

Network Working Group
Internet-Draft
Intended status: Experimental
Expires: April 2, 2013

V. Fuller
D. Lewis
V. Ermagan
A. Jain
cisco Systems
September 29, 2012

LISP Delegated Database Tree
draft-fuller-lisp-ddt-04.txt

Abstract

This draft describes the LISP Delegated Database Tree (LISP-DDT), a hierarchical, distributed database which embodies the delegation of authority to provide mappings from LISP Endpoint Identifiers (EIDs) to Routing Locators (RLOCs). It is a statically-defined distribution of the EID namespace among a set of LISP-speaking servers, called DDT nodes. Each DDT node is configured as "authoritative" for one or more EID-prefixes, along with the set of RLOCs for Map Servers or "child" DDT nodes to which more-specific EID-prefixes are delegated.

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 April 2, 2013.

Copyright Notice

Copyright (c) 2012 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	4
2.	Definition of Terms	6
3.	Database organization	8
3.1.	EID-prefix tree structure and instance IDs	8
3.2.	Configuring prefix delegation	8
3.2.1.	The root DDT node	8
4.	The Map-Referral message	10
4.1.	Action codes	10
4.2.	Referral Set	11
4.3.	Incomplete flag	11
5.	DDT network elements and their operation	12
5.1.	DDT node	12
5.1.1.	Match of a delegated prefix (or sub-prefix)	12
5.1.2.	Missing delegation from an authoritative prefix	12
5.2.	DDT Map Server	13
5.3.	DDT Map Resolver	13
5.3.1.	Queuing and sending DDT Map-Requests	13
5.3.2.	Receiving and following referrals	14
5.3.3.	Handling referral errors	15
5.3.4.	Referral loop detection	16
6.	Psuedo Code and Decision Tree diagrams	17
6.1.	Map Resolver processing of ITR Map-Request	17
6.1.1.	Pseudo-code summary	17
6.1.2.	Decision tree diagram	18
6.2.	Map Resolver processing of Map-Referral message	19
6.2.1.	Pseudo-code summary	19
6.2.2.	Decision tree diagram	21
6.3.	DDT Node processing of DDT Map-Request message	22
6.3.1.	Pseudo-code summary	22
6.3.2.	Decision tree diagram	23
7.	Example topology and request/referral following	24
7.1.	Lookup of 10.1.1.1/32 by ITR1	25
7.2.	Lookup of 10.17.8.1/32 by ITR2	26
7.3.	Lookup of 10.2.2.2/32 by ITR1	27
7.4.	Lookup of 10.16.2.1/32 by ITR2	27
7.5.	Lookup of 10.16.0.1/32 (non-existant EID) by ITR2	28
8.	Securing the database and message exchanges	29
8.1.	XEID-prefix Delegation	29
8.2.	DDT node operation	30

8.2.1.	DDT public key revocation	30
8.3.	Map Server operation	30
8.4.	Map Resolver operation	31
9.	Open Issues and Considerations	32
10.	IANA Considerations	33
11.	Security Considerations	34
12.	References	35
12.1.	Normative References	35
12.2.	Informative References	35
Appendix A.	Acknowledgments	36
Appendix B.	Map-Referral Message Format	37
B.1.	SIG section	39
Appendix C.	Encapsulated Control Message Format	41
Authors' Addresses	42

1. Introduction

[LISP] specifies an architecture and mechanism for replacing the addresses currently used by IP with two separate name spaces: relatively static Endpoint Identifiers (EIDs), used end-to-end for terminating transport-layer associations, and Routing Locators (RLOCs), which are more dynamic, are bound to topological location, and are used for routing and forwarding through the Internet infrastructure.

LISP offers a general-purpose mechanism for mapping between EIDs and RLOCs. In organizing a database of EID to RLOC mappings, this specification extends the definition of the EID numbering space by logically prepending and appending several fields for purposes of defining the database index key: Database-ID (DBID, 16 bits), Instance identifier (IID, 32-bits), Address Family Identifier (16 bits), and EID-prefix (variable, according to AFI value). The resulting concatenation of these fields is termed an "Extended EID prefix" or XEID-prefix.

The term "Extended EID" (XEID) is also used for an individual LISP EID that is further qualified through the use of an Instance ID. See [LCAF] for further discussion of the use of Instance IDs.

The DBID is provided for possible use in case a need evolves for another, higher level in the hierarchy, to allow the creation of multiple, separate database trees.

LISP-DDT is a hierarchical distributed database which embodies the delegation of authority to provide mappings, i.e. its internal structure mirrors the hierarchical delegation of address space. It also provides delegation information to Map Resolvers, which use the information to locate EID-to-RLOC mappings. A Map Resolver which needs to locate a given mapping will follow a path through the tree-structured database, contacting, one after another, the DDT nodes along that path until it reaches the leaf DDT node(s) authoritative for the mapping it is seeking.

LISP-DDT defines a new device type, the DDT node, that is configured as authoritative for one or more XEID-prefixes. It also is configured with the set of more-specific sub-prefixes that are further delegated to other DDT nodes. To delegate a sub-prefix, the "parent" DDT node is configured with the RLOCs of each child DDT node that is authoritative for the sub-prefix. Each RLOC either points to a Map Server (sometimes termed a "terminal DDT node") to which an Egress Tunnel Routers (ETRs) has registered that sub-prefix or points to another DDT node in the database tree that further delegates the sub-prefix. See [LISP-MS] for a description of the functionality of

the Map Server and Map Resolver. Note that the target of a delegation must always be an RLOC (not an EID) to avoid any circular dependency.

To provide a mechanism for traversing the database tree, LISP-DDT defines a new LISP message type, the Map-Referral, which is returned to the sender of a Map-Request when the receiving DDT node can refer the sender to another DDT node that has more detailed information. See [Section 4](#) for the definition of the Map-Referral message.

To find an EID-to-RLOC mapping, a LISP-DDT client, usually a DDT Map Resolver, starts by sending an Encapsulated Map-Request to a preconfigured DDT node RLOC. The DDT node responds with a Map-Referral message that either indicates that it will find the requested mapping to complete processing of the request or that the DDT client should contact another DDT node that has more-specific information; in the latter case, the DDT node then sends a new Encapsulated Map-Request to the next DDT node and the process repeats in an iterative manner.

Conceptually, this is similar to the way that a client of the Domain Name System (DNS) follows referrals (DNS responses that contain only NS records) from a series of DNS servers until it finds an answer.

2. Definition of Terms

Extended EID (XEID): a LISP EID, optionally extended with a non-zero Instance ID (IID) if the EID is intended for use in a context where it may not be a unique value, such as on a Virtual Private Network where [\[RFC1918\]](#) address space is used. See "Using Virtualization and Segmentation with LISP" in [\[LISP\]](#) for more discussion of Instance IDs.

XEID-prefix: a LISP EID-prefix with 16-bit LISP-DDT DBID (provided to allow the definition of multiple databases; currently always zero in this version of DDT, with other values reserved for future use), 32-bit IID and 16-bit AFI prepended. An XEID-prefix is used as a key index into the database.

DDT node: a network infrastructure component responsible for specific XEID-prefix and for delegation of more-specific sub-prefixes to other DDT nodes.

DDT client: a network infrastructure component that sends Map-Request messages and implements the iterative following of Map-Referral results. Typically, a DDT client will be a Map Resolver but it is also possible for an ITR to implement DDT client functionality.

DDT Map Server: a DDT node that also implements Map Server functionality (forwarding Map-Requests and/or returning Map-Replies if offering proxy Map-Reply service) for a subset of its delegated prefixes.

DDT Map Resolver: a network infrastructure element that accepts a Map-Request, adds the XEID to its pending request list, then queries one or more DDT nodes for the requested EID, following returned referrals until it receives one with action code MS-ACK (or an error indication). MS-ACK indicates that the Map-Request has been sent to a Map Server that will forward it to an ETR that, in turn, will provide a Map-Reply to the original sender. A DDT Map Resolver maintains both a cache of Map-Referral message results containing RLOCs for DDT nodes responsible for XEID-prefixes of interest (termed the "referral cache") a pending request list of XEIDs that are being resolved through iterative querying of DDT nodes.

Encapsulated Map-Request: a LISP Map-Request carried within an Encapsulated Control Message, which has an additional LISP header prepended. Sent to UDP destination port 4342. The "outer" addresses are globally-routable IP addresses, also known as RLOCs. Used by an ITR when sending to a Map Resolver and by a Map Server

when forwarding a Map-Request to an ETR as documented in [[LISP-MS](#)].

DDT Map-Request: an Encapsulated Map-Request sent by a DDT client to a DDT node. The "DDT-originated" flag is set in the encapsulation header indicating that the DDT node should return Map-Referral messages if the Map-Request EID matches a delegated XEID-prefix known to the DDT node. [Section 5.3.1](#) describes how DDT Map-Requests are sent.

Authoritative XEID-prefix: an XEID-prefix delegated to a DDT node and for which the DDT node may provide further delegations of more-specific sub-prefixes.

Map-Referral: a LISP message sent by a DDT node in response to a DDT Map-Request for an XEID that matches a configured XEID-prefix delegation. A non-negative Map-Referral includes a "referral", a set of RLOCs for DDT nodes that have more information about the sub-prefix; a DDT client "follows the referral" by sending another DDT Map-Request to one of those RLOCs to obtain either an answer or another referral to DDT nodes responsible for a more-specific XEID-prefix. See [Section 5.1](#) and [Section 5.3.2](#) for details on the sending and processing of Map-Referral messages.

negative Map-Referral: a Map-Referral sent in response to a DDT Map-Request that matches an authoritative XEID-prefix but for which there is no delegation configured (or no ETR registration if sent by a DDT Map-Server).

Pending Request List: the set of outstanding requests for which a DDT Map Resolver has received encapsulated Map-Requests from a DDT client for an XEID. Each entry in the list contains additional state needed by the referral following process, including the requestor(s) of the XEID (typically, one or more ITRs), saved information about the last referral received and followed (matching XEID-prefix, action code, RLOC set, index of last RLOC queried in the RLOC set), and any [[LISP-SEC](#)] information that was included in the DDT client Map-Request. An entry in the list may be interchangeably termed a "pending request list entry" or simply a "pending request".

For definitions of other terms, notably Map-Request, Map-Reply, Ingress Tunnel Router (ITR), Egress Tunnel Router (ETR), Map Server, and Map Resolver, please consult the LISP specification [[LISP](#)] and the LISP Mapping Service specification [[LISP-MS](#)].

3. Database organization

3.1. EID-prefix tree structure and instance IDs

LISP-DDT defines a tree structure that is indexed by a binary encoding of five fields, in order of significance: DBID (16 bits), Instance Identifier (IID, 32 bits), Address Family Identifier (AFI, 16 bits), and EID-prefix (variable, according to AFI value). The fields are concatenated, with the most significant fields as listed above. The index into this structure is also referred to as an Extended EID-prefix (XEID-prefix).

It is important to note that LISP-DDT does not store actual EID-to-RLOC mappings; it is, rather, a distributed index that can be used to find the devices (Map Servers and their registered EIDs) that can be queried with LISP to obtain those mappings. Changes to EID-to-RLOC mappings are made on the ETRs which define them, not to any DDT node configuration. DDT node configuration changes are only required when branches of the database hierarchy are added, removed, or modified.

3.2. Configuring prefix delegation

Every DDT node is configured with one or more XEID-prefixes for which it is authoritative along with a list of delegations of XEID-prefixes to other DDT nodes. A DDT node is required to maintain a list of delegations for all sub-prefixes of its authoritative XEID-prefixes; it also may list "hints", which are prefixes that it knows about that belong to its parents, to the root, or to any other point in the XEID-prefix hierarchy. A delegation (or hint) consists of an XEID-prefix, a set of RLOCs for DDT nodes that have more detailed knowledge of the XEID-prefix, and accompanying security information. Those RLOCs are returned in Map-Referral messages when the DDT node receives a DDT Map-Request with an xEID that matches a delegation. A DDT Map Server will also have a set of sub-prefixes for which it accepts ETR mapping registrations and for which it will forward (or answer, if it provides proxy Map-Reply service) Map-Requests. For details of security information in Map-Referrals see [Section 8](#).

3.2.1. The root DDT node

The root DDT node is the logical "top" of the database hierarchy: DBID=0, IID=0, AFI=0, EID-prefix=0/0. A DDT Map-Request that matches no configured XEID-prefix will be referred to the root node. The root node in a particular instantiation of LISP-DDT must therefore be configured with delegations for at least all defined IIDs and AFIs.

To aid in defining a "sub-root" DDT node that is responsible for all EID-prefixes within multiple IIDs (say, for using LISP to create

virtual networks that use overlapping address space), it may be useful to implement configuration language that allows for a range of IIDs to be delegated together. Additional configuration shorthand for delegating a range of IIDs (and all of the EIDs under them) may also be helpful.

4. The Map-Referral message

A Map-Referral message is sent by a DDT node to a DDT client in response to a DDT Map-Request message. The message consists of an action code along with delegation information about the XEID-prefix that matches the XEID requested.

See [Appendix B](#) for a detailed layout of the Map-Referral message fields.

4.1. Action codes

The action codes are as follows:

NODE-REFERRAL (0): indicates that the replying DDT node has delegated an XEID-prefix that matches the requested XEID to one or more other DDT nodes. The Map-Referral message contains a "map-record" with additional information, most significantly the set of RLOCs to which the prefix has been delegated, that is used by a DDT Map Resolver to "follow" the referral.

MS-REFERRAL (1): indicates that the replying DDT node has delegated an XEID-prefix that matches the requested XEID to one or more DDT Map Servers. It contains the same additional information as a NODE-REFERRAL but is handled slightly differently by the receiving DDT client (see [Section 5.3.2](#)).

MS-ACK (2): indicates that a replying DDT Map Server received a DDT Map-Request that matches an authoritative XEID-prefix for which it has one or more registered ETRs. This means that the request can be forwarded to one of those ETRs to provide an answer to the querying ITR.

MS-NOT-REGISTERED (3): indicates that the replying DDT Map Server received a Map-Request for one of its configured XEID-prefixes which has no ETRs registered.

DELEGATION-HOLE (4): indicates that the requested XEID matches a non-delegated sub-prefix of the XEID space. This is a non-LISP "hole", which has not been delegated to any DDT Map Server or ETR. See [Section 5.1.2](#) for details.

NOT-AUTHORITATIVE (5): indicates that the replying DDT node received a Map-Request for an XEID-request for which it is not authoritative. This can occur if a cached referral has become invalid due to a change in the database hierarchy.

4.2. Referral Set

For "positive" action codes (NODE-REFERRAL, MS-REFERRAL, MS-ACK), a DDT node includes in the Map-Referral message a list of RLOCs for all DDT nodes that are authoritative for the XEID-prefix being returned; a DDT Map Resolver uses this information to contact one of those DDT nodes as it "follows" a referral.

4.3. Incomplete flag

A DDT node sets the "Incomplete" flag in a Map-Referral message if the Referral Set is incomplete; this is intended to prevent a DDT Map Resolver from caching a referral with incomplete information. A DDT node must set the "incomplete" flag in the following cases:

- o If it is setting action code MS-ACK or MS-NOT-REGISTERED but does not have configuration for other "peer" DDT nodes that are also authoritative for the matched XEID-prefix.
- o If it is setting action code NOT-AUTHORITATIVE.

5. DDT network elements and their operation

As described above, DDT introduces a new network element, the "DDT node", extends the functionality of Map Servers and Map Resolvers to send and receive Map-Referral messages. The operation of each of these devices is described as follows.

5.1. DDT node

When a DDT node receives a DDT Map-Request, it compares the requested XEID against its list of XEID-prefix delegations and its list of authoritative XEID-prefixes and acts as follows:

5.1.1. Match of a delegated prefix (or sub-prefix)

If the requested XEID matches one of the DDT node's delegated prefixes, then a Map-Referral message is returned with the matching more-specific XEID-prefix and the set of RLOCs for the referral target DDT nodes including associated security information (see [Section 8](#) for details on security). If the delegation is known to be a DDT Map Server, then the Map-Referral message is sent with action code MS-REFERRAL to indicate to the receiver that LISP-SEC information (if saved in the pending request) should be included in the next DDT Map-Request; otherwise, the action code NODE-REFERRAL is used.

Note that a matched delegation does not have to be for a sub-prefix of an authoritative prefix; in addition to being configured to delegate sub-prefixes of an authoritative prefix, a DDT node may also be configured with other XEID-prefixes for which it can provide referrals to DDT nodes anywhere in the database hierarchy. This capability to define "shortcut hints" is never required to be configured but may be a useful heuristic for reducing the number of iterations needed to find an EID, particular for private network deployments.

5.1.2. Missing delegation from an authoritative prefix

If the requested XEID did not match a configured delegation but does match an authoritative XEID-prefix, then the DDT node returns a negative Map-Referral that uses the least-specific XEID-prefix that does not match any XEID-prefix delegated by the DDT node. The action code is set to DELEGATION-HOLE; this indicates that the XEID is not a LISP destination.

If the requested XEID did not match either a configured delegation or an authoritative XEID-prefix, then the request is dropped and a negative Map-Referral with action code NOT-AUTHORITATIVE is returned.

5.2. DDT Map Server

When a DDT Map Server receives a DDT Map-Request, its operation is similar to that of a DDT node with additional processing as follows:

- o If the requested XEID matches a registered XEID-prefix, then the Map-Request is forwarded to one of the destination ETR RLOCs (or the Map Server sends a Map-Reply, if it is providing proxy Map-Reply service) and a Map-Referral with the MS-ACK action is returned to the sender of the DDT Map-Request.
- o If the requested XEID matches a configured XEID-prefix for which no ETR registration has been received then a negative Map-Referral with action code MS-NOT-REGISTERED is returned to the sender of the DDT Map-Request.

5.3. DDT Map Resolver

Just as any other Map Resolver, a DDT Map Resolver accepts Map-Requests from its clients (typically, ITRs) and ensures that those Map-Requests are forwarded to the correct ETR, which generates Map-Replies. Unlike a Map Resolver that uses the ALT mapping system [[LISP-ALT](#)], however, a DDT Map Resolver uses an iterative process of following referrals to find the correct ETR to answer a Map-Request; this requires a DDT Map Resolver to maintain additional state: a Map-Referral cache and pending request list of XEIDs that are going through the iterative referral process.

5.3.1. Queuing and sending DDT Map-Requests

When a DDT Map Resolver receives an encapsulated Map-Request, it first performs a longest-match search for the XEID in its referral cache. If no match is found or if a negative cache entry is found, then the destination is not in the database; a negative Map-Reply is returned and no further processing is performed by the DDT Map Resolver.

If a match is found, the DDT Map Resolver creates a pending request list entry for the XEID and saves the original Map-Request (minus the encapsulation header) along with other information needed to track progress through the iterative referral process; the "referral XEID-prefix" is also initialized to the null value since no referral has yet been received. The Map Resolver then creates a DDT Map-Request (which is an encapsulated Map-Request with the "DDT-originated" flag set in the message header) for the XEID but without any authentication data that may have been included in the original Map-Request. It sends the DDT Map-Request to one of the RLOCs in the chosen referral cache entry. The referral cache is initially

populated with one or more statically-configured entries; additional entries are added when referrals are followed, as described below. A DDT Map Resolver is not absolutely required to cache referrals but it doing so decreases latency and reduces lookup delays.

Note that in normal use on the public Internet, the statically-configured initial referral cache for a DDT Map Resolver should include a "default" entry with RLOCs for one or more DDT nodes that can reach the DDT root node. If a Map Resolver does not have such configuration, it will return a Negative Map-Reply if it receives a query for an EID outside the subset of the mapping database known to it. While this may be desirable on private network deployments or during early transition to LISP when few sites are using it, this behavior is not appropriate when LISP is in general use on the Internet.

5.3.2. Receiving and following referrals

After sending a DDT Map-Request, a DDT Map Resolver expects to receive a Map-Referral response. If none occurs within the timeout period, the DDT Map Resolver retransmits the request, sending to the next RLOC in the referral cache entry if one is available. If the maximum number of retransmissions has occurred and all RLOCs have been tried, then the pending request list entry is dequeued.

A DDT Map Resolver processes a received Map-Referral message according to each action code:

NODE-REFERRAL: The DDT Map Resolver checks for a possible referral loop as as described in [Section 5.3.4](#). If no loop is found, the DDT Map Resolver saves the prefix returned in the Map-Referral message in the referral cache, updates the saved prefix and saved RLOCs in the pending request list entry, and follows the referral by sending a new DDT Map-Request to one of the DDT node RLOCs listed in the Referral Set; security information saved with the original Map-Request is not included.

MS-REFERRAL: The DDT Map Resolver follows an MS-REFERRAL in the same manner as a NODE-REFERRAL except that that security information saved with the original Map-Request is included in the new Map-Request sent to a Map Server (see [Section 8](#) for details on security).

MS-ACK: This is returned by a DDT Map Server to indicate that it has one or more registered ETRs that can answer a Map-Request for the XEID. If the pending request did not include saved LISP-SEC information or if that information was already included in the previous DDT Map-Request (sent by the DDT Map Resolver in response

to either an MS-REFERRAL or a previous MS-ACK referral), then the pending request for the XEID is complete and is dequeued. Otherwise, LISP-SEC information is required and has not yet been sent to the authoritative DDT Map-Server; the DDT Map Resolver re-sends the DDT Map-Request with LISP-SEC information included and the pending request queue entry remains until another Map-Referral with MS-ACK action code is received. If the "incomplete" flag is not set, the prefix is saved in the referral cache.

MS-NOT-REGISTERED: The DDT Map Server queried could not process the request because it did not have any ETRs registered for a matching, authoritative XEID-prefix. If the DDT Map Resolver has not yet tried all of the RLOCs saved with the pending request, then it sends a Map-Request to the next RLOC in that list. If all RLOCs have been tried, then the destination XEID is not registered and is unreachable. The DDT Map Resolver returns a negative Map-Reply to the original Map-Request sender; this Map-Reply contains the non-registered XEID-prefix with TTL value of one minute. A negative referral cache entry is created for the prefix (also with TTL of one minute) and the pending request is dequeued.

DELEGATION-HOLE: The DDT Map Server queried did not have an XEID-prefix defined that matched the requested XEID so it does not exist in the mapping database. The DDT Map Resolver returns a negative Map-Reply to the original Map-Request sender; this Map-Reply will indicate the least-specific XEID-prefix matching the requested XEID for which no delegations exist and will have a TTL value of 15 minutes. A negative referral cache entry is created for the prefix (also with TTL of 15 minutes) and the pending request is dequeued.

NOT-AUTHORITATIVE: The DDT Map Server queried is not authoritative for the requested XEID. This can occur if a cached referral has become invalid due to a change in the database hierarchy. If the DDT Map Resolver receiving this message can determine that it is using old cached information, it may choose to delete that cached information and re-try the original Map-Request, starting from its "root" cache entry. If this action code is received in response to a query that was not to cached referral information, then it indicates a database synchronization problem or configuration error. The pending request list entry that caused this answer is removed, with no answer returned to the original requestor.

5.3.3. Handling referral errors

Other states are possible, such as a misconfigured DDT node (acting as a proxy Map Server, for example) returning a Map-Reply to the DDT Map Resolver; they should be considered errors and logged as such.

It is not clear exactly what else the DDT Map Resolver should do in such cases; one possibility is to remove the pending request list entry and send a negative Map-Reply to the original Map-Request sender. Alternatively, if a DDT Map Resolver detects unexpected behavior by a DDT node, it could mark that node as unusable in its referral cache and update the pending request to try a different DDT node if more than one is listed in the referral cache. In any case, any prefix contained in a Map-Referral message that causes a referral error (including a referral loop) is not saved in the DDT Map-Resolver referral cache.

5.3.4. Referral loop detection

In response to a Map-Referral message with action code NODE-REFERRAL or MS-REFERRAL, a DDT Map Resolver is directed to query a new set of DDT node RLOCs that are expected to have more-specific XEID-prefix information for the requested XEID. To prevent a possible "iteration loop" (following referrals back-and-forth among a set of DDT nodes without ever finding an answer), a DDT Map Resolver saves the last received referral XEID-prefix for each pending request and checks that a newly received NODE-REFERRAL or MS-REFERRAL message contains a more-specific referral XEID-prefix; an exact or less-specific match of the saved XEID-prefix indicates a referral loop. If a loop is detected, the DDT Map Resolver handles the request as described in [Section 5.3.3](#). Otherwise, the Map Resolver saves the most recently received referral XEID-prefix with the pending request when it follows the referral.

As an extra measure to prevent referral loops, it is probably also wise to limit the total number of referrals for any request to some reasonable number; the exact value of that number will be determined during experimental deployment of LISP-DDT but is bounded by the maximum length of the XEID.

Note that when a DDT Map Resolver adds an entry to its lookup queue and sends an initial Map-Request for an XEID, the queue entry has no previous referral XEID-prefix; this means that the first DDT node contacted by a DDT Map Resolver may provide a referral to anywhere in the DDT hierarchy. This, in turn, allows a DDT Map Resolver to use essentially any DDT node RLOCs for its initial cache entries and depend on the initial referral to provide a good starting point for Map-Requests; there is no need to configure the same set of root DDT nodes on all DDT Map Resolvers.

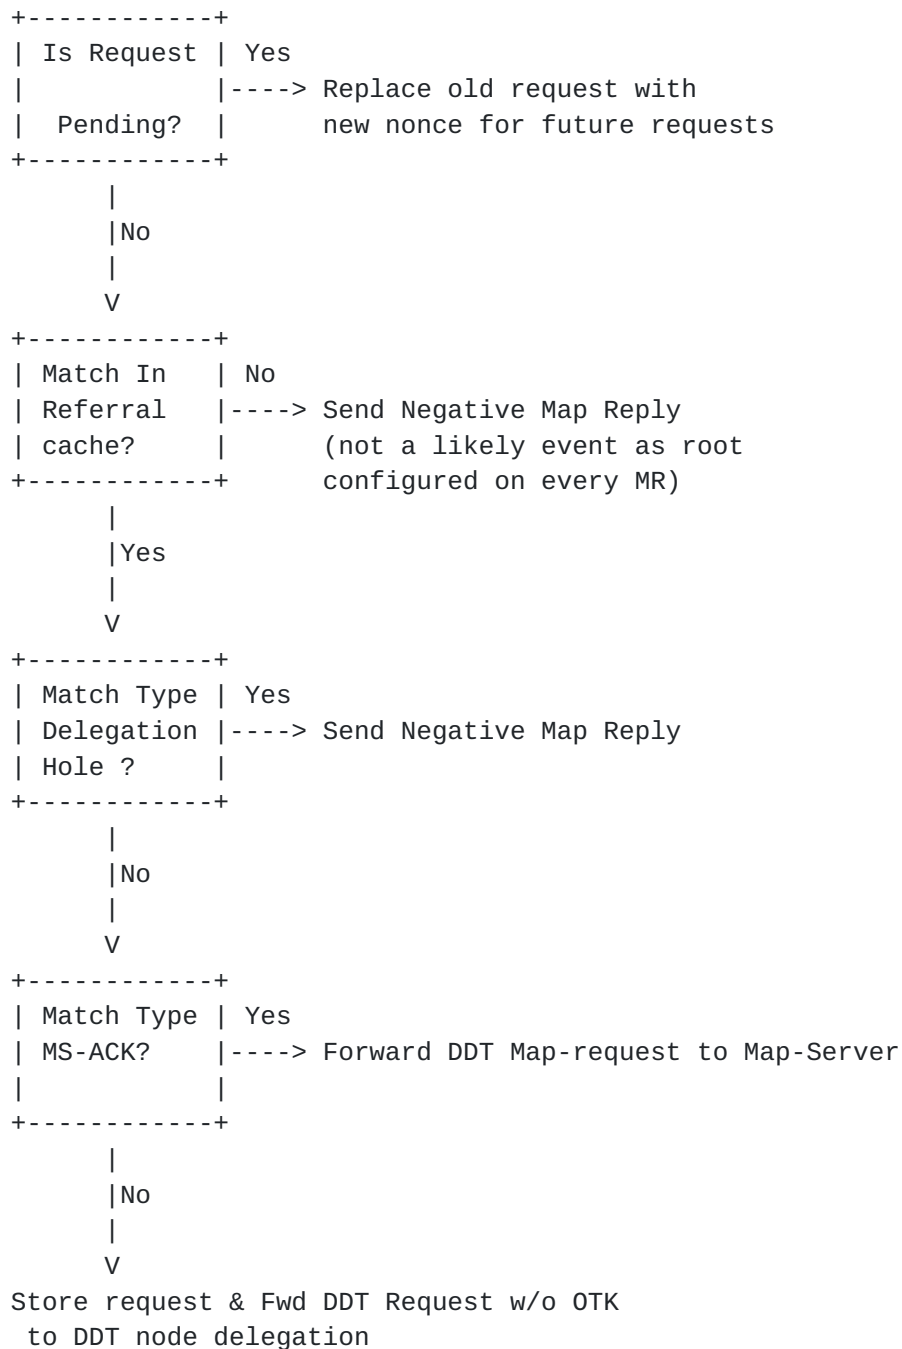
6. Psuedo Code and Decision Tree diagrams

To aid in implementation, each of the major DDT Map Server and DDT Map Resolver functions are described below, first using simple "psuedo-code" and then in the form of a decision tree.

6.1. Map Resolver processing of ITR Map-Request

6.1.1. Pseudo-code summary

```
if ( request pending i.e., (ITR,EID) of request same ) {  
    replace old request with new & use new request nonce  
    for future requests  
} else if ( no match in refcache ) {  
    return negative map-reply to ITR  
} else if ( match type delegation hole ) {  
    return negative map-reply to ITR  
} else if ( match type ms-ack ) {  
    fwd DDT request to map-server  
} else {  
    store & fwd DDT request w/o OTK to node delegation  
}
```


6.1.2. Decision tree diagram

6.2. Map Resolver processing of Map-Referral message

6.2.1. Pseudo-code summary

```
if ( no request pending matched by referral nonce ) {
    silently drop
}

if ( pfx in referral less specific than last referral used ) {
    if ( gone through root ) {
        silently drop
    } else {
        goto root
    }
}

switch (map_referral_type) {

    case NOT_AUTHORITY :
        if ( gone through root ) {
            return negative map-reply to ITR
        } else {
            goto root
        }

    case DELEGATION_HOLE:
        cache & send negative map-reply to ITR

    case MS_REFERRAL:
        if ( referral equal to last used ) {
            if ( gone through root ) {
                return negative map-reply to ITR
            } else {
                goto root
            }
        } else {
            cache & follow the referral
        }

    case NODE_REFERRAL:
        if ( referral equal to last used ) {
            if ( gone through root ) {
                return negative map-reply to ITR
            } else {
                goto root
            }
        } else {
            cache & follow the referral
        }
}
```



```
    }

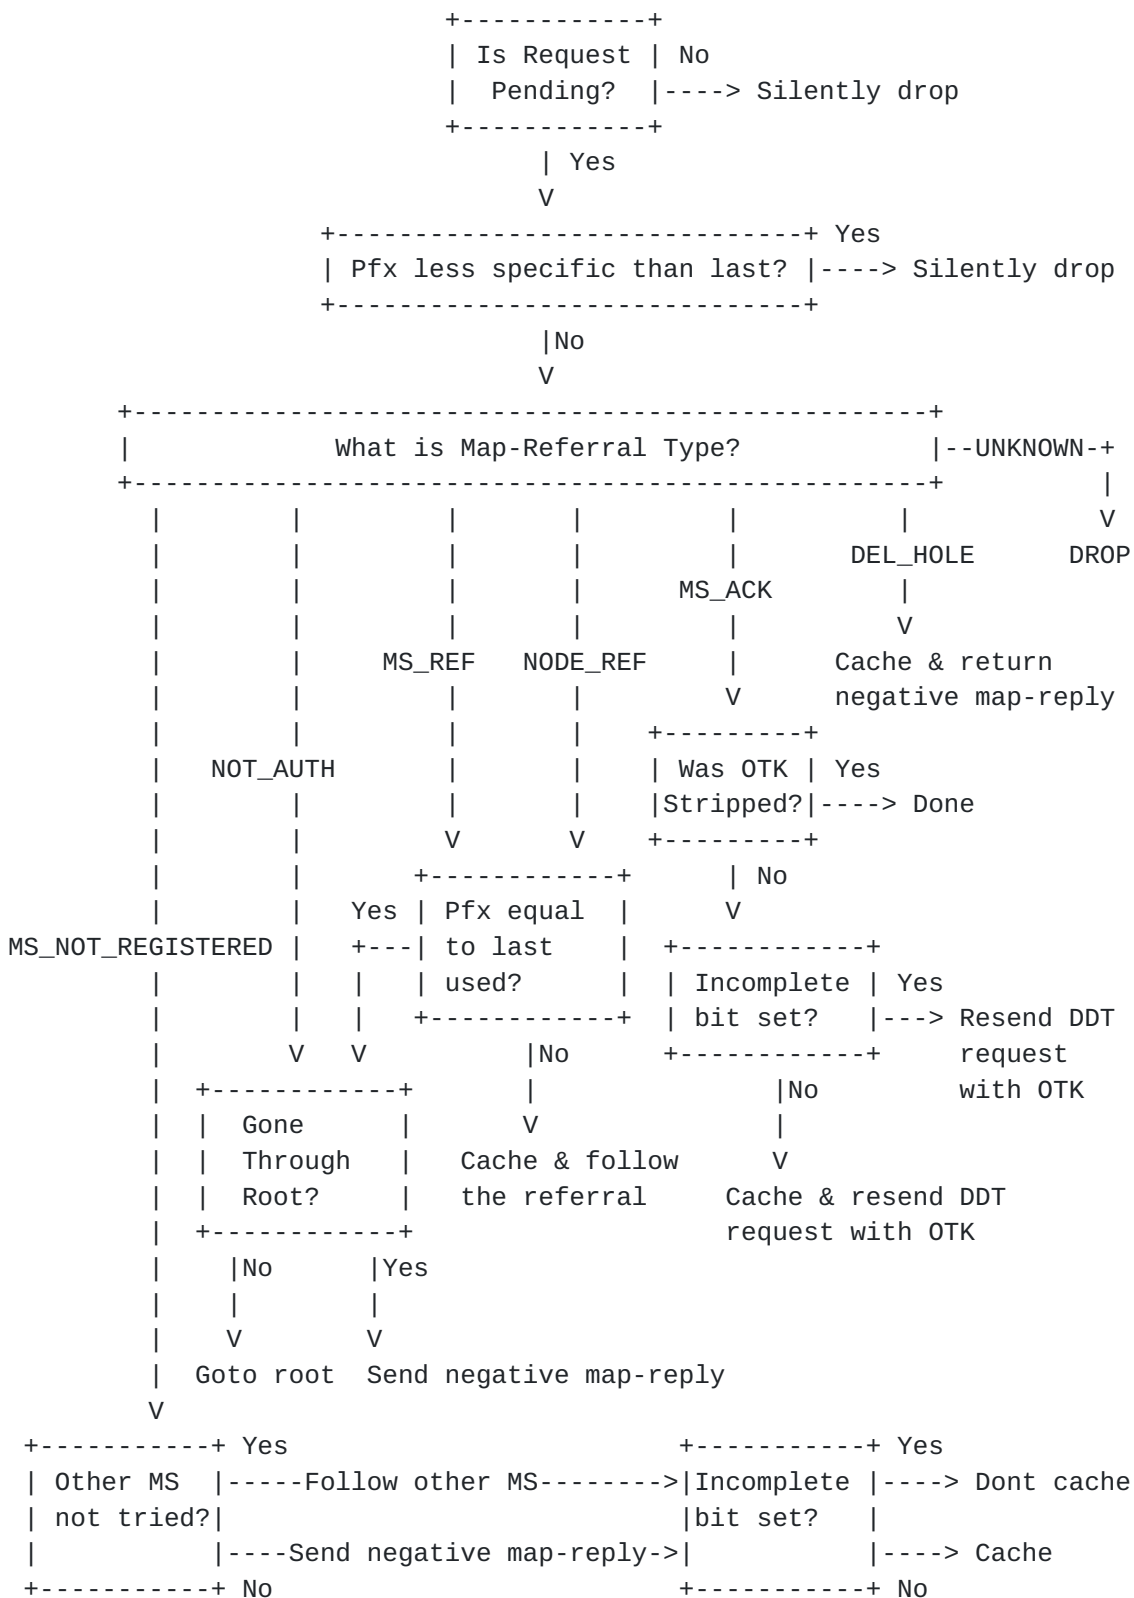
    case MS_ACKNOWLEDGEMENT:
        if ( OTK stripped ) {
            if ( incomplete ) {
                resend request with OTK
            } else {
                cache & resend request with OTK
            }
        }

    case MS_NOT_REGISTERED:
        if { all map-server delegations not tried } {
            follow delegations not tried
            if ( !incomplete ) {
                cache
            }
        } else {
            send negative map-reply to ITR
            if { !incomplete } {
                cache
            }
        }

    case DEFAULT:
        drop

}
}
```


6.2.2. Decision tree diagram



6.3. DDT Node processing of DDT Map-Request message

6.3.1. Pseudo-code summary

```
if ( I am not authoritative ) {
    send map-referral NOT_AUTHORITATIVE with
        incomplete bit set and ttl 0
} else if ( delegation exists ) {
    if ( delegated map-servers ) {
        send map-referral MS_REFERRAL with
            ttl 'Default_DdtNode_Ttl'
    } else {
        send map-referral NODE_REFERRAL with
            ttl 'Default_DdtNode_Ttl'
    }
} else {
    if ( eid in site ) {
        if ( site registered ) {
            forward map-request to etr
            if ( map-server peers configured ) {
                send map-referral MS_ACKNOWLEDGEMENT with
                    ttl 'Default_Registered_Ttl'
            } else {
                send map-referral MS_ACKNOWLEDGEMENT with
                    ttl 'Default_Registered_Ttl' and incomplete bit set
            }
        } else {
            if ( map-server peers configured ) {
                send map-referral MS_NOT_REGISTERED with
                    ttl 'Default_Configured_Not_Registered_Ttl'
            } else {
                send map-referral MS_NOT_REGISTERED with
                    ttl 'Default_Configured_Not_Registered_Ttl'
                    and incomplete bit set
            }
        }
    } else {
        send map-referral DELEGATION_HOLE with
            ttl 'Default_Negative_Referral_Ttl'
    }
}
```



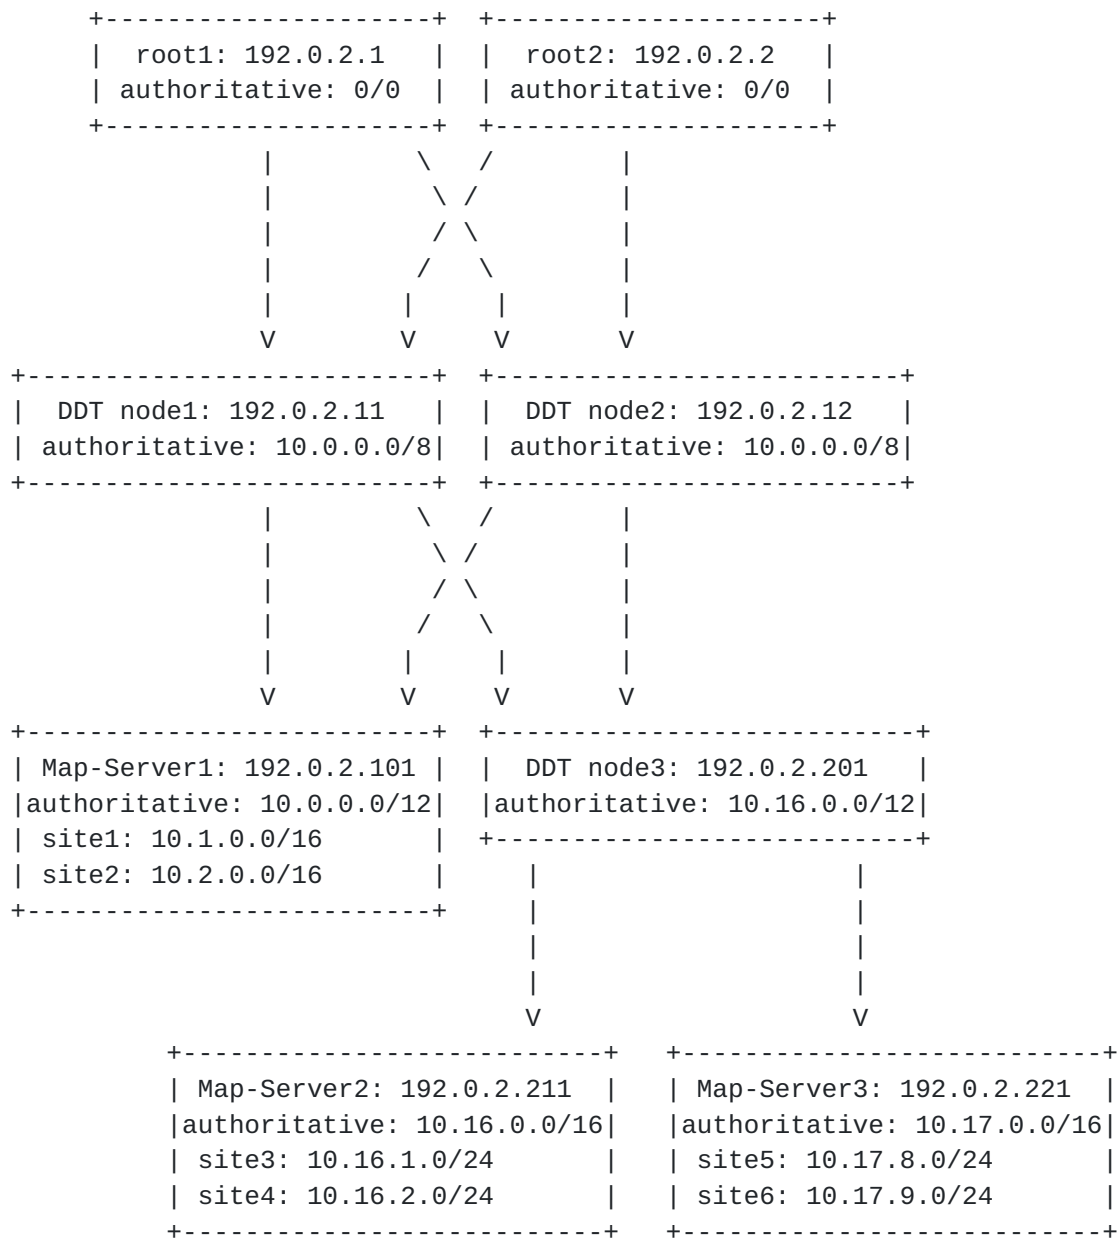
```

+-----+
|      Am I      | No
|  Authori-     | ----> Return NOT_AUTHORITATIVE
|  tative?      |      Incomplete = 1
+-----+      ttl = Default_DdtNode_Ttl
|
| Yes
|
V
+-----+      +-----+
| Delegation    | Yes | Delegations | Yes
|  Exists?      | ----> | are map     | ----> Return MS_REFERRAL
|               |      | servers?    |      ttl = Default_DdtNode_Ttl
+-----+      +-----+
|               |      \ No
| No            |      +--> Return NODE_REFERRAL
|               |      ttl = Default_DdtNode_Ttl
V
+-----+      +-----+      +-----+
| EID in        | Yes | Site         | Yes      | Map-server |
|  Site         | ----> | Registered? | ----> Forward----> | peers      |
| Config?       |      |               |      Map-request | configured? |
+-----+      +-----+      to ETR      +-----+
|               |      |               |      No      |      | Yes
|               |      | No            |      |      |      |
|               |      |               |      |      |      |
|               |      |               |      V      |      V
|               |      |               |      Return MS_ACK      Return MS_ACK
|               |      |               |      with INC=1
|               |      +-----+      ttl=Default_Registered_Ttl
|               |      | Map-server | Yes
|               |      | peers      | ----> Return MS_NOT_REGISTERED
|               |      | configured? |      ttl = Default_Negative_Referral_Ttl
+-----+      +-----+
|               |      \ No
| No            |      +--> Return MS_NOT_REGISTERED
|               |      Incomplete = 1
V               |      ttl = Default_Negative_Referral_Ttl
Return DELEGATION_HOLE
ttl = Default_Negative_Referral_Ttl

```


7. Example topology and request/referral following

To show how referrals are followed to find the RLOCs for a number of EIDs, consider the following example EID topology for DBID=0, IID=0, AFI=1, and EID=0/0



DDT nodes are configured for this "root" at IP addresses 192.0.2.1 and 192.0.2.2. DDT Map Resolvers are configured with default referral cache entries to these addresses.

The root DDT nodes delegate 10.0.0/8 to two DDT nodes with IP addresses 192.0.2.11 and 192.0.2.12.

The DDT nodes for 10.0.0.0/8 delegate 10.0.0.0/12 to a DDT Map Server with RLOC 192.0.2.101

The DDT Map Server for 10.0.0.0/12 is configured to allow ETRs to register the sub-prefixes 10.1.0.0/16 and 10.2.0.0/16

The DDT nodes for 10.0.0.0/8 also delegate 10.16.0.0/12 to a DDT node with RLOC 192.0.2.201

The DDT node for 10.16.0.0/12 is further configured to delegate 10.16.0.0/16 to a DDT Map Server with RLOC 192.0.2.211 and 10.17.0.0/16 to a DDT Map Server with RLOC 192.0.2.221

The DDT Map Server for 10.16.0.0/16 is configured to allow ETRs to register the sub-prefixes 10.16.1.0/24 and 10.16.2.0/24

The DDT Map Server for 10.17.0.0/16 is configured to allow ETRs to register the sub-prefixes 10.17.8.0/24 and 10.17.9.0/24

7.1. Lookup of 10.1.1.1/32 by ITR1

The first example shows a DDT Map Resolver following a delegation from the root to a DDT node followed by another delegation to a DDT Map Server.

ITR1 sends an Encapsulated Map-Request for 10.1.1.1 to one of its configured (DDT) Map Resolvers. The DDT Map Resolver proceeds as follows:

1. Send DDT Map-Request (for 10.1.1.1) to one of the root DDT nodes, 192.0.2.1 or 192.0.2.2
2. Receive (and save in referral cache) Map-Referral for EID-prefix 10.0.0.0/8, action code NODE-REFERRAL, RLOC set (192.0.2.11, 192.0.2.12)
3. Send DDT Map-Request to 192.0.2.11 or 192.0.2.12
4. Receive (and save in referral cache) Map-Referral for EID-prefix 10.0.0.0/12, action code MS-REFERRAL, RLOC set (192.0.2.101)

5. Send DDT Map-Request to 192.0.2.101; if the ITR-originated Encapsulated Map-Request had a LISP-SEC signature, it is included
6. DDT Map Server at 192.0.2.101 decapsulates the DDT Map-Request and forwards to a registered site1 ETR for 10.1.0.0/16
7. DDT Map Server at 192.0.2.101 sends a Map-Referral message for EID-prefix 10.1.0.0/16, action code MS-ACK to the DDT Map Resolver
8. DDT Map Resolver receives Map-Referral message and dequeues the pending request for 10.1.1.1
9. site1 ETR for 10.1.0.0/16 receives Map-Request forwarded by DDT Map Server and sends Map-Reply to ITR1

7.2. Lookup of 10.17.8.1/32 by ITR2

The next example shows a three-level delegation: root to first DDT node, first DDT node to second DDT node, second DDT node to DDT Map Server.

ITR2 sends an Encapsulated Map-Request for 10.17.8.1 to one of its configured (DDT) Map Resolvers, which are different from those for ITR1. The DDT Map Resolver proceeds as follows:

1. Send DDT Map-Request (for 10.17.8.1) to one of the root DDT nodes, 192.0.2.1 or 192.0.2.2
2. Receive (and save in referral cache) Map-Referral for EID-prefix 10.0.0.0/8, action code NODE-REFERRAL, RLOC set (192.0.2.11, 192.0.2.12)
3. Send DDT Map-Request to 192.0.2.11 or 192.0.2.12
4. Receive (and save in referral cache) Map-Referral for EID-prefix 10.16.0.0/12, action code NODE-REFERRAL, RLOC set (192.0.2.201)
5. Send DDT Map-Request to 192.0.2.201
6. Receive (and save in referral cache) Map-Referral for EID-prefix 10.17.0.0/16, action code MS-REFERRAL, RLOC set (192.0.2.221)
7. Send DDT Map-Request to 192.0.2.221; if the ITR-originated Encapsulated Map-Request had a LISP-SEC signature, it is included

8. DDT Map Server at 192.0.2.221 decapsulates the DDT Map-Request and forwards to a registered site5 ETR for 10.17.8.0/24
9. DDT Map Server at 192.0.2.221 sends a Map-Referral message for EID-prefix 10.17.8.0/24, action code MS-ACK, to the DDT Map Resolver
10. DDT Map Resolver receives Map-Referral(MS-ACK) message and dequeues the pending request for 10.17.8.1
11. site5 ETR for 10.17.8.0/24 receives Map-Request forwarded by DDT Map Server and sends Map-Reply to ITR2

7.3. Lookup of 10.2.2.2/32 by ITR1

This example shows how a DDT Map Resolver uses a saved referral cache entry to skip the referral process and go directly to a DDT Map Server for a prefix that is similar to one previously requested.

In this case, ITR1 uses the same Map Resolver used in example [Section 7.1](#). It sends an Encapsulated Map-Request for 10.2.2.2 to that (DDT) Map Resolver. The DDT Map-Resolver finds an MS-REFERRAL cache entry for 10.0.0.0/12 with RLOC set (192.0.2.101) and proceeds as follows:

1. Send DDT Map-Request (for 10.2.2.2) to 192.0.2.101; if the ITR-originated Encapsulated Map-Request had a LISP-SEC signature, it is included
2. DDT Map Server at 192.0.2.101 decapsulates the DDT Map-Request and forwards to a registered site2 ETR for 10.2.0.0/16
3. DDT Map Server at 192.0.2.101 sends a Map-Referral message for EID-prefix 10.2.0.0/16, action code MS-ACK to the DDT Map Resolver
4. DDT Map Resolver receives Map-Referral(MS-ACK) and dequeues the pending request for 10.2.2.2
5. site2 ETR for 10.2.0.0/16 receives Map-Request and sends Map-Reply to ITR1

7.4. Lookup of 10.16.2.1/32 by ITR2

This example shows how a DDT Map Resolver uses a saved referral cache entry to start the referral process at a non-root, intermediate DDT node for a prefix that is similar to one previously requested.

In this case, ITR2 asks the same Map Resolver used in example [Section 7.2](#). It sends an Encapsulated Map-Request for 10.16.2.1 to that (DDT) Map Resolver, which finds a NODE-REFERRAL cache entry for 10.16.0.0/12 with RLOC set (192.0.2.201). It proceeds as follows:

1. Send DDT Map-Request (for 10.16.2.1) to 192.0.2.201
2. Receive (and save in referral cache) Map-Referral for EID-prefix 10.16.0.0/16, action code MS-REFERRAL, RLOC set (192.0.2.211)
3. Send DDT Map-Request to 192.0.2.211; if the ITR-originated Encapsulated Map-Request had a LISP-SEC signature, it is included
4. DDT Map Server at 192.0.2.211 decapsulates the DDT Map-Request and forwards to a registered site4 ETR for 10.16.2.0/24
5. DDT Map Server at 192.0.2.211 sends a Map-Referral message for EID-prefix 10.16.2.0/24, action code MS-ACK to the DDT Map Resolver
6. DDT Map Resolver receives Map-Referral(MS-ACK) and dequeues the pending request for 10.16.2.1
7. site4 ETR for 10.16.2.0/24 receives Map-Request and sends Map-Reply to ITR2

7.5. Lookup of 10.16.0.1/32 (non-existent EID) by ITR2

This example uses the cached MS-REFERRAL for 10.16.0.0/16 learned above to start the lookup process at the DDT Map-Server at 192.0.2.211. The DDT Map Resolver proceeds as follows:

1. Send DDT Map-Request (for 10.16.0.1) to 192.0.2.211; if the ITR-originated Encapsulated Map-Request had a LISP-SEC signature, it is included
2. DDT Map Server at 192.0.2.211, which is authoritative for 10.16.0.0/16, does not have a matching delegation for 10.16.0.1. It responds with a Map-Referral message for 10.16.0.0/24, action code DELEGATION-HOLE to the DDT Map Resolver. The prefix 10.16.0.0/24 is used because it is the least-specific prefix that does match the requested EID but does not match one of configured delegations (10.16.1.0/24 and 10.16.2.0/24).
3. DDT Map Resolver receives the delegation, adds a negative referral cache entry for 10.16.0.0/24, dequeues the pending request for 10.16.0.1, and returns a negative Map-Reply to ITR2.

8. Securing the database and message exchanges

This section specifies the DDT security architecture that provides data origin authentication, data integrity protection, and XEID-prefix delegation. Global XEID-prefix authorization is out of the scope of this document.

Each DDT node is configured with one or more public/private key pair(s) that are used to digitally sign referral records for XEID-prefix(es) that the DDT node is authoritative for. In other words, each public/private key pair is associated with the combination of a DDT node and the XEID-prefix that it is authoritative for. Every DDT node is also configured with the public keys of its children DDT nodes. By including public keys of target child DDT nodes in the Map-Referral records, and signing each record with the DDT node's private key, a DDT node can securely delegate sub-prefixes of its authoritative XEID-prefixes to its children DDT nodes.

Map Resolvers are configured with one or more trusted public keys referred to as trust anchors. Trust anchors are used to authenticate the DDT security infrastructure. Map Resolvers can discover a DDT node's public key either by having it configured as a trust anchor, or by obtaining it from the node's parent as part of a signed Map-Referral. When a public key is obtained from a node's parent, it is considered trusted if it is signed by a trust anchor, or if it is signed by a key that was previously trusted. Typically, in a Map Resolver, the root DDT node public keys should be configured as trust anchors. Once a Map Resolver authenticates a public key it locally caches the key along with the associated DDT node RLOC and XEID-prefix for future use.

8.1. XEID-prefix Delegation

In order to delegate XEID sub-prefixes to its children, a parent DDT node signs its Map-Referrals. Every signed Map-Referral also includes the public keys associated with each child DDT node. Such a signature indicates that the parent node is delegating the specified XEID -prefix to a given child DDT node. The signature is also authenticating the public keys associated with the children nodes, and authorizing them to be used by the children DDT nodes to provide origin authentication and integrity protection for further delegations and mapping information of the XEID-prefix allocated to the DDT node.

As a result, for a given XEID-prefix, a Map Resolver can form an authentication chain from a configured trust anchor (typically the root DDT node) to the leaf nodes (Map Servers). Map Resolvers leverage this authentication chain to verify the Map-Referral

signatures while walking the DDT tree until they reach a Map Server authoritative for the given XEID-prefix.

8.2. DDT node operation

Upon receiving a Map-Request, the DDT node responds with a Map-Referral as specified in [Section 5](#). For every record present in the Map-Referral, the DDT node also includes the public keys associated with the record's XEID-prefix and the RLOCs of the children DDT nodes. Each record contained in the Map-Referral is signed using the DDT node's private key.

8.2.1. DDT public key revocation

The node that owns a public key can also revoke that public key. For instance if a parent node advertises a public key for one of its child DDT nodes, the child DDT node can at a later time revoke that key. Since DDT nodes do not keep track of the Map Resolvers that query them, revocation is done in a pull model, where the Map Resolver is informed of the revocation of a key only when it queries the node that owns that key. If the parent DDT is configured to advertise this key, the parent node must also be signaled to remove the key from the records it advertises for the child DDT node; this is necessary to avoid further distribution of the revoked key.

To securely revoke a key, the DDT node creates a new Record for the associated XEID-prefix and locator, including the revoked key with the R bit set. The DDT node must also include a signature in the Record that covers this record; this is computed using the private key corresponding to the key being revoked. Such a record is termed a "revocation record". By including this record in its Map-Referrals, the DDT node informs querying Map Resolvers about the revoked key. A digital signature computed with a revoked key can only be used to authenticate the revocation, and should not be used to validate any data. To prevent a compromised key from revoking other valid keys, a given key can only be used to sign a revocation for that specific key; it cannot be used to revoke other keys. This prevents the use of a compromised key to revoke other valid keys as described in [\[RFC5011\]](#). A revocation record must be advertised for a period of time equal to or greater than the TTL value of the Record that initially advertised the key, starting from the time that the advertisement of the key was stopped by removal from the parent DDT node.

8.3. Map Server operation

Similar to a DDT node, a Map Server is configured with one (or more) public/private key pairs that it must use to sign Map-Referrals.

However unlike DDT nodes, Map Servers do not delegate prefixes and as a result they do not need to include keys in the Map-Referrals they generate.

8.4. Map Resolver operation

Upon receiving a Map-Referral, the Map Resolver must first verify the signature(s) by using a trust anchor, or a previously authenticated public key, associated with the DDT node sending the Map-Referral. If multiple authenticated keys are associated with the DDT node sending this Map-Referral, the Key Tag field of the signature can be used to select the right public key for verifying the signature. If the key tag matches more than one key associated with that DDT node, the Map Resolver must try verifying the signature with all matching keys. For every matching key that is found the Map Resolver must also verify that the key is authoritative for the XEID-prefix in the Map-Referral record. If such a key is found, the Map Resolver must use it to verify the associated signature in the record. If no matching key is found, or if none of the matching keys is authoritative for the XEID-prefix in the Map-Referral record, or if such a key is found but the signature is not valid the Map-Referral record is considered corrupted and must be discarded. This may be due to expired keys. The Map Resolver can try other siblings of this node if there is an alternative node authoritative for the same prefix. If not, the Map Resolver can query the DDT node's parent to retrieve a valid key. It is good practice to use a counter or timer to avoid repeating this process if the resolver cannot verify the signature after several trials.

Once the signature is verified, the Map Resolver has verified the XEID-prefix delegation in the Map-Referral, and authenticated the public keys of the children DDT nodes. The Map Resolver must add these keys to the authenticated keys associated with each child DDT node and XEID-prefix. These keys are considered valid for the duration specified in the record's TTL field.

9. Open Issues and Considerations

There are a number of issues with the organization of the mapping database that need further investigation. Among these are:

- o Unlike in [[LISP-ALT](#)], DDT does not currently define a mechanism for propagating ETR-to-Map Server registration state. This requires DDT Map Servers to suppress returning negative Map-Reply messages for defined but unregistered XEID-prefixes to avoid loss of connectivity during partial ETR registration failures. Suppressing these messages may cause a delay for an ITR obtaining a mapping entry when such a failure is occurring.
- o Defining an interface to implement interconnection and/or interoperability with other mapping databases, such as LISP+ALT.
- o Additional key structures for use with LISP-DDT, such as to support additional EID formats as defined in [[LCAF](#)].
- o Authentication of delegations between DDT nodes.
- o Possibility of a new, more general format for the Map-Referral messages to facilitate the use of LISP-DDT with additional DBID/IID/EID combinations. Currently-defined packet formats should be considered to be preliminary and provisional until this issue has received greater attention.
- o Management of the DDT Map Resolver referral cache, in particular, detecting and removing outdated entries.

The authors expect that experimentation on the LISP pilot network will help answer open questions surrounding these and other issues.

10. IANA Considerations

This document makes no request of the IANA.

11. Security Considerations

[Section 8](#) describes a DDT security architecture that provides data origin authentication, data integrity protection, and XEID-prefix delegation within the DDT Infrastructure.

Global XEID-prefix authorization is beyond the scope of this document, but the SIDR working group [[RFC6480](#)] is developing an infrastructure to support improved security of Internet routing. Further work is required to determine if SIDR's public key infrastructure (PKI) and the distributed repository system it uses for storing and disseminating PKI data objects may also be used by DDT devices to verifiably assert that they are the legitimate holders of a set of XEID prefixes.

DDT security and [[LISP-SEC](#)] complement each other in securing the DDT infrastructure, Map-Referral messages and the Map-Request/Map-Reply protocol. In addition LISP-SEC can use the DDT public key infrastructure to secure the transport of LISP-SEC key material (the One-Time Key) from a Map-Resolver to the corresponding Map-Server. For this reason, when LISP-SEC is deployed in conjunction with a LISP-DDT mapping database and the path between Map-Resolver and Map-Server needs to be protected, DDT security should be enabled as well.

12. References

12.1. Normative References

- [LCAF] Farinacci, D., Meyer, D., and J. Snijders, "LISP Canonical Address Format", [draft-ietf-lisp-lcaf-00.txt](#) (work in progress), August 2012.
- [LISP] Farinacci, D., Fuller, V., Meyer, D., and D. Lewis, "Locator/ID Separation Protocol (LISP)", [draft-ietf-lisp-23.txt](#) (work in progress), May 2012.
- [LISP-MS] Fuller, V. and D. Farinacci, "LISP Map Server Interface", [draft-ietf-lisp-ms-16.txt](#) (work in progress), March 2012.
- [RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, [RFC 1035](#), November 1987.
- [RFC2104] Krawczyk, H., Bellare, M., and R. Canetti, "HMAC: Keyed-Hashing for Message Authentication", [RFC 2104](#), February 1997.
- [RFC4634] Eastlake, D. and T. Hansen, "US Secure Hash Algorithms (SHA and HMAC-SHA)", [RFC 4634](#), July 2006.
- [RFC5011] StJohns, M., "Automated Updates of DNS Security (DNSSEC) Trust Anchors", [RFC 5011](#), September 2007.

12.2. Informative References

- [LISP-ALT] Farinacci, D., Fuller, V., Meyer, D., and D. Lewis, "LISP Alternative Topology (LISP-ALT)", [draft-ietf-lisp-alt-10.txt](#) (work in progress), December 2011.
- [LISP-SEC] Maino, F., Ermagan, V., Cabellos, A., Sanchez, D., and O. Bonaventure, "LISP-Security", [draft-ietf-lisp-sec-03.txt](#) (work in progress), September 2012.
- [RFC1918] Rekhter, Y., Moskowitz, R., Karrenberg, D., Groot, G., and E. Lear, "Address Allocation for Private Internets", [BCP 5](#), [RFC 1918](#), February 1996.
- [RFC6480] Lepinski, M. and S. Kent, "An Infrastructure to Support Secure Internet Routing", [RFC 6480](#), February 2012.

[Appendix A](#). Acknowledgments

The authors wish to express their thanks to Damien Saucez, Lorand Jakab, and Olivier Bonaventure for work on LISP-TREE and LISP iterable mappings that inspired the hierarchical database structure and lookup iteration approach described in this document. Thanks also go to Dino Farinacci and Isidor Kouvelas for their implementation work; to Selina Heimlich and Srin Subramanian for testing; to Fabio Maino for work on security processing; and to Job Snijders, Glen Wiley, Neel Goyal, and Mike Gibbs for work on operational considerations and initial deployment of a prototype database infrastructure. Special thanks go to Jesper Skriver, Andrew Partan, and Noel Chiappa; all of whom have participated in (and put up with) seemingly endless hours of discussion of mapping database ideas, concepts, and issues.

Appendix B. Map-Referral Message Format

```

      0              1              2              3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|Type=6 |              Reserved              | Record Count |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|              Nonce . . .              |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|              . . . Nonce              |
+-> +--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| |              Record TTL              |
| +--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
R | Referral Count| EID mask-len | ACT |A|I|      Reserved      |
e +--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
c |SigCnt |  Map Version Number |              EID-AFI              |
o +--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
r |              EID-prefix ...              |
d +--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| /|  Priority  |  Weight  | M Priority  |  M Weight  |
| L +--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| o |      Unused Flags      |R|      Loc/LCAF-AFI      |
| c +--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| \|              Locator ...              |
+-> +--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

```

ACT: The "action" field of the mapping record in a Map-Referral message encodes 6 action types. The values for the action types are:

NODE-REFERRAL (0): Sent by a DDT node with a child delegation which is authoritative for the EID.

MS-REFERRAL (1): Sent by a DDT node that has information about Map Server(s) for the EID but it is not one of the Map Servers listed, i.e. the DDT-Node sending the referral is not a Map Server.

MS-ACK (2): Sent by a DDT Map Server that has one or more ETR registered for the EID.

MS-NOT-REGISTERED (3): Sent by a DDT Map Server that is configured for the EID-prefix but for which no ETRs are registered.

DELEGATION-HOLE (4): Sent by an intermediate DDT node with authoritative configuration covering the requested EID but without any child delegation for the EID. Also sent by a DDT Map Server with authoritative configuration covering the requested EID but for which no specific site ETR is configured.

NOT-AUTHORITATIVE (5): Sent by a DDT node that does not have authoritative configuration for the requested EID. The EID-prefix returned MUST be the original requested EID and the TTL MUST be set to 0. However, if such a DDT node has a child delegation covering the requested EID, it may choose to return NODE-REFERRAL or MS-REFERRAL as appropriate. A DDT Map Server with site information may choose to return of type MS-ACK or MS-NOT-REGISTERED as appropriate.

Incomplete: The "I" bit indicates that a DDT node's referral-set of locators is incomplete and the receiver of this message should not cache the referral. A DDT sets the "incomplete" flag, the TTL, and the Action Type field as follows:

Type (Action field)		Incomplete Referral-set		TTL values
0	NODE-REFERRAL	NO	YES	1440
1	MS-REFERRAL	NO	YES	1440
2	MS-ACK	*	*	1440
3	MS-NOT-REGISTERED	*	*	1
4	DELEGATION-HOLE	NO	NO	15
5	NOT-AUTHORITATIVE	YES	NO	0

*: The "Incomplete" flag setting on Map Server originated referral of MS-REFERRAL and MS-NOT-REGISTERED types depend on whether the Map Server has the full peer Map Server configuration for the same prefix and has encoded the information in the mapping record. Incomplete bit is not set when the Map Server has encoded the information, which means the referral-set includes all the RLOCs of all Map Servers that serve the prefix. It is set when the Map Server has not encoded the Map Server set information.

SigCnt: Indicates the number of signatures (sig section) present in the Record. If SigCnt is larger than 0, the signature information captured in a sig section as described in [Appendix B.1](#) will be appended to the end of the record. The number of sig sections at the end of the Record must match the SigCnt.

Loc/LCAF-AFI: If this is a Loc AFI, keys are not included in the record. If this is a LCAF AFI, the contents of the LCAF depend on

the Type field of the LCAF. Security material are stored in LCAF Type 11. DDT nodes and Map Servers can use this LCAF Type to include public keys associated with their Child DDT nodes for a XEID-prefix referral record. LCAF types and formats are defined in [LCAF].

All the field descriptions are equivalent to those in the Map-Reply message, as defined in [LISP]. Note, though, that the set of RLOCs correspond to the DDT node to be queried as a result of the referral not the RLOCs for an actual EID-to-RLOC mapping.

B.1. SIG section

If SigCnt field in the Map-Referral is not 0, the signature information is included at the end of captured in a sig section as described below. SigCnt counts the number of sig sections that appear at the end of the Record.

```

+-----+-----+-----+-----+-----+-----+-----+-----+-----+
/|                                     Original Record TTL                                     |
/ +-----+-----+-----+-----+-----+-----+-----+-----+-----+
/ |                                     Signature Expiration                                     |
| +-----+-----+-----+-----+-----+-----+-----+-----+-----+
s |                                     Signature Inception                                     |
i +-----+-----+-----+-----+-----+-----+-----+-----+-----+
g |                                     Key Tag                                     |                                     Sig Length                                     |
| +-----+-----+-----+-----+-----+-----+-----+-----+-----+
\ | Sig-Algorithm |   Reserved   |                                     Reserved                                     |
\ +-----+-----+-----+-----+-----+-----+-----+-----+-----+
\ ~                                     Signature                                     ~
+-----+-----+-----+-----+-----+-----+-----+-----+-----+

```

Original Record TTL: The original Record TTL for this record that is covered by the signature. Record TTL is in minutes.

Key Tag: An identifier to specify which key is used for this signature if more than one valid key exists for the signing DDT node.

Sig Length: The length of the Signature field.

Sig-Algorithm: The identifier of the cryptographic algorithm used for the signature. Default value is RSA-SHA1.

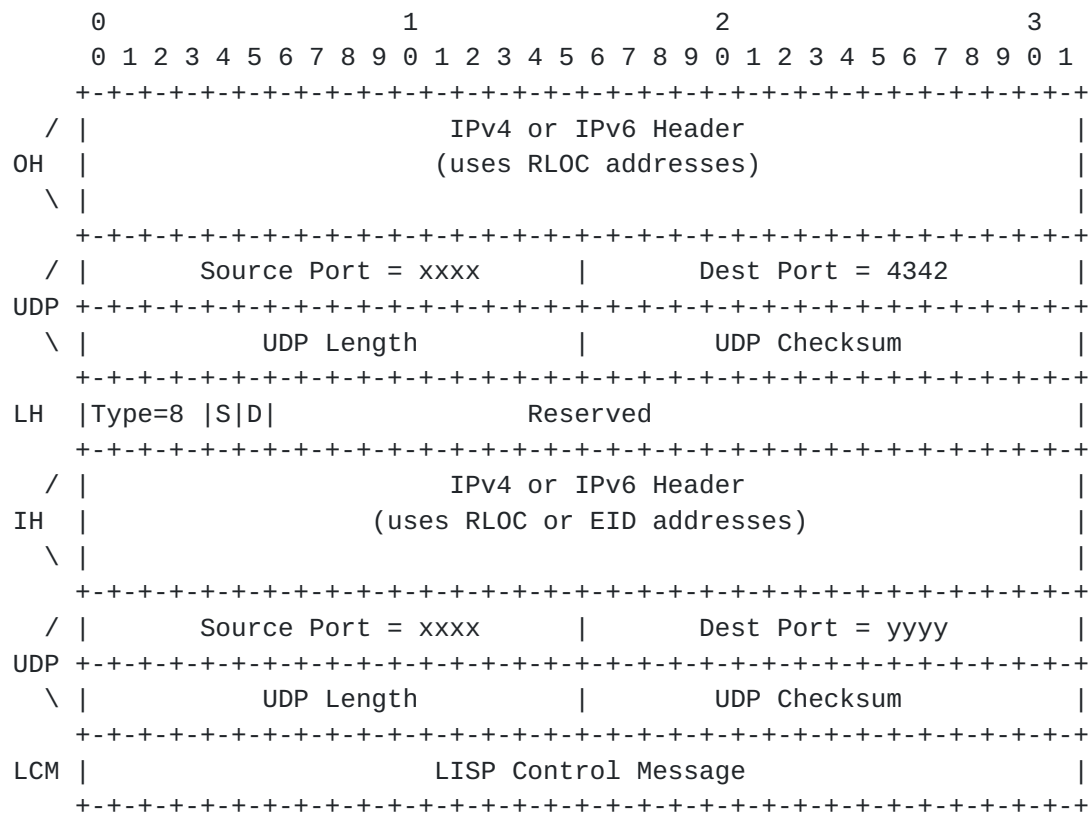
Reserved: This field must be set to 0 on transmit and must be ignored on receipt.

Signature Expiration and Inception: Specify the validity period for the signature. Each specify a date and time in the form of a 32-bit

unsigned number of seconds elapsed since 1 January 1970 00:00:00 UTC, ignoring leap seconds, in network byte order.

Signature: Contains the cryptographic signature that covers the entire record. The Record TTL and the sig fields are set to zero for the purpose of computing the Signature

Appendix C. Encapsulated Control Message Format



"D" is the "DDT-originated" flag and is set by a DDT client to indicate that the receiver can and should return Map-Referral messages as appropriate.

Authors' Addresses

Vince Fuller
cisco Systems
Tasman Drive
San Jose, CA 95134
USA

Email: vaf@cisco.com

Darrel Lewis
cisco Systems
Tasman Drive
San Jose, CA 95134
USA

Email: darlewis@cisco.com

Vina Ermagan
cisco Systems
Tasman Drive
San Jose, CA 95134
USA

Email: vermagan@cisco.com

Amit Jain
cisco Systems
Tasman Drive
San Jose, CA 95134
USA

Email: amijain@cisco.com

