CHAPTER 6

NAT-PT ENHANCEMENT TO SUPPORT MULTICAST

From chapter 4, NAT-PT was found to have been created to provide address translation and communication between IPv4 and IPv6 for unicast only. It cooperates with DNS-ALG to provide applications between the v4 and the v6 worlds. However, multicast applications such as videoconferencing, multimedia, real audio or video cannot be provided between v4 and v6 networks because NAT-PT does not support multicast. The standard for NAT-PT did not define address translation and communication for multicast.

This chapter presents the design of an enhancement to NAT-PT to support multicast address translation. The requirements for, and problems with enhancing NAT-PT are introduced in section 6.1. The plans to improve multicast function following the assumptions are presented in section 6.2. The solution to provide multicast address translation, an algorithm and a way to provide group address mapping are in section 6.3. The procedures of NAT-PT to process packet including address translation are proposed in section 6.4 and 6.5 respectively. The summary of all tasks is in section 6.6.

6.1 Introduction to NAT-PT enhancement

Multicasting is one-to-many communication. When multicast is used, data packets might need to be transmitted to every network. Multicast routers perform this task. They duplicate the original packet and send to every networks attached to its interfaces that requires the packet.

In order to provide a multicast communication between v4 and v6 realms, the same session must be identified with a unique group address for each realm. The respective addresses must be understood by v4 and v6 recipients for the same session. Since a group address can be an arbitrary number for a multicast application, the problem is how to provide group address mapping between IPv4 and IPv6 multicast addresses.

6.2 Plan to extend NAT-PT module

This section presents some assumptions and a plan to develop multicast address translation between IPv4 and IPv6. Assumptions:

1. The extension function will be created as a prototype to present the idea of multicast address translation for NAT-PT and determine its feasibility.
2. This work focuses on the enhancement NAT-PT to translate multicast address. Because we want to determine the feasibility of this, we initially limit the design to allow for only a single NAT-PT translator of multicast.
3. Only global scope multicast addresses require translation. This direction can provide multicast services in widely use of v4 and v6 recipients in the Internet.

We make these assumptions in order to determine whether multicast through NAT-PT is practical. This work is a prototype and needs implementation to test whether it is possible before focusing on performance. Having determined that we can later determine the effects of relaxing these assumptions. So we implement NAT-PT with multicast on a basic and simple isolated network in our laboratory. There is only one NAT-PT box to provide multicast communication between an IPv4-only and an IPv6-only networks. The simple testing network makes for easy checking whether our solution can provide multicast address translation on NAT-PT.

Multiple NAT-PT routers can be deployed as gateway on a network in order to perform load balancing. However, we must consider packet looping inside that network. In addition, the multiple boxes must coordinate address mapping between IPv4 and IPv6. The mapped address must not be duplicated by the other NAT-PT routers. Performance of NAT-PT or speed to process multicast packets needs to be measured and enhanced. These issues will be further studied and developed in future work.

We expect multicast through NAT-PT to be similar to unicast. NAT-PT should provide address translation from IPv4 to IPv6 and vice versa. When a v4 server provides a multicast v4 application, it sends multicast v4 packets to the world. These multicast packets must be delivered to all networks where they are required. When a v6 node wants to join this group, it needs to know an appropriate group number to join. This requires translation from IPv4 to IPv6 multicast addresses. In the inverse path, the v6 recipient uses the translated IPv6 group number to identify the application in order to send to the group. This IPv6 address must be translated back to the original IPv4 address corresponding to that session.

The address for translation can be manually added to the mapping rules by a human. In this case, the mapping list might be permanently assigned to provide the particular multicast session between v4 and v6 sites. In practice, however, group addresses are dynamically assigned for multicast applications. It is not convenient for administrator to update mapping rules following the changing of group address. So NAT-PT should provide automatic multicast address translation.

The NAT-PT enhancement to provide multicast address translation has been planed as follows:

1. IPv4 to IPv6 and inverse path

Multicast IPv4 to IPv6 address translation has been developed in the first step because it permits v6 stub networks to receive multicast data from the existing v4 Internet. This is likely to be the most required operation in the short term.

When a v4 server provides a multicast v4 application, it sends multicast v4 packets to the world. These multicast packets must be delivered to every network where they are wanted. When a v6 node wants to join this group, it needs to know the group address that it should use and, clearly, it must be a v6 address. Thus we need a v6 translation of a v4 multicast address.

The v6 recipient does not use the IPv6 group number to identify or reply to the source, for that it would use the source’s unicast address (as translated from v4 to v6). It uses the multicast address to participate in the session and send to all group members (which might not include the source of any packets being received). This IPv6 address must be translated back to the original IPv4 address corresponding to that session. The translation backwards is called inverse path address mapping of IPv4 to IPv6 as shown in Figure 6.1.

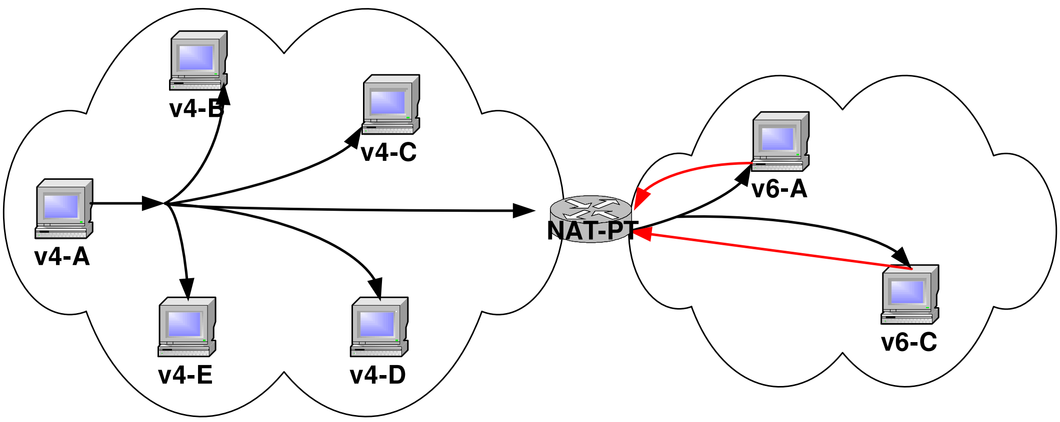

Figure 6.1 IPv4 to IPv6 communication

2. IPv6 to IPv4 and inverse path

In this direction, a v6 node would like to initiate multicast communication. Multicast data are created by the v6 node and forwarded to every network. When a v4 node would like to join this session, it must know which group address it should join. But the session is defined with IPv6 group number. The group address must be translated to IPv4 multicast address in order to be presented to that v4 node.

Whenever the v4 recipient wants to send to the group, it uses the IPv4 group number – the IPv4 address to identify that session. This group address must be translated to the original IPv6 multicast address. This backward translation is called inverse path address mapping of IPv6 to IPv4.

6.3 The method to develop NAT-PT multicast module

The enhancements to NAT-PT to add multicast can be divided into three steps as follows:

1. The solution to provide multicast packet mapping between IPv4 and IPv6.
2. The algorithm to translate multicast address.
3. The way to announce group address between v4 and v6 realms.

6.3.1 The solution to provide multicast packet mapping

The multicast address translation function is extended to NAT-PT to change an incoming packet to be the appropriate one for the recipient. This function must provide two main tasks: address and packet translation.

The function translates multicast addresses from IPv4 to IPv6 and vice versa following the mapping rules. The algorithm and method to translate IP addresses are presented in section 6.3.2.

NAT-PT changes a multicast packet from IPv4 to IPv6 and vice versa. The existing, unicast, NAT-PT changes the IP header of the original packet to the new one that contains the appropriate address for the recipient. Since multicast and unicast packets use the same formats, we re-use the packet translation method from unicast for multicast.

IPv4 and IPv6 have their own protocols to find the route for multicasting and forwarding packets. The router relies on the existing principles to process each translated packet. The packet can be treated as a normal multicast packet in the new protocol, originating from the NAT-PT. The source address, used for multicast forwarding, will also have been translated to one from a block of unicast addresses assigned to the NAT-PT.

However, address translation impacts the packet that it occurs in. The address translation impacts the upper layer protocol, TCP, UDP and ICMP. The transport layer checksum includes a pseudo-header including the network layer addresses which obviously must be updated. The port numbers might be changed. Most multicast transport protocols are layered above UDP as a basic transport layer. NAT-PT already handles UDP updates (port mapping if required, and checksum correction to allow for the altered pseudo-header.) It turns out that the most popular of the upper transport protocols used by multicast applications – the Real Time Protocol (RTP) [34] needs no further adjustment. Thus aside from the multicast destination mapping function, multicast translation needs no alteration to that used for unicast.

6.3.2 The algorithm to translate multicast address

Multicast typically uses dynamically allocated addresses. The value of the address is largely irrelevant, the only real issue is that address duplication be avoided, so traffic for distinct group is not merged. Any unused address of the correct scope would be acceptable.

The address translation principle of NAT-PT has been extended to provide multicast address translation. The procedure to translate multicast IPv4 and IPv6 address is similar to unicast address translation. The algorithm is different when translating IPv4 than when translating IPv6 addresses, so each will be presented separately.

1. Algorithm to translate IPv4 to IPv6 address

Translating unicast v4 addresses to v6 is done by simply prepending a unicast v6 prefix assigned to the NAT-PT ahead of the v4 address, 96 bits of prefix and 32 bits of v4 address generate the v6 address.

For multicast a similar approach is possible, provided that the prefix is a multicast prefix. Multicast prefixes are not allocated to specific sites or networks, however the multicast address allocation scheme from RFC3306 [37] allows sites to generate multicast addresses by using one of their unicast prefix as (a part of) a prefix for all addresses generated this way. 32 bits of address space remain available for the site to use for multiple group identifiers.

By allocating one of its subnet prefixes for NAT-PT multicast address allocation, most likely the same prefix used for generating v6 unicast addresses, the NAT-PT can easily generate v6 multicast addresses that should never clash with any other v6 multicast addresses.

While this should work, and would certainly be adequate where there is only one NAT-PT translating multicast addresses, it does have the property that v6 multicast addresses generated from the same v4 multicast address by different NAT-PT devices will differ, which may not be desirable.

An alternative is to reserve one designated prefix for use by all NAT-PT devices. This will allow a single v4 multicast group to be translated to a single v6 multicast group, regardless of how many NAT-PTs may be involved.

Which approach is best, or whether some hybrid of the two may actually perform better has not been considered in this work, and is left for further study. For our purposes either is adequate – we choose the global prefix scheme for our tests, and selected a multicast prefix to use instead of the usual unicast prefix when generating addresses. This guarantees that no site following the RFC3306 procedure for multicast address generating will accidentally duplicate an address generated by NAT-PT. The procedure for generating a translation from v4 to v6 multicast may be represented as shown in Figure 6.1

IPv4-mcast-addr(32 bits) => IPv6-mcast-prefix(96 bits) + IPv4-mcast-addr

Figure 6.1 Algorithm to translate multicast IPv4 to IPv6 address

2. Algorithm to translate IPv6 to IPv4 address

The simple algorithm used when translating v4 to v6 multicast addresses is not possible when translating v6 to v6, v6 addresses cannot be embedded into v4 addresses.

Further, since v6 multicast addresses contain 112 bits of essentially random, and certainly unpredictable value, simple taking the low N bits (for any N no greater than the maximum possible value of 28) and using that as the variable part of the v4 multicast address does not provide any expectation of uniqueness of the resulting address.

While the v4 to v6 mapping scheme cannot be directly used, it does suggest an approach that may be productive. That is, if there were a reserved block of v4 multicast address, not used by anything other than NAT-PT, then the problem of avoiding duplicate addresses has been reduced to the problem of ensuring different NAT-PT devices do not accidentally generate the same v4 multicast address. On the other hand, when multiple NAT-PTs are translating, the same v6 multicast address we anticipate that generating the same v4 group identifier would be the correct outcome. There is sufficient reserved v4 multicast address space remaining that even reserving a large block of addresses for this purpose remains feasible.

The way to choose, and the size of, the dedicated address block for NAT-PT are important issues that we must consider. This will be a future work. For our purpose the size of the reserved block is not important, nor is the prefix value of that block. It is also not important for our work what method is used to obtain an address from within the block. All these issues remain for future study. For our purpose it is sufficient to simply allocate the next available address from the block, which can be represented as shown in Figure 6.1

IPv6-mcast-addr => next-free(IPv4-mcast-addr)

Figure 6.1 Algorithm to translate multicast IPv6 to IPv4 address

6.3.3 The way to announce group address between v4 and v6 realms

When a node provides multicast services, information of these sessions should be available for the other nodes in network. Normally, multicast communication has a protocol to provide group address discovery as presented in section 5.4 SDP and SAP could be treated like a normal multicast packet. When they are advertised, they would be available for listeners in every network – v4 and v6.

SAP has two versions to deliver a session description for each listener. There is SAP for IPv4 and IPv6. It has a well known multicast IPv4 and IPv6 address for each version. However, the data content would be meaningless if either SAP IPv4 or IPv6 is advertised to the other network. In order to present multicast information to the other world, the SAP packet must be adjusted to be the appropriate packet for listeners in each IP version. Also the IP address in SAP header must be mapped to the correct well known address for the other IP version.

SDP is used to describe multicast sessions. It is carried by SAP via networks to listeners. When passing from one IP version to the other, the group address of the session must be updated to allow listener in the other IP version to join that group.

Since SAP/SDP is like the DNS and FTP which contain IP addresses in the payload, NAT-PT does not see these addresses and cannot perform address translation. An ALG for multicasting is necessary to be used to help NAT-PT like DNS-ALG and FTP-ALG. This work proposes an ALG to provide multicast address mapping to the group address in SAP/SDP. We call it the Session Announcement Protocol-Application Level Gateway (SAP-ALG). It is a function on the NAT-PT router or gateway.

An Application Layer Gateway for SAP/SDP (SAP-ALG) is added to NAT-PT as shown in Figure 6.1. SAP-ALG gets an advertisement, a SAP packet in IP version X (SAPvX). It looks for the group IPvX address in the content of SDP. Then it allocates, if required, and maps an appropriate multicast address to the original group number. SAP-ALG uses the method from section 6.3.2 to allocate multicast address for group address mapping. NAT-PT kernel code is informed of the mapping. SAP-ALG generates a new SDP message to contain the mapped group IPvY address. The new SDP holds the same data entries as the original SDP message that SAP-ALG learns from. Again SAP-ALG creates the new SAP in IP version Y (SAPvY) that used to carry the new SDP.

X and Y are the variable that represented to version of IP – 4 and 6. X and Y can be either 4 or 6 in each direction of a communication.

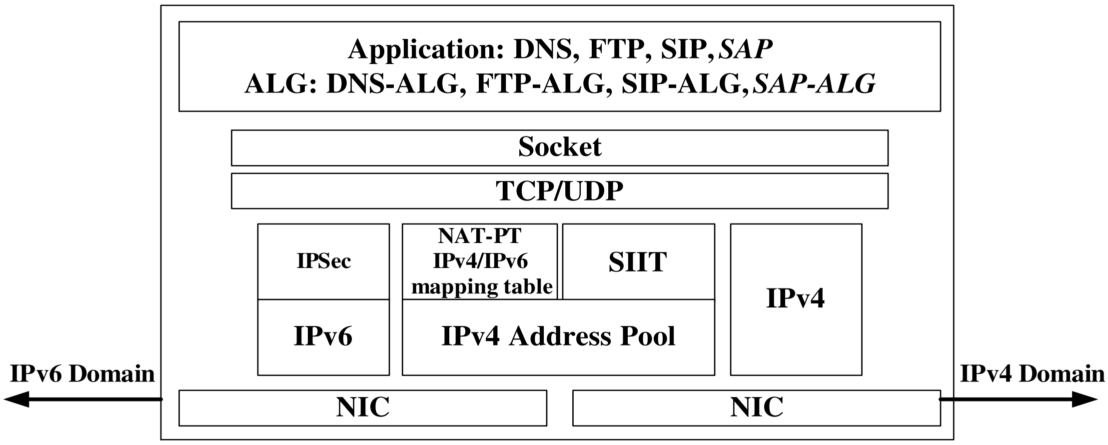

Figure 6.1 Diagram of SAP-ALG stack

SAP-ALG is the extension module and works at application layer to help NAT-PT provide multicast address translation. This module is presented in Figure 6.2 The component of SAP-ALG has been designed as follows:

1. Receiving function

This is used to get the advertisement message from v4 and v6 announcers. It listens to the well known addresses of SAP messages for IPv4 and IPv6.

2. New SAP message generation function

Sessions from the received message are translated to the other protocol including allocation of a suitable group address, and added as locally sourced SDP descriptions to the SAP transmit queue.

3. Transmit address function to NAT-PT

The address mapping that was allocated by SAP-ALG must be transmitted to NAT-PT. So packets belonging to the session can be correctly forwarded between the networks.

4. Forwarding function

SAP-ALG will send the translated packet to the participant networks. The packet is transmitted as any normal multicast packet to the destination network according to the rules for SAP packet transmission.

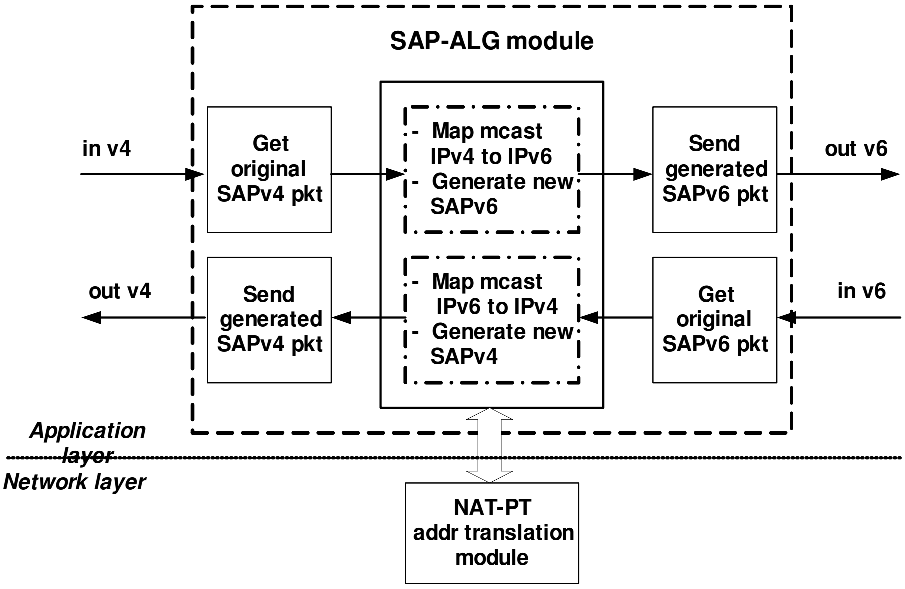

Figure 6.2 SAP-ALG module

6.4 Procedure of SAP-ALG to provide multicast address mapping

This section presents steps of SAP-ALG to provide address mapping. When an announcer advertises the schedule of a multicast application, it generates a SAP packet and sends to the network. Like any normal multicast router, NAT-PT with multicast enabled gets the SAP packet and forwards it. SAP-ALG on the NAT-PT router also receives the packet and starts the process to provide address mapping.

The procedure to provide address mapping by SAP-ALG is shown in Figure 6.1 X and Y represent the IP version, one 4, the other 6. The steps in each direction can be explained as follows:

1. When an IPvX announcer advertises SAPvX packet, SAP-ALG receives this packet.
2. The SAP packet has been treated like as normal multicast packet. It has been copied in order to be forwarded to every network. No address mapping for SAP packets will present in the NAT-PT mapping table, so the IPvX packet will not be forwarded to the IPvY network. At the same time, SAP-ALG looks for the address in the entry's session description. It must change these addresses to be the appropriate values for the listeners. Two addresses exist: the group number and the creator address.
3. For the group number, SAP-ALG compares the IPvX group address with the contents of the mapping list. If the group address exists in the list, SAP-ALG uses the existing IPvY from the list and maps to that IPvX. Otherwise, it generates the new multicast

IPvY number. The new number has been mapped to that IPvX address. And the pair of IPvX and IPvY are added to the mapping list.

4. SAP-ALG creates a new SAP packet. It is the SAP in IP version Y (SAPvY). It consists of header and payload. The packet is addressed to the well known SAP address in IP version Y. The payload is the session description obtained from the payload of SAPvX packet. Nothing is changed except the group number and address of creator.
5. The group number is assigned with the multicast IPvY address from the mapping list. The address of creator is IPvY address of SAP-ALG because this packet is created by SAP-ALG.
6. Figure 6.7 shows process to create the new SAP packet.
7. SAP-ALG sends the new SAP packet to IPvY network.

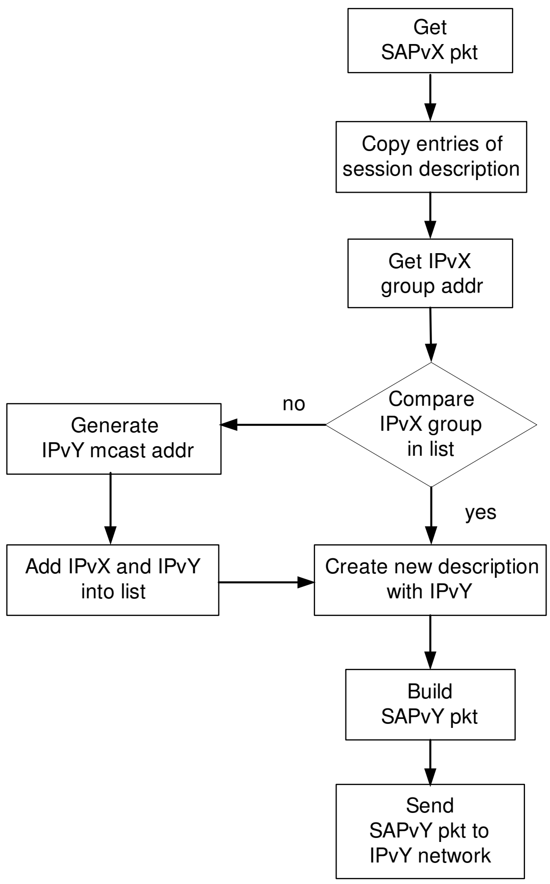

Figure 6.1 Procedure of address mapping by SAP-ALG

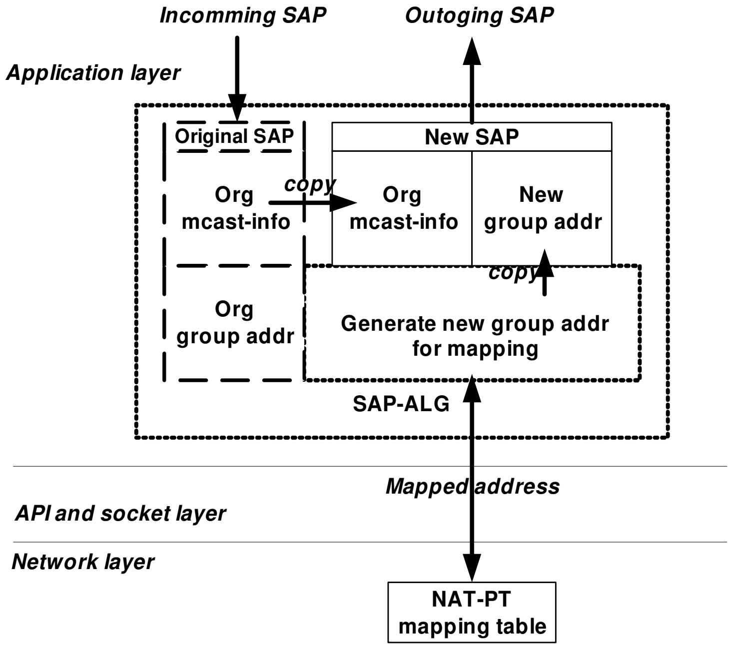

Figure 6.2 Process to create new SAP packet by SAP-ALG

6.5 Procedure to translate multicast packet

The following part presents procedure of packet translation. After the mapping rule has already been arranged by SAP-ALG, it has been passed to NAT-PT. Addresses in the rules are used to translate group address and packet in both directions. When either v4 or v6 source sends data packet of a multicast session, they are translated following the rules by NAT-PT.

The procedures to translate a packet are shown in Figure 6.1 In this diagram, the variable X and Y are used to represent number of IP version. Multicast packets are processed using the following steps:

1. When an IPvX source initiates a multicast IPvX application and sends IPvX packets, the multicast routing protocol in the IPvX network distributes the packets. NAT-PT, a function on an IPvX multicast router, participates in that multicast routing algorithm. Thus all packets it needs get delivered to it.
2. The router checks the properties of the received packets and determines destination of the packet. Since the packet is multicast, it is processed like normal multicast packets. The packet is duplicated, and the router forwards these packets via every appropriate IPvX network interface for IPvX recipients. The sending relies on multicast packet forwarding, RPF, and multicast routing protocol for IPvX to find the path.

3. NAT-PT with the multicast function looks for the IPvX group address of the packet in the mapping rules. The way to check and translate multicast packet for IPv4 and IPv6 groups are shown in Figures 6.9 and 6.10 respectively.
4. If the group address is in the list, the packet is translated to IPvY format with a new destination address from the mapping rule.
5. Packets that have been translated are treated as a normal multicast packet that arrived from any other node. They are sent via every IPvY network interface of NAT-PT to IPvY recipients. The router uses multicast routing protocol for IPvY to find the paths for sending the translated packet.
6. If the group address is not in NAT-PT's mapping rules, no packet translation is performed, and the packet is not forwarded to IPvY. There cannot be any useful purpose to dynamically generating a new mapping and using that, as there could not possibly be any numbers of the new group, that group identifier having only been created by NAT-PT after the packet arrived.

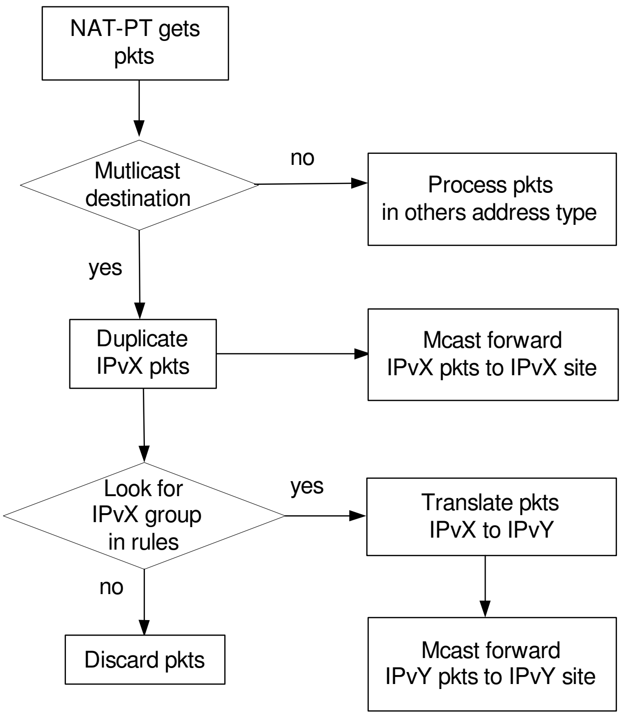

Figure 6.1 Packet processing by NAT-PT

```
if (IN_MULTICAST(IP4-dest-addr))
{
    if   IP4-mcast-addr  match in mapping rules
            assign mcast IPv6 prefix to IP4-mcast-addr
            and send translated packet
    else
            not translate to the other protocol
}
```

Figure 6.2 Packet processing to check IPv4 group address

```
if (IN_MULTICAST(IP6-dest-addr))
{
    if   IP6-mcast-addr  match in mapping rules
            assign IP4-mcast-addr to IP6-mcast-addr
            and send translated packet
    else
             not translate to the other protocol
}
```

Figure 6.3 Packet processing to check IPv6 group address

6.6 Summary

The enhancement relies on some functions of the original NAT-PT and the principles of multicast. Packet and address translation for multicast packets for both v4 and v6 are possible. An ALG for SAP has been designed to create the address mappings and make the mapped groups known to the other network.

However, the enhancement is planned for only the basic functionality in order to make sure that NAT-PT can provide address translation and communication for multicast. There are some limitations. The plan is designed for only one NAT-PT box to provide multicast application in global scope only. The way to apply multiple NAT-PT boxes in the same network and the way to solve the effect of the multiple boxes are not considered.