

Network Working Group
Internet-Draft
Intended status: Informational
Expires: July 18, 2018

T. Clausen
LIX, Ecole Polytechnique
U. Herberg
Fujitsu Laboratories of America
January 14, 2018

Snapshot of OLSRv2-Routed MANET Management
draft-clausen-manet-olsrv2-mgmt-snapshot-00

Abstract

This document describes how Mobile Ad Hoc Networks (MANETs) are typically managed, in terms of pre-deployment management, as well as rationale and means of monitoring and management of MANET routers running the Optimized Link State Routing protocol version 2 (OLSRv2) and its constituent MANET Neighborhood Discovery Protocol (NHDP). Apart from pre-deployment management for setting up IP addresses and security related credentials, OLSRv2 only needs routers to agree one single configuration parameter (called "C"). Other parameters for tweaking network performance may be determined during operation of the network, and need not be the same in all routers. This, using MIB modules and related management protocols such as SNMP (or possibly other, less "chatty", protocols). In addition, for debugging purposes, monitoring data and performance related counters, as well as notifications ("traps") can be sent to the Network Management System (NMS) via standardized management protocols.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on July 18, 2018.

Copyright Notice

Copyright (c) 2018 IETF Trust and the persons identified as the

document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	3
1.1.	Statement of Purpose	3
2.	Terminology	3
3.	Pre-Deployment Management	4
3.1.	Lower Layer Alignment	4
3.2.	Interface Addresses	4
3.3.	Security Material	5
3.4.	The Value of C	5
4.	How do we Manage OLSRv2-based MANETs?	6
4.1.	Internal Management	6
4.2.	External Management	6
5.	What and Why do we Manage and Monitor?	7
6.	Typical Communication Patterns	9
7.	This Document does not Constrain how to Manage and Monitor MANETs	11
8.	IANA Considerations	11
9.	Security Considerations	12
10.	Acknowledgments	12
11.	Informative References	13
	Authors' Addresses	14

1. Introduction

The MANET routing protocol OLSRV2 [[RFC7181](#)], as well as its constituent parts NHDP [[RFC6130](#)], [[RFC5497](#)], [[RFC5148](#)], [[RFC5444](#)], [[RFC7182](#)], [[RFC7183](#)], [[RFC7187](#)], [[RFC7188](#)] is designed to autonomously maintain routes across a dynamic network topology. OLSRV2 is designed so as to minimize operator intervention throughout the duration of a network deployment, and to allow for heterogeneous configuration of routers within the same network deployment: most configuration values are either of local significance only (e.g., message jitter parameters) or, when they are not, are carried in control signals exchanged between routers (e.g., information validity time).

All the same, a small set of configuration options must be established in each router prior to deployment, with some requiring agreement among all the routers within the same deployment. Furthermore, throughout the duration of a network deployment, external management and monitoring of a network may be desirable, e.g., for performance optimization or debugging purposes.

1.1. Statement of Purpose

Deployments of OLSRV2 are diverse, and may include community networks, constrained environments, tactical networks, etc. Each such environment may present distinctly different requirements as to management and monitoring.

This document does therefore explicitly not pretend to provide an exhaustive description of how all OLSRV2 network deployments should be managed and monitored - and does, specifically, not prescribe any management model. This document also does not address management of MANETs using any routing protocols, other than OLSRV2.

What this document does, however, is to present how typical OLSRV2 network deployments are managed and monitored, using well-established management patterns and well-known protocols. In particular, this document addresses several of the consideration from [[RFC5706](#)], in particular [Section 3](#) ("Management Considerations - How Will the Protocol Be Managed?").

2. Terminology

This document uses terminology from [[RFC7181](#)], [[RFC6130](#)], and [[RFC5497](#)].

3. Pre-Deployment Management

Prior to operation of an OLSRV2 network, or more precisely, prior to proper operation of OLSRV2 and its constituent parts, certain parameters need to be configured on each router. The following sections describe the required pre-deployment management.

3.1. Lower Layer Alignment

Interoperability between routers requires alignment of lower protocol layers below OLSRV2. In particular, all routers in the same MANET topology must be pre-configured to use the same IP address family (IPv4 or IPv6). In a single OLSRV2 topology, it is not possible to mix IPv4 and IPv6 addresses, notably because [\[RFC5444\]](#) messages can contain either IPv4 *or* IPv6 addresses, but not both at the same time. It is, however, possible to run two instances of OLSRV2, one instance for IPv4 and another one for IPv6, within the same network.

In addition to the IP address family, other lower layer parameters may also need to be aligned, e.g., MAC as well as radio channel selections. A single OLSRV2 topology may, of course, span different link layers (or the same link layer with different configuration settings such as cryptographic keys) when routers in the topology have OLSRV2 interfaces towards these different link layers.

3.2. Interface Addresses

According to [\[RFC6130\]](#), and as used by [\[RFC7181\]](#), each interface of a router must be configured with at least one IP address. [\[RFC6130\]](#) provides guidance as to the characteristics of such IP addresses, including the (limited) conditions under which a single IP address may be configured on multiple interfaces.

While automatic configuration of IP addresses on router interfaces is not excluded, currently no address autoconfiguration protocols have been standardized (in the IETF) to accomplish this. As a consequence, static configuration, or proprietary (as in: non-standardized) protocols ensure this.

Note that [\[RFC6130\]](#) and [\[RFC7181\]](#) permit to dynamically add or remove IP addresses as part of normal network operation. This applies for local MANET interfaces, as well as for local non-MANET interfaces or IP addresses from remote destinations reachable through this router (i.e., addresses for which this router serves as gateway). Interface addresses are managed by way of the Local Interface Set (as defined in [\[RFC6130\]](#)) and remote addresses by way of the Attached Network Set (as defined in [\[RFC7181\]](#)).

3.3. Security Material

Security material (keys, algorithms, etc.) must be available for generating Integrity Check Values (ICVs) for outgoing control messages, and to allow validating ICVs in incoming control messages [RFC7182] [RFC7183].

The appropriate way of making such security material available is dependent on the deployment type. For example, community networks (such as "Funkfeuer", <http://funkfeuer.at>), do currently not use any security at all. Other deployment types may use a simple manual shared key distribution mechanism, or may use a proprietary key distribution protocol. Tactical networks have much more stringent requirements for distributing key material, e.g., using manual distribution of the keys on encrypted USB flash drives, and with defensive mechanisms (up to and including mechanisms involving depleted uranium) if the keys are compromised.

In general, Automatic Key Management (AKM) as well as static/manual or other out-of-band mechanisms, can be viable options for distributing keys. Currently, no standardized AKM mechanism for MANETs exist. If the IETF standardizes such mechanisms in the future, for deployment types where such is appropriate, these can be used for distributing keys (with the obvious chicken-and-egg problem of using the routing fabric that is being constructed to distribute the keys to establish that fabric). Until such an AKM mechanism is standardized, manual key distribution as well as proprietary mechanisms prevail.

The important point to make here, however, is that by whichever method (automatic/manual, dynamic/static, ...) a key and other security material is made available, the security mechanisms of OLSRV2, as defined by [RFC7183], will be able to properly use it for generating and validating ICVs.

3.4. The Value of C

The only pre-deployment configuration parameter that directly impacts protocol operation is the value of C. This value is used by each router for calculating the representation of interval and validity time, as included in control messages. All routers in a deployment must agree on the value of C, as described in [RFC5497]. Note that since all MANET routers inside a MANET must agree to the same value of C before deployment, C is denoted "constant" in [RFC5497] rather than "parameter" as in this document. From a management perspective, C can be considered as configuration parameter prior to operation of the routing protocol.

4. How do we Manage OLSRv2-based MANETs?

A deployed OLSRv2 network is, as previously discussed, operating autonomously, but occasionally with internal or external management operations being desirable, described in the following two sections.

4.1. Internal Management

Internal management describes a local process running on a router that automatically (i.e., without external messaging or human interaction) modifies the configuration of OLSRv2 based on different environmental factors. In particular, message intervals can be updated dynamically and without external management interaction, e.g., the HELLO interval may be updated according to the rate of topology changes measured (or, inferred: after all, the 'M' in MANET is for "Mobility") locally: if the rate is high, then a more frequent HELLO update assures that routes are more accurate. At a lower rate of topology changes, network capacity and energy capacity of the router may be conserved by increasing the HELLO interval. In addition to message intervals, minimum intervals can have a significant impact on the operation of OLSRv2, and therefore need to be adjusted with care. If, for instance, the minimum interval between two successive HELLO messages (HELLO_MIN_INTERVAL) is set too low, many messages may be sent within a short timeframe, potentially leading to frame collisions or exhaustion of the available bandwidth.

Depending on the use case, many different automatic configuration agents can be envisioned. As parameters in NHDP and OLSRv2 are either only used locally or, in the case of HELLO_INTERVAL and REFRESH_INTERVAL, are selected locally and then included in the messages exchanged between adjacent routers in their HELLO messages, none of these automatic local configuration methods need necessarily to be standardized: different routers doing different things will interoperate.

4.2. External Management

For the deployments described by this document (but, see [Section 7](#)), external management operations are undertaken by a central Network Management Station (NMS).

The MIB modules developed for OLSRv2 [[RFC7184](#)] and for its constituent protocol NHDP [[RFC6779](#)] are verbose, in as much as that they expose for interrogation the complete protocol and router state, as well as enable setting all parameters (timer intervals, time-outs, jitter values etc.). They do explicitly not enable setting the value of C (as that is required to be constant and uniform across the network, see [Section 3.4](#)), nor distributing security material (see

[Section 3.3](#)).

In some deployments, the NMS communicates with individual routers by way of SNMP - and, more commonly, by way of "proprietary" simpler, less verbose and (often) less secure protocols, and over UDP. Note that this does not constitute a recommendation, but rather an observation that (apparently) SNMP has found less application in MANETs. The "Writable MIB Module IESG Statement" (<http://www.ietf.org/iesg/statement/writable-mib-module.html>) recommends to use MIB modules for read-only operations only, and to use YANG/NETCONF for read-write operations instead. While publication of the MIB modules developed for OLSRV2 and NHDP predates this statement, it may be possible to translate read-only objects from the MIB modules into YANG modules using [RFC6643]. A complete YANG model representing similar objects as in the MIB modules could, of course, be developed.

The predecessor of OLSRV2, OLSR [RFC3626] did not have any associated MIB module. Many deployments of OLSR did not support network management operations per se (i.e., configuration-on-launch was the way in which routers in these deployments were managed). Those implementations and deployments of OLSR that did support network management operations used a similar architecture to the one described in this document, but with "proprietary" protocols and APIs for parameters and router states, "proprietary" data-models, etc. It can be speculated that the "proprietary" protocols used for communication between the NMS and the MIB modules on each router also for OLSRV2, in part, exist as inherited from the protocols used for OLSR. Aligned with the recommendations from [RFC5706], management of OLSRV2 (in the form of the MIB modules for OLSRV2 and NHDP) has been developed alongside the standardization process of OLSRV2, rather than as an afterthought.

Finally, it is uncommon to see an NMS permanently active in a deployed OLSRV2 network; rather, on an "ad hoc" basis an NMS is introduced into the network, parameters configured or state interrogated, following which the NMS disappears. Part of the rationale for this is that in a MANET, network connectivity from every MANET router to an NMS cannot be guaranteed at all times due to the dynamicity of the network topology.

5. What and Why do we Manage and Monitor?

As indicated earlier, OLSRV2 and its constituent protocol NHDP, are reasonably robust with respect to parameter values: a deployment can operate with different parameters used in different routers in the same network. That being said, adapting these parameters according

to circumstances is (often) desired. For example, in a stable network (such as a wired network), TC messages may be sent infrequently and with long validity times, whereas in a highly dynamic network (such as in a vehicular network) TC messages may need to be sent more frequently and HELLO messages for discovering the local topology (almost) continuously. Note that for highly dynamic topologies, an alternative to sending control messages very frequently is to use long message intervals in combination with all of the permitted responsive mechanisms (e.g., to send an externally triggered HELLO when the local topology of a router changes) and with low minimum intervals. In this case, it is possible though that one control message may get lost, and then OLSRV2 needs to react in order to avoid a long convergence time. (One possibility is to reduce HELLO_INTERVAL to minimum for a few HELLO messages, then restore it).

In a similar vein, the message emission intervals and the information validity times should also be commensurate with the available network capacity: millisecond intervals between TC messages, for example, will consume all available network capacity whereas hourly intervals will be inappropriate even for a static and stable, wired, network (by way of simply new routers arriving in the network, which will not "learn" the network topology before undue long delays).

These adaptations may be imparted (i) autonomously by a central NMS monitoring and adopting the parameters globally, (ii) autonomously by an NMS in each router monitoring its local topology (and its stability) and adapting parameters locally, or (iii) by manual operator intervention.

Given the dynamic and evolutive topology of OLSRV2 networks, a highly desirable property of an NMS is the ability to display and offer visibility of the current network status - for example, to display a graphical map of which routers are currently part of the network. As a proactive protocol, OLSRV2 maintains, in each router, a topology map including all destinations and a subset of the links present in the network (particularly true in a very dense network). A typical feature of an NMS is to inquire as to the topology map of a single router. A less typical feature is to inquire all (or, at least, many) routers in a network, with the purpose of presenting a complete topology map.

In addition to actively monitoring an OLSRV2 network, erroneous or unusual conditions on a router can be flagged, e.g., detection of an unusually high number of 1-hop or 2-hop neighborhood changes in a short amount of time, indicating potential problems in that area of the network. [[RFC6779](#)] and [[RFC7184](#)] facilitate proactively sending "notifications" (also known as traps) from the router towards an NMS. The MIB modules defined in [[RFC6779](#)] and [[RFC7184](#)] allow for defining

both the threshold and the time window of how many times this erroneous condition may occur in the time window before the notification is sent to the NMS. Once the NMS receives a notification, a network operator may investigate if there is a problem that needs to be resolved, e.g., by changing parameters via the above-described external management.

6. Typical Communication Patterns

This section describes typical (management) communications patterns in an operating (post-startup) network. One of the key characteristics of OLSRV2 is that it enables an efficient flooding mechanism (denoted "MPR Flooding"). For some management scenarios, this facilitates better performance by (scope-limited) flooding management requests to MANET routers, rather than sending multiple consecutive unicast messages. While the MIB modules developed for OLSRV2 and NHDP do not support such broadcast operation (due to the nature of SNMP), some of the proprietary management tools mentioned in [Section 4](#) take advantage of this for increased performance.

The below list of such communication patterns is not claimed to be exhaustive, and depending on the deployment, different patterns may be used. However, these patterns have been observed in many deployments of OLSRV2 and its constituent parts, as well as of its predecessor OLSR.

- a) Inquire the state (complete topology graph, link states, and local links - even those not part of topology graph) of a router. An NMS would typically initiate that request. OLSRV2 contains a number of "Information Bases"; basically, tables with rows representing information about local interfaces, other routers in the MANET or the topology of the MANET as perceived by the MANET router. These tables are also reflected as objects in the MIB modules and can be inquired via, e.g., GETBULK for getting multiple rows in a single request. Depending on the number of MANET routers in the network and on the density of the MANET, these "Information Bases" of a router describing one-hop and two-hop routers, as well as routers farther away in the network, can contain a substantial amount of information. Therefore, an NMS inquiring for a complete copy of them will return multiple KB or more of data. Given the dynamic topology and often bandwidth-constrained wireless links between MANET routers, this is not a very commonly observed operation. Moreover, this would typically only be required in debugging situations, as during regular operations, OLSRV2 updates the state automatically and reacts to changes (e.g., by triggering control message generation). This type of operation can benefit from the optimized flooding

mechanism, by requesting the state from multiple routers in a region of the MANET in a single request.

- b) Inquire the history/statistics of a router. This request, initiated by an NMS, is typically a small inquiry, such as "how many local link changes have occurred within the past n minutes/seconds/hours". This may be a rare occurrence, or it may be occur several times per minute and per router, at least for some time: for example, an NMS may attempt to, e.g., "tune" message intervals and timers, by sending this request to a group of topologically close routers - and do so until the NMS decides that the topology has stabilized. Again, this feature of requesting performance related information is supported by the MIB modules for OLSRV2 and NHDP. While SNMP does not support sending the inquiry via optimized flooding, proprietary protocols take advantage of the optimized flooding mechanism, to reduce the number of unicast requests.
- c) Change the configuration of a router. Other than in the above case in b) (tuning), this really happens only when somehow a router gets a new uplink to an external network, and either a new gateway is added into the network, and/or a new prefix needs to be distributed to the routers. The MIB modules for OLSRV2 and NHDP allow to set all configuration parameters of each router. Optimized flooding may significantly reduce the amount of unicast requests, but are not supported by SNMP.
- d) Visualizing the network as a router sees it. As in a MANET, routers may move and link quality may vary due to link layer characteristics, the network topology may change frequently. In a naive way, this would essentially be the NMS setting up a connection to the router in question, and getting a copy of all routing protocol control messages to construct its own topology graph as would have done that router. Typically, it consists of the router sending a notification to the NMS when a topological change happens, i.e., when either of its information bases change. Even better, it consists of these notifications being "filtered" to only send for those changes that actually impact the usable topology. The latter case is supported by the MIB modules for OLSRV2 and NHDP in the form of notifications (also called "traps") that are sent from the MANET router to the NMS. While these notifications alone do not allow the NMS to visualize the topology, they may suffice to inform the NMS of an unusual change of the topology, and the NMS may inquire the current topology via the process described in a).

- e) Rekeying There is currently no (standardized) mechanism for automated key management. One of the reasons for this may be that it is difficult to come up with a single such that will satisfy the requirements for all the different deployments. However in MANET deployments rekeying is something that can be observed, e.g., as part of the parameter configuration. The particularity of this is, that it often is a "broadcast configuration operation" where new key material is supposed to be sent to everybody, and not just a single router, e.g., leveraging the optimized flooding mechanism of OLSRV2.

7. This Document does not Constrain how to Manage and Monitor MANETs

As explained in [Section 1](#), this document describes how, what, and why, some (typical) OLSRV2 networks are managed and monitored as of 2018. As such, the document is reflective, not prescriptive: it does not stipulate requirements for how to manage OLSRV2 networks, nor does it claim to be a complete list of all management patterns or protocols. Other ways of managing an OLSRV2 network are very well imaginable - now, or in future deployments of OLSRV2.

As an example of such a "future management scenario", rather than managing and monitoring routers from a central NMS, a distributed, autonomous management system between multiple routers can be envisioned. In particular, monitoring data that is used to debug network problems and to tweak inefficiencies could be distributed amongst a group of routers in the same network. This would both address problems of single point of failure when using only a single NMS, as well provide additional information about groups of multiple routers, rather than a single router. An example use for such a distributed information flow would be to identify areas of a network wherein, e.g., due to different router densities, message sending interval parameters could be exchanged and optimal values negotiated between routers, so as to obtain locally optimized performance.

While such a management model is highly interesting, it is also at present entirely fictional - at least outside the realm of research. It is included to, both, indicate directions being explored (but not exploited), and to insist that the intent of this document is not to prescribe how MANETs are to be managed, in the presence or in the future, but to describe the (known) state of how MANETs are managed, presently.

8. IANA Considerations

This document has no actions for IANA.

[This section may be removed by the RFC Editor.]

9. Security Considerations

This document does not specify a protocol, nor does it provide recommendations for how to manage an OLSRv2 deployment - rather, it reflects how some known deployments of OLSRv2 (and its predecessor, OLSR) have been known to be managed.

With that being said, managing an OLSRv2 network requires the ability to inspect and affect the internal state of the routers therein, by way of mechanisms other than the protocol signals specified for OLSRv2 [[RFC7181](#)] and NHDP [[RFC6130](#)].

When affecting the state of the OLSRv2 routing process, a management process can be considered as an "outside process" to OLSRv2 and is then expected to respect (at least) the constraints given in [Section 5.5](#), [Section 5.6](#), and in [Appendix A of \[RFC7181\]](#), as well as in [Section 5.5](#) and in [Appendix B of \[RFC6130\]](#). The example from [Section 4.1](#) of setting excessively short message intervals, leading to channel capacity exhaustion and frame collisions, demonstrates that such an outside process can harm network stability considerably when not carefully protected against unauthorized or unintended usage.

For both inspecting and affecting the state of an OLSRv2 routing process by way of a management interface, great care is necessary to avoid divulging information that should not be exposed, and in opening additional vulnerabilities by way of the management interface. In part, to be able to benefit from securing existing management interfaces, protocols, and implementations, migration to a standardized management framework is desired, and was one of the motivators for standardizing MIB modules for OLSRv2 and NHDP in the first place.

10. Acknowledgments

The authors would like to gratefully acknowledge the following people for intense technical discussions, early reviews, and comments on the documents: Alan Cullen (BAE Systems), Christopher Dearlove (BAE Systems), Adrian Farrel (Juniper), David Harrington (Comcast), and Jurgen Schoenwaelder (Jacobs University).

11. Informative References

- [RFC3626] Clausen, T. and P. Jacquet, "The Optimized Link State Routing Protocol", [RFC 3626](#), October 2003.
- [RFC5148] Clausen, T., Dearlove, C., and B. Adamson, "Jitter Considerations in Mobile Ad Hoc Networks (MANETs)", [RFC 5148](#), February 2008.
- [RFC5444] Clausen, T., Dearlove, C., Dean, J., and C. Adjih, "Generalized MANET Packet/Message Format", [RFC 5444](#), February 2009.
- [RFC5497] Clausen, T. and C. Dearlove, "Representing Multi-Value Time in Mobile Ad Hoc Networks (MANETs)", [RFC 5497](#), March 2009.
- [RFC5706] Harrington, D., "Guidelines for Considering Operations and Management of New Protocols and Protocol Extensions", [RFC 5706](#), November 2009.
- [RFC6130] Clausen, T., Dearlove, C., and J. Dean, "Mobile Ad Hoc Network (MANET) Neighborhood Discovery Protocol (NHDP)", [RFC 6130](#), April 2011.
- [RFC6643] Schoenwaelder, J., "Translation of Structure of Management Information Version 2 (SMIV2) MIB Modules to YANG Modules", [RFC 6643](#), July 2012.
- [RFC6779] Herberg, U., Cole, R., and I. Chakeres, "Definition of Managed Objects for the Neighborhood Discovery Protocol", [RFC 6779](#), May 2012.
- [RFC7181] Clausen, T., Dearlove, C., Jacquet, P., and U. Herberg, "The Optimized Link State Routing Protocol Version 2", [RFC 7181](#), April 2014.
- [RFC7182] Herberg, U., Clausen, T., and C. Dearlove, "Integrity Check Value and Timestamp TLV Definitions for Mobile Ad Hoc Networks (MANETs)", [RFC 7182](#), April 2014.
- [RFC7183] Herberg, U., Dearlove, C., and T. Clausen, "Integrity Protection for the Neighborhood Discovery Protocol (NHDP) and Optimized Link State Routing Protocol Version 2 (OLSRv2)", [RFC 7183](#), April 2014.
- [RFC7184] Herberg, U., Cole, R., and T. Clausen, "Definition of Managed Objects for the Optimized Link State Routing

Protocol Version 2", [RFC 7184](#), April 2014.

[RFC7187] Dearlove, C. and T. Clausen, "Routing Multipoint Relay Optimization for the Optimized Link State Routing Protocol Version 2 (OLSRv2)", [RFC 7187](#), April 2014.

[RFC7188] Dearlove, C. and T. Clausen, "Optimized Link State Routing Protocol Version 2 (OLSRv2) and MANET Neighborhood Discovery Protocol (NHDP) Extension TLVs", [RFC 7187](#), April 2014.

Authors' Addresses

Thomas Clausen
LIX, Ecole Polytechnique
91128 Palaiseau Cedex,
France

Phone: +33-6-6058-9349
Email: T.Clausen@computer.org
URI: <http://www.thomasclausen.org>

Ulrich Herberg
Fujitsu Laboratories of America
1240 E Arques Ave
Sunnyvale CA 94086,
US

Phone:
Email: ulrich@herberg.name
URI: <http://www.herberg.name>

