



(19) **United States**

(12) **Patent Application Publication**
GORKEMLI et al.

(10) **Pub. No.: US 2015/0312132 A1**

(43) **Pub. Date: Oct. 29, 2015**

(54) **METHOD TO CHECK HEALTH OF AUTOMATICALLY DISCOVERED CONTROLLERS IN SOFTWARE DEFINED NETWORKS (SDNS)**

Publication Classification

(51) **Int. Cl.**
H04L 12/26 (2006.01)
H04L 12/939 (2006.01)
(52) **U.S. Cl.**
CPC *H04L 43/50* (2013.01); *H04L 49/555* (2013.01)

(71) Applicant: **ARGELA YAZILIM VE BILISIM TEKNOLOJILERI SAN. VE TIC. A.S., ISTANBUL (TR)**

(72) Inventors: **BURAK GORKEMLI, ISTANBUL (TR); MURAT PARLAKISIK, ISTANBUL (TR); NAZLI GUNEY, ISTANBUL (TR); SEYHAN CIVANLAR, ISTANBUL (TR); OGUZ SUNAY, ISTANBUL (TR); UYGAR OZEN, ISTANBUL (TR); MELIH KARAMAN, ISTANBUL (TR); BULENT KAYTAZ, ISTANBUL (TR)**

(57) **ABSTRACT**

A method to check health of automatically discovered controllers comprises the steps of: (a) a first controller generating a first Controller-Advertisement message carrying the information required to identify and access the first controller when a new switch port becomes active; (b) the first controller sending the generated first Controller-Advertisement to a first switch having the newly activated port; (c) the first controller receiving at least a second Controller-Advertisement message from at least a second controller; (d) the first controller generating a third unsolicited Controller-Advertisement message with its access information to check health of the second controller; (e) the first controller receiving at least a fourth Controller-Advertisement message from the second controller indicating health information of the second controller, wherein the health information is stored in first controller's Controller Information Base (CIB).

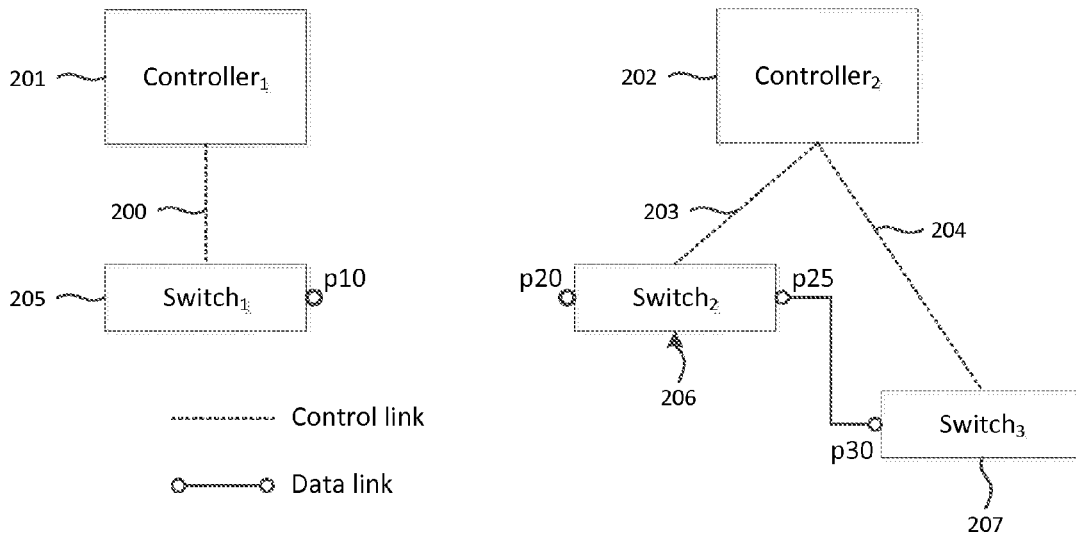
(73) Assignee: **ARGELA YAZILIM VE BILISIM TEKNOLOJILERI SAN.VE TIC. A.S., ISTANBUL (TR)**

(21) Appl. No.: **14/795,119**

(22) Filed: **Jul. 9, 2015**

Related U.S. Application Data

(62) Division of application No. 13/750,038, filed on Jan. 25, 2013, now Pat. No. 9,094,285.






FIG. 1
(PRIOR ART)




FIG. 2A




FIG. 2B

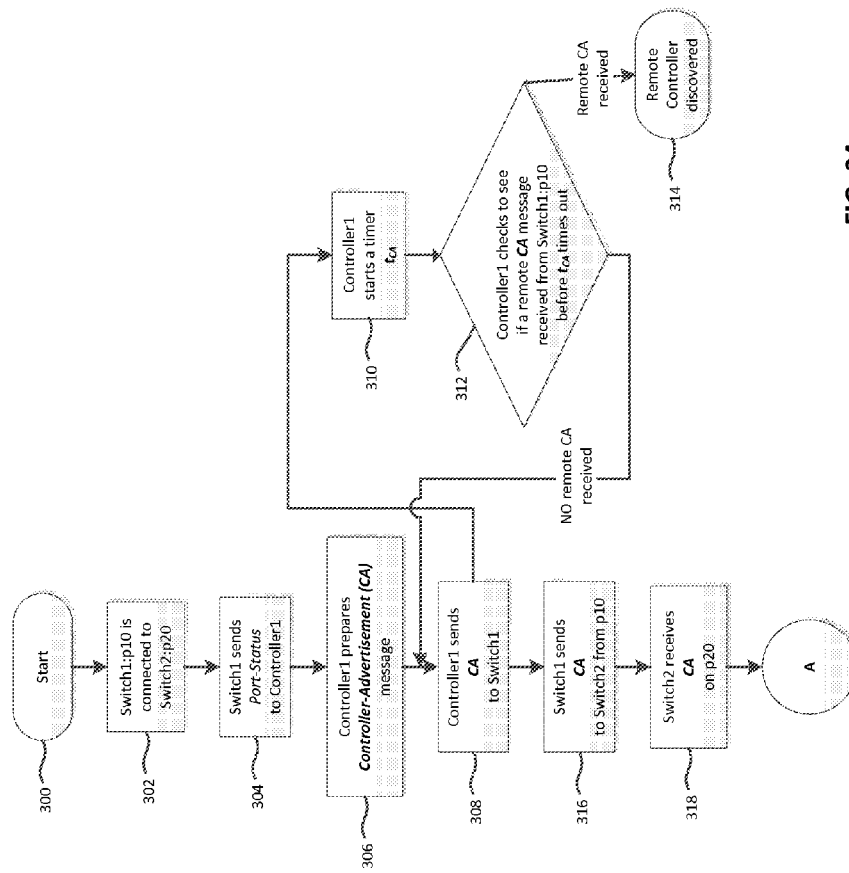


FIG. 3A

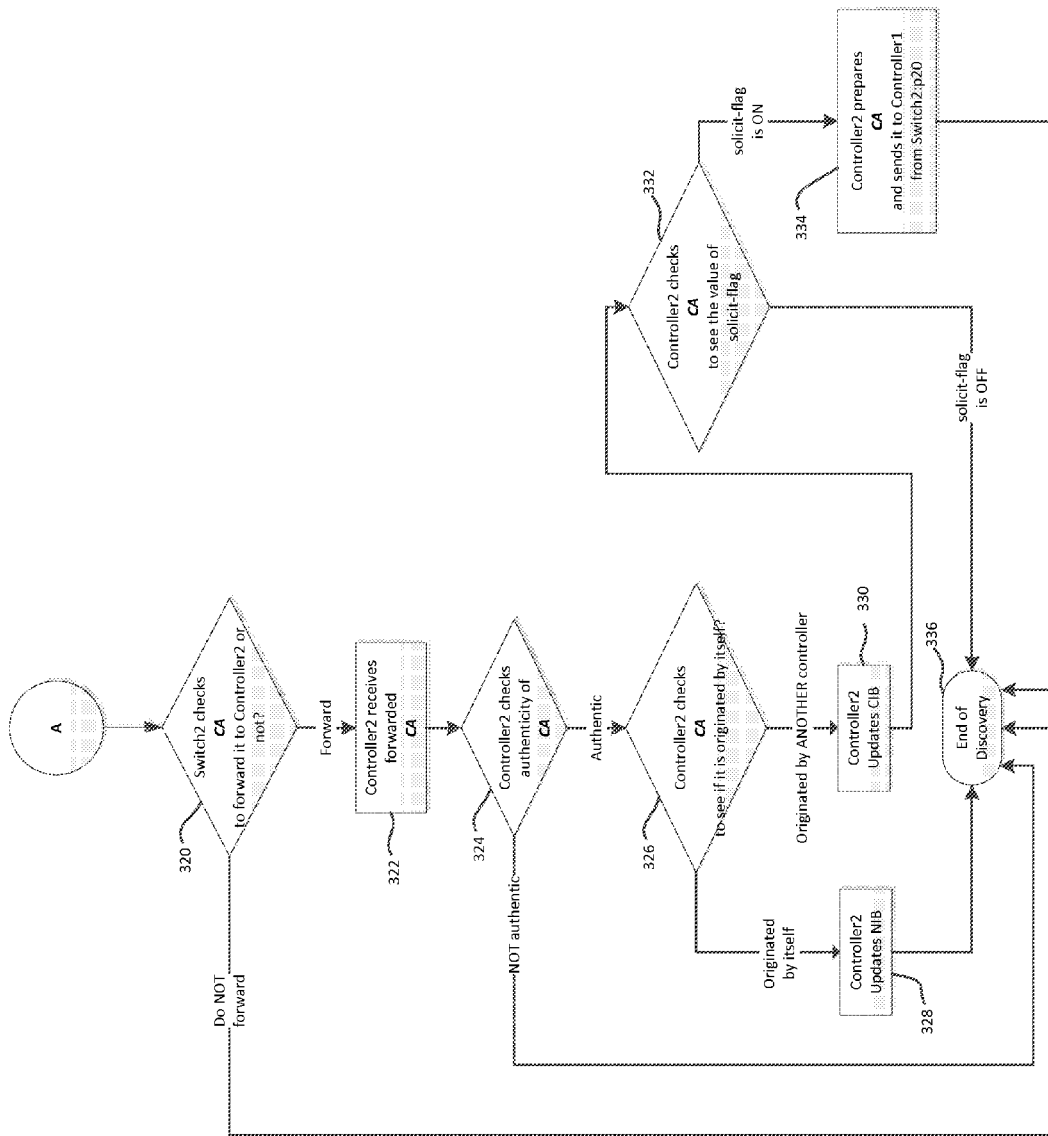


FIG. 3B




FIG. 4




FIG. 5




FIG. 6




FIG. 7

**METHOD TO CHECK HEALTH OF
AUTOMATICALLY DISCOVERED
CONTROLLERS IN SOFTWARE DEFINED
NETWORKS (SDNS)**

RELATED APPLICATIONS

[0001] This application is a divisional of U.S. application Ser. No. 13/750,038, filed Jan. 25, 2013.

BACKGROUND OF THE INVENTION

[0002] 1. Field of Invention

[0003] The disclosed invention generally relates to the controller component of a Software Defined Network (SDN). In particular, it relates to such controllers automatically discovering and authenticating each other as part of the control plane operations, without any need for manual configuration of other controllers as it is being done in prior art. The system of this invention is a new capability of the controller that receives Controller-Advertisement messages from other controllers in the same or different SDN, performing authentication and subsequently performing control plane information exchange. The Controller-Advertisement messages used for discovery are also used to periodically check the health of the pool of discovered controllers. Multiple controllers can be used for load sharing, improved reliability under failure of a controller by enabling seamless take over and recovery. Also, controllers can cooperate to enable end-to-end services to work across multiple SDNs.

[0004] 2. Discussion of Related Art

[0005] Software-Defined Networking (SDN) is a new paradigm where the control of computer and communication networks is accomplished via programmatic interfaces. There are a number of approaches to SDN. In the most popular approach, the data and control planes that typically reside in a switch are separated and the control plane is moved to a separate device, commonly referred to as the controller as described in the article "OpenFlow: Enabling Innovation in Campus Networks" by Nick McKeown et al., which appeared in ACM SIGCOMM Computer Communication Review, vol. 38, no. 2, April 2008 and also in the Pre-Grant Application by Casado et al. (2009/0138577). When triggered, the controller calculates the most appropriate route through the network between two nodes and programs all switches along this route accordingly. The controller can be a physical server, a virtual machine or an appliance. There can be multiple controllers in a computer network in this approach. It is possible to divide the network into regions and have different controllers control each region. It is also possible to have multiple controllers control a single switch for reliability and/or performance issues, where each controller may control different ports, or different flow types, or flows with different end nodes. In this approach, a well-defined protocol, OpenFlow, is used for communication between the network forwarding devices and the controller. The controller also has a northbound application programming interface so that different, custom control applications that use the OpenFlow data may be installed on the controller. This approach is advocated by the Open Networking Foundation (ONF) that aims to standardize the OpenFlow protocol. The latest version of the OpenFlow protocol may be accessed via the ONF web page. It is also possible to use a protocol other than OpenFlow to enable proper operation of this approach. FIG. 1 illustrates the OpenFlow-based SDN concept with a single controller **100** and a

plurality of switches **109, 110, 111, and 112**. In FIG. 1, dotted lines **101, 102, 103 and 104** represent control plane connections between the controller and the switches, while solid lines **105, 106, 107, and 108** represent the switch-to-switch data plane connections. All interactions in FIG. 1 between the controller and the switches use a clearly defined protocol, such as OpenFlow. This protocol allows for the control plane operations, wherein the controller programs the switches to specify the route of data paths (or flows, slices). The controller programs the switch by simply uploading one or more forwarding table(s) (or table entries).

[0006] In a second approach to SDN, virtual switches, as described in the Pre-Grant Application by Casado et al. (2010/0257263), are deployed as part of an overlay network on top of the existing network infrastructure using protocols such as VXLAN or NVGRE, which are defined in IETF Internet Drafts, draft-mahalingam-dutt-dcops-vxlan-00.txt and draft-sridharan-virtualization-nvgre-01.txt, respectively. In this approach, the data and control planes are still kept separate in the overlay, and the controller (or a group of controllers) is used to control the virtual switches.

[0007] In a third SDN approach, the data and control planes are not physically separated. Instead, every forwarding device in the network uses programmatic interfaces for network control, on which custom control applications may be installed. This vendor specific approach may be visualized as having an SDN controller in each forwarding device and using a vendor-specific protocol for communication between the controller and the forwarder.

[0008] The Pre-Grant Publication to Kato (2012/025049) describes a Load Distribution System, Load Distribution Method, and Program for SDN. In Kato, the load distribution amongst controllers is made possible using a proxy server that is situated between the switches and the controllers. The proxy server enables transparency of the different controllers to each of the switches. It also enables seamless load-balancing amongst the controllers unbeknownst to the switches in the network.

[0009] The Pre-Grant Publication to Yamato et al. (2011/0317701) describes one possible method to formulate routing decisions in a SDN architecture where multiple controllers are deployed. The routing decision, as calculated by the controller that controls the switch from which the route originates, is communicated to the other, relevant controllers in the network by this controller using a message that includes the forwarding path information.

[0010] In Kato as well as Yamata et al., as with the remainder of the prior art, when multiple controllers are used, a manual configuration step is needed to program each controller with the information of the other controller(s). The present invention eliminates this manual step by providing a system and method that allow each controller to automatically discover and authenticate other controllers in the network.

[0011] Embodiments of the present invention are an improvement over prior art systems and methods.

SUMMARY OF THE INVENTION

[0012] The present invention provides a method to check health of automatically discovered controllers comprising the steps of: (a) a first controller generating a first Controller-Advertisement message carrying the information required to identify and access the first controller when a new switch port becomes active; (b) the first controller sending the generated first Controller-Advertisement to a first switch having the

newly activated port; (c) the first controller receiving at least a second Controller-Advertisement message from at least a second controller; (d) the first controller generating a third unsolicited Controller-Advertisement message with its access information to check health of the second controller; (e) the first controller receiving at least a fourth Controller-Advertisement message from the second controller indicating health information of the second controller, wherein the health information is stored in first controller's Controller Information Base (CIB).

BRIEF DESCRIPTION OF THE DRAWINGS

[0013] FIG. 1 illustrates a prior art SDN.

[0014] FIG. 2A illustrates a SDN with two subnetworks managed by different controllers.

[0015] FIG. 2B illustrates the establishment of a data connection between a first port of a first switch and a second port of a second switch.

[0016] FIGS. 3A-B illustrate a flowchart outlining the steps describing the controller discovery process.

[0017] FIG. 4 depicts the sequence of messages involved in one example of a controller discovery process.

[0018] FIG. 5 depicts the sequence of messages involved in another example of a controller discovery process.

[0019] FIG. 6 illustrates the messaging sequence when a new connection is established between switches that are managed by the same controller.

[0020] FIG. 7 illustrates a SDN with two subnetworks managed by different controllers, where the second subnetwork is managed by two controllers for redundancy and/or performance issues.

DESCRIPTION OF THE PREFERRED EMBODIMENTS

[0021] While this invention is illustrated and described in a preferred embodiment, the invention may be produced in many different configurations. There is depicted in the drawings, and will herein be described in detail, a preferred embodiment of the invention, with the understanding that the present disclosure is to be considered as an exemplification of the principles of the invention and the associated functional specifications for its construction and is not intended to limit the invention to the embodiment illustrated. Those skilled in the art will envision many other possible variations within the scope of the present invention.

[0022] In the prior art, when multiple controllers are used in a Software Defined Network (SDN) (or several SDNs are interconnected, each with different controllers), a manual configuration step is needed to program each controller with the information of the other controller(s). The present invention eliminates that manual step by providing a system and method that allows each controller to automatically discover and authenticate other controllers in the network. An Autonomous System (AS) may be represented as a single SDN. Alternatively, it may be comprised of many interconnected SDNs. Thus, two SDNs that are interconnected may be in the same AS or in different Autonomous Systems. If the interconnected SDNs are in the same AS, their controllers may exchange all the control information they contain. Doing so, the controllers will be replicas of each other. However, policies may be defined to limit the control plane information each such controller shares with the other controllers. If the SDNs are in different Autonomous Systems, their controllers

will only exchange limited control information. For example, the internal network of the AS may be represented to the foreign controller in a more summarized or aggregated form. Also, there will be processes in place to ensure the other controller is trusted. The method of exchange of control information amongst controllers, once they discover each other, could be similar to exchange of routing information amongst routers in the classical Internet using various routing protocols, and therefore is kept outside the scope of this patent application.

[0023] The controller discovery process covered in this invention can be triggered by a change in the network (for example, when a new connection is established) and/or, periodically, at regular intervals. The former is more efficient in terms of messages injected into the network, but at the same time, it is more prone to message losses. Although some rules are defined in the following paragraphs to tolerate potential losses, it is advisable to perform controller discovery at regular intervals that are not short, besides being triggered by changes in the network.

[0024] The information gathered through the controller discovery process can be stored in the CIB, which is similar to the NIB used to store the information required to control the underlying switches. CIB can be considered as a table whose entries contain the information necessary to identify and access the discovered controllers, as shown in TABLE I below:

TABLE I

| An example CIB | | | | |
|----------------------|----------------------|-----------------|---------------|--------|
| Remote Controller-ID | Remote Controller-IP | Local Switch-ID | Local Port-No | Status |
| Controller1 | Controller1-IP | Switch2 | p20 | Live |
| Controller3 | Controller3-IP | Switch3 | p37 | Live |
| Controller4 | Controller4-IP | Switch3 | p39 | Live |

[0025] Below are the possible fields of a CIB entry that stores the access information for a discovered controller:

[0026] Controller-ID: ID of the discovered controller, which should be unique within the scope of the discovery process.

[0027] Controller-IP: IP address of the discovered controller.

[0028] Switch-ID: ID of the local switch from where the discovered controller can be accessed. This local switch is managed by the controller that maintains the CIB.

[0029] Port-no: Port number of the local switch from where the discovered controller can be accessed.

[0030] Status: Status of the discovered controller, which can take values such as 'Live', 'Dead', 'Not Accessible', or any other status values depending on the requirements.

[0031] The controllers listed in the CIB may change status under different conditions. For example, upon receiving a Controller-Advertisement message from a remote controller, the status of the remote controller will be set as 'Live', whereas not receiving a Controller-Advertisement message from a previously discovered remote controller in the health-check process will change the status to 'Dead'. Alternatively, if a data link that was previously used to interconnect the switches which enabled the control plane connectivity between the controllers is broken (i.e., the port of the switch in the CIB table above associated with a remote controller)

and access between the controllers thereby becomes unavailable, the status of the remote controller will be set as 'Not Accessible', upon receiving a message indicating the corresponding change in the network. While there are other values that can possibly be used for representing the status of a controller, it may also suffice to use only two values, which are 'Live' and 'Dead'.

[0032] The discovery process involves a single message, namely the Controller-Advertisement message, which is used by controllers to advertise their presence together with connection information. Since there is no explicit reply message of Controller-Advertisement, message losses can be tolerated by adhering to the following rules:

[0033] A controller that receives a Controller-Advertisement message with its solicit-flag set should respond with its own Controller-Advertisement message with its solicit-flag cleared, unless it does not want to participate in the controller discovery process.

[0034] A controller that has sent a Controller-Advertisement message over a connection should resend the advertisement message with its solicit-flag set, in case no remote Controller-Advertisement message is received from the same connection after a specific period.

[0035] The Controller-Advertisement message will have the following structure, in addition to the IP header, if encapsulated as an IP packet:

```

CONTROLLER_ADVERTISEMENT <
  IP-Header: <
    src-addr: Controller-IP;
    dst-addr: All-nodes multicast address; >
  controller-id: Controller-ID;
  solicit-flag: ON/OFF;
  algorithm-id: ID of the cryptographic algorithm
    that generates the signature; # Optional
  signature: cryptographically generated
    authentication code; # Optional
  switch-id: Switch-ID; # Optional
  port-no: port-ID; # Optional
  optional-fields: > # Optional

```

[0036] src-addr: Carries the access information for the controller that sends the advertisement message.

[0037] dst-addr: Multicast address to reach all nodes.

[0038] controller-id: ID of the controller, which should be unique within the multicast scope.

[0039] solicit-flag: This flag determines if the receiver of the advertisement message should respond with its own advertisement message or not. If set, the receiving controller should prepare its own Controller-Advertisement message and respond with it to the sender, unless it does not want to participate in the controller discovery process. If not set, no response should be sent. Typically, controller discovery is performed with the solicit-flag not set. On the other hand, Controller-Advertisement for a controller's health-check (towards an already discovered controller) should have its solicit-flag set to force a response from the discovered controller. Similarly, a Controller-Advertisement message, which is resent after a specific period because no remote Controller-Advertisement message is received from the newly activated switch-port, should have its solicit-flag set.

[0040] algorithm-id: ID of the cryptographic algorithm that generates the signature. This can be either a sym-

metric-key or an asymmetric-key (public-key) algorithm. This field is optional together with the following signature field.

[0041] signature: Cryptographically generated code that authenticates the advertisement message. This code can be constructed by the given cryptographic algorithm, using the src-addr, controller-id, switch-id and port-no fields. The controller that receives a Controller-Advertisement message with an existing signature value can validate the authenticity and the integrity of the received message. Since this field is optional, it is up to the controller to discard a received advertisement message with a null signature value.

[0042] switch-id: ID of the switch, from where the advertisement message is sent. The ID should be unique within the scope of the controller that originates the message. This field can be used for local topology discovery, that is, the discovery of connections between the switches managed by the same controller, together with port-no field. If a controller receives its own Controller-Advertisement message that is previously constructed, then it can utilize the switch-id and port-no fields to learn the existing connection between the sending and the receiving switches. This field is optional.

[0043] port-no: ID of the port, from where the advertisement message is sent. The ID should be unique within the scope of the sending switch. This field is optional.

[0044] optional-fields: Additional fields that can be utilized for controller-to-controller information exchange, which may be performed during the discovery process rather than being done after completing the discovery process. ID of the subnetwork that a controller manages and capability information of a controller are examples to such optional fields.

Controller Discovery Process

[0045] When a new data connection is established between two switches in an SDN network, a physical connection is established between the ports of these switches. Upon such a change in the port status (i.e., the port changing status from idle to active) of the said switches, they notify their respective controllers of the corresponding change. Then, the controllers can act accordingly to discover the change in the network.

[0046] FIG. 2A illustrates the scenario where there exist two subnetworks managed by different controllers, Controller₁ 201 and Controller₂ 202. These subnetworks may belong to the same or different SDNs. In this scenario, when two subnetworks interconnect, their controllers (i.e., Controller₁ 201 on the first subnetwork and Controller₂ 202 on the second subnetwork) will need to discover one another, and authenticate/exchange routing information to properly route packets from one subnetwork to the other.

[0047] In FIG. 2A, Switch₁ 205 is controlled by Controller₁ 201 on the first subnetwork, whereas Switch₂ 206 and Switch₃ 207 are controlled by Controller₂ 202 on the second subnetwork. The OpenFlow protocol is run on control links between the controllers and the switches, shown in dashed lines numbered with 200, 203, and 204. Initially, as shown in FIG. 2A, Switch₁ 205 and Switch₂ 206 are not connected, and Controller₁ 201 and Controller₂ 202 are not aware of each other. Thereafter, as depicted in FIG. 2B, a data connection, numbered with 208, is established between port p10 of Switch₁ 205 and port p20 of Switch₂ 206. When this connection is established, the controller discovery process is auto-

matically triggered. An exemplary flowchart, showing the discovery process initiated by Controller₁ 201 upon connecting ports p10 and p20, is depicted in FIG. 3A and FIG. 3B. It should be noted that although FIGS. 3A-B depict a discovery process initiated by Controller₁ 201, a reciprocal discovery process (not shown) is also initiated by Controller₂ 202, and is within the scope of the present invention.

[0048] The steps below describe the controller discovery process initiated by Controller₁ 201, as illustrated in FIGS. 3A-B. Reference is also made to FIG. 4 as it depicts the sequence of messages.

[0049] Step 300: Begin method.

[0050] Step 302: Switch₁ 205, port p10 is connected to Switch₂ 206, port p20.

[0051] Step 304: Upon port p10 being connected to port p20, Switch₁ 205 sends Port-Status message to Controller₁ 201. This OpenFlow message is intended to inform the controller that a previously idle port is now active.

[0052] For example, the corresponding OpenFlow (version 1.0.0) message from Switch₁ 205 to Controller₁ 201 contains:

```
OFPT_PORT_STATUS <
  reason: OFPPR_MODIFY;
  port-no: p10;
  state: OFPPS_STP_LISTEN >
```

[0053] Step 306: In response to receiving the Port-Status message sent from Switch₁ 205, Controller₁ 201 prepares a Controller-Advertisement message that carries its own ID and optionally the ID of Switch₁ 205 and port p10, according to an aspect of this invention. The solicit-flag should be cleared (i.e., its value should be OFF), since this is the initial message. This message will carry the IP address of Controller₁ 201 in its IP header. The content of the prepared Controller-Advertisement message is as follows:

```
CONTROLLER_ADVERTISEMENT <
  IP-Header: <
    src-addr: Controller1-IP;
    dst-addr: All-nodes; >
  controller-id: Controller1;
  solicit-flag: OFF;
  algorithm-id: Algorithm1;
  signature: authentication-code;
  switch-id: Switch1;
  port-no: p10; >
```

[0054] Step 308: Controller₁ 201 hands over the Controller-Advertisement message to Switch₁ 205, to be sent out from port p10. Using OpenFlow protocol, this message can be sent inside the Packet-Out message, as shown below:

```
OFPT_PACKET_OUT <
  type: OFPAT_OUTPUT;
  port: p10;
  data: CONTROLLER_ADVERTISEMENT <...> >
```

[0055] Step 310: Controller₁ 201 starts a timer t_{CA} , so as to resend the Controller-Advertisement message prepared in step 306 unless a remote Controller-Advertisement message has been received from port p10 of

Switch₁ 205 before t_{CA} times out. This timer is for tolerating potential message losses.

[0056] Step 312: Controller₁ 201 checks to see if a remote Controller-Advertisement message—in this case prepared by Controller₂ 202—has been received before t_{CA} times out. If no advertisement message is received in that period, then Controller₁ 201 resends the Controller-Advertisement message prepared in step 306, but this time with solicit-flag set (i.e., the solicit-flag is ON).

[0057] Step 314: A remote Controller-Advertisement message, which is prepared by Controller₂ 202, is received at port p10 of Switch₁ 205. When remote controller, namely Controller₂ 202, is discovered, there is no need to resend the Controller-Advertisement message prepared in Step 306.

[0058] Step 316: Switch₁ 205 sends the Controller-Advertisement message prepared by Controller₁ 201 to Switch₂ 206 from port p10.

[0059] Step 318: Switch₂ 206 receives the Controller-Advertisement message sent by Switch₁ 205 on port p20.

[0060] Step 320: Switch₂ 206 checks the Controller-Advertisement message to determine whether or not to forward the message to Controller₂ 202. The check involves looking up its forwarding table to see if there exists an entry that matches the message header. For the controller discovery to proceed, Switch₂ 206 should forward the message to Controller₂ 202, because either there is no flow entry matching the message or the matching flow entry dictates the switch to forward the message to the controller. Otherwise, the discovery process ends in step 336. This may occur because the Controller₂ 202 has a policy that restricts the controller discovery process or there may also be other reasons not covered here.

[0061] Step 322: In case Switch₂ 206 decides to forward the Controller-Advertisement message to Controller₂ 202, it does so by utilizing the Packet-In message of the OpenFlow protocol. So, Controller₂ 202 receives the advertisement message carried inside the OpenFlow (version 1.0.0) Packet-In message, as shown below:

```
OFPT_PACKET_IN <
  in_port: p20;
  data: CONTROLLER_ADVERTISEMENT <...> >
```

[0062] Step 324: Controller₂ 202 checks the authenticity of the received Controller-Advertisement message using the algorithm-id and signature fields of the message. If the message turns out to be valid, the process continues with Step 326. Otherwise, the advertisement message is discarded, and the discovery process ends in step 336. Also, if no authentication data is presented in the advertisement message, it is up to the controller—in this case, Controller₂ 202—to discard the message, or to continue with the discovery process.

[0063] Step 326: Controller₂ 202 checks the controller-id field of the Controller-Advertisement message, to determine if the message was originated by itself.

[0064] Step 328: If Controller₂ 202 determines that the received Controller-Advertisement message was its own message that is previously originated, as depicted in FIG. 6, then it can utilize the switch-id and port-no fields

of the advertisement message and the in port field of the Packet-In message to discover the connection established between the sending and receiving switches. After that, Controller₂ 202 can update its NIB accordingly, without making any changes in the CIB.

[0065] Step 330: If Controller₂ 202 determines that the received Controller-Advertisement message was originated by another controller—which is originated by Controller₁ 201 in our case, then Controller₂ 202 updates its CIB to reflect the newly discovered controller—Controller₁ 201. Assuming that the Controller₂ 202 CIB was empty at the beginning, then the status of the CIB will be as in TABLE II after discovering Controller₁ 201:

TABLE II

| Controller ₂ 202 CIB after discovering Controller ₁ 201 | | | | |
|---|----------------------|-----------------|---------------|--------|
| Remote Controller-ID | Remote Controller-IP | Local Switch-ID | Local Port-No | Status |
| Controller1 | Controller1-IP | Switch2 | p20 | Live |

[0066] Since a reciprocal discovery process is also initiated by Controller₂ 202, the Controller₁ 201 CIB will be as shown in TABLE III (shown below) after the discovery of Controller₂ 202, assuming that the CIB was empty at the beginning

TABLE III

| Controller ₁ 201 CIB after discovering Controller ₂ 202 | | | | |
|---|----------------------|-----------------|---------------|--------|
| Remote Controller-ID | Remote Controller-IP | Local Switch-ID | Local Port-No | Status |
| Controller2 | Controller2-IP | Switch1 | p10 | Live |

[0067] Step 332: Controller₂ 202 checks the solicit-flag of the received Controller-Advertisement message. If the value is OFF (i.e., if the solicit-flag is cleared), then the discovery process ends in step 336.

[0068] Step 334: If the solicit-flag of the received Controller-Advertisement message is ON (i.e., if the solicit-flag is set), then it means that Controller₁ 201 has requested a Controller-Advertisement message from Controller₂ 202. In that case, Controller₂ 202 prepares a Controller-Advertisement with its own information and solicit-flag cleared, and then sends the message to Controller₁ 201 from port p20 of Switch₂ 206.

[0069] Step 336: The discovery process ends for the established connection.

[0070] FIG. 4 depicts the Controller-Advertisement message traverses for the Controller₁ 201 initiated discovery process as follows: [Controller₁ 201]→[Switch₁ 205]→[Switch₂ 206]—[Controller₂ 202]. For the reciprocal process, where Controller₂ 202 initiates discovery, the Controller-Advertisement message traverses are as follows: [Controller₂ 202]→[Switch₂ 206]→[Switch₁ 205]→[Controller₁ 201]. This scenario is depicted in FIG. 5. In case Controller₁ 201 and Controller₂ 202 are a single controller (i.e., just Controller₁ 201), the Controller-Advertisement message traverses, as per FIG. 6, as follows: [Controller₁ 201]→[Switch₁ 205]→[Switch₂ 206]→[Controller₁ 201].

[0071] It is also possible for a switch to be managed by multiple controllers, for redundancy and/or performance

issues. FIG. 7 illustrates such a scenario, where a first sub-network is managed by Controller₁ 702 and a second sub-network is managed by Controller₂ 704 and Controller₃ 706. In this scenario, dotted lines 714, 716, 718, 720, and 722 represent various control links. It should be noted that while Switch₁ 708 has only one control link with Controller₁ 702, Switch₂ 710 has two control links 716 and 720, one for each of the controllers, Controller₂ 704 and Controller₃ 706 (i.e., Switch₂ 710 is managed by Controller₂ 704 and Controller₃ 706). Similarly, Switch₃ 712 has two control links 718 and 722, one for each of the controllers, Controller₂ 704 and Controller₃ 706 (i.e., Switch₃ 712 is also managed by Controller₂ 704 and Controller₃ 706). The controller discovery process covered in this invention covers such a scenario, such that the Controller-Advertisement message prepared by Controller₁ 702 is forwarded to both Controller₂ 704 and Controller₃ 706 by Switch₂ 710. Moreover, Controller₂ 704 and Controller₃ 706 both prepare Controller-Advertisement messages to be forwarded to Switch₁ 708 and eventually to Controller₁ 702, when notified by Switch₂ 710 of the change in the port status (i.e., the port changing status from idle to active).

[0072] Additionally, the present invention provides for an article of manufacture comprising computer readable program code contained within implementing one or more modules to automatically discover multiple controllers in software defined networks (SDNs). Furthermore, the present invention includes a computer program code-based product, which is a storage medium having program code stored therein which can be used to instruct a computer to perform any of the methods associated with the present invention. The computer storage medium includes any of, but is not limited to, the following: CD-ROM, DVD, magnetic tape, optical disc, hard drive, floppy disk, ferroelectric memory, flash memory, ferromagnetic memory, optical storage, charge coupled devices, magnetic or optical cards, smart cards, EEPROM, EPROM, RAM, ROM, DRAM, SRAM, SDRAM, or any other appropriate static or dynamic memory or data storage devices.

[0073] Implemented in computer program code based products are software modules for: a) receiving a message from said first switch identifying activation of a first port, which is due to establishment of a connection between a first port in said first switch and a second port in said second switch; b) generating a Controller-Advertisement message carrying information required to identify and access said first controller; c) forwarding said generated Controller-Advertisement message from said first controller to said first switch, said first port in said first switch further forwarding said Controller-Advertisement message to said second port in said second switch, said second switch further forwarding said Controller-Advertisement message to said second controller along with an identification of said second port over which said Controller-Advertisement message was received; wherein said second controller in said second subnetwork identifies said first controller in said first subnetwork based on information carried in said Controller-Advertisement, and said second controller generating a CIB entry containing information required to identify and access said first controller.

[0074] Implemented in computer program code based products are software modules for reciprocally (and independently of the Controller-Advertisement message originating from the first controller): (a) receiving a message from said second switch identifying activation of a second port, which

is due to establishment of a connection between a second port in said second switch and a first port in said first switch; (b) generating a Controller-Advertisement message carrying information required to identify and access said second controller; (c) forwarding said generated Controller-Advertisement message from said second controller to said second switch, said second port in said second switch further forwarding said Controller-Advertisement message to said first port in said first switch, said first switch further forwarding said Controller-Advertisement message to said first controller along with an identification of said first port over which said Controller-Advertisement message was received; wherein said first controller in said first subnetwork identifies said second controller in said second subnetwork based on information carried in said Controller-Advertisement, and said first controller generating a CIB entry containing information required to identify and access said second controller.

[0075] Also, implemented in computer program code based products are software modules for: a) receiving a message from said first switch identifying activation of a first port, which is due to establishment of a connection between a first port in said first switch and a second port in said second switch; b) generating a Controller-Advertisement message carrying information required to identify and access said first controller; c) forwarding said generated Controller-Advertisement message from said first controller to said first switch, said first port in said first switch further forwarding said Controller-Advertisement message to said second switch, said second switch further forwarding said Controller-Advertisement message to said second controller along with an identification of said second port over which said Controller-Advertisement message was received; d) if said first controller and said second controller are a single controller, then: (i) updating said single controller's NIB with a new entry showing new connection information between said first switch and said second switch, utilizing the information carried in the optional switch-id and port-no fields of said Controller-Advertisement message, and (ii) discarding any Controller-Advertisement messages that are looped to itself; e) if said first controller and said second controller are different controllers, then said second controller: (i) identifying said first controller using information carried in said Controller-Advertisement message originated from said first controller; and (ii) generating a new CIB entry containing information for said first controller.

CIB Maintenance

[0076] Once a controller has discovered a number of neighboring controllers, it should keep track of their statuses, for proper operation. This can be achieved by sending Controller-Advertisement messages, with solicit-flag set, to each of the previously discovered controllers, and waiting for them to send their Controller-Advertisement messages in return.

[0077] Returning to the discussion of FIG. 2 through FIG. 5 and considering the CIB given in TABLE IV (which is the same as the example in TABLE I), Controller₂ 202 will prepare a Controller-Advertisement message with solicit-flag set, carrying its own access information, and then send the advertisement message to each of the controllers in the CIB. As neighboring controllers receive the Controller-Advertisement message with solicit-flag set, they will respond to Controller₂ 202 with their Controller-Advertisement messages, with solicit-flag cleared. Upon receiving the advertisement messages from neighboring controllers, Controller₂ 202 will

maintain the statuses of these controllers. This process should be repeated at regular intervals for proper operation.

TABLE IV

| Controller, 202 CIB | | | | |
|----------------------|----------------------|-----------------|---------------|--------|
| Remote Controller-ID | Remote Controller-IP | Local Switch-ID | Local Port-No | Status |
| Controller1 | Controller1-IP | Switch2 | p20 | Live |
| Controller3 | Controller3-IP | Switch3 | p37 | Live |
| Controller4 | Controller4-IP | Switch3 | p39 | Live |

[0078] In addition to sending advertisement messages per CIB entry periodically, network change events should be processed to maintain the statuses of discovered controllers. For instance, upon receiving a Port-Status message showing that a port has become inactive, the accessibility states of the neighboring controllers that are accessible through that port should be changed accordingly.

Secure Controller-to-Controller Connection

[0079] After discovering each other, two controllers should build a secure connection between themselves, like Transport Layer Security (TLS), over which controller-to-controller information such as routing, service and capability information will be exchanged. Establishing the secure connection is well-defined and beyond the scope of this invention.

Controller-to-Controller Information Exchange

[0080] After completing the discovery process and establishing a secure end-to-end connection, the two controllers may need to exchange information such as:

[0081] 1—Routing information (NIB): The controllers may share a part of (aggregate) or all the routing information that they maintain, so as to enable packet routing between the networks that are in their control.

[0082] 2—Network capability information: The controllers may exchange the capabilities of their networks, such as the maximum bandwidth or minimum latency that can be offered, so as to support end-to-end quality of service (QoS).

[0083] 3—Controller capability information: Each controller may have different hardware and software-specific capabilities. Accordingly, each controller may be capable of handling only specific types of flows, and may have upper bounds on the traffic load that it can handle.

[0084] 4—Service offerings: The controller may provide certain services to the users of its subnetwork (such as location based services), which it may decide to make visible to the discovered controller so as to offer these services to the users of the other subnetwork, as well.

[0085] It is also possible to include some or all of this information in the Controller-Advertisement messages as part of the discovery process. This can be achieved by using the optional-fields defined in the Controller-Advertisement message.

CONCLUSION

[0086] A system and method has been shown in the above embodiments for the effective implementation of automatic discovery of multiple controllers in software defined networks (SDNs). While various preferred embodiments have been shown and described, it will be understood that there is

no intent to limit the invention by such disclosure, but rather, it is intended to cover all modifications falling within the spirit and scope of the invention, as defined in the appended claims. For example, the present invention should not be limited by number of sub-networks, number of controllers, number of switches, number of switches per controller, number of controllers per switch, software/program, computing environment, or specific computing hardware.

1. A method to check health of automatically discovered controllers comprising the steps of:

- a) a first controller generating a first Controller-Advertisement message carrying the information required to identify and access said first controller when a new switch port becomes active;
- b) said first controller sending said generated first Controller-Advertisement to a first switch having said newly activated port;
- c) said first controller receiving at least a second Controller-Advertisement message from at least a second controller;

d) said first controller generating a third unsolicited Controller-Advertisement message with its access information to check health of said second controller;

e) said first controller receiving at least a fourth Controller-Advertisement message from said second controller indicating health information of said second controller, wherein said health information is stored in first controller's Controller Information Base (CIB).

2. The method of claim 1, further comprising authenticating received Controller-Advertisement messages.

3. The method of claim 1, further comprising, prior to sending, setting a solicit flag in a Controller-Advertisement message.

4. The method of claim 1, further comprising generating a Controller-Advertisement message with solicit flag cleared upon receiving a Controller-Advertisement message with solicit flag set.

* * * * *