LISP Working Group
Internet-Draft
Intended status: Informational
Expires: January 10, 2013

N. Chiappa
Yorktown Museum of Asian Art

July 9, 2012

An Architectural Analysis of the LISP Location-Identity Separation System
draft-chiappa-lisp-architecture-00.txt

Abstract

LISP upgrades the architecture of the IPvN internetworking system by separating location and identity, current intermingled in IPvN addresses. This is a change which has been identified by the IRTF as a critically necessary evolutionary architectural step for the Internet. In LISP, nodes have both a 'locator' (a name which says _where_ in the network's connectivity structure the node is) and an 'identifier' (a name which serves only to provide a persistent handle for the node). A node may have more than one locator, or its locator may change over time (e.g. if the node is mobile), but it keeps the same identifier.

One of the chief novelties of LISP, compared to other proposals for the separation of location and identity, is its approach to deploying this upgrade. In general, it is comparatively easy to conceive of new network designs, but much harder to devise approaches which will actually get deployed throughout the global network. LISP aims to achieve the near-ubiquitous deployment necessary for maximum exploitation of an architectural upgrade by i) minimizing the amount of change needed (existing hosts and routers can operate unmodified); and ii) by providing significant benefits to early adopters.

This document gives additional architectural insight into LISP, and analyzes a number of aspects of LISP from a long-term perspective.

NOTE: This is an initial rough draft, a much better version will be out shortly.

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 January 10, 2013.

Table of Contents

1. Introduction

This document begins by introducing some high-level architectural frameworks which have proven useful for thinking about the LISP location-identity separation system. It then discusses some architectural aspects of LISP (e.g. its namespaces). The balance (and bulk) of the document contains architectural analysis of the LISP system; that is, it reviews from a long-term perspective various aspects of that system; e.g. its scalability, security, robustness, etc.

NOTE: This document assumes a fair degree of familiarity with LISP; in particular, the reader should have a good 'high-level' understanding of the overall LISP system architecture, such as is provided by [Introduction], "An Introduction to the LISP System".

By "architecture" above, the restricted meaning used there is: 'How the system is broken up into subsystems, and how those subsystems interact; when does information flows from one to another, and what that information is.' There is obviously somewhat more to architecture (e.g. the namespaces of a system, their syntax and semantics), and that remaining architectural content is covered here.

2. Architectual Frameworks

When considering the overall structure of the LISP system at a high level, it has proven most useful to think of it as another packet-switching layer, run on top of the original internet layer - much as the Internet first ran on top of the ARPANET.

All the functions that a normal packet switch has to undertake - such as ensuring that it can reach its neighbours, and they they are still up - the devices that make up the LISP overlay also have to do, along the 'tunnels' which connect them to other LISP devices.

There is, however, one big difference: the fanout of a typical LISP ITR will be much larger than most classic physical packet switches. (ITRs only need to be considered, as the LISP tunnels are all effectively unidirectional, from ITR to ETR - the ETR needs to keep

no per-tunnel state, etc.)

LISP is, fundamentally, a 'tunnel' based system. Tunnel system designs do have their issues (e.g. the high inter-'switch' fan-out), but it's important to realize that they also can have advantages, some of which are listed below.

2.1. 'Double-Ended' Approach

LISP may be thought of as a 'double-ended' approach to enhancing the architecture, in that it uses pairs of devices, one at each end of a communication stream. In particular, to interact with the population of 'legacy' hosts (which will be, inevitably, the vast majority, in the early stages of deployment) it requires a LISP device at both ends of the 'tunnel'.

This is in distinction to, say, NAT systems, which only need a device deployed at one end: the host at the 'other end' doesn't need a matching device at its end to massage the packets, but can simply consume them on its own, as they are fully normal packets. This allows any site which deploys such a device to get the full benefit whilst acting entirely on its own. [Wasserman]

The issue is not that LISP uses tunnels. Designs like HIP ([RFC4423]) and ILNP ([ILNP]), which do not involve tunnels, inhabit a similar space to tunnel-based designs like LISP, in that unless both ends are upgraded - or there is a proxy at the un-upgraded end - one doen't get any benefits. So it's really not the tunnel which is the key aspect, it's the 'all at one end' part which is key. Whether the system is tunnel, versus non-tunnel, is not that important.

However, the double-ended approach of LISP does have advantages, as well as costs. To put it simply, the 'feature' of the alternative approach, that there's only a box at one end, has a 'bug': there's only a box at one end. There are things which such a design cannot accomplish, because of that. To put it another way, does the fact that the packet has only a single name in it, because it is a 'normal' packet, present a limitation? Put that way, it would seem natural that it should.

To compile as complete list of such situations is beyond the scope of this document, but one example is mobility with open connections.

It is also possible to use LISP to tunnel IPv6 traffic over IPv4 infrastructure, or vice versa, invisibly to the hosts on both ends.

In the longer term, having having tunnel boxes would allow us to wrap packets in non-IP formats: perhaps to take direct advantage of the capabilities of underlying switching fabrics (e.g. MPLS); perhaps to deploy new carriage protocols, etc, where non-standard packet formats will allow extended semantics.

2.2. Critical State

LISP does have 'critical state' in the network (i.e. state which, if if lost, causes the communication to fail). However, because LISP is designed as an overall system, 'designing it in' allows for a 'systems' approach to its state issues. In LISP, this state has been designed to be maintained in an 'architected' way, so it does not produce systemic brittleness in the way that the state in NATs does.

For instance, throughout the system, provisions have been made to have redundant copies of state, in multiple devices, so that the loss of any one device does not necessarily cause a failure of an ongoing connection.

2.3. Need for a Mapping System

LISP does need to have a mapping system, which brings design, implementation, configuration and operational costs. Surely all these costs are a bad thing? However, having a mapping system have advantages, especially when there is a mapping layer which has global visibility (i.e. other entities know that it is there, and have an interface designed to be able to interact with it). This is unlike, say, the mappings in NAT, which are 'invisible' to the rest of the network.

In fact, one could argue that the mapping layer is LISP's greatest strength. Wheeler's Axiom* ('Any problem in computer science can be solved with another level of indirection') indicates that the binding layer available with the LISP mapping system will be of great value. Again, it is not the job of this document to list them all - and in any event, there is no way to forsee them all.

The author of this document has often opined that the hallmark of great architecture is not how well it does the things it was designed to do, but how well it does things it was never expected to have to handle. Providing such a powerful and generic binding layer is one sure way to achieve the sort of lasting flexibility and power that leads to that outcome.

[Footnote *: This Axiom is often mis-attributed to Butler Lampson, but Lampson himself indicated that it came from David Wheeler.]

2.4. Piggybacking of Control on User Data

LISP piggybacks control transactions on top of user data packets. This is a technique that has a long history in data networking, going back to the early ARPANET. [McQuillan] It is now apparently regarded as a somewhat dubious technique, the feeling seemingly being that control and user data should be strictly segregated.

It should be noted that _none_ of the piggybacking of control functionality in LISP is _architecturally fundamental_ to LISP. All

of the functions in LISP which are performed with piggybacking could be performed almost equally well with separate control packets.

The "almost" is solely because it would cause more overhead (i.e. control packets); neither the response time, robustness, etc would necessarily be affected - although for some functions, to match the response time observed using piggybacking on user data would need as much control traffic as user data traffic.

This technique is particularly important, however, because of the issue identified at the start of this section - the very large fanout of the typical LISP switch. Unlike a typical router, which will have control interactions with only a few neighbours, a LISP switch could eventually have control interactions with hundreds, or perhaps even thousands (for a large site) of neighbours.

Explicit control traffic, especially if good response times are desired, could amount to a great deal of overhead in such a case.

3. Namespaces

One of the key elements in any architecture, or architectural analysis, are the namespaces involved: what are their semantics and syntax, what are the kinds of things they name, etc.

LISP has two key namespace, EIDs and RLOCs, but it must be emphasized that on an architectural level, neither the syntax, or, to a lesser degree, the semantics, of either are absolutely fixed. There are certain core semantics which are unchanging (such as the notion that EIDs provide only identity, whereas RLOCs provide location), but as we will see, there is a certain amount of flexibility available for the long-term.

In particular, all of LISP's key interfaces always include an Address Family Identifier (AFI) [AFI], so that new forms can be introduced at any time the need is felt. Of course, in practise such an introduction would not be a trivial exercise - but neither is is impossibly painful, as is the case with IPv4's 32-bit addresses, which are effectively impossible to upgrade.

3.1. LISP EIDs

A 'classic' EID is defined as a subset of the possible namespaces for endpoints. [Chiappa] Like most 'proper' endpoint names, as proposed there, they contain contain no information about the location of the endpoint. EIDs are the subset of possible endpoint names which are: fixed length, 'reasonably' short', binary (i.e. not intended for direct human use), globally unique (in theory), and allocated in a top-down fashion (to achieve the former) .

LISP EIDs are, in line with the general LISP deployment philosophy, a reuse of something already existing - i.e. IPvN addresses. For

those used as in LISP as EIDs, LISP removes much (or, in some cases, all) of the location-naming function of IPvN addresses.

In addition, the goal is to have EIDs name hosts (or, more properly, their end-end communication stacks), whereas the other LISP namespace group (RLOCs) names interfaces. The idea is not just to have two namespaces (with different semantics), but also to use them to name _different classes of things_ - classes which currently do not have clearly differentiated names. This should produce even more functionality.

3.1.1. Residual Location Functionality in EIDs

LISP retains, especially in the early stages of the deployment, in many cases some residual location-naming functionality in EIDs, This is to allow the packet to be correctly routed/forwarded to the destination node, once it has been unwrapped by the ETR - and this is a direct result of LISP's deployment philosophy (see [Introduction], Section "Deployment").

Clearly, if there are one or more unmodified routers between the ETR and the desination node, those routers will have to perform a routing step on the packet, for which it will need _some_ information as to the location of the destination.

One can thus view such LISP EIDs, which retain 'stub' location information, as 'addresses' (in the definition of the generic sense of this term, as used here), but with the location information restricted to a limited, local scope.

This retention of some location functionality in LISP EIDs, in some cases, has led some people to argue that use of the name 'EID' is improper. In response, it was suggested that LISP use the term 'LEID', to distinguish LISP's 'bastardized' EIDs from 'true' EIDs, but this usage has never caught on.

It has also been suggested that one usage mode for LISP EIDs, in existing software loads, is to assign them as the address on an internal virtual interface; all the real interfaces would have RLOCs only. [Templin] This would make such LISP EIDs functionally equivalent to 'real' EIDs - they are names which are purely identity, have no location information of any kind in them, and cannot be used to make any routing decisions anywhere outside the host.

It is true that even in such cases, the EID is still not a 'pure' EID, as it names an interface, not the end-end stack directly. However, to do a perfect job here (or on separation of location and identity) is impossible without modifying existing hosts (which are, inevitably, almost always one end of an end-end communication) - and that has been ruled out, for reasons of viable deployment.

The need for interoperation with existing unmodified hosts limits the semantic changes one can impose, much as one might like to provide a cleaner separation. (Future evolution can bring us toward that state, however: see Section XXX.)

3.2. RLOCs

RLOCs are basically pure 'locators' [RFC1992], although their syntax and semantics is restricted at the moment, because in practise the only forms of RLOCs supported are IPv4 and IPv6.

3.3. Overlapping Uses of Existing Namespaces

3.4. LCAFs

--- Key-ID
--- Instance-IDs

4. Fault Discovery/Handling

Any global communication system must be robust, and to be robust, it must be able to discover and handle problems. LISP's general philosophy of robustness is usually to have overlapping, simple mechanisms to discover and repair problems.

5. Scalability

As with robustness, any global communication system must be scalable, and scalable up to almost any size. As previously mentioned, the large fanouts to be seen with LISP, due to its 'overlay' nature, present a special challenge.

One likely saving grace is that as the Internet grows, most sites will likely only interact with a limited subset of the Internet; if nothing else, the separation of the world into language blocks means that content in, e.g. Chinese, will not be of interest to most of the rest of the world. This tendency will help with a lot of things which could be problematic if constant, full, N^2 connectivity were likely on all nodes, for example the caching of mappings.

5.1. Demand Loading of Mappings

One question that many will have about LISP's design is 'why demand-load mappings - why not just load them all'? It is certainly true that with the growth of memory sizes, the size of the complete database is such that one could reasonably propose keeping the entire thing in each LISP device. (In fact, one proposed mapping system for LISP, named NERD, did just that. [NERD])

A 'pull'-based system was chosen over 'push' for several reasons; the main one being that the issue is not just the pure _size_ of the

mapping database, but its _dynamicity_. Depending on how often mappings change, the update rate of a complete database could be relatively large.

It is especially important to realize that, depending on what (probably unforseeable) uses eventually evolve for the identity->location mapping capability, the update rate could be very high indeed. E.g. if LISP is used for mobility, that will greatly increase the update rate. Such a powerful and flexible tool is likely be used in unforseen ways (Section 2.3), so it's unwise to make a choice that would preclude any which raise the update rate significantly.

Push as a mechanism is also fundamentally less desirable than pull, since the control plane overhead consumed to load and maintain information about unused destinations is entirely wasted. The only potential downside is the delay required for the demand-loading of information.

(It's also probably worth noting that many issues that some people have with the mapping approach of LISP, such as the mapping database size, etc are the same - if not worse - for push as they are for pull.)

Also, for IPv4, as the address space becomes more highly used, it will become more fragmented - i.e. there will tend to be more, smaller, entries. For a routing table, which every router has to hold, this is problematic. For a demand-loaded mapping table, it is not bad. Indeed, this was the original motivation for LISP - although many other useful and desirable uses for it have since been enumerated (see [Introduction], Section "Applications").

For all of these reasons, as long as there is locality of reference (i.e. most ITRs will use only a subset of the entire set), it makes much more sense to use the a pull model, than the classic push one heretofore seen widely at the internetwork layer (and thus somewhat novel to people who work at that layer).

It may well be that some sites (e.g. large content providers) may need non-standard mechanisms - perhaps something more of a 'push' model. This remains to be determined, but it is certainly feasible.

5.2. Caching of Mappings

It should be noted that the caching spoken of here is likely not classic caching, where there is a fixed/limited size cache, and entries have to be discarded to make room for newly needed entries. The economics of memory being what they are, there is no reason to discard mappings once they have been loaded (although of course implementations are free to chose to do so, if they wish to).

This leads to another point about the caching of mappings: the algorithms for management of the cache are purely a local issue. The algorithm in any particular ITR can be changed at will, with no need for any coordination. A change might be for purposes of experimentation, or for upgrade, or even because of environmental variations - different environments might call for different cache management strategies.

The replacability of the cache management is the architectural aspect of the design; the exact algorithm, which is engineering, is not.

5.3. Amount of State

-- Mapping cache size
--- Mention studies
-- Delegation cache size (in MRs)
--- Mention studies
-- Any others?

5.4. Scalability of The Indexing Subsystem

LISP initially used an indexing subsystem called ALT. [ALT] ALT was relatively easy to construct from existing tools (GRE, BGP, etc), but it had a number of issues that made it unsuitable for large-scale use. ALT is now being superseded by DDT. [DDT]

The basic structure and operation of DDT is identical to that of TREE, so the extensive simulation work done for TREE applies equally to DDT, as do the conclusions drawn about TREE's superiority to ALT. [Jakab]

From an architectural point of view, the main advantage of DDT is that it enables client side caching of information about intermediate nodes in the resolution hierarchy, and also enables direct communication with them. As a result, DDT has much better scaling properties than ALT.

The most important result of this change is that it avoids a concentration of resolution request traffic at the root of the indexing tree, a problem which by itself made ALT unsuitable for a global-scale system. The problem of root concentration (and thus overload) is almost unavoidable in ALT (even if masses of 'bypass' links are created).

ALT's scalability also depends on enforcing an intelligent organization that aincreases aggregation. Unfortunately, the current backbone routing BGP system shows that there is a risk of an organic growth of ALT, one which does not achieve aggregation. DDT does not display this weakness, since its organization is inherently hierarchical (and thus inherently aggregable).

The hierarchical organization of DDT also reduces the possibility for

a configuration error which interferes with the operation of the network (unlike the situation with the current BGP DFZ). DDT security mechanisms can also help produce a high degree of robustness, both against misconfiguration, and deliberate attack. The direct communication with intermediate nodes in DDT also helps to quickly locate problems when they occur, resulting in better operational characteristics.

Next, since in ALT mapping requests must be transmitted through an overlay network, a significant share of requests can see substantially increased latencies. Simulation results in the TREE work clearly showed, and quantified, this effect.

The simulations also showed that the nodes composing the ALT and DDT networks for a mapping database of full Internet size could have thousands of neighbours. This is not an issue for DDT, but would almost certainly have been problematic for ALT nodes, since handling that number of simultaneous BGP sessions would likely to be difficult.

6. Security

Security in LISP faces many of the same challenges as security for other parts of the Internet: good security usually means work for the users, but without good security, things are vulnerable.

The Internet has seen many very secure systems devised, only to see them fail to reach wide adoption; the reasons for that are complex, and vary, but being too much work to use is a common thread. It is for this reason that LISP attempts to provide 'just enough' security (see [Introduction], Section "Design-Security").

The _good_ thing about the Internet is that it brings the world to your doorstep - masses of information from all around the world are instantly available on your computing device. The _bad_ thing about the Internet is that it brings the world to your doorstep - including legions of crackers, thieves, and general scum and villainy. Thus, any node may be the target of fairly sophisticated attack - often automated (thereby reducing the effort required of the attacker to spread their attack as broadly as possible).

6.1. Basic Philosophy

To square this circle, of needing to have very good security, but of it being to difficult to use very good security, the general concept is for LISP to have a series of 'graded' security measures available, with the 'ultimate' security mechanisms being very high-grade indeed.

The concept is to devise a plan in which LISP can simultaneously attempt to have not just 'ultimate' security, but also one or more 'easier' modes, ones which will be easier to configure and use. This

'easier' mode can be both an interim system (with the full powered system available for when it it needed), as well as the system used in sections where security is less critical (following the general rule that the level of any security should generally be matched to what is being protected).

The challenge is to do this in a way that does not make the design more complex, since it has to include both the 'full strength' mechanism(s), and the 'easier to configure' mechanism(s). This is one of the fundamental tradeoffs to struggle with: it is easy to provide 'easier to configure' options, but that may make the overall design more complex.

As far as making it hard to implement to begin with (also something of a concern initially, although obviously not for the long term): we can make it 'easy' to deploy initially by simply not implementing/ configuring the heavy-duty security early on. (Provided, of course, that the packet formats, etc, are all included in the design to begin with.)

6.2. Design Guidance

In designed the security, there are a small number of key points that will guide the design:

- Design lifetime
- Threat level

How long is the design intended to last? If LISP is successful, a minimum of a 50-year lifetime is quite possible. (For comparison, IPv4 is now 34 at the time of writing this, and will be around for at least several decades yet, if not longer; DNS is 28, and will probably last indefinitely.)

How serious are the threats it needs to meet? As mentioned above, the Internet can bring the baddest actors anywhere to any location, in a flash. Their sophistication level is rising all the time: as the easier holes are plugged, they go after others. This will inevitably eventually require the most powerful security mechanisms available to counteract their attacks.

Which is not to say that LISP needs to be that secure _right away_. The threat will develop and grow over a long time period. However, the basic design has to be capable of being _securable_ to the expanded degree that will eventually be necessary. However, _eventually_ it will need to be as securable as, say, DNS - i.e. it _can_ be secured to the same level, although people may chose not to secure their LISP infrastructure as well as DNSSEC does.

In particular, it should be noted that historically many systems have been broken into, not through a weakness in the algorithms, etc, but

because of poor operational mechanics. (The well-known 'Ultra' breakins of the Allies were mostly due to failures in operational procedure.) So operational capabilities intended to reduce the chance of human operational failure are just as important as strong algorithms; making things operationally robust is a key part of 'real' security.

6.2.1. Security Mechanism Complexity

Complexity is bad for several reasons, and should always be reduced to a minimum. There are three kinds of complexity cost: protocol complexity, implementation complexity, and configuration complexity. We can further subdivide protocol complexity into packet format complexity, and algorithm complexity. (There is some overlap of algorithm complexity, and implementation complexity.)

We can, within some limits, trade off one kind of complexity for others: e.g. we can provide configuration _options_ which are simpler for the users to operate, at the cost of making the protocol and implementation complexity greater. And we can make initial (less capable) implementations simpler if we make the protocols slightly more complex (so that early implementations don't have to implement all the features of the full-blown protocol).

It's more of a question of some operational convenience/etc issues - e.g. 'How easy will it be to recover from a cryptosystem compromise'. If we have two ways to recover from a security compromise, one which is mostly manual and a lot of work, and another which is more automated but makes the protocol more complicated, if compromises really are very rare, maybe the smart call _is_ to go with the manual thing - as long as we have looked carefully at both options, and understood in some detail the costs and benefits of each.

6.3. Security Overview

First, there are two different classes of attack to be considered: denial of service (DoS, i.e. the ability of an intruder to simply cause traffic not to successfully flow) versus exploitation (i.e. the ability to cause traffic to be 'highjacked', i.e. traffic to be sent to the wrong location).

Second, one needs to look at all the places that may be attacked. Again, LISP is a relatively simple system, so there are not that many subsystems to examine.

-- Lookups
--- Nonces
-- Indexing
-- Mappings

6.4. Securing Mappings

Two approaches have been taking to securing the provision of mappings. The first, which is of course not completely satisfactory, is to secure the channel between the ITR and the entities involved in providing mappings for it. The second is to secure the mappings themselves, by signing them 'at birth' (much the same way in which DNS Security operates). [RFC4033].

Tie-in to space allocation security?

6.5. Securing the xTRs

--- Cache management
--- Unsoliticed Map-Replies are _very bad_ - must go through
    mapping system to verify that the sender is authoritative for
    that range of EIDs

7. Robustness

-- Depends on deployment as well as design
-- Replication
-- Overlapping mechanisms (ref redundancy as key for robustness)

8. Optimization

-- Philosophy
-- Piggybacking
-- 'Wiretapping' return mappings
--- Security is an issue on that

9. Open Issues

-- Provider lock-in for mapping database
-- Automated ETR synch
--- Liveness can be gathered now from some IGPs
-- EID reachability within a LISP Site
--- Existing problem with any border router

10. Additional Material

- An architectural document on LISP, starting with
    a brief motivation and problem definition, then
    a protocol overview, then more detailed discussion
    of funtional elements, tradeoffs, etc.

- A more high level document covering what Loc/ID
  separation is, why it's important, its history,
  and then why LISP is a good solution.

- A 'potential future evolution' document, covering
  what the impacts of LISP could/would be, and how it
  might evolve in the future.

- Future work
-- Better ETR sync for mappings
-- Detect and deal with gonzo ETRs

- Long-term Advantages
-- Impossible to list all the long-term uses
--- Lampson's (sic - actually Wheeler's) Law
-- EID Address space utilization
-- Allows use of class E space (240/4) for RLOCs
-- Core routing overhead reduction
--- PI space
-- Easier introduction of new names-spaces

- Routing Evolution {not sure we want this? - JNC}
-- Some short term (TE, etc)
-- Biggest long-term improvement comes inside LISP core,
    if we can get the network to that point
--- Withdrawal of EID routes in the LISP core
-- Support of non-LISP core is tricky (requires
    route reconstitution)

- Long-term Evolution {Not sure we want this? Maybe it
    should be a whole separate document. - JNC}
-- Evolution
--- Have a long term plan, but keep what's actually
    done simple to begin with
--- Leave 'hooks' for long-term evolution
--- The Internet painted itself into a corner,
    evolution-wise - LISP has opened a mouse-hole, but
    we need to make sure it doesn't just lead to another
    painted-in corner
-- Better indexing system (may be obsolete, now that
    we have DDT?)
-- 'Ringfence' of xTRs provides natural boundary between
    change domains
-- New namespaces (the semantic issues involved with
    introducing new namespaces - initially RLOCs, but
    potentially EIDs as well - the latter are
    much harder as it changes host-host semantics)
-- Separation of host/host and router/router
    interfaces / packet formats
-- More?

11. Acknowledgments

The author would like thank the core LISP group for their willingness

to allow him to add himself to their effort, and for their enthusiasm for whatever assistance he has been able to provide. He would also like to thank (in alphabetical order) Vina Ermagan, Vince Fuller, and Joel Halpern for their careful review of, and helpful suggestions for, this document. Grateful thanks also to Darrel Lewis for his help with material on non-Internet uses of LISP, and to Vince Fuller for help with XML. A final thanks is due to John Wrocklawski for the author's organizational affiliation.

12. IANA Considerations

This document makes no request of the IANA.

13. Security Considerations

14. References

14.1. Normative References

[DDT] V. Fuller, D. Lewis, and D. Farinacci, "LISP Delegated Database Tree", draft-fuller-lisp-ddt-01.txt (work in progress), March 2012.

[Introduction] J.N. Chiappa, "An Introduction to the LISP Location-Identity Separation System", draft-chiappa-lisp-introduction-00.txt (work in progress), July 2012.

[AFI] IANA, "Address Family Indicators (AFIs)", Address Family Numbers, January 2011, <http://www.iana.org/assignments/address-family-numbers>.

14.2. Informative References

[RFC1992] I. Castineyra, J. N. Chiappa, and M. Steenstrup, "The Nimrod Routing Architecture", RFC 1992, August 1996.

[RFC4033] R. Arends, R. Austein, M. Larson, D. Massey, and S. Rose, "DNS Security: Introduction and Requirements", RFC 4033, March 2005.

[RFC4423] R. Moskowitz and P. Nikander, "Host Identity Protocol (HIP) Architecture", RFC 4423, May 2006.

[ALT] D. Farinacci, V. Fuller, D. Meyer, and D. Lewis, "LISP Alternative Topology (LISP-ALT)", draft-ietf-lisp-alt-10.txt (work in progress), December 2011.

[NERD] E. Lear, "NERD: A Not-so-novel EID to RLOC Database", draft-lear-lisp-nerd-09.txt (work in progress),

April 2012.

[ILNP] R.J. Atkinson and S.N. Bhatti, "ILNP Architectural Description", draft-irtf-rrg-ilnp-arch-05.txt (work in progress), May 2012.

[Chiappa] J. N. Chiappa, "Endpoints and Endpoint Names: A Proposed Enhancement to the Internet Architecture", Personal draft (work in progress), 1999, <http://www.chiappa.net/~jnc/tech/endpoints.txt>.

[Jakab] L. Jakab, A. Cabellos-Aparicio, F. Coras, D. Saucez, and O. Bonaventure, "LISP-TREE: A DNS Hierarchy to Support the LISP Mapping System", in 'IEEE Journal on Selected Areas in Communications', Vol. 28, No. 8, pp. 1332-1343, October 2010.

[McQuillan] J. M. McQuillan, W. R. Crowther, B. P. Cosell, D. C. Walden, and F. E. Heart, "Improvements in the Design and Performance of the ARPA Network", Proceedings AFIPS 1972 FJCC, Vol. 40, pp. 741-754.

[Templin] F. Templin, "LISP WG", LISP WG list message, Message-ID: 39C363776A4E8C4A94691D2BD9D1C9A1 05B0AC71@XCH-NW-7V2.nw.nos.boeing.com, 13 March 2009,, <http://www.ietf.org/mail-archive/web/lisp/current/msg00269.html>.

[Wasserman] M. Wasserman, "IPv6 networking: Bad news for small biz", IETF list message, Message-Id: D11C4A34-7362-423E-A60E-476FC5D61D37@lilacglade.org, 5 April 2012, <https://www.ietf.org/ibin/c5i?mid=6&rid=49&gid=0&k1=933&k2=62733&tid=1340933524>.

Appendix A. RefComment

Appendix B. Glossary/Definition of Terms

- Address
- Locator
- EID
- RLOC
- ITR
- ETR
- xTR
- PITR
- PETR
- MR
- MS
- DFZ

Appendix C. Other Appendices

```
   -- Location/Identity Separation Brief History
   -- LISP History
   -- Old models (LISP 1, LISP 1.5, etc)
   -- Different mapping distribution models (e.g. LISP-NERD)
   -- Different mapping indexing models (LISP-ALT
       forwarding/overlay model),
       LISP-TREE DNS-based, LISP-CONS)
```

Author's Address


```
   J. Noel Chiappa
   Yorktown Museum of Asian Art
   Yorktown, Virginia
   USA

   EMail: jnc@mit.edu
```