

Multi-Subnet MANETs

<[draft-thaler-autoconf-multisubnet-manets-00.txt](#)>

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at
<http://www.ietf.org/lid-abstracts.html>

The list of Internet-Draft Shadow Directories can be accessed at
<http://www.ietf.org/shadow.html>.

Copyright Notice

Copyright (C) The Internet Society (2006). All Rights Reserved.

Abstract

This document describes an approach to addressing nodes in Mobile Ad-hoc Networks (MANETs) which involves assigning a separate subnet to each MANET router. This approach avoids many of the problems that arise in other approaches, and is intended to allow existing protocols and applications to work unmodified.

[1. Introduction](#)

A Mobile Ad-hoc Network (MANET) is one in which the topology may be very dynamic, such as highly mobile nodes communicating over a wireless medium. In this document we assume that there are forwarding nodes which participate in a MANET routing protocol. Furthermore, we assume that there may also be non-MANET nodes attached behind MANET forwarding nodes (e.g., via wired Ethernet). We will use the term "MANET router" to refer to the first type of nodes, and "client" to refer to the latter type of nodes.

There are three general approaches to architecting the data (forwarding) plane:

- (A) Treat the entire MANET as a single link, and do all forwarding at the datalink-layer, exposing (say) normal Ethernet to IP. The control plane thus consists of a layer-2 routing protocol. In this approach, existing protocols and applications work normally.
- (B) Treat the MANET as a single subnet composed of multiple links, such that MANET routers must do layer-3 forwarding (decrement TTL, etc.) between nodes in the same subnet. The control plane consists of a layer-3 MANET routing protocol. This approach has a number of problems as discussed in [\[MLSI\]](#), and may break existing protocols and applications. Clients are either not supported at all, or some additional proxying method is required.
- (C) Treat the MANET as a site containing many subnets, such that MANET routers act as normal routers when forwarding data between subnets. The control plane consists of a layer-3 MANET routing protocol, as with approach (B). This approach avoids most of the problems of (B), easily supports clients, and is what the remainder of this document outlines.

[2. Overview](#)

The MANET is treated as a site composed of many subnets. Each MANET router is assigned its own subnet prefix(es). Each client is assigned its own IP address(es) on a subnet assigned to a MANET router to which it is attached.

The MANET site may or may not have external connectivity. If there is, then any site gateways are configured with the site prefix(es) and advertise them into the external routing system, as normal site gateways do. If there is no external connectivity, then one or more MANET routers are configured with the site prefix(es), e.g., a Unique Local IPv6 Prefix [\[RFC4193\]](#), and 10/8 for IPv4 [\[RFC1918\]](#). It may be possible to automatically generate a site prefix for the site in a zero-configuration MANET (e.g., by electing one node to do so), but such a mechanism is outside the scope of this document. Note that a /48 IPv6 site prefix results in the ability to support up to 2^{16} MANET routers, plus any clients behind them.

Each router then uses a subnet allocation protocol to acquire one or more subnet prefixes within the site prefix. It then uses this

subnet prefix to number itself and any clients.

As a result, this approach needs only two MANET-specific protocols:

- 1) A routing protocol capable of exchanging subnet routes, and
- 2) A subnet allocation protocol capable of allocating prefixes.

Thaler	Expires August 2006	2
Draft	Multi-Subnet MANETs	February 2006

All other protocols and applications operate normally.

3. Client Behavior

Clients operate with no changes, using normal IPv4 and IPv6 mechanisms.

4. MANET Router Behavior

4.1. Obtaining a subnet prefix

The requirements for a subnet allocation protocol are:

- 1) Must be capable of allocating a subnet prefix (not just a single address).
- 2) Must prevent duplicate subnet prefix assignments.
- 3) Must handle network partitions and heals, due to the highly dynamic nature of a MANET.
- 4) Must complete rapidly, with low overhead.
- 5) Must support a bootstrap scenario where all MANET routers without a manually configured prefix attempt to allocate a prefix at the same time.
- 6) Should support the ability to detect duplicate layer-2 addresses.
- 7) Should support the ability to detect a duplicate IPv6 link-local address within the MANET site (since the dynamic nature means that any other router may eventually be on the same link, this helps avoid future conflicts).

Details of the actual subnet allocation protocol are outside the scope of this draft. At a high level, an address allocation protocol proposed for a multilink-subnet approach (approach B) can be modified to be suitable, by inclusion of a prefix length.

4.2. Using a subnet prefix

One common practice in BGP is to assign a routable address to the router's loopback interface, and use that for communication which may be routed over different physical interfaces over time (e.g., to survive failures).

In the same vein, a MANET router takes its allocated prefix, assigns it to its loopback interface, and advertises it in the MANET routing protocol. If the MANET router also has a client interface (e.g., an Ethernet link with wired clients attached), then the implementation can either allocate a second subnet prefix for that link, or may simply assign its sole subnet prefix to a client interface rather than its loopback interface. The MANET router then assigns its own

Thaler	Expires August 2006	3
Draft	Multi-Subnet MANETs	February 2006

IPv4 and/or IPv6 addresses on its loopback interface by using a fixed address suffix (e.g., .1 in IPv4 and ::1 in IPv6).

The MANET router's other interfaces (e.g., a wireless interface used to communicate with other MANET routers) is not given an IPv4 address (i.e., left unnumbered). An IPv6 link-local address is assigned per normal IPv6 stateless address autoconfiguration [[RFC2462](#)]. Note that normal IPv6 Duplicate Address Detection may occur as usual, but the MANET routers seen on that link will change over time. DAD need not be redone as long as the subnet allocation protocol is responsible for detecting duplicate IPv6 link-local addresses.

Applications running on the MANET router will use the loopback interface, since that is the one with the global address on it. All other MANET nodes will appear to the application as being off-subnet.

5. Security Considerations

Often the hardest attacks to defend against are on-link attacks (e.g., ARP attacks in IPv4). Unlike the other two approaches, this approach narrows the set of nodes which are on-link to only those nodes which are nearby, as opposed to being anywhere in the MANET. In addition to mitigations in various protocols, this may make social mitigations easier as well.

In addition, since this approach does not change the standard address allocation and forwarding mechanisms in IPv4 and IPv6, it avoids the security issues mentioned in [[MLSI](#)], resulting in a potentially more secure environment than with a multilink subnet approach.

The routing protocol in use must be secured to avoid attackers misrouting packets. The issues here are discussed in [[RFC2501](#)].

The subnet allocation protocol in use must also be secured to avoid attackers consuming all available subnet space, or preventing legitimate MANET routers from obtaining a subnet prefix.

6. IANA Considerations

This document has no actions for IANA.

7. Acknowledgements

Suresh Krishnan first suggested assigning a subnet prefix to each MANET node.

Thaler	Expires August 2006	4
Draft	Multi-Subnet MANETs	February 2006

8. Normative References

- [RFC1918] Rekhter, Y., Moskowitz, B., Karrenberg, D., de Groot, G., and E. Lear, "Address Allocation for Private Internets", [BCP 5](#), [RFC 1918](#), February 1996.
- [RFC2462] Thomson, S. and T. Narten, "IPv6 Stateless Address Autoconfiguration", [RFC 2462](#), December 1998.
- [RFC4193] Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses", [RFC 4193](#), October 2005.

9. Informative References

- [MLSI] Thaler, D., "Issues With Protocols Proposing Multilink Subnets", [draft-thaler-intarea-multilink-subnet-issues-00.txt](#), February 2006.
- [RFC2501] Corson, S. and J. Macker, "Mobile Ad hoc Networking (MANET): Routing Protocol Performance Issues and Evaluation Considerations", [RFC 2501](#), January 1999.

Authors' Addresses

Dave Thaler
Microsoft
One Microsoft Way
Redmond, WA 98052
Phone: +1 425 703 8835
Email: dthaler@microsoft.com

Thaler
Draft

Expires August 2006
Multi-Subnet MANETs

5
February 2006

Full Copyright Statement

Copyright (C) The Internet Society (2006).

This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC

documents can be found in [BCP 78](#) and [BCP 79](#).

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.