

UNITED STATES PATENT AND TRADEMARK OFFICE
BEFORE THE PATENT TRIAL AND APPEAL BOARD

CISCO SYSTEMS, INC,

Petitioner,

-vs-

CORRIGENT CORPORATION,
Patent Owner

Case No. IPR2023-00447
U.S. Patent No. 9,118,602

PETITIONER'S DEMONSTRATIVES
PTAB ORAL ARGUMENT
May 7, 2024

DEMONSTRATIVE EXHIBIT - NOT EVIDENCE




'602 Patent – Field of the Invention

FIELD OF THE INVENTION

The present invention relates generally to communication networks, and particularly to methods and systems for performing link aggregation in tunneled networks. 20

Paper 1 at 2
EX1001, Col. 1:19-21

'602 Patent – Figure 1A (Annotated)



Paper 1 at 4

'602 Patent Embodiment

Embodiments of the present invention provide tunnel pro-
30 visioning with link aggregation. Briefly described, a first
aspect of the present invention is directed to a method for
assigning and utilizing an Ethernet physical data port in an
Ethernet Link Aggregation Group (LAG) in a Multi-Protocol
Label Switching (MPLS) network. The method includes the
35 steps of selecting, by a first MPLS/LAG switch, from a plu-
rality of physical data ports in said first MPLS/LAG switch, a
single physical tunnel port which meets a bandwidth require-
ment of a network tunnel, wherein said single physical tunnel
port has a port serial number, assigning, by said first MPLS/
40 LAG switch, said single physical tunnel port to said network
tunnel, preparing, by said first MPLS/LAG switch, a data
packet label by which said single physical tunnel port may be
identified, receiving, by said first MPLS/LAG switch, a data
packet including said data packet label at said first MPLS/
45 LAG switch, and switching, by said first MPLS/LAG switch,
said data packet to said single physical tunnel port according
to said serial port number, and sending said data packet to a
second MPLS/LAG switch via said single physical tunnel
port.

Paper 1 at 3-5
EX1001, Col. 2:29-49

DEMONSTRATIVE EXHIBIT - NOT EVIDENCE

4



'602 Patent Embodiment

50 A second aspect of the present invention includes a method for tagging a packet for transport through an Ethernet physical data port in an Ethernet Link Aggregation Group (LAG) located downstream from a preceding node in a Multi-Protocol Label Switching (MPLS) network tunnel employing
55 Resource Reservation Protocol Traffic Engineering (RSVP-TE) tunnel provisioning. The method includes the steps of sending downstream to a MPLS/LAG switch, by the preceding node, an RSVP-TE PATH message, the RSVP-TE PATH message including a LABEL_REQUEST object requesting a
60 network tunnel, and a bandwidth requirement, receiving, by the preceding node, a packet label sent by the MPLS/LAG switch, said packet label including a serial number of a single physical tunnel port in the MPLS/LAG switch, and wherein the single physical tunnel port meets the bandwidth requirement,
65 attaching by the preceding node, said data packet label to a data packet sent by the preceding node, wherein the data packet is intended to be forwarded via the single physical

3

tunnel port, and sending downstream to the MPLS/LAG switch, by the preceding node, a data packet including said data packet label.

Paper 21 at 7
EX1001, Col. 2:50-3:3

DEMONSTRATIVE EXHIBIT - NOT EVIDENCE

5



'602 Patent Embodiment

A third aspect of the present invention includes an apparatus for assigning and utilizing an Ethernet physical data port in an Ethernet Link Aggregation Group (LAG) located downstream from a preceding node in a Multi-Protocol Label Switching (MPLS) network tunnel employing Resource Reservation Protocol Traffic Engineering (RSVP-TE) tunnel provisioning, the LAG having a first MPLS/LAG switch and a second MPLS/LAG switch. The apparatus includes a processor, which is configured to receive a tunnel configuration request message from the preceding node make a selection from a plurality of physical data ports in said first MPLS/LAG switch, of a single physical tunnel port which meets a bandwidth requirement of the network tunnel, wherein said single physical tunnel port has a port serial number, assign said single physical tunnel port to the network tunnel, prepare a data packet label, by which said single physical tunnel port may be identified, dedicate a sub-set of bits of said data packet label to encode said port serial number of said single physical tunnel port, and publish said packet label upstream.

Paper 21 at 7
EX1001, Col. 3:1-22

DEMONSTRATIVE EXHIBIT - NOT EVIDENCE

6



'602 Patent - MPLS

20 As part of the MPLS tunnel provisioning process (which is described in RFC 3031) each LSR along tunnel 28 attaches an MPLS label to the packets it transmits downstream to the next LSR, identifying the packets that belong to tunnel 28. Thus, in the example shown in FIG. 1A, MPLS/LAG switch A
25 receives packets to which MPLS labels have been attached at a preceding node 29, which is an LSR located in MPLS network A, upstream from MPLS/LAG switch A. In most MPLS implementations, a given MPLS label is used within a single hop along the tunnel. In this case node 29 is the LSR
30 that immediately precedes MPLS/LAG switch A in tunnel 28. In other MPLS implementations, the same label may be used for several adjacent hops of the tunnel, so that there may be one or more additional LSRs located between node 29 and switch A.

EX1001, Col. 4:20-34

'602 Patent - MPLS

MPLS/LAG switch **26** may be implemented using a network processor, which is programmed in software to carry out the functions described herein and is coupled to suitable hardware for interfacing with the MPLS network and Ethernet ports. Switch **26** may either include a standalone unit or may alternatively be integrated with other computing functions of the network processor. Some or all of the functions of switch **26** can also be implemented using a suitable general-purpose computer, a programmable logic device, an application-specific integrated circuit (ASIC) or a combination of such elements.

EX1001, Col. 5:12-22

'602 Patent - Figure 2

FIG. 2 is a block diagram that schematically illustrates an MPLS packet encapsulated into an Ethernet II format, in accordance with an embodiment of the present invention. The Ethernet frame includes a MAC (Media Access Control) destination address field 40, a MAC source address field 42, and an Ethertype identifier field 46 used to identify the specific protocol. The frame ends with a frame check sequence (FCS) field 48 used for error detection. The encapsulated MPLS packet include an MPLS header 50, which has an MPLS label field 52. An MPLS payload field 54 includes the message body, containing information transmitted by the MPLS packet. In some implementations, the frame also includes optional VLAN Ethertype and VLAN tag fields.

EX1001, Col. 5:23-35

'602 Patent – Figure 3

The method of FIG. 3 begins when the preceding node asks to establish a part of tunnel 28 (having one or more hops) for sending MPLS packets to MPLS/LAG switch 26 A. The preceding node requests and then receives the MPLS label, which it will subsequently attach to all packets that are sent to MPLS/LAG switch 26 labeled A. The preceding node sends downstream an RSVP-TE PATH message augmented with a LABEL_REQUEST object, as defined by RSVP-TE, to MPLS/LAG switch A, at a label requesting step 60. The PATH message typically includes information regarding service properties that are requested for tunnel 28. The service properties may include a guaranteed bandwidth (sometimes denoted CIR—Committed Information Rate) and a peak bandwidth (sometimes denoted PIR—Peak Information Rate), as well as a requested CoS (Class of Service—a measure of packet priority).

Paper 1 at 4
EX1001, Col. 6:3-19

DEMONSTRATIVE EXHIBIT - NOT EVIDENCE

10



The '602 Patent Addresses Bandwidth Concerns

20 CAC processor 32 of switch A receives the PATH message
and extracts the requested service properties. The CAC pro-
cessor examines the available bandwidth of all ports 24 in
LAG group 25 and selects a single physical port (“the selected
physical port”) on which to allocate bandwidth for MPLS
25 tunnel 28, responsively to the requested service properties, at
a port selection step 62. The selected physical port should be
capable of providing sufficient peak and average bandwidths,
as requested by the preceding node (and, originally, by the
ingress node).

Paper 1 at 4
EX1001, Col. 6:20-29

The '602 Patent Addresses Bandwidth Concerns

Regardless of the selection criterion used, the results of step **62** are that (1) a single physical port is explicitly selected and assigned to MPLS tunnel **28**, and (2) sufficient bandwidth is allocated to tunnel **28**, considering only the available bandwidth of the selected physical port, rather than the total available bandwidth of LAG group **25**. All packets belonging to tunnel **28** will be switched through the same selected physical port, using the port coding technique described herein below.

Paper 1 at 4
EX1001, Col. 6:46-53

'602 Patent – MPLS Packet Labeling

RSVP-TE processor of switch **26** sends the generated MPLS label upstream to the preceding node, using an RSVP- 5
TE RESV message augmented with a LABEL object, at a
label sending step **66**. At this stage, the part of tunnel **28**
between the preceding node and switch A is provisioned and
ready for use. The preceding node attaches the aforementioned MPLS label to all subsequent MPLS packets that it 10
sends downstream through tunnel **28** to MPLS/LAG switch
A, at a packet sending step **68**.

Paper 1 at 4
EX1001, Col. 7:4-12

Claim 1 of the '602 Patent

1. A method for assigning and utilizing an Ethernet physical data port in an Ethernet Link Aggregation Group (LAG) in a Multi-Protocol Label Switching (MPLS) network, the method comprising the steps of:

assigning, by a first MPLS/LAG switch, a single physical tunnel port of a LAG to a network tunnel, wherein the single physical tunnel port of the LAG meets a bandwidth requirement of the network tunnel, and wherein said single physical tunnel port of the LAG has a port serial number;

dedicating a sub-set of bits in a data packet label prepared by the first MPLS/LAG switch to encode said port serial number of said single physical tunnel port of the LAG into the data packet label;

sending, by said first MPLS/LAG switch, the data packet label, in which said port serial number of said single physical tunnel port is encoded, to a preceding node;

receiving from the preceding node, by said first MPLS/LAG switch, a data packet comprising said data packet label, in which said port serial number of said single physical tunnel port is encoded; and

sending said data packet from said first MPLS/LAG switch to a second MSPLS/LAG switch via said single physical tunnel port having the port serial number encoded in the data packet label.

Paper 1 at 12-25
EX1001, Col. 9:22-47

DEMONSTRATIVE EXHIBIT - NOT EVIDENCE

14



Claim 15 of the '602 Patent

15. An apparatus for assigning and utilizing an Ethernet physical data port in an Ethernet Link Aggregation Group (LAG) located downstream from a preceding node in a Multi-Protocol Label Switching (MPLS) network tunnel employing
45 Resource Reservation Protocol Traffic Engineering (RSVP-TE) tunnel provisioning, the LAG having a first MPLS/LAG switch and a second MPLS/LAG switch, the apparatus comprising:
a processor, which is configured to:
50 assign a single physical tunnel port of a LAG to a network tunnel, wherein said single physical tunnel port meets a bandwidth requirement of the network tunnel, and wherein said single physical tunnel port has a port serial number;
55 dedicate a sub-set of bits in a data packet label to encode said port serial number of said single physical tunnel port of the LAG into the data packet label;
publish said packet label upstream; and
a mapper, which is configured to:
60 receive a data packet from an upstream node;
detect said data packet label in said data packet; and
if said data packet contains said data packet label, send said data packet to the second MPLS/LAG switch via said single physical tunnel port.

Paper 1 at 40-48
EX1001, Col. 10:41-64

Claim 26 of the '602 Patent

26. A method for assigning and utilizing an Ethernet physical data port in an Ethernet Link Aggregation Group (LAG) in a Multi-Protocol Label Switching (MPLS) network, the method comprising the steps of:

receiving from a preceding node located upstream from a first MPLS/LAG switch, at said first MPLS/LAG switch, a request message requesting configuration of a network tunnel, the message further requesting a data packet label associated with said network tunnel, said request message specifying a bandwidth requirement;

assigning, at the first MPLS/LAG switch, a single physical tunnel port of a LAG to the network tunnel wherein the single physical tunnel port of the LAG, which meets the specified bandwidth requirement, wherein said single physical tunnel port of said LAG has a port serial number;

dedicating a sub-set of bits in a data packet label prepared by said first MPLS/LAG switch to encode said port serial number of said single physical tunnel port of the LAG into the data packet label;

sending, by said first MPLS/LAG switch, said data packet label, in which said port serial number of said single physical tunnel port is encoded, to the preceding node;

receiving from the preceding node, at said first MPLS/LAG switch, a data packet that includes said data packet label, in which said port serial number of said single physical tunnel port is encoded; and

sending said data packet from said first MPLS/LAG switch to a second MSPLS/LAG switch via said single physical tunnel port having the port serial number encoded in the data packet label.

Paper 1 at 49-52
EX1001, Col. 12:4-36



DEMONSTRATIVE EXHIBIT - NOT EVIDENCE

16

Dr. Rouskas Admits LAG was Known

10 Q. If I wanted to achieve higher than
11 10 gigabits per second in 2002, could I have used
12 a -- a LAG technology to achieve a capacity of
13 higher than 10 gigabits per second at that time?
14 A. It -- at that time it would be possible
15 to, let's say, aggregate 10 -- sorry -- four
16 10Gig links to create a 40Gig link.


EX1022 at 28:10-16

Dr. Rouskas's Deposition Testimony on Implementation of a LAG

1 Q. That's not what I asked.
2 I asked, looking at this link between
3 LSR A and LSR B. I'm not asking whether you
4 would use a LAG. I'm asking, is it technically
5 possible to implement a LAG between LSR A and
6 LSR B in order to transmit data?
7 A. You mean in general, in the abstract.
8 Q. Yeah. In the abstract, is it possible to
9 do so?
10 MR. HOLT: Objection; form.
11 A. (Continuing) Okay. So outside the context
12 of Raz, if there are two LSRs -- let's say A and
13 B. And assuming that each LSR has, you know,
14 multiple physical interfaces. Then it is
15 certainly feasible to implement a LAG between the
16 two LSRs, generally speaking.

EX1022 at 56:1-16

Raz – Figure 1 (Annotated)




Paper 1 at 13
EX1006, Figure 1

DEMONSTRATIVE EXHIBIT - NOT EVIDENCE

19



Raz – Figure 1 (Annotated)



Paper 1 at 14
EX1006, Figure 1

DEMONSTRATIVE EXHIBIT - NOT EVIDENCE

20



Raz's Invention

There is thus provided in accordance with the invention a method of forwarding a high speed packet stream to a plurality of subports of lesser speed in a label switching network, the method comprising the steps of assigning outbound packets to one of a plurality of subports in accordance with the label associated with each outbound packet, forwarding each
45 outbound packet to its assigned subport, aggregating the output of a plurality of subports into a high speed packet stream and forwarding inbound packets from the high speed packet stream to a subport corresponding to the subport the inbound
50 packet was originally sent out on.

Paper 1 at 7, 14
EX1006, Col. 2:41-51

Raz Focuses on Increasing Bandwidth

The invention provides a scheme for allocating LSP bandwidth to subports that helps in solving the problem described above in a way that provides a tradeoff between leaving subports empty for wide LSPs and spreading traffic between subports for best utilizing the bandwidth of the link (e.g., for EIR traffic). The central idea is to leave one or more subports empty (i.e. no bandwidth used) in the event a wide connection is to be allocated. The allocated bandwidth, however, is still distributed among the subports in such a way that EIR traffic is supported.

Paper 1 at 1-20
EX1006, Col. 11:3-12

Dr. Rouskas Agrees There Was A Demand for Increased Bandwidth

3 Q. Okay. And you agree that there was an
4 ever-increasing demand for more and more
5 bandwidth by users as of 2002?
6 A. I believe so.

EX1022 at 22:3-6

The Operation of Raz

In accordance with the invention, the assignment of traffic to 1G subports is based on LSPs. Thus, the MPLS label is used in order to assign an incoming packet to one of the subports. The selected subport is the same for both sides of a 10G link. 20 In the next link on the path, however, the subport may be different. This is because the next link may have other connections in addition to the 10G link and therefore it may be the case that a different subport is the only subport (or the subport best suited) that can accept the LSP with the required band- 25 width. Note that it is the downstream (i.e. receiving) node that determines the subport it wants to receive the LSP on. This decision is communicated to the upstream (i.e. sending) node via the label that is assigned to the LSP.


Paper 1 at 16, 25
EX1006, Col. 7:17-29

Dr. Rouskas Admits Raz Treats Its Subports as a LAG

21 So when he -- and -- and, by the way, he
22 is using that within the switch. So now he's
23 taking a LAG switch and he's repurposing it to
24 use it for his embodiment inside the switch.

EX1022 at 33:21-24

Raz – Figure 3



Paper 1 at 20-22
EX1006, Figure 3

DEMONSTRATIVE EXHIBIT - NOT EVIDENCE

26



Raz's Example in Figure 3

As an example, consider two LSRS, LSR A and LSR B, 30
each with 10 subports wherein each subport is mapped to an
8K range of labels. An LSP is to be provisioned from LSR A
to LSR B. When the LSP is signaled or provisioned through
management, the CAC in LSR B decides to accept the LSP on
subport 3 and thus assigns a label comprising an indication 35
that subport 3 was selected. The label is communicated to
LSR A which now knows that the output subport should be
subport 3 (i.e. the same subport is provisioned).

Paper 1 at 17, 23
EX1006, Col. 7:30-38

Raz's Example in Figure 3

When a packet for the provisioned LSP is received over a link connected to LSR A, it is forwarded to subport 3 and subsequently out over the 10G link to LSR B with the previously assigned label indicating subport 3. The port forwarding module in LSR B examines the label in the received packet and in response thereto forwards the packet to subport 3 for packet processing. Unlabeled packets do not belong to any LSP. Usually these comprise control packets so there are not many of them. These unlabeled packets will always use a specific subport, for example, subport 8.

Paper 1 at 24-25, 34
EX1006, Col. 7:39-48

Raz's Subports *May* or May Not Be Different

10 Outbound packets are assigned a subport based on their
MPLS labels. Once assigned, the packets are forwarded to a
particular subport. The outputs of all the subports are aggregated and combined into a single high-speed packet stream and transmitted over the link. Inbound packets are received by
15 the high speed PHY and forwarded to one of a plurality of
subports. The subport chosen is the same subport as the one
the packet was sent out from at the transmitting node. The
subport is determined in accordance with the label within the
packet itself. The subports used on both sides of the next hop,
20 however, are independent. Thus, the subports may be different
for each hop of the connection.

Paper 1 at 7
EX1006, Col. 2:9-20

Dr. Rouskas Acknowledges the Subports Could Be The Same

4 Q. Okay. And LSR C could have said it wants
5 to receive packets on any of the subports at
6 LSR C.

7 Correct?

8 A. That is correct.


9 Q. It could have received -- chosen to
10 receive packets on subport zero.

11 Is that correct?

12 A. LSR C could have chosen to -- to receive
13 packets from subport zero.

EX1022 at 68:4-13

Ferguson – Figure 1




Paper 1 at 12

DEMONSTRATIVE EXHIBIT - NOT EVIDENCE

31



Ferguson – Figures 2A-2C (Annotated)



Paper 1 at 54

DEMONSTRATIVE EXHIBIT - NOT EVIDENCE

32



Ferguson's Invention

FIG. 1 is a block diagram illustrating an example system 2 in which customer networks 4A, 4B ("customer networks 4") are communicatively coupled via an intermediate provider network 6. Specifically, customer network 4A includes a customer edge ("CE") router 8A that is connected to a provider edge ("PE") router 10A via an aggregated link 12A. In other words, CE router 8A and PE router 10A communicate via two or more physical interconnects that are logically associated as an aggregated link. For example, CE router 8A and PE router 10A may communicate via two or more Ethernet links, Sonet links, or other physical interconnects. Similarly, customer network 4B includes a CE router 8B that is connected to a PE router 10B via an aggregated link 12B. PE routers 10 communicate network traffic from and to customer networks 4 in accordance with a label switching protocol, such as the Multiprotocol Label Switching (MPLS) protocol, via one or more label switched paths (LSPs) 14. For example, CE router 8A receives network traffic originating from devices within customer network 4A, and pre-pends MPLS labels to the network traffic to form outbound MPLS packets. CE router 8A communicates the outbound MPLS packets to PE router 10A via aggregated link 12A.

DEMONSTRATIVE EXHIBIT - NOT EVIDENCE

Paper 1 at 11
EX1007, Col. 3:35-58

33



Ferguson's Use of MPLS

PE router **10A** receives the MPLS packets, and forwards the MPLS packets to one or more intermediate routers (not shown) within provider network **6** for communication to PE router **10A** via one of LSPs **14**. PE router **10A** may, for example, pre-pend additional MPLS labels to the traffic to form MPLS packets having multiple MPLS labels. In this manner, each MPLS packet traversing provider network **6** via LSPs **14** may be guaranteed to carry two or more MPLS labels.

Paper 1 at 53-54
EX1007, Col. 3:59-67

Ferguson Addresses Link Aggregation or LAG

The term “link aggregation” refers to a technique by which the multiple physical interconnects are logically associated and treated as a single, aggregated link. In particular, the IEEE 802.3ad Working Group recently developed a Link Aggregation specification. This approach may provide for increased link availability and bandwidth between the two devices. For example, bandwidth can be increased in linear increments, e.g., by adding individual interconnects to increase the bandwidth between the two devices. In addition, the multiple physical interconnects provide a degree of redundancy in that a failure in a single physical interconnect no longer inhibits direct traffic flow between the two devices. With link aggregation, traffic on a failed physical interconnect can be redirected to one of the other physical interconnects in the aggregated link.

Paper 1 at 55, 57
EX1007, Col. 1:28-42

DEMONSTRATIVE EXHIBIT - NOT EVIDENCE

35



Ferguson's Discussion of "Physical Interconnects"

In some environments, two network devices may be connected by multiple physical interconnects. For example, two network routers may be connected by multiple Ethernet links. Similarly, multiple physical interconnects may couple other devices, e.g., two switches, a switch and a server, two routers, a switch and an end user station, and the like. ²⁵

Paper 1 at 53
EX1007, Col. 1:21-27

Ferguson Uses “Logically Associated Physical Interconnects” to Load Balance

In general, the invention is directed to techniques for distributing label switched packets, such as multiprotocol label switched (MPLS) packets, associated with a common label switched path (LSP). The techniques may, for example, be used to load balance label switched packets of an LSP across an aggregated link having two or more logically associated physical interconnects. In addition, the techniques may be used to distribute label switched packets of a common LSP across multiple data paths within a network device, such as a router. As used herein, the term path refers to any physical path for communicating data. In this manner,

2

the term encompasses both external physical interconnects between network devices, and internal data paths within an individual network device.

Paper 1 at 10, 53
EX1007, Cols. 1:56-2:3