EXHIBIT B-1
Defendant's Preliminary Invalidity Contentions
Orckit Corporation v. Cisco Systems, Inc., 2:22-cv-00276-JRG-RSP

Chart for U.S. Patent 8,830,821 ("the '821 Patent")
Cisco IOS Multiprotocol Label Switching
Cisco IOS Release 12.2SR ("Cisco IOS System")

As shown in the chart below, all claims of the '821 Patent are invalid (1) under 35 U.S.C. §§ 102 (a) and (g) based on Cisco IOS System, (2) under 35 U.S.C. §§ 102 (a) and (b) because the references describing the Cisco IOS System disclose every limitation of every Asserted Claim, and (3) under 35 U.S.C. § 103 as obvious based on Cisco IOS Multiprotocol Label Switching Configuration Guide Release 12.2SR alone, or in combination with the knowledge of a person of ordinary skill in the art alone, and in further combination with references specifically identified in the following claim chart and/or one or more other references identified in Cisco's Preliminary Invalidity Contentions. The Cisco IOS System comprises the Cisco IOS Release 12.2SR operating system, as it ran on various Cisco switches and routers from before June 22, 2011. The Cisco IOS System was invented, known and used by others in the United States, and on sale in the United States before the claimed invention, rendering the claims invalid. Publications describing the Cisco IOS System, including those cited in this chart, were publicly disclosed and available more than a year before the purported filing date of the provisional application identified on the face of the '821 Patent, rendering the claims invalid. For example, Cisco IOS Multiprotocol Label Switching Configuration Guide Release 12.2SR was published in 2010 and discloses all elements of the Asserted Claims, thereby anticipating the Asserted Claims. Moreover, as identified in this chart, it would have been obvious to combine the disclosures in the publications regarding the Cisco IOS System with each other and with disclosures in other publications known in the art, as explained in this chart, well before the claimed invention and the effective filing date of the '821 Patent, rendering the claims invalid.

The quotations and diagrams regarding the Cisco IOS System in this chart come from the following sources:

- Cisco IOS Multiprotocol Label Switching Configuration Guide Release 12.2SR ("Cisco IOS Guide"); and
- Cisco Press Release titled "Cisco Systems' MPLS-TE 'Fast Reroute'" Function Introduced To NTT Communications' Arcstart IP-VPN" ("Cisco Press Release"); and
- Open-Shortest Path First (OSPF) Protocol, Defined in RFC 2328, https://www.ietf.org/rfc/rfc2328.txt ("RFC 2328").

Motivations to combine the disclosures in these references with disclosures in other publications known in the art, as explained in this chart, include at least the similarity in subject matter between the references to the extent they concern methods of selecting paths in a

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 1 of 698

network and/or determining the cost of potential paths in a network. Insofar as the references cite other patents or publications, or suggest additional changes, one of ordinary skill in the art would look beyond a single reference to other references in the field.

These invalidity contentions are based on Defendant's present understanding of the Asserted Claims, and Orckit's apparent construction of the claims in its November 3, 2022 Disclosure of Asserted Claims and Infringement Contentions Pursuant to P.R. 3-1, and Orckit's January 19, 2023 First Amended Disclosure of Asserted Claims and Infringement Contentions Pursuant to P.R. 3-1 (Orckit's "Infringement Disclosures"), which is deficient at least insofar as it fails to cite any documents or identify accused structures, acts, or materials in the Accused Products with particularity. Defendant does not agree with Orckit's application of the claims, or that the claims satisfy the requirements of 35 U.S.C. § 112. Defendant's contentions herein are not, and should in no way be seen as, admissions or adoptions as to any particular claim scope or construction, or as any admission that any particular element is met by any accused product in any particular way. Defendant objects to any attempt to imply claim construction from this chart. Defendant's prior art invalidity contentions are made in a variety of alternatives and do not represent Defendant's agreement or view as to the meaning, definiteness, written description support for, or enablement of any claim contained therein.

The following contentions are subject to revision and amendment pursuant to Federal Rule of Civil Procedure 26(e), the Local Rules, and the Orders of record in this matter subject to further investigation and discovery regarding the prior art and the Court's construction of the claims at issue.

No.	'821 Patent Claim 1	Cisco IOS System
1[preamble]	An entity selection method performed by a network device, comprising the steps of:	Cisco IOS System discloses an entity selection method performed by a network device, comprising the steps of. For example, Cisco IOS Guide discloses a cisco router supporting Cisco IOS Release 12.2SR which supports MPLS traffic engineering. Cisco IOS Guide further discloses that a Cisco router can act as the headend router of a TE tunnel, by controlling the path its traffic takes to a particular destination. Cisco IOS Guide at Page No. 15 ("[Cisco IOS documentation describes the tasks and commands available to configure and maintain Cisco networking devices.]") Cisco IOS Guide at Page No. 15 ("[The Cisco IOS documentation set is intended for users who configure and maintain Cisco networking devices (such as routers and switches) but who may not be familiar with the configuration and maintenance tasks, the relationship among

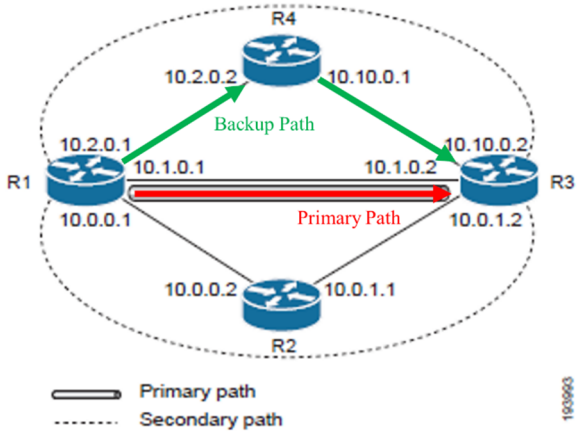
Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 2 of 698

No.	'821 Patent Claim 1	Cisco IOS System
		tasks, or the Cisco IOS commands necessary to perform particular tasks. The Cisco IOS documentation set is also intended for those users experienced with Cisco IOS software who need to know about new features, new configuration options, and new software characteristics in the current Cisco IOS release.]") Cisco IOS Guide at Page No. 15 ("[In Cisco IOS documentation, the term router may be used to refer to various Cisco products; for example, routers, access servers, and switches.]") Cisco IOS Guide at Page No. 603 ("[Traffic Engineering Tunnels MPLS TE lets you build label switched paths (LSPs) across your network for forwarding traffic. MPLS TE LSPs let the headend of a TE tunnel control the path its traffic takes to a particular destination. This method is more flexible than forwarding traffic based only on a destination address. Interarea tunnels allow you to do the following: • Build TE tunnels between areas (interarea tunnels) • Build TE tunnels that start and end in the same area, on multiple areas on a router (intra-area tunnels) Some tunnels are more important than others. For example, you may have tunnels carrying VoIP traffic and tunnels carrying data traffic that are competing for the same resources. MPLS TE allows you to have some tunnels preempt others. Each tunnel has a priority, and more-important tunnels take precedence over less-important tunnels.]") Cisco IOS Guide at Page No. 633 ("[router—A network layer device that uses one or more metrics to determine the optimal path along which network traffic should be forwarded. Routers forward packets from one network to another based on network layer information.]")
1[a]	providing a plurality of multi protocol label switching (MPLS) transport entities	Cisco IOS System discloses providing a plurality of multi protocol label switching (MPLS) transport entities between a first endpoint and a second endpoint. For example, Cisco IOS Guide discloses that Cisco IOS Release 12.2SR supports MPLS traffic engineering tunnels with enhanced path protection, including by providing support of

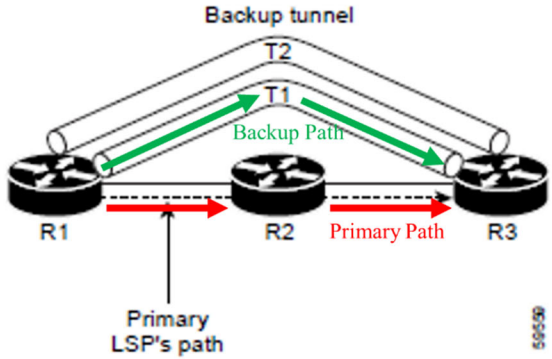
Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 3 of 698

No.	'821 Patent Claim 1	Cisco IOS System
	between a first endpoint and a second endpoint;	multiple backup path options (up to eight backup paths) per primary path option terminating at the same destination. Cisco IOS Guide at Page No. 601 ("[The MPLS Traffic Engineering (TE): Path Protection feature provides an end-to-end failure recovery mechanism (that is, full path protection) for Multiprotocol Label Switching (MPLS) traffic engineering (TE) tunnels. Note Cisco IOS Release 12.3(33)SRE and later releases support enhanced path protection, which is the ability to configure up to eight secondary path options for a given primary path option.]") Cisco IOS Guide at Page No. 603 ("[Enhanced Path Protection Enhanced path protection provides support of multiple backup path options per primary path option. You can configure up to eight backup path options for a given primary path option. Only one of the configured backup path options is actively signaled at any time. After you enter the mpls traffic-eng path-option list command, you can enter the backup path priority in the number argument of the path-option command. A lower identifier represents a higher priority. Priorities are configurable for each backup path option. Multiple backup path options and a single backup path option cannot coexist to protect a primary path option.]") Cisco IOS Guide at Page No. 623 ("[Configuration Examples for MPLS Traffic Engineering (TE): Enhanced Path Protection]")

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 4 of 698

No.	'821 Patent Claim 1	Cisco IOS System
		Figure 6 shows the network topology for enhanced path protection. ***Figure 6 Network Topology for Enhanced Path Protection***  Fig. 6 (annotation added) Cisco IOS Guide at Page No. 640 ("[Backup Tunnel Support Backup tunnel support has the following capabilities: • Backup Tunnels Can Terminate at the Next-Next Hop to Support FRR, page 6 • Multiple Backup Tunnels Can Protect the Same Interface, page 6 • Backup Tunnels Provide Scalability, page 7 Backup Tunnels Can Terminate at the Next-Next Hop to Support FRR Backup tunnels that terminate at the next-next hop protect both the downstream link and node. This provides protection for link and node failures. For more detailed information, see the "Node Protection" section on page 4. Multiple Backup Tunnels Can Protect the Same Interface There is no limit (except memory limitations) to the number of backup tunnels that can protect a given interface. In many topologies, support for node protection requires supporting multiple backup tunnels per protected interface. These backup tunnels can terminate at the same destination or at different destinations. That is, for a given protected interface, you can

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 5 of 698

No.	'821 Patent Claim 1	Cisco IOS System
		configure multiple NHOP or NNHOP backup tunnels. This allows redundancy and load balancing.]") Cisco IOS Guide at Page No. 643 ("[Backup Tunnels Terminating at the Same Destination Figure 4 shows how backup tunnels terminating at the same location can be used for redundancy and load balancing. Redundancy and load balancing work for both NHOP and NNHOP backup tunnels.]") *Figure 4 Backup Tunnels That Terminate at the Same Destination*  Fig. 4 (annotation added) Cisco IOS Guide at Page No. 644 ("[In this illustration, there are three routers: R1, R2, and R3. At R1 two NNHOP backup tunnels (T1 and T2) go from R1 to R3 without traversing R2. Redundancy—If R2 fails or the link from R1 to R2 fails, either backup tunnel can be used. If one backup tunnel is down, the other can be used. LSPs are assigned to backup tunnels when the LSPs are first established. This is done before a failure. Load balancing—If neither backup tunnel has enough bandwidth to back up all LSPs, both tunnels can be used. Some LSPs will use one backup tunnel, other LSPs will use the other backup tunnel. The router decides the best way to fit the LSPs onto the backup tunnels.]")

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 6 of 698

No.	'821 Patent Claim 1	Cisco IOS System
1[b]	determining an overall cost for each entity pair of said plurality of entities;	Cisco IOS System discloses determining an overall cost for each entity pair of said plurality of entities. For example, Cisco IOS Guide discloses that the path an LSP uses is determined by the LSP resource requirements and network resources, such as bandwidth. Cisco IOS Guide further discloses the MPLS TE SRLG feature that enhances backup tunnel path selection so a backup tunnel can avoid using links that are in the same SRLG as the interfaces it is protecting. Thus, at least under the apparent claim scope alleged by Orckit's Infringement Disclosures, this limitation is met. To the extent that the Cisco IOS System is found to not meet this limitation, determining an overall cost for each entity pair of said plurality of entities would have been obvious to a person having ordinary skill in the art, as explained below. Cisco IOS Guide at Page No. 253 ("[Path Option Selection for MPLS TE Tunnel LSPs This section contains the following topics about path option selection for MPLS TE Tunnel LSPs: • Constraint-Based Routing and Path Option Selection, page 5 • Tunnel Reoptimization and Path Option Selection, page 5 • Path Option Selection with Bandwidth Override, page 6 Constraint-Based Routing and Path Option Selection MPLS traffic engineering automatically establishes and maintains LSPs across the backbone by using the Resource Reservation Protocol (RSVP). The path that an LSP uses is determined by the LSP resource requirements and network resources, such as bandwidth. Traffic engineering tunnels are calculated at the LSP head based on a fit between required and available resources (constraint-based routing). Without the Path Option for Bandwidth Override feature, a TE tunnel establishes an LSP based on dynamic or explicit path options in order of preference. However, the bandwidth and other attributes configured on the TE tunnel allow the setup of an LSP only if LSP path options satisfy the constraints. If a path cannot be found that satisfies the configured path options, then the tunnel is not set up.]") Cisco IOS Guide at Page No. 639 ("[Bandwidth Protection

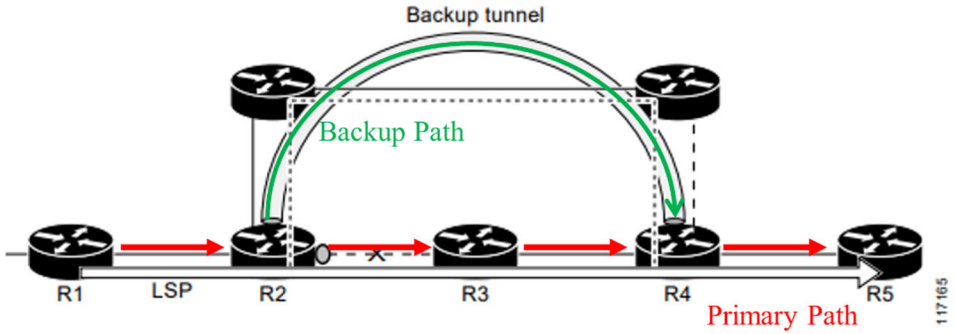
Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 7 of 698

No.	'821 Patent Claim 1	Cisco IOS System
		NHOP and NNHOP backup tunnels can be used to provide bandwidth protection for rerouted LSPs. This is referred to as backup bandwidth. You can associate backup bandwidth with NHOP or NNHOP backup tunnels. This informs the router of the amount of backup bandwidth a particular backup tunnel can protect. When a router maps LSPs to backup tunnels, bandwidth protection ensures that an LSP uses a given backup tunnel only if there is sufficient backup bandwidth. The router selects which LSPs use which backup tunnels in order to provide maximum bandwidth protection. That is, the router determines the best way to map LSPs onto backup tunnels in order to maximize the number of LSPs that can be protected. For information about mapping tunnels and assigning backup bandwidth, see the "Backup Tunnel Selection Procedure" section on page 10. LSPs that have the "bandwidth protection desired" bit set have a higher right to select backup tunnels that provide bandwidth protection; that is, those LSPs can preempt other LSPs that do not have that bit set. For more information, see the "Prioritizing Which LSPs Obtain Backup Tunnels with Bandwidth Protection" section on page 8.]") Cisco IOS Guide at Page No. 640 ("[Backup Tunnel Support Backup tunnel support has the following capabilities: • Backup Tunnels Can Terminate at the Next-Next Hop to Support FRR, page 6 • Multiple Backup Tunnels Can Protect the Same Interface, page 6 • Backup Tunnels Provide Scalability, page 7.]") Cisco IOS Guide at Page 640 ("[Multiple Backup Tunnels Can Protect the Same Interface There is no limit (except memory limitations) to the number of backup tunnels that can protect a given interface. In many topologies, support for node protection requires supporting multiple backup tunnels per protected interface. These backup tunnels can terminate at the same destination or at different destinations. That is, for a given protected interface, you can configure multiple NHOP or NNHOP backup tunnels. This allows redundancy and load balancing.]") Cisco IOS Guide at Page No. 603 ("[Traffic Engineering Tunnels MPLS TE lets you build label switched paths (LSPs) across your network for forwarding traffic.

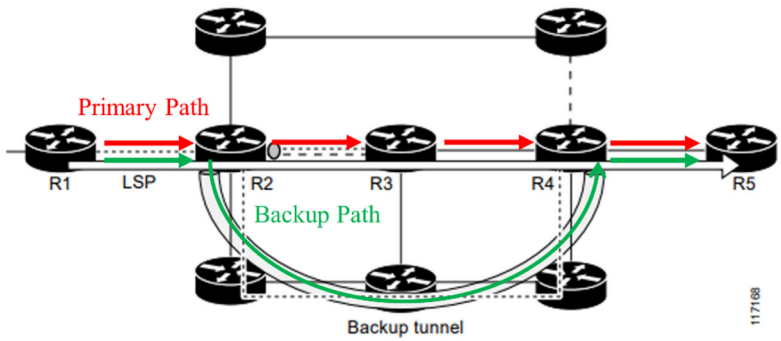
Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 8 of 698

No.	'821 Patent Claim 1	Cisco IOS System
		MPLS TE LSPs let the headend of a TE tunnel control the path its traffic takes to a particular destination. This method is more flexible than forwarding traffic based only on a destination address. Interarea tunnels allow you to do the following: • Build TE tunnels between areas (interarea tunnels) • Build TE tunnels that start and end in the same area, on multiple areas on a router (intra-area tunnels) Some tunnels are more important than others. For example, you may have tunnels carrying VoIP traffic and tunnels carrying data traffic that are competing for the same resources. MPLS TE allows you to have some tunnels preempt others. Each tunnel has a priority, and more-important tunnels take precedence over less-important tunnels.]") Cisco IOS Guide at Page No. 641 ("[Backup Bandwidth Protection Backup bandwidth protection allows you to give LSPs carrying certain kinds of data (such as voice) priority for using backup tunnels.]") Cisco IOS Guide at Page No. 650 ("[For each protected link AB with a maximum reservable subpool value of n, there may be a path from node A to node B such that the difference between the maximum reservable global and the maximum reservable subpool is at least the value of n. If it is possible to find such paths for each link in the network, you can establish all the backup tunnels along such paths without any bandwidth reservations. If there is a single link failure, only one backup tunnel will use any link on its path. Because that path has at least n available bandwidth (in the global pool), assuming that marking and scheduling is configured to classify the subpool traffic into a priority queue, the subpool bandwidth is guaranteed. This approach allows sharing of the global pool bandwidth between backup tunnels protecting independent link failures. The backup tunnels are expected to be used for only a short period of time after a failure (until the headends of affected LSPs reroute those LSPs to other paths with available subpool bandwidth). The probability of multiple unrelated link failures is very small (in the absence of node or shared risk link group (SRLG) failures, which result in multiple link failures). Therefore, it is reasonable to assume that link failures are in practice independent with high probability. This "independent failure assumption" in combination

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 9 of 698

No.	'821 Patent Claim 1	Cisco IOS System
		with backup tunnels signaled without explicit bandwidth reservation enables efficient bandwidth sharing that yields substantial bandwidth savings.]") Cisco IOS Guide at Page No. 651 ("[A similar approach can be used for node and SRLG protection. However, the decision of where to put the backup tunnels is more complicated because both node and SRLG failures effectively result in the simultaneous failure of several links. Therefore, the backup tunnels protecting traffic traversing all affected links cannot be computed independently of each other.]") Cisco IOS Guide at Page No. 579 ("[MPLS Traffic Engineering Shared Risk Link Groups SRLGs refer to situations where links in a network share a common fiber (or a common physical attribute). If one link fails, other links in the group may fail too. Links in the group have a shared risk. Backup tunnels should avoid using links in the same SRLG as interfaces they are protecting. Otherwise, when the protected link fails the backup tunnel fails too. ***Figure 1 Backup Tunnel in the Same SRLG as the Interface It Is Protecting***  Fig. 1 (annotation added) Figure 1 shows a primary label-switched path (LSP) from router R1 to router R5. The LSP protects against the failure of the R2-R3 link at R2 via a backup tunnel to R4. If the R2-R3 link fails, link protection reroutes the LSP along the backup tunnel. However, the R2-R3 link and one of the backup tunnel links are in the same SRLG. So if the R2-R3 link fails, the backup tunnel may fail too.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 10 of 698

No.	'821 Patent Claim 1	Cisco IOS System
		The MPLS TE SRLG feature enhances backup tunnel path selection so a backup tunnel can avoid using links that are in the same SRLG as the interfaces it is protecting. There are two ways for a backup tunnel to avoid the SRLGs of its protected interface: • The router does not create the backup tunnel unless it avoids SRLGs of the protected interface.]") Cisco IOS Guide at Page No. 580 ("[Open Shortest Path First (OSPF) and Intermediate System-to-Intermediate System (IS-IS) flood the SRLG membership information (including other TE link attributes such as bandwidth availability and affinity) so that all routers in the network have the SRLG information for each link. With this topology information, routers can compute backup tunnel paths that exclude links having SRLGs in common with their protected interfaces. As shown in Figure 2, the backup tunnel avoids the link between R2 and R3, which shares an SRLG with the protected interface.]") *Figure 2 Backup Tunnel That Avoids SRLG of Protected Interface*  Fig. 2 (annotation added) Cisco IOS Guide at Page No. 587 ("[Step 4 show mpls traffic-eng topology Use this command to show the SRLG link membership flooded via the Interior Gateway Protocol (IGP).]") Cisco IOS Guide at Page 632

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 11 of 698

No.	'821 Patent Claim 1	Cisco IOS System
		("[IS-IS—Intermediate System-to-Intermediate System. Link-state hierarchical routing protocol that calls for intermediate system (IS) routers to exchange routing information based on a single metric to determine network topology.]") Cisco IOS Guide at Page No. 633 ("[OSPF—Open Shortest Path First. A link-state hierarchical Interior Gateway Protocol routing algorithm, derived from the IS-IS protocol. OSPF features include least-cost routing, multipath routing, and load balancing.]") Cisco IOS Guide at Page No. 581 ("[Autotunnel Backup for MPLS TE SRLGs Autotunnel backup is the ability of routers to create backup tunnels automatically. Therefore, you do not need to preconfigure each backup tunnel and then assign the backup tunnel to the protected interface. Only automatically created backup tunnels can avoid SRLGs or their protected interfaces.]") Cisco IOS Guide at Page No. 583 ("[Configuring the Routers That Automatically Create Backup Tunnels to Avoid MPLS TE SRLGs of Their Protected Interfaces Perform the following task to configure routers that automatically create backup tunnels to avoid MPLS TE SRLGs of their protected interfaces. Backup tunnels provide link protection by rerouting traffic to the next hop bypassing failed links or in this instance by avoiding SRLGs.]") Cisco IOS Guide at Page No. 644 ("[Backup Tunnel Selection Procedure When an LSP is signaled, each node along the LSP path that provides FRR protection for the LSP selects a backup tunnel for the LSP to use if either of the following events occurs: • The link to the next hop fails. • The next hop fails. By having the node select the backup tunnel for an LSP before a failure occurs, the LSP can be rerouted onto the backup tunnel quickly if there is a failure.]") Cisco IOS Guide at Page 646 ("[NHOP Versus NNHOP Backup Tunnels More than one backup tunnel can protect a given LSP, where one backup tunnel terminates at the LSP's NNHOP, and the other terminates at the LSP's NHOP. In this case, the router

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 12 of 698

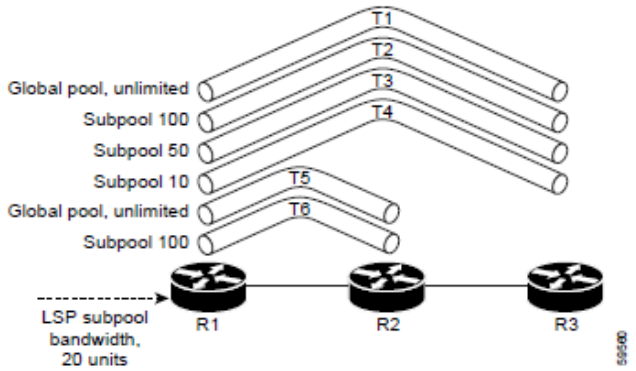
<table>
<tr><th>No.</th><th>’821 Patent Claim 1</th><th>Cisco IOS System</th></tr>
<tr><td></td><td></td><td>chooses the backup tunnel that terminates at the NNHOP (that is, FRR prefers NNHOP over NHOP backup tunnels).Table 1 lists the tunnel selection priorities. The first choice is an NNHOP backup tunnel that acquires its bandwidth from a subpool or global pool, and has limited bandwidth. If there is no such backup tunnel, the next choice (2) is a next-next hop backup tunnel that acquires a limited amount of bandwidth from any pool. The preferences go from 1 (best) to 8 (worst), where choice 3 is for an NNHOP backup tunnel with an unlimited amount of subpool or global-pool bandwidth.]”)

Table 1 Tunnel Selection Priorities

<table>
<tr><th>Preference</th><th>Backup Tunnel Destination</th><th>Bandwidth Pool</th><th>Bandwidth Amount</th></tr>
<tr><td>1 (Best)</td><td>NNHOP</td><td>Subpool or global pool</td><td>Limited</td></tr>
<tr><td>2</td><td>NNHOP</td><td>Any</td><td>Limited</td></tr>
<tr><td>3</td><td>NNHOP</td><td>Subpool or global pool</td><td>Unlimited</td></tr>
<tr><td>4</td><td>NNHOP</td><td>Any</td><td>Unlimited</td></tr>
<tr><td>5</td><td>NHOP</td><td>Subpool or global pool</td><td>Limited</td></tr>
<tr><td>6</td><td>NHOP</td><td>Any</td><td>Limited</td></tr>
<tr><td>7</td><td>NHOP</td><td>Subpool or global pool</td><td>Unlimited</td></tr>
<tr><td>8 (Worst)</td><td>NHOP</td><td>Any</td><td>Unlimited</td></tr>
</table>

Cisco IOS Guide at Page No. 647-648 (“[Figure 7 shows an example of the backup tunnel selection procedure based on the designated amount of global pool and subpool bandwidth currently available.</td></tr>
</table>

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 13 of 698

No.	'821 Patent Claim 1	Cisco IOS System
		Figure 7 Choosing from Among Multiple Backup Tunnels  In this example, an LSP requires 20 units (kilobits per second) of sub-pool backup bandwidth. The best backup tunnel is selected as follows: 1. Backup tunnels T1 through T4 are considered first because they terminate at the NNHOP. 2. Tunnel T4 is eliminated because it has only ten units of sub-pool backup bandwidth. 3. Tunnel T1 is eliminated because it protects only LSPs using global-pool bandwidth. 4. Tunnel T3 is chosen over T2 because, although both have sufficient backup bandwidth, T3 has the least backup bandwidth available (leaving the most backup bandwidth available on T2). 5. Tunnels T5 and T6 need not be considered because they terminate at an NHOP, and therefore are less desirable than T3, which terminates at an NNHOP.]") In addition, at least under the apparent claim scope alleged by Orckit's Infringement Disclosures, Cisco IOS System in combination with (1) the knowledge of a person of ordinary skill in the art, alone or in further combination with (2) each (individually, as well as one or more together) of the references identified in element 1[b] of Exhibit E-2 renders the claim, including the present limitation, obvious. Below are examples of two such references. For example, Juniper IOS Guide discloses that through fate sharing, you can configure one or more backup paths that minimize the number of shared links with the primary paths. Juniper IOS Guide further discloses that fate sharing ensures that a single point of failure will

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 14 of 698

No.	'821 Patent Claim 1	Cisco IOS System
		not affect the primary and backup paths simultaneously. At least based upon Orckit's apparent understanding of this term as reflected in its Infringement Disclosures, this limitation is met. Juniper IOS Guide at Page No. 73 ("[Through fate sharing, you can configure backup paths that minimize the number of shared links and fiber paths with the primary paths as much as possible, to ensure that if a fiber is cut, the minimum amount of data is lost and a path still exists to the destination. For a backup path to work optimally, it must not share links or physical fiber paths with the primary path, ensuring that a single point of failure will not affect the primary and backup paths simultaneously. For more information about fate sharing, see the Junos OS Routing Protocols Configuration Guide.]") Juniper IOS Guide at Page No. 97 ("[Fate-sharing groups contain three types of objects: • Point-to-point links—Identified by the IP addresses at each end of the link. Unnumbered point-to-point links are typically identified by borrowing IP addresses from other interfaces. Order is not important; from 1.2.3.4 to 1.2.3.5 and from 1.2.3.5 to 1.2.3.4 have the same meaning. • Non-point-to-point links—Include links on a LAN interface (such as Gigabit Ethernet interfaces) or nonbroadcast multiaccess (NBMA) interfaces (such as Asynchronous Transfer Mode [ATM] or Frame Relay). You identify these links by their individual interface address. For example, if the LAN interface 192.168.200.0/24 has four routers attached to it, each router link is individually identified: from 192.168.200.1; # LAN interface of router 1 from 192.168.200.2; # LAN interface of router 2 from 192.168.200.3; # LAN interface of router 3 from 192.168.200.4; # LAN interface of router 4 You can list the addresses in any order. • A router node—Identified by its configured router ID. An object can be in any number of groups, and a group can contain any number of objects. Each group has a configurable cost attributed to it, which represents the level of impact this group has on CSPF computations. The higher the cost, the less likely a backup path will share

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 15 of 698

No.	'821 Patent Claim 1	Cisco IOS System
		with the primary path any objects in the group. The cost is directly comparable to traffic engineering metrics. By default, the cost is 1. Changing the fate-sharing database does not affect established LSPs until the next reoptimization of CSPF. The fate-sharing database does influence fast-reroute computations.]") For example, RFC 4872 discloses that the plurality of MPLS LSPs, i.e., transport entities, are selected in pairs of either working and protection LSPs, or working and restoration LSPs. RFC 4872 further discloses that the plurality of MPLS LSPs, i.e., transport entities, are selected to necessarily be "resource disjoint," i.e., where multiple resources, such as links, nodes, or Shared Risk Link Groups (SRLGs) do not share common links or failure probabilities. Thus, at least under the apparent claim scope alleged by Orckit's Infringement Disclosures, this limitation is met. To the extent that the IETF MPLS-TP System is found to not meet this limitation, if said subset comprises at least two entity pairs, selecting an entity pair from said subset that minimizes an entity cost function would have been obvious to a person having ordinary skill in the art, as explained below. "When a failure is detected by one or both endpoints of the LSP, both endpoints must select traffic from the other LSP. This action must be coordinated between node A and D. From this perspective, 1+1 bidirectional protection can be seen as a coordinated protection- switching mechanism between both endpoints. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 11. "1:1 Protection with Extra-Traffic The most common case of end-to-end 1:N protection is to establish, between the same endpoints, an end-to-end working LSP (thus, N = 1) and a dedicated end-to-end protecting LSP that are mutually link/ node/SRLG disjoint. This protects against working LSP failure(s)." RFC 4872 at 13. "An approach to reduce recovery resource requirements is to have protection LSPs sharing network resources when the working LSPs that they protect are physically (i.e., link, node,

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 16 of 698

No.	'821 Patent Claim 1	Cisco IOS System
		SRLG, etc.) disjoint. This mechanism is referred to as shared mesh restoration and is described in [RFC4426]. Shared-mesh restoration can be seen as a particular case of pre-planned LSP rerouting (see Section 8) that reduces the recovery resource requirements by allowing multiple protecting LSPs to share common link and node resources. Here also, the recovery resources for the protecting LSPs are pre-reserved during the provisioning phase, thus an explicit signaling action is required to activate (i.e., commit resource allocation at the data plane) a specific protecting LSP instantiated during the (pre-) provisioning phase. This requires restoration signaling along the protecting LSP." RFC 4872 at 20. "1. Introduction Generalized Multi-Protocol Label Switching (GMPLS) extends MPLS to include support for Layer-2 Switch Capable (L2SC), Time-Division Multiplex (TDM), Lambda Switch Capable (LSC), and Fiber Switch Capable (FSC) interfaces. GMPLS recovery uses control plane mechanisms (i.e., signaling, routing, and link management mechanisms) to support data plane fault recovery. Note that the analogous (data plane) fault detection mechanisms are required to be present in support of the control plane mechanisms. In this document, the term "recovery" is generically used to denote both protection and restoration; the specific terms "protection" and "restoration" are only used when differentiation is required. The subtle distinction between protection and restoration is made based on the resource allocation done during the recovery phase (see [RFC4427]). A functional description of GMPLS recovery is provided in [RFC4426] and should be considered as a companion document. The present document describes the protocol-specific procedures for GMPLS RSVP- TE (Resource ReSerVation Protocol - Traffic Engineering) signaling (see [RFC3473]) to support end-to-end recovery. End-to-end recovery refers to the recovery of an entire LSP from its head-end (ingress node endpoint) to its tail-end (egress node endpoint). With end-to- end recovery, working LSPs are assumed to be resource-disjoint (where a resource is a link, node, or Shared Risk Link Group (SRLG)) in the network so that they do not share any failure probability, but this is not mandatory. With respect to a given set of network resources, a pair of working/protecting LSPs SHOULD be resource disjoint in case of dedicated recovery type (see below). On the other hand, in case of shared recovery (see below), a group of working LSPs SHOULD be mutually resource-disjoint in order to

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 17 of 698

No.	'821 Patent Claim 1	Cisco IOS System
		allow for a (single and commonly) shared protecting LSP, itself resource-disjoint from each of the working LSPs. Note that resource disjointness is a necessary (but not sufficient) condition to ensure LSP recoverability. The present document addresses four types of end-to-end LSP recovery: 1) 1+1 (unidirectional/bidirectional) protection, 2) 1:N (N >= 1) LSP protection with extra-traffic, 3) pre-planned LSP rerouting without extra-traffic (including shared mesh), and 4) full LSP rerouting. 1) The simplest notion of end-to-end LSP protection is 1+1 unidirectional protection. Using this type of protection, a protecting LSP is signaled over a dedicated resource-disjoint alternate path to protect an associated working LSP. Normal traffic is simultaneously sent on both LSPs and a selector is used at the egress node to receive traffic from one of the LSPs. If a failure occurs along one of the LSPs, the egress node selects the traffic from the valid LSP. No coordination is required between the end nodes when a failure/switchover occurs. In 1+1 bidirectional protection, a protecting LSP is signaled over a dedicated resource-disjoint alternate path to protect the working LSP. Normal traffic is simultaneously sent on both LSPs (in both directions), and a selector is used at both ingress/egress nodes to receive traffic from the same LSP. This requires coordination between the end-nodes when switching to the protecting LSP. 2) In 1:N (N >= 1) protection with extra-traffic, the protecting LSP is a fully provisioned and resource-disjoint LSP from the N working LSPs, that allows for carrying extra-traffic. The N working LSPs MAY be mutually resource-disjoint. Coordination between end-nodes is required when switching from one of the working LSPs to the protecting LSP. As the protecting LSP is fully provisioned, default operations during protection switching are specified for a protecting LSP carrying extra-traffic, but this is not mandatory. Note that M:N protection is out of scope of this document (though mechanisms it defines may be extended to cover it). 3) Pre-planned LSP rerouting (or restoration) relies on the establishment between the same pair of end-nodes of a working LSP and a protecting LSP that is link/node/SRLG disjoint from the working one. Here, the recovery resources for the protecting LSP are pre-reserved but explicit action is required to activate (i.e., commit resource allocation

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 18 of 698

No.	'821 Patent Claim 1	Cisco IOS System
		at the data plane) a specific protecting LSP instantiated during the (pre-)provisioning phase. Since the protecting LSP is not "active" (i.e., fully instantiated), it cannot carry any extra-traffic. This does not mean that the corresponding resources cannot be used by other LSPs. Therefore, this mechanism protects against working LSP(s) failure(s) but requires activation of the protecting LSP after working LSP failure occurrence. This requires restoration signaling along the protecting path. "Shared-mesh" restoration can be seen as a particular case of pre-planned LSP rerouting that reduces the recovery resource requirements by allowing multiple protecting LSPs to share common link and node resources. The recovery resources are pre-reserved but explicit action is required to activate (i.e., commit resource allocation at the data plane) a specific protecting LSP instantiated during the (pre-) provisioning phase. This procedure requires restoration signaling along the protecting path. Note that in both cases, bandwidth pre-reserved for a protecting (but not activated) LSP can be made available for carrying extra traffic. LSPs for extra-traffic (with lower holding priority than the protecting LSP) can then be established using the bandwidth pre-reserved for the protecting LSP. Also, any lower priority LSP that use the pre-reserved resources for the protecting LSP(s) must be preempted during the activation of the protecting LSP. 4) Full LSP rerouting (or restoration) switches normal traffic to an alternate LSP that is not even partially established until after the working LSP failure occurs. The new alternate route is selected at the LSP head-end node, it may reuse resources of the failed LSP at intermediate nodes and may include additional intermediate nodes and/or links." " RFC 4872 at 3–5. "Consider the following network topology:  The paths [A,B,C,D] and [A,E,F,G,D] are node and link disjoint, ignoring the ingress/egress nodes A and D. A 1+1 protected path is established from A to D over [A,B,C,D] and [A,E,F,G,D], and traffic is transmitted simultaneously over both component paths (i.e., LSPs).

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 19 of 698

No.	'821 Patent Claim 1	Cisco IOS System
		During the provisioning phase, both LSPs are fully instantiated (and thus activated) so that no resource sharing can be done along the protecting LSP (nor can any extra-traffic be transported). It is also RECOMMENDED to set the N bit since no protection-switching signaling is assumed in this case. When a failure occurs (say, at node B) and is detected at end-node D, the receiver at D selects the normal traffic from the other LSP. From this perspective, 1+1 unidirectional protection can be seen as an uncoordinated protection-switching mechanism acting independently at both endpoints. Also, for the LSP under failure condition, it is RECOMMENDED to not set the Path_State_Removed Flag of the ERROR_SPEC object (see [RFC3473]) upon PathErr message generation. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 9. "When a failure is detected by one or both endpoints of the LSP, both endpoints must select traffic from the other LSP. This action must be coordinated between node A and D. From this perspective, 1+1 bidirectional protection can be seen as a coordinated protection- switching mechanism between both endpoints. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 11. "(Full) LSP Rerouting LSP rerouting, on the other hand, switches normal traffic to an alternate LSP that is fully established only after failure occurrence. The new (alternate) route is selected at the LSP head- end and may reuse intermediate nodes included in the original route; it may also include additional intermediate nodes. For strict-hop routing, TE requirements can be directly applied to the route computation, and the failed node or link can be avoided. However, if the failure occurred within a loose-routed hop, the head-end node may not have enough information to reroute the LSP around the failure. Crankback signaling (see [CRANK]) and route exclusion techniques (see [RFC4874]) MAY be used in this case.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 20 of 698

No.	'821 Patent Claim 1	Cisco IOS System
		The alternate route MAY be either computed on demand (that is, when the failure occurs; this is referred to as full LSP rerouting) or pre-computed and stored for use when the failure is reported. The latter offers faster restoration time. There is, however, a risk that the alternate route will become out of date through other changes in the network; this can be mitigated to some extent by periodic recalculation of idle alternate routes. (Full) LSP rerouting will be initiated by the head-end node that has either detected the LSP failure or received a Notify message and/or a PathErr message with the new error code/sub-code "Notify Error/LSP Locally Failed" for this LSP. The new LSP resources can be established using the make-before-break mechanism, where the new LSP is set up before the old LSP is torn down. This is done by using the mechanisms of the SESSION_ATTRIBUTE object and the Shared-Explicit (SE) reservation style (see [RFC3209]). Both the new and old LSPs can share resources at common nodes. Note that the make-before-break mechanism is not used to avoid disruption to the normal traffic flow (the latter has already been broken by the failure that is being repaired). However, it is valuable to retain the resources allocated on the original LSP that will be reused by the new alternate LSP." RFC 4872 at 25. "The procedure is as follows: 1) The initiating (source) node sends the normal traffic onto both the working and the protecting LSPs. Once completed, the source node sends reliably a Notify message to the destination with the new error code/sub-code "Notify Error/LSP Recovered" (Switchback Request). This Notify message includes the MESSAGE_ID object. The ACK_Desired flag MUST be set in this object to request the receiver to send an acknowledgment for the message (see [RFC2961]). 2) Upon receipt of this message, the destination selects the traffic from the working LSP. At the same time, it transmits the traffic onto both the working and protecting LSP. The destination then sends reliably a Notify message to the source confirming the completion of the operation. This message includes the MESSAGE_ID_ACK object to acknowledge reception of the received Notify message. This Notify message also includes the MESSAGE_ID object. The ACK_Desired flag MUST be set in this object to request the receiver to send an acknowledgment for the message (see [RFC2961]).

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 21 of 698

No.	'821 Patent Claim 1	Cisco IOS System
		3) When the source node receives this Notify message, it switches to receive traffic from the working LSP. The source node then sends an Ack message to the destination node confirming that the LSP has been reverted." RFC 4872 at 27.
1[c]	selecting an entity pair from said plurality of transport entities based at least in part upon said overall cost; and	Cisco IOS System discloses selecting an entity pair from said plurality of transport entities based at least in part upon said overall cost. For example, Cisco IOS Guide discloses that the path that an LSP uses is determined by the LSP resource requirements and network resources, such as bandwidth. Cisco IOS Guide further discloses that the MPLS TE SRLG feature enhances backup tunnel path selection so a backup tunnel can avoid using links that are in the same SRLG as the interfaces it is protecting. At least based upon Orckit's apparent understanding of this term as reflected in its Infringement Disclosures, this limitation is met. Cisco IOS Guide at Page No. 253 ("[Path Option Selection for MPLS TE Tunnel LSPs This section contains the following topics about path option selection for MPLS TE Tunnel LSPs: • Constraint-Based Routing and Path Option Selection, page 5 • Tunnel Reoptimization and Path Option Selection, page 5 • Path Option Selection with Bandwidth Override, page 6 Constraint-Based Routing and Path Option Selection MPLS traffic engineering automatically establishes and maintains LSPs across the backbone by using the Resource Reservation Protocol (RSVP). The path that an LSP uses is determined by the LSP resource requirements and network resources, such as bandwidth. Traffic engineering tunnels are calculated at the LSP head based on a fit between required and available resources (constraint-based routing). Without the Path Option for Bandwidth Override feature, a TE tunnel establishes an LSP based on dynamic or explicit path options in order of preference. However, the bandwidth and other attributes configured on the TE tunnel allow the setup of an LSP only if LSP path options satisfy the constraints. If a path cannot be found that satisfies the configured path options, then the tunnel is not set up.]")

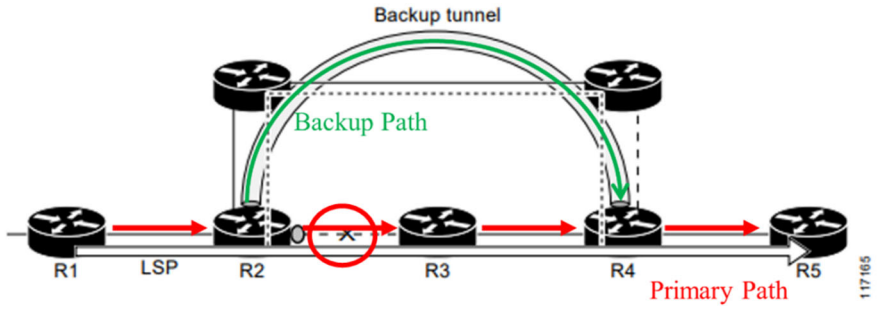
Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 22 of 698

No.	'821 Patent Claim 1	Cisco IOS System
		Cisco IOS Guide at Page No. 639 ("[Bandwidth Protection NHOP and NNHOP backup tunnels can be used to provide bandwidth protection for rerouted LSPs. This is referred to as backup bandwidth. You can associate backup bandwidth with NHOP or NNHOP backup tunnels. This informs the router of the amount of backup bandwidth a particular backup tunnel can protect. When a router maps LSPs to backup tunnels, bandwidth protection ensures that an LSP uses a given backup tunnel only if there is sufficient backup bandwidth. The router selects which LSPs use which backup tunnels in order to provide maximum bandwidth protection. That is, the router determines the best way to map LSPs onto backup tunnels in order to maximize the number of LSPs that can be protected. For information about mapping tunnels and assigning backup bandwidth, see the "Backup Tunnel Selection Procedure" section on page 10. LSPs that have the "bandwidth protection desired" bit set have a higher right to select backup tunnels that provide bandwidth protection; that is, those LSPs can preempt other LSPs that do not have that bit set. For more information, see the "Prioritizing Which LSPs Obtain Backup Tunnels with Bandwidth Protection" section on page 8.]") Cisco IOS Guide at Page No. 640 ("[Backup Tunnel Support Backup tunnel support has the following capabilities: • Backup Tunnels Can Terminate at the Next-Next Hop to Support FRR, page 6 • Multiple Backup Tunnels Can Protect the Same Interface, page 6 • Backup Tunnels Provide Scalability, page 7.]") Cisco IOS Guide at Page 640 ("[Multiple Backup Tunnels Can Protect the Same Interface There is no limit (except memory limitations) to the number of backup tunnels that can protect a given interface. In many topologies, support for node protection requires supporting multiple backup tunnels per protected interface. These backup tunnels can terminate at the same destination or at different destinations. That is, for a given protected interface, you can configure multiple NHOP or NNHOP backup tunnels. This allows redundancy and load balancing.]") Cisco IOS Guide at Page No. 603 ("[Traffic Engineering Tunnels

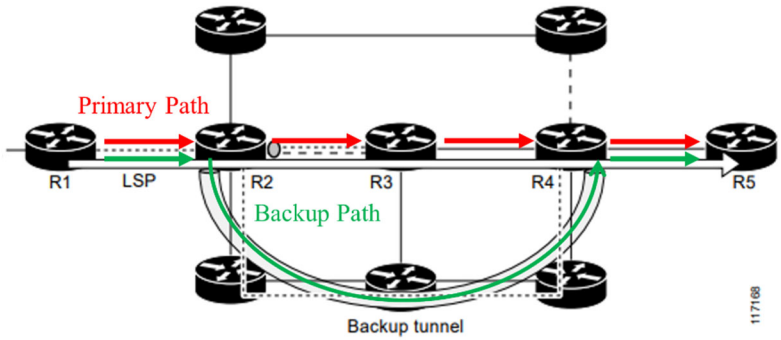
Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 23 of 698

No.	'821 Patent Claim 1	Cisco IOS System
		MPLS TE lets you build label switched paths (LSPs) across your network for forwarding traffic. MPLS TE LSPs let the headend of a TE tunnel control the path its traffic takes to a particular destination. This method is more flexible than forwarding traffic based only on a destination address. Interarea tunnels allow you to do the following: • Build TE tunnels between areas (interarea tunnels) • Build TE tunnels that start and end in the same area, on multiple areas on a router (intra-area tunnels) Some tunnels are more important than others. For example, you may have tunnels carrying VoIP traffic and tunnels carrying data traffic that are competing for the same resources. MPLS TE allows you to have some tunnels preempt others. Each tunnel has a priority, and more-important tunnels take precedence over less-important tunnels.]") Cisco IOS Guide at Page No. 641 ("[Backup Bandwidth Protection Backup bandwidth protection allows you to give LSPs carrying certain kinds of data (such as voice) priority for using backup tunnels.]") Cisco IOS Guide at Page No. 650 ("[For each protected link AB with a maximum reservable subpool value of n, there may be a path from node A to node B such that the difference between the maximum reservable global and the maximum reservable subpool is at least the value of n. If it is possible to find such paths for each link in the network, you can establish all the backup tunnels along such paths without any bandwidth reservations. If there is a single link failure, only one backup tunnel will use any link on its path. Because that path has at least n available bandwidth (in the global pool), assuming that marking and scheduling is configured to classify the subpool traffic into a priority queue, the subpool bandwidth is guaranteed. This approach allows sharing of the global pool bandwidth between backup tunnels protecting independent link failures. The backup tunnels are expected to be used for only a short period of time after a failure (until the headends of affected LSPs reroute those LSPs to other paths with available subpool bandwidth). The probability of multiple unrelated link failures is very small (in the absence of node or shared risk link group (SRLG) failures, which result in

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 24 of 698

No.	'821 Patent Claim 1	Cisco IOS System
		multiple link failures). Therefore, it is reasonable to assume that link failures are in practice independent with high probability. This "independent failure assumption" in combination with backup tunnels signaled without explicit bandwidth reservation enables efficient bandwidth sharing that yields substantial bandwidth savings.]") Cisco IOS Guide at Page No. 651 ("[A similar approach can be used for node and SRLG protection. However, the decision of where to put the backup tunnels is more complicated because both node and SRLG failures effectively result in the simultaneous failure of several links. Therefore, the backup tunnels protecting traffic traversing all affected links cannot be computed independently of each other.]") Cisco IOS Guide at Page No. 579 ("[MPLS Traffic Engineering Shared Risk Link Groups SRLGs refer to situations where links in a network share a common fiber (or a common physical attribute). If one link fails, other links in the group may fail too. Links in the group have a shared risk. Backup tunnels should avoid using links in the same SRLG as interfaces they are protecting. Otherwise, when the protected link fails the backup tunnel fails too. *Figure 1 Backup Tunnel in the Same SRLG as the Interface It Is Protecting*  Fig. 1 (annotation added) Figure 1 shows a primary label-switched path (LSP) from router R1 to router R5. The LSP protects against the failure of the R2-R3 link at R2 via a backup tunnel to R4. If the R2-R3 link fails, link protection reroutes the LSP along the backup tunnel. However, the R2-R3 link

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 25 of 698

No.	'821 Patent Claim 1	Cisco IOS System
		and one of the backup tunnel links are in the same SRLG. So if the R2-R3 link fails, the backup tunnel may fail too. The MPLS TE SRLG feature enhances backup tunnel path selection so a backup tunnel can avoid using links that are in the same SRLG as the interfaces it is protecting. There are two ways for a backup tunnel to avoid the SRLGs of its protected interface: • The router does not create the backup tunnel unless it avoids SRLGs of the protected interface.]") Cisco IOS Guide at Page No. 580 ("[Open Shortest Path First (OSPF) and Intermediate System-to-Intermediate System (IS-IS) flood the SRLG membership information (including other TE link attributes such as bandwidth availability and affinity) so that all routers in the network have the SRLG information for each link. With this topology information, routers can compute backup tunnel paths that exclude links having SRLGs in common with their protected interfaces. As shown in Figure 2, the backup tunnel avoids the link between R2 and R3, which shares an SRLG with the protected interface.]") *Figure 2 Backup Tunnel That Avoids SRLG of Protected Interface*  Fig. 2 (annotation added) Cisco IOS Guide at Page No. 587 ("[Step 4 show mpls traffic-eng topology Use this command to show the SRLG link membership flooded via the Interior Gateway Protocol (IGP).]") Cisco IOS Guide at Page 632

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 26 of 698

No.	'821 Patent Claim 1	Cisco IOS System
		("[IS-IS—Intermediate System-to-Intermediate System. Link-state hierarchical routing protocol that calls for intermediate system (IS) routers to exchange routing information based on a single metric to determine network topology.]") Cisco IOS Guide at Page No. 633 ("[OSPF—Open Shortest Path First. A link-state hierarchical Interior Gateway Protocol routing algorithm, derived from the IS-IS protocol. OSPF features include least-cost routing, multipath routing, and load balancing.]") Cisco IOS Guide at Page No. 581 ("[Autotunnel Backup for MPLS TE SRLGs Autotunnel backup is the ability of routers to create backup tunnels automatically. Therefore, you do not need to preconfigure each backup tunnel and then assign the backup tunnel to the protected interface. Only automatically created backup tunnels can avoid SRLGs or their protected interfaces.]") Cisco IOS Guide at Page No. 583 ("[Configuring the Routers That Automatically Create Backup Tunnels to Avoid MPLS TE SRLGs of Their Protected Interfaces Perform the following task to configure routers that automatically create backup tunnels to avoid MPLS TE SRLGs of their protected interfaces. Backup tunnels provide link protection by rerouting traffic to the next hop bypassing failed links or in this instance by avoiding SRLGs.]") Cisco IOS Guide at Page No. 644 ("[Backup Tunnel Selection Procedure When an LSP is signaled, each node along the LSP path that provides FRR protection for the LSP selects a backup tunnel for the LSP to use if either of the following events occurs: • The link to the next hop fails. • The next hop fails. By having the node select the backup tunnel for an LSP before a failure occurs, the LSP can be rerouted onto the backup tunnel quickly if there is a failure.]") Cisco IOS Guide at Page 646 ("[NHOP Versus NNHOP Backup Tunnels More than one backup tunnel can protect a given LSP, where one backup tunnel terminates at the LSP's NNHOP, and the other terminates at the LSP's NHOP. In this case, the router

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 27 of 698

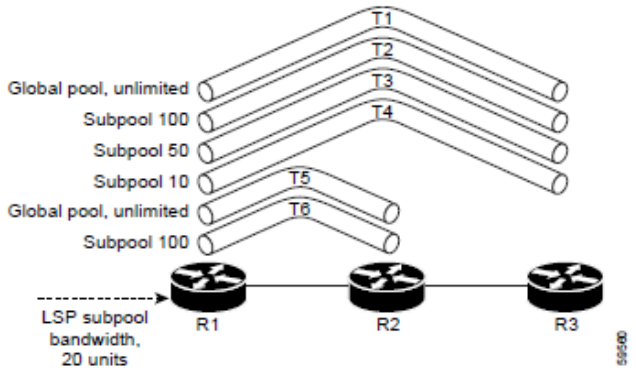
<table>
<tr><th>No.</th><th>’821 Patent Claim 1</th><th>Cisco IOS System</th></tr>
<tr><td></td><td></td><td>chooses the backup tunnel that terminates at the NNHOP (that is, FRR prefers NNHOP over NHOP backup tunnels).Table 1 lists the tunnel selection priorities. The first choice is an NNHOP backup tunnel that acquires its bandwidth from a subpool or global pool, and has limited bandwidth. If there is no such backup tunnel, the next choice (2) is a next-next hop backup tunnel that acquires a limited amount of bandwidth from any pool. The preferences go from 1 (best) to 8 (worst), where choice 3 is for an NNHOP backup tunnel with an unlimited amount of subpool or global-pool bandwidth.]”)

Table 1 Tunnel Selection Priorities

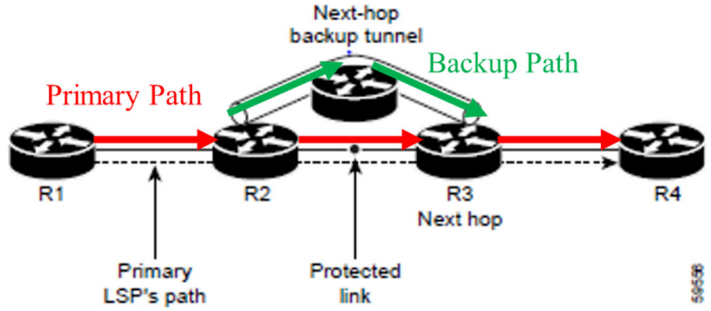
<table>
<tr><th>Preference</th><th>Backup Tunnel Destination</th><th>Bandwidth Pool</th><th>Bandwidth Amount</th></tr>
<tr><td>1 (Best)</td><td>NNHOP</td><td>Subpool or global pool</td><td>Limited</td></tr>
<tr><td>2</td><td>NNHOP</td><td>Any</td><td>Limited</td></tr>
<tr><td>3</td><td>NNHOP</td><td>Subpool or global pool</td><td>Unlimited</td></tr>
<tr><td>4</td><td>NNHOP</td><td>Any</td><td>Unlimited</td></tr>
<tr><td>5</td><td>NHOP</td><td>Subpool or global pool</td><td>Limited</td></tr>
<tr><td>6</td><td>NHOP</td><td>Any</td><td>Limited</td></tr>
<tr><td>7</td><td>NHOP</td><td>Subpool or global pool</td><td>Unlimited</td></tr>
<tr><td>8 (Worst)</td><td>NHOP</td><td>Any</td><td>Unlimited</td></tr>
</table>

Cisco IOS Guide at Page No. 647-648 (“[Figure 7 shows an example of the backup tunnel selection procedure based on the designated amount of global pool and subpool bandwidth currently available.</td></tr>
</table>

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 28 of 698

No.	'821 Patent Claim 1	Cisco IOS System
		Figure 7 Choosing from Among Multiple Backup Tunnels  In this example, an LSP requires 20 units (kilobits per second) of sub-pool backup bandwidth. The best backup tunnel is selected as follows: 1. Backup tunnels T1 through T4 are considered first because they terminate at the NNHOP. 2. Tunnel T4 is eliminated because it has only ten units of sub-pool backup bandwidth. 3. Tunnel T1 is eliminated because it protects only LSPs using global-pool bandwidth. 4. Tunnel T3 is chosen over T2 because, although both have sufficient backup bandwidth, T3 has the least backup bandwidth available (leaving the most backup bandwidth available on T2). 5. Tunnels T5 and T6 need not be considered because they terminate at an NHOP, and therefore are less desirable than T3, which terminates at an NNHOP.]")
1[d]	if an entity pair reselection event occurs, reselecting said entity pair from the group consisting of said entity pair and a replacement entity pair comprising at least one entity distinct from the	Cisco IOS System discloses if an entity pair reselection event occurs, reselecting said entity pair from the group consisting of said entity pair and a replacement entity pair comprising at least one entity distinct from the entities of said entity pair. For example, Cisco IOS Guide discloses that if an LSP is using a chosen backup tunnel and there is a network change such as a backup tunnel goes down or network bandwidth changes the backup tunnel choice should be reevaluated. Cisco IOS Guide further discloses that this reevaluation is known as promotion and a new LSP and/or backup LSP may be selected from the promotion process.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 29 of 698

No.	'821 Patent Claim 1	Cisco IOS System
	entities of said entity pair,	Cisco IOS Guide at Page No. 637 ("[Fast Reroute Fast Reroute (FRR) is a mechanism for protecting MPLS TE LSPs from link and node failures by locally repairing the LSPs at the point of failure, allowing data to continue to flow on them while their headend routers attempt to establish new end-to-end LSPs to replace them. FRR locally repairs the protected LSPs by rerouting them over backup tunnels that bypass failed links or node. Link Protection Backup tunnels that bypass only a single link of the LSP's path provide link protection. They protect LSPs if a link along their path fails by rerouting the LSP's traffic to the next hop (bypassing the failed link). These are referred to as next-hop (NHOP) backup tunnels because they terminate at the LSP's next hop beyond the point of failure. Figure 1 illustrates an NHOP backup tunnel.]") *Figure 1 NHOP Backup Tunnel*  Fig. 1 (annotation added) Cisco IOS Guide at Page No. 638 ("[Node Protection FRR provides node protection for LSPs. Backup tunnels that bypass next-hop nodes along LSP paths are called next-next-hop (NNHOP) backup tunnels because they terminate at the node following the next-hop node of the LSP paths, thereby bypassing the next-hop node. They protect LSPs if a node along their path fails by enabling the node upstream of the failure to reroute the LSPs and their traffic around the failed node to the next-next hop. FRR supports

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 30 of 698

No.	'821 Patent Claim 1	Cisco IOS System
		the use of RSVP Hellos to accelerate the detection of node failures. NNHOP backup tunnels also provide protection from link failures, because they bypass the failed link and the node.]")

Figure 2 illustrates an NNHOP backup tunnel.

Figure 2 NNHOP Backup Tunnel

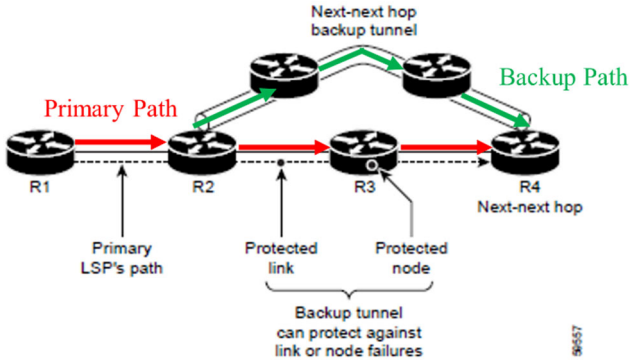


Fig. 2 (annotation added)

Cisco IOS Guide at Page No. 638 ("[If an LSP is using a backup tunnel and something changes so that the LSP is no longer appropriate for the backup tunnel, the LSP is torn down. Such changes are the following:
• Backup bandwidth of the backup tunnel is reduced.
• Backup bandwidth type of backup tunnel is changed to a type that is incompatible with the primary LSP.
• Primary LSP is modified so that FRR is disabled. (The no mpls traffic-eng fast-reroute command is entered.]")

Cisco IOS Guide at Page No. 648 ("[After a backup tunnel has been chosen for an LSP, conditions may change that will cause us to reevaluate this choice. This reevaluation, if successful, is called promotion. Such conditions may include:
1. A new backup tunnel comes up.
2. The currently chosen backup tunnel for this LSP goes down.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 31 of 698

No.	'821 Patent Claim 1	Cisco IOS System
		3. A backup tunnel's available backup bandwidth increases. For example, an LSP protected by the tunnel has been reoptimized by the headend to use another path. For cases 1 and 2, the LSP's backup tunnel is evaluated immediately. Case 3 is addressed by periodically reevaluating LSP-to-backup tunnel mappings. By default, background reevaluation is performed every 5 minutes. This interval is configurable via the mpls traffic-eng fast-reroute timers command.]") Cisco IOS Guide at Page No. 554 ("[Tunnel Reoptimization Procedure Tunnel reoptimization is the signaling of an LSP that is more optimal than the LSP a TE tunnel is currently using (for example, it may be shorter or may have a lower cost), and the switching over of the tunnel's data to use this new LSP. The new more optimal TE LSP is always established and the data moved onto it before the original LSP is torn down (so it is called the "make before break" procedure). This ensures that no data packets are lost during the transition to the new LSP.]") Cisco IOS Guide at Page No. 633 ("[Primary Tunnel Does Not Select Backup Tunnel That Is Up If a backup tunnel is up, but it is not selected as a backup tunnel by the primary tunnel (LSP), enter the following commands for the backup tunnel: • shutdown • no shutdown Note If you change the status of a backup tunnel, the backup tunnel selection algorithm is rerun for the backup tunnel. LSPs that have currently selected (that is, are ready to use) that backup tunnel will be disassociated from it, and then reassociated with that backup tunnel or another backup tunnel. This is generally harmless and usually results in mapping the same LSPs to that backup tunnel. However, if any LSPs are actively using that backup tunnel, shutting down the backup tunnel will tear down those LSPs.]")
1[e]	wherein said entity pair reselection event is selected from a group consisting of adding an entity to said plurality	Cisco IOS System discloses wherein said entity pair reselection event is selected from a group consisting of adding an entity to said plurality of transport entities, removing an entity from said plurality of transport entities, an operational status change for one of said plurality of transport entities, and a change in overall cost for one of said plurality of transport entities.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 32 of 698

No.	'821 Patent Claim 1	Cisco IOS System
	of transport entities, removing an entity from said plurality of transport entities, an operational status change for one of said plurality of transport entities, and a change in overall cost for one of said plurality of transport entities.	For example, Cisco IOS Guide discloses certain network events that can lead to LSP promotions. Cisco IOS Guide further discloses that these network events can include a new backup tunnel comes up, the currently chosen backup tunnel for this LSP goes down, the backup tunnel changing its status based upon a change in network bandwidth. Thus, at least under Orckit's apparent understanding of this claim limitation in Orckit's Infringement Disclosures, this limitation is met. To the extent that the Cisco IOS System is found to not meet this limitation, wherein said entity pair reselection event is selected from a group consisting of adding an entity to said plurality of transport entities, removing an entity from said plurality of transport entities, an operational status change for one of said plurality of transport entities, and a change in overall cost for one of said plurality of transport entities would have been obvious to a person having ordinary skill in the art, as explained below. Cisco IOS Guide at Page No. 648 ("[After a backup tunnel has been chosen for an LSP, conditions may change that will cause us to reevaluate this choice. This reevaluation, if successful, is called promotion. Such conditions may include: 1. A new backup tunnel comes up. 2. The currently chosen backup tunnel for this LSP goes down. 3. A backup tunnel's available backup bandwidth increases. For example, an LSP protected by the tunnel has been reoptimized by the headend to use another path. For cases 1 and 2, the LSP's backup tunnel is evaluated immediately. Case 3 is addressed by periodically reevaluating LSP-to-backup tunnel mappings. By default, background reevaluation is performed every 5 minutes. This interval is configurable via the mpls traffic-eng fast-reroute timers command.]") Cisco IOS Guide at Page No. 633 ("[Primary Tunnel Does Not Select Backup Tunnel That Is Up If a backup tunnel is up, but it is not selected as a backup tunnel by the primary tunnel (LSP), enter the following commands for the backup tunnel: • shutdown • no shutdown Note If you change the status of a backup tunnel, the backup tunnel selection algorithm is rerun for the backup tunnel. LSPs that have currently selected (that is, are ready to use) that

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 33 of 698

No.	'821 Patent Claim 1	Cisco IOS System
		backup tunnel will be disassociated from it, and then reassociated with that backup tunnel or another backup tunnel. This is generally harmless and usually results in mapping the same LSPs to that backup tunnel. However, if any LSPs are actively using that backup tunnel, shutting down the backup tunnel will tear down those LSPs.]") Cisco IOS Guide at Page No. 544 ("[Tunnel Reoptimization Procedure Tunnel reoptimization is the signaling of an LSP that is more optimal than the LSP a TE tunnel is currently using (for example, it may be shorter or may have a lower cost), and the switching over of the tunnel's data to use this new LSP. The new more optimal TE LSP is always established and the data moved onto it before the original LSP is torn down (so it is called the "make before break" procedure). This ensures that no data packets are lost during the transition to the new LSP.]") In addition, at least under the apparent claim scope alleged by Orckit's Infringement Disclosures, Cisco IOS System in combination with (1) the knowledge of a person of ordinary skill in the art, alone or in further combination with (2) each (individually, as well as one or more together) of the references identified in element 1[e] of Exhibit E-2 renders the claim, including the present limitation, obvious. Below are examples of two such references. For example, Juniper IOS Guide discloses that through fate sharing, you can configure one or more backup paths that minimize the number of shared links with the primary paths. Juniper IOS Guide further discloses that fate sharing ensures that a single point of failure will not affect the primary and backup paths simultaneously. At least based upon Orckit's apparent understanding of this term as reflected in its Infringement Disclosures, this limitation is met. Juniper IOS Guide at Page No. 73 ("[Through fate sharing, you can configure backup paths that minimize the number of shared links and fiber paths with the primary paths as much as possible, to ensure that if a fiber is cut, the minimum amount of data is lost and a path still exists to the destination. For a backup path to work optimally, it must not share links or physical fiber paths with the primary path, ensuring that a single point of failure will not affect the primary and backup

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 34 of 698

No.	'821 Patent Claim 1	Cisco IOS System
		paths simultaneously. For more information about fate sharing, see the Junos OS Routing Protocols Configuration Guide.]") Juniper IOS Guide at Page No. 97 ("[Fate-sharing groups contain three types of objects: • Point-to-point links—Identified by the IP addresses at each end of the link. Unnumbered point-to-point links are typically identified by borrowing IP addresses from other interfaces. Order is not important; from 1.2.3.4 to 1.2.3.5 and from 1.2.3.5 to 1.2.3.4 have the same meaning. • Non-point-to-point links—Include links on a LAN interface (such as Gigabit Ethernet interfaces) or nonbroadcast multiaccess (NBMA) interfaces (such as Asynchronous Transfer Mode [ATM] or Frame Relay). You identify these links by their individual interface address. For example, if the LAN interface 192.168.200.0/24 has four routers attached to it, each router link is individually identified: from 192.168.200.1; # LAN interface of router 1 from 192.168.200.2; # LAN interface of router 2 from 192.168.200.3; # LAN interface of router 3 from 192.168.200.4; # LAN interface of router 4 You can list the addresses in any order. • A router node—Identified by its configured router ID. An object can be in any number of groups, and a group can contain any number of objects. Each group has a configurable cost attributed to it, which represents the level of impact this group has on CSPF computations. The higher the cost, the less likely a backup path will share with the primary path any objects in the group. The cost is directly comparable to traffic engineering metrics. By default, the cost is 1. Changing the fate-sharing database does not affect established LSPs until the next reoptimization of CSPF. The fate-sharing database does influence fast-reroute computations.]") For example, RFC 4872 discloses that the plurality of MPLS LSPs, i.e., transport entities, are selected in pairs of either working and protection LSPs, or working and restoration LSPs. RFC 4872 further discloses that the plurality of MPLS LSPs, i.e., transport entities, are selected to necessarily be "resource disjoint," i.e., where multiple resources, such as links, nodes, or Shared Risk Link Groups (SRLGs) do not share common links or failure

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 35 of 698

No.	'821 Patent Claim 1	Cisco IOS System
		probabilities. Thus, at least under the apparent claim scope alleged by Orckit's Infringement Disclosures, this limitation is met. To the extent that the IETF MPLS-TP System is found to not meet this limitation, if said subset comprises at least two entity pairs, selecting an entity pair from said subset that minimizes an entity cost function would have been obvious to a person having ordinary skill in the art, as explained below. "When a failure is detected by one or both endpoints of the LSP, both endpoints must select traffic from the other LSP. This action must be coordinated between node A and D. From this perspective, 1+1 bidirectional protection can be seen as a coordinated protection- switching mechanism between both endpoints. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 11. "1:1 Protection with Extra-Traffic The most common case of end-to-end 1:N protection is to establish, between the same endpoints, an end-to-end working LSP (thus, N = 1) and a dedicated end-to-end protecting LSP that are mutually link/ node/SRLG disjoint. This protects against working LSP failure(s)." RFC 4872 at 13. "An approach to reduce recovery resource requirements is to have protection LSPs sharing network resources when the working LSPs that they protect are physically (i.e., link, node, SRLG, etc.) disjoint. This mechanism is referred to as shared mesh restoration and is described in [RFC4426]. Shared-mesh restoration can be seen as a particular case of pre-planned LSP rerouting (see Section 8) that reduces the recovery resource requirements by allowing multiple protecting LSPs to share common link and node resources. Here also, the recovery resources for the protecting LSPs are pre-reserved during the provisioning phase, thus an explicit signaling action is required to activate (i.e., commit resource allocation at the data plane) a specific protecting LSP instantiated during the (pre-) provisioning phase. This requires restoration signaling along the protecting LSP." RFC 4872 at 20. "1. Introduction

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 36 of 698

No.	'821 Patent Claim 1	Cisco IOS System
		Generalized Multi-Protocol Label Switching (GMPLS) extends MPLS to include support for Layer-2 Switch Capable (L2SC), Time-Division Multiplex (TDM), Lambda Switch Capable (LSC), and Fiber Switch Capable (FSC) interfaces. GMPLS recovery uses control plane mechanisms (i.e., signaling, routing, and link management mechanisms) to support data plane fault recovery. Note that the analogous (data plane) fault detection mechanisms are required to be present in support of the control plane mechanisms. In this document, the term "recovery" is generically used to denote both protection and restoration; the specific terms "protection" and "restoration" are only used when differentiation is required. The subtle distinction between protection and restoration is made based on the resource allocation done during the recovery phase (see [RFC4427]). A functional description of GMPLS recovery is provided in [RFC4426] and should be considered as a companion document. The present document describes the protocol-specific procedures for GMPLS RSVP- TE (Resource ReSerVation Protocol - Traffic Engineering) signaling (see [RFC3473]) to support end-to-end recovery. End-to-end recovery refers to the recovery of an entire LSP from its head-end (ingress node endpoint) to its tail-end (egress node endpoint). With end-to- end recovery, working LSPs are assumed to be resource-disjoint (where a resource is a link, node, or Shared Risk Link Group (SRLG)) in the network so that they do not share any failure probability, but this is not mandatory. With respect to a given set of network resources, a pair of working/protecting LSPs SHOULD be resource disjoint in case of dedicated recovery type (see below). On the other hand, in case of shared recovery (see below), a group of working LSPs SHOULD be mutually resource-disjoint in order to allow for a (single and commonly) shared protecting LSP, itself resource-disjoint from each of the working LSPs. Note that resource disjointness is a necessary (but not sufficient) condition to ensure LSP recoverability. The present document addresses four types of end-to-end LSP recovery: 1) 1+1 (unidirectional/bidirectional) protection, 2) 1:N (N >= 1) LSP protection with extra-traffic, 3) pre-planned LSP rerouting without extra-traffic (including shared mesh), and 4) full LSP rerouting. 5) The simplest notion of end-to-end LSP protection is 1+1 unidirectional protection. Using this type of protection, a protecting LSP is signaled over a dedicated resource-

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 37 of 698

No.	'821 Patent Claim 1	Cisco IOS System
		disjoint alternate path to protect an associated working LSP. Normal traffic is simultaneously sent on both LSPs and a selector is used at the egress node to receive traffic from one of the LSPs. If a failure occurs along one of the LSPs, the egress node selects the traffic from the valid LSP. No coordination is required between the end nodes when a failure/switchover occurs. In 1+1 bidirectional protection, a protecting LSP is signaled over a dedicated resource-disjoint alternate path to protect the working LSP. Normal traffic is simultaneously sent on both LSPs (in both directions), and a selector is used at both ingress/egress nodes to receive traffic from the same LSP. This requires coordination between the end-nodes when switching to the protecting LSP. 6) In 1:N (N >= 1) protection with extra-traffic, the protecting LSP is a fully provisioned and resource-disjoint LSP from the N working LSPs, that allows for carrying extra-traffic. The N working LSPs MAY be mutually resource-disjoint. Coordination between end-nodes is required when switching from one of the working LSPs to the protecting LSP. As the protecting LSP is fully provisioned, default operations during protection switching are specified for a protecting LSP carrying extra-traffic, but this is not mandatory. Note that M:N protection is out of scope of this document (though mechanisms it defines may be extended to cover it). 7) Pre-planned LSP rerouting (or restoration) relies on the establishment between the same pair of end-nodes of a working LSP and a protecting LSP that is link/node/SRLG disjoint from the working one. Here, the recovery resources for the protecting LSP are pre-reserved but explicit action is required to activate (i.e., commit resource allocation at the data plane) a specific protecting LSP instantiated during the (pre-)provisioning phase. Since the protecting LSP is not "active" (i.e., fully instantiated), it cannot carry any extra-traffic. This does not mean that the corresponding resources cannot be used by other LSPs. Therefore, this mechanism protects against working LSP(s) failure(s) but requires activation of the protecting LSP after working LSP failure occurrence. This requires restoration signaling along the protecting path. "Shared-mesh" restoration can be seen as a particular case of pre-planned LSP rerouting that reduces the recovery resource requirements by allowing multiple protecting LSPs to share common link and node resources. The recovery resources are pre-reserved but explicit action is required to activate (i.e., commit resource allocation at the data plane) a

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 38 of 698

No.	'821 Patent Claim 1	Cisco IOS System
		specific protecting LSP instantiated during the (pre-) provisioning phase. This procedure requires restoration signaling along the protecting path. Note that in both cases, bandwidth pre-reserved for a protecting (but not activated) LSP can be made available for carrying extra traffic. LSPs for extra-traffic (with lower holding priority than the protecting LSP) can then be established using the bandwidth pre-reserved for the protecting LSP. Also, any lower priority LSP that use the pre-reserved resources for the protecting LSP(s) must be preempted during the activation of the protecting LSP. 8) Full LSP rerouting (or restoration) switches normal traffic to an alternate LSP that is not even partially established until after the working LSP failure occurs. The new alternate route is selected at the LSP head-end node, it may reuse resources of the failed LSP at intermediate nodes and may include additional intermediate nodes and/or links." " RFC 4872 at 3–5. "Consider the following network topology:  The paths [A,B,C,D] and [A,E,F,G,D] are node and link disjoint, ignoring the ingress/egress nodes A and D. A 1+1 protected path is established from A to D over [A,B,C,D] and [A,E,F,G,D], and traffic is transmitted simultaneously over both component paths (i.e., LSPs). During the provisioning phase, both LSPs are fully instantiated (and thus activated) so that no resource sharing can be done along the protecting LSP (nor can any extra-traffic be transported). It is also RECOMMENDED to set the N bit since no protection-switching signaling is assumed in this case. When a failure occurs (say, at node B) and is detected at end-node D, the receiver at D selects the normal traffic from the other LSP. From this perspective, 1+1 unidirectional protection can be seen as an uncoordinated protection-switching mechanism acting independently at both endpoints. Also, for the LSP under failure condition, it is RECOMMENDED to not set the Path_State_Removed Flag of the ERROR_SPEC object (see [RFC3473]) upon PathErr message generation.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 39 of 698

No.	'821 Patent Claim 1	Cisco IOS System
		Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 9. "When a failure is detected by one or both endpoints of the LSP, both endpoints must select traffic from the other LSP. This action must be coordinated between node A and D. From this perspective, 1+1 bidirectional protection can be seen as a coordinated protection- switching mechanism between both endpoints. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 11. "(Full) LSP Rerouting LSP rerouting, on the other hand, switches normal traffic to an alternate LSP that is fully established only after failure occurrence. The new (alternate) route is selected at the LSP head- end and may reuse intermediate nodes included in the original route; it may also include additional intermediate nodes. For strict-hop routing, TE requirements can be directly applied to the route computation, and the failed node or link can be avoided. However, if the failure occurred within a loose-routed hop, the head-end node may not have enough information to reroute the LSP around the failure. Crankback signaling (see [CRANK]) and route exclusion techniques (see [RFC4874]) MAY be used in this case. The alternate route MAY be either computed on demand (that is, when the failure occurs; this is referred to as full LSP rerouting) or pre-computed and stored for use when the failure is reported. The latter offers faster restoration time. There is, however, a risk that the alternate route will become out of date through other changes in the network; this can be mitigated to some extent by periodic recalculation of idle alternate routes. (Full) LSP rerouting will be initiated by the head-end node that has either detected the LSP failure or received a Notify message and/or a PathErr message with the new error code/sub-code "Notify Error/LSP Locally Failed" for this LSP. The new LSP resources can be established using the make-before-break mechanism, where the new LSP is set up before the old LSP is torn down. This is done by using the mechanisms of the SESSION_ATTRIBUTE

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 40 of 698

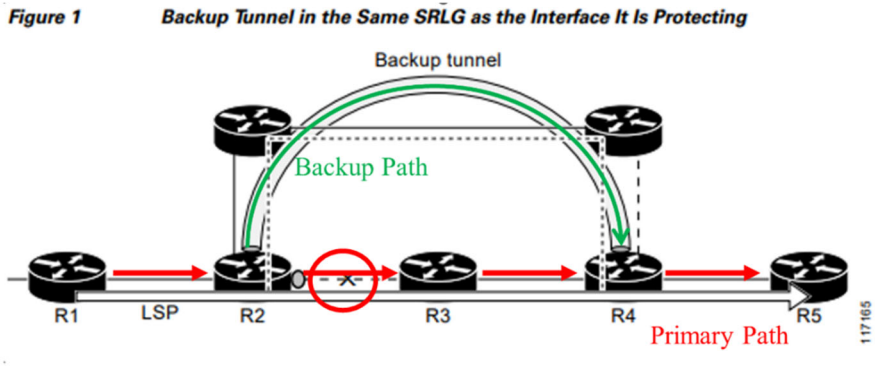
No.	'821 Patent Claim 1	Cisco IOS System
		object and the Shared-Explicit (SE) reservation style (see [RFC3209]). Both the new and old LSPs can share resources at common nodes. Note that the make-before-break mechanism is not used to avoid disruption to the normal traffic flow (the latter has already been broken by the failure that is being repaired). However, it is valuable to retain the resources allocated on the original LSP that will be reused by the new alternate LSP." RFC 4872 at 25. "The procedure is as follows: 4) The initiating (source) node sends the normal traffic onto both the working and the protecting LSPs. Once completed, the source node sends reliably a Notify message to the destination with the new error code/sub-code "Notify Error/LSP Recovered" (Switchback Request). This Notify message includes the MESSAGE_ID object. The ACK_Desired flag MUST be set in this object to request the receiver to send an acknowledgment for the message (see [RFC2961]). 5) Upon receipt of this message, the destination selects the traffic from the working LSP. At the same time, it transmits the traffic onto both the working and protecting LSP. The destination then sends reliably a Notify message to the source confirming the completion of the operation. This message includes the MESSAGE_ID_ACK object to acknowledge reception of the received Notify message. This Notify message also includes the MESSAGE_ID object. The ACK_Desired flag MUST be set in this object to request the receiver to send an acknowledgment for the message (see [RFC2961]). 6) When the source node receives this Notify message, it switches to receive traffic from the working LSP. The source node then sends an Ack message to the destination node confirming that the LSP has been reverted." RFC 4872 at 27.

No.	'821 Patent Claim 2	Cisco IOS System
2[preamble]	The method of claim 1, wherein said step of selecting an entity pair further comprises:	Cisco IOS System discloses wherein said step of selecting an entity pair further comprises. For example, Cisco IOS Guide discloses the path that an LSP uses is determined by the LSP resource requirements and network resources, such as bandwidth and the MPLS TE SRLG

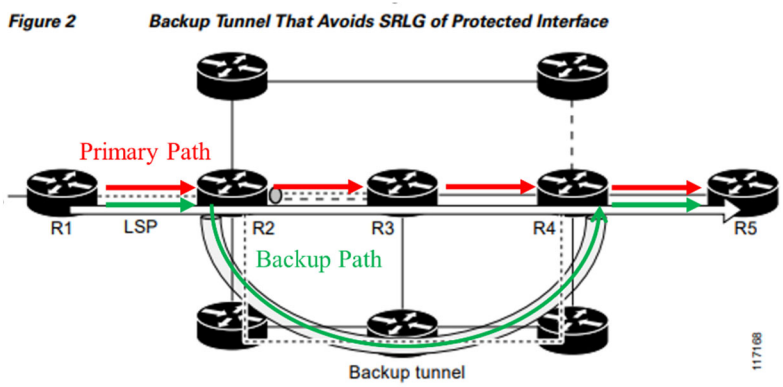
Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 41 of 698

No.	'821 Patent Claim 2	Cisco IOS System
		feature enhances backup tunnel path selection so a backup tunnel can avoid using links that are in the same SRLG as the interfaces it is protecting else when the protected link fails the backup tunnel fails too. *See Supra at* 1[c].
2[a]	selecting a working entity from said entity pair; and	Cisco IOS System discloses selecting a working entity from said entity pair. For example, Cisco IOS Guide discloses that when establishing a protected LSP, in advance, a secondary LSP is established to provide failure protection for the protected LSP. A person of ordinary skill in the art would understand the protected entity to be a working entity. Cisco IOS Guide at Page No. 603 ("[Path Protection Path protection provides an end-to-end failure recovery mechanism (that is, full path protection) for MPLS TE tunnels. A secondary LSP is established, in advance, to provide failure protection for the protected LSP that is carrying a tunnel's TE traffic. When there is a failure on the protected LSP, the headend router immediately enables the secondary LSP to temporarily carry the tunnel's traffic. If there is a failure on the secondary LSP, the tunnel no longer has path protection until the failure along the secondary path is cleared.]") *Id.* ("Some tunnels are more important than others. For example, you may have tunnels carrying VoIP traffic and tunnels carrying data traffic that are competing for the same resources. MPLS TE allows you to have some tunnels preempt others. Each tunnel has a priority, and more-important tunnels take precedence over less-important tunnels.")
2[b]	selecting a protection entity from said entity pair.	Cisco IOS System discloses selecting a protection entity from said entity pair. For example, Cisco IOS Guide discloses that when establishing a protected LSP, in advance, a secondary LSP is established to provide failure protection for the protected LSP. A person of ordinary skill in the art would understand the secondary entity to be a protection entity. Cisco IOS Guide at Page No. 603 ("[Path Protection

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 42 of 698

No.	'821 Patent Claim 2	Cisco IOS System
		Path protection provides an end-to-end failure recovery mechanism (that is, full path protection) for MPLS TE tunnels. A secondary LSP is established, in advance, to provide failure protection for the protected LSP that is carrying a tunnel's TE traffic. When there is a failure on the protected LSP, the headend router immediately enables the secondary LSP to temporarily carry the tunnel's traffic. If there is a failure on the secondary LSP, the tunnel no longer has path protection until the failure along the secondary path is cleared.]")  Fig. 1 (annotation added) Figure 1 shows a primary label-switched path (LSP) from router R1 to router R5. The LSP protects against the failure of the R2-R3 link at R2 via a backup tunnel to R4. If the R2-R3 link fails, link protection reroutes the LSP along the backup tunnel. However, the R2-R3 link and one of the backup tunnel links are in the same SRLG. So if the R2-R3 link fails, the backup tunnel may fail too. The MPLS TE SRLG feature enhances backup tunnel path selection so a backup tunnel can avoid using links that are in the same SRLG as the interfaces it is protecting. There are two ways for a backup tunnel to avoid the SRLGs of its protected interface: • The router does not create the backup tunnel unless it avoids SRLGs of the protected interface.]") Cisco IOS Guide at Page No. 580 ("[Open Shortest Path First (OSPF) and Intermediate System-to-Intermediate System (IS-IS) flood the SRLG membership information (including

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 43 of 698

No.	'821 Patent Claim 2	Cisco IOS System
		other TE link attributes such as bandwidth availability and affinity) so that all routers in the network have the SRLG information for each link. With this topology information, routers can compute backup tunnel paths that exclude links having SRLGs in common with their protected interfaces. As shown in Figure 2, the backup tunnel avoids the link between R2 and R3, which shares an SRLG with the protected interface.]")  Fig. 2 (annotation added) Cisco IOS Guide at Page No. 587 ("[Step 4 show mpls traffic-eng topology Use this command to show the SRLG link membership flooded via the Interior Gateway Protocol (IGP).]") Cisco IOS Guide at Page 632 ("[IS-IS—Intermediate System-to-Intermediate System. Link-state hierarchical routing protocol that calls for intermediate system (IS) routers to exchange routing information based on a single metric to determine network topology.]")

No.	'821 Patent Claim 3	Cisco IOS System
3	The method of claim 2, further comprising the step of selecting an active entity from the	Cisco IOS System discloses the step of selecting an active entity from the set consisting of said working entity and said protection entity.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 44 of 698

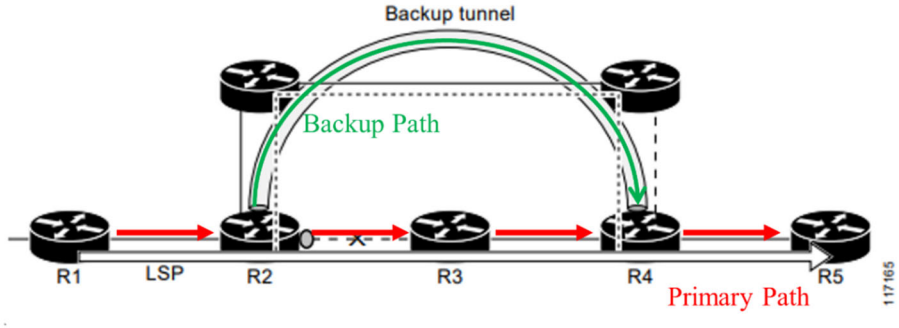
No.	'821 Patent Claim 3	Cisco IOS System
	set consisting of said working entity and said protection entity.	For example, Cisco IOS Guide discloses when there is a failure on the protected LSP, the headend router immediately enables the secondary LSP. Cisco IOS Guide further discloses that prior to a failure, the protected LSP actively moves data traffic. *See supra at* 2[preamble], 2[a], and 2[b]. Cisco IOS Guide at Page No. 603 ("[Path Protection Path protection provides an end-to-end failure recovery mechanism (that is, full path protection) for MPLS TE tunnels. A secondary LSP is established, in advance, to provide failure protection for the protected LSP that is carrying a tunnel's TE traffic. When there is a failure on the protected LSP, the headend router immediately enables the secondary LSP to temporarily carry the tunnel's traffic. If there is a failure on the secondary LSP, the tunnel no longer has path protection until the failure along the secondary path is cleared.]")

No.	'821 Patent Claim 4	Cisco IOS System
4	The method of claim 2, wherein selecting an entity pair further comprises minimizing an overall cost function.	Cisco IOS System discloses selecting an entity pair further comprises minimizing an overall cost function. For example, Cisco IOS Guide discloses the router does not create the backup tunnel unless it avoids SRLGs of the protected interface. Cisco IOS Guide further discloses that this selection method lowers the likelihood that both the protected and backup tunnel will fail. Thus, according to Orckit's apparent understanding of this limitation in its Infringement Disclosures, this limitation is met. Cisco IOS Guide at Page No. 639 ("[Bandwidth Protection NHOP and NNHOP backup tunnels can be used to provide bandwidth protection for rerouted LSPs. This is referred to as backup bandwidth. You can associate backup bandwidth with NHOP or NNHOP backup tunnels. This informs the router of the amount of backup bandwidth a particular backup tunnel can protect. When a router maps LSPs to backup tunnels, bandwidth protection ensures that an LSP uses a given backup tunnel only if there is sufficient backup bandwidth. The router selects which LSPs use which backup tunnels in

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 45 of 698

No.	'821 Patent Claim 4	Cisco IOS System
		order to provide maximum bandwidth protection. That is, the router determines the best way to map LSPs onto backup tunnels in order to maximize the number of LSPs that can be protected. For information about mapping tunnels and assigning backup bandwidth, see the "Backup Tunnel Selection Procedure" section on page 10.]") Cisco IOS Guide at Page No. 650 ("[For each protected link AB with a maximum reservable subpool value of n, there may be a path from node A to node B such that the difference between the maximum reservable global and the maximum reservable subpool is at least the value of n. If it is possible to find such paths for each link in the network, you can establish all the backup tunnels along such paths without any bandwidth reservations. If there is a single link failure, only one backup tunnel will use any link on its path. Because that path has at least n available bandwidth (in the global pool), assuming that marking and scheduling is configured to classify the subpool traffic into a priority queue, the subpool bandwidth is guaranteed. This approach allows sharing of the global pool bandwidth between backup tunnels protecting independent link failures. The backup tunnels are expected to be used for only a short period of time after a failure (until the headends of affected LSPs reroute those LSPs to other paths with available subpool bandwidth). The probability of multiple unrelated link failures is very small (in the absence of node or shared risk link group (SRLG) failures, which result in multiple link failures). Therefore, it is reasonable to assume that link failures are in practice independent with high probability. This "independent failure assumption" in combination with backup tunnels signaled without explicit bandwidth reservation enables efficient bandwidth sharing that yields substantial bandwidth savings. Backup tunnels protecting the subpool traffic do now draw bandwidth from any pool. Primary traffic using the global pool can use the entire global pool, and primary traffic using the subpool can use the entire subpool. Yet, subpool traffic has a complete bandwidth guarantee if there is a single link failure. A similar approach can be used for node and SRLG protection. However, the decision of where to put the backup tunnels is more complicated because both node and SRLG failures effectively result in the simultaneous failure of several links. Therefore, the backup tunnels protecting traffic traversing all affected links cannot be computed independently of each

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 46 of 698

No.	'821 Patent Claim 4	Cisco IOS System
		other. The backup tunnels protecting groups of links corresponding to different failures can still be computed independently of each other, which results in similar bandwidth savings.]") Cisco IOS Guide at Page No. 579 ("[MPLS Traffic Engineering Shared Risk Link Groups SRLGs refer to situations where links in a network share a common fiber (or a common physical attribute). If one link fails, other links in the group may fail too. Links in the group have a shared risk. Backup tunnels should avoid using links in the same SRLG as interfaces they are protecting. Otherwise, when the protected link fails the backup tunnel fails too. *Figure 1 Backup Tunnel in the Same SRLG as the Interface It Is Protecting*  Fig. 1 (annotation added) Figure 1 shows a primary label-switched path (LSP) from router R1 to router R5. The LSP protects against the failure of the R2-R3 link at R2 via a backup tunnel to R4. If the R2-R3 link fails, link protection reroutes the LSP along the backup tunnel. However, the R2-R3 link and one of the backup tunnel links are in the same SRLG. So if the R2-R3 link fails, the backup tunnel may fail too. The MPLS TE SRLG feature enhances backup tunnel path selection so a backup tunnel can avoid using links that are in the same SRLG as the interfaces it is protecting. There are two ways for a backup tunnel to avoid the SRLGs of its protected interface: • The router does not create the backup tunnel unless it avoids SRLGs of the protected interface.]")

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 47 of 698

No.	'821 Patent Claim 4	Cisco IOS System
		Cisco IOS Guide at Page No. 580 ("[Open Shortest Path First (OSPF) and Intermediate System-to-Intermediate System (IS-IS) flood the SRLG membership information (including other TE link attributes such as bandwidth availability and affinity) so that all routers in the network have the SRLG information for each link. With this topology information, routers can compute backup tunnel paths that exclude links having SRLGs in common with their protected interfaces.]") Cisco IOS Guide at Page No. 633 ("[OSPF—Open Shortest Path First. A link-state hierarchical Interior Gateway Protocol routing algorithm, derived from the IS-IS protocol. OSPF features include least-cost routing, multipath routing, and load balancing.]") In addition, at least under the apparent claim scope alleged by Orckit's Infringement Disclosures, Cisco IOS System in combination with (1) the knowledge of a person of ordinary skill in the art, alone or in further combination with (2) each (individually, as well as one or more together) of the references identified in element 4 of Exhibit E-2 renders the claim, including the present limitation, obvious. Below are examples of two such references. For example, Juniper IOS Guide discloses that through fate sharing, you can configure one or more backup paths that minimize the number of shared links with the primary paths. Juniper IOS Guide further discloses that fate sharing ensures that a single point of failure will not affect the primary and backup paths simultaneously. At least based upon Orckit's apparent understanding of this term as reflected in its Infringement Disclosures, this limitation is met. Juniper IOS Guide at Page No. 73 ("[Through fate sharing, you can configure backup paths that minimize the number of shared links and fiber paths with the primary paths as much as possible, to ensure that if a fiber is cut, the minimum amount of data is lost and a path still exists to the destination. For a backup path to work optimally, it must not share links or physical fiber paths with the primary path, ensuring that a single point of failure will not affect the primary and backup paths simultaneously. For more information about fate sharing, see the Junos OS Routing Protocols Configuration Guide.]")

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 48 of 698

No.	'821 Patent Claim 4	Cisco IOS System
		Juniper IOS Guide at Page No. 97 ("[Fate-sharing groups contain three types of objects: • Point-to-point links—Identified by the IP addresses at each end of the link. Unnumbered point-to-point links are typically identified by borrowing IP addresses from other interfaces. Order is not important; from 1.2.3.4 to 1.2.3.5 and from 1.2.3.5 to 1.2.3.4 have the same meaning. • Non-point-to-point links—Include links on a LAN interface (such as Gigabit Ethernet interfaces) or nonbroadcast multiaccess (NBMA) interfaces (such as Asynchronous Transfer Mode [ATM] or Frame Relay). You identify these links by their individual interface address. For example, if the LAN interface 192.168.200.0/24 has four routers attached to it, each router link is individually identified: from 192.168.200.1; # LAN interface of router 1 from 192.168.200.2; # LAN interface of router 2 from 192.168.200.3; # LAN interface of router 3 from 192.168.200.4; # LAN interface of router 4 You can list the addresses in any order. • A router node—Identified by its configured router ID. An object can be in any number of groups, and a group can contain any number of objects. Each group has a configurable cost attributed to it, which represents the level of impact this group has on CSPF computations. The higher the cost, the less likely a backup path will share with the primary path any objects in the group. The cost is directly comparable to traffic engineering metrics. By default, the cost is 1. Changing the fate-sharing database does not affect established LSPs until the next reoptimization of CSPF. The fate-sharing database does influence fast-reroute computations.]") For example, RFC 4872 discloses that the plurality of MPLS LSPs, i.e., transport entities, are selected in pairs of either working and protection LSPs, or working and restoration LSPs. RFC 4872 further discloses that the plurality of MPLS LSPs, i.e., transport entities, are selected to necessarily be "resource disjoint," i.e., where multiple resources, such as links, nodes, or Shared Risk Link Groups (SRLGs) do not share common links or failure probabilities. Thus, at least under the apparent claim scope alleged by Orckit's Infringement Disclosures, this limitation is met. To the extent that the IETF MPLS-TP System is found to

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 49 of 698

No.	'821 Patent Claim 4	Cisco IOS System
		not meet this limitation, if said subset comprises at least two entity pairs, selecting an entity pair from said subset that minimizes an entity cost function would have been obvious to a person having ordinary skill in the art, as explained below. "When a failure is detected by one or both endpoints of the LSP, both endpoints must select traffic from the other LSP. This action must be coordinated between node A and D. From this perspective, 1+1 bidirectional protection can be seen as a coordinated protection- switching mechanism between both endpoints. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 11. "1:1 Protection with Extra-Traffic The most common case of end-to-end 1:N protection is to establish, between the same endpoints, an end-to-end working LSP (thus, N = 1) and a dedicated end-to-end protecting LSP that are mutually link/ node/SRLG disjoint. This protects against working LSP failure(s)." RFC 4872 at 13. "An approach to reduce recovery resource requirements is to have protection LSPs sharing network resources when the working LSPs that they protect are physically (i.e., link, node, SRLG, etc.) disjoint. This mechanism is referred to as shared mesh restoration and is described in [RFC4426]. Shared-mesh restoration can be seen as a particular case of pre-planned LSP rerouting (see Section 8) that reduces the recovery resource requirements by allowing multiple protecting LSPs to share common link and node resources. Here also, the recovery resources for the protecting LSPs are pre-reserved during the provisioning phase, thus an explicit signaling action is required to activate (i.e., commit resource allocation at the data plane) a specific protecting LSP instantiated during the (pre-) provisioning phase. This requires restoration signaling along the protecting LSP." RFC 4872 at 20. "1. Introduction

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 50 of 698

No.	'821 Patent Claim 4	Cisco IOS System
		Generalized Multi-Protocol Label Switching (GMPLS) extends MPLS to include support for Layer-2 Switch Capable (L2SC), Time-Division Multiplex (TDM), Lambda Switch Capable (LSC), and Fiber Switch Capable (FSC) interfaces. GMPLS recovery uses control plane mechanisms (i.e., signaling, routing, and link management mechanisms) to support data plane fault recovery. Note that the analogous (data plane) fault detection mechanisms are required to be present in support of the control plane mechanisms. In this document, the term "recovery" is generically used to denote both protection and restoration; the specific terms "protection" and "restoration" are only used when differentiation is required. The subtle distinction between protection and restoration is made based on the resource allocation done during the recovery phase (see [RFC4427]). A functional description of GMPLS recovery is provided in [RFC4426] and should be considered as a companion document. The present document describes the protocol-specific procedures for GMPLS RSVP- TE (Resource ReSerVation Protocol - Traffic Engineering) signaling (see [RFC3473]) to support end-to-end recovery. End-to-end recovery refers to the recovery of an entire LSP from its head-end (ingress node endpoint) to its tail-end (egress node endpoint). With end-to- end recovery, working LSPs are assumed to be resource-disjoint (where a resource is a link, node, or Shared Risk Link Group (SRLG)) in the network so that they do not share any failure probability, but this is not mandatory. With respect to a given set of network resources, a pair of working/protecting LSPs SHOULD be resource disjoint in case of dedicated recovery type (see below). On the other hand, in case of shared recovery (see below), a group of working LSPs SHOULD be mutually resource-disjoint in order to allow for a (single and commonly) shared protecting LSP, itself resource-disjoint from each of the working LSPs. Note that resource disjointness is a necessary (but not sufficient) condition to ensure LSP recoverability. The present document addresses four types of end-to-end LSP recovery: 1) 1+1 (unidirectional/bidirectional) protection, 2) 1:N (N >= 1) LSP protection with extra-traffic, 3) pre-planned LSP rerouting without extra-traffic (including shared mesh), and 4) full LSP rerouting. 9) The simplest notion of end-to-end LSP protection is 1+1 unidirectional protection. Using this type of protection, a protecting LSP is signaled over a dedicated resource-disjoint alternate path to protect an associated working LSP. Normal traffic is

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 51 of 698

No.	'821 Patent Claim 4	Cisco IOS System
		simultaneously sent on both LSPs and a selector is used at the egress node to receive traffic from one of the LSPs. If a failure occurs along one of the LSPs, the egress node selects the traffic from the valid LSP. No coordination is required between the end nodes when a failure/switchover occurs. In 1+1 bidirectional protection, a protecting LSP is signaled over a dedicated resource-disjoint alternate path to protect the working LSP. Normal traffic is simultaneously sent on both LSPs (in both directions), and a selector is used at both ingress/egress nodes to receive traffic from the same LSP. This requires coordination between the end-nodes when switching to the protecting LSP. 10) In 1:N (N >= 1) protection with extra-traffic, the protecting LSP is a fully provisioned and resource-disjoint LSP from the N working LSPs, that allows for carrying extra-traffic. The N working LSPs MAY be mutually resource-disjoint. Coordination between end-nodes is required when switching from one of the working LSPs to the protecting LSP. As the protecting LSP is fully provisioned, default operations during protection switching are specified for a protecting LSP carrying extra-traffic, but this is not mandatory. Note that M:N protection is out of scope of this document (though mechanisms it defines may be extended to cover it). 11) Pre-planned LSP rerouting (or restoration) relies on the establishment between the same pair of end-nodes of a working LSP and a protecting LSP that is link/node/SRLG disjoint from the working one. Here, the recovery resources for the protecting LSP are pre-reserved but explicit action is required to activate (i.e., commit resource allocation at the data plane) a specific protecting LSP instantiated during the (pre-)provisioning phase. Since the protecting LSP is not "active" (i.e., fully instantiated), it cannot carry any extra-traffic. This does not mean that the corresponding resources cannot be used by other LSPs. Therefore, this mechanism protects against working LSP(s) failure(s) but requires activation of the protecting LSP after working LSP failure occurrence. This requires restoration signaling along the protecting path. "Shared-mesh" restoration can be seen as a particular case of pre-planned LSP rerouting that reduces the recovery resource requirements by allowing multiple protecting LSPs to share common link and node resources. The recovery resources are pre-reserved but explicit action is required to activate (i.e., commit resource allocation at the data plane) a specific protecting LSP instantiated during the (pre-)

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 52 of 698

No.	'821 Patent Claim 4	Cisco IOS System
		provisioning phase. This procedure requires restoration signaling along the protecting path. Note that in both cases, bandwidth pre-reserved for a protecting (but not activated) LSP can be made available for carrying extra traffic. LSPs for extra-traffic (with lower holding priority than the protecting LSP) can then be established using the bandwidth pre-reserved for the protecting LSP. Also, any lower priority LSP that use the pre-reserved resources for the protecting LSP(s) must be preempted during the activation of the protecting LSP. 12) Full LSP rerouting (or restoration) switches normal traffic to an alternate LSP that is not even partially established until after the working LSP failure occurs. The new alternate route is selected at the LSP head-end node, it may reuse resources of the failed LSP at intermediate nodes and may include additional intermediate nodes and/or links." " RFC 4872 at 3–5. "Consider the following network topology:  The paths [A,B,C,D] and [A,E,F,G,D] are node and link disjoint, ignoring the ingress/egress nodes A and D. A 1+1 protected path is established from A to D over [A,B,C,D] and [A,E,F,G,D], and traffic is transmitted simultaneously over both component paths (i.e., LSPs). During the provisioning phase, both LSPs are fully instantiated (and thus activated) so that no resource sharing can be done along the protecting LSP (nor can any extra-traffic be transported). It is also RECOMMENDED to set the N bit since no protection-switching signaling is assumed in this case. When a failure occurs (say, at node B) and is detected at end-node D, the receiver at D selects the normal traffic from the other LSP. From this perspective, 1+1 unidirectional protection can be seen as an uncoordinated protection-switching mechanism acting independently at both endpoints. Also, for the LSP under failure condition, it is RECOMMENDED to not set

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 53 of 698

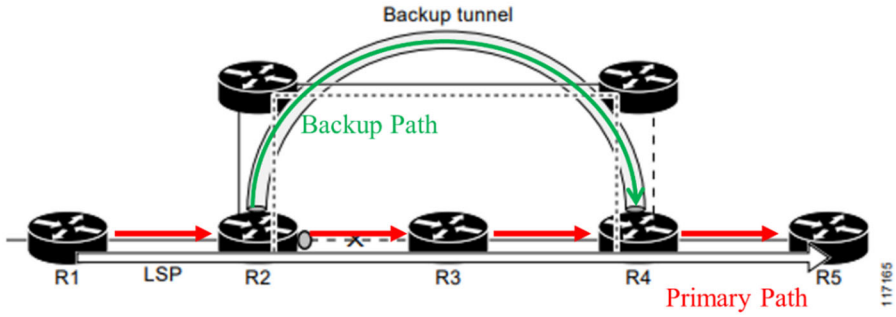
No.	'821 Patent Claim 4	Cisco IOS System
		the Path_State_Removed Flag of the ERROR_SPEC object (see [RFC3473]) upon PathErr message generation. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 9. "When a failure is detected by one or both endpoints of the LSP, both endpoints must select traffic from the other LSP. This action must be coordinated between node A and D. From this perspective, 1+1 bidirectional protection can be seen as a coordinated protection- switching mechanism between both endpoints. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 11. "(Full) LSP Rerouting LSP rerouting, on the other hand, switches normal traffic to an alternate LSP that is fully established only after failure occurrence. The new (alternate) route is selected at the LSP head- end and may reuse intermediate nodes included in the original route; it may also include additional intermediate nodes. For strict-hop routing, TE requirements can be directly applied to the route computation, and the failed node or link can be avoided. However, if the failure occurred within a loose-routed hop, the head-end node may not have enough information to reroute the LSP around the failure. Crankback signaling (see [CRANK]) and route exclusion techniques (see [RFC4874]) MAY be used in this case. The alternate route MAY be either computed on demand (that is, when the failure occurs; this is referred to as full LSP rerouting) or pre-computed and stored for use when the failure is reported. The latter offers faster restoration time. There is, however, a risk that the alternate route will become out of date through other changes in the network; this can be mitigated to some extent by periodic recalculation of idle alternate routes. (Full) LSP rerouting will be initiated by the head-end node that has either detected the LSP failure or received a Notify message and/or a PathErr message with the new error code/sub-code "Notify Error/LSP Locally Failed" for this LSP. The new LSP resources can be established using the make-before-break mechanism, where the new LSP is set up before the

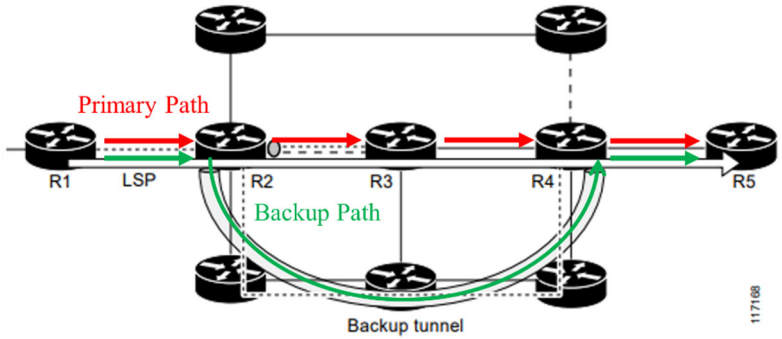
Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 54 of 698

No.	'821 Patent Claim 4	Cisco IOS System
		old LSP is torn down. This is done by using the mechanisms of the SESSION_ATTRIBUTE object and the Shared-Explicit (SE) reservation style (see [RFC3209]). Both the new and old LSPs can share resources at common nodes. Note that the make-before-break mechanism is not used to avoid disruption to the normal traffic flow (the latter has already been broken by the failure that is being repaired). However, it is valuable to retain the resources allocated on the original LSP that will be reused by the new alternate LSP." RFC 4872 at 25. "The procedure is as follows: 7) The initiating (source) node sends the normal traffic onto both the working and the protecting LSPs. Once completed, the source node sends reliably a Notify message to the destination with the new error code/sub-code "Notify Error/LSP Recovered" (Switchback Request). This Notify message includes the MESSAGE_ID object. The ACK_Desired flag MUST be set in this object to request the receiver to send an acknowledgment for the message (see [RFC2961]). 8) Upon receipt of this message, the destination selects the traffic from the working LSP. At the same time, it transmits the traffic onto both the working and protecting LSP. The destination then sends reliably a Notify message to the source confirming the completion of the operation. This message includes the MESSAGE_ID_ACK object to acknowledge reception of the received Notify message. This Notify message also includes the MESSAGE_ID object. The ACK_Desired flag MUST be set in this object to request the receiver to send an acknowledgment for the message (see [RFC2961]). 9) When the source node receives this Notify message, it switches to receive traffic from the working LSP. The source node then sends an Ack message to the destination node confirming that the LSP has been reverted." RFC 4872 at 27.

No.	'821 Patent Claim 5	Cisco IOS System
5	The method of claim 4, wherein said overall	Cisco IOS System discloses said overall cost function comprises substantially minimizing a probability of concurrent failure of said protection entity and said working entity.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 55 of 698

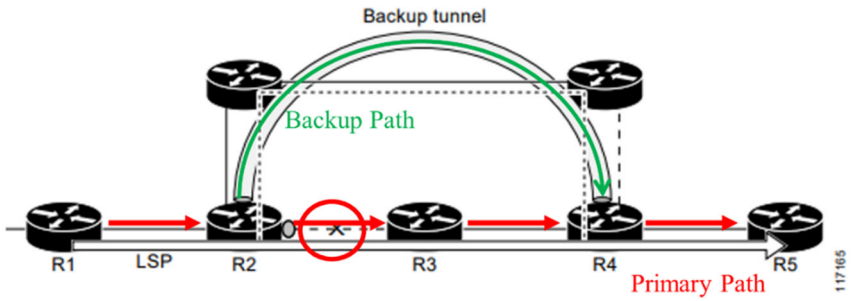
No.	'821 Patent Claim 5	Cisco IOS System
	cost function comprises substantially minimizing a probability of concurrent failure of said protection entity and said working entity.	For example, Cisco IOS Guide discloses the router does not create the backup tunnel unless it avoids SRLGs of the protected interface. Cisco IOS Guide further discloses that this selection method lowers the likelihood that both the protected and backup tunnel will fail. Thus, according to Orckit's apparent understanding of this limitation in its Infringement Disclosures, this limitation is met. Cisco IOS Guide at Page No. 651 ("[A similar approach can be used for node and SRLG protection. However, the decision of where to put the backup tunnels is more complicated because both node and SRLG failures effectively result in the simultaneous failure of several links. Therefore, the backup tunnels protecting traffic traversing all affected links cannot be computed independently of each other.]") Cisco IOS Guide at Page No. 579 ("[MPLS Traffic Engineering Shared Risk Link Groups SRLGs refer to situations where links in a network share a common fiber (or a common physical attribute). If one link fails, other links in the group may fail too. Links in the group have a shared risk. Backup tunnels should avoid using links in the same SRLG as interfaces they are protecting. Otherwise, when the protected link fails the backup tunnel fails too. *Figure 1 Backup Tunnel in the Same SRLG as the Interface It Is Protecting*  Fig. 1 (annotation added)

No.	'821 Patent Claim 5	Cisco IOS System
		Figure 1 shows a primary label-switched path (LSP) from router R1 to router R5. The LSP protects against the failure of the R2-R3 link at R2 via a backup tunnel to R4. If the R2-R3 link fails, link protection reroutes the LSP along the backup tunnel. However, the R2-R3 link and one of the backup tunnel links are in the same SRLG. So if the R2-R3 link fails, the backup tunnel may fail too. The MPLS TE SRLG feature enhances backup tunnel path selection so a backup tunnel can avoid using links that are in the same SRLG as the interfaces it is protecting. There are two ways for a backup tunnel to avoid the SRLGs of its protected interface: • The router does not create the backup tunnel unless it avoids SRLGs of the protected interface.]") Cisco IOS Guide at Page No. 580 ("[Open Shortest Path First (OSPF) and Intermediate System-to-Intermediate System (IS-IS) flood the SRLG membership information (including other TE link attributes such as bandwidth availability and affinity) so that all routers in the network have the SRLG information for each link. With this topology information, routers can compute backup tunnel paths that exclude links having SRLGs in common with their protected interfaces. As shown in Figure 2, the backup tunnel avoids the link between R2 and R3, which shares an SRLG with the protected interface.]") *Figure 2 Backup Tunnel That Avoids SRLG of Protected Interface*  Fig. 2 (annotation added)

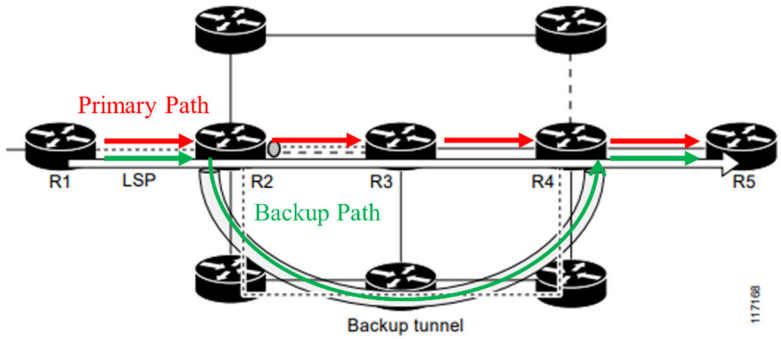
Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 57 of 698

No.	'821 Patent Claim 6	Cisco IOS System
6	The method of claim 4, wherein said overall cost function comprises a predefined entity cost metric.	Cisco IOS System discloses overall cost function comprises a predefined entity cost metric. For example, Cisco IOS Guide discloses that the best backup tunnel can be selected based on the backup bandwidth available. Thus, according to Orckit's apparent understanding of this limitation in its Infringement Disclosures, this limitation is met. Cisco IOS Guide at Page No. 253 ("[Path Option Selection for MPLS TE Tunnel LSPs This section contains the following topics about path option selection for MPLS TE Tunnel LSPs: • Constraint-Based Routing and Path Option Selection, page 5 • Tunnel Reoptimization and Path Option Selection, page 5 • Path Option Selection with Bandwidth Override, page 6 Constraint-Based Routing and Path Option Selection MPLS traffic engineering automatically establishes and maintains LSPs across the backbone by using the Resource Reservation Protocol (RSVP). The path that an LSP uses is determined by the LSP resource requirements and network resources, such as bandwidth. Traffic engineering tunnels are calculated at the LSP head based on a fit between required and available resources (constraint-based routing). Without the Path Option for Bandwidth Override feature, a TE tunnel establishes an LSP based on dynamic or explicit path options in order of preference. However, the bandwidth and other attributes configured on the TE tunnel allow the setup of an LSP only if LSP path options satisfy the constraints. If a path cannot be found that satisfies the configured path options, then the tunnel is not set up.]") Cisco IOS Guide at Page No. 639 ("[Bandwidth Protection NHOP and NNHOP backup tunnels can be used to provide bandwidth protection for rerouted LSPs. This is referred to as backup bandwidth. You can associate backup bandwidth with NHOP or NNHOP backup tunnels. This informs the router of the amount of backup bandwidth a particular backup tunnel can protect. When a router maps LSPs to backup tunnels, bandwidth protection ensures that an LSP uses a given backup tunnel only if there is sufficient backup bandwidth. The router selects which LSPs use which backup tunnels in order to provide maximum bandwidth protection. That is, the router determines the best way

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 58 of 698

No.	'821 Patent Claim 6	Cisco IOS System
		to map LSPs onto backup tunnels in order to maximize the number of LSPs that can be protected. For information about mapping tunnels and assigning backup bandwidth, see the "Backup Tunnel Selection Procedure" section on page 10.]") Cisco IOS Guide at Page No. 579 ("[MPLS Traffic Engineering Shared Risk Link Groups SRLGs refer to situations where links in a network share a common fiber (or a common physical attribute). If one link fails, other links in the group may fail too. Links in the group have a shared risk. Backup tunnels should avoid using links in the same SRLG as interfaces they are protecting. Otherwise, when the protected link fails the backup tunnel fails too. *Figure 1 Backup Tunnel in the Same SRLG as the Interface It Is Protecting*  Fig. 1 (annotation added) Figure 1 shows a primary label-switched path (LSP) from router R1 to router R5. The LSP protects against the failure of the R2-R3 link at R2 via a backup tunnel to R4. If the R2-R3 link fails, link protection reroutes the LSP along the backup tunnel. However, the R2-R3 link and one of the backup tunnel links are in the same SRLG. So if the R2-R3 link fails, the backup tunnel may fail too. The MPLS TE SRLG feature enhances backup tunnel path selection so a backup tunnel can avoid using links that are in the same SRLG as the interfaces it is protecting. There are two ways for a backup tunnel to avoid the SRLGs of its protected interface: • The router does not create the backup tunnel unless it avoids SRLGs of the protected interface.]")

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 59 of 698

No.	'821 Patent Claim 6	Cisco IOS System
		Cisco IOS Guide at Page No. 580 ("[Open Shortest Path First (OSPF) and Intermediate System-to-Intermediate System (IS-IS) flood the SRLG membership information (including other TE link attributes such as bandwidth availability and affinity) so that all routers in the network have the SRLG information for each link. With this topology information, routers can compute backup tunnel paths that exclude links having SRLGs in common with their protected interfaces. As shown in Figure 2, the backup tunnel avoids the link between R2 and R3, which shares an SRLG with the protected interface.]") *Figure 2 Backup Tunnel That Avoids SRLG of Protected Interface*  Fig. 2 (annotation added) Cisco IOS Guide at Page No. 587 ("[Step 4 show mpls traffic-eng topology "Use this command to show the SRLG link membership flooded via the Interior Gateway Protocol (IGP) .]") Cisco IOS Guide at Page No. 632 ("[IS-IS—Intermediate System-to-Intermediate System. Link-state hierarchical routing protocol that calls for intermediate system (IS) routers to exchange routing information based on a single metric to determine network topology.]") Cisco IOS Guide at Page No. 633 ("[OSPF—Open Shortest Path First. A link-state hierarchical Interior Gateway Protocol routing algorithm, derived from the IS-IS protocol. OSPF features include least-cost routing, multipath routing, and load balancing.]")

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 60 of 698

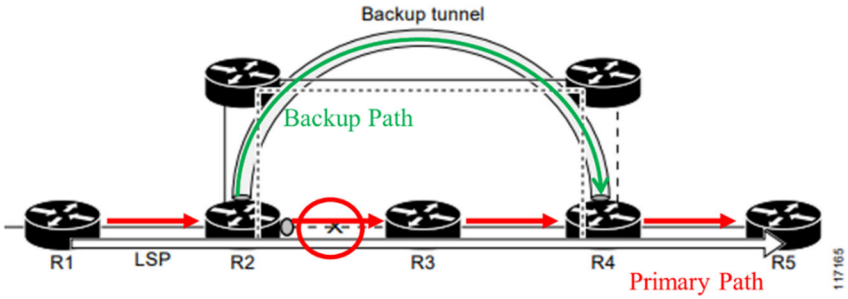
No.	'821 Patent Claim 7	Cisco IOS System
7	The method of claim 6, wherein said predefined entity cost metric is selected from the group consisting of interior gateway protocol (IGP) and traffic engineering (TE).	Cisco IOS System discloses said predefined entity cost metric is selected from the group consisting of interior gateway protocol (IGP) and traffic engineering (TE). For example, Cisco IOS Guide discloses the use of Open Shortest Path First (OSPF) and Intermediate System-to-Intermediate System (IS-IS) in order to determine membership information in the network, check network bandwidth, and to load balance among network paths. Cisco IOS Guide at Page No. 455 ("[The MPLS Traffic Engineering (TE)—Automatic Bandwidth Adjustment for TE Tunnels feature provides the means to automatically adjust the bandwidth allocation for traffic engineering tunnels based on their measured traffic load. The configured bandwidth in the running configuration is changed due to the automatic bandwidth behavior.]") Cisco IOS Guide at Page No. 580 ("[Open Shortest Path First (OSPF) and Intermediate System-to-Intermediate System (IS-IS) flood the SRLG membership information (including other TE link attributes such as bandwidth availability and affinity) so that all routers in the network have the SRLG information for each link. With this topology information, routers can compute backup tunnel paths that exclude links having SRLGs in common with their protected interfaces.]") Cisco IOS Guide at Page No. 587 ("[Step 4 show mpls traffic-eng topology Use this command to show the SRLG link membership flooded via the Interior Gateway Protocol (IGP).]") Cisco IOS Guide at Page No. 633 ("[OSPF—Open Shortest Path First. A link-state hierarchical Interior Gateway Protocol routing algorithm, derived from the IS-IS protocol. OSPF features include least-cost routing, multipath routing, and load balancing.]")

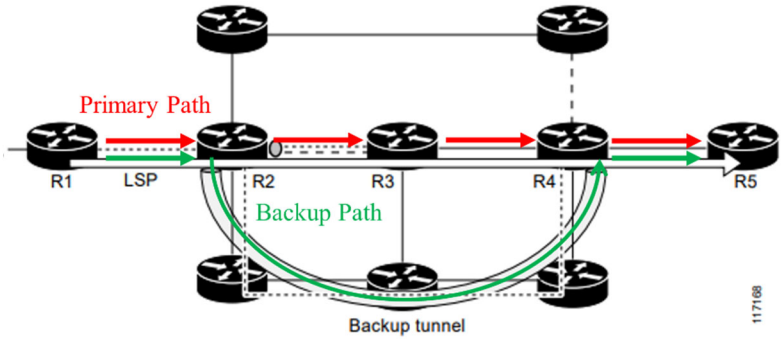
No.	'821 Patent Claim 8	Cisco IOS System
8	The method of claim 4, further comprising	Cisco IOS System discloses step of configuring said working entity as revertive.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 61 of 698

No.	'821 Patent Claim 8	Cisco IOS System
	the step of configuring said working entity as revertive.	For example, Cisco IOS Guide discloses when there is failure on the protected LSP, the headend router immediately enables the secondary LSP to temporarily carry the tunnel's traffic. Cisco IOS Guide at Page No. 603 ("[Path Protection Path protection provides an end-to-end failure recovery mechanism (that is, full path protection) for MPLS TE tunnels. A secondary LSP is established, in advance, to provide failure protection for the protected LSP that is carrying a tunnel's TE traffic. When there is a failure on the protected LSP, the headend router immediately enables the secondary LSP to temporarily carry the tunnel's traffic. If there is a failure on the secondary LSP, the tunnel no longer has path protection until the failure along the secondary path is cleared.]") Cisco IOS Guide at Page No. 603 ("Although not as fast as link or node protection, presignaling a secondary LSP is faster than configuring a secondary primary path option or allowing the tunnel's headend router to dynamically recalculate a path. The actual recovery time is topology-dependent, and affected by delay factors such as propagation delay or switch fabric latency.") Cisco IOS Guide at Page No. 604 ("The MPLS Traffic Engineering: Path Protection feature can recover after SSO. A tunnel configured for path protection may have two LSPs signaled simultaneously: the primary LSP that is carrying the traffic and the secondary LSP that carries traffic in case there is a failure along the primary path. Only information associated with one of those LSPs, the one that is currently carrying traffic, is synched to the standby RP. The standby RP, upon recovery, can determine from the checkpointed information whether the LSP was the primary or secondary. If the primary LSP was active during the switchover, only the primary LSP is recovered. The secondary LSP that was signaled and that provided path protection is resignaled after the TE recovery period is complete. This does not impact traffic on the tunnel because the secondary LSP was not carrying traffic.")

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 62 of 698

No.	'821 Patent Claim 9	Cisco IOS System
9[preamble]	The method of claim 4, wherein said overall cost function comprises:	Cisco IOS System discloses overall cost function comprises. *See Supra at* 4.
9[a]	selecting a subset of entity pairs wherein each entity pair of said subset has substantially minimum probability of a concurrent failure of said protection entity and said working entity; and	Cisco IOS System discloses selecting a subset of entity pairs wherein each entity pair of said subset has substantially minimum probability of a concurrent failure of said protection entity and said working entity. For example, Cisco IOS Guide discloses that the headend router creates backup tunnels in a manner that avoids using the same SRLGs of the protected interface. Cisco IOS Guide further discloses that this method lowers that likelihood of both the protected link and backup tunnel failing at the same time. Thus, according to Orckit's apparent understanding of this limitation in its Infringement Disclosures, this limitation is met. Cisco IOS Guide at Page No. 579 ("[MPLS Traffic Engineering Shared Risk Link Groups SRLGs refer to situations where links in a network share a common fiber (or a common physical attribute). If one link fails, other links in the group may fail too. Links in the group have a shared risk. Backup tunnels should avoid using links in the same SRLG as interfaces they are protecting. Otherwise, when the protected link fails the backup tunnel fails too. ***Figure 1 Backup Tunnel in the Same SRLG as the Interface It Is Protecting***  Fig. 1 (annotation added)

No.	'821 Patent Claim 9	Cisco IOS System
		Figure 1 shows a primary label-switched path (LSP) from router R1 to router R5. The LSP protects against the failure of the R2-R3 link at R2 via a backup tunnel to R4. If the R2-R3 link fails, link protection reroutes the LSP along the backup tunnel. However, the R2-R3 link and one of the backup tunnel links are in the same SRLG. So if the R2-R3 link fails, the backup tunnel may fail too. The MPLS TE SRLG feature enhances backup tunnel path selection so a backup tunnel can avoid using links that are in the same SRLG as the interfaces it is protecting. There are two ways for a backup tunnel to avoid the SRLGs of its protected interface: • The router does not create the backup tunnel unless it avoids SRLGs of the protected interface.]") Cisco IOS Guide at Page No. 580 ("[Open Shortest Path First (OSPF) and Intermediate System-to-Intermediate System (IS-IS) flood the SRLG membership information (including other TE link attributes such as bandwidth availability and affinity) so that all routers in the network have the SRLG information for each link. With this topology information, routers can compute backup tunnel paths that exclude links having SRLGs in common with their protected interfaces. As shown in Figure 2, the backup tunnel avoids the link between R2 and R3, which shares an SRLG with the protected interface.]") *Figure 2 Backup Tunnel That Avoids SRLG of Protected Interface*  Fig. 2 (annotation added)

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 64 of 698

No.	'821 Patent Claim 9	Cisco IOS System
		To the extent that Orckit contends, Cisco's products disclose this limitation, Cisco IOS describes that: **Using Backup Tunnels Signaled with Zero Bandwidth** Frequently it is desirable to use backup tunnels with zero signaled bandwidth, even when bandwidth protection is required. It may seem that if no bandwidth is explicitly reserved, no bandwidth guarantees can be provided. However, that is not necessarily true. In the following situation: • Only link protection is desired. • Bandwidth protection is desired only for sub-pool traffic. For each protected link AB with a maximum reservable subpool value of n, there may be a path from node A to node B such that the difference between the maximum reservable global and the maximum reservable subpool is at least the value of n. If it is possible to find such paths for each link in the network, you can establish all the backup tunnels along such paths without any bandwidth reservations. If there is a single link failure, only one backup tunnel will use any link on its path. Because that path has at least n available bandwidth (in the global pool), assuming that marking and scheduling is configured to classify the subpool traffic into a priority queue, the subpool bandwidth is guaranteed. This approach allows sharing of the global pool bandwidth between backup tunnels protecting independent link failures. The backup tunnels are expected to be used for only a short period of time after a failure (until the headends of affected LSPs reroute those LSPs to other paths with available subpool bandwidth). ***The probability of multiple unrelated link failures is very small (in the absence of node or shared risk link group (SRLG) failures, which result in multiple link failures). Therefore, it is reasonable to***

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 65 of 698

No.	'821 Patent Claim 9	Cisco IOS System
		assume that link failures are in practice independent with high probability. This "independent failure assumption" in combination with backup tunnels signaled without explicit bandwidth reservation enables efficient bandwidth sharing that yields substantial bandwidth savings. Cisco IOS Guide at Page No. 650 (emphasis added).
9[b]	if said subset comprises at least two entity pairs, selecting an entity pair from said subset that minimizes an entity cost function.	Cisco IOS System discloses if said subset comprises at least two entity pairs, selecting an entity pair from said subset that minimizes an entity cost function. For example, Cisco IOS Guide discloses that there may be more than one backup tunnel that has sufficient backup bandwidth to protect a given LSP, then the router chooses the one that has the least amount of backup bandwidth available. A person of ordinary skill in the art would understand from this disclosure that the entity pairs are selected from a subset of pairs. Thus, at least under the apparent claim scope alleged by Orckit's Infringement Disclosures, this limitation is met. To the extent that the Cisco IOS System is found to not meet this limitation, if said subset comprises at least two entity pairs, selecting an entity pair from said subset that minimizes an entity cost function would have been obvious to a person having ordinary skill in the art, as explained below. Cisco IOS Guide at Page No. 645 ("[Bandwidth Protection NHOP and NNHOP backup tunnels can be used to provide bandwidth protection for rerouted LSPs. This is referred to as backup bandwidth. You can associate backup bandwidth with NHOP or NNHOP backup tunnels. This informs the router of the amount of backup bandwidth a particular backup tunnel can protect. When a router maps LSPs to backup tunnels, bandwidth protection ensures that an LSP uses a given backup tunnel only if there is sufficient backup bandwidth. The router selects which LSPs use which backup tunnels in order to provide maximum bandwidth protection. That is, the router determines the best way to map LSPs onto backup tunnels in order to maximize the number of LSPs that can be protected. For information about mapping tunnels and assigning backup bandwidth, see the "Backup Tunnel Selection Procedure" section on page 10. Load Balancing on Limited-Bandwidth Backup Tunnels There may be more than one backup tunnel that has sufficient backup bandwidth to protect a given LSP. In this case, the router chooses the one that has the least amount of backup

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 66 of 698

No.	'821 Patent Claim 9	Cisco IOS System
		bandwidth available. This algorithm limits fragmentation, maintaining the largest amount of backup bandwidth available.]") In addition, at least under the apparent claim scope alleged by Orckit's Infringement Disclosures, Cisco IOS System in combination with (1) the knowledge of a person of ordinary skill in the art, alone or in further combination with (2) each (individually, as well as one or more together) of the references identified in element 9[b] of Exhibit E-2 renders the claim, including the present limitation, obvious. Below are examples of two such references. For example, Juniper IOS Guide discloses that through fate sharing, you can configure one or more backup paths that minimize the number of shared links with the primary paths. Juniper IOS Guide further discloses that fate sharing ensures that a single point of failure will not affect the primary and backup paths simultaneously. At least based upon Orckit's apparent understanding of this term as reflected in its Infringement Disclosures, this limitation is met. Juniper IOS Guide at Page No. 73 ("[Through fate sharing, you can configure backup paths that minimize the number of shared links and fiber paths with the primary paths as much as possible, to ensure that if a fiber is cut, the minimum amount of data is lost and a path still exists to the destination. For a backup path to work optimally, it must not share links or physical fiber paths with the primary path, ensuring that a single point of failure will not affect the primary and backup paths simultaneously. For more information about fate sharing, see the Junos OS Routing Protocols Configuration Guide.]") Juniper IOS Guide at Page No. 97 ("[Fate-sharing groups contain three types of objects: • Point-to-point links—Identified by the IP addresses at each end of the link. Unnumbered point-to-point links are typically identified by borrowing IP addresses from other interfaces. Order is not important; from 1.2.3.4 to 1.2.3.5 and from 1.2.3.5 to 1.2.3.4 have the same meaning. • Non-point-to-point links—Include links on a LAN interface (such as Gigabit Ethernet interfaces) or nonbroadcast multiaccess (NBMA) interfaces (such as Asynchronous Transfer

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 67 of 698

No.	'821 Patent Claim 9	Cisco IOS System
		Mode [ATM] or Frame Relay). You identify these links by their individual interface address. For example, if the LAN interface 192.168.200.0/24 has four routers attached to it, each router link is individually identified: from 192.168.200.1; # LAN interface of router 1 from 192.168.200.2; # LAN interface of router 2 from 192.168.200.3; # LAN interface of router 3 from 192.168.200.4; # LAN interface of router 4 You can list the addresses in any order. • A router node—Identified by its configured router ID. An object can be in any number of groups, and a group can contain any number of objects. Each group has a configurable cost attributed to it, which represents the level of impact this group has on CSPF computations. The higher the cost, the less likely a backup path will share with the primary path any objects in the group. The cost is directly comparable to traffic engineering metrics. By default, the cost is 1. Changing the fate-sharing database does not affect established LSPs until the next reoptimization of CSPF. The fate-sharing database does influence fast-reroute computations.]") For example, RFC 4872 discloses that the plurality of MPLS LSPs, i.e., transport entities, are selected in pairs of either working and protection LSPs, or working and restoration LSPs. RFC 4872 further discloses that the plurality of MPLS LSPs, i.e., transport entities, are selected to necessarily be "resource disjoint," i.e., where multiple resources, such as links, nodes, or Shared Risk Link Groups (SRLGs) do not share common links or failure probabilities. Thus, at least under the apparent claim scope alleged by Orckit's Infringement Disclosures, this limitation is met. To the extent that the IETF MPLS-TP System is found to not meet this limitation, if said subset comprises at least two entity pairs, selecting an entity pair from said subset that minimizes an entity cost function would have been obvious to a person having ordinary skill in the art, as explained below. "When a failure is detected by one or both endpoints of the LSP, both endpoints must select traffic from the other LSP. This action must be coordinated between node A and D. From this perspective, 1+1 bidirectional protection can be seen as a coordinated protection- switching mechanism between both endpoints.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 68 of 698

No.	'821 Patent Claim 9	Cisco IOS System
		Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 11. "1:1 Protection with Extra-Traffic The most common case of end-to-end 1:N protection is to establish, between the same endpoints, an end-to-end working LSP (thus, N = 1) and a dedicated end-to-end protecting LSP that are mutually link/ node/SRLG disjoint. This protects against working LSP failure(s)." RFC 4872 at 13. "An approach to reduce recovery resource requirements is to have protection LSPs sharing network resources when the working LSPs that they protect are physically (i.e., link, node, SRLG, etc.) disjoint. This mechanism is referred to as shared mesh restoration and is described in [RFC4426]. Shared-mesh restoration can be seen as a particular case of pre-planned LSP rerouting (see Section 8) that reduces the recovery resource requirements by allowing multiple protecting LSPs to share common link and node resources. Here also, the recovery resources for the protecting LSPs are pre-reserved during the provisioning phase, thus an explicit signaling action is required to activate (i.e., commit resource allocation at the data plane) a specific protecting LSP instantiated during the (pre-) provisioning phase. This requires restoration signaling along the protecting LSP." RFC 4872 at 20. "1. Introduction Generalized Multi-Protocol Label Switching (GMPLS) extends MPLS to include support for Layer-2 Switch Capable (L2SC), Time-Division Multiplex (TDM), Lambda Switch Capable (LSC), and Fiber Switch Capable (FSC) interfaces. GMPLS recovery uses control plane mechanisms (i.e., signaling, routing, and link management mechanisms) to support data plane fault recovery. Note that the analogous (data plane) fault detection mechanisms are required to be present in support of the control plane mechanisms. In this document, the term "recovery" is generically used to denote both protection and restoration; the specific terms "protection" and "restoration" are only used when differentiation is required. The subtle

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 69 of 698

No.	'821 Patent Claim 9	Cisco IOS System
		distinction between protection and restoration is made based on the resource allocation done during the recovery phase (see [RFC4427]). A functional description of GMPLS recovery is provided in [RFC4426] and should be considered as a companion document. The present document describes the protocol-specific procedures for GMPLS RSVP- TE (Resource ReSerVation Protocol - Traffic Engineering) signaling (see [RFC3473]) to support end-to-end recovery. End-to-end recovery refers to the recovery of an entire LSP from its head-end (ingress node endpoint) to its tail-end (egress node endpoint). With end-to- end recovery, working LSPs are assumed to be resource-disjoint (where a resource is a link, node, or Shared Risk Link Group (SRLG)) in the network so that they do not share any failure probability, but this is not mandatory. With respect to a given set of network resources, a pair of working/protecting LSPs SHOULD be resource disjoint in case of dedicated recovery type (see below). On the other hand, in case of shared recovery (see below), a group of working LSPs SHOULD be mutually resource-disjoint in order to allow for a (single and commonly) shared protecting LSP, itself resource-disjoint from each of the working LSPs. Note that resource disjointness is a necessary (but not sufficient) condition to ensure LSP recoverability. The present document addresses four types of end-to-end LSP recovery: 1) 1+1 (unidirectional/bidirectional) protection, 2) 1:N (N >= 1) LSP protection with extra-traffic, 3) pre-planned LSP rerouting without extra-traffic (including shared mesh), and 4) full LSP rerouting. 13) The simplest notion of end-to-end LSP protection is 1+1 unidirectional protection. Using this type of protection, a protecting LSP is signaled over a dedicated resource-disjoint alternate path to protect an associated working LSP. Normal traffic is simultaneously sent on both LSPs and a selector is used at the egress node to receive traffic from one of the LSPs. If a failure occurs along one of the LSPs, the egress node selects the traffic from the valid LSP. No coordination is required between the end nodes when a failure/switchover occurs. In 1+1 bidirectional protection, a protecting LSP is signaled over a dedicated resource-disjoint alternate path to protect the working LSP. Normal traffic is simultaneously sent on both LSPs (in both directions), and a selector is used at both

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 70 of 698

No.	'821 Patent Claim 9	Cisco IOS System
		ingress/egress nodes to receive traffic from the same LSP. This requires coordination between the end-nodes when switching to the protecting LSP. 14) In 1:N (N >= 1) protection with extra-traffic, the protecting LSP is a fully provisioned and resource-disjoint LSP from the N working LSPs, that allows for carrying extra-traffic. The N working LSPs MAY be mutually resource-disjoint. Coordination between end-nodes is required when switching from one of the working LSPs to the protecting LSP. As the protecting LSP is fully provisioned, default operations during protection switching are specified for a protecting LSP carrying extra-traffic, but this is not mandatory. Note that M:N protection is out of scope of this document (though mechanisms it defines may be extended to cover it). 15) Pre-planned LSP rerouting (or restoration) relies on the establishment between the same pair of end-nodes of a working LSP and a protecting LSP that is link/node/SRLG disjoint from the working one. Here, the recovery resources for the protecting LSP are pre-reserved but explicit action is required to activate (i.e., commit resource allocation at the data plane) a specific protecting LSP instantiated during the (pre-)provisioning phase. Since the protecting LSP is not "active" (i.e., fully instantiated), it cannot carry any extra-traffic. This does not mean that the corresponding resources cannot be used by other LSPs. Therefore, this mechanism protects against working LSP(s) failure(s) but requires activation of the protecting LSP after working LSP failure occurrence. This requires restoration signaling along the protecting path. "Shared-mesh" restoration can be seen as a particular case of pre-planned LSP rerouting that reduces the recovery resource requirements by allowing multiple protecting LSPs to share common link and node resources. The recovery resources are pre-reserved but explicit action is required to activate (i.e., commit resource allocation at the data plane) a specific protecting LSP instantiated during the (pre-) provisioning phase. This procedure requires restoration signaling along the protecting path. Note that in both cases, bandwidth pre-reserved for a protecting (but not activated) LSP can be made available for carrying extra traffic. LSPs for extra-traffic (with lower holding priority than the protecting LSP) can then be established using the bandwidth pre-reserved for the protecting LSP. Also, any lower priority LSP that use the pre-

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 71 of 698

No.	'821 Patent Claim 9	Cisco IOS System
		reserved resources for the protecting LSP(s) must be preempted during the activation of the protecting LSP. 16) Full LSP rerouting (or restoration) switches normal traffic to an alternate LSP that is not even partially established until after the working LSP failure occurs. The new alternate route is selected at the LSP head-end node, it may reuse resources of the failed LSP at intermediate nodes and may include additional intermediate nodes and/or links." " RFC 4872 at 3–5. "Consider the following network topology:  The paths [A,B,C,D] and [A,E,F,G,D] are node and link disjoint, ignoring the ingress/egress nodes A and D. A 1+1 protected path is established from A to D over [A,B,C,D] and [A,E,F,G,D], and traffic is transmitted simultaneously over both component paths (i.e., LSPs). During the provisioning phase, both LSPs are fully instantiated (and thus activated) so that no resource sharing can be done along the protecting LSP (nor can any extra-traffic be transported). It is also RECOMMENDED to set the N bit since no protection-switching signaling is assumed in this case. When a failure occurs (say, at node B) and is detected at end-node D, the receiver at D selects the normal traffic from the other LSP. From this perspective, 1+1 unidirectional protection can be seen as an uncoordinated protection-switching mechanism acting independently at both endpoints. Also, for the LSP under failure condition, it is RECOMMENDED to not set the Path_State_Removed Flag of the ERROR_SPEC object (see [RFC3473]) upon PathErr message generation. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 9.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 72 of 698

No.	'821 Patent Claim 9	Cisco IOS System
		"When a failure is detected by one or both endpoints of the LSP, both endpoints must select traffic from the other LSP. This action must be coordinated between node A and D. From this perspective, 1+1 bidirectional protection can be seen as a coordinated protection- switching mechanism between both endpoints. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 11. "(Full) LSP Rerouting LSP rerouting, on the other hand, switches normal traffic to an alternate LSP that is fully established only after failure occurrence. The new (alternate) route is selected at the LSP head- end and may reuse intermediate nodes included in the original route; it may also include additional intermediate nodes. For strict-hop routing, TE requirements can be directly applied to the route computation, and the failed node or link can be avoided. However, if the failure occurred within a loose-routed hop, the head-end node may not have enough information to reroute the LSP around the failure. Crankback signaling (see [CRANK]) and route exclusion techniques (see [RFC4874]) MAY be used in this case. The alternate route MAY be either computed on demand (that is, when the failure occurs; this is referred to as full LSP rerouting) or pre-computed and stored for use when the failure is reported. The latter offers faster restoration time. There is, however, a risk that the alternate route will become out of date through other changes in the network; this can be mitigated to some extent by periodic recalculation of idle alternate routes. (Full) LSP rerouting will be initiated by the head-end node that has either detected the LSP failure or received a Notify message and/or a PathErr message with the new error code/sub-code "Notify Error/LSP Locally Failed" for this LSP. The new LSP resources can be established using the make-before-break mechanism, where the new LSP is set up before the old LSP is torn down. This is done by using the mechanisms of the SESSION_ATTRIBUTE object and the Shared-Explicit (SE) reservation style (see [RFC3209]). Both the new and old LSPs can share resources at common nodes. Note that the make-before-break mechanism is not used to avoid disruption to the normal traffic flow (the latter has already been broken by the failure that is being repaired). However,

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 73 of 698

No.	'821 Patent Claim 9	Cisco IOS System
		it is valuable to retain the resources allocated on the original LSP that will be reused by the new alternate LSP." RFC 4872 at 25. "The procedure is as follows: 10) The initiating (source) node sends the normal traffic onto both the working and the protecting LSPs. Once completed, the source node sends reliably a Notify message to the destination with the new error code/sub-code "Notify Error/LSP Recovered" (Switchback Request). This Notify message includes the MESSAGE_ID object. The ACK_Desired flag MUST be set in this object to request the receiver to send an acknowledgment for the message (see [RFC2961]). 11) Upon receipt of this message, the destination selects the traffic from the working LSP. At the same time, it transmits the traffic onto both the working and protecting LSP. The destination then sends reliably a Notify message to the source confirming the completion of the operation. This message includes the MESSAGE_ID_ACK object to acknowledge reception of the received Notify message. This Notify message also includes the MESSAGE_ID object. The ACK_Desired flag MUST be set in this object to request the receiver to send an acknowledgment for the message (see [RFC2961]). 12) When the source node receives this Notify message, it switches to receive traffic from the working LSP. The source node then sends an Ack message to the destination node confirming that the LSP has been reverted." RFC 4872 at 27.

No.	'821 Patent Claim 10	Cisco IOS Guide
10	The method of claim 9, wherein said entity cost function comprises a predefined metric.	Cisco IOS System discloses overall cost function comprises a predefined entity cost metric. *See Supra at* 6.

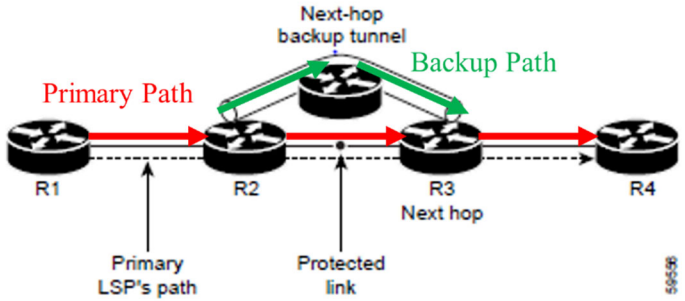
Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 74 of 698

No.	'821 Patent Claim 11	Cisco IOS System
11	The method of claim 10, wherein said predefined metric is selected from the group consisting of interior gateway protocol (IGP) and traffic engineering (TE).	Cisco IOS System discloses said predefined entity cost metric is selected from the group consisting of interior gateway protocol (IGP) and traffic engineering (TE). *See Supra at* 7.

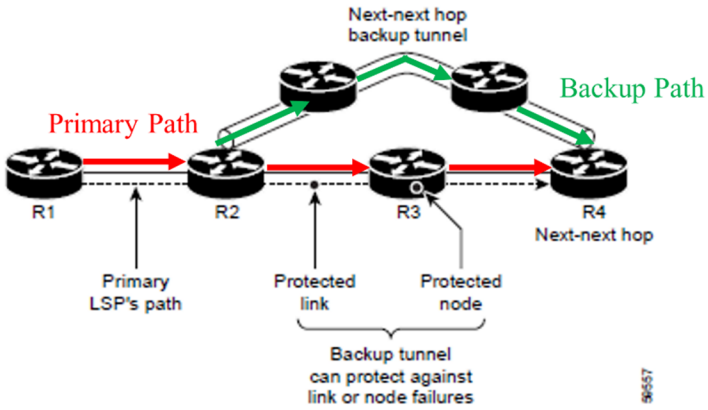
No.	'821 Patent Claim 12	Cisco IOS System
12	The method of claim 10, further comprising the step of configuring said working entity as revertive.	Cisco IOS System discloses the step of configuring said working entity as revertive. *See Supra at* 8.

No.	'821 Patent Claim 13	Cisco IOS System
13	The method of claim 1, further comprising the step of: if said entity pair reselection results in both working and protection entities being replaced, sequentially replacing said working entity and said protection entity.	Cisco IOS System discloses the method of claim 1, further comprising the step of: if said entity pair reselection results in both working and protection entities being replaced, sequentially replacing said working entity and said protection entity. For example, Cisco IOS Guide discloses Cisco IOS supports Fast reroute, link protection, and node protection in which upon failure of the LSP path, the bypass/backup tunnel is established, and when there are some changes in the previous backup tunnel a new backup tunnel is selected. Cisco IOS Guide further discloses failover within SRLG groups. Thus, according to Orckit's apparent understanding of this limitation in its Infringement Disclosures, this limitation is met. Cisco IOS Guide at Page No. 637 ("[Fast Reroute Fast Reroute (FRR) is a mechanism for protecting MPLS TE LSPs from link and node failures by locally repairing the LSPs at the point of failure, allowing data to continue to flow on them

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 75 of 698

No.	'821 Patent Claim 13	Cisco IOS System
		while their headend routers attempt to establish new end-to-end LSPs to replace them. FRR locally repairs the protected LSPs by rerouting them over backup tunnels that bypass failed links or node. Link Protection Backup tunnels that bypass only a single link of the LSP's path provide link protection. They protect LSPs if a link along their path fails by rerouting the LSP's traffic to the next hop (bypassing the failed link). These are referred to as next-hop (NHOP) backup tunnels because they terminate at the LSP's next hop beyond the point of failure. Figure 1 illustrates an NHOP backup tunnel.]") *Figure 1 NHOP Backup Tunnel*  Fig. 1 (annotation added) Cisco IOS Guide at Page No. 638 ("[Node Protection FRR provides node protection for LSPs. Backup tunnels that bypass next-hop nodes along LSP paths are called next-next-hop (NNHOP) backup tunnels because they terminate at the node following the next-hop node of the LSP paths, thereby bypassing the next-hop node. They protect LSPs if a node along their path fails by enabling the node upstream of the failure to reroute the LSPs and their traffic around the failed node to the next-next hop. FRR supports the use of RSVP Hellos to accelerate the detection of node failures. NNHOP backup tunnels also provide protection from link failures, because they bypass the failed link and the node.]")

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 76 of 698

No.	'821 Patent Claim 13	Cisco IOS System
		Figure 2 illustrates an NNHOP backup tunnel. ***Figure 2 NNHOP Backup Tunnel***  Fig. 2 (annotation added) Cisco IOS Guide at Page No. 638 ("[If an LSP is using a backup tunnel and something changes so that the LSP is no longer appropriate for the backup tunnel, the LSP is torn down. Such changes are the following: • Backup bandwidth of the backup tunnel is reduced. • Backup bandwidth type of backup tunnel is changed to a type that is incompatible with the primary LSP. • Primary LSP is modified so that FRR is disabled. (The no mpls traffic-eng fast-reroute command is entered.]") Cisco IOS Guide at Page No. 648 ("[After a backup tunnel has been chosen for an LSP, conditions may change that will cause us to reevaluate this choice. This reevaluation, if successful, is called promotion. Such conditions may include: 1. A new backup tunnel comes up. 2. The currently chosen backup tunnel for this LSP goes down.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 77 of 698

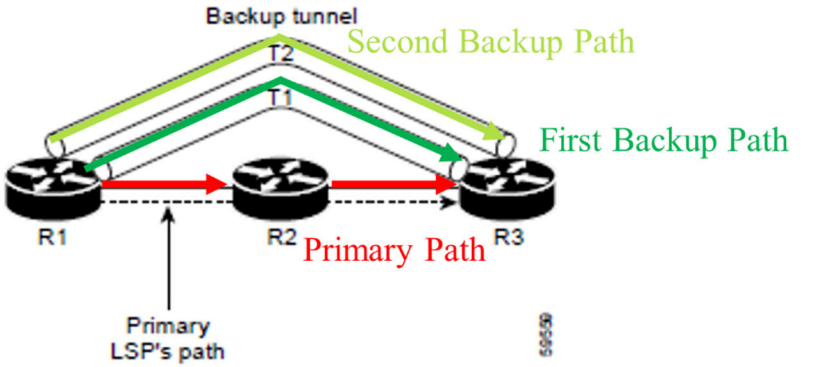
No.	'821 Patent Claim 13	Cisco IOS System
		3. A backup tunnel's available backup bandwidth increases. For example, an LSP protected by the tunnel has been reoptimized by the headend to use another path. For cases 1 and 2, the LSP's backup tunnel is evaluated immediately. Case 3 is addressed by periodically reevaluating LSP-to-backup tunnel mappings. By default, background reevaluation is performed every 5 minutes. This interval is configurable via the mpls traffic-eng fast-reroute timers command.]") Cisco IOS Guide at Page No. 654 ("[Tunnel Reoptimization Procedure Tunnel reoptimization is the signaling of an LSP that is more optimal than the LSP a TE tunnel is currently using (for example, it may be shorter or may have a lower cost), and the switching over of the tunnel's data to use this new LSP. The new more optimal TE LSP is always established and the data moved onto it before the original LSP is torn down (so it is called the "make before break" procedure). This ensures that no data packets are lost during the transition to the new LSP.]")

No.	'821 Patent Claim 14	Cisco IOS System
14[preamble]	A system for selecting entities within an MPLS network, comprising:	Cisco IOS System discloses a system for selecting entities within an MPLS network, comprising. For example, the Cisco IOS Guide discloses that the cisco router supporting Cisco IOS Release 12.2SR which supports MPLS traffic engineering and let the headend router of a TE tunnel control the path its traffic takes to a particular destination. *See Supra at* 1[preamble] and 1[a].
14[a]	a data structure comprising a plurality of transport entity descriptors;	Cisco IOS Guide discloses a plurality of transport entity descriptors. For example, Cisco IOS Guide discloses multiple backup path options are configured for each primary option. Cisco IOS further discloses that these options are used by and manipulated by a headend router with computer memory for storing data. Thus, a person of ordinary skill in the art would understand that the headend router includes a data structure for storing these entity descriptors.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 78 of 698

No.	'821 Patent Claim 14	Cisco IOS System
		Cisco IOS Guide at Page No. 640 ("[Backup Tunnel Support Backup tunnel support has the following capabilities: • Backup Tunnels Can Terminate at the Next-Next Hop to Support FRR, page 6 • Multiple Backup Tunnels Can Protect the Same Interface, page 6 • Backup Tunnels Provide Scalability, page 7 Backup Tunnels Can Terminate at the Next-Next Hop to Support FRR Backup tunnels that terminate at the next-next hop protect both the downstream link and node. This provides protection for link and node failures. For more detailed information, see the "Node Protection" section on page 4. Multiple Backup Tunnels Can Protect the Same Interface There is no limit (except memory limitations) to the number of backup tunnels that can protect a given interface. In many topologies, support for node protection requires supporting multiple backup tunnels per protected interface. These backup tunnels can terminate at the same destination or at different destinations. That is, for a given protected interface, you can configure multiple NHOP or NNHOP backup tunnels. This allows redundancy and load balancing.]") Cisco IOS Guide at Page No. 603 ("MPLS TE LSPs let the headend of a TE tunnel control the path its traffic takes to a particular destination. This method is more flexible than forwarding traffic based only on a destination address.") Cisco IOS Guide at Page No. 637 ("Fast Reroute (FRR) is a mechanism for protecting MPLS TE LSPs from link and node failures by locally repairing the LSPs at the point of failure, allowing data to continue to flow on them while their headend routers attempt to establish new end-to-end LSPs to replace them. FRR locally repairs the protected LSPs by rerouting them over backup tunnels that bypass failed links or node.") Cisco IOS Guide at Page No. 649 ("The backup bandwidth is used by the point of local repair (PLR) (that is, the headend of the backup tunnel) to decide how much primary traffic can be rerouted to this backup tunnel if there is a failure.")

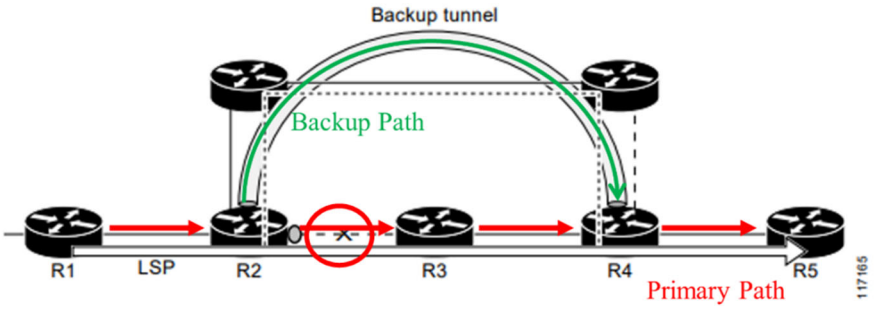
Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 79 of 698

No.	'821 Patent Claim 14	Cisco IOS System
		Cisco IOS Guide at Page No. 643 ("[Backup Tunnels Terminating at the Same Destination Figure 4 shows how backup tunnels terminating at the same location can be used for redundancy and load balancing. Redundancy and load balancing work for both NHOP and NNHOP backup tunnels.]") *Figure 4 Backup Tunnels That Terminate at the Same Destination*  Fig. 4 (annotation added) Cisco IOS Guide at Page No. 644 ("[In this illustration, there are three routers: R1, R2, and R3. At R1 two NNHOP backup tunnels (T1 and T2) go from R1 to R3 without traversing R2. Redundancy—If R2 fails or the link from R1 to R2 fails, either backup tunnel can be used. If one backup tunnel is down, the other can be used. LSPs are assigned to backup tunnels when the LSPs are first established. This is done before a failure. Load balancing—If neither backup tunnel has enough bandwidth to back up all LSPs, both tunnels can be used. Some LSPs will use one backup tunnel, other LSPs will use the other backup tunnel. The router decides the best way to fit the LSPs onto the backup tunnels.]") Cisco IOS Guide at Page No. 603 ("[Enhanced Path Protection Enhanced path protection provides support of multiple backup path options per primary path option. You can configure up to eight backup path options for a given primary path option. Only one of the configured backup path options is actively signaled at any time.

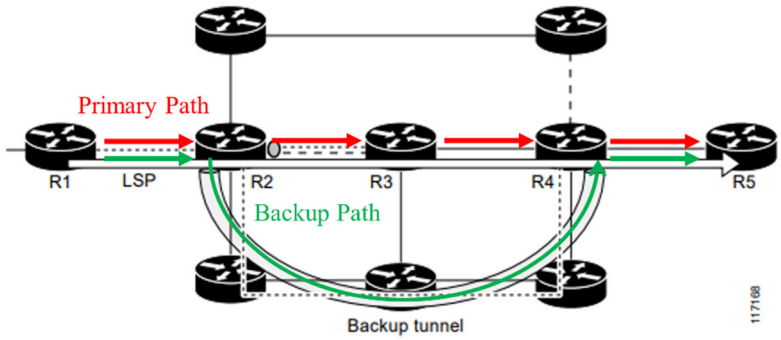
Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 80 of 698

No.	'821 Patent Claim 14	Cisco IOS System
		After you enter the mpls traffic-eng path-option list command, you can enter the backup path priority in the number argument of the path-option command. A lower identifier represents a higher priority. Priorities are configurable for each backup path option. Multiple backup path options and a single backup path option cannot coexist to protect a primary path option.]")
14[b]	an entity protection switch configured to switch between a working entity and a protection entity; and	Cisco IOS System discloses switch between a working entity and a protection entity. For example, Cisco IOS Guide discloses that when the protected LSP fails, the headend router immediately enables the secondary LSP to temporarily carry the tunnel's traffic. A person of ordinary skill in the art would understand that the corresponding logic within the headend router acts as an entity protection switch. Thus, according to Orckit's apparent understanding of this limitation in its Infringement Disclosures, this limitation is met. Cisco IOS Guide at Page No. 603 ("[Path Protection Path protection provides an end-to-end failure recovery mechanism (that is, full path protection) for MPLS TE tunnels. A secondary LSP is established, in advance, to provide failure protection for the protected LSP that is carrying a tunnel's TE traffic. When there is a failure on the protected LSP, the headend router immediately enables the secondary LSP to temporarily carry the tunnel's traffic. If there is a failure on the secondary LSP, the tunnel no longer has path protection until the failure along the secondary path is cleared.]") Cisco IOS Guide at Page No. 1101 ("Ensure that a minimum of two paths are present for the protected prefix w.x.y.z in BGP in steady state condition on the PE. The path using the protected PE should be the BGP best-path before failover occurs.") Cisco IOS Guide at Page No. 1182 ("The BGP PIC Edge for IP and MPLS-VPN feature improves convergence after a network failure. This convergence is applicable to both core and edge failures and can be used in both IP and MPLS networks. BGP PIC Edge for IP and MPLS-VPN feature creates and stores a backup/alternate path in the routing information base (RIB), forwarding information base (FIB) and in Cisco Express Forwarding, so that when a failure is detected, the backup/alternate path can immediately take over, thus enabling fast failover.")

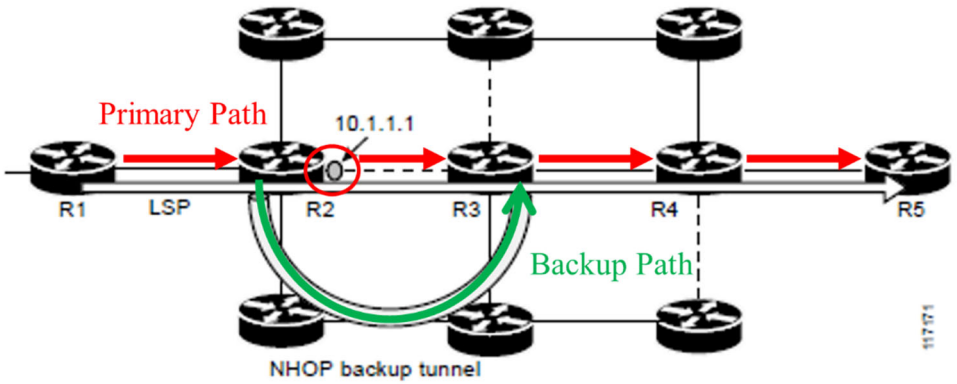
Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 81 of 698

No.	'821 Patent Claim 14	Cisco IOS System
14[c]	digital logic configured to select said working entity and said protection entity from said plurality of transport entity descriptors, comprising:	Cisco IOS System discloses digital logic configured to select said working entity and said protection entity from said plurality of transport entity descriptors, comprising. For example, Cisco IOS Guide discloses that the Cisco IOS System can determine the path that an LSP uses based upon the LSP resource requirements and network resources, such as bandwidth. Cisco IOS Guide further discloses that the MPLS TE SRLG feature tried to avoid using links that are in the same SRLG to lower the odds that both a protected entity and a backup tunnel fail at the same time. Thus, a person of ordinary skill in the art would understand that the Cisco IOS System includes digital logic configured to select said working entity and said protection entity from said plurality of transport entity descriptors. Cisco IOS Guide Page No. 579 ("[MPLS Traffic Engineering Shared Risk Link Groups SRLGs refer to situations where links in a network share a common fiber (or a common physical attribute). If one link fails, other links in the group may fail too. Links in the group have a shared risk. Backup tunnels should avoid using links in the same SRLG as interfaces they are protecting. Otherwise, when the protected link fails the backup tunnel fails too. *Figure 1 Backup Tunnel in the Same SRLG as the Interface It Is Protecting*  Fig. 1 (annotation added) Figure 1 shows a primary label-switched path (LSP) from router R1 to router R5. The LSP protects against the failure of the R2-R3 link at R2 via a backup tunnel to R4. If the R2-R3 link fails, link protection reroutes the LSP along the backup tunnel. However, the R2-R3 link

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 82 of 698

No.	'821 Patent Claim 14	Cisco IOS System
		and one of the backup tunnel links are in the same SRLG. So if the R2-R3 link fails, the backup tunnel may fail too. The MPLS TE SRLG feature enhances backup tunnel path selection so a backup tunnel can avoid using links that are in the same SRLG as the interfaces it is protecting. There are two ways for a backup tunnel to avoid the SRLGs of its protected interface: • The router does not create the backup tunnel unless it avoids SRLGs of the protected interface.]") Cisco IOS Guide at Page No. 580 ("[Open Shortest Path First (OSPF) and Intermediate System-to-Intermediate System (IS-IS) flood the SRLG membership information (including other TE link attributes such as bandwidth availability and affinity) so that all routers in the network have the SRLG information for each link. With this topology information, routers can compute backup tunnel paths that exclude links having SRLGs in common with their protected interfaces. As shown in Figure 2, the backup tunnel avoids the link between R2 and R3, which shares an SRLG with the protected interface.]") *Figure 2 Backup Tunnel That Avoids SRLG of Protected Interface*  Fig. 2 (annotation added) Cisco IOS Guide at Page No. 581 ("[Fast Reroute Protection for MPLS TE SRLGs Fast Reroute (FRR) protects MPLS TE LSPs from link and node failures by locally repairing the LSPs at the point of failure. This protection allows data to continue to flow on LSPs while

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 83 of 698

No.	'821 Patent Claim 14	Cisco IOS System
		their headend routers attempt to establish new end-to-end LSPs to replace them. FRR locally repairs the protectedLSPs by rerouting them over backup tunnels that bypass failed links or nodes. Backup tunnels that bypass only a single link of the LSP's path provide link protection. They protect LSPs if a link along their path fails by rerouting the LSP's traffic to the next hop (bypassing the failed link). These are referred to as next-hop (NHOP) backup tunnels because they terminate at the LSP's next hop beyond the point of failure. Figure 3 illustrates an NHOP backup tunnel.]") *Figure 3* *NHOP Backup Tunnel*  Fig. 3 (annotation added)
14[d]	logic configured to determine a probability of concurrent failure of said working entity and said protection entity;	Cisco IOS System discloses logic configured to determine a probability of concurrent failure of said working entity and said protection entity. *See Supra at* 14[b].

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 84 of 698

No.	'821 Patent Claim 14	Cisco IOS System
14[e]	logic configured to determine an entity cost of said plurality of transport entity descriptors: and	Cisco IOS System discloses determine an entity cost of said plurality of transport entity descriptors. For example, Cisco IOS Guide discloses a headend router to select the best backup tunnel based on the backup bandwidth available and SRLG. A person of ordinary skill in the art would understand that digital logic is configured to make such a selection. Cisco IOS Guide at Page No. 253 ("[Path Option Selection for MPLS TE Tunnel LSPs This section contains the following topics about path option selection for MPLS TE Tunnel LSPs: • Constraint-Based Routing and Path Option Selection, page 5 • Tunnel Reoptimization and Path Option Selection, page 5 • Path Option Selection with Bandwidth Override, page 6 Constraint-Based Routing and Path Option Selection MPLS traffic engineering automatically establishes and maintains LSPs across the backbone by using the Resource Reservation Protocol (RSVP). The path that an LSP uses is determined by the LSP resource requirements and network resources, such as bandwidth. Traffic engineering tunnels are calculated at the LSP head based on a fit between required and available resources (constraint-based routing). Without the Path Option for Bandwidth Override feature, a TE tunnel establishes an LSP based on dynamic or explicit path options in order of preference. However, the bandwidth and other attributes configured on the TE tunnel allow the setup of an LSP only if LSP path options satisfy the constraints. If a path cannot be found that satisfies the configured path options, then the tunnel is not set up.]") Cisco IOS Guide at Page No. 639 ("[Bandwidth Protection NHOP and NNHOP backup tunnels can be used to provide bandwidth protection for rerouted LSPs. This is referred to as backup bandwidth. You can associate backup bandwidth with NHOP or NNHOP backup tunnels. This informs the router of the amount of backup bandwidth a particular backup tunnel can protect. When a router maps LSPs to backup tunnels, bandwidth protection ensures that an LSP uses a given backup tunnel only if there is sufficient backup bandwidth. The router selects which LSPs use which backup tunnels in

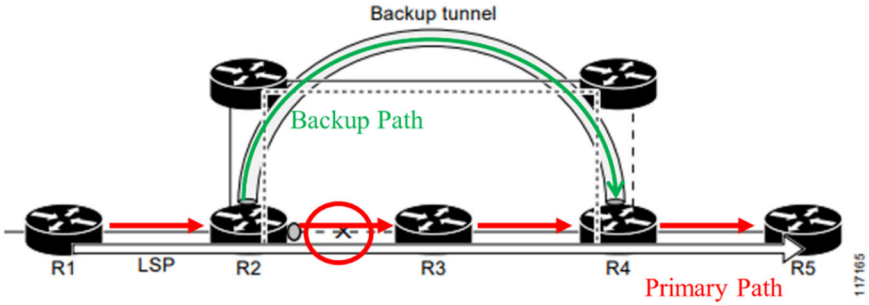
Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 85 of 698

No.	'821 Patent Claim 14	Cisco IOS System
		order to provide maximum bandwidth protection. That is, the router determines the best way to map LSPs onto backup tunnels in order to maximize the number of LSPs that can be protected. For information about mapping tunnels and assigning backup bandwidth, see the "Backup Tunnel Selection Procedure" section on page 10. LSPs that have the "bandwidth protection desired" bit set have a higher right to select backup tunnels that provide bandwidth protection; that is, those LSPs can preempt other LSPs that do not have that bit set. For more information, see the "Prioritizing Which LSPs Obtain Backup Tunnels with Bandwidth Protection" section on page 8.]") Cisco IOS Guide at Page No. 640 ("[Backup Tunnel Support Backup tunnel support has the following capabilities: • Backup Tunnels Can Terminate at the Next-Next Hop to Support FRR, page 6 • Multiple Backup Tunnels Can Protect the Same Interface, page 6 • Backup Tunnels Provide Scalability, page 7.]") Cisco IOS Guide at Page No. 640 ("[Multiple Backup Tunnels Can Protect the Same Interface There is no limit (except memory limitations) to the number of backup tunnels that can protect a given interface. In many topologies, support for node protection requires supporting multiple backup tunnels per protected interface. These backup tunnels can terminate at the same destination or at different destinations. That is, for a given protected interface, you can configure multiple NHOP or NNHOP backup tunnels. This allows redundancy and load balancing.]") Cisco IOS Guide at Page No. 603 ("[Traffic Engineering Tunnels MPLS TE lets you build label switched paths (LSPs) across your network for forwarding traffic. MPLS TE LSPs let the headend of a TE tunnel control the path its traffic takes to a particular destination. This method is more flexible than forwarding traffic based only on a destination address. Interarea tunnels allow you to do the following: • Build TE tunnels between areas (interarea tunnels)

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 86 of 698

No.	'821 Patent Claim 14	Cisco IOS System
		• Build TE tunnels that start and end in the same area, on multiple areas on a router (intra-area tunnels) Some tunnels are more important than others. For example, you may have tunnels carrying VoIP traffic and tunnels carrying data traffic that are competing for the same resources. MPLS TE allows you to have some tunnels preempt others. Each tunnel has a priority, and more-important tunnels take precedence over less-important tunnels.]") Cisco IOS Guide at Page No. 641 ("[Backup Bandwidth Protection Backup bandwidth protection allows you to give LSPs carrying certain kinds of data (such as voice) priority for using backup tunnels.]") Cisco IOS Guide at Page No. 650 ("[For each protected link AB with a maximum reservable subpool value of n, there may be a path from node A to node B such that the difference between the maximum reservable global and the maximum reservable subpool is at least the value of n. If it is possible to find such paths for each link in the network, you can establish all the backup tunnels along such paths without any bandwidth reservations. If there is a single link failure, only one backup tunnel will use any link on its path. Because that path has at least n available bandwidth (in the global pool), assuming that marking and scheduling is configured to classify the subpool traffic into a priority queue, the subpool bandwidth is guaranteed. This approach allows sharing of the global pool bandwidth between backup tunnels protecting independent link failures. The backup tunnels are expected to be used for only a short period of time after a failure (until the headends of affected LSPs reroute those LSPs to other paths with available subpool bandwidth). The probability of multiple unrelated link failures is very small (in the absence of node or shared risk link group (SRLG) failures, which result in multiple link failures). Therefore, it is reasonable to assume that link failures are in practice independent with high probability. This "independent failure assumption" in combination with backup tunnels signaled without explicit bandwidth reservation enables efficient bandwidth sharing that yields substantial bandwidth savings.]") Cisco IOS Guide at Page No. 651 ("[A similar approach can be used for node and SRLG protection. However, the decision of where to put the backup tunnels is more complicated

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 87 of 698

No.	'821 Patent Claim 14	Cisco IOS System
		because both node and SRLG failures effectively result in the simultaneous failure of several links. Therefore, the backup tunnels protecting traffic traversing all affected links cannot be computed independently of each other.]") Cisco IOS Guide at Page No. 579 ("[MPLS Traffic Engineering Shared Risk Link Groups SRLGs refer to situations where links in a network share a common fiber (or a common physical attribute). If one link fails, other links in the group may fail too. Links in the group have a shared risk. Backup tunnels should avoid using links in the same SRLG as interfaces they are protecting. Otherwise, when the protected link fails the backup tunnel fails too. *Figure 1 Backup Tunnel in the Same SRLG as the Interface It Is Protecting*  Fig. 1 (annotation added) Figure 1 shows a primary label-switched path (LSP) from router R1 to router R5. The LSP protects against the failure of the R2-R3 link at R2 via a backup tunnel to R4. If the R2-R3 link fails, link protection reroutes the LSP along the backup tunnel. However, the R2-R3 link and one of the backup tunnel links are in the same SRLG. So if the R2-R3 link fails, the backup tunnel may fail too. The MPLS TE SRLG feature enhances backup tunnel path selection so a backup tunnel can avoid using links that are in the same SRLG as the interfaces it is protecting. There are two ways for a backup tunnel to avoid the SRLGs of its protected interface: • The router does not create the backup tunnel unless it avoids SRLGs of the protected interface.]")

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 88 of 698

No.	'821 Patent Claim 14	Cisco IOS System
		Cisco IOS Guide at Page No. 587 ("[Step 4 show mpls traffic-eng topology Use this command to show the SRLG link membership flooded via the Interior Gateway Protocol (IGP). For example: Router# show mpls traffic-eng topology My_System_id:0000.0000.0003.00 (isis level-2) Signalling error holddown:10 sec Global Link Generation 9 IGP Id:0000.0000.0003.00, MPLS TE Id:10.0.3.1 Router Node (isis level-2) link[0]:Point-to-Point, Nbr IGP Id:0000.0000.0004.00, nbr_node_id:2, gen:9 frag_id 0, Intf Address:10.0.0.33, Nbr Intf Address:10.0.0.34 TE metric:10, IGP metric:10, attribute_flags:0x0 SRLGs:1 2 physical_bw:2488000 (kbps), max_reservable_bw_global:20000 (kbps) max_reservable_bw_sub:5000 (kbps) Global Pool Sub Pool Total Allocated Reservable Reservable BW (kbps) BW (kbps) BW (kbps) --- --- --- bw[0]: 0 20000 5000 bw[1]: 0 20000 5000 bw[2]: 0 20000 5000 bw[3]: 0 20000 5000 bw[4]: 0 20000 5000 bw[5]: 0 20000 5000]") Cisco IOS Guide at Page No. 598 ("[router—A network layer device that uses one or more metrics to determine the optimal path along which network traffic should be forwarded. Routers forward packets from one network to another based on network layer information.]")

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 89 of 698

No.	'821 Patent Claim 14	Cisco IOS System
14[f]	logic configured to reselect said working entity and said protection entity from said plurality of transport entity descriptors upon a reselection event,	Cisco IOS System logic configured to reselect said working entity and said protection entity from said plurality of transport entity descriptors upon a reselection event. *See Supra at* 1[d].
14[g]	wherein said reselection event is selected from a group consisting of adding an entity to said plurality of transport entities, removing an entity from said plurality of transport entities, an operational status change for one of said plurality of transport entities, and a change in overall cost for one of said plurality of transport entities.	Cisco IOS System discloses wherein said reselection event is selected from a group consisting of adding an entity to said plurality of transport entities, removing an entity from said plurality of transport entities, an operational status change for one of said plurality of transport entities, and a change in overall cost for one of said plurality of transport entities. *See Supra at* 1[e].

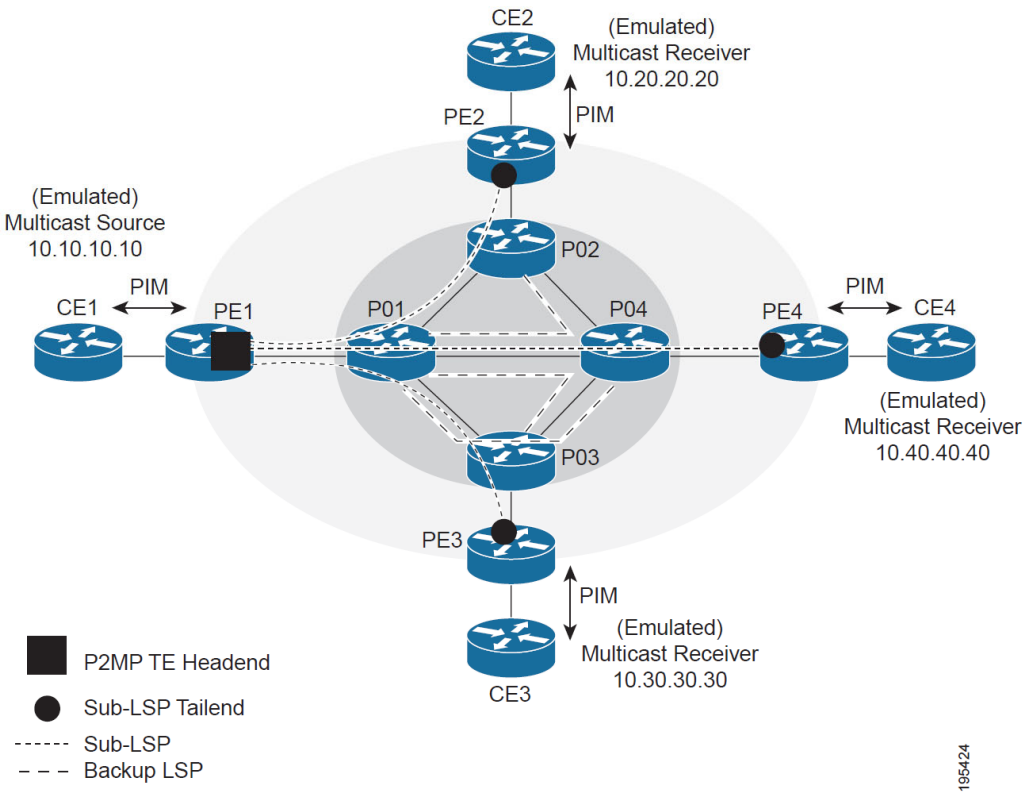
No.	'821 Patent Claim 15	Cisco IOS System
15	The system of claim 14 wherein said entity protection switch comprises a 1:1 switch.	Cisco IOS System discloses the system of claim 14 wherein said entity protection switch comprises a 1:1 switch.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 90 of 698

No.	'821 Patent Claim 15	Cisco IOS System
		For example, Cisco IOS Guide discloses when there is a failure on the protected LSP, the headend router immediately enables the secondary LSP to temporarily carry the tunnel's traffic (1:1 protection). Cisco IOS Guide Page No. 603 ("[Path Protection Path protection provides an end-to-end failure recovery mechanism (that is, full path protection) for MPLS TE tunnels. A secondary LSP is established, in advance, to provide failure protection for the protected LSP that is carrying a tunnel's TE traffic. When there is a failure on the protected LSP, the headend router immediately enables the secondary LSP to temporarily carry the tunnel's traffic. If there is a failure on the secondary LSP, the tunnel no longer has path protection until the failure along the secondary path is cleared.]") Cisco IOS Guide at Page No. 641 ("[One-to-one protection is when a separate backup tunnel must be used for each LSP needing protection. N:1 protection has significant scalability advantages over one-to-one (1:1) protection. An example of 1:1 protection is when 5000 backup tunnels protect 5000 LSPs, each router along the backup path must maintain state for an additional 5000 tunnels.]")

No.	'821 Patent Claim 16	Cisco IOS System
16	The system of claim 14 wherein said entity protection switch comprises a 1+1 switch.	Based upon Orckit's apparent understanding of this limitation set forth in its Preliminary Infringement Disclosure, the Cisco IOS System meets this limitation. To the extent that the Cisco IOS System is found to not meet this limitation, the system of claim 14 wherein said entity protection switch comprises a 1+1 switch would have been obvious to a person having ordinary skill in the art, as explained below. For example, Cisco's IOS Guide discloses: "Figure 2 shows a P2MP TE tunnel that has three destinations. • PE1 is the headend router. • P01 is a branch point router, where packet replication occurs. • PE2, PE3, and PE4 are tailend routers, where the sub-LSP ends.

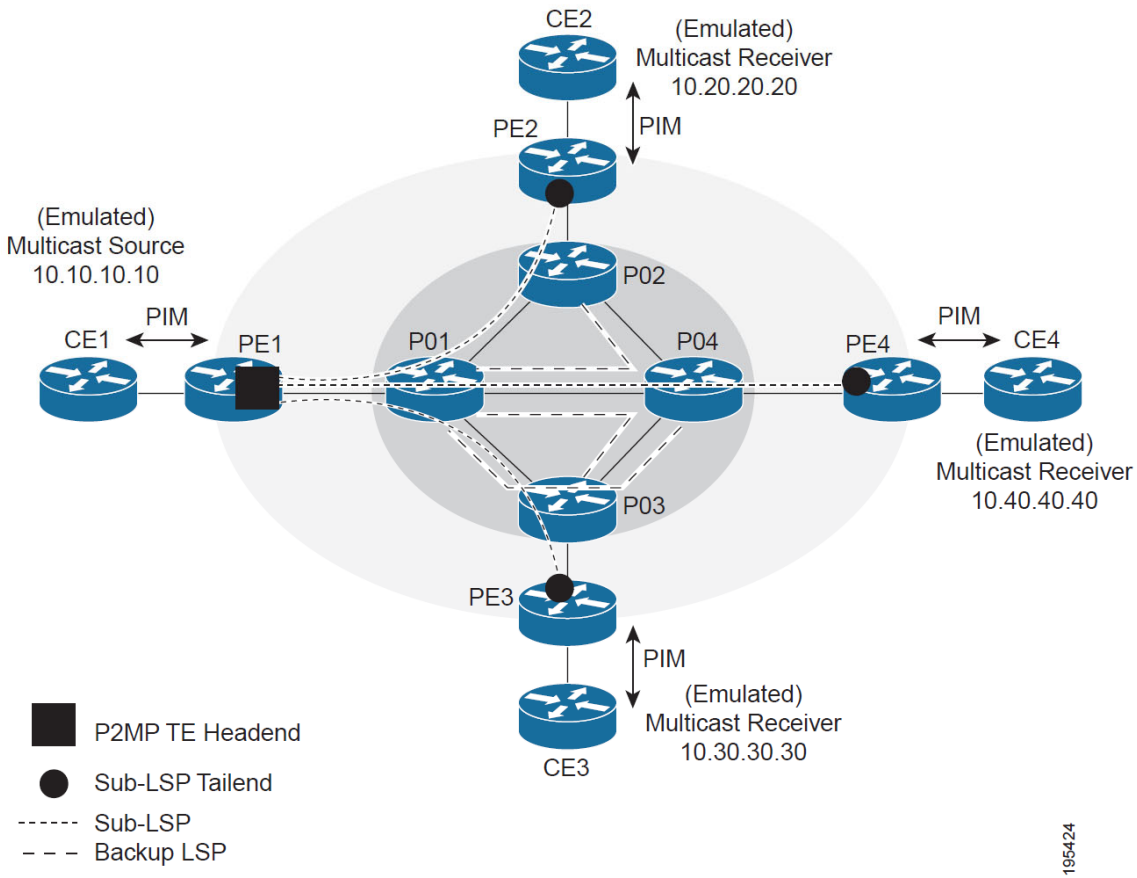
Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 91 of 698

No.	'821 Patent Claim 16	Cisco IOS System
		Between the PE and CE routers, PIM is enabled to exchange multicast routing information with the directly connected customer edge (CE) routers. PIM is not enabled across the P2MP TE tunnel." Cisco IOS Guide at Page No. 468. ***Figure 2 Network Topology with P2MP TE Tunnel***  Cisco IOS Guide at Page No. 469, FIG. 2. "RSVP TE extensions defined in RFC 4875 allow multiple sub-LSPs to be signaled from the headend router. A P2MP TE tunnel consists of multiple sub-LSPs that connect the headend router to various tailend routers.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 92 of 698

No.	'821 Patent Claim 16	Cisco IOS System
		The headend router sends one RSVP path message to each destination. The tailend router replies with a RESV message. The Label Forwarding Information Base (LFIB) is populated using the RSVP labels allocated by the RESV messages. The tailend routers allocate unreserved labels, which are greater than 15 and do not include implicit or explicit null labels. Using unreserved labels allows IP Multicast to perform a Reverse Path Forwarding (RPF) check on the tailend router. Because a sub-LSP tailend router cannot be represented as a regular interface, a special LSP virtual interface (VIF) is automatically created. The LSP VIF represents the originating interface for all IP multicast traffic originating from the P2MP TE tailend router." Cisco IOS Guide at Page No. 469. "You can configure P2P explicit backup tunnels on point of local repair (PLR) nodes for link protection of P2MP sub-LSPs, similar to LSPs for P2P TE tunnels. You can also enable automatic creation of backup tunnels using the Auto-tunnel Backup feature for P2P TE tunnels. All sibling sub-LSPs that share the same outgoing link are protected by the same backup tunnel. All cousin sub-LSPs that share the same outgoing link can be protected by multiple P2P backup tunnels. Link protection for a P2MP TE tunnel is illustrated in Figure 6, which shows PE1 as the tunnel headend router and PE2, PE3, and PE4 as tunnel tailend routers. The following sub-LSPs are signaled from PE1 in the network: • From PE1 to PE2, the sub-LSP travels the following path: PE1 -> P01 -> P02 -> PE2 • From PE1 to PE3, the sub-LSP travels the following path: PE1 -> P01 -> P03 -> PE3 • From PE1 to PE4, the sub-LSP travels the following path: PE1 -> P01 -> P04 -> PE4 Node P01 is a branch node that does packet replication in the MPLS forwarding plane; ingress traffic originating from PE1 will be replicated towards routers P02, P03, and P04." Cisco IOS Guide at Page No. 474-475.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 93 of 698

No.	'821 Patent Claim 16	Cisco IOS System
		Figure 6 P2MP TE Link Protection Example  Cisco IOS Guide at Page No. 475.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 94 of 698

No.	'821 Patent Claim 17	Cisco IOS System
17[preamble]	Non-transitory computer readable media configured to perform a method comprising the steps of:	Cisco IOS System discloses non-transitory computer readable media configured to perform a method comprising the steps of. For example, Cisco IOS Guide is intended for users who configure and maintain Cisco networking devices (such as routers and switches that contain computer memory). A person of ordinary skill in the art would understand that the devices within the Cisco IOS System comprise non-transitory computer readable media. Cisco IOS Guide at Page No. 15 ("[Cisco IOS documentation describes the tasks and commands available to configure and maintain Cisco networking devices.]") Cisco IOS Guide at Page No. 15 ("[The Cisco IOS documentation set is intended for users who configure and maintain Cisco networking devices (such as routers and switches) but who may not be familiar with the configuration and maintenance tasks, the relationship among tasks, or the Cisco IOS commands necessary to perform particular tasks. The Cisco IOS documentation set is also intended for those users experienced with Cisco IOS software who need to know about new features, new configuration options, and new software characteristics in the current Cisco IOS release.]") Cisco IOS Guide at Page No. 1652 ("Capabilities Supported by the MPLS Traffic Engineering MIB The following functionality is supported in the MPLS Traffic Engineering MIB: • The ability to generate and queue notification messages that signal changes in the operational status of MPLS TE tunnels. • Extensions to existing SNMP commands that provide the ability to enable, disable, and configure notification messages for MPLS TE tunnels. • The ability to specify the name or the IP address of a network management station (NMS) in the operating environment to which notification messages are to be sent. • The ability to write notification configurations into nonvolatile memory.")

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 95 of 698

No.	'821 Patent Claim 17	Cisco IOS System
		Cisco IOS Guide at Page No. 1959 ("RP—Route Processor. The processor module in a Cisco 7000 series router that contains the CPU, system software, and most of the memory components that are used in the router. It is sometimes called a supervisory processor.") Cisco IOS Guide at Page No. 640 ("There is no limit (except memory limitations) to the number of backup tunnels that can protect a given interface. In many topologies, support for node protection requires supporting multiple backup tunnels per protected interface. These backup tunnels can terminate at the same destination or at different destinations. That is, for a given protected interface, you can configure multiple NHOP or NNHOP backup tunnels. This allows redundancy and load balancing.")
17[a]	providing a plurality of MPLS transport entities between a first endpoint and a second endpoint;	Cisco IOS System discloses providing a plurality of MPLS transport entities between a first endpoint and a second endpoint. *See Supra at* 1[a].
17[b]	determining an overall cost for each entity pair of said plurality of entities;	Cisco IOS System discloses determining an overall cost for each entity pair of said plurality of entities. *See Supra at* 1[b]
17[c]	selecting an entity pair from said plurality of transport entities based at least in part upon said overall cost; and	Cisco IOS System discloses selecting an entity pair from said plurality of transport entities based at least in part upon said overall cost. *See Supra at* 1[c].
17[d]	if an entity pair reselection event occurs, reselecting said entity pair from the group consisting of said entity pair and a	Cisco IOS System discloses if an entity pair reselection event occurs, reselecting said entity pair from the group consisting of said entity pair and a replacement entity pair comprising at least one entity distinct from the entities of said entity pair. *See Supra at* 1[d].

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 96 of 698

No.	'821 Patent Claim 17	Cisco IOS System
	replacement entity pair comprising at least one entity distinct from the entities of said entity pair,	
17[e]	wherein said entity pair reselection event is selected from a group consisting of adding an entity to said plurality of transport entities, removing an entity from said plurality of transport entities, an operational status change for one of said plurality of transport entities, and a change in overall cost for one of said plurality of transport entities.	Cisco IOS System discloses wherein said entity pair reselection event is selected from a group consisting of adding an entity to said plurality of transport entities, removing an entity from said plurality of transport entities, an operational status change for one of said plurality of transport entities, and a change in overall cost for one of said plurality of transport entities. *See Supra at* 1[e].

No.	'821 Patent Claim 18	Cisco IOS System
18[preamble]	The non-transitory computer readable media of claim 17, wherein said step of selecting an entity pair further comprises:	Cisco IOS System discloses step of selecting an entity pair further comprises. *See Supra at* 1[Preamble] & 1[b].
18[a]	selecting a working entity from said	Cisco IOS System discloses selecting a working entity from said plurality of transport entities.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 97 of 698

No.	'821 Patent Claim 18	Cisco IOS System
	plurality of transport entities;	*See Supra at* 2[a].
18[b]	selecting a protection entity from said plurality of transport entities; and	Cisco IOS System discloses selecting a protection entity from said plurality of transport entities. *See Supra at* 2[b].
18[c]	selecting an active entity from the set consisting of said working entity and said protection entity.	Cisco IOS System discloses selecting an active entity from the set consisting of said working entity and said protection entity. *See Supra at* 3.

No.	'821 Patent Claim 19	Cisco IOS System
19	The non-transitory readable media of claim 18, wherein said step of selecting an entity pair further comprises minimizing an overall cost function.	Cisco IOS System discloses step of selecting an entity pair further comprises minimizing an overall cost function. *See Supra at* 4.

No.	'821 Patent Claim 20	Cisco IOS System
20[preamble]	The non-transitory readable media of claim 19, wherein said overall cost function comprises:	Cisco IOS System discloses said overall cost function comprises. *See Supra at* 5.
20[a]	minimizing a probability of concurrent failure of	Cisco IOS System discloses minimizing a probability of concurrent failure of said protection entity and said working entity.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 98 of 698

No.	'821 Patent Claim 20	Cisco IOS System
	said protection entity and said working entity; and	*See Supra at* 5.
20[b]	a predefined metric selected from the group consisting of interior gateway protocol (IGP) and traffic engineering (TE).	Cisco IOS System discloses said predefined entity cost metric is selected from the group consisting of interior gateway protocol (IGP) and traffic engineering (TE). *See Supra at* 7.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 99 of 698

Chart for U.S. Patent 8,830,821 ("the '821 Patent")
U.S. Patent Publication No. 2004/0205239 to Doshi ("Doshi '239")

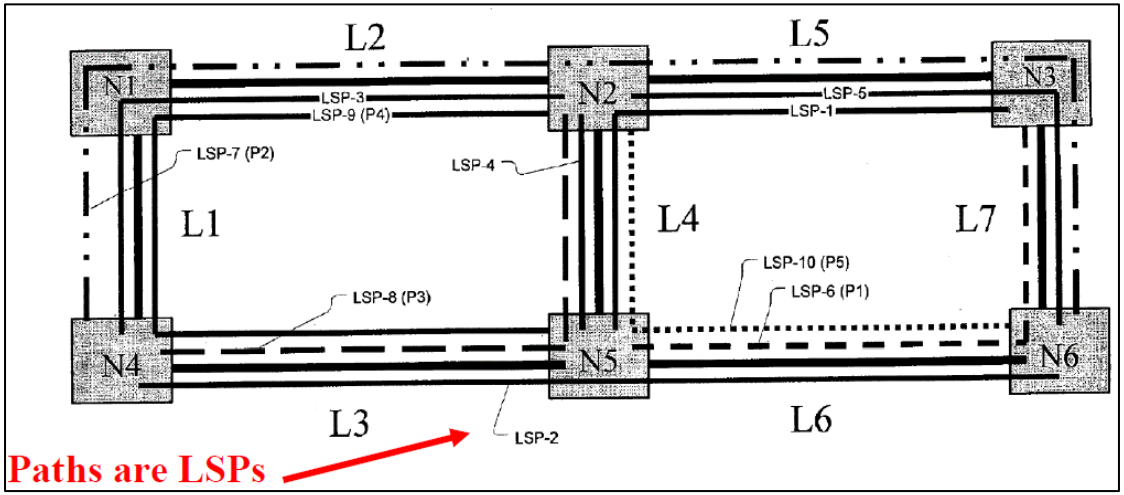
As shown in the chart below, all Asserted Claims of the '821 Patent are invalid under (1) 35 U.S.C. §§ 102 (a), (b), and (g) because Doshi '239 meets each element of those claims, and/or (2) 35 U.S.C. § 103 because Doshi '239 renders those claims obvious either alone, or in combination with the knowledge of a person having ordinary skill in the art, and in further combination with the references specifically identified below and in the following claim chart and/or one or more references identified in Defendant's Preliminary Invalidity Contentions. The following quotations and diagrams come from Doshi '239 titled "Primary/Restoration Path Calculation In Mesh Networks Based On Multiple-Cost Criteria", which was filed on September 26, 2003, and published on October 14, 2004.

Motivations to combine the disclosures in Doshi '239 with disclosures in other publications known in the art, as explained in this chart, include at least the similarity in subject matter between the references to the extent they concern methods of selecting paths in a network and/or determining the cost of potential paths in a network. Insofar as the references cite other patents or publications, or suggest additional changes, one of ordinary skill in the art would look beyond a single reference to other references in the field.

These invalidity contentions are based on Defendant's present understanding of the Asserted Claims, and Orckit's apparent construction of the claims in its November 3, 2022 Disclosure of Asserted Claims and Infringement Contentions Pursuant to P.R. 3-1, and Orckit's January 19, 2023 First Amended Disclosure of Asserted Claims and Infringement Contentions Pursuant to P.R. 3-1 (Orckit's "Infringement Disclosures"), which is deficient at least insofar as it fails to cite any documents or identify accused structures, acts, or materials in the Accused Products with particularity. Defendant does not agree with Orckit's application of the claims, or that the claims satisfy the requirements of 35 U.S.C. § 112. Defendant's contentions herein are not, and should in no way be seen as, admissions or adoptions as to any particular claim scope or construction, or as any admission that any particular element is met by any accused product in any particular way. Defendant objects to any attempt to imply claim construction from this chart. Defendant's prior art invalidity contentions are made in a variety of alternatives and do not represent Defendant's agreement or view as to the meaning, definiteness, written description support for, or enablement of any claim contained therein.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 100 of 698

The following contentions are subject to revision and amendment pursuant to Federal Rule of Civil Procedure 26(e), the Local Rules, and the Orders of record in this matter subject to further investigation and discovery regarding the prior art and the Court's construction of the claims at issue.

No.	'821 Patent Claim 1	Doshi '239
1[preamble]	An entity selection method performed by a network device, comprising the steps of:	Doshi '239 discloses an entity selection method performed by a network device, comprising the steps of. For example, Doshi '239 discloses an MPLS LSP, *i.e.*, "entity," selection method performed by a network manager. Doshi '239 discloses a shared mesh data network "SMDN 100" that is configured to "support[] the multiprotocol label switching (MPLS) architecture standard" of "RFC3031." Doshi '239 ¶0049. Doshi '239's Figure 1, reproduced below, illustrates SMDN 100 that has MPLS LSPs (*e.g.*, LSP-1 to LSP-10) between nodes (*e.g.*, N1-N6). Doshi '239 ¶¶0016, 0052-0054.  Dosh '239, Figure 1 (annotated).

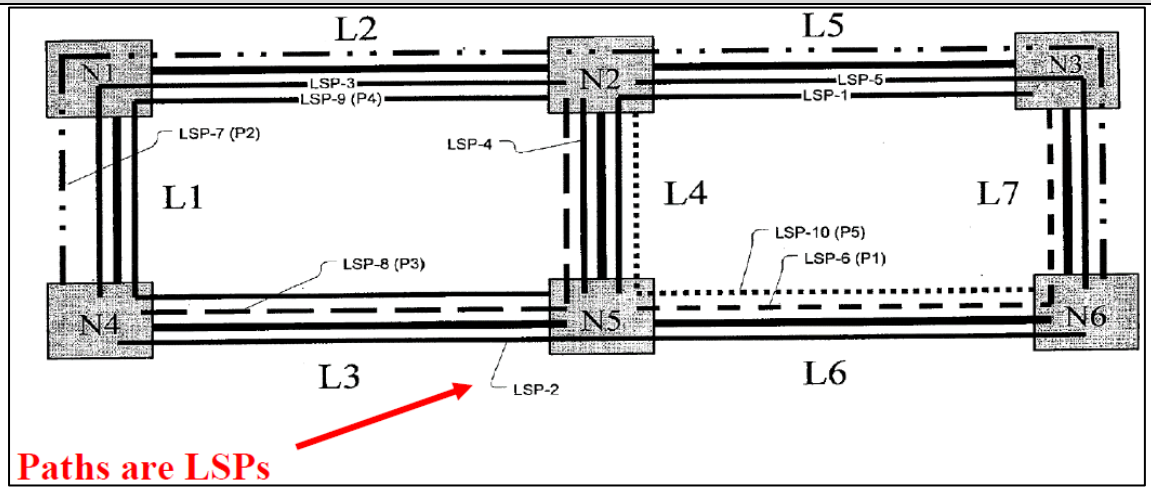
Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 101 of 698

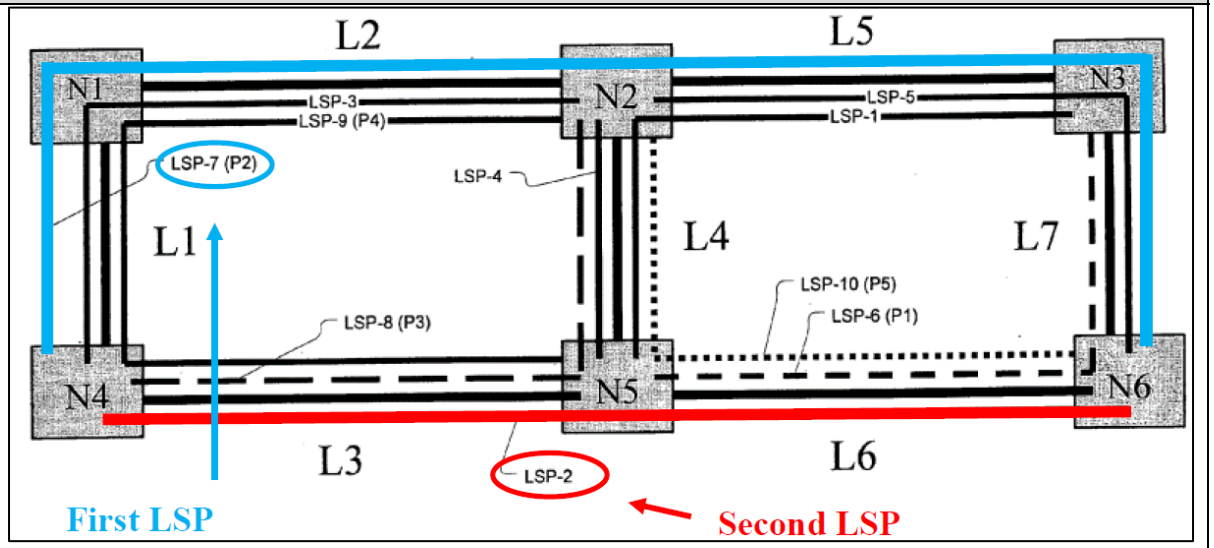
No.	'821 Patent Claim 1	Doshi '239
		In the MPLS context, Doshi '239 teaches a "method …implemented by a network manager for the mesh network." Doshi '239 ¶¶0014; 0049 ("The concepts of the present invention are discussed in the context of …[a] mesh network (*e.g.*,… MPLS)"). The network manager is implemented in a network device, such as a server or distributed partially or fully to nodes. Doshi '239 ¶0271; *see also id*. ¶0082, claim 11. The network manager performs the method, so that "primary and restoration paths … are selected from the plurality of candidate path pairs based on the path cost of each candidate path pair." *Id*. ¶0014; *see also id*., ¶0142 ("In the first flow, illustrated by steps 1102, 1104, 1106, and 1106, one minimum-cost path pair is selected from K candidate pairs by a first method."), claim 11, ¶¶0141–0145, 0156, 0270–0281, Abstract, Figs. 10, 11, 17, 18.
1[a]	providing a plurality of multi protocol label switching (MPLS) transport entities between a first endpoint and a second endpoint;	Doshi '239 discloses providing a plurality of multi protocol label switching (MPLS) transport entities between a first endpoint and a second endpoint. For example, Doshi '239 discloses that its shared mesh data network (SMDN) is configured to support MPLS. Doshi '239 further discloses that the SMDN network is provisioned with a plurality of LSPs, LSP-1 – LSP-10 between various network nodes N-1 – N-6. Based upon Orckit's apparent understanding of this limitation set forth in its Preliminary Infringement Disclosure, Doshi '239 meets this limitation. To the extent Doshi '239 is not found to disclose this limitation, providing a plurality of multiprotocol label switching (MPLS) transport entities between a first endpoint and a second endpoint would have been obvious to a person of ordinary skill in the art as explained below. Doshi '239 discloses "FIG. 1 illustrates an exemplary shared mesh data network (SMDN) according to one embodiment of this invention." Doshi '239 ¶ 0016. Doshi '239 discloses "As illustrated, SMDN 100 has been provisioned with five primary label-switched paths (LSPs) LSP-1 to LSP-5. A label-Switched path is effectively a tunnel between two nodes that carries Service traffic according to a predetermined route. For example, LSP-1 is a tunnel from node N3 to node N5 that follows the path N3-N2-N5. (LSP-1 could alternatively be described by the links it traverses, i.e., L5-L4, or by the nodes and links, i.e., N3-L5-N2-L4-N5; however, for this discussion, in cases where no ambiguity will arise from the node-only notation, that notation will be used.)

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 102 of 698

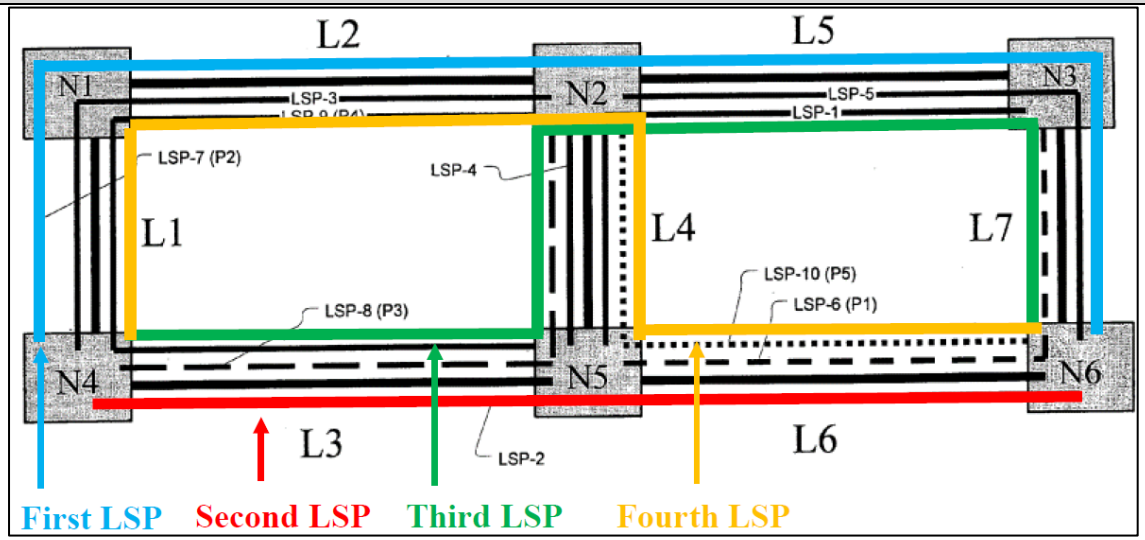
No.	'821 Patent Claim 1	Doshi '239
		An LSP is typically considered to be unidirectional; however, for clarity of illustration, each LSP in FIG. 1 is represented by a bidirectional path that could also represent the corresponding unidirectional LSP in the opposite direction. Further, each LSP may have an associated demand of any reasonable value (e.g., 2Mbps, 10 Mbps, and 100 Mbps); however, without loss of generality, for this example, each of the LSPs in FIG. 1 is assumed to be of unit demand. Extensions to pairs of unidirectional LSPs and variable demands would be understood to one skilled in the art. FIG. 1 illustrates the five primary LSPs as follows: LSP-1 carries traffic along path N3–N2–N5, LSP-2 carries traffic along path N4–N5-N6, LSP-3 carries traffic along path N2-N1-N4, LSP-4 carries traffic along path N2–N5, and LSP-5 carries traffic along path N2–N3–N6. Disjoint Primary and Protection Paths FIG. 1 also illustrates five exemplary protection paths LSP-6 to LSP-10 (also of unit bandwidth in this example) that have been reserved to use for protection of services given a failure of one of the primary paths. LSP-6 (along path N3-N6-N5) protects LSP-1, LSP-7 (along path N4-N1-N2–N3–N6) protects LSP-2, LSP-8 (along path N2–N5-N4) protects LSP-3, LSP-9 (along path N2-N1-N4 N5) protects LSP-4, and LSP-10 (along path N2–N5-N6) protects LSP-5. For clarity, in FIG. 1, the tag for each protection LSP carries a parenthetical Suffix indicating the primary LSP it protects. For example, LSP-6 (P1) indicates the protection role that LSP-6 plays for primary LSP-1." Doshi '239 ¶¶0050-0054.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 103 of 698

No.	'821 Patent Claim 1	Doshi '239
		 Doshi '239, Fig. 1 (annotated). Doshi '239 discloses that the LSPs are provisioned "between two nodes" that correspond to "source and destination nodes" or "end nodes of an LSP." Doshi '239 ¶¶ 0050, 0154–0156, 0267; *see also id.* ¶¶ 0056, 0142, 0188, 0190, 0192, 0195, Fig. 5, Fig. 8, Fig. 11, Fig. 17, Fig. 18, and Fig. 20. Since Doshi '239's LSPs are provided "between" and terminate at two end nodes, a POSITA would have understood that one end node (*e.g.*, source node) corresponds to "a first endpoint" and that the other end node (*e.g.*, destination node) corresponds to "a second end point." Any two (or more) LSPs provided between two end nodes correspond to a plurality of multi protocol label switching (MPLS) transport entities.

No.	'821 Patent Claim 1	Doshi '239
		 Doshi '239, Fig. 1 (annotated with primary LSP (blue) and protection LSP (red)). Although Figure 1 only illustrates two LSPs between N4 and N6, Doshi '239 makes clear that this figure is merely "exemplary" and expressly teaches "two or more paths between" the end nodes. Doshi '239, ¶¶0016, 0048–0055. Doshi '239 also contemplates making "[v]arious modifications" to the exemplary embodiments. *Id.* ¶0300. Doshi '239 contemplates more than two LSPs between two nodes by disclosing "two or more paths between" the end nodes and that there are a "plurality of candidate path pairs," from which to select for communication. *Id.*, Abstract, ¶¶0014, 0048. A person of ordinary skill in the art would understand this to mean that any two nodes would have at least two primary LSPs and two protection LSPs, with different primary and protection LSP combinations corresponding to candidate pairs. *See id.*, ¶0014, ("…each candidate path pair comprising a candidate primary path and a candidate restoration path…") At minimum, a person of ordinary skill in the art would have found it obvious to include two primary LSPs and two protection LSPs in the SMDN network disclosed in Doshi '239. A SMDN with two of each such pairs is illustrated below:

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 105 of 698

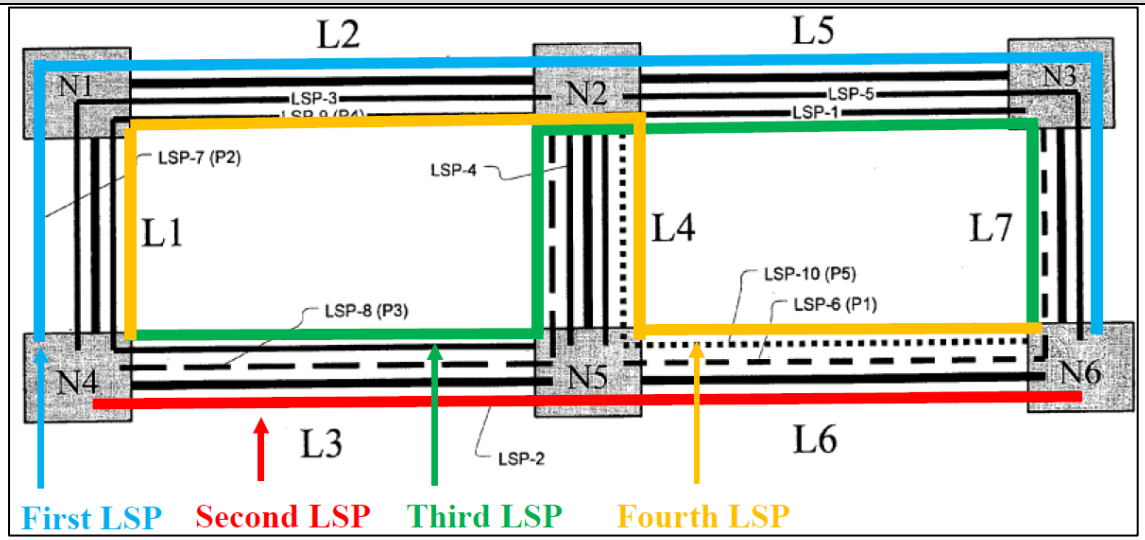
No.	'821 Patent Claim 1	Doshi '239
		 Doshi '239, Fig. 1 (annotated with first primary LSP (blue), first protection LSP (red), second primary LSP (green), and second protection LSP (yellow)). A POSITA would have recognized that providing more than two LSPs between each node pair (*e.g.*, at least two primary LSPs and two protection LSPs) would be beneficial since doing so facilitates path diversity to reduce the probability of service interruption between the nodes. For example, the above illustrated Third and Fourth LSPs between N4 and N6 offers additional protection against failures along different paths, at least in part. A concurrent failure of, for example, L2 and L6 would cause both the First and Second LSP (the original two LSPs illustrated in Doshi '239) to go down; however, such a failure would not bring down the Third LSP. Having the Third LSP available would allow traffic to be switched over with minimal interruption of service as compared to if the First and Second LSPs were the only LSPs connecting N4 and N6. Having the Fourth LSP available in the event of concurrent failure of, for example, L3 and L7 (which would affect both the First and Second LSP without affecting the fourth LSP) would similarly allow traffic to be switched over to the Fourth LSP with minimal interruption of service. Providing more than two LSPs between two end nodes is desirable because it facilitates reoptimization by providing more potential paths to select

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 106 of 698

No.	'821 Patent Claim 1	Doshi '239
		from over time. When reoptimizing, a POSITA would have understood to consider all available LSP paths and choose the optimal LSP pair (*e.g.*, primary and protection LSPs). The desirability of optimizing selection is recognized by Doshi '239, which recognizes the desirability of "optimization." Doshi '239 ¶0172; *see also id.* ¶ 0200 . Another advantage of having is that it allows traffic with different quality of service (QoS) requirements to be allocated and aggregated on different LSPs. *Id.* ¶¶ 0265,0272. For example, traffic with a high QoS requirement, *e.g.*, video traffic, would be assigned to an LSP configured to handle such high QoS traffic while traffic with a lower QoS requirement would be assigned to a different LSP. Doshi '239 in combination with (1) the knowledge of a person of ordinary skill in the art, alone or in further combination with (2) each (individually, as well as one or more together) of the references identified in element 1[a] of Exhibit E-3 renders the claim, including the present limitation, obvious. Below is an example of one such reference. Doshi '239 in combination with Guichard[1] renders this limitation obvious. Specifically, in light of Guichard, a POSITA would have found it obvious to provide additional LSPs between two nodes. That is, it would have been obvious to a person of skill to provide enough LSPs between two nodes to meet bandwidth requirements for communication between the nodes. Doshi '239 itself contemplates more than two LSPs between two nodes by disclosing "two or more paths between" the end nodes and that there are a "plurality of candidate path pairs," from which to select for communication. Doshi '239, Abstract, ¶¶0014, 0048. This would mean that any two nodes would have at least two primary LSPs and two protection LSPs, with different primary and protection LSP combinations corresponding to candidate pairs. *See id.* ¶0014, ("…each candidate path pair comprising a candidate primary path and a candidate restoration path…").

[1] "Definitive MPLS Network Designs," by Guichard et al., Cisco Press, 2005.

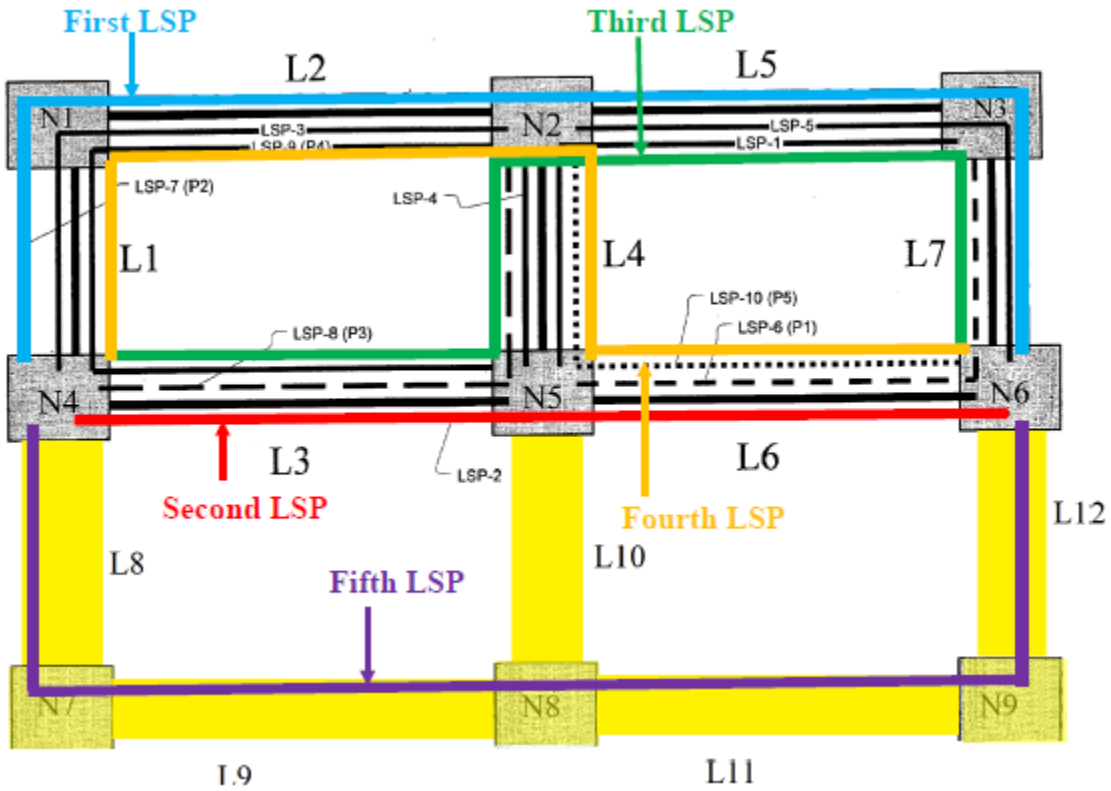
Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 107 of 698

No.	'821 Patent Claim 1	Doshi '239
		 Doshi '239, Fig. 1 (annotated with first primary LSP (blue), first protection LSP (red), second primary LSP (green), and second protection LSP (yellow)). A POSITA would have recognized that providing more than two LSPs between each node pair (e.g., at least two primary LSPs and two protection LSPs) would be beneficial since doing so facilitates path diversity to reduce the probability of service interruption between the nodes. For example, the illustrated Third and Fourth LSPs between N4 and N6 offers additional protection against failures along different paths, at least in part. A concurrent failure of, for example, L2 and L6 would cause both the First and Second LSP (the original two LSPs illustrated in Doshi '239) to go down; however, such a failure would not bring down the Third LSP. Having the Third LSP available would allow traffic to be switched over with minimal interruption of service as compared to if the First and Second LSPs were the only LSPs connecting N4 and N6. Having the Fourth LSP available in the event of concurrent failure of, for example, L3 and L7 (which would affect both the First and Second LSP without affecting the fourth LSP) would similarly allow traffic to be switched over to the Fourth LSP with minimal interruption of service.

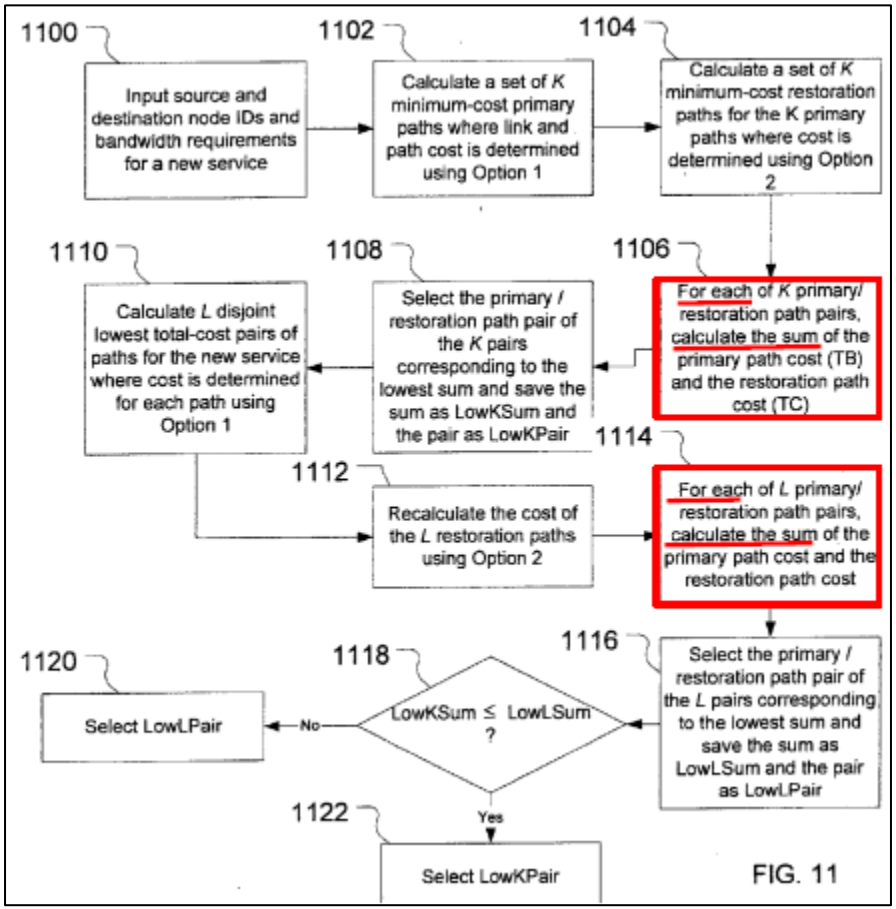
Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 108 of 698

No.	'821 Patent Claim 1	Doshi '239
		Providing more than two LSPs between two end nodes is desirable because it facilitates reoptimization by providing more potential paths to select from over time. It was known in the art that it is highly desirable to reoptimize LSP selection over time. When reoptimizing, a POSITA would have understood to consider all available LSP paths and choose the optimal LSP pair (e.g., primary and protection LSPs). The desirability of optimizing selection is recognized by Doshi '239, which recognizes the desirability of "optimization." Doshi '239 ¶0172; *see also id.* ¶0200. If only one primary LSP and one protection LSP was provided between two end nodes, then reoptimization would not be possible in some failure situations. Thus, facilitating reoptimization is another benefit of providing more than two LSPs between each node pair in Doshi. Another advantage of having is that it allows traffic with different quality of service (QoS) requirements to be allocated and aggregated on different LSPs. *See* Doshi '239 ¶¶0265, 0272. For example, traffic with a high QoS requirement, e.g., video traffic, would be assigned to an LSP configured to handle such high QoS traffic while traffic with a lower QoS requirement would be assigned to a different LSP. Moreover, it would have been understood that additional nodes with corresponding LSPs not illustrated would be present or may be added over time, as both Doshi and Guichard teach. *See, e.g.*, Doshi '239 ¶0055 (Describing that other paths and nodes may exist); Guichard at 79 ("Network state keeps changing. Links and nodes fail and recover…new TE LSPs are set up…"). A POSITA would have been motivated to provide additional LSPs along new links between new nodes in Doshi '239, as the network changes, so that the paths between two given nodes are completely disjointed, thereby further reducing the probability of service interruption. Modified Figure 1 below illustrates additional nodes N7-N9 with new links L8-L12 that provide corresponding a new Fifth LSP between nodes N4 and N6. As illustrated, the Fifth LSP is completely disjoint with respect to the First and Second LSPs:

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 109 of 698

No.	'821 Patent Claim 1	Doshi '239
		 Doshi '239 Fig. 1 (annotated to show Guichard combination). Thus, Doshi '239 in combination with Guichard discloses this limitation.
1[b]	determining an overall cost for each entity pair of said plurality of entities;	Doshi '239 discloses determining an overall cost for each entity pair of said plurality of entities. For example, Doshi '239 calculates the cost of each path (*e.g.*, each primary or protection path) as "the sum of the link costs." Doshi '239 ¶0125, Fig. 10. Then, Doshi '239's path cost is used to determine the cost of each path pair. "For each of a plurality of candidate path pairs…a path cost associated with said each candidate path pair is generated." *Id.* ¶0014; *see also id.*, Abstract, ¶¶ 0026, 0046, 0050–0054, 0076, 0082, 0130, 0124–0181, Fig. 1, Fig. 10,

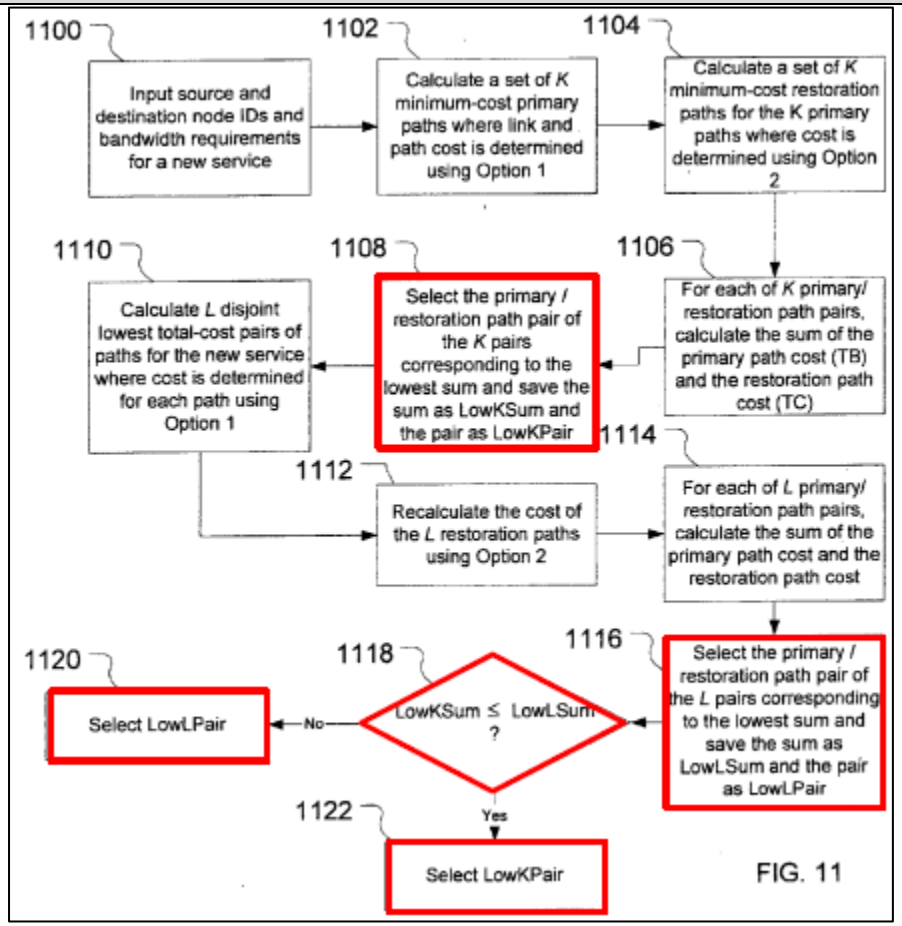
Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 110 of 698

No.	'821 Patent Claim 1	Doshi '239
		Fig. 11, Fig. 12, and Fig. 13. Specifically, Doshi '239's cost determination takes into account whether path pair is "disjoint," as well as other factors, such as the "link utilization, utilization threshold, administrative weight, and sharing degree" and "overall-minimum cost." *Id*. ¶ 0142, *see also id*., ¶¶ 0146–0156, Fig. 11. Thus, at least under the apparent claim scope alleged by Orckit's Infringement Disclosures, this limitation is met.  Doshi '239, Fig. 11 (annotated).

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 111 of 698

No.	'821 Patent Claim 1	Doshi '239
		As explained above in element 1[b] *supra*, Doshi '239 discloses that its SMDN network is implemented as an MPLS network, and thus a person of ordinary skill in the art would understand from the teachings of Doshi '239 that the "overall cost" calculation would be for a MPLS LSP. *See, e.g.*, Doshi '239 at Fig. 1.
1[c]	selecting an entity pair from said plurality of transport entities based at least in part upon said overall cost; and	Doshi '239 discloses selecting an entity pair from said plurality of transport entities based at least in part upon said overall cost. For example, Doshi '239 discloses that "primary and protection paths…are selected from the plurality of candidate path pairs based on the path cost of each candidate path pair." Doshi '239, ¶0014; *see also id.* Abstract, ¶¶ 0046, 0076, 0142–0145, Claim 1, Claim 11. As explained above with respect to Claim 1[a]-1[b], a person of ordinary skill in the art would have understood from the teachings of Doshi '239 / it would have been obvious for the "plurality of candidate path pairs" (pairs of primary and protection LSPs) to correspond to the provided plurality of LSPs of the MPLS architecture. Doshi '239 further discloses that the selection is based on the overall cost, since "the lowest-cost pair…is selected as the overall minimum-cost pair." *Id.* ¶ 0142; *see also id.* ¶¶ 0026, 0125–0139, 0143–0145, Fig. 11. Based upon Orckit's apparent understanding of this limitation set forth in its Preliminary Infringement Disclosure, Doshi '239 meets this limitation. To the extent Doshi '239 is not found to disclose this limitation, determining an overall cost for each entity pair of said plurality of entities would have been obvious to a person of ordinary skill in the art as explained below. *See* Claim 1[a], 1[b].

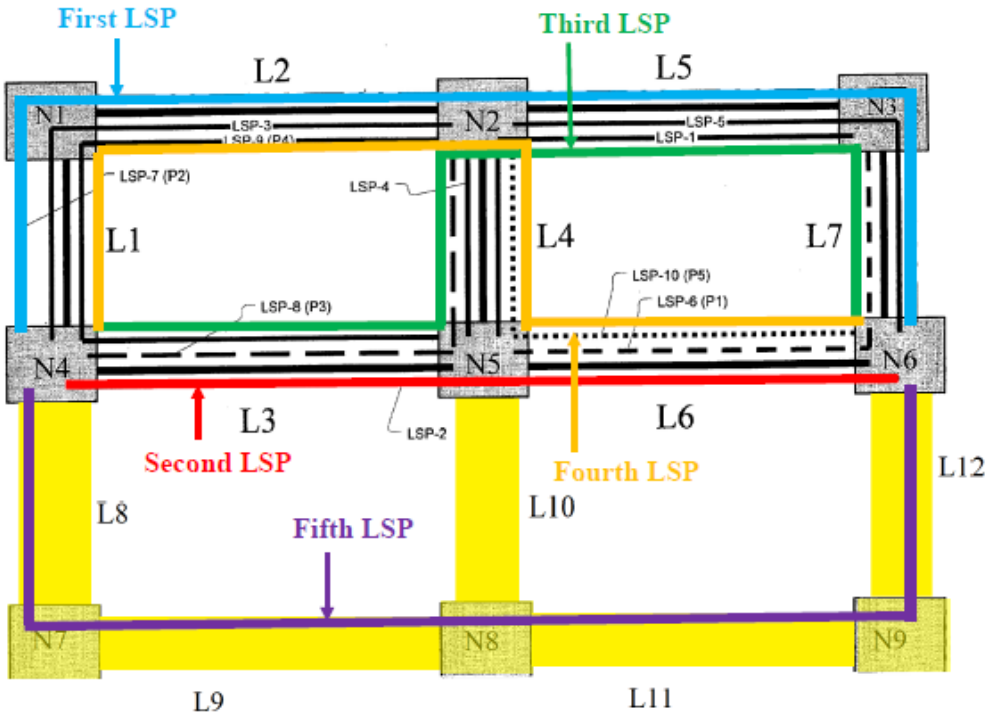
Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 112 of 698

No.	'821 Patent Claim 1	Doshi '239
		 Doshi '239, Fig. 1 (annotated). Doshi '239 in combination with (1) the knowledge of a person of ordinary skill in the art, alone or in further combination with (2) each (individually, as well as one or more together) of the references identified in element 1[c] of Exhibit E-3 renders the claim, including the present limitation, obvious.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 113 of 698

No.	'821 Patent Claim 1	Doshi '239
		Below is an example of one such reference. Doshi '239 in combination with Guichard renders this limitation obvious. It would have been obvious to a POSITA, in view of Guichard, to perform ongoing reoptimization attempts, when implementing Doshi's MPLS architecture, so that an optimum LSP pair (e.g., overall minimum-cost pair) is selected over time for communication between two end nodes. In implementing the combination, a POSITA would have evaluated all existing available LSPs (including previously selected LSPs) by repeating Doshi's candidate pair cost determination (See [1.2]) and selection techniques (discussed above) in response to a manual trigger or expiration of a timer to thereby select the LSP pair with an overall minimum-cost. A POSITA would have been motivated to perform ongoing reoptimization to identify potentially better primary and protection LSP paths (for the LSP pair between two end nodes) that minimize the overall cost. Guichard at 79 ("It is then highly desirable to detect…and reoptimize a TE LSP along a better path when it becomes available."); Doshi '239 ¶0172 ("In general, path selection is a powerful tool that can be used to achieve…optimization."); *id.* ¶0200 ("…making an optimal restoration path choice…") Below is an example of such a selection, per the combined teachings, in the context of Doshi's modified Figure 1.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 114 of 698

No.	'821 Patent Claim 1	Doshi '239
		 Doshi '239 Fig. 1 (annotation added). Take, for instance, where an original LSP pair (e.g., First LSP and Second LSP) was being utilized for communications between end nodes N4 and N6. In response to a manual trigger or expiration of a timer, as Guichard teaches, potential reoptimization of the LSPs would take place to determine if there is a better or more optimal LSP. All things being equal, for example, in the circumstance where the links of the First LSP are heavily utilized as compared to the links of the Fifth LSP, it would have been obvious for Doshi's cost determination to result in a new pair (e.g., the Fifth LSP as the primary and the Second LSP as the protection) as having the overall minimum cost. In such a circumstance, these LSPs would be selected for further communications between end nodes N4 and N6 while the First LSP would be replaced by the Fifth LSP.

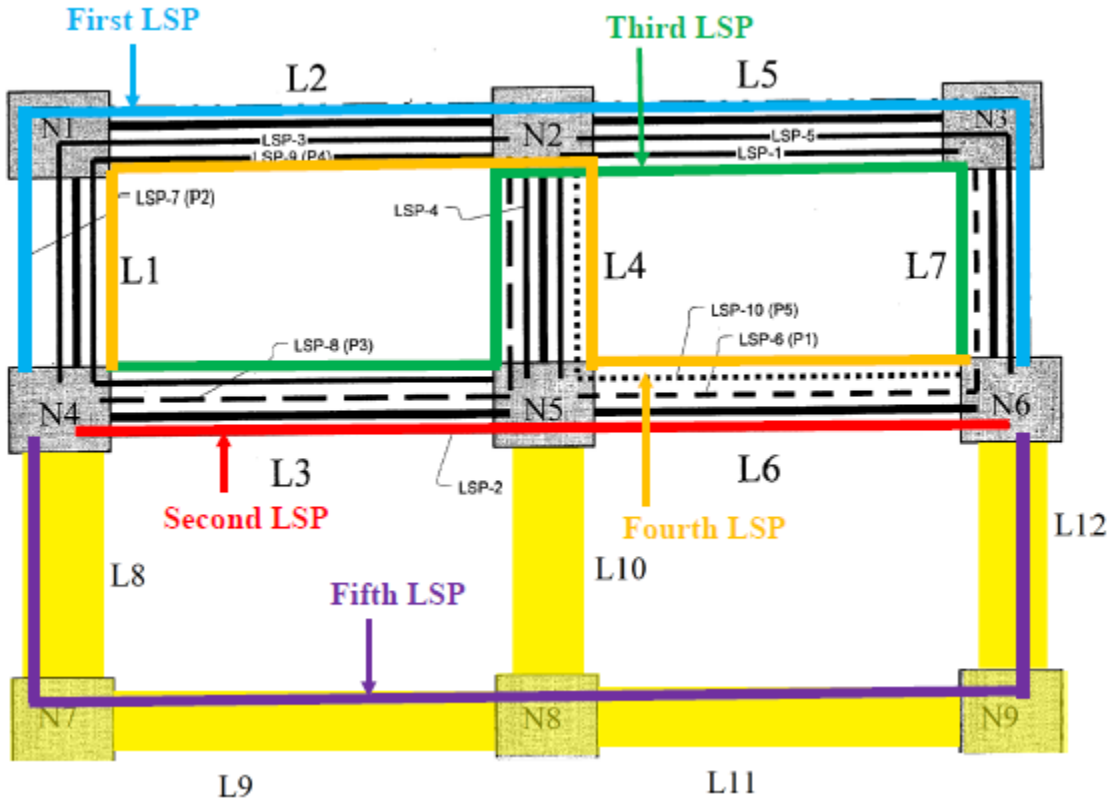
Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 115 of 698

No.	'821 Patent Claim 1	Doshi '239
1[d]	if an entity pair reselection event occurs, reselecting said entity pair from the group consisting of said entity pair and a replacement entity pair comprising at least one entity distinct from the entities of said entity pair,	Doshi '239 discloses if an entity pair reselection event occurs, reselecting said entity pair from the group consisting of said entity pair and a replacement entity pair comprising at least one entity distinct from the entities of said entity pair. For example, Doshi '239 calculates the cost of each path (*e.g.*, each primary or protection path) as "the sum of the link costs." Doshi '239 ¶0125, Fig. 10. Then, Doshi '239's path cost is used to determine the cost of each path pair. "For each of a plurality of candidate path pairs…a path cost associated with said each candidate path pair is generated." *Id*. ¶0014; *see also id.*, Abstract, ¶¶ 0026, 0046, 0050–0054, 0076, 0082, 0130, 0124–0181, Fig. 1, Fig. 10, Fig. 11, Fig. 12, and Fig. 13. Specifically, Doshi '239's cost determination takes into account whether path pair is "disjoint," as well as other factors, such as the "link utilization, utilization threshold, administrative weight, and sharing degree" and "overall-minimum cost." *Id*. ¶ 0142, *see also id*., ¶¶ 0146–0156, Fig. 11. Based upon Orckit's apparent understanding of this limitation set forth in its Preliminary Infringement Disclosure, Doshi '239 meets this limitation. To the extent Doshi '239 is not found to disclose this limitation, if an entity pair reselection event occurs, reselecting said entity pair from the group consisting of said entity pair and a replacement entity pair comprising at least one entity distinct from the entities of said entity pair would have been obvious to a person of ordinary skill in the art as explained below. *See* Claim 1[b]. Doshi '239 in combination with (1) the knowledge of a person of ordinary skill in the art, alone or in further combination with (2) each (individually, as well as one or more together) of the references identified in element 1[d] of Exhibit E-3 renders the claim, including the present limitation, obvious. Below is an example of one such reference. Doshi '239 in combination with Guichard renders this limitation obvious. First, Guichard teaches "Event-driven reoptimization" where "it may be desirable to trigger a reoptimization upon the occurrence of a particular event, such as the restoration of a link in

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 116 of 698

No.	'821 Patent Claim 1	Doshi '239
		the network." Guichard at 81. Additionally, "the evaluation of any potential reoptimization" is triggered "each time a new link is advertised as newly operational." *Id*. at 81; *see also id.* at 79-80 (triggering reoptimizing when links and nodes fail or recover and when new LSPs are set up or torn down and a better path is available). A POSITA would have understood that Guichard generally teaches that the current LSP path selection should be evaluated following a trigger event (e.g., LSP set up, LSP torn down, restoration of a link, or new link) to determine whether a more optimal LSP path is available. Also, a POSITA would have recognized that this evaluation may or may not result in a reoptimization. Guichard at 81 (referring to "evaluation of any potential reoptimization.") A person of skill would have understood that there is only a potential for reoptimization since if the evaluation indicates that the currently utilized path is still the best option, then no reoptimization would occur and the previously utilized path is reselected. Thus, Doshi '239 in combination with Guichard discloses an event triggering candidate pair cost determination and, if the cost determination indicates that the currently used LSP pair still has the overall-minimum cost, reselecting the LSP pair, which renders obvious "if an entity pair reselection event occurs, reselecting said entity pair." Second, as discussed above, Doshi '239's plurality of candidate path pairs include a plurality of LSPs (e.g., at least two primary and at least two protection LSPs) between a pair of end nodes. The currently used primary and protection LSP pair would be initially selected based on the overall minimum cost and the remaining LSPs would be available (with any restored or new LSPs in the network that generate the trigger event, per Guichard) as potential replacements during reoptimization. Thus, in the proposed combination, the reselection is "from the group consisting of said entity pair and a replacement entity pair comprising at least one entity distinct from the entities of said entity pair." Below is an annotated illustration of this combination:

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 117 of 698

No.	'821 Patent Claim 1	Doshi '239
		 Doshi '239 Fig. 1 (annotation added). The selection step selected the Fifth LSP as the primary path and the Second LSP as the protection path. When performing event-driven based reoptimization, per Guichard, an overall cost for each LSP pair permutation would be determined, per Doshi. Doshi '239 ¶0014 ("…a path cost associated with said each candidate path pair is generated…"). If the cost determination indicates that the currently used LSP pair (e.g., Fifth LSP and Second LSP) still has the overall-minimum cost, then the same LSP pair would be reselected. *Id.* ¶ 0124.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 118 of 698

No.	'821 Patent Claim 1	Doshi '239
		Since Doshi's cost determination is performed on each candidate path pair, it would have been obvious to a POSITA for the reselection to be from the group consisting of the currently used LSP pair (e.g., Fifth LSP and Second LSP) and a replacement LSP pair comprising at least one LSP distinct from the currently used LSP pair. In the present example, the reselected LSP pair (e.g., Fifth LSP and Second LSP) would correspond to "said entity pair" while each remaining permutation of LSP pairs not resected would correspond to "a replacement entity pair comprising at least one entity distinct from the entities of said entity pair." Third, to the extent argued that "reselecting said entity pair" requires selecting at least one different LSP as part of the pair instead of reselecting the same two LSPs, the proposed combination also renders that obvious. A POSITA would have found it obvious to select a different LSP pair, when the different LSP pair is determined to have a lower overall-minimum cost than the currently used LSP pair. To illustrate, in the present example, if the First LSP and Second LSP are now determined to have the lowest overall-minimum cost, then that LSP pair would be reselected.
1[e]	wherein said entity pair reselection event is selected from a group consisting of adding an entity to said plurality of transport entities, removing an entity from said plurality of transport entities, an operational status change for one of said plurality of transport entities, and a change in overall cost for one of said plurality of transport entities.	Doshi '239 discloses wherein said entity pair reselection event is selected from a group consisting of adding an entity to said plurality of transport entities, removing an entity from said plurality of transport entities, an operational status change for one of said plurality of transport entities, and a change in overall cost for one of said plurality of transport entities. For example, Doshi '239 calculates the cost of each path (*e.g.*, each primary or protection path) as "the sum of the link costs." Doshi '239 ¶0125, Fig. 10. Then, Doshi '239's path cost is used to determine the cost of each path pair. "For each of a plurality of candidate path pairs…a path cost associated with said each candidate path pair is generated." *Id*. ¶0014; *see also id.*, Abstract, ¶¶ 0026, 0046, 0050–0054, 0076, 0082, 0130, 0124–0181, Fig. 1, Fig. 10, Fig. 11, Fig. 12, and Fig. 13. Specifically, Doshi '239's cost determination takes into account whether path pair is "disjoint," as well as other factors, such as the "link utilization, utilization threshold, administrative weight, and sharing degree" and "overall-minimum cost." *Id*. ¶ 0142, *see also id*., ¶¶ 0146–0156, Fig. 11. Based upon Orckit's apparent understanding of this limitation set forth in its Preliminary Infringement Disclosure, Doshi '239 meets this limitation. To the extent Doshi '239 is not found to disclose this limitation, wherein said entity pair reselection event is selected from a group consisting of adding an entity to said plurality of

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 119 of 698

No.	'821 Patent Claim 1	Doshi '239
		transport entities, removing an entity from said plurality of transport entities, an operational status change for one of said plurality of transport entities, and a change in overall cost for one of said plurality of transport entities would have been obvious to a person of ordinary skill in the art as explained below. Thus, at least under the apparent claim scope alleged by Orckit's Infringement Disclosures, this limitation is met. *See* Claim 1[b]. Doshi '239 in combination with (1) the knowledge of a person of ordinary skill in the art, alone or in further combination with (2) each (individually, as well as one or more together) of the references identified in element 1[d] of Exhibit E-3 renders the claim, including the present limitation, obvious. Below is an example of one such reference. Doshi '239 in combination with Guichard renders this limitation obvious. Guichard teaches that reoptimization evaluation is triggered by various events ("entity pair reselection event"). For example, Guichard's event-driven reoptimization is triggered "if a link is restored in the network" and "each time a new link is advertised as newly operational." Guichard at 81. A POSITA would have understood that restoring a link restores LSPs of the link, which may provide a more optimal path. As discussed above at [1.0]-[1.4], the plurality of MPLS LSPs between end nodes correspond to "said plurality of transport entities." It follows that restoring a link that previously went down to "newly operational," means that the operational status of the link's LSPs has also changed to operational, which discloses "an operational status change for one of said plurality of transport entities." Another network change event contemplated by Guichard is when "new…LSPs are set up" and when "others may be torn down." Guichard at 79. Because these events have the potential for "a more optimal path [to] appear in the network. It is then highly desirable to detect the existence of such a path and reoptimize a TE LSP along a better path when it becomes available." Ex.1006, 79. It follows that a new LSP set up between end nodes would be added

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 120 of 698

No.	'821 Patent Claim 1	Doshi '239
		to the plurality of MPLS LSPs between the end nodes, which corresponds to "adding an entity to said plurality of transport entities." Conversely, tearing down an LSP between the end nodes would remove the LSP from the plurality of MPLS LSPs between the end nodes, which corresponds to "removing an entity from said plurality of transport entities."

No.	'821 Patent Claim 2	Doshi '239
2[preamble]	The method of claim 1, wherein said step of selecting an entity pair further comprises:	Doshi '239 discloses the method of claim 1, wherein said step of selecting an entity pair further comprises. As set forth *supra* in Claim 1, Doshi '239 performs the method of claim 1. As explained in Claim 1[c] *supra*, Doshi '239 discloses a step of selecting an entity pair, *e.g.*, a primary and protection path.
2[a]	selecting a working entity from said entity pair; and	Doshi '239 discloses selecting a working entity from said entity pair. For example, Doshi '239 discloses selecting an LSP pair from the plurality of MPLS LSPs. The selection includes selecting a working entity. Doshi '239 ¶¶0014, 0046; *see also id*. Abstract, ¶¶0142–0145, Claim 1, Claim 11. The '821 Patent describes "switch[ing] between a working entity and a protection entity" after failure. '821 Patent at 1:51-43; 2:35-36. Doshi '239 further discloses "Switching Between Working and Protection LSPs" "after detecting a failure." Doshi '239 ¶¶0266–0267. Doshi's working LSP corresponds to a "working entity" and its protection LSP corresponds to a "protection entity." Since Doshi's working and protection LSPs are part of the selected LSP pair, a POSITA would have understood that the selection of the working and protection LSPs are "from said entity pair."
2[b]	selecting a protection entity from said entity pair.	Doshi '239 discloses selecting a protection entity from said entity pair. For example, Doshi '239 discloses selecting an LSP pair from the plurality of MPLS LSPs. The selection includes selecting a protection entity. Doshi '239 ¶¶0014, 0046; *see also id*. Abstract, ¶¶0142–0145, Claim 1, Claim 11.

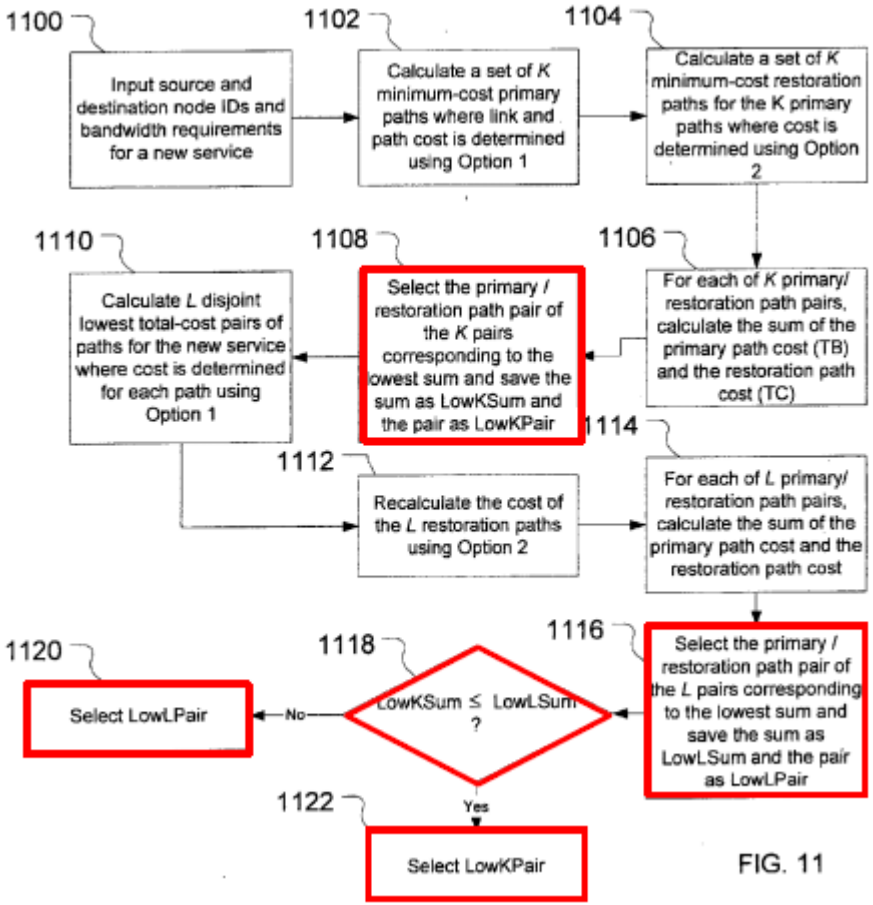
Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 121 of 698

No.	'821 Patent Claim 2	Doshi '239
		The '821 Patent describes "switch[ing] between a working entity and a protection entity" after failure. '821 Patent at 1:51-43; 2:35-36. Doshi '239 further discloses "Switching Between Working and Protection LSPs" "after detecting a failure." Doshi '239 ¶¶0266–0267. Doshi's working LSP corresponds to a "working entity" and its protection LSP corresponds to a "protection entity." Since Doshi's working and protection LSPs are part of the selected LSP pair, a POSITA would have understood that the selection of the working and protection LSPs are "from said entity pair."

No.	'821 Patent Claim 3	Doshi '239
3	The method of claim 2, further comprising the step of selecting an active entity from the set consisting of said working entity and said protection entity.	Doshi '239 discloses the method of claim 2, further comprising the step of selecting an active entity from the set consisting of said working entity and said protection entity. Doshi '239 discloses the method of Claim 2 as explained in Claim 2 *supra*. For example, as explained with respect to Claim 2 *supra*, Doshi '239 describes selecting a working LSP ("working entity") and a protection LSP ("protection entity"). Doshi further describes "Switching Between Working and Protection LSPs" and explains that "after detecting a failure, end nodes of an LSP switch traffic from a primary LSP to its corresponding protection LSP." Doshi '239¶¶0266–0267. Doshi '239's working LSP is the path actively carrying traffic absent a failure while Doshi's protection LSP becomes active when traffic is switched over following a failure. Thus, Doshi's switching decisions, which select which of the working LSP and protection LSP is in use, disclose this limitation.

No.	'821 Patent Claim 4	Doshi '239
4	The method of claim 2, wherein selecting an entity pair further comprises minimizing an overall cost function.	Doshi '239 discloses the method of claim 2, wherein selecting an entity pair further comprises minimizing an overall cost function. Doshi '239 discloses the method of Claim 2 as explained in Claim 2 *supra*. For example, as explained above in Claim 1 *supra*, Doshi '239 calculates the path cost of each path (e.g., each primary and protection LSP) in a candidate path pair and sums the path costs to find the cost of the pair together. The cost calculation considers various factors,

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 122 of 698

No.	'821 Patent Claim 4	Doshi '239
		including link utilization, utilization threshold, administrative weight, sharing degree, and disjointedness, and which represent the "overall cost" for each candidate path pair. The calculation used to determine the "overall cost" corresponds to the claimed "overall cost function." Doshi '239's Figure 11 and accompanying description teach calculating and selecting an "overall minimum-cost" pair. Doshi '239 ¶0142; *see also id.* ¶¶0026, 0125–0139, 0143–0145.  FIG. 11

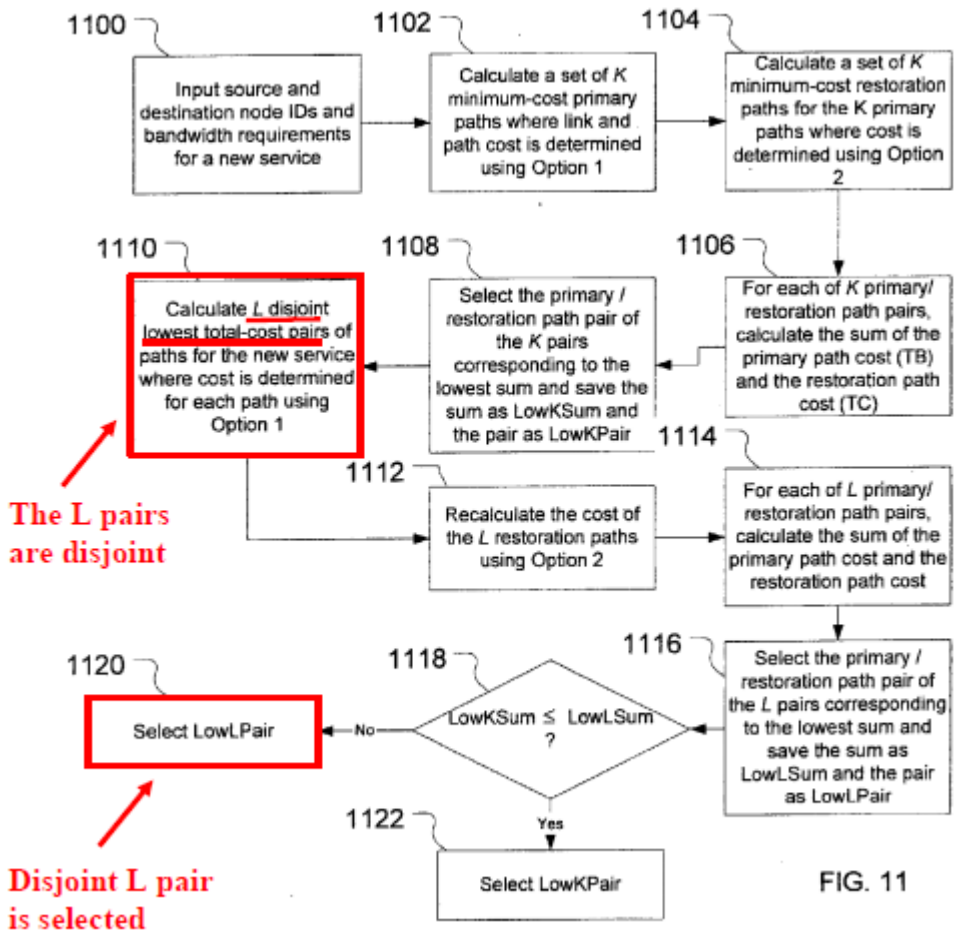
Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 123 of 698

No.	'821 Patent Claim 4	Doshi '239
		Fig. 11 (annotation added). Thus, at least under the apparent claim scope alleged by Orckit's Infringement Disclosures, this limitation is met. A POSITA would have understood that the selected "overall minimum-cost" pair is the LSP pair having the lowest overall cost function in consideration of various factors, including link utilization, utilization threshold, administrative weight, sharing degree, and disjointedness. Thus, Doshi '239 discloses this limitation.

No.	'821 Patent Claim 5	Doshi '239
5	The method of claim 4, wherein said overall cost function comprises substantially minimizing a probability of concurrent failure of said protection entity and said working entity.	Doshi '239 discloses the method of claim 4, wherein said overall cost function comprises substantially minimizing a probability of concurrent failure of said protection entity and said working entity. Doshi '239 discloses the method of Claim 4 as explained in Claim 4 *supra*. As discussed in Claim 1 *supra*, Doshi '239 calculates the path cost of each primary "working entity" and protection LSP "protection entity" in a candidate path pair and sums the path costs to find the cost of the pair together. The cost calculation considers a myriad of factors such as link utilization, sharing degree, and path disjointedness and which represent the "overall cost" for each candidate path pair. The calculation used to determine the "overall cost" corresponds to the claimed "overall cost function." Doshi '239 describes selecting primary LSP and protection LSPs that are "strictly disjoint" from each other, having "no common links or nodes other than their common ingress and egress nodes." Doshi '239 ¶¶0053,0148-0150.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 124 of 698

No.	'821 Patent Claim 5	Doshi '239



Ex.1005, Fig. 11 (annotated); Ex.1003, ¶148.

Doshi '239 Fig. 11 (annotation added).

When two paths are strictly disjoint, "then a failure affecting one of them will not affect the other." Doshi '239 ¶0040; *see also id.* ¶0155 (Failure "protection is made possible by selecting…disjoint [paths]."); ¶¶0056, 0062, 0080, 0131, 0142–0156, 0287, Fig. 1. Doshi considers path disjointedness in its cost calculation and selecting a pair of primary and

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 125 of 698

No.	'821 Patent Claim 5	Doshi '239
		protection paths that are node and links disjoint such that the failure of one does not affect the other. Thus, at least under the apparent claim scope alleged by Orckit's Infringement Disclosures, this limitation is met.

No.	'821 Patent Claim 6	Doshi '239
6	The method of claim 4, wherein said overall cost function comprises a predefined entity cost metric.	Doshi '239 discloses the method of claim 4, wherein said overall cost function comprises a predefined entity cost metric. Doshi '239 discloses the method of Claim 4 as explained in Claim 4 *supra*. As discussed at Claim 1 *supra*, Doshi '239 calculates the path cost of each path (e.g., each primary and protection LSP) in a candidate path pair and sums the path costs to find the cost of the pair together. The cost calculation considers various factors, including link utilization, utilization threshold, administrative weight, sharing degree, and disjointedness—which represent the "overall cost" for each candidate path pair. The calculation used to determine the "overall cost" corresponds to the claimed "overall cost function." Each of these factors, considered in the cost calculation, is predefined in the sense that the cost calculation is pre-configured to consider each factor. For example, before a calculation is performed, it is understood that the calculation will consider link utilization, sharing degree, and path disjointedness. Additionally, the "utilization threshold" (e.g., nominally set to 80%) and "administrative weight" (e.g., a coefficient for scaling) are predefined metrics. Doshi '239 ¶0128, 0141. Each of these metrics has an impact on the cost of the path being evaluated. *Id*. ¶0142. Thus, each of link utilization, utilization threshold, administrative weight, sharing degree, and disjointedness correspond to "a predefined entity cost metric." Additionally, traffic engineering metrics correspond to "predefined entity cost metric" in the '821 patent. Ex.1001, 5:49-51. A POSITA would have understood that Doshi '239's link utilization, utilization threshold, and administrative weight each correspond to traffic engineering data—it was well known that traffic engineering data includes such metrics.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 126 of 698

No.	'821 Patent Claim 6	Doshi '239
		Doshi '239's disclosure is in the context of "multi-protocol label switched (MPLS) traffic engineering," which would have further confirmed a POSITA's understanding that Doshi's link utilization, utilization threshold, and administrative weight each correspond to traffic engineering metrics. Doshi '239 ¶¶0298, 0010, 0093, 0191, 0257.

No.	'821 Patent Claim 7	Doshi '239
7	The method of claim 6, wherein said predefined entity cost metric is selected from the group consisting of interior gateway protocol (IGP) and traffic engineering (TE).	Doshi '239 discloses the method of claim 6, wherein said predefined entity cost metric is selected from the group consisting of interior gateway protocol (IGP) and traffic engineering (TE). *See* Claim 6.

No.	'821 Patent Claim 8	Doshi '239
8	The method of claim 4, further comprising the step of configuring said working entity as revertive.	Doshi '239 discloses the method of claim 4, further comprising the step of configuring said working entity as revertive. Doshi '239 further discloses this limitation as explained in Claim 1 *supra*. Doshi '239 in combination with (1) the knowledge of a person of ordinary skill in the art, alone or in further combination with (2) each (individually, as well as one or more together) of the references identified in Claim 8 of Exhibit E-3 renders the claim, including the present limitation, obvious. Below is an example of one such combination. Doshi '239 in combination with Guichard and Huang[2] renders this limitation obvious.

[2] U.S. Patent Publication No. 2003/0117950 to Huang.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 127 of 698

No.	'821 Patent Claim 8	Doshi '239
		Like the '821 patent and Doshi '239, Huang describes LSP selection within an MPLS network. Huang ¶¶0004, 0032, 0035, 0040. Huang's LSP selection is described in the context of "protection of connections formed through a mesh-type communication network." *Id.* ¶0001. In Huang, a request is received to set up an LSP segment between a head end node and a tail end node, where the request further identifies a backup route to the tail end node. *Id.* ¶0016. The backup routes are LSPs that "protect each of the working links" if failure occurs. *Id.* ¶0032. Huang's "working links" are also LSPs. *See id.* ¶¶0030–0031. Huang further describes "revertive switching," where "the connection will be switched back from the backup LSP to the working link once the working link has cleared the failure that caused the switchover." *Id.* ¶0053. A POSITA would have been motivated to combine the teachings of Doshi '239 and Huang for multiple reasons First, Doshi '239 and Huang are analogous art to the '821 patent. For example, just like the '821 patent, both references generally pertain to MPLS networks. '821 Patent at 3:15-18, Abstract; Doshi '239 ¶0049; Huang ¶¶0004–0007. Additionally, both references address the problem of selecting a pair of MPLS paths from available paths. '821 Patent at 4:63-5:8; Doshi '239 ¶0014; Huang ¶¶ 0004, 0032, 0035, 0040, 0048. Thus, given the similarities between the references and the fact that they address the problem of selecting LSPs, a POSITA considering Doshi would have naturally considered the teachings of Huang. Second, a POSITA would have been motivated to combine the teachings of Doshi and Huang to produce numerous predictable and beneficial results. Doshi '239 teaches "Switching Between Working and Protection LSPs" and explains that "after detecting a failure, end nodes of an LSP switch traffic from a primary LSP to its corresponding protection LSP." Doshi '239 ¶¶0266–0267. Huang complements Doshi '239 by teaching what to do when the failure affecting the working LSP has been cleared. Huang teaches "revertive switching" where "the connection will be switched back from the backup LSP to the working link once the working link has cleared the failure that caused the switchover." Huang ¶0053.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 128 of 698

No.	'821 Patent Claim 8	Doshi '239
		A POSITA would have been motivated to configure Doshi '239's working LSP as revertive, in view of Huang, to avoid the need to calculate and select a replacement working LSP, when the previously selected working LSP has been quickly restored. Reverting allows for the previously selected working LSP (which as discussed at Claim 1 *supra* is part of the overall minimum-cost pair) to be utilized without any additional computations or selections. A POSITA would have recognized that in cases of quick restoration, e.g., when a mistaken failure diagnosis is quickly resolved, it would be simpler to revert back to the working LSP instead of having the network manager calculate the costs for all LSPs to ultimately reach the same result of selecting the working LSP. Thus, Huang's teachings obviate unnecessary efforts. The combination is merely the use of a known technique (Huang's technique of configuring the working LSP as revertive) to improve a similar method (Doshi '239's method for selecting working and protection LSPs) in the same way with predictable results (reverting back to Doshi's working LSP once the failure has been cleared without performing needless calculations). The results would have been predictable and there would have been a reasonable expectation of success in the combination. A POSITA would have had a reasonable expectation of success because Huang provides implementation details, such as waiting to revert back until the working LSP has been restored. Huang ¶0053. Also, the expectation of success is evidenced by the fact that revertive switching was described in multiple RFCs discussing MPLS operation. Implementing the combination would have been within a POSITA's skillset since Doshi '239 already sets forth techniques for switching between working and protection LSPs. Doshi '239 ¶¶0266–0268. Building upon Doshi '239 to configure working LSPs as revertive, per Huang, would utilize already existing switching techniques. Doshi acknowledges that a skilled artisan would have had the necessary skill set to make "[v]arious modifications of the described embodiments." Doshi '239 ¶0300. It would have been obvious for such modifications to include reverting back to the working LSP after a failure of the working LSP has been cleared. The combination of Doshi '239 and Huang renders claim 8 obvious.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 129 of 698

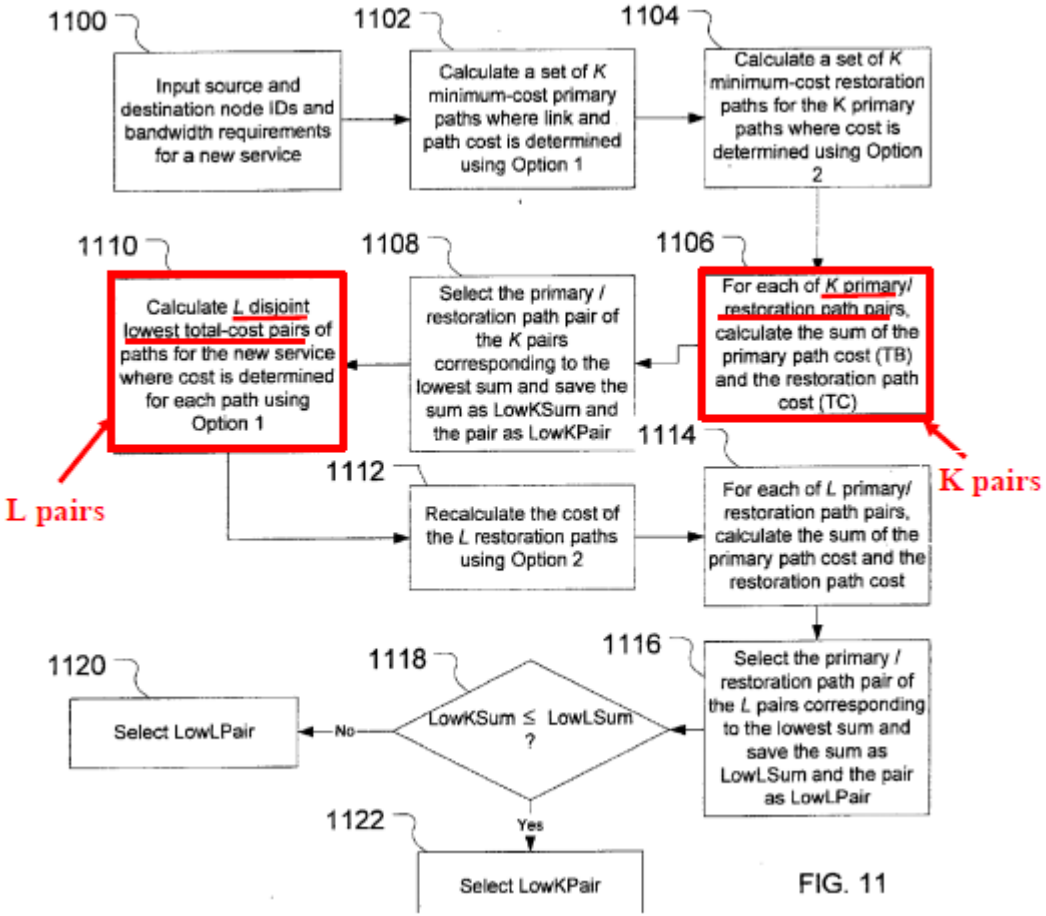
No.	'821 Patent Claim 8	Doshi '239
		First, Doshi '239 describes "Switching Between Working and Protection LSPs" and explains that "after detecting a failure, end nodes of an LSP switch traffic from a primary LSP to its corresponding protection LSP." Doshi '239 ¶¶0266–0267. Doshi '239 also incorporates by reference in its entirety RFC 3209 which describes returning to an "original path when the failed resource becomes re-activated." Since Doshi describes switching from a working LSP to a protection LSP after a failure, and since Doshi incorporates RFC 3209's disclosure of returning to an original path that is re-activated after a failure, it would have been obvious to a POSITA to configure the working LSPs as revertive. Thus, Doshi renders obvious "configuring said working entity as revertive." Second, like RFC 3209, Huang also teaches "that, in the event of a failure" of a working link, traffic "may be switched to corresponding individual backup LSPs" (e.g., a "protection LSP" in terms of Doshi). Huang ¶0032. The "working links" are LSPs. Huang ¶¶0030–0031. Huang also discloses "revertive switching," such that the "the connection will be switched back from the backup LSP to the working link once the working link has cleared the failure that caused the switchover," which corresponds to "configuring said working entity as revertive." Huang ¶0053. POSITA would have found it obvious to utilize revertive switching in Doshi '239, as taught by Huang, such that the switch configures the working LSP to be used again once the fault is cleared. A POSITA would have recognized that in some instances, e.g., in the case of a quickly resolved mistaken failure diagnosis, it would be simpler to revert back to the working LSP instead of calculating the costs for all LSPs only to again select the working LSP.

No.	'821 Patent Claim 9	Doshi '239
9[preamble]	The method of claim 4, wherein said overall cost function comprises:	Doshi '239 discloses [-]. For example, Doshi '239 discloses the method of claim 4, wherein said overall cost function comprises. *See* Claim 4.

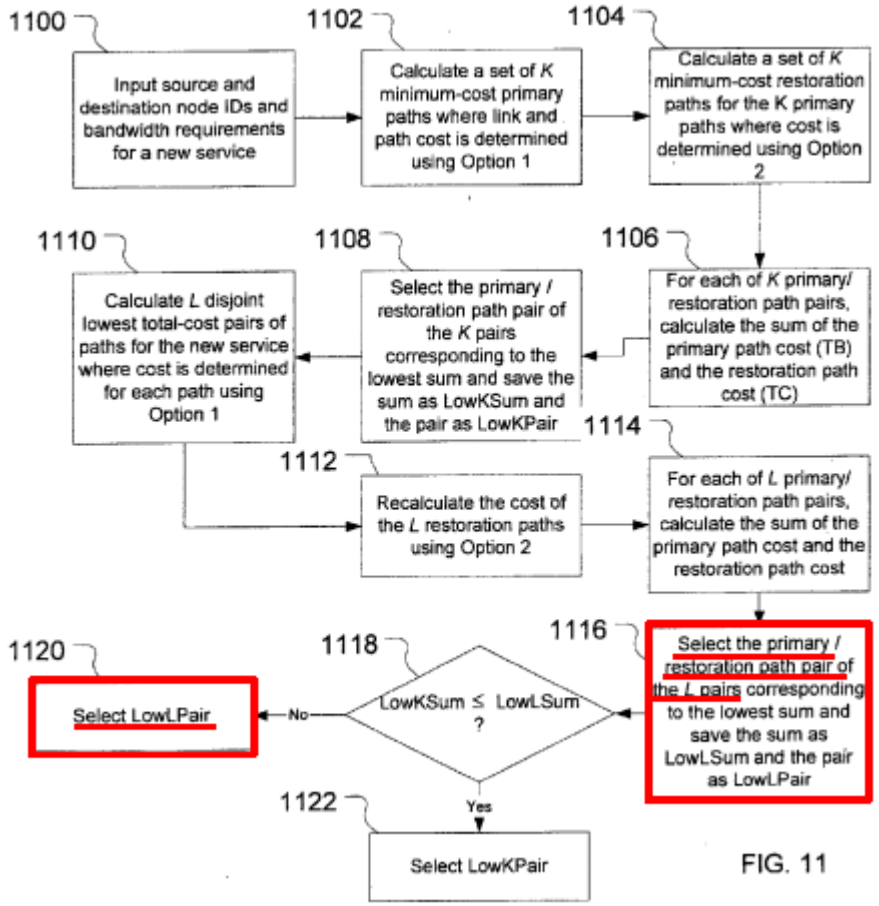
Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 130 of 698

No.	'821 Patent Claim 9	Doshi '239
9[a]	selecting a subset of entity pairs wherein each entity pair of said subset has substantially minimum probability of a concurrent failure of said protection entity and said working entity; and	Doshi '239 discloses selecting a subset of entity pairs wherein each entity pair of said subset has substantially minimum probability of a concurrent failure of said protection entity and said working entity. For example, as explained in claim 1 *supra*, Doshi '239 calculates the path cost of each path (e.g., each primary and protection LSP) in a candidate path pair and sums the path costs to find the cost of the pair together. The cost calculation considers various factors, including link utilization, utilization threshold, administrative weight, sharing degree, and disjointedness, and which represent the "overall cost" for each candidate path pair. The calculation used to determine the "overall cost" corresponds to the claimed "overall cost function." As shown at Doshi '239's Figure 11 and disclosed in the accompanying description, the cost calculation includes calculating and comparing two sets of candidate path pairs—"K candidate pairs" and "L candidate pairs," where the L pairs are disjoint pairs. Doshi '239 ¶¶ 0142–0145.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 131 of 698

No.	'821 Patent Claim 9	Doshi '239
		 Doshi '239 Fig. 11 (annotation added). The L candidate pairs and K candidate pairs are both subsets of the total number of candidate pairs available. In the case of the L candidate pairs, the subset consists of pairs that are disjoint. As discussed above, the pairs include both a working path (e.g., working LSP) and a protection path (e.g., protection LSP) that are disjoint from each other, which substantially minimizes the probability of concurrent failure. Thus, Doshi's L candidate pairs correspond

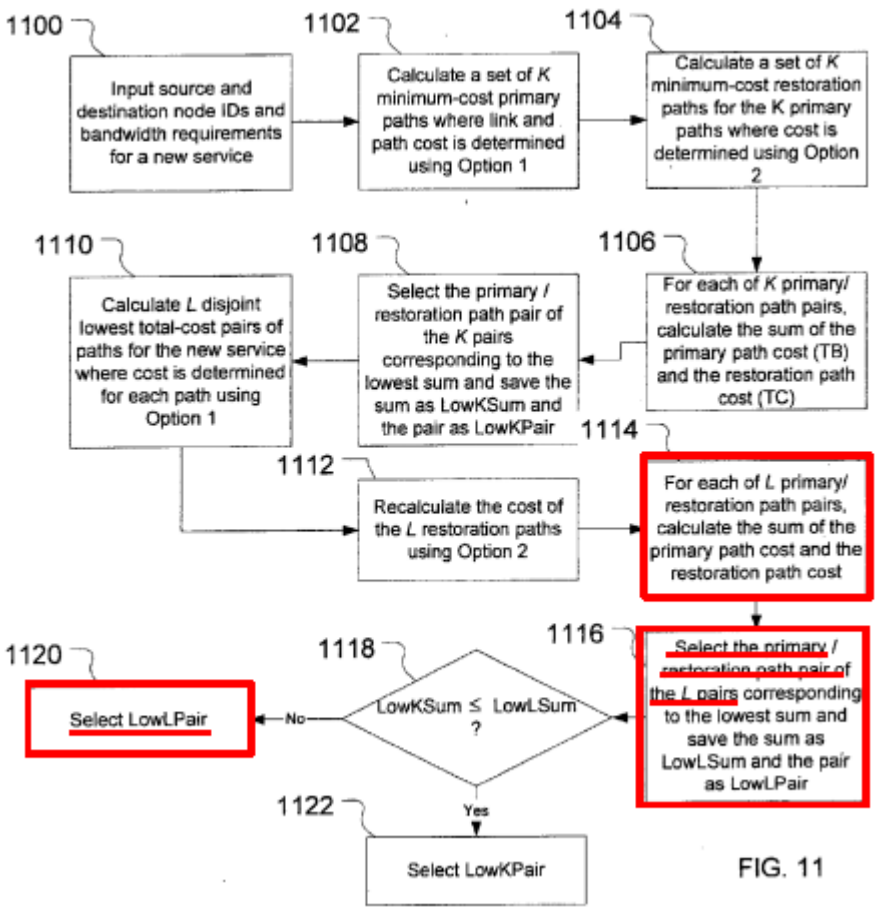
Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 132 of 698

No.	'821 Patent Claim 9	Doshi '239
		to "a subset of entity pairs." A POSITA would have understood selecting the set of L candidate pairs for calculations corresponds to "selecting" as claimed.  Doshi '239 Fig. 11 (annotation added).
9[b]	if said subset comprises at least two entity pairs, selecting	Doshi '239 discloses if said subset comprises at least two entity pairs, selecting an entity pair from said subset that minimizes an entity cost function.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 133 of 698

No.	’821 Patent Claim 9	Doshi ’239
	an entity pair from said subset that minimizes an entity cost function.	As discussed at Claim 9[a] above, Doshi ’239 calculates L “disjoint lowest total-cost pairs,” plural. Thus, Doshi ’239’s L candidate pairs include at least two pairs. Doshi ’239 goes on to explain that the L candidate pairs are assessed for cost, which as discussed at [1.2] includes assessing the cost of each path in the pair. As discussed above, Doshi ’239 calculates the path cost of each path (e.g., each primary and protection LSP) in a candidate path pair. The cost calculation considers various factors, including link utilization, utilization threshold, administrative weight, sharing degree, and disjointedness. Doshi ’230 ¶¶0124–0141. The calculation used to determine the path cost corresponds to the claimed “entity cost function.” Doshi ’239 ¶¶0144–0145, Fig. 10. Doshi ’239 describes selecting the lowest cost pair of the L candidate pairs as the lowest cost pair out of all pairs, including K candidate pairs. Doshi ’239 ¶¶0142–0145.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 134 of 698

No.	'821 Patent Claim 9	Doshi '239
		 Doshi '239 Fig. 11 (annotation added).

No.	'821 Patent Claim 10	Doshi '239
10	The method of claim 9, wherein said entity cost function	Doshi '239 discloses the method of claim 9, wherein said entity cost function comprises a predefined metric.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 135 of 698

No.	'821 Patent Claim 10	Doshi '239
	comprises a predefined metric.	For example, Doshi '239 *See* Claim 6.

No.	'821 Patent Claim 11	Doshi '239
11	The method of claim 10, wherein said predefined metric is selected from the group consisting of interior gateway protocol (IGP) and traffic engineering (TE).	Doshi '239 discloses the method of claim 10, wherein said predefined metric is selected from the group consisting of interior gateway protocol (IGP) and traffic engineering (TE). For example, Doshi '239 *See* Claim 7.

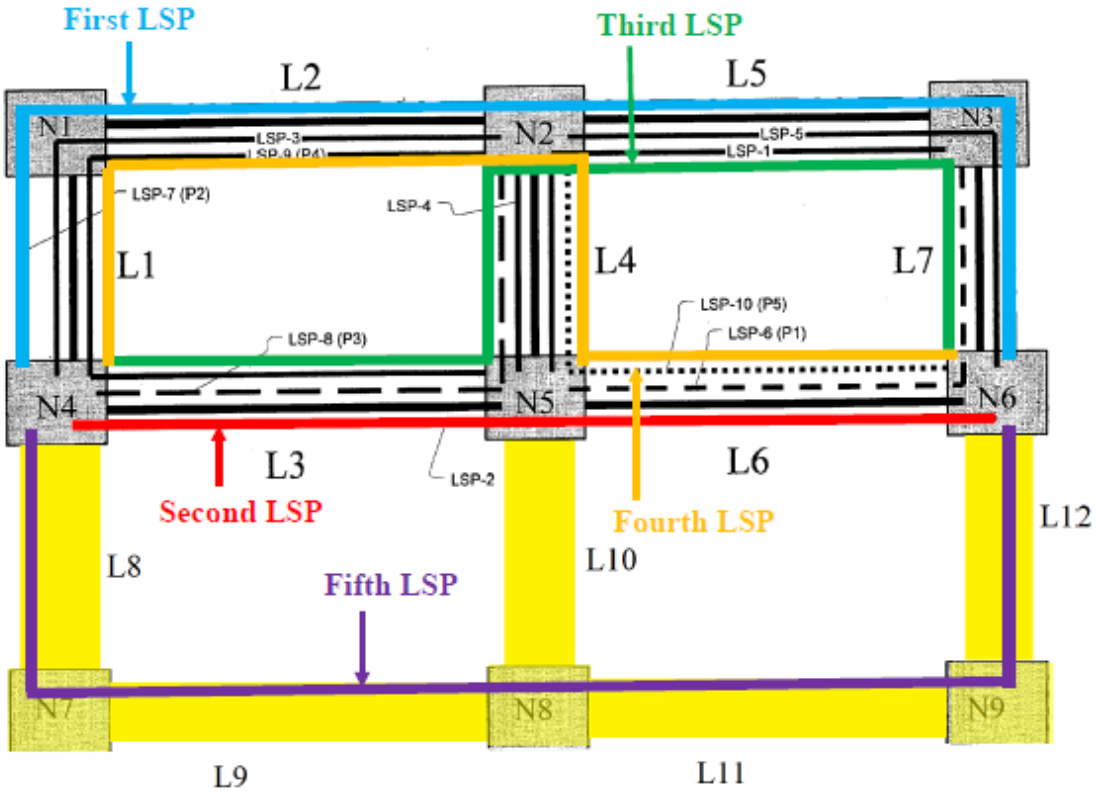
No.	'821 Patent Claim 12	Doshi '239
12	The method of claim 10, further comprising the step of configuring said working entity as revertive.	Doshi '239 discloses the method of claim 10, further comprising the step of configuring said working entity as revertive. *See* Claim 8.

No.	'821 Patent Claim 13	Doshi '239
13	The method of claim 1, further comprising the step of: if said entity pair reselection results in both working and protection entities being replaced, sequentially replacing said working entity	Doshi '239 discloses the method of claim 1, further comprising the step of: if said entity pair reselection results in both working and protection entities being replaced, sequentially replacing said working entity and said protection entity. Doshi discloses this limitation as explained in Claim 1 above. Doshi '239 in combination with (1) the knowledge of a person of ordinary skill in the art, alone or in further combination with (2) each (individually, as well as one or more together)

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 136 of 698

No.	'821 Patent Claim 13	Doshi '239
	and said protection entity.	of the references identified in Claim 13 of Exhibit E-3 renders the claim, including the present limitation, obvious. Below, one such combination is described. Regarding replacement of the claimed "working entity" and "protection entity," as discussed at Claims 2 and 5 above, Doshi '239's "primary path" is also referred to as a "working path" (e.g., working LSP), and its "restoration path" is also referred to as a "protection path" (e.g., protection LSP). Doshi '239 ¶0046. Doshi '239's primary path (aka working path) corresponds to the claimed "working entity" and Doshi '239's restoration path (aka protection path) corresponds to the claimed and "protection entity." As discussed at Claim 1 above, Doshi '239 teaches that a primary and a protection pair (e.g., LSP pair) are selected from the plurality of paths of the MPLS network and that, in view of Guichard's teaching of ongoing reoptimization efforts, it would have been an obvious to perform Doshi's cost calculations repeatedly to attempt to determine whether the originally selected LSP pair still has the overall minimum-cost or if another LSP pair should be selected. Thus, it would have been obvious for the LSPs of the originally selected pair (e.g., a primary LSP and a protection LSP) to be sequentially replaced as the ongoing reoptimization efforts identify lower cost LSPs. To build upon an earlier example, consider Doshi '239's modified and annotated Figure 1 below.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 137 of 698

No.	'821 Patent Claim 13	Doshi '239
		 Doshi '239 Fig. 1 (annotation added). As previously discussed at Claim 1 above, the Fifth LSP was selected as the primary path and the Second LSP was selected as the protection path. Upon an event trigger as taught by Guichard, attempted reoptimization of the LSPs would take place to determine if there is a better or more optimal LSP. If, for example, the reoptimization evaluation determined that the links of the Fifth LSP are heavily utilized as compared to the links of the First LSP, it would have been obvious for Doshi's cost determination to result in a new pair (e.g., the First LSP as the primary while the Second LSP remains as the protection) as having the overall minimum cost. Thus, the Fifth LSP would be replaced. A subsequent such reoptimization

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 138 of 698

No.	'821 Patent Claim 13	Doshi '239
		evaluation would find that, for example, the Second LSP has become suboptimal as a protection path and replace the Second LSP with another LSP e.g., a newly added Sixth LSP (not shown)). Of course, this is just an example, and given network status it would be understood that other LSPs may be selected as the overall minimum cost pair, but the LSPs have been replaced sequentially in this example—first the primary LSP followed by the protection LSP. It would have been obvious for the sequence of replacement to be reversed in other examples where the protection LSP becomes suboptimal before the primary LSP. It would have been similarly obvious for a single event-triggered reoptimization evaluation to find that both the Fifth and Second LSPs are suboptimal. When both LSPs are determined to be suboptimal, a POSITA would have found it obvious to replace the LSPs sequentially to minimize interruption in service. In that regard, a POSITA would have appreciated the desirability of maintaining a working and protection path throughout the replacement process. In replacing the Fifth and Second LSPs with the First and Sixth LSPs, for example, a POSITA would have found it obvious to first designate the First LSP (determined to be the optimal primary LSP) as a new protection LSP temporarily, thereby replacing the Second LSP. Next, traffic on the Fifth LSP would be switched over to the First LSP with the Fifth LSP temporarily serving as a protection LSP. The First LSP would thus become the new primary LSP as was determined to be optimal. The Fifth LSP, now acting as a protection LSP, would then be replaced by the Sixth LSP, which was determined to be the optimal protection LSP. A POSITA would have recognized that such sequential replacement would advantageously maintain two LSPs (e.g., one primary and one protection) throughout the replacement.

No.	'821 Patent Claim 14	Doshi '239
14[preamble]	A system for selecting entities within an MPLS network, comprising:	Doshi '239 discloses a system for selecting entities within an MPLS network, comprising. *See* Claim 1[preamble] and 1[a].
14[a]	a data structure comprising a plurality of transport entity descriptors;	Doshi '239 discloses a data structure comprising a plurality of transport entity descriptors. For example, Doshi '239 discloses this limitation as explained in Claim 1 above.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 139 of 698

No.	'821 Patent Claim 14	Doshi '239
		Doshi '239 in combination with (1) the knowledge of a person of ordinary skill in the art, alone or in further combination with (2) each (individually, as well as one or more together) of the references identified in element 14[a] of Exhibit E-3 renders the claim, including the present limitation, obvious. Below is an example of one such combination. Doshi '239 in combination with Guichard and Xu[3] renders this limitation obvious. Like the '821 patent and Doshi '239, Xu describes selecting paths to protect against failures in an MPLS network. Xu ¶¶0001, 0005–0006, 0030, 0035–0038, 0043. Xu discloses "selecting the end-to-end paths" based on a "set of modeled failure states." *Id.* ¶¶ 0006, 0043, 0049. The states represent "link failures that occur simultaneously" with "a predetermined probability of occurrence." Xu ¶¶ 0006–0007, 0049, 0053, 0080. A POSITA would have been motivated to combine the teachings of Doshi '239 and Xu, as discussed below. First, Doshi '239 and Xu are analogous art to the '821 patent. Just like the '821 patent and Doshi '239, Xu generally describes MPLS networks. '821 Patent, 3:15-18, Abstract; Doshi '239, ¶0049; Xu, ¶¶0036, 0038. Xu also addresses the problem of selecting paths within an MPLS network. '821 Patent, 4:63-5:8; Doshi '239, ¶0014; Xu, ¶¶ 0002, 0030, 0040, 0043, 0060. Given the similarities between the references and their addressing the same problem of selecting paths, a POSITA considering Doshi '239 would have naturally considered the teachings of Xu. Second, a POSITA would have been motivated to combine the teachings of Doshi '239 and Xu to produce numerous predictable and beneficial results. Doshi '239 teaches protecting services against link failures by considering disjointedness when selecting primary and protection paths. Doshi '239 ¶¶0155–0156. Strictly disjoint paths share no common links or

[3] U.S. Patent Publication No. 2011/0141877 to Xu et al.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 140 of 698

No.	'821 Patent Claim 14	Doshi '239
		nodes other than their common ingress and egress nodes, so the failure of a single network element would not affect both paths. Doshi '239 ¶¶0055, 0158. Doshi '239 recognizes, however, that there "will be some cases … where no strictly disjoint path exists" or where "other factors…may lead to selection of a protection path, for a given primary path, that is not strictly disjoint from the primary path." Xu ¶0055. It would be obvious, in such instances, to at least minimize the likelihood of a concurrent failure. Xu complements Doshi '239 by teaching that a probability of occurrence is determined for each of a set of failure states and that path selection decisions consider the probability of various failure states occurring. Xu ¶¶0006, 0053. A POSITA would have been motivated to determine a probability of occurrence of failure states, as taught by Xu, when implementing Doshi '239's path selection in order to further inform selection decisions and better protect services against link failures in circumstances where no strictly disjoint pair of paths exists or where other factors lead to selection of a partially disjoint pair as discussed above. Determining the probability of various failure states would further refine selection between pairs of paths that have the same disjointedness (e.g., between pairs all having a primary path that shares one link in common with its protection path). Such pairs have the same disjointedness, but there may be differences in risk of failure for different links. A pair in which the link in common between primary and protection paths of the pair has a higher risk of failure would, everything else being equal, be more likely to concurrently fail than a pair in which the link in common has a lower risk of failure. A POSITA would have been motivated to apply Xu's probability determination to differentiate between such pairs to improve reliability by reducing the probability of concurrent failure of Doshi '239's selected pair. The combination is merely the use of a known technique (Xu's probability determination) to improve a similar method (Doshi '239's method for selecting paths) in the same way with predictable results—selecting paths with lower probability of failure. The results would have been predictable, and there would have been a reasonable expectation of success in the combination since Doshi '239 and Xu are analogous art in the same field of endeavor.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 141 of 698

No.	'821 Patent Claim 14	Doshi '239
		Implementing the combination would have been within a POSITA's skillset because techniques for determining the probability of network failures were known in the art. Xu leaves implementation details of failure probability determination to a POSITA, recognizing that a POSITA would have had sufficient skills to configure the determination. Xu explains that "the present disclosure [is] not limited to any specific combination of hardware and software and the computer program code required to implement the foregoing can be developed by a person of ordinary skill in the art." Xu ¶¶0094–0096. Xu's techniques can be "practiced or carried out in a variety of applications." Xu ¶0029. Since Doshi '239 is analogous art in the same field of endeavor, it would have been obvious to carry out Xu's probability determination in the context of Doshi '239's path selection with a reasonable expectation of success. Indeed, Doshi '239 acknowledges that a POSITA would have had the necessary skill set to make "[v]arious modifications of the described embodiments." Doshi '239 ¶0300. It would have been obvious for such modifications to include Xu's failure state probability determination. The combination of Doshi '239 and Xu would render this limitation obvious. First, as discussed above, Doshi '239 in combination with Guichard discloses providing a plurality of LSPs between two end nodes where the LSPs correspond to "a plurality of [] transport entit[ies]." Second, as discussed above, Doshi '239 in combination with Guichard discloses calculating the cost of various LSP pairs and selecting the lowest cost pair. A POSITA would have found it obvious to store and reference "descriptors" of the LSPs at least temporarily to accomplish the cost calculation and selection described above. Such storage would be in a memory, entries in which a POSITA would understand corresponds to a "data structure." Moreover, Doshi '239 expressly describes saving information about a plurality of LSP pairs in a data structure during cost calculation. Doshi '239 ¶0143 ("[T]he minimum-cost pair is saved in the data structure LowKPair."); *Id.* ¶0144 ("[T]he lowest-cost pair is stored in the data structure LowLPair."). Doshi '239 also maintains "an extended link-state database" (referred to as "TE/Share database") that include information used for "path computation."

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 142 of 698

No.	'821 Patent Claim 14	Doshi '239
		Id. ¶0196. "This database contains a data structure for each link L that the node owns. The data structure for each link contains information (e.g., bandwidth, link-id) about all the other links in the network for which link L provides restoration capacity." *Id.* ¶0197; *see also id.* ¶¶ 0082, 0091–0095, 0172–0176. Doshi '239 also describes a functional architecture that includes a "Path Management" module that is "responsible for path setup, refresh, tear-down, and monitoring functions." Doshi '239 ¶0273. It would have been obvious to a POSITA to store "descriptors" of the monitored paths in a "data structure" to differentiate the paths and to provide data entries for the paths such that changes can be tracked over time and saved in association with the relevant path.
14[b]	an entity protection switch configured to switch between a working entity and a protection entity; and	Doshi '239 discloses an entity protection switch configured to switch between a working entity and a protection entity. For example, Doshi '239 discloses this limitation as explained in Claim 1 above. Doshi '239 in combination with (1) the knowledge of a person of ordinary skill in the art, alone or in further combination with (2) each (individually, as well as one or more together) of the references identified in element 14[b] of Exhibit E-3 renders the claim, including the present limitation, obvious. Below is an example of one such combination. Doshi '239 in combination with Guichard and Xu[4] renders this limitation obvious. Like the '821 patent and Doshi '239, Xu describes selecting paths to protect against failures in an MPLS network. Xu ¶¶0001, 0005–0006, 0030, 0035–0038, 0043. Xu discloses "selecting the end-to-end paths" based on a "set of modeled failure states." *Id.* ¶¶ 0006, 0043, 0049. The states represent "link failures that occur simultaneously" with "a predetermined probability of occurrence." Xu ¶¶ 0006–0007, 0049, 0053, 0080.

[4] U.S. Patent Publication No. 2011/0141877 to Xu et al.

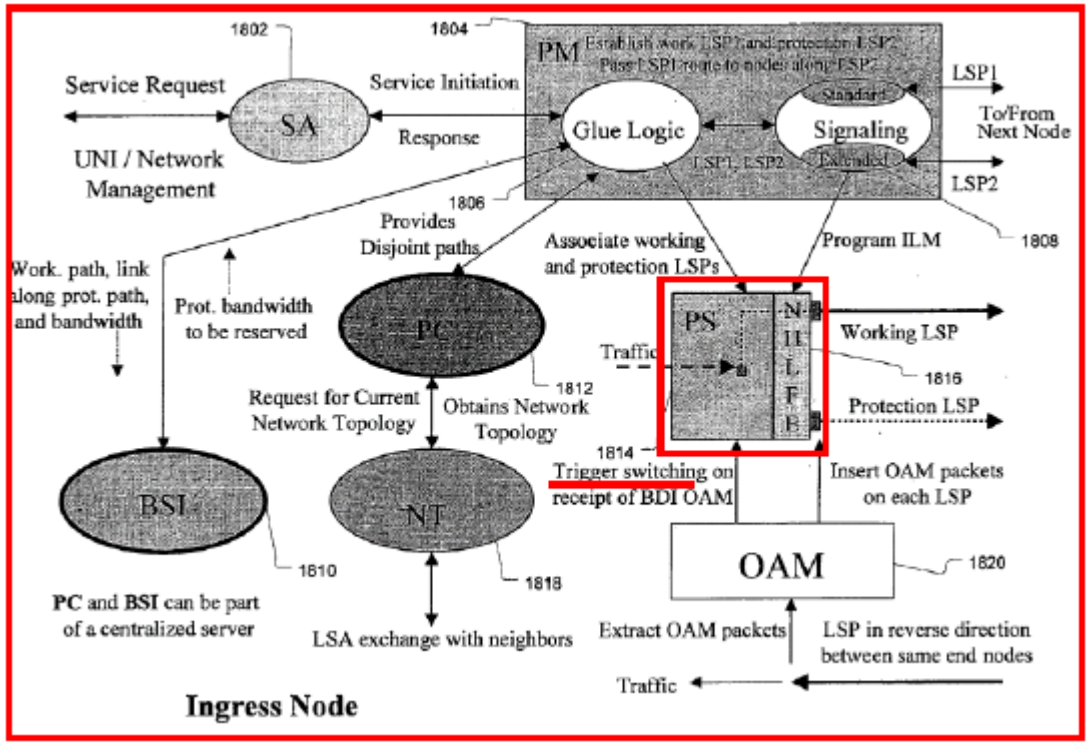
Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 143 of 698

No.	'821 Patent Claim 14	Doshi '239
		A POSITA would have been motivated to combine the teachings of Doshi '239 and Xu, as discussed below. First, Doshi '239 and Xu are analogous art to the '821 patent. Just like the '821 patent and Doshi '239, Xu generally describes MPLS networks. '821 Patent, 3:15-18, Abstract; Doshi '239, ¶0049; Xu, ¶¶0036, 0038. Xu also addresses the problem of selecting paths within an MPLS network. '821 Patent, 4:63-5:8; Doshi '239, ¶0014; Xu, ¶¶ 0002, 0030, 0040, 0043, 0060. Given the similarities between the references and their addressing the same problem of selecting paths, a POSITA considering Doshi '239 would have naturally considered the teachings of Xu. Second, a POSITA would have been motivated to combine the teachings of Doshi '239 and Xu to produce numerous predictable and beneficial results. Doshi '239 teaches protecting services against link failures by considering disjointedness when selecting primary and protection paths. Doshi '239 ¶¶0155–0156. Strictly disjoint paths share no common links or nodes other than their common ingress and egress nodes, so the failure of a single network element would not affect both paths. Doshi '239 ¶¶0055, 0158. Doshi '239 recognizes, however, that there "will be some cases … where no strictly disjoint path exists" or where "other factors…may lead to selection of a protection path, for a given primary path, that is not strictly disjoint from the primary path." Xu ¶0055. It would be obvious, in such instances, to at least minimize the likelihood of a concurrent failure. Xu complements Doshi '239 by teaching that a probability of occurrence is determined for each of a set of failure states and that path selection decisions consider the probability of various failure states occurring. Xu ¶¶0006, 0053. A POSITA would have been motivated to determine a probability of occurrence of failure states, as taught by Xu, when implementing Doshi '239's path selection in order to further inform selection decisions and better protect services against link failures in circumstances where no strictly disjoint pair of paths exists or where other factors lead to selection of a partially disjoint pair as discussed above. Determining the probability of various failure states would further refine selection between pairs of paths that have the same disjointedness (e.g., between pairs all having a primary path that shares one link in common with its protection

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 144 of 698

No.	'821 Patent Claim 14	Doshi '239
		path). Such pairs have the same disjointedness, but there may be differences in risk of failure for different links. A pair in which the link in common between primary and protection paths of the pair has a higher risk of failure would, everything else being equal, be more likely to concurrently fail than a pair in which the link in common has a lower risk of failure. A POSITA would have been motivated to apply Xu's probability determination to differentiate between such pairs to improve reliability by reducing the probability of concurrent failure of Doshi '239's selected pair. The combination is merely the use of a known technique (Xu's probability determination) to improve a similar method (Doshi '239's method for selecting paths) in the same way with predictable results—selecting paths with lower probability of failure. The results would have been predictable, and there would have been a reasonable expectation of success in the combination since Doshi '239 and Xu are analogous art in the same field of endeavor. Implementing the combination would have been within a POSITA's skillset because techniques for determining the probability of network failures were known in the art. Xu leaves implementation details of failure probability determination to a POSITA, recognizing that a POSITA would have had sufficient skills to configure the determination. Xu explains that "the present disclosure [is] not limited to any specific combination of hardware and software and the computer program code required to implement the foregoing can be developed by a person of ordinary skill in the art." Xu ¶¶0094–0096. Xu's techniques can be "practiced or carried out in a variety of applications." Xu ¶0029. Since Doshi '239 is analogous art in the same field of endeavor, it would have been obvious to carry out Xu's probability determination in the context of Doshi '239's path selection with a reasonable expectation of success. Indeed, Doshi '239 acknowledges that a POSITA would have had the necessary skill set to make "[v]arious modifications of the described embodiments." Doshi '239 ¶0300. It would have been obvious for such modifications to include Xu's failure state probability determination. The combination of Doshi '239 and Xu would render this limitation obvious.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 145 of 698

No.	'821 Patent Claim 14	Doshi '239
		Doshi '239 describes that "after detecting a failure, end nodes of an LSP switch traffic from a primary LSP to its corresponding protection LSP." Doshi '239 ¶¶0266–0267; *see also id.* ¶0266 ("Switching Between Working and Protection LSPs"). Doshi '239 further discloses a "Protection Switching (PS)" module that is "responsible for switching the affected traffic onto a protection LSP after detecting a failure or receiving failure notification." *Id.* ¶0278. Doshi '239's exemplary Protection Switching module at Figure 18 illustrates its inclusion in a node of an LSP.  Doshi '239 Fig. 18 (annotation added).

No.	'821 Patent Claim 14	Doshi '239
		Thus, Doshi '239's node (including the Protection Switching module) or the Protection Switching module itself, which is configured to switch between the working LSP and the protection LSP, renders obvious this limitation.
14[c]	digital logic configured to select said working entity and said protection entity from said plurality of transport entity descriptors, comprising:	Doshi '239 discloses digital logic configured to select said working entity and said protection entity from said plurality of transport entity descriptors, comprising. *See infra* claim14[d]-[e].
14[d]	logic configured to determine a probability of concurrent failure of said working entity and said protection entity;	Doshi '239 discloses logic configured to determine a probability of concurrent failure of said working entity and said protection entity. For example, Doshi '239 discloses this limitation as explained in Claim 1 above. Doshi '239 in combination with (1) the knowledge of a person of ordinary skill in the art, alone or in further combination with (2) each (individually, as well as one or more together) of the references identified in element 14[d] of Exhibit E-3 renders the claim, including the present limitation, obvious. Below is an example of one such combination. Doshi '239 in combination with Guichard and Xu[5] renders this limitation obvious. Like the '821 patent and Doshi '239, Xu describes selecting paths to protect against failures in an MPLS network. Xu ¶¶0001, 0005–0006, 0030, 0035–0038, 0043. Xu discloses "selecting the end-to-end paths" based on a "set of modeled failure states." *Id.* ¶¶ 0006, 0043, 0049. The states represent "link failures that occur simultaneously" with "a predetermined probability of occurrence." Xu ¶¶ 0006–0007, 0049, 0053, 0080.

[5] U.S. Patent Publication No. 2011/0141877 to Xu et al.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 147 of 698

No.	'821 Patent Claim 14	Doshi '239
		A POSITA would have been motivated to combine the teachings of Doshi '239 and Xu, as discussed below. First, Doshi '239 and Xu are analogous art to the '821 patent. Just like the '821 patent and Doshi '239, Xu generally describes MPLS networks. '821 Patent, 3:15-18, Abstract; Doshi '239, ¶0049; Xu, ¶¶0036, 0038. Xu also addresses the problem of selecting paths within an MPLS network. '821 Patent, 4:63-5:8; Doshi '239, ¶0014; Xu, ¶¶ 0002, 0030, 0040, 0043, 0060. Given the similarities between the references and their addressing the same problem of selecting paths, a POSITA considering Doshi '239 would have naturally considered the teachings of Xu. Second, a POSITA would have been motivated to combine the teachings of Doshi '239 and Xu to produce numerous predictable and beneficial results. Doshi '239 teaches protecting services against link failures by considering disjointedness when selecting primary and protection paths. Doshi '239 ¶¶0155–0156. Strictly disjoint paths share no common links or nodes other than their common ingress and egress nodes, so the failure of a single network element would not affect both paths. Doshi '239 ¶¶0055, 0158. Doshi '239 recognizes, however, that there "will be some cases … where no strictly disjoint path exists" or where "other factors…may lead to selection of a protection path, for a given primary path, that is not strictly disjoint from the primary path." Xu ¶0055. It would be obvious, in such instances, to at least minimize the likelihood of a concurrent failure. Xu complements Doshi '239 by teaching that a probability of occurrence is determined for each of a set of failure states and that path selection decisions consider the probability of various failure states occurring. Xu ¶¶0006, 0053. A POSITA would have been motivated to determine a probability of occurrence of failure states, as taught by Xu, when implementing Doshi '239's path selection in order to further inform selection decisions and better protect services against link failures in circumstances where no strictly disjoint pair of paths exists or where other factors lead to selection of a partially disjoint pair as discussed above. Determining the probability of various failure states would further refine selection between pairs of paths that have the same disjointedness (e.g., between pairs all having a primary path that shares one link in common with its protection

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 148 of 698

No.	'821 Patent Claim 14	Doshi '239
		path). Such pairs have the same disjointedness, but there may be differences in risk of failure for different links. A pair in which the link in common between primary and protection paths of the pair has a higher risk of failure would, everything else being equal, be more likely to concurrently fail than a pair in which the link in common has a lower risk of failure. A POSITA would have been motivated to apply Xu's probability determination to differentiate between such pairs to improve reliability by reducing the probability of concurrent failure of Doshi '239's selected pair. The combination is merely the use of a known technique (Xu's probability determination) to improve a similar method (Doshi '239's method for selecting paths) in the same way with predictable results—selecting paths with lower probability of failure. The results would have been predictable, and there would have been a reasonable expectation of success in the combination since Doshi '239 and Xu are analogous art in the same field of endeavor. Implementing the combination would have been within a POSITA's skillset because techniques for determining the probability of network failures were known in the art. Xu leaves implementation details of failure probability determination to a POSITA, recognizing that a POSITA would have had sufficient skills to configure the determination. Xu explains that "the present disclosure [is] not limited to any specific combination of hardware and software and the computer program code required to implement the foregoing can be developed by a person of ordinary skill in the art." Xu ¶¶0094–0096. Xu's techniques can be "practiced or carried out in a variety of applications." Xu ¶0029. Since Doshi '239 is analogous art in the same field of endeavor, it would have been obvious to carry out Xu's probability determination in the context of Doshi '239's path selection with a reasonable expectation of success. Indeed, Doshi '239 acknowledges that a POSITA would have had the necessary skill set to make "[v]arious modifications of the described embodiments." Doshi '239 ¶0300. It would have been obvious for such modifications to include Xu's failure state probability determination. The combination of Doshi '239 and Xu would render this limitation obvious.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 149 of 698

No.	'821 Patent Claim 14	Doshi '239
		First, as discussed above, Doshi '239's path calculation considers disjointedness of the working LSP ("working entity") and the protection LSP ("protection entity"). Further, a POSITA would have found it obvious to use "digital logic" to perform Doshi's flowchart logic, including path calculation and selection techniques. Doshi '239 recognizes that when two paths are disjoint, e.g., by being node- and link-disjoint, "then a failure affecting one of them will not affect the other." Doshi '239 ¶0040. A POSITA would have understood that probability of concurrent failure is inversely related to disjointedness—as paths become more disjoint, the likelihood they will fail concurrently goes down. Doshi '239 recognizes that when two paths are disjoint, e.g., by being node- and link-disjoint, "then a failure affecting one of them will not affect the other." Doshi '239 ¶0040. It would take two distinct failures (one affecting each path) for both paths to fail concurrently. Since Doshi '239 explains that the "probability of occurrence" of two such failures is "very insignificant," a POSITA would have understood that determining that two paths are disjoint with no nodes or links in common is also a determination that the paths have a "very insignificant" probability of failing concurrently. Doshi '239 ¶0040. Accordingly, determining path disjointedness (as part of Doshi '239's flowchart) corresponds to "determin[ing] a probability of concurrent failure" since knowledge of, e.g., strict disjointedness represents a determination that there is a "very insignificant" probability that the paths will fail at the same time. Second, Xu discloses a "network architecture [that] uses multiple paths between each ingress-egress router pair" in an MPLS network and provides techniques for "selecting the end-to-end paths" between the pair by "Network-Management Software (NMS)." Xu ¶¶0030–0038. To inform path selection, Xu models failure states that would cause at least one path to fail. *Id.* ¶¶0006, 0049, 0080, Claim 1, Claim 11. A POSITA would have understood that at least some of the modeled failure states would be states in which the failure affects two paths simultaneously—corresponding to a state of concurrent failure. Xu explains that each failure state has a "predetermined probability of occurrence." Xu ¶0006. A POSITA would have understood that a "predetermined" probability has been "determine[d]." When a failure state is one where the failure affects two paths

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 150 of 698

No.	'821 Patent Claim 14	Doshi '239
		simultaneously—a state of concurrent failure—Xu's predetermining a probability of occurrence for that failure state corresponds to "determin[ing] a probability of concurrent failure." Each failure state is assigned a weight in Xu's selection function based on how common the state is (probability of occurrence). *Id.* ¶0010; *see also id.* ¶¶0033, 0045, 0053, 0080. As discussed in more detail above, a POSITA would have found it obvious determine the probability of occurrence of failure states, as taught by Xu, when implementing Doshi '239's calculation and selection techniques and would have been motivated to determine the probability of various failure states to achieve a more refined selection between pairs of paths that have the same disjointedness.
14[e]	logic configured to determine an entity cost of said plurality of transport entity descriptors: and	Doshi '239 discloses logic configured to determine an entity cost of said plurality of transport entity descriptors. For example, Doshi '239 discloses this limitation as explained in Claim 1 above. Doshi '239 in combination with (1) the knowledge of a person of ordinary skill in the art, alone or in further combination with (2) each (individually, as well as one or more together) of the references identified in element 14[e] of Exhibit E-3 renders the claim, including the present limitation, obvious. Below is an example of one such combination. Doshi '239 in combination with Guichard and Xu[6] renders this limitation obvious. Like the '821 patent and Doshi '239, Xu describes selecting paths to protect against failures in an MPLS network. Xu ¶¶0001, 0005–0006, 0030, 0035–0038, 0043. Xu discloses "selecting the end-to-end paths" based on a "set of modeled failure states." *Id.* ¶¶ 0006, 0043, 0049. The states represent "link failures that occur simultaneously" with "a predetermined probability of occurrence." Xu ¶¶ 0006–0007, 0049, 0053, 0080.

[6] U.S. Patent Publication No. 2011/0141877 to Xu et al.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 151 of 698

No.	'821 Patent Claim 14	Doshi '239
		A POSITA would have been motivated to combine the teachings of Doshi '239 and Xu, as discussed below. First, Doshi '239 and Xu are analogous art to the '821 patent. Just like the '821 patent and Doshi '239, Xu generally describes MPLS networks. '821 Patent, 3:15-18, Abstract; Doshi '239, ¶0049; Xu, ¶¶0036, 0038. Xu also addresses the problem of selecting paths within an MPLS network. '821 Patent, 4:63-5:8; Doshi '239, ¶0014; Xu, ¶¶ 0002, 0030, 0040, 0043, 0060. Given the similarities between the references and their addressing the same problem of selecting paths, a POSITA considering Doshi '239 would have naturally considered the teachings of Xu. Second, a POSITA would have been motivated to combine the teachings of Doshi '239 and Xu to produce numerous predictable and beneficial results. Doshi '239 teaches protecting services against link failures by considering disjointedness when selecting primary and protection paths. Doshi '239 ¶¶0155–0156. Strictly disjoint paths share no common links or nodes other than their common ingress and egress nodes, so the failure of a single network element would not affect both paths. Doshi '239 ¶¶0055, 0158. Doshi '239 recognizes, however, that there "will be some cases … where no strictly disjoint path exists" or where "other factors…may lead to selection of a protection path, for a given primary path, that is not strictly disjoint from the primary path." Xu ¶0055. It would be obvious, in such instances, to at least minimize the likelihood of a concurrent failure. Xu complements Doshi '239 by teaching that a probability of occurrence is determined for each of a set of failure states and that path selection decisions consider the probability of various failure states occurring. Xu ¶¶0006, 0053. A POSITA would have been motivated to determine a probability of occurrence of failure states, as taught by Xu, when implementing Doshi '239's path selection in order to further inform selection decisions and better protect services against link failures in circumstances where no strictly disjoint pair of paths exists or where other factors lead to selection of a partially disjoint pair as discussed above. Determining the probability of various failure states would further refine selection between pairs of paths that have the same disjointedness (e.g., between pairs all having a primary path that shares one link in common with its protection

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 152 of 698

No.	'821 Patent Claim 14	Doshi '239
		path). Such pairs have the same disjointedness, but there may be differences in risk of failure for different links. A pair in which the link in common between primary and protection paths of the pair has a higher risk of failure would, everything else being equal, be more likely to concurrently fail than a pair in which the link in common has a lower risk of failure. A POSITA would have been motivated to apply Xu's probability determination to differentiate between such pairs to improve reliability by reducing the probability of concurrent failure of Doshi '239's selected pair. The combination is merely the use of a known technique (Xu's probability determination) to improve a similar method (Doshi '239's method for selecting paths) in the same way with predictable results—selecting paths with lower probability of failure. The results would have been predictable, and there would have been a reasonable expectation of success in the combination since Doshi '239 and Xu are analogous art in the same field of endeavor. Implementing the combination would have been within a POSITA's skillset because techniques for determining the probability of network failures were known in the art. Xu leaves implementation details of failure probability determination to a POSITA, recognizing that a POSITA would have had sufficient skills to configure the determination. Xu explains that "the present disclosure [is] not limited to any specific combination of hardware and software and the computer program code required to implement the foregoing can be developed by a person of ordinary skill in the art." Xu ¶¶0094–0096. Xu's techniques can be "practiced or carried out in a variety of applications." Xu ¶0029. Since Doshi '239 is analogous art in the same field of endeavor, it would have been obvious to carry out Xu's probability determination in the context of Doshi '239's path selection with a reasonable expectation of success. Indeed, Doshi '239 acknowledges that a POSITA would have had the necessary skill set to make "[v]arious modifications of the described embodiments." Doshi '239 ¶0300. It would have been obvious for such modifications to include Xu's failure state probability determination. The combination of Doshi '239 and Xu would render this limitation obvious.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 153 of 698

No.	'821 Patent Claim 14	Doshi '239
		First, as discussed above, Doshi '239 calculates the cost for each LSP path of candidate LSP pairs. This is done for all candidate LSP pairs, which means that a cost is calculated for each LSP path. As discussed above, a POSITA would have found it obvious to store and reference "descriptors" of the LSPs at least temporarily to accomplish cost calculation and selection of the LSPs. Further, a POSITA would have found it obvious to use "digital logic" to perform Doshi's flowchart logic, including path calculation and selection techniques.
14[f]	logic configured to reselect said working entity and said protection entity from said plurality of transport entity descriptors upon a reselection event,	Doshi '239 discloses logic configured to reselect said working entity and said protection entity from said plurality of transport entity descriptors upon a reselection event. *See* Claim 1[d].
14[g]	wherein said reselection event is selected from a group consisting of adding an entity to said plurality of transport entities, removing an entity from said plurality of transport entities, an operational status change for one of said plurality of transport entities, and a change in overall cost for one of said plurality of transport entities.	Doshi '239 discloses wherein said reselection event is selected from a group consisting of adding an entity to said plurality of transport entities, removing an entity from said plurality of transport entities, an operational status change for one of said plurality of transport entities, and a change in overall cost for one of said plurality of transport entities. *See* Claim 1[e].

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 154 of 698

No.	'821 Patent Claim 15	Doshi '239
15	The system of claim 14 wherein said entity protection switch comprises a 1:1 switch.	Doshi '239 discloses the system of claim 14 wherein said entity protection switch comprises a 1:1 switch. For example, as discussed in Claim 14 *supra*, Doshi '239 discloses the claimed "entity protection switch." Doshi '239 further discloses that "[i]f the bandwidth were allocated in advance of a failure, this would correspond, in the parlance of the field of protection and restoration for optical transport networks, to a 1:1 protection scheme." Doshi '239 ¶0060. Doshi's protection switching module in the context of a 1:1 protection scheme discloses this limitation.

No.	'821 Patent Claim 16	Doshi '239
16	The system of claim 14 wherein said entity protection switch comprises a 1+1 switch.	Doshi '239 discloses the system of claim 14 wherein said entity protection switch comprises a 1+1 switch. For example, as discussed in Claim 14 *supra*, Doshi '239 discloses the claimed "entity protection switch." Doshi further discloses that "[i]f the bandwidth were not only allocated, but additionally if a copy of the service path's data were to be duplicated to the protection path, this would correspond to a 1+1 protection scheme." Doshi '239 ¶0060. In some circumstances, "a 1+1 restoration scheme is the best option available for restoration" and therefore are "popular." *Id.* ¶¶0147, 0178. Doshi's protection switching module in the context of a 1+1 protection scheme discloses this limitation.

No.	'821 Patent Claim 17	Doshi '239
17[preamble]	Non-transitory computer readable media configured to perform a method comprising the steps of:	Doshi '239 discloses non-transitory computer readable media configured to perform a method comprising the steps of. *See* Claim 14.
17[a]	providing a plurality of MPLS transport entities between a first	Doshi '239 discloses providing a plurality of MPLS transport entities between a first endpoint and a second endpoint.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 155 of 698

No.	'821 Patent Claim 17	Doshi '239
	endpoint and a second endpoint;	*See* Claim 1[a].
17[b]	determining an overall cost for each entity pair of said plurality of entities;	Doshi '239 discloses determining an overall cost for each entity pair of said plurality of entities. *See* Claim 1[b].
17[c]	selecting an entity pair from said plurality of transport entities based at least in part upon said overall cost; and	Doshi '239 discloses selecting an entity pair from said plurality of transport entities based at least in part upon said overall cost. *See* Claim 1[c].
17[d]	if an entity pair reselection event occurs, reselecting said entity pair from the group consisting of said entity pair and a replacement entity pair comprising at least one entity distinct from the entities of said entity pair,	Doshi '239 discloses if an entity pair reselection event occurs, reselecting said entity pair from the group consisting of said entity pair and a replacement entity pair comprising at least one entity distinct from the entities of said entity pair. *See* Claim 1[d].
17[e]	wherein said entity pair reselection event is selected from a group consisting of adding an entity to said plurality of transport entities, removing an entity from said plurality of	Doshi '239 discloses wherein said entity pair reselection event is selected from a group consisting of adding an entity to said plurality of transport entities, removing an entity from said plurality of transport entities, an operational status change for one of said plurality of transport entities, and a change in overall cost for one of said plurality of transport entities. *See* Claim 1[e].

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 156 of 698

No.	'821 Patent Claim 17	Doshi '239
	transport entities, an operational status change for one of said plurality of transport entities, and a change in overall cost for one of said plurality of transport entities.	

No.	'821 Patent Claim 18	Doshi '239
18[preamble]	The non-transitory computer readable media of claim 17, wherein said step of selecting an entity pair further comprises:	Doshi '239 discloses the non-transitory computer readable media of claim 17, wherein said step of selecting an entity pair further comprises. *See* Claim 14.
18[a]	selecting a working entity from said plurality of transport entities;	Doshi '239 discloses selecting a working entity from said plurality of transport entities. *See* Claim 2[a].
18[b]	selecting a protection entity from said plurality of transport entities; and	Doshi '239 discloses selecting a protection entity from said plurality of transport entities. *See* Claim 2[b].
18[c]	selecting an active entity from the set consisting of said working entity and said protection entity.	Doshi '239 discloses selecting an active entity from the set consisting of said working entity and said protection entity. *See* Claim 3.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 157 of 698

No.	'821 Patent Claim 19	Doshi '239
19	The non-transitory readable media of claim 18, wherein said step of selecting an entity pair further comprises minimizing an overall cost function.	Doshi '239 discloses the non-transitory readable media of claim 18, wherein said step of selecting an entity pair further comprises minimizing an overall cost function. *See* Claim 4.

No.	'821 Patent Claim 20	Doshi '239
20[preamble]	The non-transitory readable media of claim 19, wherein said overall cost function comprises:	Doshi '239 discloses the non-transitory readable media of claim 19, wherein said overall cost function comprises. *See* Claim 19.
20[a]	minimizing a probability of concurrent failure of said protection entity and said working entity; and	Doshi '239 discloses minimizing a probability of concurrent failure of said protection entity and said working entity. *See* Claim 5.
20[b]	a predefined metric selected from the group consisting of interior gateway protocol (IGP) and traffic engineering (TE).	Doshi '239 discloses a predefined metric selected from the group consisting of interior gateway protocol (IGP) and traffic engineering (TE). *See* Claim 7.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 158 of 698

EXHIBIT B-5
Defendant's Preliminary Invalidity Contentions
Orckit Corporation v. Cisco Systems, Inc., 2:22-cv-00276-JRG-RSP

Chart for U.S. Patent 8,830,821 ("the '821 Patent")
U.S. Patent No. US 2005/0083928 A1 to Sivabalan, et. al. ("Sivabalan '928")

As shown in the chart below, all Asserted Claims of the '821 Patent are invalid under (1) 35 U.S.C. §§ 102 (a) (b), (e), and (g) because Sivabalan '928 meets each element of those claims, and/or (2) 35 U.S.C. § 103 because Sivabalan '928 renders those claims obvious either alone, or in combination with the knowledge of a person having ordinary skill in the art, and in further combination with the references specifically identified below and in the following claim chart and/or one or more references identified in Defendant's Preliminary Invalidity Contentions. The following quotations and diagrams come from Sivabalan '928 titled "Transparent Re-Routing Of MPLS Traffic Engineering LSPs Within A Link Bundle", which was filed on October 20, 2003, and published on April 21, 2005.

Motivations to combine the disclosures in Sivabalan '928 with disclosures in other publications known in the art, as explained in this chart, include at least the similarity in subject matter between the references to the extent they concern methods of selecting paths in a network and/or determining the cost of potential paths in a network. Insofar as the references cite other patents or publications, or suggest additional changes, one of ordinary skill in the art would look beyond a single reference to other references in the field.

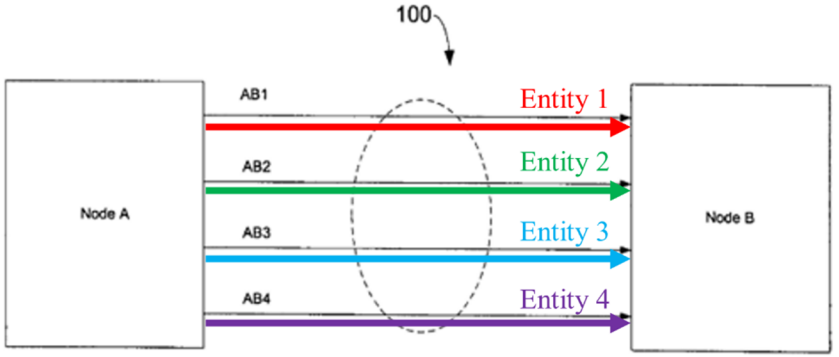
These invalidity contentions are based on Defendant's present understanding of the Asserted Claims, and Orckit's apparent construction of the claims in its November 3, 2022 Disclosure of Asserted Claims and Infringement Contentions Pursuant to P.R. 3-1, and Orckit's January 19, 2023 First Amended Disclosure of Asserted Claims and Infringement Contentions Pursuant to P.R. 3-1 (Orckit's "Infringement Disclosures"), which is deficient at least insofar as it fails to cite any documents or identify accused structures, acts, or materials in the Accused Products with particularity. Defendant does not agree with Orckit's application of the claims, or that the claims satisfy the requirements of 35 U.S.C. § 112. Defendant's contentions herein are not, and should in no way be seen as, admissions or adoptions as to any particular claim scope or construction, or as any admission that any particular element is met by any accused product in any particular way. Defendant objects to any attempt to imply claim construction from this chart. Defendant's prior art invalidity contentions are made in a variety of alternatives and do not represent Defendant's agreement or view as to the meaning, definiteness, written description support for, or enablement of any claim contained therein.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 159 of 698

The following contentions are subject to revision and amendment pursuant to Federal Rule of Civil Procedure 26(e), the Local Rules, and the Orders of record in this matter subject to further investigation and discovery regarding the prior art and the Court's construction of the claims at issue.

No.	'821 Patent Claim 1	Sivabalan '928
1[preamble]	An entity selection method performed by a network device, comprising the steps of:	Sivabalan '928 discloses an entity selection method performed by a network device. For example, Sivabalan '928 discloses a node A for transparent re-routing of MPLS traffic engineering LSPs within a link bundle and selection of primary or secondary/ backup component links. Sivabalan '928 at Para No. 0005 ("[To avoid the need to advertise information about each link between a pair of nodes, the concept of a "link bundle" has been developed to support MPLS Traffic Engineering. A link bundle combines multiple links between a pair of nodes. The links that make up a bundle are then referred to as "component links." Link bundling improves routing scalability by reducing the amount of information that has to be advertised by link state routing protocols to support signaling and computation of MPLS Traffic Engineering LSPs. Information about link bundles including bandwidth information is flooded throughout the network while information about component links is maintained only at the directly connected nodes.]")

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 160 of 698

No.	'821 Patent Claim 1	Sivabalan '928
		 FIG. 1 Fig. 1 (annotation added) Sivabalan '928 at Para No. 0033 ("[FIG. 1. depicts a representative link bundle scenario according to one embodiment of the present invention. Link bundle 100 interconnects nodes A and B. Link bundle 100 includes components links AB1, AB2, AB3, and AB4. Node A propagates information about the aggregate characteristics of link bundle 100 including the available aggregate bandwidth to other nodes in the network in accordance with the relevant link state routing protocols underlying MPLS Traffic Engineering operation. Information about the component links AB1, AB2, AB3, and AB4 is not distributed in this way but only maintained locally to node A. Node B has similar knowledge of the links extending in the reverse direction. Information requesting the establishment of an LSP through nodes A and B will thus specify link bundle 100 but not the constituent component links.]") Sivabalan '928 at Para No. 0034 ("A local admission control mechanism at node A determines whether or not to accept the LSP and picks a component link for use by the LSP. The component link should have sufficient bandwidth available to support the LSP's requirements. The component link chosen in this fashion is called a "primary" component link. The admission control mechanism further selects one or more other component links as secondary component links. As will be explained below, the secondary component links are used as backup links when the primary component link fails. The secondary component link

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 161 of 698

No.	'821 Patent Claim 1	Sivabalan '928
		information is preconfigured into the label forwarding table entry for the LSP so that the switchover can occur very quickly.]")

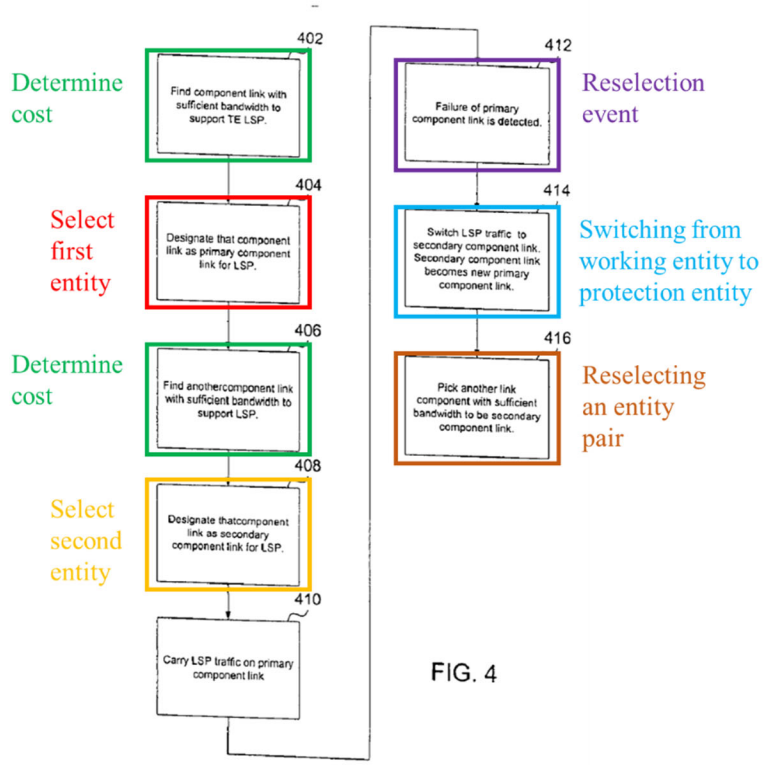
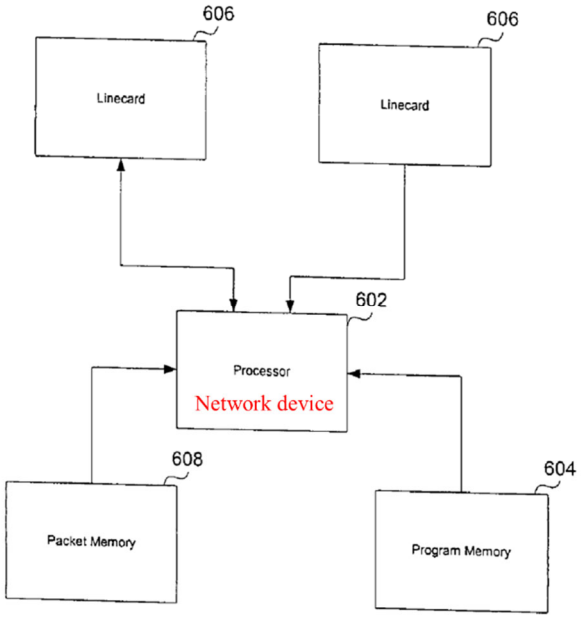


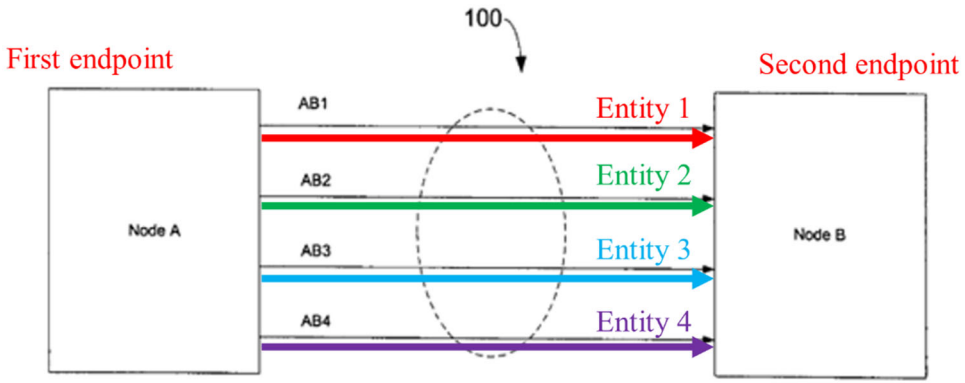
Fig. 4 (annotation added)

Sivabalan '928 at Para No. 0038 ("[FIG. 4 is a flow chart describing steps of component link operation according to the first mode as provided by embodiments of the present invention. At step 402, node A receives a request to establish a Traffic Engineering LSP using link

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 162 of 698

No.	'821 Patent Claim 1	Sivabalan '928
		bundle 100. Node A identifies a component link with sufficient bandwidth to accommodate the requirements of this Traffic Engineering LSP. At step 404, this component link is designated as the primary component link for the LSP. At step 406, node A identifies another component link with sufficient bandwidth to support the LSP. This component link is designated at step 408 as the secondary component link for the LSP. Information about the primary and secondary component links is programmed into the label forwarding table entry or entries corresponding to the LSP. At step 410, packets belonging to the LSP are carried on the primary component link.]")  FIG. 6 Fig. 6 (annotation added)

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 163 of 698

No.	'821 Patent Claim 1	Sivabalan '928
		Sivabalan '928 at Para No. 0051 ("[FIG. 6 depicts a network device 600 that may be used to implement, e.g., nodes A or B and/or perform any of the steps of FIGS. 4-5.]")
1[a]	providing a plurality of multi protocol label switching (MPLS) transport entities between a first endpoint and a second endpoint;	Sivabalan '928 discloses providing a plurality of multi protocol label switching (MPLS) transport entities between a first endpoint and a second endpoint. For example, Sivabalan '928 discloses that in MPLS traffic engineering, a link bundle combines multiple links between a pair of nodes and is referred to as a "component link." Sivabalan '928 at Para No. 0005 ("[To avoid the need to advertise information about each link between a pair of nodes, the concept of a "link bundle" has been developed to support MPLS Traffic Engineering. A link bundle combines multiple links between a pair of nodes. The links that make up a bundle are then referred to as "component links." Link bundling improves routing scalability by reducing the amount of information that has to be advertised by link state routing protocols to support signaling and computation of MPLS Traffic Engineering LSPs. Information about link bundles including bandwidth information is flooded throughout the network while information about component links is maintained only at the directly connected nodes.]")  FIG. 1 Fig. 1 (annotation added)

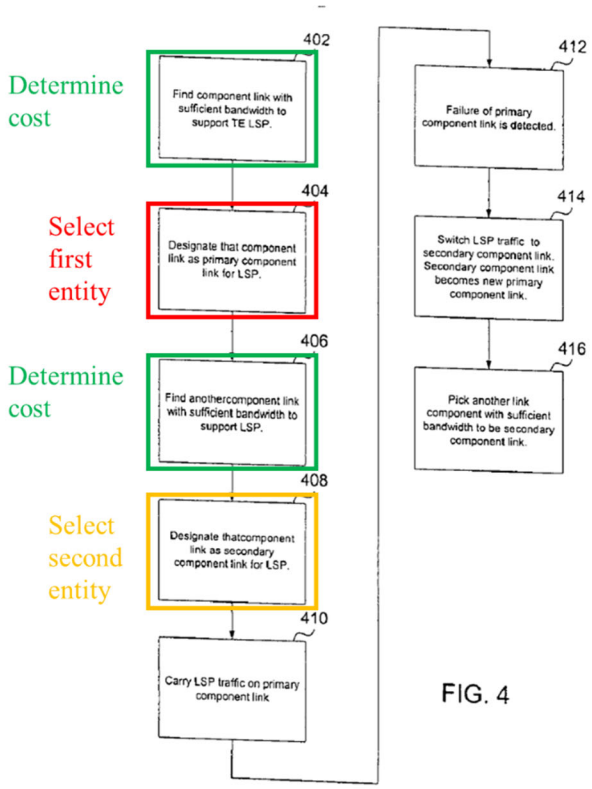
Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 164 of 698

No.	'821 Patent Claim 1	Sivabalan '928
		Sivabalan '928 at Para No. 0033 ("[FIG. 1. depicts a representative link bundle scenario according to one embodiment of the present invention. Link bundle 100 interconnects nodes A and B. Link bundle 100 includes components links AB1, AB2, AB3, and AB4. Node A propagates information about the aggregate characteristics of link bundle 100 including the available aggregate bandwidth to other nodes in the network in accordance with the relevant link state routing protocols underlying MPLS Traffic Engineering operation. Information about the component links AB1, AB2, AB3, and AB4 is not distributed in this way but only maintained locally to node A. Node B has similar knowledge of the links extending in the reverse direction. Information requesting the establishment of an LSP through nodes A and B will thus specify link bundle 100 but not the constituent component links.]") Sivabalan '928 at Claim 1 ("1. A method for operating a first node in a label switched network, said method comprising: establishing a link bundle comprising a plurality of component links between said first node and a second node; upon establishment of an LSP including said first node and said second node, selecting a first component link of said link bundle as a primary component link to assign to said LSP; selecting one or more component links of said link bundle other than said first component link as one or more secondary component links to assign to said LSP; sending traffic of said LSP over said primary component link until a failure of said primary component link; and after failure of said primary component link, sending traffic of said LSP over said one or more secondary component links instead of said primary component link.")
1[b]	determining an overall cost for each entity pair of said plurality of entities;	Sivabalan '928 discloses determining an overall cost for each entity pair of said plurality of entities. For example, Sivabalan '928 discloses that sufficient bandwidth is reserved on each link and component link in the MPLS Traffic Engineering LSP to meet a bandwidth guarantee. Sivabalan further discloses that an LSP placed on a link bundle can operate in one of two different modes: in the first mode a single secondary link is assigned as a backup for the primary component link and in the second mode, one primary component link and one or more secondary component links are assigned to the LSP. Sivabalan '928 further discloses that when node A receives a request to establish a traffic engineering LSP using a link bundle

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 165 of 698

No.	'821 Patent Claim 1	Sivabalan '928
		with the pair of nodes, it identifies primary and secondary/backup component links within the link bundle with sufficient bandwidth to accommodate the bandwidth requirement of the LSP and designates it as a primary link for LSP. At least based upon Orckit's apparent understanding of the claim scope as reflected in its Infringement Disclosures, this limitation is met. To the extent that Sivabalan '928 is found to not meet this limitation, determining an overall cost for each entity pair of said plurality of entities would have been obvious to a person having ordinary skill in the art, as explained below. Sivabalan '928 at Para No. 0035-0037 ("[According to embodiments of the present invention, an LSP placed on a link bundle can operate in one of two different modes: In the first mode, a single secondary component link is assigned as backup for the primary component link. Sufficient bandwidth to accommodate the LSP is reserved on this single secondary component link. The amount of available bandwidth on the secondary component link is reduced accordingly. FIG. 2 illustrates an example of this first mode where an MPLS Traffic Engineering LSP X employs link bundle 100 between nodes A and B. Node A ties LSP X to a primary component link AB1 and a secondary component link AB3. Further details of this first mode of operation will be described with reference to FIG. 4. In a second mode, one or more secondary component links are assigned to LSP X. Bandwidth is, however, not reserved on these secondary component links. This mode of operation is illustrated in FIG. 3 where LSP X is assigned to a primary component link AB1 and three secondary component links: AB2, AB3, and AB4. Further details of this second mode of operation will be described with reference to FIG. 5.]")

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 166 of 698

No.	'821 Patent Claim 1	Sivabalan '928
		 FIG. 4 Fig. 4 (annotation added) Sivabalan '928 at Para No. 0038 ("[FIG. 4 is a flow chart describing steps of component link operation according to the first mode as provided by embodiments of the present invention. At step 402, node A receives a request to establish a Traffic Engineering LSP using link bundle 100. Node A identifies a component link with sufficient bandwidth to accommodate the requirements of this Traffic Engineering LSP. At step 404, this component link is

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 167 of 698

No.	'821 Patent Claim 1	Sivabalan '928
		designated as the primary component link for the LSP. At step 406, node A identifies another component link with sufficient bandwidth to support the LSP. This component link is designated at step 408 as the secondary component link for the LSP. Information about the primary and secondary component links is programmed into the label forwarding table entry or entries corresponding to the LSP. At step 410, packets belonging to the LSP are carried on the primary component link.]") In addition, at least under the apparent claim scope alleged by Orckit's Infringement Disclosures, Sivabalan '928 in combination with (1) the knowledge of a person of ordinary skill in the art, alone or in further combination with (2) each (individually, as well as one or more together) of the references identified in element 1[b] of Exhibit E-2 renders the claim, including the present limitation, obvious. Below are examples of two such references. For example, Juniper IOS Guide discloses "Fate Sharing" where a Juniper router computes a primary and secondary path using CSPF, CSPF identifies all fate-sharing groups that are associated with the primary path by identifying all links that the primary path transverses and checks each link in the traffic engineering database against the compiled group list. Juniper IOS Guide further discloses that if the link is a member of a group, the cost of the link is increased by the cost of the group and higher the cost value. At least based upon Orckit's apparent understanding of the claim scope as reflected in its Infringement Disclosures, this limitation is met. Juniper IOS Guide at Page No. 73 ("[Fate Sharing Fate sharing allows you to create a database of information that CSPF uses to compute one or more backup paths to use in case the primary path becomes unstable. The database describes the relationships between elements of the network, such as routers and links. You can specify one or more elements within a group. Through fate sharing, you can configure backup paths that minimize the number of shared links and fiber paths with the primary paths as much as possible, to ensure that if a fiber is cut, the minimum amount of data is lost and a path still exists to the destination. For a backup path to work optimally, it must not share links or physical fiber paths with the primary path, ensuring that a single point of failure will not affect the primary and backup

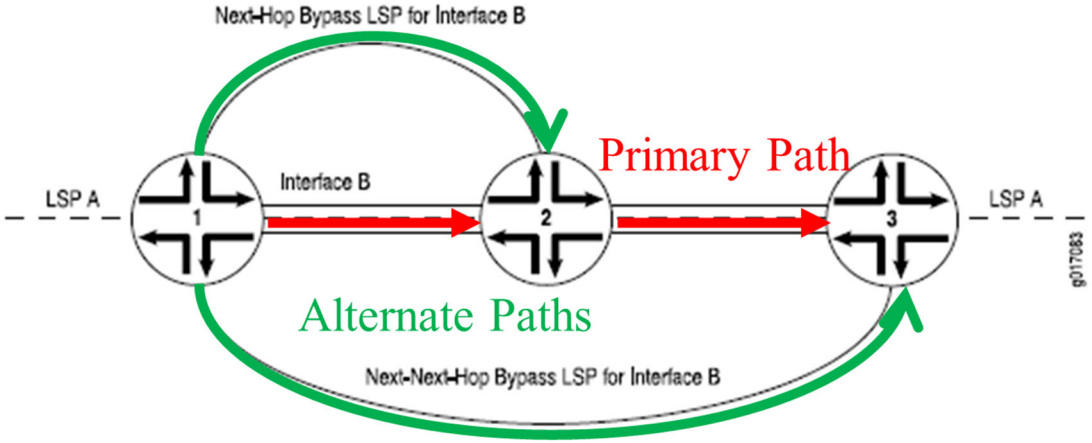
Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 168 of 698

No.	'821 Patent Claim 1	Sivabalan '928
		paths simultaneously. For more information about fate sharing, see the Junos OS Routing Protocols Configuration Guide.]") Juniper IOS Guide at Page No. 94 ("[Configuring the Ingress Router for MPLS-Signaled LSPs MPLS-signaled label-switched paths (LSPs) run from a specific ingress router to a specific egress router. For basic MPLS-signaled LSP function, you must configure the ingress router, but do not have to configure any other routers. To configure signaled LSPs, perform the following tasks on the ingress router: • Creating Named Paths on page 56 • Configuring Alternate Backup Paths Using Fate Sharing on page 58.]") Juniper IOS Guide at Page No. 96 ("[Configuring Alternate Backup Paths Using Fate Sharing You can create a database of information that Constrained Shortest Path First (CSPF) uses to compute one or more backup paths in case the primary path becomes unstable. The database describes the relationships between elements of the network, such as routers and links. Because these network elements share the same fate, this relationship is called fate sharing. You can configure backup paths that minimize the number of shared links and fiber paths with the primary paths as much as possible to ensure that, if a fiber is cut, the minimum amount of data is lost and a path still exists to the destination. For a backup path to work optimally, it must not share links or physical fiber paths with the primary path. This ensures that a single point of failure will not affect the primary and backup paths at the same time.]") Juniper IOS Guide at Page No. 97 ("[Fate-sharing groups contain three types of objects: • Point-to-point links—Identified by the IP addresses at each end of the link. Unnumbered point-to-point links are typically identified by borrowing IP addresses from other interfaces. Order is not important; from 1.2.3.4 to 1.2.3.5 and from 1.2.3.5 to 1.2.3.4 have the same meaning. • Non-point-to-point links—Include links on a LAN interface (such as Gigabit Ethernet interfaces) or nonbroadcast multiaccess (NBMA) interfaces (such as Asynchronous Transfer Mode [ATM] or Frame Relay). You identify these links by their individual interface address.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 169 of 698

No.	'821 Patent Claim 1	Sivabalan '928
		For example, if the LAN interface 192.168.200.0/24 has four routers attached to it, each router link is individually identified: from 192.168.200.1; # LAN interface of router 1 from 192.168.200.2; # LAN interface of router 2 from 192.168.200.3; # LAN interface of router 3 from 192.168.200.4; # LAN interface of router 4 You can list the addresses in any order. • A router node—Identified by its configured router ID. An object can be in any number of groups, and a group can contain any number of objects. Each group has a configurable cost attributed to it, which represents the level of impact this group has on CSPF computations. The higher the cost, the less likely a backup path will share with the primary path any objects in the group. The cost is directly comparable to traffic engineering metrics. By default, the cost is 1. Changing the fate-sharing database does not affect established LSPs until the next reoptimization of CSPF. The fate-sharing database does influence fast-reroute computations.]") Juniper IOS Guide at Page No. 97 ("[When CSPF computes the primary paths of an LSP (or secondary paths when the primary path is not active), it ignores the fate-sharing information. You always want to find the best possible path (least IGP cost) for the primary path. When CSPF computes a secondary path while the primary path (of the same LSP) is active, the following occurs: 1. CSPF identifies all fate-sharing groups that are associated with the primary path. CSPF does this by identifying all links and nodes that the primary path traverses and compiling group lists that contain at least one of the links or nodes. CSPF ignores the ingress and egress nodes in the search. 2. CSPF checks each link in the traffic engineering database against the compiled group list. If the link is a member of a group, the cost of the link is increased by the cost of the group. If a link is a member of multiple groups, all group costs are added together. 3. CSPF performs the check for every node in the traffic engineering database, except the ingress and egress node. Again, a node can belong to multiple groups, so costs are additive. 4. The router performs regular CSPF computation with the adjusted topology.]")

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 170 of 698

No.	'821 Patent Claim 1	Sivabalan '928
		Juniper IOS Guide at Page No. 330 ("In Figure 23 on page 292, both node protection and link protection are enabled on Interface B on Router 1. Both node protection and link protection are also enabled on LSP A, an LSP that traverses the link transiting Router 1, Router 2, and Router 3. If the link between Router 1 and Router 2 fails, traffic from LSP A is quickly switched to the next-hop bypass LSP generated by link protection. If Router 2 suffers a hardware or software failure, traffic from LSP A is switched to the next-next-hop bypass LSP generated by node protection.") Figure 23: Node Protection Creating a Next-Next-Hop Bypass LSP  Fig. 23 (annotation added). For example, Cisco IOS Guide discloses that the path an LSP uses is determined by the LSP resource requirements and network resources, such as bandwidth. Cisco IOS Guide further discloses the MPLS TE SRLG feature that enhances backup tunnel path selection so a backup tunnel can avoid using links that are in the same SRLG as the interfaces it is protecting. Thus, at least under the apparent claim scope alleged by Orckit's Infringement Disclosures, this limitation is met. Cisco IOS Guide at Page No. 253 ("[Path Option Selection for MPLS TE Tunnel LSPs

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 171 of 698

No.	'821 Patent Claim 1	Sivabalan '928
		This section contains the following topics about path option selection for MPLS TE Tunnel LSPs: • Constraint-Based Routing and Path Option Selection, page 5 • Tunnel Reoptimization and Path Option Selection, page 5 • Path Option Selection with Bandwidth Override, page 6 Constraint-Based Routing and Path Option Selection MPLS traffic engineering automatically establishes and maintains LSPs across the backbone by using the Resource Reservation Protocol (RSVP). The path that an LSP uses is determined by the LSP resource requirements and network resources, such as bandwidth. Traffic engineering tunnels are calculated at the LSP head based on a fit between required and available resources (constraint-based routing). Without the Path Option for Bandwidth Override feature, a TE tunnel establishes an LSP based on dynamic or explicit path options in order of preference. However, the bandwidth and other attributes configured on the TE tunnel allow the setup of an LSP only if LSP path options satisfy the constraints. If a path cannot be found that satisfies the configured path options, then the tunnel is not set up.]") Cisco IOS Guide at Page No. 639 ("[Bandwidth Protection NHOP and NNHOP backup tunnels can be used to provide bandwidth protection for rerouted LSPs. This is referred to as backup bandwidth. You can associate backup bandwidth with NHOP or NNHOP backup tunnels. This informs the router of the amount of backup bandwidth a particular backup tunnel can protect. When a router maps LSPs to backup tunnels, bandwidth protection ensures that an LSP uses a given backup tunnel only if there is sufficient backup bandwidth. The router selects which LSPs use which backup tunnels in order to provide maximum bandwidth protection. That is, the router determines the best way to map LSPs onto backup tunnels in order to maximize the number of LSPs that can be protected. For information about mapping tunnels and assigning backup bandwidth, see the "Backup Tunnel Selection Procedure" section on page 10. LSPs that have the "bandwidth protection desired" bit set have a higher right to select backup tunnels that provide bandwidth protection; that is, those LSPs can preempt other LSPs that do not have that bit set. For more information, see the "Prioritizing Which LSPs Obtain Backup Tunnels with Bandwidth Protection" section on page 8.]")

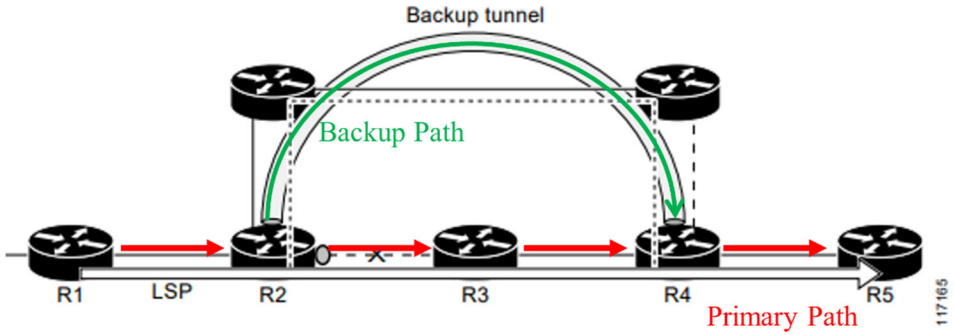
Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 172 of 698

No.	'821 Patent Claim 1	Sivabalan '928
		Cisco IOS Guide at Page No. 640 ("[Backup Tunnel Support Backup tunnel support has the following capabilities: • Backup Tunnels Can Terminate at the Next-Next Hop to Support FRR, page 6 • Multiple Backup Tunnels Can Protect the Same Interface, page 6 • Backup Tunnels Provide Scalability, page 7.]") Cisco IOS Guide at Page 640 ("[Multiple Backup Tunnels Can Protect the Same Interface There is no limit (except memory limitations) to the number of backup tunnels that can protect a given interface. In many topologies, support for node protection requires supporting multiple backup tunnels per protected interface. These backup tunnels can terminate at the same destination or at different destinations. That is, for a given protected interface, you can configure multiple NHOP or NNHOP backup tunnels. This allows redundancy and load balancing.]") Cisco IOS Guide at Page No. 603 ("[Traffic Engineering Tunnels MPLS TE lets you build label switched paths (LSPs) across your network for forwarding traffic. MPLS TE LSPs let the headend of a TE tunnel control the path its traffic takes to a particular destination. This method is more flexible than forwarding traffic based only on a destination address. Interarea tunnels allow you to do the following: • Build TE tunnels between areas (interarea tunnels) • Build TE tunnels that start and end in the same area, on multiple areas on a router (intra-area tunnels) Some tunnels are more important than others. For example, you may have tunnels carrying VoIP traffic and tunnels carrying data traffic that are competing for the same resources. MPLS TE allows you to have some tunnels preempt others. Each tunnel has a priority, and more-important tunnels take precedence over less-important tunnels.]") Cisco IOS Guide at Page No. 641 ("[Backup Bandwidth Protection

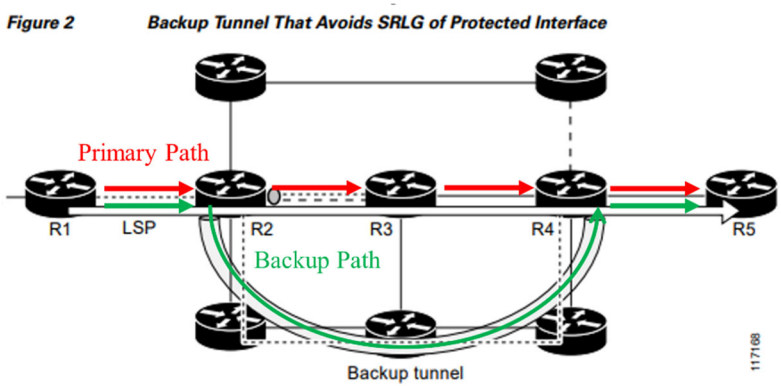
Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 173 of 698

No.	'821 Patent Claim 1	Sivabalan '928
		Backup bandwidth protection allows you to give LSPs carrying certain kinds of data (such as voice) priority for using backup tunnels.]") Cisco IOS Guide at Page No. 650 ("[For each protected link AB with a maximum reservable subpool value of n, there may be a path from node A to node B such that the difference between the maximum reservable global and the maximum reservable subpool is at least the value of n. If it is possible to find such paths for each link in the network, you can establish all the backup tunnels along such paths without any bandwidth reservations. If there is a single link failure, only one backup tunnel will use any link on its path. Because that path has at least n available bandwidth (in the global pool), assuming that marking and scheduling is configured to classify the subpool traffic into a priority queue, the subpool bandwidth is guaranteed. This approach allows sharing of the global pool bandwidth between backup tunnels protecting independent link failures. The backup tunnels are expected to be used for only a short period of time after a failure (until the headends of affected LSPs reroute those LSPs to other paths with available subpool bandwidth). The probability of multiple unrelated link failures is very small (in the absence of node or shared risk link group (SRLG) failures, which result in multiple link failures). Therefore, it is reasonable to assume that link failures are in practice independent with high probability. This "independent failure assumption" in combination with backup tunnels signaled without explicit bandwidth reservation enables efficient bandwidth sharing that yields substantial bandwidth savings.]") Cisco IOS Guide at Page No. 651 ("[A similar approach can be used for node and SRLG protection. However, the decision of where to put the backup tunnels is more complicated because both node and SRLG failures effectively result in the simultaneous failure of several links. Therefore, the backup tunnels protecting traffic traversing all affected links cannot be computed independently of each other.]") Cisco IOS Guide at Page No. 579 ("[MPLS Traffic Engineering Shared Risk Link Groups SRLGs refer to situations where links in a network share a common fiber (or a common physical attribute). If one link fails, other links in the group may fail too. Links in the group

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 174 of 698

No.	'821 Patent Claim 1	Sivabalan '928
		have a shared risk. Backup tunnels should avoid using links in the same SRLG as interfaces they are protecting. Otherwise, when the protected link fails the backup tunnel fails too. *Figure 1 Backup Tunnel in the Same SRLG as the Interface It Is Protecting*  Fig. 1 (annotation added) Figure 1 shows a primary label-switched path (LSP) from router R1 to router R5. The LSP protects against the failure of the R2-R3 link at R2 via a backup tunnel to R4. If the R2-R3 link fails, link protection reroutes the LSP along the backup tunnel. However, the R2-R3 link and one of the backup tunnel links are in the same SRLG. So if the R2-R3 link fails, the backup tunnel may fail too. The MPLS TE SRLG feature enhances backup tunnel path selection so a backup tunnel can avoid using links that are in the same SRLG as the interfaces it is protecting. There are two ways for a backup tunnel to avoid the SRLGs of its protected interface: • The router does not create the backup tunnel unless it avoids SRLGs of the protected interface.]") Cisco IOS Guide at Page No. 580 ("[Open Shortest Path First (OSPF) and Intermediate System-to-Intermediate System (IS-IS) flood the SRLG membership information (including other TE link attributes such as bandwidth availability and affinity) so that all routers in the network have the SRLG information for each link. With this topology information, routers can compute backup tunnel paths that exclude links having SRLGs in common with their protected interfaces. As shown in Figure 2, the backup tunnel avoids the link between R2 and R3, which shares an SRLG with the protected interface.]")

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 175 of 698

No.	'821 Patent Claim 1	Sivabalan '928
		 Fig. 2 (annotation added) Cisco IOS Guide at Page No. 587 ("[Step 4 show mpls traffic-eng topology Use this command to show the SRLG link membership flooded via the Interior Gateway Protocol (IGP).]") Cisco IOS Guide at Page 632 ("[IS-IS—Intermediate System-to-Intermediate System. Link-state hierarchical routing protocol that calls for intermediate system (IS) routers to exchange routing information based on a single metric to determine network topology.]") Cisco IOS Guide at Page No. 633 ("[OSPF—Open Shortest Path First. A link-state hierarchical Interior Gateway Protocol routing algorithm, derived from the IS-IS protocol. OSPF features include least-cost routing, multipath routing, and load balancing.]") Cisco IOS Guide at Page No. 581 ("[Autotunnel Backup for MPLS TE SRLGs Autotunnel backup is the ability of routers to create backup tunnels automatically. Therefore, you do not need to preconfigure each backup tunnel and then assign the backup tunnel to the protected interface. Only automatically created backup tunnels can avoid SRLGs or their protected interfaces.]")

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 176 of 698

No.	'821 Patent Claim 1	Sivabalan '928
		Cisco IOS Guide at Page No. 583 ("[Configuring the Routers That Automatically Create Backup Tunnels to Avoid MPLS TE SRLGs of Their Protected Interfaces Perform the following task to configure routers that automatically create backup tunnels to avoid MPLS TE SRLGs of their protected interfaces. Backup tunnels provide link protection by rerouting traffic to the next hop bypassing failed links or in this instance by avoiding SRLGs.]") Cisco IOS Guide at Page No. 644 ("[Backup Tunnel Selection Procedure When an LSP is signaled, each node along the LSP path that provides FRR protection for the LSP selects a backup tunnel for the LSP to use if either of the following events occurs: • The link to the next hop fails. • The next hop fails. By having the node select the backup tunnel for an LSP before a failure occurs, the LSP can be rerouted onto the backup tunnel quickly if there is a failure.]") Cisco IOS Guide at Page 646 ("[NHOP Versus NNHOP Backup Tunnels More than one backup tunnel can protect a given LSP, where one backup tunnel terminates at the LSP's NNHOP, and the other terminates at the LSP's NHOP. In this case, the router chooses the backup tunnel that terminates at the NNHOP (that is, FRR prefers NNHOP over NHOP backup tunnels).Table 1 lists the tunnel selection priorities. The first choice is an NNHOP backup tunnel that acquires its bandwidth from a subpool or global pool, and has limited bandwidth. If there is no such backup tunnel, the next choice (2) is a next-next hop backup tunnel that acquires a limited amount of bandwidth from any pool. The preferences go from 1 (best) to 8 (worst), where choice 3 is for an NNHOP backup tunnel with an unlimited amount of subpool or global-pool bandwidth.]")

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 177 of 698

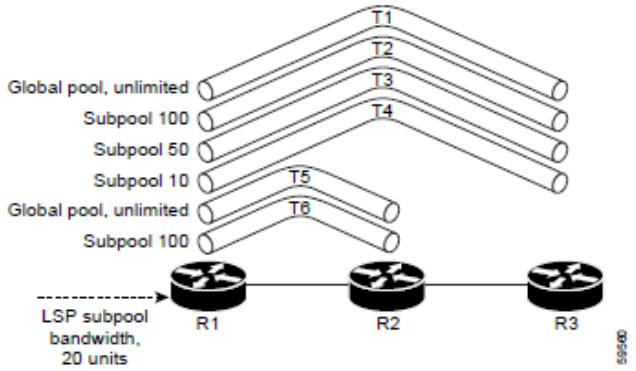
No.	'821 Patent Claim 1	Sivabalan '928

Table 1 Tunnel Selection Priorities

Preference	Backup Tunnel Destination	Bandwidth Pool	Bandwidth Amount
1 (Best)	NNHOP	Subpool or global pool	Limited
2	NNHOP	Any	Limited
3	NNHOP	Subpool or global pool	Unlimited
4	NNHOP	Any	Unlimited
5	NHOP	Subpool or global pool	Limited
6	NHOP	Any	Limited
7	NHOP	Subpool or global pool	Unlimited
8 (Worst)	NHOP	Any	Unlimited

Cisco IOS Guide at Page No. 647-648 ("[Figure 7 shows an example of the backup tunnel selection procedure based on the designated amount of global pool and subpool bandwidth currently available.

Figure 7 Choosing from Among Multiple Backup Tunnels

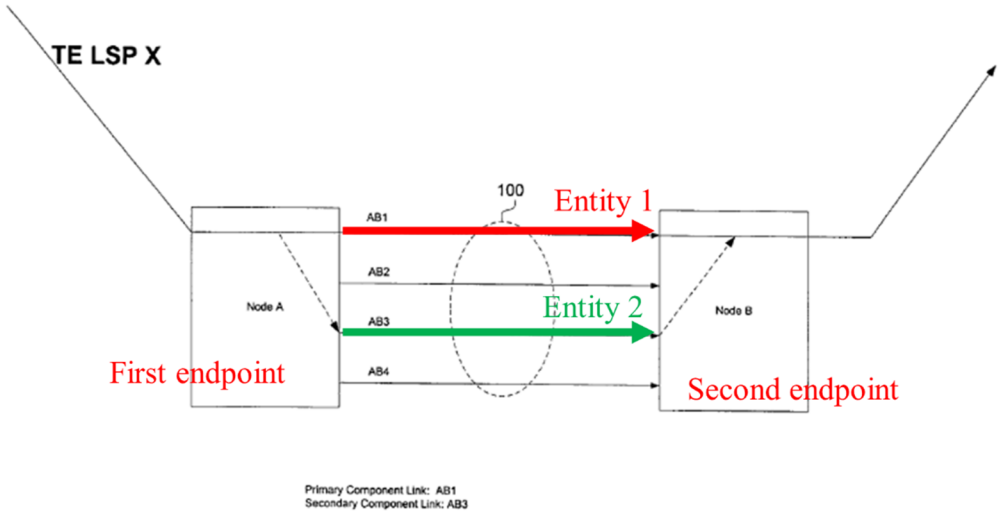


In this example, an LSP requires 20 units (kilobits per second) of sub-pool backup bandwidth. The best backup tunnel is selected as follows:
1. Backup tunnels T1 through T4 are considered first because they terminate at the NNHOP.
2. Tunnel T4 is eliminated because it has only ten units of sub-pool backup bandwidth.
3. Tunnel T1 is eliminated because it protects only LSPs using global-pool bandwidth.

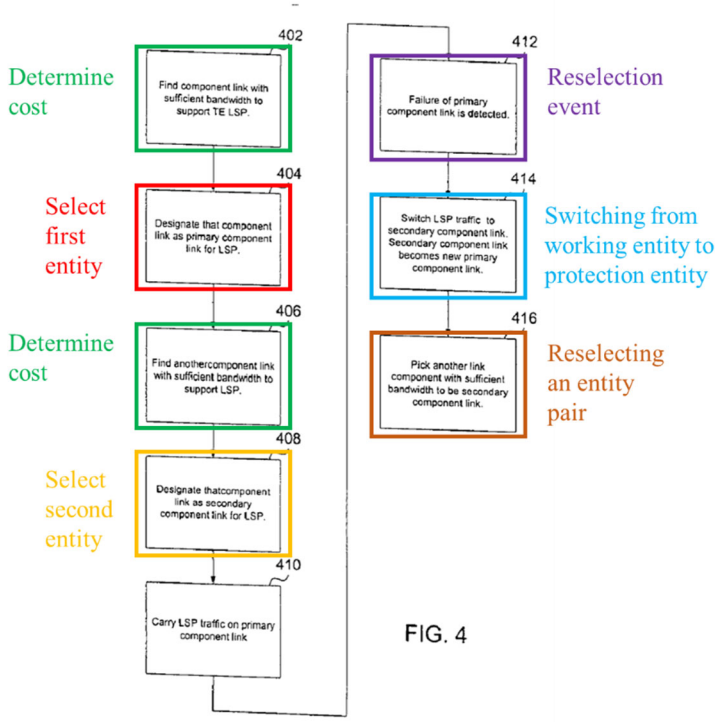
Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 178 of 698

No.	'821 Patent Claim 1	Sivabalan '928
		4. Tunnel T3 is chosen over T2 because, although both have sufficient backup bandwidth, T3 has the least backup bandwidth available (leaving the most backup bandwidth available on T2). 5. Tunnels T5 and T6 need not be considered because they terminate at an NHOP, and therefore are less desirable than T3, which terminates at an NNHOP.]")
1[c]	selecting an entity pair from said plurality of transport entities based at least in part upon said overall cost; and	Sivabalan '928 discloses selecting an entity pair from said plurality of transport entities based at least in part upon said overall cost. For example, Sivabalan '928 discloses that node A identifies a primary component link with sufficient bandwidth to accommodate the requirements of this Traffic Engineering LSP and identifies a secondary component link with sufficient bandwidth to support the LSP. At least based upon Orckit's apparent understanding of this term as reflected in its Infringement Disclosures, this limitation is met. Sivabalan '928 at Para No. 0036 ("In the first mode, a single secondary component link is assigned as backup for the primary component link. Sufficient bandwidth to accommodate the LSP is reserved on this single secondary component link. The amount of available bandwidth on the secondary component link is reduced accordingly. FIG. 2 illustrates an example of this first mode where an MPLS Traffic Engineering LSP X employs link bundle 100 between nodes A and B. Node A ties LSP X to a primary component link AB1 and a secondary component link AB3.")

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 179 of 698

No.	'821 Patent Claim 1	Sivabalan '928
		 FIG. 2 Fig. 2 (annotation added) *See Supra at* 1[b].
1[d]	if an entity pair reselection event occurs, reselecting said entity pair from the group consisting of said entity pair and a replacement entity pair comprising at least one entity distinct from the	Sivabalan '928 discloses if an entity pair reselection event occurs, reselecting said entity pair from the group consisting of said entity pair and a replacement entity pair comprising at least one entity distinct from the entities of said entity pair. For example, Sivabalan '928 discloses when failure of the primary component link is detected, node A switches over the LSP traffic to another primary component link and also selects another secondary component link. Sivabalan '928 further discloses that at least one of these component links will differ from the link that was detected to fail.

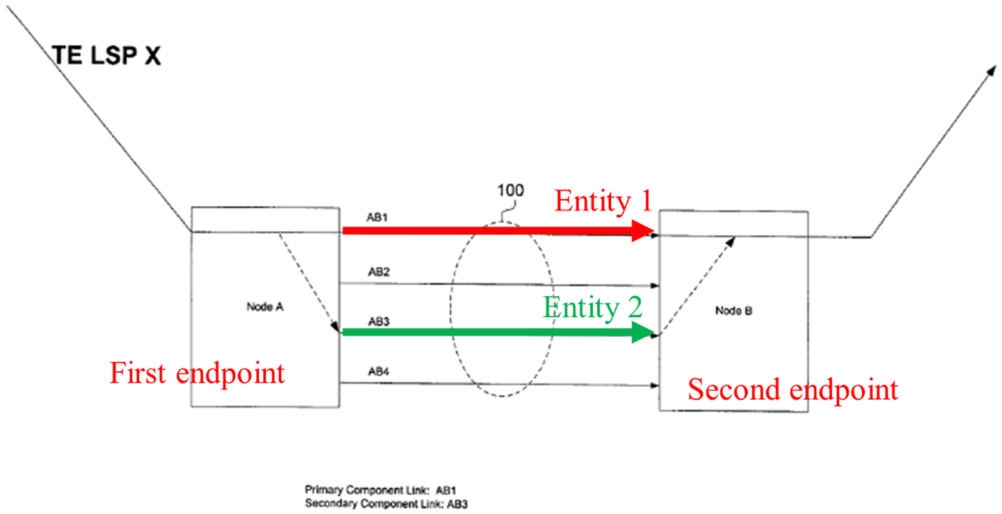
Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 180 of 698

No.	'821 Patent Claim 1	Sivabalan '928
	entities of said entity pair,	Sivabalan '928 at Para No. 0039 ("[At step 412, failure of the primary component link is detected. The response is a switchover of the LSP traffic to the secondary component link at step 414. Quality of service is maintained since the secondary component link also has sufficient guaranteed bandwidth. There is no signaling to the head-end to implement the switch-over. The secondary component link now becomes the primary component link. At step 416, node A selects another component link in the same link bundle to be the new secondary component link.]")  Fig. 4 (annotation added)

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 181 of 698

No.	'821 Patent Claim 1	Sivabalan '928
		Sivabalan '928 at Para No. 0038 ("[FIG. 4 is a flow chart describing steps of component link operation according to the first mode as provided by embodiments of the present invention. At step 402, node A receives a request to establish a Traffic Engineering LSP using link bundle 100. Node A identifies a component link with sufficient bandwidth to accommodate the requirements of this Traffic Engineering LSP. At step 404, this component link is designated as the primary component link for the LSP. At step 406, node A identifies another component link with sufficient bandwidth to support the LSP. This component link is designated at step 408 as the secondary component link for the LSP. Information about the primary and secondary component links is programmed into the label forwarding table entry or entries corresponding to the LSP. At step 410, packets belonging to the LSP are carried on the primary component link.]") Sivabalan '928 at Para No. 0039 ("At step 412, failure of the primary component link is detected. The response is a switchover of the LSP traffic to the secondary component link at step 414. Quality of service is maintained since the secondary component link also has sufficient guaranteed bandwidth. There is no signaling to the head-end to implement the switch-over. The secondary component link now becomes the primary component link. At step 416, node A selects another component link in the same link bundle to be the new secondary component link.") Sivabalan '928 at Para No. 0036 ("In the first mode, a single secondary component link is assigned as backup for the primary component link. Sufficient bandwidth to accommodate the LSP is reserved on this single secondary component link. The amount of available bandwidth on the secondary component link is reduced accordingly. FIG. 2 illustrates an example of this first mode where an MPLS Traffic Engineering LSP X employs link bundle 100 between nodes A and B. Node A ties LSP X to a primary component link AB1 and a secondary component link AB3.")

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 182 of 698

No.	'821 Patent Claim 1	Sivabalan '928
		 FIG. 2 Fig. 2 (annotation added) Sivabalan '928 at Claim 1 ("1. A method for operating a first node in a label switched network, said method comprising: establishing a link bundle comprising a plurality of component links between said first node and a second node; upon establishment of an LSP including said first node and said second node, selecting a first component link of said link bundle as a primary component link to assign to said LSP; selecting one or more component links of said link bundle other than said first component link as one or more secondary component links to assign to said LSP; sending traffic of said LSP over said primary component link until a failure of said primary component link; and after failure of said primary component link, sending traffic of said LSP over said one or more secondary component links instead of said primary component link.")

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 183 of 698

No.	'821 Patent Claim 1	Sivabalan '928
1[e]	wherein said entity pair reselection event is selected from a group consisting of adding an entity to said plurality of transport entities, removing an entity from said plurality of transport entities, an operational status change for one of said plurality of transport entities, and a change in overall cost for one of said plurality of transport entities.	Sivabalan '928 discloses "entity pair reselection event is selected from a group consisting of" & "an operational status change for one of said plurality of transport entities". For example, Sivabalan '928 discloses reselecting component links in response to a failure of the determined primary component link. A person of ordinary skill in the art would understand that the primary and secondary component links could also be reselected in response to a change in the overall cost for one of said plurality of transport entities. At least based upon Orckit's apparent understanding of the claim scope as reflected in its Infringement Disclosures, this limitation is met. To the extent that the IETF MPLS-TP System is found to not meet this limitation, wherein said entity pair reselection event is selected from a group consisting of adding an entity to said plurality of transport entities, removing an entity from said plurality of transport entities, an operational status change for one of said plurality of transport entities, and a change in overall cost for one of said plurality of transport entities would have been obvious to a person having ordinary skill in the art, as explained below. Sivabalan '928 at Para No. 0039 ("[At step 412, failure of the primary component link is detected. The response is a switchover of the LSP traffic to the secondary component link at step 414. Quality of service is maintained since the secondary component link also has sufficient guaranteed bandwidth. There is no signaling to the head-end to implement the switch-over. The secondary component link now becomes the primary component link. At step 416, node A selects another component link in the same link bundle to be the new secondary component link.]") In addition, at least under the apparent claim scope alleged by Orckit's Infringement Disclosures, Sivabalan '928 in combination with (1) the knowledge of a person of ordinary skill in the art, alone or in further combination with (2) each (individually, as well as one or more together) of the references identified in element 1[e] of Exhibit E-2 renders the claim, including the present limitation, obvious. Below are examples of two such references. For example, Juniper IOS Guide discloses "Fate Sharing" where a Juniper router computes a primary and secondary path using CSPF, CSPF identifies all fate-sharing groups that are associated with the primary path by identifying all links that the primary path transverses and

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 184 of 698

No.	'821 Patent Claim 1	Sivabalan '928
		checks each link in the traffic engineering database against the compiled group list. Juniper IOS Guide further discloses that if the link is a member of a group, the cost of the link is increased by the cost of the group and higher the cost value. At least based upon Orckit's apparent understanding of the claim scope as reflected in its Infringement Disclosures, this limitation is met. Juniper IOS Guide at Page No. 73 ("[Fate Sharing Fate sharing allows you to create a database of information that CSPF uses to compute one or more backup paths to use in case the primary path becomes unstable. The database describes the relationships between elements of the network, such as routers and links. You can specify one or more elements within a group. Through fate sharing, you can configure backup paths that minimize the number of shared links and fiber paths with the primary paths as much as possible, to ensure that if a fiber is cut, the minimum amount of data is lost and a path still exists to the destination. For a backup path to work optimally, it must not share links or physical fiber paths with the primary path, ensuring that a single point of failure will not affect the primary and backup paths simultaneously. For more information about fate sharing, see the Junos OS Routing Protocols Configuration Guide.]") Juniper IOS Guide at Page No. 94 ("[Configuring the Ingress Router for MPLS-Signaled LSPs MPLS-signaled label-switched paths (LSPs) run from a specific ingress router to a specific egress router. For basic MPLS-signaled LSP function, you must configure the ingress router, but do not have to configure any other routers. To configure signaled LSPs, perform the following tasks on the ingress router: • Creating Named Paths on page 56 • Configuring Alternate Backup Paths Using Fate Sharing on page 58.]") Juniper IOS Guide at Page No. 96 ("[Configuring Alternate Backup Paths Using Fate Sharing You can create a database of information that Constrained Shortest Path First (CSPF) uses to compute one or more backup paths in case the primary path becomes unstable. The database

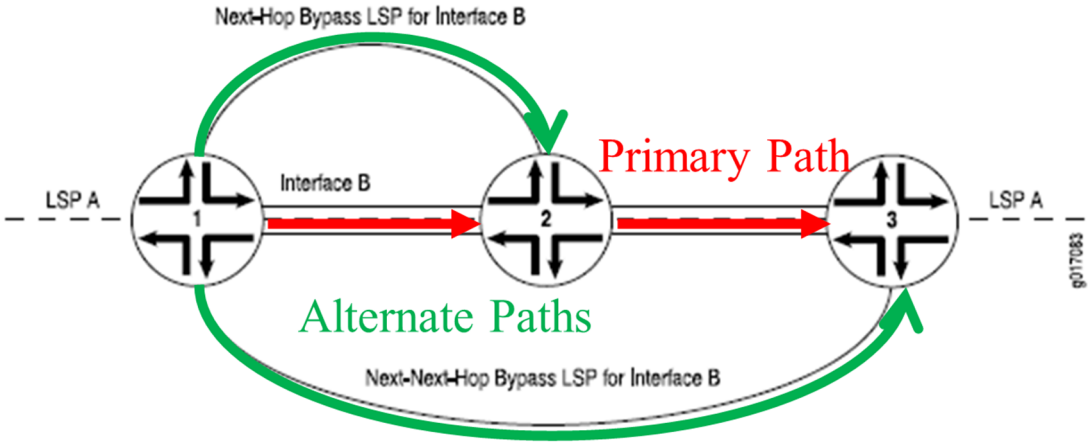
Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 185 of 698

No.	'821 Patent Claim 1	Sivabalan '928
		describes the relationships between elements of the network, such as routers and links. Because these network elements share the same fate, this relationship is called fate sharing. You can configure backup paths that minimize the number of shared links and fiber paths with the primary paths as much as possible to ensure that, if a fiber is cut, the minimum amount of data is lost and a path still exists to the destination. For a backup path to work optimally, it must not share links or physical fiber paths with the primary path. This ensures that a single point of failure will not affect the primary and backup paths at the same time.]") Juniper IOS Guide at Page No. 97 ("[Fate-sharing groups contain three types of objects: • Point-to-point links—Identified by the IP addresses at each end of the link. Unnumbered point-to-point links are typically identified by borrowing IP addresses from other interfaces. Order is not important; from 1.2.3.4 to 1.2.3.5 and from 1.2.3.5 to 1.2.3.4 have the same meaning. • Non-point-to-point links—Include links on a LAN interface (such as Gigabit Ethernet interfaces) or nonbroadcast multiaccess (NBMA) interfaces (such as Asynchronous Transfer Mode [ATM] or Frame Relay). You identify these links by their individual interface address. For example, if the LAN interface 192.168.200.0/24 has four routers attached to it, each router link is individually identified: from 192.168.200.1; # LAN interface of router 1 from 192.168.200.2; # LAN interface of router 2 from 192.168.200.3; # LAN interface of router 3 from 192.168.200.4; # LAN interface of router 4 You can list the addresses in any order. • A router node—Identified by its configured router ID. An object can be in any number of groups, and a group can contain any number of objects. Each group has a configurable cost attributed to it, which represents the level of impact this group has on CSPF computations. The higher the cost, the less likely a backup path will share with the primary path any objects in the group. The cost is directly comparable to traffic engineering metrics. By default, the cost is 1. Changing the fate-sharing database does not affect established LSPs until the next reoptimization of CSPF. The fate-sharing database does influence fast-reroute computations.]")

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 186 of 698

No.	'821 Patent Claim 1	Sivabalan '928
		Juniper IOS Guide at Page No. 97 ("[When CSPF computes the primary paths of an LSP (or secondary paths when the primary path is not active), it ignores the fate-sharing information. You always want to find the best possible path (least IGP cost) for the primary path. When CSPF computes a secondary path while the primary path (of the same LSP) is active, the following occurs: 1. CSPF identifies all fate-sharing groups that are associated with the primary path. CSPF does this by identifying all links and nodes that the primary path traverses and compiling group lists that contain at least one of the links or nodes. CSPF ignores the ingress and egress nodes in the search. 2. CSPF checks each link in the traffic engineering database against the compiled group list. If the link is a member of a group, the cost of the link is increased by the cost of the group. If a link is a member of multiple groups, all group costs are added together. 3. CSPF performs the check for every node in the traffic engineering database, except the ingress and egress node. Again, a node can belong to multiple groups, so costs are additive. 4. The router performs regular CSPF computation with the adjusted topology.]") Juniper IOS Guide at Page No. 330 ("In Figure 23 on page 292, both node protection and link protection are enabled on Interface B on Router 1. Both node protection and link protection are also enabled on LSP A, an LSP that traverses the link transiting Router 1, Router 2, and Router 3. If the link between Router 1 and Router 2 fails, traffic from LSP A is quickly switched to the next-hop bypass LSP generated by link protection. If Router 2 suffers a hardware or software failure, traffic from LSP A is switched to the next-next-hop bypass LSP generated by node protection.")

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 187 of 698

No.	'821 Patent Claim 1	Sivabalan '928
		Figure 23: Node Protection Creating a Next-Next-Hop Bypass LSP  Fig. 23 (annotation added). For example, Cisco IOS Guide discloses that the path an LSP uses is determined by the LSP resource requirements and network resources, such as bandwidth. Cisco IOS Guide further discloses the MPLS TE SRLG feature that enhances backup tunnel path selection so a backup tunnel can avoid using links that are in the same SRLG as the interfaces it is protecting. Thus, at least under the apparent claim scope alleged by Orckit's Infringement Disclosures, this limitation is met. Cisco IOS Guide at Page No. 253 ("[Path Option Selection for MPLS TE Tunnel LSPs This section contains the following topics about path option selection for MPLS TE Tunnel LSPs: • Constraint-Based Routing and Path Option Selection, page 5 • Tunnel Reoptimization and Path Option Selection, page 5 • Path Option Selection with Bandwidth Override, page 6 Constraint-Based Routing and Path Option Selection MPLS traffic engineering automatically establishes and maintains LSPs across the backbone by using the Resource Reservation Protocol (RSVP). The path that an LSP uses is determined

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 188 of 698

No.	'821 Patent Claim 1	Sivabalan '928
		by the LSP resource requirements and network resources, such as bandwidth. Traffic engineering tunnels are calculated at the LSP head based on a fit between required and available resources (constraint-based routing). Without the Path Option for Bandwidth Override feature, a TE tunnel establishes an LSP based on dynamic or explicit path options in order of preference. However, the bandwidth and other attributes configured on the TE tunnel allow the setup of an LSP only if LSP path options satisfy the constraints. If a path cannot be found that satisfies the configured path options, then the tunnel is not set up.]") Cisco IOS Guide at Page No. 639 ("[Bandwidth Protection NHOP and NNHOP backup tunnels can be used to provide bandwidth protection for rerouted LSPs. This is referred to as backup bandwidth. You can associate backup bandwidth with NHOP or NNHOP backup tunnels. This informs the router of the amount of backup bandwidth a particular backup tunnel can protect. When a router maps LSPs to backup tunnels, bandwidth protection ensures that an LSP uses a given backup tunnel only if there is sufficient backup bandwidth. The router selects which LSPs use which backup tunnels in order to provide maximum bandwidth protection. That is, the router determines the best way to map LSPs onto backup tunnels in order to maximize the number of LSPs that can be protected. For information about mapping tunnels and assigning backup bandwidth, see the "Backup Tunnel Selection Procedure" section on page 10. LSPs that have the "bandwidth protection desired" bit set have a higher right to select backup tunnels that provide bandwidth protection; that is, those LSPs can preempt other LSPs that do not have that bit set. For more information, see the "Prioritizing Which LSPs Obtain Backup Tunnels with Bandwidth Protection" section on page 8.]") Cisco IOS Guide at Page No. 640 ("[Backup Tunnel Support Backup tunnel support has the following capabilities: • Backup Tunnels Can Terminate at the Next-Next Hop to Support FRR, page 6 • Multiple Backup Tunnels Can Protect the Same Interface, page 6 • Backup Tunnels Provide Scalability, page 7.]") Cisco IOS Guide at Page 640 ("[Multiple Backup Tunnels Can Protect the Same Interface

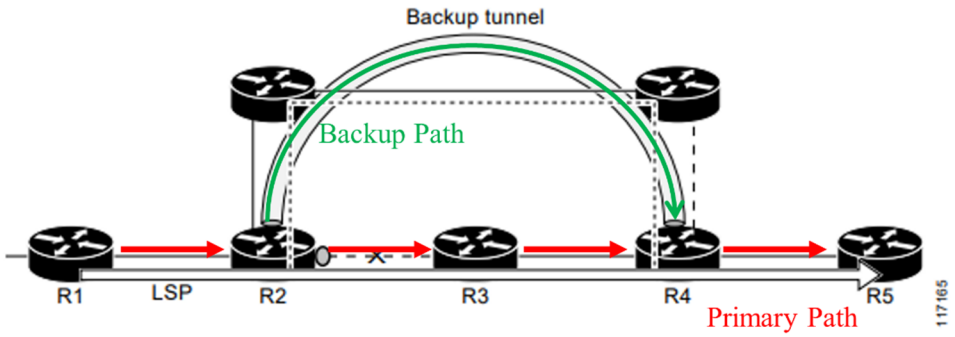
Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 189 of 698

No.	'821 Patent Claim 1	Sivabalan '928
		There is no limit (except memory limitations) to the number of backup tunnels that can protect a given interface. In many topologies, support for node protection requires supporting multiple backup tunnels per protected interface. These backup tunnels can terminate at the same destination or at different destinations. That is, for a given protected interface, you can configure multiple NHOP or NNHOP backup tunnels. This allows redundancy and load balancing.]") Cisco IOS Guide at Page No. 603 ("[Traffic Engineering Tunnels MPLS TE lets you build label switched paths (LSPs) across your network for forwarding traffic. MPLS TE LSPs let the headend of a TE tunnel control the path its traffic takes to a particular destination. This method is more flexible than forwarding traffic based only on a destination address. Interarea tunnels allow you to do the following: • Build TE tunnels between areas (interarea tunnels) • Build TE tunnels that start and end in the same area, on multiple areas on a router (intra-area tunnels) Some tunnels are more important than others. For example, you may have tunnels carrying VoIP traffic and tunnels carrying data traffic that are competing for the same resources. MPLS TE allows you to have some tunnels preempt others. Each tunnel has a priority, and more-important tunnels take precedence over less-important tunnels.]") Cisco IOS Guide at Page No. 641 ("[Backup Bandwidth Protection Backup bandwidth protection allows you to give LSPs carrying certain kinds of data (such as voice) priority for using backup tunnels.]") Cisco IOS Guide at Page No. 650 ("[For each protected link AB with a maximum reservable subpool value of n, there may be a path from node A to node B such that the difference between the maximum reservable global and the maximum reservable subpool is at least the value of n. If it is possible to find such paths for each link in the network, you can establish all the backup tunnels along such paths without any bandwidth reservations. If there is a single link failure, only one backup tunnel will use any link on its path. Because that path has at least

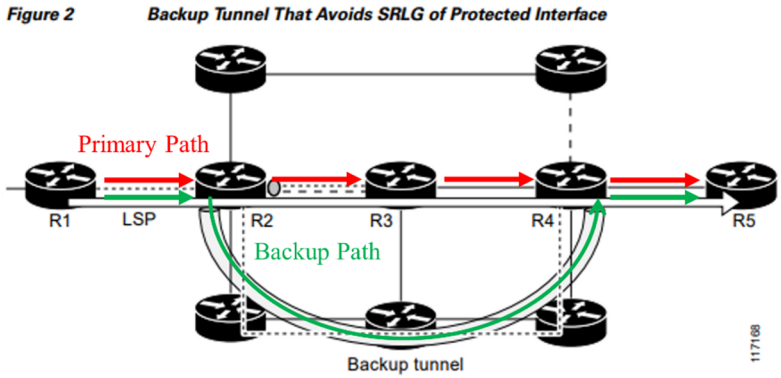
Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 190 of 698

No.	'821 Patent Claim 1	Sivabalan '928
		n available bandwidth (in the global pool), assuming that marking and scheduling is configured to classify the subpool traffic into a priority queue, the subpool bandwidth is guaranteed. This approach allows sharing of the global pool bandwidth between backup tunnels protecting independent link failures. The backup tunnels are expected to be used for only a short period of time after a failure (until the headends of affected LSPs reroute those LSPs to other paths with available subpool bandwidth). The probability of multiple unrelated link failures is very small (in the absence of node or shared risk link group (SRLG) failures, which result in multiple link failures). Therefore, it is reasonable to assume that link failures are in practice independent with high probability. This "independent failure assumption" in combination with backup tunnels signaled without explicit bandwidth reservation enables efficient bandwidth sharing that yields substantial bandwidth savings.]") Cisco IOS Guide at Page No. 651 ("[A similar approach can be used for node and SRLG protection. However, the decision of where to put the backup tunnels is more complicated because both node and SRLG failures effectively result in the simultaneous failure of several links. Therefore, the backup tunnels protecting traffic traversing all affected links cannot be computed independently of each other.]") Cisco IOS Guide at Page No. 579 ("[MPLS Traffic Engineering Shared Risk Link Groups SRLGs refer to situations where links in a network share a common fiber (or a common physical attribute). If one link fails, other links in the group may fail too. Links in the group have a shared risk. Backup tunnels should avoid using links in the same SRLG as interfaces they are protecting. Otherwise, when the protected link fails the backup tunnel fails too.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 191 of 698

No.	'821 Patent Claim 1	Sivabalan '928
		Figure 1 Backup Tunnel in the Same SRLG as the Interface It Is Protecting  Fig. 1 (annotation added) Figure 1 shows a primary label-switched path (LSP) from router R1 to router R5. The LSP protects against the failure of the R2-R3 link at R2 via a backup tunnel to R4. If the R2-R3 link fails, link protection reroutes the LSP along the backup tunnel. However, the R2-R3 link and one of the backup tunnel links are in the same SRLG. So if the R2-R3 link fails, the backup tunnel may fail too. The MPLS TE SRLG feature enhances backup tunnel path selection so a backup tunnel can avoid using links that are in the same SRLG as the interfaces it is protecting. There are two ways for a backup tunnel to avoid the SRLGs of its protected interface: • The router does not create the backup tunnel unless it avoids SRLGs of the protected interface.]") Cisco IOS Guide at Page No. 580 ("[Open Shortest Path First (OSPF) and Intermediate System-to-Intermediate System (IS-IS) flood the SRLG membership information (including other TE link attributes such as bandwidth availability and affinity) so that all routers in the network have the SRLG information for each link. With this topology information, routers can compute backup tunnel paths that exclude links having SRLGs in common with their protected interfaces. As shown in Figure 2, the backup tunnel avoids the link between R2 and R3, which shares an SRLG with the protected interface.]")

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 192 of 698

No.	'821 Patent Claim 1	Sivabalan '928
		 Fig. 2 (annotation added) Cisco IOS Guide at Page No. 587 ("[Step 4 show mpls traffic-eng topology Use this command to show the SRLG link membership flooded via the Interior Gateway Protocol (IGP).]") Cisco IOS Guide at Page 632 ("[IS-IS—Intermediate System-to-Intermediate System. Link-state hierarchical routing protocol that calls for intermediate system (IS) routers to exchange routing information based on a single metric to determine network topology.]") Cisco IOS Guide at Page No. 633 ("[OSPF—Open Shortest Path First. A link-state hierarchical Interior Gateway Protocol routing algorithm, derived from the IS-IS protocol. OSPF features include least-cost routing, multipath routing, and load balancing.]") Cisco IOS Guide at Page No. 581 ("[Autotunnel Backup for MPLS TE SRLGs Autotunnel backup is the ability of routers to create backup tunnels automatically. Therefore, you do not need to preconfigure each backup tunnel and then assign the backup tunnel to the protected interface. Only automatically created backup tunnels can avoid SRLGs or their protected interfaces.]") Cisco IOS Guide at Page No. 583 ("[Configuring the Routers That Automatically Create Backup Tunnels to Avoid MPLS TE SRLGs of Their Protected Interfaces

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 193 of 698

No.	'821 Patent Claim 1	Sivabalan '928
		Perform the following task to configure routers that automatically create backup tunnels to avoid MPLS TE SRLGs of their protected interfaces. Backup tunnels provide link protection by rerouting traffic to the next hop bypassing failed links or in this instance by avoiding SRLGs.]") Cisco IOS Guide at Page No. 644 ("[Backup Tunnel Selection Procedure When an LSP is signaled, each node along the LSP path that provides FRR protection for the LSP selects a backup tunnel for the LSP to use if either of the following events occurs: • The link to the next hop fails. • The next hop fails. By having the node select the backup tunnel for an LSP before a failure occurs, the LSP can be rerouted onto the backup tunnel quickly if there is a failure.]") Cisco IOS Guide at Page 646 ("[NHOP Versus NNHOP Backup Tunnels More than one backup tunnel can protect a given LSP, where one backup tunnel terminates at the LSP's NNHOP, and the other terminates at the LSP's NHOP. In this case, the router chooses the backup tunnel that terminates at the NNHOP (that is, FRR prefers NNHOP over NHOP backup tunnels).Table 1 lists the tunnel selection priorities. The first choice is an NNHOP backup tunnel that acquires its bandwidth from a subpool or global pool, and has limited bandwidth. If there is no such backup tunnel, the next choice (2) is a next-next hop backup tunnel that acquires a limited amount of bandwidth from any pool. The preferences go from 1 (best) to 8 (worst), where choice 3 is for an NNHOP backup tunnel with an unlimited amount of subpool or global-pool bandwidth.]")

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 194 of 698

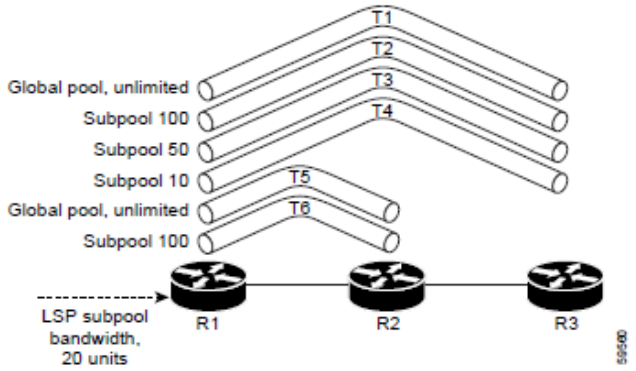
No.	'821 Patent Claim 1	Sivabalan '928

Table 1 Tunnel Selection Priorities

Preference	Backup Tunnel Destination	Bandwidth Pool	Bandwidth Amount
1 (Best)	NNHOP	Subpool or global pool	Limited
2	NNHOP	Any	Limited
3	NNHOP	Subpool or global pool	Unlimited
4	NNHOP	Any	Unlimited
5	NHOP	Subpool or global pool	Limited
6	NHOP	Any	Limited
7	NHOP	Subpool or global pool	Unlimited
8 (Worst)	NHOP	Any	Unlimited

Cisco IOS Guide at Page No. 647-648 ("[Figure 7 shows an example of the backup tunnel selection procedure based on the designated amount of global pool and subpool bandwidth currently available.

Figure 7 Choosing from Among Multiple Backup Tunnels



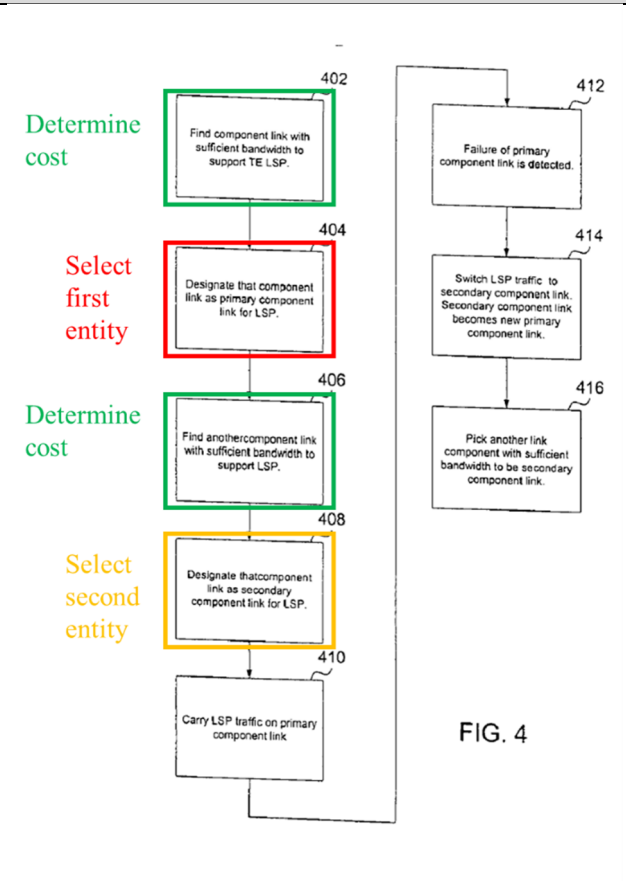
In this example, an LSP requires 20 units (kilobits per second) of sub-pool backup bandwidth. The best backup tunnel is selected as follows:
1. Backup tunnels T1 through T4 are considered first because they terminate at the NNHOP.
2. Tunnel T4 is eliminated because it has only ten units of sub-pool backup bandwidth.
3. Tunnel T1 is eliminated because it protects only LSPs using global-pool bandwidth.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 195 of 698

No.	'821 Patent Claim 1	Sivabalan '928
		4. Tunnel T3 is chosen over T2 because, although both have sufficient backup bandwidth, T3 has the least backup bandwidth available (leaving the most backup bandwidth available on T2). 5. Tunnels T5 and T6 need not be considered because they terminate at an NHOP, and therefore are less desirable than T3, which terminates at an NNHOP.]")

No.	'821 Patent Claim 2	Sivabalan '928
2[preamble]	The method of claim 1, wherein said step of selecting an entity pair further comprises:	Sivabalan '928 discloses step of selecting an entity pair further comprises. For example, Sivabalan '928 discloses selecting a primary component link and a secondary component link. *See Supra at* 1[c].
2[a]	selecting a working entity from said entity pair; and	Sivabalan '928 discloses selecting a working entity from said entity pair. For example, Sivabalan '928 discloses selecting a primary component link and a secondary component link. Sivabalan further discloses that data traffic will normally traverse the primary component link. Sivabalan '928 at Para No. 0038 ("[FIG. 4 is a flow chart describing steps of component link operation according to the first mode as provided by embodiments of the present invention. At step 402, node A receives a request to establish a Traffic Engineering LSP using link bundle 100. Node A identifies a component link with sufficient bandwidth to accommodate the requirements of this Traffic Engineering LSP. At step 404, this component link is designated as the primary component link for the LSP. At step 406, node A identifies another component link with sufficient bandwidth to support the LSP. This component link is designated at step 408 as the secondary component link for the LSP. Information about the primary and secondary component links is programmed into the label forwarding table entry or entries corresponding to the LSP. At step 410, packets belonging to the LSP are carried on the primary component link.]")

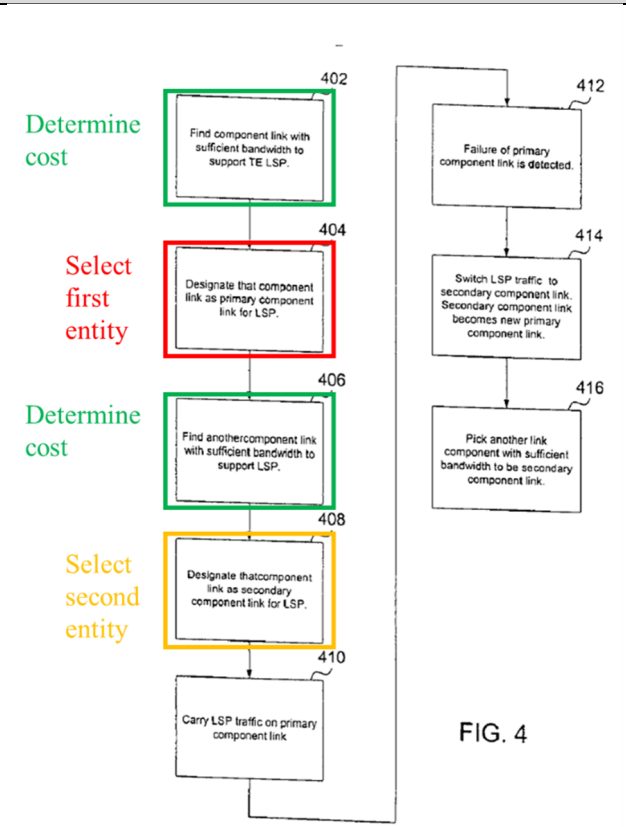
Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 196 of 698

No.	'821 Patent Claim 2	Sivabalan '928
		 Fig. 4 (annotation added) Sivabalan '928 at Para No. 0038 ("[FIG. 4 is a flow chart describing steps of component link operation according to the first mode as provided by embodiments of the present invention. At step 402, node A receives a request to establish a Traffic Engineering LSP using link bundle 100. Node A identifies a component link with sufficient bandwidth to accommodate the requirements of this Traffic Engineering LSP. At step 404, this component link is

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 197 of 698

No.	'821 Patent Claim 2	Sivabalan '928
		designated as the primary component link for the LSP. At step 406, node A identifies another component link with sufficient bandwidth to support the LSP. This component link is designated at step 408 as the secondary component link for the LSP. Information about the primary and secondary component links is programmed into the label forwarding table entry or entries corresponding to the LSP. At step 410, packets belonging to the LSP are carried on the primary component link.]") Sivabalan '928 at Para No. 0039 ("[At step 412, failure of the primary component link is detected. The response is a switchover of the LSP traffic to the secondary component link at step 414.]")
2[b]	selecting a protection entity from said entity pair.	For example, Sivabalan '928 discloses selecting a primary component link and a secondary component link. Sivabalan further discloses that data traffic will switchover to the secondary component link. For example, Sivabalan '928 discloses the response to failure of primary component link is a switchover of the LSP traffic to the secondary component link.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 198 of 698

No.	’821 Patent Claim 2	Sivabalan ’928
		 Fig. 4 (annotation added) Sivabalan ’928 at Para No. 0038 (“[FIG. 4 is a flow chart describing steps of component link operation according to the first mode as provided by embodiments of the present invention. At step 402, node A receives a request to establish a Traffic Engineering LSP using link bundle 100. Node A identifies a component link with sufficient bandwidth to accommodate the requirements of this Traffic Engineering LSP. At step 404, this component link is

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 199 of 698

No.	'821 Patent Claim 2	Sivabalan '928
		designated as the primary component link for the LSP. At step 406, node A identifies another component link with sufficient bandwidth to support the LSP. This component link is designated at step 408 as the secondary component link for the LSP. Information about the primary and secondary component links is programmed into the label forwarding table entry or entries corresponding to the LSP. At step 410, packets belonging to the LSP are carried on the primary component link.]") Sivabalan '928 at Para No. 0039 ("[At step 412, failure of the primary component link is detected. The response is a switchover of the LSP traffic to the secondary component link at step 414.]")

No.	'821 Patent Claim 3	Sivabalan '928
3	The method of claim 2, further comprising the step of selecting an active entity from the set consisting of said working entity and said protection entity.	Sivabalan '928 discloses the step of selecting an active entity from the set consisting of said working entity and said protection entity. For example, Sivabalan '928 discloses when failure of the primary component link is detected, LSP traffic is switched over to the secondary component link. Sivabalan '928 further discloses that prior to a failure, the primary component link is selected as the active link and after a failure, the secondary component link is selected as the active link. Sivabalan '928 at Para No. 0038 ("[FIG. 4 is a flow chart describing steps of component link operation according to the first mode as provided by embodiments of the present invention. At step 402, node A receives a request to establish a Traffic Engineering LSP using link bundle 100. Node A identifies a component link with sufficient bandwidth to accommodate the requirements of this Traffic Engineering LSP. At step 404, this component link is designated as the primary component link for the LSP. At step 406, node A identifies another component link with sufficient bandwidth to support the LSP. This component link is designated at step 408 as the secondary component link for the LSP. Information about the primary and secondary component links is programmed into the label forwarding table entry or entries corresponding to the LSP. At step 410, packets belonging to the LSP are carried on the primary component link.]")

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 200 of 698

No.	'821 Patent Claim 3	Sivabalan '928
		Sivabalan '928 at Para No. 0039 ("[At step 412, failure of the primary component link is detected. The response is a switchover of the LSP traffic to the secondary component link at step 414. Quality of service is maintained since the secondary component link also has sufficient guaranteed bandwidth.]")

No.	'821 Patent Claim 4	Sivabalan '928
4	The method of claim 2, wherein selecting an entity pair further comprises minimizing an overall cost function.	Sivabalan '928 discloses wherein selecting an entity pair further comprises minimizing an overall cost function. For example, Sivabalan '928 discloses node A identifies a primary component link with sufficient bandwidth to accommodate the requirements of this Traffic Engineering LSP and identifies a secondary component link with sufficient bandwidth to support the LSP. At least based upon Orckit's apparent understanding of the claim scope as reflected in its Infringement Disclosures, this limitation is met. To the extent that Sivabalan '928 is found to not meet this limitation, wherein selecting an entity pair further comprises minimizing an overall cost function would have been obvious to a person having ordinary skill in the art, as explained below. Sivabalan '928 at Para No. 0038 ("Node A identifies a component link with sufficient bandwidth to accommodate the requirements of this Traffic Engineering LSP. At step 404, this component link is designated as the primary component link for the LSP. At step 406, node A identifies another component link with sufficient bandwidth to support the LSP. This component link is designated at step 408 as the secondary component link for the LSP. Information about the primary and secondary component links is programmed into the label forwarding table entry or entries corresponding to the LSP.") *See Supra at* 1[b]. In addition, at least under the apparent claim scope alleged by Orckit's Infringement Disclosures, Sivabalan '928 in combination with (1) the knowledge of a person of ordinary skill in the art, alone or in further combination with (2) each (individually, as well as one or more

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 201 of 698

No.	'821 Patent Claim 4	Sivabalan '928
		together) of the references identified in element 4 of Exhibit E-2 renders the claim, including the present limitation, obvious. Below are examples of two such references. For example, Juniper IOS Guide discloses "Fate Sharing" where a Juniper router computes a primary and secondary path using CSPF, CSPF identifies all fate-sharing groups that are associated with the primary path by identifying all links that the primary path transverses and checks each link in the traffic engineering database against the compiled group list. Juniper IOS Guide further discloses that if the link is a member of a group, the cost of the link is increased by the cost of the group and higher the cost value. At least based upon Orckit's apparent understanding of the claim scope as reflected in its Infringement Disclosures, this limitation is met. Juniper IOS Guide at Page No. 73 ("[Fate Sharing Fate sharing allows you to create a database of information that CSPF uses to compute one or more backup paths to use in case the primary path becomes unstable. The database describes the relationships between elements of the network, such as routers and links. You can specify one or more elements within a group. Through fate sharing, you can configure backup paths that minimize the number of shared links and fiber paths with the primary paths as much as possible, to ensure that if a fiber is cut, the minimum amount of data is lost and a path still exists to the destination. For a backup path to work optimally, it must not share links or physical fiber paths with the primary path, ensuring that a single point of failure will not affect the primary and backup paths simultaneously. For more information about fate sharing, see the Junos OS Routing Protocols Configuration Guide.]") Juniper IOS Guide at Page No. 94 ("[Configuring the Ingress Router for MPLS-Signaled LSPs MPLS-signaled label-switched paths (LSPs) run from a specific ingress router to a specific egress router. For basic MPLS-signaled LSP function, you must configure the ingress router, but do not have to configure any other routers. To configure signaled LSPs, perform the following tasks on the ingress router: • Creating Named Paths on page 56

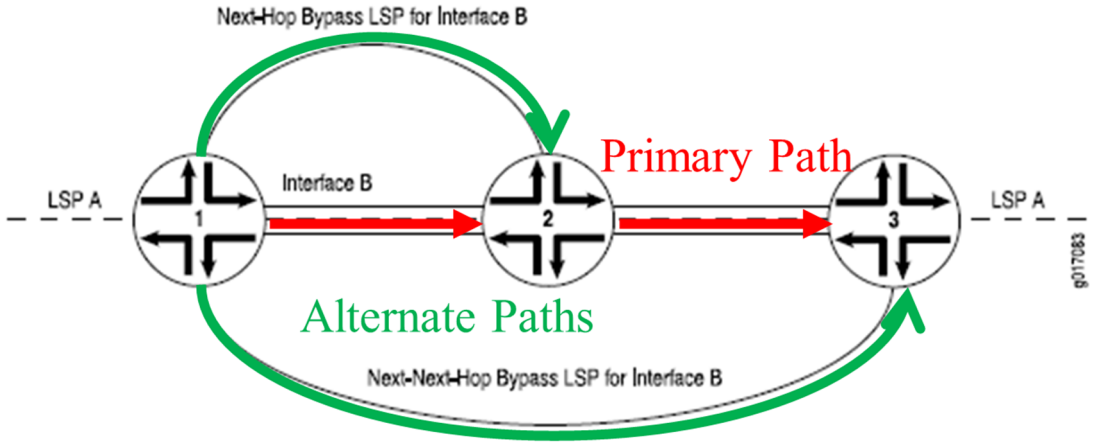
Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 202 of 698

No.	'821 Patent Claim 4	Sivabalan '928
		• Configuring Alternate Backup Paths Using Fate Sharing on page 58.]") Juniper IOS Guide at Page No. 96 ("[Configuring Alternate Backup Paths Using Fate Sharing You can create a database of information that Constrained Shortest Path First (CSPF) uses to compute one or more backup paths in case the primary path becomes unstable. The database describes the relationships between elements of the network, such as routers and links. Because these network elements share the same fate, this relationship is called fate sharing. You can configure backup paths that minimize the number of shared links and fiber paths with the primary paths as much as possible to ensure that, if a fiber is cut, the minimum amount of data is lost and a path still exists to the destination. For a backup path to work optimally, it must not share links or physical fiber paths with the primary path. This ensures that a single point of failure will not affect the primary and backup paths at the same time.]") Juniper IOS Guide at Page No. 97 ("[Fate-sharing groups contain three types of objects: • Point-to-point links—Identified by the IP addresses at each end of the link. Unnumbered point-to-point links are typically identified by borrowing IP addresses from other interfaces. Order is not important; from 1.2.3.4 to 1.2.3.5 and from 1.2.3.5 to 1.2.3.4 have the same meaning. • Non-point-to-point links—Include links on a LAN interface (such as Gigabit Ethernet interfaces) or nonbroadcast multiaccess (NBMA) interfaces (such as Asynchronous Transfer Mode [ATM] or Frame Relay). You identify these links by their individual interface address. For example, if the LAN interface 192.168.200.0/24 has four routers attached to it, each router link is individually identified: from 192.168.200.1; # LAN interface of router 1 from 192.168.200.2; # LAN interface of router 2 from 192.168.200.3; # LAN interface of router 3 from 192.168.200.4; # LAN interface of router 4 You can list the addresses in any order. • A router node—Identified by its configured router ID. An object can be in any number of groups, and a group can contain any number of objects. Each group has a configurable cost attributed to it, which represents the level of impact this

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 203 of 698

No.	'821 Patent Claim 4	Sivabalan '928
		group has on CSPF computations. The higher the cost, the less likely a backup path will share with the primary path any objects in the group. The cost is directly comparable to traffic engineering metrics. By default, the cost is 1. Changing the fate-sharing database does not affect established LSPs until the next reoptimization of CSPF. The fate-sharing database does influence fast-reroute computations.]") Juniper IOS Guide at Page No. 97 ("[When CSPF computes the primary paths of an LSP (or secondary paths when the primary path is not active), it ignores the fate-sharing information. You always want to find the best possible path (least IGP cost) for the primary path. When CSPF computes a secondary path while the primary path (of the same LSP) is active, the following occurs: 1. CSPF identifies all fate-sharing groups that are associated with the primary path. CSPF does this by identifying all links and nodes that the primary path traverses and compiling group lists that contain at least one of the links or nodes. CSPF ignores the ingress and egress nodes in the search. 2. CSPF checks each link in the traffic engineering database against the compiled group list. If the link is a member of a group, the cost of the link is increased by the cost of the group. If a link is a member of multiple groups, all group costs are added together. 3. CSPF performs the check for every node in the traffic engineering database, except the ingress and egress node. Again, a node can belong to multiple groups, so costs are additive. 4. The router performs regular CSPF computation with the adjusted topology.]") Juniper IOS Guide at Page No. 330 ("In Figure 23 on page 292, both node protection and link protection are enabled on Interface B on Router 1. Both node protection and link protection are also enabled on LSP A, an LSP that traverses the link transiting Router 1, Router 2, and Router 3. If the link between Router 1 and Router 2 fails, traffic from LSP A is quickly switched to the next-hop bypass LSP generated by link protection. If Router 2 suffers a hardware or software failure, traffic from LSP A is switched to the next-next-hop bypass LSP generated by node protection.")

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 204 of 698

No.	'821 Patent Claim 4	Sivabalan '928
		Figure 23: Node Protection Creating a Next-Next-Hop Bypass LSP  Fig. 23 (annotation added). For example, Cisco IOS Guide discloses that the path an LSP uses is determined by the LSP resource requirements and network resources, such as bandwidth. Cisco IOS Guide further discloses the MPLS TE SRLG feature that enhances backup tunnel path selection so a backup tunnel can avoid using links that are in the same SRLG as the interfaces it is protecting. Thus, at least under the apparent claim scope alleged by Orckit's Infringement Disclosures, this limitation is met. Cisco IOS Guide at Page No. 253 ("[Path Option Selection for MPLS TE Tunnel LSPs This section contains the following topics about path option selection for MPLS TE Tunnel LSPs: • Constraint-Based Routing and Path Option Selection, page 5 • Tunnel Reoptimization and Path Option Selection, page 5 • Path Option Selection with Bandwidth Override, page 6 Constraint-Based Routing and Path Option Selection MPLS traffic engineering automatically establishes and maintains LSPs across the backbone by using the Resource Reservation Protocol (RSVP). The path that an LSP uses is determined

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 205 of 698

No.	'821 Patent Claim 4	Sivabalan '928
		by the LSP resource requirements and network resources, such as bandwidth. Traffic engineering tunnels are calculated at the LSP head based on a fit between required and available resources (constraint-based routing). Without the Path Option for Bandwidth Override feature, a TE tunnel establishes an LSP based on dynamic or explicit path options in order of preference. However, the bandwidth and other attributes configured on the TE tunnel allow the setup of an LSP only if LSP path options satisfy the constraints. If a path cannot be found that satisfies the configured path options, then the tunnel is not set up.]") Cisco IOS Guide at Page No. 639 ("[Bandwidth Protection NHOP and NNHOP backup tunnels can be used to provide bandwidth protection for rerouted LSPs. This is referred to as backup bandwidth. You can associate backup bandwidth with NHOP or NNHOP backup tunnels. This informs the router of the amount of backup bandwidth a particular backup tunnel can protect. When a router maps LSPs to backup tunnels, bandwidth protection ensures that an LSP uses a given backup tunnel only if there is sufficient backup bandwidth. The router selects which LSPs use which backup tunnels in order to provide maximum bandwidth protection. That is, the router determines the best way to map LSPs onto backup tunnels in order to maximize the number of LSPs that can be protected. For information about mapping tunnels and assigning backup bandwidth, see the "Backup Tunnel Selection Procedure" section on page 10. LSPs that have the "bandwidth protection desired" bit set have a higher right to select backup tunnels that provide bandwidth protection; that is, those LSPs can preempt other LSPs that do not have that bit set. For more information, see the "Prioritizing Which LSPs Obtain Backup Tunnels with Bandwidth Protection" section on page 8.]") Cisco IOS Guide at Page No. 640 ("[Backup Tunnel Support Backup tunnel support has the following capabilities: • Backup Tunnels Can Terminate at the Next-Next Hop to Support FRR, page 6 • Multiple Backup Tunnels Can Protect the Same Interface, page 6 • Backup Tunnels Provide Scalability, page 7.]") Cisco IOS Guide at Page 640 ("[Multiple Backup Tunnels Can Protect the Same Interface

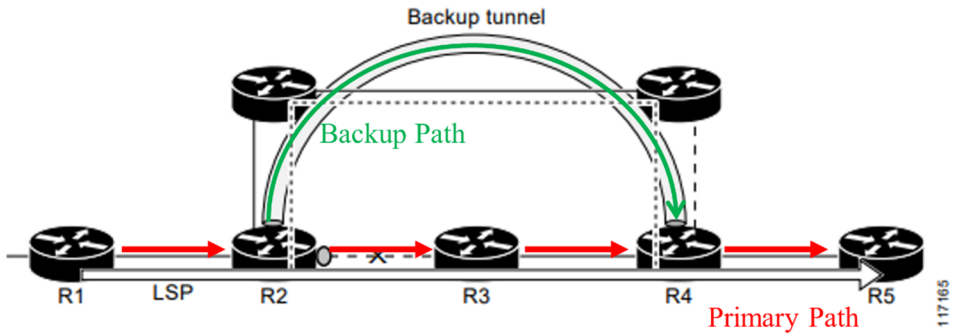
Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 206 of 698

No.	'821 Patent Claim 4	Sivabalan '928
		There is no limit (except memory limitations) to the number of backup tunnels that can protect a given interface. In many topologies, support for node protection requires supporting multiple backup tunnels per protected interface. These backup tunnels can terminate at the same destination or at different destinations. That is, for a given protected interface, you can configure multiple NHOP or NNHOP backup tunnels. This allows redundancy and load balancing.]") Cisco IOS Guide at Page No. 603 ("[Traffic Engineering Tunnels MPLS TE lets you build label switched paths (LSPs) across your network for forwarding traffic. MPLS TE LSPs let the headend of a TE tunnel control the path its traffic takes to a particular destination. This method is more flexible than forwarding traffic based only on a destination address. Interarea tunnels allow you to do the following: • Build TE tunnels between areas (interarea tunnels) • Build TE tunnels that start and end in the same area, on multiple areas on a router (intra-area tunnels) Some tunnels are more important than others. For example, you may have tunnels carrying VoIP traffic and tunnels carrying data traffic that are competing for the same resources. MPLS TE allows you to have some tunnels preempt others. Each tunnel has a priority, and more-important tunnels take precedence over less-important tunnels.]") Cisco IOS Guide at Page No. 641 ("[Backup Bandwidth Protection Backup bandwidth protection allows you to give LSPs carrying certain kinds of data (such as voice) priority for using backup tunnels.]") Cisco IOS Guide at Page No. 650 ("[For each protected link AB with a maximum reservable subpool value of n, there may be a path from node A to node B such that the difference between the maximum reservable global and the maximum reservable subpool is at least the value of n. If it is possible to find such paths for each link in the network, you can establish all the backup tunnels along such paths without any bandwidth reservations. If there is a single link failure, only one backup tunnel will use any link on its path. Because that path has at

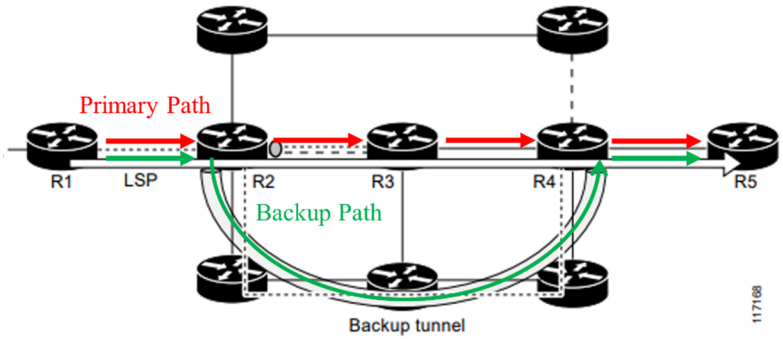
Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 207 of 698

No.	'821 Patent Claim 4	Sivabalan '928
		least n available bandwidth (in the global pool), assuming that marking and scheduling is configured to classify the subpool traffic into a priority queue, the subpool bandwidth is guaranteed. This approach allows sharing of the global pool bandwidth between backup tunnels protecting independent link failures. The backup tunnels are expected to be used for only a short period of time after a failure (until the headends of affected LSPs reroute those LSPs to other paths with available subpool bandwidth). The probability of multiple unrelated link failures is very small (in the absence of node or shared risk link group (SRLG) failures, which result in multiple link failures). Therefore, it is reasonable to assume that link failures are in practice independent with high probability. This "independent failure assumption" in combination with backup tunnels signaled without explicit bandwidth reservation enables efficient bandwidth sharing that yields substantial bandwidth savings.]") Cisco IOS Guide at Page No. 651 ("[A similar approach can be used for node and SRLG protection. However, the decision of where to put the backup tunnels is more complicated because both node and SRLG failures effectively result in the simultaneous failure of several links. Therefore, the backup tunnels protecting traffic traversing all affected links cannot be computed independently of each other.]") Cisco IOS Guide at Page No. 579 ("[MPLS Traffic Engineering Shared Risk Link Groups SRLGs refer to situations where links in a network share a common fiber (or a common physical attribute). If one link fails, other links in the group may fail too. Links in the group have a shared risk. Backup tunnels should avoid using links in the same SRLG as interfaces they are protecting. Otherwise, when the protected link fails the backup tunnel fails too.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 208 of 698

No.	'821 Patent Claim 4	Sivabalan '928
		Figure 1 Backup Tunnel in the Same SRLG as the Interface It Is Protecting  Fig. 1 (annotation added) Figure 1 shows a primary label-switched path (LSP) from router R1 to router R5. The LSP protects against the failure of the R2-R3 link at R2 via a backup tunnel to R4. If the R2-R3 link fails, link protection reroutes the LSP along the backup tunnel. However, the R2-R3 link and one of the backup tunnel links are in the same SRLG. So if the R2-R3 link fails, the backup tunnel may fail too. The MPLS TE SRLG feature enhances backup tunnel path selection so a backup tunnel can avoid using links that are in the same SRLG as the interfaces it is protecting. There are two ways for a backup tunnel to avoid the SRLGs of its protected interface: • The router does not create the backup tunnel unless it avoids SRLGs of the protected interface.]") Cisco IOS Guide at Page No. 580 ("[Open Shortest Path First (OSPF) and Intermediate System-to-Intermediate System (IS-IS) flood the SRLG membership information (including other TE link attributes such as bandwidth availability and affinity) so that all routers in the network have the SRLG information for each link. With this topology information, routers can compute backup tunnel paths that exclude links having SRLGs in common with their protected interfaces. As shown in Figure 2, the backup tunnel avoids the link between R2 and R3, which shares an SRLG with the protected interface.]")

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 209 of 698

No.	'821 Patent Claim 4	Sivabalan '928
		Figure 2 Backup Tunnel That Avoids SRLG of Protected Interface  Fig. 2 (annotation added) Cisco IOS Guide at Page No. 587 ("[Step 4 show mpls traffic-eng topology Use this command to show the SRLG link membership flooded via the Interior Gateway Protocol (IGP).]") Cisco IOS Guide at Page 632 ("[IS-IS—Intermediate System-to-Intermediate System. Link-state hierarchical routing protocol that calls for intermediate system (IS) routers to exchange routing information based on a single metric to determine network topology.]") Cisco IOS Guide at Page No. 633 ("[OSPF—Open Shortest Path First. A link-state hierarchical Interior Gateway Protocol routing algorithm, derived from the IS-IS protocol. OSPF features include least-cost routing, multipath routing, and load balancing.]") Cisco IOS Guide at Page No. 581 ("[Autotunnel Backup for MPLS TE SRLGs Autotunnel backup is the ability of routers to create backup tunnels automatically. Therefore, you do not need to preconfigure each backup tunnel and then assign the backup tunnel to the protected interface. Only automatically created backup tunnels can avoid SRLGs or their protected interfaces.]") Cisco IOS Guide at Page No. 583 ("[Configuring the Routers That Automatically Create Backup Tunnels to Avoid MPLS TE SRLGs of Their Protected Interfaces

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 210 of 698

No.	'821 Patent Claim 4	Sivabalan '928
		Perform the following task to configure routers that automatically create backup tunnels to avoid MPLS TE SRLGs of their protected interfaces. Backup tunnels provide link protection by rerouting traffic to the next hop bypassing failed links or in this instance by avoiding SRLGs.]") Cisco IOS Guide at Page No. 644 ("[Backup Tunnel Selection Procedure When an LSP is signaled, each node along the LSP path that provides FRR protection for the LSP selects a backup tunnel for the LSP to use if either of the following events occurs: • The link to the next hop fails. • The next hop fails. By having the node select the backup tunnel for an LSP before a failure occurs, the LSP can be rerouted onto the backup tunnel quickly if there is a failure.]") Cisco IOS Guide at Page 646 ("[NHOP Versus NNHOP Backup Tunnels More than one backup tunnel can protect a given LSP, where one backup tunnel terminates at the LSP's NNHOP, and the other terminates at the LSP's NHOP. In this case, the router chooses the backup tunnel that terminates at the NNHOP (that is, FRR prefers NNHOP over NHOP backup tunnels).Table 1 lists the tunnel selection priorities. The first choice is an NNHOP backup tunnel that acquires its bandwidth from a subpool or global pool, and has limited bandwidth. If there is no such backup tunnel, the next choice (2) is a next-next hop backup tunnel that acquires a limited amount of bandwidth from any pool. The preferences go from 1 (best) to 8 (worst), where choice 3 is for an NNHOP backup tunnel with an unlimited amount of subpool or global-pool bandwidth.]")

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 211 of 698

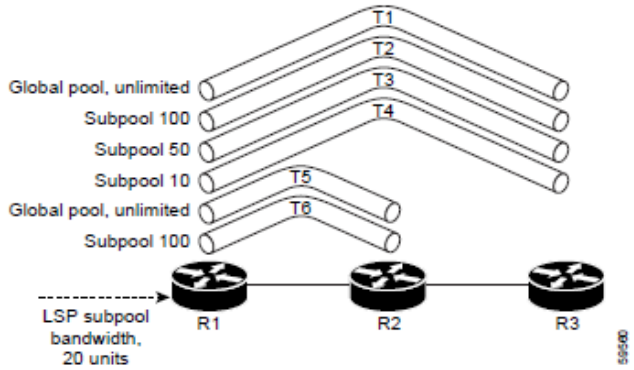
No.	'821 Patent Claim 4	Sivabalan '928

Table 1 Tunnel Selection Priorities

Preference	Backup Tunnel Destination	Bandwidth Pool	Bandwidth Amount
1 (Best)	NNHOP	Subpool or global pool	Limited
2	NNHOP	Any	Limited
3	NNHOP	Subpool or global pool	Unlimited
4	NNHOP	Any	Unlimited
5	NHOP	Subpool or global pool	Limited
6	NHOP	Any	Limited
7	NHOP	Subpool or global pool	Unlimited
8 (Worst)	NHOP	Any	Unlimited

Cisco IOS Guide at Page No. 647-648 ("[Figure 7 shows an example of the backup tunnel selection procedure based on the designated amount of global pool and subpool bandwidth currently available.

Figure 7 Choosing from Among Multiple Backup Tunnels



In this example, an LSP requires 20 units (kilobits per second) of sub-pool backup bandwidth. The best backup tunnel is selected as follows:
1. Backup tunnels T1 through T4 are considered first because they terminate at the NNHOP.
2. Tunnel T4 is eliminated because it has only ten units of sub-pool backup bandwidth.
3. Tunnel T1 is eliminated because it protects only LSPs using global-pool bandwidth.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 212 of 698

No.	'821 Patent Claim 4	Sivabalan '928
		4. Tunnel T3 is chosen over T2 because, although both have sufficient backup bandwidth, T3 has the least backup bandwidth available (leaving the most backup bandwidth available on T2). 5. Tunnels T5 and T6 need not be considered because they terminate at an NHOP, and therefore are less desirable than T3, which terminates at an NNHOP.]")

No.	'821 Patent Claim 5	Sivabalan '928
5	The method of claim 4, wherein said overall cost function comprises substantially minimizing a probability of concurrent failure of said protection entity and said working entity.	Sivabalan '928 discloses wherein said overall cost function comprises substantially minimizing a probability of concurrent failure of said protection entity and said working entity. For example, Sivabalan '928 discloses node A identifies a primary component link with sufficient bandwidth to accommodate the requirements of this Traffic Engineering LSP and identifies a secondary component link with sufficient bandwidth to support the LSP. At least based upon Orckit's apparent understanding of this term as reflected in its Infringement Disclosures, this limitation is met. *See Supra at* 1[b].

No.	'821 Patent Claim 6	Sivabalan '928
6	The method of claim 4, wherein said overall cost function comprises a predefined entity cost metric.	Sivabalan '928 discloses wherein said overall cost function comprises a predefined entity cost metric. For example, Sivabalan '928 discloses node A identifies a primary component link with sufficient bandwidth to accommodate the requirements of this Traffic Engineering LSP and identifies a secondary component link with sufficient bandwidth to support the LSP. At least based upon Orckit's apparent understanding of this term as reflected in its Infringement Disclosures, this limitation is met. *See Supra at* 1[b].

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 213 of 698

No.	'821 Patent Claim 7	Sivabalan '928
7	The method of claim 6, wherein said predefined entity cost metric is selected from the group consisting of interior gateway protocol (IGP) and traffic engineering (TE).	Sivabalan '928 discloses the method of claim 6, wherein said predefined entity cost metric is selected from the group consisting of interior gateway protocol (IGP) and traffic engineering (TE). For example, Sivabalan '928 discloses the use of MPLS traffic engineering protocols to form the primary and backup component links within its network. Sivabalan further discloses the use of other well known networking protocols for network formation, and explains that "Other Suitable protocols as known in the art may also be implemented by nodes of the example network." Sivabalan '928 ¶ 28. A person of ordinary skill in the art would understand from the teachings of Sivabalan '928 that IGP, another well-known protocol, could also be used. At least based upon Orckit's apparent understanding of this term as reflected in its Infringement Disclosures, this limitation is met. Sivabalan '928 at Para No. 0003-0005 "MPLS Traffic Engineering LSPs traverse a series of nodes and links that interconnect them. MPLS Traffic Engineering protocols define a link to be a logical construct that represents physical resources that interconnect label Switched routers. The information about a particular link including the available link bandwidth is used in determining the routes of LSPS and Signaling their establishment. single pair of nodes may be connected by a number of links including both physical links and logical links, i.e., logical connections through a network such that the nodes are considered adjacent for MPLS Traffic Engineering purposes. To place LSPs, MPLS Traffic Engineering relies on conventional Internet link State routing protocols such as OSPF and IS-IS that require each node to report information about all of its directly connected nodes to all nodes in the network. To avoid the need to advertise information about each link between a pair of nodes, the concept of a "link bundle" has been developed to support MPLS Traffic Engineering. A link bundle combines multiple links between a pair of nodes. The links that make up a bundle are then referred to as "component links." Link bundling improves routing Scalability by reducing the amount of information that has to be advertised by link State routing protocols to Support Signaling and computation of MPLS Traffic Engineering LSPs. Information about link bundles including bandwidth information is flooded throughout the network while

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 214 of 698

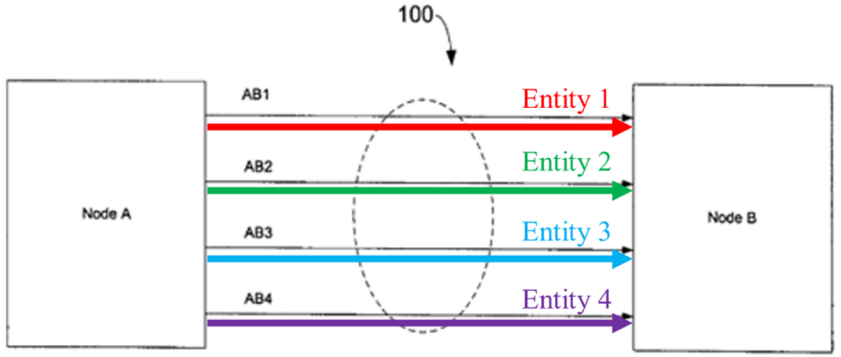
No.	'821 Patent Claim 7	Sivabalan '928
		information about component links is maintained only at the directly connected nodes. Details of link bundling are described in Kompella, et al., "Link Bundling in MPLS Traffic Engineering." IETF Internet Draft, July 2002, the contents of which are herein incorporated by reference in their entirety for all purposes." Sivabalan '928 at Para No. 0007 ("An important aspect of MPLS Traffic Engineering reliability is addressing link failures in order to minimize traffic loss. When a component link fails, all LSPs passing through that component link are broken and must be rerouted. One protection mechanism provides that the failure will be signaled to the head-end of each affected LSP and the head-ends will react by rerouting the LSP to avoid the failed component link. Another protection mechanism referred to as Fast Reroute reduces the recovery time by providing for impacted traffic to be patched through a local backup tunnel around the failure. The Fast Reroute mechanism protects the impacted traffic until such time the head-end can reroute the LSP.") Sivabalan '928 at Para No. 0032 ("MPLS Traffic Engineering reliability objectives require mechanisms for handling failures of network elements such as links and component links. Embodiments of the present invention provide for a local rerouting around a failed component link in a link bundle. Traffic on a failed component link can be shifted to one or more other component links of the same bundle without any signaling to and from the head-end of the LSP as required by prior art mechanisms.")

No.	'821 Patent Claim 8	Sivabalan '928
8	The method of claim 4, further comprising the step of configuring said working entity as revertive.	Sivabalan '928 discloses the step of configuring said working entity as revertive. For example, Sivabalan '928 discloses an immediate transition to a second path if there is an error on the first path. At least based upon Orckit's apparent understanding of this term as reflected in its Infringement Disclosures, this limitation is met. Sivabalan '928 at Claim 20 ("The computer program product of claim 16 further comprising:

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 215 of 698

No.	'821 Patent Claim 8	Sivabalan '928
		code that, while traffic of said LSP is sent over said one or more secondary component links, searches for a component link of said link bundle to serve a new primary component link; and code that, upon failure to identify a component link having sufficient available bandwidth to server as a new primary component link, signals a head-end of said LSP to reroute said LSP.") Sivabalan '928 at Para No. 0048 ("If a new primary component link cannot be found, however, the head-end of the broken LSP is signaled to initiate a rerouting of the LSP that does not include link bundle 100. This will be a soft preemption, i.e., such that the old LSP routing including link bundle 100 remains in operation until the new one that excludes link bundle 100 is operational. Traffic is thus forwarded over the secondary component links of link bundle 100 until the LSP is successfully rerouted.") Sivabalan '928 at Para No. 0033 ("FIG. 1. depicts a representative link bundle scenario according to one embodiment of the present invention. Link bundle 100 interconnects nodes A and B. Link bundle 100 includes components links AB1, AB2, AB3, and AB4. Node A propagates information about the aggregate characteristics of link bundle 100 including the available aggregate bandwidth to other nodes in the network in accordance with the relevant link state routing protocols underlying MPLS Traffic Engineering operation. Information about the component links AB1, AB2, AB3, and AB4 is not distributed in this way but only maintained locally to node A. Node B has similar knowledge of the links extending in the reverse direction. Information requesting the establishment of an LSP through nodes A and B will thus specify link bundle 100 but not the constituent component links.")

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 216 of 698

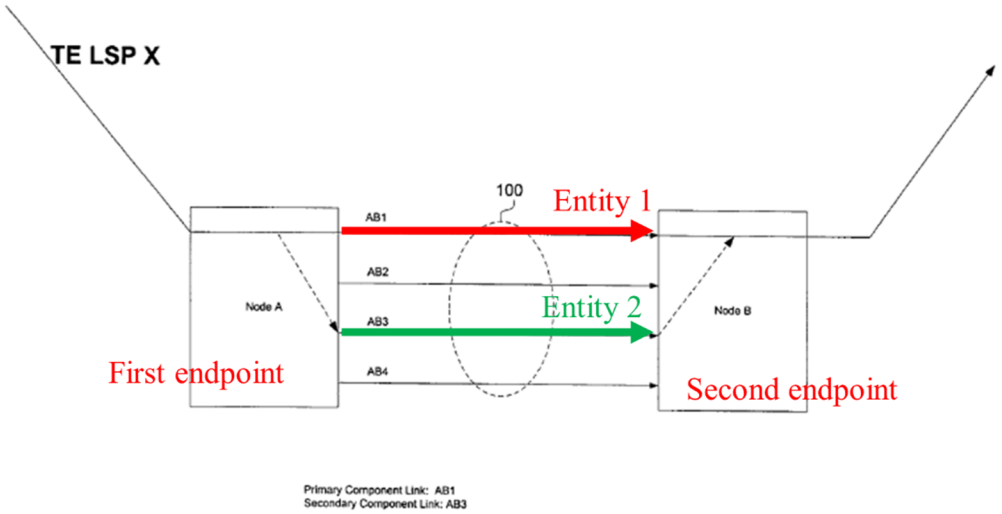
No.	'821 Patent Claim 8	Sivabalan '928
		 FIG. 1 Fig. 1 (annotation added)

No.	'821 Patent Claim 9	Sivabalan '928
9[preamble]	The method of claim 4, wherein said overall cost function comprises:	Sivabalan '928 discloses an overall cost function. For example, Sivabalan '928 discloses that sufficient bandwidth is reserved on each link and component link in the MPLS Traffic Engineering LSP to meet a bandwidth guarantee. Sivabalan further discloses that an LSP placed on a link bundle can operate in one of two different modes: in the first mode a single secondary link is assigned as a backup for the primary component link and in the second mode, one primary component link and one or more secondary component links are assigned to the LSP. Sivabalan '928 further discloses that when node A receives a request to establish a traffic engineering LSP using a link bundle with the pair of nodes, it identifies primary and secondary/backup component links within the link bundle with sufficient bandwidth to accommodate the bandwidth requirement of the LSP and designates it as a primary link for LSP. At least based upon Orckit's apparent understanding of this term as reflected in its Infringement Disclosures, this limitation is met. *See supra at* 1[b].

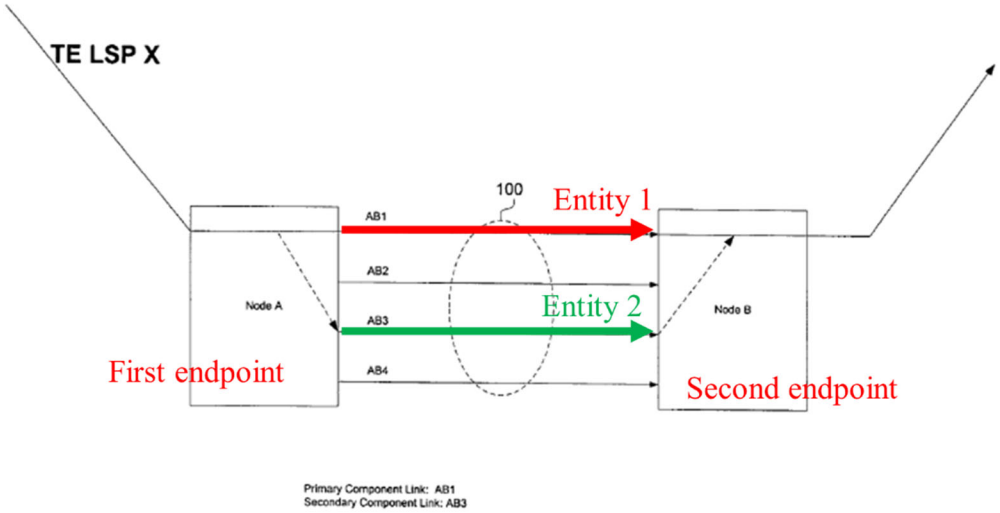
Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 217 of 698

No.	'821 Patent Claim 9	Sivabalan '928
9[a]	selecting a subset of entity pairs wherein each entity pair of said subset has substantially minimum probability of a concurrent failure of said protection entity and said working entity; and	Sivabalan '928 discloses selecting a subset of entity pairs wherein each entity pair of said subset has substantially minimum probability of a concurrent failure of said protection entity and said working entity. For example, Sivabalan '928 discloses selecting a pair of links that reduce the likelihood that they will fail together. At least based upon Orckit's apparent understanding of this term as reflected in its Infringement Disclosures, this limitation is met. Sivabalan '928 at Claim 14 ("The computer program product of claim 13 wherein, in advance of said failure, bandwidth for use by said LSP is allocated on said exactly one secondary component link responsive to a bandwidth requirement of said LSP.") Sivabalan '928 at Para No. 0036 ("In the first mode, a single secondary component link is assigned as backup for the primary component link. Sufficient bandwidth to accommodate the LSP is reserved on this single secondary component link. The amount of available bandwidth on the secondary component link is reduced accordingly. FIG. 2 illustrates an example of this first mode where an MPLS Traffic Engineering LSP X employs link bundle 100 between nodes A and B. Node A ties LSP X to a primary component link AB1 and a secondary component link AB3.")

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 218 of 698

No.	'821 Patent Claim 9	Sivabalan '928
		 FIG. 2 Fig. 2 (annotation added) *See supra at* Claim 4.
9[b]	if said subset comprises at least two entity pairs, selecting an entity pair from said subset that minimizes an entity cost function.	Sivabalan '928 discloses if said subset comprises at least two entity pairs, selecting an entity pair from said subset that minimizes an entity cost function. For example, Sivabalan '928 discloses selecting the entity pair with the lowest cost, such as by bandwidth requirements. Sivabalan '928 at Para No. 0036 ("In the first mode, a single secondary component link is assigned as backup for the primary component link. Sufficient bandwidth to accommodate the LSP is reserved on this single secondary component link. The amount of available bandwidth on the secondary component link is reduced accordingly. FIG. 2 illustrates an

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 219 of 698

No.	'821 Patent Claim 9	Sivabalan '928
		example of this first mode where an MPLS Traffic Engineering LSP X employs link bundle 100 between nodes A and B. Node A ties LSP X to a primary component link AB1 and a secondary component link AB3.")  FIG. 2 Fig. 2 (annotation added) *See supra at* Claim 4.

No.	'821 Patent Claim 10	Sivabalan '928
10	The method of claim 9, wherein said entity cost function	Sivabalan '928 discloses wherein said entity cost function comprises a predefined metric.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 220 of 698

No.	'821 Patent Claim 10	Sivabalan '928
	comprises a predefined metric.	For example, Sivabalan '928 takes into account, without limitation, metrics related to bandwidth requirements, media requirements, and/or the priority of one traffic flow versus another. At least based upon Orckit's apparent understanding of this term as reflected in its Infringement Disclosures, this limitation is met. Sivabalan '928 at Claim 4 ("The method of claim 3 wherein, in advance of said failure, bandwidth for use by said LSP is allocated on said exactly one secondary component link responsive to a bandwidth requirement of said LSP.") Sivabalan '928 at Claim 14 ("The computer program product of claim 13 wherein, in advance of said failure, bandwidth for use by said LSP is allocated on said exactly one secondary component link responsive to a bandwidth requirement of said LSP.") Sivabalan '928 at Claim 15 ("The computer program product of claim 13 further comprising: code that, while sending traffic of said LSP over said exactly one secondary component link, designates said exactly one secondary component link as a new primary component link and selects a new secondary component link responsive to a bandwidth requirement of said LSP.") Sivabalan '928 at Para 0031 ("MPLS Traffic Engineering establishes LSPs that have guaranteed bandwidth. Each MPLS Traffic Engineering LSP thus has a bandwidth requirement. Sufficient bandwidth is reserved on each link and component link in the MPLS Traffic Engineering LSP to assure that the bandwidth guarantee is met. Bandwidth reserved on a link or component link for one LSP is not available for other LSPs.") Sivabalan '928 at Para 0052 ("Network device 600 interfaces with physical media via a plurality of linecards 606. Linecards 606 may incorporate Ethernet interfaces, DSL interfaces, Gigabit Ethernet interfaces, 10-Gigabit Ethernet interfaces, SONET interfaces, etc. As packets are received, processed, and forwarded by network device 600, they may be stored in a packet memory 608. Network device 600 implements all of the network protocols and extensions thereof described above as well as the data networking features provided by the present invention.")

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 221 of 698

No.	'821 Patent Claim 10	Sivabalan '928
		Sivabalan '928 at Claim 17 ("The computer program product of claim 16 wherein said code that sends traffic of said LSP over said one or more secondary component links comprises: code that assigns lower priority to traffic of said LSP compared to other traffic carried by said one or more secondary component links.") Sivabalan '928 at Para 0044 ("In a Diff-Serv scenario that divides traffic among different queues corresponding to different priority levels, the packets from broken LSPs are sent to a "best effort" queue which is given relatively low priority in selecting packets for transmission. An alternative Diff-Serv scenario employs weighted random early detection (WRED) to randomly drop packets in congested conditions with the dependence of drop probability on congestion being set independently for each PHB. The excess traffic PHB subjects packets from broken LSPs to a very aggressive WRED profile such that these packets have a relatively high probability of being dropped as congestion increases compared to other traffic.")

No.	'821 Patent Claim 11	Sivabalan '928
11	The method of claim 10, wherein said predefined metric is selected from the group consisting of interior gateway protocol (IGP) and traffic engineering (TE).	Sivabalan '928 discloses wherein said predefined metric is selected from the group consisting of interior gateway protocol (IGP) and traffic engineering (TE). For example, Sivabalan '928 discloses the use of MPLS traffic engineering protocols to form the primary and backup component links within its network. Sivabalan further discloses the use of other well known networking protocols for network formation, and explains that "Other Suitable protocols as known in the art may also be implemented by nodes of the example network." Sivabalan '928 ¶ 28. A person of ordrinary skill in the art would understand from the teachings of Sivabalan '928 that IGP, another well-known protocol, could also be used. At least based upon Orckit's apparent understanding of this term as reflected in its Infringement Disclosures, this limitation is met. *See supra* Claim 7.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 222 of 698

No.	'821 Patent Claim 12	Sivabalan '928
12	The method of claim 10, further comprising the step of configuring said working entity as revertive.	Sivabalan '928 discloses the step of configuring said working entity as revertive. For example, Sivabalan '928 discloses an immediate transition to a second path if there is an error on the first path. At least based upon Orckit's apparent understanding of this term as reflected in its Infringement Disclosures, this limitation is met. *See supra at* Claim 8.

No.	'821 Patent Claim 13	Sivabalan '928
13	The method of claim 1, further comprising the step of: if said entity pair reselection results in both working and protection entities being replaced, sequentially replacing said working entity and said protection entity.	Sivabalan '928 discloses the method of claim 1, further comprising the step of: if said entity pair reselection results in both working and protection entities being replaced, sequentially replacing said working entity and said protection entity. For example, Sivabalan '928 discloses when failure of the primary component link is detected, secondary component link now becomes the primary component link and node A selects another component link in the same link bundle to be the new secondary component link. Sivabalan '928 at Para No. 0039 ("[At step 412, failure of the primary component link is detected. The response is a switchover of the LSP traffic to the secondary component link at step 414. Quality of service is maintained since the secondary component link also has sufficient guaranteed bandwidth. There is no signaling to the head-end to implement the switch-over. The secondary component link now becomes the primary component link. At step 416, node A selects another component link in the same link bundle to be the new secondary component link.]")

No.	'821 Patent Claim 14	Sivabalan '928
14[preamble]	A system for selecting entities within an MPLS network, comprising:	Sivabalan '928 discloses a system for selecting entities within an MPLS network.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 223 of 698

No.	'821 Patent Claim 14	Sivabalan '928
		For example, Sivabalan '928 discloses a node A/ network device for transparent re-routing of MPLS traffic engineering LSPs within a link bundle and selection of primary or secondary/ backup component links. Sivabalan '928 at Claim 21 ("[21. A network device that operates a first node in a label switched network, said network device comprising: a processor; and a memory device storing instructions executed by said processor, said instructions comprising: code that establishes a link bundle comprising a plurality of component links between said first node and a second node; code that, upon establishment of an LSP including said first node and said second node, selects a first component link of said link bundle as a primary component link to assign to said LSP; code that selects one or more component links of said link bundle other than said first component link as one or more secondary component links to assign to said LSP; code that sends traffic of said LSP over said primary component link until a failure of said primary component link; and code that, after failure of said primary component link, sends traffic of said LSP over said one or more secondary component links instead of said primary component link.]") Sivabalan '928 at Claim 22 ("[22. Apparatus for operating a first node in a label switched network, said apparatus comprising: means for establishing a link bundle comprising a plurality of component links between said first node and a second node; means for, upon establishment of an LSP including said first node and said second node, selecting a first component link of said link bundle as a primary component link to assign to said LSP; means for selecting one or more component links of said link bundle other than said first component link as one or more secondary component links to assign to said LSP; means for sending traffic of said LSP over said primary component link until a failure of said primary component link; and

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 224 of 698

No.	'821 Patent Claim 14	Sivabalan '928
		means for, after failure of said primary component link, sending traffic of said LSP over said one or more secondary component links instead of said primary component link.]")
14[a]	a data structure comprising a plurality of transport entity descriptors;	Sivabalan '928 discloses a plurality of transport entity descriptors. For example, Sivabalan '928 discloses in MPLS traffic engineering, a link bundle combines multiple links between a pair of nodes and is referred to as a "component link" and Information about the primary and secondary component links is programmed into the label forwarding table entry. Sivabalan further discloses that this information is stored within network devices, e.g., Node A and Node B. A person of ordinary skill in the art would understand from these disclosures that the information regarding component links would be stored as a data structure in digital computer memory. Sivabalan '928 at Para No. 0005 ("[To avoid the need to advertise information about each link between a pair of nodes, the concept of a "link bundle" has been developed to support MPLS Traffic Engineering. A link bundle combines multiple links between a pair of nodes. The links that make up a bundle are then referred to as "component links." Link bundling improves routing scalability by reducing the amount of information that has to be advertised by link state routing protocols to support signaling and computation of MPLS Traffic Engineering LSPs. Information about link bundles including bandwidth information is flooded throughout the network while information about component links is maintained only at the directly connected nodes.]") Sivabalan '928 at Para No. 0038 ("[FIG. 4 is a flow chart describing steps of component link operation according to the first mode as provided by embodiments of the present invention. At step 402, node A receives a request to establish a Traffic Engineering LSP using link bundle 100. Node A identifies a component link with sufficient bandwidth to accommodate the requirements of this Traffic Engineering LSP. At step 404, this component link is designated as the primary component link for the LSP. At step 406, node A identifies another component link with sufficient bandwidth to support the LSP. This component link is designated at step 408 as the secondary component link for the LSP. Information about the primary and secondary component links is programmed into the label forwarding table entry

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 225 of 698

No.	'821 Patent Claim 14	Sivabalan '928
		or entries corresponding to the LSP. At step 410, packets belonging to the LSP are carried on the primary component link.]")
14[b]	an entity protection switch configured to switch between a working entity and a protection entity; and	Sivabalan '928 discloses switch between a working entity and a protection entity. For example, Sivabalan '928 discloses when failure of the primary component link is detected, a switchover of LSP traffic to the secondary component link is performed. A person of ordinary skill in the art would understand from these disclosures that the switchover logic would be stored as a data structure in digital computer memory. Sivabalan '928 at Para No. 0039) ("[At step 412, failure of the primary component link is detected. The response is a switchover of the LSP traffic to the secondary component link at step 414. Quality of service is maintained since the secondary component link also has sufficient guaranteed bandwidth. There is no signaling to the head-end to implement the switch-over. The secondary component link now becomes the primary component link. At step 416, node A selects another component link in the same link bundle to be the new secondary component link.]")
14[c]	digital logic configured to select said working entity and said protection entity from said plurality of transport entity descriptors, comprising:	Sivabalan '928 discloses digital logic configured to select said working entity and said protection entity from said plurality of transport entity descriptors, comprising. For example, Sivabalan '928 discloses sufficient bandwidth is reserved on each link and component link in the MPLS Traffic Engineering LSP to assure that the bandwidth guarantee is met. an LSP placed on a link bundle can operate in one of two different modes: in the first mode a single secondary link is assigned as a backup for the primary component link and in the second mode, one primary component link and one or more secondary component links are assigned to LSP X. When apparatus comprises means to identify component link with sufficient bandwidth to accommodate the bandwidth requirement of LSP and designates it as a primary link for LSP. A person of ordinary skill in the art would understand from these disclosures that the switchover logic would be stored as a data structure in digital computer memory. Sivabalan '928 at Claim 11 ("[11. A computer program product for operating a first node in a label switched network, said computer program product comprising:

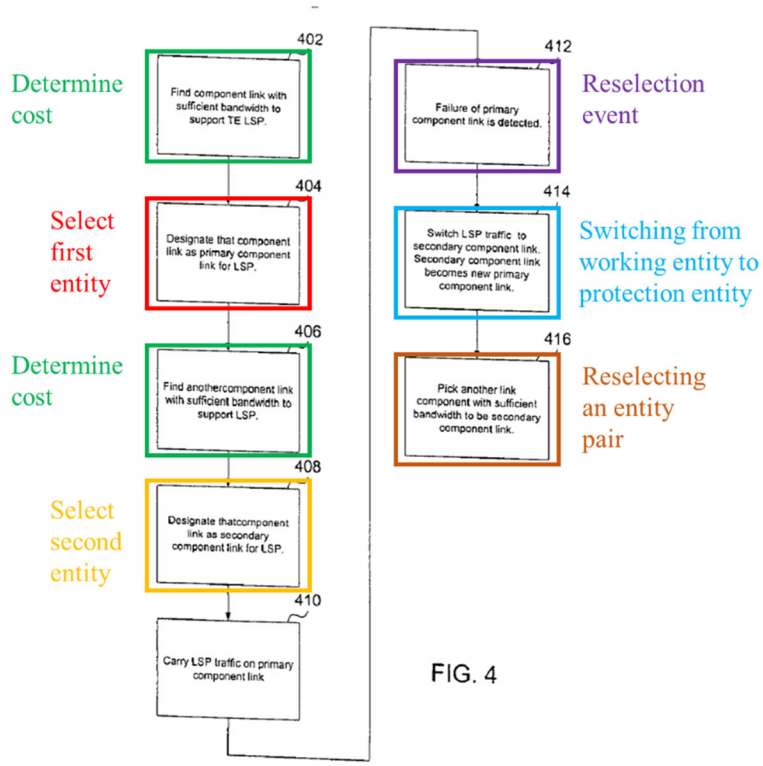
Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 226 of 698

No.	'821 Patent Claim 14	Sivabalan '928
		code that establishes a link bundle comprising a plurality of component links between said first node and a second node; code that, upon establishment of an LSP including said first node and said second node, selects a first component link of said link bundle as a primary component link to assign to said LSP; code that selects one or more component links of said link bundle other than said first component link as one or more secondary component links to assign to said LSP; code that sends traffic of said LSP over said primary component link until a failure of said primary component link; code that, after failure of said primary component link, sends traffic of said LSP over said one or more secondary component links instead of said primary component link; and a computer-readable medium that holds the codes.]") Sivabalan '928 at Claim 21 ("[21. A network device that operates a first node in a label switched network, said network device comprising: a processor; and a memory device storing instructions executed by said processor, said instructions comprising: code that establishes a link bundle comprising a plurality of component links between said first node and a second node; code that, upon establishment of an LSP including said first node and said second node, selects a first component link of said link bundle as a primary component link to assign to said LSP; code that selects one or more component links of said link bundle other than said first component link as one or more secondary component links to assign to said LSP; code that sends traffic of said LSP over said primary component link until a failure of said primary component link; and code that, after failure of said primary component link, sends traffic of said LSP over said one or more secondary component links instead of said primary component link.]") Sivabalan '928 at Claim 22 ("[22. Apparatus for operating a first node in a label switched network, said apparatus comprising:

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 227 of 698

No.	'821 Patent Claim 14	Sivabalan '928
		means for establishing a link bundle comprising a plurality of component links between said first node and a second node; means for, upon establishment of an LSP including said first node and said second node, selecting a first component link of said link bundle as a primary component link to assign to said LSP; means for selecting one or more component links of said link bundle other than said first component link as one or more secondary component links to assign to said LSP; means for sending traffic of said LSP over said primary component link until a failure of said primary component link; and means for, after failure of said primary component link, sending traffic of said LSP over said one or more secondary component links instead of said primary component link.]") Sivabalan '928 at Para No. 0035-0037 ("[According to embodiments of the present invention, an LSP placed on a link bundle can operate in one of two different modes: In the first mode, a single secondary component link is assigned as backup for the primary component link. Sufficient bandwidth to accommodate the LSP is reserved on this single secondary component link. The amount of available bandwidth on the secondary component link is reduced accordingly. FIG. 2 illustrates an example of this first mode where an MPLS Traffic Engineering LSP X employs link bundle 100 between nodes A and B. Node A ties LSP X to a primary component link AB1 and a secondary component link AB3. Further details of this first mode of operation will be described with reference to FIG. 4. In a second mode, one or more secondary component links are assigned to LSP X. Bandwidth is, however, not reserved on these secondary component links. This mode of operation is illustrated in FIG. 3 where LSP X is assigned to a primary component link AB1 and three secondary component links: AB2, AB3, and AB4. Further details of this second mode of operation will be described with reference to FIG. 5.]")

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 228 of 698

No.	'821 Patent Claim 14	Sivabalan '928
		 FIG. 4 Fig. 4 (annotation added) Sivabalan '928 at Para No. 0038 ("[FIG. 4 is a flow chart describing steps of component link operation according to the first mode as provided by embodiments of the present invention. At step 402, node A receives a request to establish a Traffic Engineering LSP using link bundle 100. Node A identifies a component link with sufficient bandwidth to accommodate the requirements of this Traffic Engineering LSP. At step 404, this component link is

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 229 of 698

No.	'821 Patent Claim 14	Sivabalan '928
		designated as the primary component link for the LSP. At step 406, node A identifies another component link with sufficient bandwidth to support the LSP. This component link is designated at step 408 as the secondary component link for the LSP. Information about the primary and secondary component links is programmed into the label forwarding table entry or entries corresponding to the LSP. At step 410, packets belonging to the LSP are carried on the primary component link.]")
14[d]	logic configured to determine a probability of concurrent failure of said working entity and said protection entity;	Sivabalan '928 discloses logic configured to determine a probability of concurrent failure of said working entity and said protection entity. *See* 14[b]-14[c] *supra*.
14[e]	logic configured to determine an entity cost of said plurality of transport entity descriptors: and	Sivabalan '928 discloses logic configured to determine an entity cost of said plurality of transport entity descriptors. Sivabalan '928 at Claim 11 ("[11. A computer program product for operating a first node in a label switched network, said computer program product comprising: code that establishes a link bundle comprising a plurality of component links between said first node and a second node; code that, upon establishment of an LSP including said first node and said second node, selects a first component link of said link bundle as a primary component link to assign to said LSP; code that selects one or more component links of said link bundle other than said first component link as one or more secondary component links to assign to said LSP; code that sends traffic of said LSP over said primary component link until a failure of said primary component link; code that, after failure of said primary component link, sends traffic of said LSP over said one or more secondary component links instead of said primary component link; and a computer-readable medium that holds the codes.]")

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 230 of 698

No.	'821 Patent Claim 14	Sivabalan '928
		Sivabalan '928 at Claim 21 ("[21. A network device that operates a first node in a label switched network, said network device comprising: a processor; and a memory device storing instructions executed by said processor, said instructions comprising: code that establishes a link bundle comprising a plurality of component links between said first node and a second node; code that, upon establishment of an LSP including said first node and said second node, selects a first component link of said link bundle as a primary component link to assign to said LSP; code that selects one or more component links of said link bundle other than said first component link as one or more secondary component links to assign to said LSP; code that sends traffic of said LSP over said primary component link until a failure of said primary component link; and code that, after failure of said primary component link, sends traffic of said LSP over said one or more secondary component links instead of said primary component link.]") Sivabalan '928 at Claim 22 ("[22. Apparatus for operating a first node in a label switched network, said apparatus comprising: means for establishing a link bundle comprising a plurality of component links between said first node and a second node; means for, upon establishment of an LSP including said first node and said second node, selecting a first component link of said link bundle as a primary component link to assign to said LSP; means for selecting one or more component links of said link bundle other than said first component link as one or more secondary component links to assign to said LSP; means for sending traffic of said LSP over said primary component link until a failure of said primary component link; and means for, after failure of said primary component link, sending traffic of said LSP over said one or more secondary component links instead of said primary component link.]") *See Supra at* 1[b].

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 231 of 698

No.	'821 Patent Claim 14	Sivabalan '928
14[f]	logic configured to reselect said working entity and said protection entity from said plurality of transport entity descriptors upon a reselection event,	Sivabalan '928 discloses logic configured to reselect said working entity and said protection entity from said plurality of transport entity descriptors upon a reselection event. Sivabalan '928 at Claim 11 ("[11. A computer program product for operating a first node in a label switched network, said computer program product comprising: code that establishes a link bundle comprising a plurality of component links between said first node and a second node; code that, upon establishment of an LSP including said first node and said second node, selects a first component link of said link bundle as a primary component link to assign to said LSP; code that selects one or more component links of said link bundle other than said first component link as one or more secondary component links to assign to said LSP; code that sends traffic of said LSP over said primary component link until a failure of said primary component link; code that, after failure of said primary component link, sends traffic of said LSP over said one or more secondary component links instead of said primary component link; and a computer-readable medium that holds the codes.]") Sivabalan '928 at Claim 21 ("[21. A network device that operates a first node in a label switched network, said network device comprising: a processor; and a memory device storing instructions executed by said processor, said instructions comprising: code that establishes a link bundle comprising a plurality of component links between said first node and a second node; code that, upon establishment of an LSP including said first node and said second node, selects a first component link of said link bundle as a primary component link to assign to said LSP; code that selects one or more component links of said link bundle other than said first component link as one or more secondary component links to assign to said LSP;

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 232 of 698

No.	'821 Patent Claim 14	Sivabalan '928
		code that sends traffic of said LSP over said primary component link until a failure of said primary component link; and code that, after failure of said primary component link, sends traffic of said LSP over said one or more secondary component links instead of said primary component link.]") Sivabalan '928 at Claim 22 ("[22. Apparatus for operating a first node in a label switched network, said apparatus comprising: means for establishing a link bundle comprising a plurality of component links between said first node and a second node; means for, upon establishment of an LSP including said first node and said second node, selecting a first component link of said link bundle as a primary component link to assign to said LSP; means for selecting one or more component links of said link bundle other than said first component link as one or more secondary component links to assign to said LSP; means for sending traffic of said LSP over said primary component link until a failure of said primary component link; and means for, after failure of said primary component link, sending traffic of said LSP over said one or more secondary component links instead of said primary component link.]") *See Supra at* 1[d].
14[g]	wherein said reselection event is selected from a group consisting of adding an entity to said plurality of transport entities, removing an entity from said plurality of transport entities, an operational status change for one	Sivabalan '928 discloses wherein said reselection event is selected from a group consisting of adding an entity to said plurality of transport entities, removing an entity from said plurality of transport entities, an operational status change for one of said plurality of transport entities, and a change in overall cost for one of said plurality of transport entities. *See Supra at* 1[e].

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 233 of 698

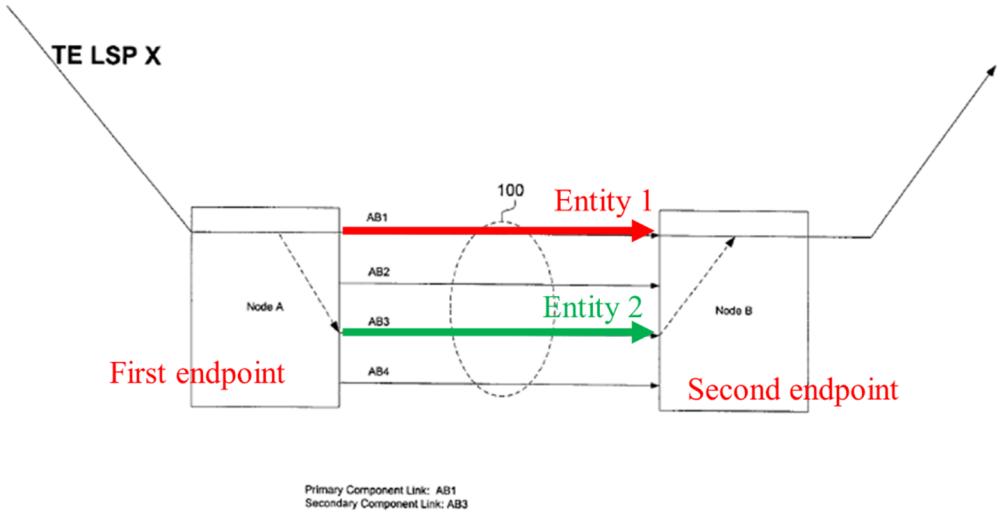
No.	'821 Patent Claim 14	Sivabalan '928
	of said plurality of transport entities, and a change in overall cost for one of said plurality of transport entities.	

No.	'821 Patent Claim 15	Sivabalan '928
15	The system of claim 14 wherein said entity protection switch comprises a 1:1 switch.	Sivabalan '928 discloses the system of claim 14 wherein said entity protection switch comprises a 1:1 switch. For example, Sivabalan '928 discloses protection by a single protection path where the bridge at the head end is not permanent and in the event of the primary path failing, it switches to the alternate path. At least based upon Orckit's apparent understanding of the claim scope as reflected in its Infringement Disclosures, this limitation is met. Sivabalan '928 at Para 0034 ("A local admission control mechanism at node A determines whether or not to accept the LSP and picks a component link for use by the LSP. The component link should have sufficient bandwidth available to support the LSP's requirements. The component link chosen in this fashion is called a "primary" component link. The admission control mechanism further selects one or more other component links as secondary component links. As will be explained below, the secondary component links are used as backup links when the primary component link fails. The secondary component link information is preconfigured into the label forwarding table entry for the LSP so that the switchover can occur very quickly.") Sivabalan '928 at Para 0039 ("At step 412, failure of the primary component link is detected. The response is a switchover of the LSP traffic to the secondary component link at step 414. Quality of service is maintained since the secondary component link also has sufficient guaranteed bandwidth. There is no signaling to the head-end to implement the switch-over. The secondary component link now becomes the primary component link. At step 416, node

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 234 of 698

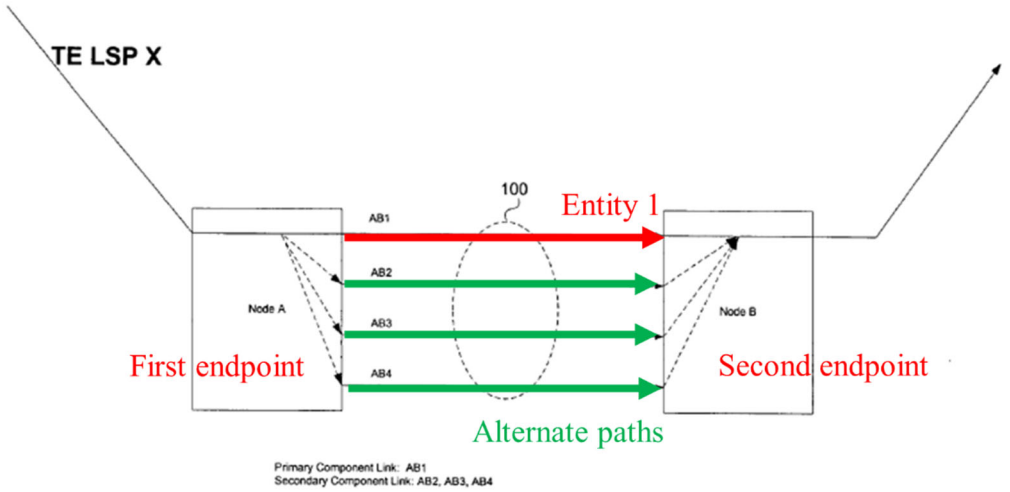
No.	'821 Patent Claim 15	Sivabalan '928
		A selects another component link in the same link bundle to be the new secondary component link.") Sivabalan '928 at Para 0042 ("At step 510, node A detects failure of the primary component link. The LSP traffic is switched immediately, at step 512, to the one or more secondary component links.") Sivabalan '928 at Para No. 0036 ("In the first mode, a single secondary component link is assigned as backup for the primary component link. Sufficient bandwidth to accommodate the LSP is reserved on this single secondary component link. The amount of available bandwidth on the secondary component link is reduced accordingly. FIG. 2 illustrates an example of this first mode where an MPLS Traffic Engineering LSP X employs link bundle 100 between nodes A and B. Node A ties LSP X to a primary component link AB1 and a secondary component link AB3.")

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 235 of 698

No.	'821 Patent Claim 15	Sivabalan '928
		 FIG. 2 Fig. 2 (annotation added)

No.	'821 Patent Claim 16	Sivabalan '928
16	The system of claim 14 wherein said entity protection switch comprises a 1+1 switch.	Sivabalan '928 discloses the system of claim 14 wherein said entity protection switch comprises a 1+1 switch. For example, Sivabalan '928 discloses protection where each working channel has a continuously active backup path. At least based upon Orckit's apparent understanding of the claim scope as reflected in its Infringement Disclosures, this limitation is met. Sivabalan '928 at Para No. 0037 ("In a second mode, one or more secondary component links are assigned to LSP X. Bandwidth is, however, not reserved on these secondary component links. This mode of operation is illustrated in FIG. 3 where LSP X is assigned to a primary component link AB1 and three secondary component links: AB2, AB3, and AB4.")

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 236 of 698

No.	'821 Patent Claim 16	Sivabalan '928
		 FIG. 3 Fig. 3 (annotation added) Sivabalan '928 at Claim 6 ("The method of claim 1 wherein no bandwidth is allocated on said one or more secondary component links for use by said LSP.")

No.	'821 Patent Claim 17	Sivabalan '928
17[preamble]	Non-transitory computer readable media configured to perform a method comprising the steps of:	Sivabalan '928 discloses non-transitory computer readable media configured to perform a method comprising the steps of. For example, Sivabalan '928 discloses "a computer program" product for operating a first node for transparent re-routing of MPLS traffic engineering LSPs within a link bundle and selection of primary or secondary/ backup component links. Sivabalan '928 at Claim 11 ("[11. A computer program product for operating a first node in a label switched network, said computer program product comprising:

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 237 of 698

No.	'821 Patent Claim 17	Sivabalan '928
		code that establishes a link bundle comprising a plurality of component links between said first node and a second node; code that, upon establishment of an LSP including said first node and said second node, selects a first component link of said link bundle as a primary component link to assign to said LSP; code that selects one or more component links of said link bundle other than said first component link as one or more secondary component links to assign to said LSP; code that sends traffic of said LSP over said primary component link until a failure of said primary component link; code that, after failure of said primary component link, sends traffic of said LSP over said one or more secondary component links instead of said primary component link; and a computer-readable medium that holds the codes.]") Sivabalan '928 at Claim 21 ("[21. A network device that operates a first node in a label switched network, said network device comprising: a processor; and a memory device storing instructions executed by said processor, said instructions comprising: code that establishes a link bundle comprising a plurality of component links between said first node and a second node; code that, upon establishment of an LSP including said first node and said second node, selects a first component link of said link bundle as a primary component link to assign to said LSP; code that selects one or more component links of said link bundle other than said first component link as one or more secondary component links to assign to said LSP; code that sends traffic of said LSP over said primary component link until a failure of said primary component link; and code that, after failure of said primary component link, sends traffic of said LSP over said one or more secondary component links instead of said primary component link.]") Sivabalan '928 at Claim 22 ("[22. Apparatus for operating a first node in a label switched network, said apparatus comprising:

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 238 of 698

No.	'821 Patent Claim 17	Sivabalan '928
		means for establishing a link bundle comprising a plurality of component links between said first node and a second node; means for, upon establishment of an LSP including said first node and said second node, selecting a first component link of said link bundle as a primary component link to assign to said LSP; means for selecting one or more component links of said link bundle other than said first component link as one or more secondary component links to assign to said LSP; means for sending traffic of said LSP over said primary component link until a failure of said primary component link; and means for, after failure of said primary component link, sending traffic of said LSP over said one or more secondary component links instead of said primary component link.]")
17[a]	providing a plurality of MPLS transport entities between a first endpoint and a second endpoint;	Sivabalan '928 discloses providing a plurality of multiprotocol label switching (MPLS) transport entities between a first endpoint and a second endpoint. *See Supra at* 1[a].
17[b]	determining an overall cost for each entity pair of said plurality of entities;	Sivabalan '928 discloses determining an overall cost for each entity pair of said plurality of entities. *See Supra at* 1[b].
17[c]	selecting an entity pair from said plurality of transport entities based at least in part upon said overall cost; and	Sivabalan '928 discloses selecting an entity pair from said plurality of transport entities based at least in part upon said overall cost. *See Supra at* 1[c].
17[d]	if an entity pair reselection event occurs, reselecting said entity pair from the group consisting of	Sivabalan '928 discloses if an entity pair reselection event occurs, reselecting said entity pair from the group consisting of said entity pair and a replacement entity pair comprising at least one entity distinct from the entities of said entity pair. *See Supra at* 1[d].

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 239 of 698

No.	'821 Patent Claim 17	Sivabalan '928
	said entity pair and a replacement entity pair comprising at least one entity distinct from the entities of said entity pair,	
17[e]	wherein said entity pair reselection event is selected from a group consisting of adding an entity to said plurality of transport entities, removing an entity from said plurality of transport entities, an operational status change for one of said plurality of transport entities, and a change in overall cost for one of said plurality of transport entities.	Sivabalan '928 discloses "entity pair reselection event is selected from a group consisting of" & "an operational status change for one of said plurality of transport entities". *See Supra at* 1[e].

No.	'821 Patent Claim 18	Sivabalan '928
18[preamble]	The non-transitory computer readable media of claim 17, wherein said step of selecting an entity pair further comprises:	Sivabalan '928 discloses a non-transitory computer readable media with a entity pair selection step. *See Supra at* 1[b] and 17.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 240 of 698

No.	'821 Patent Claim 18	Sivabalan '928
18[a]	selecting a working entity from said plurality of transport entities;	Sivabalan '928 discloses selecting a working entity from said plurality of transport entities. *See Supra at* 2[a].
18[b]	selecting a protection entity from said plurality of transport entities; and	Sivabalan '928 discloses selecting a protection entity from said plurality of transport entities. *See Supra at* 2[b].
18[c]	selecting an active entity from the set consisting of said working entity and said protection entity.	Sivabalan '928 discloses the step of selecting an active entity from the set consisting of said working entity and said protection entity. *See Supra at* 3.

No.	'821 Patent Claim 19	Sivabalan '928
19	The non-transitory readable media of claim 18, wherein said step of selecting an entity pair further comprises minimizing an overall cost function.	Sivabalan '928 discloses the step of selecting an entity pair further comprising minimalizing an overall cost function. *See Supra at 4*.

No.	'821 Patent Claim 20	Sivabalan '928
20[preamble]	The non-transitory readable media of claim 19, wherein said overall cost function comprises:	Sivabalan '928 discloses the non-transitory readable media of claim 19, wherein said overall cost function comprises. *See* Claim 19 *supra*.
20[a]	minimizing a probability of	Sivabalan '928 discloses minimizing a probability of concurrent failure of said protection entity and said working entity.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 241 of 698

No.	'821 Patent Claim 20	Sivabalan '928
	concurrent failure of said protection entity and said working entity; and	*See* Claim 5 *supra*.
20[b]	a predefined metric selected from the group consisting of interior gateway protocol (IGP) and traffic engineering (TE).	Sivabalan '928 discloses a predefined metric selected from the group consisting of interior gateway protocol (IGP) and traffic engineering (TE). See Claim 7 *supra*.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 242 of 698

EXHIBIT B-6
Defendant's Preliminary Invalidity Contentions
Orckit Corporation v. Cisco Systems, Inc., 2:22-cv-00276-JRG-RSP

Chart for U.S. Patent 8,830,821 ("the '821 Patent")
U.S. Patent No. US 2006/0250948 A1 to Zamfir, et. al. ("Zamfir '948")

As shown in the chart below, all Asserted Claims of the '821 Patent are invalid under (1) 35 U.S.C. §§ 102 (a), (b), (e), and (g) because Zamfir '948 meets each element of those claims, and/or (2) 35 U.S.C. § 103 because Zamfir '948 renders those claims obvious either alone, or in combination with the knowledge of a person having ordinary skill in the art, and in further combination with the references specifically identified below and in the following claim chart and/or one or more references identified in Defendant's Preliminary Invalidity Contentions. The following quotations and diagrams come from Zamfir '948 titled "Dynamic Path Protection In An Optical Network", which was filed on May 9, 2005, and published on November 9, 2006.

Motivations to combine the disclosures in Zamfir '948 with disclosures in other publications known in the art, as explained in this chart, include at least the similarity in subject matter between the references to the extent they concern methods of selecting paths in a network and/or determining the cost of potential paths in a network. Insofar as the references cite other patents or publications, or suggest additional changes, one of ordinary skill in the art would look beyond a single reference to other references in the field.

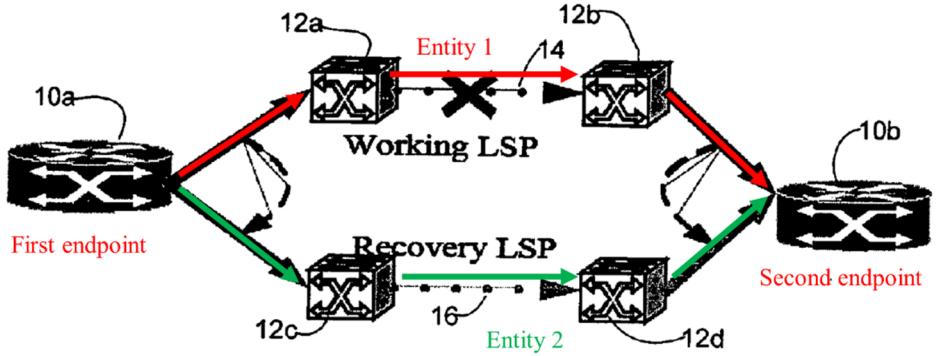
These invalidity contentions are based on Defendant's present understanding of the Asserted Claims, and Orckit's apparent construction of the claims in its November 3, 2022 Disclosure of Asserted Claims and Infringement Contentions Pursuant to P.R. 3-1, and Orckit's January 19, 2023 First Amended Disclosure of Asserted Claims and Infringement Contentions Pursuant to P.R. 3-1 (Orckit's "Infringement Disclosures"), which is deficient at least insofar as it fails to cite any documents or identify accused structures, acts, or materials in the Accused Products with particularity. Defendant does not agree with Orckit's application of the claims, or that the claims satisfy the requirements of 35 U.S.C. § 112. Defendant's contentions herein are not, and should in no way be seen as, admissions or adoptions as to any particular claim scope or construction, or as any admission that any particular element is met by any accused product in any particular way. Defendant objects to any attempt to imply claim construction from this chart. Defendant's prior art invalidity contentions are made in a variety of alternatives and do not represent Defendant's agreement or view as to the meaning, definiteness, written description support for, or enablement of any claim contained therein.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 243 of 698

The following contentions are subject to revision and amendment pursuant to Federal Rule of Civil Procedure 26(e), the Local Rules, and the Orders of record in this matter subject to further investigation and discovery regarding the prior art and the Court's construction of the claims at issue.

No.	'821 Patent Claim 1	Zamfir '948
1[preamble]	An entity selection method performed by a network device, comprising the steps of:	Zamfir '948 discloses an entity selection method performed by a network device, comprising the steps of. For example, Zamfir '948 discloses a system for dynamic end-to-end path protection in an optical network which comprises a processor on a headend router operable to create two or more LSP paths (e.g., G-LSP and TE-LSP) within the optical network. Zamfir '948 at Page No. 08, Para No. 0008 ("[In another aspect of the invention, a method for dynamic end-to-end path protection in an optical network generally comprises the establishment of a number of working and protecting paths and selecting new paths for a subset of these paths when an event causing network topology changes occurs. The new paths are selected such that the overall diversity and optimality degree are maintained to a level allowed by the new topology.]")

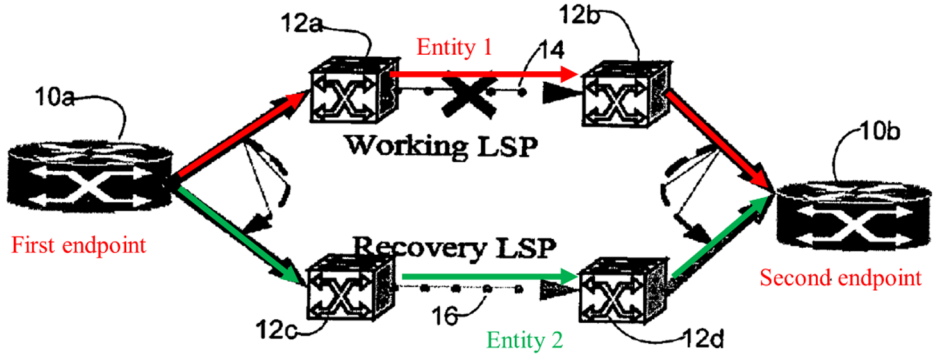
Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 244 of 698

No.	'821 Patent Claim 1	Zamfir '948
		 FIG. 1A Fig. 1A (annotation added) Zamfir '948 at Page No. 09, Para No. 0021 ("[FIG. 1A illustrates an example of a network in which data transmission occurs on generalized label-switched paths (G-LSPs). The G-LSPs are automatically setup and torn down by means of a signaling protocol, as is well known by those skilled in the art. (In the following description, the term LSP is also used to refer to generalized LSPs.) FIG. 1A illustrates end-to-end protection with one working path 14 (between nodes 12 a and 12 b) and one protecting (recovery) path 16 (between nodes 12 c and 12 d). The working and protecting paths 14, 16 are preferably signaled as two LSPs that belong to the same tunnel. The working LSP 14 transports traffic between routers 10 a and 10 b and the protecting LSP 16 is used to transport traffic when the working LSP fails. End-to-end recovery provides recovery of the entire LSP from its source (ingress router end-point 10 a) to its destination (egress router end-point 10 b).]") Zamfir '948 at Page No. 10, Para No. 0034 ('[Dynamic FRR extends link protection schemes, such as FRR (Fast Reroute), to handle GMPLS LSPs. A GMPLS LSP (G-LSP) is protected

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 245 of 698

No.	'821 Patent Claim 1	Zamfir '948
		using a backup tunnel (TE-LSP) in a similar way as to how FRR is realized for physical links.]") Zamfir '948 at Claim 17 "[17. A system for dynamic protection of virtual links for paths of an optical network in communication with an IP network, the system comprising: means for establishing two or more paths within the optical network; and means for grouping two or more of said paths in a dynamic bundle, the grouping selected at a control plane level; wherein said bundle of paths are recognized as a single routing adjacency in the IP network such that as long as one or more paths within the bundle is operating the routing adjacency in the IP network is not affected by changes within said bundle.]") Zamfir '948 at Claim 19 ("[19. A system for dynamic protection end-to-end path protection in an optical network in communication with an IP network, the system comprising: a processor operable to create two or more paths within the optical network and group two or more of said paths in a bundle, said bundle comprising at least one working path and at least one protecting path each extending from an ingress router to an egress router; and a protection module configured to dynamically identify a physical interface associated with an optical tunnel interface for a new path and replace one of said paths in said bundle with said new path.]")
1[a]	providing a plurality of multi-protocol label switching (MPLS) transport entities between a first endpoint and a second endpoint;	Zamfir '948 discloses providing a plurality of multi-protocol label switching (MPLS) transport entities between a first endpoint and a second endpoint. For example, Zamfir '948 discloses a processor on a headend router that creates two or more paths/ label switched paths (LSP) within the optical network between the same source and destination nodes.

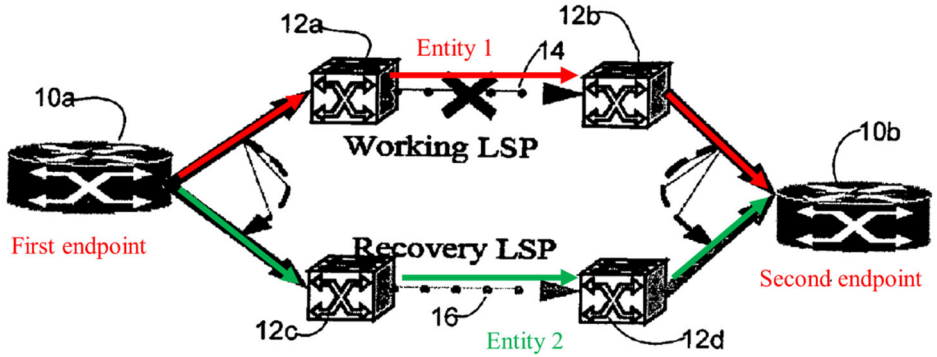
Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 246 of 698

No.	'821 Patent Claim 1	Zamfir '948
		 FIG. 1A Fig. 1A (annotation added) Zamfir '948 at Page No. 09, Para No. 0021 ("[FIG. 1A illustrates an example of a network in which data transmission occurs on generalized label-switched paths (G-LSPs). The G-LSPs are automatically setup and torn down by means of a signaling protocol, as is well known by those skilled in the art. (In the following description, the term LSP is also used to refer to generalized LSPs.) FIG. 1A illustrates end-to-end protection with one working path 14 (between nodes 12 a and 12 b) and one protecting (recovery) path 16 (between nodes 12 c and 12 d). The working and protecting paths 14, 16 are preferably signaled as two LSPs that belong to the same tunnel. The working LSP 14 transports traffic between routers 10 a and 10 b and the protecting LSP 16 is used to transport traffic when the working LSP fails. End-to-end recovery provides recovery of the entire LSP from its source (ingress router end-point 10 a) to its destination (egress router end-point 10 b).]") Zamfir '948 Page No. 09, Para No. 0022 ("[It is to be understood that the arrangement shown in FIG. 1A is only one example and the invention described herein may be applied to a network having any number of working or recovery LSPs.]")

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 247 of 698

No.	'821 Patent Claim 1	Zamfir '948
		Zamfir '948 at Page No. 09, Para No. 0027 ("[A number of LSPs are established during GMPLS tunnel setup. The traffic may be arranged or duplicated any number of ways, as is well known by those skilled in the art. The paths are grouped into one or more bundles. With link bundling of the GMPLS interface, the component membership to the bundle is decided upon by the GMPLS control plane, which brings in the data plane aspect of dynamic membership. The LSPs in a bundle are established between the same source and destination nodes and, based on configuration or enforced SLA, they have a certain degree of diversity and optimality.]") Zamfir '948 at Claim 19 ("[19. A system for dynamic protection end-to-end path protection in an optical network in communication with an IP network, the system comprising: a processor operable to create two or more paths within the optical network and group two or more of said paths in a bundle, said bundle comprising at least one working path and at least one protecting path each extending from an ingress router to an egress router; and a protection module configured to dynamically identify a physical interface associated with an optical tunnel interface for a new path and replace one of said paths in said bundle with said new path.]")
1[b]	determining an overall cost for each entity pair of said plurality of entities;	Zamfir '948 discloses determining an overall cost for each entity pair of said plurality of entities. For example, Zamfir '948 discloses that the LSPs in a bundle are established (G-LSP and TE-LSP grouped in a bundle) between the same source and destination nodes based on configuration or enforced SLA and have a certain degree of diversity e.g.., SLRG diversity and optimality, and an associated cost is determined for each LSP. Thus, at least under the apparent claim scope alleged by Orckit's Infringement Disclosures, this limitation is met. To the extent that the IETF MPLS-TP System is found to not meet this limitation, determining an overall cost for each entity pair of said plurality of entities would have been obvious to a person having ordinary skill in the art, as explained below.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 248 of 698

No.	'821 Patent Claim 1	Zamfir '948
		Zamfir '948 at Page No. 09, Para No. 0021 ("[FIG. 1A illustrates an example of a network in which data transmission occurs on generalized label-switched paths (G-LSPs). The G-LSPs are automatically setup and torn down by means of a signaling protocol, as is well known by those skilled in the art. (In the following description, the term LSP is also used to refer to generalized LSPs.) FIG. 1A illustrates end-to-end protection with one working path 14 (between nodes 12 a and 12 b) and one protecting (recovery) path 16 (between nodes 12 c and 12 d). The working and protecting paths 14, 16 are preferably signaled as two LSPs that belong to the same tunnel. The working LSP 14 transports traffic between routers 10 a and 10 b and the protecting LSP 16 is used to transport traffic when the working LSP fails. End-to-end recovery provides recovery of the entire LSP from its source (ingress router end-point 10 a) to its destination (egress router end-point 10 b).]")  FIG. 1A Fig. 1A (annotation added)

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 249 of 698

No.	'821 Patent Claim 1	Zamfir '948
		Zamfir '948 at Page No. 09, Para No. 0022 ("[It is to be understood that the arrangement shown in FIG. 1A is only one example and the invention described herein may be applied to a network having any number of working or recovery LSPs.]") Zamfir '948 at Page No. 09, Para No. 0027 ("[A number of LSPs are established during GMPLS tunnel setup. The traffic may be arranged or duplicated any number of ways, as is well known by those skilled in the art. The paths are grouped into one or more bundles. With link bundling of the GMPLS interface, the component membership to the bundle is decided upon by the GMPLS control plane, which brings in the data plane aspect of dynamic membership. The LSPs in a bundle are established between the same source and destination nodes and, based on configuration or enforced SLA, they have a certain degree of diversity and optimality.]") Zamfir '948 at(Page No. 10, Para No. 0031 ("[In order to better illustrate the make-a-set-before-break-a-set process, the following example is provided. In this example, a dynamic bundle with two components LSP1 and LSP2 is created, with the two LSPs being setup with link and node diversity (FIG. 3): LSP1−R3, R1, R6, R∝(cost=30) LSP2−R3, R5, R2, R4 (cost=30) In this example, all links are assumed to have cost equal to 10, therefore LSP1 and LSP2 each have a cost equal to 30.]") Zamfir '948 at Page No. 10, Para No. 0034 ("[Dynamic FRR extends link protection schemes, such as FRR (Fast Reroute), to handle GMPLS LSPs. A GMPLS LSP (G-LSP) is protected using a backup tunnel (TE-LSP) in a similar way as to how FRR is realized for physical links.]") Zamfir '948 at Page No. 10, Para No. 0035 ("[G-LSP and TE-LSP are preferably initially setup with a maximum diversity degree.]") In addition, at least under the apparent claim scope alleged by Orckit's Infringement Disclosures, Zamfir '948 in combination with (1) the knowledge of a person of ordinary skill in

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 250 of 698

No.	'821 Patent Claim 1	Zamfir '948
		the art, alone or in further combination with (2) each (individually, as well as one or more together) of the references identified in element 1[b] of Exhibit E-2 renders the claim, including the present limitation, obvious. Below are examples of two such references. For example, Juniper IOS Guide discloses "Fate Sharing" where a Juniper router computes a primary and secondary path using CSPF, CSPF identifies all fate-sharing groups that are associated with the primary path by identifying all links that the primary path transverses and checks each link in the traffic engineering database against the compiled group list. Juniper IOS Guide further discloses that if the link is a member of a group, the cost of the link is increased by the cost of the group and higher the cost value. At least based upon Orckit's apparent understanding of the claim scope as reflected in its Infringement Disclosures, this limitation is met. Juniper IOS Guide at Page No. 73 ("[Fate Sharing Fate sharing allows you to create a database of information that CSPF uses to compute one or more backup paths to use in case the primary path becomes unstable. The database describes the relationships between elements of the network, such as routers and links. You can specify one or more elements within a group. Through fate sharing, you can configure backup paths that minimize the number of shared links and fiber paths with the primary paths as much as possible, to ensure that if a fiber is cut, the minimum amount of data is lost and a path still exists to the destination. For a backup path to work optimally, it must not share links or physical fiber paths with the primary path, ensuring that a single point of failure will not affect the primary and backup paths simultaneously. For more information about fate sharing, see the Junos OS Routing Protocols Configuration Guide.]") Juniper IOS Guide at Page No. 94 ("[Configuring the Ingress Router for MPLS-Signaled LSPs MPLS-signaled label-switched paths (LSPs) run from a specific ingress router to a specific egress router. For basic MPLS-signaled LSP function, you must configure the ingress router, but do not have to configure any other routers. To configure signaled LSPs, perform the following tasks on the ingress router:

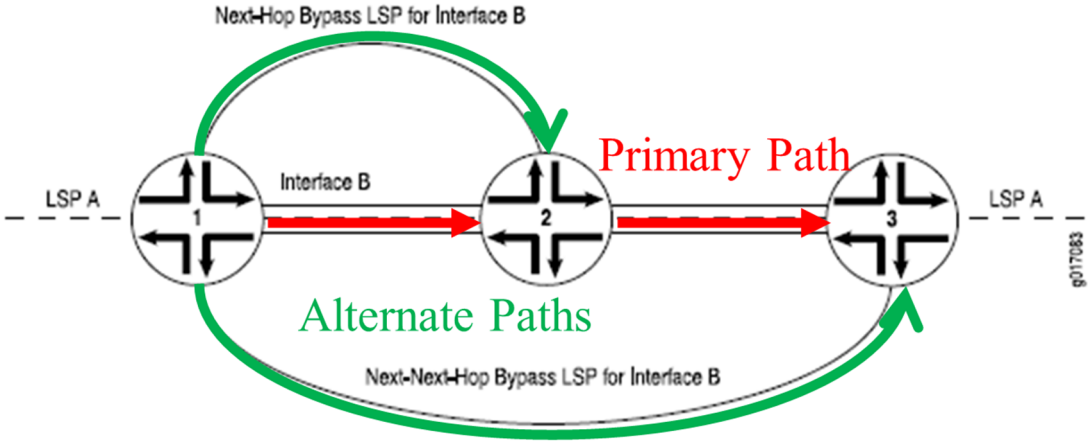
Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 251 of 698

No.	'821 Patent Claim 1	Zamfir '948
		• Creating Named Paths on page 56 • Configuring Alternate Backup Paths Using Fate Sharing on page 58.]") Juniper IOS Guide at Page No. 96 ("[Configuring Alternate Backup Paths Using Fate Sharing You can create a database of information that Constrained Shortest Path First (CSPF) uses to compute one or more backup paths in case the primary path becomes unstable. The database describes the relationships between elements of the network, such as routers and links. Because these network elements share the same fate, this relationship is called fate sharing. You can configure backup paths that minimize the number of shared links and fiber paths with the primary paths as much as possible to ensure that, if a fiber is cut, the minimum amount of data is lost and a path still exists to the destination. For a backup path to work optimally, it must not share links or physical fiber paths with the primary path. This ensures that a single point of failure will not affect the primary and backup paths at the same time.]") Juniper IOS Guide at Page No. 97 ("[Fate-sharing groups contain three types of objects: • Point-to-point links—Identified by the IP addresses at each end of the link. Unnumbered point-to-point links are typically identified by borrowing IP addresses from other interfaces. Order is not important; from 1.2.3.4 to 1.2.3.5 and from 1.2.3.5 to 1.2.3.4 have the same meaning. • Non-point-to-point links—Include links on a LAN interface (such as Gigabit Ethernet interfaces) or nonbroadcast multiaccess (NBMA) interfaces (such as Asynchronous Transfer Mode [ATM] or Frame Relay). You identify these links by their individual interface address. For example, if the LAN interface 192.168.200.0/24 has four routers attached to it, each router link is individually identified: from 192.168.200.1; # LAN interface of router 1 from 192.168.200.2; # LAN interface of router 2 from 192.168.200.3; # LAN interface of router 3 from 192.168.200.4; # LAN interface of router 4 You can list the addresses in any order. • A router node—Identified by its configured router ID.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 252 of 698

No.	'821 Patent Claim 1	Zamfir '948
		An object can be in any number of groups, and a group can contain any number of objects. Each group has a configurable cost attributed to it, which represents the level of impact this group has on CSPF computations. The higher the cost, the less likely a backup path will share with the primary path any objects in the group. The cost is directly comparable to traffic engineering metrics. By default, the cost is 1. Changing the fate-sharing database does not affect established LSPs until the next reoptimization of CSPF. The fate-sharing database does influence fast-reroute computations.]") Juniper IOS Guide at Page No. 97 ("[When CSPF computes the primary paths of an LSP (or secondary paths when the primary path is not active), it ignores the fate-sharing information. You always want to find the best possible path (least IGP cost) for the primary path. When CSPF computes a secondary path while the primary path (of the same LSP) is active, the following occurs: 1. CSPF identifies all fate-sharing groups that are associated with the primary path. CSPF does this by identifying all links and nodes that the primary path traverses and compiling group lists that contain at least one of the links or nodes. CSPF ignores the ingress and egress nodes in the search. 2. CSPF checks each link in the traffic engineering database against the compiled group list. If the link is a member of a group, the cost of the link is increased by the cost of the group. If a link is a member of multiple groups, all group costs are added together. 3. CSPF performs the check for every node in the traffic engineering database, except the ingress and egress node. Again, a node can belong to multiple groups, so costs are additive. 4. The router performs regular CSPF computation with the adjusted topology.]") Juniper IOS Guide at Page No. 330 ("In Figure 23 on page 292, both node protection and link protection are enabled on Interface B on Router 1. Both node protection and link protection are also enabled on LSP A, an LSP that traverses the link transiting Router 1, Router 2, and Router 3. If the link between Router 1 and Router 2 fails, traffic from LSP A is quickly switched to the next-hop bypass LSP generated by link protection. If Router 2 suffers a hardware or software failure, traffic from LSP A is switched to the next-next-hop bypass LSP generated by node protection.")

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 253 of 698

No.	'821 Patent Claim 1	Zamfir '948
		Figure 23: Node Protection Creating a Next-Next-Hop Bypass LSP  Fig. 23 (annotation added). For example, Cisco IOS Guide discloses that the path an LSP uses is determined by the LSP resource requirements and network resources, such as bandwidth. Cisco IOS Guide further discloses the MPLS TE SRLG feature that enhances backup tunnel path selection so a backup tunnel can avoid using links that are in the same SRLG as the interfaces it is protecting. Thus, at least under the apparent claim scope alleged by Orckit's Infringement Disclosures, this limitation is met. Cisco IOS Guide at Page No. 253 ("[Path Option Selection for MPLS TE Tunnel LSPs This section contains the following topics about path option selection for MPLS TE Tunnel LSPs: • Constraint-Based Routing and Path Option Selection, page 5 • Tunnel Reoptimization and Path Option Selection, page 5 • Path Option Selection with Bandwidth Override, page 6 Constraint-Based Routing and Path Option Selection MPLS traffic engineering automatically establishes and maintains LSPs across the backbone by using the Resource Reservation Protocol (RSVP). The path that an LSP uses is determined

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 254 of 698

No.	'821 Patent Claim 1	Zamfir '948
		by the LSP resource requirements and network resources, such as bandwidth. Traffic engineering tunnels are calculated at the LSP head based on a fit between required and available resources (constraint-based routing). Without the Path Option for Bandwidth Override feature, a TE tunnel establishes an LSP based on dynamic or explicit path options in order of preference. However, the bandwidth and other attributes configured on the TE tunnel allow the setup of an LSP only if LSP path options satisfy the constraints. If a path cannot be found that satisfies the configured path options, then the tunnel is not set up.]") Cisco IOS Guide at Page No. 639 ("[Bandwidth Protection NHOP and NNHOP backup tunnels can be used to provide bandwidth protection for rerouted LSPs. This is referred to as backup bandwidth. You can associate backup bandwidth with NHOP or NNHOP backup tunnels. This informs the router of the amount of backup bandwidth a particular backup tunnel can protect. When a router maps LSPs to backup tunnels, bandwidth protection ensures that an LSP uses a given backup tunnel only if there is sufficient backup bandwidth. The router selects which LSPs use which backup tunnels in order to provide maximum bandwidth protection. That is, the router determines the best way to map LSPs onto backup tunnels in order to maximize the number of LSPs that can be protected. For information about mapping tunnels and assigning backup bandwidth, see the "Backup Tunnel Selection Procedure" section on page 10. LSPs that have the "bandwidth protection desired" bit set have a higher right to select backup tunnels that provide bandwidth protection; that is, those LSPs can preempt other LSPs that do not have that bit set. For more information, see the "Prioritizing Which LSPs Obtain Backup Tunnels with Bandwidth Protection" section on page 8.]") Cisco IOS Guide at Page No. 640 ("[Backup Tunnel Support Backup tunnel support has the following capabilities: • Backup Tunnels Can Terminate at the Next-Next Hop to Support FRR, page 6 • Multiple Backup Tunnels Can Protect the Same Interface, page 6 • Backup Tunnels Provide Scalability, page 7.]") Cisco IOS Guide at Page 640 ("[Multiple Backup Tunnels Can Protect the Same Interface

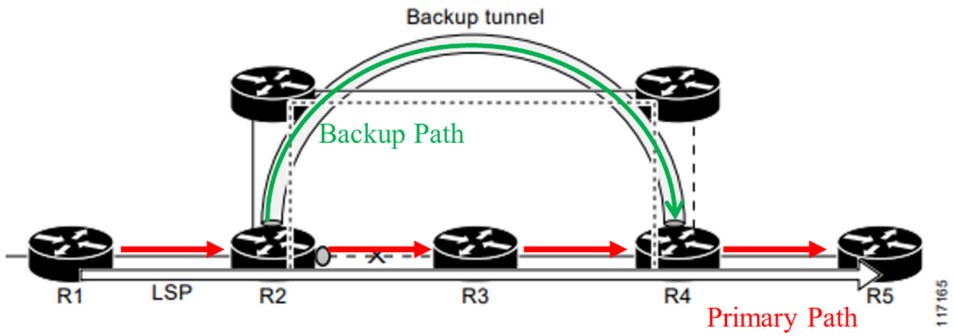
Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 255 of 698

No.	'821 Patent Claim 1	Zamfir '948
		There is no limit (except memory limitations) to the number of backup tunnels that can protect a given interface. In many topologies, support for node protection requires supporting multiple backup tunnels per protected interface. These backup tunnels can terminate at the same destination or at different destinations. That is, for a given protected interface, you can configure multiple NHOP or NNHOP backup tunnels. This allows redundancy and load balancing.]") Cisco IOS Guide at Page No. 603 ("[Traffic Engineering Tunnels MPLS TE lets you build label switched paths (LSPs) across your network for forwarding traffic. MPLS TE LSPs let the headend of a TE tunnel control the path its traffic takes to a particular destination. This method is more flexible than forwarding traffic based only on a destination address. Interarea tunnels allow you to do the following: • Build TE tunnels between areas (interarea tunnels) • Build TE tunnels that start and end in the same area, on multiple areas on a router (intra-area tunnels) Some tunnels are more important than others. For example, you may have tunnels carrying VoIP traffic and tunnels carrying data traffic that are competing for the same resources. MPLS TE allows you to have some tunnels preempt others. Each tunnel has a priority, and more-important tunnels take precedence over less-important tunnels.]") Cisco IOS Guide at Page No. 641 ("[Backup Bandwidth Protection Backup bandwidth protection allows you to give LSPs carrying certain kinds of data (such as voice) priority for using backup tunnels.]") Cisco IOS Guide at Page No. 650 ("[For each protected link AB with a maximum reservable subpool value of n, there may be a path from node A to node B such that the difference between the maximum reservable global and the maximum reservable subpool is at least the value of n. If it is possible to find such paths for each link in the network, you can establish all the backup tunnels along such paths without any bandwidth reservations. If there is a single link failure, only one backup tunnel will use any link on its path. Because that path has at least

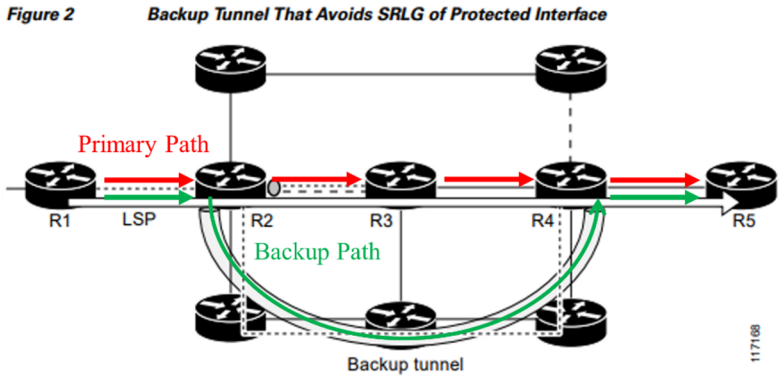
Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 256 of 698

No.	'821 Patent Claim 1	Zamfir '948
		n available bandwidth (in the global pool), assuming that marking and scheduling is configured to classify the subpool traffic into a priority queue, the subpool bandwidth is guaranteed. This approach allows sharing of the global pool bandwidth between backup tunnels protecting independent link failures. The backup tunnels are expected to be used for only a short period of time after a failure (until the headends of affected LSPs reroute those LSPs to other paths with available subpool bandwidth). The probability of multiple unrelated link failures is very small (in the absence of node or shared risk link group (SRLG) failures, which result in multiple link failures). Therefore, it is reasonable to assume that link failures are in practice independent with high probability. This "independent failure assumption" in combination with backup tunnels signaled without explicit bandwidth reservation enables efficient bandwidth sharing that yields substantial bandwidth savings.]") Cisco IOS Guide at Page No. 651 ("[A similar approach can be used for node and SRLG protection. However, the decision of where to put the backup tunnels is more complicated because both node and SRLG failures effectively result in the simultaneous failure of several links. Therefore, the backup tunnels protecting traffic traversing all affected links cannot be computed independently of each other.]") Cisco IOS Guide at Page No. 579 ("[MPLS Traffic Engineering Shared Risk Link Groups SRLGs refer to situations where links in a network share a common fiber (or a common physical attribute). If one link fails, other links in the group may fail too. Links in the group have a shared risk. Backup tunnels should avoid using links in the same SRLG as interfaces they are protecting. Otherwise, when the protected link fails the backup tunnel fails too.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 257 of 698

No.	'821 Patent Claim 1	Zamfir '948
		Figure 1 Backup Tunnel in the Same SRLG as the Interface It Is Protecting  Fig. 1 (annotation added) Figure 1 shows a primary label-switched path (LSP) from router R1 to router R5. The LSP protects against the failure of the R2-R3 link at R2 via a backup tunnel to R4. If the R2-R3 link fails, link protection reroutes the LSP along the backup tunnel. However, the R2-R3 link and one of the backup tunnel links are in the same SRLG. So if the R2-R3 link fails, the backup tunnel may fail too. The MPLS TE SRLG feature enhances backup tunnel path selection so a backup tunnel can avoid using links that are in the same SRLG as the interfaces it is protecting. There are two ways for a backup tunnel to avoid the SRLGs of its protected interface: • The router does not create the backup tunnel unless it avoids SRLGs of the protected interface.]") Cisco IOS Guide at Page No. 580 ("[Open Shortest Path First (OSPF) and Intermediate System-to-Intermediate System (IS-IS) flood the SRLG membership information (including other TE link attributes such as bandwidth availability and affinity) so that all routers in the network have the SRLG information for each link. With this topology information, routers can compute backup tunnel paths that exclude links having SRLGs in common with their protected interfaces. As shown in Figure 2, the backup tunnel avoids the link between R2 and R3, which shares an SRLG with the protected interface.]")

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 258 of 698

No.	'821 Patent Claim 1	Zamfir '948
		 Fig. 2 (annotation added) Cisco IOS Guide at Page No. 587 ("[Step 4 show mpls traffic-eng topology Use this command to show the SRLG link membership flooded via the Interior Gateway Protocol (IGP).]") Cisco IOS Guide at Page 632 ("[IS-IS—Intermediate System-to-Intermediate System. Link-state hierarchical routing protocol that calls for intermediate system (IS) routers to exchange routing information based on a single metric to determine network topology.]") Cisco IOS Guide at Page No. 633 ("[OSPF—Open Shortest Path First. A link-state hierarchical Interior Gateway Protocol routing algorithm, derived from the IS-IS protocol. OSPF features include least-cost routing, multipath routing, and load balancing.]") Cisco IOS Guide at Page No. 581 ("[Autotunnel Backup for MPLS TE SRLGs Autotunnel backup is the ability of routers to create backup tunnels automatically. Therefore, you do not need to preconfigure each backup tunnel and then assign the backup tunnel to the protected interface. Only automatically created backup tunnels can avoid SRLGs or their protected interfaces.]") Cisco IOS Guide at Page No. 583 ("[Configuring the Routers That Automatically Create Backup Tunnels to Avoid MPLS TE SRLGs of Their Protected Interfaces

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 259 of 698

No.	'821 Patent Claim 1	Zamfir '948
		Perform the following task to configure routers that automatically create backup tunnels to avoid MPLS TE SRLGs of their protected interfaces. Backup tunnels provide link protection by rerouting traffic to the next hop bypassing failed links or in this instance by avoiding SRLGs.]") Cisco IOS Guide at Page No. 644 ("[Backup Tunnel Selection Procedure When an LSP is signaled, each node along the LSP path that provides FRR protection for the LSP selects a backup tunnel for the LSP to use if either of the following events occurs: • The link to the next hop fails. • The next hop fails. By having the node select the backup tunnel for an LSP before a failure occurs, the LSP can be rerouted onto the backup tunnel quickly if there is a failure.]") Cisco IOS Guide at Page 646 ("[NHOP Versus NNHOP Backup Tunnels More than one backup tunnel can protect a given LSP, where one backup tunnel terminates at the LSP's NNHOP, and the other terminates at the LSP's NHOP. In this case, the router chooses the backup tunnel that terminates at the NNHOP (that is, FRR prefers NNHOP over NHOP backup tunnels).Table 1 lists the tunnel selection priorities. The first choice is an NNHOP backup tunnel that acquires its bandwidth from a subpool or global pool, and has limited bandwidth. If there is no such backup tunnel, the next choice (2) is a next-next hop backup tunnel that acquires a limited amount of bandwidth from any pool. The preferences go from 1 (best) to 8 (worst), where choice 3 is for an NNHOP backup tunnel with an unlimited amount of subpool or global-pool bandwidth.]")

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 260 of 698

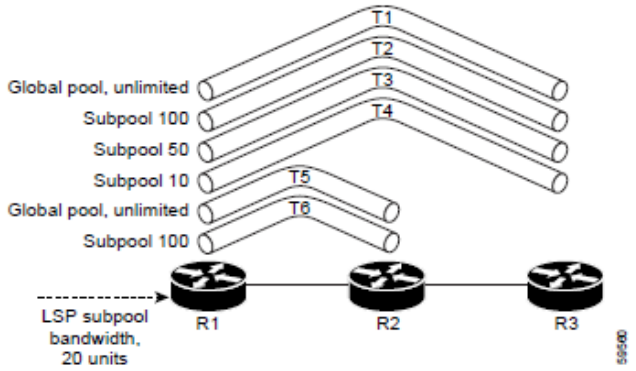
No.	'821 Patent Claim 1	Zamfir '948

Table 1 Tunnel Selection Priorities

Preference	Backup Tunnel Destination	Bandwidth Pool	Bandwidth Amount
1 (Best)	NNHOP	Subpool or global pool	Limited
2	NNHOP	Any	Limited
3	NNHOP	Subpool or global pool	Unlimited
4	NNHOP	Any	Unlimited
5	NHOP	Subpool or global pool	Limited
6	NHOP	Any	Limited
7	NHOP	Subpool or global pool	Unlimited
8 (Worst)	NHOP	Any	Unlimited

Cisco IOS Guide at Page No. 647-648 ("[Figure 7 shows an example of the backup tunnel selection procedure based on the designated amount of global pool and subpool bandwidth currently available.

Figure 7 Choosing from Among Multiple Backup Tunnels



In this example, an LSP requires 20 units (kilobits per second) of sub-pool backup bandwidth. The best backup tunnel is selected as follows:
1. Backup tunnels T1 through T4 are considered first because they terminate at the NNHOP.
2. Tunnel T4 is eliminated because it has only ten units of sub-pool backup bandwidth.
3. Tunnel T1 is eliminated because it protects only LSPs using global-pool bandwidth.

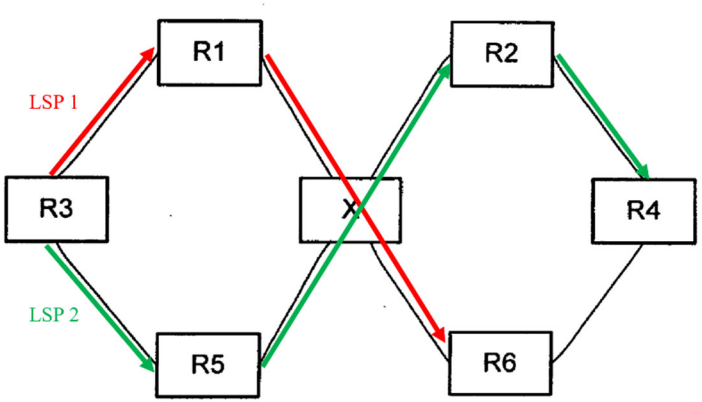
Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 261 of 698

No.	'821 Patent Claim 1	Zamfir '948
		4. Tunnel T3 is chosen over T2 because, although both have sufficient backup bandwidth, T3 has the least backup bandwidth available (leaving the most backup bandwidth available on T2). 5. Tunnels T5 and T6 need not be considered because they terminate at an NHOP, and therefore are less desirable than T3, which terminates at an NNHOP.]")
1[c]	selecting an entity pair from said plurality of transport entities based at least in part upon said overall cost; and	Zamfir '948 discloses selecting an entity pair from said plurality of transport entities based at least in part upon said overall cost. For example, Zamfir '948 discloses that LSPs in a bundle are established (e.g., G-LSP and TE-LSP) between the same source and destination nodes based on configuration or enforced SLA and have a certain degree of diversity e.g., SLRG diversity and optimality, and each LSP have an associated cost. Zamfir '948 further discloses that the cost information is used in order to select the LSPs within a bundle. *See Supra at* 1[b].
1[d]	if an entity pair reselection event occurs, reselecting said entity pair from the group consisting of said entity pair and a replacement entity pair comprising at least one entity distinct from the entities of said entity pair,	Zamfir '948 discloses if an entity pair reselection event occurs, reselecting said entity pair from the group consisting of said entity pair and a replacement entity pair comprising at least one entity distinct from the entities of said entity pair. For example, Zamfir '948 discloses that when a re-optimization operation is triggered (for example when there is a change in topology or when an LSP needs to be re-routed due to a network device failure), new LSPs are computed and an attempt is made to maintain the degree of diversity/optimality at the set level. Zamfir '948 further discloses that at least one LSP within the bundle will be different from the past LSP grouping. Zamfir '948 at Page No. 09, Para No. 0027 ("[A number of LSPs are established during GMPLS tunnel setup. The traffic may be arranged or duplicated any number of ways, as is well known by those skilled in the art. The paths are grouped into one or more bundles. With link bundling of the GMPLS interface, the component membership to the bundle is decided upon by the GMPLS control plane, which brings in the data plane aspect of dynamic

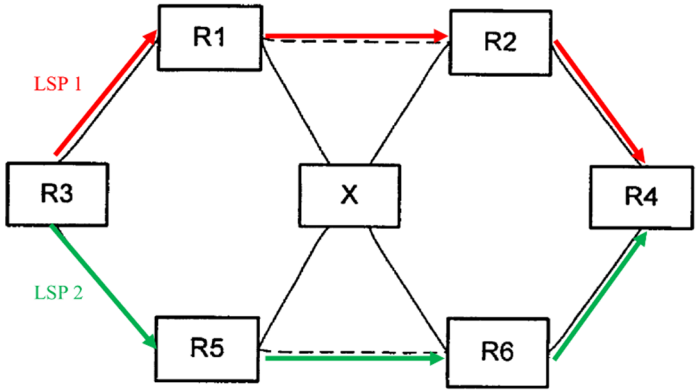
Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 262 of 698

No.	'821 Patent Claim 1	Zamfir '948
		membership. The LSPs in a bundle are established between the same source and destination nodes and, based on configuration or enforced SLA, they have a certain degree of diversity and optimality. These LSPs may be re-optimized or re-routed due to topology changes or failures in the network. As different network events change the egress physical interfaces of the LSPs, the membership in the dynamic bundle changes on-the-fly.]") Zamfir '948 at Page No. 09, Para No. 0028 ("[Another aspect of dynamic bundles is the handling of LSP bundles in the control plane. When an LSP set is created, if the routes are dynamically computed, a certain degree of diversity/optimality may be requested and realized. When a re-optimization operation is triggered (for example in the case where there is a change in topology or when an LSP needs to be re-routed due to a failure), an attempt is preferably made to maintain the degree of diversity/optimality at the set level. To realize this, path computation is used to determine a new LSP set. The LSPs in the new set can share resources with any subset of the old paths. This procedure is referred to as "make-a-set-before-break-a-set" and is described below.]") Zamfir '948 at Page No. 10, Para No. 0030 ("[When a new optical tunnel is created and added to the bundle, the physical interfaces associated with the optical tunnel interface is identified by a protection module and added to the bundle. If a failure occurs on an optical LSP which is also a component in the L2 bundle, the GMPLS control plane restores the LSP using, for example, a new physical interface and the protection module performs the translation to the new interface that will replace the old one. In other words, as GMPLS LSPs can change paths on-the-fly (e.g., due to rerouting in the event of a failure, re-optimization, etc.), TE control can change membership of the bundled interface in the event of a re-route of the GMPLS LSP.]") Zamfir '948 at Page No. 10, Para No. 0031 ("[In order to better illustrate the make-a-set-before-break-a-set process, the following example is provided. In this example, a dynamic bundle with two components LSP1 and LSP2 is created, with the two LSPs being setup with link and node diversity (FIG. 3): LSP1−R3, R1, R6, R∝(cost=30) LSP2−R3, R5, R2, R4 (cost=30)]")

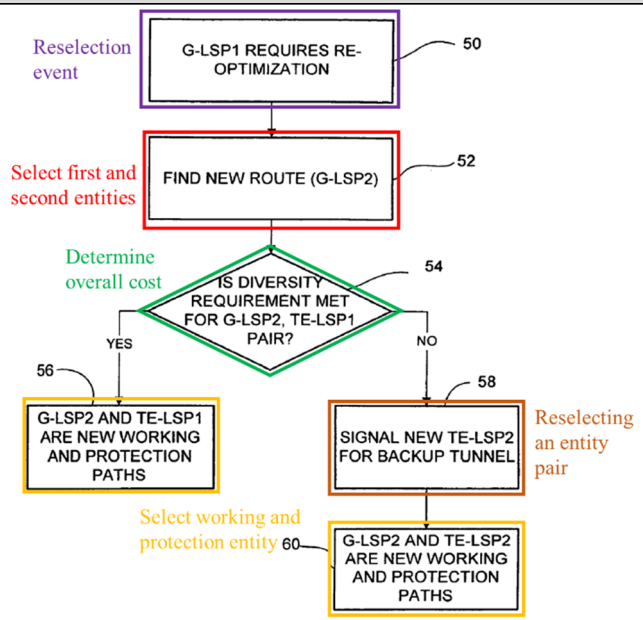
Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 263 of 698

No.	'821 Patent Claim 1	Zamfir '948
		 Fig. 3 (annotation added) Zamfir '948 at Page No. 10, Para No. 0032 ("[As shown in FIG. 4, two new links (shown with dashed lines) with cost equal to 5 are added to connect R1 with R2 and R5 with R6. If re-optimized as a set the new LSPs will be: LSP1−R3, R1, R2, R4 (cost 25) LSP2−R3, R5, R6, R4 (cost 25)]")

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 264 of 698

No.	'821 Patent Claim 1	Zamfir '948
		 Fig. 4 (annotation added) Zamfir '948 at Page No. 10, Para No. 0037 ("[FIG. 5 is a flowchart illustrating one example of a process of the present invention for dynamic FRR. At step 50, G-LSP1 needs to be re-optimized. A new (better) route is determined and G-LSP2 is ready to be signaled using a make-before-break procedure (step 52). Since G-LSP1 is protected by TE-LSP1, a determination is made if the (G-LSP2, TE-LSP1) pair still satisfies the diversity requirement (step 54). If this requirement is met, G-LSP2 and TE-LSP1 are the new working and protecting paths, respectively (step 56). If this requirement is not met, a new TE-LSP2 is signaled for the backup tunnel using a make-before-break procedure (steps 58 and 60). In order to minimize the window where a failure will have an impact on the data traffic, these two procedures are synchronized. The diversity requirement may be specified, for example, with respect to link diversity, node diversity, SLRG diversity, or any combination thereof.]")

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 265 of 698

No.	’821 Patent Claim 1	Zamfir ’948
		 Fig. 5 (annotation added) Zamfir ’948 at Claim 23 (“[23. A method for providing dynamic path protection in an optical network in communication with an IP network, the method comprising: establishing a set of working and protecting paths, each of said working paths having at least one of said protecting path designated as its backup path; following a failure or during optimization, selecting and establishing at least one new working path to replace one of said working paths; determining if said new working path and said designated backup path provide a predefined degree of diversity; and selecting a new protecting path if said predefined degree of diversity is not met.]”)

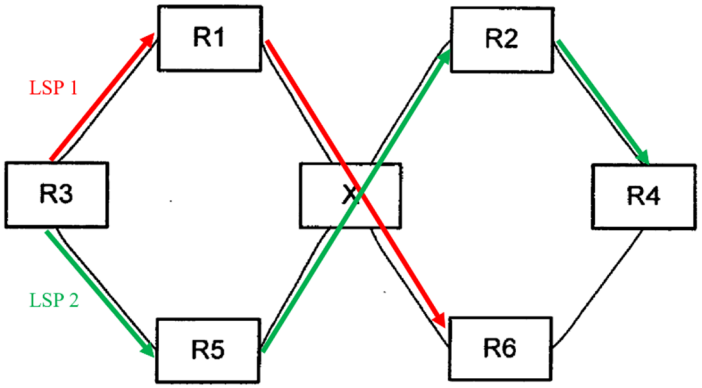
Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 266 of 698

No.	'821 Patent Claim 1	Zamfir '948
		Zamfir '948 at Claim 24 ("[24. The method of claim 23 wherein selecting a new working path comprises utilizing fast reroute.]")
1[e]	wherein said entity pair reselection event is selected from a group consisting of adding an entity to said plurality of transport entities, removing an entity from said plurality of transport entities, an operational status change for one of said plurality of transport entities, and a change in overall cost for one of said plurality of transport entities.	Zamfir '948 discloses wherein said entity pair reselection event is selected from a group consisting of adding an entity to said plurality of transport entities, removing an entity from said plurality of transport entities, an operational status change for one of said plurality of transport entities. For example, Zamfir '948 discloses that a re-optimization procedure is triggered where there is a change in topology or when an LSP needs to be re-routed due to a failure of a network device, or new links are added to the network. Thus, at least under the apparent claim scope alleged by Orckit's Infringement Disclosures, Zamfir '948 meets this limitation. To the extent that Zamfir '948 is found to not meet this limitation, wherein said entity pair reselection event is selected from a group consisting of adding an entity to said plurality of transport entities, removing an entity from said plurality of transport entities, an operational status change for one of said plurality of transport entities would have been obvious to a person having ordinary skill in the art, as explained below. Zamfir '948 at Page No. 09, Para No. 0027 ('[A number of LSPs are established during GMPLS tunnel setup. The traffic may be arranged or duplicated any number of ways, as is well known by those skilled in the art. The paths are grouped into one or more bundles. With link bundling of the GMPLS interface, the component membership to the bundle is decided upon by the GMPLS control plane, which brings in the data plane aspect of dynamic membership. The LSPs in a bundle are established between the same source and destination nodes and, based on configuration or enforced SLA, they have a certain degree of diversity and optimality. These LSPs may be re-optimized or re-routed due to topology changes or failures in the network. As different network events change the egress physical interfaces of the LSPs, the membership in the dynamic bundle changes on-the-fly.]") Zamfir '948 at Page No. 09, Para No. 0028 ("[Another aspect of dynamic bundles is the handling of LSP bundles in the control plane. When an LSP set is created, if the routes are dynamically computed, a certain degree of diversity/optimality may be requested and realized. When a re-optimization operation is triggered (for example in the case where there

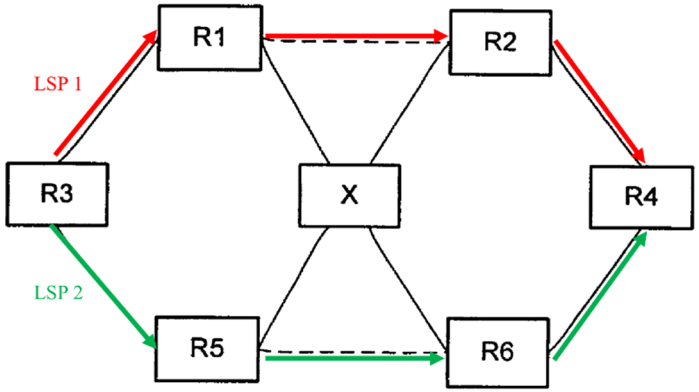
Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 267 of 698

No.	'821 Patent Claim 1	Zamfir '948
		is a change in topology or when an LSP needs to be re-routed due to a failure), an attempt is preferably made to maintain the degree of diversity/optimality at the set level. To realize this, path computation is used to determine a new LSP set. The LSPs in the new set can share resources with any subset of the old paths. This procedure is referred to as "make-a-set-before-break-a-set" and is described below.]") Zamfir '948 at Page No. 10, Para No. 0030 ("[When a new optical tunnel is created and added to the bundle, the physical interfaces associated with the optical tunnel interface is identified by a protection module and added to the bundle. If a failure occurs on an optical LSP which is also a component in the L2 bundle, the GMPLS control plane restores the LSP using, for example, a new physical interface and the protection module performs the translation to the new interface that will replace the old one. In other words, as GMPLS LSPs can change paths on-the-fly (e.g., due to rerouting in the event of a failure, re-optimization, etc.), TE control can change membership of the bundled interface in the event of a re-route of the GMPLS LSP.]") Zamfir '948 at Page No. 10, Para No. 0031 ("[In order to better illustrate the make-a-set-before-break-a-set process, the following example is provided. In this example, a dynamic bundle with two components LSP1 and LSP2 is created, with the two LSPs being setup with link and node diversity (FIG. 3): LSP1−R3, R1, R6, R∝(cost=30) LSP2−R3, R5, R2, R4 (cost=30)]")

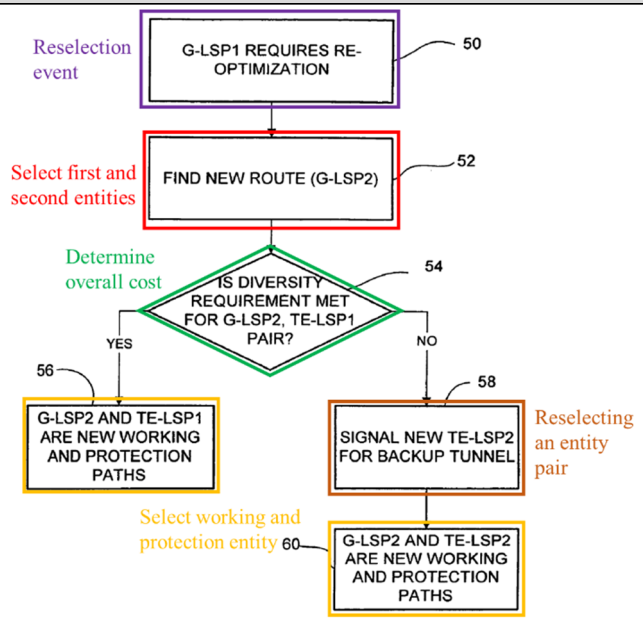
Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 268 of 698

No.	’821 Patent Claim 1	Zamfir ’948
		 Fig. 3 (annotation added) Zamfir ’948 at Page No. 10, Para No. 0032 (“[As shown in FIG. 4, two new links (shown with dashed lines) with cost equal to 5 are added to connect R1 with R2 and R5 with R6. If re-optimized as a set the new LSPs will be: LSP1−R3, R1, R2, R4 (cost 25) LSP2−R3, R5, R6, R4 (cost 25)]”)

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 269 of 698

No.	'821 Patent Claim 1	Zamfir '948
		 Fig. 4 (annotation added) Zamfir '948 at Page No. 10, Para No. 0037 ("[FIG. 5 is a flowchart illustrating one example of a process of the present invention for dynamic FRR. At step 50, G-LSP1 needs to be re-optimized. A new (better) route is determined and G-LSP2 is ready to be signaled using a make-before-break procedure (step 52). Since G-LSP1 is protected by TE-LSP1, a determination is made if the (G-LSP2, TE-LSP1) pair still satisfies the diversity requirement (step 54). If this requirement is met, G-LSP2 and TE-LSP1 are the new working and protecting paths, respectively (step 56). If this requirement is not met, a new TE-LSP2 is signaled for the backup tunnel using a make-before-break procedure (steps 58 and 60). In order to minimize the window where a failure will have an impact on the data traffic, these two procedures are synchronized. The diversity requirement may be specified, for example, with respect to link diversity, node diversity, SLRG diversity, or any combination thereof.]")

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 270 of 698

No.	'821 Patent Claim 1	Zamfir '948
		 Fig. 5 (annotation added). In addition, at least under the apparent claim scope alleged by Orckit's Infringement Disclosures, Zamfir '948 in combination with (1) the knowledge of a person of ordinary skill in the art, alone or in further combination with (2) each (individually, as well as one or more together) of the references identified in element 1[e] of Exhibit E-2 renders the claim, including the present limitation, obvious. Below are examples of two such references. For example, Juniper IOS Guide discloses "Fate Sharing" where a Juniper router computes a primary and secondary path using CSPF, CSPF identifies all fate-sharing groups that are associated with the primary path by identifying all links that the primary path transverses and checks each link in the traffic engineering database against the compiled group list. Juniper IOS Guide further discloses that if the link is a member of a group, the cost of the link is increased by the cost of the group and higher the cost value. At least based upon Orckit's

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 271 of 698

No.	'821 Patent Claim 1	Zamfir '948
		apparent understanding of the claim scope as reflected in its Infringement Disclosures, this limitation is met. Juniper IOS Guide at Page No. 73 ("[Fate Sharing Fate sharing allows you to create a database of information that CSPF uses to compute one or more backup paths to use in case the primary path becomes unstable. The database describes the relationships between elements of the network, such as routers and links. You can specify one or more elements within a group. Through fate sharing, you can configure backup paths that minimize the number of shared links and fiber paths with the primary paths as much as possible, to ensure that if a fiber is cut, the minimum amount of data is lost and a path still exists to the destination. For a backup path to work optimally, it must not share links or physical fiber paths with the primary path, ensuring that a single point of failure will not affect the primary and backup paths simultaneously. For more information about fate sharing, see the Junos OS Routing Protocols Configuration Guide.]") Juniper IOS Guide at Page No. 94 ("[Configuring the Ingress Router for MPLS-Signaled LSPs MPLS-signaled label-switched paths (LSPs) run from a specific ingress router to a specific egress router. For basic MPLS-signaled LSP function, you must configure the ingress router, but do not have to configure any other routers. To configure signaled LSPs, perform the following tasks on the ingress router: • Creating Named Paths on page 56 • Configuring Alternate Backup Paths Using Fate Sharing on page 58.]") Juniper IOS Guide at Page No. 96 ("[Configuring Alternate Backup Paths Using Fate Sharing You can create a database of information that Constrained Shortest Path First (CSPF) uses to compute one or more backup paths in case the primary path becomes unstable. The database describes the relationships between elements of the network, such as routers and links. Because these network elements share the same fate, this relationship is called fate sharing.

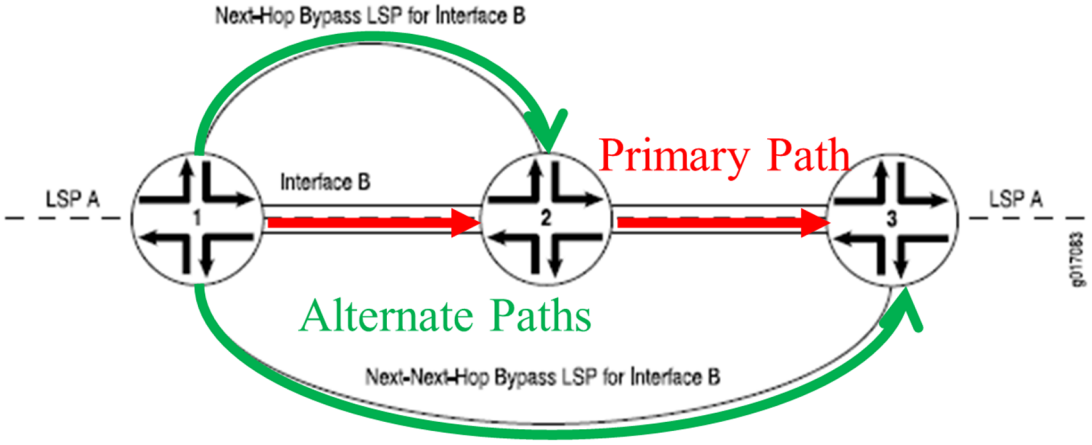
Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 272 of 698

No.	'821 Patent Claim 1	Zamfir '948
		You can configure backup paths that minimize the number of shared links and fiber paths with the primary paths as much as possible to ensure that, if a fiber is cut, the minimum amount of data is lost and a path still exists to the destination. For a backup path to work optimally, it must not share links or physical fiber paths with the primary path. This ensures that a single point of failure will not affect the primary and backup paths at the same time.]") Juniper IOS Guide at Page No. 97 ("[Fate-sharing groups contain three types of objects: • Point-to-point links—Identified by the IP addresses at each end of the link. Unnumbered point-to-point links are typically identified by borrowing IP addresses from other interfaces. Order is not important; from 1.2.3.4 to 1.2.3.5 and from 1.2.3.5 to 1.2.3.4 have the same meaning. • Non-point-to-point links—Include links on a LAN interface (such as Gigabit Ethernet interfaces) or nonbroadcast multiaccess (NBMA) interfaces (such as Asynchronous Transfer Mode [ATM] or Frame Relay). You identify these links by their individual interface address. For example, if the LAN interface 192.168.200.0/24 has four routers attached to it, each router link is individually identified: from 192.168.200.1; # LAN interface of router 1 from 192.168.200.2; # LAN interface of router 2 from 192.168.200.3; # LAN interface of router 3 from 192.168.200.4; # LAN interface of router 4 You can list the addresses in any order. • A router node—Identified by its configured router ID. An object can be in any number of groups, and a group can contain any number of objects. Each group has a configurable cost attributed to it, which represents the level of impact this group has on CSPF computations. The higher the cost, the less likely a backup path will share with the primary path any objects in the group. The cost is directly comparable to traffic engineering metrics. By default, the cost is 1. Changing the fate-sharing database does not affect established LSPs until the next reoptimization of CSPF. The fate-sharing database does influence fast-reroute computations.]")

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 273 of 698

No.	'821 Patent Claim 1	Zamfir '948
		Juniper IOS Guide at Page No. 97 ("[When CSPF computes the primary paths of an LSP (or secondary paths when the primary path is not active), it ignores the fate-sharing information. You always want to find the best possible path (least IGP cost) for the primary path. When CSPF computes a secondary path while the primary path (of the same LSP) is active, the following occurs: 1. CSPF identifies all fate-sharing groups that are associated with the primary path. CSPF does this by identifying all links and nodes that the primary path traverses and compiling group lists that contain at least one of the links or nodes. CSPF ignores the ingress and egress nodes in the search. 2. CSPF checks each link in the traffic engineering database against the compiled group list. If the link is a member of a group, the cost of the link is increased by the cost of the group. If a link is a member of multiple groups, all group costs are added together. 3. CSPF performs the check for every node in the traffic engineering database, except the ingress and egress node. Again, a node can belong to multiple groups, so costs are additive. 4. The router performs regular CSPF computation with the adjusted topology.]") Juniper IOS Guide at Page No. 330 ("In Figure 23 on page 292, both node protection and link protection are enabled on Interface B on Router 1. Both node protection and link protection are also enabled on LSP A, an LSP that traverses the link transiting Router 1, Router 2, and Router 3. If the link between Router 1 and Router 2 fails, traffic from LSP A is quickly switched to the next-hop bypass LSP generated by link protection. If Router 2 suffers a hardware or software failure, traffic from LSP A is switched to the next-next-hop bypass LSP generated by node protection.")

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 274 of 698

No.	'821 Patent Claim 1	Zamfir '948
		Figure 23: Node Protection Creating a Next-Next-Hop Bypass LSP  Fig. 23 (annotation added). For example, Cisco IOS Guide discloses that the path an LSP uses is determined by the LSP resource requirements and network resources, such as bandwidth. Cisco IOS Guide further discloses the MPLS TE SRLG feature that enhances backup tunnel path selection so a backup tunnel can avoid using links that are in the same SRLG as the interfaces it is protecting. Thus, at least under the apparent claim scope alleged by Orckit's Infringement Disclosures, this limitation is met. Cisco IOS Guide at Page No. 253 ("[Path Option Selection for MPLS TE Tunnel LSPs This section contains the following topics about path option selection for MPLS TE Tunnel LSPs: • Constraint-Based Routing and Path Option Selection, page 5 • Tunnel Reoptimization and Path Option Selection, page 5 • Path Option Selection with Bandwidth Override, page 6 Constraint-Based Routing and Path Option Selection MPLS traffic engineering automatically establishes and maintains LSPs across the backbone by using the Resource Reservation Protocol (RSVP). The path that an LSP uses is determined

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 275 of 698

No.	'821 Patent Claim 1	Zamfir '948
		by the LSP resource requirements and network resources, such as bandwidth. Traffic engineering tunnels are calculated at the LSP head based on a fit between required and available resources (constraint-based routing). Without the Path Option for Bandwidth Override feature, a TE tunnel establishes an LSP based on dynamic or explicit path options in order of preference. However, the bandwidth and other attributes configured on the TE tunnel allow the setup of an LSP only if LSP path options satisfy the constraints. If a path cannot be found that satisfies the configured path options, then the tunnel is not set up.]") Cisco IOS Guide at Page No. 639 ("[Bandwidth Protection NHOP and NNHOP backup tunnels can be used to provide bandwidth protection for rerouted LSPs. This is referred to as backup bandwidth. You can associate backup bandwidth with NHOP or NNHOP backup tunnels. This informs the router of the amount of backup bandwidth a particular backup tunnel can protect. When a router maps LSPs to backup tunnels, bandwidth protection ensures that an LSP uses a given backup tunnel only if there is sufficient backup bandwidth. The router selects which LSPs use which backup tunnels in order to provide maximum bandwidth protection. That is, the router determines the best way to map LSPs onto backup tunnels in order to maximize the number of LSPs that can be protected. For information about mapping tunnels and assigning backup bandwidth, see the "Backup Tunnel Selection Procedure" section on page 10. LSPs that have the "bandwidth protection desired" bit set have a higher right to select backup tunnels that provide bandwidth protection; that is, those LSPs can preempt other LSPs that do not have that bit set. For more information, see the "Prioritizing Which LSPs Obtain Backup Tunnels with Bandwidth Protection" section on page 8.]") Cisco IOS Guide at Page No. 640 ("[Backup Tunnel Support Backup tunnel support has the following capabilities: • Backup Tunnels Can Terminate at the Next-Next Hop to Support FRR, page 6 • Multiple Backup Tunnels Can Protect the Same Interface, page 6 • Backup Tunnels Provide Scalability, page 7.]") Cisco IOS Guide at Page 640 ("[Multiple Backup Tunnels Can Protect the Same Interface

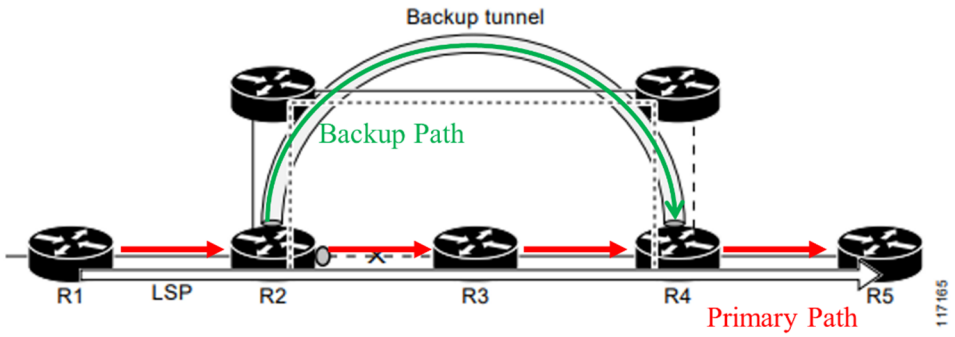
Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 276 of 698

No.	'821 Patent Claim 1	Zamfir '948
		There is no limit (except memory limitations) to the number of backup tunnels that can protect a given interface. In many topologies, support for node protection requires supporting multiple backup tunnels per protected interface. These backup tunnels can terminate at the same destination or at different destinations. That is, for a given protected interface, you can configure multiple NHOP or NNHOP backup tunnels. This allows redundancy and load balancing.]") Cisco IOS Guide at Page No. 603 ("[Traffic Engineering Tunnels MPLS TE lets you build label switched paths (LSPs) across your network for forwarding traffic. MPLS TE LSPs let the headend of a TE tunnel control the path its traffic takes to a particular destination. This method is more flexible than forwarding traffic based only on a destination address. Interarea tunnels allow you to do the following: • Build TE tunnels between areas (interarea tunnels) • Build TE tunnels that start and end in the same area, on multiple areas on a router (intra-area tunnels) Some tunnels are more important than others. For example, you may have tunnels carrying VoIP traffic and tunnels carrying data traffic that are competing for the same resources. MPLS TE allows you to have some tunnels preempt others. Each tunnel has a priority, and more-important tunnels take precedence over less-important tunnels.]") Cisco IOS Guide at Page No. 641 ("[Backup Bandwidth Protection Backup bandwidth protection allows you to give LSPs carrying certain kinds of data (such as voice) priority for using backup tunnels.]") Cisco IOS Guide at Page No. 650 ("[For each protected link AB with a maximum reservable subpool value of n, there may be a path from node A to node B such that the difference between the maximum reservable global and the maximum reservable subpool is at least the value of n. If it is possible to find such paths for each link in the network, you can establish all the backup tunnels along such paths without any bandwidth reservations. If there is a single link failure, only one backup tunnel will use any link on its path. Because that path has at least

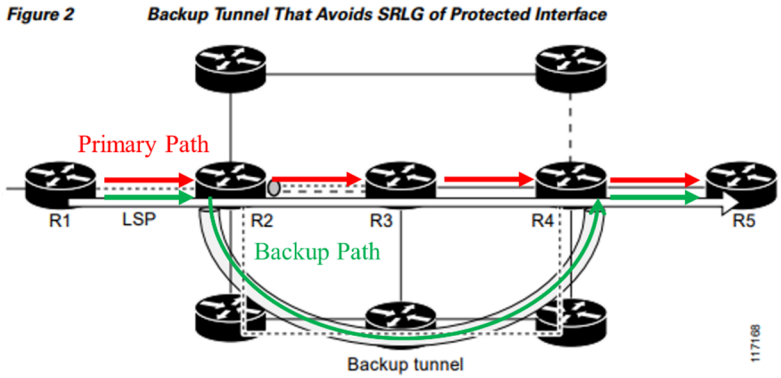
Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 277 of 698

No.	'821 Patent Claim 1	Zamfir '948
		n available bandwidth (in the global pool), assuming that marking and scheduling is configured to classify the subpool traffic into a priority queue, the subpool bandwidth is guaranteed. This approach allows sharing of the global pool bandwidth between backup tunnels protecting independent link failures. The backup tunnels are expected to be used for only a short period of time after a failure (until the headends of affected LSPs reroute those LSPs to other paths with available subpool bandwidth). The probability of multiple unrelated link failures is very small (in the absence of node or shared risk link group (SRLG) failures, which result in multiple link failures). Therefore, it is reasonable to assume that link failures are in practice independent with high probability. This "independent failure assumption" in combination with backup tunnels signaled without explicit bandwidth reservation enables efficient bandwidth sharing that yields substantial bandwidth savings.]") Cisco IOS Guide at Page No. 651 ("[A similar approach can be used for node and SRLG protection. However, the decision of where to put the backup tunnels is more complicated because both node and SRLG failures effectively result in the simultaneous failure of several links. Therefore, the backup tunnels protecting traffic traversing all affected links cannot be computed independently of each other.]") Cisco IOS Guide at Page No. 579 ("[MPLS Traffic Engineering Shared Risk Link Groups SRLGs refer to situations where links in a network share a common fiber (or a common physical attribute). If one link fails, other links in the group may fail too. Links in the group have a shared risk. Backup tunnels should avoid using links in the same SRLG as interfaces they are protecting. Otherwise, when the protected link fails the backup tunnel fails too.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 278 of 698

No.	'821 Patent Claim 1	Zamfir '948
		Figure 1 Backup Tunnel in the Same SRLG as the Interface It Is Protecting  Fig. 1 (annotation added) Figure 1 shows a primary label-switched path (LSP) from router R1 to router R5. The LSP protects against the failure of the R2-R3 link at R2 via a backup tunnel to R4. If the R2-R3 link fails, link protection reroutes the LSP along the backup tunnel. However, the R2-R3 link and one of the backup tunnel links are in the same SRLG. So if the R2-R3 link fails, the backup tunnel may fail too. The MPLS TE SRLG feature enhances backup tunnel path selection so a backup tunnel can avoid using links that are in the same SRLG as the interfaces it is protecting. There are two ways for a backup tunnel to avoid the SRLGs of its protected interface: • The router does not create the backup tunnel unless it avoids SRLGs of the protected interface.]") Cisco IOS Guide at Page No. 580 ("[Open Shortest Path First (OSPF) and Intermediate System-to-Intermediate System (IS-IS) flood the SRLG membership information (including other TE link attributes such as bandwidth availability and affinity) so that all routers in the network have the SRLG information for each link. With this topology information, routers can compute backup tunnel paths that exclude links having SRLGs in common with their protected interfaces. As shown in Figure 2, the backup tunnel avoids the link between R2 and R3, which shares an SRLG with the protected interface.]")

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 279 of 698

No.	'821 Patent Claim 1	Zamfir '948
		 Fig. 2 (annotation added) Cisco IOS Guide at Page No. 587 ("[Step 4 show mpls traffic-eng topology Use this command to show the SRLG link membership flooded via the Interior Gateway Protocol (IGP).]") Cisco IOS Guide at Page 632 ("[IS-IS—Intermediate System-to-Intermediate System. Link-state hierarchical routing protocol that calls for intermediate system (IS) routers to exchange routing information based on a single metric to determine network topology.]") Cisco IOS Guide at Page No. 633 ("[OSPF—Open Shortest Path First. A link-state hierarchical Interior Gateway Protocol routing algorithm, derived from the IS-IS protocol. OSPF features include least-cost routing, multipath routing, and load balancing.]") Cisco IOS Guide at Page No. 581 ("[Autotunnel Backup for MPLS TE SRLGs Autotunnel backup is the ability of routers to create backup tunnels automatically. Therefore, you do not need to preconfigure each backup tunnel and then assign the backup tunnel to the protected interface. Only automatically created backup tunnels can avoid SRLGs or their protected interfaces.]") Cisco IOS Guide at Page No. 583 ("[Configuring the Routers That Automatically Create Backup Tunnels to Avoid MPLS TE SRLGs of Their Protected Interfaces

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 280 of 698

No.	'821 Patent Claim 1	Zamfir '948
		Perform the following task to configure routers that automatically create backup tunnels to avoid MPLS TE SRLGs of their protected interfaces. Backup tunnels provide link protection by rerouting traffic to the next hop bypassing failed links or in this instance by avoiding SRLGs.]") Cisco IOS Guide at Page No. 644 ("[Backup Tunnel Selection Procedure When an LSP is signaled, each node along the LSP path that provides FRR protection for the LSP selects a backup tunnel for the LSP to use if either of the following events occurs: • The link to the next hop fails. • The next hop fails. By having the node select the backup tunnel for an LSP before a failure occurs, the LSP can be rerouted onto the backup tunnel quickly if there is a failure.]") Cisco IOS Guide at Page 646 ("[NHOP Versus NNHOP Backup Tunnels More than one backup tunnel can protect a given LSP, where one backup tunnel terminates at the LSP's NNHOP, and the other terminates at the LSP's NHOP. In this case, the router chooses the backup tunnel that terminates at the NNHOP (that is, FRR prefers NNHOP over NHOP backup tunnels).Table 1 lists the tunnel selection priorities. The first choice is an NNHOP backup tunnel that acquires its bandwidth from a subpool or global pool, and has limited bandwidth. If there is no such backup tunnel, the next choice (2) is a next-next hop backup tunnel that acquires a limited amount of bandwidth from any pool. The preferences go from 1 (best) to 8 (worst), where choice 3 is for an NNHOP backup tunnel with an unlimited amount of subpool or global-pool bandwidth.]")

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 281 of 698

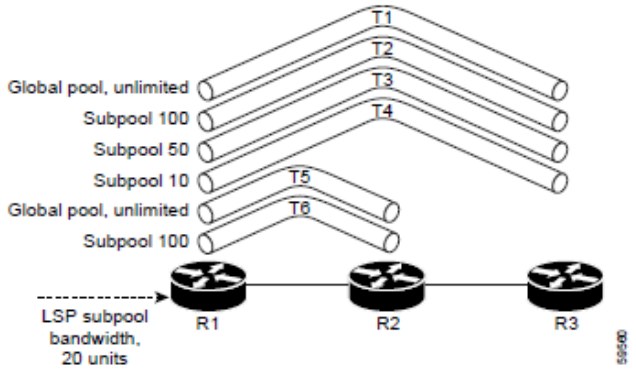
No.	'821 Patent Claim 1	Zamfir '948

Table 1 Tunnel Selection Priorities

Preference	Backup Tunnel Destination	Bandwidth Pool	Bandwidth Amount
1 (Best)	NNHOP	Subpool or global pool	Limited
2	NNHOP	Any	Limited
3	NNHOP	Subpool or global pool	Unlimited
4	NNHOP	Any	Unlimited
5	NHOP	Subpool or global pool	Limited
6	NHOP	Any	Limited
7	NHOP	Subpool or global pool	Unlimited
8 (Worst)	NHOP	Any	Unlimited

Cisco IOS Guide at Page No. 647-648 ("[Figure 7 shows an example of the backup tunnel selection procedure based on the designated amount of global pool and subpool bandwidth currently available.

Figure 7 Choosing from Among Multiple Backup Tunnels



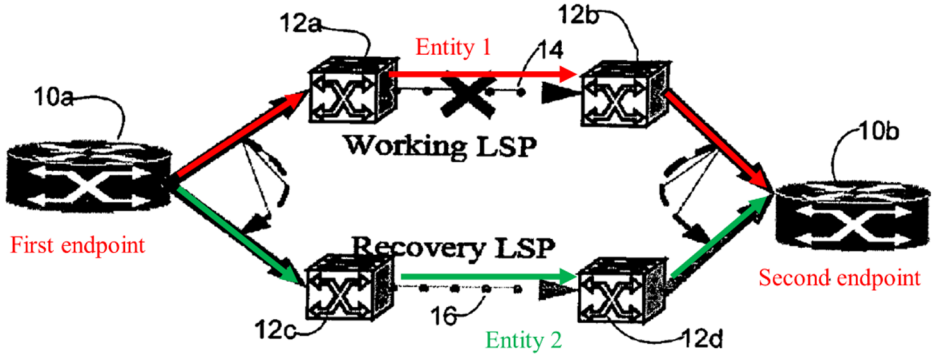
In this example, an LSP requires 20 units (kilobits per second) of sub-pool backup bandwidth. The best backup tunnel is selected as follows:
1. Backup tunnels T1 through T4 are considered first because they terminate at the NNHOP.
2. Tunnel T4 is eliminated because it has only ten units of sub-pool backup bandwidth.
3. Tunnel T1 is eliminated because it protects only LSPs using global-pool bandwidth.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 282 of 698

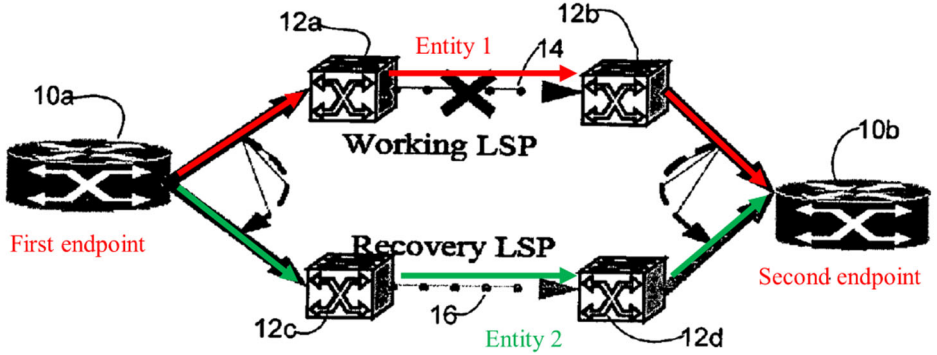
No.	'821 Patent Claim 1	Zamfir '948
		4. Tunnel T3 is chosen over T2 because, although both have sufficient backup bandwidth, T3 has the least backup bandwidth available (leaving the most backup bandwidth available on T2). 5. Tunnels T5 and T6 need not be considered because they terminate at an NHOP, and therefore are less desirable than T3, which terminates at an NNHOP.]")

No.	'821 Patent Claim 2	Zamfir '948
2[preamble]	The method of claim 1, wherein said step of selecting an entity pair further comprises:	Zamfir '948 discloses said step of selecting an entity pair further comprises. For example, Zamfir '948 discloses establishing LSPs in a bundle (e.g., G-LSP and TE-LSP) between the same source and destination nodes based on configuration or enforced SLA and have a certain degree of diversity e.g., SLRG diversity and optimality, and each LSP have an associated cost. *See Supra at* 1[c].
2[a]	selecting a working entity from said entity pair; and	Zamfir '948 discloses selecting a working entity from said entity pair. For example, Zamfir '948 discloses selecting LSP links for a bundle. Zamfir '948 further discloses that an LSP can be a working LSP and a recovery LSP. Zamfir '948 at Page No. 09, Para No. 0021 ("[FIG. 1A illustrates end-to-end protection with one working path 14 (between nodes 12 a and 12 b) and one protecting (recovery) path 16 (between nodes 12 c and 12 d). The working and protecting paths 14, 16 are preferably signaled as two LSPs that belong to the same tunnel. The working LSP 14 transports traffic between routers 10 a and 10 b and the protecting LSP 16 is used to transport traffic when the working LSP fails. End-to-end recovery provides recovery of the entire LSP from its source (ingress router end-point 10 a) to its destination (egress router end-point 10 b).]"]

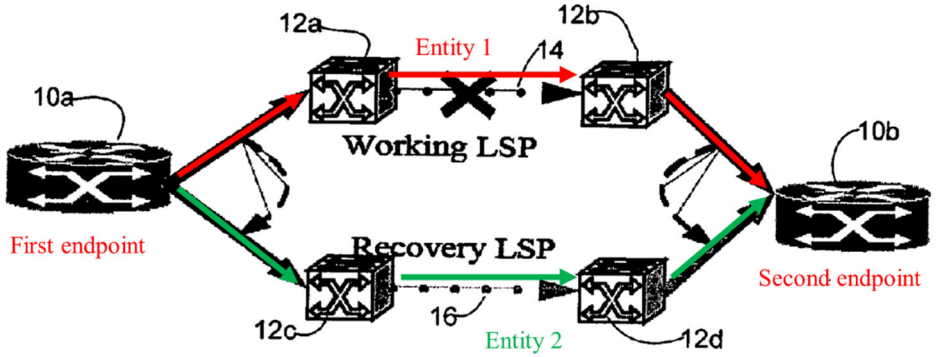
Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 283 of 698

No.	'821 Patent Claim 2	Zamfir '948
		 FIG. 1A Fig. 1A (annotation added) Zamfir '948 at Page No. 10, Para No. 0034 ("[Dynamic FRR extends link protection schemes, such as FRR (Fast Reroute), to handle GMPLS LSPs. A GMPLS LSP (G-LSP) is protected using a backup tunnel (TE-LSP) in a similar way as to how FRR is realized for physical links.]") Zamfir '948 at Page No. 10, Para No. 0035 ("[G-LSP and TE-LSP are preferably initially setup with a maximum diversity degree.]")
2[b]	selecting a protection entity from said entity pair.	Zamfir '948 discloses selecting a protection entity from said entity pair. For example, Zamfir '948 discloses selecting LSP links for a bundle. Zamfir '948 further discloses that an LSP can be a working LSP and a recovery LSP. Zamfir '948 at Page No. 09, Para No. 0021 ("[FIG. 1A illustrates end-to-end protection with one working path 14 (between nodes 12 a and 12 b) and one protecting (recovery) path 16 (between nodes 12 c and 12 d). The working and protecting paths 14, 16 are preferably signaled as two LSPs that belong to the same tunnel. The working LSP 14 transports traffic

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 284 of 698

No.	'821 Patent Claim 2	Zamfir '948
		between routers 10 a and 10 b and the protecting LSP 16 is used to transport traffic when the working LSP fails. End-to-end recovery provides recovery of the entire LSP from its source (ingress router end-point 10 a) to its destination (egress router end-point 10 b).]")  FIG. 1A Fig. 1A (annotation added) Zamfir '948 at Page No. 10, Para No. 0034 ('[Dynamic FRR extends link protection schemes, such as FRR (Fast Reroute), to handle GMPLS LSPs. A GMPLS LSP (G-LSP) is protected using a backup tunnel (TE-LSP) in a similar way as to how FRR is realized for physical links.]") Zamfir '948 at Page No. 10, Para No. 0035 ("[G-LSP and TE-LSP are preferably initially setup with a maximum diversity degree.]")

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 285 of 698

No.	'821 Patent Claim 3	Zamfir '948
3	The method of claim 2, further comprising the step of selecting an active entity from the set consisting of said working entity and said protection entity.	Zamfir '948 discloses the step of selecting an active entity from the set consisting of said working entity and said protection entity. For example, Zamfir '948 discloses that system selects the working LSP to transport data traffic between routers in normal operation, and the system selects the recovery LSP to transport data traffic when the working LSP fails. Zamfir '948 at Page No. 09, Para No. 0021 ("[FIG. 1A illustrates end-to-end protection with one working path 14 (between nodes 12 a and 12 b) and one protecting (recovery) path 16 (between nodes 12 c and 12 d). The working and protecting paths 14, 16 are preferably signaled as two LSPs that belong to the same tunnel. The working LSP 14 transports traffic between routers 10 a and 10 b and the protecting LSP 16 is used to transport traffic when the working LSP fails. End-to-end recovery provides recovery of the entire LSP from its source (ingress router end-point 10 a) to its destination (egress router end-point 10 b).]")  FIG. 1A Fig. 1A (annotation added)

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 286 of 698

No.	'821 Patent Claim 3	Zamfir '948
		Zamfir '948 at Page No. 10, Para No. 0034 ("[Dynamic FRR extends link protection schemes, such as FRR (Fast Reroute), to handle GMPLS LSPs. A GMPLS LSP (G-LSP) is protected using a backup tunnel (TE-LSP) in a similar way as to how FRR is realized for physical links.]")

No.	'821 Patent Claim 4	Zamfir '948
4	The method of claim 2, wherein selecting an entity pair further comprises minimizing an overall cost function.	Zamfir '948 discloses selecting an entity pair further comprises minimizing an overall cost function. For example, Zamfir '948 discloses establishing LSPs in a bundle (e.g., G-LSP and TE-LSP) between the same source and destination nodes based on configuration or enforced SLA and having a certain degree of diversity e.g., SLRG diversity and optimality. Thus, at least under the apparent claim scope alleged by Orckit's Infringement Disclosures, Zamfir '948 meets this limitation. To the extent that Zamfir '948 is found to not meet this limitation, wherein selecting an entity pair further comprises minimizing an overall cost function would have been obvious to a person having ordinary skill in the art, as explained below. Zamfir '948 at Page No. 09, Para No. 0021 ("[FIG. 1A illustrates an example of a network in which data transmission occurs on generalized label-switched paths (G-LSPs). The G-LSPs are automatically setup and torn down by means of a signaling protocol, as is well known by those skilled in the art. (In the following description, the term LSP is also used to refer to generalized LSPs.) FIG. 1A illustrates end-to-end protection with one working path 14 (between nodes 12 a and 12 b) and one protecting (recovery) path 16 (between nodes 12 c and 12 d). The working and protecting paths 14, 16 are preferably signaled as two LSPs that belong to the same tunnel. The working LSP 14 transports traffic between routers 10 a and 10 b and the protecting LSP 16 is used to transport traffic when the working LSP fails. End-to-end recovery provides recovery of the entire LSP from its source (ingress router end-point 10 a) to its destination (egress router end-point 10 b).]") Zamfir '948 at Page No. 09, Para No. 0022 ("[It is to be understood that the arrangement shown in FIG. 1A is only one example and the invention described herein may be applied to a network having any number of working or recovery LSPs.]")

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 287 of 698

No.	'821 Patent Claim 4	Zamfir '948
		Zamfir '948 at Page No. 09, Para No. 0027 ("[A number of LSPs are established during GMPLS tunnel setup. The traffic may be arranged or duplicated any number of ways, as is well known by those skilled in the art. The paths are grouped into one or more bundles. With link bundling of the GMPLS interface, the component membership to the bundle is decided upon by the GMPLS control plane, which brings in the data plane aspect of dynamic membership. The LSPs in a bundle are established between the same source and destination nodes and, based on configuration or enforced SLA, they have a certain degree of diversity and optimality.]") Zamfir '948 at Page No. 10, Para No. 0031 ("[In order to better illustrate the make-a-set-before-break-a-set process, the following example is provided. In this example, a dynamic bundle with two components LSP1 and LSP2 is created, with the two LSPs being setup with link and node diversity (FIG. 3): LSP1−R3, R1, R6, R∝(cost=30) LSP2−R3, R5, R2, R4 (cost=30) In this example, all links are assumed to have cost equal to 10, therefore LSP1 and LSP2 each have a cost equal to 30.]") Zamfir '948 at Page No. 10, Para No. 0034 ("[Dynamic FRR extends link protection schemes, such as FRR (Fast Reroute), to handle GMPLS LSPs. A GMPLS LSP (G-LSP) is protected using a backup tunnel (TE-LSP) in a similar way as to how FRR is realized for physical links.]") Zamfir '948 at Page No. 10, Para No. 0035 ("[G-LSP and TE-LSP are preferably initially setup with a maximum diversity degree.]") In addition, at least under the apparent claim scope alleged by Orckit's Infringement Disclosures, Zamfir '948 in combination with (1) the knowledge of a person of ordinary skill in

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 288 of 698

No.	'821 Patent Claim 4	Zamfir '948
		the art, alone or in further combination with (2) each (individually, as well as one or more together) of the references identified in element 4 of Exhibit E-2 renders the claim, including the present limitation, obvious. Below are examples of two such references. For example, Juniper IOS Guide discloses "Fate Sharing" where a Juniper router computes a primary and secondary path using CSPF, CSPF identifies all fate-sharing groups that are associated with the primary path by identifying all links that the primary path transverses and checks each link in the traffic engineering database against the compiled group list. Juniper IOS Guide further discloses that if the link is a member of a group, the cost of the link is increased by the cost of the group and higher the cost value. At least based upon Orckit's apparent understanding of the claim scope as reflected in its Infringement Disclosures, this limitation is met. Juniper IOS Guide at Page No. 73 ("[Fate Sharing Fate sharing allows you to create a database of information that CSPF uses to compute one or more backup paths to use in case the primary path becomes unstable. The database describes the relationships between elements of the network, such as routers and links. You can specify one or more elements within a group. Through fate sharing, you can configure backup paths that minimize the number of shared links and fiber paths with the primary paths as much as possible, to ensure that if a fiber is cut, the minimum amount of data is lost and a path still exists to the destination. For a backup path to work optimally, it must not share links or physical fiber paths with the primary path, ensuring that a single point of failure will not affect the primary and backup paths simultaneously. For more information about fate sharing, see the Junos OS Routing Protocols Configuration Guide.]") Juniper IOS Guide at Page No. 94 ("[Configuring the Ingress Router for MPLS-Signaled LSPs MPLS-signaled label-switched paths (LSPs) run from a specific ingress router to a specific egress router. For basic MPLS-signaled LSP function, you must configure the ingress router, but do not have to configure any other routers. To configure signaled LSPs, perform the following tasks on the ingress router:

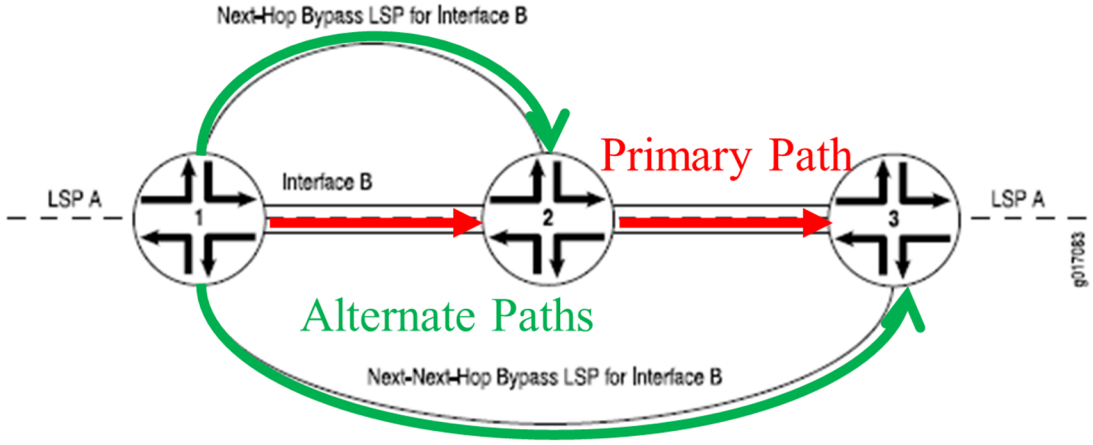
Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 289 of 698

No.	'821 Patent Claim 4	Zamfir '948
		• Creating Named Paths on page 56 • Configuring Alternate Backup Paths Using Fate Sharing on page 58.]") Juniper IOS Guide at Page No. 96 ("[Configuring Alternate Backup Paths Using Fate Sharing You can create a database of information that Constrained Shortest Path First (CSPF) uses to compute one or more backup paths in case the primary path becomes unstable. The database describes the relationships between elements of the network, such as routers and links. Because these network elements share the same fate, this relationship is called fate sharing. You can configure backup paths that minimize the number of shared links and fiber paths with the primary paths as much as possible to ensure that, if a fiber is cut, the minimum amount of data is lost and a path still exists to the destination. For a backup path to work optimally, it must not share links or physical fiber paths with the primary path. This ensures that a single point of failure will not affect the primary and backup paths at the same time.]") Juniper IOS Guide at Page No. 97 ("[Fate-sharing groups contain three types of objects: • Point-to-point links—Identified by the IP addresses at each end of the link. Unnumbered point-to-point links are typically identified by borrowing IP addresses from other interfaces. Order is not important; from 1.2.3.4 to 1.2.3.5 and from 1.2.3.5 to 1.2.3.4 have the same meaning. • Non-point-to-point links—Include links on a LAN interface (such as Gigabit Ethernet interfaces) or nonbroadcast multiaccess (NBMA) interfaces (such as Asynchronous Transfer Mode [ATM] or Frame Relay). You identify these links by their individual interface address. For example, if the LAN interface 192.168.200.0/24 has four routers attached to it, each router link is individually identified: from 192.168.200.1; # LAN interface of router 1 from 192.168.200.2; # LAN interface of router 2 from 192.168.200.3; # LAN interface of router 3 from 192.168.200.4; # LAN interface of router 4 You can list the addresses in any order. • A router node—Identified by its configured router ID.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 290 of 698

No.	'821 Patent Claim 4	Zamfir '948
		An object can be in any number of groups, and a group can contain any number of objects. Each group has a configurable cost attributed to it, which represents the level of impact this group has on CSPF computations. The higher the cost, the less likely a backup path will share with the primary path any objects in the group. The cost is directly comparable to traffic engineering metrics. By default, the cost is 1. Changing the fate-sharing database does not affect established LSPs until the next reoptimization of CSPF. The fate-sharing database does influence fast-reroute computations.]") Juniper IOS Guide at Page No. 97 ("[When CSPF computes the primary paths of an LSP (or secondary paths when the primary path is not active), it ignores the fate-sharing information. You always want to find the best possible path (least IGP cost) for the primary path. When CSPF computes a secondary path while the primary path (of the same LSP) is active, the following occurs: 1. CSPF identifies all fate-sharing groups that are associated with the primary path. CSPF does this by identifying all links and nodes that the primary path traverses and compiling group lists that contain at least one of the links or nodes. CSPF ignores the ingress and egress nodes in the search. 2. CSPF checks each link in the traffic engineering database against the compiled group list. If the link is a member of a group, the cost of the link is increased by the cost of the group. If a link is a member of multiple groups, all group costs are added together. 3. CSPF performs the check for every node in the traffic engineering database, except the ingress and egress node. Again, a node can belong to multiple groups, so costs are additive. 4. The router performs regular CSPF computation with the adjusted topology.]") Juniper IOS Guide at Page No. 330 ("In Figure 23 on page 292, both node protection and link protection are enabled on Interface B on Router 1. Both node protection and link protection are also enabled on LSP A, an LSP that traverses the link transiting Router 1, Router 2, and Router 3. If the link between Router 1 and Router 2 fails, traffic from LSP A is quickly switched to the next-hop bypass LSP generated by link protection. If Router 2 suffers a hardware or software failure, traffic from LSP A is switched to the next-next-hop bypass LSP generated by node protection.")

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 291 of 698

No.	'821 Patent Claim 4	Zamfir '948
		Figure 23: Node Protection Creating a Next-Next-Hop Bypass LSP  Fig. 23 (annotation added). For example, Cisco IOS Guide discloses that the path an LSP uses is determined by the LSP resource requirements and network resources, such as bandwidth. Cisco IOS Guide further discloses the MPLS TE SRLG feature that enhances backup tunnel path selection so a backup tunnel can avoid using links that are in the same SRLG as the interfaces it is protecting. Thus, at least under the apparent claim scope alleged by Orckit's Infringement Disclosures, this limitation is met. Cisco IOS Guide at Page No. 253 ("[Path Option Selection for MPLS TE Tunnel LSPs This section contains the following topics about path option selection for MPLS TE Tunnel LSPs: • Constraint-Based Routing and Path Option Selection, page 5 • Tunnel Reoptimization and Path Option Selection, page 5 • Path Option Selection with Bandwidth Override, page 6 Constraint-Based Routing and Path Option Selection MPLS traffic engineering automatically establishes and maintains LSPs across the backbone by using the Resource Reservation Protocol (RSVP). The path that an LSP uses is determined

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 292 of 698

No.	'821 Patent Claim 4	Zamfir '948
		by the LSP resource requirements and network resources, such as bandwidth. Traffic engineering tunnels are calculated at the LSP head based on a fit between required and available resources (constraint-based routing). Without the Path Option for Bandwidth Override feature, a TE tunnel establishes an LSP based on dynamic or explicit path options in order of preference. However, the bandwidth and other attributes configured on the TE tunnel allow the setup of an LSP only if LSP path options satisfy the constraints. If a path cannot be found that satisfies the configured path options, then the tunnel is not set up.]") Cisco IOS Guide at Page No. 639 ("[Bandwidth Protection NHOP and NNHOP backup tunnels can be used to provide bandwidth protection for rerouted LSPs. This is referred to as backup bandwidth. You can associate backup bandwidth with NHOP or NNHOP backup tunnels. This informs the router of the amount of backup bandwidth a particular backup tunnel can protect. When a router maps LSPs to backup tunnels, bandwidth protection ensures that an LSP uses a given backup tunnel only if there is sufficient backup bandwidth. The router selects which LSPs use which backup tunnels in order to provide maximum bandwidth protection. That is, the router determines the best way to map LSPs onto backup tunnels in order to maximize the number of LSPs that can be protected. For information about mapping tunnels and assigning backup bandwidth, see the "Backup Tunnel Selection Procedure" section on page 10. LSPs that have the "bandwidth protection desired" bit set have a higher right to select backup tunnels that provide bandwidth protection; that is, those LSPs can preempt other LSPs that do not have that bit set. For more information, see the "Prioritizing Which LSPs Obtain Backup Tunnels with Bandwidth Protection" section on page 8.]") Cisco IOS Guide at Page No. 640 ("[Backup Tunnel Support Backup tunnel support has the following capabilities: • Backup Tunnels Can Terminate at the Next-Next Hop to Support FRR, page 6 • Multiple Backup Tunnels Can Protect the Same Interface, page 6 • Backup Tunnels Provide Scalability, page 7.]") Cisco IOS Guide at Page 640 ("[Multiple Backup Tunnels Can Protect the Same Interface

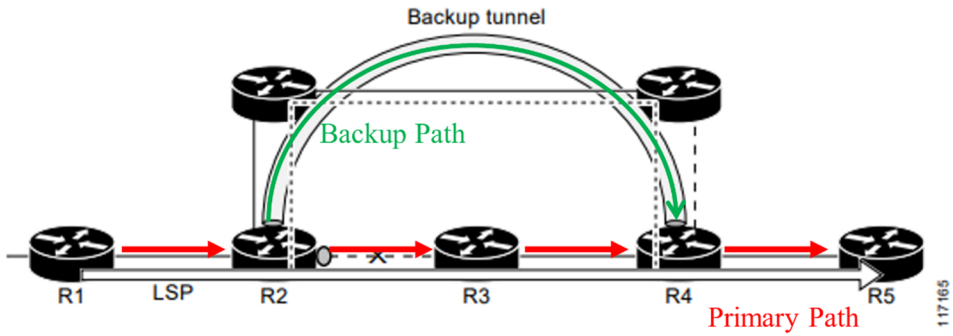
Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 293 of 698

No.	'821 Patent Claim 4	Zamfir '948
		There is no limit (except memory limitations) to the number of backup tunnels that can protect a given interface. In many topologies, support for node protection requires supporting multiple backup tunnels per protected interface. These backup tunnels can terminate at the same destination or at different destinations. That is, for a given protected interface, you can configure multiple NHOP or NNHOP backup tunnels. This allows redundancy and load balancing.]") Cisco IOS Guide at Page No. 603 ("[Traffic Engineering Tunnels MPLS TE lets you build label switched paths (LSPs) across your network for forwarding traffic. MPLS TE LSPs let the headend of a TE tunnel control the path its traffic takes to a particular destination. This method is more flexible than forwarding traffic based only on a destination address. Interarea tunnels allow you to do the following: • Build TE tunnels between areas (interarea tunnels) • Build TE tunnels that start and end in the same area, on multiple areas on a router (intra-area tunnels) Some tunnels are more important than others. For example, you may have tunnels carrying VoIP traffic and tunnels carrying data traffic that are competing for the same resources. MPLS TE allows you to have some tunnels preempt others. Each tunnel has a priority, and more-important tunnels take precedence over less-important tunnels.]") Cisco IOS Guide at Page No. 641 ("[Backup Bandwidth Protection Backup bandwidth protection allows you to give LSPs carrying certain kinds of data (such as voice) priority for using backup tunnels.]") Cisco IOS Guide at Page No. 650 ("[For each protected link AB with a maximum reservable subpool value of n, there may be a path from node A to node B such that the difference between the maximum reservable global and the maximum reservable subpool is at least the value of n. If it is possible to find such paths for each link in the network, you can establish all the backup tunnels along such paths without any bandwidth reservations. If there is a single link failure, only one backup tunnel will use any link on its path. Because that path has at

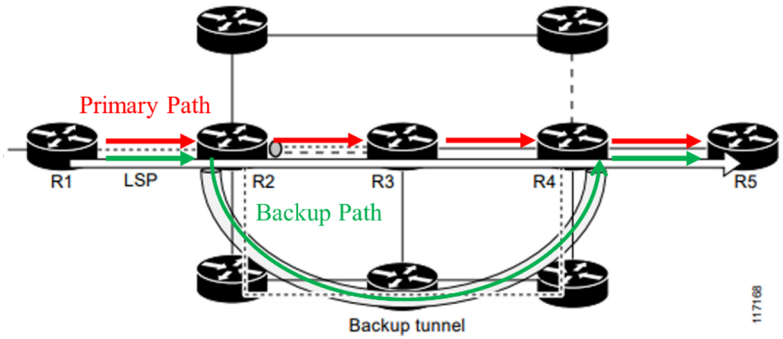
Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 294 of 698

No.	'821 Patent Claim 4	Zamfir '948
		least n available bandwidth (in the global pool), assuming that marking and scheduling is configured to classify the subpool traffic into a priority queue, the subpool bandwidth is guaranteed. This approach allows sharing of the global pool bandwidth between backup tunnels protecting independent link failures. The backup tunnels are expected to be used for only a short period of time after a failure (until the headends of affected LSPs reroute those LSPs to other paths with available subpool bandwidth). The probability of multiple unrelated link failures is very small (in the absence of node or shared risk link group (SRLG) failures, which result in multiple link failures). Therefore, it is reasonable to assume that link failures are in practice independent with high probability. This "independent failure assumption" in combination with backup tunnels signaled without explicit bandwidth reservation enables efficient bandwidth sharing that yields substantial bandwidth savings.]") Cisco IOS Guide at Page No. 651 ("[A similar approach can be used for node and SRLG protection. However, the decision of where to put the backup tunnels is more complicated because both node and SRLG failures effectively result in the simultaneous failure of several links. Therefore, the backup tunnels protecting traffic traversing all affected links cannot be computed independently of each other.]") Cisco IOS Guide at Page No. 579 ("[MPLS Traffic Engineering Shared Risk Link Groups SRLGs refer to situations where links in a network share a common fiber (or a common physical attribute). If one link fails, other links in the group may fail too. Links in the group have a shared risk. Backup tunnels should avoid using links in the same SRLG as interfaces they are protecting. Otherwise, when the protected link fails the backup tunnel fails too.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 295 of 698

No.	'821 Patent Claim 4	Zamfir '948
		Figure 1 Backup Tunnel in the Same SRLG as the Interface It Is Protecting  Fig. 1 (annotation added) Figure 1 shows a primary label-switched path (LSP) from router R1 to router R5. The LSP protects against the failure of the R2-R3 link at R2 via a backup tunnel to R4. If the R2-R3 link fails, link protection reroutes the LSP along the backup tunnel. However, the R2-R3 link and one of the backup tunnel links are in the same SRLG. So if the R2-R3 link fails, the backup tunnel may fail too. The MPLS TE SRLG feature enhances backup tunnel path selection so a backup tunnel can avoid using links that are in the same SRLG as the interfaces it is protecting. There are two ways for a backup tunnel to avoid the SRLGs of its protected interface: • The router does not create the backup tunnel unless it avoids SRLGs of the protected interface.]") Cisco IOS Guide at Page No. 580 ("[Open Shortest Path First (OSPF) and Intermediate System-to-Intermediate System (IS-IS) flood the SRLG membership information (including other TE link attributes such as bandwidth availability and affinity) so that all routers in the network have the SRLG information for each link. With this topology information, routers can compute backup tunnel paths that exclude links having SRLGs in common with their protected interfaces. As shown in Figure 2, the backup tunnel avoids the link between R2 and R3, which shares an SRLG with the protected interface.]")

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 296 of 698

No.	'821 Patent Claim 4	Zamfir '948
		Figure 2 Backup Tunnel That Avoids SRLG of Protected Interface  Fig. 2 (annotation added) Cisco IOS Guide at Page No. 587 ("[Step 4 show mpls traffic-eng topology Use this command to show the SRLG link membership flooded via the Interior Gateway Protocol (IGP).]") Cisco IOS Guide at Page 632 ("[IS-IS—Intermediate System-to-Intermediate System. Link-state hierarchical routing protocol that calls for intermediate system (IS) routers to exchange routing information based on a single metric to determine network topology.]") Cisco IOS Guide at Page No. 633 ("[OSPF—Open Shortest Path First. A link-state hierarchical Interior Gateway Protocol routing algorithm, derived from the IS-IS protocol. OSPF features include least-cost routing, multipath routing, and load balancing.]") Cisco IOS Guide at Page No. 581 ("[Autotunnel Backup for MPLS TE SRLGs Autotunnel backup is the ability of routers to create backup tunnels automatically. Therefore, you do not need to preconfigure each backup tunnel and then assign the backup tunnel to the protected interface. Only automatically created backup tunnels can avoid SRLGs or their protected interfaces.]") Cisco IOS Guide at Page No. 583 ("[Configuring the Routers That Automatically Create Backup Tunnels to Avoid MPLS TE SRLGs of Their Protected Interfaces

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 297 of 698

No.	'821 Patent Claim 4	Zamfir '948
		Perform the following task to configure routers that automatically create backup tunnels to avoid MPLS TE SRLGs of their protected interfaces. Backup tunnels provide link protection by rerouting traffic to the next hop bypassing failed links or in this instance by avoiding SRLGs.]") Cisco IOS Guide at Page No. 644 ("[Backup Tunnel Selection Procedure When an LSP is signaled, each node along the LSP path that provides FRR protection for the LSP selects a backup tunnel for the LSP to use if either of the following events occurs: • The link to the next hop fails. • The next hop fails. By having the node select the backup tunnel for an LSP before a failure occurs, the LSP can be rerouted onto the backup tunnel quickly if there is a failure.]") Cisco IOS Guide at Page 646 ("[NHOP Versus NNHOP Backup Tunnels More than one backup tunnel can protect a given LSP, where one backup tunnel terminates at the LSP's NNHOP, and the other terminates at the LSP's NHOP. In this case, the router chooses the backup tunnel that terminates at the NNHOP (that is, FRR prefers NNHOP over NHOP backup tunnels).Table 1 lists the tunnel selection priorities. The first choice is an NNHOP backup tunnel that acquires its bandwidth from a subpool or global pool, and has limited bandwidth. If there is no such backup tunnel, the next choice (2) is a next-next hop backup tunnel that acquires a limited amount of bandwidth from any pool. The preferences go from 1 (best) to 8 (worst), where choice 3 is for an NNHOP backup tunnel with an unlimited amount of subpool or global-pool bandwidth.]")

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 298 of 698

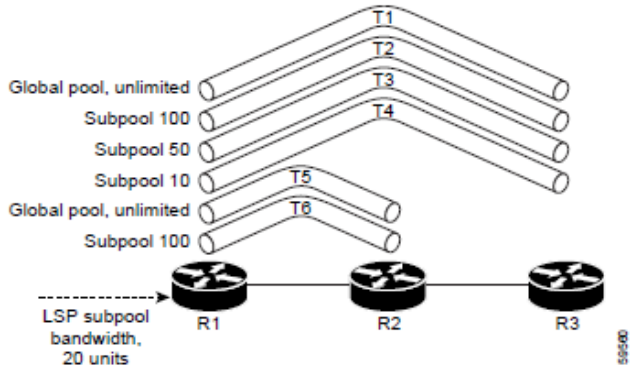
No.	'821 Patent Claim 4	Zamfir '948

Table 1 Tunnel Selection Priorities

Preference	Backup Tunnel Destination	Bandwidth Pool	Bandwidth Amount
1 (Best)	NNHOP	Subpool or global pool	Limited
2	NNHOP	Any	Limited
3	NNHOP	Subpool or global pool	Unlimited
4	NNHOP	Any	Unlimited
5	NHOP	Subpool or global pool	Limited
6	NHOP	Any	Limited
7	NHOP	Subpool or global pool	Unlimited
8 (Worst)	NHOP	Any	Unlimited

Cisco IOS Guide at Page No. 647-648 ("[Figure 7 shows an example of the backup tunnel selection procedure based on the designated amount of global pool and subpool bandwidth currently available.

Figure 7 Choosing from Among Multiple Backup Tunnels



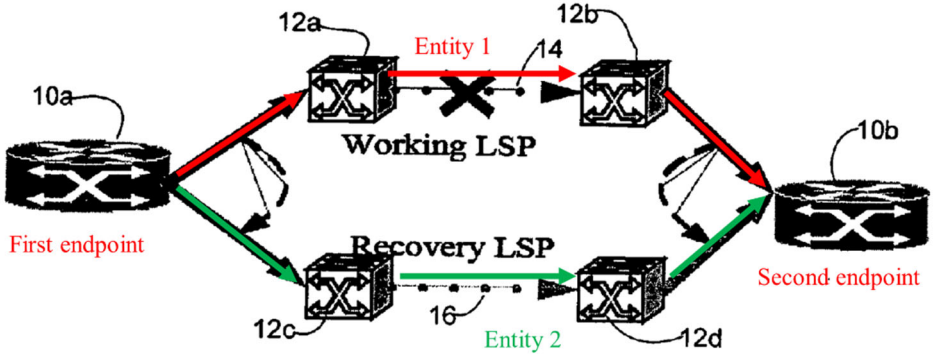
In this example, an LSP requires 20 units (kilobits per second) of sub-pool backup bandwidth. The best backup tunnel is selected as follows:
1. Backup tunnels T1 through T4 are considered first because they terminate at the NNHOP.
2. Tunnel T4 is eliminated because it has only ten units of sub-pool backup bandwidth.
3. Tunnel T1 is eliminated because it protects only LSPs using global-pool bandwidth.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 299 of 698

No.	'821 Patent Claim 4	Zamfir '948
		4. Tunnel T3 is chosen over T2 because, although both have sufficient backup bandwidth, T3 has the least backup bandwidth available (leaving the most backup bandwidth available on T2). 5. Tunnels T5 and T6 need not be considered because they terminate at an NHOP, and therefore are less desirable than T3, which terminates at an NNHOP.]")

No.	'821 Patent Claim 5	Zamfir '948
5	The method of claim 4, wherein said overall cost function comprises substantially minimizing a probability of concurrent failure of said protection entity and said working entity.	Zamfir '948 discloses wherein said overall cost function comprises substantially minimizing a probability of concurrent failure of said protection entity and said working entity. For example, Zamfir '948 discloses the establishment of a working LSP and a recovery LSP based on factors including the SRLG diversity degree. Based upon Orckit's apparent understanding of this limitation set forth in its Preliminary Infringement Disclosure, Zamfir '948 meets this limitation. To the extent that Zamfir '948 is found to not meet this limitation, wherein said overall cost function comprises substantially minimizing a probability of concurrent failure of said protection entity and said working entity would have been obvious to a person having ordinary skill in the art, as explained below. Zamfir '948 at Page No. 09, Para No. 0027 ("[A number of LSPs are established during GMPLS tunnel setup. The traffic may be arranged or duplicated any number of ways, as is well known by those skilled in the art. The paths are grouped into one or more bundles. With link bundling of the GMPLS interface, the component membership to the bundle is decided upon by the GMPLS control plane, which brings in the data plane aspect of dynamic membership. The LSPs in a bundle are established between the same source and destination nodes and, based on configuration or enforced SLA, they have a certain degree of diversity and optimality.]') Zamfir '948 at Page No. 10, Para No. 0035 ("[G-LSP and TE-LSP are preferably initially setup with a maximum diversity degree.]")

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 300 of 698

No.	'821 Patent Claim 5	Zamfir '948
		Zamfir '948 at Page No. 09, Para No. 0021 ("[FIG. 1A illustrates end-to-end protection with one working path 14 (between nodes 12 a and 12 b) and one protecting (recovery) path 16 (between nodes 12 c and 12 d). The working and protecting paths 14, 16 are preferably signaled as two LSPs that belong to the same tunnel. The working LSP 14 transports traffic between routers 10 a and 10 b and the protecting LSP 16 is used to transport traffic when the working LSP fails. End-to-end recovery provides recovery of the entire LSP from its source (ingress router end-point 10 a) to its destination (egress router end-point 10 b).]")  FIG. 1A Fig. 1A (annotation added) Zamfir '948 at Page No. 10, Para No. 0034 ("[Dynamic FRR extends link protection schemes, such as FRR (Fast Reroute), to handle GMPLS LSPs. A GMPLS LSP (G-LSP) is protected using a backup tunnel (TE-LSP) in a similar way as to how FRR is realized for physical links.]")

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 301 of 698

No.	'821 Patent Claim 6	Zamfir '948
6	The method of claim 4, wherein said overall cost function comprises a predefined entity cost metric.	Zamfir '948 discloses wherein said overall cost function comprises a predefined entity cost metric. For example, Zamfir '948 discloses the establishment of a working LSP and a recovery LSP based on factors including the SRLG diversity degree. Based upon Orckit's apparent understanding of this limitation set forth in its Preliminary Infringement Disclosure, Zamfir '948 meets this limitation. To the extent that Zamfir '948 is found to not meet this limitation, wherein said overall cost function comprises a predefined entity cost metric would have been obvious to a person having ordinary skill in the art, as explained below. Zamfir '948 at Page No. 09, Para No. 0027 ("[A number of LSPs are established during GMPLS tunnel setup. The traffic may be arranged or duplicated any number of ways, as is well known by those skilled in the art. The paths are grouped into one or more bundles. With link bundling of the GMPLS interface, the component membership to the bundle is decided upon by the GMPLS control plane, which brings in the data plane aspect of dynamic membership. The LSPs in a bundle are established between the same source and destination nodes and, based on configuration or enforced SLA, they have a certain degree of diversity and optimality.]') Zamfir '948 at Page No. 10, Para No. 0035 ("[G-LSP and TE-LSP are preferably initially setup with a maximum diversity degree.]")

No.	'821 Patent Claim 7	Zamfir '948
7	The method of claim 6, wherein said predefined entity cost metric is selected from the group consisting of interior gateway protocol (IGP) and	Zamfir '948 discloses wherein said predefined entity cost metric is selected from the group consisting of interior gateway protocol (IGP) and traffic engineering (TE). To the extent that Zamfir '948 is found to not meet this limitation, wherein said predefined entity cost metric is selected from the group consisting of interior gateway protocol (IGP) and traffic engineering (TE) would have been obvious to a person having ordinary skill in the art, as explained below.

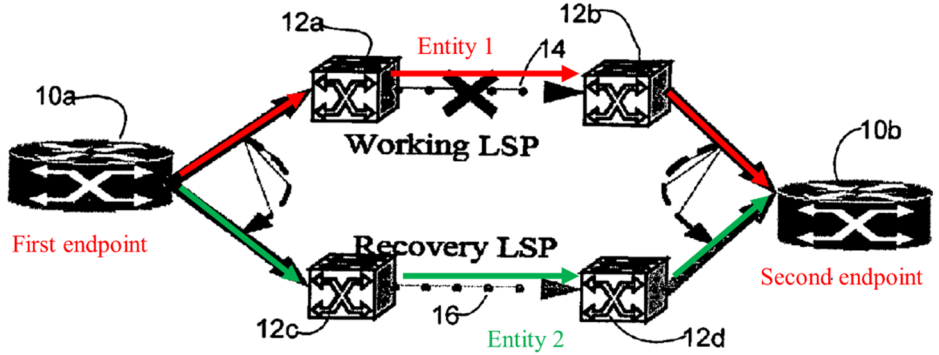
Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 302 of 698

No.	'821 Patent Claim 7	Zamfir '948
	traffic engineering (TE).	

No.	'821 Patent Claim 8	Zamfir '948
8	The method of claim 4, further comprising the step of configuring said working entity as revertive.	Zamfir '948 discloses the step of configuring said working entity as revertive. For example, Zamfir '948 discloses an immediate transition to a second path if there is an error on the first path. At least based upon Orckit's apparent understanding of this term as reflected in its Infringement Disclosures, this limitation is met. To the extent that Zamfir '948 is found to not meet this limitation, the step of configuring said working entity as revertive would have been obvious to a person having ordinary skill in the art, as explained below.

No.	'821 Patent Claim 9	Zamfir '948
9[preamble]	The method of claim 4, wherein said overall cost function comprises:	Zamfir '948 discloses an overall cost function. *See Supra at* 4.
9[a]	selecting a subset of entity pairs wherein each entity pair of said subset has substantially minimum probability of a concurrent failure of said protection entity and said working entity; and	Zamfir '948 discloses selecting a subset of entity pairs wherein each entity pair of said subset has substantially minimum probability of a concurrent failure of said protection entity and said working entity. For example, Zamfir '948 discloses the establishment of a working LSP and a recovery LSP based on factors including the SRLG diversity degree. Based upon Orckit's apparent understanding of this limitation set forth in its Preliminary Infringement Disclosure, Zamfir '948 meets this limitation. To the extent that Zamfir '948 is found to not meet this limitation, selecting a subset of entity pairs wherein each entity pair of said subset has substantially minimum probability of a concurrent failure of said protection entity and said working entity would have been obvious to a person having ordinary skill in the art, as explained below. Zamfir '948 at Page No. 09, Para No. 0027 ("[A number of LSPs are established during GMPLS tunnel setup. The traffic may be arranged or duplicated any number of ways, as is

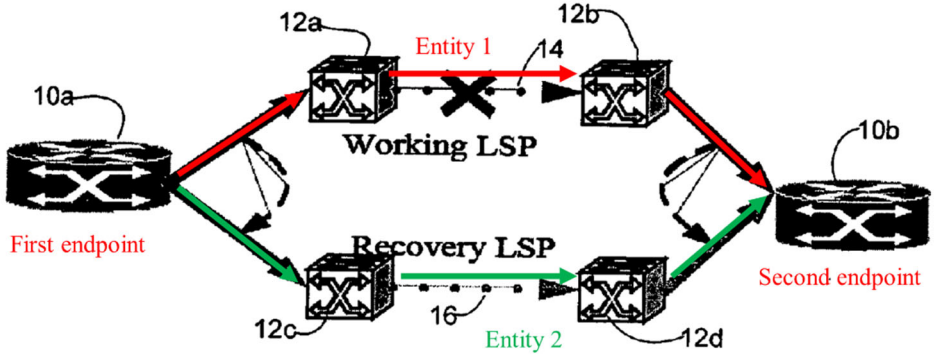
Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 303 of 698

No.	'821 Patent Claim 9	Zamfir '948
		well known by those skilled in the art. The paths are grouped into one or more bundles. With link bundling of the GMPLS interface, the component membership to the bundle is decided upon by the GMPLS control plane, which brings in the data plane aspect of dynamic membership. The LSPs in a bundle are established between the same source and destination nodes and, based on configuration or enforced SLA, they have a certain degree of diversity and optimality.]") Zamfir '948 at Page No. 10, Para No. 0035 ("[G-LSP and TE-LSP are preferably initially setup with a maximum diversity degree.]") Zamfir '948 at Page No. 09, Para No. 0021 ("[FIG. 1A illustrates end-to-end protection with one working path 14 (between nodes 12 a and 12 b) and one protecting (recovery) path 16 (between nodes 12 c and 12 d). The working and protecting paths 14, 16 are preferably signaled as two LSPs that belong to the same tunnel. The working LSP 14 transports traffic between routers 10 a and 10 b and the protecting LSP 16 is used to transport traffic when the working LSP fails. End-to-end recovery provides recovery of the entire LSP from its source (ingress router end-point 10 a) to its destination (egress router end-point 10 b).]")  FIG. 1A

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 304 of 698

No.	'821 Patent Claim 9	Zamfir '948
		Fig. 1A (annotation added) Zamfir '948 at Page No. 10, Para No. 0034 ("[Dynamic FRR extends link protection schemes, such as FRR (Fast Reroute), to handle GMPLS LSPs. A GMPLS LSP (G-LSP) is protected using a backup tunnel (TE-LSP) in a similar way as to how FRR is realized for physical links.]")
9[b]	if said subset comprises at least two entity pairs, selecting an entity pair from said subset that minimizes an entity cost function.	Zamfir '948 discloses if said subset comprises at least two entity pairs, selecting an entity pair from said subset that minimizes an entity cost function. For example, Zamfir '948 discloses establishing LSPs in a bundle (e.g., G-LSP and TE-LSP) between the same source and destination nodes based on configuration or enforced SLA and having a certain degree of diversity e.g., SLRG diversity and optimality. Based upon Orckit's apparent understanding of this limitation set forth in its Preliminary Infringement Disclosure, Zamfir '948 meets this limitation. To the extent that Zamfir '948 is found to not meet this limitation, if said subset comprises at least two entity pairs, selecting an entity pair from said subset that minimizes an entity cost function would have been obvious to a person having ordinary skill in the art, as explained below. Zamfir '948 at Page No. 09, Para No. 0027 ("[A number of LSPs are established during GMPLS tunnel setup. The traffic may be arranged or duplicated any number of ways, as is well known by those skilled in the art. The paths are grouped into one or more bundles. With link bundling of the GMPLS interface, the component membership to the bundle is decided upon by the GMPLS control plane, which brings in the data plane aspect of dynamic membership. The LSPs in a bundle are established between the same source and destination nodes and, based on configuration or enforced SLA, they have a certain degree of diversity and optimality.]") Zamfir '948 at Page No. 10, Para No. 0035 ("[G-LSP and TE-LSP are preferably initially setup with a maximum diversity degree.]")

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 305 of 698

No.	'821 Patent Claim 9	Zamfir '948
		Zamfir '948 at Page No. 09, Para No. 0021 ("[FIG. 1A illustrates end-to-end protection with one working path 14 (between nodes 12 a and 12 b) and one protecting (recovery) path 16 (between nodes 12 c and 12 d). The working and protecting paths 14, 16 are preferably signaled as two LSPs that belong to the same tunnel. The working LSP 14 transports traffic between routers 10 a and 10 b and the protecting LSP 16 is used to transport traffic when the working LSP fails. End-to-end recovery provides recovery of the entire LSP from its source (ingress router end-point 10 a) to its destination (egress router end-point 10 b).]")  FIG. 1A Fig. 1A (annotation added) Zamfir '948 at Page No. 10, Para No. 0034 ("[Dynamic FRR extends link protection schemes, such as FRR (Fast Reroute), to handle GMPLS LSPs. A GMPLS LSP (G-LSP) is protected using a backup tunnel (TE-LSP) in a similar way as to how FRR is realized for physical links.]")

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 306 of 698

No.	'821 Patent Claim 10	Zamfir '948
10	The method of claim 9, wherein said entity cost function comprises a predefined metric.	Zamfir '948 discloses wherein said entity cost function comprises a predefined metric. *See* Claim 6 supra,

No.	'821 Patent Claim 11	Zamfir '948
11	The method of claim 10, wherein said predefined metric is selected from the group consisting of interior gateway protocol (IGP) and traffic engineering (TE).	Zamfir '948 discloses wherein said predefined metric is selected from the group consisting of interior gateway protocol (IGP) and traffic engineering (TE). *See* Claim 7 supra.

No.	'821 Patent Claim 12	Zamfir '948
12	The method of claim 10, further comprising the step of configuring said working entity as revertive.	Zamfir '948 discloses The method of claim 10, further comprising the step of configuring said working entity as revertive. *See supra at* Claim 8.

No.	'821 Patent Claim 13	Zamfir '948
13	The method of claim 1, further comprising the step of: if said entity pair reselection results in both working and protection entities being replaced,	Zamfir '948 discloses the method of claim 1, further comprising the step of: if said entity pair reselection results in both working and protection entities being replaced, sequentially replacing said working entity and said protection entity. For example, Zamfir '948 discloses that if G-LSP1 needs to be re-optimized, a determination is made if the (G-LSP2, TE-LSP1) bundle pair still satisfies the SRLG diversity requirement. Zamfir '948 further discloses that if this requirement is met, G-LSP2 and TE-LSP1 are the

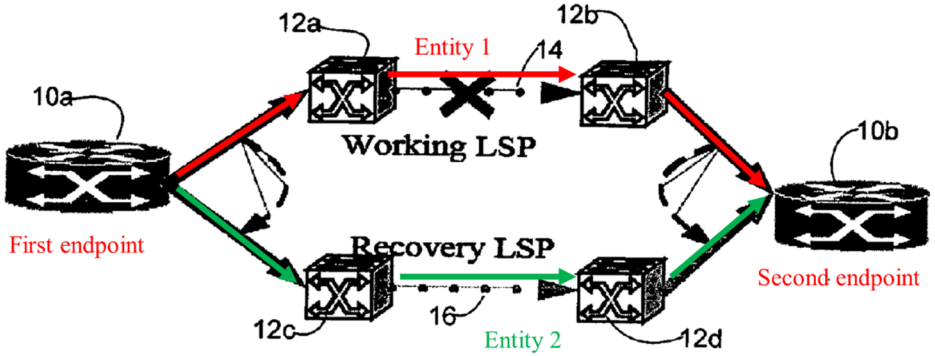
Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 307 of 698

No.	'821 Patent Claim 13	Zamfir '948
	sequentially replacing said working entity and said protection entity.	new working and protecting paths. Zamfir '948 further discloses that if this requirement is not met, a new TE-LSP2 is signaled for the backup tunnel. Zamfir '948 at Page No. 10, Para No. 0037 ("[FIG. 5 is a flowchart illustrating one example of a process of the present invention for dynamic FRR. At step 50, G-LSP1 needs to be re-optimized. A new (better) route is determined and G-LSP2 is ready to be signaled using a make-before-break procedure (step 52). Since G-LSP1 is protected by TE-LSP1, a determination is made if the (G-LSP2, TE-LSP1) pair still satisfies the diversity requirement (step 54). If this requirement is met, G-LSP2 and TE-LSP1 are the new working and protecting paths, respectively (step 56). If this requirement is not met, a new TE-LSP2 is signaled for the backup tunnel using a make-before-break procedure (steps 58 and 60). In order to minimize the window where a failure will have an impact on the data traffic, these two procedures are synchronized. The diversity requirement may be specified, for example, with respect to link diversity, node diversity, SLRG diversity, or any combination thereof.]") Zamfir '948 at Claim 23 ("[23. A method for providing dynamic path protection in an optical network in communication with an IP network, the method comprising: establishing a set of working and protecting paths, each of said working paths having at least one of said protecting path designated as its backup path; following a failure or during optimization, selecting and establishing at least one new working path to replace one of said working paths; determining if said new working path and said designated backup path provide a predefined degree of diversity; and selecting a new protecting path if said predefined degree of diversity is not met.]")

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 308 of 698

No.	'821 Patent Claim 14	Zamfir '948
14[preamble]	A system for selecting entities within an MPLS network, comprising:	Zamfir '948 discloses a system for selecting entities within an MPLS network. *See Supra at* 1[Preamble] and 1[a].
14[a]	a data structure comprising a plurality of transport entity descriptors;	Zamfir '948 discloses a plurality of transport entity descriptors. For example, Zamfir '948 discloses a processor on a headend router that creates two or more paths/label switched paths (LSP) within the optical network between the same source and destination nodes. A person of ordinary skill in the art would understand from Zamfir '948's disclosures that a data structure in computer-readable memory stores these instructions. Zamfir '948 at Page No. 09, Para No. 0021 ("[FIG. 1A illustrates an example of a network in which data transmission occurs on generalized label-switched paths (G-LSPs). The G-LSPs are automatically setup and torn down by means of a signaling protocol, as is well known by those skilled in the art. (In the following description, the term LSP is also used to refer to generalized LSPs.) FIG. 1A illustrates end-to-end protection with one working path 14 (between nodes 12 a and 12 b) and one protecting (recovery) path 16 (between nodes 12 c and 12 d). The working and protecting paths 14, 16 are preferably signaled as two LSPs that belong to the same tunnel. The working LSP 14 transports traffic between routers 10 a and 10 b and the protecting LSP 16 is used to transport traffic when the working LSP fails. End-to-end recovery provides recovery of the entire LSP from its source (ingress router end-point 10 a) to its destination (egress router end-point 10 b).]")

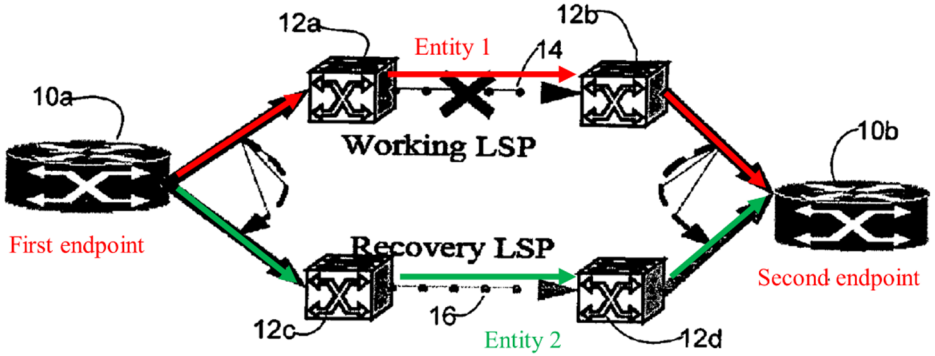
Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 309 of 698

No.	'821 Patent Claim 14	Zamfir '948
		 FIG. 1A Fig. 1A (annotation added) Zamfir '948 Page No. 09, Para No. 0022 ("[It is to be understood that the arrangement shown in FIG. 1A is only one example and the invention described herein may be applied to a network having any number of working or recovery LSPs.]") Zamfir '948 at Page No. 09, Para No. 0027 ("[A number of LSPs are established during GMPLS tunnel setup. The traffic may be arranged or duplicated any number of ways, as is well known by those skilled in the art. The paths are grouped into one or more bundles. With link bundling of the GMPLS interface, the component membership to the bundle is decided upon by the GMPLS control plane, which brings in the data plane aspect of dynamic membership. The LSPs in a bundle are established between the same source and destination nodes and, based on configuration or enforced SLA, they have a certain degree of diversity and optimality.]")

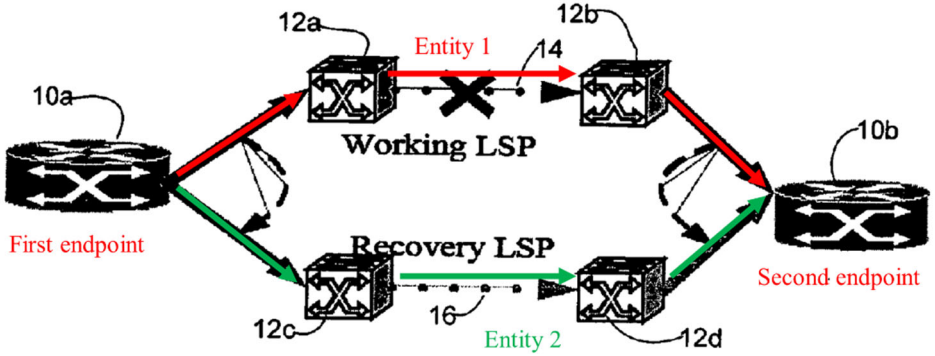
Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 310 of 698

No.	'821 Patent Claim 14	Zamfir '948
		Zamfir '948 at Claim 19 ("[19. A system for dynamic protection end-to-end path protection in an optical network in communication with an IP network, the system comprising: a processor operable to create two or more paths within the optical network and group two or more of said paths in a bundle, said bundle comprising at least one working path and at least one protecting path each extending from an ingress router to an egress router; and a protection module configured to dynamically identify a physical interface associated with an optical tunnel interface for a new path and replace one of said paths in said bundle with said new path.]")
14[b]	an entity protection switch configured to switch between a working entity and a protection entity; and	Zamfir '948 discloses an entity protection switch configured to switch between a working entity and a protection entity. For example, Zamfir '948 discloses the working LSP transports traffic between routers and the protecting LSP is used to transport traffic when the working LSP fails. Based upon Orckit's apparent understanding of this limitation set forth in its Preliminary Infringement Disclosure, Zamfir '948 meets this limitation. Zamfir '948 at Page No. 09, Para No. 0021 ("[FIG. 1A illustrates an example of a network in which data transmission occurs on generalized label-switched paths (G-LSPs). The G-LSPs are automatically setup and torn down by means of a signaling protocol, as is well known by those skilled in the art. (In the following description, the term LSP is also used to refer to generalized LSPs.) FIG. 1A illustrates end-to-end protection with one working path 14 (between nodes 12 a and 12 b) and one protecting (recovery) path 16 (between nodes 12 c and 12 d). The working and protecting paths 14, 16 are preferably signaled as two LSPs that belong to the same tunnel. The working LSP 14 transports traffic between routers 10 a and 10 b and the protecting LSP 16 is used to transport traffic when the working LSP fails. End-to-end recovery provides recovery of the entire LSP from its source (ingress router end-point 10 a) to its destination (egress router end-point 10 b).]")

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 311 of 698

No.	'821 Patent Claim 14	Zamfir '948
		 FIG. 1A Fig. 1A (annotation added) Zamfir '948 Page No. 09, Para No. 0022 ("[It is to be understood that the arrangement shown in FIG. 1A is only one example and the invention described herein may be applied to a network having any number of working or recovery LSPs.]")
14[c]	digital logic configured to select said working entity and said protection entity from said plurality of transport entity descriptors, comprising:	Zamfir '948 discloses digital logic configured to select said working entity and said protection entity from said plurality of transport entity descriptors. For example, Zamfir '948 discloses a processor on a headend router that creates two or more paths/label switched paths (LSP) within the optical network between the same source and destination nodes. A person of ordinary skill in the art would understand from Zamfir '948's disclosures that a data structure in computer-readable memory stores these instructions. Zamfir '948 at Page No. 09, Para No. 0021 ("[FIG. 1A illustrates end-to-end protection with one working path 14 (between nodes 12 a and 12 b) and one protecting (recovery) path 16 (between nodes 12 c and 12 d). The working and protecting paths 14, 16 are preferably

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 312 of 698

No.	'821 Patent Claim 14	Zamfir '948
		signaled as two LSPs that belong to the same tunnel. The working LSP 14 transports traffic between routers 10 a and 10 b and the protecting LSP 16 is used to transport traffic when the working LSP fails. End-to-end recovery provides recovery of the entire LSP from its source (ingress router end-point 10 a) to its destination (egress router end-point 10 b).]")  FIG. 1A Fig. 1A (annotation added) Zamfir '948 at Page No. 10, Para No. 0034 ('[Dynamic FRR extends link protection schemes, such as FRR (Fast Reroute), to handle GMPLS LSPs. A GMPLS LSP (G-LSP) is protected using a backup tunnel (TE-LSP) in a similar way as to how FRR is realized for physical links.]") Zamfir '948 at Page No. 10, Para No. 0035 ("[G-LSP and TE-LSP are preferably initially setup with a maximum diversity degree.]") Zamfir '948 at Claim 19 ("[19. A system for dynamic protection end-to-end path protection in an optical network in communication with an IP network, the system comprising:

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 313 of 698

No.	'821 Patent Claim 14	Zamfir '948
		a processor operable to create two or more paths within the optical network and group two or more of said paths in a bundle, said bundle comprising at least one working path and at least one protecting path each extending from an ingress router to an egress router; and a protection module configured to dynamically identify a physical interface associated with an optical tunnel interface for a new path and replace one of said paths in said bundle with said new path.]")
14[d]	logic configured to determine a probability of concurrent failure of said working entity and said protection entity;	Zamfir '948 discloses logic configured to determine a probability of concurrent failure of said working entity and said protection entity. For example, Zamfir '948 discloses a processor on a headend router that creates two or more paths/label switched paths (LSP) within the optical network between the same source and destination nodes. A person of ordinary skill in the art would understand from Zamfir '948's disclosures that a data structure in computer-readable memory stores these instructions. Based upon Orckit's apparent understanding of this limitation set forth in its Preliminary Infringement Disclosure, Zamfir '948 meets this limitation. Zamfir '948 at Page No. 09, Para No. 0027 ("[A number of LSPs are established during GMPLS tunnel setup. The traffic may be arranged or duplicated any number of ways, as is well known by those skilled in the art. The paths are grouped into one or more bundles. With link bundling of the GMPLS interface, the component membership to the bundle is decided upon by the GMPLS control plane, which brings in the data plane aspect of dynamic membership. The LSPs in a bundle are established between the same source and destination nodes and, based on configuration or enforced SLA, they have a certain degree of diversity and optimality.]') Zamfir '948 at Page No. 10, Para No. 0035 ("[G-LSP and TE-LSP are preferably initially setup with a maximum diversity degree.]")

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 314 of 698

No.	'821 Patent Claim 14	Zamfir '948
		Zamfir '948 at Claim 19 ("[19. A system for dynamic protection end-to-end path protection in an optical network in communication with an IP network, the system comprising: a processor operable to create two or more paths within the optical network and group two or more of said paths in a bundle, said bundle comprising at least one working path and at least one protecting path each extending from an ingress router to an egress router; and a protection module configured to dynamically identify a physical interface associated with an optical tunnel interface for a new path and replace one of said paths in said bundle with said new path.]")
14[e]	logic configured to determine an entity cost of said plurality of transport entity descriptors: and	Zamfir '948 discloses logic configured to determine an entity cost of said plurality of transport entity descriptors. For example, Zamfir '948 discloses the LSPs in a bundle are established (G-LSP and TE-LSP grouped in a bundle) between the same source and destination nodes based on configuration or enforced SLA and have a certain degree of diversity e.g., SLRG diversity and optimality, and each LSP have an associated cost. For example, Zamfir '948 discloses a processor on a headend router that creates two or more paths/label switched paths (LSP) within the optical network between the same source and destination nodes. A person of ordinary skill in the art would understand from Zamfir '948's disclosures that a data structure in computer-readable memory stores these instructions. Zamfir '948 at Page No. 09, Para No. 0021 ("[FIG. 1A illustrates an example of a network in which data transmission occurs on generalized label-switched paths (G-LSPs). The G-LSPs are automatically setup and torn down by means of a signaling protocol, as is well known by those skilled in the art. (In the following description, the term LSP is also used to refer to generalized LSPs.) FIG. 1A illustrates end-to-end protection with one working path 14 (between nodes 12 a and 12 b) and one protecting (recovery) path 16 (between nodes 12 c and 12 d). The working and protecting paths 14, 16 are preferably signaled as two LSPs that belong to the same tunnel. The working LSP 14 transports traffic between routers 10 a and 10 b and the protecting LSP 16 is used to transport traffic when the working LSP fails. End-

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 315 of 698

No.	'821 Patent Claim 14	Zamfir '948
		to-end recovery provides recovery of the entire LSP from its source (ingress router end-point 10 a) to its destination (egress router end-point 10 b).]")

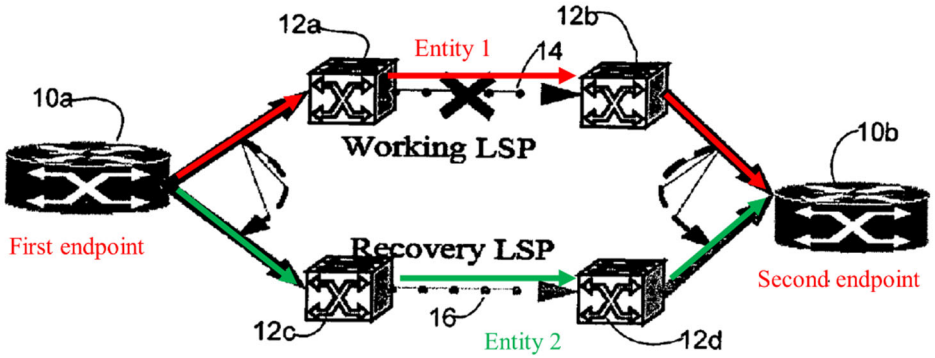


FIG. 1A

Fig. 1A (annotation added)

Zamfir '948 at Page No. 09, Para No. 0022 ("[It is to be understood that the arrangement shown in FIG. 1A is only one example and the invention described herein may be applied to a network having any number of working or recovery LSPs.]")

Zamfir '948 at Page No. 09, Para No. 0027 ("[A number of LSPs are established during GMPLS tunnel setup. The traffic may be arranged or duplicated any number of ways, as is well known by those skilled in the art. The paths are grouped into one or more bundles. With link bundling of the GMPLS interface, the component membership to the bundle is decided upon by the GMPLS control plane, which brings in the data plane aspect of dynamic membership. The LSPs in a bundle are established between the same source and destination

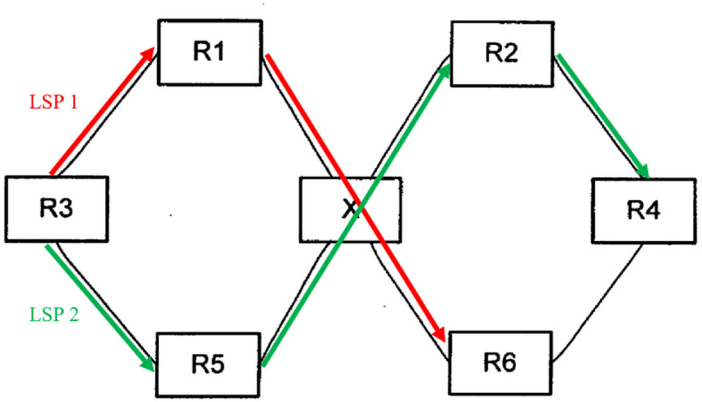
Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 316 of 698

No.	'821 Patent Claim 14	Zamfir '948
		nodes and, based on configuration or enforced SLA, they have a certain degree of diversity and optimality.]") Zamfir '948 at(Page No. 10, Para No. 0031 ("[In order to better illustrate the make-a-set-before-break-a-set process, the following example is provided. In this example, a dynamic bundle with two components LSP1 and LSP2 is created, with the two LSPs being setup with link and node diversity (FIG. 3): LSP1−R3, R1, R6, R∝(cost=30) LSP2−R3, R5, R2, R4 (cost=30) In this example, all links are assumed to have cost equal to 10, therefore LSP1 and LSP2 each have a cost equal to 30.]") Zamfir '948 at Page No. 10, Para No. 0034 ("[Dynamic FRR extends link protection schemes, such as FRR (Fast Reroute), to handle GMPLS LSPs. A GMPLS LSP (G-LSP) is protected using a backup tunnel (TE-LSP) in a similar way as to how FRR is realized for physical links.]") Zamfir '948 at Page No. 10, Para No. 0035 ("[G-LSP and TE-LSP are preferably initially setup with a maximum diversity degree.]")
14[f]	logic configured to reselect said working entity and said protection entity from said plurality of transport entity descriptors upon a reselection event,	Zamfir '948 discloses logic configured to reselect said working entity and said protection entity from said plurality of transport entity descriptors upon a reselection event. For example, Zamfir '948 discloses that when a re-optimization operation is triggered (for example in the case where there is a change in topology or when an LSP needs to be re-routed due to a network device failure), an attempt is preferably made to maintain the degree of diversity/optimality at the set level and path computation is used to determine a new LSP set (G-LSP2, TE-LSP2).

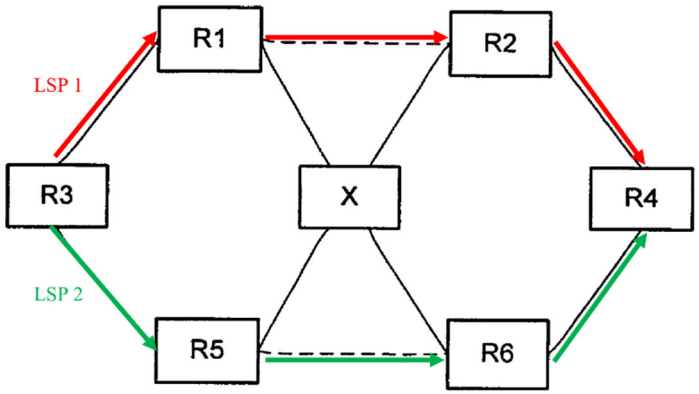
Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 317 of 698

No.	'821 Patent Claim 14	Zamfir '948
		Zamfir '948 at Page No. 09, Para No. 0027 ("[A number of LSPs are established during GMPLS tunnel setup. The traffic may be arranged or duplicated any number of ways, as is well known by those skilled in the art. The paths are grouped into one or more bundles. With link bundling of the GMPLS interface, the component membership to the bundle is decided upon by the GMPLS control plane, which brings in the data plane aspect of dynamic membership. The LSPs in a bundle are established between the same source and destination nodes and, based on configuration or enforced SLA, they have a certain degree of diversity and optimality. These LSPs may be re-optimized or re-routed due to topology changes or failures in the network. As different network events change the egress physical interfaces of the LSPs, the membership in the dynamic bundle changes on-the-fly.]") Zamfir '948 at Page No. 09, Para No. 0028 ("[Another aspect of dynamic bundles is the handling of LSP bundles in the control plane. When an LSP set is created, if the routes are dynamically computed, a certain degree of diversity/optimality may be requested and realized. When a re-optimization operation is triggered (for example in the case where there is a change in topology or when an LSP needs to be re-routed due to a failure), an attempt is preferably made to maintain the degree of diversity/optimality at the set level. To realize this, path computation is used to determine a new LSP set. The LSPs in the new set can share resources with any subset of the old paths. This procedure is referred to as "make-a-set-before-break-a-set" and is described below.]") Zamfir '948 at Page No. 10, Para No. 0030 ("[When a new optical tunnel is created and added to the bundle, the physical interfaces associated with the optical tunnel interface is identified by a protection module and added to the bundle. If a failure occurs on an optical LSP which is also a component in the L2 bundle, the GMPLS control plane restores the LSP using, for example, a new physical interface and the protection module performs the translation to the new interface that will replace the old one. In other words, as GMPLS LSPs can change paths on-the-fly (e.g., due to rerouting in the event of a failure, re-optimization, etc.), TE control

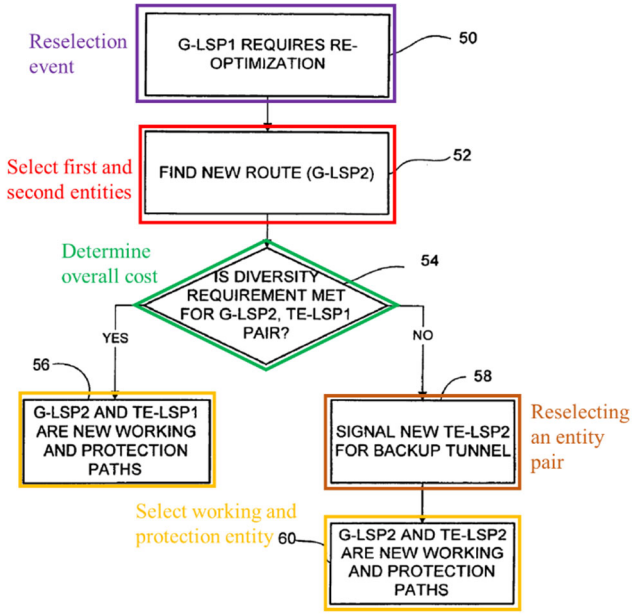
Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 318 of 698

No.	'821 Patent Claim 14	Zamfir '948
		can change membership of the bundled interface in the event of a re-route of the GMPLS LSP.]") Zamfir '948 at Page No. 10, Para No. 0031 ("[In order to better illustrate the make-a-set-before-break-a-set process, the following example is provided. In this example, a dynamic bundle with two components LSP1 and LSP2 is created, with the two LSPs being setup with link and node diversity (FIG. 3): LSP1−R3, R1, R6, R∝(cost=30) LSP2−R3, R5, R2, R4 (cost=30)]")  Fig. 3 (annotation added) Zamfir '948 at Page No. 10, Para No. 0032 ("[As shown in FIG. 4, two new links (shown with dashed lines) with cost equal to 5 are added to connect R1 with R2 and R5 with R6. If re-optimized as a set the new LSPs will be:

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 319 of 698

No.	'821 Patent Claim 14	Zamfir '948
		LSP1−R3, R1, R2, R4 (cost 25) LSP2−R3, R5, R6, R4 (cost 25)]")  Fig. 4 (annotation added) Zamfir '948 at Page No. 10, Para No. 0037 ("[FIG. 5 is a flowchart illustrating one example of a process of the present invention for dynamic FRR. At step 50, G-LSP1 needs to be re-optimized. A new (better) route is determined and G-LSP2 is ready to be signaled using a make-before-break procedure (step 52). Since G-LSP1 is protected by TE-LSP1, a determination is made if the (G-LSP2, TE-LSP1) pair still satisfies the diversity requirement (step 54). If this requirement is met, G-LSP2 and TE-LSP1 are the new working and protecting paths, respectively (step 56). If this requirement is not met, a new TE-LSP2 is signaled for the backup tunnel using a make-before-break procedure (steps 58 and 60). In order to minimize the window where a failure will have an impact on the data traffic, these two procedures are synchronized. The diversity requirement may be specified, for example, with respect to link diversity, node diversity, SLRG diversity, or any combination thereof.]")

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 320 of 698

No.	'821 Patent Claim 14	Zamfir '948
		 Fig. 5 (annotation added) Zamfir '948 at Claim 23 ("[23. A method for providing dynamic path protection in an optical network in communication with an IP network, the method comprising: establishing a set of working and protecting paths, each of said working paths having at least one of said protecting path designated as its backup path; following a failure or during optimization, selecting and establishing at least one new working path to replace one of said working paths; determining if said new working path and said designated backup path provide a predefined degree of diversity; and

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 321 of 698

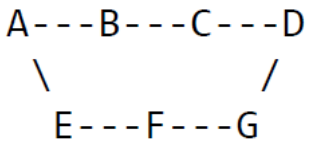
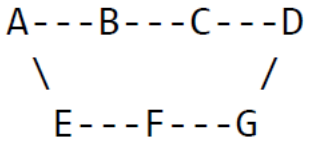
No.	'821 Patent Claim 14	Zamfir '948
		selecting a new protecting path if said predefined degree of diversity is not met.]") Zamfir '948 at Claim 24 ("[24. The method of claim 23 wherein selecting a new working path comprises utilizing fast reroute.]")
14[g]	wherein said reselection event is selected from a group consisting of adding an entity to said plurality of transport entities, removing an entity from said plurality of transport entities, an operational status change for one of said plurality of transport entities, and a change in overall cost for one of said plurality of transport entities.	Zamfir '948 discloses wherein said reselection event is selected from a group consisting of adding an entity to said plurality of transport entities, removing an entity from said plurality of transport entities, an operational status change for one of said plurality of transport entities, and a change in overall cost for one of said plurality of transport entities. *See Supra at* 1[e].

No.	'821 Patent Claim 15	Zamfir '948
15	The system of claim 14 wherein said entity protection switch comprises a 1:1 switch.	Zamfir '948 discloses the system of claim 14 wherein said entity protection switch comprises a 1:1 switch. For example, Zamfir '948 discloses protection by a single protection path where the bridge at the head end is not permanent and in the event of the primary path failing, it switches to the alternate path. To the extent that Zamfir '948 is found to not meet this limitation, wherein said

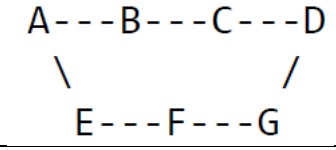
Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 322 of 698

No.	'821 Patent Claim 15	Zamfir '948
		entity protection switch comprises a 1:1 switch would have been obvious to a person having ordinary skill in the art, as explained below. In addition, at least under the apparent claim scope alleged by Orckit's Infringement Disclosures, Zamfir '948 in combination with (1) the knowledge of a person of ordinary skill in the art, alone or in further combination with (2) each (individually, as well as one or more together) of the references identified in element 15 of Exhibit E-2 renders the claim, including the present limitation, obvious. Below are examples of two such references. For example, RFC 4872 discloses entity protection switches operable under both a 1:1 and 1+1 architecture. "7. 1:1 Protection with Extra-Traffic The most common case of end-to-end 1:N protection is to establish, between the same endpoints, an end-to-end working LSP (thus, N = 1) and a dedicated end-to-end protecting LSP that are mutually link/ node/SRLG disjoint. This protects against working LSP failure(s). The protecting LSP is used for switchover when the working LSP fails. GMPLS RSVP-TE signaling allows for the pre-provisioning of protecting LSPs by indicating in the Path message (in the PROTECTION object; see Section 14) that the LSPs are of type protecting. Here, working and protecting LSPs are signaled as primary LSPs; both are fully instantiated during the provisioning phase. Although the resources for the protecting LSP are pre-allocated, preemptable traffic may be carried end-to-end using this LSP. Thus, the protecting LSP is capable of carrying extra-traffic with the caveat that this traffic will be preempted if the working LSP fails. The setup of the working LSP SHOULD indicate that the LSP head-end and tail-end node wish to receive Notify messages using the NOTIFY REQUEST object. The node upstream to the failure (upstream in terms of the direction an Path message traverses) SHOULD send a Notify message to the LSP head-end node, and the node downstream to the failure SHOULD send an Notify message to the LSP tail-end node. Upon receipt of the Notify messages, both the end-nodes MUST switch the (normal) traffic from the working LSP to the pre-configured protecting LSP (see Section 7.2). Moreover, some coordination is required if extra-traffic is carried over the end-to-end protecting LSP. Note that if the working and the protecting LSP are established between the same end-nodes, no further notification is required to indicate that the working LSPs are no longer protected.

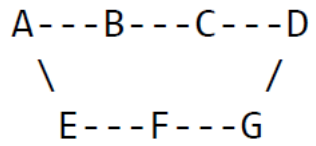
Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 323 of 698

No.	'821 Patent Claim 15	Zamfir '948
		Consider the following topology:  The working LSP [A,B,C,D] could be protected by the protecting LSP [A,E,F,G,D]. Both LSPs are fully instantiated (resources are allocated for both working and protecting LSPs) and no resource sharing can be done along the protection path since the primary protecting LSP can carry extra-traffic. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 13–14. "6. 1+1 Bidirectional Protection 1+1 bidirectional protection is a scheme that provides end-to-end protection for bidirectional LSPs. Consider the following network topology:  The LSPs [A,B,C,D] and [A,E,F,G,D] are node and link disjoint, ignoring the ingress/egress nodes A and D. A bidirectional LSP is established from A to D over each path, and traffic is transmitted simultaneously over both LSPs. In this scheme, both endpoints must receive traffic over the same LSP. Note also that both LSPs are fully instantiated (and thus activated) so that no resource sharing can be done along the protection path (nor can any extra-traffic be transported). When a failure is detected by one or both endpoints of the LSP, both endpoints must select traffic from the other LSP. This action must be coordinated between node A and D. From this perspective, 1+1 bidirectional protection can be seen as a coordinated protection- switching mechanism between both endpoints. Note: it is necessary that both paths are SRLG disjoint

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 324 of 698

No.	'821 Patent Claim 15	Zamfir '948
		to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 10–11. For example, RFC 4872 discloses entity protection switches operable under both a 1:1 and 1+1 architecture. "7. 1:1 Protection with Extra-Traffic The most common case of end-to-end 1:N protection is to establish, between the same endpoints, an end-to-end working LSP (thus, N = 1) and a dedicated end-to-end protecting LSP that are mutually link/ node/SRLG disjoint. This protects against working LSP failure(s). The protecting LSP is used for switchover when the working LSP fails. GMPLS RSVP-TE signaling allows for the pre-provisioning of protecting LSPs by indicating in the Path message (in the PROTECTION object; see Section 14) that the LSPs are of type protecting. Here, working and protecting LSPs are signaled as primary LSPs; both are fully instantiated during the provisioning phase. Although the resources for the protecting LSP are pre-allocated, preemptable traffic may be carried end-to-end using this LSP. Thus, the protecting LSP is capable of carrying extra-traffic with the caveat that this traffic will be preempted if the working LSP fails. The setup of the working LSP SHOULD indicate that the LSP head-end and tail-end node wish to receive Notify messages using the NOTIFY REQUEST object. The node upstream to the failure (upstream in terms of the direction an Path message traverses) SHOULD send a Notify message to the LSP head-end node, and the node downstream to the failure SHOULD send an Notify message to the LSP tail-end node. Upon receipt of the Notify messages, both the end-nodes MUST switch the (normal) traffic from the working LSP to the pre-configured protecting LSP (see Section 7.2). Moreover, some coordination is required if extra-traffic is carried over the end-to-end protecting LSP. Note that if the working and the protecting LSP are established between the same end-nodes, no further notification is required to indicate that the working LSPs are no longer protected. Consider the following topology: 

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 325 of 698

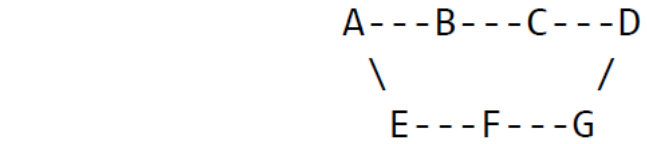
No.	'821 Patent Claim 15	Zamfir '948
		The working LSP [A,B,C,D] could be protected by the protecting LSP [A,E,F,G,D]. Both LSPs are fully instantiated (resources are allocated for both working and protecting LSPs) and no resource sharing can be done along the protection path since the primary protecting LSP can carry extra-traffic. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 13–14. "6. 1+1 Bidirectional Protection 1+1 bidirectional protection is a scheme that provides end-to-end protection for bidirectional LSPs. Consider the following network topology:  The LSPs [A,B,C,D] and [A,E,F,G,D] are node and link disjoint, ignoring the ingress/egress nodes A and D. A bidirectional LSP is established from A to D over each path, and traffic is transmitted simultaneously over both LSPs. In this scheme, both endpoints must receive traffic over the same LSP. Note also that both LSPs are fully instantiated (and thus activated) so that no resource sharing can be done along the protection path (nor can any extra-traffic be transported). When a failure is detected by one or both endpoints of the LSP, both endpoints must select traffic from the other LSP. This action must be coordinated between node A and D. From this perspective, 1+1 bidirectional protection can be seen as a coordinated protection- switching mechanism between both endpoints. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 10–11. For example, Cisco IOS Guide discloses when there is a failure on the protected LSP, the headend router immediately enables the secondary LSP to temporarily carry the tunnel's traffic (1:1 protection).

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 326 of 698

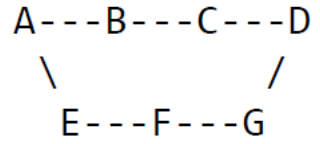
No.	'821 Patent Claim 15	Zamfir '948
		Cisco IOS Guide Page No. 603 ("[Path Protection Path protection provides an end-to-end failure recovery mechanism (that is, full path protection) for MPLS TE tunnels. A secondary LSP is established, in advance, to provide failure protection for the protected LSP that is carrying a tunnel's TE traffic. When there is a failure on the protected LSP, the headend router immediately enables the secondary LSP to temporarily carry the tunnel's traffic. If there is a failure on the secondary LSP, the tunnel no longer has path protection until the failure along the secondary path is cleared.]") Cisco IOS Guide at Page No. 641 ("[One-to-one protection is when a separate backup tunnel must be used for each LSP needing protection. N:1 protection has significant scalability advantages over one-to-one (1:1) protection. An example of 1:1 protection is when 5000 backup tunnels protect 5000 LSPs, each router along the backup path must maintain state for an additional 5000 tunnels.]")

No.	'821 Patent Claim 16	Zamfir '948
16	The system of claim 14 wherein said entity protection switch comprises a 1+1 switch.	Zamfir '948 discloses the system of claim 14 wherein said entity protection switch comprises a 1+1 switch. For example, Zamfir '948 discloses protection where each working channel has a continuously active backup path. To the extent that Zamfir '948 is found to not meet this limitation, wherein said entity protection switch comprises a 1+1 switch would have been obvious to a person having ordinary skill in the art, as explained below. In addition, at least under the apparent claim scope alleged by Orckit's Infringement Disclosures, Zamfir '948 in combination with (1) the knowledge of a person of ordinary skill in the art, alone or in further combination with (2) each (individually, as well as one or more together) of the references identified in element 16 of Exhibit E-2 renders the claim, including the present limitation, obvious. Below are examples of two such references. For example, RFC 4872 discloses entity protection switches operable under both a 1:1 and 1+1 architecture.

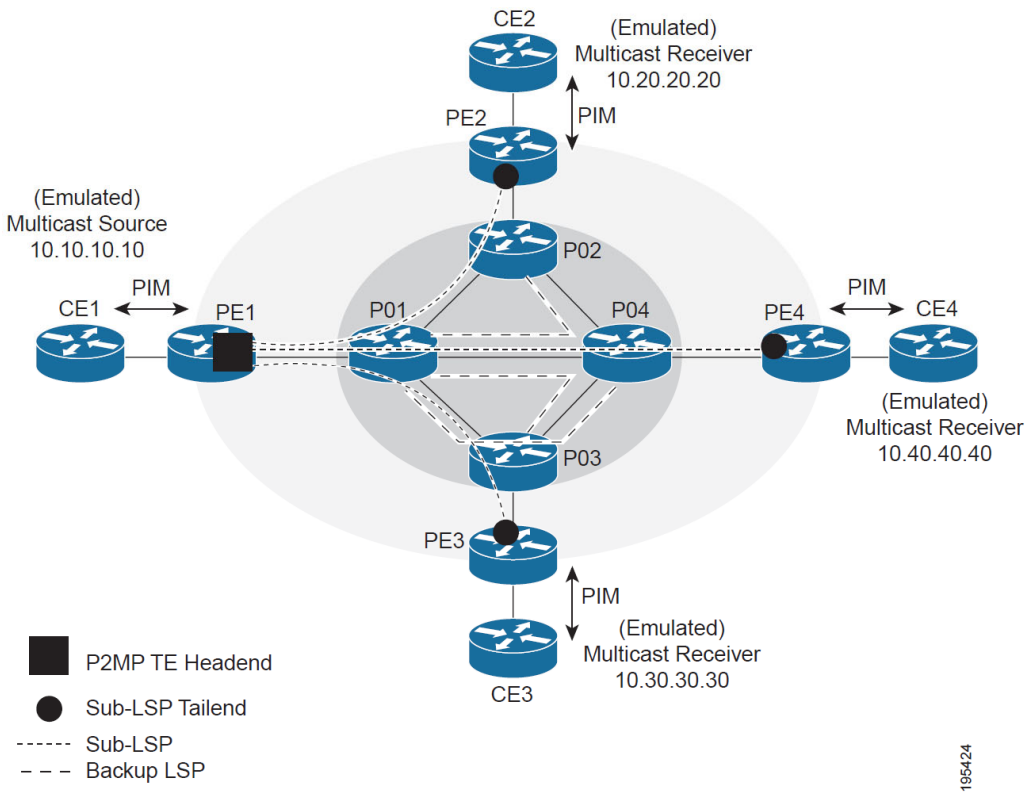
Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 327 of 698

No.	'821 Patent Claim 16	Zamfir '948
		"7. 1:1 Protection with Extra-Traffic The most common case of end-to-end 1:N protection is to establish, between the same endpoints, an end-to-end working LSP (thus, N = 1) and a dedicated end-to-end protecting LSP that are mutually link/ node/SRLG disjoint. This protects against working LSP failure(s). The protecting LSP is used for switchover when the working LSP fails. GMPLS RSVP-TE signaling allows for the pre-provisioning of protecting LSPs by indicating in the Path message (in the PROTECTION object; see Section 14) that the LSPs are of type protecting. Here, working and protecting LSPs are signaled as primary LSPs; both are fully instantiated during the provisioning phase. Although the resources for the protecting LSP are pre-allocated, preemptable traffic may be carried end-to-end using this LSP. Thus, the protecting LSP is capable of carrying extra-traffic with the caveat that this traffic will be preempted if the working LSP fails. The setup of the working LSP SHOULD indicate that the LSP head-end and tail-end node wish to receive Notify messages using the NOTIFY REQUEST object. The node upstream to the failure (upstream in terms of the direction an Path message traverses) SHOULD send a Notify message to the LSP head-end node, and the node downstream to the failure SHOULD send an Notify message to the LSP tail-end node. Upon receipt of the Notify messages, both the end-nodes MUST switch the (normal) traffic from the working LSP to the pre-configured protecting LSP (see Section 7.2). Moreover, some coordination is required if extra-traffic is carried over the end-to-end protecting LSP. Note that if the working and the protecting LSP are established between the same end-nodes, no further notification is required to indicate that the working LSPs are no longer protected. Consider the following topology:  The working LSP [A,B,C,D] could be protected by the protecting LSP [A,E,F,G,D]. Both LSPs are fully instantiated (resources are allocated for both working and protecting LSPs) and no resource sharing can be done along the protection path since the primary protecting LSP can carry extra-traffic. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 13–14.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 328 of 698

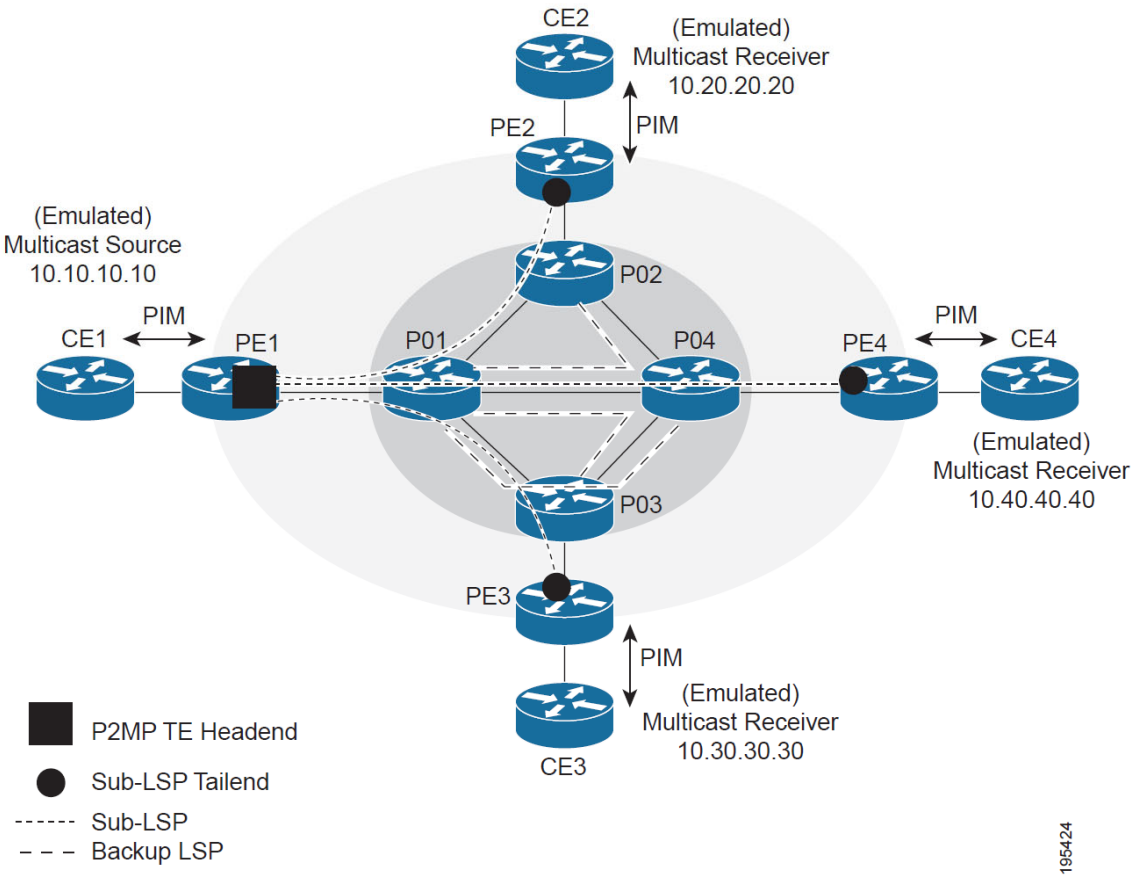
No.	'821 Patent Claim 16	Zamfir '948
		"6. 1+1 Bidirectional Protection 1+1 bidirectional protection is a scheme that provides end-to-end protection for bidirectional LSPs. Consider the following network topology:  The LSPs [A,B,C,D] and [A,E,F,G,D] are node and link disjoint, ignoring the ingress/egress nodes A and D. A bidirectional LSP is established from A to D over each path, and traffic is transmitted simultaneously over both LSPs. In this scheme, both endpoints must receive traffic over the same LSP. Note also that both LSPs are fully instantiated (and thus activated) so that no resource sharing can be done along the protection path (nor can any extra-traffic be transported). When a failure is detected by one or both endpoints of the LSP, both endpoints must select traffic from the other LSP. This action must be coordinated between node A and D. From this perspective, 1+1 bidirectional protection can be seen as a coordinated protection- switching mechanism between both endpoints. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 10–11. For example, Cisco's IOS Guide discloses: "Figure 2 shows a P2MP TE tunnel that has three destinations. • PE1 is the headend router. • P01 is a branch point router, where packet replication occurs. • PE2, PE3, and PE4 are tailend routers, where the sub-LSP ends. Between the PE and CE routers, PIM is enabled to exchange multicast routing information with the directly connected customer edge (CE) routers. PIM is not enabled across the P2MP TE tunnel." Cisco IOS Guide at Page No. 468.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 329 of 698

No.	'821 Patent Claim 16	Zamfir '948
		Figure 2 Network Topology with P2MP TE Tunnel  Cisco IOS Guide at Page No. 469, FIG. 2. "RSVP TE extensions defined in RFC 4875 allow multiple sub-LSPs to be signaled from the headend router. A P2MP TE tunnel consists of multiple sub-LSPs that connect the headend router to various tailend routers. The headend router sends one RSVP path message to each destination. The tailend router replies with a RESV message. The Label Forwarding Information Base (LFIB) is populated using the RSVP labels allocated by the RESV messages.

No.	'821 Patent Claim 16	Zamfir '948
		The tailend routers allocate unreserved labels, which are greater than 15 and do not include implicit or explicit null labels. Using unreserved labels allows IP Multicast to perform a Reverse Path Forwarding (RPF) check on the tailend router. Because a sub-LSP tailend router cannot be represented as a regular interface, a special LSP virtual interface (VIF) is automatically created. The LSP VIF represents the originating interface for all IP multicast traffic originating from the P2MP TE tailend router." Cisco IOS Guide at Page No. 469. "You can configure P2P explicit backup tunnels on point of local repair (PLR) nodes for link protection of P2MP sub-LSPs, similar to LSPs for P2P TE tunnels. You can also enable automatic creation of backup tunnels using the Auto-tunnel Backup feature for P2P TE tunnels. All sibling sub-LSPs that share the same outgoing link are protected by the same backup tunnel. All cousin sub-LSPs that share the same outgoing link can be protected by multiple P2P backup tunnels. Link protection for a P2MP TE tunnel is illustrated in Figure 6, which shows PE1 as the tunnel headend router and PE2, PE3, and PE4 as tunnel tailend routers. The following sub-LSPs are signaled from PE1 in the network: • From PE1 to PE2, the sub-LSP travels the following path: PE1 -> P01 -> P02 -> PE2 • From PE1 to PE3, the sub-LSP travels the following path: PE1 -> P01 -> P03 -> PE3 • From PE1 to PE4, the sub-LSP travels the following path: PE1 -> P01 -> P04 -> PE4 Node P01 is a branch node that does packet replication in the MPLS forwarding plane; ingress traffic originating from PE1 will be replicated towards routers P02, P03, and P04." Cisco IOS Guide at Page No. 474-475.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 331 of 698

No.	'821 Patent Claim 16	Zamfir '948
		Figure 6 P2MP TE Link Protection Example  Cisco IOS Guide at Page No. 475.

No.	'821 Patent Claim 17	Zamfir '948
17[preamble]	Non-transitory computer readable media configured to	Zamfir '948 discloses non-transitory computer readable media.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 332 of 698

No.	'821 Patent Claim 17	Zamfir '948
	perform a method comprising the steps of:	For example, Zamfir '948 discloses a computer program product and/or processor that performs dynamic protection. Zamfir '948 at Claim 14 ("[14. A computer program product for dynamic protection of virtual links for paths of an optical network in communication with an IP network; the product comprising: providing dynamic link protection in an optical network in communication with an IP network, comprising: code that establishes two or more paths within the optical network; code that groups two or more of said paths in a dynamic bundle, the groups selected at a control plane level; and a computer-readable storage medium for storing said codes; wherein said bundle of paths are recognized as a single routing adjacency in the IP network such that as long as one or more paths within the bundle is operating the routing adjacency in the IP network is not affected by changes within said bundle.]") Zamfir '948 at Claim 17 ("[17. A system for dynamic protection of virtual links for paths of an optical network in communication with an IP network, the system comprising: means for establishing two or more paths within the optical network; and means for grouping two or more of said paths in a dynamic bundle, the grouping selected at a control plane level; wherein said bundle of paths are recognized as a single routing adjacency in the IP network such that as long as one or more paths within the bundle is operating the routing adjacency in the IP network is not affected by changes within said bundle.]") Zamfir '948 at Claim 23 ("[23. A method for providing dynamic path protection in an optical network in communication with an IP network, the method comprising: establishing a set of working and protecting paths, each of said working paths having at least one of said protecting path designated as its backup path; following a failure or during optimization, selecting and establishing at least one new working path to replace one of said working paths;

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 333 of 698

No.	'821 Patent Claim 17	Zamfir '948
		determining if said new working path and said designated backup path provide a predefined degree of diversity; and selecting a new protecting path if said predefined degree of diversity is not met.]") Zamfir '948 at Claim 19 ("[19. A system for dynamic protection end-to-end path protection in an optical network in communication with an IP network, the system comprising: a processor operable to create two or more paths within the optical network and group two or more of said paths in a bundle, said bundle comprising at least one working path and at least one protecting path each extending from an ingress router to an egress router; and a protection module configured to dynamically identify a physical interface associated with an optical tunnel interface for a new path and replace one of said paths in said bundle with said new path.]")
17[a]	providing a plurality of MPLS transport entities between a first endpoint and a second endpoint;	Zamfir '948 discloses providing a plurality of multiprotocol label switching (MPLS) transport entities between a first endpoint and a second endpoint. *See Supra at* 1[a].
17[b]	determining an overall cost for each entity pair of said plurality of entities;	Zamfir '948 discloses determining an overall cost for each entity pair of said plurality of entities. *See Supra at* 1[b].
17[c]	selecting an entity pair from said plurality of transport entities based at least in part upon said overall cost; and	Zamfir '948 discloses selecting an entity pair from said plurality of transport entities based at least in part upon said overall cost. *See Supra at* 1[c].
17[d]	if an entity pair reselection event occurs, reselecting	Zamfir '948 discloses if an entity pair reselection event occurs, reselecting said entity pair from the group consisting of said entity pair and a replacement entity pair comprising at least one entity distinct from the entities of said entity pair.

No.	'821 Patent Claim 17	Zamfir '948
	said entity pair from the group consisting of said entity pair and a replacement entity pair comprising at least one entity distinct from the entities of said entity pair,	*See Supra at* 1[d].
17[e]	wherein said entity pair reselection event is selected from a group consisting of adding an entity to said plurality of transport entities, removing an entity from said plurality of transport entities, an operational status change for one of said plurality of transport entities, and a change in overall cost for one of said plurality of transport entities.	Zamfir '948 discloses wherein said entity pair reselection event is selected from a group consisting of adding an entity to said plurality of transport entities, removing an entity from said plurality of transport entities, an operational status change for one of said plurality of transport entities. *See Supra at* 1[e].

No.	'821 Patent Claim 18	Zamfir '948
18[preamble]	The non-transitory computer readable media of claim 17, wherein said step of	Zamfir '948 discloses the step of selecting an entity pair. *See Supra at* 1[b] and 17.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 335 of 698

No.	'821 Patent Claim 18	Zamfir '948
	selecting an entity pair further comprises:	
18[a]	selecting a working entity from said plurality of transport entities;	Zamfir '948 discloses selecting a working entity from said plurality of transport entities. *See Supra at* 2[a].
18[b]	selecting a protection entity from said plurality of transport entities; and	Zamfir '948 discloses selecting a protection entity from said plurality of transport entities. *See Supra at* 2[b].
18[c]	selecting an active entity from the set consisting of said working entity and said protection entity.	Zamfir '948 discloses selecting an active entity from the set consisting of said working entity and said protection entity. *See Supra at* 3.

No.	'821 Patent Claim 19	Zamfir '948
19	The non-transitory readable media of claim 18, wherein said step of selecting an entity pair further comprises minimizing an overall cost function.	Zamfir '948 discloses step of selecting an entity pair minimizing an overall cost function. *See Supra at* 4.

No.	'821 Patent Claim 20	Zamfir '948
20[preamble]	The non-transitory readable media of claim 19, wherein said overall cost function comprises:	Zamfir '948 discloses an overall cost function. For example, Zamfir '948 discloses a computer program product for dynamic protection of virtual links which groups two or more paths in a dynamic bundle by considering maximum diversity degree.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 336 of 698

No.	'821 Patent Claim 20	Zamfir '948
		Zamfir '948 at Claim 14 ("[14. A computer program product for dynamic protection of virtual links for paths of an optical network in communication with an IP network; the product comprising: providing dynamic link protection in an optical network in communication with an IP network, comprising: code that establishes two or more paths within the optical network; code that groups two or more of said paths in a dynamic bundle, the groups selected at a control plane level; and a computer-readable storage medium for storing said codes; wherein said bundle of paths are recognized as a single routing adjacency in the IP network such that as long as one or more paths within the bundle is operating the routing adjacency in the IP network is not affected by changes within said bundle.]") Zamfir '948 at Page No. 09, Para No. 0027 ("[A number of LSPs are established during GMPLS tunnel setup. The traffic may be arranged or duplicated any number of ways, as is well known by those skilled in the art. The paths are grouped into one or more bundles. With link bundling of the GMPLS interface, the component membership to the bundle is decided upon by the GMPLS control plane, which brings in the data plane aspect of dynamic membership. The LSPs in a bundle are established between the same source and destination nodes and, based on configuration or enforced SLA, they have a certain degree of diversity and optimality.]") Zamfir '948 at Page No. 10, Para No. 0035 ("[G-LSP and TE-LSP are preferably initially setup with a maximum diversity degree.]")
20[a]	minimizing a probability of concurrent failure of said protection entity and said working entity; and	Zamfir '948 discloses minimizing a probability of concurrent failure of said protection entity and said working entity. *See Supra at* 5.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 337 of 698

No.	'821 Patent Claim 20	Zamfir '948
20[b]	a predefined metric selected from the group consisting of interior gateway protocol (IGP) and traffic engineering (TE).	Zamfir '948 discloses a predefined metric selected from the group consisting of interior gateway protocol (IGP) and traffic engineering (TE). *See Supra at* 7.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 338 of 698

EXHIBIT B-2
Defendant's First Amended Invalidity Contentions
Orckit Corporation v. Cisco Systems, Inc., 2:22-cv-00276-JRG-RSP

Chart for U.S. Patent 8,830,821 ("the '821 Patent")
Juniper Networks' Junos OS Release 11.1 ("Juniper Junos OS System")

As shown in the chart below, all claims of the '821 Patent are invalid (1) under 35 U.S.C. §§ 102 (a) and (g) based on Juniper Junos OS System, (2) under 35 U.S.C. §§ 102 (a) and (b) because the references describing the Juniper Junos OS System disclose every limitation of every Asserted Claim, and (3) under 35 U.S.C. § 103 as obvious based on Juno OS Multiprotocol Label Switching Configuration Guide Release 11.1 alone, or in combination with the knowledge of a person of ordinary skill in the art alone, and in further combination with references specifically identified in the following claim chart and/or one or more other references identified in Cisco's Preliminary Invalidity Contentions. The Juniper Junos OS System comprises the Juno IOS Release 11.1 operating system, as it ran on various Juniper Networks switches and routers from before June 22, 2011. The Juniper Junos OS System was invented, known and used by others in the United States, and on sale in the United States before the claimed invention, rendering the claims invalid. Publications describing the Juniper Junos OS System, including those cited in this chart, were publicly disclosed and available more than a year before the purported filing date of the provisional application identified on the face of the '821 Patent, rendering the claims invalid. For example, Juno OS Multiprotocol Label Switching Configuration Guide Release 11.1 was published in 2010 and discloses all elements of the Asserted Claims, thereby anticipating the Asserted Claims. Moreover, as identified in this chart, it would have been obvious to combine the disclosures in the publications regarding the Juniper Junos OS System with each other and with disclosures in other publications known in the art, as explained in this chart, well before the claimed invention and the effective filing date of the '821 Patent, rendering the claims invalid.

The quotations and diagrams regarding the Juniper Junos OS System in this chart come from the following sources:

- Juniper Junos® OS Multiprotocol Label Switching Configuration Guide Release 11.1 ("Juniper Junos OS Guide");
- Juniper Networks, Inc. LSP Hierarchy support for MPLS LSPs (19230_JNPR_00000071);
- Juniper sw-projects/mpls/misc/fast-reroute/draft-ietf-mpls-rsvp-lsp-fastreroute-00.nroff (19230_JNPR_00000074);
- Juniper Networks, Inc. Flash LDP and KRT route prioritization (19230_JNPR_00000099);
- Juniper Networks, Inc. MPLS Fate-sharing (cut-set) file Functional Specification (19230_JNPR_00000102);
- Juniper Networks, Inc. Tunneling LDP LSPs in RSVP LSPs Functional Specification (19230_JNPR_00000105);
- Juniper Networks, Inc. MPLS Fate-sharing (cut-set) file Functional Specification v. 1.3 (19230_JNPR_00000110);
- Juniper Networks, Inc. Advanced Insight Solutions User Guide Release 1.3 (19230_JNPR_00000121);

- Junos® Software Broadband Subscriber Management Solutions Guide Release 9.5 (19230_JNPR_00000415);
- Junos® Channelized IQ Interfaces Feature Guide Release 9.5 (19230_JNPR_00000499);
- Junos® Class of Service Using IPv6 DiffServ Feature Guide Release 9.5 (19230_JNPR_00000599);
- Junos® Software Network Management Configuration Guide Release 9.5 (19230_JNPR_00000655);
- Junos® Software Feature Guide Release 9.5 (19230_JNPR_00001433);
- Junos® Software Glossary Release 9.3 (19230_JNPR_00002237);
- Junos® Design and Implementation Guide for J-series Services Routers Release 9.4 (19230_JNPR_00002357);
- Junos® J-series Services Routers Hardware Guide Release 9.4 (19230_JNPR_00002612);
- Junos® Migration Guide for J-series Services Routers Release 9.5 (19230_JNPR_00002892);
- WXC Integrated Services Module Installation and Configuration Guide Release 9.4 (19230_JNPR_00002984);
- JUNOS® Software Administration Guide for J-series Services Routers and SRX-series Services Gateways Release 9.5 (19230_JNPR_00003034);
- JUNOS® Software CLI Reference for J-series Services Routers and SRX-series Services Gateways Release 9.5 (19230_JNPR_00003522);
- JUNOS® Software Interfaces and Routing Configuration Guide for J-series Services Routers and SRX-series Services Gateways Release 9.5 (19230_JNPR_00004722);
- JUNOS® Software Security Configuration Guide for J-series Services Routers and SRX-series Services Gateways Release 9.5 (19230_JNPR_00005688);
- JUNOS®Software JUNOS XML API Configuration Reference Release 9.5 (19230_JNPR_00006734);
- JUNOS®Software JUNOS XML API Operational Reference Release 9.5 (19230_JNPR_00018044);
- JUNOS™ Software JUNOScope Software User Guide Release 9.5 (19230_JNPR_00020416);
- JUNOS® Software JUNOScript API Guide Release 9.5 (19230_JNPR_00021220);
- JUNOS™ Software J-Web Interface User Guide Release 9.5 (19230_JNPR_00021502);
- JUNOS® Software Multiplay Solutions Guide Release 9.5 (19230_JNPR_00021632);
- JUNOS® Software MX-series Solutions Guide Release 9.5 (19230_JNPR_00021882);
- JUNOS® Software NETCONF API Guide Release 9.5 (19230_JNPR_00022034);
- JUNOS™ Internet Software for M-series and T-series Routing Platforms Baseline Operations Guide (19230_JNPR_00022208);
- JUNOS™ Internet Software for M-series and T-series Routing Platforms Network Operations Guide Hardware (19230_JNPR_00022588);
- JUNOS™ Internet Software for M-series and T-series Routing Platforms Interfaces Operations Guide (19230_JNPR_00023308);
- JUNOS™ Internet Software MPLS Fast Reroute Solutions Network Operations Guide (19230_JNPR_00023694);

- JUNOS™ Internet Software for M-series and T-series Routing Platforms MPLS Network Operations Guide Log Reference (19230_JNPR_00023950);
- JUNOS™ Internet Software for M-series and T-series Routing Platforms MPLS Network Operations Guide (19230_JNPR_00024170);
- JUNOS® Software Protected System Domain Configuration Guide Release 9.5 (19230_JNPR_00024414);
- JUNOS Feature Guides Routing Matrix Feature Guide Release 9.5 (19230_JNPR_00024708);
- JUNOS Source Class Usage Feature Guide Release 9.5 (19230_JNPR_00024772);
- JUNOS® Software System Basics and Services Command Reference Release 9.5 (19230_JNPR_00024818);
- JUNOS® Software Interfaces Command Reference Release 9.5 (19230_JNPR_00025926);
- JUNOS® Software Routing Protocols and Policies Command Reference Release 9.5 (19230_JNPR_00026740);
- JUNOS® Software Access Privilege Configuration Guide Release 9.5 (19230_JNPR_00027476);
- JUNOS® Internet Software Configuration and Diagnostic Automation Guide Release 9.5 (19230_JNPR_00027568);
- JUNOS® Software CLI User Guide Release 9.5 (19230_JNPR_00028066);
- JUNOS® Software Class of Service Configuration Guide Release 9.5 (19230_JNPR_00028362);
- JUNOS® Software Software Installation and Upgrade Guide Release 9.5 (19230_JNPR_00028906);
- JUNOS® Software MX-series Layer 2 Configuration Guide Release 9.5 (19230_JNPR_00028998);
- JUNOS® Software MPLS Applications Configuration Guide Release 9.5 (19230_JNPR_00029168);
- JUNOS® Software Multicast Protocols Configuration Guide Release 9.5 (19230_JNPR_00029718);
- JUNIPER−ANALYZER−MIB DEFINITIONS (19230_JNPR_00030246);
- Juniper enterprise specific ATM Class Of Service MIB (19230_JNPR_00030253);
- Juniper Enterprise Specific ATM MIB (19230_JNPR_00030264);
- Juniper BFD−STD−MIB DEFINITIONS (19230_JNPR_00030280);
- JUNIPER−BFD−MIB DEFINITIONS (19230_JNPR_00030294);
- Juniper draft−ietf−idr−bgp4−mibv2−03.txt (19230_JNPR_00030297);
- Juniper Enterprise Specific MIB: Configuration Management MIB Extension (19230_JNPR_00030364);
- Juniper chassis mib definitions: OIDs used to identify various platforms and chassis components (19230_JNPR_00030369);
- Juniper Enterprise Specific MIB: J−Series FWDD (Forwarding Daemon) MIB (19230_JNPR_00030402);
- Juniper Enterprise Specific MIB: Chassis MIB (19230_JNPR_00030404);
- Juniper Enterprise Specific MIB: Collector PIC MIB (19230_JNPR_00030432);
- Juniper enterprise specific Class−Of−Service (COS) MIB (19230_JNPR_00030447);
- Juniper Enterprise Specific MIB: Destination Class Usage MIB (19230_JNPR_00030484);
- JUNIPER−DFC−MIB DEFINITIONS (19230_JNPR_00030488);

- Juniper enterprise specific Event MIB (19230_JNPR_00030502);
- Juniper Enterprise Specific MIB: Structure of Management Information (19230_JNPR_00030505);
- Juniper Networks: SNMP Experimental MIB Registry (19230_JNPR_00030506);
- Juniper Enterprise Specific MIB: Firewalls MIB (19230_JNPR_00030508);
- Juniper Enterprise Specific MIB: Host Resources MIB (19230_JNPR_00030512);
- Juniper Enterprise Specific MIB: Interface MIB Extension (19230_JNPR_00030514);
- Juniper Enterprise Specific MIB: ipForward (19230_JNPR_00030523);
- Juniper Networks IPSEC Generic Flow Monitoring object mibs (19230_JNPR_00030525);
- JNX−IPSEC−MONITOR−MIB DEFINITIONS (19230_JNPR_00030542);
- Juniper Enterprise Specific MIB: ipv4 MIB Extension (19230_JNPR_00030559);
- Juniper Enterprise Specific MIB: ipv6 MIB Extension (19230_JNPR_00030562);
- Juniper enterprise specific Access Authentication objects MIB (19230_JNPR_00030578);
- Juniper enterprise VPN Certificat objects MIB (19230_JNPR_00030582);
- Juniper Enterprise specific MIB: DNS MIB (19230_JNPR_00030585);
- Juniper Enterprise specific MIB: IDP MIB (19230_JNPR_00030587);
- Juniper enterprise security Interface extension objects MIB (19230_JNPR_00030588);
- Juniper Networks Security IPSEC VPN object mibs (19230_JNPR_00030594);
- Juniper enterprise Network Address Translation (NAT) MIB (19230_JNPR_00030597);
- Juniper enterprise specific policy objects MIB (19230_JNPR_00030603);
- Juniper enterprise security screening objects MIB (19230_JNPR_00030610);
- Juniper Enterprise Specific MIB: Structure of Management Information (19230_JNPR_00030628);
- Juniper Enterprise specific MIB: SRX5000 SPU monitoring MIB (19230_JNPR_00030629);
- Juniper enterprise security UTM MIB (19230_JNPR_00030632);
- Juniper enterprise specific Chassis Cluster objects MIB (19230_JNPR_00030636);
- Juniper Enterprise Specific MIB: Structure of Management Information (19230_JNPR_00030638);
- JUNIPER−L2CP−FEATURES−MIB DEFINITIONS (19230_JNPR_00030641);
- JNX−L2TP−MIB DEFINITIONS (19230_JNPR_00030645);
- Juniper Enterprise Specific MIB:LDP MIB Extension (19230_JNPR_00030672);
- JUNIPER−MIMSTP−MIB DEFINITIONS (19230_JNPR_00030678);
- JUNIPER−MPLS−LDP−MIB DEFINITIONS (19230_JNPR_00030717);
- Juniper Multi−Protocol Label Switched Paths MIB (19230_JNPR_00030752);

- Juniper OSPFV3−MIB−JUNIPER DEFINITIONS (19230_JNPR_00030767);
- Juniper Enterprise Specific MIB: OTN Interface Management MIB Extension (19230_JNPR_00030816);
- Juniper Enterprise Specific MIB: PAE MIB Extension (19230_JNPR_00030833);
- Juniper Enterprise Specific MIB: PFE MIB (19230_JNPR_00030839);
- Juniper Enterprise Specific MIB: Ping MIB (19230_JNPR_00030846);
- Juniper Enterprise Specific MIB: Passive Monitoring MIB (19230_JNPR_00030869);
- JUNIPER−PW−TDM−MIB DEFINITIONS (19230_JNPR_00030877);
- Juniper Enterprise Specific MIB: RMON MIB Extension (19230_JNPR_00030900);
- Juniper Enterprise Specific MIB: Reverse Path Forwarding MIB (19230_JNPR_00030903);
- Juniper Enterprise Specific MIB: RPM MIB (19230_JNPR_00030905);
- Juniper Enterprise Specific MIB: RPS (19230_JNPR_00030917);
- Juniper Enterprise specific MIB: RSVP MIB (19230_JNPR_00030923);
- Juniper Enterprise Specific MIB: Source Class Usage MIB (19230_JNPR_00030927);
- Juniper JUNIPER−SECURE−ACCESS−PORT−MIB DEFINITIONS (19230_JNPR_00030930);
- Juniper Enterprise Specific MIB: Structure of Management Information (19230_JNPR_00030937);
- Juniper Enterprise Specific MIB: Sonet Interface Management MIB Extension (19230_JNPR_00030941);
- Juniper APS−MIB DEFINITIONS (19230_JNPR_00030945);
- Juniper Enterprise Specific MIB: Service PIC MIB (19230_JNPR_00030973);
- Juniper enterprise specific syslog MIB (19230_JNPR_00030981);
- Juniper Enterprise Specific MIB: TraceRoute MIB (19230_JNPR_00030986);
- Juniper User AAA objects MIB (19230_JNPR_00030989);
- Juniper Enterprise Specific MIB: Utility MIB (19230_JNPR_00030993);
- Juniper Enterprise Specific MIB: Virtual Chassis (19230_JNPR_00030999);
- Juniper Enterprise Specific MIB: Vlan (19230_JNPR_00031003);
- JUNIPER−VPN−MIB DEFINITIONS (19230_JNPR_00031013);
- Juniper JUNOS® Software Network Interfaces Configuration Guide Release 9.5 (19230_JNPR_00031031);
- Juniper JUNOS® Software Policy Framework Configuration Guide Release 9.5 (19230_JNPR_00032419);
- Juniper JUNOS® Software Routing Protocols Configuration Guide Release 9.5 (19230_JNPR_00032931);
- Juniper JUNOS® Software Services Interfaces Configuration Guide Release 9.5 (19230_JNPR_00033801);
- Juniper JUNOS® Software Subscriber Access Configuration Guide Release 9.5 (19230_JNPR_00035011);
- Juniper JUNOS® Software System Basics Configuration Guide Release 9.5 (19230_JNPR_00035635);

- Juniper JUNOS® Software VPNs Configuration Guide Release 9.5 (19230_JNPR_00036491);
- Juniper JUNOS™ Internet Software Security Configuration Guide for Common Criteria and JUNOS-FIPS Release 8.1 (19230_JNPR_00037119);
- Juniper JUNOS® Software Hierarchy and RFC Reference Release 9.5 (19230_JNPR_00037225);
- Juniper JUNOS® Software System Log Messages Reference Release 9.5 (19230_JNPR_00037525);
- Juniper LSP Hierarchy support for MPLS LSPs Functional Specification v. 1.9 (19230_JNPR_00038137);
- Juniper Monitor Label-switched-path functional specification (19230_JNPR_00038164);
- Juniper Monitor label-switched-path unit test plan (19230_JNPR_00038168);
- Juniper sw-projects/mpls/fate-sharing/project-plan.txt Revision 1.1 (19230_JNPR_00041939);
- Juniper sw-projects/mpls/srlg/srlg (19230_JNPR_00041940);
- Juniper IGP SRLG Distribution (19230_JNPR_00041941);
- Juniper sw-projects/mpls/srlg/srlg-fs.txt (19230_JNPR_00041954);
- Juniper Implement SRLG support in JUNOS Functional Specification (19230_JNPR_00041967);
- Juniper sw-projects/mpls/srlg/srlg-mpls-fs.txt (19230_JNPR_00041981);
- Juniper sw-projects/mpls/srlg/srlg-mpls-utp.txt (19230_JNPR_00041995);
- Juniper Scoping for SRLG support in JUNOS v. 1.3 (19230_JNPR_00042011);
- Juniper sw-projects/mpls/srlg/srlg-scoping.txt (19230_JNPR_00042015);
- Juniper sw-projects/mpls/srlg/srlg-test1.cli (19230_JNPR_00042019);
- Juniper sw-projects/mpls/srlg/srlg-test2.cli (19230_JNPR_00042026);
- Juniper sw-projects/mpls/srlg/srlg-test3.cli (19230_JNPR_00042029);
- Juniper sw-projects/mpls/srlg/srlg-test4.cli (19230_JNPR_00042042);
- Juniper sw-projects/mpls/srlg/srlg.cfg (19230_JNPR_00042051);
- Juniper sw-projects/mpls/srlg/srlg.lab (19230_JNPR_00042053);
- Juniper sw-projects/mpls/srlg/srlg_utp.txt v. 1.2 (19230_JNPR_00042054);
- Juniper Tunneling LDP LSPs in RSVP LSPs Functional Specification v 1.2 (19230_JNPR_00042060);
- Juniper Tunneling LDP LSPs in RSVP LSPs Unit Test Plan v. 1.1 (19230_JNPR_00042064);
- Juniper Tunneling LDP Hierarchy support for MPLS LSPs Unit Test Plan v 1.6 (19230_JNPR_00042070); and
- Juniper Routers Sales (19230_JNPR_00042088_Attorney_Eyes_Only_Confidential).

Motivations to combine the disclosures in Juniper Junos OS Guide with disclosures in other publications known in the art, as explained in this chart, include at least the similarity in subject matter between the references to the extent they concern methods of selecting paths

in a network and/or determining the cost of potential paths in a network. Insofar as the references cite other patents or publications, or suggest additional changes, one of ordinary skill in the art would look beyond a single reference to other references in the field.

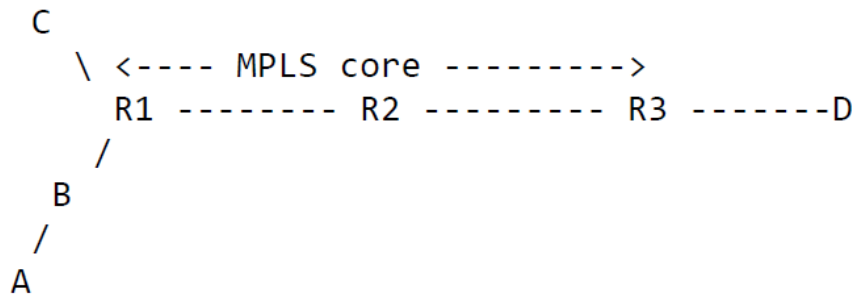
These invalidity contentions are based on Defendant's present understanding of the asserted claims, and Orckit's apparent construction of the claims in its November 3, 2022 Disclosure of Asserted Claims and Infringement Contentions Pursuant to P.R. 3-1, and Orckit's January 19, 2023 First Amended Disclosure of Asserted Claims and Infringement Contentions Pursuant to P.R. 3-1 (Orckit's "Infringement Disclosures"), which is deficient at least insofar as it fails to cite any documents or identify accused structures, acts, or materials in the Accused Products with particularity. Defendant does not agree with Orckit's application of the claims, or that the claims satisfy the requirements of 35 U.S.C. § 112. Defendant's contentions herein are not, and should in no way be seen as, admissions or adoptions as to any particular claim scope or construction, or as any admission that any particular element is met by any accused product in any particular way. Defendant objects to any attempt to imply claim construction from this chart. Defendant's prior art invalidity contentions are made in a variety of alternatives and do not represent Defendant's agreement or view as to the meaning, definiteness, written description support for, or enablement of any claim contained therein.

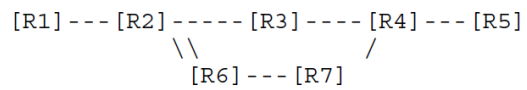
The following contentions are subject to revision and amendment pursuant to Federal Rule of Civil Procedure 26(e), the Local Rules, and the Orders of record in this matter subject to further investigation and discovery regarding the prior art and the Court's construction of the claims at issue.

No.	**'821 Patent Claim 1**	**Juniper Junos OS System**
1[preamble]	An entity selection method performed by a network device, comprising the steps of:	Juniper Junos OS System discloses an entity selection method performed by a network device, comprising the steps of. *See generally* 19230_JNPR_00000071-00042088. Juniper Networks Junos OS system was implemented and sold in a number of routing devices. *See generally* 19230_JNPR_00042088_Attorney_Eyes_Only_Confidential. For example, Juniper Junos OS Guide discloses that the Junos IOS System can run on M Series, MX Series, T Series, EX Series, or J Series Juniper Networks Routers. Juniper Junos OS Guide further discloses that these routers implement traffic engineering with MPLS and RSVP having functional component (a path selection component) through which the ingress router determines the physical path for each LSP by applying a Constrained Shortest Path First (CSPF) algorithm. For example, Juniper Networks describes using its RSVP functionality to establish LDP tunneling for LSPs in several documents. *See, e.g.*, 19230_JNPR_00042060; 19230_JNPR_00042064.

No.	'821 Patent Claim 1	Juniper Junos OS System
		Juniper Junos OS Guide at Page No. 32 ("[This guide provides an overview of the MPLS applications functions of the Junos OS and describes how to configure MPLS applications on the router.]") Juniper Junos OS Guide at Page No. 32 ("[This guide is designed for network administrators who are configuring and monitoring a Juniper Networks M Series, MX Series, T Series, EX Series, or J Series router or switch.]") Juniper Junos OS Guide at Page No. 42 ("[Components of Traffic Engineering In the Junos OS, traffic engineering is implemented with MPLS and RSVP. Traffic engineering is composed of four functional components: • Packet Forwarding Component on page 4 • Information Distribution Component on page 5 • Path Selection Component on page 6 • Signaling Component on page 7.]") Juniper Junos OS Guide at Page No. 43 ("[Information Distribution Component Traffic engineering requires detailed knowledge about the network topology as well as dynamic information about network loading. To implement the information distribution component, simple extensions to the IGPs are defined. Link attributes are included as part of each router's link-state advertisement. IS-IS extensions include the definition of new type length values (TLVs), whereas OSPF extensions are implemented with opaque link-state advertisements (LSAs). The standard flooding algorithm used by the link-state IGPs ensures that link attributes are distributed to all routers in the routing domain. Some of the traffic engineering extensions to be added to the IGP link-state advertisement include maximum link bandwidth, maximum reserved link bandwidth, current bandwidth reservation, and link coloring.]") Juniper Junos OS Guide at Page No. 43 ("[Each router maintains network link attributes and topology information in a specialized traffic engineering database. The traffic engineering database is used exclusively for calculating explicit paths for the placement of LSPs across the physical

No.	'821 Patent Claim 1	Juniper Junos OS System
		topology. A separate database is maintained so that the subsequent traffic engineering computation is independent of the IGP and the IGP's link-state database. Meanwhile, the IGP continues its operation without modification, performing the traditional shortest-path calculation based on information contained in the router's link-state database.]") Juniper Junos OS Guide at Page No. 44 ("[Path Selection Component After network link attributes and topology information are flooded by the IGP and placed in the traffic engineering database, each ingress router uses the traffic engineering database to calculate the paths for its own set of LSPs across the routing domain. The path for each LSP can be represented by either a strict or loose explicit route. An explicit route is a preconfigured sequence of routers that should be part of the physical path of the LSP. If the ingress router specifies all the routers in the LSP, the LSP is said to be identified by a strict explicit route. If the ingress router specifies only some of the routers in the LSP, the LSP is described as a loose explicit route. Support for strict and loose explicit routes allows the path selection process to be given broad latitude whenever possible, but to be constrained when necessary. The ingress router determines the physical path for each LSP by applying a Constrained Shortest Path First (CSPF) algorithm to the information in the traffic engineering database. CSPF is a shortest-path-first algorithm that has been modified to take into account specific restrictions when the shortest path across the network is calculated. Input into the CSPF algorithm includes: • Topology link-state information learned from the IGP and maintained in the traffic engineering database • Attributes associated with the state of network resources (such as total link bandwidth, reserved link bandwidth, available link bandwidth, and link color) that are carried by IGP extensions and stored in the traffic engineering database • Administrative attributes required to support traffic traversing the proposed LSP (such as bandwidth requirements, maximum hop count, and administrative policy requirements) that are obtained from user configuration.]") Juniper Networks, Inc. LSP Hierarchy support for MPLS LSPs discloses: "The goal of this document is to summarize the design aspects of LSP hierarchy feature." 19230_JNPR_00000071.

No.	'821 Patent Claim 1	Juniper Junos OS System
		"This feature adds support for LSP Hierarchy for MPLS RSVP LSPs. This will facilitate tunneling (nesting) one or more RSVP LSPs into another RSVP LSP. This will not only help scenarios where scaling of full-mesh of RSVP LSPs is an issue, but will also act as a building block to enable other applications such as inter-domain TE LSPs or tunneling of native RSVP aggregate flows over RSVP-TE LSP (for VOIP), in the future." 19230_JNPR_00000071. "The main goal of introducing the above switching capabilities in [LSP-HIERARCHY] for packet networks was to handle the case of dynamic FA setup. By assigning different switching capabilities to the links in the network, an operator could enforce a policy whereby he could control which LSPs would require dynamic FA LSP setup and which would not. Let us take an example." 19230_JNPR_00000072.  19230_JNPR_00000072. "In this example if an operator wishes to have dynamic FA LSP setup in the MPLS core and if he requires a policy that would: trigger FA LSP setup if an LSP request for A-D (A-B-R1-R2-R3) is received on R1 and NOT trigger an FA LSP if an LSP setup request for A-C (A-B-R1-C) is received then [LSP-HIERARCHY] provides one way to achieve this policy. By assigning switching capabilities to links, in the above example, such that all links in the MPLS core have switching capability PSC-2 and all other links have switching capability PSC-1 the operator has basically

No.	'821 Patent Claim 1	Juniper Junos OS System
		divided his network into regions based on link switching capability and therefore can dictate the trigger for dynamic FA based on the criteria that the LSP crosses different regions. The FA LSP itself in this example would be signaled with switching type of PSC-2 and would be advertised as a TE-link with switching capability of PSC-1." 19230_JNPR_00000072. Juniper sw-projects/mpls/misc/fast-reroute/draft-ietf-mpls-rsvp-lsp-fastreroute-00.nroff: "For each LSP which is backed up, another backup LSP is established." 19230_JNPR_00000076.  For example, suppose that in the simple topology above, R1 creates a tunnel to R5 via the path [R1->R2->R3->R4->R5]. R2 can provide user traffic protection by creating a partial backup tunnel [R2->R6->R7->R4] which merges with the original tunnel [R1->R2->R3->R4->R5] at R4. We refer a partial one-to-one backup tunnel [R2->R6->R7->R4] as a detour. To fully protect a LSP that traverses through N nodes, there could be as many as (N - 1) detours. To minimize processing overhead, it is desirable to merge detours back to a main LSP wherever possible. 19230_JNPR_00000076. "Level3 is intensely interested in making sure that, in the event of fiber cuts, the absolutely minimum amount of data is lost and all voice calls (that is right, not our typical Internet traffic) over their backbone network go uninterrupted. They have a maximum threshold of 2 second recovery time. This is achieved by pre-computing backup paths for all their existing primary paths. The backup paths might be done only at ingress (2 times redundancy), or at every hop of existing path (N times redundancy).

No.	'821 Patent Claim 1	Juniper Junos OS System
		In order for the backup path to be worth while, they must not share links/physical fiber paths with primary path as much as possible. This is to ensure that a single point of failure does not bring down both primary and backup paths at the same time." Juniper Networks, Inc. MPLS Fate-sharing (cut-set) file Functional Specification (19230_JNPR_00000102) at 1; Juniper Networks, Inc. MPLS Fate-sharing (cut-set) file Functional Specification v. 1.3 (19230_JNPR_00000110) at 1. Juniper Networks, Inc. Advanced Insight Solutions User Guide Release 1.3 (19230_JNPR_00000121) at 31-32. **Table 5: JUNOS Software Network Operations Guides** **Book** — **Description** *Baseline* — Describes the most basic tasks for running a network using Juniper Networks products. Tasks include upgrading and reinstalling JUNOS software, gathering basic system management information, verifying your network topology, and searching log messages. *Interfaces* — Describes tasks for monitoring interfaces. Tasks include using loopback testing and locating alarms. *MPLS* — Describes tasks for configuring, monitoring, and troubleshooting an example MPLS network. Tasks include verifying the correct configuration of the MPLS and RSVP protocols, displaying the status and statistics of MPLS running on all routing platforms in the network, and using the layered MPLS troubleshooting model to investigate problems with an MPLS network.

No.	'821 Patent Claim 1	Juniper Junos OS System

Table 5: JUNOS Software Network Operations Guides *(continued)*

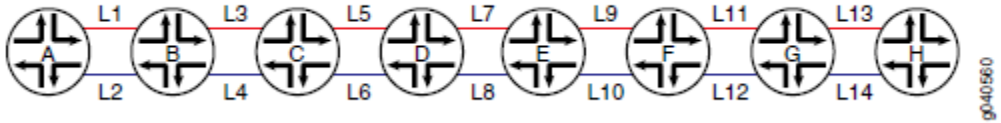
Book	Description
MPLS Log Reference	Describes MPLS status and error messages that appear in the output of the **show mpls lsp extensive** command. The guide also describes how and when to configure Constrained Shortest Path First (CSPF) and RSVP trace options, and how to examine a CSPF or RSVP failure in a sample network.
MPLS Fast Reroute	Describes operational information helpful in monitoring and troubleshooting an MPLS network configured with fast reroute (FRR) and load balancing.
Hardware	Describes tasks for monitoring M-series and T-series routing platforms.

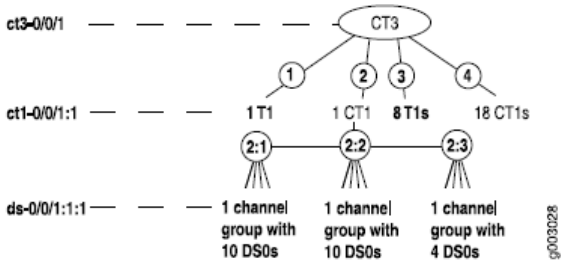
Juniper Junos OS provides network monitoring for label switched paths. *See, e.g.*, 19230_JNPR_00038164; 19230_JNPR_00038168.

Generally, Juniper's documents show how to implement its Junos OS features and provide support for the MPLS features discussed herein. *See, e.g.*, 19230_JNPR_00000415; 19230_JNPR_00000599; 19230_JNPR_00000655; 19230_JNPR_00001433; 19230_JNPR_00002237; 19230_JNPR_00002357; 19230_JNPR_00002612; 19230_JNPR_00002892; 19230_JNPR_00002984; 19230_JNPR_00003034; 19230_JNPR_00003522; 19230_JNPR_00004722; 19230_JNPR_00005688; 19230_JNPR_00006734; 19230_JNPR_00018044; 19230_JNPR_00020416; 19230_JNPR_00021220; 19230_JNPR_00021502; 19230_JNPR_00021632; 19230_JNPR_00021882; 19230_JNPR_00022034; 19230_JNPR_00022208; 19230_JNPR_00022588; 19230_JNPR_00023308; 19230_JNPR_00024414; 19230_JNPR_00024708; 19230_JNPR_00024772; 19230_JNPR_00024818; 19230_JNPR_00025926; 19230_JNPR_00026740; 19230_JNPR_00027476; 19230_JNPR_00027568; 19230_JNPR_00028066; 19230_JNPR_00028362; 19230_JNPR_00028906; 19230_JNPR_00028998; 19230_JNPR_00029168;

No.	'821 Patent Claim 1	Juniper Junos OS System
		19230_JNPR_00029718; 19230_JNPR_00031031; 19230_JNPR_00032419; 19230_JNPR_00032931; 19230_JNPR_00033801; 19230_JNPR_00035011; 19230_JNPR_00035635; 19230_JNPR_00036491; 19230_JNPR_00037119; 19230_JNPR_00037225; 19230_JNPR_00037525.
1[a]	providing a plurality of multi protocol label switching (MPLS) transport entities between a first endpoint and a second endpoint;	Juniper Junos OS System discloses providing a plurality of multi protocol label switching (MPLS) transport entities between a first endpoint and a second endpoint. For example, Juniper Junos OS Guide discloses that the Juniper Junos OS System can specify any number of LSPs as backups for a primary LSP between the ingress router and egress router. Juniper Junos OS Guide further discloses that these LSPs can be in the MPLS protocol. For example, Juniper's documents show that its Junos OS provides support for MPLS. *See, e.g.*, 19230_JNPR_00000102; 19230_JNPR_00000110; 19230_JNPR_00023694; 19230_JNPR_00023950); 19230_JNPR_00024170; 19230_JNPR_00029168; 19230_JNPR_00038137; 19230_JNPR_00042070. Further, Juniper's documents show that its Junos OS uses a management information base (MIB), which is a Simple Network Management (SNMP) database to implement its MPLS functionality. *See, e.g.*, 19230_JNPR_00030246; 19230_JNPR_00030253; 19230_JNPR_00030264; 19230_JNPR_00030280; 19230_JNPR_00030294; 19230_JNPR_00030297; 19230_JNPR_00030364; 19230_JNPR_00030369; 19230_JNPR_00030402; 19230_JNPR_00030404; 19230_JNPR_00030432; 19230_JNPR_00030447; 19230_JNPR_00030484; 19230_JNPR_00030488; 19230_JNPR_00030502; 19230_JNPR_00030505; 19230_JNPR_00030506; 19230_JNPR_00030508; 19230_JNPR_00030514; 19230_JNPR_00030523; 19230_JNPR_00030525; 19230_JNPR_00030542; 19230_JNPR_00030559; 19230_JNPR_00030562; 19230_JNPR_00030578; 19230_JNPR_00030582; 19230_JNPR_00030585; 19230_JNPR_00030587; 19230_JNPR_00030588; 19230_JNPR_00030594; 19230_JNPR_00030597; 19230_JNPR_00030603; 19230_JNPR_00030610; 19230_JNPR_00030628; 19230_JNPR_00030629; 19230_JNPR_00030632; 19230_JNPR_00030636; 19230_JNPR_00030638; 19230_JNPR_00030641;

No.	'821 Patent Claim 1	Juniper Junos OS System
		19230_JNPR_00030645; 19230_JNPR_00030672; 19230_JNPR_00030678; 19230_JNPR_00030717; 19230_JNPR_00030752; 19230_JNPR_00030767; 19230_JNPR_00030816; 19230_JNPR_00030833; 19230_JNPR_00030839; 19230_JNPR_00030846; 19230_JNPR_00030869; 19230_JNPR_00030877; 19230_JNPR_00030900; 19230_JNPR_00030903; 19230_JNPR_00030905; 19230_JNPR_00030917; 19230_JNPR_00030923; 19230_JNPR_00030927; 19230_JNPR_00030930; 19230_JNPR_00030937; 19230_JNPR_00030941; 19230_JNPR_00030945; 19230_JNPR_00030973; 19230_JNPR_00030981; 19230_JNPR_00030986; 19230_JNPR_00030989; 19230_JNPR_00030993; 19230_JNPR_00030999; 19230_JNPR_00031003; 19230_JNPR_00031013. Juniper Junos OS Guide at Page No. 46 ("[For example, consider a topology that includes two east-west paths across the United States: one in the north through Chicago and one in the south through Dallas. If you want to establish an LSP between a router in New York and one in San Francisco, you can configure the partial path for the LSP to include a single loose-routed hop of a router in Dallas. The result is an LSP routed along the southern path. The ingress router uses CSPF to compute the complete path and RSVP to install the forwarding state along the LSP. • You can configure the ingress router with no constraints whatsoever. In this case, normal IGP shortest-path routing is used to determine the path of the LSP. This configuration does not provide any value in terms of traffic engineering. However, it is easy and might be useful in situations when services such as virtual private networks (VPNs) are needed. In all these cases, you can specify any number of LSPs as backups for the primary LSP, thus allowing you to combine more than one configuration approach. For example, you might explicitly compute the primary path offline, set the secondary path to be constraint-based, and have the tertiary path be unconstrained. If a circuit on which the primary LSP is routed fails, the ingress router notices the outage from error notifications received from a downstream router or by the expiration of RSVP soft-state information. Then the router dynamically forwards traffic to a hot-standby LSP or calls on RSVP to create a forwarding state for a new backup LSP.]") Juniper Junos OS Guide at Page No. 83 ("[MPLS and Traffic Protection

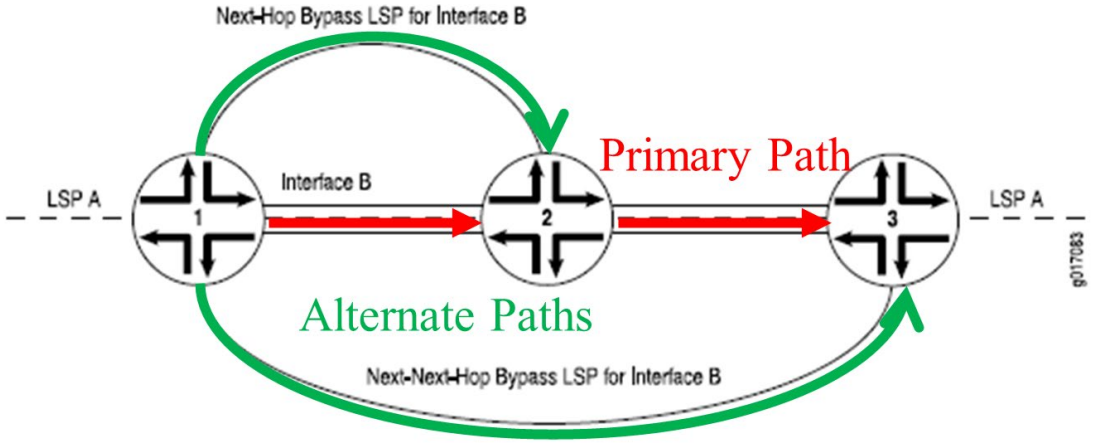
No.	'821 Patent Claim 1	Juniper Junos OS System
		Typically, when an LSP fails, the router immediately upstream from the failure signals the outage to the ingress router. The ingress router calculates a new path to the egress router, establishes the new LSP, and then directs the traffic from the failed path to the new path. This rerouting process can be time-consuming and prone to failure. For example, the outage signals to the ingress router might get lost, or the new path might take too long to come up, resulting in significant packet drops. The Junos OS provides several complementary mechanisms for protecting against LSP failures: • Standby secondary paths—You can configure primary and secondary paths. You configure secondary paths with the standby statement. To activate traffic protection, you need to configure these standby paths only on the ingress router. If the primary path fails, the ingress router immediately reroutes traffic from the failed path to the standby path, thereby eliminating the need to calculate a new route and signal a new path. For information about configuring standby LSPs, see "Configuring Hot Standby of Secondary Paths" on page 113. • Fast reroute—You configure fast reroute on an LSP to minimize the effect of a failure in the LSP. Fast reroute enables a router upstream from the failure to route around the failure quickly to the router downstream of the failure. The upstream router then signals the outage to the ingress router, thereby maintaining connectivity before a new LSP is established. For a detailed overview of fast reroute, see "Fast Reroute Overview" on page 46. For information about configuring fast reroute, see "Configuring Fast Reroute" on page 82.]") Figure 20: least-fill Load Balancing Algorithm Example  Fig. 20 Juniper Junos OS Guide at Page No. 148 ("As shown in Figure 20 on page 110, there are two potential paths for an LSP to traverse from router A to router H, the odd links from L1 through L13 and the even links from L2 through L14. Currently, the router is using the even links as the active path for the LSP.")

No.	'821 Patent Claim 1	Juniper Junos OS System
		"2.0 Operation The user indicates through configuration the RSVP LSPs that they want to be used for tunneling LDP LSPs. As a consequence, LDP will initiate extended discovery to set up adjacencies to the endpoints of these RSVP LSPs and then set up remote sessions with them. The remote LDP sessions look just like regular sessions and we will learn from and advertise labels to these sessions. If there are RSVP LSPs in both directions between a pair of routers then the LDP tunneling configuration on each LSP is independent of that on the other i.e. it is ok to have simplex tunneling. In such a scenario, labels will get advertised over the session in both directions but can become active in a given direction only if LDP tunneling is enabled on the corresponding LSP." Juniper Networks, Inc. Tunneling LDP LSPs in RSVP LSPs Functional Specification (19230_JNPR_00000105) at 1-2. Junos® Channelized IQ Interfaces Feature Guide Release 9.5 (19230_JNPR_00000499) at 59: **Example: Channelized DS3 IQ Interface Configuration** **Figure 6: Channelized DS3 IQ Interface Example** 

No.	'821 Patent Claim 1	Juniper Junos OS System
		MPLS To view information about MPLS label-switched paths (LSPs) and virtual private networks (VPNs), select **Monitor > MPLS**. Table 14 on page 31 shows a summary of the information displayed on MPLS pages and the corresponding CLI **show** commands you can enter at the J-Web CLI terminal. **Table 14: MPLS Information and Corresponding CLI show Commands** **Information Displayed** | **Corresponding CLI Command** **Interfaces** Interfaces on which MPLS is enabled, plus the operational state and any administrative groups applied to an interface. | show mpls interface **LSP Information** 19230_JNPR_00021552.
1[b]	determining an overall cost for each entity pair of said plurality of entities;	Juniper Junos OS System discloses for determining an overall cost for each entity pair of said plurality of entities. For example, Juniper Junos OS Guide discloses "Fate Sharing" where a Juniper router computes a primary and secondary path using CSPF, CSPF identifies all fate-sharing groups that are associated with the primary path by identifying all links that the primary path transverses and checks each link in the traffic engineering database against the compiled group list. Juniper Junos OS Guide further discloses that if the link is a member of a group, the cost of the link is increased by the cost of the group and higher the cost value. Thus, at least under the apparent claim scope alleged by Orckit's Infringement Disclosures, this limitation is met. *See, e.g.*, 19230_JNPR_00041939. Juniper Junos OS Guide at Page No. 73 ("[Fate Sharing Fate sharing allows you to create a database of information that CSPF uses to compute one or more backup paths to use in case the primary path becomes unstable. The database describes the relationships between elements of the network, such as routers and links. You can specify one or more elements within a group.

No.	'821 Patent Claim 1	Juniper Junos OS System
		Through fate sharing, you can configure backup paths that minimize the number of shared links and fiber paths with the primary paths as much as possible, to ensure that if a fiber is cut, the minimum amount of data is lost and a path still exists to the destination. For a backup path to work optimally, it must not share links or physical fiber paths with the primary path, ensuring that a single point of failure will not affect the primary and backup paths simultaneously. For more information about fate sharing, see the Junos OS Routing Protocols Configuration Guide.]") Juniper Junos OS Guide at Page No. 94 ("[Configuring the Ingress Router for MPLS-Signaled LSPs MPLS-signaled label-switched paths (LSPs) run from a specific ingress router to a specific egress router. For basic MPLS-signaled LSP function, you must configure the ingress router, but do not have to configure any other routers. To configure signaled LSPs, perform the following tasks on the ingress router: • Creating Named Paths on page 56 • Configuring Alternate Backup Paths Using Fate Sharing on page 58.]") Juniper Junos OS Guide at Page No. 96 ("[Configuring Alternate Backup Paths Using Fate Sharing You can create a database of information that Constrained Shortest Path First (CSPF) uses to compute one or more backup paths in case the primary path becomes unstable. The database describes the relationships between elements of the network, such as routers and links. Because these network elements share the same fate, this relationship is called fate sharing. You can configure backup paths that minimize the number of shared links and fiber paths with the primary paths as much as possible to ensure that, if a fiber is cut, the minimum amount of data is lost and a path still exists to the destination. For a backup path to work optimally, it must not share links or physical fiber paths with the primary path. This ensures that a single point of failure will not affect the primary and backup paths at the same time.]") Juniper Junos OS Guide at Page No. 97 ("[Fate-sharing groups contain three types of objects:

No.	'821 Patent Claim 1	Juniper Junos OS System
		• Point-to-point links—Identified by the IP addresses at each end of the link. Unnumbered point-to-point links are typically identified by borrowing IP addresses from other interfaces. Order is not important; from 1.2.3.4 to 1.2.3.5 and from 1.2.3.5 to 1.2.3.4 have the same meaning. • Non-point-to-point links—Include links on a LAN interface (such as Gigabit Ethernet interfaces) or nonbroadcast multiaccess (NBMA) interfaces (such as Asynchronous Transfer Mode [ATM] or Frame Relay). You identify these links by their individual interface address. For example, if the LAN interface 192.168.200.0/24 has four routers attached to it, each router link is individually identified: from 192.168.200.1; # LAN interface of router 1 from 192.168.200.2; # LAN interface of router 2 from 192.168.200.3; # LAN interface of router 3 from 192.168.200.4; # LAN interface of router 4 You can list the addresses in any order. • A router node—Identified by its configured router ID. An object can be in any number of groups, and a group can contain any number of objects. Each group has a configurable cost attributed to it, which represents the level of impact this group has on CSPF computations. The higher the cost, the less likely a backup path will share with the primary path any objects in the group. The cost is directly comparable to traffic engineering metrics. By default, the cost is 1. Changing the fate-sharing database does not affect established LSPs until the next reoptimization of CSPF. The fate-sharing database does influence fast-reroute computations.]") Juniper Junos OS Guide at Page No. 97 ("[When CSPF computes the primary paths of an LSP (or secondary paths when the primary path is not active), it ignores the fate-sharing information. You always want to find the best possible path (least IGP cost) for the primary path. When CSPF computes a secondary path while the primary path (of the same LSP) is active, the following occurs: 1. CSPF identifies all fate-sharing groups that are associated with the primary path. CSPF does this by identifying all links and nodes that the primary path traverses and compiling group lists that contain at least one of the links or nodes. CSPF ignores the ingress and egress nodes in the search.

No.	'821 Patent Claim 1	Juniper Junos OS System
		2. CSPF checks each link in the traffic engineering database against the compiled group list. If the link is a member of a group, the cost of the link is increased by the cost of the group. If a link is a member of multiple groups, all group costs are added together. 3. CSPF performs the check for every node in the traffic engineering database, except the ingress and egress node. Again, a node can belong to multiple groups, so costs are additive. 4. The router performs regular CSPF computation with the adjusted topology.]") Juniper Junos OS Guide at Page No. 330 ("In Figure 23 on page 292, both node protection and link protection are enabled on Interface B on Router 1. Both node protection and link protection are also enabled on LSP A, an LSP that traverses the link transiting Router 1, Router 2, and Router 3. If the link between Router 1 and Router 2 fails, traffic from LSP A is quickly switched to the next-hop bypass LSP generated by link protection. If Router 2 suffers a hardware or software failure, traffic from LSP A is switched to the next-next-hop bypass LSP generated by node protection.") Figure 23: Node Protection Creating a Next-Next-Hop Bypass LSP  Fig. 23 (annotation added)

No.	'821 Patent Claim 1	Juniper Junos OS System
		"Each group has a configurable 'cost' attribute to it - this represents the level of impact this group should have on CSPF computations. The higher the cost, the less likely a backup path will share (with primary path) any objects in the group. By default, the cost is 1. The 'cost' is directly comparable to TE metrics." Juniper Networks, Inc. MPLS Fate-sharing (cut-set) file Functional Specification (19230_JNPR_00000102) at 2. "When doing CSPF computation for primary paths of an LSP, (or secondary paths when the primary path is not active), the fate-sharing information is ignored. This is because we always want to find the best possible (least IGP cost) path for primary path. When doing CSPF computation for secondary (backup) path while primary path (of the same LSP) is active: First, all fate-sharing groups that primary path is associated are identified. This is done by identifying all links and all nodes that primary traverse through, and compile groups list that contain at least one of the link/node. The ingress and egress nodes are ignored in the search. Second, every link in TED is checked against the compiled group list. If the link is member of a group, the cost of the link is increased by the cost of the group. A link may be members of multiple groups, and all group costs are added together. Third, perform the same step for every node in TED, except the ingress and egress node. Again, a node may belong to multiple groups so costs are additive. Last, perform regular CSPF computation with adjusted topology." Juniper Networks, Inc. MPLS Fate-sharing (cut-set) file Functional Specification (19230_JNPR_00000102) at 3. Junos® Software Network Management Configuration Guide Release 9.5 (19230_JNPR_00000655) at 563-68: **jnxMIDot1sJuniperMstTable** The jnxMIDot1sJuniperMstTable provides MSTP module parameters for a given virtual context. Table 159 on page 564 lists the objects in the jnxMIDot1sJuniperMstTable.

Table 159: jnxMIDot1sJuniperMstTable

Object	Object Identifier	Description
jnxMIDot1sJuniperMstContextId	jnxMIDot1sJuniperMstEntry 1	Identifies the virtual context.
jnxMIMstSystemControl	jnxMIDot1sJuniperMstEntry 2	Indicates the status of MST on the ports of the device. The value start (1) indicates that MIMST is active on all ports of the device; the value shutdown (2) indicates that MIMST is shut down on all ports of the device.
jnxMIMstModuleStatus	jnxMIDot1sJuniperMstEntry 3	Indicates whether MST is enabled or disabled on the ports. When set to value 1, enabled, MST is enabled on all ports of the device; when set to value 2, disabled, MIMST is disabled on all ports. **NOTE:** The object can be set to enabled(1) only when jnxMIMstSystemControl is set to start.
jnxMIMstMaxMstInstanceNumber	jnxMIDot1sJuniperMstEntry 4	Indicates the maximum number of spanning-tree instances allowed on the bridge.
jnxMIMstNoOfMstiSupported	jnxMIDot1sJuniperMstEntry 5	Indicates the maximum number of spanning-tree instances that can be supported on the bridge.
jnxMIMstMaxHopCount	jnxMIDot1sJuniperMstEntry 6	Returns the Maximum Hop Count value.
jnxMIMstBrgAddress	jnxMIDot1sJuniperMstEntry 7	Indicates the MAC address used by the bridge, which forms a unique bridge identifier when combined with jnxMIMstCistBridgePriority or jnxMIMstMstiBridgePriority.
jnxMIMstCistRoot	jnxMIDot1sJuniperMstEntry 8	Indicates the bridge identifier of the root of the common spanning tree as determined by the Spanning Tree Protocol. This value is used as the CIST Root Identifier parameter in all configuration bridge PDUs originating at this node.
jnxMIMstCistRegionalRoot	jnxMIDot1sJuniperMstEntry 9	Indicates the bridge identifier of the root of the multiple spanning-tree region as determined by the Spanning Tree Protocol. This value is used as the CIST Regional Root Identifier parameter in all configuration bridge PDUs originating at this node.
jnxMIMstCistRootCost	jnxMIDot1sJuniperMstEntry 10	Indicates the cost of the path to the CIST root from this bridge.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 361 of 698

Table 159: jnxMIDot1sJuniperMstTable *(continued)*

Object	Object Identifier	Description
jnxMIMstCistRegionalRootCost	jnxMIDot1sJuniperMstEntry 11	Indicates the cost of the path to the CIST regional root from this bridge.
jnxMIMstCistRootPort	jnxMIDot1sJuniperMstEntry 12	Indicates the port number of the port that offers the lowest path cost from this bridge to the CIST Root Bridge.
jnxMIMstCistBridgePriority	jnxMIDot1sJuniperMstEntry 13	Indicates the value of the writable portion of the bridge identifier. The values set for Bridge Priority must be in multiples of 4096.
jnxMIMstCistBridgeMaxAge	jnxMIDot1sJuniperMstEntry 14s	Indicates the value that a bridge uses for MaxAge when the bridge is acting as the root.
jnxMIMstCistBridgeForwardDelay	jnxMIDot1sJuniperMstEntry 15	Indicates the value that the bridge uses for ForwardDelay when this bridge is acting as the root. Note that 802.1D specifies that the range for this parameter is related to the value of BridgeMaxAge.
jnxMIMstCistHoldTime	jnxMIDot1sJuniperMstEntry 16	Sets the interval between transmitting two configuration bridge PDUs.
jnxMIMstCistMaxAge	jnxMIDot1sJuniperMstEntry 17	Sets the maximum age of Spanning Tree Protocol (STP) information learned on the ports. The STP information is discarded when the age exceeds the set limit.
jnxMIMstCistForwardDelay	jnxMIDot1sJuniperMstEntry 18	Indicates the time period during which a port stays in a particular state before moving to the next state; for example, from spanning to forwarding state.
jnxMIMstMstpUpCount	jnxMIDot1sJuniperMstEntry 19	Shows the number of times the MSTP Module has been enabled on the bridge.
jnxMIMstMstpDownCount	jnxMIDot1sJuniperMstEntry 20	Shows the number of times the MSTP Module has been disabled on the bridge.
jnxMIMstPathCostDefaultType	jnxMIDot1sJuniperMstEntry 21	Shows the version of the spanning tree default path costs that are to be used by the bridge. ■ A value of 8021d1998(1) uses the 16-bit default path costs from IEEE Std. 802.1D-1998. ■ A value of stp8021t2001(2) uses the 32-bit default path costs from IEEE Std. 802.1t.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 362 of 698

Table 159: jnxMIDot1sJuniperMstTable *(continued)*

Object	Object Identifier	Description
jnxMIMstDebug	jnxMIDot1sJuniperMstEntry 23	Enables debug statements in the MSTP module. A four-byte integer can be used to set the level of debugging. The bit position maps to the following levels of debugging: ■ 0: Init and shutdown debug statements ■ 1: Management debug statements ■ 2: Memory-related debug statements ■ 3: BPDU-related debug statements ■ 4: Event handling debug statements ■ 5: Time module debug statements ■ 6: Port information SEM debug statements ■ 7: Port receive SEM debug statements (valid in the case of MSTP alone) ■ 8: Role selection SEM debug statements ■ 9: Role transition SEM debug statements ■ 10: State transition SEM debug statements ■ 11: Protocol migration SEM debug statements ■ 12: Topology change SEM debug statements ■ 13: Port transmit SEM debug statements ■ 14: Bridge detection SEM debug statements ■ 15: All failure debug statements ■ 16: Redundancy code flow debug statements The rest of the bits remain unused. You can use a combination of debug levels to generate debug statements of multiple debug levels. **NOTE:** Debug options and trace options are mutually exclusive. When the debug option is set, the trace option is set to

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 363 of 698

Table 159: jnxMIDot1sJuniperMstTable *(continued)*

Object	Object Identifier	Description
jnxMIMstForceProtocolVersion	jnxMIDot1sJuniperMstEntry 24	Indicates the version of the Spanning Tree Protocol that is running on the bridge. ■ stpCompatible(0): Represents Spanning Tree Protocol specified in IEEE 802.1D. ■ rstp(2): Represents the Rapid Spanning Tree protocol specified in IEEE 802.1w. ■ mstp(3): Represents the Multiple Spanning Tree protocol specified in IEEE 802.1s.
jnxMIMstTxHoldCount	jnxMIDot1sJuniperMstEntry 25	Indicates the value that the port transmit state machine uses to limit the maximum transmission rate.
jnxMIMstMstiConfigIdSel	jnxMIDot1sJuniperMstEntry 26	Indicates the Configuration Identifier Format Selector that the bridge uses.
jnxMIMstMstiRegionName	jnxMIDot1sJuniperMstEntry 27	Indicates the name of the region's configuration. By default, the region name and the MAC address of the bridge are the same.
jnxMIMstMstiRegionVersion	jnxMIDot1sJuniperMstEntry 28	Indicates the version of the multiple-spanning tree region.
jnxMIMstMstiConfigDigest	jnxMIDot1sJuniperMstEntry 29	Indicates the configuration digest value for the multiple-spanning tree region.
jnxMIMstBufferOverFlowCount	jnxMIDot1sJuniperMstEntry 30	Indicates the number of times buffer overflows or failures have occurred. This event generates a trap.
jnxMIMstMemAllocFailureCount	jnxMIDot1sJuniperMstEntry 31	Indicates the number of times memory allocation failures have occurred. This event generates a trap.
jnxMIMstRegionConfigChangeCount	jnxMIDot1sJuniperMstEntry 32	Indicates the number of times a Region Configuration Identifier Change was detected. This event generates a trap.
jnxMIMstCistBridgeRoleSelectionSemState	jnxMIDot1sJuniperMstEntry 33	Indicates the current state of the Port Role Selection State Machine of the bridge in a common spanning tree context
jnxMIMstCistTimeSinceTopologyChange	jnxMIDot1sJuniperMstEntry 34	Indicates the time since the tcWhile Timer for any port of the Bridge was non-zero for the common spanning-tree context.

Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 364 of 698

No.	'821 Patent Claim 1	Juniper Junos OS System

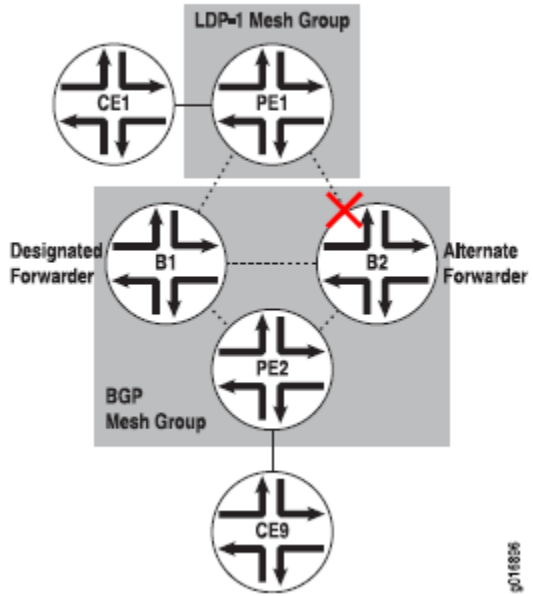
Table 159: jnxMIDot1sJuniperMstTable *(continued)*

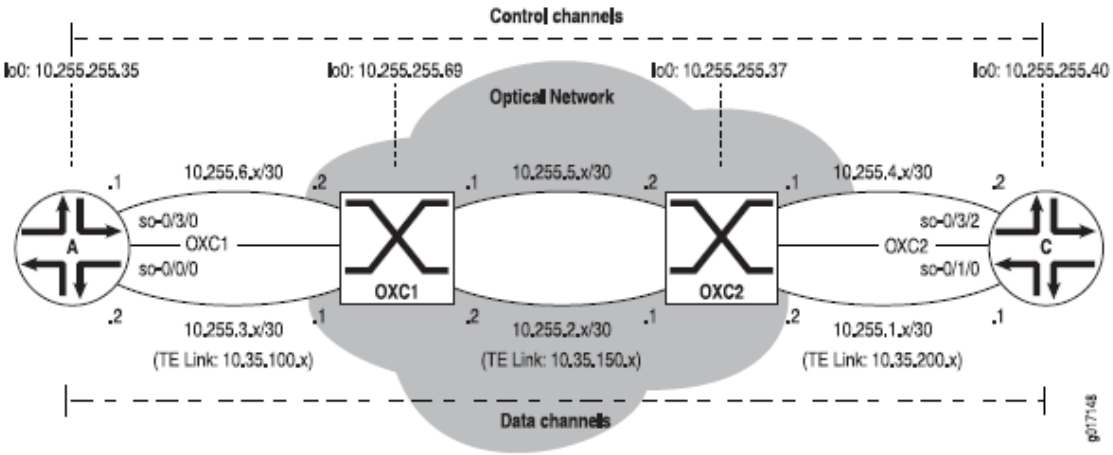
Object	Object Identifier	Description
jnxMIMstCistTopChanges	jnxMIDot1sJuniperMstEntry 35	Indicates the number of times when there was at least one non-zero TcWhile Timer on the bridge for the common spanning-tree context.
jnxMIMstCistNewRootBridgeCount	jnxMIDot1sJuniperMstEntry 36	Indicates how many times the bridge has detected a root bridge change for a common-spanning tree context. This event generates a trap.
jnxMIMstCistHelloTime	jnxMIDot1sJuniperMstEntry 37	Specifies the interval between the transmission of configuration BPDUs by this node on any port when it is either the root of the spanning tree or trying to become the root.
jnxMIMstCistBridgeHelloTime	jnxMIDot1sJuniperMstEntry 38	Specifies the interval between the transmission of configuration bridge PDUs by this node.
jnxMIMstCistDynamicPathcostCalculation	jnxMIDot1sJuniperMstEntry 39	Indicates whether the dynamic path cost calculation is allowed. If set to true, path cost is calculated dynamically from the port speed; otherwise the link speed at the time of port creation is used for calculating the path cost. In both cases, the user has configured a path cost for the port that will be used. By default, dynamic path cost calculation is set to false.

See also Id. at 568-78 (Tables 160-65).

Junos® Software Feature Guide Release 9.5 (19230_JNPR_00001433) at 682-83:

No.	'821 Patent Claim 1	Juniper Junos OS System
		Configuring Multihoming on a VPLS Border Router Configuring multihoming on VPLS border routers ensures that if one border router is unreachable, BGP/LDP PE connectivity is maintained through the other VPLS border router. With multihoming, one border router is chosen as the designated forwarder for each mesh group. The designated forwarder is chosen through either the BGP or VPLS path-selection procedure. If the designated forwarder loses connectivity with a mesh group, the alternate border router then takes over as designated forwarder for that mesh group. A VPLS instance must be configured with BGP signaling in order for multihoming to work. Figure 73 on page 683 shows a simplified example of how multihoming works with VPLS border routers. In this example, B1 is the designated forwarder and B2 is the alternate forwarder. If CE1 wanted to send data to CE9, the data would travel from CE1 to PE1, which is part of the LDP-1 mesh group. PE1 would then flood the data to B1 (the designated forwarder) which would forward the data to PE2. It would not send the data to Router B2. PE2 would then send the data to its destination, CE9. If B1 lost connectivity with the LDP-1 mesh group then B2 would become the designated forwarder. In this case, PE1 would send the data through B2, not through B1.

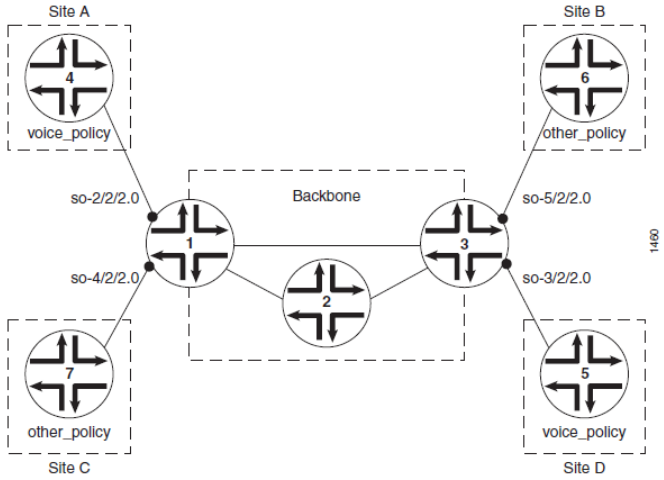
No.	'821 Patent Claim 1	Juniper Junos OS System
		Figure 73: Multihoming for Border Area Routers  Junos® Software Feature Guide Release 9.5 (19230_JNPR_00001433) at 18:

No.	'821 Patent Claim 1	Juniper Junos OS System
		Example: GMPLS Configuration **Figure 2: GMPLS Topology Diagram**  Junos supports several network management standards that account for entity costs. E.g.:

No.	'821 Patent Claim 1	Juniper Junos OS System
		Supported Network Management Features and Standards The JUNOS software supports the use of scripts and event-triggered policies that you write to automate network management. The functions you can perform with these utilities include the following: ■ Automatically detect, diagnose, and fix network problems. When it detects a problem, the script can issue a command that includes options appropriate to the current situation, and then interpret the command output to determine the next appropriate command or action. ■ Periodically check for alarms or other indicators of network or chassis problems, and perform specific actions if they exist. ■ Respond automatically to the occurrence of events and conditions that also trigger system log messages or SNMP traps. ■ Verify that the routing platform's configuration includes only statements you deem appropriate for your network, and automatically add or remove statements as necessary. 8 ■ Supported Network Management Features and Standards 19230_JNPR_00037254.

No.	'821 Patent Claim 1	Juniper Junos OS System
		▪ RFC 1213, *Management Information Base for Network Management of TCP/IP-based internets: MIB-II*. The JUNOS software supports the following features: ▪ MIB II and its SNMP version 2 derivatives, including the following: ▪ Statistics counters ▪ IP (except for the ipRouteTable object, which has been replaced by ipCidrRouteTable [RFC 2096, *IP Forwarding Table MIB*]) ▪ SNMP management ▪ Interface management ▪ SNMP version 1 Get and GetNext requests and version 2 GetBulk requests ▪ JUNOS-specific secured access list ▪ Master configuration keywords ▪ Reconfigurations upon receipt of the SIGHUP signal ▪ RFC 1215, *A Convention for Defining Traps for use with the SNMP* (only MIB II SNMP version 1 traps and version 2 notifications) ▪ RFC 1406, *Definitions of Managed Objects for the DS1 and E1 Interface Types* (T1 MIB is supported) (obsoleted by RFC 2495) ▪ RFC 1407, *Definitions of Managed Objects for the DS3/E3 Interface Type* (T3 MIB is supported) (obsoleted by RFC 2496) ▪ RFC 1472, *The Definitions of Managed Objects for the Security Protocols of the Point-to-Point Protocol* ▪ RFC 1473, *The Definitions of Managed Objects for the IP Network Control Protocol of the Point-to-Point Protocol* ▪ RFC 1657, *Definitions of Managed Objects for the Fourth Version of the Border Gateway Protocol (BGP-4) using SMIv2* (except for the bgpBackwardTransition and bgpEstablished notifications) ▪ RFC 1695, *Definitions of Managed Objects for ATM Management Version 8.0 Using SMIv2* (obsoleted by RFC 2515) ▪ RFC 1724, *RIP Version 2 MIB Extension* ▪ RFC 1850, *OSPF Version 2 Management Information Base* (except for the ospfOriginateNewLsas and ospfRxNewLsas objects, the Host Table, and the ospfLsdbApproachingOverflow, ospfLsdbOverflow, and ospfOriginateLSA traps) ▪ RFC 1901, *Introduction to Community-based SNMPv2* ▪ RFC 1905, *Protocol Operations for Version 2 of the Simple Network Management Protocol (SNMPv2)* (obsoleted by RFC 3416) ▪ RFC 1907, *Management Information Base for Version 2 of the Simple Network Management Protocol (SNMPv2)* (obsoleted by RFC 3418) ▪ RFC 2011, *SNMPv2 Management Information Base for the Internet Protocol using SMIv2* 19230_JNPR_00037256.

No.	'821 Patent Claim 1	Juniper Junos OS System
		are not supported.) MIB draft-ietf-bfd-mib-02.txt is represented by mib-jnx-bfd-exp.txt and is implemented under the Juniper Networks Enterprise jnxExperiment branch. ■ Internet draft draft-ietf-idr-bgp4-mibv2-04.txt, *Definitions of Managed Objects for the Fourth Version of Border Gateway Protocol (BGP-4), Second Version* (only the following objects: jnxBgpM2PrefixInPrefixes, jnxBgpM2PrefixInPrefixesAccepted, and jnxBgpM2PrefixInPrefixesRejected) ■ Internet draft draft-ietf-isis-wg-mib-07.txt, *Management Information Base for IS-IS* (only the following tables: isisISAdjAreaAddrTable, isisISAdjIPAddrTable, isisISAdjProtSuppTable, and isisISAdjTable) ■ Internet draft draft-ietf-msdp-mib-08.txt, *Multicast Source Discovery protocol MIB* (except for the msdpBackwardTransition, msdpEstablished, and msdpRequestsTable objects) ■ Internet draft draft-ietf-ospf-ospfv3-mib-11.txt, *Management Information Base for OSPFv3* (Represented by mib-jnx-ospfv3mib.txt and implemented under the Juniper Networks Enterprise jnxExperiment branch. Support for the ospfv3NbrTable table only. Read-only. Object names are prefixed by jnx, as in jnxOspfv3NbrAddressType, jnxOspfv3NbrPriority, and jnxOspfv3NbrTable.) ■ Internet draft draft-ietf-ppvpn-mpls-vpn-mib-05.txt, *MPLS/BGP Virtual Private Network Management Information Base Using SMIv2* (only the following objects and tables: mplsVpnScalars, mplsVpnVrfPerfTable, mplsVpnVrfRouteTargetTable, and mplsVpnVrfTable; expires May 2003) ■ Internet draft draft-reeder-snmpv3-usm-3desede-00.txt, *Extension to the User-Based Security Model (USM) to Support Triple-DES EDE in "Outside" CBC Mode* ■ Integrated Local Management Interface (ILMI) MIB (only the atmfMYIPNmAddress and atmfPortMyIfname objects). For more information about the ILMI MIB, see the *JUNOS Network Management Configuration Guide* and the ATM Forum Web site at http://www.atmforum.com/. The JUNOS software provides numerous enterprise MIBs, including the following: ■ Alarm ■ Asynchronous Transfer Mode (ATM) ■ ATM class of service (CoS) ■ BGP4 version 2 ■ Chassis ■ Chassis definitions for routing platform model ■ Chassis forwarding ■ CoS ■ Configuration management ■ Destination class usage ■ Dynamic flow capture ■ Ethernet media access control (MAC) 19230_JNPR_00037260.

No.	'821 Patent Claim 1	Juniper Junos OS System
		Example: Configuring Multiple Routing Instances of OSPF Figure 5 on page 229 shows how you can use multiple routing instances of OSPF or OSPFv3 to segregate prefixes within a large network. The network consists of three administrative entities: voice-policy, other-policy, and the default routing instance. Each entity is composed of several geographically separate sites that are connected by the backbone and managed by the backbone entity. **Figure 5: Configuration for Multiple Routing Instances**  Sites A and D belong to the voice-policy routing instance. Sites B and C belong to the other-policy instance. Router 1 and Router 3 at the edge of the backbone connect the routing instances. Each runs a separate OSPF or OSPFv3 instance (one per entity). 19230_JNPR_00033203. JUNOS® Software CLI Reference for J-series Services Routers and SRX-series Services Gateways Release 9.5 (19230_JNPR_00003522) at 83:

No.	'821 Patent Claim 1	Juniper Junos OS System
		priority **Syntax** priority *priority-number* ; **Hierarchy Level** [edit chassis cluster redundancy-group *group-number* node *node-number*] **Release Information** Statement introduced in Release 9.0 of JUNOS software. **Description** Define the priority of a node (device) in a redundancy group. Initiating a failover with the **request chassis cluster failover node** or **request chassis cluster failover redundancy-group** command overrides the priority settings. This statement is supported on J-series and SRX-series devices. **Options** *priority-number* —Priority value of the node. The eligible node with the highest priority is elected master. **Range:** 1 through 254 **Usage Guidelines** For configuration instructions and examples, see the *JUNOS Software Security Configuration Guide*. **Required Privilege Level** interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

<table>
<tr><th>No.</th><th>’821 Patent Claim 1</th><th>Juniper Junos OS System</th></tr>
<tr><td></td><td></td><td>

Table 12: Routing Information and Corresponding CLI show Commands

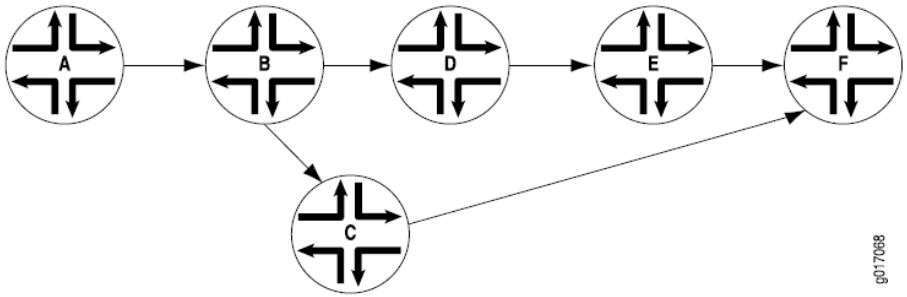
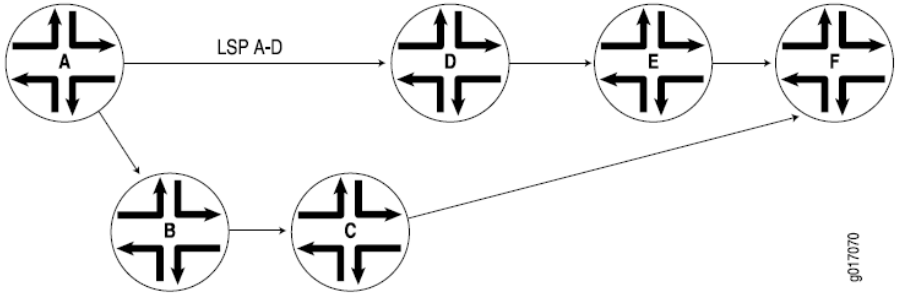
Information Displayed	Corresponding CLI Command
Route Information	
A high-level summary of the routes in the routing table.	show route terse
Detailed information about the active entries in the routing tables.	show route detail
BGP Information	
Summary about Border Gateway Protocol (BGP).	show bgp summary
BGP peers.	show bgp neighbor
OSPF Information	
Information about Open Shortest Path First (OSPF) neighbors.	show ospf neighbors
OSPF interfaces.	show ospf interfaces
OSPF statistics.	show ospf statistics

19230_JNPR_00021551.</td></tr>
<tr><td>1[c]</td><td>selecting an entity pair from said plurality of transport entities based at least in part upon said overall cost; and</td><td>Juniper Junos OS System discloses selecting an entity pair from said plurality of transport entities based at least in part upon said overall cost.

For example, Juniper Junos OS Guide discloses through Fate Sharing, a Juniper router can configure one or more backup paths that minimize the number of shared links with the primary paths to ensure that a single point of failure will be less likely to affect the primary and backup paths simultaneously. Thus, at least under the apparent claim scope alleged by Orckit’s Infringement Disclosures, this limitation is met.

See Supra at 1[b].</td></tr>
</table>

No.	'821 Patent Claim 1	Juniper Junos OS System
		"2.2 LDP route prioritization LDP will pick up IGP route priority after receiving the flash and will dictate the same priority while adding its transit route and ingress route to mpls.0 and inet.3 route table respectively. The flash list prioritization (2.1) of these tables will ensure that these routes go to KRT as per IGP priority. 2.3 KRT route prioritization There are multiple KRT queues with priorities to further classify these routes as per existing heuristics. KRT route install job which serve these queues to send route update to FIB runs at priority 3 for high priority routes and at priority 5 for other routes." Juniper Networks, Inc. Flash LDP and KRT route prioritization (19230_JNPR_00000099) at 2. Junos® Software Network Management Configuration Guide Release 9.5 (19230_JNPR_00000655) at 557: **Table 157: MPLS Traps** **Object** / **Object ID** / **Description** mplsLspInfoUp / mplsLspTraps 1 / Indicates that the LSP (mplsLspInfoName) is up. The current active path is represented by mplsPathInfoName. mplsLspInfoDown / mplsLspTraps 2 / Indicates that the LSP (mplsLspInfoName) is down because the current active path (mplsPathInfoName) has gone down. mplsLspInfoChange / mplsLspTraps 3 / Indicates that the LSP (mplsLspInfoName) has switched traffic to a new active path (mplsPathInfoName) without changing the state (up) before or after the switch. mplsLspInfoPathDown / mplsLspTraps 4 / Indicates that the specified path (mplsPathInfoName) for the LSP (mplsLspInfoName) has gone down. mplsLspInfoPathUp / mplsLspTraps 5 / Indicates that the specified path (mplsPathInfoName) for the LSP (mplsLspInfoName) has come up.

No.	'821 Patent Claim 1	Juniper Junos OS System
		Figure 5: Typical SPF Tree, Sourced from Router A  If an LSP connects Router A to Router D and if IGP shortcuts are enabled on Router A, you might have the SPF tree shown in Figure 6 on page 35. **Figure 6: Modified SPF Tree, Using LSP A–D as a Shortcut**  Note that Router D is now reachable through LSP A–D. When computing the shortest path to reach Router D, Router A has two choices:

No.	'821 Patent Claim 1	Juniper Junos OS System

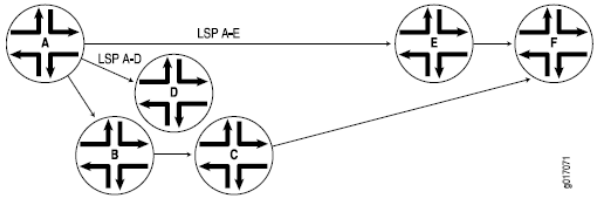
- Use IGP path A–B–D.
- Use LSP A–D.

Router A decides between the two choices by comparing the IGP metrics for path A–B–D with the LSP metrics for LSP A–D. If the IGP metric is lower, path A–B–D is chosen (Figure 5 on page 35). If the LSP metric is lower, LSP A–D is used (Figure 6 on page 35). If both metrics are equal, LSP A–D is chosen because LSP paths are preferred over IGP paths.

Note that Routers E and F are also reachable through LSP A–D, because they are downstream from Router D in the SPF tree.

Assuming that another LSP connects Router A to Router E, you might have the SPF tree shown in Figure 7 on page 36.

Figure 7: Modified SPF Tree, Using LSP A–D and LSP A–E as Shortcuts



This section discusses the following topics:

- Enabling IGP Shortcuts on page 36
- LSPs Qualified in Shortcut Computations on page 36
- IGP Shortcut Applications on page 37
- IGP Shortcuts and Routing Table on page 38
- Router Requirements on page 38
- IGP Shortcuts and VPN Environments on page 38

Enabling IGP Shortcuts

IGP shortcuts are supported for both IS-IS and OSPF. A link-state protocol is required for IGP shortcuts. Shortcuts are disabled by default. For information about enabling IGP shortcuts for IS-IS and OSPF, see the *JUNOS Routing Protocols Configuration Guide*. You can enable IGP shortcuts on a per-router basis; you do not need to enable shortcuts globally. A router's shortcut computation does not depend on another router performing similar computations, and shortcuts performed by other routers are irrelevant.

19230_JNPR_00029240-00029241.

No.	'821 Patent Claim 1	Juniper Junos OS System
		JUNOS® Software CLI Reference for J-series Services Routers and SRX-series Services Gateways Release 9.5 (19230_JNPR_00003522) at 84:

redundancy-group

Syntax

```
redundancy-group group-number {
  gratuitous-arp-count number ;
  interface-monitor interface {
    weight number ;
  }
  node node-number {
    priority priority-number ;
  }
  preempt;
}
```

Hierarchy Level [edit chassis cluster]

Release Information Statement introduced in Release 9.0 of JUNOS software.

Description Define a redundancy group. Except for redundancy group 0, a redundancy group is a logical interface consisting of two physical Ethernet interfaces, one on each chassis. One interface is active, and the other is on standby. When the active interface fails, the standby interface becomes active. The logical interface is called a redundant Ethernet (reth) interface.

Redundancy group 0 consists of the two Routing Engines in the chassis cluster and controls Routing Engine primaryship. You must define redundancy group 0 in the chassis cluster configuration.

This statement is supported on J-series and SRX-series devices.

Options *group-number* —Redundancy group identification number.
Range: 0 through 255

The remaining statements are explained separately.

Usage Guidelines For configuration instructions and examples, see the *JUNOS Software Security Configuration Guide*.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

No.	'821 Patent Claim 1	Juniper Junos OS System
1[d]	if an entity pair reselection event occurs, reselecting said entity pair from the group consisting of said entity pair and a replacement entity pair comprising at least one entity distinct from the entities of said entity pair,	Juniper Junos OS System discloses if an entity pair reselection event occurs, reselecting said entity pair. For example, Juniper Junos OS Guide discloses once an LSP has been established, topology or resources changes might, over time, make the path suboptimal. Juniper Junos OS Guide further discloses examples of when a path becomes suboptimal, such as a new path might have become available that is less congested, has a lower metric, and traverses fewer hops. Juniper Junos OS Guide further discloses that the Juniper Junos OS System is configured to recompute paths periodically to determine whether a more optimal path has become available and then use that more optimal route. Thus, at least under the apparent claim scope alleged by Orckit's Infringement Disclosures, this limitation is met. Juniper Junos OS Guide at Page No. 85 ("[If a node detects that a downstream link has failed (using a link-layer-specific liveness detection mechanism) or that a downstream node has failed (for example, using the RSVP neighbor hello protocol), the node quickly switches the traffic to the detour and, at the same time, signals the ingress router about the link or node failure. Figure 16 on page 47 illustrates the detour taken when the link between Router B and Router C fails. Note that after the node switches traffic to the detour, it might switch the traffic again to a newly calculated detour soon after. This is because the initial detour route might not be the best route. To make rerouting as fast as possible, the node switches traffic onto the initial detour without first verifying that the detour is valid. Once the switch is made, the node recomputes the detour. If the node determines that the initial detour is still valid, traffic continues to flow over this detour. If the node determines that the initial detour is no longer valid, it again switches the traffic to a newly computed detour.]") Juniper Junos OS Guide at Page No. 87 ("[Detour Computations Computing and setting up detours is done independently at each node. On a node, if an LSP has fast reroute enabled and if a downstream link or node can be identified, the router performs a Constrained Shortest Path First (CSPF) computation using the information in the local traffic engineering database. For this reason, detours rely on your IGP supporting traffic engineering extensions. Without the traffic engineering database, detours cannot be established.]")

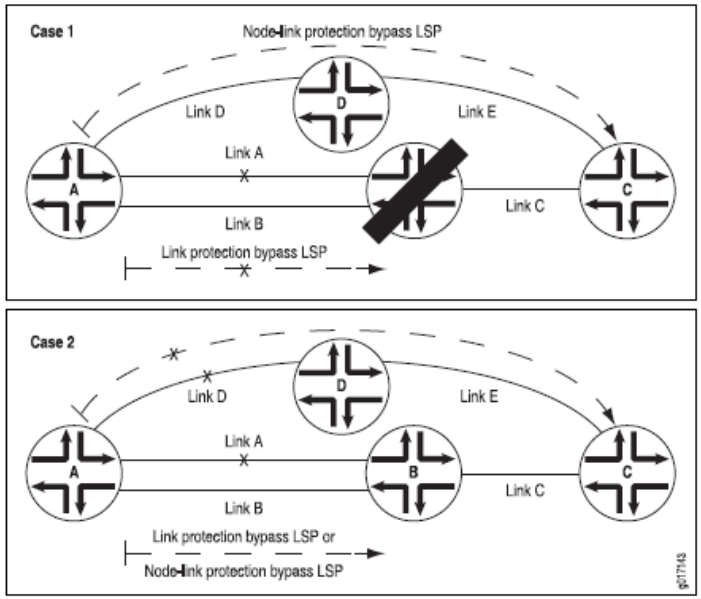
No.	'821 Patent Claim 1	Juniper Junos OS System
		Juniper Junos OS Guide at Page No. 97 ("[When CSPF computes the primary paths of an LSP (or secondary paths when the primary path is not active), it ignores the fate-sharing information. You always want to find the best possible path (least IGP cost) for the primary path. When CSPF computes a secondary path while the primary path (of the same LSP) is active, the following occurs: 1. CSPF identifies all fate-sharing groups that are associated with the primary path. CSPF does this by identifying all links and nodes that the primary path traverses and compiling group lists that contain at least one of the links or nodes. CSPF ignores the ingress and egress nodes in the search. 2. CSPF checks each link in the traffic engineering database against the compiled group list. If the link is a member of a group, the cost of the link is increased by the cost of the group. If a link is a member of multiple groups, all group costs are added together. 3. CSPF performs the check for every node in the traffic engineering database, except the ingress and egress node. Again, a node can belong to multiple groups, so costs are additive. 4. The router performs regular CSPF computation with the adjusted topology.]") Juniper Junos OS Guide at Page No. 146 ("[Optimizing Signaled LSPs Once an LSP has been established, topology or resources changes might, over time, make the path suboptimal. A new path might have become available that is less congested, has a lower metric, and traverses fewer hops. You can configure the router to recompute paths periodically to determine whether a more optimal path has become available.]") Juniper Junos OS Guide at Page No. 147 ("[If reoptimization is enabled, an LSP can be rerouted through different paths by constrained-path recomputations. However, if reoptimization is disabled, the LSP has a fixed path and cannot take advantage of newly available network resources. The LSP is fixed until the next topology change breaks the LSP and forces a recomputation. Reoptimization is not related to failover. A new path is always computed when topology failures occur that disrupt an established path.]") Junos® Software Feature Guide Release 9.5 (19230_JNPR_00001433) at 102:

Node-Link Protection

While link protection is useful for selecting an alternate path to the same router when a link fails, node-link protection establishes a bypass LSP through a different router altogether. For Case 1 in Figure 9 on page 102, link protection allows an LSP to switch to link B and immediately bypass failed link A. However, if Router B fails, link B will fail and the link-protected LSP will be lost.

With node-link protection, the backup LSP can switch to link D instead and bypass the failed links and router. Another benefit of node-link protection shown in Case 2 is that a node-link-protected LSP can act like a link-protected LSP and switch to link B if link D is unavailable.

Figure 9: Link Protection and Node-Link Protection Comparison



JUNOS software signals bypass LSPs dynamically when a protected LSP transverses the protected link. The software determines if the protected LSP needs a node bypass or a link bypass and prepares the necessary bypass LSP automatically. The bypass LSP is torn down automatically when a protected LSP does not use the link.

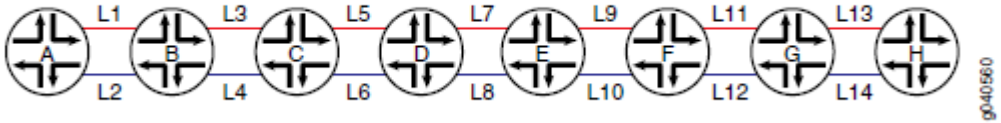
Because the creation and removal of bypass LSPs is automatic, network resources can be used for other purposes when the bypass LSP is not needed. Likewise, network administrators do not need to configure bypass LSPs statically and can focus their maintenance efforts elsewhere.

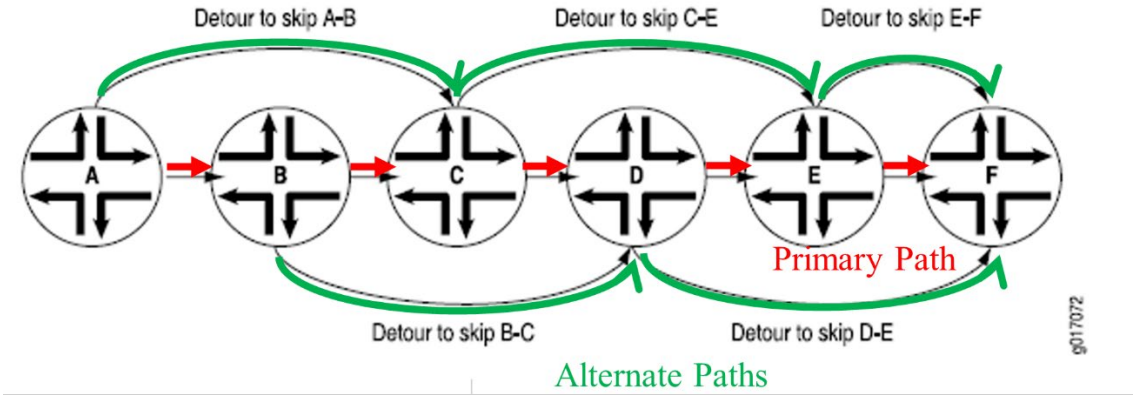
Orckit Exhibit 2017
Cisco Systems, Inc. v. Orckit Corp.
IPR2023-00402, Page 381 of 698

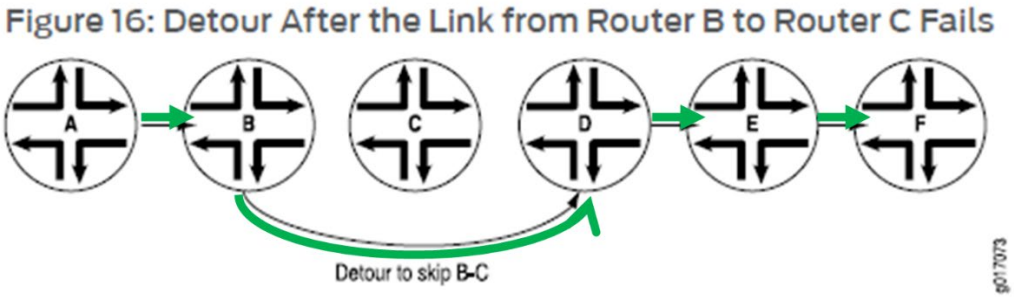
No.	'821 Patent Claim 1	Juniper Junos OS System
1[e]	wherein said entity pair reselection event is selected from a group consisting of adding an entity to said plurality of transport entities, removing an entity from said plurality of transport entities, an operational status change for one of said plurality of transport entities, and a change in overall cost for one of said plurality of transport entities.	Juniper Junos OS System discloses wherein said entity pair reselection event is selected from a group consisting of an operational status change for one of said plurality of transport entities, and a change in overall cost for one of said plurality of transport entities. For example, Juniper Junos OS Guide discloses that the Junos IOS System performs regular CSPF re-computation/ re-optimization after the LSP has been established and topology or resources change, when a new path might have become available, or a new path becomes less congested. Thus, at least under the apparent claim scope alleged by Orckit's Infringement Disclosures, this limitation is met. Juniper Junos OS Guide at Page No. 85 ("[Note that after the node switches traffic to the detour, it might switch the traffic again to a newly calculated detour soon after. This is because the initial detour route might not be the best route. To make rerouting as fast as possible, the node switches traffic onto the initial detour without first verifying that the detour is valid. Once the switch is made, the node recomputes the detour. If the node determines that the initial detour is still valid, traffic continues to flow over this detour. If the node determines that the initial detour is no longer valid, it again switches the traffic to a newly computed detour.]") Juniper Junos OS Guide at Page No. 146 ("[Optimizing Signaled LSPs Once an LSP has been established, topology or resources changes might, over time, make the path suboptimal. A new path might have become available that is less congested, has a lower metric, and traverses fewer hops. You can configure the router to recompute paths periodically to determine whether a more optimal path has become available.]") Juniper Junos OS Guide at Page No. 146 ("[If reoptimization is enabled, an LSP can be rerouted through different paths by constrained-path recomputations. However, if reoptimization is disabled, the LSP has a fixed path and cannot take advantage of newly available network resources. The LSP is fixed until the next topology change breaks the LSP and forces a recomputation. Reoptimization is not related to failover. A new path is always computed when topology failures occur that disrupt an established path.]")

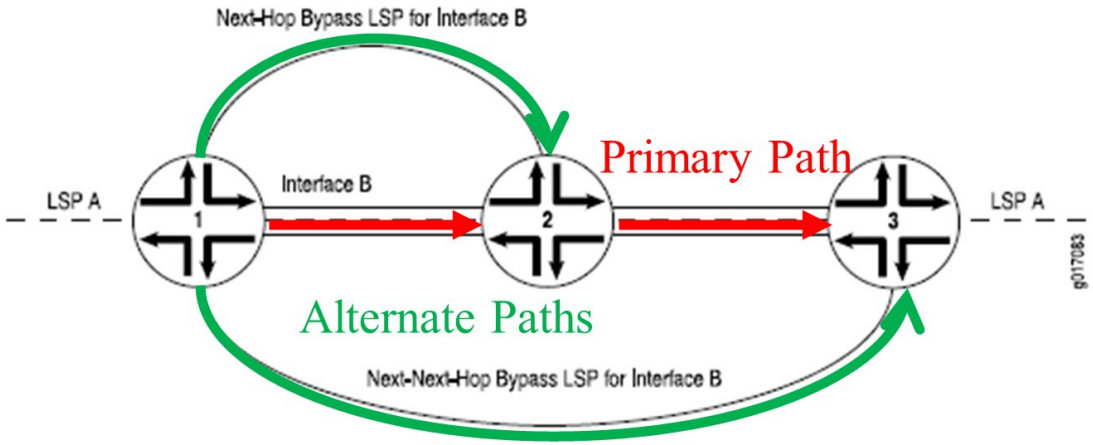
No.	'821 Patent Claim 1	Juniper Junos OS System
		Juniper Junos OS Guide at Page No. 147 ("[After reoptimization is run, the result is accepted only if it meets the following criteria: 1. The new path is not higher in IGP metric. (The metric for the old path is updated during computation, so if a recent link metric changed somewhere along the old path, it is accounted for.) 2. If the new path has the same IGP metric, it is not more hops away. 3. The new path does not cause preemption. (This is to reduce the ripple effect of preemption causing more preemption.) 4. The new path does not worsen congestion overall.]")

No.	'821 Patent Claim 2	Juniper Junos OS System
2[preamble]	The method of claim 1, wherein said step of selecting an entity pair further comprises:	Juniper Junos OS System discloses step of selecting an entity pair further comprises. *See Supra at* 1[c].
2[a]	selecting a working entity from said entity pair; and	Juniper Junos OS Guide discloses selecting a working entity from said entity pair. For example, Juniper Junos OS Guide discloses if the primary path fails, the ingress router immediately reroutes traffic from the failed path to the standby path. Juniper Junos OS Guide further discloses that the primary path sends data traffic through normal operation. Juniper Junos OS Guide at Page No. 83 ("[MPLS and Traffic Protection Typically, when an LSP fails, the router immediately upstream from the failure signals the outage to the ingress router. The ingress router calculates a new path to the egress router, establishes the new LSP, and then directs the traffic from the failed path to the new path. This rerouting process can be time-consuming and prone to failure. For example, the outage signals to the ingress router might get lost, or the new path might take too long to come up, resulting in significant packet drops. The Junos OS provides several complementary mechanisms for protecting against LSP failures: • Standby secondary paths—You can configure primary and secondary paths. You configure secondary paths with the standby statement. To activate traffic protection, you need to

No.	'821 Patent Claim 2	Juniper Junos OS System
		configure these standby paths only on the ingress router. If the primary path fails, the ingress router immediately reroutes traffic from the failed path to the standby path, thereby eliminating the need to calculate a new route and signal a new path. For information about configuring standby LSPs, see "Configuring Hot Standby of Secondary Paths" on page 113. • Fast reroute—You configure fast reroute on an LSP to minimize the effect of a failure in the LSP. Fast reroute enables a router upstream from the failure to route around the failure quickly to the router downstream of the failure. The upstream router then signals the outage to the ingress router, thereby maintaining connectivity before a new LSP is established. For a detailed overview of fast reroute, see "Fast Reroute Overview" on page 46. For information about configuring fast reroute, see "Configuring Fast Reroute" on page 82.]") Figure 20: least-fill Load Balancing Algorithm Example  Fig. 20 Juniper Junos OS Guide at Page No. 148 ("As shown in Figure 20 on page 110, there are two potential paths for an LSP to traverse from router A to router H, the odd links from L1 through L13 and the even links from L2 through L14. Currently, the router is using the even links as the active path for the LSP.") Juniper Junos OS Guide at Page No. 78 ("Link redundancy—If a link or a line card failure affects an aggregate member link, the traffic flowing across that link is immediately forwarded across one of the remaining links.") Juniper Junos OS Guide at Page No. 84 ("When standby secondary path, and fast reroute or link protection are configured on an LSP, full traffic protection is enabled. When a failure occurs in an LSP, the router upstream from the failure routes traffic around the failure and notifies the ingress router of the failure. This rerouting keeps the traffic flowing while waiting for the notification to be processed at the ingress router. After receiving the failure

No.	'821 Patent Claim 2	Juniper Junos OS System
		notification, the ingress router immediately reroutes the traffic from the patched primary path to the more optimal standby path.") Juniper Junos OS Guide at Page No. 85 ("[Fast Reroute Overview Fast reroute provides redundancy for an LSP path. When you enable fast reroute, detours are precomputed and preestablished along the LSP. In case of a network failure on the current LSP path, traffic is quickly routed to one of the detours. Figure 15 on page 47 illustrates an LSP from Router A to Router F, showing the established detours. Each detour is established by an upstream node to avoid the link toward the immediate downstream node and the immediate downstream node itself. Each detour might traverse through one or more label-switched routers that are not shown in the figure. Fast reroute protects traffic against any single point of failure between the ingress and egress routers. If there are multiple failures along an LSP, fast reroute itself might fail. Also, fast reroute does not protect against failure of the ingress or egress routers. Figure 15: Detours Established for an LSP Using Fast Reroute  Fig. 15 (annotation added) If a node detects that a downstream link has failed (using a link-layer-specific liveness detection mechanism) or that a downstream node has failed (for example, using the RSVP

No.	'821 Patent Claim 2	Juniper Junos OS System
		neighbor hello protocol), the node quickly switches the traffic to the detour and, at the same time, signals the ingress router about the link or node failure. Figure 16 on page 47 illustrates the detour taken when the link between Router B and Router C fails.]")  Fig. 16 (annotation added)
2[b]	selecting a protection entity from said entity pair.	Juniper Junos OS System discloses selecting a protection entity from said entity pair. For example, Juniper Junos OS Guide discloses if the primary path fails, the ingress router immediately reroutes traffic from the failed path to the standby path. Juniper Junos OS Guide at Page No. 83 ("[MPLS and Traffic Protection Typically, when an LSP fails, the router immediately upstream from the failure signals the outage to the ingress router. The ingress router calculates a new path to the egress router, establishes the new LSP, and then directs the traffic from the failed path to the new path. This rerouting process can be time-consuming and prone to failure. For example, the outage signals to the ingress router might get lost, or the new path might take too long to come up, resulting in significant packet drops. The Junos OS provides several complementary mechanisms for protecting against LSP failures: • Standby secondary paths—You can configure primary and secondary paths. You configure secondary paths with the standby statement. To activate traffic protection, you need to configure these standby paths only on the ingress router. If the primary path fails, the ingress router immediately reroutes traffic from the failed path to the standby path, thereby

No.	'821 Patent Claim 2	Juniper Junos OS System
		eliminating the need to calculate a new route and signal a new path. For information about configuring standby LSPs, see "Configuring Hot Standby of Secondary Paths" on page 113. • Fast reroute—You configure fast reroute on an LSP to minimize the effect of a failure in the LSP. Fast reroute enables a router upstream from the failure to route around the failure quickly to the router downstream of the failure. The upstream router then signals the outage to the ingress router, thereby maintaining connectivity before a new LSP is established. For a detailed overview of fast reroute, see "Fast Reroute Overview" on page 46. For information about configuring fast reroute, see "Configuring Fast Reroute" on page 82.]") Juniper Junos OS Guide at Page No. 330 ("In Figure 23 on page 292, both node protection and link protection are enabled on Interface B on Router 1. Both node protection and link protection are also enabled on LSP A, an LSP that traverses the link transiting Router 1, Router 2, and Router 3. If the link between Router 1 and Router 2 fails, traffic from LSP A is quickly switched to the next-hop bypass LSP generated by link protection. If Router 2 suffers a hardware or software failure, traffic from LSP A is switched to the next-next-hop bypass LSP generated by node protection.") Figure 23: Node Protection Creating a Next-Next-Hop Bypass LSP  Fig. 23 (annotation added)

No.	'821 Patent Claim 2	Juniper Junos OS System

No.	'821 Patent Claim 3	Juniper Junos OS System
3	The method of claim 2, further comprising the step of selecting an active entity from the set consisting of said working entity and said protection entity.	Juniper Junos OS System discloses step of selecting an active entity from the set consisting of said working entity and said protection entity. For example, Juniper Junos OS Guide discloses that data traffic is normally sent through a primary path, working entity. Juniper Junos OS Guide further discloses that if the primary path fails, the ingress router immediately reroutes traffic from the failed path to the standby path. Juniper Junos OS Guide at Page No. 83 ("[MPLS and Traffic Protection Typically, when an LSP fails, the router immediately upstream from the failure signals the outage to the ingress router. The ingress router calculates a new path to the egress router, establishes the new LSP, and then directs the traffic from the failed path to the new path. This rerouting process can be time-consuming and prone to failure. For example, the outage signals to the ingress router might get lost, or the new path might take too long to come up, resulting in significant packet drops. The Junos OS provides several complementary mechanisms for protecting against LSP failures: • Standby secondary paths—You can configure primary and secondary paths. You configure secondary paths with the standby statement. To activate traffic protection, you need to configure these standby paths only on the ingress router. If the primary path fails, the ingress router immediately reroutes traffic from the failed path to the standby path, thereby eliminating the need to calculate a new route and signal a new path. For information about configuring standby LSPs, see "Configuring Hot Standby of Secondary Paths" on page 113. • Fast reroute—You configure fast reroute on an LSP to minimize the effect of a failure in the LSP. Fast reroute enables a router upstream from the failure to route around the failure quickly to the router downstream of the failure. The upstream router then signals the outage to the ingress router, thereby maintaining connectivity before a new LSP is established. For a detailed overview of fast reroute, see "Fast Reroute Overview" on page 46. For information about configuring fast reroute, see "Configuring Fast Reroute" on page 82.]")

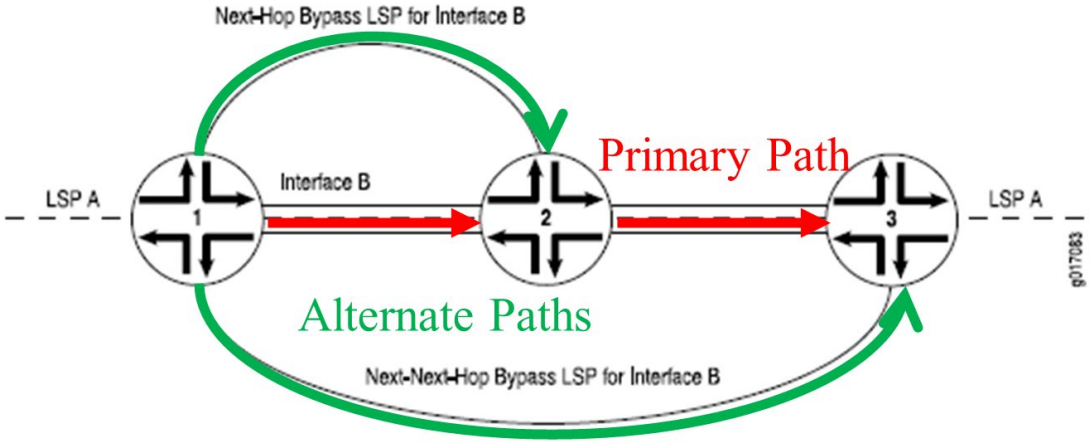
No.	'821 Patent Claim 3	Juniper Junos OS System
		Junos® Design and Implementation Guide for J-series Services Routers Release 9.4 (19230_JNPR_00002357) at 61: ■ passive Rejects traffic destined for the route. Generally, only active routes are included in the routing and forwarding tables. If a static route's next-hop address is unreachable, the route is marked ineligible, and it is not included in the routing or forwarding tables. To force a route to be included in the routing tables regardless of next-hop reachability, you can flag the route as **passive**. If a route is flagged as passive and its next-hop address is unreachable, the route is included in the routing table and all traffic destined for it is rejected.

No.	'821 Patent Claim 4	Juniper Junos OS System
4	The method of claim 2, wherein selecting an entity pair further comprises minimizing an overall cost function.	Juniper Junos OS System discloses selecting an entity pair further comprises minimizing an overall cost function. For example, Juniper Junos OS Guide discloses that through fate sharing, you can configure one or more backup paths that minimize the number of shared links with the primary paths. Juniper Junos OS further discloses that Fate Sharing ensures that a single point of failure will not affect the primary and backup paths simultaneously. Thus, at least under the apparent claim scope alleged by Orckit's Infringement Disclosures, this limitation is met. Juniper Junos OS Guide at Page No. 73 ("[Through fate sharing, you can configure backup paths that minimize the number of shared links and fiber paths with the primary paths as much as possible, to ensure that if a fiber is cut, the minimum amount of data is lost and a path still exists to the destination. For a backup path to work optimally, it must not share links or physical fiber paths with the primary path, ensuring that a single point of failure will not affect the primary and backup paths simultaneously. For more information about fate sharing, see the Junos OS Routing Protocols Configuration Guide.]") Juniper Junos OS Guide at Page No. 97 ("[Fate-sharing groups contain three types of objects:

No.	'821 Patent Claim 4	Juniper Junos OS System
		• Point-to-point links—Identified by the IP addresses at each end of the link. Unnumbered point-to-point links are typically identified by borrowing IP addresses from other interfaces. Order is not important; from 1.2.3.4 to 1.2.3.5 and from 1.2.3.5 to 1.2.3.4 have the same meaning. • Non-point-to-point links—Include links on a LAN interface (such as Gigabit Ethernet interfaces) or nonbroadcast multiaccess (NBMA) interfaces (such as Asynchronous Transfer Mode [ATM] or Frame Relay). You identify these links by their individual interface address. For example, if the LAN interface 192.168.200.0/24 has four routers attached to it, each router link is individually identified: from 192.168.200.1; # LAN interface of router 1 from 192.168.200.2; # LAN interface of router 2 from 192.168.200.3; # LAN interface of router 3 from 192.168.200.4; # LAN interface of router 4 You can list the addresses in any order. • A router node—Identified by its configured router ID. An object can be in any number of groups, and a group can contain any number of objects. Each group has a configurable cost attributed to it, which represents the level of impact this group has on CSPF computations. The higher the cost, the less likely a backup path will share with the primary path any objects in the group. The cost is directly comparable to traffic engineering metrics. By default, the cost is 1. Changing the fate-sharing database does not affect established LSPs until the next reoptimization of CSPF. The fate-sharing database does influence fast-reroute computations.]") JUNOS® Software CLI Reference for J-series Services Routers and SRX-series Services Gateways Release 9.5 (19230_JNPR_00003522) at 692:

No.	'821 Patent Claim 4	Juniper Junos OS System
		Options match (equal | greater-than | less-than | not-equal)—Match an operand. value *type-of-service-in-decimal* —The following service types are available: ■ 0000—Default ■ 0001—Minimize Cost ■ 0002—Maximize Reliability ■ 0003—Maximize Throughput ■ 0004—Minimize Delay ■ 0005—Maximize Security

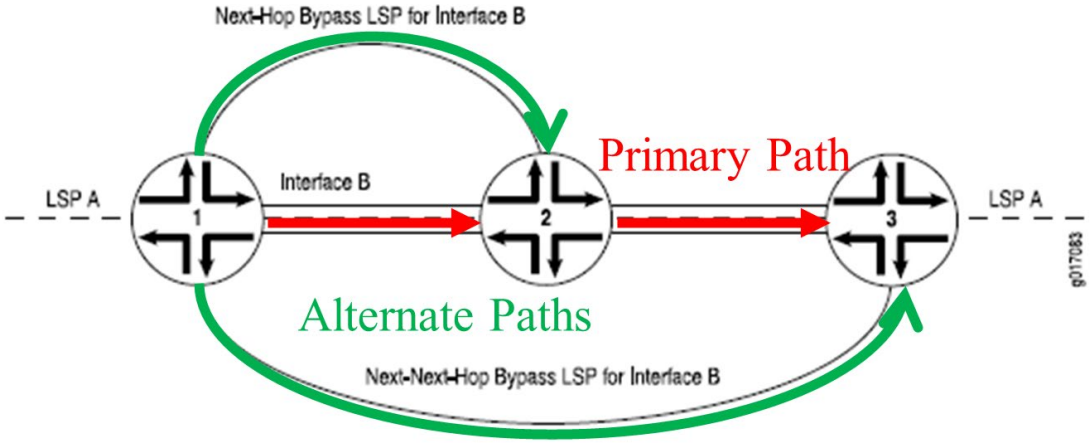
No.	'821 Patent Claim 5	Juniper Junos OS System
5	The method of claim 4, wherein said overall cost function comprises substantially minimizing a probability of concurrent failure of said protection entity and said working entity.	Juniper Junos OS System discloses overall cost function comprises substantially minimizing a probability of concurrent failure of said protection entity and said working entity. For example, Juniper Junos OS Guide discloses that through fate sharing, you can configure one or more backup paths that minimize the number of shared links with the primary paths. Juniper Junos OS further discloses that fate sharing ensures that a single point of failure will not affect the primary and backup paths simultaneously. Thus, at least under the apparent claim scope alleged by Orckit's Infringement Disclosures, this limitation is met. Juniper Junos OS Guide at Page No. 73 ("[Fate sharing allows you to create a database of information that CSPF uses to compute one or more backup paths to use in case the primary path becomes unstable. The database describes the relationships between elements of the network, such as routers and links. You can specify one or more elements within a group. Through fate sharing, you can configure backup paths that minimize the number of shared links and fiber paths with the primary paths as much as possible, to ensure that if a fiber is cut, the minimum amount of data is lost and a path still exists to the destination. For a backup path to work optimally, it must not share links or physical fiber paths with the primary path, ensuring that a single point of failure will not affect the primary and backup

No.	'821 Patent Claim 5	Juniper Junos OS System
		paths simultaneously. For more information about fate sharing, see the Junos OS Routing Protocols Configuration Guide.]") Juniper Junos OS Guide at Page No. 96 ("[For a backup path to work optimally, it must not share links or physical fiber paths with the primary path. This ensures that a single point of failure will not affect the primary and backup paths at the same time.]") Juniper Junos OS Guide at Page No. 330 ("In Figure 23 on page 292, both node protection and link protection are enabled on Interface B on Router 1. Both node protection and link protection are also enabled on LSP A, an LSP that traverses the link transiting Router 1, Router 2, and Router 3. If the link between Router 1 and Router 2 fails, traffic from LSP A is quickly switched to the next-hop bypass LSP generated by link protection. If Router 2 suffers a hardware or software failure, traffic from LSP A is switched to the next-next-hop bypass LSP generated by node protection.") Figure 23: Node Protection Creating a Next-Next-Hop Bypass LSP  Fig. 23 (annotation added)

No.	'821 Patent Claim 5	Juniper Junos OS System
		In addition, Juniper Networks discloses the use of shared risk link groups ("SRLGs"), which also minimize a number of shared links among LSPs. *See, e.g.*, 19230_JNPR_00041940; 19230_JNPR_00041941; 19230_JNPR_00041954; 19230_JNPR_00041967; 19230_JNPR_00041981; 19230_JNPR_00041995; 19230_JNPR_00042011; 19230_JNPR_00042015; 19230_JNPR_00042019; 19230_JNPR_00042026; 19230_JNPR_00042029; 19230_JNPR_00042042; 19230_JNPR_00042051; 19230_JNPR_00042053; 19230_JNPR_00042054.

No.	'821 Patent Claim 6	Juniper Junos OS System
6	The method of claim 4, wherein said overall cost function comprises a predefined entity cost metric.	Juniper Junos OS System discloses overall cost function comprises a predefined entity cost metric. For example, Juniper Junos OS Guide discloses detour computation and fast reroute optimization features that rely on IGP fate sharing. As explained above, fate sharing calculates non-overlapping routes based upon various entity cost metrics. Juniper Junos OS Guide at Page No. 87 ("[Detour Computations Computing and setting up detours is done independently at each node. On a node, if an LSP has fast reroute enabled and if a downstream link or node can be identified, the router performs a Constrained Shortest Path First (CSPF) computation using the information in the local traffic engineering database. For this reason, detours rely on your IGP supporting traffic engineering extensions. Without the traffic engineering database, detours cannot be established.]") Juniper Junos OS Guide at Page No. 97 ("[Fate-sharing groups contain three types of objects: • Point-to-point links—Identified by the IP addresses at each end of the link. Unnumbered point-to-point links are typically identified by borrowing IP addresses from other interfaces. Order is not important; from 1.2.3.4 to 1.2.3.5 and from 1.2.3.5 to 1.2.3.4 have the same meaning. • Non-point-to-point links—Include links on a LAN interface (such as Gigabit Ethernet interfaces) or nonbroadcast multiaccess (NBMA) interfaces (such as Asynchronous Transfer Mode [ATM] or Frame Relay). You identify these links by their individual interface address.

No.	'821 Patent Claim 6	Juniper Junos OS System
		For example, if the LAN interface 192.168.200.0/24 has four routers attached to it, each router link is individually identified: from 192.168.200.1; # LAN interface of router 1 from 192.168.200.2; # LAN interface of router 2 from 192.168.200.3; # LAN interface of router 3 from 192.168.200.4; # LAN interface of router 4 You can list the addresses in any order. • A router node—Identified by its configured router ID. An object can be in any number of groups, and a group can contain any number of objects. Each group has a configurable cost attributed to it, which represents the level of impact this group has on CSPF computations. The higher the cost, the less likely a backup path will share with the primary path any objects in the group. The cost is directly comparable to traffic engineering metrics. By default, the cost is 1. Changing the fate-sharing database does not affect established LSPs until the next reoptimization of CSPF. The fate-sharing database does influence fast-reroute computations.]") Juniper Junos OS Guide at Page No. 330 ("In Figure 23 on page 292, both node protection and link protection are enabled on Interface B on Router 1. Both node protection and link protection are also enabled on LSP A, an LSP that traverses the link transiting Router 1, Router 2, and Router 3. If the link between Router 1 and Router 2 fails, traffic from LSP A is quickly switched to the next-hop bypass LSP generated by link protection. If Router 2 suffers a hardware or software failure, traffic from LSP A is switched to the next-next-hop bypass LSP generated by node protection.")

No.	'821 Patent Claim 6	Juniper Junos OS System
		Figure 23: Node Protection Creating a Next-Next-Hop Bypass LSP  Fig. 23 (annotation added)

No.	'821 Patent Claim 7	Juniper Junos OS System
7	The method of claim 6, wherein said predefined entity cost metric is selected from the group consisting of interior gateway protocol (IGP) and traffic engineering (TE).	Juniper Junos OS System discloses said predefined entity cost metric is selected from the group consisting of interior gateway protocol (IGP) and traffic engineering (TE). For example, Juniper Junos OS Guide discloses detour computation and fast reroute optimization that rely on IGP supporting traffic engineering extensions and fate sharing. Junos IOS Guide further discloses that the Juniper Junos OS System uses MPLS, which includes traffic engineering. Juniper Junos OS Guide at Page No. 87 ("[Detour Computations Computing and setting up detours is done independently at each node. On a node, if an LSP has fast reroute enabled and if a downstream link or node can be identified, the router performs a Constrained Shortest Path First (CSPF) computation using the information in the local traffic engineering database. For this reason, detours rely on your IGP supporting traffic engineering extensions. Without the traffic engineering database, detours cannot be established.]")

No.	'821 Patent Claim 7	Juniper Junos OS System
		Juniper Junos OS Guide at Page No. 97 ("[Fate-sharing groups contain three types of objects: • Point-to-point links—Identified by the IP addresses at each end of the link. Unnumbered point-to-point links are typically identified by borrowing IP addresses from other interfaces. Order is not important; from 1.2.3.4 to 1.2.3.5 and from 1.2.3.5 to 1.2.3.4 have the same meaning. • Non-point-to-point links—Include links on a LAN interface (such as Gigabit Ethernet interfaces) or nonbroadcast multiaccess (NBMA) interfaces (such as Asynchronous Transfer Mode [ATM] or Frame Relay). You identify these links by their individual interface address. For example, if the LAN interface 192.168.200.0/24 has four routers attached to it, each router link is individually identified: from 192.168.200.1; # LAN interface of router 1 from 192.168.200.2; # LAN interface of router 2 from 192.168.200.3; # LAN interface of router 3 from 192.168.200.4; # LAN interface of router 4 You can list the addresses in any order. • A router node—Identified by its configured router ID. An object can be in any number of groups, and a group can contain any number of objects. Each group has a configurable cost attributed to it, which represents the level of impact this group has on CSPF computations. The higher the cost, the less likely a backup path will share with the primary path any objects in the group. The cost is directly comparable to traffic engineering metrics. By default, the cost is 1. Changing the fate-sharing database does not affect established LSPs until the next reoptimization of CSPF. The fate-sharing database does influence fast-reroute computations.]") Junos® Software Feature Guide Release 9.5 (19230_JNPR_00001433) at 17-18:

No.	'821 Patent Claim 7	Juniper Junos OS System

		Option: Configuring GMPLS Support for Unnumbered Links
		
		The JUNOS software supports RSVP traffic engineering over unnumbered interfaces. With this feature, you do not have to configure IP addresses for each interface participating in the RSVP-signaled network. Traffic engineering information about unnumbered links is carried in the IGP traffic engineering extensions for OSPF and IS-IS, as described in RFC 4203, *OSPF Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)*. Unnumbered links can also be specified in
		
		the MPLS traffic engineering signaling, as described in RFC 3477, *Signaling Unnumbered Links in Resource ReSerVation Protocol - Traffic Engineering (RSVP-TE)*.
		
		To configure RSVP for unnumbered interfaces, you must configure the router with a router ID using the **router-id** statement specified at the **[edit routing-options]** hierarchy level. The router ID must be routable (you can typically use the loopback address). The RSVP control messages for the unnumbered links are sent using the router ID address (rather than a randomly selected address). To configure link protection and fast reroute on a router with unnumbered interfaces enabled, you must configure at least two addresses. In addition to the router ID, Juniper Networks recommends that you configure a secondary interface on the loopback:
		
		[edit]
		routing-options {
		router-id *address*;
		}

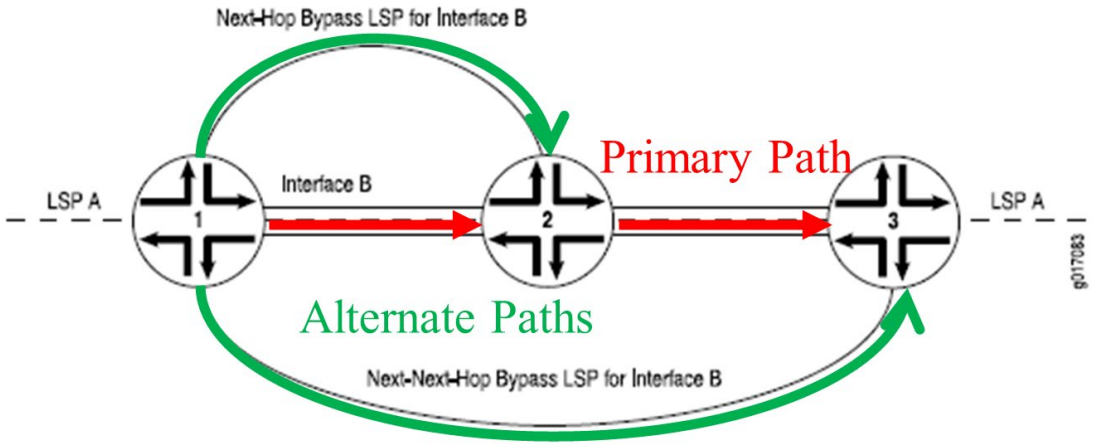
No.	'821 Patent Claim 8	Juniper Junos OS System
8	The method of claim 4, further comprising the step of configuring said working entity as revertive.	Juniper Junos OS System discloses the step of configuring said working entity as revertive. For example, Juniper Junos OS Guide discloses if a static LSP goes down and traffic is switched to the bypass LSP, the protection revert timer specifies the amount of time (in seconds) that the LSP must wait before it can revert back to the original static LSP. Juniper Junos OS Guide at Page No. 83 ("[MPLS and Traffic Protection Typically, when an LSP fails, the router immediately upstream from the failure signals the outage to the ingress router. The ingress router calculates a new path to the egress router, establishes the new LSP, and then directs the traffic from the failed path to the new path. This rerouting process can be time-consuming and prone to failure. For example, the outage

No.	'821 Patent Claim 8	Juniper Junos OS System
		signals to the ingress router might get lost, or the new path might take too long to come up, resulting in significant packet drops. The Junos OS provides several complementary mechanisms for protecting against LSP failures: • Standby secondary paths—You can configure primary and secondary paths. You configure secondary paths with the standby statement. To activate traffic protection, you need to configure these standby paths only on the ingress router. If the primary path fails, the ingress router immediately reroutes traffic from the failed path to the standby path, thereby eliminating the need to calculate a new route and signal a new path. For information about configuring standby LSPs, see "Configuring Hot Standby of Secondary Paths" on page 113. • Fast reroute—You configure fast reroute on an LSP to minimize the effect of a failure in the LSP. Fast reroute enables a router upstream from the failure to route around the failure quickly to the router downstream of the failure. The upstream router then signals the outage to the ingress router, thereby maintaining connectivity before a new LSP is established. For a detailed overview of fast reroute, see "Fast Reroute Overview" on page 46. For information about configuring fast reroute, see "Configuring Fast Reroute" on page 82.]") Juniper Junos OS Guide at Page No. 83 ("[Configuring the Protection Revert Timer for Static LSPs For static LSPs configured with a bypass static LSP, it is possible to configure the protection revert timer. If a static LSP goes down and traffic is switched to the bypass LSP, the protection revert timer specifies the amount of time (in seconds) that the LSP must wait before it can revert back to the original static LSP.]") Junos® Design and Implementation Guide for J-series Services Routers Release 9.4 (19230_JNPR_00002357) at 61: ■ retain Keeps the route in the forwarding table after the routing process shuts down or the next- hop router goes down. (By default, routes are not retained in the routing table. When the routing process starts up again, any routes configured as static routes must be added by the system to the forwarding table again. To avoid this latency, routes can be flagged as **retain**, so that they are kept in the forwarding table even after the routing process shuts down.

No.	'821 Patent Claim 8	Juniper Junos OS System
		JUNOS® Software CLI Reference for J-series Services Routers and SRX-series Services Gateways Release 9.5 (19230_JNPR_00003522) at 30: **revert-interval** See the following sections: ■ revert-interval (LDAP) on page 30 ■ revert-interval (RADIUS) on page 31 ***revert-interval (LDAP)*** **Syntax** revert-interval *seconds*; **Hierarchy Level** [edit access ldap-options], [edit access profile *profile-name* ldap-options] **Release Information** Statement introduced in Release 8.5 of JUNOS software. **Description** Specify the amount of time that elapses before the primary server is contacted if backup server is being used. This statement is supported on J-series and SRX-series devices. **Options** *seconds*—Number of seconds that elapse before the primary server is contacted. **Range:** 60 through 4294967295 seconds **Default:** 600 seconds **Usage Guidelines** For configuration instructions and examples, see the *JUNOS Software Security Configuration Guide*. **Required Privilege Level** access—To view this statement in the configuration. access-control—To add this statement to the configuration.

No.	'821 Patent Claim 9	Juniper Junos OS System
9[preamble]	The method of claim 4, wherein said overall cost function comprises:	Juniper Junos OS System discloses overall cost function comprises. See *Supra* at 4.

No.	'821 Patent Claim 9	Juniper Junos OS System
9[a]	selecting a subset of entity pairs wherein each entity pair of said subset has substantially minimum probability of a concurrent failure of said protection entity and said working entity; and	Juniper Junos OS System discloses selecting a subset of entity pairs wherein selecting a subset of entity pairs wherein each entity pair of said subset has substantially minimum probability of a concurrent failure of said protection entity and said working entity. For example, Juniper Junos OS Guide discloses that through fate sharing, you can configure one or more backup paths that minimize the number of shared links with the primary paths. Juniper Junos OS Guide further discloses that fate sharing ensures that a single point of failure will not affect the primary and backup paths simultaneously. Thus, at least under the apparent claim scope alleged by Orckit's Infringement Disclosures, this limitation is met. Juniper Junos OS Guide at Page No. 73 ("[Fate sharing allows you to create a database of information that CSPF uses to compute one or more backup paths to use in case the primary path becomes unstable. The database describes the relationships between elements of the network, such as routers and links. You can specify one or more elements within a group. Through fate sharing, you can configure backup paths that minimize the number of shared links and fiber paths with the primary paths as much as possible, to ensure that if a fiber is cut, the minimum amount of data is lost and a path still exists to the destination. For a backup path to work optimally, it must not share links or physical fiber paths with the primary path, ensuring that a single point of failure will not affect the primary and backup paths simultaneously. For more information about fate sharing, see the Junos OS Routing Protocols Configuration Guide.]") Juniper Junos OS Guide at Page No. 96 ("[For a backup path to work optimally, it must not share links or physical fiber paths with the primary path. This ensures that a single point of failure will not affect the primary and backup paths at the same time.]") Juniper Junos OS Guide at Page No. 330 ("In Figure 23 on page 292, both node protection and link protection are enabled on Interface B on Router 1. Both node protection and link protection are also enabled on LSP A, an LSP that traverses the link transiting Router 1, Router 2, and Router 3. If the link between Router 1 and Router 2 fails, traffic from LSP A is quickly switched to the next-hop bypass LSP generated by link protection. If Router 2

No.	'821 Patent Claim 9	Juniper Junos OS System
		suffers a hardware or software failure, traffic from LSP A is switched to the next-next-hop bypass LSP generated by node protection.") Figure 23: Node Protection Creating a Next-Next-Hop Bypass LSP  Fig. 23 (annotation added)
9[b]	if said subset comprises at least two entity pairs, selecting an entity pair from said subset that minimizes an entity cost function.	Juniper Junos OS System discloses if said subset comprises at least two entity pairs, selecting an entity pair from said subset that minimizes an entity cost function. For example, Juniper Junos OS Guide discloses that through fate sharing, you can configure one or more backup paths that minimize the number of shared links with the primary paths. Juniper Junos OS Guide further discloses that fate sharing ensures that a single point of failure will not affect the primary and backup paths simultaneously. Thus, at least under the apparent claim scope alleged by Orckit's Infringement Disclosures, this limitation is met. Juniper Junos OS Guide at Page No. 73 ("[Through fate sharing, you can configure backup paths that minimize the number of shared links and fiber paths with the primary paths as much as possible, to ensure that if a fiber is cut, the minimum amount of data is lost and a path still exists to the destination.

No.	'821 Patent Claim 9	Juniper Junos OS System
		For a backup path to work optimally, it must not share links or physical fiber paths with the primary path, ensuring that a single point of failure will not affect the primary and backup paths simultaneously. For more information about fate sharing, see the Junos OS Routing Protocols Configuration Guide.]") Juniper Junos OS Guide at Page No. 97 ("[Fate-sharing groups contain three types of objects: • Point-to-point links—Identified by the IP addresses at each end of the link. Unnumbered point-to-point links are typically identified by borrowing IP addresses from other interfaces. Order is not important; from 1.2.3.4 to 1.2.3.5 and from 1.2.3.5 to 1.2.3.4 have the same meaning. • Non-point-to-point links—Include links on a LAN interface (such as Gigabit Ethernet interfaces) or nonbroadcast multiaccess (NBMA) interfaces (such as Asynchronous Transfer Mode [ATM] or Frame Relay). You identify these links by their individual interface address. For example, if the LAN interface 192.168.200.0/24 has four routers attached to it, each router link is individually identified: from 192.168.200.1; # LAN interface of router 1 from 192.168.200.2; # LAN interface of router 2 from 192.168.200.3; # LAN interface of router 3 from 192.168.200.4; # LAN interface of router 4 You can list the addresses in any order. • A router node—Identified by its configured router ID. An object can be in any number of groups, and a group can contain any number of objects. Each group has a configurable cost attributed to it, which represents the level of impact this group has on CSPF computations. The higher the cost, the less likely a backup path will share with the primary path any objects in the group. The cost is directly comparable to traffic engineering metrics. By default, the cost is 1. Changing the fate-sharing database does not affect established LSPs until the next reoptimization of CSPF. The fate-sharing database does influence fast-reroute computations.]")

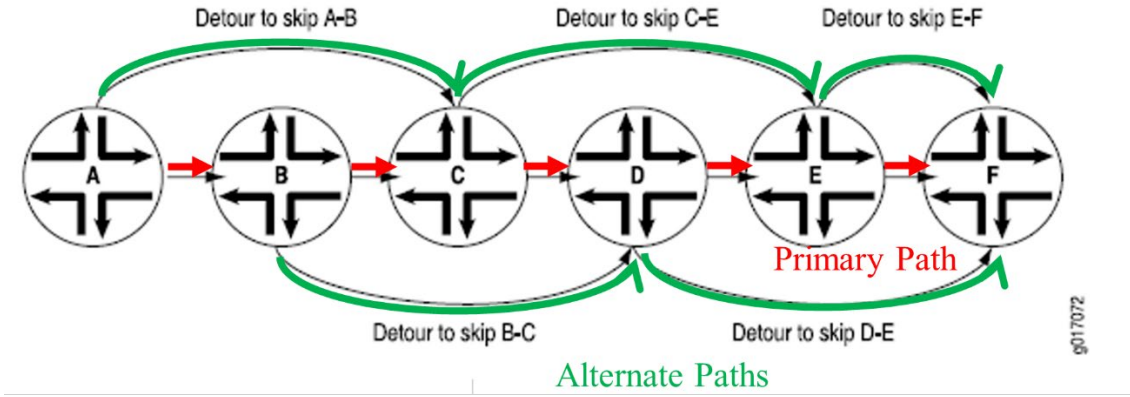
No.	'821 Patent Claim 10	Juniper Junos OS System
10	The method of claim 9, wherein said entity	Juniper Junos OS System discloses said entity cost function comprises a predefined metric.

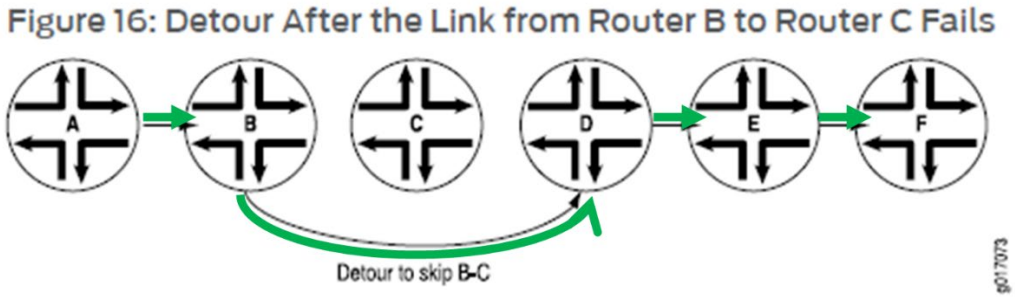
No.	'821 Patent Claim 10	Juniper Junos OS System
	cost function comprises a predefined metric.	*See Supra at* 6.

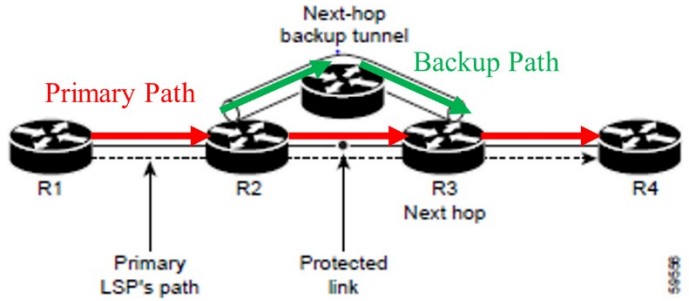
No.	'821 Patent Claim 11	Juniper Junos OS System
11	The method of claim 10, wherein said predefined metric is selected from the group consisting of interior gateway protocol (IGP) and traffic engineering (TE).	Juniper Junos OS System discloses said predefined entity cost metric is selected from the group consisting of interior gateway protocol (IGP) and traffic engineering (TE). *See Supra at* 7.

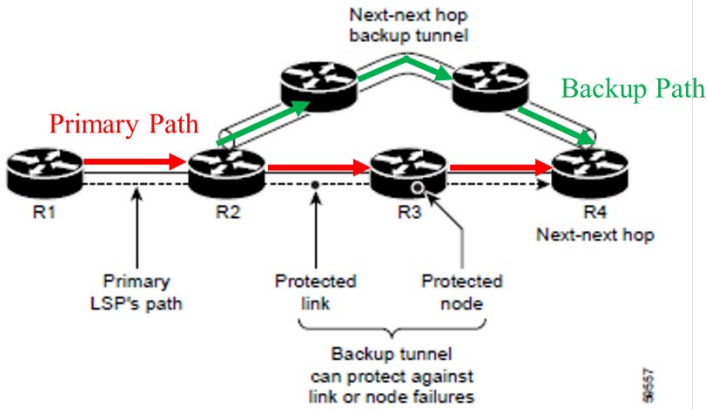
No.	'821 Patent Claim 12	Juniper Junos OS System
12	The method of claim 10, further comprising the step of configuring said working entity as revertive.	Juniper Junos OS System discloses the step of configuring said working entity as revertive. *See Supra at 8.*

No.	'821 Patent Claim 13	Juniper Junos OS System
13	The method of claim 1, further comprising the step of: if said entity pair reselection results in both working and protection entities being replaced, sequentially replacing said working entity	Juniper Junos OS System discloses the method of claim 1, further comprising the step of: the method of claim 1, further comprising the step of: if said entity pair reselection results in both working and protection entities being replaced, sequentially replacing said working entity and said protection entity. Based upon Orckit's apparent understanding of this limitation set forth in its Preliminary Infringement Disclosure, the Juniper Junos OS System meets this limitation. To the extent that the Juniper Junos OS System is found to not meet this limitation, if said entity pair reselection results in both working and protection entities being replaced, sequentially replacing said

No.	'821 Patent Claim 13	Juniper Junos OS System
	and said protection entity.	working entity and said protection entity would have been obvious to a person having ordinary skill in the art, as explained below. Juniper Junos OS Guide at Page No. 85 ("[Fast Reroute Overview Fast reroute provides redundancy for an LSP path. When you enable fast reroute, detours are precomputed and preestablished along the LSP. In case of a network failure on the current LSP path, traffic is quickly routed to one of the detours. Figure 15 on page 47 illustrates an LSP from Router A to Router F, showing the established detours. Each detour is established by an upstream node to avoid the link toward the immediate downstream node and the immediate downstream node itself. Each detour might traverse through one or more label-switched routers that are not shown in the figure. Fast reroute protects traffic against any single point of failure between the ingress and egress routers. If there are multiple failures along an LSP, fast reroute itself might fail. Also, fast reroute does not protect against failure of the ingress or egress routers. Figure 15: Detours Established for an LSP Using Fast Reroute  Fig. 15 (annotation added) If a node detects that a downstream link has failed (using a link-layer-specific liveness detection mechanism) or that a downstream node has failed (for example, using the RSVP

No.	'821 Patent Claim 13	Juniper Junos OS System
		neighbor hello protocol), the node quickly switches the traffic to the detour and, at the same time, signals the ingress router about the link or node failure. Figure 16 on page 47 illustrates the detour taken when the link between Router B and Router C fails.]")  Fig. 16 (annotation added) In addition, at least under the apparent claim scope alleged by Orckit's Infringement Disclosures, Juniper Junos OS System in combination with (1) the knowledge of a person of ordinary skill in the art, alone or in further combination with (2) each (individually, as well as one or more together) of the references identified in element 1[e] of Exhibit E-2 renders the claim, including the present limitation, obvious. Below are examples of two such references. For example, Cisco IOS Guide discloses Cisco IOS supports Fast reroute, link protection, and node protection in which upon failure of the LSP path, the bypass/backup tunnel is established, and when there are some changes in the previous backup tunnel a new backup tunnel is selected. Cisco IOS Guide further discloses failover within SRLG groups. Thus, according to Orckit's apparent understanding of this limitation in its Infringement Disclosures, this limitation is met. Cisco IOS Guide at Page No. 637 ("[Fast Reroute Fast Reroute (FRR) is a mechanism for protecting MPLS TE LSPs from link and node failures by locally repairing the LSPs at the point of failure, allowing data to continue to flow on them

No.	'821 Patent Claim 13	Juniper Junos OS System
		while their headend routers attempt to establish new end-to-end LSPs to replace them. FRR locally repairs the protected LSPs by rerouting them over backup tunnels that bypass failed links or node. Link Protection Backup tunnels that bypass only a single link of the LSP's path provide link protection. They protect LSPs if a link along their path fails by rerouting the LSP's traffic to the next hop (bypassing the failed link). These are referred to as next-hop (NHOP) backup tunnels because they terminate at the LSP's next hop beyond the point of failure. Figure 1 illustrates an NHOP backup tunnel.]") *Figure 1 NHOP Backup Tunnel*  Fig. 1 (annotation added) Cisco IOS Guide at Page No. 638 ("[Node Protection FRR provides node protection for LSPs. Backup tunnels that bypass next-hop nodes along LSP paths are called next-next-hop (NNHOP) backup tunnels because they terminate at the node following the next-hop node of the LSP paths, thereby bypassing the next-hop node. They protect LSPs if a node along their path fails by enabling the node upstream of the failure to reroute the LSPs and their traffic around the failed node to the next-next hop. FRR supports the use of RSVP Hellos to accelerate the detection of node failures. NNHOP backup tunnels also provide protection from link failures, because they bypass the failed link and the node.]")

No.	'821 Patent Claim 13	Juniper Junos OS System
		Figure 2 illustrates an NNHOP backup tunnel. ***Figure 2 NNHOP Backup Tunnel***  Fig. 2 (annotation added) Cisco IOS Guide at Page No. 638 ("[If an LSP is using a backup tunnel and something changes so that the LSP is no longer appropriate for the backup tunnel, the LSP is torn down. Such changes are the following: • Backup bandwidth of the backup tunnel is reduced. • Backup bandwidth type of backup tunnel is changed to a type that is incompatible with the primary LSP. • Primary LSP is modified so that FRR is disabled. (The no mpls traffic-eng fast-reroute command is entered.]") Cisco IOS Guide at Page No. 648 ("[After a backup tunnel has been chosen for an LSP, conditions may change that will cause us to reevaluate this choice. This reevaluation, if successful, is called promotion. Such conditions may include: 1. A new backup tunnel comes up. 2. The currently chosen backup tunnel for this LSP goes down.

No.	'821 Patent Claim 13	Juniper Junos OS System
		3. A backup tunnel's available backup bandwidth increases. For example, an LSP protected by the tunnel has been reoptimized by the headend to use another path. For cases 1 and 2, the LSP's backup tunnel is evaluated immediately. Case 3 is addressed by periodically reevaluating LSP-to-backup tunnel mappings. By default, background reevaluation is performed every 5 minutes. This interval is configurable via the mpls traffic-eng fast-reroute timers command.]") Cisco IOS Guide at Page No. 654 ("[Tunnel Reoptimization Procedure Tunnel reoptimization is the signaling of an LSP that is more optimal than the LSP a TE tunnel is currently using (for example, it may be shorter or may have a lower cost), and the switching over of the tunnel's data to use this new LSP. The new more optimal TE LSP is always established and the data moved onto it before the original LSP is torn down (so it is called the "make before break" procedure). This ensures that no data packets are lost during the transition to the new LSP.]") For example, RFC 4872 discloses a series of steps that can be used to replace a working and protection sequentially in response to a network fault. "12. Reversion Reversion refers to a recovery switching operation, where the normal traffic returns to (or remains on) the working LSP when it has recovered from the failure. Reversion implies that resources remain allocated to the LSP that was originally routed over them even after a failure. It is important to have mechanisms that allow reversion to be performed with minimal service disruption and reconfiguration. For "1+1 bidirectional Protection", reversion to the recovered LSP occurs by using the following sequence: 1. Clear the A bit of the ADMIN_STATUS object if set for the recovered LSP. 2. Then, apply the method described below to switch normal traffic back from the protecting to the recovered LSP. This is performed by using the new error code/sub-code "Notify Error/LSP Recovered" (Switchback Request). The procedure is as follows:

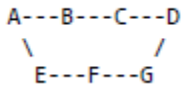
No.	'821 Patent Claim 13	Juniper Junos OS System
		1) The initiating (source) node sends the normal traffic onto both the working and the protecting LSPs. Once completed, the source node sends reliably a Notify message to the destination with the new error code/sub-code "Notify Error/LSP Recovered" (Switchback Request). This Notify message includes the MESSAGE_ID object. The ACK_Desired flag MUST be set in this object to request the receiver to send an acknowledgment for the message (see [RFC2961]). 2) Upon receipt of this message, the destination selects the traffic from the working LSP. At the same time, it transmits the traffic onto both the working and protecting LSP. The destination then sends reliably a Notify message to the source confirming the completion of the operation. This message includes the MESSAGE_ID_ACK object to acknowledge reception of the received Notify message. This Notify message also includes the MESSAGE_ID object. The ACK_Desired flag MUST be set in this object to request the receiver to send an acknowledgment for the message (see [RFC2961]). 3) When the source node receives this Notify message, it switches to receive traffic from the working LSP. The source node then sends an Ack message to the destination node confirming that the LSP has been reverted. 3. Finally, clear the O bit of the PROTECTION object sent over the protecting LSP. For "1:N Protection with Extra-traffic", reversion to the recovered LSP occurs by using the following sequence: 1. Clear the A bit of the ADMIN_STATUS object if set for the recovered LSP. 2. Then, apply the method described below to switch normal traffic back from the protecting to the recovered LSP. This is performed by using the new error code/sub-code "Notify Error/LSP Recovered" (Switchback Request). The procedure is as follows: 1) The initiating (source) node sends the normal traffic onto both the working and the protecting LSPs. Once completed, the source node sends reliably a Notify message to the destination with the new error code/sub-code "Notify Error/LSP Recovered" (Switchback Request). This Notify message includes the MESSAGE_ID object. The ACK_Desired flag MUST be set in this object to request the receiver to send an acknowledgment for the message (see [RFC2961]).

No.	'821 Patent Claim 13	Juniper Junos OS System
		2) Upon receipt of this message, the destination selects the traffic from the working LSP. At the same time, it transmits the traffic onto both the working and protecting LSP. The destination then sends reliably a Notify message to the source confirming the completion of the operation. This message includes the MESSAGE_ID_ACK object to acknowledge reception of the received Notify message. This Notify message also includes the MESSAGE_ID object. The ACK_Desired flag MUST be set in this object to request the receiver to send an acknowledgment for the message (see [RFC2961]). 3) When the source node receives this Notify message, it switches to receive traffic from the working LSP, and stops transmitting traffic on the protecting LSP. The source node then sends an Ack message to the destination node confirming that the LSP has been reverted. 4) Upon receipt of this message, the destination node stops transmitting traffic along the protecting LSP. 3. Finally, clear the O bit of the PROTECTION object sent over the protecting LSP. For "Rerouting without Extra-traffic" (including the shared recovery case), reversion implies that the formerly working LSP has not been torn down by the head-end node upon PathErr message reception, i.e., the head-end node kept refreshing the working LSP under failure condition. This ensures that the exact same resources are retrieved after reversion switching (except if the working LSP required re- signaling). Re-activation is performed using the following sequence: 1. Clear the A bit of the ADMIN_STATUS object if set for the recovered LSP. 2. Then, apply the method described below to switch normal traffic back from the protecting to the recovered LSP. This is performed by using the new error code/sub-code "Notify Error/LSP Recovered" (Switchback Request). The procedure is as follows: 1) The initiating (source) node sends the normal traffic onto both the working and the protecting LSPs. Once completed, the source node sends reliably a Notify message to the destination with the new error code/sub-code "Notify Error/LSP Recovered" (Switchback Request). This Notify message includes the MESSAGE_ID object. The ACK_Desired flag MUST be set in this object to request the receiver to send an acknowledgment for the message (see [RFC2961]). 2) Upon receipt of this message, the destination selects the traffic from the working LSP. At the same time, it transmits the traffic onto both the working and protecting LSP. The destination then sends reliably a Notify message to the source confirming the completion of

No.	'821 Patent Claim 13	Juniper Junos OS System
		the operation. This message includes the MESSAGE_ID_ACK object to acknowledge reception of the received Notify message. This Notify message also includes the MESSAGE_ID object. The ACK_Desired flag MUST be set in this object to request the receiver to send an acknowledgment for the message (see [RFC2961]). 3) When the source node receives this Notify message, it switches to receive traffic from the working LSP, and stops transmitting traffic on the protecting LSP. The source node then sends an Ack message to the destination node confirming that the LSP has been reverted. 4) Upon receipt of this message, the destination node stops transmitting traffic along the protecting LSP. 3. Finally, de-activate the protecting LSP by setting the S bit to 1 in the PROTECTION object sent over the protecting LSP." RFC 4872 at 26–29. "Note: this document assumes that Protection Type values have end-to- end significance and that the same value is sent over the protected and the protecting path. In this context, shared-mesh (for instance) appears from the end-nodes perspective as being simply an LSP rerouting without extra-traffic services. The net result of this is that a single bit (the S bit alone) does not allow determining whether resource allocation should be performed with respect to the status of the LSP within the protected entity. The introduction of the P bit solves this problem unambiguously. These bits MUST be processed on a hop-by-hop basis (independently of the LSP Protection Type context). This allows for an easier implementation of reversion signaling (see Section 12) but also facilitates the transparent delivery of protected services since any intermediate node is not required to know the semantics associated with the incoming LSP Protection Type value." RFC 4872 at 8–9. "RSVP-TE Extensions in Support of End-to-End Generalized Multi-Protocol Label Switching (GMPLS) Recovery" RFC 4872 Title. "1. Introduction Generalized Multi-Protocol Label Switching (GMPLS) extends MPLS to include support for Layer-2 Switch Capable (L2SC), Time-Division Multiplex (TDM), Lambda Switch Capable (LSC), and Fiber Switch Capable (FSC) interfaces. GMPLS recovery uses control plane mechanisms (i.e., signaling, routing, and link management mechanisms) to support data plane

No.	'821 Patent Claim 13	Juniper Junos OS System
		fault recovery. Note that the analogous (data plane) fault detection mechanisms are required to be present in support of the control plane mechanisms. In this document, the term "recovery" is generically used to denote both protection and restoration; the specific terms "protection" and "restoration" are only used when differentiation is required. The subtle distinction between protection and restoration is made based on the resource allocation done during the recovery phase (see [RFC4427]). A functional description of GMPLS recovery is provided in [RFC4426] and should be considered as a companion document. The present document describes the protocol-specific procedures for GMPLS RSVP- TE (Resource ReSerVation Protocol - Traffic Engineering) signaling (see [RFC3473]) to support end-to-end recovery. End-to-end recovery refers to the recovery of an entire LSP from its head-end (ingress node endpoint) to its tail-end (egress node endpoint). With end-to- end recovery, working LSPs are assumed to be resource-disjoint (where a resource is a link, node, or Shared Risk Link Group (SRLG)) in the network so that they do not share any failure probability, but this is not mandatory. With respect to a given set of network resources, a pair of working/protecting LSPs SHOULD be resource disjoint in case of dedicated recovery type (see below). On the other hand, in case of shared recovery (see below), a group of working LSPs SHOULD be mutually resource-disjoint in order to allow for a (single and commonly) shared protecting LSP, itself resource-disjoint from each of the working LSPs. Note that resource disjointness is a necessary (but not sufficient) condition to ensure LSP recoverability. The present document addresses four types of end-to-end LSP recovery: 1) 1+1 (unidirectional/bidirectional) protection, 2) 1:N (N >= 1) LSP protection with extra-traffic, 3) pre-planned LSP rerouting without extra-traffic (including shared mesh), and 4) full LSP rerouting. 1) The simplest notion of end-to-end LSP protection is 1+1 unidirectional protection. Using this type of protection, a protecting LSP is signaled over a dedicated resource-disjoint alternate path to protect an associated working LSP. Normal traffic is simultaneously sent on both LSPs and a selector is used at the egress node to receive traffic from one of the LSPs. If a failure occurs along one of the LSPs, the egress node selects the traffic from the valid LSP. No coordination is required between the end nodes when a failure/switchover occurs.

No.	'821 Patent Claim 13	Juniper Junos OS System
		In 1+1 bidirectional protection, a protecting LSP is signaled over a dedicated resource-disjoint alternate path to protect the working LSP. Normal traffic is simultaneously sent on both LSPs (in both directions), and a selector is used at both ingress/egress nodes to receive traffic from the same LSP. This requires coordination between the end-nodes when switching to the protecting LSP. 2) In 1:N (N >= 1) protection with extra-traffic, the protecting LSP is a fully provisioned and resource-disjoint LSP from the N working LSPs, that allows for carrying extra-traffic. The N working LSPs MAY be mutually resource-disjoint. Coordination between end-nodes is required when switching from one of the working LSPs to the protecting LSP. As the protecting LSP is fully provisioned, default operations during protection switching are specified for a protecting LSP carrying extra-traffic, but this is not mandatory. Note that M:N protection is out of scope of this document (though mechanisms it defines may be extended to cover it). 3) Pre-planned LSP rerouting (or restoration) relies on the establishment between the same pair of end-nodes of a working LSP and a protecting LSP that is link/node/SRLG disjoint from the working one. Here, the recovery resources for the protecting LSP are pre-reserved but explicit action is required to activate (i.e., commit resource allocation at the data plane) a specific protecting LSP instantiated during the (pre-)provisioning phase. Since the protecting LSP is not "active" (i.e., fully instantiated), it cannot carry any extra-traffic. This does not mean that the corresponding resources cannot be used by other LSPs. Therefore, this mechanism protects against working LSP(s) failure(s) but requires activation of the protecting LSP after working LSP failure occurrence. This requires restoration signaling along the protecting path. "Shared-mesh" restoration can be seen as a particular case of pre-planned LSP rerouting that reduces the recovery resource requirements by allowing multiple protecting LSPs to share common link and node resources. The recovery resources are pre-reserved but explicit action is required to activate (i.e., commit resource allocation at the data plane) a specific protecting LSP instantiated during the (pre-) provisioning phase. This procedure requires restoration signaling along the protecting path. Note that in both cases, bandwidth pre-reserved for a protecting (but not activated) LSP can be made available for carrying extra traffic. LSPs for extra-traffic (with lower holding priority than the protecting LSP) can then be established using the bandwidth

No.	'821 Patent Claim 13	Juniper Junos OS System
		pre-reserved for the protecting LSP. Also, any lower priority LSP that use the pre-reserved resources for the protecting LSP(s) must be preempted during the activation of the protecting LSP. 4) Full LSP rerouting (or restoration) switches normal traffic to an alternate LSP that is not even partially established until after the working LSP failure occurs. The new alternate route is selected at the LSP head-end node, it may reuse resources of the failed LSP at intermediate nodes and may include additional intermediate nodes and/or links." " RFC 4872 at 3–5. "Consider the following network topology:  The paths [A,B,C,D] and [A,E,F,G,D] are node and link disjoint, ignoring the ingress/egress nodes A and D. A 1+1 protected path is established from A to D over [A,B,C,D] and [A,E,F,G,D], and traffic is transmitted simultaneously over both component paths (i.e., LSPs). During the provisioning phase, both LSPs are fully instantiated (and thus activated) so that no resource sharing can be done along the protecting LSP (nor can any extra-traffic be transported). It is also RECOMMENDED to set the N bit since no protection-switching signaling is assumed in this case. When a failure occurs (say, at node B) and is detected at end-node D, the receiver at D selects the normal traffic from the other LSP. From this perspective, 1+1 unidirectional protection can be seen as an uncoordinated protection-switching mechanism acting independently at both endpoints. Also, for the LSP under failure condition, it is RECOMMENDED to not set the Path_State_Removed Flag of the ERROR_SPEC object (see [RFC3473]) upon PathErr message generation. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 9.

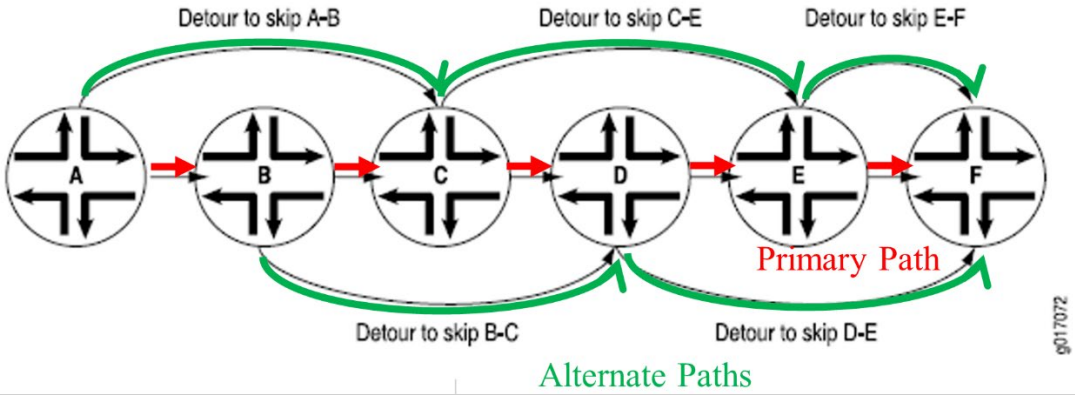
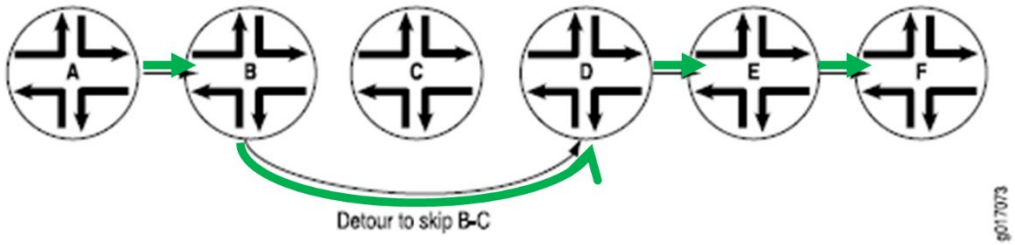
No.	'821 Patent Claim 13	Juniper Junos OS System
		"When a failure is detected by one or both endpoints of the LSP, both endpoints must select traffic from the other LSP. This action must be coordinated between node A and D. From this perspective, 1+1 bidirectional protection can be seen as a coordinated protection- switching mechanism between both endpoints. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 11. "(Full) LSP Rerouting LSP rerouting, on the other hand, switches normal traffic to an alternate LSP that is fully established only after failure occurrence. The new (alternate) route is selected at the LSP head- end and may reuse intermediate nodes included in the original route; it may also include additional intermediate nodes. For strict-hop routing, TE requirements can be directly applied to the route computation, and the failed node or link can be avoided. However, if the failure occurred within a loose-routed hop, the head-end node may not have enough information to reroute the LSP around the failure. Crankback signaling (see [CRANK]) and route exclusion techniques (see [RFC4874]) MAY be used in this case. The alternate route MAY be either computed on demand (that is, when the failure occurs; this is referred to as full LSP rerouting) or pre-computed and stored for use when the failure is reported. The latter offers faster restoration time. There is, however, a risk that the alternate route will become out of date through other changes in the network; this can be mitigated to some extent by periodic recalculation of idle alternate routes. (Full) LSP rerouting will be initiated by the head-end node that has either detected the LSP failure or received a Notify message and/or a PathErr message with the new error code/sub-code "Notify Error/LSP Locally Failed" for this LSP. The new LSP resources can be established using the make-before-break mechanism, where the new LSP is set up before the old LSP is torn down. This is done by using the mechanisms of the SESSION_ATTRIBUTE object and the Shared-Explicit (SE) reservation style (see [RFC3209]). Both the new and old LSPs can share resources at common nodes. Note that the make-before-break mechanism is not used to avoid disruption to the normal traffic flow (the latter has already been broken by the failure that is being repaired). However,

No.	'821 Patent Claim 13	Juniper Junos OS System
		it is valuable to retain the resources allocated on the original LSP that will be reused by the new alternate LSP." RFC 4872 at 25. "The procedure is as follows: 1) The initiating (source) node sends the normal traffic onto both the working and the protecting LSPs. Once completed, the source node sends reliably a Notify message to the destination with the new error code/sub-code "Notify Error/LSP Recovered" (Switchback Request). This Notify message includes the MESSAGE_ID object. The ACK_Desired flag MUST be set in this object to request the receiver to send an acknowledgment for the message (see [RFC2961]). 2) Upon receipt of this message, the destination selects the traffic from the working LSP. At the same time, it transmits the traffic onto both the working and protecting LSP. The destination then sends reliably a Notify message to the source confirming the completion of the operation. This message includes the MESSAGE_ID_ACK object to acknowledge reception of the received Notify message. This Notify message also includes the MESSAGE_ID object. The ACK_Desired flag MUST be set in this object to request the receiver to send an acknowledgment for the message (see [RFC2961]). 3) When the source node receives this Notify message, it switches to receive traffic from the working LSP. The source node then sends an Ack message to the destination node confirming that the LSP has been reverted." RFC 4872 at 27.

No.	'821 Patent Claim 14	Juniper Junos OS System
14[preamble]	A system for selecting entities within an MPLS network, comprising:	Juniper Junos OS System discloses a system for selecting entities within an MPLS network, comprising. *See generally* 19230_JNPR_00000071-00042088. *See Supra at* 1[preamble] and 1[a].
14[a]	a data structure comprising a plurality	Juniper Junos OS System discloses a data structure comprising a plurality of transport entity descriptors.

No.	'821 Patent Claim 14	Juniper Junos OS System
	of transport entity descriptors;	For example, Juniper Junos OS Guide discloses Juniper routers that generate traffic engineering through detailed knowledge about the network topology as well as dynamic information about network loading. A person of ordinary skill in the art would understand that each router maintains network link attributes and topology information in a specialized traffic engineering database which is used exclusively for calculating explicit paths for the placement of LSPs across the physical topology. Juniper Junos OS Guide at Page No. 42 ("[In the Junos OS, traffic engineering is implemented with MPLS and RSVP. Traffic engineering is composed of four functional components: • Packet Forwarding Component on page 4 • Information Distribution Component on page 5 • Path Selection Component on page 6 • Signaling Component on page 7]") Juniper Junos OS Guide at Page No. 43 ("[Information Distribution Component Traffic engineering requires detailed knowledge about the network topology as well as dynamic information about network loading. To implement the information distribution component, simple extensions to the IGPs are defined. Link attributes are included as part of each router's link-state advertisement. IS-IS extensions include the definition of new type length values (TLVs), whereas OSPF extensions are implemented with opaque link-state advertisements (LSAs). The standard flooding algorithm used by the link-state IGPs ensures that link attributes are distributed to all routers in the routing domain. Some of the traffic engineering extensions to be added to the IGP link-state advertisement include maximum link bandwidth, maximum reserved link bandwidth, current bandwidth reservation, and link coloring.]") Juniper Junos OS Guide at Page No. 43 ("[Each router maintains network link attributes and topology information in a specialized traffic engineering database. The traffic engineering database is used exclusively for calculating explicit paths for the placement of LSPs across the physical topology. A separate database is maintained so that the subsequent traffic engineering computation is independent of the IGP and the IGP's link-state database.

No.	'821 Patent Claim 14	Juniper Junos OS System
		Meanwhile, the IGP continues its operation without modification, performing the traditional shortest-path calculation based on information contained in the router's link-state database.]")
14[b]	an entity protection switch configured to switch between a working entity and a protection entity; and	Juniper Junos OS System discloses to switch between a working entity and a protection entity. For example, Juniper Junos OS Guide discloses the fast reroute feature that, when a node detects a downlink link/node has failed, the node quickly switches the traffic to the detour. A person of ordinary skill in the art would understand that the Juniper Router has digital logic configured to implement this feature. Thus, at least under the apparent claim scope alleged by Orckit's Infringement Disclosures, this limitation is met. Juniper Junos OS Guide at Page No. 85 ("[Fast Reroute Overview Fast reroute provides redundancy for an LSP path. When you enable fast reroute, detours are precomputed and preestablished along the LSP. In case of a network failure on the current LSP path, traffic is quickly routed to one of the detours. Figure 15 on page 47 illustrates an LSP from Router A to Router F, showing the established detours. Each detour is established by an upstream node to avoid the link toward the immediate downstream node and the immediate downstream node itself. Each detour might traverse through one or more label-switched routers that are not shown in the figure. Fast reroute protects traffic against any single point of failure between the ingress and egress routers. If there are multiple failures along an LSP, fast reroute itself might fail. Also, fast reroute does not protect against failure of the ingress or egress routers.

No.	'821 Patent Claim 14	Juniper Junos OS System
		Figure 15: Detours Established for an LSP Using Fast Reroute  Fig. 15 (annotation added) If a node detects that a downstream link has failed (using a link-layer-specific liveness detection mechanism) or that a downstream node has failed (for example, using the RSVP neighbor hello protocol), the node quickly switches the traffic to the detour and, at the same time, signals the ingress router about the link or node failure. Figure 16 on page 47 illustrates the detour taken when the link between Router B and Router C fails.]") Figure 16: Detour After the Link from Router B to Router C Fails  Fig. 16 (annotation added)

No.	'821 Patent Claim 14	Juniper Junos OS System
14[c]	digital logic configured to select said working entity and said protection entity from said plurality of transport entity descriptors, comprising:	Juniper Junos OS System discloses digital logic configured to select said working entity and said protection entity from said plurality of transport entity descriptors, comprising. For example, Reference discloses that the Juniper Junos OS System selects a primary path and a backup path from a plurality of potential paths. Juniper Junos OS Guide further discloses that through fate sharing, you can configure one or more backup paths that minimize the number of shared links with the primary paths to ensure that a single point of failure will not affect the primary and backup paths simultaneously. A person of ordinary skill in the art would understand that digital logic is configured to implement this functionality. Thus, at least under the apparent claim scope alleged by Orckit's Infringement Disclosures, this limitation is met. Juniper Junos OS Guide at Page No. 73 ("[Fate sharing allows you to create a database of information that CSPF uses to compute one or more backup paths to use in case the primary path becomes unstable. The database describes the relationships between elements of the network, such as routers and links. You can specify one or more elements within a group. Through fate sharing, you can configure backup paths that minimize the number of shared links and fiber paths with the primary paths as much as possible, to ensure that if a fiber is cut, the minimum amount of data is lost and a path still exists to the destination. For a backup path to work optimally, it must not share links or physical fiber paths with the primary path, ensuring that a single point of failure will not affect the primary and backup paths simultaneously. For more information about fate sharing, see the Junos OS Routing Protocols Configuration Guide.]") Juniper Junos OS Guide at Page No. 96 ("[For a backup path to work optimally, it must not share links or physical fiber paths with the primary path. This ensures that a single point of failure will not affect the primary and backup paths at the same time.]")
14[d]	logic configured to determine a probability of concurrent failure of said working entity	Juniper Junos OS System discloses determine a probability of concurrent failure of said working entity and said protection entity. *See Supra at* 14[c].

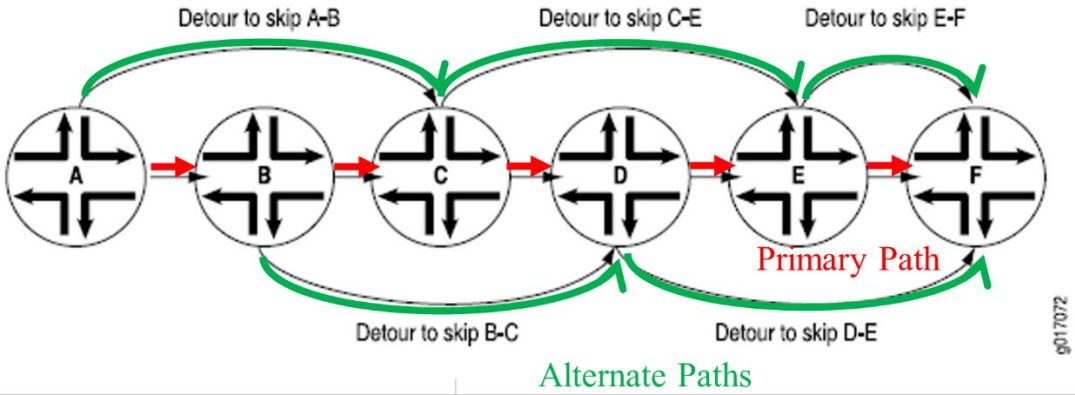
No.	'821 Patent Claim 14	Juniper Junos OS System
	and said protection entity;	
14[e]	logic configured to determine an entity cost of said plurality of transport entity descriptors: and	Juniper Junos OS System discloses determine an entity cost of said plurality of transport entity descriptors. For example, Juniper Junos OS Guide discloses detour computation and fast reroute optimization that rely on IGP supporting traffic engineering extensions and fate sharing. Juniper Junos OS guide further discloses that these functionalities rely upon the costs of transport entities. Juniper Junos OS Guide at Page No.85 ("[If a node detects that a downstream link has failed (using a link-layer-specific liveness detection mechanism) or that a downstream node has failed (for example, using the RSVP neighbor hello protocol), the node quickly switches the traffic to the detour and, at the same time, signals the ingress router about the link or node failure. Figure 16 on page 47 illustrates the detour taken when the link between Router B and Router C fails. Note that after the node switches traffic to the detour, it might switch the traffic again to a newly calculated detour soon after. This is because the initial detour route might not be the best route. To make rerouting as fast as possible, the node switches traffic onto the initial detour without first verifying that the detour is valid. Once the switch is made, the node recomputes the detour. If the node determines that the initial detour is still valid, traffic continues to flow over this detour. If the node determines that the initial detour is no longer valid, it again switches the traffic to a newly computed detour.]") Juniper Junos OS Guide at Page No. 87 ("[Detour Computations Computing and setting up detours is done independently at each node. On a node, if an LSP has fast reroute enabled and if a downstream link or node can be identified, the router performs a Constrained Shortest Path First (CSPF) computation using the information in the local traffic engineering database. For this reason, detours rely on your IGP supporting traffic engineering extensions. Without the traffic engineering database, detours cannot be established.]")

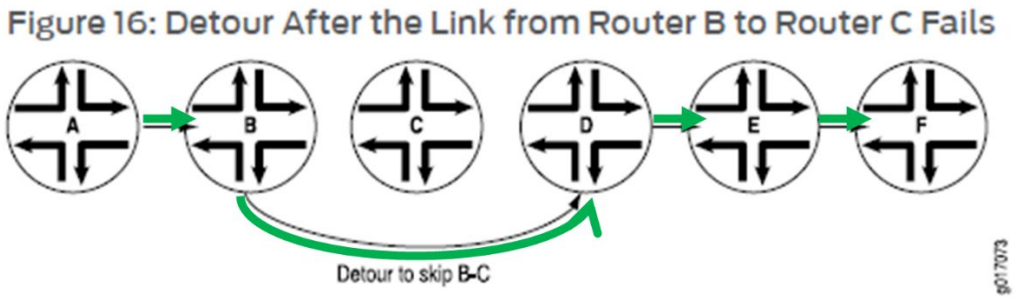
No.	'821 Patent Claim 14	Juniper Junos OS System
		Juniper Junos OS Guide at Page No. 97 ("[Fate-sharing groups contain three types of objects: • Point-to-point links—Identified by the IP addresses at each end of the link. Unnumbered point-to-point links are typically identified by borrowing IP addresses from other interfaces. Order is not important; from 1.2.3.4 to 1.2.3.5 and from 1.2.3.5 to 1.2.3.4 have the same meaning. • Non-point-to-point links—Include links on a LAN interface (such as Gigabit Ethernet interfaces) or nonbroadcast multiaccess (NBMA) interfaces (such as Asynchronous Transfer Mode [ATM] or Frame Relay). You identify these links by their individual interface address. For example, if the LAN interface 192.168.200.0/24 has four routers attached to it, each router link is individually identified: from 192.168.200.1; # LAN interface of router 1 from 192.168.200.2; # LAN interface of router 2 from 192.168.200.3; # LAN interface of router 3 from 192.168.200.4; # LAN interface of router 4 You can list the addresses in any order. • A router node—Identified by its configured router ID. An object can be in any number of groups, and a group can contain any number of objects. Each group has a configurable cost attributed to it, which represents the level of impact this group has on CSPF computations. The higher the cost, the less likely a backup path will share with the primary path any objects in the group. The cost is directly comparable to traffic engineering metrics. By default, the cost is 1. Changing the fate-sharing database does not affect established LSPs until the next reoptimization of CSPF. The fate-sharing database does influence fast-reroute computations.]") Juniper Junos OS Guide at Page No. 97 ("[When CSPF computes the primary paths of an LSP (or secondary paths when the primary path is not active), it ignores the fate-sharing information. You always want to find the best possible path (least IGP cost) for the primary path. When CSPF computes a secondary path while the primary path (of the same LSP) is active, the following occurs: 1. CSPF identifies all fate-sharing groups that are associated with the primary path. CSPF does this by identifying all links and nodes that the primary path traverses and compiling

No.	'821 Patent Claim 14	Juniper Junos OS System
		group lists that contain at least one of the links or nodes. CSPF ignores the ingress and egress nodes in the search. 2. CSPF checks each link in the traffic engineering database against the compiled group list. If the link is a member of a group, the cost of the link is increased by the cost of the group. If a link is a member of multiple groups, all group costs are added together. 3. CSPF performs the check for every node in the traffic engineering database, except the ingress and egress node. Again, a node can belong to multiple groups, so costs are additive. 4. The router performs regular CSPF computation with the adjusted topology.]") Juniper Junos OS Guide at 146 ("[Optimizing Signaled LSPs Once an LSP has been established, topology or resources changes might, over time, make the path suboptimal. A new path might have become available that is less congested, has a lower metric, and traverses fewer hops. You can configure the router to recompute paths periodically to determine whether a more optimal path has become available.]") Juniper Junos OS Guide at Page No. 147 ("[If reoptimization is enabled, an LSP can be rerouted through different paths by constrained-path recomputations. However, if reoptimization is disabled, the LSP has a fixed path and cannot take advantage of newly available network resources. The LSP is fixed until the next topology change breaks the LSP and forces a recomputation. Reoptimization is not related to failover. A new path is always computed when topology failures occur that disrupt an established path.]") Juniper Junos OS Guide at Page No.147 ("[After reoptimization is run, the result is accepted only if it meets the following criteria: 1. The new path is not higher in IGP metric. (The metric for the old path is updated during computation, so if a recent link metric changed somewhere along the old path, it is accounted for.) 2. If the new path has the same IGP metric, it is not more hops away. 3. The new path does not cause preemption. (This is to reduce the ripple effect of preemption causing more preemption.)

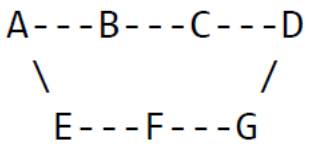
No.	'821 Patent Claim 14	Juniper Junos OS System
		4. The new path does not worsen congestion overall.]")
14[f]	logic configured to reselect said working entity and said protection entity from said plurality of transport entity descriptors upon a reselection event,	Juniper Junos OS System discloses reselect protection entity from said plurality of transport entity descriptors upon a reselection event. *See Supra at* 1[d]
14[g]	wherein said reselection event is selected from a group consisting of adding an entity to said plurality of transport entities, removing an entity from said plurality of transport entities, an operational status change for one of said plurality of transport entities, and a change in overall cost for one of said plurality of transport entities.	Juniper Junos OS System discloses wherein said reselection event is selected from a group consisting of an entity to said plurality of transport entities, removing an entity from said plurality of transport entities, an operational status change for one of said plurality of transport entities, and a change in overall cost for one of said plurality of transport entities. *See Supra at* 1[e].

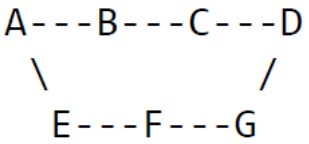
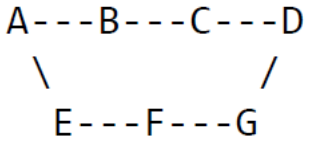
No.	'821 Patent Claim 15	Juniper Junos OS System
15	The system of claim 14 wherein said entity protection switch	Juniper Junos OS System discloses the system of claim 14 wherein said entity protection switch comprises a 1:1 switch.

No.	'821 Patent Claim 15	Juniper Junos OS System
	comprises a 1:1 switch.	At least based upon Orckit's apparent understanding of the claim scope as reflected in its Infringement Disclosures, this limitation is met. To the extent that the Juniper Junos OS System is found to not meet this limitation, wherein said entity protection switch comprises a 1:1 switch would have been obvious to a person having ordinary skill in the art, as explained below. Juniper Junos OS Guide at Page No. 85 ("[Fast Reroute Overview Fast reroute provides redundancy for an LSP path. When you enable fast reroute, detours are precomputed and preestablished along the LSP. In case of a network failure on the current LSP path, traffic is quickly routed to one of the detours. Figure 15 on page 47 illustrates an LSP from Router A to Router F, showing the established detours. Each detour is established by an upstream node to avoid the link toward the immediate downstream node and the immediate downstream node itself. Each detour might traverse through one or more label-switched routers that are not shown in the figure. Fast reroute protects traffic against any single point of failure between the ingress and egress routers. If there are multiple failures along an LSP, fast reroute itself might fail. Also, fast reroute does not protect against failure of the ingress or egress routers. Figure 15: Detours Established for an LSP Using Fast Reroute  Fig. 15 (annotation added)

No.	'821 Patent Claim 15	Juniper Junos OS System
		If a node detects that a downstream link has failed (using a link-layer-specific liveness detection mechanism) or that a downstream node has failed (for example, using the RSVP neighbor hello protocol), the node quickly switches the traffic to the detour and, at the same time, signals the ingress router about the link or node failure. Figure 16 on page 47 illustrates the detour taken when the link between Router B and Router C fails.]")  Fig. 16 (annotation added) In addition, at least under the apparent claim scope alleged by Orckit's Infringement Disclosures, Juniper Junos OS System in combination with (1) the knowledge of a person of ordinary skill in the art, alone or in further combination with (2) each (individually, as well as one or more together) of the references identified in element 15 of Exhibit E-2 renders the claim, including the present limitation, obvious. Below are examples of two such references. For example, RFC 4872 discloses entity protection switches operable under both a 1:1 and 1+1 architecture. "7. 1:1 Protection with Extra-Traffic The most common case of end-to-end 1:N protection is to establish, between the same endpoints, an end-to-end working LSP (thus, N = 1) and a dedicated end-to-end protecting LSP that are mutually link/ node/SRLG disjoint. This protects against working LSP failure(s). The protecting LSP is used for switchover when the working LSP fails. GMPLS RSVP-TE signaling allows for the pre-provisioning of protecting LSPs by indicating in the Path message

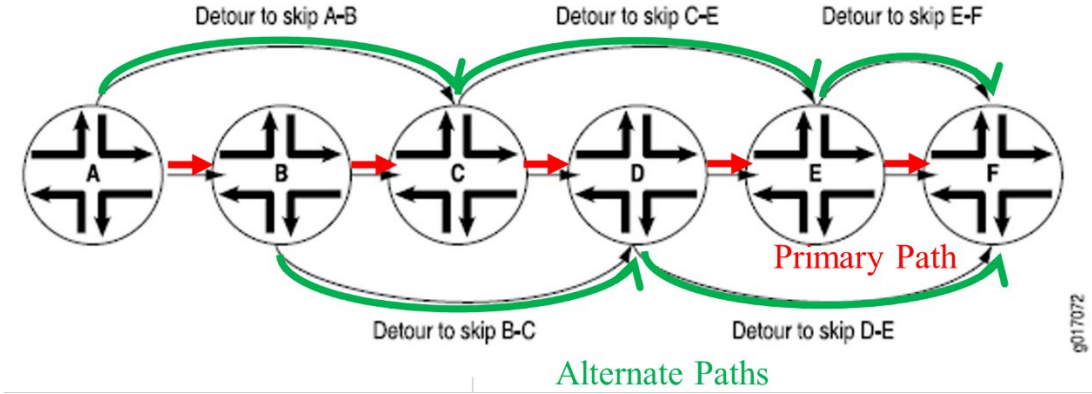
No.	'821 Patent Claim 15	Juniper Junos OS System
		(in the PROTECTION object; see Section 14) that the LSPs are of type protecting. Here, working and protecting LSPs are signaled as primary LSPs; both are fully instantiated during the provisioning phase. Although the resources for the protecting LSP are pre-allocated, preemptable traffic may be carried end-to-end using this LSP. Thus, the protecting LSP is capable of carrying extra-traffic with the caveat that this traffic will be preempted if the working LSP fails. The setup of the working LSP SHOULD indicate that the LSP head-end and tail-end node wish to receive Notify messages using the NOTIFY REQUEST object. The node upstream to the failure (upstream in terms of the direction an Path message traverses) SHOULD send a Notify message to the LSP head-end node, and the node downstream to the failure SHOULD send an Notify message to the LSP tail-end node. Upon receipt of the Notify messages, both the end-nodes MUST switch the (normal) traffic from the working LSP to the pre-configured protecting LSP (see Section 7.2). Moreover, some coordination is required if extra-traffic is carried over the end-to-end protecting LSP. Note that if the working and the protecting LSP are established between the same end-nodes, no further notification is required to indicate that the working LSPs are no longer protected. Consider the following topology: The working LSP [A,B,C,D] could be protected by the protecting LSP [A,E,F,G,D]. Both LSPs are fully instantiated (resources are allocated for both working and protecting LSPs) and no resource sharing can be done along the protection path since the primary protecting LSP can carry extra-traffic. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 13–14. "6. 1+1 Bidirectional Protection 1+1 bidirectional protection is a scheme that provides end-to-end protection for bidirectional LSPs. Consider the following network topology:

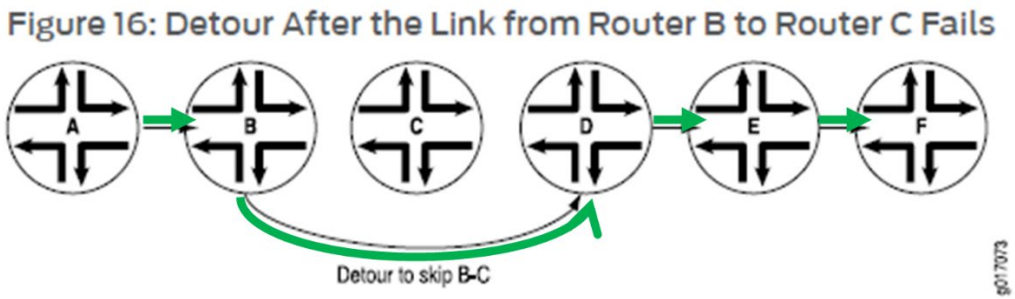
No.	'821 Patent Claim 15	Juniper Junos OS System
		 The LSPs [A,B,C,D] and [A,E,F,G,D] are node and link disjoint, ignoring the ingress/egress nodes A and D. A bidirectional LSP is established from A to D over each path, and traffic is transmitted simultaneously over both LSPs. In this scheme, both endpoints must receive traffic over the same LSP. Note also that both LSPs are fully instantiated (and thus activated) so that no resource sharing can be done along the protection path (nor can any extra-traffic be transported). When a failure is detected by one or both endpoints of the LSP, both endpoints must select traffic from the other LSP. This action must be coordinated between node A and D. From this perspective, 1+1 bidirectional protection can be seen as a coordinated protection- switching mechanism between both endpoints. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 10–11. For example, RFC 4872 discloses entity protection switches operable under both a 1:1 and 1+1 architecture. "7. 1:1 Protection with Extra-Traffic The most common case of end-to-end 1:N protection is to establish, between the same endpoints, an end-to-end working LSP (thus, N = 1) and a dedicated end-to-end protecting LSP that are mutually link/ node/SRLG disjoint. This protects against working LSP failure(s). The protecting LSP is used for switchover when the working LSP fails. GMPLS RSVP-TE signaling allows for the pre-provisioning of protecting LSPs by indicating in the Path message (in the PROTECTION object; see Section 14) that the LSPs are of type protecting. Here, working and protecting LSPs are signaled as primary LSPs; both are fully instantiated during the provisioning phase. Although the resources for the protecting LSP are pre-allocated, preemptable traffic may be carried end-to-end using this LSP. Thus, the protecting LSP is capable of carrying extra-traffic with the caveat that this traffic will be preempted if the working LSP fails. The setup of the working LSP SHOULD indicate that the LSP head-end

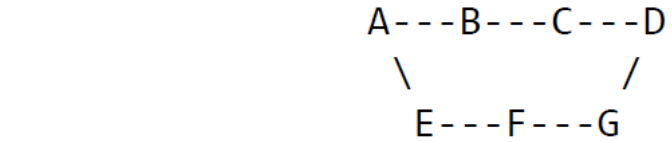
No.	'821 Patent Claim 15	Juniper Junos OS System
		and tail-end node wish to receive Notify messages using the NOTIFY REQUEST object. The node upstream to the failure (upstream in terms of the direction an Path message traverses) SHOULD send a Notify message to the LSP head-end node, and the node downstream to the failure SHOULD send an Notify message to the LSP tail-end node. Upon receipt of the Notify messages, both the end-nodes MUST switch the (normal) traffic from the working LSP to the pre-configured protecting LSP (see Section 7.2). Moreover, some coordination is required if extra-traffic is carried over the end-to-end protecting LSP. Note that if the working and the protecting LSP are established between the same end-nodes, no further notification is required to indicate that the working LSPs are no longer protected. Consider the following topology:  The working LSP [A,B,C,D] could be protected by the protecting LSP [A,E,F,G,D]. Both LSPs are fully instantiated (resources are allocated for both working and protecting LSPs) and no resource sharing can be done along the protection path since the primary protecting LSP can carry extra-traffic. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 13–14. "6. 1+1 Bidirectional Protection 1+1 bidirectional protection is a scheme that provides end-to-end protection for bidirectional LSPs. Consider the following network topology:  The LSPs [A,B,C,D] and [A,E,F,G,D] are node and link disjoint, ignoring the ingress/egress nodes A and D. A bidirectional LSP is established from A to D over each path, and traffic is

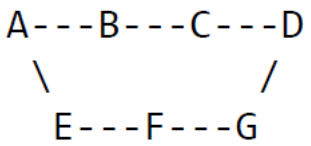
No.	'821 Patent Claim 15	Juniper Junos OS System
		transmitted simultaneously over both LSPs. In this scheme, both endpoints must receive traffic over the same LSP. Note also that both LSPs are fully instantiated (and thus activated) so that no resource sharing can be done along the protection path (nor can any extra-traffic be transported). When a failure is detected by one or both endpoints of the LSP, both endpoints must select traffic from the other LSP. This action must be coordinated between node A and D. From this perspective, 1+1 bidirectional protection can be seen as a coordinated protection- switching mechanism between both endpoints. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 10–11. For example, Cisco IOS Guide discloses when there is a failure on the protected LSP, the headend router immediately enables the secondary LSP to temporarily carry the tunnel's traffic (1:1 protection). Cisco IOS Guide Page No. 603 ("[Path Protection Path protection provides an end-to-end failure recovery mechanism (that is, full path protection) for MPLS TE tunnels. A secondary LSP is established, in advance, to provide failure protection for the protected LSP that is carrying a tunnel's TE traffic. When there is a failure on the protected LSP, the headend router immediately enables the secondary LSP to temporarily carry the tunnel's traffic. If there is a failure on the secondary LSP, the tunnel no longer has path protection until the failure along the secondary path is cleared.]") Cisco IOS Guide at Page No. 641 ("[One-to-one protection is when a separate backup tunnel must be used for each LSP needing protection. N:1 protection has significant scalability advantages over one-to-one (1:1) protection. An example of 1:1 protection is when 5000 backup tunnels protect 5000 LSPs, each router along the backup path must maintain state for an additional 5000 tunnels.]")

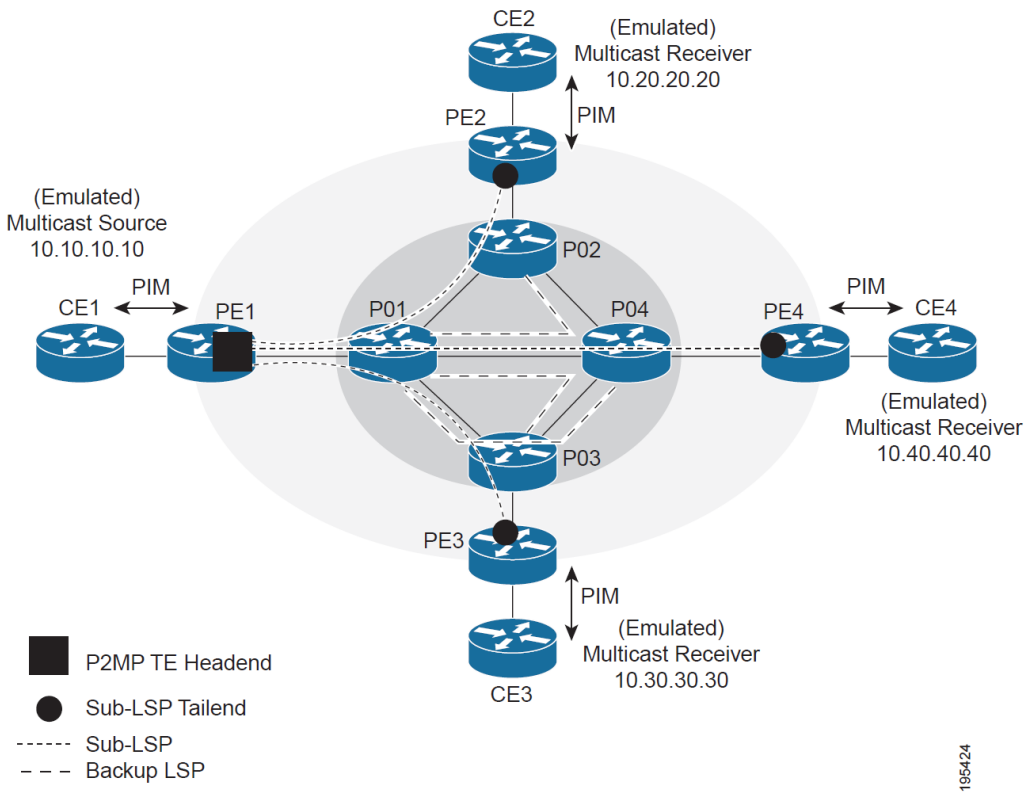
No.	'821 Patent Claim 16	Juniper Junos OS System
16	The system of claim 14 wherein said entity	Juniper Junos OS System discloses the system of claim 14 wherein said entity protection switch comprises a 1+1 switch.

No.	'821 Patent Claim 16	Juniper Junos OS System
	protection switch comprises a 1+1 switch.	At least based upon Orckit's apparent understanding of the claim scope as reflected in its Infringement Disclosures, this limitation is met. To the extent that the Juniper Junos OS System is found to not meet this limitation, wherein said entity protection switch comprises a 1+1 switch would have been obvious to a person having ordinary skill in the art, as explained below. Juniper Junos OS Guide at Page No. 85 ("[Fast Reroute Overview Fast reroute provides redundancy for an LSP path. When you enable fast reroute, detours are precomputed and preestablished along the LSP. In case of a network failure on the current LSP path, traffic is quickly routed to one of the detours. Figure 15 on page 47 illustrates an LSP from Router A to Router F, showing the established detours. Each detour is established by an upstream node to avoid the link toward the immediate downstream node and the immediate downstream node itself. Each detour might traverse through one or more label-switched routers that are not shown in the figure. Fast reroute protects traffic against any single point of failure between the ingress and egress routers. If there are multiple failures along an LSP, fast reroute itself might fail. Also, fast reroute does not protect against failure of the ingress or egress routers. Figure 15: Detours Established for an LSP Using Fast Reroute  Fig. 15 (annotation added)

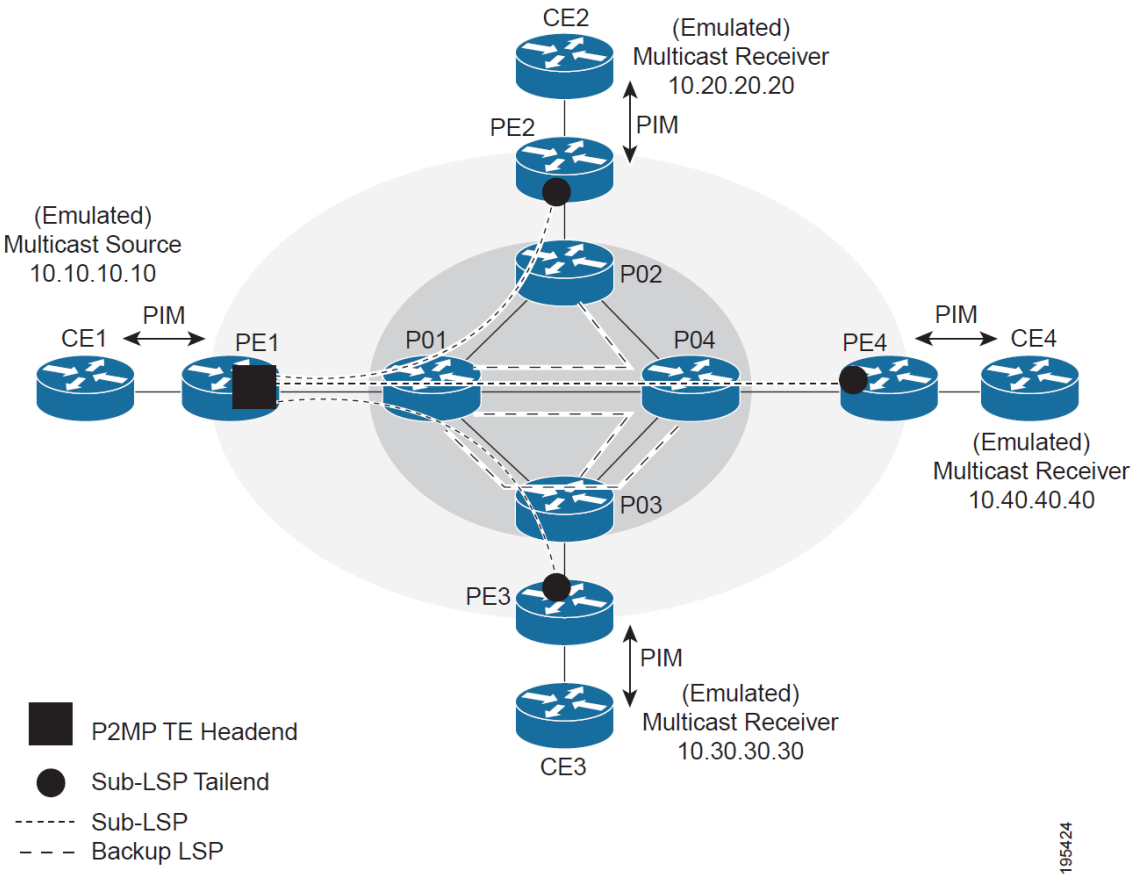
No.	'821 Patent Claim 16	Juniper Junos OS System
		If a node detects that a downstream link has failed (using a link-layer-specific liveness detection mechanism) or that a downstream node has failed (for example, using the RSVP neighbor hello protocol), the node quickly switches the traffic to the detour and, at the same time, signals the ingress router about the link or node failure. Figure 16 on page 47 illustrates the detour taken when the link between Router B and Router C fails.]")  Fig. 16 (annotation added) In addition, at least under the apparent claim scope alleged by Orckit's Infringement Disclosures, Juniper Junos OS System in combination with (1) the knowledge of a person of ordinary skill in the art, alone or in further combination with (2) each (individually, as well as one or more together) of the references identified in element 16 of Exhibit E-2 renders the claim, including the present limitation, obvious. Below are examples of two such references. For example, RFC 4872 discloses entity protection switches operable under both a 1:1 and 1+1 architecture. "7. 1:1 Protection with Extra-Traffic The most common case of end-to-end 1:N protection is to establish, between the same endpoints, an end-to-end working LSP (thus, N = 1) and a dedicated end-to-end protecting LSP that are mutually link/ node/SRLG disjoint. This protects against working LSP failure(s). The protecting LSP is used for switchover when the working LSP fails. GMPLS RSVP-TE

No.	'821 Patent Claim 16	Juniper Junos OS System
		signaling allows for the pre-provisioning of protecting LSPs by indicating in the Path message (in the PROTECTION object; see Section 14) that the LSPs are of type protecting. Here, working and protecting LSPs are signaled as primary LSPs; both are fully instantiated during the provisioning phase. Although the resources for the protecting LSP are pre-allocated, preemptable traffic may be carried end-to-end using this LSP. Thus, the protecting LSP is capable of carrying extra-traffic with the caveat that this traffic will be preempted if the working LSP fails. The setup of the working LSP SHOULD indicate that the LSP head-end and tail-end node wish to receive Notify messages using the NOTIFY REQUEST object. The node upstream to the failure (upstream in terms of the direction an Path message traverses) SHOULD send a Notify message to the LSP head-end node, and the node downstream to the failure SHOULD send an Notify message to the LSP tail-end node. Upon receipt of the Notify messages, both the end-nodes MUST switch the (normal) traffic from the working LSP to the pre-configured protecting LSP (see Section 7.2). Moreover, some coordination is required if extra-traffic is carried over the end-to-end protecting LSP. Note that if the working and the protecting LSP are established between the same end-nodes, no further notification is required to indicate that the working LSPs are no longer protected. Consider the following topology:  The working LSP [A,B,C,D] could be protected by the protecting LSP [A,E,F,G,D]. Both LSPs are fully instantiated (resources are allocated for both working and protecting LSPs) and no resource sharing can be done along the protection path since the primary protecting LSP can carry extra-traffic. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 13–14. "6. 1+1 Bidirectional Protection 1+1 bidirectional protection is a scheme that provides end-to-end protection for bidirectional LSPs. Consider the following network topology:

No.	'821 Patent Claim 16	Juniper Junos OS System
		 The LSPs [A,B,C,D] and [A,E,F,G,D] are node and link disjoint, ignoring the ingress/egress nodes A and D. A bidirectional LSP is established from A to D over each path, and traffic is transmitted simultaneously over both LSPs. In this scheme, both endpoints must receive traffic over the same LSP. Note also that both LSPs are fully instantiated (and thus activated) so that no resource sharing can be done along the protection path (nor can any extra-traffic be transported). When a failure is detected by one or both endpoints of the LSP, both endpoints must select traffic from the other LSP. This action must be coordinated between node A and D. From this perspective, 1+1 bidirectional protection can be seen as a coordinated protection- switching mechanism between both endpoints. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 10–11. For example, Cisco's IOS Guide discloses: "Figure 2 shows a P2MP TE tunnel that has three destinations. • PE1 is the headend router. • P01 is a branch point router, where packet replication occurs. • PE2, PE3, and PE4 are tailend routers, where the sub-LSP ends. Between the PE and CE routers, PIM is enabled to exchange multicast routing information with the directly connected customer edge (CE) routers. PIM is not enabled across the P2MP TE tunnel." Cisco IOS Guide at Page No. 468.

No.	'821 Patent Claim 16	Juniper Junos OS System
		Figure 2 Network Topology with P2MP TE Tunnel  Cisco IOS Guide at Page No. 469, FIG. 2. "RSVP TE extensions defined in RFC 4875 allow multiple sub-LSPs to be signaled from the headend router. A P2MP TE tunnel consists of multiple sub-LSPs that connect the headend router to various tailend routers. The headend router sends one RSVP path message to each destination. The tailend router replies with a RESV message. The Label Forwarding Information Base (LFIB) is populated using the RSVP labels allocated by the RESV messages.

No.	'821 Patent Claim 16	Juniper Junos OS System
		The tailend routers allocate unreserved labels, which are greater than 15 and do not include implicit or explicit null labels. Using unreserved labels allows IP Multicast to perform a Reverse Path Forwarding (RPF) check on the tailend router. Because a sub-LSP tailend router cannot be represented as a regular interface, a special LSP virtual interface (VIF) is automatically created. The LSP VIF represents the originating interface for all IP multicast traffic originating from the P2MP TE tailend router." Cisco IOS Guide at Page No. 469. "You can configure P2P explicit backup tunnels on point of local repair (PLR) nodes for link protection of P2MP sub-LSPs, similar to LSPs for P2P TE tunnels. You can also enable automatic creation of backup tunnels using the Auto-tunnel Backup feature for P2P TE tunnels. All sibling sub-LSPs that share the same outgoing link are protected by the same backup tunnel. All cousin sub-LSPs that share the same outgoing link can be protected by multiple P2P backup tunnels. Link protection for a P2MP TE tunnel is illustrated in Figure 6, which shows PE1 as the tunnel headend router and PE2, PE3, and PE4 as tunnel tailend routers. The following sub-LSPs are signaled from PE1 in the network: • From PE1 to PE2, the sub-LSP travels the following path: PE1 -> P01 -> P02 -> PE2 • From PE1 to PE3, the sub-LSP travels the following path: PE1 -> P01 -> P03 -> PE3 • From PE1 to PE4, the sub-LSP travels the following path: PE1 -> P01 -> P04 -> PE4 Node P01 is a branch node that does packet replication in the MPLS forwarding plane; ingress traffic originating from PE1 will be replicated towards routers P02, P03, and P04." Cisco IOS Guide at Page No. 474-475.

No.	'821 Patent Claim 16	Juniper Junos OS System
		Figure 6 P2MP TE Link Protection Example  Cisco IOS Guide at Page No. 475.

No.	'821 Patent Claim 17	Juniper Junos OS System
17[preamble]	Non-transitory computer readable media configured to	Juniper Junos OS System discloses non-transitory computer readable media configured to perform a method comprising the steps of. *See generally* 19230_JNPR_00000071-00042088.

No.	'821 Patent Claim 17	Juniper Junos OS System
	perform a method comprising the steps of:	For example, Juniper Junos OS Guide discloses Junos OS MPLS Application Configuration Guide for M Series, MX Series, T Series, EX Series, or J Series Juniper Networks routers. A person of ordinary skill in the art would understand that these routers include non-transitory computer readable media configured to perform the traffic engineering and fate sharing schemes described above. Juniper Junos OS Guide at Page No. 32 ("[This guide provides an overview of the MPLS applications functions of the Junos OS and describes how to configure MPLS applications on the router.]") Juniper Junos OS Guide at Page No. 32 ("[This guide is designed for network administrators who are configuring and monitoring a Juniper Networks M Series, MX Series, T Series, EX Series, or J Series router or switch.]") Juniper Junos OS provides network monitoring for label switched paths. *See, e.g.*, 19230_JNPR_00038164; 19230_JNPR_00038168.
17[a]	providing a plurality of MPLS transport entities between a first endpoint and a second endpoint;	Juniper Junos OS System discloses providing a plurality of MPLS transport entities between a first endpoint and a second endpoint. *See Supra at* 1[a].
17[b]	determining an overall cost for each entity pair of said plurality of entities;	Juniper Junos OS System discloses determining an overall cost for each entity pair of said plurality of entities. *See Supra at* 1[b].
17[c]	selecting an entity pair from said plurality of transport entities based at least in part upon said overall cost; and	Juniper Junos OS System discloses selecting an entity pair from said plurality of transport entities based at least in part upon said overall cost. *See Supra at* 1[c].

No.	'821 Patent Claim 17	Juniper Junos OS System
17[d]	if an entity pair reselection event occurs, reselecting said entity pair from the group consisting of said entity pair and a replacement entity pair comprising at least one entity distinct from the entities of said entity pair,	Juniper Junos OS System discloses if an entity pair reselection event occurs, reselecting said entity pair from the group consisting of said entity pair and a replacement entity pair comprising at least one entity distinct from the entities of said entity pair. For example, Juniper Junos OS Guide discloses once an LSP has been established, topology or resources changes might, over time, make the path suboptimal. A new path might have become available that is less congested, has a lower metric, and traverses fewer hops, the router is configured to recompute paths periodically to determine whether a more optimal path has become available. *See Supra at* 1[d].
17[e]	wherein said entity pair reselection event is selected from a group consisting of adding an entity to said plurality of transport entities, removing an entity from said plurality of transport entities, an operational status change for one of said plurality of transport entities, and a change in overall cost for one of said plurality of transport entities.	Juniper Junos OS System discloses wherein said entity pair reselection event is selected from a group consisting of an operational status change for one of said plurality of transport entities, and a change in overall cost for one of said plurality of transport entities. *See Supra at* 1[e].

No.	'821 Patent Claim 18	Juniper Junos OS System
18[preamble]	The non-transitory computer readable media of claim 17, wherein said step of selecting an entity pair further comprises:	Juniper Junos OS System discloses wherein said step of selecting an entity pair further comprises. *See Supra at* 1[Preamble] and 1[c].
18[a]	selecting a working entity from said plurality of transport entities;	Juniper Junos OS System discloses selecting a working entity from said plurality of transport entities. *See Supra at* 2[a].
18[b]	selecting a protection entity from said plurality of transport entities; and	Juniper Junos OS System discloses selecting a protection entity from said plurality of transport entities. *See Supra at* 2[b].
18[c]	selecting an active entity from the set consisting of said working entity and said protection entity.	Juniper Junos OS System discloses selecting an active entity from the set consisting of said working entity and said protection entity. *See Supra at* 3.

No.	'821 Patent Claim 19	Juniper Junos OS System
19	The non-transitory readable media of claim 18, wherein said step of selecting an entity pair further comprises minimizing an overall cost function.	Juniper Junos OS System discloses step of selecting an entity pair further comprises minimizing an overall cost function. *See Supra at* 4.

No.	'821 Patent Claim 20	Juniper Junos OS System
20[preamble]	The non-transitory readable media of claim 19, wherein said overall cost function comprises:	Juniper Junos OS System discloses said overall cost function comprises. *See Supra at* 5.
20[a]	minimizing a probability of concurrent failure of said protection entity and said working entity; and	Juniper Junos OS System discloses minimizing a probability of concurrent failure of said protection entity and said working entity. *See Supra at* 5.
20[b]	a predefined metric selected from the group consisting of interior gateway protocol (IGP) and traffic engineering (TE).	Juniper Junos OS System discloses a predefined metric selected from the group consisting of interior gateway protocol (IGP) and traffic engineering (TE). *See Supra at* 7.

EXHIBIT B-3
Defendant's First Amended Invalidity Contentions
Orckit Corporation v. Cisco Systems, Inc., 2:22-cv-00276-JRG-RSP

Chart for U.S. Patent 8,830,821 ("the '821 Patent")
MPLS-TP IETF Standards ("IETF MPLS-TP System")

As shown in the chart below, all claims of the '821 Patent are invalid (1) under 35 U.S.C. §§ 102 (a) and (g) based on IETF MPLS-TP System, (2) under 35 U.S.C. §§ 102 (a) and (b) because the references describing the IETF MPLS-TP System disclose every limitation of every Asserted Claim, and (3) under 35 U.S.C. § 103 as obvious based on RFC 4872, RSVP-TE Extensions in Support of End-to-End Generalized Multi-Protocol Label Switching (GMPLS) Recovery alone, or in combination with the knowledge of a person of ordinary skill in the art alone, and in further combination with references specifically identified in the following claim chart and/or one or more other references identified in Cisco's Preliminary Invalidity Contentions. The IETF MPLS-TP System comprises switches and routers from before June 22, 2011 that implemented the described standards. The IETF MPLS-TP System was invented, known and used by others in the United States, and on sale in the United States before the claimed invention, rendering the claims invalid. Publications describing the IETF MPLS-TP System, including those cited in this chart, were publicly disclosed and available more than a year before the purported filing date of the provisional application identified on the face of the '821 Patent, rendering the claims invalid. For example, RFC 4872, RSVP-TE Extensions in Support of End-to-End Generalized Multi-Protocol Label Switching (GMPLS) Recovery was published May 2007 and discloses all elements of the Asserted Claims, thereby anticipating the Asserted Claims. Moreover, as identified in this chart, it would have been obvious to combine the disclosures in the publications regarding the IETF MPLS-TP System with each other and with disclosures in other publications known in the art, as explained in this chart, well before the claimed invention and the effective filing date of the '821 Patent, rendering the claims invalid.

The quotations and diagrams regarding the IETF MPLS-TP System in this chart come from the following sources:

- RSVP-TE Extensions in Support of End-to-End Generalized Multi-Protocol Label Switching (GMPLS) Recovery ("RFC 4872");
- MPLS-TP Linear Protection draft-ietf-mpls-tp-linear-protection-01.txt ("MPLS-TP Protection"); and
- SD detection and protection triggering in MPLS-TP draft-rkhd-mpls-tp-sd-03.txt ("MPLS-TP Detection").

On January 27, 2023, Cisco received a declaration from Alexa Morris, Managing Director of the IETF Administration LLC, attesting to the public availability and authenticity of the documents cited herein.

Motivations to combine the disclosures in these references with disclosures in other publications known in the art, as explained in this chart, include at least the similarity in subject matter between the references to the extent they concern methods of selecting paths in a network and/or determining the cost of potential paths in a network. Insofar as the references cite other patents or publications, or suggest additional changes, one of ordinary skill in the art would look beyond a single reference to other references in the field.

These invalidity contentions are based on Defendant's present understanding of the Asserted Claims, and Orckit's apparent construction of the claims in its November 3, 2022 Disclosure of Asserted Claims and Infringement Contentions Pursuant to P.R. 3-1, and Orckit's January 19, 2023 First Amended Disclosure of Asserted Claims and Infringement Contentions Pursuant to P.R. 3-1 (Orckit's "Infringement Disclosures"), which is deficient at least insofar as it fails to cite any documents or identify accused structures, acts, or materials in the Accused Products with particularity. Defendant does not agree with Orckit's application of the claims, or that the claims satisfy the requirements of 35 U.S.C. § 112. Defendant's contentions herein are not, and should in no way be seen as, admissions or adoptions as to any particular claim scope or construction, or as any admission that any particular element is met by any accused product in any particular way. Defendant objects to any attempt to imply claim construction from this chart. Defendant's prior art invalidity contentions are made in a variety of alternatives and do not represent Defendant's agreement or view as to the meaning, definiteness, written description support for, or enablement of any claim contained therein.

The following contentions are subject to revision and amendment pursuant to Federal Rule of Civil Procedure 26(e), the Local Rules, and the Orders of record in this matter subject to further investigation and discovery regarding the prior art and the Court's construction of the claims at issue.

No.	'821 Patent Claim 1	IETF MPLS-TP System
1[preamble]	An entity selection method performed by a network device, comprising the steps of:	IETF MPLS-TP System discloses an entity selection method performed by a network device, comprising the steps of. For example, RFC 4872 discloses selectors at ingress and egress routers that select one or more LSPs in an MPLS network, i.e., entities, to either send or receive through a data network. RFC 4872 further discloses that this selection method can be used to implement GMPLS recovery. "1. Introduction

No.	'821 Patent Claim 1	IETF MPLS-TP System
		Generalized Multi-Protocol Label Switching (GMPLS) extends MPLS to include support for Layer-2 Switch Capable (L2SC), Time-Division Multiplex (TDM), Lambda Switch Capable (LSC), and Fiber Switch Capable (FSC) interfaces. GMPLS recovery uses control plane mechanisms (i.e., signaling, routing, and link management mechanisms) to support data plane fault recovery. Note that the analogous (data plane) fault detection mechanisms are required to be present in support of the control plane mechanisms. In this document, the term "recovery" is generically used to denote both protection and restoration; the specific terms "protection" and "restoration" are only used when differentiation is required. The subtle distinction between protection and restoration is made based on the resource allocation done during the recovery phase (see [RFC4427]). A functional description of GMPLS recovery is provided in [RFC4426] and should be considered as a companion document. The present document describes the protocol-specific procedures for GMPLS RSVP- TE (Resource ReSerVation Protocol - Traffic Engineering) signaling (see [RFC3473]) to support end-to-end recovery. End-to-end recovery refers to the recovery of an entire LSP from its head-end (ingress node endpoint) to its tail-end (egress node endpoint). With end-to- end recovery, working LSPs are assumed to be resource-disjoint (where a resource is a link, node, or Shared Risk Link Group (SRLG)) in the network so that they do not share any failure probability, but this is not mandatory. With respect to a given set of network resources, a pair of working/protecting LSPs SHOULD be resource disjoint in case of dedicated recovery type (see below). On the other hand, in case of shared recovery (see below), a group of working LSPs SHOULD be mutually resource-disjoint in order to allow for a (single and commonly) shared protecting LSP, itself resource-disjoint from each of the working LSPs. Note that resource disjointness is a necessary (but not sufficient) condition to ensure LSP recoverability. The present document addresses four types of end-to-end LSP recovery: 1) 1+1 (unidirectional/bidirectional) protection, 2) 1:N (N >= 1) LSP protection with extra-traffic, 3) pre-planned LSP rerouting without extra-traffic (including shared mesh), and 4) full LSP rerouting. 1) The simplest notion of end-to-end LSP protection is 1+1 unidirectional protection. Using this type of protection, a protecting LSP is signaled over a dedicated resource-disjoint alternate path to protect an associated working LSP. Normal traffic is

No.	'821 Patent Claim 1	IETF MPLS-TP System
		simultaneously sent on both LSPs and a selector is used at the egress node to receive traffic from one of the LSPs. If a failure occurs along one of the LSPs, the egress node selects the traffic from the valid LSP. No coordination is required between the end nodes when a failure/switchover occurs. In 1+1 bidirectional protection, a protecting LSP is signaled over a dedicated resource-disjoint alternate path to protect the working LSP. Normal traffic is simultaneously sent on both LSPs (in both directions), and a selector is used at both ingress/egress nodes to receive traffic from the same LSP. This requires coordination between the end-nodes when switching to the protecting LSP. 2) In 1:N (N >= 1) protection with extra-traffic, the protecting LSP is a fully provisioned and resource-disjoint LSP from the N working LSPs, that allows for carrying extra-traffic. The N working LSPs MAY be mutually resource-disjoint. Coordination between end-nodes is required when switching from one of the working LSPs to the protecting LSP. As the protecting LSP is fully provisioned, default operations during protection switching are specified for a protecting LSP carrying extra-traffic, but this is not mandatory. Note that M:N protection is out of scope of this document (though mechanisms it defines may be extended to cover it). 3) Pre-planned LSP rerouting (or restoration) relies on the establishment between the same pair of end-nodes of a working LSP and a protecting LSP that is link/node/SRLG disjoint from the working one. Here, the recovery resources for the protecting LSP are pre-reserved but explicit action is required to activate (i.e., commit resource allocation at the data plane) a specific protecting LSP instantiated during the (pre-)provisioning phase. Since the protecting LSP is not "active" (i.e., fully instantiated), it cannot carry any extra-traffic. This does not mean that the corresponding resources cannot be used by other LSPs. Therefore, this mechanism protects against working LSP(s) failure(s) but requires activation of the protecting LSP after working LSP failure occurrence. This requires restoration signaling along the protecting path. "Shared-mesh" restoration can be seen as a particular case of pre-planned LSP rerouting that reduces the recovery resource requirements by allowing multiple protecting LSPs to share common link and node resources. The recovery resources are pre-reserved but explicit action is required to activate (i.e., commit resource allocation at the data plane) a

No.	'821 Patent Claim 1	IETF MPLS-TP System
		specific protecting LSP instantiated during the (pre-) provisioning phase. This procedure requires restoration signaling along the protecting path. Note that in both cases, bandwidth pre-reserved for a protecting (but not activated) LSP can be made available for carrying extra traffic. LSPs for extra-traffic (with lower holding priority than the protecting LSP) can then be established using the bandwidth pre-reserved for the protecting LSP. Also, any lower priority LSP that use the pre-reserved resources for the protecting LSP(s) must be preempted during the activation of the protecting LSP. 4) Full LSP rerouting (or restoration) switches normal traffic to an alternate LSP that is not even partially established until after the working LSP failure occurs. The new alternate route is selected at the LSP head-end node, it may reuse resources of the failed LSP at intermediate nodes and may include additional intermediate nodes and/or links." " RFC 4872 at 3–5. "Consider the following network topology:  The paths [A,B,C,D] and [A,E,F,G,D] are node and link disjoint, ignoring the ingress/egress nodes A and D. A 1+1 protected path is established from A to D over [A,B,C,D] and [A,E,F,G,D], and traffic is transmitted simultaneously over both component paths (i.e., LSPs). During the provisioning phase, both LSPs are fully instantiated (and thus activated) so that no resource sharing can be done along the protecting LSP (nor can any extra-traffic be transported). It is also RECOMMENDED to set the N bit since no protection-switching signaling is assumed in this case. When a failure occurs (say, at node B) and is detected at end-node D, the receiver at D selects the normal traffic from the other LSP. From this perspective, 1+1 unidirectional protection can be seen as an uncoordinated protection-switching mechanism acting independently at both endpoints. Also, for the LSP under failure condition, it is RECOMMENDED to not set the Path_State_Removed Flag of the ERROR_SPEC object (see [RFC3473]) upon PathErr message generation.

No.	'821 Patent Claim 1	IETF MPLS-TP System
		Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 9. "When a failure is detected by one or both endpoints of the LSP, both endpoints must select traffic from the other LSP. This action must be coordinated between node A and D. From this perspective, 1+1 bidirectional protection can be seen as a coordinated protection- switching mechanism between both endpoints. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 11. "(Full) LSP Rerouting LSP rerouting, on the other hand, switches normal traffic to an alternate LSP that is fully established only after failure occurrence. The new (alternate) route is selected at the LSP head- end and may reuse intermediate nodes included in the original route; it may also include additional intermediate nodes. For strict-hop routing, TE requirements can be directly applied to the route computation, and the failed node or link can be avoided. However, if the failure occurred within a loose-routed hop, the head-end node may not have enough information to reroute the LSP around the failure. Crankback signaling (see [CRANK]) and route exclusion techniques (see [RFC4874]) MAY be used in this case. The alternate route MAY be either computed on demand (that is, when the failure occurs; this is referred to as full LSP rerouting) or pre-computed and stored for use when the failure is reported. The latter offers faster restoration time. There is, however, a risk that the alternate route will become out of date through other changes in the network; this can be mitigated to some extent by periodic recalculation of idle alternate routes. (Full) LSP rerouting will be initiated by the head-end node that has either detected the LSP failure or received a Notify message and/or a PathErr message with the new error code/sub-code "Notify Error/LSP Locally Failed" for this LSP. The new LSP resources can be established using the make-before-break mechanism, where the new LSP is set up before the old LSP is torn down. This is done by using the mechanisms of the SESSION_ATTRIBUTE

No.	'821 Patent Claim 1	IETF MPLS-TP System
		object and the Shared-Explicit (SE) reservation style (see [RFC3209]). Both the new and old LSPs can share resources at common nodes. Note that the make-before-break mechanism is not used to avoid disruption to the normal traffic flow (the latter has already been broken by the failure that is being repaired). However, it is valuable to retain the resources allocated on the original LSP that will be reused by the new alternate LSP." RFC 4872 at 25. "The procedure is as follows: 1) The initiating (source) node sends the normal traffic onto both the working and the protecting LSPs. Once completed, the source node sends reliably a Notify message to the destination with the new error code/sub-code "Notify Error/LSP Recovered" (Switchback Request). This Notify message includes the MESSAGE_ID object. The ACK_Desired flag MUST be set in this object to request the receiver to send an acknowledgment for the message (see [RFC2961]). 2) Upon receipt of this message, the destination selects the traffic from the working LSP. At the same time, it transmits the traffic onto both the working and protecting LSP. The destination then sends reliably a Notify message to the source confirming the completion of the operation. This message includes the MESSAGE_ID_ACK object to acknowledge reception of the received Notify message. This Notify message also includes the MESSAGE_ID object. The ACK_Desired flag MUST be set in this object to request the receiver to send an acknowledgment for the message (see [RFC2961]). 3) When the source node receives this Notify message, it switches to receive traffic from the working LSP. The source node then sends an Ack message to the destination node confirming that the LSP has been reverted." RFC 4872 at 27. For example, MPLS-TP Protection discloses a protection switching framework that includes the selection of a one or more "working paths" and "recovery paths," e.g. pseudo wires and/or LSPs, within an MPLS Network. MPLS-TP Protection further discloses that this method can be performed by a Label Switched Router, i.e. a network device.

No.	'821 Patent Claim 1	IETF MPLS-TP System
		"As noted in the architecture for Multi-Protocol Label Switching Transport Profile (MPLS-TP) [TPFwk], the overall architecture framework for MPLS-TP is based on a profile of the MPLS and Pseudowire (PW) procedures as specified for the MPLS and (MS-)PW architectures defined in [RFC3031], [RFC3985] and [RFC5085]. One of the basic survivability functions, pointed out by the Survivability Framework document [SurvivFwk], is that of simple and rapid protection switching mechanisms for Label Switched Paths (LSP) and Pseudo-wires (PW). Protection switching is a fully allocated survivability mechanism. It is fully allocated in the sense that the route and bandwidth of the recovery path is reserved for a selected working path or set of working paths. It provides a fast and simple survivability mechanism, that allows the network operator to easily grasp the active state of the network, compared to other survivability mechanisms. As specified in the Survivability Framework document [SurvivFwk], protection switching is applied to a protected domain. For the purposes of this document, we define the protected domain of a P2P LSP as consisting of two Label Switching Routers (LSR) and the transport paths that connect them. For a P2MP LSP the protection domain includes the root (or source) LSR, the destination (or sink) LSRs, and the transport paths that connect them." MPLS-TP Protection at 3. "3.2. Protection switching control logical architecture Protection switching processes the triggers described above together with the inputs received from the far-end LSR. These inputs cause the LSR to take certain actions, e.g. switching the Selector Bridge to select the working or recovery path, and to transmit different protocol messages.

No.	'821 Patent Claim 1	IETF MPLS-TP System

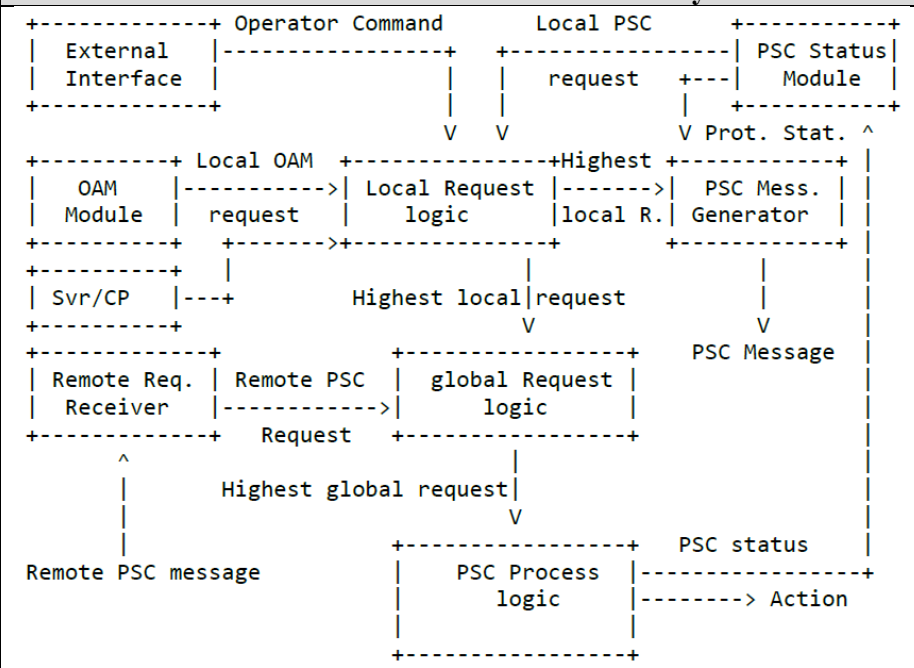
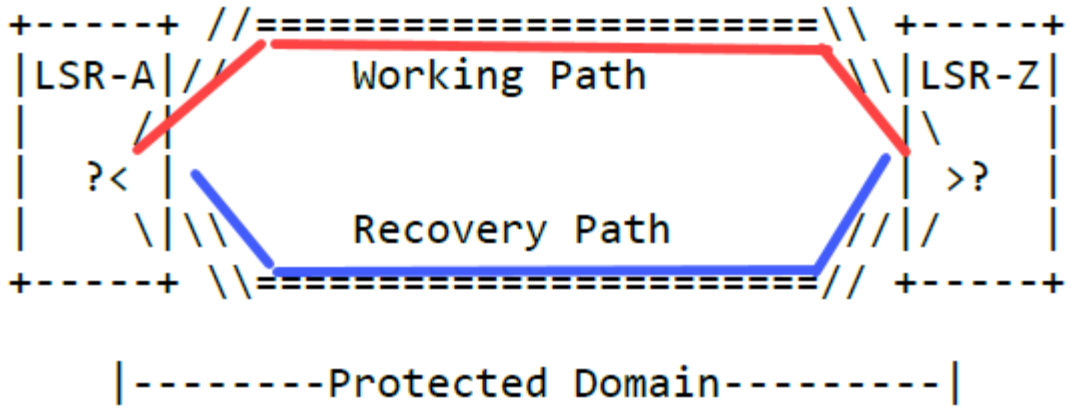


Figure 1: Protection switching control logic

Figure 1 describes the logical architecture of the protection switching control. The Local Request logic unit accepts the triggers from the OAM, external operator commands, and from the local control plane (when present) and determines the highest priority request. This high-priority request is passed to both the PSC Message generator, that will generate the appropriate protocol message to be sent to the far-end LSR, and the Global Request logic, that will cross-check this local request with the information received from the far-LSR. The Global Request logic then processes these two PSC requests that determines the highest priority request that is passed to the PSC Process logic. The PSC Process logic uses this input to determine what actions need to be taken, e.g. switching the Selector Bridge, and the current status of the protection domain." MPLS-TP Protection at 7–8.

"4.2.4. Active path indicator (Path)

No.	'821 Patent Claim 1	IETF MPLS-TP System
		The Path field of the PSC control SHALL be used to indicate which path the source MEP is currently using for data transmission. The MEP should compare the value of this bit with the path that is locally selected for data transmission to verify that there is no inconsistency between the two end-points of the protected domain. If an inconsistency is detected then an alarm should be raised. The following are the possible values: • 0: indicates that normal traffic is being transmitted on the Working path. • 1: indicates the Recovery path is being used to transmit the normal traffic from the Working path. 2-255: for future extensions" MPLS-TP Protection at 12–13. "Principles of Operation In all of the following sub-sections, assume a protected domain between LSR-A and LSR-Z, using paths W (working) and R (recovery) as shown in figure 4." MPLS-TP Protection at 13.  MPLS-TP Protection Figure 4 (annotation added). "4.3.1 PSC States 4.3.1.1 Normal State

No.	'821 Patent Claim 1	IETF MPLS-TP System
		When the protected domain has no special condition in effect, the ingress LSR SHOULD forward the user data along the working path, and, in the case of 1+1 protection, the Permanent Bridge will bridge the data to the recovery path as well. The receiving LSR SHOULD read the data from the working path. The ingress LSR MAY transmit a No Request PSC packet with the Path field set to 0 indicating that the normal data traffic should be read from the working path. 4.3.1.2. Protecting State When the protection mechanism has been triggered and the protected domain has performed a protection switch, the domain is in the protecting state. In this state the normal data traffic is transmitted and received on the recovery path. If the protection domain is currently in a protecting state, then the LSRs SHOULD NOT accept a Manual Switch request. If the protection domain is currently in a protecting state, and a Forced Switch is requested then the normal traffic SHALL continue to be transmitted on the recovery path even if the original protection trigger is cleared, and the Forced Switch condition will be signalled by the PSC messages. 4.3.1.3. Unavailable State When the recovery path is unavailable - either as a result of a Lockout operator command (see section 4.3.3), or as a result of a SF or SD detected on the recovery path (see section 4.3.4) - then the protection domain is in the unavailable state. In this state, the normal traffic is transmitted and received on the working path. While in unavailable state any event that would trigger a protection switching SHOULD be ignored with the following exception - If a Signal Degrade request is received, then protection switching will be activated only if the recovery path can guarantee a better signal than the working path. The protection domain will exit the unavailable state and revert to the normal state when, either the operator clears the Lockout command or the recovery path recovers from the signal fault or degraded situation.

No.	'821 Patent Claim 1	IETF MPLS-TP System
		Both ends will resume sending the PCS packets over the recovery path, as a result of this recovery. 4.3.2. Failure or Degraded condition (Working path) If one of the LSRs (for example, LSR-A) detects a failure condition or a serious degradation condition on the working path that warrants invoking protection switching, then it SHOULD take the following actions: • (For 1:1 protection) Switch all traffic for LSR-Z to the recovery path only. • Transmit a PCS control packet, using GACH, with the appropriate Request code (either Signal fault or Signal degrade), the Fpath set to 1, to indicate that the fault/degrade was detected on the working path, and the Path set to 1, indicating that normal traffic is now being transmitted on the recovery path. • Verify that LSR-Z replies with a PCS control packet indicating that it has switched to the recovery path. If this is not received after 2 PSC cycles then send an alarm to the management system. When the far-end LSR (in this example LSR-Z) receives the PCS packet informing it that other LSR (LSR-A) has switched, it SHOULD perform the following actions: • Check priority of the request • Switch all traffic addressed to LSR-A to the recovery path only (for 1:1 protection). • Begin transmission of a PCS control packet, using GACH, with the appropriate Request code (either Signal fault or Signal degrade), the Fpath set to 1, to indicate that the fault/degrade was detected on the working path, and the Path set to 1, indicating that traffic is now being transmitted on the recovery path." MPLS-TP Protection at 13–15. "4.3.4. Failure or Degraded condition (Recovery path) If one of the LSRs (for example, LSR-A) detects a failure condition or a serious degradation condition on the recovery path, then it SHOULD take the following actions:

No.	'821 Patent Claim 1	IETF MPLS-TP System
		• Begin transmission of a PCS control packet with the appropriate Request code (either Signal fault or Signal degrade), the Fpath set to 0, to indicate that the fault/degrade was detected on the recovery path, and the Path set to 0, indicating that traffic is now being forwarded on the working path. Note that this will actually reach the far-end if this is a unidirectional fault or recovery path is possibly in a degraded situation. • The PSC control logic should go into Unavailable state. • All traffic MUST be transmitted on the working path for the duration of the SF/SD condition. When the far-end LSR (in this example LSR-Z) receives the PCS packet informing it that other LSR (LSR-A) has become Unavailable, it SHOULD perform the following actions: • Transmit all traffic on the working path for the duration of the SF/SD condition • The PSC Control logic should go into Unavailable state." MPLS-TP Protection at 16.
1[a]	providing a plurality of multi protocol label switching (MPLS) transport entities between a first endpoint and a second endpoint;	IETF MPLS-TP System discloses providing a plurality of multi protocol label switching (MPLS) transport entities between a first endpoint and a second endpoint. For example, RFC 4872 discloses the use of Generalized Multi-Protocol Label Switching, that includes providing MPLS LSPs, i.e., transport entities, between ingress and egress routers. RFC 4872 further discloses that the MPLS LSPs can include both working LSPs and protecting LSPs. "RSVP-TE Extensions in Support of End-to-End Generalized Multi-Protocol Label Switching (GMPLS) Recovery" RFC 4872 Title. "1. Introduction Generalized Multi-Protocol Label Switching (GMPLS) extends MPLS to include support for Layer-2 Switch Capable (L2SC), Time-Division Multiplex (TDM), Lambda Switch Capable (LSC), and Fiber Switch Capable (FSC) interfaces. GMPLS recovery uses control plane mechanisms (i.e., signaling, routing, and link management mechanisms) to support data plane fault recovery. Note that the analogous (data plane) fault detection mechanisms are required to be present in support of the control plane mechanisms. In this document, the term

No.	'821 Patent Claim 1	IETF MPLS-TP System
		"recovery" is generically used to denote both protection and restoration; the specific terms "protection" and "restoration" are only used when differentiation is required. The subtle distinction between protection and restoration is made based on the resource allocation done during the recovery phase (see [RFC4427]). A functional description of GMPLS recovery is provided in [RFC4426] and should be considered as a companion document. The present document describes the protocol-specific procedures for GMPLS RSVP- TE (Resource ReSerVation Protocol - Traffic Engineering) signaling (see [RFC3473]) to support end-to-end recovery. End-to-end recovery refers to the recovery of an entire LSP from its head-end (ingress node endpoint) to its tail-end (egress node endpoint). With end-to- end recovery, working LSPs are assumed to be resource-disjoint (where a resource is a link, node, or Shared Risk Link Group (SRLG)) in the network so that they do not share any failure probability, but this is not mandatory. With respect to a given set of network resources, a pair of working/protecting LSPs SHOULD be resource disjoint in case of dedicated recovery type (see below). On the other hand, in case of shared recovery (see below), a group of working LSPs SHOULD be mutually resource-disjoint in order to allow for a (single and commonly) shared protecting LSP, itself resource-disjoint from each of the working LSPs. Note that resource disjointness is a necessary (but not sufficient) condition to ensure LSP recoverability. The present document addresses four types of end-to-end LSP recovery: 1) 1+1 (unidirectional/bidirectional) protection, 2) 1:N (N >= 1) LSP protection with extra-traffic, 3) pre-planned LSP rerouting without extra-traffic (including shared mesh), and 4) full LSP rerouting. 1) The simplest notion of end-to-end LSP protection is 1+1 unidirectional protection. Using this type of protection, a protecting LSP is signaled over a dedicated resource-disjoint alternate path to protect an associated working LSP. Normal traffic is simultaneously sent on both LSPs and a selector is used at the egress node to receive traffic from one of the LSPs. If a failure occurs along one of the LSPs, the egress node selects the traffic from the valid LSP. No coordination is required between the end nodes when a failure/switchover occurs. In 1+1 bidirectional protection, a protecting LSP is signaled over a dedicated resource-disjoint alternate path to protect the working LSP. Normal traffic is simultaneously

No.	'821 Patent Claim 1	IETF MPLS-TP System
		sent on both LSPs (in both directions), and a selector is used at both ingress/egress nodes to receive traffic from the same LSP. This requires coordination between the end-nodes when switching to the protecting LSP. 2) In 1:N (N >= 1) protection with extra-traffic, the protecting LSP is a fully provisioned and resource-disjoint LSP from the N working LSPs, that allows for carrying extra-traffic. The N working LSPs MAY be mutually resource-disjoint. Coordination between end-nodes is required when switching from one of the working LSPs to the protecting LSP. As the protecting LSP is fully provisioned, default operations during protection switching are specified for a protecting LSP carrying extra-traffic, but this is not mandatory. Note that M:N protection is out of scope of this document (though mechanisms it defines may be extended to cover it). 3) Pre-planned LSP rerouting (or restoration) relies on the establishment between the same pair of end-nodes of a working LSP and a protecting LSP that is link/node/SRLG disjoint from the working one. Here, the recovery resources for the protecting LSP are pre-reserved but explicit action is required to activate (i.e., commit resource allocation at the data plane) a specific protecting LSP instantiated during the (pre-)provisioning phase. Since the protecting LSP is not "active" (i.e., fully instantiated), it cannot carry any extra-traffic. This does not mean that the corresponding resources cannot be used by other LSPs. Therefore, this mechanism protects against working LSP(s) failure(s) but requires activation of the protecting LSP after working LSP failure occurrence. This requires restoration signaling along the protecting path. "Shared-mesh" restoration can be seen as a particular case of pre-planned LSP rerouting that reduces the recovery resource requirements by allowing multiple protecting LSPs to share common link and node resources. The recovery resources are pre-reserved but explicit action is required to activate (i.e., commit resource allocation at the data plane) a specific protecting LSP instantiated during the (pre-) provisioning phase. This procedure requires restoration signaling along the protecting path. Note that in both cases, bandwidth pre-reserved for a protecting (but not activated) LSP can be made available for carrying extra traffic. LSPs for extra-traffic (with lower holding priority than the protecting LSP) can then be established using the bandwidth pre-reserved for the protecting LSP. Also, any lower priority LSP that use the pre-

No.	'821 Patent Claim 1	IETF MPLS-TP System
		reserved resources for the protecting LSP(s) must be preempted during the activation of the protecting LSP. 4) Full LSP rerouting (or restoration) switches normal traffic to an alternate LSP that is not even partially established until after the working LSP failure occurs. The new alternate route is selected at the LSP head-end node, it may reuse resources of the failed LSP at intermediate nodes and may include additional intermediate nodes and/or links." " RFC 4872 at 3–5. "Consider the following network topology:  The paths [A,B,C,D] and [A,E,F,G,D] are node and link disjoint, ignoring the ingress/egress nodes A and D. A 1+1 protected path is established from A to D over [A,B,C,D] and [A,E,F,G,D], and traffic is transmitted simultaneously over both component paths (i.e., LSPs). During the provisioning phase, both LSPs are fully instantiated (and thus activated) so that no resource sharing can be done along the protecting LSP (nor can any extra-traffic be transported). It is also RECOMMENDED to set the N bit since no protection-switching signaling is assumed in this case. When a failure occurs (say, at node B) and is detected at end-node D, the receiver at D selects the normal traffic from the other LSP. From this perspective, 1+1 unidirectional protection can be seen as an uncoordinated protection-switching mechanism acting independently at both endpoints. Also, for the LSP under failure condition, it is RECOMMENDED to not set the Path_State_Removed Flag of the ERROR_SPEC object (see [RFC3473]) upon PathErr message generation. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 9. "When a failure is detected by one or both endpoints of the LSP, both endpoints must select traffic from the other LSP. This action must be coordinated between node A and D. From this

No.	'821 Patent Claim 1	IETF MPLS-TP System
		perspective, 1+1 bidirectional protection can be seen as a coordinated protection- switching mechanism between both endpoints. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 11. "(Full) LSP Rerouting LSP rerouting, on the other hand, switches normal traffic to an alternate LSP that is fully established only after failure occurrence. The new (alternate) route is selected at the LSP head- end and may reuse intermediate nodes included in the original route; it may also include additional intermediate nodes. For strict-hop routing, TE requirements can be directly applied to the route computation, and the failed node or link can be avoided. However, if the failure occurred within a loose-routed hop, the head-end node may not have enough information to reroute the LSP around the failure. Crankback signaling (see [CRANK]) and route exclusion techniques (see [RFC4874]) MAY be used in this case. The alternate route MAY be either computed on demand (that is, when the failure occurs; this is referred to as full LSP rerouting) or pre-computed and stored for use when the failure is reported. The latter offers faster restoration time. There is, however, a risk that the alternate route will become out of date through other changes in the network; this can be mitigated to some extent by periodic recalculation of idle alternate routes. (Full) LSP rerouting will be initiated by the head-end node that has either detected the LSP failure or received a Notify message and/or a PathErr message with the new error code/sub-code "Notify Error/LSP Locally Failed" for this LSP. The new LSP resources can be established using the make-before-break mechanism, where the new LSP is set up before the old LSP is torn down. This is done by using the mechanisms of the SESSION_ATTRIBUTE object and the Shared-Explicit (SE) reservation style (see [RFC3209]). Both the new and old LSPs can share resources at common nodes. Note that the make-before-break mechanism is not used to avoid disruption to the normal traffic flow (the latter has already been broken by the failure that is being repaired). However, it is valuable to retain the resources allocated on the original LSP that will be reused by the new alternate LSP." RFC 4872 at 25.

No.	'821 Patent Claim 1	IETF MPLS-TP System
		"The procedure is as follows: 4) The initiating (source) node sends the normal traffic onto both the working and the protecting LSPs. Once completed, the source node sends reliably a Notify message to the destination with the new error code/sub-code "Notify Error/LSP Recovered" (Switchback Request). This Notify message includes the MESSAGE_ID object. The ACK_Desired flag MUST be set in this object to request the receiver to send an acknowledgment for the message (see [RFC2961]). 5) Upon receipt of this message, the destination selects the traffic from the working LSP. At the same time, it transmits the traffic onto both the working and protecting LSP. The destination then sends reliably a Notify message to the source confirming the completion of the operation. This message includes the MESSAGE_ID_ACK object to acknowledge reception of the received Notify message. This Notify message also includes the MESSAGE_ID object. The ACK_Desired flag MUST be set in this object to request the receiver to send an acknowledgment for the message (see [RFC2961]). 6) When the source node receives this Notify message, it switches to receive traffic from the working LSP. The source node then sends an Ack message to the destination node confirming that the LSP has been reverted." RFC 4872 at 27. For example, RFC 4872 discloses selectors at ingress and egress routers that select one or more LSPs in an MPLS network, i.e., entities, to either send or receive through a data network. RFC 4872 further discloses that this selection method can be used to implement GMPLS recovery. For example, MPLS-TP Protection discloses a protection switching framework that includes the selection of a one or more "working paths" and "recovery paths," e.g. pseudo wires and/or LSPs, within an MPLS Network. A person of ordinary skill in the art would understand that these pseudo-wires and/or LSPs constitute "transport entities." "As noted in the architecture for Multi-Protocol Label Switching Transport Profile (MPLS-TP) [TPFwk], the overall architecture framework for MPLS-TP is based on a profile of the

No.	'821 Patent Claim 1	IETF MPLS-TP System
		MPLS and Pseudowire (PW) procedures as specified for the MPLS and (MS-)PW architectures defined in [RFC3031], [RFC3985] and [RFC5085]. One of the basic survivability functions, pointed out by the Survivability Framework document [SurvivFwk], is that of simple and rapid protection switching mechanisms for Label Switched Paths (LSP) and Pseudo-wires (PW). Protection switching is a fully allocated survivability mechanism. It is fully allocated in the sense that the route and bandwidth of the recovery path is reserved for a selected working path or set of working paths. It provides a fast and simple survivability mechanism, that allows the network operator to easily grasp the active state of the network, compared to other survivability mechanisms. As specified in the Survivability Framework document [SurvivFwk], protection switching is applied to a protected domain. For the purposes of this document, we define the protected domain of a P2P LSP as consisting of two Label Switching Routers (LSR) and the transport paths that connect them. For a P2MP LSP the protection domain includes the root (or source) LSR, the destination (or sink) LSRs, and the transport paths that connect them." MPLS-TP Protection at 3. "3.2. Protection switching control logical architecture Protection switching processes the triggers described above together with the inputs received from the far-end LSR. These inputs cause the LSR to take certain actions, e.g. switching the Selector Bridge to select the working or recovery path, and to transmit different protocol messages.

No.	'821 Patent Claim 1	IETF MPLS-TP System

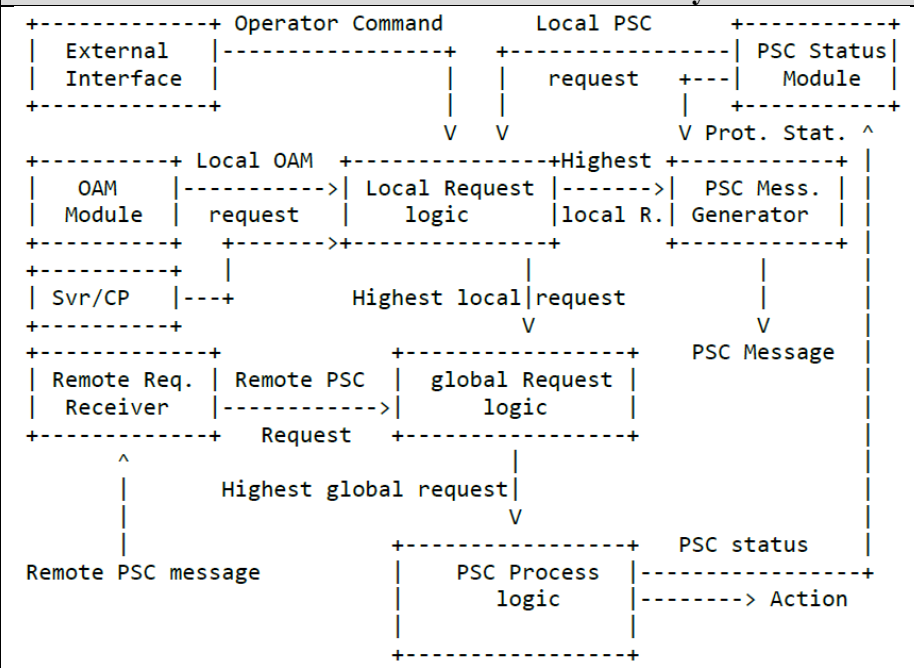
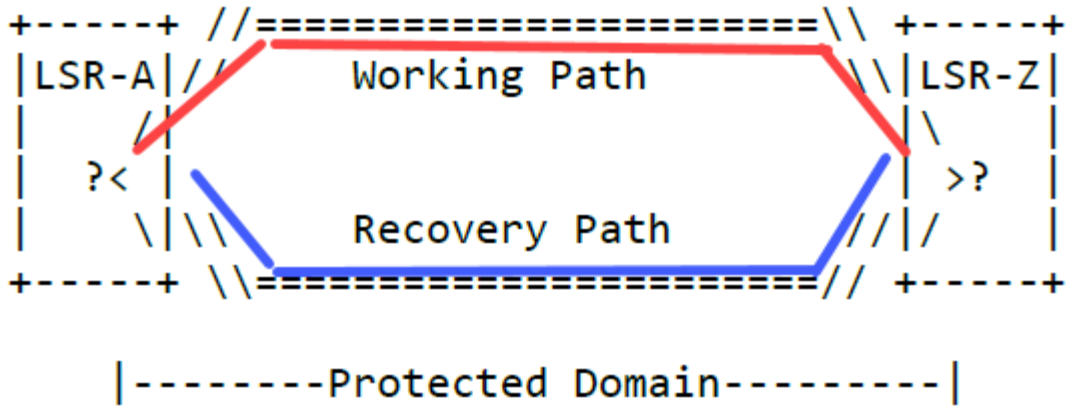


Figure 1: Protection switching control logic

Figure 1 describes the logical architecture of the protection switching control. The Local Request logic unit accepts the triggers from the OAM, external operator commands, and from the local control plane (when present) and determines the highest priority request. This high-priority request is passed to both the PSC Message generator, that will generate the appropriate protocol message to be sent to the far-end LSR, and the Global Request logic, that will cross-check this local request with the information received from the far-LSR. The Global Request logic then processes these two PSC requests that determines the highest priority request that is passed to the PSC Process logic. The PSC Process logic uses this input to determine what actions need to be taken, e.g. switching the Selector Bridge, and the current status of the protection domain." MPLS-TP Protection at 7–8.

"4.2.4. Active path indicator (Path)

No.	'821 Patent Claim 1	IETF MPLS-TP System
		The Path field of the PSC control SHALL be used to indicate which path the source MEP is currently using for data transmission. The MEP should compare the value of this bit with the path that is locally selected for data transmission to verify that there is no inconsistency between the two end-points of the protected domain. If an inconsistency is detected then an alarm should be raised. The following are the possible values: • 0: indicates that normal traffic is being transmitted on the Working path. • 1: indicates the Recovery path is being used to transmit the normal traffic from the Working path. 2-255: for future extensions" MPLS-TP Protection at 12–13. "Principles of Operation In all of the following sub-sections, assume a protected domain between LSR-A and LSR-Z, using paths W (working) and R (recovery) as shown in figure 4." MPLS-TP Protection at 13.  MPLS-TP Protection Figure 4 (annotation added). "4.3.1 PSC States 4.3.1.1 Normal State

No.	'821 Patent Claim 1	IETF MPLS-TP System
		When the protected domain has no special condition in effect, the ingress LSR SHOULD forward the user data along the working path, and, in the case of 1+1 protection, the Permanent Bridge will bridge the data to the recovery path as well. The receiving LSR SHOULD read the data from the working path. The ingress LSR MAY transmit a No Request PSC packet with the Path field set to 0 indicating that the normal data traffic should be read from the working path. 4.3.1.2. Protecting State When the protection mechanism has been triggered and the protected domain has performed a protection switch, the domain is in the protecting state. In this state the normal data traffic is transmitted and received on the recovery path. If the protection domain is currently in a protecting state, then the LSRs SHOULD NOT accept a Manual Switch request. If the protection domain is currently in a protecting state, and a Forced Switch is requested then the normal traffic SHALL continue to be transmitted on the recovery path even if the original protection trigger is cleared, and the Forced Switch condition will be signalled by the PSC messages. 4.3.1.3. Unavailable State When the recovery path is unavailable - either as a result of a Lockout operator command (see section 4.3.3), or as a result of a SF or SD detected on the recovery path (see section 4.3.4) - then the protection domain is in the unavailable state. In this state, the normal traffic is transmitted and received on the working path. While in unavailable state any event that would trigger a protection switching SHOULD be ignored with the following exception - If a Signal Degrade request is received, then protection switching will be activated only if the recovery path can guarantee a better signal than the working path. The protection domain will exit the unavailable state and revert to the normal state when, either the operator clears the Lockout command or the recovery path recovers from the signal fault or degraded situation.

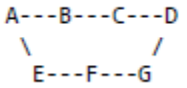
No.	'821 Patent Claim 1	IETF MPLS-TP System
		Both ends will resume sending the PCS packets over the recovery path, as a result of this recovery. 4.3.2. Failure or Degraded condition (Working path) If one of the LSRs (for example, LSR-A) detects a failure condition or a serious degradation condition on the working path that warrants invoking protection switching, then it SHOULD take the following actions: • (For 1:1 protection) Switch all traffic for LSR-Z to the recovery path only. • Transmit a PCS control packet, using GACH, with the appropriate Request code (either Signal fault or Signal degrade), the Fpath set to 1, to indicate that the fault/degrade was detected on the working path, and the Path set to 1, indicating that normal traffic is now being transmitted on the recovery path. • Verify that LSR-Z replies with a PCS control packet indicating that it has switched to the recovery path. If this is not received after 2 PSC cycles then send an alarm to the management system. When the far-end LSR (in this example LSR-Z) receives the PCS packet informing it that other LSR (LSR-A) has switched, it SHOULD perform the following actions: • Check priority of the request • Switch all traffic addressed to LSR-A to the recovery path only (for 1:1 protection). • Begin transmission of a PCS control packet, using GACH, with the appropriate Request code (either Signal fault or Signal degrade), the Fpath set to 1, to indicate that the fault/degrade was detected on the working path, and the Path set to 1, indicating that traffic is now being transmitted on the recovery path." MPLS-TP Protection at 13–15. "4.3.4. Failure or Degraded condition (Recovery path) If one of the LSRs (for example, LSR-A) detects a failure condition or a serious degradation condition on the recovery path, then it SHOULD take the following actions:

No.	'821 Patent Claim 1	IETF MPLS-TP System
		• Begin transmission of a PCS control packet with the appropriate Request code (either Signal fault or Signal degrade), the Fpath set to 0, to indicate that the fault/degrade was detected on the recovery path, and the Path set to 0, indicating that traffic is now being forwarded on the working path. Note that this will actually reach the far-end if this is a unidirectional fault or recovery path is possibly in a degraded situation. • The PSC control logic should go into Unavailable state. • All traffic MUST be transmitted on the working path for the duration of the SF/SD condition. When the far-end LSR (in this example LSR-Z) receives the PCS packet informing it that other LSR (LSR-A) has become Unavailable, it SHOULD perform the following actions: • Transmit all traffic on the working path for the duration of the SF/SD condition The PSC Control logic should go into Unavailable state." MPLS-TP Protection at 16.
1[b]	determining an overall cost for each entity pair of said plurality of entities;	IETF MPLS-TP System discloses determining an overall cost for each entity pair of said plurality of entities. For example, RFC 4872 discloses that the plurality of MPLS LSPs, i.e., transport entities, are selected to necessarily be "resource disjoint," i.e., where multiple resources, such as links, nodes, or Shared Risk Link Groups (SRLGs) do not share common links or failure probabilities. Thus, at least under the apparent claim scope alleged by Orckit's Infringement Disclosures, this limitation is met. To the extent that the IETF MPLS-TP System is found to not meet this limitation, determining an overall cost for each entity pair of said plurality of entities would have been obvious to a person having ordinary skill in the art, as explained below. "When a failure is detected by one or both endpoints of the LSP, both endpoints must select traffic from the other LSP. This action must be coordinated between node A and D. From this perspective, 1+1 bidirectional protection can be seen as a coordinated protection- switching mechanism between both endpoints. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 11.

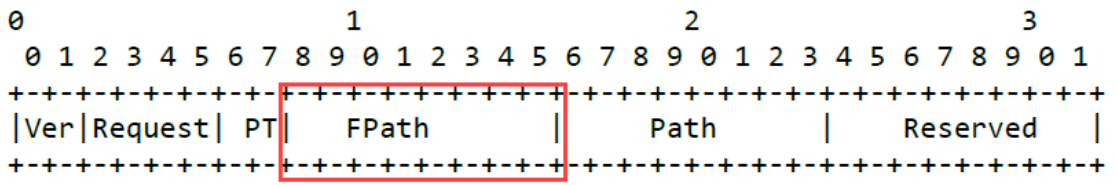
No.	'821 Patent Claim 1	IETF MPLS-TP System
		"1:1 Protection with Extra-Traffic The most common case of end-to-end 1:N protection is to establish, between the same endpoints, an end-to-end working LSP (thus, N = 1) and a dedicated end-to-end protecting LSP that are mutually link/ node/SRLG disjoint. This protects against working LSP failure(s)." RFC 4872 at 13. "An approach to reduce recovery resource requirements is to have protection LSPs sharing network resources when the working LSPs that they protect are physically (i.e., link, node, SRLG, etc.) disjoint. This mechanism is referred to as shared mesh restoration and is described in [RFC4426]. Shared-mesh restoration can be seen as a particular case of pre-planned LSP rerouting (see Section 8) that reduces the recovery resource requirements by allowing multiple protecting LSPs to share common link and node resources. Here also, the recovery resources for the protecting LSPs are pre-reserved during the provisioning phase, thus an explicit signaling action is required to activate (i.e., commit resource allocation at the data plane) a specific protecting LSP instantiated during the (pre-) provisioning phase. This requires restoration signaling along the protecting LSP." RFC 4872 at 20. "1. Introduction Generalized Multi-Protocol Label Switching (GMPLS) extends MPLS to include support for Layer-2 Switch Capable (L2SC), Time-Division Multiplex (TDM), Lambda Switch Capable (LSC), and Fiber Switch Capable (FSC) interfaces. GMPLS recovery uses control plane mechanisms (i.e., signaling, routing, and link management mechanisms) to support data plane fault recovery. Note that the analogous (data plane) fault detection mechanisms are required to be present in support of the control plane mechanisms. In this document, the term "recovery" is generically used to denote both protection and restoration; the specific terms "protection" and "restoration" are only used when differentiation is required. The subtle distinction between protection and restoration is made based on the resource allocation done during the recovery phase (see [RFC4427]). A functional description of GMPLS recovery is provided in [RFC4426] and should be considered as a companion document. The present document describes the protocol-specific

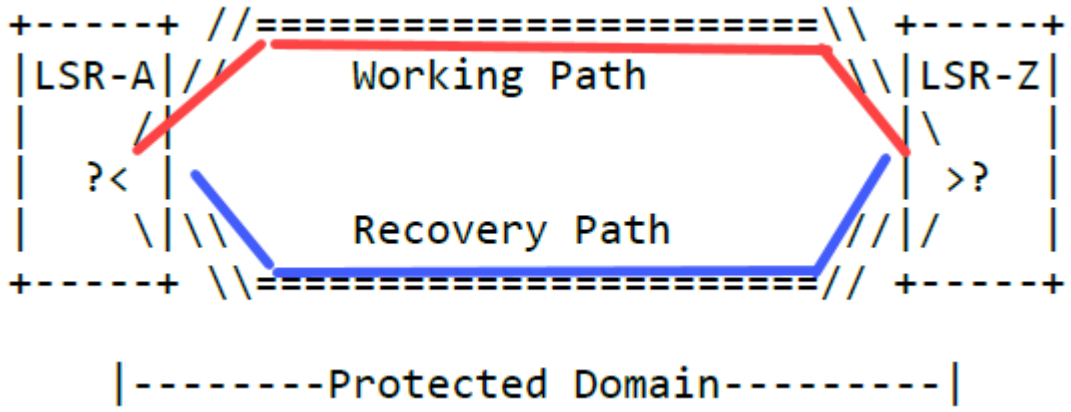
No.	'821 Patent Claim 1	IETF MPLS-TP System
		procedures for GMPLS RSVP- TE (Resource ReSerVation Protocol - Traffic Engineering) signaling (see [RFC3473]) to support end-to-end recovery. End-to-end recovery refers to the recovery of an entire LSP from its head-end (ingress node endpoint) to its tail-end (egress node endpoint). With end-to- end recovery, working LSPs are assumed to be resource-disjoint (where a resource is a link, node, or Shared Risk Link Group (SRLG)) in the network so that they do not share any failure probability, but this is not mandatory. With respect to a given set of network resources, a pair of working/protecting LSPs SHOULD be resource disjoint in case of dedicated recovery type (see below). On the other hand, in case of shared recovery (see below), a group of working LSPs SHOULD be mutually resource-disjoint in order to allow for a (single and commonly) shared protecting LSP, itself resource-disjoint from each of the working LSPs. Note that resource disjointness is a necessary (but not sufficient) condition to ensure LSP recoverability. The present document addresses four types of end-to-end LSP recovery: 1) 1+1 (unidirectional/bidirectional) protection, 2) 1:N (N >= 1) LSP protection with extra-traffic, 3) pre-planned LSP rerouting without extra-traffic (including shared mesh), and 4) full LSP rerouting. 5) The simplest notion of end-to-end LSP protection is 1+1 unidirectional protection. Using this type of protection, a protecting LSP is signaled over a dedicated resource-disjoint alternate path to protect an associated working LSP. Normal traffic is simultaneously sent on both LSPs and a selector is used at the egress node to receive traffic from one of the LSPs. If a failure occurs along one of the LSPs, the egress node selects the traffic from the valid LSP. No coordination is required between the end nodes when a failure/switchover occurs. In 1+1 bidirectional protection, a protecting LSP is signaled over a dedicated resource-disjoint alternate path to protect the working LSP. Normal traffic is simultaneously sent on both LSPs (in both directions), and a selector is used at both ingress/egress nodes to receive traffic from the same LSP. This requires coordination between the end-nodes when switching to the protecting LSP. 6) In 1:N (N >= 1) protection with extra-traffic, the protecting LSP is a fully provisioned and resource-disjoint LSP from the N working LSPs, that allows for carrying extra-traffic. The N working LSPs MAY be mutually resource-disjoint. Coordination

No.	'821 Patent Claim 1	IETF MPLS-TP System
		between end-nodes is required when switching from one of the working LSPs to the protecting LSP. As the protecting LSP is fully provisioned, default operations during protection switching are specified for a protecting LSP carrying extra-traffic, but this is not mandatory. Note that M:N protection is out of scope of this document (though mechanisms it defines may be extended to cover it). 7) Pre-planned LSP rerouting (or restoration) relies on the establishment between the same pair of end-nodes of a working LSP and a protecting LSP that is link/node/SRLG disjoint from the working one. Here, the recovery resources for the protecting LSP are pre-reserved but explicit action is required to activate (i.e., commit resource allocation at the data plane) a specific protecting LSP instantiated during the (pre-)provisioning phase. Since the protecting LSP is not "active" (i.e., fully instantiated), it cannot carry any extra-traffic. This does not mean that the corresponding resources cannot be used by other LSPs. Therefore, this mechanism protects against working LSP(s) failure(s) but requires activation of the protecting LSP after working LSP failure occurrence. This requires restoration signaling along the protecting path. "Shared-mesh" restoration can be seen as a particular case of pre-planned LSP rerouting that reduces the recovery resource requirements by allowing multiple protecting LSPs to share common link and node resources. The recovery resources are pre-reserved but explicit action is required to activate (i.e., commit resource allocation at the data plane) a specific protecting LSP instantiated during the (pre-) provisioning phase. This procedure requires restoration signaling along the protecting path. Note that in both cases, bandwidth pre-reserved for a protecting (but not activated) LSP can be made available for carrying extra traffic. LSPs for extra-traffic (with lower holding priority than the protecting LSP) can then be established using the bandwidth pre-reserved for the protecting LSP. Also, any lower priority LSP that use the pre-reserved resources for the protecting LSP(s) must be preempted during the activation of the protecting LSP. 8) Full LSP rerouting (or restoration) switches normal traffic to an alternate LSP that is not even partially established until after the working LSP failure occurs. The new alternate route is selected at the LSP head-end node, it may reuse resources of the failed LSP at intermediate nodes and may include additional intermediate nodes and/or links."

No.	'821 Patent Claim 1	IETF MPLS-TP System
		" RFC 4872 at 3–5. "Consider the following network topology:  The paths [A,B,C,D] and [A,E,F,G,D] are node and link disjoint, ignoring the ingress/egress nodes A and D. A 1+1 protected path is established from A to D over [A,B,C,D] and [A,E,F,G,D], and traffic is transmitted simultaneously over both component paths (i.e., LSPs). During the provisioning phase, both LSPs are fully instantiated (and thus activated) so that no resource sharing can be done along the protecting LSP (nor can any extra-traffic be transported). It is also RECOMMENDED to set the N bit since no protection-switching signaling is assumed in this case. When a failure occurs (say, at node B) and is detected at end-node D, the receiver at D selects the normal traffic from the other LSP. From this perspective, 1+1 unidirectional protection can be seen as an uncoordinated protection-switching mechanism acting independently at both endpoints. Also, for the LSP under failure condition, it is RECOMMENDED to not set the Path_State_Removed Flag of the ERROR_SPEC object (see [RFC3473]) upon PathErr message generation. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 9. "When a failure is detected by one or both endpoints of the LSP, both endpoints must select traffic from the other LSP. This action must be coordinated between node A and D. From this perspective, 1+1 bidirectional protection can be seen as a coordinated protection- switching mechanism between both endpoints. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 11. "(Full) LSP Rerouting

No.	'821 Patent Claim 1	IETF MPLS-TP System
		LSP rerouting, on the other hand, switches normal traffic to an alternate LSP that is fully established only after failure occurrence. The new (alternate) route is selected at the LSP head- end and may reuse intermediate nodes included in the original route; it may also include additional intermediate nodes. For strict-hop routing, TE requirements can be directly applied to the route computation, and the failed node or link can be avoided. However, if the failure occurred within a loose-routed hop, the head-end node may not have enough information to reroute the LSP around the failure. Crankback signaling (see [CRANK]) and route exclusion techniques (see [RFC4874]) MAY be used in this case. The alternate route MAY be either computed on demand (that is, when the failure occurs; this is referred to as full LSP rerouting) or pre-computed and stored for use when the failure is reported. The latter offers faster restoration time. There is, however, a risk that the alternate route will become out of date through other changes in the network; this can be mitigated to some extent by periodic recalculation of idle alternate routes. (Full) LSP rerouting will be initiated by the head-end node that has either detected the LSP failure or received a Notify message and/or a PathErr message with the new error code/sub-code "Notify Error/LSP Locally Failed" for this LSP. The new LSP resources can be established using the make-before-break mechanism, where the new LSP is set up before the old LSP is torn down. This is done by using the mechanisms of the SESSION_ATTRIBUTE object and the Shared-Explicit (SE) reservation style (see [RFC3209]). Both the new and old LSPs can share resources at common nodes. Note that the make-before-break mechanism is not used to avoid disruption to the normal traffic flow (the latter has already been broken by the failure that is being repaired). However, it is valuable to retain the resources allocated on the original LSP that will be reused by the new alternate LSP." RFC 4872 at 25. "The procedure is as follows: 7) The initiating (source) node sends the normal traffic onto both the working and the protecting LSPs. Once completed, the source node sends reliably a Notify message to the destination with the new error code/sub-code "Notify Error/LSP Recovered" (Switchback Request). This Notify message includes the MESSAGE_ID object. The

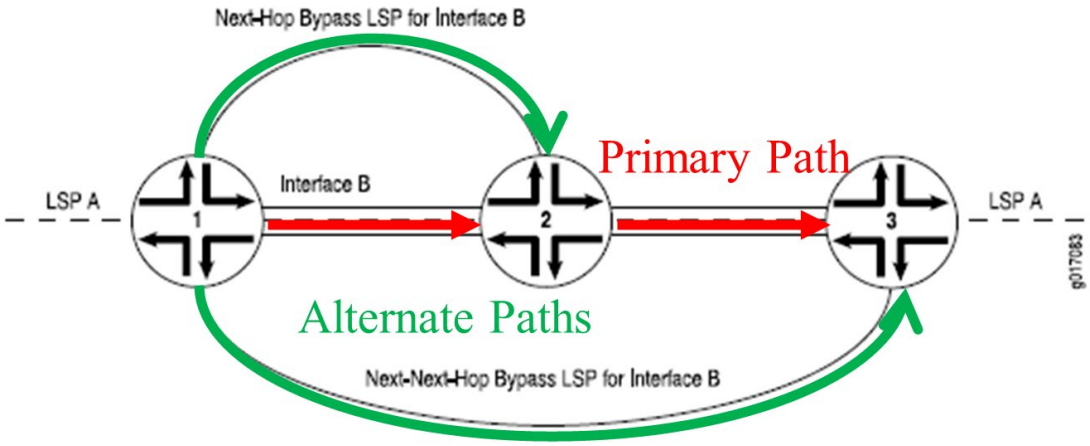
No.	'821 Patent Claim 1	IETF MPLS-TP System
		ACK_Desired flag MUST be set in this object to request the receiver to send an acknowledgment for the message (see [RFC2961]). 8) Upon receipt of this message, the destination selects the traffic from the working LSP. At the same time, it transmits the traffic onto both the working and protecting LSP. The destination then sends reliably a Notify message to the source confirming the completion of the operation. This message includes the MESSAGE_ID_ACK object to acknowledge reception of the received Notify message. This Notify message also includes the MESSAGE_ID object. The ACK_Desired flag MUST be set in this object to request the receiver to send an acknowledgment for the message (see [RFC2961]). 9) When the source node receives this Notify message, it switches to receive traffic from the working LSP. The source node then sends an Ack message to the destination node confirming that the LSP has been reverted." RFC 4872 at 27. For example, MPLS-TP Protection discloses an FPath field, that determines whether or not a specified LSP and/or pseudo-wire, i.e. network entity, is operational. At least under Orckit's Infringement Disclosures, this limitation is met. To the extent that the IETF MPLS-TP System is found to not meet this limitation, determining an overall cost for each entity pair of said plurality of entities would have been obvious to a person having ordinary skill in the art, as explained below.  Figure 3: Format of the PSC control packet MPLS-TP Protection Figure 3, annotation added.

No.	'821 Patent Claim 1	IETF MPLS-TP System
		"FPath: used to indicate the path that is reporting a failure condition, the possible values are described in section 4.2.3" MPLS-TP Protection at 11. "4.2.3. Path fault identifier (FPath) The Fpath field of the PSC control SHALL be used only in a Signal fault (0110) or Signal degrade (0101) control packet. Its value indicates on which path the signal anomaly was detected. The following are the possible values: • 0: indicates that the fault condition is on the Recovery path • 1: indicates that the fault condition is on the Working path • 2-255: for future extensions." MPLS-TP Protection at 12. "Principles of Operation In all of the following sub-sections, assume a protected domain between LSR-A and LSR-Z, using paths W (working) and R (recovery) as shown in figure 4." MPLS-TP Protection at 13.  MPLS-TP Protection Figure 4 (annotation added).

No.	'821 Patent Claim 1	IETF MPLS-TP System
		In addition, at least under the apparent claim scope alleged by Orckit's Infringement Disclosures, IETF MPLS System in combination with (1) the knowledge of a person of ordinary skill in the art, alone or in further combination with (2) each (individually, as well as one or more together) of the references identified in element 1[b] of Exhibit E-2 renders the claim, including the present limitation, obvious. Below are examples of two such references. For example, Juniper IOS Guide discloses "Fate Sharing" where a Juniper router computes a primary and secondary path using CSPF, CSPF identifies all fate-sharing groups that are associated with the primary path by identifying all links that the primary path transverses and checks each link in the traffic engineering database against the compiled group list. Juniper IOS Guide further discloses that if the link is a member of a group, the cost of the link is increased by the cost of the group and higher the cost value. At least based upon Orckit's apparent understanding of the claim scope as reflected in its Infringement Disclosures, this limitation is met. Juniper IOS Guide at Page No. 73 ("[Fate Sharing Fate sharing allows you to create a database of information that CSPF uses to compute one or more backup paths to use in case the primary path becomes unstable. The database describes the relationships between elements of the network, such as routers and links. You can specify one or more elements within a group. Through fate sharing, you can configure backup paths that minimize the number of shared links and fiber paths with the primary paths as much as possible, to ensure that if a fiber is cut, the minimum amount of data is lost and a path still exists to the destination. For a backup path to work optimally, it must not share links or physical fiber paths with the primary path, ensuring that a single point of failure will not affect the primary and backup paths simultaneously. For more information about fate sharing, see the Junos OS Routing Protocols Configuration Guide.]") Juniper IOS Guide at Page No. 94 ("[Configuring the Ingress Router for MPLS-Signaled LSPs

No.	'821 Patent Claim 1	IETF MPLS-TP System
		MPLS-signaled label-switched paths (LSPs) run from a specific ingress router to a specific egress router. For basic MPLS-signaled LSP function, you must configure the ingress router, but do not have to configure any other routers. To configure signaled LSPs, perform the following tasks on the ingress router: • Creating Named Paths on page 56 • Configuring Alternate Backup Paths Using Fate Sharing on page 58.]") Juniper IOS Guide at Page No. 96 ("[Configuring Alternate Backup Paths Using Fate Sharing You can create a database of information that Constrained Shortest Path First (CSPF) uses to compute one or more backup paths in case the primary path becomes unstable. The database describes the relationships between elements of the network, such as routers and links. Because these network elements share the same fate, this relationship is called fate sharing. You can configure backup paths that minimize the number of shared links and fiber paths with the primary paths as much as possible to ensure that, if a fiber is cut, the minimum amount of data is lost and a path still exists to the destination. For a backup path to work optimally, it must not share links or physical fiber paths with the primary path. This ensures that a single point of failure will not affect the primary and backup paths at the same time.]") Juniper IOS Guide at Page No. 97 ("[Fate-sharing groups contain three types of objects: • Point-to-point links—Identified by the IP addresses at each end of the link. Unnumbered point-to-point links are typically identified by borrowing IP addresses from other interfaces. Order is not important; from 1.2.3.4 to 1.2.3.5 and from 1.2.3.5 to 1.2.3.4 have the same meaning. • Non-point-to-point links—Include links on a LAN interface (such as Gigabit Ethernet interfaces) or nonbroadcast multiaccess (NBMA) interfaces (such as Asynchronous Transfer Mode [ATM] or Frame Relay). You identify these links by their individual interface address. For example, if the LAN interface 192.168.200.0/24 has four routers attached to it, each router link is individually identified: from 192.168.200.1; # LAN interface of router 1 from 192.168.200.2; # LAN interface of router 2 from 192.168.200.3; # LAN interface of router 3

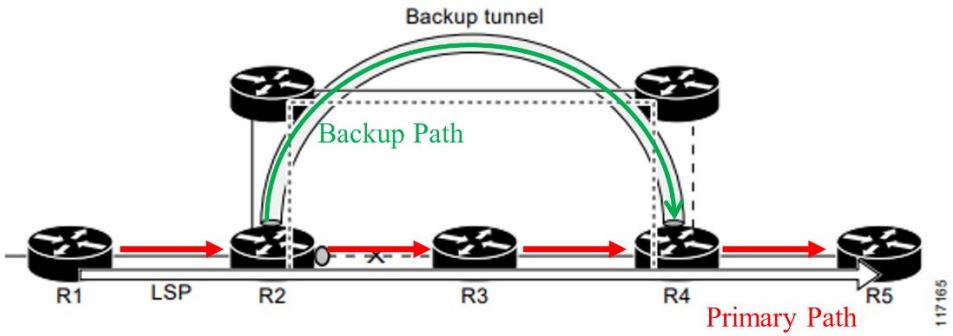
No.	'821 Patent Claim 1	IETF MPLS-TP System
		from 192.168.200.4; # LAN interface of router 4 You can list the addresses in any order. • A router node—Identified by its configured router ID. An object can be in any number of groups, and a group can contain any number of objects. Each group has a configurable cost attributed to it, which represents the level of impact this group has on CSPF computations. The higher the cost, the less likely a backup path will share with the primary path any objects in the group. The cost is directly comparable to traffic engineering metrics. By default, the cost is 1. Changing the fate-sharing database does not affect established LSPs until the next reoptimization of CSPF. The fate-sharing database does influence fast-reroute computations.]") Juniper IOS Guide at Page No. 97 ("[When CSPF computes the primary paths of an LSP (or secondary paths when the primary path is not active), it ignores the fate-sharing information. You always want to find the best possible path (least IGP cost) for the primary path. When CSPF computes a secondary path while the primary path (of the same LSP) is active, the following occurs: 1. CSPF identifies all fate-sharing groups that are associated with the primary path. CSPF does this by identifying all links and nodes that the primary path traverses and compiling group lists that contain at least one of the links or nodes. CSPF ignores the ingress and egress nodes in the search. 2. CSPF checks each link in the traffic engineering database against the compiled group list. If the link is a member of a group, the cost of the link is increased by the cost of the group. If a link is a member of multiple groups, all group costs are added together. 3. CSPF performs the check for every node in the traffic engineering database, except the ingress and egress node. Again, a node can belong to multiple groups, so costs are additive. 4. The router performs regular CSPF computation with the adjusted topology.]") Juniper IOS Guide at Page No. 330 ("In Figure 23 on page 292, both node protection and link protection are enabled on Interface B on Router 1. Both node protection and link protection are also enabled on LSP A, an LSP that traverses the link transiting Router 1, Router 2, and Router 3. If the link between Router 1 and Router 2 fails, traffic from LSP A is quickly switched to the next-hop bypass LSP generated by link protection. If Router 2

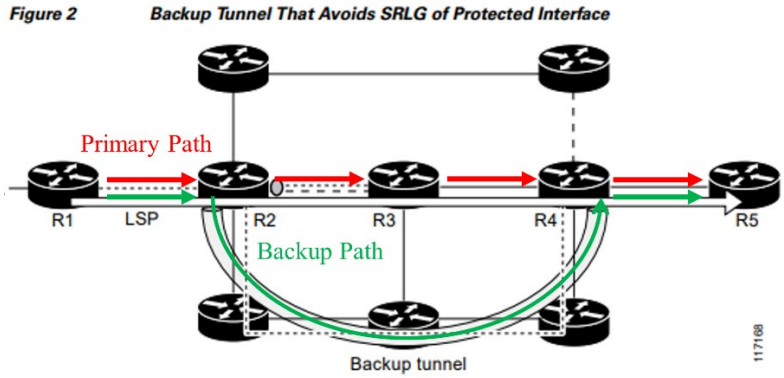
No.	'821 Patent Claim 1	IETF MPLS-TP System
		suffers a hardware or software failure, traffic from LSP A is switched to the next-next-hop bypass LSP generated by node protection.") Figure 23: Node Protection Creating a Next-Next-Hop Bypass LSP  Fig. 23 (annotation added). For example, Cisco IOS Guide discloses that the path an LSP uses is determined by the LSP resource requirements and network resources, such as bandwidth. Cisco IOS Guide further discloses the MPLS TE SRLG feature that enhances backup tunnel path selection so a backup tunnel can avoid using links that are in the same SRLG as the interfaces it is protecting. Thus, at least under the apparent claim scope alleged by Orckit's Infringement Disclosures, this limitation is met. Cisco IOS Guide at Page No. 253 ("[Path Option Selection for MPLS TE Tunnel LSPs This section contains the following topics about path option selection for MPLS TE Tunnel LSPs: • Constraint-Based Routing and Path Option Selection, page 5 • Tunnel Reoptimization and Path Option Selection, page 5 • Path Option Selection with Bandwidth Override, page 6

No.	'821 Patent Claim 1	IETF MPLS-TP System
		Constraint-Based Routing and Path Option Selection MPLS traffic engineering automatically establishes and maintains LSPs across the backbone by using the Resource Reservation Protocol (RSVP). The path that an LSP uses is determined by the LSP resource requirements and network resources, such as bandwidth. Traffic engineering tunnels are calculated at the LSP head based on a fit between required and available resources (constraint-based routing). Without the Path Option for Bandwidth Override feature, a TE tunnel establishes an LSP based on dynamic or explicit path options in order of preference. However, the bandwidth and other attributes configured on the TE tunnel allow the setup of an LSP only if LSP path options satisfy the constraints. If a path cannot be found that satisfies the configured path options, then the tunnel is not set up.]") Cisco IOS Guide at Page No. 639 ("[Bandwidth Protection NHOP and NNHOP backup tunnels can be used to provide bandwidth protection for rerouted LSPs. This is referred to as backup bandwidth. You can associate backup bandwidth with NHOP or NNHOP backup tunnels. This informs the router of the amount of backup bandwidth a particular backup tunnel can protect. When a router maps LSPs to backup tunnels, bandwidth protection ensures that an LSP uses a given backup tunnel only if there is sufficient backup bandwidth. The router selects which LSPs use which backup tunnels in order to provide maximum bandwidth protection. That is, the router determines the best way to map LSPs onto backup tunnels in order to maximize the number of LSPs that can be protected. For information about mapping tunnels and assigning backup bandwidth, see the "Backup Tunnel Selection Procedure" section on page 10. LSPs that have the "bandwidth protection desired" bit set have a higher right to select backup tunnels that provide bandwidth protection; that is, those LSPs can preempt other LSPs that do not have that bit set. For more information, see the "Prioritizing Which LSPs Obtain Backup Tunnels with Bandwidth Protection" section on page 8.]") Cisco IOS Guide at Page No. 640 ("[Backup Tunnel Support Backup tunnel support has the following capabilities: • Backup Tunnels Can Terminate at the Next-Next Hop to Support FRR, page 6 • Multiple Backup Tunnels Can Protect the Same Interface, page 6

No.	'821 Patent Claim 1	IETF MPLS-TP System
		• Backup Tunnels Provide Scalability, page 7.]") Cisco IOS Guide at Page 640 ("[Multiple Backup Tunnels Can Protect the Same Interface There is no limit (except memory limitations) to the number of backup tunnels that can protect a given interface. In many topologies, support for node protection requires supporting multiple backup tunnels per protected interface. These backup tunnels can terminate at the same destination or at different destinations. That is, for a given protected interface, you can configure multiple NHOP or NNHOP backup tunnels. This allows redundancy and load balancing.]") Cisco IOS Guide at Page No. 603 ("[Traffic Engineering Tunnels MPLS TE lets you build label switched paths (LSPs) across your network for forwarding traffic. MPLS TE LSPs let the headend of a TE tunnel control the path its traffic takes to a particular destination. This method is more flexible than forwarding traffic based only on a destination address. Interarea tunnels allow you to do the following: • Build TE tunnels between areas (interarea tunnels) • Build TE tunnels that start and end in the same area, on multiple areas on a router (intra-area tunnels) Some tunnels are more important than others. For example, you may have tunnels carrying VoIP traffic and tunnels carrying data traffic that are competing for the same resources. MPLS TE allows you to have some tunnels preempt others. Each tunnel has a priority, and more-important tunnels take precedence over less-important tunnels.]") Cisco IOS Guide at Page No. 641 ("[Backup Bandwidth Protection Backup bandwidth protection allows you to give LSPs carrying certain kinds of data (such as voice) priority for using backup tunnels.]") Cisco IOS Guide at Page No. 650 ("[For each protected link AB with a maximum reservable subpool value of n, there may be a path from node A to node B such that the difference between the maximum reservable global and the maximum reservable subpool is at least the

No.	'821 Patent Claim 1	IETF MPLS-TP System
		value of n. If it is possible to find such paths for each link in the network, you can establish all the backup tunnels along such paths without any bandwidth reservations. If there is a single link failure, only one backup tunnel will use any link on its path. Because that path has at least n available bandwidth (in the global pool), assuming that marking and scheduling is configured to classify the subpool traffic into a priority queue, the subpool bandwidth is guaranteed. This approach allows sharing of the global pool bandwidth between backup tunnels protecting independent link failures. The backup tunnels are expected to be used for only a short period of time after a failure (until the headends of affected LSPs reroute those LSPs to other paths with available subpool bandwidth). The probability of multiple unrelated link failures is very small (in the absence of node or shared risk link group (SRLG) failures, which result in multiple link failures). Therefore, it is reasonable to assume that link failures are in practice independent with high probability. This "independent failure assumption" in combination with backup tunnels signaled without explicit bandwidth reservation enables efficient bandwidth sharing that yields substantial bandwidth savings.]") Cisco IOS Guide at Page No. 651 ("[A similar approach can be used for node and SRLG protection. However, the decision of where to put the backup tunnels is more complicated because both node and SRLG failures effectively result in the simultaneous failure of several links. Therefore, the backup tunnels protecting traffic traversing all affected links cannot be computed independently of each other.]") Cisco IOS Guide at Page No. 579 ("[MPLS Traffic Engineering Shared Risk Link Groups SRLGs refer to situations where links in a network share a common fiber (or a common physical attribute). If one link fails, other links in the group may fail too. Links in the group have a shared risk. Backup tunnels should avoid using links in the same SRLG as interfaces they are protecting. Otherwise, when the protected link fails the backup tunnel fails too.

No.	'821 Patent Claim 1	IETF MPLS-TP System
		Figure 1 Backup Tunnel in the Same SRLG as the Interface It Is Protecting  Fig. 1 (annotation added) Figure 1 shows a primary label-switched path (LSP) from router R1 to router R5. The LSP protects against the failure of the R2-R3 link at R2 via a backup tunnel to R4. If the R2-R3 link fails, link protection reroutes the LSP along the backup tunnel. However, the R2-R3 link and one of the backup tunnel links are in the same SRLG. So if the R2-R3 link fails, the backup tunnel may fail too. The MPLS TE SRLG feature enhances backup tunnel path selection so a backup tunnel can avoid using links that are in the same SRLG as the interfaces it is protecting. There are two ways for a backup tunnel to avoid the SRLGs of its protected interface: • The router does not create the backup tunnel unless it avoids SRLGs of the protected interface.]") Cisco IOS Guide at Page No. 580 ("[Open Shortest Path First (OSPF) and Intermediate System-to-Intermediate System (IS-IS) flood the SRLG membership information (including other TE link attributes such as bandwidth availability and affinity) so that all routers in the network have the SRLG information for each link. With this topology information, routers can compute backup tunnel paths that exclude links having SRLGs in common with their protected interfaces. As shown in Figure 2, the backup tunnel avoids the link between R2 and R3, which shares an SRLG with the protected interface.]")

No.	'821 Patent Claim 1	IETF MPLS-TP System
		 Fig. 2 (annotation added) Cisco IOS Guide at Page No. 587 ("[Step 4 show mpls traffic-eng topology Use this command to show the SRLG link membership flooded via the Interior Gateway Protocol (IGP).]") Cisco IOS Guide at Page 632 ("[IS-IS—Intermediate System-to-Intermediate System. Link-state hierarchical routing protocol that calls for intermediate system (IS) routers to exchange routing information based on a single metric to determine network topology.]") Cisco IOS Guide at Page No. 633 ("[OSPF—Open Shortest Path First. A link-state hierarchical Interior Gateway Protocol routing algorithm, derived from the IS-IS protocol. OSPF features include least-cost routing, multipath routing, and load balancing.]") Cisco IOS Guide at Page No. 581 ("[Autotunnel Backup for MPLS TE SRLGs Autotunnel backup is the ability of routers to create backup tunnels automatically. Therefore, you do not need to preconfigure each backup tunnel and then assign the backup tunnel to the protected interface. Only automatically created backup tunnels can avoid SRLGs or their protected interfaces.]") Cisco IOS Guide at Page No. 583 ("[Configuring the Routers That Automatically Create Backup Tunnels to Avoid MPLS TE SRLGs of Their Protected Interfaces

No.	'821 Patent Claim 1	IETF MPLS-TP System
		Perform the following task to configure routers that automatically create backup tunnels to avoid MPLS TE SRLGs of their protected interfaces. Backup tunnels provide link protection by rerouting traffic to the next hop bypassing failed links or in this instance by avoiding SRLGs.]") Cisco IOS Guide at Page No. 644 ("[Backup Tunnel Selection Procedure When an LSP is signaled, each node along the LSP path that provides FRR protection for the LSP selects a backup tunnel for the LSP to use if either of the following events occurs: • The link to the next hop fails. • The next hop fails. By having the node select the backup tunnel for an LSP before a failure occurs, the LSP can be rerouted onto the backup tunnel quickly if there is a failure.]") Cisco IOS Guide at Page 646 ("[NHOP Versus NNHOP Backup Tunnels More than one backup tunnel can protect a given LSP, where one backup tunnel terminates at the LSP's NNHOP, and the other terminates at the LSP's NHOP. In this case, the router chooses the backup tunnel that terminates at the NNHOP (that is, FRR prefers NNHOP over NHOP backup tunnels).Table 1 lists the tunnel selection priorities. The first choice is an NNHOP backup tunnel that acquires its bandwidth from a subpool or global pool, and has limited bandwidth. If there is no such backup tunnel, the next choice (2) is a next-next hop backup tunnel that acquires a limited amount of bandwidth from any pool. The preferences go from 1 (best) to 8 (worst), where choice 3 is for an NNHOP backup tunnel with an unlimited amount of subpool or global-pool bandwidth.]")

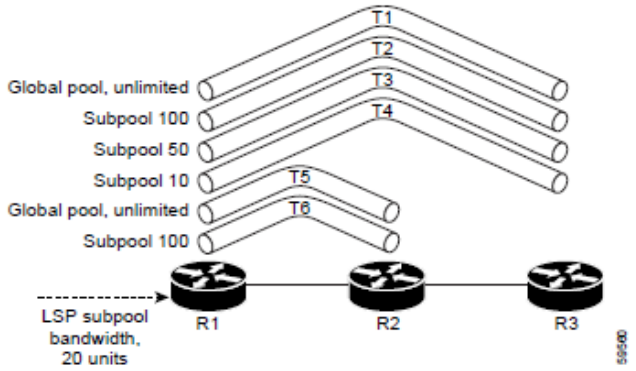
No.	'821 Patent Claim 1	IETF MPLS-TP System

Table 1 Tunnel Selection Priorities

Preference	Backup Tunnel Destination	Bandwidth Pool	Bandwidth Amount
1 (Best)	NNHOP	Subpool or global pool	Limited
2	NNHOP	Any	Limited
3	NNHOP	Subpool or global pool	Unlimited
4	NNHOP	Any	Unlimited
5	NHOP	Subpool or global pool	Limited
6	NHOP	Any	Limited
7	NHOP	Subpool or global pool	Unlimited
8 (Worst)	NHOP	Any	Unlimited

Cisco IOS Guide at Page No. 647-648 ("[Figure 7 shows an example of the backup tunnel selection procedure based on the designated amount of global pool and subpool bandwidth currently available.

Figure 7 Choosing from Among Multiple Backup Tunnels



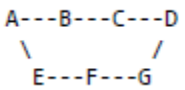
In this example, an LSP requires 20 units (kilobits per second) of sub-pool backup bandwidth. The best backup tunnel is selected as follows:
1. Backup tunnels T1 through T4 are considered first because they terminate at the NNHOP.
2. Tunnel T4 is eliminated because it has only ten units of sub-pool backup bandwidth.
3. Tunnel T1 is eliminated because it protects only LSPs using global-pool bandwidth.

No.	'821 Patent Claim 1	IETF MPLS-TP System
		4. Tunnel T3 is chosen over T2 because, although both have sufficient backup bandwidth, T3 has the least backup bandwidth available (leaving the most backup bandwidth available on T2). 5. Tunnels T5 and T6 need not be considered because they terminate at an NHOP, and therefore are less desirable than T3, which terminates at an NNHOP.]")
1[c]	selecting an entity pair from said plurality of transport entities based at least in part upon said overall cost; and	IETF MPLS-TP System discloses selecting an entity pair from said plurality of transport entities based at least in part upon said overall cost. For example, RFC 4872 discloses that the plurality of MPLS LSPs, i.e., transport entities, are selected in pairs of either working and protection LSPs, or working and restoration LSPs. RFC 4872 further discloses that the plurality of MPLS LSPs, i.e., transport entities, are selected to necessarily be "resource disjoint," i.e., where multiple resources, such as links, nodes, or Shared Risk Link Groups (SRLGs) do not share common links or failure probabilities. Thus, at least under the apparent claim scope alleged by Orckit's Infringement Disclosures, this limitation is met. To the extent that the IETF MPLS-TP System is found to not meet this limitation, selecting an entity pair from said plurality of transport entities based at least in part upon said overall cost would have been obvious to a person having ordinary skill in the art, as explained below. "When a failure is detected by one or both endpoints of the LSP, both endpoints must select traffic from the other LSP. This action must be coordinated between node A and D. From this perspective, 1+1 bidirectional protection can be seen as a coordinated protection- switching mechanism between both endpoints. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 11. "1:1 Protection with Extra-Traffic The most common case of end-to-end 1:N protection is to establish, between the same endpoints, an end-to-end working LSP (thus, N = 1) and a dedicated end-to-end protecting LSP that are mutually link/ node/SRLG disjoint. This protects against working LSP failure(s)." RFC 4872 at 13.

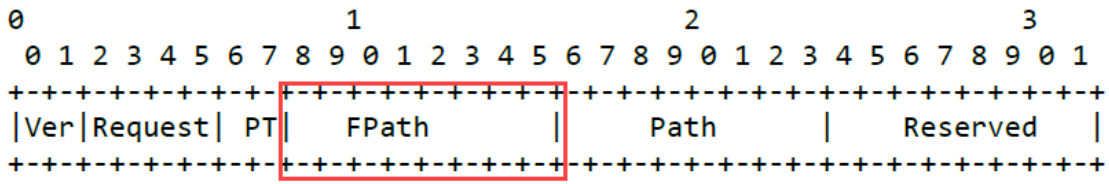
No.	'821 Patent Claim 1	IETF MPLS-TP System
		"An approach to reduce recovery resource requirements is to have protection LSPs sharing network resources when the working LSPs that they protect are physically (i.e., link, node, SRLG, etc.) disjoint. This mechanism is referred to as shared mesh restoration and is described in [RFC4426]. Shared-mesh restoration can be seen as a particular case of pre-planned LSP rerouting (see Section 8) that reduces the recovery resource requirements by allowing multiple protecting LSPs to share common link and node resources. Here also, the recovery resources for the protecting LSPs are pre-reserved during the provisioning phase, thus an explicit signaling action is required to activate (i.e., commit resource allocation at the data plane) a specific protecting LSP instantiated during the (pre-) provisioning phase. This requires restoration signaling along the protecting LSP." RFC 4872 at 20. "1. Introduction Generalized Multi-Protocol Label Switching (GMPLS) extends MPLS to include support for Layer-2 Switch Capable (L2SC), Time-Division Multiplex (TDM), Lambda Switch Capable (LSC), and Fiber Switch Capable (FSC) interfaces. GMPLS recovery uses control plane mechanisms (i.e., signaling, routing, and link management mechanisms) to support data plane fault recovery. Note that the analogous (data plane) fault detection mechanisms are required to be present in support of the control plane mechanisms. In this document, the term "recovery" is generically used to denote both protection and restoration; the specific terms "protection" and "restoration" are only used when differentiation is required. The subtle distinction between protection and restoration is made based on the resource allocation done during the recovery phase (see [RFC4427]). A functional description of GMPLS recovery is provided in [RFC4426] and should be considered as a companion document. The present document describes the protocol-specific procedures for GMPLS RSVP- TE (Resource ReSerVation Protocol - Traffic Engineering) signaling (see [RFC3473]) to support end-to-end recovery. End-to-end recovery refers to the recovery of an entire LSP from its head-end (ingress node endpoint) to its tail-end (egress node endpoint). With end-to- end recovery, working LSPs are assumed to be resource-disjoint (where a resource is a link, node, or Shared Risk Link Group (SRLG)) in the network so that they do not share any failure probability, but this is not mandatory. With respect to a given

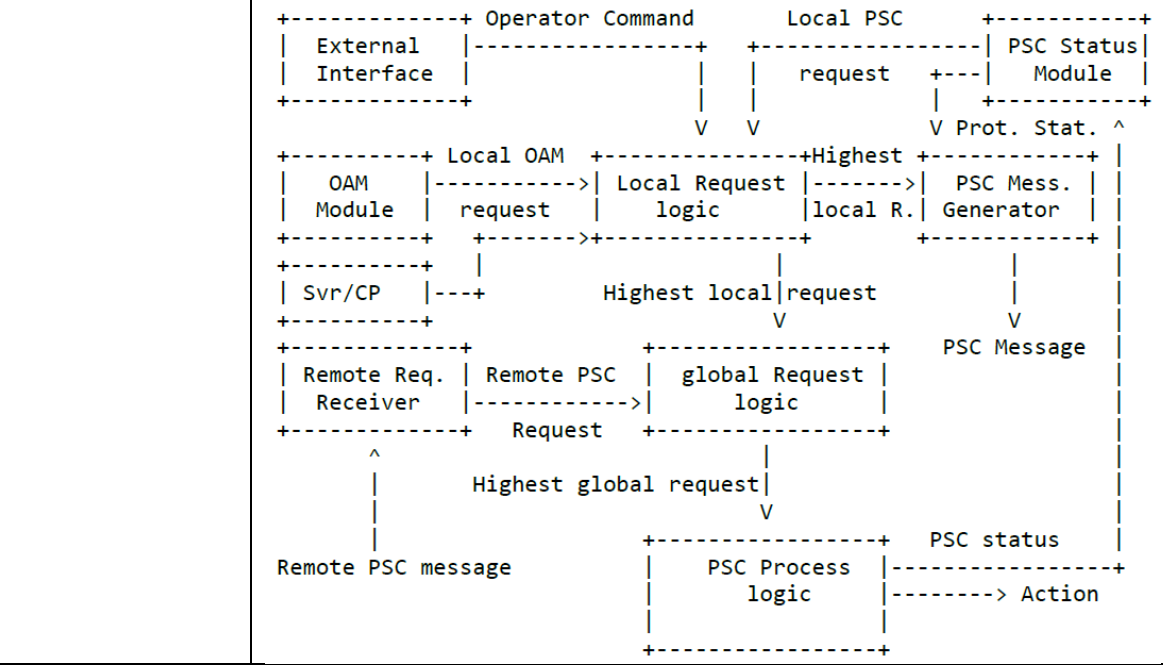
No.	'821 Patent Claim 1	IETF MPLS-TP System
		set of network resources, a pair of working/protecting LSPs SHOULD be resource disjoint in case of dedicated recovery type (see below). On the other hand, in case of shared recovery (see below), a group of working LSPs SHOULD be mutually resource-disjoint in order to allow for a (single and commonly) shared protecting LSP, itself resource-disjoint from each of the working LSPs. Note that resource disjointness is a necessary (but not sufficient) condition to ensure LSP recoverability. The present document addresses four types of end-to-end LSP recovery: 1) 1+1 (unidirectional/bidirectional) protection, 2) 1:N (N >= 1) LSP protection with extra-traffic, 3) pre-planned LSP rerouting without extra-traffic (including shared mesh), and 4) full LSP rerouting. 9) The simplest notion of end-to-end LSP protection is 1+1 unidirectional protection. Using this type of protection, a protecting LSP is signaled over a dedicated resource-disjoint alternate path to protect an associated working LSP. Normal traffic is simultaneously sent on both LSPs and a selector is used at the egress node to receive traffic from one of the LSPs. If a failure occurs along one of the LSPs, the egress node selects the traffic from the valid LSP. No coordination is required between the end nodes when a failure/switchover occurs. In 1+1 bidirectional protection, a protecting LSP is signaled over a dedicated resource-disjoint alternate path to protect the working LSP. Normal traffic is simultaneously sent on both LSPs (in both directions), and a selector is used at both ingress/egress nodes to receive traffic from the same LSP. This requires coordination between the end-nodes when switching to the protecting LSP. 10) In 1:N (N >= 1) protection with extra-traffic, the protecting LSP is a fully provisioned and resource-disjoint LSP from the N working LSPs, that allows for carrying extra-traffic. The N working LSPs MAY be mutually resource-disjoint. Coordination between end-nodes is required when switching from one of the working LSPs to the protecting LSP. As the protecting LSP is fully provisioned, default operations during protection switching are specified for a protecting LSP carrying extra-traffic, but this is not mandatory. Note that M:N protection is out of scope of this document (though mechanisms it defines may be extended to cover it).

No.	'821 Patent Claim 1	IETF MPLS-TP System
		11) Pre-planned LSP rerouting (or restoration) relies on the establishment between the same pair of end-nodes of a working LSP and a protecting LSP that is link/node/SRLG disjoint from the working one. Here, the recovery resources for the protecting LSP are pre-reserved but explicit action is required to activate (i.e., commit resource allocation at the data plane) a specific protecting LSP instantiated during the (pre-)provisioning phase. Since the protecting LSP is not "active" (i.e., fully instantiated), it cannot carry any extra-traffic. This does not mean that the corresponding resources cannot be used by other LSPs. Therefore, this mechanism protects against working LSP(s) failure(s) but requires activation of the protecting LSP after working LSP failure occurrence. This requires restoration signaling along the protecting path. "Shared-mesh" restoration can be seen as a particular case of pre-planned LSP rerouting that reduces the recovery resource requirements by allowing multiple protecting LSPs to share common link and node resources. The recovery resources are pre-reserved but explicit action is required to activate (i.e., commit resource allocation at the data plane) a specific protecting LSP instantiated during the (pre-) provisioning phase. This procedure requires restoration signaling along the protecting path. Note that in both cases, bandwidth pre-reserved for a protecting (but not activated) LSP can be made available for carrying extra traffic. LSPs for extra-traffic (with lower holding priority than the protecting LSP) can then be established using the bandwidth pre-reserved for the protecting LSP. Also, any lower priority LSP that use the pre-reserved resources for the protecting LSP(s) must be preempted during the activation of the protecting LSP. 12) Full LSP rerouting (or restoration) switches normal traffic to an alternate LSP that is not even partially established until after the working LSP failure occurs. The new alternate route is selected at the LSP head-end node, it may reuse resources of the failed LSP at intermediate nodes and may include additional intermediate nodes and/or links." " RFC 4872 at 3–5. "Consider the following network topology:

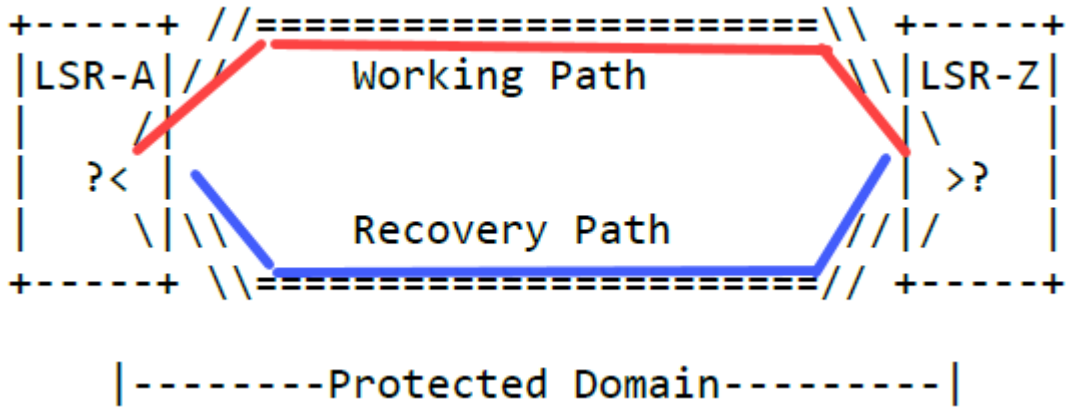
No.	'821 Patent Claim 1	IETF MPLS-TP System
		 The paths [A,B,C,D] and [A,E,F,G,D] are node and link disjoint, ignoring the ingress/egress nodes A and D. A 1+1 protected path is established from A to D over [A,B,C,D] and [A,E,F,G,D], and traffic is transmitted simultaneously over both component paths (i.e., LSPs). During the provisioning phase, both LSPs are fully instantiated (and thus activated) so that no resource sharing can be done along the protecting LSP (nor can any extra-traffic be transported). It is also RECOMMENDED to set the N bit since no protection-switching signaling is assumed in this case. When a failure occurs (say, at node B) and is detected at end-node D, the receiver at D selects the normal traffic from the other LSP. From this perspective, 1+1 unidirectional protection can be seen as an uncoordinated protection-switching mechanism acting independently at both endpoints. Also, for the LSP under failure condition, it is RECOMMENDED to not set the Path_State_Removed Flag of the ERROR_SPEC object (see [RFC3473]) upon PathErr message generation. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 9. "When a failure is detected by one or both endpoints of the LSP, both endpoints must select traffic from the other LSP. This action must be coordinated between node A and D. From this perspective, 1+1 bidirectional protection can be seen as a coordinated protection- switching mechanism between both endpoints. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 11. "(Full) LSP Rerouting LSP rerouting, on the other hand, switches normal traffic to an alternate LSP that is fully established only after failure occurrence. The new (alternate) route is selected at the LSP head- end and may reuse intermediate nodes included in the original route; it may also include

No.	'821 Patent Claim 1	IETF MPLS-TP System
		additional intermediate nodes. For strict-hop routing, TE requirements can be directly applied to the route computation, and the failed node or link can be avoided. However, if the failure occurred within a loose-routed hop, the head-end node may not have enough information to reroute the LSP around the failure. Crankback signaling (see [CRANK]) and route exclusion techniques (see [RFC4874]) MAY be used in this case. The alternate route MAY be either computed on demand (that is, when the failure occurs; this is referred to as full LSP rerouting) or pre-computed and stored for use when the failure is reported. The latter offers faster restoration time. There is, however, a risk that the alternate route will become out of date through other changes in the network; this can be mitigated to some extent by periodic recalculation of idle alternate routes. (Full) LSP rerouting will be initiated by the head-end node that has either detected the LSP failure or received a Notify message and/or a PathErr message with the new error code/sub-code "Notify Error/LSP Locally Failed" for this LSP. The new LSP resources can be established using the make-before-break mechanism, where the new LSP is set up before the old LSP is torn down. This is done by using the mechanisms of the SESSION_ATTRIBUTE object and the Shared-Explicit (SE) reservation style (see [RFC3209]). Both the new and old LSPs can share resources at common nodes. Note that the make-before-break mechanism is not used to avoid disruption to the normal traffic flow (the latter has already been broken by the failure that is being repaired). However, it is valuable to retain the resources allocated on the original LSP that will be reused by the new alternate LSP." RFC 4872 at 25. "The procedure is as follows: 10) The initiating (source) node sends the normal traffic onto both the working and the protecting LSPs. Once completed, the source node sends reliably a Notify message to the destination with the new error code/sub-code "Notify Error/LSP Recovered" (Switchback Request). This Notify message includes the MESSAGE_ID object. The ACK_Desired flag MUST be set in this object to request the receiver to send an acknowledgment for the message (see [RFC2961]). 11) Upon receipt of this message, the destination selects the traffic from the working LSP. At the same time, it transmits the traffic onto both the working and protecting LSP.

No.	'821 Patent Claim 1	IETF MPLS-TP System
		The destination then sends reliably a Notify message to the source confirming the completion of the operation. This message includes the MESSAGE_ID_ACK object to acknowledge reception of the received Notify message. This Notify message also includes the MESSAGE_ID object. The ACK_Desired flag MUST be set in this object to request the receiver to send an acknowledgment for the message (see [RFC2961]). 12) When the source node receives this Notify message, it switches to receive traffic from the working LSP. The source node then sends an Ack message to the destination node confirming that the LSP has been reverted." RFC 4872 at 27. For example, MPLS-TP Protection discloses an Fpath field, that determines whether or not a specified LSP and/or pseudo-wire, i.e. network entity, is operational. MPLS-TP Protection further discloses that a network path, *e.g.*, a working path or a recovery path, is selected based upon this FPath information. At least under Orckit's Infringement Disclosures, this limitation is met. To the extent that the IETF MPLS-TP System is found to not meet this limitation, selecting an entity pair from said plurality of transport entities based at least in part upon said overall cost would have been obvious to a person having ordinary skill in the art, as explained below.  0 1 2 3 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 |Ver|Request| PT| FPath | Path | Reserved | Figure 3: Format of the PSC control packet MPLS-TP Protection Figure 3, annotation added. "FPath: used to indicate the path that is reporting a failure condition, the possible values are described in section 4.2.3" MPLS-TP Protection at 11. "4.2.3. Path fault identifier (FPath)

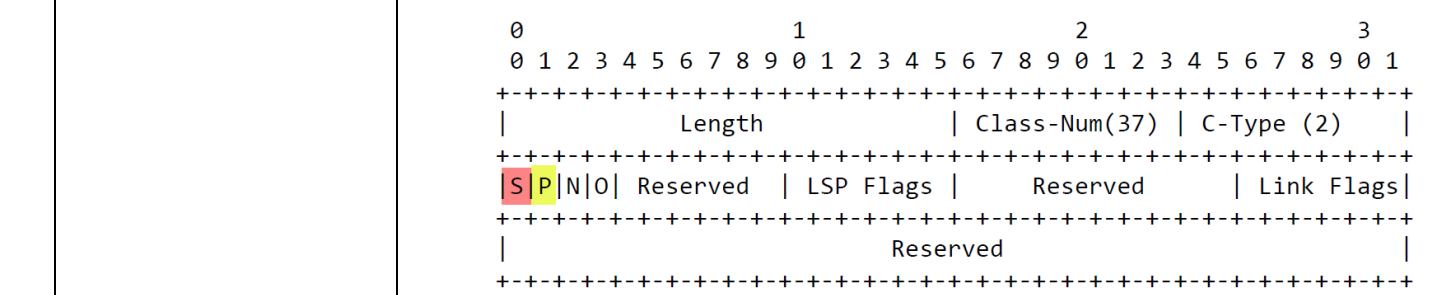
No.	'821 Patent Claim 1	IETF MPLS-TP System
		The Fpath field of the PSC control SHALL be used only in a Signal fault (0110) or Signal degrade (0101) control packet. Its value indicates on which path the signal anomaly was detected. The following are the possible values: • 0: indicates that the fault condition is on the Recovery path • 1: indicates that the fault condition is on the Working path • 2-255: for future extensions." MPLS-TP Protection at 12. "3.2. Protection switching control logical architecture Protection switching processes the triggers described above together with the inputs received from the far-end LSR. These inputs cause the LSR to take certain actions, e.g. switching the Selector Bridge to select the working or recovery path, and to transmit different protocol messages. 

No.	'821 Patent Claim 1	IETF MPLS-TP System
		Figure 1: Protection switching control logic Figure 1 describes the logical architecture of the protection switching control. The Local Request logic unit accepts the triggers from the OAM, external operator commands, and from the local control plane (when present) and determines the highest priority request. This high-priority request is passed to both the PSC Message generator, that will generate the appropriate protocol message to be sent to the far-end LSR, and the Global Request logic, that will cross-check this local request with the information received from the far-LSR. The Global Request logic then processes these two PSC requests that determines the highest priority request that is passed to the PSC Process logic. The PSC Process logic uses this input to determine what actions need to be taken, e.g. switching the Selector Bridge, and the current status of the protection domain." MPLS-TP Protection at 7–8. "4.2.4. Active path indicator (Path) The Path field of the PSC control SHALL be used to indicate which path the source MEP is currently using for data transmission. The MEP should compare the value of this bit with the path that is locally selected for data transmission to verify that there is no inconsistency between the two end-points of the protected domain. If an inconsistency is detected then an alarm should be raised. The following are the possible values: • 0: indicates that normal traffic is being transmitted on the Working path. • 1: indicates the Recovery path is being used to transmit the normal traffic from the Working path. 2-255: for future extensions" MPLS-TP Protection at 12–13. "Principles of Operation In all of the following sub-sections, assume a protected domain between LSR-A and LSR-Z, using paths W (working) and R (recovery) as shown in figure 4." MPLS-TP Protection at 13.

No.	'821 Patent Claim 1	IETF MPLS-TP System
		 MPLS-TP Protection Figure 4 (annotation added). "4.3.1 PSC States 4.3.1.1 Normal State When the protected domain has no special condition in effect, the ingress LSR SHOULD forward the user data along the working path, and, in the case of 1+1 protection, the Permanent Bridge will bridge the data to the recovery path as well. The receiving LSR SHOULD read the data from the working path. The ingress LSR MAY transmit a No Request PSC packet with the Path field set to 0 indicating that the normal data traffic should be read from the working path. 4.3.1.2. Protecting State When the protection mechanism has been triggered and the protected domain has performed a protection switch, the domain is in the protecting state. In this state the normal data traffic is transmitted and received on the recovery path. If the protection domain is currently in a protecting state, then the LSRs SHOULD NOT accept a Manual Switch request. If the protection domain is currently in a protecting state, and a Forced Switch is requested then the

No.	'821 Patent Claim 1	IETF MPLS-TP System
		normal traffic SHALL continue to be transmitted on the recovery path even if the original protection trigger is cleared, and the Forced Switch condition will be signalled by the PSC messages. 4.3.1.3. Unavailable State When the recovery path is unavailable - either as a result of a Lockout operator command (see section 4.3.3), or as a result of a SF or SD detected on the recovery path (see section 4.3.4) - then the protection domain is in the unavailable state. In this state, the normal traffic is transmitted and received on the working path. While in unavailable state any event that would trigger a protection switching SHOULD be ignored with the following exception - If a Signal Degrade request is received, then protection switching will be activated only if the recovery path can guarantee a better signal than the working path. The protection domain will exit the unavailable state and revert to the normal state when, either the operator clears the Lockout command or the recovery path recovers from the signal fault or degraded situation. Both ends will resume sending the PCS packets over the recovery path, as a result of this recovery. 4.3.2. Failure or Degraded condition (Working path) If one of the LSRs (for example, LSR-A) detects a failure condition or a serious degradation condition on the working path that warrants invoking protection switching, then it SHOULD take the following actions: • (For 1:1 protection) Switch all traffic for LSR-Z to the recovery path only. • Transmit a PCS control packet, using GACH, with the appropriate Request code (either Signal fault or Signal degrade), the Fpath set to 1, to indicate that the fault/degrade was detected on the working path, and the Path set to 1, indicating that normal traffic is now being transmitted on the recovery path. • Verify that LSR-Z replies with a PCS control packet indicating that it has switched to the recovery path. If this is not received after 2 PSC cycles then send an alarm to the management system.

No.	'821 Patent Claim 1	IETF MPLS-TP System
		When the far-end LSR (in this example LSR-Z) receives the PCS packet informing it that other LSR (LSR-A) has switched, it SHOULD perform the following actions: • Check priority of the request • Switch all traffic addressed to LSR-A to the recovery path only (for 1:1 protection). • Begin transmission of a PCS control packet, using GACH, with the appropriate Request code (either Signal fault or Signal degrade), the Fpath set to 1, to indicate that the fault/degrade was detected on the working path, and the Path set to 1, indicating that traffic is now being transmitted on the recovery path." MPLS-TP Protection at 13–15. "4.3.4. Failure or Degraded condition (Recovery path) If one of the LSRs (for example, LSR-A) detects a failure condition or a serious degradation condition on the recovery path, then it SHOULD take the following actions: • Begin transmission of a PCS control packet with the appropriate Request code (either Signal fault or Signal degrade), the Fpath set to 0, to indicate that the fault/degrade was detected on the recovery path, and the Path set to 0, indicating that traffic is now being forwarded on the working path. Note that this will actually reach the far-end if this is a unidirectional fault or recovery path is possibly in a degraded situation. • The PSC control logic should go into Unavailable state. • All traffic MUST be transmitted on the working path for the duration of the SF/SD condition. When the far-end LSR (in this example LSR-Z) receives the PCS packet informing it that other LSR (LSR-A) has become Unavailable, it SHOULD perform the following actions: • Transmit all traffic on the working path for the duration of the SF/SD condition The PSC Control logic should go into Unavailable state." MPLS-TP Protection at 16.
1[d]	if an entity pair reselection event occurs, reselecting said entity pair from the group consisting of said	IETF MPLS-TP System discloses if an entity pair reselection event occurs, reselecting said entity pair from the group consisting of said entity pair and a replacement entity pair comprising at least one entity distinct from the entities of said entity pair.

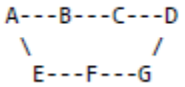
No.	'821 Patent Claim 1	IETF MPLS-TP System
	entity pair and a replacement entity pair comprising at least one entity distinct from the entities of said entity pair,	For example, RFC 4872 discloses that the plurality of MPLS LSPs, i.e., transport entities, are selected in pairs of either working and protection LSPs, or working and restoration LSPs. RFC 4872 further discloses the use of recovery signaling to reroute traffic to a different LSP for a LSP pair in response to a fault. For example, RFC 4872 discloses a secondary LSP designated by a "S" secondary bit, that can replace the protection and/or restoration LSP in response to a fault or command. Thus, at least under the apparent claim scope alleged by Orckit's Infringement Disclosures, this limitation is met. **14.1. Format** The format of the PROTECTION Object (Class-Num = 37, C-Type = 2) is as follows:  Secondary (S): 1 bit When set to 1, this bit indicates that the requested LSP is a secondary LSP. When set to 0 (default), it indicates that the requested LSP is a primary LSP. Protecting (P): 1 bit When set to 1, this bit indicates that the requested LSP is a protecting LSP. When set to 0 (default), it indicates that the requested LSP is a working LSP. The combination, S set to 1 with P set to 0 is not valid. RFC4872 at 31 (annotation added).

No.	'821 Patent Claim 1	IETF MPLS-TP System
		"4.2.1. LSP Status The following bits are used in determining resource allocation and status of the LSP within the group of LSPs forming the protected entity: - S (Secondary) bit: enables distinction between primary and secondary LSPs. A primary LSP is a fully established LSP for which the resource allocation has been committed at the data plane (i.e., full cross-connection has been performed). Both working and protecting LSPs can be primary LSPs. A secondary LSP is an LSP that has been provisioned in the control plane only, and for which resource selection MAY have been done but for which the resource allocation has not been committed at the data plane (for instance, no cross-connection has been performed). Therefore, a secondary LSP is not immediately available to carry any traffic (thus requiring additional signaling to be available). A secondary LSP can only be a protecting LSP. The (data plane) resources allocated for a secondary LSP MAY be used by other LSPs until the primary LSP fails over to the secondary LSP. - P (Protecting) bit: enables distinction between working and protecting LSPs. A working LSP must be a primary LSP whilst a protecting LSP can be either a primary or a secondary LSP. When protecting LSP(s) are associated with working LSP(s), one also refers to the latter as protected LSPs." RFC 4872 at 7. "D. Requested switch for normal traffic: Recovery signaling is initiated that switches normal traffic to the recovery LSP following the procedures defined in Section 6, 7, 8, and 9. This happens unless a fault condition exists on other LSPs or spans (including the recovery LSP), or a switch command of equal or higher priority is in effect. E. Requested switch for recovery LSP: Recovery signaling is initiated that switches normal traffic to the working LSP following the procedure defined in Section 12. This request is executed except if a fault condition exists on the working LSP or an equal or higher priority switch command is in effect." RFC 4872 at 31. "When a failure is detected by one or both endpoints of the LSP, both endpoints must select traffic from the other LSP. This action must be coordinated between node A and D. From this perspective, 1+1 bidirectional protection can be seen as a coordinated protection- switching mechanism between both endpoints.

No.	'821 Patent Claim 1	IETF MPLS-TP System
		Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 11. "1:1 Protection with Extra-Traffic The most common case of end-to-end 1:N protection is to establish, between the same endpoints, an end-to-end working LSP (thus, N = 1) and a dedicated end-to-end protecting LSP that are mutually link/ node/SRLG disjoint. This protects against working LSP failure(s)." RFC 4872 at 13. "An approach to reduce recovery resource requirements is to have protection LSPs sharing network resources when the working LSPs that they protect are physically (i.e., link, node, SRLG, etc.) disjoint. This mechanism is referred to as shared mesh restoration and is described in [RFC4426]. Shared-mesh restoration can be seen as a particular case of pre-planned LSP rerouting (see Section 8) that reduces the recovery resource requirements by allowing multiple protecting LSPs to share common link and node resources. Here also, the recovery resources for the protecting LSPs are pre-reserved during the provisioning phase, thus an explicit signaling action is required to activate (i.e., commit resource allocation at the data plane) a specific protecting LSP instantiated during the (pre-) provisioning phase. This requires restoration signaling along the protecting LSP." RFC 4872 at 20. "1. Introduction Generalized Multi-Protocol Label Switching (GMPLS) extends MPLS to include support for Layer-2 Switch Capable (L2SC), Time-Division Multiplex (TDM), Lambda Switch Capable (LSC), and Fiber Switch Capable (FSC) interfaces. GMPLS recovery uses control plane mechanisms (i.e., signaling, routing, and link management mechanisms) to support data plane fault recovery. Note that the analogous (data plane) fault detection mechanisms are required to be present in support of the control plane mechanisms. In this document, the term "recovery" is generically used to denote both protection and restoration; the specific terms "protection" and "restoration" are only used when differentiation is required. The subtle distinction between protection and restoration is made based on the resource allocation done during the recovery phase (see [RFC4427]).

No.	'821 Patent Claim 1	IETF MPLS-TP System
		A functional description of GMPLS recovery is provided in [RFC4426] and should be considered as a companion document. The present document describes the protocol-specific procedures for GMPLS RSVP- TE (Resource ReSerVation Protocol - Traffic Engineering) signaling (see [RFC3473]) to support end-to-end recovery. End-to-end recovery refers to the recovery of an entire LSP from its head-end (ingress node endpoint) to its tail-end (egress node endpoint). With end-to- end recovery, working LSPs are assumed to be resource-disjoint (where a resource is a link, node, or Shared Risk Link Group (SRLG)) in the network so that they do not share any failure probability, but this is not mandatory. With respect to a given set of network resources, a pair of working/protecting LSPs SHOULD be resource disjoint in case of dedicated recovery type (see below). On the other hand, in case of shared recovery (see below), a group of working LSPs SHOULD be mutually resource-disjoint in order to allow for a (single and commonly) shared protecting LSP, itself resource-disjoint from each of the working LSPs. Note that resource disjointness is a necessary (but not sufficient) condition to ensure LSP recoverability. The present document addresses four types of end-to-end LSP recovery: 1) 1+1 (unidirectional/bidirectional) protection, 2) 1:N (N >= 1) LSP protection with extra-traffic, 3) pre-planned LSP rerouting without extra-traffic (including shared mesh), and 4) full LSP rerouting. 13) The simplest notion of end-to-end LSP protection is 1+1 unidirectional protection. Using this type of protection, a protecting LSP is signaled over a dedicated resource-disjoint alternate path to protect an associated working LSP. Normal traffic is simultaneously sent on both LSPs and a selector is used at the egress node to receive traffic from one of the LSPs. If a failure occurs along one of the LSPs, the egress node selects the traffic from the valid LSP. No coordination is required between the end nodes when a failure/switchover occurs. In 1+1 bidirectional protection, a protecting LSP is signaled over a dedicated resource-disjoint alternate path to protect the working LSP. Normal traffic is simultaneously sent on both LSPs (in both directions), and a selector is used at both ingress/egress nodes to receive traffic from the same LSP. This requires coordination between the end-nodes when switching to the protecting LSP.

No.	'821 Patent Claim 1	IETF MPLS-TP System
		14) In 1:N (N >= 1) protection with extra-traffic, the protecting LSP is a fully provisioned and resource-disjoint LSP from the N working LSPs, that allows for carrying extra-traffic. The N working LSPs MAY be mutually resource-disjoint. Coordination between end-nodes is required when switching from one of the working LSPs to the protecting LSP. As the protecting LSP is fully provisioned, default operations during protection switching are specified for a protecting LSP carrying extra-traffic, but this is not mandatory. Note that M:N protection is out of scope of this document (though mechanisms it defines may be extended to cover it). 15) Pre-planned LSP rerouting (or restoration) relies on the establishment between the same pair of end-nodes of a working LSP and a protecting LSP that is link/node/SRLG disjoint from the working one. Here, the recovery resources for the protecting LSP are pre-reserved but explicit action is required to activate (i.e., commit resource allocation at the data plane) a specific protecting LSP instantiated during the (pre-)provisioning phase. Since the protecting LSP is not "active" (i.e., fully instantiated), it cannot carry any extra-traffic. This does not mean that the corresponding resources cannot be used by other LSPs. Therefore, this mechanism protects against working LSP(s) failure(s) but requires activation of the protecting LSP after working LSP failure occurrence. This requires restoration signaling along the protecting path. "Shared-mesh" restoration can be seen as a particular case of pre-planned LSP rerouting that reduces the recovery resource requirements by allowing multiple protecting LSPs to share common link and node resources. The recovery resources are pre-reserved but explicit action is required to activate (i.e., commit resource allocation at the data plane) a specific protecting LSP instantiated during the (pre-) provisioning phase. This procedure requires restoration signaling along the protecting path. Note that in both cases, bandwidth pre-reserved for a protecting (but not activated) LSP can be made available for carrying extra traffic. LSPs for extra-traffic (with lower holding priority than the protecting LSP) can then be established using the bandwidth pre-reserved for the protecting LSP. Also, any lower priority LSP that use the pre-reserved resources for the protecting LSP(s) must be preempted during the activation of the protecting LSP. 16) Full LSP rerouting (or restoration) switches normal traffic to an alternate LSP that is not even partially established until after the working LSP failure occurs. The new

No.	'821 Patent Claim 1	IETF MPLS-TP System
		alternate route is selected at the LSP head-end node, it may reuse resources of the failed LSP at intermediate nodes and may include additional intermediate nodes and/or links." " RFC 4872 at 3–5. "Consider the following network topology:  The paths [A,B,C,D] and [A,E,F,G,D] are node and link disjoint, ignoring the ingress/egress nodes A and D. A 1+1 protected path is established from A to D over [A,B,C,D] and [A,E,F,G,D], and traffic is transmitted simultaneously over both component paths (i.e., LSPs). During the provisioning phase, both LSPs are fully instantiated (and thus activated) so that no resource sharing can be done along the protecting LSP (nor can any extra-traffic be transported). It is also RECOMMENDED to set the N bit since no protection-switching signaling is assumed in this case. When a failure occurs (say, at node B) and is detected at end-node D, the receiver at D selects the normal traffic from the other LSP. From this perspective, 1+1 unidirectional protection can be seen as an uncoordinated protection-switching mechanism acting independently at both endpoints. Also, for the LSP under failure condition, it is RECOMMENDED to not set the Path_State_Removed Flag of the ERROR_SPEC object (see [RFC3473]) upon PathErr message generation. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 9. "When a failure is detected by one or both endpoints of the LSP, both endpoints must select traffic from the other LSP. This action must be coordinated between node A and D. From this perspective, 1+1 bidirectional protection can be seen as a coordinated protection- switching mechanism between both endpoints. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 11.

No.	'821 Patent Claim 1	IETF MPLS-TP System
		"(Full) LSP Rerouting LSP rerouting, on the other hand, switches normal traffic to an alternate LSP that is fully established only after failure occurrence. The new (alternate) route is selected at the LSP head- end and may reuse intermediate nodes included in the original route; it may also include additional intermediate nodes. For strict-hop routing, TE requirements can be directly applied to the route computation, and the failed node or link can be avoided. However, if the failure occurred within a loose-routed hop, the head-end node may not have enough information to reroute the LSP around the failure. Crankback signaling (see [CRANK]) and route exclusion techniques (see [RFC4874]) MAY be used in this case. The alternate route MAY be either computed on demand (that is, when the failure occurs; this is referred to as full LSP rerouting) or pre-computed and stored for use when the failure is reported. The latter offers faster restoration time. There is, however, a risk that the alternate route will become out of date through other changes in the network; this can be mitigated to some extent by periodic recalculation of idle alternate routes. (Full) LSP rerouting will be initiated by the head-end node that has either detected the LSP failure or received a Notify message and/or a PathErr message with the new error code/sub-code "Notify Error/LSP Locally Failed" for this LSP. The new LSP resources can be established using the make-before-break mechanism, where the new LSP is set up before the old LSP is torn down. This is done by using the mechanisms of the SESSION_ATTRIBUTE object and the Shared-Explicit (SE) reservation style (see [RFC3209]). Both the new and old LSPs can share resources at common nodes. Note that the make-before-break mechanism is not used to avoid disruption to the normal traffic flow (the latter has already been broken by the failure that is being repaired). However, it is valuable to retain the resources allocated on the original LSP that will be reused by the new alternate LSP." RFC 4872 at 25. "The procedure is as follows: 13) The initiating (source) node sends the normal traffic onto both the working and the protecting LSPs. Once completed, the source node sends reliably a Notify message to

No.	'821 Patent Claim 1	IETF MPLS-TP System
		the destination with the new error code/sub-code "Notify Error/LSP Recovered" (Switchback Request). This Notify message includes the MESSAGE_ID object. The ACK_Desired flag MUST be set in this object to request the receiver to send an acknowledgment for the message (see [RFC2961]). 14) Upon receipt of this message, the destination selects the traffic from the working LSP. At the same time, it transmits the traffic onto both the working and protecting LSP. The destination then sends reliably a Notify message to the source confirming the completion of the operation. This message includes the MESSAGE_ID_ACK object to acknowledge reception of the received Notify message. This Notify message also includes the MESSAGE_ID object. The ACK_Desired flag MUST be set in this object to request the receiver to send an acknowledgment for the message (see [RFC2961]). 15) When the source node receives this Notify message, it switches to receive traffic from the working LSP. The source node then sends an Ack message to the destination node confirming that the LSP has been reverted." RFC 4872 at 27. For example, MPLS-TP Protection discloses that responsive to certain protection switching triggers, *e.g.*, a server layer failure, a loss of OAM connectivity, or signal degradation, switching from either a working path to a recovery path or vice versa. At least under Orckit's Infringement Disclosures, this limitation is met. To the extent that the IETF MPLS-TP System is found to not meet this limitation, if an entity pair reselection event occurs, reselecting said entity pair from the group consisting of said entity pair and a replacement entity pair comprising at least one entity distinct from the entities of said entity pair would have been obvious to a person having ordinary skill in the art, as explained below. "3. Protection switching logic 3.1. Protection switching trigger mechanisms The protection switching should be initiated in reaction to any of the following triggers: • Server layer indication - if the MPLS-TP server layer detects a failure within its own layer, or due to a failure of its server layer (e.g. the physical layer) notifies the MPLS-

No.	'821 Patent Claim 1	IETF MPLS-TP System
		TP layer that a failure has been detected. It should be noted that this trigger will be controlled by a holdoff-timer that should be configured to allow the server layer to apply any protection procedures that may be relevant. • OAM signalling - if, for example, OAM continuity and connectivity verification tools detect that there is a loss of continuity or mis-connectivity or performance monitoring indicates a degradation of the utility of the working path for the current transport path. In cases of signal degradation, switching to the recovery path SHOULD only be activated if the recovery path can guarantee better conditions than the degraded working path. • Control plane - if there is a control plane active in the network (either signaling or routing), it MAY trigger protection switching based on conditions detected by the control plane. If the control-plane is based on GMPLS [RFC3945] then the recovery process should comply with the process described in [RFC4872]. • Operator command - the network operator may issue commands that trigger protection switching. The commands that are supported include - Forced Switch, Manual Switch, Clear, Lockout of Protection, (see definitions in [RFC4427]).” MPLS-TP at 6–7. “3.2. Protection switching control logical architecture Protection switching processes the triggers described above together with the inputs received from the far-end LSR. These inputs cause the LSR to take certain actions, e.g. switching the Selector Bridge to select the working or recovery path, and to transmit different protocol messages.

No.	'821 Patent Claim 1	IETF MPLS-TP System

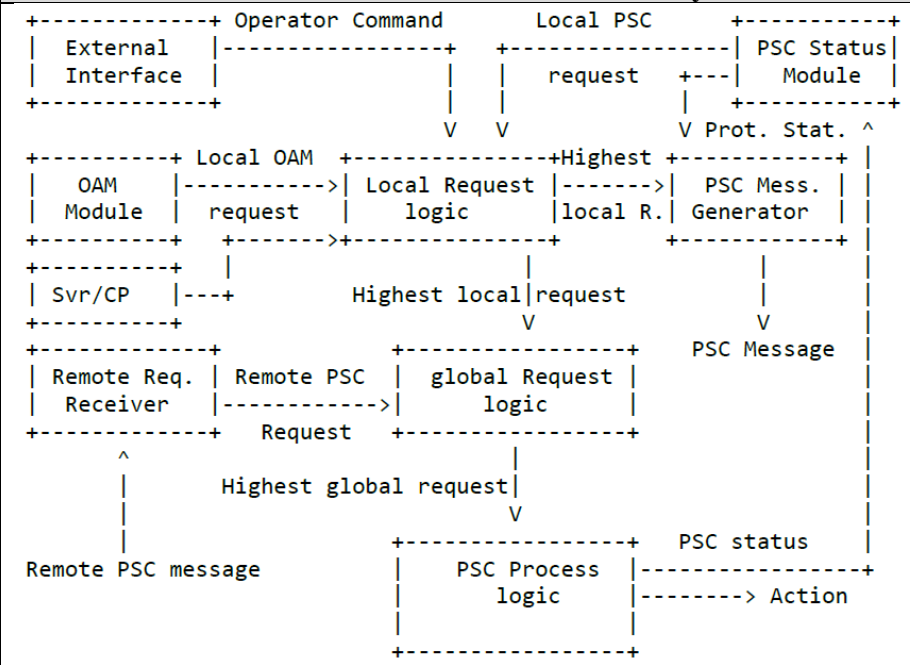
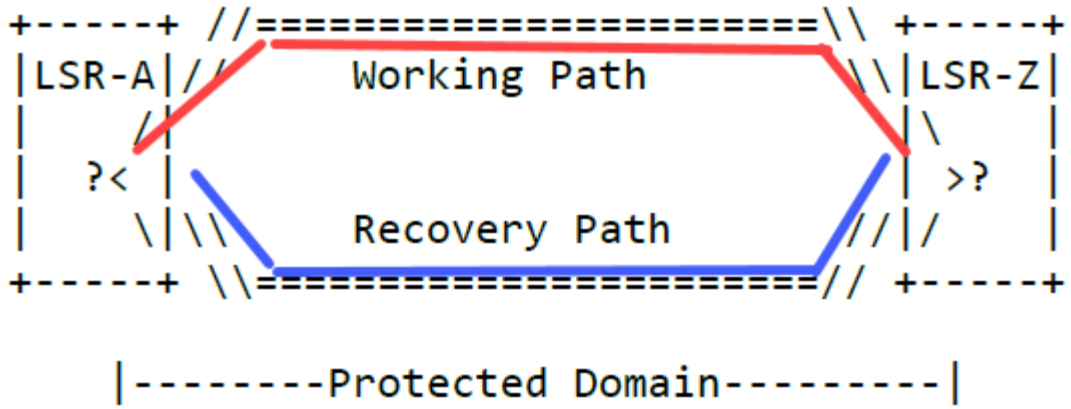


Figure 1: Protection switching control logic

Figure 1 describes the logical architecture of the protection switching control. The Local Request logic unit accepts the triggers from the OAM, external operator commands, and from the local control plane (when present) and determines the highest priority request. This high-priority request is passed to both the PSC Message generator, that will generate the appropriate protocol message to be sent to the far-end LSR, and the Global Request logic, that will cross-check this local request with the information received from the far-LSR. The Global Request logic then processes these two PSC requests that determines the highest priority request that is passed to the PSC Process logic. The PSC Process logic uses this input to determine what actions need to be taken, e.g. switching the Selector Bridge, and the current status of the protection domain." MPLS-TP Protection at 7–8.

"4.2.4. Active path indicator (Path)

No.	'821 Patent Claim 1	IETF MPLS-TP System
		The Path field of the PSC control SHALL be used to indicate which path the source MEP is currently using for data transmission. The MEP should compare the value of this bit with the path that is locally selected for data transmission to verify that there is no inconsistency between the two end-points of the protected domain. If an inconsistency is detected then an alarm should be raised. The following are the possible values: • 0: indicates that normal traffic is being transmitted on the Working path. • 1: indicates the Recovery path is being used to transmit the normal traffic from the Working path. 2-255: for future extensions" MPLS-TP Protection at 12–13. "Principles of Operation In all of the following sub-sections, assume a protected domain between LSR-A and LSR-Z, using paths W (working) and R (recovery) as shown in figure 4." MPLS-TP Protection at 13.  MPLS-TP Protection Figure 4 (annotation added). "4.3.1 PSC States 4.3.1.1 Normal State

No.	'821 Patent Claim 1	IETF MPLS-TP System
		When the protected domain has no special condition in effect, the ingress LSR SHOULD forward the user data along the working path, and, in the case of 1+1 protection, the Permanent Bridge will bridge the data to the recovery path as well. The receiving LSR SHOULD read the data from the working path. The ingress LSR MAY transmit a No Request PSC packet with the Path field set to 0 indicating that the normal data traffic should be read from the working path. 4.3.1.2. Protecting State When the protection mechanism has been triggered and the protected domain has performed a protection switch, the domain is in the protecting state. In this state the normal data traffic is transmitted and received on the recovery path. If the protection domain is currently in a protecting state, then the LSRs SHOULD NOT accept a Manual Switch request. If the protection domain is currently in a protecting state, and a Forced Switch is requested then the normal traffic SHALL continue to be transmitted on the recovery path even if the original protection trigger is cleared, and the Forced Switch condition will be signalled by the PSC messages. 4.3.1.3. Unavailable State When the recovery path is unavailable - either as a result of a Lockout operator command (see section 4.3.3), or as a result of a SF or SD detected on the recovery path (see section 4.3.4) - then the protection domain is in the unavailable state. In this state, the normal traffic is transmitted and received on the working path. While in unavailable state any event that would trigger a protection switching SHOULD be ignored with the following exception - If a Signal Degrade request is received, then protection switching will be activated only if the recovery path can guarantee a better signal than the working path. The protection domain will exit the unavailable state and revert to the normal state when, either the operator clears the Lockout command or the recovery path recovers from the signal fault or degraded situation.

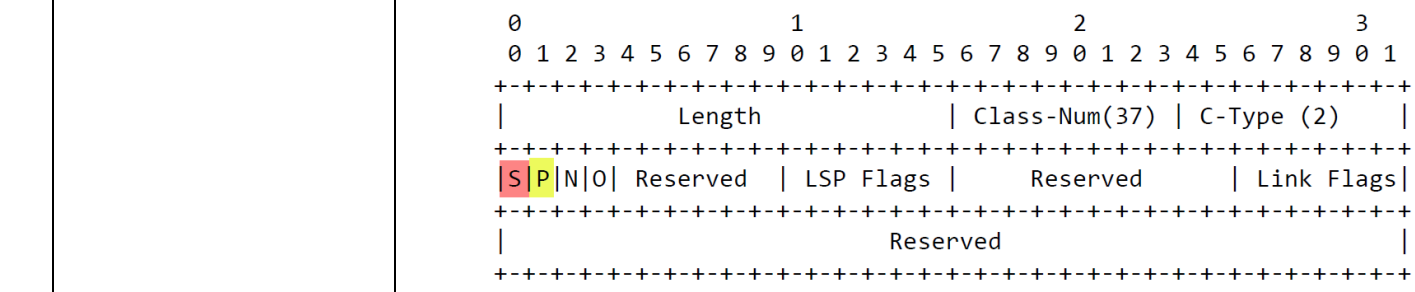
No.	'821 Patent Claim 1	IETF MPLS-TP System
		Both ends will resume sending the PCS packets over the recovery path, as a result of this recovery. 4.3.2. Failure or Degraded condition (Working path) If one of the LSRs (for example, LSR-A) detects a failure condition or a serious degradation condition on the working path that warrants invoking protection switching, then it SHOULD take the following actions: • (For 1:1 protection) Switch all traffic for LSR-Z to the recovery path only. • Transmit a PCS control packet, using GACH, with the appropriate Request code (either Signal fault or Signal degrade), the Fpath set to 1, to indicate that the fault/degrade was detected on the working path, and the Path set to 1, indicating that normal traffic is now being transmitted on the recovery path. • Verify that LSR-Z replies with a PCS control packet indicating that it has switched to the recovery path. If this is not received after 2 PSC cycles then send an alarm to the management system. When the far-end LSR (in this example LSR-Z) receives the PCS packet informing it that other LSR (LSR-A) has switched, it SHOULD perform the following actions: • Check priority of the request • Switch all traffic addressed to LSR-A to the recovery path only (for 1:1 protection). • Begin transmission of a PCS control packet, using GACH, with the appropriate Request code (either Signal fault or Signal degrade), the Fpath set to 1, to indicate that the fault/degrade was detected on the working path, and the Path set to 1, indicating that traffic is now being transmitted on the recovery path." MPLS-TP Protection at 13–15. "4.3.4. Failure or Degraded condition (Recovery path) If one of the LSRs (for example, LSR-A) detects a failure condition or a serious degradation condition on the recovery path, then it SHOULD take the following actions:

No.	'821 Patent Claim 1	IETF MPLS-TP System
		• Begin transmission of a PCS control packet with the appropriate Request code (either Signal fault or Signal degrade), the Fpath set to 0, to indicate that the fault/degrade was detected on the recovery path, and the Path set to 0, indicating that traffic is now being forwarded on the working path. Note that this will actually reach the far-end if this is a unidirectional fault or recovery path is possibly in a degraded situation. • The PSC control logic should go into Unavailable state. • All traffic MUST be transmitted on the working path for the duration of the SF/SD condition. When the far-end LSR (in this example LSR-Z) receives the PCS packet informing it that other LSR (LSR-A) has become Unavailable, it SHOULD perform the following actions: • Transmit all traffic on the working path for the duration of the SF/SD condition The PSC Control logic should go into Unavailable state." MPLS-TP Protection at 16.
1[e]	wherein said entity pair reselection event is selected from a group consisting of adding an entity to said plurality of transport entities, removing an entity from said plurality of transport entities, an operational status change for one of said plurality of transport entities, and a change in overall cost for one of said plurality of transport entities.	IETF MPLS-TP System discloses wherein said entity pair reselection event is selected from a group consisting of adding an entity to said plurality of transport entities, removing an entity from said plurality of transport entities, an operational status change for one of said plurality of transport entities, and a change in overall cost for one of said plurality of transport entities. For example, RFC 4872 discloses that the plurality of MPLS LSPs, i.e., transport entities, are selected in pairs of either working and protection LSPs, or working and restoration LSPs. RFC 4872 further discloses the use of recovery signaling to reroute traffic to a different LSP for a LSP pair in response to a fault. RFC 4872 further discloses a "secondary" LSP that can be used a replacement to the protection and/or restoration LSP. Thus, at least under the apparent claim scope alleged by Orckit's Infringement Disclosures, this limitation is met. To the extent that the IETF MPLS-TP System is found to not meet this limitation, selecting an entity pair from said plurality of transport entities based at least in part upon said overall cost would have been obvious to a person having ordinary skill in the art, as explained below.

No.	'821 Patent Claim 1	IETF MPLS-TP System

14.1. Format

The format of the PROTECTION Object (Class-Num = 37, C-Type = 2) is as follows:



Secondary (S): 1 bit

When set to 1, this bit indicates that the requested LSP is a secondary LSP. When set to 0 (default), it indicates that the requested LSP is a primary LSP.

Protecting (P): 1 bit

When set to 1, this bit indicates that the requested LSP is a protecting LSP. When set to 0 (default), it indicates that the requested LSP is a working LSP. The combination, S set to 1 with P set to 0 is not valid.

RFC4872 at 31 (annotation added).

"4.2.1. LSP Status
The following bits are used in determining resource allocation and status of the LSP within the group of LSPs forming the protected entity:
- S (Secondary) bit: enables distinction between primary and secondary LSPs. A primary LSP is a fully established LSP for which the resource allocation has been committed at the data plane (i.e., full cross-connection has been performed). Both working and protecting LSPs can be primary LSPs. A secondary LSP is an LSP that has been provisioned in the control plane only,

No.	'821 Patent Claim 1	IETF MPLS-TP System
		and for which resource selection MAY have been done but for which the resource allocation has not been committed at the data plane (for instance, no cross-connection has been performed). Therefore, a secondary LSP is not immediately available to carry any traffic (thus requiring additional signaling to be available). A secondary LSP can only be a protecting LSP. The (data plane) resources allocated for a secondary LSP MAY be used by other LSPs until the primary LSP fails over to the secondary LSP. - P (Protecting) bit: enables distinction between working and protecting LSPs. A working LSP must be a primary LSP whilst a protecting LSP can be either a primary or a secondary LSP. When protecting LSP(s) are associated with working LSP(s), one also refers to the latter as protected LSPs." RFC 4872 at 7. "D. Requested switch for normal traffic: Recovery signaling is initiated that switches normal traffic to the recovery LSP following the procedures defined in Section 6, 7, 8, and 9. This happens unless a fault condition exists on other LSPs or spans (including the recovery LSP), or a switch command of equal or higher priority is in effect. E. Requested switch for recovery LSP: Recovery signaling is initiated that switches normal traffic to the working LSP following the procedure defined in Section 12. This request is executed except if a fault condition exists on the working LSP or an equal or higher priority switch command is in effect." RFC 4872 at 31. "When a failure is detected by one or both endpoints of the LSP, both endpoints must select traffic from the other LSP. This action must be coordinated between node A and D. From this perspective, 1+1 bidirectional protection can be seen as a coordinated protection- switching mechanism between both endpoints. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 11. "1:1 Protection with Extra-Traffic The most common case of end-to-end 1:N protection is to establish, between the same endpoints, an end-to-end working LSP (thus, N = 1) and a dedicated end-to-end protecting LSP that are mutually link/ node/SRLG disjoint. This protects against working LSP failure(s)." RFC 4872 at 13.

No.	'821 Patent Claim 1	IETF MPLS-TP System
		"An approach to reduce recovery resource requirements is to have protection LSPs sharing network resources when the working LSPs that they protect are physically (i.e., link, node, SRLG, etc.) disjoint. This mechanism is referred to as shared mesh restoration and is described in [RFC4426]. Shared-mesh restoration can be seen as a particular case of pre-planned LSP rerouting (see Section 8) that reduces the recovery resource requirements by allowing multiple protecting LSPs to share common link and node resources. Here also, the recovery resources for the protecting LSPs are pre-reserved during the provisioning phase, thus an explicit signaling action is required to activate (i.e., commit resource allocation at the data plane) a specific protecting LSP instantiated during the (pre-) provisioning phase. This requires restoration signaling along the protecting LSP." RFC 4872 at 20. "1. Introduction Generalized Multi-Protocol Label Switching (GMPLS) extends MPLS to include support for Layer-2 Switch Capable (L2SC), Time-Division Multiplex (TDM), Lambda Switch Capable (LSC), and Fiber Switch Capable (FSC) interfaces. GMPLS recovery uses control plane mechanisms (i.e., signaling, routing, and link management mechanisms) to support data plane fault recovery. Note that the analogous (data plane) fault detection mechanisms are required to be present in support of the control plane mechanisms. In this document, the term "recovery" is generically used to denote both protection and restoration; the specific terms "protection" and "restoration" are only used when differentiation is required. The subtle distinction between protection and restoration is made based on the resource allocation done during the recovery phase (see [RFC4427]). A functional description of GMPLS recovery is provided in [RFC4426] and should be considered as a companion document. The present document describes the protocol-specific procedures for GMPLS RSVP- TE (Resource ReSerVation Protocol - Traffic Engineering) signaling (see [RFC3473]) to support end-to-end recovery. End-to-end recovery refers to the recovery of an entire LSP from its head-end (ingress node endpoint) to its tail-end (egress node endpoint). With end-to- end recovery, working LSPs are assumed to be resource-disjoint (where a resource is a link, node, or Shared Risk Link Group (SRLG)) in the network so that they do not share any failure probability, but this is not mandatory. With respect to a given

No.	'821 Patent Claim 1	IETF MPLS-TP System
		set of network resources, a pair of working/protecting LSPs SHOULD be resource disjoint in case of dedicated recovery type (see below). On the other hand, in case of shared recovery (see below), a group of working LSPs SHOULD be mutually resource-disjoint in order to allow for a (single and commonly) shared protecting LSP, itself resource-disjoint from each of the working LSPs. Note that resource disjointness is a necessary (but not sufficient) condition to ensure LSP recoverability. The present document addresses four types of end-to-end LSP recovery: 1) 1+1 (unidirectional/bidirectional) protection, 2) 1:N (N >= 1) LSP protection with extra-traffic, 3) pre-planned LSP rerouting without extra-traffic (including shared mesh), and 4) full LSP rerouting. 17) The simplest notion of end-to-end LSP protection is 1+1 unidirectional protection. Using this type of protection, a protecting LSP is signaled over a dedicated resource-disjoint alternate path to protect an associated working LSP. Normal traffic is simultaneously sent on both LSPs and a selector is used at the egress node to receive traffic from one of the LSPs. If a failure occurs along one of the LSPs, the egress node selects the traffic from the valid LSP. No coordination is required between the end nodes when a failure/switchover occurs. In 1+1 bidirectional protection, a protecting LSP is signaled over a dedicated resource-disjoint alternate path to protect the working LSP. Normal traffic is simultaneously sent on both LSPs (in both directions), and a selector is used at both ingress/egress nodes to receive traffic from the same LSP. This requires coordination between the end-nodes when switching to the protecting LSP. 18) In 1:N (N >= 1) protection with extra-traffic, the protecting LSP is a fully provisioned and resource-disjoint LSP from the N working LSPs, that allows for carrying extra-traffic. The N working LSPs MAY be mutually resource-disjoint. Coordination between end-nodes is required when switching from one of the working LSPs to the protecting LSP. As the protecting LSP is fully provisioned, default operations during protection switching are specified for a protecting LSP carrying extra-traffic, but this is not mandatory. Note that M:N protection is out of scope of this document (though mechanisms it defines may be extended to cover it).

No.	'821 Patent Claim 1	IETF MPLS-TP System
		19) Pre-planned LSP rerouting (or restoration) relies on the establishment between the same pair of end-nodes of a working LSP and a protecting LSP that is link/node/SRLG disjoint from the working one. Here, the recovery resources for the protecting LSP are pre-reserved but explicit action is required to activate (i.e., commit resource allocation at the data plane) a specific protecting LSP instantiated during the (pre-)provisioning phase. Since the protecting LSP is not "active" (i.e., fully instantiated), it cannot carry any extra-traffic. This does not mean that the corresponding resources cannot be used by other LSPs. Therefore, this mechanism protects against working LSP(s) failure(s) but requires activation of the protecting LSP after working LSP failure occurrence. This requires restoration signaling along the protecting path. "Shared-mesh" restoration can be seen as a particular case of pre-planned LSP rerouting that reduces the recovery resource requirements by allowing multiple protecting LSPs to share common link and node resources. The recovery resources are pre-reserved but explicit action is required to activate (i.e., commit resource allocation at the data plane) a specific protecting LSP instantiated during the (pre-) provisioning phase. This procedure requires restoration signaling along the protecting path. Note that in both cases, bandwidth pre-reserved for a protecting (but not activated) LSP can be made available for carrying extra traffic. LSPs for extra-traffic (with lower holding priority than the protecting LSP) can then be established using the bandwidth pre-reserved for the protecting LSP. Also, any lower priority LSP that use the pre-reserved resources for the protecting LSP(s) must be preempted during the activation of the protecting LSP. 20) Full LSP rerouting (or restoration) switches normal traffic to an alternate LSP that is not even partially established until after the working LSP failure occurs. The new alternate route is selected at the LSP head-end node, it may reuse resources of the failed LSP at intermediate nodes and may include additional intermediate nodes and/or links." " RFC 4872 at 3–5. "Consider the following network topology:

No.	'821 Patent Claim 1	IETF MPLS-TP System
		 The paths [A,B,C,D] and [A,E,F,G,D] are node and link disjoint, ignoring the ingress/egress nodes A and D. A 1+1 protected path is established from A to D over [A,B,C,D] and [A,E,F,G,D], and traffic is transmitted simultaneously over both component paths (i.e., LSPs). During the provisioning phase, both LSPs are fully instantiated (and thus activated) so that no resource sharing can be done along the protecting LSP (nor can any extra-traffic be transported). It is also RECOMMENDED to set the N bit since no protection-switching signaling is assumed in this case. When a failure occurs (say, at node B) and is detected at end-node D, the receiver at D selects the normal traffic from the other LSP. From this perspective, 1+1 unidirectional protection can be seen as an uncoordinated protection-switching mechanism acting independently at both endpoints. Also, for the LSP under failure condition, it is RECOMMENDED to not set the Path_State_Removed Flag of the ERROR_SPEC object (see [RFC3473]) upon PathErr message generation. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 9. "When a failure is detected by one or both endpoints of the LSP, both endpoints must select traffic from the other LSP. This action must be coordinated between node A and D. From this perspective, 1+1 bidirectional protection can be seen as a coordinated protection- switching mechanism between both endpoints. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 11. "(Full) LSP Rerouting LSP rerouting, on the other hand, switches normal traffic to an alternate LSP that is fully established only after failure occurrence. The new (alternate) route is selected at the LSP head- end and may reuse intermediate nodes included in the original route; it may also include

No.	'821 Patent Claim 1	IETF MPLS-TP System
		additional intermediate nodes. For strict-hop routing, TE requirements can be directly applied to the route computation, and the failed node or link can be avoided. However, if the failure occurred within a loose-routed hop, the head-end node may not have enough information to reroute the LSP around the failure. Crankback signaling (see [CRANK]) and route exclusion techniques (see [RFC4874]) MAY be used in this case. The alternate route MAY be either computed on demand (that is, when the failure occurs; this is referred to as full LSP rerouting) or pre-computed and stored for use when the failure is reported. The latter offers faster restoration time. There is, however, a risk that the alternate route will become out of date through other changes in the network; this can be mitigated to some extent by periodic recalculation of idle alternate routes. (Full) LSP rerouting will be initiated by the head-end node that has either detected the LSP failure or received a Notify message and/or a PathErr message with the new error code/sub-code "Notify Error/LSP Locally Failed" for this LSP. The new LSP resources can be established using the make-before-break mechanism, where the new LSP is set up before the old LSP is torn down. This is done by using the mechanisms of the SESSION_ATTRIBUTE object and the Shared-Explicit (SE) reservation style (see [RFC3209]). Both the new and old LSPs can share resources at common nodes. Note that the make-before-break mechanism is not used to avoid disruption to the normal traffic flow (the latter has already been broken by the failure that is being repaired). However, it is valuable to retain the resources allocated on the original LSP that will be reused by the new alternate LSP." RFC 4872 at 25. "The procedure is as follows: 16) The initiating (source) node sends the normal traffic onto both the working and the protecting LSPs. Once completed, the source node sends reliably a Notify message to the destination with the new error code/sub-code "Notify Error/LSP Recovered" (Switchback Request). This Notify message includes the MESSAGE_ID object. The ACK_Desired flag MUST be set in this object to request the receiver to send an acknowledgment for the message (see [RFC2961]). 17) Upon receipt of this message, the destination selects the traffic from the working LSP. At the same time, it transmits the traffic onto both the working and protecting LSP.

No.	'821 Patent Claim 1	IETF MPLS-TP System

The destination then sends reliably a Notify message to the source confirming the completion of the operation. This message includes the MESSAGE_ID_ACK object to acknowledge reception of the received Notify message. This Notify message also includes the MESSAGE_ID object. The ACK_Desired flag MUST be set in this object to request the receiver to send an acknowledgment for the message (see [RFC2961]).

18) When the source node receives this Notify message, it switches to receive traffic from the working LSP.

The source node then sends an Ack message to the destination node confirming that the LSP has been reverted." RFC 4872 at 27.

For example, MPLS-TP Protection discloses an active path indicator field that indicates which LSP and/or pseudo-wire, i.e. network entity, will carry network traffic. MPLS-TP Protection further discloses that the active path indicator will change in response to certain protection switching triggers. MPLS-TP Protection further discloses that the protection switching triggers include, *e.g.*, a server layer failure, a loss of OAM connectivity, or signal degradation, switching from either a working path to a recovery path or vice versa. At least under Orckit's Infringement Disclosures, this limitation is met. To the extent that the IETF MPLS-TP System is found to not meet this limitation, wherein said entity pair reselection event is selected from a group consisting of adding an entity to said plurality of transport entities, removing an entity from said plurality of transport entities, an operational status change for one of said plurality of transport entities, and a change in overall cost for one of said plurality of transport entities would have been obvious to a person having ordinary skill in the art, as explained below.

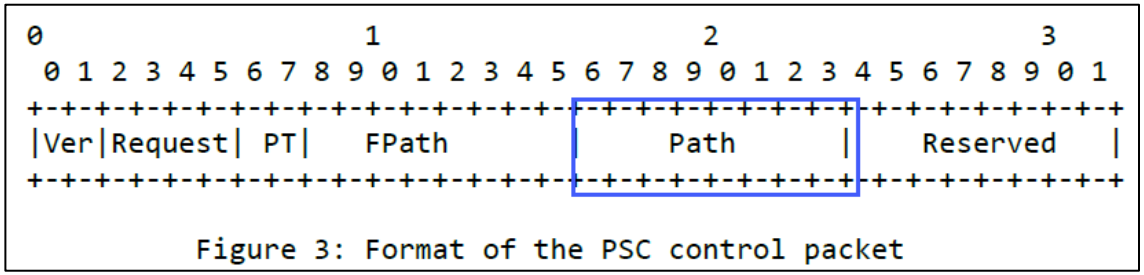


Figure 3: Format of the PSC control packet

MPLS-TP Figure 3 (annotation added).

No.	'821 Patent Claim 1	IETF MPLS-TP System
		"3. Protection switching logic 3.1. Protection switching trigger mechanisms The protection switching should be initiated in reaction to any of the following triggers: • Server layer indication - if the MPLS-TP server layer detects a failure within its own layer, or due to a failure of its server layer (e.g. the physical layer) notifies the MPLS-TP layer that a failure has been detected. It should be noted that this trigger will be controlled by a holdoff-timer that should be configured to allow the server layer to apply any protection procedures that may be relevant. • OAM signalling - if, for example, OAM continuity and connectivity verification tools detect that there is a loss of continuity or mis-connectivity or performance monitoring indicates a degradation of the utility of the working path for the current transport path. In cases of signal degradation, switching to the recovery path SHOULD only be activated if the recovery path can guarantee better conditions than the degraded working path. • Control plane - if there is a control plane active in the network (either signaling or routing), it MAY trigger protection switching based on conditions detected by the control plane. If the control-plane is based on GMPLS [RFC3945] then the recovery process should comply with the process described in [RFC4872]. • Operator command - the network operator may issue commands that trigger protection switching. The commands that are supported include - Forced Switch, Manual Switch, Clear, Lockout of Protection, (see definitions in [RFC4427])." MPLS-TP at 6–7. "3.2. Protection switching control logical architecture Protection switching processes the triggers described above together with the inputs received from the far-end LSR. These inputs cause the LSR to take certain actions, e.g. switching the Selector Bridge to select the working or recovery path, and to transmit different protocol messages.

No.	'821 Patent Claim 1	IETF MPLS-TP System

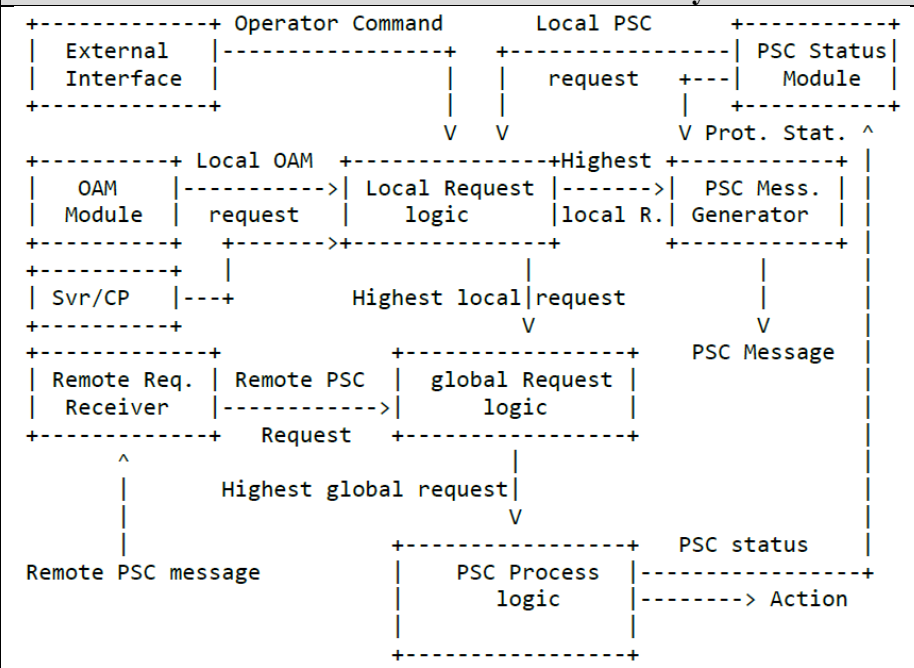
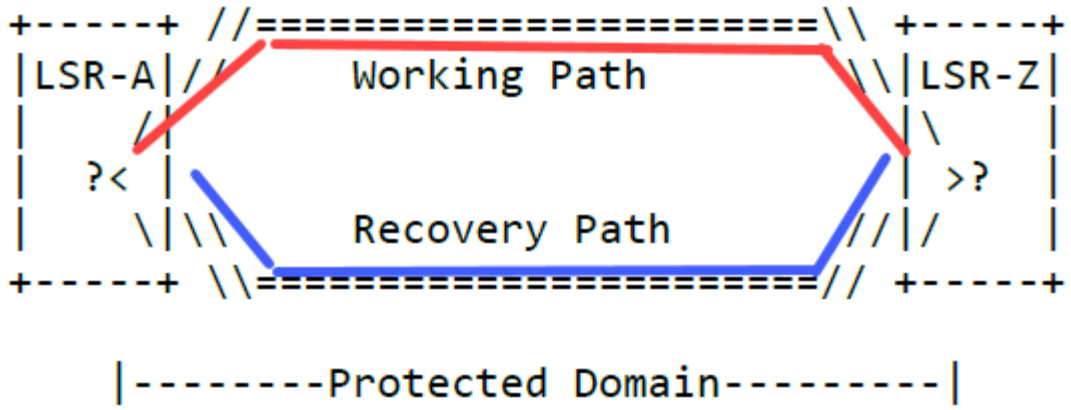


Figure 1: Protection switching control logic

Figure 1 describes the logical architecture of the protection switching control. The Local Request logic unit accepts the triggers from the OAM, external operator commands, and from the local control plane (when present) and determines the highest priority request. This high-priority request is passed to both the PSC Message generator, that will generate the appropriate protocol message to be sent to the far-end LSR, and the Global Request logic, that will cross-check this local request with the information received from the far-LSR. The Global Request logic then processes these two PSC requests that determines the highest priority request that is passed to the PSC Process logic. The PSC Process logic uses this input to determine what actions need to be taken, e.g. switching the Selector Bridge, and the current status of the protection domain." MPLS-TP Protection at 7–8.

"4.2.4. Active path indicator (Path)

No.	'821 Patent Claim 1	IETF MPLS-TP System
		The Path field of the PSC control SHALL be used to indicate which path the source MEP is currently using for data transmission. The MEP should compare the value of this bit with the path that is locally selected for data transmission to verify that there is no inconsistency between the two end-points of the protected domain. If an inconsistency is detected then an alarm should be raised. The following are the possible values: • 0: indicates that normal traffic is being transmitted on the Working path. • 1: indicates the Recovery path is being used to transmit the normal traffic from the Working path. 2-255: for future extensions" MPLS-TP Protection at 12–13. "Principles of Operation In all of the following sub-sections, assume a protected domain between LSR-A and LSR-Z, using paths W (working) and R (recovery) as shown in figure 4." MPLS-TP Protection at 13.  MPLS-TP Protection Figure 4 (annotation added). "4.3.1 PSC States 4.3.1.1 Normal State

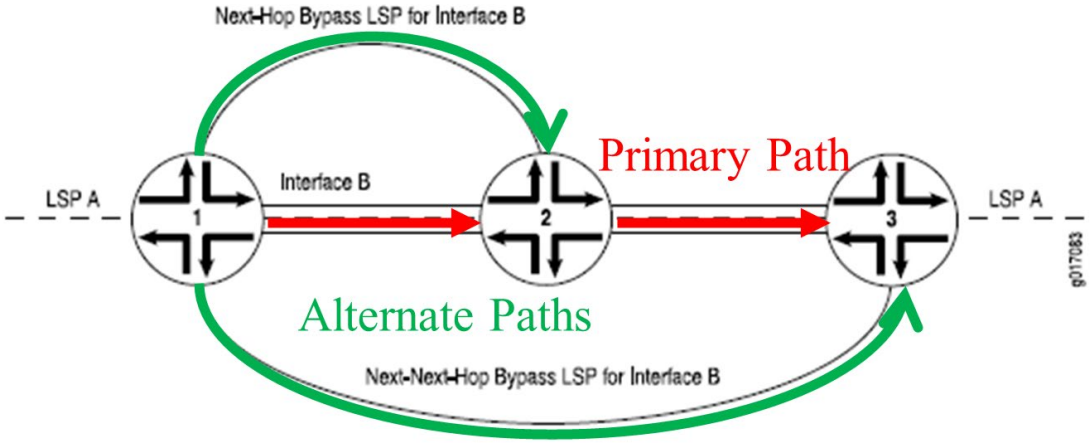
No.	'821 Patent Claim 1	IETF MPLS-TP System
		When the protected domain has no special condition in effect, the ingress LSR SHOULD forward the user data along the working path, and, in the case of 1+1 protection, the Permanent Bridge will bridge the data to the recovery path as well. The receiving LSR SHOULD read the data from the working path. The ingress LSR MAY transmit a No Request PSC packet with the Path field set to 0 indicating that the normal data traffic should be read from the working path. 4.3.1.2. Protecting State When the protection mechanism has been triggered and the protected domain has performed a protection switch, the domain is in the protecting state. In this state the normal data traffic is transmitted and received on the recovery path. If the protection domain is currently in a protecting state, then the LSRs SHOULD NOT accept a Manual Switch request. If the protection domain is currently in a protecting state, and a Forced Switch is requested then the normal traffic SHALL continue to be transmitted on the recovery path even if the original protection trigger is cleared, and the Forced Switch condition will be signalled by the PSC messages. 4.3.1.3. Unavailable State When the recovery path is unavailable - either as a result of a Lockout operator command (see section 4.3.3), or as a result of a SF or SD detected on the recovery path (see section 4.3.4) - then the protection domain is in the unavailable state. In this state, the normal traffic is transmitted and received on the working path. While in unavailable state any event that would trigger a protection switching SHOULD be ignored with the following exception - If a Signal Degrade request is received, then protection switching will be activated only if the recovery path can guarantee a better signal than the working path. The protection domain will exit the unavailable state and revert to the normal state when, either the operator clears the Lockout command or the recovery path recovers from the signal fault or degraded situation.

No.	'821 Patent Claim 1	IETF MPLS-TP System
		Both ends will resume sending the PCS packets over the recovery path, as a result of this recovery. 4.3.2. Failure or Degraded condition (Working path) If one of the LSRs (for example, LSR-A) detects a failure condition or a serious degradation condition on the working path that warrants invoking protection switching, then it SHOULD take the following actions: • (For 1:1 protection) Switch all traffic for LSR-Z to the recovery path only. • Transmit a PCS control packet, using GACH, with the appropriate Request code (either Signal fault or Signal degrade), the Fpath set to 1, to indicate that the fault/degrade was detected on the working path, and the Path set to 1, indicating that normal traffic is now being transmitted on the recovery path. • Verify that LSR-Z replies with a PCS control packet indicating that it has switched to the recovery path. If this is not received after 2 PSC cycles then send an alarm to the management system. When the far-end LSR (in this example LSR-Z) receives the PCS packet informing it that other LSR (LSR-A) has switched, it SHOULD perform the following actions: • Check priority of the request • Switch all traffic addressed to LSR-A to the recovery path only (for 1:1 protection). • Begin transmission of a PCS control packet, using GACH, with the appropriate Request code (either Signal fault or Signal degrade), the Fpath set to 1, to indicate that the fault/degrade was detected on the working path, and the Path set to 1, indicating that traffic is now being transmitted on the recovery path." MPLS-TP Protection at 13–15. "4.3.4. Failure or Degraded condition (Recovery path) If one of the LSRs (for example, LSR-A) detects a failure condition or a serious degradation condition on the recovery path, then it SHOULD take the following actions:

No.	'821 Patent Claim 1	IETF MPLS-TP System
		• Begin transmission of a PCS control packet with the appropriate Request code (either Signal fault or Signal degrade), the Fpath set to 0, to indicate that the fault/degrade was detected on the recovery path, and the Path set to 0, indicating that traffic is now being forwarded on the working path. Note that this will actually reach the far-end if this is a unidirectional fault or recovery path is possibly in a degraded situation. • The PSC control logic should go into Unavailable state. • All traffic MUST be transmitted on the working path for the duration of the SF/SD condition. When the far-end LSR (in this example LSR-Z) receives the PCS packet informing it that other LSR (LSR-A) has become Unavailable, it SHOULD perform the following actions: • Transmit all traffic on the working path for the duration of the SF/SD condition The PSC Control logic should go into Unavailable state." MPLS-TP Protection at 16. In addition, at least under the apparent claim scope alleged by Orckit's Infringement Disclosures, IETF MPLS System in combination with (1) the knowledge of a person of ordinary skill in the art, alone or in further combination with (2) each (individually, as well as one or more together) of the references identified in element 1[e] of Exhibit E-2 renders the claim, including the present limitation, obvious. Below are examples of two such references. For example, Juniper IOS Guide discloses "Fate Sharing" where a Juniper router computes a primary and secondary path using CSPF, CSPF identifies all fate-sharing groups that are associated with the primary path by identifying all links that the primary path transverses and checks each link in the traffic engineering database against the compiled group list. Juniper IOS Guide further discloses that if the link is a member of a group, the cost of the link is increased by the cost of the group and higher the cost value. At least based upon Orckit's apparent understanding of the claim scope as reflected in its Infringement Disclosures, this limitation is met. Juniper IOS Guide at Page No. 73 ("[Fate Sharing Fate sharing allows you to create a database of information that CSPF uses to compute one or more backup paths to use in case the primary path becomes unstable. The database

No.	'821 Patent Claim 1	IETF MPLS-TP System
		describes the relationships between elements of the network, such as routers and links. You can specify one or more elements within a group. Through fate sharing, you can configure backup paths that minimize the number of shared links and fiber paths with the primary paths as much as possible, to ensure that if a fiber is cut, the minimum amount of data is lost and a path still exists to the destination. For a backup path to work optimally, it must not share links or physical fiber paths with the primary path, ensuring that a single point of failure will not affect the primary and backup paths simultaneously. For more information about fate sharing, see the Junos OS Routing Protocols Configuration Guide.]") Juniper IOS Guide at Page No. 94 ("[Configuring the Ingress Router for MPLS-Signaled LSPs MPLS-signaled label-switched paths (LSPs) run from a specific ingress router to a specific egress router. For basic MPLS-signaled LSP function, you must configure the ingress router, but do not have to configure any other routers. To configure signaled LSPs, perform the following tasks on the ingress router: • Creating Named Paths on page 56 • Configuring Alternate Backup Paths Using Fate Sharing on page 58.]") Juniper IOS Guide at Page No. 96 ("[Configuring Alternate Backup Paths Using Fate Sharing You can create a database of information that Constrained Shortest Path First (CSPF) uses to compute one or more backup paths in case the primary path becomes unstable. The database describes the relationships between elements of the network, such as routers and links. Because these network elements share the same fate, this relationship is called fate sharing. You can configure backup paths that minimize the number of shared links and fiber paths with the primary paths as much as possible to ensure that, if a fiber is cut, the minimum amount of data is lost and a path still exists to the destination. For a backup path to work optimally, it must not share links or physical fiber paths with the primary path. This ensures that a single point of failure will not affect the primary and backup paths at the same time.]") Juniper IOS Guide at Page No. 97 ("[Fate-sharing groups contain three types of objects:

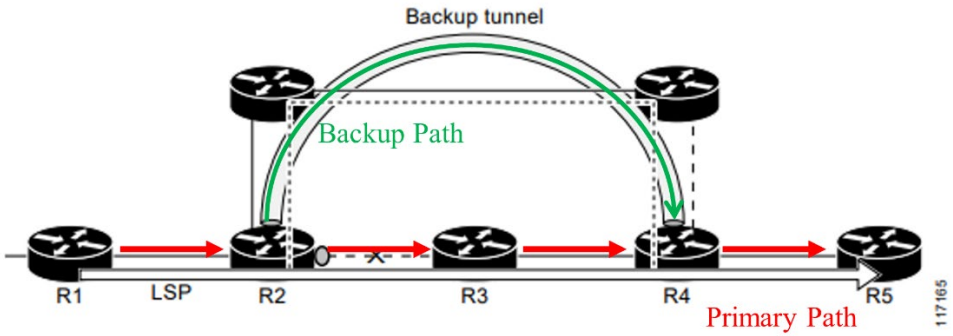
No.	'821 Patent Claim 1	IETF MPLS-TP System
		• Point-to-point links—Identified by the IP addresses at each end of the link. Unnumbered point-to-point links are typically identified by borrowing IP addresses from other interfaces. Order is not important; from 1.2.3.4 to 1.2.3.5 and from 1.2.3.5 to 1.2.3.4 have the same meaning. • Non-point-to-point links—Include links on a LAN interface (such as Gigabit Ethernet interfaces) or nonbroadcast multiaccess (NBMA) interfaces (such as Asynchronous Transfer Mode [ATM] or Frame Relay). You identify these links by their individual interface address. For example, if the LAN interface 192.168.200.0/24 has four routers attached to it, each router link is individually identified: from 192.168.200.1; # LAN interface of router 1 from 192.168.200.2; # LAN interface of router 2 from 192.168.200.3; # LAN interface of router 3 from 192.168.200.4; # LAN interface of router 4 You can list the addresses in any order. • A router node—Identified by its configured router ID. An object can be in any number of groups, and a group can contain any number of objects. Each group has a configurable cost attributed to it, which represents the level of impact this group has on CSPF computations. The higher the cost, the less likely a backup path will share with the primary path any objects in the group. The cost is directly comparable to traffic engineering metrics. By default, the cost is 1. Changing the fate-sharing database does not affect established LSPs until the next reoptimization of CSPF. The fate-sharing database does influence fast-reroute computations.]") Juniper IOS Guide at Page No. 97 ("[When CSPF computes the primary paths of an LSP (or secondary paths when the primary path is not active), it ignores the fate-sharing information. You always want to find the best possible path (least IGP cost) for the primary path. When CSPF computes a secondary path while the primary path (of the same LSP) is active, the following occurs: 1. CSPF identifies all fate-sharing groups that are associated with the primary path. CSPF does this by identifying all links and nodes that the primary path traverses and compiling group lists that contain at least one of the links or nodes. CSPF ignores the ingress and egress nodes in the search.

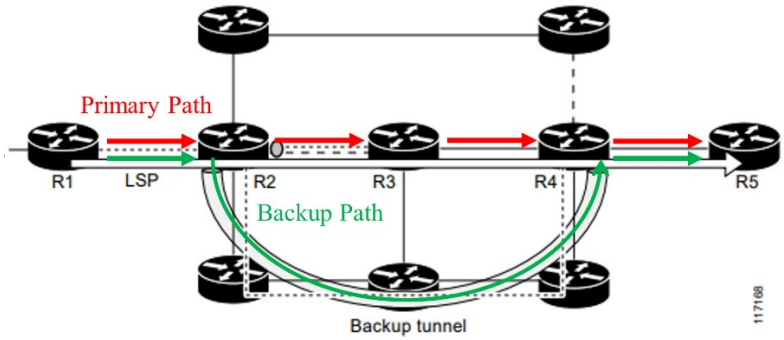
No.	'821 Patent Claim 1	IETF MPLS-TP System
		2. CSPF checks each link in the traffic engineering database against the compiled group list. If the link is a member of a group, the cost of the link is increased by the cost of the group. If a link is a member of multiple groups, all group costs are added together. 3. CSPF performs the check for every node in the traffic engineering database, except the ingress and egress node. Again, a node can belong to multiple groups, so costs are additive. 4. The router performs regular CSPF computation with the adjusted topology.]") Juniper IOS Guide at Page No. 330 ("In Figure 23 on page 292, both node protection and link protection are enabled on Interface B on Router 1. Both node protection and link protection are also enabled on LSP A, an LSP that traverses the link transiting Router 1, Router 2, and Router 3. If the link between Router 1 and Router 2 fails, traffic from LSP A is quickly switched to the next-hop bypass LSP generated by link protection. If Router 2 suffers a hardware or software failure, traffic from LSP A is switched to the next-next-hop bypass LSP generated by node protection.") Figure 23: Node Protection Creating a Next-Next-Hop Bypass LSP  Fig. 23 (annotation added).

No.	'821 Patent Claim 1	IETF MPLS-TP System
		For example, Cisco IOS Guide discloses that the path an LSP uses is determined by the LSP resource requirements and network resources, such as bandwidth. Cisco IOS Guide further discloses the MPLS TE SRLG feature that enhances backup tunnel path selection so a backup tunnel can avoid using links that are in the same SRLG as the interfaces it is protecting. Thus, at least under the apparent claim scope alleged by Orckit's Infringement Disclosures, this limitation is met. Cisco IOS Guide at Page No. 253 ("[Path Option Selection for MPLS TE Tunnel LSPs This section contains the following topics about path option selection for MPLS TE Tunnel LSPs: • Constraint-Based Routing and Path Option Selection, page 5 • Tunnel Reoptimization and Path Option Selection, page 5 • Path Option Selection with Bandwidth Override, page 6 Constraint-Based Routing and Path Option Selection MPLS traffic engineering automatically establishes and maintains LSPs across the backbone by using the Resource Reservation Protocol (RSVP). The path that an LSP uses is determined by the LSP resource requirements and network resources, such as bandwidth. Traffic engineering tunnels are calculated at the LSP head based on a fit between required and available resources (constraint-based routing). Without the Path Option for Bandwidth Override feature, a TE tunnel establishes an LSP based on dynamic or explicit path options in order of preference. However, the bandwidth and other attributes configured on the TE tunnel allow the setup of an LSP only if LSP path options satisfy the constraints. If a path cannot be found that satisfies the configured path options, then the tunnel is not set up.]") Cisco IOS Guide at Page No. 639 ("[Bandwidth Protection NHOP and NNHOP backup tunnels can be used to provide bandwidth protection for rerouted LSPs. This is referred to as backup bandwidth. You can associate backup bandwidth with NHOP or NNHOP backup tunnels. This informs the router of the amount of backup bandwidth a particular backup tunnel can protect. When a router maps LSPs to backup tunnels, bandwidth protection ensures that an LSP uses a given backup tunnel only if there is sufficient backup bandwidth. The router selects which LSPs use which backup tunnels in

No.	'821 Patent Claim 1	IETF MPLS-TP System
		order to provide maximum bandwidth protection. That is, the router determines the best way to map LSPs onto backup tunnels in order to maximize the number of LSPs that can be protected. For information about mapping tunnels and assigning backup bandwidth, see the "Backup Tunnel Selection Procedure" section on page 10. LSPs that have the "bandwidth protection desired" bit set have a higher right to select backup tunnels that provide bandwidth protection; that is, those LSPs can preempt other LSPs that do not have that bit set. For more information, see the "Prioritizing Which LSPs Obtain Backup Tunnels with Bandwidth Protection" section on page 8.]") Cisco IOS Guide at Page No. 640 ("[Backup Tunnel Support Backup tunnel support has the following capabilities: • Backup Tunnels Can Terminate at the Next-Next Hop to Support FRR, page 6 • Multiple Backup Tunnels Can Protect the Same Interface, page 6 • Backup Tunnels Provide Scalability, page 7.]") Cisco IOS Guide at Page 640 ("[Multiple Backup Tunnels Can Protect the Same Interface There is no limit (except memory limitations) to the number of backup tunnels that can protect a given interface. In many topologies, support for node protection requires supporting multiple backup tunnels per protected interface. These backup tunnels can terminate at the same destination or at different destinations. That is, for a given protected interface, you can configure multiple NHOP or NNHOP backup tunnels. This allows redundancy and load balancing.]") Cisco IOS Guide at Page No. 603 ("[Traffic Engineering Tunnels MPLS TE lets you build label switched paths (LSPs) across your network for forwarding traffic. MPLS TE LSPs let the headend of a TE tunnel control the path its traffic takes to a particular destination. This method is more flexible than forwarding traffic based only on a destination address. Interarea tunnels allow you to do the following: • Build TE tunnels between areas (interarea tunnels)

No.	'821 Patent Claim 1	IETF MPLS-TP System
		• Build TE tunnels that start and end in the same area, on multiple areas on a router (intra-area tunnels) Some tunnels are more important than others. For example, you may have tunnels carrying VoIP traffic and tunnels carrying data traffic that are competing for the same resources. MPLS TE allows you to have some tunnels preempt others. Each tunnel has a priority, and more-important tunnels take precedence over less-important tunnels.]") Cisco IOS Guide at Page No. 641 ("[Backup Bandwidth Protection Backup bandwidth protection allows you to give LSPs carrying certain kinds of data (such as voice) priority for using backup tunnels.]") Cisco IOS Guide at Page No. 650 ("[For each protected link AB with a maximum reservable subpool value of n, there may be a path from node A to node B such that the difference between the maximum reservable global and the maximum reservable subpool is at least the value of n. If it is possible to find such paths for each link in the network, you can establish all the backup tunnels along such paths without any bandwidth reservations. If there is a single link failure, only one backup tunnel will use any link on its path. Because that path has at least n available bandwidth (in the global pool), assuming that marking and scheduling is configured to classify the subpool traffic into a priority queue, the subpool bandwidth is guaranteed. This approach allows sharing of the global pool bandwidth between backup tunnels protecting independent link failures. The backup tunnels are expected to be used for only a short period of time after a failure (until the headends of affected LSPs reroute those LSPs to other paths with available subpool bandwidth). The probability of multiple unrelated link failures is very small (in the absence of node or shared risk link group (SRLG) failures, which result in multiple link failures). Therefore, it is reasonable to assume that link failures are in practice independent with high probability. This "independent failure assumption" in combination with backup tunnels signaled without explicit bandwidth reservation enables efficient bandwidth sharing that yields substantial bandwidth savings.]") Cisco IOS Guide at Page No. 651 ("[A similar approach can be used for node and SRLG protection. However, the decision of where to put the backup tunnels is more complicated

No.	'821 Patent Claim 1	IETF MPLS-TP System
		because both node and SRLG failures effectively result in the simultaneous failure of several links. Therefore, the backup tunnels protecting traffic traversing all affected links cannot be computed independently of each other.]") Cisco IOS Guide at Page No. 579 ("[MPLS Traffic Engineering Shared Risk Link Groups SRLGs refer to situations where links in a network share a common fiber (or a common physical attribute). If one link fails, other links in the group may fail too. Links in the group have a shared risk. Backup tunnels should avoid using links in the same SRLG as interfaces they are protecting. Otherwise, when the protected link fails the backup tunnel fails too. ***Figure 1 Backup Tunnel in the Same SRLG as the Interface It Is Protecting***  Fig. 1 (annotation added) Figure 1 shows a primary label-switched path (LSP) from router R1 to router R5. The LSP protects against the failure of the R2-R3 link at R2 via a backup tunnel to R4. If the R2-R3 link fails, link protection reroutes the LSP along the backup tunnel. However, the R2-R3 link and one of the backup tunnel links are in the same SRLG. So if the R2-R3 link fails, the backup tunnel may fail too. The MPLS TE SRLG feature enhances backup tunnel path selection so a backup tunnel can avoid using links that are in the same SRLG as the interfaces it is protecting. There are two ways for a backup tunnel to avoid the SRLGs of its protected interface: • The router does not create the backup tunnel unless it avoids SRLGs of the protected interface.]")

No.	'821 Patent Claim 1	IETF MPLS-TP System
		Cisco IOS Guide at Page No. 580 ("[Open Shortest Path First (OSPF) and Intermediate System-to-Intermediate System (IS-IS) flood the SRLG membership information (including other TE link attributes such as bandwidth availability and affinity) so that all routers in the network have the SRLG information for each link. With this topology information, routers can compute backup tunnel paths that exclude links having SRLGs in common with their protected interfaces. As shown in Figure 2, the backup tunnel avoids the link between R2 and R3, which shares an SRLG with the protected interface.]") *Figure 2 Backup Tunnel That Avoids SRLG of Protected Interface*  Fig. 2 (annotation added) Cisco IOS Guide at Page No. 587 ("[Step 4 show mpls traffic-eng topology Use this command to show the SRLG link membership flooded via the Interior Gateway Protocol (IGP).]") Cisco IOS Guide at Page 632 ("[IS-IS—Intermediate System-to-Intermediate System. Link-state hierarchical routing protocol that calls for intermediate system (IS) routers to exchange routing information based on a single metric to determine network topology.]") Cisco IOS Guide at Page No. 633 ("[OSPF—Open Shortest Path First. A link-state hierarchical Interior Gateway Protocol routing algorithm, derived from the IS-IS protocol. OSPF features include least-cost routing, multipath routing, and load balancing.]")

No.	'821 Patent Claim 1	IETF MPLS-TP System
		Cisco IOS Guide at Page No. 581 ("[Autotunnel Backup for MPLS TE SRLGs Autotunnel backup is the ability of routers to create backup tunnels automatically. Therefore, you do not need to preconfigure each backup tunnel and then assign the backup tunnel to the protected interface. Only automatically created backup tunnels can avoid SRLGs or their protected interfaces.]") Cisco IOS Guide at Page No. 583 ("[Configuring the Routers That Automatically Create Backup Tunnels to Avoid MPLS TE SRLGs of Their Protected Interfaces Perform the following task to configure routers that automatically create backup tunnels to avoid MPLS TE SRLGs of their protected interfaces. Backup tunnels provide link protection by rerouting traffic to the next hop bypassing failed links or in this instance by avoiding SRLGs.]") Cisco IOS Guide at Page No. 644 ("[Backup Tunnel Selection Procedure When an LSP is signaled, each node along the LSP path that provides FRR protection for the LSP selects a backup tunnel for the LSP to use if either of the following events occurs: • The link to the next hop fails. • The next hop fails. By having the node select the backup tunnel for an LSP before a failure occurs, the LSP can be rerouted onto the backup tunnel quickly if there is a failure.]") Cisco IOS Guide at Page 646 ("[NHOP Versus NNHOP Backup Tunnels More than one backup tunnel can protect a given LSP, where one backup tunnel terminates at the LSP's NNHOP, and the other terminates at the LSP's NHOP. In this case, the router chooses the backup tunnel that terminates at the NNHOP (that is, FRR prefers NNHOP over NHOP backup tunnels).Table 1 lists the tunnel selection priorities. The first choice is an NNHOP backup tunnel that acquires its bandwidth from a subpool or global pool, and has limited bandwidth. If there is no such backup tunnel, the next choice (2) is a next-next hop backup tunnel that acquires a limited amount of bandwidth from any pool. The preferences go from 1 (best) to 8 (worst), where choice 3 is for an NNHOP backup tunnel with an unlimited amount of subpool or global-pool bandwidth.]")

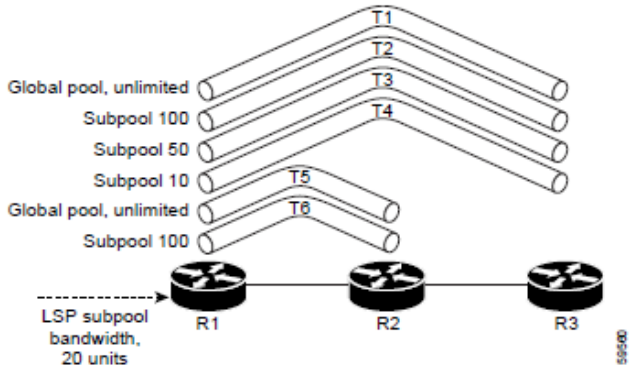
No.	'821 Patent Claim 1	IETF MPLS-TP System

Table 1 Tunnel Selection Priorities

Preference	Backup Tunnel Destination	Bandwidth Pool	Bandwidth Amount
1 (Best)	NNHOP	Subpool or global pool	Limited
2	NNHOP	Any	Limited
3	NNHOP	Subpool or global pool	Unlimited
4	NNHOP	Any	Unlimited
5	NHOP	Subpool or global pool	Limited
6	NHOP	Any	Limited
7	NHOP	Subpool or global pool	Unlimited
8 (Worst)	NHOP	Any	Unlimited

Cisco IOS Guide at Page No. 647-648 ("[Figure 7 shows an example of the backup tunnel selection procedure based on the designated amount of global pool and subpool bandwidth currently available.

Figure 7 Choosing from Among Multiple Backup Tunnels



In this example, an LSP requires 20 units (kilobits per second) of sub-pool backup bandwidth. The best backup tunnel is selected as follows:
1. Backup tunnels T1 through T4 are considered first because they terminate at the NNHOP.
2. Tunnel T4 is eliminated because it has only ten units of sub-pool backup bandwidth.
3. Tunnel T1 is eliminated because it protects only LSPs using global-pool bandwidth.

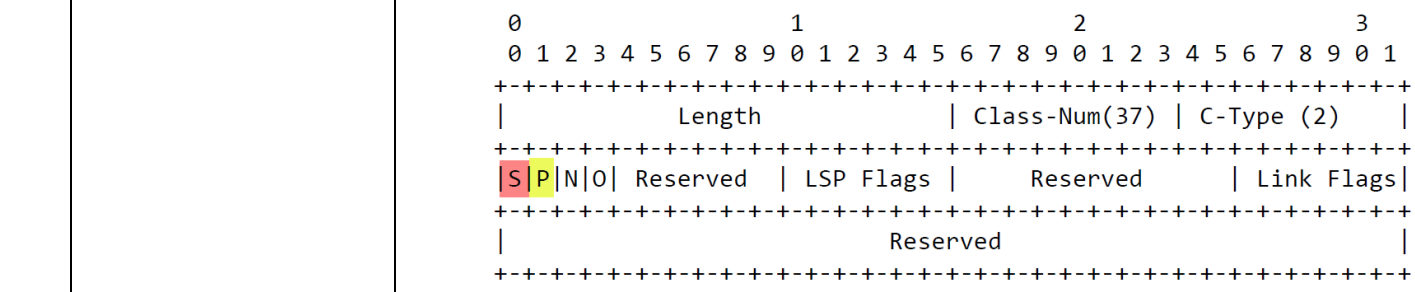
No.	'821 Patent Claim 1	IETF MPLS-TP System
		4. Tunnel T3 is chosen over T2 because, although both have sufficient backup bandwidth, T3 has the least backup bandwidth available (leaving the most backup bandwidth available on T2). 5. Tunnels T5 and T6 need not be considered because they terminate at an NHOP, and therefore are less desirable than T3, which terminates at an NNHOP.]")

No.	'821 Patent Claim 2	IETF MPLS-TP System
2[preamble]	The method of claim 1, wherein said step of selecting an entity pair further comprises:	IETF MPLS-TP System discloses the method of claim 1, wherein said step of selecting an entity pair further comprises. For example, MPLS-TP Protection *See* Claim 1[c] above.
2[a]	selecting a working entity from said entity pair; and	IETF MPLS-TP System discloses selecting a working entity from said entity pair. For example, RFC 4872 discloses that the plurality of MPLS LSPs, i.e., transport entities, are selected in pairs of either working and protection LSPs, or working and restoration LSPs. Specifically, RFC 4872 further discloses the use of a "protection bit" that designates a LSP as a working LSP when it is set to zero.

No.	'821 Patent Claim 2	IETF MPLS-TP System

14.1. Format

The format of the PROTECTION Object (Class-Num = 37, C-Type = 2) is as follows:



Secondary (S): 1 bit

When set to 1, this bit indicates that the requested LSP is a secondary LSP. When set to 0 (default), it indicates that the requested LSP is a primary LSP.

Protecting (P): 1 bit

When set to 1, this bit indicates that the requested LSP is a protecting LSP. When set to 0 (default), it indicates that the requested LSP is a working LSP. The combination, S set to 1 with P set to 0 is not valid.

RFC4872 at 31 (annotation added).

"4.2.1. LSP Status
The following bits are used in determining resource allocation and status of the LSP within the group of LSPs forming the protected entity:
- S (Secondary) bit: enables distinction between primary and secondary LSPs. A primary LSP is a fully established LSP for which the resource allocation has been committed at the data plane (i.e., full cross-connection has been performed). Both working and protecting LSPs can be primary LSPs. A secondary LSP is an LSP that has been provisioned in the control plane only,

No.	'821 Patent Claim 2	IETF MPLS-TP System
		and for which resource selection MAY have been done but for which the resource allocation has not been committed at the data plane (for instance, no cross-connection has been performed). Therefore, a secondary LSP is not immediately available to carry any traffic (thus requiring additional signaling to be available). A secondary LSP can only be a protecting LSP. The (data plane) resources allocated for a secondary LSP MAY be used by other LSPs until the primary LSP fails over to the secondary LSP. - P (Protecting) bit: enables distinction between working and protecting LSPs. A working LSP must be a primary LSP whilst a protecting LSP can be either a primary or a secondary LSP. When protecting LSP(s) are associated with working LSP(s), one also refers to the latter as protected LSPs." RFC 4872 at 7. "D. Requested switch for normal traffic: Recovery signaling is initiated that switches normal traffic to the recovery LSP following the procedures defined in Section 6, 7, 8, and 9. This happens unless a fault condition exists on other LSPs or spans (including the recovery LSP), or a switch command of equal or higher priority is in effect. E. Requested switch for recovery LSP: Recovery signaling is initiated that switches normal traffic to the working LSP following the procedure defined in Section 12. This request is executed except if a fault condition exists on the working LSP or an equal or higher priority switch command is in effect." RFC 4872 at 31. "When a failure is detected by one or both endpoints of the LSP, both endpoints must select traffic from the other LSP. This action must be coordinated between node A and D. From this perspective, 1+1 bidirectional protection can be seen as a coordinated protection- switching mechanism between both endpoints. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 11. "1:1 Protection with Extra-Traffic The most common case of end-to-end 1:N protection is to establish, between the same endpoints, an end-to-end working LSP (thus, N = 1) and a dedicated end-to-end protecting LSP that are mutually link/ node/SRLG disjoint. This protects against working LSP failure(s)." RFC 4872 at 13.

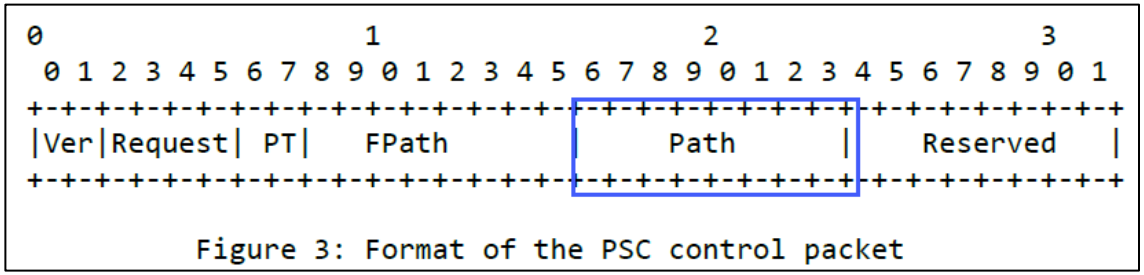
No.	'821 Patent Claim 2	IETF MPLS-TP System
		"An approach to reduce recovery resource requirements is to have protection LSPs sharing network resources when the working LSPs that they protect are physically (i.e., link, node, SRLG, etc.) disjoint. This mechanism is referred to as shared mesh restoration and is described in [RFC4426]. Shared-mesh restoration can be seen as a particular case of pre-planned LSP rerouting (see Section 8) that reduces the recovery resource requirements by allowing multiple protecting LSPs to share common link and node resources. Here also, the recovery resources for the protecting LSPs are pre-reserved during the provisioning phase, thus an explicit signaling action is required to activate (i.e., commit resource allocation at the data plane) a specific protecting LSP instantiated during the (pre-) provisioning phase. This requires restoration signaling along the protecting LSP." RFC 4872 at 20. "1. Introduction Generalized Multi-Protocol Label Switching (GMPLS) extends MPLS to include support for Layer-2 Switch Capable (L2SC), Time-Division Multiplex (TDM), Lambda Switch Capable (LSC), and Fiber Switch Capable (FSC) interfaces. GMPLS recovery uses control plane mechanisms (i.e., signaling, routing, and link management mechanisms) to support data plane fault recovery. Note that the analogous (data plane) fault detection mechanisms are required to be present in support of the control plane mechanisms. In this document, the term "recovery" is generically used to denote both protection and restoration; the specific terms "protection" and "restoration" are only used when differentiation is required. The subtle distinction between protection and restoration is made based on the resource allocation done during the recovery phase (see [RFC4427]). A functional description of GMPLS recovery is provided in [RFC4426] and should be considered as a companion document. The present document describes the protocol-specific procedures for GMPLS RSVP- TE (Resource ReSerVation Protocol - Traffic Engineering) signaling (see [RFC3473]) to support end-to-end recovery. End-to-end recovery refers to the recovery of an entire LSP from its head-end (ingress node endpoint) to its tail-end (egress node endpoint). With end-to- end recovery, working LSPs are assumed to be resource-disjoint (where a resource is a link, node, or Shared Risk Link Group (SRLG)) in the network so that they do not share any failure probability, but this is not mandatory. With respect to a given set of network resources, a pair of working/protecting LSPs SHOULD be resource disjoint in

No.	'821 Patent Claim 2	IETF MPLS-TP System
		case of dedicated recovery type (see below). On the other hand, in case of shared recovery (see below), a group of working LSPs SHOULD be mutually resource-disjoint in order to allow for a (single and commonly) shared protecting LSP, itself resource-disjoint from each of the working LSPs. Note that resource disjointness is a necessary (but not sufficient) condition to ensure LSP recoverability. The present document addresses four types of end-to-end LSP recovery: 1) 1+1 (unidirectional/bidirectional) protection, 2) 1:N (N >= 1) LSP protection with extra-traffic, 3) pre-planned LSP rerouting without extra-traffic (including shared mesh), and 4) full LSP rerouting. 21) The simplest notion of end-to-end LSP protection is 1+1 unidirectional protection. Using this type of protection, a protecting LSP is signaled over a dedicated resource-disjoint alternate path to protect an associated working LSP. Normal traffic is simultaneously sent on both LSPs and a selector is used at the egress node to receive traffic from one of the LSPs. If a failure occurs along one of the LSPs, the egress node selects the traffic from the valid LSP. No coordination is required between the end nodes when a failure/switchover occurs. In 1+1 bidirectional protection, a protecting LSP is signaled over a dedicated resource-disjoint alternate path to protect the working LSP. Normal traffic is simultaneously sent on both LSPs (in both directions), and a selector is used at both ingress/egress nodes to receive traffic from the same LSP. This requires coordination between the end-nodes when switching to the protecting LSP. 22) In 1:N (N >= 1) protection with extra-traffic, the protecting LSP is a fully provisioned and resource-disjoint LSP from the N working LSPs, that allows for carrying extra-traffic. The N working LSPs MAY be mutually resource-disjoint. Coordination between end-nodes is required when switching from one of the working LSPs to the protecting LSP. As the protecting LSP is fully provisioned, default operations during protection switching are specified for a protecting LSP carrying extra-traffic, but this is not mandatory. Note that M:N protection is out of scope of this document (though mechanisms it defines may be extended to cover it). 23) Pre-planned LSP rerouting (or restoration) relies on the establishment between the same pair of end-nodes of a working LSP and a protecting LSP that is link/node/SRLG

No.	'821 Patent Claim 2	IETF MPLS-TP System
		disjoint from the working one. Here, the recovery resources for the protecting LSP are pre-reserved but explicit action is required to activate (i.e., commit resource allocation at the data plane) a specific protecting LSP instantiated during the (pre-)provisioning phase. Since the protecting LSP is not "active" (i.e., fully instantiated), it cannot carry any extra-traffic. This does not mean that the corresponding resources cannot be used by other LSPs. Therefore, this mechanism protects against working LSP(s) failure(s) but requires activation of the protecting LSP after working LSP failure occurrence. This requires restoration signaling along the protecting path. "Shared-mesh" restoration can be seen as a particular case of pre-planned LSP rerouting that reduces the recovery resource requirements by allowing multiple protecting LSPs to share common link and node resources. The recovery resources are pre-reserved but explicit action is required to activate (i.e., commit resource allocation at the data plane) a specific protecting LSP instantiated during the (pre-) provisioning phase. This procedure requires restoration signaling along the protecting path. Note that in both cases, bandwidth pre-reserved for a protecting (but not activated) LSP can be made available for carrying extra traffic. LSPs for extra-traffic (with lower holding priority than the protecting LSP) can then be established using the bandwidth pre-reserved for the protecting LSP. Also, any lower priority LSP that use the pre-reserved resources for the protecting LSP(s) must be preempted during the activation of the protecting LSP. 24) Full LSP rerouting (or restoration) switches normal traffic to an alternate LSP that is not even partially established until after the working LSP failure occurs. The new alternate route is selected at the LSP head-end node, it may reuse resources of the failed LSP at intermediate nodes and may include additional intermediate nodes and/or links." " RFC 4872 at 3–5. "Consider the following network topology: 

No.	'821 Patent Claim 2	IETF MPLS-TP System
		The paths [A,B,C,D] and [A,E,F,G,D] are node and link disjoint, ignoring the ingress/egress nodes A and D. A 1+1 protected path is established from A to D over [A,B,C,D] and [A,E,F,G,D], and traffic is transmitted simultaneously over both component paths (i.e., LSPs). During the provisioning phase, both LSPs are fully instantiated (and thus activated) so that no resource sharing can be done along the protecting LSP (nor can any extra-traffic be transported). It is also RECOMMENDED to set the N bit since no protection-switching signaling is assumed in this case. When a failure occurs (say, at node B) and is detected at end-node D, the receiver at D selects the normal traffic from the other LSP. From this perspective, 1+1 unidirectional protection can be seen as an uncoordinated protection-switching mechanism acting independently at both endpoints. Also, for the LSP under failure condition, it is RECOMMENDED to not set the Path_State_Removed Flag of the ERROR_SPEC object (see [RFC3473]) upon PathErr message generation. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 9. "When a failure is detected by one or both endpoints of the LSP, both endpoints must select traffic from the other LSP. This action must be coordinated between node A and D. From this perspective, 1+1 bidirectional protection can be seen as a coordinated protection- switching mechanism between both endpoints. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 11. "(Full) LSP Rerouting LSP rerouting, on the other hand, switches normal traffic to an alternate LSP that is fully established only after failure occurrence. The new (alternate) route is selected at the LSP head- end and may reuse intermediate nodes included in the original route; it may also include additional intermediate nodes. For strict-hop routing, TE requirements can be directly applied to the route computation, and the failed node or link can be avoided. However, if the failure

No.	'821 Patent Claim 2	IETF MPLS-TP System
		occurred within a loose-routed hop, the head-end node may not have enough information to reroute the LSP around the failure. Crankback signaling (see [CRANK]) and route exclusion techniques (see [RFC4874]) MAY be used in this case. The alternate route MAY be either computed on demand (that is, when the failure occurs; this is referred to as full LSP rerouting) or pre-computed and stored for use when the failure is reported. The latter offers faster restoration time. There is, however, a risk that the alternate route will become out of date through other changes in the network; this can be mitigated to some extent by periodic recalculation of idle alternate routes. (Full) LSP rerouting will be initiated by the head-end node that has either detected the LSP failure or received a Notify message and/or a PathErr message with the new error code/sub-code "Notify Error/LSP Locally Failed" for this LSP. The new LSP resources can be established using the make-before-break mechanism, where the new LSP is set up before the old LSP is torn down. This is done by using the mechanisms of the SESSION_ATTRIBUTE object and the Shared-Explicit (SE) reservation style (see [RFC3209]). Both the new and old LSPs can share resources at common nodes. Note that the make-before-break mechanism is not used to avoid disruption to the normal traffic flow (the latter has already been broken by the failure that is being repaired). However, it is valuable to retain the resources allocated on the original LSP that will be reused by the new alternate LSP." RFC 4872 at 25. "The procedure is as follows: 19) The initiating (source) node sends the normal traffic onto both the working and the protecting LSPs. Once completed, the source node sends reliably a Notify message to the destination with the new error code/sub-code "Notify Error/LSP Recovered" (Switchback Request). This Notify message includes the MESSAGE_ID object. The ACK_Desired flag MUST be set in this object to request the receiver to send an acknowledgment for the message (see [RFC2961]). 20) Upon receipt of this message, the destination selects the traffic from the working LSP. At the same time, it transmits the traffic onto both the working and protecting LSP. The destination then sends reliably a Notify message to the source confirming the completion of the operation. This message includes the MESSAGE_ID_ACK object

No.	'821 Patent Claim 2	IETF MPLS-TP System
		to acknowledge reception of the received Notify message. This Notify message also includes the MESSAGE_ID object. The ACK_Desired flag MUST be set in this object to request the receiver to send an acknowledgment for the message (see [RFC2961]). 21) When the source node receives this Notify message, it switches to receive traffic from the working LSP. The source node then sends an Ack message to the destination node confirming that the LSP has been reverted." RFC 4872 at 27. For example, MPLS-TP Protection discloses an active path indicator field that indicates which LSP and/or pseudo-wire, i.e. network entity, will carry network traffic. MPLS-TP Protection further discloses that the active path indicator can select a working path from the entity pair.  0 1 2 3 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 |Ver|Request| PT| FPath | Path | Reserved | Figure 3: Format of the PSC control packet MPLS-TP Figure 3 (annotation added). "3. Protection switching logic 3.1. Protection switching trigger mechanisms The protection switching should be initiated in reaction to any of the following triggers: • Server layer indication - if the MPLS-TP server layer detects a failure within its own layer, or due to a failure of its server layer (e.g. the physical layer) notifies the MPLS-TP layer that a failure has been detected. It should be noted that this trigger will be

No.	'821 Patent Claim 2	IETF MPLS-TP System
		controlled by a holdoff-timer that should be configured to allow the server layer to apply any protection procedures that may be relevant. • OAM signalling - if, for example, OAM continuity and connectivity verification tools detect that there is a loss of continuity or mis-connectivity or performance monitoring indicates a degradation of the utility of the working path for the current transport path. In cases of signal degradation, switching to the recovery path SHOULD only be activated if the recovery path can guarantee better conditions than the degraded working path. • Control plane - if there is a control plane active in the network (either signaling or routing), it MAY trigger protection switching based on conditions detected by the control plane. If the control-plane is based on GMPLS [RFC3945] then the recovery process should comply with the process described in [RFC4872]. • Operator command - the network operator may issue commands that trigger protection switching. The commands that are supported include - Forced Switch, Manual Switch, Clear, Lockout of Protection, (see definitions in [RFC4427]).” MPLS-TP at 6–7. “3.2. Protection switching control logical architecture Protection switching processes the triggers described above together with the inputs received from the far-end LSR. These inputs cause the LSR to take certain actions, e.g. switching the Selector Bridge to select the working or recovery path, and to transmit different protocol messages.

No.	'821 Patent Claim 2	IETF MPLS-TP System

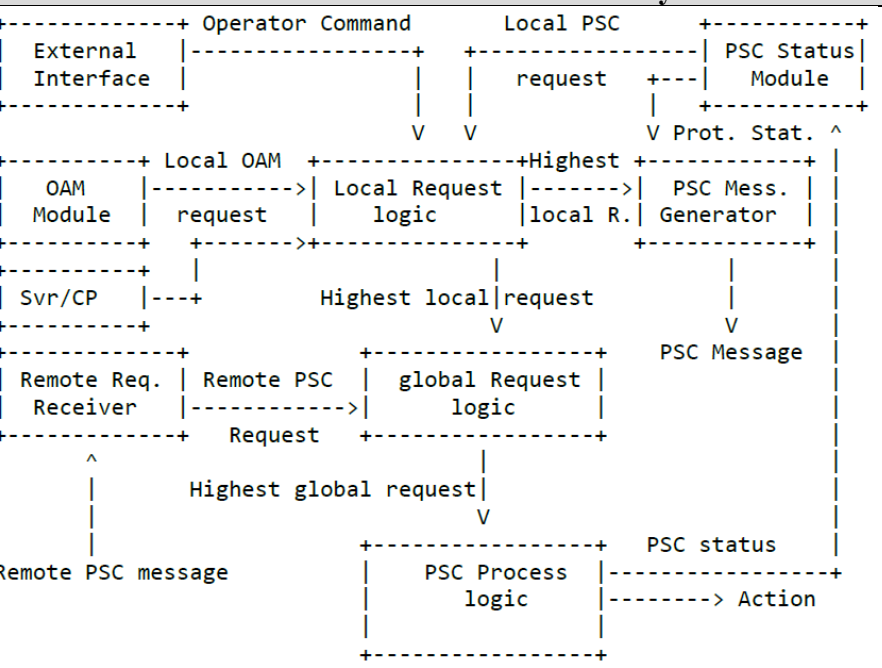
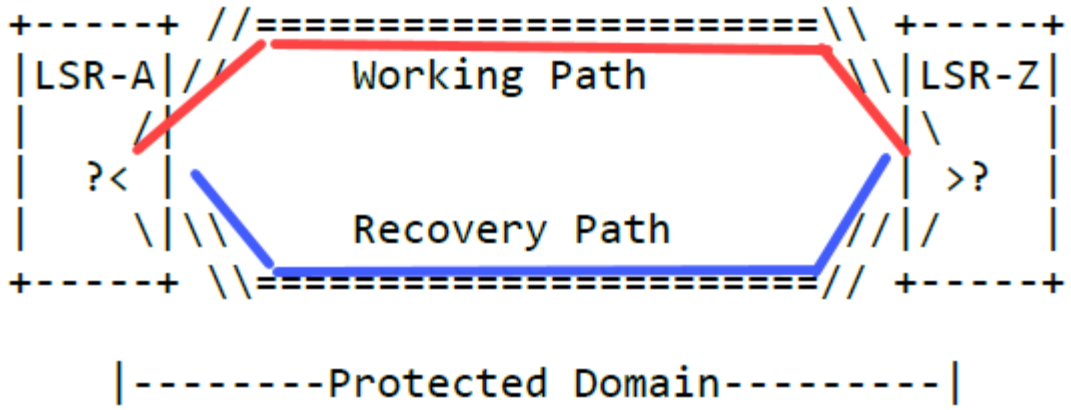


Figure 1: Protection switching control logic

Figure 1 describes the logical architecture of the protection switching control. The Local Request logic unit accepts the triggers from the OAM, external operator commands, and from the local control plane (when present) and determines the highest priority request. This high-priority request is passed to both the PSC Message generator, that will generate the appropriate protocol message to be sent to the far-end LSR, and the Global Request logic, that will cross-check this local request with the information received from the far-LSR. The Global Request logic then processes these two PSC requests that determines the highest priority request that is passed to the PSC Process logic. The PSC Process logic uses this input to determine what actions need to be taken, e.g. switching the Selector Bridge, and the current status of the protection domain." MPLS-TP Protection at 7–8.

"4.2.4. Active path indicator (Path)

No.	'821 Patent Claim 2	IETF MPLS-TP System
		The Path field of the PSC control SHALL be used to indicate which path the source MEP is currently using for data transmission. The MEP should compare the value of this bit with the path that is locally selected for data transmission to verify that there is no inconsistency between the two end-points of the protected domain. If an inconsistency is detected then an alarm should be raised. The following are the possible values: • 0: indicates that normal traffic is being transmitted on the Working path. • 1: indicates the Recovery path is being used to transmit the normal traffic from the Working path. 2-255: for future extensions" MPLS-TP Protection at 12–13. "Principles of Operation In all of the following sub-sections, assume a protected domain between LSR-A and LSR-Z, using paths W (working) and R (recovery) as shown in figure 4." MPLS-TP Protection at 13.  MPLS-TP Protection Figure 4 (annotation added). "4.3.1 PSC States 4.3.1.1 Normal State

No.	'821 Patent Claim 2	IETF MPLS-TP System
		When the protected domain has no special condition in effect, the ingress LSR SHOULD forward the user data along the working path, and, in the case of 1+1 protection, the Permanent Bridge will bridge the data to the recovery path as well. The receiving LSR SHOULD read the data from the working path. The ingress LSR MAY transmit a No Request PSC packet with the Path field set to 0 indicating that the normal data traffic should be read from the working path. 4.3.1.2. Protecting State When the protection mechanism has been triggered and the protected domain has performed a protection switch, the domain is in the protecting state. In this state the normal data traffic is transmitted and received on the recovery path. If the protection domain is currently in a protecting state, then the LSRs SHOULD NOT accept a Manual Switch request. If the protection domain is currently in a protecting state, and a Forced Switch is requested then the normal traffic SHALL continue to be transmitted on the recovery path even if the original protection trigger is cleared, and the Forced Switch condition will be signalled by the PSC messages. 4.3.1.3. Unavailable State When the recovery path is unavailable - either as a result of a Lockout operator command (see section 4.3.3), or as a result of a SF or SD detected on the recovery path (see section 4.3.4) - then the protection domain is in the unavailable state. In this state, the normal traffic is transmitted and received on the working path. While in unavailable state any event that would trigger a protection switching SHOULD be ignored with the following exception - If a Signal Degrade request is received, then protection switching will be activated only if the recovery path can guarantee a better signal than the working path. The protection domain will exit the unavailable state and revert to the normal state when, either the operator clears the Lockout command or the recovery path recovers from the signal fault or degraded situation.

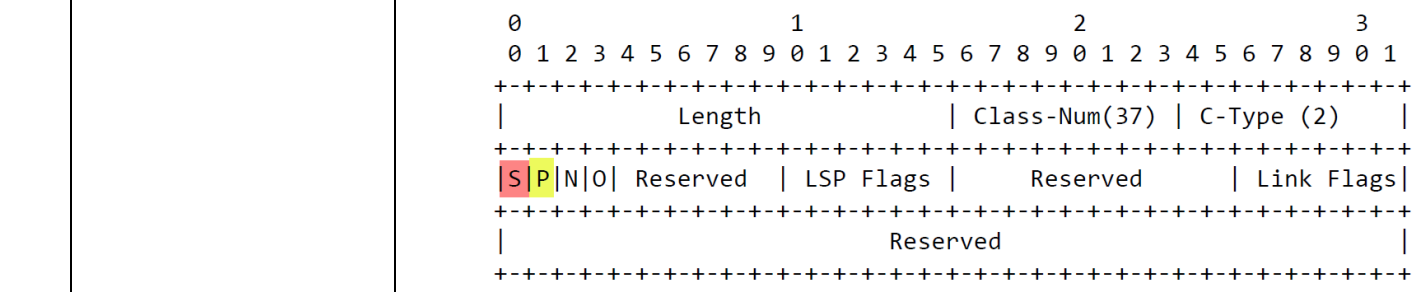
No.	'821 Patent Claim 2	IETF MPLS-TP System
		Both ends will resume sending the PCS packets over the recovery path, as a result of this recovery. 4.3.2. Failure or Degraded condition (Working path) If one of the LSRs (for example, LSR-A) detects a failure condition or a serious degradation condition on the working path that warrants invoking protection switching, then it SHOULD take the following actions: • (For 1:1 protection) Switch all traffic for LSR-Z to the recovery path only. • Transmit a PCS control packet, using GACH, with the appropriate Request code (either Signal fault or Signal degrade), the Fpath set to 1, to indicate that the fault/degrade was detected on the working path, and the Path set to 1, indicating that normal traffic is now being transmitted on the recovery path. • Verify that LSR-Z replies with a PCS control packet indicating that it has switched to the recovery path. If this is not received after 2 PSC cycles then send an alarm to the management system. When the far-end LSR (in this example LSR-Z) receives the PCS packet informing it that other LSR (LSR-A) has switched, it SHOULD perform the following actions: • Check priority of the request • Switch all traffic addressed to LSR-A to the recovery path only (for 1:1 protection). • Begin transmission of a PCS control packet, using GACH, with the appropriate Request code (either Signal fault or Signal degrade), the Fpath set to 1, to indicate that the fault/degrade was detected on the working path, and the Path set to 1, indicating that traffic is now being transmitted on the recovery path." MPLS-TP Protection at 13–15. "4.3.4. Failure or Degraded condition (Recovery path) If one of the LSRs (for example, LSR-A) detects a failure condition or a serious degradation condition on the recovery path, then it SHOULD take the following actions:

No.	'821 Patent Claim 2	IETF MPLS-TP System
		• Begin transmission of a PCS control packet with the appropriate Request code (either Signal fault or Signal degrade), the Fpath set to 0, to indicate that the fault/degrade was detected on the recovery path, and the Path set to 0, indicating that traffic is now being forwarded on the working path. Note that this will actually reach the far-end if this is a unidirectional fault or recovery path is possibly in a degraded situation. • The PSC control logic should go into Unavailable state. • All traffic MUST be transmitted on the working path for the duration of the SF/SD condition. When the far-end LSR (in this example LSR-Z) receives the PCS packet informing it that other LSR (LSR-A) has become Unavailable, it SHOULD perform the following actions: • Transmit all traffic on the working path for the duration of the SF/SD condition The PSC Control logic should go into Unavailable state." MPLS-TP Protection at 16.
2[b]	selecting a protection entity from said entity pair.	IETF MPLS-TP System discloses selecting a protection entity from said entity pair. For example, RFC 4872 discloses that the plurality of MPLS LSPs, i.e., transport entities, are selected in pairs of either working and protection LSPs, or working and restoration LSPs. Specifically, RFC 4872 further discloses the use of a "protection bit" that designates a LSP as a protection and/or restoration LSP when it is set to one.

No.	'821 Patent Claim 2	IETF MPLS-TP System

14.1. Format

The format of the PROTECTION Object (Class-Num = 37, C-Type = 2) is as follows:



Secondary (S): 1 bit

When set to 1, this bit indicates that the requested LSP is a secondary LSP. When set to 0 (default), it indicates that the requested LSP is a primary LSP.

Protecting (P): 1 bit

When set to 1, this bit indicates that the requested LSP is a protecting LSP. When set to 0 (default), it indicates that the requested LSP is a working LSP. The combination, S set to 1 with P set to 0 is not valid.

RFC4872 at 31 (annotation added).

"4.2.1. LSP Status
The following bits are used in determining resource allocation and status of the LSP within the group of LSPs forming the protected entity:
- S (Secondary) bit: enables distinction between primary and secondary LSPs. A primary LSP is a fully established LSP for which the resource allocation has been committed at the data plane (i.e., full cross-connection has been performed). Both working and protecting LSPs can be primary LSPs. A secondary LSP is an LSP that has been provisioned in the control plane only,

No.	'821 Patent Claim 2	IETF MPLS-TP System
		and for which resource selection MAY have been done but for which the resource allocation has not been committed at the data plane (for instance, no cross-connection has been performed). Therefore, a secondary LSP is not immediately available to carry any traffic (thus requiring additional signaling to be available). A secondary LSP can only be a protecting LSP. The (data plane) resources allocated for a secondary LSP MAY be used by other LSPs until the primary LSP fails over to the secondary LSP. - P (Protecting) bit: enables distinction between working and protecting LSPs. A working LSP must be a primary LSP whilst a protecting LSP can be either a primary or a secondary LSP. When protecting LSP(s) are associated with working LSP(s), one also refers to the latter as protected LSPs." RFC 4872 at 7. "D. Requested switch for normal traffic: Recovery signaling is initiated that switches normal traffic to the recovery LSP following the procedures defined in Section 6, 7, 8, and 9. This happens unless a fault condition exists on other LSPs or spans (including the recovery LSP), or a switch command of equal or higher priority is in effect. E. Requested switch for recovery LSP: Recovery signaling is initiated that switches normal traffic to the working LSP following the procedure defined in Section 12. This request is executed except if a fault condition exists on the working LSP or an equal or higher priority switch command is in effect." RFC 4872 at 31. "When a failure is detected by one or both endpoints of the LSP, both endpoints must select traffic from the other LSP. This action must be coordinated between node A and D. From this perspective, 1+1 bidirectional protection can be seen as a coordinated protection- switching mechanism between both endpoints. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 11. "1:1 Protection with Extra-Traffic The most common case of end-to-end 1:N protection is to establish, between the same endpoints, an end-to-end working LSP (thus, N = 1) and a dedicated end-to-end protecting LSP that are mutually link/ node/SRLG disjoint. This protects against working LSP failure(s)." RFC 4872 at 13.

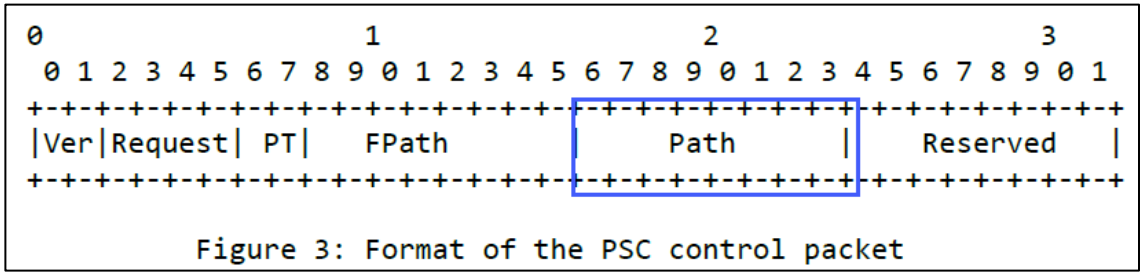
No.	'821 Patent Claim 2	IETF MPLS-TP System
		"An approach to reduce recovery resource requirements is to have protection LSPs sharing network resources when the working LSPs that they protect are physically (i.e., link, node, SRLG, etc.) disjoint. This mechanism is referred to as shared mesh restoration and is described in [RFC4426]. Shared-mesh restoration can be seen as a particular case of pre-planned LSP rerouting (see Section 8) that reduces the recovery resource requirements by allowing multiple protecting LSPs to share common link and node resources. Here also, the recovery resources for the protecting LSPs are pre-reserved during the provisioning phase, thus an explicit signaling action is required to activate (i.e., commit resource allocation at the data plane) a specific protecting LSP instantiated during the (pre-) provisioning phase. This requires restoration signaling along the protecting LSP." RFC 4872 at 20. "1. Introduction Generalized Multi-Protocol Label Switching (GMPLS) extends MPLS to include support for Layer-2 Switch Capable (L2SC), Time-Division Multiplex (TDM), Lambda Switch Capable (LSC), and Fiber Switch Capable (FSC) interfaces. GMPLS recovery uses control plane mechanisms (i.e., signaling, routing, and link management mechanisms) to support data plane fault recovery. Note that the analogous (data plane) fault detection mechanisms are required to be present in support of the control plane mechanisms. In this document, the term "recovery" is generically used to denote both protection and restoration; the specific terms "protection" and "restoration" are only used when differentiation is required. The subtle distinction between protection and restoration is made based on the resource allocation done during the recovery phase (see [RFC4427]). A functional description of GMPLS recovery is provided in [RFC4426] and should be considered as a companion document. The present document describes the protocol-specific procedures for GMPLS RSVP- TE (Resource ReSerVation Protocol - Traffic Engineering) signaling (see [RFC3473]) to support end-to-end recovery. End-to-end recovery refers to the recovery of an entire LSP from its head-end (ingress node endpoint) to its tail-end (egress node endpoint). With end-to- end recovery, working LSPs are assumed to be resource-disjoint (where a resource is a link, node, or Shared Risk Link Group (SRLG)) in the network so that they do not share any failure probability, but this is not mandatory. With respect to a given set of network resources, a pair of working/protecting LSPs SHOULD be resource disjoint in

No.	'821 Patent Claim 2	IETF MPLS-TP System
		case of dedicated recovery type (see below). On the other hand, in case of shared recovery (see below), a group of working LSPs SHOULD be mutually resource-disjoint in order to allow for a (single and commonly) shared protecting LSP, itself resource-disjoint from each of the working LSPs. Note that resource disjointness is a necessary (but not sufficient) condition to ensure LSP recoverability. The present document addresses four types of end-to-end LSP recovery: 1) 1+1 (unidirectional/bidirectional) protection, 2) 1:N (N >= 1) LSP protection with extra-traffic, 3) pre-planned LSP rerouting without extra-traffic (including shared mesh), and 4) full LSP rerouting. 25) The simplest notion of end-to-end LSP protection is 1+1 unidirectional protection. Using this type of protection, a protecting LSP is signaled over a dedicated resource-disjoint alternate path to protect an associated working LSP. Normal traffic is simultaneously sent on both LSPs and a selector is used at the egress node to receive traffic from one of the LSPs. If a failure occurs along one of the LSPs, the egress node selects the traffic from the valid LSP. No coordination is required between the end nodes when a failure/switchover occurs. In 1+1 bidirectional protection, a protecting LSP is signaled over a dedicated resource-disjoint alternate path to protect the working LSP. Normal traffic is simultaneously sent on both LSPs (in both directions), and a selector is used at both ingress/egress nodes to receive traffic from the same LSP. This requires coordination between the end-nodes when switching to the protecting LSP. 26) In 1:N (N >= 1) protection with extra-traffic, the protecting LSP is a fully provisioned and resource-disjoint LSP from the N working LSPs, that allows for carrying extra-traffic. The N working LSPs MAY be mutually resource-disjoint. Coordination between end-nodes is required when switching from one of the working LSPs to the protecting LSP. As the protecting LSP is fully provisioned, default operations during protection switching are specified for a protecting LSP carrying extra-traffic, but this is not mandatory. Note that M:N protection is out of scope of this document (though mechanisms it defines may be extended to cover it). 27) Pre-planned LSP rerouting (or restoration) relies on the establishment between the same pair of end-nodes of a working LSP and a protecting LSP that is link/node/SRLG

No.	'821 Patent Claim 2	IETF MPLS-TP System
		disjoint from the working one. Here, the recovery resources for the protecting LSP are pre-reserved but explicit action is required to activate (i.e., commit resource allocation at the data plane) a specific protecting LSP instantiated during the (pre-)provisioning phase. Since the protecting LSP is not "active" (i.e., fully instantiated), it cannot carry any extra-traffic. This does not mean that the corresponding resources cannot be used by other LSPs. Therefore, this mechanism protects against working LSP(s) failure(s) but requires activation of the protecting LSP after working LSP failure occurrence. This requires restoration signaling along the protecting path. "Shared-mesh" restoration can be seen as a particular case of pre-planned LSP rerouting that reduces the recovery resource requirements by allowing multiple protecting LSPs to share common link and node resources. The recovery resources are pre-reserved but explicit action is required to activate (i.e., commit resource allocation at the data plane) a specific protecting LSP instantiated during the (pre-) provisioning phase. This procedure requires restoration signaling along the protecting path. Note that in both cases, bandwidth pre-reserved for a protecting (but not activated) LSP can be made available for carrying extra traffic. LSPs for extra-traffic (with lower holding priority than the protecting LSP) can then be established using the bandwidth pre-reserved for the protecting LSP. Also, any lower priority LSP that use the pre-reserved resources for the protecting LSP(s) must be preempted during the activation of the protecting LSP. 28) Full LSP rerouting (or restoration) switches normal traffic to an alternate LSP that is not even partially established until after the working LSP failure occurs. The new alternate route is selected at the LSP head-end node, it may reuse resources of the failed LSP at intermediate nodes and may include additional intermediate nodes and/or links." " RFC 4872 at 3–5. "Consider the following network topology: 

No.	'821 Patent Claim 2	IETF MPLS-TP System
		The paths [A,B,C,D] and [A,E,F,G,D] are node and link disjoint, ignoring the ingress/egress nodes A and D. A 1+1 protected path is established from A to D over [A,B,C,D] and [A,E,F,G,D], and traffic is transmitted simultaneously over both component paths (i.e., LSPs). During the provisioning phase, both LSPs are fully instantiated (and thus activated) so that no resource sharing can be done along the protecting LSP (nor can any extra-traffic be transported). It is also RECOMMENDED to set the N bit since no protection-switching signaling is assumed in this case. When a failure occurs (say, at node B) and is detected at end-node D, the receiver at D selects the normal traffic from the other LSP. From this perspective, 1+1 unidirectional protection can be seen as an uncoordinated protection-switching mechanism acting independently at both endpoints. Also, for the LSP under failure condition, it is RECOMMENDED to not set the Path_State_Removed Flag of the ERROR_SPEC object (see [RFC3473]) upon PathErr message generation. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 9. "When a failure is detected by one or both endpoints of the LSP, both endpoints must select traffic from the other LSP. This action must be coordinated between node A and D. From this perspective, 1+1 bidirectional protection can be seen as a coordinated protection- switching mechanism between both endpoints. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 11. "(Full) LSP Rerouting LSP rerouting, on the other hand, switches normal traffic to an alternate LSP that is fully established only after failure occurrence. The new (alternate) route is selected at the LSP head- end and may reuse intermediate nodes included in the original route; it may also include additional intermediate nodes. For strict-hop routing, TE requirements can be directly applied to the route computation, and the failed node or link can be avoided. However, if the failure

No.	'821 Patent Claim 2	IETF MPLS-TP System
		occurred within a loose-routed hop, the head-end node may not have enough information to reroute the LSP around the failure. Crankback signaling (see [CRANK]) and route exclusion techniques (see [RFC4874]) MAY be used in this case. The alternate route MAY be either computed on demand (that is, when the failure occurs; this is referred to as full LSP rerouting) or pre-computed and stored for use when the failure is reported. The latter offers faster restoration time. There is, however, a risk that the alternate route will become out of date through other changes in the network; this can be mitigated to some extent by periodic recalculation of idle alternate routes. (Full) LSP rerouting will be initiated by the head-end node that has either detected the LSP failure or received a Notify message and/or a PathErr message with the new error code/sub-code "Notify Error/LSP Locally Failed" for this LSP. The new LSP resources can be established using the make-before-break mechanism, where the new LSP is set up before the old LSP is torn down. This is done by using the mechanisms of the SESSION_ATTRIBUTE object and the Shared-Explicit (SE) reservation style (see [RFC3209]). Both the new and old LSPs can share resources at common nodes. Note that the make-before-break mechanism is not used to avoid disruption to the normal traffic flow (the latter has already been broken by the failure that is being repaired). However, it is valuable to retain the resources allocated on the original LSP that will be reused by the new alternate LSP." RFC 4872 at 25. "The procedure is as follows: 22) The initiating (source) node sends the normal traffic onto both the working and the protecting LSPs. Once completed, the source node sends reliably a Notify message to the destination with the new error code/sub-code "Notify Error/LSP Recovered" (Switchback Request). This Notify message includes the MESSAGE_ID object. The ACK_Desired flag MUST be set in this object to request the receiver to send an acknowledgment for the message (see [RFC2961]). 23) Upon receipt of this message, the destination selects the traffic from the working LSP. At the same time, it transmits the traffic onto both the working and protecting LSP. The destination then sends reliably a Notify message to the source confirming the completion of the operation. This message includes the MESSAGE_ID_ACK object

No.	'821 Patent Claim 2	IETF MPLS-TP System
		to acknowledge reception of the received Notify message. This Notify message also includes the MESSAGE_ID object. The ACK_Desired flag MUST be set in this object to request the receiver to send an acknowledgment for the message (see [RFC2961]). 24) When the source node receives this Notify message, it switches to receive traffic from the working LSP. The source node then sends an Ack message to the destination node confirming that the LSP has been reverted." RFC 4872 at 27. For example, MPLS-TP Protection discloses an active path indicator field that indicates which LSP and/or pseudo-wire, i.e. network entity, will carry network traffic. MPLS-TP Protection further discloses that the active path indicator can select a recovery path from the entity pair.  Figure 3: Format of the PSC control packet MPLS-TP Figure 3 (annotation added). "3. Protection switching logic 3.1. Protection switching trigger mechanisms The protection switching should be initiated in reaction to any of the following triggers: • Server layer indication - if the MPLS-TP server layer detects a failure within its own layer, or due to a failure of its server layer (e.g. the physical layer) notifies the MPLS-TP layer that a failure has been detected. It should be noted that this trigger will be

No.	'821 Patent Claim 2	IETF MPLS-TP System
		controlled by a holdoff-timer that should be configured to allow the server layer to apply any protection procedures that may be relevant. • OAM signalling - if, for example, OAM continuity and connectivity verification tools detect that there is a loss of continuity or mis-connectivity or performance monitoring indicates a degradation of the utility of the working path for the current transport path. In cases of signal degradation, switching to the recovery path SHOULD only be activated if the recovery path can guarantee better conditions than the degraded working path. • Control plane - if there is a control plane active in the network (either signaling or routing), it MAY trigger protection switching based on conditions detected by the control plane. If the control-plane is based on GMPLS [RFC3945] then the recovery process should comply with the process described in [RFC4872]. • Operator command - the network operator may issue commands that trigger protection switching. The commands that are supported include - Forced Switch, Manual Switch, Clear, Lockout of Protection, (see definitions in [RFC4427]).” MPLS-TP at 6–7. “3.2. Protection switching control logical architecture Protection switching processes the triggers described above together with the inputs received from the far-end LSR. These inputs cause the LSR to take certain actions, e.g. switching the Selector Bridge to select the working or recovery path, and to transmit different protocol messages.

No.	'821 Patent Claim 2	IETF MPLS-TP System

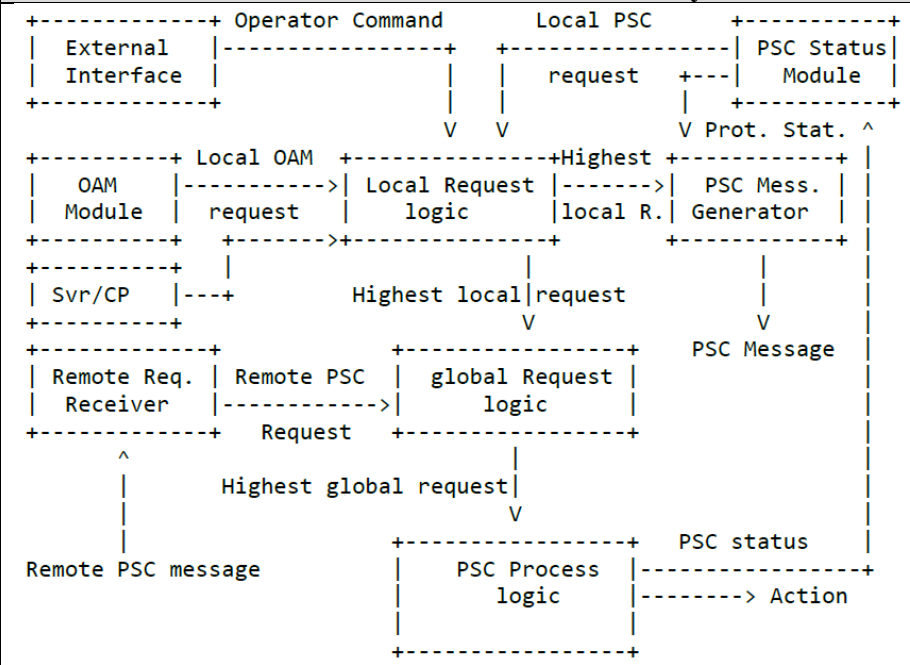
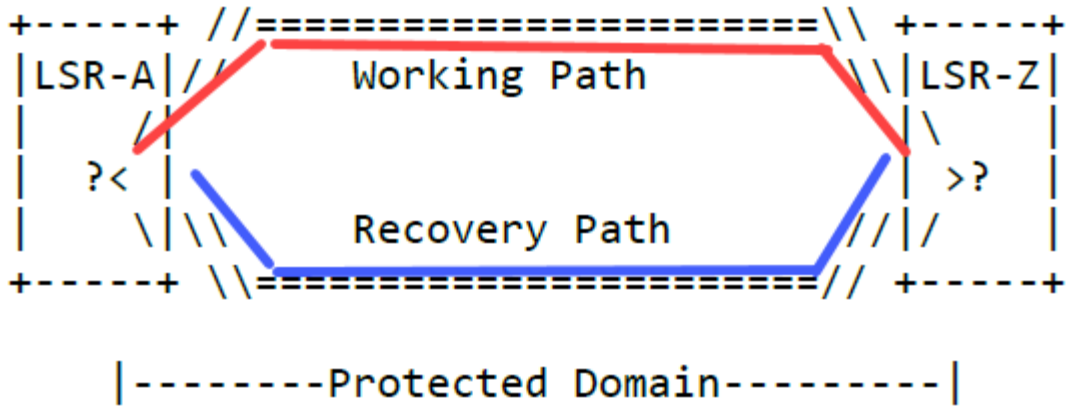


Figure 1: Protection switching control logic

Figure 1 describes the logical architecture of the protection switching control. The Local Request logic unit accepts the triggers from the OAM, external operator commands, and from the local control plane (when present) and determines the highest priority request. This high-priority request is passed to both the PSC Message generator, that will generate the appropriate protocol message to be sent to the far-end LSR, and the Global Request logic, that will cross-check this local request with the information received from the far-LSR. The Global Request logic then processes these two PSC requests that determines the highest priority request that is passed to the PSC Process logic. The PSC Process logic uses this input to determine what actions need to be taken, e.g. switching the Selector Bridge, and the current status of the protection domain." MPLS-TP Protection at 7–8.

"4.2.4. Active path indicator (Path)

No.	'821 Patent Claim 2	IETF MPLS-TP System
		The Path field of the PSC control SHALL be used to indicate which path the source MEP is currently using for data transmission. The MEP should compare the value of this bit with the path that is locally selected for data transmission to verify that there is no inconsistency between the two end-points of the protected domain. If an inconsistency is detected then an alarm should be raised. The following are the possible values: • 0: indicates that normal traffic is being transmitted on the Working path. • 1: indicates the Recovery path is being used to transmit the normal traffic from the Working path. 2-255: for future extensions" MPLS-TP Protection at 12–13. "Principles of Operation In all of the following sub-sections, assume a protected domain between LSR-A and LSR-Z, using paths W (working) and R (recovery) as shown in figure 4." MPLS-TP Protection at 13.  MPLS-TP Protection Figure 4 (annotation added). "4.3.1 PSC States 4.3.1.1 Normal State

No.	'821 Patent Claim 2	IETF MPLS-TP System
		When the protected domain has no special condition in effect, the ingress LSR SHOULD forward the user data along the working path, and, in the case of 1+1 protection, the Permanent Bridge will bridge the data to the recovery path as well. The receiving LSR SHOULD read the data from the working path. The ingress LSR MAY transmit a No Request PSC packet with the Path field set to 0 indicating that the normal data traffic should be read from the working path. 4.3.1.2. Protecting State When the protection mechanism has been triggered and the protected domain has performed a protection switch, the domain is in the protecting state. In this state the normal data traffic is transmitted and received on the recovery path. If the protection domain is currently in a protecting state, then the LSRs SHOULD NOT accept a Manual Switch request. If the protection domain is currently in a protecting state, and a Forced Switch is requested then the normal traffic SHALL continue to be transmitted on the recovery path even if the original protection trigger is cleared, and the Forced Switch condition will be signalled by the PSC messages. 4.3.1.3. Unavailable State When the recovery path is unavailable - either as a result of a Lockout operator command (see section 4.3.3), or as a result of a SF or SD detected on the recovery path (see section 4.3.4) - then the protection domain is in the unavailable state. In this state, the normal traffic is transmitted and received on the working path. While in unavailable state any event that would trigger a protection switching SHOULD be ignored with the following exception - If a Signal Degrade request is received, then protection switching will be activated only if the recovery path can guarantee a better signal than the working path. The protection domain will exit the unavailable state and revert to the normal state when, either the operator clears the Lockout command or the recovery path recovers from the signal fault or degraded situation.

No.	'821 Patent Claim 2	IETF MPLS-TP System
		Both ends will resume sending the PCS packets over the recovery path, as a result of this recovery. 4.3.2. Failure or Degraded condition (Working path) If one of the LSRs (for example, LSR-A) detects a failure condition or a serious degradation condition on the working path that warrants invoking protection switching, then it SHOULD take the following actions: • (For 1:1 protection) Switch all traffic for LSR-Z to the recovery path only. • Transmit a PCS control packet, using GACH, with the appropriate Request code (either Signal fault or Signal degrade), the Fpath set to 1, to indicate that the fault/degrade was detected on the working path, and the Path set to 1, indicating that normal traffic is now being transmitted on the recovery path. • Verify that LSR-Z replies with a PCS control packet indicating that it has switched to the recovery path. If this is not received after 2 PSC cycles then send an alarm to the management system. When the far-end LSR (in this example LSR-Z) receives the PCS packet informing it that other LSR (LSR-A) has switched, it SHOULD perform the following actions: • Check priority of the request • Switch all traffic addressed to LSR-A to the recovery path only (for 1:1 protection). • Begin transmission of a PCS control packet, using GACH, with the appropriate Request code (either Signal fault or Signal degrade), the Fpath set to 1, to indicate that the fault/degrade was detected on the working path, and the Path set to 1, indicating that traffic is now being transmitted on the recovery path." MPLS-TP Protection at 13–15. "4.3.4. Failure or Degraded condition (Recovery path) If one of the LSRs (for example, LSR-A) detects a failure condition or a serious degradation condition on the recovery path, then it SHOULD take the following actions:

No.	'821 Patent Claim 2	IETF MPLS-TP System
		• Begin transmission of a PCS control packet with the appropriate Request code (either Signal fault or Signal degrade), the Fpath set to 0, to indicate that the fault/degrade was detected on the recovery path, and the Path set to 0, indicating that traffic is now being forwarded on the working path. Note that this will actually reach the far-end if this is a unidirectional fault or recovery path is possibly in a degraded situation. • The PSC control logic should go into Unavailable state. • All traffic MUST be transmitted on the working path for the duration of the SF/SD condition. When the far-end LSR (in this example LSR-Z) receives the PCS packet informing it that other LSR (LSR-A) has become Unavailable, it SHOULD perform the following actions: • Transmit all traffic on the working path for the duration of the SF/SD condition The PSC Control logic should go into Unavailable state." MPLS-TP Protection at 16.

No.	'821 Patent Claim 3	IETF MPLS-TP System
3	The method of claim 2, further comprising the step of selecting an active entity from the set consisting of said working entity and said protection entity.	IETF MPLS-TP System discloses the method of claim 2, further comprising the step of selecting an active entity from the set consisting of said working entity and said protection entity. For example, RFC 4872 discloses that the plurality of MPLS LSPs, i.e., transport entities, are selected in pairs of either working and protection LSPs, or working and restoration LSPs. RFC 4872 further discloses the recovery signaling feature that commands traffic to either traverse the working or protection / restoration LSP. "D. Requested switch for normal traffic: Recovery signaling is initiated that switches normal traffic to the recovery LSP following the procedures defined in Section 6, 7, 8, and 9. This happens unless a fault condition exists on other LSPs or spans (including the recovery LSP), or a switch command of equal or higher priority is in effect. E. Requested switch for recovery LSP: Recovery signaling is initiated that switches normal traffic to the working LSP following the procedure defined in Section 12. This request is executed except if a fault condition exists on the working LSP or an equal or higher priority switch command is in effect." RFC 4872 at 31.

No.	'821 Patent Claim 3	IETF MPLS-TP System
		"When a failure is detected by one or both endpoints of the LSP, both endpoints must select traffic from the other LSP. This action must be coordinated between node A and D. From this perspective, 1+1 bidirectional protection can be seen as a coordinated protection- switching mechanism between both endpoints. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 11. "1:1 Protection with Extra-Traffic The most common case of end-to-end 1:N protection is to establish, between the same endpoints, an end-to-end working LSP (thus, N = 1) and a dedicated end-to-end protecting LSP that are mutually link/ node/SRLG disjoint. This protects against working LSP failure(s)." RFC 4872 at 13. "An approach to reduce recovery resource requirements is to have protection LSPs sharing network resources when the working LSPs that they protect are physically (i.e., link, node, SRLG, etc.) disjoint. This mechanism is referred to as shared mesh restoration and is described in [RFC4426]. Shared-mesh restoration can be seen as a particular case of pre-planned LSP rerouting (see Section 8) that reduces the recovery resource requirements by allowing multiple protecting LSPs to share common link and node resources. Here also, the recovery resources for the protecting LSPs are pre-reserved during the provisioning phase, thus an explicit signaling action is required to activate (i.e., commit resource allocation at the data plane) a specific protecting LSP instantiated during the (pre-) provisioning phase. This requires restoration signaling along the protecting LSP." RFC 4872 at 20. "1. Introduction Generalized Multi-Protocol Label Switching (GMPLS) extends MPLS to include support for Layer-2 Switch Capable (L2SC), Time-Division Multiplex (TDM), Lambda Switch Capable (LSC), and Fiber Switch Capable (FSC) interfaces. GMPLS recovery uses control plane mechanisms (i.e., signaling, routing, and link management mechanisms) to support data plane fault recovery. Note that the analogous (data plane) fault detection mechanisms are required

No.	'821 Patent Claim 3	IETF MPLS-TP System
		to be present in support of the control plane mechanisms. In this document, the term "recovery" is generically used to denote both protection and restoration; the specific terms "protection" and "restoration" are only used when differentiation is required. The subtle distinction between protection and restoration is made based on the resource allocation done during the recovery phase (see [RFC4427]). A functional description of GMPLS recovery is provided in [RFC4426] and should be considered as a companion document. The present document describes the protocol-specific procedures for GMPLS RSVP- TE (Resource ReSerVation Protocol - Traffic Engineering) signaling (see [RFC3473]) to support end-to-end recovery. End-to-end recovery refers to the recovery of an entire LSP from its head-end (ingress node endpoint) to its tail-end (egress node endpoint). With end-to- end recovery, working LSPs are assumed to be resource-disjoint (where a resource is a link, node, or Shared Risk Link Group (SRLG)) in the network so that they do not share any failure probability, but this is not mandatory. With respect to a given set of network resources, a pair of working/protecting LSPs SHOULD be resource disjoint in case of dedicated recovery type (see below). On the other hand, in case of shared recovery (see below), a group of working LSPs SHOULD be mutually resource-disjoint in order to allow for a (single and commonly) shared protecting LSP, itself resource-disjoint from each of the working LSPs. Note that resource disjointness is a necessary (but not sufficient) condition to ensure LSP recoverability. The present document addresses four types of end-to-end LSP recovery: 1) 1+1 (unidirectional/bidirectional) protection, 2) 1:N (N >= 1) LSP protection with extra-traffic, 3) pre-planned LSP rerouting without extra-traffic (including shared mesh), and 4) full LSP rerouting. 29) The simplest notion of end-to-end LSP protection is 1+1 unidirectional protection. Using this type of protection, a protecting LSP is signaled over a dedicated resource-disjoint alternate path to protect an associated working LSP. Normal traffic is simultaneously sent on both LSPs and a selector is used at the egress node to receive traffic from one of the LSPs. If a failure occurs along one of the LSPs, the egress node selects the traffic from the valid LSP. No coordination is required between the end nodes when a failure/switchover occurs.

No.	'821 Patent Claim 3	IETF MPLS-TP System
		In 1+1 bidirectional protection, a protecting LSP is signaled over a dedicated resource-disjoint alternate path to protect the working LSP. Normal traffic is simultaneously sent on both LSPs (in both directions), and a selector is used at both ingress/egress nodes to receive traffic from the same LSP. This requires coordination between the end-nodes when switching to the protecting LSP. 30) In 1:N (N >= 1) protection with extra-traffic, the protecting LSP is a fully provisioned and resource-disjoint LSP from the N working LSPs, that allows for carrying extra-traffic. The N working LSPs MAY be mutually resource-disjoint. Coordination between end-nodes is required when switching from one of the working LSPs to the protecting LSP. As the protecting LSP is fully provisioned, default operations during protection switching are specified for a protecting LSP carrying extra-traffic, but this is not mandatory. Note that M:N protection is out of scope of this document (though mechanisms it defines may be extended to cover it). 31) Pre-planned LSP rerouting (or restoration) relies on the establishment between the same pair of end-nodes of a working LSP and a protecting LSP that is link/node/SRLG disjoint from the working one. Here, the recovery resources for the protecting LSP are pre-reserved but explicit action is required to activate (i.e., commit resource allocation at the data plane) a specific protecting LSP instantiated during the (pre-)provisioning phase. Since the protecting LSP is not "active" (i.e., fully instantiated), it cannot carry any extra-traffic. This does not mean that the corresponding resources cannot be used by other LSPs. Therefore, this mechanism protects against working LSP(s) failure(s) but requires activation of the protecting LSP after working LSP failure occurrence. This requires restoration signaling along the protecting path. "Shared-mesh" restoration can be seen as a particular case of pre-planned LSP rerouting that reduces the recovery resource requirements by allowing multiple protecting LSPs to share common link and node resources. The recovery resources are pre-reserved but explicit action is required to activate (i.e., commit resource allocation at the data plane) a specific protecting LSP instantiated during the (pre-) provisioning phase. This procedure requires restoration signaling along the protecting path. Note that in both cases, bandwidth pre-reserved for a protecting (but not activated) LSP can be made available for carrying extra traffic. LSPs for extra-traffic (with lower

No.	'821 Patent Claim 3	IETF MPLS-TP System
		holding priority than the protecting LSP) can then be established using the bandwidth pre-reserved for the protecting LSP. Also, any lower priority LSP that use the pre-reserved resources for the protecting LSP(s) must be preempted during the activation of the protecting LSP. 32) Full LSP rerouting (or restoration) switches normal traffic to an alternate LSP that is not even partially established until after the working LSP failure occurs. The new alternate route is selected at the LSP head-end node, it may reuse resources of the failed LSP at intermediate nodes and may include additional intermediate nodes and/or links." " RFC 4872 at 3–5. "Consider the following network topology:  The paths [A,B,C,D] and [A,E,F,G,D] are node and link disjoint, ignoring the ingress/egress nodes A and D. A 1+1 protected path is established from A to D over [A,B,C,D] and [A,E,F,G,D], and traffic is transmitted simultaneously over both component paths (i.e., LSPs). During the provisioning phase, both LSPs are fully instantiated (and thus activated) so that no resource sharing can be done along the protecting LSP (nor can any extra-traffic be transported). It is also RECOMMENDED to set the N bit since no protection-switching signaling is assumed in this case. When a failure occurs (say, at node B) and is detected at end-node D, the receiver at D selects the normal traffic from the other LSP. From this perspective, 1+1 unidirectional protection can be seen as an uncoordinated protection-switching mechanism acting independently at both endpoints. Also, for the LSP under failure condition, it is RECOMMENDED to not set the Path_State_Removed Flag of the ERROR_SPEC object (see [RFC3473]) upon PathErr message generation. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 9.

No.	'821 Patent Claim 3	IETF MPLS-TP System
		"When a failure is detected by one or both endpoints of the LSP, both endpoints must select traffic from the other LSP. This action must be coordinated between node A and D. From this perspective, 1+1 bidirectional protection can be seen as a coordinated protection- switching mechanism between both endpoints. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 11. "(Full) LSP Rerouting LSP rerouting, on the other hand, switches normal traffic to an alternate LSP that is fully established only after failure occurrence. The new (alternate) route is selected at the LSP head- end and may reuse intermediate nodes included in the original route; it may also include additional intermediate nodes. For strict-hop routing, TE requirements can be directly applied to the route computation, and the failed node or link can be avoided. However, if the failure occurred within a loose-routed hop, the head-end node may not have enough information to reroute the LSP around the failure. Crankback signaling (see [CRANK]) and route exclusion techniques (see [RFC4874]) MAY be used in this case. The alternate route MAY be either computed on demand (that is, when the failure occurs; this is referred to as full LSP rerouting) or pre-computed and stored for use when the failure is reported. The latter offers faster restoration time. There is, however, a risk that the alternate route will become out of date through other changes in the network; this can be mitigated to some extent by periodic recalculation of idle alternate routes. (Full) LSP rerouting will be initiated by the head-end node that has either detected the LSP failure or received a Notify message and/or a PathErr message with the new error code/sub-code "Notify Error/LSP Locally Failed" for this LSP. The new LSP resources can be established using the make-before-break mechanism, where the new LSP is set up before the old LSP is torn down. This is done by using the mechanisms of the SESSION_ATTRIBUTE object and the Shared-Explicit (SE) reservation style (see [RFC3209]). Both the new and old LSPs can share resources at common nodes. Note that the make-before-break mechanism is not used to avoid disruption to the normal traffic flow (the latter has already been broken by the failure that is being repaired). However,

No.	'821 Patent Claim 3	IETF MPLS-TP System
		it is valuable to retain the resources allocated on the original LSP that will be reused by the new alternate LSP." RFC 4872 at 25. "The procedure is as follows: 25) The initiating (source) node sends the normal traffic onto both the working and the protecting LSPs. Once completed, the source node sends reliably a Notify message to the destination with the new error code/sub-code "Notify Error/LSP Recovered" (Switchback Request). This Notify message includes the MESSAGE_ID object. The ACK_Desired flag MUST be set in this object to request the receiver to send an acknowledgment for the message (see [RFC2961]). 26) Upon receipt of this message, the destination selects the traffic from the working LSP. At the same time, it transmits the traffic onto both the working and protecting LSP. The destination then sends reliably a Notify message to the source confirming the completion of the operation. This message includes the MESSAGE_ID_ACK object to acknowledge reception of the received Notify message. This Notify message also includes the MESSAGE_ID object. The ACK_Desired flag MUST be set in this object to request the receiver to send an acknowledgment for the message (see [RFC2961]). 27) When the source node receives this Notify message, it switches to receive traffic from the working LSP. The source node then sends an Ack message to the destination node confirming that the LSP has been reverted." RFC 4872 at 27. For example, MPLS-TP Protection discloses an active path indicator field that indicates which LSP and/or pseudo-wire, i.e. network entity, will carry network traffic. MPLS-TP Protection further discloses that the active path indicator can select a working path from the entity pair.

No.	'821 Patent Claim 3	IETF MPLS-TP System

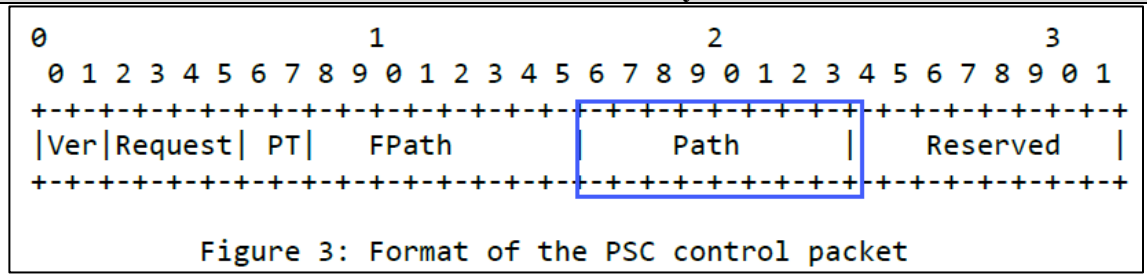


Figure 3: Format of the PSC control packet

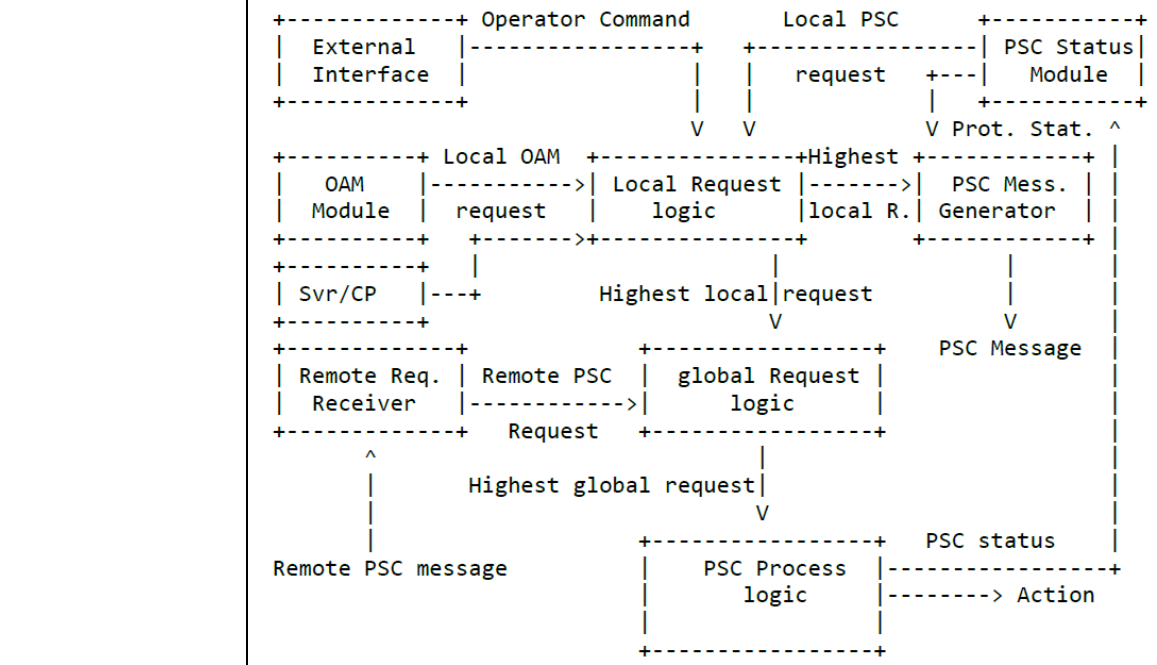
MPLS-TP Figure 3 (annotation added).

"3. Protection switching logic

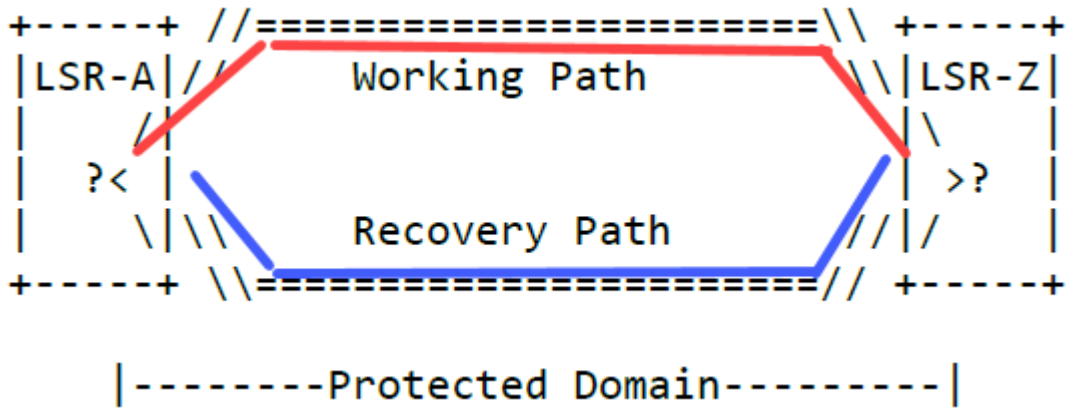
3.1. Protection switching trigger mechanisms

The protection switching should be initiated in reaction to any of the following triggers:

- Server layer indication - if the MPLS-TP server layer detects a failure within its own layer, or due to a failure of its server layer (e.g. the physical layer) notifies the MPLS-TP layer that a failure has been detected. It should be noted that this trigger will be controlled by a holdoff-timer that should be configured to allow the server layer to apply any protection procedures that may be relevant.
- OAM signalling - if, for example, OAM continuity and connectivity verification tools detect that there is a loss of continuity or mis-connectivity or performance monitoring indicates a degradation of the utility of the working path for the current transport path. In cases of signal degradation, switching to the recovery path SHOULD only be activated if the recovery path can guarantee better conditions than the degraded working path.
- Control plane - if there is a control plane active in the network (either signaling or routing), it MAY trigger protection switching based on conditions detected by the control plane. If the control-plane is based on GMPLS [RFC3945] then the recovery process should comply with the process described in [RFC4872].
- Operator command - the network operator may issue commands that trigger protection switching. The commands that are supported include - Forced Switch,

No.	'821 Patent Claim 3	IETF MPLS-TP System
		Manual Switch, Clear, Lockout of Protection, (see definitions in [RFC4427]).” MPLS-TP at 6–7. “3.2. Protection switching control logical architecture Protection switching processes the triggers described above together with the inputs received from the far-end LSR. These inputs cause the LSR to take certain actions, e.g. switching the Selector Bridge to select the working or recovery path, and to transmit different protocol messages.  Figure 1: Protection switching control logic Figure 1 describes the logical architecture of the protection switching control. The Local Request logic unit accepts the triggers from the OAM, external operator commands, and from the local control plane (when present) and determines the highest priority request. This high-

No.	'821 Patent Claim 3	IETF MPLS-TP System
		priority request is passed to both the PSC Message generator, that will generate the appropriate protocol message to be sent to the far-end LSR, and the Global Request logic, that will cross-check this local request with the information received from the far-LSR. The Global Request logic then processes these two PSC requests that determines the highest priority request that is passed to the PSC Process logic. The PSC Process logic uses this input to determine what actions need to be taken, e.g. switching the Selector Bridge, and the current status of the protection domain." MPLS-TP Protection at 7–8. "4.2.4. Active path indicator (Path) The Path field of the PSC control SHALL be used to indicate which path the source MEP is currently using for data transmission. The MEP should compare the value of this bit with the path that is locally selected for data transmission to verify that there is no inconsistency between the two end-points of the protected domain. If an inconsistency is detected then an alarm should be raised. The following are the possible values: • 0: indicates that normal traffic is being transmitted on the Working path. • 1: indicates the Recovery path is being used to transmit the normal traffic from the Working path. 2-255: for future extensions" MPLS-TP Protection at 12–13. "Principles of Operation In all of the following sub-sections, assume a protected domain between LSR-A and LSR-Z, using paths W (working) and R (recovery) as shown in figure 4." MPLS-TP Protection at 13.

No.	'821 Patent Claim 3	IETF MPLS-TP System
		 MPLS-TP Protection Figure 4 (annotation added). "4.3.1 PSC States 4.3.1.1 Normal State When the protected domain has no special condition in effect, the ingress LSR SHOULD forward the user data along the working path, and, in the case of 1+1 protection, the Permanent Bridge will bridge the data to the recovery path as well. The receiving LSR SHOULD read the data from the working path. The ingress LSR MAY transmit a No Request PSC packet with the Path field set to 0 indicating that the normal data traffic should be read from the working path. 4.3.1.2. Protecting State When the protection mechanism has been triggered and the protected domain has performed a protection switch, the domain is in the protecting state. In this state the normal data traffic is transmitted and received on the recovery path. If the protection domain is currently in a protecting state, then the LSRs SHOULD NOT accept a Manual Switch request. If the protection domain is currently in a protecting state, and a Forced Switch is requested then the

No.	'821 Patent Claim 3	IETF MPLS-TP System
		normal traffic SHALL continue to be transmitted on the recovery path even if the original protection trigger is cleared, and the Forced Switch condition will be signalled by the PSC messages. 4.3.1.3. Unavailable State When the recovery path is unavailable - either as a result of a Lockout operator command (see section 4.3.3), or as a result of a SF or SD detected on the recovery path (see section 4.3.4) - then the protection domain is in the unavailable state. In this state, the normal traffic is transmitted and received on the working path. While in unavailable state any event that would trigger a protection switching SHOULD be ignored with the following exception - If a Signal Degrade request is received, then protection switching will be activated only if the recovery path can guarantee a better signal than the working path. The protection domain will exit the unavailable state and revert to the normal state when, either the operator clears the Lockout command or the recovery path recovers from the signal fault or degraded situation. Both ends will resume sending the PCS packets over the recovery path, as a result of this recovery. 4.3.2. Failure or Degraded condition (Working path) If one of the LSRs (for example, LSR-A) detects a failure condition or a serious degradation condition on the working path that warrants invoking protection switching, then it SHOULD take the following actions: • (For 1:1 protection) Switch all traffic for LSR-Z to the recovery path only. • Transmit a PCS control packet, using GACH, with the appropriate Request code (either Signal fault or Signal degrade), the Fpath set to 1, to indicate that the fault/degrade was detected on the working path, and the Path set to 1, indicating that normal traffic is now being transmitted on the recovery path. • Verify that LSR-Z replies with a PCS control packet indicating that it has switched to the recovery path. If this is not received after 2 PSC cycles then send an alarm to the management system.

No.	'821 Patent Claim 3	IETF MPLS-TP System
		When the far-end LSR (in this example LSR-Z) receives the PCS packet informing it that other LSR (LSR-A) has switched, it SHOULD perform the following actions: • Check priority of the request • Switch all traffic addressed to LSR-A to the recovery path only (for 1:1 protection). • Begin transmission of a PCS control packet, using GACH, with the appropriate Request code (either Signal fault or Signal degrade), the Fpath set to 1, to indicate that the fault/degrade was detected on the working path, and the Path set to 1, indicating that traffic is now being transmitted on the recovery path." MPLS-TP Protection at 13–15. "4.3.4. Failure or Degraded condition (Recovery path) If one of the LSRs (for example, LSR-A) detects a failure condition or a serious degradation condition on the recovery path, then it SHOULD take the following actions: • Begin transmission of a PCS control packet with the appropriate Request code (either Signal fault or Signal degrade), the Fpath set to 0, to indicate that the fault/degrade was detected on the recovery path, and the Path set to 0, indicating that traffic is now being forwarded on the working path. Note that this will actually reach the far-end if this is a unidirectional fault or recovery path is possibly in a degraded situation. • The PSC control logic should go into Unavailable state. • All traffic MUST be transmitted on the working path for the duration of the SF/SD condition. When the far-end LSR (in this example LSR-Z) receives the PCS packet informing it that other LSR (LSR-A) has become Unavailable, it SHOULD perform the following actions: • Transmit all traffic on the working path for the duration of the SF/SD condition The PSC Control logic should go into Unavailable state." MPLS-TP Protection at 16.

No.	'821 Patent Claim 4	IETF MPLS-TP System
4	The method of claim 2, wherein selecting an entity pair further	IETF MPLS-TP System discloses the method of claim 2, wherein selecting an entity pair further comprises minimizing an overall cost function.

No.	'821 Patent Claim 4	IETF MPLS-TP System
	comprises minimizing an overall cost function.	For example, RFC 4872 discloses that the plurality of MPLS LSPs, i.e., transport entities, are selected to necessarily be "resource disjoint," i.e., where multiple resources, such as links, nodes, or Shared Risk Link Groups (SRLGs) do not share common links or failure probabilities. Thus, at least under the apparent claim scope alleged by Orckit's Infringement Disclosures, this limitation is met. To the extent that the IETF MPLS-TP System is found to not meet this limitation, wherein selecting an entity pair further comprises minimizing an overall cost function would have been obvious to a person having ordinary skill in the art, as explained below. "When a failure is detected by one or both endpoints of the LSP, both endpoints must select traffic from the other LSP. This action must be coordinated between node A and D. From this perspective, 1+1 bidirectional protection can be seen as a coordinated protection- switching mechanism between both endpoints. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 11. "1:1 Protection with Extra-Traffic The most common case of end-to-end 1:N protection is to establish, between the same endpoints, an end-to-end working LSP (thus, N = 1) and a dedicated end-to-end protecting LSP that are mutually link/ node/SRLG disjoint. This protects against working LSP failure(s)." RFC 4872 at 13. "An approach to reduce recovery resource requirements is to have protection LSPs sharing network resources when the working LSPs that they protect are physically (i.e., link, node, SRLG, etc.) disjoint. This mechanism is referred to as shared mesh restoration and is described in [RFC4426]. Shared-mesh restoration can be seen as a particular case of pre-planned LSP rerouting (see Section 8) that reduces the recovery resource requirements by allowing multiple protecting LSPs to share common link and node resources. Here also, the recovery resources for the protecting LSPs are pre-reserved during the provisioning phase, thus an explicit signaling action is required to activate (i.e., commit resource allocation at the

No.	'821 Patent Claim 4	IETF MPLS-TP System
		data plane) a specific protecting LSP instantiated during the (pre-) provisioning phase. This requires restoration signaling along the protecting LSP." RFC 4872 at 20. "1. Introduction Generalized Multi-Protocol Label Switching (GMPLS) extends MPLS to include support for Layer-2 Switch Capable (L2SC), Time-Division Multiplex (TDM), Lambda Switch Capable (LSC), and Fiber Switch Capable (FSC) interfaces. GMPLS recovery uses control plane mechanisms (i.e., signaling, routing, and link management mechanisms) to support data plane fault recovery. Note that the analogous (data plane) fault detection mechanisms are required to be present in support of the control plane mechanisms. In this document, the term "recovery" is generically used to denote both protection and restoration; the specific terms "protection" and "restoration" are only used when differentiation is required. The subtle distinction between protection and restoration is made based on the resource allocation done during the recovery phase (see [RFC4427]). A functional description of GMPLS recovery is provided in [RFC4426] and should be considered as a companion document. The present document describes the protocol-specific procedures for GMPLS RSVP- TE (Resource ReSerVation Protocol - Traffic Engineering) signaling (see [RFC3473]) to support end-to-end recovery. End-to-end recovery refers to the recovery of an entire LSP from its head-end (ingress node endpoint) to its tail-end (egress node endpoint). With end-to- end recovery, working LSPs are assumed to be resource-disjoint (where a resource is a link, node, or Shared Risk Link Group (SRLG)) in the network so that they do not share any failure probability, but this is not mandatory. With respect to a given set of network resources, a pair of working/protecting LSPs SHOULD be resource disjoint in case of dedicated recovery type (see below). On the other hand, in case of shared recovery (see below), a group of working LSPs SHOULD be mutually resource-disjoint in order to allow for a (single and commonly) shared protecting LSP, itself resource-disjoint from each of the working LSPs. Note that resource disjointness is a necessary (but not sufficient) condition to ensure LSP recoverability. The present document addresses four types of end-to-end LSP recovery: 1) 1+1 (unidirectional/bidirectional) protection, 2) 1:N (N >= 1) LSP protection with extra-traffic,

No.	'821 Patent Claim 4	IETF MPLS-TP System
		3) pre-planned LSP rerouting without extra-traffic (including shared mesh), and 4) full LSP rerouting. 33) The simplest notion of end-to-end LSP protection is 1+1 unidirectional protection. Using this type of protection, a protecting LSP is signaled over a dedicated resource-disjoint alternate path to protect an associated working LSP. Normal traffic is simultaneously sent on both LSPs and a selector is used at the egress node to receive traffic from one of the LSPs. If a failure occurs along one of the LSPs, the egress node selects the traffic from the valid LSP. No coordination is required between the end nodes when a failure/switchover occurs. In 1+1 bidirectional protection, a protecting LSP is signaled over a dedicated resource-disjoint alternate path to protect the working LSP. Normal traffic is simultaneously sent on both LSPs (in both directions), and a selector is used at both ingress/egress nodes to receive traffic from the same LSP. This requires coordination between the end-nodes when switching to the protecting LSP. 34) In 1:N (N >= 1) protection with extra-traffic, the protecting LSP is a fully provisioned and resource-disjoint LSP from the N working LSPs, that allows for carrying extra-traffic. The N working LSPs MAY be mutually resource-disjoint. Coordination between end-nodes is required when switching from one of the working LSPs to the protecting LSP. As the protecting LSP is fully provisioned, default operations during protection switching are specified for a protecting LSP carrying extra-traffic, but this is not mandatory. Note that M:N protection is out of scope of this document (though mechanisms it defines may be extended to cover it). 35) Pre-planned LSP rerouting (or restoration) relies on the establishment between the same pair of end-nodes of a working LSP and a protecting LSP that is link/node/SRLG disjoint from the working one. Here, the recovery resources for the protecting LSP are pre-reserved but explicit action is required to activate (i.e., commit resource allocation at the data plane) a specific protecting LSP instantiated during the (pre-)provisioning phase. Since the protecting LSP is not "active" (i.e., fully instantiated), it cannot carry any extra-traffic. This does not mean that the corresponding resources cannot be used by other LSPs. Therefore, this mechanism protects against working LSP(s) failure(s) but requires activation of the protecting LSP after working LSP

No.	'821 Patent Claim 4	IETF MPLS-TP System
		failure occurrence. This requires restoration signaling along the protecting path. "Shared-mesh" restoration can be seen as a particular case of pre-planned LSP rerouting that reduces the recovery resource requirements by allowing multiple protecting LSPs to share common link and node resources. The recovery resources are pre-reserved but explicit action is required to activate (i.e., commit resource allocation at the data plane) a specific protecting LSP instantiated during the (pre-) provisioning phase. This procedure requires restoration signaling along the protecting path. Note that in both cases, bandwidth pre-reserved for a protecting (but not activated) LSP can be made available for carrying extra traffic. LSPs for extra-traffic (with lower holding priority than the protecting LSP) can then be established using the bandwidth pre-reserved for the protecting LSP. Also, any lower priority LSP that use the pre-reserved resources for the protecting LSP(s) must be preempted during the activation of the protecting LSP. 36) Full LSP rerouting (or restoration) switches normal traffic to an alternate LSP that is not even partially established until after the working LSP failure occurs. The new alternate route is selected at the LSP head-end node, it may reuse resources of the failed LSP at intermediate nodes and may include additional intermediate nodes and/or links." " RFC 4872 at 3–5. "Consider the following network topology:  The paths [A,B,C,D] and [A,E,F,G,D] are node and link disjoint, ignoring the ingress/egress nodes A and D. A 1+1 protected path is established from A to D over [A,B,C,D] and [A,E,F,G,D], and traffic is transmitted simultaneously over both component paths (i.e., LSPs). During the provisioning phase, both LSPs are fully instantiated (and thus activated) so that no resource sharing can be done along the protecting LSP (nor can any extra-traffic be

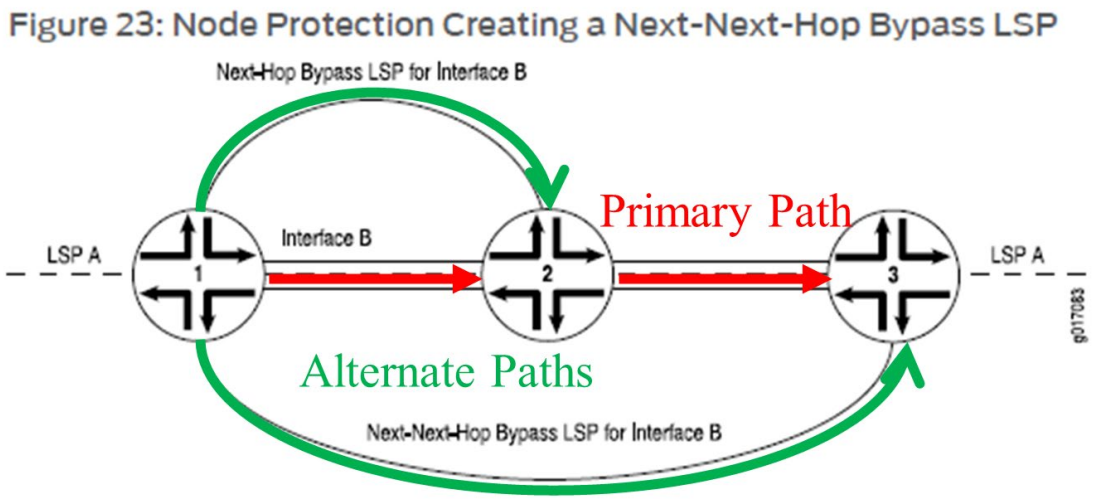
No.	'821 Patent Claim 4	IETF MPLS-TP System
		transported). It is also RECOMMENDED to set the N bit since no protection-switching signaling is assumed in this case. When a failure occurs (say, at node B) and is detected at end-node D, the receiver at D selects the normal traffic from the other LSP. From this perspective, 1+1 unidirectional protection can be seen as an uncoordinated protection-switching mechanism acting independently at both endpoints. Also, for the LSP under failure condition, it is RECOMMENDED to not set the Path_State_Removed Flag of the ERROR_SPEC object (see [RFC3473]) upon PathErr message generation. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 9. "When a failure is detected by one or both endpoints of the LSP, both endpoints must select traffic from the other LSP. This action must be coordinated between node A and D. From this perspective, 1+1 bidirectional protection can be seen as a coordinated protection- switching mechanism between both endpoints. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 11. "(Full) LSP Rerouting LSP rerouting, on the other hand, switches normal traffic to an alternate LSP that is fully established only after failure occurrence. The new (alternate) route is selected at the LSP head- end and may reuse intermediate nodes included in the original route; it may also include additional intermediate nodes. For strict-hop routing, TE requirements can be directly applied to the route computation, and the failed node or link can be avoided. However, if the failure occurred within a loose-routed hop, the head-end node may not have enough information to reroute the LSP around the failure. Crankback signaling (see [CRANK]) and route exclusion techniques (see [RFC4874]) MAY be used in this case. The alternate route MAY be either computed on demand (that is, when the failure occurs; this is referred to as full LSP rerouting) or pre-computed and stored for use when the failure is reported. The latter offers faster restoration time. There is, however, a risk that the alternate

No.	'821 Patent Claim 4	IETF MPLS-TP System
		route will become out of date through other changes in the network; this can be mitigated to some extent by periodic recalculation of idle alternate routes. (Full) LSP rerouting will be initiated by the head-end node that has either detected the LSP failure or received a Notify message and/or a PathErr message with the new error code/sub-code "Notify Error/LSP Locally Failed" for this LSP. The new LSP resources can be established using the make-before-break mechanism, where the new LSP is set up before the old LSP is torn down. This is done by using the mechanisms of the SESSION_ATTRIBUTE object and the Shared-Explicit (SE) reservation style (see [RFC3209]). Both the new and old LSPs can share resources at common nodes. Note that the make-before-break mechanism is not used to avoid disruption to the normal traffic flow (the latter has already been broken by the failure that is being repaired). However, it is valuable to retain the resources allocated on the original LSP that will be reused by the new alternate LSP." RFC 4872 at 25. "The procedure is as follows: 28) The initiating (source) node sends the normal traffic onto both the working and the protecting LSPs. Once completed, the source node sends reliably a Notify message to the destination with the new error code/sub-code "Notify Error/LSP Recovered" (Switchback Request). This Notify message includes the MESSAGE_ID object. The ACK_Desired flag MUST be set in this object to request the receiver to send an acknowledgment for the message (see [RFC2961]). 29) Upon receipt of this message, the destination selects the traffic from the working LSP. At the same time, it transmits the traffic onto both the working and protecting LSP. The destination then sends reliably a Notify message to the source confirming the completion of the operation. This message includes the MESSAGE_ID_ACK object to acknowledge reception of the received Notify message. This Notify message also includes the MESSAGE_ID object. The ACK_Desired flag MUST be set in this object to request the receiver to send an acknowledgment for the message (see [RFC2961]). 30) When the source node receives this Notify message, it switches to receive traffic from the working LSP.

No.	'821 Patent Claim 4	IETF MPLS-TP System
		The source node then sends an Ack message to the destination node confirming that the LSP has been reverted." RFC 4872 at 27. In addition, at least under the apparent claim scope alleged by Orckit's Infringement Disclosures, IETF MPLS System in combination with (1) the knowledge of a person of ordinary skill in the art, alone or in further combination with (2) each (individually, as well as one or more together) of the references identified in element 4 of Exhibit E-2 renders the claim, including the present limitation, obvious. Below are examples of two such references. For example, Juniper IOS Guide discloses "Fate Sharing" where a Juniper router computes a primary and secondary path using CSPF, CSPF identifies all fate-sharing groups that are associated with the primary path by identifying all links that the primary path transverses and checks each link in the traffic engineering database against the compiled group list. Juniper IOS Guide further discloses that if the link is a member of a group, the cost of the link is increased by the cost of the group and higher the cost value. At least based upon Orckit's apparent understanding of the claim scope as reflected in its Infringement Disclosures, this limitation is met. Juniper IOS Guide at Page No. 73 ("[Fate Sharing Fate sharing allows you to create a database of information that CSPF uses to compute one or more backup paths to use in case the primary path becomes unstable. The database describes the relationships between elements of the network, such as routers and links. You can specify one or more elements within a group. Through fate sharing, you can configure backup paths that minimize the number of shared links and fiber paths with the primary paths as much as possible, to ensure that if a fiber is cut, the minimum amount of data is lost and a path still exists to the destination. For a backup path to work optimally, it must not share links or physical fiber paths with the primary path, ensuring that a single point of failure will not affect the primary and backup paths simultaneously. For more information about fate sharing, see the Junos OS Routing Protocols Configuration Guide.]")

No.	'821 Patent Claim 4	IETF MPLS-TP System
		Juniper IOS Guide at Page No. 94 ("[Configuring the Ingress Router for MPLS-Signaled LSPs MPLS-signaled label-switched paths (LSPs) run from a specific ingress router to a specific egress router. For basic MPLS-signaled LSP function, you must configure the ingress router, but do not have to configure any other routers. To configure signaled LSPs, perform the following tasks on the ingress router: • Creating Named Paths on page 56 • Configuring Alternate Backup Paths Using Fate Sharing on page 58.]") Juniper IOS Guide at Page No. 96 ("[Configuring Alternate Backup Paths Using Fate Sharing You can create a database of information that Constrained Shortest Path First (CSPF) uses to compute one or more backup paths in case the primary path becomes unstable. The database describes the relationships between elements of the network, such as routers and links. Because these network elements share the same fate, this relationship is called fate sharing. You can configure backup paths that minimize the number of shared links and fiber paths with the primary paths as much as possible to ensure that, if a fiber is cut, the minimum amount of data is lost and a path still exists to the destination. For a backup path to work optimally, it must not share links or physical fiber paths with the primary path. This ensures that a single point of failure will not affect the primary and backup paths at the same time.]") Juniper IOS Guide at Page No. 97 ("[Fate-sharing groups contain three types of objects: • Point-to-point links—Identified by the IP addresses at each end of the link. Unnumbered point-to-point links are typically identified by borrowing IP addresses from other interfaces. Order is not important; from 1.2.3.4 to 1.2.3.5 and from 1.2.3.5 to 1.2.3.4 have the same meaning. • Non-point-to-point links—Include links on a LAN interface (such as Gigabit Ethernet interfaces) or nonbroadcast multiaccess (NBMA) interfaces (such as Asynchronous Transfer Mode [ATM] or Frame Relay). You identify these links by their individual interface address. For example, if the LAN interface 192.168.200.0/24 has four routers attached to it, each router link is individually identified: from 192.168.200.1; # LAN interface of router 1

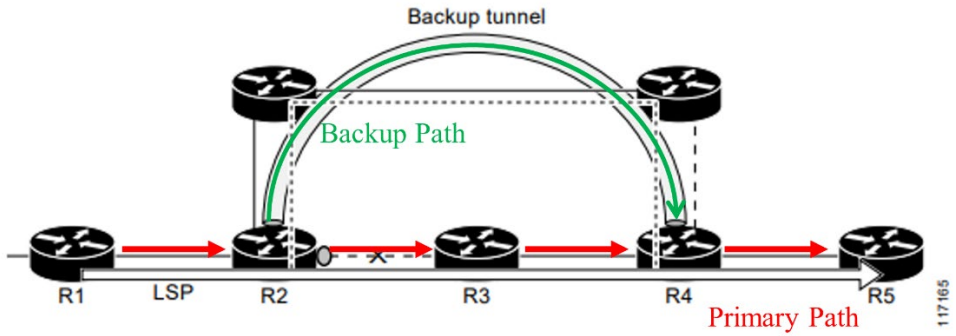
No.	'821 Patent Claim 4	IETF MPLS-TP System
		from 192.168.200.2; # LAN interface of router 2 from 192.168.200.3; # LAN interface of router 3 from 192.168.200.4; # LAN interface of router 4 You can list the addresses in any order. • A router node—Identified by its configured router ID. An object can be in any number of groups, and a group can contain any number of objects. Each group has a configurable cost attributed to it, which represents the level of impact this group has on CSPF computations. The higher the cost, the less likely a backup path will share with the primary path any objects in the group. The cost is directly comparable to traffic engineering metrics. By default, the cost is 1. Changing the fate-sharing database does not affect established LSPs until the next reoptimization of CSPF. The fate-sharing database does influence fast-reroute computations.]") Juniper IOS Guide at Page No. 97 ("[When CSPF computes the primary paths of an LSP (or secondary paths when the primary path is not active), it ignores the fate-sharing information. You always want to find the best possible path (least IGP cost) for the primary path. When CSPF computes a secondary path while the primary path (of the same LSP) is active, the following occurs: 1. CSPF identifies all fate-sharing groups that are associated with the primary path. CSPF does this by identifying all links and nodes that the primary path traverses and compiling group lists that contain at least one of the links or nodes. CSPF ignores the ingress and egress nodes in the search. 2. CSPF checks each link in the traffic engineering database against the compiled group list. If the link is a member of a group, the cost of the link is increased by the cost of the group. If a link is a member of multiple groups, all group costs are added together. 3. CSPF performs the check for every node in the traffic engineering database, except the ingress and egress node. Again, a node can belong to multiple groups, so costs are additive. 4. The router performs regular CSPF computation with the adjusted topology.]") Juniper IOS Guide at Page No. 330 ("In Figure 23 on page 292, both node protection and link protection are enabled on Interface B on Router 1. Both node protection and link protection are also enabled on LSP A, an LSP that traverses the link transiting Router 1,

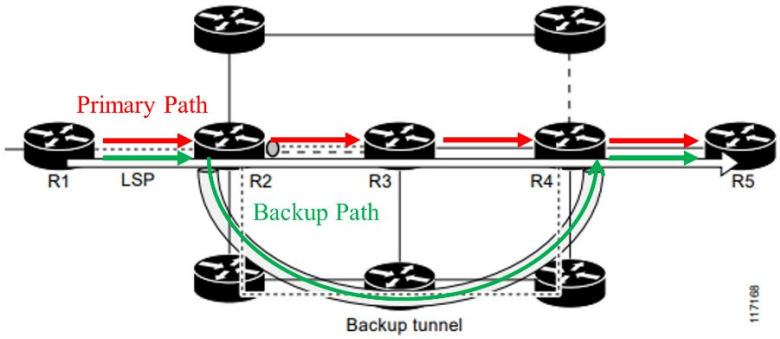
No.	'821 Patent Claim 4	IETF MPLS-TP System
		Router 2, and Router 3. If the link between Router 1 and Router 2 fails, traffic from LSP A is quickly switched to the next-hop bypass LSP generated by link protection. If Router 2 suffers a hardware or software failure, traffic from LSP A is switched to the next-next-hop bypass LSP generated by node protection.")  Fig. 23 (annotation added). For example, Cisco IOS Guide discloses that the path an LSP uses is determined by the LSP resource requirements and network resources, such as bandwidth. Cisco IOS Guide further discloses the MPLS TE SRLG feature that enhances backup tunnel path selection so a backup tunnel can avoid using links that are in the same SRLG as the interfaces it is protecting. Thus, at least under the apparent claim scope alleged by Orckit's Infringement Disclosures, this limitation is met. Cisco IOS Guide at Page No. 253 ("[Path Option Selection for MPLS TE Tunnel LSPs This section contains the following topics about path option selection for MPLS TE Tunnel LSPs: • Constraint-Based Routing and Path Option Selection, page 5

No.	'821 Patent Claim 4	IETF MPLS-TP System
		• Tunnel Reoptimization and Path Option Selection, page 5 • Path Option Selection with Bandwidth Override, page 6 Constraint-Based Routing and Path Option Selection MPLS traffic engineering automatically establishes and maintains LSPs across the backbone by using the Resource Reservation Protocol (RSVP). The path that an LSP uses is determined by the LSP resource requirements and network resources, such as bandwidth. Traffic engineering tunnels are calculated at the LSP head based on a fit between required and available resources (constraint-based routing). Without the Path Option for Bandwidth Override feature, a TE tunnel establishes an LSP based on dynamic or explicit path options in order of preference. However, the bandwidth and other attributes configured on the TE tunnel allow the setup of an LSP only if LSP path options satisfy the constraints. If a path cannot be found that satisfies the configured path options, then the tunnel is not set up.]") Cisco IOS Guide at Page No. 639 ("[Bandwidth Protection NHOP and NNHOP backup tunnels can be used to provide bandwidth protection for rerouted LSPs. This is referred to as backup bandwidth. You can associate backup bandwidth with NHOP or NNHOP backup tunnels. This informs the router of the amount of backup bandwidth a particular backup tunnel can protect. When a router maps LSPs to backup tunnels, bandwidth protection ensures that an LSP uses a given backup tunnel only if there is sufficient backup bandwidth. The router selects which LSPs use which backup tunnels in order to provide maximum bandwidth protection. That is, the router determines the best way to map LSPs onto backup tunnels in order to maximize the number of LSPs that can be protected. For information about mapping tunnels and assigning backup bandwidth, see the "Backup Tunnel Selection Procedure" section on page 10. LSPs that have the "bandwidth protection desired" bit set have a higher right to select backup tunnels that provide bandwidth protection; that is, those LSPs can preempt other LSPs that do not have that bit set. For more information, see the "Prioritizing Which LSPs Obtain Backup Tunnels with Bandwidth Protection" section on page 8.]") Cisco IOS Guide at Page No. 640 ("[Backup Tunnel Support Backup tunnel support has the following capabilities:

No.	'821 Patent Claim 4	IETF MPLS-TP System
		• Backup Tunnels Can Terminate at the Next-Next Hop to Support FRR, page 6 • Multiple Backup Tunnels Can Protect the Same Interface, page 6 • Backup Tunnels Provide Scalability, page 7.]") Cisco IOS Guide at Page 640 ("[Multiple Backup Tunnels Can Protect the Same Interface There is no limit (except memory limitations) to the number of backup tunnels that can protect a given interface. In many topologies, support for node protection requires supporting multiple backup tunnels per protected interface. These backup tunnels can terminate at the same destination or at different destinations. That is, for a given protected interface, you can configure multiple NHOP or NNHOP backup tunnels. This allows redundancy and load balancing.]") Cisco IOS Guide at Page No. 603 ("[Traffic Engineering Tunnels MPLS TE lets you build label switched paths (LSPs) across your network for forwarding traffic. MPLS TE LSPs let the headend of a TE tunnel control the path its traffic takes to a particular destination. This method is more flexible than forwarding traffic based only on a destination address. Interarea tunnels allow you to do the following: • Build TE tunnels between areas (interarea tunnels) • Build TE tunnels that start and end in the same area, on multiple areas on a router (intra-area tunnels) Some tunnels are more important than others. For example, you may have tunnels carrying VoIP traffic and tunnels carrying data traffic that are competing for the same resources. MPLS TE allows you to have some tunnels preempt others. Each tunnel has a priority, and more-important tunnels take precedence over less-important tunnels.]") Cisco IOS Guide at Page No. 641 ("[Backup Bandwidth Protection Backup bandwidth protection allows you to give LSPs carrying certain kinds of data (such as voice) priority for using backup tunnels.]")

No.	'821 Patent Claim 4	IETF MPLS-TP System
		Cisco IOS Guide at Page No. 650 ("[For each protected link AB with a maximum reservable subpool value of n, there may be a path from node A to node B such that the difference between the maximum reservable global and the maximum reservable subpool is at least the value of n. If it is possible to find such paths for each link in the network, you can establish all the backup tunnels along such paths without any bandwidth reservations. If there is a single link failure, only one backup tunnel will use any link on its path. Because that path has at least n available bandwidth (in the global pool), assuming that marking and scheduling is configured to classify the subpool traffic into a priority queue, the subpool bandwidth is guaranteed. This approach allows sharing of the global pool bandwidth between backup tunnels protecting independent link failures. The backup tunnels are expected to be used for only a short period of time after a failure (until the headends of affected LSPs reroute those LSPs to other paths with available subpool bandwidth). The probability of multiple unrelated link failures is very small (in the absence of node or shared risk link group (SRLG) failures, which result in multiple link failures). Therefore, it is reasonable to assume that link failures are in practice independent with high probability. This "independent failure assumption" in combination with backup tunnels signaled without explicit bandwidth reservation enables efficient bandwidth sharing that yields substantial bandwidth savings.]") Cisco IOS Guide at Page No. 651 ("[A similar approach can be used for node and SRLG protection. However, the decision of where to put the backup tunnels is more complicated because both node and SRLG failures effectively result in the simultaneous failure of several links. Therefore, the backup tunnels protecting traffic traversing all affected links cannot be computed independently of each other.]") Cisco IOS Guide at Page No. 579 ("[MPLS Traffic Engineering Shared Risk Link Groups SRLGs refer to situations where links in a network share a common fiber (or a common physical attribute). If one link fails, other links in the group may fail too. Links in the group have a shared risk. Backup tunnels should avoid using links in the same SRLG as interfaces they are protecting. Otherwise, when the protected link fails the backup tunnel fails too.

No.	'821 Patent Claim 4	IETF MPLS-TP System
		Figure 1 Backup Tunnel in the Same SRLG as the Interface It Is Protecting  Fig. 1 (annotation added) Figure 1 shows a primary label-switched path (LSP) from router R1 to router R5. The LSP protects against the failure of the R2-R3 link at R2 via a backup tunnel to R4. If the R2-R3 link fails, link protection reroutes the LSP along the backup tunnel. However, the R2-R3 link and one of the backup tunnel links are in the same SRLG. So if the R2-R3 link fails, the backup tunnel may fail too. The MPLS TE SRLG feature enhances backup tunnel path selection so a backup tunnel can avoid using links that are in the same SRLG as the interfaces it is protecting. There are two ways for a backup tunnel to avoid the SRLGs of its protected interface: • The router does not create the backup tunnel unless it avoids SRLGs of the protected interface.]") Cisco IOS Guide at Page No. 580 ("[Open Shortest Path First (OSPF) and Intermediate System-to-Intermediate System (IS-IS) flood the SRLG membership information (including other TE link attributes such as bandwidth availability and affinity) so that all routers in the network have the SRLG information for each link. With this topology information, routers can compute backup tunnel paths that exclude links having SRLGs in common with their protected interfaces. As shown in Figure 2, the backup tunnel avoids the link between R2 and R3, which shares an SRLG with the protected interface.]")

No.	'821 Patent Claim 4	IETF MPLS-TP System
		Figure 2 Backup Tunnel That Avoids SRLG of Protected Interface  Fig. 2 (annotation added) Cisco IOS Guide at Page No. 587 ("[Step 4 show mpls traffic-eng topology Use this command to show the SRLG link membership flooded via the Interior Gateway Protocol (IGP).]") Cisco IOS Guide at Page 632 ("[IS-IS—Intermediate System-to-Intermediate System. Link-state hierarchical routing protocol that calls for intermediate system (IS) routers to exchange routing information based on a single metric to determine network topology.]") Cisco IOS Guide at Page No. 633 ("[OSPF—Open Shortest Path First. A link-state hierarchical Interior Gateway Protocol routing algorithm, derived from the IS-IS protocol. OSPF features include least-cost routing, multipath routing, and load balancing.]") Cisco IOS Guide at Page No. 581 ("[Autotunnel Backup for MPLS TE SRLGs Autotunnel backup is the ability of routers to create backup tunnels automatically. Therefore, you do not need to preconfigure each backup tunnel and then assign the backup tunnel to the protected interface. Only automatically created backup tunnels can avoid SRLGs or their protected interfaces.]") Cisco IOS Guide at Page No. 583 ("[Configuring the Routers That Automatically Create Backup Tunnels to Avoid MPLS TE SRLGs of Their Protected Interfaces

No.	'821 Patent Claim 4	IETF MPLS-TP System
		Perform the following task to configure routers that automatically create backup tunnels to avoid MPLS TE SRLGs of their protected interfaces. Backup tunnels provide link protection by rerouting traffic to the next hop bypassing failed links or in this instance by avoiding SRLGs.]") Cisco IOS Guide at Page No. 644 ("[Backup Tunnel Selection Procedure When an LSP is signaled, each node along the LSP path that provides FRR protection for the LSP selects a backup tunnel for the LSP to use if either of the following events occurs: • The link to the next hop fails. • The next hop fails. By having the node select the backup tunnel for an LSP before a failure occurs, the LSP can be rerouted onto the backup tunnel quickly if there is a failure.]") Cisco IOS Guide at Page 646 ("[NHOP Versus NNHOP Backup Tunnels More than one backup tunnel can protect a given LSP, where one backup tunnel terminates at the LSP's NNHOP, and the other terminates at the LSP's NHOP. In this case, the router chooses the backup tunnel that terminates at the NNHOP (that is, FRR prefers NNHOP over NHOP backup tunnels).Table 1 lists the tunnel selection priorities. The first choice is an NNHOP backup tunnel that acquires its bandwidth from a subpool or global pool, and has limited bandwidth. If there is no such backup tunnel, the next choice (2) is a next-next hop backup tunnel that acquires a limited amount of bandwidth from any pool. The preferences go from 1 (best) to 8 (worst), where choice 3 is for an NNHOP backup tunnel with an unlimited amount of subpool or global-pool bandwidth.]")

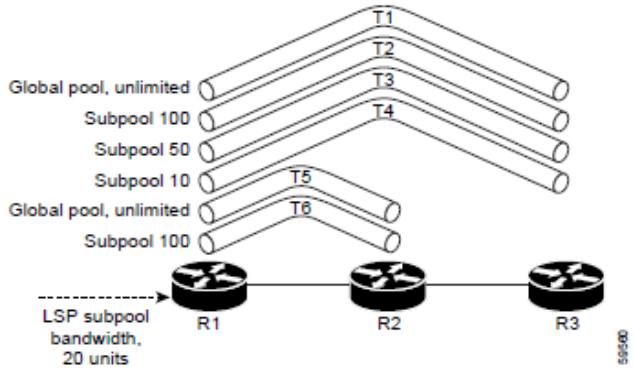
No.	'821 Patent Claim 4	IETF MPLS-TP System

Table 1 Tunnel Selection Priorities

Preference	Backup Tunnel Destination	Bandwidth Pool	Bandwidth Amount
1 (Best)	NNHOP	Subpool or global pool	Limited
2	NNHOP	Any	Limited
3	NNHOP	Subpool or global pool	Unlimited
4	NNHOP	Any	Unlimited
5	NHOP	Subpool or global pool	Limited
6	NHOP	Any	Limited
7	NHOP	Subpool or global pool	Unlimited
8 (Worst)	NHOP	Any	Unlimited

Cisco IOS Guide at Page No. 647-648 ("[Figure 7 shows an example of the backup tunnel selection procedure based on the designated amount of global pool and subpool bandwidth currently available.

Figure 7 Choosing from Among Multiple Backup Tunnels



In this example, an LSP requires 20 units (kilobits per second) of sub-pool backup bandwidth. The best backup tunnel is selected as follows:
1. Backup tunnels T1 through T4 are considered first because they terminate at the NNHOP.
2. Tunnel T4 is eliminated because it has only ten units of sub-pool backup bandwidth.
3. Tunnel T1 is eliminated because it protects only LSPs using global-pool bandwidth.

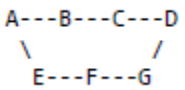
No.	'821 Patent Claim 4	IETF MPLS-TP System
		4. Tunnel T3 is chosen over T2 because, although both have sufficient backup bandwidth, T3 has the least backup bandwidth available (leaving the most backup bandwidth available on T2). 5. Tunnels T5 and T6 need not be considered because they terminate at an NHOP, and therefore are less desirable than T3, which terminates at an NNHOP.]")

No.	'821 Patent Claim 5	IETF MPLS-TP System
5	The method of claim 4, wherein said overall cost function comprises substantially minimizing a probability of concurrent failure of said protection entity and said working entity.	IETF MPLS-TP System discloses the method of claim 4, wherein said overall cost function comprises substantially minimizing a probability of concurrent failure of said protection entity and said working entity. For example, RFC 4872 discloses that the plurality of MPLS LSPs, i.e., transport entities, are selected to necessarily be "resource disjoint," i.e., where multiple resources, such as links, nodes, or Shared Risk Link Groups (SRLGs) do not share common links or failure probabilities. Thus, at least under the apparent claim scope alleged by Orckit's Infringement Disclosures, this limitation is met. To the extent that the IETF MPLS-TP System is found to not meet this limitation, wherein said overall cost function comprises substantially minimizing a probability of concurrent failure of said protection entity and said working entity would have been obvious to a person having ordinary skill in the art, as explained below. "When a failure is detected by one or both endpoints of the LSP, both endpoints must select traffic from the other LSP. This action must be coordinated between node A and D. From this perspective, 1+1 bidirectional protection can be seen as a coordinated protection- switching mechanism between both endpoints. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 11. "1:1 Protection with Extra-Traffic The most common case of end-to-end 1:N protection is to establish, between the same endpoints, an end-to-end working LSP (thus, N = 1) and a dedicated end-to-end protecting

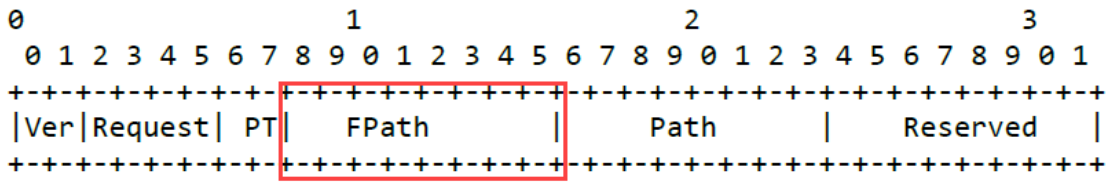
No.	'821 Patent Claim 5	IETF MPLS-TP System
		LSP that are mutually link/ node/SRLG disjoint. This protects against working LSP failure(s)." RFC 4872 at 13. "An approach to reduce recovery resource requirements is to have protection LSPs sharing network resources when the working LSPs that they protect are physically (i.e., link, node, SRLG, etc.) disjoint. This mechanism is referred to as shared mesh restoration and is described in [RFC4426]. Shared-mesh restoration can be seen as a particular case of pre-planned LSP rerouting (see Section 8) that reduces the recovery resource requirements by allowing multiple protecting LSPs to share common link and node resources. Here also, the recovery resources for the protecting LSPs are pre-reserved during the provisioning phase, thus an explicit signaling action is required to activate (i.e., commit resource allocation at the data plane) a specific protecting LSP instantiated during the (pre-) provisioning phase. This requires restoration signaling along the protecting LSP." RFC 4872 at 20. "1. Introduction Generalized Multi-Protocol Label Switching (GMPLS) extends MPLS to include support for Layer-2 Switch Capable (L2SC), Time-Division Multiplex (TDM), Lambda Switch Capable (LSC), and Fiber Switch Capable (FSC) interfaces. GMPLS recovery uses control plane mechanisms (i.e., signaling, routing, and link management mechanisms) to support data plane fault recovery. Note that the analogous (data plane) fault detection mechanisms are required to be present in support of the control plane mechanisms. In this document, the term "recovery" is generically used to denote both protection and restoration; the specific terms "protection" and "restoration" are only used when differentiation is required. The subtle distinction between protection and restoration is made based on the resource allocation done during the recovery phase (see [RFC4427]). A functional description of GMPLS recovery is provided in [RFC4426] and should be considered as a companion document. The present document describes the protocol-specific procedures for GMPLS RSVP- TE (Resource ReSerVation Protocol - Traffic Engineering) signaling (see [RFC3473]) to support end-to-end recovery. End-to-end recovery refers to the recovery of an entire LSP from its head-end (ingress node endpoint) to its tail-end (egress node endpoint). With end-to- end recovery, working LSPs are assumed to be resource-disjoint

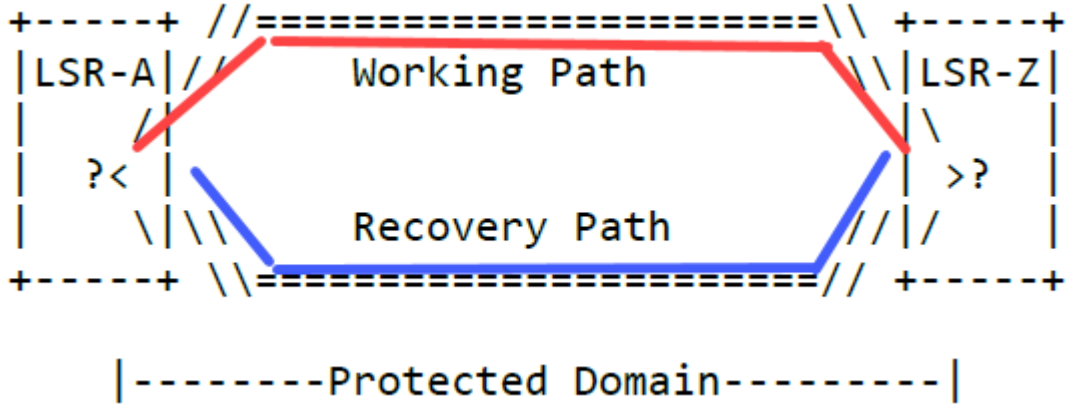
No.	'821 Patent Claim 5	IETF MPLS-TP System
		(where a resource is a link, node, or Shared Risk Link Group (SRLG)) in the network so that they do not share any failure probability, but this is not mandatory. With respect to a given set of network resources, a pair of working/protecting LSPs SHOULD be resource disjoint in case of dedicated recovery type (see below). On the other hand, in case of shared recovery (see below), a group of working LSPs SHOULD be mutually resource-disjoint in order to allow for a (single and commonly) shared protecting LSP, itself resource-disjoint from each of the working LSPs. Note that resource disjointness is a necessary (but not sufficient) condition to ensure LSP recoverability. The present document addresses four types of end-to-end LSP recovery: 1) 1+1 (unidirectional/bidirectional) protection, 2) 1:N (N >= 1) LSP protection with extra-traffic, 3) pre-planned LSP rerouting without extra-traffic (including shared mesh), and 4) full LSP rerouting. 37) The simplest notion of end-to-end LSP protection is 1+1 unidirectional protection. Using this type of protection, a protecting LSP is signaled over a dedicated resource-disjoint alternate path to protect an associated working LSP. Normal traffic is simultaneously sent on both LSPs and a selector is used at the egress node to receive traffic from one of the LSPs. If a failure occurs along one of the LSPs, the egress node selects the traffic from the valid LSP. No coordination is required between the end nodes when a failure/switchover occurs. In 1+1 bidirectional protection, a protecting LSP is signaled over a dedicated resource-disjoint alternate path to protect the working LSP. Normal traffic is simultaneously sent on both LSPs (in both directions), and a selector is used at both ingress/egress nodes to receive traffic from the same LSP. This requires coordination between the end-nodes when switching to the protecting LSP. 38) In 1:N (N >= 1) protection with extra-traffic, the protecting LSP is a fully provisioned and resource-disjoint LSP from the N working LSPs, that allows for carrying extra-traffic. The N working LSPs MAY be mutually resource-disjoint. Coordination between end-nodes is required when switching from one of the working LSPs to the protecting LSP. As the protecting LSP is fully provisioned, default operations during protection switching are specified for a protecting LSP carrying extra-traffic, but this

No.	'821 Patent Claim 5	IETF MPLS-TP System
		is not mandatory. Note that M:N protection is out of scope of this document (though mechanisms it defines may be extended to cover it). 39) Pre-planned LSP rerouting (or restoration) relies on the establishment between the same pair of end-nodes of a working LSP and a protecting LSP that is link/node/SRLG disjoint from the working one. Here, the recovery resources for the protecting LSP are pre-reserved but explicit action is required to activate (i.e., commit resource allocation at the data plane) a specific protecting LSP instantiated during the (pre-)provisioning phase. Since the protecting LSP is not "active" (i.e., fully instantiated), it cannot carry any extra-traffic. This does not mean that the corresponding resources cannot be used by other LSPs. Therefore, this mechanism protects against working LSP(s) failure(s) but requires activation of the protecting LSP after working LSP failure occurrence. This requires restoration signaling along the protecting path. "Shared-mesh" restoration can be seen as a particular case of pre-planned LSP rerouting that reduces the recovery resource requirements by allowing multiple protecting LSPs to share common link and node resources. The recovery resources are pre-reserved but explicit action is required to activate (i.e., commit resource allocation at the data plane) a specific protecting LSP instantiated during the (pre-) provisioning phase. This procedure requires restoration signaling along the protecting path. Note that in both cases, bandwidth pre-reserved for a protecting (but not activated) LSP can be made available for carrying extra traffic. LSPs for extra-traffic (with lower holding priority than the protecting LSP) can then be established using the bandwidth pre-reserved for the protecting LSP. Also, any lower priority LSP that use the pre-reserved resources for the protecting LSP(s) must be preempted during the activation of the protecting LSP. 40) Full LSP rerouting (or restoration) switches normal traffic to an alternate LSP that is not even partially established until after the working LSP failure occurs. The new alternate route is selected at the LSP head-end node, it may reuse resources of the failed LSP at intermediate nodes and may include additional intermediate nodes and/or links." " RFC 4872 at 3–5. "Consider the following network topology:

No.	'821 Patent Claim 5	IETF MPLS-TP System
		 The paths [A,B,C,D] and [A,E,F,G,D] are node and link disjoint, ignoring the ingress/egress nodes A and D. A 1+1 protected path is established from A to D over [A,B,C,D] and [A,E,F,G,D], and traffic is transmitted simultaneously over both component paths (i.e., LSPs). During the provisioning phase, both LSPs are fully instantiated (and thus activated) so that no resource sharing can be done along the protecting LSP (nor can any extra-traffic be transported). It is also RECOMMENDED to set the N bit since no protection-switching signaling is assumed in this case. When a failure occurs (say, at node B) and is detected at end-node D, the receiver at D selects the normal traffic from the other LSP. From this perspective, 1+1 unidirectional protection can be seen as an uncoordinated protection-switching mechanism acting independently at both endpoints. Also, for the LSP under failure condition, it is RECOMMENDED to not set the Path_State_Removed Flag of the ERROR_SPEC object (see [RFC3473]) upon PathErr message generation. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 9. "When a failure is detected by one or both endpoints of the LSP, both endpoints must select traffic from the other LSP. This action must be coordinated between node A and D. From this perspective, 1+1 bidirectional protection can be seen as a coordinated protection- switching mechanism between both endpoints. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 11. "(Full) LSP Rerouting LSP rerouting, on the other hand, switches normal traffic to an alternate LSP that is fully established only after failure occurrence. The new (alternate) route is selected at the LSP head- end and may reuse intermediate nodes included in the original route; it may also include

No.	'821 Patent Claim 5	IETF MPLS-TP System
		additional intermediate nodes. For strict-hop routing, TE requirements can be directly applied to the route computation, and the failed node or link can be avoided. However, if the failure occurred within a loose-routed hop, the head-end node may not have enough information to reroute the LSP around the failure. Crankback signaling (see [CRANK]) and route exclusion techniques (see [RFC4874]) MAY be used in this case. The alternate route MAY be either computed on demand (that is, when the failure occurs; this is referred to as full LSP rerouting) or pre-computed and stored for use when the failure is reported. The latter offers faster restoration time. There is, however, a risk that the alternate route will become out of date through other changes in the network; this can be mitigated to some extent by periodic recalculation of idle alternate routes. (Full) LSP rerouting will be initiated by the head-end node that has either detected the LSP failure or received a Notify message and/or a PathErr message with the new error code/sub-code "Notify Error/LSP Locally Failed" for this LSP. The new LSP resources can be established using the make-before-break mechanism, where the new LSP is set up before the old LSP is torn down. This is done by using the mechanisms of the SESSION_ATTRIBUTE object and the Shared-Explicit (SE) reservation style (see [RFC3209]). Both the new and old LSPs can share resources at common nodes. Note that the make-before-break mechanism is not used to avoid disruption to the normal traffic flow (the latter has already been broken by the failure that is being repaired). However, it is valuable to retain the resources allocated on the original LSP that will be reused by the new alternate LSP." RFC 4872 at 25. "The procedure is as follows: 31) The initiating (source) node sends the normal traffic onto both the working and the protecting LSPs. Once completed, the source node sends reliably a Notify message to the destination with the new error code/sub-code "Notify Error/LSP Recovered" (Switchback Request). This Notify message includes the MESSAGE_ID object. The ACK_Desired flag MUST be set in this object to request the receiver to send an acknowledgment for the message (see [RFC2961]). 32) Upon receipt of this message, the destination selects the traffic from the working LSP. At the same time, it transmits the traffic onto both the working and protecting LSP.

No.	'821 Patent Claim 5	IETF MPLS-TP System
		The destination then sends reliably a Notify message to the source confirming the completion of the operation. This message includes the MESSAGE_ID_ACK object to acknowledge reception of the received Notify message. This Notify message also includes the MESSAGE_ID object. The ACK_Desired flag MUST be set in this object to request the receiver to send an acknowledgment for the message (see [RFC2961]). 33) When the source node receives this Notify message, it switches to receive traffic from the working LSP. The source node then sends an Ack message to the destination node confirming that the LSP has been reverted." RFC 4872 at 27. For example, MPLS-TP Protection discloses an FPath field, that determines whether or not a specified LSP and/or pseudo-wire, i.e. network entity, is operational. At least under Orckit's Infringement Disclosures, this limitation is met. To the extent that the IETF MPLS-TP System is found to not meet this limitation, wherein said overall cost function comprises substantially minimizing a probability of concurrent failure of said protection entity and said working entity would have been obvious to a person having ordinary skill in the art, as explained below.  Figure 3: Format of the PSC control packet MPLS-TP Protection Figure 3, annotation added. "FPath: used to indicate the path that is reporting a failure condition, the possible values are described in section 4.2.3" MPLS-TP Protection at 11. "4.2.3. Path fault identifier (FPath)

No.	'821 Patent Claim 5	IETF MPLS-TP System
		The Fpath field of the PSC control SHALL be used only in a Signal fault (0110) or Signal degrade (0101) control packet. Its value indicates on which path the signal anomaly was detected. The following are the possible values: • 0: indicates that the fault condition is on the Recovery path • 1: indicates that the fault condition is on the Working path • 2-255: for future extensions." MPLS-TP Protection at 12. "Principles of Operation In all of the following sub-sections, assume a protected domain between LSR-A and LSR-Z, using paths W (working) and R (recovery) as shown in figure 4." MPLS-TP Protection at 13.  MPLS-TP Protection Figure 4 (annotation added).

No.	'821 Patent Claim 6	IETF MPLS-TP System
6	The method of claim 4, wherein said overall cost function comprises a	IETF MPLS-TP System discloses the method of claim 4, wherein said overall cost function comprises a predefined entity cost metric. For example, RFC 4872 discloses that the plurality of MPLS LSPs, i.e., transport entities, are selected to necessarily be "resource disjoint," i.e., where multiple resources, such as links,

No.	'821 Patent Claim 6	IETF MPLS-TP System
	predefined entity cost metric.	nodes, or Shared Risk Link Groups (SRLGs) do not share common links or failure probabilities. Thus, at least under the apparent claim scope alleged by Orckit's Infringement Disclosures, this limitation is met. To the extent that the IETF MPLS-TP System is found to not meet this limitation, wherein said overall cost function comprises a predefined entity cost metric would have been obvious to a person having ordinary skill in the art, as explained below. "When a failure is detected by one or both endpoints of the LSP, both endpoints must select traffic from the other LSP. This action must be coordinated between node A and D. From this perspective, 1+1 bidirectional protection can be seen as a coordinated protection- switching mechanism between both endpoints. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 11. "1:1 Protection with Extra-Traffic The most common case of end-to-end 1:N protection is to establish, between the same endpoints, an end-to-end working LSP (thus, N = 1) and a dedicated end-to-end protecting LSP that are mutually link/ node/SRLG disjoint. This protects against working LSP failure(s)." RFC 4872 at 13. "An approach to reduce recovery resource requirements is to have protection LSPs sharing network resources when the working LSPs that they protect are physically (i.e., link, node, SRLG, etc.) disjoint. This mechanism is referred to as shared mesh restoration and is described in [RFC4426]. Shared-mesh restoration can be seen as a particular case of pre-planned LSP rerouting (see Section 8) that reduces the recovery resource requirements by allowing multiple protecting LSPs to share common link and node resources. Here also, the recovery resources for the protecting LSPs are pre-reserved during the provisioning phase, thus an explicit signaling action is required to activate (i.e., commit resource allocation at the data plane) a specific protecting LSP instantiated during the (pre-) provisioning phase. This requires restoration signaling along the protecting LSP." RFC 4872 at 20. "1. Introduction

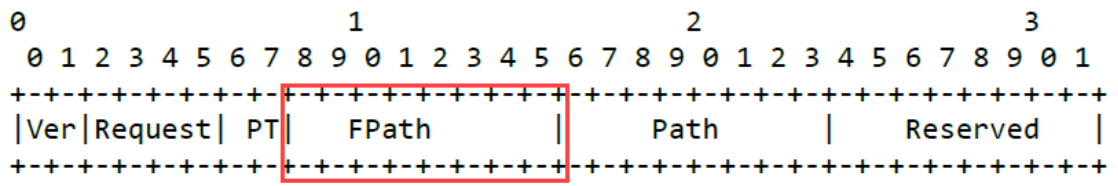
No.	'821 Patent Claim 6	IETF MPLS-TP System
		Generalized Multi-Protocol Label Switching (GMPLS) extends MPLS to include support for Layer-2 Switch Capable (L2SC), Time-Division Multiplex (TDM), Lambda Switch Capable (LSC), and Fiber Switch Capable (FSC) interfaces. GMPLS recovery uses control plane mechanisms (i.e., signaling, routing, and link management mechanisms) to support data plane fault recovery. Note that the analogous (data plane) fault detection mechanisms are required to be present in support of the control plane mechanisms. In this document, the term "recovery" is generically used to denote both protection and restoration; the specific terms "protection" and "restoration" are only used when differentiation is required. The subtle distinction between protection and restoration is made based on the resource allocation done during the recovery phase (see [RFC4427]). A functional description of GMPLS recovery is provided in [RFC4426] and should be considered as a companion document. The present document describes the protocol-specific procedures for GMPLS RSVP- TE (Resource ReSerVation Protocol - Traffic Engineering) signaling (see [RFC3473]) to support end-to-end recovery. End-to-end recovery refers to the recovery of an entire LSP from its head-end (ingress node endpoint) to its tail-end (egress node endpoint). With end-to- end recovery, working LSPs are assumed to be resource-disjoint (where a resource is a link, node, or Shared Risk Link Group (SRLG)) in the network so that they do not share any failure probability, but this is not mandatory. With respect to a given set of network resources, a pair of working/protecting LSPs SHOULD be resource disjoint in case of dedicated recovery type (see below). On the other hand, in case of shared recovery (see below), a group of working LSPs SHOULD be mutually resource-disjoint in order to allow for a (single and commonly) shared protecting LSP, itself resource-disjoint from each of the working LSPs. Note that resource disjointness is a necessary (but not sufficient) condition to ensure LSP recoverability. The present document addresses four types of end-to-end LSP recovery: 1) 1+1 (unidirectional/bidirectional) protection, 2) 1:N (N >= 1) LSP protection with extra-traffic, 3) pre-planned LSP rerouting without extra-traffic (including shared mesh), and 4) full LSP rerouting. 41) The simplest notion of end-to-end LSP protection is 1+1 unidirectional protection. Using this type of protection, a protecting LSP is signaled over a dedicated resource-

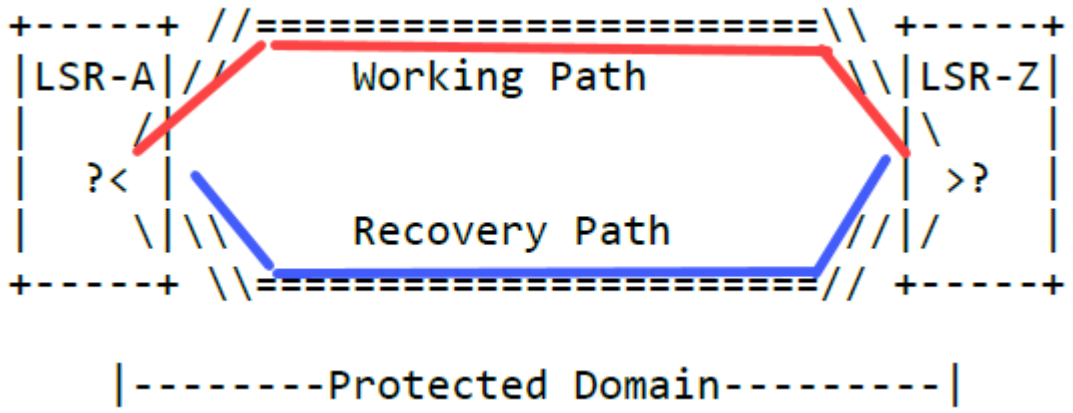
No.	'821 Patent Claim 6	IETF MPLS-TP System
		disjoint alternate path to protect an associated working LSP. Normal traffic is simultaneously sent on both LSPs and a selector is used at the egress node to receive traffic from one of the LSPs. If a failure occurs along one of the LSPs, the egress node selects the traffic from the valid LSP. No coordination is required between the end nodes when a failure/switchover occurs. In 1+1 bidirectional protection, a protecting LSP is signaled over a dedicated resource-disjoint alternate path to protect the working LSP. Normal traffic is simultaneously sent on both LSPs (in both directions), and a selector is used at both ingress/egress nodes to receive traffic from the same LSP. This requires coordination between the end-nodes when switching to the protecting LSP. 42) In 1:N (N >= 1) protection with extra-traffic, the protecting LSP is a fully provisioned and resource-disjoint LSP from the N working LSPs, that allows for carrying extra-traffic. The N working LSPs MAY be mutually resource-disjoint. Coordination between end-nodes is required when switching from one of the working LSPs to the protecting LSP. As the protecting LSP is fully provisioned, default operations during protection switching are specified for a protecting LSP carrying extra-traffic, but this is not mandatory. Note that M:N protection is out of scope of this document (though mechanisms it defines may be extended to cover it). 43) Pre-planned LSP rerouting (or restoration) relies on the establishment between the same pair of end-nodes of a working LSP and a protecting LSP that is link/node/SRLG disjoint from the working one. Here, the recovery resources for the protecting LSP are pre-reserved but explicit action is required to activate (i.e., commit resource allocation at the data plane) a specific protecting LSP instantiated during the (pre-)provisioning phase. Since the protecting LSP is not "active" (i.e., fully instantiated), it cannot carry any extra-traffic. This does not mean that the corresponding resources cannot be used by other LSPs. Therefore, this mechanism protects against working LSP(s) failure(s) but requires activation of the protecting LSP after working LSP failure occurrence. This requires restoration signaling along the protecting path. "Shared-mesh" restoration can be seen as a particular case of pre-planned LSP rerouting that reduces the recovery resource requirements by allowing multiple protecting LSPs to share common link and node resources. The recovery resources are pre-reserved but explicit action is required to activate (i.e., commit resource allocation at the data plane) a

No.	'821 Patent Claim 6	IETF MPLS-TP System
		specific protecting LSP instantiated during the (pre-) provisioning phase. This procedure requires restoration signaling along the protecting path. Note that in both cases, bandwidth pre-reserved for a protecting (but not activated) LSP can be made available for carrying extra traffic. LSPs for extra-traffic (with lower holding priority than the protecting LSP) can then be established using the bandwidth pre-reserved for the protecting LSP. Also, any lower priority LSP that use the pre-reserved resources for the protecting LSP(s) must be preempted during the activation of the protecting LSP. 44) Full LSP rerouting (or restoration) switches normal traffic to an alternate LSP that is not even partially established until after the working LSP failure occurs. The new alternate route is selected at the LSP head-end node, it may reuse resources of the failed LSP at intermediate nodes and may include additional intermediate nodes and/or links." " RFC 4872 at 3–5. "Consider the following network topology:  The paths [A,B,C,D] and [A,E,F,G,D] are node and link disjoint, ignoring the ingress/egress nodes A and D. A 1+1 protected path is established from A to D over [A,B,C,D] and [A,E,F,G,D], and traffic is transmitted simultaneously over both component paths (i.e., LSPs). During the provisioning phase, both LSPs are fully instantiated (and thus activated) so that no resource sharing can be done along the protecting LSP (nor can any extra-traffic be transported). It is also RECOMMENDED to set the N bit since no protection-switching signaling is assumed in this case. When a failure occurs (say, at node B) and is detected at end-node D, the receiver at D selects the normal traffic from the other LSP. From this perspective, 1+1 unidirectional protection can be seen as an uncoordinated protection-switching mechanism acting independently at both endpoints. Also, for the LSP under failure condition, it is RECOMMENDED to not set the Path_State_Removed Flag of the ERROR_SPEC object (see [RFC3473]) upon PathErr message generation.

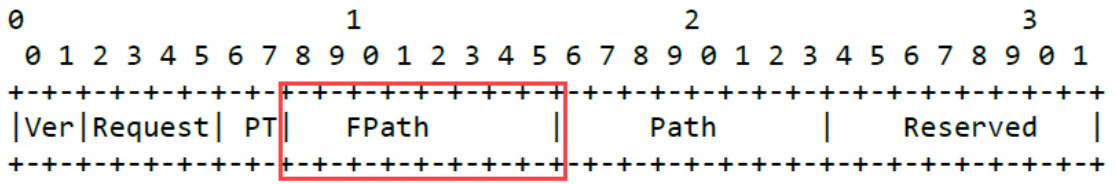
No.	'821 Patent Claim 6	IETF MPLS-TP System
		Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 9. "When a failure is detected by one or both endpoints of the LSP, both endpoints must select traffic from the other LSP. This action must be coordinated between node A and D. From this perspective, 1+1 bidirectional protection can be seen as a coordinated protection- switching mechanism between both endpoints. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 11. "(Full) LSP Rerouting LSP rerouting, on the other hand, switches normal traffic to an alternate LSP that is fully established only after failure occurrence. The new (alternate) route is selected at the LSP head- end and may reuse intermediate nodes included in the original route; it may also include additional intermediate nodes. For strict-hop routing, TE requirements can be directly applied to the route computation, and the failed node or link can be avoided. However, if the failure occurred within a loose-routed hop, the head-end node may not have enough information to reroute the LSP around the failure. Crankback signaling (see [CRANK]) and route exclusion techniques (see [RFC4874]) MAY be used in this case. The alternate route MAY be either computed on demand (that is, when the failure occurs; this is referred to as full LSP rerouting) or pre-computed and stored for use when the failure is reported. The latter offers faster restoration time. There is, however, a risk that the alternate route will become out of date through other changes in the network; this can be mitigated to some extent by periodic recalculation of idle alternate routes. (Full) LSP rerouting will be initiated by the head-end node that has either detected the LSP failure or received a Notify message and/or a PathErr message with the new error code/sub-code "Notify Error/LSP Locally Failed" for this LSP. The new LSP resources can be established using the make-before-break mechanism, where the new LSP is set up before the old LSP is torn down. This is done by using the mechanisms of the SESSION_ATTRIBUTE

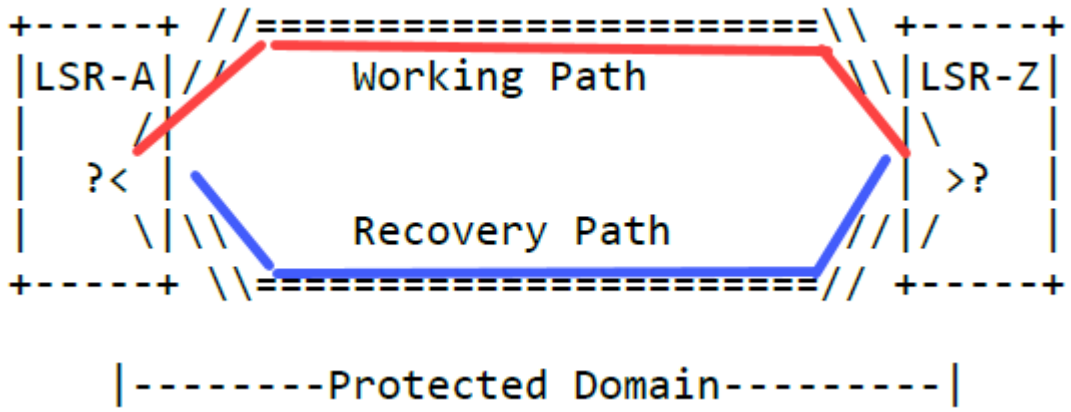
No.	'821 Patent Claim 6	IETF MPLS-TP System
		object and the Shared-Explicit (SE) reservation style (see [RFC3209]). Both the new and old LSPs can share resources at common nodes. Note that the make-before-break mechanism is not used to avoid disruption to the normal traffic flow (the latter has already been broken by the failure that is being repaired). However, it is valuable to retain the resources allocated on the original LSP that will be reused by the new alternate LSP." RFC 4872 at 25. "The procedure is as follows: 34) The initiating (source) node sends the normal traffic onto both the working and the protecting LSPs. Once completed, the source node sends reliably a Notify message to the destination with the new error code/sub-code "Notify Error/LSP Recovered" (Switchback Request). This Notify message includes the MESSAGE_ID object. The ACK_Desired flag MUST be set in this object to request the receiver to send an acknowledgment for the message (see [RFC2961]). 35) Upon receipt of this message, the destination selects the traffic from the working LSP. At the same time, it transmits the traffic onto both the working and protecting LSP. The destination then sends reliably a Notify message to the source confirming the completion of the operation. This message includes the MESSAGE_ID_ACK object to acknowledge reception of the received Notify message. This Notify message also includes the MESSAGE_ID object. The ACK_Desired flag MUST be set in this object to request the receiver to send an acknowledgment for the message (see [RFC2961]). 36) When the source node receives this Notify message, it switches to receive traffic from the working LSP. The source node then sends an Ack message to the destination node confirming that the LSP has been reverted." RFC 4872 at 27. For example, MPLS-TP Protection discloses an FPath field, that determines whether or not a specified LSP and/or pseudo-wire, i.e. network entity, is operational. At least under Orckit's Infringement Disclosures, this limitation is met. To the extent that the IETF MPLS-TP System is found to not meet this limitation, wherein said overall cost function comprises a predefined

No.	'821 Patent Claim 6	IETF MPLS-TP System
		entity cost metric would have been obvious to a person having ordinary skill in the art, as explained below.  Figure 3: Format of the PSC control packet MPLS-TP Protection Figure 3, annotation added. "FPath: used to indicate the path that is reporting a failure condition, the possible values are described in section 4.2.3" MPLS-TP Protection at 11. "4.2.3. Path fault identifier (FPath) The Fpath field of the PSC control SHALL be used only in a Signal fault (0110) or Signal degrade (0101) control packet. Its value indicates on which path the signal anomaly was detected. The following are the possible values: • 0: indicates that the fault condition is on the Recovery path • 1: indicates that the fault condition is on the Working path • 2-255: for future extensions." MPLS-TP Protection at 12. "Principles of Operation In all of the following sub-sections, assume a protected domain between LSR-A and LSR-Z, using paths W (working) and R (recovery) as shown in figure 4." MPLS-TP Protection at 13.

No.	'821 Patent Claim 6	IETF MPLS-TP System
		 MPLS-TP Protection Figure 4 (annotation added).

No.	'821 Patent Claim 7	IETF MPLS-TP System
7	The method of claim 6, wherein said predefined entity cost metric is selected from the group consisting of interior gateway protocol (IGP) and traffic engineering (TE).	IETF MPLS-TP System discloses the method of claim 6, wherein said predefined entity cost metric is selected from the group consisting of interior gateway protocol (IGP) and traffic engineering (TE). For example, RFC 4872 discloses that the SRLG identifiers can be determined. A person of ordinary skill in the art would understand that an IGP-TE database includes metrics consisting of Interior Gateway Protocol (IGP) and Traffic Engineering (TE). "Note: an intermediate node processing a PPRO can derive SRLG identifiers from the local IGP-TE database using its Type 1, 2, or 4 subobject values as pointers to the corresponding TE Links (assuming each of them has an associated SRLG TE attribute)." RFC 4872 at 35. "Abstract This document describes protocol-specific procedures and extensions for Generalized Multi-Protocol Label Switching (GMPLS) Resource Reservation Protocol - Traffic Engineering (RSVP-TE) signaling to support end-to-end Label Switched Path (LSP) recovery that denotes

No.	'821 Patent Claim 7	IETF MPLS-TP System
		protection and restoration. A generic functional description of GMPLS recovery can be found in a companion document, RFC 4426." RFC 4872 at 1. "A functional description of GMPLS recovery is provided in [RFC4426] and should be considered as a companion document. The present document describes the protocol-specific procedures for GMPLS RSVP- TE (Resource Reservation Protocol - Traffic Engineering) signaling (see [RFC3473]) to support end-to-end recovery. End-to-end recovery refers to the recovery of an entire LSP from its head-end (ingress node endpoint) to its tail-end (egress node endpoint). With end-to- end recovery, working LSPs are assumed to be resource-disjoint (where a resource is a link, node, or Shared Risk Link Group (SRLG)) in the network so that they do not share any failure probability, but this is not mandatory. With respect to a given set of network resources, a pair of working/protecting LSPs SHOULD be resource disjoint in case of dedicated recovery type (see below). On the other hand, in case of shared recovery (see below), a group of working LSPs SHOULD be mutually resource-disjoint in order to allow for a (single and commonly) shared protecting LSP, itself resource-disjoint from each of the working LSPs. Note that resource disjointness is a necessary (but not sufficient) condition to ensure LSP recoverability." RFC 4872 at 3. For example, MPLS-TP Protection discloses an FPath field, that determines whether or not a specified LSP and/or pseudo-wire, i.e. network entity, is operational. At least under Orckit's Infringement Disclosures, this limitation is met. To the extent that the IETF MPLS-TP System is found to not meet this limitation, wherein said predefined entity cost metric is selected from the group consisting of interior gateway protocol (IGP) and traffic engineering (TE) would have been obvious to a person having ordinary skill in the art, as explained below.  0 1 2 3 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 +-+ |Ver|Request| PT| FPath | Path | Reserved | +-+ Figure 3: Format of the PSC control packet

No.	'821 Patent Claim 7	IETF MPLS-TP System
		MPLS-TP Protection Figure 3, annotation added. "FPath: used to indicate the path that is reporting a failure condition, the possible values are described in section 4.2.3" MPLS-TP Protection at 11. "4.2.3. Path fault identifier (FPath) The Fpath field of the PSC control SHALL be used only in a Signal fault (0110) or Signal degrade (0101) control packet. Its value indicates on which path the signal anomaly was detected. The following are the possible values: • 0: indicates that the fault condition is on the Recovery path • 1: indicates that the fault condition is on the Working path • 2-255: for future extensions." MPLS-TP Protection at 12. "Principles of Operation In all of the following sub-sections, assume a protected domain between LSR-A and LSR-Z, using paths W (working) and R (recovery) as shown in figure 4." MPLS-TP Protection at 13.  MPLS-TP Protection Figure 4 (annotation added).

No.	'821 Patent Claim 8	IETF MPLS-TP System
8	The method of claim 4, further comprising the step of configuring said working entity as revertive.	IETF MPLS-TP System discloses the method of claim 4, further comprising the step of configuring said working entity as revertive. For example, RFC 4872 discloses reversion signaling, where traffic reverts back to a working LSP after it has recovered from a failure. RFC 4872 further discloses that the reversion signaling can be configured on a LSP, i.e. an entity, by following a series of steps. "12. Reversion Reversion refers to a recovery switching operation, where the normal traffic returns to (or remains on) the working LSP when it has recovered from the failure. Reversion implies that resources remain allocated to the LSP that was originally routed over them even after a failure. It is important to have mechanisms that allow reversion to be performed with minimal service disruption and reconfiguration. For "1+1 bidirectional Protection", reversion to the recovered LSP occurs by using the following sequence: 1. Clear the A bit of the ADMIN_STATUS object if set for the recovered LSP. 2. Then, apply the method described below to switch normal traffic back from the protecting to the recovered LSP. This is performed by using the new error code/sub-code "Notify Error/LSP Recovered" (Switchback Request). The procedure is as follows: 1) The initiating (source) node sends the normal traffic onto both the working and the protecting LSPs. Once completed, the source node sends reliably a Notify message to the destination with the new error code/sub-code "Notify Error/LSP Recovered" (Switchback Request). This Notify message includes the MESSAGE_ID object. The ACK_Desired flag MUST be set in this object to request the receiver to send an acknowledgment for the message (see [RFC2961]). 2) Upon receipt of this message, the destination selects the traffic from the working LSP. At the same time, it transmits the traffic onto both the working and protecting LSP. The destination then sends reliably a Notify message to the source confirming the completion of the operation. This message includes the MESSAGE_ID_ACK object to acknowledge

No.	'821 Patent Claim 8	IETF MPLS-TP System
		reception of the received Notify message. This Notify message also includes the MESSAGE_ID object. The ACK_Desired flag MUST be set in this object to request the receiver to send an acknowledgment for the message (see [RFC2961]). 3) When the source node receives this Notify message, it switches to receive traffic from the working LSP. The source node then sends an Ack message to the destination node confirming that the LSP has been reverted. 3. Finally, clear the O bit of the PROTECTION object sent over the protecting LSP. For "1:N Protection with Extra-traffic", reversion to the recovered LSP occurs by using the following sequence: 1. Clear the A bit of the ADMIN_STATUS object if set for the recovered LSP. 2. Then, apply the method described below to switch normal traffic back from the protecting to the recovered LSP. This is performed by using the new error code/sub-code "Notify Error/LSP Recovered" (Switchback Request). The procedure is as follows: 1) The initiating (source) node sends the normal traffic onto both the working and the protecting LSPs. Once completed, the source node sends reliably a Notify message to the destination with the new error code/sub-code "Notify Error/LSP Recovered" (Switchback Request). This Notify message includes the MESSAGE_ID object. The ACK_Desired flag MUST be set in this object to request the receiver to send an acknowledgment for the message (see [RFC2961]). 2) Upon receipt of this message, the destination selects the traffic from the working LSP. At the same time, it transmits the traffic onto both the working and protecting LSP. The destination then sends reliably a Notify message to the source confirming the completion of the operation. This message includes the MESSAGE_ID_ACK object to acknowledge reception of the received Notify message. This Notify message also includes the MESSAGE_ID object. The ACK_Desired flag MUST be set in this object to request the receiver to send an acknowledgment for the message (see [RFC2961]). 3) When the source node receives this Notify message, it switches to receive traffic from the working LSP, and stops transmitting traffic on the protecting LSP. The source node then sends an Ack message to the destination node confirming that the LSP has been reverted.

No.	'821 Patent Claim 8	IETF MPLS-TP System
		4) Upon receipt of this message, the destination node stops transmitting traffic along the protecting LSP. 3. Finally, clear the O bit of the PROTECTION object sent over the protecting LSP. For "Rerouting without Extra-traffic" (including the shared recovery case), reversion implies that the formerly working LSP has not been torn down by the head-end node upon PathErr message reception, i.e., the head-end node kept refreshing the working LSP under failure condition. This ensures that the exact same resources are retrieved after reversion switching (except if the working LSP required re- signaling). Re-activation is performed using the following sequence: 1. Clear the A bit of the ADMIN_STATUS object if set for the recovered LSP. 2. Then, apply the method described below to switch normal traffic back from the protecting to the recovered LSP. This is performed by using the new error code/sub-code "Notify Error/LSP Recovered" (Switchback Request). The procedure is as follows: 1) The initiating (source) node sends the normal traffic onto both the working and the protecting LSPs. Once completed, the source node sends reliably a Notify message to the destination with the new error code/sub-code "Notify Error/LSP Recovered" (Switchback Request). This Notify message includes the MESSAGE_ID object. The ACK_Desired flag MUST be set in this object to request the receiver to send an acknowledgment for the message (see [RFC2961]). 2) Upon receipt of this message, the destination selects the traffic from the working LSP. At the same time, it transmits the traffic onto both the working and protecting LSP. The destination then sends reliably a Notify message to the source confirming the completion of the operation. This message includes the MESSAGE_ID_ACK object to acknowledge reception of the received Notify message. This Notify message also includes the MESSAGE_ID object. The ACK_Desired flag MUST be set in this object to request the receiver to send an acknowledgment for the message (see [RFC2961]). 3) When the source node receives this Notify message, it switches to receive traffic from the working LSP, and stops transmitting traffic on the protecting LSP. The source node then sends an Ack message to the destination node confirming that the LSP has been reverted. 4) Upon receipt of this message, the destination node stops transmitting traffic along the protecting LSP. 3. Finally, de-activate the protecting LSP by setting the S bit to 1 in the PROTECTION object sent over the protecting LSP." RFC 4872 at 26–29.

No.	'821 Patent Claim 8	IETF MPLS-TP System
		"Note: this document assumes that Protection Type values have end-to- end significance and that the same value is sent over the protected and the protecting path. In this context, shared-mesh (for instance) appears from the end-nodes perspective as being simply an LSP rerouting without extra-traffic services. The net result of this is that a single bit (the S bit alone) does not allow determining whether resource allocation should be performed with respect to the status of the LSP within the protected entity. The introduction of the P bit solves this problem unambiguously. These bits MUST be processed on a hop-by-hop basis (independently of the LSP Protection Type context). This allows for an easier implementation of reversion signaling (see Section 12) but also facilitates the transparent delivery of protected services since any intermediate node is not required to know the semantics associated with the incoming LSP Protection Type value." RFC 4872 at 8–9. For example, MPLS-TP Protection discloses that the protection switching method can be done both with a revertive and non-revertive network. "4.3.6. Recovery from switching When the condition that triggered the protection switching clears, e.g. the cause of the failure condition has been corrected, or the operator clears a Manual Switch, then the protection domain SHOULD follow the following procedures: • If the network is configured for non-revertive behaviour, then the two LSRs SHOULD transmit DNR (Request code 0010) messages. • If the network is recovering from an operator switching command (in revertive mode), then both LSRs SHOULD return to using the working transport path and transmit No request (Request code 0000) messages. • If the network is recovering from a failure or degraded condition (in revertive mode), then the LSR that detects this recovery SHALL activate a local Wait-to-restore (WTR) timer (see section 4.3.6.1) to verify that there is not an intermittent failure. After the WTR expires, the LSR SHOULD return to using the working transport path and transmit No request (Request code 0000) messages.

No.	'821 Patent Claim 8	IETF MPLS-TP System
		4.3.6.1. Wait-to-restore timer In revertive mode, in order to prevent frequent activation of protection switching due to an intermittent defect, the working transport path must become stable and fault-free before reverting to the normal condition. In order to verify that this is the case a fixed period of time must elapse before the normal traffic uses the working transport path. This period, called the Wait-to-restore (WTR) period, should be configurable by the operator in 1-minute intervals within the range 1-12 minutes. The default value is 5 minutes. During this period, if a failure condition is detected on the working transport path, then the WTR timer is cleared and the normal traffic SHALL continue to be transported over the recovery transport path. If the WTR timer expires without being pre-empted by a failure, then the traffic SHOULD be returned to use the working transport path (as above)." MPLS-TP at 17–19.

No.	'821 Patent Claim 9	IETF MPLS-TP System
9[preamble]	The method of claim 4, wherein said overall cost function comprises:	IETF MPLS-TP System discloses the method of claim 4, wherein said overall cost function comprises. *See* Claim 4.
9[a]	selecting a subset of entity pairs wherein each entity pair of said subset has substantially minimum probability of a concurrent failure of said protection entity and said working entity; and	IETF MPLS-TP System discloses selecting a subset of entity pairs wherein each entity pair of said subset has substantially minimum probability of a concurrent failure of said protection entity and said working entity. For example, RFC 4872 discloses that the plurality of MPLS LSPs, i.e., transport entities, are selected in pairs of either working and protection LSPs, or working and restoration LSPs. RFC 4872 further discloses that the plurality of MPLS LSPs, i.e., transport entities, are selected to necessarily be "resource disjoint," i.e., where multiple resources, such as links, nodes, or Shared Risk Link Groups (SRLGs) do not share common links or failure probabilities. Thus, at least under the apparent claim scope alleged by Orckit's Infringement Disclosures, this limitation is met. To the extent that the IETF MPLS-TP System is found to

No.	'821 Patent Claim 9	IETF MPLS-TP System
		not meet this limitation, selecting a subset of entity pairs wherein each entity pair of said subset has substantially minimum probability of a concurrent failure of said protection entity and said working entity would have been obvious to a person having ordinary skill in the art, as explained below. "When a failure is detected by one or both endpoints of the LSP, both endpoints must select traffic from the other LSP. This action must be coordinated between node A and D. From this perspective, 1+1 bidirectional protection can be seen as a coordinated protection- switching mechanism between both endpoints. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 11. "1:1 Protection with Extra-Traffic The most common case of end-to-end 1:N protection is to establish, between the same endpoints, an end-to-end working LSP (thus, N = 1) and a dedicated end-to-end protecting LSP that are mutually link/ node/SRLG disjoint. This protects against working LSP failure(s)." RFC 4872 at 13. "An approach to reduce recovery resource requirements is to have protection LSPs sharing network resources when the working LSPs that they protect are physically (i.e., link, node, SRLG, etc.) disjoint. This mechanism is referred to as shared mesh restoration and is described in [RFC4426]. Shared-mesh restoration can be seen as a particular case of pre-planned LSP rerouting (see Section 8) that reduces the recovery resource requirements by allowing multiple protecting LSPs to share common link and node resources. Here also, the recovery resources for the protecting LSPs are pre-reserved during the provisioning phase, thus an explicit signaling action is required to activate (i.e., commit resource allocation at the data plane) a specific protecting LSP instantiated during the (pre-) provisioning phase. This requires restoration signaling along the protecting LSP." RFC 4872 at 20. "1. Introduction

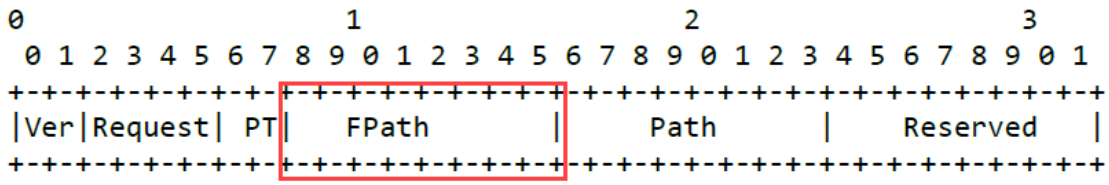
No.	'821 Patent Claim 9	IETF MPLS-TP System
		Generalized Multi-Protocol Label Switching (GMPLS) extends MPLS to include support for Layer-2 Switch Capable (L2SC), Time-Division Multiplex (TDM), Lambda Switch Capable (LSC), and Fiber Switch Capable (FSC) interfaces. GMPLS recovery uses control plane mechanisms (i.e., signaling, routing, and link management mechanisms) to support data plane fault recovery. Note that the analogous (data plane) fault detection mechanisms are required to be present in support of the control plane mechanisms. In this document, the term "recovery" is generically used to denote both protection and restoration; the specific terms "protection" and "restoration" are only used when differentiation is required. The subtle distinction between protection and restoration is made based on the resource allocation done during the recovery phase (see [RFC4427]). A functional description of GMPLS recovery is provided in [RFC4426] and should be considered as a companion document. The present document describes the protocol-specific procedures for GMPLS RSVP- TE (Resource ReSerVation Protocol - Traffic Engineering) signaling (see [RFC3473]) to support end-to-end recovery. End-to-end recovery refers to the recovery of an entire LSP from its head-end (ingress node endpoint) to its tail-end (egress node endpoint). With end-to- end recovery, working LSPs are assumed to be resource-disjoint (where a resource is a link, node, or Shared Risk Link Group (SRLG)) in the network so that they do not share any failure probability, but this is not mandatory. With respect to a given set of network resources, a pair of working/protecting LSPs SHOULD be resource disjoint in case of dedicated recovery type (see below). On the other hand, in case of shared recovery (see below), a group of working LSPs SHOULD be mutually resource-disjoint in order to allow for a (single and commonly) shared protecting LSP, itself resource-disjoint from each of the working LSPs. Note that resource disjointness is a necessary (but not sufficient) condition to ensure LSP recoverability. The present document addresses four types of end-to-end LSP recovery: 1) 1+1 (unidirectional/bidirectional) protection, 2) 1:N (N >= 1) LSP protection with extra-traffic, 3) pre-planned LSP rerouting without extra-traffic (including shared mesh), and 4) full LSP rerouting. 45) The simplest notion of end-to-end LSP protection is 1+1 unidirectional protection. Using this type of protection, a protecting LSP is signaled over a dedicated resource-disjoint alternate path to protect an associated working LSP. Normal traffic is

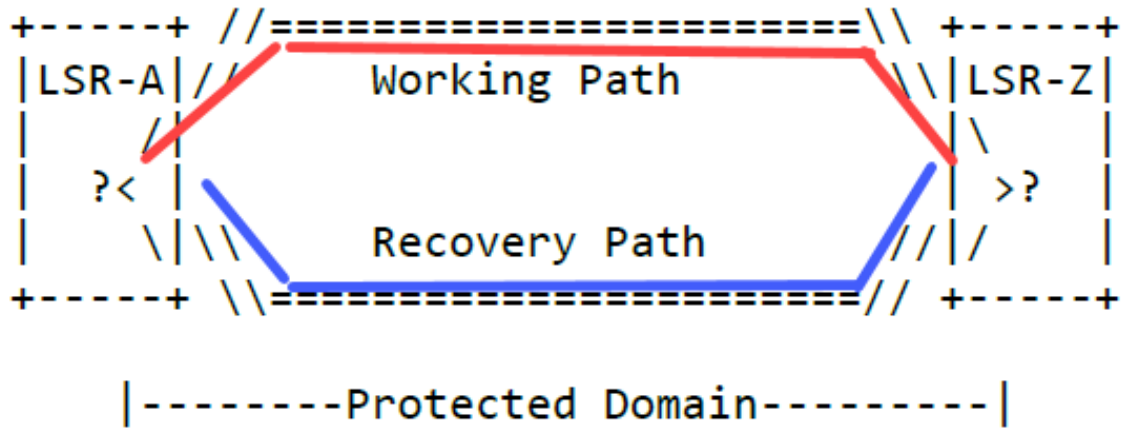
No.	'821 Patent Claim 9	IETF MPLS-TP System
		simultaneously sent on both LSPs and a selector is used at the egress node to receive traffic from one of the LSPs. If a failure occurs along one of the LSPs, the egress node selects the traffic from the valid LSP. No coordination is required between the end nodes when a failure/switchover occurs. In 1+1 bidirectional protection, a protecting LSP is signaled over a dedicated resource-disjoint alternate path to protect the working LSP. Normal traffic is simultaneously sent on both LSPs (in both directions), and a selector is used at both ingress/egress nodes to receive traffic from the same LSP. This requires coordination between the end-nodes when switching to the protecting LSP. 46) In 1:N (N >= 1) protection with extra-traffic, the protecting LSP is a fully provisioned and resource-disjoint LSP from the N working LSPs, that allows for carrying extra-traffic. The N working LSPs MAY be mutually resource-disjoint. Coordination between end-nodes is required when switching from one of the working LSPs to the protecting LSP. As the protecting LSP is fully provisioned, default operations during protection switching are specified for a protecting LSP carrying extra-traffic, but this is not mandatory. Note that M:N protection is out of scope of this document (though mechanisms it defines may be extended to cover it). 47) Pre-planned LSP rerouting (or restoration) relies on the establishment between the same pair of end-nodes of a working LSP and a protecting LSP that is link/node/SRLG disjoint from the working one. Here, the recovery resources for the protecting LSP are pre-reserved but explicit action is required to activate (i.e., commit resource allocation at the data plane) a specific protecting LSP instantiated during the (pre-)provisioning phase. Since the protecting LSP is not "active" (i.e., fully instantiated), it cannot carry any extra-traffic. This does not mean that the corresponding resources cannot be used by other LSPs. Therefore, this mechanism protects against working LSP(s) failure(s) but requires activation of the protecting LSP after working LSP failure occurrence. This requires restoration signaling along the protecting path. "Shared-mesh" restoration can be seen as a particular case of pre-planned LSP rerouting that reduces the recovery resource requirements by allowing multiple protecting LSPs to share common link and node resources. The recovery resources are pre-reserved but explicit action is required to activate (i.e., commit resource allocation at the data plane) a specific protecting LSP instantiated during the (pre-)

No.	'821 Patent Claim 9	IETF MPLS-TP System
		provisioning phase. This procedure requires restoration signaling along the protecting path. Note that in both cases, bandwidth pre-reserved for a protecting (but not activated) LSP can be made available for carrying extra traffic. LSPs for extra-traffic (with lower holding priority than the protecting LSP) can then be established using the bandwidth pre-reserved for the protecting LSP. Also, any lower priority LSP that use the pre-reserved resources for the protecting LSP(s) must be preempted during the activation of the protecting LSP. 48) Full LSP rerouting (or restoration) switches normal traffic to an alternate LSP that is not even partially established until after the working LSP failure occurs. The new alternate route is selected at the LSP head-end node, it may reuse resources of the failed LSP at intermediate nodes and may include additional intermediate nodes and/or links." " RFC 4872 at 3–5. "Consider the following network topology:  The paths [A,B,C,D] and [A,E,F,G,D] are node and link disjoint, ignoring the ingress/egress nodes A and D. A 1+1 protected path is established from A to D over [A,B,C,D] and [A,E,F,G,D], and traffic is transmitted simultaneously over both component paths (i.e., LSPs). During the provisioning phase, both LSPs are fully instantiated (and thus activated) so that no resource sharing can be done along the protecting LSP (nor can any extra-traffic be transported). It is also RECOMMENDED to set the N bit since no protection-switching signaling is assumed in this case. When a failure occurs (say, at node B) and is detected at end-node D, the receiver at D selects the normal traffic from the other LSP. From this perspective, 1+1 unidirectional protection can be seen as an uncoordinated protection-switching mechanism acting independently at both endpoints. Also, for the LSP under failure condition, it is RECOMMENDED to not set

No.	'821 Patent Claim 9	IETF MPLS-TP System
		the Path_State_Removed Flag of the ERROR_SPEC object (see [RFC3473]) upon PathErr message generation. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 9. "When a failure is detected by one or both endpoints of the LSP, both endpoints must select traffic from the other LSP. This action must be coordinated between node A and D. From this perspective, 1+1 bidirectional protection can be seen as a coordinated protection- switching mechanism between both endpoints. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 11. "(Full) LSP Rerouting LSP rerouting, on the other hand, switches normal traffic to an alternate LSP that is fully established only after failure occurrence. The new (alternate) route is selected at the LSP head- end and may reuse intermediate nodes included in the original route; it may also include additional intermediate nodes. For strict-hop routing, TE requirements can be directly applied to the route computation, and the failed node or link can be avoided. However, if the failure occurred within a loose-routed hop, the head-end node may not have enough information to reroute the LSP around the failure. Crankback signaling (see [CRANK]) and route exclusion techniques (see [RFC4874]) MAY be used in this case. The alternate route MAY be either computed on demand (that is, when the failure occurs; this is referred to as full LSP rerouting) or pre-computed and stored for use when the failure is reported. The latter offers faster restoration time. There is, however, a risk that the alternate route will become out of date through other changes in the network; this can be mitigated to some extent by periodic recalculation of idle alternate routes. (Full) LSP rerouting will be initiated by the head-end node that has either detected the LSP failure or received a Notify message and/or a PathErr message with the new error code/sub-code "Notify Error/LSP Locally Failed" for this LSP. The new LSP resources can be established using the make-before-break mechanism, where the new LSP is set up before the

No.	'821 Patent Claim 9	IETF MPLS-TP System
		old LSP is torn down. This is done by using the mechanisms of the SESSION_ATTRIBUTE object and the Shared-Explicit (SE) reservation style (see [RFC3209]). Both the new and old LSPs can share resources at common nodes. Note that the make-before-break mechanism is not used to avoid disruption to the normal traffic flow (the latter has already been broken by the failure that is being repaired). However, it is valuable to retain the resources allocated on the original LSP that will be reused by the new alternate LSP." RFC 4872 at 25. "The procedure is as follows: 37) The initiating (source) node sends the normal traffic onto both the working and the protecting LSPs. Once completed, the source node sends reliably a Notify message to the destination with the new error code/sub-code "Notify Error/LSP Recovered" (Switchback Request). This Notify message includes the MESSAGE_ID object. The ACK_Desired flag MUST be set in this object to request the receiver to send an acknowledgment for the message (see [RFC2961]). 38) Upon receipt of this message, the destination selects the traffic from the working LSP. At the same time, it transmits the traffic onto both the working and protecting LSP. The destination then sends reliably a Notify message to the source confirming the completion of the operation. This message includes the MESSAGE_ID_ACK object to acknowledge reception of the received Notify message. This Notify message also includes the MESSAGE_ID object. The ACK_Desired flag MUST be set in this object to request the receiver to send an acknowledgment for the message (see [RFC2961]). 39) When the source node receives this Notify message, it switches to receive traffic from the working LSP. The source node then sends an Ack message to the destination node confirming that the LSP has been reverted." RFC 4872 at 27. For example, MPLS-TP Protection discloses an FPath field, that determines whether or not a specified LSP and/or pseudo-wire, i.e. network entity, is operational. At least under Orckit's Infringement Disclosures, this limitation is met. To the extent that the IETF MPLS-TP System

No.	'821 Patent Claim 9	IETF MPLS-TP System
		is found to not meet this limitation, selecting a subset of entity pairs wherein each entity pair of said subset has substantially minimum probability of a concurrent failure of said protection entity and said working entity would have been obvious to a person having ordinary skill in the art, as explained below.  Figure 3: Format of the PSC control packet MPLS-TP Protection Figure 3, annotation added. "FPath: used to indicate the path that is reporting a failure condition, the possible values are described in section 4.2.3" MPLS-TP Protection at 11. "4.2.3. Path fault identifier (FPath) The Fpath field of the PSC control SHALL be used only in a Signal fault (0110) or Signal degrade (0101) control packet. Its value indicates on which path the signal anomaly was detected. The following are the possible values: • 0: indicates that the fault condition is on the Recovery path • 1: indicates that the fault condition is on the Working path • 2-255: for future extensions." MPLS-TP Protection at 12. "Principles of Operation In all of the following sub-sections, assume a protected domain between LSR-A and LSR-Z, using paths W (working) and R (recovery) as shown in figure 4." MPLS-TP Protection at 13.

No.	'821 Patent Claim 9	IETF MPLS-TP System
		 MPLS-TP Protection Figure 4 (annotation added).

No.	'821 Patent Claim 9	IETF MPLS-TP System
9[b]	if said subset comprises at least two entity pairs, selecting an entity pair from said subset that minimizes an entity cost function.	IETF MPLS-TP System discloses if said subset comprises at least two entity pairs, selecting an entity pair from said subset that minimizes an entity cost function. For example, RFC 4872 discloses that the plurality of MPLS LSPs, i.e., transport entities, are selected in pairs of either working and protection LSPs, or working and restoration LSPs. RFC 4872 further discloses that the plurality of MPLS LSPs, i.e., transport entities, are selected to necessarily be "resource disjoint," i.e., where multiple resources, such as links, nodes, or Shared Risk Link Groups (SRLGs) do not share common links or failure probabilities. Thus, at least under the apparent claim scope alleged by Orckit's Infringement Disclosures, this limitation is met. To the extent that the IETF MPLS-TP System is found to not meet this limitation, if said subset comprises at least two entity pairs, selecting an entity pair from said subset that minimizes an entity cost function would have been obvious to a person having ordinary skill in the art, as explained below. "When a failure is detected by one or both endpoints of the LSP, both endpoints must select traffic from the other LSP. This action must be coordinated between node A and D. From this perspective, 1+1 bidirectional protection can be seen as a coordinated protection- switching mechanism between both endpoints. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 11. "1:1 Protection with Extra-Traffic The most common case of end-to-end 1:N protection is to establish, between the same endpoints, an end-to-end working LSP (thus, N = 1) and a dedicated end-to-end protecting LSP that are mutually link/ node/SRLG disjoint. This protects against working LSP failure(s)." RFC 4872 at 13. "An approach to reduce recovery resource requirements is to have protection LSPs sharing network resources when the working LSPs that they protect are physically (i.e., link, node, SRLG, etc.) disjoint. This mechanism is referred to as shared mesh restoration and is described in [RFC4426]. Shared-mesh restoration can be seen as a particular case of pre-

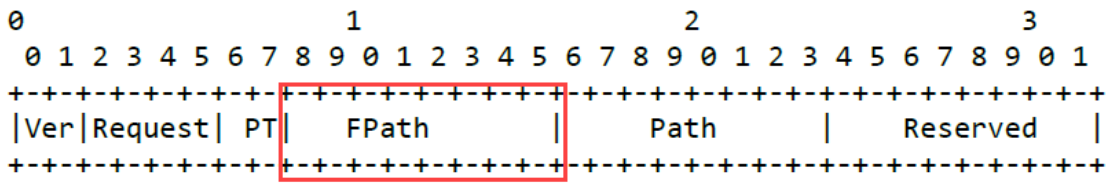
No.	'821 Patent Claim 9	IETF MPLS-TP System
		planned LSP rerouting (see Section 8) that reduces the recovery resource requirements by allowing multiple protecting LSPs to share common link and node resources. Here also, the recovery resources for the protecting LSPs are pre-reserved during the provisioning phase, thus an explicit signaling action is required to activate (i.e., commit resource allocation at the data plane) a specific protecting LSP instantiated during the (pre-) provisioning phase. This requires restoration signaling along the protecting LSP." RFC 4872 at 20. "1. Introduction Generalized Multi-Protocol Label Switching (GMPLS) extends MPLS to include support for Layer-2 Switch Capable (L2SC), Time-Division Multiplex (TDM), Lambda Switch Capable (LSC), and Fiber Switch Capable (FSC) interfaces. GMPLS recovery uses control plane mechanisms (i.e., signaling, routing, and link management mechanisms) to support data plane fault recovery. Note that the analogous (data plane) fault detection mechanisms are required to be present in support of the control plane mechanisms. In this document, the term "recovery" is generically used to denote both protection and restoration; the specific terms "protection" and "restoration" are only used when differentiation is required. The subtle distinction between protection and restoration is made based on the resource allocation done during the recovery phase (see [RFC4427]). A functional description of GMPLS recovery is provided in [RFC4426] and should be considered as a companion document. The present document describes the protocol-specific procedures for GMPLS RSVP- TE (Resource ReSerVation Protocol - Traffic Engineering) signaling (see [RFC3473]) to support end-to-end recovery. End-to-end recovery refers to the recovery of an entire LSP from its head-end (ingress node endpoint) to its tail-end (egress node endpoint). With end-to- end recovery, working LSPs are assumed to be resource-disjoint (where a resource is a link, node, or Shared Risk Link Group (SRLG)) in the network so that they do not share any failure probability, but this is not mandatory. With respect to a given set of network resources, a pair of working/protecting LSPs SHOULD be resource disjoint in case of dedicated recovery type (see below). On the other hand, in case of shared recovery (see below), a group of working LSPs SHOULD be mutually resource-disjoint in order to allow for a (single and commonly) shared protecting LSP, itself resource-disjoint from each

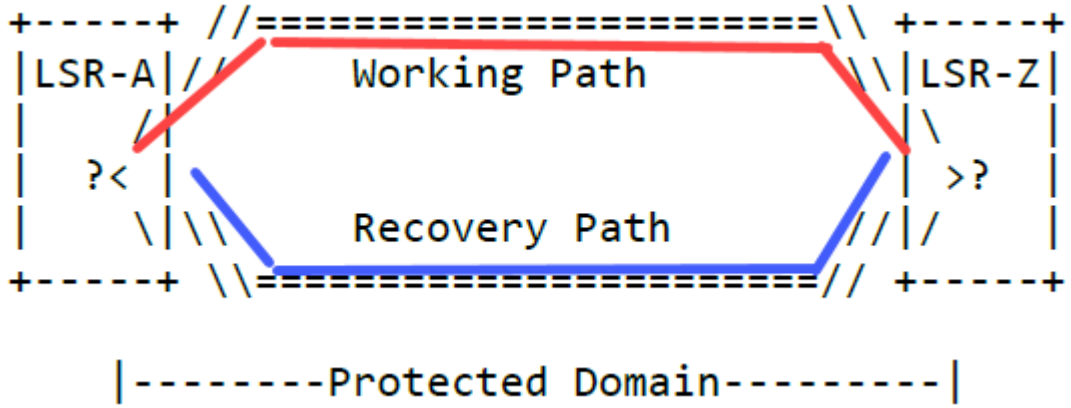
No.	'821 Patent Claim 9	IETF MPLS-TP System
		of the working LSPs. Note that resource disjointness is a necessary (but not sufficient) condition to ensure LSP recoverability. The present document addresses four types of end-to-end LSP recovery: 1) 1+1 (unidirectional/bidirectional) protection, 2) 1:N (N >= 1) LSP protection with extra-traffic, 3) pre-planned LSP rerouting without extra-traffic (including shared mesh), and 4) full LSP rerouting. 49) The simplest notion of end-to-end LSP protection is 1+1 unidirectional protection. Using this type of protection, a protecting LSP is signaled over a dedicated resource-disjoint alternate path to protect an associated working LSP. Normal traffic is simultaneously sent on both LSPs and a selector is used at the egress node to receive traffic from one of the LSPs. If a failure occurs along one of the LSPs, the egress node selects the traffic from the valid LSP. No coordination is required between the end nodes when a failure/switchover occurs. In 1+1 bidirectional protection, a protecting LSP is signaled over a dedicated resource-disjoint alternate path to protect the working LSP. Normal traffic is simultaneously sent on both LSPs (in both directions), and a selector is used at both ingress/egress nodes to receive traffic from the same LSP. This requires coordination between the end-nodes when switching to the protecting LSP. 50) In 1:N (N >= 1) protection with extra-traffic, the protecting LSP is a fully provisioned and resource-disjoint LSP from the N working LSPs, that allows for carrying extra-traffic. The N working LSPs MAY be mutually resource-disjoint. Coordination between end-nodes is required when switching from one of the working LSPs to the protecting LSP. As the protecting LSP is fully provisioned, default operations during protection switching are specified for a protecting LSP carrying extra-traffic, but this is not mandatory. Note that M:N protection is out of scope of this document (though mechanisms it defines may be extended to cover it). 51) Pre-planned LSP rerouting (or restoration) relies on the establishment between the same pair of end-nodes of a working LSP and a protecting LSP that is link/node/SRLG disjoint from the working one. Here, the recovery resources for the protecting LSP are pre-reserved but explicit action is required to activate (i.e., commit resource allocation at the data plane) a specific protecting LSP instantiated during the (pre-

No.	'821 Patent Claim 9	IETF MPLS-TP System
		)provisioning phase. Since the protecting LSP is not "active" (i.e., fully instantiated), it cannot carry any extra-traffic. This does not mean that the corresponding resources cannot be used by other LSPs. Therefore, this mechanism protects against working LSP(s) failure(s) but requires activation of the protecting LSP after working LSP failure occurrence. This requires restoration signaling along the protecting path. "Shared-mesh" restoration can be seen as a particular case of pre-planned LSP rerouting that reduces the recovery resource requirements by allowing multiple protecting LSPs to share common link and node resources. The recovery resources are pre-reserved but explicit action is required to activate (i.e., commit resource allocation at the data plane) a specific protecting LSP instantiated during the (pre-) provisioning phase. This procedure requires restoration signaling along the protecting path. Note that in both cases, bandwidth pre-reserved for a protecting (but not activated) LSP can be made available for carrying extra traffic. LSPs for extra-traffic (with lower holding priority than the protecting LSP) can then be established using the bandwidth pre-reserved for the protecting LSP. Also, any lower priority LSP that use the pre-reserved resources for the protecting LSP(s) must be preempted during the activation of the protecting LSP. 52) Full LSP rerouting (or restoration) switches normal traffic to an alternate LSP that is not even partially established until after the working LSP failure occurs. The new alternate route is selected at the LSP head-end node, it may reuse resources of the failed LSP at intermediate nodes and may include additional intermediate nodes and/or links." " RFC 4872 at 3–5. "Consider the following network topology:  The paths [A,B,C,D] and [A,E,F,G,D] are node and link disjoint, ignoring the ingress/egress nodes A and D. A 1+1 protected path is established from A to D over [A,B,C,D] and

No.	'821 Patent Claim 9	IETF MPLS-TP System
		[A,E,F,G,D], and traffic is transmitted simultaneously over both component paths (i.e., LSPs). During the provisioning phase, both LSPs are fully instantiated (and thus activated) so that no resource sharing can be done along the protecting LSP (nor can any extra-traffic be transported). It is also RECOMMENDED to set the N bit since no protection-switching signaling is assumed in this case. When a failure occurs (say, at node B) and is detected at end-node D, the receiver at D selects the normal traffic from the other LSP. From this perspective, 1+1 unidirectional protection can be seen as an uncoordinated protection-switching mechanism acting independently at both endpoints. Also, for the LSP under failure condition, it is RECOMMENDED to not set the Path_State_Removed Flag of the ERROR_SPEC object (see [RFC3473]) upon PathErr message generation. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 9. "When a failure is detected by one or both endpoints of the LSP, both endpoints must select traffic from the other LSP. This action must be coordinated between node A and D. From this perspective, 1+1 bidirectional protection can be seen as a coordinated protection- switching mechanism between both endpoints. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 11. "(Full) LSP Rerouting LSP rerouting, on the other hand, switches normal traffic to an alternate LSP that is fully established only after failure occurrence. The new (alternate) route is selected at the LSP head- end and may reuse intermediate nodes included in the original route; it may also include additional intermediate nodes. For strict-hop routing, TE requirements can be directly applied to the route computation, and the failed node or link can be avoided. However, if the failure occurred within a loose-routed hop, the head-end node may not have enough information to

No.	'821 Patent Claim 9	IETF MPLS-TP System
		reroute the LSP around the failure. Crankback signaling (see [CRANK]) and route exclusion techniques (see [RFC4874]) MAY be used in this case. The alternate route MAY be either computed on demand (that is, when the failure occurs; this is referred to as full LSP rerouting) or pre-computed and stored for use when the failure is reported. The latter offers faster restoration time. There is, however, a risk that the alternate route will become out of date through other changes in the network; this can be mitigated to some extent by periodic recalculation of idle alternate routes. (Full) LSP rerouting will be initiated by the head-end node that has either detected the LSP failure or received a Notify message and/or a PathErr message with the new error code/sub-code "Notify Error/LSP Locally Failed" for this LSP. The new LSP resources can be established using the make-before-break mechanism, where the new LSP is set up before the old LSP is torn down. This is done by using the mechanisms of the SESSION_ATTRIBUTE object and the Shared-Explicit (SE) reservation style (see [RFC3209]). Both the new and old LSPs can share resources at common nodes. Note that the make-before-break mechanism is not used to avoid disruption to the normal traffic flow (the latter has already been broken by the failure that is being repaired). However, it is valuable to retain the resources allocated on the original LSP that will be reused by the new alternate LSP." RFC 4872 at 25. "The procedure is as follows: 40) The initiating (source) node sends the normal traffic onto both the working and the protecting LSPs. Once completed, the source node sends reliably a Notify message to the destination with the new error code/sub-code "Notify Error/LSP Recovered" (Switchback Request). This Notify message includes the MESSAGE_ID object. The ACK_Desired flag MUST be set in this object to request the receiver to send an acknowledgment for the message (see [RFC2961]). 41) Upon receipt of this message, the destination selects the traffic from the working LSP. At the same time, it transmits the traffic onto both the working and protecting LSP. The destination then sends reliably a Notify message to the source confirming the completion of the operation. This message includes the MESSAGE_ID_ACK object to acknowledge reception of the received Notify message. This Notify message also

No.	'821 Patent Claim 9	IETF MPLS-TP System
		includes the MESSAGE_ID object. The ACK_Desired flag MUST be set in this object to request the receiver to send an acknowledgment for the message (see [RFC2961]). 42) When the source node receives this Notify message, it switches to receive traffic from the working LSP. The source node then sends an Ack message to the destination node confirming that the LSP has been reverted." RFC 4872 at 27. For example, MPLS-TP Protection discloses an FPath field, that determines whether or not a specified LSP and/or pseudo-wire, i.e. network entity, is operational. At least under Orckit's Infringement Disclosures, this limitation is met. To the extent that the IETF MPLS-TP System is found to not meet this limitation, if said subset comprises at least two entity pairs, selecting an entity pair from said subset that minimizes an entity cost function would have been obvious to a person having ordinary skill in the art, as explained below.  Figure 3: Format of the PSC control packet MPLS-TP Protection Figure 3, annotation added. "FPath: used to indicate the path that is reporting a failure condition, the possible values are described in section 4.2.3" MPLS-TP Protection at 11. "4.2.3. Path fault identifier (FPath) The Fpath field of the PSC control SHALL be used only in a Signal fault (0110) or Signal degrade (0101) control packet. Its value indicates on which path the signal anomaly was detected. The following are the possible values:

No.	'821 Patent Claim 9	IETF MPLS-TP System
		• 0: indicates that the fault condition is on the Recovery path • 1: indicates that the fault condition is on the Working path • 2-255: for future extensions." MPLS-TP Protection at 12. "Principles of Operation In all of the following sub-sections, assume a protected domain between LSR-A and LSR-Z, using paths W (working) and R (recovery) as shown in figure 4." MPLS-TP Protection at 13.  MPLS-TP Protection Figure 4 (annotation added).

No.	'821 Patent Claim 10	IETF MPLS-TP System
10	The method of claim 9, wherein said entity cost function comprises a predefined metric.	IETF MPLS-TP System discloses the method of claim 9, wherein said entity cost function comprises a predefined metric. *See* Claim 6.

No.	'821 Patent Claim 11	IETF MPLS-TP System
11	The method of claim 10, wherein said predefined metric is selected from the group consisting of interior gateway protocol (IGP) and traffic engineering (TE).	IETF MPLS-TP System discloses the method of claim 10, wherein said predefined metric is selected from the group consisting of interior gateway protocol (IGP) and traffic engineering (TE). *See* Claim 7.

No.	'821 Patent Claim 12	IETF MPLS-TP System
12	The method of claim 10, further comprising the step of configuring said working entity as revertive.	IETF MPLS-TP System discloses the method of claim 10, further comprising the step of configuring said working entity as revertive. *See* Claim 8.

No.	'821 Patent Claim 13	IETF MPLS-TP System
13	The method of claim 1, further comprising the step of: if said entity pair reselection results in both working and protection entities being replaced,	IETF MPLS-TP System discloses the method of claim 1, further comprising the step of: if said entity pair reselection results in both working and protection entities being replaced, sequentially replacing said working entity and said protection entity. For example, RFC 4872 discloses a series of steps that can be used to replace a working and protection sequentially in response to a network fault.

No.	'821 Patent Claim 13	IETF MPLS-TP System
	sequentially replacing said working entity and said protection entity.	"12. Reversion Reversion refers to a recovery switching operation, where the normal traffic returns to (or remains on) the working LSP when it has recovered from the failure. Reversion implies that resources remain allocated to the LSP that was originally routed over them even after a failure. It is important to have mechanisms that allow reversion to be performed with minimal service disruption and reconfiguration. For "1+1 bidirectional Protection", reversion to the recovered LSP occurs by using the following sequence: 1. Clear the A bit of the ADMIN_STATUS object if set for the recovered LSP. 2. Then, apply the method described below to switch normal traffic back from the protecting to the recovered LSP. This is performed by using the new error code/sub-code "Notify Error/LSP Recovered" (Switchback Request). The procedure is as follows: 1) The initiating (source) node sends the normal traffic onto both the working and the protecting LSPs. Once completed, the source node sends reliably a Notify message to the destination with the new error code/sub-code "Notify Error/LSP Recovered" (Switchback Request). This Notify message includes the MESSAGE_ID object. The ACK_Desired flag MUST be set in this object to request the receiver to send an acknowledgment for the message (see [RFC2961]). 2) Upon receipt of this message, the destination selects the traffic from the working LSP. At the same time, it transmits the traffic onto both the working and protecting LSP. The destination then sends reliably a Notify message to the source confirming the completion of the operation. This message includes the MESSAGE_ID_ACK object to acknowledge reception of the received Notify message. This Notify message also includes the MESSAGE_ID object. The ACK_Desired flag MUST be set in this object to request the receiver to send an acknowledgment for the message (see [RFC2961]). 3) When the source node receives this Notify message, it switches to receive traffic from the working LSP. The source node then sends an Ack message to the destination node confirming that the LSP has been reverted. 3. Finally, clear the O bit of the PROTECTION object sent over the protecting LSP.

No.	'821 Patent Claim 13	IETF MPLS-TP System
		For "1:N Protection with Extra-traffic", reversion to the recovered LSP occurs by using the following sequence: 1. Clear the A bit of the ADMIN_STATUS object if set for the recovered LSP. 2. Then, apply the method described below to switch normal traffic back from the protecting to the recovered LSP. This is performed by using the new error code/sub-code "Notify Error/LSP Recovered" (Switchback Request). The procedure is as follows: 1) The initiating (source) node sends the normal traffic onto both the working and the protecting LSPs. Once completed, the source node sends reliably a Notify message to the destination with the new error code/sub-code "Notify Error/LSP Recovered" (Switchback Request). This Notify message includes the MESSAGE_ID object. The ACK_Desired flag MUST be set in this object to request the receiver to send an acknowledgment for the message (see [RFC2961]). 2) Upon receipt of this message, the destination selects the traffic from the working LSP. At the same time, it transmits the traffic onto both the working and protecting LSP. The destination then sends reliably a Notify message to the source confirming the completion of the operation. This message includes the MESSAGE_ID_ACK object to acknowledge reception of the received Notify message. This Notify message also includes the MESSAGE_ID object. The ACK_Desired flag MUST be set in this object to request the receiver to send an acknowledgment for the message (see [RFC2961]). 3) When the source node receives this Notify message, it switches to receive traffic from the working LSP, and stops transmitting traffic on the protecting LSP. The source node then sends an Ack message to the destination node confirming that the LSP has been reverted. 4) Upon receipt of this message, the destination node stops transmitting traffic along the protecting LSP. 3. Finally, clear the O bit of the PROTECTION object sent over the protecting LSP. For "Rerouting without Extra-traffic" (including the shared recovery case), reversion implies that the formerly working LSP has not been torn down by the head-end node upon PathErr message reception, i.e., the head-end node kept refreshing the working LSP under failure condition. This ensures that the exact same resources are retrieved after reversion switching (except if the working LSP required re- signaling). Re-activation is performed using the following sequence: 1. Clear the A bit of the ADMIN_STATUS object if set for the recovered LSP.

No.	'821 Patent Claim 13	IETF MPLS-TP System
		2. Then, apply the method described below to switch normal traffic back from the protecting to the recovered LSP. This is performed by using the new error code/sub-code "Notify Error/LSP Recovered" (Switchback Request). The procedure is as follows: 1) The initiating (source) node sends the normal traffic onto both the working and the protecting LSPs. Once completed, the source node sends reliably a Notify message to the destination with the new error code/sub-code "Notify Error/LSP Recovered" (Switchback Request). This Notify message includes the MESSAGE_ID object. The ACK_Desired flag MUST be set in this object to request the receiver to send an acknowledgment for the message (see [RFC2961]). 2) Upon receipt of this message, the destination selects the traffic from the working LSP. At the same time, it transmits the traffic onto both the working and protecting LSP. The destination then sends reliably a Notify message to the source confirming the completion of the operation. This message includes the MESSAGE_ID_ACK object to acknowledge reception of the received Notify message. This Notify message also includes the MESSAGE_ID object. The ACK_Desired flag MUST be set in this object to request the receiver to send an acknowledgment for the message (see [RFC2961]). 3) When the source node receives this Notify message, it switches to receive traffic from the working LSP, and stops transmitting traffic on the protecting LSP. The source node then sends an Ack message to the destination node confirming that the LSP has been reverted. 4) Upon receipt of this message, the destination node stops transmitting traffic along the protecting LSP. 3. Finally, de-activate the protecting LSP by setting the S bit to 1 in the PROTECTION object sent over the protecting LSP." RFC 4872 at 26–29. "Note: this document assumes that Protection Type values have end-to- end significance and that the same value is sent over the protected and the protecting path. In this context, shared-mesh (for instance) appears from the end-nodes perspective as being simply an LSP rerouting without extra-traffic services. The net result of this is that a single bit (the S bit alone) does not allow determining whether resource allocation should be performed with respect to the status of the LSP within the protected entity. The introduction of the P bit solves this problem unambiguously. These bits MUST be processed on a hop-by-hop basis (independently of the LSP Protection Type context). This allows for an easier implementation of reversion signaling

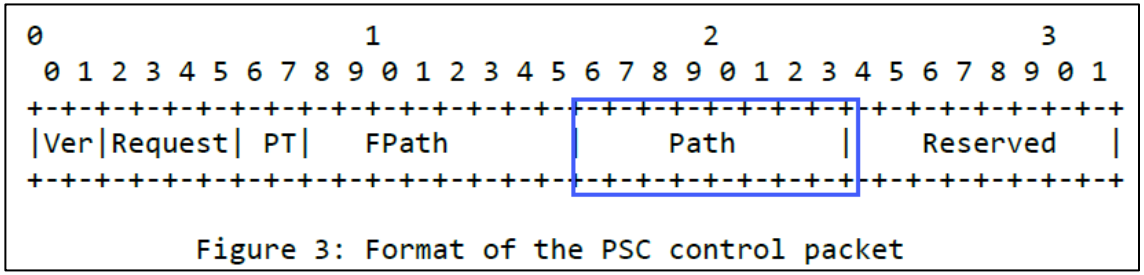
No.	'821 Patent Claim 13	IETF MPLS-TP System
		(see Section 12) but also facilitates the transparent delivery of protected services since any intermediate node is not required to know the semantics associated with the incoming LSP Protection Type value." RFC 4872 at 8–9. "RSVP-TE Extensions in Support of End-to-End Generalized Multi-Protocol Label Switching (GMPLS) Recovery" RFC 4872 Title. "1. Introduction Generalized Multi-Protocol Label Switching (GMPLS) extends MPLS to include support for Layer-2 Switch Capable (L2SC), Time-Division Multiplex (TDM), Lambda Switch Capable (LSC), and Fiber Switch Capable (FSC) interfaces. GMPLS recovery uses control plane mechanisms (i.e., signaling, routing, and link management mechanisms) to support data plane fault recovery. Note that the analogous (data plane) fault detection mechanisms are required to be present in support of the control plane mechanisms. In this document, the term "recovery" is generically used to denote both protection and restoration; the specific terms "protection" and "restoration" are only used when differentiation is required. The subtle distinction between protection and restoration is made based on the resource allocation done during the recovery phase (see [RFC4427]). A functional description of GMPLS recovery is provided in [RFC4426] and should be considered as a companion document. The present document describes the protocol-specific procedures for GMPLS RSVP- TE (Resource ReSerVation Protocol - Traffic Engineering) signaling (see [RFC3473]) to support end-to-end recovery. End-to-end recovery refers to the recovery of an entire LSP from its head-end (ingress node endpoint) to its tail-end (egress node endpoint). With end-to- end recovery, working LSPs are assumed to be resource-disjoint (where a resource is a link, node, or Shared Risk Link Group (SRLG)) in the network so that they do not share any failure probability, but this is not mandatory. With respect to a given set of network resources, a pair of working/protecting LSPs SHOULD be resource disjoint in case of dedicated recovery type (see below). On the other hand, in case of shared recovery (see below), a group of working LSPs SHOULD be mutually resource-disjoint in order to allow for a (single and commonly) shared protecting LSP, itself resource-disjoint from each

No.	'821 Patent Claim 13	IETF MPLS-TP System
		of the working LSPs. Note that resource disjointness is a necessary (but not sufficient) condition to ensure LSP recoverability. The present document addresses four types of end-to-end LSP recovery: 1) 1+1 (unidirectional/bidirectional) protection, 2) 1:N (N >= 1) LSP protection with extra-traffic, 3) pre-planned LSP rerouting without extra-traffic (including shared mesh), and 4) full LSP rerouting. 5) The simplest notion of end-to-end LSP protection is 1+1 unidirectional protection. Using this type of protection, a protecting LSP is signaled over a dedicated resource-disjoint alternate path to protect an associated working LSP. Normal traffic is simultaneously sent on both LSPs and a selector is used at the egress node to receive traffic from one of the LSPs. If a failure occurs along one of the LSPs, the egress node selects the traffic from the valid LSP. No coordination is required between the end nodes when a failure/switchover occurs. In 1+1 bidirectional protection, a protecting LSP is signaled over a dedicated resource-disjoint alternate path to protect the working LSP. Normal traffic is simultaneously sent on both LSPs (in both directions), and a selector is used at both ingress/egress nodes to receive traffic from the same LSP. This requires coordination between the end-nodes when switching to the protecting LSP. 6) In 1:N (N >= 1) protection with extra-traffic, the protecting LSP is a fully provisioned and resource-disjoint LSP from the N working LSPs, that allows for carrying extra-traffic. The N working LSPs MAY be mutually resource-disjoint. Coordination between end-nodes is required when switching from one of the working LSPs to the protecting LSP. As the protecting LSP is fully provisioned, default operations during protection switching are specified for a protecting LSP carrying extra-traffic, but this is not mandatory. Note that M:N protection is out of scope of this document (though mechanisms it defines may be extended to cover it). 7) Pre-planned LSP rerouting (or restoration) relies on the establishment between the same pair of end-nodes of a working LSP and a protecting LSP that is link/node/SRLG disjoint from the working one. Here, the recovery resources for the protecting LSP are pre-reserved but explicit action is required to activate (i.e., commit resource allocation at the data plane) a specific protecting LSP instantiated during the (pre-)provisioning

No.	'821 Patent Claim 13	IETF MPLS-TP System
		phase. Since the protecting LSP is not "active" (i.e., fully instantiated), it cannot carry any extra-traffic. This does not mean that the corresponding resources cannot be used by other LSPs. Therefore, this mechanism protects against working LSP(s) failure(s) but requires activation of the protecting LSP after working LSP failure occurrence. This requires restoration signaling along the protecting path. "Shared-mesh" restoration can be seen as a particular case of pre-planned LSP rerouting that reduces the recovery resource requirements by allowing multiple protecting LSPs to share common link and node resources. The recovery resources are pre-reserved but explicit action is required to activate (i.e., commit resource allocation at the data plane) a specific protecting LSP instantiated during the (pre-) provisioning phase. This procedure requires restoration signaling along the protecting path. Note that in both cases, bandwidth pre-reserved for a protecting (but not activated) LSP can be made available for carrying extra traffic. LSPs for extra-traffic (with lower holding priority than the protecting LSP) can then be established using the bandwidth pre-reserved for the protecting LSP. Also, any lower priority LSP that use the pre-reserved resources for the protecting LSP(s) must be preempted during the activation of the protecting LSP. 8) Full LSP rerouting (or restoration) switches normal traffic to an alternate LSP that is not even partially established until after the working LSP failure occurs. The new alternate route is selected at the LSP head-end node, it may reuse resources of the failed LSP at intermediate nodes and may include additional intermediate nodes and/or links." " RFC 4872 at 3–5. "Consider the following network topology:  The paths [A,B,C,D] and [A,E,F,G,D] are node and link disjoint, ignoring the ingress/egress nodes A and D. A 1+1 protected path is established from A to D over [A,B,C,D] and [A,E,F,G,D], and traffic is transmitted simultaneously over both component paths (i.e., LSPs).

No.	'821 Patent Claim 13	IETF MPLS-TP System
		During the provisioning phase, both LSPs are fully instantiated (and thus activated) so that no resource sharing can be done along the protecting LSP (nor can any extra-traffic be transported). It is also RECOMMENDED to set the N bit since no protection-switching signaling is assumed in this case. When a failure occurs (say, at node B) and is detected at end-node D, the receiver at D selects the normal traffic from the other LSP. From this perspective, 1+1 unidirectional protection can be seen as an uncoordinated protection-switching mechanism acting independently at both endpoints. Also, for the LSP under failure condition, it is RECOMMENDED to not set the Path_State_Removed Flag of the ERROR_SPEC object (see [RFC3473]) upon PathErr message generation. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 9. "When a failure is detected by one or both endpoints of the LSP, both endpoints must select traffic from the other LSP. This action must be coordinated between node A and D. From this perspective, 1+1 bidirectional protection can be seen as a coordinated protection- switching mechanism between both endpoints. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 11. "(Full) LSP Rerouting LSP rerouting, on the other hand, switches normal traffic to an alternate LSP that is fully established only after failure occurrence. The new (alternate) route is selected at the LSP head- end and may reuse intermediate nodes included in the original route; it may also include additional intermediate nodes. For strict-hop routing, TE requirements can be directly applied to the route computation, and the failed node or link can be avoided. However, if the failure occurred within a loose-routed hop, the head-end node may not have enough information to reroute the LSP around the failure. Crankback signaling (see [CRANK]) and route exclusion techniques (see [RFC4874]) MAY be used in this case.

No.	'821 Patent Claim 13	IETF MPLS-TP System
		The alternate route MAY be either computed on demand (that is, when the failure occurs; this is referred to as full LSP rerouting) or pre-computed and stored for use when the failure is reported. The latter offers faster restoration time. There is, however, a risk that the alternate route will become out of date through other changes in the network; this can be mitigated to some extent by periodic recalculation of idle alternate routes. (Full) LSP rerouting will be initiated by the head-end node that has either detected the LSP failure or received a Notify message and/or a PathErr message with the new error code/sub-code "Notify Error/LSP Locally Failed" for this LSP. The new LSP resources can be established using the make-before-break mechanism, where the new LSP is set up before the old LSP is torn down. This is done by using the mechanisms of the SESSION_ATTRIBUTE object and the Shared-Explicit (SE) reservation style (see [RFC3209]). Both the new and old LSPs can share resources at common nodes. Note that the make-before-break mechanism is not used to avoid disruption to the normal traffic flow (the latter has already been broken by the failure that is being repaired). However, it is valuable to retain the resources allocated on the original LSP that will be reused by the new alternate LSP." RFC 4872 at 25. "The procedure is as follows: 43) The initiating (source) node sends the normal traffic onto both the working and the protecting LSPs. Once completed, the source node sends reliably a Notify message to the destination with the new error code/sub-code "Notify Error/LSP Recovered" (Switchback Request). This Notify message includes the MESSAGE_ID object. The ACK_Desired flag MUST be set in this object to request the receiver to send an acknowledgment for the message (see [RFC2961]). 44) Upon receipt of this message, the destination selects the traffic from the working LSP. At the same time, it transmits the traffic onto both the working and protecting LSP. The destination then sends reliably a Notify message to the source confirming the completion of the operation. This message includes the MESSAGE_ID_ACK object to acknowledge reception of the received Notify message. This Notify message also includes the MESSAGE_ID object. The ACK_Desired flag MUST be set in this object to request the receiver to send an acknowledgment for the message (see [RFC2961]).

No.	'821 Patent Claim 13	IETF MPLS-TP System
		45) When the source node receives this Notify message, it switches to receive traffic from the working LSP. The source node then sends an Ack message to the destination node confirming that the LSP has been reverted." RFC 4872 at 27. For example, MPLS-TP Protection discloses an active path indicator field that indicates which LSP and/or pseudo-wire, i.e. network entity, will carry network traffic. MPLS-TP Protection further discloses that the active path indicator can sequentially select a working path and a protection path.  0 1 2 3 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 |Ver|Request| PT| FPath | Path | Reserved | Figure 3: Format of the PSC control packet MPLS-TP Figure 3 (annotation added). "3. Protection switching logic 3.1. Protection switching trigger mechanisms The protection switching should be initiated in reaction to any of the following triggers: • Server layer indication - if the MPLS-TP server layer detects a failure within its own layer, or due to a failure of its server layer (e.g. the physical layer) notifies the MPLS-TP layer that a failure has been detected. It should be noted that this trigger will be controlled by a holdoff-timer that should be configured to allow the server layer to apply any protection procedures that may be relevant. • OAM signalling - if, for example, OAM continuity and connectivity verification tools detect that there is a loss of continuity or mis-connectivity or performance monitoring indicates a degradation of the utility of the working path for the current transport path.

No.	'821 Patent Claim 13	IETF MPLS-TP System
		In cases of signal degradation, switching to the recovery path SHOULD only be activated if the recovery path can guarantee better conditions than the degraded working path. • Control plane - if there is a control plane active in the network (either signaling or routing), it MAY trigger protection switching based on conditions detected by the control plane. If the control-plane is based on GMPLS [RFC3945] then the recovery process should comply with the process described in [RFC4872]. • Operator command - the network operator may issue commands that trigger protection switching. The commands that are supported include - Forced Switch, Manual Switch, Clear, Lockout of Protection, (see definitions in [RFC4427]).” MPLS-TP at 6–7. “3.2. Protection switching control logical architecture Protection switching processes the triggers described above together with the inputs received from the far-end LSR. These inputs cause the LSR to take certain actions, e.g. switching the Selector Bridge to select the working or recovery path, and to transmit different protocol messages.

No.	'821 Patent Claim 13	IETF MPLS-TP System

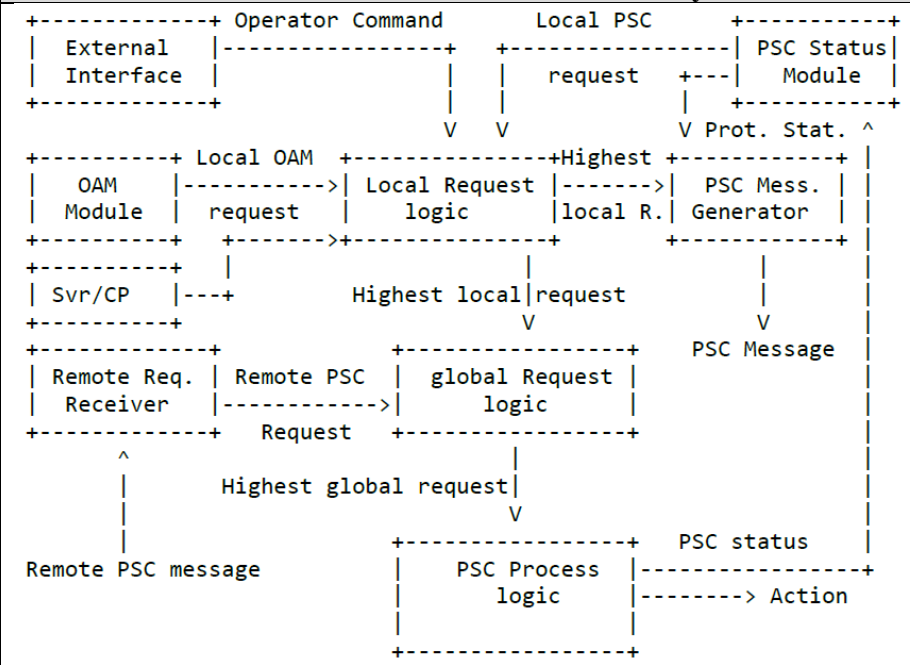
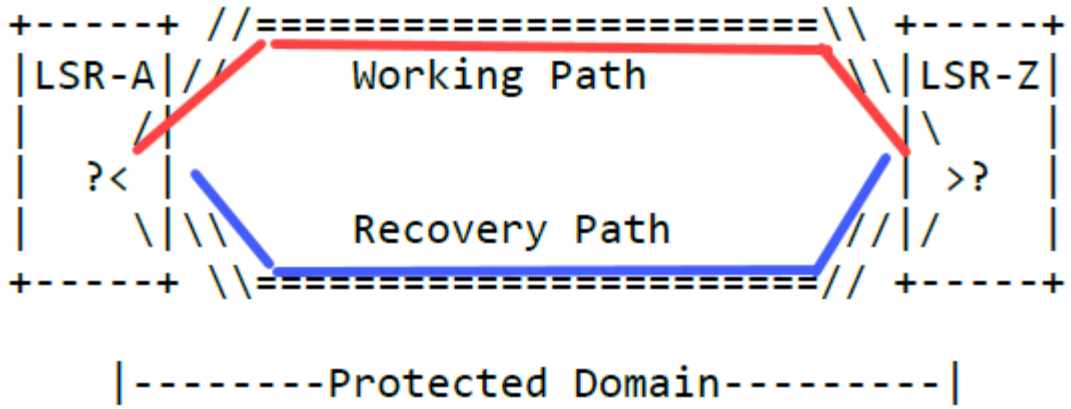


Figure 1: Protection switching control logic

Figure 1 describes the logical architecture of the protection switching control. The Local Request logic unit accepts the triggers from the OAM, external operator commands, and from the local control plane (when present) and determines the highest priority request. This high-priority request is passed to both the PSC Message generator, that will generate the appropriate protocol message to be sent to the far-end LSR, and the Global Request logic, that will cross-check this local request with the information received from the far-LSR. The Global Request logic then processes these two PSC requests that determines the highest priority request that is passed to the PSC Process logic. The PSC Process logic uses this input to determine what actions need to be taken, e.g. switching the Selector Bridge, and the current status of the protection domain." MPLS-TP Protection at 7–8.

"4.2.4. Active path indicator (Path)

No.	'821 Patent Claim 13	IETF MPLS-TP System
		The Path field of the PSC control SHALL be used to indicate which path the source MEP is currently using for data transmission. The MEP should compare the value of this bit with the path that is locally selected for data transmission to verify that there is no inconsistency between the two end-points of the protected domain. If an inconsistency is detected then an alarm should be raised. The following are the possible values: • 0: indicates that normal traffic is being transmitted on the Working path. • 1: indicates the Recovery path is being used to transmit the normal traffic from the Working path. 2-255: for future extensions" MPLS-TP Protection at 12–13. "Principles of Operation In all of the following sub-sections, assume a protected domain between LSR-A and LSR-Z, using paths W (working) and R (recovery) as shown in figure 4." MPLS-TP Protection at 13.  MPLS-TP Protection Figure 4 (annotation added). "4.3.1 PSC States 4.3.1.1 Normal State

No.	'821 Patent Claim 13	IETF MPLS-TP System
		When the protected domain has no special condition in effect, the ingress LSR SHOULD forward the user data along the working path, and, in the case of 1+1 protection, the Permanent Bridge will bridge the data to the recovery path as well. The receiving LSR SHOULD read the data from the working path. The ingress LSR MAY transmit a No Request PSC packet with the Path field set to 0 indicating that the normal data traffic should be read from the working path. 4.3.1.2. Protecting State When the protection mechanism has been triggered and the protected domain has performed a protection switch, the domain is in the protecting state. In this state the normal data traffic is transmitted and received on the recovery path. If the protection domain is currently in a protecting state, then the LSRs SHOULD NOT accept a Manual Switch request. If the protection domain is currently in a protecting state, and a Forced Switch is requested then the normal traffic SHALL continue to be transmitted on the recovery path even if the original protection trigger is cleared, and the Forced Switch condition will be signalled by the PSC messages. 4.3.1.3. Unavailable State When the recovery path is unavailable - either as a result of a Lockout operator command (see section 4.3.3), or as a result of a SF or SD detected on the recovery path (see section 4.3.4) - then the protection domain is in the unavailable state. In this state, the normal traffic is transmitted and received on the working path. While in unavailable state any event that would trigger a protection switching SHOULD be ignored with the following exception - If a Signal Degrade request is received, then protection switching will be activated only if the recovery path can guarantee a better signal than the working path. The protection domain will exit the unavailable state and revert to the normal state when, either the operator clears the Lockout command or the recovery path recovers from the signal fault or degraded situation.

No.	'821 Patent Claim 13	IETF MPLS-TP System
		Both ends will resume sending the PCS packets over the recovery path, as a result of this recovery. 4.3.2. Failure or Degraded condition (Working path) If one of the LSRs (for example, LSR-A) detects a failure condition or a serious degradation condition on the working path that warrants invoking protection switching, then it SHOULD take the following actions: • (For 1:1 protection) Switch all traffic for LSR-Z to the recovery path only. • Transmit a PCS control packet, using GACH, with the appropriate Request code (either Signal fault or Signal degrade), the Fpath set to 1, to indicate that the fault/degrade was detected on the working path, and the Path set to 1, indicating that normal traffic is now being transmitted on the recovery path. • Verify that LSR-Z replies with a PCS control packet indicating that it has switched to the recovery path. If this is not received after 2 PSC cycles then send an alarm to the management system. When the far-end LSR (in this example LSR-Z) receives the PCS packet informing it that other LSR (LSR-A) has switched, it SHOULD perform the following actions: • Check priority of the request • Switch all traffic addressed to LSR-A to the recovery path only (for 1:1 protection). • Begin transmission of a PCS control packet, using GACH, with the appropriate Request code (either Signal fault or Signal degrade), the Fpath set to 1, to indicate that the fault/degrade was detected on the working path, and the Path set to 1, indicating that traffic is now being transmitted on the recovery path." MPLS-TP Protection at 13–15. "4.3.4. Failure or Degraded condition (Recovery path) If one of the LSRs (for example, LSR-A) detects a failure condition or a serious degradation condition on the recovery path, then it SHOULD take the following actions:

No.	'821 Patent Claim 13	IETF MPLS-TP System
		• Begin transmission of a PCS control packet with the appropriate Request code (either Signal fault or Signal degrade), the Fpath set to 0, to indicate that the fault/degrade was detected on the recovery path, and the Path set to 0, indicating that traffic is now being forwarded on the working path. Note that this will actually reach the far-end if this is a unidirectional fault or recovery path is possibly in a degraded situation. • The PSC control logic should go into Unavailable state. • All traffic MUST be transmitted on the working path for the duration of the SF/SD condition. When the far-end LSR (in this example LSR-Z) receives the PCS packet informing it that other LSR (LSR-A) has become Unavailable, it SHOULD perform the following actions: • Transmit all traffic on the working path for the duration of the SF/SD condition The PSC Control logic should go into Unavailable state." MPLS-TP Protection at 16.

No.	'821 Patent Claim 14	IETF MPLS-TP System
14[preamble]	A system for selecting entities within an MPLS network, comprising:	IETF MPLS-TP System discloses a system for selecting entities within an MPLS network, comprising. *See* Claim 1[preamble] and 1[a].
14[a]	a data structure comprising a plurality of transport entity descriptors;	IETF MPLS-TP System discloses a data structure comprising a plurality of transport entity descriptors. For example, RFC 4872 discloses that the SRLG identifiers can be determined. A person of ordinary skill in the art would understand that an IGP-TE database includes metrics consisting of Interior Gateway Protocol (IGP) and Traffic Engineering (TE). "Note: an intermediate node processing a PPRO can derive SRLG identifiers from the local IGP-TE database using its Type 1, 2, or 4 subobject values as pointers to the corresponding TE Links (assuming each of them has an associated SRLG TE attribute)." RFC 4872 at 35. "Abstract

No.	'821 Patent Claim 14	IETF MPLS-TP System
		This document describes protocol-specific procedures and extensions for Generalized Multi-Protocol Label Switching (GMPLS) Resource Reservation Protocol - Traffic Engineering (RSVP-TE) signaling to support end-to-end Label Switched Path (LSP) recovery that denotes protection and restoration. A generic functional description of GMPLS recovery can be found in a companion document, RFC 4426." RFC 4872 at 1. "A functional description of GMPLS recovery is provided in [RFC4426] and should be considered as a companion document. The present document describes the protocol-specific procedures for GMPLS RSVP- TE (Resource Reservation Protocol - Traffic Engineering) signaling (see [RFC3473]) to support end-to-end recovery. End-to-end recovery refers to the recovery of an entire LSP from its head-end (ingress node endpoint) to its tail-end (egress node endpoint). With end-to- end recovery, working LSPs are assumed to be resource-disjoint (where a resource is a link, node, or Shared Risk Link Group (SRLG)) in the network so that they do not share any failure probability, but this is not mandatory. With respect to a given set of network resources, a pair of working/protecting LSPs SHOULD be resource disjoint in case of dedicated recovery type (see below). On the other hand, in case of shared recovery (see below), a group of working LSPs SHOULD be mutually resource-disjoint in order to allow for a (single and commonly) shared protecting LSP, itself resource-disjoint from each of the working LSPs. Note that resource disjointness is a necessary (but not sufficient) condition to ensure LSP recoverability." RFC 4872 at 3. For example, MPLS-TP Protection discloses the use of a Label Switched Router to perform the described standard. A person of ordinary skill in the art would understand the disclosed Label Switched Router to include computer-readable data structures. "As noted in the architecture for Multi-Protocol Label Switching Transport Profile (MPLS-TP) [TPFwk], the overall architecture framework for MPLS-TP is based on a profile of the MPLS and Pseudowire (PW) procedures as specified for the MPLS and (MS-)PW architectures defined in [RFC3031], [RFC3985] and [RFC5085]. One of the basic survivability functions, pointed out by the Survivability Framework document [SurvivFwk], is that of simple and rapid protection switching mechanisms for Label Switched Paths (LSP) and Pseudo-wires (PW).

No.	'821 Patent Claim 14	IETF MPLS-TP System
		Protection switching is a fully allocated survivability mechanism. It is fully allocated in the sense that the route and bandwidth of the recovery path is reserved for a selected working path or set of working paths. It provides a fast and simple survivability mechanism, that allows the network operator to easily grasp the active state of the network, compared to other survivability mechanisms. As specified in the Survivability Framework document [SurvivFwk], protection switching is applied to a protected domain. For the purposes of this document, we define the protected domain of a P2P LSP as consisting of two Label Switching Routers (LSR) and the transport paths that connect them. For a P2MP LSP the protection domain includes the root (or source) LSR, the destination (or sink) LSRs, and the transport paths that connect them." MPLS-TP Protection at 3.
14[b]	an entity protection switch configured to switch between a working entity and a protection entity; and	IETF MPLS-TP System discloses an entity protection switch configured to switch between a working entity and a protection entity. For example, RFC 4872 discloses protection switching, to switch between working and protection and/or restoration LSPs. "In 1:N (N >= 1) protection with extra-traffic, the protecting LSP is a fully provisioned and resource-disjoint LSP from the N working LSPs, that allows for carrying extra-traffic. The N working LSPs MAY be mutually resource-disjoint. Coordination between end-nodes is required when switching from one of the working LSPs to the protecting LSP. As the protecting LSP is fully provisioned, default operations during protection switching are specified for a protecting LSP carrying extra-traffic, but this is not mandatory. Note that M:N protection is out of scope of this document (though mechanisms it defines may be extended to cover it)." RFC 4872 at 4. "O (Operational) bit: this bit is set when a protecting LSP is carrying the normal traffic after protection switching (i.e., applies only in case of dedicated LSP protection or LSP protection with extra-traffic; see Section 4.2.2)." RFC 4872 at 7. "For LSP protection, in particular, when the data plane provides automated protection-switching capability (see for instance ITU-T [G.841] Recommendation), a Notification (N)

No.	'821 Patent Claim 14	IETF MPLS-TP System
		bit is defined in the PROTECTION object. It allows for distinction between protection switching signaling via the control plane or the data plane." RFC 4872 at 8. "When a failure occurs (say, at node B) and is detected at end-node D, the receiver at D selects the normal traffic from the other LSP. From this perspective, 1+1 unidirectional protection can be seen as an uncoordinated protection-switching mechanism acting independently at both endpoints. Also, for the LSP under failure condition, it is RECOMMENDED to not set the Path_State_Removed Flag of the ERROR_SPEC object (see [RFC3473]) upon PathErr message generation." RFC 4872 at 9. "After protection switching completes, and after reception of the PathErr message, to keep track of the LSP from which the signal is taken, the protecting LSP SHOULD be signaled with the O bit set. The formerly working LSP MAY be signaled with the A bit set in the ADMIN_STATUS object (see [RFC3473]). This process assumes the tail- end node has notified the head-end node that traffic selection switchover has occurred. When a failure is detected by one or both endpoints of the LSP, both endpoints must select traffic from the other LSP. This action must be coordinated between node A and D. From this perspective, 1+1 bidirectional protection can be seen as a coordinated protection- switching mechanism between both endpoints." RFC 4872 at 10–11. "After protection switching completes (step 2), and after reception of the PathErr message, to keep track of the LSP from which the signal is taken, the protecting LSP SHOULD be signaled with the O bit set. The formerly working LSP MAY be signaled with the A bit set in the ADMIN_STATUS object (see [RFC3473])." RFC 4872 at 13. For example, MPLS-TP Protection discloses the use of a Label Switched Router to perform the described standard. MPLS-TP Protection further discloses protection switching logic that automatically deploys in response to a protection trigger. "As noted in the architecture for Multi-Protocol Label Switching Transport Profile (MPLS-TP) [TPFwk], the overall architecture framework for MPLS-TP is based on a profile of the MPLS and Pseudowire (PW) procedures as specified for the MPLS and (MS-)PW

No.	'821 Patent Claim 14	IETF MPLS-TP System
		architectures defined in [RFC3031], [RFC3985] and [RFC5085]. One of the basic survivability functions, pointed out by the Survivability Framework document [SurvivFwk], is that of simple and rapid protection switching mechanisms for Label Switched Paths (LSP) and Pseudo-wires (PW). Protection switching is a fully allocated survivability mechanism. It is fully allocated in the sense that the route and bandwidth of the recovery path is reserved for a selected working path or set of working paths. It provides a fast and simple survivability mechanism, that allows the network operator to easily grasp the active state of the network, compared to other survivability mechanisms. As specified in the Survivability Framework document [SurvivFwk], protection switching is applied to a protected domain. For the purposes of this document, we define the protected domain of a P2P LSP as consisting of two Label Switching Routers (LSR) and the transport paths that connect them. For a P2MP LSP the protection domain includes the root (or source) LSR, the destination (or sink) LSRs, and the transport paths that connect them." MPLS-TP Protection at 3. "4.3.1 PSC States 4.3.1.1 Normal State When the protected domain has no special condition in effect, the ingress LSR SHOULD forward the user data along the working path, and, in the case of 1+1 protection, the Permanent Bridge will bridge the data to the recovery path as well. The receiving LSR SHOULD read the data from the working path. The ingress LSR MAY transmit a No Request PSC packet with the Path field set to 0 indicating that the normal data traffic should be read from the working path. 4.3.1.2. Protecting State When the protection mechanism has been triggered and the protected domain has performed a protection switch, the domain is in the protecting state. In this state the normal data traffic

No.	'821 Patent Claim 14	IETF MPLS-TP System
		is transmitted and received on the recovery path. If the protection domain is currently in a protecting state, then the LSRs SHOULD NOT accept a Manual Switch request. If the protection domain is currently in a protecting state, and a Forced Switch is requested then the normal traffic SHALL continue to be transmitted on the recovery path even if the original protection trigger is cleared, and the Forced Switch condition will be signalled by the PSC messages. 4.3.1.3. Unavailable State When the recovery path is unavailable - either as a result of a Lockout operator command (see section 4.3.3), or as a result of a SF or SD detected on the recovery path (see section 4.3.4) - then the protection domain is in the unavailable state. In this state, the normal traffic is transmitted and received on the working path. While in unavailable state any event that would trigger a protection switching SHOULD be ignored with the following exception - If a Signal Degrade request is received, then protection switching will be activated only if the recovery path can guarantee a better signal than the working path. The protection domain will exit the unavailable state and revert to the normal state when, either the operator clears the Lockout command or the recovery path recovers from the signal fault or degraded situation. Both ends will resume sending the PCS packets over the recovery path, as a result of this recovery. 4.3.2. Failure or Degraded condition (Working path) If one of the LSRs (for example, LSR-A) detects a failure condition or a serious degradation condition on the working path that warrants invoking protection switching, then it SHOULD take the following actions: • (For 1:1 protection) Switch all traffic for LSR-Z to the recovery path only. • Transmit a PCS control packet, using GACH, with the appropriate Request code (either Signal fault or Signal degrade), the Fpath set to 1, to indicate that the fault/degrade was detected on the working path, and the Path set to 1, indicating that normal traffic is now being transmitted on the recovery path.

No.	'821 Patent Claim 14	IETF MPLS-TP System
		• Verify that LSR-Z replies with a PCS control packet indicating that it has switched to the recovery path. If this is not received after 2 PSC cycles then send an alarm to the management system. When the far-end LSR (in this example LSR-Z) receives the PCS packet informing it that other LSR (LSR-A) has switched, it SHOULD perform the following actions: • Check priority of the request • Switch all traffic addressed to LSR-A to the recovery path only (for 1:1 protection). • Begin transmission of a PCS control packet, using GACH, with the appropriate Request code (either Signal fault or Signal degrade), the Fpath set to 1, to indicate that the fault/degrade was detected on the working path, and the Path set to 1, indicating that traffic is now being transmitted on the recovery path." MPLS-TP Protection at 13–15. "4.3.4. Failure or Degraded condition (Recovery path) If one of the LSRs (for example, LSR-A) detects a failure condition or a serious degradation condition on the recovery path, then it SHOULD take the following actions: • Begin transmission of a PCS control packet with the appropriate Request code (either Signal fault or Signal degrade), the Fpath set to 0, to indicate that the fault/degrade was detected on the recovery path, and the Path set to 0, indicating that traffic is now being forwarded on the working path. Note that this will actually reach the far-end if this is a unidirectional fault or recovery path is possibly in a degraded situation. • The PSC control logic should go into Unavailable state. • All traffic MUST be transmitted on the working path for the duration of the SF/SD condition. When the far-end LSR (in this example LSR-Z) receives the PCS packet informing it that other LSR (LSR-A) has become Unavailable, it SHOULD perform the following actions: • Transmit all traffic on the working path for the duration of the SF/SD condition The PSC Control logic should go into Unavailable state." MPLS-TP Protection at 16.

No.	'821 Patent Claim 14	IETF MPLS-TP System
14[c]	digital logic configured to select said working entity and said protection entity from said plurality of transport entity descriptors, comprising:	IETF MPLS-TP System discloses digital logic configured to select said working entity and said protection entity from said plurality of transport entity descriptors, comprising. For example, RFC 4872 discloses that the plurality of MPLS LSPs, i.e., transport entities, are selected in pairs of either working and protection LSPs or working and restoration LSPs. "When a failure is detected by one or both endpoints of the LSP, both endpoints must select traffic from the other LSP. This action must be coordinated between node A and D. From this perspective, 1+1 bidirectional protection can be seen as a coordinated protection- switching mechanism between both endpoints. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 11. "1:1 Protection with Extra-Traffic The most common case of end-to-end 1:N protection is to establish, between the same endpoints, an end-to-end working LSP (thus, N = 1) and a dedicated end-to-end protecting LSP that are mutually link/ node/SRLG disjoint. This protects against working LSP failure(s)." RFC 4872 at 13. "An approach to reduce recovery resource requirements is to have protection LSPs sharing network resources when the working LSPs that they protect are physically (i.e., link, node, SRLG, etc.) disjoint. This mechanism is referred to as shared mesh restoration and is described in [RFC4426]. Shared-mesh restoration can be seen as a particular case of pre-planned LSP rerouting (see Section 8) that reduces the recovery resource requirements by allowing multiple protecting LSPs to share common link and node resources. Here also, the recovery resources for the protecting LSPs are pre-reserved during the provisioning phase, thus an explicit signaling action is required to activate (i.e., commit resource allocation at the data plane) a specific protecting LSP instantiated during the (pre-) provisioning phase. This requires restoration signaling along the protecting LSP." RFC 4872 at 20. "1. Introduction

No.	'821 Patent Claim 14	IETF MPLS-TP System
		Generalized Multi-Protocol Label Switching (GMPLS) extends MPLS to include support for Layer-2 Switch Capable (L2SC), Time-Division Multiplex (TDM), Lambda Switch Capable (LSC), and Fiber Switch Capable (FSC) interfaces. GMPLS recovery uses control plane mechanisms (i.e., signaling, routing, and link management mechanisms) to support data plane fault recovery. Note that the analogous (data plane) fault detection mechanisms are required to be present in support of the control plane mechanisms. In this document, the term "recovery" is generically used to denote both protection and restoration; the specific terms "protection" and "restoration" are only used when differentiation is required. The subtle distinction between protection and restoration is made based on the resource allocation done during the recovery phase (see [RFC4427]). A functional description of GMPLS recovery is provided in [RFC4426] and should be considered as a companion document. The present document describes the protocol-specific procedures for GMPLS RSVP- TE (Resource ReSerVation Protocol - Traffic Engineering) signaling (see [RFC3473]) to support end-to-end recovery. End-to-end recovery refers to the recovery of an entire LSP from its head-end (ingress node endpoint) to its tail-end (egress node endpoint). With end-to- end recovery, working LSPs are assumed to be resource-disjoint (where a resource is a link, node, or Shared Risk Link Group (SRLG)) in the network so that they do not share any failure probability, but this is not mandatory. With respect to a given set of network resources, a pair of working/protecting LSPs SHOULD be resource disjoint in case of dedicated recovery type (see below). On the other hand, in case of shared recovery (see below), a group of working LSPs SHOULD be mutually resource-disjoint in order to allow for a (single and commonly) shared protecting LSP, itself resource-disjoint from each of the working LSPs. Note that resource disjointness is a necessary (but not sufficient) condition to ensure LSP recoverability. The present document addresses four types of end-to-end LSP recovery: 1) 1+1 (unidirectional/bidirectional) protection, 2) 1:N (N >= 1) LSP protection with extra-traffic, 3) pre-planned LSP rerouting without extra-traffic (including shared mesh), and 4) full LSP rerouting. 53) The simplest notion of end-to-end LSP protection is 1+1 unidirectional protection. Using this type of protection, a protecting LSP is signaled over a dedicated resource-

No.	'821 Patent Claim 14	IETF MPLS-TP System
		disjoint alternate path to protect an associated working LSP. Normal traffic is simultaneously sent on both LSPs and a selector is used at the egress node to receive traffic from one of the LSPs. If a failure occurs along one of the LSPs, the egress node selects the traffic from the valid LSP. No coordination is required between the end nodes when a failure/switchover occurs. In 1+1 bidirectional protection, a protecting LSP is signaled over a dedicated resource-disjoint alternate path to protect the working LSP. Normal traffic is simultaneously sent on both LSPs (in both directions), and a selector is used at both ingress/egress nodes to receive traffic from the same LSP. This requires coordination between the end-nodes when switching to the protecting LSP. 54) In 1:N (N >= 1) protection with extra-traffic, the protecting LSP is a fully provisioned and resource-disjoint LSP from the N working LSPs, that allows for carrying extra-traffic. The N working LSPs MAY be mutually resource-disjoint. Coordination between end-nodes is required when switching from one of the working LSPs to the protecting LSP. As the protecting LSP is fully provisioned, default operations during protection switching are specified for a protecting LSP carrying extra-traffic, but this is not mandatory. Note that M:N protection is out of scope of this document (though mechanisms it defines may be extended to cover it). 55) Pre-planned LSP rerouting (or restoration) relies on the establishment between the same pair of end-nodes of a working LSP and a protecting LSP that is link/node/SRLG disjoint from the working one. Here, the recovery resources for the protecting LSP are pre-reserved but explicit action is required to activate (i.e., commit resource allocation at the data plane) a specific protecting LSP instantiated during the (pre-)provisioning phase. Since the protecting LSP is not "active" (i.e., fully instantiated), it cannot carry any extra-traffic. This does not mean that the corresponding resources cannot be used by other LSPs. Therefore, this mechanism protects against working LSP(s) failure(s) but requires activation of the protecting LSP after working LSP failure occurrence. This requires restoration signaling along the protecting path. "Shared-mesh" restoration can be seen as a particular case of pre-planned LSP rerouting that reduces the recovery resource requirements by allowing multiple protecting LSPs to share common link and node resources. The recovery resources are pre-reserved but explicit action is required to activate (i.e., commit resource allocation at the data plane) a

No.	'821 Patent Claim 14	IETF MPLS-TP System
		specific protecting LSP instantiated during the (pre-) provisioning phase. This procedure requires restoration signaling along the protecting path. Note that in both cases, bandwidth pre-reserved for a protecting (but not activated) LSP can be made available for carrying extra traffic. LSPs for extra-traffic (with lower holding priority than the protecting LSP) can then be established using the bandwidth pre-reserved for the protecting LSP. Also, any lower priority LSP that use the pre-reserved resources for the protecting LSP(s) must be preempted during the activation of the protecting LSP. 56) Full LSP rerouting (or restoration) switches normal traffic to an alternate LSP that is not even partially established until after the working LSP failure occurs. The new alternate route is selected at the LSP head-end node, it may reuse resources of the failed LSP at intermediate nodes and may include additional intermediate nodes and/or links." " RFC 4872 at 3–5. "Consider the following network topology:  The paths [A,B,C,D] and [A,E,F,G,D] are node and link disjoint, ignoring the ingress/egress nodes A and D. A 1+1 protected path is established from A to D over [A,B,C,D] and [A,E,F,G,D], and traffic is transmitted simultaneously over both component paths (i.e., LSPs). During the provisioning phase, both LSPs are fully instantiated (and thus activated) so that no resource sharing can be done along the protecting LSP (nor can any extra-traffic be transported). It is also RECOMMENDED to set the N bit since no protection-switching signaling is assumed in this case. When a failure occurs (say, at node B) and is detected at end-node D, the receiver at D selects the normal traffic from the other LSP. From this perspective, 1+1 unidirectional protection can be seen as an uncoordinated protection-switching mechanism acting independently at both endpoints. Also, for the LSP under failure condition, it is RECOMMENDED to not set the Path_State_Removed Flag of the ERROR_SPEC object (see [RFC3473]) upon PathErr message generation.

No.	'821 Patent Claim 14	IETF MPLS-TP System
		Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 9. "When a failure is detected by one or both endpoints of the LSP, both endpoints must select traffic from the other LSP. This action must be coordinated between node A and D. From this perspective, 1+1 bidirectional protection can be seen as a coordinated protection- switching mechanism between both endpoints. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 11. "(Full) LSP Rerouting LSP rerouting, on the other hand, switches normal traffic to an alternate LSP that is fully established only after failure occurrence. The new (alternate) route is selected at the LSP head- end and may reuse intermediate nodes included in the original route; it may also include additional intermediate nodes. For strict-hop routing, TE requirements can be directly applied to the route computation, and the failed node or link can be avoided. However, if the failure occurred within a loose-routed hop, the head-end node may not have enough information to reroute the LSP around the failure. Crankback signaling (see [CRANK]) and route exclusion techniques (see [RFC4874]) MAY be used in this case. The alternate route MAY be either computed on demand (that is, when the failure occurs; this is referred to as full LSP rerouting) or pre-computed and stored for use when the failure is reported. The latter offers faster restoration time. There is, however, a risk that the alternate route will become out of date through other changes in the network; this can be mitigated to some extent by periodic recalculation of idle alternate routes. (Full) LSP rerouting will be initiated by the head-end node that has either detected the LSP failure or received a Notify message and/or a PathErr message with the new error code/sub-code "Notify Error/LSP Locally Failed" for this LSP. The new LSP resources can be established using the make-before-break mechanism, where the new LSP is set up before the old LSP is torn down. This is done by using the mechanisms of the SESSION_ATTRIBUTE

No.	'821 Patent Claim 14	IETF MPLS-TP System
		object and the Shared-Explicit (SE) reservation style (see [RFC3209]). Both the new and old LSPs can share resources at common nodes. Note that the make-before-break mechanism is not used to avoid disruption to the normal traffic flow (the latter has already been broken by the failure that is being repaired). However, it is valuable to retain the resources allocated on the original LSP that will be reused by the new alternate LSP." RFC 4872 at 25. "The procedure is as follows: 46) The initiating (source) node sends the normal traffic onto both the working and the protecting LSPs. Once completed, the source node sends reliably a Notify message to the destination with the new error code/sub-code "Notify Error/LSP Recovered" (Switchback Request). This Notify message includes the MESSAGE_ID object. The ACK_Desired flag MUST be set in this object to request the receiver to send an acknowledgment for the message (see [RFC2961]). 47) Upon receipt of this message, the destination selects the traffic from the working LSP. At the same time, it transmits the traffic onto both the working and protecting LSP. The destination then sends reliably a Notify message to the source confirming the completion of the operation. This message includes the MESSAGE_ID_ACK object to acknowledge reception of the received Notify message. This Notify message also includes the MESSAGE_ID object. The ACK_Desired flag MUST be set in this object to request the receiver to send an acknowledgment for the message (see [RFC2961]). 48) When the source node receives this Notify message, it switches to receive traffic from the working LSP. The source node then sends an Ack message to the destination node confirming that the LSP has been reverted." RFC 4872 at 27. For example, MPLS-TP Protection discloses a protection switching framework that includes the selection of a one or more "working paths" and "recovery paths," e.g. pseudo wires and/or LSPs, within an MPLS Network. MPLS-TP Protection further discloses that this method can be performed by a Label Switched Router, i.e. a network device.

No.	'821 Patent Claim 14	IETF MPLS-TP System
		"As noted in the architecture for Multi-Protocol Label Switching Transport Profile (MPLS-TP) [TPFwk], the overall architecture framework for MPLS-TP is based on a profile of the MPLS and Pseudowire (PW) procedures as specified for the MPLS and (MS-)PW architectures defined in [RFC3031], [RFC3985] and [RFC5085]. One of the basic survivability functions, pointed out by the Survivability Framework document [SurvivFwk], is that of simple and rapid protection switching mechanisms for Label Switched Paths (LSP) and Pseudo-wires (PW). Protection switching is a fully allocated survivability mechanism. It is fully allocated in the sense that the route and bandwidth of the recovery path is reserved for a selected working path or set of working paths. It provides a fast and simple survivability mechanism, that allows the network operator to easily grasp the active state of the network, compared to other survivability mechanisms. As specified in the Survivability Framework document [SurvivFwk], protection switching is applied to a protected domain. For the purposes of this document, we define the protected domain of a P2P LSP as consisting of two Label Switching Routers (LSR) and the transport paths that connect them. For a P2MP LSP the protection domain includes the root (or source) LSR, the destination (or sink) LSRs, and the transport paths that connect them." MPLS-TP Protection at 3. "3.2. Protection switching control logical architecture Protection switching processes the triggers described above together with the inputs received from the far-end LSR. These inputs cause the LSR to take certain actions, e.g. switching the Selector Bridge to select the working or recovery path, and to transmit different protocol messages.

No.	'821 Patent Claim 14	IETF MPLS-TP System

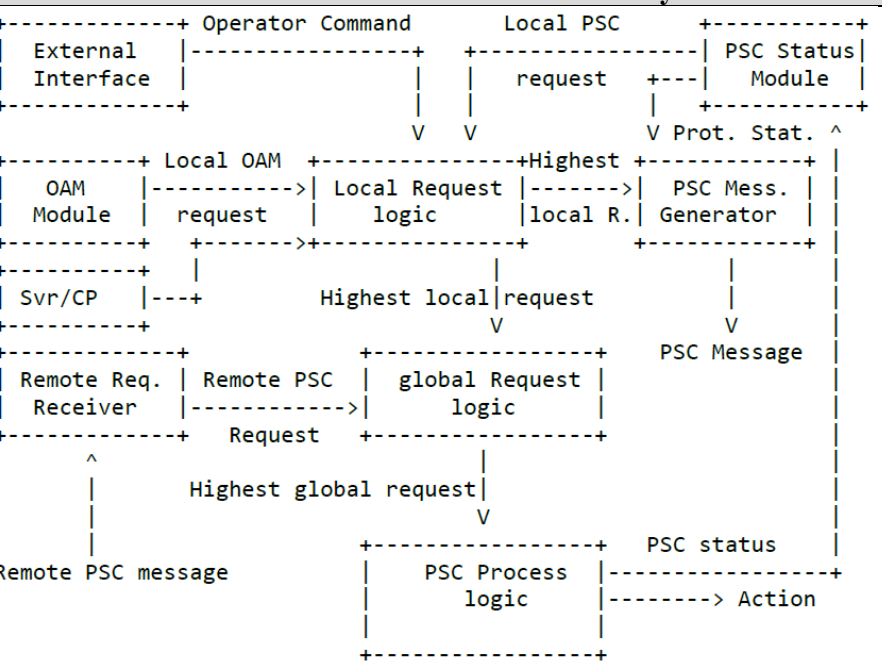
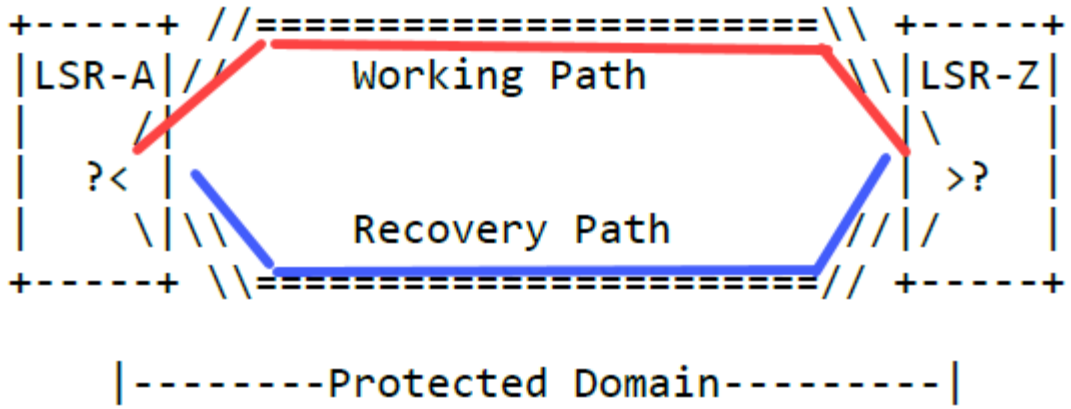


Figure 1: Protection switching control logic

Figure 1 describes the logical architecture of the protection switching control. The Local Request logic unit accepts the triggers from the OAM, external operator commands, and from the local control plane (when present) and determines the highest priority request. This high-priority request is passed to both the PSC Message generator, that will generate the appropriate protocol message to be sent to the far-end LSR, and the Global Request logic, that will cross-check this local request with the information received from the far-LSR. The Global Request logic then processes these two PSC requests that determines the highest priority request that is passed to the PSC Process logic. The PSC Process logic uses this input to determine what actions need to be taken, e.g. switching the Selector Bridge, and the current status of the protection domain." MPLS-TP Protection at 7–8.

"4.2.4. Active path indicator (Path)

No.	'821 Patent Claim 14	IETF MPLS-TP System
		The Path field of the PSC control SHALL be used to indicate which path the source MEP is currently using for data transmission. The MEP should compare the value of this bit with the path that is locally selected for data transmission to verify that there is no inconsistency between the two end-points of the protected domain. If an inconsistency is detected then an alarm should be raised. The following are the possible values: • 0: indicates that normal traffic is being transmitted on the Working path. • 1: indicates the Recovery path is being used to transmit the normal traffic from the Working path. 2-255: for future extensions" MPLS-TP Protection at 12–13. "Principles of Operation In all of the following sub-sections, assume a protected domain between LSR-A and LSR-Z, using paths W (working) and R (recovery) as shown in figure 4." MPLS-TP Protection at 13.  MPLS-TP Protection Figure 4 (annotation added). "4.3.1 PSC States 4.3.1.1 Normal State

No.	'821 Patent Claim 14	IETF MPLS-TP System
		When the protected domain has no special condition in effect, the ingress LSR SHOULD forward the user data along the working path, and, in the case of 1+1 protection, the Permanent Bridge will bridge the data to the recovery path as well. The receiving LSR SHOULD read the data from the working path. The ingress LSR MAY transmit a No Request PSC packet with the Path field set to 0 indicating that the normal data traffic should be read from the working path. 4.3.1.2. Protecting State When the protection mechanism has been triggered and the protected domain has performed a protection switch, the domain is in the protecting state. In this state the normal data traffic is transmitted and received on the recovery path. If the protection domain is currently in a protecting state, then the LSRs SHOULD NOT accept a Manual Switch request. If the protection domain is currently in a protecting state, and a Forced Switch is requested then the normal traffic SHALL continue to be transmitted on the recovery path even if the original protection trigger is cleared, and the Forced Switch condition will be signalled by the PSC messages. 4.3.1.3. Unavailable State When the recovery path is unavailable - either as a result of a Lockout operator command (see section 4.3.3), or as a result of a SF or SD detected on the recovery path (see section 4.3.4) - then the protection domain is in the unavailable state. In this state, the normal traffic is transmitted and received on the working path. While in unavailable state any event that would trigger a protection switching SHOULD be ignored with the following exception - If a Signal Degrade request is received, then protection switching will be activated only if the recovery path can guarantee a better signal than the working path. The protection domain will exit the unavailable state and revert to the normal state when, either the operator clears the Lockout command or the recovery path recovers from the signal fault or degraded situation.

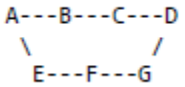
No.	'821 Patent Claim 14	IETF MPLS-TP System
		Both ends will resume sending the PCS packets over the recovery path, as a result of this recovery. 4.3.2. Failure or Degraded condition (Working path) If one of the LSRs (for example, LSR-A) detects a failure condition or a serious degradation condition on the working path that warrants invoking protection switching, then it SHOULD take the following actions: • (For 1:1 protection) Switch all traffic for LSR-Z to the recovery path only. • Transmit a PCS control packet, using GACH, with the appropriate Request code (either Signal fault or Signal degrade), the Fpath set to 1, to indicate that the fault/degrade was detected on the working path, and the Path set to 1, indicating that normal traffic is now being transmitted on the recovery path. • Verify that LSR-Z replies with a PCS control packet indicating that it has switched to the recovery path. If this is not received after 2 PSC cycles then send an alarm to the management system. When the far-end LSR (in this example LSR-Z) receives the PCS packet informing it that other LSR (LSR-A) has switched, it SHOULD perform the following actions: • Check priority of the request • Switch all traffic addressed to LSR-A to the recovery path only (for 1:1 protection). • Begin transmission of a PCS control packet, using GACH, with the appropriate Request code (either Signal fault or Signal degrade), the Fpath set to 1, to indicate that the fault/degrade was detected on the working path, and the Path set to 1, indicating that traffic is now being transmitted on the recovery path." MPLS-TP Protection at 13–15. "4.3.4. Failure or Degraded condition (Recovery path) If one of the LSRs (for example, LSR-A) detects a failure condition or a serious degradation condition on the recovery path, then it SHOULD take the following actions:

No.	'821 Patent Claim 14	IETF MPLS-TP System
		• Begin transmission of a PCS control packet with the appropriate Request code (either Signal fault or Signal degrade), the Fpath set to 0, to indicate that the fault/degrade was detected on the recovery path, and the Path set to 0, indicating that traffic is now being forwarded on the working path. Note that this will actually reach the far-end if this is a unidirectional fault or recovery path is possibly in a degraded situation. • The PSC control logic should go into Unavailable state. • All traffic MUST be transmitted on the working path for the duration of the SF/SD condition. When the far-end LSR (in this example LSR-Z) receives the PCS packet informing it that other LSR (LSR-A) has become Unavailable, it SHOULD perform the following actions: • Transmit all traffic on the working path for the duration of the SF/SD condition The PSC Control logic should go into Unavailable state." MPLS-TP Protection at 16.
14[d]	logic configured to determine a probability of concurrent failure of said working entity and said protection entity;	IETF MPLS-TP System discloses logic configured to determine a probability of concurrent failure of said working entity and said protection entity. For example, RFC 4872 discloses that the plurality of MPLS LSPs, i.e., transport entities, are selected in pairs of either working and protection LSPs or working and restoration LSPs. RFC 4872 further discloses that the plurality of MPLS LSPs, i.e., transport entities, are selected to necessarily be "resource disjoint," i.e., where multiple resources, such as links, nodes, or Shared Risk Link Groups (SRLGs) do not share common links or failure probabilities. Thus, at least under the apparent claim scope alleged by Orckit's Infringement Disclosures, this limitation is met. To the extent that the IETF MPLS-TP System is found to not meet this limitation, logic configured to determine a probability of concurrent failure of said working entity and said protection entity would have been obvious to a person having ordinary skill in the art, as explained below. "When a failure is detected by one or both endpoints of the LSP, both endpoints must select traffic from the other LSP. This action must be coordinated between node A and D. From this perspective, 1+1 bidirectional protection can be seen as a coordinated protection- switching mechanism between both endpoints.

No.	'821 Patent Claim 14	IETF MPLS-TP System
		Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 11. "1:1 Protection with Extra-Traffic The most common case of end-to-end 1:N protection is to establish, between the same endpoints, an end-to-end working LSP (thus, N = 1) and a dedicated end-to-end protecting LSP that are mutually link/ node/SRLG disjoint. This protects against working LSP failure(s)." RFC 4872 at 13. "An approach to reduce recovery resource requirements is to have protection LSPs sharing network resources when the working LSPs that they protect are physically (i.e., link, node, SRLG, etc.) disjoint. This mechanism is referred to as shared mesh restoration and is described in [RFC4426]. Shared-mesh restoration can be seen as a particular case of pre-planned LSP rerouting (see Section 8) that reduces the recovery resource requirements by allowing multiple protecting LSPs to share common link and node resources. Here also, the recovery resources for the protecting LSPs are pre-reserved during the provisioning phase, thus an explicit signaling action is required to activate (i.e., commit resource allocation at the data plane) a specific protecting LSP instantiated during the (pre-) provisioning phase. This requires restoration signaling along the protecting LSP." RFC 4872 at 20. "1. Introduction Generalized Multi-Protocol Label Switching (GMPLS) extends MPLS to include support for Layer-2 Switch Capable (L2SC), Time-Division Multiplex (TDM), Lambda Switch Capable (LSC), and Fiber Switch Capable (FSC) interfaces. GMPLS recovery uses control plane mechanisms (i.e., signaling, routing, and link management mechanisms) to support data plane fault recovery. Note that the analogous (data plane) fault detection mechanisms are required to be present in support of the control plane mechanisms. In this document, the term "recovery" is generically used to denote both protection and restoration; the specific terms "protection" and "restoration" are only used when differentiation is required. The subtle distinction between protection and restoration is made based on the resource allocation done during the recovery phase (see [RFC4427]).

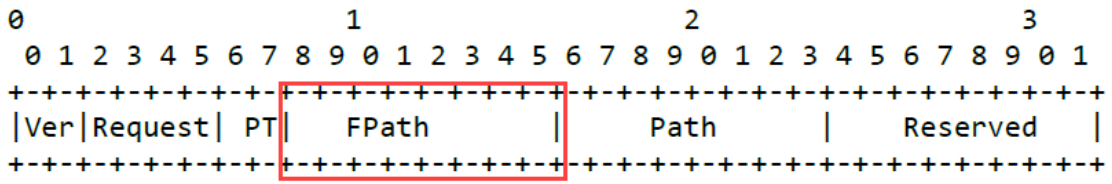
No.	'821 Patent Claim 14	IETF MPLS-TP System
		A functional description of GMPLS recovery is provided in [RFC4426] and should be considered as a companion document. The present document describes the protocol-specific procedures for GMPLS RSVP- TE (Resource ReSerVation Protocol - Traffic Engineering) signaling (see [RFC3473]) to support end-to-end recovery. End-to-end recovery refers to the recovery of an entire LSP from its head-end (ingress node endpoint) to its tail-end (egress node endpoint). With end-to- end recovery, working LSPs are assumed to be resource-disjoint (where a resource is a link, node, or Shared Risk Link Group (SRLG)) in the network so that they do not share any failure probability, but this is not mandatory. With respect to a given set of network resources, a pair of working/protecting LSPs SHOULD be resource disjoint in case of dedicated recovery type (see below). On the other hand, in case of shared recovery (see below), a group of working LSPs SHOULD be mutually resource-disjoint in order to allow for a (single and commonly) shared protecting LSP, itself resource-disjoint from each of the working LSPs. Note that resource disjointness is a necessary (but not sufficient) condition to ensure LSP recoverability. The present document addresses four types of end-to-end LSP recovery: 1) 1+1 (unidirectional/bidirectional) protection, 2) 1:N (N >= 1) LSP protection with extra-traffic, 3) pre-planned LSP rerouting without extra-traffic (including shared mesh), and 4) full LSP rerouting. 57) The simplest notion of end-to-end LSP protection is 1+1 unidirectional protection. Using this type of protection, a protecting LSP is signaled over a dedicated resource-disjoint alternate path to protect an associated working LSP. Normal traffic is simultaneously sent on both LSPs and a selector is used at the egress node to receive traffic from one of the LSPs. If a failure occurs along one of the LSPs, the egress node selects the traffic from the valid LSP. No coordination is required between the end nodes when a failure/switchover occurs. In 1+1 bidirectional protection, a protecting LSP is signaled over a dedicated resource-disjoint alternate path to protect the working LSP. Normal traffic is simultaneously sent on both LSPs (in both directions), and a selector is used at both ingress/egress nodes to receive traffic from the same LSP. This requires coordination between the end-nodes when switching to the protecting LSP.

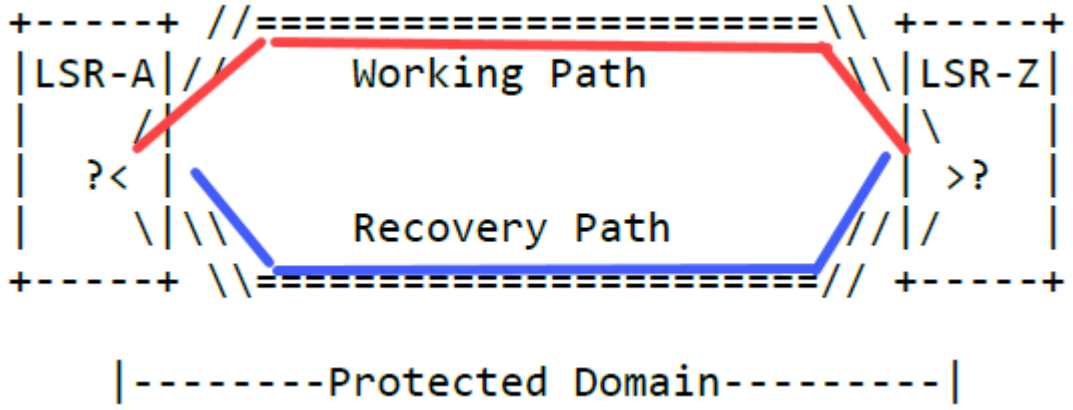
No.	'821 Patent Claim 14	IETF MPLS-TP System
		58) In 1:N (N >= 1) protection with extra-traffic, the protecting LSP is a fully provisioned and resource-disjoint LSP from the N working LSPs, that allows for carrying extra-traffic. The N working LSPs MAY be mutually resource-disjoint. Coordination between end-nodes is required when switching from one of the working LSPs to the protecting LSP. As the protecting LSP is fully provisioned, default operations during protection switching are specified for a protecting LSP carrying extra-traffic, but this is not mandatory. Note that M:N protection is out of scope of this document (though mechanisms it defines may be extended to cover it). 59) Pre-planned LSP rerouting (or restoration) relies on the establishment between the same pair of end-nodes of a working LSP and a protecting LSP that is link/node/SRLG disjoint from the working one. Here, the recovery resources for the protecting LSP are pre-reserved but explicit action is required to activate (i.e., commit resource allocation at the data plane) a specific protecting LSP instantiated during the (pre-)provisioning phase. Since the protecting LSP is not "active" (i.e., fully instantiated), it cannot carry any extra-traffic. This does not mean that the corresponding resources cannot be used by other LSPs. Therefore, this mechanism protects against working LSP(s) failure(s) but requires activation of the protecting LSP after working LSP failure occurrence. This requires restoration signaling along the protecting path. "Shared-mesh" restoration can be seen as a particular case of pre-planned LSP rerouting that reduces the recovery resource requirements by allowing multiple protecting LSPs to share common link and node resources. The recovery resources are pre-reserved but explicit action is required to activate (i.e., commit resource allocation at the data plane) a specific protecting LSP instantiated during the (pre-) provisioning phase. This procedure requires restoration signaling along the protecting path. Note that in both cases, bandwidth pre-reserved for a protecting (but not activated) LSP can be made available for carrying extra traffic. LSPs for extra-traffic (with lower holding priority than the protecting LSP) can then be established using the bandwidth pre-reserved for the protecting LSP. Also, any lower priority LSP that use the pre-reserved resources for the protecting LSP(s) must be preempted during the activation of the protecting LSP. 60) Full LSP rerouting (or restoration) switches normal traffic to an alternate LSP that is not even partially established until after the working LSP failure occurs. The new

No.	'821 Patent Claim 14	IETF MPLS-TP System
		alternate route is selected at the LSP head-end node, it may reuse resources of the failed LSP at intermediate nodes and may include additional intermediate nodes and/or links." " RFC 4872 at 3–5. "Consider the following network topology:  The paths [A,B,C,D] and [A,E,F,G,D] are node and link disjoint, ignoring the ingress/egress nodes A and D. A 1+1 protected path is established from A to D over [A,B,C,D] and [A,E,F,G,D], and traffic is transmitted simultaneously over both component paths (i.e., LSPs). During the provisioning phase, both LSPs are fully instantiated (and thus activated) so that no resource sharing can be done along the protecting LSP (nor can any extra-traffic be transported). It is also RECOMMENDED to set the N bit since no protection-switching signaling is assumed in this case. When a failure occurs (say, at node B) and is detected at end-node D, the receiver at D selects the normal traffic from the other LSP. From this perspective, 1+1 unidirectional protection can be seen as an uncoordinated protection-switching mechanism acting independently at both endpoints. Also, for the LSP under failure condition, it is RECOMMENDED to not set the Path_State_Removed Flag of the ERROR_SPEC object (see [RFC3473]) upon PathErr message generation. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 9. "When a failure is detected by one or both endpoints of the LSP, both endpoints must select traffic from the other LSP. This action must be coordinated between node A and D. From this perspective, 1+1 bidirectional protection can be seen as a coordinated protection- switching mechanism between both endpoints. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 11.

No.	'821 Patent Claim 14	IETF MPLS-TP System
		"(Full) LSP Rerouting LSP rerouting, on the other hand, switches normal traffic to an alternate LSP that is fully established only after failure occurrence. The new (alternate) route is selected at the LSP head- end and may reuse intermediate nodes included in the original route; it may also include additional intermediate nodes. For strict-hop routing, TE requirements can be directly applied to the route computation, and the failed node or link can be avoided. However, if the failure occurred within a loose-routed hop, the head-end node may not have enough information to reroute the LSP around the failure. Crankback signaling (see [CRANK]) and route exclusion techniques (see [RFC4874]) MAY be used in this case. The alternate route MAY be either computed on demand (that is, when the failure occurs; this is referred to as full LSP rerouting) or pre-computed and stored for use when the failure is reported. The latter offers faster restoration time. There is, however, a risk that the alternate route will become out of date through other changes in the network; this can be mitigated to some extent by periodic recalculation of idle alternate routes. (Full) LSP rerouting will be initiated by the head-end node that has either detected the LSP failure or received a Notify message and/or a PathErr message with the new error code/sub-code "Notify Error/LSP Locally Failed" for this LSP. The new LSP resources can be established using the make-before-break mechanism, where the new LSP is set up before the old LSP is torn down. This is done by using the mechanisms of the SESSION_ATTRIBUTE object and the Shared-Explicit (SE) reservation style (see [RFC3209]). Both the new and old LSPs can share resources at common nodes. Note that the make-before-break mechanism is not used to avoid disruption to the normal traffic flow (the latter has already been broken by the failure that is being repaired). However, it is valuable to retain the resources allocated on the original LSP that will be reused by the new alternate LSP." RFC 4872 at 25. "The procedure is as follows: 49) The initiating (source) node sends the normal traffic onto both the working and the protecting LSPs. Once completed, the source node sends reliably a Notify message to

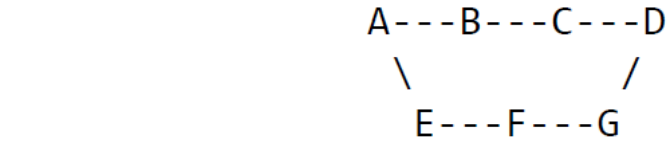
No.	'821 Patent Claim 14	IETF MPLS-TP System
		the destination with the new error code/sub-code "Notify Error/LSP Recovered" (Switchback Request). This Notify message includes the MESSAGE_ID object. The ACK_Desired flag MUST be set in this object to request the receiver to send an acknowledgment for the message (see [RFC2961]). 50) Upon receipt of this message, the destination selects the traffic from the working LSP. At the same time, it transmits the traffic onto both the working and protecting LSP. The destination then sends reliably a Notify message to the source confirming the completion of the operation. This message includes the MESSAGE_ID_ACK object to acknowledge reception of the received Notify message. This Notify message also includes the MESSAGE_ID object. The ACK_Desired flag MUST be set in this object to request the receiver to send an acknowledgment for the message (see [RFC2961]). 51) When the source node receives this Notify message, it switches to receive traffic from the working LSP. The source node then sends an Ack message to the destination node confirming that the LSP has been reverted." RFC 4872 at 27. For example, MPLS-TP Protection discloses an FPath field, that determines whether or not a specified LSP and/or pseudo-wire, i.e. network entity, is operational. MPLS-TP Protection discloses the use of a Label Switched Router to perform the described standard. A person of ordinary skill in the art would understand the disclosed Label Switched Router to include computer-readable data structures. At least under Orckit's Infringement Disclosures, this limitation is met. To the extent that the IETF MPLS-TP System is found to not meet this limitation, logic configured to determine a probability of concurrent failure of said working entity and said protection entity would have been obvious to a person having ordinary skill in the art, as explained below. "As noted in the architecture for Multi-Protocol Label Switching Transport Profile (MPLS-TP) [TPFwk], the overall architecture framework for MPLS-TP is based on a profile of the MPLS and Pseudowire (PW) procedures as specified for the MPLS and (MS-)PW architectures defined in [RFC3031], [RFC3985] and [RFC5085]. One of the basic survivability functions, pointed out by the Survivability Framework document [SurvivFwk],

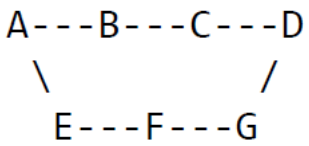
No.	'821 Patent Claim 14	IETF MPLS-TP System
		is that of simple and rapid protection switching mechanisms for Label Switched Paths (LSP) and Pseudo-wires (PW). Protection switching is a fully allocated survivability mechanism. It is fully allocated in the sense that the route and bandwidth of the recovery path is reserved for a selected working path or set of working paths. It provides a fast and simple survivability mechanism, that allows the network operator to easily grasp the active state of the network, compared to other survivability mechanisms. As specified in the Survivability Framework document [SurvivFwk], protection switching is applied to a protected domain. For the purposes of this document, we define the protected domain of a P2P LSP as consisting of two Label Switching Routers (LSR) and the transport paths that connect them. For a P2MP LSP the protection domain includes the root (or source) LSR, the destination (or sink) LSRs, and the transport paths that connect them." MPLS-TP Protection at 3.  0 1 2 3 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 |Ver|Request| PT| FPath | Path | Reserved | Figure 3: Format of the PSC control packet MPLS-TP Protection Figure 3, annotation added. "FPath: used to indicate the path that is reporting a failure condition, the possible values are described in section 4.2.3" MPLS-TP Protection at 11. "4.2.3. Path fault identifier (FPath) The Fpath field of the PSC control SHALL be used only in a Signal fault (0110) or Signal degrade (0101) control packet. Its value indicates on which path the signal anomaly was detected. The following are the possible values: • 0: indicates that the fault condition is on the Recovery path

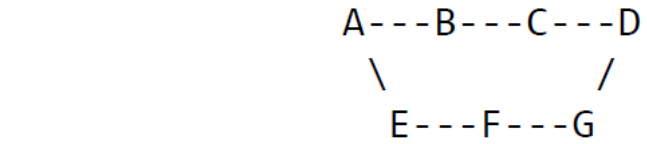
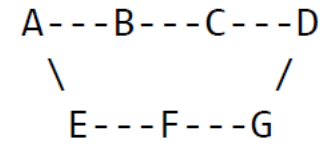
No.	'821 Patent Claim 14	IETF MPLS-TP System
		• 1: indicates that the fault condition is on the Working path • 2-255: for future extensions." MPLS-TP Protection at 12. "Principles of Operation In all of the following sub-sections, assume a protected domain between LSR-A and LSR-Z, using paths W (working) and R (recovery) as shown in figure 4." MPLS-TP Protection at 13.  MPLS-TP Protection Figure 4 (annotation added).
14[e]	logic configured to determine an entity cost of said plurality of transport entity descriptors: and	IETF MPLS-TP System discloses logic configured to determine an entity cost of said plurality of transport entity descriptors. *See* Claim 6.
14[f]	logic configured to reselect said working entity and said protection entity from said plurality of transport entity	IETF MPLS-TP System discloses logic configured to reselect said working entity and said protection entity from said plurality of transport entity descriptors upon a reselection event. *See* Claim 1[d].

No.	'821 Patent Claim 14	IETF MPLS-TP System
	descriptors upon a reselection event,	
14[g]	wherein said reselection event is selected from a group consisting of adding an entity to said plurality of transport entities, removing an entity from said plurality of transport entities, an operational status change for one of said plurality of transport entities, and a change in overall cost for one of said plurality of transport entities.	IETF MPLS-TP System discloses wherein said reselection event is selected from a group consisting of adding an entity to said plurality of transport entities, removing an entity from said plurality of transport entities, an operational status change for one of said plurality of transport entities, and a change in overall cost for one of said plurality of transport entities. *See* Claim 1[e].

No.	'821 Patent Claim 15	IETF MPLS-TP System
15	The system of claim 14 wherein said entity protection switch comprises a 1:1 switch.	IETF MPLS-TP System discloses the system of claim 14 wherein said entity protection switch comprises a 1:1 switch. For example, RFC 4872 discloses entity protection switches operable under both a 1:1 and 1+1 architecture. "7. 1:1 Protection with Extra-Traffic The most common case of end-to-end 1:N protection is to establish, between the same endpoints, an end-to-end working LSP (thus, N = 1) and a dedicated end-to-end protecting LSP that are mutually link/ node/SRLG disjoint. This protects against working LSP failure(s). The protecting LSP is used for switchover when the working LSP fails. GMPLS RSVP-TE

No.	'821 Patent Claim 15	IETF MPLS-TP System
		signaling allows for the pre-provisioning of protecting LSPs by indicating in the Path message (in the PROTECTION object; see Section 14) that the LSPs are of type protecting. Here, working and protecting LSPs are signaled as primary LSPs; both are fully instantiated during the provisioning phase. Although the resources for the protecting LSP are pre-allocated, preemptable traffic may be carried end-to-end using this LSP. Thus, the protecting LSP is capable of carrying extra-traffic with the caveat that this traffic will be preempted if the working LSP fails. The setup of the working LSP SHOULD indicate that the LSP head-end and tail-end node wish to receive Notify messages using the NOTIFY REQUEST object. The node upstream to the failure (upstream in terms of the direction an Path message traverses) SHOULD send a Notify message to the LSP head-end node, and the node downstream to the failure SHOULD send an Notify message to the LSP tail-end node. Upon receipt of the Notify messages, both the end-nodes MUST switch the (normal) traffic from the working LSP to the pre-configured protecting LSP (see Section 7.2). Moreover, some coordination is required if extra-traffic is carried over the end-to-end protecting LSP. Note that if the working and the protecting LSP are established between the same end-nodes, no further notification is required to indicate that the working LSPs are no longer protected. Consider the following topology:  The working LSP [A,B,C,D] could be protected by the protecting LSP [A,E,F,G,D]. Both LSPs are fully instantiated (resources are allocated for both working and protecting LSPs) and no resource sharing can be done along the protection path since the primary protecting LSP can carry extra-traffic. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 13–14. "6. 1+1 Bidirectional Protection 1+1 bidirectional protection is a scheme that provides end-to-end protection for bidirectional LSPs. Consider the following network topology:

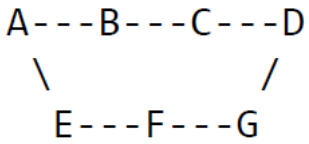
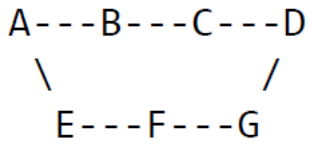
No.	'821 Patent Claim 15	IETF MPLS-TP System
		 The LSPs [A,B,C,D] and [A,E,F,G,D] are node and link disjoint, ignoring the ingress/egress nodes A and D. A bidirectional LSP is established from A to D over each path, and traffic is transmitted simultaneously over both LSPs. In this scheme, both endpoints must receive traffic over the same LSP. Note also that both LSPs are fully instantiated (and thus activated) so that no resource sharing can be done along the protection path (nor can any extra-traffic be transported). When a failure is detected by one or both endpoints of the LSP, both endpoints must select traffic from the other LSP. This action must be coordinated between node A and D. From this perspective, 1+1 bidirectional protection can be seen as a coordinated protection- switching mechanism between both endpoints. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 10–11. For example, RFC 4872 discloses entity protection switches operable under both a 1:1 and 1+1 architecture. "7. 1:1 Protection with Extra-Traffic The most common case of end-to-end 1:N protection is to establish, between the same endpoints, an end-to-end working LSP (thus, N = 1) and a dedicated end-to-end protecting LSP that are mutually link/ node/SRLG disjoint. This protects against working LSP failure(s). The protecting LSP is used for switchover when the working LSP fails. GMPLS RSVP-TE signaling allows for the pre-provisioning of protecting LSPs by indicating in the Path message (in the PROTECTION object; see Section 14) that the LSPs are of type protecting. Here, working and protecting LSPs are signaled as primary LSPs; both are fully instantiated during the provisioning phase. Although the resources for the protecting LSP are pre-allocated, preemptable traffic may be carried end-to-end using this LSP. Thus, the protecting LSP is capable of carrying extra-traffic with the caveat that this traffic will be preempted if the working LSP fails. The setup of the working LSP SHOULD indicate that the LSP head-end

No.	'821 Patent Claim 15	IETF MPLS-TP System
		and tail-end node wish to receive Notify messages using the NOTIFY REQUEST object. The node upstream to the failure (upstream in terms of the direction an Path message traverses) SHOULD send a Notify message to the LSP head-end node, and the node downstream to the failure SHOULD send an Notify message to the LSP tail-end node. Upon receipt of the Notify messages, both the end-nodes MUST switch the (normal) traffic from the working LSP to the pre-configured protecting LSP (see Section 7.2). Moreover, some coordination is required if extra-traffic is carried over the end-to-end protecting LSP. Note that if the working and the protecting LSP are established between the same end-nodes, no further notification is required to indicate that the working LSPs are no longer protected. Consider the following topology:  The working LSP [A,B,C,D] could be protected by the protecting LSP [A,E,F,G,D]. Both LSPs are fully instantiated (resources are allocated for both working and protecting LSPs) and no resource sharing can be done along the protection path since the primary protecting LSP can carry extra-traffic. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 13–14. "6. 1+1 Bidirectional Protection 1+1 bidirectional protection is a scheme that provides end-to-end protection for bidirectional LSPs. Consider the following network topology:  The LSPs [A,B,C,D] and [A,E,F,G,D] are node and link disjoint, ignoring the ingress/egress nodes A and D. A bidirectional LSP is established from A to D over each path, and traffic is

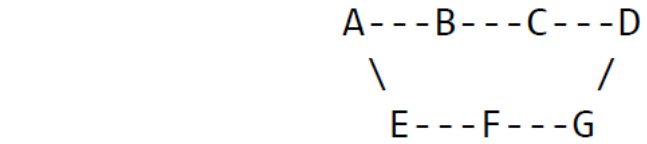
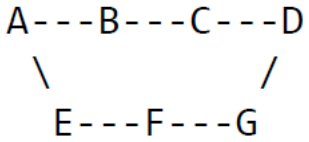
No.	'821 Patent Claim 15	IETF MPLS-TP System
		transmitted simultaneously over both LSPs. In this scheme, both endpoints must receive traffic over the same LSP. Note also that both LSPs are fully instantiated (and thus activated) so that no resource sharing can be done along the protection path (nor can any extra-traffic be transported). When a failure is detected by one or both endpoints of the LSP, both endpoints must select traffic from the other LSP. This action must be coordinated between node A and D. From this perspective, 1+1 bidirectional protection can be seen as a coordinated protection- switching mechanism between both endpoints. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 10–11. For example, MPLS-TP Protection discloses that its protection method can be used in both a 1:1 and a 1+1 architecture. "In 1+1 unidirectional architecture as presented in [SurvivFwk], a recovery transport path is dedicated to each working transport path. Normal traffic is bridged (as defined in [RFC4427])and fed to both the working and the recovery transport entities by a permanent bridge at the source of the protection domain. The sink of the protection domain selects which of the working or recovery entities to receive the traffic from, based on a predetermined criteria, e.g. server defect indication. When used for bidirectional switching the 1+1 protection architecture must also support a Protection State Coordination (PSC) protocol. This protocol is used to help synchronize the decisions of both ends of the protection domain in selecting the proper traffic flow. In the 1:1 architecture, a recovery transport path is dedicated to the working transport path of a single service. However, the normal traffic is transmitted only once, on either the working or the recovery path, by using a selector bridge at the source of the protection domain. A selector at the sink of the protection domain then selects the path that carries the normal traffic. Since the source and sink need to be coordinated to ensure that the selector bridge at both ends select the same path, this architecture must support a PSC protocol. The 1:n protection architecture extends this last architecture by sharing the recovery path amongst n services. Again, the recovery path is fully allocated and disjoint from any of the n working transport paths that it is being used to protect. The normal data traffic for each service

No.	'821 Patent Claim 15	IETF MPLS-TP System
		is transmitted only once, either on the normal working path for that service or, in cases that trigger protection switching (as defined in [SurvivFwk]), may be sent on the recovery path. It should be noted that in cases where multiple working path services have triggered protection switching that some services, dependent upon their Service Level Agreement (SLA), may not be transmitted as a result of limited resources on the recovery path. In this architecture there is a need for coordination of the protection switching, and in addition there is need for resource allocation negotiation. Due to the added complexity of this architecture, the procedures for this will be delayed to a different document and further study." MPLS-TP Protection at 4–5.

No.	'821 Patent Claim 16	IETF MPLS-TP System
16	The system of claim 14 wherein said entity protection switch comprises a 1+1 switch.	IETF MPLS-TP System discloses the system of claim 14 wherein said entity protection switch comprises a 1+1 switch. For example, RFC 4872 discloses entity protection switches operable under both a 1:1 and 1+1 architecture. "7. 1:1 Protection with Extra-Traffic The most common case of end-to-end 1:N protection is to establish, between the same endpoints, an end-to-end working LSP (thus, N = 1) and a dedicated end-to-end protecting LSP that are mutually link/ node/SRLG disjoint. This protects against working LSP failure(s). The protecting LSP is used for switchover when the working LSP fails. GMPLS RSVP-TE signaling allows for the pre-provisioning of protecting LSPs by indicating in the Path message (in the PROTECTION object; see Section 14) that the LSPs are of type protecting. Here, working and protecting LSPs are signaled as primary LSPs; both are fully instantiated during the provisioning phase. Although the resources for the protecting LSP are pre-allocated, preemptable traffic may be carried end-to-end using this LSP. Thus, the protecting LSP is capable of carrying extra-traffic with the caveat that this traffic will be preempted if the working LSP fails. The setup of the working LSP SHOULD indicate that the LSP head-end and tail-end node wish to receive Notify messages using the NOTIFY REQUEST object. The node upstream to the failure (upstream in terms of the direction an Path message traverses) SHOULD send a Notify message to the LSP head-end node, and the node downstream to the

No.	'821 Patent Claim 16	IETF MPLS-TP System
		failure SHOULD send an Notify message to the LSP tail-end node. Upon receipt of the Notify messages, both the end-nodes MUST switch the (normal) traffic from the working LSP to the pre-configured protecting LSP (see Section 7.2). Moreover, some coordination is required if extra-traffic is carried over the end-to-end protecting LSP. Note that if the working and the protecting LSP are established between the same end-nodes, no further notification is required to indicate that the working LSPs are no longer protected. Consider the following topology:  The working LSP [A,B,C,D] could be protected by the protecting LSP [A,E,F,G,D]. Both LSPs are fully instantiated (resources are allocated for both working and protecting LSPs) and no resource sharing can be done along the protection path since the primary protecting LSP can carry extra-traffic. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 13–14. "6. 1+1 Bidirectional Protection 1+1 bidirectional protection is a scheme that provides end-to-end protection for bidirectional LSPs. Consider the following network topology:  The LSPs [A,B,C,D] and [A,E,F,G,D] are node and link disjoint, ignoring the ingress/egress nodes A and D. A bidirectional LSP is established from A to D over each path, and traffic is transmitted simultaneously over both LSPs. In this scheme, both endpoints must receive traffic over the same LSP. Note also that both LSPs are fully instantiated (and thus activated)

No.	'821 Patent Claim 16	IETF MPLS-TP System
		so that no resource sharing can be done along the protection path (nor can any extra-traffic be transported). When a failure is detected by one or both endpoints of the LSP, both endpoints must select traffic from the other LSP. This action must be coordinated between node A and D. From this perspective, 1+1 bidirectional protection can be seen as a coordinated protection- switching mechanism between both endpoints. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 10–11. For example, RFC 4872 discloses entity protection switches operable under both a 1:1 and 1+1 architecture. "7. 1:1 Protection with Extra-Traffic The most common case of end-to-end 1:N protection is to establish, between the same endpoints, an end-to-end working LSP (thus, N = 1) and a dedicated end-to-end protecting LSP that are mutually link/ node/SRLG disjoint. This protects against working LSP failure(s). The protecting LSP is used for switchover when the working LSP fails. GMPLS RSVP-TE signaling allows for the pre-provisioning of protecting LSPs by indicating in the Path message (in the PROTECTION object; see Section 14) that the LSPs are of type protecting. Here, working and protecting LSPs are signaled as primary LSPs; both are fully instantiated during the provisioning phase. Although the resources for the protecting LSP are pre-allocated, preemptable traffic may be carried end-to-end using this LSP. Thus, the protecting LSP is capable of carrying extra-traffic with the caveat that this traffic will be preempted if the working LSP fails. The setup of the working LSP SHOULD indicate that the LSP head-end and tail-end node wish to receive Notify messages using the NOTIFY REQUEST object. The node upstream to the failure (upstream in terms of the direction an Path message traverses) SHOULD send a Notify message to the LSP head-end node, and the node downstream to the failure SHOULD send an Notify message to the LSP tail-end node. Upon receipt of the Notify messages, both the end-nodes MUST switch the (normal) traffic from the working LSP to the pre-configured protecting LSP (see Section 7.2). Moreover, some coordination is required if extra-traffic is carried over the end-to-end protecting LSP. Note that if the working and the

No.	'821 Patent Claim 16	IETF MPLS-TP System
		protecting LSP are established between the same end-nodes, no further notification is required to indicate that the working LSPs are no longer protected. Consider the following topology:  The working LSP [A,B,C,D] could be protected by the protecting LSP [A,E,F,G,D]. Both LSPs are fully instantiated (resources are allocated for both working and protecting LSPs) and no resource sharing can be done along the protection path since the primary protecting LSP can carry extra-traffic. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 13–14. "6. 1+1 Bidirectional Protection 1+1 bidirectional protection is a scheme that provides end-to-end protection for bidirectional LSPs. Consider the following network topology:  The LSPs [A,B,C,D] and [A,E,F,G,D] are node and link disjoint, ignoring the ingress/egress nodes A and D. A bidirectional LSP is established from A to D over each path, and traffic is transmitted simultaneously over both LSPs. In this scheme, both endpoints must receive traffic over the same LSP. Note also that both LSPs are fully instantiated (and thus activated) so that no resource sharing can be done along the protection path (nor can any extra-traffic be transported). When a failure is detected by one or both endpoints of the LSP, both endpoints must select traffic from the other LSP. This action must be coordinated between node A and D. From this perspective, 1+1 bidirectional protection can be seen as a coordinated protection- switching

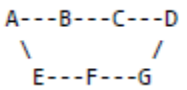
No.	'821 Patent Claim 16	IETF MPLS-TP System
		mechanism between both endpoints. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 10–11. For example, MPLS-TP Protection discloses that its protection method can be used in both a 1:1 and a 1+1 architecture. "In 1+1 unidirectional architecture as presented in [SurvivFwk], a recovery transport path is dedicated to each working transport path. Normal traffic is bridged (as defined in [RFC4427])and fed to both the working and the recovery transport entities by a permanent bridge at the source of the protection domain. The sink of the protection domain selects which of the working or recovery entities to receive the traffic from, based on a predetermined criteria, e.g. server defect indication. When used for bidirectional switching the 1+1 protection architecture must also support a Protection State Coordination (PSC) protocol. This protocol is used to help synchronize the decisions of both ends of the protection domain in selecting the proper traffic flow. In the 1:1 architecture, a recovery transport path is dedicated to the working transport path of a single service. However, the normal traffic is transmitted only once, on either the working or the recovery path, by using a selector bridge at the source of the protection domain. A selector at the sink of the protection domain then selects the path that carries the normal traffic. Since the source and sink need to be coordinated to ensure that the selector bridge at both ends select the same path, this architecture must support a PSC protocol. The 1:n protection architecture extends this last architecture by sharing the recovery path amongst n services. Again, the recovery path is fully allocated and disjoint from any of the n working transport paths that it is being used to protect. The normal data traffic for each service is transmitted only once, either on the normal working path for that service or, in cases that trigger protection switching (as defined in [SurvivFwk]), may be sent on the recovery path. It should be noted that in cases where multiple working path services have triggered protection switching that some services, dependent upon their Service Level Agreement (SLA), may not be transmitted as a result of limited resources on the recovery path. In this architecture there is a need for coordination of the protection switching, and in addition there is need for resource allocation negotiation. Due to the added complexity of this architecture, the

No.	'821 Patent Claim 16	IETF MPLS-TP System
		procedures for this will be delayed to a different document and further study." MPLS-TP Protection at 4–5.

No.	'821 Patent Claim 17	IETF MPLS-TP System
17[preamble]	Non-transitory computer readable media configured to perform a method comprising the steps of:	IETF MPLS-TP System discloses non-transitory computer readable media configured to perform a method comprising the steps of. For example, RFC 4872 discloses the use of network routers and data structures to perform its method. A person of ordinary skill in the art would understand the disclosed systems to include computer-readable data structures. "Note: an intermediate node processing a PPRO can derive SRLG identifiers from the local IGP-TE database using its Type 1, 2, or 4 subobject values as pointers to the corresponding TE Links (assuming each of them has an associated SRLG TE attribute)." RFC 4872 at 35. "1. Introduction Generalized Multi-Protocol Label Switching (GMPLS) extends MPLS to include support for Layer-2 Switch Capable (L2SC), Time-Division Multiplex (TDM), Lambda Switch Capable (LSC), and Fiber Switch Capable (FSC) interfaces. GMPLS recovery uses control plane mechanisms (i.e., signaling, routing, and link management mechanisms) to support data plane fault recovery. Note that the analogous (data plane) fault detection mechanisms are required to be present in support of the control plane mechanisms. In this document, the term "recovery" is generically used to denote both protection and restoration; the specific terms "protection" and "restoration" are only used when differentiation is required. The subtle distinction between protection and restoration is made based on the resource allocation done during the recovery phase (see [RFC4427]). A functional description of GMPLS recovery is provided in [RFC4426] and should be considered as a companion document. The present document describes the protocol-specific procedures for GMPLS RSVP- TE (Resource ReSerVation Protocol - Traffic Engineering) signaling (see [RFC3473]) to support end-to-end recovery. End-to-end recovery refers to the recovery of an entire LSP from its head-end (ingress node endpoint) to its tail-end (egress

No.	'821 Patent Claim 17	IETF MPLS-TP System
		node endpoint). With end-to- end recovery, working LSPs are assumed to be resource-disjoint (where a resource is a link, node, or Shared Risk Link Group (SRLG)) in the network so that they do not share any failure probability, but this is not mandatory. With respect to a given set of network resources, a pair of working/protecting LSPs SHOULD be resource disjoint in case of dedicated recovery type (see below). On the other hand, in case of shared recovery (see below), a group of working LSPs SHOULD be mutually resource-disjoint in order to allow for a (single and commonly) shared protecting LSP, itself resource-disjoint from each of the working LSPs. Note that resource disjointness is a necessary (but not sufficient) condition to ensure LSP recoverability. The present document addresses four types of end-to-end LSP recovery: 1) 1+1 (unidirectional/bidirectional) protection, 2) 1:N (N >= 1) LSP protection with extra-traffic, 3) pre-planned LSP rerouting without extra-traffic (including shared mesh), and 4) full LSP rerouting. 1) The simplest notion of end-to-end LSP protection is 1+1 unidirectional protection. Using this type of protection, a protecting LSP is signaled over a dedicated resource-disjoint alternate path to protect an associated working LSP. Normal traffic is simultaneously sent on both LSPs and a selector is used at the egress node to receive traffic from one of the LSPs. If a failure occurs along one of the LSPs, the egress node selects the traffic from the valid LSP. No coordination is required between the end nodes when a failure/switchover occurs. In 1+1 bidirectional protection, a protecting LSP is signaled over a dedicated resource-disjoint alternate path to protect the working LSP. Normal traffic is simultaneously sent on both LSPs (in both directions), and a selector is used at both ingress/egress nodes to receive traffic from the same LSP. This requires coordination between the end-nodes when switching to the protecting LSP. 2) In 1:N (N >= 1) protection with extra-traffic, the protecting LSP is a fully provisioned and resource-disjoint LSP from the N working LSPs, that allows for carrying extra-traffic. The N working LSPs MAY be mutually resource-disjoint. Coordination between end-nodes is required when switching from one of the working LSPs to the protecting LSP. As the protecting LSP is fully provisioned, default operations during protection switching are specified for a protecting LSP carrying extra-traffic, but this

No.	'821 Patent Claim 17	IETF MPLS-TP System
		is not mandatory. Note that M:N protection is out of scope of this document (though mechanisms it defines may be extended to cover it). 3) Pre-planned LSP rerouting (or restoration) relies on the establishment between the same pair of end-nodes of a working LSP and a protecting LSP that is link/node/SRLG disjoint from the working one. Here, the recovery resources for the protecting LSP are pre-reserved but explicit action is required to activate (i.e., commit resource allocation at the data plane) a specific protecting LSP instantiated during the (pre-)provisioning phase. Since the protecting LSP is not "active" (i.e., fully instantiated), it cannot carry any extra-traffic. This does not mean that the corresponding resources cannot be used by other LSPs. Therefore, this mechanism protects against working LSP(s) failure(s) but requires activation of the protecting LSP after working LSP failure occurrence. This requires restoration signaling along the protecting path. "Shared-mesh" restoration can be seen as a particular case of pre-planned LSP rerouting that reduces the recovery resource requirements by allowing multiple protecting LSPs to share common link and node resources. The recovery resources are pre-reserved but explicit action is required to activate (i.e., commit resource allocation at the data plane) a specific protecting LSP instantiated during the (pre-) provisioning phase. This procedure requires restoration signaling along the protecting path. Note that in both cases, bandwidth pre-reserved for a protecting (but not activated) LSP can be made available for carrying extra traffic. LSPs for extra-traffic (with lower holding priority than the protecting LSP) can then be established using the bandwidth pre-reserved for the protecting LSP. Also, any lower priority LSP that use the pre-reserved resources for the protecting LSP(s) must be preempted during the activation of the protecting LSP. 4) Full LSP rerouting (or restoration) switches normal traffic to an alternate LSP that is not even partially established until after the working LSP failure occurs. The new alternate route is selected at the LSP head-end node, it may reuse resources of the failed LSP at intermediate nodes and may include additional intermediate nodes and/or links." " RFC 4872 at 3–5. "Consider the following network topology:

No.	'821 Patent Claim 17	IETF MPLS-TP System
		 The paths [A,B,C,D] and [A,E,F,G,D] are node and link disjoint, ignoring the ingress/egress nodes A and D. A 1+1 protected path is established from A to D over [A,B,C,D] and [A,E,F,G,D], and traffic is transmitted simultaneously over both component paths (i.e., LSPs). During the provisioning phase, both LSPs are fully instantiated (and thus activated) so that no resource sharing can be done along the protecting LSP (nor can any extra-traffic be transported). It is also RECOMMENDED to set the N bit since no protection-switching signaling is assumed in this case. When a failure occurs (say, at node B) and is detected at end-node D, the receiver at D selects the normal traffic from the other LSP. From this perspective, 1+1 unidirectional protection can be seen as an uncoordinated protection-switching mechanism acting independently at both endpoints. Also, for the LSP under failure condition, it is RECOMMENDED to not set the Path_State_Removed Flag of the ERROR_SPEC object (see [RFC3473]) upon PathErr message generation. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 9. For example, MPLS-TP Protection discloses the use of a Label Switched Router to perform the described standard. A person of ordinary skill in the art would understand the disclosed Label Switched Router to include computer-readable data structures. "As noted in the architecture for Multi-Protocol Label Switching Transport Profile (MPLS-TP) [TPFwk], the overall architecture framework for MPLS-TP is based on a profile of the MPLS and Pseudowire (PW) procedures as specified for the MPLS and (MS-)PW architectures defined in [RFC3031], [RFC3985] and [RFC5085]. One of the basic survivability functions, pointed out by the Survivability Framework document [SurvivFwk], is that of simple and rapid protection switching mechanisms for Label Switched Paths (LSP) and Pseudo-wires (PW).

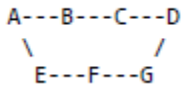
No.	'821 Patent Claim 17	IETF MPLS-TP System
		Protection switching is a fully allocated survivability mechanism. It is fully allocated in the sense that the route and bandwidth of the recovery path is reserved for a selected working path or set of working paths. It provides a fast and simple survivability mechanism, that allows the network operator to easily grasp the active state of the network, compared to other survivability mechanisms. As specified in the Survivability Framework document [SurvivFwk], protection switching is applied to a protected domain. For the purposes of this document, we define the protected domain of a P2P LSP as consisting of two Label Switching Routers (LSR) and the transport paths that connect them. For a P2MP LSP the protection domain includes the root (or source) LSR, the destination (or sink) LSRs, and the transport paths that connect them." MPLS-TP Protection at 3.
17[a]	providing a plurality of MPLS transport entities between a first endpoint and a second endpoint;	IETF MPLS-TP System discloses providing a plurality of MPLS transport entities between a first endpoint and a second endpoint. *See* Claim 1[a].
17[b]	determining an overall cost for each entity pair of said plurality of entities;	IETF MPLS-TP System discloses determining an overall cost for each entity pair of said plurality of entities. *See* Claim 1[b].
17[c]	selecting an entity pair from said plurality of transport entities based at least in part upon said overall cost; and	IETF MPLS-TP System discloses selecting an entity pair from said plurality of transport entities based at least in part upon said overall cost. *See* Claim 1[c].
17[d]	if an entity pair reselection event occurs, reselecting said entity pair from the group consisting of said entity pair and a	IETF MPLS-TP System discloses if an entity pair reselection event occurs, reselecting said entity pair from the group consisting of said entity pair and a replacement entity pair comprising at least one entity distinct from the entities of said entity pair. *See* Claim 1[d].

No.	'821 Patent Claim 17	IETF MPLS-TP System
	replacement entity pair comprising at least one entity distinct from the entities of said entity pair,	
17[e]	wherein said entity pair reselection event is selected from a group consisting of adding an entity to said plurality of transport entities, removing an entity from said plurality of transport entities, an operational status change for one of said plurality of transport entities, and a change in overall cost for one of said plurality of transport entities.	IETF MPLS-TP System discloses wherein said entity pair reselection event is selected from a group consisting of adding an entity to said plurality of transport entities, removing an entity from said plurality of transport entities, an operational status change for one of said plurality of transport entities, and a change in overall cost for one of said plurality of transport entities. *See* Claim 1[e].

No.	'821 Patent Claim 18	IETF MPLS-TP System
18[preamble]	The non-transitory computer readable media of claim 17, wherein said step of selecting an entity pair further comprises:	IETF MPLS-TP System discloses the non-transitory computer readable media of claim 17, wherein said step of selecting an entity pair further comprises. For example, RFC 4872 discloses the use of network routers and data structures to perform its method. A person of ordinary skill in the art would understand the disclosed systems to include computer-readable data structures.

No.	'821 Patent Claim 18	IETF MPLS-TP System
		"Note: an intermediate node processing a PPRO can derive SRLG identifiers from the local IGP-TE database using its Type 1, 2, or 4 subobject values as pointers to the corresponding TE Links (assuming each of them has an associated SRLG TE attribute)." RFC 4872 at 35. "1. Introduction Generalized Multi-Protocol Label Switching (GMPLS) extends MPLS to include support for Layer-2 Switch Capable (L2SC), Time-Division Multiplex (TDM), Lambda Switch Capable (LSC), and Fiber Switch Capable (FSC) interfaces. GMPLS recovery uses control plane mechanisms (i.e., signaling, routing, and link management mechanisms) to support data plane fault recovery. Note that the analogous (data plane) fault detection mechanisms are required to be present in support of the control plane mechanisms. In this document, the term "recovery" is generically used to denote both protection and restoration; the specific terms "protection" and "restoration" are only used when differentiation is required. The subtle distinction between protection and restoration is made based on the resource allocation done during the recovery phase (see [RFC4427]). A functional description of GMPLS recovery is provided in [RFC4426] and should be considered as a companion document. The present document describes the protocol-specific procedures for GMPLS RSVP- TE (Resource ReSerVation Protocol - Traffic Engineering) signaling (see [RFC3473]) to support end-to-end recovery. End-to-end recovery refers to the recovery of an entire LSP from its head-end (ingress node endpoint) to its tail-end (egress node endpoint). With end-to- end recovery, working LSPs are assumed to be resource-disjoint (where a resource is a link, node, or Shared Risk Link Group (SRLG)) in the network so that they do not share any failure probability, but this is not mandatory. With respect to a given set of network resources, a pair of working/protecting LSPs SHOULD be resource disjoint in case of dedicated recovery type (see below). On the other hand, in case of shared recovery (see below), a group of working LSPs SHOULD be mutually resource-disjoint in order to allow for a (single and commonly) shared protecting LSP, itself resource-disjoint from each of the working LSPs. Note that resource disjointness is a necessary (but not sufficient) condition to ensure LSP recoverability. The present document addresses four types of end-to-end LSP recovery: 1) 1+1 (unidirectional/bidirectional) protection, 2) 1:N (N >= 1) LSP protection with extra-traffic, 3)

No.	'821 Patent Claim 18	IETF MPLS-TP System
		pre-planned LSP rerouting without extra-traffic (including shared mesh), and 4) full LSP rerouting. 1) The simplest notion of end-to-end LSP protection is 1+1 unidirectional protection. Using this type of protection, a protecting LSP is signaled over a dedicated resource-disjoint alternate path to protect an associated working LSP. Normal traffic is simultaneously sent on both LSPs and a selector is used at the egress node to receive traffic from one of the LSPs. If a failure occurs along one of the LSPs, the egress node selects the traffic from the valid LSP. No coordination is required between the end nodes when a failure/switchover occurs. In 1+1 bidirectional protection, a protecting LSP is signaled over a dedicated resource-disjoint alternate path to protect the working LSP. Normal traffic is simultaneously sent on both LSPs (in both directions), and a selector is used at both ingress/egress nodes to receive traffic from the same LSP. This requires coordination between the end-nodes when switching to the protecting LSP. 2) In 1:N (N >= 1) protection with extra-traffic, the protecting LSP is a fully provisioned and resource-disjoint LSP from the N working LSPs, that allows for carrying extra-traffic. The N working LSPs MAY be mutually resource-disjoint. Coordination between end-nodes is required when switching from one of the working LSPs to the protecting LSP. As the protecting LSP is fully provisioned, default operations during protection switching are specified for a protecting LSP carrying extra-traffic, but this is not mandatory. Note that M:N protection is out of scope of this document (though mechanisms it defines may be extended to cover it). 3) Pre-planned LSP rerouting (or restoration) relies on the establishment between the same pair of end-nodes of a working LSP and a protecting LSP that is link/node/SRLG disjoint from the working one. Here, the recovery resources for the protecting LSP are pre-reserved but explicit action is required to activate (i.e., commit resource allocation at the data plane) a specific protecting LSP instantiated during the (pre-)provisioning phase. Since the protecting LSP is not "active" (i.e., fully instantiated), it cannot carry any extra-traffic. This does not mean that the corresponding resources cannot be used by other LSPs. Therefore, this mechanism protects against working LSP(s) failure(s) but requires activation of the protecting LSP after working LSP failure occurrence.

No.	'821 Patent Claim 18	IETF MPLS-TP System
		This requires restoration signaling along the protecting path. "Shared-mesh" restoration can be seen as a particular case of pre-planned LSP rerouting that reduces the recovery resource requirements by allowing multiple protecting LSPs to share common link and node resources. The recovery resources are pre-reserved but explicit action is required to activate (i.e., commit resource allocation at the data plane) a specific protecting LSP instantiated during the (pre-) provisioning phase. This procedure requires restoration signaling along the protecting path. Note that in both cases, bandwidth pre-reserved for a protecting (but not activated) LSP can be made available for carrying extra traffic. LSPs for extra-traffic (with lower holding priority than the protecting LSP) can then be established using the bandwidth pre-reserved for the protecting LSP. Also, any lower priority LSP that use the pre-reserved resources for the protecting LSP(s) must be preempted during the activation of the protecting LSP. 4) Full LSP rerouting (or restoration) switches normal traffic to an alternate LSP that is not even partially established until after the working LSP failure occurs. The new alternate route is selected at the LSP head-end node, it may reuse resources of the failed LSP at intermediate nodes and may include additional intermediate nodes and/or links." " RFC 4872 at 3–5. "Consider the following network topology:  The paths [A,B,C,D] and [A,E,F,G,D] are node and link disjoint, ignoring the ingress/egress nodes A and D. A 1+1 protected path is established from A to D over [A,B,C,D] and [A,E,F,G,D], and traffic is transmitted simultaneously over both component paths (i.e., LSPs). During the provisioning phase, both LSPs are fully instantiated (and thus activated) so that no resource sharing can be done along the protecting LSP (nor can any extra-traffic be transported). It is also RECOMMENDED to set the N bit since no protection-switching signaling is assumed in this case.

No.	'821 Patent Claim 18	IETF MPLS-TP System
		When a failure occurs (say, at node B) and is detected at end-node D, the receiver at D selects the normal traffic from the other LSP. From this perspective, 1+1 unidirectional protection can be seen as an uncoordinated protection-switching mechanism acting independently at both endpoints. Also, for the LSP under failure condition, it is RECOMMENDED to not set the Path_State_Removed Flag of the ERROR_SPEC object (see [RFC3473]) upon PathErr message generation. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 9. For example, MPLS-TP Protection discloses the use of a Label Switched Router to perform the described standard. A person of ordinary skill in the art would understand the disclosed Label Switched Router to include computer-readable data structures. "As noted in the architecture for Multi-Protocol Label Switching Transport Profile (MPLS-TP) [TPFwk], the overall architecture framework for MPLS-TP is based on a profile of the MPLS and Pseudowire (PW) procedures as specified for the MPLS and (MS-)PW architectures defined in [RFC3031], [RFC3985] and [RFC5085]. One of the basic survivability functions, pointed out by the Survivability Framework document [SurvivFwk], is that of simple and rapid protection switching mechanisms for Label Switched Paths (LSP) and Pseudo-wires (PW). Protection switching is a fully allocated survivability mechanism. It is fully allocated in the sense that the route and bandwidth of the recovery path is reserved for a selected working path or set of working paths. It provides a fast and simple survivability mechanism, that allows the network operator to easily grasp the active state of the network, compared to other survivability mechanisms. As specified in the Survivability Framework document [SurvivFwk], protection switching is applied to a protected domain. For the purposes of this document, we define the protected domain of a P2P LSP as consisting of two Label Switching Routers (LSR) and the transport paths that connect them. For a P2MP LSP the protection domain includes the root (or source) LSR, the destination (or sink) LSRs, and the transport paths that connect them." MPLS-TP Protection at 3.

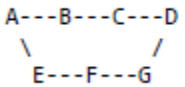
No.	'821 Patent Claim 18	IETF MPLS-TP System
18[a]	selecting a working entity from said plurality of transport entities;	IETF MPLS-TP System discloses [-]. *See* Claim 2[a].
18[b]	selecting a protection entity from said plurality of transport entities; and	IETF MPLS-TP System discloses [-]. *See* Claim 2[b].
18[c]	selecting an active entity from the set consisting of said working entity and said protection entity.	IETF MPLS-TP System discloses [-]. *See* Claim 3.

No.	'821 Patent Claim 19	IETF MPLS-TP System
19	The non-transitory readable media of claim 18, wherein said step of selecting an entity pair further comprises minimizing an overall cost function.	IETF MPLS-TP System discloses [-]. *See* Claim 4.

No.	'821 Patent Claim 20	IETF MPLS-TP System
20[preamble]	The non-transitory readable media of claim 19, wherein said overall cost function comprises:	IETF MPLS-TP System discloses [-]. For example, RFC 4872 discloses the use of network routers and data structures to perform its method. A person of ordinary skill in the art would understand the disclosed systems to include computer-readable data structures.

No.	'821 Patent Claim 20	IETF MPLS-TP System
		"Note: an intermediate node processing a PPRO can derive SRLG identifiers from the local IGP-TE database using its Type 1, 2, or 4 subobject values as pointers to the corresponding TE Links (assuming each of them has an associated SRLG TE attribute)." RFC 4872 at 35. "1. Introduction Generalized Multi-Protocol Label Switching (GMPLS) extends MPLS to include support for Layer-2 Switch Capable (L2SC), Time-Division Multiplex (TDM), Lambda Switch Capable (LSC), and Fiber Switch Capable (FSC) interfaces. GMPLS recovery uses control plane mechanisms (i.e., signaling, routing, and link management mechanisms) to support data plane fault recovery. Note that the analogous (data plane) fault detection mechanisms are required to be present in support of the control plane mechanisms. In this document, the term "recovery" is generically used to denote both protection and restoration; the specific terms "protection" and "restoration" are only used when differentiation is required. The subtle distinction between protection and restoration is made based on the resource allocation done during the recovery phase (see [RFC4427]). A functional description of GMPLS recovery is provided in [RFC4426] and should be considered as a companion document. The present document describes the protocol-specific procedures for GMPLS RSVP- TE (Resource ReSerVation Protocol - Traffic Engineering) signaling (see [RFC3473]) to support end-to-end recovery. End-to-end recovery refers to the recovery of an entire LSP from its head-end (ingress node endpoint) to its tail-end (egress node endpoint). With end-to- end recovery, working LSPs are assumed to be resource-disjoint (where a resource is a link, node, or Shared Risk Link Group (SRLG)) in the network so that they do not share any failure probability, but this is not mandatory. With respect to a given set of network resources, a pair of working/protecting LSPs SHOULD be resource disjoint in case of dedicated recovery type (see below). On the other hand, in case of shared recovery (see below), a group of working LSPs SHOULD be mutually resource-disjoint in order to allow for a (single and commonly) shared protecting LSP, itself resource-disjoint from each of the working LSPs. Note that resource disjointness is a necessary (but not sufficient) condition to ensure LSP recoverability. The present document addresses four types of end-to-end LSP recovery: 1) 1+1 (unidirectional/bidirectional) protection, 2) 1:N (N >= 1) LSP protection with extra-traffic, 3)

No.	'821 Patent Claim 20	IETF MPLS-TP System
		pre-planned LSP rerouting without extra-traffic (including shared mesh), and 4) full LSP rerouting. 1) The simplest notion of end-to-end LSP protection is 1+1 unidirectional protection. Using this type of protection, a protecting LSP is signaled over a dedicated resource-disjoint alternate path to protect an associated working LSP. Normal traffic is simultaneously sent on both LSPs and a selector is used at the egress node to receive traffic from one of the LSPs. If a failure occurs along one of the LSPs, the egress node selects the traffic from the valid LSP. No coordination is required between the end nodes when a failure/switchover occurs. In 1+1 bidirectional protection, a protecting LSP is signaled over a dedicated resource-disjoint alternate path to protect the working LSP. Normal traffic is simultaneously sent on both LSPs (in both directions), and a selector is used at both ingress/egress nodes to receive traffic from the same LSP. This requires coordination between the end-nodes when switching to the protecting LSP. 2) In 1:N (N >= 1) protection with extra-traffic, the protecting LSP is a fully provisioned and resource-disjoint LSP from the N working LSPs, that allows for carrying extra-traffic. The N working LSPs MAY be mutually resource-disjoint. Coordination between end-nodes is required when switching from one of the working LSPs to the protecting LSP. As the protecting LSP is fully provisioned, default operations during protection switching are specified for a protecting LSP carrying extra-traffic, but this is not mandatory. Note that M:N protection is out of scope of this document (though mechanisms it defines may be extended to cover it). 3) Pre-planned LSP rerouting (or restoration) relies on the establishment between the same pair of end-nodes of a working LSP and a protecting LSP that is link/node/SRLG disjoint from the working one. Here, the recovery resources for the protecting LSP are pre-reserved but explicit action is required to activate (i.e., commit resource allocation at the data plane) a specific protecting LSP instantiated during the (pre-)provisioning phase. Since the protecting LSP is not "active" (i.e., fully instantiated), it cannot carry any extra-traffic. This does not mean that the corresponding resources cannot be used by other LSPs. Therefore, this mechanism protects against working LSP(s) failure(s) but requires activation of the protecting LSP after working LSP failure occurrence.

No.	'821 Patent Claim 20	IETF MPLS-TP System
		This requires restoration signaling along the protecting path. "Shared-mesh" restoration can be seen as a particular case of pre-planned LSP rerouting that reduces the recovery resource requirements by allowing multiple protecting LSPs to share common link and node resources. The recovery resources are pre-reserved but explicit action is required to activate (i.e., commit resource allocation at the data plane) a specific protecting LSP instantiated during the (pre-) provisioning phase. This procedure requires restoration signaling along the protecting path. Note that in both cases, bandwidth pre-reserved for a protecting (but not activated) LSP can be made available for carrying extra traffic. LSPs for extra-traffic (with lower holding priority than the protecting LSP) can then be established using the bandwidth pre-reserved for the protecting LSP. Also, any lower priority LSP that use the pre-reserved resources for the protecting LSP(s) must be preempted during the activation of the protecting LSP. 4) Full LSP rerouting (or restoration) switches normal traffic to an alternate LSP that is not even partially established until after the working LSP failure occurs. The new alternate route is selected at the LSP head-end node, it may reuse resources of the failed LSP at intermediate nodes and may include additional intermediate nodes and/or links." " RFC 4872 at 3–5. "Consider the following network topology:  The paths [A,B,C,D] and [A,E,F,G,D] are node and link disjoint, ignoring the ingress/egress nodes A and D. A 1+1 protected path is established from A to D over [A,B,C,D] and [A,E,F,G,D], and traffic is transmitted simultaneously over both component paths (i.e., LSPs). During the provisioning phase, both LSPs are fully instantiated (and thus activated) so that no resource sharing can be done along the protecting LSP (nor can any extra-traffic be transported). It is also RECOMMENDED to set the N bit since no protection-switching signaling is assumed in this case.

No.	'821 Patent Claim 20	IETF MPLS-TP System
		When a failure occurs (say, at node B) and is detected at end-node D, the receiver at D selects the normal traffic from the other LSP. From this perspective, 1+1 unidirectional protection can be seen as an uncoordinated protection-switching mechanism acting independently at both endpoints. Also, for the LSP under failure condition, it is RECOMMENDED to not set the Path_State_Removed Flag of the ERROR_SPEC object (see [RFC3473]) upon PathErr message generation. Note: it is necessary that both paths are SRLG disjoint to ensure recoverability; otherwise, a single failure may impact both working and protecting LSPs." RFC 4872 at 9. For example, MPLS-TP Protection discloses the use of a Label Switched Router to perform the described standard. A person of ordinary skill in the art would understand the disclosed Label Switched Router to include computer-readable data structures. "As noted in the architecture for Multi-Protocol Label Switching Transport Profile (MPLS-TP) [TPFwk], the overall architecture framework for MPLS-TP is based on a profile of the MPLS and Pseudowire (PW) procedures as specified for the MPLS and (MS-)PW architectures defined in [RFC3031], [RFC3985] and [RFC5085]. One of the basic survivability functions, pointed out by the Survivability Framework document [SurvivFwk], is that of simple and rapid protection switching mechanisms for Label Switched Paths (LSP) and Pseudo-wires (PW). Protection switching is a fully allocated survivability mechanism. It is fully allocated in the sense that the route and bandwidth of the recovery path is reserved for a selected working path or set of working paths. It provides a fast and simple survivability mechanism, that allows the network operator to easily grasp the active state of the network, compared to other survivability mechanisms. As specified in the Survivability Framework document [SurvivFwk], protection switching is applied to a protected domain. For the purposes of this document, we define the protected domain of a P2P LSP as consisting of two Label Switching Routers (LSR) and the transport paths that connect them. For a P2MP LSP the protection domain includes the root (or source) LSR, the destination (or sink) LSRs, and the transport paths that connect them." MPLS-TP Protection at 3.

No.	'821 Patent Claim 20	IETF MPLS-TP System
20[a]	minimizing a probability of concurrent failure of said protection entity and said working entity; and	IETF MPLS-TP System discloses [-]. *See* Claim 5.
20[b]	a predefined metric selected from the group consisting of interior gateway protocol (IGP) and traffic engineering (TE).	IETF MPLS-TP System discloses [-]. *See* Claim 7.