

Description

Title of Invention: CONNECTION PATH MANAGEMENT SYSTEM, CONNECTION
PATH MANAGEMENT METHOD, AND CONNECTION PATH MANAGEMENT
PROGRAM

5

Technical Field

[0001] The present invention relates to a connection path management system, a
connection path management method, and a connection path management program

10 Background Art

[0002] Conventionally, there is a need for establishing a flow in order to perform
communication such as the one using Openflow (registered trademark). There is
provided a system in which flow control is performed by a controller in charge of a
network system (see, e.g., Non-Patent Literature 1). In recent years, this system is
15 used by an architecture called as an SDN (Software Defined Network).

[0003] There is provided a system in which a flow is identified based on a flow
identifier in a data frame received through an interface, and when the flow does not
match a flow that is originally to be received through the interface, the data frame is
determined to have been received through an unprescribed interface (see, e.g., Patent
20 Literature 1).

Citation List

Patent Literature

[0004] Patent Literature 1: JP 2013-115695 A

25 Non-patent Literature

[0005] Non-patent Literature 1: "OpenFlow Switch Specification Version 1.0.0",
OPEN NETWORKING FOUNDATION, December 31, 2009

Summary of Invention

5 Technical Problem

[0006] In the conventional systems, it cannot be prevented that, with respect to one connection between terminals, separate flows are set as an outgoing path and a return path. It is enough to set one bidirectional flow for one connection. When two bidirectional flows are set, network resources that are twice as those in the case of the
10 one bidirectional flow are consumed, thus leading to a waste. When two unidirectional flows are set, a management burden increases.

[0007] An object of the present invention is to reduce an amount of network resource consumption.

Solution to Problem

15 [0008] A connection path management system according to one aspect of the present invention includes:

a setting unit to set, in association with a connection between terminals, a bidirectional connection path for forwarding packets including information for identifying the connection, between interfaces of a network to which the packets are
20 input or from which the packets are output; and

an aggregation unit to set a connection path between a first input interface and a second input interface in association with one connection when a second packet including information for identifying the one connection is input to the second input interface after a connection path between the first input interface and a first output
25 interface has been set by the setting unit in association with the one connection, and

delete the other connection path associated with the one connection, the first input interface being an interface to which a first packet including the information for identifying the one connection has been input, the first output interface being an interface from which the first packet has been output, the second input interface being
5 an interface different from the first output interface, the second packet having been transmitted in a direction opposite to the first packet.

Advantageous Effects of Invention

[0009] In the present invention, when a situation occurs where a plurality of connection paths are set for one connection between the terminals, those plurality of
10 connection paths are aggregated into one connection path. Therefore, according to the present invention, it becomes possible to reduce a consumption amount of network resources. The “connection paths” are same as the flows mentioned above.

Brief Description of Drawings

15 [0010] Fig. 1 is a diagram illustrating a typical example of a network configuration.

Fig. 2 is a diagram illustrating an example of a network configuration for comparison with a first embodiment.

Fig. 3 is a diagram illustrating an example of the network configuration for comparison with the first embodiment.

20 Fig. 4 is a diagram illustrating an example of the network configuration for comparison with the first embodiment.

Fig. 5 is a block diagram illustrating a configuration of a connection path management system according to the first embodiment.

Fig. 6 is a flowchart illustrating operations of the connection path
25 management system according to the first embodiment.

Fig. 7 is a diagram illustrating an example of a network configuration according to the first embodiment.

Fig. 8 is a diagram illustrating an example of a network configuration according to a variation example of the first embodiment.

5 Fig. 9 is a diagram illustrating an example of the network configuration according to the first embodiment.

Fig. 10 is a diagram illustrating an example of the network configuration according to the first embodiment.

10 Fig. 11 is a diagram illustrating an example of the network configuration according to the first embodiment.

Fig. 12 is a diagram illustrating an example of the network configuration according to the first embodiment.

Fig. 13 is a diagram illustrating an example of the network configuration according to the first embodiment.

15 Fig. 14 is a diagram illustrating an example of the network configuration according to the first embodiment.

Fig. 15 is a diagram illustrating a hardware configuration of a network control apparatus according to the embodiment of the present invention.

20 Description of Embodiments

[0011] Hereinafter, an embodiment of the present invention will be described using drawings. Same or corresponding portions in the respective drawings are given same reference signs. In the description of the embodiment, descriptions of the same or corresponding portions are omitted or simplified as appropriate.

25 [0012] First Embodiment

This embodiment relates to connection path control when a network system capable of establishing a connection path based on detailed identification information of a packet on a connection and a network system where a plurality of apparatuses being in a redundant relationship operate are connected through a plurality of redundant

5 interfaces and apparatuses. In the network system where the plurality of apparatuses being in the redundant relationship operate, different outgoing and return paths may be used as communication paths of the connection. “A path being the same” or a “same path” means that a packet is transmitted or received through a same node or a same apparatus, regardless of a communication direction. “A path being different” or a

10 “different path” means opposite to the “the path being the same” or the “same path”. [0013] First, a description will be given about a phenomenon in which passage of different outgoing and return paths is made in a network having a redundant configuration. A description will be directed to a network configuration where that phenomenon occurs, with reference to Fig. 1. An example of Fig. 1 is a typical

15 example.

[0014] A Router-A, a Router-B, a Router-E, and a Router-F are routers. A Switch-C, Switch-D, a Switch-G, and a Switch-H are layer 2 switches. A Subnet-Z is an IP (Internet Protocol) subnet.

[0015] An IP-1 is an interface of the Router-E connected to the Switch-C. An IP-2

20 is an interface of the Router-F connected to the Switch-D. An IP-3 is an interface of the Router-E connected to the Switch-D. An IP-4 is an interface of the Router-F connected to the Switch-C. An “interface” means a port to which a packet is input or from which the packet is output.

[0016] A Path-1 is a path through which a packet sequentially passes through the

25 Router-B, the Switch-C, the Router-E, and the Switch-G to reach the Subnet-Z. A path

2 is a path through which a packet is transmitted from the Subnet-Z and sequentially passes through the Switch-H, the Router-F, the Switch-C, and the Router-B.

[0017] Herein, the Router-A and the Router-B are in a redundant relationship to each other, while operating normally. That is, during a period in which both of the Router-A and the Router-B operate normally, a packet forwarding load is shared. When one of the Router-A and the Router-B has failed, the router in operation wholly undertakes forwarding of a packet that will pass through the router that has failed. The Switch-C and the switch-D are in the same relationship to each other, the Router-E and the Router-F are in the same relationship to each other, and the Switch-G and the Switch-H are in the same relationship to each other.

[0018] When dynamic routing to automatically search a subsequent destination is applied in an IP, it is possible to automatically identify a path fault and not to select a destination with the fault. In that case, a plurality of destinations may be searched in a normal state. Then, it may also be so arranged that by setting "cost values", the plurality of destinations are distinguished to constantly select a top-priority destination. In the case of the redundant configuration with equality as illustrated in Fig. 1, however, it is so configured that no difference is set for cost values, and that load distribution is made at a normal time. In that case, for each of the Router-A and the Router-B, four equivalent destinations are registered as follows, as IP route entries for the Subnet Z:

- (Destination 1) Subnet-Z NextHop: IP-1
- (Destination 2) Subnet-Z NextHop: IP-2
- (Destination 3) Subnet-Z NextHop: IP-3
- (Destination 4) Subnet-Z NextHop: IP-4

[0019] Normally, each of the Router A and the Router B performs packet forwarding so that packet assignment is equally performed to the above-mentioned four addresses

according to an algorithm independently implemented by a vendor. Normally, a destination is computed, based on connection identification information for identifying an IP connection such as transmission and reception IP addresses, a protocol, and TCP (Transmission Control Protocol) port numbers. By doing so, destination routing is performed based on each connection rather than each packet, so that the order of packets of a same connection is not interchanged in this case. An IP router, however, is not conscious of the connection when performing forwarding. The IP router just forwards a packet input through an arbitrary interface according to a routing table. Accordingly, the IP router does not identify bidirectional packets of one connection and does not forwards the packets so that the packets mutually pass through a same interface. Even if such a situation occurs when a plurality of forwarding destinations are present, the situation occurs just by chance. Accordingly, when a response for a packet that has passed through the Path-2 as an outgoing path is forwarded to the Router-B, the Router-B selects the IP-4 as a destination, and is not conscious that passage through the Router-F is in a direction opposite to the Path-2. The Router-B just determines the destination by performing computation so that packet assignment is equally performed to the above-mentioned four destinations (of destination 1 to destination 4). In Fig. 1, the IP-1 is selected. The response consequently passes through the Router-E along the Path-1. In other words, the result is that different paths are used as the outgoing path and a return path. Further, a packet of the return path arrives at the Router-B in the example illustrated herein. This arrival occurs by chance.

[0020] Subsequently, a configuration will be given in which a network system constituted from a network having such a redundant configuration as mentioned above, network apparatuses to set connection paths for packet forwarding, and a control apparatus to perform control of those connection paths is connected using a plurality of

redundant interfaces. An example of the configuration will be described, with reference to Fig. 2. The example in Fig. 2 is an example for comparison with this embodiment.

[0021] A Router-1 and a Router-2 are routers. An SDN-SW1, an SDN-SW2, an SDN-SW3, and an SDN-SW4 are SDN switches. A Term-0, a Term-1, a Term-2, and a Term-3 are terminals. The terminals may be server apparatuses, not being limited to client computers. The Term-0 also indicates an interface of the Term-0 that is the terminal, which is connected to a network 121. A Term-2-1 is an interface of the Term-2, which is connected to the SDN-SW3. A Term-2-2 is an interface of the Term-2, which is connected to the SDN-SW4. The Term-3 also indicates a bundle of interfaces of the Term-3 that is the terminal, which is connected to the SDN-SW4.

[0022] Herein, a network 110 constituted from the SDN switches that are apparatuses whereby the connection paths are set will be referred to as an SDN. The network 121 such as a LAN (Local Area Network) or a WAN (Wide Area Network) and a different network 120 connected to the network 121 or included in the network 121 and constituted from routers and layer 2 switches illustrated in Fig. 1 are referred to as an existing network. A network control apparatus 150 being an SDN controller is connected to the SDN.

[0023] The plurality of interfaces that connect the existing network and the SDN are in a redundant relationship. When a packet is input to one of the routers and when the packet is forwarded from the router to the SDN, equal packet assignment is performed. Accordingly, a phenomenon may occur in which an outgoing path and a return path become different paths and the interface through which a packet is input to the SDN from the existing network and the interface through which a packet is output to the existing network from the SDN differ. A side of the SDN opposite to the existing

network is connected to the terminals. There also exists the terminal beyond a side of the existing network opposite to the SDN. Fig. 2 illustrates, as the terminals, the Term-0 and the Term-1 each including one interface, the Term-2 including two interfaces with equal redundancy, and the Term-3 to which link aggregation causing a plurality of links to appear as if they were one link has been applied. Herein, in both of the cases of the equal redundancy and the link aggregation, reciprocating packets of one connection do not necessarily pass through a same interface.

[0024] When the existing network or the terminals connected to the SDN using two or more of the interfaces is or are present as mentioned above, a situation may occur in which reciprocation is performed using different paths. A connection for the Term-1 may be made by using different paths that pass through the interfaces which are different on a side of the existing network.

[0025] Communication requests of the following four types will be considered:
 (Type 1) a communication request from the Term-0 to the Term-1
 (Type 2) a communication request from the Term-1 to the Term-0
 (Type 3) a communication request from the Term-0 to the Term-2
 (Type 4) a communication request from the Term-2 to the Term-0

[0026] Since the Term-2 and the Term-3 are regarded to be equivalent in the sense that the Term-2 and the Term-3 are connected to the SDN through a plurality of interfaces that are equal, the Term-2 and the Term-3 are classified to be the same. Even if a connection similar to the Term-2 or the Term-3 is present in place of the existing network, it follows that the connection of a network or a terminal is regarded to be equivalent because the network or the terminal is the one connected through the interfaces regarded to be equal. A situation in which reciprocation is performed using different paths will be considered, for each of the above-mentioned four types.

[0027] It is assumed herein that each of the routers such as the Router-1 and the Router-2 and the terminals such as the Term-2 and the Term-3 performs an operation of outputting packets of a same transmission destination through a same interface. This operation is usually performed in order to maintain the order of the packets.

5 [0028] In the (Type 1), a request packet is transmitted from the Term-0 to the Term-1. If the SDN constructs a bidirectional connection path, a response packet from the Term-1 is forwarded to the side of the existing network through the connection path of an opposite direction with no problem. Thus, the problem of different routing is solved by a basic operation of establishing a bidirectional connection when the
10 communication request arrives.

[0029] In the (Type 2), a request packet is transmitted from the Term-1 to the Term-0. A response packet from the Term-0 may be forwarded by the router provided immediately before the SDN. Thus, the response packet may pass through a different path. As illustrated in Fig. 3, after a CP-1 has been established, a CP-2 may be
15 established. In this case, the request packet passes through a Path-1. The response packet passes through a Path-2.

[0030] In the example in Fig. 3, the CP-1 is a connection path between the SDN-SW3 and the SDN-SW2. The CP-2 is a connection path between the SDN-SW1 and the SDN-SW3. The Path-1 is a path whereby the packet is transmitted from the
20 Term-1 and reaches the Term-0 through the CP-1, the Router-2, and the network 121, sequentially. The Path-2 is a path whereby the packet is transmitted from the Term-0 and reaches the Term-1 through the network 121, the Router-2, and the CP-2.

[0031] Assume that a request packet is transmitted from the Term-0 to the Term-2 in the (Type 3). Then, even if the request packet has passed through the Term-2-1, a
25 response packet from the Term-2 may pass through the Term-2-2, or the reverse may

occur. On the side of the existing network, no problem will occur for the same reason as in the case of the “Type 1” if the same interface is used on the side of the Term-2.

However, if the different interfaces are used on the side of the Term-2, the packets may pass through the different interfaces even on the side of the existing network. As

5 illustrated in Fig. 4, two connection paths of a CP-1 and a CP-2 are constructed, and either ends of these two connection paths may not be the same. In this case, a request packet passes through the Path-2. A response packet passes through the Path-1.

[0032] In the example in Fig. 4, the CP-1 is a connection path between the SDN-SW4 and the SDN-SW2. The CP-2 is a connection path between the SDN-SW1
10 and the SDN-SW3. The Path-1 is a path whereby the packet is transmitted from the Term-2-2 and reaches the Term-0 through the CP-1, the Router-2, and the network 121, sequentially. The Path-2 is a path whereby the packet is transmitted from the Term-0 and reaches the Term-2-1 through the network 121, the Router-2, and CP-2.

[0033] In the (Type 4), a request packet is transmitted from the Term-2 to the Term-0.
15 A response packet from the Term-0 may pass through a different path for the same reason as in the case of the “Type 2”. It depends on connection path control of the network control apparatus 150 whether the response packet passes through the Term-2-1 or the Term-2-2 for connection to the Term-2 after the response packet has been input to a different interface from the existing network.

20 [0034] If an outgoing path and a return path are different, connection paths that are different between the outgoing path and the return path are constructed in the SDN in a basic SDN implementation. The two connection paths are handled by management of the network control apparatus 150.

[0035] If a packet is input to one of the interfaces, each SDN switch determines
25 whether the packet is the one associated with the connection path set for the same

interface. If that packet is a new one not associated with the connection path that has been set, the SDN switch sends a new connection request to the network control apparatus 150. The network control apparatus 150 sets a new connection path according to the new connection request. For handling the problem of different routing, an operation of leaving the connection path to be newly constructed and deleting the connection path that is old by the network control apparatus 150 may be conceived when the different routing has become clear.

[0036] In the (Type 2), the interface at one of both ends of the connection paths is common. Therefore, a new connection request is generated just for the interface that is not common. Accordingly, once the network control apparatus 150 takes an action of leaving the new connection path and deleting the other connection path, a new connection request is not further generated so long as other condition change or the like is not made.

[0037] The (Type 3), however, has a problem that will be given below.

[0038] As illustrated in Fig. 4, the interfaces may not be the same at the either ends of the connection paths. That is, the CP-2 is first constructed as a connection path for the Path-2 of an outgoing path. Then, the CP-1 is constructed as a connection path for the Path-1 of a return path. In the above-mentioned operation, the CP-2 is to be deleted. However, the existing network tries to maintain the Path-2. Therefore, when a packet is transmitted from the Term-0, the packet is input to the SDN-SW1 in the SDN. If the CP-2 is constructed again, the CP-1 is to be deleted in the above-mentioned operation. Then, the process of constructing and deleting the CP-2 and CP-1 is repeated until the state of the SDN is changed and a connection path connected to the Term-2-2 is constructed in place of the CP-2 connected to the Term-2-1. This makes it difficult to perform communication, depending on a load state of the network control

apparatus

[0039] The (Type 4) also has the same problem as the (Type 3).

[0040] Assume that a plurality of connection paths are detected which establish a condition (a) of being within a same tenant and a condition (b) of being a same

5 connection determined on an identifiable highest-order layer. Then, in this embodiment, a procedure of newly constructing a connection path between the interfaces to which packets causing establishing of the connection path have been input and deleting the old connection paths is taken.

[0041] The “same tenant” means a same network for each service. In the SDN, a
10 logically independent network may be configured by network virtualization, irrespective to a physical configuration. Duplication of IP addresses is permitted between independent networks. A plurality of VLANs (Virtual Local Area Networks) may also be constructed in each independent network. Accordingly, a condition regarding (a) indicates that it is necessary for the condition regarding (b) to be established in one
15 logically independent network.

[0042] The “same connection determined on the identifiable highest-order layer” means connections having a same pair of transmission and reception IP addresses and a same pair of transmission and reception port numbers in a TCP/IP. It is assumed that, in a connection path entry table of the network control apparatus 150 that is the SDN
20 controller, information indicating a direction of a packet that has triggered construction of each connection path and an interface to which the packet has been input are described in predetermined fields. As the information indicating the direction of the packet, transmission and reception IP addresses may be used.

[0043] A configuration of a system according to this embodiment, operations of the
25 system according to this embodiment, and effects of this embodiment will be

sequentially described.

[0044] *** Description of Configuration ***

A description will be given about the configuration of a connection path management system 100, which is the system according to this embodiment, will be
5 described with reference to Fig. 5.

[0045] The connection path management system 100 includes the network 110, a setting unit 151, and an aggregation unit 152. In this embodiment, the setting unit 151 and the aggregation unit 152 are included in the network control apparatus 150.
However, as in a variation example that will be described later, the setting unit 151 and
10 the aggregation unit 152 may be provided at separate apparatuses.

[0046] The network 110 includes the interfaces to which packets including information for identifying a connection between the terminals are input from outside the network 110 or from which the packets are output to outside the network 110. The “information for identifying the connection” is specifically connection identification
15 information for identifying an IP connection described above. The different network 120 and a different network 130 are connected to the network 110. Each terminal is connected to one of the different networks 120 and 130. In this embodiment, as a technique for the network 110, the SDN is employed. However, any technique may be employed if a connection path or a flow is set by the technique.

[0047] The setting unit 151 sets, in association with a connection between terminals, a bidirectional connection path for forwarding packets including information for identifying the connection, between the interfaces of the network 110 to which the packets are input or from which the packets are output. It is assumed herein that a first packet including information for identifying one connection C1 has been input to one
25 interface of the network 110. In this case, the setting unit 151 sets a connection path

between a first input interface, which is the interface to which the first packet has been input, and a first output interface, which is the interface from which the first packet has been output, in association with the connection C1.

[0048] The aggregation unit 152 sets a connection path between the first input
5 interface and a second input interface, which is the interface different from the first output interface, in association with the connection C1 when a second packet including the information for identifying the connection C1 and transmitted in a direction opposite to the first packet is input to the second input interface after the connection path between the first input interface and the first output interface has been set by the setting
10 unit 151 in association with the connection C1. The aggregation unit 152 deletes the other connection path associated with the connection C1. With this arrangement, when a problem occurs that an outgoing path for transmitting the first packet and a return path for transmitting the second packet will be different paths or have become the different paths with respect to the connection C1, the aggregation unit 152 may cope with the
15 problem. In this embodiment, the aggregation unit 152 detects occurrence of the problem at a timing when the second packet has been input to the second input interface. Therefore, the aggregation unit 152 may cope with the problem in advance. “In advance” means before setting of the connection path that will cause the different paths. “Cope with the problem” means to aggregate the connection paths into one connection
20 path, without wasting resources. As will be described later, the aggregation unit 152 may operate at an arbitrary timing to cope with the problem later instead of operating at a timing when the second packet has been input to the second input interface. “Later” means after setting of the connection path that will cause the different paths.

[0049] In this embodiment, the aggregation unit 152 causes the setting unit 151 to
25 perform connection path setting and connection path deletion. However, the

aggregation unit 152 may perform the connection path setting and the connection path deletion by itself. In both cases, the aggregation unit 152 substantially performs the connection path setting and the connection path deletion.

[0050] In this embodiment, a packet that requests establishment of the connection C1 corresponds to the first packet, and a packet that responds to the first packet corresponds to the second packet. Specifically, a TCP-SYN packet corresponds to the first packet, and a TCP-SYN-ACK packet corresponds to the second packet. The “TCP-SYN packet” is a packet in which a SYN flag in a TCP header indicates “1” and an ACK flag indicates “0”. The “TCP-SYN-ACK packet” is a packet in which both SYN flag and ACK flag in a TCP header indicate “1”.

[0051] Each time a packet including the information for identifying the connection C1 and satisfying a setting condition D1 is input to an interface IF-X different from interfaces IF-A and IF-B at both ends of a connection path already set by the setting unit 151 in association with the connection C1, the aggregation unit 152 may set a connection path between the interface IF-X and one of the interface IF-A and the interface IF-B, in association with the connection C1 and may delete the other connection path associated with the connection C1. With this arrangement, whenever a problem occurs that an outgoing path and a return path of the connection C1 become different paths, the aggregation unit 152 may cope with the problem. Specifically, the aggregation unit 152 may cope with the problem that the outgoing path and the return path become the different paths after establishment of the connection C1 that is a TCP connection as well as the problem that the transmission path of the TCP-SYN packet and the transmission path of the TCP-SYN-ACK packet for the connection C1 become the different paths.

[0052] The setting condition D1 may be a condition that is constantly true. That is,

it may also be so arranged that any packet including the information for identifying the connection C1 corresponds to the packet satisfying the setting condition D1.

[0053] It is desirable that, in terms of processing efficiency improvement and security strengthening, a packet input to the interface IF-X within a certain period of time after
5 detection of a change in a state of one of the different networks 120 and 130 connected to the network 110 correspond to the packet satisfying the setting condition D1.

[0054] A change in a routing table of the different network 120 or 130 corresponds to the change in the state. A link failure or a link recovery in the different network 120 or 130 also corresponds to the change in the state.

10 [0055] Assume a case where the aggregation unit 152 operates at the arbitrary timing, as a variation example of this embodiment. Then, each time a packet including the information for identifying the connection C1 is input to the interface IF-X different from the interfaces IF-A and IF-B at the both ends of the connection path already set by the setting unit 151 in association with the connection C1, the setting unit 151
15 additionally sets a connection path between the interface IF-X and another interface of the network 110 in association with the connection C1. The aggregation unit 152 determines whether a plurality of connection paths have been set by the setting unit 151 in association with the connection C1. When the aggregation unit 152 determines that the plurality of connection paths have been set by the setting unit 151 in association
20 with the connection C1, the aggregation unit 152 sets the connection path between the first input interface and the second input interface in association with the connection C1. The aggregation unit 152 deletes the other connection paths associated with the connection C1. As described above, the first input interface is the interface to which the first packet has been input. The second input interface is the interface to which the
25 second packet has been input. In this variation example, a packet that has caused a

latest one of the plurality of connection paths to be set by the setting unit 151 corresponds to the second packet. A last input one of the packets, each of which has been transmitted in a direction opposite to the second packet and has caused one of the plurality of connection paths to be set by the setting unit 151, corresponds to the first packet. Accordingly, the packet that requests establishment of the connection C1 does not necessarily correspond to the first packet, and the packet that responds to the first packet does not necessarily correspond to the second packet.

[0056] *** Description of Operations ***

Operations of the connection path management system 100 will be described, with reference to Fig. 6. The operations of the connection path management system 100 correspond to a connection path management method according to this embodiment. The operations of the connection path management system 100 correspond to a processing procedure of a connection path management program according to this embodiment.

[0057] In S11, a packet including the information for identifying one connection C1 is input to one interface of the network 110 from outside the network 110.

[0058] In S12, it is determined whether a connection path associated with the connection C1 to be identified by the information included in the packet input in S11 is already set by the setting unit 151. If the connection path is not set, the flow proceeds to S13. If the connection path is already set, the flow proceeds to S17. In this embodiment, an apparatus in the network 110 including the interface to which the packet has been input in S11 performs the process in S12. However, the setting unit 151 or the aggregation unit 152 may perform the process in S12.

[0059] In S13, the aggregation unit 152 determines whether the packet input in S11 is the packet requesting establishment of the connection C1. If the packet is not the one

requesting the establishment of the connection C1, the flow proceeds to S14. If the packet is the one requesting the establishment of the connection C1, the flow proceeds to S15. The setting unit 151 or the apparatus in the network 110 including the interface to which the packet has been input in S11 may perform the process in S13.

- 5 [0060] In S14, the aggregation unit 152 determines whether a connection path associated with the connection C1 to be identified by the information included in the packet input in S11 is already set by the setting unit 151. If the connection path is not set, the flow proceeds to S15. If the connection path is set, the flow proceeds to S16.
- [0061] In S15, the setting unit 151 sets a bidirectional connection path between the
10 interface to which the packet has been input in S11 and another interface of the network 110, as the connection path associated with the connection C1 to be identified by the information included in the packet input in S11. It is mainly determined from the transmission destination of the packet input in S11 which one of the interfaces in the network 110 is set to the different interface.
- 15 [0062] In S16, the aggregation unit 152 sets a bidirectional connection path between the interface to which the packet has been input from outside the network 110 in S11 and the interface to which the packet that has caused setting of the connection path already set by the setting unit 151 and associated with the connection C1 was input from outside the network 110, as a connection path associated with the connection C1 to be
20 identified by the information included in the packet input in S11. Then, the aggregation unit 152 deletes the other connection path associated with the connection C1. This aggregates the connection paths into one connection path. Due to occurrence of a change in a state of the different network 120 or 130, a packet to be transmitted in a same direction as the packet that has caused setting of the connection
25 path already set by the setting unit and associated with the connection C1 may be input

in S11. In that case, the aggregation unit 152 simply deletes the connection path already set, and causes the setting unit 151 to perform the process in S15. Accordingly, the setting unit 151 sets a bidirectional connection path between the interface to which the packet has been input in S11 and another interface of the network 110 as a
 5 connection path associated with the connection C1.

[0063] In S17, the packet input in S11 is forwarded inside the network 110 along the connection path associated with the connection C1 to be identified by the information included in that packet. That packet is output to outside the network 110 from the interface on one of both sides of that connection path opposite to the interface to which
 10 the packet has been input in S11. Then, the flow is finished.

[0064] An example of a network configuration according to this embodiment corresponding to the example in Fig. 2 will be described, with reference to Fig. 7.

[0065] The Router-1 and the Router-2 are routers that are the same as those in the example in Fig. 2. A Router-3 and a Router-4 are also routers. The SDN-SW1, the
 15 SDN-SW2, the SDN-SW3, and the SDN-SW4 are SDN switches that are the same as those in the example in Fig. 2. The Term-0 and the Term-1 are terminals. The Term-0 also indicates an interface of the Term-0 that is the terminal, which is connected to the network 121. The Term-1 also indicates an interface of the Term-1 that is the terminal, which is connected to a network 131.

20 [0066] It is herein also assumed that the network 110 constituted from the SDN switches is referred to as the SDN. The network 121 and the different network 120 constituted from the routers, the layer 2 switches, and so forth connected to the network 121 or included in the network 121 are referred to as the existing network. The network 131 similar to the network 121 and the different network 130 constituted from
 25 routers, layer 2 switches, and so forth connected to the network 131 or included in the

network 131 are also referred to the existing network. The network control apparatus 150 to perform the operations illustrated in Fig. 6 is connected to the SDN.

[0067] The network control apparatus 150 includes the setting unit 151 and the aggregation unit 152. The setting unit 151 functions as the SDN controller or a part of the SDN controller. The aggregation unit 152 implements a connection path reconfiguration and connection path aggregation function as one or a plurality of software processes. Generally, an API (Application Programming Interface) for function addition is disclosed for a vendor or an open-source SDN controller that provides a basic function. Function addition of the aggregation unit 152 may be conceived, using such an API.

[0068] Herein, a variation example of this embodiment will be illustrated in Fig. 8.

[0069] If the network control apparatus 150 is a vendor product and is provided as an appliance or the like, a management apparatus 160 including the aggregation unit 152 is separately installed in an environment in which the management apparatus 160 can access the API of the network control apparatus 150. In this variation example, the management apparatus 160 is a PC (Personal Computer). When the API is provided in a format assuming TCP/IP communication using an HTTP (HyperText Transfer Protocol) or the like, a physical positional relationship does not need to be cared of for software generation. Accordingly, functions that can be implemented are the same in both of the configurations in Figs. 7 and 8.

[0070] A description will be given about an example of operations of the connection path management system 100 when a communication request is transmitted from the Term-1 to the Term-0 in the configuration illustrated in Fig. 7, with reference to Figs. 9 to 13.

[0071] In the configuration illustrated in Fig. 7, interfaces of the Router-3 and the

Router-4 connected to the SDN may be regarded to correspond to the interfaces of the terminals such as the Term-2-1 and the Term-2-2 in the example in Fig. 2.

[0072] First, the following processes are performed, as illustrated in Fig. 9.

(Process 1-1) If a request packet that is the first packet of a connection transmitted by
 5 the Term-1 reaches the SDN, the request packet is input to an IF-4-2 that is an interface of the SDN-SW4. The SDN-SW4 and the IF-4-2 are selected due to assignment by the router in the existing network. The SDN is not involved in the selection. The request packet may be input to the SDN-SW3.

(Process 1-2) The SDN-SW4 determines whether connection identification information
 10 of the request packet is associated with a connection path already set in the SDN-SW4. If there is the connection path associated, the SDN-SW4 forwards the request packet according to the connection path. If there is not the connection path associated, the SDN-SW4 transmits a connection request 201 including the request packet to the network control apparatus 150.

(Process 1-3) If the setting unit 151 of the network control apparatus 150 receives the
 15 connection request 201, the setting unit 151 makes an inquiry to the aggregation unit 152. The aggregation unit 152 searches existing connection paths for a connection path associated with the connection identification information of the request packet included in the connection request 201 received, according to the above-mentioned
 20 conditions (a) and (b). Herein, there is no matching connection path. Therefore, the aggregation unit 152 notifies the setting unit 151 to perform subsequent processes as usual. It may be so arranged that depending on implementation, the search herein is performed by the basic function of the setting unit 151 through the API, and a result is received by the aggregation unit 152.

25 [0073] Then, the following processes are performed, as illustrated in Fig. 10.

- (Process 1-4) In the network control apparatus 150, the setting unit 151 performs a basic process as the SDN controller. First, the setting unit 151 computes the connection path. As a result, an output destination is determined to be an IF-1-2 that is an interface of the SDN-SW1. The request packet included in the connection request 201 is forwarded to
- 5 the SDN-SW1 that has become an exit of the SDN, as a forwarded packet 202. On the other hand, each SDN switch on the path is instructed to set a connection path by a connection path setting instruction 203. The SDN-SW1 outputs the request packet received from the network control apparatus 150 from the IF-1-2 that is the output destination.
- 10 (Process 1-5) The request packet is received by the Term-0 through a Path-1. The Path-1 is a path passing through the SDN-SW1 and the Router-2.
- [0074] A CP-1 is constructed as a result of the (Process 1-4), as illustrated in Fig. 11. This allows a communication from the Term-0 to the Term-1 through the SDN. A communication in an opposite direction using the same path also becomes possible.
- 15 The connection path is bidirectional and is constructed between the interfaces that serve as an entrance and an exit of the SDN. The CP-1 is set between the IF-4-2 and the IF-1-2.
- [0075] Subsequently, the following processes are performed, as illustrated in Fig. 12.
- (Process 2-1) A response packet that is the first packet of the connection responded by
- 20 the Term-0 is transmitted from the Term-0. The response packet is input to the SDN through a Path-2. The Path-2 is a path passing through the Router-2 and the SDN-SW2. The response packet is input to an IF-2-1 that is an interface of the SDN-SW2. The SDN-SW2 and the IF-2-1 are selected by assignment by the router in the exiting network. The SDN is not involved in the selection. The response packet
- 25 may be input to the SDN-SW1. In that case, the CP-1 may be used.

(Process 2-2) The SDN-SW2 determines whether connection identification information of the response packet is associated with a connection path already set in the SDN-SW2. If there is the connection path associated, the SDN-SW2 forwards the response packet according to that connection path. If there is no connection path associated, the

5 SDN-SW2 transmits a connection request 204 including the response packet to the network control apparatus 150. The input interface is also included as a condition for the determination herein. Thus, also if the response packet from the Term-0 has been input to an IF-1-1, the result is obtained that there is no connection path associated.

(Process 2-3) If the setting unit 151 receives the connection request 204 in the
10 network control apparatus 150, the setting unit 151 makes an inquiry to the aggregation unit 152. The aggregation unit 152 searches the existing connection paths for a connection path associated with the connection identification information of the response packet included in the connection request 204 received, according to the above-mentioned conditions (a) and (b). Herein, the CP-1 is found as the matching
15 connection path.

[0076] Then, the following processes are performed, as illustrated in Fig. 13.

(Process 2-4) In the network control apparatus 150, the aggregation unit 152 confirms that the input interface of the packet that has triggered construction of the CP-1 is the IF-4-2 and that the IF-4-2 can be the destination of the request packet included in the
20 connection request 204, with reference to information on the CP-1. Then, the aggregation unit 152 notifies the setting unit 151 to construct (i) a new connection path between the IF-2-1 that is the input interface of the response packet and the IF-4-2 that is the interface of the SDN-SW4, and (ii) and delete the matching connection path obtained by the search.

25 (Process 2-5) In the network control apparatus 150, the setting unit 151 computes the

connection path between the IF-2-1 and the IF-4-2. Each path in the SDN is determined according to various conditions set in advance. The response packet included in the connection request 204 is forwarded to the SDN-SW4 that has become the exit of the SDN, as a forwarded packet 205. On the other hand, connection path setting and deletion is instructed to each SDN switch on the path, by a connection path setting and deleting instruction 206. The SDN-SW4 outputs the response packet received from the network control apparatus 150 from the IF-4-2 that is the output destination

[0077] As the result of the (Process 2-5), a CP-2 is constructed, and the CP-1 is deleted. This allows a bidirectional communication between the Term-0 and the Term-1 passing through the SDN. The CP-2 is set between the IF-4-2 and the IF-2-1.

[0078] By the processes as described above, it becomes possible to reconfigure the connection path to an optical state when reciprocation using different routing is performed by the packets which have been transmitted first and respectively pass thorough the outgoing path and the return path.

[0079] In the operations in Figs. 9 to 13, the request packet corresponds to the first packet, while the response packet corresponds to the second packet. The IF-4-2 corresponds to the first input interface, the IF-1-2 corresponds to the first output interface, and the IF-2-1 corresponds to the second input interface.

[0080] The network configuration and the pattern for connection establishment described herein correspond to the (Type 3) and the (Type 4) of the (Type 1) to the (Type 4).

[0081] In the example illustrated with respect to the (Type 3) and the (Type 4), the Term-2 is connected to the SDN through two interfaces. The two interfaces are treated to be equal, which is similar to a situation where switches or routers for performing

assignment are logically present in the Term-2. When a configuration being the same as that of the example in Fig. 2 is applied to this embodiment, it can be regarded that the Term-2 is constituted from a logical network 301 and virtual PCs or blade PCs 302 connected to the network 301. That is, as seen from the SDN, this configuration is the same as connection of an existing network having a redundant configuration, and the configuration in Fig. 14 can be regarded to be equivalent to the configuration in Fig. 7. The number of interfaces for connection herein is not limited to two.

[0082] The Term-3 is connected to the SDN-SW4 by link aggregation. When packets are forwarded from the SDN to the Term-3 and when packets are forwarded from the Term-3 to the SDN, it is guaranteed that the packets of a same communication destination are forwarded using a same interface. This is a usual operation for maintaining the order of the packets, as described above. However, it is not guaranteed that a packet in a direction opposite to a direction of a packet received through a certain interface is forwarded using the interface. The link aggregation is a technology by which a plurality of interfaces are used to be virtually seen as a broadband line, and each of the apparatuses opposed to each other may perform assignment. Thus, it is not necessary that bidirectional packets of a connection pass through a same interface, and such implementation is not performed because a lot of processes and resources are needed. An Ethernet (registered trademark) originally uses a half-duplex line. If the half-duplex line is used, it may be necessary to become aware of a transmission state from an opposed apparatus in order to perform more efficient link aggregation. However, in an Ethernet using a twisted pair cable such as UTP (Unshielded Twisted Pair) or STP (Shielded Twisted Pair) that has generally become widely used in recent years, full-duplex communication is possible. A Gigabit Ethernet, in particular, assumes the full-duplex communication. Thus, interfaces that

will be passed through by reciprocating packets of one connection may be different, which may be the same as a situation where the Term-2 or the existing network is connected. Accordingly, different routing may occur even with the link aggregation, and the configuration of the Gigabit Ethernet may be regarded to be equivalent to the
 5 configuration in Fig. 7.

[0083] From the above description, the (Type 3) and the (Type 4) can be handled by the operations illustrated from Figs. 9 to 13.

[0084] With respect to the (Type 2), the operation is performed where the connection request is transmitted from the Term-1 and the response from the Term-0 reaches the
 10 single interface of the Term-1 using the different path. Accordingly, it only indicates that the IF-4-2, which is the input interface on the side of the Term-1 for the SDN is constantly fixed in the operations from Figs. 9 to 13. An operation of storing the interface to which a packet has first input is performed, regardless of whether the interface is fixed or not. Thus, the pattern for the (Type 2) may be also handled by the
 15 operations from Figs. 9 to 13.

[0085] As given above, the object of this embodiment may be achieved for various network configurations as well by the above-mentioned operations. An additional function will be given below.

[0086] As described above, TCP connections account for most of communications,
 20 and there is a clear connection establishment procedure in the TCP. Accordingly, by recognizing the establishment procedure and determining whether it is necessary or unnecessary to execute the procedure, a load on the network control apparatus 150 may be expected to be reduced.

[0087] In the case of the TCP-SYN packet, a new connection is basically to be made.
 25 Thus, no search is made even if information has been sent to the aggregation unit 152,

and an immediate response is made so that the process is performed as usual. In a case where the TCP connection is disconnected for some reason and is to be immediately reconnected or the like, it is seldom that same port numbers are applied and a reconnected connection can be recognized to be quite same as the disconnected connection. Thus, reuse of a connection path may make the process heavier.

[0088] In the case of the TCP-SYN-ACK packet, it is highly likely that different routing is to be used. Thus, the process by the aggregation unit 152 is applied. In addition to the case of the TCP-SYN-ACK packet, the process by the aggregation unit 152 may also be applied to a packet that is neither the TCP-SYN packet nor the TCP-SYN-ACK packet or a packet without a TCP SYN flag set. It is because, specifically, the following two situations and so forth may be considered, so that a new connection request may be generated even for the packet without the TCP SYN flag set.

(Situation 1) In an assignment rule of each router in the existing network connected to the SDN for the interfaces that are equal, assignment may be changed by a change in the order in an internal table or the like though no entry change is made. Such a case may include a change in a route table due to a failure in a portion not directly connected to the router, or the like.

(Situation 2) An assignment rule for the link aggregation changes if a link failure occurs.

[0089] If the connection path management system 100 handles all the TCP packets, however, the connection path management system 100 may readily become an attack target, or a failure may readily occur in the connection path management system 100. Then, in order to handle the packet without the TCP SYN flag set, the handling may be considered in which after some triggers have been recognized, a non-permission state in a normal case enters into a permission state. As such a trigger, notification of a change

in a dynamic routing table may be used. A link failure or a link recovery when connection with a terminal or an external network using link aggregation or the like is made may also be used. A period of time of the permission is also related to network design including a relationship with a TCP keep-alive timer and so on. The period of
 5 time of the permission is also related to setting of a holding period of time of a connection path in the SDN. Accordingly, determination of the period of time of the permission is entrusted to each designer.

[0090] Now, a relationship between this embodiment and a unidirectional connection path will be described.

10 [0091] Most of communications in an IP network in recent years use the TCP, and it is necessary to implement bidirectional communication for the TCP. Thus, for the control apparatus, it is more efficient to construct a bidirectional path at a point of time when a packet transmitted first has been input and a path for the packet is determined than to perform a process of receiving a connection path for each direction. The
 15 reason for this is as follows. In the unidirectional case, it is necessary to compute a path and compute input/output interfaces of each SDN switch on the path, for each request. In the case of simultaneously setting a connection path in an opposite direction, however, it is enough to conversely follow the input/output interfaces of each switch, and path computation is therefore substantially unneeded. Handling one
 20 bidirectional connection path is convenient for various managements as well in the following respects and so on: (i) bidirectional deletion can be performed by one process at a time of the deletion; (ii) bidirectional deletion is immediately possible even if the deletion by a TCP-RST is performed; (iii) each switch can identify bidirectional passages by TCP-FINs, so that a deletion process may readily be made to be more
 25 efficient; and (iv) with respect to a process of reconnecting a flow at a time of a failure

as well, the process may be performed using the flow as one connection path. With this arrangement, the (Type 1) among the (Type 1) to the (Type 4) can be handled with no additional process. Further, by performing the bidirectional setting, the effects of this embodiment are also greatly exhibited.

5 [0092] Handling a connection path for each direction eliminates unnecessary consumption of resources of each switch for setting the path even in the case of different paths. However, the different paths are to be handled as separate paths. Therefore, consumption of computational resources and resources in the entry table for a search by the SDN controller, and entry resources for hardware processing by the
 10 SDN switches is almost doubled at most. SDN performance is thereby greatly affected. When a large volume of new connections occur, when a failure has occurred to make a connection change, or the like, it may take time or the processing may not be able to be performed, depending on a case. Further, in order to solve inconvenience in management, implementation of a bidirectional correspondence and display of the
 15 bidirectional correspondence on a management screen, or the like is performed. Computation processing is thereby increased.

[0093] It is considered that use of the computational resources can be reduced by providing a bidirectional connection path entry table in addition to a unidirectional connection path entry table, by making association between each unidirectional entry
 20 and the bidirectional connection path entry table, and by using the bidirectional table for the search by the SDN controller or the like. Meanwhile, however, it remains that a processing load of receiving each connection path is twice as that in the case of each bidirectional path. Computation for making the association is added, and processing for the additional table is generated when a connection path change or deletion, a
 25 change due to a failure, or the like is performed. Thus, a lot of the computational

resources are consumed.

[0094] Thus, in order to achieve a state in a network configuration where use of various SDN resources is minimized and a burden is reduced by preventing operation management from becoming complex, the following state should be achieved. That is, connection paths should be aggregated into one connection path, and bidirectional path setting should be performed for each switch on the connection path. Then, when an existing network that has been made to be redundant, a terminal, and the SDN are connected, this embodiment is applied. Effective reduction of resource consumption thereby becomes possible.

[0095] Currently, it is a common practice to apply the SDN to a data center and issue a connection request from a terminal. Nevertheless, a redundant interface is often provided for a terminal at the data center or a server in recent years via a physical or logical switch or the like in order to accommodate a blade and virtualization. That case corresponds to the (Type 3) or the (Type 4), so that different routing occurs.

However, when the server is connected through one interface, an access from the terminal corresponds to the (Type 1). Thus, bidirectional setting mostly leads to solution. However, in recent years, achievement of a thin client of a terminal is under way. Thus, a connection between the thin client and a thin client server is mainly becoming a case corresponding to the (Type 1). A connection before achievement of the thin client remains to be made from the thin client server to a different server, almost as it is. Thus, even in the same data center, more accesses to an existing network tend to be made to pass through the existing network or to pass out to a WAN due to the large size of the data center or division of floors of the data center, so that a connection corresponding to the (Type 2) tends to be more increased. This embodiment is therefore effective.

[0096] When the SDN and the existing network are connected using four equal redundant interfaces to cause different routing, a probability that an outgoing path and a return path pass through a same one of the interfaces is one quarter. Accordingly, in 75% of connections in the situations of the (Type 2), the (Type 3), and the (Type 4), the different routing is used. If two interfaces are used, the probability is a half, and in 50% of the connections, different routing is used. This embodiment is effective for such an environment.

[0097] When connection paths are aggregated in this embodiment, it is desirable to preferentially leave a latest connection path. It is therefore desirable to store a time with a precision that is sufficient for indicating an order where each connection path has been established or a number indicating the order, as information on each connection path established.

[0098] During operation of the connection path management system 100, a manager may regularly reduce connection path information manually or automatically. When there are a plurality of connection paths having same connection identification information, interfaces at both ends of the SDN for the plurality of connection paths and interfaces that are respectively in a redundant relationship with those interfaces at the both ends are identified, a latest one of the connection paths established by packets input to the interfaces at one end and a latest one of the connection paths established by packets input to the interfaces at the other end are identified to finally select two connection paths. Then, a connection path is set, having at its both ends the interfaces of the selected connection paths to which the packets have been input, and the other connection paths are deleted.

[0099] *** Description of Effects ***

In this embodiment, when a situation occurs that a plurality of connection

paths are set for one connection between the terminals, the plurality of connection paths are aggregated into one connection path. Thus, according to this embodiment, a consumption amount of network resources may be reduced.

[0100] This embodiment prevents a situation where, by connection with the existing
 5 network having a redundant configuration through a plurality of interfaces, an outgoing path and a return path become different paths to cause unnecessary consumption of various SDN resources.

[0101] According to this embodiment, it becomes possible to aggregate two pieces of connection path information to be registered by the SDN controller into one connection
 10 path information. It thereby becomes possible to reduce consumption of table resources and a computation load in a process of search by the SDN controller in response to a new request, or the like.

[0102] When a TCP-RST packet passes, connection path deletion can be performed. However, only one unidirectional connection path can be deleted and a load of further
 15 searching and processing the other connection path is imposed. According to this embodiment, the load is eliminated.

[0103] According to this embodiment, the following situation may be prevented. That is, TCP-FIN packets themselves pass through different connection paths and ACKs from those connection paths pass through different connection paths. Thus, when the
 20 TCP-FIN packets are detected to perform connection path deletion, efficient deletion cannot be performed unless information detected at the respective different paths is aggregated.

[0104] According to this embodiment, also when a connection path is changed due to a failure or the like, it becomes possible to reduce a computation load because
 25 connection path entries are aggregated.

[0105] According to this embodiment, when the SDN is applied, design can be performed using a small number of resources of SDN-related apparatuses. Thus, cost reduction is achieved, and slow-start introduction of the SDN is readily achieved.

[0106] Hereinafter, an example of a hardware configuration of the network control apparatus 150 according to this embodiment will be described, with reference to Fig. 15.

[0107] The network control apparatus 150 is a computer. The network control apparatus 150 includes hardware such as a processor 901, an auxiliary storage device 902, a memory 903, a communication device 904, an input interface 905, and a display interface 906. The processor 901 is connected to the other hardware via a signal line 910, and controls these other hardware. The input interface 905 is connected to an input device 907. The display interface 906 is connected to a display 908.

[0108] The processor 901 is an IC (Integrated Circuit) to perform processing. The processor 901 is a CPU (Central Processing Unit), a DSP (Digital Signal Processor), or a GPU (Graphics Processing Unit), for example.

[0109] The auxiliary storage device 902 is a ROM (Read Only Memory), a flash memory, or an HDD (Hard Disk Drive), for example.

[0110] The memory 903 is a RAM (Random Access Memory), for example.

[0111] The communication device 904 includes a receiver 921 to receive data and a transmitter 922 to transmit data. The communication device 904 is a communication chip or a NIC (Network Interface Card), for example.

[0112] The input interface 905 is a port to which a cable 911 of the input device 907 is connected. The input interface 905 is a USB (Universal Serial Bus) terminal, for example.

[0113] The display interface 906 is a port to which a cable 912 of the display 908 is

connected. The display interface 906 is a USB terminal or an HDMI (registered trademark) (High Definition Multimedia Interface) terminal, for example.

[0114] The input device 907 is a mouse, a stylus pen, a keyboard, or a touch panel, for example.

5 [0115] The display 908 is an LCD (Liquid Crystal Display), for example.

[0116] Programs to implement functions of “units” such as the setting unit 151 and the aggregation unit 152 are stored in the auxiliary storage device 902. These programs are loaded into the memory 903, are read into the processor 901, and are executed by the processor 901. An OS (Operating System) is also stored in the

10 auxiliary storage device 902. At least some of the OS are loaded into the memory 903, and the processor 901 executes the program to implement the function of each “unit”.

[0117] Though Fig. 15 illustrates one processor 901, the network control apparatus 150 may include a plurality of processors 901. Then, the plurality of processors 901 may cooperate and execute the programs to implements the functions of the “units”.

15 [0118] Information, data, signal values, and variable values indicating results of processes by the “units” are stored in the auxiliary storage device 902, the memory 903, or a register or a cache memory in the processor 901.

[0119] Each “unit” may be provided by “circuitry”. Further, the “unit” may be replaced by a “circuit”, a “step”, a “procedure”, or a “process”. The “circuit” and the
20 “circuitry” are each a concept including not only the processor 901 but also a processing circuit of a different type such as a logic IC, GA (Gate Array), an ASIC (Application Specific Integrated Circuit), or an FPGA (Field-Programmable Gate Array).

[0120] A hardware configuration which is the same as that in the example in Fig. 15 may be applied to the management apparatus 160 as well.

25 [0121] The embodiment of the present invention has been described above; however,

this embodiment may be partially executed. Only one of the “units” described in the description of this embodiment may be adopted, or some of an arbitrary combination of the units may be adopted. Note that the present invention is not limited to this embodiment, and various modifications may be made as necessary.

5

Reference Signs List

- [0122] 100: connection path management system, 110: network, 120: different network, 121: network, 130: different network, 131: network, 150: network control apparatus, 151: setting unit, 152: aggregation unit, 160: management apparatus, 201: connection request, 202: forwarded packet, 203: connection path setting instruction, 204: connection request, 205: forwarded packet, 206: connection path setting and deleting instruction, 301: network, 302: virtual PC or blade PC, 901: processor, 902: auxiliary storage device, 903: memory, 904: communication device, 905: input interface, 906: display interface, 907: input device, 908: display, 910: signal line, 911: cable, 912: cable, 921: receiver, 922: transmitter

Claims

[Claim 1] A connection path management system comprising:

a setting unit to set, in association with a connection between terminals, a
bidirectional connection path for forwarding packets including information for
5 identifying the connection, between interfaces of a network to which the packets are
input or from which the packets are output; and

an aggregation unit to set a connection path between a first input interface and
a second input interface in association with one connection when a second packet
including information for identifying the one connection is input to the second input
10 interface after a connection path between the first input interface and a first output
interface has been set by the setting unit in association with the one connection, and
delete the other connection path associated with the one connection, the first input
interface being an interface to which a first packet including the information for
identifying the one connection has been input, the first output interface being an
15 interface from which the first packet has been output, the second input interface being
an interface different from the first output interface, the second packet having been
transmitted in a direction opposite to the first packet.

[Claim 2] The connection path management system according to claim 1,

wherein a packet that requests establishment of the one connection
20 corresponds to the first packet, and
wherein a packet that responds to the first packet corresponds to the second
packet.

[Claim 3] The connection path management system according to claim 1 or 2,

wherein, each time a packet including the information for identifying the one
25 connection is input to an interface different from the interfaces at both ends of the

connection path already set by the setting unit in association with the one connection, the aggregation unit sets a connection path between the different interface and one of the interfaces at the both ends, in association with the one connection, and deletes the other connection path associated with the one connection.

5 [Claim 4] The connection path management system according to claim 1 or 2,

wherein, each time a packet including the information for identifying the one connection and satisfying a setting condition is input to an interface different from the interfaces at both ends of the connection path already set by the setting unit in association with the one connection, the aggregation unit sets a connection path between
10 the different interface and one of the interfaces at the both ends, in association with the one connection, and deletes the other connection path associated with the one connection.

[Claim 5] The connection path management system according to claim 4,

wherein a packet that has been input to the different interface within a certain
15 period of time after detection of a change in a state of a different network connected to the network corresponds to the packet satisfying the setting condition.

[Claim 6] The connection path management system according to claim 5,

wherein a change in a routing table of the different network corresponds to the change in the state.

20 [Claim 7] The connection path management system according to claim 5,

wherein a link failure or a link recovery in the different network corresponds to the change in the state.

[Claim 8] The connection path management system according to claim 1,

wherein, each time a packet including the information for identifying the one
25 connection is input to an interface different from the interfaces at both ends of the

connection path already set by the setting unit in association with the one connection, the setting unit additionally sets a connection path between the different interface and another interface of the network in association with the one connection,

wherein the aggregation unit determines whether a plurality of connection
 5 paths have been set by the setting unit in association with the one connection, and when the aggregation unit determines that the plurality of connection paths have been set, the aggregation unit sets the connection path between the first input interface and the second input interface in association with the one connection and deletes the other connection path associated with the one connection,

10 wherein a packet that has caused a latest one of the plurality of connection paths to be set by the setting unit corresponds to the second packet, and

wherein a last input one of the packets, each of which has been transmitted in a direction opposite to the second packet and has caused one of the plurality of connection paths to be set by the setting unit, corresponds to the first packet.

15 [Claim 9] The connection path management system according to any one of claims 1 to 8, further comprising the network.

[Claim 10] A connection path management method comprising:

by a setting unit, setting, in association with a connection between terminals, a bidirectional connection path for forwarding packets including information for
 20 identifying the connection, between interfaces of a network to which the packets are input or from which the packets are output; and

by an aggregation unit, setting a connection path between a first input interface and a second input interface in association with one connection when a second packet including information for identifying the one connection is input to the second
 25 input interface after a connection path between the first input interface and a first output

interface has been set by the setting unit in association with the one connection, and deleting the other connection path associated with the one connection, the first input interface being an interface to which a first packet including the information for identifying the one connection has been input, the first output interface being an interface from which the first packet has been output, the second input interface being an interface different from the first output interface, the second packet having been transmitted in a direction opposite to the first packet.

[Claim 11] A connection management program to cause a computer to execute:

a process of setting, in association with a connection between terminals, a bidirectional connection path for forwarding packets including information for identifying the connection, between interfaces of a network to which the packets are input or from which the packets are output; and

a process of setting a connection path between a first input interface and a second input interface in association with one connection when a second packet including information for identifying the one connection is input to the second input interface after a connection path between the first input interface and a first output interface has been set in association with the one connection, and deleting the other connection path associated with the one connection, the first input interface being an interface to which a first packet including the information for identifying the one connection has been input, the first output interface being an interface from which the first packet has been output, the second input interface being an interface different from the first output interface, the second packet having been transmitted in a direction opposite to the first packet.