



- (51) International Patent Classification:
H04L 12/755 (2013.01)
- (21) International Application Number:
PCT/CN2013/077369
- (22) International Filing Date:
18 June 2013 (18.06.2013)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
201210230484.X 3 July 2012 (03.07.2012) CN
- (71) Applicant: **HANGZHOU H3C TECHNOLOGIES CO., LTD.** [CN/CN]; Huawei Hangzhou Manufacturing Base, 310 Liuhe Road, Zhijiang Hi-Tech Park, Hangzhou Hi-Tech Industry Development Zone, Hangzhou, Zhejiang 310053 (CN).
- (72) Inventor: **LIN, Changwang**; Room 730, Oriental Electronic BLD., No.2, Chuangye Road, Shangdi Information Industry Base, Haidian District, Beijing 100085 (CN).
- (74) Agent: **CHINA PATENT AGENT (HK) LTD.**; 22/F., Great Eagle Center, 23 Harbour Road, Wanchai, Hong Kong (CN).
- (81) Designated States (*unless otherwise indicated, for every kind of national protection available*): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.
- (84) Designated States (*unless otherwise indicated, for every kind of regional protection available*): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM,

[Continued on next page]

(54) Title: REDUCING FLOODING OF LINK STATE INFORMATION




Fig. 2

(57) Abstract: The present disclosure describes reducing flooding of link state information of a link state protocol in a network with multiple routers. A filtering strategy is pre-configured on a router interface for filtering link state information generated by a first router. The filtering strategy is to enable link state database (LSDB) isolation between a second router associated with the router interface and the first router. Link state information generated by the first router is received or sent according to the pre-configured filtering strategy via the router interface.

WO 2014/005479 A1



TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, **Published:**
KM, ML, MR, NE, SN, TD, TG).

— *with international search report (Art. 21(3))*

Reducing Flooding of Link State Information

Background

[0001] Link state routing protocols are used for communication in packet-switched networks. Routers exchange link state information with each other using a flooding process, where each router creates link state information and floods it to their neighbours. The link state information is used to maintain a link state database (LSDB) that reflects the topology of the autonomous system.

Brief Description of Drawings

[0002] The present disclosure will refer to examples in the following drawings, in which:

[0003] Fig. 1 is a block diagram of a first example network environment for reducing flooding of link state information;

[0004] Fig. 2 is a flowchart of an example method for reducing flooding of link state information;

[0005] Fig. 3A is a block diagram of part of the first example network environment in Fig. 1, Fig. 3B shows example global filtering strategies and Fig. 3C shows example interface filtering strategies;

[0006] Fig. 4 is a flowchart of an example process for receiving link state information according to a pre-configured filtering strategy;

[0007] Fig. 5 is a flowchart of an example process for transmitting link state information according to a pre-configured filtering strategy;

[0008] Fig. 6 is a block diagram of a second example network environment for reducing flooding of link state information;

[0009] Fig. 7 is a block diagram of a third example network environment for reducing flooding of link state information; and

[0010] Fig. 8 is a block diagram of an example structure of a routing device for reducing flooding of link state information.

Detailed Description

[0011] Link state information may be exchanged by routers for the purpose of LSDB synchronization. For instance, each router may receive link state information generated by its neighbours, so that each router may synchronize its LSDB with that of its neighbours. Since a flooding process is used to exchange the link state information, this creates a lot of traffic in the network and a burden on the routers to process the information.

[0012] The present disclosure describes a method for reducing flooding of link state information of a link state protocol in a network with multiple routers. A filtering strategy is pre-configured on a router interface for filtering link state information generated by a first router. The filtering strategy is to enable LSDB isolation between a second router associated with the router interface and the first router. Link state information generated by the first router is received or sent according to the pre-configured filtering strategy via the router interface.

[0013] Throughout the present disclosure, the term “LSDB isolation” broadly describes the situation where the LSDB of a router is not synchronized with the LSDB of at least one neighbouring router. By pre-configuring the filtering strategy, unnecessary link state information is filtered so that the information is not received or sent via a router interface, thus enabling LSDB isolation. This in turn reduces flooding of link state information in the network. From the routers' perspective, this reduces the processing burden associated with route calculation and LSDB update, as well as the size of their LSDB.

[0014] Examples will be described with reference to the accompanying drawings.

[0015] Fig. 1 is a schematic diagram of an example network environment 100 for reducing flooding of link state information of a link state protocol. The network environment 100 in Fig. 1 includes three layers: core layer 110, convergence layer 120 and access layer 130 on which routers 112, 122, 132 are deployed. For example:

[0016] The core layer 110 serves as the backbone for the network 100 and is deployed with Level 2 (L2) routers 112 (e.g. Core1 to Core4).

[0017] The convergence layer 120 (also known as a distribution layer) is deployed with Level 1-2 (L12) routers 122 (e.g. MTR1 to MTR6) to facilitate traffic routing within the access layer 130. L12 routers have both L1 and L2 routing functions.

[0018] The access layer 130 connects to the convergence layer 120, and is deployed with Level 1 (L1) routers 132 (e.g. RT1 to RTn). The access layer 130 is generally where end users or end systems (ES) connect to the network 100.

[0019] Throughout the present disclosure, the terms 'L2', 'L12' and 'L1' broadly refer to different levels of routers deployed in the core 110, convergence 120 and access 130 layers. Use of these terms should not be confused with the layers of an OSI (Open Systems Interconnection) protocol stack, such as Layer 1 (physical layer), Layer 2 (link layer) etc. For example, the core 110, convergence 120 and access 130 layers may be logical layers in the network environment 100 with the functions described above. Further, the term "router" is used broadly in the present disclosure to refer to a network device with routing functionality, such as a network router, or a switch with routing functionality etc.

[0020] A group of routers using the same protocol to exchange routing information is generally referred to as a "routing domain", which can be further divided into different "areas" or "sections" 134. For example, the L1 routers 132 deployed in the access layer 130 belong to different areas 134. In Fig. 1, a first group of routers (i.e. RT1, RT2 and RT3) belong to a first area, while a second group (i.e. RT4, RT5 and RT5) belong to a different L1 area.

[0021] Each router 112, 122, 132 runs a link state protocol and maintains a LSDB that stores the router's local state, such as the router's interfaces and how to reach its neighbours etc. L2 routers may establish neighbour relationship with L2 routers and L12 routers in the same or different areas. L12 routers may establish L1 neighbour relationship with L1 and L12 routers in the same area, or L2 neighbour relationships with L2 and L12 routers in different areas. L1 routers 132 only establish neighbour relationship with L1 and L12 routers in the same area.

[0022] If no filtering strategy is pre-configured, neighbouring routers (e.g. RT1, RT2 and RT3) exchange link state information to synchronize their LSDBs. Each router receives link state information from its neighbours, performs route calculation and updates its LSDB. For example, using a flooding process, RT1 sends its link state information to RT2 and RT3, RT2 sends to RT1 and RT3, and RT3 sends to RT1 and RT2. L1 routers RT1, RT2, RT3 exchange the link state information via a L2 router, e.g. MTR1. The flooding process is generally repeated from time to time.

[0023] As the number of neighbouring routers increases, the traffic caused by the flooding process also increases. In this case, since there is more link state information to process, this also increases the burden on router CPUs to maintain a larger LSDB.

[0024] In the example in Fig. 1, LSDB isolation is enabled between routers to reduce flooding of link state information in the network 100. For example, when LSDB isolation is enabled between routers RT1 and RT2, they do not synchronize their LSDBs, i.e. RT1 does not have to process the link state information generated by RT2 to update its LSDB, and vice versa.

[0025] Fig. 2 shows an example method 200 for reducing flooding of link state information of a link state protocol. LSDB isolation is enabled by pre-configuring a filtering strategy 150. The example method 200 includes:

[0026] At block 210, a filtering strategy or route filtering strategy (see 150 in Fig. 1) is pre-configured on a router interface (e.g. an interface on MTR1) for filtering link state information generated by a first router (e.g. RT1). The filtering strategy 150 enables LSDB isolation between (i) the first router (e.g. RT1), and (ii) a second router (e.g. RT2) associated with the router interface.

[0027] At block 220, link state information generated by the first router is received or sent via the router interface (e.g. interface on MTR1) according to the pre-configured filtering strategy 150.

[0028] Throughout the present disclosure, the term “link state information” refers to any link state information generated by a link state protocol to facilitate routing in the network.

Examples of the “link state protocol” include the Intermediate System-to-Intermediate System (IS-IS), Open Shortest Path First (OSPF), etc. The link state information may be in any suitable format, such as Link State Packets (LSPs) when IS-IS is used, and Link State Advertisements (LSAs) when Open Shortest Path First (OSPF) is used. LSPs may also be Link State Protocol data units.

[0029] The term “link state information” may also refer to other types of link state information such as complete sequence number PDUs (CSNPs) and partial sequence number PDUs (PSNPs) etc. The CSNPs and PSNPs are exchanged among IS-IS routers in order to maintain a correct LSDB. CSNPs contain a list of link state information from the current database to inform other routers that may have outdated or missing information. This ensures all routers have the same information and their LSDBs are synchronized. In general, when a router receives CSNP, it compares header information of the CSNP with its local link state information. If the local link state information is incomplete, the router will send a PSNP to request any missing link state information. PSNPs contain partial link state information and are used to request link state information or acknowledge receipt of link state information.

[0030] Referring also to the examples in Figs. 3A, 3B and 3C, the LSDB isolation between the routers may be complete or partial. For simplicity, Fig. 3A shows part of the network environment 100 in Fig. 1 that includes MTR1 in the convergence layer 120, and RT1, RT2 and RT3 in the access layer 130.

[0031] Complete LSDB isolation broadly refers to the case where a router (e.g. RT2) does not process link state information generated by all of its neighbours (e.g. RT1 and RT3), except its upstream (UP) neighbour (e.g. MTR1). The term “upstream” in the present disclosure refers to the direction from the access layer 130 to the convergence layer 120 to the core layer 110. For example, MTR1 in the convergence layer 120 is an upstream neighbour of RT1, RT2 and RT3 in the access layer 130, and RT1, RT2 and RT3 are downstream neighbours of MTR1. Similarly, ‘Core 3’ in the core layer 110 is an upstream neighbour of MTR1, and MTR1 is a downstream neighbour of ‘Core 3’ etc. See Fig. 1 again.

[0032] Referring to 310 and 320 in Fig. 3B, complete LSDB isolation may be achieved by pre-configuring a “global” filtering strategy. In this example, 310 or 320, or both, may be pre-configured to enable complete LSDB isolation.

[0033] A first global filtering strategy (see 310) may be pre-configured on L12 router MTR1's interface P1 (to which RT2 is connected) to filter link state information generated by RT2's L1 neighbours, i.e. RT1 and RT3. Once a filtering strategy has been pre-configured, MTR1 does not send link state information generated by RT1 and RT3 to RT2 via interface P1. Since MTR1 is an UP neighbour of RT2, MTR1 only sends link state information generated by itself (i.e. MTR1) to RT2 via interface P1.

[0034] Alternatively or additionally, a second global filtering strategy (see 320) may be pre-configured on RT2's router interface R2 to filter link state information generated by RT1 and RT3. Once pre-configured, RT2 discards or does not receive link state information generated by RT1 and RT3 via interface R2. Since MTR1 is an UP neighbour of RT2 (i.e. MTR1 is a neighbour of interface R2), RT2 only receives link state information generated by MTR1 via interface R2.

[0035] Partial LSDB isolation broadly refers to the case where a router (e.g. RT2) does not process link state information generated by at least one neighbour (e.g. RT1), but continues to do so for link state information generated by at least one other neighbour (e.g. RT3). Referring to 330 and 340 in FIG. 3C, partial LSDB isolation may be achieved by pre-configuring an "interface" filtering strategy.

[0036] A first interface filtering strategy (see 330) may be pre-configured on L12 router MTR1's interface P1 (to which RT2 is connected) to filter link state information generated by RT1 but not RT3. Once pre-configured, MTR1 does not send link state information generated by RT1 to RT2 via interface P1 to enable LSDB isolation between RT1 and RT2. However, MTR1 will continue sending link state information by RT3 to RT2 via P1.

[0037] Alternatively or additionally, a second interface filtering strategy (see 340) may be pre-configured on RT2's router interface R2 to filter link state information generated by RT1. Once pre-configured, RT2 discards link state information generated by RT1 via interface R2 to enable LSDB isolation between RT1 and RT2. However, RT2 will continue receiving link state information generated by RT3.

[0038] Based on the above, the example method may be implemented on a router in the core layer 110, convergence layer 120 (e.g. MTR1) or access layer 130 (e.g. RT2) to reduce flooding of link state information. Also, the term “associated with” in block 210 broadly describes any suitable association, such as the case where the second router is connected to the router interface (e.g. RT2 is connected to P1 of MTR1), or the router interface is an interface of the second router (e.g. R2 is an interface of RT2).

[0039] The link state information may include the system ID of the source router, i.e. the router that generated the link state information. The system ID may be a single value or a range of values, and configured as required. As such, the strategy of whether to receive or send is made by checking the system ID in the link state information.

[0040] Fig. 4 shows an example implementation of block 220 in Fig. 2 when sending link state information via a router interface according to a pre-configured filtering strategy. For ease of understanding, the terms “sender router” (e.g. MTR1) and “recipient router” (e.g. RT2) below respectively refer to the router sending and router (potentially) receiving the link state information generated by a “source router”. The “source router” may be the “sender router” itself (e.g. MTR1), or a different router (e.g. RT1).

[0041] At 410, when a sender router (e.g. MTR1) is sending link state information to a recipient router (e.g. RT2) via a router interface (e.g. P1), the sender router checks whether the router interface is pre-configured with a filtering strategy.

[0042] At 415, if no filtering strategy is pre-configured, the sender router (e.g. MTR1) sends the link state information to the recipient router (e.g. RT2) via the router interface (e.g. P1) and performs the necessary processing.

[0043] At 420, the sender router (e.g. MTR1) determines whether the filtering strategy is a global or interface filtering strategy. 430 is performed if a global filtering strategy is pre-configured, but otherwise (interface filtering strategy) 460 is performed.

[0044] At 430, if the router interface (e.g. P1) is pre-configured with a global routing filtering strategy (e.g. 310 in Fig. 3B), the sender router (e.g. MTR1) checks

whether the source router of the link state information is an UP neighbour of the recipient router (e.g. MTR1 is an UP neighbour of RT2).

[0045] At 440, if yes (e.g. source router is MTR1), the link state information is sent by the sender router (e.g. MTR1) to the recipient router (e.g. RT2) via the router interface (e.g. P1).

[0046] Otherwise, at 450, the link state information (e.g. source router is RT1 or RT3) will be discarded. This enables LSDB isolation between RT1 and RT2, as well as between RT3 and RT2.

[0047] At 460, if the router interface (e.g. P1 of MTR1) is pre-configured with an interface filtering strategy (e.g. 330 in Fig. 3C), the sender router (e.g. MTR1) checks whether the system ID of the source router (e.g. RT1) generating the link state information satisfies the interface filtering strategy.

[0048] At 440, if yes (e.g. source router is RT3), the link state information will be sent via the router interface (e.g. P1) to the recipient router (e.g. RT2). Otherwise at 450 (e.g. source router is RT1), the link state information will not be sent. This enables LSDB isolation between RT1 and RT2, but not RT3 and RT2.

[0049] Fig. 5 shows an example implementation of block 220 in Fig. 2 when receiving link state information via a router interface according to a pre-configured filtering strategy.

[0050] At 510, when a router (e.g. RT2) is receiving link state information via a router interface (e.g. R2 of RT2), the router (e.g. RT2) checks whether the router interface is pre-configured with a filtering strategy.

[0051] At 515, if no filtering strategy pre-configured, the router (e.g. RT2) receives the link state information via the router interface (e.g. R2) and performs the necessary processing of the link state information.

[0052] At 520, the router (e.g. RT2) determines whether the filtering strategy is a global or interface filtering strategy. 530 is performed if a global filtering strategy is pre-configured, but otherwise (interface filtering strategy) 560 is performed.

[0053] At 530, if the router interface (e.g. R2) is pre-configured with a global routing filtering strategy (e.g. 320 in Fig. 3B), the router (e.g. RT2) checks whether the source router of the link state information is an UP neighbour (e.g. MTR1) of the interface.

[0054] At 540, if yes, the link state information generated by its UP neighbour (e.g. MTR1) is received via the router interface (e.g. R2).

[0055] Otherwise, at 550, the link state information is discarded. In other words, link state information generated by other routers (e.g. RT1 and RT3) will be discarded. This enables LSDB isolation between RT1 and RT2, as well as between RT3 and RT2.

[0056] At 560, if the router interface (e.g. R2 of RT2) is pre-configured with an interface filtering strategy (e.g. 340 in Fig. 3C), the router (e.g. RT2) checks whether the system ID of the source router (e.g. RT1) generating the link state information satisfies the interface filtering strategy.

[0057] At 540, if yes (e.g. source router is RT3), the link state information is received via the router interface (e.g. R2). Otherwise, at 550 (i.e. source router is RT1), the link state information is discarded. This enables LSDB isolation between RT1 and RT2, but not RT3 and RT2.

[0058] Based on the above examples, the interface filtering strategy (e.g. 330, 340 in Fig. 3C) is more flexible than the global filtering strategy (e.g. 310, 320 in Fig. 3B) because the interface filtering strategy is based on system ID of the source router generating the link state information. On the other hand, once the global filtering strategy is pre-configured, complete LSDB isolation is implemented. In the example in Fig. 3A, the interface filtering strategy may be used if RT2 only requires LSDB isolation from RT3 but not RT1. Otherwise, the

global filtering strategy may be preferred if RT2 requires LSDB isolation from both RT1 and RT3, i.e. complete LSDB isolation.

[0059] It should be noted that the system ID(s) configured by the interface filtering strategy may be a blacklist or a whitelist. For example, the interface filtering strategy 330 in Fig. 3C includes the system ID of RT1 in a blacklist. Link state information generated by routers in the blacklist will not be received or sent. Alternatively, the system ID of RT1 may be excluded from a whitelist. Link state information generated by routers in the whitelist will be received or sent.

[0060] Example Network 600

[0061] Using pre-configuration of filtering strategy, a network 600 having a flat hub-and-spoke topology may be implemented. An example is shown in Fig. 6 in which an L12 router 612 (i.e. MTR1) in the convergence layer 610 is connected to multiple neighbouring L1 routers 622 (i.e. RT1, RT2 to RTn) in the access layer 520 via different interfaces (i.e. P1, P2 to Pn). The L12 router 612 (e.g. MTR1) may be connected to a L2 router in the core layer (not shown for simplicity). Compared with the network topology in Fig. 1, the L1 routers 622 in Fig. 6 do not have to be separated into different areas (e.g. 134 in Fig. 1) using multiple L12 routers. Instead, according to the example hub-and-spoke topology 600, MTR1 is connected to L1 routers RT1 to RTn, e.g. MTR1 may be interpreted as a 'hub' that is connected to L1 routers RT1 to RTn via links or 'spokes'.

[0062] If no filtering strategy is pre-configured, LSDB isolation cannot be implemented among the L1 routers 622 in the network 600 because the L1 routers 622 are within the same area. Thus, link state information of an L1 router (e.g. RT1) is sent to every other neighbour L1 router (e.g. RT2 to RTn). For example, RT1 generates and sends link state information to MTR1 612, which forwards it to neighbour routers RT2 to RTn. This flooding process is repeated for link state information generated by other routers (RT2 to RTn) to synchronize their LSDBs.

[0063] Similar to Fig. 1, a filtering strategy (e.g. 630, 632, 640) may be pre-configured on an interface between an L12 router (i.e. MTR) and an L1 router (e.g. RT2) to enable complete or partial LSDB isolation.

[0064] In a first example, if it is not necessary for RT2 to process link state information generated by RT1, an interface filtering strategy 630 is pre-configured on a router interface between MTR and RT2 (i.e. P2). The interface filtering strategy 630 includes the system ID of RT1. When MTR receives link state information generated by RT1, MTR will not send the link state information to RT2. As such, the interface filtering strategy 630 enables LSDB isolation between RT1 and RT2.

[0065] In a second example, if it is not necessary for RT2 to receive and store link state information generated by all other L1 routers (i.e. RT1, RT3,..., RTn), a global filtering strategy 632 is pre-configured on an interface between MTR and RT2 (i.e. P2). According to the global filtering strategy 632, MTR only sends link state information generated by its UP neighbour MTR to RT2, but link state information generated by neighbour routers RT1, RT3,..., RTn will not be forwarded to RT2. As such, the global filtering strategy 632 enables complete LSDB isolation between RT2 each of its neighbour routers RT1, RT3,..., RTn within the access layer 130.

[0066] In a third example, if LSDB isolation is required between RT3 and its neighbour routers in the access layer 130, a global filtering strategy 640 may be pre-configured on an interface between MTR and RT3 (i.e. P3). This process may be repeated for any other router which requires complete LSDB isolation.

[0067] Although the filtering strategy is pre-configured on MTR's interface in the above examples, it should be understood that the filtering strategy may be configured on an L1 router in the access layer. For example, instead of pre-configuring a filtering strategy on router interface P2 of MTR, the strategy may be pre-configured on a router interface of RT2 (not labelled in Fig. 6 for simplicity).

[0068] Similarly, consider the situation where RT2 also sends link state information generated by RT3 to RT1. To reduce flooding of link state information between RT1 and RT2, an interface filtering strategy is configured on the interface between RT1 and RT2 such that RT2 does not send link state information generated by RT3 to RT1. In other words, once RT2 recognizes that the source router of the link state information is the system ID of RT3, RT2 will not send it to RT1.

[0069] In Fig. 1, L12 routers (i.e. MTR1 to MTR6) deployed in the convergence layer 120 allow L1 routers (i.e. RT1 to RTn) to be isolated into different areas (see 134) such that L1 routers from different areas do not have to exchange link state information. Since the example method in Fig. 2 enables LSDB isolation through pre-configuration of filtering strategy, a hub-and-spoke topology may be used in the convergence layer.

[0070] Example Network 700

[0071] Using pre-configuration of filtering strategy, a network topology may also be simplified, therefore reducing the number of routers required and corresponding deployment cost and efforts. For example in Fig. 1, L12 routers (i.e. MTR1 to MTR6) allow physical separation of L1 routers (i.e. RT1 to RTn) into different areas (see 134) such that L1 routers in different areas (e.g. RT1 and RT5) do not have to exchange link state information. However, the use of a filtering strategy enables LSDB isolation between routers (e.g. RT1 and RT2) in the same area. Therefore when a filtering strategy is used the network may require less L12 routers than would be the case if no filtering strategy was used.

[0072] Fig. 7 is a schematic diagram of a third example network environment 700 in which link state information flooding may be reduced. Compared with Fig. 1, the network environment 700 is simplified with less L12 routers deployed in the convergence layer 720. In particular, two routers MTR1 and MTR2 are deployed in Fig. 7 to support L1 routers RT1 to RTn compared to six routers MTR1 to MTR6 in the corresponding convergence layer 120 in Fig. 1. In this example, less L12 routers are required because a filtering strategy may be configured to enable complete or partial LSDB isolation, effectively separating the L1 routers into different logical areas to reduce link state information flooding.

[0073] Similar to Fig. 1, LSDB isolation between routers may be achieved by pre-configuring a filtering strategy (e.g. see 740 at MTR1). For example, as described, an interface filtering strategy 740 that includes the system ID of RT1 may be pre-configured on a router interface of MTR1. When link state information generated by RT1 is received, MTR1 does not send the link state information to RT2. As the filtering strategy 740 achieves LSDB isolation, less L1 routers need to be physically isolated or divided into different areas to achieve LSDB isolation. This in turn reduces the number of L12 routers required in the convergence layer 720.

[0074] Default link state information

[0075] Default link state information that includes default route information may be sent to a router from time to time.

[0076] Using the example in Fig. 6, when a global filtering strategy is pre-configured on a router interface (e.g. P2) between MTR and RT2, RT2 does not receive link state information of neighbour routers RT1, RT3 to RTn. Therefore, RT2 is unable to calculate routes to these routers.

[0077] To remedy this, MTR generates and sends default link state information that includes default route information of other routers to RT2. The default link state information specifies that RT2 can reach other parts of the network via MTR. For example, if RT2 wishes to send packets to a neighbour L1 router (e.g. RT1) but cannot find matching routing information to RT1 in its database, RT2 uses the default route information to reach RT1 via MTR.

[0078] In practice, default route '0.0.0.0/0' is used to indicate that any section of the network may be reached. When RT2 wishes to forward packets to a destination address but cannot find any routing information associated with the destination address, RT2 performs route calculation based on '0.0.0.0/0'. The calculated next hop will be MTR, i.e. the sender of the default route information. As such, RT2 will send all packets via MTR to reach another part in the network.

[0079] Network device 800

[0080] The above examples can be implemented by hardware, software or firmware or a combination thereof. Referring to Fig. 8, an example network device 800 capable of acting as a router (e.g. 112, 122, 132 in Fig. 1) is shown.

[0081] The example network device 800 includes a processor 810, a memory 820 and a network interface device 840 that communicate with each other via bus 830. The memory 820 stores any necessary data 822 and machine-readable instructions 824 to perform any of

the processes described in the present disclosure. For example, the data 822 may include information relating to filtering strategy to enable LSDB isolation between routers.

[0082] The processor 810 is to perform processes described herein. In one example, the processor 810 is to:

[0083] Pre-configure a filtering strategy on a router interface for filtering link state information generated by a first router. The filtering strategy is to enable LSDB isolation between a second router associated with the router interface and the first router.

[0084] Receive or send link state information generated by the first router according to the pre-configured filtering strategy via the router interface.

[0085] The memory 820 may store machine-readable instructions 824 to cause the processor 810 to perform processes described herein. In one example, the instructions 824 may include:

[0086] Filtering strategy instructions to cause the processor 810 to pre-configure a filtering strategy on a router interface for filtering link state information generated by a first router. The filtering strategy is to enable link state database (LSDB) isolation between a second router associated with the router interface and the first router.

[0087] Processing instructions to cause the processor 810 to receive or send link state information generated by the first router according to the pre-configured filtering strategy via the router interface.

[0088] The network device 800 in Fig. 8 may include units to perform the processes described herein. In one example, the network device 800 may include the following modules:

[0089] Filtering strategy module to pre-configure a filtering strategy on a router interface for filtering link state information generated by a first router. The filtering

strategy is to enable link state database (LSDB) isolation between a second router associated with the router interface and the first router.

[0090] The filtering strategy module may include a global filtering strategy unit and an interface routing strategy unit to implement the global filtering strategy and interface routing strategy respectively.

[0091] Processing module to receive or send link state information generated by the first router according to the pre-configured filtering strategy via the router interface.

[0092] The methods, processes and functional units described herein may be implemented by hardware (including hardware logic circuitry), software or firmware or a combination thereof. The term ‘processor’ is to be interpreted broadly to include a processing unit, ASIC, logic unit, or programmable gate array etc. The processes, methods and functional units may all be performed by the one or more processors 810; reference in this disclosure or the claims to a ‘processor’ should thus be interpreted to mean ‘one or more processors’.

[0093] Although one network interface device 840 is shown in Fig. 8, processes performed by the network interface device 840 may be split among multiple network interface devices (not shown for simplicity). As such, reference in this disclosure to a ‘network interface device’ should be interpreted to mean ‘one or more network interface devices’.

[0094] Further, the processes, methods and functional units described in this disclosure may be implemented in the form of a computer software product. The computer software product is stored in a storage medium and comprises a plurality of instructions for making a processor to implement the methods recited in the examples of the present disclosure.

[0095] The figures are only illustrations of an example, wherein the units or procedure shown in the figures are not necessarily essential for implementing the present disclosure. Those skilled in the art will understand that the units in the device in the example can be arranged in the device in the examples as described, or can be alternatively located in one or more devices different from that in the examples. The units in the examples described can be combined into one module or further divided into a plurality of sub-units.

[0096] Although the flowcharts described show a specific order of execution, the order of execution may differ from that which is depicted. For example, the order of execution of two or more blocks may be changed relative to the order shown. Also, two or more blocks shown in succession may be executed concurrently or with partial concurrence. All such variations are within the scope of the present disclosure.

[0097] It will be appreciated by persons skilled in the art that numerous variations and/or modifications may be made to the above-described embodiments, without departing from the broad general scope of the present disclosure. The present embodiments are, therefore, to be considered in all respects as illustrative and not restrictive.

Claims

1. A method for reducing flooding of link state information of a link state protocol in a network with multiple routers, comprising:

pre-configuring a filtering strategy on a router interface for filtering link state information generated by a first router, wherein the filtering strategy is to enable link state database (LSDB) isolation between a second router associated with the router interface and the first router; and

receiving or sending link state information generated by the first router according to the pre-configured filtering strategy via the router interface.

2. The method of claim 1, wherein the filtering strategy is a global filtering strategy to enable LSDB isolation between the second router and first router, and the first router is not an upstream neighbour of the second router.

3. The method of claim 2, wherein receiving or sending the link state information generated by the first router according to the global filtering strategy comprises:

checking whether the first router is an upstream neighbour of the second router;

if yes, receiving or sending the link state information via the router interface; but otherwise, discarding or not sending the link state information respectively.

4. The method of claim 2, the method further comprises:

sending, to the second router, default link state information to generate default routes to the first router.

5. The method of claim 1, wherein the filtering strategy is an interface filtering strategy to enable LSDB isolation between the second router and first router based on a system ID of the first router.

6. The method of claim 5, wherein receiving or sending the link state information generated by the first router according to the interface filtering strategy comprises:

checking whether the system ID of the first router satisfies the interface filtering strategy; and

if yes, receiving or sending the link state information via the router interface; but otherwise, discarding or not sending the link state information respectively.

7. The method of claim 1, wherein the router interface is a router interface of the second router, or a router interface of an upstream neighbour of the second router.

8. The method of claim 1, wherein the link state protocol is Intermediate System-to-Intermediate System (ISIS) and the link state information is a link state packet (LSP).

9. A network device for reducing flooding of link state information of a link state protocol in a network with multiple routers, wherein the network device is capable of acting as a router and comprises a processor to:

pre-configure a filtering strategy on a router interface for filtering link state information generated by a first router, wherein the filtering strategy is to enable link state database (LSDB) isolation between a second router associated with the router interface and the first router; and

receive or send link state information generated by the first router according to the pre-configured filtering strategy via the router interface.

10. The network device of claim 8, wherein the processor is to pre-configure a global filtering strategy to enable LSDB isolation between the second router and first router, and the first router is not an upstream neighbour of the second router.

11. The network device of claim 9, wherein when receiving or sending the link state information generated by the first router according to the global filtering strategy, the processor is to:

check whether the first router is an upstream neighbour of the second router;

if yes, receive or send the link state information via the router interface; but otherwise, discard or not send the link state information respectively.

12. The network device of claim 9, the processor is further to:

send, to the second router, default link state information to generate default routes to the first router.

13. The network device of claim 8, wherein the processor is to pre-configure an interface filtering strategy to enable LSDB isolation between the second router and first router based on a system ID of the first router.

14. The network device of claim 12, wherein when receiving or sending the link state information generated by the first router according to the interface filtering strategy, the processor is to:

check whether the system ID of the first router satisfies the interface filtering strategy;
and

if yes, receive or send the link state information via the router interface; but otherwise, discard or not send the link state information respectively.

15. The network device of claim 8, wherein the link state protocol is Intermediate System-to-Intermediate System (ISIS) and the link state information is a link state packet (LSP).

1/8




Fig. 1

2/8

200




Fig. 2

3/8




Fig. 3A

Filtering Strategy at P1 ~ 310

Type	Action	Source router	Yes/No
Global	Transmit to RT2	RT1, RT3	No

Filtering Strategy at R2 ~ 320

Type	Action	Source router	Yes/No
Global	Receive from MTR1	RT1, RT3	No

Fig. 3B

Filtering Strategy at P1 ~ 330

Type	Action	Source router	Yes/No
Interface	Transmit to RT2	RT1	No

Filtering Strategy at R2 ~ 340

Type	Action	Source router	Yes/No
Interface	Receive from MTR1	RT1	No

Fig. 3C

4/8




Fig. 4

5/8




Fig. 5

6/8




Fig. 6

718

700




Fig. 7

8/8




Fig. 8

INTERNATIONAL SEARCH REPORT

International application No.

PCT/CN2013/077369

A. CLASSIFICATION OF SUBJECT MATTER

H04L12/755 (2013.01) i

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC: H04L

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

CPRSABS, DWPI, CNKI: link, state, information, advertis+, flood+, filter+, upstream+, discard+

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2009041037 A1 (CISCO TECHNOLOGY INC) 12 February 2009 (12.02.2009) see claims 1 and 2	1, 5-9, 13-15
A	see the whole document	2-4, 10-12
A	US 7334047 B1 (CISCO TECHNOLOGY INC) 19 February 2008 (19.02.2008) see the whole document	1-15
A	US 6820134 B1 (CISCO TECHNOLOGY INC) 16 November 2004 (16.11.2004) see the whole document	1-15

☐ Further documents are listed in the continuation of Box C.☒ See patent family annex.

* Special categories of cited documents:	“T” later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
“A” document defining the general state of the art which is not considered to be of particular relevance	“X” document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
“E” earlier application or patent but published on or after the international filing date	“Y” document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
“L” document which may throw doubts on priority claim (S) or which is cited to establish the publication date of another citation or other special reason (as specified)	“&” document member of the same patent family
“O” document referring to an oral disclosure, use, exhibition or other means	
“P” document published prior to the international filing date but later than the priority date claimed	

Date of the actual completion of the international search 05 September 2013 (05.09.2013)	Date of mailing of the international search report 19 Sep. 2013 (19.09.2013)
Name and mailing address of the ISA/CN The State Intellectual Property Office, the P.R.China 6 Xitucheng Rd., Jimen Bridge, Haidian District, Beijing, China 100088 Facsimile No. 86-10-62019451	Authorized officer QIN, Sheng Telephone No. (86-10)62411225

International application No.
PCT/CN2013/077369

Form PCT/ISA /210 (patent family annex) (July 2009)