




(86) Date de dépôt PCT/PCT Filing Date: 2005/05/25
 (87) Date publication PCT/PCT Publication Date: 2005/12/15
 (45) Date de délivrance/Issue Date: 2014/11/18
 (85) Entrée phase nationale/National Entry: 2006/09/22
 (86) N° demande PCT/PCT Application No.: IB 2005/002421
 (87) N° publication PCT/PCT Publication No.: 2005/117519
 (30) Priorités/Priorities: 2004/05/25 (US60/574,253);
 2004/05/28 (US60/575,772)

(51) Cl.Int./Int.Cl. *H04L 12/24* (2006.01),
G08B 23/00 (2006.01), *H04L 12/26* (2006.01),
H04L 12/28 (2006.01)
 (72) Inventeurs/Inventors:
 HOLNESS, MARC, CA;
 MOHAN, DINESH, CA
 (73) Propriétaire/Owner:
 ROCKSTAR CONSORTIUM US LP, US
 (74) Agent: BORDEN LADNER GERVAIS LLP

(54) Titre : NOTIFICATION D'ERREUR DE CONNECTIVITE
 (54) Title: CONNECTIVITY FAULT NOTIFICATION



(57) **Abrégé/Abstract:**

Connectivity fault notification is provided by generating an alarm indication signal at a device that is logically adjacent to the fault, and forwarding the alarm indication signal upward through various levels to at least one client level entity. The alarm indication signal may be suppressed at any level for a service instance if service is restored at that level, or if a protection path prevents disruption of the service instance at that level, or auto-suppressed at an originating node based on number of times transmitted or elapsed time. The alarm indication signal may include a point of failure indicator such as the MAC address of the device that generates the alarm indication signal, or a failed resource identity such as an IEEE 802.1 AB LLDP MAC Service Access Point ("MSAP"). Further, the alarm indication signal may be employed to trigger use of the protection path.

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
15 December 2005 (15.12.2005)

PCT

(10) International Publication Number
WO 2005/117519 A3

(51) International Patent Classification:

H04L 12/24 (2006.01) *G08B 23/00* (2006.01)
H04L 12/26 (2006.01) *H04L 12/28* (2006.01)

(21) International Application Number:

PCT/IB2005/002421

(22) International Filing Date: 25 May 2005 (25.05.2005)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:

60/574,253 25 May 2004 (25.05.2004) US
60/575,772 28 May 2004 (28.05.2004) US

(71) Applicant (for all designated States except US): **NORTEL NETWORKS LIMITED** [CA/CA]; 2351 Boulevard Alfred-Nobel, St. Laurent, Quebec H4S 2A9 (CA).

(72) Inventors: **MOHAN, Dinesh**; 89 Kenins Cres, Kanata, Ontario K2K 3E5 (CA). **HOLNESS, Marc**; 928 Fisher Avenue, Ottawa, Ontario K1Z 6P4 (CA).

(74) Agent: **BORDEN LADNER GERVAIS LLP**; World Exchange Plaza, 100 Queen Street, Suite 1100, Ottawa, Ontario K1P 1J9 (CA).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM,

AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NA, NG, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, SM, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, YU, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IS, IT, LT, LU, MC, NL, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:


- with international search report
- before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments

(88) Date of publication of the international search report:

29 June 2006

[Continued on next page]

(54) Title: CONNECTIVITY FAULT NOTIFICATION



(57) Abstract: Connectivity fault notification is provided by generating an alarm indication signal at a device that is logically adjacent to the fault, and forwarding the alarm indication signal upward through various levels to at least one client level entity. The alarm indication signal may be suppressed at any level for a service instance if service is restored at that level, or if a protection path prevents disruption of the service instance at that level, or auto-suppressed at an originating node based on number of times transmitted or elapsed time. The alarm indication signal may include a point of failure indicator such as the MAC address of the device that generates the alarm indication signal, or a failed resource identity such as an IEEE 802.1 AB LLDP MAC Service Access Point ("MSAP"). Further, the alarm indication signal may be employed to trigger use of the protection path.

WO 2005/117519 A3

WO 2005/117519 A3



For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

CONNECTIVITY FAULT NOTIFICATION

FIELD OF THE INVENTION

This invention is generally related to the field of network communications, and
5 more particularly to fault notifications associated with connectivity failures.

BACKGROUND OF THE INVENTION

Ethernet and other communications protocols that were once used only in Local
Area Network ("LAN") environments are now being used in Metropolitan Area Networks
10 ("MANs") and Wide Area Networks ("WANs"). One reason for this development is that
Enterprises and other customers of communications Service Providers find it convenient to
operate, end-to-end, in an environment which is native to their LANs and understood by
their IT professionals. However, the extension of such protocols to environments in which
they were not originally intended presents some problems.

15 One problem is that a single failure can trigger a cascade of alarms that overwhelm
an NMS/OSS. A single node or link in a WAN, for example, is likely to support many
more different services than a LAN device. Each supported service may also rely on a
greater number of network nodes between endpoints. If an intermediate node or link fails,
an alarm is generated for each failed service. However, the other network nodes that
20 support those services also generate alarms. Hence, in comparison to operation on a LAN,
an overwhelming number of alarms can be generated as a result of a single failure.

Another problem is fault localization. In the case where communications services
require use of the networks of multiple different service providers, for example, service
faults may trigger provisions in Service Level Agreements ("SLAs"), which are contracts
25 between different service providers, and also between service providers and their
customers. An enterprise customer or service provider may be entitled to a credit, or
invocation of a service cancellation clause, depending on the cause, frequency and
duration of faults. It is therefore desirable to have technology to produce accurate, perhaps
auditable, notification regarding a service fault. The notification may also include an
30 indication of the cause of the fault and fault localization information.

SUMMARY OF THE INVENTION

In accordance with the invention, a method for providing connectivity fault notification in a network includes the steps of: detecting a connectivity fault or failure, which could include loss of continuity and misconnections, at a node logically adjacent to the failure; generating an alarm indication signal in response to detecting the connectivity fault or failure; and forwarding the alarm indication signal to at least one client level entity. In particular, the alarm indication signal is forwarded from a server layer (which has detected the fault or failure) toward the client layer, possibly through multiple intermediate levels. The alarm indication signal is periodically retransmitted until it is suppressed. For example, the alarm indication signal may be suppressed following expiration of a timer. Alternatively, the alarm indication signal may be suppressed for a client connection if a protection path prevents disruption of the client connection. It should be noted that the alarm indication signal may be suppressed at any level.

The alarm indication signal may include a point of failure indicator. For example, the alarm indication signal may include the MAC address of the device that generates the alarm indication signal, or a failed resource identity such as an IEEE 802.1AB LLDP MAC Service Access Point ("MSAP"). Hence, the client can advantageously determine the fault origin from the alarm indication signal. In the case of a service provided by multiple Service Provider networks this advantageously provides evidence of fault origin. The alarm indication signal may also be employed to trigger use of a protection path.

A network device in accordance with the invention may include circuitry operable to detect a connectivity fault or failure, which could include loss of continuity and misconnections, logically adjacent to the device; logic operable to generate an alarm indication signal in response to detection of the connectivity failure; and at least one port operable to forward the alarm indication signal to at least one client level entity. The device is operable to periodically retransmit the alarm indication signal, e.g., once per second. The device may also include a memory operable to maintain a record including bindings of Destination Addresses ("DAs"), associated ports, and an indication of client connections, e.g., Virtual LAN ("VLAN") identification associated with each DA, and Maintenance End Point ("MEP") of the client level entity. When a fault is detected the memory can be employed to determine which client connections are affected, and which ports are associated with those services. The alarm indication signal is then forwarded

only via those ports associated with service instances affected by the connectivity failure. As will be discussed in greater detail below, the DA is most useful when the alarm indication signal is unicast, rather than multicast.

5 BRIEF DESCRIPTION OF THE FIGURES

Figure 1 illustrates a technique for delivering an indication of a network fault from a lower level to an higher level.

Figure 2 illustrates a format of the AIS of Figure 1.

Figure 3 illustrates the technique of Figure 1 in the context of a network diagram.

10 Figure 4 illustrates alarm suppression and fault recovery.

Figure 5 illustrates selective and non-selective AIS suppression.

DETAILED DESCRIPTION

Figures 1 and 2 illustrate a technique for delivering an indication of a network fault
 15 that affects an Ethernet-based service such as VLANs. The end-to-end network includes a first service provider sub-network (100), a second service provider sub-network (102), and client level entities (104, 106) associated with customer networks. Initially, a fault (108) is detected at the Ethernet physical layer ("ETY") by nodes such as bridges (110, 112) that are logically adjacent to the fault. The fault may be a link failure, node failure or
 20 misconnection. Further, the fault could be detected at any layer of the OSI model. In response to detection of the fault, the nodes that are logically adjacent to the failed link (108) prompt generation of an Alarm Indication Signal ("AIS") (200). In particular, Maintenance End Points ("MEPs") (116) on either side of the fault each generate an AIS (114) which is forwarded by Maintenance Intermediate Points ("MIPs") (118) at the next
 25 higher ME level towards the MEPs at the same ME level as the MIP (120) using either unicast or multicast. Upon receiving the respective AISs, each higher level MEP (120) may generate a corresponding higher level AIS which is forwarded by MIPs towards MEPs (122) at the next higher level. Similarly, upon receiving the respective AISs, those MEPs (122) may generate higher level AISs which are forwarded by corresponding MIPs
 30 towards MEPs (124) associated with the client level entities (104, 106).

The alarm indication signal is periodically retransmitted until it is suppressed via one of various techniques. For example, the alarm indication signal is suppressed if the

fault is repaired. A client level device assumes a fault condition has ended when X
 consecutive AISs are not received. Alternatively, the alarm indication signal may be auto-
 suppressed by the originating MEP. In particular, the originating MEP may be
 programmed to transmit the AIS a predetermined number of times, or for only a
 5 predetermined period of time in accordance with a countdown timer, following which the
 AIS is suppressed regardless of whether the fault condition remains. Alternatively, the
 alarm indication signal may be suppressed for a client connection if a protection path
 prevents disruption of the client connection. It should be noted that the alarm indication
 signal may be suppressed at any level. In particular, at any given level the MEP may be
 10 operable to suppress the AIS by not generating a corresponding AIS to be forwarded to the
 next higher level. Such an action could be taken when, for example, the connection has
 been restored, or does not appear to have failed, at that level. However, that receiving
 MEP may forward the AIS to its higher ME level upon confirming the failure on its own
 basis, e.g. determining failure by loss of continuity. Alarm suppression may also be
 15 associated with a particular client level entity. Further, an alarm may be suppressed by
 manual, administrative action.

Referring now to Figures 1 and 2, the AIS includes a Destination MAC address
 field (200), a source MAC address field (202), an EtherType (VLAN) field (204), a VLAN
 tag field (206), an EtherType (OAM) field (208), Version field (210), ME Level field
 20 (212), OpCode field (214), Header Length field (216), transaction/sequence identifier field
 (218), transmission timestamp field (220), Fixed header future extension field (222),
 Service ID TLV (Type Length Value) field (224), and other TLVs field (226). The
 transaction identifier (218) and transmission timestamp field (220) are not particularly
 useful for the AIS, but may still be present in the generic frame format. Further, field (222)
 25 may be used to extend the fixed header in future versions of the OAM, where the version
 is indicated by a different version field value. The other TLVs field (226) may include a
 point of failure indicator TLV (228), e.g., the link ID (232) of the failed link, or an
 identification of which device generated the AIS, in order to provide fault isolation
 capabilities. In particular, the AIS initiating node can insert location information such as
 30 the node's own unique MAC address TLV (234), or failed resource identity TLV (236),
 e.g., IEEE 802.1ab LLDP MAC Service Access Point (MSAP) (238). AIS initiating node

may also insert a cause indicator TLV (230) into the AIS if the cause of the fault can be determined.

Referring to Figure 3, each node may maintain a record (300) in the form of a table or database to facilitate generation and selective transmission of the AIS. Each record in the table may include a Destination Address ("DA") field (302) in the form of a MAC address, a field (304) indicating of the port associated with that DA, a field (306) indicating the client entities identified by VLANs associated with that port, and a field (308) indicating associated client level MEPs. When a fault is detected it is associated with a particular port of the node. The node can then index into the table based on port number to identify the affected VLANs. Hence, the node may employ the table to identify all ports and DAs associated with the affected VLANs. The node will then transmit the MS via each port associated with each VLAN that is associated with the port via which the fault is detected. Furthermore, the node may indicate via the AIS which MEPs have been isolated at the client ME level while forwarding the AIS towards the client ME level MEPs on the other side of the failure. Such transmission of the MS in response to a fault mitigates superfluous alarm suppression at the client ME level MEPs since the MEPs are able to determine which peer MEPs have been isolated as a result of fault for which the AIS is received. Such an AIS is called selective AIS. Information about the client level MEPs that are associated with a MEP at the server level may be communicated via a specific OAM message that is used when AIS functionality with selective alarm suppression capability is desirable. Further, the node may be configured not to maintain the records and hence may not have enough information to allow selective alarm suppression information in the transmitted AIS. Such AIS is called non-selective AIS. In the cases where network topology associated with the affected VLAN is known it may be possible to utilize a unicast AIS, and hence to utilize the DA. Such AIS is called as unary AIS and is directed for a specific MEP at the client ME level. Otherwise, a multicast AIS is employed for selective or non-selective purposes. In particular, a multicast AIS is sent for each affected VLAN because the AIS includes VLAN-specific fields.

In the illustrated example the network supports three different VLANs: S12, S13 and S14. VLAN S12 provides connectivity between client level entity CE1 and client level entity CE2. VLAN S13 provides connectivity between client level entity CE1 and client level entity CE3. VLAN S14 provides connectivity between client level entity CE1 and

client level entity CE4. When a fault occurs in link P32-PE21, only VLAN S12 is affected because functional paths remain between client level entity CE1 and both client level entities CE4 and CE3. The fault in link P32-PE21 is initially detected on nodes PE2 and P3. Node PE2 associates the fault with its port 1, and node P3 associates the fault with its port 2. Rather than transmit a corresponding AIS via all ports in response to detection of the link failure, nodes PE2 and P3 transmit the AIS only via those ports associated with the affected VLAN. For example, node P3 determines from its table that a fault associated with its port 2 affects VLAN S12, and that VLAN S12 is also associated with its port 1, and hence transmits an AIS via its port 1. Similarly, node PE2 determines that a fault associated with its port 1 affects VLAN S12, and that VLAN S12 is also associated with its port 2, and hence transmits an AIS via its port 2. Client level entity CE2 receives the AIS from node PE2, thereby providing timely fault notification. Node P1 is operative in response to the AIS received via its port 2 from node P3 to determine that VLAN S12 is affected by a fault associated with its port 2, and that VLAN S12 is also associated with its port 1. Hence, node P1 forwards the corresponding AIS via its port 1. In response to the AIS from node P1, node PE1 determines that VLAN S12 is affected by the fault associated with its port 2, and that its port 1 is also associated with VLAN S12. Hence, node PE1 forwards the corresponding AIS via its port 1 to client level entity CE1, thereby providing timely fault notification. Consequently, the AIS is propagated only to the client level entities CE1 and CE2 associated with VLANs affected by the fault. It will be appreciated by those skilled in the art that a single fault could affect more than one VLAN, and hence execution of the technique described above at each intermediate node may be desirable.

Figure 4 illustrates a fault scenario in which an alternate, protection path (400) is available to restore the affected VLANs in a single Service Provider network. Initially, link P12-P31 is employed to support both VLAN S12 and VLAN S14. When link P12-P31 fails both VLAN S12 and VLAN S14 are affected. However, traffic associated with VLANs S12 and S14 may be redirected along a protection path (400) from node PE1 to node P2 to node P3 in order to avoid the failed link. In the case where the protection path is implemented in a manner which avoids disruption of the affected VLANs it may be desirable to suppress the AIS messages. However, if the protection path is implemented in a manner which does disrupt the VLANs then the AIS messages may be forwarded to the

client level entities even though the protection path eventually restores service. The network may also be implemented such that MS generation is employed to trigger implementation of a protection path. In such a case it might be worthwhile to transmit the AIS for a short, predetermined period of time.

5 Referring to Figure 5, a VLAN may be associated with more than two client level entities. For example, VLAN S124 is associated with client level entities CE1, CE2 and CE4. When link P32-PE21 fails, client level entity CE2 is isolated from client level entities CE1 and CE4. AIS messages are generated in response to detection of the fault as described above. However, the alarms may be partially or entirely suppressed because the
10 VLAN has not failed entirely. In particular, if the client level entities have insufficient information in the received AIS upon which to selectively suppress the alarms, then all of the alarms are suppressed. If, however, the identification of the client level entity or entities isolated on the other side of the fault is available in the received AIS then the alarms are selectively suppressed. For example, when AIS related to link P32-PE21 failure
15 is received at client level entity CE1 and provides information that client level entity CE2 is isolated, client level entity CE1 can suppress alarms related to client level entity CE2 however can still report alarms related to client level entity CE4.

While the invention is described through the above exemplary embodiments, it will be understood by those of ordinary skill in the art that modification to and variation of the
20 illustrated embodiments may be made without departing from the inventive concepts herein disclosed. Moreover, while the preferred embodiments are described in connection with various illustrative structures, one skilled in the art will recognize that the system may be embodied using a variety of specific structures. Accordingly, the invention should not be viewed as limited except by the scope and spirit of the appended claims.

CLAIMS:

1. A method for providing connectivity fault notification in a network, the method comprising:
 - 5 detecting a fault at a node logically adjacent to the fault;
 - generating an alarm indication signal in response to detecting the fault; and
 - forwarding the alarm indication signal to only client level entities associated with a particular virtual local area network (VLAN) affected by the fault.
- 10 2. The method of claim 1, further comprising:
 - forwarding the alarm indication signal to the client level entities via a higher maintenance entity (ME) level.
3. The method of claim 1, further comprising:
 - 15 suppressing a transmission of the alarm indication signal for a service instance at a level if the service instance is not disrupted at that level.
4. The method of claim 1, further comprising:
 - 20 suppressing a transmission of the alarm indication signal for a service instance if the fault is repaired.
5. The method of claim 1, further comprising:
 - 25 suppressing a transmission of the alarm indication signal for a service instance in response to administrative input.
6. The method of claim 1, further comprising:
 - suppressing a transmission of the alarm indication signal for a service instance after a predetermined period of time.

7. The method of claim 1, further comprising:
suppressing a transmission of the alarm indication signal for a service instance after periodically transmitting a predetermined number of the alarm indication signal.
- 5 8. The method of claim 1, further comprising:
forwarding the alarm indication signal via each port of the node associated with a service instance affected by the fault.
9. The method of claim 1, further comprising:
10 forwarding the alarm indication signal via each port of the node.
10. The method of claim 1, further comprising:
inserting a point of failure indicator in the alarm indication signal.
- 15 11. The method of claim 10 wherein the point of failure indicator includes a media access control (MAC) address.
12. The method of claim 10 wherein the point of failure indicator includes a failed resource identity.
- 20 13. The method of claim 1 wherein the fault relates to a service instance, the service instance being the Virtual Local Area Network ("VLAN").
14. The method of claim 13 wherein a different multicast alarm indication signal is
25 generated for each VLAN.
15. The method of claim 13 wherein a different unicast alarm indication signal is generated for each VLAN.
- 30 16. The method of claim 1, further comprising:
inserting a cause of failure indicator in the alarm indication signal.

17. The method of claim 1, further comprising:
triggering use of a protection path in response to receipt of the alarm indication signal.

5 18. The method of claim 1 wherein, if the client level entities have insufficient information in the alarm indication signal upon which to selectively suppress the alarms, then all of the alarms are suppressed.

10 19. The method of claim 1 wherein if an identification of a client level entity isolated on an other side of the fault is available in the alarm indication signal then the alarms for which identification is known are selectively suppressed.

20. A network device operable to provide connectivity fault notification in a network, comprising:

15 circuitry configured to detect a fault logically adjacent to the device;
logic configured to generate an alarm indication signal in response to detection of the fault; and
at least one port configured to forward the alarm indication signal to only client level entities associated with a particular virtual local area network (VLAN) affected by
20 the fault.

21. The device of claim 20 further comprising:
logic configured to suppress the alarm indication signal for a service instance at a level if a the service instance is not disrupted at that level.

25

22. The device of claim 20 further comprising:
logic configured to suppress the alarm indication signal for a service instance if the fault is repaired.

30 23. The device of claim 20 further comprising:
logic configured to suppress the alarm indication signal for a service instance in response to administrative input.

24. The device of claim 20 further comprising:
logic configured to suppress the alarm indication signal for a service instance after a predetermined period of time.

5

25. The device of claim 20 further comprising:
logic configured to suppress the alarm indication signal for a service instance after transmitting a predetermined number of copies of the alarm indication signal.

10

26. The device of claim 20 further comprising:
logic configured to forward the alarm indication signal via each port of the node associated with a service instance affected by the fault.

27. The device of claim 20 further comprising:

15

logic configured to forward the alarm indication signal via each port of the device.

28. The device of claim 20 further comprising:

logic configured to insert a point of failure indicator in the alarm indication signal.

20

29. The device of claim 28 wherein the point of failure indicator includes a MAC address.

30. The device of claim 28 wherein the point of failure indicator includes a failed resource identity.

25

31. The device of claim 20 further comprising:

logic configured to insert a cause of failure indicator in the alarm indication signal.

32. The device of claim 20 further comprising:

30

logic configured to trigger use of a protection path in response to receipt of the alarm indication signal.

33. The device of claim 20 wherein a different multicast alarm indication signal is generated for each affected VLAN.

34. The device of claim 20 wherein a different unicast alarm indication signal is
5 generated for each affected VLAN.

35. The device of claim 20 wherein, if the client level entities have insufficient information in the alarm indication signal upon which to selectively suppress the alarms, then all of the alarms are suppressed.

10

36. The device of claim 20 wherein if the identification of the client level entity isolated on the other side of the fault is available in the alarm indication signal then the alarms for which identification is known are selectively suppressed.

123-003
16904RO




Figure 1

123-003
16904RO




Figure 2

123-003
16904RO

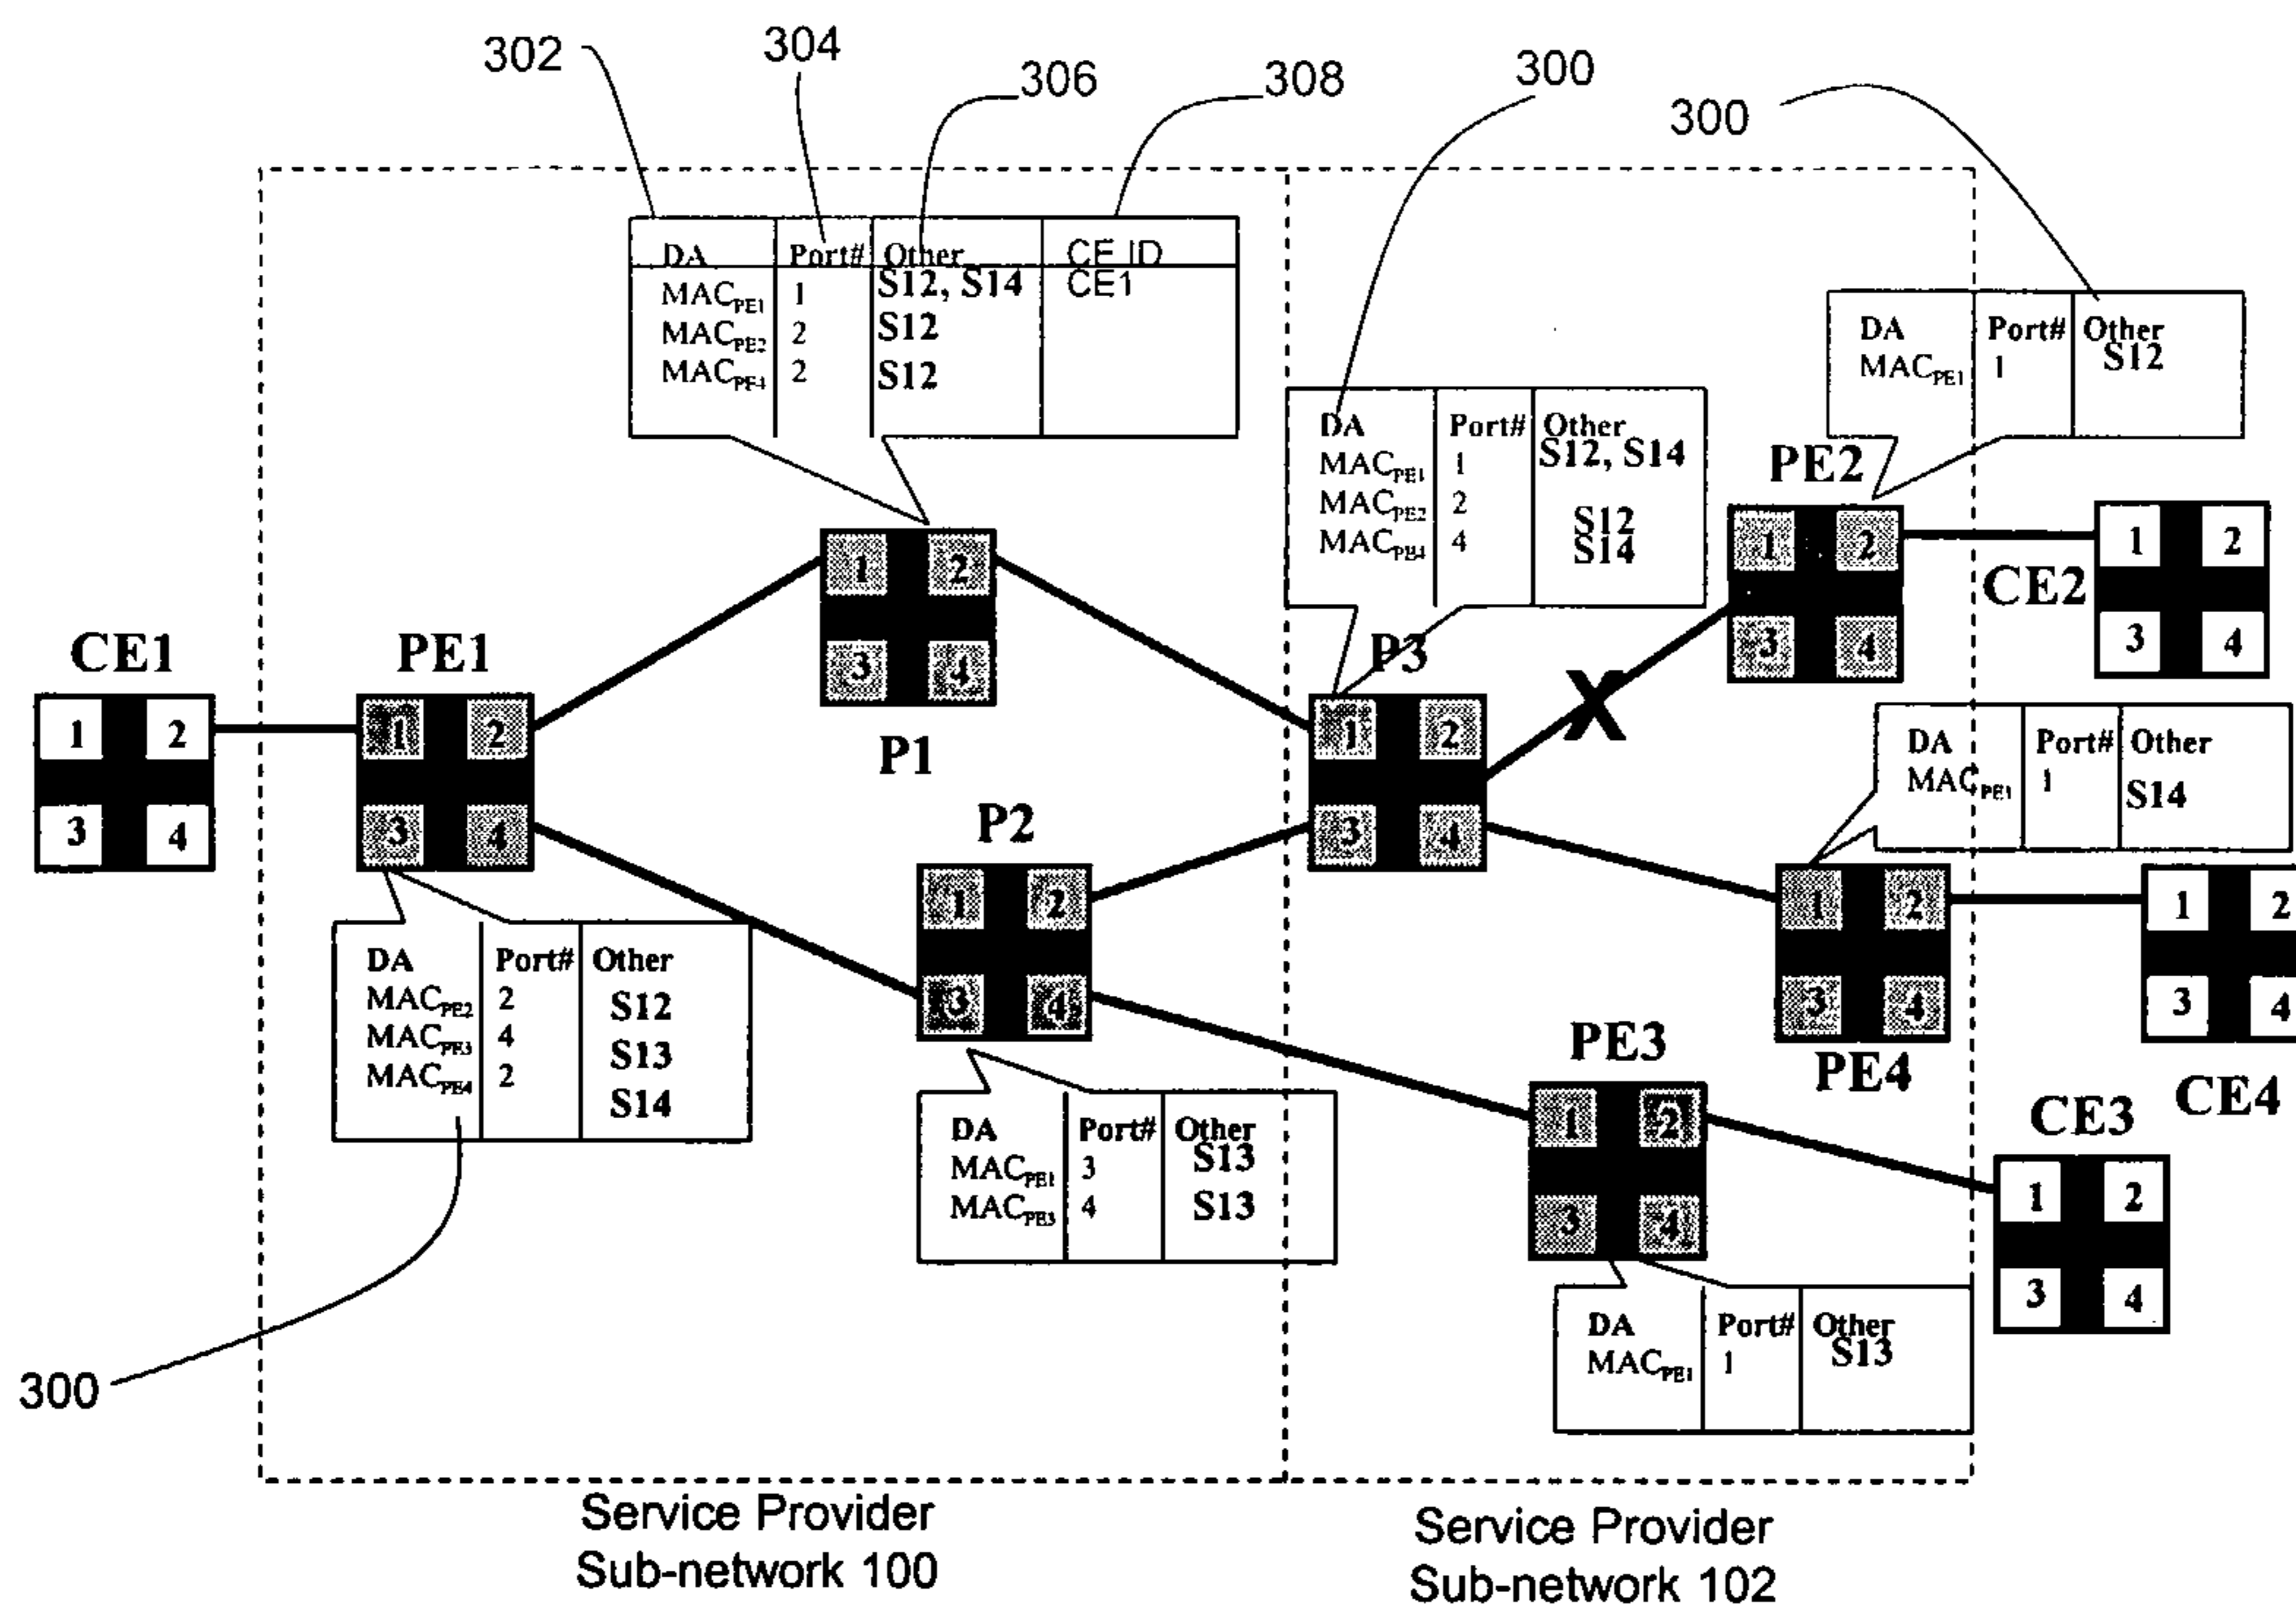


Figure 3

123-003
16904RO

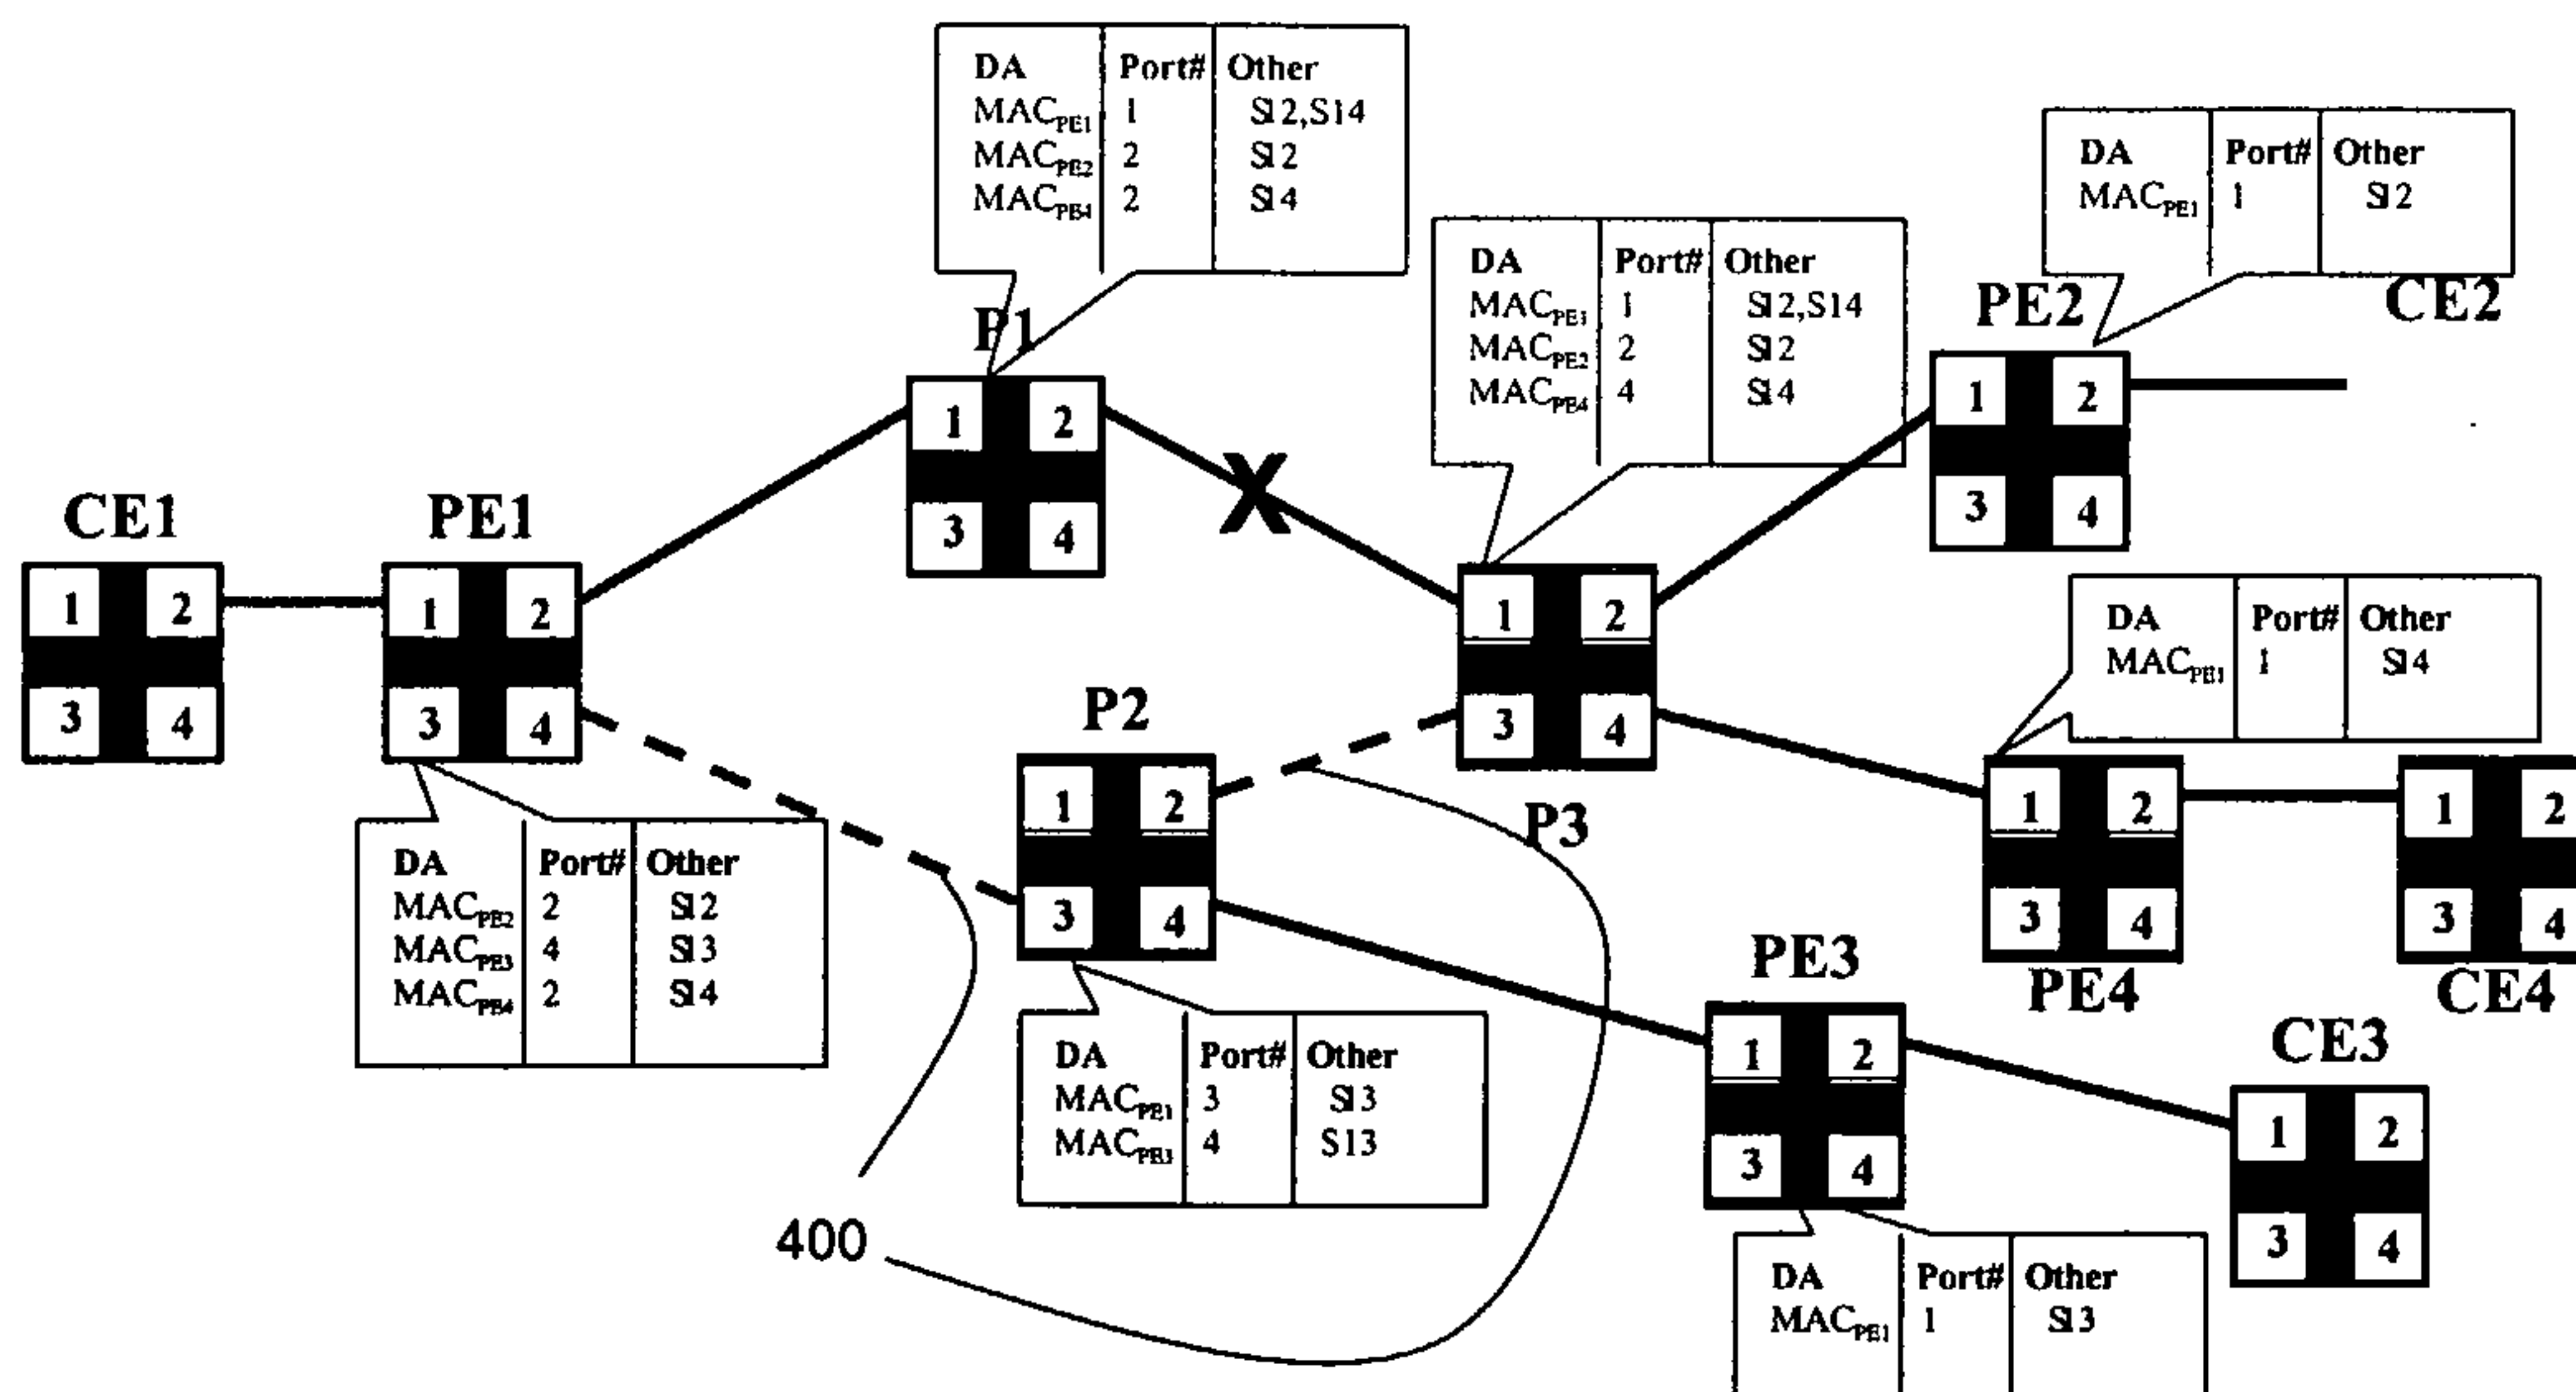


Figure 4

123-003
16904RO

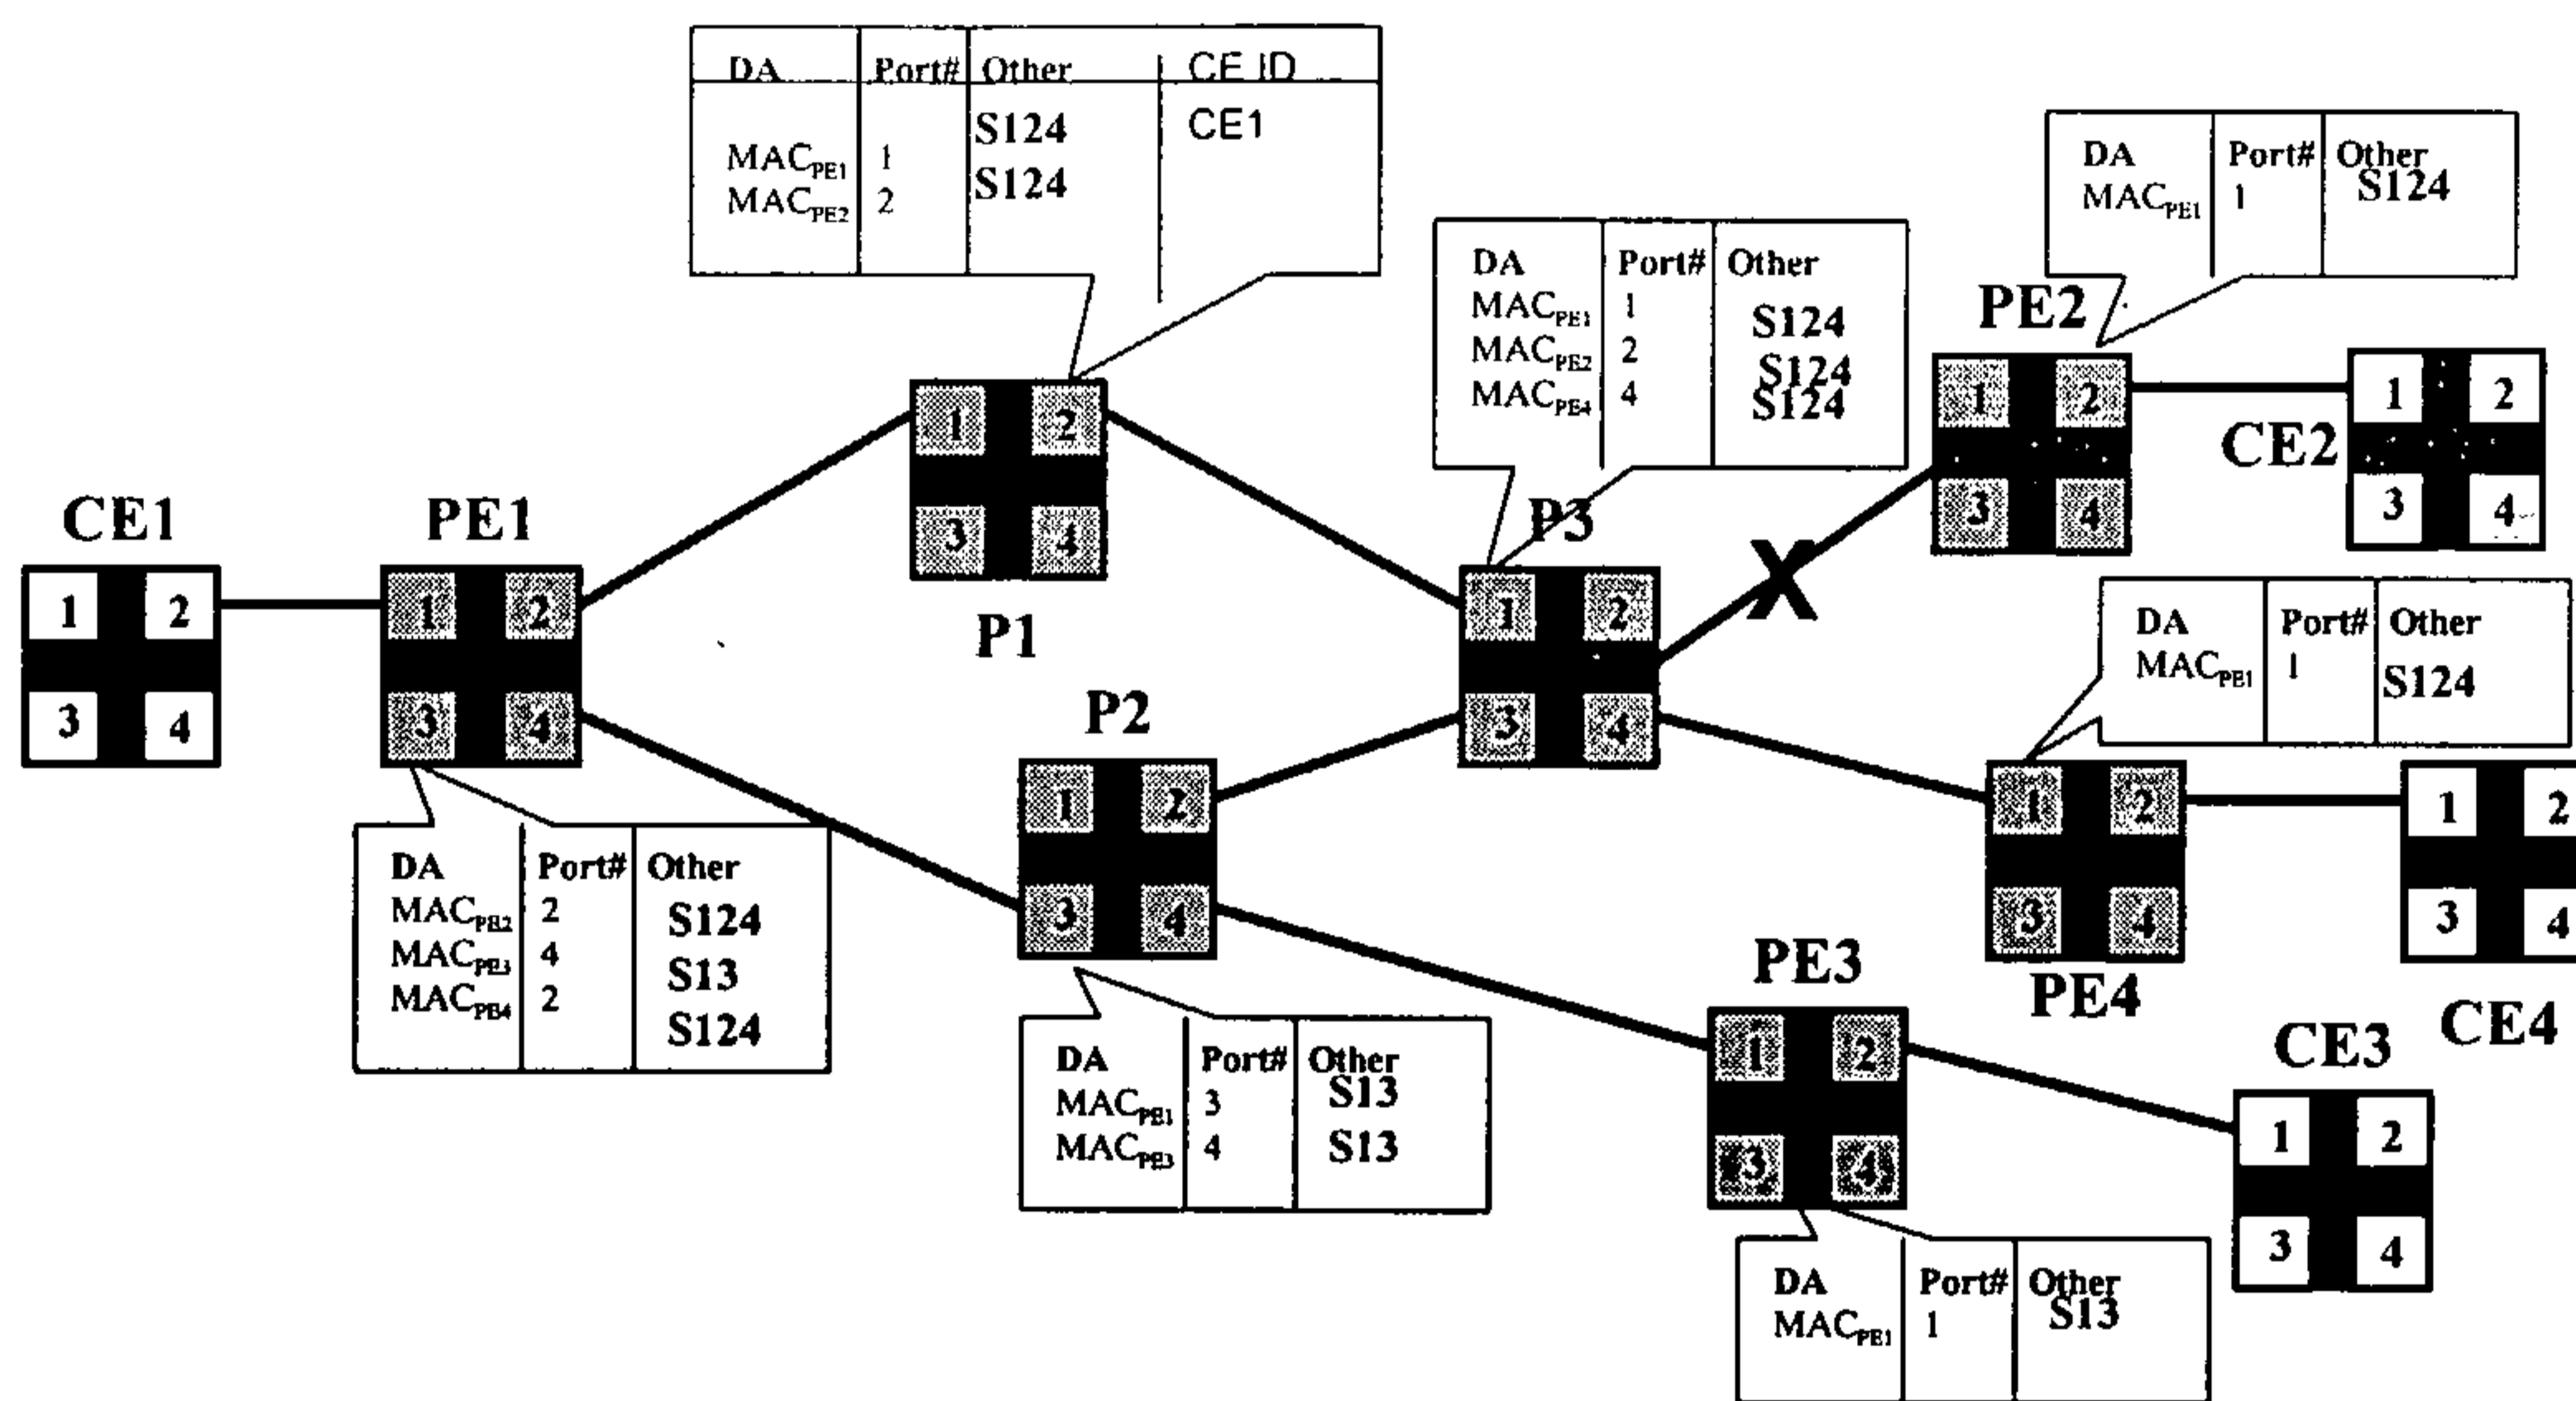


Figure 5

