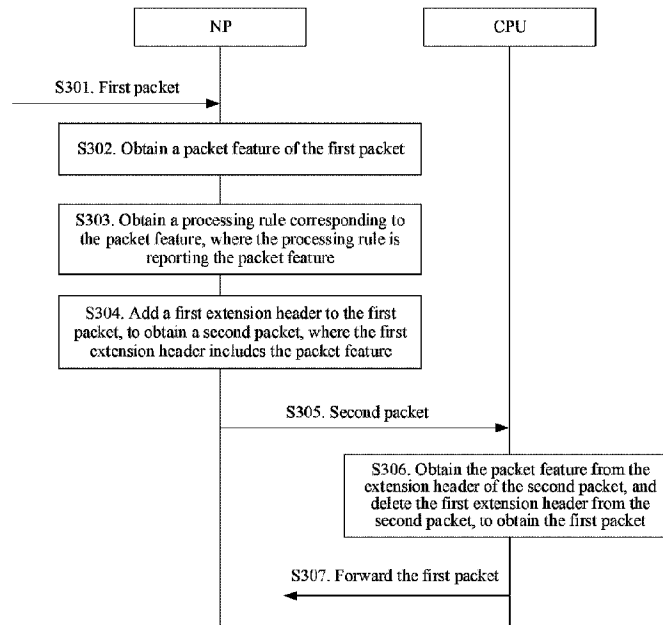




(22) Date de dépôt/Filing Date: 2020/04/17
(41) Mise à la disp. pub./Open to Public Insp.: 2020/10/19
(45) Date de délivrance/Issue Date: 2024/03/19
(30) Priorité/Priority: 2019/04/19 (CN201910320133.X)

(51) Cl.Int./Int.Cl. *H04L 49/9057* (2022.01),
H04L 45/74 (2022.01)
(72) Inventeurs/Inventors:
GUO, ZUN, CN;
TIAN, XIANGYUAN, CN
(73) Propriétaire/Owner:
HUAWEI TECHNOLOGIES CO., LTD., CN
(74) Agent: GOWLING WLG (CANADA) LLP

(54) Titre : PROCÉDE POUR IMPUTATION DES COMMANDES ET PÉRIPHÉRIQUE RESEAU
(54) Title: SERVICE PROCESSING METHOD AND NETWORK DEVICE



(57) **Abrégé/Abstract:**

This application provides a service processing method and a network device. The network device includes an NP and a CPU. The NP is configured to: receive a first packet; obtain a packet feature of the first packet; obtain a processing rule corresponding to the packet feature; process the first packet based on the processing rule, to obtain a second packet; and send the second packet to the CPU. The CPU is configured to: receive the second packet, and process the second packet. Before the CPU processes the first packet, the NP processes the first packet, and sends the second packet obtained after processing to the CPU, so that the CPU may directly process the second packet, to simplify a process in which the CPU processes the packet received by the CPU, thereby improving service processing efficiency of the network device.

ABSTRACT

This application provides a service processing method and a network device. The network device includes an NP and a CPU. The NP is configured to: receive a first packet; obtain a packet feature of the first packet; obtain a processing rule
5 corresponding to the packet feature; process the first packet based on the processing rule, to obtain a second packet; and send the second packet to the CPU. The CPU is configured to: receive the second packet, and process the second packet. Before the CPU processes the first packet, the NP processes the first packet, and sends the second packet obtained after processing to the CPU, so that the CPU may directly process the
10 second packet, to simplify a process in which the CPU processes the packet received by the CPU, thereby improving service processing efficiency of the network device.

SERVICE PROCESSING METHOD AND NETWORK DEVICE

TECHNICAL FIELD

[0001] This application relates to the field of communications technologies, and in
5 particular, to a service processing method and a network device.

BACKGROUND

[0002] With the popularity of the Internet, an edge router plays an important role in
a network system. The edge router is oriented to transfer a user from a local area
network to a wide area network. Nowadays, there is still a significant difference
10 between the local area network technology and the wide area network technology, and
the edge router needs to meet a variety of service requirements of users, for example,
from simple networking to a complex multimedia service and a (virtual private network,
VPN) service.

[0003] Currently, services related to the edge router are all processed by a central
15 processing unit (central processing unit, CPU) of the edge router. However, services
related to the edge router are of a wide variety, and forwarding and processing
capabilities of the CPU are difficult to improve. Consequently, service processing
efficiency of the edge router is low.

SUMMARY

20 [0004] This application provide a service processing method and a network device,
to simplify a process of processing a packet by a CPU in the network device, thereby
improving service processing efficiency of the network device.

[0005] According to a first aspect, this application provides a network device,

where the network device includes a network processor (network processor, NP) and a CPU. The NP is configured to: receive a first packet; obtain a packet feature of the first packet; obtain a processing rule corresponding to the packet feature; process the first packet based on the processing rule, to obtain a second packet; and send the second
5 packet to the CPU. The CPU is configured to: receive the second packet, and process the second packet.

[0006] In a possible design, the CPU is further configured to send the processing rule to the NP.

[0007] In a possible design, the processing rule is reporting the packet feature. The
10 NP is configured to: add a first extension header to the first packet, to obtain the second packet, where the first extension header includes the packet feature. The CPU is configured to: obtain the packet feature from the first extension header of the second packet; delete the first extension header from the second packet, to obtain the first packet; and forward the first packet based on the packet feature.

[0008] In a possible design, the processing rule is obtaining routing information. The NP is configured to: add a first extension header to the first packet, to obtain the second packet, where the first extension header includes the routing information. The CPU is configured to: obtain the routing information from the first extension header of the second packet; delete the first extension header from the second packet, to obtain
20 the first packet; and forward the first packet based on the routing information.

[0009] In a possible design, the processing rule is reporting service information. The NP is configured to: obtain a service configured on an interface that receives the first packet, and add a first extension header to the first packet, to obtain the second packet, where the first extension header includes information about the service. The
25 CPU is configured to: obtain the information about the service from the first extension header of the second packet; delete the first extension header from the second packet, to obtain the first packet; and perform the service on the first packet.

[0010] In a possible design, when obtaining the processing rule corresponding to the packet feature, the NP looks up in a rule table stored on the NP based on the packet
30 feature; obtains the processing rule from the rule table when the rule table includes the

processing rule; and when the rule table does not include the processing rule, sends the first packet or the packet feature to the CPU, and receives the processing rule returned by the CPU.

[0011] In a possible design, the NP further adds the processing rule to the rule table.

5 [0012] In a possible design, the CPU sends a third packet to the NP, where the third packet includes a second extension header, and the second extension header includes processing information and a layer 2 header.

[0013] The NP is configured to: obtain the processing information and the layer 2 header in the second extension header; delete the second extension header from the
10 third packet, to obtain a fourth packet; process the fourth packet based on the processing information; and add the layer 2 header to the fourth packet, to obtain a fifth packet.

[0014] According to a second aspect, this application provides a service processing method, applied to a network device, where the network device includes an NP and a CPU, and the method includes: receiving, by the NP, a first packet, and obtaining a
15 packet feature of the first packet; obtaining, by the NP, a processing rule corresponding to the packet feature, and processing the first packet based on the processing rule, to obtain a second packet; sending, by the NP, the second packet to the CPU; and receiving and processing, by the CPU, the second packet.

[0015] In a possible design, the CPU sends the processing rule to the NP.

20 [0016] In a possible design, the processing rule is reporting the packet feature. The processing, by the NP, the first packet based on the processing rule, to obtain a second packet includes: adding, by the NP, a first extension header to the first packet, to obtain the second packet, where the first extension header includes the packet feature; and the processing, by the CPU, the second packet includes: obtaining, by the CPU, the packet
25 feature from the first extension header of the second packet; deleting the first extension header from the second packet, to obtain the first packet; and forwarding the first packet based on the packet feature.

[0017] In a possible design, the processing rule is obtaining routing information. The processing, by the NP, the first packet based on the processing rule, to obtain a
30 second packet includes: adding, by the NP, a first extension header to the first packet,

to obtain the second packet, where the first extension header includes the routing information; and the processing, by the CPU, the second packet includes: obtaining, by the CPU, the routing information from the first extension header of the second packet; deleting the first extension header from the second packet, to obtain the first packet;
5 and forwarding the first packet based on the routing information.

[0018] In a possible design, the processing rule is reporting service information. The processing, by the NP, the first packet based on the processing rule, to obtain a second packet includes: obtaining, by the NP, a service configured on an interface that receives the first packet, and adding a first extension header to the first packet, to obtain
10 the second packet, where the first extension header includes information about the service; and the processing, by the CPU, the second packet includes: obtaining, by the CPU, the information about the service from the first extension header of the second packet; deleting the first extension header from the second packet, to obtain the first packet; and performing the service on the first packet.

[0019] In a possible design, the obtaining, by the NP, a processing rule corresponding to the packet feature includes:

looking up, by the NP, in a rule table stored on the NP based on the packet feature;

obtaining the processing rule from the rule table when the rule table includes
20 the processing rule;

sending the first packet or the packet feature to the CPU when the rule table does not include the processing rule; and

receiving the processing rule returned by the CPU.

[0020] In a possible design, after the receiving, by the NP, the processing rule
25 returned by the CPU, the method further includes:

adding, by the NP, the processing rule to the rule table.

[0021] In a possible design, the method further includes:

sending, by the CPU, a third packet to the NP, where the third packet includes a second extension header, and the second extension header includes
30 processing information and a layer 2 header; and

obtaining, by the NP, the processing information and the layer 2 header in the second extension header; deleting the second extension header from the third packet, to obtain a fourth packet; processing the fourth packet based on the processing information; and adding the layer 2 header to the fourth packet, to obtain a fifth packet.

5 [0022] According to a third aspect, this application provides a network device, where the network device includes a processor and a memory, and

the memory stores a computer-executable instruction; and the processor executes the computer-executable instruction stored in the memory, to implement any one of the methods according to the second aspect of this application.

10 [0023] According to a fourth aspect, this application provides a computer-readable storage medium, where the computer-readable storage medium stores program code, and when executed by a processor, the program code implements any one of the methods according to the second aspect of this application.

[0024] According to a fifth aspect, this application provides a program product, 15 where the program product includes a computer program, the computer program is stored in a readable storage medium, a processor of a network device may read the computer program from the readable storage medium, and the processor executes the computer program, to enable the network device to perform any one of the methods according to the second aspect of this application.

20 [0025] According to the service processing method and the network device provided by this application, before the CPU of the network device processes the first packet, the NP of the network device first processes the first packet, and sends the second packet obtained after processing to the CPU, so that the CPU may directly process the second packet, to simplify the process of processing the received packet by 25 the CPU, thereby improving the service processing efficiency of the network device.

BRIEF DESCRIPTION OF DRAWINGS

[0026] FIG. 1 is a schematic diagram of a communications system according to an embodiment of this application;

[0027] FIG. 2 is a flowchart of a service processing method according to an embodiment of this application;

[0028] FIG. 3 is a flowchart of a service processing method according to another embodiment of this application;

5 [0029] FIG. 4 is a flowchart of a service processing method according to another embodiment of this application;

[0030] FIG. 5 is a flowchart of a service processing method according to another embodiment of this application;

[0031] FIG. 6 is a flowchart of a service processing method according to another
10 embodiment of this application;

[0032] FIG. 7 is a schematic structural diagram of a network device according to an embodiment of this application; and

[0033] FIG. 8 is a schematic structural diagram of a network device according to another embodiment of this application.

15

DESCRIPTION OF EMBODIMENTS

[0034] The term "at least one" in this application refers to one or more, and the term "a plurality of" refers to two or more. The term "and/or" describes an association relationship between associated objects and represents that three relationships may exist. For example, A and/or B may represent the following three cases: Only A exists, both
20 A and B exist, and only B exists. A and B may be singular or plural. The character "/" generally indicates an "or" relationship between the associated objects. "At least one of the following items" or other similar expressions refer to any combination of the items, including any combination of a single item or a plurality of items. For example, at least one of a, b, or c may be expressed as: a, b, c, a-b, a-c, b-c, or a-b-c, where a, b, and c
25 may be singular or plural.

[0035] FIG. 1 is a schematic diagram of a communications system according to an embodiment of this application. As shown in FIG. 1, the communications system includes a plurality of network devices. Some network devices are located in a local

area network, and some network devices are located in a wide area network, and a packet may be forwarded between the network devices, to be transmitted from a source network device to a target network device. By using the network device, the packet may be forwarded between the local area network and the wide area network.

5 [0036] Edge network devices among the network devices may use solutions in embodiments of this application. It should be noted that the network device in the embodiments of this application is not limited to the edge network device.

[0037] FIG. 2 is a flowchart of a service processing method according to an embodiment of this application. The service processing method is applied to a network
10 device, where the network device includes an NP and a CPU. As shown in FIG. 2, the method of this embodiment may include the following steps.

[0038] S201. The NP receives a first packet.

[0039] S202. The NP obtains a packet feature of the first packet.

[0040] In this embodiment, the packet received by the NP may be any packet, the
15 packet herein is referred to as the first packet, and after the first packet is received, the packet feature of the first packet is obtained from the first packet. The packet feature may be, for example, at least one of the following: a media access control (media access control, MAC) address, virtual local area network (virtual local area network, VLAN) information, a destination port number, a source port number, a protocol number, an
20 Internet Protocol (Internet Protocol, IP) address and other information.

[0041] Optionally, the NP may further obtain other information corresponding to the first packet from the first packet, for example, information about a to-be-processed service included in the first packet, routing information of the first packet, and a service configured on an interface that receives the first packet.

25 [0042] S203. The NP obtains a processing rule corresponding to the packet feature.

[0043] In a possible design, the CPU sends the processing rule to the NP. The CPU may actively send the processing rule to the NP, or after receiving a request that is for obtaining the processing rule and that is sent by the NP, the CPU sends the processing rule to the NP. The NP processes the received packet based on the processing rule, and
30 a manner in which the NP processes the packet is standardized, thereby improving

efficiency of processing the packet by the NP.

[0044] In a possible design, one possible implementation of S203 may include the following:

5 looking up, by the NP, in a rule table stored on the NP based on the packet feature;

obtaining the processing rule from the rule table when the rule table includes the processing rule; and sending the first packet or the packet feature to the CPU when the rule table does not include the processing rule; and

receiving the processing rule returned by the CPU.

10 [0045] In this embodiment, the NP stores the processing rule corresponding to the packet feature of the packet, and the processing rule is stored in a form of the rule table. Therefore, after obtaining the packet feature of the first packet, the NP looks up in the rule table based on the packet feature of the first packet, to obtain the processing rule corresponding to the packet feature of the first packet. When the rule table of the NP
15 already includes the processing rule corresponding to the packet feature of the first packet, the NP directly obtains the processing rule corresponding to the packet feature of the first packet from the rule table. When the rule table does not include the processing rule corresponding to the packet feature of the first packet, the NP needs to obtain the processing rule corresponding to the packet feature of the first packet from
20 the CPU. In an implementation, the NP may add an extension header to the first packet, to obtain a packet, where the extension header includes first information, and the first information is used to indicate that the NP does not include the processing rule, and then the NP sends the packet to the CPU, to request to obtain the processing rule from the CPU. The CPU receives the packet sent by the NP, obtains the first information from
25 the extension header of the packet, and sends the processing rule to the NP based on the first information. In another implementation, the NP may alternatively send the packet feature of the first packet to the CPU, and the CPU receives the packet feature of the first packet sent by the NP, and sends the processing rule corresponding to the packet feature of the first packet to the NP based on the packet feature of the first packet.
30 Correspondingly, the NP receives the processing rule that corresponds to the packet

feature of the first packet and that is sent by the CPU; and processes the first packet based on the processing rule, to obtain a second packet.

[0046] In this embodiment, after obtaining the packet feature, the NP in the network device first looks up in the rule table for the processing rule corresponding to the packet feature, and when the rule table does not include the processing rule corresponding to the packet feature, the NP sends the first packet or the packet feature to the CPU, to obtain the processing rule corresponding to the packet feature. In this way, a processing amount of the CPU is reduced, thereby improving processing efficiency of the CPU.

[0047] In a possible design, after the obtaining, by the NP, a processing rule corresponding to the packet feature, the method further includes: adding, by the NP, the processing rule to the rule table.

[0048] In this embodiment, after receiving the processing rule that corresponds to the packet feature of the first packet and that is sent by the CPU, the NP adds the processing rule corresponding to the packet feature of the first packet to the rule table, so that when subsequently receiving a packet that has a same packet feature as that of the first packet, the NP may directly find the processing rule corresponding to the packet feature from the rule table, to avoid a case in which processing amounts of the NP and the CPU are increased because the NP needs to obtain the processing rule from the CPU each time the NP receives the packet, and improve efficiency of processing the packet by the CPU, thereby improving service processing efficiency of the network device.

[0049] S204. The NP processes the first packet based on the processing rule, to obtain a second packet.

[0050] In this embodiment, after obtaining the packet feature of the first packet, the NP obtains the processing rule corresponding to the packet feature, where the processing rule is content agreed between the NP and the CPU. The processing rule, for example, may be that the NP obtains the packet feature of the packet received by the NP and sends the packet feature to the CPU, or may be that the NP obtains routing information of the packet received by the NP and sends the routing information to the CPU, or may be that the NP obtains a task configured on an interface that receives the

packet, and sends the task to the CPU, where the interface that receives the packet is, for example, a physical interface or a virtual interface on the network device. The NP processes the first packet based on the processing rule, to obtain the second packet, and the second packet may include information included in the first packet that is parsed by
5 the NP.

[0051] S205. The NP sends the second packet to the CPU, and correspondingly, the CPU receives the second packet.

[0052] S206. The CPU processes the second packet.

[0053] In this embodiment, the NP sends the second packet to the CPU, and the
10 CPU receives the second packet. Then the CPU processes the second packet. Because the second packet is obtained after processing by the NP based on the processing rule, the CPU does not need to process the received packet based on the processing rule, to simplify a process of processing the received packet by the CPU.

[0054] In this embodiment, the network device includes the NP and the CPU, where
15 the NP receives the first packet; obtains the packet feature of the first packet; obtains the processing rule corresponding to the packet feature; processes the first packet based on the processing rule, to obtain the second packet; and sends the second packet to the CPU; and the CPU receives and processes the second packet. Before the CPU processes the first packet, the NP processes the first packet, and sends the second packet obtained
20 after processing to the CPU, so that the CPU may directly process the second packet, to simplify the process of processing, by the CPU, the packet received by the CPU, thereby improving service processing efficiency of the network device.

[0055] FIG. 3 is a flowchart of a service processing method according to another
25 embodiment of this application. The service processing method is applied to a network device, where the network device includes an NP and a CPU. As shown in FIG. 3, the method in this embodiment is described by using an example in which a processing rule includes reporting a packet feature, and the method in this embodiment may include the following steps.

[0056] S301. The NP receives a first packet.

30 [0057] S302. The NP obtains a packet feature of the first packet.

[0058] In this embodiment, for specific implementations of S301 and S302, refer to S201 and S202, and details are not described herein again.

[0059] S303. The NP obtains a processing rule corresponding to the packet feature, where the processing rule is reporting the packet feature.

5 [0060] S304. The NP adds a first extension header to the first packet, to obtain a second packet, where the first extension header includes the packet feature.

[0061] S305. The NP sends the second packet to the CPU, and correspondingly, the CPU receives the second packet.

[0062] In this embodiment, the processing rule obtained by the NP may be reporting
10 the packet feature, and the reporting the packet feature is used to instruct the NP to send both the first packet and the packet feature of the first packet to the CPU. Therefore, the NP may add an extension header to the first packet, the extension header herein is referred to as the first extension header, and the first extension header includes the packet feature, to obtain the second packet that includes the packet feature of the first
15 packet and the first packet. The NP sends the second packet obtained after processing based on the reported packet feature to the CPU.

[0063] S306. The CPU obtains the packet feature from the first extension header of the second packet, and deletes the first extension header from the second packet, to obtain the first packet.

20 [0064] S307. The CPU forwards the first packet based on the packet feature.

[0065] In this embodiment, when the CPU receives the second packet that includes the packet feature of the first packet, because the NP already adds the packet feature of the first packet to the first extension header of the second packet, the CPU directly obtains the packet feature of the first packet from the first extension header of the
25 second packet. Then, the CPU deletes the first extension header of the second packet, to obtain the first packet, and forwards the first packet based on information included in the packet feature of the first packet. For example, the packet feature includes a MAC address, and the CPU forwards the first packet based on the MAC address at which the first packet is received.

30 [0066] In this embodiment, the NP in the network device adds the first extension

header that includes the packet feature to the first packet based on the processing rule of reporting the packet feature, to obtain the second packet, and sends the second packet to the CPU, so that the CPU may directly obtain the packet feature of the first packet from the first extension header of the second packet, and does not need to analyze the
5 obtained first packet to obtain the packet feature of the first packet, to simplify a process of processing the packet by the CPU, reduce time of processing the packet by the CPU, and improve efficiency of forwarding and processing the packet by the CPU, thereby improving service processing efficiency of the network device.

[0067] FIG. 4 is a flowchart of a service processing method according to another
10 embodiment of this application. The service processing method is applied to a network device, where the network device includes an NP and a CPU. As shown in FIG. 4, the method in this embodiment is described by using an example in which a processing rule includes obtaining routing information, and the method in this embodiment may include the following steps.

15 [0068] S401. The NP receives a first packet.

[0069] S402. The NP obtains a packet feature of the first packet.

[0070] In this embodiment, for specific implementations of S401 and S402, refer to S201 and S202, and details are not described herein again.

[0071] S403. The NP obtains a processing rule corresponding to the packet feature,
20 where the processing rule is obtaining routing information.

[0072] S404. The NP adds a first extension header to the first packet, to obtain a second packet, where the first extension header includes the routing information.

[0073] S405. The NP sends the second packet to the CPU, and correspondingly, the CPU receives the second packet.

25 [0074] In this embodiment, the processing rule obtained by the NP may be obtaining the routing information, and the obtaining the routing information is that the CPU obtains, from the NP, the routing information of the first packet obtained by the NP, and in this case, the NP needs to send both the first packet and the routing information of the first packet to the CPU. Therefore, the NP may add an extension
30 header to the first packet, the extension header herein is referred to as the first extension

header, and the first extension header includes the routing information, to obtain the second packet that includes the routing information of the first packet and the first packet. The routing information may include, for example, a route index, where the route index is an identifier of a router entry, and the identifier of the router entry is used
5 to indicate a storage location of the router entry in a routing table. The NP sends the second packet obtained after processing based on the processing rule to the CPU.

[0075] S406. The CPU obtains the routing information from the first extension header of the second packet, and deletes the first extension header from the second packet, to obtain the first packet.

10 [0076] S407. The CPU forwards the first packet based on the routing information.

[0077] In this embodiment, when the CPU receives the second packet that includes the routing information of the first packet, because the NP already adds the routing information of the first packet to the first extension header of the second packet, the CPU directly obtains the routing information of the first packet from the first extension
15 header of the second packet. Then, the CPU deletes the first extension header of the second packet, to obtain the first packet, and forwards the first packet based on the routing information of the first packet. For example, if the routing information includes the route index, after obtaining the route index, the CPU finds the storage location of the router entry in the routing table based on the route index, so that the router entry
20 stored in the storage location indicated by the route index may be obtained.

[0078] In this embodiment, the NP in the network device adds the first extension header that includes the routing information to the first packet based on the obtaining the routing information, to obtain the second packet, and sends the second packet to the CPU, so that the CPU may directly obtain the routing information of the first packet
25 from the first extension header of the second packet, and forward the first packet based on the routing information, and does not need to analyze the first packet to obtain the routing information of the first packet, to simplify a process of processing a packet by the CPU, reduce time of processing the packet by the CPU, and improve efficiency of forwarding and processing the packet by the CPU, thereby improving service
30 processing efficiency of the network device.

[0079] In some embodiments, based on the embodiment shown in FIG. 3 or FIG. 4, when the CPU forwards the first packet, the CPU may forward the first packet to the NP, so that the CPU may hand over, to the NP, subsequent forwarding in a process required to complete forwarding of the first packet. For example, when forwarding the first packet, the CPU needs to encapsulate the first packet. Currently, the encapsulation of the first packet is completed by the CPU. In this application, the CPU hands over the process of encapsulating the first packet to the NP, to reduce a processing amount of the CPU, and improve forwarding and processing efficiency of the CPU, thereby improving service processing efficiency of the network device, and reducing processing time of the packet.

[0080] FIG. 5 is a flowchart of a service processing method according to another embodiment of this application. The service processing method is applied to a network device, where the network device includes an NP and a CPU. As shown in FIG. 5, the method in this embodiment is described by using an example in which a processing rule includes reporting service information, and the method in this embodiment may include the following steps.

[0081] S501. The NP receives a first packet.

[0082] S502. The NP obtains a packet feature of the first packet.

[0083] In this embodiment, for specific implementations of S501 and S502, refer to S201 and S202, and details are not described herein again.

[0084] S503. The NP obtains a processing rule corresponding to the packet feature, where the processing rule is reporting service information.

[0085] S504. The NP obtains a service configured on an interface that receives the first packet, and adds a first extension header to the first packet, to obtain a second packet, where the first extension header includes information about the service.

[0086] S505. The NP sends the second packet to the CPU, and correspondingly, the CPU receives the second packet.

[0087] In this embodiment, the processing rule obtained by the NP may be reporting the service information, and the reporting the service information is used to instruct the NP to send both the first packet and the obtained service information of the first packet

to the CPU. Therefore, the NP obtains the service configured on the interface that receives the first packet, forms the service information by using the obtained service configured on the interface that receives the first packet, and adds an extension header to the first packet, the extension header herein is referred to as the first extension header, and the first extension header includes the service information, to obtain the second packet that includes the service information of the first packet and the first packet. The NP sends the second packet obtained after processing based on the processing rule to the CPU.

[0088] There is at least one service that is obtained by the NP and that is configured on the interface that receives the first packet. Therefore, the service information includes information about at least one service, and when the service information includes a plurality of pieces of information about services, the service information is a service stack formed by information about a plurality of services.

[0089] Optionally, in the at least one service that is obtained by the NP and that is configured on the interface that receives the first packet, a service that the NP can process may exist, and a service that the NP cannot process may also exist. Therefore, the NP may add information about the service that the NP cannot process to the first extension header, or the NP may alternatively add information about the at least one service configured on the interface that receives the first packet to the extension header. This is not limited in this embodiment of this application.

[0090] S506. The CPU obtains the information about the service from the first extension header of the second packet, and deletes the first extension header from the second packet, to obtain the first packet.

[0091] S507. The CPU performs the service on the first packet.

[0092] In this embodiment, when the CPU receives the information about the service configured on the interface that receives the first packet, because the NP already adds the information about the service configured on the interface that receives the first packet to the first extension header of the second packet, the CPU directly obtains, from the first extension header of the second packet, the information about the service configured on the interface that receives the first packet. Then, the CPU deletes the first

extension header of the second packet, to obtain the first packet. Then, the CPU performs the service based on the service information. For example, if the service that is configured on the interface receiving the first packet and that is obtained by the CPU by using the service information is collecting statistics on traffic, the CPU performs a
5 service of collecting statistics on traffic.

[0093] In this embodiment, the NP in the network device adds, based on the processing rule, the service configured on the interface that receives the first packet to the first extension header of the second packet, and sends the service to the CPU, so that the CPU may directly obtain, from the first extension header of the second packet,
10 the information about the service configured on the interface that receives the first packet, to directly process the service, and does not need to analyze the received packet to obtain the information about the service configured on the interface that receives the packet, to reduce time of processing the packet by the CPU, reduce the processing amount of the CPU, and improve efficiency of forwarding and processing the packet
15 by the CPU, thereby improving service processing efficiency of the network device.

[0094] FIG. 6 is a flowchart of a service processing method according to another embodiment of this application. As shown in FIG. 6, based on the embodiments shown in FIG. 2 to FIG. 5, the service processing method provided in this application may further include the following steps.

20 [0095] S601. The CPU sends a third packet to the NP, where the third packet includes a second extension header, and the second extension header includes processing information and a layer 2 header.

[0096] In this embodiment, the packet that should have been processed by the CPU may be handed over to the NP for processing, and in this case, when there is such a
25 packet that needs to be processed, the NP may directly process the packet, without occupying processing resources of the CPU. For example, an implementation may be that: The CPU sends a third packet to the NP, where the third packet includes a second extension header, and the second extension header includes processing information and a layer 2 header. The processing information is used to notify the NP of a manner of
30 processing the packet and information required for processing the packet. For example,

the CPU notifies, by using the processing information, the NP of whether to encapsulate the layer 2 header, or whether to change fields in the packet. The CPU instructs, by using the processing information, the NP to add a field to the packet, and sends the to-be-added field to the NP by using the processing information, and the layer 2 header
5 includes, for example, a MAC address, or VLAN information.

[0097] S602. The NP obtains the processing information and the layer 2 header in the second extension header.

[0098] S603. The NP deletes the second extension header from the third packet, to obtain a fourth packet.

10 [0099] S604. The NP processes the fourth packet based on the processing information.

[0100] S605. The NP adds the layer 2 header to the fourth packet, to obtain a fifth packet.

[0101] In this embodiment, the NP receives the third packet, and after obtaining the
15 processing information and the layer 2 header from the second extension header of the third packet, the NP deletes the second extension header of the third packet, to obtain the fourth packet. Then, the NP adds the obtained layer 2 header to the fourth packet, to obtain the fifth packet.

[0102] This embodiment is described by using an example in which layer 2 virtual
20 extensible local area network (virtual extensible local area network, VXLAN) processing that should have been performed by the CPU is handed over to the NP for processing.

[0103] The CPU configures the layer 2 VXLAN, and sends information related to the layer 2 VXLAN to the NP, so that the NP performs encapsulation and decapsulation
25 of the layer 2 VXLAN on the packet based on the information related to the layer 2 VXLAN. The information related to the layer 2 VXLAN may include, for example, any one or more of an instruction of performing encapsulation and decapsulation of VXLAN on the packet, a MAC address, a VLAN identifier, a source IP address used for VXLAN encapsulation, a destination IP address used for VXLAN encapsulation,
30 and a VXLAN network identifier used for VXLAN encapsulation.

[0104] In this embodiment, the CPU hands over the packet that should have been processed by the CPU to the NP for processing, to reduce a quantity of packets processed by the CPU, and improve efficiency of forwarding and processing the packet by the CPU, thereby improving service processing efficiency of the network device,
5 and reducing time of forwarding and processing the packet.

[0105] Based on the above, in this embodiment of this application, the CPU and the NP may be integrated to collaboratively process the packet, or the packet processed by the CPU may be handed over to the NP for independent processing, so that the NP undertakes a part of work of the CPU, to improve efficiency of forwarding and
10 processing the packet by the CPU, thereby improving service processing efficiency of the network device, and reducing time of processing the packet.

[0106] FIG. 7 is a schematic structural diagram of a network device according to an embodiment of this application. As shown in FIG. 7, the network device 700 provided in this embodiment includes an NP 701 and a CPU 702. The NP 701 is
15 configured to: receive a first packet; obtain a packet feature of the first packet; obtain a processing rule corresponding to the packet feature; process the first packet based on the processing rule, to obtain a second packet; and send the second packet to the CPU 702; and the CPU 702 is configured to: receive the second packet, and process the second packet.

20 [0107] In a possible design, the CPU 702 is further configured to send the processing rule to the NP 701.

[0108] In a possible design, the processing rule is reporting the packet feature, where

the NP 701 is configured to: add a first extension header to the first packet,
25 to obtain the second packet, where the first extension header includes the packet feature; and

the CPU 702 is configured to: obtain the packet feature from the first extension header of the second packet; delete the extension header from the second packet, to obtain the first packet; and forward the first packet based on the packet feature.

30 The processing rule is obtaining routing information. The NP 701 is configured to add

a first extension header to the first packet, to obtain the second packet, and the first extension header includes the routing information. The CPU 702 is configured to: obtain the routing information from the first extension header of the second packet; delete the first extension header from the second packet, to obtain the first packet; and
5 forward the first packet based on the routing information.

[0109] In a possible design, the processing rule is reporting service information. The NP 701 is configured to: obtain a service configured on an interface that receives the first packet, and add a first extension header to the first packet, to obtain the second packet, and the first extension header includes information about the service; and the
10 CPU 702 is configured to: obtain the information about the service from the first extension header of the second packet; delete the first extension header from the second packet, to obtain the first packet; and perform the service on the first packet.

[0110] In a possible design, when obtaining the processing rule corresponding to the packet feature, the NP 701 is configured to: look up in a rule table stored on the NP
15 701 based on the packet feature; obtain the processing rule from the rule table when the rule table includes the processing rule; and when the rule table does not include the processing rule, send the packet or the packet feature to the CPU 702; and receive the processing rule returned by the CPU 702.

[0111] In a possible design, the NP 701 is further configured to add the processing
20 rule to the rule table.

[0112] In a possible design, the CPU 702 is configured to send a third packet to the NP 701, where the third packet includes a second extension header, and the second extension header includes processing information and a layer 2 header.

[0113] The NP 701 is configured to: obtain the processing information and the layer
25 2 header in the second extension header; delete the second extension header from the third packet, to obtain a fourth packet; process the fourth packet based on the processing information; and add the layer 2 header to the fourth packet, to obtain a fifth packet.

[0114] The network device provided in this embodiment may be configured to perform the technical solutions of the foregoing method embodiments. The
30 implementation principles and technical effects are similar, and are not further

described herein.

[0115] FIG. 8 is a schematic structural diagram of a network device according to another embodiment of this application. As shown in FIG. 8, the network device 800 of this embodiment may include: at least one processor 801 and a memory 802. A network
5 device having one processor is used as an example in FIG. 8.

[0116] The memory 802 is configured to store a program. Specifically, the program may include program code, and the program code includes a computer-executable instruction. The memory 802 may include a high speed random access memory (random access memory, RAM), or may further include a non-volatile memory (non-
10 volatile memory), such as at least one magnetic disk memory.

[0117] The processor 801 is configured to execute the computer-executable instruction stored in the memory 802, to implement the service processing method shown in the foregoing method embodiments and performed in the foregoing
embodiments.

[0118] Optionally, in a specific implementation, if the memory 802 and the
15 processor 801 are independently implemented, the memory 802 and the processor 801 may be connected to each other by using a bus and complete communication between each other. The bus may be an Industry Standard Architecture (Industry Standard Architecture, ISA) bus, a Peripheral Component Interconnect (Peripheral Component
20 Interconnect, PCI) bus or an Extended Industry Standard Architecture (Extended Industry Standard Architecture, EISA) bus, or the like. The bus may be categorized as an address bus, a data bus, a control bus, or the like, but this does not mean that there is only one bus or a bus of one type.

[0119] Optionally, in a specific implementation, if the memory 802 and the
25 processor 801 are integrated on one chip for implementation, the memory 802 and the processor 801 may complete communication between each other by using an internal interface.

[0120] The network device in this embodiment may be configured to perform the technical solutions of the foregoing method embodiments. The implementation
30 principles and technical effects are similar, and are not described herein again.

[0121] When the functions are implemented in the form of a software functional unit and sold or used as an independent product, the functions may be stored in one computer-readable storage medium. Based on such an understanding, the technical solutions of this application essentially, or the part contributing to the prior art, or some
5 of the technical solutions may be implemented in a form of a software product. The software product is stored in one storage medium, and includes several instructions for instructing a computer device (which may be a personal computer, a server, a network device, or the like) to perform all or some of the steps of the methods described in the embodiments of this application.

10 [0122] The foregoing descriptions are merely specific implementations of this application, but are not intended to limit the protection scope of this application. Any variation or replacement readily figured out by a person skilled in the art within the technical scope disclosed in this application shall fall within the protection scope of this application. Therefore, the protection scope of this application shall be subject to the
15 protection scope of the claims.

CLAIMS

1. A network device, comprising a network processor (NP) and a central processing unit (CPU), wherein the network device is an edge router, and wherein:

the NP is configured to:

receive a first packet;

obtain a packet feature of the first packet;

obtain a processing rule corresponding to the packet feature;

process the first packet based on the processing rule, to obtain a second packet, wherein the NP is configured to add a first extension header to the first packet to obtain the second packet, and wherein the first extension header comprises the packet feature; and

send the second packet to the CPU; and

the CPU is configured to: receive the second packet, and process the second packet.

2. The network device according to claim 1, wherein when obtaining the processing rule corresponding to the packet feature, the NP is configured to:

look up in a rule table stored on the NP based on the packet feature;

obtain the processing rule from the rule table when the rule table comprises the processing rule; and

when the rule table does not comprise the processing rule, send the first packet or the packet feature to the CPU, and receive the processing rule returned by the CPU.

3. The network device according to claim 1 or 2, wherein

the processing rule is reporting the packet feature;

the CPU is configured to:

obtain the packet feature from the first extension header of the second packet;

delete the first extension header from the second packet, to obtain the first packet;

and

forward the first packet based on the packet feature.

4. The network device according to claim 1 or 2, wherein the NP is further configured to: obtain a third packet;

obtain a second packet feature of the third packet;

obtain a second processing rule that is obtaining routing information corresponding to the second packet feature;

process the third packet to obtain a fourth packet with a second extension header having the routing information; and
send the fourth packet to the CPU;

wherein the CPU is further configured to:

receive the fourth packet with the second extension header;

obtain the routing information from the second extension header of the fourth packet;

delete the second extension header from the fourth packet, to obtain the third packet; and

forward the third packet based on the routing information.

5. The network device according to claim 1 or 2,

wherein the NP is further configured to:

obtain a service configured on an interface that receives a third packet;

process the third packet to obtain a fourth packet with a second extension header having information about the service; and

send the fourth packet to the CPU; and

wherein the CPU is further configured to:

receive the fourth packet with the second extension header;

obtain the information about the service from the second extension header of the fourth packet;

delete the second extension header from the fourth packet, to obtain the third packet; and
perform the service on the third packet.

6. The network device according to any one of claims 1 to 3, wherein

the CPU is configured to: send a third packet to the NP, wherein the third packet comprises a second extension header, and the second extension header comprises processing information and a layer 2 header; and

the NP is configured to:

obtain the processing information and the layer 2 header in the second extension header; delete the second extension header from the third packet, to obtain a fourth packet;

process the fourth packet based on the processing information; and

add the layer 2 header to the fourth packet, to obtain a fifth packet.

7. A service processing method, applied to a network device, the network device being an edge router, wherein the network device comprises a network processor (NP) and a central

processing unit (CPU), and the method comprises:

receiving, by the NP, a first packet, and obtaining a packet feature of the first packet;

obtaining, by the NP, a processing rule corresponding to the packet feature, and processing the first packet based on the processing rule, to obtain a second packet, wherein the processing the first packet based on the processing rule comprises adding a first extension header to the first packet to obtain the second packet, wherein the first extension header comprises the packet feature;

sending, by the NP, the second packet to the CPU; and

receiving and processing, by the CPU, the second packet.

8. The method according to claim 7, wherein the obtaining, by the NP, a processing rule corresponding to the packet feature comprises:

looking up in a rule table stored on the NP based on the packet feature;

obtaining the processing rule from the rule table when the rule table comprises the processing rule;

sending the first packet or the packet feature to the CPU when the rule table does not comprise the processing rule; and

receiving the processing rule returned by the CPU.

9. The method according to claim 7 or 8, wherein the processing rule is reporting the packet feature;

the processing, by the CPU, the second packet comprises: obtaining, by the CPU, the packet feature from the first extension header of the second packet;

deleting the first extension header from the second packet, to obtain the first packet;

and

forwarding the first packet based on the packet feature.

10. The method according to claim 7 or 8, further comprising:

obtaining, by the NP, a third packet;

obtaining, by the NP, a second packet feature of the third packet,

obtaining, by the NP, a second processing rule that is obtaining routing information corresponding to the second packet feature;

processing, by the NP, the third packet to obtain a fourth packet with a second extension header having the routing information;

sending, by the NP, the fourth packet to the CPU;

receiving, by the CPU, the fourth packet with the second extension header;

obtaining, by the CPU, the routing information from the second extension header of the

fourth packet;

deleting, by the CPU, the second extension header from the fourth packet, to obtain the third packet; and

forwarding the third packet based on the routing information.

11. The method according to claim 7 or 8, further comprising:

obtaining, by the NP, a service configured on an interface that receives a third packet;

processing the third packet to obtain a fourth packet with a second extension header having information about the service;

sending, by the NP, the fourth packet to the CPU;

obtaining, by the CPU, the information about the service from the second extension header of the fourth packet;

deleting, by the CPU, the second extension header from the fourth packet, to obtain the third packet; and

performing, by the CPU, the service on the third packet.

12. The method according to any one of claims 7 to 9, further comprising:

sending, by the CPU, a third packet to the NP, wherein the third packet comprises a second extension header, and the second extension header comprises processing information and a layer 2 header; and

obtaining, by the NP, the processing information and the layer 2 header in the second extension header; deleting the second extension header from the third packet, to obtain a fourth packet; processing the fourth packet based on the processing information; and adding the layer 2 header to the fourth packet, to obtain a fifth packet.

13. A network device, wherein the network device is an edge router, and wherein the network device comprises a network processor (NP) and a central processing unit (CPU), the network device is configured to perform the service processing method according to any one of claims 7 to 12.

14. A readable storage medium, wherein the readable storage medium stores a computer program, and when executed by an edge router, the computer program implements the service processing method according to any one of claims 7 to 12.




FIG. 1

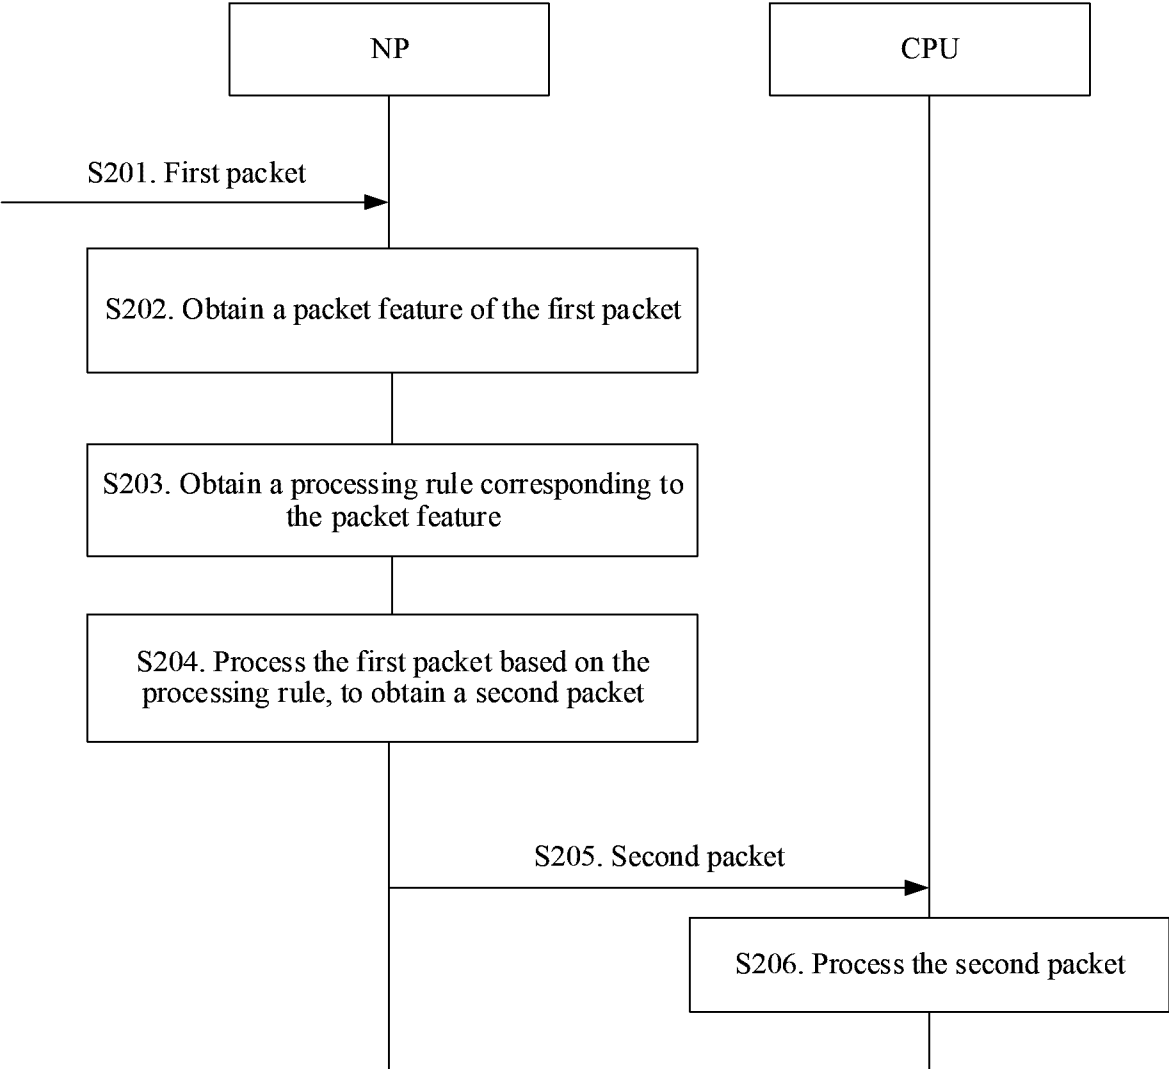


FIG. 2

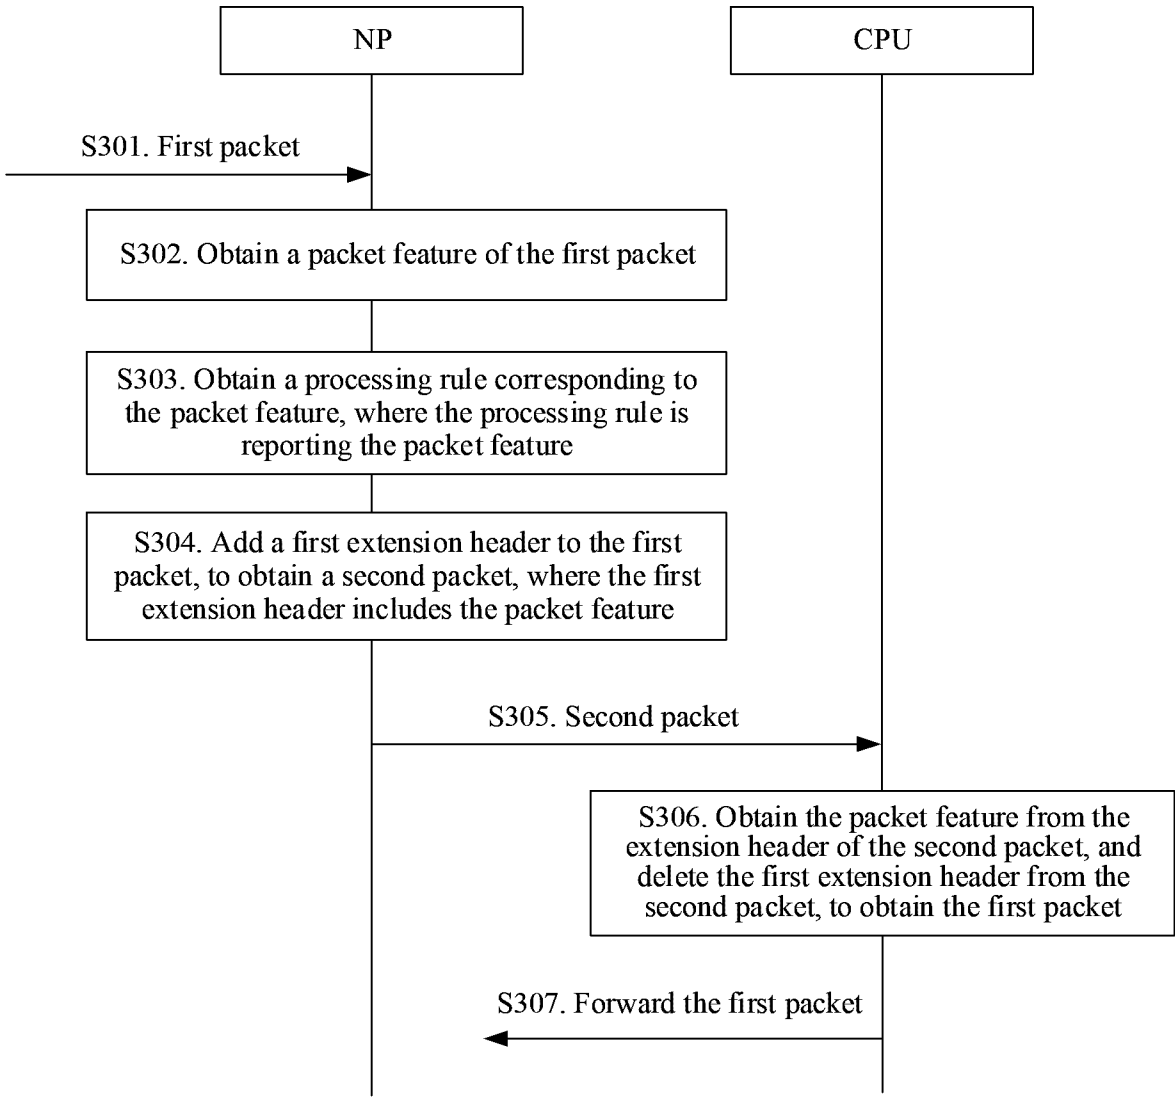


FIG. 3

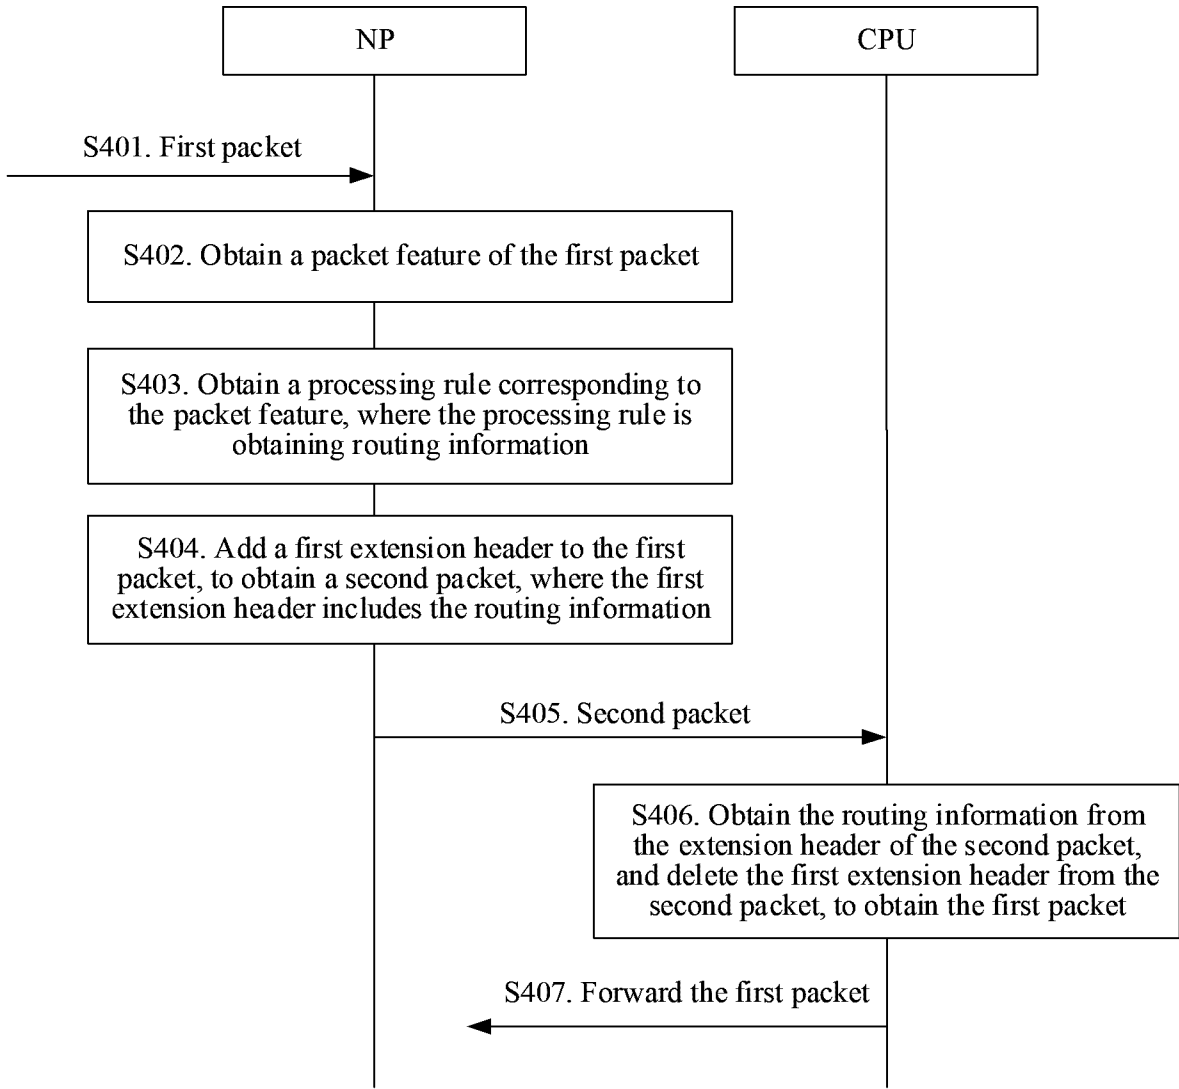


FIG. 4

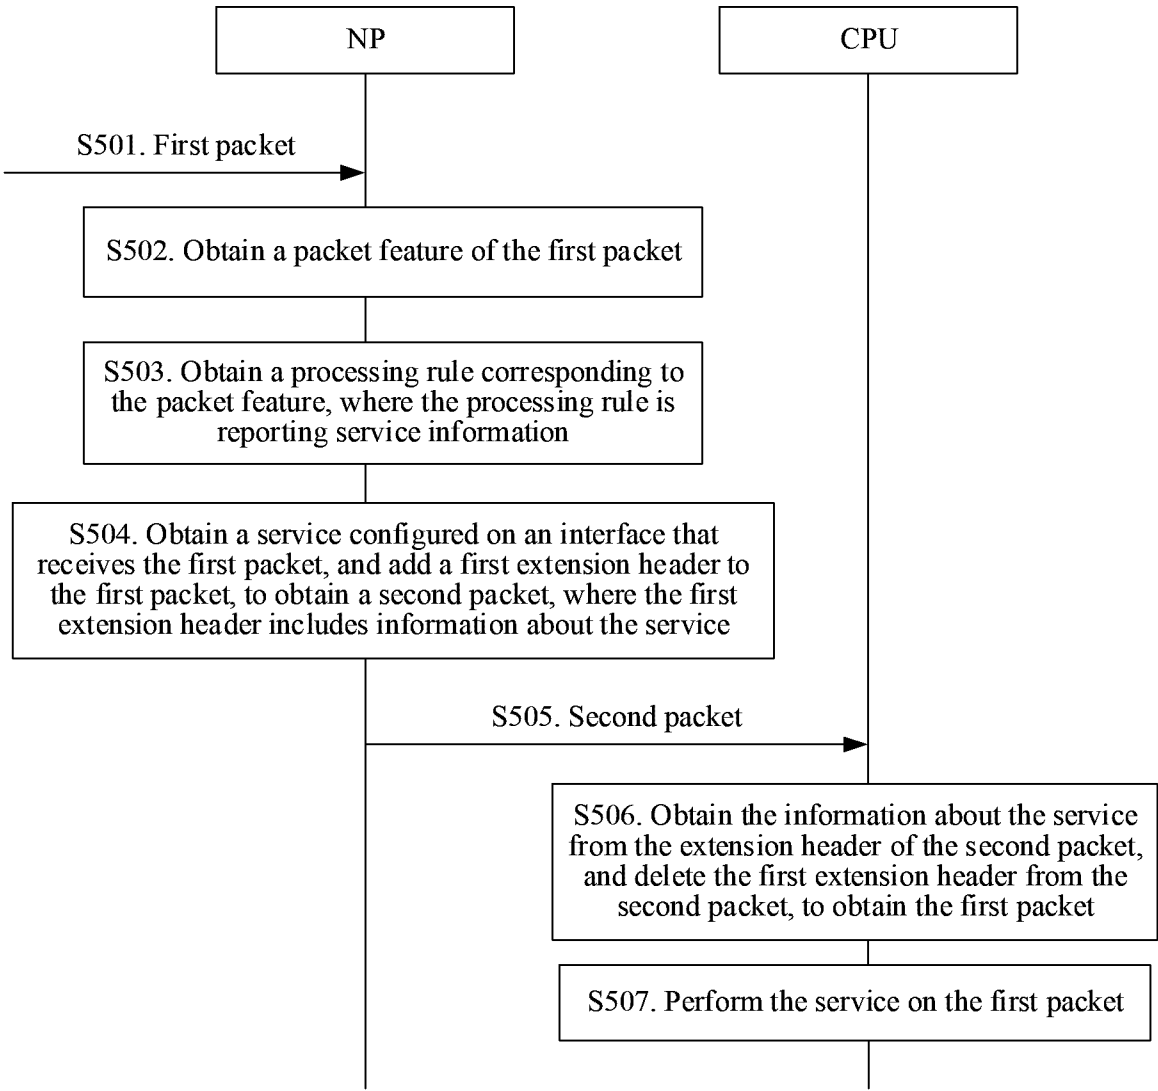


FIG. 5

5/6

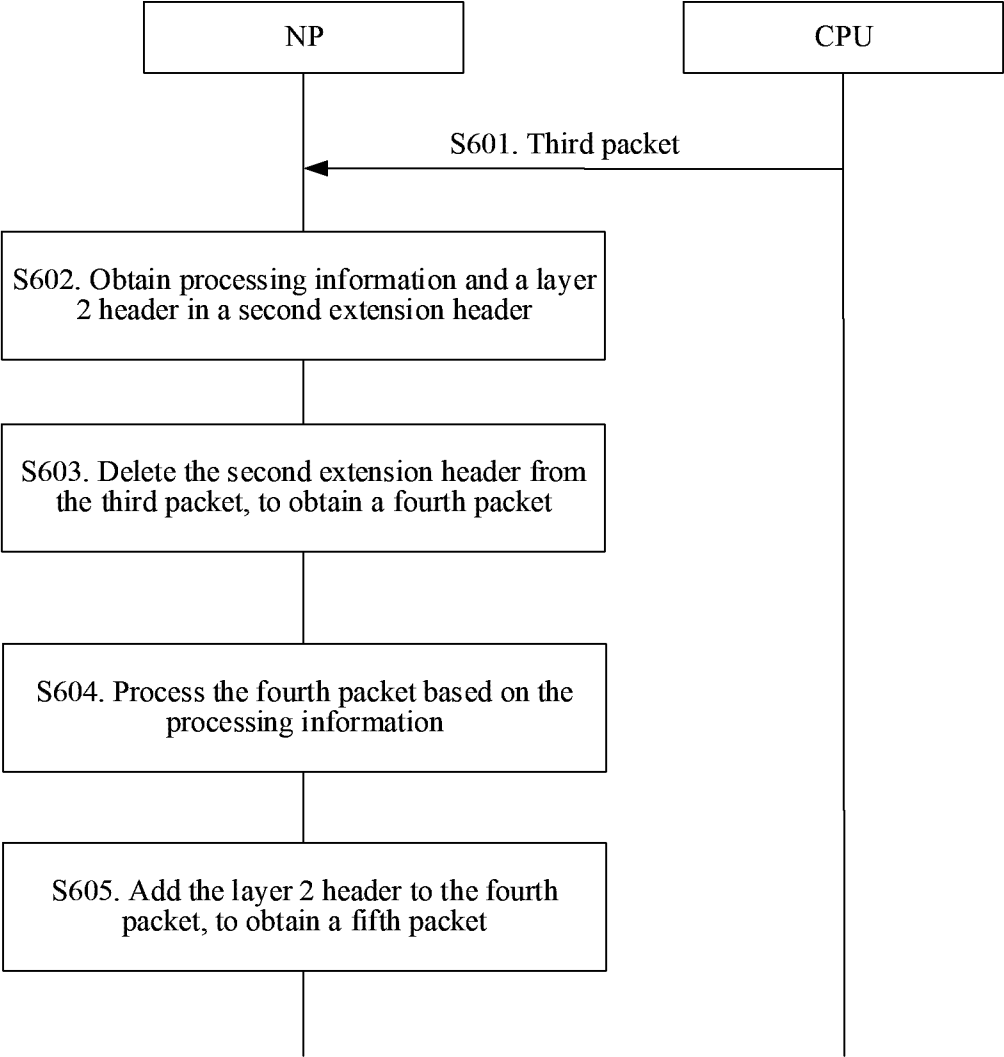


FIG. 6

6/6




FIG. 7




FIG. 8

