

Method of Efficiently Routing Multicast Traffic

Technical field

5 The invention relates to a method of routing Internet Protocol version 6, IPv6, multicast traffic in a utility metering system.

Background

10 IPv6 (Internet Protocol Version 6) is the successor to IPv4 (Internet Protocol Version 4), and provides devices that communicate over the Internet with a unique, numerical IP address. IPv4, which was the original IP address standard, provides devices with unique 32-bit IP addresses, and therefore can provide support for 2³² IP addresses (in total, around 4.29 billion). However, due to the widespread usage of the Internet, the
15 addresses available to be allocated under the 32-bit address system are rapidly reducing. IPv6 differs from IPv4 in that it uses unique 128-bit IP addresses, and therefore can provide 2¹²⁸ IP addresses. This should allow allocation of 128-bit IP addresses to Internet devices for a long time to come.

20 The 128-bit IPv6 addresses are represented using hexadecimal notation. The 128-bits are divided into eight 16-bit segments, with each 16-bit segment being converted into a 4-digit hexadecimal number separated by a colon. For example:

2001:0db8:0001:0000:0000:0000:0000:0001

25 IPv6 can facilitate one-to-many communications using multicast. Multicasting allows a single device to send identical communications to all devices within a multicast group using a IPv6 multicast address. Multicasting is therefore a useful communication tool that may be used in a range of systems, such as utility metering systems, to allow
30 efficient communication to hundreds, or even thousands, of devices within the system (e.g. to metering devices within a metering system).

 However, routing multicast traffic can be expensive, and increase system memory and processing demands, because each multicast address needs to be tracked by the
35 operating system, and large systems may comprise thousands or more multicast

addresses. For example, using the multicast address format defined in RFC-3306 (Unicast-Prefix-based IPv6 Multicast Addresses) provides the ability to use up to 4,294,967,295 possible multicast addresses.

- 5 There exists a need to provide a way of more efficiently routing multicast traffic.

Summary

10 According to the invention in a first aspect, there is provided a method of routing Internet Protocol version 6, IPV6, multicast traffic in a utility metering system, the method comprising: creating tunnels between each of a plurality of devices of the utility metering system on a first network and at least one device of the utility metering system on a second network to allow two way communication therebetween, wherein each of the plurality of devices on the first network facilitate communication to at least
15 one multicast group; generating and assigning a unique network prefix to each of the plurality of devices on the first network; prepending to each unique network prefix, a multicast prefix, to generate a unique device multicast prefix for each of the plurality of devices on the first network comprising only the unique network prefix and the multicast prefix and not a group ID identifying the at least one multicast group; populating a
20 multicast routing table using the unique device multicast prefixes of each of the plurality of devices on the first network to allow communication routes between each of the plurality of devices on the first network and the at least one device on the second network to be established; and using the multicast routing table to route multicast traffic from the at least one device on the second network a relevant one of the plurality of
25 devices on the first network.

Optionally, the unique device multicast prefixes within the multicast routing table are /96 prefixes.

30 Optionally, each unique network prefix comprises a /64 prefix.

Optionally, each multicast prefix comprises a /32 prefix.

Optionally, each unique network prefix comprises a length of 64-bits or less.

35

Optionally, the second network is a subnet of the first network.

Optionally, the multicast routing table comprises only a single entry for each of the plurality of devices on the first network regardless of how many multicast groups the
5 respective device on the first network facilitates communication with.

Optionally, each of the plurality of devices on the first network comprise a border router configured for communication with a plurality of utility metering devices, and the at least one device on the second network comprises a tunnel management server configured
10 to facilitate creation of the tunnels.

Optionally, the utility metering system further comprises a head end system configured to communicate with the plurality of border routers via the tunnel management server using the unique device multicast prefixes of each of the plurality of border routers.
15

Optionally, the method further comprises assigning the tunnel management server a network prefix, and wherein the unique network prefix of each of the border routers is a /64 prefix, the /64 prefix being a subnet of the network prefix of the tunnel management server.
20

Brief description of drawings

Figure 1 shows an exemplary utility metering system;
Figure 2 shows a schematic view of the format of an IPv6 multicast address;
25 Figure 3 shows a schematic view of the format of an IPv6 multicast address;
Figure 4 shows an exemplary utility metering system;
Figure 5 shows a flow chart of an exemplary method of routing multicast traffic in an IPv6 system; and
Figure 6 shows an exemplary utility metering system.

30

Detailed description

Generally disclosed herein is a method for efficiently routing multicast traffic in utility metering systems. Multicasting comprises a method wherein data is sent from one
35 source device to multiple devices within a specified multicast group. A multicast group

may be defined as comprising a plurality of devices identified by a single multicast address. As such, when data is sent from the source device to the single multicast address, it is delivered to all devices that are members of the multicast group.

5 As shown schematically in Figure 1, a utility metering system 100 may typically comprise a head end system 102, which functions as a central processing system, and which communicates with a plurality of metering devices 108a-n, 110a-n and 112a-n via border routers 106a-c (also known as a collectors or root nodes). The plurality of utility metering devices may comprise electric, gas or water utility meters configured to
10 measure consumption of a respective utility. The plurality of utility metering devices are connected via a wireless mesh network. The border routers 106a-c may facilitate the communication of data from the wireless mesh network to the head end system 102 via one or more networks (or by cellular communications or power line communications).

15 As shown in Figure 1, typical utility metering systems 100 comprise a plurality of border routers 106a-c, each in communication with the head end system 102, and each in communication with a plurality of metering devices. In the arrangement shown in Figure 1, the border router 106a communicates with the metering devices 108a-n, the
20 border router 106b communicates with the metering devices 110a-n and the border router 106c communicates with the metering devices 112a-n. The border routers 106a-c may comprise a personal area network (PAN) coordinator, a gateway, or any other device capable of routing communications from the head end system 102 to the plurality of metering devices 108a-n, 110a-n and 112a-n and vice versa.

25 Within the context of a utility metering system, multicasting may be used to allow efficient communications between, for example, the head end system 102 and the relevant plurality of metering devices 108a-n, 110a-n and/or 112a-n.

30 For example, in the utility metering system 100, each utility metering device connected to a given border router, 106a, 106b or 106c, may be a member of one or more multicast groups. For example, a first multicast group may be defined comprising a first set of utility metering devices (e.g. electrical metering devices) connected to the border router 106a, and a second multicast group may be defined comprising a second
35 set of utility metering device (e.g. gas metering devices) connected to the border router

106a etc. As such, communications from the head end system can be directed to relevant utility metering devices using the relevant multicast address defining the relevant multicast group.

5 However, as the number of multicast groups increases, so too does the number of multicast addresses that need to be tracked. This increases the expense and complexity of the system.

10 The invention disclosed herein reduces the number of multicast addresses that need to be tracked, while still allowing multicast communications, and therefore providing the advantages associated therewith.

15 To aid understanding of the invention, the format of an Internet Protocol version 6 (IPv6) multicast address, specifically a unicast-prefix-based IPv6 multicast address, is firstly described with reference to Figure 2. The format of a unicast-prefix-based IPv6 multicast address is defined in RFC-3306. Unicast-prefix-based IPv6 multicast addresses are a specific type of multicast address wherein a unicast prefix (typically a 64-bit network prefix) is embedded inside the multicast address itself. This allows multicast groups to be created that are tied to a particular set of network addresses/prefixes assigned to devices within the system.

20 The first 8 bits of IPv6 multicast addresses are always “11111111”, represented as “FF” in hexadecimal, and shown in Figure 2 as field 202.

25 The next 4 bits are flags, shown schematically in Figure 2 as field 204. The following 4 bits after the flags are the scope, shown schematically in Figure 2 as field 206. The scope field identifies the application scope of the multicast group, as below.

Value	Meaning
1	Interface-local scope
2	Link-local scope
4	Admin-local scope
5	Site-local scope
8	Organisation-local scope
E	Global scope

The next 8 bits are reserved, and have a default value of “0”, as designated by the “reserved” field 208 in Figure 2.

5 The next 8 bits represent the plen or prefix length field, shown as field 210 in Figure 2. This field 210 indicates the effective length of the embedded unicast prefix, or network prefix.

10 The next 64 bits represent the embedded unicast prefix, or network prefix, shown as field 212 in Figure 2.

The final 32 bits represent the Group ID, and are used to distinguish different multicast groups that share the same network prefix. This is designated as field 214 in Figure 2.

15 As such, a 128-bit IPv6 multicast address may be defined. Figure 3 shows an example 128-bit IPv6 multicast address, mapping the hexadecimal digits to the fields 202-214 described above to aid understanding.

20 A method of routing IPv6 multicast traffic in a utility metering system is now described with reference to Figures 1 to 5.

25 502: As shown in Figure 1, each border router 106a-c is in communication with a tunnel manager, or tunnel management server, 104. In the exemplary arrangement shown in Figure 1, each border router 106a-c is in communication with the same tunnel manager 104, however the skilled person will appreciate that in alternative arrangements, one or more of the border routers 106a-c may be in communication with different tunnel managers, as will be described in more detail below.

30 For the purposes of this example, the method is described with respect to a single border router 106a, as shown in Figure 4. However the skilled person will understand that similar methods are undertaken to facilitate routing of multicast traffic to different border routers, e.g. 106b and 106c shown in Figure 1, as applicable.

35

When the border router 106a, boots up, it initiates a tunnel creation sequence to establish a tunnel between the tunnel manager 104 and the border router 106a to allow two way communication therebetween.

5 The tunnel created between the border router 106a and the tunnel manager 104 may be configured to support multicast traffic, for example, by enabling the relevant multicast routing protocols. The methods for configuring a tunnel to support multicast traffic will be familiar to the skilled person.

10 504: A unique network prefix is generated and assigned to the border router 106a.
The unique network prefix generated and assigned to the border router 106a
may comprise a 64-bit prefix, denoted as /64. The skilled person will appreciate
that in alternative arrangements, a different number of bits may be used for the
network prefix, for example a network prefix of 64-bits or less may be used.
15 Typically however, the network prefix assigned to the border router 106a will not
be more than 64-bits, in accordance with RFC-3306.

The unique /64 prefix may be generated as a subnet of a network prefix assigned to the tunnel manager 104. For example, as shown in Figure 4, the tunnel manager 104 is assigned the /48 prefix 2001:DB8:1::/48. The /48 prefix means that the first 48 bits of the 128-bit IPv6 address are designated for the network prefix, in this case the network prefix of the tunnel manager 104. The remaining bits of the 128 bit IPv6 address can be used for subnets and host addresses. Typically subnets are assigned a /64 size, and as such a very large number of /64 subnets can be created within its assigned /48 prefix. The skilled person will appreciate that in alternative arrangements, the tunnel manager 104 may comprise a network prefix that is a different number of bits in length. In other words, the tunnel manager network prefix need not be 48 bits in length, and in other examples may be less than 48 bits or more than 48 bits, and this may be configurable by the user in dependence on the system requirements.

In the example shown in Figure 4, the border router 106a is assigned a /64 network prefix of:

35 **2001:DB8:1:1::/64.**

As will be evident to the skilled person, the part of the /64 network prefix of the border router 106a denoted in bold and underlined corresponds to the /48 prefix of the tunnel manager 104, identifying the border router 106a as a subnet of the tunnel manager 104, while the remaining 16-bits of the 64-bit network prefix of the border router 106a uniquely identify the border router 106a.

The skilled person will also understand that in the example shown in Figure 4, the /48 network prefix of the tunnel manager 104 is a subnet of the network prefix of the headend system 102.

506: A border router multicast prefix is created using the /64 network prefix assigned to the border router 106a. The border router multicast prefix is distinct from a multicast address as outlined in RFC 3306 (i.e. a unicast-prefix-based IPv6 multicast address).

Creating the border router multicast prefix comprises prepending a multicast prefix to the /64 network prefix of the border router 106a to form a /96 border router multicast prefix. As used herein, the term "multicast prefix" is used to refer to the first 32-bits of a multicast address that immediately precede the /64 network prefix. Referring to Figure 2, the multicast prefix comprises the fields 202-210 comprising the FF designation 202, the flags 204, the scope 206, the reserved field 208 and the plen field 210. For example, a multicast prefix may be FF38:40::, which when prepended to the /64 network prefix of the border router 106a creates a border router multicast prefix of:

FF38:40:2001:DB8:1:1::/96.

As used herein, the term "device multicast prefix" or "border router multicast prefix", refers to the /96 prefix assigned to a device comprising the network prefix and the multicast prefix prepended thereto.

As mentioned above, the border router multicast prefix is distinct from a multicast address, as defined within RFC-3306. Using the multicast address format outlined in RFC-3306 (and shown in Figures 2 and 3), a group ID

identifying the multicast groups served by the border router 106a would be appended to the multicast prefix and the network prefix of the border router 106a, as shown in field 214 in Figure 2, to form a 128-bit multicast address. For example, where three multicast groups are served by the border router 106a,
5 three multicast addresses may be created, for example:

FF38:40:2001:DB8:1:1::1;
FF38:40:2001:DB8:1:1::2;
FF38:40:2001:DB8:1:1::3.

10

In contrast, the border router multicast prefix omits the group ID field. Therefore a single border router multicast prefix for the border router 106a may be used to route traffic to from the headend system 102 and to the border router 106a, regardless of how many multicast groups the border router 106a facilitates
15 communication with.

508: A multicast routing table is then populated using the /96 border router multicast prefix. The border router 106a uses the multicast routing table to manage and direct received data (e.g. data received from the tunnel manager 104). The
20 multicast routing table comprises data entries that can be used by devices, such as the tunnel manager 104 and the border router 106a, to determine the optimal path for forwarding multicast data from a source to a multicast group and over the IPv6 network.

25 Typically, in known multicast systems, each data entry in the multicast routing table comprises a source address (indicating where the multicast data originated from) and the destination group address in the form of a multicast address as shown in Figures 2 and 3, comprising the multicast prefix, the network prefix and the group ID, i.e. the 128-bit multicast address. This means
30 that where there are three multicast groups served by the border router 106a, there are three entries within the multicast routing table. The skilled person will appreciate that exemplary data entries in multicast routing tables may comprise further information, such as interface lists etc.

A utility metering system may comprise hundreds or thousands of border routers each serving multiple multicast groups. This means that the multicast routing tables are required to manage a large number of multicast addresses – i.e. if there are one thousand border routers and each border router serves three multicast groups, the multicast routing table is required to manage three thousand multicast addresses. This increases memory and processing demands, and can result in scalability challenges as the amount of information to be maintained in the multicast routing tables can grow rapidly as more border routers are implemented and/or as more multicast groups are defined.

The inventors have realised that these problems can be ameliorated when routing data from the head end system 102 to a given border router (e.g. border router 106a) via the tunnel manager 104 by using the /96 border router multicast prefix within the multicast routing table instead of the 128-bit multicast addresses corresponding to each multicast group. In other words, the 32-bit group ID field is omitted, and the multicast routing table comprises a /96 device multicast prefix (in this example, a border router multicast prefix) for each border router. As explained above the border router multicast prefix comprises the multicast prefix and the network prefix but not the 32-bit group ID field.

Continuing with the above example, in which the border router 106a serves three multicast groups, with addresses FF38:40:2001:DB8:1:1::1; FF38:40:2001:DB8:1:1::2 and FF38:40:2001:DB8:1:1::3, instead of there being three entries within the multicast routing table for border router 106a, there is instead only a single /96 border router prefix, which in this case would be FF38:40:2001:DB8:1:1::/96. The /96 border router multicast prefix covers all group IDs served by the corresponding border router, since the first 96 bits of the multicast addresses for each group served by the border router are the same. As such, if there are one thousand border routers in the utility metering system, then regardless of how many multicast groups are served by each border router, the multicast routing table is only required to manage one entry per border router.

510: Multicast data is routed from the head end system 102, via the tunnel manager 104, to the relevant border router that facilitates communication with the utility

metering devices within a given multicast group. The data is routed using the multicast routing table comprising the /96 border router multicast prefix. The use of a /96 multicast prefix within the multicast routing tables represents a departure from the known standards in the art, which specify populating a multicast routing table with the 128-bit multicast group address.

For example, any data destined for a multicast group served by the border router 106a is directed from the head end system 102 to tunnel manger 104 and then to border router 106a by virtue of the /96 border router multicast prefix (since all multicast addresses for the multicast groups served by the border router 106a will begin with the same 96 bits as defined in the first /96 border router multicast prefix).

The traffic may then be forwarded on to the relevant multicast group, by the border router, using further communications standards.

As mentioned above, exemplary utility metering systems may comprise a plurality of tunnel managers, each in communication with one or more border routers. A method of routing multicast traffic in a system comprising a plurality of tunnel managers is described below, however the skilled person will appreciate that a similar method is followed to that outlined in Figure 5.

Figure 6 shows an exemplary utility metering system comprising two tunnel managers 104a, 104b in communication with a head end system 102, though the skilled person will appreciate that in alternative systems may comprise a different number of tunnel managers, sometimes many thousands.

Tunnel manager 104a has a /48 prefix of 2001:DB8:1::/48, and tunnel manager 104b has a /48 prefix of 2001:DB8:2::/48. These prefixes are examples only to aid understanding, and as mentioned above, in alternative arrangements may be a different length to 48-bits (/48). As part of the tunnel creation sequences, /64 network prefixes are created for the border routers 106a, 106b and 106c. Again, the skilled person will appreciate that each tunnel manager 104a, 104b may communicate with a large number of border routers, and that the arrangement shown in Figure 6 is an example only.

Respective tunnels are created between border router 106a and the tunnel manager 104a; border router 106b and the tunnel manager 104a; and the border router 106c and the tunnel manager 104b. As such, the /64 network prefixes assigned to border
5 routers 106a and 106b are subnets of the /48 prefix of the tunnel manager 104a, and the /64 network prefix assigned to border router 106c is a subnet of the /48 prefix of the tunnel manager 104b. The skilled person will also understand that in the example shown in Figure 6, the /48 network prefix of the tunnel manager 104a is a first subnet of the network prefix of the headend system 102, and the /48 network prefix of the tunnel
10 manager 104b is a second subnet of the network prefix of the headend system 102.

/96 border router multicast prefixes may then be created in respect of each border router of the utility metering system, similarly to as described above in respect of Figure 5: in this case a first /96 border router multicast prefix is created for border router 106a,
15 a second /96 border router multicast prefix is created in respect of border router 106b and a third /96 border router multicast prefix is created in respect of border router 106c. As described above, the /96 border router multicast prefixes are created by prepending the multicast prefixes to each of the /64 network prefixes of the border routers 106a-c.

20 For example:

First /96 border router multicast prefix: FF38:40:2001:DB8:1:1::/96

Second /96 border router multicast prefix: FF38:40:2001:DB8:1:2::/96

Third /96 border router multicast prefix: FF38:40:2001:DB8:2:2::/96

25 As such, any data destined for a multicast group served by the border router 106a is directed from the head end system 102 to tunnel manager 104a and then to the border router 106a by virtue of the first /96 border router multicast prefix (since all multicast addresses for the multicast groups served by the border router 106a will begin with the
30 same 96 bits as defined in the first /96 border router multicast prefix). Any data destined for a multicast group served by the border router 106b is directed from the head end system 102 to tunnel manager 104a and then to the border router 106b by virtue of the second /96 border router multicast prefix (since all multicast addresses for the multicast groups served by the border router 106b will begin with the same 96 bits
35 as defined in the second /96 border router multicast prefix). Any data destined for a

multicast group served by the border router 106c is directed from the head end system 102 to tunnel manager 104b and then to the border router 106c by virtue of the third /96 border router multicast prefix (since all multicast addresses for the multicast groups served by the border router 106c will begin with the same 96 bits as defined in the third /96 border router multicast prefix).

Multicast traffic from the head end system 102 may be directed to the relevant border router 106a-c for distribution to a multicast group, using a multicast routing table containing one entry per border router, each entry comprising the /96 border router multicast prefix.

The skilled person will appreciate that the above described methods allow IPv6 multicast traffic to be routed efficiently, while decreasing the demands placed on memory and processing in comparison to known methods, even when additional multicast groups are defined or when additional border routers are added to the system.

It will be appreciated by the person of skill in the art that various modifications may be made to the above described embodiments without departing from the scope of the invention. The word “exemplary” is used herein to mean “an example”. Any embodiment described herein as “exemplary” is not necessarily to be construed as preferred or advantageous over other embodiments.

A computer program may be configured to provide any of the above described methods. The computer program may be provided on a computer readable medium. The computer program may be a computer program product. The product may comprise a non-transitory computer usable storage medium. The computer program product may have computer-readable program code embodied in the medium configured to perform the method. The computer program product may be configured to cause at least one processor to perform some or all of the method.

Various methods and apparatus are described herein with reference to block diagrams or flowchart illustrations of computer-implemented methods, apparatus (systems and/or devices) and/or computer program products. It is understood that a block of the block diagrams and/or flowchart illustrations, and combinations of blocks in the block

diagrams and/or flowchart illustrations, can be implemented by computer program instructions that are performed by one or more computer circuits. These computer program instructions may be provided to a processor circuit of a general purpose computer circuit, special purpose computer circuit, and/or other programmable data processing circuit to produce a machine, such that the instructions, which execute via the processor of the computer and/or other programmable data processing apparatus, transform and control transistors, values stored in memory locations, and other hardware components within such circuitry to implement the functions/acts specified in the block diagrams and/or flowchart block or blocks, and thereby create means (functionality) and/or structure for implementing the functions/acts specified in the block diagrams and/or flowchart block(s).

Computer program instructions may also be stored in a computer-readable medium that can direct a computer or other programmable data processing apparatus to function in a particular manner, such that the instructions stored in the computer-readable medium produce an article of manufacture including instructions which implement the functions/acts specified in the block diagrams and/or flowchart block or blocks.

A tangible, non-transitory computer-readable medium may include an electronic, magnetic, optical, electromagnetic, or semiconductor data storage system, apparatus, or device. More specific examples of the computer-readable medium would include the following: a portable computer diskette, a random access memory (RAM) circuit, a read-only memory (ROM) circuit, an erasable programmable read-only memory (EPROM or Flash memory) circuit, a portable compact disc read-only memory (CD-ROM), and a portable digital video disc read-only memory (DVD/Blu-ray).

The computer program instructions may also be loaded onto a computer and/or other programmable data processing apparatus to cause a series of operational steps to be performed on the computer and/or other programmable apparatus to produce a computer-implemented process such that the instructions which execute on the computer or other programmable apparatus provide steps for implementing the functions/acts specified in the block diagrams and/or flowchart block or blocks.

Accordingly, the invention may be embodied in hardware and/or in software (including firmware, resident software, micro-code, etc.) that runs on a processor, which may collectively be referred to as "circuitry," "a module" or variants thereof.

5 It should also be noted that in some alternate implementations, the functions/acts noted in the blocks may occur out of the order noted in the flowcharts. For example, two blocks shown in succession may in fact be executed substantially concurrently or the blocks may sometimes be executed in the reverse order, depending upon the functionality/acts involved. Moreover, the functionality of a given block of the
10 flowcharts and/or block diagrams may be separated into multiple blocks and/or the functionality of two or more blocks of the flowcharts and/or block diagrams may be at least partially integrated. Finally, other blocks may be added/inserted between the blocks that are illustrated.

15

20

CLAIMS:

1. A method of routing Internet Protocol version 6, IPV6, multicast traffic in a utility metering system, the method comprising:
 - 5 creating tunnels between each of a plurality of devices of the utility metering system on a first network and at least one device of the utility metering system on a second network to allow two way communication therebetween, wherein each of the plurality of devices on the first network facilitate communication to at least one multicast group;
 - 10 generating and assigning a unique network prefix to each of the plurality of devices on the first network;
prepending to each unique network prefix, a multicast prefix, to generate a unique device multicast prefix for each of the plurality of devices on the first network comprising only the unique network prefix and the multicast prefix and not a group ID
 - 15 identifying the at least one multicast group;
populating a multicast routing table using the unique device multicast prefixes of each of the plurality of devices on the first network to allow communication routes between each of the plurality of devices on the first network and the at least one device on the second network to be established; and
 - 20 using the multicast routing table to route multicast traffic from the at least one device on the second network a relevant one of the plurality of devices on the first network.
2. A method according to claim 1, wherein the unique device multicast prefixes
25 within the multicast routing table are /96 prefixes.
3. A method according to claim 1 or claim 2, wherein each unique network prefix comprises a /64 prefix.
- 30 4. A method according to claim 3, wherein each multicast prefix comprises a /32 prefix.
5. A method according to claim 1 or 2, wherein each unique network prefix comprises a length of 64-bits or less.

6. A method according to any preceding claim, wherein the second network is a subnet of the first network.
7. A method according to any preceding claim, wherein the multicast routing table
5 comprises only a single entry for each of the plurality of devices on the first network regardless of how many multicast groups the respective device on the first network facilitates communication with.
8. A method according to any preceding claim, wherein each of the plurality of
10 devices on the first network comprise a border router configured for communication with a plurality of utility metering devices, and wherein the at least one device on the second network comprises a tunnel management server configured to facilitate creation of the tunnels.
9. A method according to claim 8, wherein the utility metering system further
15 comprises a head end system configured to communicate with the plurality of border routers via the tunnel management server using the unique device multicast prefixes of each of the plurality of border routers.
10. A method according to claim 8 or 9, comprising assigning the tunnel
20 management server a network prefix, and wherein the unique network prefix of each of the border routers is a /64 prefix, the /64 prefix being a subnet of the network prefix of the tunnel management server.