



Office de la Propriété

Intellectuelle
du Canada

Un organisme
d'Industrie Canada

Canadian
Intellectual Property
Office

An agency of
Industry Canada

CA 2481340 C 2012/12/18

(11)(21) **2 481 340**

(12) **BREVET CANADIEN
CANADIAN PATENT**

(13) **C**

(86) Date de dépôt PCT/PCT Filing Date: 2003/03/25
(87) Date publication PCT/PCT Publication Date: 2003/10/23
(45) Date de délivrance/Issue Date: 2012/12/18
(85) Entrée phase nationale/National Entry: 2004/10/04
(86) N° demande PCT/PCT Application No.: US 2003/009332
(87) N° publication PCT/PCT Publication No.: 2003/088050
(30) Priorité/Priority: 2002/04/05 (US10/118,386)

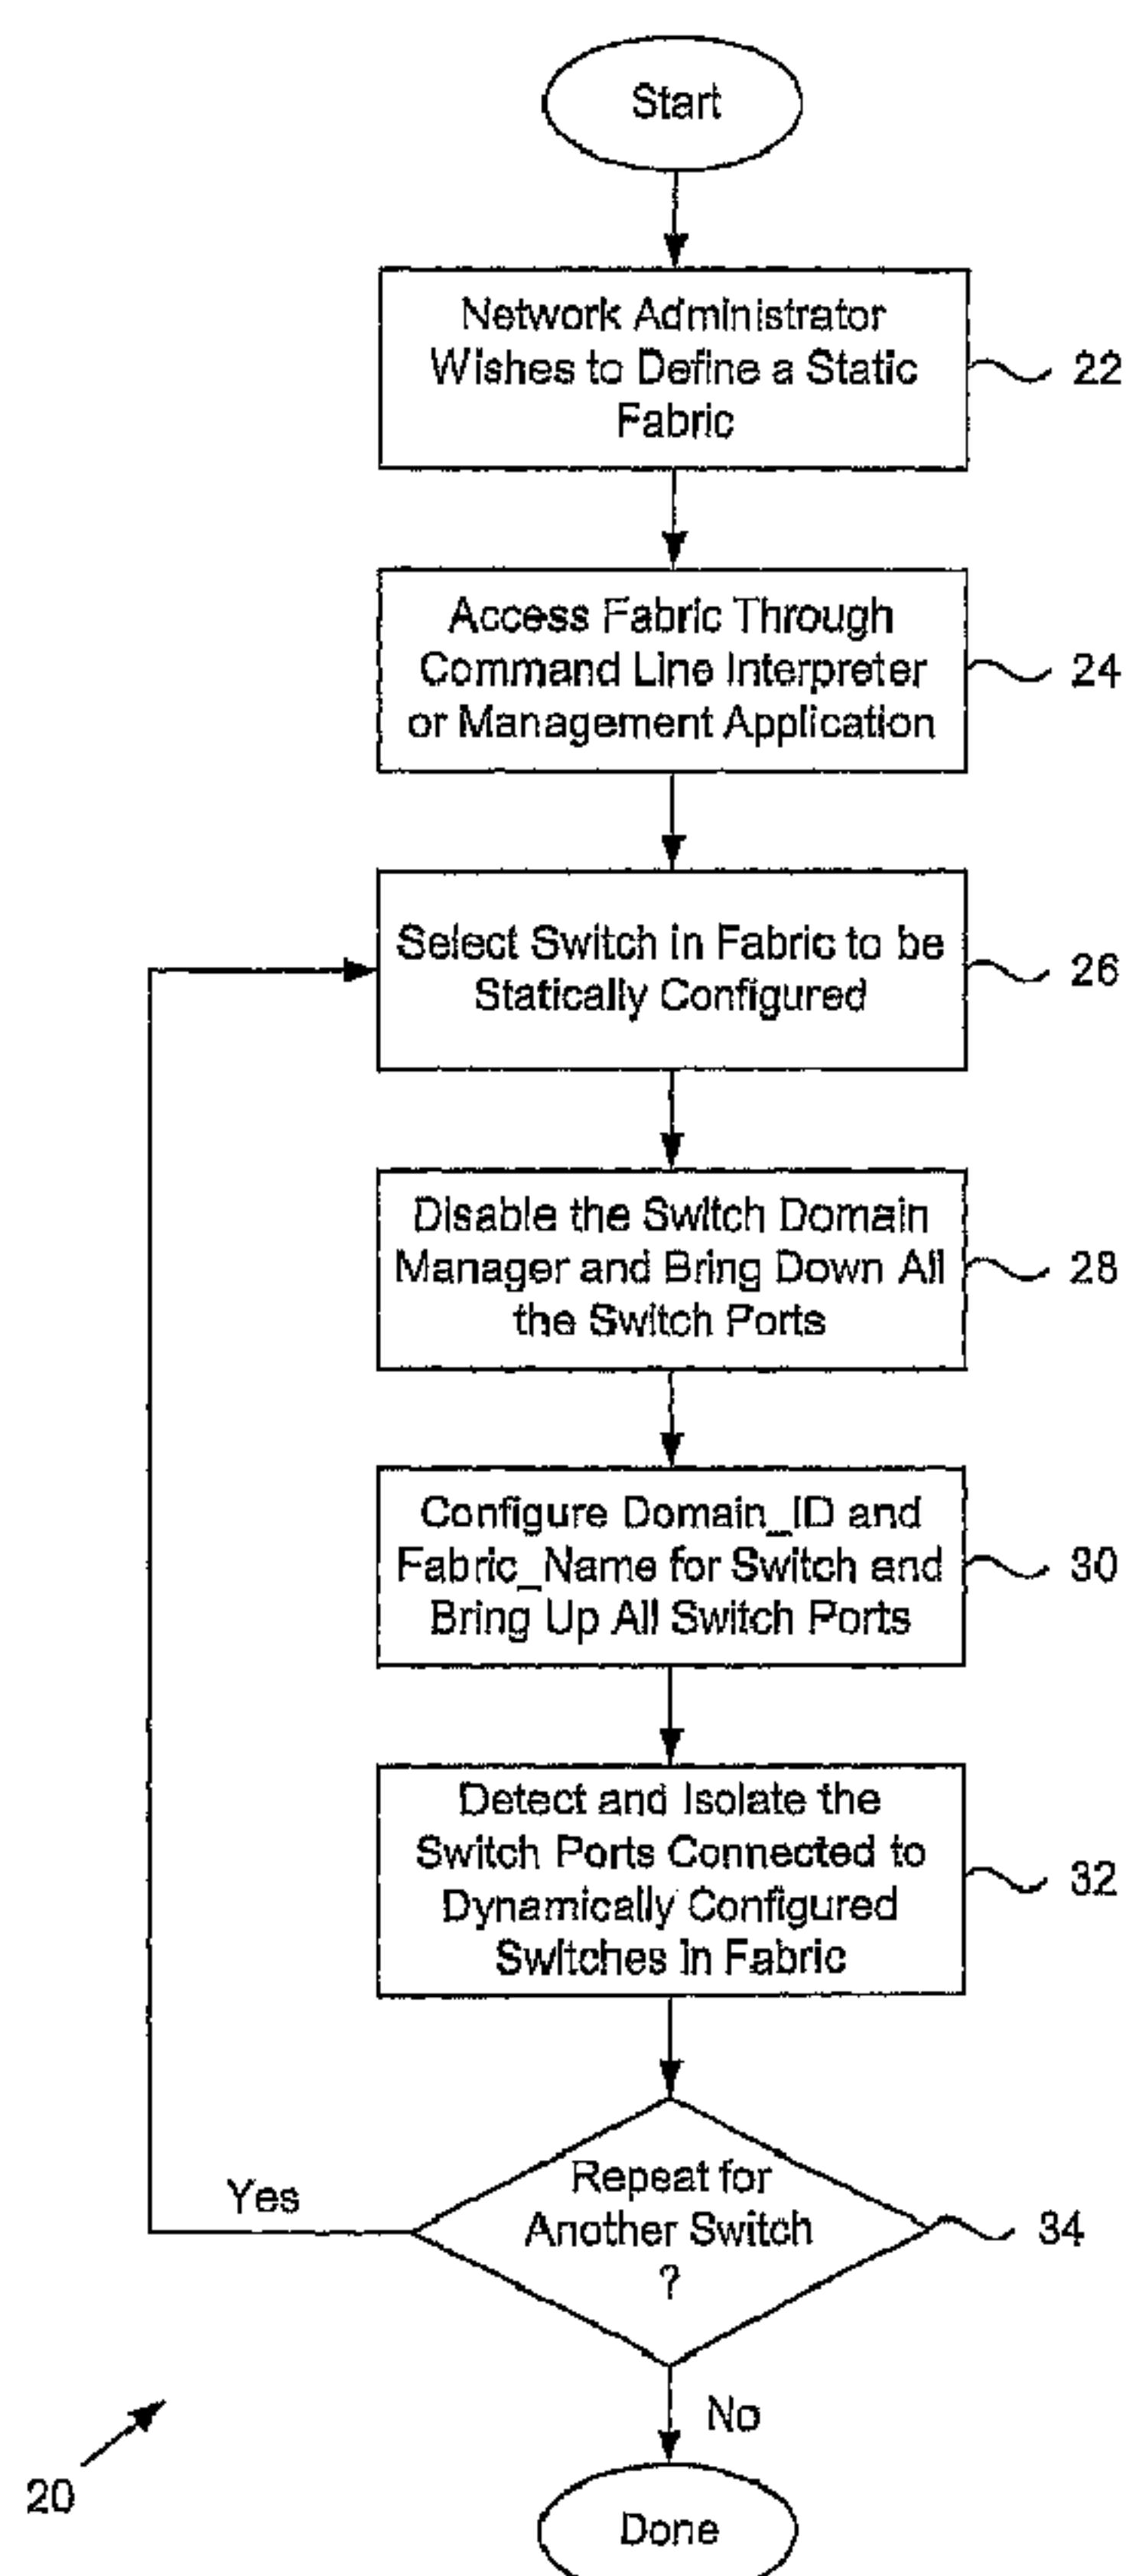
(51) Cl.Int./Int.Cl. *H04L 12/24* (2006.01),
H04L 12/56 (2006.01)

(72) Inventeurs/Inventors:
DESAINTI, CLAUDIO, US;
DIBENEDETTO, MARCO, US

(73) Propriétaire/Owner:
CISCO TECHNOLOGY, INC., US

(74) Agent: GOWLING LAFLEUR HENDERSON LLP

(54) Titre : APPAREIL ET PROCEDE PERMETTANT DE DEFINIR UNE MATRICE DE CANAL DE FIBRES OPTIQUES
STATIQUE
(54) Title: APPARATUS AND METHOD FOR DEFINING A STATIC FIBRE CHANNEL FABRIC



(57) Abrégé/Abstract:

A storage area network and method for defining a static Fibre Channel Fabric that does not require a Principal Switch. The storage area network comprises one or more hosts, one or more storage devices, and a static Fabric connecting the one or more hosts

(57) Abrégé(suite)/Abstract(continued):

and storage devices. Within the static Fabric, the Switches have their Domain_ID and Fabric_Name statically set. The method comprises accessing the Fabric, selecting a Switch in the Fabric, and statically configuring the Domain_ID and Fabric_Name for the selected Switch. The above sequence is repeated for each Switch in the static Fabric. In one embodiment, after being statically configured, the Switch is isolated from any dynamically set Switches in the Fabric. The Switch detects which of its Ports are connected to dynamically set Switches, and then isolates them, while maintaining operational the Ports connected to statically configured Switches.

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

(19) World Intellectual Property Organization
International Bureau(43) International Publication Date
23 October 2003 (23.10.2003)

PCT

(10) International Publication Number
WO 03/088050 A1

(51) International Patent Classification⁷: G06F 12/08, 9/46, H04L 12/56, 12/24

(21) International Application Number: PCT/US03/09332

(22) International Filing Date: 25 March 2003 (25.03.2003)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data: 10/118,386 5 April 2002 (05.04.2002) US

(71) Applicant (for all designated States except US): ANDI-AMO SYSTEMS, INC. [US/US]; 375 East Tasman Drive, San Jose, CA 95134 (US).

(72) Inventors; and

(75) Inventors/Applicants (for US only): DESANTI, Claudio [IT/US]; 4141 Olga Drive, San Jose, CA 95117 (US). DIBENEDETTO, Marco [IT/US]; 3534 Shafer Drive, Santa Clara, CA 95051 (US).

(74) Agent: WEAVER, Jeffrey, K.; Beyer Weaver & Thomas LLP, P.O.Box 778, Berkeley, CA 94704-0778 (US).

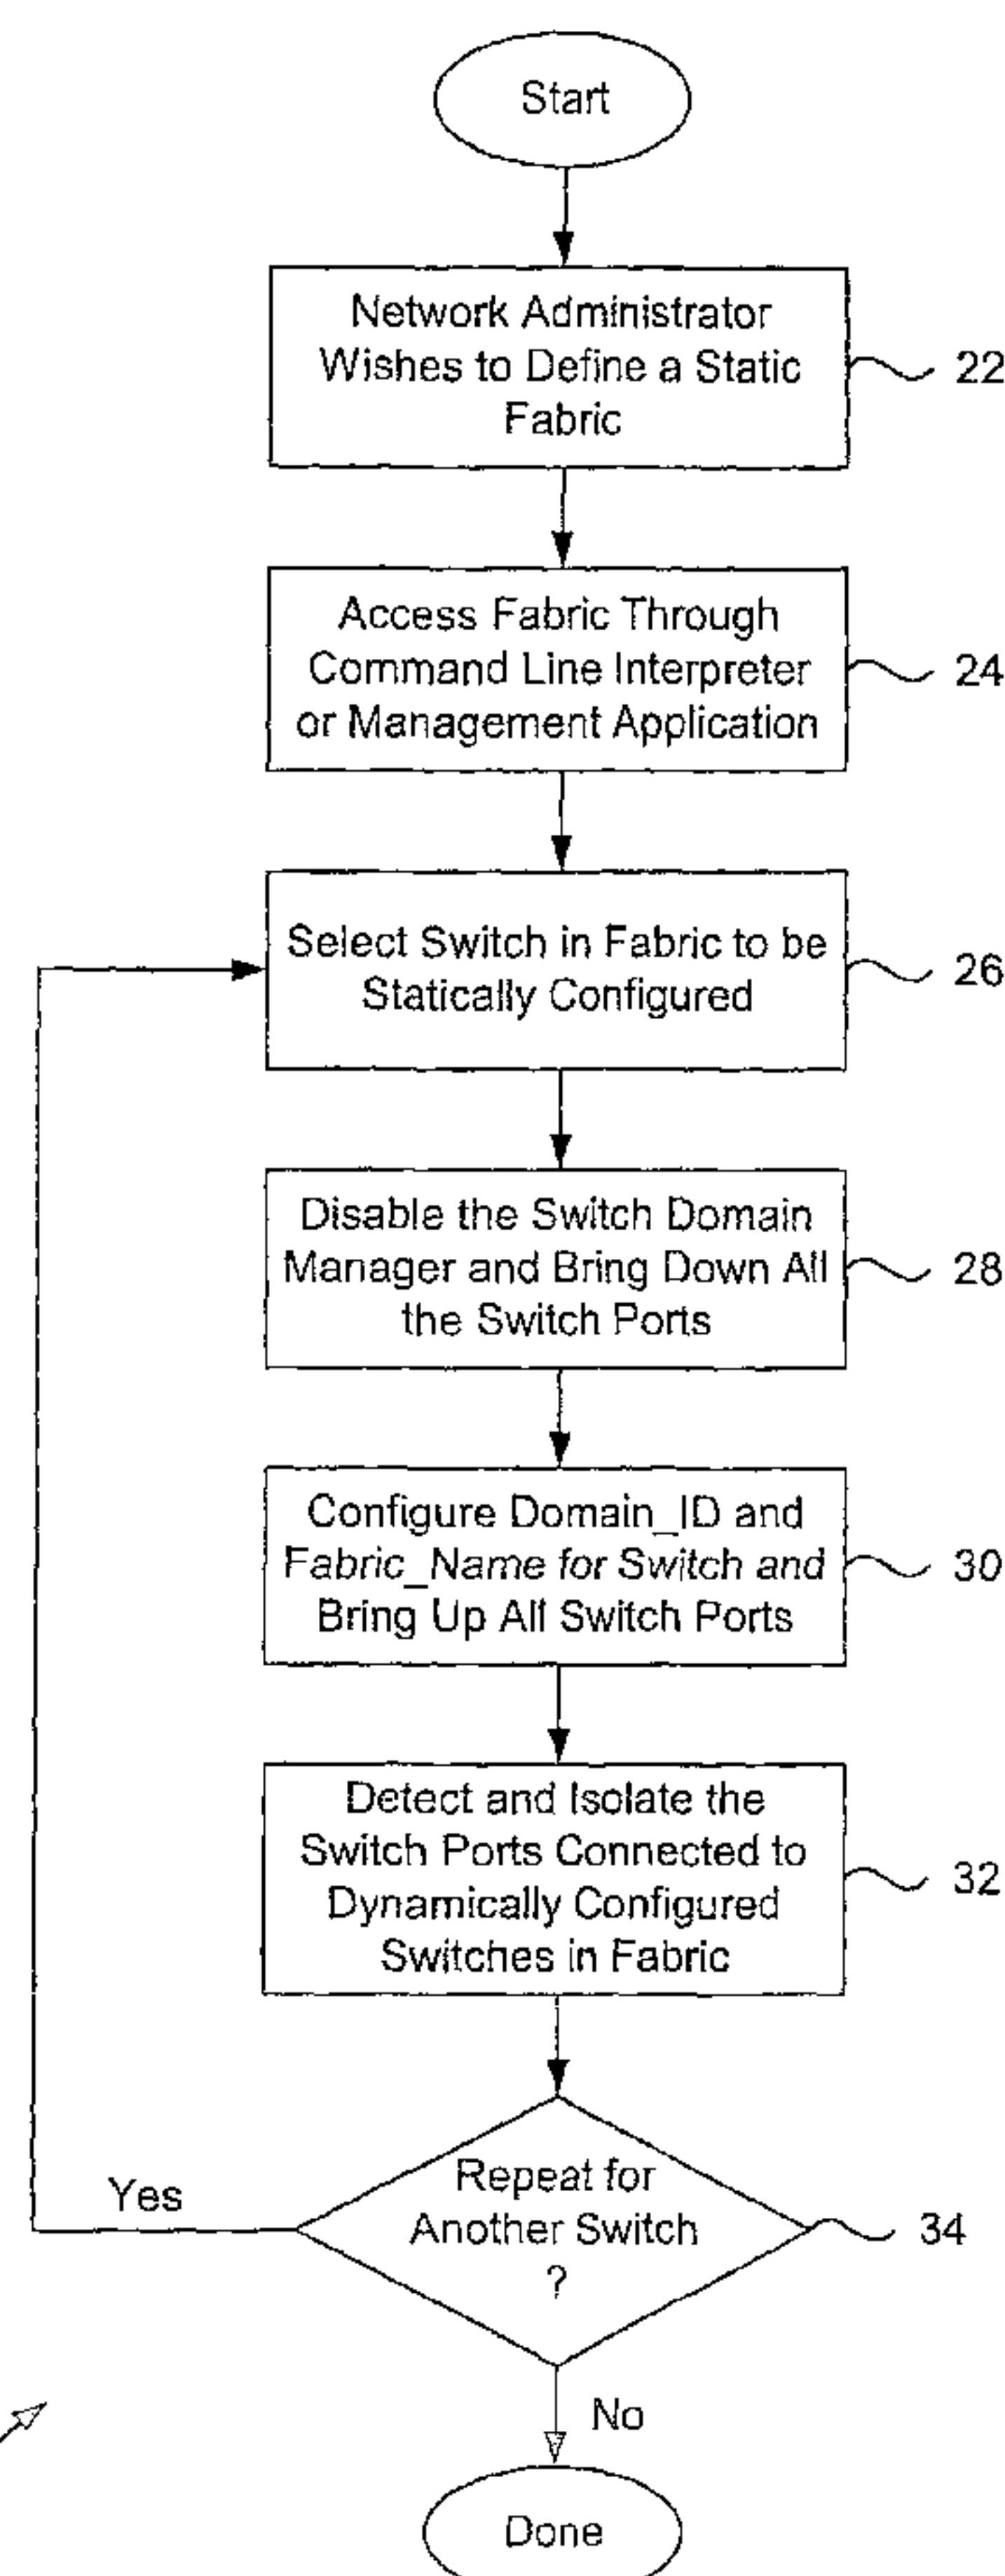
(81) Designated States (national): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NI, NO, NZ, OM, PH, PL, PT,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

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

Published:
— with international search report

[Continued on next page]

(54) Title: APPARATUS AND METHOD FOR DEFINING A STATIC FIBRE CHANNEL FABRIC



(57) Abstract: A storage area network and method for defining a static Fibre Channel Fabric that does not require a Principal Switch. The storage area network comprises one or more hosts, one or more storage devices, and a static Fabric connecting the one or more hosts and storage devices. Within the static Fabric, the Switches have their Domain_ID and Fabric_Name statically set. The method comprises accessing the Fabric, selecting a Switch in the Fabric, and statically configuring the Domain_ID and Fabric_Name for the selected Switch. The above sequence is repeated for each Switch in the static Fabric. In one embodiment, after being statically configured, the Switch is isolated from any dynamically set Switches in the Fabric. The Switch detects which of its Ports are connected to dynamically set Switches, and then isolates them, while maintaining operational the Ports connected to statically configured Switches.

WO 03/088050 A1

WO 03/088050 A1



— before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments

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.

PATENT APPLICATION**APPARATUS AND METHOD
FOR
DEFINING A STATIC FIBRE CHANNEL FABRIC**

5

BACKGROUND OF THE INVENTION**1. Field of the Invention**

The present invention relates to storage area networks, and more particularly, 10 to an apparatus and method for defining a static Fibre Channel Fabric which does not require a Principal Switch.

2. Background of the Invention

With the increasing popularity of Internet commerce and network centric computing, businesses and other organizations are becoming more and more reliant on 15 information. To handle all of this data, storage area networks or SANs have become very popular. A SAN typically includes a number of storage devices, a number of hosts, and a plurality of Switches arranged in a Switching Fabric. The Switches selectively connect the storage devices and the hosts within the SAN.

Most SANs rely on the Fibre Channel protocol for communication within the 20 Fabric. For a detailed explanation of the Fibre Channel protocol and Fibre Channel Switching Fabrics, see FC-FS (Fibre Channel Framing and Signaling) and FC-SW-2 (Fibre Channel Switch Fabric - 2), incorporated by reference herein for all purposes.

In Fibre Channel, each device (hosts, storage devices and Switches) is identified by an unique eight (8) byte wide Node_Name assigned by the manufacturer. 25 When the Fibre Channel devices are interconnected to form a SAN, the Node_Name (together with other parameters) is used to identify each device. Fibre Channel frames are used for communication among the devices in the SAN. The Node_Name, however, is not used by the frames. Instead the Fibre Channel Port of each end device (hosts and storage devices) is addressed via a three (3) byte Fibre Channel address (or 30 FC_ID), allocated dynamically to the end devices by the Fabric. Each end device

acquires its FC_ID by performing a Fabric Login procedure with the Switching Fabric. In this procedure, the end device and the Fabric exchange their credentials and the operating parameters required for a successful communication across the SAN. Initially the Fabric identifies itself by an unique Fabric_Name and the end device by its 5 unique Node_Name. Thereafter the Fabric assigns the FC_IDs to the Ports of the end devices.

The three byte wide Fibre Channel addresses are hierarchically structured in three fields, each one byte long: Domain_ID, Area_ID, and Port_ID. Each Switch within the Fabric is assigned a Domain_ID. The end devices attached to a particular 10 Switch are assigned the Domain_ID of that Switch. The Switch manages the allocation of the Area_ID and Port_ID fields for each end device to guarantee the uniqueness of the assigned addresses in that Domain. For example, if a Switch is assigned a Domain number five and the Switch subdivides its address space in two areas each having three connected end devices, then a possible Fibre Channel address allocation is: 5:1:1, 15 5:1:2, 5:1:3, 5:2:1, 5:2:2, and 5:2:3.

When the Switching Fabric initializes, one of the Switches is selected as the Principal Switch. The Principal Switch assigns the Domain_IDs to all the Switches in the Fabric and its Node_Name becomes the Fabric_Name of the Switching Fabric. To 20 select the Principal Switch, all the Switches exchange with each other a message called Exchange Fabric Parameters (EFP). The EFP contains, among other parameters, the Node_Name of the sending Switch. The Switch with the lowest Node_Name is designated as the Principal Switch. All the other Switches are referred to as non-principal Switches. Once the Principal Switch is selected, it sends to its neighbor 25 Switches a Domain Identifier Assigned (DIA) message, which informs the neighbor Switches that it has been assigned a Domain_ID by the Principal Switch. In reply, the neighbor Switches send a Request Domain Identifier (RDI) message to the Principal Switch. The Principal Switch allocates the Domain_IDs and responds by sending each Switch its Domain_ID. Thereafter, the Switches that received a Domain_ID send a DIA to their neighbor Switches, receive an RDI in reply, and forward the RDI to the 30 Principal Switch, which assigns the Domain_IDs to the requesting Switches. This process continues until all the Switches received a Domain_ID. After having received a Domain_ID, the individual Switches assign the Area_IDs and Port_IDs for each end

device in its Domain. The Fabric configuration is considered completed when all the Switches have been assigned a Domain_ID. Consequently the end devices are all assigned their Area_IDs and Port_IDs.

Fibre Channel allows the merging of two separate Switching Fabrics into one.

5 This happens when a connection is established between two Switches each belonging to a different Fabric. When such an event occurs, the Domain_ID of some of the Switches and the FC_ID of their end devices of the merged Fabric may need to be reassigned. For example, if a Fabric A which includes Domain_IDs one, two and three (1, 2 and 3) is to be merged with a second Fabric B which includes Domain_IDs one 10 and two (1 and 2), then the overlapping Domain_IDs (1 and 2) of one of the Fabrics must be reassigned. When two Fabrics are connected, an EFP message is exchanged across the link that connects them to determine if there is any overlap among the Domain_IDs. Depending on the outcome, one of two things may happen.

If there is any overlap of Domain_ID assignments among the Switches, the link 15 that connects the two original Fabrics is isolated. The link is logically disconnected and is not recognized by the devices, although the physical link still remains. A SAN administrator may then request a disruptive reconfiguration of the joined Fabrics to resolve the Domain_ID conflict. In this case a ReConfigure Fabric (RCF) message is flooded across all the Switches of the two original Fabrics. This stops the delivery of 20 all the data traffic, and each Switch revokes its Domain_ID. Thereafter, a Principal Switch is selected, new Domain_IDs are assigned to the Switches, and new FC_IDs are assigned to the end devices in the same way as described above. In this manner, the two Fabrics are merged into one.

If there is no Domain_ID overlap among the Switches, then a non-disruptive 25 Fabric reconfiguration is automatically performed. A Build Fabric (BF) message is flooded across all the Switches of the two original Fabrics. Data frames delivery is not stopped, and each Switch keeps its Domain_ID. Since the two Fabrics each have a Principal Switch, one of the two has to “resign” from its principal status leaving only one Principal Switch for the merged Fabric. Consequently the Principal Switch 30 selection process described above takes place. Each non-Principal Switch then makes an RDI request to the surviving Principal Switch asking for the same Domain_ID that

it had before the BF message. In this way, the two Fabrics are merged without changing any Switch Domain_ID assignments or any FC_IDs assigned to the end devices.

For one of the two original Fabric, however, the Principal Switch is changed. 5 Consequently the Fabric_Name for those Switches needs to be updated. Given that the Fabric_Name is part of the Fabric Login state information that each end device maintains, the Switches of the loosing Fabric have to re-initialize their end devices to update their status. This process causes a disruption of the data traffic in the Fabric with the losing Principal Switch.

10 A number of problems are associated with the way Domain_IDs are assigned among the Switches of a Fabric under the current Fibre Channel standard. Foremost, the Principal Switch is selected dynamically. In other words, the Principal Switch is selected “on the fly” when the Fabric is initially configured or whenever a change to the Fabric is implemented. Further the RCF and BF processes are disruptive to the 15 Fabric. When the RCF process is invoked, all the traffic across the merged Fabric is halted while the Principal Switch is identified and the Domain_IDs are re-assigned. The BF process is also partially disruptive. Traffic continues in the Fabric where the original Principal Switch continues to be the Principal for the merged Fabric. With the other Fabric, however, the Switches must be updated to reflect that they are included in 20 a new Fabric. Traffic is thus disrupted until the update is completed. Further, Fabrics with Principal Switches are prone to catastrophic problems. For example, if the Principal Switch goes down, it may render the entire Fabric inoperable. Also if a cable is mistakenly plugged into the wrong connector, inadvertently connecting together two Fabrics, it may cause the entire Fabrics to be reconfigured.

25 An apparatus and method for defining a static Fibre Channel Fabric that does not require a Principal Switch is therefore needed.

SUMMARY OF THE INVENTION

To achieve the foregoing, and in accordance with the purpose of the present invention, a storage area network and method for defining a static Fibre Channel 30 Fabric that does not require a Principal Switch is disclosed. The storage area network comprises one or more hosts, one or more storage devices, and a static Fabric

connecting the one or more hosts and storage devices. Within the static Fabric, the Switches have their Domain_ID and Fabric_Name manually set. The method comprises accessing the Fabric, selecting a Switch in the Fabric, and manually setting the Domain_ID and Fabric_Name for the selected Switch. The above sequence is 5 repeated for each Switch in the static Fabric. In one embodiment, after being statically configured, the Switch is isolated from any dynamically set Switches in the Fabric. The Switch detects which of its Ports are connected to dynamically set Switches, and then isolates them, while maintaining operational the Ports connected to statically configured Switches.

10

BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 is a Switching Fabric of a storage area network.

Figure 2 is a flowchart illustrating the sequence for defining a static Fibre Channel Fabric according to the present invention.

15 Figure 3 illustrates an exemplary static Fibre Channel Fabric isolated from a dynamically set Fibre Channel Fabric according to the present invention.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

In the following description, numerous specific details are set forth in order to provide a thorough understanding of the present invention. It will be obvious, however, to one skilled in the art, that the present invention may be practiced without 20 some or all of these specific details. In other instances, well known process steps have not been described in detail in order not to unnecessarily obscure the present invention.

Referring to Figure 1, a Switching Fabric of a storage area network is shown. The storage area network (SAN) 10 includes a Switching Fabric 12 that includes a plurality of Fibre Channel Switches SW1 through SW6. Also included in the SAN 10 25 are a plurality of hosts H1 through H6 and a plurality of storage devices D1 through D5. A tool to manage the Fabric, such as a command line interpreter or a graphic management application 14, is connected to the Fabric 12 through the host H4. The command line interpreter or management application 14 enables a network administrator 16 to manage the Fabric 12 through the host H4 and Switch SW4.

According to various embodiments of the invention, the hosts H1-H6 can be any type of host including but not limited to servers or personal computers running on either the Unix, Windows, or any other computing platform, or a combination thereof. Similarly, the storage devices D1-D5 can be any type of storage device including but not limited to tape back-up systems, emulated tape back-up systems, CD-ROM storage arrays, or one or more disks such as a Redundant Array of Independent Disks (RAID), or a combination thereof. The Switches SW1-SW6 can be any type of Fibre channel Switch such as those commercially available from Brocade of San Jose, CA or Andiamo Systems, the assignee of the present application. It should also be noted that the Fabric 12 as shown in the figure is merely illustrative of a SAN useful for describing the present invention. In no way should its simplicity be construed as limiting the present invention which may be used in any SAN configuration.

The present invention provides a way to manually configure the Switches SW of the Fabric 12. This technique can be used either when the Fabric 12 is initially configured or if a network administrator would like to implement a change across the Fabric 12 after it is in operation.

Referring to Figure 2, a flowchart 20 illustrating the sequence for defining a static Fibre Channel Fabric according to the present invention is shown. When a network administrator 16 wishes to statically define the Fabric 12 (Box 22), the administrator first accesses the Fabric 12. In various embodiments, this is accomplished through either the command line interpreter or management application 14 (Box 24).

Once access to the Fabric 12 is established, the administrator selects a Switch SW to be statically configured (Box 26). The domain manager (the component of the Switch that implements the Principal Switch selection and Domain_ID assignment state machine) of the Switch SW is then disabled by the administrator (Box 28). Disabling the domain manager necessitates a significant change in the working mode of the Switch. Thus when the domain manager is disabled, the Switch brings down all its Ports, including those connecting to other Switches and those connecting to the end devices. Once the domain manager is disabled, the administrator statically configures (i.e. manually writes) in the memory of the Switch an appropriate Domain_ID and

Fabric_Name. After the Domain_ID and Fabric_Name are statically configured, the Switch then brings up all its Ports (Box 30). End devices connected to Ports are assigned their FC_IDs using the statically configured information using the standard Fabric_Login procedure. In an alternative embodiment, the administrator may 5 configure the static Domain_ID and Fabric_Name before disabling the domain manager. In this situation, the Switch will not use the statically configured Domain_ID and Fabric_Name until the domain manager is disabled.

After being statically configured, the Switch is isolated from any dynamically set Switches in the Fabric 12. The Switch detects which of its Ports are connected to 10 dynamically set Switches, and then isolates them, while maintaining operational the Ports connected to statically configured Switches. The detection algorithm is based on the Switch Ports behavior as defined by FC-SW-2. Any Port that receives a message used for a Principal Switch selection, such as EFP, BF, RCF, DIA, RDI, indicates that the Switch that sent that message was dynamically configured. The statically 15 configured Switch in response sends a reject (SW_RJT) message to the dynamic Switch with reason code explanation “E_Port is Isolated”, and then isolates that Port. On receiving this message, the dynamically configured Switch is required to isolate its Port connected to the statically configured Switch. In this manner, the manually configured Switch is logically isolated from any other Switches in the Fabric 12 that 20 are dynamically configured. The Ports linked to any other statically configured Switches do not receive the Principal Switch selection messages (EFP, BF, RCF, DIA, RDI) and are therefore not isolated (box 32).

In decision diamond 34, the administrator decides if another Switch is to be 25 statically configured. If not, the sequence is complete. If yes, then the administrator selects another Switch (Box 26) and the above sequence is repeated. It is responsibility of the administrator to define the same Fabric_Name and non conflicting Domain_IDs for each of the statically configured Switches.

It is useful to clarify the actions performed by the administrator versus the 30 actions performed by the Switches. The administrator is required to disable the domain manager and write the static configuration information for each statically configured Switch. The statically configured Switches automatically performs the Port

actions as described above to logically isolate the statically configured Switches from dynamically set Switches. No other action is required on the part of the administrator to statically define the Fabric 12.

The present invention thus provides a way to statically configure the 5 Domain_ID and Fabric_Name for each Switch in the Fabric 12. Additionally, the present invention provides a method for the Switches to determine if a statically configured Switch is connected to a dynamically configured Switch or to another statically configured Switch. Furthermore, the Switches that are statically set are considered peer-to-peer equals. In other words, the present invention can be used to 10 define a static Switching Fabric that does not require a Principal Switch.

Referring to Figure 3, an exemplary Fabric 40 having several Switches SW that have been statically configured is shown. In this example, three of the Switches SW4, SW5 and SW6 have been statically configured as described above. The remaining Switches SW1, SW2 and SW3, however, have been dynamically 15 configured. As described above, whenever a Switch SW is statically configured, it is logically isolated from any dynamically configured Switch SW. The dashed lines 42 and 44 are representative of the isolated links between Switches SW1 and SW6 and SW3 and SW4 respectively. It should be noted that the isolation is logical and not necessarily physical. In other words, neither data traffic nor configuration messages 20 flow over the isolated link, except for the link initialization message (Exchange Link Parameter, ELP), which may be used to exit from the isolated state in the case that the neighboring Switch becomes a statically configured Switch. Thus, in this example, the Fabric 40 has effectively been divided into two fabrics that do not communicate with one another. Fabric A includes dynamically set Switches SW1-SW3. Fabric B 25 includes statically set Switches SW4-SW6.

Alternatively, the present invention also enables an administrator to dynamically configure a statically configured Fabric 12. This procedure is essentially the same as that described above but in reverse. Initially the administrator is required to select a Switch and enable the domain manager of that Switch. According to one 30 embodiment, when the above occurs the Switch brings down all of its Ports and then brings them up again in accordance with the FC-SW-2 standard. Thus the Ports behave

in accordance with the FC-SW-2 standard and receive and recognize messages used for the selection of a Principal Switch, such as EFP, BF, RCF, DIA and RDI. The Switch thus becomes dynamically configured and is able to communicate with other dynamically configured Switches in the Fabric. The aforementioned process is 5 repeated for each Switch in the Fabric that is to be dynamically configured.

The embodiments of the present invention described above are to be considered as illustrative and not restrictive. For example, the present invention does not necessarily have to be used with a SAN. Rather, it can be used by any type of network with Fibre Channel Switches that are dynamically set. The invention is not to be 10 limited to the details given herein, but may be modified within the scope and equivalents of the appended claims.

CLAIMS

1. A method, comprising:

receiving fabric name and domain name information at a first switch in a

5 fibre channel fabric;

statically configuring the first switch using the fabric name and domain name information;

determining that a second switch in the fibre channel fabric that is connected to the first switch was dynamically configured; and

10 isolating the first switch from the second switch upon determining that the second switch was dynamically configured.

2. The method of claim 1, further comprising disabling a domain manager for the first switch before statically configuring the first switch using the fabric name

15 and domain name information.

3. The method of either claim 1 or claim 2, wherein the second switch is determined to be dynamically configured upon receiving at the first switch a principal switch selection message by the second switch.

20

4. The method of claim 3, wherein a reject message is sent to the second switch after receipt of the principal switch selection message.

5. The method of claim 4, wherein the first switch is isolated from the second switch when a port associated with the first switch ignores principal switch selection messages from the second switch.

6. The method of any one of claims 3, 4 or 5, wherein the principal switch selection message comprises one of the following types of messages: Exchange
30 Fabric Parameter (EFP), Build Fabric (BF), Reconfigure Fabric (RCF), Domain Identifier Assigned (DIA) or Request Domain Identifier (RDI).

7. The method of any one of claims 1 to 6, wherein statically configuring the first switch comprises accessing the fabric through a management application or command line interpreter.
- 5 8. The method of any one of claims 1 to 7, wherein the fabric name and domain name information correspond to a Fabric_Name field and a Domain_ID field.
9. The method of claim 8, further comprising assigning Fibre Channel
10 Identifiers (FC_IDs) to an end device connected to the first switch after the Fabric_Name field and the Domain_ID field have been written into memory of the first switch.
10. The method of any one of claims 1 to 9, further comprising:
15 determining that a third switch in the fibre channel fabric that is connected to the first switch was dynamically configured; and
isolating the first switch from the third switch upon determining that the second switch was dynamically configured.
- 20 11. A computer readable medium having stored thereon instructions, which when executed by a processor cause the processor to implement the method of any one of claims 1 to 10.
12. A switch in a fibre channel fabric, the switch comprising means for
25 executing the instructions of claim 11.
13. A switch in a fibre channel fabric, the switch comprising means for carrying out the method of any one of claims 1 to 10.
- 30 14. A storage area network, comprising:
one or more hosts;

one or more storage devices; and
a switching fabric connects the one or more hosts and storage devices,
the switching fabric having at least one switch according to either claim 12 or
claim 13.

5

15. The storage area network of claim 14, wherein the at least one switch is a
peer-to-peer equal within the static fabric.

16. The storage area network of either claim 14 or claim 15, further
10 comprising a management application or command line interpreter coupled to the
fabric for accessing the fabric.

17. The storage area network of claim 16, wherein the management
application or command line interpreter selects the first switch from a plurality of
15 switches in the fabric.

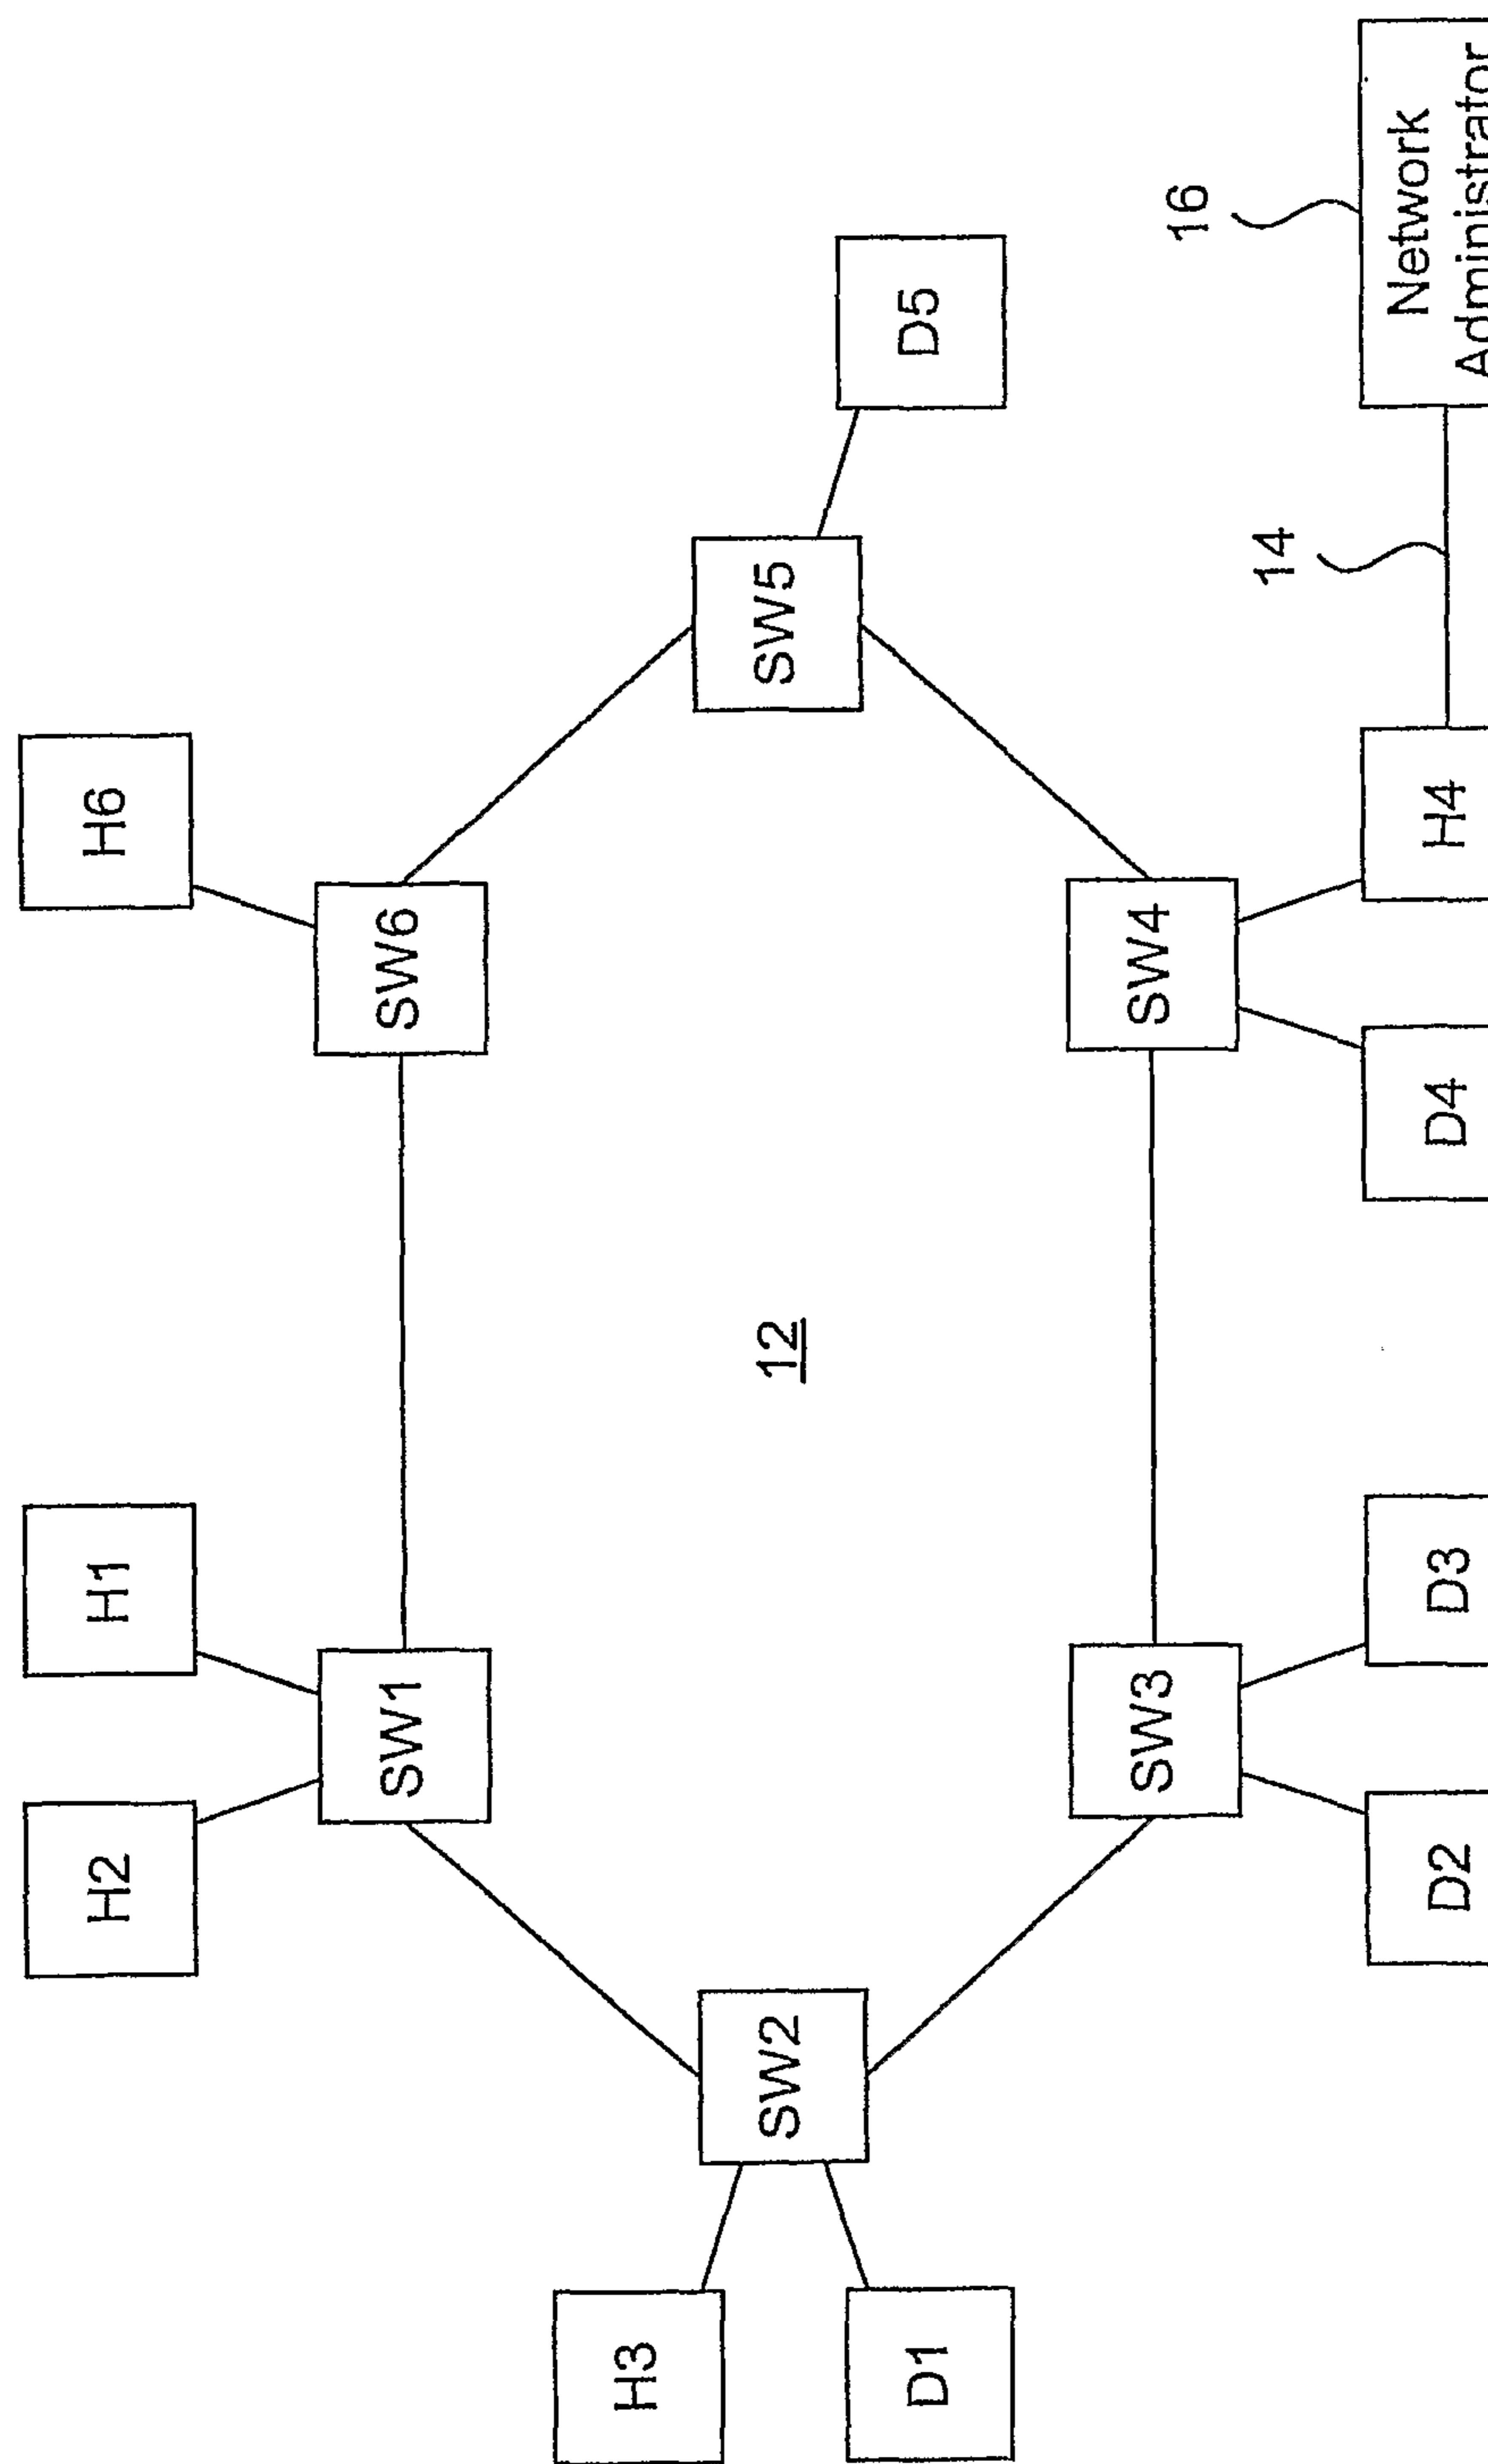


Figure 1

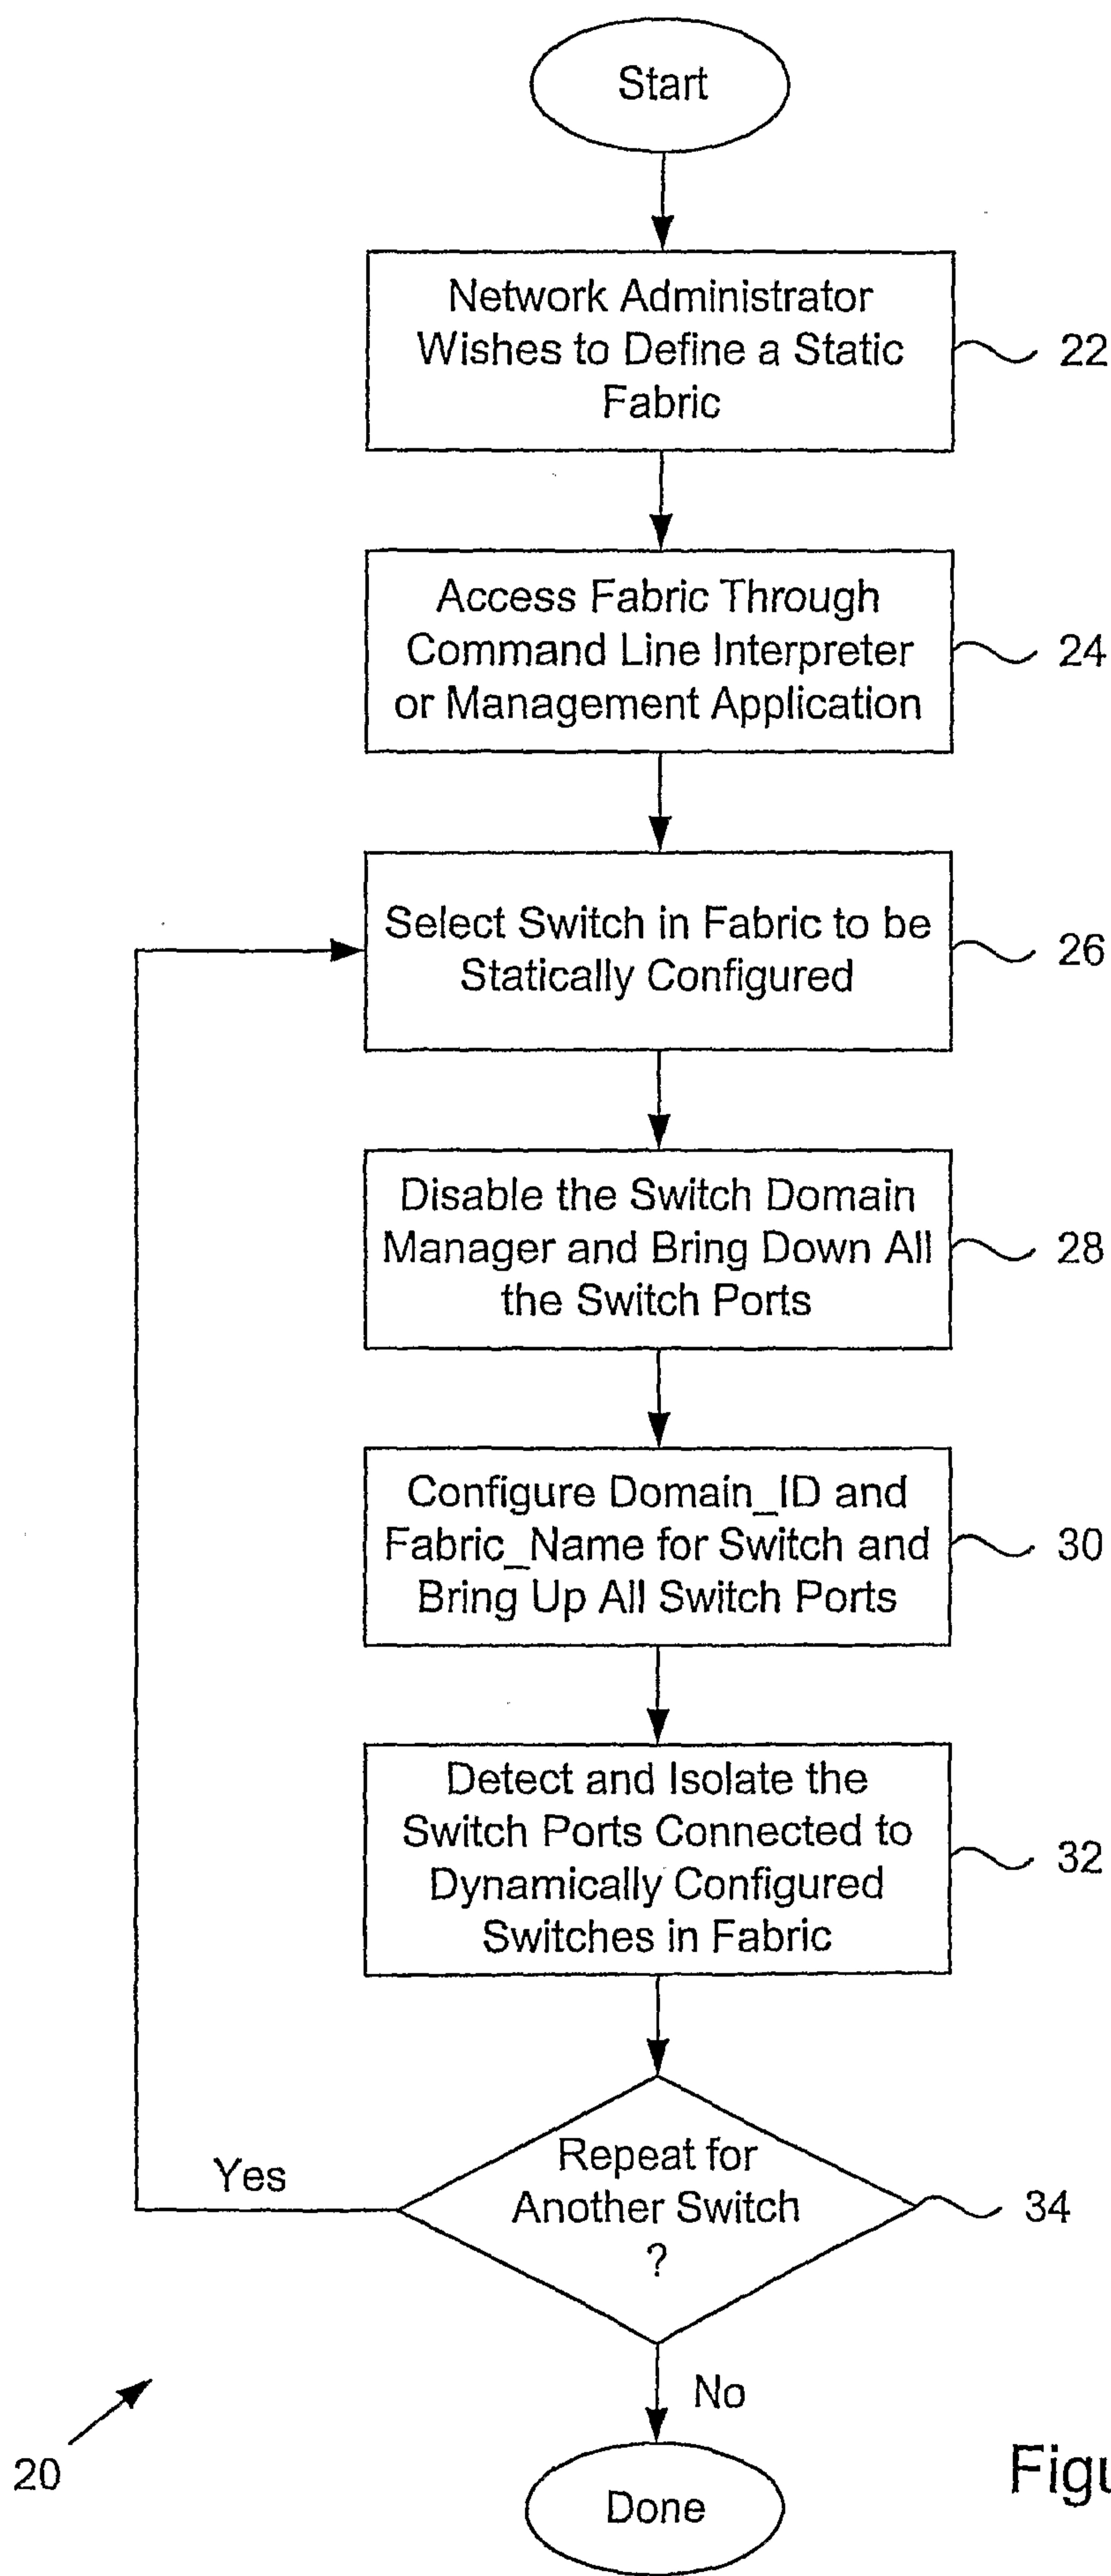


Figure 2

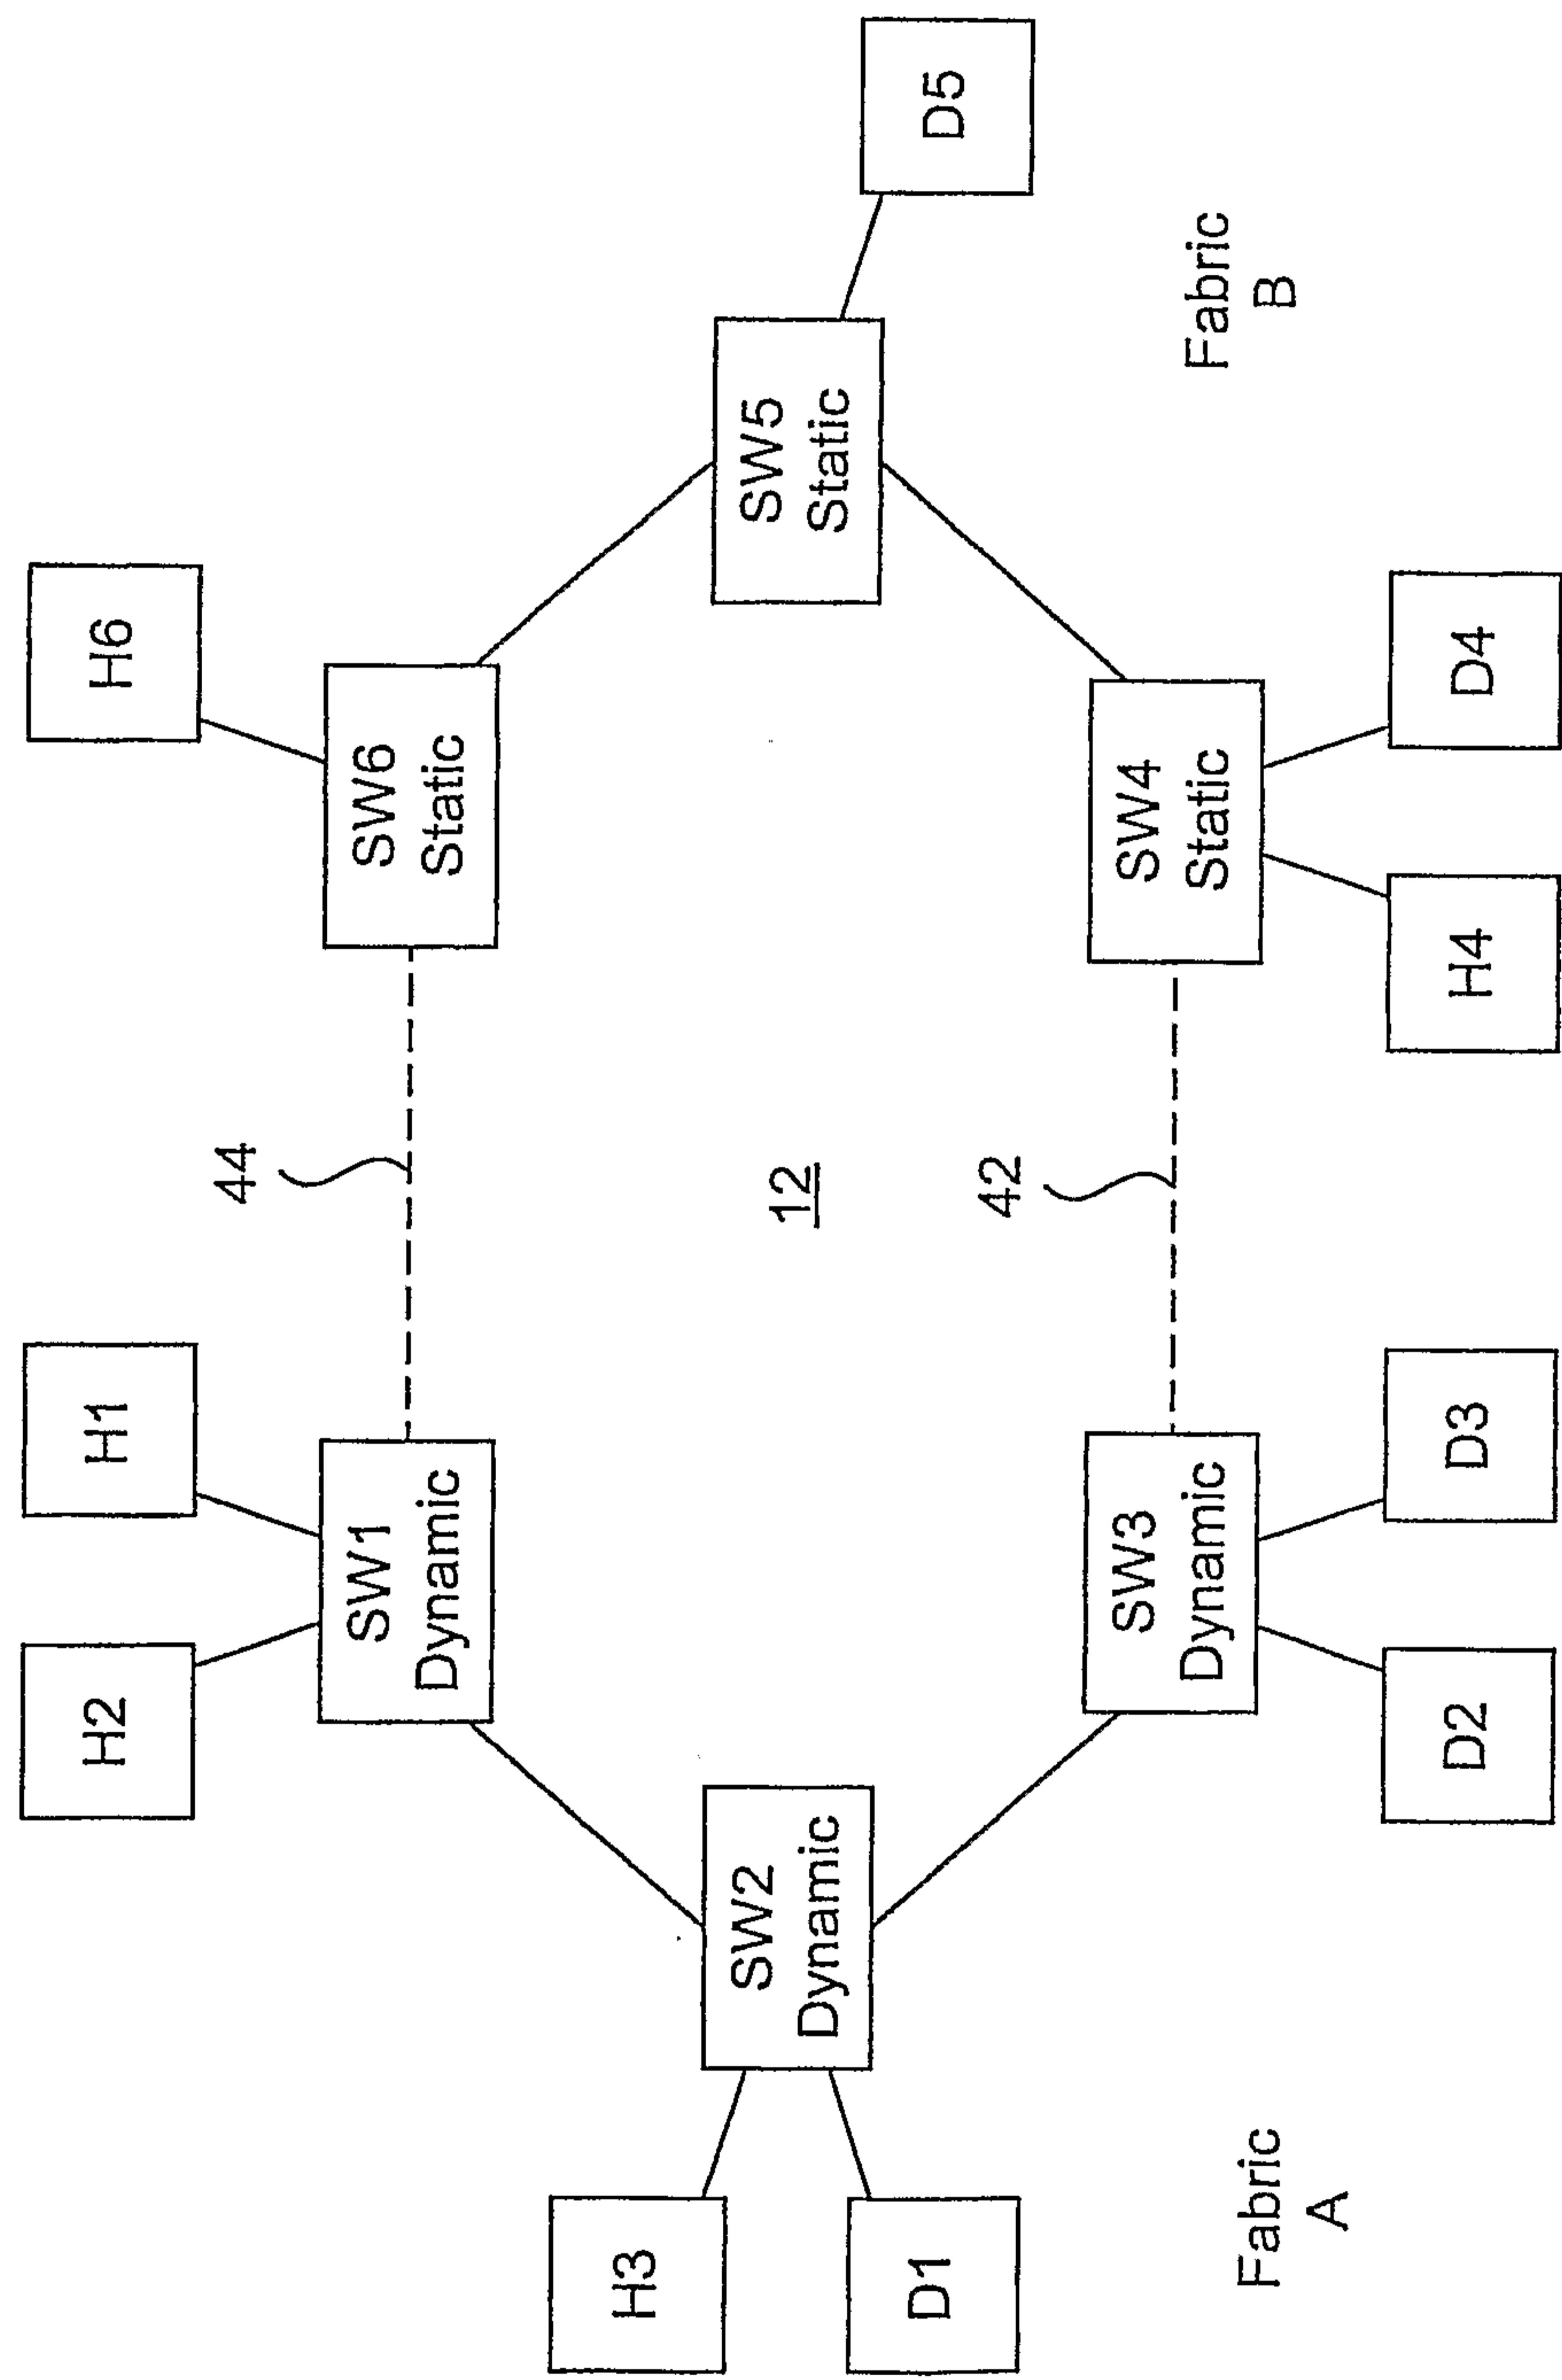


Figure 3

40

