

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
25 January 2007 (25.01.2007)


PCT

(10) International Publication Number
WO 2007/010354 A1

- (51) International Patent Classification:
H04L 12/56 (2006.01)
- (21) International Application Number:
PCT/IB2006/001950
- (22) International Filing Date: 14 July 2006 (14.07.2006)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
11/182,967 15 July 2005 (15.07.2005) US
- (71) Applicant (for all designated States except US): **TELEFONAKTIEBOLAGET LM ERICSSON (PUBL)** [SE/SE]; Telefonplan, S-164 83 Stockholm (SE).
- (72) Inventors: **SIVAKUMAR, Kaarthik**; 2347 Summerville Drive, #102, Santa Rosa, CA 95404 (US). **BHOGAVILLI, Suresh**; 113 Bates Avenue, Gaithersburg, MD 20877 (US). **SHAH, Samir**; 4470 Openmeadow Court, San Jose, CA 95129 (US). **SHERWIN, Christopher**; 9905 Lambertina Lane, Rockville, MD 20850 (US).
- (74) Agents: **SMITH, Steven, W.** et al.; Ericsson Inc., 6300 Legacy, MS EVR 1-C-11, Plano, TX 75024 (US).
- (81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LV, LY, MA, MD, MG, MK, MN, MW, MX, MZ, NA, NG, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, SY, 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, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IS, IT, LT, LU, LV, MC, NL, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

[Continued on next page]

(54) Title: ROUTER AND METHOD FOR PROTOCOL PROCESS MIGRATION



(57) Abstract: A router and a method for migrating routing protocol processes or Virtual Routers (VRs) from one Route Processor (RP) to another using graceful restart procedures for maintaining packet flow to the router and assisting the router to obtain restart information. The router includes first and second RPs and a forwarding engine for forwarding packets to neighbor routers in the network. When the routing protocol process is terminated on the first RP, the neighbor routers detect a router failure and initiate the graceful restart procedures. A filter rule in the forwarding engine causes it to forward packets to the new routing protocol process on the second RP. The new process learns the network topology from the neighbor routers, and the migration is completed without packet loss and without requiring complex changes to the network protocol stack.

WO 2007/010354 A1



Published:

- with international search report
- before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

ROUTER AND METHOD FOR PROTOCOL PROCESS MIGRATION

BACKGROUND

5 The present invention relates to communication systems. More particularly, and not by way of limitation, the present invention is directed to a router and a method in an Internet Protocol (IP)-based network for migrating protocol processes from one route processor to another route processor using a graceful restart procedure.

10 In state-of-the-art IP routers, different processors are implemented to handle control functions and packet-forwarding functions. When routing and signaling protocols running on the control processor fail, forwarding of payload traffic is interrupted even though all the information required to perform such forwarding is available to the packet-forwarding processor. This interruption
15 occurs because neighbor routers detect the failure of the routing and signaling protocols and assume that the entire router has failed. Consequently, the neighbor routers compute alternate paths bypassing the "failed" router. During this time, called routing convergence time, there is a potential for traffic loss.

 To address this problem, the Internet Engineering Task Force (IETF) has
20 standardized a set of extensions to the routing and signaling protocols to gracefully handle the restart of a failed protocol process on a neighbor router. When these extensions are implemented and a router's control software must be restarted, the router's neighbors continue to use it for forwarding traffic. Neighbors also help the restarted router software relearn the state that was known
25 prior to the failure.

 It is also known for IP network operators to build, manage, and provision virtual private networks (VPNs) on top of their existing infrastructure. These networks are typically used by enterprises that need interconnectivity between geographically distributed sites. Using a private network is also appealing
30 because it offers a level of protection from intruders. Telecom network operators also use VPNs to provide traffic separation between various classes of telecom traffic. This is useful for providing different quality of service (QoS) and security services to these traffic classes.

-2-

With the growth in the size and speed of IP networks, routers or packet processing nodes must be scalable. Otherwise as the demands for processing power increase, or as more customer VPNs are configured, more and more routers must be deployed with added operational complexity and expenditure. To handle the resiliency and scalability needs of IP and telecom networks, routers or packet processing nodes are increasingly designed using a cluster of processors. To address scalability needs, middleware known as cluster management software distributes the processing load across multiple processors. The increase in processing demands is addressed by adding more processors to the cluster and migrating processes to the new processors. Any state needed by the process is also migrated to its new location by the cluster management software. Although effective, the use of processor clusters increases the complexity and cost of implementation of the routers.

An alternative that offers resiliency against node failures is to maintain one or more control processors (usually known as Route Processors or RPs) in hot-standby state to backup a primary RP. The protocol state is replicated between the primary and backup RPs. If the primary RP fails, the backup RP takes over and masks the failure from the router's neighbors. The complexity in this approach is in synchronizing state information between the primary and backup RPs. For protocols like Border Gateway Protocol (BGP) and Label Distribution Protocol (LDP) that run over the Transmission Control Protocol (TCP), TCP session state (such as sequence numbers, congestion window parameters, and the like) must also be replicated.

Thus, what is needed in the art is a more efficient way to handle the resiliency and scalability needs of IP and telecom networks that overcomes the deficiencies of conventional systems and methods. The present invention provides such a router and method.

SUMMARY

The present invention provides a router and a method of migrating routing protocol processes from one Route Processor (RP) to another using graceful restart procedures. The RPs may be in the same router or different routers. A

-3-

similar mechanism is utilized to migrate a Virtual Router (VR) and all its associated processes from one RP to another.

Thus, in one aspect, the present invention is directed to a method in a router in an Internet Protocol (IP)-based network for migrating routing protocol
5 processes from a first route processor to a second route processor. The router includes a forwarding engine for forwarding packets to neighbor routers in the network, and the neighbor routers include a graceful restart procedure for maintaining packet flow to the router and assisting the router if the router fails. The method includes terminating the routing protocol processes on the first route
10 processor, thereby indicating to the neighbor routers that the router has failed, and causing the neighbor routers to start the graceful restart procedure; and restarting the routing protocol processes on the second route processor. The method also includes diverting by the forwarding engine, packets destined for the routing protocol processes on the first route processor to the restarted routing protocol
15 processes on the second route processor; and receiving by the restarted routing protocol processes on the second route processor, information from the neighbor routers regarding the network's topology. In this aspect, the routing protocol processes may be, for example, Border Gateway Protocol (BGP) processes.

In another aspect, the method may include terminating the routing protocol
20 processes on the first route processor; and adding a filter rule to the forwarding engine after terminating the routing protocol processes on the first route processor. The filter rule causes the forwarding engine to divert packets destined for the routing protocol processes on the first route processor to the restarted routing protocol processes on the second route processor. The method also
25 includes restarting the routing protocol processes on the second route processor; sending a message to each of the neighbor routers requesting the neighbor routers to start the graceful restart procedure; and receiving by the restarted routing protocol processes on the second route processor, information from the neighbor routers regarding the network's topology. In this aspect, the routing
30 protocol processes may be, for example, Open Shortest Path First (OSPF) processes or Intermediate System to Intermediate System Protocol (IS-IS) processes.

In yet another aspect, when the routing protocol processes do not have a defined graceful restart procedure, said method may include terminating the routing protocol processes on the first route processor; adding the filter rule to the forwarding engine; restarting the routing protocol processes on the second route processor; and receiving periodic refresh messages from each of the neighbor routers regarding a protocol state. The refresh messages include information on all routes within the network's routing domain. The router then sends a message to each of the neighbor routers advertising that the router is operational. In this aspect, the routing protocol processes may be, for example, Routing Information Protocol (RIP).

In yet another aspect, the present invention is directed to a method in a router in an IP-based network for migrating a Virtual Router (VR) from a first route processor to a second route processor. The router includes a forwarding engine for forwarding packets to neighbor routers in the network; the VR includes routing and signaling protocol processes and a Route Table Manager (RTM) that computes a forwarding information base utilized by the forwarding engine; and the neighbor routers include a graceful restart procedure for maintaining packet flow to the router and assisting the router if the router fails. The method includes terminating the VR's routing and signaling protocol processes on the first route processor, thereby indicating to the neighbor routers that the router has failed, and causing the neighbor routers to start the graceful restart procedure; restarting the VR's routing and signaling protocol processes on the second route processor; and diverting by the forwarding engine, packets destined for the VR's routing and signaling protocol processes on the first route processor to the VR's restarted routing and signaling protocol processes on the second route processor. The method also includes receiving by the VR's restarted routing and signaling protocol processes on the second route processor, information from the neighbor routers regarding the network's topology; restarting the VR's RTM on the second route processor; and receiving by the VR's restarted RTM, information from the VR's restarted routing and signaling protocol processes regarding the network's topology, thereby establishing a routing state. The method also includes receiving by the VR's restarted RTM, information from the forwarding engine regarding the

-5-

current forwarding information base, thereby establishing a forwarding information state; and synchronizing by the RTM, the forwarding information state and the routing state.

In yet another aspect, the present invention is directed to a router in an IP-based network in which neighbor routers include a graceful restart procedure for maintaining packet flow to the router and assisting the router if the router fails. The router includes a forwarding engine for forwarding packets to the neighbor routers in the network; a first route processor connected to the forwarding engine for running routing protocol processes; a second route processor connected to the forwarding engine for running routing protocol processes; and means for migrating a routing protocol process from the first route processor to the second route processor, said migrating means including means for preventing packet loss while migrating the routing protocol process, without requiring changes to the network protocol stack.

15

BRIEF DESCRIPTION OF THE SEVERAL VIEWS OF THE DRAWINGS

In the following section, the invention will be described with reference to exemplary embodiments illustrated in the figures, in which:

FIG. 1 is a simplified block diagram of a router in an embodiment of the present invention;

FIG. 2 is a flow chart illustrating the steps of an embodiment of the method of the present invention when migrating a BGP process from a first Route Processor to another;

FIG. 3 is a flow chart illustrating the steps of an embodiment of the method of the present invention when migrating an OSPF or IS-IS process from a first Route Processor to another;

FIG. 4 is a flow chart illustrating the steps of an embodiment of the method of the present invention when migrating an RIP process from a first Route Processor to another; and

FIG. 5 is a flow chart illustrating the steps of an embodiment of the method of the present invention when migrating a Virtual Router (VR) from a first Route Processor to another.

DETAILED DESCRIPTION

The present invention provides a router and a method for migrating routing protocol processes from one Route Processor (RP) to another using graceful restart procedures. The RPs may be in the same router or different routers. A similar mechanism is utilized to migrate a Virtual Router (VR) and all its associated processes from one RP to another. Routing protocols, such as the Border Gateway Protocol (BGP), Open Shortest Path First (OSPF), Intermediate System to Intermediate System Protocol (IS-IS), and Routing Information Protocol (RIP), communicate with a Route Table Manager (RTM) using an Application Programming Interface (API) that hides the communication between these processes. In such a communication method, it is unknown to each process exactly where the other end of the communication channel is located. This provides flexibility by allowing routing protocols to reside on a first RP while the RTM resides on a different RP or, conversely, the routing protocols and RTM may all reside on the same RP.

FIG. 1 is a simplified block diagram of a router (Router-A) 10 in an embodiment of the present invention. Router-A includes one or more Route Processor (RP) boards 11a-11n that house the control plane of the router. The RPs may include routing protocols 12a-12n and RTMs 13a-13n. A plurality of forwarding engines 14 in the router perform the packet-forwarding functions of the router. The control plane uses interfaces on the forwarding engines to communicate with specific router neighbors and to set up the forwarding engines to provide packet forwarding.

In the embodiment described herein, the routing protocols 12 and RTM processes 13 run on one or more of the RP boards 11. If only one RP board is present, all the processes run on the single RP. If two or more RPs are present, the routing protocols and RTM may run on separate available RPs. Interprocess Communications (IPCs) between the routing protocols and the RTMs ensure that all the processes can talk to each other irrespective of the RP board on which they are running. Similarly one or more VRs may be instantiated on multiple RP boards depending on the processing requirements of the VR's processes.

BGP Graceful Restart for Process Migration

BGP is a path-vector-based routing protocol that is primarily used to set up routing connections across routing domains. It is a very popular and important protocol used in routers. BGP works by setting up a TCP socket connection
5 between two neighbor routers. The entire BGP connection is dependent on the continued existence of the TCP connection. Failure of the TCP connection in any way results in the loss of the BGP session and a further loss of forwarding capabilities, which leads to severe damage of the network functionality.

10 The Graceful Restart procedure for BGP requires that a loss of the TCP connection be a signal to the receiving router that the neighbor router has failed. The procedure requires the receiving router to continue to use the failed router for forwarding until told otherwise by the restarting router. This procedure ensures that if only the BGP process on a router dies, then traffic will continue to be
15 forwarded while the BGP process restarts itself, thus ensuring that the network traffic is not lost.

Migration of BGP processes requires that the TCP connection continue to be maintained during the time that the process moves from one RP to another. This is a very complex problem that is still being researched and there are very
20 few successful implementations of this technique. Existing techniques require that the entire TCP/IP network stack of the router be revamped to deal with a TCP connection migration.

FIG. 2 is a flow chart illustrating the steps of an embodiment of the method of the present invention when migrating a BGP process from a first Route
25 Processor (RP-1) 11a in Router-A 10 to a second Route Processor (RP-2) 11b. In the present invention, no TCP connection migration is necessary. Instead, when the BGP process is ready to migrate to RP-2, it terminates itself at step 21. At step 22, Router-A's neighbor routers detect the failure of Router-A. This causes the neighbor routers to enter the Graceful Restart state at step 23. At step 24, a
30 filter rule is added to the forwarding engines 14 to detect TCP packets destined for the terminated BGP process, and to divert the packets to a new BGP process on RP-2. At step 25, the BGP process then restarts on RP-2 and sets up peering

-8-

relationships with its neighbor routers. At step 26, the new BGP process learns the network topology with the help of the neighbor routers.

[0010] Thus, the invention's use of Graceful Restart for BGP process migration completely removes the need to do any complex manipulation of the TCP/IP
5 network stack. The invention may additionally be utilized on any stock, off-the-shelf, network stack. The invention therefore drastically reduces the complexity of implementing process migration for BGP processes. This reduces developer costs for the vendor and purchasing costs for the customer.

10 Migration of Protocols other than BGP

Protocols such as OSPF, IS-IS, and RIP do not run over TCP and hence there is no session state maintained inside the operating system running over the RP. However, the protocol processes do maintain protocol state, and this state needs to be relearned once the protocol process is restarted on the new RP. The
15 OSPF and IS-IS protocols have graceful restart mechanisms defined, but RIP does not.

FIG. 3 is a flow chart illustrating the steps of an embodiment of the method of the present invention when migrating an OSPF or IS-IS process from a first RP-1 to a second RP-2. When the protocol process is ready to migrate to RP-2, it is
20 terminated on RP-1 at step 31. At step 32, a filter rule is added to the forwarding engines to detect the protocol packets destined for the old OSPF/IS-IS process and to divert the packets to the new OSPF/IS-IS process on RP-2. At step 33, the protocol process is then restarted on RP-2. If the process is an OSPF process, the method moves to step 34 where the restarted OSPF process sends out
25 Opaque Link State Advertisements (LSAs) requesting graceful restart support from neighbor routers. If the process is an IS-IS process, the method moves instead to step 35 where the restarted IS-IS process sends a Restart Type, Length and Value (TLV) field in its hello messages, with the Restart Request (RR) flag set, thus indicating to its neighbors that it is restarting. At step 36, the new
30 OSPF/IS-IS process learns the network topology with the help of the neighbor routers. Thus, all routing adjacencies are brought up using Graceful Restart mechanisms without affecting traffic forwarding.

FIG. 4 is a flow chart illustrating the steps of an embodiment of the method of the present invention when migrating a RIP process from a first RP-1 to a second RP-2. The RIP protocol does not have any graceful restart mechanisms defined. However RIP is another distance-vector protocol, and it uses periodic refreshes of all the routes across the routing domain. A restarting RIP process on a new RP simply waits long enough to relearn its protocol state prior to sending out its own advertisements. That achieves process migration without affecting packet forwarding.

Thus, when the protocol process is ready to migrate to RP-2, it is terminated on RP-1 at step 41. At step 42, a filter rule is added to the forwarding engines to detect the protocol packets destined for the old RIP process and to divert the packets to the new RIP process on RP-2. At step 43, the RIP process is then restarted on RP-2. At step 44, the new RIP process receives periodic refreshes of all the routes across the routing domain. At step 45, the new RIP process sends advertisements to neighbor routers after the RIP process has relearned its protocol state.

Migration of VRs

FIG. 5 is a flow chart illustrating the steps of an embodiment of the method of the present invention when migrating a Virtual Router (VR) from a first RP-1 to a second RP-2. A VR typically consists of routing and signaling protocol processes and an RTM that computes the forwarding information base utilized by the forwarding engine. To migrate the VR, all these processes must be migrated to the new RP. The routing and signaling protocols of the VR are terminated and restarted on the new RP as described above. Migrating the RTM involves restarting it on the new RP and allowing it to relearn the routing state from the routing protocols. It learns the current forwarding information base from the forwarding engines, and synchronizes the two states.

Thus, when the VR is ready to migrate to RP-2, its routing and signaling protocol processes are terminated on RP-1 at step 51. At step 52, assuming the routing and signaling protocol processes have graceful restart mechanisms defined, Router-A's neighbor routers detect the failure of Router-A and enter the

Graceful Restart state at step 53. At step 54, a filter rule is added to the forwarding engines 14 to detect TCP packets destined for the VR's terminated routing and signaling protocol processes, and to divert the packets to new routing and signaling protocol processes on RP-2. At step 55, the RTM is restarted on
5 RP-2. At step 56, the routing and signaling protocol processes restart on RP-2 and set up peering relationships with neighbor routers. At step 57, the new routing and signaling protocol processes learn the network topology with the help of the neighbor routers. At step 58, the RTM learns the routing state from the routing and signaling protocol processes, and at step 59, the RTM learns the
10 current forwarding information base from the forwarding engines and synchronizes the forwarding information state with the routing state.

As will be recognized by those skilled in the art, the innovative concepts described in the present application can be modified and varied over a wide range of applications. Accordingly, the scope of patented subject matter should not be
15 limited to any of the specific exemplary teachings discussed above, but is instead defined by the following claims.

WHAT IS CLAIMED IS:

1. A method in a router in an Internet Protocol (IP)-based network for migrating routing protocol processes from a first route processor to a second route processor, wherein the router includes a forwarding engine for forwarding packets
5 to neighbor routers in the network, and the neighbor routers include a graceful restart procedure for maintaining packet flow to the router and assisting the router if the router fails, said method comprising:
 - terminating the routing protocol processes on the first route processor, thereby indicating to the neighbor routers that the router has failed, and causing
10 the neighbor routers to start the graceful restart procedure;
 - restarting the routing protocol processes on the second route processor;
 - diverting by the forwarding engine, packets destined for the routing protocol processes on the first route processor to the restarted routing protocol processes on the second route processor; and
15 receiving by the restarted routing protocol processes on the second route processor, information from the neighbor routers regarding the network's topology.
2. The method of claim 1, wherein the first route processor and the second route processor are in a single router.
20
3. The method of claim 1, wherein the first route processor and the second route processor are in different routers.
4. The method of claim 1, wherein the method is preceded by the step of
25 determining that a processing load, shared between multiple route processors, needs to be redistributed to increase router scalability.
5. The method of claim 1, wherein the step of diverting packets by the forwarding engine includes adding a filter rule to the forwarding engine after
30 terminating the routing protocol processes on the first route processor, the filter rule causing the forwarding engine to divert packets destined for the routing

-12-

protocol processes on the first route processor to the restarted routing protocol processes on the second route processor.

6. The method of claim 1, wherein the step of restarting the routing protocol processes on the second route processor includes setting up peering relationships with the neighbor routers.

7. The method of claim 1, wherein the routing protocol processes are Border Gateway Protocol (BGP) processes.

10

8. A method in a router in an Internet Protocol (IP)-based network for migrating routing protocol processes from a first route processor to a second route processor, wherein the router includes a forwarding engine for forwarding packets to neighbor routers in the network, and the neighbor routers include a graceful restart procedure for maintaining packet flow to the router and assisting the router if the router fails, said method comprising:

terminating the routing protocol processes on the first route processor;
adding a filter rule to the forwarding engine after terminating the routing protocol processes on the first route processor, the filter rule causing the forwarding engine to divert packets destined for the routing protocol processes on the first route processor to the restarted routing protocol processes on the second route processor;

restarting the routing protocol processes on the second route processor;
sending a message to each of the neighbor routers requesting the neighbor routers to start the graceful restart procedure; and

receiving by the restarted routing protocol processes on the second route processor, information from the neighbor routers regarding the network's topology.

9. The method of claim 8, wherein the routing protocol processes are Open Shortest Path First (OSPF) processes, and the step of sending a message to each of the neighbor routers includes sending an Opaque Link State Advertisement (LSA) to each of the neighbor routers.

30

10. The method of claim 8, wherein the routing protocol processes are Open Shortest Path First (OSPF) processes, and the step of sending a message to each of the neighbor routers includes sending the message to the neighbor routers before terminating the processes on the first route processor.

5

11. The method of claim 8, wherein the routing protocol processes are Intermediate System to Intermediate System Protocol (IS-IS) processes, and the step of sending a message to each of the neighbor routers includes sending a message with a Restart Type, Length and Value (TLV) field to each of the neighbor routers indicating that the router is restarting.

10

12. A method in a router in an Internet Protocol (IP)-based network for migrating routing protocol processes from a first route processor to a second route processor, wherein the router includes a forwarding engine for forwarding packets to neighbor routers in the network, and wherein the routing protocol processes do not have a defined graceful restart procedure, said method comprising:

15

terminating the routing protocol processes on the first route processor;

adding a filter rule to the forwarding engine after terminating the routing protocol processes on the first route processor, the filter rule causing the forwarding engine to divert packets destined for the routing protocol processes on the first route processor to the restarted routing protocol processes on the second route processor;

20

restarting the routing protocol processes on the second route processor;

receiving periodic refresh messages from each of the neighbor routers regarding a protocol state, said refresh messages including information on all routes within the network's routing domain; and

25

sending a message to each of the neighbor routers advertising that the router is operational.

30

13. The method of claim 12, wherein the routing protocol processes are Routing Information Protocol (RIP) processes.

-14-

14. A method in a router in an Internet Protocol (IP)-based network for migrating a Virtual Router (VR) from a first route processor to a second route processor, wherein the router includes a forwarding engine for forwarding packets to neighbor routers in the network, the VR includes routing and signaling protocol processes and a Route Table Manager (RTM) that computes a forwarding information base utilized by the forwarding engine, and the neighbor routers include a graceful restart procedure for maintaining packet flow to the router and assisting the router if the router fails, said method comprising:

5
10 terminating the VR's routing and signaling protocol processes on the first route processor, thereby indicating to the neighbor routers that the router has failed, and causing the neighbor routers to start the graceful restart procedure;

restarting the VR's routing and signaling protocol processes on the second route processor;

15 diverting by the forwarding engine, packets destined for the VR's routing and signaling protocol processes on the first route processor to the VR's restarted routing and signaling protocol processes on the second route processor;

receiving by the VR's restarted routing and signaling protocol processes on the second route processor, information from the neighbor routers regarding the network's topology;

20 restarting the VR's RTM on the second route processor;

receiving by the VR's restarted RTM, information from the VR's restarted routing and signaling protocol processes regarding the network's topology, thereby establishing a routing state;

25 receiving by the VR's restarted RTM, information from the forwarding engine regarding the current forwarding information base, thereby establishing a forwarding information state; and

synchronizing by the RTM, the forwarding information state and the routing state.

30 15. The method of claim 14, wherein the first route processor and the second route processor are in a single router.

16. The method of claim 14, wherein the first route processor and the second route processor are in different routers.
17. The method of claim 14, wherein the step of diverting packets by the forwarding engine includes adding a filter rule to the forwarding engine after terminating the VR's routing and signaling protocol processes on the first route processor, the filter rule causing the forwarding engine to divert packets destined for the VR's routing and signaling protocol processes on the first route processor to the VR's restarted routing and signaling protocol processes on the second route processor.
18. The method of claim 14, wherein the step of restarting the routing protocol processes on the second route processor includes setting up peering relationships with the neighbor routers.
19. A router in an Internet Protocol (IP)-based network in which neighbor routers include a graceful restart procedure for maintaining packet flow to the router and assisting the router if the router fails, said router comprising:
- a forwarding engine for forwarding packets to the neighbor routers in the network;
 - a first route processor connected to the forwarding engine for running routing protocol processes;
 - a second route processor connected to the forwarding engine for running routing protocol processes; and
- means for migrating a routing protocol process from the first route processor to the second route processor, said migrating means including means for preventing packet loss while migrating the routing protocol process, without requiring changes to the network protocol stack.
20. The router of claim 19, wherein the means for preventing packet loss while migrating the routing protocol process includes means for initiating a graceful

-16-

restart procedure defined for the routing protocol process to maintain packet flow destined for the router while the routing protocol process is migrated.

21. The router of claim 20, wherein the means for initiating a graceful restart
5 procedure includes means for terminating the routing protocol process on the first route processor, thereby indicating to the neighbor routers that the router has failed, and causing the neighbor routers to start the graceful restart procedure.

22. The router of claim 21, wherein the means for migrating the routing protocol
10 process includes restarting the routing protocol process on the second route processor.

23. The router of claim 22, wherein the means for preventing packet loss while
15 migrating the routing protocol process includes means for adding a filter rule to the forwarding engine, said filter rule causing the forwarding engine to divert packets destined for the routing protocol process on the first route processor to the restarted routing protocol process on the second route processor while migrating the routing protocol process.




FIG. 1

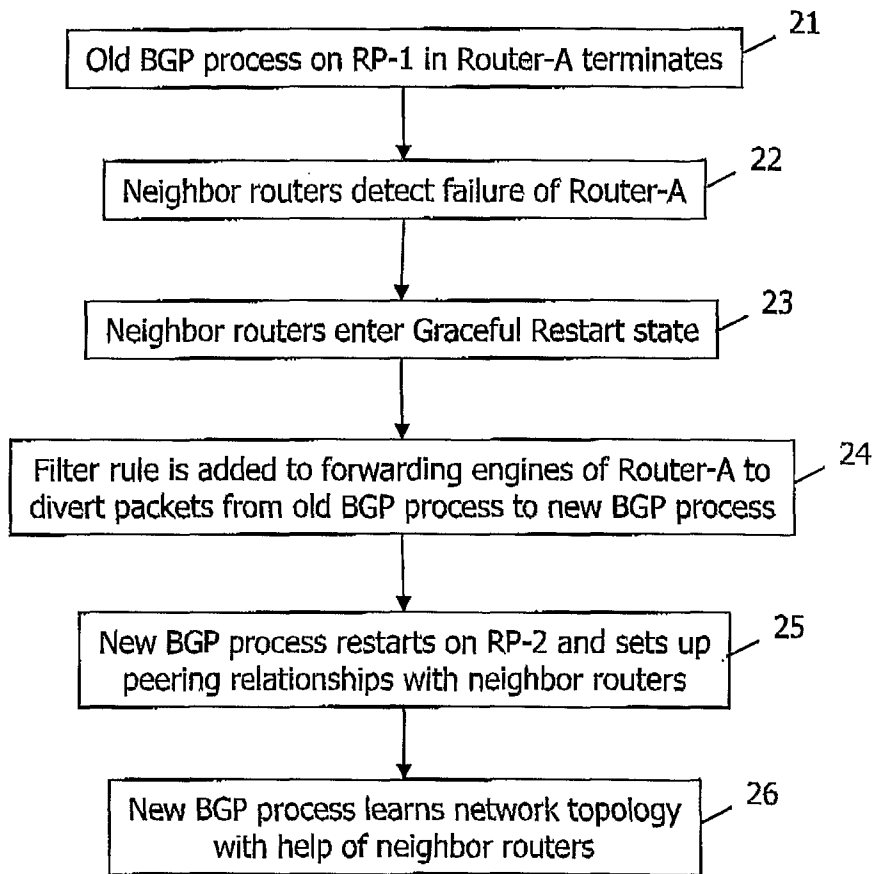


FIG. 2

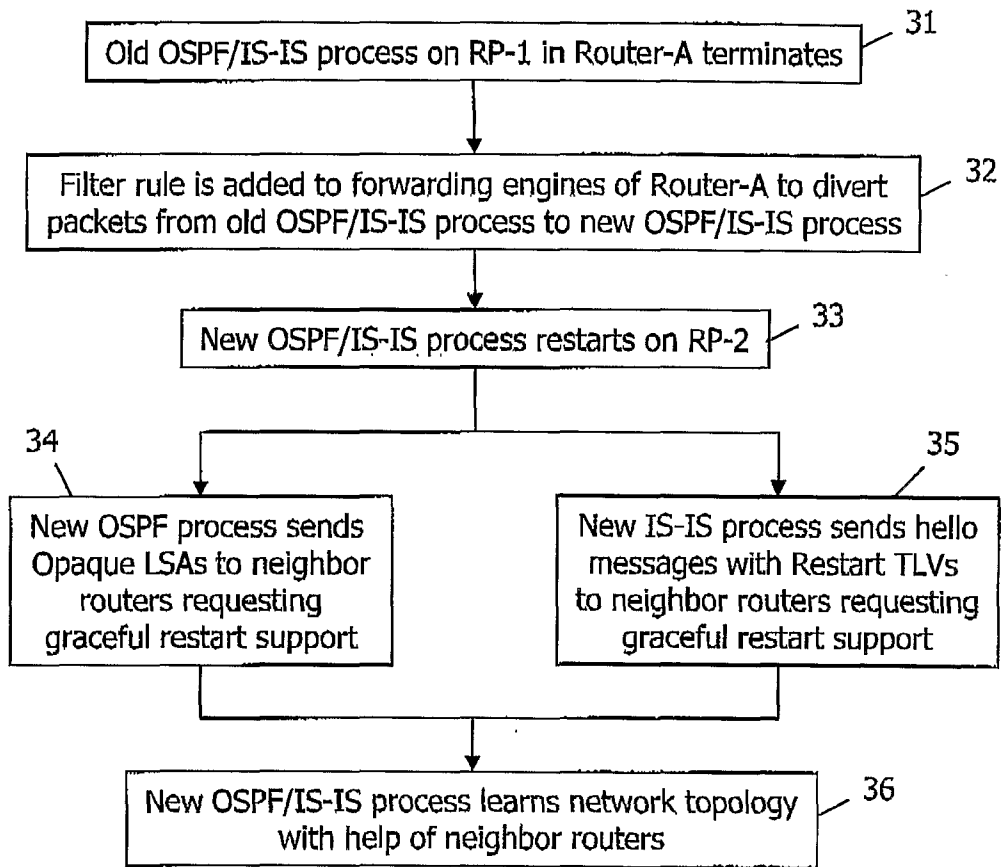


FIG. 3

4/5

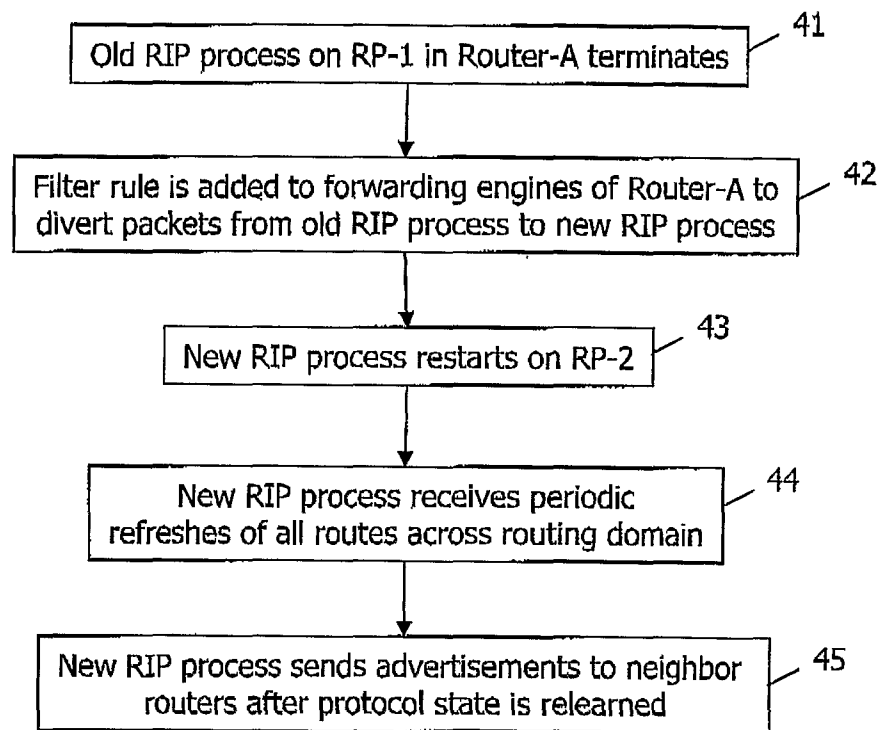


FIG. 4

5/5

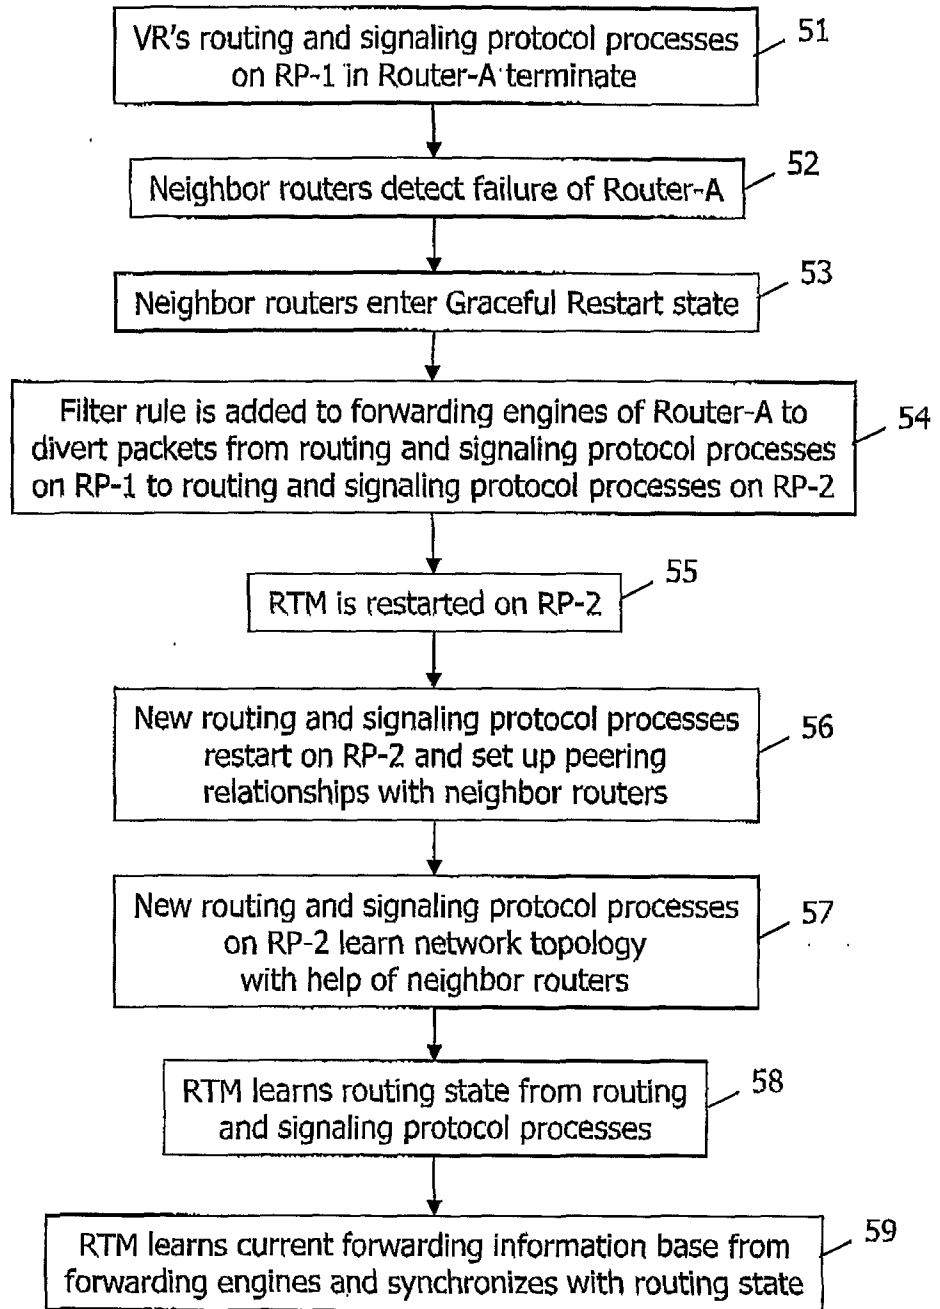


FIG. 5

INTERNATIONAL SEARCH REPORT

International application No

PCT/IB2006/001950

A. CLASSIFICATION OF SUBJECT MATTER
INV. H04L12/56

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)
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 practical, search terms used)

EPO-Internal, INSPEC

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2005/135233 A1 (MAY ROBERT [CA]) 23 June 2005 (2005-06-23) paragraph [0011] - paragraph [0076]	1-23
A	US 6 397 260 B1 (WILS JORIS JOHANNES MARIA [US] ET AL) 28 May 2002 (2002-05-28) column 2, line 29 - column 8, line 11	1-23
A	EP 1 006 702 A2 (NORTEL NETWORKS CORP [CA]) 7 June 2000 (2000-06-07) paragraph [0005] - paragraph [0021]	1-23
A	EP 1 309 135 A1 (CIT ALCATEL [FR]) 7 May 2003 (2003-05-07) paragraph [0001] - paragraph [0023]	1-23
A	US 2002/080752 A1 (JOHANSSON FREDRIK [SE] ET AL) 27 June 2002 (2002-06-27) the whole document	1-23

Further documents are listed in the continuation of Box C.

See patent family annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance
"E" earlier document but published on or after the international filing date
"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)
"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

"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
"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
"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.
"&" document member of the same patent family

Date of the actual completion of the international search

2 November 2006

Date of mailing of the international search report

13/11/2006

Name and mailing address of the ISA/

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

Ciurel, Cristian

INTERNATIONAL SEARCH REPORT

International application No
PCT/IB2006/001950

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2005135233	A1	23-06-2005	NONE
		AU 2004306913 A1	28-04-2005
		CA 2542645 A1	28-04-2005
		EP 1678894 A1	12-07-2006
		US 2005083953 A1	21-04-2005
		WO 2005039129 A1	28-04-2005
US 6397260	B1	28-05-2002	NONE
EP 1006702	A2	07-06-2000	CA 2290267 A1
			04-06-2000
EP 1309135	A1	07-05-2003	AT 290285 T
			15-03-2005
		DE 60109177 D1	07-04-2005
		DE 60109177 T2	28-07-2005
		US 2003093557 A1	15-05-2003
US 2002080752	A1	27-06-2002	NONE