(12) UK Patent (19) GB (11) 2465017 (13) B

(45) Date of B Publication 16.05.2012

(54) Title of the Invention: Many-to-many embedded multicast addressing and access control method

(51) INT CL: ***H04L 29/06*** (2006.01) ***H04L 12/56*** (2006.01) ***H04L 29/12*** (2006.01)

(21) Application No: **0820316.8**

(22) Date of Filing: **06.11.2008**

(43) Date of A Publication **12.05.2010**

(56) Documents Cited:
KR 030052452 A **US 6243758 B1**
US 20090067360 A1 **US 20080165783 A1**
RFC 4291

(58) Field of Search:
As for published application 2465017 A viz:
INT CL **H04L**
Other: **WPI EPODOC INSPEC**
updated as appropriate

(72) Inventor(s):
Imed Romdhani

(73) Proprietor(s):
The Court of Edinburgh Napier University
(Incorporated in the United Kingdom)
Craiglockhart Campus, 219 Colinton Road,
EDINBURGH, Lothian, EH14 1DJ, United Kingdom

(74) Agent and/or Address for Service:
Murgitroyd & Company
Scotland House, 165-169 Scotland Street,
GLASGOW, G5 8PL, United Kingdom

GB 2465017 B

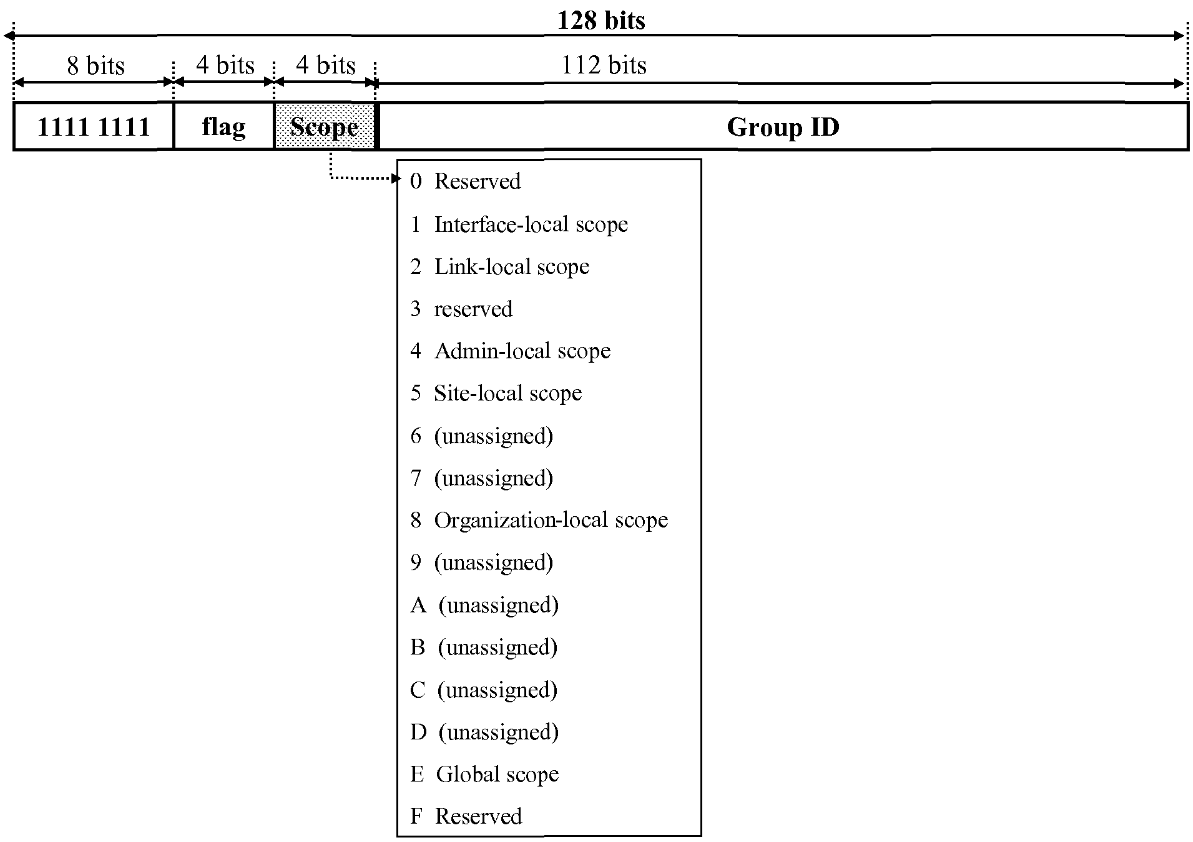
128 bits
8 bits
4 bits
4 bits
112 bits
1111 1111
flag
Scope
Group ID
0 Reserved
1 Interface-local scope
2 Link-local scope
3 reserved
4 Admin-local scope
5 Site-local scope
6 (unassigned)
7 (unassigned)
8 Organization-local scope
9 (unassigned)
A (unassigned)
B (unassigned)
C (unassigned)
D (unassigned)
E Global scope
F Reserved

Fig. 1

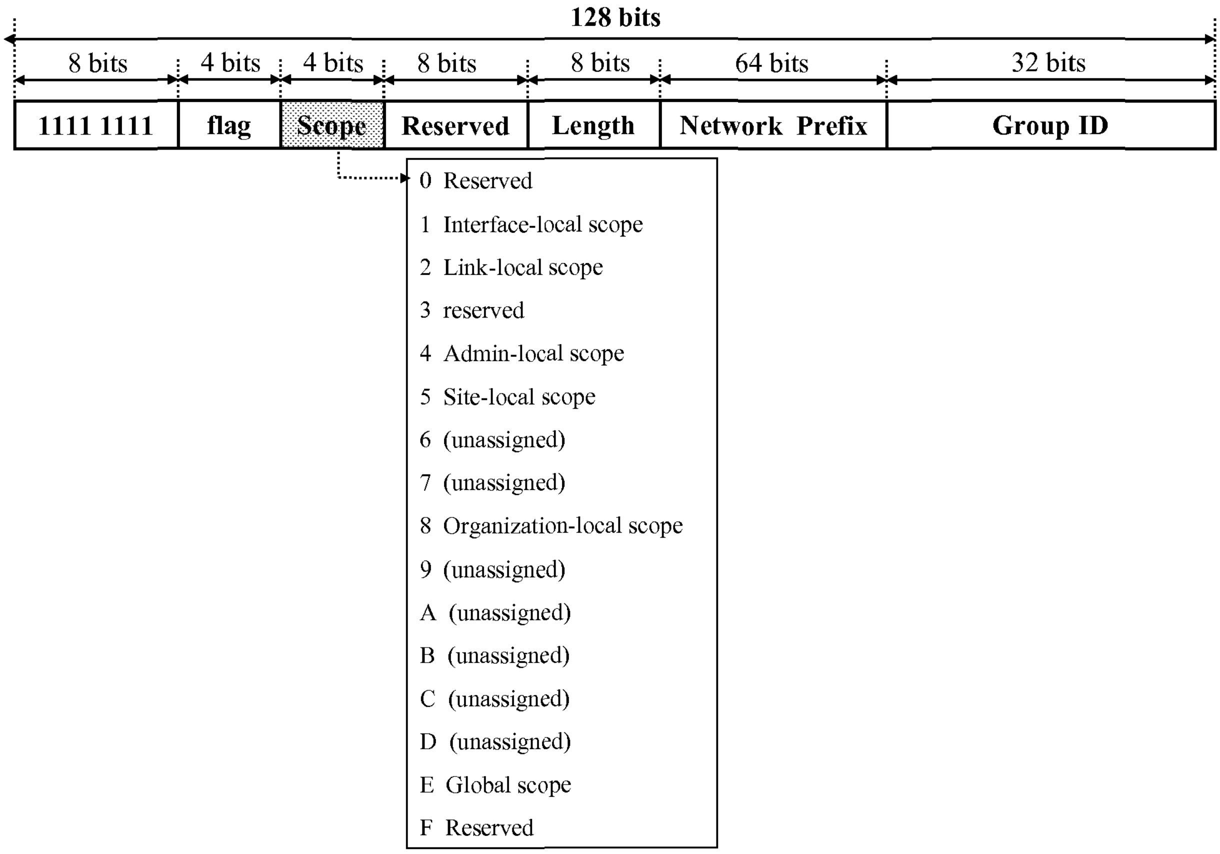
128 bits
8 bits
4 bits
4 bits
8 bits
8 bits
64 bits
32 bits
1111 1111
flag
Scope
Reserved
Length
Network Prefix
Group ID
0 Reserved
1 Interface-local scope
2 Link-local scope
3 reserved
4 Admin-local scope
5 Site-local scope
6 (unassigned)
7 (unassigned)
8 Organization-local scope
9 (unassigned)
A (unassigned)
B (unassigned)
C (unassigned)
D (unassigned)
E Global scope
F Reserved

Fig. 2

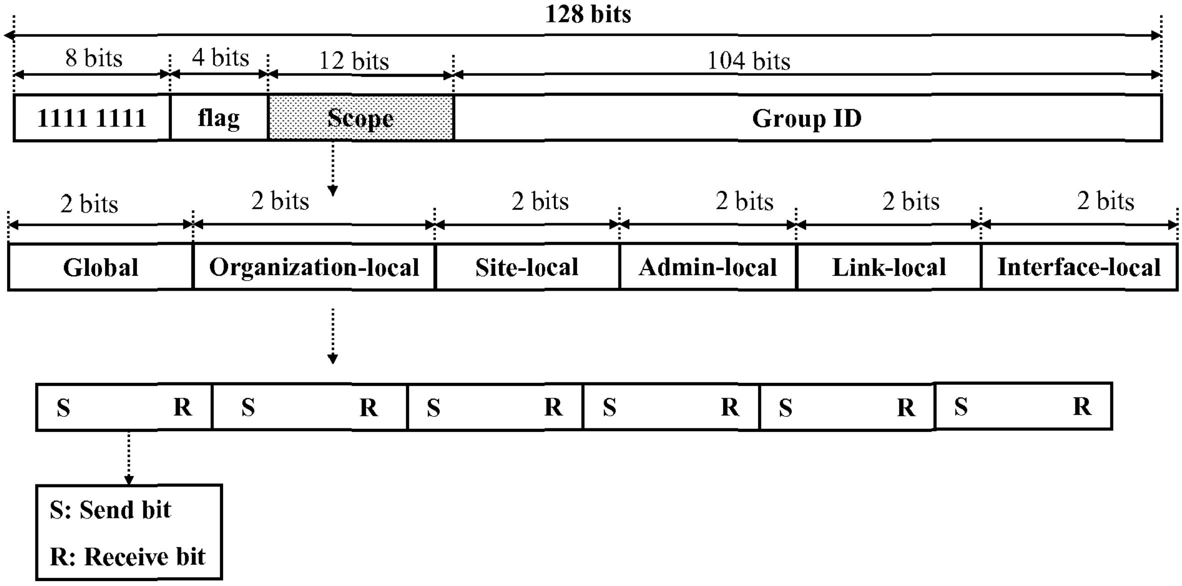
128 bits
8 bits
4 bits
12 bits
104 bits
1111 1111
flag
Scope
Group ID
2 bits
2 bits
2 bits
2 bits
2 bits
2 bits
Global
Organization-local
Site-local
Admin-local
Link-local
Interface-local
S R
S R
S R
S R
S R
S R
S: Send bit
R: Receive bit

Fig. 3

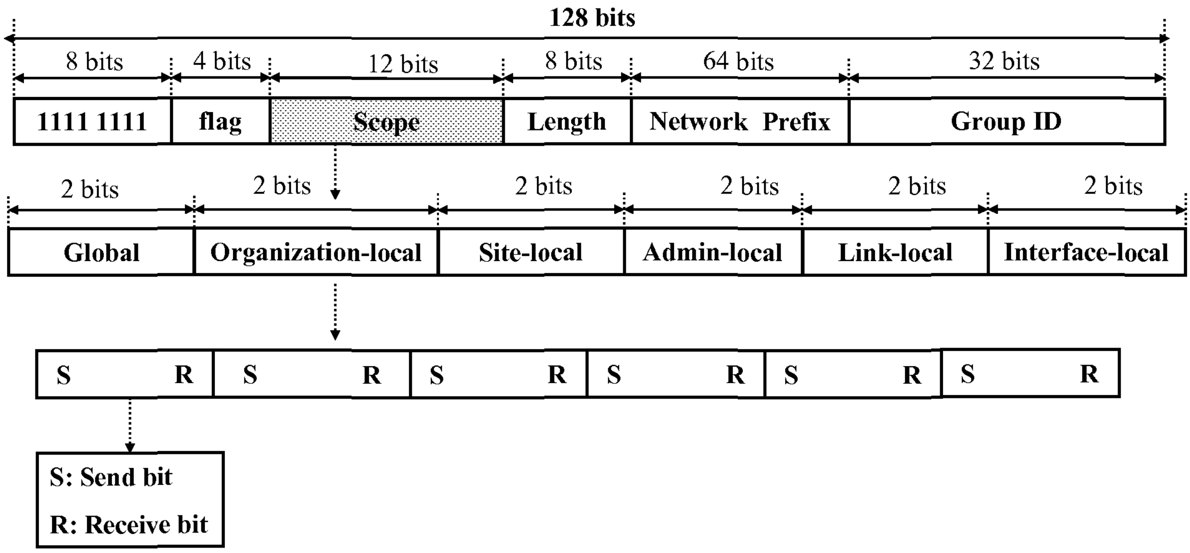
128 bits
8 bits
4 bits
12 bits
8 bits
64 bits
32 bits
1111 1111
flag
Scope
Length
Network Prefix
Group ID
2 bits
2 bits
2 bits
2 bits
2 bits
2 bits
Global
Organization-local
Site-local
Admin-local
Link-local
Interface-local
S R
S R
S R
S R
S R
S R
S: Send bit
R: Receive bit

Fig. 4

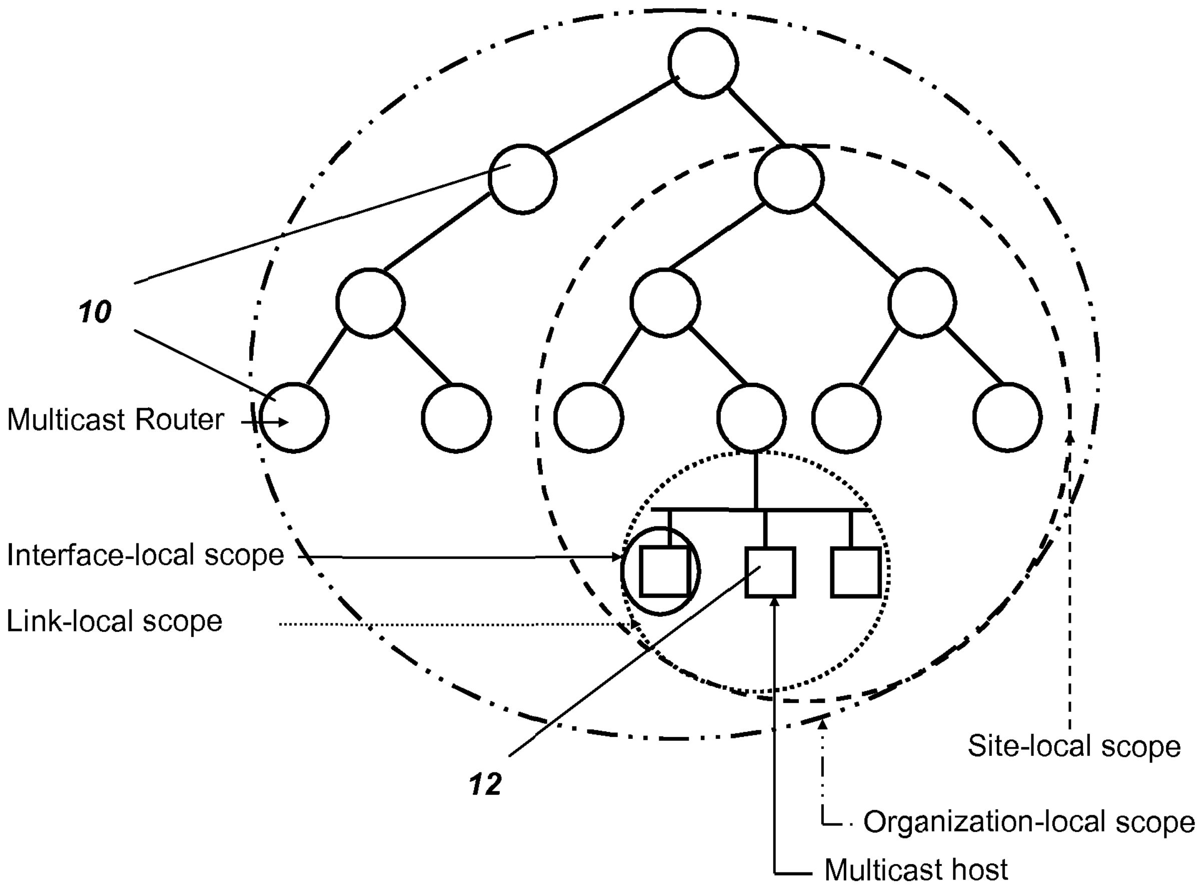
10
Multicast Router
Interface-local scope
Link-local scope
12
Site-local scope
Organization-local scope
Multicast host

Fig. 5

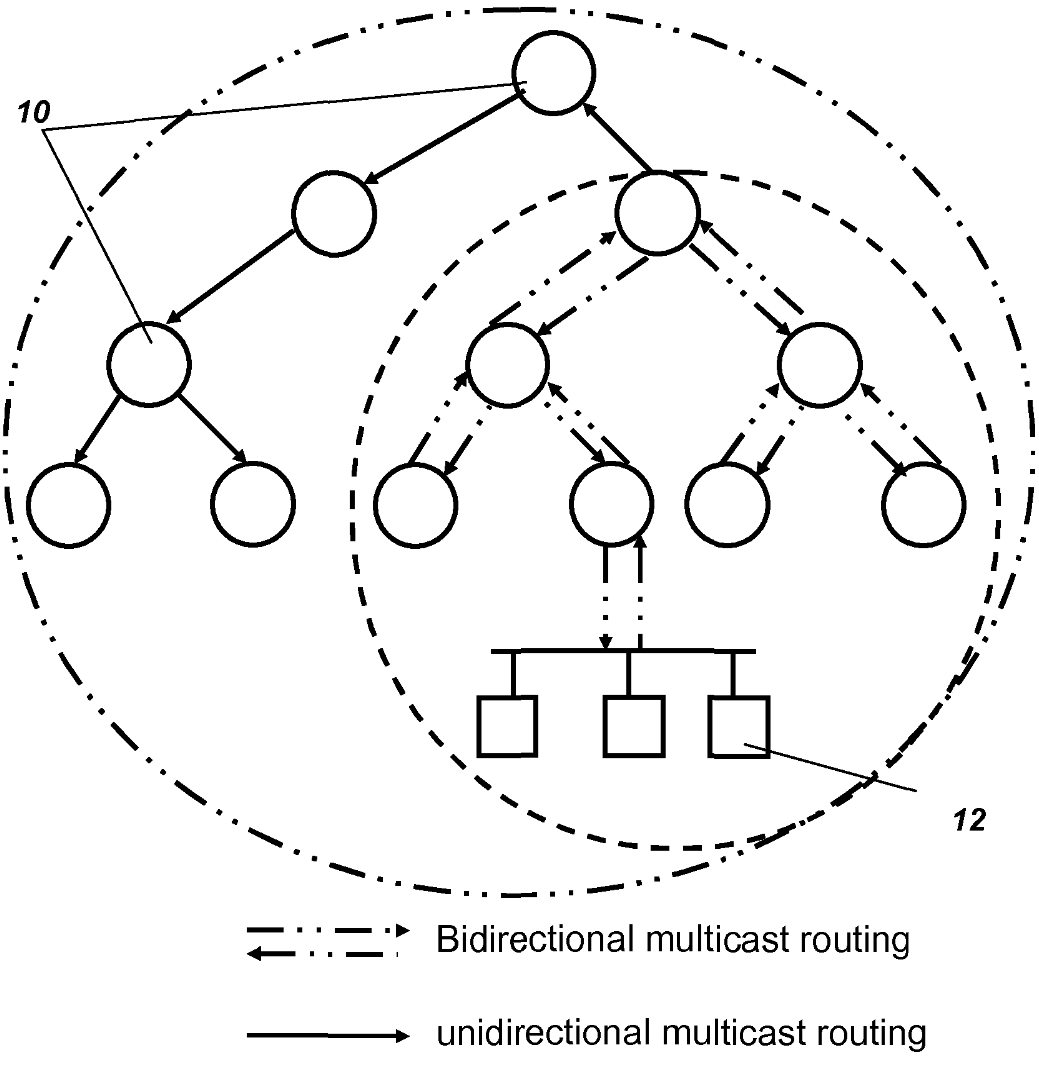
10
12
Bidirectional multicast routing
unidirectional multicast routing

Fig. 6

Many-to-Many Embedded Multicast Addressing and Access Control Method

The present invention relates to a many-to-many embedded multicast addressing and access control method, an important example being multicast addressing for packet switched networks such as the Internet.

When the various components that make up computer networks for the communication of data were designed, the primary objective was to enable point-to-point traffic. However there is now an increasing need for multipoint or group communications between different parties, for which purpose various multicast technologies have been developed. In the context of the Internet, IP multicast was designed as an efficient network layer protocol that handles group communication requirements.

Within the context of multicast, various types of multicast communication can be differentiated, depending on the number of senders and receivers. Two main general categories of multicast applications exist: one-to-many and many-to-many. In one-to-many multicast applications, a single host sends multicast data to a set of receivers. Examples include scheduled audio/video distribution, push media (news headlines, weather updates, sports scores, etc.), file distribution and caching. However, in many-to-many multicast applications, any number of hosts are able to send to the same multicast group address, as well as receiving from it. Examples include multimedia conferencing, shared distributed databases, distributed parallel processing, shared document editing, distance learning, chat groups, and multi-player games. In each of these applications, each member of the multicast group may receive data from multiple senders while it can also send data to all of the other members of the multicast group.

Existing multicast address allocation schemes do not provide or define a common strategy to allocate one-to-many and many-to-many multicast addresses. Moreover, there are no mechanisms for hosts to indicate which of these types of addresses they have to allocate to fulfil their multicast application requirements so that a multicast address allocation system responds accordingly. Presently, IP multicast relies on two main membership models: Any-Source Model (ASM) and Source-Specific Model (SSM). These models have different multicast address ranges and different terminologies. In fact, IP multicast defines a special IP multicast address to identify the group of interested receivers. Senders (multicast sources) send to the multicast address without prior knowledge of the multicast receivers. IP multicast does not require senders to a group to be members of the group.

To set-up a multicast session and distribute the multicast data, the group of interested receivers should be identified. A multicast address is a specific IP address used to identify a set of hosts to deliver IP packets to. The multicast address can be allocated from a specific range of IP addresses dedicated for sending to groups.

The IETF has defined guidelines that explain how to assign and allocate IP multicast address for both the ASM and the SSM models. A large number of multicast applications are restricted either to a link or a site, and it is extremely undesirable to propagate them further (beyond the link or the site). To deal with this, each IP address has a specific scope, which limits the flooding of multicast packets. An address scope defines a region or span where an address can be defined as a unique identifier of a network interface.

Within IPv4, the Class D address range (224.0.0.0 - 239.255.255.255), is reserved for multicast IP addresses. Multicast address scoping is achieved in IPv4 either by using the “Time To Live” field in the IPv4 packet header, or by an administratively scoped multicast address allocation strategy under 239.0.0.0/8, as set out in RFC 2365.

Within IPv6, the scope field is an integrated part of the multicast address itself. The most commonly used scopes are,link-local, site-local and global. The encoding of the scopes in the IPv6 multicast address prefixes is set out in RFC4291.

To allocate a multicast address, two mechanisms can be used. The first mechanism is a centralized one where the allocation is carried out by an authorized entity. Hence, a multicast address has to be requested from this authority and cancelled when the multicast session ends.

The second mechanism is a distributed one. The distributed allocation is done locally through multicast address allocation servers and requires specific protocol like the Multicast Address Dynamic Client Allocation Protocol (MADCAP) (RFC 2730). Compared to the centralized mechanism, the distributed mechanism does not guarantee the uniqueness of the multicast address, but it is more flexible than the centralized mechanism.

To avoid confusion and reduce the probability of IP multicast address collision, the IPv6 multicast address architecture has been revised and new extensions have been introduced in order to simplify the dynamic allocation of multicast addresses and embed IPv6 unicast prefixes in multicast addresses (see RFC 3306 and RFC 3307). By delegating multicast addresses at the same time as unicast prefixes, network operators will be able to identify their multicast addresses without running an inter-domain allocation protocol. Once a multicast address is allocated, the multicast source or the group manager describes the multicast session to be launched using the multicast Session Description Protocol (SDP) (RFC 3266 and RFC 4566) and can advertise this description using the multicast Session Announcement Protocol (SAM) (RFC 2974) or other advertisement methods.

To receive multicast traffic, an interested receiver requires a mechanism to join the multicast group. The receiver accomplishes this task by using a membership protocol such as IGMP (Internet Group Management Protocol for IPv4 capable hosts) (RFC 3376) or MLD (Multicast Listener Discovery Protocol for IPv6 capable hosts) (RFC 3810). To build a multicast distribution tree that spreads from the senders to all receivers, multicast capable routers need to deploy a multicast routing protocol to handle the duplication of multicast traffic and conveying multicast packet across the built tree (RFC 5110).

Several multicast routing protocols have been proposed for use on the Internet. Since the early routing protocols such as DVMRP and MOSPF were designed to handle dense multicast groups, new other protocols have been proposed to offer better scalability. Sparse-mode protocols like PIM-SM provide efficient multicast communication between members that are sparsely distributed. Such protocols use a single unidirectional shared tree that spans all members of a group. Consequently, multicast traffic for each group is sent and received over the same delivery tree, regardless of the source. Compared to shared tree protocols, SSM routing protocols construct a specific delivery tree per source and group address pair. For collaborative many-to-many applications where the receivers could be sources as well, multicast routers require to build a bidirectional multicast delivery tree to optimize the multicast routing and avoid using multiple instances of unidirectional delivery trees. Currently, bidirectional PIM and CBT protocols are designed to construct bidirectional multicast delivery trees (RFC 5110).

In order to distinguish between one-to-many and many-to-many addresses, it is necessary to use a higher level protocol, such as SDP or SAP. For instance, using the SDP protocol a source can describe for multicast receivers whether they could interactively participate in the multicast group or not. For this purpose the source can set a send-receive attribute on its session description messages ("a=sendrecv"). The session description will be dispatched to all the multicast receivers by the SAP protocol. Unfortunately, multicast routers are not notified about this multicast group feature as they don't process SAP announcement messages that contain session descriptions. To overcome this shortcut, a network administrator currently has to use additional security access lists to associate multicast ranges to either unidirectional multicast routing protocols (PIM-SM, PIM-SSM, etc.) or bi-directional routing protocols in order to avoid unauthorized access. However this is a cumbersome and administratively complex method and it would be advantageous to have a more elegant way of distinguishing between one-to-many and many-to-many addresses.

STATEMENTS OF INVENTION

According to the present invention there is provided a multicast address scheme comprising a scope field that simultaneously includes information relating to a plurality of distribution scopes, and wherein the scope field further includes access rights for each distribution scope.

16 03 12

Optionally, the multicast address is an IP multicast address.

Optionally, the multicast address is an IPv6 multicast address.

Optionally, the multicast address uses extra bits from the reserved field as set out in RFC 4291; alternatively, the multicast address uses extra bits from the Group ID field set out in RFC 3306.

Optionally, the multicast address is an IPv4 multicast address.

Optionally, the plurality of distribution scopes are each defined by a plurality of bits of the scope field.

Optionally, the scope field is of a 12-bit size.

Optionally, the scope field comprises six ordered hierarchies, which optionally are defined as interface-local, link-local, admin-local, site-local, organisation-local, and global.

Optionally, each hierarchy is identified by two consecutive bits.

Optionally the set of hierarchies and their associated access rights are carried separately in specific routing control packets or embedded in the multicast address format.

According to a second aspect of the present invention there is provided a multicast controller which is capable to interpret or assign datagram destinations according to any of the above methods.

Optionally, the controller can dynamically auto-configure itself to distinguish between one-to-many and many-to-many IP multicast addresses.

Optionally, the controller can dynamically auto-configure itself to use either a unidirectional or bidirectional multicast routing protocol per distribution scope basis or per integrated multicast group address basis.

The multicast controller is optionally a router.

In further aspects a computer programme product is provided which carries instructions that, when run on a computer, enable implementation of the above.

BRIEF DESCRIPTION OF DRAWINGS

The present invention will now be described, by way of example only, with reference to the accompanying drawings, in which:

Fig. 1 illustrates an IPv6 multicast addresses format with respect to RFC 4291;

Fig. 2 illustrates a new address scope, which is compliant with RFC 3306;

Fig. 3 illustrates a unicast-prefix-based IPv6 multicast address format with respect to RFC 4291;

Fig. 4 illustrates a new address scope, which is compliant with RFC 3306;

Fig. 5 illustrates part of an embodiment of new hierarchical distribution address scopes; and

Fig. 6 illustrates an example of hierarchical multicast routing for the same multicast group.

DETAILED DESCRIPTION

This disclosure discusses innovations that are applicable to any multicast addressing scheme, so long as the set of receivers can be clustered into sub-groups or scopes and where the whole group can be identified by a unique address. The multicast addressing scheme can be IP based or non-IP based. For the purpose of explaining the addressing scheme of the invention, the present disclosure will discuss IPv6 and IPv4 embodiments of an addressing scheme. However these specific examples are not intended to limit the scope of protection afforded by any claims to be granted on this application. It is to be understood that the principles discussed herein can be applied to non-IP multicast addresses, to all different versions and implementations of IPv4 and IPv6 addressing schemes and to future versions of IP addressing schemes, and to addressing schemes for any other packet-switched networks. Furthermore, the disclosure and examples contained herein provides a person skilled in the art with sufficient teaching and direction to implement the invention in any multicast addressing scheme.

A new multicast addressing scheme is disclosed which has an expanded scope field, which enables a plurality of distribution scopes to be embedded simultaneously within the multicast address. They can also be embedded consecutively.

Within IPv6, the scope field is defined by the standards RFC 4291, RFC 3306 and RFC 3307. With respect to IPv6, a new multicast addressing scheme can include an IPv6 scope field that is extended beyond the standard 4-bit value.

Fig. 1 shows the conventional IPv6 multicast address format with respect to the specification of RFC 4291. The sixteen values of the scope field are shown.

Fig. 2 shows the conventional IPv6 multicast address format with respect to the specification of RFC 3306. The sixteen values of the scope field are shown.

Fig. 3 shows an embodiment wherein the scope field is extended to a 12-bit value in accordance with RFC 4291, with the additional eight bits being borrowed from the Group ID field. It is to be appreciated that in other embodiments the format, the length, and the placement of the scope field could be changed, and distribution scopes could be added, removed, reordered, or changed.

Fig. 4 shows an embodiment wherein the scope field is extended to a 12-bit value in accordance with RFC 3306, with the additional eight bits being borrowed from the reserved field. Again it is to be appreciated that in other embodiments the format, the length, and the placement of the scope field could be changed, and distribution scopes could be added, removed, reordered, or changed.

As mentioned above, the expanded scope field means that a plurality of distribution scopes can be embedded simultaneously within the multicast address. The size of the scope field will govern how many hierarchies can be embedded. As an example, a 12-bit value scope field can be used to provide a basis for the definition of six ordered hierarchies, with two bits being allocated for the identification of each of the six hierarchies.

The definition of the hierarchies will depend on network topology and the structures of user groups for a given organisation, but as an example where there are six hierarchies these could relate to *interface-local*, *link-local*, *admin-local*, *site-local*, *organization-local*, and *global* scopes (these terms have a self-evident meaning, for example an "*interface-local*" scope is intended to mean a scope that is local to a particular interface).

The standard IPv6 addressing format defines one distribution scope per multicast address, which is explicitly coded in the multicast address by using all the 4 bits. However in the new format, a plurality of distribution scopes are consecutively and simultaneously coded in each multicast address. The order of coding defines and differentiates implicitly each hierarchical scope. This method improves the multicast address allocation mechanism by avoiding allocating a multicast address per a single scope basis. In the conventional allocation method, the distribution scopes are mutually exclusive, while in the new method the distribution scopes could collocate and they are independent.

An expanded scope field can also be implemented in IPv4, by modifying the format of the Time to Live (TTL) field of the IPv4 packet header. Other methods could be used to embed or carry the new expanded scope format either in IPv4 data packets or in any specific multicast routing control messages. Examples of implementation could include carrying the new format as options or attributes in the IPv4 packet header. Alternatively, administratively scoped IPv4 multicast addresses can be used in conjunction with either a modified Time to Live field of the IPv4 data multicast packet header or any specific other options, attributes or separate multicast routing control messages.

The expanded scope field could be implemented either in the host part or in a multicast address allocation system such as the Multicast Address Dynamic Client Allocation Protocol (MADCAP, RFC 2730). In the first scenario, multicast hosts can dynamically auto-configure themselves and manage the multicast access right of each scope they create without a third party involvement. When a multicast address allocation system is used an extra negotiation signalling protocol is needed between the client and the server. The client is the host requesting a multicast address and the server is the entity responsible for allocating the multicast address. The client and the server require confirming the distribution scopes to be created and the access rights for each scope.

It is also possible in the new multicast addressing scheme to define access right modes for each scope. Conventionally, the access right for a given multicast group is defined by external mechanisms such as security access lists and authentication and authorizations methods which are configured by the network administrator. However in the new format, the access rights are encoded in the multicast address itself.

There are many ways in which the access rights can be encoded. In one embodiment, the access rights can be encoded by using 2 bits that define send and receive permissions. The binary combination of these two bits allows four permission rights. One bit (for example the highest) can be reserved for a “transmission” right and one bit (for example the lowest) can be reserved for a “reception” right. Therefore, within each scope multicast receivers may be assigned with one of the following access rights: denial of access, receive only, send only, or send and receive simultaneously. The numeric binary value of each bit is either one or zero.

These are illustrated in Figs. 3 and 4, where each of the scopes is of a 2 bit size with the bits encoding data regarding the send and receive permissions for each of the scopes.

Like the expanded scope field, the access rights definition could be implemented either in the host part or in a multicast address allocation system, as discussed above.

A new multicast controller can be provided which implements the new multicast addressing scheme. The multicast controller can for example be a router or a switch for forwarding packetized data.

A controller thus modified can distinguish between one-to-many and many-to-many multicast groups by analysing the format of the scope field. When the “transmission” and the “reception” access rights are set simultaneously for a specific distribution scope or for all scopes (i.e. each bit is set to one), the controller can differentiate whether the multicast data require to be transmitted using unidirectional or bidirectional delivery method or not for that specific distribution scope or for all scopes.

Furthermore, a controller thus modified can dynamically auto-configure itself by using one or multiple optimal multicast routing protocols, either per integral multicast group basis or per scope basis. When different multicast routing protocols are used simultaneously for the same multicast group address, controllers within a specific scope can choose the same routing protocol.

The selection of a given multicast routing protocol within a given scope depends on the access rights for this scope. The multicast routing protocols for all the distribution scopes could be also identical.

Therefore the modified controller introduces hierarchical scope-based multicast routing where unidirectional and bidirectional multicast routing protocols can be combined and used simultaneously to construct a multicast delivery tree for a given multicast group. Multicast routing protocols can be triggered per scope basis or per whole address basis. Thanks to the embedded access rights, multicast controllers can set the appropriate security access rights to the multicast group.

Examples of such hierarchical scopes are shown in the accompanying figures, where Fig. 5 shows some example hierarchical multicast scopes, and Fig. 6 shows an example of hierarchical multicast routing for the same multicast group.

Referring to Fig. 5, a network comprises a number of multicast controllers 10, each of which can be in connection with one or more multicast hosts 12 (the hosts 12 are shown for only one of the routers 10 for ease of illustration). As shown in this example, an “interface-local” scope is unique to an individual host 12; a “link-local” scope is unique to a set of hosts 12 which are associated with a single multicast controller 10; a “site-local” scope is unique to a set of controllers 10 which are linked as having a common site (the definition of a “site” in the real world will depend on the particular infrastructure of a given organization); an “organization-local” scope is unique to a set of controllers 10 which are linked as having a common organization, and optionally also that set of controllers 10 can also correspond to a set of sites (again, the definition of an “organization” in the real world will depend on the particular network structure). The “administration-local” and “global” scopes are not shown for ease of illustration, but it will be appreciated that these scopes can refer to various others points on the network.

Fig. 6 shows an example of hierarchical multicast routing for the same multicast group as shown in Fig. 5. In one of the sites that constitute the organisation, a bidirectional multicast routing protocol is used. However in the rest of the organisation, a unidirectional multicast routing protocol is used.

When the multicast controller is a router, it can auto-decide by itself whether it will trigger a multicast routing protocol per scope basis or per whole multicast address basis. For this purpose, the multicast router could be preconfigured to enable one of these modes (distribution scope basis or whole group basis).

In other cases, the multicast controller can be ordered by a third party to trigger a specific multicast routing protocol per scope or per group basis. The third party could be another multicast controller, a configuration server, etc.

It will be appreciated that the features discussed above give rise to many advantages. The expanded scope field means that a standardized allocation or assignment method of multicast addresses that differentiates between one-to-many and many-to-many address ranges can be provided. In addition, the encoding of all scopes in the expanded scope field avoids the need to implement complex security mechanisms such as security access lists and avoids the need for a third party to filter unauthorized multicast receivers and sources for the matching and binding of specific ranges of multicast addresses to specific multicast routing protocols, as is required by multicast routing policies.

Furthermore, by encoding both distribution scopes and the associated access rights in the multicast address itself, the invention helps to auto-configure multicast capable routers with the appropriate multicast routing protocols and the access rights and therefore ease network management. Multiple multicast routing protocols can be deployed for the same multicast group instead of using a single multicast routing protocol for all hierarchical distribution scopes in association with many confusing security access lists to perform the same results. Also, multicast routers within an administrative domain or in the Internet will share the same view of what type of multicast routing protocol to be used (for the scope and for the whole multicast group) and what are the associated access rights to be verified. The new methods (and associated apparatus as mentioned above) are simple to implement and deploy and they ease autonomic multicast capable router configuration, in contrast to conventional methods are complex to deploy and they are not homogenous.

In the present disclosure, references to “RFC” documents refer to the Request for Comments library published by the IETF (Internet Engineering Task Force), which is well known to those skilled in the art. The contents of all such documents referred to are herein incorporated by reference, to the fullest extent as allowed by law.

Various improvements and modifications can be made to the above without departing from the scope of the invention. For example, it is to be appreciated that the extension, the modification, the elimination, or the alteration of either the access right modes or the distribution scopes will not change the basic invention. In addition the implementation of this invention in any IP or non IP based multicast addressing scheme will not affect the present invention.

CLAIMS

1. A multicast address scheme comprising a scope field that simultaneously includes information relating to a plurality of distribution scopes, and wherein the scope field further includes access rights for each distribution scope.

2. The multicast address scheme of claim 1, wherein the multicast address is an IP multicast address.

3. The multicast address scheme of any preceding claim, wherein the multicast address is an IPv6 multicast address.

4. The multicast address scheme of any preceding claim, wherein the multicast address uses extra bits from the reserved field as set out in RFC 4291.

5. The multicast address scheme of any of claims 1 to 3, wherein the multicast address uses extra bits from the Group ID field set out in RFC 3306.

6. The multicast address scheme of claim 1 or claim 2, wherein the multicast address is an IPv4 multicast address.

7. The multicast address scheme of any preceding claim, wherein the plurality of distribution scopes are each defined by a plurality of bits of the scope field.

8. The multicast address scheme of any preceding claim, wherein the scope field is of a 12-bit size.

9. The multicast address scheme of any preceding claim, wherein the scope field comprises six ordered hierarchies.

10. The multicast address scheme of claim 9, wherein the ordered hierarchies are defined as interface-local, link-local, admin-local, site-local, organisation-local, and global.

11. The multicast address scheme of claim 9 or claim 10, wherein each hierarchy is identified by two consecutive bits.

12. The multicast address scheme of any of claims 9 to 11, wherein the set of hierarchies and their associated access rights are carried separately in specific routing control packets or embedded in the multicast address format.

13. A multicast controller which is capable to interpret or assign datagram destinations according to any of the above methods.

14. The multicast controller of claim 13, being able to dynamically autoconfigure itself to distinguish between one-to-many and many-to-many IP multicast addresses.

15. The multicast controller of claim 13 or claim 14, being able to dynamically auto-configure itself to use either a unidirectional or bidirectional multicast routing protocol per distribution scope basis or per integrated multicast group address basis.

16. The multicast controller of any of claims 13 to 15, being a router.

17. A computer programme product that carries instructions that, when run on a computer, enable implementation of any preceding claim.

16 03 12