



(51) International Patent Classification:
Not classified

(21) International Application Number:
PCT/US2023/081291

(22) International Filing Date:
28 November 2023 (28.11.2023)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
63/443,253 03 February 2023 (03.02.2023) US

(71) Applicant: **FUTUREWEI TECHNOLOGIES, INC.**
[US/US]; 15851 Dallas Parkway, Suite 650, Addison, Texas
75001 (US).


(72) Inventor: **CHEN, Huaimo**; 15851 Dallas Parkway, Suite
650, Addison, Texas 75001 (US).

(74) Agent: **DIETRICH, William H.** et al.; Conley Rose, P.C.,
4965 Preston Park Boulevard, Suite 195E, Plano, Texas
75093 (US).

(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, CV, CZ, DE, DJ, DK, DM,
DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT,
HN, HR, HU, ID, IL, IN, IQ, IR, IS, IT, JM, JO, JP, KE, KG,
KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY,
MA, MD, MG, MK, MN, MU, MW, MX, MY, MZ, NA,
NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO,
RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, ST, SV, SY, TH,
TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, WS,
ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every
kind of regional protection available): ARIPO (BW, CV,
GH, GM, KE, LR, LS, MW, MZ, NA, RW, SC, SD, SL, ST,
SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ,

(54) Title: GENERAL FAST EGRESS PROTECTION



(57) Abstract: A method implemented by a network node neighboring a primary egress node in a domain, comprising receiving a piece of egress protection information that identifies a backup egress node for the primary egress node; making a determination, in response to receiving the egress protection information, to protect the primary egress node with the backup egress node when the primary egress node fails; detecting whether the primary egress node has failed when forwarding a packet to the primary egress node; and forwarding, in response to the determination and in response to detecting the primary egress node has failed, the packet to the backup egress node to transmit the packet to a customer edge (CE) receiver.

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, ME, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Declarations under Rule 4.17:

- *as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))*
- *as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))*

Published:

- *upon request of the applicant, before the expiration of the time limit referred to in Article 21(2)(a)*
- *without international search report and to be republished upon receipt of that report (Rule 48.2(g))*

General Fast Egress Protection

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This patent application claims the benefit of U.S. Provisional Patent Application No. 63/443,253 filed on February 3, 2023 by Huaimo Chen and titled “General Fast Egress Protection,” which is hereby incorporated by reference.

TECHNICAL FIELD

[0002] The present disclosure is generally related to the field of egress protection and, in particular, to egress protection against the failure of a node or link in a domain or a network.

BACKGROUND

[0003] There is an existing method for fast protection against the failure of an egress node of a bit index explicit replication (BIER) domain using a BIER protocol. BIER egress protection is described in further detail in Internet Engineering Task Force (IETF) document entitled “BIER Egress Protection” by H. Chen, et al., published July 2023. However, there are no techniques for fast protection against the failure of an egress node in another domain. Therefore, a customer edge (CE) receiver will not receive a requested packet when the egress node in the other domain fails.

SUMMARY

[0004] The disclosed aspects/embodiments provide fast egress protection (EP) against the failure of an egress node of various types of domains or networks. These domains or networks may include a Border Gateway Protocol (BGP) domain, a BIER domain, a segment routing version 6 (SRv6) domain, or an interior gateway protocol (IGP) domain. To facilitate the fast egress protection procedure, a backup egress node of an egress node distributes information about egress protection to neighboring nodes of the egress node. The information is used by the neighboring nodes to forward packets to the backup egress node when the egress node has failed. Therefore, packet routing within various domains is improved.

[0005] A first aspect relates to a method implemented by a network node neighboring a primary egress node in a domain, comprising receiving a piece of egress protection information that identifies a backup egress node for the primary egress node; making a determination, in response to receiving the egress protection information, to protect the primary egress node with the backup egress node when the primary egress node fails; detecting whether the primary egress

node has failed when forwarding a packet to the primary egress node; and forwarding, in response to the determination and in response to detecting the primary egress node has failed, the packet to the backup egress node to transmit the packet to a customer edge (CE) receiver.

[0006] Optionally, in any of the preceding aspects, another implementation of the aspect provides that the egress protection information is received from the backup egress node or from both the primary egress node and the backup egress node.

[0007] Optionally, in any of the preceding aspects, another implementation of the aspect further comprising detecting whether the primary egress node has failed via a bi-directional forwarding detection (BFD) session between the network node and the primary egress node.

[0008] Optionally, in any of the preceding aspects, another implementation of the aspect provides that the domain is a Border Gateway Protocol (BGP) domain, an interior gateway protocol (IGP) domain, a bit index explicit replication (BIER) domain, or a segment routing version 6 (SRv6) domain, wherein the method further comprises receiving the egress protection information that comprises an IP address prefix for an IP anycast address from the primary egress node and the backup egress node, and wherein the IP anycast address is configured on the primary egress node and the backup egress node by a network operator or tools.

[0009] Optionally, in any of the preceding aspects, another implementation of the aspect provides that when the domain is the BGP domain, the method further comprises identifying a remote node using information in a routing information base (RIB); and forwarding, after detecting the primary egress node has failed, the packet with IP anycast address to the remote node through a tunnel.

[0010] Optionally, in any of the preceding aspects, another implementation of the aspect provides that when the domain is the IGP domain, the method further comprises identifying, in response to the determination, an intermediate node that is a loop free alternative (LFA) to the backup egress node; and after detecting the primary egress node has failed, either forwarding the packet with IP anycast address to the intermediate node when the intermediate node is directly connected to the network node, or forwarding the packet with IP anycast address, through a tunnel, to the intermediate node when the intermediate node is not directly connected to the network node.

[0011] Optionally, in any of the preceding aspects, another implementation of the aspect provides that when the domain is the BIER domain, the method further comprises receiving the egress protection information that comprises the IP address prefix for the IP anycast address and a Bit-Forwarding Router Identifier (BFR-ID) for the IP address prefix; identifying, in response to the receiving, an intermediate node that is a loop free alternative (LFA) to the backup egress

node; and after detecting the primary egress node has failed, either forwarding the packet with the BFR-ID to the intermediate node when the intermediate node is directly connected to the network node, or forwarding the packet with IP anycast address and the BFR-ID, through a tunnel, to the intermediate node when the intermediate node is not directly connected to the network node.

[0012] Optionally, in any of the preceding aspects, another implementation of the aspect provides that when the domain is the SRv6 domain, the method further comprises identifying, in response to the determination, an intermediate node that is a loop free alternative (LFA) to the backup egress node; and forwarding, after detecting the primary egress node has failed, the packet with the IP anycast address to the intermediate node.

[0013] Optionally, in any of the preceding aspects, another implementation of the aspect provides that the LFA comprises a basic LFA or a topology independent (TI) LFA.

[0014] Optionally, in any of the preceding aspects, another implementation of the aspect further comprising forwarding, in response to detecting the primary egress node has not failed, the packet with the IP anycast address to the primary egress node according to forwarding information base (FIB) or bit index forwarding table (BIFT).

[0015] A second aspect relates to a network node neighboring a primary egress node in a domain, comprising: a memory storing instructions; and one or more processors coupled to the memory, wherein the one or more processors are configured to execute the instructions to cause the network node to receive a piece of egress protection information that identifies a backup egress node for the primary egress node; make a determination, in response to receiving the egress protection information, to protect the primary egress node with the backup egress node when the primary egress node fails; detect whether the primary egress node has failed when forwarding a packet to the primary egress node; and forward, in response to the determination and in response to detecting the primary egress node has failed, the packet to the backup egress node to transmit the packet to a customer edge (CE) receiver.

[0016] Optionally, in any of the preceding aspects, another implementation of the aspect provides that the egress protection information is received from the backup egress node or from both the primary egress node and the backup egress node.

[0017] Optionally, in any of the preceding aspects, another implementation of the aspect provides that one or more processors are configured to execute the instructions to further cause the network node to detect whether the primary egress node has failed via a bi-directional forwarding detection (BFD) session between the network node and the primary egress node.

[0018] Optionally, in any of the preceding aspects, another implementation of the aspect provides that the domain is a Border Gateway Protocol (BGP) domain, an interior gateway protocol (IGP) domain, a bit index explicit replication (BIER) domain, or a segment routing version 6 (SRv6) domain, wherein the one or more processors are configured to execute the instructions to further cause the network node to receive the egress protection information that comprises an IP address prefix for an IP anycast address from the primary egress node and the backup egress node, and wherein the IP anycast address is configured on the primary egress node and the backup egress node by a network operator or tools.

[0019] Optionally, in any of the preceding aspects, another implementation of the aspect provides that when the domain is the BGP domain, the one or more processors are configured to execute the instructions to further cause the network node to identify a remote node using information in a routing information base (RIB); and forward, after detecting the primary egress node has failed, the packet with IP anycast address to the remote node through a tunnel.

[0020] Optionally, in any of the preceding aspects, another implementation of the aspect provides that when the domain is the IGP domain, the one or more processors are configured to execute the instructions to further cause the network node to identify, in response to the determination, an intermediate node that is a loop free alternative (LFA) to the backup egress node; and after detecting the primary egress node has failed, either forward the packet with IP anycast address to the intermediate node when the intermediate node is directly connected to the network node, or forward the packet with IP anycast address, through a tunnel, to the intermediate node when the intermediate node is not directly connected to the network node.

[0021] Optionally, in any of the preceding aspects, another implementation of the aspect provides that when the domain is the BIER domain, the one or more processors are configured to execute the instructions to further cause the network node to receive the egress protection information that comprises the IP address prefix for the IP anycast address and a Bit-Forwarding Router Identifier (BFR-ID) for the IP address prefix; identify, in response to the receiving, an intermediate node that is a loop free alternative (LFA) to the backup egress node; and after detecting the primary egress node has failed, either forward the packet with the BFR-ID to the intermediate node when the intermediate node is directly connected to the network node, or forward the packet with the IP anycast address and the BFR-ID, through a tunnel, to the intermediate node when the intermediate node is not directly connected to the network node.

[0022] Optionally, in any of the preceding aspects, another implementation of the aspect provides that when the domain is the SRv6 domain, the one or more processors are configured to execute the instructions to further cause the network node to identify, in response to the

determination, an intermediate node that is a loop free alternative (LFA) to the backup egress node; and forward, after detecting the primary egress node has failed, the packet with the IP anycast address to the intermediate node.

[0023] Optionally, in any of the preceding aspects, another implementation of the aspect provides that the LFA comprises a basic LFA or a topology independent (TI) LFA.

[0024] Optionally, in any of the preceding aspects, another implementation of the aspect provides that one or more processors are configured to execute the instructions to further cause the network node to forward, in response to detecting the primary egress node has not failed, the packet with the IP anycast address to the primary egress node according to forwarding information base (FIB) or bit index forwarding table (BIFT).

[0025] A third aspect relates to a non-transitory computer readable medium comprising a computer program product for use by a network node, the computer program product comprising computer executable instructions stored on the non-transitory computer readable medium that, when executed by one or more processors, cause the first network node to execute the method of any of the first aspect.

[0026] A fourth aspect relates to a network node neighboring a primary egress node in a domain, comprising receiving means configured to receive a piece of egress protection information that identifies a backup egress node for the primary egress node; processing means configured to make a determination, in response to receiving the egress protection information, to protect the primary egress node with the backup egress node when the primary egress node fails; and detect whether the primary egress node has failed when forwarding a packet to the primary egress node; and transmitting means configured to forward, in response to the determination and in response to detecting the primary egress node has failed, the packet to the backup egress node to transmit the packet to a customer edge (CE) receiver.

[0027] For the purpose of clarity, any one of the foregoing embodiments may be combined with any one or more of the other foregoing embodiments to create a new embodiment within the scope of the present disclosure.

[0028] These and other features will be more clearly understood from the following detailed description taken in conjunction with the accompanying drawings and claims.

BRIEF DESCRIPTION OF THE DRAWINGS

[0029] For a more complete understanding of this disclosure, reference is now made to the following brief description, taken in connection with the accompanying drawings and detailed description, wherein like reference numerals represent like parts.

[0030] FIG. 1 is a schematic diagram of a network topology for fast egress protection according to an embodiment of the disclosure.

[0031] FIG. 2A is a schematic diagram of a BGP topology including a BGP domain according to an embodiment of the disclosure.

[0032] FIG. 2B is a schematic diagram of fast egress protection in a BGP domain according to an embodiment of the disclosure.

[0033] FIG. 3A is a schematic diagram of an IGP topology including an IGP domain according to an embodiment of the disclosure.

[0034] FIG. 3B is a schematic diagram of fast egress protection in an IGP domain according to an embodiment of the disclosure.

[0035] FIG. 4A is a schematic diagram of a BIER topology including a BIER domain according to an embodiment of the disclosure.

[0036] FIG. 4B is a schematic diagram of fast egress protection in a BIER domain according to an embodiment of the disclosure.

[0037] FIG. 5A is a schematic diagram of an SRv6 topology including an SRv6 domain according to an embodiment of the disclosure.

[0038] FIG. 5B is a schematic diagram of fast egress protection in an SRv6 domain according to an embodiment of the disclosure.

[0039] FIG. 6 is a method implemented by a network node in a domain according to an embodiment of the disclosure.

[0040] FIG. 7 is a schematic diagram of a network apparatus according to an embodiment of the disclosure.

DETAILED DESCRIPTION

[0041] It should be understood at the outset that although an illustrative implementation of one or more embodiments is provided below, the disclosed systems and/or methods may be implemented using any number of techniques, whether currently known or in existence. The disclosure should in no way be limited to the illustrative implementations, drawings, and techniques illustrated below, including the exemplary designs and implementations illustrated and described herein, but may be modified within the scope of the appended claims along with their full scope of equivalents.

[0042] There are some existing methods for fast protection against the failure of an egress node. For example, there is fast protection against failure of an egress node in a BIER domain using a BIER protocol. However, for some domains such as a BGP domain or an IGP domain,

there are no techniques specified for the fast egress protection. Therefore, a customer edge (CE) receiver will not receive a requested packet when the egress node in a given domain fails.

[0043] Disclosed herein is a fast egress protection (EP) for various types of domains or networks including a BGP domain, an IGP domain, a BIER domain, and an SRv6 domain.

[0044] To facilitate the fast egress protection procedure, a backup egress node of an egress node distributes information about egress protection to neighboring nodes of the egress node. The information is used by the neighboring nodes to forward packets to the backup egress node when the egress node has failed. Therefore, packet routing and network reliability within various domains are improved. The disclosed embodiments can be deployed in any router, switch, and controller, which are used by service providers around the world.

[0045] FIG. 1 is a schematic diagram of a network topology 100 including a domain 102. The domain 102 may be a BGP domain, an IGP domain, a BIER domain, or an SRv6 domain. The domain 102 comprises a plurality of network nodes 104, 106, 108, 110, 112, 114, 116, and 118. While eight network nodes 104-118 are shown in the domain 102, more or fewer nodes may be included in practical applications.

[0046] For ease of discussion, all of the network nodes 104-118 have been given a letter designation. For example, the network node 104 has the designation A, the network node 106 has the designation B, the network node 108 has the designation C, the network node 110 has the designation D, the network node 112 has the designation E, the network node 114 has the designation F, the network node 116 has the designation G, and the network node 118 has the designation H.

[0047] Each of the network nodes 104-118 may comprise a router, switch, or other telecommunications device configured to receive, route, store, and transmit packets. Some of the network nodes, namely the network nodes 104, 106, 112, and 118, are disposed at an edge of the domain 102. The network nodes 104, 106, 112, and 118 receiving packets from outside the domain 102 may be referred to as ingress network nodes (or simply, ingress nodes). The network nodes 104, 106, 112, and 118 transmitting packets out of the domain 102 may be referred to as egress network nodes (or simply, egress nodes). Depending on the direction of packet traffic, each of the network nodes 104, 106, 112, and 118 may function as an ingress network node and/or an egress network node. The network nodes 108, 110, 114, and 116 forwarding packets within the domain 102 may be referred to as transit network nodes.

[0048] Each of the network nodes 104, 106, 112, and 118 may be referred to herein as a destination network node. The network nodes 104, 106, 112, and 118 have each been assigned a bit forwarding router identifier (BFR-id) or an internet protocol (IP) anycast address. Anycast

is a well-known network addressing and routing methodology in which a single destination IP address is shared by a group of network nodes in multiple locations. The routing algorithm selects the single receiver from the group based on which is the nearest according to some distance or cost measure. For example, each of the network nodes 104, and 106 is configured with an IP anycast address of 10.1.1.0.

[0049] Each of the network nodes 104-118 has one or more neighbor nodes. As used herein, a neighbor node refers to a network node that is only one hop away from the network node. For example, network node 104 has two neighbor nodes in FIG. 1, namely network node 108 and network node 116.

[0050] The network nodes 104-118 in FIG. 1 are coupled to, and communicate with each other, via links 122. The links 122 may be wired, wireless, or some combination thereof. Each of the links 122 has a cost. As shown, the network node 104 and the network node 106 are each coupled to a customer edge (CE) receiver 120 outside the domain 102. Thus, the CE receiver 120 is configured to receive packets from, and send packets to, the network node 104 and the network node 106.

[0051] The network node 104 may be referred to herein as a primary egress node (or simply an egress node), and the network node 106 may be referred to herein as a backup egress node for the primary egress node 104. In an embodiment, a primary egress node is a network node that functions as the primary point of egress for traffic received from a neighbor node while the network node is functioning normally. In an embodiment, a backup egress node is a network node that functions as the backup point of egress for traffic received from a neighbor node when the primary egress node is functioning abnormally or has failed.

[0052] The disclosed embodiments include mechanisms for fast protection against the failure of an egress node in various types of domains or networks including a BGP domain, an IGP domain, a BIER domain, and an SRv6 domain. In an embodiment, for fast protection of an egress node of a given domain, when a backup egress node (e.g., node 106) is selected to protect the primary egress node (e.g., node 104), egress protection information is configured on the backup egress node to protect the primary egress node (e.g., node 104). After the configuration, the backup egress node distributes the egress protection information, through IGP, to the direct neighbors of the primary egress node. The egress protection information indicates that the backup egress node protects the primary egress node.

[0053] For a packet to the primary egress node 104 of the domain 102, in normal operations (i.e., the primary egress node is functioning normally), an upstream hop (e.g., node 108 or node 116) sends the packet to the egress node 104 according to its forwarding information base (FIB)

(as shown by an arrow 124 in FIG. 1). In an embodiment, when the primary egress node 104 fails, an upstream hop (e.g., node 108 or node 116) detects the failure of the primary egress node 104 using a fast failure detection method such as a bi-directional forwarding detection (BFD) method and as a point of local repair (PLR) extend its forwarding information base (FIB) for protecting the primary egress node with the backup egress node 106. The upstream hop, using the egress protection information, sends the packet to the backup egress node 106 configured to protect the primary egress node once the PLR detects the failure (as shown by an arrow 126 in FIG. 1). The upstream hop of the primary egress node is its direct neighbor. The backup egress node 106 then sends the packet to the CE receiver 120.

[0054] FIG. 2A is a schematic diagram of a BGP topology 200A including a plurality of BGP domains according to an embodiment of the disclosure. In general, redundancy in a BGP network is provided at the routing level using an IP anycast address. However, using this redundancy for egress protection is slow. When an egress node fails, the traffic to the egress node continues flowing toward the egress node and gets lost until the network converges on the failure. Furthermore, the network convergence on BGP takes seconds.

[0055] Disclosed herein is an implementation of the configuration and distribution of fast egress protection information for the BGP topology 200A. The BGP topology 200A comprises a plurality of BGP domains (or autonomous systems (ASs)) 202, 204, and 206. While three BGP domains/ASs are shown in the BGP topology 200A, more or fewer domains may be included in practical applications. For ease of discussion, all of the BGP domains/ASs 202-206 have been given letter designation. For example, the BGP domain 202 has the designation AS1, the BGP domain 204 has the designation AS2, and the BGP domain 206 has the designation AS3.

[0056] In an embodiment, the network node 208 (or node A) may be referred to herein as a primary egress node (or simply an egress node), and the network node 210 (or node B) may be referred to herein as a backup egress node for the primary egress node 208. The network nodes 208 and 210 are configured with an IP anycast address of 10.1.1.0. In an embodiment, when a neighboring node such as node C (or node 212) receives a packet with the IP anycast address as its destination, the node C sends the packet toward the closer node with the IP anycast address. Once the node C receives the packet, node C has at least two paths from the node C to the destination (i.e., the IP anycast address). In an embodiment, a first path is from node C to node A and a second path is from node C to node B. The BGP routing algorithm selects the path from the at least two paths based on which is the nearest according to cost measure. In a first case, when the first path from node C to node A has the least cost (or the shortest), node C sends

the packet toward node A. In a second case, when the path from node C to node B has the least cost, node C sends the packet toward node B. In a third case, when the path from node C to node A has the same cost as the path from node C to node B, node C sends the packet to node A (or node B) depending on whether node C sends the last packet with the same address to node B (or node A).

[0057] In an embodiment, when one of the two nodes A and B fails and the network converges on the failure, there is only one path from node C to the destination (i.e., the IP anycast address). If node A fails, there is one path from node C to node B and node C sends the packet to node B. If node B fails, there is one path from node C to node A and node C sends the packet to node A. As such, when the IP anycast address is configured on nodes A and B for redundancy, the BGP routes, for the IP anycast address on nodes A and B, are distributed in the BGP topology 200A. The configuration and distribution are reused for fast protection of egress node A (or B).

[0058] A mechanism of determination of fast egress protection in the BGP topology 200A is described as follows. In an embodiment, the primary egress node 208 (or node A) and the backup egress node 210 (or node B) are configured with an IP anycast address of 10.1.1.0. After the configuration, node A and node B distribute the egress protection information about the backup egress to the direct neighbors of node A (as shown by arrow 214 in FIG. 2A). In an embodiment, the egress protection information is the IP anycast address distributed by node A and node B. Then, Node C determines whether to provide fast egress protection for node A based on the received egress protection information and pre-defined conditions. The pre-defined conditions are: 1) Node C is connected to node A, 2) Node C can fast detect the failure of node A using BFD, 3) Node C has two paths to the prefix corresponding to the same IP anycast address configured on node A and node B (e.g., node C has two paths to prefix 10.1.1.0/24 corresponding to IP anycast address 10.1.1.0 configured on nodes A and B as shown in FIG. 2A), and 4) The path from node C to the prefix through node A has less cost than the path from node C to the prefix through node B (e.g., the path from node C to 10.1.1.0/24 through node A has less cost than the path from node C to 10.1.1.0/24 through node B as shown in FIG. 2A). Thus, node C determines that it provides the fast protection for egress node A when the four conditions are met.

[0059] FIG. 2B is a schematic diagram of fast egress protection in a BGP network topology 200B according to an embodiment of the disclosure. In an embodiment, after node C determines to protect the primary egress node A with the backup egress node B, node C provides fast protection against the failure of the egress node as follows. In an embodiment, in normal operations (i.e., node A is functioning normally), after receiving a packet with IP anycast address 10.1.1.0 as destination, node C sends the packet to node A according to its FIB (as shown by an

arrow 218 in FIG. 2B). In an embodiment, when the node A is functioning abnormally or has failed, node C detects the failure of node A using a fast failure detection method such as a bi-directional forwarding detection (BFD) method. After detecting the failure of node A, node C sends the packet with IP anycast address 10.1.1.0 as the destination to a remote node (e.g., node F 216) through a tunnel such as an internet protocol (IP) tunnel (as shown by an arrow 220 in FIG. 2B). Node F is far enough away from 10.1.1.0 through node A such that the path from the node F to 10.1.1.0 through node B is shorter than the path from the node F to 10.1.1.0 through node A. In an embodiment, Node C finds the remote node using the information in Routing Information Base (RIB) such as Adj-RIBs-In and local (LOC)-RIB, and builds the IP tunnel to the remote node in advance. In an embodiment, when node F receives the packet with IP anycast address 10.1.1.0 as a destination from the IP tunnel, node F sends the packet towards node B instead of node A since node F has a shorter path to 10.1.1.0 through node B.

[0060] FIG. 3A is a schematic diagram of an IGP topology 300A including an IGP domain 302 according to an embodiment of the disclosure. As shown in FIG. 3A, an IGP domain 302 comprises an area 304 with an implementation of configuration and distribution of fast egress protection information. The network nodes and links depicted in FIG. 3A are similar to the network nodes and links depicted in FIG. 1. For the sake of brevity, a detailed description of these elements is not repeated herein. In an embodiment, the network node 104 (or node A) may be referred to herein as a primary egress node (or simply an egress node), and the network node 106 (or node B) may be referred to herein as a backup egress node for the primary egress node 104. The network nodes 104 and 106 are configured with an IP anycast address of 10.1.1.0. In an embodiment, when a neighboring node such as node 108 (or node C) receives a packet with the IP anycast address as its destination, the node 108 sends the packet toward the closer node with the IP anycast address. Once the node C receives the packet, node C has at least two paths from the node C to the destination (i.e., the IP anycast address). In an embodiment, a first path is from node C to node A and a second path is from node C to node B. The IGP routing algorithm selects the path from the at least two paths based on which is the nearest according to cost measure. Assume that the cost from node B to the IP anycast address makes the path from node C to the address via node A is shorter than the path from the node C to the address via node B.

[0061] In an embodiment, when the IP anycast address is configured on nodes A and B, the IGP running on A and B distribute the IP address prefix 10.1.1.0/32 for the IP anycast address configured on nodes A and B. Every node in the area 304 receives the prefix. Thus, the egress protection information about the prefix distributed in the area 304 is used by the

neighbors of the node A for fast protection against the failure of node A. The configuration and distribution are reused for fast protection of egress node A (or B).

[0062] A mechanism of determination of fast egress protection in the IGP domain 302 is described as follows. In an embodiment, the primary egress node 104 (or node A) and the backup egress node 106 (or node B) are configured with an IP anycast address of 10.1.1.0. After the configuration, node A and node B distribute the egress protection information about the backup egress to the direct neighbors of the node A (as shown by an arrow 306 in FIG. 3A). Then, Node C determines whether to provide fast egress protection for node A based on the received egress protection information and pre-defined conditions. The pre-defined conditions are: 1) Node C is connected to node A, 2) Node C can fast detect the failure of node A using BFD, 3) Node C has two paths to the IP address prefix corresponding to the same IP anycast address configured on node A and node B (e.g., node C has two paths to prefix 10.1.1.0/32 corresponding to IP anycast address 10.1.1.0 configured on nodes A and B as shown in FIG. 3A), and 4) The path from node C to the IP address prefix through node A is shorter than the path from node C to the IP address prefix through another node B since the cost from node B to the address configured makes the former path shorter than the latter. For example, the path from node C to the prefix through node A has less cost than the path from node C to the prefix through node B. Thus, node C determines that it can provide the fast protection for egress node A when the four conditions are met.

[0063] FIG. 3B is a schematic diagram of fast egress protection in an IGP network topology 300B according to an embodiment of the disclosure. In an embodiment, after determining to protect the primary egress node A with the backup egress node B, node C provides fast protection against the failure of the egress node as follows. In an embodiment, in normal operations (i.e., node A is functioning normally), after receiving a packet with IP anycast address 10.1.1.0 as destination, node C sends the packet to node A according to its FIB (as shown by an arrow 308 in FIG. 3B). In an embodiment, when the node A is functioning abnormally or has failed, node C detects the failure of node A using a fast failure detection method such as a bi-directional forwarding detection (BFD) method. After detecting the failure of node A, node C sends the packet with IP anycast address 10.1.1.0 as destination to a loop free alternative (LFA). In an embodiment, the LFA comprises a basic LFA or a topology independent (TI) LFA. The LFA node is far enough away from 10.1.1.0 through node A such that the path from the LFA to 10.1.1.0 through node B is shorter than the path from the LFA to 10.1.1.0 through node A (or the path from the LFA to B is shorter than the path from the LFA to node A).

[0064] In an embodiment, assume LFA is node B and node C is not directly connected to the LFA (or node B), node C sends the packet with 10.1.1.0 as destination to the LFA through a tunnel to the LFA. For example, node C sends the packet to node B through an IP tunnel or an SR tunnel from node C to node B. In an embodiment, assume LFA is node D and node C is directly connected to the LFA, node C sends the packet to the LFA without using any tunnel (as shown by an arrow 310 in FIG. 3A). In an embodiment, when the LFA such as node D receives the packet with 10.1.1.0 as destination, the LFA sends the packet towards node B instead of node A since the LFA has a shorter path to 10.1.1.0 through node B. Node B then sends the packet to the CE receiver 120.

[0065] FIG. 4A is a schematic diagram of a BIER topology 400A including a BIER domain 402 according to an embodiment of the disclosure. As shown in FIG. 4A, a BIER domain 402 comprises an area 404 with an implementation of configuration and distribution of fast egress protection information. The network nodes and links depicted in FIG. 4A are similar to the network nodes and links depicted in FIG. 1. For the sake of brevity, a detailed description of these elements is not repeated herein. In an embodiment, the network node 104 (or node A) may be referred to herein as a primary egress node (or simply an egress node), and the network node 106 (or node B) may be referred to herein as a backup egress node for the primary egress node 104. The network nodes 104 and 106 are configured with an IP anycast address of 10.1.1.0. In an embodiment, when a neighboring node such as node 108 (or node C) receives a packet with the IP anycast address as its destination, the node 108 sends the packet toward the closer node with the IP anycast address. Once the node C receives the packet, node C has at least two paths from the node C to the destination (i.e., the IP anycast address). In an embodiment, a first path is from node C to node A and a second path is from node C to node B. The routing algorithm selects the path from the at least two paths based on which is the nearest according to cost measure. Assume that the cost from node B to the address configured makes the path from node C to the address via A shorter than the path from the node C to the address via B.

[0066] In an embodiment, when the IP anycast address is configured on nodes A and B, the IGP running on A and B distributes the IP address prefix 10.1.1.0/32 for the IP anycast address configured on nodes A and B. Every node in the area receives the prefix. In addition, a Bit-Forwarding Router Identifier (BFR-ID) is configured for each IP anycast address prefix as a unicast IP address prefix. The mapping between the BFR-ID and the IP anycast address prefix is distributed in the area 404. Thus, the egress protection information about the prefix and BFR-ID distributed in the area 404 is used by the neighbors of the node A for fast protection against

the failure of node A. The configuration and distribution are reused for fast protection of egress node A (or B).

[0067] A mechanism of determination of fast egress protection in the BIER domain 402 is described as follows. In an embodiment, the primary egress node 104 (or node A) and the backup egress node 106 (or node B) are configured with an IP anycast address of 10.1.1.0. After the configuration, node A and node B distributes the egress protection information about the backup egress to the direct neighbors of the node A (as shown by an arrow 406 in FIG. 3A). Then, Node C as a direct neighbor of the node A determines whether to provide fast egress protection for node A based on the received egress protection information and pre-defined conditions. The pre-defined conditions are: 1) Node C is connected to node A, 2) Node C can fast detect the failure of node A using a fast failure detection method such as BFD, 3) Node C has two paths to the IP address prefix corresponding to the same IP anycast address configured on node A and node B (e.g., node C has two paths to prefix 10.1.1.0/32 corresponding to IP anycast address 10.1.1.0 configured on nodes A and B as shown in FIG. 4A), 4) The path from node C to the IP address prefix 10.1.1.0/32 through node A is shorter than the path from node C to 10.1.1.0/32 through node B. For example, the path from node C to the prefix through node A has less cost than the path from node C to the prefix through node B, and 5) The prefix 10.1.1.0/32 for anycast address 10.1.1.0 on node A has BFR-ID 101, which is the same as the BFR-ID 101 that the prefix 10.1.1.0/32 for the anycast address 10.1.1.0 on node B has. Thus, node C determines that it can provide the fast protection for egress node A when the five conditions are met. In an embodiment, the pre-defined conditions are the first four conditions of the five conditions: 1) to 4) as described above. In this case, the prefix 10.1.1.0/32 for anycast address 10.1.1.0 on node A has BFR-ID 101, which may be different from the BFR-ID that the prefix 10.1.1.0/32 for the anycast address 10.1.1.0 on node B has. For example, the prefix 10.1.1.0/32 for anycast address 10.1.1.0 on node B has BFR-ID 102.

[0068] FIG. 4B is a schematic diagram of fast egress protection in a BIER network 400B according to an embodiment of the disclosure. In an embodiment, after determining to protect the primary egress node A with the backup egress node B, node C provides fast protection against the failure of the egress node as follows. In an embodiment, in normal operations (i.e., node A is functioning normally), after receiving a packet with bitstring comprising a bit of value one for the BFR-ID of prefix 10.1.1.0/32 as destination, node C sends the packet to node A according to its Bit Index Forwarding Table (BIFT) (as shown by an arrow 408 in FIG. 4B). In an embodiment, when the node A is functioning abnormally or has failed, node C detects the failure of node A using a fast failure detection method such as a bi-directional forwarding detection

(BFD) method. After detecting the failure of node A, node C sends the packet with bitstring having a bit of value 1 for 10.1.1.0/32 to a loop free alternative (LFA). In an embodiment, the LFA comprises a basic LFA or a topology independent (TI) LFA. The LFA node is far enough away from 10.1.1.0 through node A such that the path from the LFA to 10.1.1.0 through node B is shorter than the path from the LFA to 10.1.1.0 through node A (or the path from the LFA to B is shorter than the path from the LFA to node A). When the prefix 10.1.1.0/32 for anycast address 10.1.1.0 on node A has BFR-ID 101, which may be different from the BFR-ID such as 102 that the prefix 10.1.1.0/32 for the anycast address 10.1.1.0 on node B has, node C sets the bit for BFR-ID 101 to 0 (zero) and sets the bit for BFR-ID 102 to 1 (one) in the packet with the bitstring before sending the packet to the loop free alternative (LFA).

[0069] In an embodiment, assume LFA is node B and node C is not directly connected to the LFA (or node B), then node C sends the packet with 10.1.1.0 as destination to the LFA through a tunnel to the LFA. For example, node C sends the packet to node B through an IP tunnel or segment routing (SR) tunnel from node C to node B (as shown by an arrow 410 in FIG. 4B). In an embodiment, when the LFA such as node B receives the packet with bitstring having a bit of value 1 for 10.1.1.0/32, the LFA forwards the packet according to its BIFT. In an embodiment, assume LFA is node D and node C is directly connected to the LFA, node C sends the packet to the LFA without using any tunnel. In an embodiment, when the LFA such as node D receives the packet with bitstring having a bit of value 1 for 10.1.1.0/32 as destination, the LFA sends the packet towards node B instead of node A since the LFA has a shorter path through node B. Node B then sends the packet to the CE receiver 120.

[0070] FIG. 5A is a schematic diagram of an SRv6 topology 500A including an SRv6 domain 502 according to an embodiment of the disclosure. As shown in FIG. 5A, an Srv6 domain 502 comprises an area 504 with an implementation of configuration and distribution of fast egress protection information. The network nodes and links depicted in FIG. 5A are similar to the network nodes and links depicted in FIG. 1. For the sake of brevity, a detailed description of these elements is not repeated herein. In an embodiment, the network node 104 (or node A) may be referred to herein as a primary egress node (or simply an egress node), and the network node 106 (or node B) may be referred to herein as a backup egress node for the primary egress node 104. The network nodes 104 and 106 are configured with an IP anycast address of 2001:db8:a136::/128. In an embodiment, when a neighboring node such as node 108 (or node C) receives a packet with the IP anycast address as its destination, the node 108 sends the packet toward the closer node with the IP anycast address. Once the node C receives the packet, node C has at least two paths from the node C to the destination (i.e., the IP anycast address). In an

embodiment, a first path is from node C to node A and a second path is from node C to node B. The SRv6 routing algorithm selects the path from the at least two paths based on which is the nearest according to cost measure. Assume that the cost from node B to the address configured makes the path from a neighboring node (i.e. node C) to the address via A shorter than the path from the neighboring node to the address via B.

[0071] In an embodiment, when the IP anycast address is configured on nodes A and B, the IGP's running on A and B distribute the IP address prefix 2001:db8:a136::/128 for the IP anycast address configured on nodes A and B. Every node in the area 504 receives the prefix. Thus, the egress protection information about the prefix distributed in the area 504 is used by the neighbors of the node A for fast protection against the failure of node A. The configuration and distribution are reused for fast protection of egress node A (or B).

[0072] A mechanism of determination of fast egress protection in the SRV6 domain 502 is described as follows. In an embodiment, the primary egress node 104 (or node A) and the backup egress node 106 (or node B) are configured with an IP anycast address of 2001:db8:a136::/128. After the configuration, node A and node B distributes the egress protection information about the backup egress to the direct neighbors of the node A (as shown by an arrow 506 in FIG. 5A). In an embodiment, the egress protection information is the IP anycast address distributed by node A and node B. Then, Node C determines whether it provides fast egress protection for node A based on the received egress protection information. Node C provides fast protection for node A based on pre-defined conditions. The pre-defined conditions are: 1) Node C connects to node A, 2) Node C can fast detects the failure of node A using a method such as BFD, 3) Node C has two paths to the IP address prefix corresponding to the same IP anycast address configured on node A and node B (e.g., node C has two paths to prefix 2001:db8:a136::/128 corresponding to IP anycast address configured on nodes A and B as shown in FIG. 5A), and 4) The path from node C to the IP address prefix 2001:db8:a136::/128 through node A is shorter than the path from node C to the IP address prefix through another node B. For example, the path from node C to the prefix 2001:db8:a136::/128 through node A has less cost than the path from node C to the prefix through node B. Thus, node C determines that it provides the fast protection for egress node A when the four conditions are met.

[0073] FIG. 5B is a schematic diagram of fast egress protection in a SRv6 network 500B according to an embodiment of the disclosure. In an embodiment, after determining to protect the primary egress node A with the backup egress node B of, node C provides fast protection against the failure of the egress node as follows. In an embodiment, in normal operations (i.e., node A is functioning normally), after receiving a packet with IP anycast address

2001:db8:a136::/128 as destination, node C sends node A the packet according to its FIB (as shown by an arrow 508 in FIG. 5B). In an embodiment, when the node A is functioning abnormally or has failed, node C detects the failure of node A using a fast failure detection method such as a bi-directional forwarding detection (BFD) method. After detecting the failure of node A, node C sends the packet with 2001:db8:a136::/128 as destination to a loop free alternative (LFA). In an embodiment, the LFA comprises a basic LFA or a topology independent (TI) LFA. The LFA node is far enough away from 2001:db8:a136::/128 through node A such that the path from the LFA to 2001:db8:a136::/128 through node B is shorter than the path from the LFA to 2001:db8:a136::/128 through node A (or the path from the LFA to B is shorter than the path from the LFA to node A).

[0074] In an embodiment, assume LFA is node D and node C is directly connected to the LFA, node C sends the packet to the LFA without using any tunnel (as shown by an arrow 510 in FIG. 5A). In an embodiment, when the LFA such as node D receives the packet with 2001:db8:a136::/128 as destination, the LFA sends the packet towards node B instead of node A since the LFA has a shorter path to 2001:db8:a136::/128 through node B. Node B then sends the packet to the CE receiver 120.

[0075] FIG. 6 is a method 600 implemented by a network node (e.g., network node 108) neighboring a primary egress node in a domain according to an embodiment of the disclosure. The method may be performed by the network node.

[0076] In block 602, the network node receives a piece of egress protection information that identifies a backup egress node for the primary egress node. In an embodiment, the TLV structure is received from a primary egress node. In an embodiment, the domain comprises a Border Gateway Protocol (BGP) domain, an interior gateway protocol (IGP) domain, a bit index explicit replication (BIER) domain, or a segment routing version 6 (SRv6) domain. In an embodiment, the egress protection information is received from the primary egress node and the backup egress node. In an embodiment, a network operator or tools configures an internet protocol (IP) anycast address on the primary egress node and the backup egress node; and the network node receives the egress protection information that comprises an IP address prefix for the IP anycast address from the primary egress node and the backup egress node.

[0077] In block 604, the network node make a determination, in response to receiving the egress protection information, to protect the primary egress node with the backup egress node when the primary egress node fails.

[0078] In block 606, the network node detects whether the primary egress node has failed when forwarding a packet to the primary egress node.

[0079] In block 608, the network node forwards, in response to the determination and in response to detecting the primary egress node has failed, the packet to the backup egress node to transmit the packet to a customer edge (CE) receiver. In an embodiment, the network node forwards, in response to detecting the primary egress node has not failed, the packet with the IP anycast address to the primary egress node according to forwarding information base (FIB) or bit index forwarding table (BIFT).

[0080] In an embodiment, when the domain is the BGP domain, the network node further identifies a remote node using information in a routing information base (RIB); and forward, after detecting the primary egress node has failed, the packet with IP anycast address to the remote node through a tunnel such as an IP tunnel. In an embodiment, when the domain is the IGP domain, the network node further identifies, in response to the determination, an intermediate node that is a loop free alternative (LFA) to the backup egress node; and after detecting the primary egress node has failed, either forwarding the packet with IP anycast address to the intermediate node when the intermediate node is directly connected to the network node, or forwarding the packet with IP anycast address, through a tunnel, to the intermediate node when the intermediate node is not directly connected to the network node. In an embodiment, when the domain is the BIER domain, the network node further receives the egress protection information that comprises the IP address prefix for the IP anycast address and a Bit-Forwarding Router Identifier (BFR-ID) for the IP address prefix; identifies, in response to the receiving, an intermediate node that is a loop free alternative (LFA) to the backup egress node; and after detecting the primary egress node has failed, either forward the packet with the BFR-ID to the intermediate node when the intermediate node is directly connected to the network node, or forward the packet with IP anycast address and the BFR-ID, through a tunnel, to the intermediate node when the intermediate node is not directly connected to the network node. In an embodiment, when the domain is the SRv6 domain, the network node further identifies, in response to the determination, an intermediate node that is a loop free alternative (LFA) to the backup egress node; and forward, after detecting the primary egress node has failed, the packet with the IP anycast address to the intermediate node. In an embodiment, the LFA comprises a basic LFA or a topology independent (TI) LFA.

[0081] FIG. 7 is a schematic diagram of a network apparatus 700 (e.g., a network node, a destination node, a neighbor node, etc.). The network apparatus 700 is suitable for implementing the disclosed embodiments as described herein. The network apparatus 700 comprises ingress ports/ingress means 710 and receiver units (Rx)/receiving means 720 for receiving data; a processor, logic unit, or central processing unit (CPU)/processing means 730 to process the data;

transmitter units (Tx)/transmitting means 740 and egress ports/egress means 750 for transmitting the data; and a memory/memory means 760 for storing the data. The network apparatus 700 may also comprise optical-to-electrical (OE) components and electrical-to-optical (EO) components coupled to the ingress ports/ingress means 710, the receiver units/receiving means 720, the transmitter units/transmitting means 740, and the egress ports/egress means 750 for egress or ingress of optical or electrical signals.

[0082] The processor/processing means 730 is implemented by hardware and software. The processor/processing means 730 may be implemented as one or more CPU chips, cores (e.g., as a multi-core processor), field-programmable gate arrays (FPGAs), application specific integrated circuits (ASICs), and digital signal processors (DSPs). The processor/processing means 730 is in communication with the ingress ports/ingress means 710, receiver units/receiving means 720, transmitter units/transmitting means 740, egress ports/egress means 750, and memory/memory means 760. The processor/processing means 730 comprises a fast egress protection module 770. The fast egress protection module 770 is able to implement the methods disclosed herein. The inclusion of the fast egress protection module 770 therefore provides a substantial improvement to the functionality of the network apparatus 700 and effects a transformation of the network apparatus 700 to a different state. Alternatively, the fast egress protection module 770 is implemented as instructions stored in the memory/memory means 760 and executed by the processor/processing means 730.

[0083] The network apparatus 700 may also include input and/or output (I/O) or devices/I/O means 780 for communicating data to and from a user. The I/O devices or I/O means 780 may include output devices such as a display for displaying video data, speakers for outputting audio data, etc. The I/O devices or I/O means 780 may also include input devices, such as a keyboard, mouse, trackball, etc., and/or corresponding interfaces for interacting with such output devices.

[0084] The memory/memory means 760 comprises one or more disks, tape drives, and solid-state drives and may be used as an over-flow data storage device, to store programs when such programs are selected for execution, and to store instructions and data that are read during program execution. The memory/memory means 760 may be volatile and/or non-volatile and may be read-only memory (ROM), random access memory (RAM), ternary content-addressable memory (TCAM), and/or static random-access memory (SRAM).

[0085] While several embodiments have been provided in the present disclosure, it may be understood that the disclosed systems and methods might be embodied in many other specific forms without departing from the spirit or scope of the present disclosure. The present examples are to be considered as illustrative and not restrictive, and the intention is not to be limited to the

details given herein. For example, the various elements or components may be combined or integrated in another system or certain features may be omitted, or not implemented.

[0086] In addition, techniques, systems, subsystems, and methods described and illustrated in the various embodiments as discrete or separate may be combined or integrated with other systems, components, techniques, or methods without departing from the scope of the present disclosure. Other examples of changes, substitutions, and alterations are ascertainable by one skilled in the art and may be made without departing from the spirit and scope disclosed herein.

CLAIMS

What is claimed is:

1. A method implemented by a network node neighboring a primary egress node in a domain, comprising:
 - receiving a piece of egress protection information that identifies a backup egress node for the primary egress node;
 - making a determination, in response to receiving the egress protection information, to protect the primary egress node with the backup egress node when the primary egress node fails;
 - detecting whether the primary egress node has failed when forwarding a packet to the primary egress node; and
 - forwarding, in response to the determination and in response to detecting the primary egress node has failed, the packet to the backup egress node to transmit the packet to a customer edge (CE) receiver.
2. The method of claim 1, wherein the egress protection information is received from the backup egress node or from both the primary egress node and the backup egress node.
3. The method of any of claims 1-2, further comprising detecting whether the primary egress node has failed via a bi-directional forwarding detection (BFD) session between the network node and the primary egress node.
4. The method of any of claims 1-3, wherein the domain is a Border Gateway Protocol (BGP) domain, an interior gateway protocol (IGP) domain, a bit index explicit replication (BIER) domain, or a segment routing version 6 (SRv6) domain, wherein the method further comprises receiving the egress protection information that comprises an IP address prefix for an IP anycast address from the primary egress node and the backup egress node, and wherein the IP anycast address is configured on the primary egress node and the backup egress node by a network operator or tools.
5. The method of any of claims 1-4, wherein when the domain is the BGP domain, the method further comprises:
 - identifying a remote node using information in a routing information base (RIB); and
 - forwarding, after detecting the primary egress node has failed, the packet with IP anycast address to the remote node through a tunnel.

6. The method of any of claims 1-4, wherein when the domain is the IGP domain, the method further comprises:

identifying, in response to the determination, an intermediate node that is a loop free alternative (LFA) to the backup egress node; and

after detecting the primary egress node has failed, either forwarding the packet with IP anycast address to the intermediate node when the intermediate node is directly connected to the network node, or forwarding the packet with IP anycast address, through a tunnel, to the intermediate node when the intermediate node is not directly connected to the network node.

7. The method of any of claims 1-4, wherein when the domain is the BIER domain, the method further comprises:

receiving the egress protection information that comprises the IP address prefix for the IP anycast address and a Bit-Forwarding Router Identifier (BFR-ID) for the IP address prefix;

identifying, in response to the receiving, an intermediate node that is a loop free alternative (LFA) to the backup egress node; and

after detecting the primary egress node has failed, either forwarding the packet with the BFR-ID to the intermediate node when the intermediate node is directly connected to the network node, or forwarding the packet with the IP anycast address and the BFR-ID, through a tunnel, to the intermediate node when the intermediate node is not directly connected to the network node.

8. The method of any of claims 1-4, wherein when the domain is the SRv6 domain, the method further comprises:

identifying, in response to the determination, an intermediate node that is a loop free alternative (LFA) to the backup egress node; and

forwarding, after detecting the primary egress node has failed, the packet with the IP anycast address to the intermediate node.

9. The method of any of claims 1-8, wherein the LFA comprises a basic LFA or a topology independent (TI) LFA.

10. The method of any of claims 1-9, further comprising forwarding, in response to detecting the primary egress node has not failed, the packet with the IP anycast address to the primary egress node according to forwarding information base (FIB) or bit index forwarding table (BIFT).

11. A network node neighboring a primary egress node in a domain, comprising:
 - a memory storing instructions; and
 - one or more processors coupled to the memory, wherein the one or more processors are configured to execute the instructions to cause the network node to:
 - receive a piece of egress protection information that identifies a backup egress node for the primary egress node;
 - make a determination, in response to receiving the egress protection information, to protect the primary egress node with the backup egress node when the primary egress node fails;
 - detect whether the primary egress node has failed when forwarding a packet to the primary egress node; and
 - forward, in response to the determination and in response to detecting the primary egress node has failed, the packet to the backup egress node to transmit the packet to a customer edge (CE) receiver.
12. The network node of claim 11, wherein the egress protection information is received from the backup egress node or from both the primary egress node and the backup egress node.
13. The network node of any of claims 11-12, wherein the one or more processors are configured to execute the instructions to further cause the network node to detect whether the primary egress node has failed via a bi-directional forwarding detection (BFD) session between the network node and the primary egress node.
14. The network node of any of claims 11-13, wherein the domain is a Border Gateway Protocol (BGP) domain, an interior gateway protocol (IGP) domain, a bit index explicit replication (BIER) domain, or a segment routing version 6 (SRv6) domain, wherein the one or more processors are configured to execute the instructions to further cause the network node to receive the egress protection information that comprises an IP address prefix for an IP anycast address from the primary egress node and the backup egress node, and wherein the IP anycast address is configured on the primary egress node and the backup egress node by a network operator or tools.

15. The network node of any of claims 11-14, wherein when the domain is the BGP domain, the one or more processors are configured to execute the instructions to further cause the network node to:

- identify a remote node using information in a routing information base (RIB); and
- forward, after detecting the primary egress node has failed, the packet with IP anycast address to the remote node through a tunnel.

16. The network node of any of claims 11-14, wherein when the domain is the IGP domain, the one or more processors are configured to execute the instructions to further cause the network node to:

- identify, in response to the determination, an intermediate node that is a loop free alternative (LFA) to the backup egress node; and
- after detecting the primary egress node has failed, either forward the packet with IP anycast address to the intermediate node when the intermediate node is directly connected to the network node, or forward the packet with IP anycast address, through a tunnel, to the intermediate node when the intermediate node is not directly connected to the network node.

17. The network node of any of claims 11-14, wherein when the domain is the BIER domain, the one or more processors are configured to execute the instructions to further cause the network node to:

- receive the egress protection information that comprises the IP address prefix for the IP anycast address and a Bit-Forwarding Router Identifier (BFR-ID) for the IP address prefix;
- identify, in response to the receiving, an intermediate node that is a loop free alternative (LFA) to the backup egress node; and
- after detecting the primary egress node has failed, either forward the packet with the BFR-ID to the intermediate node when the intermediate node is directly connected to the network node, or forward the packet with the IP anycast address and the BFR-ID, through a tunnel, to the intermediate node when the intermediate node is not directly connected to the network node.

18. The network node of any of claims 11-14, wherein when the domain is the SRv6 domain, the one or more processors are configured to execute the instructions to further cause the network node to:

- identify, in response to the determination, an intermediate node that is a loop free alternative (LFA) to the backup egress node; and

forward, after detecting the primary egress node has failed, the packet with the IP anycast address to the intermediate node.

19. The method of any of claims 11-18, wherein the LFA comprises a basic LFA or a topology independent (TI) LFA.

20. The network node of any of claims 11-19, wherein the one or more processors are configured to execute the instructions to further cause the network node to forward, in response to detecting the primary egress node has not failed, the packet with the IP anycast address to the primary egress node according to forwarding information base (FIB) or bit index forwarding table (BIFT).

21. A non-transitory computer readable medium comprising a computer program product for use by a network node, the computer program product comprising computer executable instructions stored on the non-transitory computer readable medium that, when executed by one or more processors, cause the network node to execute the method of any of claims 1-10.

22. A network node neighboring a primary egress node in a domain, comprising:
receiving means configured to receive a piece of egress protection information that identifies a backup egress node for the primary egress node;

processing means configured to:

make a determination, in response to receiving the egress protection information, to protect the primary egress node with the backup egress node when the primary egress node fails; and

detect whether the primary egress node has failed when forwarding a packet to the primary egress node; and

transmitting means configured to forward, in response to the determination and in response to detecting the primary egress node has failed, the packet to the backup egress node to transmit the packet to a customer edge (CE) receiver.




FIG. 1




FIG. 2A




FIG. 2B




FIG. 3A




FIG. 3B




FIG. 4A




FIG. 4B




FIG. 5A




FIG. 5B




FIG. 6




FIG. 7