Realizing the Benefits of Virtual LANs by Using IPv6

Thorsten Kurz   Jean-Yves Le Boudec
Hans Joachim Einsiedler

March 22, 1997

Abstract

The benefits that Virtual LANs offer can be realized by using features of an IPv6 network along with small enhancements the IPv6 and DHCPv6 protocol stacks.

Contents

1 Introduction

Virtual LANs are a widespread special form of LANs (Local Area Networks) enhanced by support for workgroups. A LAN is a collections of hosts which communicate directly on layer 2 without a router between them. All hosts on a LAN share the same layer 3 subnet address, which means communication between the hosts of a LAN remains in the LAN. Thus the layer 3 subnet address forms a broadcast scope which contains all hosts on the LAN.

The performance, security and broadcast scope offered by LANs are used to build workgroups i.e. groups of hosts sharing the same servers and resources. Therefore all hosts of a workgroup are attached to the same LAN segment, so that broadcasting can be used for server detection, name resolution and name reservation like as is done by B-nodes in the NetBIOS Protocol [17].

In order to overcome the limits of traditional LANs, switched LANs appeared, which use a switch infrastructure to connect several LAN segments even over high speed backbones. Switched LANs continue to share the same layer 3 subnet address, but offer an increased performance compared to traditional LANs, because not all hosts of a switched LAN have to share the bandwidth of the same LAN segment. The possibility to connect the LAN segments over backbones makes it feasible to distribute hosts over larger areas than a single LAN segment could cover.

In environments with several different workgroups, running on different LANs, using traditional switched LANs requires a separate switch infrastructure for each LAN. Virtual LANs are switched LANs with a software configurable switch infrastructure. This makes it possible to operate several different LANs over the same switch infrastructure and to change easily the LAN membership of single segments. Workgroups can then be formed and maintained by a central administration.

The disadvantage of virtual LANs is however, that a special switch infrastructure is needed and administration includes layer 2 as well as layer 3. A solution which offers the same features but involves only layer 3 and does not require special hardware is desirable.

We show here that with the help of features of the new Internet Protocol version 6 (IPv6) [1] it is possible to form a flexible broadcast scope based on layer 3. The features used are multicast addressing, mobility support and the dynamic host configuration protocol for IPv6 (DHCPv6). These features exist also for the current Internet Protocol (IPv4), but as they are all optional for IPv4, their availability cannot be assumed in an IPv4 network. Another drawback of the IPv4 multicasting is, that it can only be scaled using the Time To Live field, whereas in IPv6 there are different multicast scopes [2].

The focus of this proposal is to realize a flexible broadcast scope. Security can be achieved using authentication [13] and encryption [14] mechanisms for the Internet Protocol (IP). Regarding the performance of virtual LANs, it can be expected that there will be routers which rival the performance of switches. Tag switching [15] is one possible way to implement the packet forwarding of a router in hardware.

2 Flexible Broadcast Scope with IPv6

IPv6 has enhanced the broadcast of IPv4 into scalable multicast [2]. In IPv4 there is only one broadcast address for a particular scope and broadcasts are always received by all hosts in this scope. In IPv6, on the other hand, there is a special address range reserved for multicast addresses for each scope and a multicast is only received by those hosts in this scope, which are configured to listen to this specific multicast address. To address all hosts in a scope with a multicast, the multicast must be made to the predefined *all nodes address* [5], which all hosts must listen to. When existing software using IPv4 is migrated to IPv6 it is expected that the IPv4 broadcasts are changed to multicasts to the *all nodes address*, as this is the simplest way to maintain the complete functionality of the software.

We propose to use IPv6 multicasting to form the broadcast scope of a workgroup (WG). This means a workgroup is the multicast (MC) group, whose hosts listen all to the same multicast address, the workgroup address. Since a host can listen to several multicast addresses at the same time, a host can even be a member of several workgroups.

While in a virtual LAN the workgroup membership of a host is determined by the configuration of the switches, in our proposal a host has to determine its workgroups and their corresponding addresses. The separation of different workgroups takes place on layer 3 since, with multicasting, each host has the possibility to address a specified subset of hosts of the network. Thus all hosts can be connected to routers directly, even members of different workgroups can share the same LAN segment.

To make the administration of the workgroups easy, we propose that the correspondence between hosts and workgroups be stored in a central database and the information be distributed using the Dynamic Host Configuration Protocol version 6 (DHCPv6) [8].

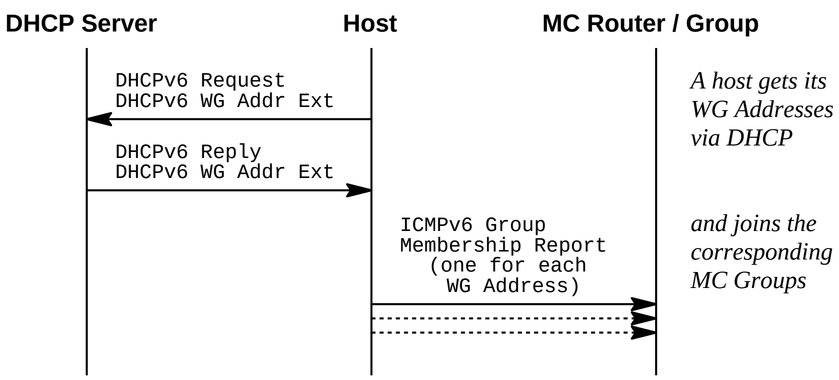

Figure 1: Configuration of Workgroup Addresses

Figure 1 shows the proposed workgroup (WG) address configuration for a host. When starting up, the host must send a DHCP Request with a Workgroup Address Extension, as described in the following section, to its DHCP Server. The DHCP Server must reply with a Workgroup Address Extension containing all workgroup addresses assigned to this host. After receiving its workgroup addresses, the host has to send an *ICMPv6 Group Membership Report* [3] to each of its workgroup addresses to inform the multicast routers about its new membership in these multicast groups.

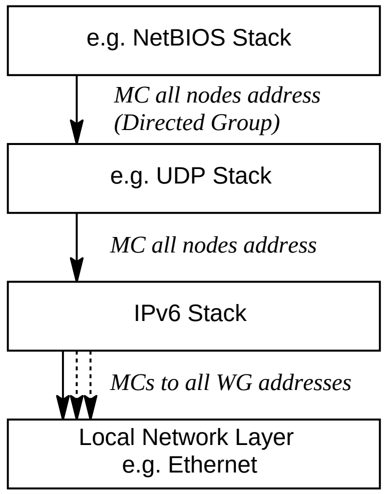

Figure 2: Translation to Workgroup Addresses

After learning its workgroup (WG) addresses, the host also has to configure its interfaces to listen to these multicast (MC) addresses. Additionally it is required that all outgoing multicasts to the *all nodes address*, which are equivalent to broadcasts, are changed to multicasts to the workgroup addresses of the host. This is accomplished by patching the IPv6 stack to intercept all outgoing multicasts to the *all nodes address* and to change this address to the workgroup addresses of the host as shown in figure 2. If the host is a member of several workgroups the multicast has to be sent to all workgroup addresses of the host.

3 DHCP Workgroup Address Extension

The purpose of DHCP [8] is to provide hosts with addresses and other configuration information. DHCP delivers the configuration data in extensions [9] which are embedded in request, reply or reconfigure messages. The request message is used by the client to request configuration data from the server and the reply message is used by the server to return the requested information to the client. If there is a change in the DHCP database, the server uses the reconfigure message to notify the client about the change and to start a new request-reply cycle.

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---------------+---------------+---------------+---------------+
|   Ext Type    |  Ext Length   |   Reserved    |Workgroup Count|
+---------------+---------------+---------------+---------------+
|                                                               |
+--                     Workgroup Addresses                   --+
|          (one for each Workgroup, 16 octets each)             |
+--                                                           --+
|                                                               |
+---------------------------------------------------------------+
|               Node Name (variable Length) ...
+-----------------------------------------------------------
```

Figure 3: DHCP Workgroup Address Extension

We propose a DHCP Extension, shown in figure 3, in order to deliver to a host its workgroup addresses.

Ext Type Identifier for the extension.

Ext Length The length of the extension.

Reserved Must be zero.

Workgroup Count Number of workgroup multicast addresses contained in this extension. For each workgroup the client belongs to there is one multicast address.

Workgroup Addresses The 16 octet multicast address of each workgroup the client is member of.

Node Name The node name is the DNS domain name of the client which is used as an unique identifier to look up client specific information in the server databases.

The DHCP Workgroup Address Extension has to be used the following way:

- In a DHCP Request the client
  - must set the workgroup count to zero.
  - must not specify any workgroup addresses.
  - must specify its node name.
- In a DHCP Reply the server
  - must set the workgroup count to number of workgroup addresses existing for this client.
  - must include all workgroup addresses existing for this client.
  - must use the clients node name.
- In a DHCP Reconfigure the server
  - must set the workgroup count to zero.
  - must not specify any workgroup addresses.
  - must use the clients node name.

4 Mobility for Multicast Group Members

In the future we have to deal with more and more mobile hosts, some of which are members of workgroups. The Internet draft *Mobility Support in IPv6* [10] proposes that a mobile host attached to a network segment other than its home segment keeps its home address on the home segment and forms a global care-of address for its new location. Then binding update options included in IPv6 packets are used to inform correspondent hosts as well as the home agent, a router which is on the same segment as the home address of the mobile host, about its new care-of address. After being informed about a new care-of address of the mobile host the home agent intercepts packets on the home segment addressed to the mobile host and tunnels [11] them to the care-of address of the mobile host. There is no specified way as to how a mobile host could send or receive multicast packets from its home network. We suggest the following enhancements to the Internet draft *Mobility Support in IPv6* [10], so that mobile multicast group members can continue to participate in the multicast traffic of their group.

If a mobile host leaves the scope of a multicast group it joined, the home agent must not only forward packets sent to the home address of the mobile host, but also all packets sent to the concerned multicast address.

Furthermore, the mobile host has to be able to send packets to the multicast address of its workgroup, even though it is outside the scope of this address. This can only be done by tunneling [11] the packets to a host inside the scope of the multicast address and resending them from there. Since the home agent is on the segment associated to the home address of the mobile host, the task of resending multicasts of a mobile host can also be taken over by the home agent.

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
                                +---------------+---------------+
                                |  Option Type  | Option Length |
+-+-+-+-+-----------------------+---------------+---------------+
|A|H|L|M|        Reserved       |           Lifetime            |
+-+-+-+-+-----------------------+-------------------------------+
|                         Identification                        |
+--                                                           --+
|                                                               |
+---------------------------------------------------------------+
|                                                               |
+--                                                           --+
|                        Care-of Address                        |
+--                                                           --+
|                                                               |
+--                                                           --+
|                                                               |
+---------------------------------------------------------------+
|                                                               |
+--                                                           --+
|                    Home Link-Local Address                    |
+--                                                           --+
|                  (only present if L bit set)                  |
+--                                                           --+
|                                                               |
+---------------------------------------------------------------+
|                                                               |
+--                                                           --+
|                       Multicast Address                       |
+--                                                           --+
|                  (only present if M bit set)                  |
+--                                                           --+
|                                                               |
+---------------------------------------------------------------+
```

Figure 4: Enhanced Binding Update Option

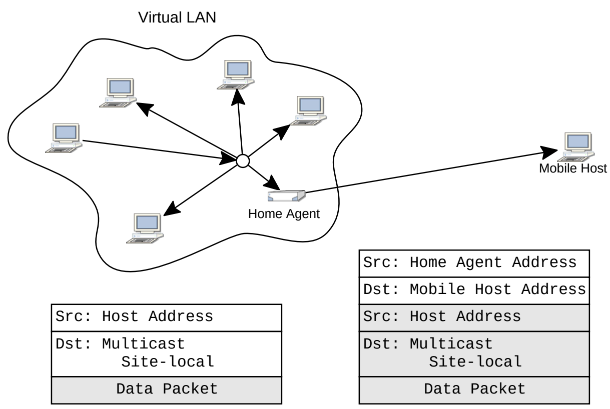

Figure 5: Mobile Host Receiving Multicast

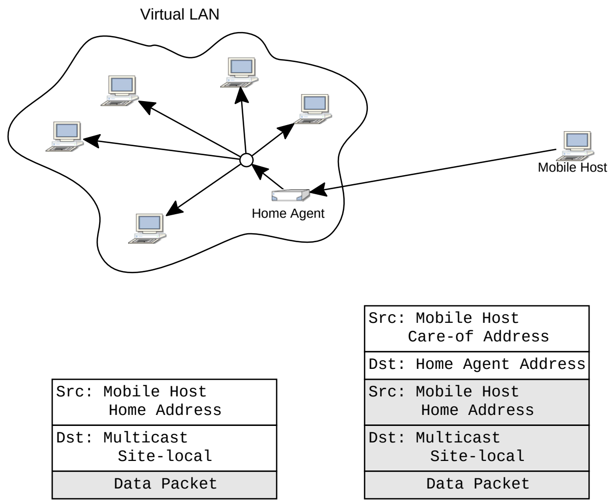

Figure 6: Mobile Host Sending Multicast

The Internet draft *Mobility Support in IPv6* [10] proposes a binding update option, which is used to notify the home agent and other hosts about a new care-of address of a mobile host. The original home address of the mobile host has to be specified in the source address field in the IP header of the packet containing the binding update option, but it is not allowed to specify a multicast address at this place. Hence there must be an optional field for a multicast address in the binding update option. Figure 4 shows the binding update option enhanced by a field for the workgroup address and an M-Bit. The M-bit is used to indicate that there is a multicast group address specified in the option.

A mobile host, which has left the scope of one of its multicast groups, sends a binding update option to its home agent in order to inform it about a new care-of address and has to specify its multicast group address in the binding update option and set the M-bit in this option. In case the mobile host is a member of several multicast groups, it has to send a binding update option for each of its multicast groups.

A home agent notified by a binding update option about a multicast address for a mobile host must join this multicast group, if it has not already done so, and handle packets with this multicast address in the destination address field in the same way as packets with the home address of the mobile node in this field (Figure 5). See section 7.3 of *Mobility Support in IPv6* [10] for further details.

The mobile host must treat a received encapsulated multicast packet the same way as if it received the decapsulated packet directly. It must not send a binding update option to the address specified in the source address field of an encapsulated multicast packet.

When sending a multicast packet to its multicast group, the mobile host has to use its home address in the source address field of the multicast packet and tunnel this packet to its home agent (Figure 6).

When a home agent receives an encapsulated multicast packet in which the source address field is the same as the home address of a mobile host served by it, then the home agent has to act like a router, receiving this multicast packet from the home segment of the mobile host and additionally forwarding it to the home segment of the mobile host.

This way of providing mobile workgroup members with the possibility to leave the scope of the multicast address has the drawback that it might not scale very well in the case of broadcast intensive workgroup protocol stacks, since all the broadcasting traffic, which was intended to remain in a limited area, has to be forwarded to the mobile node. Assuming that a lot of workgroup members use the possibility of global mobility, there is a risk of overloading the Internet with workgroup broadcasting traffic.

5 Conclusion

Under the described conditions equivalent features of a Virtual LAN can be realized by using IPv6. Today the latency and the throughput of routers are still a deterrent to a solution based on layer 3 routing, but tag switching [15] for example can increase the performance of routers considerably. The IPv6 flowlabel field can be used to place the tag information [16].

Nevertheless, the use of authentication and encryption mechanisms in the end nodes raises the latency and possibly also impacts the throughput of the end nodes.

Considering the usage of Virtual LANs and IPv6, VLANs enhance the flexibility of currently available software without requiring any changes of the software. Software, which is being adapted to the new IPv6 address space in the future, can be changed to use the *all nodes multicast address* instead of IPv4 broadcast. Using IPv6 no special VLAN protocols and hardware is required and only small enhancements in the IPv6 protocol stack must be done. IPv6 can offer a viable software alternative to Virtual LANs as soon as faster solutions for routing are available.

Abbreviations

DHCP Dynamic Host Configuration Protocol.

ICMPv6 Internet Control Message Protocol for IPv6.

IGMP Internet Group Management Protocol.

IPv4 Internet Protocol version 4.

IPv6 Internet Protocol version 6.

LAN Local Area Network.

MC MultiCast

WG WorkGroup

References

[1] *Internet Protocol Version 6*, S. Deering, R. Hinden, December 1995, RFC 1883

[2] *IP Version 6 Addressing Architecture*, R. Hinden, S. Deering, December 1995, RFC 1884

[3] *Internet Control Message Protocol version 6*, A. Conta, S. Deering, December 1995, RFC 1885

[4] *Host Extensions for IP Multicasting*, S. Deering, August 1989, RFC 1112

[5] *IPv6 Multicast Address Assignments*, R. Hinden, November 1996, draft-ietf-ipngwg-multicast-assgn-01.txt

[6] *Neighbor Discovery for IPv6*, T. Narten, E. Nordmark, August 1996, RFC 1970

[7] *IPv6 Stateless Address Autoconfiguration*, S. Thomson, T. Narten, August 1996, RFC 1971

[8] *Dynamic Host Configuration Protocol for IPv6*, J. Bound, C. Perkins, November 1996, draft-ietf-dhc-dhcpv6-08.txt

[9] *Extensions for DHCPv6*, C. Perkins, November 1996, draft-ietf-dhc-v6exts-04.txt

[10] *Mobility Support in IPv6*, C. Perkins, D. Johnson, November 1996, draft-ietf-mobileip-ipv6-02.txt

[11] *Generic Packet Tunneling in IPv6*, A. Conta, S. Deering, November 1996, draft-ietf-ipngwg-ipv6-tunnel-05.txt

[12] *Security Architecture for the Internet Protocol*, R. Atkinson, August 1995, RFC 1825

[13] *IP Authentication Header*, R. Atkinson, August 1995, RFC 1826

[14] *IP Encapsulation Security Payload*, R. Atkinson, August 1995, RFC 1827

[15] *Tag Switching Architecture Overview*, Y. Rekhter, B. Davie, D. Katz, E. Rosen, G. Swallow, September 1996, draft-rfced-info-rekhter-00.txt

[16] *Use of Flow Label for Tag Switching*, F. Baker, Y. Rekhter, October 1996, draft-baker-flow-label-00.txt

[17] *IBM PC Network Technical Reference*, Document Number 6322916, September 1984