G.8032 (B-1), 2008 p. 18.

11.1 Static and dynamic VLAN configuration

The combined functionality provided by the ability to configure Static VLAN Registration Entries in the Filtering Database, coupled with the use of the Restricted_VLAN_Registration control (11.2.3.2.3) and the ability of GVRP to dynamically create and update Dynamic VLAN Registration Entries, offers the following possibilities with respect to how VLANs are configured on a given Port:

c) Combined static and dynamic configuration. The static configuration mechanisms are used in order to configure some VLAN membership information; for other VLANs, GVRP is relied on to establish the configuration. The GVRP management controls are set to enable the operation of the GVRP protocol on that Port.

Prior Art G.8032 (B-1), IEEE 2005 p. 82. *See also:*

When a node detects a link failure on a certain port, the node, according to its maintained Port-Ringlet table, queries for the Ethernet ring instance ID information corresponding to that port, obtains the ID information of the Ethernet ring instance corresponding to that port that is affected by the failed link, and carries this ID information in this field. Their total number is saved in the Number of Ringlets field.

Then, the node clears forwarding table entries related to all VLANs in the service VLAN group protected by the aforementioned Ethernet ring instance affected by the failed link, blocks the failed port of that Ethernet ring instance, and sends the aforementioned R-APS packet to an adjacent node.

Long (B-2), [0067-0068].

In some embodiments, the control packets include a topology change flag. The first or second master bridge sets the topology change flag in at least one control packet transmitted following conversion of the first or second master bridge first port from a blocking state to a non-blocking state or from a non-blocking state to a blocking state. The other bridges in the first and second rings each store bridging tables, and, upon receipt of a control packet having a set topology change flag, each of the other bridges flushes its stored bridging table.

Hsu (B-3), [0016].

ITU-T G.8032 is being developed to define the APS (automatic protection switching) protocol data units for protection switching in Ethernet ring topologies. In order to form a loop-free topology according to Ethernet networking requirements, an Ethernet ring requires one or a couple of ports of a link forced in the 'block' state. On failure, a well-defined OAM functional set such as in ITU-T Y.1731 can be used to detect failed links and ports. To recover the ring from a failure on a link or on a node in a ring topology, the APS functions change blocking port to the failed node, forming a new topology of the network. Subsequently, all entries in FDBs in all ring nodes are 'flushed' so as to have no information. Subsequently, ring nodes newly start MAC learning processes. During these processes, Ethernet ring nodes broadcast received packets with unknown DA (Destination Address) through all ports. This protection method requires over-provisioning due to introduction of overshoots on protection.

Im (B-4), p. 1.

- *Utilizing VLAN identifiers to block or forward traffic in major rings and/or sub rings*

Major rings, sub rings, and utilizing VLAN identifiers to block or forward traffic in major rings and sub rings was well known in the art before the alleged invention of the '807 Patent. For example, Prior Art G.8032 discloses:

For the sub-ring, the connectivity at the interconnection node is provided between a sub-ring link and the domain of another network; in the example illustrated in Figure 9-5, this network corresponds to ring ERP1. An R-APS virtual channel provides R-APS connectivity between this interconnection node and the other interconnection node of the same sub-ring, over the network. The R-APS virtual channel may follow the same path as the traffic channel over the network. ERP control function of the sub-ring is capable to receive and insert R-APS messages over the R-APS virtual channel.

R-APS messages of this sub-ring that are forwarded over its R-APS virtual channel are broadcasted or multicasted over the interconnected network. For this reason, the broadcast/multicast domain of the R-APS virtual channel could be limited to the necessary links and nodes. For example, the R-APS virtual channel could span only the interconnecting rings or sub-rings that are necessary for forwarding R-APS messages of this sub-ring. Care must be taken to ensure that the local R-APS messages of the sub-ring being transported over the R-APS virtual channel into the interconnected network can be uniquely disambiguated from those of other interconnected ring R-APS messages. This can be achieved, for example, using separate VIDs for R-APS virtual channels of different sub-rings.

Prior Art G.8032 (B-1), Amd1 p. 5.




Figure VI.1 – Interconnected rings example

Prior Art G.8032 (B-1), Amd1 FIG. VI.1.

Figure 9-6 illustrates a situation where protection switching has occurred due to a signal-fail condition on the link between interconnection nodes C and D. The failure of this link triggers protection only on the ring it belongs to, in this case ERP1. The traffic and R-APS channels are blocked bidirectionally on the ports where the failure is detected and bidirectionally unblocked at the RPL connection point on ERP1. The traffic channels remain bidirectionally blocked at the RPL connection point on ERP2. This prevents the formation of a loop.

Prior Art G.8032 (B-1), Amd1 p. 3.




Figure 9-6 – Ethernet ring interconnection architecture – Signal fail condition on a link between interconnection nodes (multi-ring/ladder network)

Prior Art G.8032 (B-1), Amd1 FIG. 9-6.

9.5 R-APS channel blocking

R-APS channel VLAN traffic forwarding is always blocked at the same ports where traffic channel is blocked. It is supported by excluding the connection point from the ETH_FF function for the LAN ID of the R-APS traffic and is equivalent to perform VID filtering as defined in IEEE [802.1Q], c8.13.10. This only prevents R-APS messages received at one ring port from being forwarded to the other ring port. This, however, does not prevent R-APS messages, locally generated at the ERP control process, from being transmitted over both ring ports, and also allows R-APS messages received at each port to be delivered to the ERP control process.

Prior Art G.8032 (B-1), 2008 p. 11.

- *Flushing filtering databases*

Flushing filtering databases was well known in the art before the alleged invention of the '807 Patent. For example, Im discloses:

Ethernet ring protection (ERP) is a new technology based on OAM (operations, administration, and maintenance) being standardized by the ITU-T G.8032 working group. In this paper, we present the recent development of Ethernet ring protection which is called FDB (filtering database) flush scheme and propose a new Ethernet

ring protection technique introducing a managed FDB using APS to deliver information how to fix FDB selectively. We discuss the current development of the ERP technology at ITU-T and performance comparisons between different proposals.

Im (B-4), 1.

ITU-T G.8032 is being developed to define the APS (automatic protection switching) protocol data units for protection switching in Ethernet ring topologies. In order to form a loop-free topology according to Ethernet networking requirements, an Ethernet ring requires one or a couple of ports of a link forced in the ‘block’ state. On failure, a well-defined OAM functional set such as in ITU-T Y.1731 can be used to detect failed links and ports. To recover the ring from a failure on a link or on a node in a ring topology, the APS functions change blocking port to the failed node, forming a new topology of the network. Subsequently, all entries in FDBs in all ring nodes are ‘flushed’ so as to have no information. Subsequently, ring nodes newly start MAC learning processes. During these processes, Ethernet ring nodes broadcast received packets with unknown DA (Destination Address) through all ports. This protection method requires over-provisioning due to introduction of overshoots on protection.

Im (B-4), 1. *See also:*

9.6 FDB flush

A filtering database (FDB) flush consists of removing the learned MAC addresses of the ring ports from the node's filtering database. NOTE – The inclusion of the completion of FDB flush operation within the transfer time is for further study.

Prior Art G.8032 (B-1), 2008 p. 12.

A POSITA would have understood the disclosure of each of Long, Hsu, and Im in view of Ethernet and Ethernet Ring Protection (“ERP”) terminology, concepts and procedures generally known in the art, including those demonstrated by Prior Art G.8032.

Defendant is not aware of any secondary considerations supporting non-obviousness of the Asserted Claims of the ’807 Patent. To the extent Plaintiff seeks to introduce any alleged secondary considerations in the future, Defendant reserves the right to respond at that time.

For at least the reasons described herein as well as in the combinations of the prior art identified above and charted in the claim charts found at Exhibits B-1 to B-4, it would have been

obvious to one of ordinary skill in the art to combine any of a number of prior art references, including any combination of those identified above.

Defendant further identifies the commercial embodiments of the inventions described in the prior art charted in Exhibits B-1 to B-4, to the extent those systems or services were on sale or in public use before the priority dates of the '807 Patent.

A person having ordinary skill in the art would have been motivated to combine any of the references charted in attached Exhibits B-1 to B-4 for at least the reasons described above. A person having ordinary skill in the art at the time of filing of the '807 Patent would also have understood the references listed above, alone or in combination, to contain explicit and/or implicit teaching, suggestion, and/or rationales to combine them, including as further described above.

To the extent one of the above references does not disclose a specific limitation, it would have been obvious to combine that reference with one or more of the other references. A person of ordinary skill in the art at the time of the alleged invention would have combined the references because each generally relates to similar technology (e.g., Ethernet ring protection methods and associated systems), and the references could be combined, or substituted, according to known methods that would yield predictable results. Moreover, to the extent the '807 Patent identifies any needed improvement in the art, the references provide known techniques that would yield the required improvement, and given the finite number of techniques identified, it would have been obvious to try those approaches with a reasonable expectation of success.

The combinations of references disclosed in these contentions are exemplary and are not intended to be exhaustive. Additional obviousness combinations of the references identified here are possible, and Defendant may rely on such combination(s) in this litigation. In particular,

Defendant is currently unaware of Plaintiff's allegations with respect to the level of skill in the art and the qualifications of the typical person of ordinary skill in the art. Defendant is also unaware of the extent, if any, to which Plaintiff may contend that limitations of the claims at issue are not disclosed in the prior art identified by Defendant, and the extent to which Plaintiff will contend that elements not disclosed in the '807 Patent specifications would have been known to persons of skill in the art. And Defendant does not yet know how the Court will construe terms in the Asserted Claims of the '807 Patent. Defendant is also continuing their investigation of the large universe of prior art to identify potential prior art systems, publications related to those systems, and third parties that may have information about those systems. Plaintiff may also be in possession of prior art that Defendant may receive as discovery progresses in this case. Defendant reserves the right to amend and supplement these contentions to identify other prior art and combinations rendering the Asserted Claims of the '807 Patent anticipated and/or obvious.

2. Invalidity Under 35 U.S.C. § 112

Defendant identifies the following claim limitations that render the Asserted Claims of the '807 Patent invalid based on lack of enablement and written description under pre-AIA 35 U.S.C. § 112 ¶ 1 and/or are indefinite under pre-AIA 35 U.S.C. § 112 ¶ 2. All claims that depend from a claim containing a limitation identified below are invalid by way of their dependence. Defendant reserves the right to modify these contentions or further develop them as discovery progresses, including through expert disclosure and discovery.

a) Indefiniteness

The following claim elements are indefinite and, accordingly, the claims in which they appear are invalid:

- “setting ring ports of an Ethernet ring node constituting an Ethernet ring” (claim 13)

- “an Ethernet ring node constituting an Ethernet ring” (claim 13)
- “performing the traffic channel blocking and the R-APS channel blocking in a major ring for which the Ethernet ring node constituting in the FDBs” (claim 13)
- “for which the Ethernet ring node constituting” (Claim 13)
- “traffic channel blocking” (claim 13)
- “Ring Automatic Protection Switching (R-APS) channel blocking” (claim 13)
- “channel blocking” (claim 13)
- “registering VIDs for performing the traffic channel blocking and the R-APS channel blocking in a major ring for which the Ethernet ring node constituting” (claim 14)
- “for which the Ethernet ring node constituting” (claim 14)

b) Lack of Enablement and Written Description

Defendant contends that the specification as originally filed (including any specification to which priority is sought) fails to describe the full scope of the claimed inventions in sufficient detail such that one skilled in the art could reasonably conclude that the inventor had possession of the claimed invention, and/or lack enablement, including:

- “determining whether the Ethernet ring node is a ring protection link (RPL) owner node or an RPL neighbor node, wherein an RPL is located between the RPL owner node and the RPL neighbor node;” (claim 13)

For example, Defendant contends that the specification (including any specification to which priority is sought) fails to describe the full scope of the claimed inventions in sufficient detail such that one skilled in the art could reasonably conclude that the inventor had possession of an invention involving determining whether an Ethernet ring node is a RPL owner node or an RPL neighbor node, and/or fails to enable an invention where such a determination is made. No

determining operation, or manner of making such a determination, between two such nodes is described in the specification. Indeed, the term “RPL neighbor” or the like is not even used in the specification.

- “setting ring ports of an Ethernet ring node constituting an Ethernet ring, wherein a ring ID of the Ethernet ring is set” (claim 13) and “Ethernet ring node constituting an Ethernet ring” (claim 13)

For example, Defendant contends that the specification (including any specification to which priority is sought) fails to describe the full scope of the claimed inventions in sufficient detail such that one skilled in the art could reasonably conclude that the inventor had possession of an invention involving setting ring ports of an Ethernet ring node constituting an Ethernet ring, and/or fails to enable an invention where ring ports of an Ethernet ring node are set constituting an Ethernet ring. That is, the specification does not describe or enable a node constituting an Ethernet ring or ring ports of a node constituting an Ethernet, nor for analogous reasons operations for performing such a setting of ring ports of an Ethernet ring node constituting an Ethernet ring.

- “registering VLAN IDs (VIDs) for performing a traffic channel blocking and a Ring Automatic Protection Switching (R-APS) channel blocking in Filtering Databases (FDBs) of the ring ports of the Ethernet ring node” (claim 13)

For example, Defendant contends that the specification (including any specification to which priority is sought) fails to describe the full scope of the claimed inventions in sufficient detail such that one skilled in the art could reasonably conclude that the inventor had possession of an invention involving traffic channel blocking and Ring Automatic Protection Switching (R-APS) channel blocking, and/or fails to enable an invention where a traffic channel blocking and

an R-APS channel blocking are performed. That is, the specification does not describe a “traffic channel,” “traffic channel blocking,” an “R-APS channel,” or “R-APS channel blocking.”

- “performing a topology update of the Filtering Databases, wherein the topology update includes: registering VIDs for performing the traffic channel blocking and the R-APS channel blocking in a major ring for which the Ethernet ring node constituting in the FDBs” (claim 13)

For example, Defendant contends that the specification (including any specification to which priority is sought) fails to describe the full scope of the claimed inventions in sufficient detail such that one skilled in the art could reasonably conclude that the inventor had possession of an invention involving performing a topology update of Filtering Databases including registering VIDs for performing traffic channel blocking and R-APS channel blocking in a sub ring, and/or fails to enable an invention where a topology update is performed that includes registering VIDs for performing traffic channel blocking and an R-APS channel blocking in a sub ring in the FDBs. That is, the specification does not describe a “traffic channel” in a sub ring, “traffic channel blocking” in a sub ring in the FDBs, an “R-APS channel” in a sub ring in the FDBs, or “R-APS channel blocking” in a sub ring in the FDBs.

- “performing a topology update of the Filtering Databases, wherein the topology update includes: . . . registering VIDs for performing the traffic channel blocking and the R-APS channel blocking in a sub ring, a lower Ethernet ring of the major ring, in the FDBs” (claim 13)

For example, Defendant contends that the specification (including any specification to which priority is sought) fails to describe the full scope of the claimed inventions in sufficient detail such that one skilled in the art could reasonably conclude that the inventor had possession

of an invention involving performing a topology update of Filtering Databases including registering VID's for performing traffic channel blocking and R-APS channel blocking in a sub ring, and/or fails to enable an invention where a topology update is performed that includes registering VID's for performing traffic channel blocking and an R-APS channel blocking in a sub ring in the FDBs. That is, the specification does not describe a "traffic channel" in a sub ring, "traffic channel blocking" in a sub ring in the FDBs, an "R-APS channel" in a sub ring in the FDBs, or "R-APS channel blocking" in a sub ring in the FDBs.

- "registering VID's for performing the traffic channel blocking and the R-APS channel blocking in a major ring for which the Ethernet ring node constituting, and in a sub ring, a lower Ethernet ring of the major ring, in the FDBs" (claim 14)

For example, Defendant contends that the specification (including any specification to which priority is sought) fails to describe the full scope of the claimed inventions in sufficient detail such that one skilled in the art could reasonably conclude that the inventor had possession of an invention involving performing traffic channel blocking and R-APS channel blocking in a major ring for which the Ethernet ring node constituting and in a sub ring in the FDBs, and/or fails to enable an invention where VID's are registered for performing traffic channel blocking and R-APS channel blocking in a major ring for which the Ethernet ring node constituting and a sub ring in the FDBs. That is, the specification does not describe a "traffic channel" in a major ring or a sub ring in the FDBs, "traffic channel blocking" in a major ring or sub ring in the FDBs, an "R-APS channel" in a major ring or a sub ring in the FDBs, or "R-APS channel blocking" in a major ring or sub ring in the FDBs.

- “when a status of a ring port is changed, blocking or unblocking the VIDs of the traffic channel and the R-APS channel of a subset of VLANs assigned for protection, related to the ring port” (claim 15)

For example, Defendant contends that the specification (including any specification to which priority is sought) fails to describe the full scope of the claimed inventions in sufficient detail such that one skilled in the art could reasonably conclude that the inventor had possession of an invention involving a traffic channel and an R-APS channel, and/or fails to enable an invention where VIDs of a traffic channel and an R-APS channel are blocked or unblocked. That is, the specification does not describe a “traffic channel,” “traffic channel blocking,” an “R-APS channel,” or “R-APS channel blocking.” Furthermore, the specification does not describe blocking or unblocking VIDs of a traffic channel and an R-APS channel.

C. '887 Patent

1. Prior Art Patents and Publications/Additional Prior Art

Reference	Issue/Publication Date
Chinese Patent Publication No. 1001873359A (“Yang”)	October 27, 2010
U.S. Patent Application Publ. No. 2008/0148379 (“Xu”)	June 19, 2008
U.S. Patent Application Publ. No. 2004/0100953 (“Chen”)	May 27, 2004
U.S. Patent Application Publ. No. 2005/0238034 (“Gillespie”)	October 27, 2005
U.S. Patent Application Publ. No. 2002/0026531 (“Keane”)	February 28, 2002
Bryan Ford et. al., <i>Peer-to-Peer Communication Across Network Address Translators</i> (“Ford”)	April 10, 2005
IETF RFC2401 (Security Architecture for the Internet Protocol) (“Kent”)	November, 1998
U.S. Patent Application Publ. No. 2005/0235352 (“Staats”)	October 20, 2005
U.S. Patent Application Publ. No. 2008/0080524 (“Tsushima”)	April 3, 2008
U.S. Patent Application Publ. No. 2008/0198858 (“Townesley”)	August 21, 2008

Reference	Issue/Publication Date
U.S. Patent Application Publ. No. 2007/0157303 (“Pankratov”)	July 5, 2007
U.S. Patent Application Publ. No. 2008/0285575 (“Biswas”)	November 20, 2008
IETF RFC4110 (A Framework for Layer 3 Provider-Provisioned Virtual Private Networks (PPVPNs)) (“RFC4110”)	July, 2005
IETF RFC4026 (Provider Provisioned Virtual Private Network (VPN) Terminology) (“RFC4026”)	March, 2005
IETF RFC5766 (Traversal Using Relay NAT (TURN)) (“RFC5766”)	April, 2010
IETF RFC5245 (Interactive Connectivity Establishment (ICE)) (“RFC5245”)	April, 2010
IETF RFC5128 (State of Peer-to-Peer (P2P) Communication across Network Address Translators (NATs)) (“RFC5128”)	March, 2008
ITEF RFC3261 (Session Initiation Protocol (SIP)) (“RFC3261”)	June, 2002
IETF RFC5389 (Simple Traversal of User Datagram Protocol (UDP) (STUN)) (“RFC3261”)	October, 2008
Denise Donahue, <i>CCNP ROUTE 642-902 Quick Reference</i> (“Donahue”)	January, 2010

In addition, before the alleged invention of the ’887 Patent, the invention was made in this country by another inventor who had not abandoned, suppressed, or concealed it under 35 U.S.C. § 271(g)(2). The invention, through Meraki Auto VPN products embodying the invention, again under the apparent claim construction and application to Cisco’s accused products, as best they can be deciphered, was also known or used by others in this country under 35 U.S.C. pre-AIA § 102(a), to the extent the ’887 Patent is not entitled to its claimed priority date of October 29, 2010.

The Meraki Auto VPN products include:

- Meraki Teleworker VPN
- Meraki AP to AP VPN
- Meraki MX router with site-to-site (peer-to-peer) VPN

The Teleworker VPN and the AP to AP VPNs are L2 products. The Teleworker VPN included the Teleworker VPN Concentrator could be deployed as a software appliance virtual

machine designed to run on a hypervisor in the corporate data center. A Meraki MX Security Appliance could also serve as the concentrator for connected tunnels. The AP to AP VPNs could be deployed between MR series routers. Meraki conceived of its Teleworker VPN and AP to AP VPN solution, including the use of hole punching to allow for the formation of tunnels, before October 29, 2010 and worked diligently towards reducing the invention to practice, as corroborated by the product releases in 2011. By way of example only, □

“Requirements/Design/Braindump: VPN Service” (VPN.docx (v. 9.27.2010)) corroborates this earlier conception of the VPN solutions for these products. *See also* Presentation: Meraki VPN (September 2010) at 6 (describing tunnel formation between remote AP (VPN client) and VPN server/concentrator), at 7-8 (control dashboard) and 10-13 (use of hole punching in tunnel establishment and NAT traversal). An email string, dated January 31, 2011 and February 4, 2011, titled “Re: VPN read for testing” from Meraki (From Greg W. (gw@meraki.net)) with cc to Patrick Verkaik (patrick.verkaik@meraki.net) to Lucas Clara at Foster Garvey (clarl@foster.com) referenced the “wireless VPN feature” that “[c]reates Layer 2 bridge” was available for customer testing. An email dated January 28, 2011, titled “VPN feature turned on for your network”, from Meraki (From Greg W. (gw@meraki.net)) with cc to Patrick Verkaik (patrick.verkaik@meraki.net) to Chris Franks at American Baptist Home of the West² (cfranks@abhow.com) describes the VPN feature for the customer’s network being enabled, and including the VPN concentrator.

Meraki conceived of its site-to-site VPN solution, including the use of hole punching to allow for the formation of tunnels, for the router products before October 29, 2010 and worked diligently towards reducing the invention to practice, as corroborated by the MX product releases

² <https://web.archive.org/web/20110128083010/http://www.abhow.com/home.php>

in 2011. By way of example only, a July 2010 Board of Directors Meeting slide deck (board for new.pdf) is dated July 22, 2010 and corroborates this earlier conception of the site-to-site VPN functionality. The device that would use this functionality is referred to as a Network Appliance in the July 22, 2010 Board of Directors slide deck and would become the MX series router. This site-to-site VPN solution was L3 routing product that would use the cloud dashboard to distribute subnet information to the MX devices. Meraki announced the site-to-site VPN feature in late 2010 or early 2011 while working diligently towards reducing it to practice. For example, code for site-to-site (peer-to-peer) VPN is found in the cloud dashboard code dated February 24, 2011, which reflects the use of L3 routing and peer subnet information being provided by the dashboard to the MX routers. Router hardware code to consume that cloud information can be found in a March 3, 2011 router code update by Robert Shanks. The router hardware hole punching code, which can be traced back to the hole punching functionality of Meraki Teleworker VPN and AP to AP VPN solutions, was updated on March 4, 2011 by Patrick Verkaik to support the site-to-site (peer-to-peer) VPN solution for the MX routers. Work on the routing coding in general for the MX router is reflected in an initial code commit on June 2, 2010, with Robert Shanks handling. The coding for the dashboard to configure configurations in the router began at least by September 3, 2010, as reflected by an initial code commit, with Dan Aguayo handling. The coding for an initial implementation of the PunchNode and PunchRegistry, as well as defining the communication protocol between VPN mode (MX or MR) and the VPN registry, is reflected in code commits dated October 7, 2010 and updated in the aforementioned March 4, 2011 code commit to support L3 site-to-site (peer-to-peer) VPN.

In summary, code for hole punching to support tunnel formation was created by October 7, 2010 to support the Meraki Teleworker and AP to AP VPN products, and ultimately this code

progressed from the Teleworker solution to the broader use case in site-to-site VPN by March 2011. MX router and cloud dashboard code work began before Oct 29, 2010, and ultimately added the feature for site-to-site VPN in February and March 2011. This is the natural and diligent coding progression for a product solution like the MX router with site-to-site VPN functionality.

Documents and other evidence corroborating and providing further details regarding the prior art Meraki Auto VPN products and solution are referenced in Ex. C-11 and produced with these contentions at Bates CSI-OPT-PA_0005348 - CSI-OPT-PA_0005454, CSI-OPT-PA_0008235 - CSI-OPT-PA_0008248 and CSI-OPT-PA_0008282 - CSI-OPT-PA_0015013 , and as Cisco's investigation is ongoing, will be produced throughout discovery in accordance with the applicable rules. These documents and other evidence include, by way of example, deployment guides, release notes, product release notices, command references, guides such as deployment and configuration guides, press releases, on-line videos (e.g., YouTube), FAQ documents, third party articles and reviews, internal engineering design documents, source code, user manuals, sales invoices, advertisements and witness testimony.

To the extent the Asserted Claims of the '887 Patent are not anticipated by the prior art references and additional prior art identified above and in the claim charts attached as Exhibits C-1 to C-11, the Asserted Claims are rendered obvious by those references, either alone or in combination with other prior art references as identified below and in the attached claim charts and/or the knowledge of one skilled in the art.

In *KSR International Co. v. Teleflex Inc.*, 500 U.S. 398 (2007), the U.S. Supreme Court rejected the Federal Circuit's rigid "teaching, suggestion, or motivation" requirement in favor of a flexible, functional approach in which an explicit finding of a "motivation" to combine prior art references is not required to establish obviousness. The Supreme Court held that it is sufficient

that a combination of elements was “obvious to try” holding that, “[w]hen there is a design need or market pressure to solve a problem and there are a finite number of identified, predictable solutions, a person of ordinary skill has good reason to pursue the known option within his or her technical grasp.” *Id.* at 402; *see also Dystar Textilfarben GmbH v. C.H. Patrick Co.*, 464 F.3d 1356, 1368 (Fed. Cir. 2006) (explaining that when the “combination of references results in a product or process that is more desirable, for example because it is stronger, cheaper, cleaner, faster, lighter, smaller, more durable, or more efficient,” there exists a motivation to combine prior art references even when there is no explicit suggestion in the references themselves “[b]ecause the desire to enhance commercial opportunities by improving a product or process is universal—and even commonsensical”); *LeapFrog Enters., Inc. v. Fisher-Price, Inc.*, 485 F.3d 1157 (Fed. Cir. 2007) (applying *KSR* and holding that “one of ordinary skill in the art of children’s learning toys would have found it obvious to combine the Bevan device with the SSR to update it using modern electronic components in order to gain the commonly understood benefits of such adaption, such as decreased size, increased reliability, simplified operation, and reduced cost”).

Defendant contends that where the Asserted Claims are anticipated by a particular reference they are also rendered obvious by that same reference alone, or in combination with the other references, as show in the table below:

Primary Reference	In Combination with One or More of
Yang	Xu
Yang	Chen
Yang	Gillespie
Yang	Keane
Yang	Ford
Yang	Kent
Yang	Staats
Yang	Tsushima
Yang	Townsley
Yang	Pankratov

Primary Reference	In Combination with One or More of
Yang	Biswas
Yang	RFC4110
Yang	RFC4026
Yang	RFC5766
Yang	RFC5245
Yang	RFC5128
Yang	RFC3261
Yang	Donahue
Yang	Meraki Auto VPN Products
Xu	Yang
Xu	Chen
Xu	Gillespie
Xu	Keane
Xu	Ford
Xu	Kent
Xu	Staats
Xu	Tsushima
Xu	Townsley
Xu	Pankratov
Xu	Biswas
Xu	RFC4110
Xu	RFC4026
Xu	RFC5766
Xu	RFC5245
Xu	RFC5128
Xu	RFC3261
Xu	Donahue
Xu	Meraki Auto VPN Products
Chen	Yang
Chen	Xu
Chen	Gillespie
Chen	Keane
Chen	Ford
Chen	Kent
Chen	Staats
Chen	Tsushima
Chen	Townsley
Chen	Pankratov
Chen	Biswas
Chen	RFC4110

Primary Reference	In Combination with One or More of
Chen	RFC4026
Chen	RFC5766
Chen	RFC5245
Chen	RFC5128
Chen	RFC3261
Chen	Donahue
Chen	Meraki Auto VPN Products
Gillespie	Yang
Gillespie	Xu
Gillespie	Chen
Gillespie	Keane
Gillespie	Ford
Gillespie	Kent
Gillespie	Staats
Gillespie	Tsushima
Gillespie	Townsley
Gillespie	Pankratov
Gillespie	Biswas
Gillespie	RFC4110
Gillespie	RFC4026
Gillespie	RFC5766
Gillespie	RFC5245
Gillespie	RFC5128
Gillespie	RFC3261
Gillespie	Donahue
Gillespie	Meraki Auto VPN Products
Keane	Yang
Keane	Xu
Keane	Chen
Keane	Gillespie
Keane	Ford
Keane	Kent
Keane	Staats
Keane	Tsushima
Keane	Townsley
Keane	Pankratov
Keane	Biswas
Keane	RFC4110
Keane	RFC4026
Keane	RFC5766

Primary Reference	In Combination with One or More of
Keane	RFC5245
Keane	RFC5128
Keane	RFC3261
Keane	Donahue
Keane	Meraki Auto VPN Products
Ford	Yang
Ford	Xu
Ford	Chen
Ford	Gillespie
Ford	Keane
Ford	Kent
Ford	Staats
Ford	Tsushima
Ford	Townsley
Ford	Pankratov
Ford	Biswas
Ford	RFC4110
Ford	RFC4026
Ford	RFC5766
Ford	RFC5245
Ford	RFC5128
Ford	RFC3261
Ford	Donahue
Ford	Meraki Auto VPN Products
Staats	Yang
Staats	Xu
Staats	Chen
Staats	Gillespie
Staats	Keane
Staats	Ford
Staats	Kent
Staats	Tsushima
Staats	Townsley
Staats	Pankratov
Staats	Biswas
Staats	RFC4110
Staats	RFC4026
Staats	RFC5766
Staats	RFC5245
Staats	RFC5128

Primary Reference	In Combination with One or More of
Staats	RFC3261
Staats	Donahue
Staats	Meraki Auto VPN Products
Pankratov	Yang
Pankratov	Xu
Pankratov	Chen
Pankratov	Gillespie
Pankratov	Keane
Pankratov	Ford
Pankratov	Kent
Pankratov	Staats
Pankratov	Tsushima
Pankratov	Townsley
Pankratov	Biswas
Pankratov	RFC4110
Pankratov	RFC4026
Pankratov	RFC5766
Pankratov	RFC5245
Pankratov	RFC5128
Pankratov	RFC3261
Pankratov	Donahue
Pankratov	Meraki Auto VPN Products
Biswas	Yang
Biswas	Xu
Biswas	Chen
Biswas	Gillespie
Biswas	Keane
Biswas	Ford
Biswas	Kent
Biswas	Staats
Biswas	Tsushima
Biswas	Townsley
Biswas	Pankratov
Biswas	RFC4110
Biswas	RFC4026
Biswas	RFC5766
Biswas	RFC5245
Biswas	RFC5128
Biswas	RFC3261
Biswas	Donahue

Primary Reference	In Combination with One or More of
Biswas	Meraki Auto VPN Products
RFC4110	Yang
RFC4110	Xu
RFC4110	Chen
RFC4110	Gillespie
RFC4110	Keane
RFC4110	Ford
RFC4110	Kent
RFC4110	Staats
RFC4110	Tsushima
RFC4110	Townsley
RFC4110	Pankratov
RFC4110	Biswas
RFC4110	RFC4026
RFC4110	RFC5766
RFC4110	RFC5245
RFC4110	RFC5128
RFC4110	RFC3261
RFC4110	Donahue
RFC4110	Meraki Auto VPN Products
Meraki Auto VPN Products	Yang
Meraki Auto VPN Products	Xu
Meraki Auto VPN Products	Chen
Meraki Auto VPN Products	Gillespie
Meraki Auto VPN Products	Keane
Meraki Auto VPN Products	Ford
Meraki Auto VPN Products	Kent
Meraki Auto VPN Products	Staats
Meraki Auto VPN Products	Tsushima
Meraki Auto VPN Products	Townsley
Meraki Auto VPN Products	Pankratov
Meraki Auto VPN Products	Biswas
Meraki Auto VPN Products	RFC4026
Meraki Auto VPN Products	RFC5766
Meraki Auto VPN Products	RFC5245
Meraki Auto VPN Products	RFC5128
Meraki Auto VPN Products	RFC3261
Meraki Auto VPN Products	Donahue
Meraki Auto VPN Products	RFC4110

Defendant further contends that one of ordinary skill in the art, at the time of the alleged invention of the Asserted Claims, would have been motivated to combine the references disclosed herein in such a way to reach the alleged inventions. The teaching, suggestion, or motivation to combine these references, although not required, is explicitly or implicitly found in one or more of the following: the knowledge or common sense of one of ordinary skill in the art; the prior art references themselves and/or the prior art as a whole, including interrelated teachings of multiple prior art references; the subject matter acknowledged as prior art in the '887 Patent; the nature of the problem to be solved and the existence of similar improvements in similar applications; design incentives and other market forces, including the advantages of creating a superior and more desirable product and the effects of demands known to the design community or present in the marketplace; the ability to implement the alleged invention as a predictable variation of the prior art; improvements in similar devices; the interrelated teachings of multiple prior art references; any needs or problems known in the field addressed by the '887 Patent; and the number of identified, predictable solutions to the problem allegedly addressed by these patents. In addition, the simultaneous (and/or prior) inventions described above, and elsewhere in these contentions, is evidence that motivation to combine the concepts described in the various prior art references did, in fact, exist, and they were, in fact, combined. The combinations evidencing teachings, suggestions, and/or motivations to combine the prior art references in a way that renders the Asserted Claim obvious are merely exemplary. A person of ordinary skill would have access to the materials found in the cited references and attached exhibits and would have at least the ordinary creativity and skill to combine the cited references in ways not explicitly recited herein.

Further, Defendant hereby identifies additional motivation and reasons to combine the cited art: one of skill in the art would have been motivated to combine these references, because

these references relate to common objectives and subject matter. The identified prior art references relate to common objectives and share commonalities in terms of their general subject matter as well as the types of equipment, products, and approaches used, including in their descriptions of network-based communication in virtual network environments and/or in networks which may employ firewalls and/or or network address translation (NAT). *See, e.g.*, Yang at [0015]; Xu at [0009], Chen at [0035], Gillespie at [0010], [0011]; Keane at [0017], Ford p.1; Kent pp. 1-2; Staats at [0005], [0058], Tsushima at [0007], and Townsley at [0020].

Further, the Asserted Claims of the '887 Patent contain merely a duplication of concepts known in the prior art having no patentable significance given that no unexpected results occur (i.e., the patentee merely combines known prior art elements according to known methods to yield predictable results). For example, the prior art references of record describe well-known systems, devices and methods and standards relating to network-based communication in virtual network environments and/or in networks which may employ firewalls and/or NAT (e.g., Internet Engineering Task Force (IETF) RFC5128 (State of Peer-to-Peer (P2P) Communication across Network Address Translators (NATs)), IETF RFC3489 and RFC 5389 (Simple Traversal of User Datagram Protocol (UDP) (STUN)), IETF RFC761 (Transmission Control Protocol (TCP)), IETF RFC768 (UDP), IETF RFC3261 (Session Initiation Protocol (SIP)), RFC5766 (Traversal Using Relay NAT (TURN)), IETF RFC3931 (Layer 2 Tunneling Protocol v 3 (L2TPv3)), IETF RFC2401 and RFC2412 (Internet Protocol Security (IPsec)), IETF RFC5245 (Interactive Connectivity Establishment (ICE)), IETF RFC4026 (Provider Provisioned Virtual Private Network (VPN) Terminology), and Internet Protocol Version 4 (IPv4) and IPv6)) and their associated messaging (e.g., for synchronization, discovery, grant and the like), which are incorporated by reference herein as knowledge of a person of ordinary skill in the art (POSITA).

Additional evidence that there would have been a motivation to combine the prior art references identified above includes the interrelated teachings of multiple prior art references; the effects of demands known to the design community or present in the marketplace; the existence of a known problem for which there was an obvious solution encompassed by the Asserted Claims; the existence of a known need or problem in the field of the endeavor at the time of the alleged invention(s); and the background knowledge that would have been possessed by a person having ordinary skill in the art. For example, the prior art references are directed to solving the same problem. Thus, a skilled artisan seeking to solve this problem, would have looked to these cited references alone or in combination. Accordingly, the motivation to combine the teachings of the prior art references disclosed herein is found in the references themselves and: (1) the nature of the problem being solved, (2) the express, implied and inherent teachings of the prior art, (3) the knowledge of persons of ordinary skill in the art, (4) the prior art as of the priority date is generally within the same field of endeavor, or (5) the predictable results obtained in combining the different elements of the prior art.

To the extent Plaintiff contends that any reference contains multiple distinct embodiments, it would be obvious to combine elements of the distinct embodiments. A person would be motivated to make such a combination because the elements are found in the same reference and the reference as a whole is directed to the same topic or topics.

Numerous prior art references, including those identified above, reflect common knowledge and the state, scope, and content of the prior art before the alleged priority dates of the Asserted Patents. *See Graham v. John Deere Co.*, 383 U.S. 1, 35-36 (1966).

Examples of technical features that were well-known in the prior art as of the priority date of the Asserted Claims, including at least under Plaintiff's apparent interpretation of the claims as asserted in its December 3, 2025 contentions, as best as they can be understood, include:

Systems and methods for enabling end-to-end communication in virtual network environments were well-known and standardized. For example, IETF RFC4026 defines a virtual network:

VPN is a generic term that covers the use of public or private networks to create groups of users that are separated from other network users and that may communicate among them as if they were on a private network. It is possible to enhance the level of separation (e.g., by end-to-end encryption), but this is outside the scope of IETF VPN working group charters. This VPN definition is from [RFC2764]. In the [L3VPN-FRAME], the term VPN is used to refer to a specific set of sites as either an intranet or an extranet that have been configured to allow communication. Note that a site is a member of at least one VPN and may be a member of many. In this document, "VPN" is also used as a generic name for all services listed in Section 3.

IETF RFC4026 p. 5

Additionally, Donahue explains how it was well-known to connect spatially disparate routers using a managed VPN:

Cisco Easy VPN can simplify the deployment of all these services and policies to remote users. Easy VPN enables a server to push down VPN configuration to a client. It is a way to create site-to-site VPNs without manually configuring each remote router.

Cisco Easy VPN dynamically handles the following items:

- Negotiating VPN tunnel parameters
- Establishing the VPN tunnel based on those parameters

Donahue p. 115

The names and functions of the well-known components used in such systems included, for example, tunnel routers and tunnel mapping controllers. The functions of tunnel routers were also well-known to be executed by components termed tunnel peers, as explained by Chen:


In detail, the present embodiment provides a scheme to connect IPv6 sites over IPv4 networks via tunnels with dynamic performance optimisation. So-called peers (or Tunnel Peers) act as the endpoints of IPv6-to-IPv4 tunnels and as the routers in the IPv6 virtual network.

Chen at [0006]

The functions of tunnel routers were also well-known to be executed by components termed tunnel peers, as explained by Keane:

In detail FIG. 1 shows a general block diagram of a network 100, in accordance with an embodiment of the present invention. The network 100 may include a control system 175 with one or more network operations centers 170, a communication channel 120, one or more gateways 150-153, one or more local networks 160, 161, one or more hosts 154, 155, and a computer 101. The communication channel 120 may include a shared or base network, such as the Internet to facilitate communication and exchanges between the various entities depicted in the network 100 of FIG. 1.

Keane at [0105]



Keane at Fig. 1.

The functions of tunnel routers were also well-known to be executed by components termed edge routers, as explained by IETF RFC4026:

A CE-R is a router in the customer network interfacing the provider network. There are many reasons to use a router in the customer network; e.g., in an L3VPN using private IP addressing, this is the router that is able

to do forwarding based on the private addresses. Another reason to require the use of a CE-R on the customer side is that one wants to limit the number of MAC-addresses that need to be learned in the provider network.

A CE-R could be used to access both L2 and L3 services.

IETF RFC4026 p. 8

A PE-R is a L3 device that participates in the PSN (see Section 8) routing and forwards packets based on the routing information.

IETF RFC4026 p. 9

The functions of tunnel routers were also well-known to be executed by components termed nodes (*see, e.g.*, Townsley at [0004]), clients (*see, e.g.*, Ford p. 2; Pankratov at [0023]; and Gillespie at [0021] (includes “servers”)), peers (*see, e.g.*, Yang at [0013]), devices (*see, e.g.*, Staats at [0027]), SIP-enabled devices (*see, e.g.*, Xu at [0052]), and/or intermediate drivers (*see, e.g.*, Xu at [0056]), although many other terms for components performing the functions of tunnel routers were well-known.

The functions of tunnel mapping controllers were also well-known to be executed by components termed tunnel arbiters, as explained by Chen:

A dedicated server (which is also referred to as a network configuration control element), named Tunnel Arbiter (TA), is defined as the core component of the architecture, which makes decision for tunnelling and routing so that congested (or poorly behaving) paths among the peers would be excluded from the set of logical links. All the tunnel end points are in peer relationship, i.e. there is not server-client differentiation.

Chen at [0041]

The functions of tunnel routers were also well-known to be executed by components termed tunnel peers, as explained by Keane:

The control system 175 may also enable a third encrypted information flow through the communication channel 120 between the first gateway 150 and the second gateway 151. The control system 175 may enable the third encrypted information flow after the first gateway 150 and the second gateway 151 consent to enabling the third encrypted information flow.

Keane at [0107]

The functions of tunnel mapping controllers were also well-known to be executed by components termed VPN managers (*see, e.g.*, Townsley at [0036]), servers (*see, e.g.*, Ford p. 2; Pankratov at [0023]; and Gillespie at [0012] (“session control server”)), third-party devices (*see, e.g.*, Yang at [0063]), management centers (*see, e.g.*, Staats at [0049]), network controllers (*see, e.g.*, Xu at [0046]), and/or routers (*see, e.g.*, Tsushima at [0007]), although many other terms for components performing the functions of tunnel routers were well-known.

It was additionally well-known for components that perform the functions of tunnel routers to receive a message containing virtual network prefix information from components performing the functions of tunnel mapping controllers. In some systems and methods, virtual network prefix information was known to include IP addresses and lengths. This is explained by Pankratov and Keane:

When the client computer is initially configured, it also generates a cryptographic identity. It then establishes a control session with the server, secures it using cryptography mechanisms, authenticates the server and submits its public identity (public key) for registration. In return, preferably the server allocates a virtual IP address for the client computer and sends it to the client computer.

Once the client has obtained the virtual IP address, it uses it to configure a virtual network adapter. The virtual network adapter is an operating system component that appears as a physical network card to the system, but does not correspond to any actual hardware. Instead, preferably all traffic sent by the system to this adapter is forwarded to an application process. This process then inspects the traffic and either drops it (discards it) or forwards it to the other VPN member(s) via the tunnel(s).

The virtual adapter preferably is configured with a Virtual IP address and a non-trivial subnet that covers the entire Virtual IP address range as configured on the server.

Pankratov at [0109], [0110], [0114]

For example, to enable the third tunnel (step 540), the control system 175 may perform one or more of the following: update the partner lists of the first gateway 450 and the second gateway 451 to reflect mutual consent; provide an indication that a tunnel between the first and second gateways 450, 451 is authorized; provide real IP addresses for each of the gateways to permit a connection through a base network, such as the Internet; provide the virtual IP address of each gateway to the other gateway to enable a tunnel between the gateways; facilitate the establishment of one or more tunnels by providing out-of-band signaling to the first gateway 450 and the second gateway 451 through the first tunnel 425 and the second tunnel 426, respectively; determine one or more partner lists for one or more gateways 450, 451; provide configuration information for the network and/or for each gateway; exchange control information with the first gateway 450 and the second gateway 451 on the first tunnel 425 and the second tunnel 426, respectively; negotiate an encryption algorithm with each gateway; and negotiate an authentication technique. Moreover, the control system 175 may also monitor the status and performance of the tunnels established through the communication channel 120 (step 550).

Keane at [0137]; *see, also* Keane at [0144], [0110]

It was additionally well-known for components that perform the functions of tunnel routers to request virtual network prefix information from the components that perform the functions of tunnel mapping controllers. This is explained by Keane and Staats:

A registration entry for each TP include at least: a) a unique identifier (i.e., the TP identifier); b) the IPv4 address of the TP; c) an IPv6 address of the TP; d) the IPv6 address prefix that the TP holds; etc. The IPv6 address prefix (or address block) is a set of contiguous IPv6 addresses. An example for such an address prefix is 3ffe:3211::/32.

c) In step S 3, each TP gets the IPv4 addresses of the other TPs from the TA.

Keane at [0057] and [0059]

The process starts at block 200, where the management center 12 receives an activation key automatically transmitted from a device 14 connected to the network. . . .

Once the license key is received by the web-based management portal 90, the configuration of the device 14 may be provisioned via the web-based management portal 90. Setting the configuration of the device 14 may

include setting the IP address of the device 14, and setting the configurations for the firewall configuration, the intrusion prevention configuration, the anti-virus configuration, the content filtering configuration, the anti-spam configuration, the VPN configuration, the DHCP server configuration, the network management configuration, the network interface configuration, the VLAN configuration, the QOS configuration and any other device configurations. Each configuration provisioned for the device 14 may be stored in the database cluster 82. According to various embodiments, a default configuration may be selected for the device 14.

Staats at [0076]-[0077]

It was additionally well-known for components that perform the functions of tunnel mapping controllers to generate virtual network prefix information and transmit it along with a list of components that perform the functions of tunnel routers. In some systems and embodiments, the list of components included information as required to create tunnels, such as IP addresses. This is explained by Keane and Staats:

Accordingly, when the control system 175 determines that the first gateway 150 and the second gateway mutually consent to the third tunnel, the control system may then provide to the first and second gateways through the first and second tunnels, respectively, sufficient information to enable the third tunnel. The third tunnel may be enabled, for example, when the first and second gateways are provided sufficient information allowing them to establish this third tunnel through the communication channel 120. In one embodiment, the sufficient information includes the partner list for the first gateway and the partner list for the second gateway. Moreover, for each gateway listed on the partner list, the partner list may include, for example, a virtual IP address, a real IP address, and/or other information describing each gateway.

Keane at [0110]; *see, also* Townsley at [0021]

The DHCP server module 66 of the device 14 is operable to provide the device 14 with DHCP server functionality. For example, according to various embodiments, the DHCP server module 66 allows the device 14 to provide IP addresses and configuration parameters to network devices requesting this information using the DHCP protocol. IP address pools with characteristics such as default gateways, domain names, and DNS servers can be defined. Static assignments can also be defined based on MAC address.

Staats at [0039]

The IP and network interface configuration tool module 152 allows for the centralized management of the network configuration for each device 14. The centralized management of the network configuration may include, for example, managing IP Address, IP Types (static IP, DHCP, PPPOE), IP routing, Ethernet Trunking, VLANs, and QOS configuration.

Staats at [0067]

It was additionally well-known for components of systems enabling end-to-end communication in virtual network environments to determine if hole punching messages had been received by components performing the functions of tunnel routers. In some systems and methods, hole punching messages were used to create tunnels between components and traverse NATs. Many methods for hole punching were well known, such as UDP and TCP hole punching (*see, e.g.,* Ford p. 1). This is further explained by IETF RFC5766 and Ford:

A host behind a NAT may wish to exchange packets with other hosts, some of which may also be behind NATs. To do this, the hosts involved can use "hole punching" techniques (see [RFC5128]) in an attempt to discover a direct communication path; that is, a communication path that goes from one host to another through intervening NATs and routers, but does not traverse any relays.

IETF RFC5766 p. 3; *see also*, Xu at [0009]

A and B wait for outgoing connection attempts to succeed, and/or for incoming connections to appear. If one of the outgoing connection attempts fails due to a network error such as "connection reset" or "host unreachable," the host simply re-tries that connection attempt after a short delay (e.g., one second), up to an application-defined maximum timeout period.

Ford p. 8; *see also*, Pankratov at [0095], Gillespie at [0024]-[0025]

It was additionally well-known for components performing the functions of tunnel routers to establish tunnels, including tunnels based on received virtual network information. This is explained by IETF RFC4026 and Staats:

A tunnel is connectivity through a PSN that is used to send traffic across the network from one PE to another. The tunnel provides a means to transport packets from one PE to another. Separation of one customer's traffic from another customer's traffic is done based on tunnel multiplexers (see Section 8.5). How the tunnel is established depends on the tunnelling mechanisms provided by the PSN; e.g., the tunnel could be based on the IP-header, an MPLS label, the L2TP Session ID, or the GRE Key field.

IETF RFC4026 p. 14; *see also*, IETF RFC5245 p. 8, Pankratov at [0110], Kent p. 6, Yang at [0005]

The VPN module 64 of the device 14 is operable to provide the device 14 with VPN functionality. For example, according to various embodiments, the VPN module 64 provides the encryption protocol for the automatic building of a site to site VPN which is implemented as a secure tunnel that connects two different devices 14.

Staats at [0038]; *see also*, Staats at [0053], Townsley at [0029], Chen at [0003]

It was additionally well-known for components performing the functions of tunnel routers to perform direct communication through established tunnels. This is explained by IETF RFC5766 and Pankratov:

A host behind a NAT may wish to exchange packets with other hosts, some of which may also be behind NATs. To do this, the hosts involved can use "hole punching" techniques (see [RFC5128]) in an attempt to discover a direct communication path; that is, a communication path that goes from one host to another through intervening NATs and routers, but does not traverse any relays.

ITEF RFC5766 p. 3; *see also*, Ford pp. 3-4

The following provides a more detailed explanation of a server-mediated setup and maintenance of a tunnel between or among a set of peer computers according to the present invention. In one embodiment, it is assumed that a service provider provides a virtual private networking service whereby at least first and second peer nodes (typically, end user machines) use the help of a third node (a mediation server) to locate each other and to facilitate the connection between themselves. Once established, the connection between the peer nodes is direct; preferably, no end user traffic flows through the service provider's servers.

Pankratov at [0070]; *see also*, Townsley at [0030], [0036], Chen at [0129], [0135]

It was additionally well-known for components performing the functions of tunnel routers to perform direct communication through established tunnels by linking received virtual network prefix information. In some systems and methods, linking virtual network prefix information to established tunnels provided for direct communications between terminals located on disparate networks. *See, e.g.,* Keane at [0117]. (“the gateway 153 may permit a user to participate in one or more virtual private networks established over communication channel. . .”). Communication through established tunnels is further explained by Townsley, IETF RFC4026, and Donahue:

The LANs 152 include end nodes 154 a, 154 b, 154 c (among others, collectively referenced hereinafter as end node 154) and routing gateway nodes. The LANs 152 are connected to corresponding edge nodes 120 on subnetwork 110 via corresponding routing gateway nodes 150 a, 150 b, 150 c (among others, collectively referenced hereinafter as routing gateway, RG 150). A RG 150 is an intermediate network node on customer premises, which is connected to SP equipment edge nodes 120 via one or more attachment circuits. RG 150 includes a network address translation (NAT) process to communicate L3 traffic between subnetwork 110 and end nodes 154. Commercial off the shelf devices suitable to serve as RGs 150 include home wired and wireless routers, such as those sold by LINKSYS™, of Irvine, Calif., a division of CISCO SYSTEMS™, Inc., of San Jose, Calif., for wireless and Ethernet networking for home. Small Office, Home Office (SOHO) and small business users.

Townsley at [0035]; *see also*, Pankratov at [0109]-[0110], Staats at [0074], Keane at [0117], Tsushima at [0007]

A CE-R is a router in the customer network interfacing the provider network. There are many reasons to use a router in the customer network; e.g., in an L3VPN using private IP addressing, this is the router that is able to do forwarding based on the private addresses. Another reason to require the use of a CE-R on the customer side is that one wants to limit the number of MAC-addresses that need to be learned in the provider network.

A CE-R could be used to access both L2 and L3 services.

IETF RFC4026 p. 8

Cisco Easy VPN has two components: a server and a remote client. The *Easy VPN Server* can be a Cisco router, ASA Firewall, or Cisco VPN concentrator. It contains security policies and pushes those to remote clients. The *Easy VPN Remote* can be a Cisco router, ASA Firewall, a hardware client, or a software client. It contacts the server and receives policies from it to establish the IPsec tunnel. The steps to configure the headend for Easy VPN are

1. Allow IPsec traffic through the edge firewall or access list.
2. Create the IP address pool for the VPN clients.
3. Verify the IPsec VPN configuration.
4. Ensure that corporate devices have routes to the VPN subnets.

Donahue p. 116

Defendant is not aware of any secondary considerations supporting non-obviousness of the Asserted Claims of the '887 Patent. To the extent Plaintiff seeks to introduce any alleged secondary considerations in the future, Defendant reserves the right to respond at that time.

For at least the reasons described herein as well as in the combinations of the prior art identified above and charted in the claim charts found at Exhibits C-1 to C-11, it would have been obvious to one of ordinary skill in the art to combine any of a number of prior art references, including any combination of those identified above.

Defendant further identifies the commercial embodiments of the inventions described in the prior art charted in Exhibits C-1 to C-11, to the extent those systems or services were on sale or in public use before the priority dates of the '887 Patent.

A person having ordinary skill in the art would have been motivated to combine any of the references charted in attached Exhibits C-1 to C-11 for at least the reasons described above. A person having ordinary skill in the art at the time of filing of the '887 Patent would also have understood the references listed above, alone or in combination, to contain explicit and/or implicit teaching, suggestion, and/or rationales to combine them, including as further described above.

To the extent one of the above references does not disclose a specific limitation, it would have been obvious to combine that reference with one or more of the other references. A person of ordinary skill in the art at the time of the alleged invention would have combined the references because each generally relates to similar technology (e.g., systems and methods for enabling end-to-end communication in virtual network environments), and the references could be combined, or substituted, according to known methods that would yield predictable results. Moreover, to the extent the '887 Patent identifies any needed improvement in the art, the references provide known techniques that would yield the required improvement, and given the finite number of techniques identified, it would have been obvious to try those approaches with a reasonable expectation of success.

The combinations of references disclosed in these contentions are exemplary and are not intended to be exhaustive. Additional obviousness combinations of the references identified here are possible, and Defendant may rely on such combination(s) in this litigation. In particular, Defendant is currently unaware of Plaintiff's allegations with respect to the level of skill in the art and the qualifications of the typical person of ordinary skill in the art. Defendant is also unaware of the extent, if any, to which Plaintiff may contend that limitations of the claims at issue are not disclosed in the prior art identified by Defendant, and the extent to which Plaintiff will contend that elements not disclosed in the '887 Patent specifications would have been known to persons of skill in the art. And Defendant does not yet know how the Court will construe terms in the Asserted Claims of the '887 Patent. Defendant is also continuing their investigation of the large universe of prior art to identify potential prior art systems, publications related to those systems, and third parties that may have information about those systems. Plaintiff may also be in possession of prior art that Defendant may receive as discovery progresses in this case. Defendant reserves the right

to amend and supplement these contentions to identify other prior art and combinations rendering the Asserted Claims of the '887 Patent anticipated and/or obvious.

2. Invalidity Under 35 U.S.C. § 112

Defendant identifies the following claim limitations that render the Asserted Claims of the '887 Patent invalid based on lack of enablement and written description under post-AIA 35 U.S.C. § 112 ¶ 1. All claims that depend from a claim containing a limitation identified below are invalid by way of their dependence. Defendant reserves the right to modify these contentions or further develop them as discovery progresses, including through expert disclosure and discovery.

a) Lack of Enablement and Written Description

Defendant contends that the specification (including any specification to which priority is sought) fails to describe the full scope of the claimed inventions in sufficient detail such that one skilled in the art could reasonably conclude that the inventor had possession of the claimed invention, and/or lack enablement, including:

- “if the first virtual tunnel router has received a hole punching message, the first virtual tunnel router establishing a tunnel with the second virtual tunnel router based on the virtual network prefix information” (claim 1)

For example, in the '887 Patent, tunnels are not established between the first and second virtual tunnel routers based on virtual network prefix information. Rather, virtual network prefix information is linked to or associated with an already established tunnel. *See, e.g.*, '887 Patent, 5:39-47, 6:10-14, 8:26-35, FIG. 4.

- “the first virtual tunnel router receiving a response message containing virtual network prefix information created from the first virtual tunnel router from the tunnel mapping controller” (claim 2)

There is no written description support for or enablement of virtual network prefix information created “from the first virtual tunnel router.”

D. '101 Patent

1. Prior Art Patents and Publications/Additional Prior Art

Reference	Issue/Publication Date
U.S. Patent Application Publ. No. 2011/0225280 (“Delsesto”)	September 15, 2011
U.S. Patent Application Publ. No. 2013/0231081 (“Mo”)	September 5, 2013
WO 2010/094188 (“ZTE WO ‘188”)	August 26, 2010
Japanese Patent Application Publication No. JP 2002-374294 (“Hitachi JP ‘294”)	December 26, 2002
EP 1068574 (“Nortel EP ‘574”)	October 6, 2010
U.S. Patent No. 6,167,445 (“Gai”)	December 26, 2000
Time Based Access Lists Using Time Ranges	No later than April 26, 2010
U.S. Patent No. 7,676,559 (“Cuervo”)	March 9, 2010

In addition, the Cisco Secure Access Control Server (ACS), Cisco Security Manager (CSM), and Cisco Service Control Engine/Service Exchange Framework (SCE/SEF) systems and products (Exs. D-08, D-09, and D-10) were in public use, on sale, or otherwise available to the public before the effective filing date of the '101 Patent. *See* pre-AIA 35 U.S.C. § 102(a) and/or § 102(b). These systems include, for example, ACS Release 4.0 and 5.x providing time-of-day/day-of-week authorization, downloadable IP ACLs with Network Access Filters and ordered content, Network Access Restrictions; CSM 3.3.1 and 4.x providing centralized definition of time-range policy objects and tuple-based access rules, device clock and NTP settings; and SCE/SEF providing, time-of-day and peak/off-peak policies, parental-control time windows, and per-subscriber allowance/quota/time-based services.

These products constitute the items offered for sale, publicly used or otherwise available to the public. These items were offered for sale to Cisco customers (companies, Universities, and

governmental entities implementing data centers) and/or used and known or made known by Cisco Systems Inc. and/or to its customers (companies, Universities, and governmental entities implementing data centers) no later than their release dates shown below:

- Cisco Secure ACS 4.0: No later than 2006. *See, e.g.*, Cisco Secure ACS for Windows, Release 4.0 User Guide (© 2002–2006).
- Cisco Secure ACS 5.3: No later than September 2012. *See, e.g.*, Cisco Secure Access Control System 5.3 Data Sheet (Document C78-683481-02, dated 09/12).
- Cisco Security Manager 4.0: No later than June 17, 2010. *See, e.g.*, Release Notes for Cisco Security Manager 4.0 (published June 17, 2010; revised October 6, 2010).
- Cisco Security Manager 4.2sp1: No later than February 29, 2012. *See, e.g.*, Security Manager Release 4.2sp1 availability (release date 02/29/2012).
- Cisco Security Manager 3.3.1: No later than September 9, 2011. *See, e.g.*, End-of-Sale and End-of-Life Announcement for Cisco Security Manager 3.3 (EoL dated 09/09/2011) and User Guide for Cisco Security Manager 3.3.1.
- Cisco Service Control Engine 8000 (SCE 8000) and related SEF components: No later than 2008–2010. *See, e.g.*, Cisco SCE 8000 Data Sheet (© 2008–2010).
- Cisco Service Control Engine solution overview: No later than 2007. *See, e.g.*, Cisco Service Control Engine presentation materials (© 2007).
- Cisco IP NGN Mobility/Service Exchange Framework: No later than 2005. *See, e.g.*, Cisco IP NGN Mobility/SEF deck (© 2005).

Documents and other evidence corroborating and providing further details regarding these prior art Cisco systems and products are referenced in Exs. D-08, D-09, and D-10 and produced with these contentions at Bates CSI-OPT-PA_0000902 - CSI-OPT-PA_0004264, and, as Cisco's

investigation is ongoing, will be produced throughout discovery in accordance with the applicable rules. These documents and other evidence include, by way of example, deployment guides, release notes, product release notices, command references, guides such as deployment and configuration guides, press releases, on-line videos (e.g., YouTube), FAQ documents, third party articles and reviews, internal engineering design documents, source code, user manuals, sales invoices, advertisements and witness testimony.

To the extent the Asserted Claims of the '101 Patent are not anticipated by the prior art references identified above and in the claim charts attached as Exhibits D-1 to D-10, the Asserted Claims are rendered obvious by those references, either alone or in combination with other prior art references as identified below and in the attached claim charts and/or the knowledge of one skilled in the art.

In *KSR International Co. v. Teleflex Inc.*, 500 U.S. 398 (2007), the U.S. Supreme Court rejected the Federal Circuit's rigid "teaching, suggestion, or motivation" requirement in favor of a flexible, functional approach in which an explicit finding of a "motivation" to combine prior art references is not required to establish obviousness. The Supreme Court held that it is sufficient that a combination of elements was "obvious to try" holding that, "[w]hen there is a design need or market pressure to solve a problem and there are a finite number of identified, predictable solutions, a person of ordinary skill has good reason to pursue the known option within his or her technical grasp." *Id.* at 402; *see also Dystar Textilfarben GmbH v. C.H. Patrick Co.*, 464 F.3d 1356, 1368 (Fed. Cir. 2006) (explaining that when the "combination of references results in a product or process that is more desirable, for example because it is stronger, cheaper, cleaner, faster, lighter, smaller, more durable, or more efficient," there exists a motivation to combine prior art references even when there is no explicit suggestion in the references themselves "[b]ecause

the desire to enhance commercial opportunities by improving a product or process is universal—and even commonsensical”); *LeapFrog Enters., Inc. v. Fisher-Price, Inc.*, 485 F.3d 1157 (Fed. Cir. 2007) (applying *KSR* and holding that “one of ordinary skill in the art of children’s learning toys would have found it obvious to combine the Bevan device with the SSR to update it using modern electronic components in order to gain the commonly understood benefits of such adaption, such as decreased size, increased reliability, simplified operation, and reduced cost”).

Defendant contends that where the Asserted Claims are anticipated by a particular reference they are also rendered obvious by that same reference alone, or in combination with the other references, as show in the table below:

Primary Reference	In Combination with One or More of
Delsesto	Mo
Delsesto	ZTE WO ‘188
Delsesto	Hitachi JP ‘294
Delsesto	Nortel EP ‘574
Delsesto	Gai
Delsesto	Cuervo
Delsesto	Cisco Secure Access Control Server (ACS) System
Delsesto	Cisco Security Manager (CSM)
Delsesto	Cisco SCE/SEF (IMS)
Mo	Delsesto
Mo	ZTE WO ‘188
Mo	Hitachi JP ‘294
Mo	Nortel EP ‘574
Mo	Gai
Mo	Cuervo
Mo	Cisco Secure Access Control Server (ACS) System
Mo	Cisco Security Manager (CSM)
Mo	Cisco SCE/SEF (IMS)
ZTE WO ‘188	Delsesto
ZTE WO ‘188	Mo
ZTE WO ‘188	Hitachi JP ‘294
ZTE WO ‘188	Nortel EP ‘574
ZTE WO ‘188	Gai
ZTE WO ‘188	Cuervo

Primary Reference	In Combination with One or More of
ZTE WO '188	Cisco Secure Access Control Server (ACS) System
ZTE WO '188	Cisco Security Manager (CSM)
ZTE WO '188	Cisco SCE/SEF (IMS)
Nortel EP '574	Delsesto
Nortel EP '574	Mo
Nortel EP '574	ZTE WO '188
Nortel EP '574	Hitachi JP '294
Nortel EP '574	Gai
Nortel EP '574	Cuervo
Nortel EP '574	Cisco Secure Access Control Server (ACS) System
Nortel EP '574	Cisco Security Manager (CSM)
Nortel EP '574	Cisco SCE/SEF (IMS)
Gai	Delsesto
Gai	Mo
Gai	ZTE WO '188
Gai	Hitachi JP '294
Gai	Nortel EP '574
Gai	Cuervo
Gai	Cisco IOS Release 12.0(1) and/or its accompanying documentation, including "Time-Based Access Lists Using Time Ranges" (collectively and individually, "Cisco Time-Based Access Lists")
Gai	Cisco Secure Access Control Server (ACS) System
Gai	Cisco Security Manager (CSM)
Gai	Cisco SCE/SEF (IMS)
Cuervo	Delsesto
Cuervo	Mo
Cuervo	ZTE WO '188
Cuervo	Hitachi JP '294
Cuervo	Nortel EP '574
Cuervo	Gai
Cuervo	Cisco Secure Access Control Server (ACS) System
Cuervo	Cisco Security Manager (CSM)
Cuervo	Cisco SCE/SEF (IMS)
Cisco Secure Access Control Server (ACS) System	Delsesto
Cisco Secure Access Control Server (ACS) System	Mo

Primary Reference	In Combination with One or More of
Cisco Secure Access Control Server (ACS) System	ZTE WO '188
Cisco Secure Access Control Server (ACS) System	Hitachi JP '294
Cisco Secure Access Control Server (ACS) System	Nortel EP '574
Cisco Secure Access Control Server (ACS) System	Gai
Cisco Secure Access Control Server (ACS) System	Cuervo
Cisco Secure Access Control Server (ACS) System	Cisco Security Manager (CSM)
Cisco Secure Access Control Server (ACS) System	Cisco SCE/SEF (IMS)
Cisco Security Manager (CSM)	Delsesto
Cisco Security Manager (CSM)	Mo
Cisco Security Manager (CSM)	ZTE WO '188
Cisco Security Manager (CSM)	Hitachi JP '294
Cisco Security Manager (CSM)	Nortel EP '574
Cisco Security Manager (CSM)	Gai
Cisco Security Manager (CSM)	Cuervo
Cisco Security Manager (CSM)	Cisco Secure Access Control Server (ACS) System
Cisco Security Manager (CSM)	Cisco SCE/SEF (IMS)
Cisco SCE/SEF (IMS)	Delsesto
Cisco SCE/SEF (IMS)	Mo
Cisco SCE/SEF (IMS)	ZTE WO '188
Cisco SCE/SEF (IMS)	Hitachi JP '294
Cisco SCE/SEF (IMS)	Nortel EP '574
Cisco SCE/SEF (IMS)	Gai
Cisco SCE/SEF (IMS)	Cuervo
Cisco SCE/SEF (IMS)	Cisco Secure Access Control Server (ACS) System
Cisco SCE/SEF (IMS)	Cisco Security Manager (CSM)

Defendant further contends that one of ordinary skill in the art, at the time of the alleged invention of the Asserted Claims, would have been motivated to combine the references disclosed herein in such a way to reach the alleged inventions. The teaching, suggestion, or motivation to combine these references, although not required, is explicitly or implicitly found in one or more

of the following: the knowledge or common sense of one of ordinary skill in the art; the prior art references themselves and/or the prior art as a whole, including interrelated teachings of multiple prior art references; the subject matter acknowledged as prior art in the '101 Patent; the nature of the problem to be solved and the existence of similar improvements in similar applications; design incentives and other market forces, including the advantages of creating a superior and more desirable product and the effects of demands known to the design community or present in the marketplace; the ability to implement the alleged invention as a predictable variation of the prior art; improvements in similar devices; the interrelated teachings of multiple prior art references; any needs or problems known in the field addressed by the '101 Patent; and the number of identified, predictable solutions to the problem allegedly addressed by these patents. In addition, the simultaneous (and/or prior) inventions described above, and elsewhere in these contentions, is evidence that motivation to combine the concepts described in the various prior art references did, in fact, exist, and they were, in fact, combined. The combinations evidencing teachings, suggestions, and/or motivations to combine the prior art references in a way that renders the Asserted Claim obvious are merely exemplary. A person of ordinary skill would have access to the materials found in the cited references and attached exhibits and would have at least the ordinary creativity and skill to combine the cited references in ways not explicitly recited herein.

Further, Defendant hereby identifies additional motivation and reasons to combine the cited art: one of skill in the art would have been motivated to combine these references, because these references relate to common objectives and subject matter. The identified prior art references relate to common objectives and share commonalities in terms of their general subject matter as well as the types of equipment, products, and approaches used, including their descriptions regarding controlling packet traffic, including through use of time-based policies and rules. *See*,

e.g., Delsesto at [0003], [0004], [0022], [0052] and [0058]; Mo at Abstract, [0028], [0035], [0036], [0041]; ZTE WO '188 at Abstract, pp. 3-4, 7-8 and 10; Hitachi JP '294 at [0026], and [0034]; Nortel EP '574 at [0019], [0051], [0084], [0132]; Gai at Abstract, FIGS. 3-5; Time-Based Access Lists Using Time Ranges at pp. 1-12; Cuervo at Abstract, 4:4-55, 5:36-6:14, FIG. 2,

Further, the Asserted Claims of the '101 Patent contain merely a duplication of concepts known in the prior art and having no patentable significance given that no unexpected results occur (i.e., the patentee merely combines known prior art elements according to known methods to yield predictable results). For example, the '101 Patent itself admits that: (a) service provider networks control traffic flowing through an apparatus or system using a policy that is set to operate a network, (b) the traffic includes packets with 5-tuple (source IP address, destination IP address, source port, destination port, and protocol) and 7-tuple (source IP address, destination IP address, source port, destination port, protocol, input port number, and signature) information that is stored; (c) policies for controlling traffic are set based on this 5-tuple and 7-tuple information; (d) policies are set at a policy server; and (e) policies are executed in a policy execution point (PEP):

Traffic generally indicates a load that is applied to a communication apparatus or a system, and in a network, traffic indicates a quantity of data that are exchanged through the network. In order to provide a smooth network service or to apply a reasonable accounting model, a service provider monitors and controls a traffic situation. In this case, a policy is set to operate a network, and when transmission of data is requested from the outside of a server through the network, a service provider controls traffic according to the set policy.

Such traffic includes a packet in which information such as a 5-tuple (source IP address, destination IP address, source port, destination port, and protocol) and a 7-tuple (source IP address, destination IP address, source port, destination port, protocol, input port number, and signature) is stored, and a general policy is set based on information such as a 7-tuple that is included in a packet of traffic. A policy that is set at a policy server is executed in a policy execution point (PEP) and controls traffic.

'101 Patent, 1:21-39.

Additional evidence that there would have been a motivation to combine the prior art references identified above includes the interrelated teachings of multiple prior art references; the effects of demands known to the design community or present in the marketplace; the existence of a known problem for which there was an obvious solution encompassed by the Asserted Claims; the existence of a known need or problem in the field of the endeavor at the time of the alleged invention(s); and the background knowledge that would have been possessed by a person having ordinary skill in the art. For example, the prior art references are directed to solving the same problem. Thus, a skilled artisan seeking to solve this problem, would have looked to these cited references alone or in combination. Accordingly, the motivation to combine the teachings of the prior art references disclosed herein is found in the references themselves and: (1) the nature of the problem being solved, (2) the express, implied and inherent teachings of the prior art, (3) the knowledge of persons of ordinary skill in the art, (4) the prior art as of the priority date is generally within the same field of endeavor, or (5) the predictable results obtained in combining the different elements of the prior art.

To the extent Plaintiff contends that any reference contains multiple distinct embodiments, it would be obvious to combine elements of the distinct embodiments. A person would be motivated to make such a combination because the elements are found in the same reference and the reference as a whole is directed to the same topic or topics.

Numerous prior art references, including those identified above, reflect common knowledge and the state, scope, and content of the prior art before the alleged priority dates of the Asserted Patents. *See Graham v. John Deere Co.*, 383 U.S. 1, 35-36 (1966).

Non-Limiting examples of technical features that were well-known in the prior art as of the priority date of the Asserted Claims include:

- *Time-based policies that control traffic using a time condition that is provided independently from the traffic/determining and continuing to determine whether an execution time point of the policy arrives*

Time-based policies that control traffic using a time condition (e.g., time of day, month, etc.) that is provided independently from the traffic were well-known in the art prior to the alleged invention of the '101 Patent. It was also known to check this time condition, and re-check it, to determine if the execution time of the policy had arrived. For example, Delsesto (D-01) provides:

FIG. 5 is a message flow diagram illustrating an exemplary policy rule push for communicating policy information according to an embodiment of the subject matter described herein. Referring to FIG. 5, an IP CAN session may exist between UE 104 and PCEF node 116 enabling UE 104 to communicate with a host in Internet 110. User data plane 500, carrying traffic associated with UE 104's IP CAN session, may traverse service node 120. Upon the occurrence of specified conditions, PCRF node 114 may desire to push a policy rule update associated with UE 104's IP CAN session to service node 120. Service node 120 may not implement the full blown Gx protocol and thus utilization of the Gx protocol may not be possible. In accordance with embodiments of the subject matter described herein, PCRF node 114 may push a policy rule update associated with UE 104's IP CAN session to service node 120. For example, at step 1, a triggering event may occur at PCRF node 114. For example, a time-of-day condition may occur at PCRF node 114. In response to the triggering event, at step 2, PCRF node 114 may send a RAR message to service node 120. The message may include a session ID identifying the session between service node 120 and PCRF node 114 associated with UE 104's IP CAN session. The message may further include a charging rule(s) update for service node 120 to implement with respect to UE 104's IP CAN session. The message may be sent, for example, via a Gx-Lite RAR Diameter message.

Delsesto (D-01), [0058]. *See Also:*

Abstract: Embodiments of the invention provide a Policy and Charging Rules Function (PCRF) node and a method for authorizing a Quality of Service (QoS) for a wireless subscriber based on subscriber usage and time-of-day for a 3GPP-compliant packet data network.

Mo (D-02), Abstract.

FIG. 2 illustrates a process flow diagram illustrating the basic control flow when a PCC rule is created to handle time-of-day events. For example, rule tables can specify a QoS parameter that is generally applicable, and another rule table that has effect during a specific block of time, such as peak network load time for example 8:00 AM to 5:00 PM.

Mo (D-02), [0034].

Abstract: A service policy control method and system, wherein said method includes: a Policy and Charging Rules Function (PCRF) or a Policy and Charging Execution Function (PCEF) makes decisions on the basis of the time condition of service control policy of a service and the cumulative time for using the service, and obtains the decision information for service policy control. A PCEF implements the service policy control on the basis of the decision information.

ZTE WO '188 (D-03), Abstract.

That is, the policy control information (300) is composed of host information (301) indicating the IP addresses of the source host and destination host of the traffic to be subjected to QoS control, flow type information (302) indicating the flow characteristics of the traffic to be subjected to QoS control, QoS control information (303) indicating the QoS control content to be set for the communication device, setting time information (304) indicating the time to set and release QoS control for the communication device, and transfer information (305) indicating the transfer history of the policy control information (300). Here, the policy control information includes host information (301) indicating the IP addresses of the source host and destination host of the communication traffic, and is stored in the main memory (203) of the policy server for each piece of communication traffic. The host information (301) includes the IP address (301a) of the host sending the traffic and the IP address (301b) of the host receiving the traffic. The flow type information (302) includes a protocol number (302a) indicating a protocol such as TCP, UDP, or ICMP as the flow type of traffic to be subjected to QoS control, a source port number (302b) and a destination port number (302c) used in TCP or UDP, and a DSCP detected value (302d) in Diffserv (differentiated services, one of the QoS control technologies). The data format of the IP addresses (301a, 301b), protocol number (302a), and port numbers (302b, 302c) follows the notation specified in the Internet Standard Specifications (RFC). The DSCP detected value (302d) is expressed as a value between 0 and 63, and is used to associate with transmission priority in Diffserv. The QoS control information (303) includes transmission priority information (303a), guaranteed bandwidth information (303b), limited bandwidth information (303c), and a DSCP rewrite value (303d) in Diffserv, which are QoS control details. The transmission priority information (303a) is indicated by numbers 1 to 8, with 1 meaning the highest transmission priority and 8 meaning the lowest transmission priority. The guaranteed bandwidth information (303b) indicates, in Kbit/s, the value that guarantees the bandwidth of the line to be used. A value of 0 means no bandwidth is guaranteed. The limited bandwidth information (303c) indicates, in Kbit/s, the upper limit value for the bandwidth of the line to be used. A value of 0 means no bandwidth limitations are imposed. The DSCP rewrite value (303d) is expressed as a value between 0 and 63, and is used to associate with transmission priority in Diffserv. Of the policy control information (300), the host information (301), flow type information (302), and QoS control information (303) correspond to setting values for the communication device. Once this information

is set, the communication device performs QoS control according to the QoS control information (303) on traffic that matches the conditions of the host information (301) and the flow type information (302). The setting time information (304) includes QoS setting start time information (304a), which is the time when the host information (301), the flow type information (302), and the QoS control information (303) are set in the communication device, and QoS setting release time information (304b), which is the time when the set host information (301), the flow type information (302), and the QoS control information (303) are released. The QoS setting start release time information (304a, 304b) is expressed in years, months, days, hours, minutes, and seconds. The transfer information (305) includes transfer history information (305a) when the policy control information (300) is transferred to another policy server. The transfer history information (305a) indicates the IP address of the policy server that forwarded the policy control information (300). Of the policy control information (300), the host information (301), the flow type information (302), the QoS control information (303), and the setting time information (304) are set by the network administrator via the user interface of the policy server. Of the policy control information (300), the transfer information (305) is set by the control unit (202) of the policy server when transferring the policy control information.

Hitachi JP '294 (D-04), [0026].

The configuration manager may translate the service policies into a set of processing rules, with each processing rule containing a classifier and an associated action. The classifier generally specifies the data flow(s) and any conditions under which the action can be applied to a set of data bits transferred on the data flow(s). In the internet protocol (IP) environment, the source/ destination IP addresses, source/destination ports and the protocol type field together generally define an IP data flow supporting a specific remote access application.

Conditions may include matching of a specific variable defining the service policy. For example, a service policy may specify that data bits be treated a specific way between 9 AM - 5PM, in which case TIME is a variable and the condition is TIME = 9 AM - 5PM. As another example, data bits for a subscriber may be given lower priority if the aggregate bandwidth used by the subscriber is greater than T1, in which case BANDWIDTH is the variable and the condition is BANDWIDTH > T1.

Nortel EP '574 (D-05), [0018-0019].

Time-Based Access Lists Using Time Ranges

Feature Summary


It is now possible to implement access lists based on the time of day. To do so, you create a time range that defines specific times of the day and week. The time range is identified by a name and then referenced by a function, so that those time restrictions are imposed on the function itself.

Currently, IP and IPX extended access lists are the only functions that can use time ranges. The time range allows the network administrator to define when the **permit** or **deny** statements in the access list are in effect. Prior to this feature, access list statements were always in effect once they were applied. Both named or numbered access lists can reference a time range.

Time-Based Access Lists Using Time Ranges, at p. 1.

Also making reference to FIG. 3, the membership of each individual policy is determined with respect to a corresponding policy execution context, also referred to as Policy Equivalency Class (PEC), associated with a particular triggering event **212** and ensuing policy decisions **256** by specifying: policy equivalency classes, policy activation triggers, and policy equivalence class memberships. Policies **254** with the same triggering condition(s) and the same decisions(s) **256**, are aggregated into a single policy equivalency class. Therefore a policy equivalency class relates a target group of policy-managed entities **260** to at least one corresponding decision **256** to be applied upon the occurrence of a triggering event that satisfies the corresponding triggering condition(s). Determining a policy equivalency class aggregate to which a: policy-managed entity, policy, policy-managed flow of information, or any other part of a communications network belongs to, may depend on several factors, such as but not limited to: time-of-day, Quality-of-Service (QoS), network address, user's service prepaid status, etc.

Cuervo (D-07), 4:4-23.



Cuervo (D-07), FIG. 2.

- *Setting policies with a policy server (PS)/transmitting policies to a policy management system (PMS) and policy execution equipment (PEE) for execution*

It was well-known in the art at the time of the alleged invention to set policies with a policy server (PS) and to transmit policies to a policy management system (PMS) and connected policy execution equipment (PEE). For example, Delsesto (D-01) provides:

Methods, systems, and computer readable media for communicating policy information are provided. FIG. 1 is a network diagram illustrating an exemplary network environment for communicating policy information according to an embodiment of the subject matter described herein. Referring to FIG. 1, network environment 100 may include access network 102. Access network 102 may include nodes, functions, devices, and/or components for providing user equipment (UE) 104 access to services, functions, or devices in one or more networks. In one embodiment, access network 102 may be a radio access network (RAN). For example, access network 102 may be a global system for mobile communications (GSM) RAN (GRAN), a general packet radio service (GPRS) access network, a universal mobile telecommunications system (UMTS) RAN (UTRAN), an evolved UTRAN (eUTRAN), an Internet protocol (IP) connectivity access network (IPCAN), a code division multiple access (CDMA) network, an evolution-data optimized (EV-DO) network, a wideband CDMA (WCDMA) network, a high speed packet access (HSPA) network, an evolved HSPA (EHSPA+) network, or a

long term evolution (LTE) access network. Access network 102 may include one or more transceiver nodes 106 for communicating with UE 104. UE 104 may include a computer, a pager, a mobile phone, a smartphone, a wireless modem, or other devices through which a subscriber accesses network services.

Network environment 100 may further include a carrier network 108. Carrier network 108 may be utilized by UE 104 to access Internet 110. Carrier network 108 may include a bearer binding and event reporting function (BBERF) node 112. BBERF node 112 may be, for example, a service gateway (SGW) or a serving general packet radio service (GPRS) support node (SGSN). Carrier network 108 may further include a PCRF node 114. PCRF node 114 is a centralized node that can act as a policy decision point for carrier network 108. PCRF node 114 may take operator defined service policies, subscription information pertaining to a user, and other data into account to build policy decisions. Policy decisions may be formulated as policy control and charging (PCC) rules. PCC rules may contain information about user plane traffic expressed as a packet filter. A packet filter may take the form of an IP five-tuple specifying: (1) source IP address(es), (2) destination IP address(es), (3) source port number(s), (4) destination port number(s), and (5) application protocol(s) (e.g., transmission control protocol (TCP), user datagram protocol (UDP)). All IP packets matching a packet filter of a PCC rule may be designated an SDF.

Flow-based charging models may introduce the ability to charge for SDFs identified by service data flow filters according to specified charging rules. Charging rules may contain information that allows the filtering of traffic to identify packets belonging to a particular SDF (e.g., IP multimedia subsystem (IMS), file transfer protocol (FTP), browsing) and allow an operator to define how a particular SDF is to be charged (e.g., different media streams within a single packet data protocol (PDP) context.) Charging rules may be requested by a policy and charging enforcement function (PCEF) node (e.g., by a packet data network (PDN) gateway in an evolved packet system (EPS)), at bearer establishment, upon a specified trigger event, and/or upon bearer termination. Such a request may be made using a Gx reference point towards a PCRF.

Carrier network 108 may also include PCEF node 116. PCEF node 116 may serve as a policy enforcement point and may be placed in line between access network 102 and PCRF node 114. PCEF node 116 may be, for example, a gateway GPRS support node (GGSN) or a PDN gateway. As an enforcement point, PCEF node 116 may request and receive policy rules from PCRF node 114. Policy rules may take the form of, for example, Gx rules contained in credit control messages.

Carrier network 108 may also include subscriber data management (SDM) node 118. SDM node 118 may contain a comprehensive subscriber database, including information pertaining to subscribers' locations and Internet protocol information. SDM node 118 may be, for example, a home subscriber server (HSS), a subscription profile repository (SPR), or a user profile serving function (UPSF).

Carrier network 108 may further include service node 120. Service node 120 may include functionality for DPI, content-filtering, and/or web-optimization. DPI is the use of a packet's non-header information by a network entity that is not an endpoint for that packet. For example, service node 120 may include functionality enabling it to examine packets originating in Internet 110 and destined for UE 104. Content-filtering is the blocking of specified content based on analysis of the content itself rather than other criteria such as its source. For example, service node 120 may include functionality enabling it to analyze the content of packets originating in Internet 110 and destined for UE 104. Based on such content-analysis, service node 120 may filter or prevent the content from reaching UE 104. Web-optimization is provided to enhance a user's experience and may involve refining and/or altering content to better suit the hardware and/or software utilized by a particular user. For example, service node 120 may include functionality enabling it to optimize content originating in Internet 110 and destined for UE 104 based on, for example, the type of device UE 104 is. For example, service node 120 may detect video streaming from a source located in Internet 110 and destined for UE 104. In response, service node 120 may, for example, resize the video for optimal display on UE 104. A service node may or may not support the full blown Gx protocol as specified in 3GPP 29.212. For example, service node 120 may not include support for the full blown Gx protocol. In accordance with embodiments of the subject matter described herein, a service node that does not include support for the full blown Gx protocol may communicate with a PCRF node, enabling it to subscribe to SDF event notifications and install policy rules.

Delsesto (D-01), [0021-0026]. *See also:*

One such network architecture is shown ZTE WO '188:




Figure 1

ZTE WO '188 (D-03), FIG. 1.

The policy control process executed by the control unit (202) of the policy server 1 (100a) is started by the network administrator of the policy-based QoS management domain 1 (105a) inputting policy control information (300) for HTTP traffic to be sent from the host 1 (102a) to the host 2 (102b) via the user interface (204) of the policy server 1 (100a) (F401). This process may also be started by receiving policy control information (300) from a policy server managing another network via the network interface (205) (F402), or by receiving a policy control information setting request from the host 1 (102a) (F408).

Hitachi JP '294 (D-04), [0024].

When the time comes, the control units (202) of the policy server 1 (100a) and the policy server 2 (100b) read out the host information (301), the flow type information (302), and the QoS control information (303) from the policy control information (300) saved in the main memory (203), and configure the settings via the network interface (205) for the route of the network, that is, for the router 1 (101a) and the router 2 (101b) managed by the policy server 1 (100a), and for the router 4 (101d) and the router 6 (101f) managed by the policy server 2 (100b). The control unit (1002) of the communication devices (101a, 101b, 101d, 101f) that received the settings stores the host information (301), the flow type information (302), and the QoS control information (303) in the main memory (1003). The control unit (1002) also sets the same information in the routing processing unit (1004). The routing processing unit (1004) of the communication devices (101a, 101b, 101d, 101f) configured with this information initiates QoS control according to the QoS control information (303) for traffic matching the conditions of the host information (301) and the flow type information (302). Next, the control units (202) of the policy server 1 (100a) and the policy server 2 (100b) each read out the QoS setting release time (304b) from the policy control information (300) saved in the main memory (203), and wait until the clock unit (206) reaches the release time (304b). When the time arrives, the control unit (202) then releases the QoS settings for the communication devices (101a, 101b, 101d, 101f) present on the route of the network via the network interface (205) (F404).

Hitachi JP '294 (D-04), [0036].

FIG. 3 is a highly schematic block diagram of a computer network 300. The network 300 may be segregated into one or more network domains by a network administrator, such as network domains 302 and 304, as described below. The network 300 includes a plurality of entities, such as end stations and servers, interconnected by a plurality of intermediate devices, such as bridges, switches and routers. In particular, network 300 includes a plurality of switches 306-314 and routers 316-318 interconnected by a number of links 320a to 320f, which may be high-speed point-to-point or shared links. Each domain 302 and 304, moreover, includes at least one policy server 322 that is preferably connected to one or more intermediate devices, such as switch 308 and routers 316-318, by links 324a-324c.

The network 300 also includes one or more repositories, such as repository 326, that is preferably connected to each policy server 322 by a link 328. The repository 326 may be an organization-based directory server.




FIG. 3

Gai (D-06), 7:55-8:5 and FIG. 3.

FIG. 4 is a highly schematic, partial block diagram of policy server 322 in accordance with the preferred embodiment of the present invention. The policy server 322 is comprised of several components, including a policy translator 410 having one or more storage devices 412a-412c. Policy server 322 also includes a policy validation tool (PVT) 413 and a policy rule generating engine 414 that are each in communication with the policy translator 410, a device-specific filter entity 416 and a communication engine 418. As shown, the device-specific filter entity 416 communicates with both the policy rule generating engine 414 and the communication engine 418. The communication engine 418, moreover, is preferably configured to exchange messages with the intermediate devices (e.g., switches 306-314 and routers 316-318) of network 300. That is, communication engine 418 is connected to or includes conventional circuitry for transmitting and receiving messages across network links, such as links 324a-324c.




FIG. 4

Gai (D-06), 9:59-10:9 and FIG. 4.

FIG. 5 is a partial block diagram of an intermediate device, such as router 318, in accordance with the preferred embodiment of the present invention. Router 318 preferably includes a communication engine 510 that is coupled to a traffic management controller 512. The communication engine 510 is configured to exchange messages with the policy server 322. That is, communication engine 510, like the communication engine 418 at policy server 322, is similarly connected to or includes conventional circuitry for transmitting and receiving messages across the network 300. The traffic management controller 512, which includes a policy rule decoder 514, is coupled to several components and mechanisms. In particular, traffic management controller 512 is coupled to a packet/frame classifier 516, a traffic conditioner entity 518, a queue selector/mapping entity 520 and a scheduler 522. The traffic conditioner 518 also includes several sub-components, including one or more metering entities 524, one or more marker entities 526 and one or more shaper/dropper entities 528. The queue selector/mapping entity 520 and scheduler 518 operate on the various queues established by router 318 for its ports and/or interfaces, such as queues 530a-530e corresponding to an interface 532.





FIG. 5

Gai (D-06), 10:12-34 and FIG. 5.



Cuervo (D-07), FIG. 2.

- *converting, by the PS, the policy to PMS transfer syntax and transmitting the PMS transfer syntax to the PMS*

Converting policies and rule to PMS transfer syntax and transmitting the PMS transfer syntax to the PMS was well-known at the time of the alleged invention of the '101 Patent. For

example, several references described the use of the 3GPP-defined Gx interface and protocol³ for communication of policies and rules between nodes. *See, e.g.,* Delsesto (D-01), [0024], [0026], [0052], [0058]; Mo (D-02), [0035], [0041], [0045]; ZTE WO '188 (D-03), pp. 7-9. *See also* Gai (G-06) 7:11-29 (describing sending of rules as Common Open Policy Service (COPS) objects).

- *parsing, before the transmitting of the policy to the PEE, by the PMS, the policy to a structure that is appropriate for the PEE*

It was also well known in the art before the alleged invention of the '101 Patent to parse rules to ensure that they were in a structure that is appropriate for policy execution equipment. For example, Gai (D-06) discloses:

A computer network having multiple, dissimilar network devices includes a system for implementing high-level, network policies. The high-level policies, which are generally device-independent, are translated by one or more policy servers into a set of rules that can be put into effect by specific network devices. Preferably, a network administrator selects an overall traffic template for a given domain and may assign various applications and/or users to the corresponding traffic types of the template. Location-specific policies may also be established by the network administrator. The policy server translates the high-level policies inherent in the selected traffic template and location-specific policies into a set of rules, which may include one or more access control lists, and may combine several related rules into a single transaction. Intermediate network devices, which may have one or more roles assigned to their interfaces, are configured to request traffic management information from the policy server which replies with a particular set of transactions and rules. The rules, which may correspond to the particular roles assigned to the interfaces, are then utilized by the intermediate devices to configure their particular services and traffic management mechanisms. Other rules are utilized by the intermediate devices to classify packets with a particular priority and/or service value and to treat classified packets in a particular manner so as to realize the selected high-level policies within the domain.

Gai, Abstract.

It is a further object of the present invention to provide a method and apparatus for translating high-level policies into a form that may be understood and applied by numerous dissimilar network devices.

³ 3GPP Specifications and RFCs related to the Gx interface and protocol are incorporated herein, including 3GPP TS 29.212 (v.8.15.0 and v.7.13.0 both dated 2012-02-16), 3GPP TS 23.203 (v.8.13.00 and v.7.1.3.0 both dated 2012-03-09) and 3GPP 29.210 (v. 6.7.0. dated 2006-12-18) and any earlier and/or related releases and versions.

Gai, 5:55-59.

“FIG. 4 is a highly schematic, partial block diagram of policy server 322 in accordance with the preferred embodiment of the present invention. The policy server 322 is comprised of several components, including a policy translator 410 having one or more storage devices 412a-412c. Policy server 322 also includes a policy validation tool (PVT) 413 and a policy rule generating engine 414 that are each in communication with the policy translator 410, a device-specific filter entity 416 and a communication engine 418. As shown, the device-specific filter entity 416 communicates with both the policy rule generating engine 414 and the communication engine 418. The communication engine 418, moreover, is preferably configured to exchange messages with the intermediate devices (e.g., switches 306-314 and routers 316-318) of network 300. That is, communication engine 418 is connected to or includes conventional circuitry for transmitting and receiving messages across network links, such as links 324a-324c.”




FIG. 4

Gai, 9:59-10:9.

- *synchronizing a time-based policy that is stored at the PMS with the downloaded policy*

It was also well known in the art before the alleged invention of the '101 Patent to synchronize a time-based policy that is stored at the PMS with the downloaded policy. For example, it was known that various nodes in a network architecture stored policies and updated these policies with new policies when received. *See, e.g., Delsesto (D-01), [0005], [0009], [0021]-*

[0026], [0052], 0058]; Mo (D-02), [0021], [0029], [0035]; ZTE WO '188 (D-03), pp. 6-8, 10-12; Hitachi JP '294 (D-04), [0024], [0032], [0034]-[0036], [0040], [0043], [0044], FIGS. 1, 2, 4; Nortel EP '574 (D-05)), [0056], [0057], [0084]-[0089], [0102], [0103], [0110], [0111], [0118], [0119], [0135]; Gai (D-06), 3:61-4:12, 7:55-8:5, 10:12-10:34, 14:64-15:4, 17:49-62, 19:25-36, FIGS. 3-5; Cuervo (D-07), 5:27-35, 5:36-6:14, FIG. 2; *see also* U.S. Patent Application Publication No. 2008/0316922 to Riddle et al., [0063]; May 9, 2014 Non-Final Rejection in U.S. Application No. 13/645,090 at 3 (rejecting claim 6, which issued in as claim 5 of the '101 Patent, over Riddle et al.).

Defendant is not aware of any secondary considerations supporting non-obviousness of the Asserted Claims of the '101 Patent. To the extent Plaintiff seeks to introduce any alleged secondary considerations in the future, Defendant reserves the right to respond at that time.

For at least the reasons described herein as well as in the combinations of the prior art identified above and charted in the claim charts found at Exhibits D-1 to D-10, it would have been obvious to one of ordinary skill in the art to combine any of a number of prior art references, including any combination of those identified above.

Defendant further identifies the commercial embodiments of the inventions described in the prior art charted in Exhibits D-1 to D-10, to the extent those systems or services were on sale or in public use before the priority dates of the '101 Patent.

A person having ordinary skill in the art would have been motivated to combine any of the references charted in attached Exhibits D-1 to D-10 for at least the reasons described above. A person having ordinary skill in the art at the time of filing of the '101 Patent would also have understood the references listed above, alone or in combination, to contain explicit and/or

implicit teaching, suggestion, and/or rationales to combine them, including as further described above.

To the extent one of the above references does not disclose a specific limitation, it would have been obvious to combine that reference with one or more of the other references. A person of ordinary skill in the art at the time of the alleged invention would have combined the references because each generally relates to similar technology (e.g., packet filtering rules and policies and systems and associated equipment and methods using the same), and the references could be combined, or substituted, according to known methods that would yield predictable results. Moreover, to the extent the '101 Patent identifies any needed improvement in the art, the references provide known techniques that would yield the required improvement, and given the finite number of techniques identified, it would have been obvious to try those approaches with a reasonable expectation of success.

The combinations of references disclosed in these contentions are exemplary and are not intended to be exhaustive. Additional obviousness combinations of the references identified here are possible, and Defendant may rely on such combination(s) in this litigation. In particular, Defendant is currently unaware of Plaintiff's allegations with respect to the level of skill in the art and the qualifications of the typical person of ordinary skill in the art. Defendant is also unaware of the extent, if any, to which Plaintiff may contend that limitations of the claims at issue are not disclosed in the prior art identified by Defendant, and the extent to which Plaintiff will contend that elements not disclosed in the '101 Patent specifications would have been known to persons of skill in the art. And Defendant does not yet know how the Court will construe terms in the Asserted Claims of the '101 Patent. Defendant is also continuing their investigation of the large universe of prior art to identify potential prior art systems, publications related to those systems, and third

parties that may have information about those systems. Plaintiff may also be in possession of prior art that Defendant may receive as discovery progresses in this case. Defendant reserves the right to amend and supplement these contentions to identify other prior art and combinations rendering the Asserted Claims of the '101 Patent anticipated and/or obvious.

2. Invalidity Under 35 U.S.C. § 112

Defendant identifies the following claim limitations that render the Asserted Claims of the '101 Patent invalid based on lack of enablement and written description under pre-AIA 35 U.S.C. § 112 ¶ 1. All claims that depend from a claim containing a limitation identified below are invalid by way of their dependence. Defendant reserves the right to modify these contentions or further develop them as discovery progresses, including through expert disclosure and discovery.

a) Lack of Enablement and Written Description

Defendant contends that the specification (including any specification to which priority is sought) fails to describe the full scope of the claimed inventions in sufficient detail such that one skilled in the art could reasonably conclude that the inventor had possession of the claimed invention, and/or lack enablement, including:

- “determining ... whether an execution time point of the policy arrives” and “continuing, if an execution time point of the policy does not arrive, to determine an execution time point of the policy” (claims 1, 3, 6)

For example, to the extent it is Plaintiff's assertion in its infringement allegations that these limitations of “determining an execution time point of the policy arrives” and then “continuing, if an execution time point of the policy does not arrive to determine an execution time point of the policy” encompass an operation of simply waiting for an execution or scheduling a deployment time for a policy, then Defendant contends that the specification (including any specification to

which priority is sought) fails to describe the full scope of the claimed inventions in sufficient detail such that one skilled in the art could reasonably conclude that the inventor had possession of the claimed invention, and/or fails to enable the full scope of the claimed invention.

- “wherein the execution time point reflects a time condition that is provided independently from the traffic” (claims 1, 3, 6)

For example, Defendant contends that the specification (including any specification to which priority is sought) fails to describe the full scope of the claimed inventions in sufficient detail such that one skilled in the art could reasonably conclude that the inventor had possession of an invention involving an execution time point that reflects a time condition that is provided independently from the traffic, and/or fails to enable an invention where a time condition is provided independently from the traffic. No invention involving such a feature is disclosed in the specification. *See generally* ’101 Patent.

- “parsing ... by the PMS, the policy to a structure that is appropriate for the PEE” (claim 4)

For example, Defendant contends that the specification (including any specification to which priority is sought) fails to describe the full scope of the claimed inventions in sufficient detail such that one skilled in the art could reasonably conclude that the inventor had possession of an invention involving parsing a policy to a structure that is appropriate for the PEE, and/or fails to enable an invention where a policy is parsed to a structure that is appropriate for the PEE. That is, the specification describes no specific parsing operations or any structures that are the result of parsing.

- “synchronizing a time-based policy that is stored at the PMS with the downloaded policy” (claim 5)

For example, to the extent it is Plaintiff’s assertion in its infringement allegations that the limitation of “synchronizing a time-based policy that is stored at the PMS with the downloaded policy” encompass an operation of redeploying a modified policy, then Defendant contends that the specification (including any specification to which priority is sought) fails to describe the full scope of the claimed inventions in sufficient detail such that one skilled in the art could reasonably conclude that the inventor had possession of the claimed invention, and/or fails to enable the full scope of the claimed invention.

E. ’961 Patent

1. Prior Art Patents and Publications/Additional Prior Art

Reference	Issue/Publication Date
U.S. Patent Application Publ. No. 2014/0086253 to Yong et al. (“Yong”)	March 27, 2014
U.S. Patent Application Publ. No. 2008/0247406 to Figueira et al. (“Figueira”)	October 9, 2008
U.S. Patent Application Publ. No. 2015/0146727 to Ye et al. (“Ye”)	May 28, 2015
U.S. Patent No. 7,649,904 to Ghosh et al. (“Ghosh”)	January 19, 2010
U.S. Patent Application Publ. No. 2013/0266019 to Qu et al. (“Qu”)	October 10, 2013
U.S. Patent Application Publ. No. 2013/0311637 to Kamath et al. (“Kamath”)	November 21, 2013
U.S. Patent Application Publ. No. 2013/0250951 to Koganti et al. (“Koganti”)	September 26, 2013
U.S. Patent No. 8,166,205 to Farinacci et al. (“Farinacci”)	April 24, 2012
U.S. Patent No. 8,213,435 to Sajassi (“Sajassi”)	July 3, 2012
Internet Engineering Task Force (IETF) Internet-Draft, titled Overlay Transport Virtualization (draft-hasmit-otv-04), by Grover et al. (“Grover”)	February 23, 2013

In addition, the Cisco Data Center solution and products (Ex. E-8) was in public use, on sale or otherwise available to the public before the effective filing date of the ’961 Patent. *See* AIA 35 U.S.C. § 102(a)(1). Cisco Data Center includes the Cisco Nexus 7000 with NX-OS 5.0(3) or later) supporting Overlay Transport Virtualization (OTV) in combination with FabricPath (NX-

OS 5.1(1) or later) and/or or Nexus 1000V virtual distributed switch. These products constitute the items offered for sale, publicly used or otherwise available to the public. These items were offered for sale to Cisco customers (companies, Universities, and governmental entities implementing data centers) and/or used and known or made known by Cisco Systems Inc. and/or to its customers (companies, Universities, and governmental entities implementing data centers) no later than their release dates shown below:

- Cisco Nexus 7000 with NX-OS 5.0(3) (or later) supporting Overlay Transport Virtualization (OTV): No later than July 25, 2010. *See* Cisco Nexus 7000 Series NX-OS Release Notes, Release 5.0 (November 22, 2010).
- FabricPath (NX-OS 5.1(1) or later): No later than October 25, 2010. *See* Cisco Nexus 7000 Series NX-OS Release Notes, 5.1 (February 7, 2012).
- Nexus 1000V virtual distributed switch: No later than August 29, 2011. *See* Data Center Interconnect Design Guide for Virtualized Workload Mobility with Cisco, EMC and VMware (August 29, 2011).

Documents and other evidence corroborating and providing further details regarding the prior art Cisco Data Center product and solution are referenced in Ex. E-8 and produced with these contentions at Bates CSI-OPT-PA_0006886 - CSI-OPT-PA_0007505, CSI-OPT-PA_0007692 - CSI-OPT-PA_0008147, and CSI-OPT-PA_0008249 - CSI-OPT-PA_0008281 , and as Cisco's investigation is ongoing, will be produced throughout discovery in accordance with the applicable rules. These documents and other evidence include, by way of example, deployment guides, release notes, product release notices, command references, guides such as deployment and configuration guides, press releases, on-line videos (e.g., YouTube), FAQ documents, third party

articles and reviews, internal engineering design documents, source code, user manuals, sales invoices, advertisements and witness testimony.

To the *extent* the Asserted Claims of the '961 Patent are not anticipated by the prior art references identified above and in the claim charts attached as Exhibits E-1 to E-8, the Asserted Claims are rendered obvious by those references, either alone or in combination with other prior art references as identified below and in the attached claim charts and/or the knowledge of one skilled in the art.

In *KSR International Co. v. Teleflex Inc.*, 500 U.S. 398 (2007), the U.S. Supreme Court rejected the Federal Circuit's rigid "teaching, suggestion, or motivation" requirement in favor of a flexible, functional approach in which an explicit finding of a "motivation" to combine prior art references is not required to establish obviousness. The Supreme Court held that it is sufficient that a combination of elements was "obvious to try" holding that, "[w]hen there is a design need or market pressure to solve a problem and there are a finite number of identified, predictable solutions, a person of ordinary skill has good reason to pursue the known option within his or her technical grasp." *Id.* at 402; *see also Dystar Textilfarben GmbH v. C.H. Patrick Co.*, 464 F.3d 1356, 1368 (Fed. Cir. 2006) (explaining that when the "combination of references results in a product or process that is more desirable, for example because it is stronger, cheaper, cleaner, faster, lighter, smaller, more durable, or more efficient," there exists a motivation to combine prior art references even when there is no explicit suggestion in the references themselves "[b]ecause the desire to enhance commercial opportunities by improving a product or process is universal—and even commonsensical"); *LeapFrog Enters., Inc. v. Fisher-Price, Inc.*, 485 F.3d 1157 (Fed. Cir. 2007) (applying *KSR* and holding that "one of ordinary skill in the art of children's learning toys would have found it obvious to combine the Bevan device with the SSR to update it using

modern electronic components in order to gain the commonly understood benefits of such adaption, such as decreased size, increased reliability, simplified operation, and reduced cost”).

Defendant contends that where the Asserted Claims are anticipated by a particular reference they are also rendered obvious by that same reference alone, or in combination with the other references, as show in the table below:

Primary Reference	In Combination with One or More of
Yong	Figueira
Yong	Ye
Yong	Ghosh
Yong	Qu
Yong	Kamath
Yong	Koganti
Yong	Cisco Data Center
Yong	Farinacci
Yong	Sajassi
Yong	Grover
Figueira	Yong
Figueira	Ye
Figueira	Ghosh
Figueira	Qu
Figueira	Kamath
Figueira	Koganti
Figueira	Cisco Data Center
Figueira	Farinacci
Figueira	Sajassi
Figueira	Grover
Ye	Yong
Ye	Figueira
Ye	Ghosh
Ye	Qu
Ye	Kamath
Ye	Koganti
Ye	Cisco Data Center
Ye	Farinacci
Ye	Sajassi
Ye	Grover
Ghosh	Yong
Ghosh	Figueira
Ghosh	Ye
Ghosh	Qu

Primary Reference	In Combination with One or More of
Ghosh	Kamath
Ghosh	Koganti
Ghosh	Cisco Data Center
Ghosh	Farinacci
Ghosh	Sajassi
Ghosh	Grover
Qu	Yong
Qu	Figueira
Qu	Ye
Qu	Ghosh
Qu	Kamath
Qu	Koganti
Qu	Cisco Data Center
Qu	Farinacci
Qu	Sajassi
Qu	Grover
Kamath	Yong
Kamath	Figueira
Kamath	Ye
Kamath	Ghosh
Kamath	Qu
Kamath	Koganti
Kamath	Cisco Data Center
Kamath	Farinacci
Kamath	Sajassi
Kamath	Grover
Koganti	Yong
Koganti	Figueira
Koganti	Ye
Koganti	Ghosh
Koganti	Qu
Koganti	Kamath
Koganti	Cisco Data Center
Koganti	Farinacci
Koganti	Sajassi
Koganti	Grover
Cisco Data Center	Yong
Cisco Data Center	Figueira
Cisco Data Center	Ye
Cisco Data Center	Ghosh
Cisco Data Center	Qu
Cisco Data Center	Kamath
Cisco Data Center	Koganti

Primary Reference	In Combination with One or More of
Cisco Data Center	Farinacci
Cisco Data Center	Sajassi
Cisco Data Center	Grover
Farinacci	Yong
Farinacci	Figueira
Farinacci	Ye
Farinacci	Ghosh
Farinacci	Qu
Farinacci	Kamath
Farinacci	Koganti
Farinacci	Cisco Data Center
Farinacci	Sajassi
Farinacci	Grover
Sajassi	Yong
Sajassi	Figueira
Sajassi	Ye
Sajassi	Ghosh
Sajassi	Qu
Sajassi	Kamath
Sajassi	Koganti
Sajassi	Cisco Data Center
Sajassi	Farinacci
Sajassi	Grover
Grover	Yong
Grover	Figueira
Grover	Ye
Grover	Ghosh
Grover	Qu
Grover	Kamath
Grover	Koganti
Grover	Cisco Data Center
Grover	Farinacci
Grover	Sajassi

Defendant further contends that one of ordinary skill in the art, at the time of the alleged invention of the Asserted Claims, would have been motivated to combine the references disclosed herein in such a way to reach the alleged inventions. The teaching, suggestion, or motivation to combine these references, although not required, is explicitly or implicitly found in one or more of the following: the knowledge or common sense of one of ordinary skill in the art; the prior art

references themselves and/or the prior art as a whole, including interrelated teachings of multiple prior art references; the subject matter acknowledged as prior art in the '961 Patent; the nature of the problem to be solved and the existence of similar improvements in similar applications; design incentives and other market forces, including the advantages of creating a superior and more desirable product and the effects of demands known to the design community or present in the marketplace; the ability to implement the alleged invention as a predictable variation of the prior art; improvements in similar devices; the interrelated teachings of multiple prior art references; any needs or problems known in the field addressed by the '961 Patent; and the number of identified, predictable solutions to the problem allegedly addressed by these patents. In addition, the simultaneous (and/or prior) inventions described above, and elsewhere in these contentions, is evidence that motivation to combine the concepts described in the various prior art references did, in fact, exist, and they were, in fact, combined. The combinations evidencing teachings, suggestions, and/or motivations to combine the prior art references in a way that renders the Asserted Claim obvious are merely exemplary. A person of ordinary skill would have access to the materials found in the cited references and attached exhibits and would have at least the ordinary creativity and skill to combine the cited references in ways not explicitly recited herein.

Further, Defendant hereby identifies additional motivation and reasons to combine the cited art: one of skill in the art would have been motivated to combine these references, because these references relate to common objectives and subject matter. The identified prior art references relate to common objectives and share commonalities in terms of their general subject matter as well as the types of equipment, products, and approaches used, including their descriptions of virtual machines, virtual networks, and the associated systems and methods required to implementing them, including the use of overlay networks, encapsulated packets, tunnels, and

virtual switches. *See, e.g.*, Yong at [0004], [0005]; Figueira at Abstract, [0002]; Ye at Abstract, [0001], [0002]; Ghosh at 1:6-2:32; Qu at Abstract, [0004], [0005], [0006]; Kamath at Abstract, [0001]-[0007]; Koganti at Abstract, [0010], [0013], [0015], [0016]; Farinacci at Abstract, 1:5-7; Sajassi at 1:17-20, 1:28-2:62; Grover at pp. 3-6.

Further, the Asserted Claims of the '961 Patent contain merely a duplication of concepts known in the prior art and having no patentable significance given that no unexpected results occur (i.e., the patentee merely combines known prior art elements according to known methods to yield predictable results). For example, the prior art references of record describe and the '961 Patent itself admits as prior art well-known networks, systems, devices, methods, concepts, and components which are incorporated by reference herein as knowledge of a person of ordinary skill in the art (POSITA). These include, but are not limited to: server virtualization, VM (virtual machine)-based computing, network virtualization, overlay-based virtual networks, tunneling, VXLAN, NVGRE, tunnel end points, network virtualization managers, layer 2 and layer 3 switching, virtual switching instances, virtual network identification, encapsulation, domain/edge virtual ports, unicast packets, multicast packets, broadcast packets, flooding, and network virtualization domains. The '961 Patent itself explains in its Description of Related Art:

Recent trends of IT (information technology) include server virtualization, VM (virtual machine)-based computing, concentration of IT resources on data center, and network virtualization. Network virtualization is a technology where a plurality of virtual networks may be created on a physical network when necessary and can be deleted. This technology may be utilized to configure a network between virtual machines. A generally used overlay-based virtual network is a L2 network that may be realized on top of a L3 network in an overlay method. In the overlay-based virtual network, tunneling may be performed that is attaching a L3 IP (internet protocol) header to a L2 Ethernet frame being transmitted between virtual machines.

Most representative overlay-based network virtualization includes a VXLAN (virtual extensible local area network) and NVGRE (network virtualization using generic routing encapsulation). In a VXLAN or NVGRE, a virtual L2 network may

be created by connecting a tunnel between servers. For this purpose, a full-mesh tunnel must be set between TEPs (Tunnel End Points) being realized in a server or gateway.

'961 Patent at 1:20-40.

Additional evidence that there would have been a motivation to combine the prior art references identified above includes the interrelated teachings of multiple prior art references; the effects of demands known to the design community or present in the marketplace; the existence of a known problem for which there was an obvious solution encompassed by the Asserted Claims; the existence of a known need or problem in the field of the endeavor at the time of the alleged invention(s); and the background knowledge that would have been possessed by a person having ordinary skill in the art. For example, the prior art references are directed to solving the same problem. Thus, a skilled artisan seeking to solve this problem, would have looked to these cited references alone or in combination. Accordingly, the motivation to combine the teachings of the prior art references disclosed herein is found in the references themselves and: (1) the nature of the problem being solved, (2) the express, implied and inherent teachings of the prior art, (3) the knowledge of persons of ordinary skill in the art, (4) the prior art as of the priority date is generally within the same field of endeavor, or (5) the predictable results obtained in combining the different elements of the prior art.

To the extent Plaintiff contends that any reference contains multiple distinct embodiments, it would be obvious to combine elements of the distinct embodiments. A person would be motivated to make such a combination because the elements are found in the same reference and the reference as a whole is directed to the same topic or topics.

Numerous prior art references, including those identified above, reflect common knowledge and the state, scope, and content of the prior art before the alleged priority dates of the Asserted Patents. *See Graham v. John Deere Co.*, 383 U.S. 1, 35-36 (1966).

Examples of technical features that were well-known in the prior art as of the priority date of the Asserted Claims include:

- *Overlay-based network virtualization using L3 tunnels connecting tunnel end points in a full-mesh topology to create virtual L2 networks*

The practice of overlaying Layer 2 virtual networks onto Layer 3 infrastructure via tunneling was well known prior to the alleged invention of the '961 Patent. Such overlay networks were known under various names, including VXLAN (Virtual eXtensible Local Area Network), VPLS (Virtual Private LAN Service), NVGRE, and Overlay Transport Virtualization (OTV). These systems employed components that perform the functions of the “tunnel end point[s]” recited in claims 1, 7, and 13 of the '961 Patent. Such tunnel end points were variously referred to as VTEPs (VXLAN Tunnel End Points), NVEs (Network Virtualization Edges), PE devices (Provider Edge devices), and edge devices. For example, this is disclosed by Kamath, Yong, Figueira, and Cisco Data Center:

Overlay networks like Virtual eXtensible Local Area Network (VXLAN) connect geographically separated Layer-2 (L2) networks using tunnels. These are L2 over Layer-3 (L3) tunnels. L2 packets originated by a virtual machine (VM) in a VXLAN and destined to another VM or group of VMs in same VXLAN in another physical location are carried over L3 tunnels.

Kamath at [0003]; *see also* Kamath FIG. 3.

An overlay virtual network may provide communication among a set of tenant systems (TSs), where TSs may be VMs on a server or physical servers. An overlay virtual network may provide Layer 2 (L2) or Layer 3 (L3) services to the connected TSs via network virtualization edges (NVEs), where NVEs may be implemented as part of a virtual switch within a hypervisor, and/or physical switch or router. An NVE encapsulates ingress tenant traffic and sends the encapsulated traffic over a tunnel across an underlying network toward an egress NVE. An egress NVE at the

tunnel remote end point decapsulates [sic] the traffic prior to delivering the original data packet to the appropriate TS.

Yong at [0005].

In an overlay virtual network instance, any pair of NVEs 120 may be connected directly by an overlay tunnel 140, which may be a point-to-point (P2P), or point-to-multipoint (P2MP), or multipoint-to-point (MP2P) connection. The overlay tunnel 140 may transport encapsulated data traffic across the underlying network 130 between the pair of NVEs 120.

Yong at [0032]; *see also id.* at FIG. 1.

Another prior solution called Virtual Private LAN Services (VPLS) allows for the creation of L2VPNs over MPLS/IP networks.

Figueira at [0020].


This document introduces a Cisco innovative LAN extension technology called Overlay Transport Virtualization (OTV). OTV is an IP-based functionality that has been designed from the ground up to provide Layer 2 extension capabilities over any transport infrastructure: Layer 2 based, Layer 3 based, IP switched, label switched, and so on. The only requirement from the transport infrastructure is providing IP connectivity between remote data center sites. In addition, OTV provides an overlay that enables Layer 2 connectivity between separate Layer 2 domains while keeping these domains independent and preserving the fault-isolation, resiliency, and load-balancing benefits of an IP-based interconnection.

As of this writing, the Nexus 7000 is the only Cisco platform supporting OTV. All the technology and deployment considerations contained in this paper focus on positioning the Nexus 7000 platforms inside the data center to establish Layer 2 connectivity between remote sites. OTV support on Nexus 7000 platforms has been introduced from the NX-OS 5.0(3) software release. When necessary, available OTV features will be identified in the current release or mentioned as a future roadmap function. This document will be periodically updated every time a software release introduces significant new functionality....

Edge Device

The edge device ([Figure 1-1](#)) performs OTV functions: it receives the Layer 2 traffic for all VLANs that need to be extended to remote locations and dynamically encapsulates the Ethernet frames into IP packets that are then sent across the transport infrastructure.

Figure 1-1 OTV Edge Device



It is expected that at least two OTV edge devices are deployed at each data center site to improve the resiliency, as discussed more fully in [Multi-Homing, page 1-21](#).

Finally, the OTV edge device can be positioned in different parts of the data center. The choice depends on the site network topology. [Figure 1-1](#) shows edge device deployment at the aggregation layer.

Cisco Overlay Transport Virtualization Technology Introduction and Deployment Considerations (January 17, 2012) at 1-1, 1-2; *see also id.* at 1-33.

It was additionally well-known for such overlay networks to enable devices at geographically separated sites to communicate as if they were connected to the same local area network. In these systems, L3 tunnels transported L2 Ethernet frames across an intermediate IP network, preserving L2 functions such as MAC address learning and broadcast functionality. For example, this is disclosed by Ye and Ghosh:

Virtual private LAN service (VPLS): VPLS is a layer 2 multipoint virtual private network (VPN) that connects multiple customer sites in a bridged domain over a backbone network, such as a MPLS network. All customer sites in a VPLS instance (i.e., a particular VPN) appear to be on the same local area network (LAN) regardless of locations. VPLS emulates an Ethernet bridge over the backbone network and forwards Ethernet packets based on media access control (MAC) address or MAC address and virtual LAN (VLAN) tag.

Ye at [0026].

In some instances, a layer three (L3) network is used as an intermediate transport network between two or more L2 networks in order to allow communication between the L2 networks. In this type of configuration, the L3 network

transparently transports L2 communication between the L2 networks, thereby allowing the L2 networks to share an L2 service. Common protocols for transporting the L2 service through the intermediate L3 network are label switching protocols, such as Multi-protocol Label Switching (MPLS) protocols like Resource Reservation Protocol (RSVP) and the Label Distribution Protocol (LDP). In accordance with MPLS, a source device, such as a router connected to one of the L2 networks, can request a path through the intermediate network. This path, referred to as a Label Switched Path (LSP), defines one or more distinct, dedicated, and guaranteed paths through the network to carry MPLS packets from the source to the destination. The MPLS packets encapsulate the L2 communications, thereby effectively shielding the L3 network from the transported L2 information.

Ghosh at 1:26-44; *see also id.* at FIG. 1.

Cisco Overlay Transport Virtualization (OTV) is a MAC-in-IP method that extends Layer 2 connectivity across a core network infrastructure. OTV provides an overlay that enables Layer 2 connectivity between separate Layer 2 domains. Using MAC address-based routing and IP-encapsulated forwarding across a core network, OTV provides support for applications that require Layer 2 adjacency, such as clusters and virtualization. You deploy OTV on the edge devices in each site. OTV requires no other changes to the sites or the core network.

Cisco Nexus 7000 Series NX-OS Release Notes, Release 5.0 (November 22, 2010) at 10.

It was additionally well-known for tunnel end points in such networks to be connected in a full-mesh topology, enabling L2 connectivity between sites. This full-mesh architecture corresponds to the network topology recited in the '961 Patent, wherein each tunnel end point is connected with other tunnel end points via tunnels. For example, this is disclosed by Ye, Sajassi, and Koganti:

As backbone network 104 has a full mesh of tunnels (MPLS or GRE) between all the PE devices, a device may select at least one tunnel for each VSI at the PE device to communicate with a VSI of the same VPLS instance at an opposing PE device.

Ye at [0078]; *see also* Ye at FIG. 1.

Virtual Private LAN Service (VPLS) has recently emerged to meet the need to connect geographically dispersed locations with a protocol-transparent, any-to-any connectivity service. VPLS is an architecture that delivers Layer 2 service that in all respects emulates an Ethernet LAN across a WAN and inherits the scaling characteristics of a LAN. All sites in a VPLS instance appear to be on the same LAN, regardless of location. In other words, with VPLS, customers can

communicate as if they were connected via a private Ethernet LAN segment. Basically, VPLS offers a MPLS-based approach with multipoint connectivity for L2 services, i.e., multipoint Ethernet LAN services, often referred to as Transparent LAN Service (TLS). VPLS thus supports the connection of multiple sites in a single broadcast domain over a managed IP/MPLS network. Since a VPLS is normally provided over a service provider MAN/WAN network, it therefore needs to scale to accommodate a very large number of VPNs (e.g., a large number of customers, numerous services for each customer, and a large number of customer sites).

Sajassi at 2:13-31; *see also* Sajassi at FIGS. 3, 4.

As Internet traffic is becoming more diverse, virtual computing in a network is becoming progressively more important as a value proposition for network architects. In addition, the evolution of virtual computing has placed additional requirements on the network. For example, as the locations of virtual servers become more mobile and dynamic, it is often desirable that the network infrastructure can provide network overlay tunnels to assist the location changes of the virtual servers.

Koganti at [0013]; *see also* Koganti at FIG. 1A.

It was additionally well-known for tunnel end points to encapsulate outgoing L2 packets in L3 tunnel headers and to decapsulate incoming tunnel packets to recover the original L2 frames. This encapsulation and decapsulation functionality—which enables L2 packets to traverse an L3 network—corresponds to the “tunnel packet” processing recited throughout the claims of the '961 Patent. For example, this is disclosed by Qu and Farinacci:

As neither layer 2 nor layer 3 systems provide a completely adequate solution for a large cloud computing environment, various hybrid layer routing schemes have been proposed. These schemes are unworkable for various reasons. One example of such a hybrid layer routing scheme is the recently proposed virtual extensible local area network (VXLAN), which attempts to provide VLAN functionality over a hybrid layer 2 layer 3 network. Components on a VXLAN may send a layer 2 VLAN packet to a VXLAN tunnel endpoint (VTEP). The VTEP may encapsulate the VLAN packet in a VXLAN packet and transmit the resulting VXLAN packet over a layer 3 network. A VTEP in the destination network may receive the VXLAN packet, decapsulate the packet to obtain the original VLAN packet, and forward the VLAN packet over the layer 2 network. This system allows VLAN components in a layer 2 network to reach layer 2 components in a different layer 2 network.

Qu at [0005]; *see also* Qu at FIG. 1.

A method and system described herein support Layer 2 (L2) and Layer 3 (L3) virtual private networks (VPNs) over a L2 or Layer 3 (L3) infrastructure. The method and system are referred to herein as Overlay Transport Virtualization (OTV) since they provide an “overlay” method of doing virtualization versus traditional “in-the-network” type systems, in which multiple routing and forwarding tables are maintained in every device between a source and a destination. With OTV, state is maintained at the network edges, but is not required at other devices in a network site or in a core network. OTV operates at edge devices interposed between the network sites and the core network. The edge devices perform L2 learning and forwarding functions (similar to a traditional L2 switch) on their site-facing interfaces (internal interfaces) and perform IP based virtualization functions on their core-facing interfaces, for which an overlay network is created. The dual functionality of the edge device provides the ability to connect L2 networks, L3 networks, or hybrid (L2 and L3) networks together. OTV can also run at the same time as the network sites and the core network are operating “in-the-network” virtualization via VRF (Virtual Routing and Forwarding) or VNET (Virtual Network).

Farinacci at 2:65-3:19.

It was additionally well-known for such tunnel end points to reside at network edges, performing L2 switching and MAC address learning locally while transmitting encapsulated traffic through a network core that routes IP packets without awareness of the virtual network topology.

For example, this is disclosed by Grover:

OTV is a new “MAC in IP” technique for supporting L2 VPNs over an L2/L3 infrastructure. OTV provides an “over-the-top” method of doing virtualization among a large number of sites where the routing and forwarding state is maintained at the network edges, but not within the site or in the core....

Layer-2 traffic which requires traversing the overlay to reach its destination, is prepended with an IP header which ensures the packet is delivered to the edge boxes that provide connectivity to the Layer-2 destination in the original MAC header. As shown in figure 1, if a destination is reachable via Edge Device X2 (with a core facing IP address of IPB), other Edge Devices forwarding traffic to such destination will add an IP header with a destination IP address of IPB and forward the traffic into the core. The core will forward traffic based on IP address IPB, once the traffic makes it to Edge Device X2 it will be stripped of the overlay IP header and it will be forwarded into the site in the same way a regular bridge would forward a packet at layer-2. Broadcast or multicast traffic is encapsulated with a multicast header and follows a similar process.

Grover at p. 3.

- *A tunnel manager that collects tunnel end point information within a network and configures/establishes tunnel connections based on the collected information*

Components performing tunnel management functions were well-known and variously referred to as tunnel management modules, external control planes, join interfaces, and overlay control planes in the prior art. These components gathered information about tunnel end points within a network—including their IP addresses and virtual network associations—and used that information to establish tunnel connections between them. For example, this is disclosed by Koganti and Cisco Data Center:


One embodiment of the present invention provides a switch. The switch includes a tunnel management module, a packet processor, and a forwarding module. The tunnel management module operates the switch as a tunnel gateway capable of terminating an overlay tunnel. During operation, the packet processor, which is coupled to the tunnel management module, identifies in a data packet a virtual Internet Protocol (IP) address associated with a virtual tunnel gateway. This virtual tunnel gateway is associated with the switch and the data packet is associated with the overlay tunnel. The forwarding module determines an output port for an inner packet in the data packet based on a destination address of the inner packet.

Koganti at Abstract, [0008], [0100]; *see also* Koganti at FIG. 3A.

Join Interface

The Join interface (Figure 1-3) is used to source the OTV encapsulated traffic and send it to the Layer 3 domain of the data center network.

Figure 1-3 OTV Join Interface



The Join interface is a Layer 3 entity and with the current NX-OS release can only be defined as a physical interface (or subinterface) or as a logical one (i.e. Layer 3 port channel or Layer 3 port channel subinterface). A single Join interface can be defined and associated with a given OTV overlay. Multiple overlays can also share the same Join interface.



Note Support for loopback interfaces as OTV Join interfaces is planned for a future NX-OS release.

The Join interface is used by the edge device for different purposes:

- "Join" the Overlay network and discover the other remote OTV edge devices.
- Form OTV adjacencies with the other OTV edge devices belonging to the same VPN.
- Send/receive MAC reachability information.
- Send/receive unicast and multicast traffic.

Cisco Overlay Transport Virtualization Technology Introduction and Deployment Considerations (January 17, 2012) at 1-3 to 1-4; *see also id.* at 1-5 to 1-6.

OTV builds Layer 2 reachability information by communicating between edge devices with the overlay protocol. The overlay protocol forms adjacencies with all edge devices. Once each edge device is adjacent with all its peers on the overlay, the edge devices share MAC address reachability information with other edge devices that participate in the same overlay network. OTV discovers edge devices through dynamic neighbor discovery which leverages the multicast support of the core.

Cisco Nexus 7000 Series NX-OS OTV Configuration Guide (First Published 2010-08-06; Modified 09-08-2014) at 6; *see also id.* at 13, 14, 34, 35.

FabricPath Layer 2 IS-IS is a dynamic link-state routing protocol that detects changes in the network topology and calculates loop-free paths to other nodes in the network. Each FabricPath device maintains a link-state database (LSDB) that describes the state of the network; each device updates the status of the links that are adjacent to the device. The FabricPath device sends advertisements and updates to the LSDB through all the existing adjacencies. FabricPath Layer 2 IS-IS protocol packets do not conflict with standard Layer 2 IS-IS packets because the FabricPath

packets go to a different Layer 2 destination MAC address than that used by standard IS-IS for IPv4/IPv6 address families.

The system sends hello packets on the FabricPath core ports to form adjacencies. After the system forms IS-IS adjacencies, the FabricPath unicast traffic uses the equal-cost multipathing (ECMP) feature of Layer 2 IS-IS to forward traffic, which provides up to 16 paths for unicast traffic.

Within the FabricPath network, you use a single control plane protocol, Layer 2 IS-IS, for all unicast, multicast, and broadcast traffic. To use the basic FabricPath functionality, you do not need to configure Layer 2 IS-IS because you can use the default topology. The control plane Layer 2 IS-IS comes up and runs automatically when you enable FabricPath on the device.

Cisco Nexus 7000 Series NX-OS FabricPath Configuration Guide (December 7, 2011) at 2-4.

It was additionally well-known for such components to collect tunnel end point information—including addresses, virtual machine associations, and virtual network identifiers—and to use that information to establish tunnel connections. For example, this is disclosed by Kamath:

[A] system includes logic adapted for receiving, at a first end point station, an information exchange packet from each end point station in a virtual network having a specified virtual network identifier (VNID) and logic adapted for processing each received information exchange packet to retrieve information about connections at each end point station in the virtual network having the specified VNID, wherein each end point station either terminates or originates a tunnel shared by the first end point station in an overlay network.

Kamath at Abstract, [0008], [0032].

In a further approach, the method 500 may further include optional operation 510, where a Tunnel Information Base (TIB) may be assembled (such as from information in the information exchange packet, and other information available to or determinable by a point end station or a hypervisor operating on a point end station). In this way, the MAC address for each VM connected to each end point station and other information may be included in the TIB, which may be maintained in various forms. In one embodiment, the TIB may be a table of information related to tunnel end points which may comprise any or all of the following: a Tunnel ID, the specified VNID, MAC addresses of member VMs of the virtual network having the specified VNID, virtual port numbers corresponding to a hosting hypervisor, ARP entries learnt by the member VMs of the virtual network having the specified VNID, timers related to tunnel end points, a tunnel end point source IP Address, a point-to-point or multipoint nature of the tunnel, an IP Address of a peer tunnel end

point, shared forwarding tables of all the member VMs of the virtual network having the specified VNID, a default Gateway MAC address of each of the member VMs, multicast MAC and IP addresses used by all the member VMs of the virtual network having the specified VNID, and/or Virtual Port members or a Mask of member virtual ports of the virtual network having the specified VNID.

Kamath at [0070]; *see also* Kamath FIG. 5.

It was additionally well-known for such components to enable automatic discovery of tunnel end points and to exchange information used to configure tunnel connections-including address reachability and virtual network instance information. For example, this is disclosed by Ye, Figueira, Farinacci, Grover, and Sajassi:

A PE device uses a signaling protocol, such as label distribution protocol (LDP) or BGP, to establish a PW with another PE device in the same VPLS instance. Under manual or automatic configuration, each PE device allocates a PW label for an inbound VC LSP that the other PE device is to use to send packets to the PE device.

Ye at [0034].

In block 504, VPLS instance 107 is created. PE devices 108 may automatically discover each other using an automatic discovery protocol, such as BGP, over the first MP2MP LSP. PE devices 108 create VSIs 112 that are part of VPLS instance 107, associate VSIs 112 with VPLS instance 107 (for example to its VPLS ID), and allocate VSI labels 118, 122, and 124 (FIG. 1) to VSIs 112-1, 112-2, and 112-4 (FIG. 1), respectively. ACs may be manually associated with VSIs 112 on PE devices 108.

PE devices 108 may communicate over the first MP2MP LSP to establish PWs between VSIs 112. For example, PE devices 108-1 and 108-2 establish PW 114 (FIG. 1) between VSI 112-1 and 112-2, PE devices 108-1 and 108-4 establish PW 116 (FIG. 1) between VSI 112-1 and 112-4, and PE devices 108-2 and 108-4 establish PW 120 (FIG. 1) between VSI 112-2 and 110-2.

Ye at [0061], [0062]; *see also* Ye at FIG. 5.

FIG. 4 is a diagram illustrating an embodiment of an external control plane. In the example shown, external control plane 400 is used to provision the VSIs and the interconnections of VSIs with a full mesh of customer Service Instances (which are transported over PBT trunks), which are necessary to support multipoint-to-multipoint service. External control plane 400 allows service providers to automate the provisioning of L2VPNs over PBB, which would otherwise have to be manually and/or locally provisioned on each device.

Figueira at [0025]; *see also* Figueira at FIG. 4.

In one embodiment, the control plane uses an overlay IGP (Interior Gateway Protocol) to carry MAC addresses and IP addresses. The MAC addresses are those of hosts connected to the network and the IP next hops are the addresses of the edge devices through which the hosts are reachable through the core.

Farinacci at 3:40-45.

The overlay control plane provides auto-discovery of the Edge Devices that are members of an Overlay VPN. It also conveys Layer-2 unicast and multicast reachability information from a site to Edge Devices in other sites and the VLANs or layer-2 bridge domains being extended.

The MAC addresses that are locally connected to an Edge Device are advertised in the overlay URP to other Edge Devices in the VPN. Thus, MAC learning on the overlay is not based on data plane flooding, but is based on explicit advertisements of MAC addresses done by the overlay control plane. Similarly, the multicast groups that a site has receivers or sources for are advertised in the overlay MRP to other Edge Devices in the VPN.

Grover at pp. 9-10.

To limit the broadcast domain associated with a provider VLAN (or VPLS instance) over a set of PWs, a modified version of GARP (Generic Attribute Registration Protocol) VLAN Registration Protocol (GVRP) is run, which modified version is herein referred to as Extended GVRP (E-GVRP), among the PE devices (or n-PE devices only). GVRP is a known application defined in the IEEE 802.1Q standard that allows for the control of 802.1Q VLANs, i.e., 802.1Q-compliant VLAN pruning and dynamic VLAN creation on 802.1Q trunk ports. GVRP basically allows a switch to exchange VLAN configuration information with other GVRP switches, prune unwanted VLANs and their associated broadcast, multicast, and unicast traffic, and dynamically create and manage VLANs on switches connected through 802.1Q trunk ports.

The E-GVRP is essentially a compact GVRP that has a coding used to compress the VLAN information for up to 4K VLANs into a single Ethernet frame. In the E-GVRP, a jumbo Ethernet frame is utilized to carry the information needed for a large number of E-VLANs. Like GVRP, E-GVRP only runs among PE devices (or n-PEs); it does not interfere with and runs independently from each access domain's GVRP. Using E-GVRP each PE device may indicate to other PE devices what VPLS instances (E-VLANs) it is interested in on a given PW. Thus, the other PE devices may operate by only sending traffic from those VLANs over that PW, thereby limiting the scope of broadcast domain for each E-VLAN.

Sajassi at 5:18-44.

It was additionally well-known for such components to maintain and distribute forwarding information used to direct traffic through tunnel connections. For example, this is disclosed by Ghosh, Qu, and Yong:

FIGS. 3A-3B are block diagrams illustrating routing engine 26 and packet forwarding engine 30A of routing node 10 of FIG. 2 in further detail. Routing engine 26 includes high-level, control plane routing protocols 40A-40N (“routing protocols 40”). In the example of FIGS. 3A and 3B, routing protocols 40 include VPLS 40A and BGP 40N. Routing engine 26 may include other routing protocols not shown in FIG. 2. Routing protocols 40 interact with kernel 42 to update routing information base (RIB) 44 based on routing protocol messages received by routing node 20. In response, kernel 42 generates forwarding information in the form of FIBs 46A-46N based on the network topology represented in RIB 44. Kernel 42 then programs PFEs 30 to install copies of the FIBs as FIBs 48A-48N within their respective data plane 50, i.e., packet forwarding components having specialized forwarding chipsets.

Ghosh at 13:12-27.

FIG. 6 is a connection diagram of an embodiment of a hybrid layer 2/layer 3 network 600 with hypervisors. The network 600 comprises substantially the same components as network 400 in a different configuration. The network 400 comprises a plurality of host components. Host 1 may further comprise a hypervisor 620. Hypervisor 620 may be a process that manages multiple VMs. Hypervisor 620 may perform routing for its VMs and may have a routing table similar to the routing table discussed in FIG. 7 below. Hypervisor 620 may comprise VM 601, VM 602, and a virtual switch. Host component 2 may comprise a hypervisor 621, which may be substantially similar to hypervisor 620, and may comprise VM 604, VM 605, and a virtual switch. Host 1 may be connected to an attached routing component 603. Host 1 and Host 2 may be connected by a network of other components/nodes.

Qu at [0031].

FIG. 7 illustrates a flowchart of an example embodiment of a method 700 for performing control plane functions in an OVN. Method 700 may be implemented by an NVE 120 or an OVG 220 as discussed in FIG. 2. At step 710, method 700 may advertise the supported data plane encapsulation types and routes. The advertised message may be received by any peer (e.g. NVEs 120 or OVGs 220 from FIG. 2) that belongs to the same OVN. In an overlay virtual network instance, other peers may also advertise their supported encapsulation types and routes. At step 720, method 700 may check if a packet has been received from the OVN. At step 730, method 700 may check if the received packet is an advertised message from

other peers in regards to the supported data encapsulation types and routes by other peers. Upon the reception of an advertised message from a peer, method 700 may update a forwarding table with the peer's encapsulation types and routes as shown in step 740.

Yong at [0048].

- *Domain virtual switching instances (D-VSIs) corresponding to virtual network identifiers (VNIDs) that perform L2 switching operations, including obtaining source address information and forwarding packets based on destination addresses*

The functions of the claimed D-VSI were well known prior to the alleged invention of the '961 Patent. Virtual switching instances corresponding to virtual networks that could perform Layer-2 switching functions, including obtaining source address information and forwarding packets based on destination address information, as recited in the claims were well-known. For example, this is disclosed by Ye, Figueira, Sajassi, and Cisco Data Center:

Virtual switch instance (VSI): VSI is an Ethernet bridge function entity of a VPLS instance on a PE device. It forwards packets based on MAC address or both MAC destination and VLAN tag.

Ye at [0030].

A VSI is an IETF term for a bridge used to construct a VPLS service over an MPLS network. As defined in IETF RFC 4762 and IETF RFC 4761, a VSI provides standard MAC address learning, filtering, and forwarding actions, as defined in IEEE 802.1D and IEEE 802.1Q standards, which define how learning bridges work. A VSI is a learning bridge and can be used to bridge Ethernet traffic from any source. Thus, Ethernet traffic being carried by a PBT (PBB-TE) trunk can be forwarded by a VSI.

Figueira at [0018].


Conceptually, VPLS can be thought of as an emulated Ethernet LAN network with each Virtual Switch Instance (VSI) being analogous to a virtual Ethernet switch.

Sajassi at 2:32-34.

Internal Interfaces

To perform OTV functionality, the edge device must receive the Layer 2 traffic for all VLANs that need to be extended to remote locations. The Layer 2 interfaces, where the Layer 2 traffic is usually received, are named internal interfaces (Figure 1-2).

Figure 1-2 OTV Internal Interfaces



Internal interfaces are regular Layer 2 interfaces configured as access or trunk ports. Trunk configuration is typical given the need to concurrently extend more than one VLAN across the overlay. There is no need to apply OTV-specific configuration to these interfaces. Also, typical Layer 2 functions (like local switching, spanning-tree operation, data plane learning, and flooding) are performed on the internal interfaces. Figure 1-2 shows Layer 2 trunks that are considered internal interfaces which are usually deployed between the edge devices also.

Cisco Overlay Transport Virtualization Technology Introduction and Deployment Considerations
(January 17, 2012) at FIG. 1-2.

Nexus 1000V allows policy to move with a virtual machine during live migration, ensuring persistent network, security, and storage compliance, resulting in improved business continuance, performance management, and security compliance. Last but not least, it aligns management of the operational environment for virtual machines and physical server connectivity in the data center, reducing the total cost of ownership (TCO) by providing operational consistency and visibility throughout the network. It offers flexible collaboration between the server, network, security, and storage teams while supporting various organizational boundaries and individual team autonomy.

The Cisco Nexus 1000V Series comprises two components:

- Virtual Ethernet Module (VEM), a software switch embedded in the VMware ESX hypervisor. When comparing the N1KV virtual switch to a physical modular switch (like the Nexus 7000), each VEM can be thought as a physical linecard. Unlike multiple line cards within a single chassis, each VEM acts as an independent switch from a forwarding perspective. For what concerns scalability values, at the time of writing of this document up to 64 VEMs and 2048 virtual Ethernet (vEth) ports can be supported as part of the same virtual switch instance. It is worth noticing that this does not necessarily mean that 2048 virtual machines can be “attached” to a virtual switch instance, since usually more than one vNIC is

defined on each VM (and a 1:1 static mapping between vEth and vNIC is performed when the VM is connected to the N1KV switch).

- Virtual Supervisor Module (VSM), which manages networking policies and quality of service for virtual machines in concert with the VEM. Repeating the previous analogy, the VSM represents the switch supervisor establishing control plane communication with the various VEMs (the “linecards”). The Cisco Nexus 1000V Series requires a VSM high-availability deployment model much like dual supervisors in a physical chassis. Two VSMs are deployed in an active-standby configuration, with the first VSM functioning in the primary role and the other VSM functioning in a secondary role. If the primary VSM fails, the secondary VSM will take over. Note that unlike in cross-bar-based modular switching platforms, the VSM is not in the data path. General data packets are not forwarded to the VSM to be processed, but rather switched by the VEM directly.

Data Center Interconnect Design Guide for Virtualized Workload Mobility with Cisco, EMC and VMware (Last Update August 29, 2011) at 2-23.

It was additionally well-known for components performing the functions of a D-VSI to associate packets with a virtual network identifier and to use that identifier to select the correct switching path and the tunnel used for forwarding. For example, this is disclosed by Kamath, Koganti, and Qu:

Components of an overlay network 300 typically identify where to route packets based on a virtual network identifier, referred to as a VNI or VNID. This is typically a 24-bit code or number, which excludes 0x0 and 0xFFFFFFFF. The overlay network 300 has the capability of tunneling layer-2 packets over the layer-3 network by encapsulating the layer-2 packets into an overlay header. This may be performed using Virtual eXtensible Local Area Network (VXLAN) or some other overlay capable protocol, such as locator/ID separation protocol (LISP), overlay transport virtualization (OTV), etc. The packet may also be encapsulated in a user datagram protocol (UDP) and internet protocol (IP) UDP/IP header. The overlay network 300 may include one or more point-to-point tunnels, and/or point-to-multipoint tunnels. In addition, any of these tunnels may be created, removed, altered and modified based on any number of factors, such as new devices being added to the overlay network 300, removal of devices from the overlay network 300, startup of any end devices, i.e., devices managing tunnel end points, such as virtual overlay network gateways, Hypervisors, switches capable of overlay functionality, etc. In order for a device to manage a tunnel, there needs to be a mapping between an original packet's source address, destination address, and a tunnel identifier. In this way, a physical server is capable of forwarding the encapsulated original packet to the proper destination device.

Kamath at [0047].

In addition, the information exchange packet may comprise, at least, a media access control (MAC) address for each virtual machine (VM) connected to each end point station that is in the virtual network having the specified VNID. Furthermore, the system 400 may further include logic adapted for mapping tunnel identifiers for each end point station and the MAC address for each VM connected to each end point station.

Kamath at [0051].

If the frame has a tunnel encapsulation, the destination MAC address of the decapsulated frame is a virtual MAC address associated with the virtual tunnel gateway. The switch checks whether the destination MAC address is a local address (operation 458). For example, if the switch is a member tunnel gateway, the virtual MAC address is a local address. If the destination MAC address is not local, the frame is destined for a locally coupled external device, and the switch forwards the decapsulated frame to the locally coupled external device (operation 470). If the MAC address is local, the switch promotes the frame to the upper layer and extracts the internal encapsulated packet (operation 460), and obtains the IP address of the extracted packet (operation 462).

Koganti at [0084].

The VXLAN packet 800 may further comprise a VXLAN ID field 807, which may indicate the VXLAN network that the packet 800 belongs to. The VXLAN packet 800 may comprise a VM MAC DA field 808 and a VM MAC SA field 809, which comprise the encapsulated VLAN destination and source MAC addresses.

Qu at [0033].

The active/standby VSMs and the VEMs are linked together as part of the same virtual switch instance leveraging the concept of “domain ID”. A domain ID is a parameter of the Cisco Nexus 1000V Series Switch that is used to identify a VSM and VEM as relating to one another. Each command sent by the VSM to any associated VEMs is tagged with this domain ID and if the same domain ID is shared by a VSM and a VEM, the VEM will accept and respond to requests and commands from the VSM. If the VEM receives a command or configuration request that is not tagged with the correct domain ID, that request is ignored. Similarly, if the VSM receives a packet from a VEM that is tagged with the wrong domain ID, the packet will be ignored.

Data Center Interconnect Design Guide for Virtualized Workload Mobility with Cisco, EMC and VMware (Last Update August 29, 2011) at 2-24.

It was additionally well-known for such components to obtain source address information to populate a forwarding table within a virtual switching instance, consistent with the claim language. For example, this is disclosed by Ye, Farinacci, and Grover:

When a PE device receives a unicast packet from an AC, the packet is an Ethernet frame that may include a MAC source address, a MAC destination address, a VLAN tag with a VLAN ID, and a payload. Based on the AC's port number, the PE device maps the packet to a VSI of a particular VPLS instance and its MAC address table. When a PE device receives a unicast packet from an inbound VC LSP of a PW, the packet is a MPLS packet that may include a tunnel label (an outer label), a PW label (an inner label), and an Ethernet frame. The PE device removes the tunnel label to expose the PW label. Based on the PW label, the PE device maps the packet to a VSI of a particular VPLS instance and its MAC address table.

When the MAC source address of the packet is unknown, the PE device creates a new entry in the MAC address table that maps the MAC address to the AC (and VLAN) or an outbound VC LSP of the PW. The PE device also starts an aging timer for the entry. When the aging timer expires, the PE device withdraws the entry from the MAC address table.

When the MAC source address of the packet is known, the PE device resets the aging timer for the entry.

When the MAC destination address of the packet is known, the PE device forwards the packet to an AC or a PW specified in the entry. When the packet is forwarded through a PW, the PE device adds a PW label and a tunnel label to encapsulate and forward the packet. Based on the PW label, the opposing PE device maps the packet to a local VSI of the same VPLS instance and its MAC address table.

Ye at [0036]-[0039].

The MAC table 30 is populated using the address information received in LSPs and GM-LSPs at the overlay interface 32 and packets received at internal interfaces E1, E2, as previously described with respect to advertising addresses.... The table 30 provides information used by the edge device 12 in forwarding packets received at its interfaces.

Farinacci at 9:54-66.

MAC Table - This is a forwarding table of 48-bit MAC addresses. The table can contain unicast or multicast MAC addresses. The table is populated by two sources. One being traditional data-plane learning on internal interfaces and the other by the URP/MRP at the control-plane on the overlay interface. A MAC table is scoped

by VLAN therefore allowing the same MAC address to be used in different VLANs, and potentially in different VPNs.

Grover at p. 7.

It was additionally well-known for such components to forward packets based on destination address information—including by decapsulating received packets, switching according to Layer-2 information, and encapsulating and sending the packet through a tunnel—based on information contained in such forwarding tables within virtual switching instances, as recited in the claims. For example, this is disclosed by Ye, Grover, Ghosh, Yong, and Cisco Data Center:

When a PE device receives a unicast packet from an AC, the packet is an Ethernet frame that may include a MAC source address, a MAC destination address, a VLAN tag with a VLAN ID, and a payload. Based on the AC's port number, the PE device maps the packet to a VSI of a particular VPLS instance and its MAC address table. When a PE device receives a unicast packet from an inbound VC LSP of a PW, the packet is a MPLS packet that may include a tunnel label (an outer label), a PW label (an inner label), and an Ethernet frame. The PE device removes the tunnel label to expose the PW label. Based on the PW label, the PE device maps the packet to a VSI of a particular VPLS instance and its MAC address table.

When the MAC source address of the packet is unknown, the PE device creates a new entry in the MAC address table that maps the MAC address to the AC (and VLAN) or an outbound VC LSP of the PW. The PE device also starts an aging timer for the entry. When the aging timer expires, the PE device withdraws the entry from the MAC address table.

When the MAC source address of the packet is known, the PE device resets the aging timer for the entry.

When the MAC destination address of the packet is known, the PE device forwards the packet to an AC or a PW specified in the entry. When the packet is forwarded through a PW, the PE device adds a PW label and a tunnel label to encapsulate and forward the packet. Based on the PW label, the opposing PE device maps the packet to a local VSI of the same VPLS instance and its MAC address table.

Ye at [0036]-[0039].

When a unicast packet is received on the overlay interface, the outer OTV IP header is removed, and the VLAN-ID and the MAC DA from the inner header is used to

do the MAC table lookup. Here onwards, this is a regular bridging operation, whether the MAC address entry is present or not.

Grover at p. 19.

In a typical configuration, VPLS-enabled routers that are associated with the customer networks define LSPs within the intermediate network to carry encapsulated L2 communications as if these customer networks were directly attached to the same LAN. To properly communicate via these LSPs, each of these VPLS-enabled routers store L2 information, such as Media Access Control (MAC) addresses, as well as VPLS information, such as local and remote VPLS site information. In this manner, these VPLS-enabled routers provide transparent L2 connectivity across the intermediate network and simulate a direct LAN.

Ghosh at 1:51-62.

A network virtualization instance may be functioned as L2 or L3 as described in the IETF document draft-narten-nvo3-arch-00, where tenant traffic may be tunneled to remote NVEs 120 based on the Media Access Control (MAC) address of the TSs 110 or the IP addresses of the TSs 110, respectively.

Yong at [0033].

When an edge device receives a Layer 2 frame on an internal interface, OTV performs the MAC table lookup based on the destination address of the Layer 2 frame. If the frame is destined to a MAC address that is reachable through another internal interface, the frame is forwarded out on that internal interface. OTV performs no other actions and the processing of the frame is complete. If the frame is destined to a MAC address that was learned over an overlay interface, OTV performs the following tasks:

1. Strips off the preamble and frame check sequence (FCS) from the Layer 2 frame.
2. Adds an OTV header to the Layer 2 frame and copies the 802.1Q information into the OTV header.
3. Adds the IP address to the packet, based on the initial MAC address table lookup. This IP address is used as a destination address for the IP packet that is sent into the core switch.

Cisco Nexus 7000 Series NX-OS OTV Configuration Guide (First Published 2010-08-06; Modified 09-08-2014) at 6-7.

- *Tunnel packet processing including receiving and decapsulating tunnel packets for L2 switching, and encapsulating packets with tunnel headers for transmission through the overlay network*

The tunnel packet processing functions recited in the '961 Patent-including receiving and decapsulating tunnel packets for L2 switching, and encapsulating packets with tunnel headers for transmission-were well known prior to the alleged invention. For example, this is disclosed by Kamath and Ghosh:

The tunnel is actually implemented by encapsulating packets transmitted by a source end station into an overlay header that transports the packet from the source switch to a target switch in user datagram protocol (UDP) transport via an internet protocol (IP)-based network. The overlay header includes an identifier (ID) that uniquely identifies the virtual network. The target switch strips off the overlay header encapsulation, UDP transport header, and IP header, and delivers the original packet to the destination end station. In addition to this tunneling mechanism, the edge switches participate in an address discovery protocol, which may be learning/flooding based, or lookup-based.

Kamath at [0002].

The MPLS packets encapsulate the L2 communications, thereby effectively shielding the L3 network from the transported L2 information. One example of an L2 service is the Virtual Private LAN Service (VPLS) that simulates a local area network. In general, VPLS allows two or more remote customer networks to be extended through the intermediate network as if the intermediate network does not exist. In particular, L2 communications, such as Ethernet packets, are transported between customer networks via the intermediate network.

Ghosh at 1:42-51.

It was additionally well-known for components performing these functions to receive tunnel packets and decapsulate them for L2 switching, as recited in the claims. For example, this is disclosed by Yong, Farinacci, and Grover:

Returning to step 820, an NVE may also receive a data packet destined to one of its associated TSs via an overlay tunnel either from an OVG or a peer NVE. In this case, method 800 may operate as the egress NVE. At step 851, method 800 may remove the tunnel header from the received data packet. At step 852, method 800

may decapsulate the received data packet (i.e. removing the encapsulation header). At step 853, method 800 may deliver the data packet to the destination TS.

Yong at [0054].

When a packet 40 is received on the overlay interface 32, it is decapsulated to reveal the inner MAC header for forwarding. The inner MAC SA (source address) and DA (destination address) are used for the MAC table lookup described below.

Farinacci at 9:40-43.

When a packet is received on the overlay interface, it will need to be IP decapsulated to reveal the inner MAC header for forwarding. The inner MAC header SA and DA addresses and VLAN-ID will be used for forwarding actions. For any type of packet received on the overlay interface, it will be accepted only if the Edge Device is the Authoritative Edge Device as determined by an inspection of the received packet header.

Grover at p. 14.

An overlay interface is a logical interface which is associated with an IP address in the provider/core address space. Traffic out of these interfaces is encapsulated with an IP header, and traffic received on these interfaces must be de-capsulated to produce a L2 frame. The encapsulated packets exit the Edge Device on one or more underlying physical or logical L3 interfaces.

Grover at p. 19.

It was additionally well-known for such components to perform L2 switching on decapsulated packets based on destination address information, including forwarding locally or re-encapsulating for transmission through the overlay network. For example, this is disclosed by Ye,

Koganti and Cisco Data Center:

In block 1202, PE device 108-1 receives a unicast packet from PW 114 and removes the tunnel label from the packet to expose VSI label 118 (FIG. 1). Block 1202 may be followed by block 1204.

In block 1204, PE device 108-1 maps VSI label 118 to VSI 112-1 and its MAC address table. Block 1204 may be followed by block 1206.

In block 1206, PE device 108-1 determines if the packet has a known MAC destination address. For example, PE device 108-1 looks for a matching entry with the same MAC address in the MAC address table of VSI 112-1. If the packet has a

known MAC destination address, block 1206 may be followed by block 1208. Otherwise block 1208 may be followed by block 1210.

In block 1208, PE device 108-1 removes the VSI label and forwards the packet to the path or port (and VLAN) identified in the matching entry of the MAC address table of VSI 112-1. For example, device 108-1 adds a VSI label and a tunnel label to encapsulate and forward the packet. After block 1208, method 1200 may end.

Ye at [0099]-[0102]; *see also* Ye at FIG. 12.

Upon receiving the frame, egress switch 103 identifies the destination MAC address to be associated with virtual gateway switch 150. Switch 103 considers virtual gateway switch 150 to be another member switch and forwards the frame to switch 101. Upon receiving the frame, switch 101 recognizes the virtual IP and MAC addresses to be local addresses, extracts the inner packet, and forwards the inner packet to router 134 based on the forwarding information of the inner packet. Similarly, if virtual machine 124 in host machine 114 sends a frame toward switch 136, the hypervisor in host machine 114 tunnels the frame by encapsulating the frame in a layer-3 packet with the virtual IP address as the destination IP address. Switch 103 receives the frame, recognizes the virtual IP and MAC addresses to be local addresses, extracts the inner packet, and forwards the inner packet to switch 136 based on the forwarding information of the inner packet.

Koganti at [0058]; *see also* Koganti at FIGS. 4A, 4B.

When an edge device receives a Layer 2 frame on an internal interface, OTV performs the MAC table lookup based on the destination address of the Layer 2 frame. If the frame is destined to a MAC address that is reachable through another internal interface, the frame is forwarded out on that internal interface. OTV performs no other actions and the processing of the frame is complete. If the frame is destined to a MAC address that was learned over an overlay interface, OTV performs the following tasks:

1. Strips off the preamble and frame check sequence (FCS) from the Layer 2 frame.
2. Adds an OTV header to the Layer 2 frame and copies the 802.1Q information into the OTV header.
3. Adds the IP address to the packet, based on the initial MAC address table lookup. This IP address is used as a destination address for the IP packet that is sent into the core switch.

Cisco Nexus 7000 Series NX-OS OTV Configuration Guide (First Published 2010-08-06; Modified 09-08-2014) at 6-7.

It was additionally well-known for such components to encapsulate packets with tunnel headers for transmission, including a VN header with a virtual network identifier and an L3 tunnel header with source and destination address information, as recited in the claims. For example, this is disclosed by Yong, Qu, Farinacci, and Cisco Data Center:

FIG. 8 illustrates a flowchart of an example embodiment of a method 800 for routing data traffic at an NVE, which may be implemented by an NVE 120, as shown in FIG. 2. Method 800 may begin with receiving a data packet at step 810. At step 820, method 800 may check if the data packet is received from an attached TS or from an overlay tunnel. If the data packet is from an attached TS, method 800 may operate as an ingress NVE. At step 831, method 800 may retrieve a destination TS address from the data packet. At step 832, method 800 may retrieve an address mapping table entry with a mapping of the destination TS address to its associated NVE address (i.e. egress NVE address). The address mapping table may be built previously from control plane as described in method 700 of FIG. 7, a network node, and/or some other address resolution protocols. At step 833, method 800 may retrieve a forwarding table entry with the egress NVE's route and encapsulation type. The forwarding table may be built previously from control plane as described in method 700 of FIG. 7, a network node, and/or some other central authority. At step 834, method 800 may encapsulate the data packet according to the encapsulation type supported at the ingress NVE (e.g. adding an encapsulation header as shown in FIG. 4).

At step 835, method 800 may check if the egress NVE supports the same encapsulation type as the ingress NVE. If the egress NVE supports the same encapsulation type, method 800 may proceed to step 836 to add a tunnel header to the encapsulated data packet. The tunnel header may comprise the egress NVE's IP address and the ingress NVE's IP address, as described in packet header 300 with respect to FIG. 3. At step 837, method 800 may send the encapsulated data packet directly to the egress NVE.

Yong at [0051]-[0052]; *see also* Yong at FIGS. 8, 9.

FIG. 8 is a diagram of an embodiment of a VXLAN communications packet 800 as used in some of the embodiments discussed herein. The VXLAN packet 800 was discussed in draft-mahalingam-dutt-dcops-vxlan-00, which is incorporated by reference as if completely set forth herein. The VXLAN packet 800 may encapsulate a VLAN packet. The VXLAN packet 800 may comprise a Virtual Ethernet Module (VEM) MAC DA field 801 and a VEM MAC SA field 802, which comprise the packet's outer next hop address and source hop mac address respectively. Each of these values may be updated each time the packet 800 makes a hop to a new node. The VXLAN packet 800 may further comprise a transport VLAN 802.1 Q tag 803, which may be used to encode information about the

VXLAN the packet 800 belongs to. The VXLAN packet 800 may further comprise a VEM IP DA field 804 and a VEM IP SA field 805, which may comprise the IP addresses of the destination VTEP and the source VTEP. The VXLAN packet 800 may further comprise a UDP field 806, which may comprise VTEP port information. The VXLAN packet 800 may further comprise a VXLAN ID field 807, which may indicate the VXLAN network that the packet 800 belongs to. The VXLAN packet 800 may comprise a VM MAC DA field 808 and a VM MAC SA field 809, which comprise the encapsulated VLAN destination and source MAC addresses. The VXLAN packet 800 may comprise a L3 field 810, which may comprise a VLAN identification tag and a payload of the information to be communicated. The VXLAN packet 800 may comprise a CRC field 811, which may contain information used for transmission error checking purposes.

Qu at [0033].

A processor 28 encapsulates L2 (MAC level) packets 38 from the internal interfaces 34, 36 in IP packets (with unicast or multicast headers) 40 before forwarding the packets to the core network, and decapsulates IP packets received from the core network. For example, in one embodiment a unicast frame is encapsulated in an IP unicast packet; a broadcast frame is encapsulated in an ASM/Bidir IP multicast packet; a link-local multicast frame is encapsulated in an ASM/Bidir IP multicast packet; and a customer multicast frame is encapsulated in a SSM IP multicast packet. The ASM/Bidir multicast packet is also used for IS-IS peering.

Farinacci at 8:9-19.

Difference#5: Virtual Extensible VLAN (VXLAN)

Generically speaking, VLANs offer a nice model for the separation of security domains within a physical network. Through the isolation of multi-destination traffic (broadcast and floods), two physical hosts connected to different VLANs depend on other network functions, such as routing, to communicate.

While they are adequate for most data center environments, VLANs present the following limitations against the high scalability of server virtualization environments:

- VLANs permit only the creation of 4094 different domains
- The creation of a VLAN usually depends on configurations on every Layer 2 device in the network, and sometimes, on a Layer 2 extension technology such as Overlay Transport Virtualization (OTV) or Virtual Private LAN Service (VPLS)
- Virtualized servers with hundreds of virtual machines can quickly consume hardware resources (such as MAC address table entries) from the physical network.

To overcome these challenges, Cisco, VMware and other manufacturers have developed a technology called *Virtual eXtensible LAN* (or simply VXLAN), which was submitted to IETF as a draft in 2010 (<http://tools.ietf.org/html/draft-mahalingam-dutt-dcops-vxlan-00>). In summary, this technology allows that two different VMs exchange Ethernet frames through the dynamic encapsulation of these frames into UDP datagrams marked with 24-bit VXLAN segment identifiers.

Nexus 1000V currently deploys VXLAN according to the Figure 3.




Figure 3: Virtual eXtensible LAN Implementation

In Figure 3, two VEMs share two VXLAN segments (50000 and 60000). For VMs A and B, VEM West performs transparent bridging exactly like it would for a VLAN. Differences arise when hosts from the same VXLAN are connected to different VEMs. In this case, frames are encapsulated into UDP datagrams

Chalk Talk: The Top 5 Differences between Physical and Virtual Networking (07-18-2013) at 6.

It was additionally well-known for such tunnel packet processing to be performed at network edge devices that connected customer sites to an intermediate network. For example, this is disclosed by Figueira and Sajassi:

The architecture in FIG. 3 can be thought of as a type of hierarchical L2VPN service where a L2VPN over PBB (L2VPN over PBB) is used at the metro network, while VPLS is used at the MPLS WAN (core) network. This architecture decouples the

full mesh of Pseudowires from the full mesh of customer Service Instances (which are transported by PBT trunks), allowing for increased scalability of the overall end-to-end L2VPN service.

Figueira at [0022].

FIG. 3 is a simplified example of a SP network in accordance with one embodiment of the present invention in which a core network 20 is shown connected to a pair of Ethernet access networks 21 & 22 via network provider edge (n-PE) devices 32 & 33, respectively. User-facing provider edge (u-PE) devices 31 & 34 connect respective customer edge (CE) devices 41 & 42 to Ethernet access networks 21 & 22. In this illustration, n-PE devices 32 and 33 are implemented in accordance with the PE device model shown in FIG. 1. That is, for traffic passing from left to right in FIG. 3, n-PE device 32 maps the 12-bit provider VLAN ID associated with customer edge device 41 into a 20-bit E-VLAN tag identifier that is then sent over core network 20. At the other end, n-PE device 33 translates the 20-bit E-VLAN tag back to a standard 12-bit provider VLAN ID, which eventually passes to CE device 42 at the destination end.

Sajassi at 4:52-67.

- *Edge virtual ports and domain virtual ports within a D-VSI, each associated with corresponding tunnel end points, with modified flooding rules based on input port type to control traffic forwarding and prevent loops*

The concept of a D-VSI having both edge virtual ports and domain virtual ports, with modified flooding rules based on input port type, was well known prior to the alleged invention of the '961 Patent. Components performing the functions of the claimed "edge virtual port" were variously referred to as edge ports, customer-facing ports, access circuits (ACs), internal interfaces, and customer interface connections. Components performing the functions of the claimed "domain virtual port" were variously referred to as inter-switch ports, network-facing ports, pseudowires (PWs), overlay interfaces, and service instance connections. For example, this is disclosed by Koganti and Farinacci:

Member switches in fabric switch 100 use edge ports to communicate to external devices and inter-switch ports to communicate to other member switches. For example, switch 102 is coupled to external devices, such as router 134 and switch 136, via edge ports and to switches 101, 103, 104, and 105 via inter-switch ports.

Koganti at [0054].

Referring again to the example shown in FIG. 2, edge device 12 includes one or more internal interfaces 34, 36 in communication with network devices at the same network site as the edge device, and an overlay interface 32 in communication with the overlay network 26. The internal interfaces 34, 36 are L2 interfaces connected to site-based switches or site-based routers and learn the source MAC addresses of traffic they receive. The internal interfaces 34, 36 are L2 interfaces regardless if they connect to a switch or a router. Traffic received on an internal interface may trigger oUMRP advertisements or pMRP group joins, as previously described. Traffic received on the internal interfaces 34, 36 is forwarded according to a MAC table 30 onto another internal interface (regular bridging (or routing in the case of L3 edge device)) or onto the overlay interface (OTV forwarding).

In one embodiment, the forwarding table 30 is a MAC table of 48-bit MAC addresses. The table 30 may contain unicast MAC addresses or multicast MAC addresses. The table 30 is populated by conventional data-plane learning on internal interfaces 34, 36 and by the IS-IS protocol (or other IGP protocol) at the control plane on the overlay interface 32. Arrow 39 illustrates data-plane learning on internal interface 34 and arrow 43 illustrates control-plane learning on the overlay interface 32. The table 30 allows the same MAC address to be used in different VLANs and potentially in different VPNs. The table 30 is described further below with respect to forwarding.

The overlay interface 32 is a logical multi-access multicast-capable interface and is realized by one or more physical core facing interfaces. The core facing interface 32 is assigned an IP address (e.g., IPB in FIG. 2) out of the provider/core address space. Even though the overlay interface 32 has an IP address, it does not participate in the provider IGP or UMRP, as noted above. The overlay interface 32 operates as a host connected to the core network 10. The IP address assigned to the overlay interface 32 is used as a next hop address by the oUMRP, therefore, table 30 for the overlay interface includes a remote IP address as the next hop information for remote MAC addresses. As described below, traffic transmitted out of the overlay interface 32 is encapsulated with an IP header and traffic received on this interface is decapsulated to produce an L2 frame.

In the example shown in FIG. 2, there is only one overlay interface 32. It is to be understood, however, that there may be any number of overlay interfaces.

Farinacci at 7:31-8:8.

It was additionally well-known for each domain virtual port to be associated with a corresponding tunnel end point, such that traffic destined for a remote site would be forwarded through the domain virtual port to that tunnel end point. For example, this is disclosed by Sajassi:

FIG. 3 is a simplified example of a SP network in accordance with one embodiment of the present invention in which a core network 20 is shown connected to a pair of Ethernet access networks 21 & 22 via network provider edge (n-PE) devices 32 & 33, respectively. User-facing provider edge (u-PE) devices 31 & 34 connect respective customer edge (CE) devices 41 & 42 to Ethernet access networks 21 & 22. In this illustration, n-PE devices 32 and 33 are implemented in accordance with the PE device model shown in FIG. 1. That is, for traffic passing from left to right in FIG. 3, n-PE device 32 maps the 12-bit provider VLAN ID associated with customer edge device 41 into a 20-bit E-VLAN tag identifier that is then sent over core network 20. At the other end, n-PE device 33 translates the 20-bit E-VLAN tag back to a standard 12-bit provider VLAN ID, which eventually passes to CE device 42 at the destination end.

Sajassi at 4:52-67; see also Sajassi at FIGS. 1, 3, 4.

It was additionally well-known for D-VSIs to apply modified flooding rules based on whether the input port was an edge virtual port or a domain virtual port, as recited in claims 2 and 8 of the '961 Patent. When a packet with an unknown destination was received on an edge virtual port, it would be flooded to domain virtual ports; when received on a domain virtual port, it would be flooded only to edge virtual ports. For example, this is disclosed by Ye, Figueira, and Cisco Data Center:

PE devices may be fully meshed logically and they implements [sic] horizontal split forwarding to avoid loops. When a PE device receives from a PW associated with a VSI (1) a unicast packet with an unknown MAC destination address, (2) a multicast packet, or (3) a broadcast packet, the PE device floods the packet to all ACs associated with the VSI but not PWs associated with the VSI. When the PE device receives the same from an AC associated with the VSI, the PE device floods the packet to all the ACs and all the PWs associated with the VSI.

When a PE device receives a broadcast or multicast packet at an AC, which is indicated by the MAC destination address in the packet, the PE device broadcasts the packet to other PE devices having VSIs in the same VPLS instance. The PE device makes a copy of the packet and adds the appropriate PW labels and tunnel labels to send the packet to each of the other PE devices.

Ye at [0040], [0041].

FIG. 5 is a diagram illustrating an embodiment of frame forwarding rules collectively called split horizon. A form of split horizon is also used with VPLS, as described in RFC 4762. In some embodiments, split horizon is used to prevent or

avoid loops so that frames do not unnecessarily flood the L2VPN. A split horizon rule constrains the normal forwarding feature of a VSI (which is a learning bridge), which would normally allow a frame received on a VSI connection to be forwarded to any one of the other VSI connections. New split horizon rules are defined here for use with L2VPN over PBB networks, as is described next. FIG. 5 illustrates an end-to-end L2VPN for a customer using an MPLS Core network and two PBB metro networks. The end-to-end L2VPN is constructed by interconnecting L2VPN1 530 and L2VPN2 531, which are both instances of L2VPN over PBB, and VPLS 532 using the common VSIs 510 and 520 at the border of the PBB metro networks and the MPLS core network. To avoid forwarding loops, the split horizon rule 500 is first defined. Split horizon rule 500 is conceptual in nature and has two sides, each corresponding to a different type of network. On the lower right hand side of split horizon rule 500 are Service Instance connections 502 and 503. On the other side (i.e., upper left hand side) of split horizon rule 500 are the customer interface connections 504 and 505. A split horizon rule such as 500 applies to VSIs connected to Service Instances over PBT trunks and customer bound interfaces. Split horizon rule 500 states that: (1) whenever a frame is received from a Service Instance (which is transported over a PBT trunk) it can only be forwarded to customer bound interfaces; and (2) whenever a frame is received from a customer bound interface, it can be forwarded to any other VSI connection, i.e., the VSI uses its normal forwarding behavior. For example, a frame received from Service Instance 502 can only be forwarded to interface connections 504 and 505 (i.e., the frame can be forwarded to either one of these interface connection or both) and will not be forwarded to another Service Instance, such as Service Instance connection 503. In addition, a frame received from customer interface 504 can be forwarded as usual by the VSI forwarding process, i.e., it can be forwarded to interface connection 505 and Service Instances 502 and 503. Similarly, the split horizon rule 501 will avoid forwarding loops. Split horizon rule 501 is also conceptual in nature and has three sides (i.e., this is a three-dimensional split horizon rule), each corresponding to a different type of network. On the upper hand side of split horizon rule 501 are Service Instance connections 506 and 507. On the lower hand of split horizon rule 501 are Pseudowire connections 508 and 509. On the left hand side is customer interface connection 550. Split horizon rule 501 applies to the VSIs that provide interworking between the PBB and MPLS networks, such as VSI 510. Split horizon rule 501 states that: (1) whenever a frame is received from a Service Instance (which is transported over a PBT trunk), it can only be forwarded to Pseudowires over the MPLS network and customer interface connections; (2) whenever a frame is received from a Pseudowire (which is transported over an MPLS network), it can only be forwarded to Service Instances (which are transported over PBT trunks) and customer interface connections; and (3) whenever a frame is received from a customer bound interface, it can be forwarded to any other VSI connection, i.e., the VSI uses its normal forwarding behavior. For example, a frame received from Service Instance 506 can only be forwarded to Pseudowires 508 and 509 and customer interface 550 and will not be forwarded to another Service Instance, such as Service Instance connection 507. Conversely, a frame received from Pseudowire 508 can only be forwarded to Service Instances

506 and 507 and customer interface connection 550 and will not be forwarded to another Pseudowire, such as Pseudowire connection 509. Note that customer interface connections like 550 are optional. If no customer interface connection such as 550 is present, the split horizon rule 501 can be simplified. In this case, it would state that: (1) whenever a frame is received from a Service Instance (which is transported over a PBT trunk), it can only be forwarded to Pseudowires over the MPLS network; and (2) whenever a frame is received from a Pseudowire (which is transported over an MPLS network), it can only be forwarded to Service Instances (which are transported over PBT trunks). Note that split horizon rules 500 and 501 are new types of split horizon rules defined here to avoid loops due to the L2VPNs over PBB. The split horizon rule 511 exemplified here is the common well-known split horizon rule defined in RFC 4762 to avoid loops in the VPLS network. Split horizon rule 511 is explained here to highlight how it differs from split horizon rules 500 and 501. Split horizon rule 511 states that: (1) whenever a frame is received from a Pseudowire, it can only be forwarded to customer bound interfaces; and (2) whenever a frame is received from a customer interface, it can be forwarded to any one of the other VSI connections. For example, a frame received from a Pseudowire will never be forwarded to another Pseudowire. In one embodiment, one or multiple PBB metro networks can be connected to a core MPLS network to provide an end-to-end L2VPN for a customer. The fact that FIG. 5 shows two PBB metro networks is exemplary and should not be construed as limiting.

Figuiera at [0030].

Normally, unknown unicast Layer 2 frames are not flooded between OTV sites, and MAC addresses are not learned across the overlay interface. Any unknown unicast messages that reach the OTV edge device are blocked from crossing the logical overlay, allowing OTV to prevent Layer 2 faults from spreading to remote sites.

Cisco Nexus 7000 Series NX-OS OTV Configuration Guide (First Published 2010-08-06; Modified 09-08-2014) at 37.

When NLB is used in the unicast mode and deployed concurrently with OTV, it relies on unknown unicast flooding. OTV as a technology does not allow unknown unicast flooding across sites. Unknown unicast Layer 2 frames are not flooded between OTV sites, and MAC addresses are not learned across the overlay interface. Any unknown unicast messages that reach the OTV edge device are blocked from crossing the logical overlay. To use NLB in unicast mode across OTV sites, you must enable selective unicast flooding in OTV. This feature enables flooding for the specified destination MAC address to all other edge devices in the OTV overlay network.

Overlay Transport Virtualization Best Practice Guide – Configuration Guide (November 2017) at 11.

It was additionally well-known to use such modified flooding rules, such as the split-horizon rule, or Authoritative Edge Devices to prevent forwarding loops in L2 networks, consistent with the claims of the '961 Patent. For example, this is disclosed by Ghosh, Grover, and Cisco Data Center:

In L2 networks, an L2 device responsible for forwarding network traffic, such as an L2 switch or an VPLS-enabled router, operates to “flood” (i.e., forward a copy) network traffic to a set of output ports of the device in a situation where the network traffic is destined for a network address that has not previously been seen, and therefore not learned, by the device. In this manner the L2 device can ensure that at least the destination receives the network traffic regardless of the destination's location within the L2 network. The L2 device continues to flood to the set of output ports all subsequent network traffic bound for the destination address until the L2 device receives some network traffic originated from the address, at which time the L2 device is able to “learn” the specific port by which the L2 address is reachable. As a result, the L2 device forwards future network traffic bound for the address to the particular output port and no longer needs to flood the network traffic. In this way, L2 network devices rely on network traffic flooding to ensure that L2 traffic reach their destinations. However, the use of network flooding requires that the L2 devices operate to reduce or eliminate any “data loops” in the L2 network, i.e., situations where network traffic is flooded in a manner that results in the flooding device subsequently receiving the same traffic it flooded. The presence of such a data loop may have severe adverse affects on a L2 network. For example, a data loop may result in consumption of significant bandwidth and resources as the flooded traffic within the data loop is continuously flooded by the devices. Moreover, a data loop may also result in the final destination receiving multiple copies of the same packet.

Thus, it is often necessary to eliminate data loops from the L2 network topology. As a result, many L2 switches or VPLS-enabled routers often attempt to operate in accordance with a rule known as the “split-horizon rule” to avoid data loops. The split-horizon rule prohibits a router flooding a copy of an L2 packet back out an interface from which the original L2 packet was received.

Ghosh at 1:63-2:32.


A broadcast packet, sourced in a site, gets to all Edge Devices because each Edge Device is on the site spanning tree. However, duplicates must not be allowed to appear on the overlay network when there are multiple Edge Devices, so the Authoritative Edge Device for the VLAN is the only Edge Device that forwards the packet on the overlay network. All edge devices at a remote site will receive the broadcast packet over the core multicast group. To prevent duplicates going into the site, only the Authoritative Edge Device in that site will forward the packet into

the site. And once sent into the site, the packet gets to all switches on the site spanning tree. Because only the AED can forward broadcast packets in or out of the site, broadcast loops are avoided.

Grover at p. 21.

Figure 1-24 shows how the definition of the AED role in each site prevents the creation of end-to-end STP loops.

Figure 1-24 *Prevention of End-to-End STP Loops*



Assume, for example, a Layer 2 broadcast frame is generated in the left data center. The frame is received by both OTV edge devices however, only the AED is allowed to encapsulate the frame and send it over the OTV overlay. All OTV edge devices in remote sites will also receive the frame via the overlay, since broadcast traffic is delivered via the multicast group used for the OTV control protocol, but only the AED is allowed to decapsulate the frame and send it into the site.

Cisco Overlay Transport Virtualization Technology Introduction and Deployment

Considerations (January 17, 2012) at 1-23.

Defendant is not aware of any secondary considerations supporting non-obviousness of the Asserted Claims of the '961 Patent. To the extent Plaintiff seeks to introduce any alleged secondary considerations in the future, Defendant reserves the right to respond at that time.

For at least the reasons described herein as well as in the combinations of the prior art identified above and charted in the claim charts found at Exhibits E-1 to E-11, it would have been obvious to one of ordinary skill in the art to combine any of a number of prior art references, including any combination of those identified above.

Defendant further identifies the commercial embodiments of the inventions described in the prior art charted in Exhibits E-1 to E-11, to the extent those systems or services were on sale or in public use before the priority dates of the '961 Patent.

A person having ordinary skill in the art would have been motivated to combine any of the references charted in attached Exhibits E-1 to E-11 for at least the reasons described above. A person having ordinary skill in the art at the time of filing of the '961 Patent would also have understood the references listed above, alone or in combination, to contain explicit and/or implicit teaching, suggestion, and/or rationales to combine them, including as further described above.

To the extent one of the above references does not disclose a specific limitation, it would have been obvious to combine that reference with one or more of the other references. A person of ordinary skill in the art at the time of the alleged invention would have combined the references because each generally relates to similar technology (e.g., overlay network virtualization, Layer 2 virtual private networks, and associated tunneling systems and methods), and the references could be combined, or substituted, according to known methods that would yield predictable results. Moreover, to the extent the '961 Patent identifies any needed improvement in the art, the references provide known techniques that would yield the required improvement, and given the finite number of techniques identified, it would have been obvious to try those approaches with a reasonable expectation of success.

The combinations of references disclosed in these contentions are exemplary and are not intended to be exhaustive. Additional obviousness combinations of the references identified here are possible, and Defendant may rely on such combination(s) in this litigation. In particular, Defendant is currently unaware of Plaintiff's allegations with respect to the level of skill in the art and the qualifications of the typical person of ordinary skill in the art. Defendant is also unaware of the extent, if any, to which Plaintiff may contend that limitations of the claims at issue are not disclosed in the prior art identified by Defendant, and the extent to which Plaintiff will contend that

elements not disclosed in the '961 Patent specifications would have been known to persons of skill in the art. And Defendant does not yet know how the Court will construe terms in the Asserted Claims of the '961 Patent. Defendant is also continuing their investigation of the large universe of prior art to identify potential prior art systems, publications related to those systems, and third parties that may have information about those systems. Plaintiff may also be in possession of prior art that Defendant may receive as discovery progresses in this case. Defendant reserves the right to amend and supplement these contentions to identify other prior art and combinations rendering the Asserted Claims of the '961 Patent anticipated and/or obvious.

2. Invalidity Under 35 U.S.C. § 112

Defendant identifies the following claim limitations that render the Asserted Claims of the '961 Patent invalid based on lack of enablement and written description under post-AIA 35 U.S.C. § 112(a). All claims that depend from a claim containing a limitation identified below are invalid by way of their dependence. Defendant reserves the right to modify these contentions or further develop them as discovery progresses, including through expert disclosure and discovery.

a) Lack of Enablement and Written Description

Defendant contends that the specification (including any specification to which priority is sought) fails to describe the full scope of the claimed inventions in sufficient detail such that one skilled in the art could reasonably conclude that the inventor had possession of the claimed subject matter, and the claims are not enabled without undue experimentation. In particular, the specification does not provide adequate description or enabling guidance for the following limitations as recited and used in the '961 claims:

- “connecting tunnel end points within the network virtualization domain to other, tunnel end points within another network virtualization network domain with a unicast tunnel using the new tunnel packet” (claim 7)

Defendant asserts the specification (including any specification to which priority is sought) fails to describe the full scope of the claimed inventions in sufficient detail such that one skilled in the art could reasonably conclude that the inventor had possession of the claimed invention, i.e., an operating method of a network virtualization apparatus including connecting tunnel end points within the network virtualization domain to other, tunnel end points within another network virtualization network domain with a unicast tunnel using the new tunnel packet. The specification does not enable or disclose how to “connect” tunnel end points “with a unicast tunnel using the new tunnel packet.” The disclosure describes adding headers to create a “new tunnel packet” for forwarding, and separately describes that a tunnel manager may configure tunnels, but it does not explain how a “new tunnel packet” is used to create or establish a unicast tunnel between domains.

- “process the packet such that an L2 switching may be performed” (claims 1, 7, 13)

Defendant asserts the specification (including any specification to which priority is sought) fails to describe the full scope of the claimed inventions in sufficient detail such that one skilled in the art could reasonably conclude that the inventor had possession of the claimed invention, i.e., one that involves processing the packet such that an L2 switching may be performed. The specification does not enable or disclose what is involved in “processing the packet” “such that” L2 switching “may be performed.” At most, the specification describes that the packet is “modified” in some undescribed way to such that L2 switching may be performed. ’961 Patent, 6:56-59.

IV. SUBJECT MATTER ELIGIBILITY CONTENTIONS

35 U.S.C. § 101 provides that “any new and useful process, machine, manufacture, or composition of matter” is eligible for patenting, subject to long-recognized exceptions for laws of nature, natural phenomena, and abstract ideas. *Alice Corp. v. CLS Bank Int’l*, 573 U.S. 208, 216 (2014); *Mayo Collaborative Servs. v. Prometheus Labs., Inc.*, 566 U.S. 66, 70-73 (2012).

Under the first step of the *Alice* analysis, the relevant inquiry is whether a patent claim is directed to an abstract idea. *Alice*, 573 U.S. at 218. The Federal Circuit and the Supreme Court have repeatedly held that claims that state results in purely functional language and do not explain how to achieve those results are directed to abstract ideas at *Alice* Step One. *See, e.g., Interval Licensing LLC v. AOL, Inc.*, 896 F.3d 1335, 1342-43 (Fed. Cir. 2018) (explaining that claims reciting a “result” or “effect” without specifying how to achieve it in a non-abstract way are directed to an abstract idea).

Under *Alice* step two, a claim that is directed to an abstract idea can only survive if it contains an inventive concept, an element or combination of elements that transforms the claims into something “significantly more” than a claim on the patent-ineligible concept itself. *Alice*, 573 U.S. at 222. Merely introducing a computer into the claims does not satisfy the second step. *See Alice*, 573 U.S. at 220 (distinguishing between simply implementing an abstract idea on a computer (patent ineligible) and using a computer to improve an existing technological process that requires additional steps (patent eligible)); *see also BASCOM Glob. Internet Servs., Inc. v. AT&T Mobility LLC*, 827 F.3d 1341, 1352 (Fed. Cir. 2016) (claims to “an abstract idea implemented on generic computer components, without providing a specific technical solution beyond simply using generic computer concepts in a conventional way” do not suffice at step two.); *Intell. Ventures I LLC v. Capital One Bank (USA)*, 792 F.3d 1363, 1367 (Fed. Cir. 2015) (merely “claiming the

improved speed or efficiency inherent with applying the abstract idea on a computer [does not] provide a sufficient inventive concept.” (citation omitted)). A claim that recites an abstract idea must include “additional features” to ensure “that the [claim] is more than a drafting effort designed to monopolize the [abstract idea].” *Alice*, 537 U.S. at 221.

Furthermore, the transformation of ineligible claims into a patent-eligible application requires “more than simply stat[ing] the [abstract idea] while adding the words ‘apply it.’” *Id.*

A. The ’101 Patent Is Not Directed to Eligible Subject Matter

The Asserted Claims of the ’101 Patent are directed to patent-ineligible subject matter under 35 U.S.C. § 101 because (1) they are directed to the abstract idea of determining if a time has arrived for a time-dependent action, and (2) they lack an inventive concept. Applying this abstract idea in a particular technological environment, involving generic network architecture with generic components (policy server (PS), policy management system (PMS) and policy execution equipment (PEE)) all operating according to their ordinary functions, does not make it patent eligible.

Defendant contends that asserted claim 1 is representative of all Asserted Claims of the ’101 Patent, as all claims of the ’101 Patent are “substantially similar and linked to the same abstract idea.” *Intell. Ventures I LLC v. Symantec Corp.*, 838 F.3d 1307, 1317 (Fed. Cir. 2016). Plaintiff failed to assert distinctiveness between each claim of the ’101 Patent in its infringement contentions or in the Complaint. *See Wireless Discovery LLC v. Coffee Meets Bagel, Inc.*, 654 F. Supp. 3d 347, 355 (D. Del. 2023) (holding asserted claim to be representative because *plaintiff* failed to present meaningful argument regarding distinctiveness). Moreover, claim 2 depends from claim 1 and is directed to the same abstract idea and merely adds limitations that were well understood, routine, and conventional in the relevant industry at the relevant time. Claims 3-6 are

directed to the similar abstract idea of determining if a time has arrived for a time-dependent action, and similarly do not recite additional limitations amounting to significantly more than the abstract idea itself. Plaintiff's infringement contentions confirm this by relying on the same analysis and evidence for claims 3-6 as for claim 1. Should Plaintiff seek and obtain leave to amend their Infringement Contentions, Defendant reserves its right to respond accordingly by amending these Subject Matter Eligibility Contentions, including to address the representativeness of claim 1.

As explained above, prior art and the '101 Patent's own specification and prior art cited in these invalidity contentions show that service provider networks, components, and procedures were well-understood, routine, and conventional as of the alleged priority date. Each element of claim 1 was well-understood, routine, and conventional as of the alleged priority date, as explained in greater detail in the chart below:

[1pre] A method of determining an execution time point of a time-based policy that controls traffic, the method comprising:	The preamble describes the abstract concept of determining if a time has arrived for a time-dependent action (i.e., “determining an execution time point of a time-based policy that controls traffic”). Moreover, controlling traffic with policies, including time-based policies, and also the necessity to determine whether the execution time of a time-based policy has arrived, was well-understood, routine, and conventional as of the alleged priority date. By way of non-limiting examples, '101 Patent itself and the prior art demonstrate this: '101 Patent at 1:22-30: “Traffic generally indicates a load that is applied to a communication apparatus or a system, and in a network, traffic indicates a quantity of data that are exchanged through the network. In order to provide a smooth network service or to apply a reasonable accounting model, a service provider monitors and controls a traffic situation. In this case, a policy is set to operate a network, and when transmission of data is requested from the outside of a server through the network, a service provider controls traffic according to the set policy.” Delsesto (D-01) at [0003], [0005], [0052], [0058] Mo (D-02) at Abstract, [0028], FIG. 4
---	--

	ZTE WO ‘188 (D-03) at Abstract, p. 13 Hitachi JP ‘294 (D-04) at [0002], [0017], [0026], Nortel EP ‘574 (D-05) at [0018], [0019], [0051], [0103], [0132], [0135] Gai (D-06) at Abstract, 1:12-63, 3:61-4:12 Time-Based Access Lists Using Time Ranges, at pp. 1-2 Cuervo (D-07) at Abstract, 4:4-23, FIG. 2 Corresponding element at Ex. D-08 (Cisco Secure Access Control Server (ACS) System Corresponding element at Ex. D-09 (Cisco Security Manager (CSM)) Corresponding element at Ex. D-10 (Cisco SCE/SEF (IMS))
[1a] determining, by a policy server (PS) to which the policy is set, whether an execution time point of the policy arrives;	This element essentially recites the abstract concept of determining if a time has arrived for a time-dependent action (i.e., “determining ... whether an execution time point of the policy arrives”). Moreover, setting a policy with a policy server was well-understood, routine, and conventional as of the alleged priority date, as were time-based policies and thus determining whether an execution point of the policy arrives. By way of non-limiting examples, the ’101 Patent itself and the prior art demonstrate this: ’101 Patent at 1:37-39: “A policy that is set at a policy server is executed in a policy execution point (PEP) and controls traffic.” Delsesto (D-01) at [0003]-[0005], [0009], [0021]-[0026], [0052], [0058], FIGS 1-2 Mo (D-02) at [0029]-[0036], [0038]-[0041], [0045], [0048], FIGS. 1-4 ZTE WO ‘188 (D-03) at pp. 1-8, 10 Hitachi JP ‘294 (D-04) at [0002], [0004], [0008], [0013]- [0017], [0024], [0026], [0034], [0036], Claim 4 Nortel EP ‘574 (D-05) at [0019], [0056], [0057], [0074]-[0077], [0084], [0102], [0103], [0132], [0135] Gai (D-06) at 3:61-4:12, 7:55-8:5, 9:59-10:9, FIGS. 3-4 Time-Based Access Lists Using Time Ranges, at pp. 6-10, 12 Cuervo (D-07) at 2:54-58, 4:4-23, 4:24-55, 4:56-67, 5:1-35, 5:36-6:14; 6:45-52, FIG. 2 Corresponding element at Ex. D-08 (Cisco Secure Access Control Server (ACS) System

	Corresponding element at Ex. D-09 (Cisco Security Manager (CSM)) Corresponding element at Ex. D-10 (Cisco SCE/SEF (IMS))
[1b] continuing, if an execution time point of the policy does not arrive, to determine an execution time point of the policy;	This element essentially recites the abstract concept of determining if a time has arrived for a time-dependent action (i.e., “continuing, if an execution time point of the policy does not arrive, to determine an execution time point of the policy”). Moreover, time-based policies and thus determining and continuing to determine whether an execution point of the policy arrives was well-understood, routine, and conventional as of the alleged priority date. By way of non-limiting examples, the prior art demonstrates this: Delsesto (D-01) at [0052], [0058], Mo (D-02) at [0034]-[0036], [0038]-[0041], [0044]-[0045], [0048] ZTE WO ‘188 (D-03) at pp. 3, 7 Hitachi JP ‘294 (D-04) at [0034], [0036] Nortel EP ‘574 (D-05) at [0019], [0084], [0132], Gai (D-06) at 3:61-4:12, 7:55-8:5, 9:59-10:9, FIGS. 3-4 Time-Based Access Lists Using Time Ranges, at pp. 1-2 Cuervo (D-07) at 6:45-52, 7:21-30, FIG. 2 Corresponding element at Ex. D-08 (Cisco Secure Access Control Server (ACS) System Corresponding element at Ex. D-09 (Cisco Security Manager (CSM)) Corresponding element at Ex. D-10 (Cisco SCE/SEF (IMS))
[1c-1] transmitting, if an execution time point of the policy arrives, the policy to a policy management system (PMS), wherein the PMS is connected to policy execution equipment (PEE) that executes the policy, and wherein the execution time point reflects a time condition that is provided independently from the traffic,	Transmitting policies withing a service provider network, including to a policy management system connected to policy execution equipment that executes the policy was well-understood, routine, and conventional as of the alleged priority date, as was condition based (i.e., triggered) transmission of policies. By way of non-limiting examples, the ‘101 Patent itself and the prior art demonstrate this: ’101 Patent at 1:37-39: “A policy that is set at a policy server is executed in a policy execution point (PEP) and controls traffic.” ’101 Patent at 3:11-13: “Referring to FIG. 1, a time-based policy is set at a policy server (PS) 100 and is executed in policy execution equipment (PEE) 300 via a policy management system (PMS) 200.”

	Delsesto (D-01) at [0009], [0052], [0058], Mo (D-02) at [0029]-[0033], [0035], [0041], [0045] ZTE WO ‘188 (D-03) at pp. 6-8, 10-11 Hitachi JP ‘294 (D-04) at [0024], [0026], [0032], [0034], [0040], [0043], [0044], FIGS. 1, 4 Nortel EP ‘574 (D-05) at [001], [0020], [0021], [0056], [0057], [0075]-[0077], [0088]-[0090], [0102], [0103], [0118], [0119], [0002], [0132], [0135]. Gai (D-06) at 1:12-64, 3:61-4:12, 5:55-59, 5:63-6:26, 7:11-29, 7:55-8:5, 9:36-58, 9:59-10:9, 10:12-34, 13:60-15:14 17:49-62, 19:25-36, FIGS. 3-5 Time-Based Access Lists Using Time Ranges, at pp. 6-10, 12 Cuervo (D-07) at 2:54-58, 4:4-23, 4:24-55, 5:27-35, 6:45-52, 5:36-6:13, FIG. 2 Corresponding element at Ex. D-08 (Cisco Secure Access Control Server (ACS) System Corresponding element at Ex. D-09 (Cisco Security Manager (CSM)) Corresponding element at Ex. D-10 (Cisco SCE/SEF (IMS)
[1c-2] wherein the time-based policy is set based on tuple information that is included in a packet from the traffic and the time condition by the policy server.	Setting a policy based on tuple information that is included in a packet from traffic, and by a time condition set by a policy server, was well-understood, routine, and conventional as of the alleged priority date, as was condition based (i.e., triggered) transmission of policies. By way of non-limiting examples, the ‘101 Patent itself and the prior art demonstrate this: ’101 Patent at 1:23-39: “In this case, a policy is set to operate a network, and when transmission of data is requested from the outside of a server through the network, a service provider controls traffic according to the set policy. Such traffic includes a packet in which information Such as a 5-tuple (source IP address, destination IP address, source port, destination port, and protocol) and a 7-tuple (source IP address, destination IP address, source port, destination port, protocol, input port number, and signature) is stored, and a general policy is set based on information Such as a 7-tuple that is included in a packet of traffic. A policy that is set at a policy server is executed in a policy execution point (PEP) and controls traffic.” Delsesto (D-01) at [0007]-[0009], [0022]-[0029], [0052], [0058], [0074], [0089], Claim 11

	Mo (D-02) at [0001], [0002], [0006], [0007], [0029], [0034], [0035], [0043], Claim 1 ZTE WO ‘188 (D-03) at pp. 6, 9, 10 Hitachi JP ‘294 (D-04) at [0010], [0026], FIG. 3 Nortel EP ‘574 (D-05) at [0018]-[0019], [0067], [0082]-[0084], [0110]-[0111], [0126]. [0132]. [0135] Gai (D-06) at 1:13-64, 1:66-2:10, 3:6-29, 12:41-67, 13:60-15:14, 15:38-42 Time-Based Access Lists Using Time Ranges, at pp. 14-17 Cuervo (D-07) at 4:24-55 Corresponding element at Ex. D-08 (Cisco Secure Access Control Server (ACS) System Corresponding element at Ex. D-09 (Cisco Security Manager (CSM)) Corresponding element at Ex. D-10 (Cisco SCE/SEF (IMS)
--	--

These citations are illustrative only; for each element, Defendant may rely on any evidence charted or otherwise cited for that element in these contentions or Defendant’s Invalidity Contentions and related exhibits, or in the arguments and statements made in the prosecution history of the ’101 Patent.

Defendant may also rely on percipient and/or expert testimony to supplement, summarize, and/or explain the state of the industry at the time the earliest application to which the ’101 Patent claims priority was filed. Any such witnesses will be disclosed pursuant to the Federal Rules of Civil Procedure and the Court’s Docket Control Order in this case for initial, additional, and expert disclosures.

Setting aside the routine, conventional, and well-understood components and operations identified above, what remains of the Asserted Claims is the abstract concept of determining if a time has arrived for a time-dependent action. The claims 1 and 3 on their face are directed to “a method of determining an execution time point of a time-based policy that control traffic.” All three independent claims recite “determining ... whether an execution time point of the policy arrives” and “continuing, if an execution time point of the policy does not arrive, to determine an

execution time point of the policy.” These limitations do no more than recite the abstract idea itself, which is one that could be done mentally by a human, as humans have done since the concept of time was known (e.g., look at clock and evaluate if it is time to do something (e.g., time to go to work); look at the sun and determine if it is time to do something (e.g., time to milk the cows)). The remaining claim element offer on only generic, routine, conventional, and well understood examples of components (i.e., a policy server, policy management system and policy execution system) and operations (i.e., transmitting a policy), and the claim simply recite that the abstract idea is performed at these conventional components (i.e., at the policy server (claim 1), at the policy management system (claim 3) or at the policy execution equipment (claim 6). Indeed, the transformation of ineligible claims into a patent-eligible application requires “more than simply stat[ing] the [abstract idea] while adding the words ‘apply it.’” *Alice*, 573 U.S. at 221. But that is precisely what claim 1 does.⁴ Beyond that, the claims recite only generic information-handling functions transmitting (claim 1) or downloading (claims 3 and 6) policies and storing (claim 3) policies., and executing the policy (claim 6). These functions are classic markers of abstractness. *See e.g. Electric Power Group*, 830 F.3d 1350, 1353-55 (Fed. Cir. 2016); *and Univ. of Florida Rsch. Found., Inc v. General Elec. Co.*, 916 F.3d 1363, 1367-69 (Fed. Cir. 2019). Even the conditions of the policy (“time condition that is provided independently of the traffic” and “tuple

⁴ Moreover, merely introducing a computer into the claims does not satisfy the second step. *See Alice*, 573 U.S. at 220 (distinguishing between simply implementing an abstract idea on a computer (patent ineligible) and using a computer to improve an existing technological process that requires additional steps (patent eligible)); *see also BASCOM Glob. Internet Servs., Inc. v. AT&T Mobility LLC*, 827 F.3d 1341, 1352 (Fed. Cir. 2016) (Claims to “an abstract idea implemented on generic computer components, without providing a specific technical solution beyond simply using generic computer concepts in a conventional way” do not suffice at step two.); *Intellectual Ventures I LLC v. Capital One Bank (USA)*, 792 F.3d 1363, 1367 (Fed. Cir. 2015) (citation omitted) (merely “claiming the improved speed or efficiency inherent with applying the abstract idea on a computer [does not] provide a sufficient inventive concept.”). A claim that recites an abstract idea must include “additional features” to ensure “that the [claim] is more than a drafting effort designed to monopolize the [abstract idea].” *Alice*, 537 U.S. at 221.

information that is included in a packet”) are routine, conventional and well-understood and not-inventive as set forth in these contentions.

The legal bases for the ineligibility of the Asserted Claims include that the claims are similar to ones that several other courts have found to be abstract, including but not limited to:

1. *Affinity Labs of Texas, LLC v. DIRECTV, LLC*, 838 F.3d 1253 (Fed. Cir. 2016): holding that claims “entirely functional in nature” directed to out-of-region delivery of broadcast content to a cellphone, implemented with a generic GUI and downloadable app, were abstract because they recited desired results without explaining how to achieve them.
2. *First-Class Monitoring, LLC v. United Parcel Service of Am.*, 389 F. Supp. 3d 456 (E.D. Tex. 2019): finding and dismissing claims as abstract that recited only “requesting, collecting, analyzing, and transmitting information” via conventional SMS because they were purely functional steps that claim a result, not a specific means.
3. *Specialized Monitoring Sol’ns, LLC v. ADT LLC*, 367 F. Supp. 3d 575 (E.D. Tex. 2019): holding claims were abstract and dismissing because they recited outcome-focused operations—detecting events, encoding packet messages, storing data in categorized areas, and displaying results—without any specific technique for how to implement those results beyond using generic technology.
4. *VoIP-Pal.com, Inc. v. Apple Inc.*, 375 F. Supp. 3d 1110 (N.D. Cal. 2019): holding claims related to call routing based on caller/callee characteristics abstract where written in “vague, functional terms” (“locating,” “determining,” “classifying”), analogous to human switchboard operations, and no inventive concept beyond generic computing.
5. *Intellectual Ventures I LLC v. Symantec Corp.*, 838 F.3d 1307, 1316-1318, 1320-21 (Fed. Cir. 2016). In this case, the Federal Circuit relied upon the specification when explaining that the claimed electronic post office, which recited limitations describing how the system would receive, screen and distribute email on a computer network, was analogous to how a person decides whether to read or dispose of a particular piece of mail and that “with the exception of generic computer-implemented steps, there is nothing in the claims themselves that foreclose them from being performed by a human, mentally or with pen and paper”. 838 F.3d at 1318.
6. *Realtime Tracker, Inc. v. RELX, Inc.*, 659 F.Supp.3d 384, 397 (S.D.N.Y. March 7, 2023), *aff’d*, 2024 WL 4746162 (Fed. Cir. 2024): “Courts have also often found that processes that can be accomplished mentally or within the human mind draw upon abstract ideas. That is because the ‘application of only human intelligence to the solution of practical problems is no more than a claim to a fundamental

principle.’ *CyberSource Corp. v. Retail Decisions, Inc.*, 654 F.3d 1366, 1371 (Fed. Cir. 2011) (alterations omitted) (quoting *Bilski*, 545 F.3d at 965); *see also, e.g., Mortg. Grader*, 811 F.3d at 1324–25. In this vein, the Federal Circuit has ‘treated analyzing information by steps people go through in their minds, or by mathematical algorithms, without more, as essentially mental processes within the abstract-idea category.’ *Elec. Power Grp., LLC v. Alstom S.A.*, 830 F.3d 1350, 1354 (Fed. Cir. 2016).” *See also id.* at 402: “In essence, the Patent recites the abstract concept of timekeeping for compensation. Such is a ‘fundamental economic practice long prevalent in our system of commerce,’ *Alice*, 573 U.S. at 219, 134 S.Ct. 2347 (citation omitted), to which professionals in numerous fields in which client billing is commonly based on unit of time worked (for example, lawyers, accountants, and architects) could attest.”

7. *FairWarning IP, LLC v. Iatric Sys., Inc.*, 839 F.3d 1089 (Fed. Cir. 2016). The patentee in *FairWarning* claimed a system and method of detecting fraud and/or misuse in a computer environment, in which information regarding accesses of a patient’s personal health information was analyzed according to one of several rules (*i.e.*, related to accesses in excess of a specific volume, accesses during a pre-determined time interval, or accesses by a specific user) to determine if the activity indicates improper access. 839 F.3d. at 1092. The court determined that these claims were directed to a mental process of detecting misuse, and that the claimed rules here were “the same questions (though perhaps phrased with different words) that humans in analogous situations detecting fraud have asked for decades, if not centuries.” 839 F.3d. at 1094-95.

Consistent with these authorities, the ’101 Patent claims are abstract because they do not recite any particularized technical means and instead only describe their desired results. Defendant further incorporates by reference MPEP § 2106 (and all subsections) regarding patent subject matter eligibility, including the specific examples and analyses of abstract ideas.

Under *Alice* Step Two, a claim that is directed to an abstract idea can only survive if it contains an inventive concept, an element or combination of elements that transforms the claims into something “significantly more” than a claim on the patent-ineligible concept itself. *Alice*, 573 U.S. at 222. Claim 1 fails to recite an inventive concept sufficient to transform the claimed abstract idea into a patent-eligible application of the abstract idea. As discussed above, Claim 1 merely recites the generic, conventional components that comprise a service provider network, including a policy server, policy management system and policy execution equipment, and their

conventional functionality (e.g., using a policy server to set a policy and transmit a policy to a policy management system, using a policy management system to store and manage policies for policy execution equipment, and using a policy execution equipment to execute a policy). Indeed, the names of these components themselves reflect their function. Well-understood, routine, and conventional activities cannot supply the inventive concept. *BSG Tech LLC v. BuySeasons, Inc.*, 899 F.3d 1281, 1290-91 (Fed. Cir. 2018). The claim contains nothing that explains how the abstract idea of determining if a time has arrived for a time-dependent action is achieved in an inventive way. As shown in the chart above, and in Defendant's invalidity charts for the remaining claims, each element of Claim 1 contains routine, conventional, and well-understood concepts. Thus, the claim lacks an inventive concept and fails at Alice Step Two.

V. TECHNICAL DOCUMENT PRODUCTION

Pursuant to Local Patent Rule 3-4(a), Defendant has diligently collected technical documents relating to the products identified in the Infringement Contentions produced at CSI-OPT_00000001 - CSI-OPT_00982971 contemporaneously with these Invalidity Contentions.

VI. PRIOR ART DOCUMENT PRODUCTION

Pursuant to Local Patent Rule 3-4(b), Defendant is producing prior art and related evidence concurrent with these contentions, including at CSI-OPT-PA_00000001 - CSI-OPT-PA_0015013.

Dated: February 11, 2026

/s/ Joseph A. Powers

Matthew C. Gaudet

DUANE MORRIS LLP

1075 Peachtree Street NE, Suite 1700

Atlanta, GA 30309-3929

Tel: (404) 253-6902

mcgaudet@duanemorris.com

Tia D. Fenton

Tel: (202) 776-7825

tdfenton@duanemorris.com

Elissa L. Sanford

Tel: (202) 776-5231

esanford@duanemorris.com

DUANE MORRIS LLP

901 New York Ave. N.W.

Suite 700 East

Washington, DC 20001-4795

Joseph A. Powers

DUANE MORRIS LLP

30 South 17th Street

Philadelphia, PA 19103-4196

Tel: (215) 979-1842

japowers@duanemorris.com

Melissa R. Smith (TBN 24001351)

GILLAM & SMITH, LLP

303 South Washington Avenue

Marshall, Texas 75670

Tel: (903) 934-8450

Facsimile: (903) 934-9257

Counsel for Defendant Cisco Systems, Inc.

CERTIFICATE OF SERVICE

I hereby certify that on February 11, 2026, the foregoing document was served via e-mail on all counsel of record for Plaintiff.

/s/ Joseph A. Powers
Joseph A. Powers