

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
6 March 2003 (06.03.2003)

PCT

(10) International Publication Number
WO 03/019906 A2

(51) International Patent Classification⁷: **H04L 29/12**

(21) International Application Number: PCT/GB02/03957

(22) International Filing Date: 27 August 2002 (27.08.2002)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
01307239.2 24 August 2001 (24.08.2001) EP

(71) Applicant (for all designated States except US): **BRITISH TELECOMMUNICATIONS PUBLIC LIMITED COMPANY** [GB/GB]; 81 Newgate Street, London EC1A 7AJ (GB).

(72) Inventors; and

(75) Inventors/Applicants (for US only): **KING, John, Robert** [GB/GB]; 2 Hertfords Place, Chillesford, Woodbridge, Suffolk IP12 3SD (GB). **MOTTRAM, Stephen, Willis** [GB/GB]; 15 Brooks Hall Road, Ipswich, Suffolk IP1 4BY (GB).

(74) Agent: **WALLIN, Nicholas, James**; BT Group Legal Services, Intellectual Property Department, Holborn Centre, 8th floor, 120 Holborn, London EC1N 2TE (GB).

(81) Designated States (*national*): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, OM, PH, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, US, UZ, VN, YU, ZA, ZM, ZW.


(84) Designated States (*regional*): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, SK, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

— without international search report and to be republished upon receipt of that report

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

(54) Title: APPARATUS AND METHOD OF COORDINATING NETWORK EVENTS



(57) Abstract: Embodiments of the present invention relate to coordinating address allocation events at a routing device and at a host machine, where the routing device is arranged to route data between the host machine and at least one other host machine. The routing device may be a router and the host machine may be a computer located in a first network, the first network operating in accordance with a first transmission protocol. The other host machine may be a computer, located in a second network, the second network operating in accordance with a second transmission protocol. The method is particularly suited to situations where the host machine is running one or more applications in accordance with the second transmission protocol. Embodiments of the invention are concerned with coordinating address allocation events that are required to enable data of the second transmission protocol to be transmitted through the first network and on to the other host machine. The method comprises the following steps: receiving an allocated address, for use by the host machine in communicating with the other host machine; receiving an indicator representative of allocation status of the allocated address; sending, to the routing device, the allocated address together with the address of the host machine; storing a mapping between the allocated address and the address of the host machine; monitoring the indicator, and, if the indicator indicates that the allocated address is no longer valid for the host machine, rendering the stored mapping unusable.

APPARATUS AND METHOD OF COORDINATING NETWORK EVENTS

This invention relates to coordination of network events, and has particular, but not exclusive, application in coordinating Internet Protocol address management
5 events.

Currently all commercial Internet Protocol (IP) networks are IP version 4 (IPv4) networks; however, at some point in the future, commercial IP networks will be IP version 6 (IPv6) networks. In the meantime there will be a transitory period, during which commercial IP networks will comprise a mixture of IPv4 and IPv6 networks.
10 IPv6 is a totally different protocol to IPv4 and is fundamentally incompatible with IPv4. Therefore, during the transitory period at least, network devices and/or networks will require mechanisms to enable a node and/or host in an IPv4 network, having an IPv4 address, to communicate with a node and/or host in an IPv6 network, having an IPv6 address.

15 Several migration mechanisms have been developed; see for example a document published in November 2000 by the Internet Engineering Task Force (IETF) and available from the IETF at <http://www.ietf.org/internet-drafts/draft-ietf-ngtrans-introduction-to-ipv6-transition-05.txt>, entitled "An Overview of the Introduction of IPv6 in the Internet", authors: W. Biemolt et al, IETF Status: Draft working towards
20 Informational RFC.

One particular migration method, known as Dual Stack Transition Mechanism (DSTM), enables hosts to send and receive both IPv4 and IPv6 data. This mechanism is typically used when a host, located in an IPv6 network, is running one or more IPv4 applications, and thus needs to communicate with hosts located in an
25 IPv4 network. Such a host is hereinafter referred to as a dual stack host. A feature of the DSTM method is that IPv4 packets are encapsulated into IPv6 packets when they leave the dual stack host, and are subsequently un-encapsulated by a router, which is located on the boundary of the IPv4 and IPv6 networks.

The dual stack host is assigned an IPv4 address, which is used as an alias
30 for the dual stack host and forms the source address of the encapsulated packet. This assigned IPv4 address is "leased" to the dual stack host for a specified period, after which it can either be renewed or released for use by another dual stack host.

A problem with the DSTM method is that the renewal and release information is typically not conveyed to the router in a timely fashion, so that packets can get misdirected, and/or dropped and snooped en route for their intended dual stack host destination.

5 According to a first aspect of the present invention there is provided a method of coordinating address allocation events at a routing device and at a host machine, where the routing device is arranged to route data between the host machine and at least one other host machine, in which the host machine is located in a first network, the first network operating in accordance with a first transmission
10 protocol, and in which the other host machine is located in a second network, the second network operating in accordance with a second transmission protocol, and wherein the host machine is operable to process packets of data corresponding to either transmission protocols,. The method comprises the steps of:

 receiving an allocated address, for use by the host machine in
15 communicating with the other host machine;

 receiving an indicator representative of allocation status of the allocated address;

 sending, to the routing device, the allocated address together with the address of the host machine;

20 storing a mapping between the allocated address and the address of the host machine;

 monitoring the indicator, and, if the indicator indicates that the allocated address is no longer valid for the host machine,

 rendering the stored mapping unusable.

25 Preferably the method includes using the mapping to establish a tunnel between the routing device and the host machine, the tunnel being used in the said routing of data between the host machine and the other host machine.

 Conveniently the method further includes requesting renewal of the allocated address if the indicator indicates that the allocated address is no longer
30 valid for the host machine.

 Advantageously the method includes checking at least one interface of the host in order to determine whether the allocated address is to be renewed. This

essentially provides a means of establishing whether the host is in the course of sending and/or receiving data.

Conveniently the rendering step causes the stored mapping to be deleted from the routing device, so that the tunnel is removed.

5 According to a second aspect of the invention there is provided apparatus for carrying out the afore-described method. The apparatus is preferably distributed between devices and the routing device, so that at least some of the functionality provided by the apparatus may be located on the host machine; on the routing
10 device; on a bespoke device, which is located in the first network; or distributed between a device that is arranged to issue the allocated address and the host machine. The device that is arranged to allocate addresses to requesting host machines is preferably a dynamic host configuration protocol server.

In the following description the term "host" is used; this is defined as follows:

15 "host" – any computer that has two-way access to other computers in a network such as the Internet or an Intranet. Examples of hosts include clients, routers, switches, and servers.

Further aspects and advantages of the present invention will be apparent from the following description of preferred embodiments of the invention, which are given by
20 way of example only and with reference to the accompanying drawings, in which

Figure 1 is a schematic diagram of a network, within which embodiments of the invention operate;

Figure 2 is a flow diagram showing operation of the known Dual Stack Transition Mechanism;

25 Figure 3 is a schematic diagram of components of a device comprising part of the network of Figure 1;

Figure 4 is a schematic flow diagram showing processes involved in communication between a dual stack host, located in a first network, and a host located in a second network, incorporating aspects of a first embodiment of the
30 invention;

Figure 5 is a schematic flow diagram showing processes involved in coordination aspects of a first embodiment of the invention;

Figure 6 is a schematic flow diagram showing further processes involved in coordination aspects of a first embodiment of the invention;

Figure 7 is a schematic flow diagram showing yet further processes involved in coordination aspects of a first embodiment of the invention;

5 Figure 8 is a schematic flow diagram showing further processes involved in coordination aspects of a first embodiment of the invention;

Figure 9 is a schematic flow diagram showing processes involved in coordination aspects of a first embodiment of the invention;

Figure 10 is a schematic flow diagram showing processes involved in further
10 coordination aspects of a first embodiment of the invention;

Figure 11 is a schematic flow diagram showing processes involved in aspects of address allocation without embodiments of the invention;

Figure 12 is a schematic flow diagram showing processes involved in further aspects of address allocation without embodiments of the invention,

15 Figure 13 is a schematic diagram showing the components of Figure 3 arranged according to another embodiment of the invention;

Figure 14 is a schematic flow diagram showing processes involved in communication between a dual stack host, located in a first network, and a host located in a second network, incorporating aspects of the embodiment of Figure 13;

20 Figure 15 is a schematic flow diagram showing further processes involved in coordination aspects of the embodiment of Figure 13;

Figure 16 is a schematic flow diagram showing yet further processes involved in coordination aspects of the embodiment of Figure 13; and

Figure 17 is a schematic diagram showing the components of Figure 3
25 arranged according to yet another embodiment of the invention.

Embodiments of the invention are concerned with issues relating to migration from IPv4 to IPv6 networks. Specifically, embodiments of the invention are concerned with enabling communications between an IPv4 application, which is
30 running on a host located within an IPv6 network, and a corresponding IPv4 application running within an IPv4 network.

One embodiment of the invention is concerned with a transition mechanism known as Dual Stack Transition Mechanism (DSTM), which is documented by the

Internet Engineering Task Force (IETF) in document draft-ietf-ngtrans-dstm-04.txt (Note that this document is referenced by its given title at August 2001; as is common with IETF publications, the title may subsequently change, so the reader should refer to the IETF in case of doubt), available from the IETF. Hosts that run
5 DSTM have both an IPv4 and an IPv6 stack therein, meaning that data to be transported in accordance with the IPv4 protocol passes through the IPv4 stack (or layer), and data to be transported in accordance with the IPv6 protocol passes through the IPv6 stack, for processing thereby.

Accordingly, incoming and outgoing IPv4 and IPv6 packets are respectively
10 routed via an IPv4 or an IPv6 part of the stack. Thus when operating in accordance with DSTM, a host running DSTM can be located in an IPv6 network and retain its IPv4 functionality.

Figure 1 shows an implementation of DSTM in operation between a first network NW1, operating in accordance with a first transmission protocol and a
15 second network NW2, operating in accordance with a second transmission protocol. In this example the first network is an IPv6 network and the second network is an IPv4 network.

Conventional operation of DSTM is shown in Figure 2. As is known in the art, when a dual stack host H1 in the IPv6 network NW1 participates in IPv4
20 communication, an IPv4 address is dynamically allocated to the dual stack host H1 by an address pool. The address pool is provided by IPv6 Dynamic Host Configuration Protocol (DHCPv6) server 103, which stores a pool of IPv4 addresses that it leases to other machines in the network upon request.

Accordingly, at step S 2.1 the dual stack host H1 makes a request to the
25 DHCPv6 server for an IPv4 address. At step S 2.2 the DHCPv6 server 103 responds by allocating an IPv4 address to the dual stack host H1, and then, at step S 2.3, sends the allocated IPv4 address to the dual stack host H1. This allocated address is accompanied by two timeout values: *preferred* and *valid*, which are respectively renewable and non-renewable address timeouts. Upon expiry of a *preferred* type
30 timeout, if the dual stack host H1 wishes to continue participating in outgoing communications, it must request renewal of the address allocated thereto; upon expiry of the *valid* type timeout the allocated address cannot be renewed for any

type of communication. If the preferred timer is renewed, both the *preferred* and the *valid* timers are simultaneously refreshed.

For the dual stack host H1 to send information to, or receive information from, one or more hosts located in one or more IPv4 networks, a tunnel, connecting
5 the dual stack host H1 with a Border Router BR, which is located between the two types of networks NW1, NW2, has to be created. As is known in the art, the phrase "tunnelling" is used to refer to a process whereby a packet is encapsulated within another packet; in the known method, as indeed in embodiments of the present invention, IPv4 packets are encapsulated within IPv6 packets. Thus a tunnel is the
10 means for performing the tunnelling process, and essentially comprises access points, termed end points and described below, for performing packet-encapsulation.

As is known to those in the art, a tunnel can be either manually or automatically configured; there are several known tunnelling methods, including 6to4, 6 over 4, dynamic tunnelling, tunnel brokering and a method developed by the
15 applicant, which is described in copending published patent application WO01/22683 (applicants' case ref A25800). For more information the reader is referred to the following documents "Request for comments number RFC2529"; draft-ietf-ngtrans-6to4-02.txt (or RFC3056); draft-ietf-ngtrans-dstm-04.txt; draft-ietf-ngtrans-broker-00.txt (or RFC3053), all available from the IETF. (Note that these
20 documents are referenced by their given titles at August 2001; as is common with IETF publications, the title may subsequently change, so the reader should refer to the IETF in case of doubt).

Accordingly at step S 2.4 the dual stack host H1 configures a first endpoint EP1 of an IPv6 tunnel 105. At step S 2.5 the dual stack host H1 sends its IPv6
25 address, together with the IPv4 address allocated by the DHCPv6 at step S 2.2, to the Border Router BR, to enable the Border Router BR to configure a second endpoint EP2 of the IPv6 tunnel 105 (when the dual stack host H1 sends its first packet destined for an IPv4 host H2).

Then at step S 2.6 the Border Router BR configures the second endpoint
30 EP2 of the IPv6 tunnel 105, so that IPv4 packets can be tunnelled through the IPv6 network NW1 between the two endpoints EP1, EP2.

Management of address allocation is essentially performed by the dual stack host H1, so that, if the preferred timeout has expired, and the dual stack host H1 is

still in the process of communicating with an IPv4 host H2 in the IPv4 network NW2, the dual stack host H1 has to initiate renewal of the address and timer.

The current DSTM draft (draft-ietf-ngtrans-dstm-04.txt) does not address timing issues, in particular coordinating the renewal and release of address allocation
5 between the dual stack host and the Border Router BR. A problem can arise if the dual stack host H1 has finished communicating with the IPv4 host H2 and has released the IPv4 address allocated at step S 2.2 without informing the Border Router BR. If the DHCPv6 server 103 thereafter allocates the IPv4 address to another dual stack host Hd, and the Border Router BR has not modified the second
10 tunnel endpoint EP2, the second endpoint EP2 will be "pointing" to the wrong dual stack host (because the tunnel will point to the IPv6 address of host H1, instead of to the IPv6 address of the newly allocated host Hd).

Embodiments of the invention are therefore concerned with coordinating allocation and release of IPv4 addresses, specifically ensuring that tunnel endpoints
15 EP1, EP2 only exist for the lifetime of the IPv4 address allocation. This has a first advantage that, in the event of expiry of a timer or explicit cancellation of the address allocation on the part of the dual stack host H1, the tunnel endpoints EP1, EP2 are released, and a second advantage that, in the event that the dual stack host H1 renews the IPv4 address allocation upon expiry of the timer, embodiments ensure
20 that correct tunnel endpoints EP1, EP2 are retained.

In a first embodiment the coordination of address allocation between the dual stack host and Border Router BR, each forming a respective end of the tunnel
105, is controlled by the dual stack host at the first tunnel endpoint EP1. This has the benefit of minimising network traffic, because information is only sent to the
25 Border Router BR when there is a change to the tunnel. In other embodiments the coordination of address allocation between the dual stack host and Border Router BR is controlled by the Border Router BR or by an independent device located in the first network NW1.

Referring to Figure 3, a first embodiment of the invention will now be
30 discussed in more detail.

Figure 3 shows a dual stack host H1 comprising a central processing unit (CPU) 301, a memory unit 303, an input/output device 305 for connecting the host H1 to the network NW1, storage 307, and a suite of operating system programs

309, which control and coordinate low level operation of the host H1, and in particular handle operation of the IPv4 and IPv6 stacks. Such a configuration is well known in the art.

Generally embodiments of the invention are referred to as a coordinator 300,
5 and comprise at least some of programs 311, 313, 315. These programs are stored on storage 307 and are processable by the CPU 301.

The programs include a program 311 for monitoring timer status; a program 313 for requesting renewal, or re-allocation, of allocated address; and a program 315 for updating the Border router BR with address allocation information.

10 The monitoring program 311 monitors the preferred and/or valid timers to identify expiry thereof, the renewal program 313 interacts with the DHCPv6 server 103 to request renewal of a previously allocated IPv4 address; and the updating program 315 informs the Border Router BR of any changes that have been made to address allocation and/or timer information.

15 As stated above, in one embodiment the coordinator 300 could be located on the dual stack host H1 that is requesting and receiving the address allocation information from the DHCPv6 server 103, in a second embodiment the coordinator 300 could be located on a controller device 107 (see Figure 1) dedicated to managing address allocation and re-allocation, which acts as a kind of broker
20 between the dual stack host H1, the Border Router BR and the DHCPv6 server 103, and in a third embodiment the coordinator 300 could be located in the Border Router BR.

The operation of the coordinator 300, according to a first embodiment of the invention, will now be described with reference to the flowcharts shown in Figures 4
25 to 9. In the first embodiment the coordinator 300 is assumed to be located on the dual stack host H1.

The functionality of the invention is easiest to describe in the context of an end-to-end communications process, so Figure 4 shows the steps involved in establishing communications between a dual stack host H1 in an IPv6 network and
30 an IPv4 host H2, when communications are initiated by the dual stack host H1. Most of these steps are standard and are disclosed in document draft-ietf-ngtrans-dstm-04.txt referred to above, while others utilise the functionality of the coordinator 300. Figures 5 and 6 show steps involved in subsequent management, by the coordinator

300, of the established communications, and illustrate the inventive concept of the present invention.

In Figures 4, 5 and 6 (and in subsequent figures presenting information in this form) the steps are enumerated, some being accompanied by an arrow. The start and finish of an arrow indicates respectively the initiator and recipient of the communication, and the vertical position of an arrow indicates the position of a step corresponding thereto in the temporal sequence of events.

Considering firstly Figure 4, dual stack host H1 sends 402, in IPv6 format, an IPv4 Domain Name Server (DNS) request for the IP address of the host H2 to a DNSv6 server 111 located in the IPv6 network NW1. The DNS request is forwarded onto the Border Router BR, which sends 404 the DNS request to a DNSv4 server 109. This step 404 involves "relaying" the packet from the Border Router to the DNSv4 server 109.

Relaying can be used in any of the following example situations: where IPv6 DNS records are held on an IPv4 DNS server; where IPv4 DNS records are held on an IPv6 DNS server; or where an IPv6 DNS message requires passage over an IPv4 network in order to reach an answering DNS server (i.e. relaying enables processing of disparate types of DNS requests). In the latter scenario, the DNS server could be an IPv4 server or an IPv6 server. In the current example, shown in Figure 4, relaying is being used so that a dual stack host can, via IPv6 DNS messages, access DNS records being held on an IPv4 DNS server 109.

The DNS request arriving at the Border Router BR from the dual stack host H1 has originated from an IPv4 application. Thus the DNS request has an IPv4 DNS payload (containing the domain name of the IPv4 host H2 with which the dual stack host H1 wishes to communicate) and an IPv6 header, which is required to route the packet through the first network NW1. The Border Router BR translates the received IPv6 packet header into an IPv4 header for onward routing in the second network NW2 to the IPv4 DNS server 109. During this translation process the payload remains unaltered.

Once the DNS message has arrived at the DNSv4 server 109, the DNSv4 server 109 retrieves 406 an IPv4 address corresponding to the IPv4 payload. The IPv4 DNS server 109 then creates and sends 408 a DNS response, which comprises the retrieved IPv4 payload (containing the IPv4 address of host H2 retrieved at step

406), stored within a packet containing an IPv4 header. Upon arrival at the Border Router BR, the header of the IPv4 DNS response packet is converted into to an IPv6 header for onward routing within the first network NW1. Again, during the translation process, the payload of the packet is unchanged. Once the IPv4 DNS
5 response arrives at the dual stack host H1, the IPv4 application that initiated the DNS request will interpret the IPv4 DNS payload.

The dual stack host H1 then sends 410 a request to the DHCPv6 server 103 for an IPv4 address to be allocated to it, whereupon the DHCPv6 server 103 returns 412 an allocated IPv4 address together with a preferred and a valid timer. Typically
10 the dual stack host H1 and DHCPv6 server 103 send request and response messages (steps 410 and 412 respectively) in accordance with the User Datagram Protocol (UDP), although an alternative protocol, such as Transmission Control Protocol (TCP), could be used.

Next the updating program 315 running on the dual stack host H1 sends
15 414 to the Border Router BR details of the IPv4 address that was returned from the DHCPv6 server 103 at step 412. This is the first operation performed by the coordinator 300, and essentially involves the updating program 315 sending a packet, preferably using the UDP protocol, to the Border Router BR. For example, the updating program 315 may send a copy of the response packet received at step 412
20 to the Border router BR.

The existing documentation relating to the DSTM procedure (draft-ietf-ngtrans-dstm-04.txt, referred to above) is silent as to whether timing of IPv4 address allocation information is an issue, and certainly does not suggest proactive communication of IPv4 address allocation information. Thus, if one were to operate a
25 DSTM communications as per the standard documentation, the Border Router BR learns of the IPv4 allocated address when the dual stack host H1 sends a packet destined for an IPv4 host H2 in the IPv4 network NW2.

In embodiments of the present invention, because the dual stack host H1 sends (as described above at step 414) the Border Router BR a packet detailing the
30 allocated IPv4 address in advance of sending a packet destined for the IPv4 host H2, the Border Router BR can set up the tunnel 105 in preparation for packets passing between the two hosts H1, H2.

On receipt of the address allocation message sent at step 414, the Border Router BR retrieves 416 address information therein and creates a mapping between the allocated IPv4 address and the IPv6 address of the dual stack host H1. This mapping is subsequently used to create the second endpoint EP2 of the tunnel 105.

5 After the dual stack host H1 has sent the address information to the Border Router (step 414), the host H1 creates and sends 418 an encapsulated IPv4 data packet destined for host H2 in the IPv4 network. This data packet contains data from an IPv4 application running on the dual stack host H1, which has been encapsulated into an IPv6 packet for transmission over the IPv6 network. Such
10 packet encapsulation is well known to those skilled in the art, and the reader is referred to RFC2893, available from the IETF, for more details.

Accordingly the dual stack host H1 sets, as a source address of the outgoing encapsulated packet, the IPv6 address of the dual stack host H1, and as destination address of the encapsulated packet, the IPv6 address of the Border
15 Router BR. The payload of this packet comprises an IPv4 packet having the IPv4 address allocated at step 412 as source address, and the IPv4 address returned by the DNSv4 server 109 at step 408 as destination address (i.e. IPv4 address of H2). Upon receipt of the encapsulated packet, the Border Router BR un-encapsulates
20 the received packet, and sends 422 the now IPv4 packet into the IPv4 network NW2.

Assuming the IPv4 host H2 sends 424 a reply packet having, as source address, the IPv4 address of the IPv4 host H2 and as destination address the IPv4 allocated address, such a reply packet will be received at the Border Router BR. In response to receipt of the reply packet, the Border Router BR retrieves 426 an IPv6
25 address corresponding to the IPv4 destination address of the received packet, by accessing tunnel information stored at step 416.

The Border Router BR subsequently encapsulates 428 the received packet using the retrieved IPv6 address as destination address and the IPv6 address of the Border Router BR as source address. The packet is then sent 430 into the IPv6
30 network whereupon it is received by the dual stack host H1.

Data can be continually transmitted between the dual stack host H1 and the IPv4 host H2, according to steps 418 – 430, provided the tunnel 105, and thus the IPv4 address allocated to the dual stack host H1, is active. As stated above, when

the DHCPv6 server 103 allocates an IPv4 address to a requesting dual stack host, it sends a DHCPv6 response message, which includes the allocated IPv4 address and preferred and valid timers. The preferred timer starts decrementing as soon as the address is allocated to the dual stack host H1 (step 412), so for communications to
5 continue past expiry of the preferred timer, the corresponding address allocation and timers require renewing. Failure to renew the address allocation and timers results in expiry of the IPv4 address allocation and removal of the tunnel 105 (this is described later).

This renewal process is handled by the coordinator 300 as is shown in
10 Figure 5. Typically the allocated IPv4 address is applied to a particular interface on the dual stack host H1 (referred to herein as the encapsulating interface), so that all outgoing and incoming IPv4 encapsulated packets pass through this interface. The monitoring program 311 is arranged to monitor 502 both the count down of the preferred timer and packets arriving at the encapsulating interface, so that, if the
15 preferred timer has expired, and, for example an outgoing IPv4 packet arrives at the encapsulating interface, this triggers the renewal program 313 to request 504 renewal of the previously allocated IPv4 address from the DHCPv6 server 103.

Alternatively the allocated address can be stored in an address allocation database (not shown), with a flag indicating the status thereof. At step 502, the
20 monitoring program 311 could decrement both the preferred timer and valid timer, whilst continually checking whether or not they have expired. Once the preferred timer is identified to have timed out, the monitoring program 311 could check the status of the allocated address in the address allocation database, and, should the status of the flag indicate that communications are to continue, the renewal program
25 313, at step 504, could request renewal of the said previously allocated IPv4 address. The flag can be modified by applications running on the dual stack host H1. If the flag does not indicate that communication are to continue, the flag may be monitored by 311 continuously until the valid timer is identified to have timed out. If
30 at any point prior to valid timer expiry, the flag indicates that communications are to continue, the renewal program 313, at step 504, could request renewal of the said previously allocated IPv4 address.

The flag could be modified in response to receipt of an IPv4 packet arriving at the encapsulating interface or in response to creation of an IPv4 TCP or UDP socket on the host H1.

In response to receipt of a renewal request, the DHCPv6 server 103 sends
5 506 the renewed allocated IPv4 address, together with reset preferred and valid timers, to the dual stack host H1. These are received by the monitoring program 311, which sets 508 the preferred and valid timers held on the dual stack host H1 to the timer values sent at step 506. This process (steps 502 – 508) continues indefinitely, provided at least one application on the dual stack host H1 requires a
10 communication path to the IPv4 network (or the status of the flag in the address allocation database indicates that communications should continue).

In this embodiment, the dual stack host H1 only needs to communicate with the Border Router BR if the allocated IPv4 address is not renewed – i.e. if the tunnel
105 is to be broken. Thus for this embodiment, the renewal process does not involve
15 communication with the Border Router BR.

If none of the application programs require renewal of the allocated IPv4 address (or if the flag in the address allocation database is set to null for whatever reason), then when the preferred timer expires, the monitoring program 311 triggers
602 release of the allocated IPv4 address. This release can be effected in one of two
20 ways: referring to Figure 6, either the renewal program 313 sends 604 a release message to the DHCPv6 server 103, or the dual stack host H1 waits for the valid timer to expire (recall that on expiry of the *valid* timer, allocation of the allocated IPv4 address is automatically annulled).

In either case, the updating program 315 sends 606 a release message to
25 the Border Router BR, whereupon the Border Router BR removes 608 the mapping stored at step 416. This has the effect of removing the tunnel 105 between the dual stack host H1 and the Border Router BR.

As stated above, the coordinator 300 could alternatively be stored on, and run from, either an independent controller device 107 or the Border Router BR.
30 Figures 7 – 10 show allocation, renewal and release processes when controlled by the controller device 107. This embodiment is generally similar to the first embodiment described with reference to Figures 1 – 6, so that, in Figures 7 – 10,

like parts have been given like reference numerals and will not be described further in detail.

Referring first to Figure 7, steps 402 – 412 are identical to those described above. Thereafter dual stack host H1 sends 414 the controller 107 details of the allocated IPv4 address and the preferred and valid timers, whereupon the controller 107 sends 415 the allocated IPv4 address to the Border Router BR to enable the Border Router BR to create the tunnel 105, at step 416. Subsequent steps progress as per the first embodiment.

Renewal of the preferred timer is shown in Figure 8: the monitoring program 311 (running on the controller 107) monitors 502 the count down of the preferred timer. When the preferred timer has expired, the controller 107 sends 503a a message to the dual stack host H1 asking whether the IPv4 address should be renewed. If it receives 503b a positive response from the dual stack host H1, the renewal program 313 requests renewal 504 of the previously allocated IPv4 address from the DHCPv6 server 103. As for the first embodiment, the dual stack host H1 could have an address allocation database where the allocated address and a flag, indicating the status of the allocation, are stored. Thus step 503a could involve querying the flag status stored in the said address allocation database.

In response, the DHCPv6 server 103 sends 506 the renewed allocated IPv4 address, together with new preferred and valid timers, to the controller 107. These are received by the monitoring program 311, which sets 508 the preferred and valid timers held on the controller 107 to the timer values sent at step 506. This process (steps 502 – 508) continues indefinitely, provided at least one application on the dual stack host H1 requires a communication path to the IPv4 network.

Referring to Figure 9, if the response sent at step 503b is negative, meaning, for example, that none of the applications running on the dual stack host H1 require communications with the IPv4 network, the monitoring program 311 allows the valid timer to elapse 605. Once the valid timer has elapsed, the updating program 315 sends 606 a release message to the Border Router BR, whereupon the Border Router BR removes 608 the mapping stored at step 416. In addition, the updating program 315 sends 610 a message to the dual stack host H1, informing the dual stack host H1 that the allocated IPv4 address is no longer valid, whereupon the host H1 removes the allocated IPv4 address from its internally maintained

address tables. This has the effect of removing the tunnel 105 between the dual stack host H1 and the Border Router BR.

Alternatively, as is shown in Figure 10, the dual stack host H1 can proactively inform 1002, 1004 the DHCPv6 server that the allocated IPv4 address is
5 no longer required. In this case, the DHCPv6 server 103 informs 1006 the controller 107 of the said expiry, whereupon the updating program 315 sends 1008 a release message to the Border Router BR. As shown in Figures 8 and 9, on receipt of the release message, the Border Router BR removes 1010 the mapping stored at step 416.

10 If the coordinator 300 were to be located on the Border Router BR (not shown), step 414 would involve sending IPv4 address allocation and timer information directly to the Border Router BR, as shown in Figure 4, and management of address and timer information would be monitored and renewed by the coordinator 300 as shown in Figures 8, 9 and 10, but with the relevant processes
15 running on the border router BR.

As a further alternative, the coordinator 300 could run on the DHCPv6 server 103, so that the DHCPv6 server 103 essentially acts as controller 107. In this arrangement the steps performed by the renewal program 313, such as step 504 shown in Figure 8, would be internally run processes.

20 As a yet further alternative, the coordinator 300 could be distributed between the host H1 and the DHCPv6 server 103, as shown in Figure 13; the process steps associated with allocation and renewal of an allocated IPv4 address then progress as shown in Figures 14 and 15 (which correspond, respectively, to Figures 4 and 6). Referring firstly to Figure 13, the monitoring and renewing
25 programs 311, 313 are run on the host H1, while the updating program 315 runs on the DHCPv6 server 103. Thus the only device that interacts with the Border Router BR is the DHCPv6 server 103.

Turning now to Figure 14, steps 402 – 410 progress as per the first embodiment. Then, at step 1401, the updating program 315 running on the DHCPv6
30 server 103 sends the Border Router BR a packet comprising the IPv4 address that the DHCPv6 server 103 sent to the host H1 at step 412, together with the IPv6 address of the host H1. The Border Router BR then retrieves 416 address information in the packet and creates a mapping between the allocated IPv4 address

and the IPv6 address of the dual stack host H1. Steps 418 – 430 progress as per the first embodiment.

As stated above, the monitoring and renewing programs 311, 313 run on the host H1 in this embodiment; once the host H1 has received, at step 412, its allocated IPv4 address and timer information, the monitoring program 311 monitors
5 countdown of the preferred timer, as per the first embodiment, so that renewal of the timers proceeds as described in Figure 5 (steps 502 – 508). However, when the host is ready to release the allocated IPv4 address, the DHCPv6 server 103 communicates a release message to the Border Router. This is shown in Figure 15,
10 where it can be seen that, in response to receipt of a release message from the host H1 (step 604), the updating program 315 running on the DHCPv6 server 103 informs 1501 the Border Router BR of the address expiry.

Figure 16 shows an alternative scenario, where the allocated IPv4 address is released by default due to expiry of the valid timer. Since the timers are sent out
15 from the DHCPv6 server 103 when allocating an IPv4 address (step 412), the status of the preferred and valid timers is also observed by a process running on the DHCPv6 server 103. Accordingly, in the event that the DHCPv6 server 103 does not receive a request for renewal of the allocated address (as per Figure 5), it waits for the valid time to expire and then informs 1601 the Border Router BR that the
20 allocated address is no longer valid for host H1. The Border Router BR then deletes the mapping corresponding thereto, as described at step 608 in the context of the first embodiment.

Figure 17 shows an arrangement whereby the Border Router BR and DHCPv6 server 103 are combined within a single device. Since the complexity of a
25 network manager's job scales with number of devices that require managing, this arrangement has the advantage of reducing the overhead associated with DSTM network management. In particular, as the DHCPv6 server 103 is essentially a pool of IPv4 addresses, combining the functionality of the Border Router BR and DHCPv6 server 103 in one device reduces the overhead associated with management of the
30 address pool. In addition, this arrangement is convenient for processing of packets by the Border Router BR because it needs to store, or have access to, a list of IPv4 addresses that have been allocated by the DHCPv6 server 103 in order to detect (and process) packets using these allocated addresses.

Additional advantages of embodiments of the invention can be seen with reference to Figures 11 and 12, which show operation of the dual stack host H1, DHCPv6 server 103, and border router BR in the absence of the coordinator 300. As for Figures 7 – 10, parts and steps that are common between the first embodiment
5 and Figures 11 and 12 (communications in the absence of coordination) are not described further in detail.

Figure 11 shows the situation post expiry of the valid timer, where the IPv4 address allocated to the dual stack host H1 has been released 1102. If the IPv4 host H2 subsequently sends, at step 424, a packet to the dual stack host H1, using the
10 previously valid allocated IPv4 address as destination address, the Border Router BR will proceed to encapsulate 426 the received packet (as described above with reference to Figure 4), and send 430 the encapsulated packet to the dual stack host H1. Once the packet arrives at the dual stack host H1, then because the dual stack host H1 has released this allocated IPv4 address, none of the interfaces on the dual
15 stack host H1 recognise the IPv4 address of the packet, and the packet will thus be dropped.

Conversely, with embodiments of the invention, the Border Router BR will be informed of the release of allocated IPv4 addresses, whereupon the tunnel 105 will be removed. Any packets subsequently arriving at the Border Router BR with a
20 destination address of the released IPv4 address will thereafter be dropped at the router BR.

In a second scenario, also shown in Figure 11, the released IPv4 address is subsequently allocated 1104 to another dual stack host Hd by the DHCPv6 server 103, whereupon the other dual stack host Hd creates an encapsulated IPv4 packet
25 (step 418) using allocated IPv4 address, and sends the encapsulated packet into the IPv6 network NW1 (recall that without functionality of embodiments of the invention the Border Router BR is a) unaware of the release of the allocated IPv4 address by the dual stack host H1 and is b) unaware of the subsequent allocation to the other dual stack host Hd).

30 When the packet arrives at the Border Router BR, there is confusion because the IPv6 address (source address of the encapsulated packet, here the other dual stack host Hd) does not match the source address expected by the Border Router (the IPv6 address of dual stack host H1 that was stored previously). With standard

DSTM methods, the Border Router BR is likely to replace any existing endpoint EP2 mapping with a new mapping between this IPv6 address and the allocated IPv4 address. This approach can have serious security implications, not least because the Border Router BR has no means of assessing "correct" IPv4 address allocation.

5 Essentially, the Border Router BR has no way of knowing whether the dual stack host H1 has in fact released the allocated IPv4 address. It could be, e.g. that the other dual stack host Hd is spoofing the allocated IPv4 address in an attempt to intercept packets destined for another dual stack host in the network NW1. Clearly in this situation the Border Router BR should not alter the endpoint EP2 mapping,
10 but, without embodiments of the invention, the Border Router BR is likely to alter the endpoint EP2 mapping, thereby allowing dual stack host Hd access to data that is destined for another dual stack host.

A third scenario, shown in Figure 12, is an alternative to the situation described above, where the Border Router BR modifies the endpoint EP2 mapping
15 upon receipt of a packet from the other dual stack host Hd. Figure 12 is thus representative of a case where, upon receipt of a packet from the other dual stack host Hd, the Border Router BR does not change the mapping between dual stack host and allocated IPv4 address. However, this action, on the part of the Border Router BR, also incurs problems – clearly the Border Router BR should modify the
20 endpoint mapping when other dual stack host Hd has legitimately been allocated that IPv4 address.

As a result of the Border Router BR failing to modify the endpoint EP2 mapping, when packets arrive at the Border Router BR from the IPv4 host in the second network NW2, the Border Router BR will encapsulate the received packets
25 using an out of date IPv6 destination address, so that the packets will arrive at the wrong dual stack host H1.

The scenarios discussed above and shown in Figures 11 and 12 thus show that without some kind of address allocation coordination mechanism, packets can get dropped, snooped, and misdirected, which generates additional network traffic
30 and increases the possibility of security problems. Embodiments of the present provide such a coordination mechanism, and thereby provide an improved solution to management of packet routing for the DSTM migration method.

In fact, embodiments of the invention could be applied to any communications method in which:

- addresses are allocated dynamically;
- the allocated address can be allocated to one of a plurality of machines; and
- 5 • the allocated address is used to route data within a network.

As will be understood by those skilled in the art, the invention described above may be embodied in one or more computer programs. These programs can be contained on various transmission and/or storage mediums such as a floppy disc, CD-ROM, or other optically readable medium, or magnetic tape so that the programs
10 can be loaded onto one or more general purpose computers or could be downloaded over a computer network using a suitable transmission medium.

The programs 311, 313, 315 of the present invention are conveniently written using the C programming language, but it is to be understood that this is inessential to the invention.

CLAIMS

1. A method of coordinating address allocation events at a routing device and at a host machine, the routing device being arranged to route data between the host
5 machine and at least one other host machine, in which the host machine is located in a first network, the first network operating in accordance with a first transmission protocol, and in which the other host machine is located in a second network, the second network operating in accordance with a second transmission protocol, and wherein the host machine is operable to process packets of data corresponding to
10 either transmission protocols,
the method comprising the steps of:
- (i) receiving an allocated address, for use by the host machine in communicating with the other host machine;
 - (ii) receiving an indicator representative of allocation status of the allocated
15 address;
 - (iii) sending, to the routing device, the allocated address together with the address of the host machine;
 - (iv) storing a mapping between the allocated address and the address of the host machine;
 - 20 (v) monitoring the indicator, and, if the indicator indicates that the allocated address is no longer valid for the host machine,
 - (vi) rendering the stored mapping unusable.
2. A method according to claim 1, including using the mapping to establish a
25 tunnel between the routing device and the host machine, the tunnel being used in the said routing of data between the host machine and the other host machine.
3. A method according to any one of the preceding claims, wherein the indicator includes a timer and the method includes requesting renewal of the allocated address
30 in the event that the indicator indicates that allocation of the allocated address to the host machine has expired.

4. A method according to claim 3, including checking at least one interface of the host in order to determine whether the allocated address is to be renewed.
5. A method according to any one of claims 1 to 3, wherein the indicator includes a communication status and the method includes checking the communication status in order to determine whether the allocated address is to be renewed.
6. A method according to any one of the preceding claims, in which the rendering step causes the stored mapping to be deleted.
7. Apparatus for coordinating address allocation events at a routing device and at a host machine, the routing device being arranged to route data between the host machine and at least one other host machine, in which the host machine is located in a first network, the first network operating in accordance with a first transmission protocol, and in which the other host machine is located in a second network, the second network operating in accordance with a second transmission protocol, and wherein the host machine is operable to process packets of data corresponding to either transmission protocols,
- the apparatus comprising
- (a) receiving means arranged to receive an allocated address, for use by the host machine in communicating with the other host machine, and to receive an indicator representative of allocation status of the allocated address;
 - (b) means arranged to send the allocated address together with the address of the host machine to the routing device;
 - (c) monitoring means arranged to monitor the indicator, and operable to alert the routing device of changes to the indicator,
- wherein the routing means is arranged, upon receipt of the allocated address, to store a mapping between the allocated address and the address of the host machine, and upon receipt of the alert, to render the stored mapping unusable.
8. Apparatus for coordinating address allocation events at a routing device and at a host machine, the routing device being arranged to route data between the host

- machine and at least one other host machine, in which the host machine is located in a first network, the first network operating in accordance with a first transmission protocol, and in which the other host machine is located in a second network, the second network operating in accordance with a second transmission protocol, and
- 5 wherein the host machine is operable to process packets of data corresponding to either transmission protocols,
- the apparatus comprising
- (a) allocating means arranged to allocate an address to the host machine, for use by the host machine in communicating with the other host machine,
- 10 (b) means arranged to send the allocated address together with the address of the host machine to the routing device;
- (c) updating means arranged to receive an indicator representative of allocation status of the allocated address, and, if the indicator indicates that the allocated address is to be renewed, the updating means is arranged to send data identifying
- 15 the said renewal to the host machine, and if the indicator indicates that the allocated address is not to be renewed, the updating means is arranged to send an alert to the routing device;
- wherein the host is arranged to monitor the indicator, and is operable to send data identifying changes to the indicator to the apparatus;
- 20 and wherein routing means is arranged, upon receipt of the allocated address, to store a mapping between the allocated address and the address of the host machine, and upon receipt of the alert, to render the stored mapping unusable.
9. Apparatus according to claims 7 or 8, wherein the routing device is
- 25 arranged to remove the stored mapping upon receipt of the alert.
10. Apparatus according to any one of claims 7 to 9, in which the indicator involves a timer, and the allocated address is deemed to be invalid for the host machine when the timer has expired.
- 30
11. Apparatus according to any one of claims 7 to 10, in which the host machine is operable to process at least one application in accordance with the second transmission protocol.

12. A device for managing networks including apparatus according to any one of claims 7 to 11.

- 5 13. A computer program, or a suite of computer programs, which, when loaded on a computer, or a suite of computers, is operable to carry out the method according to any one of claims 1 to 6.





Fig 1

**Fig 2**




Fig 3

4/15




Fig 4




Fig 5




Fig 6




Fig 7




Fig 8




Fig 9




Fig 10

10/15




Fig 11





Fig 12

12/15

**Fig 13**

13/15






Fig 14

14/15

**Fig 15****Fig 16**




Fig 17